Approved for release by NSA on 12-01-2011, Transparency Case# 63852

The Fiber Distributed Data Interface

STATUTORILY EXEMPT

Editor's Note: This paper was awarded first prize in the Implementing and Supporting Computer Systems category of the 1987 Computer and Information Sciences Institute Essay Competition.

The Fiber Distributed Data Interface is a newly proposed standard for fiber-based computer networks. The FDDI will operate at up to two hundred megabits per second and is designed to be fault tolerant. A large number of vendors are developing a standard FDDI product for their machines. This document explains how it works, how it can be used, and how it can be expected to develop and improve.

1. INTRODUCTION

Over three years ago, the American National Standards Institute (ANSI) began work in committee X3T9.5 to specify a standard token ring computer network based on optical fibers. The Fiber Distributed Data Interface – hereafter referred to as FDDI – is the emerging standard which sets out the results of their work. As the product of a group effort by many commercial vendors, the FDDI document can be expected to standardize all commercially developed, fiber-based computer networks (much as Ethernet has done for traditional cable networks). R531 has followed the work of the ANSI committee and has begun efforts to see the FDDI realized as a product. There are many potential uses for such a network. What follows is a brief outline of how the FDDI functions, possible applications of the FDDI to high-speed networks, the status of its development, and a summary of future directions in FDDI development.

2. AN OVERVIEW OF THE FDDI

FDDI is a token passing ring based on fiber-optic technology. The operation of the network is similar to that of the IEEE 802.5 token ring standard, with some important differences. First, FDDI is based on light fiber. It is a ring network which consists of a circular set of point-to-point fiber connections between nodes. Unlike the 802.5 ring, FDDI is counter-rotating. That means that the network actually contains two rings which pass information in opposite directions. A node may be connected to one or both rings. A node which is connected to both rings is known as a Class A (or double duplex connection) node. A node which is connected to only one ring is known as a Class B (or single duplex connection) node. Class B nodes connect to the ring via a Class A station known as a Wiring Concentrator. Class B nodes are cheaper and simpler to wire, but are less robust and reliable. As many as one thousand nodes can be connected to the network at once. Stations can be separated by up to two kilometers as long as the total ring length is less than two hundred kilometers.

~~FOR OFFICIAL USE ONLY~~

Each ring consists of a single light fiber. Since each fiber can carry one hundred megabits of data per second, the effective bandwidth of the network is two hundred megabits per second. It should also be realized that each fiber carries one hundred megabits of *data* per second. The fibers actually run at 125 megabits per second with 80 percent efficiency (the rest of the bandwidth is taken up by clocking signals). The coding scheme used to achieve 80 percent efficiency also makes it possible to use LED diodes instead of laser diodes to drive the fiber. This helps to keep costs low without sacrificing reliability. FDDI even has a special "wrap" mode, which automatically causes it to reconfigure and continue running if a link in the ring breaks. FDDI is being designed as a versatile, high performance networking standard. [1]

3. FDDI OPERATION

The most important aspect of the FDDI design is that the network does not have a central management or "watchdog" facility. Instead, each node in an FDDI network maintains its own network status model. Since control is decentralized, it is easiest to explain how the network functions by examining it during normal operation and then contrasting that with what happens when a link breaks. Other general aspects of the network which also need to be examined include the structure of the protocol layers, transmissions synchronization, bandwidth allocation, and transmission coding. Finally, the use of the second ring needs to be clarified.

3.1. Normal Network Configuration

Normal ring operation is characterized by the fact that all links are operational and no station is currently in "wrap" mode.

3.1.1. FDDI Network Structure

A fully functioning FDDI ring network will consist of the two counter-rotating fiber rings, Class A stations, Wiring Concentrators, and optional Class B stations. (See fig. 1.) The Class A stations use two duplex interfaces: one interface contains the receiver for the first ring and the transmitter for the second, the other interface contains the transmitter for the first ring and the receiver for the second. The connector cable which plugs into an interface holds two fibers, one strand each for both rings. The Class A station is known as "double duplex" because it connects to two such cables. Since a Class B station connects to only one ring, it requires only one pair of fibers. One cable is sufficient to carry two fibers, so Class B stations are called "single duplex." Both classes of stations contain optical bypass relays which passively connect the fibers entering and leaving the node in the event that the node is not operating. Notice, however, that a Class B station requires the two strands in its cable be from the same ring. This is why Class B stations must connect to the FDDI via a Wiring Concentrator, which manages the interfaces for the Class B station. The Wiring Concentrator is also able to drop any of the Class B nodes connected to it if their link breaks, thus saving the FDDI network from having to reconfigure (a Class B node connected via a single link which has broken would be disconnected from the FDDI network despite any reconfiguration). The Wiring Concentrator is also designed to cut costs by allowing connections by Class B stations through other (slower) media, such as coaxial cable.

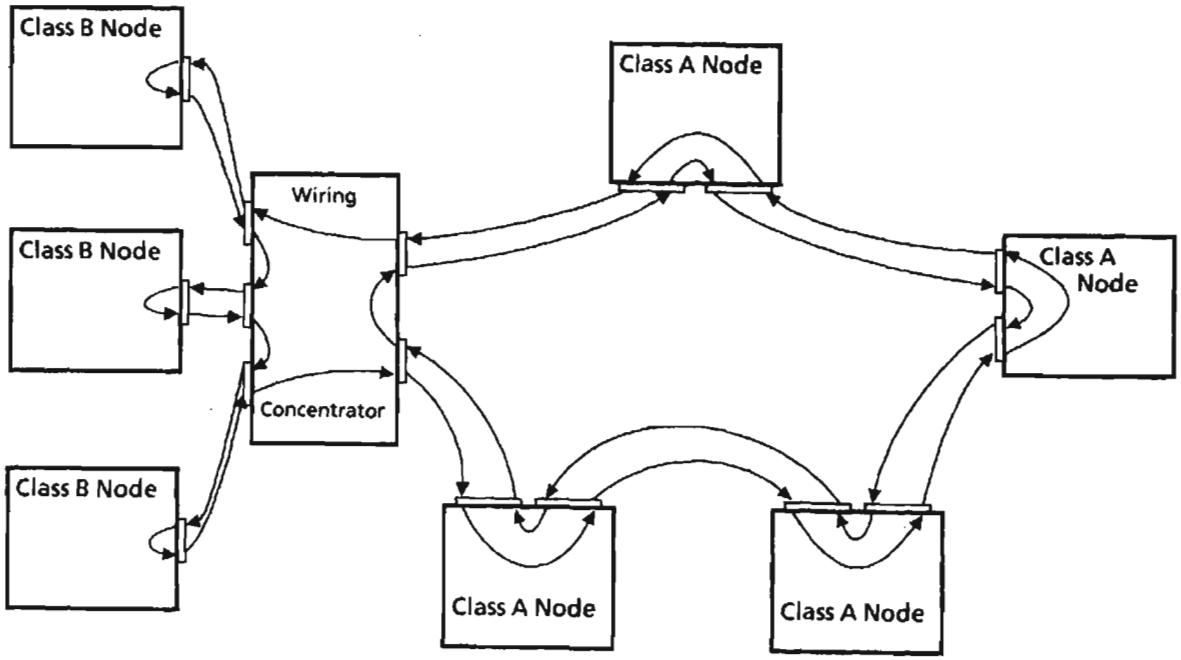

Fig. 1. A Normal FDDI Network

FOR OFFICIAL USE ONLY

3.1.2. Token Passing Operation

Under normal conditions, the FDDI functions in the following manner. A circulating token passes from one node to the next, around the ring. A node may not transmit except when it has the token. The token is a unique set of bits, and to begin transmitting, the node merely grabs the token off of the network as it passes. After transmitting the message, the node ends transmission by appending a new token to the message. When the intended destination receives the message, it flips a bit at the end of the message, indicating that the message was received and passes the rest of the message along (unless it wishes to grab the token for its own use). Notice that each node is responsible only for its own messages. When the received message is detected at the originating node, it is that node's responsibility to strip the message off of the ring. If any other station did this, the originator would have no means of determining whether its message had been received. The only duty other stations have with respect to messages they did not originate is to acknowledge those which are addressed to them and to transmit to the next node everything they receive from their upstream neighbor.

3.2. Network Reconfiguration after a Ring Fault

A ring fault occurs whenever a link between two nodes is no longer usable. (See fig. 2.) If this is not quickly detected, data will be rapidly lost and the network will break down. When a node detects such an occurrence, it begins transmitting special *beacon* packets, informing every other node to begin reconfiguration. Each node reconfigures its interfaces based upon their perceived functionality (i.e., if messages or IDLE signals can be detected on the line, it is probably working). If a link has failed, the interfaces of both nodes connected to the faulty line are "wrapped." Wrapping means reconnecting the interfaces, in effect folding the two rings together into one large ring which excludes the broken link. Two broken links cause the network to reconfigure as two separate rings, both of which continue functioning. (See fig. 3.) [2]

3.3. Distributed Management

Each station is responsible for a number of network functions which would ordinarily reside in a central control unit. Since these functions are implicit to the operation of the token ring, outlining their operation enhances understanding of how the network operates.

3.3.1. Protocol Layers

For the sake of modularity, reliability, and clarity, the FDDI network protocols have been broken down into the following layers:

a. Physical Layer (PL). The Physical Layer specifies the physical parameters of the network, including the fibers, connectors, driver and receiver devices, encode/decode methods, and clock recovery and synchronization mechanisms. The Physical Layer is divided into two sub-layers: the Physical Medium Dependent (PMD) and the Physical Layer Protocol (PHY). PMD specifies the digital baseband communications facilities up through the transceivers. [3] PHY provides clocking, encoding/decoding, and medium initialization. Also, it is PHY which provides connection to the Data Link layer. [4]

DOCID: 3928967

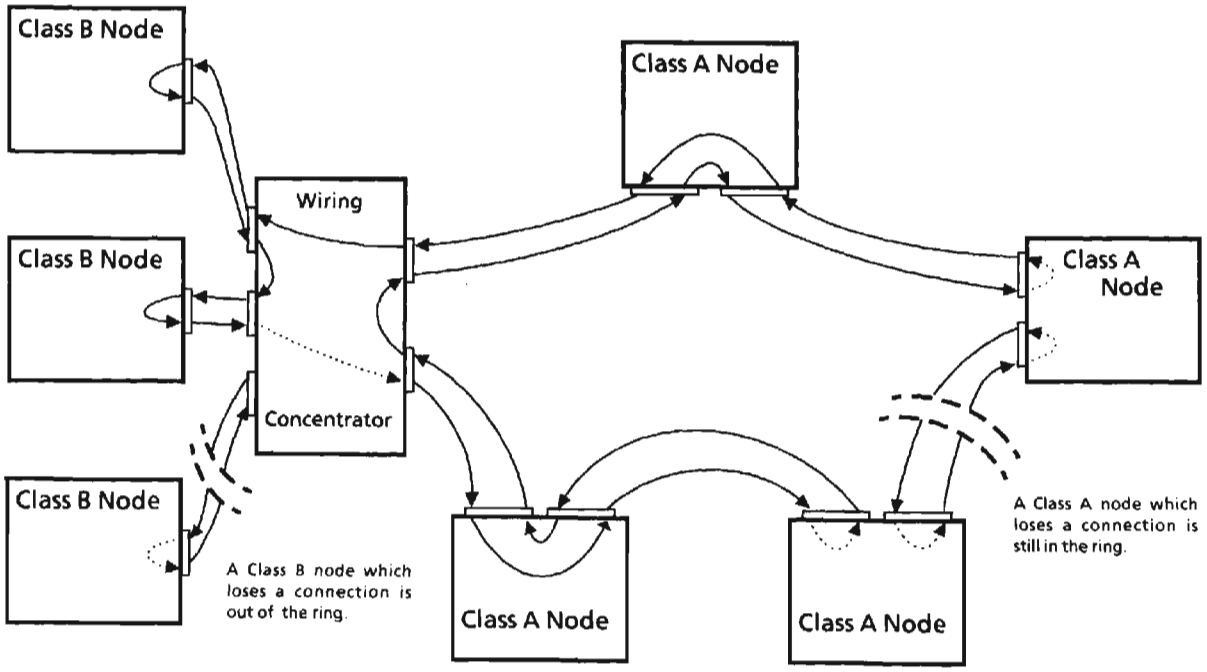

Fig. 2. An FDDI with Ring Faults

FOR OFFICIAL USE ONLY

DOCID: 3928967

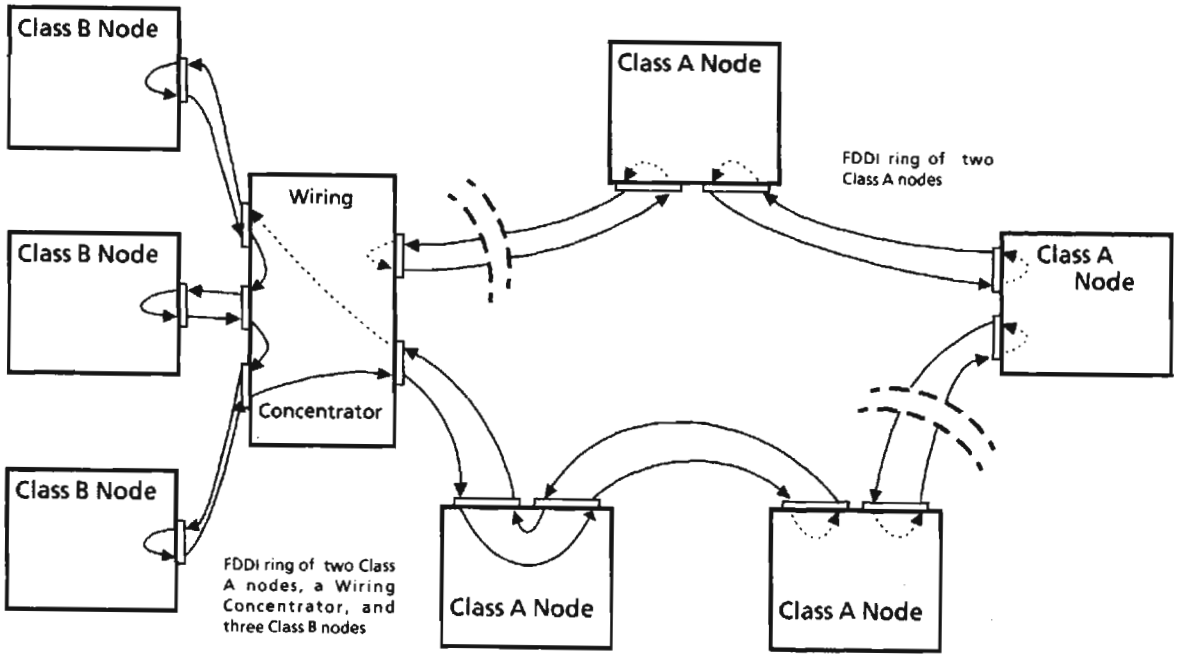

Fig. 3. FDDI with Two Ring Faults

b. Data Link Layer (DLL). The Data Link Layer is concerned with establishing reliable connections, keeping track of the addresses of the nodes on the network, controlling access to the medium, and assuring that data is delivered reliably. Data Link builds on the facilities provided by the Physical Layer, and (in turn) offers its services for use by the Station Management protocol. The Data Link Layer is divided into Media Access Control (MAC) and Logical Link Control (LLC). MAC controls the delivery and reception of frames, including the generation and verification of Cyclic Redundancy Codes. [5] LLC interfaces MAC services to the Station Management layer.

c. Station Management (SMT). Station Management is the section of the FDDI protocols which deals with getting the nodes on the ring to cooperate as a network. This includes the operations of adding new stations to the network and removing stations when they wish to disconnect. It is here that nodes decide whether to wrap when they detect error conditions on their links. In addition, the Station Management protocol is in charge of bringing links back up into the network as they come back into service. [6]

3.3.2. Ring Synchronization

Remember that each link in the ring operates as a point-to-point connection. There is no central facility for clocking. This implies that each link may be running at a slightly different speed. Not only that, but any two links are not guaranteed to have their clocks synchronized. A node could conceivably be receiving data off of one ring connection at one speed, and sending it down the next at a different speed at a later time. To do this, each node maintains a latency buffer, where bits can be held pending transmission. It also must negotiate with each of its neighbors (on both rings, in the case of Class A stations) to agree on a transmission speed. This is handled by the Station Management layer.

3.3.3. Bandwidth Allocation

Even when the network is connected and operating, there still must be a way to resolve how the network resources are to be shared. The token is the primary means of determining which node has access to the network, but there are further divisions of the allocation technique. Primarily, the bandwidth of the network (as a number of bits/second that get transmitted) is divided between synchronous and asynchronous traffic. Synchronous traffic has a guaranteed data rate set aside. A node with a synchronous bandwidth allocation is guaranteed to be able to transmit when the token arrives. Applications such as voice or video data are good examples which require synchronous bandwidth. After nodes have negotiated their synchronous allocations, whatever bandwidth is left over is open for asynchronous use.

Asynchronous bandwidth is first distributed on a priority basis. FDDI supports eight levels of message priority. As a message circulates around the network, the stations through which it passes mark off the priority of their next message in three of the header bits. It is then easy to pick up the next message with the highest priority. This may still be inadequate for some facilities, however. A program which needs to stream data at high rates may not be able to wait for the token to return. FDDI supports stream data transfer through its "restricted tokens" facility. Two nodes which wish to hog the asynchronous bandwidth merely need to set up the transfer and can then use it as long as is necessary. The "restricted tokens" facility is intended for use mostly in back-end networks, where a computer needs to communicate with a disk drive over a high bandwidth, streaming link.

The leftover bandwidth that is available between restricted sessions is available to any other application that can grab it. To make access fair, the FDDI has built into it a Token Rotation Timer (TRT), along with an optimal Target TRT (TTRT) value. A node which grabs the token can only transmit on the condition that the TRT is less than the TTRT. Otherwise, it must assume that the network is experiencing high loading and relinquish the token.

3.3.4. Coding

FDDI makes a significant reduction in cost by using a new coding technique. Most networks use a technique known as Manchester encoding, which is 50 percent efficient. A more complicated technique known as 4B/5B – which is 80 percent efficient – is being applied to FDDI. Coding is necessary in order to send data and clock information in the same signal. Manchester encoding accomplishes this by sending a clock transition for each "1" data bit, and an inverted clock transition for each "0" data bit. 4B/5B maintains that the clock can be recovered without being sent with each bit. Instead, groups of four bits are encoded with the clock to form five-bit "symbols," which are transmitted. In order to send one hundred megabits of data per second using Manchester encoding requires an LED switching speed of two hundred megabits per second. LED technology cannot run at this speed without prohibitive cost. An LED speed of one hundred twenty-five megabits per second is sufficient to send one hundred megabits of 4B/5B encoded data per second. [7]

3.4. Use of the Second Ring

One area that has yet to be resolved is the use of the second ring. The current standard leaves the use of the second ring open to the hardware or applications developer. If the second ring is fully used in a manner similar to the first ring, the aggregate bandwidth of the FDDI network can be doubled to two hundred megabits per second. Notice, though, that if the ring ever faults, the network which is created after wrapping around the broken link will effectively only have a single fiber, with one hundred megabits of bandwidth (meaning that half of the applications which had previously been running would no longer be supportable). At the opposite extreme, the second ring can be left unused until the ring breaks. It is used, in effect, merely as a backup for the primary ring. There are no traffic problems, since the primary ring only carries one hundred megabits of traffic. In between the extremes are designs which stipulate that the first ring will run primary nodes, while the second ring will only run secondary nodes. Secondary nodes would be expendable: in the event that a ring breaks, the secondary nodes get off of the net, and the primary nodes take over whichever ring still operates. It is expected that the standard will not stipulate a single usage for the second ring, but leave the field open for creative solutions.

4. APPLICATIONS OF THE FDDI

The outstanding feature of FDDI networks is their tremendous bandwidth. They have the ability to carry large amounts of data in small amounts of time. They have deterministic load performance. That is, every node gets a fair share of the bandwidth even when the network is extremely loaded. The cabling is much more TEMPEST secure (fiber does not radiate like coaxial cable). Finally, an FDDI network can be up to two hundred kilometers in circumference. It would not, however, be as cheap to install as

Ethernet. The type of networks which could best benefit from FDDI include back-end networks, LANs which connect high performance workstations, and backbone facilities for an entire business or academic community. (See fig. 4.)

4.1. Back-End Network

A back-end network is the connecting medium between a computer and its peripherals (disk and tape drives, printers, other computers, etc.). The requirements of back-end networks include the following: they must be highly reliable, they must have a very high bandwidth for rapid data access, and they must be capable of extended serial transmissions (such as writing to and from the disk). FDDI meets the first requirement because its optical bypass relays keep the network connected even when some nodes are turned off. FDDI was designed to meet the speed requirement. The third requirement, that of extended serial transmissions, can be accomplished by using FDDI's restricted token access. FDDI is capable of connecting a large number of peripherals, most of which will be Class A stations. Distance is usually not a problem, since a computer and its peripherals are normally located together at one site.

4.2. High Performance LAN

A high performance local area network (LAN) is any network which requires a high data rate for graphics, text, signal, or other data transmissions. Most nodes on this type of network will be engineering or high performance workstations, with a small number of data storage and computation nodes supporting their operation. The primary requirement for a network in this configuration is high data throughput; the workstations may require multi-megabyte file updates from the storage or computation nodes every few seconds. While the number of workstations may not be large, there will be a lot of traffic between them and their servers. Again, most of the nodes will be double duplex connections for reliability. FDDI meets these requirements.

4.3. High Speed Backbone

The idea of having a backbone springs from the normal course of development that computer networks seem to follow. An Ethernet may be installed as an inexpensive way to connect an office. Engineering workstations may be put on a token ring network to share resources. The main computer complex may be several floors down. Somehow, all of these networks should be connected: the engineering workstations need to draw on the computational power of the main computer, the office machines need to store their data on a central storage facility, and the main computer operators need to be able to contact all of the machines in the building. The solution may be to run an FDDI backbone through the entire building. Every network can maintain one machine as the gateway to the FDDI backbone. That is, one machine on each network is also connected to the FDDI backbone, and traffic which needs to go to another network goes through that gateway, onto the backbone, through another gateway, and onto the desired network. For simplicity in wiring, a number of gateways on each floor of the building can be connected to a Wiring Concentrator, which is connected to the FDDI network. The only real requirement here is that a variety of interfaces will be needed. Each network must have a machine for which an FDDI (or Wiring Concentrator) interface has been (or can be) constructed. There is also the touchy problem of protocols and protocol conversion, but that goes beyond the

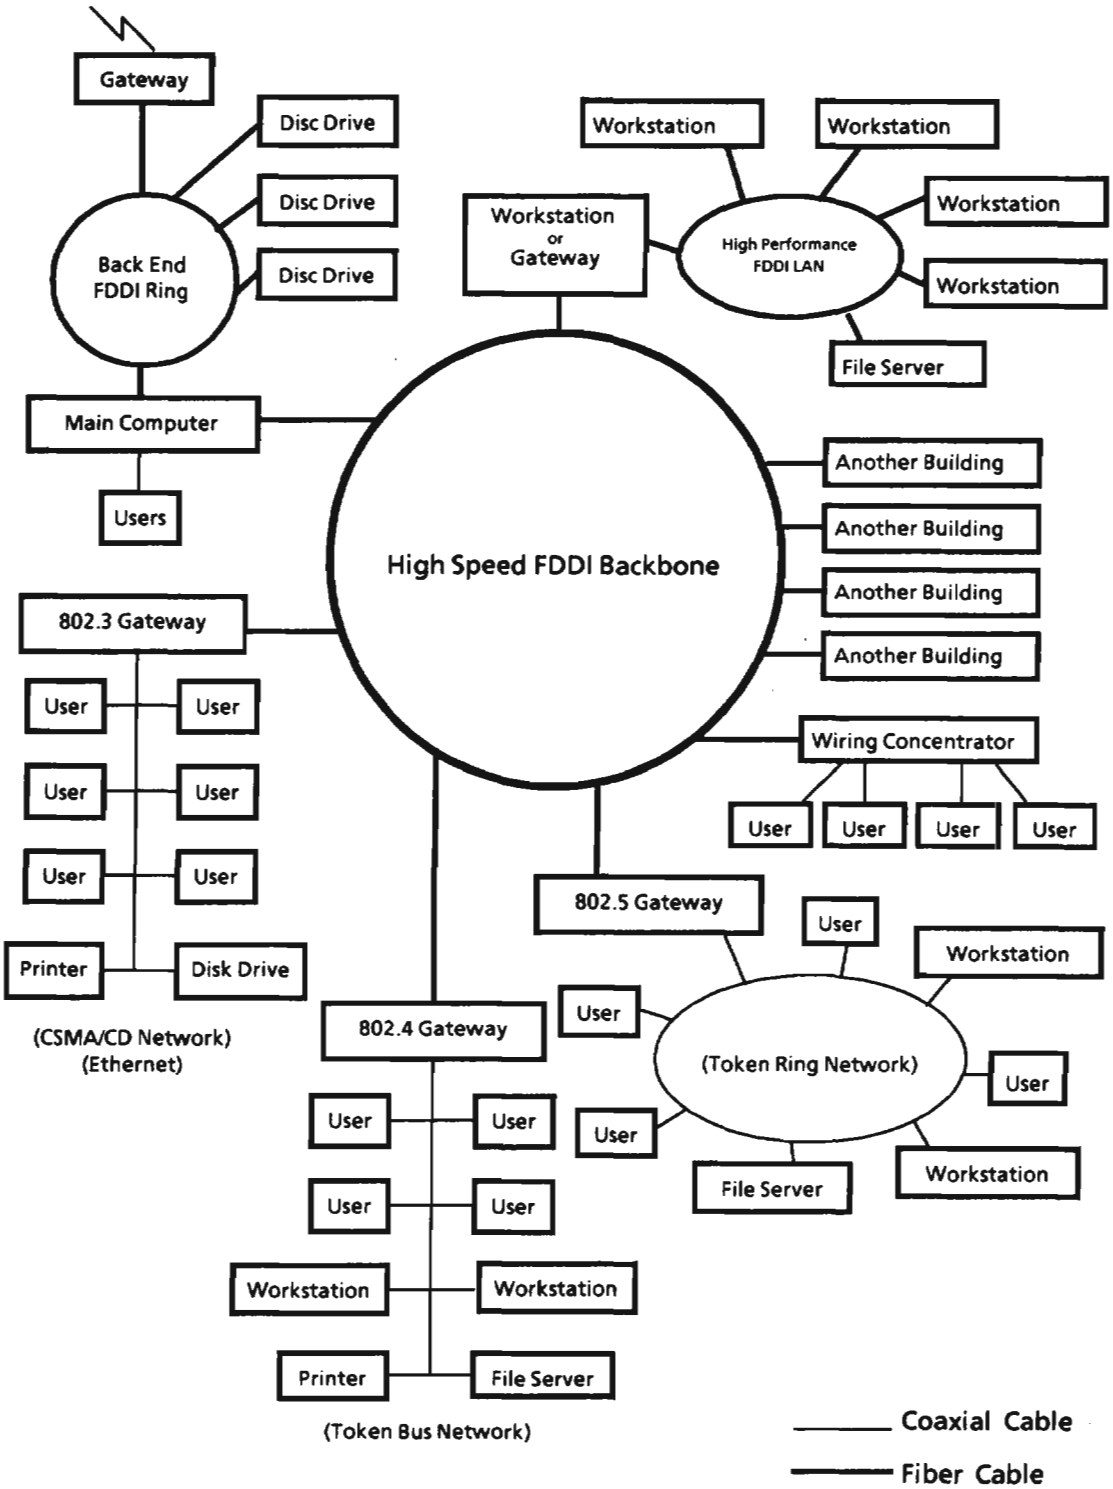

Fig. 4. Uses of FDDI Networks

scope of this document. Suffice it to say that interconnecting heterogeneous networks is a problem which FDDI does not address but may help solve.

5. THE STATUS OF THE FDDI

At present, FDDI only exists as an incomplete ANSI standard. Everything except the Station Management protocol has been finished. The finished protocols are now in various stages of review in preparation for acceptance as international standards. It could be as long as a year, though, before the Station Management protocols are completed. The reasons for this are that many controversial issues were deferred for inclusion in the Station Management document. Many vendors are willing to begin hardware development work now but can not get the support because the protocols are not finished. However, many of them refuse to agree on the protocols because they want to skew the protocols to favor their own development. Also, they often cooperate less than they should for fear of their competitors. The idea of establishing an exclusive de facto standard by finishing development first may be one of the few motivations of committee members.

There is, therefore, a good possibility that FDDI products will come to market before the standard is complete and accepted by the International Standards Organization (ISO). These first products will most likely be based on a chip set being worked on by Advanced Micro Devices (despite some skepticism about slipping AMD production schedules). The serious flaw of these early products will be that they incorporate an incomplete or extrapolated Station Management protocol. The vendors will disagree about just what constitutes a minimal operating set of Station Management controls. For this reason, the early FDDI products will most likely not be able to interconnect with the FDDI products of other vendors. They will also be limited as to how many stations they can support and the configurations that will be allowed. The special issues of Station Management in Wiring Concentrators will probably push back their availability to some later date after the initial FDDI product is brought to market. The early products will quickly converge to a common Station Management specification, and this, most likely, will be what emerges as the standard Station Management. Interest in FDDI will increase among third-party hardware vendors, and as the technology reaches greater acceptance, an increasing number of lower-cost products will become available. This scenario is taken from the way Ethernet reached maturity.

In the future, we will also see the emergence of FDDI-II, on which work has already begun. FDDI-II embodies all of normal FDDI services, with the added ability to support circuit switching for time-sensitive data (such as voice or video). FDDI-II dynamically partitions the available bandwidth between packet and circuit switched services, for optimal bandwidth utilization. National Semiconductor seems to be leading development of FDDI-II, but no implementations are expected for several years.

6. CONCLUSION

FDDI is still in the future, but it may be available on a developmental basis within a year. That is, for organizations which are willing to pay for development costs, most vendors will claim that a working prototype could be delivered within a year. Commercial availability is still about two years in the future. For the Agency, this means that a very high bandwidth network will be commercially available to interconnect almost all computers within two years. R531 is beginning investigations with the hope of making the technology available before it becomes commercial. FDDI holds great

DOCID: 3928967

STATUTORILY EXEMPT

promise for solving the Agency's long term computer networking and communications needs. In addition, it offers the chance to migrate from traditional, proprietary networking solutions which are still in use (such as Network Systems Corporation's HYPERchannel) to less expensive and more reliable means of communication and interconnection.

REFERENCES

[1] Ross, Floyd E. and William E. Burr. "The Fiber Distributed Data Interface: A Proposal for a Standard 100 Mbit/s Fiber Optic Token Ring Network," prepublication draft: presented at FOC/LAN 1984.

[2] Iyer, Venkatraman and Sunil Joshi. "FDDI's 100M-bps Protocol Improves on 802.5 Spec's 4M-bps Limit," EDN, 2 May 1985.

[3] American National Standards Institute (X3T9.5 Committee), Draft Proposed American National Standard: "FDDI Token Ring Physical Layer Medium Dependent (PMD)," 10 July 1986.

[4] American National Standards Institute (X3T9.5 Committee), Draft Proposed American National Standard: "FDDI Token Ring Physical Layer Protocol (PHY)," 22 August 1986.

[5] American National Standards Institute (X3T9.5 Committee), Draft Proposed American National Standard: "FDDI Token Ring Media Access Control (MAC)," 28 February 1986.

[6] American National Standards Institute (X3T9.5 Committee), Draft Proposed American National Standard: "FDDI Station Management (SMT)," 5 September 1986.

[7] Joshi, Sunil and Venkatraman Iyer. "New Standards for Local Networks Push Upper Limits for Lightwave Data," *Data Communications*, July 1984.