

N 9 1 - 2 4 0 7 2 !

Fiber Optic Tactical Local Network (FOTLAN)

L. A. Bergman, R. Hartmayer, W. H. Wu, P. Cassell, G. Edgar,
J. Lambert, R. Mancini, J. Jeng, and C. Pardo

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109

ABSTRACT

A 100 Mbit/s FDDI network interface unit (NIU) is described that supports real-time data, voice and video. Its high-speed interrupt-driven hardware architecture efficiently manages stream and packet data transfers to the FDDI network. Other enhancements include modular single-mode laser-diode fiber optic links to maximize node spacing, optic bypass switches for increased fault tolerance, and a hardware performance monitor to gather real-time network diagnostics.

INTRODUCTION

The FDDI token ring standard provides for a high bandwidth (100 Mbit/s) general purpose interconnection among computers and peripheral equipment using 1.3 μm optical transceivers and 62.5 μm optical fibers as the transmission medium. The network stations are connected in a dual counter-rotating ring configuration with a maximum node-to-node spacing of 2-km [1]. While intended mainly for data traffic, other forms of real-time traffic, such as voice and video, may also be transported over the asynchronous FDDI network using a novel double elastic buffer interface reported earlier in [2] for a lower-speed 80-Mbit/s network. Ordinary data packet traffic was simultaneously transmitted with synchronous T1 voice traffic using a novel voice interface that required only a deterministic network protocol.

Our present effort expands on the previous achievements by migrating the data/voice interface used previously to a FDDI network interface unit (NIU) and additionally providing real-time video capability using a technique similar to that used for voice. Furthermore, single mode fiber links, laser diode transmitters and optical bypass switches have been added to increase the distance between adjacent nodes and to introduce network fault tolerance and redundancy. Finally, a network latency performance monitor has been implemented using custom logic.

The NIU system architecture will be described in further detail in Section 2, including the design of the FDDI protocol logic, node processor, voice/video interface, and fiber optic links. The theoretical and experimental performance measurements will be discussed in Section 3, followed by an overview of potential applications in Section 4.

NIU SYSTEM ARCHITECTURE

Concept

The general design provides for computer asynchronous data packets, as well as real-time data packets to be transported by each FDDI node. The major blocks of a typical FDDI node are shown in Fig. 1. The real-time data stream may be accepted from the host computer's bus or more typically from an internal peripheral bus, allowing the operating system not to be degraded by steady real-time traffic loads. FDDI, being an asynchronous network, cannot directly interface to real-time traffic streams without appropriate buffering and resynchronizing logic. In order to interface the real-time voice and video traffic to the LAN elastic buffers, described elsewhere [3], double elastic buffers are added to the input/output sides of each FDDI node with real-time services.

FDDI Chip Set

The NIU design is based on the AMD SUPERNET chip set [4], shown in Fig. 2. The five-chip SUPERNET family meets the ANSI X3T9.5 FDDI standard and acts as an interface between a host computer and the network medium, transferring data and converting it between parallel form at the host and serial form at the media. The chip set consists of the RAM buffer controller (RBC), the data path controller (DPC), the fiber optic ring media access controller (FORMAC), the encoder/decoder (ENDEC), and the ENDEC data separator (EDS). The RBC generates addresses to buffer memory for received and transmitted frames, handles buffer management, and provides interrupts to the Node Processor (NP). The DPC converts data in received frames from byte-wide to 32-bit word formats and vice versa in transmitted frames, performs parity checks, generates frame and node status, and provides interrupts to the NP. The FORMAC performs the media access control layer protocol for the FDDI networking scheme. It determines when a node can get access to the network, implements the logic required for token handling and address recognition, and generates status bits which identify node conditions and frame status. The ENDEC performs the 4B/5B encoding, converts data from parallel to serial format and sends it to the fiber optic transmitter. The EDS extracts the clock from the received bit stream, converts the serial data into a five-bit parallel data format, and performs the 4B/5B decoding.

Node Processor

The node processor (NP) is implemented using a microprocessor based system. Its function is to initialize and oversee the operation of the FDDI chip set, manage the internal NIU bus, perform data conversion/management functions, act as a fault/diagnostics watchdog, and setup external interfaces. Since it functions primarily as a hardware controller, the processor has been designed to respond quickly to packet level interrupt signals. A goal is that NP respond to and process interrupts within a 4 μ s window. This will allow most real-time interface buffers to be monitored with sufficient accuracy. The 32-bit AMD 29000 RISC CPU [5] was selected from among the several available candidates. It has a high transfer rate, operates at speeds up to 33 MHz, handles DMAC functions, and has strong software and hardware support. The NP operates with a 25-MHz clock input and incorporates four state machines in the design. The first one generates the CPU wait states, the second and third ones provide odd/even word generation and the fourth one provides decoding for the microprocessor's I/O and addressing functions. The NP communicates with the FDDI chip set

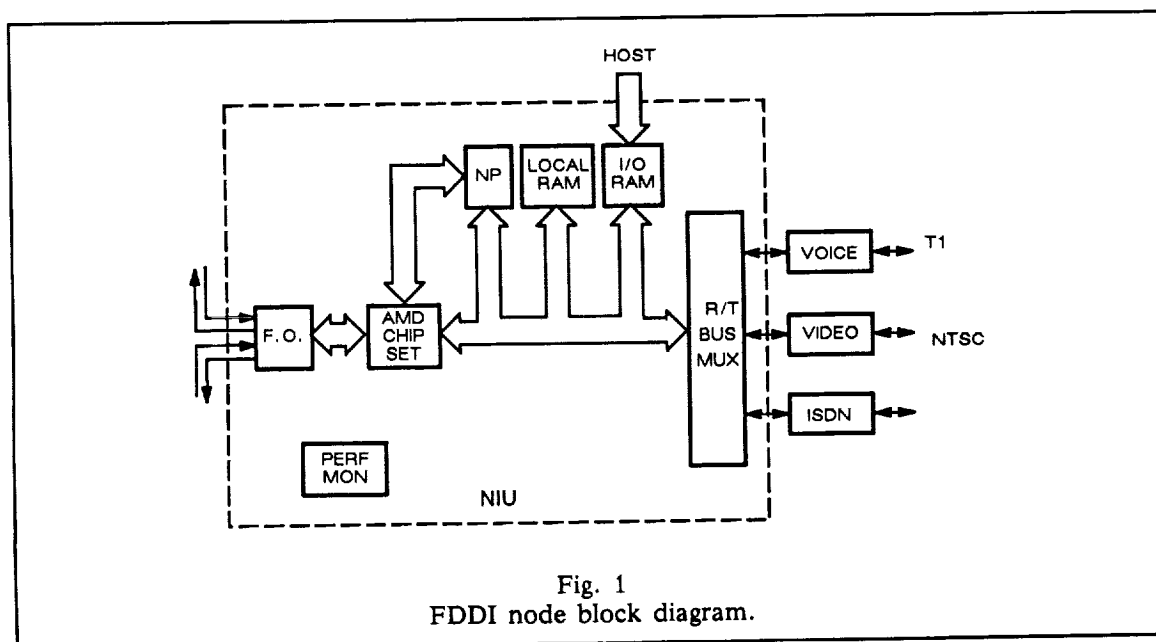


Fig. 1
FDDI node block diagram.

using the 16-bit Node Data and Node Address busses. The SUPERNET chip-set and the buffer memory make their status available to the NP, which has complete control over them.

Voice/Video Interface

The voice and video traffic are both transferred throughout the system in real-time. The voice call setup and switching functions are performed external to the network with PABX equipment [6]. The T1 voice traffic transmitted at 1.544 Mbit/s is stored in the first-in-first-out (FIFO) memory. The 8-bit FIFO busses are then multiplexed and transferred, at 12.5 MB/s, to the NIU data bus via the node path interface circuit. Multiple FIFO channels may be interleaved into one LAN packet. The video signal, fully compatible with the NTSC 525-line, 30-Hz input/output standard, requires a 88 Mbit/s capacity to be transmitted in real-time. Once the pictures are taken with a TV camera, they are sampled at 10 MHz and stored in a 525x525x8 (275,625 bytes/frame) frame store. The frame field position is then added via a header, multiplexed and transferred, at 12.5 MB/s, to the NIU data bus via the node path interface circuit. Since each FDDI packet can only hold approximately 4000 bytes, each video frame is partitioned into 75 packets of 3675 bytes. The frames will be sent via the reliable TCP/IP protocol or via UDP, and corrupted frames will be discarded. In the event that network traffic is high, pictures will be updated less frequently, resulting in "freeze frame" TV pictures. This feature allows the video to act as a free-space bandwidth hog on the FDDI network, giving it considerably more operational flexibility than its voice counterpart.

Fiber Optic Links

Each FDDI node will be configured in a dual-attach station mode, requiring two fiber optic transceiver pairs and four fiber optic links for a dual counter-rotating network topology. The ring configuration, regular versus loop-back, is controlled by the FDDI chip set. In addition, optical switches allow for the node to be entirely bypassed in the event of optical transceiver failure. Two nodes will utilize laser diode transmitters, extending the distance between them to at least 10 km using single-mode fiber optic cable.

Transceivers. The fiber optic transceivers for normal FDDI operation, are hybrid modules [7] operating at 125-Mbit/s NRZ. The transmitter utilizes a high-speed LED having an average rise/fall time of 2 ns and output power of -15 to -20 dBm at $\lambda=1305-1380$ nm. The receiver uses a PIN




Fig. 2
Detailed FDDI node/system block diagram.




Fig. 3
Fiber optic cable assembly.

photodiode and has a sensitivity of -34 dBm [8]. Both the transmitter and the receiver are terminated with an ST type connector. For extended distances greater than 3 Km, a single-mode laser diode transmitter, having an output power of 0 to -3 dBm will be used. The receivers, having sufficient sensitivity, remain unchanged.

Dual Fiber Cable. The ruggedized two-fiber tactical fiber optic cable contains two tightly buffered 50/125- μm multi-mode optical fibers [9] and are terminated with hermaphroditic biconic connectors [10], which facilitate field deployment and retrieval without physical or optical degradation. The fibers are radiation hard and fully militarized to withstand the tactical field environment. The fibers are cabled in a ruggedized all-dielectric structure with kevlar yarns and surrounded by reinforcing elements and an outer jacket. Their attenuation is 1.0 dB/km at 1300 nm. The outside diameter is 6 mm and its weight is 30 kg/km. The link optical power budget is 10 dB, yielding a maximum transmission distance of 3 km. A typical fiber optic cable assembly is shown in Fig. 3.

Single Mode Fiber Cable. The long distance transmission links will utilize ruggedized single mode tactical fiber cables [11]. These cables, similar in packaging to the multi-mode version, contain a radiation hardened, single mode optical fiber. The fiber has a core diameter of 8.3 μm and has an attenuation of 0.35 - 0.50 dB/km at $\lambda=1310$ nm. At this wavelength, which is also the zero dispersion wavelength, the maximum fiber dispersion is 3.2 ps/nm-km.

Optical Bypass Switches. The optical bypass switches enable the network to be reconfigured in two different ways. A single-node will be bypassed in the event of an optical transceiver failure. The data will enter the node and exit immediately, without reaching the FDDI card. In the second bypass mode, an entire portion of the ring can be disconnected, thereby bypassing several nodes. This scheme would be useful if a larger portion of the ring is not functional or is even destroyed. Ideally, a low attenuation 3x3 optical crossbar switch would be used in order to implement this scheme. Since no commercial product is available at this time, two 2x2 switches connected in series are being used. Since the increased loss exceeds the available power budget, the optical bypass will only be used with the laser diode transmitter links.

Packaging

The FDDI boards will be implemented on a PCB, which will conform to the VME standard. For the prototype version, the design will be broken down into three 6U high boards operating in a standard VIM/VME cardcage. The final version will include all the different FDDI functions on one 9U high card, which can be used in a SUN workstation setup [12]. Test points for signal monitoring will be included, as well as jumpers isolating the separate board functions, if needed. LED indicators will also be included to monitor the link operation.

PERFORMANCE

Theory

The predicted network performance is illustrated in Fig. 4 [13]. Time delay versus throughput for several token protocols are shown. Observe that FDDI offers a capacity of 768 voice channels at time delays of ~ 1 ms or more. In practice, fewer voice channels would be used to give reasonable access delays for data-only packets. For example, in a network file system (NFS) service, an access delay of $<250 \mu\text{s}$ would be desired thereby restricting the throughput to about 30 Mbit/s or 240 voice channels. The capacity of low-speed LANs is too low to support T1 traffic. However, single voice channels might be supported using the same technique with 802.4 or 802.5 LAN protocols.

Simulations

Real and Non-Real-Time Performance Estimation. A simulation of the data flow within the NIU can provide valuable feedback to designers on how to optimize the overall hardware performance. For example, to maximize end-to-end throughput one might intuitively allocate an equal bandwidth among each of the functional blocks (e.g., bus, memory). However, in a system where real-time and non-real-time traffic are freely mixed together, the message sizes, frequencies, and statistical distributions will vary dynamically, thereby making such an estimation considerably more difficult. In this instance, packets may be forced to be randomly stored (or buffered) at numerous locations creating bottleneck points. For example, a low bandwidth introduces high queue populations for even moderate network/host loads and consequently high latencies. Conversely, a very high-bandwidth maintains a low-Q population but at very low internal bus utilization, and at great cost. Thus, an optimum exists.

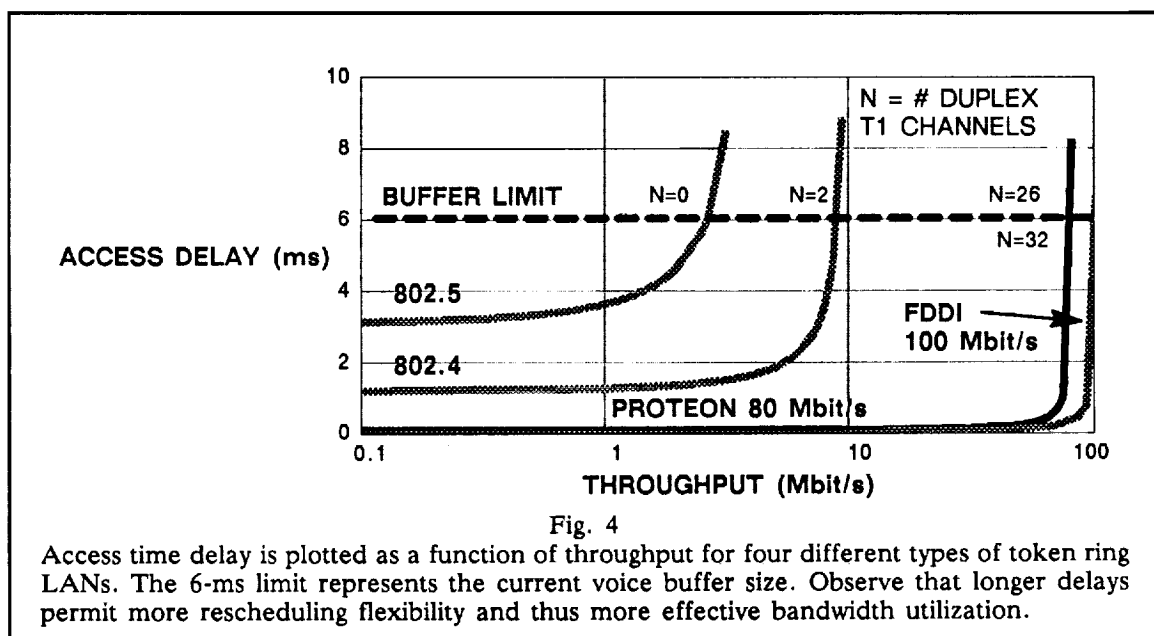



Fig. 4

Access time delay is plotted as a function of throughput for four different types of token ring LANs. The 6-ms limit represents the current voice buffer size. Observe that longer delays permit more rescheduling flexibility and thus more effective bandwidth utilization.



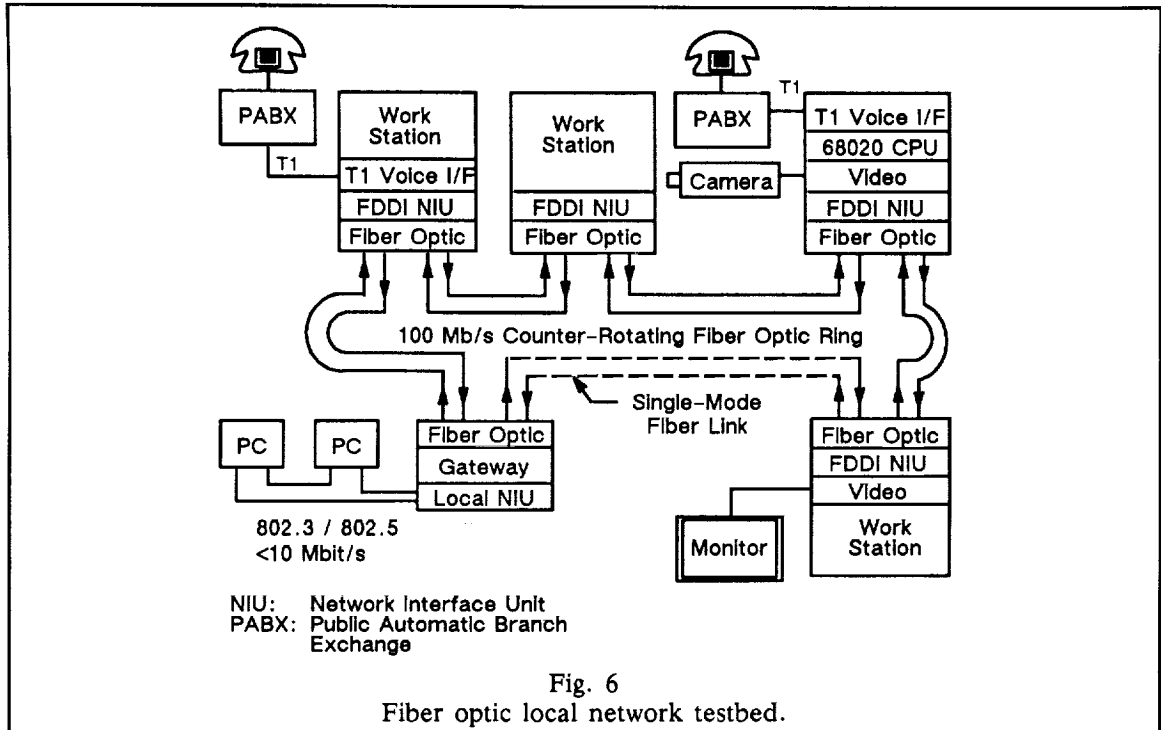
Discrete Event Simulator. An object-oriented graphical discrete-event modeling environment such as *SES/Workbench* can provide an effective means of estimating the minimum buffer depths and bus bandwidths to control queue buildup and excess latency. The objective was to study the effect of buffer memory bandwidth on overall performance throughput and latency, and suggest reasonable performance goals to the designers. The RAM buffer design suggested by AMD has a bandwidth of 200 Mbits/s. In the case of input distributions of 3T1, 2 ISDN, 2T2 and Poisson processes for computer data, the simulation results show that a buffer memory bandwidth of under 200 Mbits/s can cause a severe bottleneck (as expected), but that a buffer memory bandwidth of around 300 Mbits/s is optimal (and less obvious). If the portion of the network bandwidth allocated to video is increased from 12 Mbit/s to 40 Mbit/s, a very high Q population for communication specific bus will result even with a bandwidth of 350 Mbit/s for buffer memory. Thus, even larger bus/buffer memory bandwidths may be desirable in these cases.

Measurements

Performance Monitor. The purpose of the performance monitor is to provide a means to measure the end to end network latency and throughput at the TCP and IP layers, and the queue timing at the NP. In measuring these parameters, we need to minimize the influence of the protocol as well as the software processing. In order to synchronize timing measurements, a global reference time needs to be established throughout the LAN. Each NP module will include a global clock module (GCM), which will read a reference time with 100-ns resolution from an external common reference source. One master GCM will orchestrate all timing measurements throughout the network and synchronize the reference time. The real-time measurements will be stored in a FIFO, to be collected by the CPU at its convenience. A block diagram of the GCM distribution is shown in Fig. 5.

Workstation Throughput. The throughput is measured with the User Datagram Protocol (UDP) and Transmission Control Protocol/Internet Protocol (TCP/IP) provided the workstations [12]. UDP is considerably more efficient, but does not perform flow control, acknowledgments, and error checking. The test software generates long blocks of random data in RAM and directly sends this to the network, bypassing the disk subsystem. At the receiving workstation, the same data is discarded on a byte-by-byte basis.

For a TCP/IP (non-Van Jacobson) protocol and unloaded LAN, the throughput between two workstations was measured to be 2.2 Mbit/s for a 80 Mbit/s LAN and 1.8 Mbit/s for an Ethernet network (Sun 3/160 workstations). The 20% improvement is limited by the TCP/IP software overhead. For the simpler UDP protocol, the token ring provides a two to three-fold improvement over Ethernet. Still, the throughput is much less than expected from theory—a limitation imposed by the operating system overhead and hardware platform. Throughputs exceeding 30 Mbit/s is expected between current RISC-based workstations.



Voice. Bit error rate (BER) tests are performed at the T1 interface of two units attached to the ring. Error injection tests indicate that BERs on the order of 10^{-4} were not discernible.

APPLICATIONS

Army. The testbed, shown in Fig. 6, will be used to demonstrate the feasibility of serving a variety of Command, Control, and Communication (C³) needs within the tactical Army using a FDDI based LAN for high-speed integrated voice/video/data communications. Modular network concepts, rapid deployment capability with fiber optics, and compatibility among several types of Army stations and LAN programs have already been investigated [14].

NASA: The techniques developed on FOTLAN for streaming voice/video traffic over a packet network may also be used many future NASA applications in both the spaceborne and terrestrial arenas. Spaceborne examples include spacecraft data buses for combining packet data and high-rate instruments such as HIRIS and SAR onto one fiber optic media. Within a decade, other stream-oriented instruments may also be added to space vehicles such as optic storage systems, optical processors, and systolic array processors. One of the important advantages of such a network is that it would be rapidly reconfigurable, making it relatively easy to relocate and replace subsystem modules in spacecraft assembly or for repair/upgrading by the shuttle. For ground based applications, the rapid deployment capability of FOTLAN may be important in simplifying cable layout for data/voice/video in some mobile tracking stations.

CONCLUSIONS

With the emergence of FDDI compatible components, it is only a matter of time before high-speed networks will be in common use. While the conventional FDDI standard only requires multi-mode data transmission over 2 km spans, the single-mode fiber and real-time interface augmentations offered here bring the standard to a more advanced plateau with added functionality, more diverse applications, and broader geographical coverage. Finally, the multi-level communication fabric proposed here provides, for the first time, a basis for a single integrated service workstation with seamlessly merged voice, video, and data graphics capability. Such end applications as voice/video E-mail, multi-station video conferencing, and distributed tactical command posts are but a few of

the possibilities. We expect stream-oriented ISDN to eventually provide similar capability, but its hierarchical fixed address topology may limit its usefulness in environments where rapid and frequent network reconfiguration is a requisite.

ACKNOWLEDGMENTS

The authors would like to acknowledge the contributions of team members G. Bolotin and L. J. Paul.

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the U. S. Army CECOM through an agreement with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government, the U. S. Army CECOM, or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

- [1] FDDI 2 (X3T9.5) Working Paper, Feb.1, 1987.
- [2] Bergman, L. A.; Hartmayer, R.; Marelid, S.; Wu, W. H.; Edgar, G.; Cassell, P.; Mancini, R.; Kiernicki, J.; Paul, L. J.; Jeng, J.; Pardo, C.; Halloran, F.; and Martinez, J.: A Fiber Optic Tactical Voice/Data Network Based on FDDI. Armed Forces Communications and Electronics Association (AFCEA) DoD Fiber Optic Conference '88, McLean, VA, pp. 123-125, March 22-25, 1988.
- [3] Bergman, L. A.; Hartmayer, R.; Marelid, S.; Wu, W. H.; Edgar, G.; Cassell, P.; Mancini, R.; Kiernicki, J.; Paul, L. J.; Jeng, J.; Pardo, C.; Halloran, F.; and Martinez, J.: A Fiber Optic Tactical Voice/Data Network Based on FDDI. INFOCOM'88, New Orleans, LA, pp. 45-54, March 29-31, 1988.
- [4] Advanced Micro Devices - The SUPERNET Family for FDDI. [EDS: Am7984A, ENDEC: Am7985A, FORMAC: Am79C83, DPC: Am79C82, RBC: Am79C81]
- [5] Advanced Micro Devices - Am29000 32-bit streamlined instruction processor.
- [6] Wang WBX public automatic branch exchange.
- [7] AT&T ODL-125TR (105711923) Light Wave Data Links.
- [8] Specified at 125 Mb/s data rate, using 62.5 μ m fibers.
- [9] AT&T 10524331 tactical dual-fiber cable.
- [10] AT&T 105243596 dual-fiber bulkhead connectors.
- [11] AT&T radiation hardened fiber.
- [12] Sun Microsystems, Inc., 3/160M.
- [13] Bergman, L. A.; Halloran, F.; and Martinez, J.; Implementation of a Tactical Voice/Data Network Over FDDI. MILCOM'88, San Diego, CA, October 23-26, 1988
- [14] Frankel, M. S.; et al: An Overview of the Army/ DARPA Distributed Communications and Processing Experiment. IEEE J. Select. Areas Communications, Vol. SAC-4, No. 2, pp. 207-215, Mar 1986.