Norman C. Strole

A Local Communications Network Based on Interconnected Token-Access Rings: A Tutorial

Local area networks are expected to provide the communications base for interconnecting computer equipment and terminals over the next decade. The primary objective of a local area network (LAN) is to provide high-speed data transfer among a group of nodes consisting of data-processing terminals, controllers, or computers within the confines of a building or campus environment. The network should be easily accessible, extremely reliable, and extendible in both function and physical size. The rapid advances in computing and communications technology over the last two decades have led to several different transmission schemes and media types that could be used in these networks. The star/ring wiring topology with token-access control has emerged as a technology that can meet all of these objectives. The requirements of small networks with just a few nodes, as well as those of very large networks with thousands of nodes, can be achieved through this one architecture. This paper is a tutorial of the fundamental aspects of the architecture, physical components, and operation of a token-ring LAN. Particular emphasis is placed on the fault detection and isolation capabilities that are possible, as well as the aspects that allow for network expansion and growth. The role of the LAN relative to IBM's Systems Network Architecture (SNA) is also discussed.

Introduction

A local area network (LAN) can be defined as an information transport system for high-speed data transfer among a group of nodes consisting of office or industrial system terminals and peripherals, cluster controllers, or computers, via a common interconnecting medium within the bounds of a single office building, building complex, or campus [1]. The geographical constraints together with advanced transmission technologies enable data transfer rates of many millions of bits per second within the local network. These transmission rates, coupled with a reliable access control scheme, permit a large number of devices to share a common physical interconnection link with minimum interference to one another, while allowing large blocks of data to be transferred with simple error-recovery procedures and data management protocols.

This paper is a tutorial of a LAN communication system based on a ring topology with token-access control. The concepts presented here have been the subjects of research investigations on token-ring local area networks conducted at the IBM Research Laboratory in Zurich, Switzerland [2–4] and at IBM's Research Triangle Park laboratory in North Carolina [5–9]. A brief historical perspective on the evolution of teleprocessing and a description of some basic LAN topologies, control schemes, and transmission techniques are presented. A discussion of the criteria that should be considered in designing a LAN is followed by a functional description of several of the key physical components that comprise a token-ring network. The token-access control protocol for regulating data flow on the ring is explained, including the data frame format and addressing structure, and mechanisms for ensuring token integrity and uniform token access to all attached nodes. The implementation of a LAN in the context of IBM's Systems Network Architecture (SNA) [10] is also discussed. Finally, some of the fault detection and isolation capabilities that enhance the overall reliability of the token-access ring LAN are presented.

Historical perspective

The factors influencing the present surge in office automation and the use of computer terminal equipment in industry have arisen as a natural outgrowth of the revolution in technology that has occurred during the last twenty-five years in both computing and communications. Through the

1960s, large host computers were used primarily in batch processing modes, with punched cards or magnetic tape as the primary input medium. In 1964, one of the first commercial interactive systems, known as the Semi-Automatic Business-Related Environment (SABRE), became fully operational for handling airline passenger information and reservations. This project introduced a new era in telecommunications and contained many significant innovations in the areas of line control, multiplexors, and real-time operating systems [11]. In 1968, one of the earliest and largest nationwide nonmilitary computer networks was established for air-route traffic control using the IBM 9020 system [12]. Data concerning each flight are established at the point of takeoff and are passed from one control sector to the next in the path as the flight progresses. This system introduced new concepts in hardware and software reliability, having such features as a distributed data base, a failsafe system design with gradual degradation, and highly interactive application programs.

The 1970s brought new technologies to both computing and communications that enabled more and more users to have direct access to a computing facility. Advances in microelectronics led to the development of minicomputers which offered good performance capabilities at substantially lower cost. Also, interactive processing came into general use with the advent of the keyboard/printer and keyboard/display terminals as costs decreased and as operating systems shifted to allow both batch mode and interactive processing. The introduction of commercial distributed processing permitted multiple users to share access to common data bases and remote input/output equipment. Distributed processing has now expanded to a global level with the advent of satellite communications and packet-switching. The ARPANET is one of the better-known examples of a nationwide computer communications network [13]. Structured communications functions, such as those implemented in SNA [10, 14], provide well-defined protocols for the reliable exchange of information between network nodes. At the local level, the increased installation of terminals and workstations has resulted in a steady growth in the volume of information that must be transferred over the communications network. High-speed local networks are emerging to meet this growing demand as it becomes economically feasible to implement networking within an establishment. New technologies, such as very-large-scale integrated (VLSI) circuits, have greatly reduced the cost of implementing high-speed transmission and sophisticated communication functions that are necessary to interface a large number of devices to the local area networks.

Current LAN concepts

• *LAN configurations and control schemes*

A fundamental understanding of the basic LAN configurations and control schemes that are prevalent today will enable the reader to comprehend the more detailed discussions in subsequent sections of this paper. There are basically three network configurations that are particularly suited to LANs, namely the *ring*, the *bus*, and the *star* [15]. There are also a number of control schemes that can be implemented with each of these network configurations for regulating access among the network nodes.

Ring A ring network configuration consists of a series of nodes that are connected by unidirectional transmission links to form a closed path. Information signals on the ring pass from node to node and are regenerated as they pass through each node. The application of ring communication systems as local networks has been the object of significant research investigations including such projects as the Distributed Computing System (DCS) [16], the Pierce ring [17], the MIT ring [18], the Cambridge ring [19], and the experimental ring network at the IBM Research Laboratory in Zurich, Switzerland [3].

The Pierce ring and the Cambridge ring utilize a slotted-ring control scheme in which several fixed-length data "slots" flow continuously around the ring. Any node can place a data packet in one of the empty slots, along with the appropriate address information. Each node on the ring examines the address information in each packet and copies those frames that are directed to that node. A "message received bit" can be set by the receiving node to indicate to the originator that the message was received. The originating node must remove the packet from the ring when it returns, thus freeing that slot on the ring.

A second control scheme applicable to ring systems is known as *register insertion*. Each node must load the frame it wishes to transmit into a shift register and insert the entire contents of the register into the ring whenever the ring becomes idle. The shift register, in effect, becomes a part of the ring. The frame is shifted out onto the ring and any incoming frame is shifted into the register temporarily until it too is shifted onto the ring. A frame may be purged (removed from the ring) by either the sending or the receiving node. The sending node can only remove its register from the ring whenever the register contains only idle characters (i.e., no frame data) or the returned original frame itself. A new message cannot be loaded into the shift register until it has been removed from the ring. The IBM local communications controller for Series/1 uses a control scheme based on the register insertion concept [20].

A third control scheme, known as *token-access* control, has been implemented in the DCS, the MIT ring, and the Zurich ring. A unique bit sequence, called a *token*, is passed from one node to another. If a node has no data to transmit, the token is simply passed on to the next node. The receipt of a token gives permission to the receiving node to initiate a transmission. Upon completion of the transmission, the token is passed implicitly (without addressing information) to the next node on the ring.

In all three of the above ring control schemes, the sending node can determine on its own when it may begin transmitting based on the status of the ring at the time. This is contrasted with other ring (or loop) control schemes in which a single master node is responsible for initiating all data transfers. An example of this is the IBM 8100 communication loop. Here, the master node polls the nodes around the loop, allowing each of them in turn to send data if they have any waiting when they are polled.

Bus The *bus* provides a bidirectional transmission facility to which all nodes are attached [15]. Information signals propagate away from the originating node in both directions to the terminated ends of the bus. Each node is tapped into the bus and copies the message as it passes that point in the cable. One bus control scheme that is particularly suited for LAN systems is carrier sense, multiple access, with collision detection (CSMA/CD) [21]. With this scheme, any node can begin transmitting data whenever it detects that the bus is idle. The node continues to monitor the bus for interference from another node that may have begun transmitting at about the same time. Any "collisions" will be detected by all transmitting nodes, causing those nodes to halt transmission for a short random time period before attempting to transmit again.

The token control scheme can also be employed on a bus system. The operation is similar to that for a ring except that an explicit token (containing a specific node address) is employed for a bus, resulting in a logical ordering of the nodes that resembles a ring [5]. Polling can also be used in multidrop-bus configurations where a master controller polls each node to initiate data transfers.

Star The *star* topology, which derives its name from the radial or star-like connection of the various nodes to a centralized controller or computer, is implemented in point-to-point communication schemes that enable each node to exchange data with the central node. This topology is also used in private branch exchanges (PBXs) or computerized branch exchanges (CBXs) where the central node acts as a high-speed switch to establish direct connections between pairs of attached nodes. CBX systems are well suited for voice communications, and they can also accommodate the transfer of data between two nodes. Transmission speeds are presently in the 56-kb/s or 64-kb/s range, which is much less than the multiple-megabit-per-second rates achievable with ring and bus LAN systems.

The bus and ring network topologies are examples of what are called "shared-access links." That is, all nodes share access to a common communications facility, and any signal that is generated at a node propagates to all other active nodes. However, for a meaningful and reliable exchange of data to take place, two nodes must establish a logical, point-to-point link with one another [1]. The physical network thus provides a mechanism for moving the information between nodes that have established logical connections.

• *Transmission techniques*

The various transmission techniques that are implemented within LANs can be generally categorized by the signaling scheme used to transfer the electrical energy onto the medium. The digital information to be transmitted over the medium must first be electrically encoded such that the bits (*1*s and *0*s) are distinguishable at the receiving node(s). The rate at which the encoded bit information is applied to the medium by a sending node is referred to as the transmission speed, expressed in bits per second. The encoded information is applied to the medium in one of two basic methods, commonly referred to as *baseband* and *broadband* signaling.

Baseband In *baseband* signaling, the simpler of the two methods, the encoded signal [22] is applied directly to the medium as a continuous stream of voltage transitions on a copper medium or as a stream of light pulses on an optical fiber medium. One node at a time may apply signals to the medium, resulting in a single channel over which signals from multiple nodes must be time-multiplexed to separate the energy. Baseband data rates exceeding 100 million bits per second (Mb/s) are possible. However, practical limitations, such as the rates at which the attached nodes can continuously send or receive information and the maximum signal drive distance for a given data rate and medium, result in typical LAN data rates of up to 16 Mb/s. Baseband signals must be periodically repeated over a long distance to avoid data loss or interference due to signal degradation. The maximum distance between repeaters is a function of the properties of the transmission medium, use of intermediate connectors, and the data rate. In general, a decrease in the distance between repeaters is needed as the data rate is increased [7].

Broadband *Broadband* transmission schemes, unlike baseband, employ analog signals and multiplexing techniques on the LAN medium to permit more than one node to transmit at a time. Multiple channels or frequency bands can be created by a technique known as frequency division multiplexing (FDM). A typical broadband system has a bandwidth of 300 megahertz (MHz) which can be divided into multiple 6-MHz channels (as used in cable television signal distribution) with pairs of channels designated for bidirectional communications over a single cable. A standard 6-MHz channel can readily accommodate data rates up to 5 Mb/s. Two adjacent 6-MHz channels can be used to provide a single 12-MHz channel for data rates up to 10 Mb/s [23]. Access to a channel that is shared by a large number of nodes must be regulated by some type of control scheme, such as token passing or CSMA/CD. Broadband operation requires that a modulate/demodulate function be performed by radio-frequency modulator/demodulators (rf modems) at the sender and receiver, respectively, resulting in a higher cost per attachment than with baseband schemes. However, this enables broadband signals to be transmitted for longer distances between repeaters over a specially designed broadband cable.

• *Data frames or packets*
The control scheme that is selected for a particular topology governs the access of the various nodes to the network. In all of the transmission control schemes, the data being transferred from one node to another must be transmitted in a particular format that is recognizable by the receiving node. The actual format of a packet or frame varies depending upon the specific communication requirements. In addition to the data (information) field, the frame contains some type of delimiter field that is distinguishable from all other fields to mark the beginning and end of the frame. There may also be some control fields that contain status information, specific commands, format qualifiers, or the frame length. Most of the schemes require that the frame contain address information to identify the sender (source) and intended receiver (destination). Finally, one or more check bits are normally included to enable the receiver to detect possible errors that may have occurred after the sender transmitted the frame. A more complete description of a frame format is given later.

Selection criteria for a LAN

There are a number of factors that must be considered when selecting or designing a LAN. Each of the LAN configurations and control schemes described earlier has both strong and weak points, requiring the designer or user to make trade-offs in selecting the appropriate system design. One criterion that may be used is a comparison of the information transfer characteristics of the various schemes. The performance characteristics of several LAN transmission schemes, including token-access rings, slotted rings, and CSMA/CD buses, were evaluated by Bux [24]. He compared their data-throughput characteristics as affected by system parameters such as transmission rate, cable length, packet lengths, and control overhead. He concluded that the token ring performs well at high throughputs and transmission speeds of 5 and 10 Mb/s. The CSMA/CD bus also performed well at these speeds as long as the ratio of propagation delay to mean packet transmission time was on the order of 2 to 5%. The requirement to minimize propagation delay on a bus restricts the length of the bus in a LAN system, a factor that is not as critical in a token ring system. Another performance study by Stuck [25] showed the token ring to be the least sensitive to workload, offering short delays under light load and controlled delays under heavy load. The token bus was shown to be not as efficient as the token ring under heavy load and to have greater delay under light load. Stuck also found the CSMA/CD bus to have the shortest delay under light load, but to be quite sensitive to heavy workload, and to be sensitive to the bus length and message length, performing better with a shorter bus and longer messages.

A primary criterion in selecting a LAN is to minimize the cost of installing, operating, and maintaining the system without sacrificing network performance and reliability. In general, a baseband token-access communication scheme can be implemented with low-cost interface adapters and does not require the more sensitive rf modems associated with broadband systems. Also, the token-access ring wiring medium can be installed within a building using a radial (star) wiring scheme. Star wiring implies that the transmission cable is installed from concentration points within a building to the various user work areas, such as offices or laboratories. Wall outlets in each of these areas can provide physical interfaces to the network to permit fast, reliable attachment or relocation of the nodes.

Another design criterion is that any network faults or disruptions should be quickly detected and isolated to restore normal operation. The wiring concentration points provide centralized access to some of the primary network components to enhance the maintenance and reliability aspects of the overall system [18]. In addition, the flow of data and control signals around a closed ring provides an inherent capability for monitoring normal token operation from a single point on the ring and enables ring faults (such as a break in a segment of the ring) to be quickly detected. A more comprehensive discussion of the fault detection and isolation capabilities within a token-ring LAN is presented later in this paper.

The above LAN design criteria and others are incorporated into the token-ring architecture described here. This description is intended to give the casual reader an understanding of the fundamental aspects of a token-ring system while providing enough detail to permit further comparison with other LAN architectures.

Token-ring topology and network components

A star/ring wiring scheme combines the basic star wiring topology with the unidirectional, point-to-point signal propagation of the ring configuration. A major benefit of radial wiring is the capability of isolating those links that have failed, or are not in use, from the network [18]. Also, additional links may be added at any time. Major components of a two-ring system are represented in Fig. 1 and described next.

• *Ring interface adapter*

The term *node* encompasses a wide variety of machine types that can attach to and communicate over the local network. The primary functions associated with token recognition and data transmission are distributed to each node within the network. Advances in VLSI technology make it possible to delegate a large portion of this communication function to a ring interface adapter within each node, thus freeing the node from this processing. The adapter handles the basic transmission functions that are described later, including frame recognition, token generation, address decoding, error checking, buffering of frames, and link fault detection.

There may be instances where a device with an incompatible communication interface is to be attached to the local network. In this case a separate interface converter may be attached between the device and the token ring to perform the protocol and timing conversions that are required to send and receive data over the local network.

• *Wiring concentrators*

Wiring concentrators which contain a series of electronic relays are the central elements for structuring the star/ring wiring layout. Wiring lobes, consisting of two pairs of conductors for separate send and receive paths, emanate from the wiring concentrators to the various network interface points (e.g., wall outlets) at each of the node locations throughout a building. The lobes are physically connected within the concentrators to form a serial link, with the wiring concentrators then being interconnected in a serial fashion to complete the ring (Fig. 1). A lobe is only included in the ring path when the node is active; otherwise, an electronic relay within the wiring concentrator causes that lobe to be bypassed [4, 26]. Nodes may easily be moved from one location to another without requiring the installation of a new cable. The wiring is segmented at the wiring concentrators rather than being a continuous cable, thus permitting the intermixing of transmission media. For example, shielded twisted-pair wire can be used to interconnect the wiring concentrators to nodes, while optical fibers can be used for transmission links between wiring concentrators.

The wiring concentrators provide points within the star/ring network that facilitate reconfiguration and maintenance, thus enhancing network reliability. As described earlier, electronic relays within the wiring concentrators are activated whenever a node is powered on to bring that node and the associated wiring lobe into the active ring. Should the attached adapter detect a fault within either its own components or the wiring lobe between itself and the wiring concentrator, it can deactivate the relay and remove itself

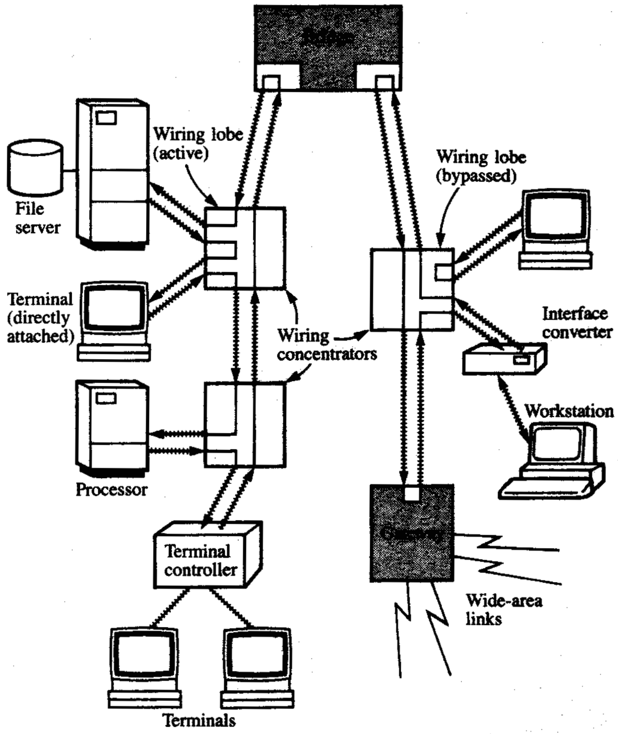

Figure 1 Wiring topology and components of a two-ring network. *Wiring concentrators* are the central elements for structuring a star/ring wiring layout. *Bridges* provide high-speed links between multiple rings. *Gateways* are used to access wide-area networks for long-distance communications.

from the ring. This lobe-bypass function may also be accomplished by manually deactivating the faulty node in case a faulty adapter is unable to remove itself from the ring. Passive concentrators contain electronic relays but no active elements, such as processing logic or power supplies, and require only enough power from an attached node to activate the relays to insert a node into the ring [18, 26]. Active wiring concentrators, on the other hand, contain processing logic and their own power supplies, and thus have the ability to detect and bypass faults that occur in the ring segments between concentrators. Active concentrators may also be remotely activated through the receipt of appropriate network management commands to bypass faults.

• *Bridge function*

Multiple rings may be required in a LAN when the data transfer requirements exceed the capacity of a single ring or when the attached nodes are widely dispersed, as in a multifloor building or campus environment. Therefore, large networks may typically have several rings with 100 to 200 nodes per ring. Two rings can be linked together by a high-speed switching device known as a *bridge* (Fig. 1). A

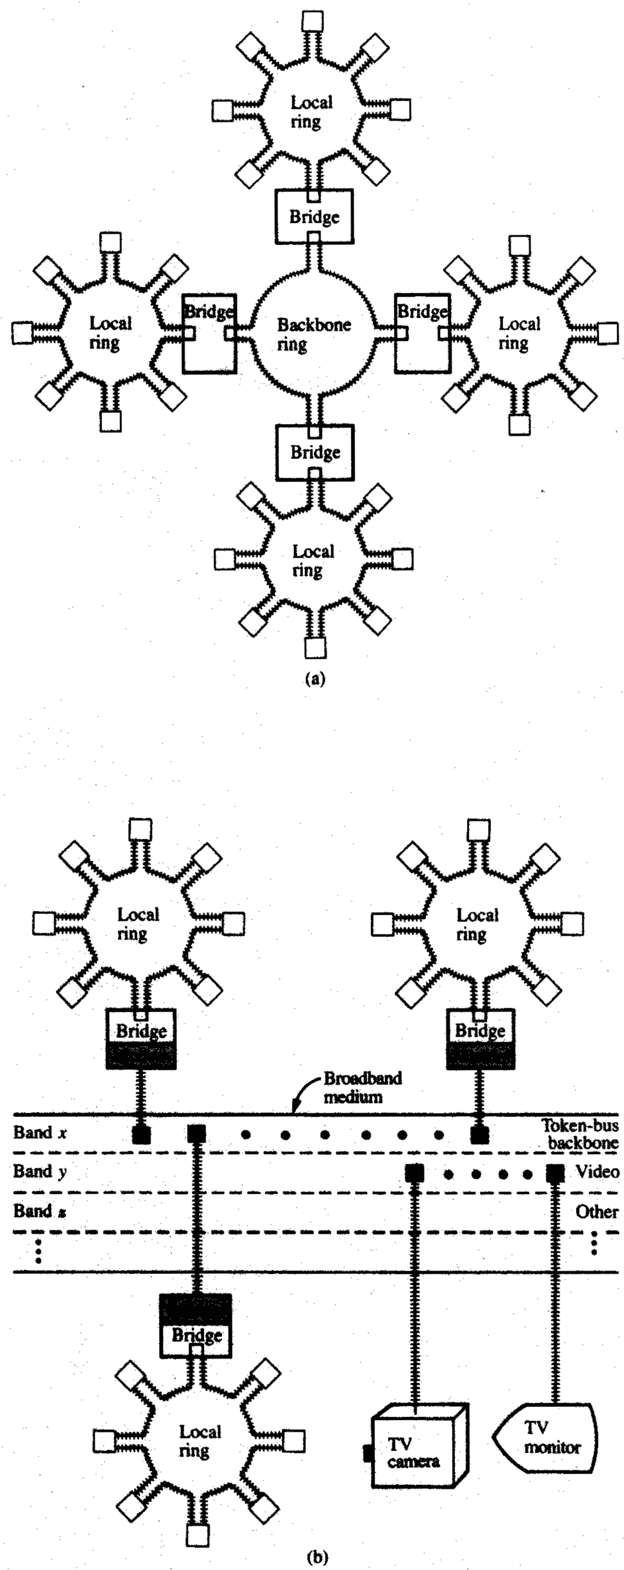

Figure 2 Multiple-bridge local area network: (a) A token-ring backbone wherein the bridges provide a logical routing of frames among rings and perform a speed conversion between the local rings and the backbone ring; (b) A broadband-bus backbone in which the bridges switch the ring data onto a channel of the broadband medium.

bridge is capable of providing a logical routing of frames between the rings based on the destination address information contained in the headers of the frames. An additional capability of the bridge is to perform transmission speed changes from one ring to another. Each ring retains its individual identity and token mechanism, and can therefore stand alone in the event the bridge or another ring is disrupted. The bridge's interface to a ring is the same as any other node's, except that it must recognize and copy frames with a destination address for one of the other rings within the network. Also, several frames may be temporarily buffered in the bridge while awaiting transfer to the next ring.

The local network can be further expanded to meet larger data capacity requirements by interconnecting multiple bridges. This results in a hierarchical network in which multiple rings are interconnected via bridges to a separate high-speed link known as a *backbone*. The backbone itself may be a high-speed token ring, Fig. 2(a), or it may be a token-access bus link, such as a channel within a broadband CATV system, Fig. 2(b). The address field format described later is structured to designate the specific ring to which a node is attached, thereby facilitating the routing of packets through bridges [17].

• *Gateway function*

In today's networking environment, a growing number of users require access to nationwide communication links as well as to the local network. For example, the primary host processing system may be located in another city or state, requiring that access be provided over a geographically dispersed area. A *gateway* node provides an interface between a LAN and a wide-area network for establishing long-distance communications between nodes within the LAN and nodes that are within other LANs or that are accessible directly on the wide-area network. Wide-area networks include private and commercial satellite links, packet-switching networks, leased lines, or other terrestrial links. They generally operate at lower transmission rates than most LANs, usually in the kb/s transmission range. A gateway can perform the necessary address translations as well as provide the speed and protocol conversions that are required to interface the LAN to these various transmission facilities. There may also be applications where a gateway could be used as an intermediate node between a token-ring LAN and a node in either a CSMA/CD- or PBX-based LAN.

• *Transmission media in a ring network*

The most prevalent LAN wiring media today are twisted-pair copper wire and coaxial cable. However, the use of optical fiber media in LAN systems will certainly increase in the coming years. Unshielded twisted-pair copper medium is used extensively for normal voice communications. However, voice transmission frequencies are much lower than those of local networks, and the impact of environmental noise is not as detrimental to analog voice signals as it could be to digital LAN traffic. Thus, "voice-grade" twisted-pair media cannot satisfy all of the requirements of a high-speed LAN. A high-quality "data-grade" twisted-pair medium has a higher characteristic impedance and is shielded to reduce the external electromagnetic interference (EMI), as well as the outward electromagnetic radiation from the cable itself. Also, the data-grade pair is a balanced medium (i.e., both wires have the same impedance to ground) which, in combination with the shielding, reduces the crosstalk interference from other twisted pairs within the same cable. (See Park and Love [27].) Thus, a well-designed twisted-pair cable can provide a reliable transmission medium for a local network system operating at baseband data rates of 1 to 10 Mb/s [7].

Optical fibers offer several advantages over copper media for use in LANs, including low susceptibility to electromagnetic interference and low signal attenuation over long distances at speeds greater than 100 Mb/s [7]. As this technology matures, lower-cost splicing techniques, as well as lower production costs, will make the use of optical fibers more economical than it is today. As noted earlier, the token-ring control scheme and the radial wiring topology enable optical fibers to be used within the LAN. For example, optical fibers could be used between buildings where lightning poses a hazard that requires additional protection with metallic media but not with optical fiber. Fibers could also be used in industrial environments where higher levels of EMI are found.

Broadband coaxial cable systems can be incorporated within a LAN for transmitting both analog video and digital data signals. While television systems do not play a major role in the office today, they may do so in the future. As noted earlier, standard CATV channels could be used for data transmission between rings by means of broadband bridges or for the interconnection of LANs via appropriate gateways.

- *Differential Manchester encoding scheme*

The binary data that originate within a node must be encoded for effective transmission over the ring. Many different encoding schemes are possible, including 8B/10B encoding [28] and differential Manchester encoding [22]. The differential Manchester code is being considered by the IEEE Project 802 Local Area Network Committee as the standard scheme for baseband ring transmission [29]. This code allows for simpler receive/transmit and timing recovery circuitry and offers a smaller delay per station than do block codes. Also, the differential Manchester code allows for the interchanging of the two wires of a twisted pair without introducing data errors [4]. A series of contiguous *1*s and *0*s

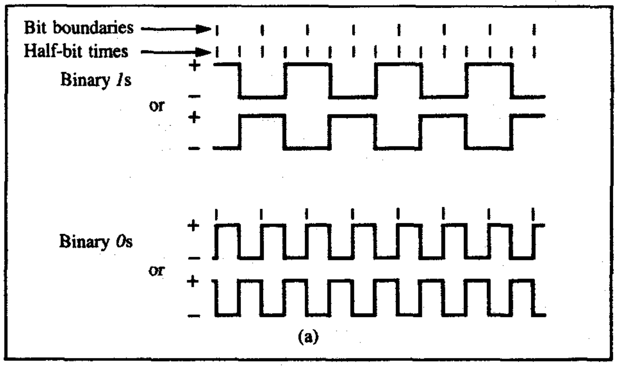

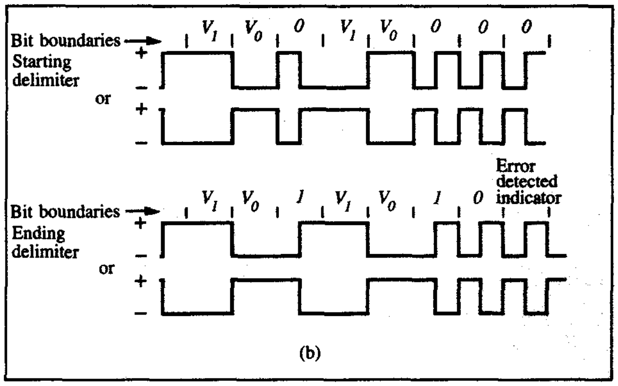

Figure 3 Differential Manchester encoding: (a) A *1* is distinguished from a *0* by the presence or absence of a polarity change at the bit boundary followed by a polarity change at the half-bit time of the synchronous signal on the medium; (b) Starting and ending delimiters contain code violations to distinguish them from data for delineating a frame.

(bits) would appear as shown in Fig. 3 when encoded. Each bit is encoded as a two-segment signal, with a signal transition (polarity change) at the middle or half-bit time. This transition at the half-bit time effectively cancels the dc component of the signal and provides a signal transition for clock synchronization at each adapter. The *1*s are differentiated from the *0*s at the leading bit boundary; a value of *1* has no signal transition at the bit boundary, while a value of *0* does. In decoding the signal, only the presence or absence of the signal transition, and not the actual polarity (positive or negative), is detected.

A code violation results if no signal transition occurs at the half-bit time. Code violations can be intentionally created to form a unique signal pattern that can be distinguished from normal data (*1*s and *0*s) by the receiving adapter(s). Such unique signal patterns can be inserted to mark the start and end of a valid data frame. Four code violations are used in pairs to maintain the dc balance, as well as to prevent spurious signals from being recognized as valid delimiters.

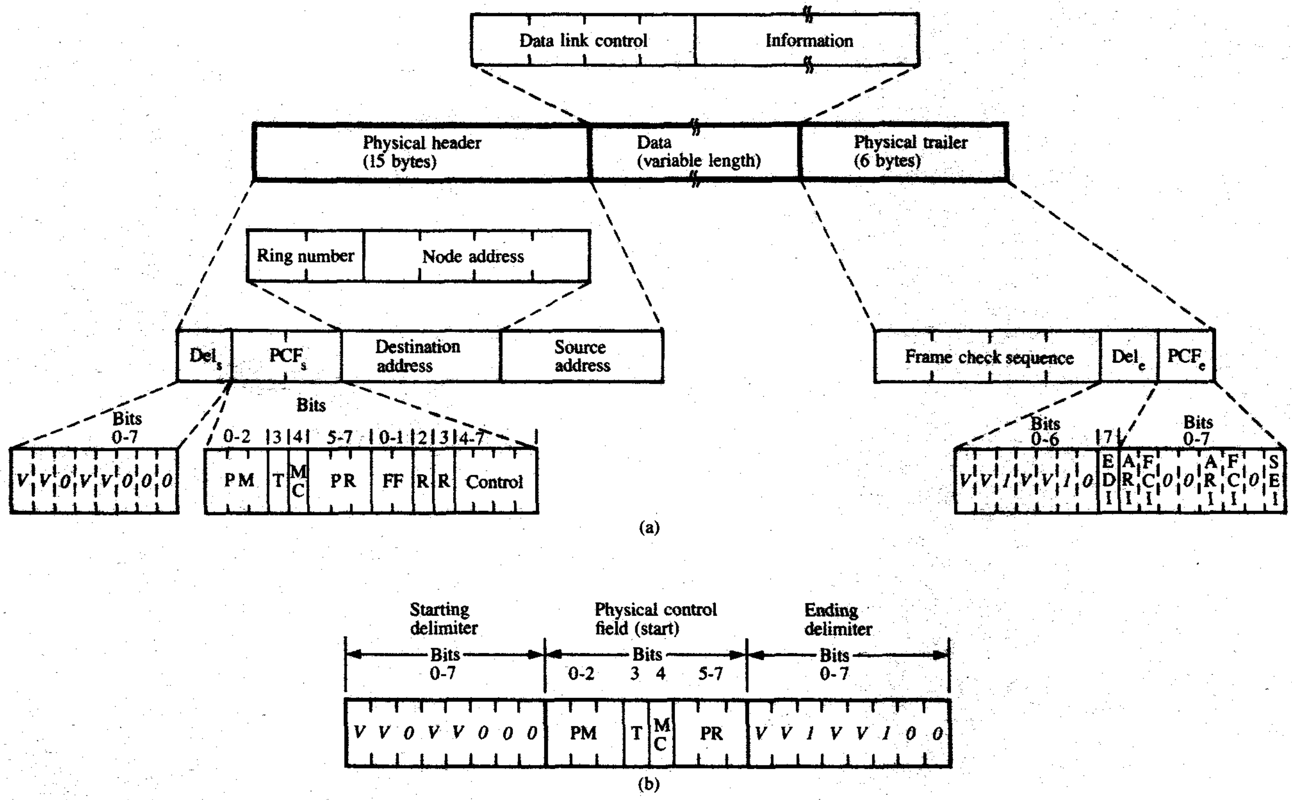

Figure 4 (a) Overall frame format. (b) Free token frame format. The abbreviations used are as follows: ARI = address-recognized indicator; Del_e = ending delimiter; Del_s = starting delimiter; EDI = error-detected indicator; FCI = frame-copied indicator; FF = frame format; MC = monitor count; PCF_e = ending physical control field; PCF_s = starting physical control field; PM = priority mode; PR = priority reservation; R = reserved; SEI = soft error indicator; T = token: *1* means busy, *0* means free; V = code violation.

Also, within each violation pair, the first is a violation of the normal *1* bit, V_1 (i.e., no signal transition at the leading bit boundary), while the second is a violation of the normal *0* bit, V_0 (i.e., signal transition occurs at the leading bit boundary). Thus the bit pattern $V_1V_0 0 V_1V_0 0$ is used to denote the starting delimiter, while the pattern $V_1V_0 1 V_1V_0 1$ is used to denote the ending delimiter (Fig. 3). These code patterns comprise the first six bit-times of either delimiter, leaving two bit-times within the byte (eight bit-times) for other uses. Any other code violations that occur are assumed to be transmission errors.

Token-ring architecture

• *Data frame format*

The general format for transmitting information on the ring, called a frame, is shown in Fig. 4(a). The "data" portion of the frame is variable in length (up to a fixed maximum) and contains the information that the sender is transferring to the receiver. The data field is preceded by a physical header, which contains three subfields. The first is a starting delimiter (Del_s) that identifies the start of the frame. The starting delimiter is a unique signal pattern that includes pairs of code violations of the differential Manchester encoding scheme as described earlier. Next, a starting physical control field (PCF_s) is defined for controlling the access to the transmission facility and for passing encoded information to the adapters. This two-byte field includes a one-bit token indicator that indicates whether the token is free (*0*) or busy (*1*). The frame format shown in Fig. 4(a) would contain a busy token indication. A free token, on the other hand, contains only the first byte of the PCF_s and the starting and ending delimiters [Fig. 4(b)]. A token priority mode, in conjunction with the priority reservation indicators within the first byte of the PCF_s, provides different priority levels of access to the ring. The monitor count bit is used in conjunction with a token monitor function to maintain the validity of the token. Both the priority access scheme and the token monitor function are described subsequently. The second byte of the PCF_s contains a two-bit frame format (FF), two reserved bits, and a four-bit control indicator. The FF bits enable the receiving node to determine whether the information within the data field of the frame contains medium-access control (MAC) information (FF = *00*) or user data (FF = *01*). MAC frames may optionally include frame status information or other urgent ring management information within the control indicator subfield. Finally, the header includes a field containing the address of the node that originated the information and the address of the node (or nodes) destined to receive the information. Both address fields contain six bytes, with the first two bytes indicating the specific ring number (in multiple ring networks) and the last four bytes indicating the unique node address.

The data field itself may be subdivided to include a data link control subfield as well as the user information field. Data link control information is necessary for the higher-level protocols that are normally used within data communication networks [14].

The data field is followed by a physical trailer which is also composed of three subfields. The first portion of the trailer contains a four-byte frame check sequence (FCS) that is calculated by the source node and 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. The FCS is a 32-bit cyclic redundancy check (CRC) that is calculated using a standard generator polynomial [29]. Next, an ending delimiter (Del_e) is provided to identify the end of the frame. This delimiter also contains pairs of code violations such as were found in the starting delimiter, but with *1*s following the violations to distinguish the Del_e from the Del_s. The last bit of the Del_e is designated as the error-detected indicator (EDI). This indicator is always *0* during error-free ring operation, but is set to *1* by any intermediate ring adapter that detects an error with the FCS of a frame. The Del_e is followed by an ending physical control field (PCF_e), which is also employed for certain physical control functions. The PCF_e contains bits that can be modified while the frame is traversing the ring and is therefore not included in the calculation of the FCS character. For this reason the frame-copied indicator (FCI) and the address-recognized indicator (ARI) bits are duplicated to provide a redundancy check to detect erroneous settings. The uses of the FCI, the ARI, and the soft error indicator (SEI) within the PCF_e are discussed later.

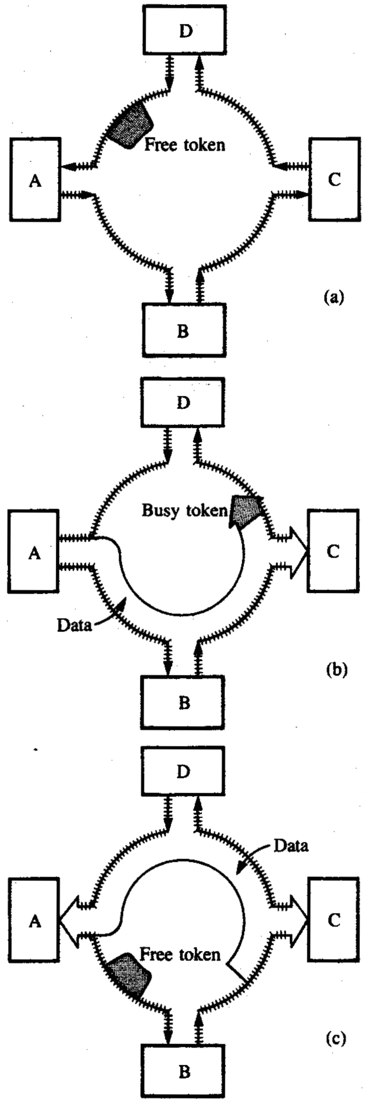

Figure 5 Token-ring access control protocol. (a) Sender (node A) looks for a *free* token, then changes *free* token to *busy* and appends data. (b) Receiver (node C) copies data addressed to it; token and data continue around the ring. (c) Sender (node A) generates free token upon receipt of physical header and completion of transmission, and continues to remove data until receipt of the physical trailer.

• *Token-access control protocol*

The token-access control mechanism for regulating data flow in a ring topology is based on the principle that permission to use the communications link, in the form of a "free" token, is passed sequentially from node to node around the ring. The "free" token, as described earlier, contains a one-bit indication that the token is "free" (T = *0*). With the token-access control scheme, a single free token circulates on the ring (Fig. 5), giving each node in turn an opportunity to transmit data when it receives the token. Each node introduces into the ring approximately a one-bit delay as the time to examine, copy, or change a bit as necessary. A node having data to transmit can change the token indicator from free (*0*) to busy (*1*) and begin data transmission by appending the remainder of the physical header field (second byte of PCF_s and destination/source address fields), the data field, and the physical trailer. The node that initiates a frame transfer must remove that frame from the ring and issue a free token upon receipt of the physical header, allowing other nodes an opportunity to transmit. If a node finishes transmitting the entire frame prior to receiving its own physical header, it continues to transmit idle characters (contiguous *0*s) until the header is recognized. This ensures that only one token (free or busy) is on the ring at any given time. The single-token protocol is thus distinguished from a multiple-token protocol in which a new free token is released immediately

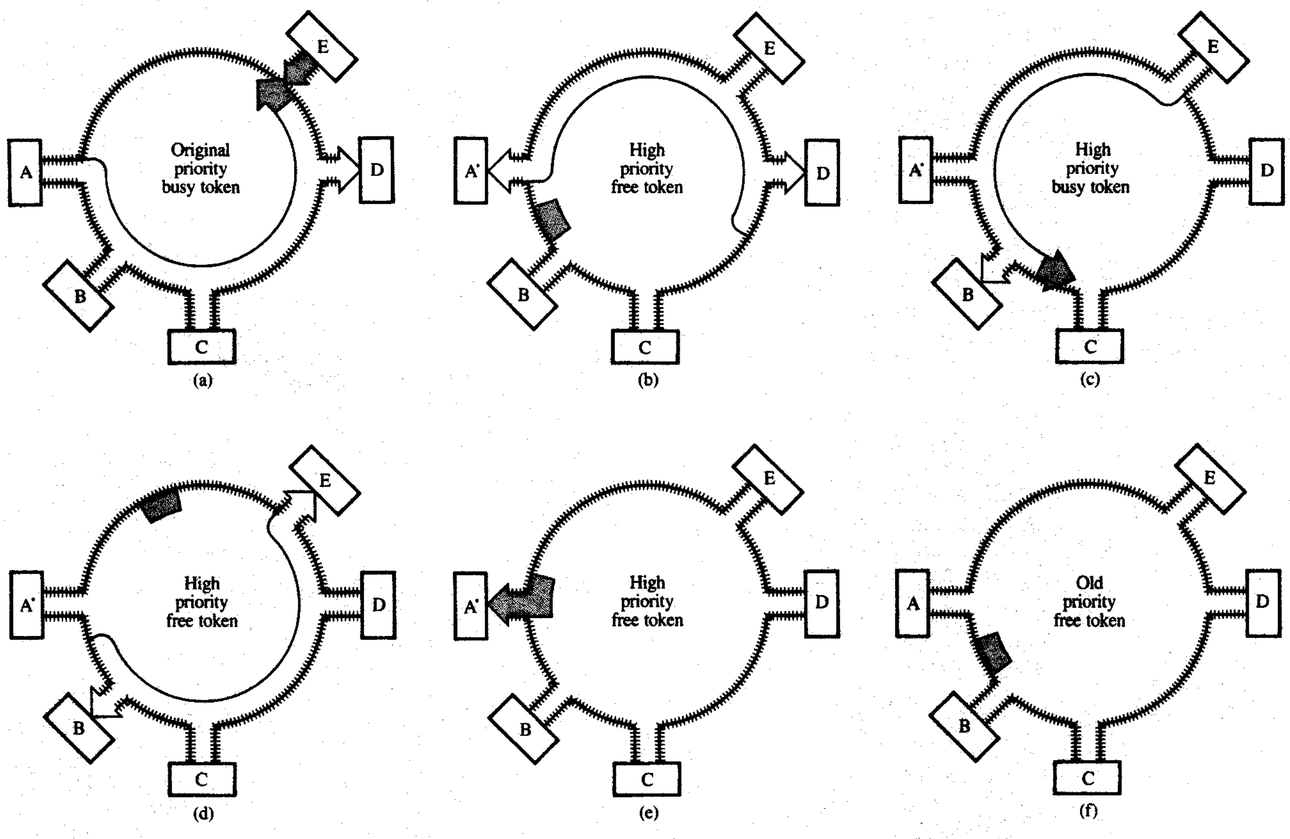

Figure 6 Priority with "uniform access." (a) "A" is in the process of sending to "D." "E" makes a higher priority level reservation. (b) "A" generates a higher priority level free token and saves the original priority level. Since "B" is at a lower priority level, it cannot use the higher level token. (c) "E" uses the free token to send data to "B." (d) "E" generates a free token at the current (higher) priority level. (e) "A" sees the higher priority level free token and generates a free token at the saved preempted priority level. (f) Now "B" is given the chance it missed earlier to use the free token.

following the last byte of the trailer. Only one free token is permitted in either case, with multiple busy tokens (frames) possible in the latter case if the length of a data frame is short compared to the time that it takes the leading busy token to travel completely around the ring. However, multiple-token error recovery is much more complex, while performance improvements of a multiple-token protocol over a single-token protocol are marginal when adapter delays are small [24].

Note that the ARI bits in the physical trailer of a frame are set by a node whenever it recognizes its own address within the destination field of that frame. The FCI bits are set by a node whenever it copies a frame from the ring. If a node cannot copy a frame from the ring, only the ARI bits are set, indicating to the sending node that the destination node is active but did not copy the frame. The ARI and FCI bits are used only for passing status information on a ring and do not provide an end-to-end acknowledgment of frame transfer in a multiple ring network.

The token-access control protocol provides uniform access to the ring for all nodes. A node must release a free token after each transmission and is not allowed to transmit continuously (beyond a maximum frame size or preset time limit) on a single token. All other nodes on the ring will have a chance to capture a free token before that node can capture the token again. In some system configurations, it may be necessary for selected nodes, such as bridges or synchronous devices, to have priority access to the free tokens. The priority mode and reservation indicators in the PCF_e are used in regulating this access mechanism. Various nodes may be assigned priority levels for gaining access to the ring, with the lowest priority being *000*. This means that a selected node can capture any free token that has a priority mode setting equal to or less than its assigned priority. The requesting node can set its priority request in the reservation field of a frame if the priority of that node is higher than any current reservation request (Fig. 6). The current transmitting node must examine the reservation field and release the next free token with the new priority mode indication, but retain the interrupted priority level for later release. A requesting node uses the priority token and releases a new token at the same priority so that any other nodes assigned that priority can also have an opportunity to transmit. When the node that originally released the priority free token recognizes a free token at that priority, it then releases a new free token at the level that was interrupted by the original request. Thus, the lower-priority token resumes circulation at the point of interruption.

Network addressing and frame routing The local network environment is expected to be changing frequently as nodes are added to a ring, moved from one ring to another, or removed completely [6]. The addressing and routing schemes for transferring frames from the source node to the destination node must be efficient and reliable in this changing environment. There are three types of addresses that could be used to identify a particular node. One type of identification is an identifier address that could be preset by hardware components within a node and remain unique no matter where the node is located within the network. Another type of address could be assigned to the specific wall outlet or wiring concentrator lobe to which the node is attached, denoting the physical location of the node within the network. A third logical address could be assigned to a node by a separate address server whenever the node becomes active in the network. The important point is that an address exists that distinguishes one particular node from all the others. This address is appended to the ring number to form the complete source or destination address that is included in each frame [Fig. 4(a)].

Certain addresses can be defined to be "all-stations" or "all rings" addresses that can be used when a frame is to be sent to all nodes in the network or on a particular ring. Otherwise the sender must determine at least the unique portion of the intended receiver's address before the first message can be sent. This information could be obtained from a central address management function or from a published source (similar to a telephone book). The sender then transmits a special MAC frame, known as an *address resolution request* frame, containing the unique portion of the destination address and an "all-rings" address, since the sender does not yet know which ring the destination node is attached to, or even if that node is currently active. If the target node is active, it will receive the address resolution request and respond to the requesting source node with another MAC frame containing its full address information (including ring number). Both nodes then know the complete address of the other and can continue communicating. Saltzer describes how a similar scheme could be used to determine the exact route through the network that the resolution request frame traversed in reaching the destination node [30].

The ring number portion of the destination address within frames is examined by a bridge and those frames that are directed to another ring are transferred within the bridge to the appropriate destination ring (if multiple rings are attached to a single bridge) or to the backbone. A frame is removed from the backbone by a bridge only if the destination ring number can be reached directly through that bridge or if the frame contains a general "all-rings" broadcast address.

A communication path between a node, A, within the local network and another node, B, in a separate network is established through an intermediate gateway node. Node A can communicate directly with the gateway node using the addressing scheme described above. However, the address information that identifies node B must be contained within the information portion of the frames which the gateway node examines to route the frames. Thus, routing through a gateway node can be associated with those functions performed at the path control layer of an SNA node or the network control layer of the ISO model [10].

Token monitor function Normal token operation is monitored by a special function, known as the token monitor, that is always active in a single node on each ring. This function can be performed by any node on the ring and is necessary to initiate the proper error recovery procedure if normal token operation is disrupted. This includes the loss of a free token or the continuous circulation of a busy token, both of which prevent further access to the ring [3]. 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 a 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 to *1* as the frame passes by. The failure of the transmitting node to remove the frame causes it to pass the token monitor a second time. The token monitor, observing the monitor count flag set, removes the frame from the ring and issues a free token. The active monitor also maintains a timer which is reset upon the passage of either a busy token or a free token. Loss of the token due to interference or noise causes the timer to expire, prompting the token monitor to reinitialize the ring with a free token.

The capability to be an active monitor exists in all active nodes attached to a ring. These other nodes maintain a standby monitor status and are prepared to become the new active monitor should a failure in the current active monitor occur. The standby monitors are essentially monitoring the ring to detect abnormal ring operation that could occur whenever the active monitor has failed. The detection of an error condition by any standby monitor initiates a recovery procedure that allows it or one of the other nodes to become the new monitor [3].

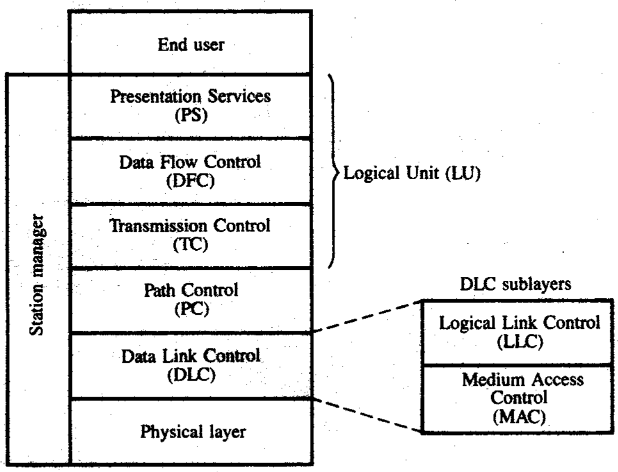

Figure 7 Systems Network Architecture (SNA) layers. Local network protocols and transmissions are performed at the data link control (DLC) and physical layer functional layers of SNA. DLC functions can be further subdivided into logical link control and medium access control sublayers.

Extension of token-access protocol for synchronous applications The token-access control protocol, in conjunction with the priority reservation scheme described earlier, provides a basis for incorporating synchronous data transfer over a token ring [1, 3, 6]. Synchronous operations require that data packets be transmitted and/or received at periodic intervals rather than asynchronously. However, the token-access mechanism is basically an asynchronous control scheme. Synchronous operations, such as real-time digitized voice transfer, can be accomplished by periodically making the ring available only for synchronous data transfer. A special function node, known as a synchronous bandwidth manager, can maintain an internal clock to provide the correct time interval. At the end of each time interval, the synchronous bandwidth manager raises the priority of the circulating free token via the reservation scheme described earlier. This unique priority frame is available only to those nodes requiring synchronous data transfer. The token remains at this priority until all synchronous nodes have had an opportunity to transmit one data frame. The synchronous bandwidth manager then releases a free token at the original priority. The synchronous bandwidth manager must be involved in the establishment of any synchronous communication connections to regulate the synchronous load added to the network so as to avoid excessive delays in the asynchronous traffic.

Token-ring architecture relative to SNA

• *Overview*

A LAN provides a basic transport mechanism for data transfer among nodes within the network. However, the LAN by itself does not provide all of the functions that are necessary for two nodes to manage and conduct a meaningful two-way exchange of information. The same higher-level communication protocols that are implemented to control data transfer across public data networks [31] are also applicable to data transfer across a LAN. IBM's Systems Network Architecture (SNA) and the International Organization for Standardization (ISO) reference model separate network functions into layers to facilitate the description of the protocols as well as their implementation [10, 32, 33]. SNA protocols can be implemented for managing the flow of information within a local network as in any other SNA environment [8]. The basic SNA layers are depicted in Fig. 7. The functions of the lowest layer, known as the *physical layer*, are unique to the particular transport mechanism that is implemented, whether it be a communications loop, a multidrop bus, or a token ring. The next higher layer, *data link control* (DLC), is traditionally independent of the actual physical transport mechanism and performs the functions that are necessary to ensure the integrity of the data that reach the layers above DLC. A modification of this concept that would allow the DLC layer to share physical access functions with the physical layer is presented later. Network flow control and message unit routing functions are provided at the next higher layer, known as *path control* (PC). The composite of the functions at the DLC and PC layers of all of the nodes within an SNA network comprises the path control subnetwork, the fundamental transport mechanism for transferring data from a source SNA node to a destination SNA node.

An end user within an SNA environment may be a person engaged in interactive work at a display terminal or an application software function that is active within a host. The end user in an SNA node is represented to the network by a *logical unit* (LU). A logical unit comprises the functions of the upper three layers of SNA, transmission control (TC), data flow control (DFC), and presentation services (PS). Two end users communicate within an SNA network by establishing a logical path, called a session, between their respective LUs. The LU-LU session provides a temporary connection for moving the data between the end users, utilizing the services of the PC subnetwork. A more detailed discussion of the logical unit functions relative to local area networks can be found in [8]. The functions of the physical layer and the data link control layer are discussed in more detail next.

• *Physical layer*

The physical layer of the model encompasses the basic functions associated with placing the electrical signals onto the transmission medium. This includes such fundamental operations as signal generation, phase timing along the ring, and the encoding of the signal information using the differential Manchester scheme. These operations are performed within the ring interface adapter at each active node in the network.

• *Data link control*
Data link control (DLC) is the next layer above the physical layer in SNA. The IEEE Project 802 Committee on Local Area Networks and the European Computer Manufacturers Association (ECMA) have proposed that for local networks, the data link layer of the ISO model, which corresponds to the DLC layer of SNA, be further subdivided into two functional sublayers: logical link control (LLC) and medium-access control (MAC), as depicted in Fig. 7 [29, 34]. This functional decomposition essentially separates those DLC functions that are hardware dependent from those that are hardware independent, thereby possibly reducing the cost of developing data equipment to interconnect to different types of physical transmission interfaces [35].

Medium-access control sublayer The medium-access control (MAC) sublayer of DLC includes those functions associated with frame and token transmission which can be, but are not necessarily, performed by the interface adapter in each node. The actual functions performed at this sublevel depend upon the LAN architecture that is being implemented. The functions for a token-ring LAN include the following:

- Token protocols—basic token access and generation, including token priority reservation scheme.
- Address recognition—copying of frames from ring based on destination address in physical header.
- Frame delimiters—identifying beginning and ending of frames (transmit and receive) via unique delimiters.
- Frame check sequence (FCS)—calculation and verification of cyclic redundancy check (CRC) for each frame.
- Token monitor—performance of all functions associated with active or standby token monitors.

Logical link control sublayer The second sublayer, logical link control (LLC), includes those functions unique to the particular link control procedures that are associated with the attached node and not the medium access. This permits various logical link protocols to coexist on a common network without interfering with one another. Logical links are established primarily to ensure data integrity to the higher layers. For example, HDLC employs a frame-sequencing protocol for verifying that frames have successfully reached their destination, with the appropriate retransmission of lost frames [14].

Multiple appearances of LLC may exist within each node. In these instances, a link multiplex function within the LLC layer directs incoming frames to the appropriate LLC task and also provides the correct address information for outgoing frames. Each LLC appearance is logically associated with one other LLC appearance in another node. Models of four nodes are shown in Fig. 8, with examples of some of the logical links that may exist. Each link within the figure 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 is provided by the LLC sublayer of the DLC layer of the architecture.

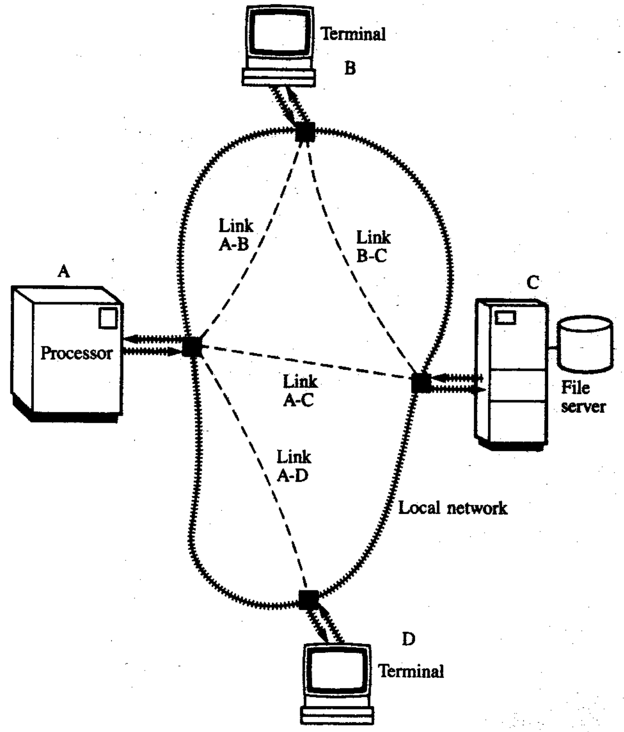

Figure 8 Logical link connectivity. The local network provides the physical transport mechanism for the transfer of data between nodes that have established a logical link connection.

• *Station manager*
The station manager function interfaces to both sublayers of the DLC layer, as well as to the higher layers of SNA. The overall function of the station manager can be partitioned according to the various levels of the architecture, such as PC manager or MAC manager functions. The MAC manager performs operations relative to fault diagnosis and statistics collection within the LAN that are described next.

Ring fault detection and isolation
The topological structure of a star/ring configuration, in conjunction with the token-access control protocol, provides an extremely reliable communication link [5]. The unidirectional propagation of information (electrical signals and data frames) from node to node provides a basis for detecting

certain types of network faults which can be categorized into two types: hard faults and soft faults.

A *hard* fault is defined as a complete disruption of the electrical signal path at some point on the ring. This disruption may be caused by a break in the wiring, either in the wiring lobe between a node and the wiring concentrator or in the wiring segment interconnecting two wiring concentrators. The failure of the receiver and/or transmitter of an active node can also cause a total disruption of the signal path. A hard fault is detected immediately by the next active node downstream from the fault as a loss of signal transitions at its receiver. The node that detects the presence of a fault then begins transmitting a unique MAC frame known as a *beacon*. Contained within the beacon frame is the address of the beaconing node as well as the last received address of the active node immediately upstream. (A special function for determining the address of the upstream node is discussed later.) Even though the ring is disrupted at one point, all of the active nodes, except the beaconing node, can receive and transmit normally. Thus, the beacon frame is received by all active nodes and may even be received by the node immediately adjacent to and upstream from the break. Recovery actions at the wiring concentrator might include removing from the ring the beaconing node or the node upstream from the beaconing node, or else switching to an alternate ring, as discussed later.

A *soft* fault is characterized by intermittent errors, usually caused by a degradation in the electrical signal or environmental electromagnetic interference. The frame check sequence (FCS) of all frames is calculated and verified by all intermediate nodes as the frames are repeated. The first node on the ring that detects an FCS error can set the error-detected indicator (EDI) bit in the physical trailer of the frame as an indication to all other nodes that the error has been detected. That node also logs the occurrence of the error by maintaining a count of the errors it has detected. If a predetermined threshold of FCS errors is reached over a given time interval at any one node, an indication of the condition can be reported to a special ring network management (RNM) function. The error report message that is sent to the ring network manager contains information that can be used in locating the source of the soft errors. The node detecting the FCS errors can also send a frame with the SEI bit set to *1* to all other nodes on the ring to inform them that a soft error condition exists on the ring.

The RNM node is a special function node within the local network that is capable of monitoring the network operation and configuration. For example, the RNM node could compile error statistics on the occurrence of FCS errors throughout the network or notify an operator of an excessive number of soft errors at a particular node [6]. With this information, the operator or a software application within the RNM node could initiate an appropriate ring reconfiguration action via a command to an active wiring concentrator to bypass a faulty ring segment or node. The RNM node can also maintain an up-to-date list of what nodes are active on any given ring within the network. Network management functions may be distributed over several RNM nodes in large networks if required.

As described earlier, the ARI bits are set by a node whenever it recognizes its own address within the destination field of that frame. If two nodes on a ring were inadvertently assigned the same address, one of the nodes would detect that the ARI bits had been set by the other node. This condition is reported as an error to the RNM node for action.

The isolation of a fault is enhanced if the node detecting and reporting it can also identify the address of the next active upstream node from the fault. This address information can be considered in conjunction with the address of the reporting node to expedite the isolation of a fault, such as was described in the beaconing and soft error procedures earlier. Thus, periodically, the active monitor issues a broadcast frame called a roll-call-poll MAC frame [1]. The first active node downstream from the monitor node sets the ARI bits and copies the source address. Other nodes on the ring do not copy the address information in this particular broadcast frame when the ARI bits are set. The node that received the roll-call-poll frame then issues a roll-call-repeat MAC frame containing its own source address whenever a free token is observed. This frame is recognized and copied by the next downstream active node. This process continues around the ring until the active monitor receives the roll-call-repeat MAC frame without the ARI bits set. At that time, each node has copied and saved the specific address of the adjacent node immediately upstream. This information is transmitted with all beacon frames and soft error report frames, thereby allowing the RNM node to log the general location of the fault.

Once the general location of a fault (hard or soft) has been determined, some action is necessary to eliminate the faulty segment(s) from the ring so that normal operation can resume. The wiring concentrators provide concentration points for bypassing such faults, as was discussed earlier with lobe bypass. A separate technique is necessary if reconfiguration is required as a result of faults occurring within the ring segments interconnecting wiring concentrators. Alternate backup links can be installed between the wiring concentrators in parallel with the principal links (Fig. 9). 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 restores the physical path of the ring. This wrapping function, like the lobe bypass function, may be activated manually or via automatic switching logic, or may be command initiated from an RNM node. The figure shows four wiring concentrators as they would be configured to bypass such a fault with both a principal and an alternate ring. The signals on the alternate ring are propagated in the direction opposite to those on the principal ring, thus maintaining the logical order of the nodes on the ring. 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 [1]. In addition to the procedure for bypassing faults, each node can perform self-diagnostic tests of its own circuitry and wiring lobe to ensure that it does not disrupt the signal path when it is inserted into the ring. If these tests indicate a potential problem, the node is not inserted into the ring until after the situation is remedied [6].

Summary

This tutorial has highlighted many of the key aspects of a token-ring local area network. The star/ring topology with token-access control offers a network architecture that can satisfy today's requirements while allowing for future growth and expansion. Wiring concentrators provide points throughout a building for a radial wiring scheme in which any node can be isolated from the network in case of failure. This concept forms the basis for a network management scheme that should provide quick isolation and recovery from both hard faults and soft faults. Token-ring networks are extendible in that additional token-controlled rings may be added to the network through the use of bridges. Furthermore, a hierarchical network structure can be implemented in a large building or a campus environment by interconnecting multiple bridges via a high-speed backbone. The backbone link may utilize the token-access control scheme on either another baseband ring or a channel within a broadband communication system (e.g., CATV). A token-ring LAN provides the basic transport mechanism for the transfer of data among nodes within the LAN. Higher-level communication protocols, such as those within SNA, must be implemented to ensure meaningful and reliable information exchange across the network.

Acknowledgments

The author wishes to acknowledge Roy C. Dixon, Don W. Andrews, and James D. Markov for their contributions and suggestions in preparing this tutorial, and Frank P. Corr and Edward H. Sussenguth for their technical review and comments. Several IBM researchers have contributed to the overall token-ring concepts presented here. The basic research within IBM that led to the development of this token-ring LAN architecture was conducted at the IBM Research Laboratory in Zurich, Switzerland.

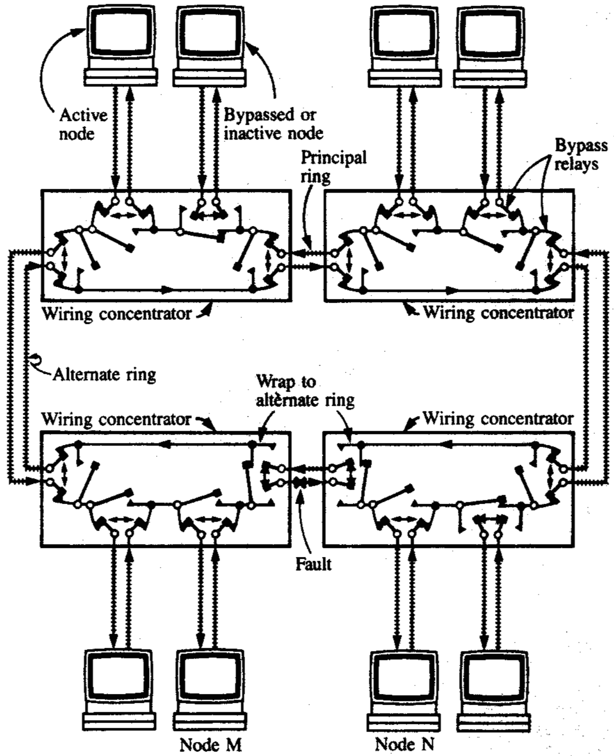

Figure 9 Ring reconfiguration with alternate ring. A fault on the ring between two wiring concentrators can be bypassed by wrapping the connections within the wiring concentrators, thereby using the alternate path and maintaining the same node ordering that existed before the fault occurred.

References

1. R. C. Dixon, N. C. Strole, and J. D. Markov, "A Token-Ring Network for Local Data Communications," *IBM Syst. J.* **22,** 47–62 (1983).
2. W. Bux, F. Closs, P. A. Janson, K. Kümmerle, and H. R. Müller, "A Reliable Token Ring System for Local-Area Communication," *Proceedings of the National Telecommunications Conference,* November 1981, pp. A.2.2.1–A.2.2.6.
3. W. Bux, F. Closs, P. A. Janson, K. Kümmerle, H. R. Müller, 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.
4. H. R. Müller, 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. R. C. Dixon, "Ring Network Topology for Local Data Communications," *Proc. COMPCON Fall '82,* Washington, DC, IEEE Catalog No. 82CH1796-2, pp. 591–605, available through the IEEE, 445 Hoes Lane, Piscataway, NJ 08854.
6. D. W. Andrews and G. D. Schultz, "A Token-Ring Architecture for Local-Area Networks—An Update," *Proc. COMPCON Fall '82,* Washington, DC, IEEE Catalog No. 82CH1796-2, pp. 615–624, available through the IEEE, 445 Hoes Lane, Piscataway, NJ 08854.
7. P. Abramson and F. E. Noel, "Matching the Media to Local Network Requirements," *Data Communications* **12,** No. 7, pp. 115–123 (July 1983).

8. J. G. Rusnak, "Local-Area Networking and Higher-Level Protocols: An SNA Example," *Contribution of Working Papers to IEEE Project 802 on Local Area Networks*, IBM CPD Publication, March 8, 1982; available through IBM branch offices.
9. J. D. Markov and N. C. Strole, "Token Ring Local Area Networks—A Perspective," *Proc. COMPCON Fall '82*, Washington, DC, IEEE Catalog No. 82CH1796-2, pp. 606–614, available through the IEEE, 445 Hoes Lane, Piscataway, NJ 08854.
10. *Systems Network Architecture: Concepts and Products*, Order No. GC30-3072, available through IBM branch offices.
11. David R. Jarema and Edward H. Sussenguth, "IBM Data Communications: A Quarter Century of Evolution and Progress," *IBM J. Res. Develop.* **25**, 391–404 (1981).
12. J. F. Keeley et al., "An Application-Oriented Multiprocessing System," *IBM Syst. J.* **6**, 78–133 (1967).
13. P. M. Karp, "Origin, Development and Current Status of the ARPA Network," *Proceedings of the Seventh Annual IEEE Computer Society International Conference*, San Francisco, 1973, pp. 49–52.
14. *Data Communication—High-Level Data Link Control Procedures—Frame Structure*, ISO Draft Proposal ISO/DP 3309 (Revision of ISO Standard ISO/3309-1976), October 1981; available through the American National Standards Institute, 1430 Broadway, New York, NY 10018.
15. D. D. Clark, K. T. Pogran, and D. P. Reed, "An Introduction to Local Area Networks," *Proc. IEEE* **66**, 1497–1517 (1978).
16. David J. Farber, Julian Feldman, Frank R. Heinrich, Marsha D. Hopwood, Kenneth C. Larson, Donald C. Loomis, and Lawrence A. Rowe, "The Distributed Computing System," *Proceedings of the Seventh Annual IEEE Computer Society International Conference*, San Francisco, 1973, pp. 31–34.
17. J. R. Pierce, "Network for Block Switching of Data," *Bell Syst. Tech. J.* **51**, 1133–1145 (1972).
18. 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.
19. M. V. Wilkes and R. M. Needham, "The Cambridge Model Distributed System," *ACM Oper. Syst. Rev.* **14**, 21–29 (1980).
20. *IBM Series/1 Local Communications Control, General Information*, Order No. GA27-3093, available through IBM branch offices.
21. R. M. Metcalfe and D. R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," *Commun. ACM* **19**, 395–404 (1976).
22. "A Survey of Digital Baseband Signaling Techniques," *NASA Technical Memorandum X-64615*, University of Tennessee, June 1971.
23. "Token-Passing Bus Access Method and Physical Layer Specifications," *IEEE Standard 802.4*, Draft D, Sections 14–15, December 1982; available from the Director of Standards, IEEE, 345 East 47th St., New York, NY 10017.
24. W. Bux, "Local-Area Subnetworks: A Performance Comparison," *IEEE Trans. Commun.* **29**, 1465–1473 (1981).
25. B. W. Stuck, "Calculating the Maximum Mean Data Rate in Local Area Networks," *Computer* **16**, No. 5, 72–76 (1983).
26. K. Kümmerle and M. Reiser, "Local-Area Communication Networks—An Overview," *J. Telecommun. Networks* **1**, 349–370 (1982).
27. D. I.-H. Park and R. D. Love, "An Analysis of the Tolerance to Crosstalk Noise of a Pulse Width Modulation System," *IBM J. Res. Develop.* **27**, 432–439 (1983, this issue).
28. A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM J. Res. Develop.* **27**, 440–451 (1983, this issue).
29. *IEEE Project 802, Local Network Standard, Draft C*, May 1982; available from the Director of Standards, IEEE, 345 East 47th St., New York, NY 10017.
30. J. H. Saltzer, D. P. Reed, and D. D. Clark, "Source Routing for Campus-Wide Internet Transport," *Local Networks for Computer Communications*, North-Holland Publishing Co., New York, 1981, pp. 1–23.
31. F. P. Corr and D. H. Neal, "SNA and Emerging International Standards," *IBM Syst. J.* **18**, 244–262 (1979).
32. *ISO Reference Model of Open-Systems Interconnection* ISO/TC97/SC16, DP 7498; available from the American National Standards Institute, 1430 Broadway, New York, NY 10018 (Rev., January 1983).
33. P. François and A. Potocki, "Some Methods for Providing OSI Transport in SNA," *IBM J. Res. Develop.* **27**, 452–463 (1983, this issue).
34. Standard ECMA-80-81-82, Local Area Networks: Coaxial Cable System—Physical Layer—Link Layer, September 1982.
35. K. N. Larson and W. R. Chestnut, "Adding Another Layer to the ISO Net Architecture Reduces Costs," *Data Communications*, 215–222 (1983).

Received March 29, 1983; revised June 28, 1983

Norman C. Strole *IBM Communication Products Division, P.O. Box 12195, 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. His research activities have included the development of a simulation methodology for the performance evaluation of shared-resource computer architectures. Dr. Strole received a B.S. in electrical engineering from North Carolina State University, Raleigh, in 1973. He received the M.S. and Ph.D. 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 and is a member of Eta Kappa Nu, the Institute of Electrical and Electronics Engineers, and Tau Beta Pi.