TOKEN RING LOCAL AREA NETWORKS

TOKEN RING LOCAL AREA NETWORK, TOKEN RING LAN

INTRODUCTION

Local area networks (*LAN*s) have evolved since the late 1970s to become central elements in corporate networks around the world. Many of today's LANs are based on one or more network technologies that originated as one of the early standards developed within the Institute of Electrical and Electronic Engineer's (*IEEE*'s) Project 802. The Token Ring Working Group, IEEE 802.5, which produced a family of token ring standards throughout the 1980s and 1990s, played a central role in the evolution of one highly popular LAN technology, particularly in commercial environments. Both the IEEE 802.5 and the ANSI X3T9.5 FDDI token ring standards were developed through the collaboration of many individuals representing a broad segment of the network industry. These token ring standards were subsequently endorsed as worldwide standards by the International Organization for Standardization (*ISO*).

BRIEF HISTORY OF RING NETWORKS

A token ring network is distinguished from other networks by a combination of network topology and an access method that allows hundreds of devices to reliably share the total available bandwidth. Token ring and the predecessor networks were based on evolving time division multiple access (*TDMA*) schemes where a common medium was used for both transmit and receive operations by tens or hundreds of devices. Researchers continued their quest for the ideal protocol that was both efficient and reliable. The loop systems of the 1970s were important precursors to ring systems. In both systems, network nodes are interconnected in a serial fashion forming a closed loop (or ring) on which encoded digital information flows in one direction. Loop and ring systems can be distinguished by their respective access control schemes. Loops operate in a master/slave fashion in which access to the medium is governed by a single master control node via the periodic issuance of a special control message known as the poll. Upon receipt of the poll, selected network nodes are permitted to send data to the master control node. In contrast, all nodes on a ring system are peers and autonomously determine when to transmit, based on the state of the ring. Early token ring prototypes also demonstrated that data transmission rates of 1 Mbps and greater were achievable, representing a significant advancement over the 56 kbps link rates that were prevalent at the time.

One of the first accounts of a ring-based communication system was presented by Farmer and Newhall (1); other significant ring networks include the distributed computing system (*DCS*) (2), the Pierce ring (3), the ring built at MIT (4), the Cambridge ring (5), and the ring network at the IBM Zurich Research Laboratory (6).

Both the Pierce ring and the Cambridge ring used a slotted-ring technique where multiple fixed-length data slots continuously circulate around the ring. Any node can place a data packet (or a packet fragment) in one of the empty slots, along with the appropriate address information. Each node examines the address information and copies those slot contents destined to that node.

In a second type of ring system, the buffer or register insertion ring, contention between data ready to be transmitted by a node and the data stream flowing on the ring is resolved by allowing the transmitting node to dynamically insert sufficient buffer space into the ring to avoid loss of data (7, 8).

A third scheme known as token-access control was first implemented in the DCS and the MIT rings and was the basis for the ring system built at the IBM Zurich Research Laboratory. It also underlies two important LAN standards, the IEEE 802.5 Token Ring (9) and the American National Standards Institute (*ANSI*) X3T9.5 Fibre Distributed Data Interface (*FDDI*) (10–12). In a token ring, access to the transmission medium is controlled by passing a unique digital signal, the permission token, around the ring. Each time the ring is initialized, a designated node generates a token that travels around the ring until a node with data to transmit captures the token and transmits its data. At the end of its transmission, the node passes the access opportunity to the next node downstream by generating a new token.

Standardization has played an important role in the evolution of the various LAN technologies. LAN standards have been developed primarily by the IEEE, the European Computer Manufacturers Association (*ECMA*), and ANSI. Token ring standardization was pursued by the IEEE 802.5 Committee, which produced its first standard in 1984. ANSI ratified it as an American National Standard in 1985 and forwarded it to ISO in 1985, which approved it as an International Standard (*IS*) in 1986. In Europe, ECMA issued its first token ring standard, ECMA-89, in 1982.

The IEEE 802.5 Token Ring standard was to a considerable extent based on contributions from IBM Corporation, which in the late 1970s, had investigated various LAN techniques at its Zurich Research Laboratory. The token ring topology and protocol was found to be particularly applicable to commercial applications, with several advantages over other LANs. Performance studies showed that the token protocol was more efficient as the network load increased and that a token ring was not subject to the same distance constraints as collision-based access protocols (13). IBM's research also represented an advancement beyond earlier ring concepts in an effort to define a system architecture that was reliable, easily deployed, and relatively simple to recover from errors or faults. One important innovation was the introduction of a central wiring concentrator unit and the star-ring topology. The robustness of the token ring was improved through the concept of a backup ring path that would permit self-healing when a break in the primary ring occurred. A method to ensure that there was always a token on the ring was introduced with the development of the token monitor concept. The design of a priority access mechanism provided the basis needed to support real-time applications such as voice (6,

J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*.

14).

The first network adapter and concentrator products to support the IEEE standard were shipped by IBM and Texas Instruments in 1985. Several other companies joined them to provide a wide range of token ring products throughout the 1990s. IEEE standard-based token ring deployment declined and, in many instances, was replaced by higher speed 100 Mbps and 1 Gbps switched Ethernet products in the early 2000s. However, the original IBM Cabling System remains in use as the transmission media for these 100 Mbps and 1 Gbps systems over 20 years after it was initially deployed for 4 Mbps LANs.

IEEE 802.5 Standards Evolution

IBM's technical contribution to the IEEE Project 802 committee in March 1982 formed the basis for the initial IEEE 802.5 Token-Ring Standard. This standard incorporates both the Physical (*PHY*) and Media Access Control (*MAC*) layers, which are Layers 1 and 2 of the Open Systems Interconnection (*OSI*) reference model (15). Token ring was initially standardized at the PHY layer as a 4 Mbps data transmit rate over 150 ohm shielded twisted pair (*STP*) cables known as the IBM Cabling System. Shortly after the issuance of that standard, subsequent releases expanded the cabling options to include 100 ohm unshielded twisted pair (*UTP*) cabling (i.e., telephone grade wire) and optical fiber. These were followed by standards for 16 Mbps token ring, first on 150 ohm STP cabling and optical fiber and later on 100 ohm UTP. Support of the 100 ohm UTP cabling required the introduction of active concentrators. The migration from a shared ring to a dedicated, switched link per station, known as dedicated token ring (*DTR*), provided a transition path for 16 Mbps operation as well as for 100 Mbps token ring operation. With each of these changes, the frame format remained the same so that any token ring formatted frame could be easily and economically bridged between token ring segments operating at different speeds.

TOKEN RING TECHNOLOGY

Basic Protocol

Information on a token ring is transferred sequentially from one node to the next. The token is a control signal composed of a unique signaling sequence which any node may capture (Fig. 1a). The node having control of the token and, thus sole access to the medium, transmits information onto the ring (Node A in Fig. 1b). For IEEE 802.5 operation, capturing the token is accomplished by simply modifying a single bit on-the-fly to form a start-of-frame sequence and then appending appropriate control information, address fields, the user information, frame-check sequence, and a frame-ending delimiter. All other nodes repeat, and thus redrive, each bit received. The addressed destination node copies the information from the ring as it passes (Node C in Fig. 1b). After completion of its information transfer and after checking for the correct return of its frame header, the sending node generates a new token that provides other nodes with the opportunity to gain access to the ring (Node A in Fig. 1c). The transmitting node keeps the ring open by transmitting idle characters until its complete frame has returned to be removed. The transmit opportunity passes with the token to all other nodes on the ring before a node can seize the token again to send additional data.

The maximum frame size is bounded by the maximum transmit time when a token is captured. The IEEE 802.5 standard defines this time as 9.1 milliseconds (e.g., 0.0091 seconds). As the number of 8-bit octets transmitted in a fixed time period is dependent on the data transmit rate, the upper limit becomes 4550 octets at 4 Mbps and 18,200 octets at 16 Mbps (16).

An important and unique characteristic of ring networks is that each node becomes an active participant in all ring communications, because each node must forward or retransmit the data signal to the next downstream node. This fundamental property is reflected in the wiring and transmission techniques that have been chosen for token rings, as described below.

Star Wiring

As information flows sequentially from node to node around the ring, a failure in a single node can disrupt the operation of the entire ring. This potential problem is addressed by the star-wiring topology in which each node is wired to a so-called wiring concentrator or multistation access unit, whereas the wiring concentrators are interconnected with point-to-point links (Fig. 2). Wiring lobes, consisting of two distinct send and receive signal paths, radiate from the wiring concentrators to the various network interface points, typically wall outlets, in a building. The wall outlets provide physical interfaces to the network to allow fast, reliable, and convenient attachment or relocation of workstations or servers.

The lobes are physically interconnected within the concentrators via electromechanical relays to form a serial link. A lobe is only included in the ring path when the node is active; otherwise, the bypass mechanism in the wiring concentrator causes that lobe to be skipped. If the bypass mechanism were positioned at the node itself, the inactive lobe would cause an undesirable increase in the distance between the active nodes on the ring. The wiring concentrators can be completely passive, i.e., contain 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 when a node needs to get inserted into the ring. The concentrator design and interconnect scheme also provides an alternative ring path that can be used to bypass a break or disruption in the primary ring path (Fig. 2).

As an alternative to the simple passive wiring concentrators, active concentrators or "hubs" later became very popular among users who required more stringent availability and manageability features. Active concentrators incorporate additional processing capability and contain either one or two complete token ring nodes; i.e., they represent addressable entities, which enable them to provide powerful management, security, and reconfiguration functions. Furthermore, active concentrators may be combined with a bridging function to other rings or a high-speed "uplink" to a LAN switch.

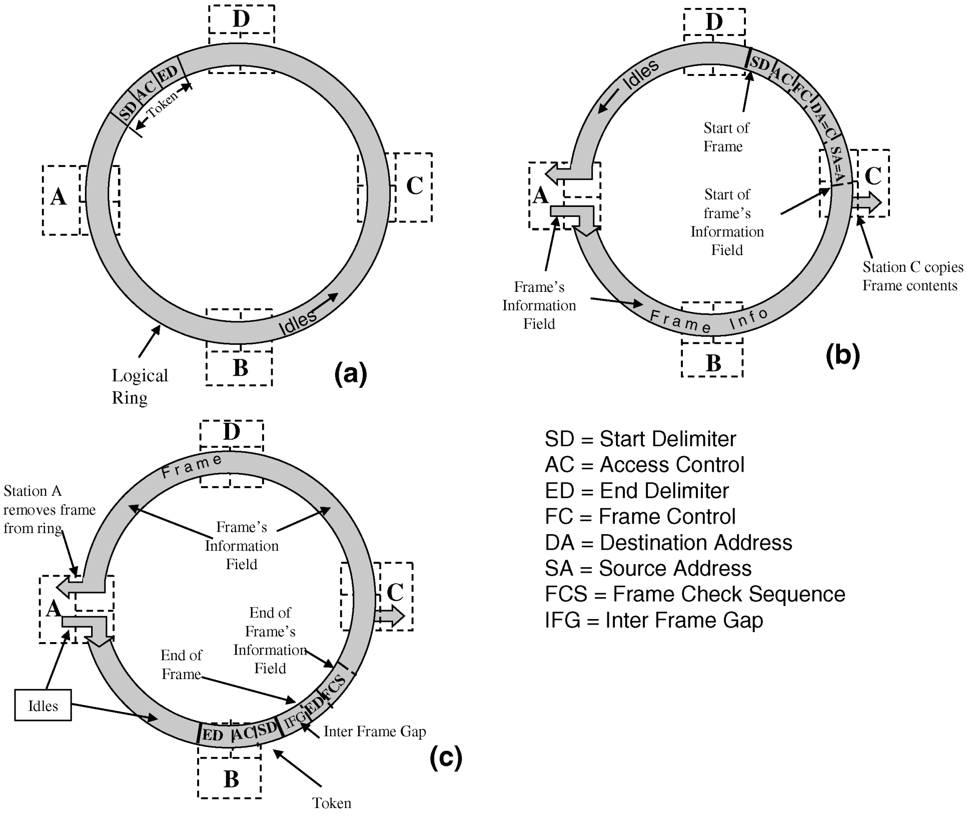

Figure 1. Token protocol overview.

When optical fibers are used to interconnect wiring concentrators and the nodes, insertion/removal signaling is accomplished via optical rather than electrical signals. The control information for ring insertion can be carried by special MAC control frames or unused code symbols; an alternative is to use out-of-band signaling with a suitable form of multiplexing. Wiring concentrators need to be active in this case.

Pre-cabling an office area or building with star wiring is practiced for most LAN installations and has several additional advantages:

1. It provides centralized points for reconfiguration management.
2. Workstations can easily be moved from one location to another without requiring installation of a new cable.
3. The wiring is segmented at the wiring concentrators rather than being a continuous cable, thus permitting the intermixing of transmission media. For example, twisted-pair wire can be used to interconnect the wiring concentrators and nodes, whereas optical fiber is employed for the transmission links between wiring concentrators.
4. As long as a node is in passive state (i.e., not inserted in the ring), its lobe is wrapped around in the wiring concentrator (Fig. 2), which enables the node to perform a self-check before inserting itself into the ring. Should an active node detect a fault within either its own components or its wiring lobe, it can remove itself from the ring.

Transmission Media

When token ring was initially under development in the early 1980s, it was thought that telephone-grade unshielded twisted-pair cabling was incapable of carrying high-speed (4 Mbps) signals for sufficient distances to be practical for commercial LAN applications. Therefore, a specially designed shielded twisted-pair cabling was simultaneously developed by IBM. In 1984 the IBM Cabling System using 150 ohm shielded twisted-pair copper media for deployment in a star topology was introduced. Commercial customers embraced the star-wired topology but were unsatisfied with the large diameter of the 150 ohm cabling and subsequent cable expense. In that same time frame, the telephony industry was looking at ways to use thinner telephone wiring for high-speed data transmission and realized that, to do so, telephone wire would have to be substantially improved. Those improvements had not reached the market when the first token ring products became available in 1985. At that time, complex wiring rules were published for robust operation of the token ring protocol over large networks spanning multiple wiring closets and containing up to 260 devices (See tables 2.3–2.8 on pp 64–69 of Reference 15). As a result of customer demand to use existing telephone wires, alternative rules were also published for more modest networks of up to 100 nodes, all cabled to the same wiring closet (16). Minimum requirements were placed on the telephone wire, which was dubbed “Type 3 media.” During this time frame, the telephony industry was also developing and standardizing improved telephone wiring and connectors for the express purpose of carrying high-speed LAN data for both token ring and Ethernet applications, first in North America under the auspices of TIA, and later internationally, under JTC/1 SC25/WG3. The North American Standard for telecommunication grade cabling was TIA/EIA 568. The first edition of that standard included the 150 ohm cabling of the IBM Cabling System and a more rigorous standard for data grade telephone wire called Category 3 cabling. Later editions of that standard specified a Category 5 (and later Category 5e) telephone wiring, which was crucial for supporting token ring's next generation of hardware operating at 16 Mbps and later at 100 Mbps. The international standard, IS11801, specified the 150 ohm IBM data grade cable as well as both Category 3 and Category 5 twisted-pair cabling. Category 5 (and later Category 5e) cabling was more advanced, with better transmission characteristics and became the de facto twisted-pair data cable. As token ring increased its operating speed, first from 4 to 16 Mbps and then to 100 Mbps, the requirement to operate over Category 5 cabling at distances of up to 100 m had to be addressed. The solution at 16 Mbps was to require the concentrators to be active devices, repeating the signal between each pair of wiring concentrators but not between the nodes themselves. The solution at 100 Mbps was to use a more efficient line coding and to operate in dedicated token ring (*DTR*) mode where each signal received by the concentrator from any of its attachment ports was regenerated and retransmitted to the next lobe. DTR design required one active device in the concentrator for each end station that attached to it. These topology and signal reclocking changes were required to overcome the signal attenuation that occurs over the copper media as the signal clock rate increases, thus decreasing the maximum distance before the signal must be reclocked.

Although two principle choices existed for cabling from the wiring closet to the active devices in the offices and within the wiring closet, the requirement to be able to transmit over multi-100 m lengths between wiring closets and between buildings was addressed with optical fiber transmission media. Although optical fiber media could have been used to attach devices located in offices, it never gained significant market share for that application because of attachment and media cost.

Advances in signal processing technology have allowed transmission rates to increase from 100 Mbps to over 1000 Mbps while maintaining a clock rate that allows the 100 m length to be maintained. Subsequent advances in the copper cable design and digital encoding schemes since the 1990s have resulted in transmission capacities of up to 10 Gbps over short distances (up to 15 m), with the promise of longer distances in the future. Much of this increase in transmission capacity over copper media can be attributed to the advancements in Application Specific Integrated Circuit (*ASIC*) technology and digital encoding schemes.

Line Coding

The data generated by a node must be encoded for transmission over a ring network. The IEEE 802.5 standard specifies differential Manchester encoding for both the 4 and 16 Mb/s token ring transmission rates (Fig. 3a). Differential Manchester encoding is characterized by the transmission of two line signal elements per bit, which results in a link clock rate that is double the bit transmission rate. In the case of a binary one or zero, a signal element of one polarity is transmitted for one half-bit time, followed by the transmission of a signal element of the opposite polarity. This line coding has two advantages: 1) The resulting signal has no dc component and can readily be inductively coupled, and 2) the mid-bit transition conveys inherent timing information. The ones are different from the zeros at the leading bit boundary; a value of one has no signal transition at the bit boundary, whereas a value of zero does. In decoding the signal, only the presence or absence of the signal transition and not the actual polarity is detected; thus, interchanging the two wires of the twisted pair introduces no data errors. A code violation results if no signal transition occurs at the half-bit position. Code violations can be intentionally created to form a unique non-data signal pattern that can be distinguished from normal data (Fig. 3b). These so-called J/K signals can be inserted to mark the start or end of a valid data frame. The J/K code violations are used in pairs to maintain the dc balance of the Manchester signaling.

Manchester coding is a very simple and robust technology but, at the same time, is also very bandwidth-inefficient and therefore not suitable for transmission rates significantly higher than 16 Mbps. For example, in FDDI, information on the medium is transmitted in a 4-out-of-5 group code (4B/5B) with each 5-bit code group, called a symbol, used to represent 4 bits of data. The symbols are transmitted in a non-return to zero inverted (*NRZI*) line transmission format (11). NRZI is distinguished from Manchester in that

1. There are no transitions at the half-bit boundary.
2. Transitions occur at the beginning of a binary '1'.
3. No transition occurs at the beginning of a binary '0'
4. The 4B/5B encoding scheme has excess code groups that can be used as non-data symbols to distinguish the start and end of a valid data frame.

Synchronization

Synchronization of the link clocking among stations is a key technical problem in the design of any ring system. Rings built according to the IEEE 802.5 standard employ a centralized clocking technique. In normal operation, one node on the ring is automatically designated as the active monitor during ring initialization. The monitor plays a crucial role in the supervision of the ring, as will be described below. In addition, it provides the ring master clock. All other nodes on the ring are frequency and phase-locked to the monitor. They extract timing from the received data by means of a phase-locked loop while redriving the digital signal to reach the next node. Each port of an active concentrator can also reclock and redrive the signal as well.

Although the mean transmission rate on the ring is controlled by the active monitor node, segments of the ring can, instantaneously, operate at rates slightly different from the

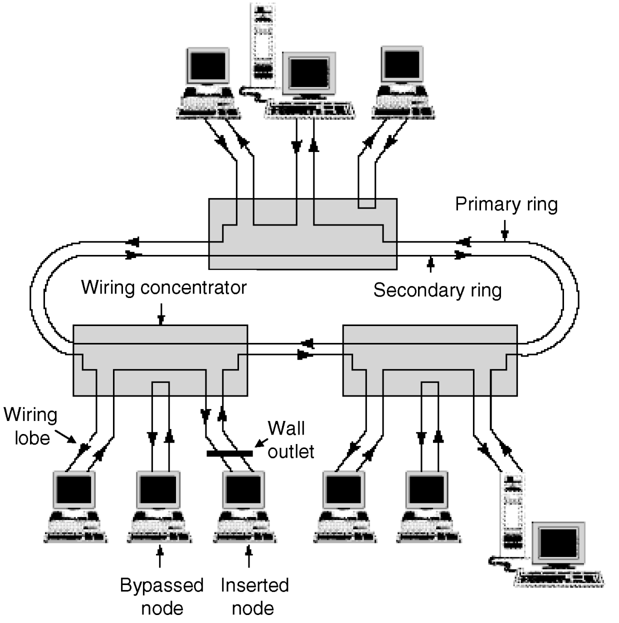

Figure 2. Star-wiring topology with dual-ring example.

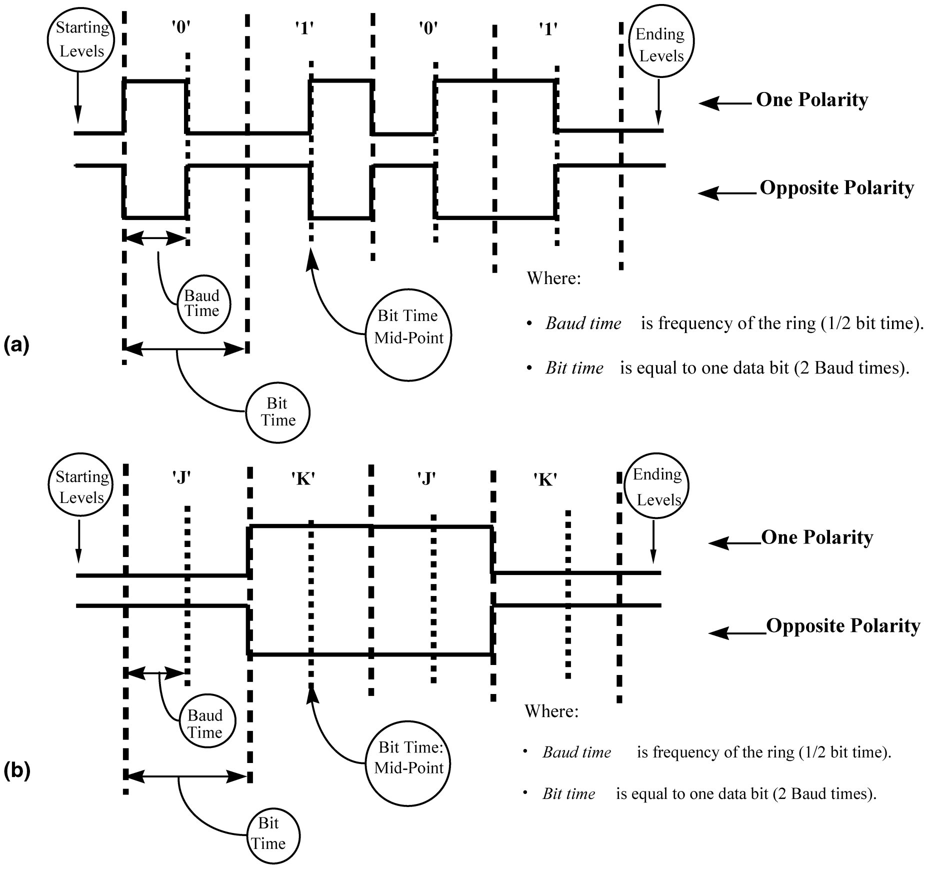

Where:

- *Baud time* is frequency of the ring (1/2 bit time).
- *Bit time* is equal to one data bit (2 Baud times).

Where:

- *Baud time* is frequency of the ring (1/2 bit time).
- *Bit time* is equal to one data bit (2 Baud times).

Figure 3. Differential Manchester encoding.

frequency of the master oscillator. The cumulative effect of such rate variations is sufficient to cause variations of a few bits in the latency of a ring. Unless the latency of the ring remains constant, bits would have to be either dropped or added. To maintain a constant ring latency, an elastic buffer is provided in the monitor. If the received signal at the active monitor node is slightly faster than the master oscillator, the buffer will fill up to avoid dropping bits. If the received signal is slow, the buffer will be emptied to avoid adding bits to the repeated bit stream. Detailed discussion and analysis of this clocking scheme are given in References 16 and 17.

A major advantage of the centralized clocking approach is that it minimizes the total latency of the ring and thus allows use of the IEEE 802.5 protocol, which for optimum performance, requires the ring latency to be small. An alternative synchronization technique that introduces greater latency but is easier to implement at high transmission rates is employed in FDDI, where information is transmitted between nodes asynchronously; that is, each node uses its own autonomous clock source to transmit or repeat information on the ring (11). This type of operation requires the use of an elastic buffer at every node. Information is clocked into the buffer at the clock rate recovered from the incoming bit stream, but it is clocked out of the buffer at the autonomous fixed clock rate of the node. A preamble that precedes each frame enables the buffer to be reset to its midpoint before frame reception. The reset operation increases, or decreases, the length of the preamble. For the IEEE 802.5 100 Mbps token ring operation, the issue of ring latency was sidestepped by defining only a switched token ring operation with one active repeater node in the wiring closet for each attached station. With this configuration, the attached station and its active repeater form a two-station ring. The active repeater acts like a bridge sending information onto this two station ring and broadcasting information received from the attached station to the other direct-attached stations.

TOKEN RING ACCESS PROTOCOL, MONITORING, AND RECOVERY

Data Frame Format

The token format and the general format for transmitting information on the ring, called a frame, are shown in Fig. 4. A token contains only the access control (*AC*) subfield and the starting and ending delimiters (Fig. 4a). The one-byte AC field includes a one-bit token (*T*) indicator that indicates whether this is a token (0) or a frame (1). A token priority mode that uses the priority reservation indicators provides different priority levels of access to the ring (see below). The monitor (*M*) bit is used in connection with the token monitor function to maintain the validity of the token.

The "data" portion of the frame is variable in length and contains the information that the sender is transferring to the receiver. The information (*INFO*) field is preceded by a header, which contains several subfields (Fig. 4b). The first is a starting delimiter (*SD*) 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 (Fig. 3b). Next, the AC subfield, with the token bit (T) set to 1, is defined for controlling transmit access to the shared media as described above. The frame control (*FC*) subfield contains a two-bit frame format (*FF*) and a three-bit frame priority subfield. The frame format enables the receiving node to determine whether the information within the data field of the frame contains media access control (*MAC*) information (FF=00) or user data (FF=01). MAC frames may optionally include frame status information within the control indicator subfield. Finally, the header includes the source address (*SA*) of the node that originated the information and the destination address (*DA*) of the node (or nodes) destined to receive the information. Both address fields contain six bytes, with the first two bits of the DA designating that the address is intended for multiple destination nodes (Group bit) or that the address has been assigned by the user (Local Administered Address bit). Use of the routing information field (*RIF*) will be described in the section on multiring networks.

The information field is followed by a trailer that is 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 by downstream nodes for detecting bit errors that may occur within the frame during transmission bounded by the FC subfield and the last bit of the information field. Next, an ending delimiter (*ED*) is provided to identify the end of the frame. This delimiter also contains a unique, although slightly different, bit combination along with pairs of code violations as were found in the starting delimiter. The last bit of the ending delimiter is designated as the error-detected indicator (*EDI*). This indicator will always be zero during error-free ring operation. The ending delimiter is followed by a one-byte frame status (*FS*) field. The FS field contains bits that can be modified while the frame is traversing the ring by nodes that match the destination address and/or copy the frame. The FS field is therefore not included in the calculation of the FCS character. For this reason, these bits are defined as pairs to avoid erroneous conditions caused by single-bit errors on the wire.

Priority Protocol

In some applications, it may be necessary for selected nodes to gain priority access to the ring. A priority scheme was designed specifically for the token ring protocol that was initially one of the distinguishing features versus other access control schemes. The priority (PPP) and reservation (RRR) indicators in the AC field are used to facilitate this access scheme (Fig. 4b). Various nodes may be assigned priority levels for gaining access to the ring, with the lowest priority being '000' and the highest being '111'. This assignment allows up to eight protocol levels to be defined. A selected node can seize any token that has a priority setting (Bits 0–3) equal to or lower than its assigned priority. The requesting node can set its priority request in the AC reservation field (Bits 5–7) of a frame as it is being repeated if that node's priority is higher than any current reservation request. The current transmitting node must examine the reservation request in the returning frame and release the next token with the new priority indication (Bits 0–3) but retain the previous priority level within its MAC state information for later release. A requesting node uses the priority token and releases the 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 token recognizes a token at that priority, it then releases a new token at the level that was interrupted by the original request. Thus,

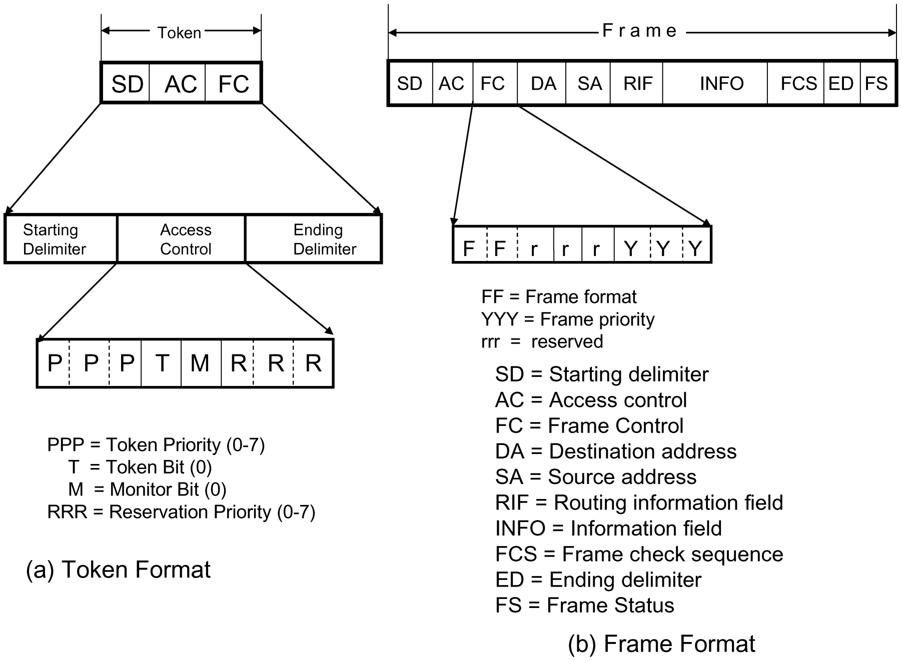

(a) Token Format

(b) Frame Format

Figure 4. Token and frame formats.

the lower priority token resumes circulation at the point of interruption.

In 1995, the IEEE 802.5 standard committee published a set of guidelines for the use of previously reserved priority levels 5 and 6. Priority 5 is recommended for delay-sensitive, high-bandwidth data streams such as video applications. Priority 6 is recommended for delay-sensitive, low-bandwidth data streams, such as interactive voice communication. Priority 7 is reserved for ring management and error recovery frames. Priority 4 is generally recognized for bridge access.

Token Protocol Performance Issues

From a performance point of view, token rings have two distinct advantages over other access protocols:

1. From the cyclic operation (also sometimes referred to as "round-robin") enforced by the rotating token, all users of the ring are serviced in a perfectly fair fashion within a given priority class. The priority mechanism, however, may be used to give a subset of the user's preferential service as described above.
2. From its deterministic behavior, the token protocol scales better with respect to network latency than random access protocols such as CSMA/CD.

However, the original IEEE 802.5 protocol is not totally insensitive to ring latency because it requires that idle characters be inserted by a transmitting node until it has recognized its source address in the header of the returning frame. This requirement leads to improved error robustness of the operation and is necessary for a higher priority token to be released, but it results in decreased efficiency as the physical ring length, number of active nodes, and/or the ring speed increase.

Ring protocol efficiency can be maximized at very high speeds and/or long ring lengths if nodes release the token immediately after finishing frame transmission. In 1987, the IEEE 802.5 standard was enhanced by an early token release (*ETR*) option that allows a transmitting node to release a new token as soon as it has completed frame transmission, whether or not the frame header has returned to that node. This enhancement was necessary to allow the support of a 16 Mbps operation over large campus networks. One impact of ETR is that priority reservations applied to "short" frames are lost until the circulation of a subsequent frame that exceeds the ring length. This trade-off was shown to be acceptable for typical campus-wide 16 Mbps rings. A detailed discussion of token ring performance issues can be found in Reference 18.

Ring Monitor Function

In token ring networks, error detection and recovery mechanisms are provided to restore network operation in the event that transmission errors or medium transients, for example, those resulting from node insertion or removal, cause the ring to deviate from normal operation. The IEEE 802.5 token rings use a network monitoring function that is performed in a specific token monitor node with back-up capability in all other nodes attached to the ring. The monitoring function is based on the scheme developed by the IBM research team in Zurich (14). Through an arbitration process, the nodes on the ring select one node to be the active monitor. As described, this node also provides the master clock for the ring. The remaining nodes function as standby monitors. The active monitor keeps watch over the health of ring and token, activating recovery procedures when necessary.

Ring errors can be quickly isolated to a specific ring segment if an accurate map of the ring stations is maintained. Periodically, the active monitor will issue a broadcast frame called an active monitor present (*AMP*) frame. The first active node downstream from the monitor node will set the address recognized indicator (*ARI*) bits in the FS subfield and save the source address. Other nodes on the ring will ignore this particular broadcast frame when the ARI bits are set. The node that received the AMP frame will then issue a standby monitor present (*SMP*) frame containing its own source address whenever a 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 SMP frame without the address-recognized flag set. At that time, each node will have the specific address of the adjacent node immediately upstream, which is known as the nearest active upstream neighbor (*NAUN*). The NAUN information is transmitted with all beacon frames and soft error report frames, thereby allowing a network management node to log the logical location of the fault. The AMP and SMP frames are transmitted at the highest ring priority to ensure that the process completes in the least amount of time, even during periods of peak ring utilization.

Ring Fault Detection and Isolation

The topological structure of a star-ring configuration, in conjunction with the token-access control protocol, permitted the development of additional protocols for rapid detection and isolation of network faults (15, 16). The unidirectional propagation of information (electrical signals and data frames) from node to node provides a basis for detecting certain types of network faults. Network faults can be categorized into two types: hard faults and soft faults.

A hard fault occurs when there is a complete break in the ring wiring between two adjacent nodes or wiring concentrators or a failure in the transmitter or receiver elements of a node. A node that detects loss of signal at its receiver will begin transmitting a unique series of contiguous MAC 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 and downstream of the fault will exit the beacon state as they begin receiving beacon frames from their upstream neighbors. Thus, the location of the fault will be isolated to the particular ring segment and the last known NAUN that is immediately upstream from the node that is transmitting beacon type frames.

A soft fault is characterized by a high frame error rate, 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 sets the error detected indicator (*EDI*) in the ending delimiter field as an indication to all other nodes that the error has been logged. If a predetermined threshold of FCS errors is reached over a given time interval, an indication of the condition can be reported to a network management application. The location of the soft fault can be readily determined from the information in the error report message and isolated to the ring segment immediately upstream of the reporting node.

Once the location of a fault (hard or soft) has been determined, several options are available for eliminating 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. Also, alternative backup links are normally available between the wiring concentrators in parallel with the principal links. If a fault occurs in the ring segment between two wiring concentrators or if a concentrator failure occurs, wrapping of the principal ring to the alternative ring within the two wiring concentrators on either side of the fault will restore the physical path of the ring (Fig. 5). This wrapping function, like the lobe bypass function, is automatic in active concentrators. Figure 5 shows four wiring concentrators as they would be configured with both a principal and an alternative ring. The signals on the alternative ring are propagated in the direction opposite to those on the principal ring, thus maintaining the logical order of the nodes on the ring.

MULTI-RING NETWORKS

Multiple rings are required in a campus or building LAN when the aggregate data transfer requirements or total number of stations exceed the capacity of a single ring or when a large number of attached nodes is spread over a broad area (15). Two rings can be linked together by a high-speed interconnect mechanism known as a bridge (Fig. 6). A bridge is capable of providing a logical forwarding of frames between the rings based on the SA, DA, and/or RIF inserted by the source node. 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 could therefore stand alone in the event the bridge or another ring was to be disrupted. The bridge interface to a ring is the same as any other node, except that it must recognize and copy frames with a destination address or RIF subfield 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, which results in a hierarchical network in which multiple rings are interconnected via bridges and multiple bridges are interconnected via a separate high-speed link known as a backbone (Fig. 6). The backbone may be a high-speed ring, such as FDDI, or it may be another network type, such as an asynchronous transfer mode (*ATM*) network.

Most token ring and FDDI network devices support a MAC-level bridging scheme known as source route bridging (*SRB*). With this scheme, intermediate bridges, switches, or routers and the associated ring segments that form the path between a source and a destination node are uniquely and explicitly identified within the RIF within the frame header (Fig. 4b). The RIF is created via a discovery protocol at the beginning of a session that allows the source node to designate a unique path to the destination

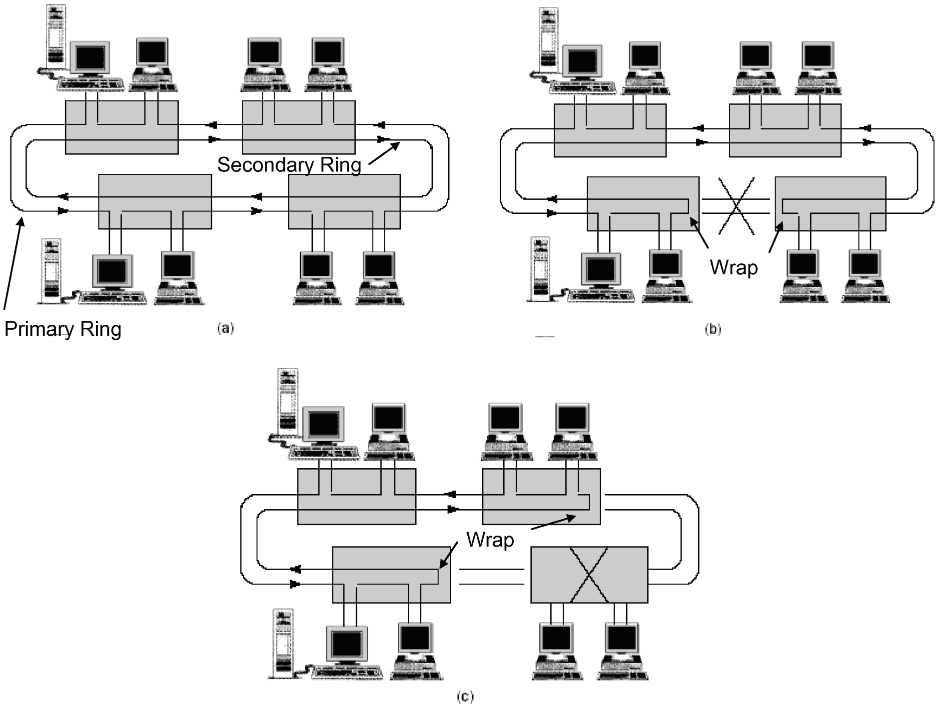

Figure 5. Fault detection and isolation with wiring concentrators.

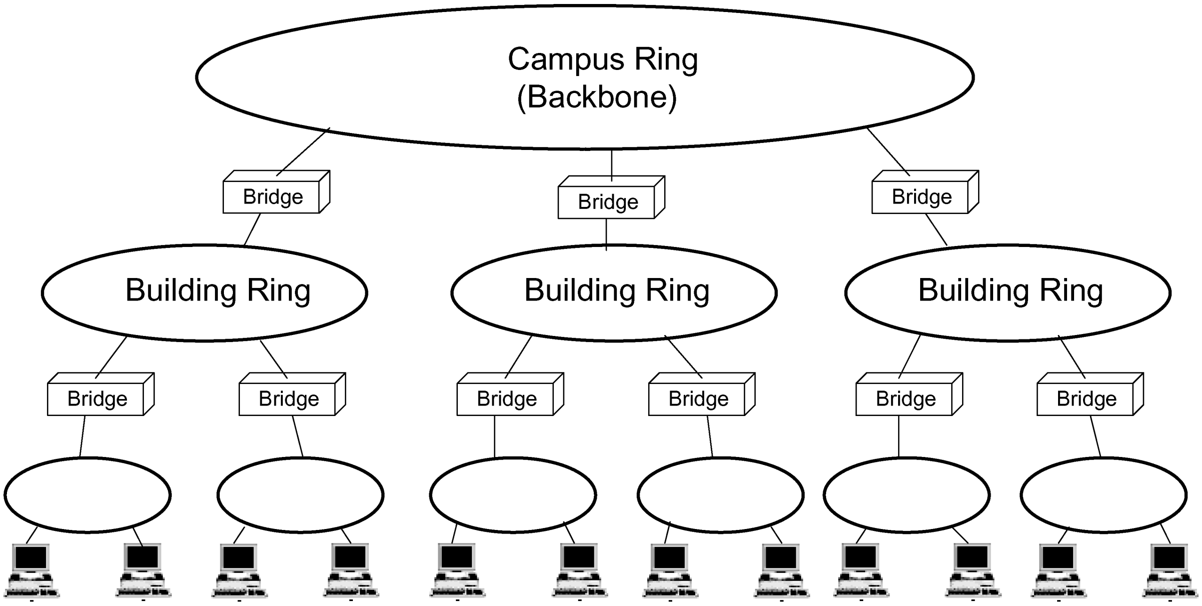

Figure 6. Multi-ring network topology.

node, enumerated as a sequence of bridge and ring segment IDs (16). This scheme simplifies the bridge processing that is required at each intermediate device, while also providing a mechanism that allows multiple active data paths between two points of the network.

The IEEE 802.1 committee developed standards for LAN bridging. This committee developed an alternative scheme, known as transparent bridging (*TB*), that required the bridge devices to create and maintain bridge tables to determine which frames to forward (or drop). This scheme was more applicable to the existing IEEE 802.3 Ethernet standard-based products, without requiring changes to the existing adapter hardware and could also be applied to token rings as well. The IEEE 802.1 bridge standard incorporated the criteria for a combination TB/SRB bridge a few years later.

Traditional bridged networks were gradually replaced by switched networks beginning in 1995 (19). Fewer nodes per ring segment allow individual nodes access to more bandwidth. Dedicated switched links allow one node to use all available bandwidth without contention with other nodes.

DEDICATED ("SWITCHED") TOKEN RINGS

As long as a token ring is operated as a shared medium, the total transmission capacity available to all users can obviously not exceed the ring's transmission rate. FDDI extended token ring speed to 100 Mbps by using a different token-based media access protocol and different frame format in a separate standard effort that was completed in the late 1980s (10–12). The IEEE 802.5 protocol could have been extended to 100 Mbps, but this was not considered to be a commercially viable option with the completion of the FDDI standard.

Overcoming the limitations of a shared-media protocol required the introduction of a high-speed switching function that became technically and economically possible in the early 1990s with the advancement of ASIC technology. In 1993, the IEEE 802.5 standard committee began looking at options to extend the token ring standard to meet the demand for additional bandwidth. As a result, the DTR standard was completed in 1997. DTR increased the number of ring segments with the introduction of the DTR Concentrator by allowing a ring segment to contain one or more stations supported by one active node in the wiring concentrator, and it introduced the concept of full-duplex operation for directly attached stations (15).

One catalyst for the DTR effort was to leverage the beacon transmit mode that was already present in the hardware design of millions of token ring adapters. In this mode, the node adapter is simultaneously transmitting the beacon frame and receiving frames on the inbound side in order to determine whether it is the station nearest to the fault. The existing token ring adapter firmware was modified to create the new full-duplex mode defined by the DTR standard, thus allowing existing adapters to migrate to the new mode with a firmware update combined with the introduction of a multi-port token ring packet switch to replace the classic wiring concentrator. DTR also allowed a node to be the only station on a ring shared with the DTR port as the other station. In this case, no change to the station interface is required. The DTR standard and technology was expanded in 1998 to enable 100 Mbps token ring to use Category 5 data cabling, which had become standard in many commercial businesses.

With DTR, a token ring node is allocated the full bandwidth via a dedicated segment between the node and the DTR concentrator (Fig. 7) (16). A new mode, full-duplex operation, is also supported. With the full-duplex mode, the token is no longer required. Instead, two dedicated parallel paths are established between the two nodes. For 100 Mbps token ring, up to 200 Mbps of data transfer (100 transmit and 100 receive) can be achieved per link. For 16 Mbps operation, up to 32 Mbps of data transfer (16 transmit, 16 receive) can be achieved per link. Data frames are forwarded among the dedicated segments by a high-speed data transfer unit within a DTR (Fig. 7) or a packet switch (19). For commercial applications, the token ring switch uses the existing RIF within the frame header to accelerate the packet forwarding. The effective aggregate bandwidth of the DTR system is determined by the switch capacity rather than by the clock speed of the shared media, thus providing much greater bandwidth than the shared-media configuration. Devices attached to a dedicated link have access to the full-duplex bandwidth, thereby providing significantly more application growth potential than with a shared-media access control scheme.

FDDI TOKEN PROTOCOL

A discussion of token ring would not be complete without a more detailed discussion of the FDDI token protocol, which is significantly different from the IEEE 802.5 operation in several fundamental areas (20). Unlike the IEEE 802.5 standard, FDDI defines two classes of data traffic: synchronous and asynchronous (11). The synchronous class is applicable to traffic that requires regular intervals between consecutive frame transmissions, such as real-time voice or video, for example. Synchronous traffic is given the highest priority, and the protocol is designed to guarantee frames of this class a transmit opportunity on each revolution of the token within predetermined bounds on the transmit intervals. Stations requiring synchronous access are assigned reserved bandwidth in advance via a distributed control scheme. Asynchronous frames have up to eight priority levels or thresholds but with no guarantee of access. The protocol allocates asynchronous bandwidth based on the priority after synchronous demand has been satisfied.

FDDI Token Timers and Operation

Unlike the IEEE 802.5 protocol, control of the FDDI ring under normal operation is decentralized; i.e., there is no master station. The algorithm that each station executes allocates use of the ring based on a fixed value that is the same for all MAC entities on the ring and on the contents of two timers present in every MAC (11). The fixed value is the target token rotation time (*TTRT*), and the timers are the token rotation timer (*TRT*) and the token hold timer (*THT*).

As the network load increases, the TTRT defines the average time for the token to complete one rotation around the ring, which in turn determines the response time that the network's users need for their synchronous traffic. The stations determine the value for TTRT during ring initialization. The FDDI protocol guarantees that the maximum token wait time for any station on the ring will never exceed two times the TTRT value.

One of the timers present in every station's MAC entity is the TRT. In conjunction with a counter called Late_Ct (the "late counter"), it indicates the amount of time that has elapsed since the station last received a token. By examining its TRT and Late_Ct, a station knows whether the token is taking more or less than the TTRT to complete a rotation. Stations can transmit asynchronous traffic only if the token is received when the Late_Ct is zero.

The Late_Ct is set to zero by the MAC each time the token is received, and the TRT is initialized to the TTRT value each time the token is received early. TRT is a decrementing counter that measures the time required for the token to circulate around the ring. If the TRT expires before the token returns, the Late_Ct is incremented and TRT is reset to TTRT and continues to decrement. A late to-

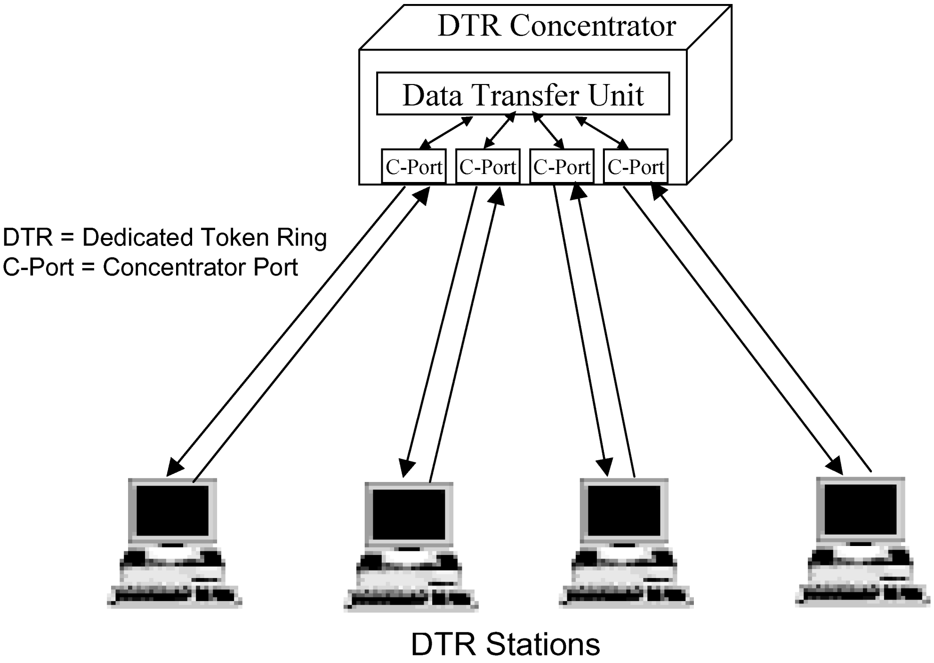

Figure 7. Dedicated token ring example.

ken, or one that arrives when Late_Ct = 1, does not reset the TRT but allows it to continue decrementing, thus carrying forward the lateness of the current token rotation into the next token rotation time. This process may restrict a station's ability to transmit asynchronous data frames for multiple successive token rotations. If the TRT expires again while Late_Ct = 1, an error condition recovery procedure is initiated by that station.

The second timer used in bandwidth allocation is the THT. Each THT indicates the amount of time that the MAC may use for asynchronous frame transmission. The value of THT for each station will vary from one token revolution to the next, depending on the network load.

A station may use a "late" token only for synchronous transmission, because the token has taken more than TTRT to complete a rotation. However, if Late_Ct equals zero when the station receives the token, the station may transmit asynchronous frames as well. In this case, the THT will determine the amount of time the station may transmit asynchronous frames. A station transmits all pending synchronous frames first. The time required for synchronous transmission has already been factored into the TTRT value and is thus not subject to THT limits. THT is initialized as the residual value of the TRT (e.g., as the difference between TTRT and the amount of time that the early token took to rotate). THT is decremented by the MAC only during the transmission of asynchronous frames. A station may transmit multiple asynchronous frames until the THT expires.

The FDDI priority scheme is based on an array of priority threshold values called T_Pri. These values indicate the length of time that the station may transmit frames at a given priority. A station can only begin transmitting frames at a given priority if the remaining THT is greater than the threshold value for that priority. Thus, under elevated ring loads, it is possible that a given station will be allowed to transmit synchronous frames and only a few of the higher priority asynchronous frames, but it will then need to wait additional token rotations to transmit the lower priority asynchronous frames. Dykeman and Bux (21) provide an in-depth analysis of the FDDI token-access and priority schemes.

Additional details regarding FDDI operation can be found in referenced FDDI standard documentation (10–12).

LAN EVOLUTION

The full-duplex, star configuration continued to evolve from the early 1990s, but the basic principals remained the same. The current generation of LAN switches provides in excess of 100 Gbps internal switch capacity, with the Ethernet packet format being the most widely deployed. These advances are enabled by high-speed ASIC switch technology that allows high-density 1 Gbps and 10 Gbps ports on a single chip. Fiber media are also more pervasive now than in the past, and the fiber connector technology has improved significantly. As pointed out in the Transmission Media section, dedicated link speeds of 1 and 10 Gbps are possible today on high-quality copper media as well.

CONCLUSION

This artile provides both a historical perspective and an in-depth technical review of the IEEE and ANSI token ring LAN protocols that emerged in the 1980s. Interested readers are encouraged to refer to Reference 16 as the most comprehensive source on the IEEE 802.5 token ring that is still available today.

BIBLIOGRAPHY

1. Farmer, W. D.; Newhall, E. E. An Experimental Distributed Switching System to Handle Bursty Computer Traffic; *Proc. ACM Symposium on Problems in the Optimization of Data Communications*; Pine Mountain, GA, 1963, pp 31–34.
2. Farber, D. J., et al. The Distributed Computing System; *Proc. of COMPCON '73*; 1973, pp 31–34.
3. Pierce, J. R. Network for Block Switching of Data. *Bell Syst. Tech. J.* 1972, **51**, pp 1133–1145.
4. Saltzer, J. H.; Pogran, K. T. A Star-Shaped Ring Network with High Maintainability; *Proc. of the Local Area Communications Network Symposium*; Boston, MA, 1979, pp 179–189.
5. Hopper, A. Data Ring at Computer Laboratory, University of Cambridge. In *Computer Science and Technology: Local Area Networking*. Nat. Bur. Stand., NBS Special Pub.: Washington, D. C., 500-31, 1977;pp 11–16.
6. Bux, W.; Closs, F.; Janson, P. A.; Kümmerle, K.; Müller, H. R.; Rothauser, E. H. A Reliable Token Ring System for Local-Area Communication; *Conf. Rec. NTC '81*; Piscataway, NJ, 1981, pp A.2.2.1–A.2.2.6.
7. Hafner, E. R.; Nenadal, Z.; Tschanz, M. A Digital Loop Communications System. *IEEE Trans. Commun.* 1974, **COM-22**, pp 877–881.
8. Reames, C. C.; Liu, M. T. A Loop Network for Simultaneous Transmission of Variable Length Messages; *Proc. 2nd Annual Symposium Comput. Architecture*; Houston, TX, 1975, pp 7–12.
9. ISO/IEC 8802-5:1998. Part 5. Token Ring Access Method and Physical Layer Specifications. www.iso.org.
10. ISO 9314-2:1989. Information Processing Systems–Fibre Distributed Data Interface (FDDI) – Part 2: Token Ring Media Access Control (MAC). www.iso.org.
11. ISO 9314-1:1989. Information Processing Systems–Fibre Distributed Data Interface (FDDI) – Part 1: Token Ring Physical Layer Protocol (PHY). www.iso.org.
12. ISO/IEC 9314-3:1990. Information Processing Systems–Fibre Distributed Data Interface (FDDI) – Part 3: Physical Layer Medium Dependent (PMD). www.iso.org.
13. Bux, W. Local-Area Subnetworks: A Performance Comparison. *IEEE Trans. Commun.* 1981, **29**, pp 1465–1473.
14. Bux, W.; Closs, F.; Kümmerle, K.; Keller, H.; Müller H. R. A Reliable Token Ring for Local Communications. *IEEE J. Select. Areas Commun.* 1983, **SAC-1**, pp 756–765.
15. Strole, N. C. A Local Communications Network Based on Interconnected Token-Access Rings: A Tutorial. *IBM J. Res. Develop.* 1983, **27**, pp 481–496.
16. Carlo, J. T.; Love, R. D.; Siegel, M. S.; Wilson, K. T. *Understanding Token Ring Protocols and Standards*. Artech House: Norwood, MA, 1998.
17. Müller H. R.; Keller, H.; Meyr, H. Transmission Design Criteria for a Synchronous Token Ring. *IEEE J. Select. Areas Commun.* 1983, **SAC-1**, pp 721–733.
18. Bux, W. Token-Ring Local-Area Networks and Their Performance. *Proc. IEEE* 1989, **77**, pp 238–256.
19. Christensen, K. J.; Haas, L. C.; Noel, F. E.; Strole, N. C. Local Area Networks—Evolving from Shared to Switched Access. *IBM Syst. J.* 1995, **34**, pp 347–374.
20. Ross, F. E. FDDI–A Tutorial. *IEEE Commun. Mag.* 1986, **24**, pp 10–17.
21. Dykeman, D.; Bux, W. Analysis and Tuning of the FDDI Media Access Control Protocol. *IEEE J. Selected Areas Commun.* 1998, **6**, pp 997–1010.

NORMAN C. STROLE, Ph.D
WERNER BUX, Ph.D
ROBERT D. LOVE
IBM Corporation, Research Triangle Park, NC
IBM Corporation, Zurich, Switzerland
LAN Connect Consultants, Raleigh, NC