The Ring-Based Local Networks Report

May 1986

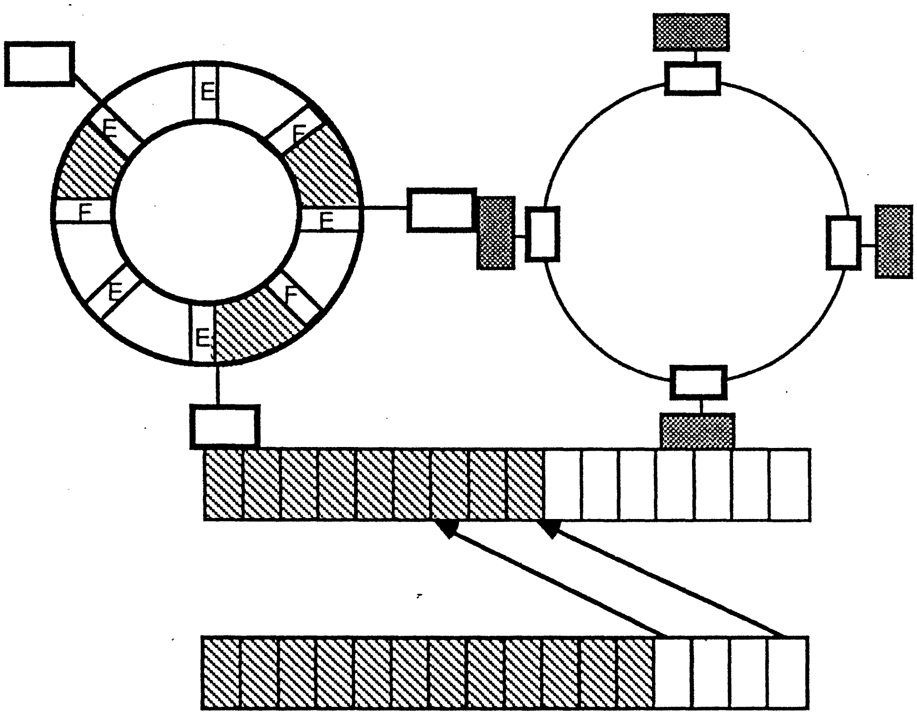

ARCHITECTURE
TECHNOLOGY
CORPORATION
SPECIALISTS IN COMPUTER ARCHITECTURE

P.O. BOX 24344 • MINNEAPOLIS, MINNESOTA 55424 • (612) 935-2035

Foreword

Now that IBM has introduced its Token-Ring Network, it has been theorized that all local networks will soon fall into two categories: TokenNet and TokenNot. Because token-passing ring technology has the potential to shape the future of the LAN industry, LAN vendors, users, and prospective buyers alike can profit from becoming aware of its history, standards, methodology, and performance, and the various commercially available systems that employ it.

This report examines not only IBM's "strategic" local area network offering, but each of the three basic kinds of ring-based LANs: slotted-ring (also called empty-slot), token-passing, and register-insertion. All three kinds are examined at length, in terms of origins, methodology, advantages and disadvantages, and systems currently in use. The slotted-ring approach, best embodied in the Cambridge Ring LAN, is the ring-based LAN of choice in England and Europe; token-passing is the most widely used ring-control technique in the United States, primarily because of IBM's Token-Ring Network and the IEEE 802.5 standard; register-insertion technology offers the most efficient use of a ring of the three methods discussed. The token-ring type of network will likely become the predominant LAN technology for the balance of this century, driven to this pre-eminence by its espousal by both manufacturing giant IBM and standards maker IEEE 802.

The Ring-Based Local Networks Report consists of three sections. Section One examines LANs in general and ring-based LANs in particular, in terms of history, standards, methodology, performance, reliability, types, and systems currently in use.

Section Two of this report is a survey of the key ring-based LAN systems available on the market. Each system description in this section features technical data and a price breakdown on the systems described therein.

This report concludes with Section Three, which consists of a copy of the reissued Soderblom patent on token-passing ring technology, a bibliography of sources used in the preparation of this report, and a listing of the addresses of the vendors covered in this report who are still selling ring based local networks.

Table of Contents

Section I

Section II

Section III

List of Figures

Section I

Section II

List of Tables

Section I

Section II

SECTION I

Introduction to the Technology

1. Introduction

1.1. Overview - Local Networks In General

1.1.1. Origins

1.1.1.1. Definition Of A Local Area Network (LAN)

Within computer circles, the decade of the 60s is looked back upon as the decade of the mainframe computer and the 70s as the decade wherein timesharing and the minicomputer came of age. There is little doubt that when the curtain draws down upon the final year of the 80s, this decade will have earned the label "the decade of the local area network."

What, exactly, is a local area network (LAN)? Oddly enough, mid-way through this, its decade of prominence, there is still some disagreement as to what, precisely, a LAN is.

One elemental definition of a LAN is that it is a communications network that supplies interconnection of a variety of data communicating devices (computers, terminals, peripherals, etc.) within a small area. (The small area referred to here is typically a single building, although LANs also commonly span several buildings -- such as on a military base or college campus -- and, not as often, sites "a few" miles apart.)

This basic definition has been expanded by some to include mention of a LAN's generally being owned by a single organization, and its employment of some type of switching element technology; others still would add some mention of a LAN's use of a single cable to this definition, and stress its dependence on a communications channel of moderate to high data rate, low delay, and low error rate.

What is most interesting here is that all the various embellishments to the basic definition of a LAN aptly serve to underscore the very reason for conflicts among the various definitions: most people define a local network in terms of what it does for them, or how it goes about performing these tasks. Since the technology is still evolving, so are its definitions and terminology.

1.1.1.2. LANs - Response To A Need

There is, however, a good deal more agreement within this same community as to how LANs came into existence. In the computer field, too, necessity proved to be the mother of invention.

As the numbers of mainframes, minicomputers, terminals, and peripherals within the office proliferated, the costs of adding new units to enlarge a company's growing resource needs became prohibitive, and the growing complexity of wiring schemes served as a deterrent to expansion.

Out of this potential bottleneck came the idea to distribute one's data processing capabilities: specialized functions, it was discovered, could be offloaded from a single host, developed as separate subsystems with optimized architectures, and used by other nodes on "the network." Such a distributed network -- especially a local network -- has the potential for resource-sharing, which gives users access to specialized facilities without having to directly own them. Further, distributed networks extend the useful life of separate computers and improve their performance by allowing them to share mass-storage and communications devices, at the front end, to obtain more processing capabilities.

Two other benefits of LANs that were also soon discovered is that there is now the possibility to interconnect components from different manufacturers, and that LANs offer a greater degree of fault-tolerance than standalone units: redundant elements in a network can assume the tasks of a failed component.

1.1.1.3. LAN Components

A local area network is composed of three basic hardware elements: 1) a transmission medium (e.g., twisted-pair wiring, coaxial cable, or fiber optics); (2) a control mechanism, to manage transmission over the medium; and (3) an interface to the network for the host computers or other devices -- i.e., the nodes of the network -- that are connected to the network.

Of late, however, a fourth basic LAN element has been added: a software element, a set of software protocols which control the transmission of information from one host or device to another via the hardware elements of the network. These software protocols are implemented in the host computers or attached devices, and they function at various levels of network operation.

1.1.2. Types of LANs

Local area networks can be based on any one of several technologies -- e.g., PBX/CBX (private branch exchange/ computerized private branch exchange), baseband, and broadband -- and they can incorporate different architectures -- such as centralized, bus, or loop. These terms will be explained in greater detail in the sections below.

1.1.2.1. Baseband Systems

Within local area network circles, a baseband LAN is generally defined as one that uses digital signaling (see Figure I-1). (More generally, "baseband" refers to the transmission of an analog or digital signal in its original form, without modulation.) With baseband LANs, digital signals are inserted on the line as voltage pulses, usually using a form of Manchester encoding. The total frequency spectrum of the medium is used to form the signal, and, consequently, frequency division multiplexing (FDM) cannot be used. Transmission is bidirectional: i.e., a signal inserted at any point on the medium propagates in both directions to the ends, where it is absorbed.

The digital signaling requires a bus topology: unlike analog signals, digital signals cannot be easily propagated through the splitters and joiners required for a tree topology. Baseband systems can extend only a limited distance -- about one mile at most -- because the attenuation of the signal, which is most pronounced at higher frequencies, causes a blurring of the pulses and a weakening of the signal such that communication over larger distances is impractical. Most baseband LANs use coaxial cable, though it can be used on twisted-pair wire as well. Baseband transmission is generally viewed as "data only" transmission.

The elemental baseband system, then, is the bus system. It consists of a bus (i.e., a shared communications medium) such as an Ethernet bus, and thus would have some form of coaxial cable. Connected along the cable are interface units which, in turn, are connected to a device such as a terminal, printer, workstation, CPU, or even a gateway which could be connected to additional devices or networks.

An interface unit can be connected to a single terminal or multiple terminals. Manufacturers' products can be distinguished by the number of terminals or number of devices that multiplex onto an interface unit. Multiple connections provide a low per-port cost when compared to a stand-alone unit.

In terms of protocol conversion, where a variety of devices exist, a variety of electrical protocols are available. If the protocol does not meet a particular standard, it is necessary to have a specialized connection or additional software. Thus, to create a baseband system, consideration must be given to the cabling plan for the interface unit, the properties of the interface unit, and any software that is available.

1.1.2.2. Broadband Systems

The concept of broadband resembles that of baseband: both consist primarily of bus structures; the bus is broken down into

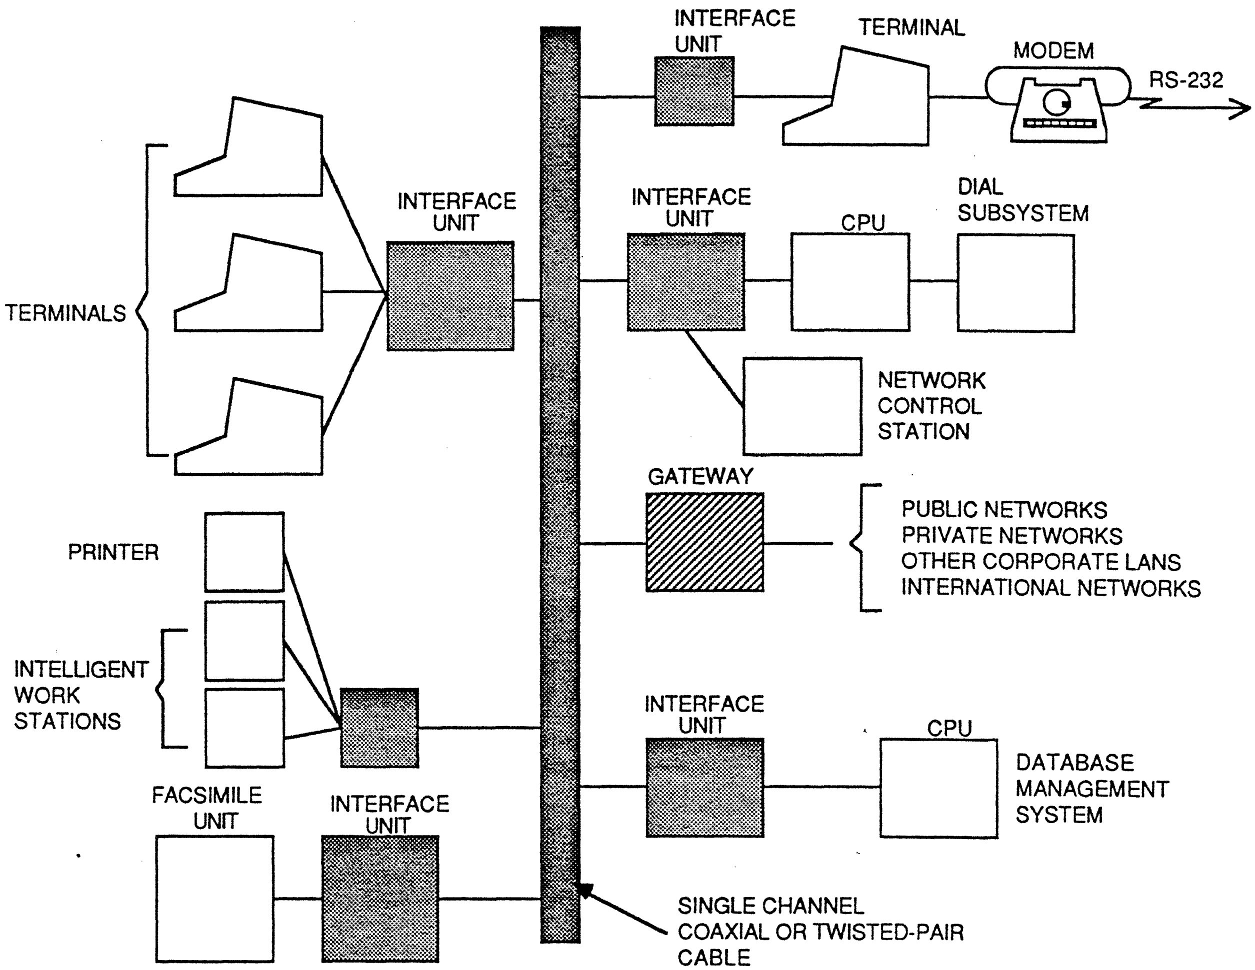

Figure I-1: Generic (Coaxial) Baseband LAN Model several different frequency spectrums; and there are interface units connected to terminals (see Figure I-2). The more popular applications are oriented toward systems in which terminals are connected together, or terminals are connected to minicomputers or hosts.

In general, broadband refers to any channel with a bandwidth greater than a voice-grade channel (4 KHz); within the realm of networking, it applies to coaxial cable on which analog signaling is used. Transmissions can be data, voice, and audio, while baseband is data only.

Whereas baseband transmission signals consume the entire bandwidth, with broadband transmission, frequency-division multiplexing (FDM) is possible -- i.e., multiple data channels, audio, and video. Broadband signals are unidirectional, and these signals can be transmitted up to distances of 10 kilometers -- a greater distance by far than that possible with baseband. Broadband can employ both bus and tree topologies, because its components permit joining and splitting operations.

1.1.2.3. PBX Systems

A private branch exchange (PBX) system is yet another means of networking information (see Figure I-3). The PBX merges two technologies: telephone exchange and digital switching. Basically, a PBX is an on-site facility, either leased or owned by an organization, which interconnects telephone stations within a facility and supplies access to the public telephone system.

Typical "buzzwords" within the PBX realm are "first-generation," "second generation," and "third-generation" systems -- terms coined to describe the three major evolutionary stages PBX systems are said to have experienced.

The "first-generation" PBXs were manual in nature, in that an operator at a switchboard made all connections; in the 1920s automatic systems replaced these operators (hence the name PABX (private automatic branch exchange). These systems employed electromechanical technology and analog signaling, and data connections were made by means of modems.

"Second-generation" PBXs came into prominence in the mid-70s; they employ electronic technology and digital switching and are sometimes called digital PBXs, or computerized branch exchanges (CBXs). These systems are essentially engineered to accommodate analog voice traffic, though they are also capable of handling digital data connections without use of a modem.

"Third-generation" systems claim to be "integrated voice/data systems," but, in reality, the distinctions between upgraded second-generation and third-generation PBXs are very blurred indeed. Typical characteristics ascribed to this newer

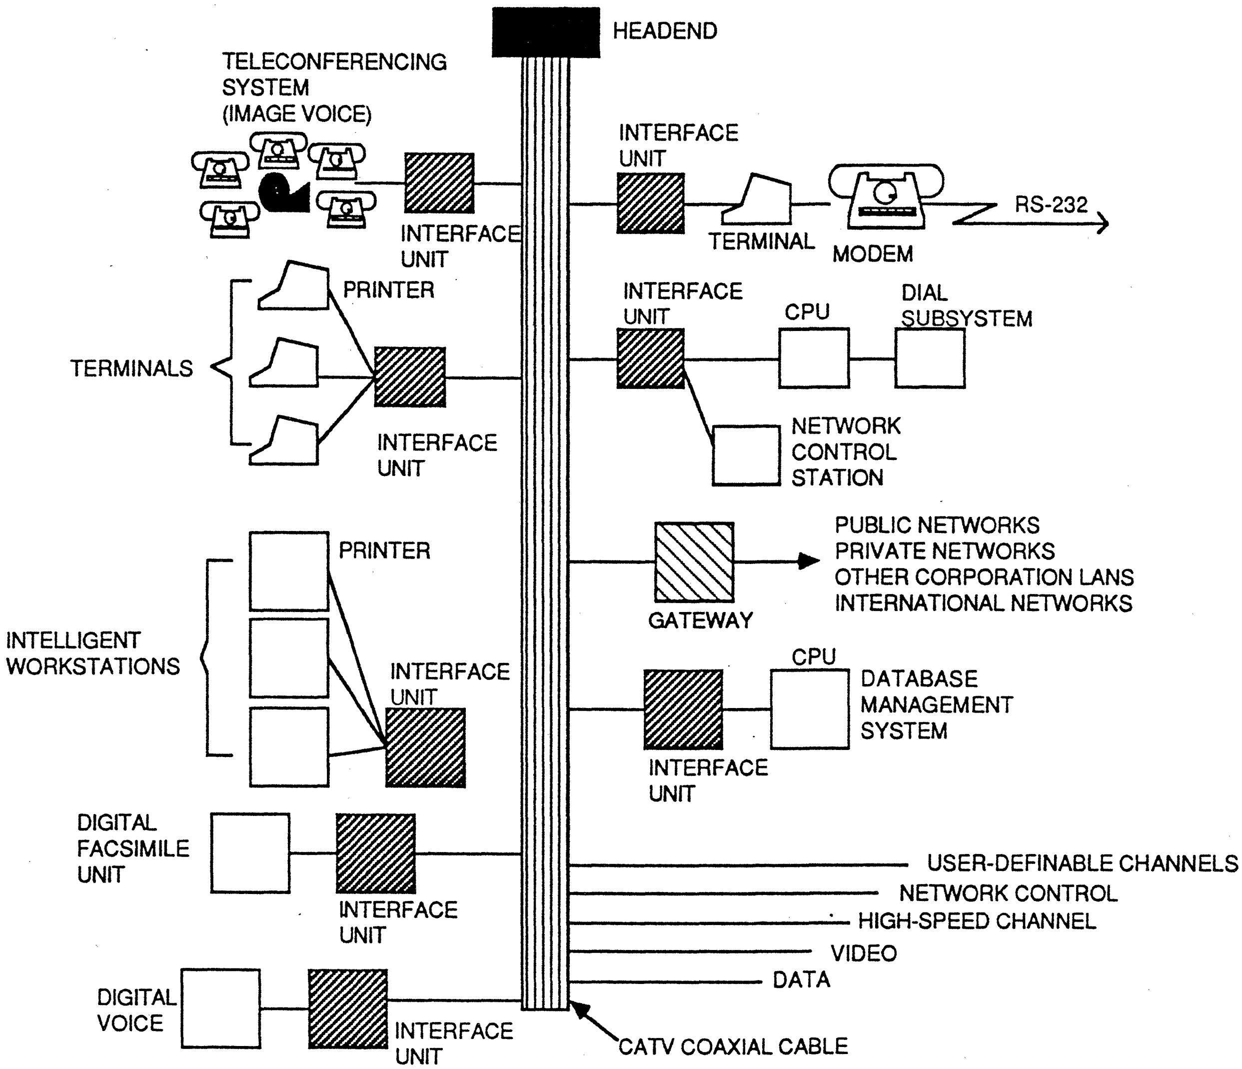

Figure I-2: Generic (Coaxial) Broadband LAN Model

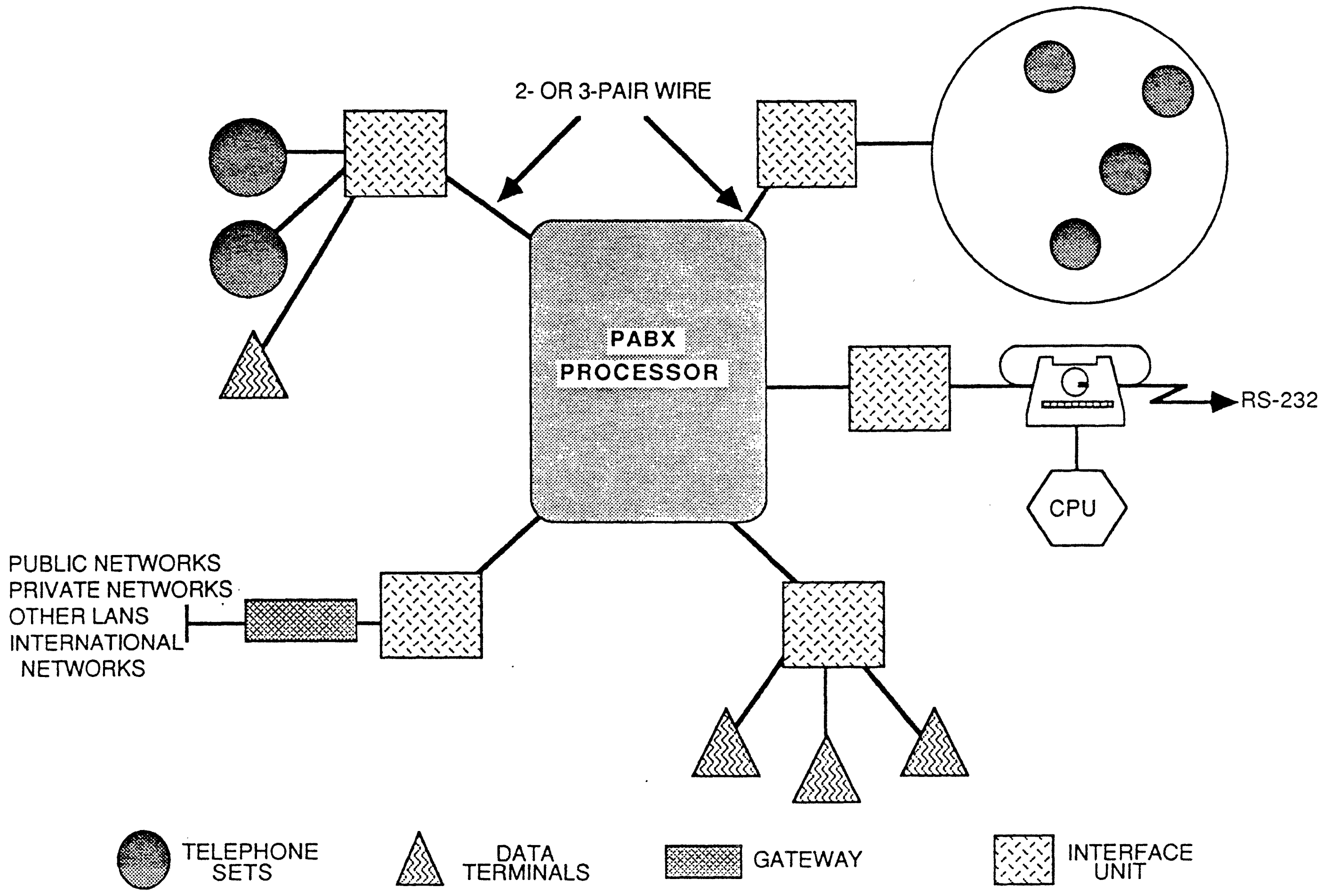

Figure I-3: Generic PBX Model

generation of PBXs include distributed architecture -- i.e., multiple switches in a hierarchical configuration, with distributed systems -- and the use of digital telephones, which allows for integrated voice/data workstations. PBX interface units are almost exclusively RS-232.

With modern PBXs, the internal integration of digitized voice and data are of paramount importance. Also, the capabilities of LANs and PBXs do tend to overlap. PBXs are generally more economical to install than LANs, as they can build upon existing telephone wire, and they are superior to LANs when it comes to handling voice. LANs, however, have a definite edge in terms of higher data rates. The current PBX versus LAN debate is still an ongoing issue, and the issue of differences and advantages may become even more clouded as a new "fourth-generation" LAN -- best embodied in Ztel's PNX and CXC's Rose PBXs -- attempt to use a token-ring LAN in their architecture to truly integrate voice and data in a PBX application.

1.1.2.4. Ring-Based LANs

Ring-based LANs are relatively new to the United States, though ring networks are the prevalent type of LAN used in Europe. The token-ring approach to local area networking has, however, begun to receive an enormous espousal in the U.S. ever since IBM announced its plan to adopt the token-ring as its standard for all local area networking applications.

The basic ring configuration consists of a number of repeaters, each linked to two others via unidirectional transmission links to form a single closed path (see Figure I-4). Bit by bit, sequentially, data is transferred around the ring, from one repeater to the next, as each repeater regenerates and retransmits each bit.

In order for a ring to function as a communications network, messages (data) must be inserted, received, and transmitted, and removed. These tasks fall to the repeaters, which serve as device attachment points for message insertion. In a ring network, messages are conveyed in packets, each of which contains a destination address field; as a packet travels past a repeater, the packet's address field is copied to the attached station -- if the station recognizes that address, the remainder of the packet is copied.

In terms of message removal, within a ring LAN data will circulate endlessly unless removed, because the ring is, after all, a closed loop. A packet may be removed by the addressed repeater, or, as an alternative, it could be removed by the transmitting repeater after it has made one complete trip around the loop. This latter approach permits automatic acknowledgement and multicast addressing -- i.e., one packet sent simultaneously to multiple stations.

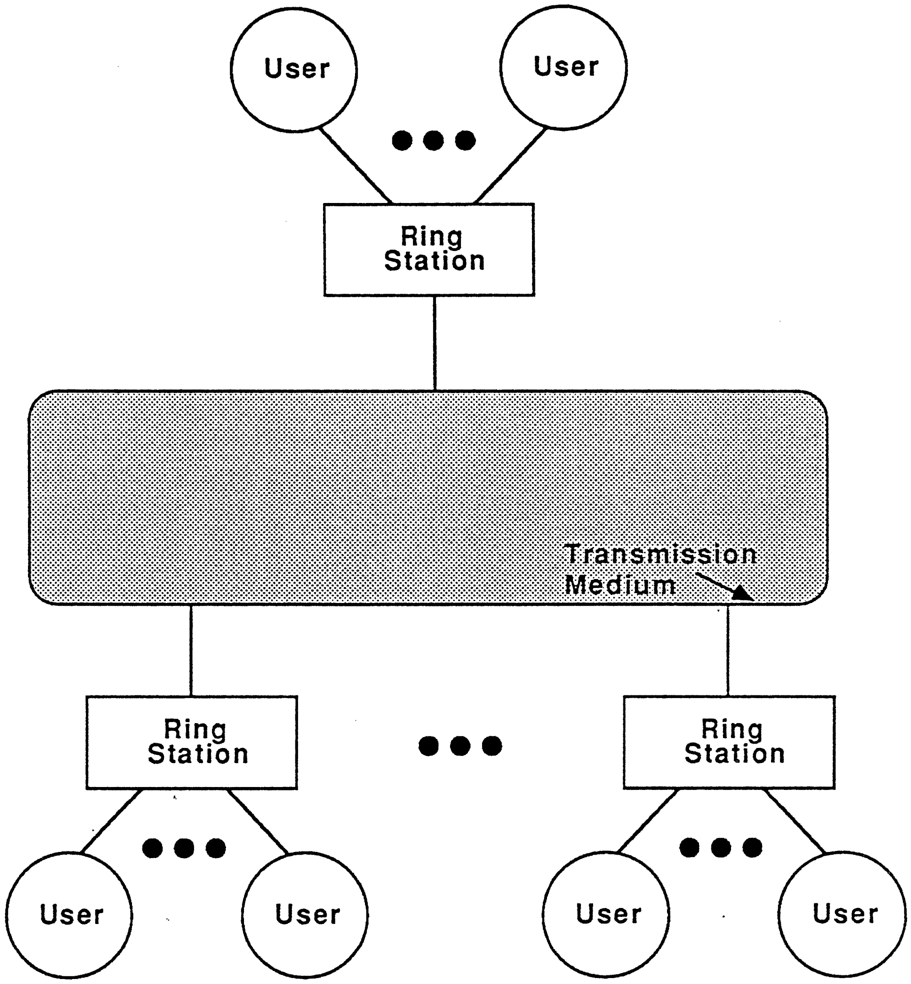

Figure I-4: Basic Ring Structure

Repeaters on a ring network exhibit two states when a ring is operating: the listen state and the transmit state. In the listen state, each bit received by a repeater is retransmitted with but a small delay; this delay is neceessary for the repeater to perform its necessary functions. (These functions include the scanning of the passing bit stream for pertinent patterns -- the addresses of attached devices, permission to transmit, e.g. -- copying of each incoming bit and transmission of it to the attached station, and modification of a bit as it passes by (e.g., acknowledgement that a packet has been copied).) When a repeater's attached station has data to send, and the repeater has permission to send, the repeater enters the transmit state, wherein it receives bits from the station and retransmits them on its outgoing link.

A third state, the bypass state, activated via a bypass relay, can be used to eliminate repeater delay for those stations that are not active on the network. In this state, signals propagate past the repeater with no delay other than medium propagation.

The benefits of a ring-based LAN are drawn from its use of point-to-point communication links. In that a transmitted signal is regenerated at each node, transmission errors are lessened and greater distances can be reached; this gives it an edge in distance over baseband and greater data integrity at high data rates than broadband. Also, the ring can accommodate optical fiber links, which yield very high data rates and desirable EMI (electro-magnetic interference) characteristics. Fault isolation and recovery on a ring is easier than on a bus tree (see 1.1.3.3 below), and duplicate address problems are avoided by a ring station's ability to modify a bit in a packet to acknowledge to other stations that it has already been received. A ring-based LAN also boasts greater throughput than a comparable tree or bus LAN. And again, the impact of IBM's endorsing the token-ring LAN as its standard LAN cannot be overlooked.

On the debit side, cable failure in a ring network brings down the whole network, as does failure of a repeater. Also, in both cases, the whole ring must be inspected to locate the failed link or repeater -- often a time-consuming task.

1.1.2.4.1. Token Rings

There are, basically, two kinds of token rings: those with fixed-length tokens and those with variable-length tokens. (Basically, a token is nothing more than a pattern that has two states, one busy and one free.) Fixed-length tokens are, obviously, tokens of some fixed, predetermined length; when manufacturers say their tokens are variable-length, they mean they are of variable length to some maximum size. The advantage of the variable-length token is that it is basically insensitive to ring size: one does not have to have enough capacity in the ring to have a fixed number of bits on the ring.

Thus, IBM proposes that, in a token ring, the token be of a small fixed size, that is, a free token that is very small: when it is sized and marked to designate a frame, one then has available a frame of variable size, that can be filled to a fixed maximum size.

Also, with variable length set to a maximum size and a fixed-ring-wire-length kind of mechanism, it is possible to begin performance calculation: if one knows the number of devices, and if one has just completed a transmission, the worst-case time before reentering the ring can be determined. The ring performance becomes deterministic in the sense that it is bounded.

1.1.2.4.2. Slotted Rings

A slotted ring can be viewed as no more than a giant shift register with a series of fixed bit positions that allows for the division of a fixed number, fixed-length, identical slots (refer to Figure I-5). There are used and unused slots on the ring which continually circulate, much like the cylinders on a revolver; they constantly and synchronously move ahead. A device can transmit in any free slot as it moves by.

Bandwidth on the ring, then, can be divided up into "slots." Each bit area can be called a slot, and each of these slots is then a token, and they are packed on after another; i.e., they are stacked up around the ring. In handling rings of varying sizes, the slot probably has to be small; if it is large, one will be unable to get enough slots on the ring. Given a minimum size of the ring, it must be possible to fit at least one slot on it. The length of variable-length tokens can be 1,000, 2,000, or 4,000 bytes; the slotted ring has slots in the size of 4 and 5 bytes.

There is a significant difference between 5 bytes and 1,000 bytes, which gives rise to a philosophical issue: in a system, does one prefer infrequent opportunity at long transmissions, or frequent opportunity at short transmissions? With a great amount of overhead bits, there are very frequent opportunities to transmit, and the idea of a slotted ring is that one adjusts the ring to be of variable length, to fit a certain specific number of these slots.

Another question that naturally comes to mind is this: is there an environment in which a slotted ring or a token ring would be preferable over the other? As an example, in real-time voice, there would be quantities on the size of 16 bits; if a series of these quantities has to be transmitted at a very rapid rate, it would probably be preferable to have a ring with many little slots around it to provide frequent opportunities to

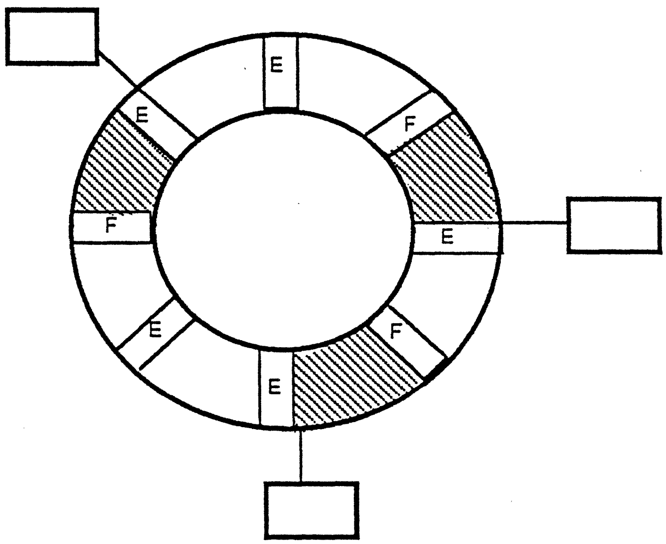

Figure I-5: Generic Slotted Ring Configuration transmit this real-time voice as it comes out, versus a ring with giant variable-length tokens resulting in a long wait time before another 16 bits are transmitted.

1.1.3. LAN Topologies

The concept of topology refers to the various methods whereby computers, terminals, peripherals, etc., are to be physically interconnected; stated another way, it is the pattern of interconnection used among the various nodes of the network, i.e., the overall shape of the LAN. The predominant LAN topologies are the star, the loop (ring), the bus, and the tree, all of which are described in the following subsections. The choice of topology depends on a variety of factors (e.g., performance, reliability, expandability) and the tradeoffs a LAN designer is willing to make among the various approaches.

1.1.3.1. Star Topology

In the star topology, each station on the perimeter is connected to a common central switch via a point-to-point link, resulting in a configuration that resembles a star (see Figure I-6). The circuit-switching technology employed in a star network resides at the center of the star, through which all communication must pass. For a station on the perimeter to transmit data, it must first request the central switch to connect it to some destination station; once the requested circuit is opened, data may be exchanged between the two stations.

The potential problem with the star network is that if the central point should fail, all communications fail: this would be the classic central point of failure in a network. What is good about it is that while the central switch node is complex, the communications processing burden on the stations is minimal, other than for basic connection-requesting and accepting considerations.

1.1.3.2. Ring Topology

In the ring topology, as alluded to earlier, the LAN consists of a set of repeaters connected by point-to-point links in a closed loop (ring) (refer to Figure I-7). The repeater is capable of receiving data on one link and transmitting it bit-by-bit on the other link, as fast as it is received. Data is transmitted in one direction only, circulating around the ring in either a clockwise or counterclockwise direction.

Each station attaches to the network at a repeater, and data is then transmitted in packets. Thus, if station A wants to transmit a message to station B, it breaks the message up into

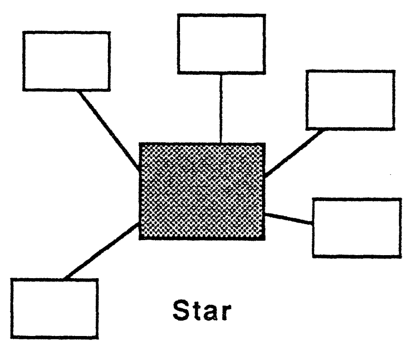

Figure I-6: Star Topology

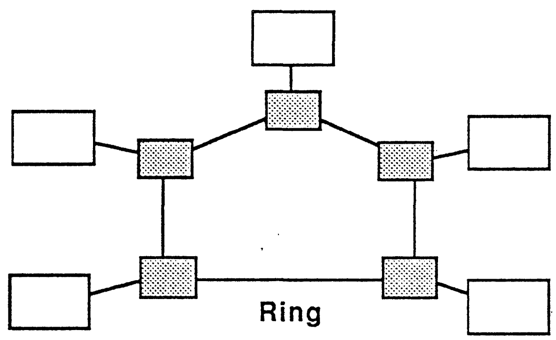

Figure I-7: Ring Topology

packets: each packet contains a portion of the data plus some control information -- including B's address. The packets are inserted into the ring one-by-one and circulate through the other repeaters; station B recognizes its address and copies the packets as they go by. Since multiple devices share the ring, control is needed to determine at what time each station may insert packets, and this is almost universally done with some form of distributed control. Each station contains access logic that controls transmission and reception.

1.1.3.3. Bus/Tree Topology

The bus topology is the topological approach to local networking most commonly used in the United States (see Figure I-8). In this topology the communications network is the transmission medium -- there are no repeaters and no switches. All stations, in such a scheme, attach directly to a linear transmission medium -- bus -- via hardware interfacing. The bus topology features multipoint communications: i.e., a transmission from any station propagates the full length of the medium and can be received by all other stations along the bus. In that multiple devices share a solitary data path, only one device may transmit at any one time.

The tree topology is an extension of the bus, in which several buses are linked by passive splitters or active repeaters (see Figure I-9). The transmission medium is a branching cable with no closed loops, and, again, a transmission from any station traverses the length of the medium and can be received by all other stations. Messages transmitted on both buses and trees are in the familiar packet format, and the medium is referred to as multipoint or broadcast.

1.1.4. Network Control Strategies

Of the toplogies discussed above, the bus and ring are the two that require a control mechanism for proper network operation. Whereas a star network's central node either handles messages for every node simultaneously or polls each node in turn to see if it wishes to transmit, the ring and bus lack a central node and must use some distributed mechanism to determine which node may transmit at any given time.

1.1.4.1. Node-To-Node Permission

In ring networks, there are three basic control strategies predicated on the premise that permission to use the network is passed sequentially around the ring from node to node. These are message slots, control token, and daisy chain.

In the message slot approach -- such a slot being a sequence

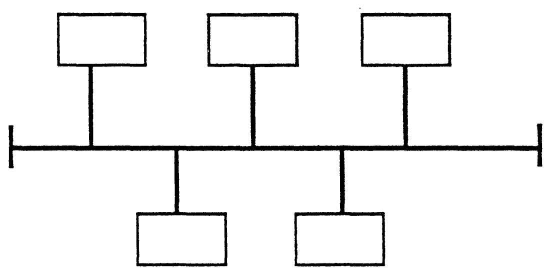

Figure I-8: Bus Topology

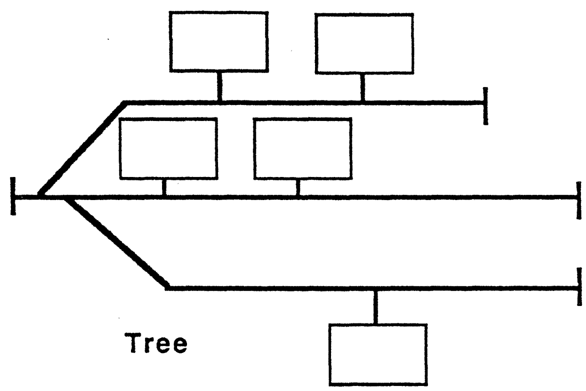

Figure I-9: Tree Topology

of bits sufficient to hold a complete message -- a series of slots is continually circulated around the network. A slot may be either empty or full, and any node on the network, upon noticing an empty slot passing by, may mark that slot as full and insert a message in it. This approach is not completely decentralized, as one node must initially generate the slot program.

In the control token strategy, a control token -- a special bit pattern -- is passed sequentially around a ring network. Any node, upon receipt of the control token, may remove the token from the ring, send a message, and then pass the control token onto the next node.

In a daisy-chain network, dedicated wires are used to pass the control information along from one node to the next.

1.1.4.2. Contention Control

A simple control strategy often used on bus networks is that of contention. In a contention network, any node wishing to transmit a message does so; because there is no control mechanism at work in this type of network, collisions of messages invariably occur, as nothing prevents two nodes from transmitting simultaneously.

With a contention control strategy, a node is given the ability to detect a collision, whereupon it waits a random amount of time and then retransmits its message. If network traffic consumes but a small percentage of the available bandwidth, the number of collisions and retransmissions is reasonably small. The timer needed to generate a random distribution and the mechanism needed to detect collisions are very simply mechanisms.

1.1.4.3. Register Insertion

The register insertion control strategy is one especially well-suited to the ring topology. In this approach, a transmission message is first loaded into a shift register. The network loop is then broken, in order to insert this shift register into the network, either between two adjacent messages or when the network is at rest. The message is then shifted out onto the network, with any message arriving during this period being shifted into the register behind the message being sent. In that the shift register has thus become an active component of the network, no additional messages can be transmitted by this node until the register is switched back out of the ring.

The register insertion technique is said to ensure a fair, balanced allocation of network bandwidth, as a node that has transmitted one message cannot transmit a second message until the first message has passed completely around the ring and the node has removed its register from the network. Every other node has a chance to send one message before any given node may transmit a second. Register insertion is treated in greater detail in section 2.1.3.

1.1.5. LAN Standards - The OSI Reference Model

The 1970s witnessed both a proliferation of heterogeneous computing devices and a spiraling upward growth in the complexity of communications technology. In order to foster compatibility among machines built by different manufacturers, and in order to simplify the mechanics of creating open systems interconnections, a movement began within computing circles to specify and implement standards. If standards were defined that computer vendors could adopt, and purchasers could demand adherence to these standards in the products they buy -- so the reasoning went -- some semblance of order could be imposed on the chaos threatening to choke-off the evolving technology.

In accordance with this line of reasoning, the Computers and Information Processing Committee of the International Standards Organization authorized a new subcommittee to develop a model of the communications process within which new standards could be developed. Through this model, it was hoped, a definition of the scope of future standards would evolve.

The subcommittee, known as ISO/TC97/SC16, produced a report (ISO/TC97/SC16/N227) in June of 1979 which laid down the Reference Model for Open Systems Interconnection (OSI) -- usually known as "the seven-layer model". The model was subsequently published in a draft proposal in 1980. The OSI model provides both a framework through which existing systems can be analyzed, and a guideline for the definition of new systems and standards.

The seven layers of the OSI model (refer to Figure I-10) are all involved in any actual communication and are arranged in a strict hierarchy; i.e., each layer provides services to the layer above it and uses services from the layer below it. The interfaces between the layers were not intended to be standardized -- standardization was intended to control the protocols that exist between corresponding layers in communicating devices. ("Protocol" is a word borrowed from the terminology of political diplomacy to describe the rules governing communication between different pieces of computing equipment; conventions are necessary for data exchange to proceed in a predetermined, orderly fashion.)

ISO intends to develop an open systems architecture that consists of one or more protocols at each level; once the full set of protocols is defined, any two devices that conform to the full set of protocols will be able to communicate with one another despite differences in internal interfaces, construction, function, etc. This openness should put an end to the present

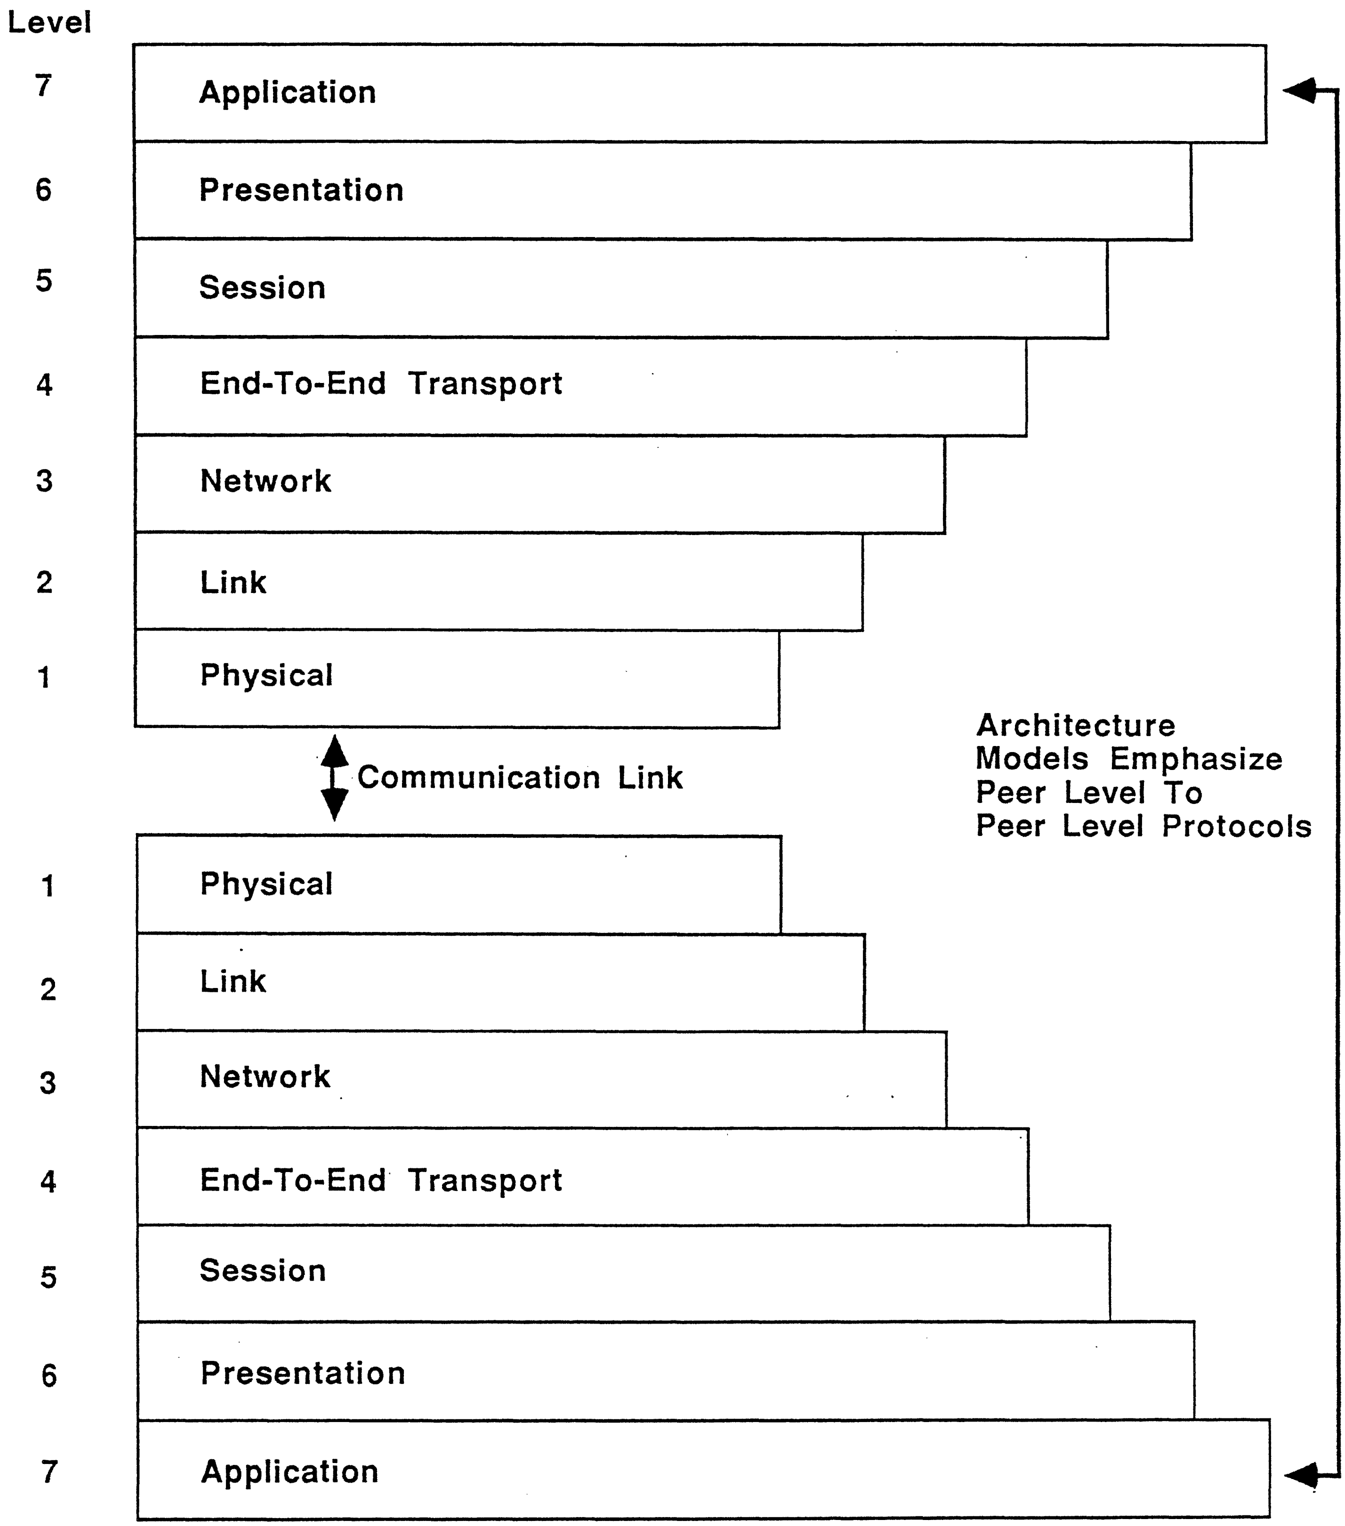

Figure I-10: ISO Reference Model

pitfall of users sometimes becoming locked in to single vendors.

ISO chose a widely accepted structuring technique -- layering. All communications functions are partitioned into a hierarchical set of layers. Each layer performs a related subset of the functions required to communicate with another system at an analogous layer. It relies on the next lower layer to perform more primitive functions and to conceal the details of those functions; similarly, it provides services to the next higher layer. Ideally, the layer should be defined so that changes in one layer do not require changes in other layers. In this way, a single monolithic problem can be broken down into a number of more manageable subproblems.

The task of the OSI subcommittee was to define a set of layers and the services to be performed by each layer. By breaking the communications process into tasks, the OSI subcommittee hoped to group functions logically and provide enough layers so that each layer was a manageable size, yet all the layers together did not require an unrealistic amount of processing overhead. The resulting OSI reference model has seven layers, which are listed in ascending order and described below (refer also to Figure I-11).

Physical Layer - The physical protocol layer, which is the lowest layer of the model, has to do with electrical and mechanical specifications, e.g., signal levels and connectors. The function performed by the physical layer is the transmission and reception of an unstructured stream of bits over the network communication medium.

Link Layer - The function performed by the link protocol layer (also called the data link layer) is the transmission and reception of a structured stream of bits over the network medium. Whereas the physical layer provides for the transmission of a raw bit stream, the data link layer provides the ability to subdivide this bit stream into structured blocks of information commonly called frames or packets. Frames are typically delimited by reserved bit patterns or reserved character sequences, or through the inclusion of an explicit bit or byte count within the frame.

Network Layer - The function of the network layer (sometimes called the communication subnet layer) is to facilitate the transmission and reception of a packet from a source node to a destination node within the network. In traditional point-to-point, store-and-forward networks, packets are typically passed from source-to-destination through a series of intermediate nodes; hence, some strategy for routing a packet to its eventual destination must exist. This strategy is implemented within the network layer.

Transport Layer - The function of the transport layer (also called the host-to-host layer) is to provide for reliable end-to-end or host-to-host communication over the network. Two

Layers	Function	Layers
User Program	Application programs (not part of the OSI model)	Server Machine
Layer 7 Application	Provides all services directly comprehensible to application programs	Layer 7 Application
Layer 6 Presentation	Restructures data to/from standardized format used within the network	Layer 6 Presentation
Layer 5 Session	Name/address translation, access security, and synchronize & manage data	Layer 5 Session
Layer 4 Transport	Provides transparent, reliable data transfer from end node to end node	Layer 4 Transport
Layer 3 Network	Performs message routing for data transfer between non-adjacent nodes	Layer 3 Network
Layer 2 Data Link	Improves error rate for messages moved between adjacent nodes	Layer 2 Data Link
Layer 1 Physical	Encodes and physically transfers messages between adjacent nodes	Layer 1 Physical

Physical Link

Figure I-11: OSI Reference Model, Layers, Functions differing strategies have evolved to support this reliable communication -- datagrams and virtual circuits. Datagrams are self-contained messages which include explicit source and destination addressing information, and which are delivered reliably under the control of transport layer protocols. Virtual circuits allow logical (or physical) connections to be established between source and destination nodes; they also provide for a reliable data stream to be passed over the "circuit" from end to end.

Session Layer - The function of the session layer is to manage end-to-end communication between processes running on network hosts. Typically, the session layer provides facilities to map logical processes or port names (character strings) into network address information that is meaningful to the transport layer. The session layer manages interprocess communication within the network by opening, closing, and sending data over transport layer virtual circuits, or by sending transport layer datagrams, or both, depending upon the facilities which are available at the transport layer.

Presentation Layer - The function of the presentation layer is to transform the data to be sent to, or received from, the session layer. For example, one might wish to use a text compression algorithm in order to minimize the amount of data to be sent through the network. Text would be compressed within the presentation layer of the transmitting host and expanded again in the presentation layer of the receiving host. Similarly, data encryption algorithms for enhanced security might be implemented at the presentation layer.

Application Layer - The function of the application layer is to provide a variety of application-specific protocols to application layer programs that might include services to facilitate funds transfer, information retrieval, electronic mail, text editing, and remote job entry. The variety of services provided at the application layer is theoretically limited only by the variety of possible application programs to be run in the network environment.

1.1.6. Transmission Media

There are many tradeoffs to be considered among the various types of transmission media available for use with local area networks. Each of the choices is designed to do away with the "rat's nest" problem of wiring and replace it with a single, shared network medium that is inexpensive, reliable, simple, durable, noise-free, and high-speed. Such media should also be relatively simple to install, service, and reconfigure.

1.1.6.1. Twisted-Pair Wiring

Twisted-pair wiring is one of the more popular media used in local area networks. It has been used by the phone companies for years, and is, therefore, well understood. With it, it is a simple matter to add nodes between existing nodes, and further, it is easy to tap onto that twisted-pair.

While it is easy to install, it is, unfortunately, also very susceptible to tapping, so there is a security problem involved, and this is one of its disadvantages. It is also relatively susceptible to noise, which accounts for its very short distances: most such systems are 1,000 to 2,000 feet in length, because they are running at very high data rates. If these local networks are run at 300 baud to 1,200 baud, noise is no problem, but, when one has high data rates, noise does become a factor.

1.1.6.2. Coaxial Cable

A coaxial cable, on the other hand, does support higher bit rates than twisted pair. Most of the twisted-pair systems are limited to 1 Mbps, while coaxial cable systems tend to be from 1 to 10 to 50 Mbps.

With a coax cable medium, users can extend the distances, because it is a better shared-medium; in that it is less susceptible to noise, distances can be extended, typically to 2,000 to 3,000 to 4,000 feet. And it is relatively inexpensive; like twisted-pair, it varies from 10 cents a foot all the way up to $3 a foot.

Like twisted-pair, coaxial cable is susceptible to tapping, and, because of its size and weight, it may be difficult to install in some environments, because often, for example, building codes require installers to pull the cable through conduit -- and Ethernet cable is not very flexible but instead thick and difficult to manipulate.

Also, coax requires high driving power, so that the transceivers on an interface have to be higher-powered.

1.1.6.3. Fiber-Optic Cable

Fiber-optic cable has a potentially unlimited data rate that could exceed 50 Mbps; unfortunately, there is at the present no interface or integrated circuit technology that can drive that high a data rate, and the limitation here -- unlike the coax and twisted pair -- is not in the medium but in the interface itself.

Fiber-optic cable is highly immune to noise; there is no electrical signal because it works on light waves; and it is very small and flexible and hence easy to work with, as it is very easy to pull through conduit, etc.

On the minus side, fiber-optic usages seems handicapped by a lack of standard components, and, also, transmission using fiber-optic technology is inherently unidirectional -- which eliminates the bus topology, as things now stand.

(One other potention candidate for LAN data transmission is CATV (cable television) technology. CATV offers a voluminous bandwidth, and its wide-spread usage has driven down the costs of system components. Also, in many instances CATV is already installed, and a LAN could use some of the channels of existing equipment.)

2. Ring-Based LANs

A ring LAN, as has already been seen, consists of a number of repeaters each connected to two others via unidirectional transmission links to form a single closed path. Data is transmitted around the ring sequentially, bit by bit, from one repeater to the next. Each repeater regenerates and retransmits each of these bits.

In a ring operating as a communications network, three distinct functions must be performed: message insertion, message reception, and message removal. These tasks are performed by the repeaters. Each repeater on the ring serves as both an active ring element and a device attachment point.

Messages transmitted around the ring are in the form of packets, each of which contains a destination address field. As a packet passes by a repeater, its address field is copied; if the attached station recognizes the address as its own, the remainder of the packet is also copied. If the attached station does not recognize the address field, it passes the packet along to the next station's repeater.

2.1. Major Types

A variety of ring control strategies can be used for determining how and when packets are added to and removed from the ring. These various strategies generally serve to categorize a ring network as belonging to one of three kinds: slotted ring, token ring, or register insertion ring. These three basic types of ring networks are described below.

2.1.1. Slotted-Ring Networks

Slotted-ring networks are those that are "slotted" into a constant number of fixed-length packet slots that circulate continuously about the ring (see Figure I-12). Each of these packet slots contain a leading bit that designates the slot as empty or full. Initially all slots are marked empty. (This method is also known as the empty-slot technique.)

A station with data that is to be transmitted must break the data up into fixed-length frames. The station then waits until an empty slot comes along, whereupon it marks that slot as full and inserts a frame of data as the slot goes by. That station then cannot transmit another frame until this slot returns.

The slot also contains two response bits, which can be set on the fly by the addressee station to indicate that the data frame was accepted or rejected, or that it, the addressed station, was busy. The full slot travels completely around the ring back to its point of origination, to again be marked empty

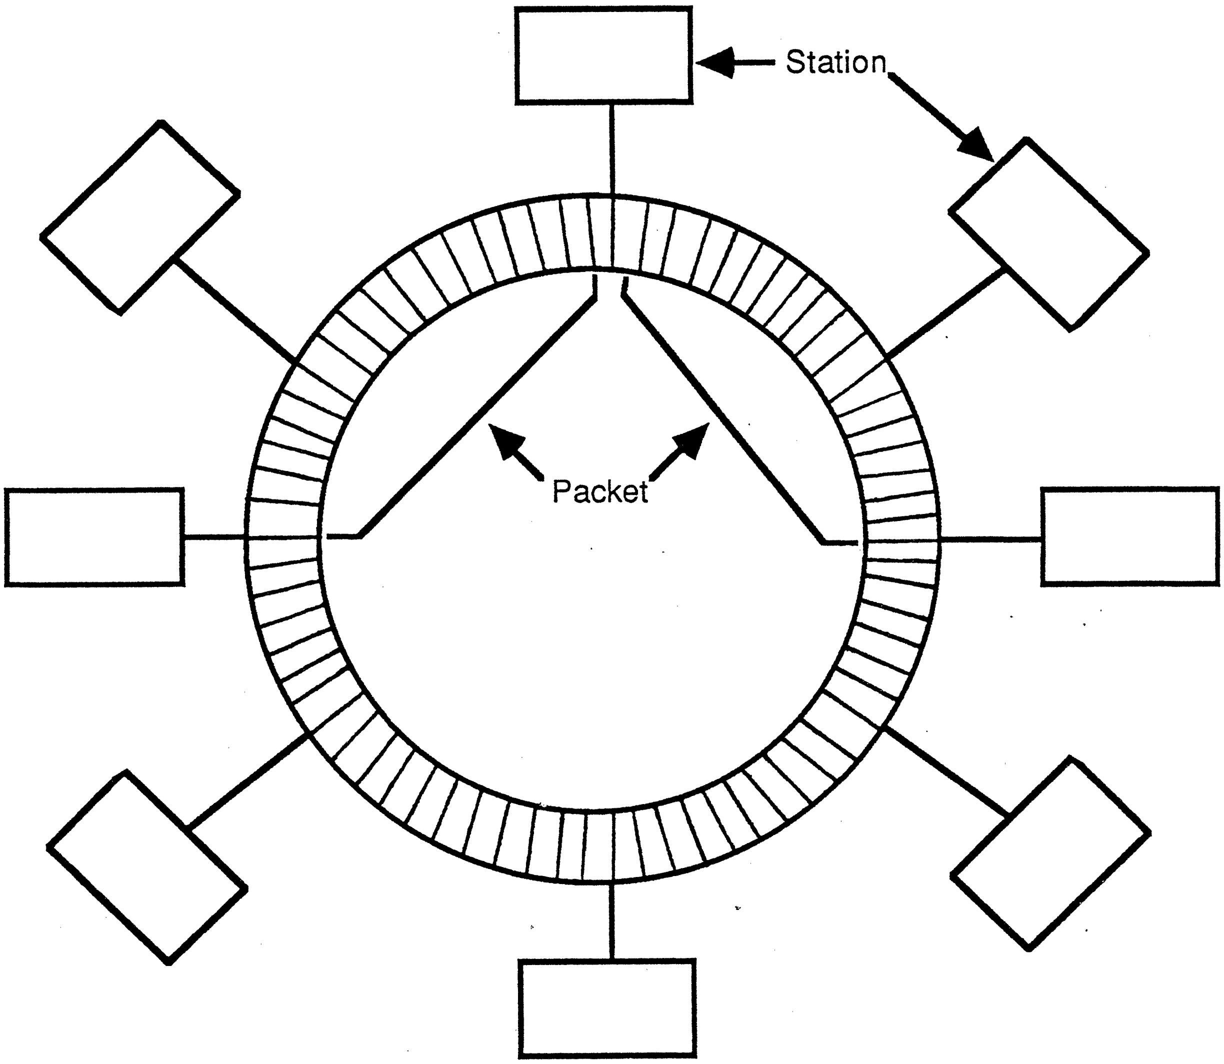

Figure I-12: Slotted Ring (Empty Slot) LAN by its source. Each station knows the number of slots on the ring and consequently can clear the full/empty bit as it goes by. Once the now-empty slot goes by its originating station, that station is free to transmit again.

This slotted-ring strategy was first developed by J.R. Pierce and is sometimes referred to as the Pierce Loop, but most of the development work on it was performed at the University of Cambridge in England. The Cambridge Ring -- as the resultant network came to be called -- is that LAN which likely has given the slotted-ring approach to networking its wide-spread recognition within the industry. Several British firms market commercial versions of the Cambridge Ring. (The Cambridge Ring is discussed in greater detail in Section 3.1 below.)

Industry acceptance of the slotted-ring strategy seems to be based upon the attractiveness of its simplicity: interaction with the ring at each node is minimized, which equates to a high coefficient of reliability.

The primary rap against the slotted-ring network is that it is wasteful of bandwidth. In the Cambridge Ring, each slot consists of room for one source address byte, one destination address byte, two data bytes, and five control bits, for a total length of 37 bits; yet in a slotted ring each frame contains only 16 bits of data out of 37 bits total -- a very sizable amount of overhead! Also, a station may send only one frame per round-trip ring time -- if only one or a few stations have frames to transmit, many slots will circulate empty.

2.1.2. Token-Ring Networks

Token-passing on a ring is, in all probability, the oldest ring control technique, dating back as it does to 1969 and E.E. Newhall, after whom it was dubbed the Newhall Ring, and D.J. Farber's Distributed Computing System at the University of California at Irvine in 1972. While the Cambridge ring/slotted-ring technique is the prevalent ring control strategy in England and Europe, the token-ring technique is the most popular means of ring control in the U.S. This U.S. popularity may be attributable in large part to IBM's long-stated commitment to the token-ring method -- and the fact that it is, after all, the one ring access method selected for standardization by the IEEE 802 Local Network Standards Committee, in its 802.5 token-ring specification.

In a token-ring network, a special bit pattern (or data structure) called the token circulates about the ring whenever all stations are idle. This token can be viewed as an access-granting message, and, basically it consists of a delimiter signal and a free indicator.

When a station wishes to transmit (refer to Figure I-13), it

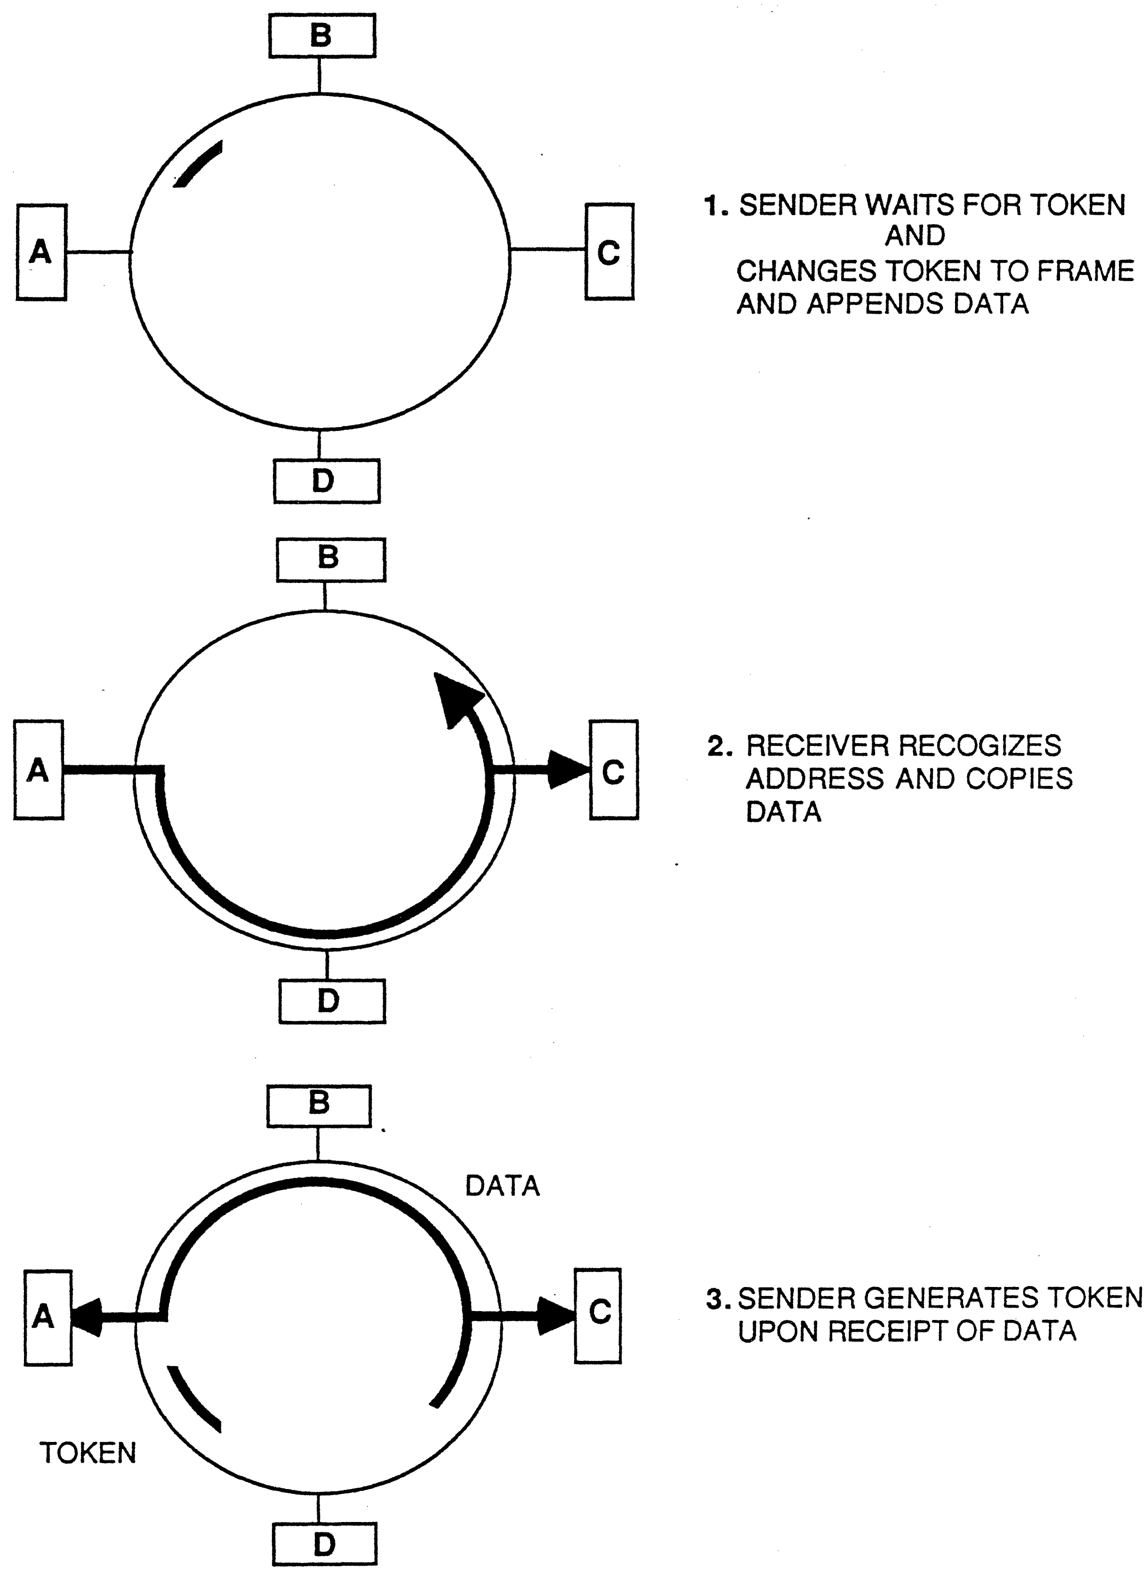

Figure I-13: Token Ring Transmission Process must wait until it detects a token-passing by, at which point it sizes it and removes it from the ring. (This removal of the token is managed by the ring interface, which connects the station to the ring.) The ring interface, which monitors all bits passing by, inverts the last bit of the token it wishes to seize, changing the token into another pattern known as a connector. Immediately after such a token transformation has been made, the station making the transformation is permitted to begin transmitting, and it then transmits a frame immediately following the busy token.

(In that there is now no free token on the ring, other stations wanting to transmit must wait until there is a free token available; thus the use of a token guarantees that only one station at a time may transmit, which is to say that there is never any contention, as in Ethernet.)

The ring interfaces -- the implementers of the basic token-seizure process -- have two operating modes: listen and transmit (see Figure I-14). In the listen mode, the input bits are but copied to output, with a delay of one bit time; in the transmit mode -- which is entered into only after the token has been seized -- the ring interface severs the connection between input and output and enters its own data onto the ring.

As bits that have been sent around the ring return, they are removed from the ring by the sender. At this point the sending station can either save them (to compare them with the original data sent and thus monitor the reliability of the ring) or discard them. After a station has transmitted the last bit of its packet, it must then regenerate the token and remove the packet, and its interface must immediately switch back into the listen mode, to ensure that it does not remove the token or connector that follows. When a transmitting station has released a new free token, the next station downstream with data to send will be able to seize the token and transmit.

Once data transmission is made, the frame -- which can be of arbitrary length -- passes all ring interfaces until it comes around to the ring interface of its intended destination station. The destination station's ring interface recognizes its address as part of the frame and copies the information into its receive buffer and then passes the frame along. The frame ultimately returns to the sending station, where, again, it is saved or erased and a new token is issued and passed along.

Error situations can occur on a token-ring network: principally they are lost token, persistent busy token, and duplicate token. These error conditions can be addressed via two different approaches: decentralized control and centralized control.

The decentralized approach -- suggested by the IEEE 802 Committee -- employs two timers: a relatively long no-token

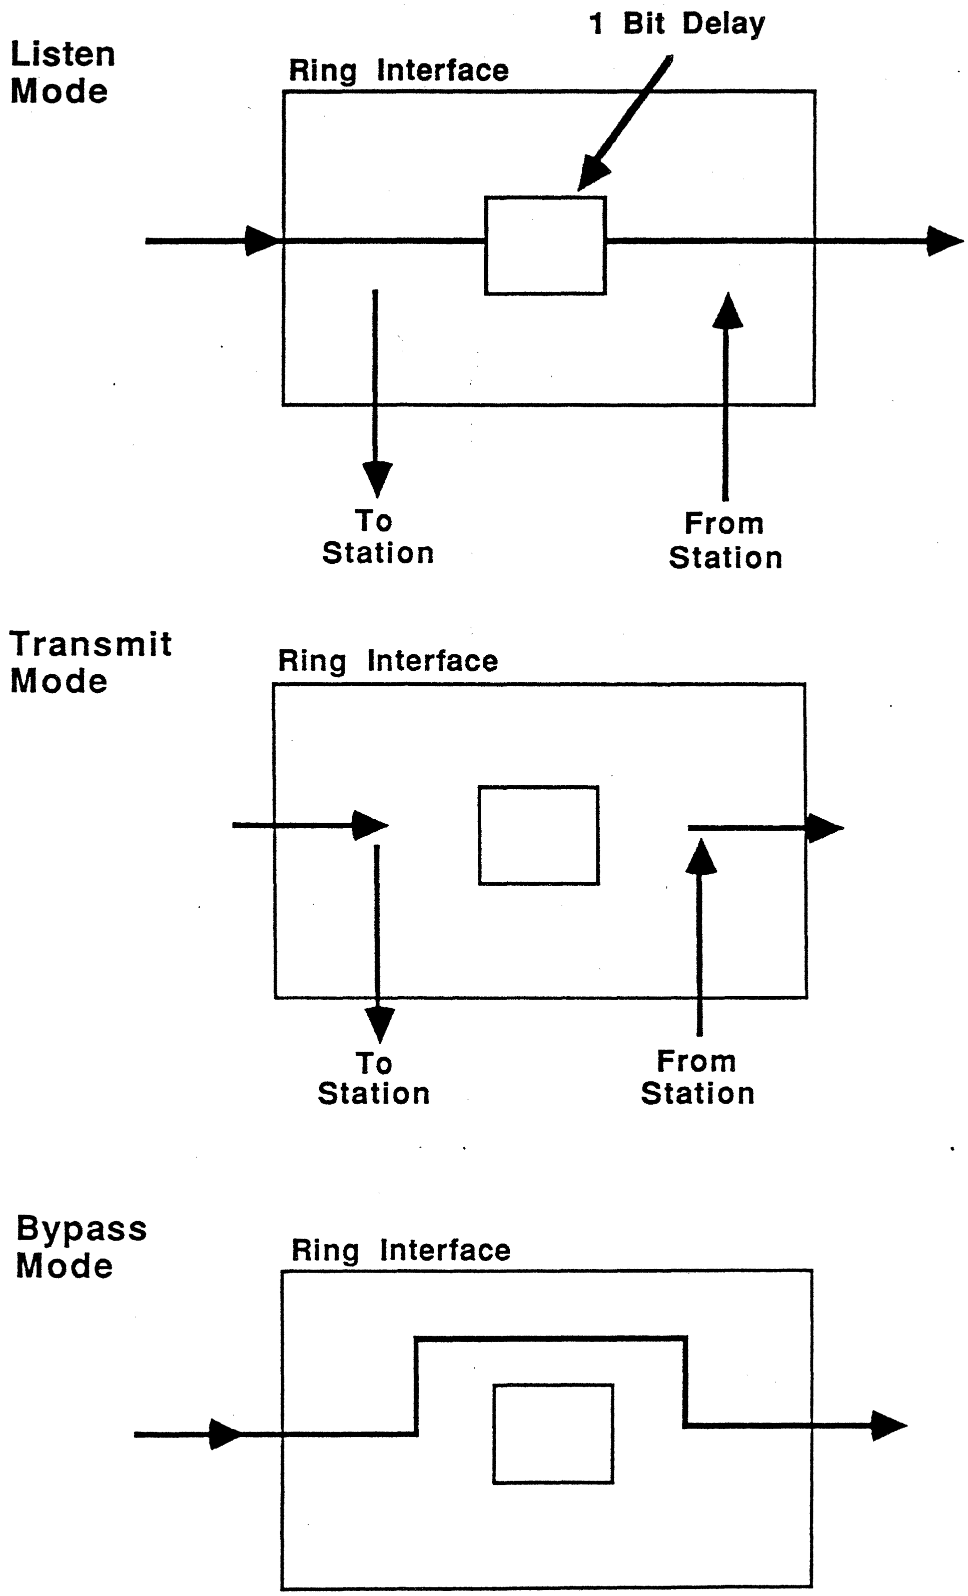

Figure I-14: Ring Interface Listen, Transmit, Bypass Modes timer (TNT) and a shorter valid frame timer (TVX). With the former, when a station has a frame to transmit and has not seen a token for a time TNT, it assumes that the token is lost, purges the ring, and issues a token; with the latter, a station also keeps track of the amount of time since the last valid data frame or token and, when this time exceeds TVX, it assumes that the token is lost and issues a new token.

IBM, however, prefers a centralized approach which uses a station designated as the active token monitor. This monitor detects the lost-token condition by using a time-out greater than the time required for the longest frame to completely travel around the ring -- to recover, the monitor purges the ring of any residual data and issues a free token. To detect a circulating busy token, the monitor sets the monitor bit to 1 on any passing busy token: if it sees a busy token with a bit already set, it knows that the transmitting station failed to purge its frame, and the monitor changes the busy token to a free token.

Other stations on the ring have the role of passive monitor: their primary responsibility is to detect failure of the active monitor and assume that role. A contention-resolution algorithm is used to determine which station takes over.

As alluded to above, the token-ring approach to networking has the benefit of having behind it the driving force of the IEEE working to establish it as a standard, via the IEEE 802.5 token-ring specification. The most salient elements of the IEEE 802.5 specification in relation to the mechanics of the ring are briefly listed below:

1) Single-token protocol: a station that has completed transmission will not issue a new token until the busy token returns. (This is not as efficient, for small frames, as a multiple-token strategy of issuing a free token at the end of a frame; nonetheless, the single-token system does simplify priority and error-recovery functions.)

2) Priority bits: these indicate the priority of a token and thus which stations are allowed to use the token; in a multiple-priority scheme, priorities may be set by station or by message.

3) Monitor bit: may be used if a central ring monitor is employed. (The operation of the monitor is not covered in the standard.)

4) Reservation indicators: these may be used to permit stations with high priority messages to indicate in a frame that the next token be issued at the requested priority.

5) Token-holding timer: begun at the onset of data

transfer, it controls the length of time a station may occupy the medium before transmitting a token.

6) Acknowledgement bits (3): error detected (E), address recognized (A), and frame copied (C); these are reset to 0 by the transmitting station. Any station may set the E bit, while addressed stations may set the A and C bits.

As per the specification and Figure I-15, the individual fields of the frame formats are:

- Starting delimiter (SD): a unique 8-bit pattern used to start each frame
- Access control (AC): has the format 'PPPTMRRR', where PPP and RRR are 3-bit priority and reservation variables, M is the monitor bit, and T indicates whether this is a token or data frame. In the case of a token frame, the only additional field is ED.
- Frame control (FC): indicates whether this is an LLC (Logical Link Control) data frame; if not, bits in this field control operation of the token-ring MAC protocol.
- Destination address (DA)
- Source address (SA)
- LLC: Logical Link Address
- FCS: Frame Check Sequence
- Ending delimiter (ED): contains the error detection (E) bit
- Frame status (FS): contains the address recognized (A) and frame copied (C) bits

The issue of delay in the operation of a token ring network merits consideration. When a ring interface has read all but the last bit of a "potential" token, it must be able to read the last bit and ascertain whether or not it should invert it before forwarding any part of that bit: if the preceding bit pattern is a token, the final bit is to be inverted; if the final bit shows that the bit pattern is not a token but rather just some data, this final bit must not be inverted. Consequently, each arriving bit must be stored in its entirety and a new bit generated -- and this process of evaluation/identification and bit-generation creates a 1-bit delay in each ring interface. The consideration to weigh here is that if there are a good number of stations on the ring, the cumulative effect of all the 1-bit delays has a

Starting Delimiter 1 Byte	Access Control 1 Byte	Frame Control 1 Byte	Destination Address 6 Bytes	Source Address 6 Bytes	Information Field	Frame Check Sequence 4 Bytes	Ending Delimiter 1 Byte	Frame Status 1 Byte

Figure I-15: Frame Format Fields

very significant impact on the ring's performance.

Also, the ring itself must have a sufficient delay to permit a complete token to circulate when all stations on the ring are idle. Basically the delay has two components: the 1-bit delay of each station described above, and the signal propagation delay.

The wiring scheme for a token-ring network (refer to Figures I-16 and I-17) employs a main ring cable that is installed in conduits or ceilings. Wall outlets are connected to distribution panels through extension wires, or "local lobes." Relays -- housed in the distribution panels -- insert these local lobes into the ring when activated by "phantom circuits"; if the relays are not powered, the lobes will not be inserted into the ring.

In attaching a station to the ring, sender and receiver circuits can be packaged with a protocol handler on the same card, which resides in and is powered from the station -- and this makes for low-cost packaging. DC-free coupling of both in-bound and outbound links and the use of the transformers to balance transmission give such a system immunity to ground-loop problems. The electronic on electromechanical relays by which a station is switched into a ring are remotely situated for purposes of reducing cabling length between repeaters and increasing reliability. Cable options available include shielded twisted-pair, coaxial, and optical fibers.

(IBM announced a Cabling System on May 8, 1984, that employs many of the cabling principles and elements described above. This IBM Cabling System was presented as a system that would reduce the cost and complexity of installing and/or moving computer devices within a building. IBM also emphasized that this Cabling System would form the electronic backbone of its "soon-to-be-announced" token ring network, and it stressed the importance of being pre-wired for the Token Ring Network once it arrived. The IBM Cabling System is examined at greater length in subsection 4.1.4.)

The Token-Ring Network is the ring network of choice in the U.S. One of its major advantages over other LANs is that it offers the fairest of all access schemes. Another is that failure or breakdown of one station will not bring down the entire network. Also attractive is the flexibility it gives users to mix media and transmission technologies, by virtue of the point-to-point nature of its transmission scheme, and the easy availability it gives to points for centralized reconfiguration and repair via its use of distribution panels. Furthermore, in that rings do not require transceivers to be located next to the transmission medium, the prewiring of buildings and the installation of wall plugs in every office can be justified on the basis of cost.

Disadvantages associated with a ring network are its

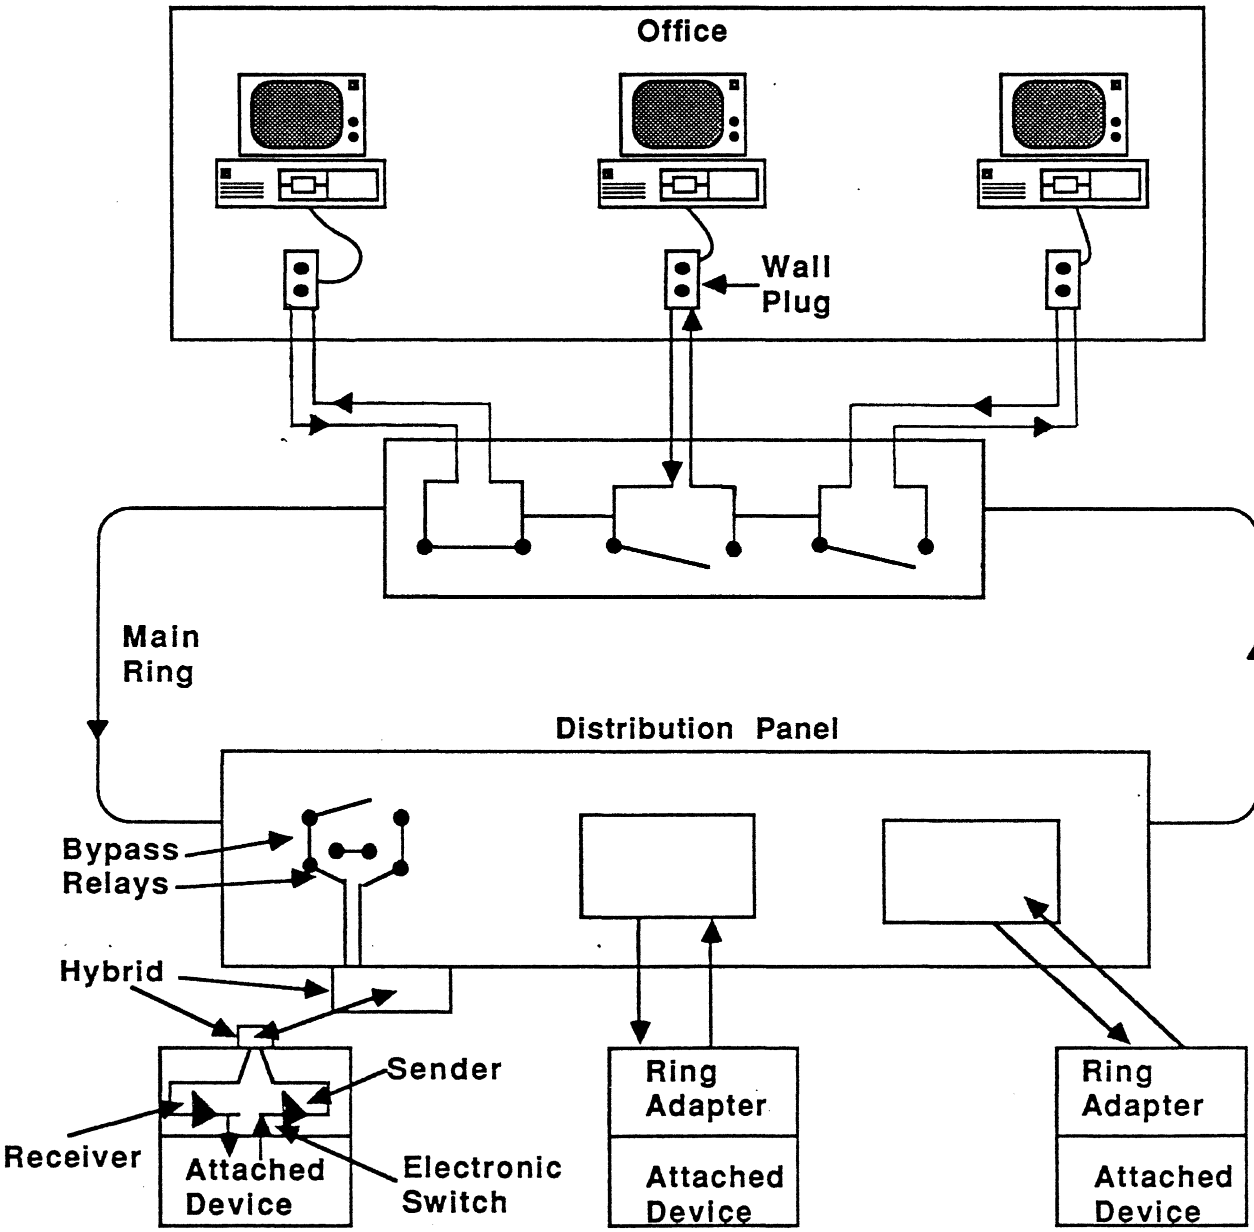

Figure I-16: Ring LAN With Distribution Panels

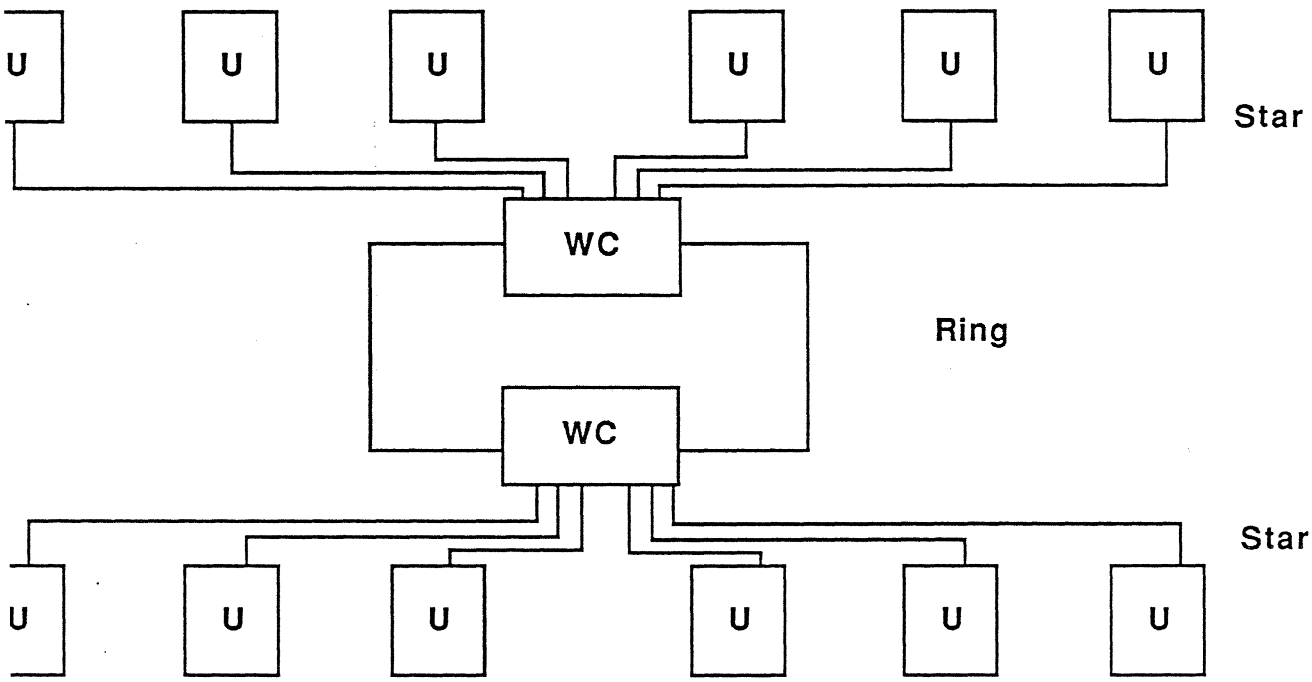

Legend:

WC - wiring center, such as a data wiring panel, a patch panel, or a wiring concentrator
U - user location, such as a data terminal device or a telephone, or a personal computer

Figure I-17: Star-Shaped Ring Wiring Scheme inefficiency under lightly loaded conditions -- because a station wishing to transmit must wait for the token to come around before transmitting -- and its requirement for token maintenance -- loss of the free token prevents further use of the ring, and duplication of the token can wreak havoc on ring operation -- one station must be appointed monitor to ensure that only one token is on the ring and to reinsert a free token if necessary.

2.1.3. Register Insertion Ring Networks

The third major ring control strategy, register insertion, could be called a more sophisticated version of the slotted ring. Originally proposed in 1974 by E. Hafner, Z. Nenedal, and M. Tschanz and developed by researchers at Ohio State University from 1975 - 1978 as the Distributed Loop Computing Network (DLCN), the technique is also used in the IBM Series 1 product.

The register insertion strategy derives its name from the shift register associated with each node on the ring (see Figure I-18). The ring interface of a register insertion ring contains two registers: a shift register and an output buffer. When a station wishes to transmit a packet, it loads the packet into the output buffer. These packets may be variable in length but cannot exceed the size of the output buffer.

When the ring is started up, an input pointer points to the rightmost bit position in the shift register, which means that all the bit slots including and to the left of where it is pointing are empty. When a bit arrives from the ring, it is placed at the position pointed to by the input pointer, and the pointer is moved one bit to the left. Once the address field of the packet has arrived, the interface can ascertain whether or not the packet is addressed to it; if it is, the rest of the packet is diverted to the station, removing the packet from the ring, and the input pointer is then reset to the extreme right.

However, if the packet is not addressed to the local station, the interface begins forwarding it: as each new bit arrives, it is inserted in the place pointed to by the input pointer. The total contents of the shift register is then shifted right one position, pushing the rightmost bit out onto the ring -- the input pointer is not advanced. If no new input arrives, the contents of the shift register can be reduced by one bit and the input pointer moved right one position.

In terms of output, whenever the shift register has pushed out the last bit of a packet, it checks to see if there is an output packet waiting, and if the number of empty slots in the shift register is at least as large as the output packet. Output can proceed only if both conditions are met, in which case the output switch is flipped and the output buffer is then shifted out onto the ring, one bit at a time, in synchronization with the

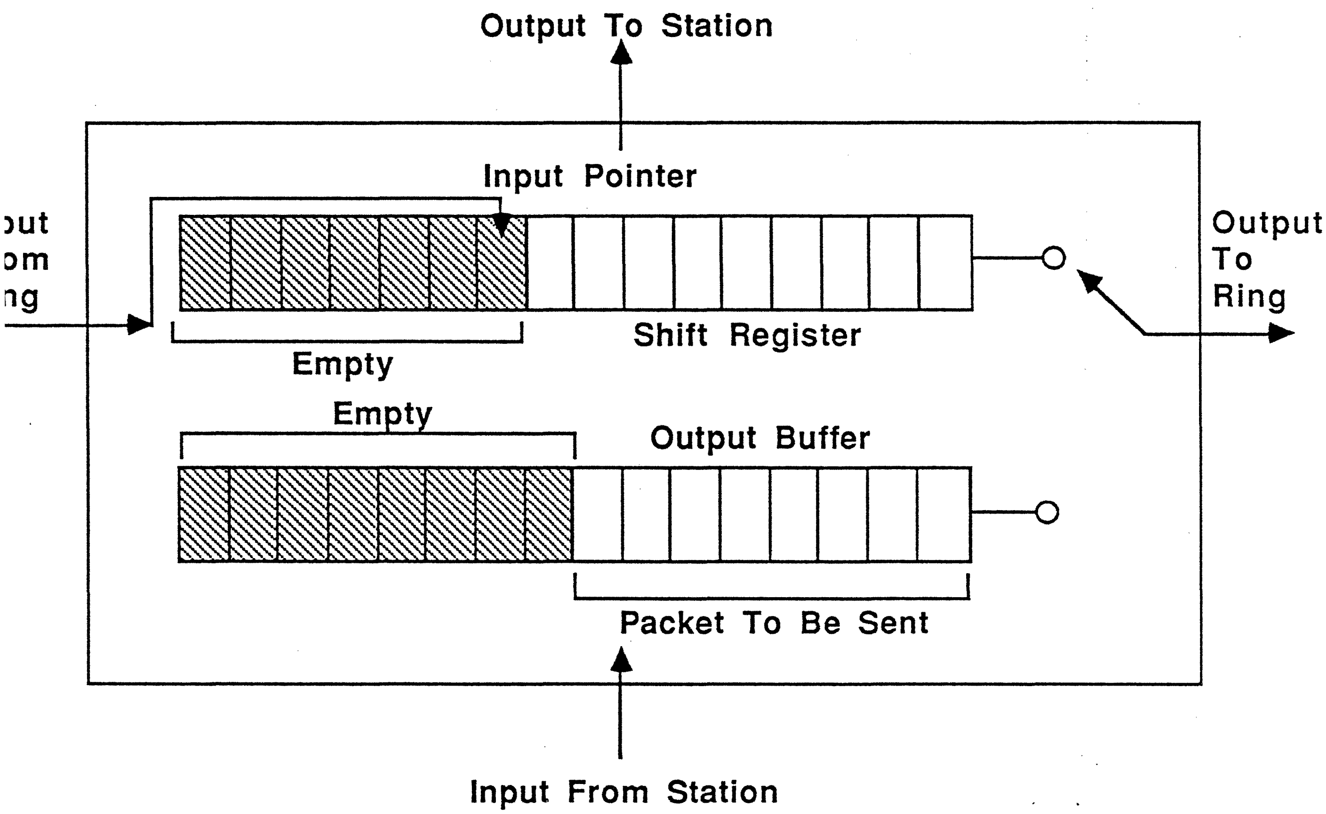

Figure I-18: Interface of a Register Insertion Ring input. New input is accumulated in the shift register -- the very reason why there must be enough empty slots in the shift register to accommodate all the input while the output buffer is being emptied. While not all the slots may be needed, the interface must be prepared for the worst case. As soon as the output buffer has been emptied, the switch is thrown again, and the shift register is emptied (if it has anything in it).

Such a ring architecture prevents a station from monopolizing the ring. If the ring is idle, the shift register will be empty when an output packet has been finished, and thus the next output packet can be sent as soon as it can be loaded by the station in the output buffer. In essence, if the ring is idle, any station can have the entire bandwidth if it so wishes -- which is to say that the protocol permits multiple frames to circulate; however, if the ring is busy, after sending a packet the station will not be allowed to send another one because most often there will not be enough empty slots in the shift register: only when enough interpacket gaps have been saved up can output occur again. Thus the register insertion technique enforces an efficient means of fairness.

Also, once a message is inserted into the ring via the shift register, the incoming traffic that was diverted into another register and so held is re-inserted onto the ring after the new message.

The delay at each station of a register insertion ring network has as a minimum value the length of the address field and as a maximum value the length of the shift register -- a marked contrast to the slotted and token ring networks whose delays at each station are but the repeater delay -- generally one or two bit times.

The primary advantage of the register insertion ring control strategy is that of the three methods treated here it achieves the maximum ring utilization. Also, like the token system, it efficiently allows variable-length frames, and, like the slotted ring, it permits multiple frames to be on the ring, an efficient use of bandwidth.

Its primary disadvantage is its purge mechanism: allowing multiple frames on the ring requires the recognition of an address prior to removal of a frame, regardless of whether it is removed by the sender or the receiver. However, if a frame's address field is damaged, it could circulate indefinitely. Also the requirement for address field recognition mandates that each frame be delayed at each node by the length of that field -- no such requirement exists in the other two methods.

2.2. Token-Ring Media Access Control

The token-ring LAN employs a media access control scheme to ensure orderly, collision-free access to the ring: this control strategy is embodied in the Media Access Control (MAC) frame format.

As has been seen, a single "free" token circulates around the ring when no data is being transmitted. (In the IEEE 802.5 spec, this token is 24 bits in length (refer to Figure I-19) and consists of three octets: starting delimiter (1 byte), access control (1 byte), and ending delimiter (1 byte).) A node (station) wishing to transmit waits for a free token, captures it, and marks it "busy" (thereby converting it into a connector); the node then transmits a frame containing data onto the ring.

The frame, then, is the basic transmission unit on the (IBM) ring. Per the IEEE 802.5 spec, it consists of nine fields (see Figure I-20): starting delimiter (1 byte), access control (1 byte), frame control (1 byte), destination address (6 bytes), source address (6 bytes), Logical Link Control (LLC) -- also called the information field, i.e., the data (greater than zero bytes), frame check sequence (FCS) (4 bytes), ending delimiter (1 byte), and frame status (1 byte). The source address, LLC, and FCS fields comprise the actual "data packet."

Running through these various fields, they can be explained thusly: the starting delimiter (SD) is a unique bit pattern that indicates the start of a frame; the access control (AC) field identifies the whole frame as being available for media access control -- i.e., "the token," or logical link control data (the token is passed using a particular bit pattern in the field); the destination address (DA) and source address (SA) fields are self-explanatory -- bits within them identify the sending node and the intended receiving (addressed) node; the information (or LLC) field is of variable length and is/contains the information to be exchanged; the frame check sequence (FCS) field employs a 32-bit cyclic redundancy check (CRC); the ending delimiter (ED) indicates the end of a frame; and the frame status (FS) field contains the address recognized (A) and the frame copied (C) bits.

When the destination node (station) copies data addressed to it and acknowledges receipt of the frame, it sets (or turns on) the address-recognized and frame-copied bits in the frame status (FS) field. The originating node then removes the "receipt-acknowledged" frame from the ring when it circulates back to it, and it releases a free token -- as has already been seen.

The concept of access to the ring also brings up the issue of priority usage: obviously designers of a token-ring LAN will want to address the assignment of various levels of priority to data to be transmitted over the ring.

The IEEE 802.5 spec provides three bits in the AC field for specifying up to eight levels of priority. This priority mode

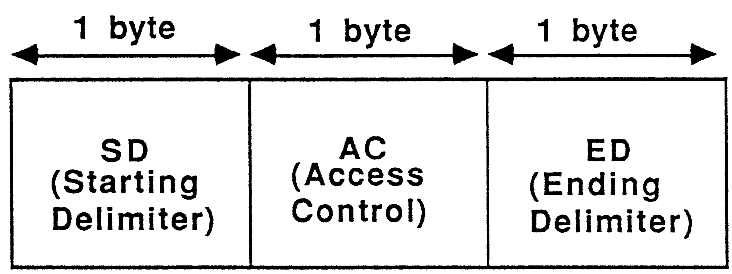

Figure I-19: Free Token Format

←SFS→		←F C S Coverage				→	←EFS→	
SD	AC	FC	DA	SA	INFO (LLC)	FCS	ED	FS

Where

SD = Starting Delimiter (1 octet)
AC = Access Control (1 octet)
FC = Frame Control (1 octet)
DA = Destination Address (2 or 6 octets)
SA = Source Address (2 or 6 octets)
INFO = Information (0 or more octets) (Also called LLC -- Logical Link Control)
FCS = Frame Check Sequence (4 octets)
ED = Ending Delimiter (1 octet)
FS = Frame Status (1 octet)
SFS = Start of Frame Sequence
EFS = End of Frame Sequence

Figure I-20: IBM Token-Ring Frame Format indicates the priority level of a free token and which stations are allowed to use it. In a multiple-priority system, different stations can be authorized to use different priorities; in a basic ring only the lowest priority is defined.

The priority scheme employed can be summarized thusly: a station that has a higher priority than that of the current busy token can reserve the next free token for its priority level as the busy token passes by. When the current transmitting station is finished with its transmission, it issues a free token at that higher priority. Because stations of lower priority cannot seize the token, it passes to the requesting station or an intermediate station of equal or higher priority with data to send. Also, the station that upgraded the priority level is responsible for downgrading it to its former level when all higher-priority stations are finished: when the station sees a free token at the higher priority, it can assume that there is no more higher-priority traffic waiting and downgrade the token before passing it on.

2.3. Comparative Performance Of Ring Protocols

Very few studies have been conducted to compare the performance abilities of the three types of ring protocols: slotted ring, token ring, and register inseration ring. M.T. Liu, W. Bux, and W. Cheng, and P.S. Yu have performed studies independent of one another on protocol performance comparisons, from which several basic conclusions can be gleaned.

Liu, et. al., conducted performance comparisons based on analytic models of slotted and token rings developed by other parties plus his own formulations for register insertion rings; his findings showed that the slotted ring performs poorest of the three, and that the register-insertion ring was the best performer -- because its protocol allows multiple frames to circulate.

Bux's analysis efforts led to the finding that the delay characteristics of the token ring are superior to those of the slotted ring, and, like Liu, he found results of register insertion to exceed those of the token ring.

Cheng's results, that the delay of the slotted ring exceeds that of the token ring, agreed with Bux's. Yu found that the delays of the token ring were longer than those of the register-insertion ring.

What does all this mean? It means that far more analysis work is needed before hard-and-fast findings become available. What can be surmised thus far is that the slotted ring -- by virtue of the high overhead of each small packet -- seems to be the least desirable of the three ring protocols (the Cambridge Ring, as shall be seen, uses but 16 data bits in a 37-bit slot) and seems to base most of its value on its reliability.

Of the other two, evidence thus far suggests that register insertion offers superior delay performance. The primary advantage of register insertion seems to be the potentially high utilization it can achieve, in that multiple stations can be transmitting and, further, that a station can transmit as soon as a gap opens up on the ring and need not wait for a token. On the other hand, though, the propagation time around the ring is not constant -- it depends on the amount of traffic.

Why, then, is the token ring taking giant strides toward becoming the dominant ring network in America? Because under heavy traffic conditions almost 100% utilization is achieved on a token ring -- and with IBM and IEEE 802.5 pushing for a token ring standard, there would be no highly identifiable advantage to using register insertion.

2.4. Ring Reliability Issues

While ring networks are more flexible and less complex than bus-oriented networks, and they provide higher throughputs with a distributed means of control, they have been criticized as being less reliable than both the star- and the bus-based networks. Also, while the ring is a decentralized network that avoids the perils of central mode failures, its various access and control strategies are complex and have been said to reduce both its reliability and maintainability.

Cable vulnerability is a problem often used as a knock against ring-based LANs: after all, a break on any of the links between repeaters brings down the whole network. Further, the failure of a single repeater crashes the entire network as well. And once such a single-entity breakdown occurs, the entire ring may have to be inspected ("perambulation of the ring" is another term for this problem).

The point to be considered here is that these faults apply to the most basic ring design; as with all evolving technologies, as problems have arisen, solutions often follow. Ring designers have bypassed these problems by using either fail-safe bypassing or redundant paths.

The wire center approach (refer to Figure I-21) is one type of simple circuit that can be employed to bypass failed nodes. The wire center acts as the hub of a star-configured network. While it is a passive device, the wire center provides a means for users to break into the ring.

Attachment to the ring is most often by means of a relay. When the user host is functioning properly and wishes to join the ring, it energizes its relay at the wire center. If the cable linking the host to the wiring center is broken, for instance, or

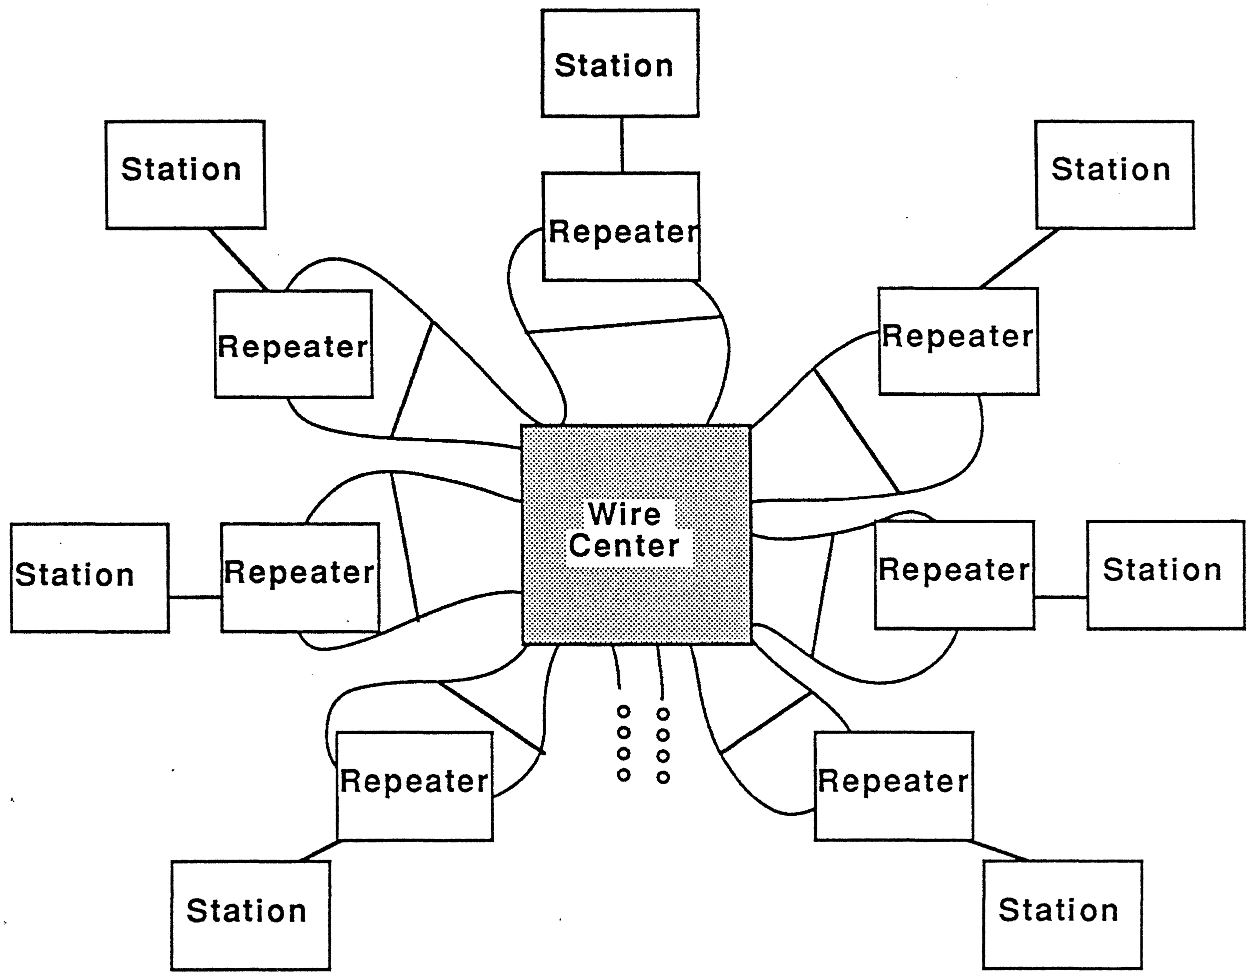

Figure I-21: Ring Wiring Center

there is a power supply failure, the relay (which is energized by the user hardware through cable from the user site to the wire center) simply returns to its normal de-energized state, bypassing the user. Thus defective nodes are automatically bypassed without destroying the entire system.

This wire center scheme does more than just provide such a "graceful" failure mode: it also reduces the maintenance burden, in that not all failures demand an immediate reaction. The wire center concept also enables rapid fault isolation, playing a role in the network identical to that a circuit breaker panel plays in a power distribution system -- i.e., bypassing a defective node without disrupting the whole ring. Also, many wire centers provide centralized sites where trouble signals can be monitored, and some manual means for adding or dropping nodes can be implemented -- all of which constitutes a substantial reduction in the mean time to repair the network.

In that effective communication is another measure of reliability, it should be noted that synchronous operation of the ring ensures more reliable performance at data rates of 10 Mbps and above than does asynchronous operation. Also noteworthy is the fact that the ring design has a high degree of adaptability to existing fiber-optic technology: fiber-optic technology permits the use of longer transmission links with fewer repeaters than does coaxial cable technology.

3. The Empty Slot Approach

3.1. The Cambridge Digital Communication Ring

3.1.1. Overview

The Cambridge Digital Communication Ring (see Figure I-22), was developed at the Cambridge University Computer Laboratory in England in the 1970s and presented to the industry in a paper by Maurice V. Wilkes and D.J. Wheeler in 1979. It is, essentially, a baseband ring predicated upon the empty slot technique. It was designed in response to a need to improve upon the interconnectivity of varied computing devices and the data rate available from existing telecommunication and time-sharing systems. It transmits data, in the form of minipackets, over two twisted pairs of wire, each providing one channel, at a raw data rate of 10 Mbps. It can accommodate up to 256 stations, when the separation between repeaters does not exceed 100 meters. The repeaters are powered directly from the ring.

Each node on the ring (refer to Figure I-23) consists of a repeater, a station unit,and an access box. Each station is fully duplex, thus permitting transmission and reception to occur concurrently and independently. The number of bits delay on each station is about 1.5 bits, and the minimum ring delay is about 5 microseconds. Whereas repeaters are powered directly from the ring, it is generally the associated computer or peripheral that powers the access box and station unit; the advantage inherent in such a scheme is that the repeaters continue to work when the station units are de-energized.

One unique station on the ring is designated a monitor station: it is not connected to any host, and its responsibilities include the establishment of slot structure during turn-on, monitoring of the ring, clearing of lost packets, and accumulation of error statistics. Immediately before the monitor station is an error-logging station, which uses a normal station and receives packets containing error information sent from active ring stations and the monitor station.

Data on the Cambridge Ring is transmitted between stations in bit-serial form as 5V pulses. Channels are coupled to the station logic by transformers, and the electronics scheme employed is such that if a station fails, the repeater then acts as a passive coupler.

Power on the Cambridge Ring is conveyed as 50 Vdc either on the signal wires or other wires provided for this purpose. In that the Cambridge Ring is a point-to-point system, it is not necessary to use the same medium throughout, and thus fiber-optic or coaxial cable may be used as well (particularly where a long link is needed) -- though fiber-optic cable can but carry a

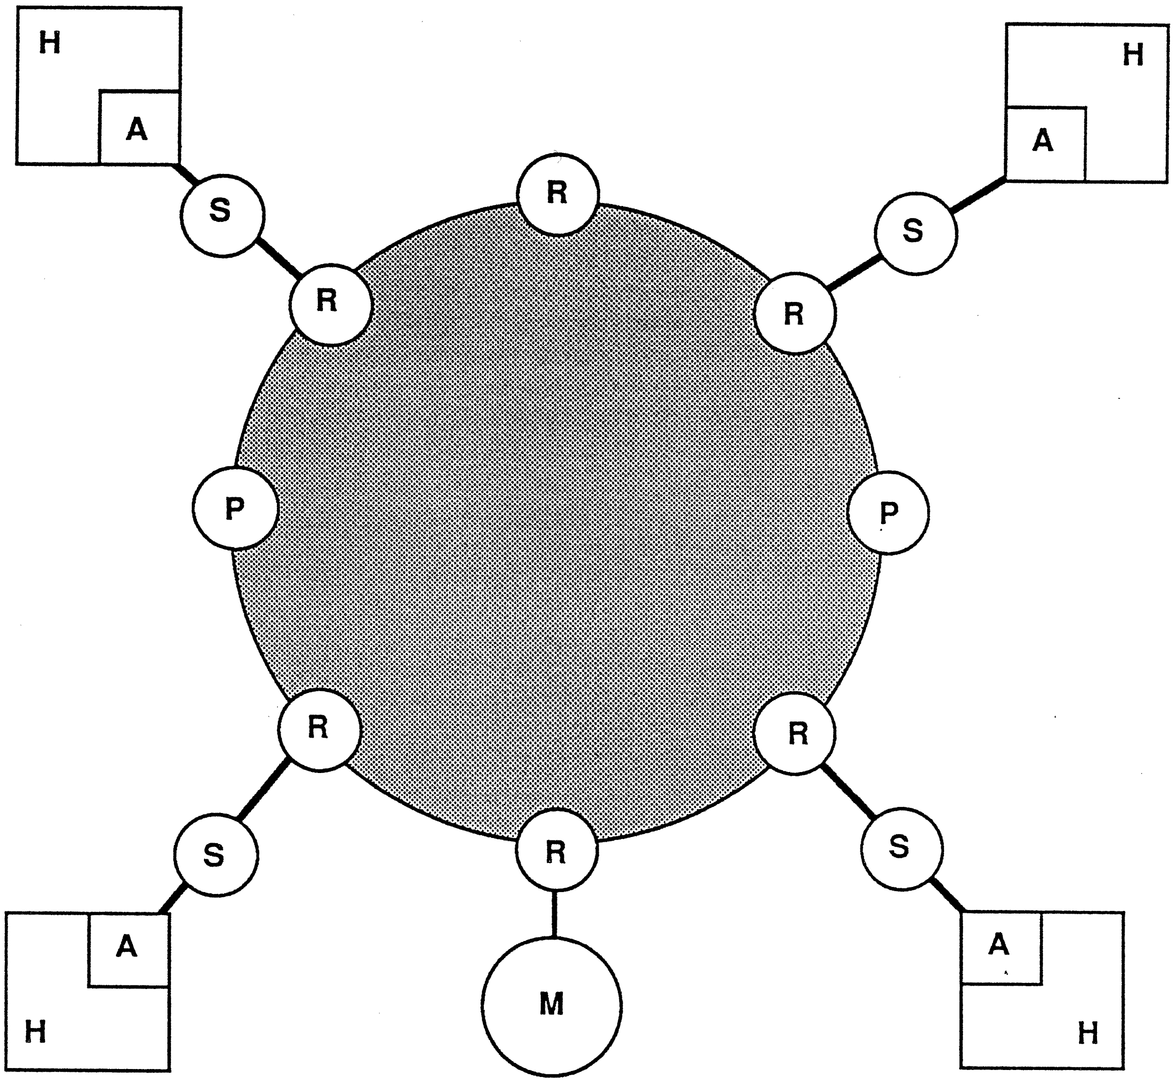

R = Repeaters
S = Station Unit
A = Access Circuits
H = Host Computers
M = Monitor Station
P = Power Supplies

Figure I-22: Typical Cambridge Ring Configuration

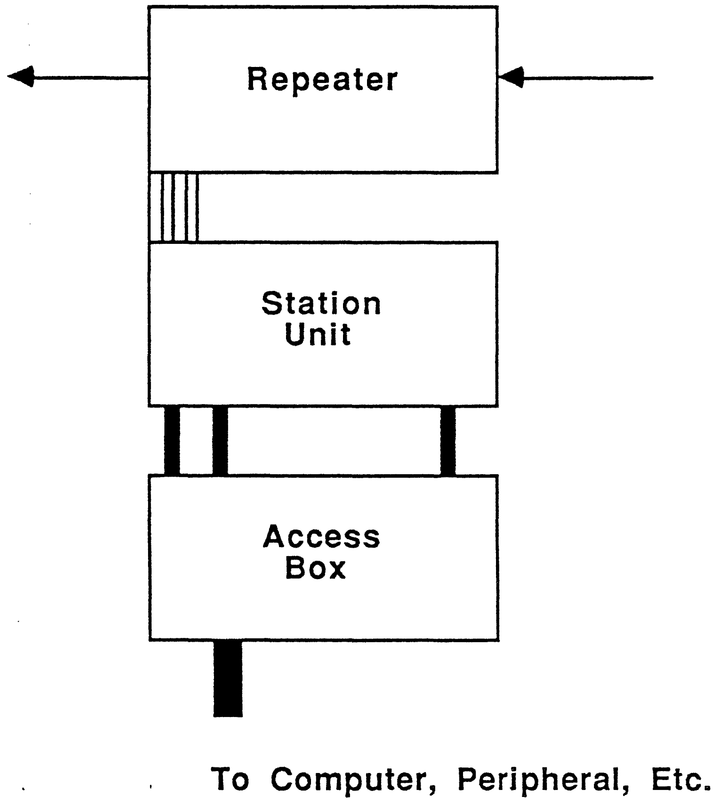

Figure I-23: Typical Node, Cambridge Ring signal, as it cannot carry electrical power for the stations.

Ironically enough, the Cambridge Ring was originally designed on the register-insertion principle, which was selected for its guaranteed anti-hogging properties; by the time a broadband model of a station working on the principle of register insertion had been constructed, however, it was realized, by its developers, that the empty-slot control strategy offered equally effective anti-hogging properties while avoiding two reliability-related shortcomings of register-insertion: a fault developing in a shift register can disrupt the entire action of the ring, and, if the fault develops in an inactive station, this fault will not become apparent until that station begins to trasnmit.

Also, with the empty slot technique it was found that an additional shift register permanently connected in series with the ring at each station could be eliminated, thereby reducing to a minimum the amount of equipment in series with the ring and proportionally increasing the ring's reliability.

It was subsequently found that the original broadband could readily be modified for the empty slot technique and that the total count of TTL package remained the same, and the changeover was implemented.

3.1.2. Hardware Components, System Configuration

The Cambridge Digital Ring -- or Cambridge Ring, for short -- transmits signals at 10 MHz, in bit-serial form as 5V pulses, along a ring (or loop) of two twisted pairs of wire. Repeaters along the ring receive their power directly from the ring and must operate reliably whether connected to a station or not; this power is injected into the ring by power supply units (PSUs) located at a number of independent points.

Each node on the ring consists of a repeater, a station, and an access box. Each station operates in full-duplex mode, enabling transmission and reception to occur concurrently and independently. The number of bits delay on each station is, again, about 1.5 bits, and the maximum ring delay is approximately 5 ms. Stations are identical except for the logging and monitor stations: the former records errors for purposes of maintenance, while the latter initializes the ring and maintains the integrity of the slot structure.

In its most essential configuration, each node -- or station -- consists of a repeater, a station unit, and an access box. It is the task of the repeater to manage the flow of bits on the ring; the station unit implements the minipacket protocol (see 3.1.3 below); and the access box provides conversion between the network-specific functions of the station unit and the attached computing devices. (The computing devices -- computers, terminals, servers, etc. -- are connected to the ring via the access box-station unit-repeater chain.) The interface between the computer and access box is generally provided on the computer's own internal bus -- although on the first Cambridge Ring model, it was the access box which handled this interfacing, having to be specially built to interface a particular kind of computer or device to the ring.

The original Cambridge Digital Communication Ring (see Figure I-24) set out to provide -- by means of interconnected minicomputers and microcomputers -- facilities similar to those commonly provided by a conventional time-sharing system. In it, terminals called visual display units (VDUs) were linked to the ring via concentrators. On log-in a user was assigned a computer from among the pool of computers, servers, terminals, etc., available for assignment. The user had exclusive use of that computer for the duration of his/her session, i.e., until log-off.

To support the time-sharing environment, special-purpose elements were also included on the ring and are described below.

- VDU Concentrator - the VDU concentrator communicates with other parts of the ring; there is enough memory space in the Z80A microcomputer in the VDU concentrator to buffer a line being typed on each VDU and also to buffer lines received for display. The VDU also provides local line-editing facilities.

- Name Server - the name server, when presented with the name of a service, a process, or a computer anywhere in the system, returns the location of that item (it can also perform the converse translation). The name server provides the user with a flexible system for naming and completely frees him/her from the need to know the physical placement on the ring of services, processes, etc.

- File Server - the file server provides a complete filing system which holds both user and system files in the normal manner. It can further be used as a backing store: i.e., it can be used for swapping or overlaying the contents of the high-speed memory of a computer.

- Printing Server - the printing server provides spooling facilities and facilities for printing files. It makes use of the file server for storage, as the printer itself is not connected directly to the printing server but is separately interfaced to the ring.

- Time Server - the time server is a Z80A microcomputer linked to a radio receiver that receives continuous date and time signals (including pulses at one-second intervals) transmitted on 60 KHz from a nearby radio

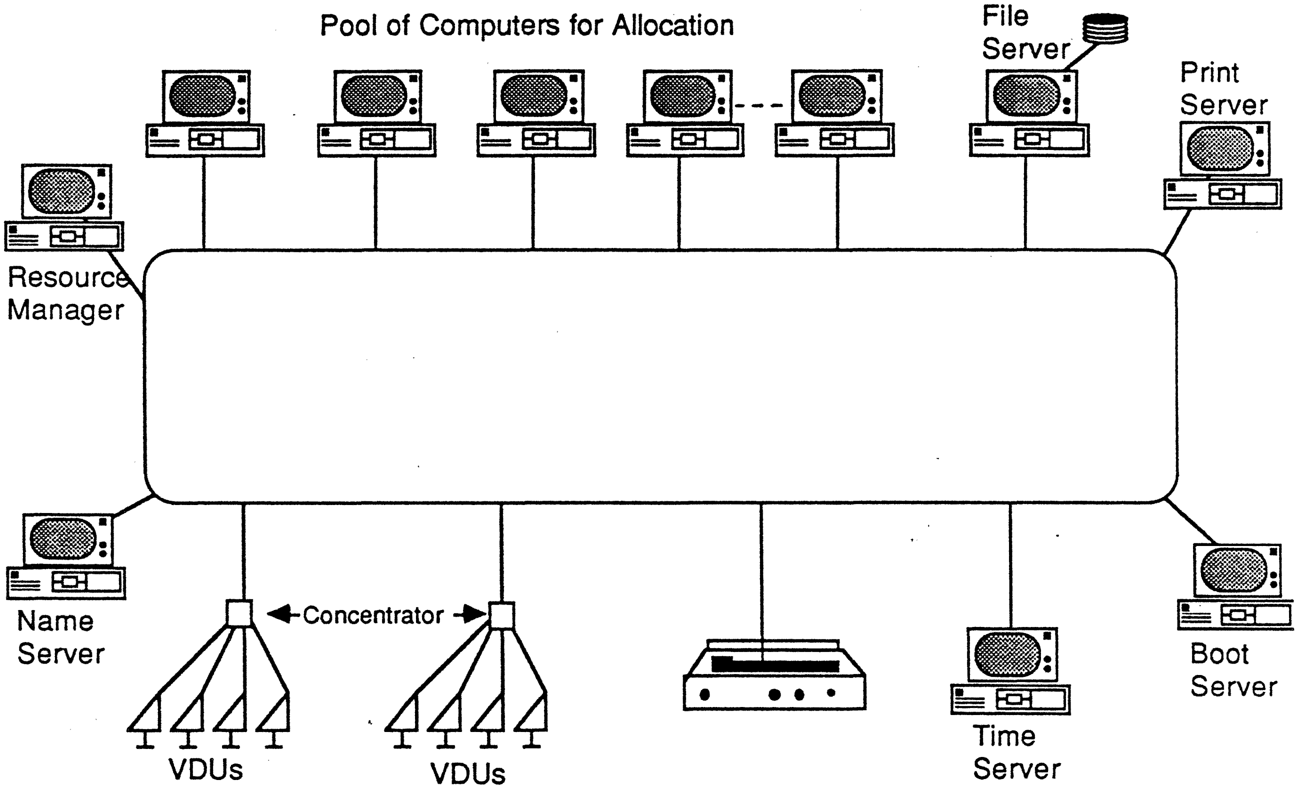

Figure I-24: Original Cambridge Digital Communication Ring station; upon request, the time server supplies the date and time to the nearest millisecond.

- Boot Server - all the servers in the system are provided with a bootstrap routine in ROM which is automatically activated when a server is switched on and causes an application to be made to the boot server, which stores the programs required for the initial loading of the various computers on the ring; the boot server recognizes what the function of each server is and delivers to it along the ring a program (normally, the initial program required), which it then loads into its memory.

- Resource Manager - the resource manager validates user log-in information, assigns servers to users, and maintains integrity of the system.

When the Cambridge Digital Communications Ring was defined in a paper by Wilkes and Wheeler in May of 1979 at the LACN Symposium, the devices attached to the ring was defined as follows: a PDP7 (used for initial testing), two PDP11s, the Cambridge CAP computer (an experimental computer built in connection with a memory protection project), a NOVA computer, a Computer Automation LS14 with two 80-Mbyte disks, and a free-standing plotter. A VDU concentrator based on a Z80A microprocessor and capable of providing a buffered interface for four VDUs had been constructed, and various other Z80A-based units were also in use or under construction.

In that same paper, Wilkes and Wheeler acknowledged that the prototype Cambridge Ring would inevitably undergo change and be improved to keep abreast of the burgeoning local area network technology. While it has, the basic configuration of the ring system has remained essentially the same, with sophistication being added primarily through new attached devices and new features incorporated into the attached devices. The system hardware described here was included by virtue of its historicity, to present an accurate overview of the configuration of the first Cambridge Ring; while the attached machines have changed, the operational principles remain basically the same.

3.1.3. Minipacket Protocol

Data is transmitted around the ring in 38-bit minipackets (see Figure I-25). The first bit is always a "1"; the second bit indicates whether the minipacket is empty or full, and the third is the monitor bit. Following these are the destination and source bytes (8 bits each), followed by two bytes (16 bits) of data. These are in turn followed by two response bits, which are both set to "1" when the minipacket is sent: if the destination is absent, they will be unchanged (i.e., 11) when the minipacket returns to its originating point; otherwise they are changed by

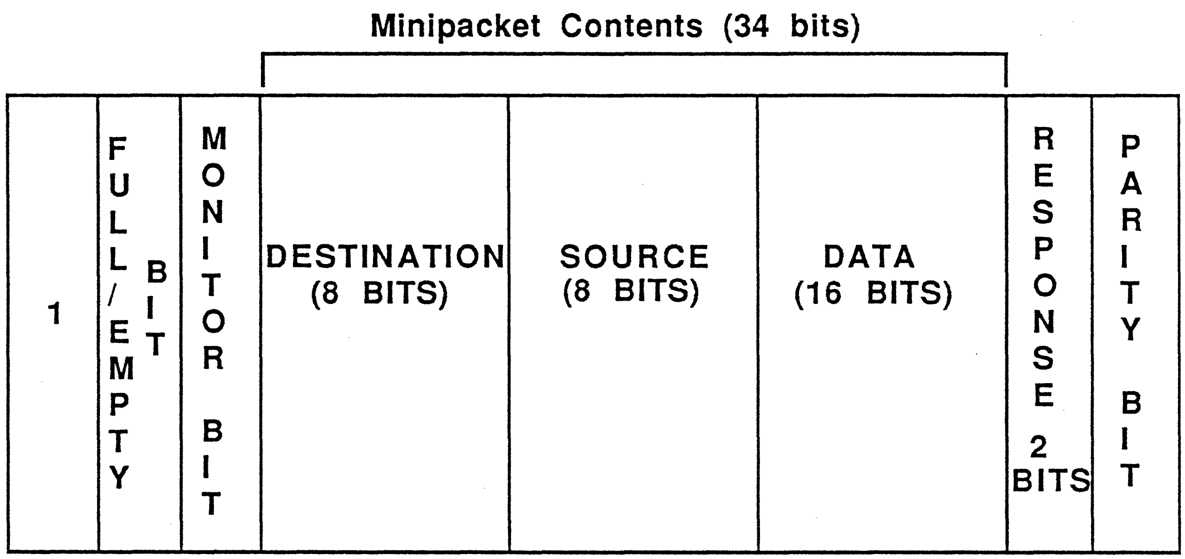

Figure I-25: Cambridge Minipacket Format the destination to indicate whether the minipacket was accepted (01), the destination was deaf (10), or the destination was busy (00). Last on the minipacket is a parity bit, which is not copied at the destination and is used only by the error-reporting service.

With the Cambridge Ring -- an empty-slot network -- a constant number of slots circulate around the ring head-to-tail, up to the capacity of the ring -- which is determined by the length of the wire and the number of repeaters. Each slot is able to contain one minipacket. The train of minipackets is terminated by two or three zeros -- called gap digits -- which, together with the leading ones in the minipackets, are used for minipacket framing and counting. (When Wilkes and Wheeler presented their paper in 1979, the model of the Cambridge Ring then in operation contained but two minipackets -- and even this required a shift register of about 20 digits in length in the monitor station to increase the ring's length artificially. A later model of the Cambridge Ring, dated 1982, was about one kilometer long and had 50 devices attached to it.)

3.1.4. Modus Operandi

To transmit a minipacket on a Cambridge Ring, it is necessary to load the data bytes and a destination byte into the station and issue a transmit signal. The station then watches for an empty slot; when one arrives, it is marked full and the source, destination, and data fields are filled in. The minipacket is then sent along its way, eventually returning to the sending station.

On the way it should have passed its destination, at which time one of a variety of possibilities will have occurred: No Action -- the receiving station was switched off (or non-existent); Accepted -- the data and source bytes were copied into the receiving station; Busy -- the receiving station was switched on but its reception register was not empty; Unselected -- the selection register in the receiving station had been set to exclude reception from this source, or from all sources but one, and this is not that source; or returned minipacket in error. The status is then marked in the minipacket and reported to the originating station.

If the destination was busy, the minipacket may be sent again immediately; if it was deaf, the minipacket may either be re-sent immediately or after a delay; or the attempt may be altogether abandoned. The appropriate action is determined by a higher-level protocol, to which all such rejection information is passed.

In order to touch upon another important related point, this method of operation may be restated thusly: when the minipacket has returned, the five bits available for reading by the access box correspond, as mentioned above, to busy, unselected, accepted, unrecognized, and returned packet in error. These bits -- or some of them -- must be inspected before the decision can be made whether to resend the minipacket that has just returned or send a new one; therefore, it is not possible for the access box to fill the minipacket immediately after the one that has returned -- even if it is empty. When allowance is made for this and for the fact that a minipacket carries nearly 60% overhead, it is found that the maximum data rate at which a station can transmit is approximately 0.4/(n + 2) Mpbs, where n is the number of minipackets in the ring; for the special case of two minipackets circulating, this comes out to slightly over 1 Mbps.

Once the minipacket returns to the sending station, it sets the "full-empty" flag to "empty," making it available to other stations. In order to prevent one station from hogging all the available transmission capacity, the sending station is not permitted to reuse the packet itself. The sending station may recognize its own packet either by counting the packets on the ring, or by reading the source address -- the former option was used in the original system at the University of Cambridge, but it requires that either the number of packets on the ring is fixed or that all stations must always know the actual number.

3.1.5. Maintainability, Reliability

When the first version of the Cambridge Ring was being developed, the issues of reliability and maintainability had to be addressed. These issues centered, respectively, about the removal from the ring of faulty or indefinitely circulating minipackets and the detection of breaks in the ring, and were addressed by the inclusion of a monitor station and a logging station, respectively.

In the empty packet system, it is necessary to provide a monitor station which will, in essence, remove from the ring minipackets that circulate indefinitely as a result of some fault condition. Thus, a monitor bit was included in the minipacket which is unset when the minipacket is launched and is set at the monitor station. The monitor station marks as empty any minipacket which reaches it with the monitor bit already set.

Errors which can cause minipackets to circulate indefinitely are very rare in the form of the empty packet system adopted: an example is the full/empty bit in an empty minipacket becoming accidentally reversed. Corruption of the source byte in a minipacket would not cause the minipacket to circulate indefinitely, since this byte plays no part in the recognition by the source of the returning minipacket.

The monitor station may also provide the means whereby the ring is initialized upon switching on or re-initialized after a failure, and it can also be designed to keep statistics of ring performance and loading.

The designers also desired that there be some feature whereby a station may render itself deaf to stations from which it does not wish to receive packets, and so it was decided to incorporate in the station unit a single byte register known as the source select register. When this register contains all ones, the station will receive a packet addressed to it from any source; when the register contains all zeros, the station will not receive packets from any source -- otherwise the station is open to receive packets only from the source whose address is specified by the number standing in the register.

As to the problem of readily locating breaks in the ring, the first ring to be built had an echo feature: that is, it was possible to launch from the monitor station an interrogating packet which would be recognized as such by the station to which it was addressed; the station would then respond by sending a signal along a logically separate low bandwidth channel to the monitor station. Valuable, however, as this feature might be in the location of hard faults, the designers realized it would not help with the location of intermittent or rarely occurring faults, or with the identification of places at which noise was finding its way into the ring. A proposal for continuously monitoring the ring and reporting errors to a logging station was made, and this proposal was so attractive that the ring was rebuilt to implement it.

In the most recent Cambridge Ring system, each packet has an additional trailing bit used as a parity bit. As the packet leaves each station -- including the station at which it is launched -- this bit is set so that the packet has even parity. When a packet enters a station, the parity is checked by separate circuits: if it is found to be correct no action is taken; if, however, it is incorrect, the station that has discovered the error proceeds autonomously to launch a reporting packet to a logging station, which has address zero. Thus, if one of the repeaters develops an intermittent or marginal fault, or if noise enters the ring from outside, the logging station will begin to receive reporting minipackets from the station immediately downstream of the trouble. The faulty packet continues on its way, but, insofar as the parity bit is reset as it leaves each station, the stations downstream of that at which the parity was found to be wrong will not give rise to further reporting packets.

The logging station is placed immediately ahead of the monitor station, and reporting packets are launched with the monitor bit already set, so that they are removed by the monitor station without making a complete circuit of the ring. As an alternative to providing a separate logging station, facilities for recording information arriving in reporting packets may be built into the monitor station.

The monitor station fills empty minipackets with random digits, and these minipackets are subjected to the procedure just described in exactly the same way as the full ones. The functioning of the ring is thus continuously monitored whether or not it is carrying traffic. The staff responsible for the maintenance of the ring can thus receive immediate notification that trouble is occurring and is in no way dependent on reports received from users.

The leading digit in a packet is always a "one," so that in normal operation a station can never receive a long unbroken stream of zeros: if this situation does occur, the design of the station is such that it will emit a sequence of reporting packets as long as the condition persists. At the same time, it passes downstream what will appear to the next station to be packets of normal parity. Thus the reporting system enables breaks in the ring -- and intermittent faults -- to be located.

3.1.6. Performance

Discussion of Cambridge Ring maintainability and reliability invariably leads to accompanying consideration of the ring's performance. Much of the industry's evaluation of the Cambridge Ring's performance is expressed in terms of comparison between its performance and that of Ethernet -- an evaluation which does much to explain why the Cambridge Ring is the most common method of networking in Europe, while Ethernet is one of the more prevalent means of networking in America.

Several comparative evaluations of the performance of the Ethernet and the Cambridge Ring have concluded that Ethernet functions well under light loads, but that with increasing traffic loads time is wasted in collision avoidance. It has further been pointed out that in Ethernet, as the cable length increases, a further deterioration in performance occurs as the result of a longer contention interval, whereas the ring is said to maintain high system throughput even under heavy loads.

Conversely, some researchers have concluded that the minipacket protocol employed by the Cambridge Ring causes higher expected delays compared to the Ethernet for most configurations. Yet it has also been shown that if the LAN size increases, the performance of the Ethernet is subject to degradation, while the Cambridge Ring remains relatively unaffected. This is said to be due to an increase in the contention interval in Ethernet, whereas the ring compensates for the increased ring delay by the accommodation of more circulating slots.

Research has also shown that Ethernet LANs have a higher variance of delay than rings at all load levels, as the back-off-in-case-of-collision algorithm employed by Ethernet increases the retransmission delay with the increasing number of collisions (i.e., it takes much longer to transmit for stations which have attempted to retransmit several times than for a station transmitting for the first time). The Cambridge Ring, on the other hand, guarantees transmission time not to exceed a set duration -- a maximum time relative to the number of stations and minipackets on the ring.

Pursuing this comparison even further, it has been shown that both possess similar performance characteristics and that delay is related to the number of nodes in the LAN. While the ring has the largest theoretical delay, in contention networks errors are assumed to occur frequently: rings can be made relatively error-free, whereas in Ethernet error occurrence is assumed to be frequent, demanding powerful error detection facilities.

Also, Ethernet was not intended for use in hostile environments or applications wherein a real-time response must be guaranteed, while the Cambridge Ring seems able to handle these environments quite readily.

Zeroing in on the strengths and allure of the Cambridge Ring, it can be said that its strongest features are its high bandwidth and its very low error rate (one error in 5 x 10**11 bits). Such a low error rate is significant, as it allows protocols to be simple, in that they need not be so concerned with the transmission of correct packets should a fault occur; instead, this can be handled by higher-level protocols which would merely request that the packet be retransmitted. The bandwidth itself is a strong function of ring size and -- more particularly -- the number of circulating slots.

Because of these and other considerations, England and Europe espoused the concept of the Cambridge Ring over that of Ethernet and so set themselves apart from local area networking trends prevalent in America. The prototype Cambridge Ring, and its successor models, gave rise to a number of privately built ring-based LANs being marketed in England and on the continent, several of which will be examined below.

3.2. Other Empty-Slot LANs

3.2.1. SEEL's TransRing 3000

TransRing 3000 is Scientific & Electronic Enterprises Ltd.'s (SEEL's) version of the Cambridge Ring. It uses two twisted pairs for cabling (or fiber-optic) and offers a raw data rate of 10 Mbps (see Figure I-26).

Among the interfaces available for TransRing 3000 are an Interface Interrupt Unibus, an Interrupt Q-Bus, an S-100 bus, a DMA Unibus interface, and a DMA Q-bus interface. A Z80A-based terminal concentrator is capable of supplying RS-232-C channels

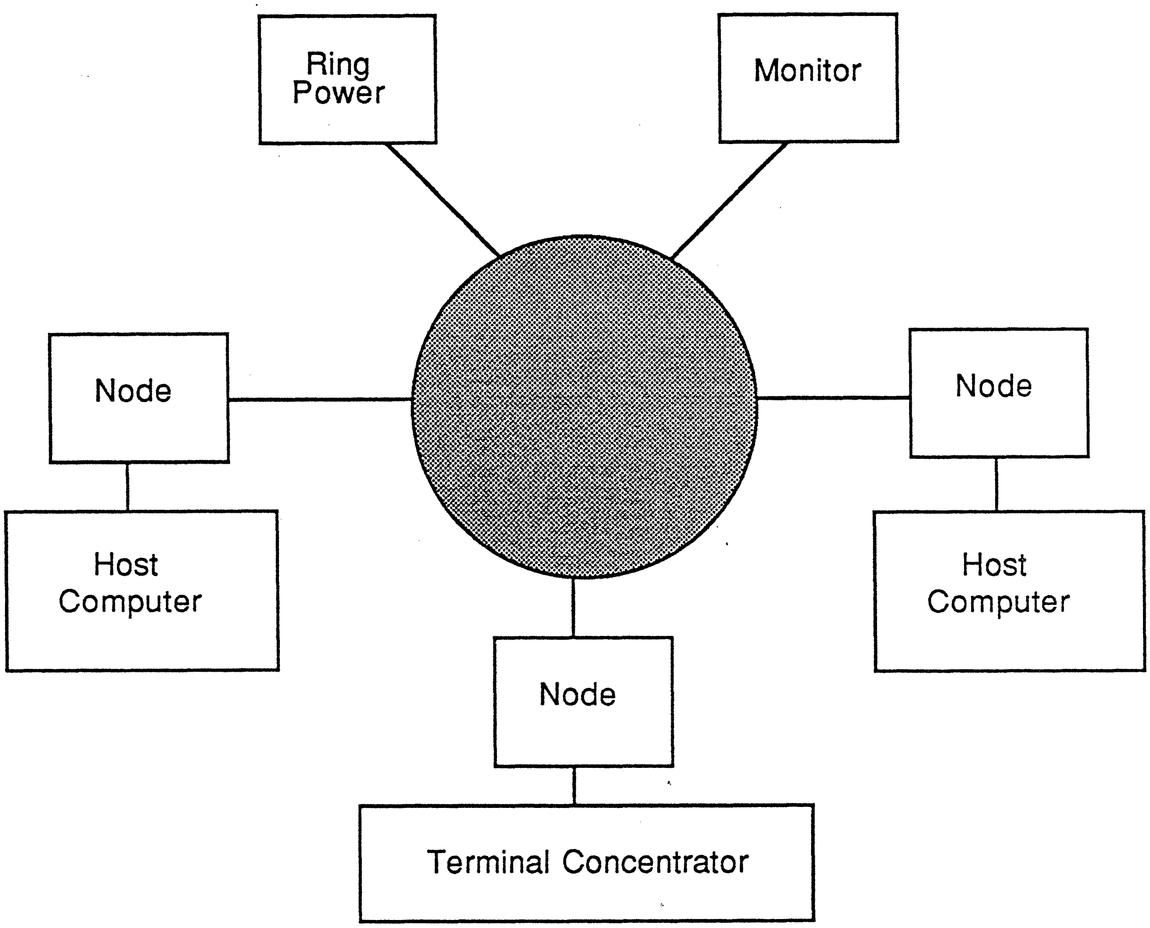

Figure I-26: SEEL TransRing 3000 Configuration for terminals. It supports a maximum of 254 stations, with 300 meters maximum between stations.

SEEL says that the TransRing data protocol makes possible the addition of a new processor to the LAN: communication with all other processors attached to the ring is possible, as the only requirement is a hardware and software interface between the new host and the ring.

The main modules for the system are the monitor station, repeater, and access logic; versions of TransRing are based on Eurocard-sized circuit boards designed by SEEL using MSI chips and a Ferranti ULA two-chip.

While SEEL has systems installed at its offices and the Edinburgh Regional Computing Centre, Hatfield Polytechnic and a group of southern universities have also installed their own TransRings.

3.2.2. Xionics' Xinet

Xionics, Ltd. offers a ring-based system (refer to Figure I-27) that contains three principal elements: Xinet, the Xibus master node, and multifunction workstations.

Xinet is a ring-based LAN operating at about 10 Mbps. It is formed of "intelligent sockets" interconnected by multicore cable. These intelligent sockets are called communications adapters and contain the many software routines which are necessary for various network protocol levels. Each such socket possesses a packet of its own, and thus transmission does not necessitate the sockets competing for a limited number of slots.

The second element of the system is the Xibus master node. This consists of a multiple microprocessor resource combined with high capacity Winchester-technology disk drives. The Xibus master node performs three main functions: it facilitates the sharing of a database; it presents a location index for any data, resource, or user forming part of the architecture; and it supervises the network by providing the database administrator with facilities to gather statistics and interrogate or alter configuration details.

The third system component of Xinet is the multifunction workstation, which is connected to one of the intelligent sockets. It allows users access to information sources and processing resources within the system.

The storage, retrieval, and processing of information is achieved by Xibus downloading the necessary software routines to the Z80 microprocessor-based workstations. Programming languages available include BASIC and COBOL.

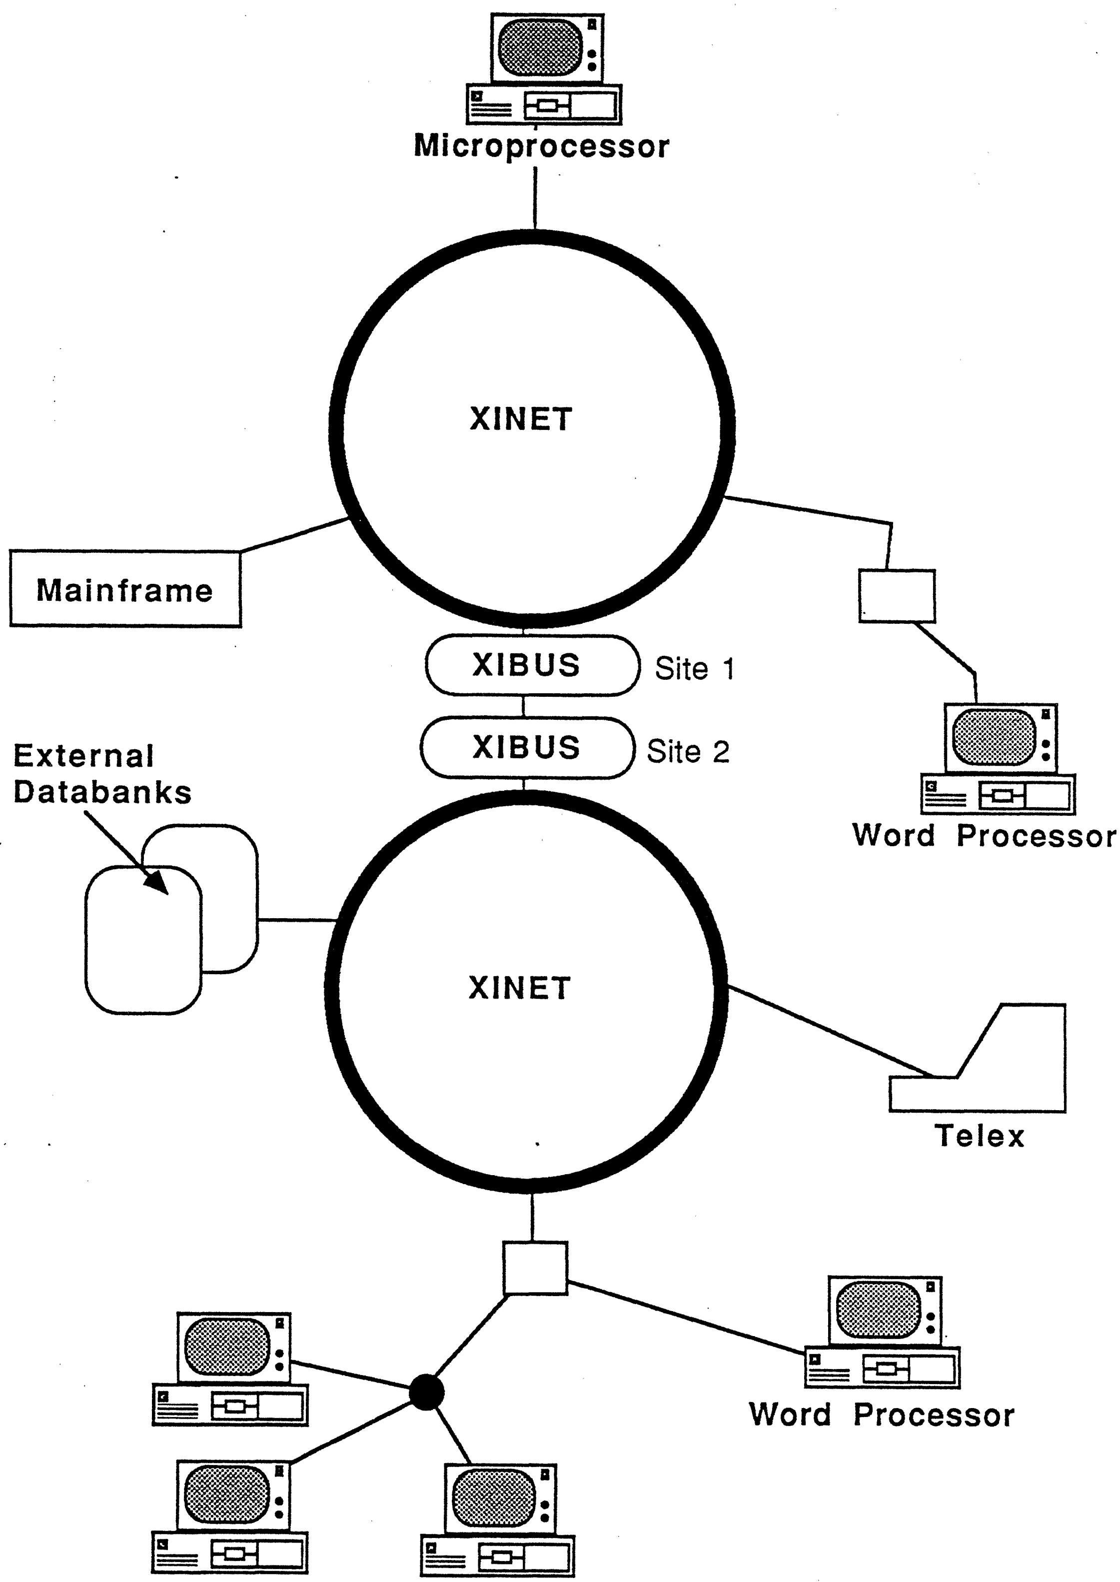

Figure I-27: Xionics Xinet, Multiple Site

Every functional module within Xibus is at least duplicated. Duplicated rings operate in parallel in Xinet, while socket logic is duplicated in the intelligent sockets.

Xibux/Xinet architecture may be used for conventional computer transactions, documents, messages, and digitized speech.

3.2.3. Toltec's Data Ring

Data Ring, offered by Toltec, Ltd., is yet another circulating slot Cambridge Ring (refer to Figure I-28). Toltec can supply a full range of software for Data Ring, most of which is written in high-level language BCPL. It can run under the native operating system, or under TROUT -- Toltec's own portable packet-oriented operating system also written in BCPL.

Drivers and handlers comprise the basic software; they control data transfer to and from a given station, implement the basic block protocol (BBP), and maintain a list of reception requests relating to port numbers. (For particular devices, these functions may be assigned to the access logic.)

The single-shot protocol (SSP) employed on Data Ring is a combination of transmitted and received basic blocks and is the fundamental unit of conversation for standard interfaces to the ring handler. File transfer between dissimilar operating systems is accomplished using the byte stream protocol (BSP).

Data Ring employs fiber-optic cabling and transmits data at a rate of 10 Mbps. Software consists of file transfer, file server, print server, and remote terminal server.

A wide variety of applications is claimed for Data Ring, covering industrial and commercial communications requirements such as process control, office automation, test equipment, and management information. A file server allows the implementation of shared logic word processing; the output can be through shared printers connected via terminal concentrators in a manner similar to the terminals themselves. A "gateway" machine also allows geographically remote networks to communicate with each other.

3.2.4. Racal-Milgo's Planet

Racal-Milgo Limited is yet another United Kingdom firm that offers a Cambridge Ring/empty slot technique type of network, in its Planet Token-Ring Series of LANs (see Figure I-29). There are, however, several modifications incorporated into the Planet Series that differentiate it from the Cambridge Ring.

The Planet Token-Ring Series of LANs uses baseband coaxial (RG-59) cable (not twisted-pair) in a twin-ring configuration. It transmits data in full- or half-duplex at

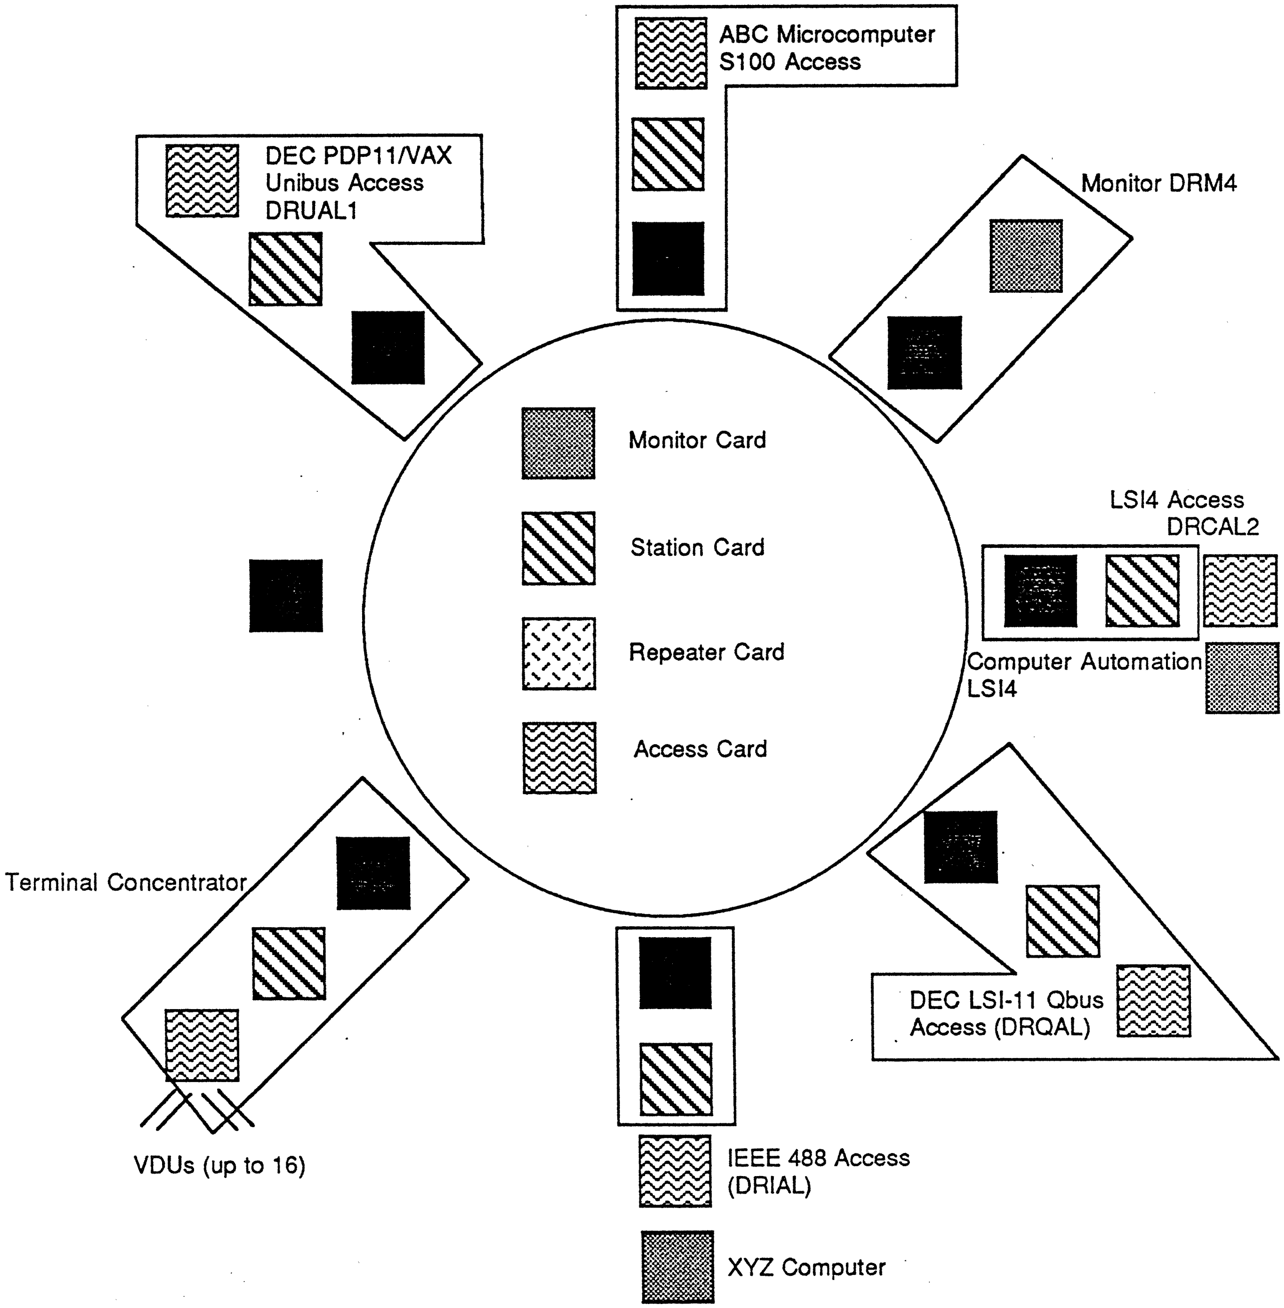

Figure I-28: Toltec Data Ring

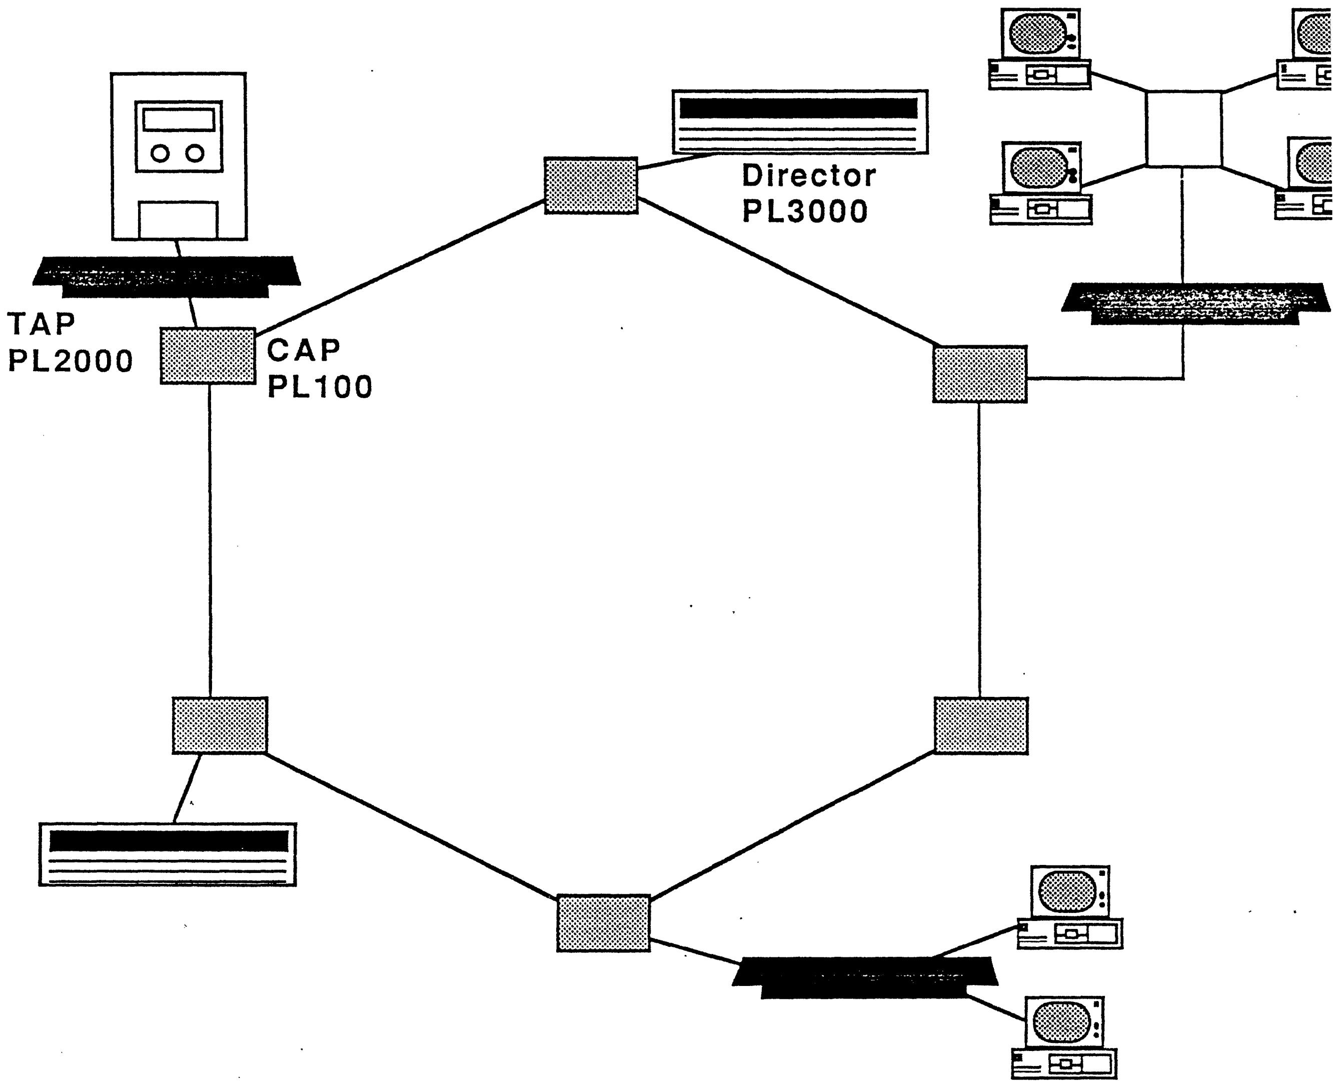

Figure I-29: Racal-Milgo Planet System

rates from 110 bps to 19.2 Kbps and can support up to 500 user ports, per ring, on up to 250 TAPs.

The system consists of three components: (1) the Director, an intelligent network processor; (2) Terminal Access Points (TAPs), up to 250 of them, each of which has two RS-232 ports; and (3) the Cable Access Point (CAP), which acts as an outlet to the cable. These components correspond roughly to the Cambridge Ring's monitor station, stations, and repeaters, respectively. Racal-Milgo says the Planet Series accommodates communicating devices from different vendors, regardless of protocol. Also, the Director features a failsafe capability, in that it immediately detects any CAP or TAP failures or cable breaks and automatically reconfigures the ring system by isolating the failed component and bypassing the problem area.

Differences worth noting are: (1) in the Planet Series, the minipacket is 42 bits in length, with 16 bits of data and one 16-bit destination address contained -- in the Cambridge Ring, minipackets include both 8-bit source and destination addresses; and (2) Planet assumes the form of a double ring of coaxial cable and does not use twisted pairs of wires as do other Cambridge Ring products.

On the Planet Series, the CAPs are passive units and connect directly to the cable; the TAPs -- which are microprocessor-based -- interface user devices and connect to the CAPs and possess the resident software which is responsible for packet transmission and reception over the network. The administrator station is functionally more complex than the monitor station on the Cambridge Ring implementations: included in the administrator station on Planet is a name server (which would be provided by a functionally separate device on other Cambridge Rings); essentially, the function of the name server is to establish a call between two nodes by checking that the destination node exists and determining the location of the node.

The Planet Series is marketed in America by Racal-Milgo Information Systems, based in Plantation, FL.

3.2.5. Logica VTS' Polynet

Still another United Kigndom company to base a LAN on the Cambridge Ring is Logica VTS Ltd.

Its ring-based LAN, dubbed Polynet and first announced in February 1981, interconnects up to 254 nodes using twisted-pair telephone grade cable (refer to Figure I-30). Operation is at 10 Mbps with distances of up to 100 meters between nodes. A typical Polynet system consists of network nodes, a monitor station used for network initialization and monitoring, wall sockets that allow connections between the cable and the nodes, and an appropriate number of power supply units. A number of interface

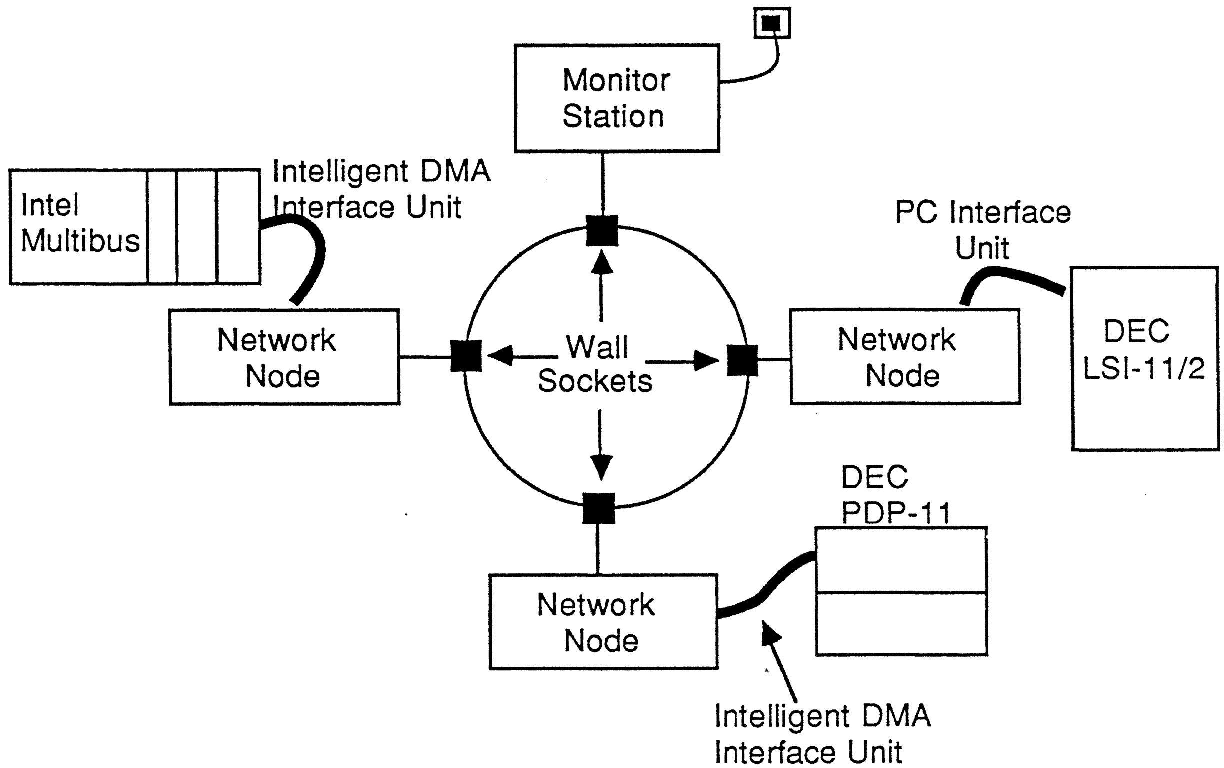

Figure I-30: Logica VTS Polynet

units -- including intelligent microprocessor-based units -- are currently available to connect computers to the Polynet: these are for the DEC PDP-11 or LSI-11, IBM PC, the Intel Multibus systems, Q-bus, Unibus, Multibus, and Logica Kenner.

With Polynet, two forms of cabling can be employed. First, it is possible to configure a ring by installing a set of wall sockets and interconnecting them using triple twisted-pair cable; a plug that resembles a telephone jack plug is used to link Polynet components at convenient points, and the nodes themselves may be attached or removed from the wall sockets by their plugs without affecting the operation of Polynet. Alternatively, where increased flexibility is necessary, daisy-chaining is also possible: this is made possible by the provision of an XLR plug and socket on each Polynet unit.

The network node has been designed such that it supports two bytes of user data either in a 38 or 40 bit circulating slot and is attached directly to the network medium. A single 50-conductor ribbon cable links the interface unit of the attached device to the node. Logica provides two kinds of Polynet interfaces, DMA and interrupt: the former provides point-to-point connection between attached devices at a rate of 1 Mbps, while the latter provides a lower data rate of 100-300 Kbps.

Software includes VAX-VMS Driver and RSX-IIM Driver.

4. The Token-Ring Approach

4.1. IBM's Token-Ring LAN

When, on October 15, 1985, IBM finally announced its long-awaited Token-Ring Network, the effect was almost one of anticlimax: after all, IBM's interest in and involvement with token rings dates back at least to early 1981 and IEEE 802, and its sporadic Token-Ring-related announcements -- along with its calculated silence on the matter and resultant industry-wide conjecture -- did much to keep the LAN industry on edge with uncertainty. The seldom-documented progress of the IBM Token Ring made for most interesting speculations on the part of both its competitors and its hangers-on, and it served to rivet industry interest on IBM, for, as IBM turns, so, invariably, does the rest of the industry.

In being mindful, then, of the sense of suspense that pervaded the networking arena for the preceding four years, this report will chronicle, step-by-step, the events leading up to IBM's announcement of the Token-Ring Network, and will then examine the network in detail with its treatment of the October 15, 1985, formal announcement by IBM of the Token Ring.

4.1.1. The Soderblom Patent

While IBM first publicly expressed interest in token rings in early 1981 and began its support of IEEE 802, it is now known that at some time during 1980 IBM quietly acquired a non-exclusive license to a patent governing token-ring networks from one Olof Soderblom. What is not known is the amount of money IBM paid Mr. Soderblom for the license to his patent, though it is generally speculated to have been in the $5 million - $10 million range. The full text of the Soderblom Patent -- reissued March 19, 1985 -- is reprinted in Section III of this report.

When Soderblom first attempted to file for a patent on token-ring technology in Europe in the late 1960s, no one within the networking industry paid him much heed. Soderblom filed for his patent in the U.S. in October of 1974, and it was issued in October 1981. Soderblom later refiled to broaden the patent with 33 extended claims, and, as alluded to above, a re-issued patent, number 31,852, dated March 19, 1985, was granted to Soderblom for a data transmission system.

As the LAN industry and IEEE 802 became aware of Soderblom's patent and IBM's esoteric foresight in acquiring a license to the patent, a substantial uproar ensued, with Soderblom being referred to as a "pirate" and his patent being likened to "blackmail," and many firms saying they refused to deal with Soderblom. However, it currently seems that as the terms (costs) of licensing agreements have come down, manufacturers are beginning to sign up: for example, Ungermann-Bass, Inc., one of the industry's pre-eminent LAN manufacturers and a pioneer in LAN development, was the first major vendor to sign a licensing agreement for the patent, in order to continue producing its own line of token ring products. The basic areas of pricing will be for the chip manufacturers and board or NIU manufacturers.

The signing of manufacturers may have been accelerated by a lawsuit filed in mid-year of 1985, in London, against Logica VTS, a Cambridge Ring manufacturer, on the grounds of an infringement of the token-ring patent. The suit was filed by Willemijn Holding BV, a Rotterdam-based company that owns the Soderblom patent rights. (It is also said that Proteon, Inc., is being scrutinized for possible patent infringement.) George Vande Sande, a partner in the Washington law firm of Pollock, Vande Sande, and Priddy and counsel for the patent, has stated that there has never been a serious patent examiner challenge to either the original or reissued patent, adding that in no case did the examiner ever ask that claims be limited, and that in all cases all of Soderblom's claims were allowed.

According to Vande Sande, board-level pricing has been quite firmly established. Pricing is based on the cost of the board or NIU. For network interface nodes, a flat 2.75% of the selling price of the node is being asked, up to a $275 maximum charge per node. Circuit board pricing is a bit more complex: the fee for a circuit board priced under $400 is a straight 6%, but, for boards over $400, a flat fee of $13.60, plus 2.6% of the cost over $400 (which is not to exceed the total price of $275) is assessed.

Vande Sande has said that negotiations are currently being conducted with a number of manufacturers in the U.S. and abroad, with varying degrees of success. He noted that in the U.S., manufacturers generally talk a tough line in public, but in private, negotiations are proceeding very well, and a number of agreements are close to being signed. It should be noted here that, compared to prices asked for in the original patent, the current prices seem to be very reasonable.

From Vande Sande's perspective, the key issues at stake are the exact details of the terms, the rate at which the manufacturers sign up for the patent, and the establishment of a pricing strategy for the chip manufacturers. He believes there will be no substantial legal battles because, for most manufacturers, the amount of royalties would be small in comparison to the stakes of a lawsuit.

Vande Sande also noted that the patent in question is basically for a token ring technology; thus, he dismisses buses as not relevant to the patent. Vande Sande feels his client's claims cover token rings in full, including both the slotted ring and the asynchronous token ring. Mr. Vande Sande also feels that the patent covers all aspects of token rings as defined in IEEE 802.5.

4.1.2. First Inklings -- Papers, Conferences

IBM's first public dalliance with token rings was embodied in a series of four papers IBM gave to IEEE 802 in March of 1982 and at a conference (IFIP) in Florence, Italy, in April 1982. The IEEE 802 presentation described a new architecture; the Florence presentation described some implementations of that architecture. In August of that same year, at an IBM users' group meeting, IBM presented a fifth paper, the subject of which was wiring.

A month later, IBM took the IEEE 802 papers and presented them at the IEEE Computer Society's semi-annual conference, COMPCON. (The five papers pared down to three for that Fall '82 COMPCON). At that meeting, IBM's vice-president of communications gave a speech in which he said these developments represented "...the most significant architectural announcements since the architecture of the IBM 360." This took place in the same month IBM published its first wiring book. IBM also presented a paper at an annual conference called the TI Mix Users' Group, on similar developments in architecture. The net effect of all this was that customers could finally start to believe that the token ring was indeed IBM's approach.

The three papers presented at the Fall '82 COMPCON were: "Ring Network Topology For Local Data Communications," by R. C. Dixon; "Token-Ring Local Area Networks: A Perspective," by J.D. Markov and N.C. Strole; and "A Token-Ring Architecture For Local Area Networks: An Update," by D.W. Andrews and G.D. Schultz.

The first paper discusses unique characteristics of ring topology that ideally suit it to data communications within a LAN; it covers the physical wiring of ring networks, the routing of data packets, mechanisms for ring networks, the routing of data packets, mechanisms for ring access control, and token access control. The second paper stresses the flexibility, reliability, and future expansibility of a ring topology, and looks into the physical wiring of a building, via the use of wiring concentrators and high-speed bridges for linking multiple rings. The Andrews-Schultz paper discusses a token-ring architecture that meets LAN requirements of a long-range wiring strategy for linking large numbers of office terminals, an integrated approach to handling data, voice, and facsimile applications, and a flexibility that can exploit high-tech advances in transmission media.

4.1.3. The Geneva Token-Ring Demo

One of the truest indicators that IBM was having a difficult time getting all the components of the Token-Ring Network in place was the fact that it publicly displayed a prototype of the LAN but once -- at Telecom '83 in Geneva, Switzerland, October 26 through November 1, 1983. And even then, IBM officials declined to reveal a date when the network would be commercially available.

Normally, the IBM Token-Ring Network is comprised of a series of star-shaped networks, each controlled from a wiring concentrator; a backbone ring connects these concentrators with each other and with a bridge through which one backbone can be interfaced to others. The backbone is comprised of two parallel, separate message paths. Yet the particular network displayed by IBM at Telecom '83 consisted of but one ring: a wiring concentrator connected to several terminals -- 3270 CRT terminals, 8775 display terminal, a 3275 front-end processor, and 8100 distributed processors -- within IBM's exhibit area.

In the demo, an attached device gained access to the network by changing the status of a perpetually circulating 1-bit token from "free" to "busy"; the token is in the header of a message frame, and this frame is then filled with all or part of the message itself.

Once the message frame is delivered to the receiving station, it is returned to the sending station -- thereby confirming delivery, and it is then removed from the network.

Digital voice and other kinds of synchronous traffic were transmitted with the assistance of a Synchronous Bandwidth Manager (SBM), a special station which, upon request, reserves in advance a series of message frames. The bits making up the synchronous message are then sent through the network in bursts, at regular intervals. These are buffered at the receiving point and then output.

By keeping the delay between bursts short, relative to the speed at which the buffer is emptied, a voice message can be transmitted in a state that sounds to the listener as if it is being spoken in real time.

IBM representatives distributed literature at the Telecom '83 booth which concurred with the contention expressed in a paper delivered a short while earlier by engineers at IBM's Zurich laboratories: that reliability is the major advantage of the new token ring network.

Both sources attributed this reliability to the system's star-shaped subnetwork architecture, wherein each node is connected by separate send/receive links to the wiring concentrator: this permits bypassing any malfunctioning node without disconnecting the others.

Further, the double message path within the backbone ring (that interconnects the wiring concentrators) allows a faulty backbone link or a faulty concentrator to be bypassed without bringing down the entire system.

IBM further maintained that the constantly circulating token offers a means of monitoring the system continuously; this, coupled with the relays with the wiring concentrators which enable automatic bypass of defective components, provides excellent fault isolation and recovery, according to IBM.

The new network, IBM said at that time, can be deployed on coaxial cable, fiber-optic cable, twisted copper pairs, or any mix of these media. The cable used in the Telecom '83 display network included two shielded twisted pairs -- for the network -- plus four wires for telephone communication.

Company representatives at Telecom '83 hastened to point out that though the Token-Ring Network transmits digital voice systems, it is nonetheless not a private automated branch exchange, and that more work would be required on it before it can handle real-time voice conversations as well as telephone networks can.

IBM also disclosed at Telecom '83 that scientists at its Zurich laboratory were working on a 16-Mbps version of the network; the model shown at Geneva had a nominal speed of 4 Mbps -- though the actual throughput is less than that.

4.1.4. The IBM Cabling System

On May 8, 1984, IBM released a cabling system intended for prospective token-ring customers, a system whereby parties could wire their buildings in anticipation of the network. The IBM Cabling System, as it is called, permits traditional types of connections to be made between devices with a common cable consisting of twisted pairs of copper conductors. The cabling system connects wall outlets in offices to wiring closets; a star-wired system, the cable fans out from each closet to special outlets. The outlets are installed in walls to connect to data devices and, optionally, telephones. Most currently available IBM products -- including workstations and small and intermediate computers -- can be plugged into the outlets.

Another major IBM selling point for the cabling system -- in addition to its accommodation of the Token-Ring -- is its cost-savings capability, in terms of the simplification it offers for installing or moving computer devices within a building. The IBM Cabling System is permanently wired, just as phone lines are run in buildings today, with connections made to outlet plates in office walls. By eliminating most of the expense of moving a workstation, IBM estimates the cabling can pay for itself within a matter of a few years. The single cabling system can be used instead of coaxial, twinaxial, twisted-pair, and other special cabling.

With the IBM Cabling System, each office wall outlet is connected by standard cable to a distribution panel (see Figure I-31) located in a wiring closet. A panel can accept up to 64 cables from different devices, and any two devices can easily be connected using patch cables at the closets where the cables converge. If an IBM 3270 Personal Computer, for example, is moved from one office to another, it is simply plugged into the wall outlet in the new office, and the patch cable in the wiring closet is reconnected. Wiring closets in the same building or different buildings on a campus can be connected with either twisted-pair or optical fiber cable.

The basic elements of the IBM Cabling System announced in 1984 were as follow.

Transmission Cables:

- Type 1: For use between faceplates in work areas and wiring closets, or between two wiring closets in the same or different buildings; contains two balanced, twisted pairs for data transmission.

- Type 2: For use between faceplates in work areas and wiring closets in the same building; contains two twisted pairs for data transmission and four twisted pairs for voice transmission.

- Type 5: For use between wiring closets in the same or different buildings; contains two optical fibers for data transmission.

- Type 6: For use as patch cables in wiring closets; contains two twisted pairs for data transmission. (The various cable types are show in Figure I-32).

Faceplates and Connectors:

- Faceplates for mounting on electrical outlet boxes installed in work areas.

- Data Connector and telephone connector for installation in the faceplate; data connector allows termination of two data-grade twisted pairs. A telephone jack connector allows termination of three voice-grade twisted pairs.

Distribution Panel:

- Cable junction panel (see Figure I-33) mounted on a rack in a wiring closet; each panel allows connection of up to 64 data cables.

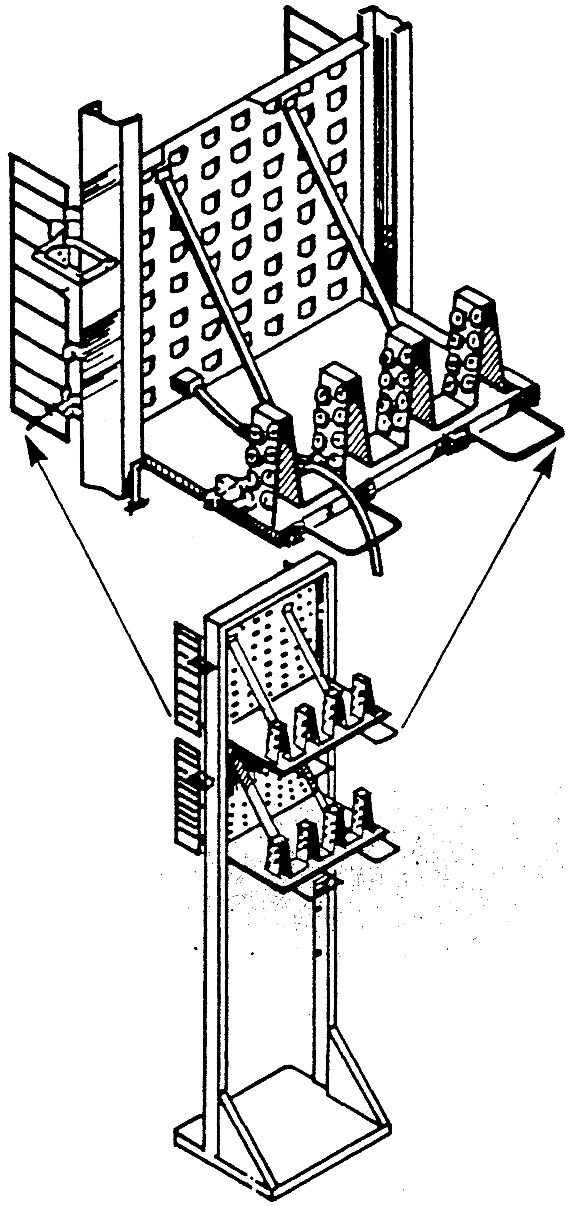

Figure I-31: Distribution Panel for IBM Cabling System

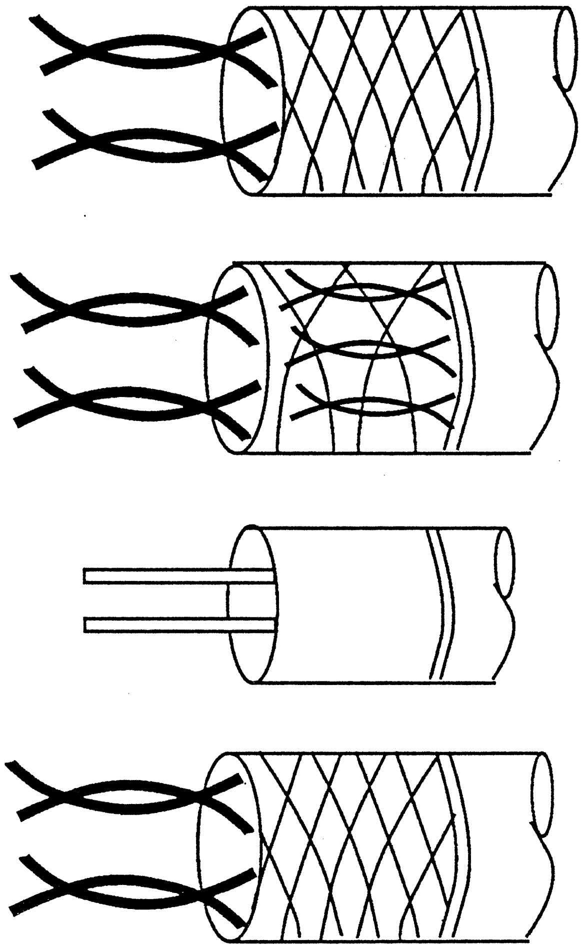

Figure I-32: Original IBM Cabling System Cable Types

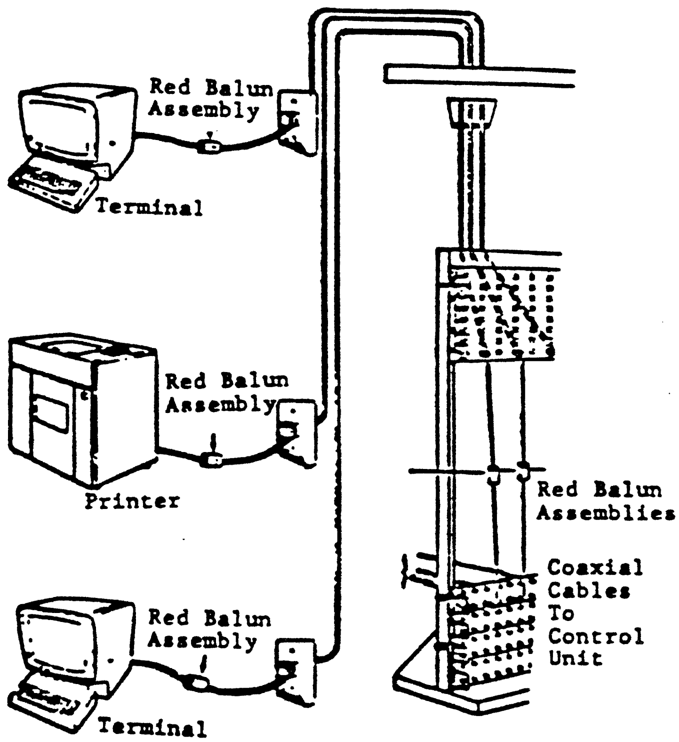

Figure I-33: 3270 Attachment, Cable Junction Panel

- Patch cable for making panel cross connections.

Device Attachment Cable and Accessories:

- For use in attaching a device to the cabling system; the attachment cable is terminated with a workstation connector at one end and a cabling system connector at the other end; it includes an impedance matching device, when required.

4.1.5. TI And Ring Chips -- Problems, Delays

Texas Instruments, Inc., has played a very instrumental role in the history of the IBM Token-Ring Network: it was, after all, TI's inability to deliver ring chips to IBM on schedule that delayed IBM's development of the Token-Ring, and its secrecy on the matter that lent an extensive air of uncertainty to IBM's Token-Ring progress and strategy. It is necessary to back-track in time to pick up the initial involvement of TI in the token-ring history, and to make some conjectures as to how IBM saw fit to approach TI for assistance in product development.

In early 1981, the formation of IEEE 802 became a source of grave concern for IBM: IEEE 802 was proceeding along a course centered about Ethernet, yet IBM had expertise on token rings.

In late 1980, IBM had initiated a comprehensive study of LAN technology, with the intent of developing a corporate strategy based upon its findings. In or near spring of 1981, the strategy began to unfold; nonetheless, IBM was completely taken by surprise by IEEE 802 and Xerox/DEC/Intel standards efforts. Not wanting to be trapped again, as it was earlier in the personal computer arena, IBM decided that it must strongly influence or even dominate the standards efforts. Yet it did not fully understand the market or technology, and further, it did not understand Ethernet -- IBM understood rings.

IBM thus squared off with 802 after coming to the conclusion that Ethernet was clumsy for a number of reasons: wiring complexity, transceiver complexity, and lack of guaranteed responses. IBM selected the token-ring over Ethernet because it gives a predictable bounded time response necessary to provide for connections of synchronous equipments. When this conclusion was reached, in mid-1981, IBM decided to forge ahead full-tilt on token-based systems; however, its primary expertise in this area lay in Zurich. The first steps to be taken, then, were to develop this expertise in the U.S., particularly in Raleigh, and to stop 802 from adopting Ethernet.

Eventually, IBM negotiated a compromise with IEEE 802 -- acceptance of the concept of a family of standards -- and IBM then took over the ring subcommittee.

As 1982 approached, in Raleigh, the groups were being put into place -- perhaps as many as two to three hundred employees. Ethernet was firmly rejected, and the corporate strategy began to take shape: a single LAN for all products.

Finally, as 1981 drew to a close, at some time IBM negotiated a contract with TI (at first code-named Skywalker) to develop a set of chips for its token ring system, which was revealed to the IEEE in March of 1982 and divulged to the public in The Wall Street Journal in September of 1982.

Once IBM had decided upon its strategy, for political reasons it chose to move ahead at as fast a pace as possible. In effect, IBM seemed forced to bring out a ring on a time schedule matching or preceding the "competition" of Ethernet: indeed, IBM seemed to feel that it should have prototypes in 1982. To perform at this level, IBM decided to have a semiconductor house build the essential chips: TI. Clearly, if a token system could be developed prior to Ethernet and, at the same time, be an approved IEEE standard, it might dominate the industry.

TI provided IBM with a credible approach: it understood the requirements for a chip set, and it had a viable architecture based upon existing parts. A problem unforeseen at the start of the IBM-TI liaison was that the key TI players, who were from the microprocessor-oriented groups, sold IBM a solution based on the 16-bit TI microprocessor architecture; the solution designed by IBM required a great deal of complexity for the interface chip and a generous amount of space for all necessary chip functions -- more complexity and room than TI's microprocessor expertise had afforded it.

Essentially, TI was under contract with IBM to develop two basic parts: an analog-like chip to accommodate the physical interface, and the token engine.

The interface chip was to perform a variety of functions; however, its most essential feature was its variable speed characteristics -- 4, 10, 16, 32, and, possibly, 64 Mbps. Flexibility, then, is the key issue. In that it also must interface to a variety of media, this chip was far more complex than anticipated, and TI had problems fitting the chip onto the die sizes necessary to boost yields while at the same still using standard cell building blocks for the chip.

The token engine had to be capable of processing complex link protocols while buffering host and ring interfaces; the chip thus required high performance in its processing sequencing, i.e., more processing capabilities then originally anticipated. Again, TI had problems fitting all the functions into the chip's performance levels while still using the standard microprocessor chip as the building block.

When these chips were put in cascade, the overall performance was quite low. Simply summarized, the problem was this: TI underestimated the complexity of the interface chip and consequently did not design adequate space onto the chip for all required functions. Early test chips underscored this problem, but the problem was never resolved; further, TI also underestimated the token protocol complexity, and the microprocessor was undersized. This caused a delay of nearly eighteen months in the IBM token ring project and ultimately forced IBM to develop its own ring chips.

As a sort of footnote, it should be added that because of its involvement with IBM on development of the infamous rejected chip set, TI is currently the only "third-party" vendor to offer an IEEE 802.2/802.5 chip set. The TI chip set -- the TMS380 LAN Adapter Chipset -- was formally announced on October 15, 1985, the same day as IBM's Token Ring announcement. Because of its work with IBM on the chip set, TI was privy to information not available to the other vendors. Thus a system developed with TI's chip set will be fully compatible with IBM-developed Token-Ring adapter kits. In fact, within 1-1/2 months of its announcement, over 200 vendors requested TI chip set evaluation kits.

4.1.6. IBM's Token-Ring Strategy

4.1.6.1. Overall IBM Strategy

It is important to note that IBM is not obsessively interested in LANs. While the corporation feels it needs a LAN technology to evolve products that will compete in the marketplace, the basic interfaces themselves are of little interest to IBM: it is the equipment that connects to LANs that is of extreme importance to IBM. Hence, from a strategic viewpoint, IBM's involvement in LANs is strictly a tactical maneuver to ensure it a firm foothold in future systems markets.

Thus, when discussing IBM strategy, one must differentiate between the need to control equipment versus the need to provide equipment. IBM will provide local networks, but it does not intend to contol local network equipment, and it will -- in a limited context -- allow an open system. Nonetheless, as more sophisticated devices connect to the LAN and users proceed into the higher-level protocols, IBM will then become more concerned about control.

With this consideration in mind, the year 1982 can be seen as the year of the IBM warning: IBM told everyone in 1982 that changes were coming -- it even defined the exact form and sequence. Figure I-34 and Table I-1 illustrate the IBM strategy.

First is the wiring -- the idea is to get an installed base

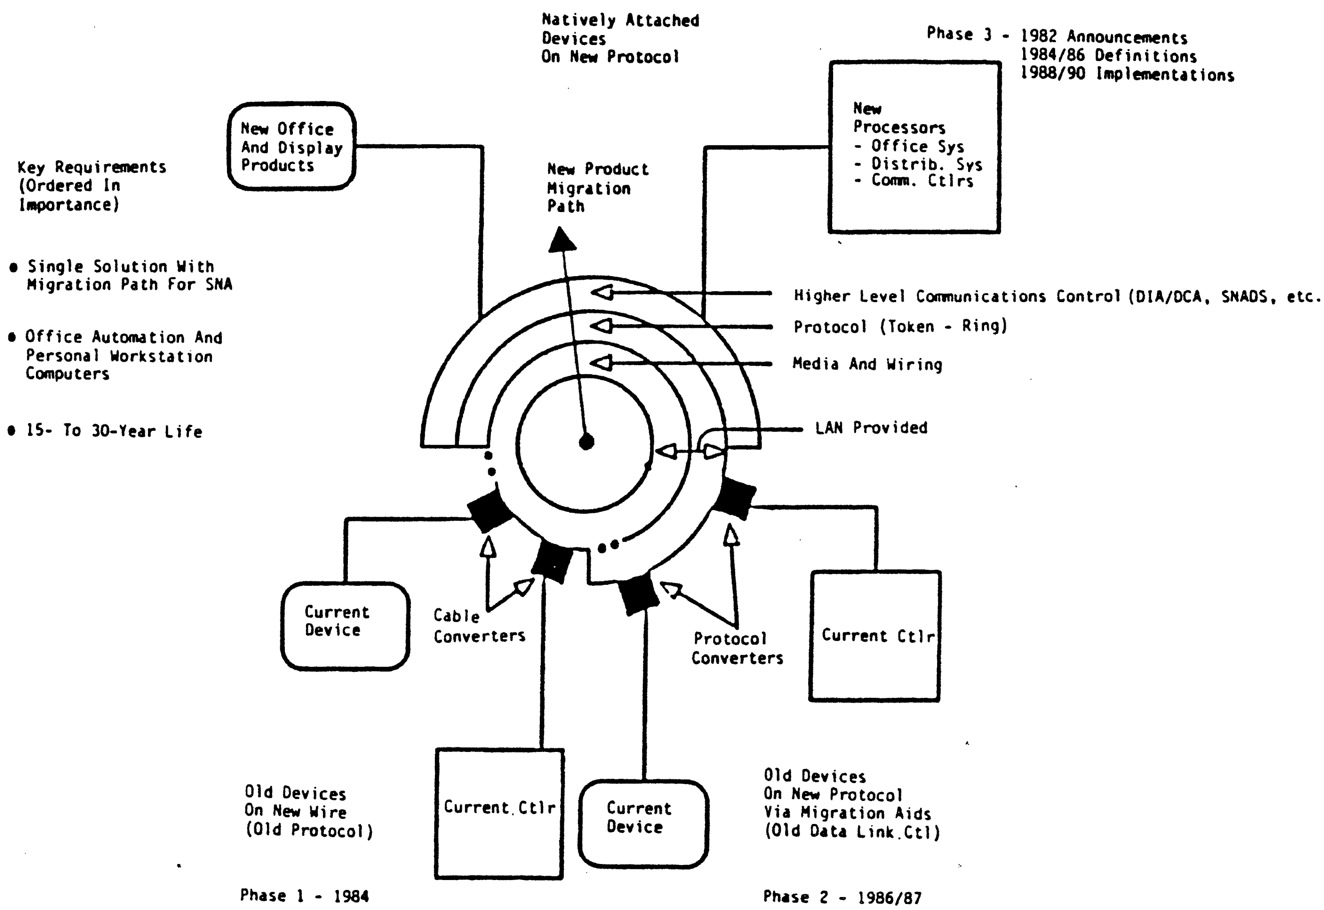

Figure I-34: IBM's Overall Strategy

LAN Type	Advantages	Disadvantages	Project State
Prong 1 - SNA Support Direct From IBM			
Token Ring	SNA Open Architecture	Expensive Wiring Lack Of 327x-Type Support	PCs Are First Products To Link To Token-Ring
Prong 2 - Arbitrary LANs			
Integrated Voice/Data	ROLM Installed Base Available Now Gateways Available	Low Speed Asynchronous Traffic	CBX Now Available
Ethernet	Multiple Vendors Easy To Get Parts	Not Taken	S/1 Development Groups Are Heavy Buyers
PC Cluster	Simple Cheap Available	Poor Design No Security Limited Volumes Per User Not Open Architecture Slow	Now Available
PC Network	Simple Cheap Large Numbers Of PCs Fast	Expensive Head-End Not Open Architecture	Now Available
Token Bus	Response Time Guarantees Broadband GM MAP Support	Asynchronous Traffic	From Factory/ Industrial Oriented Groups

Prong 3 Gateways And Prong 4 SNA Software - No fully Defined Strategy Yet

Table I-1: Prongs in IBM Strategy

of wiring in place, because with nothing to attach to one has nothing to transmit on. Second is the token protocol -- bring along the overall transport system so as to provide a basis for the long-term product, but with an evolution process clearly in mind; however, by leaking the strategy details as the wiring goes into facilities, the progress of other vendors is frozen. Lastly, new software -- but because of the time required for software development, the announcements of some software pieces can occur at the same time as -- or even before -- the hardware; and thus is seen the introduction of DIA/DCA, SNADs, and LU6.2. The importance to IBM of the token and SNA, as parts of this overall strategy, is discussed below.

4.1.6.2. The Openness Of The Token

The importance of the token -- at least with IBM -- is that all features of its architecture, ranging from the link layer and the access layer up to the connection of SNA, will be open to any product groups within IBM that want to connect to the token.

The impact of the IBM corporate direction -- the token -- is that it allows all the IBM divisions to set various-length strategic and tactical options in place within their departments, to live within the mainstream software and hardware departments of the company without having to be bothered with the details of setting up their own strategic communications rings.

The impact of the token is long-term: it is to create and put in place wiring schemes and access methods that will allow all IBM products to continue to enter into large markets. Thus, product developers within IBM have a 10-year view and up to a 30-year window within which to push their products, to develop and begin to migrate into a new strategy.

It is IBM's intention that the token will become the standard in the office-of-the-future. If IBM did not have as a corporate direction the Token-Ring, the proliferation of networks from IBM would hamstring the corporation with redundancy, and incompatibility, and confusion.

4.1.6.3. SNA And The Token-Ring

SNA is the key to the IBM LAN strategy. The company should be viewed as though it is broken into two parts: people who use SNA -- and thus, by default, the token ring -- and people who use other local network technologies such as the PC Cluster and who, by default, do not get IBM supported access to the SNA protocols.

In general, SNA is the description of the rules that enable IBM's customers to transmit and receive information through their computer networks.

SNA may also be viewed as three distinct but related entities: a specification, a plan for structuring a network, and a set of products.

First, SNA is a specification governing the design of IBM products that are to communicate with one another in a network; it is called an architecture because it specifies the operating relationships of those products as part of a system.

Second, SNA provides a coherent structure that enables users to establish and manage their networks and, in response to new requirements and technologies, to change or expand them.

Third, SNA may be viewed as a set of products: combinations of hardware and programming designed in accordance with the specifications of SNA. In addition to a large number of computer terminals for both specific industries and general applications, IBM's SNA product line includes host processors, communication controllers and adapters, modems, and data encryption units.

The SNA product line also includes a variety of programs and programming subsystems. Telecommunication access methods, network management programs, distributed applications programming, and the network control program are examples.

According to Edward H. Sussenguth, one of the early developers of SNA and now an IBM Fellow at Research Triangle Park: "One of the primary motivations behind the creation of SNA in the early 1970s was the advent of distributed data processing and so-called "intelligent" terminals. Previously, all application programming and instructions for device control resided exclusively in the host computer.

"But now, for the first time, the customer's applications, as well as device control, could be performed in any one of several places within a network -- at the host, at a local controller, or even at the terminal itself. Obviously, there was a need to create definitions and protocols, or rules, to control such a system," Sussenguth said.

"We plainly needed a blueprint, something like SNA, to unify IBM's communication products ...a single, extendable architecture to allow all our products to connect via the same protocol.

"Since SNA was designed to be evolutionary, it allowed for the non-disruptive introduction of new system capabilities and individual products for years to come. SNA enabled our customers to protect their investment in products by permitting functional extensions and enhancements without making current models obsolete."

One indicator of the clout SNA commands within IBM is the fact that in 1984, when the Entry Systems Division (PC Group) at Boca Raton was embarrassing the Token-Ring Group at Raleigh, NC, by beating it to market with new products (e.g., the PC Cluster and the PC Network -- the latter was made for the Entry Systems Division by Sytek, Inc.), these products were not accepted as products by the Mainframe Group -- and it is the Mainframe Group which controls access to the SNA protocol.

With IBM's announcement of DIA/DCA (Document Interchange Architecture/Document Content Architecture), SNADS (SNA Distributed Services), LU6.2 (Logical Unit 6.2), and the Token-Ring LAN, it is anticipated here that SNA will now undergo a 10-year migration from being a hierarchical system to being a decentralized system based on hierarchically structured LANs.

4.1.7. Token-Ring Transmission Rates Paper

The final event of noteworthy importance prior to the official October 1985 announcement of the Token-Ring LAN was the presentation -- at the IEEE Computer Society-sponsored Fifth International Conference on Distributed Computing Systems in Denver, CO, May 13-17, 1985 -- of a paper entitled, "File Transfer In Local-Area Networks: A Performance Study."

Co-authored by Messrs. B. Meister, P. Janson, and L. Svobodova (all of whom work at IBM's Zurich Research Laboratory), the paper describes a simulation of a 4-Mbps token-ring, with particular emphasis on the throughput which can be achieved between two stations on a LAN, measured either from memory-to-memory of the two stations, or from disk-to-disk, where the disks are attached to the two stations. One station can be conceived of as being a file server, with the other being a workstation, or it can be assumed that there are two processors on the system.

Several important points were made in this paper regarding the issue of simulation. It begins by noting that one of its main objectives was to measure the performance benefits of a LAN. In that there are not many extant LAN applications which actually use a data transfer rate in the Mbps range, the authors felt it would be interesting to see what the performance limitations of the local network would be if one were to test the performance bounds (assuming one had applications). (The authors did note that by modeling the limitations of the local network, one is, in essence, removing any limits that might be placed by the hardware and software context in which real applications would perform.) In particular, what the investigators were attempting to do was to see what impact the connection-oriented link protocol would have on data transfer throughput at the application level, and thus the simulation was constructed.

The article assumes that the reader is familiar with OSI terminology, and so it is further assumed that a connection-type protocol would have an establishment, transfer, and release phase. Therefore, the connection protocol makes a logical channel, in contrast to the connection-less protocol where each data unit -- generally referred to as a datagram -- is provided to the network as a self-contained, unrelated unit with respect to any unit sent before or after it.

IBM's view -- via its three researchers -- is that simplicity and low overhead are arguments in favor of connection-less protocols. However, connection-type protocols are necessary due to three primary issues: loss of packets, compatibility with existing systems, and the fact that LANs are not necessarily stand-alone items and need to be viewed in a more global context. With these three factors weighing against the simplicity and low overhead advantages of a connection-less protocol, one is prompted to ask, what is the penalty for having a connection-oriented protocol? This is, in fact, the issue the paper addresses.

The paper goes into great detail in describing the experiments conducted, the means whereby the data is derived, the kinds of computers used, the simulation models, etc., before addressing the central issue of performance. The key points made in the article are based on several sets of figures: Figures 4 through 6 (in the paper) illustrate what happens in five different cases of memory-to-memory transfers; Figures 7 through 9 illustrate what happens with disk-to-disk transfers; and Figures 10 through 12 depict the results when one has disk-to-disk transfers with parallel disk access. For each of the figures or for each of the simulation runs, IBM, in essence, looked at three different processor speeds: .25 Mips, 1 Mips, and 10 Mips.

In perusing these figures and the best- to worst-case results, several surprising statistics emerge. Figure 4 of the paper illustrates what happens with a quarter-of-a-Mips processor using memory-to-memory transfers: it shows that the throughput basically ranges between 250 Kbps and 80 Kbps for an optimum-size packet mechanism on a 4-Mbyte ring. The 10-Mbps processor goes all the way up to nearly 2 Mbps, and it begins at around 800 Kbps.

In order to present a more convenient comparison of the performance figures for these processes, the IBM data is redrawn below as three separate figures (Figures I-35, I-36, and I-37) featuring envelopes of performance. The problem with the manner in which IBM laid out the data is that one tends to not see the information relative to the processor speed.

In the first redrawn figure (Figure I-35), the processor speeds are matched, because the information is going between memory; the only bottleneck in the system is the transmission subsystem. This means that the slow-down from the 4-Mbps range down to as low as 80-Kbps is due not to anything associated with the processor, but, essentially, only with the speed of the LAN; these, then, are maximum theoretical speeds on the LAN.

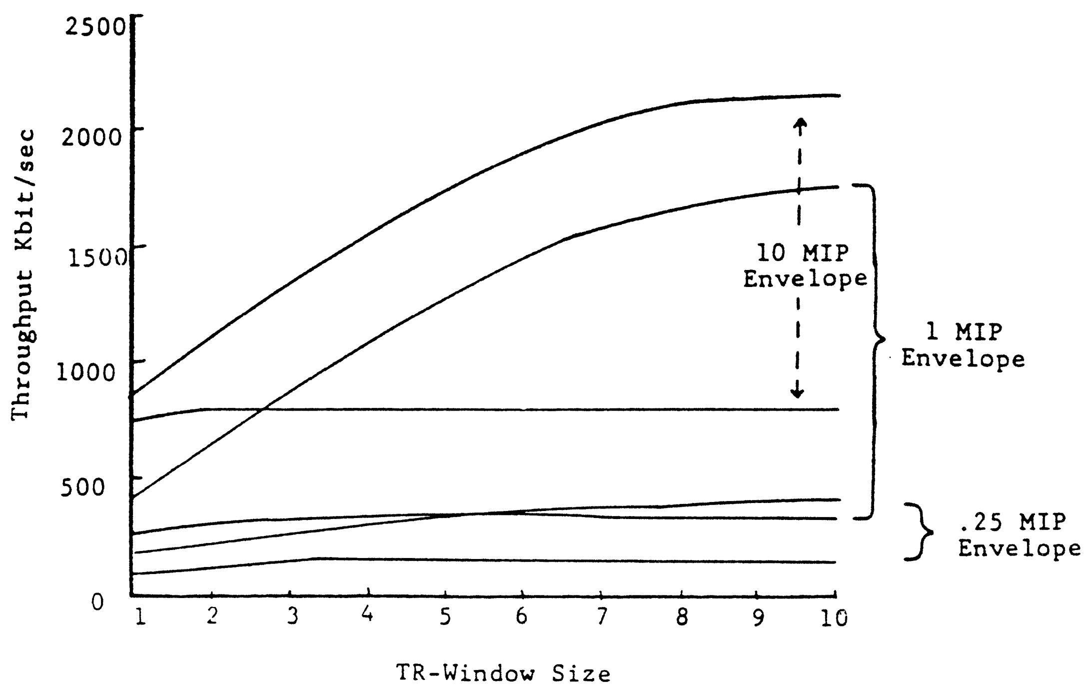

Figure I-35: 0.25-, 1-, 10-Mips Processors, Memory-to-Memory

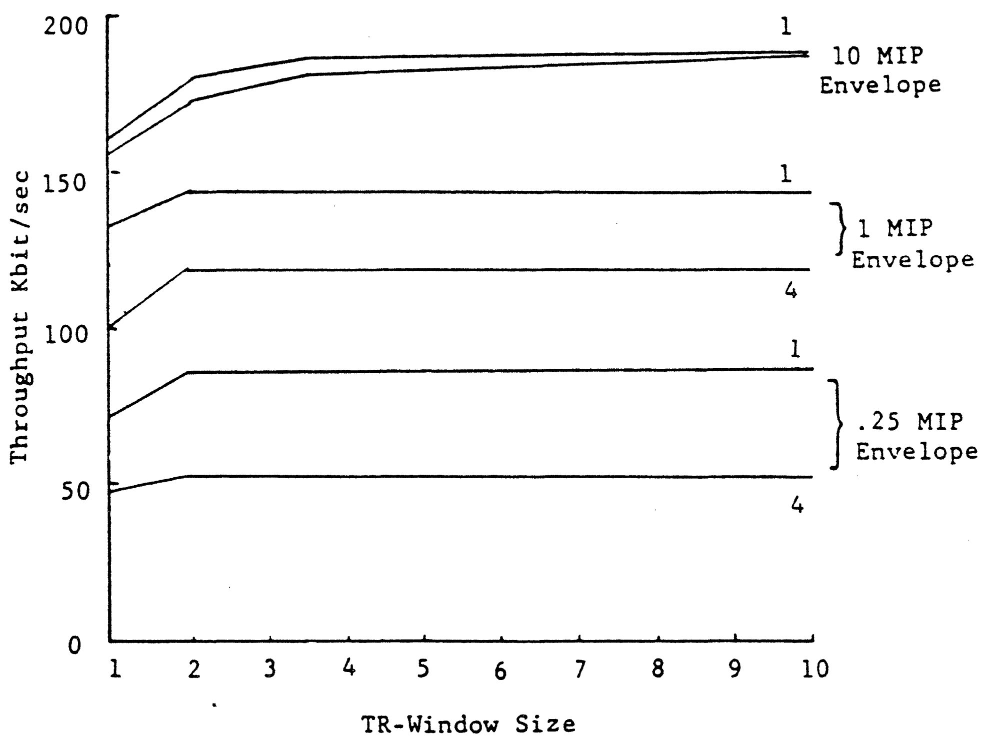

Figure I-36: 0.25-, 1-, 10-Mips Processors, Disk-to-Disk, Disk Access 10 ms Per Frame

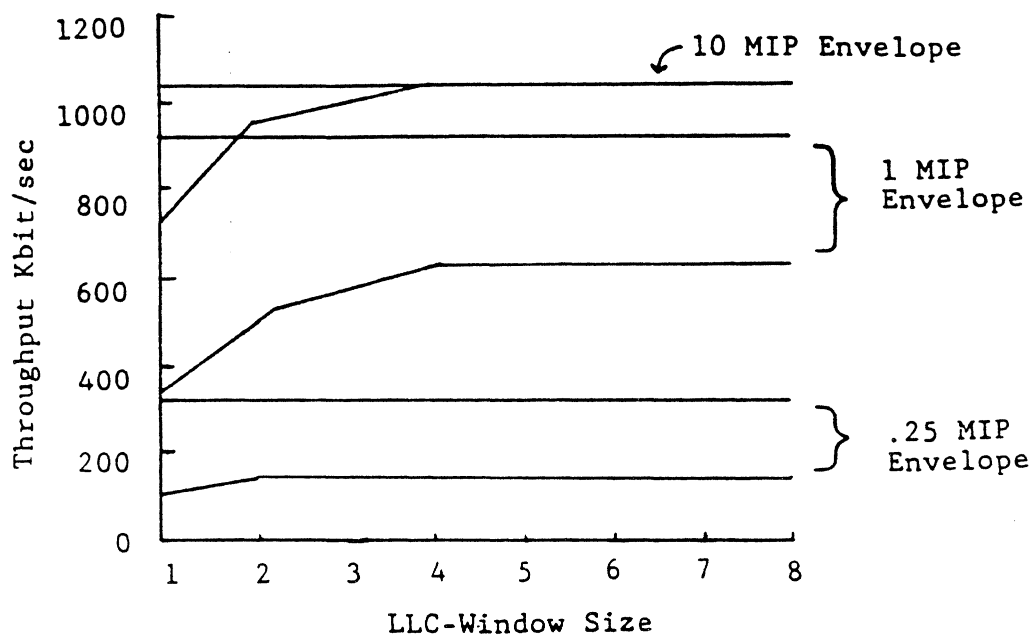

Figure I-37: 0.25-, 1-, 10-Mips Processor, TR-Window Size=8, Parallel Disk Access

Figure I-36 is a composite of IBM's Figures 7 through 9; it depicts the performance attainable due basically to what is the performance associated with disk-to-disk transfers. Here the system seems to have a low of 50 Kbps up to a peak of about 200 Kbps, and these are with balanced processors. This falling performance appears to be due to disk access, so when one then goes into a system in which there are parallel disk accesses (which is shown in composite Figure I-37), one sees that the throughput goes all the way up to 1 Mbps.

It is important to note the key IBM conclusion drawn in this paper: "...it appears that in the absence of packet losses, the use of a connection-oriented protocol at the link control level on a LAN may represent a performance penalty for file transfers as high as 50%, given a straightforward software implementation." IBM notes that if one had either smarter network interface units or more careful software development, one might be able to drop this overhead from 50% into the range of 10-20%.

While the paper states that the transmission rate on a 4-Mbps Token-Ring LAN will not likely exceed 200 Kbps, it is likely that IBM will follow with a 16-Mbps version; to further play the seer, one can also expect, shortly thereafter, an even higher performance 32- or 64-Mbps protocol board to ensure that the processor is able to see that throughput. If end-to-end speed is low, IBM will build a faster link and add hardware when necessary.

IBM has a very simple procedure in mind here: what it intends to do is simply encapsulate an SDLC frame into one of the MAC frames on the Token-Ring. That is why gateways are seen as so simple and why the packets are so simple: it leaves the software easy to make.

IEEE 802 specifies connection and connectionless type protocols, yet IBM essentially provides a connection-oriented protocol. In one sense, this really puts the transport layer into the frame or at the 802 link level, and it makes Levels 3 and 4 null from an IBM standpoint.

IBM drops the SDLC frame into the physical and data link mechanism, where it is shipped across a connection; thus the system hardware at the ends bypasses the network and transport layers. SNA will provide the correct intelligence to perform the upper-level functions.

The implications of IBM's procedure are ease of evolution and ease of fielding the equipment for the company in the short term; the disadvantage is that it will lead to more of the same IBM equipment consumers are seeing now. IBM is oriented to support of host-based systems, and thus its LANs will be, too.

4.1.8. Announcement Of The IBM Token-Ring LAN

On October 15, 1985, IBM finally, officially announced its Token-Ring Network, as part of a massive series of new work group connectivity products announced on that same day. Introduced initially as a network for personal computers, when finally equipped with interfaces for all the other IBM products the Token-Ring LAN will provide complete local communications at an IBM-outfitted site.

4.1.8.1. Token-Ring Network Overview

The IBM Token-Ring LAN is a 4-Mbps baseband PC network that conforms to both IEEE 802.5 and ECMA 89 standards. It can use either the IBM Cabling System or IBM's later announced Cabling System specification for telephone twisted-pair wiring. If the data-grade medium is used, up to 260 coaxial, twinaxial, loop, and token-ring devices can be attached to the Token-Ring; if the telephone medium is used, up to 72 coaxial and token ring devices can be attached to the network.

A token-ring access protocol is employed for network traffic control. The various devices to be attached to the network are cabled together through a network access unit to form a logical ring. Access to the shared ring is controlled by a token. Each networked device regenerates the signal as information is passed around the ring. The access protocol and baseband signaling are implemented by a network adapter in each attached device.

Hardware components of the Token-Ring Network include an adapter card with logical link control for the attachment of the IBM PC to the network, a multistation access unit, and cabling. Software components include a program that provides an IBM PC Network NETBIOS interface to the network, an SNA programming interface for the PC, a program that links the PC Network and the Token-Ring Network, and a communications server.

4.1.8.2. Token-Ring Conceptual Design

The protocols and formats employed by the IBM Token-Ring medium access control sublayer, the physical layer, and the means of attachment to the token-ring physical medium are those specified by the IEEE 802.5 token-passing ring standard.

The IBM Token-Ring LAN comprises a set of stations serially connected by a transmission medium. Information is transferred sequentially, bit by bit, from one active station to the next. Each station typically regenerates and repeats each bit and serves as the means for attaching one or more devices to the ring for the purpose of communicating with other devices on the network. A given station (the one that has access to the medium) transfers information onto the ring, whereupon the information circulates from one station to the next. The addressed destination station(s) "copies" the information as it passes. Finally, the station which transmitted the information removes the information from the ring.

A station gains the right to transmit its information onto the medium when it detects a token passing along the medium. (The token is a control signal comprised of a unique signaling sequence that circulates on the medium following each information transfer.) Any station, upon detection of such a token, may capture the token by modifying it to a start-of-frame sequence and appending appropriate control and status fields, address fields, information field, frame check sequence, and the end-of-frame sequence. At the completion of its information transfer and after appropriate checking for proper operation, the station initiates a new token, which provides other stations the opportunity to gain access to the ring.

A token-holding timer (THT) controls the maximum period of time a station may occupy the medium before passing the token. Multiple levels of priority are available for independent and dynamic assignment, depending upon the relative class of service required for any given message, e.g., synchronous (real-time voice), asynchronous (interactive), or immediate (network recovery). The allocation of priorities is by mutual agreement among users of the network.

IBM's Token-Ring Network incorporates error detection and recovery mechanisms to restore network operation in the event that transmission errors or medium transients (e.g., those resulting from station insertion or removal) cause the access method to deviate from normal operation. These detection and recovery procedures use a network-monitoring function that is performed in a specific station, with back-up capability in all other stations which are attached to the ring.

4.1.8.2.1. Token Ring Formats

Token rings employ two basic formats: tokens and frames. Figures I-38 and I-39 depict the formats of the fields in the sequence they are transmitted on the medium, with the left-most bit or symbol transmitted first.

Those processes which require comparison of fields or bits perform that comparison upon those fields or bits as depicted, with the left-most bit or symbol compared first and, for the purpose of comparison, considered most significant.

The token is the means whereby the right to transmit (as opposed to the normal process of repeating) is passed from one station to another.

The Token Format is:

SD	AC	ED

where

SD = Starting Delimiter (1 octet)
AC = Access Control (1 octet)
ED = Ending Delimiter (1 octet)

Figure I-38: Token Format

The Frame Format is:

<-- SFS --> < ------- F C S Coverage -----> <-- EFS -->

SD	AC	FC	DA	SA	INFO	FCS	ED	FS

where

SFS = Start of Frame Sequence
SD = Starting Delimiter (1 octet)
AC = Access Control (1 octet)
FC = Frame Control (1 octet)
DA = Destination Address (2 or 6 octets)
SA = Source Address (2 or 6 octets)
INFO = Information (0 or more octets)
FCS = Frame Check Sequence (4 octets)
EFS = End of Frame Sequence
ED = Ending Delimiter (1 octet)
FS = Frame Status (1 octet)

Figure I-39: Frame Format

The frame format is used for transmitting both Medium Access Control (MAC) and Logical Link Control (LLC) messages to the destination station(s); it may or may not have an information field.

4.1.8.2.2. Token Ring Protocol

Within the Medium Access Control (MAC) layer, access to the physical medium -- i.e., the ring -- is controlled via the passing of a token around the ring. This token gives the receiving station (i.e., a station downstream from the station passing the token) the opportunity to transmit a frame or a sequence of frames. Upon request for transmission of the Logical Link Control (LLC), Protocol Data Unit (PDU), or Network Management (NMT) PDU, MAC prefixes the PDU with the appropriate FC, DA, and SA fields and enqueues it to await the reception of a token that may be used for transmission.

Such a token has a priority less than or equal to the priority of the PDU(s) that is/are to be sent. Upon queuing the PDU for transmission and prior to receiving a usable token, in the instance where a frame or an unusable token is repeated on the ring the station requests a token of appropriate priority in the reservation (RRR) bits of the repeated AC field; upon receipt of a usable token, it is changed to a start of frame sequence by setting the token bit.

At this point, the station stops repeating the incoming signal and begins to transmit a frame; during this transmission, the FCS for the frame is accumulated and appended to the end of the information field.

Once transmission of the frame(s) has been completed, the station checks to see if the station's address has returned in the SA field, as indicated by the MA_FLAG; if it has not, the station transmits until the MA_FLAG is set, at which time the station transmits a token.

Upon transmission of the token, the station will remain in the transmit state until all of the frames that it originated are removed from the ring; this is done in order to avoid unnecessary recovery action that would be caused were a frame allowed to continuously circulate on the ring.

Stations, in addition to repeating the incoming signal stream, check it for frames which they are to copy or act upon. If the frame-type bits indicate a MAC frame, the control bits are interpreted by all the stations on the ring. Furthermore, if the frame's DA field corresponds to (matches) the station's individual address, relevant group address, or broadcast address, the FC, DA, SA, INFO, and FS fields are copied into a receive buffer and subsequently forwarded to the appropriate sublayer.

The proprietary bits (PPP) are the reservation bits (RRR) in the Access Control field which work together to match the service priority of the ring to the highest priority PDU that is ready for transmission on the ring. The current ring service priority is indicated by the priority bits in the Access Control field, which is circulated on the ring.

The priority mechanism functions in such a manner that "fairness" (i.e, equal access to the ring) is maintained for all stations within a priority level; this fairness is accomplished by having the same station that raised the service priority level of the ring return the ring to the original service priority.

When a station has a priority (a value greater than zero) PDU (or PDUs) ready to transmit, it requests a priority token, by changing the reservation bits (RRR) as the station repeats the Access Control field. If the priority level (PM) of the PDU that is ready for transmission is greater than the RRR bits, the station increases the value of RRR field to the value PM; if the value of the RRR bits is equal to or greater than PM, the reservation bits (RRR) are repeated unchanged.

Once a station has claimed the token, it transmits PDUs which are at or above the present ring service priority level until it has completely transmitted those PDUs, or until the transmission of another frame would not be completed before the token holding timer (THT) would expire. The priority of all the PDUs that are transmitted should be at the present ring service priority value, whereupon the station will then generate a new token for transmission on the ring.

When a hardware failure is detected in a token-ring LAN, its cause must be isolated to the proper "failure domain," so that appropriate recovery actions can be initiated. A failure domain consists of the station reporting the failure, the station upstream of the reporting station, and the ring medium between the two.

To perform accurate problem determination, all elements of the failure domain must be known at the time that a failure is detected -- and thus, at any given time, each station should know the identity of its upstream neighbor station. A process known as Neighbor Notification is used to obtain this identity information.

Neighbor Notification is founded upon the Address Recognized and Frame Copied bits (the A and C bits) of the Frame Status field, which are transmitted as zeros. If a station recognizes the destination address of the frame as one of its own, the station sets the A bits to 1 in the passing frame; if a station also copies the frame, then the C bits are also set to a 1.

When a frame is broadcast to all the stations on a ring, the first station downstream of the broadcaster will ascertain that the A and C bits are all zeros. Since a broadcast frame will have its destination address recognized by all of the stations on the ring, the first station downstream will, in particular, set the A bits to 1; therefore, all stations further downstream will not see the A and C bits as all zeros. This process continues in a circular, daisy-chained fashion, to inform every station of the identity of its upstream neighbor.

The monitor begins Neighbor Notification by broadcasting the Active Monitor Present (AMP) MAC frame. The station immediately downstream from it then resets its timer TSM, based on setting the AMP value in the Frame Control field; if possible, it then copies the broadcast AMP MAC frame and stores the upstream station's identity in an "Upstream Neighbor's Address" memory location, sets the A bits (and C bits if the frame was copied) of the passing frame to ones, and, at a suitable transmit opportunity, broadcasts a similar Standby Monitor Present (SMP) MAC frame.

One by one, each station receives an SMP frame with the A and C bits set to zeros, stores its upstream neighbor's address, and continues the process by broadcasting such a frame itself.

Since frame AMP must pass each station on a regular basis (the Active Monitor Present MAC frame sent by the monitor), the continuous transmission of tokens onto a ring can thereby be detected. In addition to the timer TAM in the active monitor, each standby station has a timer TSM that is reset each time an Active Monitor Present (AMP) MAC frame passes; if it expires, that standby monitor station begins transmitting Claim Token frames.

4.1.8.2.3. <u>The Physical Layer</u>

The data signaling rate of the Token-Ring LAN is 1 or 4 Mbps, with a tolerance of plus or minus 0.01%.

The physical layer recovers the symbol timing information inherent in the transitions between levels of the received signal, minimizing the phase jitter in this recovered timing signal to provide suitable timing at the data signaling rate for internal use and for the transmission of symbols on the ring. The rate at which symbols are transmitted is adjusted continuously in order to remain in phase with the receive signal. Under normal operating conditions, there is one station on the ring that is the active monitor; all other stations on the ring are frequency- and phase-locked to this station and extract timing from the received data by means of a phase-locked loop.

In order for the token to continuously circulate around the ring when all the stations are in repeat mode, the ring must have a latency (i.e., time, expressed in number of bits transmitted, for a signal element to proceed around the entire ring) of at least the number of bits in the token sequence, i.e., 24. In that the latency of the ring varies from one system to another and no a priori knowledge is available, a buffer of at least 24 bits is provided by the active monitor.

The source timing (or master) oscillator of the ring is supplied by the active monitor station; all other stations in the ring track the frequency and phase of the incoming signal it receives. While the mean data signaling rate around the ring is controlled by the active monitor station, segments of the ring can instantaneously operate at speeds slightly higher or lower than the frequency of the master oscillator: the cumulative effect of these variations in speed is sufficient to cause effective variations of up to plus or minus 3 bits in the latency of a ring that has been configured with a maximum number of stations (i.e., 250).

Nonetheless, unless the latency of the ring remains constant, bits will either be dropped (not retransmitted) as the latency of the ring decreases, or added as the latency increases. In order to maintain a constant ring latency, an elastic buffer with a length of 6 bits (12 signal elements) is added to the fixed 24-bit buffer; the resulting 30-bit buffer is initialized to 27 bits. If the received signal at the active monitor station is slightly faster than the master oscillator, the buffer will expand, as required, to 28, 29, or 30 bits to avoid dropping bits; if the received signal is slow, the buffer will contract to 26, 25, or 24 bits in order to avoid adding bits to the repeated bit stream.

4.1.8.3. Hardware Components Of The Token-Ring

4.1.8.3.1. The IBM Token-Ring Network PC Adapter Board

The central core of the IBM Token-Ring LAN, in its current configuration, is the IBM Token-Ring Network PC Adapter Board. The adapter contains a microprocessor operating under the control of microcode resident on the adapter. It transmits and receives data at 4 Mbps, using protocols that conform with IEEE 802.5 and ECMA 89 standards. An adapter handler support program that runs in PC memory (requires 7 Kbytes) provides data link control in accordance with the IEEE 802.2 standard.

As noted earlier in this section, IBM and Texas Instruments worked together on an adapter chip set that many believed would be at the heart of the new network. However, IBM ultimately decided to use its own components, which appear to employ a denser and more sophisticated technology than the jointly developed chip set. (Texas Instruments itself introduced the TMS380 chip set, shortly after the IBM Token-Ring announcement.)

Reliability, availability, and serviceability functions are built into the Token-Ring adapter and microcode. Diagnostics invoked during adapter initialization verify the adapter operation and examine the cabling to the access unit. The adapter detects permanent errors -- e.g., the loss of a receive signal -- and generates a notification signal to initiate automatic network recovery. Recoverable errors -- such as bit errors in the transmitted message -- are detected by the adapter and then reported to a ring diagnostic program.

An attachment cable is necessary to connect the adapter to the network cabling system. The IBM Token-Ring Network PC

Adapter Cable is a 2.4-meter (8-foot) cable for attaching the adapter to data-grade media; the Type 3 Media Filter includes a 2.4-meter cable to connect the adapter to Type 3 telephone media. An optional advanced diagnostics program, the IBM Token-Ring Network PC Adapter Hardware Maintenance and Service, aids in isolating faults in the adapter, the adapter attachment cable, and the cable between the work area and the wiring closet.

4.1.8.3.2. Token-Ring Multistation Access Unit

The IBM Token-Ring Network Multistation Access Unit is a passive access unit that permits up to eight devices to be attached to the network to form a logical ring. Cables from the various devices are simply plugged into the unit. Access units can be connected to each other to expand the size of the network (see Figure I-40). Also, an attached device may be bypassed if the number of soft errors exceeds an established minimum, or if there is a wire or adapter error.

The access unit may be installed in a standard 19-inch rack (not available from IBM) or in a component housing available as an option for installing an access unit on a wall or tabletop. A cable bracket is included with the access unit for organizing and identifying the cables attached to a rack-mounted access unit.

4.1.8.4. Software Components Of The Token Ring

4.1.8.4.1. The IBM Token-Ring Network NETBIOS Program

Application programs written to the NETBIOS high-level programming interface for IBM PCs can operate on both the Token-Ring Network and the IBM PC Network: in the PC Network, NETBIOS runs in microcode, whereas in the Token-Ring LAN it runs in PC memory and requires 46 Kbytes. Application program requests for NETBIOS functions are translated into token-ring protocol requests which are then passed to the Token-Ring Network PC Adapter for execution.

The IBM PC Network SNA 3270 Emulation Program, Asynchronous Communications Server Program, and Token-Ring Network/PC Network Interconnect Program all use the NETBIOS programming interface and have been tested on the Token-Ring. Programs that operate in the IBM PC Network environment using the server functions of the IBM PC Network Program should also operate in the IBM Token-Ring environment using the server functions of the IBM PC Network Program. The IBM PC Network Program, using an IBM PC with the IBM 3812 Pageprinter as a print server, has been successfully tested on the IBM Token-Ring LAN.

On April 2, 1986, among a series of announcements introducing the IBM PC Convertible was one that described an

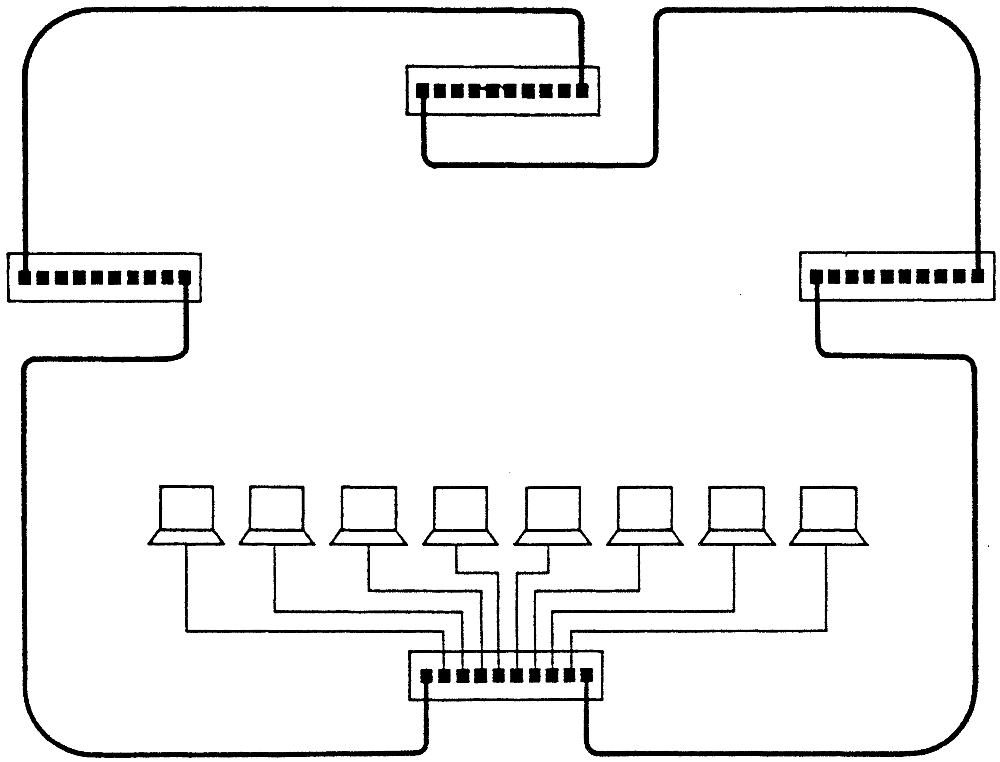

Figure I-40: IBM Token-Ring Network Multistation Access Unit Configuration enhanced edition of the PC Network Program, renamed the IBM PC Local Area Network Program Version 1.1. The name change, according to IBM, reflected the support for IBM PCs on both the IBM PC Network and the IBM Token-Ring Network.

The program, which requires DOS 3.2, enables IBM PC users to share applications, common databases, and printers on the network. The enhanced version of the program works with the IBM Personal Computer TopView Program Version 1.10. The user can send messages through the network using the PC Local Area Network Program while running multiple applications, such as DisplayWrite 3 and IBM Drawing Assistant.

4.1.8.4.2. Advanced Program-To-Program Communications

The Advanced Program-To-Program Communications For IBM PC (APPC/PC) supports the SNA application interface (LU 6.2) and allows program-to-program communication over the IBM Token-Ring Network and SDLC communications links. The APPC/PC allows applications programs to be developed for the IBM PC using communications protocols common to other SNA devices, such as the System/370, System/36, System/38, and Series/1.

The addition of an SNA interface promises to be a key selling point for the Token-Ring. An implementation of SNA LU Type 6.2 provides enhanced Systems Network Architecture support for distributed processing. As a communications protocol for a growing number of IBM products, APPC can improve connectivity among distributed transaction programs.

APPC/PC can be used by IBM PCs to help satisfy the requirements to attach to other SNA products using the SNA APPC architecture (LU 6.2). It provides communication services for applications in much the same way that IBM PC-DOS provides disk input/output and file management services. The subsystem will allow a program developer to concentrate on the application function rather than on the implementation of the communication protocols.

APPC/PC supports PU2.1 architecture and provides a peer-to-peer relationship between the IBM PC and the System/36, System/38, and Series/1; it attaches to a System/370 as a PU 2.0 node. It supports PC-to-PC connectivity using a Token-Ring communications link or an SDLC link, and PC to System/370, 36, 38, or Series/1 connectivity, using an SDLC communications link.

APPC/PC supports both the IBM Token-Ring Network PC Adapter and SDLC adapters. It provides a common program-to-program protocol that allows multiple conversations between applications running on an IBM PC and a System/370 (CICS/OS/VS Version 1, Release 7), System/36, System/38, Series/1 (Realtime Programming System Version 7.1), or another IBM PC. APPC/PC does not provide direct connectivity between sessions on the IBM Token-Ring

Network link and sessions on the SDLC link, but it does provide the application programming interface to allow a user-written program to communicate between sessions on the two links.

The APPC/PC program has been designed to have an open application program interface -- that is, it provides an interface which may be used by other products. APPC/PC will appear as a set of communication services provided by the operating systems.

APPC/PC supports a program interface through a set of verbs that allows application transaction programs to communicate at a conversational level with no session awareness. Application programs can engage in either mapped conversation -- for use in communications between user-written programs -- or basic conversation -- intended to provide a lower-level interface such as that required by logical unit service programs.

A PC can use both the APPC/PC and the NETBIOS program for concurrent access to the Token-Ring. Figure I-41 is a schematic representation of the two high-level programming interfaces; only one Personal Computer Adapter is required. The program requires 195 Kbytes of memory in addition to the 7 Kbytes required for the adapter handler program and the ordinary DOS requirement.

4.1.8.4.3. Token-Ring/PC Network Interconnect Program

A dedicated IBM PC running only the Token-Ring Network/PC Network Interconnect Program can link the two kinds of networks so that devices can intercommunicate. Network application programs to be linked by the program must use NETBIOS protocols. Many programs that operate on either network using the NETBIOS programming interface will operate across both networks.

Up to 16 names for each network can be defined to the interconnect program; therefore, the maximum number of attached devices that can be identified to the other network is 16. (Some applications, such as the file server function of the IBM PC Network Program, require the use of multiple names.)

4.1.8.4.4. Asynchronous Communications Server Program

The IBM Asynchronous Communications Server program provides IBM PCs attached to the IBM Token-Ring Network or IBM PC Network access to asynchronous communications applications via circuit-switched connections through a ROLM CBX II, a PBX, or a public switched network. With the IBM Asynchronous Communications Server Program running on an IBM PC, XT, or AT, modems are not required at each individual IBM PC station in order for users to be able to access dial-up applications such as ASCII host systems, other IBM PCs, or information services (e.g., the IBM Information Network, the Dow Jones News/Retrieval

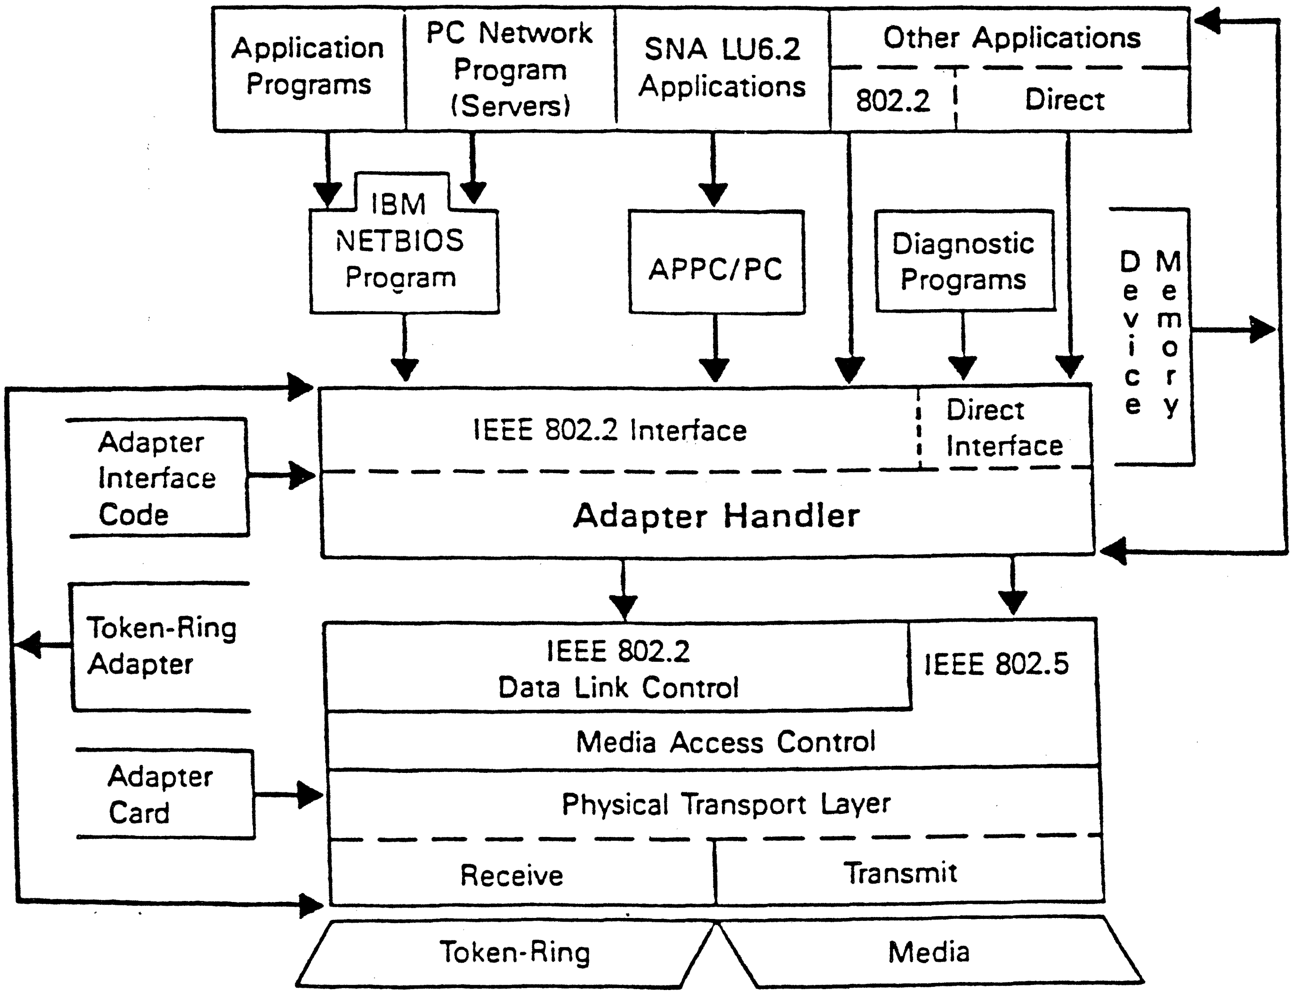

Figure I-41: APPC/PC and NETBIOS

Services, CompuServe).

When an IBM PC user wishes to establish a switched connection, he/she executes an asynchronous communications program that establishes a connection with the IBM Asynchronous Communications Server Program. The asynchronous communications program must use the IBM Asynchronous Communications Server Program protocol via NETBIOS. The server program establishes outbound calls to ASCII applications and accepts inbound calls destined for asynchronous programs running on a networked PC. ASCII data is transferred transparently.

The server is not dedicated, and multiple servers can be attached to either type of network.

4.1.8.5. Physical Interconnectivity For The Token Ring

It is, of course, the IBM Cabling System described earlier (see section 4.1.4) that provides the physical interconnection for the Token-Ring Network. As has already been seen, components for the prewiring of a building include bulk cable, connectors, wall faceplates, and wiring closet distribution panels. Four cable assemblies (path cables) of varying lengths allow the network to be installed in a building that has not been prewired. The cable assembly lengths are 2.4 meters (8 feet), 9.1 meters (30 feet), 22.9 meters (75 feet), and 45.7 meters (150 feet).

The structured wiring technique of connecting devices to wiring closets allows new devices to be added to the network without disrupting the other networked devices. Type 1 or 2 Cable specified in the initial IBM Cabling System allows the data transmission rate to be up to 16 Mbps and the distance between workstation and wiring closet to be up to 300 meters (990 feet), with 200 meters (660 feet) between wiring closets in a multiple closet environment.

In an apparent concession to cost pressure from AT&T's STARLAN, IBM announced an extension to the IBM Cabling System called Type 3 Specified Media -- or, in lay terms, telephone wire -- on October 15. While Token-Ring networks using Type 3 Specified Media will be limited in size and performance compared to those using Type 1 or 2 Cable, some users may find the price advantage worth the performance tradeoff. Table I-2 compares IBM Cabling System Type 1 and 2 Cable and Type 3 Specified Media.

Type 3 Media Filter is used to connect the adapter card to the twisted-pair network: one end of the cable has a D shell connector at the PC end and a standard G-pin modular jack at the other end. Type 3 Media Jumper Cables are used to attach the Type 3 Media to the Token-Ring Network Multistation Access Unit. Figure I-42 provides a more detailed look at Type 3 telephone twisted-pair media. Four pair of Type 3 Specified Media can be included in one sheath -- two pair for the Token-Ring attachment,

	IBM Cabling System Type 1 /Type 2 Cable	IBM Cabling System Type 3 Specified Media
Other Media Applications:	Coax, Twinax, Loop	Coax
Token-Ring Media Operational Parameters:		
Tested Data Rates (Manchester Encoded)	16 Mbps	4 Mbps
Signal Drive Distance (Using 4 Mbps)		
Work Area To 1 Wiring Closet	Max. 300 M (990 ft.)	Max. 100 M (330 ft.)
Wiring Closets Per Ring	Max. 12	Max. 2
Work Area To Wiring Closet If Multiple Wiring Closets	Max. 100M (330 ft.)	Max. 45 M (150 ft.)
Wiring Closet To Wiring Closet (Cabling System Type 1 Only)	Max. 200 M (660 ft.)	Max. 120 M (400 ft.)
Electronic Interference Protection	See GA27-3361	See GA27-3714
Configurations Supported:		
Devices Per Ring	Max. 260	Max. 72
Undercarpet Cable Support	Yes	No
System Price Considerations:		
Media And Accessories	Yes	Yes
Installation Labor	Yes	Yes
IBM Token-Ring Network PC Adapter Kits	Yes	Yes
IBM Token-Ring Network Multistation Access Units	Yes	Yes
Devices Per Ring (Including Devices Used To Connect Rings)	260	72
Type 3 Media Filter	N/A	Yes
Type 3 Media Jumper Cables	N/A	Yes
IBM Token-Ring Network PC Adapter Cable	Yes	N/A
IBM Cabling System Patch Cables	Yes	Yes
Telephone Connectors/ Punch-Down Blocks	No	Yes
3270 Parameters		
From Workstation To Controller Or IBM 3299 Model 1 Multiplexer	606 M (2000 ft.)	273 M (900 ft.)
From Workstation or Controller To IBM 3299 Model 2 Multiplexer	1000 M (3,280 ft.)	273 M (900 ft.)

Table I-2: IBM Cabling System Types 1 & 2 Vs. Type 3 Specified Media

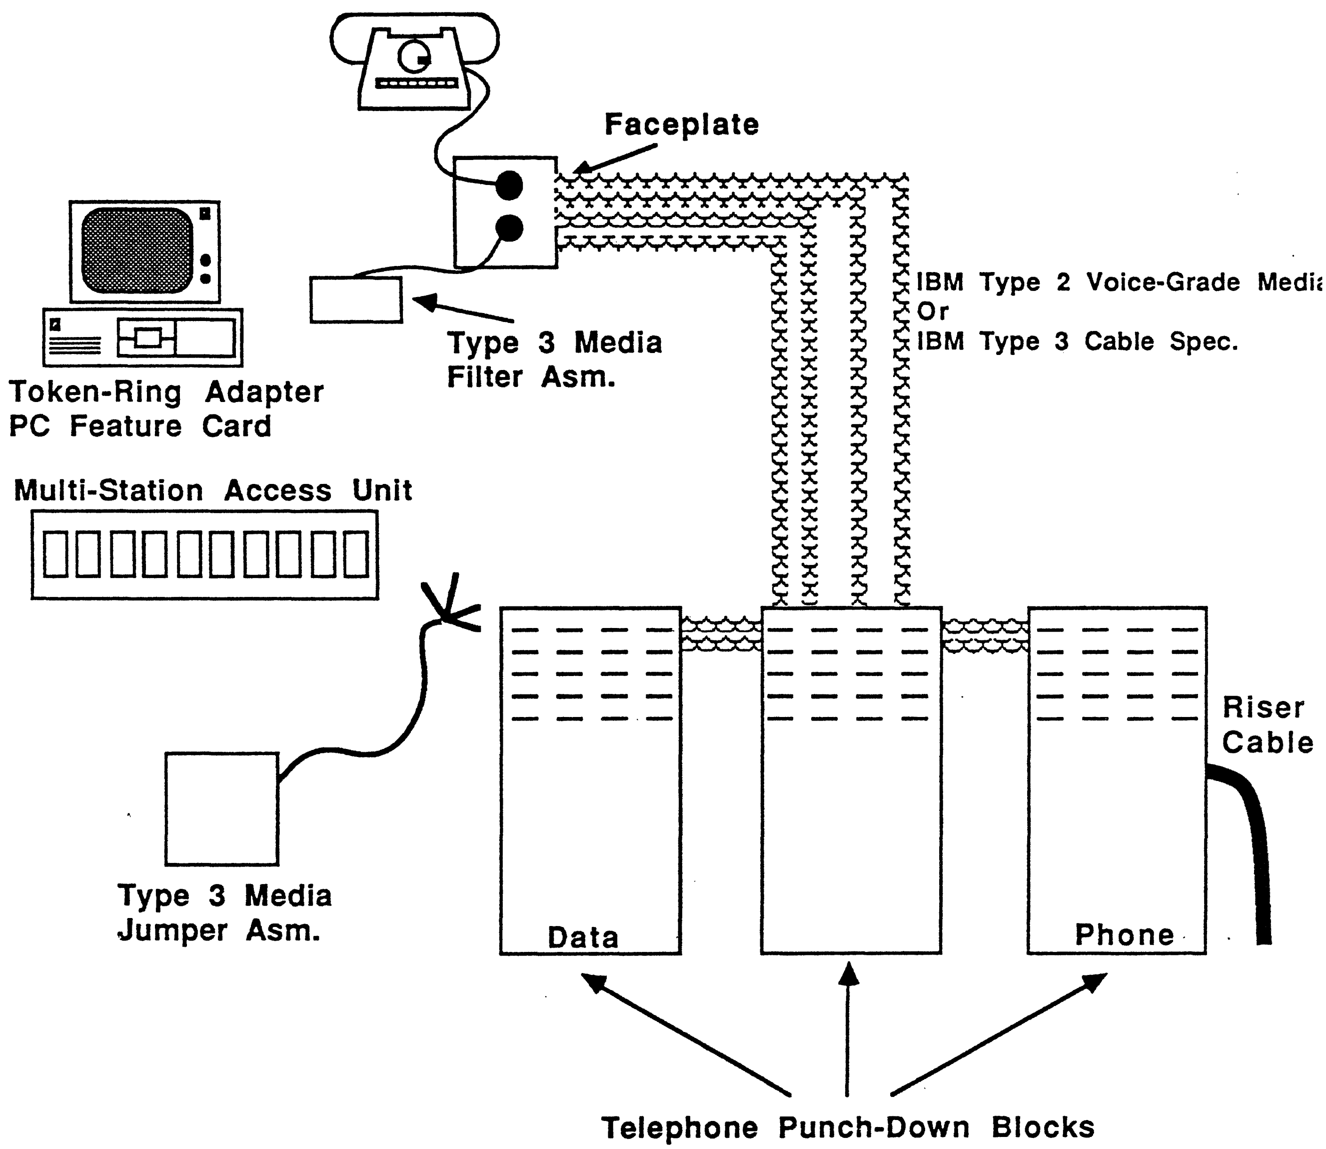

Figure I-42: IBM Cabling System Type 3 Specified Media one pair for 3270 coax-to-twisted pair attachment, and one pair for telephone attachment.

4.1.8.6. Implications Of The IBM Token-Ring LAN

In its October 15, 1985, media blitz, IBM very carefully delineated the differences between the IBM PC Network and the new Token-Ring Network. While both LANs support attachment of IBM PCs, XTs, ATs, and Portables, as well as the NETBIOS interface and interconnecting PCs and related devices in a work group or departmental environment, the Token-Ring is intended to connect a variety of systems and devices across an entire establishment. The interconnect program previously described will provide a gateway between the two networks in an organization.

While the PC Network is a broadband network using coaxial cable, the Token-Ring Network is a baseband implementation that can use either shielded or telephone twisted-pair wiring. Broadband has the advantage of being able to support voice, data, and motion video on one medium, drawing on standard CATV components, and it also has lower attachment costs than baseband, as well as high-bandwidth transmission capabilities for supporting OEM devices on multiple channels. However, broadband networks can be difficult to expand.

On the other hand, baseband transmission, such as that used by the Token-Ring, allows for smooth system expansion: new devices can be added without disturbing the remainder of the network. However, baseband may also mean that networks cannot span the same long distances as broadband networks.

The star-wired ring configuration that IBM has chosen to adopt is a logical choice, in view of its Cabling System wiring closets: the star-wired ring topology can support large network configurations, and because the data path is in a single direction, it is relatively simple to administer both on a day-to-day basis and in the event of a malfunction.

The Token-Ring improves on the PC Network's 2-Mbps transmission rate by offering 4 Mbps, and it can support up to 260 devices using shielded twisted pair or 72 devices using telephone twisted-pair, as compared to the PC Network's 72 devices. In addition to NETBIOS, the Token-Ring supports the SNA LU 6.2 programming interface, Advanced Program-to-Program Communications for the IBM PC (APPC/PC). This support indicates that the network is likely to be the cornerstone for all of IBM's future office networking plans.

IBM has stated that it intends to provide an interconnection between the IBM PC Network, the Token-Ring Network, and a future industrial local area network. The IBM industrial LAN will use a token-passing protocol rather than broadband transmission, because token-passing protocol is more sensitive to the response time requirements of a process control environment. IBM has also indicated that its industrial LAN will be an open system allowing the attachment of IBM PCs in addition to third-party process controllers.

IBM caught the industry experts off-guard when it announced a Token-Ring LAN that was essentially a PC LAN, rather than a total token-ring system solution incorporating terminals and mainframes. When all the necessary interfaces are available, IBM's Token-Ring will provide complete local communications at IBM customer sites. In the interim period, look for other vendors to provide Token-Ring interfaces to non-IBM devices, as well as gateways and applications software for the Token-Ring Network, as the new LAN catches on with large IBM accounts.

Six months plus one day from the original IBM Token-Ring Network introduction October 15, 1985, IBM announced major additions and enhancements to its strategic network. On April 16, 1986, IBM significantly enhanced its LANs with new LAN host, departmental systems, and Personal Computer attachments, and with products and programs that support IBM's Token-Ring Network and broadband PC Network.

The following is a summary of the highlights of the announcements:

- New special hardware features support direct attachment of the 3725 Communication Controller to the Token-Ring Network using the IBM Cabling System, including Type 3 Specified Media.

- IBM Token-Ring Network PC Adapter II, when operating with the IBM Token-Ring Network Bridge program, provides bridging of Token-Ring Network rings.

- 8218 Token-Ring Network Copper Repeater is designed to operate on IBM Cabling System data grade media and extend the distance between Multistation Access Units up to 2500 feet.

- 8219 Token-Ring Network Optical Fiber Repeater extends the allowable distance between Multistation Access Units up to 6,600 feet.

- Type 9 specified media contains two twisted-pair of #26 AWG copper wire and is offered as a potentially lower cost alternative to Type 1 plenum cable.

- System/36 can be connected to the Token-Ring through a directly attached PC AT functioning as a LAN communications controller.

- Token-Ring Network Bridge Program is a licensed program for the PC AT and Industrial Computer AT that supports message directing between IBM Token-Ring Network rings.

- Token-Ring Network Manager, another IBM licensed program, assists the user in problem determination and error recovery for the Token-Ring Network.
- The IBM Token-Ring Network/IBM PC Network Interconnect Program enhancements allow the exchange of information between IBM PCs attached to two IBM PC Networks.

With these new Token-Ring product additions, with the original products introduced on October 15, 1985, and with additional products (e.g., direct 327x-Token-Ring connectivity) whose introductions are anticipated for the last half of 1986, the IBM Token-Ring Network could well be the major driving force within the local area networking field far into the next decade.

4.2. Other Token-Ring-Based LANs

4.2.1. Ztel PNX

Ztel, Inc.'s Private Network Exchange (PNX) is an integrated voice/data system which combines the data capabilities of a local area network with the voice capabilities of a PBX. The Ztel PNX was one of the first products for which token-ring compatibility was claimed.

While aggressive watchers of IBM had detected by as early as 1980 that the LAN trend for IBM was to be rings, as it became more and more obvious during 1982 that IBM was going to pursue token rings, a number of more venturesome companies decided that they should commit their products to the IBM strategy: Ztel, Inc., was one of these companies.

Ztel was actually incorporated before 1982 -- in January 1981, with financing provided by Adler & Company and GE. A second round of financing in the spring of 1983 produced an additional $23.5 million. (One of the new investors was NCR, whose share of Ztel subsequently rose to 19%.)

4.2.1.1. PNX System Concept

The PNX is aimed at medium to large installations of 200 to several thousand ports; it is said to be price-competitive down to 200 ports with a repackaged version. The first release of the PNX concentrates on a structure of telephone and data switching, with later releases to support office applications and token ring local networks.

The PNX intends to use an industry-standard 802.5 ring architecture. The rings provide universal communications paths between users, terminals, hosts, resources, and other devices. Two distinct types of rings are employed -- circuit and packet.

The circuit ring is a time-slotted ring which transmits 64 Kbps digitized voice and up to 56 Kbaud data. Each circuit ring supports 228 simultaneous voice conversations, or data sessions. The rings can be resident, i.e., entirely within central site cabinets if a traditional voice/data PBX solution is desired; or cable and/or fiber may be brought directly to the areas where true LAN operations are required. Circuit rings operate at 18.4 Mbps. The circuit rings are best suited to continuous, high-speed transmissions, digitized voice, or any transparent circuit data, and can handle any protocol. Data or digitized analog information, such as voice, is switched transparently by these rings.

Any data device with an IEEE 802.5 token-passing ring LAN interface will be able to directly attach to the PNX ring. The packet ring also networks data from RS-232/449 devices. Data is packetized at the network entry point, such as at the telephone sets, to achieve significant savings in bandwidth. The PNX packet ring is to be plug-to-plug compatible with IBM's Token-Ring LAN offerings. Packet rings in the network operate at a rate of 4 Mbps. The ring transmission media may be either coaxial cable or optical fiber, which is used in severe environments or to cover distances longer than five miles.

A load-sharing feature of the system allows use of the system while providing back-up: if one ring fails for any reason, its traffic is automatically load-shared within the remaining network or rerouted along a back-up ring.

Analog information is digitized upon entering the network, such as at the telephone set. Data is typically packetized, that is, broken into discrete blocks of data which contain information about the sender and the receiver. These packets are intermixed with packets from other transmissions and routed through the token ring to their proper destinations, allowing efficient use of the network. Data which is a nonstandard protocol, or non-bursty and high-speed, is typically transmitted over the time-slotted circuit rings at aggregate speeds that can be as high as 4 Mbps.

The PNX has a distributed architecture, typically with multiple system processing units; units can be added or taken away to make a system larger or smaller. A PNX can be cost-effective from as low as 200 lines up to the largest possible installation, due to the extreme modularity of the architecture.

4.2.1.2. PNX System Components

A variable number of System Processing Units (SPUs) are linked by the circuit rings and packet rings; these units provide the switching and processing capabilities of the PNX. Figure I-43 illustrates Ztel's envisioned PNX System.

Three types of subsystems can comprise an SPU: these are the Applications Processing Subsystem, Switch Processing Subsystem, and Data Conversion Subsystem. The actual percentage of the physical SPU dedicated to each of these functions can vary widely, per customer requirements. In most systems, each SPU would incorporate all three subsystems, but in some situations -- particularly large systems -- it is conceivable that an entire SPU could be dedicated to one subsystem function, such as applications processing. The Switch Processing Subsystem handles all call processing and switching; the Applications Processing Subsystem handles all system control processing, as well as applications such as on-line directory and least-cost routing; and the Data Conversion Subsystem, as an example, processes protocol conversions within the system to ensure transparency to the user.

Three digital telephones were introduced with the PNX: the Z/28 telephone set has a 40-character display and 28-user-definable feature keys; the Z/2 is a telephone set with 12 feature keys; and the Z/4 contains 4 feature keys. Standard 2500/500 analog stations are also supported.

4.2.1.3. PNX Software

The Enhanced Networking software package for Ztel's PNX allows a company to easily integrate the system into existing private networks and, in many cases, use the PNX's advanced capabilities to improve network performance and increase network control. By integrating the capabilities for tandem calling, least-cost networking (LCN), uniform dialing, and main/satellite operation, the Enhanced Networking software package permits PNX participation in a wide-area network.

Enhanced Networking also provides PNX users with a controlled environment for expanding single-site systems to multi-site capabilities: used in conjunction with Ztel's LCN software, this package permits the network to route calls outside the network and perform digit insertion and translation to more efficiently route calls. Its operation is transparent to users.

Enhanced Networking software provides PNX compatibility with eastablished site code numbering plans, which are required for inter-site dialing on a private network. The PNX will easily accommodate any numbering plan a customer chooses.

In addition, the ability of the PNX to function as a satellite system off another PBX or to serve as the main system

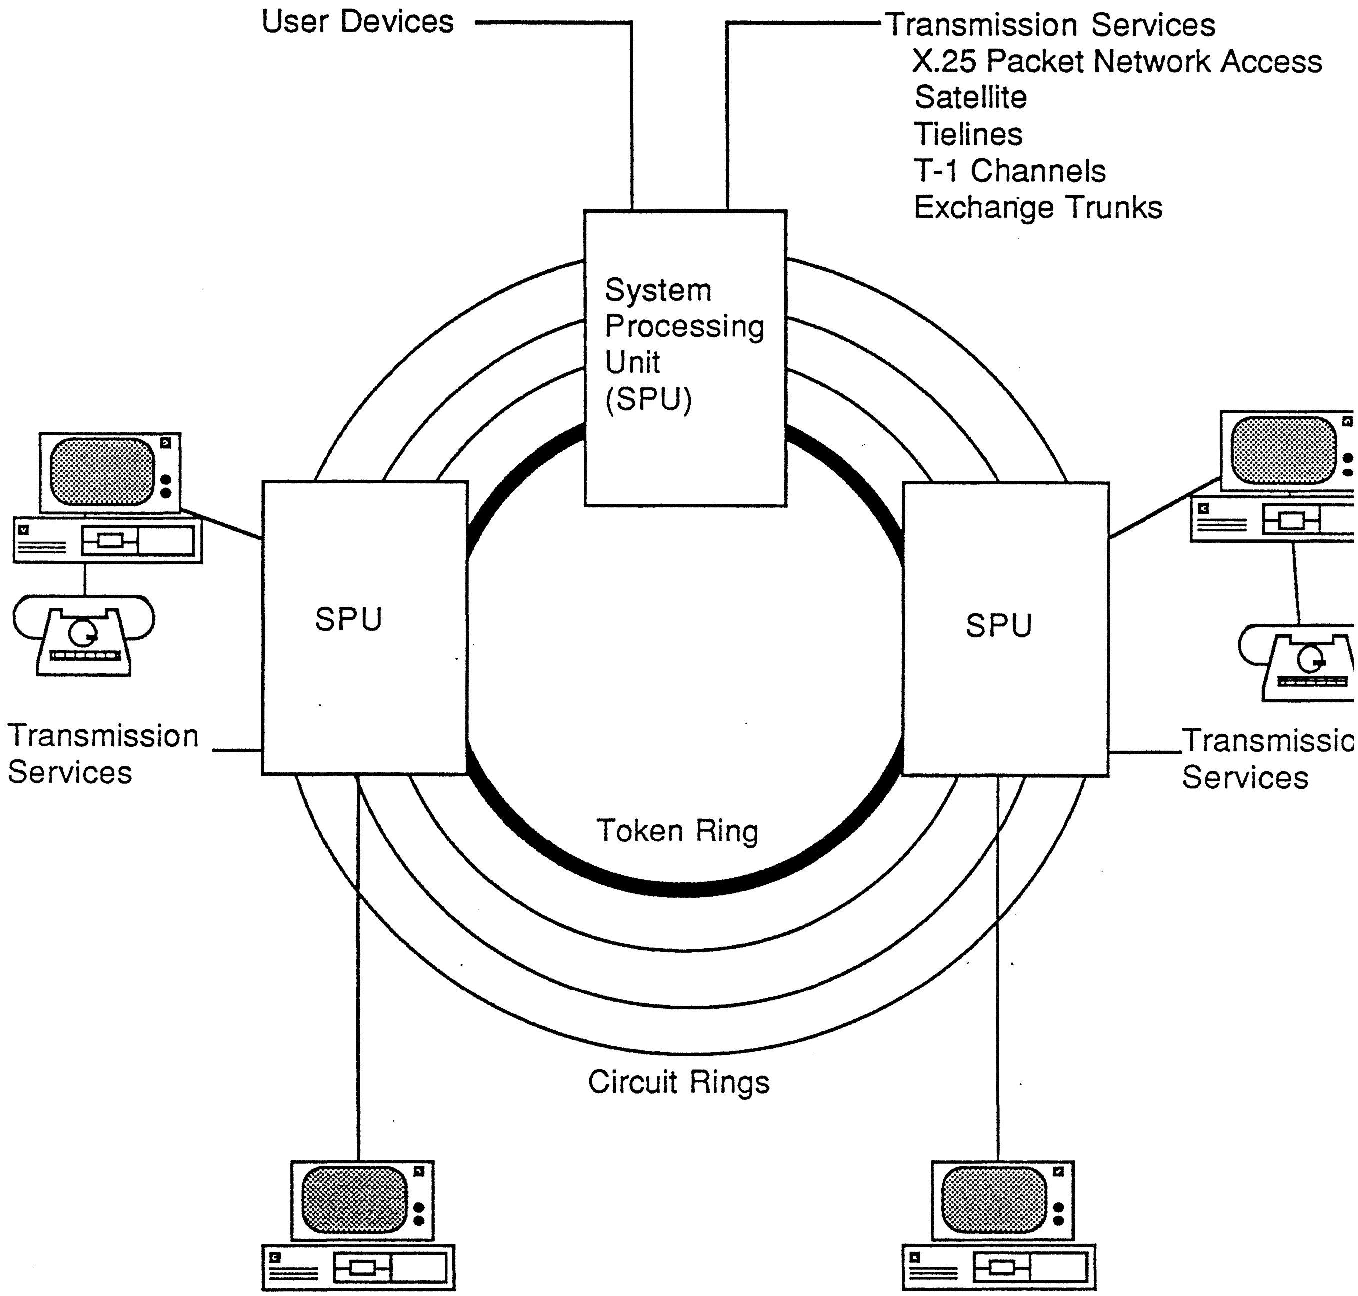

Figure I-43: Original Ztel PNX System Concept for other satellite PBXs is provided through the Enhanced Networking package. The PNX eliminates the need for users to dial tie line codes in main/satellite operation.

4.2.1.4. <u>PNX System Maintenance</u>

For reliability, the Ztel PNX incorporates on-board maintenance processors into its architecture. The backbone of Ztel's maintenance philosophy is the Maintenance Thread: this consists of a subnetwork of maintenance processors dedicated to enhancing reliability, availability, and serviceability. A maintenance processor on each circuit board monitors the board for critical function faults. In addition, maintenance processors can be activated by a maintenance person, either on-site or from a remote location, for additional testing or to recheck original results.

The system schedules on-line testing on a regular basis. When a fault is suspected, additional testing is activated that can isolate the problem to one board. The maintenance subsystem is independent of primary communications in the PNX and does not affect its operations.

The three levels of maintainability provided by the PNX include the software which detects faults (generally to a specific board), remote diagnostics, and an on-site maintenance person who can request repeated or additional tests.

Hardware will detect a ring malfunction; the maintenance system provides a report after isolating the fault and automatically repairing it.

All processors that are part of the system but not part of the Maintenance Thread are reporting to the Thread on a continual basis. For example, the microprocessor in the digital telephone station is not part of the Thread; it reports conditions back to a higher-level processor dedicated to maintenance which, in turn, reports to an executive maintenance console. Critical problems are reported on the attendant console, and non-critical problems may appear on the system printer for an audit trail.

4.2.1.5. <u>PNX History, Status</u>

As pointed out above, the Ztel PNX was one of the first products to claim token-ring capability.

Ztel first demonstrated the local area network capabilities of its Private Network Exchange (PNX) at the 37th Annual Conference and Exhibition of the International Communications Association (ICA), in Las Vegas on May 8-10, 1984. The PNX integrated a LAN with a private branch exchange (PBX) for voice and data communications.

The PNX LAN demonstrated at ICA incorporated multiple nodes through which IBM, Apple, and DEC personal computers accessed LAN-based processing and data file resources. Basic PNX nodes, called System Processing Units (SPUs) (see 4.2.1.2 above), can each provide for the attachment of up to 1500 voice and/or data terminals. Devices are attached to the ring via RS-232-C connections, and information is transmitted via an SPU-based ASCII PAD (packet assembly and disassembly). Up to 256 SPU nodes can be incorporated into a single PNX system. Token rings interconnect the distributed SPUs and are the vehicles for moving information between SPUs and user-attached equipment.

At that same ICA Conference and Exhibition, Ztel announced the signing of a $25 million, multi-year contract with ASI Telesystems, Inc.

In mid- to late-1984, Ztel completed the Beta test installation of a one-node PNX at Compugraphic Corporation. (Compugraphic is one of the world's largest suppliers of computerized composition systems for publishing, printing, advertising, and other graphic arts operations.) The Compugraphic PNX provided integrated voice and data communications for its Research and Development and Engineering departments. The configuration included approximately 480 ports for support of lines and trunks, as well as over 200 Z/28 full feature telsets which have 28 programmable features keys and a 40-character display for messages. The Data Adapter option on 51 of the Z/28 telsets allowed users to connect terminals, personal computers, workstations, and most other EIA-compatible devices to the PNX. Compugraphic users then had access to DEC VAX computers via connections provided by 49 Ztel data servers.

In 1984, Ztel also signed a $35 million, multi-year letter-of-intent with Consolidated Data Systems, Ltd., calling for Consolidated to market Ztel's PNX. Consolidated Data is a Paramus, NJ-based company which provides communications systems research, planning, and design to customers throughout the United States and in a number of foreign countries.

In July of 1984 Ztel then signed a three-year, $35 million contract with Tel-Matic Systems, a Toledo, OH-based interconnect distributor.

During this same period, FirsTel Information Systems, Inc. (a US West company) sold a Ztel PNX system to the State of Utah: the system was scheduled to be operational October 1, 1984. However, FirsTel Information Systems then signed an additional agreement to also distribute the InteCom IBX. Originally, FirsTel was to market only the Ztel PNX and the NEC NEAX 2400 switch; as a result of this second agreement, the company was then marketing the PNX, the NEAX 2400, and the IBX. FirsTel's agreement with InteCom seemed to represent a break by

FirsTel from Ztel, on the basis that InteCom's IBX competes more directly with the PNX than with the NEC products.

The PNX system sold to the State of Utah by FirsTel and scheduled to be operational by October 1, 1984, constituted the first commercial shipment of the PNX. This PNX was to integrate all office automation equipment (telephones, personal computers, word processors, workstations, facsimile machines, and mainframe computer) at participating state agencies into one central communications systems; approximately 250 stations were to be included in the scheduled October 1 cutover to the PNX, and all equipment and services were located on-site at the North Plaza Office Building.

Troubles with this State of Utah PNX signaled the beginning of a downward turn of events at Ztel. Within eight months of its installation, the PNX was pulled out of the North Plaza Office Building. Craig Jorgenson, telecommunications manager for the State of Utah, said the plug was pulled because the 255-line, voice-only PNX installation did not deliver on feature's Ztel has promised. When Ztel subsequently offered to provide the missing features, the State of Utah asked FirsTel -- as the contractor -- for guarantees that FirsTel decided it could not give. These guarantees were a lifetime warranty and free spare parts.

Jorgensen stated that the PNX never met the specifications required in the original RFP, and that Ztel missed several due dates on the required upgrades -- especially in mid-March for software release 1.1. (Software release 1.1 provides single-node data transmission, direct-inward-dialing, tie-line support, and an improved maintenance thread.) Jorgensen condensed Utah's problem down to one basic premise: from the standpoint of dealing with public funds, it needed specific guarantees from both FirsTel and Ztel, and that in looking at those guarantees, neither firm felt they could live up to the stringent requirements if Utah accepted 1.1 beyond the due date.

Jorgensen went on to say that when the state wanted to move stations or reconfigure the system, FirsTel first had to bring the entire system down at night at its own expense, and that this happended frequently.

During the bidding process in 1984, FirsTel had offered a NEC Telephones, Inc., NEAX 2400, with the PNX as a secondary bid; at that time FirsTel said it would replace the PNX with a NEAX 2400 at its own expense if the state was dissatisifed with the PNX.

Thus, at Ztel, by the spring of 1985 the optimism of 1984 had vanished. Smarting from a December 1984 takeover by its investors and struggling under the burden of liabilities totaling $25.3 million, Ztel in mid-May announced a reorganization plan under which it appointed Murray H. Bolt to the position of president and chief executive officer. (Mr. Bolt joined Ztel from Codex, where he was vice-president and GM of the LAN Products Business Unit: in this capacity, he was responsible for LAN, terminal, and switching products. Prior to joining Codex, Mr. Bolt had been employed for 18 years with IBM in various development and marketing management positions, including some dealing with the IBM Token Ring.) A new management team reporting to Bolt was also announced at that time.

The company further announced that its investors had committed to provide $17,000,000 in new equity financing under a reorganization plan. This financing was to enable Ztel -- under Bolt's direction -- to continue to accelerate the enhancement of its products and expand its marketing capabilities.

The startling portion of Ztel's mid-May announcement was its disclosure that it had concurrently filed a petition for protection under Chapter 11 of the Bankruptcy Code. (This reorganization plan was quickly approved in principle by Ztel's creditors and the U.S. Bankruptcy Court in Boston.) The simultaneous filing of Chapter 11 and a highly detailed reorganization plan was unusual and indicated that Ztel had been planning the move for some time: in normal circumstances a firm files for Chapter 11 and is then given 90 days to submit a plan to repay some or all of its debts.

The Chapter 11 filing by Ztel was directly attributable to the refusal of one of its principal investors -- Adler & Company -- to take part in a new round of funding and earlier efforts this year by the company to raise $25 million through the issuance of new securities.

Approval of the reorganization plan -- which also involved re-incorporation of the firm -- opened the way for the $17 million new cash infusion in exchange for stock, which reduced Adler & Company's holdings from 20% to 2.5% and increased the holdings of NCR, General Electric, Capital Corporation, and Hillman Company to about 28% each; previously NCR owned about 18.5% of Ztel.

The reason for Adler & Company's refusal to continue its participation in funding of Ztel was never made public; of the $52 million invested in Ztel to that date, over $5 million was contributed by Adler, head of the venture firm. Legal representatives of Adler & Company were not present at any of the court proceedings.

Ztel also disclosed at this time that Fred Adler and James Long had resigned from the firm; both had been directors, and Long also had been acting chief operating officer since the December takeover of the firm by its investors.

In its Chapter 11 filing, Ztel said it had expected to begin commercial shipments of its PNX in the fourth quarter of 1984, but due "...to overly ambitious hiring and delayed development of some features necessary for widespread market acceptance, sales and production did not increase to anticipated levels, and the cost of sales incurred by Ztel exceeded the level required to support the manufacture of systems sold in 1984."

While the entire top management team left the firm in the wake of the December takeover by the investors, co-founder Henry Zannini was dismissed on May 23, 1985; Zannini -- who was the "Z" in Ztel -- was vice-president of development until the December takeover, when he was appointed vice-chairman.

On August 2, 1985 -- eleven weeks after it filed its original petition -- Ztel emerged from Chapter 11 reorganization with a court-approved recapitalization plan. G.E., NCR, and Hillman Company -- with the addition of Harvest Ventures -- had agreed to a funding proposal of approximately $17 million; and while in Chapter 11 proceedings, the company, under the direction of Murray H. Bolt, Ztel's new president and CEO, had completely restructured its entire operation, putting in place a new management team, a new marketing plan, and a targeted, market-driven product focus and direction. The reorganization plan also provided for 100% payment to all Ztel creditors.

Ztel seems to have emerged from its bleak times with its cadre of distributors still in place. For example, ASI Telesystems and Tel-Matic Systems, Inc., have remained committed to selling and supporting Ztel's PNX, from the first site tests to the first commercial product offering. In addition to existing distributor agreements, the company is now negotiating long-term contracts with new distributors to further extend its distribution channel.

The ability to add custom software to the PNX and the fact that the system is OSI-structured should attract OEMs. Also, since PNX software has been written in the C language, and since the OSI model is rigorously followed, OEM buyers can easily program specific applications. Furthermore, the token ring architecture is particularly appropriate for real-time processes such as interactive voice, since token ring LANs are deterministic: worst-case delay for any access is fixed, so performance under any load conditions is know beforehand and can be accommodated in the design.

Ztel has now delivered a number of PNXs to customers. At present, the PNXs being delivered are standard PBX nodes with no token-ring LAN capabilities; Ztel seems to be having problems implementing its token-ring configuration, though it has not abandoned efforts to do so.

The original Ztel PNX is now designated the PNX 5300; two more recent PNX configurations, the PNX 5200 and PNX 5400, support voice and data line sizes of up to 300 and 1,500 lines, respectively. All three models have been well received by their customers, and offer features such as a CRT-based attendant console; an on-line directory; integrated system-wide messaging; full data transmission capabilities; least cost networking; and service and diagnostic tools.

4.2.2. Apollo DOMAIN

The Apollo DOMAIN system is an integrated, high-speed local network of personal workstations and server computers manufactured by Apollo Computer, Inc., of Chelmsford, MA. The DOMAIN (an acronym for Distributed Operating Multi-Access Interactive Network) was designed to be an integrated distributed computing environment that would serve as an alternative to dedicated time-sharing minicomputers.

DOMAIN is a baseband, coaxial-cable-based, token-passing ring that operates at 12 Mbps. The maximum permissible distance between nodes (Apollo workstations and servers and gateway-connected devices from other manufacturers) is 1000 meters. Media access is achieved by one of these nodes capturing a token -- the only means whereby transmission may occur. When transmission is completed, the node passes on the token. Only one token must be on the ring at any instant.

Each node's ring controller provides the node with a unique node ID, which is assigned at the factory and contained in the controller's microcode PROMs. The maximum packet size supported by the controller is 2048 bytes.

Apollo Computer designed the DOMAIN system for applications of a highly interactive or computational nature, based on the assumption that a dedicated computer for each user provides better service than conventional time-sharing systems do.

The standard computational node structure of the Apollo DOMAIN consists of four major parts: the central processing unit, the memory management unit, the I/O system, and the display system (refer to Figure I-44).

The CPU of Apollo's low-end systems is built from dual Motorola MC68000 microprocessor chips, with one of the chips dedicated solely to handling page faults. While these chips are internally organized around a 32-bit architecture, the external bus is only 16 bits wide. Apollo's high-end systems (DN460/DN660) use a proprietary chip design which supports full 32-bit architecture.

The memory management unit is used to translate 24-bit process virtual address spaces into 22-bit physical memory address spaces. The memory is organized into 1/2-Mbyte units and is available with full error correction code protection in sizes up to 3-1/2 Mbytes. This limit is set by the density (16K) of the memory chip currently used in the system.

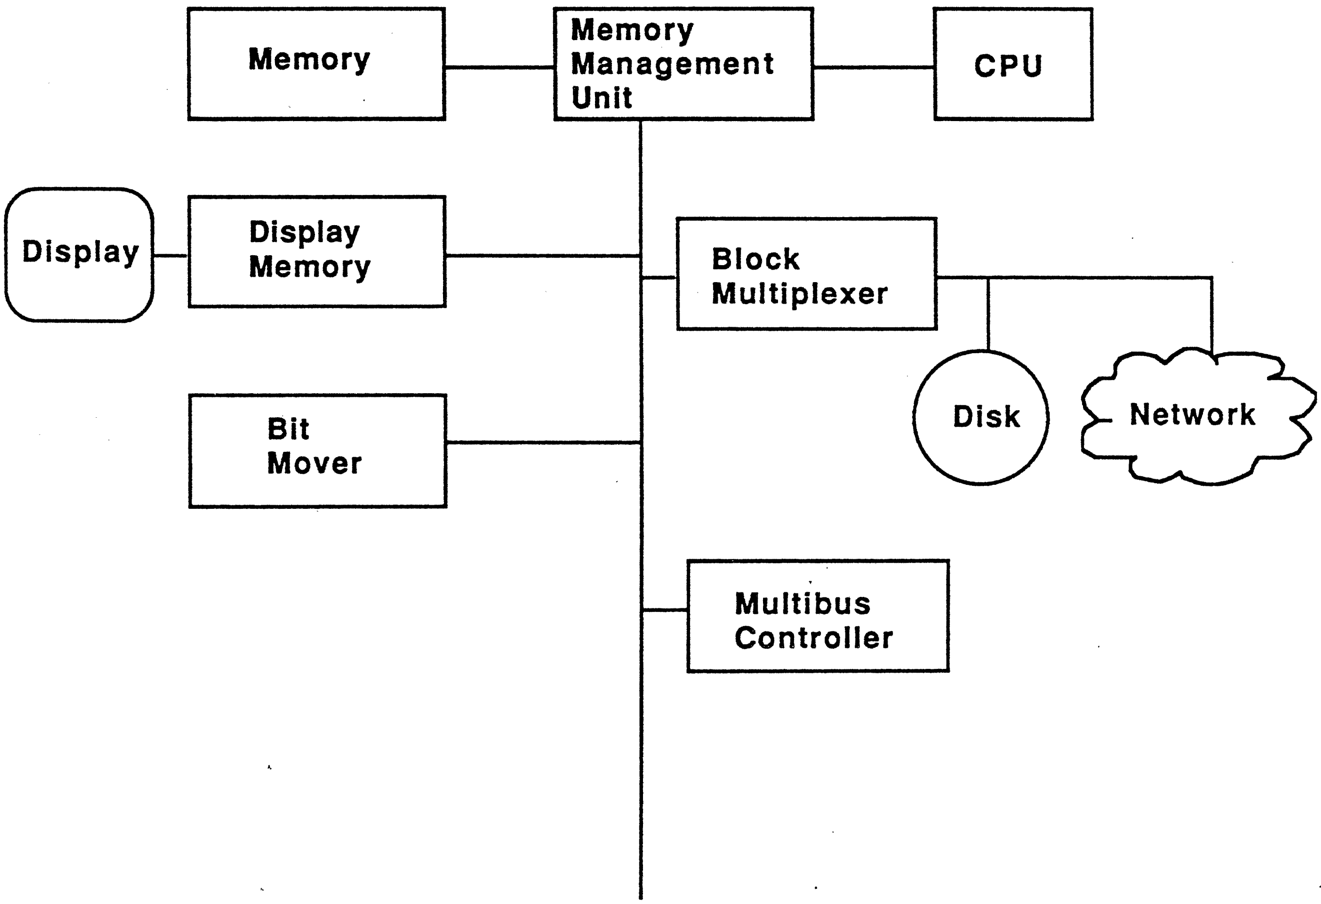

Figure I-44: Apollo DOMAIN Node Organization

The block multiplexer portion of the I/O system handles all disk and network traffic. Designed to maximize node-to-node processing across the network, it allows packet transfer at near memory speeds and without superfluous buffering. Full bandwidth access to primary memory is available to the network, with all other block transfers -- such as the 34-Mbyte Winchester disk -- disallowed. A Peripheral Node Adapter is available to connect line printers, magnetic tapes, and other devices to the node. Because Apollo uses the Intel MULTIBUS for this adapter, hardware was easily added.

In the DOMAIN display system, a separate, autonomous 1/8-Mbyte bit map memory, organized into a square array of 1024 bits per side, is used to refresh an 800 by 1024 bit map CRT; the remaining display memory area is used as temporary storage for character font tables. Rectangular areas can be moved by the 32-Mbps hardware bit mover from any one place on the screen to any other place on the screen. The high-resolution CRT terminal can display the output from multiple programs concurrently through multiple windows. The display manager allows the user to independently control what is physically shown, through the use of function keys on the user keyboard and optional mouse or touchpad.

The operating system (AEGIS) for DOMAIN is used in controlling the mapping of object address spaces into process address spaces, translation of process virtual address spaces into physical memory address spaces, transfer of pages of physical memory into and out of the memory to either local or remote disks, and relating of object onto disk data blocks. AEGIS supports FORTRAN 77, PASCAL, and C and LISP in addition to an editor, and it also strongly supports the "Software Tools" package of standardized program modules. Finally, the network-wide operating system coordinates the user's access to network-wide facilities in a location-transparent manner.

Network reliability is achieved by bypassing each node's ring repeater with a relay, which closes in the event of a node crash or power failure. Network maintainability is enhanced by using quick-disconnect hardware for attaching nodes to the ring, so that they can easily be removed if they fail and easily be reconnected after they are repaired. Also, if the token is destroyed (because of entry or exit of nodes from the ring), software timeouts will detect the condition and recover by forcing a packet transmission, thereby reinserting a token on the ring.

4.2.3. Prime Computer Ringnet

Ringnet is a commercially available token-ring-based LAN developed by Prime Computer, Inc., of Natick, MA. Developed by the firm for the express purpose of connecting a number of Prime minicomputers at a single site, Ringnet forms a portion of Prime

Computer's more comprehensive network architecture, Primenet, which consists of a set of communications protocols and supporting products.

The physical medium for Ringnet is a ring of coaxial cable; attachment points along the ring are generally within 250 meters of one another -- although larger distances are possible by means of adding amplifiers or repeaters to the ring. The ring operates at a speed of 10 Mbps. Control of the network is completely decentralized and depends exclusively upon a circulating token for arbitrary control of the communication medium. Ringnet is a unidirectional loop network, and each operational node normally retransmits the incoming bit stream.

A Ringnet LAN consists of two or more host nodes (a host node can be any one of several Prime superminicomputers -- see Figure I-45), each of which contains (1) an interface that is connected to the cable, (2) a node controller that plugs into a standard Prime I/O bus, and (3) network software that interfaces to PRIMOS (Prime Operating System, used on all Prime computers). (The node interface -- or junction box -- comprises an electronically operated relay controlled by the node controller; when deselected by the controller or by a failure of system power, it provides a straight-through connection that bypasses the node -- otherwise, it links the transmission lines to transmitting and receiving circuits in the node controller.)

In order to facilitate running of application software on three different kinds of networks (local, long-haul, and public data networks), PRIMENET software provides the user with three network services common to all three types of networks: Interactive Terminal Support (ITS), an Interprocess Communication Facility (IPCF), and a transparent network File Access Method (FAM). The interprocess communication facility provides a virtual circuit capability at the program level, (i.e., a set of operating system calls for X.25 services) upon which the other two network services are based. The full use of these services allows any terminal connected to a node of a local network to be used to login to any other node, and the free sharing of all host files by all local network users, though the level of sharing may be modified by a system administrator.

Each message on the network is transmitted as one or more packets, where a packet comprises a header field, a data field, and a trailer field. Data bits and some of the header information are passed to the node controller (under program control) by the host, while the remainder of the header bits and trailer bits are handled exclusively by the node controller hardware. Packets may vary in size from a minimum of 4 bytes (header-only) up to a maximum of 2048 bytes (header plus 2044 data bytes). A transmitting node will interrupt its host when a packet has been transmitted and acknowledged by one or more nodes if its buffer is full, or if a network fault has occurred; a receiving node will interrupt the host only when a packet has

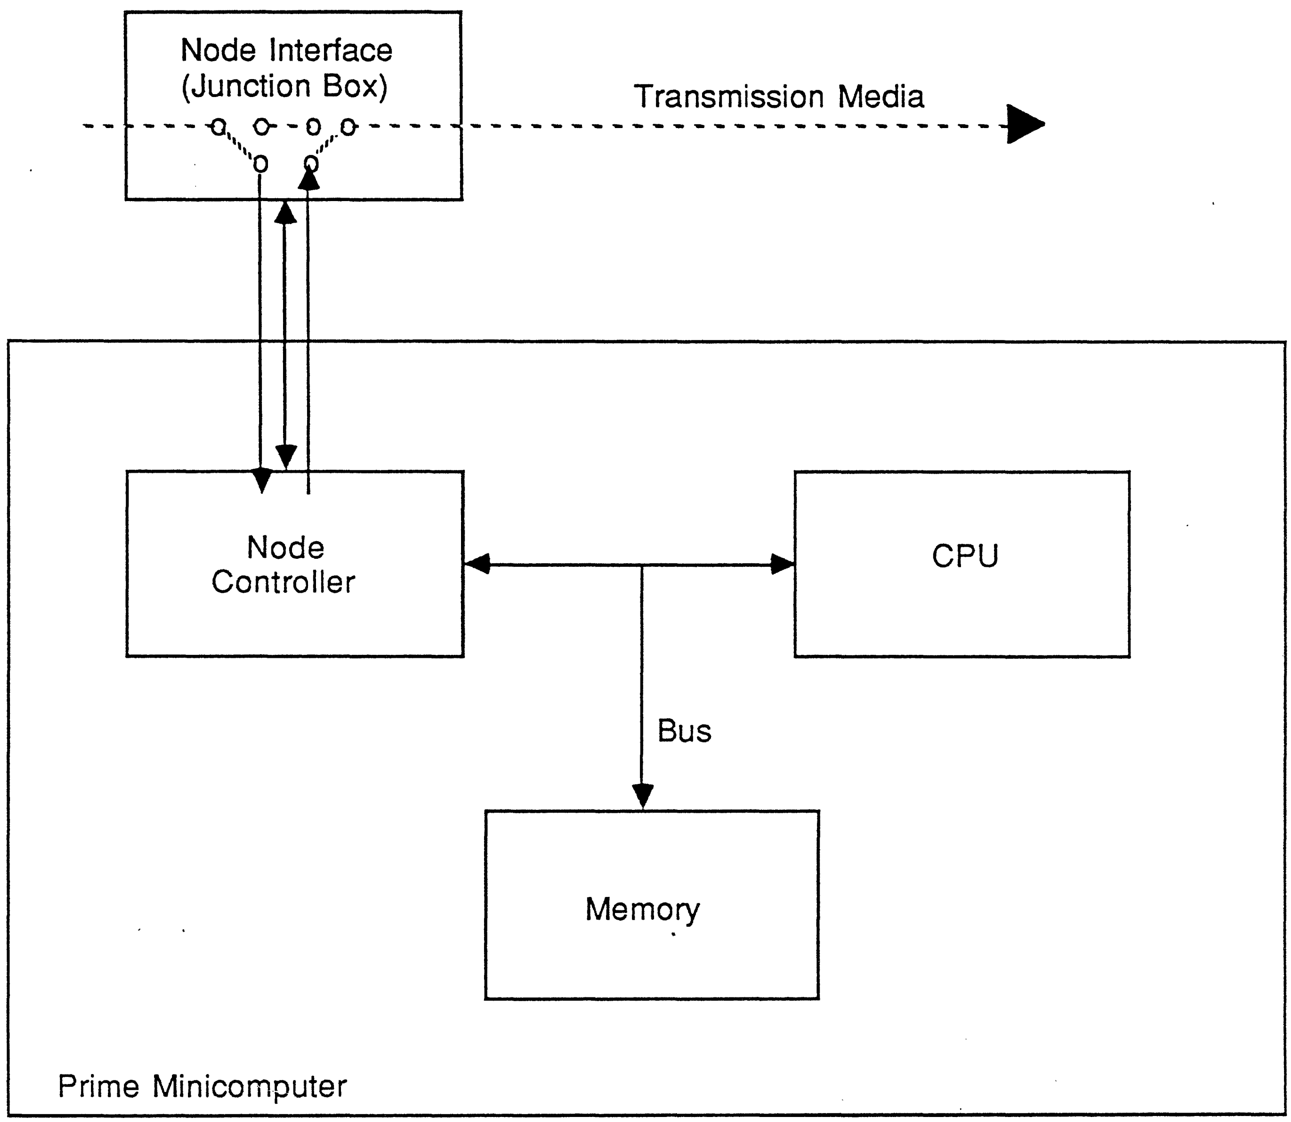

Figure I-45: Ringnet Components

been correctly received.

True to its token-ring-based nature, control of the network is maintained by passing a special bit pattern known as a token around the ring. If a node controller has a packet pending for transmission and detects a "free" token, it may remove the token and put a packet on the ring. Packets can be placed onto and removed from the network only by the transmitting node; during this time the other active nodes operate in a "transceive" mode and are responsible for receiving, retransmitting, and copying messages if necessary.

In terms of error recovery, the loss of a token is detected by a finding that there has been no traffic on the ring for a long period of time. This diagnosis may be made by any controller, in that token management is fully distributed. This condition is usually detected during an unsuccessful attempt to transmit a packet; in such an event, the controller creates a new token and retransmits the packet. If two hosts simultaneously execute this procedure, they may introduce two tokens to the network, although the likelihood of this happening is lessened by the fact that the hosts will wait different periods before creating a new token. Nonetheless, if it does happen, the duplicate tokens will be identified and suppressed.

If a node detects persistent failures, it will disconnect itself from the ring, using the relay in the junction box, and initiate various self-diagnostic procedures. A defective node will stay disconnected, while nodes that pass their own tests will reconnect to the ring. This procedure is intended to eliminate defective nodes.

As with ring-based LANs in general, the decentralized control employed by Ringnet allows autonomous operation of each node on the ring, which avoids the risk of central failure and facilitates smooth system expandability.

4.2.4. Proteon ProNET

Proteon, Inc., of Natick, MA, is the manufactuer of the ProNET line of token-passing, star-shaped ring LANs; the line consists at present of the ProNET-10, ProNET-80, and ProNET-4. ProNET host interfaces have been designed for each particular network, while other network components -- cables, wire centers, etc. -- can be shared by the three networks.

The fact that Proteon's first system, the ProNET-10, was installed at a customer site in June of 1981 shows that Proteon is no latecomer to the technology of ring-based LANs. (It has since installed over 12,000 ProNET nodes). The firm designed its product to resolve incompatibilities between communicating devices, and to be modular in design for ease of network expansion. Also, its star-shaped design eliminates the inherent wiring diagnosis problems of a ring used to wire a building.

ProNET host interface systems are available for UNIBUS, MULTIBUS, VME Bus, Q-Bus, and Universal Bus. These interfaces provide user-transparent connection services with minimal delay across a high-bandwidth LAN; in fact, during host-to-host file transfer, the rate is limited only by the speed of the computer's bus.

The ProNET systems operate over a variety of media, including shielded twisted-pair, fiber-optic cable, coax, or infrared link; a combination of media may be employed within and between networks (see Figure I-46).

With a passive ProNET Wire Center containing the physical ring, a ProNET system is physically arranged as a series of stars (see Figure I-47). Each Wire Center can connect 4, 8, 12, or 16 nodes and provides local expansion connections for adjacent Wire Centers. Via a relay, each node is connected to the ring and operates only when the host interface gives the "join ring" command. Thus a power failure at one node will not bring the ring down, and nodes which have found themselves faulty through the digital or analog loopback test will not join the ring; similarly, nodes will immediately disconnect from the network in the event of a cable break. The Wire Center requires no power except for the small current from the host interface necessary to activiate the relays.

Wire Centers allow easy addition or deletion of nodes without disrupting network operation. The Wire Center also provides automatic bypassing of malfunctioning nodes and facilitates fault isolation.

Wire Centers within the 10-Mbps network may be connected to one another via shielded twisted pair, infrared, coax, or fiber optic cable; the maximum distance between Wire Centers or from host-to-Wire Centers is 480 meters using shielded twisted pair with repeaters. (The maximum unrepeatered distance between a single Wire Center and a host is 240 meters.) ProNET fiber-optic links allow distances of 2.5 kilometers between Wire Centers and/or hosts, while the infrared link provides a 350 meter distance between infrared devices.

Repeaters allow longer distances between nodes by amplifying the ring signal; when used in adjacent Wire Centers, repeaters allow the star-shaped ring to partition into smaller rings upon a cable break.

Fiber-optic links are available for those users requiring security and immunity to RFI (Radio Frequency Interference) and EMI (Electro-Magnetic Interference). Each fiber-optic unit contains a microprocessor to control the system, such that when the ring fails, a node will disconnect from the network. If fiber is used between Wire Centers, the ring will divide into two

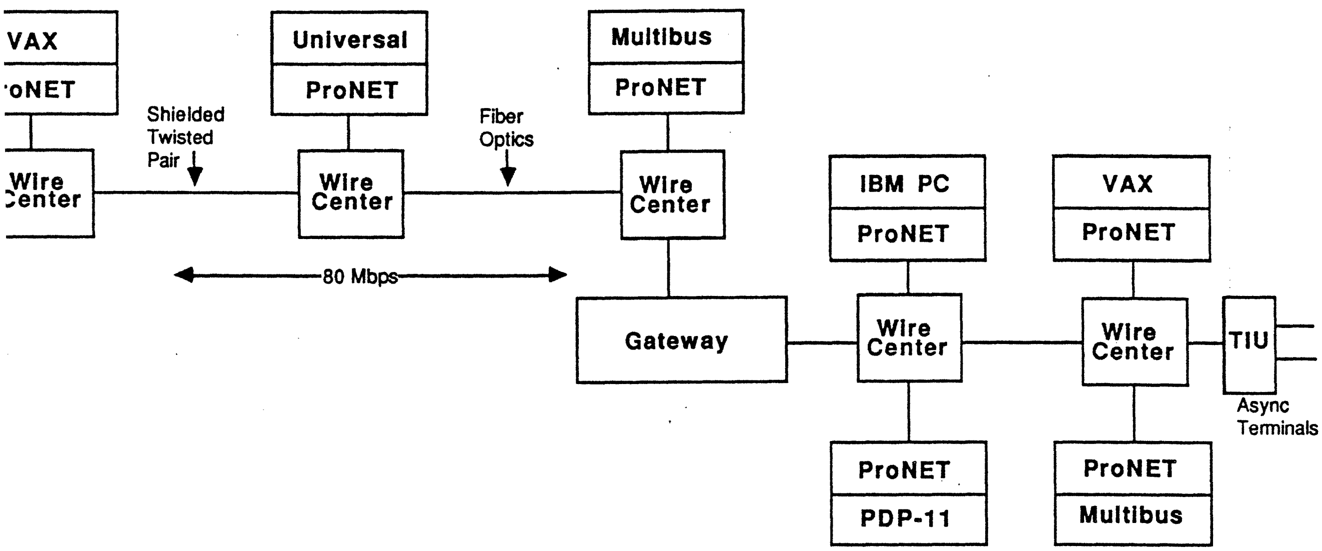

Figure I-46: Typical ProNET Network

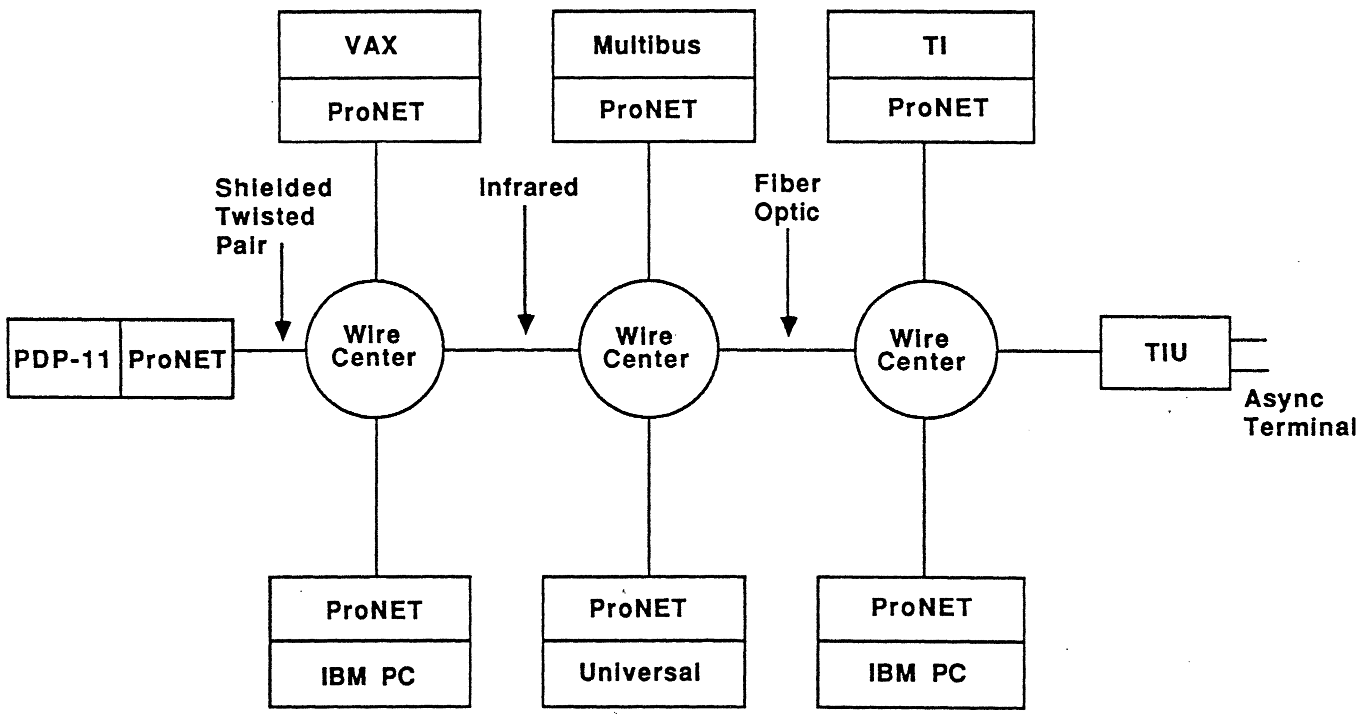

Figure I-47: ProNET System -- Series of Stars functional subrings should the signal fail.

The ProNET-80 employs the same shielded twisted-pair cable as the 10-Mbps system, and provides distances of 100 meters between Wire Centers and/or hosts; fiber-optic links extend this distance to two kilometers.

From each Wire Center relay, connections are made to host interface systems, IBM PC interfaces, or terminal interface units (TIUs); each of these interfaces contains the logic required to interconnect an information-processing device to the ProNET LAN.

Each host interface system consists of two boards: a host-specific board (HSB) and a control (CTL) module. The HSB performs the communications processing for the host, including sequencing, buffering, and controlling; the CTL board manages all the network transmission functions, including modem, serial-to-parallel conversion, repeaters, token management, bit stuffing, and error timeouts. The CTL also activates Wire Center relays and features circuitry for self-testing under program control.

The IBM Personal Computer Interface is a single card which plugs into a large expansion slot of an IBM PC, XT, AT, or compatible and provides a high-performance link between the PC and the ProNET token-ring technology. It is compatible with all other ProNET host interfaces.

The Terminal Interface Units (TIUs) interface asynchronous ASCII terminals to the star-shaped ring network.

The Host Specific Board (HSB) alluded to above serves as the physical interface between a user's host processor or computer and the data link layer; its design is intended to minimize the software burden on the host. ProNET's HSB features a full-duplex direct memory access (DMA): two separate 2046-byte packet buffers, control and status logic, and full-duplex operation provide for concurrent transmitter and receiver functionality.

Proteon offers a system that monitors and manages its ProNET LANs: the Network Monitoring System measures network performance, debugs network software and hardware, and generates long-term records of LAN traffic.

Proteon's newest ring LAN is the ProNET-4, a 4-Mbps IBM-Token-Ring-compatible LAN that offers as an alternative medium of connection an Infrared Communication Link, which can link hosts or Wire Centers up to 350 meters apart. ProNET-4 users can connect to the faster ProNET-10 and ProNET-80 LANs via Proteon gateways to form a mixed-vendor, multi-speed LAN system.

Like the other ProNET LANs, through the UNIBUS and Q-bus interfaces, ProNET-4 links DEC VAXes, MicroVAXes, PDP-11s, and LSI-11s, plus all computers employing MULTIBUS and VMEbus, and IBM PCs and compatibles. ProNET-4 supports PC connectivity to IBM's SNA and works with TCP/IP, XNS, DECnet, NetWare/P, and NETBIOS.

4.2.5. Ungermann-Bass Net/One

Ungermann-Bass, Inc., of Santa Clara, CA, is a longstanding LAN manufacturer whose first local network, Net/One, dates back to July of 1980. It is also a firm which has, since its inception, professed a commitment to "providing networking solutions for all leading local area network standards." Ungermann-Bass has now extended its Net/One functionality to incorporate Token Ring standards, via the Ungermann-Bass Net/One Token Ring Intro/Net System.

The Net/One Token Ring Intro/Net is compatible with the IEEE 802.5 Token-Ring Access Method Specifications and IBM's Token Ring implementation. Designed to complement IBM's Token-Ring offering, the Ungermann-Bass system offers the capability to network non-IBM equipment and IBM products for which an IBM networking solution is unavailable, and provides networking alternatives to existing IBM solutions. The Net/One Token Ring Intro/Net provides connectivity support for asynchronous devices, the IBM 3270 family, and devices based on the IBM PC Bus -- as well as host support from a pool of vendors that includes DEC and IBM.

Hardware components of the Ungermann-Bass Net/One Token-Ring Intro/Net include the NIU-180, the Personal NIU, NIU-74 and NIU-72, the Distributed Wiring Concentrator, and a ring bridge.

The NIU-180 (Network Interface Unit, Model 180) family was expanded by Ungermann-Bass to incorporate Token-Ring capabilities. As a result, up to eight vendor-independent asynchronous devices with RS-232 interfaces can physically and logically connect to an Ungermann-Bass Token-Ring LAN via the NIU-180.

The Personal NIU (Personal Network Interface Unit) provides the board-level connection for the IBM family of PCs onto the Net/One Token-Ring network. The Personal NIU enables the PC to share peripherals attached to other PCs and NIU-180s, as well as establish connections to other hosts. Software with the Personal NIU provides the PC user with file-sharing, print-sharing, 3270 terminal emulation, and PC-to-mainframe file transfer. Distributed applications written to both the NETBIOS and the Microsoft Networks (MS-Net) Session Interface will operate on the Ungermann-Bass Token Ring Network through the Personal NIU.

The NIU-74 and NIU-78 are NIUs designed specifically to connect the IBM 3270 Information System to the Net/One Token Ring

Network. These devices permit users of 3270 system terminals and printers to communicate with IBM 3274 controller via the Net/One LAN. In addition, the NIU-74 provides high-speed access to an IBM mainframe for IBM PCs using the 3270 Personal Connection software program.

The Distributed Wiring Concentrator enables a combination of up to ten 802.5 Network Interface Units -- including the NIU-180, NIU74/78, PC-NIU, Ring Bridge, and/or IBM PC Adapter Cards and their associated devices -- to be connected to the Token Ring Network. The DWC functions in both an active and a passive state: in a passive state, the DWC provides the connection hub of the logical ring, physical star topology. With the addition of a power source, the DWC functions in an active state: it allows the Network Management Console (NMC) to perform problem determination and fault isolation, and, when appropriate, remove stations from the ring. DWCs may be connected in a daisy-chain fashion to increase ring size to the allowable maximum of the media selected.

The Ring Bridge provides the capability to configure multiple rings. The Ring Bridge reinforces the Ungermann-Bass modular approach, in that segments of a network can be isolated to improve problem determination and network security. In addition, the Ring Bridge allows the user to effectively manage load balancing and increase available network bandwidth.

Net/One Token Ring Intro/Net software falls into three general categories: Virtual Connection Service; Personal Connection Services; and Net/one Operating Software and Network Management.

The Virtual Circuit Service (VCS) provides point-to-point communication between two devices or user programs attached to a Net/One LAN. It is implemented by software that routes packets of data along the shared Net/one communications medium. VCS provides terminal device users with features such as: interactive command interpreter; distributed name service; rotoring on assigned names; host password; automatic baud rate detection; idle circuit time-out; automatic device mismatch resolution; and sequencing and error recovery.

The Net/One 3270 Personal Connection is a high-performance local connection to the IBM mainframe from a Token Ring Intro/Net-networked IBM PC or compatible. When running Model I of the 3270 Personal Connection in a Net/One Token Ring Intro/Net environment, the PC emulates an IBM 3278 or 3279 terminal. Emulation is transparent to host software. The PC becomes a universal workstation, able to use multiple resources without having to disconnect from the host.

Network file and printer sharing is accomplished by implementing the MS-Net Version 1.0 software on Net/One: this software allows Token Ring Intro/Net users to share system resources such as data files, programs, hard disks, and printers within a networked MS-DOS or PC-DOS operating system. Since MS-Net is an extension of the MS-DOS/PC-DOS operating system, it is transparent both to network users and to existing MS-DOS/PC-DOS application software.

Token Ring Intro/net users can send programs, reports, memos, and other information to any other user on the network, and users can also spool print files to shared network printers. MS-Net gives all network users the benefits of hard disk storage and expensive printers without requiring these devices at every workstation. Increased productivity is provided through the sharing of information and application programs.

The Net/One Personal Connection includes full support for the IBM NETBIOS interface: almost any program written to run on IBM's PC Network can be used without modification on Intro/Net with MS-Net. In addition, the IBM PC Network Program can be run on the Personal Connection networking hardware. Both IBM and Ungermann-Bass network operating systems can be mixed on the same network, and any workstation can access any server.

Net/One OPS (Net/One Operating Software) includes the network operating kernel, data link protocol software, and network management services. The network management services allow a network administrator to configure, monitor, and debug a network, and to communicate conveniently to all network users. The network management services include: Configure - an easy to use program that allows a network administrator to prepare configuration information for the network NIUs; Download Server (SLS) - the program that downloads configuration files and operational software, such as Virtual Circuit Service, across the network to the NIUs; Broadcast/Multicase Messages - the Network Administrator can send a message to all users, or multicast the message to users belonging to a pre-defined group; and Media and MAC Management - the Media and MAC management program allows the network administrator to overview his/her network configuration, display error data from a ring, display ring condition messages, and, when appropriate, remove or restore a station on a ring.

A sample Net/One Token Ring Intro/Net configuration is depicted in Figure I-48.

All personal computer workstations attached to the Net/One LAN by the Personal NIU must have the following minimum configuration: IBM PC or IBM PC XT with keyboard and display; one available expansion slot for the Personal NIU; 512 Kbytes of memory for maximum system capability; MS-DOS/PC-DOS 3.1 licensed to run on the personal computer; and one double-sided diskette drive.

Approved terminals are: IBM 3278 - Models 2, 3, 4, 5; IBM 3279 (color); IBM 3178; IBM 3179 (color); IBM 3180, all models; ITT Courier 1778; and Telex 078.

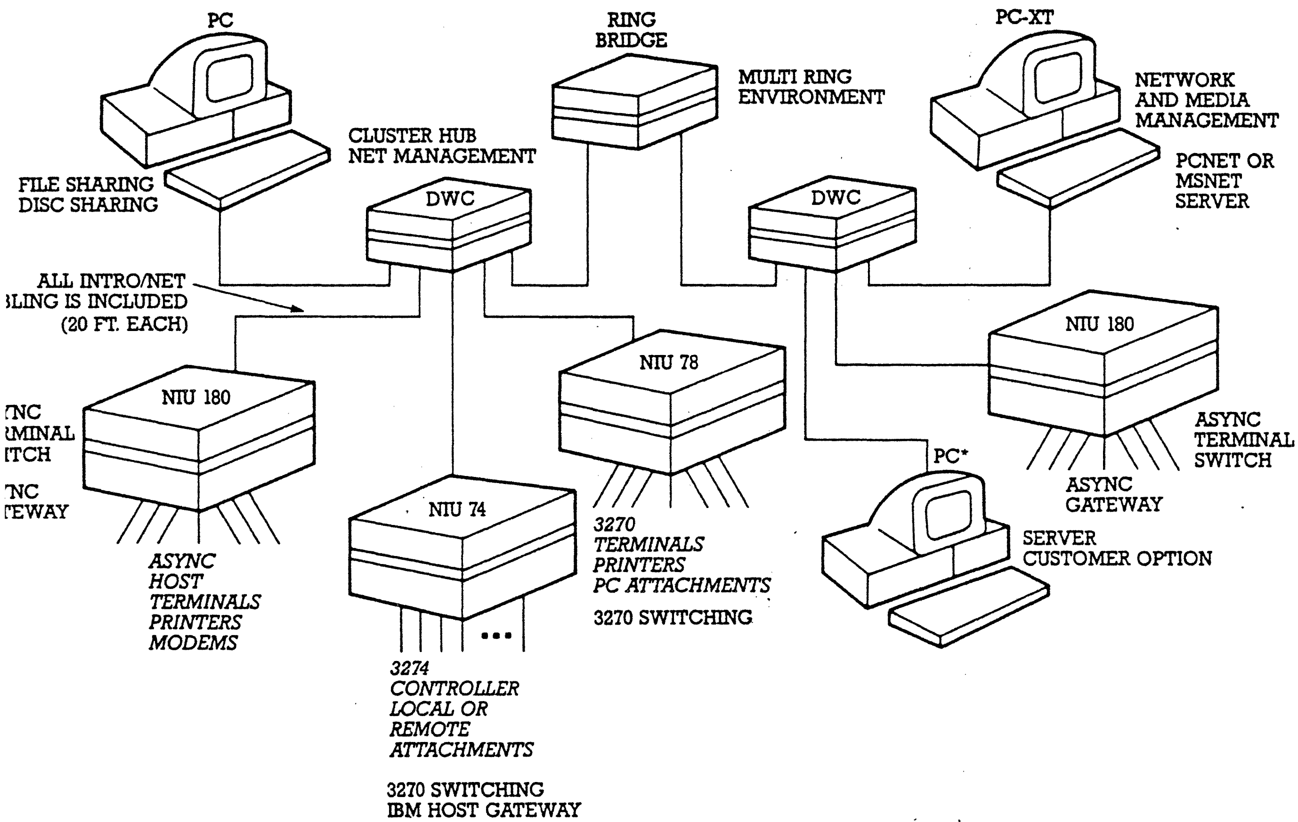

Figure I-48: Typical Net/One Token Ring Intro/Net Configuration

Personal Computer Terminals suitable for use on the Ungermann-Bass Net/One Token Ring Intro/Net are the 3270 PC in Control Unit Terminal (CUT) Mode only, and the following 3278/3279 emulation coax boards and associated control program software: IRMA Board, Forte Board, and IBM PC 3278/79 Emulation Adapter Card.

5. Other Approaches

5.1. Register Insertion Rings

Register insertion rings -- discussed earlier (see section 2.1.3) -- were originally developed for use in systems that integrate speech and voice/data applications but have since found use in distributed processing systems as well.

As was seen earlier, in such a system a ring station includes a register (or buffer) placed in parallel with the ring. During normal operation, the register is shifted out of the ring; the attached device may place a new packet in the register whenever it is free, whereupon the station will wait for an inter-packet gap on the ring. As this gap passes by the station, the register is shifted into the ring and its contents transmitted.

Each station reads the packets circulating on the ring and determines which packets are intended for it; these packets are then passed to the attached device and are either removed from the ring or designated "empty," so that they will be removed by a central monitor station.

In order to guarantee that data received from the ring is not lost during transmission, the station must either buffer incoming packets or transmit only when it detects an incoming empty packet.

In that the destination of a packet cannot be known until the address has been read, the ring is required to include a buffer at least as long as the address.

Generally speaking, all register insertion rings vary in length over a given period of time, as they increase in length as their loads increase, and thus these rings must be kept short to sustain low delay packets. Also, because packets are emptied or removed by the receiving station, the same slot on a register insertion ring may be used more than once per cycle, thus permitting these rings to transfer data at a rate in excess of their raw speed.

5.1.1. SILK

SILK -- System for Integrated Local Communication (in German, "Kommunikation") -- was developed at the Hasler Research Laboratory in Berne, Switzerland. (SILK is a registered trademark of Hasler, Ltd.) It was developed by Hasler, Ltd., specifically to provide both data and voice transmission on a single system.

A preliminary experimental system that used fixed-length shift registers as insertion buffers was tested at Hasler's Research Laboratory and reported on in 1974 by E.R. Hafner, Z. Nenadal, and M. Tschanz. The current version of SILK, which is an enhancement of the early experimental system, was presented in a 1983 paper jointly authored by D.E. Huber, W. Steinlin, and P.J. Wild.

SILK is a coaxial-cable-based LAN which accommodates packetized voice and data traffic at a bit transfer rate of 16.896 Mbps; the effective data rate is approximately 10 Mbps, depending upon the type of traffic it is handling. The loop (ring) is made of either fiber-optic cable or coaxial cable, or a combination of the two.

Each SILK loop (see Figure I-49) may accommodate as many as 150 access devices -- i.e., local blocks (LBLs) incorporating three high-speed registers interlinked by switches and control logic -- at up to 100 meters apart, and to each of these it is possible to connect seven user devices each attached through appropriate data communications equipment (DEC) or -- in the case of Hasler's own digital telephones -- directly; thus a SILK system may consist of as many as 1,050 devices. The local blocks and main blocks (MBLs) manage SILK traffic transmission and reception. While hardwiring is employed in LBLs to speed LAN traffic rates, microprocessor-based equipment is used when software (programmable) control is deemed appropriate by SILK users.

The media access control protocol employed by SILK is, of course, register insertion, which is performed by the LBLs. All user devices on the network transmit their packets whenever they have them ready, enabling the loop to simultaneously handle many packets. These packets are addressed to their intended destinations, where they are subsequently removed from the loop and reassembled as an exact copy of the original transmitted data.

In that packet flow on SILK is unidirectional, the packet stream passing any one point on the loop means that all packets in the stream will travel past that point in moving from their source to their destination; thus traffic on the LAN varies from point to point, because user information only travels along that portion of the loop between its origin and the destination.

With the SILK system, the signals on the ring take the form of minipackets of up to 16 bytes (see Figure I-50), separated by idle bytes.

When a station on the ring accepts a minipacket, it replaces the received minipacket on the ring with an idle minipacket (an idle minipacket is so called because it has an address but includes no data). Because idle minipackets take up valuable space on the ring, they are suppressed by a monitor station dubbed the main block (the MBL discussed above). It is the main

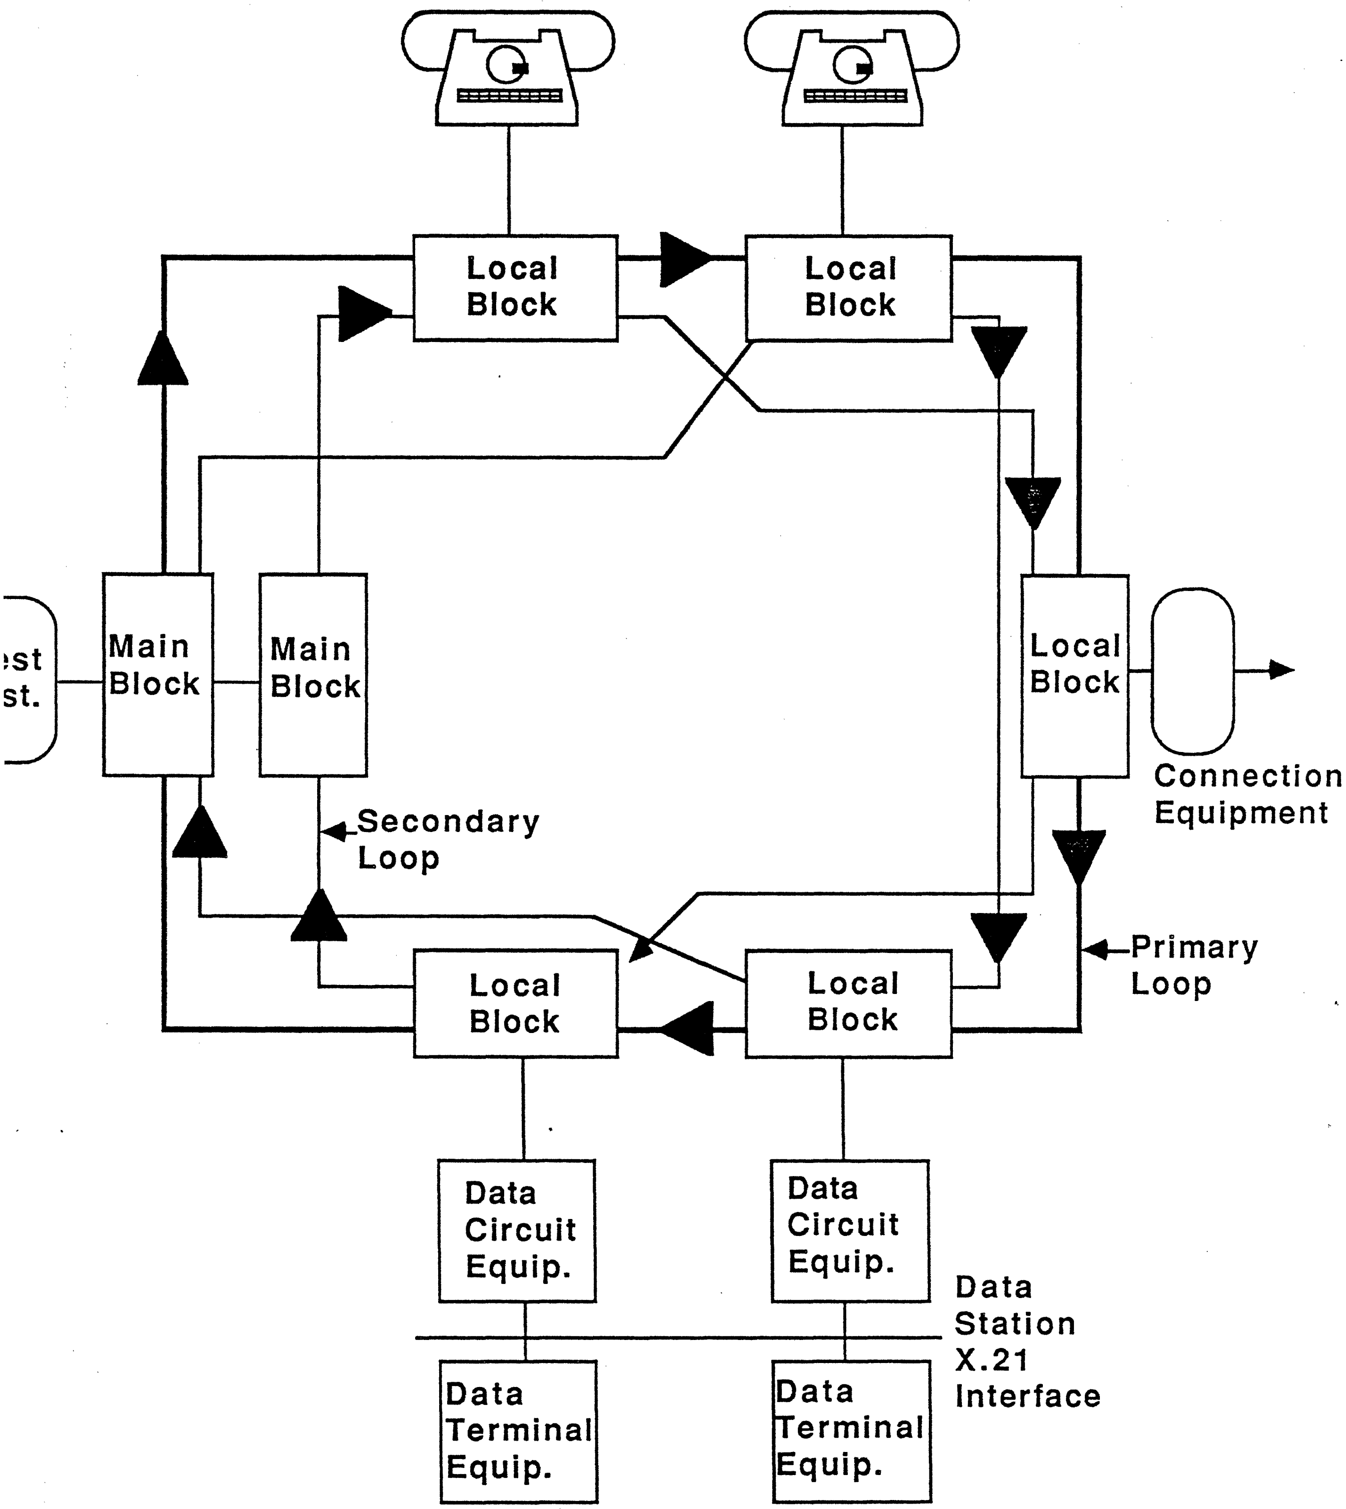

Figure I-49: Typical SILK Register Insertion Ring Configuration

Header	Destination Address	Data
2 Bytes	3 Bytes	0-11 Bytes

Figure I-50: SILK Minipacket Format

block that maintains the integrity of the ring by preserving bit synchronization, removing minipackets that have completely circled the loop, and sending test minipackets to inspect ring integrity and ring station (i.e., local block -- LBL) functionality. This MBL also accepts error report packets from LBLs and supports local diagnostic and alarm equipment.

The actual register insertion method employed by'the SILK system is as follows.

To transmit over the ring, the DCE (Data Circuit Equipment) places a correctly formatted minipacket in the output register. As soon as the control logic detects either idle bytes or the end of a packet at the output switch, it connects the output register to the line and transmits its contents. (During this period of time, a transient packet may be building up in the shift register.) When the output register is empty, the output switch will connect the shift register to the ring output so that this minipacket can be sent. In that minipackets are of variable length and are placed in the output register at the convenience of the DCE, transient packets are subject to varying delays.

Incoming packets are placed into the shift and receiving registers one byte at a time. The address is read as packets enter the shift register and is compared with the address of the attached devices; if the address matches, the minipacket is permitted to build up in the receiving register for transfer to the attached device. The ring is disconnected from the shift register and idle bytes are fed to the shift register, but they are not stored by it. Thus the incoming minipacket is truncated into an idle packet and transmitted in that form.

Whenever the shift register is empty and there is no data in the output register, the system transmits idle bytes on the ring.

For telephony purposes, SILK digitizes data at a rate of 64 Kbytes, following the usual pulse code modulation (PCM) scheme of

8,000 eight-bit samples per second; samples are sent out in groups of four per minipacket, and no packets are sent during silent periods.

For data transmission purposes, DCEs are available that support CCITT at speeds up to 96 Kbps, CCITT X.21 transmission with HDLC framing and error control, datagram capability, and multiplexing of up to eight V.24 circuits in one X.21 circuit. An electronic mail system has also been implemented on SILK.

Hasler Ltd. of Great Britain -- the manufacturer of SILK -- sells the system as a LAN product only, choosing not to become involved in software development matters or attached hardware concerns.

One special property of SILK is its lack of susceptibility to a single point of failure disabling the entire system: the reason for this is that SILK employs a braided system of interconnection: in this approach, a continuous loop may be maintained via the use of secondary and tertiary cable paths which permit bypassing of failed points.

5.1.2. Clearway

Real Time Developments Ltd. is a United Kingdom company that manufactures a ring-based LAN called Clearway. The Clearway LAN interconnects devices with standards RS-232-C asynchronous interfaces over coaxial cable at a maximum distance of 800 meters. Real Time Developments says Clearway's network access is analogous to register insertion and is performed by junction boxes.

These junction boxes link user devices to the ring; they feature buffer storage for up to 1,500 characters, accommodating the varying baud rates of the communicating devices at both ends of an interconnection. Each junction box collects a maximum of 33 characters and inserts a destination address and a checksum at the front and rear of the packet, repsectively. The packet is then transmitted along the coaxial cable, being forwarded along by each subsequent junction incurring a two-character delay until it reaches its destination. (Clearway's transmission rate is said to be approximately 35 Kbaud.) The destination on the ring checks the packet and acknowledges receipt of it. Clearway employs a stop-and-wait protocol wherein each packet that is received is acknowledged, to make certain that all packets are received in the correct order.

While in theory up to 99 units may be addressed on Clearway, Real Time Developments recommends a maximum of 50 units. In addition, each attached unit is informed by Clearway as to whether it can establish a direct connection with a destination, or if it must proceed to do so via a third unit. Clearway is intended for small configurations of about 12 VDUs (Video Display

Units) each operating at 9.6 Kbaud, but not requiring continuous data.

Clearway accommodates two alternative forms of wiring: users can either establish a series of junction boxes to which devices can be attached as necessary, or daisy-chain each of the nodes together.

5.1.3. DDLCN

The DDLCN (Distributed Double-Loop Computer Network) traces its origin to Ohio State University, where it was developed as a research tool by M.T. Lin, J.J. Wolf, and coworkers in 1978. It is, in essence, an extension of an earlier network -- the Distributed Loop Computer Network, also worked on by Liu -- which also employed register insertion.

In that DDLCN is primarily intended for intercomputer communication, it allows much larger packets than does SILK. On DDLCN, packets are removed by the destination station, and this means that the station must have two buffers. The ring is double (see Figure I-51), in order to provide a transmission path in the event of the failure of one segment of the ring.

The elemental configuration of a ring station is shown in Figure I-52. At the station, packets are received into the input buffers, and their addresses are checked by the microprocessor; packets for this station are then removed from the ring and passed along to the attached device. In turn, other packets are passed along to the output buffer -- when it is free -- and transmitted to the line. Also, the output buffers may be filled by the attached system.

For purposes of efficiency, a station on the loop does not just send a packet by the physically shortest path, because this might entail involvement of heavily congested stations; rather the station selects the best route dynamically, via means of traffic flow measurement.

The DDLCN places great emphasis on its ability to recover from the failure of individual stations or sections of the loop. By means of special circuits, each separate ring segment can be made to carry signals in the reverse of its usual direction. Also, each output buffer can be connected to the input buffer on the other loop, thereby permitting it to send its contents over the other ring in the event that its own outgoing segment fails.

Detection of failures is made possible by each station's incorporation of a bit map that indicates the present status of every segment of both loops; this map is updated on the basis of status change messages broadcast by stations that detect a change.

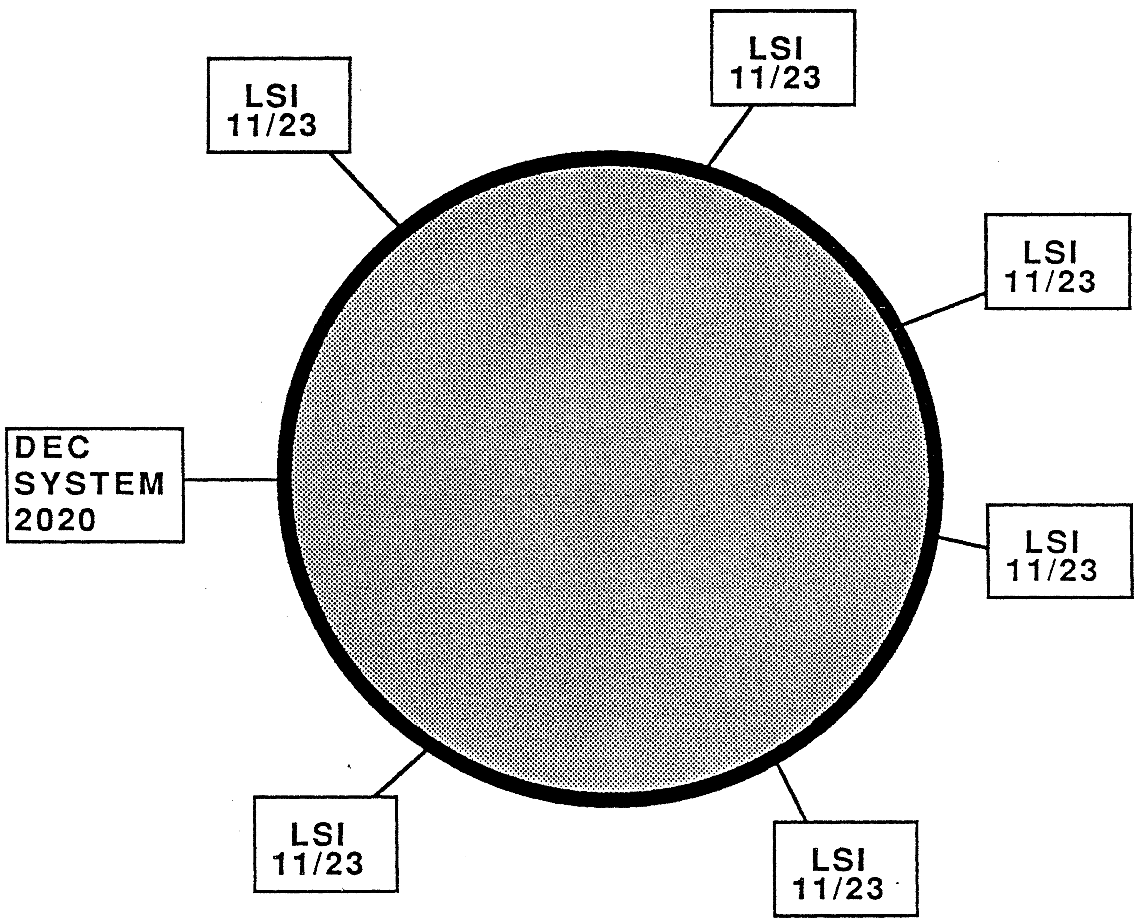

Figure I-51: Basic DDLCN Configuration

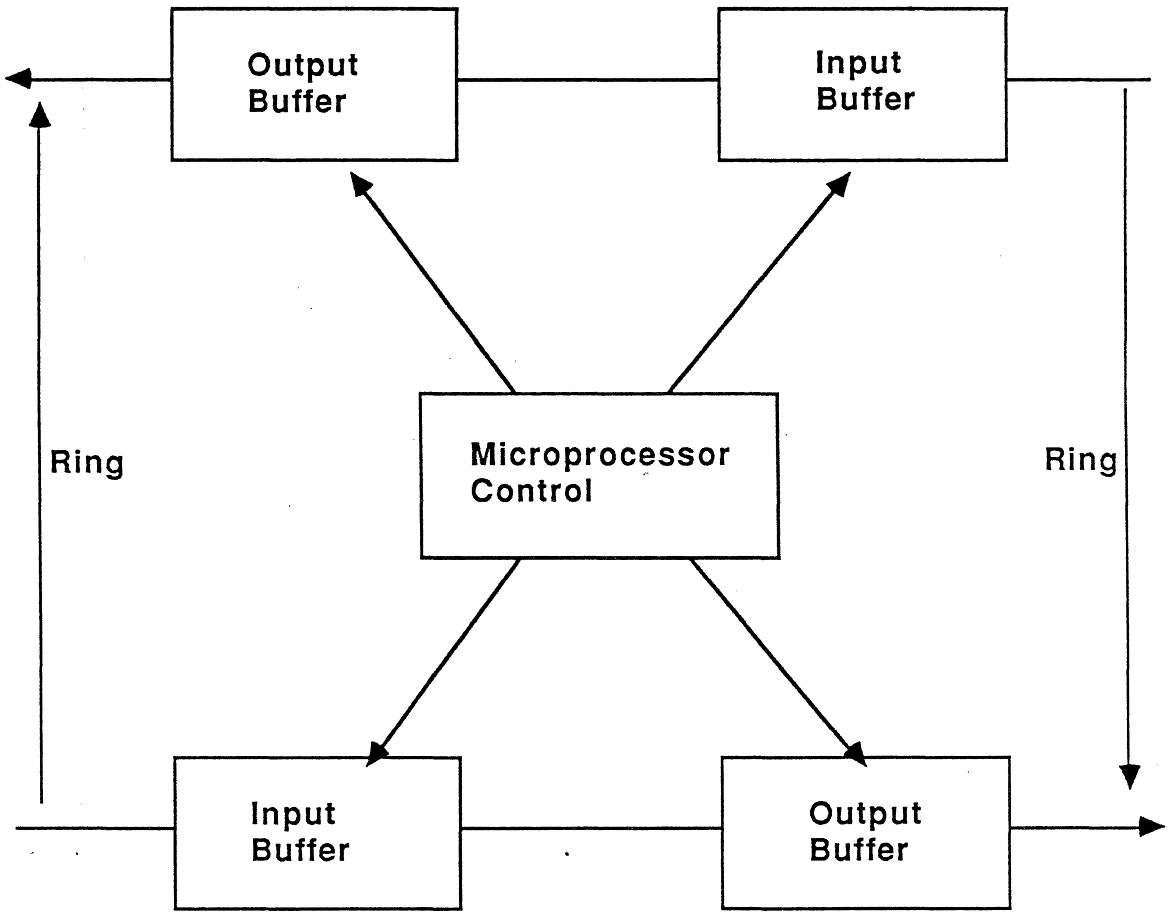

Figure I-52: DDLCN Ring Station

Impressively enough, DDLCN is able to provide degraded service even after several loop segments have failed, and can survive the combined failure of both rings in one segment and segments elsewhere, by means of operating its links in the half-duplex mode.

A distributed operating system for DDLCN, MIKE (Multicomputer Integrator Kernel), leaves the operating systems of the attached computers largely unaltered and runs primarily in the loop station microprocessors. It provides a system-wide view of all network resources and also supplies protocols for cooperation among the attached computers. (The first implementation of DDLCN used twisted pairs to interlink six PDP-11/23s and one DEC System 20.)

DDLCN also employs the concept of a guardian, i.e., a permanent process that provides other processes with controlled access to certain resources. Each attached computer's operating system is treated as the guardian of that computer.

5.2. Fiber-Optic-Based Rings -- ODR-1, ODR-2

Syscon Ltd. of the United Kingdom produces a fiber-optic based LAN called ODR-1. This network is made up of one controller and up to 15 slaves and employs a time division multiplexing technique with a self-synchronizing, self-clocking code developed by the firm. In this code, the width of a pulse indicates whether a bit or data bit is being used for synchronizaton purposes.

The ODR-1 has RS-232-C compatibility, and the slave units are designed as DCE (data communication equipment), which enables them to interface directly to terminals. Designed as DTE (data terminal equipment), the controller interfaces directly to computer I/O ports, which results in an RS-232-C transparent link.

The fiber-optic cable used in ODR-1 is a plastic cable that can provide up to a 24-meter link at 10 MHz; this cable length can be increased up to 40 meters, but the bandwidth must be reduced accordingly. The system uses low-cost electrical jack plugs rather than optical connectors. CMOS-technology is employed in the LAN.

All units on ODR-1 must function in order for the ring to be operational. Recognizing the need for fault-tolerance, Syscon also offers ODR-2, a dual-channel, full-duplex LAN capable of withstanding the failure of a single station.

6. Conclusions

There are, then, three elemental types of ring-based local area networks: slotted (also known as empty-slot), token-passing, and register insertion. While ring network technology is a relative newcomer to the United States, ring-based networks are the prevalent LAN type in England and Europe and have been for quite some time. Much of the impetus for European ring-network development is traceable to the initial success of the Cambridge Ring, while the present swing in momentum in America toward token-passing ring networks is primarily attributable to IBM's interest in rings and its introduction of the Token-Ring Network.

As has been seen, the benefits of ring-based local networking are heavily based upon its use of point-to-point communications links. In such a system, the transmitted signal is regenerated at each node, thereby lessening the chances of transmission errors and giving it a greater data integrity, at high transmission rates, than broadband -- and greater distance.

Also, it is easier to isolate and recover from an error on a ring than on a bus or tree: failed nodes on a ring can be bypassed. The throughput on a ring is greater than that of a bus or tree, and it avoids the problem of duplicate addresses.

Additionally, ring networks accommodate optical fiber links, which yield very high data rates and desirable EMI characteristics. The use of fiber-optic cable in ring networks allows longer transmission links with fewer repeaters than does coax cable.

The slotted-ring (or empty-slot) ring, best embodied in the Cambridge Ring LAN, has gained widespread acceptance in England and Europe because of its simplicity: interaction on such a ring at each node is minimized, which yields high reliability. However, slotted-ring LANs are wasteful of bandwidth (e.g., in the Cambridge Ring each frame contains only 16 bits of actual data out of a total of 37 bits). Also, each station can send only one frame per round-trip ring time, and thus if only one or a few stations have frames to transmit, many slots will circulate in an empty state.

Token-passing on a ring network is the most popular ring-control technique in the United States, due in great part to IBM's espousal of it in its Token-Ring Network; also contributing to its popularity is the fact that it is the only ring access method picked for standardization by the IEEE 802 Local Network Standards Committee.

The fairest of all access schemes, use of a token guarantees that only one station at a time may transmit -- and thus does away with the contention familiar to Ethernet. Also, the constant circulation of a token ensures continuous monitoring of a system.

Also because of its point-to-point nature, the failure of one station on a token-ring LAN will not bring down the entire system (i.e., failed nodes can be bypassed until repaired). Further, token-ring LANs feature points of centralized reconfiguration/repair via their use of distribution panels; offer the flexibility to mix media and transmission technologies; and allow prewiring of buildings and use of wall plugs because they do not require the immediate juxtaposition of transceivers and transmission media.

The register-insertion ring technology offers the most efficient utilization of a ring of the three methods of ring access discussed in this report, in that it allows both variable-length frames and (for purposes of efficient use of bandwidth) multiple frames on the ring.

Working against register insertion are its requirement that each frame be delayed at each node by the length of its address field for purposes of address field recognition, and the fact that its purge mechanism requires recognition of an address prior to frame removal (e.g., damaged address fields could then circulate indefinitely).

The token-ring network is currently taking giant strides toward becoming the dominant type of local area network within the United States. Bolstered in its increasing usage by its almost 100% utilization rate under heavy traffic conditions and IBM and IEEE 802.5's push for a token-ring standard, the token-ring network -- particularly as embodied in the IBM Token-Ring Network -- will likely become the driving force within LAN technology for the balance of this century. The fact that prominent LAN manufacturers such as Ungermann-Bass, Texas Instruments, Excelan, Bridge, Nestar, and 3Com have already jumped on the IBM Token-Ring Network -- by announcing similar or complementary products -- serves but to point out the direction in which the entire LAN industry will soon move. The token-ring network has become the ring-based technology of choice in America, and with its champion IBM having international clout as well, it could soon become the LAN technology of choice the world over.

SECTION II

A Closer Look at the Key Systems

1. Apollo Computer DOMAIN

Apollo Computer, Inc., was formed early in 1980 by seven highly experienced individuals from Prime, Digital Equipment, and Data General, and announced the DOMAIN system by mid-November. DOMAIN, Distributed Operating Multi-Access Interactive Network, is an integrated hardware/software data access system based on a high-performance local network. Intended for applications involving large amounts of program development, design activities, and calculations or "number crunching," DOMAIN offers performance not found in multi-terminal time-sharing networks and resource sharing capabilities not found in dedicated minicomputer systems. DOMAIN appear to have borrowed heavily from the network techniques developed under Dave Nelson at Prime, software concepts from Bell Lab's UNIX, and display techniques from Xerox's ALTO. Initial shipments were made to Harvard, MIT, Brown University, and the Rutherford Labs in Cambridge, England.

1.1. Network Structure

The DOMAIN system, as shown in Figure II-1, is based upon a 12 million bits per second token-based ring network. Although coaxial cable is the medium shipped to date, the essentially point-to-point communications structure of the ring will allow fiber optic cable, microwave links, etc., to be intermixed.

The single token, a specific bit pattern (011111100), circulates continuously around the ring. When a specific node desires to use the network, it must first recognize the token and then change one bit to obtain exclusive network access. This is easily accomplished in the 1 1/2 bit time delay for a message passing through a node. Conventional bit stuffing techniques are used to ensure the uniqueness of the token and other special characters (Frame, Message Headers, and Message Separator). Token replacement at the end of the message is Apollo's method of preventing network "hogging."

The message format consists of the following fields:

1. Frame Character
2. Destination Node Address
3. Source Node Address
4. Header Information (0-1018 bytes)
5. Data (0-1024) bytes
6. Cyclic Redundancy Check Code
7. Acknowledgement
8. Cyclic Redundancy Check Code
9. Token

One bit of the Destination Address is used for early acknowledgement. If the destinaion node is listening when a message arrives, that bit is set. The first CRC field is used for the message proper; i.e., destination address through data

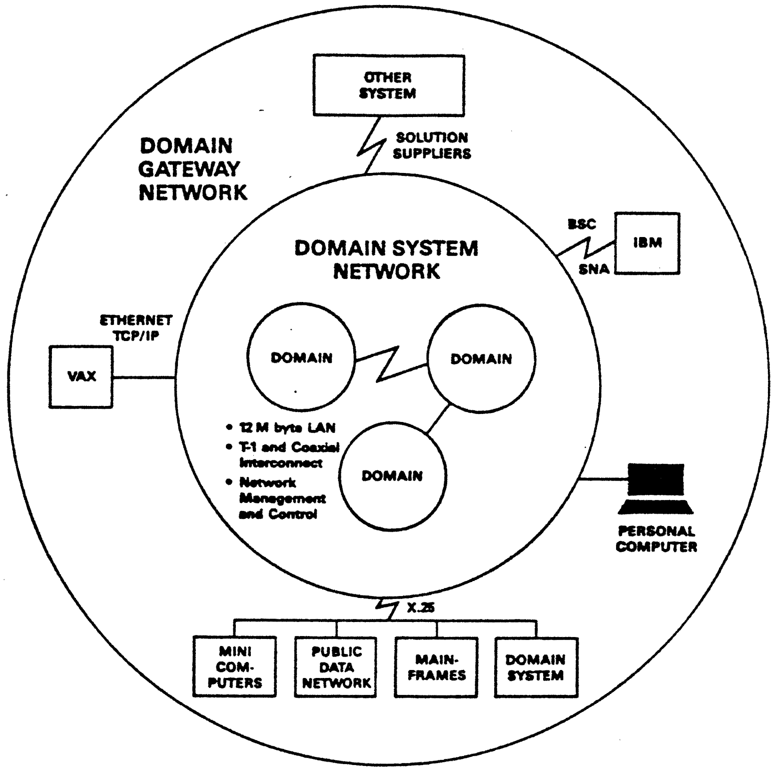

Figure II-1: Apollo DOMAIN System

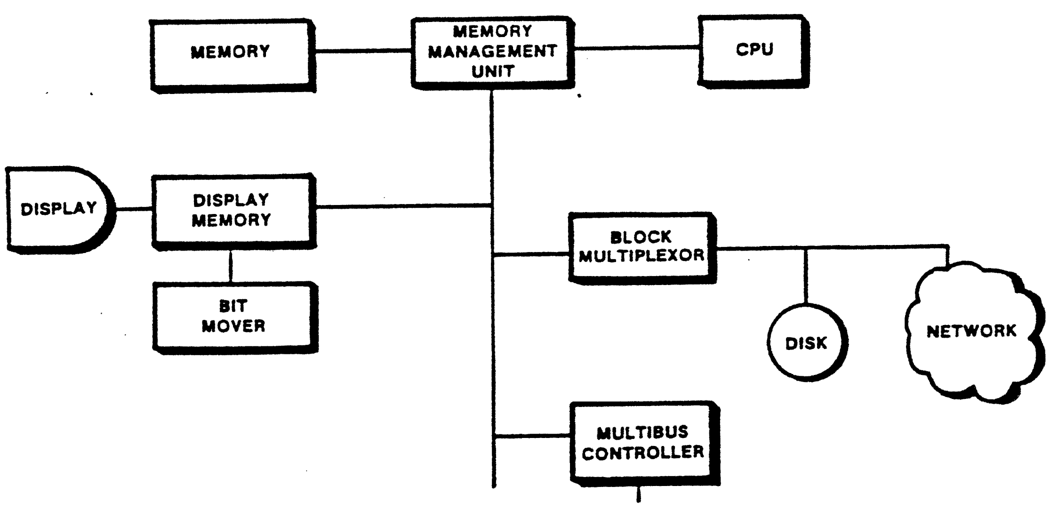

Figure II-2: DOMAIN Node Organization

fields. The second CRC field checks the Acknowledgement. The Message Header is used to indicate the start of the Header information and the Message Separator for the start of the Data. Total fixed overhead for every message is 22 bytes.

The company now offers a Peripheral to Node Adapter, to connect a group of high speed devices, like 300 or 600 MB disks, to the network for independent access from any node. It will not include the full range of capabilities as a computational node but will have its own operating system.

1.2. Node Structure

The standard computational node structure shown in Figure II-2 consists of four major parts: the central processing unit, the memory management unit, the I/O system, and the display system. The CPU of Apollo's low-end systems is built from dual Motorola MC68000 microprocessor chips, with one of the chips dedicated solely to handling page faults. Although these chips are internally organized around a 32-bit architecture, the external bus is only 16 bits wide. The high end systems (DN460/DN660) use a proprietary chip design which supports full 32-bit architecture.

The memory management unit is used to translate 24 bit process virtual address spaces into 22 bit physical memory address spaces. The memory is organized into 1/2 megabyte units and is available with full error correction code protection in sizes up to 3 1/2 megabyte. This limit is set by the density (16K) of the memory chip currently used in the system.

The block multiplexer portion of the I/O system handles all disk and network traffic. Designed to maximize node-to-node processing across the network, it allows packet transfer at near memory speeds and without superfluous buffering. Full bandwidth access to primary memory is available to the network, with all other block transfers, such as the 34 Mb Winchester disk, disallowed. A Peripheral Node Adapter is available to connect line printers, magnetic tapes, and other devices to the node. Because Apollo uses the Intel Multibus for this adapter, hardware was easily added.

The display system is one of the key features of DOMAIN. A separate autonomous 1/8 megabyte bit map memory, organized into a square array of 1024 bits per side, is used to refresh an 800 by 1024 bit map CRT. The remaining display memory area is used as temporary storage for character font tables. Rectangular areas can be moved by the 32 Mbps hardware bit mover from any one place on the screen to any other place on the screen. The high-resolution CRT terminal can display the output from multiple programs concurrently through multiple windows. The display manager allows the user to independently control what is physically shown through the use of function keys on the user keyboard and optional mouse or touchpad.

The extremely powerful operating system (AEGIS) for DOMAIN is used in controlling the mapping of object address spaces into process address spaces, translation of process virtual address spaces into physical memory address spaces, transfer of pages of physical memory into and out of the memory to either local or remote disks, and relating of object onto disk data blocks. AEGIS supports FORTRAN 77, PASCAL, and C and LISP in addition to an editor. It strongly supports the new "Software Tools" package of standardized program modules (Kernigan and Plauger, Software Tools, Addison Wesley publishers) which is gaining considerable attention as a fresh approach to simplified programming. Finally, the network wide operating system coordinates the user's access to network-wide facilities in a location transparent manner. The DOMAIN system software environment is shown in Figure II-3.

AEGIS is an object-oriented operating system in that it treats each element in the system environment--programs, data files, records, peripherals, etc.--as a unique object. The system assigns to each object a unique 64 bit identifier. This identifier code allows the assignment of unique names to every object every created on any DOMAIN network.

1.3. Distributed Data Management

DOMAIN Distributed Data Management, D3M, allows users to organize and access information located anywhere in a DOMAIN network. Users may combine whole or partial views of many individual databases into a single, logical database for both query and update purposes. D3M integrates the runtime efficiency of a CODASYL-compliant design with the personal productivity advantages of a relational interface to span a spectrum of data management applications from simple, file drawer chores to CAD/CAM, engineering, scientific and software development applications.

D3M is optimized for performance in the networking environment, taking full advantage of DOMAIN's distributed file architecture, large address space, high resolution graphics, and high-speed network communications for its operation. Because D3M areas (data files) can be located anywhere throughout the DOMAIN network, users avoid time-consuming bottlenecks in competing for the centralized disk, CPU and memory resources of conventional mainframe computers. D3M's distributed data structures also simplify system expansion and maintenance because the database can grow dynamically - without the need to unload or reload whenever a change is made.

D3M, as a family of software components, provides users with all the database management tools they need to create, maintain and update both small and large databases. These components

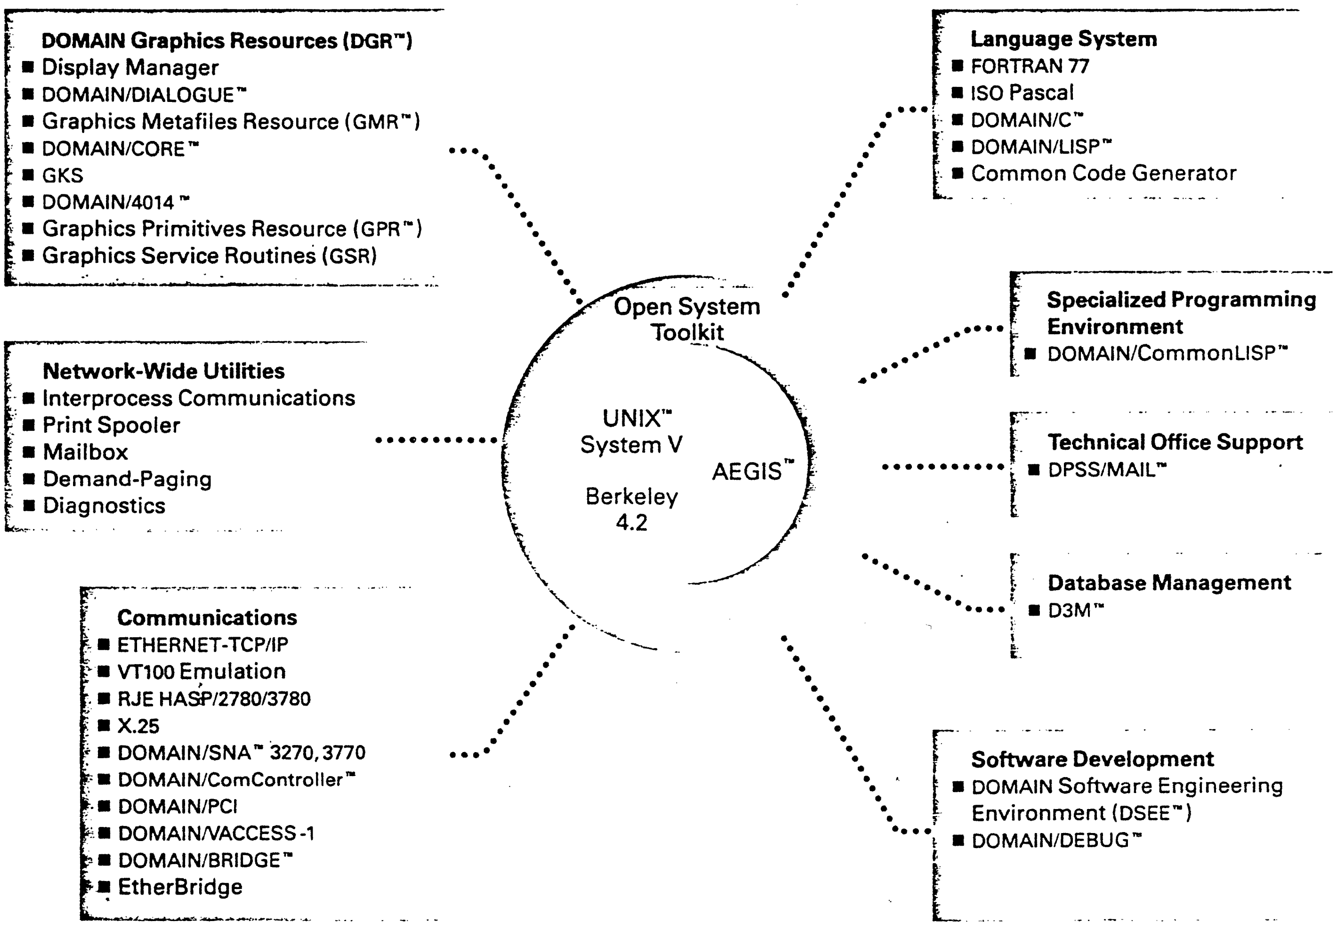

Figure II-3: DOMAIN System Software Environment include:

- D3M/Dataview, a query/update language that provides easy-to-use, relational capabilities for both queries and forms-oriented updates.

- D3M/Describe, a fully interactive database description tool, which can be driven using either forms or commands.

- D3M/Unite, an aggregate schema compiler to create logical combinations of multiple databases located anywhere in a DOMAIN local-area network.

- D3M/Formatter, a complete report writing package specifically tailored for the non-programmer.

- D3M Runtime Library, resides in the shared virtual memory with the rest of the DOMAIN distributed operating system. It is bound to user programs at execution time to provide D3M services.

- Schema and Subschema Compilers which process the CODASYL-standard data description language to generate database descriptions.

- Database Maintenance Utilities, such as Index, Collect Free-space, and Initialize Diskspace, which work on an on-going basis with the shared runtime library.

1.4. DOMAIN System Nodes

The DN300 Computational System Nodes includes a high performance 32-bit VLSI processor, memory management unit, support for up to 4 Mbytes of main memory with ECC, interface to the 12 Mbps, token-passing local area DOMAIN network, 3 asynchronous RS232-C I/O ports, integrated high resolution bit-mapped graphics display with detachable keyboard and locating device, and license to use Aegis network-wide virtual memory operating system with display manager software, font editor, graphics primitives, high-level language debugger, software support for IBM 3270 and HASP communications (requires outboard hardware devices), and network management utilities. The DN320 Computational Node maintains these features, but adds special floating point hardware.

In addition to the basic computational nodes, Apollo recently announced two new high-end processors, DN460 and DN660, claimed to offer roughly the same performance as DEC's VAX-11/780 superminicomputer. The processors are identical to one another, except the DN460 offers a monochrome display while the DN660 features a color display.

Though both processors mark a departure from the firm's use of Motorola, Inc.'s 68000 microprocessor chip in favor of its own proprietary 32-bit chip, the proprietary chip is compatible with 68000-based applications. It is the chip that boosts the systems' internal performance to the equivalent of the VAX-11/780. The processors perform at approximately 1.2 million instructions per second -- about a three-fold increase over Apollo's earlier high-end processor, the DN420.

The two processors provide up to 4 Mbytes of main memory and a virtual address space of up to 256 Mbytes. Each is housed in a 10-slot chassis; an optional five-slot chassis is available as an add-on for users who want to use Intel Multibus boards.

The DN460, with 1 Mbyte of main memory, a 19" 800- by 1,024-pixel bit-mapped monochromatic graphics display, an integrated hardware floating-point processor, three asynchronous I/O ports, detachable keyboard, local area network interface, and a license to use the firm's networkwide virtual memory operating system, costs $39,500; the DN660, with 2 Mbytes of main memory, a 19" 1,024- by 1,024-pixel bit-mapped color graphics display, 1 Mbyte of dedicated display memory, an integrated floating-point processor, detachable keyboard, local area network interface, and a license to use the firm's operating system, is priced at $39,500.

1.5. DN550 Color Graphics Workstation

The DN550 is a powerful, 32-bit color graphics workstation in the $30,000 range. It can accommodate a broad array of applications such as mechanical design, integrated circuit design, and logic simulation. Another model, the RM550, is a modular version of the DN550 suitable for customers who have diverse packaging requirements, such as the Value-Added Supplier (VAS).

The node is compatible with Apollo's entire family of workstations and software, and can run over 250 existing scientific and engineering application packages from over 80 third-party vendors. The DN550 is designed with full 32-bit internal architecture. The dedicated, bit-sliced graphics processor, coupled with extensive microcode, provides fast 2D operations, and allows greater than 10,000 transformed, clipped vectors per second throughput.

Based on a Motorola 68010 processor, the DN550 features a dedicated, bit-sliced graphics processor, a 1024 by 800 resolution color monitor, and a low-profile detachable keyboard with optional touchpad or mouse. The system offers up to 3 Mbytes of main memory and up to 2 Mbytes of double-buffered display memory capable of displaying up to 256 colors simultaneously. The DN550 provides flicker-free color graphics, a multiwindow, multitasking environment allowing up to 24 concurrent processes with 16 Mbytes of virtual address space per process, and an integral interface to Apollo's 12 Mbps DOMAIN local area network.

An OEM version of the DN550, the RM550, is designed specifically to suit the diverse packaging requirements of the value-added supplier (VAS), also providing industry standards such as a MULTIBUS peripheral adapter, and an optional UNIX System III operating environment with Berkeley extensions. The RM550 incorporates the standard DOMAIN local area network, and offers extensive off-net communications using a variety of standard protocols such as VT100 emulation, X.25, and an Ethernet gateway, allowing communication with a wide range of other networks, mainframes, and superminis.

1.6. DOMAIN-Ethernet Products

In June 1985, Apollo demonstrated the compatibility of its UNIX software implementation -- the DOMAIN/IX operating system -- with other vendors' superminicomputers and workstations. The compatibility was demonstrated over Apollo's enhanced version of the DOMAIN System's Ethernet gateway, also unveiled at that time.

DOMAIN/IX software is a twin port of the UNIX operating system to a professional 32 bit workstation. It allows Berkeley 4.2 and System V UNIX standards to operate as co-resident operating systems on Apollo's DOMAIN workstations, so users can run either standard, or both simultaneously, on the same workstation.

DOMAIN/IX software runs on all Apollo servers and workstations and includes a version of TCP/IP that supports Berkeley 4.2 communications on a local DOMAIN ring.

The price of the full DOMAIN/IX software (both Berkeley 4.2 and System V UNIX standards) is $425 per node and $9,100 per site (up to 100 nodes). Apollo's implementations of Berkeley 4.2 or System V UNIX can be purchased separately for $300 per node and $6,500 per site.

The enhanced version of the DOMAIN to Ethernet gateway includes support for Berkeley 4.2 network extensions and performance upgrades. The gateway consists of a transceiver and a floppy diskette with TCP/IP protocol software.

The transceiver can reside on Apollo's DSP80A servers, DSP160 computational servers, and on the DN460, DN550, and DN660 workstations.

Network enhancements include support for the following Berkeley 4.2 capabilities on either side of the DOMAIN System's Ethernet gateway: bi-directional remote log-in, interprocess communications, remote command execution, and remote file transfer among other users connected to Berkeley 4.2. The enhancements also allow any user on the DOMAIN-Ethernet to list all users logged into a remote Berkeley 4.2 host. With the enhancements, Berkeley 4.2 users outside the local DOMAIN ring can gain transparent access to files and compatible applications on any DOMAIN/IX-TCP/IP node.

Like the previous version of the product, one gateway services a community of nodes, including diskless nodes. The enhancement distributes TCP and FTP/TELNET protocol services to DOMAIN/IX nodes. In the earlier version, only IP resided on the gateway. Now that TCP will be distributed to the nodes, users can expect at least twice the throughput of the earlier gateway.

The DOMAIN to Ethernet gateway provides access to the DEC VAX supermini. The DOMAIN System also supports the X.25 gateway and RJE communications with IBM host computers. All DOMAIN workstations have RS-232-C ports and provide VT100 terminal emulation.

The price of DOMAIN System's Ethernet gateway is $3,500 per gateway.

In March 1986 Apollo introduced an internetwork system that forms an Ethernet-based link between DOMAIN network rings.

EtherBridge, as the new product is called, lets technical organizations build a community of networks, linking multiple, independent DOMAIN 12 Mbps token rings into a single DOMAIN network using industry-standard Ethernet cable as a backbone.

With EtherBridge, technical professionals can divide their network environment into multiple rings connected by an Ethernet backbone to create smaller, more managable networks. Every user on all of the linked rings has transparent access to every other workstation's data and resources. Also, if any network or link in an internetwork system becomes unavailable for any reason, the EtherBridge will automatically seek a new path over which to send the information -- whether it is a file, message, or other form of data.

The EtherBridge internetwork router system includes a MULTIBUS-compatible EtherController-MB Ethernet communications controller board, a transceiver and cable for connecting to the Ethernet backbone, and the EtherBridge software package. An Apollo DSP80A-HMB or DSP90 server processor are the minimum requirements for mounting the internetwork router system.

EtherBridge software is priced at $1,000 (per node); EtherController-MB hardware is priced at $3,300.

1.7. RJE Access Gateway

Apollo has expanded its communications gateway facilities with the DOMAIN Networking Remote Job Entry (RJE) Access Gateway. This new product offers Apollo users bi-directional text file transfer capabilities with a variety of remote host mainframes. The new gateway is one of many off-net communications options offered through Apollo's Gateway Program.

The RJE Access Gateway software emulates several RJE subsystems: IBM's HASP and the standard 2780 and 3780 environments. Offering enhanced queue management facilities, the RJE Access Gateway allows all users on an Apollo DOMAIN network access to remote environments using a KMW protocol converter (available from KMW, Inc., of Austin, TX). The RJE Access Gateway software and converter need only be implemented on one workstation within the network.

Apollo DOMAIN users, using the RJE Access Gateway in conjunction with the KMW equipment, can transmit and receive text files in a batch mode to and from a wide variety of remote computer systems.

1.8. DOMAIN X.25 Gateway

DOMAIN X.25 provides DOMAIN users with shared access to efficient long-distance communications via international public packet switching networks or private X.25 networks. DOMAIN X.25 conforms to CCITT recommendations for communications protocols and is compatible with the ISO Opens System Interconnection Reference Model. In addition, DOMAIN X.25 includes a highly reliable file transfer service for use between remote DOMAIN computing systems. The physical protocol, CCITT Recommendation X.21 bis (equivalent to EIA RS-232C at the lowest level), is performed in hardware by the DOMAIN X.25 intelligent communications controller. The higher protocol levels are provided in controller firmware and main processor software: LAPB and LAP at link level, the 1980 CCITT standard for virtual circuit support at the packet level, and both the X.3 and X.29 Recommendations for remote virtual terminal service at the application level. The DOMAIN-to-DOMAIN file transfer service (FTS) also operates at the application level. DOMAIN X.25 operates under AEGIS, the DOMAIN system's network-wide virtual memory operating system.

DOMAIN X.25 provides services that are shareable across a DOMAIN local-area network. Either a DSP80 DOMAIN Server Processor or a DN600, DN420, or DN440 Computational Node equipped with a Peripheral-to-Node Adapter can support the DOMAIN X.25 intelligent communications controller, performing as a gateway server to a community of DOMAIN users. The gateway server configuration removes the synchronous line-handling burden from users' Nodes, allowing them to devote their full resources to engineering and technical problem solving.

DOMAIN X.25 conforms to protocols X.29 and X.3 to provide line-oriented terminal communications that are built upon the lower level, packet-oriented services of X.25. This facility lets users write programs that call upon the computing services of conventional mainframes. X.29 and X.3 dynamically assign virtual circuits to application programs upon their request. Access to X.25 commands and events, plus additional line-oriented commands and events, are provided to the application program by X.29 and X.3. These protocols also includes character padding and parameter passing to X.25 in conformance with CCITT recommendations. An interactive terminal emulator built upon X.29 and X.3 provides DOMAIN users with access to conventional computing facilities. A user opens a window on DOMAIN's high-resolution display and runs the terminal emulator within it. To the remote computer, this emulation window appears to be an ordinary interactive terminal. Under the emulator, DOMAIN users can "log on" with conventional timesharing systems that support X.25 to run programs, query databases, initiate file transfers, etc. A full transcript of the interactive session can be saved on the DOMAIN system for later review.

DOMAIN X.25 FTS lets users on one DOMAIN network send and receive files to and from another DOMAIN system. Built upon X.25, FTS offers the data rate and transmission cost advantages of that protocol. And use of X.25's extensive error checking and recovery mechanisms by FTS to ensures data integrity. In case of temporary communication line failure, FTS automatically detects the problem, suspends operations, then retransmits the current file in its entireity once the line is restored. These features allow FTS to operate unattended at night or at other times when public carriers may be less heavily loaded or may offer lower transmission charges.

DOMAIN X.25 is certified for international operation using the public packet network services of Euronet, Transpac, Telenet, British Telecom, Datex-P, Datapac, and Tymnet.

1.9. DOMAIN Server Processor

The DOMAIN server processor, the DSP160, features a 32-bit architecture with up to 4 Mbytes of main memory and a 256-Mbyte virtual address space. The unit can support a 68- or 158-Mbyte Winchester disk drive, plus 1.2-Mbyte floppy disk drive.

The server processor is designed to perform computationally intensive applications; it can optionally use a 300-Mbyte disk drive, a 1,600-bps magnetic tape drive, or a Versatec, Inc. V80 plotter.

The DSP160, with 2 Mbytes of main memory, integrated hardware floating-point unit, three asynchronous I/O ports, license to use the Apollo operating system, and a network interface, costs $37,500.

Apollo ushered in 1985 with the announcement of a new DSP80A server processor. The DSP80A server processor offers any DOMAIN network a low-cost resource for peripheral and communications support. An enhanced model of Apollo's DSP80, the DSP80A has an expanded memory capacity from 1 Mbyte to 3 Mbytes, an enhanced cooling system for increased reliability, and mechanical improvements for easy access and maintainability.

The DSP80A is a 32-bit intelligent peripheral and communications server that expands Apollo's network-wide resource-sharing capabilities. The DSP80A can improve the performance of a DOMAIN network by allowing user-nodes to off-load peripheral and communications functions. The DSP80A is completely software-compatible with the entire family of Apollo workstations and servers.

The DSP80A, with two RS-232 I/O ports and a 5-slot MULTIBUS card cage, can support a broad range of peripheral devices, including mass-storage subsystems, a variety of hardcopy devices, and additional hardware such as array processors and other special-purpose computers. The DSP80A, configured with 1 Mbyte of main memory, two RS-232-C serial ports, a subset of the AEGIS network-wide virtual memory operating system, and a five-slot MULTIBUS card cage, is priced at $12,000; configured with 3 Mbytes, it is priced at $21,000.

1.10. Enhancement To Mass Storage Capabilities

Apollo's DOMAIN File Server (DFS) is avialable with a 500 or 1,000 Mbyte mass storage sub-system, and provides file server capabilities for an entire DOMAIN network. Apollo is also offering a new disk configuration featuring 500 to 1,000 Mbytes of local storage for the firm's family of workstations and servers. These new products represent a significant enhancement to Apollo's mass storage capabilities, affording diskless network nodes a more effective file service and storage facility.

The DFS, in addition to its mass storage capabilities, contains a high-performance 68010 processor, 1 Mbyte of main memory, and a 4-slot MULTIBUS peripheral adapter. The MULTIBUS interface lets users connect a wide variety of peripherals to a DOMAIN network, and permits these resources to be transparently shared by all users. The DFS can also provide communications gateway services by supporting such communication standards as X.25 and Ethernet, and can be configured with a magnetic tape drive for system backup.

1.11. Graphics Architecture

Graphics Metafile Resource (GMR), a new graphics architecture for the DOMAIN family, is designed to accommodate emerging industry standards such as GKS, PHIGS, and VDM, and is a standard feature on all Apollo workstations.

GMR combines both graphics database and advanced graphic display routines into one integrated package. While other graphics systems are limited by a division between compute power and display subsystems, GMR takes advantage of the integration offered by Apollo's graphics workstations. With GMR, graphics entities are stored in a Graphics Metafile, which may be shared among applications and viewed on any Apollo workstation within a DOMAIN local area network.

The Graphics Metafile, an editable, tree-structured database, is the focal point of GMR. GMR provides a comprehensive set of routines and lets the application developer build and edit complex models. Graphics data is stored directly in world coordinates, and may be viewed on all DOMAIN networks. GMR's segmented graphics database supports nesting and instantiation. Nesting permits segments to contain other segments, while instancing reduces storage requirements for repetitive data. In addition, segments may store non-graphic data or reference data in other files.

Metafile data may be displayed in multiple viewports within a given window, and GMR handles all scaling, translation, windowing, and clipping. Changes to the Metafile are reflected in all viewports simultaneously. Attributes may be associated with viewports or instances to control a wide range of display parameters including color, line style, and fill pattern.

1.12. Apollo Open-Architecture Program

At the end of July 1985, Apollo officially launched its open-architecture program with an initial set of products creating integrated networks of computers from multiple vendors. The program exploits and extends the resource-sharing architecture of Apollo's DOMAIN System.

These products offer DOMAIN users transparent access to data and resources from geographically dispersed DOMAIN Systems, IBM PC and hardware-compatible personal computers, and Digital Equipment Corporation VAX/VMS systems, and explicit access to IBM computers through IBM SNA networks. The products were implemented in accordance with industry and de facto communications standards.

Apollo said its open-architecture program provides: adoption of industry standards (and eventual support of the International Standards Organization's Open Systems Interconnect (ISO/OSI) model) for multi-vendor interconnects to the DOMAIN System; extension of DOMAIN architecture's transparent data and resource access to mixed-vendor environments; and interconnection of multiple DOMAIN System networks to one another over a variety of industry-standard networks such as Ethernet.

Apollo's architecture enables DOMAIN workstation users to access resources and files from a single, shared, network-wide environment as if the resources were residing on their local workstations. This results in a transparent user interface, which the open-architecture program extends to other systems outside a local DOMAIN System (i.e., network).

The communications products give DOMAIN users easy access to a variety of productive resources (e.g., programs, files, input/output devices, electronic mail) from interconnected DOMAIN Systems and computers from multiple vendors. The products allow the sharing of information among workgroup members who were previously isolated by location and by the dissimilarity of the computers they used.

DOMAIN/BRIDGE is a hardware/software communications package which extends the functionality of the DOMAIN System across long distances and facilitates the management of extremely large DOMAIN Systems. The product connects multiple DOMAIN System (i.e., multiple networks of Apollo workstations), making them one integrated DOMAIN System.

This product has two options. DOMAIN/BRIDGE-A -- an industry-standard TI class carrier link -- extends the DOMAIN single-system image across long distances and can give corporations new information-sharing and project-management efficiency. DOMAIN/BRIDGE-B -- a standard coaxial cable link -- partitions large networks, making them more manageable and greatly increasing the potential size of Apollo networks which have the DOMAIN single-system image.

DOMAIN System features are supported across the new communications bridges. These include the network-wide single file system and demand-paged virtual memory, file-locking, DOMAIN/DIALOGUE and DSEE software development tools, third-party applications software, electronic mail, and the DOMAIN/IX operating system.

Hardware for the DOMAIN/BRIDGE product consists of one or more cards inserted in MULTIBUS slots of a DOMAIN DSP80A, which acts as a dedicated bridge to other DOMAIN systems. A single DOMAIN System can have direct interconnects to more than one other DOMAIN System. The initial release of the product supports five alternate configurations of interconnected DOMAIN Systems.

DOMAIN/BRIDGE-A hardware consists of an internet interface controller card. The card provides an RS-449 interface to a modem, which connects to a full-duplex data communications link, such as microwave, local area network, and T1 Accunet services.

DOMAIN/BRIDGE-B consists of an internet interface controller card connected to an internet media adapter via its RS-449 interface. Requiring two MULTIBUS slots, this package directly connects to other DOMAIN Systems via coaxial cable.

DOMAIN/BRIDGE software is based upon the internet transport protocol layer of the Xerox Networking System (XNS); it is compatible with all DOMAIN workstations, servers, and operating systems, including DOMAIN/IX, Apollo's twin-port of industry standard Berkeley 4.2 and System V UNIX operating systems.

Apollo's DOMAIN/VACCESS-1 package is the first in a series of hardware/software products that gives DOMAIN users transparent access to files residing on DEC VAX/VMS computers. DOMAIN/VACCESS-1 is implemented over the industry-standard Ethernet 802.3 transport and transmission control protocol/internet protocol (TCP/IP) software.

Using this package, DOMAIN users can transfer VAX/VMS System 730, 750, 780, and 8600 files to their workstations at high speeds and can execute Apollo directory and file-manipulation functions on VAX/VMS files as if they were local DOMAIN files. The DOMAIN/VACCESS-1 product downloads files from VAX/VMS applications to the same applications on the DOMAIN System.

The package supports file transfer rates up to .4 Mbps. The hardware consists of an industry-standard, IEEE 802.3-compatible Ethernet transceiver/controller card for the DOMAIN System. The card resides in a MULTIBUS slot of DOMAIN DSP80, DSP80A, and DSP90 servers, DSP160 computational servers, and DN460, DN550, DN560, and DN660 workstations.

The Apollo package includes floppy disk-based communications software for the DOMAIN System gateway. Once configured with third-party software, the VAX host directly connects to the DOMAIN gateway node via industry-standard 802.3 Ethernet transport.

By providing bi-directional remote log-in, the product allows interactive terminal support for both DOMAIN and VAX users.

Users will be able to upgrade DOMAIN/VACCESS-1 to future DOMAIN/VACCESS products which will fully comply with the International Standards Organization's file transfer/access method (FTAM) protocol.

Apollo's DOMAIN/SNA hardware and software communications package allows Apollo's DOMAIN System users to share resources of IBM computers through IBM SNA networks. The package distributes SNA resources across Apollo's integrated network of technical workstations. Now, on a single Apollo workstation, a user can simultaneously work with DOMAIN and IBM resources by virtue of the DOMAIN System's multi-window environment.

The DOMAIN/SNA 3270 emualtor -- one of two optional packages

-- lets DOMAIN users log-on to a remote IBM host computer and work with such applications as text editors, program compilers, and database managers.

The 3270 option distributes up to 32 IBM 3278 Model 2 display stations and IBM 3287 Model 2 printer emulators to DOMAIN workstations. A single DOMAIN workstation can operate more than one display and printer emulator at a time.

The package's second option, the DOMAIN/SNA 3770 emulator, is a remote job entry (RJE) service allowing DOMAIN users to send batch jobs to an IBM host's job entry subsystem. This option lets a DOMAIN gateway workstation user send jobs to any combination of six IBM 2502 card readers, IBM 3521 card punches, and printers.

The DOMAIN/SNA communications panel connects directly to an IBM communications controller, directly to an IBM host, or to the host through a modem eliminator. A single DOMAIN gateway node supports up to two SNA lines connected to the same or two different hosts. And the same node can run 3270 emulation on one line and 3770 emulation on the other -- or either emulation on both lines at once. DOMAIN/SNA software supports data transmission rates of 1200 to 9600 bps (when both ports are used) and 1200 bps to 19.2 Kbps (when one port is used).

Depending upon the connection to the host, standard RS-232 null-modem or straight-through modem cables are used.

Apollo's DOMAIN/PCI communications package integrates IBM PCs, PC XTs, PC ATs, and hardware-compatible personal computers (such as COMPAQ computers) into Apollo's DOMAIN System. The new personal computer interconect (PCI) allows typical PC users -- such as management, support, and administrative staff -- to more productively share information and resources with their technical staff who use DOMAIN Systems.

Using DOMAIN/PC virtual terminal emulation, users can execute DOMAIN operating system shell commands to locate DOMAIN workstation or DOMAIN/PCI users and communicate with them by electronic mail.

DOMAIN/PCI users can also store text and binary files from DOS 3.10 applications programs, including Lotus 1-2-3 and WordStar programs, on the DOMAIN distributed file system. These files can then be accessed by any DOMAIN/PCI-connected personal computer. Personal computer ASCII files stored on the DOMAIN System are automatically translated into standard DOMAIN format when accessed by any DOMAIN workstation on the network.

PC users execute the same file-access and output-device command formats on the DOMAIN System as they use on their personal computers.

The package has two optional configurations: DOMAIN/PCI-1 and DOMAIN/PCI-8. DOMAIN/PCI-1, a software- only package, allows a direct RS-232 serial port connection from a single personal computer to any Apollo DOMAIN node via an RS-232-C cable. DOMAIN/PCI-8 includes: communications software; an eight-port RS-232 card for insertion in Apollo DSP80, DSP80A, and DN550 MULTIBUS slots; and an attached distribution panel connecting up to eight personal computers.

Both DOMAIN/PCI configurations require personal computers with 256 Kbytes of RAM and the DOS 3.10 operating system. All standard IBM PC, PC XT, PC AT, and compatible computers speeds up to 9600 bps are supported. For remote access to the DOMAIN System, both options also require a standard modem with control lines.

DOMAIN/PCI software is based upon Microsoft's MS-NET communications software.

All the above products are compatible with DOMAIN operating systems, including DOMAIN/IX, Apollo's twin-port to industry-standard System V and Berkeley 4.2 UNIX operating systems.

A summary of Apollo's major DOMAIN Systems offerings is given in the Table.

1.13. MAP

In November 1985, Apollo announced its commitment to support GM MAP. The company's GM MAP product development program is part of its Open Architecture Program unveiled in July 1985, and the MAP program -- along with recent announcements of Apollo-based, third-party manufacturing information software -- signals the company's move to further penetrate the growing manufacturing automation marketplace.

1.14. Open System Toolkit

In March 1986, Apollo introduced the Open System Toolkit, designed to let Apollo users add new peripherals and customize the DOMAIN System without having to modify the operating system source code.

The Open System Toolkit enables technical professionals to quickly and easily add new devices, build transparent connections to other vendor file systems, and define new types of files. This creates an open operating system, since the simplicity of the Toolkit means that it is not limited to only those individuals with extensive experience in operating system internals.

Workstations Optional touch pad and mouse	Desktop Supermini		Mid-Range Color Workstation	High Performance Workstation	High Performance Color Graphics Workstation	
	DN300	**DN330**	**DN560**	**DN460**	**DN660**	**DN660A**
Central Processor	32-bit, VLSI MC68010 Processor N/A	32-bit VLSI MC68020 Processor Dedicated MC68881 Floating Point Processor		32-bit, bit-slice processor 4 Kb bipolar instruction Cache memory 16 Kb bipolar data cache memory Integral single and double precision hardware, floating point		
Main Memory Supported	1 to 3 Mb	2 to 3 Mb		4 to 16 Mb ECC	4 to 8 Mb ECC	
Virtual Address Space 24 Concurrent Processes	Up to 16 Mb per process	Up to 64 per process		Up to 256 Mb per process		
Display Units High resolution, bit-mapped with detached keyboard	12 Mbit/sec bit-blt 17-in., 1024 × 800 monochromatic display		320 Mbit/sec bit-blt 19-in., 1024 × 800 color display	32 Mbit/sec bit-blt 19-in., 1024 × 800 monochromatic display	320 Mbit/sec bit-blt 19-in. color display 1024×1024 (DN660)	1024 × 800 (DN660A)
DOMAIN System Network	12 Mbit/sec token passing baseband communications network; Up to 3,280 feet (1 kilometer) separation between active nodes Several thousand nodes per network					
Wide Area Network	DOMAIN/BRIDGE multiple networks via high speed common carriers					
DSP80A DOMAIN Server Processor: Network-Wide Peripheral & File Server Support	32-bit VLSI MC68010 server processor. 16 Mb virtual address space. Up to 24 concurrent processes. .5 Mb to 3 Mb main memory. 5-slot IEEE-P796 MULTIBUS™ adaptor, 2 RS232C serial ports.					
DSP90 DOMAIN Server Processor: Network-Wide Compute & Peripheral Support	32-bit VLSI MC68020 server processor. Dedicated floating point co-processor. 64 Mb virtual address space per process. Up to 24 concurrent processes. 2 or 3 Mb main memory. 5-slot IEEE-P796 MULTIBUS adaptor, 2 RS232C serial ports.					
DSP160 DOMAIN Server Processor: Computational Server	32-bit bit-slice processor, 4 to 16 Mb ECC Memory 24 processes of 256 Mb virtual memory, 5-slot IEEE-P796 MULTIBUS adaptor, 3 RS232C serial ports Optional 80, or 167 Mb Winchester drives, 1.2 Mb diskette					
Software Support and DOMAIN Gateway Network	**Operating Systems** DOMAIN/IX Twin port of UNIX: Berkeley 4.2, System V Release 2 AEGIS, object-oriented, network wide, demand-paged virtual memory operating system		**DOMAIN Graphics Resources** DOMAIN/DIALOGUE™ Graphics Metafile Resource (GMR™) DOMAIN/CORE Level 3C GKS Level 2B DOMAIN/4014 Display Manager Graphics Primitives Resources	**Communications** ETHERNET TCP/IP VT100 Emulation X.25, DOMAIN/ ComController RJE: HASP/2780/3780 DOMAIN/PCI DOMAIN/VACCESS-1 DOMAIN/SNA **Development Tools** DOMAIN Software Engineering Environment (DSEE™)	**Language System** FORTRAN 77 ISO Pascal DOMAIN/C DOMAIN/LISP DOMAIN/DEBUG **DOMAIN Professional Support Service (DPSS)** Mail	
	More than 350 software application solutions					
	DN300	**DN330**	**DN560**	**DN460**	**DN660**	
Node-Based Peripheral Support	70 Mb Winchester disk (63 Mb formatted) 1.2 Mb diskette 2 RS232C serial ports		86 Mb 5¼" Winchester disk (69 Mb formatted) Dual 86 Mb 5¼" Winchester disk option 60 Mb ¼" cartridge tape 4-slot IEEE-796 MULTIBUS adaptor 2 RS232C serial ports	80 Mb Winchester disks (74 Mb formatted) 167 Mb Winchester disk (147 Mb formatted) 1.2 Mb diskette 5-slot IEEE-P796 MULTIBUS adaptor 3 RS232C serial ports		
			*300 Mb storage module disk, 500, 1000, and 2000 Mb fixed storage drive *1600 bpi and 6250 bpi 9-track tape drive Interface to Versatec printer/plotter			
Hard Copy Devices	Multi-mode printer for graphics, letter quality and draft modes DOMAIN/ LASER-26 for professional quality text and graphics					

*These are also available for all DSP server processors.

Table II-1: DOMAIN Systems Summary Chart

Open System Toolkit features a well-defined environment for users to extend the operating system; it also insulates user modifications from new operating system releases; requires no operating system source code modifications; and allows users to add new devices, connect to foreign file systems, and define new file types.

All Apollo operating systems -- including DOMAIN/IX, the company's twin-port of the Berkeley 4.2 and UNIX System V operating system -- are supported by the new Open System Toolkit. With DOMAIN/IX, a user can run either UNIX standard or both simultaneously on the same workstation.

1.15. Next Generation Products

In March 1986, Apollo removed the wraps from the next generation of its DOMAIN System -- an entirely new line of workstation products -- that offers the technical marketplace a broad range of fully compatible, UNIX-based computing resources and high standards of price/performance. At the same time, Apollo also introduced a new parallel processor/computer server and two new link-up products.

1.15.1. DOMAIN Series 3000 Personal Workstation

The DOMAIN Series 3000 Personal Workstations offer the performance of high-end graphics workstations for less than half the price. The new Series 3000 -- available in high-resolution color and monochrome -- bridges the gap between personal computers (PCs) and high-performance workstations by combining the best attributes of PCs with the functionality, advanced graphics, and 32-bit computing power necessary for efficiently running engineering and scientific applications.

These next-generation DOMAIN workstations combine the desktop, ergonomic packaging of a PC with the flexible, open architecture of an IBM PC AT-compatible bus with the high-performance 32-bit Motorola MC68020 processor and MC68881 floating-point co-processor.

In addition, the Series 3000 systems are fully functional DOMAIN workstations, with integrated networking and graphics, virtual memory for executing large engineering and scientific programs, and full, upward compatibility with the entire DOMAIN series of workstations, compute servers, and more than 500 software applications.

Apollo also announced that it will open the DOMAIN architecture further by offering a PC co-processor option for the Series 3000 later in 1986. The co-processor -- now being jointly developed by Apollo and Phoenix Software Associates, Ltd., of Norwood, MA -- will let Series 3000 users run widely used PC programs in a display window while concurrently running technical applications in windows under the much higher-performance DOMAIN environment.

Also immediately available is the DSP3000, a peripheral server based upon the Series 3000 system. Users across DOMAIN networks will be able to share low-cost, PC-compatible peripheral devices which are integrated with the new server's PC/AT-compatible bus.

The Series 3000 color graphics version is immediately available in configurations starting at $14,900; the Series 3000 monochrome version will be available in configurations starting at $9,900 in May 1986.

1.15.2. DN570 and DN580 Color Workstations

The new DN580 and DN570 high-resolution color workstations combine the dedicated processing power of Motorola's 32-bit, 16 MHz MC68020 microprocessor with the MC68881 floating point co-processor -- and a high-performance floating point co-processor accelerator (FPX) option that improves the performance over the MC68881 by a factor of three.

Both systems also bring advanced computer graphics technology to the workstation market. With Apollo's new 3D Graphics Accelerator (3DGA) option, the DN580 is capable of real-time, 3-D graphics performance at rates exceeding 100,000 transformed and clipped vectors per second, and introduces virtual memory techniques to graphics processing. The integrated display processor in the DN570 performs up to 40,000, 2-D transformed and clipped vectors per second.

Standard equipment on both the DN580 and DN570 includes ergonomic packaging, dedicated video memory with eight color planes allowing the display of 256 simultaneous colors from a palette of 16 million, a low-profile, detachable, programmable keyboard and optional mouse, two RS-232-C ports, an integrated 12 Mbps LAN interface, and licenses for AEGIS and DOMAIN/IX (twin port of Berkeley 4.2 and UNIX System V operating systems) software.

The DN570 sells for $29,900 with a 15" monitor and $32,900 with a 19" monitor. The DN580, with its 19" standard display, sells for $43,900. The 3-D Graphics Accelerator Option for the DN580 is $9,900. The FPX floating point accelerator option for both models is $5,900.

Apollo also introduced for its family of workstations a new 5-1/4" 190 Mbyte (154 Mbytes formatted) Winchester disk, which is also available in a dual Winchester disk (308 Mbytes formatted) configuration. This new disk provides an increase of 123% in formatted data storage capacity on a single volume for the new

DN580 and DN570, as well as for the DN560 and DN550 color workstations. The disk subsystem is also offered in a configuration with a 1/4" 60 Mbyte streaming tape cartridge back-up.

1.16. Computer Server

Apollo boosted the computational power of its DOMAIN System in March 1986 into near-supercomputer ranges by announcing the DSP9000, an advanced parallel processor and compute server. The DSP9000 servers are based upon the FX/Series mini-supercomputers from Alliant Computer Systems Corporation of Acton, MA. This technology's high-speed computing capability is the result of parallel processing technology -- or the ability of a computer to devote more than one computational element (CE) to a single task.

Unlike other parallel processors, the DSP9000 has concurrency -- i.e., the ability to automatically run existing ANSI FORTRAN 77 application programs without having to modify the source code (i.e., the software itself) -- and this saves the time and expense of re-programming existing applications. DOMAIN users will directly benefit from the DSP9000 server's concurrency because many programs for the DOMAIN System are written in FORTRAN 77, the industry-standard language for scientific/engineering software applications.

In addition, the DSP9000 FORTRAN implementation includes language extensions based upon VMS. Pascal and C compilers are also available for this compute server.

The entry-level DSP9000 configuration, the DSP9011, $195,750 with just one CE, delivers 11.8 MFLOPS and 4.5 MIPS peak performance, or up to eight times the performance of a 1 MIP workstation; eight times the performance is achieved because special floating-point hardware provides performance measured in MFLOPS in addition to the performance measured in MIPS. The FORTRAN compiler automatically invokes this additional facility using vectorizing optimization techniques.

The DSP9081, expandable to 8 CEs, delivers the maximum power of 94 MFLOPS and 35 MIPs while selling for $325,250.

The DSP9000 server is available in several pre-packaged configurations. All configurations include hardware and software for interfacing with the DOMAIN network. DSP9000 servers have an air-cooled cabinet 30" wide, suitable for office environments.

1.17. Link-Up Products

Apollo extended the capabilities of the DOMAIN Network in March 1986 with the introduction of a low-cost, fiber-optic link for interconnecting multiple DOMAIN networks up to 3 kilometers apart into a single token ring network, and the announcement of the DOMAIN Quick Connect System, which allows for the quick and easy connection of Apollo nodes to a token ring network.

The DOMAIN/DFL-100 Fiber Link is a low-cost, fiber-optic-based network device that significantly enhances the maintainability and operation of Apollo's 12 Mbps token-passing ring network. The new link supports failure detection and recovery, as well as automatic secondary link recovery, providing highly available uptime for inter-building network links. In addition, network reliability and flexibility are greatly enhanced with the new optical link, in that it is immune from electrical interference, lightning strikes, and harsh environments.

When the DOMAIN/DFL-100 senses a failure within itself or in its optical interface media, it immediately closes its coaxial ring interface until the link is repaired and returned to operation. Also, parallel DOMAIN/DFL-100 links can be installed between locations, so that if one pair ceases to operate or is shut off, the second, redundant pair automatically takes over.

The DOMAIN/DFL-100 Fiber Link is available 90 days ARO and is priced at $1,200 in its primary version and $1,250 in its redundant link version.

Apollo's DOMAIN Quick Connect System facilitates the connection and management of the Apollo token ring, and provides a network management tool for today's high uptime network environments.

DOMAIN/DQC-100 Quick Connect is a low-cost, durable unit that, for the first time, allows for the quick connection of Apollo nodes to a DOMAIN Network without shutting down the network. It can be mounted in a standard electrical wall outlet box or surface-mounted using hardware available from third-parties. The connector uses no power and is transparent to electrical signals in the DOMAIN Network.

DOMAIN/DQC-101 Quick Connect Cable supports the attachment of DOMAIN nodes to the DOMAIN/DQC-100 Quick Connect Outlet. The cable is ten feet long and attaches to DOMAIN nodes via industry-standard BNC connectors and to the Quick Connect Outlet through a durable posi-lock connection.

The DOMAIN/DQC-100 and DOMAIN/DQC-101 are priced at $35 and $45, respectively, and are available through Apollo's DOMAIN System User Catalog.

1.18. Pricing

Model Number	Description	List Price

DN300-HMB	DOMAIN DN300 DESKTOP COMPUTATIONAL NODE Includes: High-Resolution 1024 x 800 Bit-Mapped Black and White 17" Display, 1 Mbyte Main Memory	$ 9,900
DN320-1.5MB	DOMAIN DN320 DESKTOP COMPUTATIONAL NODE Includes: High-Resolution 1024 x 800 Bit-Mapped Black and White 17" Display, 1.5 Mbyte Main Memory, Hardware Floating Point	$18,900
DN330	1 MIP DESKTOP WORKSTATION with 2 MB Main Memory	$15,900
DN330-70MB	Same As DN330 with 70MB Disk	$26,800
DN460	DOMAIN DN460 COMPUTATIONAL NODE Includes: High-Resolution 800 x 1024 Bip-Mapped Monochrome Graphics Display, 19" Landscape Orientation 1 Mbyte Main Memory Hardware Floating Point 5-Slot Chassis	$39,500
DN550	DOMAIN DN550 COMPUTATIONAL NODE Includes: 1 Mbyte Main Memory 19" Color Display Low-Profile Keyboard Electronics Enclosure 1 MB Display Memory Available 120 Days ARO	$31,500
DN560	1 MIP, Mid-Range Color Workstation with 2MB Main Memory	$35,500
DN560-86MB	DN560 With 86MB Disk	$46,400
DN570	Color Graphics Workstation 15" Color Display	$29,900
	Color Graphics Workstation 19" Color Display	$32,900
DN580	Color Graphics Workstation Includes: 2 Mbyte Main Memory High-Performance Display Processor	$43,900

	19", 60Hz, Non-Interlaced, 1280 x 1024 Bit-Mapped Color Display	
DN660	DOMAIN DN660 COMPUTATIONAL NODE Includes: High-Resolution 1024 x 1024 Bit-Mapped Color Graphics Display, 19" Landscape Orientation 2 Mbyte Main Memory 1 Mbyte Dedicated Display Memory, 4/8 Plane Hardware Floating Point 5-Slot Chassis	$59,500
DSP160	DOMAIN DSP160 SERVER NODE Includes: 2 Mbyte Main Memory Hardware Floating Point	$37,500
	DOMAIN FILE SERVER Includes: 500 MB Disk, Including Cabinet Peripheral Server/Processor	$36,000
DSP80A	DOMAIN SERVER PROCESSOR Includes: 1 Mbyte Main Memory 2 RS-232 Serial Ports AEGIS Subset 5-Slot Multibus Card Cage	$12,000
DSP90	DOMAIN DSP90 SERVER PROCESSOR Includes: 2 Mbyte Main Memory Hardware Floating Point	$18,000
DSP3000	SERVER PROCESSOR Includes: MC68020 Processor and MC68881 Floating Point Coprocessor; 2 Mbyte Main Memory; 86 MB 5-1/4" Winchester Disk; 60 MB 1/4" Cartridge Tape Drive and Controller; 2 RS-232-C Ports; IBM PC/AT Compatible Bus; DOMAIN Ring Interface; DOMAIN/IX and AEGIS Software Licenses	$15,500
DSP9011	COMPUTER SERVER Includes: 1 Computational Element; Appropriate I/O Subsystem; 2 Gbytes Virtual Memory Per Process; Operating	$195,750

	System Including C Compiler and Emacs Editor; Documentation; Hardware and Software for Interfacing with DOMAIN Network; Air-Cooled Cabinet; FORTRAN 77 Compiler for Single CPU; Plus 8 Mbyte Main Memory (expandable to 16 Mbyte); 268 MB Disk and Controller; Cartridge Tape Drive; 2 MULTIBUS Chassis; Expansion Cabinet; Cartridge Tape Drive	
DSP9081	COMPUTER SERVER Includes: 1 Computational Element; Appropriate I/O Subsystem; 2 Gbytes Virtual Memory Per Process; Operating System Including C Compiler and Emacs Editor; Documentation; Hardware and Software for Interfacing with DOMAIN Network; Air-Cooled Cabinet; FORTRAN 77 Compiler for Single CPU; 8 MB Main Memory (expandable to 64 Mbyte); 379 Mbyte Disk; 800/1600/6250 Tri-Density Tape Drive, Controller, and Cabinet.	$325,250
COM-ETH	DOMAIN/ETHERNET GATEWAY Includes: Hardware Controller Transceiver And Cable TCP/IP Software	$3,500
COM-RJE	RJE ACCESS GATEWAY	$1,800
COM-X.25	X.25 GATEWAY Includes: Hardware Controller Dual Synchronous Lines Full X.25 Software Protocol 2 Modem Cables	$6,950
	DOMAIN/BRIDGE-A, Industry Standard T1 Class Carrier Link	$8,400
	DOMAIN/BRIDGE-B, Industry Standard Coaxial Link	$8,900
	DOMAIN VACCESS-1. Hardware/ Software Communications Product That Transparently Shares Files From DEC VAX VMS computers Across A Network Of DOMAIN Workstations (License Per Site - Up to 100 Nodes)	$5,000

DOMAIN/SNA 3270 EMULATOR. Lets DOMAIN Users Log On To A Remote IBM Host Computer And Work With Text Editors, Program Compilers, And Database Manager	
- License Per Node	$1,700
- License Per Site	$7,000
DOMAIN/SNA 3770 EMULATOR. Lets DOMAIN Users Send Batch Jobs To An IBM Host's Job Entry Subsystem	
- License Per Node	$1,700
- License Per Site	$7,000
DOMAIN/PCI-1. Connects One Personal Computer Directly To A DOMAIN Workstation. (License Per Node)	$500
DOMAIN/PCI-8. Connects Up To Eight IBM PC's And Hardware-Compatible PCs To A Single DOMAIN Workstation. (License Per Note)	$6,400

DOMAIN SERIES 3000 PERSONAL WORKSTATION

Features include: Integrated MC68020 processor and MC68881 floating point coprocessor; integrated IBM PC/AT-compatible peripherals bus; 2 Mbytes or 4 Mbytes main memory; 1 RS-232-C port; programmable keyboard and mouse.

COLOR PERSONAL WORKSTATION Includes: 2 MB Main Memory; 15", 60 Hz, Non-Interlaced 1024 x 800 Bit-Mapped Color Display; 4-Bit Planes Support Simultaneous Display of 16 Colors From Palette of 4096	$14,900
COLOR PERSONAL WORKSTATION Includes: Fully Configured Includes all of the above Plus Disk Controller; 5-1/4", 86 MB Winchester Disk With 1.2 MB Floppy Disk	$19,400
COLOR PERSONAL WORKSTATION Includes: Same as above except with 5-1/4", 86 MB Winchester disk with 60 MB 1/4" cartridge tape	$21,400
MONOCHROME PERSONAL WORKSTATION	$9,900

	Includes: 2 MB Main Memory; 19", 64 Hz, Non-Interlaced 1280 x 1024, Bit-Mapped Monochrome Display	
	MONOCHROME PERSONAL WORKSTATION Includes: Fully Configured - Includes all of the above Plus Disk Controller; 5-1/4" 86 MB Winchester Disk with 1.2 MB Floppy Disk	$14,400
	MONOCHROME PERSONAL WORKSTATION Includes: Same as above except with 5-1/4", 86 MB Winchester Disk with 60 MB 1/4" Cartridge Tape	$16,400
	2MB Add-On Memory	$2,000

System Options

PEB	PERFORMANCE ENHANCEMENT BOARD Hardware Floating Point 4 Kbyte Cache Memory	$4,500
PNA PNA-50	PERIPHERAL-TO-NODE ADAPTER Including: 5-Slot IEEE-796 [MULTIBUS (tm)] Card Cage Power Supply	$3,000
3D Graphics Accelerator Option (DN580)		$9,900
FPX Floating Point Accelerator Option (DN570, DN580)		$5,900

Magnetic Storage Devices

	MASS STORAGE EXPANSION UNITS HOUSED WITHIN THE DOMAIN COMPUTATIONAL NODE	
MSD-1.2M	1.2 Mbyte Diskette Drive (Maximum of One Per Node)	$2,500
	14-inch Winchester Disk (Maximum of One Per Node):	
MSD-34M	-- 34 Mbyte	$10,000
MSD-68M	-- 68 Mbyte	$11,000
MSD-86M	-- 86 Mbyte 5-1/4" Winchester Disk Subsystem	$5,900

Ring-Based Local Networks

MSD-86M-TC	-- 86 Mbyte 5-1/4" Winchester Disk With Integrated 60 MB 1/4" Tape Cartridge Unit	$8,900
MSD-158M	-- 158 Mbyte	$16,800
MSD-190M	-- 190 MB 5-1/4" Winchester Disk Subsystem	$9,900
MSD-190M-TC	-- 190 MB 5-1/4" Winchester Disk With Integrated 60 MB 1/4" Tape Cartridge Unit	$12,900
MSD-190M-190M	-- Dual 190 MB 5-1/4" Winchester Disk Drives	$5,900
MSD-1600 MSD-1600-50	MAGNETIC TAPE DRIVE AND CONTROLLER (Maximum of One Per Node): 1600 BPI, 25 IPS, 9-Track Tape Drive Peripheral-to-Node Adapter Required	$12,500

Hard Copy Devices

HCD-LQP HCD-LQP-50	LETTER QUALITY PRINTER Uses one of the Asynchronous RS232-C I/O Ports Available on the Computational Node	$4,000
	HIGH-SPEED LINE PRINTERS WITH GRAPHICS CAPABILITY (Maximum of One Per Node, Peripheral-to-Node Adapter Required)	
HCD-300 HCD-300-50	300 Lines/Minute, with Controller	$8,500
HCD-600 HCD-600-50	600 Lines/Minute, with Controller	$11,000
	SOFTWARE LICENSES (Per Node Basis Unless Otherwise Specified) Includes Documentation and Distribution Diskette	
SFW-FTN	FORTRAN-77	$1,250
SFW-PAS	Pascal	$1,250
SFW-C	C	$1,250
SFW-D3M	DOMAIN Distributed Data Management System (D3M)	$2,500
SFW-CORE	SIGGRAPH Core Graphics Software Package	$1,000

SFW-GPIO	General Purpose I/O Software (Only a single GPIO License is required at a single site, regardless of the number of Nodes at that site) -- Used for Interfacing IEEE-796 [Multibus (tm)] Peripherals to Peripheral-to-Node Adapter -- Peripheral-to-Node Adapter Required	$3,500

1.19. Evaluation

Apollo has done a fine job of integrating a local network, screen management, a file system, and virtual memory into a coherent system that appears easy to use. New processors and peripheral devices can be connected with minimal changes in hardware (one cable) and software (one line of code).

Apollo has chosen to enter scientific and engineering markets initially, because of the DOMAIN system's ability to handle large computational problems. Although the system can address "a very large number" nodes, 1,000 is probably a practical limit, with 10 to 20 the most likely configurations.

The company is located at 330 Billerica Road, Chelmsford, MA 01824; (617) 256-6600. The company has over 70 sales and service offices worldwide. Its international headquarters are in Geneva, Switzerland.

2. IBM Token-Ring

On October 15, IBM put an end to years of speculation and announced major enhancements to its work group connectivity product line. Among the new products announced was a new local area network, the long-awaited IBM Token-Ring Network. Also, the IBM Cabling System was extended with a new media, the Type-3 Telephone Twisted Pair Specification, which supports the IBM Token-Ring Network and 3270-to-3274 attachment.

Expected by many observers though missing from the announcement were products to connect 327x-type devices. The interfaces to these and other types of IBM devices are not as clean as for the PC; therefore, the software development effort, which does not lend itself to the economies of scale that hardware development does, probably has not yet been completed. IBM is expected to announce interfaces for these products, most likely in 1986.

As a network for personal computers, IBM's Token-Ring joins the 70 or 80 other baseband personal computer local networks from as many vendors. The significance of this announcement is that it is just the tip of an iceberg. When fully populated with interfaces for the rest of IBM's products, the Token-Ring will provide complete local communications in an IBM customer site.

The IBM Token-Ring Network is a 4 Mbps baseband PC network that conforms to IEEE 802.5 and ECMA 89 standards. It can use either the IBM Cabling System or IBM's newly announced Cabling System specification for telephone twisted pair wiring. If the data-grade medium is used, up to 260 devices can be attached to the Token-Ring; if the telephone medium is used, up to 72 token ring devices can be attached to the network.

A token ring access protocol is employed for network traffic control. The various devices to be attached to the network are cabled together through a network access unit to form a logical ring. Access to the shared ring is controlled by a token. Each networked device regenerates the signal as information is passed around the ring. The access protocol and baseband signaling are implemented by a network adapter in each attached device.

Hardware components of the Token-Ring Network include an adapter card with logical link control for the attachment of the IBM PC to the network, a multistation access unit, and cabling. Software components include a program that provides an IBM PC Network NETBIOS interface to the network, an SNA programming interface for the PC, a program that links the PC Network and the Token-Ring Network, and a communications server.

The IBM Token-Ring Connectivity is illustrated in Figure II-4.

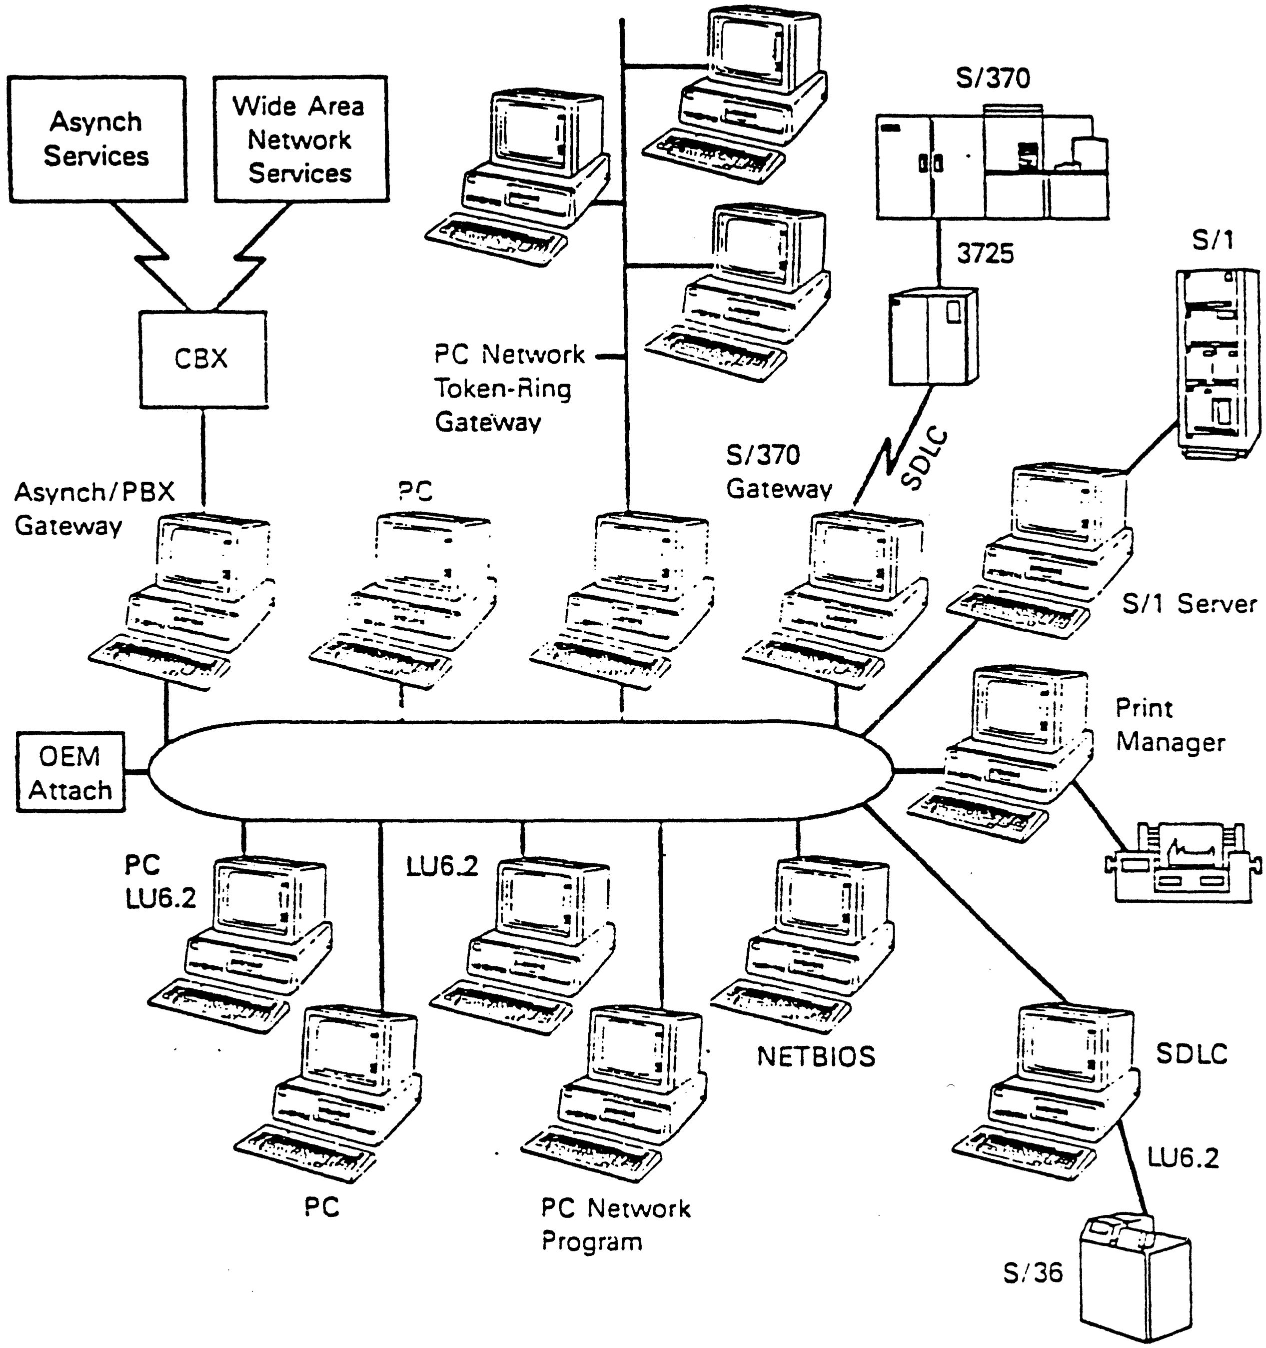

Figure II-4: IBM Token-Ring Network Connectivity

2.1. Token-Ring Concept

IEEE 802.5, the token-passing ring standard, specifies the format and protocols used by the IBM Token-Ring medium access control (MAC) sublayer, the physical (PHY) layer, and the means of attachment to the Token-Ring physical medium.

The Token-Ring consists of a set of stations serially connected by a transmission medium. Information is transferred sequentially, bit by bit, from one active station to the next. Each station generally regenerates and repeats each bit and serves as the means for attaching one or more devices (terminals, workstations) to the ring for the purpose of communicating with other devices on the network. A given station (the one that has access to the medium) transfers information onto the ring, where the information circulates from one station to the next. The addressed destination station(s) "copies" the information as it passes. Finally, the station which transmitted the information effectively removes the information from the ring.

A station gains the right to transmit its information onto the medium when it detects a token passing on the medium. The token is a control signal comprised of a unique signalling sequence that circulates on the medium following each information transfer. Any station, upon detection of an appropriate token, may capture the token by modifying it to a start-of-frame sequence and appending appropriate control and status fields, address fields, information field, frame check sequence, and the end-of-frame sequence. At the completion of its information transfer and after appropriate checking for proper operation, the station initiates a new token, which provides other stations the opportunity to gain access to the ring.

A token holding timer controls the maximum period of time a station will use (occupy) the medium before passing the token.

Multiple levels of priority are available for independent and dynamic assignment depending upon the relative class of service required for any given message, e.g., synchronous (real-time voice), asychronous (interactive), immediate (network recovery). The allocation of priorities is by mutual agreement among users of the network.

Error detection and recovery mechanisms are provided to restore network operation in the event that transmission errors or medium transients (e.g., those resulting from station insertion or removal) cause the access method to deviate from normal operation. Detection and recovery for these cases uses a network monitoring function that is performed in a specific station with back-up capability in all other stations which are attached to the ring.

2.2. Token Ring Formats

There are two basic formats used in token rings: tokens and frames. In the following discussion, the Figures depict the formats of the fields in the sequence they are transmitted on the medium, with the left-most bit or symbol transmitted first.

Processes, which require comparison of fields or bits, perform that comparison upon those fields or bits as depicted, with the left most bit or symbol compared first and for the purpose of comparison, considered most significant.

```
The Token Format is:    SD    AC    ED
```

where

SD = Starting Delimiter (1 octet),
AC = Access Control (1 octet), and
ED = Ending Delimiter (1 octet).

Figure II-5: Token Format

The token is the means by which the right to transmit (as opposed to the normal process of repeating) is passed from one station to another.

The Frame Format is:

```
 < - SFS - >  < ------- F C S Coverage -----> < - EFS - >

  SD    AC      FC     DA    SA      INFO    FCS      ED    FS
```

where

SFS = Start of Frame Sequence,
SD = Starting Delimiter (1 octet),
AC = Access Control (1 octet),
FC = Frame Control (1 octet),
DA = Destination Address (2 or 6 octets),
SA = Source Address (2 or 6 octets),
INFO = Information (0 or more octets),
FCS = Frame Check Sequence (4 octets),
EFS = End of Frame Sequence,
ED = Ending Delimiter (1 octet), and
FS = Frame Status (1 octet).

Figure II-6: Frame Format

The frame format is used for transmitting both Medium Access Control (MAC) and Logical Link Control (LLC) messages to the destination station(s). It may or may not have an information field.

2.3. IBM Token-Ring Network Hardware Components

2.3.1. The IBM Token-Ring Network PC Adapter

The IBM Token-Ring Network PC Adapter board is at the heart of IBM's token ring network as it exists today. The adapter contains a microprocessor operating under the control of microcode resident on the adapter. The adapter transmits and receives data at 4 Mbps, using protocols that conform with IEEE 802.5 and ECMA 89 standards. An adapter handler support program that runs in PC memory (requires 7 Kbytes) provides data link control in accordance with the IEEE 802.2 standard.

Before the Token-Ring was announced, IBM and Texas Instruments worked together on an adapter chip set that many believed would be at the heart of the new network. IBM eventually decided to use its own components, however, which appear to employ a more dense and sophisticated technology than the jointly developed chip set. Texas Instruments introduced the TMS380 chip set itself, shortly after the IBM Token-Ring announcement.

Reliability, availability, and serviceability functions are built into the Token-Ring adapter and microcode. Diagnostics invoked during adapter initialization verify the adapter operation and examine the cabling to the access unit. The adapter detects permanent errors, such as the loss of a receive signal, and generates a notification signal to initiate automatic network recovery. Recoverable errors, such as bit errors in the transmitted message, are detected by the adapter and then reported to a ring diagnostic program.

An attachment cable is required for connecting the adapter to the network cabling system. The IBM Token-Ring Network PC Adapter Cable ($35) is a 2.4 meter (8 foot) cable for attaching the adapter to data-grade media; the Type 3 Media Filter includes a 2.4 meter cable to connect the adapter to Type 3 telephone media. An optional advanced diagnostics program, the IBM Token-Ring Network PC Adapter Hardware Maintenance and Service, aids in isolating faults in the adapter, the adapter attachment cable, and the cable between the work area and the wiring closet.

Each PC Adapter costs $695, a significantly lower figure than was originally expected. The first shipment of this product was March 18, 1986.

2.3.2. The IBM Token-Ring Network Multistation Access Unit

The Token-Ring Multistation Access Unit is a passive access unit that permits up to eight devices to be attached to the network to form a logical ring (see Figure II-7). Cables from the various devices are simply plugged into the unit. Access units can be connected to each other to expand the size of the network (see Figure II-7). An attached device may be bypassed if the number of soft errors exceed an established minimum or if there is a wire or adapter error. The access unit may be installed in a standard 19-inch rack (not available from IBM) or in a component housing (available as an option for installing access unit on a wall or table top - $99). A cable bracket is included with the access unit for organizing and identifying the cables attached to a rack-mounted access unit. Price is $660 per unit.

2.4. IBM Token-Ring Network Software Components

2.4.1. The IBM Token-Ring Network NETBIOS Program

The NETBIOS program provides a high-level programming interface for IBM PCs. Application programs written to the NETBIOS programming interface can operate on both the Token-Ring and the IBM PC Network. (In the PC Network, NETBIOS runs in microcode; in the Token-Ring, it runs in PC memory and requires 46 Kbytes.) Application program requests for NETBIOS functions are translated into token-ring protocol requests which are passed to the Token-Ring Network PC Adapter for execution. There is a one-time program charge of $35.

The IBM PC Network SNA 3270 Emulation Program, Asynchronous Communications Server Program, and Token-Ring Network/PC Network Interconnect Program all use the NETBIOS programming interface and have been tested on the Token-Ring. Programs that operate in the IBM PC Network environment using the server functions of the IBM PC Network Program should also operate in the IBM Token-Ring environment using the server functions of the IBM PC Network Program.

The IBM PC Network Program, using an IBM PC with the IBM 3812 Pageprinter as a print server, has been successfully tested on the IBM Token-Ring Network. The IBM Series/1 PC Connect Program is expected to be available for use with the IBM Token-Ring Network in the second quarter of 1986. The IBM Series/1 may be positioned as the Token-Ring Network file server.

2.4.2. APPC/PC

The Advanced Program-To-Program Communications for IBM PC

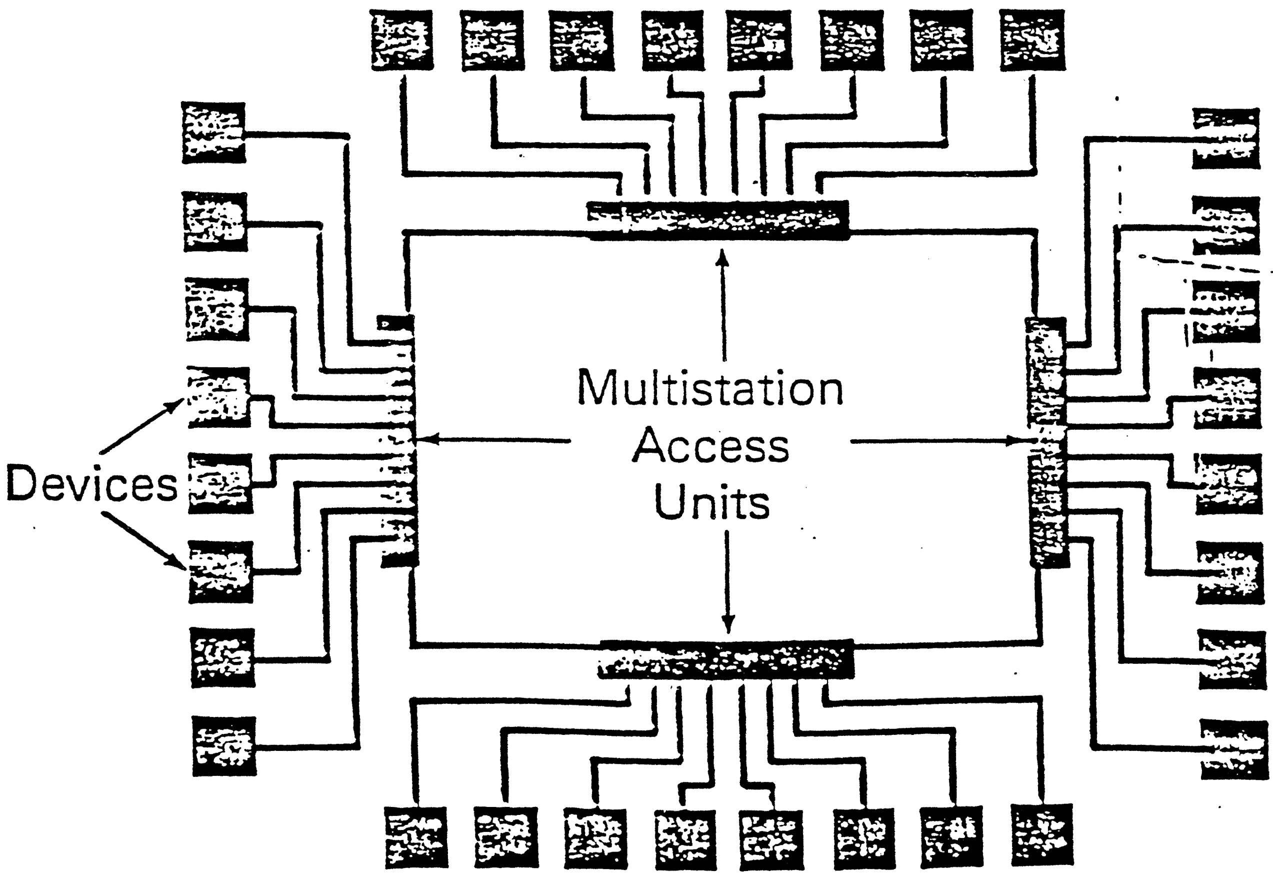

Figure II-7: Interconnecting Multistation Access Units

(APPC/PC) supports the SNA application interface (LU 6.2) and allows program-to-program communication over the IBM Token-Ring Network and SDLC communications links. The APPC/PC allows applications programs to be developed for the IBM PC using communications protocols common to other SNA devices, such as the System/370, System/36, System/38, and Series/1.

The addition of an SNA interface will be a very key selling point for the Token-Ring. An implementation of SNA LU Type 6.2 provides enhanced Systems Network Architecture support for distributed processing. As a communications protocol for a growing number of IBM products, APPC can improve connectivity among distributed transaction programs.

APPC/PC can be used by IBM PCs to help satisfy the requirements to attach to other SNA products using the SNA APPC architecture (LU 6.2). It provides communication services for applications in much the same way that IBM PC-DOS provides disk input/output and file management services. The subsystem will allow a program developer to concentrate on the application function rather than on the implementation of the communication protocols.

APPC/PC supports PU2.1 architecture and provides a peer-to-peer relationship between the IBM PC and the System/36, System/38, and Series/1; it attaches to a System/370 as a PU 2.0 node. It supports PC-to-PC connectivity using a Token-Ring communications link or an SDLC link, and PC to System/370, 36, 38 or Series/1 connectivity, using an SDLC communications link.

APPC/PC supports both the IBM Token-Ring Network PC Adapter and SDLC adapters. It provides a common program-to-program protocol that allows multiple conversations between applications running on an IBM PC and a System/370 (CICS/OS/VS Version 1, Release 7), System/36, System/38, Series/1 (Realtime Programming System Version 7.1), or another IBM PC. APPC/PC does not provide direct connectivity between sessions on the IBM Token-Ring Network link and sessions on the SDLC link, but it does provide the application programming interface to allow a user-written program to communicate between sessions on the two links.

The APPC/PC program has been designed to have an open application program interface; that is, it provides an interface which may be used by other products. APPC/PC will appear as a set of communication services provided by the operating systems.

APPC/PC supports a program interface through a set of verbs that allows application transaction programs to communicate at a conversational level with no session awareness. Application programs can engage in either mapped conversation, for use in communications between user-written programs, or basic conversation, intended to provide a lower level interface such as that required by logical unit service programs.

A PC can use both the APPC/PC and the NETBIOS program described earlier for concurrent access to the Token-Ring. Figure II-8 is a schematic representation of the two high-level programming interfaces. Only one Personal Computer Adapter is required. The program requires 195 Kbytes of memory in addition to the 7 Kbytes required for the adapter handler program and the ordinary DOS requirement. There is a one-time charge of $150 for licensing.

2.4.3. IBM Token-Ring Network/PC Network Interconnect Program

A dedicated IBM PC running only the Token-Ring Network/PC Network Interconnect Program can link the two kinds of networks so that devices can intercommunicate. Network application programs to be linked by the program must use NETBIOS protocols. Many programs that operate on either network using the NETBIOS programming interface will operate across both networks.

Up to 16 names for each network can be defined to the interconnect program; therefore, the maximum number of attached devices that can be identified to the other network is 16. (Some applications, such as the file server function of the IBM PC Network Program, require the use of multiple names.) The interconnect program is priced at $495 and scheduled for release in the second quarter of 1986.

2.4.4. IBM Asynchronous Communications Server Program

The IBM Asynchronous Communications Server program provides IBM PCs attached to the IBM Token-Ring Network or IBM PC Network access to asynchronous communications applications via circuit-switched connections through a CBX II, a PBX, or a public switched network. With the IBM Asynchronous Communications Server Program running on an IBM PC, XT, or AT, modems are not required at each individual IBM PC station in order for users to be able to access dial-up applications such as ASCII host systems, other IBM PCs, or information services (e.g., the IBM Information Network, the Dow Jones News/Retrieval Services, CompuServe).

When an IBM PC user wishes to establish a switched connection, he/she executes an asynchronous communications program that establishes a connection with the IBM Asynchronous Communications Server Program. The asynchronous communications program must use the IBM Asynchronous Communications Server Program protocol via NETBIOS. The server program establishes outbound calls to ASCII applications and accepts inbound calls destined for async programs running on a networked PC. ASCII data is transferred transparently.

The server is not dedicated, and multiple servers can be attached to either type of network. The communications server program is priced at $495.

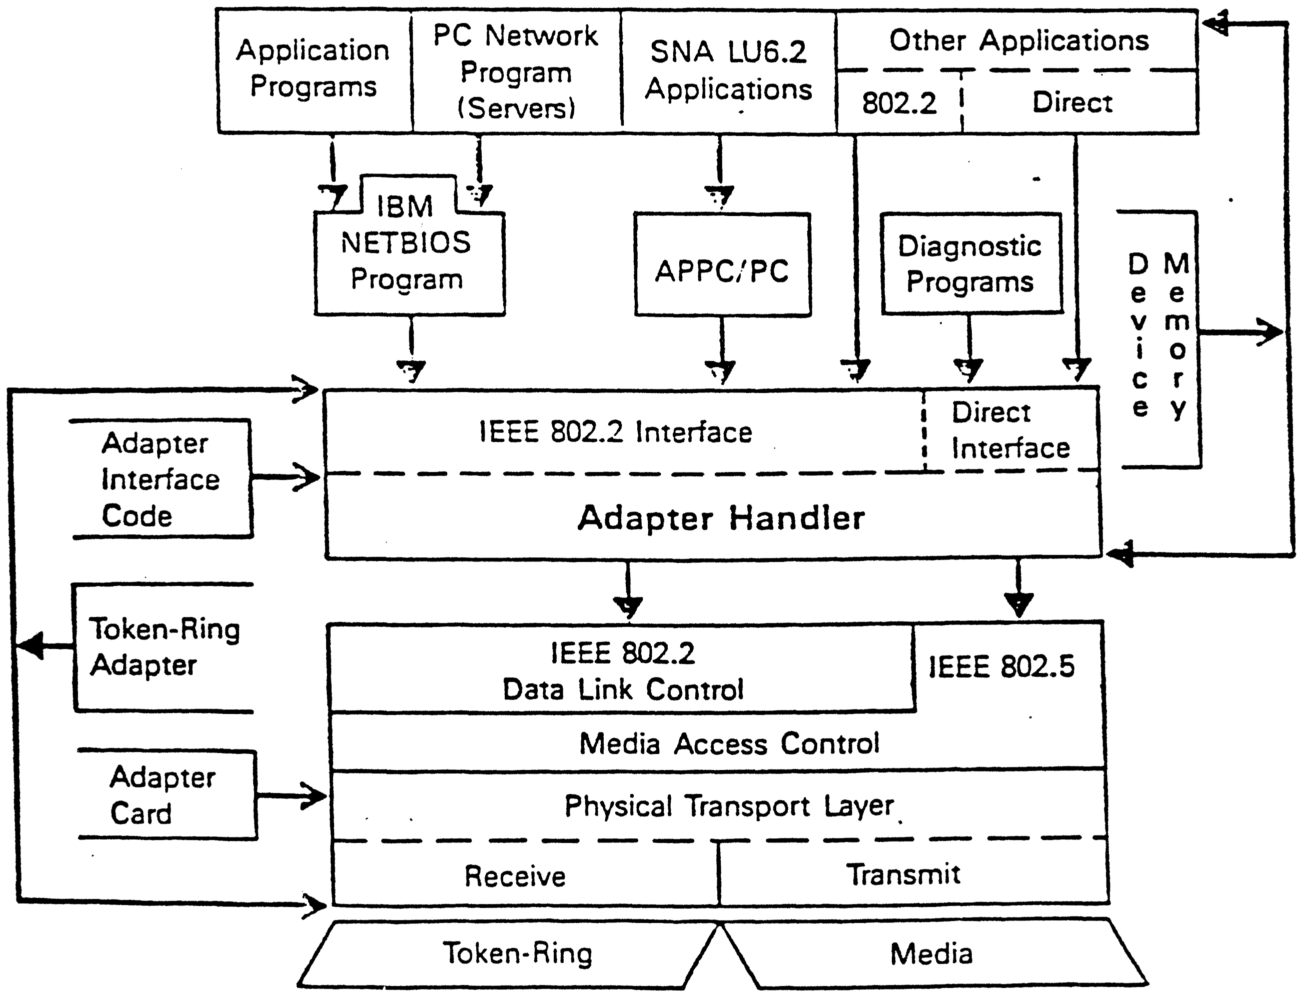

Figure II-8: Program Interfaces

2.5. The IBM Cabling System

The IBM Cabling System provides the physical interconnection for the Token-Ring. Components for the prewiring of a building include bulk cable, connectors, wall faceplates, and wiring closet distribution panels. Four cable assemblies (patch cables) of varying lengths allow the network to be installed in a building that has not been prewired. The cable assembly lengths are 2.4 meters (8 feet), 9.1 meters (30 feet), 22.9 meters (75 feet), and 45.7 meters (150 feet).

The structured wiring technique of connecting devices to wiring closets allows new devices to be added to the network without disrupting the other networked devices. Type 1 or 2 Cable specified in the initial IBM Cabling System allows the data transmission rate to be up to 16 Mbps and the distance between workstation and wiring closet to be up to 300 meters (990 feet), with 200 meters (660 feet) between wiring closets in a multiple closet environment. In an apparent bow to cost pressure from AT&T's STARLAN, IBM announced an extension the IBM Cabling System, called Type 3 Specified Media -- or, in lay terms, telephone wire -- on October 15. While Token-Ring networks using Type 3 Specified Media will be limited in size and performance compared to those using Type 1 or 2 Cable, many users will find the price advantage worth the performance tradeoff. Table II-2 compares IBM Cabling System Type 1 and 2 Cable and Type 3 Specified Media.

Type 3 Media Filter is used to connect the adapter card to the twisted pair network. One end of the cable has a D shell connector at the PC end and a standard G-pin modular jack at the other end. Type 3 Media Jumper Cables are used to attach the Type 3 Media to the Token-Ring Network Multistation Access Unit. Figure II-9 provides a more detailed look at Type 3 telephone twisted pair media. Four pair of Type 3 Specified Media can be included in one sheath -- two pair for the Token-Ring attachment, one pair for 3270 coax-to-twisted pair attachment, and one pair for telephone attachment.

2.5.1. Basic Elements of The IBM Cabling System

Transmission Cables:

- Type 1: For use between faceplates in work areas and wiring closets, or between two wiring closets in the same or different buildings; contains two balanced, twisted pairs for data transmission.

- Type 2: For use between faceplates in work areas and wiring closets in the same building; contains two twisted pairs for data transmission and four twisted pairs for voice transmission.

	IBM Cabling System Type 1 /Type 2 Cable	IBM Cabling System-Type-3 Specified Media
Other media applications:	Coax, Twinax, Loop	Coax
Token-Ring Media Operational Parameters:		
Tested data rates (Manchester encoded)	16Mbps	4Mbps
Signal drive distance (Using 4Mbps)		
Work area to 1 wiring closet	Max. 300 M (990 ft.)	Max. 100 M (330 ft.)
Wiring closets per ring	Max. 12	Max. 2
Work area to wiring closet if multiple wiring closets	Max. 100 M (330 ft.)	Max. 45 M (150 ft.)
Wiring closet to wiring closet (Cabling System Type 1 only)	Max. 200 M (660 ft.)	Max. 120 M (400 ft.)
Electronic interference protection	See GA27-3361	See GA27-3714
Configurations supported:		
Devices per ring	Max. 260	Max. 72
Undercarpet cable Support	Yes	No
System price considerations:		
Media and accessories	Yes	Yes
Installation labor	Yes	Yes
IBM Token-Ring Network PC Adapter Kits	Yes	Yes
IBM Token-Ring Network Multistation Access Units	Yes	Yes
Devices per ring (including devices used to connect rings)	260	72
Type 3 media filter	N/A	Yes
Type 3 media jumper cables	N/A	Yes
IBM Token-Ring Network PC Adapter Cable	Yes	N/A
IBM Cabling System patch cables	Yes	Yes
Telephone connectors/ punch-down blocks	No	Yes
3270 parameters		
From workstation to controller or IBM 3299 Model 1 Multiplexer	606 M (2,000 feet)	273 M (900 feet)
From workstation or controller to IBM 3299 Model 2 Multiplexer	1000 M (3,280 feet)	273 M (900 feet)

Table II-2: Comparison of IBM Specified Media

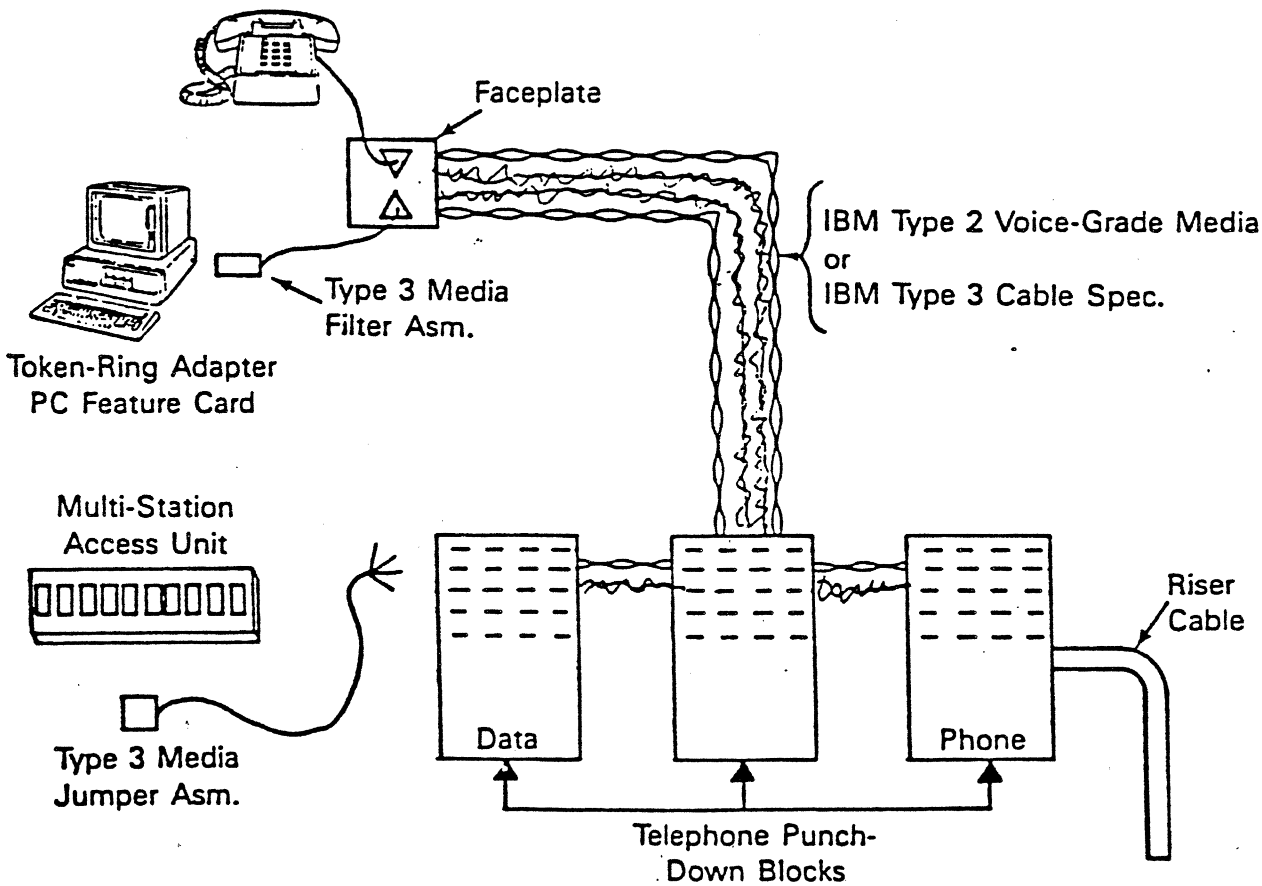

Figure II-9: IBM Cabling System Type 3 Specified Media

- Type 3: For use between faceplates in work areas and wiring closets in the same building; four unshielded, twisted pairs is recommended.
- Type 5: For use between wiring closets in the same or different buildings; contains two optical fibers for data transmission.
- Type 6: For use as patch cables in wiring closets; contains two twisted pairs for data transmission.
- Type 8: For use under carpet; contains two balanced, twisted pairs for data transmission; to be used with Type 1 or Type 2.
- Type 9: Lower cost alternative to Type 1; contains 2 twisted pair, #26 AWG copper wire.

Faceplates and Connectors:

- Faceplates for mounting on electrical outlet boxes installed in work areas.
- Data Connector and telephone connector for installation in the faceplate; data connector allows termination of two data-grade twisted pairs. A telephone jack connector allows termination of three voice-grade twisted pairs.

Distribution Panel:

- Cable junction panel mounted on a rack in a wiring closet; each panel allows connection of up to 64 data cables.
- Patch cable for making panel cross connections.

Device Attachment Cable and Accessories:

For use in attaching a device to the cabling system; the attachment cable is terminated with a workstation connector at one end and a cabling system connector at the other end; it includes an impedance matching device, when required.

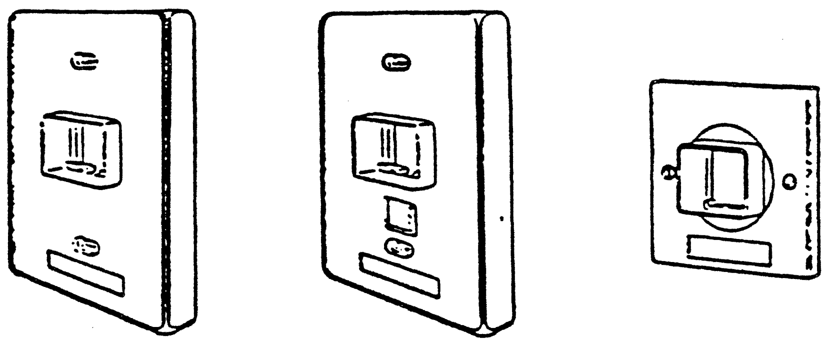

Figure II-10: Faceplates, IBM Cabling System

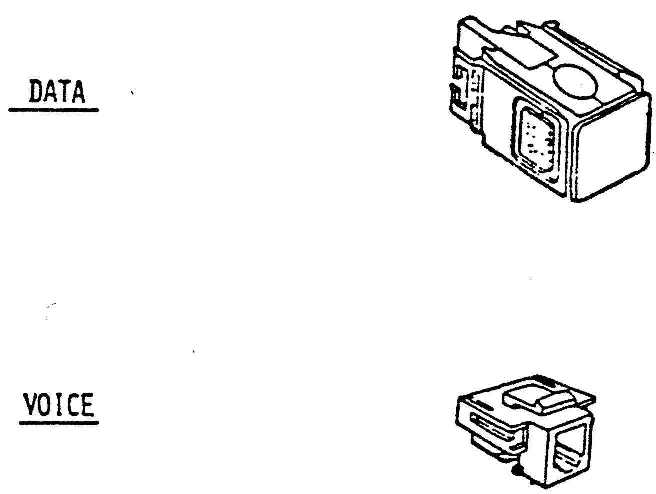

Figure II-11: Data and Voice Connectors, IBM Cabling System

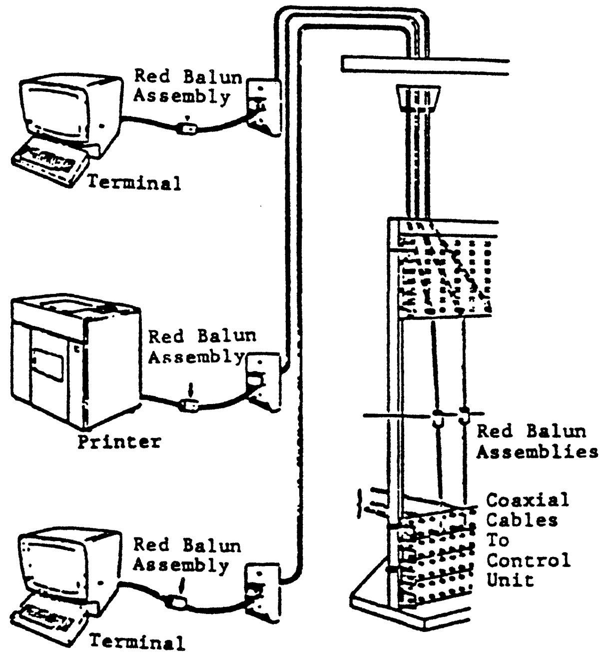

Figure II-12: 3270 Attachment, Cable Junction Panel

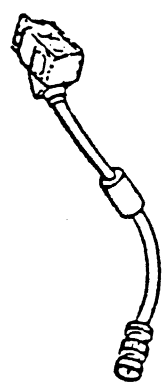

Figure II-13: Patch Cable for Panel Cross-Connection
(93:150 Balun)

2.6. Additional Token-Ring Products

On April 16, 1986, IBM announced several new offerings related to the Token-Ring. Included were several hardware, software, and cabling enhancements. These new offerings are summarized here.

In the area of PC enhancements, the IBM Token-Ring Network PC Adapter II adds the capability to function as a bridge between rings and is also qualified by IBM for use in the IBM industrial versions of the PC XT and AT. PC Adapter II can be used in a PC connected as a workstation to the ring, but due to its higher cost, its use should be limited to bridging rings.

The IBM Token-Ring Network Manager adds enhanced capabilities to the diagnostics software to better manage a Token-Ring. Up to 260 devices can be monitored. The program monitors the ring for hard and soft failures, logs network error and status information to disk, and can identify the probable cause and source of network failures. It can also remove a station from operation on the ring, if requested by an operator.

The Interconnect Program has been enhanced to include support for PC Network-to-PC Network communication.

The IBM Token-Ring Network Bridge Program supports message passing between two rings. Rings can be connected to other rings to form series, parallel, or hierarchical ring configurations.

With respect to products for larger machines, IBM still has not addressed the problem of supporting all of its larger machines on the Token-Ring (i.e., there are no direct-attach products). However, the company is now offering products that more directly support the System/36 and System/370.

The IBM 3725 Attachment to the Token-Ring Network is special hardware that supports direct attachment of the 3725 Communication Controller to the Token-Ring. 3275 Models 1 and 2 can attach to up to four separate rings. With an IBM 3276 Communication Controller Expansion Unit, a 3275 Model 1 can attach to up to eight rings. Communication between host applications on 370 systems and IBM PCs attached to the ring can be achieved with APPC, Personal Services/PC Version 1 Release 2, or the 3270 Emulation Program Version 3.

The IBM System/36 5360/5362 and 53/64 LAN Communications Licensed Program allows a System/36 to channel-attach to an IBM PC AT to connect to the Token-Ring. The package includes two software applications, one to run in the System/36, the other to run in the PC AT. The AT functions as a dedicated communications controller. Two adapter cards can be used in the AT to allow a System/36 to communicate with two separate rings. This product works in conjunction with the IBM System/36 LAN Attachment Feature. Planned availability is second quarter, 1987.

PC Support/36 has been enhanced to support a PC attached to the Token-Ring. All PC Support/36 functions that are available via direct twinaxial cable attached to the IBM Systems/36 are now available on the Token-Ring. System/36 Models 5360, 5362, and 5364 are supported. Planned availability is second quarter, 1987.

The IBM Series/1 Office Connect Version 1 Program provides document distribution services and document library services (similar to DISOSS) for Series/1 Realtime Programming System customers through DIA and SNADS. It also supports Personal Services/PC Version 1 Release 2 LAN support, via an interconnected PC (channel-attached between the Series/1 and the Token-Ring).

The IBM PC 3270 Emulation Program has been enhanced to support 3270 functions from a PC via the Token-Ring attached to the new 3725. It has also been enhanced to support BSC communication in a stand-alone configuration.

Personal Services/PC has been enhanced (Version 1, Release 2) to provide access to the IBM DISOSS/370 library via the 3725 Token- Ring connection. This gives users of IBM PCs, System/36s, and System/370s, access to office-oriented mail and document management facilities of DISOSS.

As for cabling enhancements, IBM is officially supporting the Type 5 fiber media with the IBM 8219 Token-Ring Network Optical Fiber Repeater. It attaches to the 8228 MAU to extend the distance between MAUs to 2.0 kilometers (6,600 feet).

A repeater for copper-type media, the IBM 8218 Token-Ring Copper Repeater, extends the distance between MAUs to 750 meters (2,500 feet), up from 200 meters.

A Type 9 media specification has been added that is designed to offer a lower cost alternative to Type 1 plenum cable. It consists of two twisted pairs of #26 AWG copper wire and can operate up to two-thirds of the distance specifications of Type 1. Type 9 is certified for use up to 16 Mbps. IBM will not manufacture this type of wire.

Figure II-14 summarizes the products announced on April 16, 1986.

2.7. Evaluation

In its October 15, 1985 announcement, IBM was careful to outline the differences between the IBM PC Network and the new Token-Ring Network. Though both LANs support attachment of IBM PCs, XTs, ATs, and Portables as well as the NETBIOS interface and IBM PC Network Program, the PC Network is intended primarily for

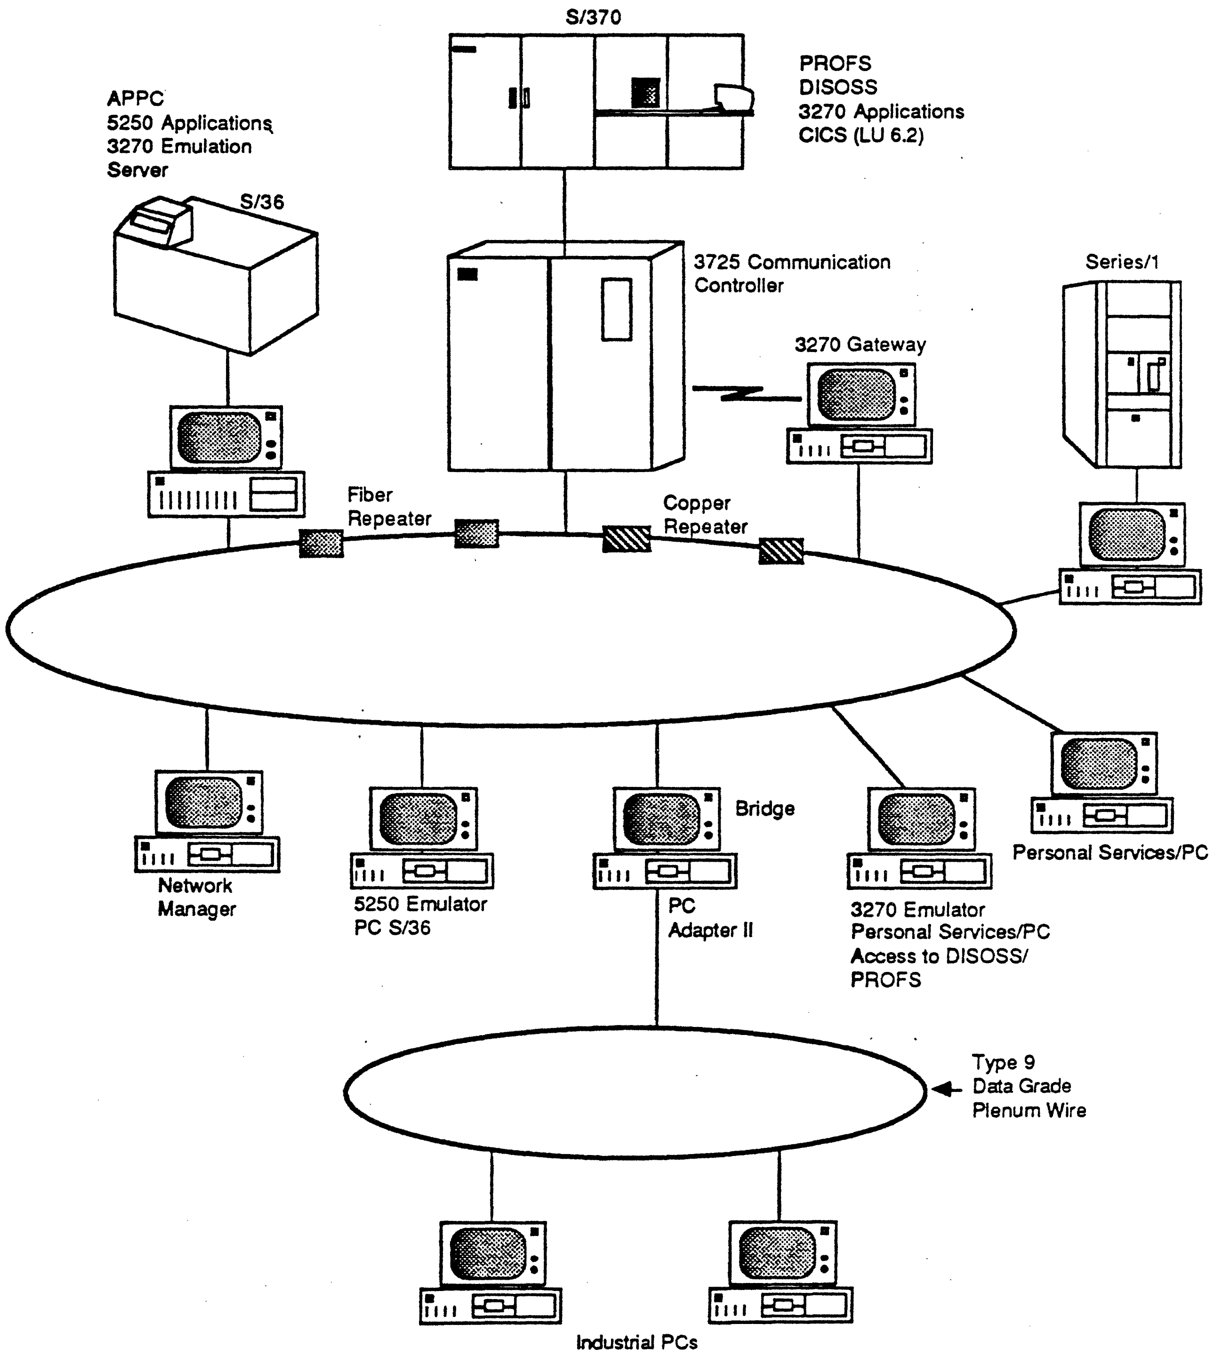

Figure II-14: IBM Token Ring Products - April 16, 1986 interconnecting PCs and related devices in a work group or departmental environment, while the Token-Ring is intended to connect a variety of systems and devices across an entire establishment. The interconnect program described previously will provide a gateway between the two networks in an organization.

Where the PC Network is a broadband network using coaxial cable, the Token-Ring Network is a baseband implementation that can use either shielded or telephone twisted-pair wiring. Broadband has the advantage of being able to support voice, data, and motion video on one medium, drawing on standard CATV components. It also has lower attachment costs than baseband, as well as high bandwidth transmission capabilities for supporting OEM devices on multiple channels. Broadband networks can be difficult to expand, however: they may require electrical design using a customized network to balance and tune the electrical characteristics of each of the drops and main line transmission.

Baseband transmission, such as that used by the Token-Ring, allows for smooth system expansion: new devices can be added without disturbing the rest of the network. Baseband may mean that networks cannot span the same long distances as broadband networks, however.

The star-wired ring configuration that IBM has chosen to adopt is a logical choice in view of its Cabling System wiring closets. The star-wired ring topology can support large network configurations, and because the data path is in a single direction, it is relatively simple to administer both on a day-to-day basis and in the event of a malfunction.

The Token-Ring improves on the PC Network's 2 Mbps transmission rate by offering 4 Mbps, and it can support up to 260 devices using shielded twisted pair or 72 devices using telephone twisted pair, as compared to the PC Network's 72 devices. In addition to NETBIOS, the Token-Ring supports the SNA LU 6.2 programming interface, Advanced Program-to-Program Communications for the IBM PC (APPC/PC). This support indicates that the network is likely to be the cornerstone for all of IBM's future office networking plans.

On May 8, 1984, IBM released a statement of direction in which it announced its intention to implement a star-wired, token ring LAN using the IBM Cabling System within two to three years. On October 15, 1985, the company solidified the statement of direction by introducing a token ring PC LAN. Not all came off according to plans -- IBM was forced to make a concession to other PC local network vendors whose products run on inexpensive telephone wires by introducing an alternative medium to the Type 1 and 2 cables of the IBM Cabling System. Furthermore, products to connect 327x devices to the Token-Ring were conspicuous in their absence.

Six months plus one day from the original IBM Token-Ring Network introduction October 15, 1985, IBM announced major additions and enhancements to its strategic network. On April 16, 1986, IBM significantly enhanced its LANs with new LAN host, departmental systems, and Personal Computer attachments, and with products and programs that support IBM's Token-Ring Network and broadband PC Network.

With these new Token-Ring product additions, with the original products introduced on October 15, 1985, and with additional products (e.g., direct 327x-Token-Ring connectivity) whose introductions are anticipated for the last half of 1986, the IBM Token-Ring Network could well be the major driving force within the local area networking field far into the next decade.

3. INCOMNET

Intelligent Communications Networks, Inc. (INCOMNET) was founded in 1974, as a communications consulting firm. The company designed an on-line data collection and distribution system for Arbitron in 1975 and created a global communications network for Lockheed in 1977. INCOMNET initiated efforts to establish its own network in 1979, but the lack of a properly designed, commercially available communications computer for expansion was the impetus for the company's undertaking the development of the INCOMNET 3000.

INCOMNET went public in 1981, as it began developing the INCOMNET 3000 under the leadership of Charan S. Lohara, INCOMNET's president and chairman of the board. Today, the company is in the process of making the transition from a research and development orientation to that of manufacturing, sales, and providing network services.

During 1985 INCOMNET has started to attract the interest of some major companies. INCOMNET has signed a contract with Pacific Bell to deliver over $5 million of hardware and software support during the next two years, and it has also signed a contract with the Taj Hotel Group to develop and install a hotel management and communications system. Further, INCOMNET recently reached an agreement with Telecommunications Consultants India, Ltd. (TCIL) and Pacific Telesis International to explore opportunities for joint marketing and sales.

The INCOMNET 3000 family includes two types of machines: FNP and INP. The Front-End Network Processor (FNP) connects to one or two personal computers or terminals, and the FNP(s) are connected to the closest Intelligent Network Processor (INP) in a single star or multi-star configuration. By adding FNP and INP systems to the network as users are added, INCOMNET can create a modular system for thousands of users.

The INP provides network control plus processing and routing of all messages to the appropriate stations. As the need develops, the unit can accommodate expansion blocks (controller board sets, containing additional communications lines, memory, and processing power). Whether expansion blocks are added to empty slots within the same cabinet or an additional cabinet is added, they can function as one integrated system through INCOMNET's Tightly Coupled Local Area Network (TC-LAN). Stations such as FNP(s), IBM PC(s), or compatibles, terminals, and printers may be added at will.

Each INP contains up to 126 Distributed Processing Controllers (DPC(s)) with CPU, memory, and disk storage. A basic configuration contains three 1200 baud communications lines. Through expansion, this number can grow to over 300 lines. The INP can also route messages, employing the most economical means available (direct dial, WATS, leased line, or satellite).

The FNP is used in each remote site and provides the connection between a user's terminal or microcomputer and INP. The FNP contains a Z-8000 microprocessor, up to 192 Kbytes of main memory, two RS-232 serial ports, a parallel port, and an RJ-11 telephone jack.

A three-day memory power back-up is included. Messages prepared and sent are time- and date-stamped. The FNP is equipped to handle automatic dialing and log-ons, and automatic answering, and it can send and receive messages. It also acknowledges receipt of messages and can operate unattended.

The FNP makes use of windows, formatting and immediate terminal response during message preparation and reception. An instant hard copy is available to any user with a printer.

Both the FNP and INP can send messages immediately, with a forced delivery to each other, or can store messages and send them when telephone rates are lower. This gives INCOMNET an advantage over those competing electronic mail systems that hold messages in a mailbox until the addressee calls for them, which lessens the systems' immediacy.

Much of INCOMNET's software was developed by its subsidiary in India. INCOMNET has a staff of 83 in Northern India, developing software and manufacturing its hardware. The Indian government has a one-third interest in INCOMNET India Ltd.

Software for the INCOMNET 3000 includes the DCOS operating system, advanced message switching software, an interactive forms package, microcomputer interface package, network management software, and a distributed database system. The commonly used programming languages -- Pascal, C, and Ada -- are provided.

In addition, the INCOMNET 3000 can run the standard UNIX operating system. This allows a variety of business applications to run on the system and will allow INCOMNET to use the same hardware to sell UNIX systems.

3.1. TC-LAN

INCOMNET's TC-LAN distributes communications and information processing resources across a collection of interconnected, cooperating, and independent microprocessors to provide distributed processing. The high performance communications capabilities necessary to effectively implement a distributed processing communications network are obtained through the use of Zilog VLSI Chip Sets from the Z8, Z80, and Z8000 product lines. They provide communications speeds ranging from 50 bps asynchronous to 56 Kbps synchronous and support multiple protocols including HDLC and SDLC.

The high-performance communication microprocessors are interconnected via the TC-LAN, which combines a dual-ring topology with a token-passing networking access method over twisted-pair wired (see Figure II-15). Messages travel between microprocessors through each ring, providing a shared resource structure for microprocessors, memory, and peripherals. The system architecture provides simultaneous communications between microprocessors with immediate point-to-point acknowledgement.

Multiprocessing and multitasking is achieved by combining the high performance communication microprocessor, the TC-LAN, and INCOMNET's Distributed Computer Operating System (DCOS). This combination allows system configurations to grow from a single stand-alone microprocessor-based networking system to a multiprocessing system that is capable of interconnecting up to 126 microprocessors that provide multitasking functions.

INCOMNET's TC-LAN is built around controller boards that act as independent microprocessors. There are seven controller boards manufactured by INCOMNET that can be used to configure a TC-LAN system. In addition, there are any number of peripheral controller boards that can be installed. These controller boards can be combined with disk drives, displays, desktop computers, and other peripherals to provide information delivery and information processing solutions.

The two controller boards that are directly associated with the operation of the TC-LAN are the Distributed Processing Controller (DPC) and the Expansion Subsystem Controller (ESC).

The DPC features the Zilog 16-bit Z8001 microprocessor and a full array of support devices from the Z8000 family. The DPC includes a 16-bit on-board data bus with 24-bit main memory addressing and is configured with 256 Kbytes RAM, four serial ports, and three parallel I/O ports.

The ESC connects directly to the DPC board, forming a two-board set. The ESC is divided into two parts: a memory expansion area and the loop interface subsystem. Memory can be expanded to 512 Kbytes using 64K chips, or up to 2 Mbytes using 256K chips. The loop interface subsystem transmits and receives data at speeds up to 1.2 Mbps over each of the two interprocessor rings. The ESC contains the dual-ring transfer logic and therefore does not burden the host processor on the DPC (see Figure II-16).

INCOMNET builds three other controllers that allow interfacing with the TC-LAN.

The Local Network Processor (LNP) connects one of the DPC's parallel ports to an IBM PC-compatible LAN, providing a 19.2 Kbps link between it and the TC-LAN.

The Low-Speed Line Termination Controller (LTC) is composed

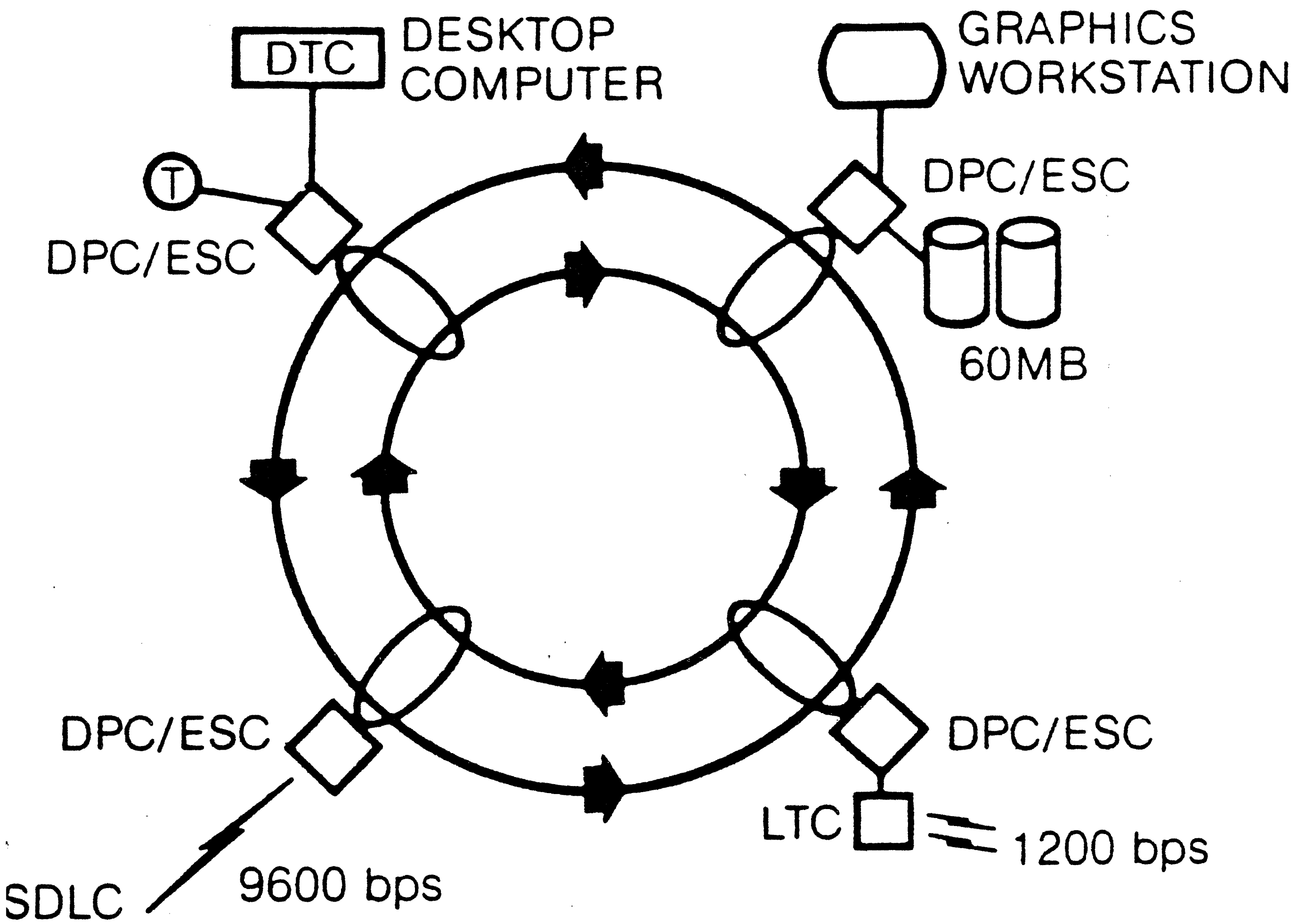

Figure II-15: TC-LAN Topology

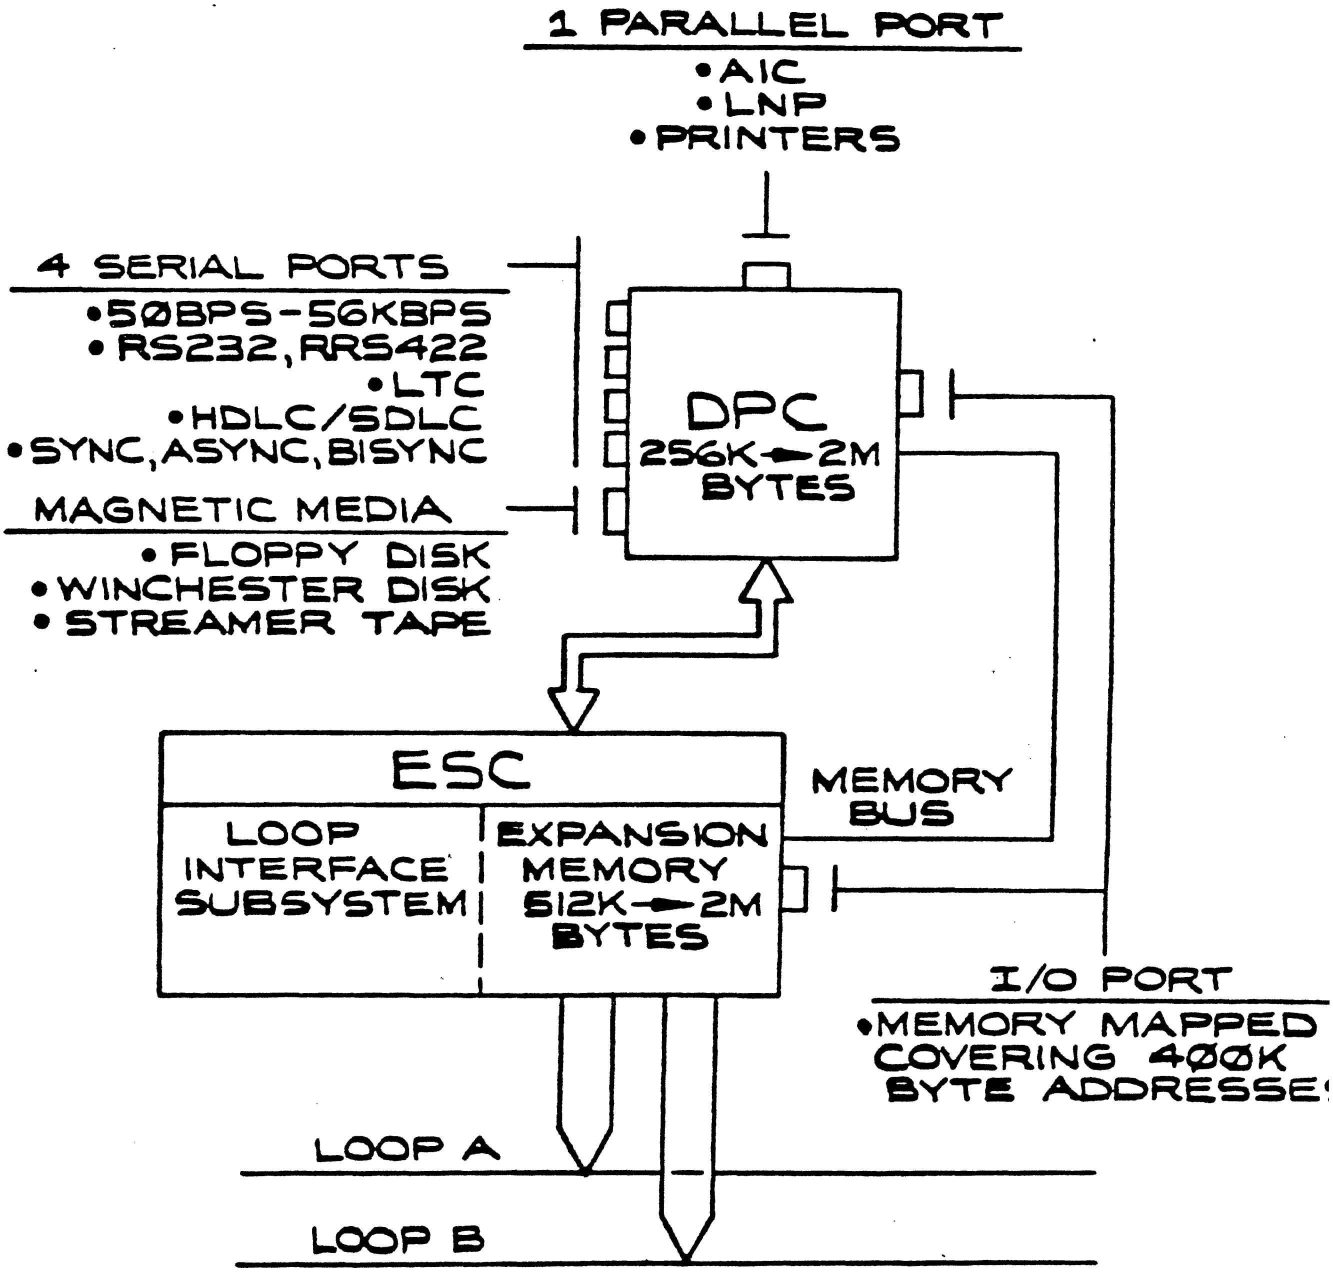

Figure II-16: DPC/ESC Block Diagram

of three sections: the processor section has two RS-232 ports and additional memory; the test and control section has a DTMF, a digital/analog and an analog/digital converter, and a call progress tone detector; and the modem section has an SCC modem, which is Bell 103- and 202-compatible, and an RJ-11 jack.

The Analog Interface Controller (AIC) is much like the LTC but has an analog switching section instead of a modem section. It also has a high-speed parallel port for communication with the ESC. The analog switching arrangement allows up to four daughter-boards to be attached. These boards can be an asynchronous Bell 103- or 202-compatible modem, a synchronous modem, additional memory up to 128K RAM plus a clock, telephone handset interface, or other analog devices.

The LTC and the AIC can be used to store, forward, and control message transfer between users hooked up through analog devices and those attached directly to the TC-LAN.

3.2. Operating System

The heart of the INCOMNET TC-LAN is its Distributed Computer Operating System (DCOS). DCOS is a UNIX-like software set implemented in Pascal. DCOS includes both information processing capabilities and information delivery links. DCOS, unlike other operating systems that support LAN services, also distributes processes and balances resources that include CPUs and memory on the network.

DCOS treats a collection of autonomous computers as a single resource. All networking activities are transparently managed by DCOS; the user never needs to know where a process is running.

DCOS controls a multiprocessing/multitasking software environment and is designed to be distributed across a collection of interconnected microprocessors. One of the distinguishing characteristics of the DCOS environment is that single microprocessor systems can be expanded to multiprocessor distributed processing systems without reprogramming. The DCOS integrated file handling provides consistent file access within a single microprocessor as well as in a distributed multiprocessor environment.

DCOS is a process-oriented operating system. A process is the basic unit of activity in DCOS, which is similar to a "job" or "task." The system and all application programs are represented as processes, since the process is the owner of system resources such as memory, file access, processor time, channel access, and peripherals. A process contains its own address space, execution context, working registers, and program counter. The process structure permits modular construction and expansion within the operating system. New features may be installed or removed by adding or deleting the appropriate process dynamically, while the system is in operation. Location independence of processes permits DCOS to execute them anywhere in the network, making possible microprocessor load-sharing, and it also permits data to be distributed over the network in file resource balancing.

The overall structure of DCOS can be viewed as a sequence of layers. The innermost layer is the kernel of the operating system. The kernel creates copies of processes, schedules their use of microprocessor time and memory, performs control functions, and provides interprocess communication.

The second layer includes all processes for I/O and support services. These processes include low-level I/O drivers and physical/logical device I/O interaction with the intermediate/high-level access method processes.

The third layer provides network services, including interprocess communication for processes on different network nodes; it also manages the microprocessor and file resources on a network-wide basis.

The fourth and outermost layer provides the direct services to the user and to the application programs. This layer includes command processing, spooling, and other support services such as program development and text processing (see Figure II-17).

3.3. Interpersonal Messaging System

INCOMNET's Interpersonal Messaging System software has a number of very attractive features such as forced delivery, combined text, graphics, voice, and image, group addressing, and transparent local or remote delivery. It is designed to allow the user to create his own forms, store them, fill them in, file, and forward them. A great deal of flexibility is provided to allow the user to tailor his/her communications to his/her own needs. Plans to be compatible with the X.400 electronic mail standard as it is developed will allow compatibility with other systems and networks.

INCOMNET offers a Mail/Message Starter Series (MMS) which is designed to send and receive electronic mail and other data on a store-and-forward basis, providing the capabilities of a large mail and message system at entry-level configurations. Intelligent Network Processors (INPs) accept messages from desktop computers and terminals and store them until they can be delivered to their destinations. The message is taken from the sender, validated on-line for message format and addressee existence, acknowledged to the sender with a unique message identification tag, and then routed to the addressed receiver(s). If requested, reports of the message delivery can be generated and routed to the sender.

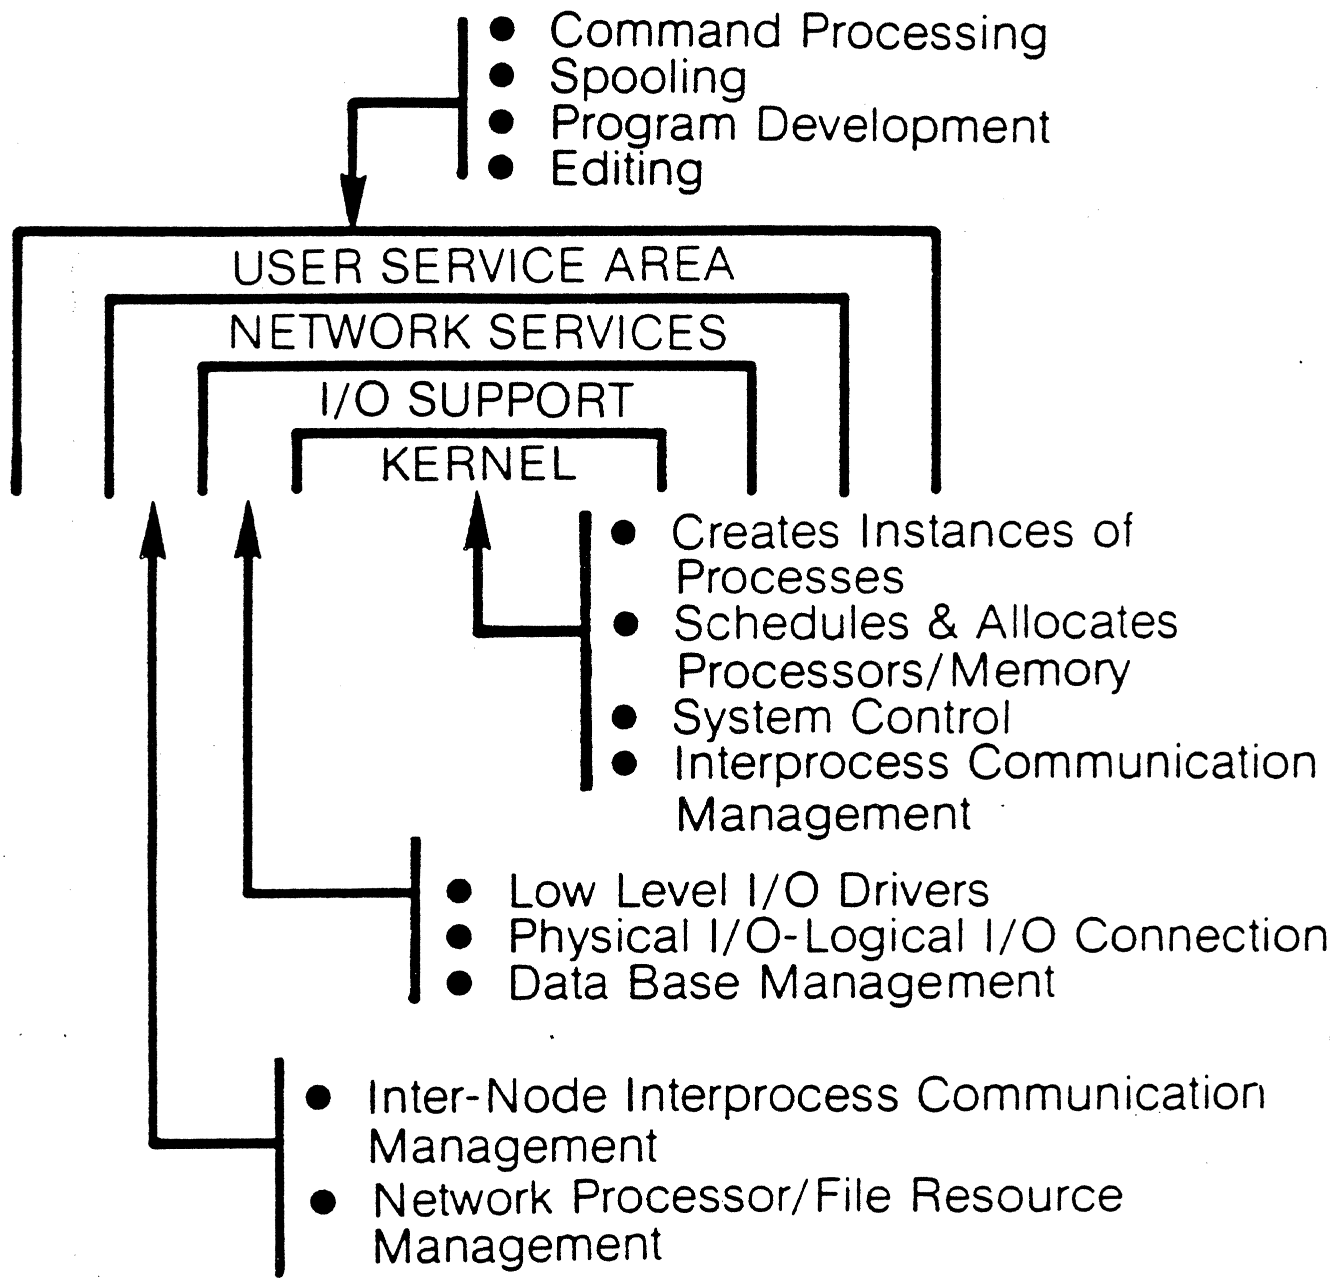

Figure II-17: DCOS Structure

The INP's integrated mail box provides both a central and distributed capability so that the user can match communication requirements to the available nodes. The INP is able, using diagnostic hardware features, to diagnose a problem with a minimal interruption to operations and provide detailed and hierarchical summarization reporting.

One of the strengths of the INCOMNET MMS is the building-block structure. A user can gain economies of scale while building more power into the MSS by connecting multiple INP systems with the INCOMNET TC-LAN. Should the need for the MMS go beyond the local environment, INPs can be dispersed geographically and interconnected through switched or leased line communication services.

The combination of INP hardware, operating software, MMS application software, and the TC-LAN provides the user with a uniquely distributed set of mail and message system capabilities with growth from a single INP system to a multiple INP global network. The distributed processing architecture allows for distribution of files and programs operating within each INP, while providing load and resource sharing across multiple INPs.

The communications solution of the INP can provide protocol conversion between dissimilar devices and speed matching for the various grades of communications facilities.

INCOMNET makes its Interpersonal Messaging System available in two ways: a product called YourNet, which involves selling the equipment directly to end-users who then operate the network themselves, and OurNet, which is a custom-built network owned by INCOMNET but subscribed to by end-users.

OurNet is basically for mail users with the desire to have private data communications in specific vertical markets. YourNet, on the other hand, uses the same hardware and software, but INCOMNET sells and installs the equipment and then lets the users operate the system. In addition, INCOMNET will OEM/private-label OEM its hardware and software.

3.4. Pricing

Prices for the INCOMNET 3000 Series range from $23,995 to $250,000, depending on the size of the system. OEM prices are subject to negotiation.

OurNet prices are based on the number of subscribers, message traffic, equipment provided, and a few other variables, and can run as low as $250 per month.

3.5. Evaluation

INCOMNET has targeted the rapidly growing information communications market and has developed a well-conceived and modular product line to be a strong player in this market. The key to the long term success of any new company in the computer or communications market is the ability to follow-up on its initial product with a continuously evolving and growing product line. By building the INCOMNET series of network products around a standard processor, on a standard processor board, in uniform and adaptable cabinets, and the like, INCOMNET has positioned itself to do just that. By implementing distributed processing concepts to which others just pay lip service, INCOMNET can put together the components to meet a wide range of application and user requirements and environments.

The INCOMNET 3000 distributed system is positioned in an area that has not yet seen many products. Coupling a wide area network with a tightly coupled local network provides the type of communications environment that most Fortune 1000 companies need. Many vendors have targeted the local network market or the wide area market, but not both. INCOMNET's strength is the compatible coupling of these two communications areas.

INCOMNET has chosen to go with its proprietary token ring local network for local communications. Under control of its Distributed Computer Operating System, the local network affords the true control of processes anywhere in the system. Although the current thrust of the local network market is toward local networks that follow one of the developing international standards such as IEEE 802.5, INCOMNET can easily supply a gateway to any of these standard networks as well as offer the traditional RS-232 connection into an INCOMNET Front-End Network Processor (FNP) or Intelligent Network Processor (INP).

The distributed database software that will soon be available for the INCOMNET 3000 will significantly enhance this product. More and more attention now is being given to the database system as the key to most applications. By being able to offer a truly distributed database system in a local or wide area network environment, INCOMNET will appeal to those who have identified applications that need to be implemented but for which the tools to do that do not exist.

4. Logica VTS POLYNET

Logica VTS Ltd., one of the Logica Group of Companies, was formed as a separate Logica Limited group in 1975 and became a separate operating company in February, 1979. Logica VTS's charter is to design, develop, and manufacture word processing products and future office systems. Headquartered in London, the company currently operates with a staff of 220 and a L5.5M budget. It also has acquired a 60,000 square foot factory and warehouse in Swindon, England, in which modern production and test equipment have been installed, to meet anticipated product demand. At the present time, Logica VTS is negotiating with a U.S. company for distribution, maintenance, and support of its products in the North American Market.

The POLYNET local network offered by Logica VTS is based on the Cambridge Ring technology developed at the Computing Laboratory at Cambridge University. Logica VTS's experience with the prototype network based on this technology and used for distributed word processing and for software development became the basis for development of the POLYNET product line that was announced in February 1981.

4.1. Network Structure

POLYNET is a ring network that interconnects up to 254 nodes using twisted pair telephone grade cable. Operation is at 10 Mbps with distances of up to 350 meters between nodes. A typical POLYNET system as shown in Figure II-18 consists of network nodes, a monitor station used for network initialization and monitoring, wall sockets that allow connection between the cable and the nodes, and an appropriate number of power supply units. A number of interface units, including intelligent microprocessor-based units, are currently available to connect computers to the POLYNET. Currently, interface units exist for the DEC PDP-11, LSI-11, VAX, and the Intel MultiBus systems. Two kinds of POLYNET interface units are available and are described later in this report: DMA variety offering point-to-point data rates of 1 Mbps, and a program interrupt interface unit with a data rate of between 100 and 300 Kbps.

In the original Cambridge ring design, three components were identified: a repeater, a station unit, and an access box. Logica VTS combined the repeater and station unit onto a single card and called this the POLYNET network node. The Cambridge ring access box is the POLYNET interface unit. Logica VTS's contribution to the Cambridge ring approach is the new interface. This interface unit design provides for easier customer attachment of equipment as well as allowing Logica VTS to build and easily attach network products.

POLYNET uses the empty slot technique, in which a small number of bits in packets or slots continuously circulate around

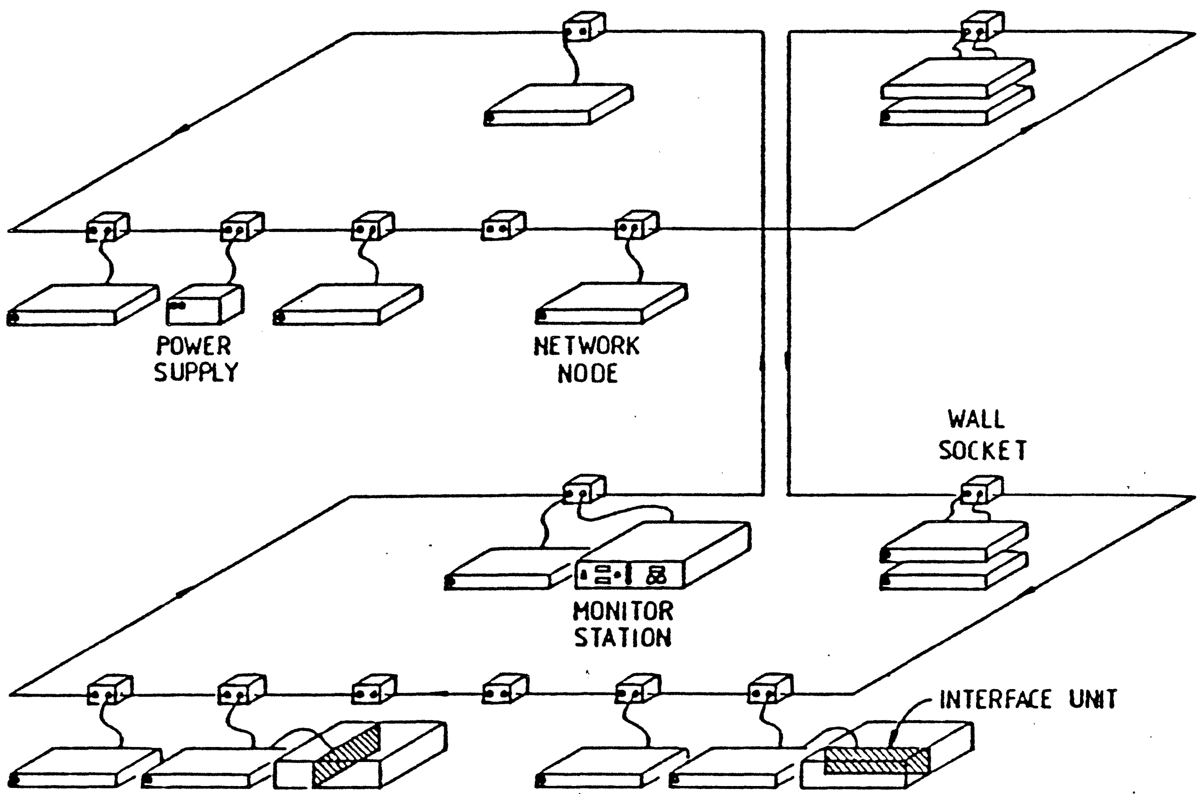

Figure II-18: POLYNET System

the ring. POLYNET nodes are responsible for filling and emptying these slots in order to provide the appropriate communication and delivery of byte streams. Given the bit delay through a node and the bit delay passing through the cable connecting two nodes, a maximum of 16 packets or 40 bits each can be supported on POLYNET. Although POLYNET has been designed with 40 bit slots, it is possible to configure the network to match the older 38 bit slot of the Cambridge ring. These slots, or "mini-packets" in POLYNET terminology, contain a start-of-slot bit, a packet full/empty bit, and a monitor pass bit to guard against packets circulating indefinitely around the ring. There are eight bits for a destination address, eight bits for a source address and two bytes or one 16-bit word for user data. At the end of each mini-packet there are two bits with which the user can identify type of data, two bits that the receiving node can use for a packet acknowledgement, and one maintenance bit that is used for a parity check on the packet as it passes through each node.

A node that wishes to transmit on the POLYNET network monitors the passing packets until a packet marked "empty" is encountered. The empty bit is changed to full, the address and data bits are entered, and the packet is passed to the next node. A receiving node monitors the passing data stream, and when a packet addressed to it arrives, it is copied into registers within the node. The response bits in the passing packet are sent to indicate that the packet has been accepted. If the destination node is busy or has decided to refuse that packet from that particular source, the response bits are set to indicate that the packet has not been accepted. The packet continues around the ring until it reaches the source node where it is marked as empty. The source node also checks to see if the packet has been accepted and whether the address and data fields are identical to those originally transmitted. A node must wait until another slot has passed in order to transmit another packet or retransmit the same packet. This mechanism prevents a node from "hogging" the ring. Nodes can transmit a mini-packet into one slot and receive packets addressed to them in the next slot, thus making the ring effectively a full duplex communication medium.

For IBM PCs on Polynet, each IBM Personal Computer is attached to the Polynet network by means of a standard Polynet network node together with a Polynet IBM PC interface card. This card will drop into any available expansion slot in the IBM PC, and uses DMA Channel 1 or 3.

4.2. Node Structure

POLYNET's network node is implemented in a single desk-top or rack- or wall-mounted box, which is directly connected to the network cable and the interface unit of the attached device or computer by a single ribbon cable. The POLYNET node, shown in Figure II-19, consists of two components, a repeater and a

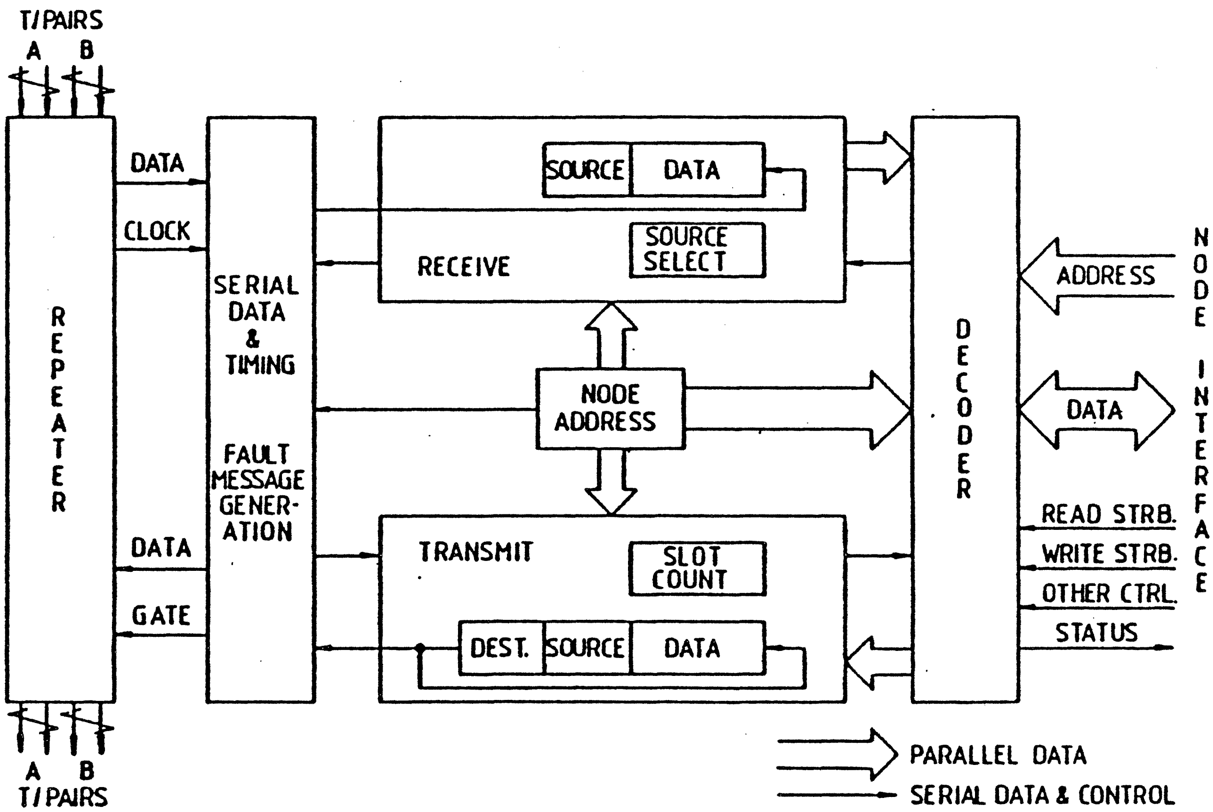

Figure II-19: POLYNET Node

transceiver. The repeater, connected to the network cable at each node, demodulates and remodulates data around the network, delivers serial data and clock signals to the node logic, and accepts gating and data signals from the node logic. The repeater, powered from the network via an integral DC-DC converter, is independent of local power applied to the node logic. Thus the network can continue to operate in the presence of a local power failure. The transceiver portion is responsible for mini-packet reception and transmission under the control of the attached device via the interface unit. On the receiving side, the node transceiver continuously monitors the incoming data stream. It looks for and recognizes packets corresponding to that particular node address, copies the packet into its buffer, adjusts the appropriate response bits, and signals the interface units as required. A source selection register is included to enable each network node to select one transmitting node from which it will accept packets. This register can also be set to a value that will enable that node to accept all packets addressed to it or all packets transmitted in a broadcast node. In the transmission mode, the transceiver logic accepts packets from the interface unit and inserts them into the first available empty slot. Packets that circulate around the ring and are returned are compared with the original transmitted message, their response bits are copies, the slot is marked empty, and the interface unit is notified as required.

A parity checking capability is included with each node logic in order to determine and locate faults in the network. Each fault detected causes the transmission of a fault packet to destination zero in the next empty slot. If the user has installed a logging station at node zero, sophisticated error analysis may be performed. If no logging station exists, such messages do not have any effect.

The POLYNET network node provides a number of improvements on the original repeater and station unit design of the Cambridge ring. These include automatic retransmission upon receiving a busy response, automatic slot counting, and broadcast node addressing.

Another improvement on the Cambridge design is in the interface unit. This interface is an 8-bit bi-directional data bus, a 4-bit wide address bus, and a set of command and status lines. This node interface may be mapped onto a conventional microprocessor input/output structure by means of a very simple interface unit. This node interface allows attached station access to load and examine the source select register, examine the node address, load and transmit a packet, receive and examine an incoming packet, perform various control functions, and examine status indicators.

4.3. Monitor Station

A monitor station connected into the network is used to provide a number of network initialization and maintenance functions. These include POLYNET initialization, full packet counting detection of and recovery from framing errors, automatic adjustment to the insertion and removal of nodes while POLYNET is in operation, automatic restarting of the ring when heavy error bursts are detected, continuous circulation of random data in empty slots to provide continuous error checking capabilities, and power for repeaters on a small network. Up to 8 additional nodes can be powered from the monitor station. The monitor station may be configured for a system containing only POLYNET components, or for an open system containing other equipment.

4.4. Power Supply

If more than 8 nodes are installed or a higher degree of reliability is desired, extra power supplies can be added. The power supplies are connected to the electrical supply and are switched on and off, through an internal relay, in tandem with the monitor station. Each power supply includes a standard POLYNET repeater that can be powered from the network in the usual manner. Thus the power supply can be placed at any convenient point on the network and can be used to increase the distances between nodes to several hundred meters. Distances in POLYNET can also be extended using higher quality tables.

4.5. DMA Interface Unit

Logica VTS offers a POLYNET DMA interface unit for the PDP-11 based on a proprietary front-end processor manufactured by Associated Computer Consultants of Santa Barbara, California. This interface unit allows low-level POLYNET protocols and some of the PDP-11 input-output queuing to be performed by this front-end processor. The front-end processor uses direct memory access (DMA) to transfer data between its memory and a network and between its memory and that of a PDP-11. Data transfer between the front-end processor and the network occurs at POLYNET speeds. The interface unit consists of two boards which plug into a standard PDP-11 back plane.

4.6. PI Interface Unit

POLYNET offers a program interrupt (PI) interface unit designed to provide a simple, flexible interface between the DEC LSI-11 bus and a POLYNET node. The PI interface unit allows program controlled data transfers between an LSI-11 and POLYNET, and provides bus interface and control logic for interrupt processing. This particular unit can be either polled or used in the interrupt node. Polling is performed in the standard LSI-11 fashion, with the interface status information being byte or word accessible. In the interrupt mode, an interrupt occurs when a received packet is available at the node and ready for input to the processor, or when a packet has been transmitted and the transmitter is once again ready. Both interrupts function under program control. This PI interface unit for the LSI-11 is constructed on a single board.

4.7. Multibus Intelligent Interface Unit

The principal addition to the Polynet product range during 1982 was the Multibus intelligent controller. This intelligent DMA interface consists of an Intel 8086 processor board with its own ROM and RAM that connects to an Intel Multibus. It has a Polynet interface designed to support high-speed data transfer and a pair of Multibus registers to support interprocessor communication.

In the form of an intelligent DMA interface, the interface unit connects a system using the Multibus to a Polynet node. Controller firmware provides a Basic Block Protocol (BBP) service as defined for Polynet. BBP data transfers may take place at up to 1 Mbps, corresponding to the maximum point-to-point data rate of a two-slot Polynet network.

The Multibus intelligent interface unit will send and receive one or a stream of BBP packets on command from the host. A large number of commands may be outstanding at any time. The interface unit will enforce receive time-outs. Additional operations allow considerable flexibility in the use of the interface unit.

4.8. Multibus PI Interface Unit

The Multibus program interface unit is designed to provide a simple, flexible interface between a Multibus and two Polynet nodes. It allows program controlled data transfers between a Multibus processor and two nodes and provides bus interface and control logic for interrupt processing.

The interfaces are independent and can function in either polled, delayed acknowledge, or interrupt mode. Polling is accomplished by examining the interface status information available from a Multibus I/O port. Delayed acknowledge mode is essentially a hardware implementation of the software polling mode; the current instruction is halted until the interface status is correct or a transmission error occurs. In the interrupt mode, an interrupt occurs when a received mini-packet has arrived at the network node. Interrupt mode may be selectively enabled and disabled. The interface operates only in the 8-bit Multibus mode.

4.9. Polynet File Server

The Polynet File Server is a software package that can run on a Personal Computer equipped with at least 128K memory and suitable disk storage. It provides a Unix-like file structure and incorporates multiple level hierarchical directory structure, mountable-demountable volumes, special files for direct volume access, and time-stamped file access.

Files are referenced by the standard Unix pathname technique. Each file may be up to 1 Mbyte; each independent file system volume can be up to 32 Mbytes.

Included with the File Server is a comprehensive set of administrative tools to allow file system creation, integrity checking of individual file system volumes, and volume copying for back-up and archiving.

The Polynet File Server is designed to provide very high performance file access by user work-stations. It incorporates a sophisticated cacheing strategy to minimize disk access for frequently referenced data.

The Polynet File Server is a "drop-in" software package that operates concurrently with the standard MS-DOS system in the IBM Personal Computer. There is no need to dedicate a PC to act as the network File Server.

Since the Polynet File Server accesses the disks via standard MS-DOS system calls, it will operate with any type of disk storage attached to the IBM Personal Computer provided that a suitable MS-DOS software driver is available.

Using currently available hard disks for the IBM Personal Computer, a File Server can be configured for up to 100 Mbytes. For users having requirements for very large file storage, or higher performance, a version of the Polynet File Server is available for 8086-based Multibus systems. This allows industry standard SMD disks and controllers to be used offering up to 1200 Mbytes per file server.

4.10. Polynet MS-DOS Upgrade Package

The MS-DOS Upgrade Package provides the link between MS-DOS running on each IBM Personal Computer and the Polynet File Server. It provides totally transparent access to the File Server for file creation, read-write access etc, and loading of MS-DOS utility and application programs stored on the File Server.

Since the File Server offers a Unix-like multi-level directory structure, the MS-DOS upgrade package provides up to 16 drive specifiers, each of which can be mapped onto any part of the File Server directory hierarchy or local floppy disks attached to the IBM Personal Computer. Simple commands are included to display and alter the drive specifier mapping at any time. The mapping establishes a logical connection between a drive specifier and a File Server "pathname". In cases where multiple Polynet File Servers are located on the same network, then each mapping would also specify the associated File Server.

The Polynet MS-DOS Upgrade Package is loaded from floppy disk along with the standard MS-DOS system and requires no special installation or configuration procedures. The package currently operates with MS-DOS version 1.1; an enhancement to work with MS-DOS version 2.0 is being developed and is expected to be available shortly.

4.11. Polynet Optlyne Fiber Optic Ring Extenders

Optlyne units are Polynet repeaters which, in the process of retransmitting the ring signal, convert it from bi-polar code on twin twisted-pair copper cable to self-clocking code on a fiber optic cable, or vice versa. They are used to create mixed optical and copper networks with the most appropriate transmission medium being used for each part of the network. Polynet optical ring segments enable network distances to be considerably increased; using the standard-grade (4 dB/km) fiber optic cable supplied, distances up to 2 km may be spanned between a pair of optlyne units. Lower attenuation cable is available if greater distances are required. These ring segments are also totally immune to interference. Whatever the length of the segment, network data transmitted via optical cable cannot be corrupted by external influences. This provides complete freedom to route the cable via existing power or telephone ducts and may therefore provide quicker and lower cost installation.

To provide full flexibility in configuring mixed networks, optlyne units are available in three versions:

Optlyne-T: This simplex transmitter provides a termination for the twisted-pair network cable, demodulates the received signals to extract serial data and clock, and re-encodes them into a self-clocking code which is transmitted optically via a single fiber.

Optlyne-R is a simplex unit which performs the complementary function to Optlyne-T, receiving the optical signal and retransmitting the data in Polynet bipolar format via the copper network cable.

Optlyne-D is a duplex unit performing both transmit and receive functions.

For simplex (one-way) ring segments, Optlyne-T and -R units are used together in pairs. For duplex (two-way) spurs, Optlyne-D units are used in pairs. More complex configurations can use mixtures of different units.

4.12. Pricing

Description	Price List
VNN Network Node	L 833
VMS Monitor Station (including power supply)	L2,800
VPS Slave Power Supply Unit	L 480
VQP LSI-11 Interface Unit	L 390
VUI PDP-11 DMA Interface Unit	L3,500
VMI Multibus Intelligent DMA Controller	L1,500
MOD Duplex Opthyne	L2,190
Polynet File Server	
Polynet MS-DOS Upgrade Package	
Other attached station interface units	by quotation
Planning and Installation Advice, Cabling, Wall Sockets, and Installation	by quotation

Prices are quoted exclusive of sales taxes and delivery charges and are subject to change without notice.

4.13. Evaluation

Logica VTS now has a mature range of POLYNET networking components. The node interface has now been adopted as the Type 1 CR82 standard. Interfaces for various personal computers have led to additional applications and provide the means for interworking of 16-bit micro computers and larger processors. The file server allows true file sharing capability between various network machines. The Unix-like file structure is both powerful and easy to use.

5. Proteon ProNET-10

Proteon, Inc., is a privately held corporation that markets a line of token arbitrated, star-shaped ring local area networks. Founded in 1972, Proteon primarily developed advanced communication and computer systems for government and private clients during its early organization. In June 1981 the first ProNET-10 system, a 10 Mbps LAN, was installed.

In February 1984, Proteon received $2.35 million in financing from major high-technology venture capital firms. In May 1985 it received $6.5 million in its second round of financing. These new funds are being used to expand marketing activities and for research and development of Proteon's ProNET-10 and its ProNET-80, 80 Mbps LAN, and its new ProNET-4, IEEE 802.5-compatible, 4 Mbps LAN. The first ProNET-80 sale was made in February 1985.

Proteon started 1986 with Francis M. Scricco in place as the new president and CEO. Scricco oversees the company's total operations, including sales, marketing, manufacturing, and development of the expanding line of ProNET token-ring LANs and networking equipment.

The ProNET local area network interconnects information processing equipment within and between facilities, such as office, industrial, educational, institutional, and campus environments. ProNET is designed to resolve incompatibilities between communicating devices as well as be modular in design, allowing easy network expansion. ProNET host interfaces have been designed specifically for each network, the ProNET-10 at 10 Mbps, the ProNET-80 at 80 Mbps, and the ProNET-4 at 4 Mbps. Other network components, such as cable and Wire Centers, may be shared by the two networks.

ProNET is based upon a token passing, star-shaped ring network architecture. One of the primary benefits of ProNET's star-shaped ring topology is its immunity to single component failures, allowing graceful and predictable degradation in the event of host (node) malfunctions.

ProNET-10 host interface systems are available for UNIBUS, Multibus, VME Bus, Q-bus, and Universal Bus. These interfaces provide user-transparent connection services with minimal delay across a high bandwidth local area network. During host-to-host file transfer, the rate is limited only by the speed of the computer's bus.

The ProNET-10 and ProNET-4 operate over a variety of media including shielded twisted pair, coax, infrared link, microwave, and fiber optic cable. A combination of media may be employed within and between networks.

5.1. Wire Centers

With the passive ProNET Wire Center containing the physical ring, a ProNET system is physically arranged as a series of stars (see Figure II-20). The star shaped ring provides increased reliability and maintainability, while also simplifying cabling. Each Wire Center can connect 4, 8, or 12 nodes and provides local expansion connections for adjacent Wire Centers. Through a quality sealed relay, each node is connected to the ring and operates only when the host interface gives the "join ring" command. Thus a power failure at one node will not bring the ring down, and nodes will not join the ring which have found themselves faulty through the digital or analog loopback test. Similarly, nodes will immediately disconnect from the network in the event of a cable break. The Wire Center requires no power except for the small current from the host interface necessary to activate the relays.

Wire Centers allow easy addition or deletion of nodes without disrupting network operation. The Wire Center also provides automatic bypassing of malfunctioning nodes and facilitates fault isolation -- all adding up to low cost of ownership.

Wire Centers within the 10 Mbps network may be connected to one another via any of the aforementioned media (shielded twisted pair, infrared, coax, or fiber optic cable). The maximum distance between Wire Centers or from host-to-Wire Centers is 480 meters using shielded twisted pair with repeaters. The maximum unrepeatered distance between a single Wire Center and a host is 240 meters. ProNET Fiber Optic Links allow distances of 2.5 kilometers between Wire Centers and/or hosts, while the infrared link provides a 350 meter distance between infrared devices.

Repeaters allow longer distances between nodes by amplifying the ring signal. When used in adjacent Wire Centers, repeaters allow the star-shaped ring to partition into smaller rings upon a cable break.

Fiber optic links are available for those users requiring security and immunity to RFI (Radio Frequency Interference) and EMI (Electro-Magnetic Interference). Each fiber optic unit contains a microprocessor to control the system such that when the ring fails a node will disconnect from the network. If fiber is used between Wire Centers, the ring will divide into two functional subrings should the signal fail. There is no need for exotic couplers or mixers.

5.2. ProNET-10 System

From each Wire Center relay, connections are made to ProNET-10 host interface systems, IBM PC interface, or a Terminal Interface unit (TIU). Each interface or TIU contains the logic

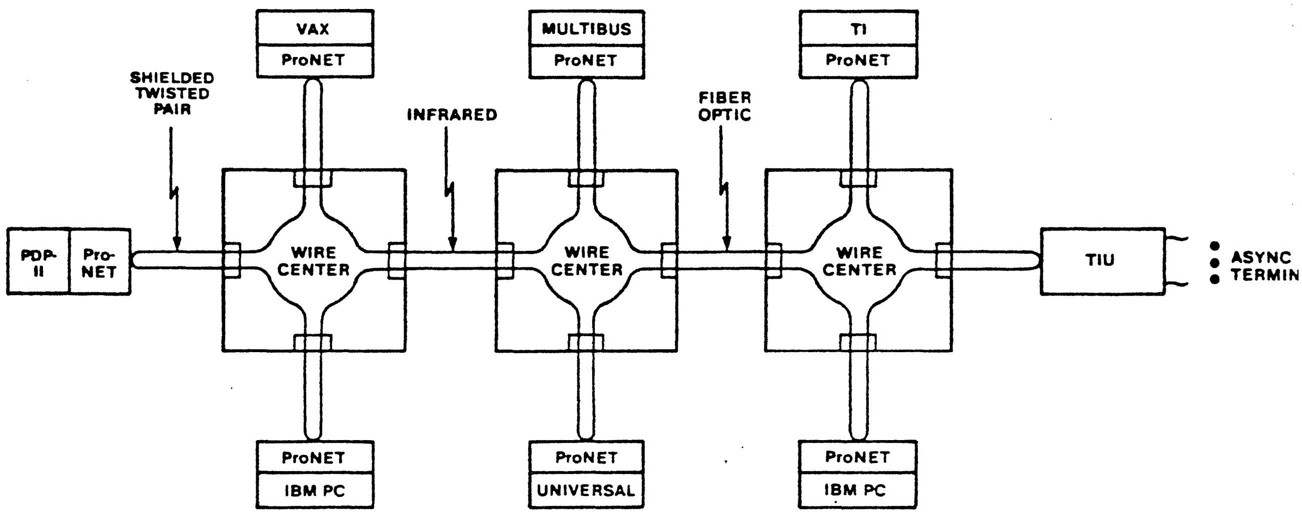

Figure II-20: ProNET System - String of Stars necessary to interconnect an information processing device and/or computer to the ProNET local area network.

5.2.1. Host Interface Systems

The host interface system plugs directly into the computer's bus and provides a high speed link to the token passing, star-shaped ring network. The systems support Digital Equipment Corporation's VAXes, PDP-11s, and LSI-11s using UNIBUS and Q-bus; the entire spectrum of computers employing Multibus; VME-based machines, and computers such as Gould, Perkin-Elmer, and Texas instruments using the Universal bus.

Each host interface system consists of two boards called a host specific board (HSB) and a control module (CTL).

The HSB peforms the communications processing for the host, including sequencing, buffering, and controlling. Thus providing for a low software overhead burden in the host. The HSB contains packet buffers, direct memory access (DMA) logic, control logic, and maintains two 2046 byte (1023 word) buffers, one for transmit and one for receive. Full duplex DMA circuits allow independent DMA transfers between the host's memory and the two buffers. The DMA circuit eliminates the need for any copy loops and keeps the entire addressing space available for operating system use.

Separate transmit and receive control and status registers (CSRs) regulate these DMA operations, and also control the transmission and reception of data between the packet buffers and the ProNET LAN.

The CTL board manages all the network transmission functions for ProNET including modem, serial-to-parallel conversion, repeater, address recognition, parity check, token management, bit stuffing, and error timeouts. The CTL also activates Wire Center relays and has circuitry for self-testing under program control.

When commanded to send by the HSB card, the CTL card waits for the first available token and then serializes the message into the ring, inserting the source address automatically.

The CTL performs hardware address matching, checking all passing messages against the 8-bit node address set by a dip switch. Should the packet destination address of a circulating message match the node address (or the broadcast address) the CTL notifies the HSB. If the HSB enabled COPY, the CTL routes the message to the packet buffers on the HSB. The CTL also repeats the message to the originator clearing the refused bit. This allows the sender to see that the message was passed around the ring and was accepted by the destination node. The sender then drains the message. If the HSB has not enabled COPY, the CTL repeats the message without clearing the refused bit, i.e., receiver did not accept the message. Upon receipt of the message, the sender will drain the message and issue a new token.

The CTL card provides a message parity bit from the source to destination. The parity bit detects rare errors caused by noise. Check-sums may be provided by software on a memory-to-memory transfer level.

The modem is a fully synchronous phase locked loop Manchester system with low error rate and high noise immunity. It is transformer coupled to the transmission medium to reject all signals below 1 MHz. Synchronous operation allows narrowband phase-locked tracking to ensure accurate decoding. The synchronization method does not limit the strength of the ring.

5.2.2. IBM Personal Computer Interface

The IBM Personal Computer Interface consists of a single card which plugs into a large expansion slot of an IBM PC, PC XT, PC AT, or IBM compatible PC. It provides a high performance connection between the PC and the ProNET token ring technology.

Additionally, the PC interface is compatible with all other ProNET host interfaces (UNIBUS, Q-bus, Multibus, VMEbus, and Universal bus). Thus, PCs may communicate with other hosts on the network and if desired have, access to remote host applications. For example, IBM PCs may communicate with another host, such as DEC VAX, by operating with the TCP/IP software.

Additionally, by using the Netware/P software up to 54 IBM PCs (and compatibles) are able to share disks, printers, plotters, and other peripheral devices. Off the shelf PC application software packages can be loaded and immediately put into use. File and record locking prevent unauthorized access to information. Electronic mail is also included.

5.2.3. Terminal Interface Units (TIUs)

The TIU interfaces asynchronous ASCII terminals to the star-shaped ring network. It is available with eight or sixteen RS-232-C connections each operating at speeds up to 9.6 Kbps. TIUs are compatible with and work in conjunction with ProNET host interfaces operating with TCP/IP Server TELNET software.

5.3. Software

There are a variety of software packages available to drive the ProNET host interfaces, depending upon the user's application. Broadly, they may be classified as HDLC (Higher Level Data Link Control) software, TCP/IP (Transmission Control Protocol/Internet Protocol) software, DECNet software, or

NetWare/P.

The HDLC package provides an error-free data stream between two cooperating processes on separate machines: that is, a task on one machine may communicate with a task on another machine. The following DEC operating systems are supported: RT-11 V4.0, RSX-11M V3.2 V4.0, RSX-11M Plus V2.0, and VAX/VMS V3.0.

TCP/IP software provides full file transfer, electronic mail, and remote login functions for several host operating systems, including UNXI V6 and 4.1 BSD VAX UNIX and V7 VAX/VMS. Additionally, UNIX 2.9 BSD and 4.2 BSD are distributed with full ProNET support for TCP/IP protocols. As a part of the TCP/IP protocols, TELNET software allows a terminal user to "log on" to another machine, eliminating the need for complex and expensive RS-232 port switching devices.

The current release of Berkeley UNIX includes a complete implementation of the TCP/IP protocols, which are integrated into the socket communications facilities in the UNIX kernel. Device drivers are contained within 4.2 BSD for the ProNET local area network.

The SUN-2 workstation from Sun Microsystems includes the full 4.2BSD networking system, plus some remote file system enhancements. Proteon offers a device driver for the ProNET P1200 Multibus interface, which allows customers the ability to operate ProNET hardware. Installation requires only a simple system reconfiguration.

A package supporting a useful subset of the TCP/IP protocols is available for IBM Personal Computers and compatibles operating the MS-DOS or PC-DOS V2.0 (or higher) operating system. This package provides User TELNET, allowing the user to log onto other TCP/IP machines on the network. The user then appears to the foreign system as a VT52 terminal. A server TELNET is not provided, as the MS-DOS system is not multi-user.

The package includes both User and Server TFTP (Trivial File Transfer Protocol) programs. This is a very simple file transfer protocol. TFTP can be run as a user program, and there is also a Server TFTP operating under the User TELNET program. Thus one can transfer files to and from his/her PC while logged into another machine on the network.

The DECNet software consists of device drivers which allow the ProNET local area network to be used with Digital Equipment's DECnet software on RSX-11 and VAX/VMS operating systems. Thus the full functionality of Phase II and Phase IV (VMS only) DECnet is available within the ProNET network.

The NetWare/P Operating System allows IBM PCs, IBM XTs, IBM ATs, or IBM compatibles to share one or multiple file servers, each within as much as 252 Mbytes of disk storage. Each file server on the network can support 54 personal computers. Three printers may be connected to the file server and may be used by all network stations.

Advanced NetWare/P expands the NetWare capabilities to embrace multiple file servers coexisting on a potentially heterogenous internetwork. Advanced NetWare retains all of the features that make NetWare such a powerful personal computer local area network environment plus additional features:

- Multiple NetWare File Servers
- NetWare Bridges

Now a user positioned at a workstation is provided a simple, consistent method of accessing network resources (file server, print servers, etc.) regardless of the physical position of the workstation or the resource on the network. Underlying network topologies and interconnections are transparent to user. A workstation may simultaneously use files from multiple file servers, allowing data integration across logical and/or physical boundaries.

The NetWare Bridge consists of a PC which is capable of containing two printed circuit boards in addition to the PC mother board. Those boards consist of one network interface from each type of network. The network interface boards may be of the same type or of dissimilar types. The NetWare Bridge must contain the appropriate NetWare software needed to function with the two network interface boards. Bridges are available for ProNET-to-ProNET networks, or ProNET to some other type of NetWare supported network.

5.4. Infrared Link

ProNET-10 may use infrared communications as a media connection. The infrared (IR) link is an alternative to physical cable connections which provides the same high speed and reliability of ProNET systems.

The infrared devices can link ProNET token-passing networks together or act as a connecting medium in a single star-shaped ring. Information can be transmitted at a speed of 10 Mbps at unobstructed distances up to 350 meters.

Typical applications for the infrared link include connecting ProNET LANs on different floors in adjacent buildings, across highways, and networking equipment in electrically noisy environments, such as factory or automated warehouse. Up to 255 nodes can be connected, even between two buildings without stringing cable.

A built-in microprocessor analyzes the performance of the link and ensures the integrity of the data transmitted in the network. If the signal level falls below a threshold -- for example when the pathway becomes obstructed -- the link will cease transmission and the ProNET system will bypass the node. Thus, a blockage of the infrared link will not affect other network users.

No special permit is required to use the link, nor do the Federal Communications Commission or Bureau of Radiological Health place any restrictions on the device. Transmission distances between 100 and 350 meters are possible depending upon weather conditions. The infrared link may be minimally affected by rain or snow. The infrared device is lightweight, small sized, and easy to install, requiring only a stable platform.

5.5. Microwave Links

The Proteon Microwave Links allow ProNET-10 networks to operate over line of sight distances of up to ten miles. (Greater distances may be achieved by connecting several microwave links back-to-back.) Since each link in a ProNET is independent of the others, the Microwave System may be used with any combination of other ProNET media (such as shielded twisted pair, fiber optics, or infrared).

Two Interface Units are available:

- Single Microwave Link: A single microwave connection between nodes or wire centers on a ProNET.

- Redundant/Hot Standby Interface Unit: Also a single microwave connection, but with two sets of transmit and receive electronics. Logic in the microwave system automatically switches between receivers in the event of a failure, allowing the system to continue operation without loss of any capability.

The Microwave Links connect hosts and/or wire centers. Each link can span a distance of up to 16 km (10 miles). The Microwave Links operate at full network speed.

Once connected, the Microwave Links function automatically. Network electrical signals are converted to microwave signals, transmitted across the link, and reconverted at the other end. No adjustments or monitoring are necessary. No alteration of network software is required.

While operating, the units monitor the quality of the microwave signal and the function of transmit and receive electronics. Should the signal level fall below a set threshold, the transmitter or receiver is automatically removed from service. If the microwave system is a single link, the network then partitions into two independent subnets, as no alternative electronics are available; the entire microwave path is removed from service. If the system is a Dual Link, the secondary receiver is used; the network is not disrupted. Both systems automatically restore the failed components to service as soon as the signal is acceptable.

5.6. Gateways

The Proteon p42000 ProNET-Gateways are a family of gateways designed to support high throughput interconnection of two or more networks of the same or different types. ProNET-Gateways accommodate a variety of communication protocols and provide reliable local and wide area networking in multi-vendor computing environments.

The networks supported by ProNET-Gateways include ProNET-10, ProNET-80, Ethernet, ARPANET networks, and point-to-point telephone lines. Network protocols supported include IP (Internet Protocol), XNS (Xerox Network System), and DEC's DNA (Digital Network Architecture).

ProNET-Gateways provide the following major benfits: 1) they accommodate multiple communication protocols at the network level and thus provide a wide range of network connections while reducing the requirement for additional gateway hardware; 2) they transfer information between networks in excess of 1000 packets per second, supporting high traffic conditions; and 3) they provide connectivity between multiple networks, making ProNET-Gateways very cost effective.

A ProNET-Gateway is based upon a Multibus 68000 processor board with RAM, ROM, timers, and serial lines. This is supplemented with appropriate network interface hardware for connection to the desired local or wide area networks. For example, by configuring the gateway with ProNET-10 and ProNET-80 interfaces, a high-performance connection is established between the ProNET 10 and 80 Mbps networks.

ProNET-Gateway software is highly modular and provides complete packet-switching functionality. The software supports the following:

- device drivers for various network interface hardware
- data link processing for various network-dependent and network-independent data link functions
- forwarders for the supported network-level protocols
- other network level functions, required by various network protocols for routing, error reporting, and maintenance

5.7. ProNET Architecture

The ProNET System's architecture consists of 3 layers as shown in Figure II-21. The client layer encompasses programs and routines that access the network through the data-link and physical signaling layers and the Host Specific Board (HSB), which serves as the physical interface between the user's host computer or processor and the data link layer.

Data encapsulation, network management and error recovery functions are performed by the data link layer. The actual transmission and reception of data on the network transmission media is performed by the physical signaling layer which encompasses signal modulation and demodulation circuitry and the medium itself.

5.7.1. The Client Layer

ProNET's HSB design is intended to minimize the software burden on the host. Message transmission, reception and all communication protocols are executed automatically by the HSB/CTL hardware. PRONET's HSB features a full duplex direct memory access (DMA). The HSB contains two separate 2046 byte packet buffers (4096 bytes on the VMEbus card), plus control and status logic. Full duplex operation provides for concurrent transmitter and receiver operations. The design supports scatter-gather operation and operates at the speed limit imposed by the host bus interface specification.

5.7.2. The Data-Link Layer

The following functions are performed by the data link layer. In addition, the simple design of the control and status register commands, the minimal number of interrupts, and the use of hardware finite-state machines on the HSB, have caused the ProNET-10 hardware to be a consistent speed winner in benchmark tests.

Data Encapsulation
- Message Framing
- Source Addressing
- Bit Stuffing/Destuffing

Network Management and Error Recovery
- Transmission Queueing
- Control Character Detection
- Ring Initialization
- Diagnostic Operation
- Message Error Detection

Figure II-22 illustrates the ProNET-10 data link format and control characters.

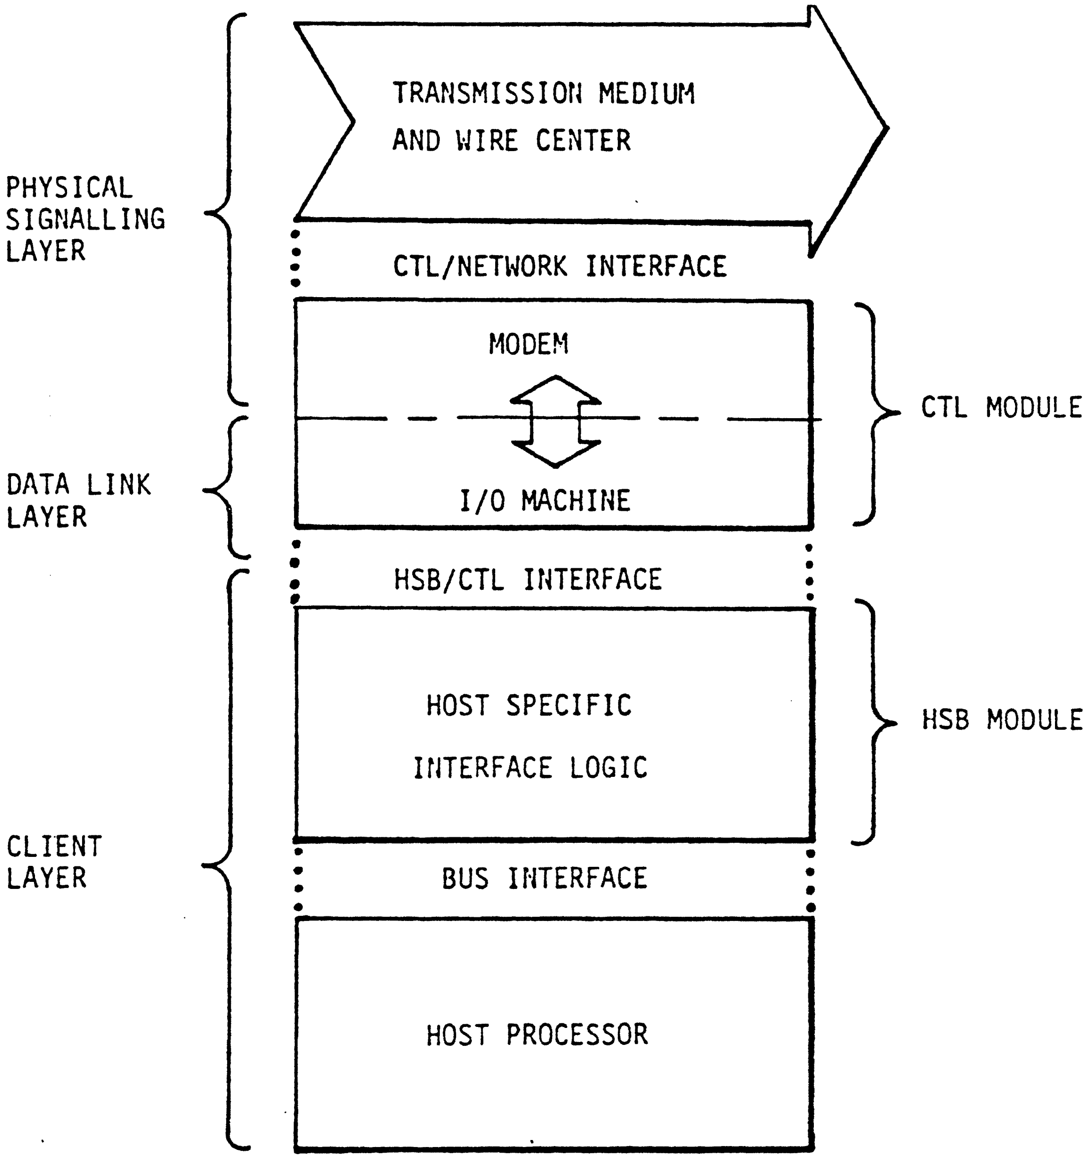

Figure II-21: ProNET Architectural Layering

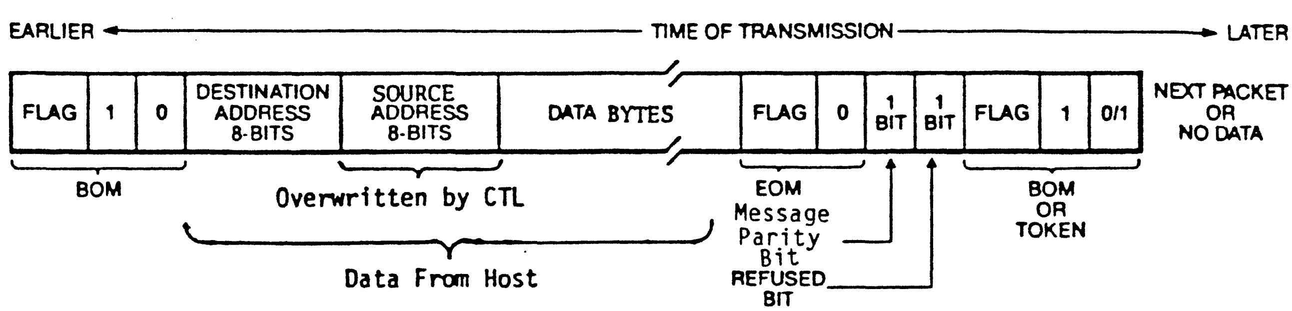

Figure II-22: ProNET-10 Data Link Format and Control Characters

An important feature of the data encapsulation format is the absence of a word count field. Each packet may have a variable number of data bytes, since the EOM control character serves as a data field boundary.

5.7.3. The Physical Signaling Layer

Analog signal transmission and reception is handled at the physical signaling layer.

Physical Signaling Layer Functions

- Biphase Encoding/Decoding
- Transmission Line Interface
- Clock Synchronization
- Network Frequency Control

The biphase modulation technique allows the nodes to be AC coupled. Frequency components below 1 MHz are rejected by filtering at each node so that the network is relatively immune to ambient noise encountered in the operational environment.

5.8. Operating Modes

Four operating modes are available: Repeat, Copy, Originate, and Initialize. In addition, there are two possible test modes, Digital loopback and Analog loopback, as well as the Normal mode. Most of the time, the CTL will find itself in the Repeat Mode. In this mode, the CTL acts as a repeater of data on the ring. The data received from the adjacent node upstream is passed on to the next node downstream on the ring.

The CTL will automatically switch from the repeat mode to the Copy Mode if it determines that the message passing through it is intended for its associated host processor. Specifically, the message will be copied if the destination address in the message corresponds to the address of the CTL. It will also be copied if the message has the broadcast address.

When the host desires to transmit a message to another node, it carries out an Originate operation. This begins with the transfer of the packet to be sent from the host memory to an output packet buffer on the HSB via DMA. Then, upon recognition of the token, the CTL changes it to a BOM, transmits the message and ends it with a token. When the message returns around the ring, the CTL removes it, leaving the following message or token intact.

If the ring network does not have a valid TOKEN control character circulating, the CTL module may create a TOKEN by initializing the ring in response to an Initialize ring request from the host processor. The Initialize Mode is essentially the same as the Originate Mode, with one important exception: The CTL does not wait until an access control TOKEN passes into it before sending the message. In this way, the CTL module forces a message onto the ring. It removes the returning message, leaving only the valid TOKEN circulating on the ring.

The input and output portions of the HSB module are capable of independent operation. Thus, the following scenarios are possible: a node may originate a message while simultaneously emptying its input packet buffer. A node may copy a message while waiting to originate one. A node may originate a message to itself.

In the Normal mode of operation, the CTL receives data from the next adjacent node upstream. It transmits data to the next node downstream.

Two test modes are also provided. These are Digital Loopback and Analog Loopback.

Digital Loopback can be used to test all portions of the CTL/HSB, except for the modem and cable interconnection. All four modes of operation can be checked out in Digital Loopback and all hardware associated with the operation can be tested without actually disturbing the existing network.

The modem portion of the CTL may be checked out in the Analog Loopback mode. It allows the user to test the analog circuitry and the transmission medium to a Wire Center.

5.9. Monitoring System

Proteon now offers a system that monitors and manages its ProNET line of token passing, star-shaped ring LANs.

The Network Monitoring System measures network performance, debugs network hardware and software, and produces long-term records of network traffic. Information from the system can then be used to determine the most efficient ways to use networking hardware and software.

Network performance can be measured by the number of packets on the network or by the percentage of bandwidth utilization. Thus, a communications manager can determine peak operating periods and which department needs more computing devices. Error rates, network uptime, and a listing of nodes on the network can also be used to track the efficiency of the network.

The P2302 ProNET Network Monitor is based on an IBM PC or compatible operating with a ProNET IBM PC Network Monitor software. The PC must be equipped with 128K of memory and at least one floppy disk drive.

5.10. Pricing

ProNET-10 Local Network Interfaces (p1000 series)

Model	Description	1-9 Units
p1000	Unibus System (HSBU/CTL)	$3150
p1100	Q-bus System (HSBQ/CTL)	$3150
p1200	Multibus System (HSBM/CTLM)	$3150
p1300	IBM PC System (software not included)	$799
p1400	Univesal Bus Interface System (UBI/CTLU)	$950

Wire Centers, Repeaters, and Accessories (p2000 series)

p2000	8 Node Wire Center	$1220
p2100	16 Node Wire Center	$2040
p2200	Repeater, CTL-to-Wire Center	$300
p2210	Repeater, Wire Center-to-Wire Center (2 units)	$600
P2302	Network Monitor	$1394
p2900-03	Network Cable Assembly, 3 meters	$50
p2900-10	Network Cable Assembly, 10 meters	$65
p2900-100	Network Cable Assembly, 100 meters	$450
p2900-300	Network Cable Assembly, 300 meters	$1180
p2910	Coax Adapter	$50

Fiber Optic Products (p3000 series)

p3000	Fiber Optic Link, CTL-to-CTL	$2300
p3001	Fiber Optic Link, CTL-to-Wire Center	$2300
p3002	Fiber Optic Link, Wire Center-to-Wire Center	$2300
p3100	Redundant Fiber Optic Link CTL-CTL	$2900
p3101	Redundant Fiber Optic Link CTL-WC	$2900
p3102	Redundant Fiber Optic Link WC-WC	$2900
p3200	Dual Ring Fiber Optic Unit, CTL Node	$1450
p3201	Dual Ring Fiber Optic Unit, Wire Center Node	$1450

ProNET-10 Terminal Interface Units (p4000 series)

p4015	8 Line Terminal Interface Unit with TCP/IP/TELNET Software	$9700
p4025	16 Line Terminal Interface Unit TCP/IP/TELNET Software	$11000

ProNET-10 Software (p5000 series)

p5000-YA	Diagnostic & Echo Server for PDP-11 and LSI-11	$100
p5002-GA	ProNET-10 Diagnostic and Echo Server for VAX-11/750, 730, Source and System Image (RT-11 format on TU58)	$100
p5002-YA	ProNET-10 Diagnostic and Echo Server for VAX-11/780, Source and System Image (RT-11	$100

	format on RX01)	
p5003-22	ProNET-10 Diagnostic and Echo Server for IBM PC, Source and Executeable (MS-DOS 2.0 double-sided 9-sector floppy)	$100
p5100-DA	RT-11 HDLC ProNET-10 Driver (800 bpi)	$700/ne
p5110-DC	RSX-11M,S HDLC ProNET-10 Driver (800 bpi)	$900/no
p5120-GA	VAX/VMS HDLC ProNET-10 Driver for VAX-11/ 750 730 (TU58)	$1000/no
p5130-DC	RSX-11M-PLUS HDLC ProNET-10 Driver (800 bpi)	$1000/ne
p5230	MS-DOS V2.0 TCP/IP Package for IBM PC	$250
p5300-YA	RINGWAY for RSX-11M,S (RX01), copies 1-3 ea.	$2500
p5300-YA	RINGWAY for RSX-11M,S (RX01), copies 4-20 ea.	$1200
p5310-GA	RINGWAY for VAX/VMS (TU58), copies 1-3 ea.	$2500
p5310-GA	RINGWAY for VAX/VMS (TU58), copies 4-20 ea.	$1200
p5310-YA	RINGWAY for VAX/VMS (RX01), copies 1-3 ea.	$2500
p5310-YA	RINGWAY for VAX/VMS (RX01), copies 4-20 ea.	$1200
p5400-MG	UNIX V7 ProNET-10 packet device driver (1600 bpi)	$150
p5500	NetWare/P for MS-DOS V1.1 & V2.0	$1495/ne
p5505	Advanced NetWare/P Upgrade	$295

All prices are F.O.B. Waltham, MA, and subject to change without notice. Some software is available on other media to special order, however the price may be higher and delivery longer.

5.11. Evaluation

Proteon has been successfully selling its ProNET-10 token ring network since June of 1981. While the spotlight has been focused on IBM's grand token ring scheme, Proteon's base of 10,000 installed nodes illustrates that commercial token ring systems have been practical for some time now.

Proteon's ProNET-80 network will allow Proteon to continue to do well even after IBM's 4 Mbps network becomes available in 1986. ProNET-80 offers an order of magnitude performance improvement and will provide connections to the IBM token ring. Its current connections to existing ProNet-10s and Ethernets allow it also to be used as a backbone network in appropriate environments.

6. Proteon ProNET-80

The ProNET local area network interconnects information processing equipment within and between facilities, such as office, industrial, educational, institutional, and campus environments. ProNET is designed to resolve incompatibilities between communicating devices as well as be modular in design, allowing easy network expansion. ProNET host interfaces have been designed specifically for each network, the ProNET-10 at 10 Mbps and the ProNET-80 at 80 Mbps, and the ProNET-4 at 4 Mbps. Other network components, such as cable and Wire Centers, may be shared by the two networks.

ProNET is based upon a token passing, star-shaped ring network architecture. One of the primary benefits of ProNET's star-shaped ring topology is its immunity to single component failures, allowing graceful and predictable degradation in the event of host (node) malfunctions.

ProNET-80 host interface systems are available for UNIBUS, Multibus, VME Bus, Q-Bus, and Universal bus. These interfaces provide user-transparent connection services with minimal delay across a high bandwidth local area network. During host-to-host file transfer, the rate is limited only by the speed of the computer's bus.

The ProNET-80 operates over shielded twisted pair or fiber optic cable. A combination of media may be employed within and between networks.

6.1. Wire Centers

With the passive ProNET Wire Center containing the physical ring, a ProNET system is physically arranged as a series of stars. The star shaped ring provides increased reliability and maintainability, while also simplifying cabling. Each Wire Center can connect 4, 8, or 12 nodes and provides local expansion connections for adjacent Wire Centers. Through a quality sealed relay, each node is connected to the ring and operates only when the host interface gives the "join ring" command. Thus a power failure at one node will not bring the ring down, and nodes will not join the ring which have found themselves faulty through the digital or analog loopback test. Similarly, nodes will immediately disconnect from the network in the event of a cable break. The Wire Center requires no power except for the small current from the host interface necessary to activate the relays.

Wire Centers allow easy addition or deletion of nodes without disrupting network operation. The Wire Center also provides automatic bypassing of malfunctioning nodes and facilitates fault isolation -- all adding up to low cost of ownership.

The ProNET-80 employs the same shielded twisted pair cable as the 10 Mbps system, and provides distances of 100 meters between Wire Centers and/or hosts. Fiber-optic links extend this distance to two kilometers.

Repeaters allow longer distances between nodes by amplifying the ring signal. When used in adjacent Wire Centers, repeaters allow the star-shaped ring to partition into smaller rings upon a cable break.

Fiber optic links are available for those users requiring security and immunity to RFI (Radio Frequency Interference) and EMI (Electro-Magnetic Interference). Each fiber optic-unit contains a microprocessor to control the system such that when the ring signal fails a node will disconnect from the network. If fiber is used between Wire Centers, the ring will divide into two functional subrings should the signal fail. There is no need for exotic couplers or mixers.

6.2. ProNET-80 Architecture

The ProNET-80 local area network employs many of the components that are used in Proteon's ProNET-10. Wire Centers, cables, and Host Specific Boards (HSBs) are compatible. The difference is in the controller (CTL) boards. ProNET-80 supports UNIBUS, Q-bus, Multibus, VME bus, and Universal bus.

6.2.1. Host Interface Systems

The ProNET-80 CTL board manages all the network transmission functions for ProNET including serial-to-parallel conversion, repeater, address recognition, parity check, token management, modem, bit stuffing, and error timeouts. The ProNET-80 CTL is complete by compatible with the ProNET-10 HSB (Host Specific Board). Therefore, an 80 Mbps network may be built using the 10 Mbps Host Specific Boards.

The CTL contains two 4092 byte receive FIFO packet buffers. The FIFO buffer ensures that simultaneous back-to-back packets will be received without loss of a packet. Additionally, the CTL contains a 4092 byte transmit buffer.

The HSB performs all communications processing work for the host including sequencing, buffering, and controlling, thus providing for a low software overhead burden in the host. The HSB contains packet buffers and direct memory access (DMA) logic, and it maintains two 2046 byte (1023 word) buffers, one for transmit and one for receive. Full-duplex DMA circuits allow independent DMA transfers between the host's memory and the two buffers.

In total, the ProNET-80 system contains three receive and two transmit buffers.

The actual token circulation process is identical to that of the ProNET 10 Mbps network. When commanded to send by the HSB card, the CTL card waits for the first available token and then serializes the message onto the ring, inserting the source address automatically. Should the packet destination address of a circulating message match the node address (or the broadcast address), the CTL notifies the HSB; if the HSB enabled COPY, the CTL routes the message to the packet buffers on the HSB. If the HSB has not enabled COPY, the CTL repeats the message without clearing the refused bit, i.e., receiver did not accept the message.

6.2.2. Error Detection And Addressing

As a function of the 4 into 6 bit coding technique, the ProNET-80 provides a highly reliable error detection mechanism providing fault isolation between adjacent nodes. Should a fault occur between the transmitting and receiving node, the first intermediate node to recognize the fault will correct the parity error. Upon acceptance of the packet, the receiving node obtains a format error, meaning there was a parity error but it was corrected. When the packet returns to the sender, the transmitting node receives a bad format message allowing the host to take appropriate action, such as retransmission.

The 80 Mbps ProNET has an addressing system tailored for the distributed processing environment. Each node in the ring has an eight bit address, set by switches in hardware. The eight bit address, set by switches in hardware. The eight bit address space is divided into two four-bit spaces. Thus, sixteen groups can be selected using four bits, and within each group fifteen users are supported by the next four bits. One of the groups has only fourteen users. This allows for a unique broadcast address to all users on the network. Thus messages can be addressed to specific groups or users within a group.

Additionally, each group and each member of a group receives a message acknowledgement providing the capability to retransmit to those members of a group which did not receive the original message. Those members that did receive the original message will not receive the retransmission.

6.2.3. Coding

The 10 Mbps ring uses the well known Manchester encoding scheme, useful for its regular transitions and perfect zero dc power. At 10 Mbps, the bandwidth expansion required by the biphase technique is bearable; at 80 Mbps, copying this scheme would lead to unnecessarily fast and less reliable circuitry.

The code chosen avoids half of that bandwidth expansion but is still strictly balanced. The 4-into-6 code used by Proteon is balanced over every six bits.

The hardware implementation of this code is relatively simple compared, for example, to the implementation of a 16-bit CRC, and yet the 4-into-6 code has fairly good error detecting properties. For example, with Gaussian error rates of 1 x 10 to the minus 9th, a two node network filled at all times with maximum length packets would make an error once every 176 years. The checking of parity over every six bits was so much simpler than inserting a CRC that some additional time between errors was sacrificed. More important is the performance of the 4-into-6 error conditions. These occur during reconfiguration in this system and with collisions in a CSMA/CD system. In that case, the 4 into 6 error detection system is an order of magnitude better than a 32 bit CRC.

The code with Proteon's format produces a maximum of four bits at a level, and for a maximum data rate of 80 Mbps, the code bit rate must be 120 Mbps. Data transitions are thus separated by a maximum of 33.3 nanoseconds, still small enough to allow reliable extraction of clock from the data stream.

6.2.4. Application

Figure II-23 shows a typical configuration of a ProNET 80 Mbps star-shaped ring network. Potential applications for the ProNET-80 include host-to-host communication, image processing, and high-speed backbone for connecting smaller subnetworks.

The ProNET-80 operates on the same shielded twisted pair as the ProNET-10. Additionally, fiber optic cable may be used.

The 80 Mbps ProNET host interface system derives all of the inherent ProNET local area network benefits, such as high performance and reliability, easy fault isolation and maintenance, and expandability for future growth.

Another improvement over the 10 Mbps network is that the maximum packet length has been doubled, to nearly four Kbytes. Although none of the presently available host interfaces take advantage of this feature, it will leave the system upwards compatible for the future. For the graphics workstation, sending packets twice as long usually means interrupting the host half as often, and so decreasing the host overhead when sending the large files associated with images. The graphics user ideally would like an even larger buffer, say 32 Kbytes, but these sizes will not be practical until memory costs have decreased substantially; again, when that occurs, this variation of the 80 Mbit interface will be available.

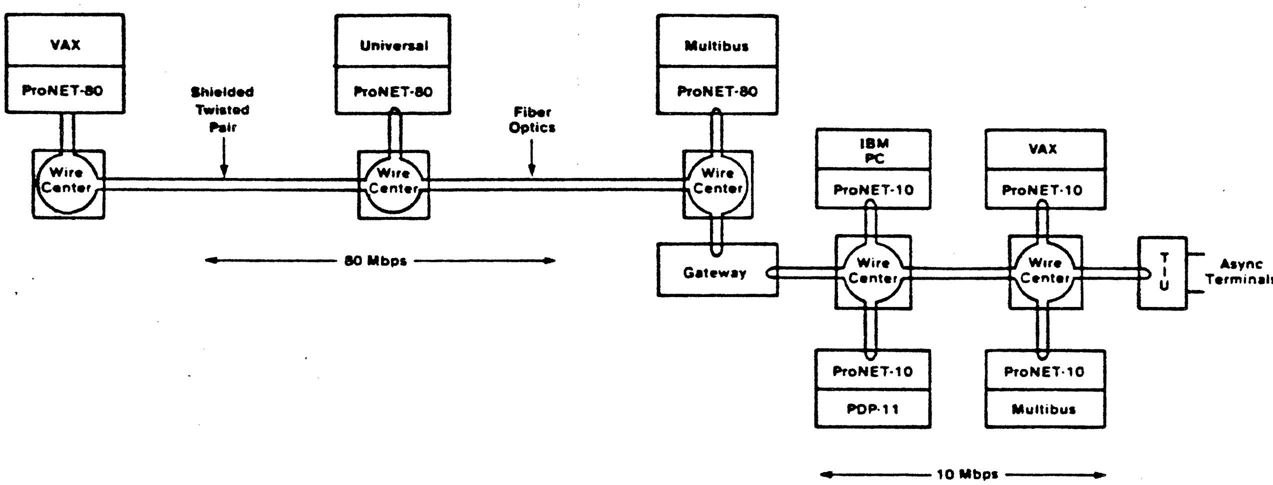

Figure II-23: Typical ProNET Network

6.3. Software

There are a variety of software packages available to drive the ProNET host interfaces, depending upon the user's application. Broadly, they may be classified as HDLC (Higher Level Data Link Control) software, TCP/IP (Transmission Control Protocol/Internet Protocol) software, or DECNet software.

The HDLC package provides an error-free data stream between two cooperating processes on separate machines: that is, a task on one machine may communicate with a task on another machine. The following DEC operating systems are supported: RT-11 V4.0, RSX-11M V3.2 and V4.0, RSX-11M Plus V2.0, and VAX/VMS V3.0.

TCP/IP software provides file transfer, electronic mail, and remote login functions for several host operating systems, including UNIX V6 and 4.1 BSD VAX UNIX and V7 VAX/VMS. Additionally, UNIX 2.9 BSD and 4.2 BSD are distributed with full ProNET support for TCP/IP protocols. As a part of the TCP/IP protocols, TELNET software allows a terminal user to "log on" to another machine, eliminating the need for complex and expensive RS-232 port switching devices.

The current release of Berkeley UNIX includes a complete implementation of the TCP/IP protocols, which are integrated into the socket communications facilities in the UNIX kernel. Device drivers are contained within 4.2 BSD for the ProNET local area network.

The SUN-2 workstation from Sun Microsystems includes the full 4.2BSD networking system, plus some remote file system enhancements. Proteon offers a device driver for the ProNET P1280 Multibus interface, which allows customers the ability to operate ProNET hardware. Installation requires only a simple system reconfiguration.

The DECNet software consists of device drivers which allow the ProNET local area network to be used with Digital Equipment's DECnet software on RSX-11 and VAX/VMS operating systems. Thus the full functionality of Phase II and Phase IV (VMS only) DECnet is available within the ProNET network.

6.4. Gateways

The Proteon p4200 ProNET-Gateways are a family of gateways designed to support high throughput interconnection of two or more networks of the same or different types. ProNET-Linkways accommodate a variety of communication protocols and provide reliable local and wide area networking in multi-vendor computing environments.

The networks supported by ProNET-Gateways include ProNET-10, ProNET-4, ProNET-80, Ethernet, ARPANET networks, and point-to-point telephone lines. Network protocols supported include IP (Internet Protocol), XNS (Xerox Network System), and DEC's DNA (Digital Network Architecture).

ProNET-Gateways provide the following major benefits: 1) they accommodate multiple communication protocols at the network level and thus provide a wide range of network connections while reducing the requirement for additional gateway hardware; 2) they transfer information between networks in excess of 1000 packets per second, supporting high traffic conditions; and 3) they provide connectivity between multiple networks, making ProNET-Gateways very cost effective.

A ProNET-Gateway is based upon a Multibus 68000 processor board with RAM, ROM, timers, and serial lines. This is supplemented with appropriate network interface hardware for connection to the desired local or wide area networks. For example, by configuring the gateway with ProNET-10 and ProNET-80 interfaces, a high-performance connection is established between the ProNET 10 and 80 Mbps networks.

ProNET-Gateway software is highly modular and provides complete packet-switching functionality. The software supports the following:

- device drivers for various network interface hardware
- data link processing for various network-dependent and network-independent data link functions
- forwarders for the supported network-level protocols
- other network level functions required by various network protocols for routing, error reporting, and maintenance

6.5. Pricing

Model	Description	Price
P1080	Unibus System (HSBU/CTL-80)	$ 8,000
P1180	Q-bus System (HSBQ/CTL-80)	$ 8,000
P1280	Multibus System (HSBM/CTL-80)	$ 8,000
P1480	Universal Bus Interface System (UBI/CTL-80)	$ 8,000
P1580	VME Bus System (HSBV/CTL-80)	$ 8,000
P2412	8 Node Wire Center	$ 630
P2413	12 Node Wire Center	$ 925
P3080	Fiber Optic Link, Wire Center-	$ 4,500

	to-Wire Center	
P4200	Gateway Series	
	ProNET-80-to-ProNET-10 Gateway	$14,345
	ProNET-80-to-Ethernet Gateway	$13,345
	ProNET-80-to-ProNET-80 Gateway	$19,195
	ProNET-80-to-9.6 Kbps Wide Area Network	$11,195

6.6. Evaluation

Proteon has been successfully selling its ProNET-10 token ring network since June of 1981. While the spotlight has been focused on IBM's grand token ring scheme, Proteon's base of 10,000 installed nodes illustrates that commercial token ring systems are practical.

The ProNET-80 network will allow Proteon to continue to do well even after the introduction of IBM's 4 Mbps network. ProNET-80 will offer an order of magnitude performance improvement and will provide connections to the IBM token ring. Its current connections to existing ProNET-10s and Ethernets allow it also to be used as a backbone network in appropriate environments.

7. Racal-Milgo Planet

Racal-Milgo Electronics Group is an international company with over 60 firms around the globe. It has been in the communications business for more than 25 years. Planet (Private Local Area NETwork) is the first local area network product in Racal-Milgo's line of data communication products.

7.1. System Concept

Planet is based on an empty-slot, fixed-packet protocol operating on a twin ring. The protocol also incorporates an anti-hogging scheme in which the packet may be used only once and then passed on, regardless of the recipient's response.

The system is heterogeneous, in that all terminals or equipment connected to the network need to have an RS-232 interface. These systems may be 3270s, personal computers, major CPUs, or slow-speed terminals. It should be noted that although all systems can be connected to the network, this does not mean that all can communicate with one another. For example, a 3270 terminal could not communicate with a teletype on the network without some protocol conversion through a third station. The network controller determines whether or not a connection between points is allowable by checking terminal profiles: these profiles must match in terms of baud rate, protocol, and other characteristics before the connection will be made.

7.2. System Components

The system is built on three major components: a Cable Access Point (CAP-PL1000), the Terminal Access Point (TAP-PL2000), and the Director (PL3000). The CAP connects directly to the coaxial cable and is a passive component; the TAP contains the intelligence of the connection, handling full/empty bit setting; the Director is the network controller. One CAP supports one TAP which supports two RS-232 connections. The Director may support up to 250 TAPs on one ring. Figure II-24 illustrates an example Planet network.

7.3. The CAP

The CAP is an outlet-type device which can be pre-installed when the building is wired, or installed as necessary. It acts as the interface between the ring and the TAP and provides user-access to the transmission media. This device is passive, routing the cable signal either to a TAP or, if no TAP is connected, on through to the ring. Four BNC coaxial connectors are located on the back of the box: in and out for the outgoing direction, and out and in for the incoming direction. Each CAP

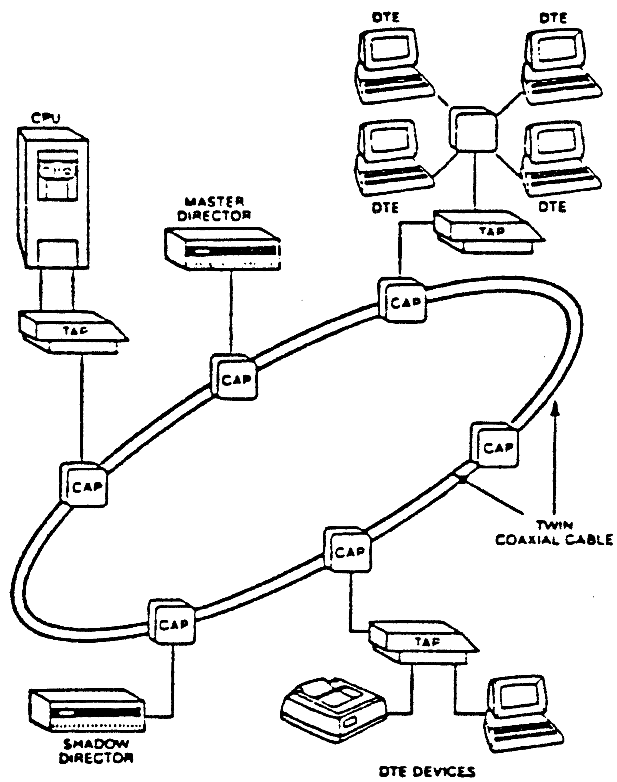

Figure II-24: Example Planet Network

can accommodate one TAP.

The twin coaxial cable is used to provide fail-safe system operation in the event of cable break or equipment malfunction. The twin cable ring enables signal loopback, which supports the fail-safe feature.

7.4. The TAP

The TAP is analogous to a modem in analog systems: it acts as the interface between the Data Terminal Equipment (DTE) and the CAP. Each TAP can support two RS-232-C DTE or DCE devices. A maximum of 250 active TAPs may be supported on one ring within a maximum 12-mile circumference. The TAP is responsible for signal regeneration and contains the logic circuitry to enable communication of the attached DTE with the Director. The address of each TAP is determined by the dipswitch on the bottom of each unit. The longest distance allowable between TAPs is 650 feet for RG62 cable, 950 feet for RG59 cable, and 1,350 feet for RG11 cable. If longer distances are needed, a TAP may be used as a regenerative repeater. The various cables may be mixed within a single ring.

7.5. The Director

The Director is the "brain" of the network. It controls all connect, disconnect, and transfer functions and provides management information, system statistics, and error logging. The Director also generates and monitors the data packets, determining the total number of simultaneous packets placed on the ring; keeps track of network configurations, or Switch Plans; and maintains DTE operating parameters and addresses. User priority assignment and security password protection are also maintained through the Director. Any combination of independent circuit switch plans is permitted. The system control may be centralized by the Director in one ring or distributed by communicating rings with separate Directors.

To determine whether or not a connection is permissible, the Director checks the caller port profile against that of the destinations: the port profile, which may be one of 30 defined by the user, contains the operational characteristics of a particular port; if the profiles do not match, the stations cannot communicate.

For reliability, it is suggested that a redundant Director be purchased for the network and located away from the first Director. Its operation on the ring will appear as a shadow Director, monitoring the primary Director's functions; the updating of system data tables and records of the system configuration is completed by constant communication between the master and the shadow. When the shadow Director detects a failure, it starts an internal timer. If this timer times out, the shadow assumes that the master has failed to reconfigure the ring and takes over the master's responsibilities. The shadow notifies the operator at its terminal that the master has failed. The only connections lost during this operation are those that are being set up by the master Director at the time of its failure; all others are maintained.

A shadow Director may be added to the system at any time. As soon as the shadow Director is powered up, the master sends a copy of the system data tables to the shadow's memory and cassette.

For non-volatile storage, the Director uses a 6 Kbyte CMOS RAM cassette which plugs into a slot in the Director front panel. This cassette is usually used to store configuration information and can be written to or read from the Director software. Once the cassette is removed from the Director, it can be stored at room temperature for at least two years without loss of data.

7.6. The Packet

When a newly created packet leaves the Director, the full/empty bit has been reset, indicating packet availability. The packet travels along the network to be examined by each TAP. When an operator initiates a transmission, the first available packet passing the associated TAP is loaded with the operator's message, and the destination TAP's address. The originating TAP also sets the full/empty bit. The packet traverses the ring, with each TAP inspecting the destination address. When the packet reaches its destination, the TAP recognizes the address and accepts the message. This TAP also sets the acknowledge bit to indicate that it was received. When the packet passes through the Director, the monitor bit is set to indicate that the packet in use has passed through the Director once. When the packet once again returns to the originating TAP, the TAP recognizes the acknowledge bit and resets the full/empty and monitor bits.

If an occupied packet tries to pass through the Director a second time, the Director clears the packet and resets the control and full/empty bits. The Director keeps track of all packets not receiving acknowledgement by the destination and passing through the Director twice. The Director then clears the full/empty bit, which in effect destroys the packet.

7.7. System Configurations

The high-speed data packets generated by the Director allow virtual circuits to be established between any of the DTEs connected to the ring. A virtual circuit (a path which appears between two or more devices for the purpose of communication) appears to the user as a direct connection. The types of virtual circuits supported by Planet are:

1. Permanent virtual circuits, both point-to-point and multipoint

2. Switched-plan virtual circuits

3. Switched virtual circuits

Each ring can support up to 250 virtual circuit combinations.

A permanent virtual circuit can be set up between any of the ports on the ring. Once set up during initialization, these connections are long term and are preserved over power failures or Director shutdown. A direct permanent virtual circuit is a point-to-point connection between two network devices. A multicast virtual circuit consists of a master port which broadcasts to a number of slave ports, with each slave transmitting only to the master. One or more switch plans may be implemented simultaneously through this type of connection.

A switched-plan virtual circuit which is defined at system initialization consists of a group of connections between two or more ports. Up to four switched-plan virtual circuits may be implemented at the same time. The plans themselves are preserved over power failures or Director shutdown, but the connections must be initiated at the authorized terminal. A switched plan may be a single multicast plan, or it may consist of a number of direct connections. If several types of plans are combined, a PLEX plan is formed. If a caller changes a switched plan already in use, the previous switched plan is disconnected before the new connection is made.

A switched virtual circuit is a short-term connection between two devices, set up by the user through an asynchronous ASCII terminal. Since this type of connection is not stored on non-volatile memory, it will be cleared on occurrence of power failure or Director shutdown. To implement a switched virtual circuit, the originating operator, or caller, requests the connection by entering the required connect commands directly at his/her terminal. A third-party, or an operator not directly involved in the communication, may also request a connection by entering the required commands; a third-party connect is used where unattended devices need to be connected, e.g., a CPU to a printer. After a call is established, it may be transferred to another station by either party.

A conference circuit may also be implemented, allowing a number of users to be connected as a group, with each transmission being received by all other users in the conference group. Any device communicating with the Director may set up a conference, which is initiated on request. The conference is protected by assigning a user or system password so that connection to the conference would require both a CONNECT COMMAND and a PASSWORD. Any authorized terminal can expand or reduce the number of conference connections, provided that the number of connections and passwords is known. The system allows up to six simultaneous conferences.

7.8. Reliability

Planet incorporates a "fail-safe" feature to ensure network operation in the event of a cable break and/or failure of a TAP or a CAP. If a problem of this type occurs, the cable break or failure is detected by the Director and isolated to the TAPs located on either side of the system malfunction. The TAPs are looped back, maintaining system operation. It should be noted that when the system is functioning normally, it in and of itself is looped back on both sides. Ring reconfiguration is shown in Figure II-25.

This feature can be illustrated by unplugging a transmitting TAP. Immediately the Director will display the error message on a CRT (which is acting as a diagnostic printer) and initiat the loopback with no noticeable delay. When the TAP is plugged back in, the system returns to the original configuration, printing another diagnostic message on the CRT. All this can be done in real time with no problem.

7.9. System Features

The system features include queuing, message service, and systems statistics for network management.

If a destination is busy, the call request is automatically queued. The Director records the call request and completes the connection as soon as the destination is available. Each sending port may only have one request queued at any given time. A destination may have up to twenty queued call requests. One potential hindrance in some applications is that when a call is queued, the caller does not know what his/her position is in the first-in/first-out (FIFO) queue and cannot estimate the delay. A caller can remove himself/herself from the queue with a disconnect command, or define a maximum time to remain in queue to a specific device.

The message service supports system and broadcast messages. A system message may contain up to 60 characters of text, which is output to a terminal after the operator has entered the data recall character (CTRL-P). The message is inserted after the sign-on sequence and will be preserved over power failures. A broadcast message, which also contains up to 60 characters, is sent to all asynchronous ASCII ports at a specified time and date. These messages are removed from the Director's memory following the lapse of the specified date and time and completion

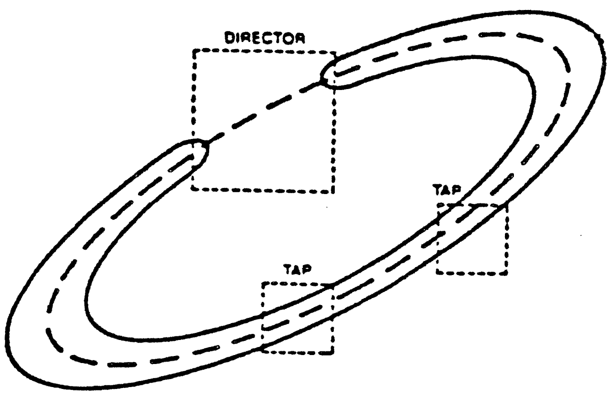

A) Normal Configuration

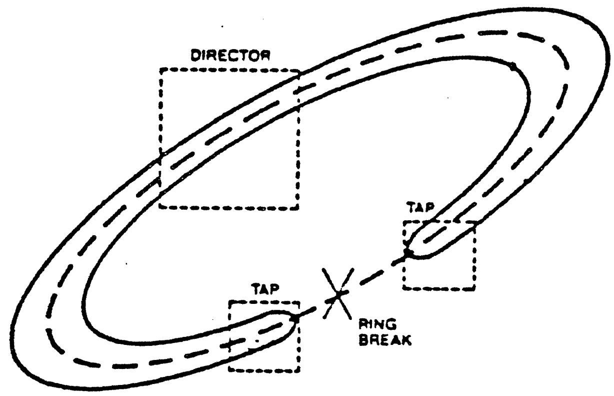

B) Reconfigured After Failure

Figure II-25: Reliability Reconfiguration of the broadcast. All uncompleted broadcast messages are preserved over a power failure.

System statistics are printed out at the management terminal to indicate the use and loading of the ring. The two categories of system statistics are call statistics and network statistics. The information provided by call statistics includes number of attempted connections, number of successful connections, number of unsuccessful connections, and total connection time in hours. The network statistics provide information about the number of ring breaks detected, including all ring breaks of the system to date, and the number of parity errors detected, including all those detected during system operation. Once the system has been configured, the Director is responsible for network management, which enables the Director to notify the operator regarding a TAP failure in response to a Director poll or message, a removal of a TAP, insertion of a TAP, ring breaks, parity errors, or failure of the "master" Director or the "shadow" Director.

7.10. Port Features

Port features are assigned to individual ports at system initialization and may include a port name, a contention group name, a password, or a closed user group.

Each port in the system has a default name of up to four digits which is determined by its TAP and its port off the TAP. For example, in the case of TAP 201, port 1 has the default name of 2011, whereas port 2 at the same TAP is 2012. A descriptive name can be assigned to each port. The assignment of the name "PRINTER" to TAP 201, port 1, enables connection to the port by using the name "PRINTER" instead of 2011. The assigned name may contain up to eight alphanumeric characters.

A port can be assigned to one of 14 contention groups, each having an assigned name. The name of the group can then be specified when contending for a service. For example, a request for connection to the group named BASIC will enable a connection to any free port in that group, or the request will queue the call if no port is available. Group names can also contain up to eight alphanumeric characters.

A password restricts access to certain ports in the system, requiring the caller to give both the port address and the password. Each password may also have up to eight alphanumeric characters.

A closed user group is a group of ports assigned so that only members of that group have access to the port. Members of the closed user group have access only to those ports assigned to that particular group. Each port may have one of 14 user groups. This feature may be useful in areas where security is a concern.

7.11. High Density Planet

The High Density Planet (HDP), is a very compact unit containing 1 Dual External Ring Interface Card and up to 16 Terminal Access Point (TAP) cards. The unit is rack mountable and designed for use in Planet networks having a large amount of DTE and DCE ports concentrated in one area.

HDP uses the latest LSI microprocessor technology to achieve advanced LAN features such as port selection, queuing, resource sharing, closed user groups, password protection and priority access. There are also a broad range of diagnostic aids and network control features including automatic network restoration in the event of a cable failure, power failure or equipment malfunction.

7.11.1. System Components

The basic system consists of a PL4000 nest or chassis, 1 to 16 - PL2001 Terminal Access Point Cards, and 1 - PL1001 External Interface Card, with system control being administered by a Racal-Milgo PL3000 Planet Director.

The compact unit measures 12" high x 19" wide x 19" deep and may be mounted in a standard 19" rack or set on a desk or table top. The 32 port nest includes 2 power supplies, one of which is used for backup in the event of a power failure.

The PL2001 TAP Card provides the interface between the Planet network and any standard EIA RS-232 device, thus allowing terminals to communicate over the high speed digital ring. Once access to the network has been obtained, the TAP card will act as a packet assembler/disassembler passing data to and from the network. Each TAP Card enables two RS-232-C DCE and DTE connections to operate in a synchronous, asynchronous, or mixed mode at speeds up to 19.2 Kbps.

The PL1001 External Interface Card enables flexible configuration of the HDP for a variety of communication needs. Simple pushbuttons permit interfacing up to 16 TAP Cards to the ring via a single CAP (Cable Access Point), interfacing the HDP nest to a Director; interfacing a number of serially connected or daisy chained nests with individual "nest by-pass" in the event of a power failure. Note that a nest power failure cannot bring down the entire network.

7.11.2. Applications

Some High Density Planet applications include a Hunt Group feature; Front-End Port Switching; Matrix Switching; and Multiplexer Switching. When using the Hunt Group feature, a port can be assigned to one of fourteen hunt contention groups. If a connection is requested to a busy port and that port is in a hunt group, the system will search for a free port in the group. If none are free, the call may be queued.

Front-End Port Switching allows the central site TAPs to be configured through multiple switch plans so as to switch groups of circuit connections from one processor to a spare processor, thus allowing the individual ports to be switched for back-up or maintenance.

The HDP becomes a powerful data switch, when used in a Matrix Switching application, allowing interconnection of any of the ports on the ring network. For example when the HDP nest is used in conjunction with Planets' closed user group feature, individual user groups become interconnected "data switching modules."

When combined with the Racal-Milgo Omnimux the nest becomes a central switch for remote multiplexers. And when used with modems, HDPs flexibility extends beyond the local environment allowing resource sharing of expensive distance facilities to become a reality.

7.11.3. Enhancements

In addition to High Density Planet Racal-Milgo has a Starter Pack identified as the PL15-SP. Effectively reconfigured and competitively priced for the fledgling local area network customer, the PL15-SP contains 1-Director, 5-TAPs, 6-CAPs and 6 cable segments plus full documentation.

7.12. Pricing

Planet

PL15-SP Starter Pack	$9,995
PL1000 Cable Access Point (CAP)	$95
PL2000 Terminal Access Point (TAP)	$1,645
PL3000 Director	$5,975

High Density Planet

PL10001 Dual External Ring Interface Card (DERIC)	$375
PL4000 High Density Nest	$2,750
PL2001 TAP Card	$1,405

Quantity and OEM discounts available. Leasing is available in 12, 24, 36, and 60 month terms. Installment purchases are also available for 12, 24, and 36 months.

7.13. Evaluation

Racal-Milgo installed the first Planet system in the United States in April of 1982; over 500 systems have been installed worldwide since late 1981.

As a communications product, Planet will do what it claims to do - transmit protocol-independent data over a circuit at speeds up to 19.2 Kbps. It provides the connection over which DTEs may communicate, but it supports no protocol conversion. The reliability of the system is a major advantage: cable breaks, CAP failures, the TAP failures are all instantly recoverable by the Director. The only central point of failure is the Director which, if duplicated by the shadow, would not bring down the system.

8. Ungermann-Bass Net/One Token Ring

Ungermann-Bass continues to be one of the major forces in the local network area. It has concentrated on offering general purpose LAN systems capable of interconnecting information processing devices from different manufacturers.

Ungermann-Bass is a publicly-held company headquartered in Santa Clara, California. It has research facilities at its Santa Clara headquarters and in Boca Raton, Florida, and Bohemia, New York; the latter two facilities were acquired when U-B purchased Amdax in December 1982.

The Ungermann-Bass Net/One is a flexible local network system for heterogeneous device interconnection. The first version of Net/One -- a 4 Mbps baseband system -- was delivered in July 1980. An Ethernet-compatible version (10 Mbps baseband) was delivered one year later. In July 1982, the first CATV-compatible broadband version of Net/One, which operates at 5 Mbps per channel, was installed. In April 1983, as a result of the Amdax acquisition, the Series 600 Products were added to the broadband product offerings.

In October 1985, Ungermann-Bass announced the expansion of its Net/One product line to include a token-ring implementation. Called Net/One Token Ring, the IEEE 802.5-compatible system includes a group of Net/One products that allow users to configure a 4 Mbps multiple-ring system.

The 4 Mbps token-ring system includes server support for hosts, terminals, printers, and modems; workstation support for devices based on the Intel MULTIBUS and the IBM PC bus; distributed wiring concentrator (DWC) support; and network management in a multiple-ring environment in which connectivity is achieved via Net/One Ring-to-Ring Bridges. The system's wiring concentrators provide configuration flexibility and support for alternative wiring strategies. Network bridges allow nodes located in different workgroup rings to communicate with each other. U-B's Net/One Token Ring offers the industry's first bridge capable of supporting the IEEE 802.5 specification.

In order to bring a full token-ring implementation to the marketplace, Ungermann-Bass developed its own custom VLSI semiconductor chip to handle the token-passing functions defined by the 802.5 specification. This chip, in conjunction with an Intel 80186 microprocessor, provides the silicon environment required to support the features and performance of the Net/One Token-Ring LAN. Because it is using its own developed chip, U-B was the first to ship (in January 1986) an 802.5 token-ring product. Also, unlike its previous action in allowing Fujitsu to later sell the CSMA/CD chip Fujitsu had developed for the U-B 10 Mbps Net/One, U-B may not let anyone else use this one.

8.1. Intro/Net System

Ungermann-Bass introduced a token-ring starter system, called Intro/Net, to allow users to become familiar with the new technology in a systems environment. It also is being used to highlight U-B's technical expertise and product line breadth.

Designed to complement the IBM Token-Ring product offering, Token Ring Intro/Net demonstrates the capability to network non-IBM equipment and IBM products for which an IBM networking solution is not available, and it also provides value-added networking alternatives to existing IBM solutions. The system provides connectivity support for asynchronous devices, the IBM 3270 family of products, and devices based on the IBM PC Bus, as well as host support from a multitude of vendors including IBM and DEC.

Net/One hardware and software included in Token Ring Intro/Net enables the IBM PC family to not only share disk and printer resources, but also to simultaneously maintain high-speed connections to asynchronous hosts and IBM mainframes.

The Token Ring Intro/Net is available in several forms. Option I includes two Distributed Wiring Concentrators (DWCs), one Ring Bridge, two NIU-180s, three PC-NIUs, one NIU-74, and one NIU-78. An example system using this option is shown in Figure II-26. Option II includes all of Option I plus an IBM PC XT for Network Management. Option III includes one Distributed Wiring Concentrator, one NIU-180, one NIU-74, and six PC-NIUs.

Also included with each option are cables and a network software package which features Net/One operating software, virtual circuit service, Personal Connection software (including NETBIOS compatibility), network management software, and cables.

8.2. Intro/Net Hardware

8.2.1. Network Interface Unit, Model 180 (NIU-180)

The NIU-180 family is expanded to incorporate Token Ring capabilities. The NIU-180, as well as the NIU-74 and NIU-78, are the same as the existing U-B products of the same designation, except they contain a token ring Personal Network Interface Unit board instead of an Ethernet one. Up to eight vendor-independent asynchronous devices with RS-232 interfaces can physically and logically connect to an Ungermann-Bass Token Ring LAN through the NIU-180.

8.2.2. Personal Network Interface Units

The Personal Network Interface Unit (Personal NIU) provides

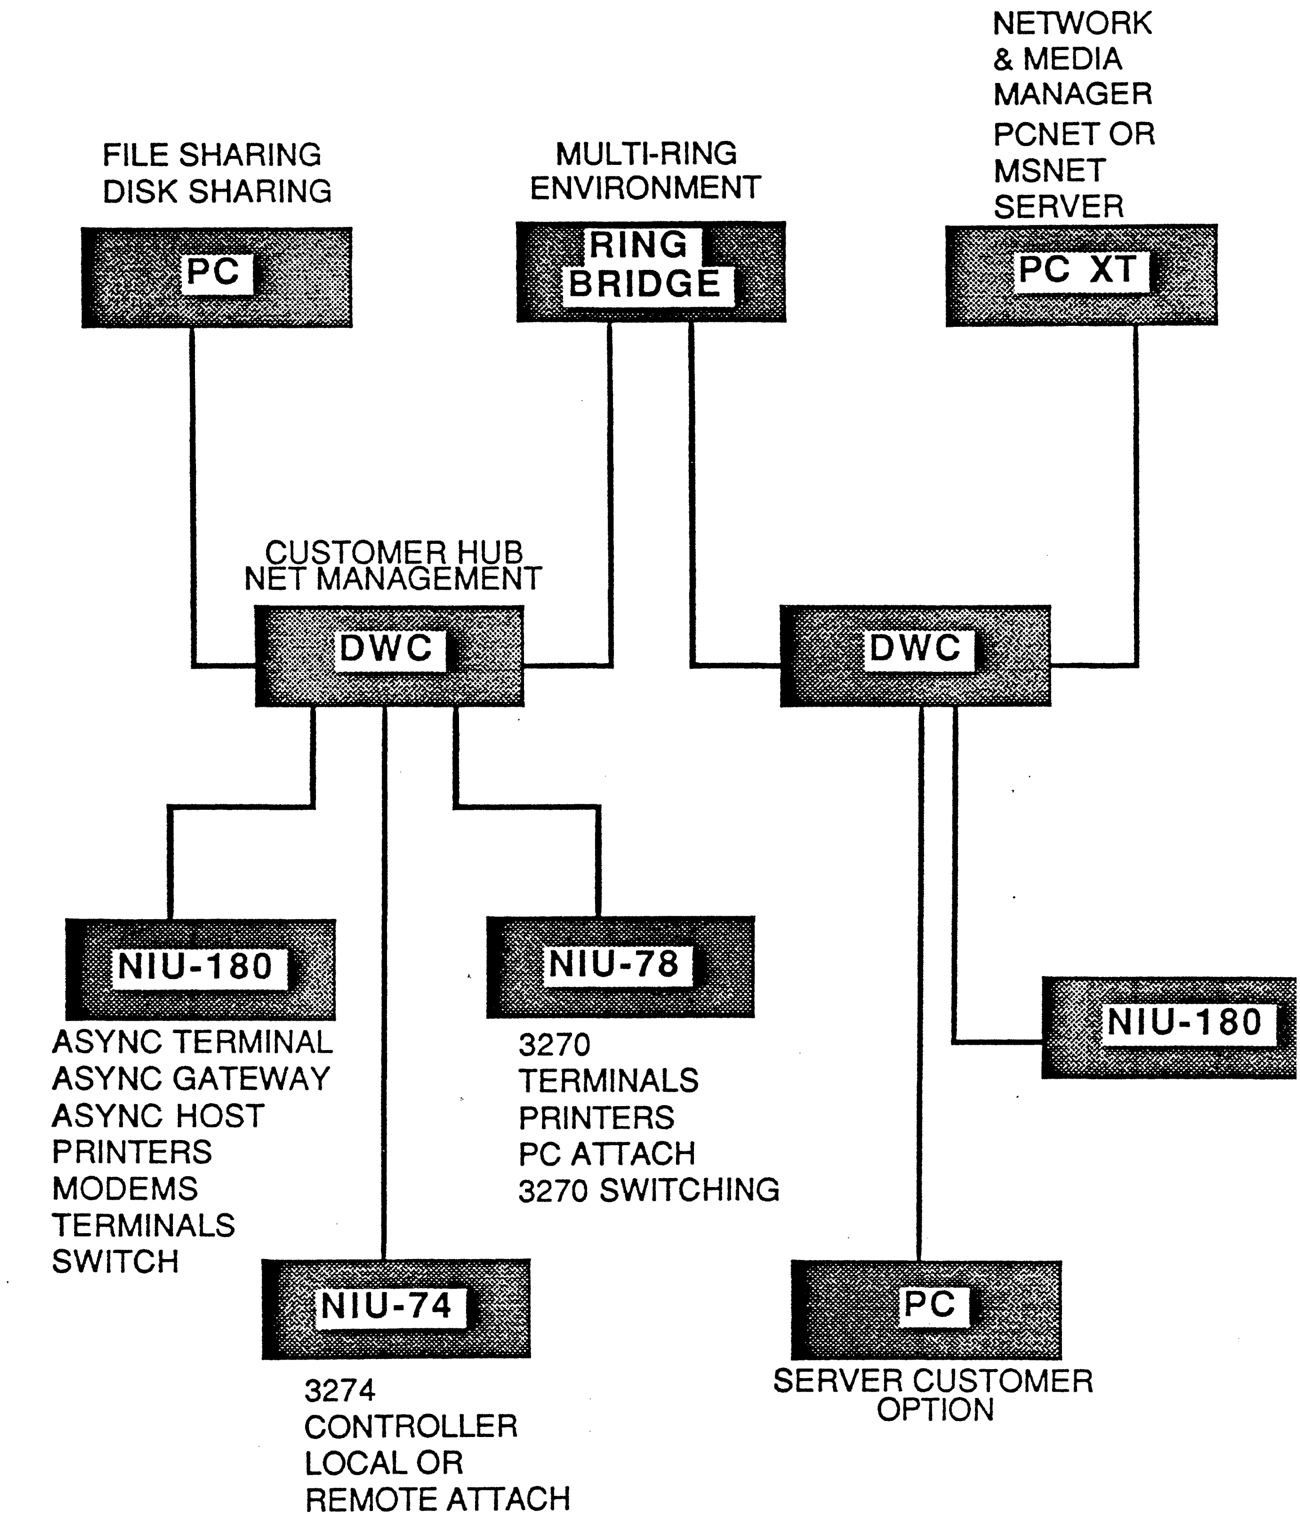

Figure II-26: Example Token Ring Intro/Net the board-level connection for the IBM family of PCs onto the Net/One Token Ring network. The Personal NIU enables the PC to share peripherals attached to other PCs and NIU-180s, as well as establish connections to other hosts. Software with the Personal NIU provides the PC users with file sharing, print sharing, 3270 terminal emulation, and PC-to-mainframe file transfer. Distributed applications written to both the NETBIOS and the Microsoft Networks (MS-Net) Session Interface will operate on the Ungermann-Bass Token Ring Network through the Personal NIU.

The Personal NIU Token Ring Adapter is an Intel 80186 microprocessor-based design that includes the following major items: an 8 MHz 80186 Microprocessor; 16 Kbytes of EPROM expandable to 32 Kbytes; 256 Kbytes of Local/Shared RAM, with parity, between the 80186, host, and the Token Ring Handler Gate Array; a Storage Arbitration and Control Unit; a PC Command Channel Interface; and the Token Ring Handler Gate Array (TRH), U-B's proprietary chip. It also has an expansion interface and a front end. The Token Ring Adapter hardware is mounted on a card measuring 4.15" in height and 13.115" in length and is constructed such that it can be attached to one of the host feature slots.

The Intel 80186 microprocessor in the Personal NIU performs the following functions: bootstrap, where it initializes all board hardware after reset and performs power-up diagnostics; control of the Token Ring gate array; providing data formatting and manipulation functions on request for the gate array; and control of peripheral hardware attached to the expansion connector.

8.2.3. Network Interface Unit 74/78 (NIU-74)(NIU-78)

The NIU-74 and NIU-78 are NIUs designed specifically to connect the IBM 3270 Information System to the Net/One Token Ring Network. These devices permit users of 3270 system terminals and printers to communicate with an IBM 3274 controller via the Net/One LAN. In addition, the NIU-74 provides high-speed access to an IBM mainframe for IBM PCs using the 3270 Personal Connection software program.

8.2.4. Distributed Wiring Concentrator (DWC)

The Distributed Wiring Concentrator enables a combination of up to ten 802.5 Network Interface Units -- including the NIU-180, NIU 74/78, PC-NIU, Ring Bridge, and/or IBM PC Adapter Cards and their associated devices -- to be connected to the Token Ring Network. The DWC functions in both an active and a passive state: in a passive state, the DWC provides the connection hub of the logical ring, physical star topology; with the addition of a power source, the DWC functions in an active state, allowing the Network Management Console (NMC) to perform problem determination, fault isolation, and, when appropriate, to remove stations from the ring. DWCs may be connected in a daisy chain fashion to increase ring size to the allowable maximum of the media selected.

8.2.5. Ring Bridge

The Ring Bridge provides the capability to configure multiple rings. It bypasses more than the 260- and 72-station limitation of the current IBM Token-Ring Network. The Ring Bridge reinforces the Ungermann-Bass modular approach, wherein segments of a network can be isolated to improve problem determination and network security. In addition, the Ring Bridge allows the user to effectively manage load balancing and increase available network bandwidth.

The Token Ring Bridge has the following features: open-ended architecture to accommodate LAN management functions and expansion; flexible architecture to accommodate user configuration changes and growth; and software that can be updated via a token ring server. It also has 16 Kbytes of EPROM expandable to 128 Kbytes, 32 Kbytes or 128 Kbytes of parity protected RAM, and two IEEE 802.5 Token Ring interfaces. The Token Ring Bridge is user-installable, can be located in either an office or a wiring closet, and has a small footprint (9.5" by 15.5").

Additionally, the Token Ring Bridge has architecture expansion capabilities to include a LAN management coprocessor internal to the Bridge that will add processing power for Bridge application programs such as Network Servers. It also can be expanded to include and control a Distributed Wiring Concentrator.

Ungermann-Bass is using the same source routing technique as IBM. In this approach, the route is inserted into the frame at the source or point of origin; this provides the necessary routing information required by bridges. Bridges then build the route through a broadcast process. The destination station receives the frame with route built in; following this, the route can be used for further communication.

8.3. Intro/Net Software

The Token Ring Intro/Net software provides users with the various services that comprise the essence of Net/One functionality, features, and benefits. These services can be grouped into three general categories: Virtual Connection Service, Personal Connection Services, and Net/One Operating Software and Network Management.

8.3.1. Virtual Connection Service

Virtual Circuit Service (VCS) is the most widely used communications service offered on Net/One. It provides reliable point-to-point communication between two devices or user programs attached to a Net/One LAN. It derives its name from the physical circuit-switching systems it emulates, such as the telephone network. However, instead of using switches to set up an unshared physical link between two points for the duration of a session, VCS is implemented by software that routes packets of data along the shared Net/One communications medium.

Net/One Virtual Circuit Service provides terminal device users with the following features: Interactive Command Interpreter; Distributed Name Service; Rotoring on Assigned Names; Port Password; Automatic Baud Rate Detection; Idle Circuit Time-Out; Automatic Device Mismatch Resolution; and Sequencing and Error Recovery.

8.3.2. Personal Connection Services

The Net/One 3270 Personal Connection is a high-performance local connection to the IBM mainframe from a Token Ring Intro/Net-networked IBM PC or compatible. When running Model I of the 3270 Personal Connection in a Net/One Token Ring Intro/Net environment, the PC emulates an IBM 3278 or 3279 terminal. Emulation is transparent to host software. The PC becomes a universal workstation, able to use multiple resources without having to disconnect from the host.

Network file and printer sharing is accomplished by implementing the Microsoft Networks MS-Net Version 1.0 software on Net/One. This software allows Token Ring Intro/Net users to share system resources such as data files, programs, hard disks, and printers within a networked MS-DOS or PC-DOS operating system.

Since MS/Net is an extension of the MS-DOS/PC-DOS operating system, it is transparent to both network users and to existing MS-DOS/PC-DOS application software.

Token Ring Intro/Net users can send programs, reports, memos, and other information to any other user on the network. Users can also spool print files to shared network printers. MS-Net gives all network users the benefits of hard disk storage and expensive printers without requiring these devices at every workstation. Also, the Net/One Personal Connection includes full support for the IBM NETBIOS interface: almost any program written to run on IBM's PC Network can be used without modification on Intro/Net with MS-Net. In addition, the IBM PC Network Program can be run on the Personal Connection networking hardware. Both IBM and Ungermann-Bass network operating systems can be mixed on the same network, and any workstation can access any server.

8.3.3. Net/One Operating Software (OPS) and Network Management

Net/One OPS includes the network operating kernel, data link protocol software, and network management services. The network management services allow a network administrator to configure, monitor, and debug a network, and to communicate conveniently to all network users. Network management services available in Intro/Net include:

Configure - an easy to use program that allows a network administrator to prepare configuration information for the network NIUs.

Download Server - the program that actually downloads configuration files and operational software, such as Virtual Circuit Service, across the network to the NIUs.

Broadcast/Multicast Messages - a Network Administrator can send a message to all users, or multicast the message to users belonging to a pre-defined group.

Media and MAC Management - this program allows the network administrator to overview his/her network configuration, display error data from a ring, display ring condition messages, and, when appropriate, remove or restore a station on a ring.

8.4. Pricing

Prices for the Token Ring Intro/Net starter systems are given below. Ungermann-Bass expects to offer unbundled unit-priced products around June 1, 1986.

Option	Description	Price
I	2 DWCs, 1 Ring Bridge, 2 NIU-180s, 3 PC-NIUs, 1 NIU-74, 1 NIU-78	$34,750
II	Option I, 1 IBM PC XT	$41,940
III	1 DWC, 1 NIU-180, 1 NIU-74, 6 PC-NIUs	$15,000

8.5. Evaluation

The goal for Net/One Token Ring products is to coexist on the cable and also communicate with IBM equipment running U-B's software immediately, to participate in IBM's future network management offerings, and to support the transparent communication of IBM equipment using IBM-supplied software and program products and U-B equipment on the network in the future. Compatibility is difficult to obtain prior to IBM product availability, because there is yet no independent agency to test different vendors' implementations of the 802.5 specification. However, Ungermann-Bass is confident that its compatibility goals are realistic, because of the large investment it has made in token-passing technology: it has already conducted well over 1,000 test cases to validate its media-access definition. In addition, normal- and stress-condition tests have been conducted in a 100-node system using technology supplied by Ungermann-Bass as well as other vendors.

Ungermann-Bass' Token Ring's impact on its 1986 revenues will depend to a large extent on how quickly its customer base embraces the 802.5 standard: its view is that demand for the token-ring technology will grow rapidly throughout 1986, making the technology one of the most important in the industry. The company sees this technology forming a major portion of its 1987 revenues.

Ungermann-Bass' token-ring product offerings are designed to complement the product offerings of IBM. It believes that IBM's primary market thrust will be to use the token-ring product to interconnect IBM word- and data-processing equipment. As a result, customers who choose to install the IBM Token-Ring Network will still need to network non-IBM equipment, network IBM products for which an IBM networking solution is not readily available, or have value-added networking alternatives to the existing IBM solution. Ungermann-Bass feels that these customers will turn to it for help in these areas.

In January 1986, Ungermann-Bass announced it had signed a licensing agreement for worldwide rights under the token-ring patents of Olof Soderblom, thus clearing the way for the company's continued implementation of its own token ring products. The terms of the agreement were not disclosed, though IBM previously paid $5 million to Soderblom for similar rights.

The agreement with Willemijn Holding BV -- a Netherlands-based company that owns the Soderblom patent rights -- secures for Ungermann-Bass the rights to implement token-ring LAN technology. The company said that after carefully reviewing the patent, it felt that its best course of action was to obtain the license and thereby protect itself and its customers from any risk of patent infringement.

9. Ztel PNX

Ztel was incorporated in January, 1981, with financing provided by Adler & Co. and GE. Second and third rounds of financing produced an additional $33.5 million. One of the new investors was NCR, whose share subsequently came to 19% of Ztel. Ztel's Private Network Exchange (PNX) is an integrated voice/data system which combines the data capabilities of a local area network with the full feature voice capabilities of a PBX. Developed by Ztel, voice-only production shipments took place in the third and fourth quarters of 1984.

Ztel is headed now by Murray H. Bolt, president and chief executive officer. Mr. Bolt joined Ztel from Codex, where he was vice-president and GM of the LAN Products Business Unit.

On May 15, 1985, Ztel petitioned the United States Bankruptcy Court for financial reorganization under Chapter 11 of the Federal Statutes.

Ztel had made a voluntary business decision to file Chapter 11; it was used, according to the firm, as a means to successfully recapitalize Ztel so it could accelerate its business objectives of developing, manufacturing, marketing, and supporting a superior communications system.

Since its inception through March 30, 1985, Ztel had invested approximately $49,000,000 in research, development, manufacturing, and marketing. In April 1985 Ztel approached its four primary financial backers -- the General Electric Venture Capital Corporation, NCR Corporation, The Hillman company and its affiliates, and The Adler Group -- to discuss a recapitalization of the company to fund a continued ramp-up of its product manufacturing and marketing operations.

G.E., NCR, and Hillman Co., with the addition of Harvest Ventures, agreed to a funding proposal of approximately $17 million. The reorganization of the company under Chapter 11 was necessary to permit Ztel's financial backers to continue their investment in, and commitment to, the company.

While in Chapter 11 proceedings, the company, under the direction of Murray H. Bolt, completely restructured its entire operation, putting in place a new management team, a new marketing plan, and a targeted, market-driven product focus and direction. The reorganization plan also provided for 100% payement to all Ztel creditors.

Effective August 2, 1985 -- only 11 weeks after filing its original petition -- Ztel emerged from Chapter 11 reorganization with a court-approved recapitalization plan.

Now that these plans have been approved by the court, the company's investors, and creditors, Ztel intends to continue the manufacturing and marketing of its PNX, as well as to fund the development of further product enhancements. NCR, G.E., Hillman Co., and Harvest Ventures have been issued shares of New Class A Preferred Stock; in addition, NCR, G.E., Hillman Co., and The Adler Group have received shares of Common Stock and Class B Redeemable Preferred Stock based on their prior debt investment funding levels.

CEO Bolt has begun building a new management team. Recent appointments include Kevin P. Moersch, vice president of distribution; Joseph W. Milacci, vice-president of manufacturing; and Elder N. (Kip) Ripper, vice-president of marketing. Under the direction of this team, the company has begun to more aggressively market the PNX.

Essential to Ztel's success are the continuing relationships that have been secured with distributors. For example, ASI Telesystems has remained committed to selling and supporting Ztel's PNX from the first site tests to the first commercial product offering. In addition to existing distributor agreements, the company is now negotiating long-term contracts with new distributors to further extend its distribution channel. Present distributor include Allwest Telecommunications, ATS, Chadwick Telephone, Computer Telephone, Manitoba Telephone Systems, Miller's Wireline, Nobell Telecommunications, Northstar Communications, Skaggs Telephone, and Tenant Network Systems.

At the Tele-Communications Association Conference (TCA) September 1985, Ztel announced the signing of two new regional distributor agreements, revealed an unprecedented two year warranty and made public its installed customer sites.

9.1. The Private Network Exchange

Ztel designs, manufactures, markets and supports the Private Network Exchange (PNX), a highly advanced voice/data communications controller for the business environment. Ztel's PNX serves a broad spectrum of market line size requirements. The PNX 5200 serves the small market (50-300) lines. The PNX 5300 serves the medium size market (300-800) lines; and the PNX 5400 serves the large market (800-1500) lines. The PNX simultaneously transmits both voice and data communications over two twisted pair wire. The PNX combines digital transmission technology with distributed processing to negate the possibility of single point failure and provide economic and reliable service.

The PNX software uses the C programming language and follows the OSI/ISO model. The modularity inherent in both the software and the hardware allows for gradual, cost-effective system growth. Current release system configuration designs can grow from less than 512 ports lines to 2048 ports. Several software options, such as Least Cost Networking, Messaging, and On-Line

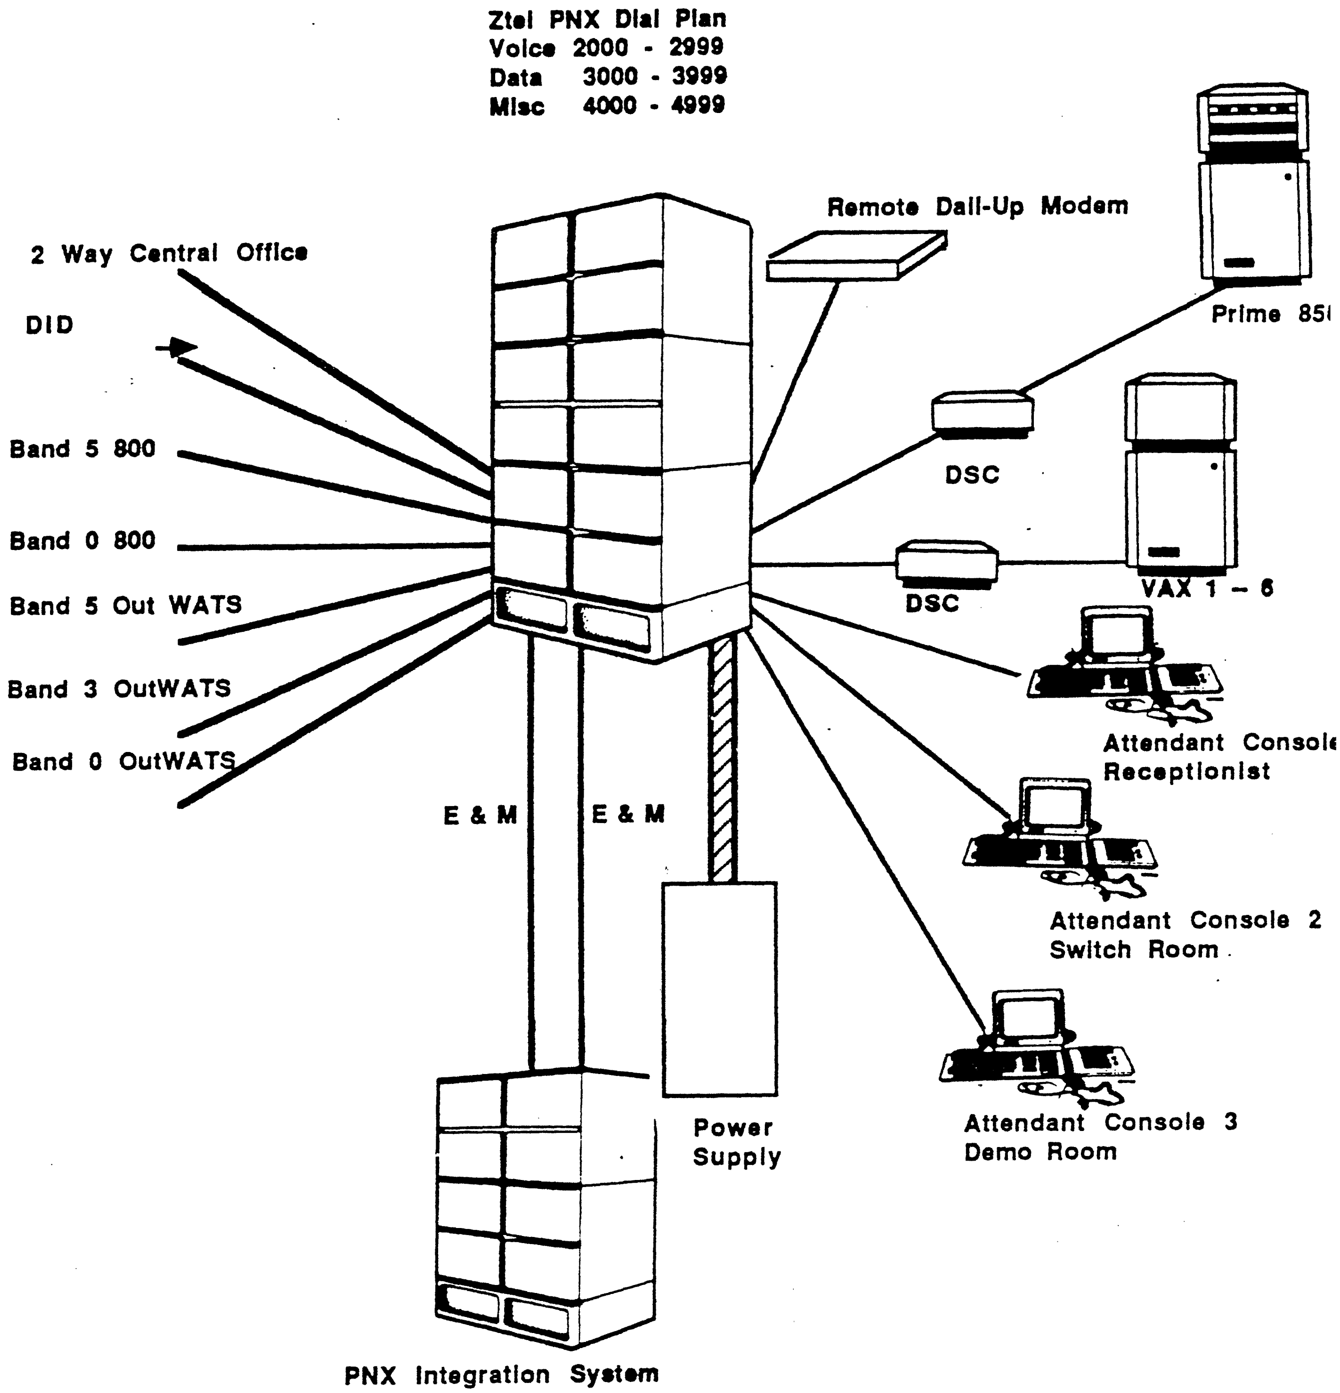

Figure II-27: Example PNX 5300 Configuration

Directory can be purchased to upgrade the basic software configuration.

The PNX system includes full menu-driven system management functions to provide on-line management capability. Use of the system's state-of-the-art attendant console can substantially increase operator productivity. Ztel's complete line of digital telephone sets and data servers provide users with highly advanced calling features such as integral messaging, speed-calling, multi-line conferencing and hands-free calling.

9.2. System Concept

The PNX automatically selects the most cost-effective routes for communication, redistributes traffic loads for better performance, and automatically converts operating characteristics between dissimilar devices. It also provides pertinent information to help determine when to add more facilities to maintain or improve a desired level of service. For qualitative decision making, the PNX provides information on a timely basis, organizing it so the user can make low-risk decisions regarding his business communications.

At present, the PNX is aimed at a wide spectrum of installations with a base of less than 100 to over 2,000 ports. The PNX (at all line/port requirements) is a full feature voice and data communiction system that is currently being used in a variety of applications including shared tenant service, financial and college environments.

Ztel's tenant support capability will enable a property's telecommunication manager to provide shared voice and data communication service on a single PNX system to multiple tenants on a property. The tenant support capability can be partitioned and tailored to meet the unique needs of each tenant.

Many of the standard functions of the PNX can be offered to tenants as add-on services, or enhancements to basic service. These include centralized messaging center service, inter-tenant messaging, and speed calling. Ztel demonstrated the tenant support capability of its PNX in June 1985.

9.3. System Components

The PNX is composed of switching, processing and line elements. Processing elements perform all call processing system management functions under the control of the operating system software. Switching elements contain the hardware required to perform time division and space division switching for ports in the system. Line elements establish the specific communication lines and trunks configured in the system. The Mass Storage Unit (MSU) hardware assembly contains a tape cartridge drive, a 40 MB

Winchester disk drive, controlling modules and a dc/dc power supply.

9.3.1. Processing Element Module

The Processing Element Module is a 16 x 16 inch printed circuit card that plugs into the processing element backplane slot. The Processing Element contains the following types of modules:

1. Basic Processing Module (BPM) - a MC68010-based Ztel CPU module that incorporates a Memory Management Unit (MMU), Direct Memory Access (DMA) circuitry, and four serial ports.

3. Line Communications Controller (LCC) - packetizes and depacketizes control information between a switch element and terminal equipment. It is a companion to the TSI module.

4. Time Slot Interchange (TSI) - a time division switch that switches time-slotted TDM data under memory mapped control. Up to 512 ports are supported by a single TSI.

9.3.2. Switching Element Module

The Switching Element Module contains the following modules:

1. Line Communications Controller (LCC) - packetizes and depacketizes control information between a switch element and terminal equipment. It is a companion to the TSI module.

2. Local Memory Module (LMM) - Identical to that used in processing elements and used as shared memory.

3. Shared Bus Adapter (SBA) - Connects a switching element to a shared bus from processing elements.

4. Time Slot Interchange (TSI) - a time division switch that switches time-slotted TDM data under memory mapped control. Up to 512 ports are supported by a single TSI.

9.3.3. Line Element Modules

The Line Element Modules contain the following five modules in the Line/Trunk Interfaces:

1. Analog Line Card (ALC) - Interfaces 16 type 2500 DTMF phones to the PNX.

2. Digital Line Card (DLC) - Interfaces 16 lines from Ztel digital telsets or data servers to the PNX. Accepts voice-only, data only, or voice-data devices.

3. Ground Start/Loop Start (GS/LS) Trunk Card (TC) - Interfaces 8 telephone company Central Office trunks to the PNX.

4. E&M (2-wire, 4-wire, type 1 & 2) Trunk Card - Interfaces 8 E&M trunks to the PNX.

5. Direct Inward Dial Card (DID) - Allows outside lines to access the PNX phones without attendant assistance.

In addition, the Line Element Modules contain the following four modules in the Service Cards:

1. Clock Card (CLK) - Outputs the timing signals required for circuit switched operation in an SPU.

2. Conference Card (CONF) - Enables conference calls on the PNX. Each Conference Card module provides 16 conferencing ports that can be mixed in any combination.

3. DTMF Card (DTMF) - Generates and decodes PCM encoded DTMF tone pairs. This module is required in PNX systems that interface with analog circuits and equipment.

4. Tone Card (TONE) - Supplies up to 16 signal tones for use in the SPU. Interfaces to 4 independently controlled paging zones. It also provides control for 4 UNA relays.

9.3.4. Digital Telsets

The PNX system features three digital telephone sets (Telsets): the Z/28, Z/12, and Z/4 telsets. The Z/28 is a 28 key programmable telset with a 40 character LCD; the Z/12 is a 12 programmable feature key telset. Both the Z/28 and Z/12 can be configured to support data and voice communication. The Z/4 contains four definable feature keys. The PNX also supports standard 500/2500 analog station telephones.

The telsets feature keys can be programmed to let users send and receive messages. When a busy signal or no answer is received, a message can be sent to any other telset by pushing the feature key with an appropriate message (such as Please call back, returned your call, etc.). Messaging automatically records the caller's name and extension number, and the time and date of calling. A centralized message center or any data terminal with the Terminal User Interface (TUI) option also can send and receive both preprogrammed or free-form text messages to individual telsets on the system. The Z/28 40 character LCD display allows users to read messaging directly from the telset. All voice and data calling features not assigned to feature keys can be activated by entering simple access codes.

The TUI option allows a data terminal to access other terminals and devices through the PNX at transmission speeds up to 19.2 Kbps asynchronous or 56 Kbps synchronous.

9.3.5. Data Servers

Ztel also provides Basic Data Servers for uses who do not need voice communication devices. In addition, a data server cluster can contain up to 16 data server modules. Data Server Clusters can be used in a variety of ways including modem pooling and computer room applications.

A Data Server connects to the SPU over a separate line of 2 twisted pairs or quad wire. The Data Server has interface connections for asynchronous (supports EIA standard RS-232C, CCITT standard V.24, V.28, X.21 bis) or synchronous type lines.

9.3.6. Attendant Console

The CRT-based attendant console includes a custom keypad and alphanumeric keyboard provides efficient and courteous connection of incoming and internal telephone calls. The CRT displays a variety of information, including source and destination of calls, the current saved (or holding calls), the status of calls and alarms on the system. The console also displays On-Line Directory, Messaging, Alarm Status, Trunk Group Status, Station Feature Editor, and Attendant Notepad displays when these features are part of the console configuration.

9.3.7. Other Components

For administration purposes, the PNX encorporates a full featured menu-driven system management console. The console provides access to all system management functions as well as detailed traffic and statistical reporting and on-line maintenance and diagnostics. Through the system, full network management and control functions provide control over system configuration of all voice and data ports, and provide for assigning all voice and data call processing features. The system provides for management of all voice and data hunt groups, telset voice and data line and station profiles, trunk group management and least cost routing.

For maintenance, the system incorporates on-board maintenance processors. This consists of a network of processors, called the Maintenance Thread, which monitors the system to enhance system reliability, availability and serviceability. The Maintenance Thread provides continual testing. Faults can be detected and isolated. Additional maintenance diagnostics can be activated by on-site or from a remote location.

9.4. PNX Family

In January 1986, Ztel introduced two new PNX system offerings, the PNX 5200 and PNX 5400, which support voice and data line size configurations of up to 300 and 1,500 lines, respectively.

The new PNX 5200, which supports up to 512 ports or 300 lines, is aimed directly at small businesses. It offers the same features and functionality of the original PNX, now called the PNX 5300, but at a lower price.

The features that now apply across the family of PNX systems include a CRT-based attendant console; an on-line directory; integrated system-wide messaging; full data transmission capabilities; least cost networking; advanced system administration; service and diagnostic tools; and product warranties that exceed industry norms.

The PNX 5400 is essentially a doubling of the original system, which means it supports up to 2,048 ports or 1,500 lines.

The PNX 5200 is ready for shipment immediately, and the 5400 is scheduled for shipment in first quarter 1986.

9.5. Pricing

Installed pricing for a system depends on the system configuration, number of digital telsets used, trunking and a variety of parameters. The range of price per line can be estimated at between $600-$1000 per line.

9.6. Evaluation

Ztel appears again to be a stable company with secured financing producing a highly advanced product. It has a substantial customer base running numerous applications involving a varied mixture of voice and data communications. Its aggressive distribution plan provides for territorial exclusivity within regional areas.

Ztel is enhancing the PNX's networking capability by incorporating means to gateway to industry standard local area (LAN) and wide area (WAN) networks. The PNX will gateway to 802.3 (Ethernet), 802.4 (MAP) and 802.5 (Token passing ring) networks.

SECTION III

Supplemental Information

1. Soderblom Patent

The following is a copy of U.S. Patent Number Re.31,852, reissued March 19, 1985, to Olof Soderblom. It expands Soderblom's original claims on his token-passing ring invention and forms the basis for his demand for royalties from any vendor building token ring local networks.

United States Patent [19] [11] E **Patent Number: Re. 31,852**

Soderblom [45] **Reissued Date of Patent: Mar. 19, 1985**

[54] **DATA TRANSMISSION SYSTEM**

[75] Inventor: **Olof Soderblom**, Leatherhead, United Kingdom

[73] Assignee: **Willemijn Houdstermaatschappij BV**, Rotterdam, Netherlands

[21] Appl. No.: **394,099**

[22] Filed: **Jul. 1, 1982**

Related U.S. Patent Documents

Reissue of:
[64] Patent No.: **4,293,948**
Issued: **Oct. 6, 1981**
Appl. No.: **518,450**
Filed: **Oct. 29, 1974**

U.S. Applications:
[63] Continuation of Ser. No. 391,717, Aug. 27, 1973, abandoned, which is a continuation of Ser. No. 773,056, Nov. 4, 1968, abandoned.

[30] **Foreign Application Priority Data**

Nov. 23, 1967 [SE] Sweden 16077/67
May 15, 1968 [SE] Sweden 6589/68

[51] **Int. Cl.**[3] **H04J 3/16**
[52] **U.S. Cl.** **370/90**; 370/86; 370/89; 370/97
[58] **Field of Search** 370/90, 89, 87, 86, 370/97; 340/825.05, 825.5

[56] **References Cited**

U.S. PATENT DOCUMENTS

Re. 21,284	12/1939	Boswau	340/825.14
2,091,301	8/1937	Boswau	340/825.14
2,177,609	10/1939	Boose	178/23 R
2,185,851	1/1940	Kinkead	178/2 D
2,238,142	4/1941	Weaver	178/2 D

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

702947	of 0000	Fed. Rep. of Germany .
1041104	of 0000	Fed. Rep. of Germany .
1053360	of 0000	Fed. Rep. of Germany .
1121107	of 0000	Fed. Rep. of Germany .
1169009	of 0000	Fed. Rep. of Germany .
1187263	of 0000	Fed. Rep. of Germany .
1197513	of 0000	Fed. Rep. of Germany .
1206183	of 0000	Fed. Rep. of Germany .
1448803	of 0000	France .
62-4921	of 0000	Japan .
188775	of 0000	Sweden .
882207	of 0000	United Kingdom .
903806	of 0000	United Kingdom .
930128	of 0000	United Kingdom .
987293	of 0000	United Kingdom .
993818	of 0000	United Kingdom .

OTHER PUBLICATIONS

Report of Study Made by the French Firm Le Materiel, Telephonique in LMT 441.632-A, issue 2 of Dec. 27, 1962.

(List continued on next page.)

Primary Examiner—Douglas W. Olms
Attorney, Agent, or Firm—Pollock, Vande Sande & Priddy

[57] **ABSTRACT**

The main features of the present invention reside in the fact that one or more remote groups of subordinated terminal units are serially interconnected and arranged in a closed loop which is terminated at the output and input of a master unit. Each of said groups and said master unit is connected to the transmission line via a modulator and a demodulator. In order to selectively connect a pulse equipment associated with each subordinated terminal unit, and adapted for transmission of data, with the transmission line, each of said subordinated terminal units includes a switching unit, which either connects or disconnects the associated pulse equipment from the transmission line. According to the invention, the novelty resides both in the arrangement of the network connecting the subordinated terminal units with the master pulse equipment, and the procedure for the exchange of information between the master pulse equipment and the subordinated terminal unit, and vice versa, through the closed loop, and by selectively actuating said switching unit.

33 Claims, 4 Drawing Figures

Re. 31,852

Page 2

U.S. PATENT DOCUMENTS

2,332,645	10/1943	Keithley et al.	178/2 C
2,374,908	5/1945	Weaver	178/2 C
2,502,654	4/1950	Keyes	178/2 D
2,505,728	4/1950	Zenner	178/4
2,694,802	11/1954	Terry et al.	340/825.51
2,793,357	5/1957	Wynn	340/825.09
2,840,797	6/1958	Derr	340/825.05
2,954,431	9/1960	Royden	178/17 R
2,974,187	3/1961	Barrett et al.	178/3
2,982,809	5/1961	Light et al.	178/2 R
2,986,602	5/1961	Tubinis	370/55
3,016,516	1/1962	Doersam	370/18
3,043,905	7/1962	Semancik et al.	178/3
3,078,336	2/1963	Cameron et al.	178/3
3,120,606	2/1964	Eckert et al.	364/736
3,208,049	9/1965	Doty et al.	340/825.26
3,214,733	10/1965	Spencer	370/85
3,221,307	11/1965	Manning	364/200
3,239,819	3/1966	Masters	364/200
3,242,467	3/1966	Lamy	364/200
3,244,804	5/1966	Wittenberg	178/3
3,245,043	5/1966	Gaffney et al.	340/825.5
3,281,797	10/1966	Harris	364/900
3,289,165	11/1966	Hawley et al.	370/86
3,303,470	2/1967	Brixner et al.	340/825.08
3,303,476	2/1967	Moyer et al.	364/200
3,336,577	8/1967	Frielinghaus	340/870.11
3,336,582	8/1967	Beausoleil et al.	364/200
3,377,619	4/1968	Marsh et al.	364/200
3,445,822	5/1969	Driscoll	364/200
3,456,242	7/1969	Lubkin et al.	364/900
3,480,914	11/1969	Schlaeppi	364/200
3,483,329	12/1969	Hunkins et al.	370/86
3,517,130	6/1970	Rynders	370/85
3,519,750	7/1970	Beresin	370/16
3,534,337	10/1970	Martin et al.	364/900

OTHER PUBLICATIONS

"A Fast-Response Data Communications System for Airline Reservations" by Levine et al. Trans. A.I.E.D., Pt. 1 (Communication & Electronics) Jan. 1963, pp. 590-600.

"Data Transmission and Switching Equipment for the Seat Reservation System of United Airlines", Philips Telecommunications Review, vol. 24, No. 1, Feb. 1963, pp. 13-24.

IBM Technical Disclosure Bulletin No. 7 (1964/65) pp. 592-593.

Teknisk Tidskrift 1959, pp. 339-344 by J. M. Unk.

Elektroniskt Fjarrstyrningssystem JZA 41 SIB 486/1964 (Leaflet from Swedish Firm L. M. Ericsson AB).

"Design of Real-Time Computer Systems"; J. Martain; Prentice-Hall; published Oct. 9, 1967; pp. 291-294.

Fig.1

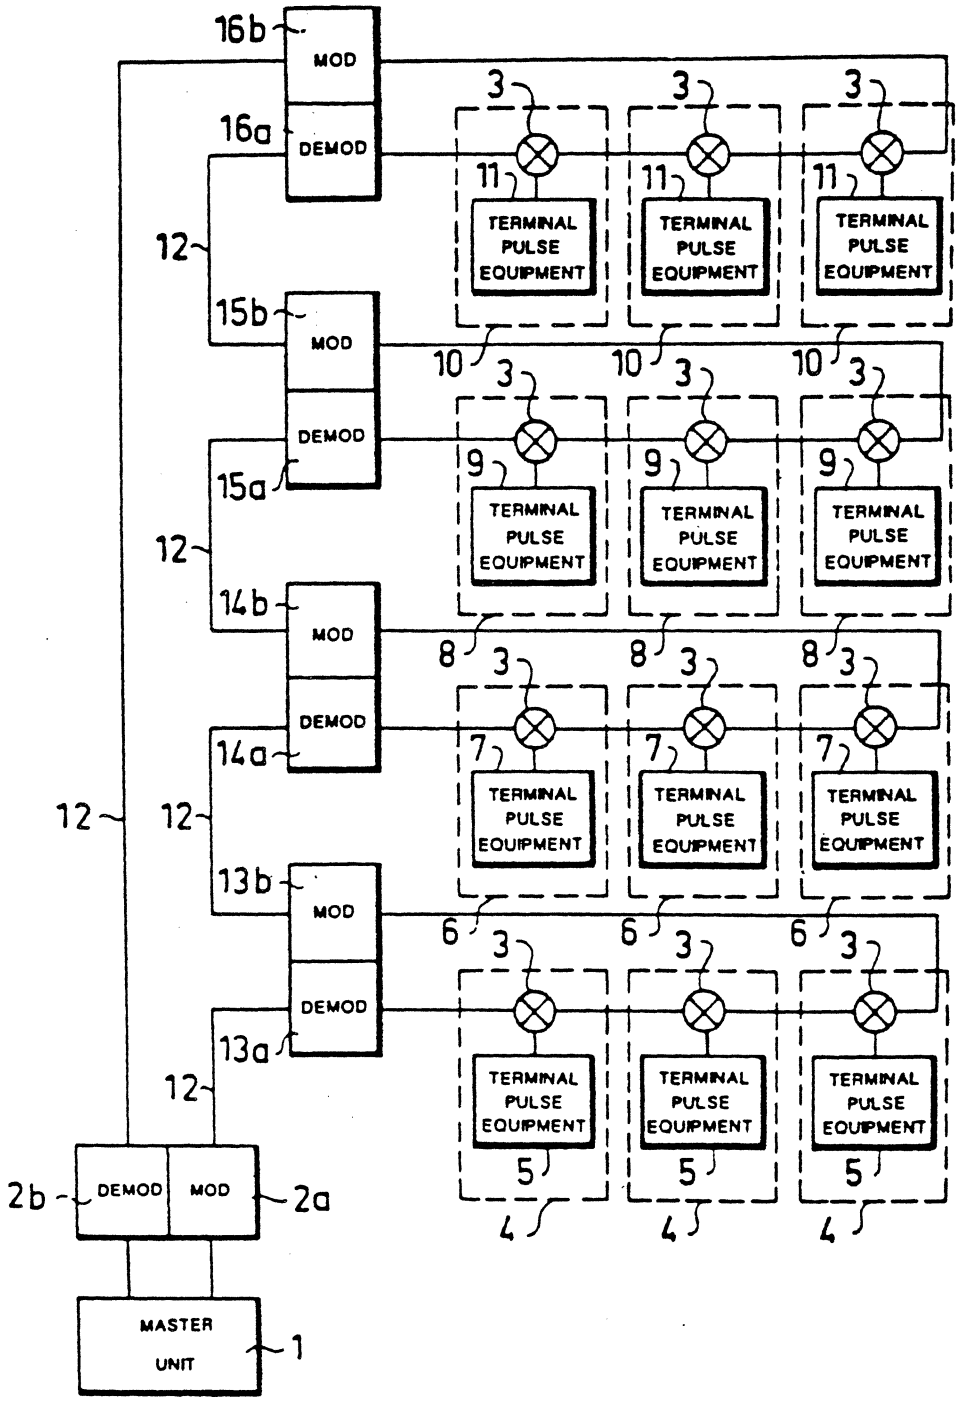

16b
MOD
16a
DEMOD
15b
MOD
DEMOD
15a
14b
MOD
DEMOD
14a
13b
MOD
DEMOD
13a
12
3
11
9
7
10
8
6
5
4
TERMINAL PULSE EQUIPMENT
2b
DEMOD
MOD
2a
MASTER UNIT
1

Fig.2

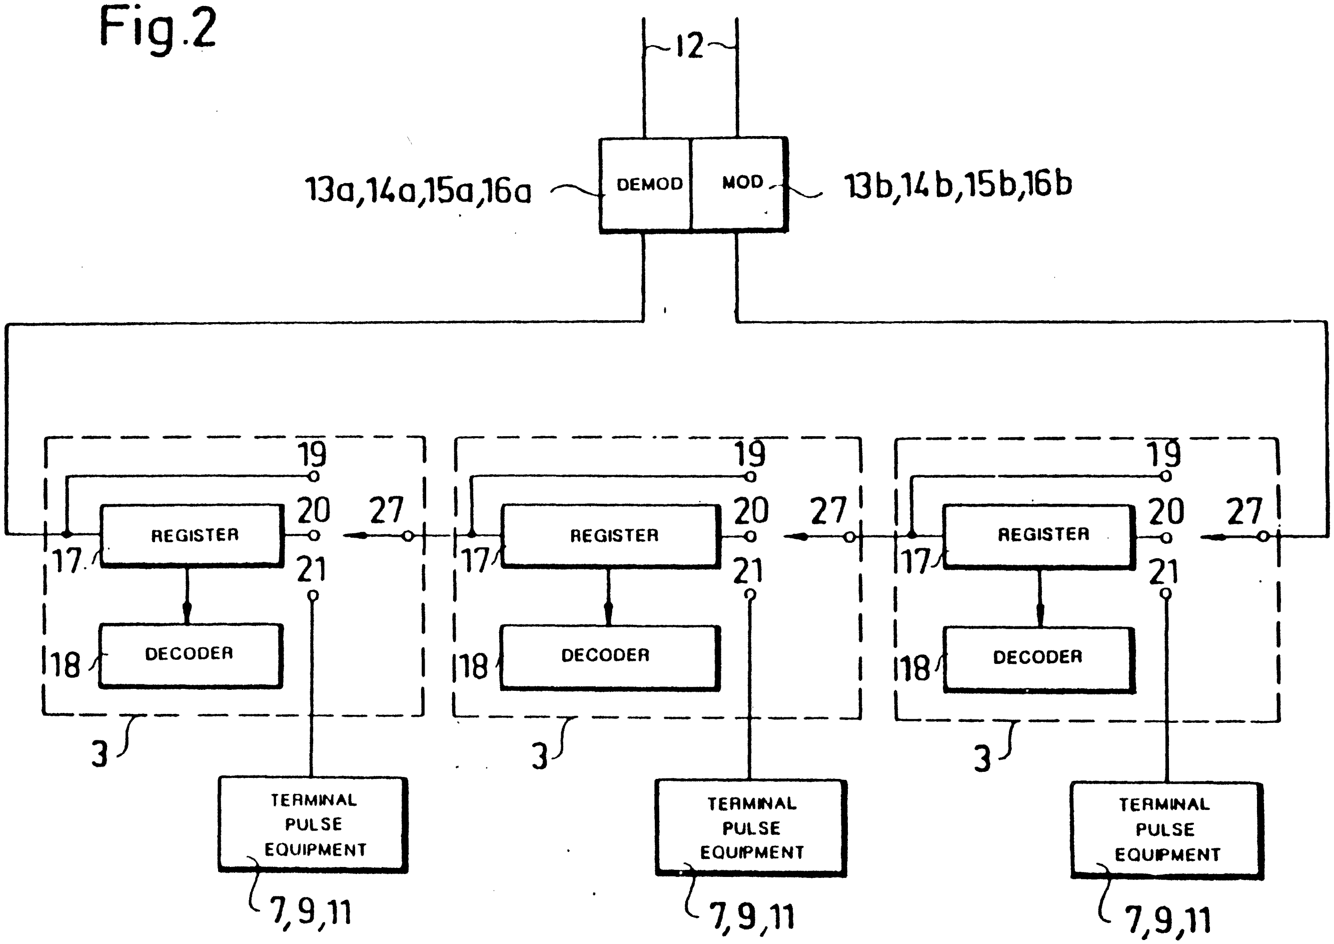
12
DEMOD
MOD
13a,14a,15a,16a
13b,14b,15b,16b
19
20
21
27
17
18
REGISTER
DECODER
3
TERMINAL PULSE EQUIPMENT
7,9,11

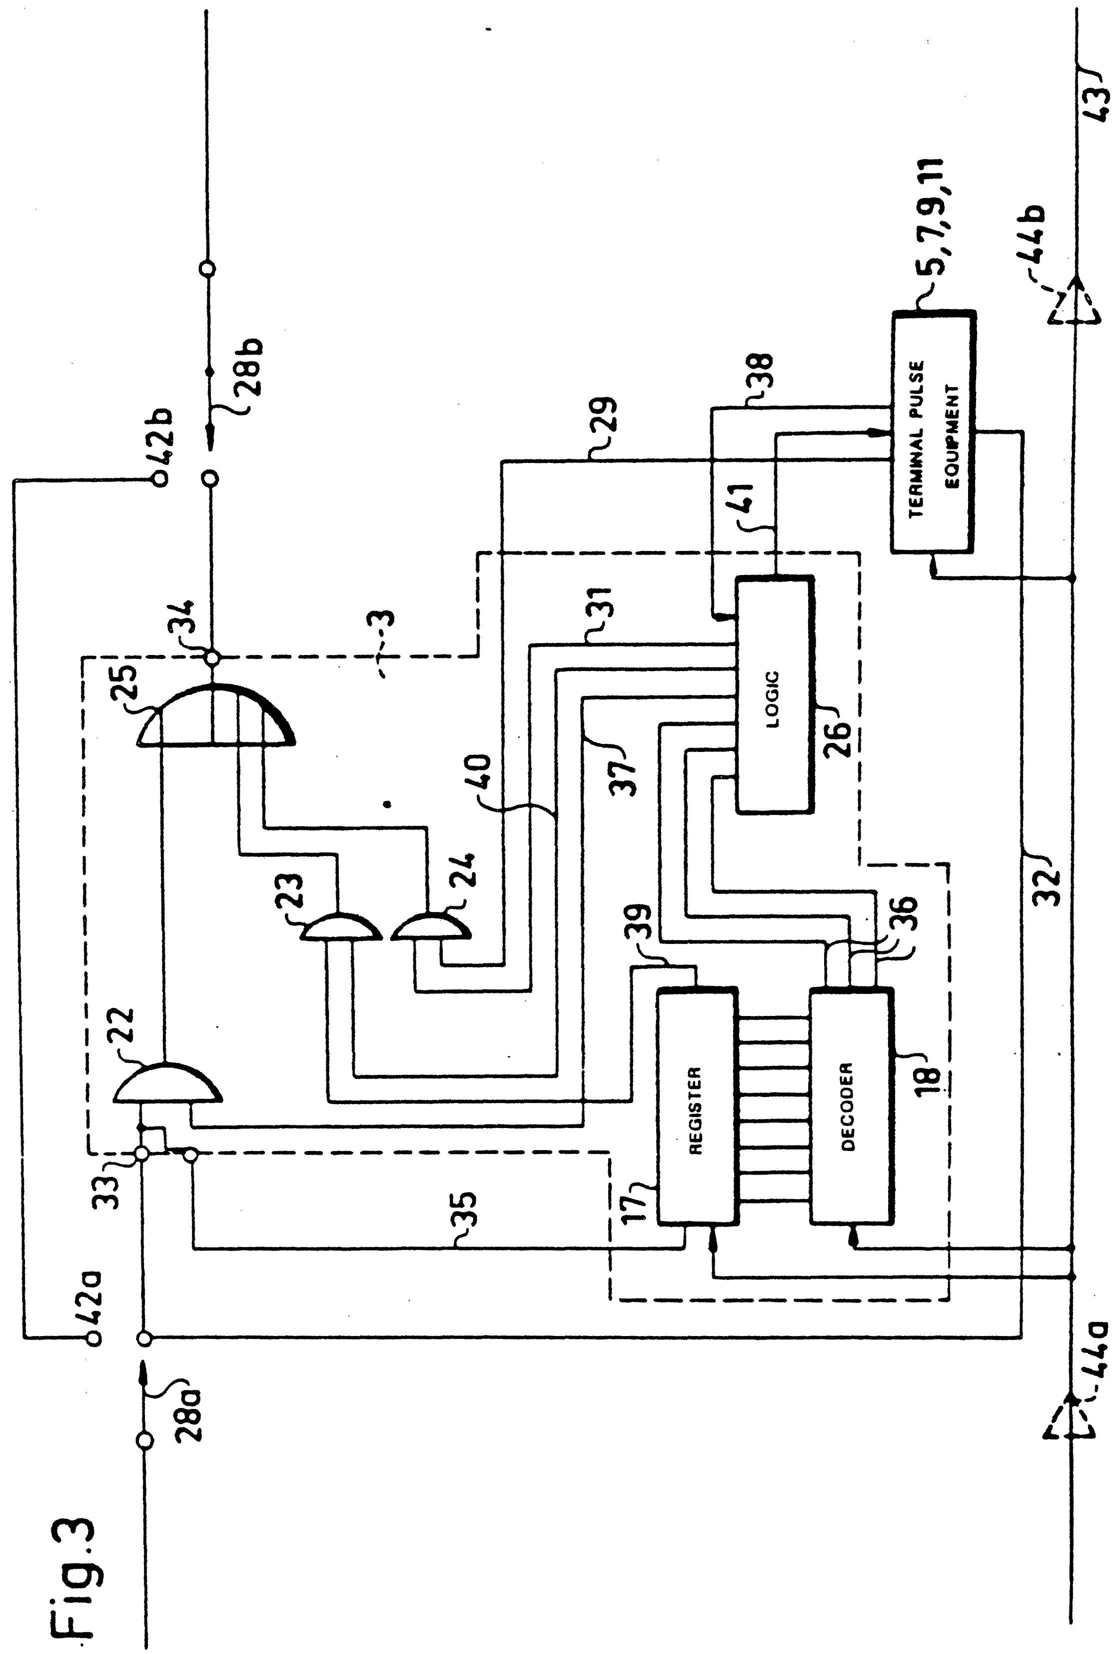
TERMINAL PULSE EQUIPMENT
LOGIC
REGISTER
DECODER
5,7,9,11
43
44b
44a
42b
42a
28b
28a
29
38
41
31
34
3
25
26
37
40
23
24
39
36
22
32
33
35
17
18

Fig.3

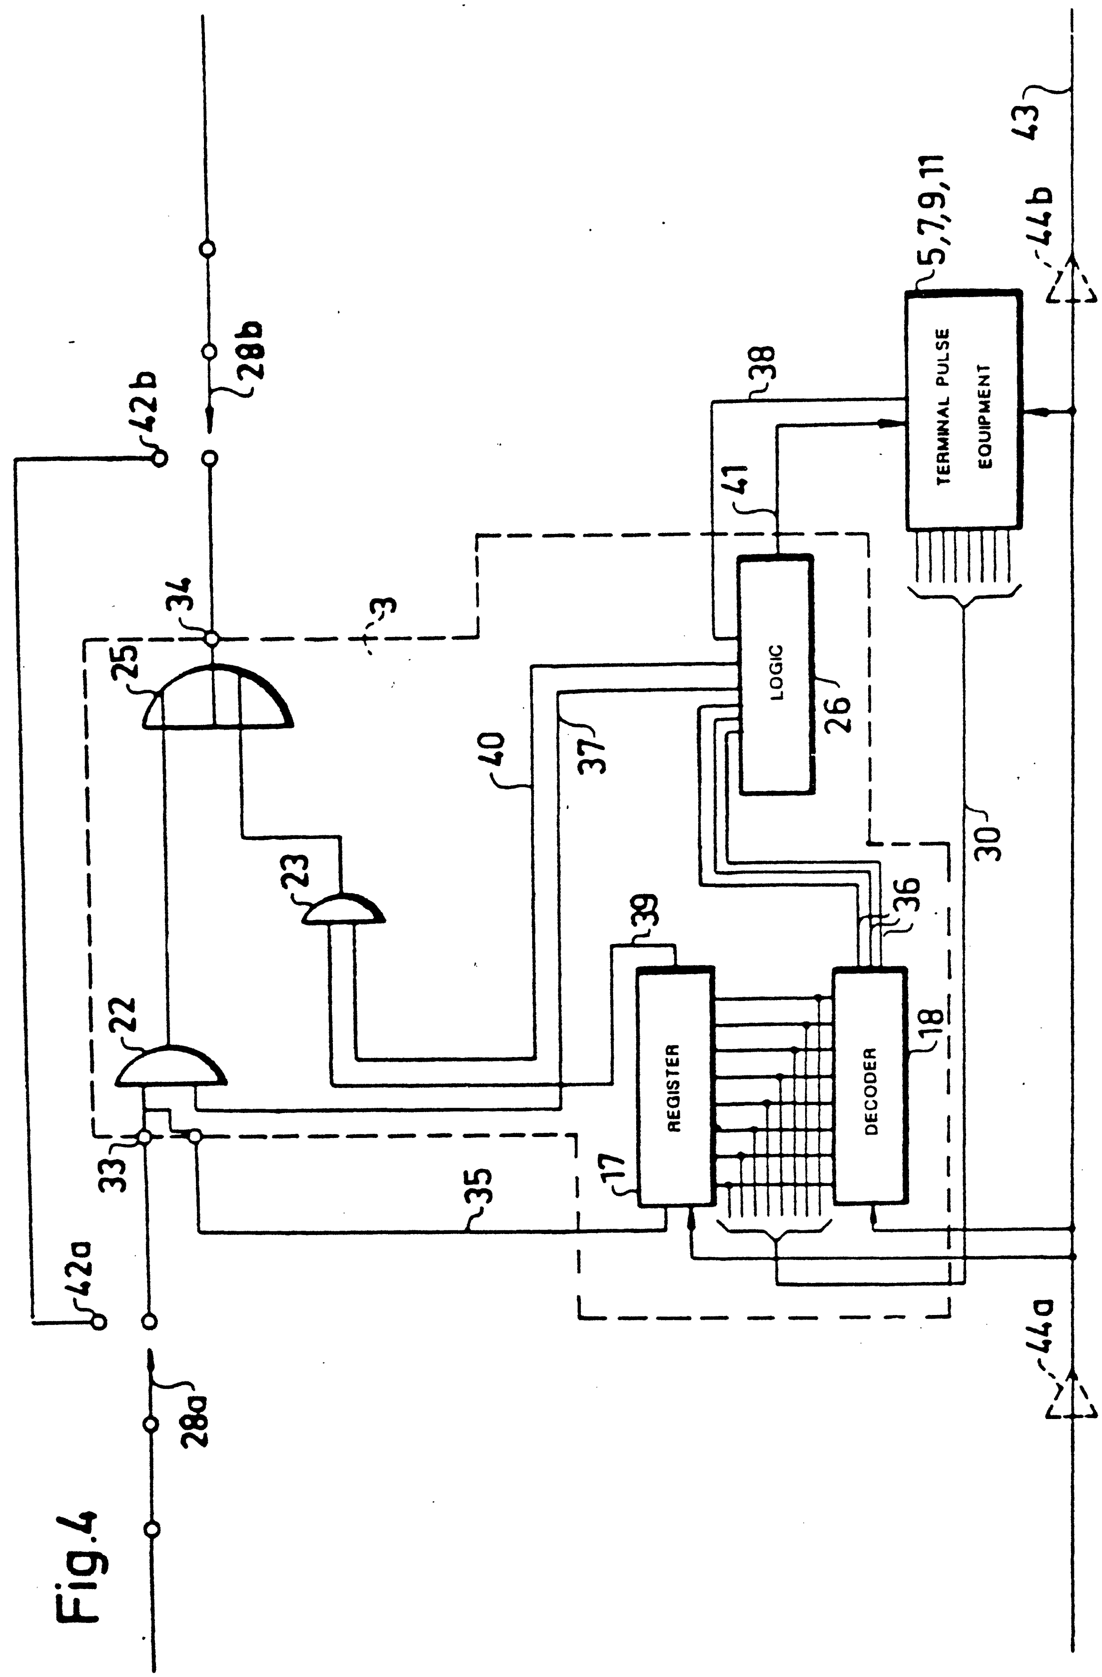
TERMINAL PULSE EQUIPMENT
LOGIC
REGISTER
DECODER
5,7,9,11
43
44b
44a
38
41
42b
42a
28b
28a
34
33
35
3
25
26
40
37
23
39
22
36
30
18
17

Fig.4

DATA TRANSMISSION SYSTEM

Matter enclosed in heavy brackets [ ] appears in the original patent but forms no part of this reissue specification; matter printed in italics indicates the additions made by reissue.

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 391,717 filed Aug. 27, 1973. Application Ser. No. 391,717 *now abandoned* is a continuation of U.S. patent application Ser. No. 773,056 filed Nov. 4, 1968 and now abandoned.

BACKGROUND OF THE INVENTION

The present invention relates to a data transmission system for transmitting order-characters, address-characters and information-characters between a master unit and one or more remote groups of subordinated terminal units, the several units in a group normally being disposed at a common location such as a branch bank or store. Each of said groups and said master unit includes a modulator and a demodulator, i.e. a so-called "modem", adapting pulses from the master unit and the subordinated terminal units to the transmission characteristics of the transmission lines which connect said master unit and said subordinated terminal units to each other.

To interconnect an equipment built up of pulse circuits with another such equipment over a relatively long distance, it is necessary to convert the pulses into signals which are adapted to the frequency characteristics of the transmission lines. For this purpose, it is common practice to employ at one end of the line a modulator for converting a pulse into either a phase shifted or frequency shifted pulse relative to a reference signal, and to connect the other end of the line to a demodulator detecting this phase or frequency shift and producing a pulse in correspondence thereto.

In systems where a plurality of pulse equipments disposed at geographically remote places are to be connected to a master pulse equipment, it is common practice to connect these remote pulse equipments and their associated modulator/demodulator units in parallel to one single transmission line for the purpose of shortening the overall length of the transmission lines required. Further, it is common practice to concentrate all pulse equipments which have a common location to one single modulator/demodulator unit by means of a traffic concentrator which, in addition to its concentrating function, effects the control functions required for the above-mentioned parallel connection. This equipment, which effects the interconnection between a plurality of pulse equipments and one single modulator/demodulator unit and enables the connection of a plurality of such units of pulse equipments in parallel to one single transmission line, tends to become rather complicated.

A further disadvantage of connecting a plurality of such units in parallel to one single transmission line resides in that the most efficient mode of transmission, i.e. so called synchronous transmission, cannot be utilized, due to the fact that one demodulator will have to receive signals in alternating order from a plurality of modulators.

A further drawback of the above-mentioned parallel connection resides in that a great number of signals are

required for maintaining the order of succession between the units connected to the line, and such sequence signals will intrude upon the time available for the transmission of information.

SUMMARY OF THE INVENTION

The object of the present invention is to provide a system permitting a plurality of commonly located pulse equipments to be connected to one modulator/-demodulator unit, and for several modulator/-demodulator units disposed at geographically spaced places to be connected to one single transmission line without using any complicated equipments therefor and without involving the drawbacks of a conventional parallel system as indicated hereinbefore, while still maintaining the advantages of the parallel system.

This object of the invention is substantially realized by the fact that the transmission lines connecting the master unit with the modems and subordinated terminal units included in the system are arranged in a closed loop in which modulators, demodulators and subordinated terminal units are serially interconnected. All signals comprising order-characters, address-characters and information-characters are transmitted in the closed loop which is arranged such that the output from the modulator of the master unit is connected to the input of the first demodulator appearing in the loop associated with a first group of subordinated terminal units, which are serially connected with each other and form a closed loop between the output of said first demodulator and the input of the associated first modulator. The output of the first modulator may be connected to the input of a succeeding second demodulator in the series circuit which is associated with a second group of subordinated terminal units and so on, until finally, the last modulator occurring in the loop has its output connected to the input of the demodulator in the master unit, thereby closing the loop.

According to the invention, all modulators, demodulators and subordinated terminal units are serially interconnected between the data output and data input of the master pulse equipment which normally transmits a stream of synchronization characters and, at times, transmits order-characters or information-characters. The first subordinated terminal unit having information to send to the master pulse equipment and receiving order-characters from said master pulse equipment, interrupts the series circuit and transmits the information to the data input of the master pulse equipment via succeeding modulators, demodulators, and subordinated terminal units. The transmission ends with distinctive order-characters which are transmitted around the rest of the loop downstream of such first station and which cause the succeeding subordinated terminal unit with information to send to once again interrupt the series circuit at such succeeding station and send the information to the data input of the master pulse equipment. The said succeeding station similarly ends the transmission with order-characters, and so on, until the last subordinated terminal unit with information to send has, in turn, emitted order-characters which are then finally received by the master pulse equipment, thereby establishing that all subordinated terminal units have sent their information. Information from the master pulse equipment to a particular subordinated terminal unit includes address-characters.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will now be described more in detail with reference to the accompanying drawings, in which:

FIG. 1 shows a block diagram illustrating the closed loop of the data transmission system according to the invention;

FIG. 2 shows a functional diagram illustrating the function of the switching unit in the transmission system according to FIG. 1;

FIG. 3 shows a detailed block diagram of the switching unit; and

FIG. 4 shows an alternative embodiment of the switching unit.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

According to the invention, a plurality of remotely disposed groups of pulse equipments **5, 7, 9, 11** are interconnected and connected to a master pulse equipment **1** via modulator/demodulator units **2a, 2b, 13a, 13b, 14a, 14b, 15a, 15b, 16a, 16b** and transmission lines **12**. These modulator/demodulator units operate in a manner such that the modulators **2a, 13b, 14b, 15b,** and **16b** accept pulses at the same rate as pulses are transmitted from the output of the associated demodulators **2b, 13a, 14a, 15a** and **16a**, respectively.

The system is arranged so that modulator **2a** of the master pulse equipment **1** is connected to demodulator **13a** via a portionn of line **12**. The pulse output of demodulator **13a** is connected to the pulse input of modulator **13b** via a number of switching units **3**. Each switching unit **3** selectively connects a pulse equipment **5** to the master pulse equipment **1**. Each pulse equipment **5** and corresponding switching unit **3** form a subordinated terminal unit **4**. Modulator **13b** is connected to demodulator **14a**, associated with the succeeding group of subordinated terminal units, via another portion of line **12**. Demodulator **14a** is connected to modulator **14b** via switching units **3** each of which is connected to a pulse equipment **7** which together form a subordinated terminal unit **6**. The system can include an arbitrary number of groups of subordinated terminal units. In the embodiment shown in FIG. **1**, the system is terminated at the modulator **16b**, which is connected to demodulator **2b** at the master pulse equipment. The system according to the invention is not restricted to the number of modulator/demodulator units and switching units **3** as specifically shown in FIG. **1**, but is operable with an arbitrary number of both modulator/demodulator units and switching units.

The function of the switching unit **3** is, in principle, illustrated in FIG. **2**, where it is shown as a three-way switch, including switching means **27**, a by-pass line, a shift pulse register **17** and a decoding unit **18**. In position **19**, the switching unit **3** is by-passed in the sense that any characters appearing at the output of the demodulator are not only shifted into the register **17** particularly associated with such switch but are also simultaneously applied to and shifted into the register **17** of the next downstream terminal unit. Obviously, therefore, if all switches **17** are in position **19**, any character transmitted by the master pulse equipment will be shifted substantially simultaneously into every register **17**.

In position **20**, the output of the switch **3** is connected to the input through the pulse register **17** which has a capacity equal to the number of bits forming one unit of information, here called a character. Under these circumstances, the register **17** of the next-downstream terminal unit cannot be directly responsive to a character provided by the demodulator output; instead, such further register can only be responsive to whatever character is shifted out of the first-mentioned register **17** directly associated with the switch which is in position **20**.

In the third position **21**, the output of the switch **3** is connected directly to the associated pulse equipment **5, 7, 9** or **11**. In this circumstance, the associated terminal pulse equipment is effectively connected to the loop, as will be shown, so that this terminal unit can not transmit data back to the master pulse equipment.

To the pulse register **17**, which is permanently connected to the input of the switch unit is also connected a decoding unit **18** which identifies certain distinctive characters in the pulse register **17**. The manner in which switching unit **3** affects the function of the system according to the invention is as follows:

The system possesses three states of operation, i.e. an idle state, a second state in which information can be transmitted from pulse equipments **5, 7, 9, 11** to the master pulse equipment **1**, and, finally, a third state in which information can be transmitted from the master pulse equipment **1** to pulse equipments **5, 7, 9** and **11**.

In said idle state, the switching units **3** are all by-passed in the sense described above and the master pulse equipment **1** transmits a continuous flow of characters, each character comprising a pulse combination or character A. The decoder **18** in each terminal unit recognizes the presence of the A character in the associated register **17** and, in response thereto, causes switch **27** to assume position **19**. As a result, all the switching units **3** are by-passed, and the pulse output of each demodulator is connected directly to the pulse input of the associated modulator, the flow of characters transmitted by pulse equipment **1** being regenerated in each modulator/demodulator unit and finally returning to the pulse equipment **1** through the demodulator **2b**. This flow of characters has two purposes, viz.

(a) to establish synchronism between the oscillator of a demodulator and the oscillator of the preceding modulator, such as, for example, between demodulator **14a** and modulator **13b**;

(b) to establish synchronism between the pulses transmitted as a character by pulse equipment **1**, and the pulses detected as a character by pulse register **17**. When the characters received by pulse equipment **1** through demodulator **2b** contain the same pulse combination as the characters transmitted from pulse equipment **1**, the entire system has become synchronized.

Another method of synchronizing a demodulator with the preceding modulator resides in maintaining a two-way connection between them i.e. to each demodulator an auxiliary modulator is connected, which transmits pulses to an auxiliary demodulator connected to the preceding modulator in the system. By this method, signaling will indicate whether synchronism is being maintained separately between each modulator/demodulator connection.

When the master pulse equipment **1** is ready to receive information from pulse equipments **5, 7, 9, 11**, pulse equipment **1** will interrupt the flow of A-characters and transmit two distinctive characters in succession, B and C, and then recommence its transmission of A-characters. The purpose of said B- and C-characters is the following:

The B-character is shifted into the register 17 of each terminal unit substantially simultaneously since each switch 27 is still in position 19. When a B-character has been identified in each pulse register 17 by decoding unit 18, the following procedure will be initiated, provided that the associated pulse equipment 5, 7, 9, 11 has information to transmit. (1) The switching means 27 is set to position 20. (2) The next succeeding character, i.e. the C-character, can now be received only in the register 17 of the first upstream terminal having data to transmit; the reason for this, of course, is that the switch 27 of such first upstream station was operated to position 20 by the B-character, thereby effectively opening the loop so as to prevent any other register 17 from responding to the C-character. (3) Assuming that such next character is a C-character, the switching means 27 of such first upstream unit with data to transmit is set to position 21, thereby permitting only the characters transmitted from the associated pulse equipment to be put on the line. Since the character C only appears in the register 17 of such one unit, it cannot be advanced to the next succeeding terminal unit in the series circuit in response to the next character placed on the loop by the master pulse equipment because with switch 27 in position 21, the register 17 storing the C-character is effectively disconnected at its output from all other downstream registers. As a consequence, no succeeding pulse register 17 will receive character C directly after having received character B. This will cause the associated switching means 27 to be reset to the idle position.

If the terminal unit has no information to transmit when the B-character is detected in pulse register 17, the switching means 27 will retain the idle position. The information transmitted from each of the pulse equipments 5, 7, 9 or 11 will pass through the succeeding switching units in the series circuit to the input of the master pulse equipment 1.

Each transmitting pulse equipment 5, 7, 9 or 11 ends its transmission with the characters B and C and will then reset switching means 27 to the idle position, whereby the sequence of procedures described above will be repeated in the succeeding terminal units in turn.

When characters B and C are received in the master pulse equipment 1, this provides an indication that all pulse equipments 5, 7, 9, 11 in the closed loop have had an opportunity to transmit information, and pulse equipment 1 will respond by transmitting a character, D, indicating that the information transmitted from the subordinated terminal units has been properly received. However, if the information received by pulse equipment 1 contained errors, characters B and C will be transmitted once more and the procedure just described will be repeated. Thus, pulse equipments 5, 7, 9, 11 which have transmitted information but received no D-character will retransmit their information.

Transmission of information from the master pulse equipment 1 to any of the pulse equipments 5, 7, 9, 11 is performed by transmitting two characters E and X before each message. Character E indicates that the next succeeding character X, contains the address to the specific pulse equipment for which the message is meant. This message will pass through all switching units 3, but only the pulse equipment assigned number X will respond to the information.

The switching unit 3 according to the invention has as its main purpose to enable the series connection between the terminal units and the master pulse equipment.

The unit 3, shown in FIG. 3, comprises three AND-gates 22, 23 and 24, an OR-gate 25 and a control logic 26 having connected thereto a decoding unit 18 and a pulse equipment 5, 7, 9, 11. In the idle position of the switching unit 3, the input 33 and output 34 are effectively connected together through the enabled AND-gate 22 and the OR-gate 25. In the idle state, synchronism is maintained by the decoding unit 18 via a pulse register 17, being operative to decode information appearing at the input 33 of the switching unit 3 via line 35.

When the decoding unit 18 identifies a B-character in the pulse register 17, it emits a signal via lines 36 to the control logic 26 to remove the enabling input to AND-gate 22 via line 37. The pre-requisite is, however, that pulse equipment 5, 7, 9, 11 has informed control logic 26 via line 38, that it has information to transmit. At the same time, the control logic 26 enables AND-gate 23, whereby pulse register 17 is connected to the output 34 of the switching unit 3 through AND-gate 23 and OR-gate 25, corresponding to the second state of the switching unit. If the next-following character received in the pulse register 17 and decoded by decoding unit 18 is a C-character, control logic 26 will disable AND-gate 23 and enable AND-gate 24 via line 31. At the same time, control logic 26 will emit a pulse to the pulse equipment 5, 7, 9, 11 via line 41 informing the latter that it may transmit information on line 29, which now connects output 34 of the switching unit 3 with the pulse equipment via AND-gate 24 and OR-gate 25. This corresponds to the third state of the switching unit. The transmission ends with a signal from pulse equipment 5, 7, 9, 11 to the control logic 26 via line 38 informing same of the fact that all information has been transmitted. Control logic 26 then enables AND-gate 24 and disables AND-gate 22, whereupon the switching unit has been reset to its idle state.

When the switching unit 3 is in the idle state, and decoding unit 18 identifies an address-character in pulse register 17 informing that the information to follow is meant for a specific pulse equipment 5 or 7 or 9 or 11, then decoder 18 will transmit a signal to control logic 26 via lines 36. Control logic 26 then informs the pulse equipment via line 41, whereupon the information is received via line 32. When the decoder has detected a character in the pulse register 17 signifying the end of the message, the pulse equipment is informed by control logic 26. The switches 28a, 28b are meant for shunting of the switching unit 3 and the pulse equipment, for instance in case of break-down or absence of primary power. These switches may consist of a relay taking the positions 42a, 42b, respectively in case of absence of primary power or being controlled by a logic circuit actuating both switches 28a and 28b.

According to an alternative embodiment of the invention shown in FIG. 4, the switching unit operates with only two switching states. Lines 29 and 31 of FIG. 3 are replaced by parallel lines 30 in FIG. 4 between the pulse equipment 5, 7, 9, 11 on one hand, and pulse register 17 and decoder 18, on the other. Thus, if the character received in the pulse register 17 and decoded by decoding unit 18 is a C-character, the control logic 26 will inform the pulse equipment that it may transmit information on lines 30 via pulse register 17, which is connected with the output 34 of the switching unit via line 39, AND-gate 23 and OR-gate 25. Thus, in this case the conditions of the AND-and OR-gates will not be changed, i.e. that the switching unit will maintain its second switching state.

Information transmitted from the master pulse equipment 1 to a specific pulse equipment 5 or 7 or 9 or 11 is received in said equipment via lines 30 after the control logic 26 has informed the equipment via line 4 that it has information to receive appearing on lines 30.

Which of these two embodiments of the switching unit is preferred depends on whether the pulse equipment is arranged to receive and transmit the bits which constitute a character in serial or parallel.

The rate of all information transmitted and received is preferably determined by a clock included in the demodulator and adapted to control the components of the terminal units 4, 6, 8, 10 via line 43. If required, the clock signal can be amplified at the in- and outputs of each terminal unit by means of amplifiers 44a and 44b, respectively.

The terminal units 4, 5, 8, 10 forming part of a transmission system of the kind here referred to, include, for example, data terminals or computers or any other equipment operating with pulses. In its most extreme form, the data transmission system according to the invention may include only one subordinated terminal unit interconnected with the master pulse equipment via modulator/demodulator units.

What I claim is:

1. Apparatus for the transmission of data characters in pulse form from a plurality of terminal units to a master unit comprising in combination:

pulse input and pulse output means for each said terminal unit and for said master unit,

a single series loop connecting said terminal units in series along said loop between the pulse output means and pulse input means of said master unit and over which pulses originating either with said master unit or any terminal unit are transmitted always in the same predetermined direction,

each said terminal unit including:

(a) pulse responsive means connected to said pulse input means for receiving pulses appearing at said pulse input means,

(b) decoding means distinctively responsive to pulses received by said pulse responsive means,

(c) and logic means responsive to said decoding means for normally interconnecting said pulse input and pulse output means to thereby complete the loop at said terminal unit,

(d) said logic means in response to a distinctive pattern of said pulses appearing at the associated pulse input means interrupting the loop between said pulse input and pulse output means only at a terminal unit having data to transmit and enabling the application to the associated pulse output means of stored data of arbitrary variable length as required by said terminal unit for transmission over the loop to said master unit,

(e) said logic means being effective upon the opening of said loop at the first upstream terminal unit having data to send for inhibiting the receipt at any downstream terminal unit of pulses otherwise effective to control said downstream terminal unit to transmit data.

2. The apparatus of claim 1 wherein each said terminal unit at the conclusion of its data transmission transmitting over said loop pulses to thereby control the downstream terminal units having data to send to serially transmit their data to said master unit.

3. The apparatus of claim 1 wherein said logic means in response to a first distinctive order character appearing at the associated pulse input means opens the normal said interconnection between said pulse input and pulse output means to thereby enable only the first upstream terminal unit having data to send to respond to a second distinctive order character,

said logic means being responsive to said second distinctive order character at said pulse input means to enable the application to said pulse output means of the stored data.

4. The apparatus of claim 3 wherein said pulse responsive means comprises a shift register and wherein said logic means when rendered responsive to said second distinctive order character operatively connects said pulse output means to the output of said register to thereby cause the contents of said register to be shifted out of said register and into the register of the next downstream terminal unit in response to the appearance of pulses at said pulse input means.

5. The apparatus of claim 1 in which said master unit includes means for transmitting data characters over said loop in said predetermined direction to one or more selected terminal units.

6. Apparatus for the transmission of data characters in pulse form from a plurality of terminal units to a master unit comprising in combination:

pulse input and pulse output means for each said terminal unit and for said master unit,

a single series loop connecting said terminal units in series along said loop between the pulse output means and pulse input means of said master unit and over which pulses originating either with said master unit or any terminal unit are transmitted always in the same predetermined direction,

each said terminal including:

(a) a pulse responsive means connected to said pulse input means for receiving pulses appearing at said pulse input means,

(b) decoding means distinctively responsive to the pulses received by said pulse responsive means,

(c) and logic means being distinctively responsive to at least one distinctive pattern of pulses detected by said decoding means which is the same for each terminal unit and appearing on said loop and effective only when said terminal unit has data to send,

(d) said logic means when so distinctively rendered responsive enabling the application to said pulse output means of data of any desired variable length as required by said terminal unit for transmission over said loop in said predetermined direction to said master unit,

(e) the first upstream terminal unit having data to send acting to inhibit the receipt at any downstream terminal unit of a pattern of pulses otherwise effective to control said downstream terminal unit to transmit data.

7. Apparatus for the transmission of data in pulse form a plurality of which constitute a character from a plurality of terminal units to a master unit comprising in combination:

pulse input and pulse output means for each said terminal unit and for said master unit,

a single series loop connecting said terminal units in series along said loop between the pulse output means and pulse input means of said master unit and over which pulses originating either with said master unit or any terminal unit are transmitted always in the same predetermined direction, a pulse responsive means connected to said pulse input means at each said terminal unit and responsive to each coded character appearing at said pulse input means,

decoding means at each terminal unit distinctively responsive to different characters received by said pulse responsive means,

logic means at each said terminal unit controlled by said decoding means and responsive to a distinctive synchronizing character transmitted from said master unit for operatively interconnecting said pulse input means to said pulse output means to thereby cause any character transmitted over said loop to be substantially simultaneously received at each terminal unit and be received by said pulse responsive means,

said logic means at said plurality of terminal units being further responsive to the reception of at least one distinctive character simultaneously at said plurality of terminal units to operatively disconnect said pulse input means from said pulse output means only at each terminal unit having data to send,

and means at said plurality of terminal units and effective only at those terminal units which responded to said distinctive character for controlling said terminal units sequentially to transmit data over said loop and via successive downstream terminal units to said master unit.

8. The apparatus of claim 7 which includes means controlled by said logic means and effective when the associated terminal unit has data to send and receives said first distinctive character for connecting said pulse input means to said pulse output means only through said pulse responsive means, whereby each said pulse responsive means included in a terminal unit having data to send is included in series in said loop.

9. Apparatus for the transmission of data in pulse form a plurality of which constitute a character from a plurality of terminal units to a master unit comprising in commbination:

pulse input and puse output means for each said terminal unit and for said master unit,

a single series loop connecting said terminal units in series along said loop between the pulse output means and input means of the master unit and over which pulses originating either with said master unit or any terminal unit are transmitted always in the same predetermined direction,

first means at each terminal unit responsive to a first distinctive signal appearing at the associated pulse input means only if said terminal unit has data to send for opening said loop,

second means at each terminal unit which is effective in response to a second distinctive signal applied to the associated pulse input means only at those terminal units which responded to said first signal for applying to the associated pulse output means the data stored at said terminal unit for transmission over said series loop to said master unit,

means at each said terminal unit responsive to a signal other than said second signal and received from the immediately preceding upstream station for reclosing the loop at said terminal unit.

10. Apparatus for the transmission of data from a plurality of terminal units to a master unit comprising in combination:

pulse input and pulse output means for each said terminal unit and for said master unit,

a single series loop connecting said terminal units in series along said loop between the pulse output means and input means of the master unit and over which pulses originating either with said master unit or any terminal unit are transmitted always in the same predetermined direction,

and means at said plurality of terminal units, and initiated into operation in response to signals transmitted over said loop, said signals acting serially only on said means at a single terminal unit at a time and only on those terminal units having data to transmit and in the order in which they appear on said loop in said predetermined direction to interrupt said loop and to apply to said loop to said master unit for transmission over the remainder of the loop the data of arbitrary variable length stored at the respective terminal unit.

[11. The apparatus of claim 10 in which said master unit includes means for transmitting data characters over said loop in said predetermined direction to one or more selected terminal units.]

12. Apparatus for the transmission of data characters in pulse form from a plurality of remote stations to a master unit comprising in combination:

at least one terminal unit included in each said remote station;

a modulator and a demodulator for each said remote station and for said master unit;

said at least one terminal unit at each said remote station being connected serially between the demodulator and the modulator for the associated station, said remote stations being serially connected to each other over a series loop extending between the modulator and demodulator of said master unit;

each said terminal unit including,

(a) a register for receiving and storing distinctive characters transmitted over said series loop both from said master unit and from any upstream remote station,

(b) decoding means distinctively responsive to the characters stored in the associated register,

(c) logic means responsive to said decoder,

(d) data storage means for storing data for selective transmission over said loop to said master unit, and,

(e) pulse input means and pulse output means,

(f) a gate connected between said pulse input and pulse output means,

said logic means being at times operable in response to distinctive pulse signals appearing on said loop and in the associated register to open said gate and thereby prevent the application of pulses to said pulse output means from said pulse input means and to connect said data storage means to said pulse output means for transmission of said data over said loop to said master unit,

said logic means being effective upon the opening of said loop at the first upstream terminal unit having data to send for inhibiting the receipt at any downstream terminal unit of an order character otherwise effective to control said downstream terminal unit to transmit data.

13. Apparatus for the transmission of data from a plurality of remote stations to a master unit comprising in combination:

at least one terminal unit included in each said remote station;

a modulator and a demodulator for each said remote station and for said master unit;

said at least one terminal unit at each said remote station being connected serially between the demodulator and the modulator for the associated station, said remote stations being serially connected to each other over a series loop extending between the modulator and demodulator of said master unit;

each said terminal unit including,

(a) a register for receiving and storing distinctive characters transmitted over said series loop both from said master unit and from any upstream remote station,

(b) decoding means distinctively responsive to the characters stored in the associated register,

(c) logic means controlled by said decoder,

(d) data storage means for storing data for selective transmission over said loop to said master unit, and

(e) pulse input means and pulse output means,

said logic means being normally in a first condition wherein said logic means is effective to connect said pulse input means to said pulse output means,

said logic means being further operable to a second condition in response to a distinctive signal appearing in said register at a time when said data storage means is storing data for transmission to said master unit to disconnect said pulse output means from said pulse input means to thereby open said loop at each terminal unit having data to send,

said logic means when in its said second condition being additionally responsive to a particular distinctive signal appearing in said register to operatively connect said data storage means to said [data] *pulse* output means only at the first upstream terminal unit which had data to send, said logic means when in its said second condition being responsive to any further signal in said register other than said particular distinctive signal to restore said logic means to its said first condition and thereby connect said [data] *pulse* input means to said [data] *pulse* output means at all terminal units other than said first upstream terminal unit.

14. The apparatus of claim **13** wherein said logic means when in its said second condition connects said pulse input means to said pulse output means only through said register.

15. Apparatus for the transmission of data characters in the form of pulses from a plurality of remote stations to a master unit comprising in combination:

at least one terminal unit included in each said remote station;

a modulator and a demodulator for each said remote station and for said master unit;

said at least one terminal unit at each said remote station being connected serially between the demodulator and the modulator for the associated station, said remote stations being serially connected to each other over a series loop extending between the modulator and demodulator of said master unit and over which pulses originating with either said master unit or any terminal unit are transmitted always in the same predetermined direction; each said terminal unit including means normally effective to maintain said series loop closed at said terminal unit but being responsive to

distinctive signals appearing on said loop to interrupt said loop and to apply to said loop pulses of any number as determined solely by said terminal unit for transmission to said master unit.

16. Apparatus for the transmission of data from a plurality of terminal units to a master unit comprising in combination:

pulse input and pulse output means for each said terminal unit and for said master unit,

a single series loop connecting said terminal units in series along said loop between the pulse output means and pulse input means of said master unit and over which pulses originating either with said master unit or any terminal unit are transmitted always in the same predetermined direction,

first means at each said terminal unit responsive to a distinctive first signal appearing at the associated pulse input means only when said terminal unit has data to send for operatively disconnecting said pulse input means from said pulse output means,

second means at each said terminal unit responsive to a distinctive second signal appearing at the associated pulse input means only provided that said first means has been rendered responsive to said distinctive first signal for applying the data stored at said terminal unit to the associated pulse output means,

each said terminal unit at the termination of its transmission of data over said loop to said master unit transmitting both said distinctive first and second characters and also operatively reconnecting the associated pulse input and pulse output means,

whereby the transmission of said distinctive first and second characters by said master unit causes the transmission of data only from the first upstream station in said loop having data to send and the subsequent transmission of said distinctive first and second characters by said first upstream station causes the transmission of data only from the next station downstream of said first upstream station which had data to send, and so on, until every station in said loop which had data to send has transmitted its data in turn over said loop to said master unit.

17. A method of data transmission between a master unit and a plurality of remote terminal units each having a pulse input and a pulse output which are normally operatively interconnected and with said terminal units being serially connected in a loop extending from the pulse output to the pulse input of the master unit, said method comprising the steps of:

transmitting from said master unit to each said terminal unit over said loop a first distinctive signal,

decoding said first signal at each terminal unit and in response thereto interrupting said loop provided only that said terminal unit has data to transmit,

transmitting from said master unit over said now interrupted loop a second distinctive signal,

decoding said second signal only at the first upstream terminal unit which had data to transmit and in response thereto applying to said loop the data stored at said terminal unit for transmission to said master unit,

transmitting again at the termination of the data transmission from said first upstream terminal unit said first and second distinctive signals.

18. A method of data transmission between a master unit and a plurality of terminal units each having a pulse input and a pulse output and with said terminal units being serially interconnected in a loop extending from the pulse output to the pulse input of the master unit, said method comprising the steps of:

at times transmitting at least one distinctive order character over said loop from said pulse output of said master unit to initiate a cycle of data transmission from all those terminal units, one at a time, which are ready to transmit data to said master unit,

in response to said at least one distinctive order character interrupting said loop only at a terminal unit having data to send and applying to said loop at said terminal unit the stored data of arbitrary variable length for transmission to said master unit,

said transmission of data from said master unit and also from each said terminal unit always occurring in the same direction around said loop.

19. A method of synchronous data transmission over a loop comprising a serial transmission channel extending between the output and input of a master unit and having a plurality of terminal units each including an input and an output which are serially connected in said loop, each terminal unit being selectively operable between two distinctive conditions in a first of which its input is effectively connected to its output and in the second of which the input is not connected to the output and in response to signals appearing at the input said terminal unit is operated between its two said distinctive conditions, the transmission of data from the master unit or from any terminal unit proceeding always in the same direction over the loop, the method steps comprising:

controlling each terminal unit to its said first condition in response to the receipt of a first distinctive signal at its input as monitored by the associated register,

controlling each terminal unit having data to send to its second distinctive condition in response to the reception of a second distinctive signal at its input as monitored by the associated register,

and thereafter transmitting stored data over said loop only from the first upstream terminal unit which is now in its second distinctive condition following the resetting of each downstream terminal unit which was operated to its said second distinctive condition back to its said first distinctive condition.

20. The method of claim **19** wherein said first upstream terminal unit at the conclusion of its data transmission transmits over the loop a further second distinctive character to thereby control all downstream terminal units having data to transmit to their second distinctive condition, and thereafter transmitting over the loop the stored data from the first downstream terminal unit of the group of terminal units which was operated to its second distinctive condition in response to said further second distinctive character while concurrently restoring to its first distinctive condition each remaining terminal unit.

21. Apparatus for the transmission of data from a plurality of terminal units to a master unit comprising in combination:

pulse input and pulse output means for each said terminal unit and for said master unit,

a single series loop connecting said terminal units in series along said loop between the pulse output means and input means of the master unit and over which pulses originating either with said master unit or any terminal unit are transmitted always in the same predetermined direction,

control means at any one of said terminal units which is rendered effective, only when said terminal unit has data to transmit, in response to a single occurrence on said loop of distinctive signals applied to the associated pulse input means, said control means when so rendered effective opening at least at the first upstream terminal unit having data to transmit a normally closed by-pass means which when in its normally closed condition permits the retransmission from the output means of said one terminal unit of any signals appearing at its input means, and further means effective at the first upstream terminal unit with data to transmit and when its said by-pass means is open for applying to said output means the data of whatever length is then stored at said one terminal unit,

and additional means at each said terminal unit for applying to said pulse output means at the conclusion of data transmission said distinctive signals to initiate the transmission of data from the next downstream terminal unit having data to transmit,

whereby the single occurrence of a transmission of said distinctive signals from said master unit initiates a cycle of data transmission sequentially from all terminal units having data to transmit.

22. A method of data transmission between a master unit and a plurality of terminal units each having a pulse input and a pulse output and with said terminal units being serially interconnected in a loop extending from the pulse output to the pulse input of the master unit, said method comprising the steps of:

transmitting onto said loop from said master unit distinctive signals when it is desired to initiate a cycle of data transmission from all terminal units having data to transmit,

responding to each occurrence of said distinctive signals at least at the first upstream terminal unit ready to transmit data to the master unit by opening a normally closed by-pass which is effective only when closed to repeat at the output means of said terminal unit the signals received from said loop at its input means, applying to the output means the data of whatever length is then stored at said terminal unit while said by-pass is open,

transmitting from the output means of said terminal unit at the conclusion of data transmission said distinctive signals to initiate the transmission of data from the next downstream station ready to transmit data to the master unit,

and initiating a cycle of data transmission from all terminal units having data to transmit by a single transmission from said master unit of said distinctive signals.

23. A method for controlling the transmission of data over a common loop communications channel from a plurality of terminal units each connected to said loop comprising the steps of:

transmitting on said loop a distinctive signal to initiate data transmission over said loop from those terminal units having data to send;

responding at a terminal unit on said loop which has data to send and which is responsive to said distinctive signal by interrupting said loop at said terminal unit, by placing on said loop the data at said terminal unit, and by thereafter placing on said loop at said terminal unit said distinctive signal at the termination of the *transmission of the stored data for similarly effecting the transmission of data from a further terminal unit having data to send,*

said transmission of said distinctive signal and of said data always occurring in the same direction around said loop.

24. The method of claim 23 in which said distinctive signal and said data on said loop comprise electrical signals.

25. A method for transmitting data from a plurality of terminal units over a unidirectional communications loop, the method comprising for each individual terminal unit of said plurality the steps of:

- *a. transmitting data from the terminal unit over the loop in response to reception at that terminal unit, at a time when it has data to send, of an order signal propagated over the loop;*
- *b. blocking at the terminal unit onward passage of the received order signal over the loop when the terminal unit has data to send; and*
- *c. repropagating said order signal over the loop for onward passage from the terminal unit following each transmission of data from that terminal unit.*

26. In a method of data transmission from terminal units that are connected in a communications loop, in which an order signal for initiating such transmission is transmitted on the loop to reach the terminal units sequentially, and in which each terminal unit having data to transmit sends data over the loop in response to the condition in which the order signal reaches that terminal unit over the loop, the improvement wherein the response of the terminal unit to said condition includes blocking the onward transmission of the order signal past that terminal unit in the loop, and transmitting from the terminal unit itself said order signal onward in the loop following transmission of the data from that terminal unit.

27. The method of claim 25 or claim 26 wherein said order signal is a predetermined signal combination having first and second parts and said blocking of the order signal in its onward transmission is effected by altering said second part.

28. Apparatus for transmitting data over a communications loop, said apparatus comprising a plurality of terminal means for individually propagating data and order signals over the loop, and a unidirectional signal-transmission path, said path intercoupling the terminal means serially in the loop for communicating data and order signals onward in the loop between the successive ones of the terminal means, each said terminal means comprising:

- *(i) first means coupled in said path, said first means being selectively operable for interrupting onward transmission in the loop from the respective terminal means of signals received from said path at that terminal means;*
- *(ii) second means coupled to said path, said second means being conditioned only when said terminal means has data to send over the loop to provide a distinctive response upon reception from said path of a predetermined order signal;*
- *(iii) third means for operating said first means in response to said distinctive response of said second means, to interrupt onward transmission beyond the respective terminal means of said predetermined order signal;*
- *(iv) fourth means responsive to said distinctive response of said second means to propagate in said path for transmission onwardly in the loop from the respective terminal means, signals in accordance with said data at that means; and*
- *(v) fifth means operative to propagate said order signal in said path for transmission onwardly in the loop from the respective terminal means following each transmission for that terminal means of the data signals.*

29. Apparatus for the transmission of data over a unidirectional communications loop, said apparatus comprising a plurality of terminal units intercoupled with one another serially in said loop for inputting data to the loop, each said terminal unit comprising first means providing as part of said loop a selectively-blockable signal pathway through the respective terminal unit for signals received at that terminal unit in the loop, said first means being conditioned in the event that the respective terminal unit has data to send to respond to reception at the respective terminal unit of a predetermined order signal received at that terminal unit in the loop to block said pathway to onward transmission of said order signal, second means responsive to said order signal when the respective terminal unit has data to send for propagating in said loop from the respective terminal unit signals in accordance with the data for transmission onward from that terminal unit in the loop, and third means for propagating said order signal for transmission onward in the loop from the respective terminal unit following transmission from that terminal unit of said data signals.

30. Control apparatus for operatively connecting each of a plurality of terminal units to a loop-type unidirectional communications channel, said apparatus at each terminal unit comprising:

first means for detecting signals on said loop, said first means being rendered distinctively responsive to an order signal on said loop only in the instance when the terminal unit has data to send,

second means normally permitting the onward passage of signals on said loop but being effective to block the onward transmission of said order signal in the event that said first means has been rendered distinctively responsive as aforesaid,

third means also responsive to said order signal on said loop for transmitting signals onto said loop in accordance with said data,

and fourth means effective after the transmission of said data signals to transmit said order signal onto said loop.

31. Control apparatus for coupling data and order signals between a terminal means and a unidirectional communications loop to which a plurality of such terminal means are connected; said control apparatus comprising:

first means for providing a normally unblocked but selectively blockable signal pathway through said control appparatus for signals on said loop received at such control apparatus, said signal pathway forming a part of said loop;

second means effective when the associated terminal means has data to send to respond to the reception at such control apparatus of a predetermined order signal transmitted on the loop to operate said first means so as to block said pathway to onward transmission of said order signal;

third means also responsive to said order signal, and effective when said first means has been operated by said second means to block said pathway, to enable the propagation over said loop of signals in accordance with data inputted to the associated terminal means;

and fourth means rendered effective following each transmission of said data signals for propagating said order signal over said loop.

32. The data transmission apparatus of claim 31 wherein said predetermined order signal comprises first and second parts;

said second means being responsive to said first part of said order signal to operate said first means to block said pathway to said second part, and

said third means being rendered effective by said second part of said order signal to propagate said data signals.

33. The data transmission apparatus of claim 32 wherein said third means comprises means responsive to the condition in which said second part of said order signal is not received immediately following the reception of said first part of said order signal, to operate said first means to its normal condition to unblock said signal pathway.

34. The apparatus of claim 10 in which said master unit includes means for transmitting data characters over said loop in said predetermined direction to one or more selected terminal units.

* * * * *

2. Bibliography

Andrews, Don W., Schultz, Gary D. "A Token-Ring Architecture For Local Area Networks: An Update." Proceedings, Computer Networks, COMPCON Fall 82, pp. 615-624. Silver Spring: IEEE Computer Society Press, 1982.

Bartee, Thomas C., Data Communications, Networks, and Systems. Indianapolis: Howard W. Sams & Co., Inc., 1985.

Bux, Werner, Felix H. Closs, Kuemmerle, Karl, Keller, Heinz J., and Mueller, Hans R., "Architecture And Design Of A Reliable Token-Ring Network." Local Network Equipment (Tutorial), pp. 166-175. H.A. Freeman and K.J. Thurber, Editors. Silver Spring: IEEE Computer Society Press, 1985.

Cheong, V.E., and Hirschheim, R.A., Local Area Networks -- Issues, Products, And Development. Great Britain: John Wiley & Sons Ltd., 1983.

Clark, David D., Pogran, Kenneth T., and David P. Reed. "An Introduction To Local Area Networks." Local Computer Networks, Second Edition (Tutorial), pp. 16-36. Kenneth J. Thurber and Harvey A. Freeman, Editors. Silver Spring: IEEE Computer Society Press, 1981.

Cotton, Ira W. "Technologies For Local Area Computer Networks." Proceedings Of The Local Area Communications Network Symposium, pp. 25-44. Boston, 1979.

Dixon, R.C. "Ring Network Topology For Local Data Communications." Proceedings, Computer Networks, COMPCON Fall 1982, pp. 591-605. Silver Spring: IEEE Computer Society Press, 1982.

Dowros, Bryan P., Hamilton, James A., Nelson, David A., and Stumpf, Bernard L. "The Architecture Of An Integrated Local Network." Local Network Equipment, (Tutorial), pp. 145-159. H.A. Freeman and K.J. Thurber, Editors. Silver Spring, MD: IEEE Computer Society Press, 1985.

Ethernet-Type Local Networks Report (The). Minneapolis: Architecture Technology Corporation, 1985.

Farber, David J. "A Ring Network," Datamation, February 1975, pp. 44-46.

Flint, David C. The Data Ring Main. Chicester: John Wiley & Sons, 1983.

Franta, W.R., and Chlamtac, I. Local Networks. Lexington, MA: Lexington Books, 1981.

Gordon, R.L., Farr, W.W., and Levine, P. "Ringnet: A Packet

Switched Local Network With Decentralized Control." Proceedings, 4th Conference On Local Computer Networks, pp. 13-19. New York: IEEE Computer Society, 1979.

Hopper, Andrew, and Williamson, Robin C. "Design And Use Of An Integrated Cambridge Ring." Local Network Equipment (Tutorial), pp. 195-204. H.A. Freeman And K.J. Thurber, Editors. Silver Spring: IEEE Computer Society Press, 1985.

Huber, Daniel E., Steinlin, Walter, and Wild, Peter J. "SILK: An Implementation Of A Buffer Insertion Ring." Local Network Equipment (Tutorial), pp. 186-194. H.A. Freeman and K.J. Thurber, Editors. Silver Spring: IEEE Computer Society Press, 1985.

Janson, P., Meister, B., and Svobodova, L. "File Transfer In Local-Area Networks: A Performance Study." Proceedings, 5th International Conference On Distributed Computing Systems, pp. 338-349. Silver Spring: IEEE Computer Society Press, 1985.

Kuemmerle, K., and Reiser, M. "Local-Area Communication Networks -- An Overview." Local Network Equipment (Tutorial), pp. 3-24. H.A. Freeman and K.J. Thurber, Editors. Silver Spring: IEEE Computer Society Press, 1985.

Local Network Vendors' Strategies Report (The). Minneapolis: Architecture Technology Corporation, 1985.

LOCALNetter Designer's Handbook (The), 4th Edition. Minneapolis: Architecture Technology Corporation, 1985.

Markov, J.D., Strobe, N.C. "Token Ring Local Area Networks: A Perspective." Proceedings, Computer Networks, COMPCON Fall 1982, pp. 606-614. Silver Spring: IEEE Computer Society Press, 1982.

Meijer, Anton, and Peeters, Paul. Computer Network Architectures. Rockville, MD: Computer Science Press, 1983.

Micro To Mainframe Communications Directory (The), 2nd Edition. Minneapolis: Architecture Technology Corporation, 1985.

Needham, R.M., and Herbert, A.J. The Cambridge Distributed Computing System. London: Addison-Wesley Publishing Company, 1982.

Personal Computer Local Networks Report (The). Minneapolis: Architecture Technology Corporation, 1985.

Salwen, Howard C. "In Praise Of Ring Architecture For Local Area Networks." Local Network Equipment (Tutorial) pp. 160-165. H.A. Freeman and K.J. Thurber, Editors. Silver Spring: IEEE Computer Society Press, 1985.

Saltzer, Jerome H., and Pogran, Kenneth T. "A Star-Shaped Ring Network With High Maintainability." Proceedings Of The LACN Symposium, pp. 179-189. Boston, 1979.

Stallings, William. Local Networks, An Introduction. New York: Macmillan Publishing Company, 1984.

Stallings, William. Tutorial: Local Network Technology. Silver Spring: IEEE Computer Society Press, 1983.

Tanenbaum, Andrew S. Computer Networks. Englewood Cliffs, NJ: Prentice-Hall, Inc., 1981.

Token Perspectives Newsletter (The). Minneapolis: Architecture Technology Corporation, Vol. 1-4, 1983-1986.

Weitzman, Cay. Distributed Micro/Minicomputer Systems. Englewood Cliffs, NJ: Prentice-Hall, Inc., 1980.

Wilkes, M.V., and Wheeler, D.J. "The Cambridge Digital Communications Ring." Proceedings Of The LACN Symposium, pp. 47-61. Norman B. Meisner and Robert Rosenthal, Editors. Boston, 1979.

3. Ring-Based Local Networks Manufacturers

Apollo Computer, Inc.
15 Elizabeth Drive
Chelmsford, MA 01824

Hasler A.G., Ltd.
Belpstr. 23
CH-3000 Berne 14
SWITZERLAND

IBM Corporation
1113 Westchester Avenue
White Plains, NY 10604

INCOMNET, Inc.
5795 Lindero Canyon Rd.
West Lake Village, CA 91362

Logica VTS Ltd.
86 Newman St.
W1A 4SE, London
ENGLAND

Prime Computer Inc.
Prime Park
Natick, MA 01760

Proteon
4 Tech Circle
Natick, MA 01760

Racal Data Comms Group
Richmond Ct.
309 Fleet Rd.
Fleet, Hants
UNITED KINGDOM

Racal-Milgo
8600 N.W. 41st Street
Miami, FL 33166

Ungermann-Bass, Inc.
2560 Mission College Blvd.
Santa Clara, CA 95050

Xionics
68 Oxford Street
W1N 9LA, London
ENGLAND

Ztel
181 Ballardvale St.
Wilmington, MA 01887