Fiber Distributed Data Interface attachment to System/390

by J. J. Coleman
C. B. Meltzer
J. L. Weiner

IBM has articulated a commitment to open systems. An important element in this direction is connectivity based upon pervasive industry standards. A fundamental standard that enables System/390® to participate in heterogeneous systems environments is the Fiber Distributed Data Interface (FDDI), defined by ANSI. The initial IBM offering in support of FDDI is attachment to System/390 machines via the 3172 Interconnect Controller. FDDI provides a high-performance alternative to lower-speed local area networks (LANs) for attachment of workstations to System/390 mainframes. A key feature of the 3172 Micro Channel® (MC) controller is its internal bus structure, derived from PS/2® technology. The 3172 FDDI adapter is capable of data rates up to 80 megabytes per second (MBps). This should be sufficient to support multiple FDDI LANs at their rated speed of 10 MBps. Also, because of its MC orientation, the 3172 FDDI adapter is potentially extendable to other platforms derived from PS/2 and RISC System/6000™ technology.

Introduction

The Fiber Distributed Data Interface (FDDI) is an important technology for enabling a wide variety of products to communicate with System/390®. FDDI is an ANSI standard for fiber optic token ring communication at 100 megabits per second. The standard has been defined to provide both dual- and single-ring attachment options; thus, FDDI nodes are referred to respectively as A-stations or B-stations. The dual-ring structure has been integrated into the standard in support of high-availability environments. Accordingly, an A-station dynamically switches to the second ring upon failure of the first. Many vendors have been active in FDDI definition, and numerous products have already been announced.

Interest in FDDI can be viewed from two distinct vantage points. First, due to the apparent pervasiveness of FDDI technology, many solutions in the area of interoperability can be based upon FDDI. Second, because of the large bandwidth afforded by FDDI relative to other industry-standard local area networks (LANs), it provides a basis for performance-oriented solutions.

The initial IBM support for FDDI on System/390 is via the 3172 control unit. Because the Micro Channel® (MC) is the internal bus in the 3172 unit, the interface to the

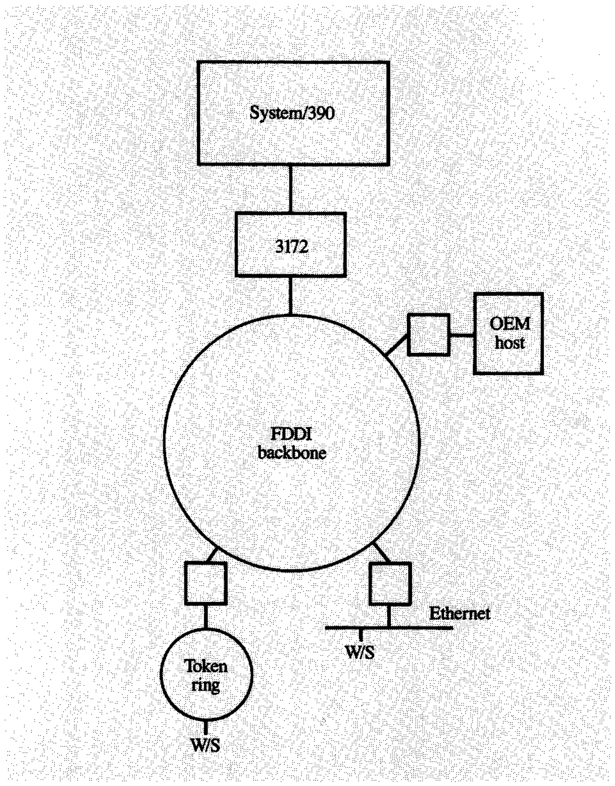

Figure 1

Industry standard backbone configuration (W/S — workstation).

MC is fundamental to the structure of the FDDI adapter described in this paper.

Please note that in the discussions that follow, FDDI bandwidth is referred to both as 100 Mbps (megabits per second) and 10 MBps (megabytes per second). The former is the raw bandwidth of FDDI after its 4/5 coding scheme is factored. The latter is the approximate usable bandwidth available to a single adapter, considering header and other adapter software overhead.

Role of FDDI

- *Interoperability*

Probably the most fundamental contributor to the attractiveness of FDDI technology is its widespread acceptance in the industry. Most major vendors have been active in defining the FDDI standard, and it is evident from product announcements and statements of direction that its pervasiveness is all but ensured. Therefore, it is important to any serious interoperability strategy for the IBM System/390.

In today's world of heterogeneous environments, interoperability makes possible a broad range of customer solutions. As an example, take mainframe-to-mainframe, i.e., host-to-host, communication. In environments where IBM and non-IBM hosts coexist, customers often have the requirement to access the data on one from the other. If both IBM and the other vendor support FDDI, the infrastructure is in place for the two to communicate. Note that cooperating access methods and applications are also required for the total system. The former is typically facilitated via the use of a pervasive network architecture such as that provided by Systems Network Architecture (SNA) or Transmission Control Protocol/Internet Protocol (TCP/IP).

FDDI also enables the introduction of System/390 into non-IBM environments. When potential customers (in this example) consider doing business with IBM, a question that frequently arises pertains to communication with hosts that already exist in their environments. FDDI is an important component in the overall manner in which that requirement is addressed.

An increasingly important interoperability focus for System/390 is host-to-workstation communication. As workstations bring more and more processing power to the desk top, the distinctions between host-to-host and host-to-workstation requirements begin to blur. Furthermore, since the IBM RISC System/6000™ (RS/6000™) processor is a relative newcomer to the high-performance workstation marketplace, many IBM customers have already invested in an assortment of non-IBM workstations.

Because of the existence of non-IBM workstations in customer environments, it is imperative that we define the appropriate communication facility. FDDI supplies an ideal base. The goal is to provide solutions based upon cooperation between the host and the workstation. Given the processing power of the workstation, users will surely want to run applications there. The role of the host then becomes to provide services to the workstation. Examples are database access, network management, software distribution, and print services. An effective communication facility is inherent in that relationship.

The interoperability afforded by FDDI is also important with respect to attachment of other media types to System/390. Assume a channel-attached FDDI facility. Also assume that a cluster of, say, Ethernet®-attached workstations must connect to System/390. By means of a gateway that supports both FDDI and Ethernet, that workstation cluster can attach to the channel-attached FDDI facility and, in so doing, attach to System/390. Several vendors offer gateways supporting both FDDI and Ethernet. Other relevant media types include token ring and some of the telephone company (telco) services.

- *Industry standard backbone configuration*

Two important applications drive the FDDI requirement. The first is the "industry standard backbone" (**Figure 1**).

Given the fact that FDDI accommodates up to two-kilometer spacing between adjacent nodes, an FDDI LAN can traverse hundreds of kilometers. This is particularly applicable in environments where right-of-way considerations are unimportant, such as office buildings and college campuses.

From the perspective of System/390, the industry standard backbone represents an element in a three-tiered hierarchy consisting of System/390, FDDI, and lower-speed LANs. The functions provided by the industry standard backbone are

- Attachment of non-IBM hosts to System/390.
- Aggregation of lower-speed LANs.
- Provision of standard interface for multivendor gateway attachment.
- Management of distributed resources.
- Facilitation of LAN-to-LAN communication.

- *High-speed LAN configuration*

The other important FDDI application is as a high-speed LAN for attachment of high-performance workstations and other devices (see **Figure 2**). This is represented by a two-tiered structure.

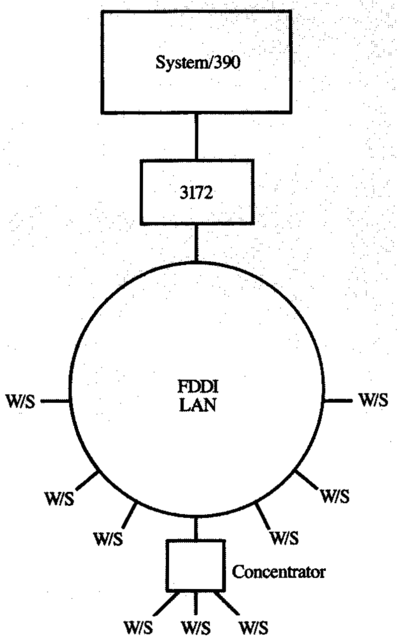

Figure 2

High-speed LAN configuration.

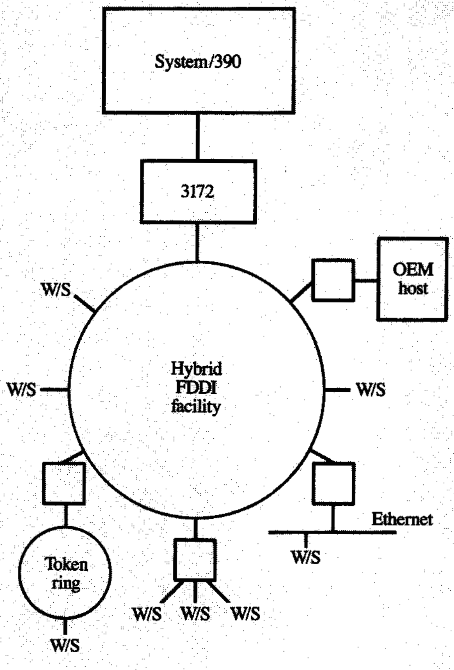

Figure 3

Hybrid FDDI facility.

FDDI as a high-speed LAN is not intended for widespread replacement of the prevalent lower-speed LANs such as token ring and Ethernet. While the FDDI bandwidth of 100 Mbps is an order of magnitude greater than those of token ring and Ethernet, attachment to FDDI is several times more expensive. This is particularly true when the user wants to take advantage of the FDDI dual-attachment capability for high availability. The decision whether to attach a workstation to System/390 via FDDI or a lower-speed LAN is based primarily upon price/performance considerations. FDDI attachment should be more prevalent with higher-priced, higher-performing workstations such as the RS/6000. The following functions are provided by the high-speed LAN:

- High-speed attachment of workstations to System/390.
- Standard interface for multivendor workstation attachment.
- High availability via the dual-attachment option.
- Facility for workstation-to-workstation communication.

- *Hybrid configuration*

In actuality, many customers are likely to merge the industry standard backbone and high-speed LAN applications into a single FDDI facility (see **Figure 3**).

Assuming that the requirement for both exists and that the bandwidth of one FDDI ring is sufficient, a single facility should suffice.

Development philosophy

- *System components*

An important element in a total Enterprise System is the connectivity media. Other major entities that must be addressed are hardware platforms, access methods, and, of course, applications. While it is impractical to take a totally linear approach (i.e., develop the media accesses, then the hardware platforms, then the access method support, and, finally, the applications), a "bottom-up" approach that puts early focus on access media while concurrently setting in motion other product development activities is beneficial. That philosophy prompted the timely development of FDDI technology by IBM.

Of course, the fact of deciding to pursue FDDI was the springboard for numerous other important considerations. A critical one concerned the characteristics of the hardware platform that would support the FDDI adapter. What should be the base for that platform? Should it be predicated on existing hardware, or is a new structure required? What are the power and cooling requirements for that platform to effectively support FDDI? What other media must coexist with FDDI in terms of a total system concept? What internal connectivity structure should that platform provide? Should that connectivity structure be a switch, a bus, or some other interface?

- *Complementary interfaces*

From a System/390 perspective, the basis for any exploitation of FDDI technology requires FDDI to coexist with a System/390 channel interface on a common hardware platform. Since the Enterprise Systems Connection (ESCON™) channel and the older original equipment manufacturer's information (OEMI) channel are both supported on System/390, either is sufficient to provide the required function. That is, specifically, what constitutes the initial FDDI offering on the 3172 control unit.

However, consideration of the industry standard backbone and high-speed LAN applications, described earlier, introduces the desirability of ensuring that the hardware platform is sufficiently flexible to accommodate other media interfaces. As illustrated with the industry standard backbone, token ring and Ethernet LANs are fundamental to that application. Furthermore, a smaller set of customers are likely to require that other LAN types be integrated into that environment. A good example is token bus, an IEEE LAN.

Another set of interfaces to be considered relative to a total FDDI system involves telco services. There are two important functions that rely on telco: *routers* and *split bridges*. A router is part of a more global network of which FDDI is only one component. It exists to provide the appropriate media interfaces to enable data to find its way through a network of routers to the destination node. The role of the split bridge is to concatenate two discrete FDDI rings via some medium, typically telco, into a single logical FDDI facility. Many customers require both routers and bridges in their environments.

Various services are provided by the telephone companies, and several interfaces relate to a total FDDI system. The higher-bandwidth telco services (T1, T2, and T3) are important because of the bandwidth provided by FDDI. A full-duplex T3 link represents about 90% of the bandwidth afforded by FDDI. Since there are environments where most of the traffic will be local to a single physical FDDI ring, some customers may require interfaces to lower-speed telco media.

- *Internal connectivity structure*

Because of the numbers of LAN, telco, and other interfaces that potentially may need to coexist with FDDI, care is needed in defining the hardware platform. Probably the most fundamental consideration is the structure internal to that hardware platform for interconnecting the FDDI adapter with those providing other interfaces. The key attributes required of that internal connectivity structure are flexibility and high bandwidth; the MC provides both.

On the subject of flexibility, assume for the moment that the hardware platform chosen for FDDI attachment to System/390 includes a new internal bus. From the perspective of developing an FDDI adapter, an interface to that new bus must be developed. This obviates consideration of using any existing technology for that purpose. Additionally, there is significant impact on the timeliness in responding to requirements for adding other adapters to that hardware platform for the purpose of providing LAN and telco interfaces to coexist with FDDI. As an example, assume that the decision has been made to support token ring. Since an interface to the new internal bus is required, use of any existing token ring adapters is impossible. Therefore, a new adapter must be developed. Clearly, this cannot be done as expeditiously as when "off-the-shelf" technology is available.

An alternative to the new internal bus, theorized above, is the Micro Channel (MC). It is tantamount to a *de facto* standard; there are literally hundreds (if not thousands) of Micro Channel-based adapters offered by a variety of vendors. There is great opportunity in supporting some of those adapters on the same hardware platform that supports FDDI. The MC is fundamental to the 3172 control unit.

With respect to bandwidth, the MC is ideal for attaching FDDI to the ESCON channel. Consider that the bandwidth of both media is 10 MBps. Allowing for flexibility in the 3172 to accommodate multiple ESCON channel adapters and multiple FDDI adapters, the required internal bandwidth is of the order of tens of megabytes per second. The 3172 FDDI adapter is capable of transferring data on the MC at 80 MBps. That mode of operation provides more than sufficient bandwidth to multiplex communication between several pairs of adapters at the rated speeds of the media with which they interface.

- *FDDI standard*

The Fiber Distributed Data Interface (FDDI) is a high-speed LAN which has been developed and approved by the ANSI X3T9.5 Standards Committee. The LAN enables the attachment of up to 500 stations through 1000 physical connections via pairs of optical fiber for a total distance of 100 km. FDDI provides a bandwidth of 100 Mbps for use as a backbone network for lower-speed LANs, a front-end network for connection of workstations or a back-end network for linking host computers and I/O devices. The FDDI standard is split into four compatible "layers":

- Physical medium dependent (PMD)—definition of medium and opto-electronic signals.
- Physical layer protocol (PHY)—definition of serial/parallel digital signals.
- Media access control (MAC)—definition of LAN data and routing frames.
- Station management (SMT)—definition of control frames and link control signals.

The SMT layer can be viewed as an umbrella over the other three layers (see **Figure 4**). Details of the FDDI standard are available in many publications, and various papers characterizing the layers have been written [1–5].

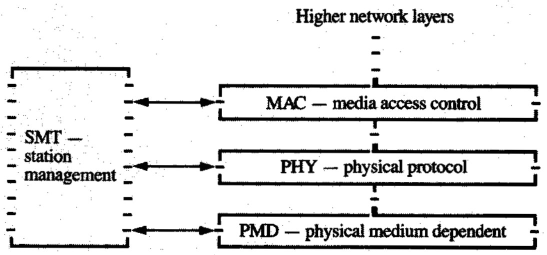

Figure 4

FDDI standard layers.

FDDI Micro Channel gateway adapter

The adapter is a card that implements an FDDI dual attachment in accordance with the ANSI standard. The adapter is intended to be used with other adapters in a PS/2® or RS/6000 gateway to perform internet protocol routing and other gateway functions. The major elements of the adapter (see **Figure 5**) are the following:

- An Intel® 80960CA processor clocked at 25 MHz. This high-performance RISC processor sets up and monitors all activity on the adapter [6, 7].
- A PHY chip is used to provide the physical sublayer functions of the standard. Two chips of this type are needed for the dual attachment. The chip is supplied by AT&T [8].
- A MAC function is provided for the media access control layer of the standard by chips supplied by the Sumitomo Electric Industries [9].
- Two optical transceivers produce the signals required for the physical medium for the dual attachment supplied by Sumitomo Electric Industries [10].
- A Micro Channel interface chip (MIC) allows the card to act as a bus master or slave with all the standard MC options through a 32-bit interface, and provides 64-bit data streaming at 8.33 MHz. MSI support logic is used between the MIC and the MC to provide extra drive capability [11]. This is an IBM-manufactured VLSI chip.
- Memories
  - 80960CA built-in 1KB RAM.
  - Program 1MB store static (PSS) fast access; zero wait state, pipelined for critical instructions and data.
  - Program 256KB store slower access, noncritical instructions and data for MIC queues, control tables, etc.
  - Data 256KB store, data store for frame data; accessed by the MIC chip, the 80960CA, and the direct memory access (DMA) hardware that moves data to/from the FDDI RAM.
  - FDDI 256KB store RAM. There is a direct mapping from the 80960CA address space to the FDDI store.
  - 128KB ROM. The IEEE 48-bit MAC address is personalized in the ROM.
- Two VLSI support chips—FDDI to/from 80960CA and MIC to/from 80960CA.
- DMA controls, clocking, error checking and self-test facilities.
- Address multiplexer and address register (DS_ADR) for the data store and FDDI chip set.
- Data multiplexer and other handling (DS_DATA) for the data store and the data interface to the FDDI chip set.

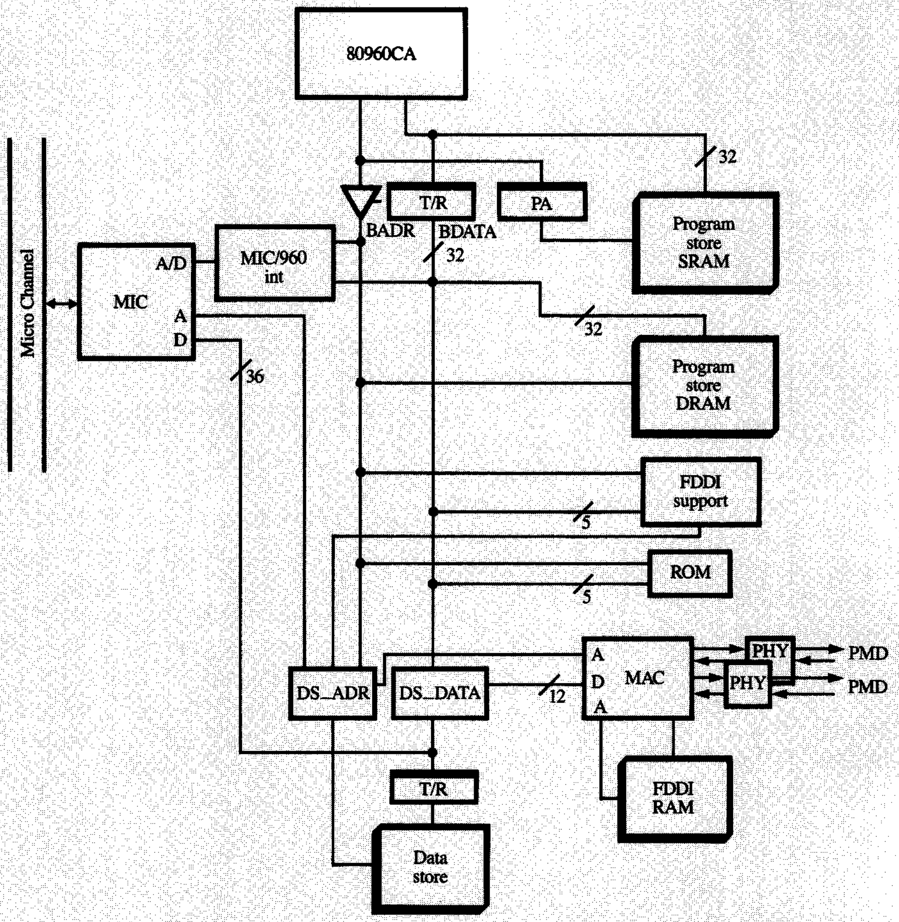

Figure 5

FDDI adapter structure.

- Address register (PA) allowing pipelined access to the program store.
- Bus transceivers (T/R).

Micro Channel exploitation

- *History*

The Micro Channel was introduced in April 1987, with the PS/2 family of personal computers. It is a high-speed bus defined by IBM, providing a general-purpose I/O attachment facility for both the PS/2 and RS/6000 product line. Accordingly, it is integral to the system structure of those two products. At first, it received strong competition from the IBM PC/AT® bus, which had become a *de facto* industry standard because of the overwhelming acceptance of the original IBM Personal Computer products (PC/XT™, PC/AT, etc.). An entire industry evolved producing add-on boards for the PC/AT bus. Many of the architectural advantages of the MC were not fully appreciated, and the original products did not exploit its capabilities in a clear and consistent manner.

Adapters which could take advantage of the MC-enhanced performance were slow to appear. At first, the only tangible benefit appeared to be the programmable option setup (POS) registers, which eliminated the need for custom select switches on the adapters and provided an automated way for the user to configure his system to avoid memory conflicts, interrupt level conflicts, etc. However, the real power behind the MC, namely alternate bus masters, would not appear in force until 1989.

- *MC operation overview*

Simplified, a normal bus cycle on the MC consists of a bus master directing one, two, or four bytes of data to a destination by first placing the "address" of the destination on the bus, waiting for an acknowledgment, and then placing the "data" on the bus. This bus cycle normally requires 200 nanoseconds (billionths of a second) and could result in 4 bytes of data being transferred in that 200 ns. A bus master directly controls transfer of data across the MC without requiring assistance from the system unit microprocessor. The resulting maximum sustained data rate, here, is 20 MBps.

When the system microprocessor is acting as the bus master, it must first obtain the data from the source memory by issuing a "read" on the MC. This cycle also requires a minimum of 200 ns, so that when the system processor is moving a block of data from main memory to an adapter, it really takes 400 ns for each 4-byte move (best case). Now, the effective throughput rate is 10 MBps. While this movement is occurring, 100% of the CPU resources are being consumed. Hence, applications which involve moving large amounts of data between memory and I/O adapters and require the microprocessor to perform these data moves make very inefficient use of the processor.

The MC offers significantly better ways to move data between adapters and memory while using potentially a small amount of the CPU to control the flow of data. The best technique involves using an alternate bus master. An alternate bus master arbitrates for control of the MC and, once granted, may burst data across the bus at speeds up to 80 MBps. This performance improvement is realized over the normal bus cycle by two techniques. The first is a concept known as *data streaming*. Data streaming allows a bus master to place the address of the destination on the bus for the first cycle; thereafter, the destination automatically increments the address. This allows a transfer of 4 bytes to occur every 100 ns, resulting in burst throughput of 40 MBps. The second technique capitalizes on the first and uses the now-idle address lines on the bus to send an additional 4 bytes of data during the same 100-ns cycle. This results in 8 bytes being transferred every 100 ns, for a burst throughput of 80 MBps.

- *MC-based hardware platforms*

Both of these significant performance techniques were exploited in the development of the custom IBM VLSI chip (MIC), described earlier. This was developed in cooperation with IBM Manassas, IBM Poughkeepsie, IBM Kingston, IBM Austin, and IBM Research at Hawthorne, New York. This chip is used in all of the high-speed adapter cards of the 3172 Interconnect Controller, enabling it to achieve significantly higher throughput rates. Its custom VLSI design enables high-speed bus master operation on the MC.

The ability to burst data between adapters at 80 MBps, plus the availability of a wide range of MC adapters, made the MC the overwhelming choice as the bus upon which to base a new high-speed interconnect controller. Additionally, since the RS/6000 processor incorporates the MC as its I/O bus, adapters that were originally developed for a single use now have the potential to be shared across a range of product families.

One of the more interesting technical considerations relates to which MC-based hardware platform is best for a specific set of applications. It would seem that trade-offs in the areas of function, performance, and cost all but ensure that both MC platforms, PS/2 and RS/6000, will assume complementary roles in supporting a broad range of IBM offerings in the heterogeneous networking marketplace.

Summary

The Fiber Distributed Data Interface is an important component in the communication strategy for System/390. As an overwhelmingly accepted industry standard, it provides the ideal base for interoperability between System/390 and a wide variety of other equipment. While not limited to mainframes and workstations, the early focus for System/390 is host-to-host and host-to-workstation communication.

Another attractive feature of FDDI is its high bandwidth relative to other industry-standard LAN technologies. Its 100Mbps data rate is an order of magnitude greater than the 16 and 10 Mbps provided by token ring and Ethernet, respectively. When the higher cost of FDDI can be justified, it provides a high-bandwidth alternative to the other two prevalent LAN technologies for workstation and other attachments.

The interoperability and performance characteristics of FDDI allow it to solve several important customer problems. One is the industry-standard backbone network, where FDDI provides the middle tier in a three-tiered attachment structure. Here, token rings and Ethernets attach to FDDI which, in turn, attaches to System/390.

Another important application for FDDI is the high-speed attachment of workstations to System/390. Here, a two-tiered structure is required because performance is critical. The workstation attaches to FDDI, which attaches to System/390. In actuality, many customers are likely to implement a hybrid environment, where a given FDDI facility will accommodate both three-tiered and two-tiered attachment to System/390.

The 3172 Interconnect Controller is the initial IBM offering for attaching FDDI to System/390. Perhaps the most important attribute of that product is its internal connectivity structure. Since the 3172 is based upon PS/2 technology, that internal structure is the Micro Channel.

The MC is ideally suited for satisfying two important requirements, flexibility and performance. This should enable a timely solution to customer requirements via "off-the-shelf" technology. With respect to performance, IBM has already developed the technology required to exploit the MC with data rates of 80 MBps (or 640 Mbps). That technology is integral to the 3172 FDDI adapter. This is very flexible, especially when viewed in relation to the 10MBps capability of FDDI. Accordingly, there is excess MC bandwidth to accommodate multiple FDDI adapters and/or to address future requirements involving higher-speed media interfaces.

One final point is the extensibility of Micro Channel-based solutions. While the FDDI adapter described in this paper was developed specifically for the 3172 Interconnect Controller, because of its MC orientation it lends itself to other platforms derived from either PS/2 or RS/6000 technology. Given the importance of both of those technologies in the marketplace, it is likely that they will evolve complementary roles in satisfying customer communication requirements.

Acknowledgments

The authors would like to acknowledge the contributions of several people to the success of the very advanced project discussed in this paper. We thank everyone in the Kingston FDDI development group. Specifically, Kurt Olin, Jeff Warren, Roger McMillan, Bruce Ratcliff, Binh Hua, Brian Smith, and Mike Wiegand were instrumental in this effort. We further wish to acknowledge key contributors from the Manassas System Interconnect Development group—Jeff Swarts, Dennis Smetana, Guy Guthrie, and Steve Fields—and from the Research Laboratory for Advanced Connectivity Engineering—Jeff Kravitz, Todd Fellela, and Nick Dono.

References

1. F. Ross, "FDDI—A Tutorial," *IEEE Commun. Mag.* **24**, 10–17 (May 1986).
2. D. Dykeman and W. Bux, "Analysis and Tuning of the FDDI Media Access Control Protocol," *IEEE J. Sel. Areas Commun.* **6**, 997 (July 1988).
3. M. Cerqueiro and K. Jabbour, "Model and Simulation of FDDI Token Ring," *Proceedings of the 21st Modeling and Simulation Conference*, Pittsburgh, May 1990, pp. 1497–1505.
4. G. Makhoul and K. Jabbour, "A Simulation Model for the PMD Sublayer of FDDI," *Proceedings of the 20th Modeling and Simulation Conference*, Pittsburgh, May 1989, pp. 825–836.
5. M. Cerqueiro and K. Jabbour, "Modeling and Performance Evaluation of MAC Layer of FDDI," *Proceedings of the 23rd Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, November 1989, pp. 914–919.
6. *80960CA User's Manual*, Intel Corp., Hillsboro, OR, 1989.
7. *80960CA Data Sheet*, Intel Corp., Hillsboro, OR, 1989.
8. *T73151 FDDI Physical Layer Device Data Sheet*, AT&T Corp., Basking Ridge, NJ, 1990.
9. Y. Kida, S. Tsuzuki, K. Takayama, A. Natsume, N. Gotoh, I. Yoshida, and K. Kitajima, "FDDI Based 100 Mbps Fiber Optic LAN," *Sumitomo Electr. Tech. Rev.*, No. 29, 109 (January 1990).
10. F. Nguyen, H. Goh, Y. Matsumura, K. Kubota, S. Hayashi, M. Kamakura, O. Akita, and K. Tanida, "Fully FDDI-Compliant Optical Transceiver Module Using Plastic Mold Package," *Proceedings of Broadband 90*, Baltimore, September 1990, p. 17.
11. *Personal System/2 Hardware Interface Technical Reference—Architectures*, Order No. S84F-98008-00, October 1990; available through IBM branch offices.

Received April 18, 1991; accepted for publication March 10, 1992

John J. Coleman *IBM Enterprise Systems, Neighborhood Road, Kingston, New York 12401 (retired).* Mr. Coleman was an IBM Senior Technical Staff Member and manager of the group that developed the subject FDDI technology. He managed many advanced projects in IBM and also served as a member of the IBM Corporate Technical Committee, reporting to the IBM Chief Scientist. Mr. Coleman received his electrical engineering B.S. degree from City College of New York, and his M.S. degree from Syracuse University. He is a Senior Member of the IEEE and has received an Outstanding Paper Award from the IEEE.

Clifford B. Meltzer *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (CBMELT at PKSMRVM).* Mr. Meltzer is a Research Staff Member and Senior Manager of the Laboratory for Advanced Connectivity Engineering in the Computer Sciences Department. He joined IBM Biomedical Systems in 1977 to work on computer-assisted analysis of electrocardiograms. In 1983, he joined the Research Division to work on an experimental videotex system. In 1984, Mr. Meltzer became interested in the close coupling of personal computers to IBM mainframes. In 1987, he received an IBM Outstanding Technical Achievement Award for his work on the PC Channel Attach adapter used in the IBM 8232 LAN Channel Station. His research in recent years has continued to focus on high-speed communications and systems. Mr. Meltzer holds a B.A. degree in mathematics and an M.S. degree in computer sciences, both from the University of Rochester.

Jeffrey L. Weiner *IBM Enterprise Systems, P.O. Box 950, Poughkeepsie, New York 12602 (JWEINER at TDCSYS2).* Mr. Weiner is a Senior Programmer in the System Architecture and Design organization, where he has worked since 1984. His interests are in System/390 communications, interoperability, networking, and related areas. Mr. Weiner joined IBM in 1981 as an employee of the Systems Communication Division in Kingston, where he worked in high-speed communications relating to local area networks, wide-band communications, and operating system extensions. Mr. Weiner holds a B.A. in economics and an M.S. in education, both from Brooklyn College, and an M.S. in computer science from the State University of New York at Albany. He has served as an adjunct faculty member in the Computer Science Department at Siena College. Prior to joining IBM, he was Manager of Data Processing at the Albany Housing Authority and led a software development project at the New York State Department of Social Services. He has also served as a mathematics teacher and is a certified guidance counselor.