LOCAL AREA NETWORKS

Local area networks (LAN) are data communications networks that are restricted in extent to an office, home, building, or, in some cases, areas as large as a campus. Due to the spectacular growth in networking, LANs can be found deployed in almost every organization. There are a variety of established and evolving technologies that are used in LANs, based on physical facilities ranging from copper and optical fiber to radio. The characteristics of commonly used LAN technologies will be discussed in this article.

LAN TOPOLOGIES

LANs can be logically organized in several topologies, the most popular of which are the bus, star, and ring.

In the bus structure, illustrated in Fig. 1, nodes (computers, printers, or similar devices) are interconnected to a common, shared physical resource, typically a wire or cable. This topology is inexpensive, since wiring expenses are shared among the nodes. Unfortunately, this scheme involves sharing the limited bandwidth resources of the bus and can also be somewhat unreliable, as bus failures in the vicinity of one node can affect the others on the same bus. IEEE 802.3 10base5 and 10base5 Ethernet are examples of networking standards based on a bus topology at the physical layer. This remains the most common topology in use, however, due to the simplicity of deployment.

An alternative is the star topology, shown in Fig. 2, in which each node has dedicated resources to some central switching site. This has the advantage of dedicated bandwidth to the interconnection point, but the attendant cabling costs are often higher than the bus topologies. Asynchronous transfer mode (ATM) is an example of a networking standard based on a star topology. There is increasing interest in star topologies (switched Ethernet, for another example) because the limited bandwidth on a cable is not shared and traffic is

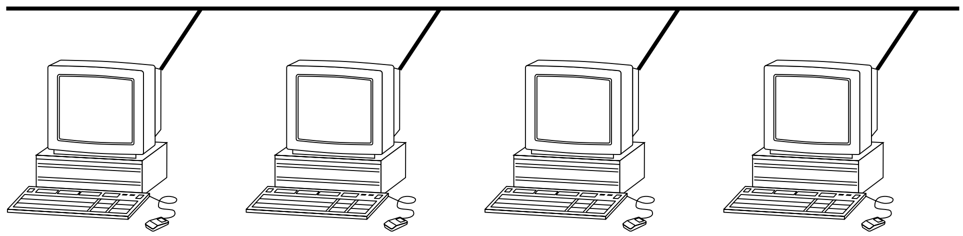

Figure 1. Bus topology. The transmission medium is shared among stations in this configuration.

not subject to internode arbitration delays for access to the medium.

Another option is the ring topology, shown in Fig. 3, in which each node is interconnected to its neighbor. The IEEE 802.5 Token Ring and the Fiber Distributed Data Interface (FDDI) are examples of networking standards based on a ring topology. This topology shares many of the advantages and disadvantages of the bus topology—inexpensive wiring, but with reliability problems if the ring should be broken. Ring-based LANs have been designed to overcome the reliability issues by using counter-rotating rings (FDDI, for example). Depending on the protocols in use, bandwidth in a ring-based network can be reused (since the ring is not physically contiguous) and hence such a topology can have a capacity greater than the equivalent bus.

IEEE 802 LAN STANDARDS

Much of the growth in deployment of LAN technology can be attributed to the standardization of selected technology options, which has enabled multivendor interoperability and has spawned a highly competitive market. The IEEE 802 LAN standards are among the most widely used data protocols yet developed.

The IEEE 802.2 standard specifies the Logical Link Control (LLC) protocols used by the other IEEE LAN standards. IEEE 802.2 allows the lower-level protocols to interface with higher-level protocols in a consistent manner. Using this approach, for example, the Internet Protocol (IP) need not know the type of underlying hardware being used on a particular host, which implies that software can be simplified and made more reliable. Note that certain other protocol suites (IP over ATM, for example) use the IEEE LLC SAP (service access point) codes for protocol multiplexing and demultiplexing, so that similar benefits can be obtained. IEEE 802.2 provides several services; which services are used and the extent to which they are used depends on the needs of the other protocols involved.

The IEEE 802.3 standard has been one of the most successful in the IEEE LAN suite. This standard describes the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol, which forms the basis for the Ethernet family (note that the 802.3 standard and Ethernet differ slightly but can be made to interoperate). The IEEE 802.3 standard is comprised of several related protocols, for different physical media. Included are the original 10base5 standard for CSMA/CD on 50 Ω thick coaxial cable, the 10base2 standard for lighter 50 Ω coaxial cable, and the 10baseT standard for unshielded twisted pair cables. Less commonly used today are the 1base5 StarLAN standard and the 10broad36 standard for more widely dispersed networks. In addition, fiber extension options are available for distributed site interconnection (within protocol distance limits). Ethernets can be found in almost all corporate data networks. The primary data rate is 10 megabits/s, although higher rate Ethernet protocols are becoming available, particularly 100 Mbits/s Fast Ethernet and ongoing work on Gigabit Ethernet.

The IEEE 802.4 standard specifies the Token Bus protocol. This protocol has been the basis for several networking technologies, including the MAP/TOP (Manufacturing Automa-

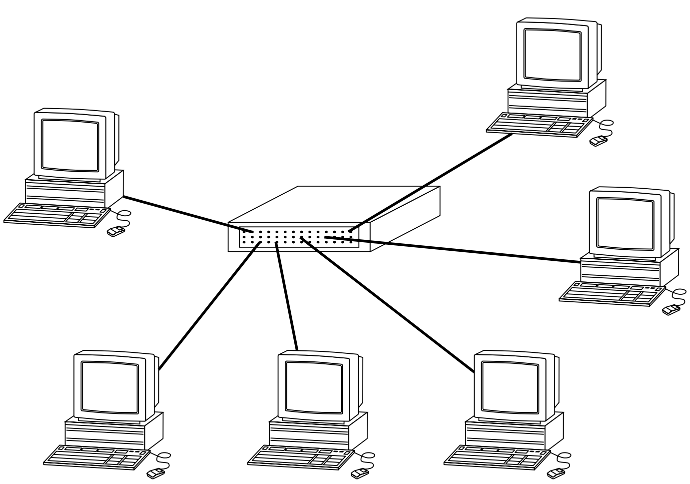

Figure 2. Star topology. This scheme is based on a central interconnection point for the transmission medium.

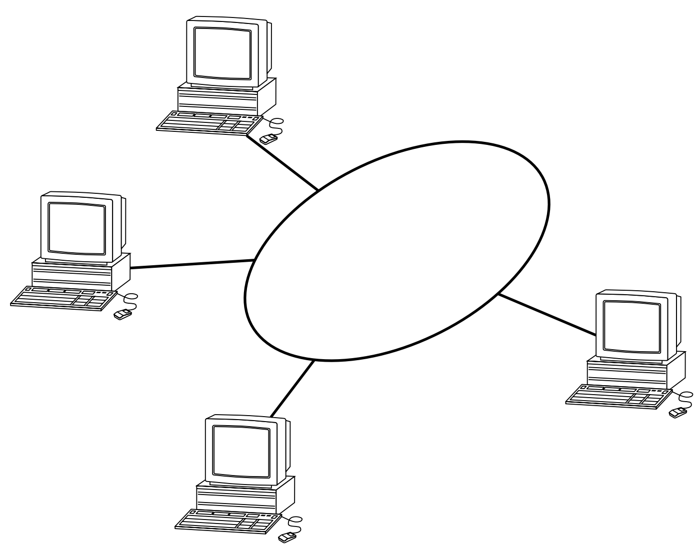

Figure 3. Ring topology. The ring is based on a loop configuration for the medium.

tion Protocol/Technical and Office Protocol) suite. Multiple physical layers are defined for token bus on 75 Ω coaxial cable, including systems at 1 Mbit/s, 5 Mbits/s, and 10 Mbits/s. These are all broadband systems. The original 1 Mbit/s system has been quite popular due to its low cost and relative simplicity.

The IEEE 802.5 standard specifies the Token Ring protocol. This standard has been widely deployed in PC-based networks and is second only to Ethernet in ubiquity. It uses unshielded twisted pair cabling, with data rates at 4 Mbits/s and 16 Mbits/s. It has several very desirable features, including robust behavior in the presence of high traffic loads and bounded delay (to transmit) times.

PROTOCOL LAYERING

For standardization purposes, networking protocols are most often conceptually partitioned into several layers. In the case of LAN technologies, the physical layer (PHY), media access layer (MAC), and logical link layer (LLC) are commonly specified. The latter two are often grouped together to form the data link layer in standard layering schemes.

PHY Layer

The PHY, or physical layer, is the lowest layer of a protocol stack. The standards for this layer typically describe the medium to be used (e.g., cable, fiber, wireless), modulation schemes, and encoding schemes used to transmit information across the medium.

The PHY layer of LAN protocols generally fall into two categories, baseband and broadband. A baseband PHY layer is one in which the information bearing signals are digital signals, typically encoded using simple level-based keying, Manchester encoding, or differential Manchester encoding. This is the most common type of PHY layer in current LANs, being relatively inexpensive and sufficiently robust for most local environments. The disadvantages are distance limitations, typically 100 m to at most 1000 m on copper, and bandwidth, no more than about 155 Mbits/s over copper using current technologies. Baseband techniques may be used over optical fiber at much greater distances and rates, but with the attendant installation and network equipment costs. For typical LAN installations, however, baseband systems on copper are sufficient.

Encoding schemes are another key element of the PHY layer. A variety of schemes, tailored to the physical medium for a given protocol, have been developed. Some typical encoding schemes are depicted in Fig. 4. These can be broadly classed as non-return-to-zero (NRZ) techniques and biphase techniques.

The conceptually simplest schemes are the NRZ methods. In the NRZ-level approach, for example, zeros are encoded as low voltage level, and ones are encoded as a high voltage level. In optical fiber systems, the corresponding schemes may be that ones are the presence of optical power, zeros the lack of light. In the NRZI (NRZ with invert on ones) approach, a transition (either falling or rising edge) denotes a one, and the lack of a transition signifies a zero.

While simple, the NRZ schemes have several shortcomings. Most significantly, recovery of bit timing at the receiver can be difficult—the moment in time at which to sample a bit to determine if it is a zero or one is often not apparent in the presence of noise and other such impairments. A technique that provides an unambiguous timing reference is highly desirable. Furthermore, the occurrence of a long string of zeros or ones can result in an undesirable dc voltage bias on the transmission medium, which may cause threshold-related errors and problems with the use of transformers. There are several approaches to resolving these related problems, which center around the need for signal transitions.

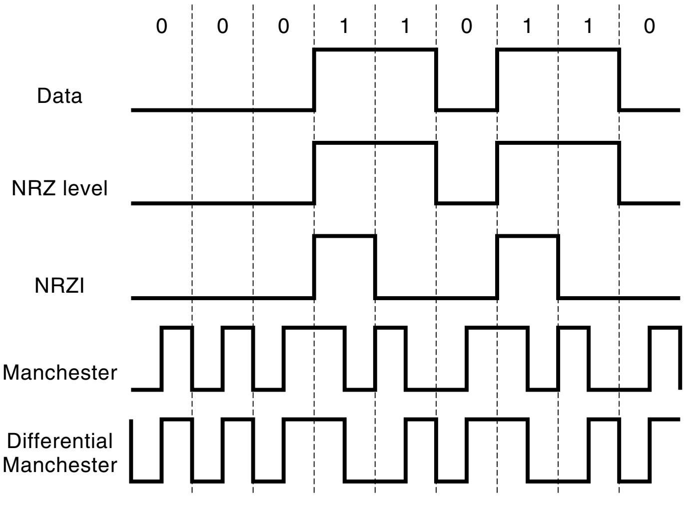

Figure 4. Baseband PHY encoding schemes. This illustrates the relationship between data bits and the signal (optical or electrical) sent across the physical medium over time.

The 4B/5B and related techniques (4B/6B and 8B/10B are also common) involve guaranteeing sufficient transitions by inserting extra bits into the signal stream. Data symbols, 4 bits in this case, are mapped into a 5 bit code, which is then transmitted using NRZI, for example. This is illustrated in Table 1. An inspection of this table will prove that strings with a maximum of three consecutive zeros are possible, even when code words are concatenated. Multiple ones are not an issue if NRZI is used for transmission, as ones force a transition to occur. The cost of a 4B/5B mapping, of course, is that only 80% efficiency is possible.

Table 1. 4B/5B Encoding

Data Symbol	Code Word
0000	11110
0001	01001
0010	10100
0011	10101
0100	01010
0101	01011
0110	01110
0111	01111
1000	10010
1001	10011
1010	10110
1011	10111
1100	11010
1101	11011
1110	11100
1111	11101

The biphase encoding is beneficial with respect to signal balance maintenance, and bit timing recovery is particularly easy to implement in this scheme. It is based on signal transitions at a rate double that of the bit rate. A transition (rising edge or falling edge) is guaranteed to occur at the center of the bit period. The absence of such a transition can be used as an error detection method.

In Manchester encoding, a zero is encoded as a rising edge at the center of the bit period, and a one as a falling edge at such a time. The encoding mechanism can be implemented as an exclusive-or operation between the data and the clock. This is the encoding used for most of the common IEEE 802.3 protocols (10base5, 10base2, 10baseT). Differential Manchester encoding uses the midperiod transition for a clocking reference only, and uses the presence (denoting a zero) or absence (denoting a one) of a transition at the beginning of the bit period to encode the information. This is the method used for the IEEE 802.5 token ring standard.

The primary disadvantage of biphase signaling is that transitions happen at twice the data rate, which means that the bandwidth required is greater than that of the equivalent NRZ system, and the hardware must operate twice as fast. The former is particularly critical in wireless systems.

A broadband PHY layer is one in which the information is coupled into the medium by analog signals, which are modulated by some carrier, and encoded using frequency shift keying (FSK), amplitude shift keying (ASK), phase shift keying (PSK), or some similar scheme. This type of PHY layer is most often used in situations where longer distances need be served or additional bandwidth is required. Much greater bandwidths may be supported on one cable using broadband schemes, as multiple frequencies can be used. The primary disadvantage of the broadband approach is the cost of modulators, demodulators, and the associated analog hardware.

MAC Layer

The MAC, or media access layer, is used to arbitrate access to the PHY layer. For example, in the case of Ethernet, there is a shared medium (cable) that must be used by several nodes, and only one of the nodes can be permitted to access the cable at a particular time.

The MAC layer influences the effective throughput over a given physical layer and should be efficient in its use of the available bandwidth. This includes minimizing the overhead due to factors such as protocol headers and dead time between transmissions, while at the same time maximizing the successful transmissions on a busy shared-medium network. In addition, the MAC layer is often designed to ensure that errors are not propagated to the higher-layer protocols.

Various MAC schemes have been developed for the LAN protocols. The three most common are the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol used in Ethernet, the token ring protocols, and the token bus protocol.

The CSMA/CD MAC protocol involves detecting the use of the medium by another station by checking the state of the carrier. If a station has data to transmit, it first attempts to verify if the medium is unused. If it is available, the station transmits. If the medium is not available, the station waits until the medium goes idle and then immediately begins to transmit (note that this is the IEEE 802.3 solution, but other options are possible in the general case of CSMA). The success or failure of transmissions is monitored on the shared medium, and if a transmission is unsuccessful—that is, a collision is detected—the station waits a prescribed random amount of time (binary exponential back-off) and attempts to transmit again. This procedure is repeated until the transmission is successful, or the limit to the number of transmission attempts (16 in IEEE 802.3) is reached. CSMA/CD is simple, inexpensive, and performs well under light loads. Unfortunately, it can perform poorly under heavy loads and be sensitive to physical layer errors.

Token ring protocols use a "token" to arbitrate access to the transmission medium. A token is a small frame that is exchanged between stations to gain the right to transmit. If a station has data to transmit, it waits until a token is seen on the medium. This station then modifies the token and appends the necessary fields as well as its data. When this frame returns around the ring to the originating station, it is purged from the medium. When data transmission is complete, the station inserts a new token onto the ring. Token rings support fair, controlled access to the medium and perform well under heavy load conditions. A disadvantage is the need for careful token maintenance, particularly in the presence of errors. Several varieties of token ring exist; some of these will be discussed in subsequent sections.

The token bus protocol is closely related to the token ring, but with an underlying physical bus topology. The token exchange mechanism, however, does in fact use a logical ring for token passing. This logical ring is simply an ordering of stations on the bus. Once the logical ring is in place, token passing can proceed as in a ring-based system. This system provides controlled access to the bus and is robust under heavy loads. One of the disadvantages of this approach is that ring initialization and maintenance is more complex than in a physical ring—the ordering of stations must be determined through some algorithm, and station additions and deletions must be managed.

LLC Layer

The LLC, or logical link control layer, can be viewed as the upper part of the data link layer. It is used to provide data services to the higher layers. In particular, two types of services, connectionless and connection oriented, are defined. In LANs supporting complex higher-layer protocols, such as TCP/IP, only the simplest LLC services are commonly used.

An example of a LLC protocol is the IEEE 802.2 layer. This provides both connectionless and connection-oriented services. The unacknowledged connectionless service provides simple datagram support for the multiplexing and demultiplexing of higher-layer protocols. In addition, a connectionless service with acknowledgments (for monitoring systems, for

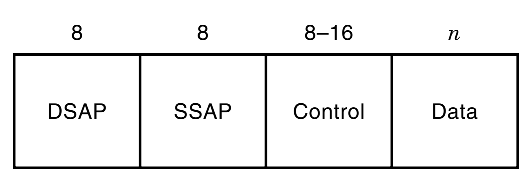

Figure 5. IEEE 802.2 LLC frame format. The SAP (service access point) fields are used to select the appropriate protocol handler on reception of a packet.

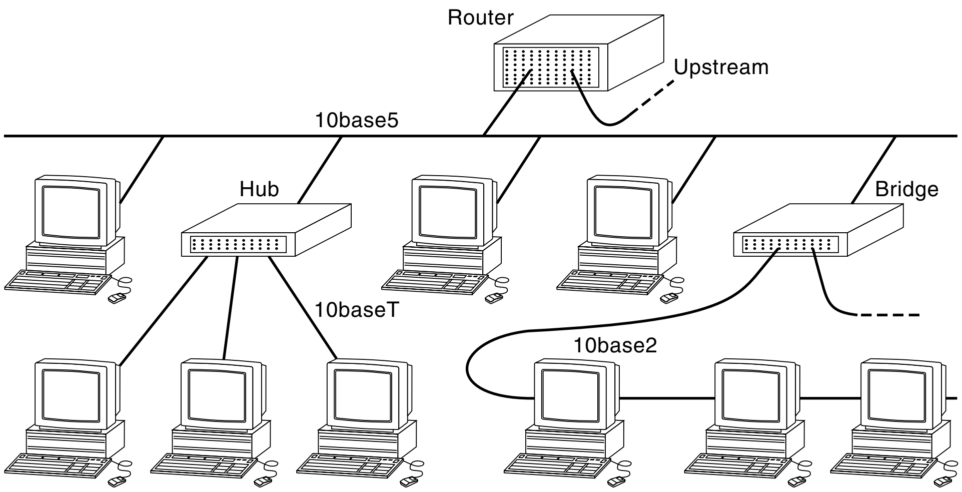

Figure 6. Typical Ethernet installation. This illustrates the interconnection of various physical and protocol devices in a typical LAN. Hubs are devices used to concentrate the physical media from several Ethernet stations, and are often used as physical layer translation devices (10baseT to 10base2, for example). Bridges provide isolation between Ethernets and allow more complex LANs to be built.

example) is supported, as well as a connection-oriented service that furnishes flow control and error recovery capabilities based on the lower-layer CRC and a "go-back-N" strategy. The IEEE 802.2 LLC frame format is depicted in Fig. 5. The destination service access point (DSAP) and source service access point (SSAP) fields are used to indicate the service type (IP or IPX, for example) to higher layers. The control field is used for the LLC service support, including indication of the type of service.

ETHERNET (IEEE 802.3)

The Ethernet, and the closely related IEEE 802.3 standard, has been one of the most successful LAN protocols developed to date. This technology is based on CSMA/CD and takes a variety of forms at the PHY layer. A typical Ethernet installation is depicted in Fig. 6.

The Ethernet frame format is illustrated in Fig. 7. The preamble is used for frame delineation. The destination and source address fields (48 bits each) are globally unique identifiers for each Ethernet adapter and are used for station to station communication, as well as broadcast (all ones) and multicast (first bit is one). It should be noted that the 48 bit addresses used in Ethernet have become a common feature in IEEE 802-based LANs. The type field (Ethernet) can be used for higher-layer demultiplexing, as in an LLC protocol. The length field (IEEE 802.3) is used to aid in end of frame detection. A 32-bit CRC is used for error detection and is followed by a postamble for end of frame detection.

TOKEN PASSING BUS (IEEE 802.4)

The IEEE 802.4 Token Bus standard has been widely used in manufacturing systems and early office automation products. Because it is based on a broadband physical medium, it is somewhat more resistant to the low-frequency electromagnetic (EM) noise that might arise on a factory floor.

A token passing bus is a LAN with a bus topology that operates on the principle that a token will be received prior to the transmission of data by a station. The token bus format includes a preamble, frame control byte for denoting whether a particular frame is a token or data, the destination and source addresses (48 bits, as in 802.3), the data (an LLC frame), an error detection field (CRC-32, as in 802.3), and the postamble.

The token bus operates by first establishing a logical ring that overlays the physical bus topology. Station additions and deletions require reconfiguration of the logical ring. When a token is received, a station is permitted to transmit multiple packets, until its token holding time has expired. The token bus offers optional support for multiple classes of service through the use of complex timer specifications that enable per-class bandwidth guarantees. Support for simpler nontoken stations is included to allow low-cost devices to respond to polling requests using this medium.

MAP/TOP

MAP is the Manufacturing Automation Protocol developed by General Motors Corporation for communication among automated manufacturing devices, including robotic equipment and the associated controllers. It was primarily designed to support communication between very different sorts of devices, in real time with low, predictable delays. It supports applications as varied as word processing and equipment telemetry (temperature measurement, for example).

TOP is the Technical and Office Protocol developed by Boeing Corporation for communication between office automation

64	48	48	16		32	8
Preamble	Destination address	Source address	Type or length	Data	CRC	Postamble

Figure 7. Ethernet frame format. Ethernet and IEEE 802.3 differ in the fourth field, which is a type field in Ethernet and a length field in IEEE 802.3.

devices such as word processing systems and printers. Interoperability between devices from a variety of manufacturers was a key design goal of this protocol.

The MAP/TOP protocol suite is based on the IEEE 802.4 Token Bus protocols. As such, MAP/TOP networks are often interconnected with some variety of token passing network for ease of interface design.

TOKEN RING (IEEE 802.5)

The Token Ring protocol has been widely deployed in networks based on PCs. Token Ring operates on the principle of the exchange of a "token" to a station before it is permitted to transmit. Only one token is allowed on the ring at one time.

The IEEE 802.5 Token Ring frame formats are illustrated in Fig. 8. The first format is used for token frames and only includes start and end delimiters and the access control field, with priorities and reservation information. The second format includes start and end delimiters, a frame control word for optional LLC support, source and destination addresses (in 802.3 format), the LLC (data) frame, a CRC-32, and a frame status word used by transmitting stations to verify reception.

OTHER TOKEN RINGS

Another example of token ring technology in wide use today is the Fiber Distributed Data Interface (FDDI) standard. This technology supports multiple packets on the ring at one time, with rates of 100 megabits/s. Provisions are made for multiple service classes (synchronous and asynchronous) with differing throughput and delay requirements. Further, reliability support is provided through the capability for optional redundant counterrotating rings, which can mask a station or fiber failure.

Slotted Ring

Slotted ring technology uses multiple "slots" that rotate around the ring to arbitrate access. Each slot is a small frame that can be marked empty or full. When an empty slot arrives at a station with data to transmit, the slot is marked full and data is injected. The slot is marked empty when it returns around the ring to its source. A given station cannot transmit again when it has an outstanding slot. The provision for multiple packets from different sources on the ring at one time assists in fair utilization and quality of service support.

The Cambridge Ring is an early example of such technology (some claim it is the ancestor of the ATM protocols, also based on small fixed frame sizes). A slot contains one octet each for the source and destination addresses, five control bits, and two, four, six, or eight data octets, and thus slot sizes are extremely small. This implies that higher-layer packet data is almost always segmented into small units prior to transmission. Stations could choose not to receive packets from particular sources; some of the control bits support this through response codes. The Cambridge Ring was simple to implement, but was somewhat wasteful of bandwidth due to the header overhead in such small datagrams.

Register Insertion Ring

Register insertion rings are a common LAN technology and can be used to provide high performance through their support for multiple packets on the ring at one time. The register insertion ring uses a small shift register at each station to control forwarding and insertion onto the ring. The shift register is at least as large as the maximum frame size. This allows a station to store a frame as it passes. If the station has no data to send, a passing frame is buffered long enough to determine if it is destined for the local station. If it is destined locally, a typical implementation will both copy the frame into adapter memory and forward the frame back around the ring to support acknowledgments. Transmission when the medium is available is handled by simply copying the data onto the ring. If a frame arrives during this time, it is buffered in the insertion register. The register insertion method provides excellent ring utilization due to the multiple simultaneous packets on the ring without the overhead penalty of the slotted ring. The disadvantage of this technology is that the purge mechanism—that is, the technique used to remove problematic packets from the ring—is generally more complex than in other systems.

HYPERchannel AND HIPPI

A number of LAN protocols are designed for very high-speed interconnection of computers and their peripherals. HYPERchannel, developed by Network Systems Corporation, is one of these. This protocol was developed in the mid-1980s for the interconnection of supercomputers and high-performance peripherals, and has been used with Cray and Amdahl systems, among others. It supports data rates of up to 275 megabits/s over a variety of physical layers.

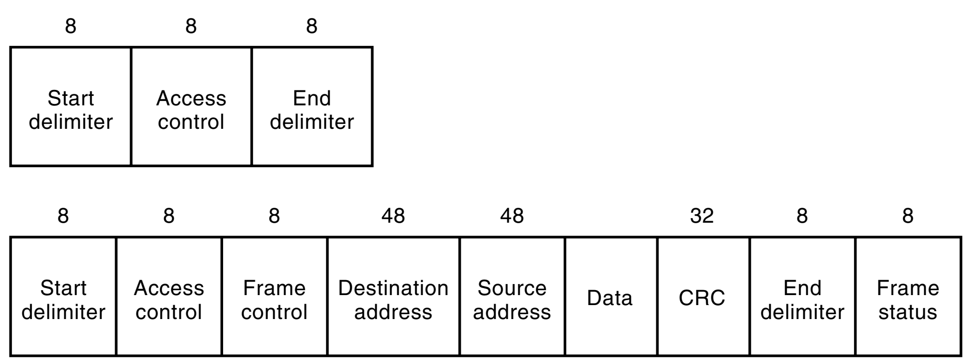

Figure 8. Token ring frame format. The different formats used for the control token and data frames are depicted.

HIPPI, or High Performance Parallel Interface, is another of the protocols developed primarily for interconnection of supercomputers. This protocol supports 800 Mbits/s or 1.6 gigabits/s links over a large parallel cable, which is either 32 lines or 64 lines wide and runs at 25 MHz. The distances over which HIPPI can be used are quite limited, but large enough for a typical supercomputer center equipment floor. Interconnection of sites can be accomplished using fiber extension options. Simple flow control features are provided to lessen problems with computers and peripherals of widely different I/O (input/output) bandwidth. Although simple and effective, this flow control scheme does contribute to the problems of extending HIPPI networks over larger distances while maintaining high throughput. To build HIPPI networks of nontrivial size, simple switches are used to interconnect devices. These switches are typically not designed to switch between sources and destinations at high rates, as with routers and packet switches, but rather act as interconnection panels that may be reconfigured at reasonable rates for sharing peripherals.

OTHER LAN PROTOCOLS

A number of new, higher-performance LAN protocols have been developed in recent years. FiberChannel is a LAN protocol suite designed for high-speed communication between nodes using optical fiber. Rates of up to 800 Mbits/s are supported, with systems up to 4 Gbits/s under design. Other developments include FireWire (IEEE 1394) and universal serial bus (USB), high-speed protocol suites based on serial interconnection technology. FireWire, for example, supports bandwidths of up to 400 Mbits/s with up to 63 devices (with no more than 16 cable hops) per bus. USB, a 12 Mbits/s serial protocol with chaining support, is designed primarily as an improvement over traditional serial port technologies.

Asynchronous Transfer Mode (ATM) networks are also being widely deployed in LANs. ATM is a switch-based technology that uses small packets (53 bytes) called cells. Interconnection of nodes is through virtual circuits, which are analogous to circuits in voice telephony. Multiple physical layers are supported, including both copper and fiber infrastructure options. Although ATM is often viewed as a wide area networking technology, it does provide support for features that are not available in other technologies. For example, ATM allows the definition of virtual LANs, which provide network administrators with options that are not available in less sophisticated technologies. Virtual or emulated LANs are interconnections of LANs, perhaps widely separated, which are configured to emulate a single local area network. Furthermore, multiple logical local area networks can be supported over a single physical infrastructure using this capability.

WIRELESS LANS

Wireless LANs use radio or infrared as the transmission medium, as opposed to the traditional wire or fiber. This has significant advantages, particularly for deployment in older buildings where wiring costs are high, as well as environments in which workers may be moving frequently.

Many of the initial wireless LANs in the United States have used radio frequencies in one of the ISM (Instrumentation, Scientific, Medical) bands, which generally may be used without individual site licensing subject to restrictions on power output. The data rates on these systems range from tens of kilobits per second to a few megabits per second, with typical ranges of a few hundreds of meters. Some early European products in this area were based on the digital enhanced cordless telephony (DECT) standard for digital telephony. These systems used multiple channels to provide data rates of hundreds of kilobits per second.

Wireless LANs are still evolving, but instances of several standards are now being seen as products. The most significant development in this area is the IEEE 802.11 standard, which will provide data rates of 1 Mbit/s to 2 Mbits/s over a range of approximately 100 m in typical radio configurations. This is based on a CSMA/CA (CSMA with Collision Avoidance) MAC layer with multiple physical layers—in particular, direct sequence spread spectrum, frequency hopping spread spectrum, and infrared.

Work on indoor wireless ATM and the European ETSI RES10 standards are focusing on systems with data rates of up to 25 megabits/s with ranges on the order of 100–200 m. While they are still early in the development cycle, these systems promise to deliver multimedia services over wireless links with quality of service support.

BIBLIOGRAPHY

P. T. Davis and C. R. McGuffin, *Wireless Local Area Networks,* New York: McGraw-Hill, 1995.

L. L. Peterson and B. S. Davie, *Computer Networks—A Systems Approach,* San Francisco: Morgan Kaufmann, 1996.

S. Saunders, *The McGraw-Hill High-Speed LANs Handbook,* New York: McGraw-Hill, 1996.

W. Stallings, *Local Networks,* 5th ed., New York: Macmillan, 1996.

J. Walrand and P. Varaiya, *High Performance Communication Networks,* San Francisco: Morgan Kaufmann, 1996.

J. Wobus, *LAN Technology Scorecard* [Online], 1996. Available: http://web.syr.edu/~jmwobus/comfaqs/lantechnology.html

JOSEPH B. EVANS
University of Kansas

LOCAL AREA NETWORKS. See ETHERNET.

LOCATION SYSTEMS, VEHICLE. See VEHICLE NAVIGATION AND INFORMATION SYSTEMS.