Development of the DECbridge 500 Product

By Robert C. Kochem, James S. Hiscock, and Brian T. Mayo

Abstract

The DECbridge 500 product connects Ethernet/802.3 local area networks (LANs) to fiber distributed data interface (FDDI) LANs and is, therefore, a fundamental element of an extended LAN. Developers of this product encountered many technical hurdles. The higher data rate and token ring topology inherent in the FDDI technology impose several demands on any bridging product connected to an FDDI LAN. The differences in formats and size of frames on the two types of LANs introduce further requirements. The development team met these requirements and delivered a high-performance product that provides seamless integration of both LAN types.

Introduction

Bridges are essential to the creation of extended local area networks (LANs) because they provide transparent forwarding of traffic between adjacent LANs; local traffic is confined to its home LAN.

One important function of bridges is the ability, under network management control, to block traffic of selected protocol types or traffic from specific sources. Restricting unnecessary traffic, especially multicast or broadcast, significantly improves the utilization of LAN bandwidth.

In this paper we first discuss the role of the DECbridge 500 product in an FDDI and Ethernet /802.3 extended LAN and outline the design of the bridge. We then describe the operation of the bridge by tracing the flow of LAN traffic through it. This description gives insight into many of the complex tasks that a bridge must perform to connect two dissimilar LANs. Key points of the development methodology are also presented.

LANs.[1] Traffic may be forwarded to or from individual destinations, to groups of destinations (multicast), or to all destinations (broadcast).

Development of the DECbridge 500 Product

DECbridge 500 Design Considerations

The DECbridge 500 device serves as the point of connection between a new family of LAN products based on the fiber distributed data interface (FDDI) technology and a large installed base of Ethernet/802.3 LANs. The DECbridge 500 product must meet the requirements of both LANs to provide a smooth migration path for Digital's customers. Note that Ethernet and 802.3 have media access control (MAC) frame formats that may be used on the same 10-megabit (Mb)-per-second LAN. Throughout this paper, the expression Ethernet/802.3 is used to identify such LANs and to distinguish them from 100-Mb-per-second FDDI LANs. The terms Ethernet, 802.3, and FDDI are used when discussing the specific MAC frame formats.

System Description

Two typical extended LAN applications involving DECbridge 500 devices are shown in Figure 1. The backbone application employs an FDDI LAN to provide a high-bandwidth interconnect of multiple Ethernet/802.3 LANs. The DECbridge 500 device is the point of connection between requirements. File servers and other common resources may also be part of the local FDDI LAN. Here, the role of the DECbridge 500 product is to provide a path from the local work group to other parts of the extended LAN via Ethernet /802.3 LANs.

the Ethernet/802.3 LAN and the FDDI backbone LAN. In the work group application, FDDI LANs provide localized connectivity of users, such as DECstation 5000 workstations, that have high throughput

In either application, the bridge must perform the following functions:

- Forward traffic between nodes residing on two different LANs
- Prevent (i.e., filter or not forward) traffic between nodes on the same side of the bridge from getting to the LAN on the other side of the bridge
- Be responsive to host-based network management, provided by, for example, Digital's extended LAN management software (DECelms) and Digital's management control center (DECmcc) product
- Be a proper participant as an end station on both LANs to which it is connected
- Interact with other bridges in the topology of the extended LAN to prevent redundant paths or loops[2]

Hardware Description

Figure 2 shows a block diagram of the DECbridge 500 hardware design. The applications processor (AP), a subsystem based

on a 68020 microprocessor, performs initialization and maintenance of the bridge hardware as well as some steps involved in processing frames. The AP also acts as the management entity for the bridge.

The operating programs for the AP as well as two other processors, the queue manager and the translation processor, are stored in a nonvolatile electrically erasable programmable read-only memory (EEPROM). At initialization, the AP distributes the programs to random access memory (RAM) in the other two processors' subsystems. The AP executes much of its own program directly from the nonvolatile memory, although some high-performance operations are executed from static RAM.

The DECbridge 500 device may have an entire new operating program downloaded over the network and stored in the nonvolatile memory. This allows rapid updates of functionality without the need to perform a hardware upgrade on-site. Program updates are received via either of the attached LANs and stored in an area of RAM referred to as the "landing pad." The AP then transfers the new program into the nonvolatile memory and initiates a firmware reset.

The FDDI and Ethernet/802.3 chip sets and some analog interface circuitry provide connection to the two LANs. The bridge represents a single attachment station

Each chip set checks every incoming frame for integrity. Also, some rudimentary identity, or address, tests are applied. Frames that meet the integrity and identity requirements are then placed in a packet memory. The bridge maintains a table of learned MAC addresses. The table contains data for each address that is used to decide if a frame should be forwarded or filtered. The queue manager is a subsystem dedicated to checking each frame received in FDDI packet memory against the information contained in the learned address table. Based on this information, the queue manager decides whether to filter the frame, forward the frame to the Ethernet/802.3, or deliver the frame to the bridge entity for action.

The FDDI and Ethernet/802.3 LANs employ different data link protocols. The translation processor, a second 68020 subsystem, examines frames to be forwarded from one side of the bridge to the other. Each frame is reformatted

to the appropriate outbound protocol and moved from the incoming packet memory to the outbound packet memory. The two chip sets examine their respective packet

(SAS) on the FDDI ring. On the Ethernet/802.3 side of the bridge, switch-selectable ThinWire and attachment unit interface (AUI) connections are provided.

memories for outbound frames and transmit them onto their LANs.

Physical Description The DECbridge 500 product is shown in Figure 3. The hardware is approximately

Development of the DECbridge 500 Product

7 inches high by 17 inches wide by 14 inches deep

and may be rack-mounted or installed on a tabletop. It operates over the range of

100 to 240 voltage AC (VAC) at 50 or 60 hertz (Hz).

Figure 3, a photograph, (DECbridge 500 Product) is inserted here in the bound version only.

External signal connectors are located on the front edge of the two network interface cards, FI and NI. Each module has light-emitting diodes (LEDs) for various status functions and also diagnostics. In addition, the AP has a bank of switches for setting certain bridge operating functions.

The power and packaging were designed to simplify the swap out of field-replaceable units (FRUs). The four logic module FRUs can be replaced through the

front of the box, without opening it. By taking out only two screws, the outer

An exploded view of the bridge is shown in Figure 4. The electronics is implemented by the following four logic modules:

- AP, the applications processor
- QM, the queue manager subsystem including the learned address table
- FI, the FDDI chip set and the FDDI packet memory
- NI, the Ethernet /802.3 chip set and packet memory and the translation processor

frequency interference (EMI /RFI).

Operation

As mentioned previously, the DECbridge 500 device forwards traffic between two different LAN types. Consequently, the product development team faced several challenges beyond those encountered in previous bridges that connect similar Ethernet /802.3 LANs. The principal sources of these new challenges were:

- Higher data rates on the FDDI LAN. Ethernet/802.3 operates at 10Mb per

shell of the case can be removed. This gives access to the three other FRUs, namely, the power supply, the passive backplane, and a fan assembly. The five-sided design of the outer shell results in a product that is mechanically strong and provides shielding from electromagnetic and radio

second and has a minimum MAC frame size of 64 bytes. The maximum frame arrival rate is 14,880 frames per second (fps). FDDI operates at a rate of 100Mb per second and has a minimum MAC frame size of 17 bytes. The maximum frame arrival rate is 446,429 fps, a

rate 30 times greater than that of Ethernet /802.3.

- Different frame formats. Ethernet, 802.3, and FDDI have different MAC frame formats. Traffic entering an FDDI LAN from an Ethernet/802.3 LAN must be properly translated to an FDDI frame format. This translation must be performed in such a way that passage through a second bridge back to a different Ethernet /802.3 LAN results in a frame that recovers its original frame format.

- Different frame sizes. Ethernet and 802.3 differ from FDDI in both maximum and minimum frame sizes. FDDI frames shorter than

  the Ethernet and 802.3 minimum must be padded. FDDI frames longer than the Ethernet and 802.3 maximum cannot be forwarded, with the exception of special protocol types, which must be broken into multiple, smaller frames.

Objectives

The bridge can only forward frames from the FDDI LAN to itself, and frames to be discarded.

To comply with Digital's bridge architecture specification and the IEEE standard 802.1d for bridges, the bridge must examine all incoming frames.[2] It must identify, set aside, and process frames of each protocol type directed to itself, in the order received. To meet product performance requirements, the bridge must be able to forward frames at the full Ethernet/802.3 rate. A best effort must be made to buffer frames received

in bursts exceeding that rate. Frames should not be erroneously discarded. Results of this compliance visible to the network user are:

- Transparency. Nodes across the extended LAN operate as if they were connected to the same LAN.
- Stability. The paths in the LAN remain constant yet can reconfigure around equipment changes with a minimum loss of connectivity. Frames

  are not duplicated; nor

  are they received out of order.

the Ethernet/802.3 LAN at the maximum rate accepted by that LAN, i.e., 14,880 fps. Yet the arrival rate of frames from the FDDI LAN may be in excess of 440,000 fps. The incoming frames consist of an unknown mixture of frames that need to be forwarded, frames directed to the bridge

- Manageability. Network management can always observe and control the components of the extended LAN.

The operation of the DECbridge 500 device is best described by examining the progress made through the bridge by frames received from the FDDI LAN. Tracing this flow of traffic also gives insight into many of the challenges faced by the product's development team. The subsystems that process these frames and the flow of frames through logical queues in these subsystems are shown in Figure 5. The operation of the subsystems as the frames progress through them is described sequentially in the following sections.

Receiving FDDI Frames

The FDDI chip set in the bridge places all received frames in the FDDI packet memory on the receive queue. Frames in the FDDI packet memory can be accessed by subsystems in the bridge by using a virtual address method. A page table memory is used to assign a physical 512-byte buffer to each of 16K virtual buffers. Queues contain sequential sets of virtual buffers. Data frames are "moved" from one queue to another by moving the virtual address buffer pointers from one queue to another.

Frames received from the FDDI LAN may be as long as 4500 bytes. Frames longer than 512 bytes are chained, that is, stored in multiple buffers. Each buffer has an associated descriptor longword containing status information about the frame such as error conditions, frame length, and flags indicating the start and end of multibuffer frames. The ability of the bridge to chain small buffers to handle frames of various sizes increases the efficiency of the packet memory by minimizing the amount of unused buffer space. (Statistically, LAN traffic has a higher

and used by the queue manager and the translation processor subsystems.

Queue Manager Process

The queue manager subsystem operates on all frames in the receive queue to determine if they should be discarded, forwarded to the Ethernet/802.3 LAN, or received and processed by the bridge management entity. Discarded frames are returned to the receive queue; the remaining frames are placed on the forward or bridge queues. The queue manager constantly makes updates to the table of learned addresses based on source addresses observed on the FDDI LAN.

The queue manager's operational decisions are based on the following data:

- The frame descriptor containing assorted status information such as transmission errors and frame length
- The frame control field specifying the type of frame
- The type and quantity of frames previously received (used to prevent a flood of any one type of frame from blocking out other types)
- A learned database

content of shorter frames.) Thus, more buffers are made available to handle bursts of traffic. Additional information about buffer status is contained in the page table memory. This information is generated

containing addresses indicating on which side of the bridge each MAC is located and special filtering status information assigned to each address by network management

The gate array burster (GAB) allows the queue manager to access the FDDI packet memory. This application-specific integrated circuit (ASIC) is a specialized direct memory access (DMA) device. It is capable of moving selected fields or large sections of frames into or out of FDDI packet memory. The objects may be moved either into internal holding registers for examination by the queue manager engine or directly to destinations such as registers in the table lookup engine (TLU). Note that the GAB used in the queue manager subsystem is the same device used in the translation processor, which is discussed later in this paper. These two subsystems have many similar requirements, but each also has unique requirements. Using one GAB design for both subsystems reduced the overall development effort.

The table address RAM and the TLU are key components of the queue manager. The RAM contains a table of up to 16K 48-bit addresses. Each address also has status bits that determine what action the bridge should take when a frame's source or destination address matches address is found, the TLU presents that status to the queue manager processor. Otherwise, the TLU gives the queue manager processor a programmable status indicating whether to forward or to discard the frame. A second TLU port allows the TLU and the table address RAM to serve as slaves to the AP. Thus, destination address filtering for traffic received from the Ethernet/802.3 LAN and table maintenance can be performed.

To keep up with the packet arrival rate, the queue manager subsystem makes extensive use of pipelining. The queue manager engine operates concurrently on six packets. The TLU unit performs three searches concurrently: one each for the source and destination addresses on FDDI packets and one source or destination search on Ethernet/802.3 packets.

Discarding and Keeping Frames

The decision to discard a frame is based principally on the frame's address or its contents. The following are typical of frames that are discarded:

- Frames destined for nodes that the bridge

a particular address. The TLU is an ASIC with a port that is a slave to the queue manager engine. The queue manager engine inputs an address to the TLU which scans the RAM for that address. If the recognizes as not on the Ethernet/802.3 side of the LAN. Also, network management may specify addresses to be discarded regardless of location in the topology.

- Frames of either a reserved or undefined frame control type.
- Frames that are either too long or too short.

When a frame is discarded, its buffers are returned to the end of the receive queue by reassigning them in the page table.

Frames that are kept are placed on either the forward or bridge queues. Frames ultimately destined for the Ethernet/802.3 LAN are placed on the forward queue. Frames placed on the bridge queue, to be processed internally by the bridge, are of the following types:

- FDDI station management (SMT) frames
- Digital's extended LAN management sofware (DECelms) frames or maintenance operation protocol (MOP) frames
- Spanning tree frames, containing messages used to determine the network topology and turn individual bridge ports on or off to eliminate path redundancy
- Frames containing errors
- Frames placed on the bridge queue

transmit additional SMT frames on the FDDI LAN.

Counters

Each frame type is guaranteed a minimum amount of processing time by the bridge. If at any time the bridge holds too many of any one frame type, it discards further frames of that type. The queue manager uses allocation counters to keep track of the number of forwarded, FDDI SMT, bridge management, spanning tree, and error frames.

The queue manager also has counters that summarize its activity. These counters are periodically dumped to the AP and are used to calculate LAN utilization statistics required by network management.

Translation

Bridges operate at and below the data link level in the seven-layer International Standards Organization (ISO)/Open System Interconnection (OSI) reference model shown in Figure 6. The data link layer is divided into a lower MAC sublayer and an upper logical link control (LLC) sublayer. The LLC

not forwarded to the Ethernet/802.3 LAN.

protocol is specified in ANSI/IEEE standard 802.2.[3]

However, after receiving and processing these frames, the bridge may generate one or more frames on either or both LANs. For example, received SMT frames are never forwarded, but a given SMT frame may cause the bridge to

When forwarding frames from one LAN to another, the DECbridge 500 device converts the outgoing frame to the MAC frame format of that LAN. This process is called translation. Also, when a frame is generated by the DECbridge 500 product on either LAN, the data link frame format of that LAN is employed.

By performing translation, the DECbridge 500 product complies with the IEEE 802.1d requirements for transparent bridging. This enables end nodes to communicate across the extended LAN as if the nodes are directly connected to the same LAN. An alternative to translation, called encapsulation, is possible, but it does not comply with the IEEE 802.1d requirements. Further, using encapsulation puts restrictions on the configuration of the network.

The Process

Ethernet, 802.3, and FDDI have different MAC frame formats. When Ethernet or 802.3 frames are bridged to an FDDI LAN, they are reformatted to the FDDI MAC frame format. The original MAC type (Ethernet or 802.3) is indicated by 802.3 MAC protocol. IEEE standard 802.1 defines a mechanism for translating Ethernet frames into an IEEE 802.2 format (as is used on FDDI LANs). Figure 7 illustrates how Ethernet frames and two types of 802.3 LLC frames are translated into three different types of FDDI data link frames.

Maximum and minimum frame sizes of the LANs also impose requirements on the translation process. Ethernet and 802.3 MAC protocols require a minimum data field length of 46 bytes. The FDDI MAC protocol supports zero-length data fields. When a bridge forwards frames that originated at nodes on an FDDI LAN to an Ethernet /802.3 LAN, the translation process must add padding (null bytes) to any short data fields to bring them up to the 46-byte minimum size.

FDDI has a maximum frame size of 4500 bytes. The Ethernet/802.3 maximum frame size is 1518 bytes. Frames received from the FDDI ring that are longer than 1518 bytes after translation are discarded with the exception of frames discussed in the text that follows.

The internet protocol

setting information in the LLC header. If the frame passes through a second FDDI-to-Ethernet/802.3 translation at another bridge, the LLC information is used to determine if the bridge should translate the frame into Ethernet or (IP) is a widely used, network-layer protocol. Nodes on FDDI rings may generate IP frames longer than the Ethernet/802.3 maximum size. The DECbridge 500 product performs one function beyond the process

of transparent bridging. The bridge breaks up large IP packets into smaller ones. This function is supported by IP and is called fragmentation. Without fragmentation, the queue manager would discard these long IP frames, preventing communication between nodes on separate FDDI rings that are linked by Ethernet/802.3 LANs.

Since the translation process alters frames, the original cyclic redundancy check (CRC) field is no longer valid. In the DECbridge 500 device, the translation process concurrently verifies the received CRC, translates the frame, and generates a new CRC. This concurrent processing results in a high degree of data integrity.

Address Bit-ordering

The bits of the destination and source addresses are transmitted in the reverse order on FDDI from that on Ethernet/802.3 data links. Digital's FDDI chip set performs a bit-reversal operation on receive and transmit for only the MAC frames' destination and source address fields. Since these MAC fields are stored inside the bridge in IEEE 802.1, canonical bit-ordering, only one version of each address needs to be kept in the forwarding database. Also, when generating management messages, this method of frame storage allows the bridge to move an address from the source or destination field of a received frame into the data field of the management message without modification.

The translation processor consists of principally a GAB and a translation engine (based on a 68020 subsystem). The GAB and translation engine interactively copy frames from FDDI frame memory to Ethernet/802.3 frame memory. Concurrently, the translation engine makes changes to the frame format, and the GAB calculates both the CRC of the incoming frame, using old bit-ordering, and generates the CRC of the translated frame, using both new bit-ordering and new frame format.

Frames from the forward queue in FDDI frame memory are thus translated and moved to the forward queue in Ethernet frame memory. Frames from the bridge queue are separated into management and spanning tree queues in Ethernet /802.3 frame memory. The translation processor returns buffers from the forward and bridge queues to the free queue in FDDI frame memory. The queue manager returns buffers from the free queue to the receive queue, making them available to store newly received frames.

NI-side Processing

Frames placed in output queues by the translation

When calculating the CRC on incoming packets, or generating a new CRC on forwarded packets, the translation process must take into account the bit-ordering.

Implementation

processor are processed next by the AP. Frames in the spanning tree and management queues are destined for the bridge as a manageable entity on the extended LAN. The AP processes these frames

and may generate response traffic on either the FDDI or the Ethernet/802.3 LAN. Frames in the forward queue are subjected to additional match or nonmatch filtering by the AP. These frames are checked against a list of protocol types loaded by network management (e.g., TCP/IP and AppleTalk protocols). Protocol filtering is often a useful mechanism to prevent all frames of one or more protocol types from propagating across the extended LAN. The AP also uses the table lookup engine to check frames against a list of source addresses loaded in the address table RAM for filter /forward requirements. Source address filtering may be used to contain traffic from nodes with an unusually high transmit rate. In additon, this filtering may be used as a form of security to deny access to nodes that are masquerading, that is, transmitting by using another node's address.

The bridge checks frames against protocol and source address lists after the frames have been filtered by destination address in the queue manager. The rate of frames here is lower, not exceeding the

The frames in the forward queue that pass the protocol and source address filters are placed on the transmit queue of the Ethernet/802.3 chip set. The AP must merge these frames into the transmit queue with management traffic that the AP has generated for the Ethernet /802.3 LAN.

NI-to-FDDI Forwarding

The bridge processes traffic received from the Ethernet/802.3 LAN in much the same way as FDDI LAN traffic processing was described in the preceding material. It is essentially a mirror-image process, but a few significant differences exist.

The lower arrival rate of frames from the Ethernet /802.3 LAN does not require a dedicated frame-processing engine such as the queue manager. Thus, destination address filtering is performed by the AP, which shares the TLU engine and table address RAM with the queue manager.

Also, allocation counters are not used on Ethernet /802.3 traffic. The AP directs all incoming traffic into different queues at full rate. Unusually high bursts of a particular frame type could

Ethernet/802.3 LAN rate. Performing such checking on all incoming traffic from the FDDI LAN would require significant additional computational work by the queue manager subsystem.

overflow a given queue. Another difference is a requirement for stations placing traffic on the FDDI ring. On token ring networks, the transmitting station is responsible for

removing its own frames from the ring. A typical station knows which frames to strip by recognizing its own address as the source address. When a bridge transmits a frame, the source address is that of the originating node. In the DECbridge 500 product, the stripping function is handled in the FDDI chip set. A bridge strip algorithm is implemented that generates frames marking the end of a block of transmitted frames and also makes use of counters for sent and stripped frames.

Development Methodology

It was important to get the DECbridge 500 product to market in as short a time as possible. The solidification of the ANSI FDDI specifications, coupled with the appearance of products from different vendors, created a finite window of opportunity. At the same time, the requirements to be met were significant. The next three sections present brief descriptions of some elements of the development methodology that were employed to meet the design requirements while optimizing the development schedule.

Utilization of Existing Electrical, Firmware, and Mechanical and Power. System-level Design

The AP and the Ethernet /802.3 packet memory of the DECbridge 500 product correspond approximately to the processor and memory of its most recent predecessor, the LAN Bridge 200. The high FDDI packet rate required the use of a separate processor to filter incoming FDDI traffic. Also, another dedicated processor was necessary to perform the translation function. (Ethernet-to-Ethernet bridges do not require a translation function.) The resulting increase in the rate at which a processor accesses frame data required the development of a separate packet memory for the FDDI LAN.

Electrical Design

The Ethernet interface and packet memory designs again were borrowed from the LAN Bridge 200 product, but several extensions were needed. The AP design is very similar to the design of the processor in the LAN Bridge 200, but it has several new features, namely, a down-line, loadable program memory, a bus system for communicating over the backplane with other

Technology

The DECbridge 500 developers combined technology from existing products with their own new technology in the following design areas: system level, modules, and a distributed interrupt system. The queue manager, the translation processor, and the FDDI chip set with packet memory are new designs. The additional circuitry

resulted in a multimodule system with a backplane.

Firmware Design

The DECbridge 500 product uses the same operating system as other Telecommunications and Network products. Much of the firmware associated with the bridge entity and with Ethernet-side processing was modified from the LAN Bridge 200 product. The queue manager and translation processor required all new code.

Mechanical and Power Designs

Previous products typically consisted of a single module mounted in a box. The DECbridge 500 device required developing a multimodule system with a backplane. The initial goals were to install two, or at most three, logic modules. To minimize the risk to the module development schedule, a four-board approach was adopted, which closely follows the block diagram shown in Figure 2. The box, the backplane, and the power supply are all new designs.

Integration of FDDI Products and Chip Set Development

A strategy was adopted to maximize the commonality of effort in the development test bed for the FDDI chip set. The test bed design was expanded midstream so that separate modules could be added, turning it into a breadboard for either a DECconcentrator or a DECbridge device. The two DECbridge modules contained the queue manager, the translation processor, and the Ethernet/802.3 chip set and packet memory. The test bed provided the FDDI interface, an FDDI packet memory, and an application processor, as well as a power/packaging platform. While evaluation of the breadboards was still taking place, activities were accelerated to develop the products.

Technical Risk Analysis

Different approaches were adopted for various parts of the bridge design based on technical risk. Completely new technology, e.g., the queue manager and the translation processor, were simulated, breadboarded, and tested. Areas that were understood but still new, e.g., packet memory designs, were evaluated largely by gate-level simulation. High-confidence areas, such as designs taken from previous products, were evaluated in the prototype products.

The DECbridge 500 product employs three processors.

of the DECbridge 500 and the DECconcentrator 500 products, and in the evaluation of the FDDI chip set. When a product set was defined, plans were in place to develop a hardware

Thus, a lot of the bridge functionality was in firmware, and changes could be made with relatively little impact on the schedule. Also, in several

instances, deficiencies found in the system-level design could be corrected in the firmware.

Use of Parallel Activities

Several parts of the usual development process were overlapped to minimize time. A combined functional and design specification was generated instead of going through two serial stages to produce separate specifications. In the hardware design, module layout started once a confidence factor was achieved through simulation. Design reviews were held concurrently with module layout, and performance simulation continued throughout the process. There was a close interaction of the printed circuit board layout group and the electrical designers.

In product qualification, a pipelined system of reliability qualification testing (RQT), process qualification testing (PQT), and internal /external field test was set up to accommodate a phased release of firmware. RQT and PQT started with a functional, subset release of the firmware. Hardware confidence grew. After electrical design verification testing, corrections to recognized problems. A process was developed whereby new releases were tested for a few days each in RQT and at internal field test sites and then released to external field test sites. The down-line-upgrade ability was instrumental in allowing us to use this process.

Conclusions

Differences in frame format, frame length, and transmission speed place requirements on an Ethernet /802.3-to-FDDI bridge that are not encountered in bridges between like data links. The DECbridge 500 product met these requirements by dedicating one processor subsystem to the translation process and another to the process of filtering and sorting incoming FDDI frames. By adhering to the requirements of the IEEE standard for transparent bridging, the DECbridge 500 device allows the problem-free interconnection of FDDI LANs to the large existing base of Ethernet /802.3 LANs.

The development team concluded that by performing risk analysis and having backup plans in place, several parts of

firmware with the minimal functionality for field test was tested briefly in RQT and PQT and then shipped to the field. New firmware releases were developed with increased functionality as well as the standard design process could be compressed or overlapped. Fundamental to the design was the ability to make remote, nonvolatile upgrades to the product's operating firmware.

Acknowledgments

A large number of people contributed to the success of the DECbridge 500 project. It is not possible to name all of these important contributors in the space provided. Therefore, the authors wish to acknowledge that the DECbridge 500 product could not have been developed without significant contributions from members of the Telecommunications and Networks (T&N) Communications Systems Engineering Group responsible for the diagnostics, firmware, hardware, product assurance, and software and from the following T&N organizations: Architecture, Customer Services Systems Engineering, Marketing, Manufacturing, Product Management, and Publications.

References

- F. Backes, "Transparent Bridges for Interconnection of IEEE LANs," IEEE Network, vol. 2, no. 1 (January 1988).
- Local Area Network MAC (Media Access Control) Bridges, IEEE Standard 802.1(d) (New York: The Institute of Electrical and Electronic Engineers, Inc., 1990).
- Logical Link Control, ANSI/IEEE Standard 802.2-1985, ISO/DIS 8802/2 (New York: The Institute of Electrical and Electronics Engineers, Inc., 1985).