The DECconcentrator 500 Product

By William J. Tiffany, G. Paul Koning, and James E. Kuenzel

Abstract

Digital's decision to implement the fiber distributed data interface (FDDI) physical topology with a dual ring of trees, as opposed to a dual ring only, resulted in the development of

the DECconcentrator 500 product. The dual ring of trees topology provides high availability, manageability, and support for building wiring standards. The function of the concentrator demanded that the product be reliable, provide for remote management and control, and allow a low cost per connection. The use of common FDDI hardware and software components developed by Digital helped the product team to meet these goals.

Concentrators in the ANSI FDDI X3T9.5 Standard

In the initial stages of its development, the American National Standards Institute (ANSI) standard Interconnect (CI) bus. All stations were to be dual attachment stations (DASs) and the interconnections between the stations were to be wired directly, without patch panels or similar structured wiring schemes.

Using this dual ring topology is feasible for a small number of stations in a single, tightly controlled room. However, soon the emphasis of FDDI shifted primarily to local area networks (LANs). A LAN consists of many nodes, spread over a large area and with potentially many individuals able to connect and disconnect stations. To accomodate LAN topology requirements, ANSI chose to add the concept of concentrators midway in the development of

the FDDI standard.[1] In the simplest case, a concentrator is a device that attaches to a dual ring (via A and B ports) and provides additional ports (M ports) to which for the fiber distributed data interface (FDDI) technology was intended for a computer room system interconnect, similar to Digital's Computer

stations can be connected by means of radially wired cables. These additional stations can be single attachment stations (SASs) with a single port (S port) rather than the pair of ports required by a DAS. This simple topology was soon generalized, allowing concentrators to be nested to any depth in a dual ring of trees. Concentrators may be singly attached (by using an S port plus M ports rather than A and B ports plus M ports), and DASs may be connected to concentrators. Figure 1 illustrates the basic FDDI topologies.

FDDI concentrators are more than wiring hubs, unlike certain other LAN technologies. They also perform two functions that are key to network integrity. When a station's connection to a concentrator is activated, a multistep initialization procedure called physical connection management (PCM) takes place, using physical layer (PHY) signaling. In this procedure, the station and the concentrator exchange some topology information, and a link confidence test is performed to verify that the data integrity on the link is acceptable. Once the PCM initialization is complete, the connection becomes part of the ring.

The ANSI standard specifies some topology rules to reduce this problem.

Also, concentrators continuously perform link error monitoring (LEM). Each active link is monitored for data errors, and a link found to have excessive data errors is disabled. In this way concentrators ensure that the ring error rate, and therefore the packet loss rate, remains acceptably low.

Given the many choices for concentrator interconnection allowed

by the ANSI standard, it is possible to construct highly complex topologies, including many that have "bad" properties. When a station is physically plugged in and that connection is operating properly, the station should be able to communicate with all other stations in its own network. This property is often stated as "Physical connectivity equals logical connectivity," or, in other words, "Being plugged in implies being able to communicate." In bad topologies, this property does not hold. Such topologies are very confusing to network managers and are therefore undesirable.

therefore depends to some extent on the competence of the network manager to

However, the decision to accept or reject offered connections is based only on local knowledge (i.e., information held locally in each station or concentrator), and it is not possible to detect all the bad topologies. FDDI avoid bad topologies.

One special case allowed by the topology rules is called dual homing, as shown in Figure 2. With a dual homing configuration, the A and B ports of a concentrator or a DAS are connected to M ports, usually of two different concentrators. In this case, the B to M connection becomes active, and the A to M connection remains in a "standby" state. The standby connection is not part of the ring, but it can quickly change to the active state if the B to M connection breaks. In this way, connectivity is maintained when failures occur.

We next present the reasons Digital chose the dual ring of trees topology over the dual ring topology and examine the resultant need for a concentrator. A detailed discussion of the development of the DECconcentrator 500 product follows.

Digital's Choice of the Dual Ring of Trees Topology

A dual ring topology consisting of dual attachment stations may be appropriate in very small networks that do not use structured and permanent cable plants. But a dual ring of trees topology, using single attachment stations as the end-user devices and concentrators as hubs, provides the optimal solution for a FDDI networks has been recognized by many of our key customers. Digital chose the tree /dual ring of trees architecture implemented with concentrators over a purely dual ring approach because this architecture offers:

- Support for industry-standard radial wiring practices
- Manageability
- Configuration flexibility
- A definable demarcation point between the end user and the backbone
- Scalability

Although the specific behavior of an FDDI concentrator is relatively new, the concept is not. Most system vendors and users of large systems have adopted the use of manageable hubs (multiport repeaters) for Ethernet networks and media access units (a type of passive concentrator composed of bypass relays) in token ring networks.

Initially, cost was a major concern in the decision to implement a tree-type architecture. Some users saw the addition of the concentrator as an added highly flexible, reliable, available, maintainable, and robust FDDI LAN. As described in the preceding section, the concentrator is the cornerstone of an FDDI network. Its significance in building large, robust, and most importantly manageable cost burden. However, the cost increase, which can be amortized over the entire network, is greatly outweighed by the added advantages. With a concentrator, stations can be separated into two categories, end-user devices and backbone devices. This is especially important in large networks where the functions of network administration and users of the LAN are totally disparate. The concentrator becomes a tool that simplifies the role of network managers. As a demarcation point between end-user devices and the backbone, the concentrator protects the backbone from inadvertent disruption caused by the end user.

As shown in Figure 3, the actual number of components that are required to connect eight FDDI stations, whether into the dual ring or in the tree, is approximately the same; only the distribution of these components is different.

In both topologies we have eight physical connections. A physical connection, whether in the tree or in the dual ring, is a point-to-point, full-duplex path between adjacent physical layers. The initial reasoning for a dual, counter rotating ring was to create a bidirectional data path between adjacent stations in which the secondary path's main purpose is to assist in startup, initialization, and reconfiguration of the primary ring.

Either topology requires roughly the same number of components, i.e., PHY entities, optical transceivers, cables. With the tree, the PHYs are rearranged so that conceptually we have taken a PHY entity from each DAS, in conjunction with a data path switch, to create the concentrator. This approach does incur the additional network, made possible by the concentrator.

A large network of many stations has a high probability of disruptions. Although a DAS and a concentrator are fundamentally different, they have in common the role of controlling the network topology through PHY-level signaling. In the case of a disruption, whether an operator-initiated function (i.e., a station powered down, installed, or removed) or a failure of the station or cable, the token path is modified according to the station management /connection management (SMT /CMT) algorithm to maintain a continuous logical ring.

The main difference between the two approaches to solving the disruption problem is in the way a station is bypassed. With the dual ring approach shown in Figure 4, when cost of the power and packaging. However, from the perspective of the network administrator, this small incremental cost is offset by the increased ability to manage and control the

a disruption occurs, the stations adjacent to the disruption bypass the offending station(s) and reconfigure the ring by wrapping the secondary and primary rings to form a new single continuous ring.

This provides a degree of fault tolerance but is limited to only a single disruption.
 In the case of multiple failures or disruptions, all dual attachment

stations adjacent to the faults reconfigure, thereby creating multiple disjoint rings. Even though the majority of the stations in the network might be operational, they would operate over several disjoint networks. The potential loss of the service access point would effectively leave the network nonoperational from the client/server perspective. Management of such a situation would also be an ordeal, since access to fault information would be limited to the stations remaining on the portion of the ring to which the management station was directly attached.

 An FDDI concentrator provides fault tolerance in a different way, as illustrated in Figure 5. When a station connected to a concentrator is removed or powered off, the failure is bypassed through the concentrator data path switch at the PHY level. Any one or all of the stations can be effectively bypassed through the concentrator without affecting the connectivity of the other stations or the global topology of the FDDI network.

Structured Wiring

To fully appreciate the benefits of the concentrator, let us consider an FDDI network implemented in an office building environment in conjunction with structured wiring.

A typical building environment includes wiring between offices and equipment rooms on the same floor and in between floors. The wiring is permanent and involves a relatively large number of end-user devices as well as backbone devices over moderate distances. Moreover, frequent adds /moves/changes occur in this environment, and the ability to move from one location to another without manual intervention or network disruption is desirable. A clear demarcation between end-user devices and backbone is required to maintain the integrity of the backbone and to minimize disruption to, or manipulation of, the backbone cabling.

The Telecommunication Industries Association (TIA), together with the Electronics Industries Association (EIA), is defining a commercial building wiring standard; draft EIA/TIA 568 has the framework as designated in Figure 6.[2] According to this standard, end-user devices in offices should be wired from a telecommunications closet (TC). All closets in each building then connect to the intermediate cross connect (IC) or to the building hub for that building. In turn, all building hubs connect to a single main cross connect (MC) or to the campus hub in the campus.

Single Attachment Stations

The tree topology, as illustrated in Figure 7, facilitates structured (or radial) wiring as prescribed by the draft EIA/TIA 568 standard. The end-user devices, implemented as SASs or DASs, connect to the so that a disruption in an end-user device, such as disconnecting a station, does not affect the operation of the network. The concentrators in the various closets connect to root concentrators in the building hubs. If other backbone devices are concentrators located in telecommunications closets, which are maintained and controlled by the network administrator. When concentrators are used, the most cost-effective user stations are SASs. Connection to concentrators keeps the end-user devices separate from the backbone present in the building hubs or communication closets, such as bridges, they can also be connected to concentrators in the building hubs or the campus hub, as appropriate. The tree topology offers the fault tolerance as well as configuration flexibility required in a structured wiring system. Also, it allows for adds /moves/changes without disrupting or manipulating the backbone cabling.

Dual Attachment Stations

End-user devices directly attached to the dual ring, however, are not easy to isolate from the backbone LAN. Both the end-user devices and the backbone devices are part of the same physical loop, as shown in Figure 8. To a network administrator, management and control of the backbone becomes an ever-increasing ordeal because each end-user station is now considered part of the backbone. Even though rules for the end-user behavior can be established, they cannot easily be enforced.

The availability of the backbone is increased by the use of concentrators, since these are the only devices that form the dual ring backbone. This benefit is very important for a large network. For example, in a network supporting 200 end-user stations on a dual ring of trees topology, if 8-port concentrators are used to connect them to the dual ring backbone, only 25 concentrators reside on the dual ring backbone. The reconfiguration of the backbone is dependent on 200 devices. The probability of having two or more disjoint rings is much higher in the latter case. Also, with DAS stations, the network administrator is faced with the impractical task of directly controlling 200 devices.

Fault-tolerant Configuration Options

Two fault-tolerant configuration options are available: bypass relays and dual homing. Bypass relays may be used with DASs directly attached to the dual ring to provide fault tolerance in addition to the single fault protection provided by wrapping to the secondary ring. Dual homing is an alternative mechanism which allows dual attachment devices to have a redundant connection to a concentrator when installed in a tree topology. These two alternatives are examined in this section.

Bypass Relays

To avoid the aforementioned reconfiguration problems with DAS implementations, the FDDI standard offers an option of using an optical bypass relay. While such relays are the backbone is dependent on only 25 devices. Also, the network administrator needs to control only 25 devices. In contrast, if the same 200 stations are DASs directly attached to the dual ring, the

envisioned to alleviate some of the reconfiguration problems, they may induce more problems than they solve. The inclusion of relays in the network means added cost of components, cables, and connectors, loss of optical power, reduction in interstation distance, and an additional failure mechanism. These factors limit the use of such relays to possibly very small, physically collocated work group LANs and make the relay an unattractive solution for a large network environment.

Let us assume that the fiber-optic cable has a loss of 1.5 dB per kilometer (according to the EIA/TIA-492), and the distance between the communication closet and each station is 50 meters.[3] Since the cable length between stations A and C is 200 m, the power loss is 0.3 dB. A connector has a loss of 0.7 dB (according to the EIA). Since there are eight connectors (labeled 1 through 8 in Figure 9) between A and C, the power loss is 5.6 dB. A bypass relay has a loss of 2.5 dB. Since there are three relays between A and C, the power loss is 7.5 dB. The total link loss between stations A and C, therefore, is 13.4 dB, which is in excess of the maximum allowed by the FDDI standard. Note that with the use of optical bypass

As shown in Figure 9, the end-user stations A, B, and C are dual attachment with bypass relays, and station B has failed. With B bypassed, stations A and C become adjacent. The total link loss between these two stations, using a fiber-optic cable of 62.5-micron core diameter and 125-micron cladding diameter, must not exceed 11 decibels (dB) to comply with the FDDI standard.

For some applications a tree or dual ring of trees may not meet the customer's requirements. The dual homing topology, shown in Figure 2 and described earlier in the paper, has none of the limitations or problems that can be imposed by the dual ring. This topology is beneficial in a large campus to connect remote buildings into the FDDI backbone. It allows standard radial wiring practices, with up to 2 km between the campus hub and any given building through multimode fiber (MMF) links. The dual homing topology is especially useful with long distance links utilizing single mode fiber (SMF), which span distances of up to 40 km. With the dual ring topology each span has to be counted four times towards the fiber-optic path length maximum relays, the effective distance is limited to less than 200 m, which is far below the maximum of 2 km allowed by the standard.

Dual Homing

of 200 km allowed by the standard, in the event that a wrap occurs. For a link of 40 km, 160 km has to be subtracted from the maximum of 200 km, thus leaving only 40 km for the rest of the network. Using a tree topology, the span is counted twice as it has only one active fiber pair.

Product Development

The following sections examine the DECconcentrator 500 product development process from beginning to end, elaborating on details of the product functionality as they were refined along the way.

Several key factors provided a smooth product development process. First, the architecture for the DECconcentrator 500 product was chosen as a subset of the generality allowed by the ANSI FDDI standard.

Several features which would have significantly complicated the product without greatly enhancing functionality were not included. Examples of these are "roving MAC" and the ability to allow stations to select the ring (primary or secondary) to which to attach. Second, product management established a clear list of priorities, requirements, and goals that allowed the development team members to focus their efforts. The absence of significant changes in architecture or product requirements during the development simplicity and reliability enabled us to keep product transfer cost to a minimum and to nearly meet the time-to-market goal. And finally, the relatively small and tightly knit development team stayed together from conception through field test and first product shipment.

Hardware Partitioning

As the hardware block diagrams were developed, several concepts for partitioning the hardware into modules were evaluated. The electrical, mechanical, and power supply designers worked closely together to choose a suitable partitioning.

The initial high cost of the FDDI fiber-optic transceivers led the designers to select modular hardware partitioning. A modular design allows ports to be added according to the customer's needs, thereby minimizing both the initial cost and the number of unused ports. Since Digital's networking products have traditionally used side-to-side airflow for cooling, a card cage that supported horizontal modules was chosen. Four ports per module was considered to be a reasonable number by which a customer could helped the team stay on schedule. The priorities selected were to design for low cost and high reliability, provide for ease of firmware upgrades, and strive for quick time to market. The emphasis on

increment a system. This module, referred to as the port module, contains four sets of the FDDI physical layer chip sets, one status light-emitting diode (LED) per port, one module field-replaceable unit (FRU) fault LED, and a small amount of support logic.

One function of the DECconcentrator 500 product is to provide support for network management; this requires data link layer hardware. In addition, the DECconcentrator 500 device must connect to the FDDI dual ring; this requires the A and B port types. The DECconcentrator 500 management module was designed to meet these needs. By combining the data link and A and B PHY port hardware, we ensured that any DECconcentrator 500 device installed in a dual ring of trees would be manageable.

The DECconcentrator 500 product also includes a microprocessor to execute diagnostic and operational firmware. To minimize the number of modules, we specified that the microprocessor and its support logic fit on the backplane. The use of an active backplane eliminates the need for a separate control module in the card cage, thereby reducing both cost and the vertical height of the box. The backplane also provides the token ring data path function which interconnects any allowable configuration of port and management modules. The number of end-user configurations. Two basic configurations are supported in the DECconcentrator 500 product. A concentrator configured with one to three port modules (4 to 12 ports) can support a standalone work group but cannot connect in the dual ring of trees topology. A concentrator configured with a management module and one or two port modules supports the dual ring of trees topology and is remotely manageable.

Another goal of the hardware team was to eliminate the use of cables within the box. This goal was consistent with minimizing cost and maximizing reliability. The use of modular port and management cards led the team to believe that the power supply could also be modular and plug into the backplane in a similar manner. To avoid potential safety hazards, the power supply module is not accessible without opening up the box; however, the interconnection of the supply with the backplane is achieved with the same type of connector used on the logic option modules. The only cable used in the modules supported by the backplane is based on our

evaluation of customer need. We decided that 8 to 12 ports per concentrator is sufficient for most

DECconcentrator 500 device provides power to the fans.

Figure 10 is a diagram of the various modules that compose the DECconcentrator 500 hardware.

Power and Packaging Trade-offs

Once we decided that three modules could support 8 to 12 end users, we focused on the selection of packaging. Two basic proposals were examined. The first was to modify the Digital's NAC common box which has been used in many of Digital's products. The second proposal was for a new box design, which allowed improved serviceability via quick access to all FRUs (field-replaceable units). The existing common box design was chosen to minimize risk to the product development schedule. The necessary modifications were the addition of a card cage for the port and management modules, provision for mounting the power supply and fans, and improvement of the airflow characteristics.

Cooling was also seen as a potential problem in product development. The bipolar logic used in the FDDI physical medium dependent layer (PMD) dissipates a considerable amount of power in a small area. Evaluation of the NAC common box showed that the grill area in each side had approximately 35 percent open area. Analysis open area in each side of the box. When prototypes were tested, we found the improvements in cooling followed predictions. The modification also yielded considerable reduction in acoustic noise levels, which allowed the use of off-the-shelf ball bearing fans with good reliability.

Mechanical and electrical requirements could not be met by any of Digital's existing power supplies. The power supply height was limited; a unique connector was required to interface to the backplane; and the available area was determined by the size of the card cage on top of which the supply was mounted. Electrically, the supply had to provide a relatively large amount of minus5.2-volt power to support the emitter-coupled logic (ECL) in the FDDI PMD. Fortunately the total power requirement

was similar to that of Digital's Ethernet-to-Ethernet bridge product, the LAN Bridge 200. The power supply group agreed to modify this existing high-volume supply to meet the requirements of the DECconcentrator 500 product. The use of an existing design was expected to result in fewer bugs, and this proved to showed that significant improvements in airflow could be achieved by increasing this percentage. Mechanical rigidity was traded off against airflow improvement to reach a compromise of 50 percent

be the case. Only one bug was found in system stress testing, and it was easily corrected with a minor design change.

Card Handles

The module handle design was probably the single most complex part of the mechanical design, and it had the potential to impact the cost of the product. The mechanical design team recommended modifying a handle design from an existing Digital product. This essentially meant stretching the handle and making openings as required for I/O connectors and LED displays. Handles for both DECbridge 500 and DECconcentrator 500 products are processed from the same mold since only I/O connector and switch/indicator openings are different for the two products.

Traditionally, these handles have been made from a plated cast alloy. Plastic offered the potential of significant cost savings and weight reduction. However, there was concern about the quality of plating with plastic, as well as about the structural strength. The decision was made to start the development effort for plastic handles while using machined metal parts for the interim. When the final plastic product was available, it met all requirements.

Another critical factor in handle design radiation. A "waveguide-beyond-cutoff" structure was evaluated. This structure is a rectangular extension to the handle with an opening the size of the FDDI connector. The design attenuates all emissions below the cutoff frequency and was predicted to provide excellent attenuation performance for all harmonics of the FDDI signals. Testing of prototypes verified the emissions problem created by the connector opening and proved the effectiveness of the waveguide structure in eliminating the emissions. This structure was then included in the design of the handles for each FDDI port.

Logic Design

The logic design team developed prototype hardware as quickly as possible so that the diagnostic, operational firmware, and common node software (CNS) firmware teams could proceed with hardware-based debugging. First-pass prototypes of the controller/backplane module and the four-port module were in the laboratory within six months of the start of the project. The GenRad HILO simulation software was used in the module design was electromagnetic interference (EMI). Each FDDI duplex connector required a large, 1.4 cm by 3.8 cm, opening in the handle which posed the risk of emitting

process.

One type of bug was discovered in the first-pass prototypes that was not caught in simulation. A through-hole component body was used in schematic capture instead of a surface mount body. As a result, the layout was wired according to the through-hole pinout. This error was not caught by any of the software that checks design rules. Thus, the controller/backplane modules required engineering change order (ECO) wires to mount a component onto the back of the module in a "dead bug" fashion.

On the four-port module, the differential ECL signal detect lines from the fiber-optic receiver to the clock and data conversion receiver (CDCR) component were crossed. This logic was not included in the simulation due to its analog nature. The problem in this case was an inconsistency in nomenclature between the CDCR and fiber-optic receiver chip bodies used in schematic capture.

The strategy of building first-pass hardware prototypes as quickly as possible to support early firmware debugging was successful. Simulation played a key role in guaranteeing functionally correct designs. When second-pass prototypes were tested, only a single ECO wire was required in the product, and the

Concentrator Port. An FDDI concentrator consists of a group of serially connected ports, each of which implements the FDDI PHY functionality. Key to our product was Digital's PHY chip, which implements the physical layer functionality and supports the physical connection management requirements for station management. In addition to the PHY chip, the CDC chips (one each for transmit and receive) provide the serial/parallel conversion, clock recovery on the receive side, and nonreturn to zero inverted (NRZI) encoding/decoding. The fiber-optic transceiver, however, had to be purchased from an outside vendor. Since many mechanically incompatible devices are on the market, we tested the products of many vendors. Fortunately for the product development teams, our optics group was able to identify pin-compatible transceivers from two vendors. Dual sourcing of the tranceivers used in our concentrator ports reduced the risk of shipping products based on this relatively new fiber-optic technology.

Internal Token Ring Data Path. The FDDI PCM process provides fault coverage and topology rule checking between any product shipped with the second-pass designs on all modules. The following sections examine several areas of the logic design that are unique to the concentrator function.

two connected FDDI ports. This is essential to ring stability since a token ring is made up of a series of physical connections, any one of which can bring down the entire ring.

Because connections between ports within a concentrator are not specified by the FDDI standard, they are a critical design area for ring availability. The DECconcentrator 500 design addressed this aspect of the product as described below.

- Data path (port-to-port) integrity. Adjacent ports in the concentrator interface with a dual-symbol wide data path of 10 bits plus parity for a total of 11 bits. This PHY-to-PHY interconnection is the same interface used to connect PHY to media access control (MAC) in a station. Parity checking was added to Digital's PHY chip to ensure that intermittent or hard faults could easily be detected. If a parity error in this data path occurs, the DECconcentrator 500 device is taken off line to ensure that the entire token ring is not corrupted. Without parity protection, a hard failure on this internal path stops ring traffic altogether, in which case the SMT trace function might isolate the fault. However, an intermittent fault in a concentrator's internal data path that the mechanisms built into SMT.
- Fault detection /isolation. The controller/backplane and the PHY chip design allow the DECconcentrator 500 device to offer continuous service in the presence of hardware faults by isolating the faulty hardware from the data path. The diagnostics that are invoked at power up or on command from firmware (as in the SMT trace function) have the ability to isolate faults very close to the component level. The fault report is then passed to initialization firmware which configures the DECconcentrator 500 product so that the faulty hardware is not included in the data path. Two levels of bypassing are provided, one at the port level and one at the option module level. Bypassing is always performed one level of hardware away from the detected fault. Thus if a fault is detected at the CDC component level (using a CDC loopback test), then tha single port is bypassed in the PHY chip. If a fault

is not protected with parity could arbitrarily reduce ring performance and increase the risk of undetected data errors and would not be isolated by any of

is detected at the PHY level or between PHY chips within a module, the entire module is bypassed on the backplane. Note that power-up diagnostics do not provide a PMD external loopback test except in a special manufacturing mode. Fault coverage of this hardware is provided by the SMT PCM process, which prevents a faulty connection from being included into the ring. Fiber loopback connectors are included with each product for isolating media faults between the fiber-optic transceivers and the fiber-optic cables.

- Internal MAC attachment. A MAC is not required for an FDDI concentrator to function; however, it is included as an option to provide remote management. The internal MAC can be thought of as a "management station" attached to one of the concentrator ports whose job is to provide control /status of the local concentrator function. This attachment is internal to the concentrator, but must provide the same basic service as a physical connection to an external station. This service is provided by logic in the data path referred to as the "null PHY." The null PHY provides a means of

  It also provides ring scrubbing in case the MAC should have to leave or enter the ring while the ring is operational.

Upgradeable Nonvolatile Program Memory. To support firmware upgrades over the network, the FDDI products require electrically erasable programmable read-only memory (EEPROM). All code in the DECconcentrator 500 product is executed directly out of EEPROM since the microprocessor's clock rate is relatively slow. In the first-pass design we used conventional 32K by 8 dual in-line package (DIP) devices as they were qualified within Digital. In order for the controller/backplane module to accomodate sufficient program memory, we needed a denser technology. At the time, component engineering was evaluating flash EEPROM technology. Flash devices became available in surface mount packages with a density of 128KB that met our needs. The flash memory proved to be a robust technology; however, development of a flash programming algorithm was challenging and required extensive testing. The older EEPROM technology had built-in logic to handle the details of the erasing and programming steps, bypassing the internal MAC if diagnostics should find a fault in any of the hardware [MAC, CAM, ring memory controller (RMC), and packet memory] related to the data link layer.

but with the flash memory these details were directly controlled by software.

In order to upgrade firmware over the network, a well-controlled procedure was developed. A firmware image plus the flash programming code is transferred over the network through multiple packets and stored in packet memory. This down-line upgrade is provided by a network device upgrade (NDU) utility that was developed for the FDDI products. Once the entire image is received in packet memory, it is checked against a cyclic redundancy check (CRC) included in the image. If the CRC is correct and the firmware image is of the correct type (destined for this product), the DECconcentrator 500 product takes itself off line. The 68000 microprocessor then executes the flash programming firmware directly from packet memory to load the new image. Once the load is complete, the firmware forces a reset, and a power-up self-test is run that includes a CRC check of the contents of flash memory.

Software Design Issues

Essential to the completion of the development process was the use of common software and the field testing of the DECconcentrator 500 product.

Common Software. Early in the development process it was clear that the aggressive time-to-identify ways to shorten the software development cycle. Code that could potentially be shared among the products, and code that could be ported from previous projects was identified. These early efforts resulted in the common use of the real time operating system (RTOS), common FDDI chip diagnostics, diagnostic error logger, diagnostic dispatcher, and common node software (CNS). In addition to the code that is common among the FDDI product set, much time was saved by porting portions of the remote bridge management software (RBMS) responder and maintenance operation protocol (MOP) from the LAN Bridge 200 product. The management model for the data link and physical layer entities for both the DECbridge 500 and DECconcentrator 500 products was developed to ensure commonality between the two products.

Field Test. The field test provided valuable information regarding the quality of the products. Several of the sites chosen to field test the FDDI backbone products were technically knowledgeable about networking. They were able to perform specific testing while monitoring their networks.

market goals for the FDDI product set would require the development team to be resourceful. In the beginning of the development cycle a significant effort was made

As a result, detailed test information was provided to engineering. One engineer was assigned to each field test site as a "site parent." The site parents monitored their sites and channeled the information back to engineering. This structure for supporting the field test enabled engineering to react quickly to the needs of the sites as well as act on problems found. This testing and feedback, coupled with the capability to load new firmware revisions over the network, was crucial to achieving quality prior to first customer shipment.

Conclusions

The dual ring of trees topology is well suited for all FDDI environments that require fault tolerance, scalability, and flexibility of configuration. This topology is the right choice for managing the ever-growing local area networks through the 1990s.

Several factors proved to be crucial to meeting both functional and time-to-market requirements with a quality DECconcentrator 500 product.

- Establishing architecture and product feature requirements early and maintaining these with minimal changes.
- Providing a thorough testing process for both hardware and software which tested the product in realistic environments with a process in place to correct problems and verify fixes quickly. We had quite robust products at the time of first customer shipment as a result of our test /fix/verify process.

Acknowledgments

People from many organizations worked together to achieve the development of the product on time. The key members of the engineering design team were Dave Benson, Gerry Capone, Stan Chmielecki, Paul Ciarfella, Alison Coolidge, Janis Roth Cooper, Joe Dagdigian, Tom Ertel, Cheryl Galvin, Dave Hyre, John Iannarone, Bill McCarthy, Charlie McDonald, Dick Muse, Luc Pariseau, Dave Sawyer, and Karen Shay.

References

- FDDI Station Management (SMT) Preliminary Draft Proposed American National Standard,ANSI X3T9/90-078, Rev. 6.2

o Establishing and maintaining a close-knit product development team with good communication channels.

o Leveraging wherever possible from proven and available designs, making incremental improvements as needed.

(May 1990).

o Building Wiring Standard for Industrial and Commercial Use, EIA /TIA 568, PN1907-

B (Washington, DC: Electronics Industries Association, Engineering Department, forthcoming 1990).

o Detailed Specification for 62.5-m Core Diameter /125-m Cladding Diameter Class Ia Multimode, Graded Index Optical Waveguide Fibers, ANSI/EIA/TIA-492AAAA-1989 (Washington, DC: Electronics Industries Association, Engineering Department, 1989).