

Fiber Distributed Data Interface Overview

By William R. Hawe, Richard Graham and Peter C. Hayden

Abstract

After exploring various alternatives to second-generation local area networks (LANs), Digital selected the fiber distributed data interface (FDDI) system. FDDI implements the International Standards Organization (ISO) physical layer and the media access control sublayer of the data link layer. This system is based on a 100-megabit-per-second fiber-optic ring network and uses a timed-token protocol to coordinate station access to the network. Digital has developed the FDDI base technology, including very large-scale integration (VLSI) chips and software. These chips, licensed to Advanced Micro Devices and Motorola, Inc., provide high-quality alternatives in the market and foster cost reduction. Digital's implementation of FDDI, including backbones in extended LANs, as well as high-speed interconnection of workstations, servers,

As the use of local area networks (LANs) continues to grow at an exponential rate, many large networks with Ethernet backbones are reaching their usable capacity. In addition, the explosion in the use of high-performance workstations is placing increasing demands on network performance as larger volumes of data pass from station to station. Several years ago, Digital recognized this growth trend and began to plan and develop a second-generation LAN that would follow Ethernet and provide an evolutionary path to higher performance. The selection of FDDI as the second-generation LAN was made with great deliberation. This paper explores the criteria for that choice and the history of the FDDI system to the present. The theory of FDDI operation, the development of the FDDI technology's role in Digital's networks, and the resulting products are also presented and discussed.

and central computers,
makes available a complete
range of system products.

Introduction

Digital Technical Journal Vol. 3 No. 2 Spring 1991

Fiber Distributed Data Interface Overview

Selection of FDDI

Many of the same criteria originally used to select Ethernet were again used to evaluate the application environment for the new

LAN. The need to consider migration from the popular LANs currently in use presented the only new concern. Paramount among the reasons for selecting the FDDI technology were its tenfold increase in bandwidth over Ethernet, its consistency with other IEEE 802 LANs, and the standardization effort already begun in the American National Standards Institute (ANSI).

It is important when developing a new LAN technology to be sure the differentiation from current capabilities is

sufficient to warrant the necessary investment. Moreover, a new LAN is a significant investment for a customer and should offer a large increase in capabilities such as speed and throughput. Without this increase, the technology will have a short life span (a few years) and technology such as parallel use of existing LANs to double capacity will be a realistic alternative to a wholesale replacement of the LAN. However, it is

solution will not integrate well in the computer interconnect environment and will not be cost effective for wide-scale use.

FDDI, with its tenfold increase in speed, provides significant differentiation from Ethernet/802.3 and current token ring and bus technology to justify the new investment. Examination of the clocking, buffering, and state machine needs of the media access control (MAC) sublayer of the data link layer also showed that the FDDI technology could be implemented in several VLSI chip components. Further, as silicon technology improves, cost reduction is possible, enhancing the longevity of the FDDI LAN technology.

important not to take such
a large technological step
that exotic and complex
implementation constraints
become necessary. A LAN
that does not lend itself
to a very large-scale
integration (VLSI) logic

2 Digital Technical Journal Vol. 3 No. 2 Spring 1991

Overview

Migration is another important factor in the selection of a new LAN. Many devices exist with embedded LAN interfaces that will never be directly connected to any other LAN. Still other devices will not benefit from the added capabilities of a higher speed interconnect. Examples of such devices include a processor or workstation too slow to send or to receive data any faster from the network, an output-limited print server, or communications servers with other, more constraining, I/O ports. It would never prove cost effective to upgrade these devices to a new higher speed LAN interface but, as a group, they would need to obtain unconstrained (no bottleneck) connectivity to the services of the new LAN for smooth migration allowing protection of the investment in devices and LANs already in place.

Standards are important for networks as a way to ensure consistent interface compatibility

and interoperability for communications services. As the IEEE 802.1d standard readily demonstrates by attempting to interconnect dissimilar LANs at the MAC sublayer, some standards are more compatible than

between LANs. Ethernet, the forerunner of IEEE 802.3, does not use the IEEE 802.2 formatted LLC and, therefore, migration of those frames is more challenging.

Lastly, the media selected for the new LAN has to be consistent with current and projected future needs. Even for slower speed LANs, fiber-optic media is gaining in popularity because of its superior qualities in spanning greater distance, its noise immunity, and its declining user cost.

The FDDI technology meets the necessary selection criteria as an emerging American National Standards Institute (ANSI) standard using fiber and allowing other media in place of fiber while providing a tenfold increase in speed. Migration of some devices could be affected directly by changing controllers, and bridging between LANs could allow smooth migration of all existing devices.

FDDI History

Both the ANSI FDDI standards and the industry-wide implementations of these standards have evolved slowly. A variety of factors

others. A common logical link control (LLC) format or the format within the MAC frame allows a smooth migration between LANs by allowing a transparent bridge to provide protocol-independent translation

have contributed to this course of development. The FDDI ring was originally invented at Sperry and Burroughs Corporation. The ring was to be used as a machine room interconnect between processors and

Fiber Distributed Data Interface Overview

storage systems, much like the Computer Interconnect components are used in Digital's VAXcluster systems.[1] The timed-token, media access control protocol itself was first publicized in 1982 by Bob Grow while he was at Burroughs.[2]

As a machine room interconnect between processors and storage systems, the initial ANSI standard requirements on the FDDI technology were quite different from today's needs. In particular, as a machine room network, the number of stations was assumed to be relatively small compared to a LAN and unstructured

cabling was to be used. Since most machines were always running, fault recovery could be accomplished by having a dual ring with failover to the secondary ring. Thus, a failed station or cable could be isolated without partitioning the ring. The FDDI technology retains this property today. However, that basic operation capability is insufficient in a LAN environment with structured cabling requirements and a large number of stations, any number of which might be unplugged or turned off by users. Therefore, we have expanded the

In 1982 the FDDI technology was brought to the attention of the ANSI X3T9 committee, which develops standards for I/O interconnects and channels. Since FDDI was intended to be used as a machine room interconnect, this committee was the appropriate arena for study. Over time, however, as the need for a 100-megabit-per-second LAN emerged, the FDDI technology evolved into a local area network. Some classic standards territory conflicts developed between IEEE 802, the group that defines all the LAN standards, and this ANSI committee.

While FDDI was evolving from a machine room interconnect into a general-purpose LAN, the requirements changed. For a machine room interconnect, some management operation to install and initialize the network might reasonably be allowed. For example, the manager might set the values of various parameters to control the operation and performance of the interconnect network. However, in a general-purpose LAN, manager involvement is unacceptable. For simplicity, robustness, and ease of management, the industry widely

definition of the FDDI technology to include such products as concentrators and adapters.

accepts that LANs must autoconfigure, also called "plug-and-play." Inevitably, FDDI was required to exhibit the attributes of a true local area network.

Overview

Since the FDDI technology and standards were already in development when this evolution of requirements occurred, the ANSI committee made an attempt to accommodate the following two views of the network: first, the network should be completely configurable with almost every parameter and policy controlled by management; and second, the network should be a local area network with the corresponding attributes of simplicity and autoconfiguration. Incorporating both models into the ANSI FDDI standards made the standards complex and was one factor contributing to the eight-year-long time period to completion.

Theory of Operation

FDDI stations are composed of the basic elements defined by the FDDI standards. The physical medium dependent (PMD) layer specifies the fiber-optic interface and data driver and receiver operation for FDDI stations.[3] The physical layer protocol (PHY) specifies the encoding and framing of data and control information exchanged

between stations.[4] The control information exchanged varies with the physical layer protocol type, which is either PHY-A, PHY-B, PHY-M, or PHY-S. The MAC sublayer specifies the protocols for logical ring formation and control, for the transmission and

reception of packets at a station, and for the repetition and stripping of packets on the ring.[5,6] Station management (SMT) provides n-layer management and a local management interface to the PMD, PHY, and MAC layers.[7] Together these components support an IEEE 802.2-compatible logical link control capable of supporting client protocols such as the Digital networking (DECnet) protocol, open systems interconnection (OSI), local area transport (LAT), and the transmission control protocol/internet protocol (TCP/IP). Stations utilizing the FDDI components can take several forms such as single attachment stations (SASs), dual attachment stations (DASs), and dual attachment concentrators (DACs). An architectural model is shown in Figure 1.

Fiber Distributed Data Interface Overview

Configurations of FDDI Components

A single attachment station is the simplest configuration and consists of the fundamental FDDI components arranged as shown in Figure 2. There is a single incoming data path and a single outgoing data path with a MAC in between.

Dual attachment stations, as illustrated in Figure 3, include a second physical layer and provide connections to a secondary ring for use in the event of breakage on the primary ring. Under fault-free operating conditions represented by the THRU STATE area of Figure 3, the primary data path enters through PHY-A, travels through the MAC, and exits through PHY-B. The secondary data path enters through PHY-B and exits directly through PHY-A. If a discontinuity is detected in the primary data path, either within the station or on one of the PHYs, the station wraps the two data paths, thereby providing an alternate route through the secondary data path. This situation is shown in the WRAP A STATE area of Figure 3.

A dual attachment concentrator builds on the dual attachment station

the additional PHY-Ms in the concentrator.

An FDDI LAN is formed by joining multiple stations to form a logical ring topology. The logical ring can take two physical forms, a dual trunk ring and a tree ring.

As depicted in the upper portion of Figure 5,

the dual trunk ring is formed by connecting dual attachment stations and concentrators to form a LAN. This portion of the LAN consists of two data paths in opposite directions, called the primary and secondary rings. Under normal operation, data flows on the primary ring from station to station. In the event of a cable or a station breakage, the stations adjacent to the fault join the primary and secondary rings and then use the secondary path to reestablish a logical ring.

A tree ring can be formed by connecting stations or concentrators to the PHY-Ms of a concentrator as shown in Figure 5. In this formation, the primary data path descends down each branch of the tree passing through each station in

the tree, until it finally reemerges into the dual trunk ring.

by adding additional
master PHYs (PHY-M) in
the primary data path as
shown in Figure 4. Single
attachment stations can
then be included in the
ring by connecting them to

Overview

Media Access Control
Sublayer Operation

As mentioned previously, the MAC sublayer provides the protocols for logical ring formation and data packet operations. To initialize the ring, all MACs first enter the claim process to determine which MAC will generate the token and to establish the token rotation time for the ring. Each station continuously transmits claim frames that contain the station's requested token rotation time. When a MAC receives claim frames with times shorter than its own, or equal to its own but from a station with a numerically larger address, it yields, stops sending claims, and repeats the claims received from its neighboring station. Eventually the station with the shortest or "winning" time will receive its own claim. This station then generates the token and the ring enters the operational state.

If the claim process does not complete within approximately 100 milliseconds, the MACs in the ring perform a beacon process to confirm continuity of the ring. Special beacon frames are transmitted continuously

To transmit data packets, the MAC first waits for the token to arrive, holds it, and, then, transmits the packets, reissuing a new token at the end of the transmitted packet stream. The time allowance a station has to transmit packets after receiving a token is equal to the token rotation time established by the claim process.

Packets received by a MAC are either repeated for reception by the next station or stripped from the ring. In addition, a MAC may store a copy of a packet for use by the station. After transmitting a frame, a MAC is responsible for stripping that frame from the ring after the frame makes exactly one traversal of the ring. Frames left unstripped are considered no-owner frames and can circulate the ring forever. This condition floods the station to which the frame is addressed and is thus detrimental to ring performance. A MAC

typically strips frames by comparing the source address in the frame with the MAC's own address. The MAC strips any frame it has previously sent but repeats the frame, otherwise.

FDDI Ring Formation

by all MACs until a beacon is received, at which point it stops transmitting. This process continues until one MAC transmits and receives its own beacon, indicating ring continuity.

An FDDI ring is formed in several stages, beginning with the successful establishment of point-to-point links between all adjacent PHYs. These link connections are made by the connection management

Fiber Distributed Data Interface Overview

protocol (CMT).[7] This protocol defines control signal exchanges to synchronize the two ends of the link, to exchange information about the PHY type (i.e., A, B, M, S) and link testing requirements of each end to perform link quality testing, and to finalize the connection for normal operation. Before the link establishment enters its final phase, the PHY types of the two ends of the connection are compared, and the connection is allowed if the end types conform to specific connection rules. These rules are carefully established to ensure that rings are configured correctly and to prevent miscabling, which could cause partitioned and unnecessarily wrapped rings. An established link is continually monitored for errors indicated by the reception of improperly encoded data and is shut down if those errors exceed a predetermined threshold.[8]

After the PHY connection is fully established, the station's MAC is inserted into the ring, and the claim process begins. Digital's stations all use a default requested token rotation time of eight milliseconds in their claim process to ensure a token

essential to many LAN applications such as disk I/O, interactive graphics, and remote procedure calls. Higher settings of the requested token rotation time result in very large delays on the ring while improving the efficiency only slightly.[9]

Therefore, after extensive performance modeling, Digital decided to use this default value for the requested token rotation time. The resulting network operates with low delay and high bandwidth as the default and does not need complicated network management procedures to achieve this level of operation.

Several unique but harmful conditions in FDDI stations must be addressed. Among them are the prevention of and protection against duplicate addresses, no-owner frames, and the stripping of frames sent by bridges or end stations that have multiple addresses.[10] Because several of the

algorithms fundamental to the operation of the FDDI technology use the stations' addresses, the presence of two stations with the same address causes numerous malfunctions, ranging from beacon-claim oscillations to blocked communications

rotation time, and hence traffic latency, similar to that experienced in other LANs like Ethernet. FDDI provides high aggregate bandwidth and, thus, delivers the low delay

between stations. Frames not properly stripped from the ring can circulate forever, flooding the stations that copy these frames. To protect against strip errors, Digital's

Overview

chip set has several built-in mechanisms. Digital also greatly improved the data integrity of the ring. This improvement is particularly important to token ring architecture where messages traverse virtually all of the links in the network before arriving at their destinations. Robustness in the face of link bit errors becomes extremely important. Digital designed several improvements to the basic FDDI algorithms, and the ANSI committee adopted them to improve the undetected error rates on a network.[11]

Role of FDDI in Digital's Networks

The FDDI technology is more than just another data link interface that allows the use of transmit and receive fibers between devices. Digital's decision to embark on the FDDI development effort was a major program undertaking involving the development of VLSI chips and, subsequently, FDDI software and hardware products. Although the development of chips may seem to be at the heart of the FDDI program, chips are certainly not the products that help customers solve problems. Chips comprise only a small portion of

Before focusing on chip development, Digital carried out a large simulation effort to ensure the ANSI standards were correct and complete. Once the standards were verified, modeling was performed to produce chips that met these standards. Allowances and trade-offs were made for unfinished sections and future standards migration. Real products were then planned around the chips as the use of FDDI was threaded into Digital's network architectures and existing products. An implementation strategy for each product was then formulated to balance the risks, the resources, and the timeliness of customer needs.

It was important for Digital to understand migration of current products and to think forward to the needs beyond the initial program goals. LAN management and host connections take the FDDI technology beyond a simple high-speed backbone for the extension of bandwidth-limited existing LANs. Because Digital considered the FDDI technology beyond its use as just another new data link, this technology is the natural choice as the next step in network evolution.

each large printed circuit board, but this portion is an important one.

At the onset of Digital's FDDI program, and at several points during its development, a number of key program-wide decisions and policies were adopted. The commitment to

Digital Technical Journal Vol. 3 No. 2 Spring 1991

Fiber Distributed Data Interface Overview

simulation, both in breadth (e.g., chip design, board design, and software) and extent (e.g., behavioral models, gate-level models, integration of operational software with simulation models) was essential to the success of the FDDI technology development. All the chips developed were fully functional in their first pass in silicon, and the integration of the controlling and test software with the chips was fast and smooth.

Given the extensive investment in ensuring correctness of the FDDI technology, all FDDI products were to use the same FDDI building blocks, including the chips and controlling software.[12] Such sharing and reuse of investment eliminated the duplication of effort, guaranteed the consistent operation of all products, and shortened the time to market. To ensure an even broader use of this technology, major portions were made available on the open market.

Another key program decision was the adoption of electrically alterable read-only memory (EAROM) in place of traditional read-only memory (ROM) to store the firmware in each product. SMT, which is implemented almost entirely

shipment, EAROM was included in all products to allow firmware to be updated remotely over the network. As a result, EAROM reduced the cost of product enhancement by eliminating the need to change ROM in the field or swap out boards and rework them at the factory.

Deployment of FDDI

As a baseline effort in deploying FDDI in products, Digital developed the FDDI design corner shown in Figure 6, consisting of chips and FDDI control firmware for use in all applications of FDDI. In addition to using the design corner in all of Digital's products, two of the chips critical to FDDI interoperability, the MAC and the ELM, together with the SMT firmware, which controls the chips and the behavior of the station, were licensed to Advanced Micro Devices and to Motorola, Inc. for manufacture and sale on the open market. Widespread availability of this technology will foster competition to drive down cost and increase levels of interoperability and consistency among FDDI implementations.

in firmware, was a rapidly
evolving specification
while the products were
under development. To
accommodate these ongoing
changes and the threat
of change after product

Fiber Distributed Data Interface

Overview

FDDI Concentrator

FDDI-To-Ethernet Bridge

The Ethernet/802.3 bridge to the FDDI network is an important device in the first product set offered for FDDI connectivity. The DECbridge 500 product provides the smooth migration of all Ethernet and IEEE 802.3 devices currently in the marketplace to the increased backbone bandwidth of the FDDI technology. Most local networks are LAN-based with the LANs being interconnected to form extended LANs by the use of bridges. The bridges are transparent to all users' protocols on the LANs. This capability for transparent interconnection is one of the keys to the instant use of and easy migration to the added bandwidth and other benefits of FDDI. However, providing the transparent interconnect is difficult. Frame formats, bit ordering, padding fields, and even the length of the packets differ between the Ethernet and IEEE 802.3 LANs. These differences force the frames to be translated

as they pass through the DECbridge 500 device. This translation, coupled with a stiff requirement for

To make the FDDI technology widely acceptable as the next generation LAN, several of the logistical obstacles inherent in a ring topology had to be overcome. Connecting new stations to a LAN can be frequent events. These events must not be restricted or cause a disruption if the LAN is to be truly usable. The nature of a ring topology causes the network to break whenever a station is removed or while a station is being inserted. The dual ring accommodates one such event by wrapping at the adjacent stations, but two or more such events are disastrous because they partition the network into isolated parts. A concentrator can solve this problem by connecting to a station, testing the link between the station and the concentrator, and then splicing the new station into the ring with minimal disruption. At the same time, a concentrator can automatically drop a station out of the ring should the station malfunction or lose power.

In addition to the obstacle of connecting new stations, it is inconvenient to map a ring topology into

data integrity for the
MAC bridging standard
IEEE 802.1d and the
performance needs of
dealing efficiently with
packet rates approaching
500,000 per second,
created a formidable design
task.[13]

star wiring, which is
well known and widely
used for its convenience
and manageability.
Concentrators can turn
a logical ring into a
hierarchical star-wired
network, thus solving this
mapping problem.

Digital Technical Journal Vol. 3 No. 2 Spring 1991

Fiber Distributed Data Interface Overview

Digital's DECconcentrator 500 product was built to support simple installation and network configuration, existing building cabling plants, and the cost savings and simplicity of single attachment stations. This FDDI concentrator is low cost and flexible, with a dual attachment and up to eight master ports for connecting single attachment stations or additional concentrators into the ring. With the concentrator, any of the described FDDI topologies can be created to best suit a customer's cabling, network management, and availability needs.[14]

FDDI for Workstations

Digital designed a simple high-performance FDDI adapter for TURBOchannel workstations and servers in conjunction with the design of the DECstation 5000 reduced instruction set computer (RISC)-based system.[15] The aim of this product was to produce an FDDI interface that was both inexpensive and high performance. With this adapter and ULTRIX operating system support, Digital attempted to address many of the shortcomings that have plagued most adapters and operating systems attempting to achieve high

applications and on the network configuration. One application is the use of the FDDI technology for a work group of high-performance workstations and servers. Some high-performance I/O-intensive applications need more bandwidth than a large, shared Ethernet can handle. As a result, a work group may form around an FDDI network. One application example is called visualization, where a large volume of real-time graphics or imaging information is transmitted over the network from a compute server. Various applications of digital image transmission have few real-time latency requirements, but still require a large bandwidth.[17]

Another application of the FDDI technology in workstation environments is to achieve higher bandwidth for server/server and backbone interconnection. A shared Ethernet for client /server as well as server /server communications can be overloaded, particularly if the population of servers or workstations is large. FDDI can be used as a backbone to provide higher bandwidth between the servers. Alternatively, the network may be a hybrid, where the FDDI network is used for high

performance.[16]

The applications of the FDDI technology for workstations and servers depend on the particular choice of user

client/server bandwidth by

some systems and for high server/server bandwidth by others. Hybrid networks often have multiple FDDI networks and multiple

Overview

Ethernet networks. This configuration is typical of evolving extended LANs, as new technology such as FDDI is added into the existing network infrastructure.

kilometers such as a campus backbone. Other medium types are currently under investigation to provide optimization in areas such as cost and ease of installation.

FDDI in the Future

As a new technology,

FDDI has a long future with opportunities for greater levels of circuit integration, lower cost designs, and alternative media types to match the wide range of possible applications. Significant progress has already been made in this direction with the introduction of the DECcontroller 500 adapter, the most compact, least expensive, and one of the highest performing workstation adapters in the industry today. In addition, the ground-breaking technical work and standards committee activity regarding various types of copper media offers tremendous opportunities for cost reduction which will enable broader utilization of the performance offered by the FDDI technology.[18]

In addition to the original physical medium dependent standard already specified by ANSI for FDDI, an additional standard

Acknowledgments

We would like to thank the efforts of the Telecommunications and Networking LAN Architecture and Development teams. Without their hard work none of the many technical advances made in bringing the FDDI technology to market would have been possible. We would also like to thank Bob Krueger for his role in successfully leading the entire program. Finally, we would like to thank Mark Kempf, Paul Koning, and Tony Lauck for their timely review of this paper.

Reference

1. N. Kronenberg, H. Levy, W. Strecker, and R. Merewood, "The VAXcluster Concept: An Overview of a Distributed System," Digital Technical Journal vol. 1, no. 5 (September 1987): 7-21.
2. R. Grow, "A Timed Token Protocol for Local Area

has been developed for operation on single-mode fiber over distances up to 30 kilometers. This standard will be particularly useful in deploying an FDDI network that must span one or more distances greater than 2

Networks," Proceedings of IEEE Electro/82 Conference, Boston, MA (May 25-27, 1982).

Fiber Distributed Data Interface Overview

3. Token Ring Physical Layer, Medium Dependent (PMD), (International Standards Organization, reference no. ISO 9314-3, 1990).
4. Token Ring Physical Layer Protocol,

(International Standards Organization, reference no. ISO 9314-1, 1989).
5. Token Ring Media Access Control (MAC), (International Standards Organization, reference no. ISO 9314-2, 1989).
6. H. Yang, B. Spinney, and S. Towing, "FDDI Data Link Development," Digital Technical Journal, vol. 3, no. 2 (Spring 1991, this issue): 31-41.
7. FDDI Station Management (SMT) Preliminary Draft Proposed American National Standard, ANSI X3T9/90-X3T9.5/84-49, REV 6.2 (May 1990).
8. J. Hutchison, C. Baldwin, and B. Thompson, "Development of the FDDI Physical Layer," Digital Technical Journal, vol. 3, no. 2 (Spring 1991, this issue): 19-30.
9. R. Jain, "Performance Analysis of FDDI Token Ring Networks: Effect of Parameters and Guidelines for

10. H. Yang and K. Ramakrishnan, "Frame Content Independent Stripping for Token Rings," Proceedings of the ACM SIGCOM '90 Symposium (1990): 276-286.
11. R. Jain, "Error Characteristics of Fiber Distributed Data Interface (FDDI)," IEEE Transactions on Communications, vol. 38, no. 8 (August 1990).
12. P. Ciarfella, D. Benson, and D. Sawyer, "An Overview of the Common Node Software," Digital Technical Journal, vol. 3, no. 2 (Spring 1991, this issue): 42-52.
13. R. Kochem, J. Hiscock, and B. Mayo, "Development of the DECbridge 500 Product," Digital Technical Journal, vol. 3, no. 2 (Spring 1991, this issue): 53-63.
14. W. Tiffany, P. Koning, and J. Kuenzel, "The DECconcentrator 500 Product," Digital Technical Journal, vol. 3, no. 2 (Spring 1991, this issue): 64-75.
15. U. Sinkewicz, C. Chang, L. Palmer, C. Smelser, and F. Templin, "ULTRIX Fiber Distributed Data Interface

Setting TTRT," DEC-TR
655 (Maynard: Digital
Equipment Corporation,
September 1989).
(internal)

Networking Subsystem
Implementation," Digital
Technical Journal, vol.
3, no. 2 (Spring 1991,
this issue): 85-93.

14 Digital Technical Journal Vol. 3 No. 2 Spring 1991

Fiber Distributed Data Interface

Overview

16. W. Hawe and K. Ramakrishnan, "The Workstation on the Network: Performance Considerations for the Communications Interface," Proceedings of the Second Workshop on Workstation Operating Systems, Asilomar Conference Center, CA (September 1989).
17. K. Ramakrishnan and W. Hawe, "Performance of an Extended Local Area Network for Image Applications," Proceedings of the Fifth Annual Phoenix International Conference on Computers and Communications, Scottsdale, AZ (March 26-28, 1986).
18. S. Ginzburg, W. Mallard, and D. Newman, "FDDI over Unshielded Twisted Pairs," IEEE Proceedings, Eighteenth Conference on Local Computer Networks (October 1990).

=====
Copyright 1991 Digital Equipment Corporation. Forwarding and copying of this
article is permitted for personal and educational purposes without fee
provided that Digital Equipment Corporation's copyright is retained with the
article and that the content is not modified. This article is not to be
distributed for commercial advantage. Abstracting with credit of Digital
Equipment Corporation's authorship is permitted. All rights reserved.
=====