A token-ring network for local data communications

by R. C. Dixon
N. C. Strole
J. D. Markov

Technical innovations such as large-scale integrated circuit technology and distributed operating systems have respectively reduced the cost of computing and provided a basis for large networks within the confines of a single building or cluster of buildings in close proximity to one another. Local area networks can provide a systematic approach for interconnecting personal workstations, control units, and central processing units, thereby providing a means for these machines to pass information from one to the other. This paper describes a local area network based on the fundamental concepts of a token-ring. Two main ideas are presented. The first idea concerns the physical topology of the wiring network and its star-ring organization. Next, the logical data flows are overlaid on the physical network to provide control procedures for exchanging data through the network. The resulting system has unique features that produce a local area network with good performance and reliability characteristics.

The present trend toward distributed data processing, coupled with the increased emphasis on office automation and the widespread use of automated office terminals, displays, and computers, has created a demand for a more comprehensive approach to terminal interconnection systems. Personal workstations can increase overall productivity by allowing the individual office workers, both clerical and professional, to share access to host systems, common data bases, peripheral printers, and remote computer networks. The workstation itself may range from a low-cost keyboard/display device to a small computer. The industry response to this need has resulted in several proposals and offerings of communication systems specifically designed for local data communications.

A local area network (LAN) can be defined as an information transport system for data transfer among office system terminals and peripherals, clustered controllers, or host systems, via an interconnecting medium within the bounds of a single office building, building complex, or campus. The geographical constraints eliminate the need to use common-carrier facilities, thus increasing the data transfer capacity of the LAN by allowing economical data transmission rates of many millions of bits per second. Today's network transmission technology permits the transfer of large blocks of data at these rates with simple error recovery procedures and control protocols. These transmission rates also permit a large number of data terminals to share the common physical interconnection link with a minimum of interference from one another.

Figure 1 Shared access links

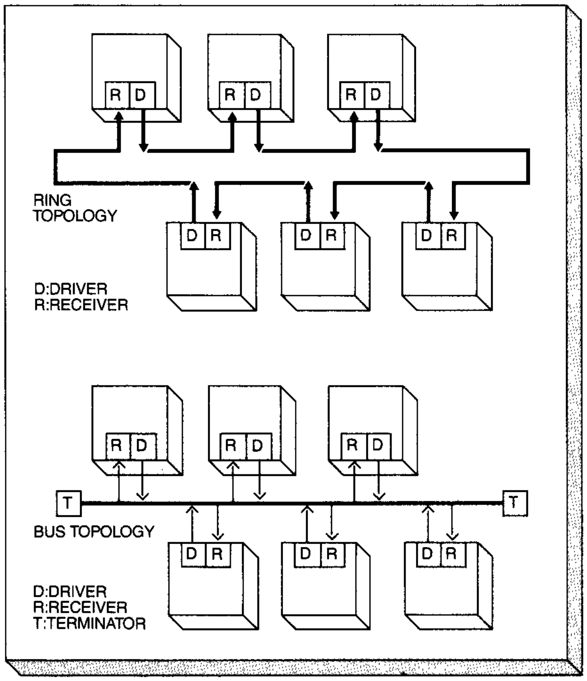

This paper describes an LAN based on a ring topology with token-access control. The concepts presented here resulted from research investigations concerning token-ring LANs conducted by the IBM research group in Zurich, Switzerland, along with efforts at IBM's laboratory in Research Triangle Park, North Carolina. Some of the considerations that go into the design of an LAN are discussed. A functional description of several of the key physical components that comprise a token-ring network, and factors that influence the selection of a transmission cable for the LAN environment, are presented. Various control mechanisms for regulating the flow of data within shared-access links are described, along with a discussion of how data are routed between links that comprise an LAN. These concepts are discussed in the context of a layered communication architecture model, specifically concentrating on the two lower levels, data link control and physical transmission. Particular emphasis is placed on the fault detection and isolation capabilities that are available with the ring topology.

Local area network considerations

A number of requirements must be considered when designing a local network. A primary goal in the design of an LAN is to maintain compatible physical interfaces and control protocols to allow users to intermix different equipment types from multiple manufacturers over the common network. Physical interfaces to the network (i.e., wall outlets) located throughout the building or campus should permit fast, reliable attachment or relocation of these products. The LAN must be flexible to satisfy the many requirements imposed on it and to provide an enduring base for incorporating new technologies over the next 15 to 30 years. In addition, the LAN should allow for communications with geographically remote locations through wide-area networks. Fault detection and isolation capabilities are also essential to ensure reliability and availability of the network. Expansion capability should permit mixed application environments. Digitally encoded data and text information, as well as noncoded data such as facsimile, should be considered. Some applications may include both analog and digital voice data. Analog video transmission may also be required in the local network environment.

An examination of the various applications reveals major technical differences. Examples of these differences are the bandwidth and data rates required and the information integrity requirements. Analog voice requires four kilohertz of bandwidth, whereas standard video channels require six megahertz. Digital voice is normally 64 kilobits per second, whereas many keyboard/display terminals operate at one or two megabits per second. Digitally encoded data normally have sequence numbers and error checking, with automatic retransmission when errors occur. However, typical facsimile, voice, and video applications do not require sequence numbers or error checking. It is expected that these systems will, to some degree, find application in the office and pass information over LANs. These different applications have several common aspects. For example, the cable routing, the concentration points, and the cables themselves can take into account certain aspects that are common among the different applications.

To satisfy many of these requirements, network structures have been developed that allow for the attachment of many nodes to a common physical link.[1] (Node is used in this paper as a generic term to refer to any machine that attaches to and uses the LAN.) The connectivity of such shared-access links has two aspects: the physical connections achieved through propagation of electrical signals and the logical connections established by the link protocols. Physically, the shared-access link permits any node to communicate directly with any other node attached to the link. This connectivity allows a single node to transfer information to all nodes, or to a subset of all nodes, that are physically attached to the link. Thus, the physical topology can be described as "any-to-any." Each node attached to the shared-access link can communicate on a one-to-one basis with any other node within the network. The logical data flow can therefore be described as "point-to-point," whereby each node may simultaneously support multiple logical links. Data are assumed to be transferred as packets consisting of a variable-length information field preceded by appropriate addressing field(s).

Two general classes of physical link topologies are employed for shared-access links: the ring and the bus.

Two general classes of physical link topologies are employed for shared-access links: the ring and the bus (Figure 1). The ring consists of a series of nodes that are connected by unidirectional transmission links to form a single closed path. Information signals on the ring pass from node to node and are regenerated as they pass through each node. The bus provides a bidirectional transmission facility to which all nodes are attached. The transmission facility may consist of a single cable, with either single or multiple channels, or multiple cables, depending on the technology employed. Information signals propagate away from the originating node in both directions to the terminated ends of the bus.

The following sections discuss the interconnection of nodes in a ring configuration with a token-access control mechanism. An examination of the total spectrum of characteristics of token-ring networks, rather than any one key attribute, provides a strong basis for considering the token-ring approach.

Ring network configuration

Two principal objectives are normally considered when configuring a local network. First, the total length of cable should be minimized to achieve low system cost and to limit overall transmission distances. Second, concentration points should be provided in the network to facilitate cable installation and network configuration and maintenance.

These two objectives are not independent and must be balanced to reach a single solution for a ring network. In order to keep cable lengths to a minimum, a pure serial interconnection of nodes would be appropriate. However, the costs associated with maintaining and reconfiguring the network would be intolerable. In the other case, a pure star-radial cabling scheme, where all nodes are cabled to a single concentration point, provides a configuration that overcomes these problems. For this case, however, the total length of cable required is obviously quite extensive. A hybrid solution that combines the key advantages of both approaches is desirable.

The star-ring hierarchical topology described below combines many of the desirable features of both the star and ring configurations. Major components of this system are shown in Figure 2. They are the nodes, wiring concentrators, bridges, and gateways. These components are organized to form the composite local network. Personal workstations and other nodes each attach to the network through a communications adapter. Transmission media, such as shielded twisted-pair wire or optical fibers, carry the signals.

Wiring concentrators provide a key element for structuring a system to provide both flexibility and reliability.

Figure 2 Wiring topology and components of a ring network

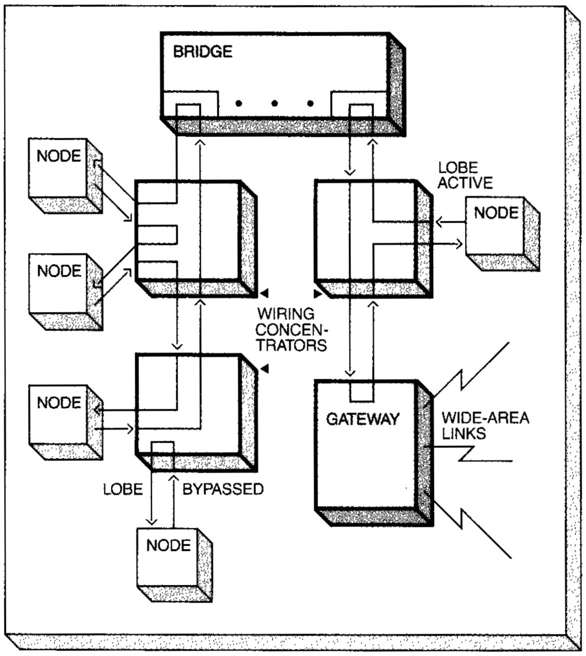

Wiring concentrators. Wiring concentrators provide a key element for structuring a system to provide both flexibility and reliability.[2] Flexibility in system configuration is provided via a star-ring hierarchical wiring scheme that is implemented by connecting distributed nodes to wiring concentrators. Each node is attached to a single lobe of a wiring concentrator, with the lobes being physically connected within the concentrators to form a serial link (Figure 2). The wiring concentrators are then interconnected in a serial fashion to complete the ring. The wiring is segmented at the wiring concentrators, thus providing the ability to preplan wiring, facilitate workstation moves, and intermix transmission media. For example, shielded twisted-pair wire can be used to interconnect the wiring concentrators to the workstations, whereas optical fibers can be used for transmission links between wiring concentrators.

Wiring concentrators also provide points within the ring network that facilitate reconfiguration and maintenance of the network to enhance reliability.[3] Switching elements exist within the wiring concentrators for bypassing each lobe selectively. This lobe-bypass function may be implemented with manually controlled switches or automatically controlled circuits, or may be controlled remotely through the receipt of appropriate network management commands.

Bridge operation. The electrical signals in a ring configuration are regenerated as the signal passes through an active node. Thus, the capability exists for a ring network to span a local area of considerable size since the electrical signal generated at one extreme point in the network does not propagate to another extreme point in the network, but only to the next active node. It could be theoretically possible to construct a single ring in which an unlimited number of nodes are linked through wiring concentrators. However, studies have indicated that systematic low-frequency jitter caused by a very large number of repeaters in the link can lead to a loss of bit synchronization.[4] Also, such a configuration would be unwieldy from a network management viewpoint if several hundred nodes were attached to a single ring. The data capacity of a ring must also be considered when configuring a ring system. Thus, a typical configuration will normally allow for the connection of 100 to 200 nodes to a single ring.

Gateways extend the access of nodes within the LAN.

Multiple rings may be employed when the capacity of a single ring is exceeded or when the attached nodes are spread over a large area. A high-speed digital switching mechanism known as a bridge can link multiple rings to provide a logical routing that is transparent to the attached nodes. Individual rings that are connected through the bridge operate autonomously and, therefore, may stand alone. This capability allows several independent rings to be installed and then interconnected to form a larger LAN. The network can be further expanded by interconnecting multiple bridges. An additional function of the bridge can be to perform transmission speed changes from one ring to another. A very large installation could then be supported by many rings, all interconnected via a hierarchical network structure with a separate high-speed backbone link between bridges. The backbone itself may be a high-speed ring (Figure 3A) or it may be a bus network, such as a channel within a broadband cable television system (Figure 3B). The ring and the backbone must share a common addressing structure to achieve the routing of data between them. The address of a given node can be structured to designate the ring to which the node is attached, thereby facilitating the routing of packets through bridges.[5]

Gateway operation. Gateways extend the access of nodes within the LAN by providing long-distance communication through wide-area networks. Wide-area networks generally operate in the kilobits-per-second transmission range, with a variety of protocols and standards governing their interfaces over private or commercial satellite links, packet-switched networks, or terrestrial links. A gateway can provide the speed and protocol conversions that are required to interface the LAN to these various transmission facilities. Thus, nationwide or global communication among multiple LANs can be established. Such interconnections allow for a more economic allocation of computing resources within major locations of an organization, while still allowing access to these resources over a geographically dispersed area.

Optical fibers will play an important role in future LANs.

Data packets are routed unchanged through a single LAN. This implies a single address space for the network that normally must be extended to access destinations outside a given LAN. One function of a gateway is to provide the proper address translation between the interconnected networks. Consequently, routing of a data packet through a gateway results in disassembly of the data packet and reassembly into a new data packet with the appropriate address translation.

Physical wiring for a ring network. The wiring of buildings for data transmission on a local level will

Figure 3 Multiple-bridge local area network: (A) Token-ring backbone, (B) Broadband-bus backbone

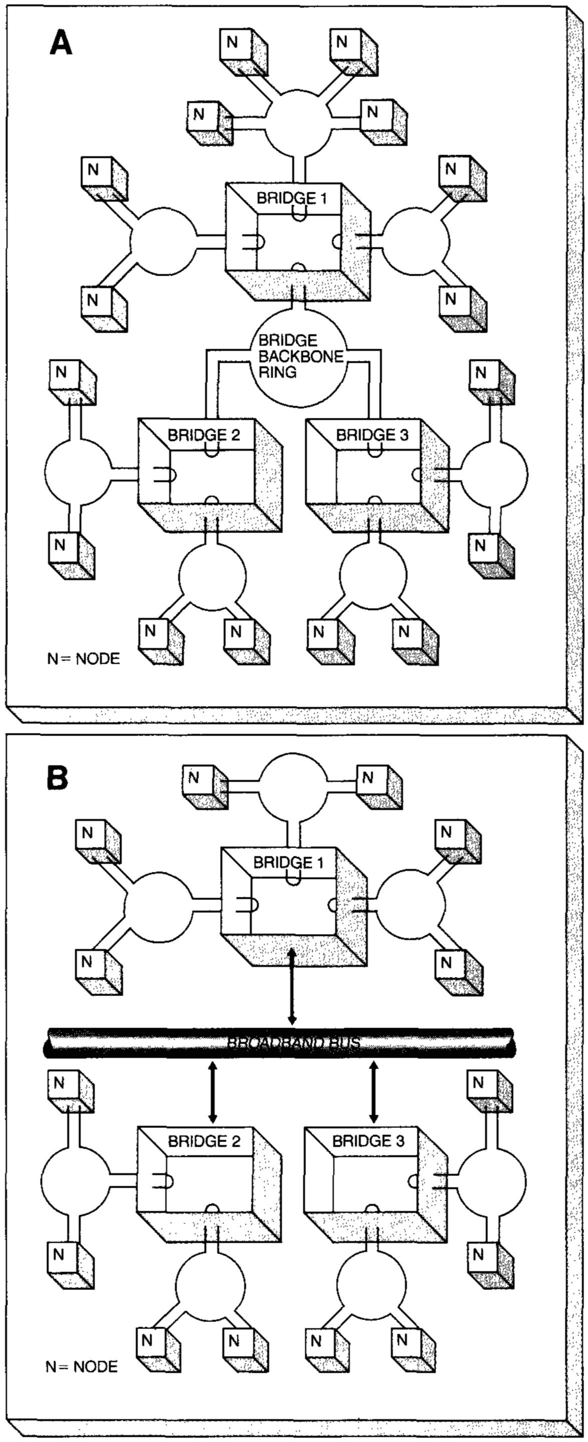

Figure 4 Time division multiplexing

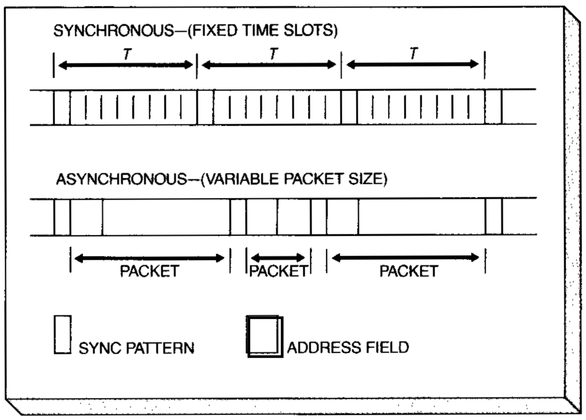

become common practice in the near future. The cost and configuration limitations associated with wiring schemes that require direct cable connections between physical devices (e.g., terminal to controller) can be overcome by a wiring system that can be installed by users planning new facilities, as well as by those with existing facilities. Such a wiring system may also encompass voice and video transmission media, thus allowing each office in a building or complex to share access to voice, video, and data networks.

Several factors must be considered in selecting the best transmission medium for an LAN. A large percentage of a building's total cable runs will be between the wiring concentrators and the office wall outlets. An efficient cost/performance wiring for these rather short distances at data rates of 1 to 10 megabits per second has been shown to be twisted-pair copper wire.[6] There is a trade-off between drive distance and data rate such that four megabits per second is a good design point for this medium. Copper medium is a well-understood technology, with readily available low-cost connectors and splicing techniques, as well as proven installation procedures. Also, a balanced twisted-pair cable with a grounded shield is less susceptible to environmental noise and outward electromagnetic radiation than an unbalanced coaxial cable.

Optical fibers will play an important role in future LANs because of their low susceptibility to electromagnetic interference and their high-speed transmission characteristics. This technology is well-suited for high-speed links with data rates greater than 10 megabits per second, such as the links between bridges. Optical fibers can also be used between buildings, thereby eliminating the additional circuits required for lightning protection on metallic media.

Broadband coaxial cable systems, such as cable television, can be incorporated within the LAN for transmitting both video and digital data signals using standard amplifiers, taps, and connectors. While television systems do not play a major role in the office today, they may in the future. The nature of the cable provides a total bandwidth spectrum of 300 to 400 megahertz, which can be subdivided through frequency division multiplexing into multiple transmission channels. Six-megahertz-bandwidth channels can accommodate standard television signals. Other channels could be used for data transmission between rings by means of bridges (Figure 3B) or for interconnection of LANs via appropriate gateways.

Three basic access control mechanisms have been employed in conjunction with bus and ring topologies.

Ring control protocols

Access control structures. The physical topology of a shared-access link permits simultaneous access to the transmission facility by all attached nodes. Thus it is possible for two or more nodes to transmit at the same time, resulting in interference between the electrical signals. Control mechanisms, both physical and logical, must be introduced to resolve the contention.

Electrical signals can be multiplexed on the transmission facilities by two methods, both of which separate the electromagnetic energy. One, called *frequency division multiplexing* (FDM), separates the energy into different frequency bands; the other, called *time division multiplexing* (TDM), separates it into different periods of time. A single frequency band in an FDM network can also be further subdivided using TDM. Two general categories for the time multiplexing of digital information are synchronous TDM and asynchronous TDM. As shown in Figure 4, synchronous TDM allows each node to use a portion of the bandwidth at periodic intervals. The time allocated is usually relatively short and is determined by a set time interval from a known synchronization pattern. Asynchronous TDM allows each node to transmit for a variable length of time. The data are blocked into packets that contain appropriate addressing information for routing the data.

A logical control mechanism must be employed in conjunction with asynchronous TDM to regulate each node's transmission opportunity. Three basic access control mechanisms have been employed in conjunction with bus and ring topologies.[7] These mechanisms are random access and two versions of controlled access: centralized and distributed. Table 1 lists six link-access mechanisms used for the transmission of packets on a shared-access link.

For a bus topology, one class of random-access method is called Carrier Sense Multiple Access with Collision Detection (CSMA/CD). Prior to initiating transmission, a node senses whether or not a carrier signal is present. If so, another node is transmitting, and origination of a new transmission is delayed until the carrier signal is removed. Once the transmission has started, the node continues to monitor the transmission for a collision (destruction of the transmitted electrical signals), indicating the presence of another transmitting node. Transmission is halted, and the node either waits a random period of time before attempting to transmit again,[8] or waits for a predetermined (node-specific) time interval.[9]

One random-access method for a ring topology is called register insertion. A node on a ring can initiate a transmission whenever an idle state exists on the ring. A data packet received while the node is transmitting is held in a register by the node, and is transmitted following the end of the current data packet. The IBM Local Communications Controller for Series/1 employs a register insertion ring that operates at two megabits per second.[10]

Selective polling is one of the most common centralized-control access mechanisms for a bus topology. This mechanism has been extensively employed for control of multipoint configurations of telecommunications links. For a ring topology with centralized control (often referred to as a loop), the access mechanism is based on a group poll. In both cases, a single control node transmits a unique poll frame to the other nodes that grants them permission to transmit. For a bus topology, each poll frame is addressed to a particular node on the bus. That node will either initiate the transmission of data or return a negative response to the poll. Loop operation uses a single broadcast poll to all nodes. Each node has the opportunity to respond in turn according to its physical ordering on the loop. The IBM 8100 Information System offers a multipurpose loop controller for the attachment of terminal products to an 8100 processor. The Synchronous Data Link Control (SDLC) loop protocols are employed by this system.[11]

Table 1 Access methods for packet transmission

Access Method	Network Topology: Bus	Network Topology: Ring
Random access	CSMA/CD	Register insertion
Centralized control	Selective poll	Group poll
Distributed control	Token (explicit)	Token (implicit)

Figure 5 Architecture reference model

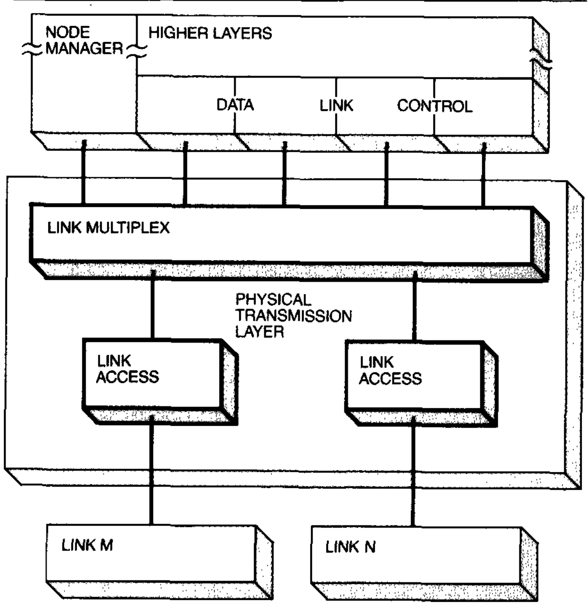

Figure 6 Frame format

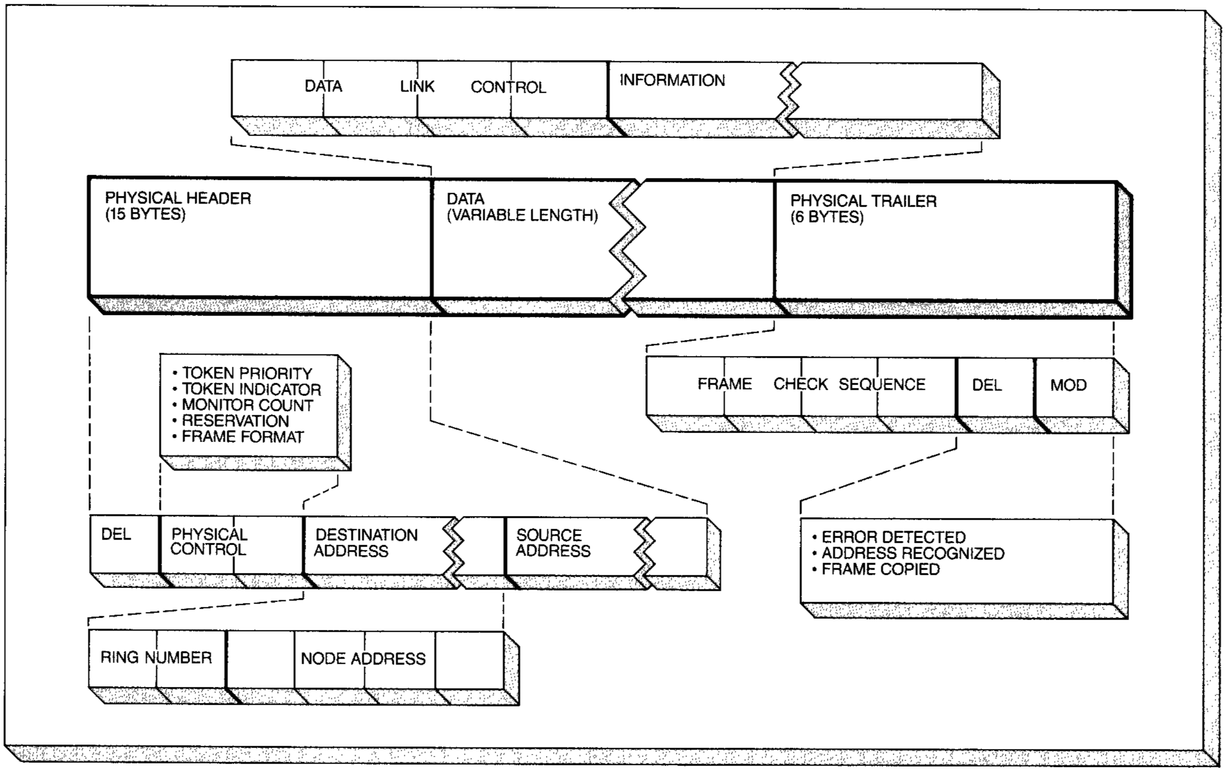

The token control scheme is a type of distributed-control access method. A unique sequence, called a token, is passed from one node to another. The receipt of the token gives permission to the receiving node to initiate a transmission. Upon completion of the transmission, the token is passed to another node. For a ring, this token is passed implicitly (without addressing information) to the next node on the ring.[12] An explicit token (containing specific node addresses) is employed for a bus. This latter approach causes an ordering of the nodes that is often described as a logical ring.

Logical connectivity. Systems Network Architecture (SNA) and the International Organization for Standardization (ISO), as well as other network architectures,[13] employ multiple-layer reference models.[14] Separation of network functions into layers facilitates the description of the protocols as well as the implementation. The two lowest layers of such a model, data link control and the physical transmission layer, are depicted in Figure 5, along with a control element labeled node manager. The logical connectivity associated with a ring network, including the routing of data frames and the control of access to the ring, will be discussed with reference to this model. The higher architectural layers are not addressed in this paper.

Data frame format. Associated with an architecture model is an assumed structure for the information as it appears on the transmission link. The general structure for transmitting the information, called a frame, is shown in Figure 6. The portion of the frame labeled data is, in general, variable in length and contains information processed by functions above the physical transmission layer. Specifically, the first portion of the data field contains information associated with data link control.

The data field is preceded by a physical header, which contains three subfields. The first is a delimiter (DEL) that identifies the start of the frame. Next, a field is defined for controlling the access to the transmission facility (Physical Control Field). This two-byte field includes the token indicator that can exist in one of two states, busy or free. A token priority field, in conjunction with the reservation field, provides different priority levels of access to the ring.[15] The monitor count will be discussed later in connection with the token monitor function, a mechanism that maintains the validity of the token. Finally, the header includes a field that contains the address of the node that originated the information and the address of the node (or nodes) destined to receive the information.

The data field is followed by a physical trailer. The first portion of the trailer contains a frame check sequence (FCS), used for detecting errors that occur during transmission within the second byte of the physical control field, the address fields, or the data field itself. Next, a delimiter is provided to identify the end of the frame. This is followed by a modifier (MOD) field, which, in conjunction with the ending delimiter, is employed for certain physical control functions.[16] Two of these functions, the error-detected and the address-recognized indicators, will be discussed later.

Data link control. Specific functions allocated to data link control (DLC) include the sequencing of frames between two nodes attached to the network, along with appropriate recovery from transmission errors. These functions, as well as those in higher layers, are essentially independent of the transmission facility.

Multiple appearances of DLC may exist within each data processing node. Each DLC appearance is logically associated with one other DLC appearance in another node. Models of four nodes are shown in Figure 7, with examples of some of the logical DLC links that may exist. Each link is identified by the pair of addresses of the nodes at the ends of the logical link. Thus, the point-to-point connectivity discussed earlier in this paper is provided by the DLC layer of the architecture.

Figure 7 Logical link structure

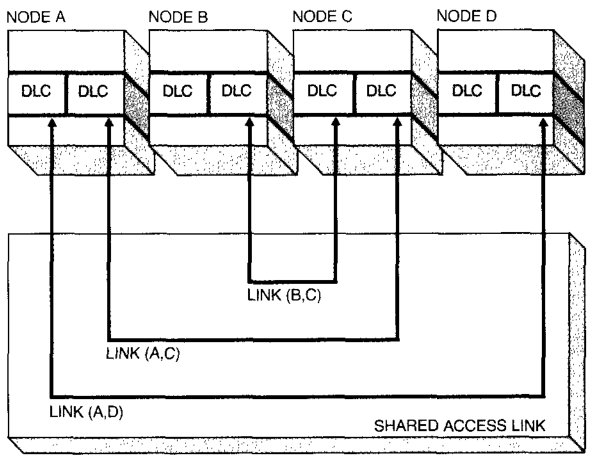

The management of the address fields that appear in the header of a frame is usually considered to be a function of DLC. However, the use of such address fields is strongly dependent on the specific access method being employed. For instance, synchronous TDM does not require any explicit addressing information to be transmitted by the node, and multipoint or loop facilities require only a single node address. Also, validation of the frame check sequence is not a function of DLC. The particular checking algorithm that is employed is dependent on the transmission medium and the specific error environment. Consequently, these functions of addressing and checking are independent of the DLC protocols for sequencing of frames.

Physical transmission. Differences between network topologies are reflected at the physical transmission layer of the architecture model. In particular, the functions associated with the generation and interpretation of the fields within the physical header and physical trailer occur within this layer. To illustrate these functions, the physical transmission layer can be further subdivided into two sublayers as depicted in Figure 5.

- *Link multiplex*—This function provides for routing of the data frames through the network to the appropriate appearance of DLC within a node. For a shared-access link without centralized control, both a destination address and a source

Figure 8 Local network routing

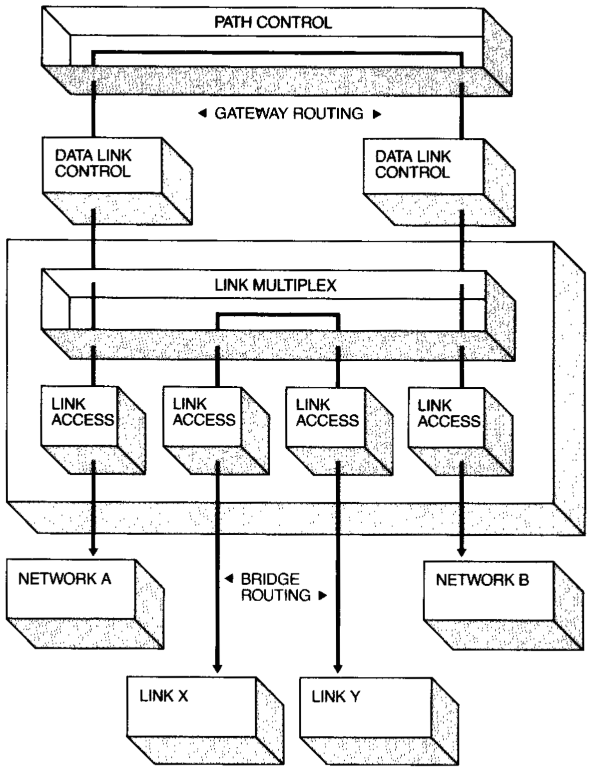

address are required for each data frame. In addition, a frame check sequence for error control is derived prior to transmission. Although the link multiplex sublayer is shown as part of the physical transmission layer, functions associated with these address parameters are shared with the DLC elements of the node.

- *Link access*—This function regulates the transfer of data frames to the transmission facility. Upon receipt of a frame, the link access protocols perform appropriate validity checks and set flags in the modifier field. In addition, the generation and detection of delimiter fields are performed. The delimiters contain a sequence of signal elements that is unique compared to the sequence of signal elements occurring within the body of the frame.

Routing of data frames within a ring network is a function associated with the physical transmission layer. No routing decision by the intermediate ring nodes is required as a frame is propagated from node to node on a single ring. A node must only recognize messages addressed to it.

The routing of frames through a bridge takes place within the link multiplex sublayer (Figure 8). Data frames are routed through the bridge based on the address information contained in the physical transmission header of the frame. In particular, the first two bytes of each address field specify the ring numbers of the destination node and the source node respectively. As stated previously, the two rings run independent link-access mechanisms. Routing in a gateway is performed above DLC in a layer of the architecture called path control by SNA or network control by ISO. This type of routing is typical of the function performed by intermediate nodes within a mesh network. Note that a gateway interfaces to two logical DLC links. In contrast, a bridge provides no DLC function. Therefore, data integrity across the local network must be ensured by an end-to-end protocol at or above the DLC level within the source and destination nodes.

The node manager is responsible for establishing logical links.

Node manager and higher layers. The node manager is responsible for establishing logical links, as well as performing certain network management functions that have been distributed to all nodes. The latter includes not only the token monitor functions and the priority access functions but also those functions required for detection and isolation of ring fault conditions. (The characteristics of the ring network that facilitate ring network management are discussed in the next section.)

Logical link connections are readily established because of the inherent broadcast capability of a shared-access network. Each node in the network has a unique name or identifier. A single frame can then be routed to all rings within the network requesting a specific connection. This frame contains the unique name of the target node as well as other connection parameters. The node recognizing its unique name will respond to the connection request, thereby establishing a logical link between the two nodes. Initialization of DLC procedures can then proceed as if the nodes were connected to a point-to-point transmission facility.

Once a logical link between two nodes has been established, protocols associated with the higher architectural layers may be employed to establish and control logical sessions between application programs and end users. The functions that exist in these layers are defined to be largely independent of the communications network topology. The DLC procedures associated with each logical link ensure that error-free information is passed to the higher-layer functions within each node. Therefore, the currently defined set of protocols within SNA, as well as other higher-layer network protocols, can be employed for data communication in LANs.

Token-access control. The control mechanism for regulating data flow in a ring topology is generally based on the idea that permission to use the communications link is passed sequentially from node to node around the ring. The access function is distributed to all nodes attached to the ring, in contrast to the centrally controlled loop access. An implicit token consists of a unique sequence of information bits that contains an indication of whether the token is free or busy. With the token-access control scheme, the free token circulates on the ring (Figure 9), giving each node in turn an opportunity to transmit data when it receives the token. A node having data to transmit can seize the free token, change the token status to busy, and begin data transmission. The node that initiates a data transfer must remove those data from the ring and issue a free token upon receipt of the physical header, allowing other nodes an opportunity to transmit.

The token-control mechanism of each ring attached to a bridge operates independently of all other attached rings. The bridge appears as a normal node on the ring, except that it must recognize frames whose destinations are other rings and perform the appropriate routing. The primary functions associated with token recognition and data transmission are performed by a ring interface adapter at each node. This adapter handles the basic transmission functions, including frame recognition, token generation, address recognition, error checking, buffering of frames, and link fault detection.

Figure 9 Ring with token-access control protocol

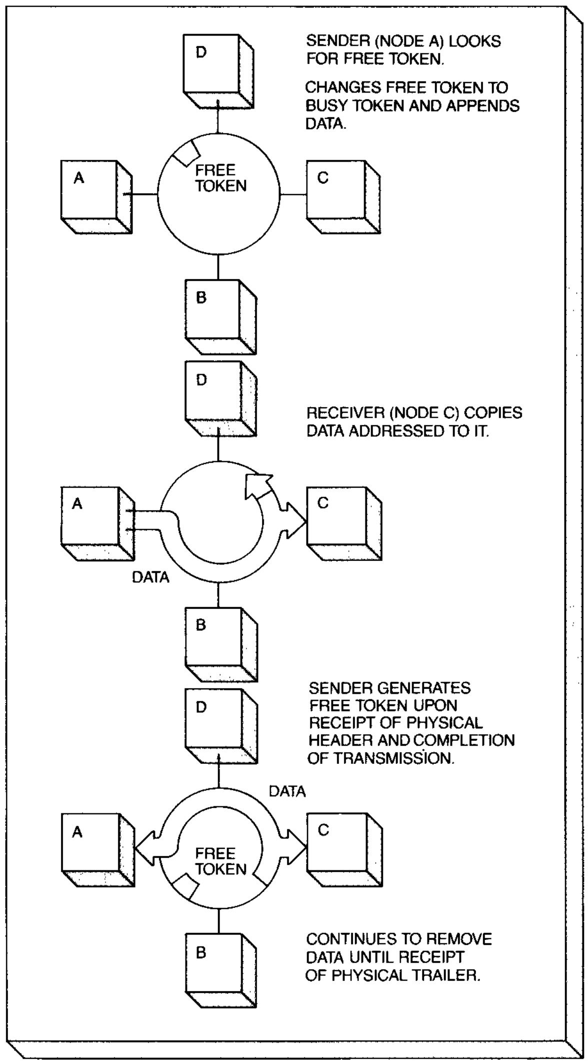

Performance characteristics. The time required to transfer a data frame from one node to another is affected by several factors, including the ring speed, the number of active stations on the ring, the physical length of the ring, the number of bytes in the data frame, and the ring utilization. A plot of the average frame transfer time as a function of the total number of frames per second for frame lengths of 512 and 256 bytes on a four-megabit-per-second ring is shown in Figure 10. These results were based on an analytic model for comparing the perform-

Figure 10 Token-ring frame transfer at four megabits per second

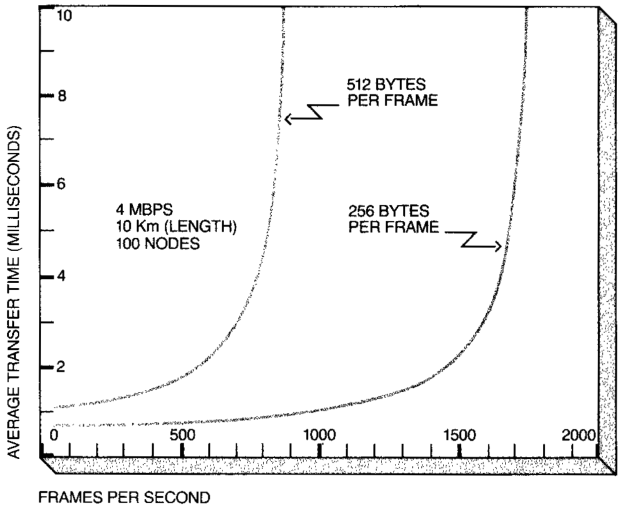

ance of various access control schemes. This analytic model was developed by Bux.[17]

The token-ring network provides extremely good throughput and response time characteristics. For example, a throughput of 1000 frames per second can be obtained with a transfer time of approximately one millisecond per frame. This assumes that the average frame length is 256 bytes on a 10-km, four-megabit-per-second ring with 100 stations active. The throughput may be increased with corresponding increases in transfer times as shown in the graph.

Token monitor function. In addition to the basic access control functions, one ring interface adapter on each ring acts as an active token monitor to perform error recovery if normal token operation is disrupted. This disruption includes the loss of a free token or the continuous circulation of a busy token, both of which prevent further utilization of the ring. Such a monitor function has been defined in conjunction with an experimental ring network being implemented at the IBM Zurich Research Laboratory.[18] It is important to note that this monitor function exists only for token recovery and does not play an active role in the normal exchange of data frames.

The monitor count flag in the physical control field of the frame is employed by the active token monitor to detect the continuous circulation of a busy token. When a busy token is first observed by the active monitor, the monitor count flag is set. Normal ring operation results in the removal of this token by the transmitting node. A failure in this operation will cause the busy token to pass the token monitor a second time. The token monitor, observing the monitor count flag set, will remove the busy token from the ring and issue a free token. For

A star-ring provides an extremely reliable communication link.

detection of a lost token, the active monitor maintains a timer that is reset upon the passage of either a busy token or a free token. If the timer expires, the token monitor will reinitialize the ring with a free token.

The capability to be an active token monitor exists in all active nodes attached to a specific ring. These other nodes maintain a passive monitor status and are prepared to become the new active monitor should a failure in the current active monitor occur. The node with the highest station address will assume this role and become the new monitor.[15]

Synchronous traffic integration. In order to allocate a portion of the available bandwidth for the transfer of synchronous traffic (for example, real-time digitized voice), a mechanism is required that will interrupt the normal flow of asynchronous data traffic. The token-access control scheme allows such an approach. A node called a synchronous bandwidth manager can prevent the issuance of a free token by requesting synchronous priority via the reservation bits of the physical control field of a busy token. This action would be performed at some periodic time interval, *T* (Figure 11). Prior to the release of a free token, the node currently transmitting must check the state of this field. If a synchronous priority request is present, further asynchronous traffic is suspended, thus making the bandwidth available for synchronous traffic. On the average, it will require one frame time to suspend the transfer of asynchronous traffic, thereby allowing the manager node to initiate the transfer of synchronous traffic. If a free token exists on the ring at a given clock interval, no interruption of asynchronous traffic is required; the manager node simply captures the free token to initiate the synchronous traffic.

Once the asynchronous traffic is suspended, a synchronous priority token is issued that can be used only by those nodes that require synchronous transfer of data. Each synchronous node receives one opportunity to transmit during each clock interval. To control the amount of synchronous traffic that exists on a ring, the synchronous bandwidth manager must be involved in the establishment of logical data links used for synchronous data. Control frames sent using normal asynchronous protocols are employed for this function.

Figure 11 Synchronous integration

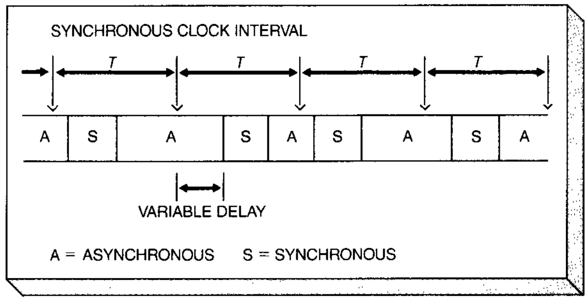

Ring network management

An LAN should provide a data transmission facility with a high degree of availability, resilient to the occurrence of any single fault. Furthermore, the structure of the network should permit comprehensive management of the network facilities and bandwidth. Simple integration of these management functions is imperative to achieve reasonable cost and performance objectives without jeopardizing the reliability of the network. The star-ring network topology in conjunction with a token-access control satisfies these objectives.

Ring fault detection. The topological structure of a ring is often suggested as being inherently vulnerable to faults because of the serial interconnection of nodes. To the contrary, a star-ring provides an extremely reliable communication link because of two properties that can be capitalized upon to offer unique network management functions:

- *Resilient topology*—A single break occurring between two nodes still leaves every node physically attached to the communication medium in a serial link. Therefore, transmission from any node downstream of the break can still be received by all other downstream nodes. The node adjacent to the break will immediately recognize the loss of signal at its receiver and can indicate this condition to all other nodes. The integrity of the ring can then be restored by reconfiguration to bypass the break.
- *Purge mechanism*—A data frame that is transmitted on the ring passes sequentially from one node to another until removed from the ring. On a token-controlled ring, a frame is transmitted completely around the ring, being purged from the ring by the node that originated the frame. As frames traverse the ring, flags can be set by the intermediate nodes to be interpreted by downstream nodes. This allows different types of control information to be propagated around the ring.

Fault detection mechanisms for a ring network are based on the unidirectional propagation of information (electrical signals and data frames) from node to node. Network faults can be categorized into two types: hard faults and soft faults. The resilient topology characteristic of a star-ring provides for efficient hard fault detection, whereas the ring purge mechanism provides for efficient soft fault detection.

A hard fault results from a complete break in the ring segment between two adjacent nodes. (It also includes failures in transmitter or receiver elements located at each end of the ring segment.) When a node on the ring detects the loss of a signal at its receiver, it will transmit a unique series of contiguous frames. Such a transmit state is called "beaconing." A hard fault may initially cause more than one node to enter the beacon state, but eventually all nodes but the one immediately adjacent to the fault will have their received signal restored and will exit the beacon state. Thus, the location of the fault will be isolated to a particular ring segment, the one immediately adjacent to the node that is transmitting beacon-type frames.

A soft fault is characterized by a high error rate, usually caused by a degradation in the electrical signal, as opposed to a complete loss of signal. Soft fault detection by a node on a ring is accomplished by monitoring all frames and verifying the validity of the accompanying frame check sequence (FCS) (Figure 6). A count of the detected FCS errors is kept by each node, and if a predetermined threshold is reached over a given time interval, an indication of the condition can be reported.

Since an invalid FCS will be observed by all nodes downstream from the degraded ring segment, an additional function is required to facilitate the isolation of soft faults. The first node on the ring that detects and logs the error can set the error-detected flag in the physical trailer at the end of the frame. The nodes downstream from this node, observing that the flag has been set, will not log the error. As a result, the location of the fault can be readily isolated to a particular ring segment.

Another flag in the modifier field may be employed to detect the loss or failure of an intended receiving node. The address-recognized flag is set by a node when it recognizes its address within the destination address field, independent of whether the frame is actually copied by the node. The transmitting node will check the state of this flag as the frame is purged from the ring. The flag not being set is an indication that the intended receiver is disconnected from the ring. The logical link that may have been established between the two nodes should then be tested prior to reporting the disconnection to the node manager within the transmitting node.

It is important to note that a bridge will set the address-recognized flag if a frame is to be routed to another ring. Consequently, this flag cannot be employed as an acknowledgment that the frame has successfully reached its destination. Such acknowledgments are considered to be DLC functions and are performed above the physical transport functions of the ring. The receiving node should send a separate frame to convey acknowledgment status back to the original transmitting node.

Ring fault isolation and reconfiguration. The address-recognized flag will also be set by the first node on a ring to receive a frame that contains a broadcast destination address. Downstream nodes can be conditioned to ignore specific broadcast frames if the address-recognized flag has been set. Such a technique can be employed to determine the logical order of active nodes on a ring. Correlation of this information with the source address contained within beacon-type frames or soft error report frames provides the required parameters to isolate the location of the fault.

Periodically, the active monitor will issue a broadcast frame called a roll-call poll. The first active node downstream from the monitor node will set the address-recognized flag and save the source address. Other nodes on the ring will not take any further action at this time. The node that received the frame will then reissue a new roll-call-poll frame containing its own source address whenever a free token is observed. This process continues around the ring until the active monitor receives a roll-call-poll frame without the address-recognized flag set. At that time, each node will have the specific address of the adjacent node immediately upstream. This information is transmitted with all beacon frames and soft error report frames, thereby allowing a network management node to log the location of the fault.

The simple addition of an alternate ring path through each wiring concentrator provides a level of availability not readily apparent with other network topologies.

The use of wiring concentrators to configure a ring network provides the additional capability to remove faulty elements from the network. Lobe bypass switching for isolating faults in the lobe wiring elements has already been discussed in a previous section. A separate technique is necessary if reconfiguration is required as a result of faults occurring within the ring segments interconnecting wiring concentrators. Reconfiguration of this serial portion of the ring network is obtained by use of an alternate ring that parallels the path of the principal ring through each wiring concentrator (Figure 12). The signals on this alternate ring are propagated in the direction opposite to those on the principal ring.[19]

Figure 12 shows four wiring concentrators as they would be configured with both a principal and an alternate ring. If a fault occurs in the ring segment between two wiring concentrators, wrapping of the principal ring to the alternate ring within the two wiring concentrators will restore the physical path of the ring. Note that the logical order of the nodes upon the ring has not been altered. In other words, Node N remains immediately downstream from Node M. Consequently, configuration tables associated with the network management functions do not have to be altered.

The simple addition of an alternate ring path through each wiring concentrator provides a level of availability not readily apparent with other network topologies. This wrapping function, like the lobe bypass function, may be implemented with manual switches or automatic switching logic, or be command-initiated from a remote management facility. The IBM 8100 loop system employs a manual form of both the wrap and bypass functions.

Figure 12 Ring reconfiguration (after wrapping)

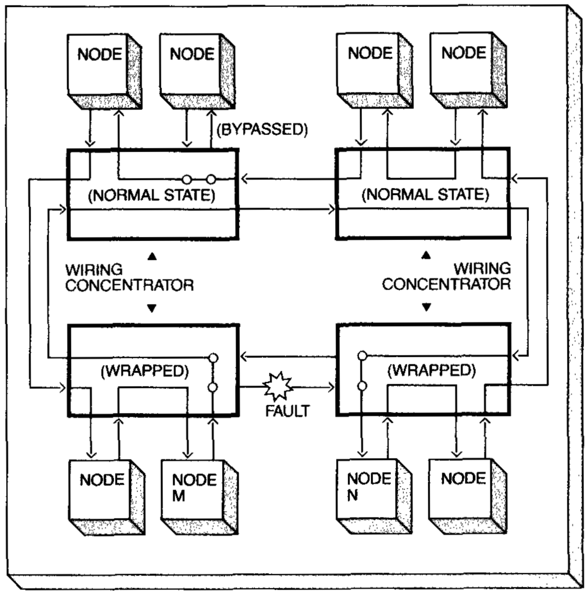

Summary

The rapid development and user acceptance of computer equipment in the office environment has caused an explosive growth in personal workstations and computer terminals. This, in turn, creates a need for a systematic approach to interconnect this equipment by way of a local area network (LAN). One of the key factors distinguishing LANs from other networks is the geographical constraint that allows the freedom to choose innovative transmission techniques for the network. For example, LANs can operate at data rates exceeding several megabits per second as opposed to wide-area networks which operate at data rates of kilobits per second.

An LAN based on a star-ring wiring topology has been described. Token-access control protocols operate on top of the physical wiring system to resolve contention among the attached nodes. As the number of terminals increases, more rings can be added to increase the data-carrying capacity. The rings are interconnected by high-speed switching units called bridges, thus allowing full network interconnectivity among all nodes. Additionally, a degree of network granularity is provided by allowing several rings to be installed concurrently and then interconnected at a later time. Furthermore, these rings maintain autonomy and thereby offer advantages in reliability and performance. Since each ring operates independently of the others, faults on one ring do not affect the others.

The overall system architecture provides a comprehensive approach to a wide variety of network configurations. Small configurations of 10 to 20 terminals can be handled by one ring with little concern for the management problems of a large network. The same system architecture provides the capability to attach several thousand terminals to multiple rings that are bridged together to form a single network. For such large networks, management of the configuration becomes an important aspect. Through the hierarchical wiring structure and the control procedures, it is possible to quickly isolate faults and reconfigure the network around the failed components. The star connection of nodes into the wiring concentrators provides for adding and deleting single nodes. In addition, an alternate ring between wiring concentrators, along with techniques for wrapping to this ring, allows a network to continue operating in cases of failure between wiring concentrators.

Cited references

1. D. D. Clark, K. T. Pogran, and D. P. Reed, "An introduction to local area networks," *Proceedings of the IEEE* **66**, No. 11, 1497–1517 (November 1978).

2. J. D. Markov and N. C. Strole, "Token ring local area networks—A perspective," *Proceedings of COMPCON Fall '82*, Washington, DC (September 1982), pp. 606–614.
3. J. H. Saltzer and K. T. Pogran, "A star-shaped ring network with high maintainability," *Proceedings of the Local Area Communications Network Symposium*, Boston (May 1979), pp. 179–189.
4. H. R. Muller, H. Keller, and H. Meyer, "Transmission in a synchronous token ring," *Proceedings of the International Symposium on Local Computer Networks*, Florence, Italy (April 1982), pp. 125–147.
5. J. R. Pierce, "Network for block switching of data," *Bell System Technical Journal* **51**, No. 6, 1133–1145 (July–August 1972).
6. P. Abramson and F. E. Noel, "Local area network media selection for ring topologies," presented at *COMPCON Fall '82*, Washington, DC (September 1982).
7. R. C. Dixon, "Ring network topology for local data communications," *Proceedings of COMPCON Fall '82*, Washington, DC (September 1982), pp. 591–605.
8. R. M. Metcalfe and D. R. Boggs, "Ethernet: Distributed packet switching for local computer networks," *Communications of the ACM* **19,** No. 7, 395–404 (July 1976).
9. J. E. Donnelley and J. W. Yeh, "Interaction between protocol levels in a prioritized CMSA broadcast network," *Computer Networks* **3,** No. 1, 9–23 (February 1979).
10. *IBM Series/1 Local Communications Controller Feature Description*, GA34-0142, IBM Corporation; available through IBM branch offices.
11. *IBM Synchronous Data Link Control, General Information*, GA27-3093, IBM Corporation; available through IBM branch offices.
12. W. D. Farmer and E. E. Newhall, "An experimental distributed switching system to handle bursty computer traffic," *ACM Symposium on Problems in the Optimization of Data Communications Systems*, Pine Mountain, GA (October 1969).
13. K. E. Coates, "A comparison of the Bell Labs network architecture and the OSI reference model," *Proceedings of COMPCON Fall '82*, Washington, DC (September 1982), pp. 67–75.
14. F. P. Corr and D. H. Neal, "SNA and emerging international standards," *IBM Systems Journal* **18,** No. 2, 244–262 (1979).
15. D. W. Andrews and G. D. Schultz, "A token-ring architecture for local-area networks—An update," *Proceedings of COMPCON Fall '82*, Washington, DC (September 1982), pp. 615–624.
16. M. V. Wilkes and D. J. Wheeler, "The Cambridge Digital Communication Ring," *Proceedings of the Local Area Communications Network Symposium*, Boston (May 1979), pp. 47–61.
17. W. Bux, "Local-area subnetworks: A performance comparison," *IEEE Transactions on Communications* **COM-29,** No. 10, 1465–1473 (October 1981).
18. W. Bux, F. Closs, P. A. Janson, K. Kummerle, H. R. Muller, and E. H. Rothauser, "A local-area communication network based on a reliable token ring system," *Proceedings of the International Symposium on Local Computer Networks*, Florence, Italy (April 1982), pp. 69–82.
19. P. Zafiropulo, "Performance evaluation of reliability improvement techniques for single-loop communication systems," *IEEE Transactions on Communications* **COM-22,** No. 6, 742–751 (June 1974).

Reprint Order No. G321-5182.

Roy C. Dixon *IBM Communication Products Division, P.O. Box 12275, Research Triangle Park, North Carolina 27709.* Dr. Dixon is a senior engineer in the network products system design group at the Research Triangle Park laboratory. He joined IBM in 1964 at Lexington, Kentucky, where he worked on technology activities associated with the IBM copier. He transferred to Raleigh in 1966 and joined the advanced technology organization, working on speaker verification and line control reliability. From 1968 through 1974, he worked on the development of Synchronous Data Link Control and various loop communication systems. He was manager of the Communication Systems Technology department from 1973 to 1979, with key responsibility for the processor architecture of the IBM 8100 Information System. He is currently involved with the architecture and system design of local area networks. Dr. Dixon received the B.S. and M.S. degrees in engineering physics from The Ohio State University in 1958 and a Ph.D. degree in engineering science from Purdue University in 1966.

Norman C. Strole *IBM Communication Products Division, P.O. Box 12275, Research Triangle Park, North Carolina 27709.* Dr. Strole is a staff engineer at the Research Triangle Park laboratory, where he is currently involved in the performance evaluation and analysis of local area network systems. He joined IBM in 1973, and prior to his current assignment, worked in the development and performance evaluation of the IBM supermarket/retail store systems. Dr. Strole received a B.S. degree in electrical engineering from North Carolina State University in 1973. He received the M.S. and Ph.D. degrees in electrical engineering from Duke University in 1975 and 1980, respectively. He currently holds an adjunct appointment in the School of Electrical Engineering at Duke University.

James D. Markov *IBM Communication Products Division, P.O. Box 12275, Research Triangle Park, North Carolina 27709.* Mr. Markov is a senior engineering manager at the Research Triangle Park laboratory. He is currently the manager of system projects for local networks. Previously he was manager of Networking Architecture, a position he held from 1975 to 1978. Since 1970, he has been working in the field of communication systems architecture with a major emphasis on the development of SNA. During the 1960s, he worked on the development of the guidance and control system for the Saturn launch vehicle as part of the Apollo program. Mr. Markov received a B.S. degree in electrical engineering from the University of Akron in 1961 and an M.S. degree in engineering sciences from the University of Alabama in 1970. He is a member of Sigma Tau.