Fault-tolerance schemes for WDM-based multiprocessor networks

Poornima Lalwaney and Israel Koren
Department of Electrical and Computer Engineering
University of Massachusetts
Amherst, MA 01003

Abstract

Due to the difference in implementing network topologies in the optical and electronic domain, schemes to tolerate link failures in electronic implementations may not be the most efficient ones for optical implementations of multiprocessor networks. Rerouting messages along alternate paths in the network topology is commonly used in electronically implemented networks to tolerate link failures. In this paper, we propose two schemes that use the properties of wavelength division multiplexing to tolerate link failures in optically interconnected multiprocessor networks. These are referred to as wavelength reassignment *and* time division multiplexing. *We show that optically implemented networks have better fault-tolerance properties than electronic ones.*

1 Introduction

Driven primarily by bandwidth advantages, optical interconnects are being considered as alternatives to electronic interconnects in computing systems. The advantage of fiber-optics in local and wide area networks is well established. Due to advances in semiconductor optoelectronic technology over the past decade, devices with increased optical to electrical (and vice-versa) conversion efficiency, low power requirements and small physical dimensions are currently available. Optical interconnects are therefore being considered in multiprocessing and distributed processing environments [7] [8].

In optical link implementations, every unidirectional link requires a transmitter at the source node and a receiver at the destination node, transmitting and receiving data at a common wavelength over a communication medium. Optical interconnects are being experimented on in several industry and research multiprocessor prototypes to test their capabilities from system design and implementation considerations. The high bandwidth of optics has been used to reduce wiring density in a Connection Machine prototype [10] and to demonstrate system scalability in Intel's Touchtone Delta Supercomputer [6]. In both efforts, point-to-point fiber-optic links were used to overcome problems due to the limited bandwidth of the electronic interconnects. To incorporate network flexibility, some coupling/switching device is required. In the fiber-optic domain, wavelength division multiplexing (WDM) achieved by using passive star couplers is used in local area networks. In WDM-based networks, the bandwidth of the optical fiber is split into many channels, each channel carrying data at a particular wavelength [15]. The logical connectivity is obtained by assigning wavelengths to the system's transmitters and receivers. Reconfiguring the interconnection network to a different topology is a simple matter of wavelength reassignment, if the transmitters and/or receivers are tunable over the entire range of wavelengths used. In this paper, the independence of the logical and physical network topology in WDM networks is used to propose two schemes to tolerate link failures in multiprocessor networks. These schemes are referred to as wavelength reassignment and time-division multiplexing. Some reliability issues for optical networks within multiprocessors have been addressed in [13].

In Section 2, the principle of realizing a network topology using wavelength division multiplexing is outlined. The proposed fault-tolerance schemes for tolerating link failures are presented in Section 3 and their performance is compared in Section 4.

2 Wavelength-Division Multiplexing

In wavelength division multiplexing the large capacity of the optical fiber is utilized by splitting the available bandwidth into independent, noninteracting channels, thereby supporting simultaneous communication between many source-destination pairs on a common medium. In WDM-based systems, wavelength encoded signals from the transmitting nodes are multiplexed onto the fiber using a device called passive star coupler. The transmitters at the net-

0-8186-7101-7/95 $4.00 © 1995 IEEE

Proceedings of the Second Workshop on Massively Parallel Processing Using Optical Interconnections (MMPOI '95)
0-8186-7101-7/95 $10.00 © 1995 IEEE

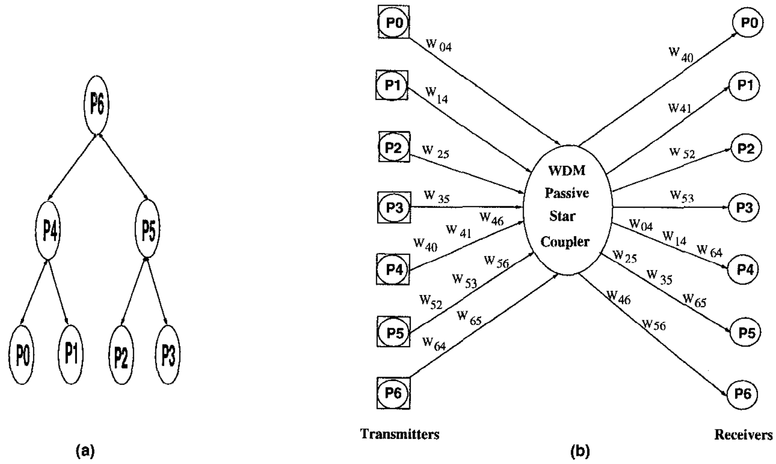

Figure 1: (a) A binary tree with seven nodes. (b) WDM star embedding of the tree. W_{ij} is the wavelength assigned for communicating from node i to node j. W_{ij} is assigned to the transmitter at processor i and the receiver at node j.

work nodes are connected to the input ports of the star coupler. The signal power at each input port is divided uniformly amongst the output ports. Thus, wavelengths from all the transmitters appear at each output port. Demultiplexing is performed at the receivers by recovering the desired input port signal from the common output port signal. The desired connectivity amongst the system nodes can be obtained by appropriate wavelength assignments to the transmitters and receivers. This is illustrated for a seven node binary tree in Figure 1. In the figure, all links are bidirectional and are implemented using fixed wavelength transmitters and receivers. Thus, the degree of a node determines the number of transmitters and receivers at that node.

Consider an example of the unidirectional link from node $P0$ to $P4$. The transmitter at node $P0$ transmits data at wavelength W_{04}. The transmitter is connected to an input port of the star coupler. The data signal on wavelength W_{04} is seen at all output ports. The receivers at network node are connected to the output ports of the coupler. The node $P4$ has 3 receivers, as it is required to receive data from nodes $P0$, $P1$ and $P6$. One of the receivers at node $P4$ is tuned to wavelength W_{04}, so that it can retrieve node $P0$'s data from output port of the coupler. In other words, W_{04} is the wavelength assigned to the transmitter at node 0 and the receiver at node 4 to realize the logical link from node 0 to 4. All the transmitters (receivers) at a network node may share a single fiber to the star coupler port. Alternately, every transmitter (receiver) at a node may have a dedicated fiber to the star coupler port. Depending on the implementation constraints, one of the above two schemes may be selected.

By using tunable transmitters and receivers, versatile interconnections can be realized. As the logical connectivity between nodes can be achieved by tuning the transmitters and receivers to the desired wavelengths, reconfiguration of the interconnection network to support a different topology can be achieved by wavelength reassignment. With the availability of narrow line width semiconductor lasers and inexpensive passive star couplers [2], WDM-based optical interconnects that support reconfiguration seem feasible for multiprocessors. Rapid progress is being made in the development of tunable devices, both in the range over which they are tunable, and their tuning times [4] [12]. Limited tuning range restricts the kinds of topologies that can be supported, and the tuning times of devices constitute the reconfiguration overheads. Both these parameters are decided by the method of tuning used. Current tuning ranges are in the $4 - 10nm$ range and the tuning times vary from nanoseconds to milliseconds [4]. Currently available passive stars can handle a maximum of a hundred wavelengths. Star couplers with a maximum of 128×128 ports are feasible with currently available technology. An experimental ISDN switch architec-

Proceedings of the Second Workshop on Massively Parallel Processing Using Optical Interconnections (MMPOI '95)
0-8186-7101-7/95 $10.00

ture using eight 128×128 multiple passive stars to handle over ten thousand input port lines has been reported [4].

3 Fault-tolerance schemes

In this section, we describe the message rerouting scheme, time-division multiplexing and the wavelength reassignment scheme for tolerating link failures in optical implementations of multiprocessor networks based on wavelength-division multiplexing. The fault model considered here assumes that a link failure is caused by a transmitter failure or a receiver failure. The time-division multiplexing and the wavelength reassignment schemes make use of the properties of wavelength-division multiplexing. They have been discussed in this section by assuming that a link failure is caused by a transmitter failure. The schemes are equally applicable to receiver failures. The performance of these schemes is compared in Section 4.

3.1 Rerouting

The commonly used technique to tolerate link failures in multiprocessor networks is to reroute messages between the source and destination along alternate paths so as to avoid the use of the failed links. These alternate paths may be present in the original network topology as with the case of the mesh topology. Several variations of network topologies that cannot provide alternate paths have been proposed by using additional links that allow for message rerouting. The tree topology is one such example. Variations of the tree topology to tolerate single link/single node failures include full-ringed and half-ringed trees [14]. The above-mentioned variations use additional horizontal links at various levels of the binary tree topology. In a full-ringed binary tree, the adjacent nodes at a level are connected. In a half-ringed tree, alternate pairs of nodes at a level are connected. Rerouting of messages affects both the average distance between nodes and the average delay between them. The increase in average distance is due to the additional hops on the alternate path used for rerouting.

3.2 Time-Division Multiplexing

In this scheme, we assume that every node has at least one *fault-free* tunable transmitter that can tune to the transmitting wavelengths of all the outgoing links at that node. In the event of a transmitter failure at the node, the tunable transmitter switches periodically between its normal operating wavelength and the wavelength of the failed transmitter. As the tunable transmitter time multiplexes the outgoing communication in two directions, we refer to this scheme as time-division multiplexing (TDM). In the event of a link failure, the TDM scheme covers the link failure by logically restoring the failed link by wavelength tuning. Upon the occurence of a transmitter failure, the messages in two directions are likely to be affected (that is, the direction of the faulty link and the direction of the tunable link). There is no increase in average distance. However, additional delays are seen along the two affected directions. Besides these delays, there is the overhead equal to the transmitter tuning time for every time the tunable transmitter switches between the two wavelengths. The above assumes that each node requires one fault-free tunable transmitter while other transmitters at the node may be fixed wavelength transmitters. The same assumptions hold for the receivers at the node so that link failures caused by receiver failures may be covered by receiver tuning.

3.3 Wavelength Reassignment

In this scheme, redundancy is introduced in the network by selectively incorporating spare transmitters and receivers at network nodes. In the event of a link failure due to a failed transmitter, the outdegree at the affected node decreases by 1. Assuming that all node in the network are identical in functionality, the node with the failed transmitter could function as a node of lower degree in the network topology considered. A node of lower degree with unused spare transmitters, having as many functioning transmitters as required at the position of the failed node, could be used to replace the failed node. This change of logical connectivity between the failed node and its replacement can be achieved by wavelength reassignment. For a topology with nonuniform degree, spares may be placed at lower degree nodes, so as to logically change its functionality with that of a higher degree node with a failed link. The number of link failures that may be tolerated depends on the number of available spare transmitters.

Let us illustrate the use of this technique by considering the binary tree topology. The degrees of the leaf node, root node and an internal node are 1, 2 and 3, respectively. We consider an optical implementation of the tree network in which all nodes have three transmitters and receivers. Upon the occurence of a link failure at an internal node due to a transmitter failure, a fault-free leaf node takes over the functions of the internal node. This is done in a single wavelength reassignment phase in which the three transmitters at the leaf node are tuned to the three distinct wavelengths that were used by the internal node. One operational transmitter at the failed internal node is set to the wavelength of the leaf node. The network structure is restored. The number of single link failures that may

Proceedings of the Second Workshop on Massively Parallel Processing Using Optical Interconnections (MMPOI '95)
0-8186-7101-7/95 $10.00 © 1995 IEEE

be covered by this scheme equals the number of fault-free leaf nodes (with spare transmitters and receivers) available.

If wavelength reassignment is possible in a network to tolerate a link failure, the network is restored to its fault-free state in a single reassignment (wavelength tuning) phase. For any two nodes to be able to swap functionality through wavelength reassignment, the transmitters at all nodes should be capable of tuning to all the wavelengths used for network implementation. Note that in the electronic implementations, the use of redundant spares is prevalent to cover single link failures [1, 3]. However, the method of reconfiguration and restoring the network to its full functioning capacity is not as easily accomplished. In optical WDM implementation, the logical topology may be different from the physical star topology. In the event of a link failure and the availability of spare wavelengths and transmitters, the wavelength reassignment phase restores the network to its fault-free state. Thus a few spare links (transmitters/receivers) at low degree network nodes can provide a high level of fault tolerance. The efficient use of a few global spares to provide a high degree of fault tolerance can be applied in networks in which the network topology is of non-uniform degree. If all nodes in the network have the same degree then spares have to be added to every network node for tolerating link failures at that node.

In WDM implementations, the amount of redundancy that may be placed in the network depends on the number of available spare wavelengths and spare input ports available on the star coupler. For small networks, all the network links may be realized on a single passive star coupler. As there are limitations on the number of available wavelengths and the number of input and output ports on a star coupler, multiple star couplers are required to implement large networks. The mapping scheme used to assign nodes to multiple star couplers does not affect the performance of the application being executed on the system. It does affect the number of link failures that may be tolerated and the number of stars required to implement this network.

The effect of the mapping scheme on network fault-tolerance is illustrated by considering two mappings for the mesh topology.

Scheme 1: Row and column mapping

In this scheme, separate stars are used for row communication and column communication in the mesh [5]. A row or a set of rows is mapped onto a single star. A similar scheme is used for the columns. Let us assume that all nodes in the network have 4 transmitters and receivers. For every row mapped onto a star, the two nodes at the edge have a spare transmitter/wavelength to cover for link failures along the row. Note that every internal node in the mesh has 4 outgoing links. Thus every node has 4 transmitters. Two of these are mapped onto a row star, and the other two are mapped onto a column star. We refer to a star used for row(column) communication as a row(column) star. The connection of transmitters to stars is a physical one and is therefore a rigid constraint. The flexibility on the logical links, comes from wavelength reassignments to the transmitters mapped onto the same star. To replace the failed internal node by its replacement along the periphery (node of degree 3), one has to reassign wavelengths to all the physical connections (transmitters) of the two nodes. In the wavelength reassignment phase, the column star on the periphery has to be assigned the wavelengths of the column star with the failed node. In turn, the column star with the failed node has to function as the column star on the periphery. The interchange of columns in the presence of a link failure along the row is necessary, as the physical connections in the column direction of the failed node and the replacement node are mapped onto different column stars. All the available redundancy is along the edge of the mesh. From the nature of partitioning the network using this mapping, to cover a single link failure along any row, a column along the periphery with redundant spares is used up. The number of link failures along rows that can be tolerated by wavelength reassignment in the entire network is therefore limited to 2. The same limitation applies to link failures along columns. The total number of link failures that may be tolerated in the network with wavelength reassignment using this mapping scheme, is therefore, limited to 4.

Scheme 2: Submesh mapping

In this scheme, the mesh network is partitioned into equal sized submeshes. Each submesh is mapped to a star. The maximum submesh size that can be mapped onto a star is determined from the wavelength and input port limitations of the star coupler. As each node in the submesh is of degree 4 (except the submeshes along the mesh periphery), additional transmitters are needed at the internal submeshes to tolerate link failures. For a topology with uniform degree, additional transmitters are required at each node of the network to cover for link failures at that node. The maximum number of additional transmitters that may be added equals the number of unused wavelengths per star. The number of unused wavelengths available per star depends on the submesh size mapped onto the star.

Proceedings of the Second Workshop on Massively Parallel Processing Using Optical Interconnections (MMPOI '95)
0-8186-7101-7/95 $10.00 © 1995 IEEE

Table 1: Number of stars required and the number of faults that may be tolerated per star and per network in the two mapping schemes as a function of the mesh size with a wavelength limit of 100 and an input port limit of 128 per star.

Mesh Size	Row/Star				Submesh/Star				
	Stars	*%Red--dancy*	*Link flts /star*	*Link flts /nw*	*Submesh size*	*Stars*	*%Red--dancy*	*Link flts /star*	*Link flts /nw*
20×20	20	8%	2	4	4×4	25	44%	36	900
22×22	22	12%	2	4	4×4	36	44%	36	1296
24×24	24	25%	2	4	4×4	36	44%	36	1296
26×26	52	4%	2	4	4×4	49	44%	36	1764
28×28	56	4%	2	4	4×4	49	44%	36	1764
30×30	60	4%	2	4	4×4	64	44%	36	2304

The performance of the two mapping schemes is compared by considering the number of link faults that may be tolerated in the network. Table 1 shows the number of unidirectional link faults that may be tolerated per star, and over the entire network for mesh sizes varying between 20×20 and 30×30. The limitation on the number of input ports and the number of wavelengths is taken as 128 and 100, respectively. In the table, the percentage of redundancy represents the percentage of the maximum available redundancy that can be used by the mapping scheme. As discussed earlier, in the row/star mapping scheme, the number of link failures that may be covered per network is limited to 4, although the number of spare wavelengths available is greater than this. Here, the maximum available redundancy cannot be used efficiently because of the placement of the spares along the periphery, and by the fact that all the transmitters from a node are not mapped onto the same star. In the submesh scheme, the efficient submesh size with a wavelength restriction of 100 is a 5×5 mesh. However, for this submesh size there are no spare wavelengths available for tolerating link failures. We therefore assume that a 4×4 submesh is mapped onto a single star. This implementation requires 64 wavelengths. The maximum number of available spare wavelengths per star is 36. As there are 16 nodes per star, the amount of redundancy that can be used to tolerate single link failures at every node of the mesh is 44%. Thus, a single link failure at every network node may be covered, because the amount of redundancy available can be utilized efficiently. The number of stars required to implement the topology in the row or column/star scheme is generally less than the submesh scheme. However, in the submesh partitioning scheme, the number of link errors that may be tolerated is much higher. Depending on the reliability of the transmitter/receiver components and the hardware cost per star, one can decide on the mapping scheme.

4 Performance Comparisons

In this section, we compare the wavelength reassignment, time-division multiplexing and the rerouting schemes by considering the effect of a single link failure on network performance.

Using the wavelength reassignment scheme, the number of single link failures that may be tolerated depends on the number of available spare wavelengths per star and the mapping scheme used. There is no effect on the average distance or the average delay. The overhead in implementing this scheme is the wavelength reassignment phase. The tuning time of the interface devices varies from nanoseconds to milliseconds, depending on the method used to achieve tuning [4]. A maximum overhead of a few milliseconds will be incurred for every link failure.

In the event of a link failure that cannot be covered by wavelength reassignment, either the TDM or the rerouting scheme may be used. When using the TDM scheme, there is no increase in average distance due to the link failure. However, there is an increase in delay attributed to the message waiting time at the affected node for the logical link to be periodically restored by the tunable transmitter for a duration referred to as the TDM time slot. In the rerouting scheme, both the average distance and average delay increase due to the presence of a link failure.

Overall network performance of the fault-free network and the effect of the TDM and the rerouting scheme on network delay was measured through simu-

Proceedings of the Second Workshop on Massively Parallel Processing Using Optical Interconnections (MMPOI '95)
0-8186-7101-7/95 $10.00 © 1995 IEEE

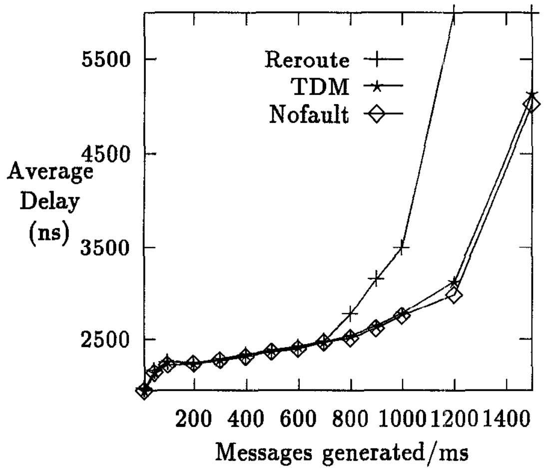

Figure 2: Variation in average delay on a 64-node mesh when using the TDM and rerouting schemes in the presence of a single link failure. Average delay in the absence of failures plotted for comparison. $\alpha = 400$ ns, $\beta = 8$ ns, and TDM slot= 1 μs.

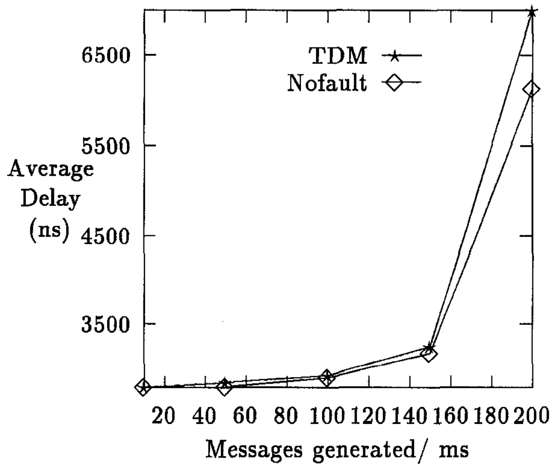

Figure 3: Variation in average delay with message generation rate per node for the fault-free binary tree and measured delays when using the TDM scheme to cover a link failure on a 63-node tree. $\alpha = 400$ ns, $\beta = 8$ ns, and TDM slot= 1μs.

lations. The mesh, binary tree, full-ring and half-ring tree topologies were simulated by the simulator developed for the purpose. Link implementations issues used in the simulator are now discussed. The time to traverse a network link comprises of the link set-up time and the time to transfer a message. The link-setup time denoted by α is assumed to be 400ns. This was the set-up time reported in Intel's effort in incorporating fiber-optic interconnects in a Touchtone Delta Supercomputer prototype. We assume that the message size is 32 bits. Assuming a 4 Gbps fiber-optic link, a 32 bit word may be transferred in 8 ns. The message transfer time denoted by β, is therefore 8 ns. For the TDM scheme, the reconfiguration overhead of the tunable transmitter or receiver denoted by t_r is assumed to be 100 ns (it may vary between nanoseconds to milliseconds depending on the method of tuning used). In using the TDM scheme, the value of the TDM time slot value is assumed to be 1 μs. Comparisons of the TDM and the rerouting schemes based on an analysis may be found in [11].

The variation in simulated values of average network delay with message generation rate on a 64-node mesh for the TDM and rerouting schemes in the presence of a single link failure is shown in Figure 2. The message generation rate is represented as the number of messages generated per node per millisecond. As seen from the figure, the two schemes have a comparable performance in covering link failures at low message generation rates. When the message generation rate exceeds 700 messages per millisecond, one observes that the TDM scheme provides a lower average delay than the rerouting scheme.

In Figure 3, we depict the variation in simulated delay for the fault-free binary tree and effect of using the TDM scheme to cover a single link failure on a 63 node tree. As there is a unique path between a source and destination in a binary tree topology, the rerouting scheme cannot be used. As can be seen from the above figure, the performance of the TDM scheme when used to cover a single link failure is close to that of the fault-free binary tree.

The full-ring tree and half-ring tree are two variations of the binary tree topology that may use rerouting to tolerate link failures. Both the topologies were simulated. Figures 4 and 5 depict the performance of the rerouting and TDM schemes in the presence of a single link failure on the full-ring and half-ring trees, respectively. The average delay in the fault-free network is plotted for comparison in the two figures. The

Proceedings of the Second Workshop on Massively Parallel Processing Using Optical Interconnections (MMPOI '95)
0-8186-7101-7/95 $10.00 © 1995 IEEE

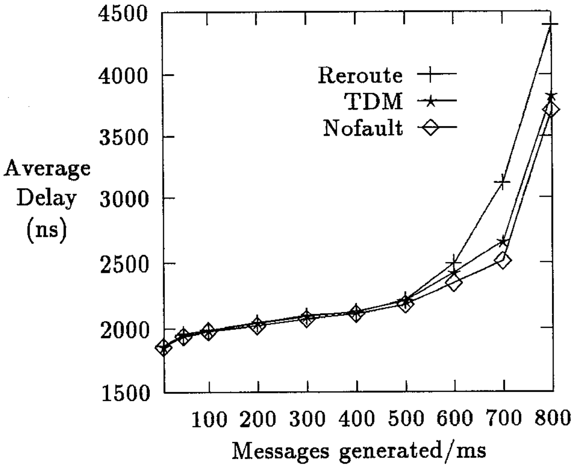

Figure 4: Variation in average delay with message generation rate when using the rerouting and TDM scheme to cover a link failure on a 63 node full-ring tree. $\alpha = 400$ ns, $\beta = 8$ ns, and TDM slot $= 1$ μs.

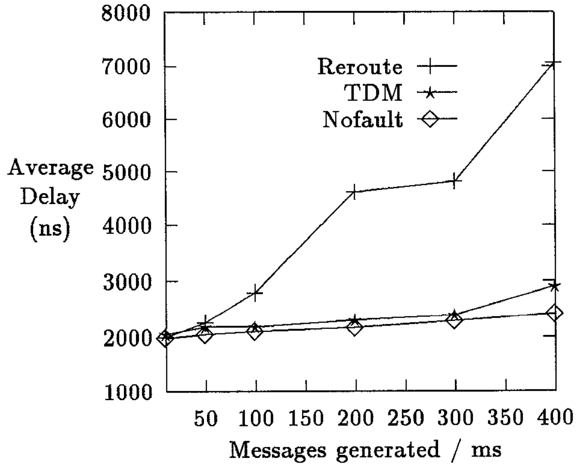

Figure 5: Variation in average delay with message generation rate when using the rerouting and TDM scheme to cover a link failure on a 63 node half-ring tree. $\alpha = 400$ ns, $\beta = 8$ ns, and TDM slot$= 1$ μs.

number of nodes in the full-ring tree and half-ring tree used for the above simulations is 63. As seen from the figures, the fault-free full-ring tree has a lower average delay than a half-ring tree. This is because the number of additional horizontal links in a full-ring tree is higher than that in a half-ring tree. The length of the alternate path used for rerouting is therefore less than that in half-ring trees. The length of the alternate path used in the rerouting scheme depends on the network topology and the routing scheme. The topology is therefore an important factor in deciding the relative performance of the TDM and rerouting scheme over the fault-free network. As an example, for a message generation rate of 300 messages per millisecond per node, the average delay increases by 1% over the fault-free network when the TDM or the rerouting scheme is used to cover a link failure. For the same message generation rate, on a half-ring tree, using the TDM scheme to cover a link failure results in a 20% increase in delay and the rerouting scheme results in a 109% increase in average delay over the delay in the fault-free half-ring network. It should be noted that 300 messages per millisecond is close to link saturation limit for a half-ring tree whereas 300 messages a millisecond is not close to the link saturation limit in a full-ring tree. As seen from Figures 2, 3, 4 and 5, the maximum permissible message generation rate (with the assumed link implementation parameters) is 1500, 200, 800 and 400 for the 64 node mesh, 63 node binary, full-ring and half-ring trees, respectively. From the above figures, one notes that the performance of the TDM scheme is better than the rerouting scheme at high message generation rates.

We now compare the wavelength reassignment, time-division multiplexing and rerouting schemes from an implementation point of view. Each one of these schemes to tolerate link failures requires different capabilities for the transmitters and receivers. In the rerouting scheme, no device tuning capability is required. This scheme could therefore be used in networks implemented with fixed wavelength transmitters and receivers. To tolerate link failures through time-division multiplexing, each node requires at least one fault-free tunable transmitter(receiver) capable of tuning to all the outgoing(incoming) wavelengths at the node. In the wavelength reassignment scheme, every transmitter (receiver) in the system has to be capable of tuning over all the wavelengths used in the network implementation. In wavelength reassignment, a single tuning/wavelength reassignment phase restores the network to the fault-free state. In the TDM and rerouting scheme, there is a degradation in network

Proceedings of the Second Workshop on Massively Parallel Processing Using Optical Interconnections (MMPOI '95)
0-8186-7101-7/95 $10.00 © 1995 IEEE

performance due to the link failure. The magnitude of the performance degradation depends on the message generation rate and the routing scheme.

5 Summary

Fault-tolerance schemes used in electronic implementations of multiprocessor networks may not be the best for optical implementations. We analyzed the fault tolerance properties of optical interconnects by considering fiber-optic networks based on wavelength-division multiplexing (WDM). Using the property of wavelength-division multiplexing, we presented two schemes, namely, wavelength reassignment and time-division multiplexing, for tolerating link failures in optical interconnects. In the wavelength reassignment scheme, redundant spares may be used to cover link failures with no degradation in network performance. In electronic networks, spare links are used to provide local redundancy. For networks of non-uniform degree, the spares provide global redundancy when using wavelength reassignment to tolerate link failures. Therefore, a few spare links can provide a high level of fault-tolerance. If wavelength reassignment cannot be used, a time-division multiplexing scheme that uses the tunability of optical devices to implement logical links was presented. An important advantage of the time-division multiplexing scheme is that one or more link failures per node may be tolerated (we assume that a tunable transmitter or receiver is fault-free). The degradation in network performance does not depend on the position of the failing link. This is significantly different from the rerouting scheme where a number of link failures at a node may not be tolerated. Wavelength reassignment and time-division multiplexing therefore have better fault-tolerance properties than rerouting schemes. The wavelength reassignment has a one-time reconfiguration cost upon the occurence of a link failure. Furthermore, the transmitters/receivers used for network implementation should be capable of tuning over a wide range of wavelengths. The TDM scheme requires tansmitters/receivers with limited tunability. The rerouting scheme places no restrictions on device tunability. The appropriate scheme to tolerate link failures may therefore be selected depending on the implementation limitations. Based on the capabilities of the wavelength reassignment and the TDM scheme, we conclude that from the point of view of network fault-tolerance optical interconnects have significant advantages over electronic interconnects.

References

[1] M. S. Alam and R. G. Melhem, "Routing in Modular Fault-Tolerant Multiprocessor Systems," *Proceedings of the Twenty-Second IEEE International Symposium on Fault-Tolerant Computing,* pp. 185-193, July 1992.

[2] C. A. Brackett, "Dense Wavelength Division Multiplexing Networks: Principles and Applications," *IEEE Journal on Selected Areas in Communications,* Vol. 8, No. 6, pp. 948-964, August 1990.

[3] J. Bruck, R. Cypher and C. T. Ho, "Wildcard Dimensions, Coding Theory and Fault-Tolerant Meshes and Hypercubes," *Proceedings of the Twenty-Third IEEE International Symposium on Fault-Tolerant Computing,* pp. 260-267, June 1993.

[4] A. Cisneros and C. A. Brackett, "A large ATM Switch Based on Memory Switches and Optical Star Couplers," *IEEE Journal on Selected Areas in Communications,* Vol. 9, No. 8, pp. 1348-1360, October 1991.

[5] P. W. Dowd, "High Performance Interprocessor Communication Through Optical Wavelength Division Multiple Access Channels," *Proceedings of the Computer Architecture Conference,* pp. 96-105, May 1991.

[6] B. E. Floren et al, "Optical Interconnects in the Touchtone Supercomputer Program," *Integrated Optoelectronics for Communication and Processing,* Proc. SPIE 1582, pp. 46-54, 1991.

[7] J. W. Goodman, F. J. Leonberger, S. Y. Kung and R. A. Athale, "Optical Interconnections for VLSI Systems," *Proceedings of the IEEE,* Vol. 72, No. 7, July 1984.

[8] A. Guha, J. Bristow, C Sullivan and A. Husain, "Optical Interconnections for Massively Parallel Architectures," *Applied Optics,* Vol. 29, No. 8, March 1990.

[9] H. S. Hinton, "Architectural Considerations for Photonic Switching Networks," *IEEE Journal on Selected Areas in Communications,* August 1988.

[10] B.O. Kahle, E.C. Parish, T.A. Lane and J. A. Quam, "Optical Interconnects for Interprocessor Communications in the Connection Machine," *IEEE Conference on Computer Design,* Cambridge, MA, October 1989.

[11] P. Lalwaney, and I. Koren, "Fault-tolerance in Optically Interconnected Multiprocessor Networks", *IEEE Workshop on Fault-Tolerant Parallel and Distributed Systems,* June 1994.

[12] P.F. Moulton, "Tunable Solid State Lasers," *Proceedings of the IEEE,* Vol. 80, No. 3, pp. 348-364, March 1992.

[13] J. J. Olsen, "Control and Reliability of Optical Networks in Multiprocessors," Ph.D dissertation, Massachussets Institute of Technology, May 1993.

[14] D. K. Pradhan," Fault-Tolerant Computing : Theory and Techniques", Vol II, Prentice Hall, 1986.

[15] Special Issue on "Dense Wavelength Division Multiplexing Techniques for High Capacity and Multiple Access Communication Systems," *IEEE Journal on Selected Areas in Communications,* August 1990.

Proceedings of the Second Workshop on Massively Parallel Processing Using Optical Interconnections (MMPOI '95)
0-8186-7101-7/95 $10.00 © 1995 IEEE