Tolerating Multiple Faults in Multistage Interconnection Networks with Minimal Extra Stages

Chenggong Charles Fan, *Student Member*, *IEEE*, and Jehoshua Bruck, *Senior Member*, *IEEE*

Abstract—In their 1982 paper, Adams and Siegel proposed an Extra Stage Cube Interconnection Network that tolerates one switch failure with one extra stage. We extend their results and discover a class of Extra Stage Interconnection Networks that tolerate multiple switch failures with a minimal number of extra stages. Adopting the same fault model as Adams and Siegel, the faulty switches can be bypassed by a pair of demultiplexer/multiplexer combinations. It is easy to show that, to maintain point to point and broadcast connectivities, there must be at least f extra stages to tolerate f switch failures. We present the first known construction of an Extra Stage Interconnection Network that meets this lower-bound. This n-dimensional Multistage Interconnection Network has $n+f$ stages and tolerates f switch failures. An n-bit label called mask is used for each stage that indicates the bit differences between the two inputs coming into a common switch. We designed the fault-tolerant construction such that it repeatedly uses the singleton basis of the n-dimensional vector space as the stage mask vectors. This construction is further generalized and we prove that an n-dimensional Multistage Interconnection Network is optimally fault-tolerant if and only if the mask vectors of every n consecutive stages span the n-dimensional vector space.

Index Terms—Multistage Interconnection Networks (MIN), fault tolerance, extra-stage, switch faults, stage masks.

1 Introduction

MULTISTAGE Interconnection Network (MIN) has enjoyed important applications in fields such as telecommunications and parallel computing in the past decades [1], [4], [10], [12], [13]. The fault-tolerant capabilities of MIN have been widely studied. In 1982, Adams and Siegel introduced the Extra Stage Cube [2], a construction that tolerates one switch failure with one additional switching stage. In this paper, we study how to construct Multistage Interconnection Networks to tolerate multiple switch faults with extra stages. The main result of this paper is the discovery of a class of constructions that tolerates multiple switch faults with a minimal number of redundant switching stages.

Many types of cube-type Multistage Interconnection Networks have been proposed, such as the baseline, delta, generalized cube, indirect binary n-cube, omega, and shuffle-exchange. It has been proven that these cube-type MINs are topologically equivalent [3]. We will focus our attention on one particular construction, shown in Fig. 1. This MIN allows point-to-point connection from an input node to any output nodes. There is only one path between a pair of nodes. When there is a fault on that path, the communication will fail.

A 2×2 switch is a common building block of this MIN. For point to point connections, the 2×2 switch operates in either the *straight* mode or the *exchange* mode, as illustrated by Fig. 2. Two additional broadcast modes of operations exist to enable one node to send a message simultaneously to all other nodes (Fig. 3). Two connectivity models are considered in this paper, namely, the point-to-point connectivity between any two nodes and the broadcast connectivity. In this paper, we will first prove the results for the point-to-point connections. The results will then be extended to the broadcast case.

We assume the same fault model as the Adams and Siegel paper. When a switch is at fault, it is stuck in the *straight* mode. This fault model can be implemented by using two demultiplexor/multiplexor pairs for each 2×2 switch. The multiplexors and the demultiplexors are assumed to be fault-free. In the solution proposed by Adams and Siegel, not all switches need to be accompanied by these demultiplexor/multiplexor pairs. Only the switches in the first and last stages need them. Similarly, for the solutions proposed in this paper, they may not be needed for all switches. Only the switches in the first and last f stages need them to tolerate f switch faults.

For the MIN in Fig. 1, there are three stages of 2×2 switches that interconnect the eight nodes. Each node is labeled by a binary vector. The length of this vector, n, is the dimension of the MIN. Clearly, $n = log_2 N$, where N is the number of nodes in the MIN. Each switch is also characterized by an n-bit vector, called a "mask." The mask indicates the difference between the two input nodes, $B - A$. This difference is obtained by modular-2 vector subtraction. Notice that all switches in the same stage have the same masks, therefore we can associate the entire stage with a single stage mask, shown above each stage in the figure.

In this example, The MIN has the singleton mask set: $\{m_1 = 001, m_2 = 010, m_3 = 100\}$. This mask set forms a basis of the three-dimensional space, therefore all vectors in this space can be represented as a linear combination of the masks. In other words, this mask set spans the three-dimensional vector space. Consequently, we can find a path between any pair of nodes by using the following routing algorithm: To route a connection between node A and node B, we decompose the difference between A and B into a linear combination of the masks.

$$B - A = \sum_{i=1}^{n} c_i m_i. \tag{1}$$

The switches in stage i operates in the *straight* mode if $c_i = 0$ and in the *exchange* mode if $c_i \neq 0$.

Shown in the bottom half of Fig. 1 is the Bar Diagram [8] representation of the same MIN. Each node in the MIN is represented by a horizontal bar in the Bar Diagram and each switch is represented by a vertical bar. A broken vertical bar in the diagram indicates a faulty switch in the MIN. Connectivity exists between two nodes if and only if a path can be found between these two nodes. Such a path must use at most one switch at each stage and must not change direction inside the MIN, as shown in the figure. Tolerating f switch faults in the MIN is equivalent to tolerating f broken vertical bars in the Bar Diagram.

To tolerate broken vertical bars in the Bar Diagram, we need to find disjoint paths between any pair of nodes. Two paths are disjoint in a Bar Diagram if they share no vertical bars. To tolerate f broken vertical bars, it is sufficient and necessary to find $f+1$ mutually disjoint paths between all pairs of nodes. It is sufficient because f broken vertical bars can at most break f disjoint paths and there is at least one path left between each pair of nodes. It is necessary because if only f disjoint paths can be found between some pair, f broken vertical bars can break all of them, and destroy the connectivity between that pair.

In the MIN shown in Fig. 1, one and only one path can be found between any pair of nodes. Therefore, it cannot tolerate switch faults. To make this MIN single-fault-tolerant, redundant stages need to be added. This problem of tolerating a single switch fault with extra stages has been investigated extensively in the past. Adams and Siegel first proposed a solution, called Extra Stage Cube (ESC) [2]. Xu and Shen proposed a solution, which can be viewed as adding an extra stage with an all-1 mask [16]. An

• *The authors are with the Department of Electrical Engineering, 136-93 California Institute of Technology, Pasadena, CA 91125. E-mail: {fan, bruck}@paradise.caltech.edu.*

Manuscript received 27 Oct. 1997; revised 15 Mar. 1999; accepted 21 June 2000.
For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number 105827.

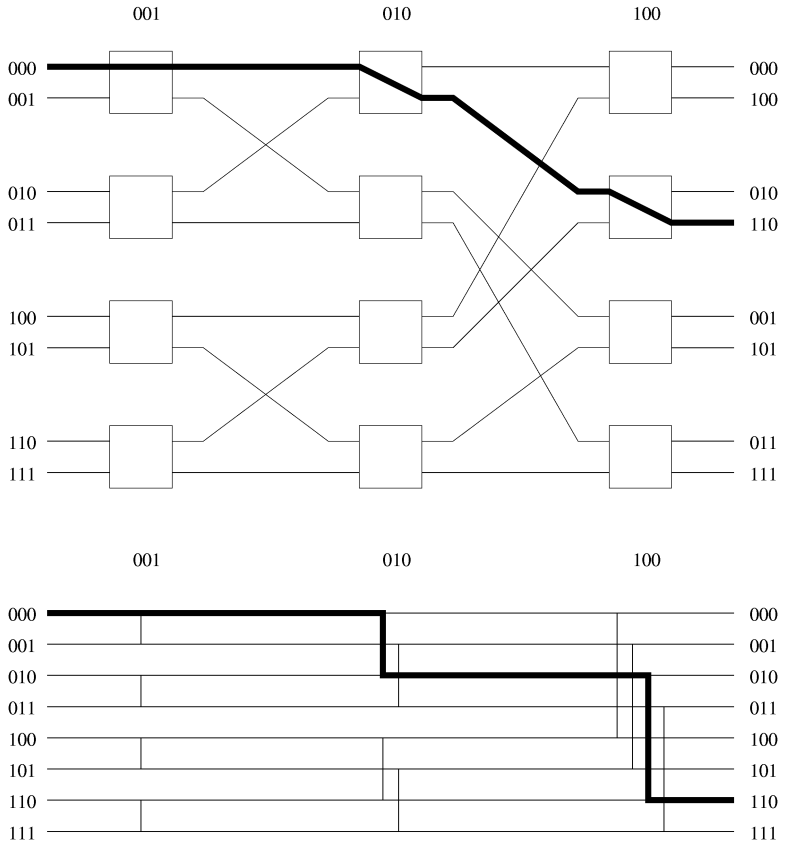

Fig. 1. A three-dimensional Multistage Interconnection Network using 2×2 switches.

example of single-fault-tolerant MIN with eight nodes is shown in Fig. 4.

Adams and Siegel showed that the ESC tolerates one switch fault by finding two disjoint paths between all pairs of nodes. There is another simple way to see that ESC is one-fault-tolerant. After a switch fault occurs, even if we discard the entire stage that contains the faulty switch, the masks of the remaining three stages still span the space. Therefore, the difference between any pair of nodes can still be decomposed into a linear combination of the remaining three masks and a correct routing is therefore available by using the three surviving stages. Two disjoint paths between 000 and 110 are outlined in the figure. This scheme, in essence, tolerates a stage fault, i.e., it tolerates any number of switch faults if all of them occur in the same stage. The ESC solution does not, however, tolerate two switch faults when they occur in different stages.

This is not a unique solution to the single-fault-tolerant problem. There exist other solutions that tolerate a single switch fault. We present one of these solutions in Fig. 5. This is also a one-extra-stage construction and the extra stage is masked 001. This MIN does not tolerate a stage fault since erasing stage 010 or stage 100, the masks of the three remaining stages, does not span the space. But, this MIN can indeed tolerate a single switch fault. The two disjoint paths between 000 and 110 are outlined in the figure as an example.

The problem of tolerating stage faults has been investigated in previous research works [5], [14]. Bruck and Ho correlated the problem of fault-tolerant MIN to the results in the field of Error-Correcting Codes and proved that a MIN constructed according to a (n, k, d) code can tolerate d-1 switch faults, as well as stage faults [5]. It showed that a fault-tolerant MIN constructed according to an MDS code uses an optimum number of extra stages to tolerate f stage faults. Despite extensive research in the field [6], [7], [9], [11], [12], [15], constructions that use a minimal number of extra stages to tolerate f switch faults, $f > 1$, had not been proposed.

The two examples we showed led us to consider the following questions: Are the existing solutions the best we can do in tolerating switch faults? If not, what is? Furthermore, if we are able to find optimal constructions, are those constructions the only solutions? The answers to all of these questions are the main contributions of this paper.

In Section 2, we propose a construction of fault-tolerant Multistage Interconnection Networks that uses an optimal number of extra stages to tolerate f switch faults. In that section, we first prove that, to tolerate f switch faults, at least f extra stages must be added. None of the previously proposed constructions meets this

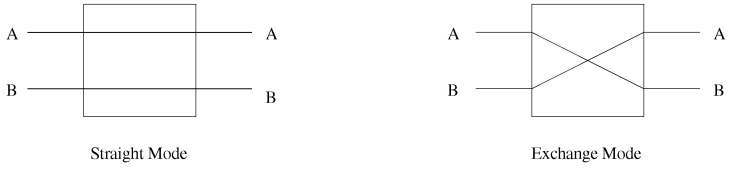

Fig. 2. The modes of operation for point to point connection on a 2×2 switch.

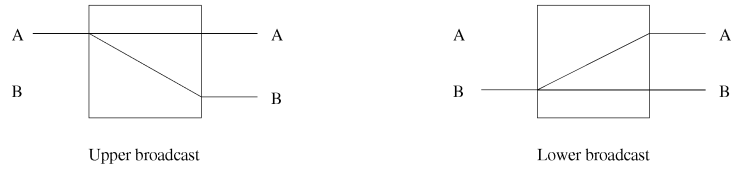

Fig. 3. The modes of operation for broadcast connection on a 2×2 switch.

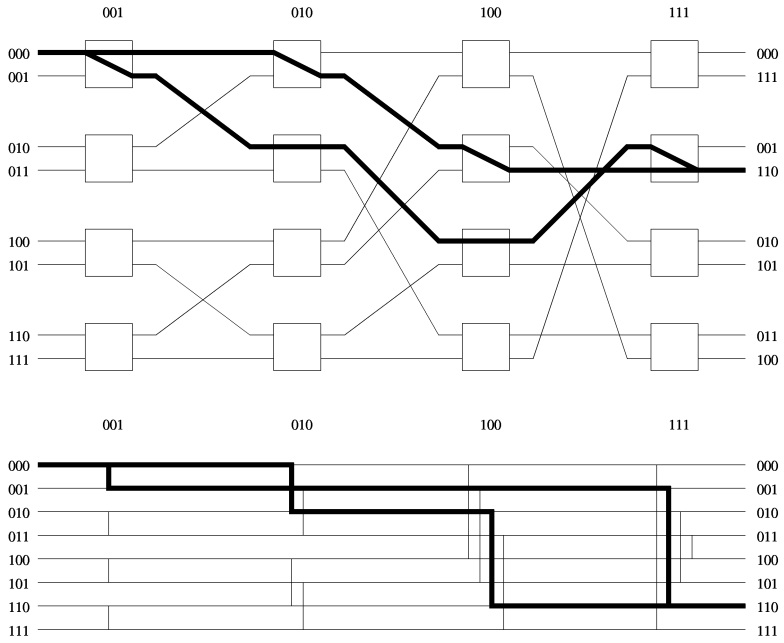

Fig. 4. Three-dimensional one-extra-stage Extra Stage Cube (ESC) Network.

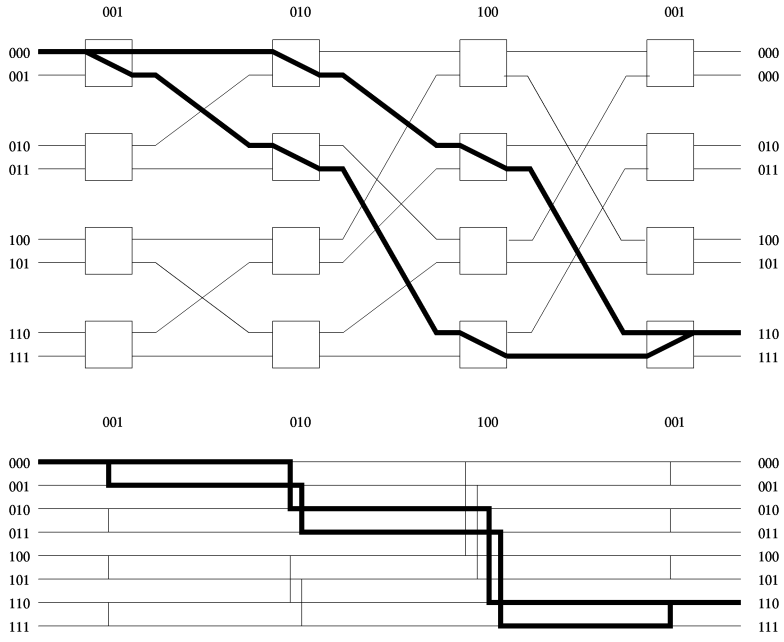

Fig. 5. Three-dimensional one-extra-stage Cyclic Multistage Interconnection Network.

lower bound when f is greater than 1. We then propose a new construction that meets this lower bound. The routing algorithm is also given in that section. We will also show that this construction is easy to implement in practice. In Section 3, we generalize the construction proposed in Section 2 and prove a necessary and sufficient condition for MINs to tolerate any given number of switch faults with an optimal number of extra stages. While we focus on the MINs that use 2×2 switches under the point to point connection model in Section 2 and Section 3, we extend the results to the Mulstistage Interconnection Networks that use $t \times t$ switches and MINs under the broadcast model in Section 4. In Section 5, we conclude.

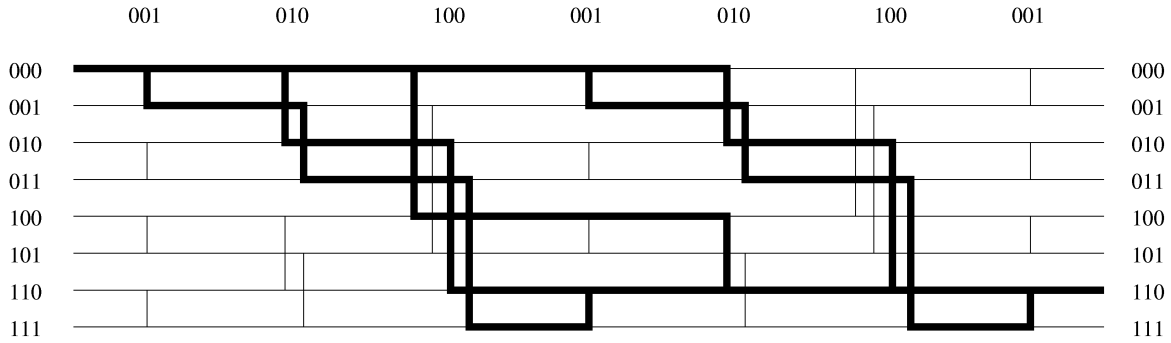

Fig. 6. Three-dimensional four-extra-stage Cyclic Multistage Interconnection Network.

2 An Optimal Construction

We first present the following theorem which states the lower bound on the number of extra stages required to tolerate f switch faults for MINs with $t \times t$ switches.

Theorem 1. *To tolerate f switch faults in an n-dimensional Multistage Interconnection Network with $t \times t$ switches, at least f extra stages must be added.*

Proof. (by contradiction) Suppose only $f-1$ extra stages were added. When f failures occur at the switches that are connected to node 0 in the first f stages, the first f stages are completely paralyzed in connecting node 0 to other nodes. Only $n-1$ stages can be used to connect node 0 to the other $t^n - 1$ nodes. It is clearly not possibl, since, with $n-1$ stages, at most t^{n-1} nodes can be reached. □

None of the previously proposed MINs meets this lower bound for any given f. For example, the ESC solution only works for $f = 1$ [2]; the number of switch faults that the Error-Correcting Code solutions tolerate is in general less than the redundant stages required [5].

Now, we present a new construction of MINs with 2×2 switches that meets this lower bound.

Definition 1 (Cyclic Multistage Interconnection Networks). *An (n, f) Cyclic Multistage Interconnection Network is an n-dimensional f-extra-stage MIN which has the singleton basis of the n-dimensional binary vector space as the masks of its first n stages and $m_{n+i} = m_i$ for $1 \le i \le f$, where m_i is the mask vector of stage i.*

A (3, 4) cyclic Multistage Interconnection Network is illustrated in Fig. 6. The following theorem implies that this MIN tolerates four faults, therefore meets the lower bound stated in Theorem 1. The five mutually disjoint paths can be found between any pair of nodes. In the figure, the paths between node 000 and node 110 are outlined.

Theorem 2. *An (n, f) Cyclic Multistage Interconnection Network with 2×2 switches tolerates f switch faults.*

Proof. We will prove the theorem by explicitly showing that between any two nodes, A and B, $A \neq B$, there are $f+1$ mutually disjoint paths in the Bar Diagram. Two paths are disjoint if they share no vertical bars in the Bar diagram. In other words, there is a conflict if and only if one switch operates in the *exchange* mode in two different paths. Please note that, in this proof, the nodes A and B and the masks $\{m_1, m_2, \ldots, m_{n+f}\}$ are n-bit binary vectors and all arithmatic operations between them are bitwise mod-2 additions.

We construct the $f+1$ paths as follows: In path i, $1 \le i \le f$, the switch in stage i operates in the *exchange* mode. Stages $i+1$ through $i+n$ are used to route the path to destination B. This is possible because every n consecutive masks in the Cyclic MIN span the n-dimensional binary vector space by definition. The switches in all other stages operate in the *straight* mode.

In the last path, path $f+1$, the switches in stages 1 through f operate in the *straight* mode and the switches in stages $f+1$ through $n+f$ route the path from node A to node B. Fig. 7 illustrates this construction.

Now, we need to show that these paths are mutually disjoint from each other. To prove that, we will prove that path i, $1 \le i \le f+1$, is disjoint from path j, $1 \le j < i$. Three cases are considered:

Case 1: $j < i - n$, there are no common stages in which both paths operate in the *exchange* mode. Therefore, path i and path j are disjoint.

Case 2: $j = i - n$, the only stage of possible conflict is in stage i if both path j and path i perform the *exchange* operation on the same switch in this stage. We know that, in stage i, by construction, path i is switching from node A to some node and path j is switching from some node to node B. Given $A \neq B$, the only possibility of conflict is that both paths are switching from A to B. It is not possible. Path j switches from A to some other node at stage j and, because of the fact that n consecutive masks in a Cyclic MIN are linearly independant, path j will not reach A until after stage i. Therefore, path i and path j are disjoint.

For $i - n < j < i$, the two paths are disjoint until stage $j+n$ since path j *exchanges* at stage j while path i goes *straight* and they differ at bit $(j \bmod n)$. After stage $j+n$, they must agree on bit $(j \bmod n)$ since they must reach the same destination B. Therefore, only one of the two paths will use the switch at stage $j+n$. Consequently, path i and path j are disjoint.

Hence, the $f+1$ paths from A to B are mutually disjoint and a (n, f) Cyclic Multistage Interconnection Networks tolerates f switch faults. □

Theorem 2 shows that the performance of Cyclic Multistage Interconnection Networks meets the lower bound stated in Theorem 1. In other words, this construction is optimal in the number of extra stages used to tolerate any given number of switch faults.

Notice that, in the construction of the Cyclic Multistage Interconnection Networks, the $f+1$ paths share horizontal lines only in the first f and last f stages. This indicates that the demultiplexor/multiplexor pairs that enable the *stuck in straight* fault model are only needed for those stages.

Routing in the Cyclic Multistage Interconnection Network is performed by routing tags. [2]. The routing tag is a binary vector composed of the coefficients obatained by decomposing the difference between the source and destination into the linear combinations of the stage masks. The tag is computed at the source node only once and is attached to each message. The ith bit in the tag in turn determines whether the switch in ith stage operates in the *exchange* or the *straight* mode. The proof for Theorem 2

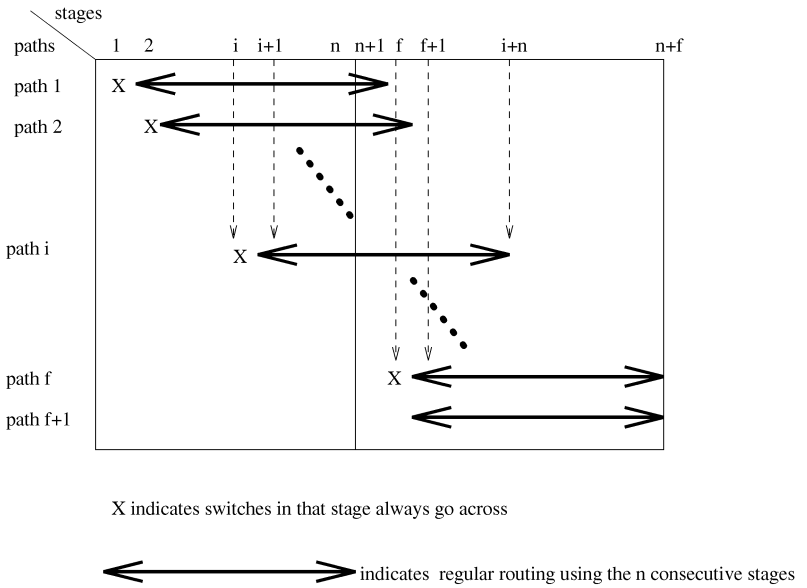

Fig. 7. Construction of $f+1$ disjoint paths.

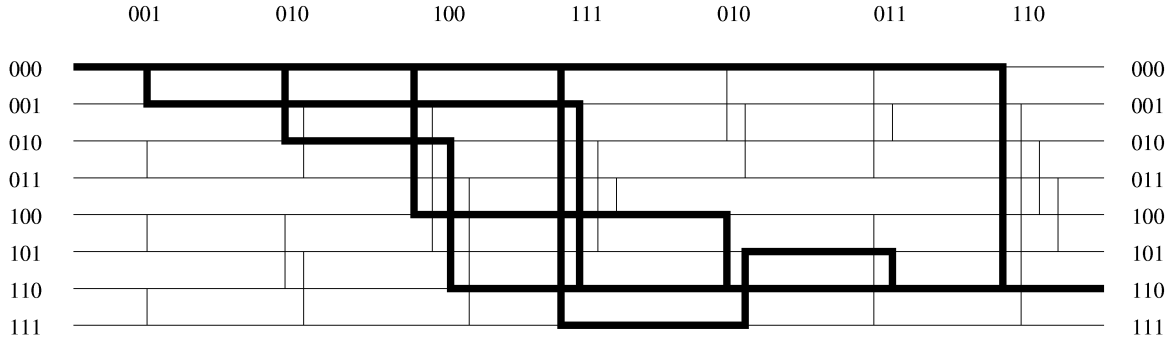

Fig. 8. Three-dimensional 4-extra-stage Generalized Cyclic Multistage Interconnection Network.

explicitly gives the construction of $f+1$ disjoint paths between any two nodes. Therefore, for a source node in a failure-free situation, there are $f+1$ sets of routing tags to each destination. A message can be sent by selecting any one out of the $f+1$ tags. When a switch failure occurs, it should be detected by the source node and the route that uses the faulty switch is eliminated from the tag selections. This tag elimination process only occurs once for each failure, thus it is not computationally intensive for the source node.

In practice, it is easy to use the advantage of the Cyclic Multistage Interconnection Networks. For example, the manufacturer can build demultiplexor/multiplexor pairs around each 2×2 switch for the n-dimensional MIN, then, by putting two of them together head-to-tail, you have an n-fault-tolerant MIN.

3 A Necessary and Sufficient Condition for Optimal Fault Tolerance

In Section 2, we introduced the Cyclic Multistage Interconnection Network that demonstrates optimal performance in tolerating any number of switch faults. The construction, however, is not unique. In this section, we extend the results to a more general class of fault-tolerant Multistage Interconnection Networks, named Generalized Cyclic Multistage Interconnection Networks.

Definition 2 (The Generalized Cyclic Multistage Interconnection Network). *An (n, f) Generalized Cyclic Multistage Interconnection Network is an n-dimensional f-extra-stage Multistage Interconnection Network which has the property that the masks of every n consecutive stages span the n-dimensional vector space.*

Fig. 8 illustrates a (3, 4) generalized cyclic MIN using a nonsingleton and nonrepetitive mask set. The five disjoint paths between node 000 and node 110 are shown in the illustration.

Clearly, the Cyclic MINs is a subclass of the Generalized Cyclic MINs. We will prove that the Generalized Cyclic MINs have the same fault tolerance capabilities as the Cyclic MINs, namely, they tolerate f faults with f extra stages. In addition, it is the necessary condition for a Multistage Interconnection Networks to demonstrate the optimal fault-tolerance capability.

Theorem 3. *An n-dimensional f-extra-stage Multistage Interconnection Network with 2×2 switches tolerates f switch faults if and only if the masks of every n consecutive stages span the n-dimensional vector space.*

Proof. We prove the forward direction of the theorem by contradiction. Suppose an n-dimensional f-extra-stage Multistage Interconnection Network does not have the property that the masks of any n consecutive stages span the space. There exists n consecutive stages in the MIN whose masks do not span the n-dimensional space. We can find node A and B, between which a path cannot be found in the nonspanning n stages. Suppose the faults happen on the switches of the remaining f stages at both sides of these n nonspanning stages in such a way that all the faults before the n stages happen at switches connected to node A and all the faults after the n stages

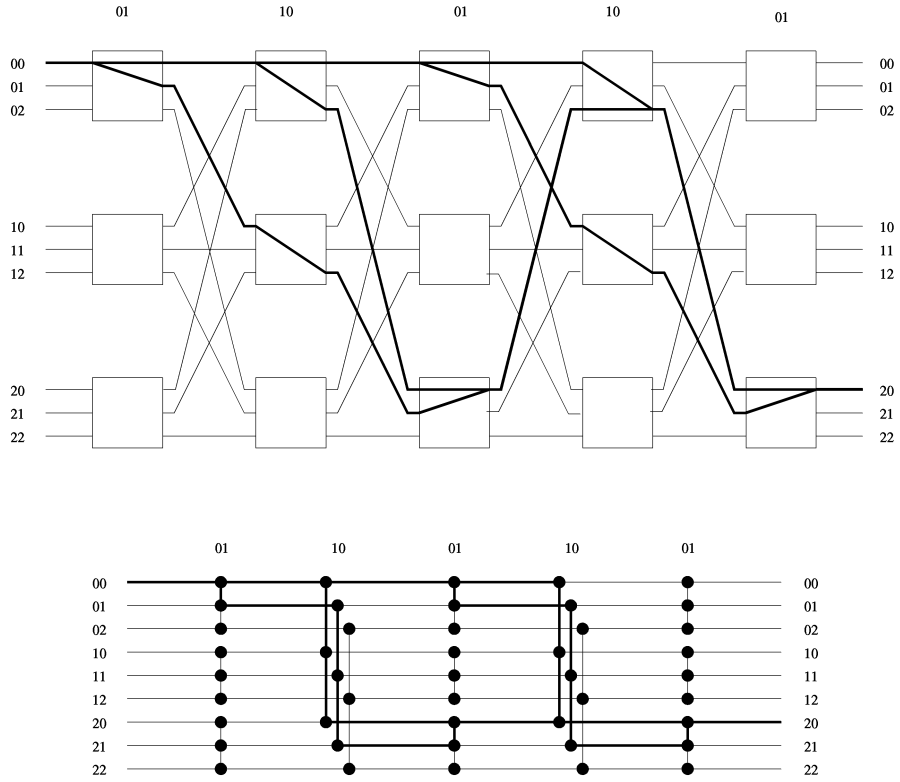

Fig. 9. Extension to the MINs with 3×3 Switching Elements.

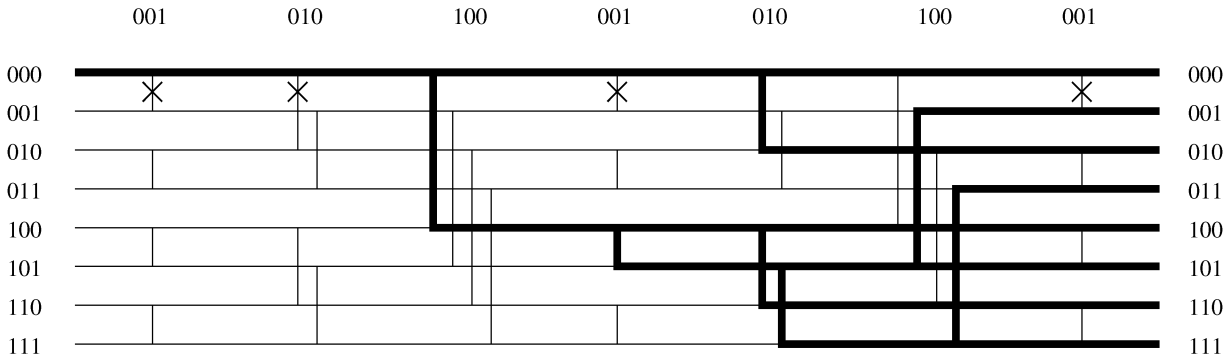

Fig. 10. Survived broadcast tree in the presence of four faults in a (3,4) Cyclic MIN.

happen at switches connected to node B. The communication between A and B fails. Therefore, an n-dimensional f-extra-stage MIN tolerates f switch faults only if the masks of every n consecutive stages span the n-dimensional vector space.

The proof of the backward direction is similar to the proof of Theorem 2. The construction of the $f+1$ paths from node A to node B are the same. We need to show that these paths are all mutually disjoint from each other. Again, we prove that path i, $i \leq f+1$, is disjoint from path j, $j < i$, by considering three cases:

Case 1: $j < i - n$, there are no common stages that the switch in the stage operates in *exchange* mode for both paths. Therefore, path i and path j are disjoint.

Case 2: $j = i - n$, the two paths share stage i. We know that path j *exchanges* at stage j, while path i goes *straight*. Since any n consecutive masks are linearly independent, m_j cannot be represented by a linear combination of m_{j+1} through m_{i-1}. Therefore, the two paths are disjoint until stage i and the only way that the two paths conflict is that, at stage i, path j *exchanges* from $A - m_i$ to A while path i *exchanges* from A to $A - m_i$. But, it is not possible since path j must reach B after ith stage and $A \neq B$. Therefore path i and path j are disjoint.

For $i - n < j < i$, since path j *exchanges* at stage j while path i goes *straight*, the two paths are disjoint until stage $j+n$ with the same reasoning as the previous case. At stage $j+n$, only one of the two paths *exchanges* since they must reach the same destination. Therefore, path i and path j are disjoint. □

4 Extensions

In this section, we will make two extensions to the results presented in the previous sections. First, instead of looking at Multistage Interconnection Networks with 2×2 Switching Elements, we will show that the theorems presented in the previous sections also apply to the MINs consisting of $t \times t$ Switching Elements. Following that, we will show that the results are also valid if we are to guarantee the broadcast capabilities of the network.

Let us look at a 9-node 3-extra stage $(2,3)$ generalized cyclic Multistage Interconnection Network consisting of 3×3 switches.

Fig. 9 shows the four mutually disjoint paths from node 00 to node 20.

Theorem 4. *An n-dimensional f-extra-stage Multistage Interconnection Network with $t \times t$ switches tolerates f switch faults if and only if the masks of every n consecutive stages span the n-dimensional vector space.*

Proof. The proof of the forward direction is the same as the proof for the 2×2 case. To prove the backward direction, we similarly construct $f+1$ paths from node A to node B and prove that they are disjoint. The difference lies in the construction of the first f paths. The reason for the modification is that a 2×2 switch can only go *straight* or *exchange*, while a $t \times t$ switch has t ways of switching. We say a switch is in mode s if, for that switch:

$$output = input + s \times mask \qquad 0 \le s \le t-1. \tag{2}$$

In this proof, all vector operations are mod-t. In the construction of path i, $i \le f$, we decompose the n-dimensional vector $B - A$ into a linear combinations of the masks $\{m_i, m_{i+1}, \ldots, m_{i+n-1}\}$:

$$B - A = \sum_{j=i}^{i+n-1} c_j m_j. \tag{3}$$

Since $\{m_i, m_{i+1}, \ldots, m_{i+n-1}\}$ span the space, such a decomposition is always possible. If the coefficient $c_i \neq 0$, the switches in stage i are forced to be in mode c_i, i.e., $output = input + c_i \times m_i$; if $c_i = 0$, we only need to make sure that the switches in stage i *exchange* to some output, as long as they do not go *straight*. The path i will reach the destination B by a regular routing in the next n stages, i.e., stages $i+1$ through $i+n$. The construction of path $f+1$ and the proof that these $f+1$ paths are disjoint to each other are the same as the proof for the 2×2 case. □

In the previous sections, we have shown that, in an f-extra-stage Cyclic MIN, there exist $f+1$ disjoint paths between any pair of nodes. It follows that, in the presense of f faults, at least one path remains intact between any pair of nodes. The broadcast from any node A is achieved by picking a survived path between node A and every other node. If a switch is used by more than one path and the two paths enter the switch from two different inputs, we collapse the part of the two paths before the switch so that the switch will operate in one of the legal modes. It is obvious that, after such collapses, the connections between the node pairs remain. Therefore, an f-extra-stage Cyclic Multistage Interconnection Network guarantees broadcast connectivity in the presense of f switch faults. As an example, Fig. 10 shows the survived broadcast switch faults.

5 Conclusion

In this paper, we studied the fault tolerance capabilities of Multistage Interconnection Networks. The fault model we used is the *stuck-in straight* model and the fault tolerance criterion is to guarantee point-to-point and broadcast connections. We constructed the first known fault-tolerant Multistage Interconnection Network that is optimal in the number of redundant stages. In addition, we proved the general condition that is both sufficient and necessary for MINs to achieve this optimal performance in tolerating switch faults.

Acknowledgments

The authors wish to thank the editor and the reviewers for their valuable input to this paper. This research was supported in part by US National Science Foundation (NSF) Young Investigator Award CCR-9457811, by the NSF Graduate Fellowship, by the Sloan Research Fellowship, and by DARPA and BMDO through an agreement with NASA/OSAT.

References

[1] G. Adams, D. Agrawal, and H. Siefel, "A Survey and Comparison of Fault-Tolerant Multistage Interconnection Networks," *Computer,* pp. 14-27, June 1987.

[2] G. Adams and H. Siegel, "The Extra Stage Cube: A Fault-Tolerant Interconnection Network for Supersystems," *IEEE Trans. Computers,* vol. 31, no. 5, pp. 443-454, May 1982.

[3] D.P. Agrawal, "Graph Theoretical Analysis and Design of Multistage Interconnection Networks," *IEEE Trans. Computers,* vol. 32, no. 7, pp. 637-648, July 1983.

[4] V.E. Benes, "Optimal Rearrangeable Multistage Connecting Networks," *Bell System Technical J.,* no. 4, II, July 1964.

[5] J. Bruck and C. Ho, "Fault-Tolerant Cube Graphs and Coding Theory," *IEEE Trans. Information Theory,* Nov. 1996.

[6] H. Cam and J. Fortes, "Rearrangeability of Shuffle-Exchange Networks," *Proc. Third Symp. Frontiers of Massively Parallel Computation,* pp. 303-314, Oct. 1990.

[7] P.J. Chuang, "CGIN: A Fault Tolerant Modified Gamma Interconnection Network," *IEEE Trans. Parallel and Distributed Systems,* vol. 7, no. 12, pp. 1,303-1,308, Dec. 1996.

[8] D. Knuth, *The Art of Computer Programming: Sorting and Search.* Reading, Mass.: Addison-Wesley, 1973.

[9] V.P. Kumar and S.M. Reddy, "Augmented Shuffle-Exchange Multistage Interconnection Networks," *Computer,* pp. 30-40, June 1987.

[10] F. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes.* San Mateo, Calif.: Morgan Kaufmann, 1992.

[11] Y. Liao, M. Lu, and N.F. Tzeng, "The Palindrome Network for Fault-Tolerant Interconnection," *Proc. Eighth IEEE Symp. Parallel and Distributed Processing,* pp. 556-561, Oct. 1996.

[12] N. Linial and M. Tarsi, "Interpolation between Bases and the Shuffle Exchange Network," *European J. Combinatorics,* pp. 29-39, Oct. 1989.

[13] M. Pease, "The Indirect Binary n-Cube Microprocessor Array," *IEEE Trans. Computers,* vol. 26, no. 5, pp. 458-473, May 1977.

[14] C. Shih and K. Batcher, "Adding Multiple-Fault Tolerance to Generalized Cube Networks," *IEEE Trans. Parallel and Distributed Systems,* vol. 5, no. 8, pp. 785-792, Aug. 1994.

[15] A. Varma and C.S. Raghavendra, "Fault-Tolerant Routing in Multistage Interconnection Networks," *IEEE Trans. Computers,* vol. 38, no. 3, pp. 185-393, Mar. 1989.

[16] M. Xu and X. Shen, "A New Fault-Tolerant Generalized Cube with an Extra Stage," *Proc. Int'l Computer Science Conference (ICSC '92),* pp. 99-105, Dec. 1992.