Hierarchical FDDI—An Approach for FFOL

Bernhard Albert and Anura P. Jayasumana
Department of Electrical Engineering
Colorado State University
Fort Collins, CO 80523
e-mail: ba606909@lance.colostat.edu

Abstract

Looking toward the future of computer networking, we see that the need for increased interconnectivity will require continual growth in speed and features. Today there are several successful LANs competing in the 100 Mbps arena. Now we are considering the next generation of network. Ideally, users need a solution that provides the necessary bandwidth and connectivity while preserving their investments in hardware and software. This paper will present an architecturally different design for a new more effective LAN. The new architecture will provide a convenient solution to the need for higher total network capacity by interconnecting multiple FDDI networks (and/or network segments). The underlying network architecture is based on the FDDI protocols, which is an attractive feature for those who have invested resources in FDDI. Scalability is a required feature for any future network design and is addressed in detail by this architecture. This alternative next-generation networking solution has been presented to the ANSI standards committee X3T12 as a possible FDDI follow-on standard (FFOL).

1. Introduction

Local area networks (LANs) are everywhere in today's computing environments and are considered to be commodity products. Ethernet, with over 30 million installed nodes, has been the most successful LAN of the 80s and will continue to play a dominant role well into the 90s. Ethernet has a limited bandwidth. In many applications, such as high-definition video, the network cannot operate in a satisfactory manner. This issue has led to several Ethernet spin-offs and fostered the new generation of LANs competing in the 100 Mbps range. Fiber distributed data interface (FDDI), FDDI-II, dual queue dual bus (DQDB), asynchronous transfer mode (ATM), and two variants of faster Ethernet are the main players in this 100 Mbps market with standards and products generally available for most of these technologies.

Network traffic continually increases over time; customers also are demanding increased connectivity from their LAN providers. This has led to much more effort applied to inter-networking and connectivity. One way to expand a network throughput is to scale the operating speed of the network. ATM and DQDB are scalable, and as the push for speed continues, they probably will be the beneficiary. Ethernet is not realistically scalable beyond a few hundred Mbps and has a very limited geographical reach at that speed. The FDDI community has not chosen an expansion path beyond the current 100 Mbps standards and they are soliciting new ideas at this point. This proposal is in response to their solicitation.

It is important to consider the types of services a network must provide. Multimedia is poised to become a major user of advanced LAN services and will significantly contribute to LAN traffic levels. Typical multimedia applications are video and sound clips along with text, graphics, and internet access. Other users of bandwidth include high-definition television (HDTV) connections, medical and other imaging, e-mail, and voice-mail. The real success of multimedia will come when it is possible seamlessly to make all these information exchanges with a large number of users, similar to today's telephone network. At the same time the network also must be efficient for packet-type data services, possibly requiring multiple priority classes.

Multiple priority classes have long been used to provide services at a variety of quality levels to different users. By using priorities, it is possible to control access to the network during times of peak load and guarantee that service is provided to users in the intended sequence. In the case of severe congestion, service may be denied altogether to users with low priority. Many of today's networks, including token ring and FDDI, provide prioritized access control.

As network capacity demands grow, there are two basic directions in which to expand: increase the operating speed or interconnect multiple networks. The choice is not obvious, and involves trade-offs and limitations on the range of scalability and its applicability to different traffic patterns. In most designs the increase in operating speed approach is taken. A major drawback to this approach is that all of the equipment on the network has to be upgraded to the higher operating speed. The main drawback of the interconnected or parallel network scheme is that more routing and circuit switching are needed. Most of the work in this research project will focus on the parallel network scheme and the need for routing and circuit switching.

There are several major benefits to an architecture of interconnected parallel LANs. First, The network can be expanded incrementally without having to upgrade all nodes. Second, in many cases, much of the traffic can be isolated to a single network; this lowers the overall network traffic. Also, a single network failure does not have to cripple the global network traffic if the fault can be circumvented or isolated. Together, these features increase the performance and reliability of LAN environment as a whole.

The goal of this research project is to propose a next-generation high-speed network architecture that will allow for scalable and flexible network designs at aggregate data rates of up to several Gbps. The use of an industry standard network, FDDI, as an underlying network protocol, would give the project a realistic operating environment and a chance at widespread acceptance.

Prior work at Colorado State University has dealt with timed token protocols such as FDDI & FDDI-II [1, 2, 3, 10, 12, 17], including extensive work on priority schemes [8, 9, 11, 23, 24]. Other work has been produced focusing on performance of a network with multiple classes.

The better defined networks, such as FDDI [7, 12, 20], DQDB [22], and ATM [6, 16, 21], have been extensively covered in the literature, and a few general overview documents are listed for each. FastEthernet is in the project approval status within IEEE 802 committee and has splintered into two separate subgroups at this point. There has been some architectural work printed on FDDI follow-on network [18]. This project is too early in the development stage to be able to do any performance analysis, the subcommittee is dormant at this point, and there are no corporate sponsors. Some interesting work on characteristics of FDDI protocol at gigabit speeds has been produced [13]. The thrust of FFOL seems to be dropping the timed-token protocol in favor of a cell-based protocol, although this is not final.

Some studies have been performed on internetworking [4, 5, 14, 14], and this is an important research field that is not yet fully understood. Not everybody agrees on what multimedia is, but they do agree that it will be an important factor in the future. High-speed networks will be needed to support the increasing demand presented by multimedia applications [3, 19]. These will include but not be limited to voice, video, graphics, animation, and data.

2. Network Architecture

The network architecture presented here defines an interconnection scheme for regular FDDI and FDDI-II networks. Each individual network remains separate but has direct access to all other networks. This creates an expandable supernetwork made up of independent subnetworks. This approach has many benefits, including:

- Uses Existing and Approved Standards
- Preserves Investment in Hardware and Software
- Presents an Expandable and Scalable Architecture
- Provides Built-in Fault Tolerance
- Provides Independently Tuneable Sub-Networks
- Allows Mixing of FDDI and FDDI Networks

The new architecture is based on using the FDDI standards and defining the inter-connection topology. This creates an expandable and scalable architecture which preserves the investment that has been made in FDDI hardware and software. A big advantage of having separate networks, rather than one large one, is that a failure of a subnetwork does not cripple the entire network and that each of the subnetworks can be tuned to meet the requirements of the users it connects. Another benefit of this architecture is that it permits the mixing of FDDI and FDDI networks and the interconnection of traffic between them.

Whenever networks are interconnected, data traffic has the possibility of accumulating in the bridge or router. This could happen either if the outgoing network is busy or if the interface is between two networks of different speeds. This problem is addressed by providing an amount of buffer space, although it is difficult to determine the optimum size.

The interconnection network can be thought of as a backbone network with the exception that the subnetworks remain independent. An example with four subnetworks is shown in Figure 1. Each FDDI or FDDI-II ring (subnetwork) interfaces to the backbone through a backbone interface node, which is responsible for connecting a given endnetwork to the backbone and all bridging functions from other networks. All routing is provided by the endnetwork, which means that each backbone interface node must monitor all other channels on the backbone for packets coming its way. The backbone network transports the multiple subnetworks in an unaltered form, as channels. The channels are interleaved at a low level (bit or byte), which minimizes

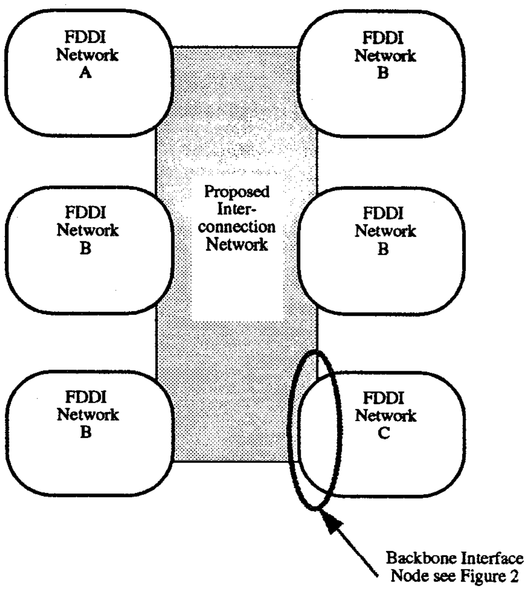

Figure 1: Example of three Connected FDDI Networks

the latency introduced into the channels. A high-level diagram of a backbone interface node is shown in Figure 2.

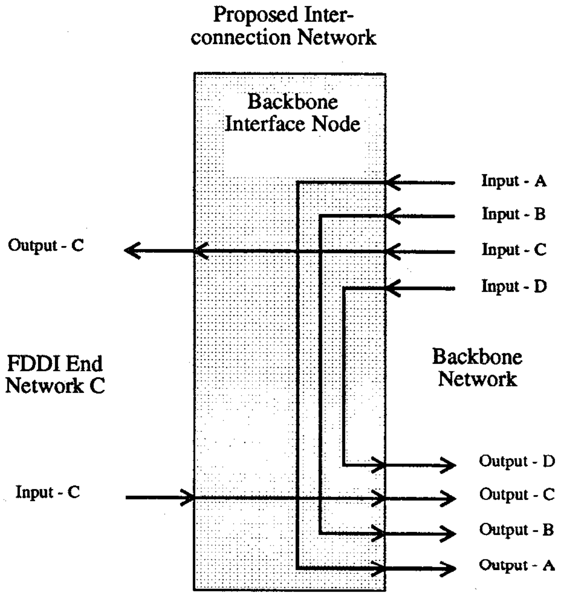

Figure 2: Network Data Flow in Backbone Interface Node

The four individual FDDI subnetworks are shown unmultiplexed in this figure to emphasize that the networks remain separate on the backbone. This figure also does not show the bridging function performed inside the node and the monitoring of the other three networks for traffic destined for this network. Since there are four FDDI networks involved, we refer to it as a 4-X network. This concept can be expanded in a hierarchical manner as shown in Figure 3.

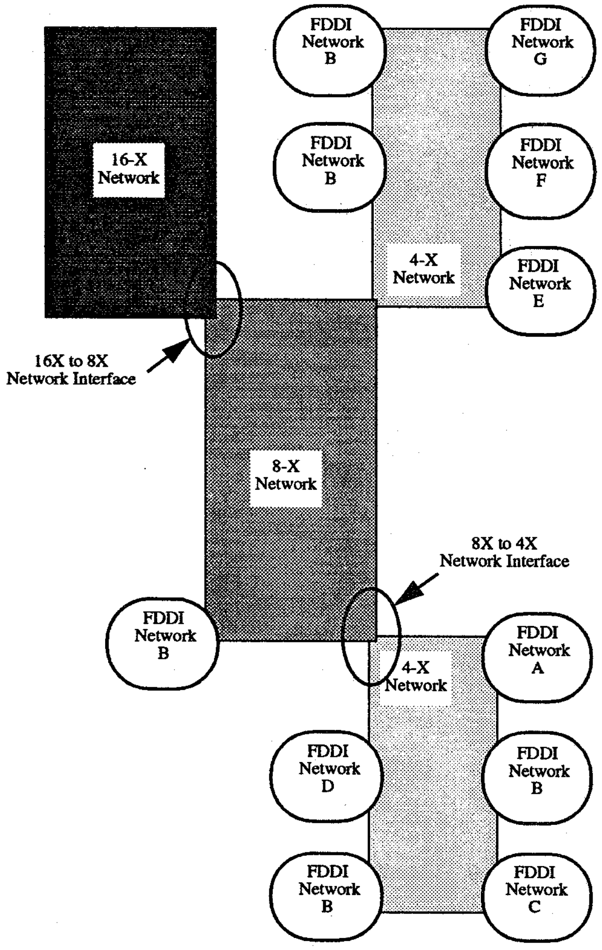

Figure 3: Hierarchical network structure.

In this example the 4-X network from Figure 1 is expanded by becoming part of the 8-X network which, in turn, becomes part of a larger 16-X network. If the 16-X network represents the most recent addition, then only the interface node to the 8-X network would be new; the rest of the networks would be existing.

Let us carry this example further. Again, there are four independent FDDI networks; to the left of the backbone interface node (Figure 2) is one plain FDDI network (Network C), and to the right is the backbone, which carries all four of the FDDI networks. By keeping the networks independent on the backbone, a failure in any one FDDI network will not affect the operation of the other networks.

The network interface nodes do not provide any routing on outgoing packets since this is the responsibility of the receiving end. A block diagram of a network interface node is shown in Figure 4. The proposed backbone network

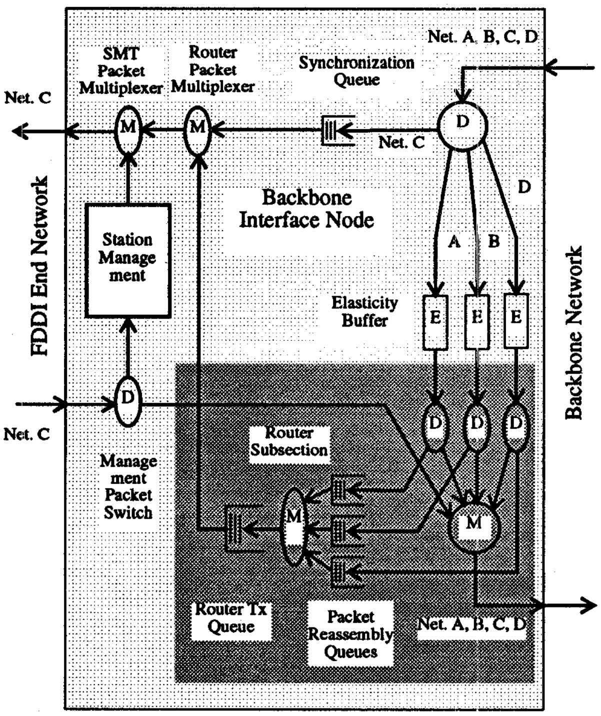

Figure 4: Network Interface Node

would be able to operate in the dual-ring architecture that FDDI requires on the main ring, but for simplicity, only the primary ring is shown. In this example the backbone network has four channels, each assigned to one of the FDDI networks. The number of networks on the backbone can be increased to at least 16 (2 GHz design limit). Path control is achieved by using intelligent multiplexers (M) and demultiplexer (D) elements.The interface node from Figure 4 is shown represented in its FDDI functional structure form in Figure 5. It shows the signal flow of the individual FDDI channels through the FDDI layers in the backbone interface node. A 4-X backbone interface node will consist of at least four FDDI nodes. One of the FDDI nodes is circled in Figure 5; this is the interface to the endnetwork (Network C), and the other three FDDI nodes monitor the other FDDI channels. In each node the FDDI signal is decoded by the PMD (physical medium dependent layer), then processed through the PHY (physical layer) and passed to the MAC (media access control layer). If the packet is to be received by this node, then it is copied to logical link controller (LLC). Otherwise it is only copied to the outgoing PHY and PMD.

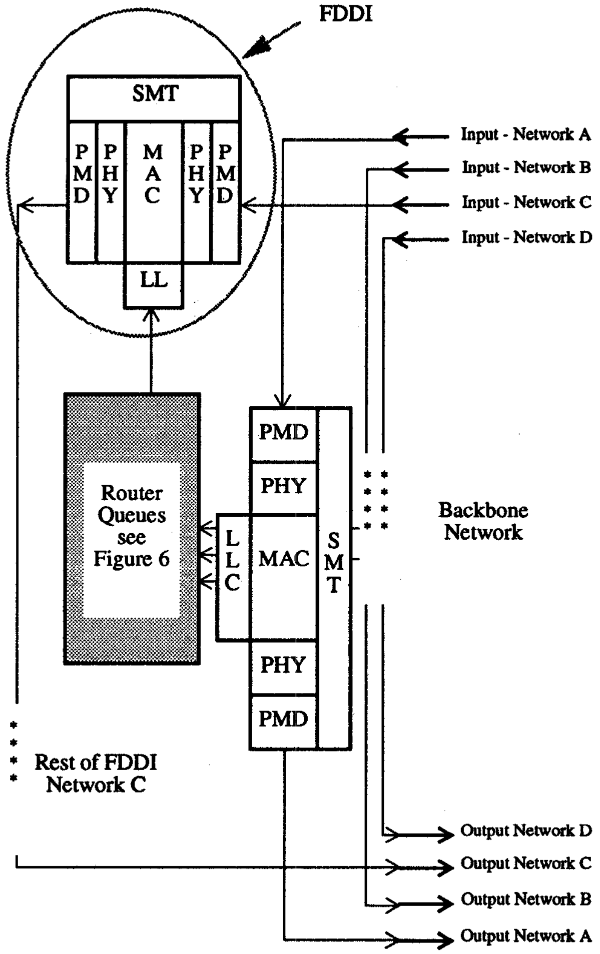

Figure 5: Backbone Interface Node

A possible implementation of the router subsection queues is shown in Figure 6. The incoming data streams are independent and possibly concurrent. This must be dealt with in the buffer allocation scheme and access control schemes. Another simplification would be the eliminatation of either the synchronous queues, multiple asynchronous queues, or both; this would be very implementation dependent.

The other detail, not shown in Figure 5, is the need to multiplex the independent FDDI streams across a single wire between network interface nodes. This does not affect the rest of the architecture and can be solved separately. Several high-speed multiplexing schemes are outlined in Figure 7. Three multiplexing methods are outlined, two electrical and one optical. In method A, the four channels are time-slice multiplexed and then demultiplexed at the

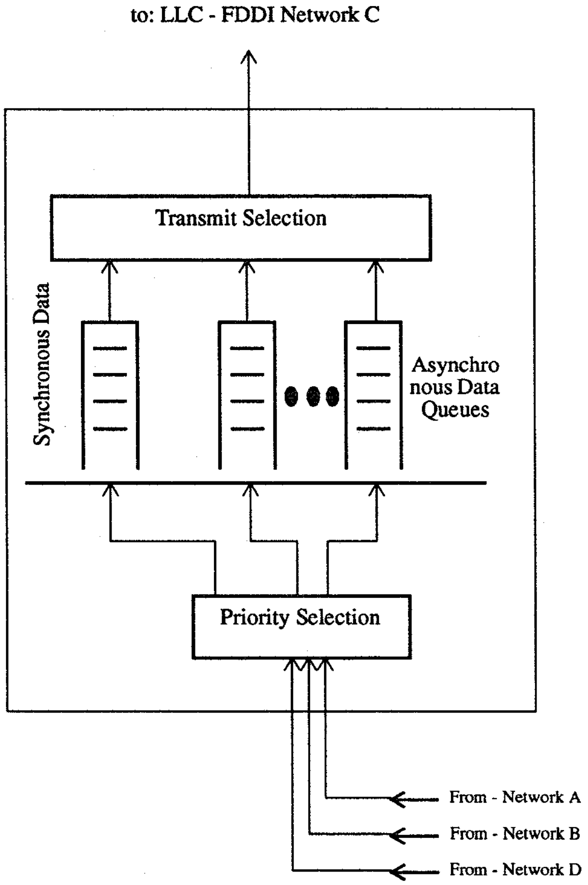

Figure 6: Router Queues

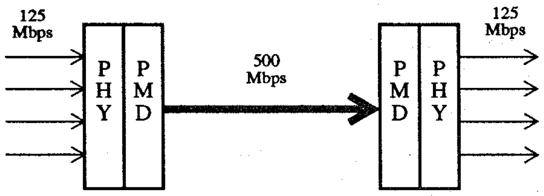

a: Using Dedicated PHY/PMD Design

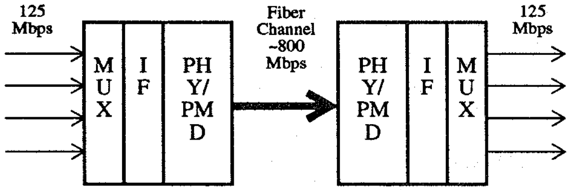

b: Using Existing PHY/PMD Design

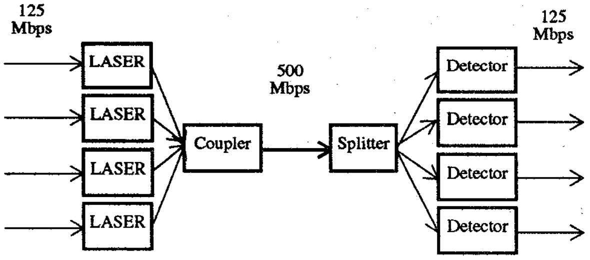

c: Using WDM Techniques

Figure 7: High-Speed Multiplexing Schemes

other end. The channel data rate would be 500 Mbps. This would require the design of a dedicated PMD for each speed variant but would otherwise be relatively straightforward since a common clock is used. The receiving end would need a PLL (phase locked loop), running at a multiple of the input signal, for local clock regeneration and a method for selecting the proper channel. With method B, an existing PMD is used that has a higher bandwidth than required. This has the advantage that no new PMD must be designed within the limit of existing PMDs. The last method, C, takes advantage of the emerging field of WDM (wavelength division multiplexing) and keeps the signal completely in the optical realm. Each FDDI channel is assigned a different wavelength and thus requires independent.

Referring back to Figure 3, we have the interface point where eight network segments get split into two paths of four segments each. At a high level (similar to Figure 2) this inner backbone node can be represented as shown in Figure 8. At the interface point four of eight FDDI channels are split off from the incoming channel streams (Networks E-H); the other four channels (Networks A-D) are passed directly to the output section. At the output the four bypass channels are recombined with the four returning channels from the 4-X network. This 8-X to 4-X interface node only needs FDDI layers PMD and PHY and can be implemented by a switch as shown in Figure 9. This switch has 12 inputs and 12 outputs with the 8 input channels and 4 return channels forming the inputs. The switch is unconstrained in its basic form but can be significantly simplified due to implementation constraints. If fixed routing is acceptable, then the entire switching circuit can be implemented on a patch panel.

To be truly FDDI compliant, the backbone portion of the network must be a dual counterrotating ring. An example involving only two channels on the backbone is shown

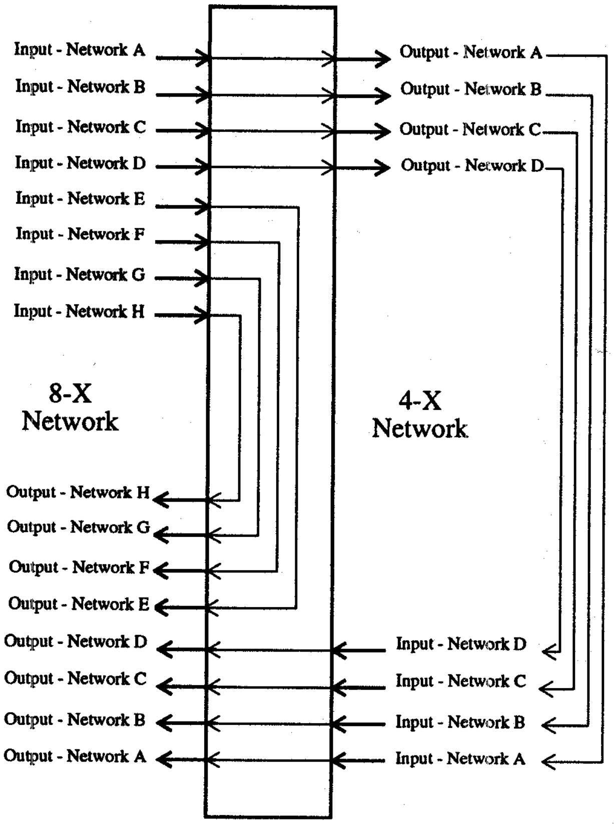

Figure 8: Inner Backbone Interface Node

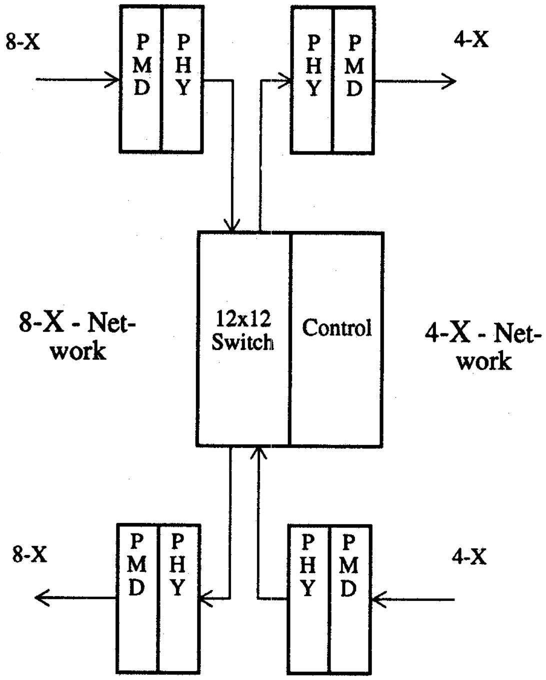

Figure 9: Inner Backbone Interface Node

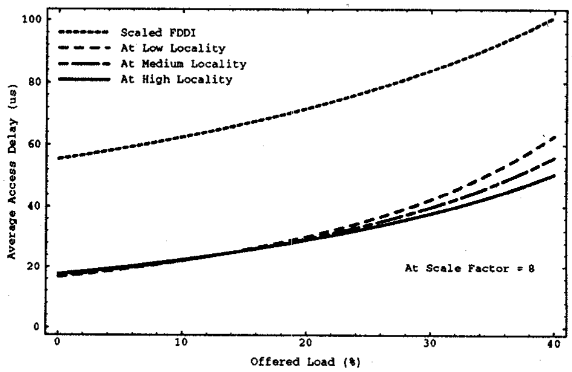

Figure 11: Average Access Delay with 8 Networks

in Figure 10. This presents some interesting opportunities to research on spacial reuse, selective routing primary and secondary channels, and more complex error recovery scheme. In standard FDDI the secondary ring does not transmit any data. It is used only as a hot standby and in case the ring is severed. In case of failure, the token is routed backwards around the failure, and network operation is restored.

Extensive work has been done on the performance of the proposed network. In the following example an 8-X hierarchical FDDI network is compared to a linearly scaled FDDI network of equal size. One of the disadvantages of the hierarchical FDDI scheme is that all network traffic that does not terminate on the originating network must be transmitted on a second network, thereby increasing the overall network traffic. This is taken into account with the locality factor. A high locality factor means that most of the traffic stays local, whereas a low locality factor implies that much of the traffic must be repeated on the second network. The average access delay versus offered load is shown in Figure 11, and the average end-to-end delay versus offered load is shown in Figure 12

The value of the proposed interconnection network is further enhanced by supporting FDDI-II compatible isochronous traffic. To implement this, the bridging section needs to be expanded by adding the hardware to perform the isochronous path routing. Although this approach is more complicated, it allows FDDI and FDDI-II networks on the same backbone. The cycle buffer replaces the packet reassembly queue from the FDDI design. The cycle buffer will have to be double buffered. This double buffering permits the unconstrained routing of isochronous streams between networks. Without double buffering, it would not be possible to route data backwards in the cycle structure, and the overall routing timing would be much more constrained. It

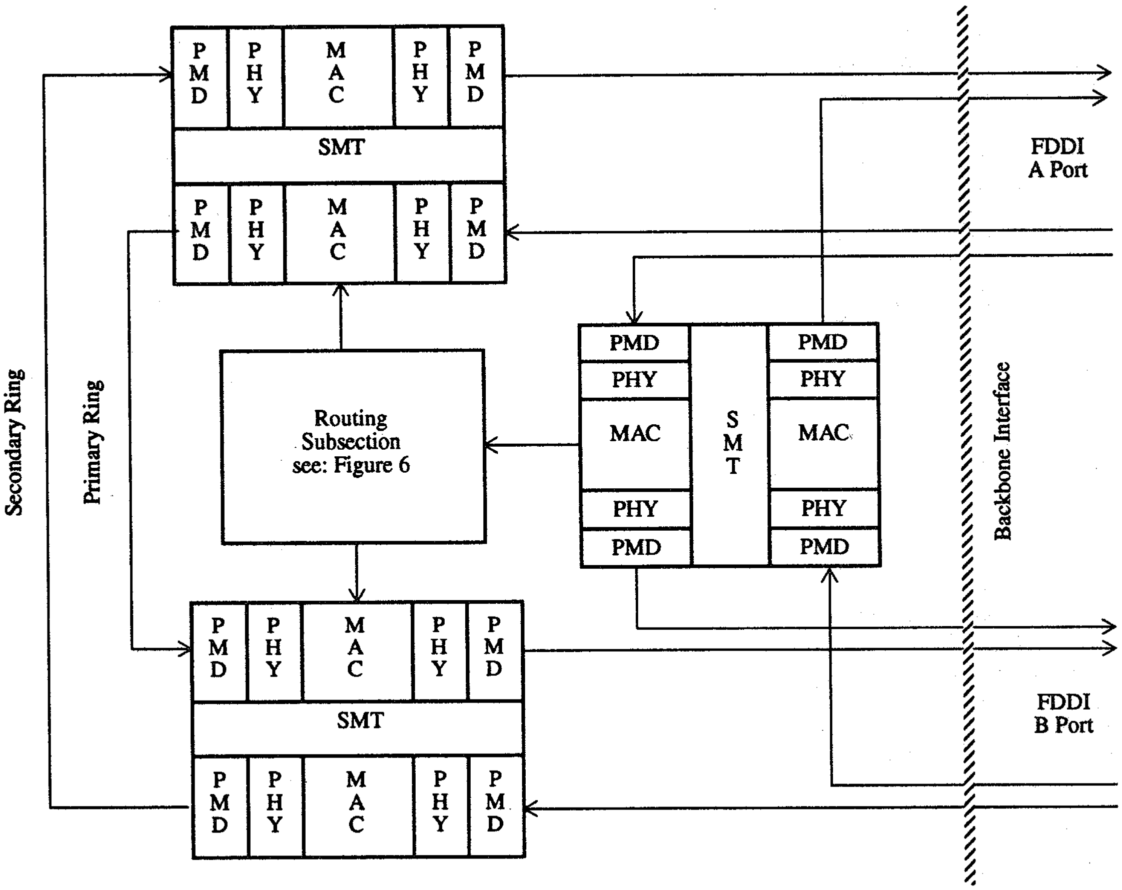

Figure 10: Dual Ring Example

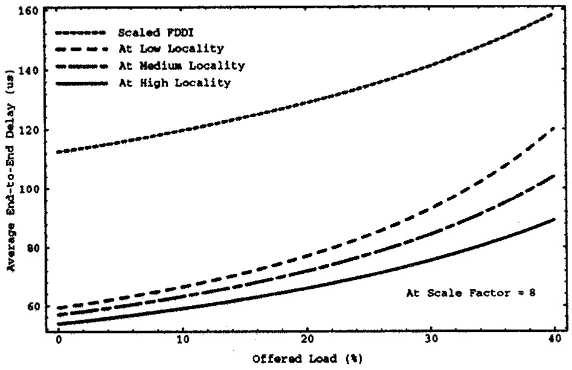

Figure 12: Average End-To-End Delay with 8 Networks

is necessary to buffer one full frame in order to route paths backwards in the cycle structure. These paths are set up at the beginning of a dialog and are maintained until the dialog is completed. It is important to note that whenever multiple networks are used on the path between source and destination, bandwidth is used in all of these networks, thereby effectively multiplying the used network bandwidth.

The protocol for managing these isochronous streams does not yet exist but is being developed by the synchronous bandwidth management forum. A detailed analysis of the block diagram for a FDDI-II compatible network interface node, the FDDI functional layer representation, and the isochronous routing diagram is given in [1].

The only queue of consequential size will be the router transmit queue. It also might be desirable to add a queue after the packet reassembly queue. This would allow the node to buffer network traffic in case of a minor network disturbance and prevent the loss of data packets. It also would be easy to implement multiple priorities in these queues, either with multiple queues or by flagging the entries in the queue and then processing them in the priority-adjusted sequence.

The station management block is needed to perform the required protocols for maintaining compatibility with the FDDI standards. An important feature of this architecture is that it is scalable at the backbone level. This allows for easy expansion since only the backbone interface nodes must be upgraded and it is not necessary to upgrade the end network nodes.

3. Summary

This paper has presented the architectural outline for an advanced high-speed LAN. The architecture presented allows any number of FDDI and FDDI-II networks to be interconnected using a backbone network. Internetwork communication (both packets and circuits) is controlled by an intelligent router section. This is a distributed approach with all routing for any given FDDI segment produced by that network segment's backbone interface node.

Much work remains to verify this architecture; this is the goal of future research. Any comments or suggestions should be sent to the authors. Thank you in advance for any comments and they are much appreciated.

References

1. Albert, B., "Hierarchical FDDI--An Advanced Local Area Network," Ph.D. Dissertation, Colorado State University, 1996.
2. Albert, B., Jayasumana, A. P., "FDDI & FDDI-II - Architecture, Protocols and Performance," Artech House, 1994.
3. Albert, B., Jayasumana, A. P., "FDDI - A Local-Area Network for Distributed Real-Time Applications," International Symposium of Intelligent Control, 1990, pp. 279-284.
4. Albert, B., Jayasumana, A., "Performance of an Integrated Voice, Data, and Video FDDI Network with Error Conditions," Journal of Data and Computer Communications, Vol. 4. Num. 3, Winter 1992, pp. 18-28.
5. Berg, B., "End-to-end Performance of Interconnected LANs," PA, 3-91, 8p.
6. Chae, K., Nilsson, A. A., "Performance Evaluation of FDDI Network and Interconnected Heterogeneous Networks," IEEE Conference on Local Computer Networks, Sept. 1990, pp. 75-83.
7. Eliazov, T. E., ..., "Performance of an ATM Switch: Simulation Study," INFOCOM '90, 6-90, pp. 644-659.
8. Jain, R., "Performance Analysis of FDDI Token Ring Networks: Effect of Parameters and Guidelines for Setting TTRT," LTS, 5-91, pp. 16-22.
9. Jayasumana, A., "Throughput Analysis of the IEEE 802.4 Priority Scheme," IEEE Transactions in Communications, Vol. 37, No. 6, June 1989, pp. 565-571.
10. Jayasumana, A., "Simulation and Performance Evaluation of 802.4 Priority Schemes," IEEE Computer Society, Technical Committee on Simulation, Los Alamitos, CA, Aug. 1987, pp. 94-101.
11. Jayasumana, A., Fisher, P. D., "TSPS: A Token-Skipping Scheme for Bus Networks," IEEE CH2149-3/85, pp. 56-63.
12. Jayasumana, A. P., Werahera, P. N., "Performance of Fiber Distributed Data Interface for Multiple Classes of Traffic," to appear in Proceedings of IEE, Part E: Computers and Digital Techniques.
13. Jayasumana, A., Werahera, P., Albert, B., "Performance of FDDI Networks under Normal and Faulty Coditions," IEEE International Conference on Communications, ICC'91, Denver Colorado, June 23-26 1991, Vol. 2, pp. 748-752.
14. LaMaire, R.O., "FDDI Performance at 1 Gigabit/s," IEEE International Conference on Communications, ICC'91, Denver Colorado, June 23-26 1991, Vol. 2, pp. 1043-1047. [HHBC]
15. Mehmet-Ali, M., "The Performance of Interconnected Ring Networks with Priority," Globecom'88.
16. Ng, J.K.Y., Liu, J.W.S. "Performance of Multiple-Ring Networks for Real-Time Communications," IEEE Proceedings of the 17th Conference on LCN, Minneapolis Minnesota, Sept 13-16 1992, pp. 426-435.
17. Ohnishi, H., ..., "ATM Ring Protocol and Performance," ICC, 9-89, pp. 384-398.
18. Ott, J., "RIB: A Register Insertion Bus Fiber Optic LAN," IEEE Proceedings of the 17th Conference on Local Computer Networks, Minneapolis Minnesota, Sept 13-16 1992, pp. 612-621.
19. Ross, F. E., Fink, R. L., "Overview of FFOL- FDDI Follow-On LAN," Computer Communications, v. 15 n. 1, 1-92, pp. 5-10.
20. Sohrsby, K. A., Austin, G. P., "Performance Evaluation of Integrated Voice, Data, Image on CSMA/CD Networks," ?, 2-90, pp. 1431-1441.
21. Takahashi, K., Suda, T., "Performance Analysis of an FDDI LAN with Synchronous Traffic," IEEE International Conference on Communications, ICC'91, Denver Colorado, June 23-26 1991, Vol. 2, pp. 736-740.
22. Takahashi, K., "New Concepts of ATM Network Performance Specification," ICC '90, 2-90, pp. 532-536.
23. Tari, F., Frost, V.S., "Performance Comparison of DQDB and FDDI for Integrated Networks," IEEE Conference on Local Computer Networks, Oct. 1991, pp. 96-105.
24. Werahera, P., Jayasumana, A., "Throughput Evaluation of an Asymmetrical FDDI Token Ring Network with Multiple Classes," Proceedings IEEE INFOCOM '90, June 1990, pp. 997-1004.
25. Werahera, P., Jayasumana, A., "Real-Time Communication in CIM Systems using FDDI," to be presented as International Conference on Technical Management, 4-93, 20p.