Local area networks—Evolving from shared to switched access

by K. J. Christensen
L. C. Haas
F. E. Noel
N. C. Strole

Local area networks (LANs) have become pervasive in their application to business, government, and academic data communications. In the United States alone, by 1993 there were over three million LANs installed. This paper describes how advances in structured cabling systems and electronics concentrated in the wiring closet have enabled the evolution of LANs toward a dedicated bandwidth model. Due to the growing number of users, LANs are required to provide increased data capacity and improved management capabilities. These user requirements are being met by increased functions in the electronic wiring closet that include fault isolation, remote network management, increased aggregate bandwidth, and virtual LAN capabilities. This paper focuses on present and future user requirements and the influences of emerging LAN switching technologies on already installed LANs.

Shared-access local area networks (LANs) have matured over the last decade to where they now provide the predominant interconnection fabric for local enterprise networks around the world. In the United States alone, there were 3.1 million LANs in 1993, up from 0.5 million in 1991.[1] This proliferation of LAN installations has led to a network-centric paradigm for the 1980s and 1990s, versus the host-centric paradigm of the 1960s and 1970s. LAN access speeds have increased from early 1 megabit per second (Mbps) data rates to in excess of 100 Mbps today. As the number and importance of LAN-attached workstations expands, customers are demanding increased bandwidth[2] and improved management. Existing and emerging applications, including client/server computing and full-motion video-on-demand, are now pushing the bandwidth limits of shared-access LAN technology. At the same time, the personnel costs to configure and manage LANs have grown to become a major portion of the information systems budget for many large companies.[3]

A recent trend in network design is to reduce the number of stations that share the bandwidth of a single physical LAN segment (i.e., that are "shared media" by virtue of all stations sharing the common bandwidth of a single transmission medium), even to the point where only a single station occupies a segment. The interconnect hub, normally installed in the LAN wiring closet, thus becomes the "traffic cop" that steers frames (data structures that consist of protocol, control, and user data) to the appropriate destinations. This rather simple concept, known as *LAN switching*, can result in significant increases in available bandwidth and improved management of the individual workstation. The LAN wiring closet and LAN switching are described later in this paper (and are also shown in Figure 1 and later in Figure 12). LAN switching is

0018-8670/95/$3.00 © 1995 IBM

Figure 1 First-generation direct-wired LAN and second-generation star-wired LAN

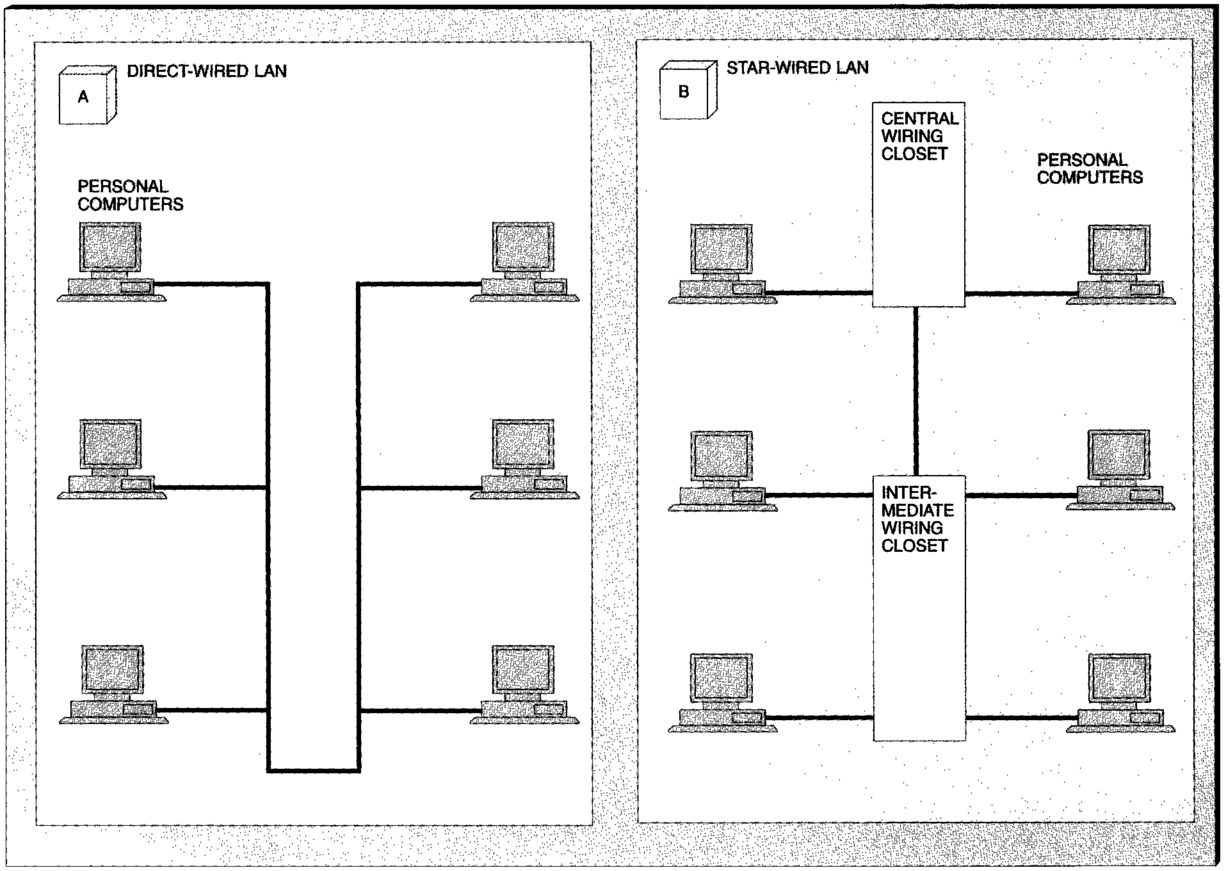

expected to fuel both the expansion and evolution of LAN technology through the next decade.

This paper describes the evolution of LANs starting from early shared-media departmental LANs, through today's interconnected LANs, and into the emerging future of switched LANs. Four generations of LANs can be identified. The first generation was typified by small shared-media LANs such as early coaxial bus-topology Ethernet.[4] First-generation LANs did not readily grow in size due to a lack of planned cabling methods, thus driving the development of systematic "star-wired" cabling systems and the second generation of LANs. Figure 1 shows both the first-generation direct-wired cabling (Figure 1A) and the second-generation star-wired cabling (Figure 1B). In star-wired cabling all workstations are attached to a common point—the wiring closet.

The second generation of LANs was typified by star-wired 4- and 16-Mbps token ring and 10-Mbps Ethernet on unshielded twisted pair cabling.[5] The second generation of LANs also saw the emergence of internetworking—the connection of local and remote LANs with bridges and routers. Increasing electronics in the wiring closet and increasing segmentation of LANs typifies the present third generation of LANs. The emerging fourth generation is based on switching technology. Switching technology increases network bandwidth and enables virtual LANs. This increase in bandwidth can enable continued growth of existing networks and enable new high-bandwidth applications. In a virtual LAN, software definition and not geographic locality determines the LAN segment to which a workstation is attached. A LAN segment is the subnetwork to which a user is attached. For this paper, a LAN segment contains the users attached to a sin-

Figure 2 Four generations of LANs

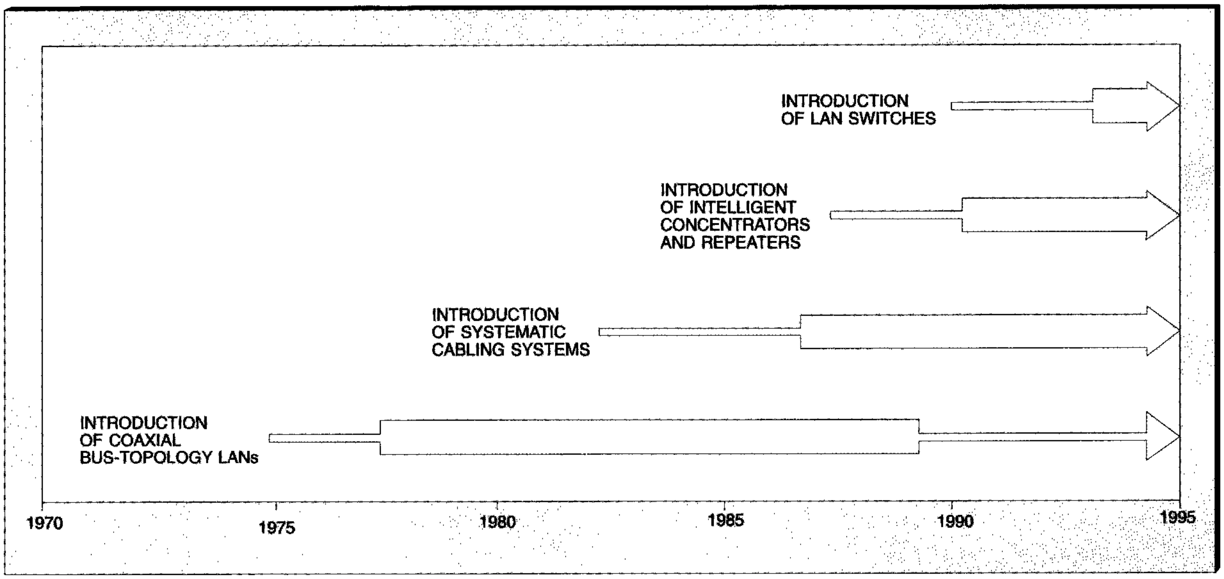

gle address domain, for example a single token ring or Ethernet. LAN segments are interconnected, for example by LAN bridges, to create larger networks. Virtual LANs simplify user relocation and other aspects of network management. Switching technology is also the means of accessing both higher-speed LAN technologies and future asynchronous transfer mode (ATM) networks. Figure 2 shows the four overlapping generations of LANs. The thinning arrow for coaxial-based LANs indicates that these LAN types are decreasing in number. The remainder of this paper traces the evolution of LANs with an emphasis on the emerging fourth generation.

First- and second-generation LANs

The first LANs addressed the need to communicate among the various distributed computers appearing on the desktop and to communicate with a centralized mini-computer installation. These first LANs evolved from first-generation ad hoc direct-wired installations to second-generation systematically planned star-wired installations.

First-generation LANs. The centralized mini-computer installations of the late 1970s were primarily department-level processors, supporting less than 100 users, who were, for the most part, located in close proximity to one another. The key desires were for simple installation, low cost, and low maintenance. The systems that were developed to support these distributed computing installations can be illustrated by briefly describing several different choices that became available between 1976 and 1984. The exact timing of these offerings is less important than overall system characterizations. An early predecessor technology of first-generation LANs was IBM's R-Loop[6] communications system. This system transmitted at 19.2 kilobits per second (Kbps), was based on telephone twisted pair cabling, and was first introduced as a direct-wired system (see Figure 3). This system had one master, the IBM 8100, and all attaching terminals were slave units communicating with the IBM 8100 in sequence based on their position in the loop. R-Loop used a Synchronous Data Link Control (SDLC) frame format. It is important to note that within about six months of its introduction, the R-Loop was enhanced by the introduction of a wiring concentrator unit that allowed for a star-wiring configuration instead of direct wiring. Star wiring is described in the next section of this paper.

A second technology in the first generation of LANs was Ethernet—the first true LAN.[4,7] Ethernet is a 10-Mbps bus-topology architecture based on a station "listening" to the medium before transmitting.

Figure 3 Direct-wired IBM R-Loop

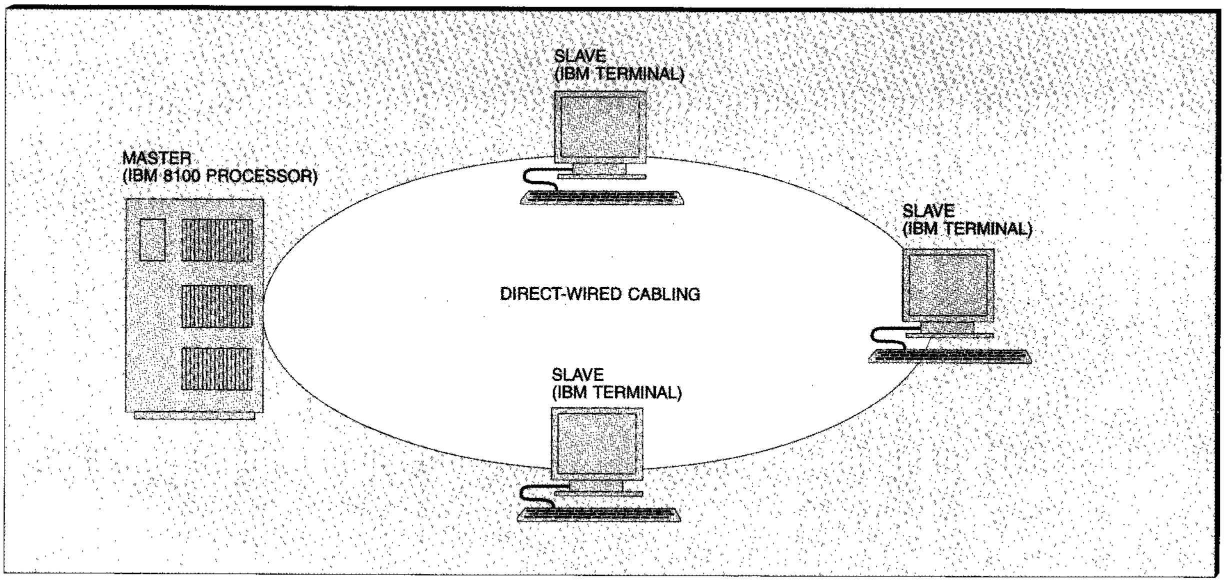

If the medium is silent, then no other station is transmitting and a station with a queued frame can transmit. If two or more stations start transmitting simultaneously, a "collision" occurs and all stations stop transmitting. This access protocol is called carrier sense multiple access with collision detection or CSMA/CD. References 8 and 9 describe the CSMA/CD protocol and its standard forms as Ethernet and the IEEE Standard 802.3. Ethernet in its classic coaxial bus-topology configuration supports up to five coaxial cable segments of 500 meters each, which is sufficient to allow for cabling one to two floors of most commercial buildings. Each terminal attaches via an electronic unit, a transceiver that is tapped into the coaxial cable and interfaces to the terminal, as shown in Figure 4. The 2500 meter cable limit was based on the data rate of the system (10 Mbps) and the minimum frame size (64 bytes). If the cable were longer, then the Ethernet collision-based protocol would not allow stations to accurately detect collisions. A cable that is too long could result in two simultaneously transmitting stations not detecting each other before assuming "no collision" and continuing with their transmissions. One of the issues with Ethernet was that troubleshooting required manually following the 2500 meters of coaxial cable through cabling ducts, ceiling crawl spaces, and other places. This method of troubleshooting is called "walking the wire" and is not acceptable today due to the high labor cost and time needed to repair failures.

A third technology was the IBM PC Network, a 2-Mbps tree-topology collision-based protocol.[10] This technology utilized either baseband coaxial cable or a broadband channel. The PC Network architecture had the advantage of being a tree topology with a physical *head-end* that allowed for relatively large configurations, but with most maintenance being performed at a central location. For broadband operation, the head-end translates between the forward (all stations transmit "up" toward the head-end) and return (the head-end translates the "up" frequency to a "down" frequency for transmission "down" the network) channel frequencies and the splitter distributes the signal to all stations. Figure 5 shows a sample of an IBM PC Network configuration.

Second-generation LANs. Systematic enterprise-level cabling systems were introduced by both IBM and AT&T in the mid 1980s. The IBM Cabling System used a shielded twisted pair (STP) cable, a high-performance cable with specifications intended to support all of IBM's existing terminal families, as well as allowing for future higher speed LANs. The AT&T Premises Distribution System offering used

Figure 4 Coaxial Ethernet attachment

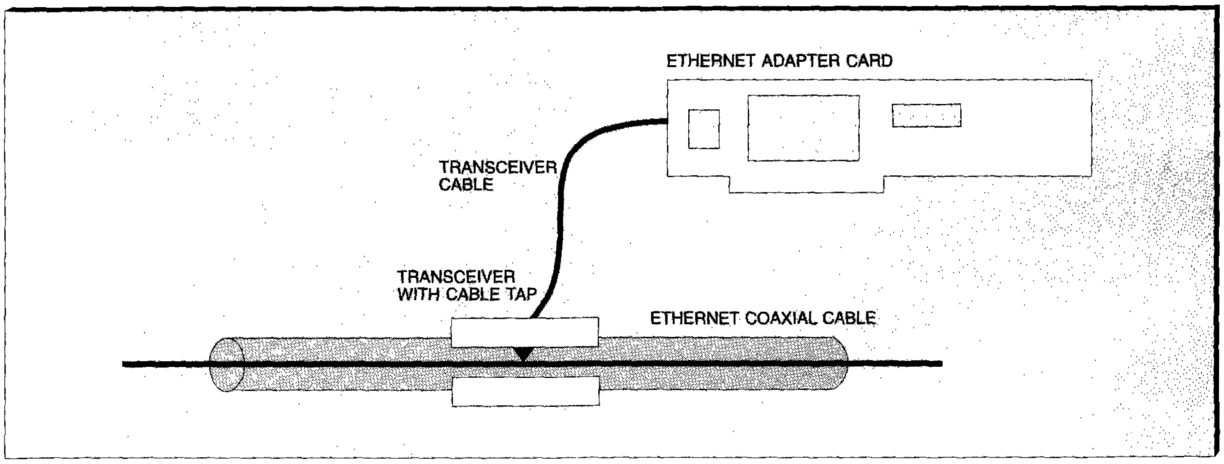

an unshielded twisted pair (UTP) cable that provided an appropriate level of performance with published specifications. The IBM Cabling System was designed to support the IBM Token-Ring LAN that was under development at that time[11] and also to provide support for the existing local telecommunications systems. These two cabling systems offered a known cabling environment that customers, not just the common carrier companies, could both install and maintain. Equally important, the cabling systems allowed manufacturers to design and test their LAN equipment over known and controlled cable links without guessing or making risky assumptions about the quality of cabling that could possibly have been installed from the 1920s to the 1970s.

Prior to the IBM and AT&T cabling systems, an equipment manufacturer had no feasible means to establish the performance parameters of the installed telecommunications cabling. Each manufacturer had to choose (or guess), in the absence of reliable data, a set of specifications for equipment design that included cable type, length, attenuation, and crosstalk. This set of choices then dictated the outcome of the design. If the choices were overly optimistic, the manufacturer's equipment would perform poorly or not at all. This poor performance often led to financial losses to the manufacturer because of field maintenance calls, return of installed products, and the general wrath of unsatisfied customers. Alternatively, if the choices were overly pessimistic, the manufacturer's equipment could be "over-designed," resulting in noncompetitive performance specifications. These risks were one of the major motivators driving the telecommunications industry to create standards.

Figure 5 IBM PC Network configuration

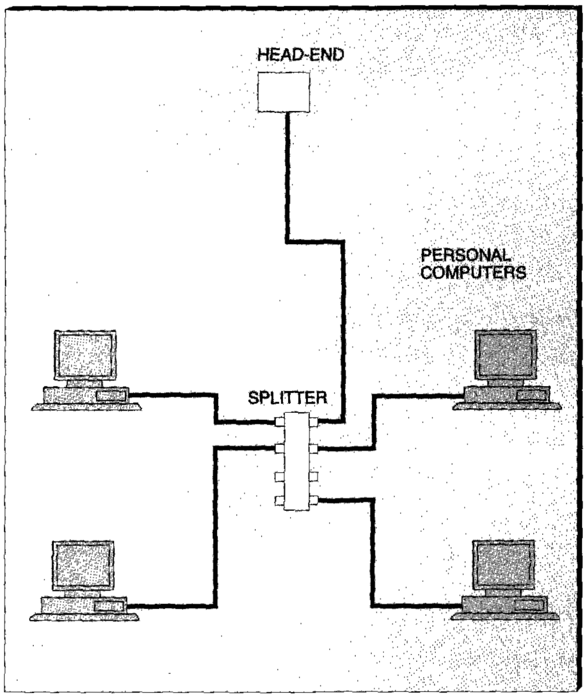

Another key motivation for a cabling standard was the gradual recognition by commercial building designers and owners that telecommunications is a critical element of virtually every business that would wish to lease or buy office space in a commercial building. The building owners and designers saw the benefit of being able to install a cabling system that would support a variety of potential occupants and avoid the need to re-cable for every new occupant. Just as there are standards or guides for lighting, plumbing, heating, and air conditioning in commercial buildings, there was a need to provide standards for telecommunications infrastructure within commercial buildings. The Electronic Industry Association (EIA) formed committees to address building cabling standards. After six years of work, two key standards were created. The two standards are *EIA/TIA-568, Commercial Building Telecommunications Wiring Standard* (TIA 568-1991), and *EIA/TIA-569, Commercial Building Standard for Telecommunications Pathways and Spaces* (TIA 569-1991).

Key ingredients of the cabling standards that spurred LAN deployment include:

- Specification of a star topology for telecommunications cabling.
- A limit to four cable types (UTP, STP, coaxial, and optical fiber) with electrical and mechanical specifications for each of the cable types.
- Maximum allowed lengths for cables used in each part of a system.
- Specifications for size and placement of telecommunication wiring closets for housing equipment that serves offices and work areas.
- Specifications for cable pathways to service offices, work areas, telecommunications equipment rooms, and entrance facilities.

The choice of a star topology for the cabling standard was easily made. Figure 1 shows a star cabling distribution for a LAN. The choice of the star topology was natural to those accustomed to the traditional practices of the telephony industry. The benefits of a star topology had also become obvious to those in the data transmission sector who had more recently gained painful experiences in the operation and maintenance of LANs that had been cabled by direct runs between stations.

Direct cable runs between stations offered shorter cables that delivered stronger signals and resulted in lower initial installation costs because less cable was needed. This was the initial installation practice for Ethernet LANs and IBM's early R-Loop system. However, early dissatisfying experiences in problem identification, isolation, and correction in direct-wired systems prepared the industry for quick acceptance of a star-wired topology. With star wiring, telecommunications wiring closets offered key locations for almost all of the testing necessary to locate and correct a system fault that might have affected hundreds of stations. Not only was the testing made easier, much more importantly, it was made immensely faster. The availability of the network was increased by being able to correct a fault and return the network to operation in less time. The phrase, "walking the wire" became part of history.

IBM's and AT&T's cabling systems were star topologies that preceded the standard. These star-wired cabling systems were easily configured to support either ring or bus configuration as required by token ring and Ethernet, respectively. The cabling standards included example configurations accommodating bus, ring, and tree connections.

The industry successfully provided products for 4-Mbps token-ring and 10-Mbps Ethernet operation over the UTP cable specified in the initial EIA standard. 10-Mbps Ethernet on UTP is called "10BaseT" for 10-Mbps baseband on twisted pair. 10BaseT is star wired, where each workstation is cabled directly to a wiring closet.

It was recognized that the initial UTP cable performance characteristics were inadequate to accommodate future, higher rate, data transmission systems. For example, operation over UTP was not included in the original 16-Mbps token-ring LAN standard. While UTP was cost-competitive with STP, it was obvious that UTP performance must be improved in order to be recognized as suitable for future LAN applications. In response, the cable manufacturers developed and marketed a series of enhanced UTP cable products. To help control a growing level of confusion over different UTP cable offerings, the EIA committee issued *Technical Systems Bulletin 36, Additional Cabling Specifications for Unshielded Twisted Pair Cables* in November 1991.

This bulletin specified five categories of UTP performance and gave typical usage. The improved cables were designated UTP category-4 and UTP category-5. UTP category-3 was the designation for the specifications in the earlier TIA 568-1991 standard. A later bulletin was issued that specified performance levels for corresponding connecting hardware.

The creation of networks much larger than a few kilometers in size requires additional interconnection equipment. Repeaters that operate on the LAN physical layer can extend single LAN segments, but only over a geographically small area and for a limited number of stations. To form needed extended networks of LANs with thousands of stations and to interconnect geographically remote LANs, bridges and routers are needed.

Where we are today—third-generation LANs

LANs today have progressed to a third generation, typified by three things:

1. Intelligence in the wiring closet for remote and automatic management
2. Increased internetworking for connectivity, building the enterprise network by interconnecting LANs with bridges and routers (although bridges and routers are similar in concept, routers usually have more function for traffic management of large networks)
3. The reduction of large LAN segments into multiple smaller LAN segments and providing high-speed backbone LANs for improved performance and reliability (typically the backbone is the primary high-speed link connecting the nodes of the network)

Intelligence in the wiring closet. Systematic cabling systems brought access to the LAN into one place—the wiring closet. For token ring, the access point is designated a concentrator and for Ethernet based on UTP, a repeater. A combination of concentrators, repeaters, and possibly other functions (such as bridging) combined into a single chassis is a *hub*. Each card in a hub is called a *module*—one module for each function. Figure 6 shows an example of a heterogeneous LAN hub with modules for two token-ring concentrators, an Ethernet repeater, and a bridge. Intelligence is added in the form of microprocessors programmed for specific management and control functions. Two such functions typical in the electronic wiring closet are network management and port switching.

By distributing network management functions to the wiring closet, remote and automatic management of LANs is possible. Remote management functions include the ability to observe key measurements of the LAN and the ability to enable and disable ports in a concentrator or repeater. Stations attach to a concentrator or repeater through a port. Key measurements include error and performance statistics for a LAN segment and for individual attached stations. Error statistics, such as the percentage of frames with errors, allow a network administrator to isolate faulty components in a LAN. Automatic management functions include the ability to prevent access to a LAN segment by unauthorized or invalid stations. Unauthorized stations could include stations whose addresses are not contained in a management database. For token ring, invalid stations would include stations configured for the wrong data rate. For example, an intelligent concentrator for token ring will prevent a 4-Mbps station from inserting into and disrupting a 16-Mbps token-ring LAN segment.

Figure 6 Example of a heterogeneous LAN hub

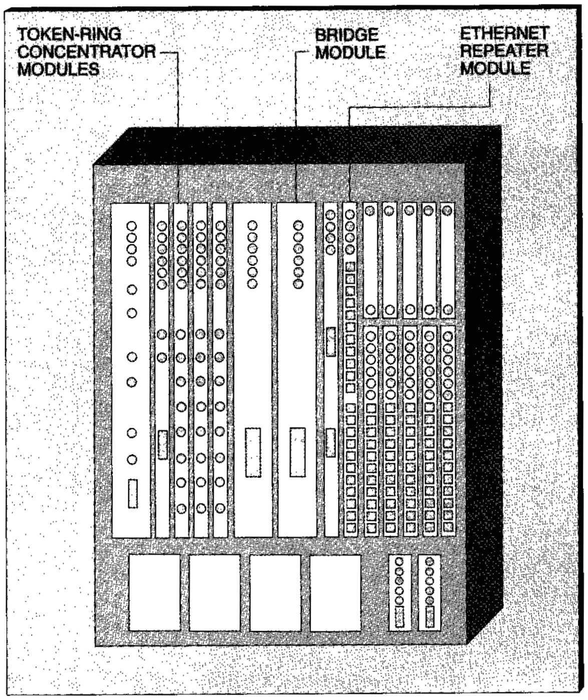

Remote management of network devices, including hubs in wiring closets, is possible via the Simple Network Management Protocol (SNMP). SNMP is an Internet standard defined by Request for Comments (RFC) 1157 and is the most commonly used management protocol for LANs.[12] In SNMP, a device maintains a database of management information called the Management Information Base (MIB). The MIB contains variables for specific items of interest such as number of frames received and transmitted, number of frames in error, and so on. These variables are remotely accessed via a network manager. The IBM NetView/6000* is one such SNMP-based network manager. There are standard MIBs defined for common network devices, including concentrators, repeaters, and bridges.

The remote SNMP network manager can perform three operations on a MIB variable (e.g., an error counter)—GET, SET, and TRAP. GET allows a manager to retrieve a variable, SET sets the value of a variable, and TRAP allows the device to alert a network manager when a specific defined event has occurred (e.g., when an error counter overflows). In addition to the MIB, each remote device contains an SNMP agent and protocol stack. The SNMP agent is the software that reads and writes the MIB in response to received GET and SET commands. The protocol stack, typically Transmission Control Protocol/Internet Protocol (TCP/IP), is used as the transport protocol for SNMP between the managed device and the SNMP network manager. The SNMP network manager can display MIB data directly on the screen as text or via a visual application. A visual application could show a hub module with enabled and disabled ports in different colors, support point-and-click interfaces to access information, and so on. Another Internet standard is Remote Monitor (RMON), described in RFC 1271. Devices that support RMON (e.g., supporting RFC 1513 for token ring) have capabilities for frame tracing, detailed traffic statistics, and so on. RMON data are accessed via a remote network manager.

Port switching is the ability to remotely change the LAN segment to which a workstation is attached. A port switching hub contains multiple high-speed buses where each bus is a separate LAN segment. Via a network manager it is possible to define the attachment of each workstation to a specific LAN segment within the hub. Without port switching, to change the LAN segment to which a workstation is attached would require a manual cabling change in a wiring closet. Thus, port switching can save labor costs for reconfiguring users. Port switching can also be used to localize traffic to individual segments and to isolate stations for troubleshooting. The introduction of port switching

SOME INTERNET STANDARDS FOR MANAGEMENT OF NETWORK DEVICES:

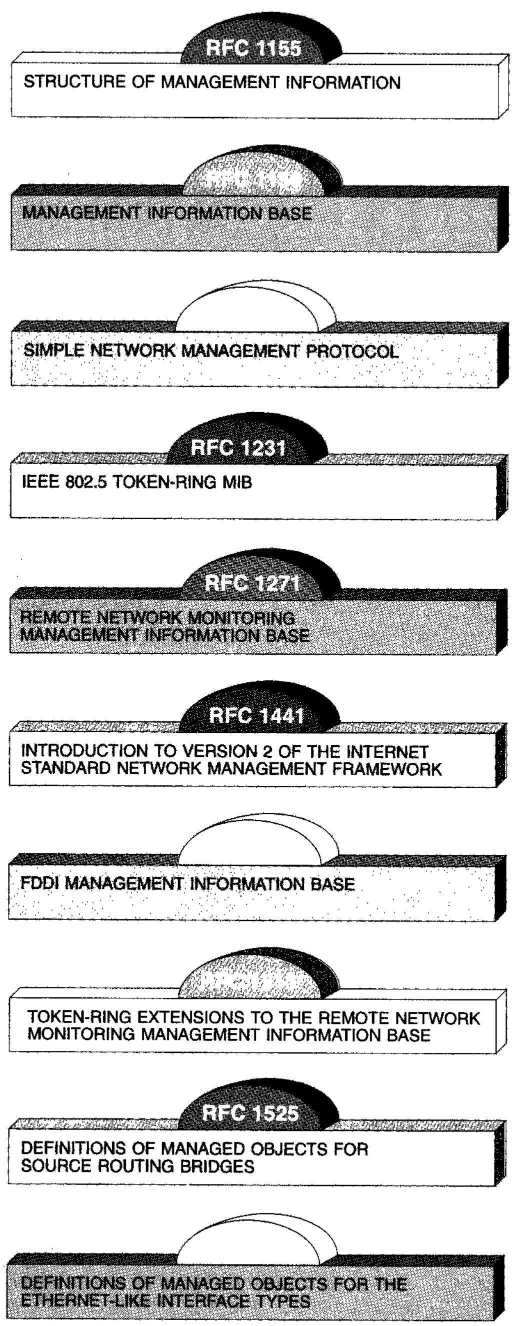

FDDI = FIBER DISTRIBUTED DATA INTERFACE
MIB = MANAGEMENT INFORMATION BASE
RFC = REQUEST FOR COMMENTS

Figure 7 Internetwork of connected LANs

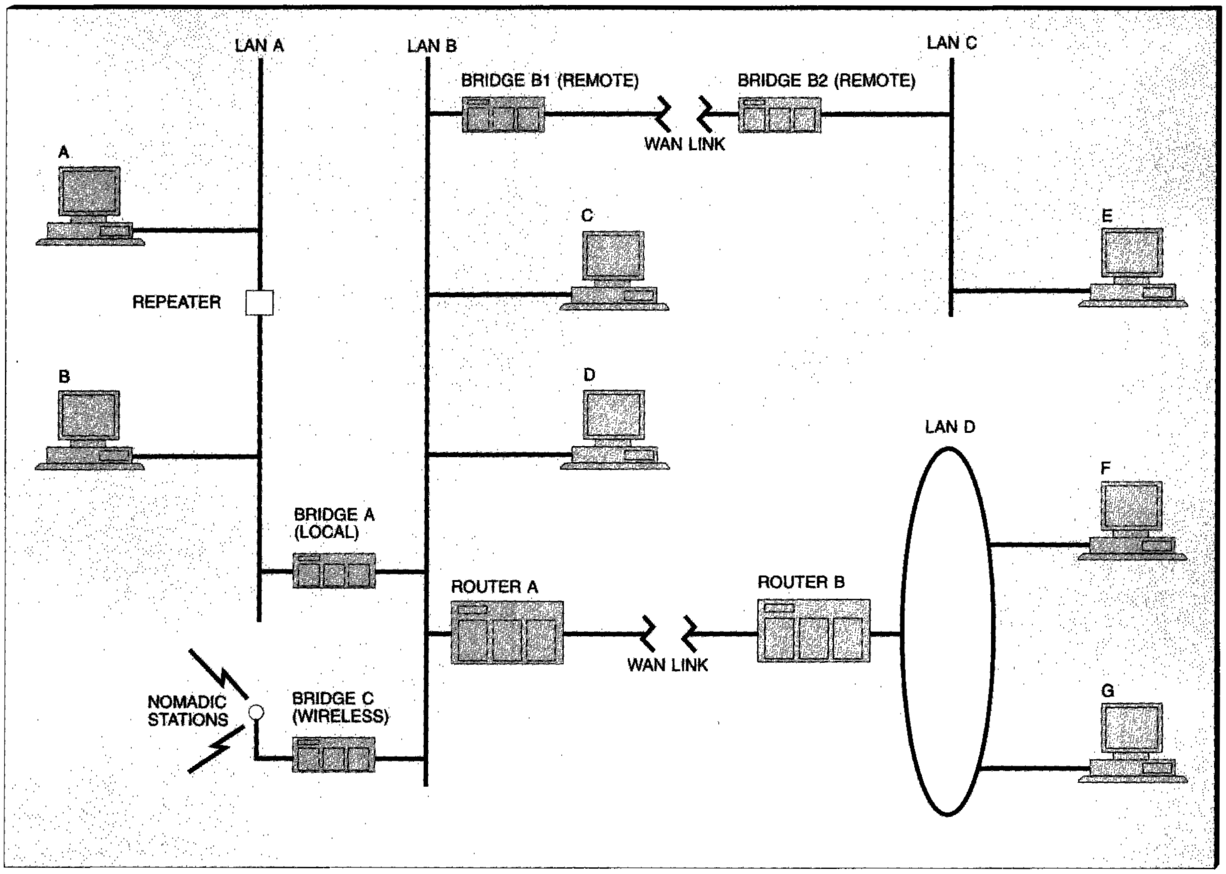

was the first step toward a "virtual LAN." In a virtual LAN, the LAN segment to which a station is attached is entirely independent of the physical location of the workstation or its physical attachment in a wiring closet. Virtual LANs are described later in this paper.

Other capabilities in the wiring closet include bridging and routing. A bridge can be packaged in a hub module while routers, being more complex than bridges, are usually stand-alone units. Bridging and routing are often used to connect LAN segments between, or even within, wiring closets. For example, if a single wiring closet supports token-ring and Ethernet LAN segments, an Ethernet-to-token-ring translating bridge (or router) is necessary to provide connectivity between the two different LAN types.

Internetworking for connectivity. Bridges and routers are used to build large enterprise networks. Bridges act as address filters, thus allowing LANs to be interconnected with only nonlocal traffic being forwarded between LANs. Traffic that cannot be identified as local to a LAN segment is transferred across a bridge. Bridges also forward all traffic sent to a broadcast, or "all stations" address. Routers also act as filters, but at a higher layer protocol, for example the Internet Protocol layer of TCP/IP.[13] Not all protocols can be routed. Routers can filter broadcast traffic and are often used as "fire walls" to contain broadcast traffic to specific LAN segments. Figure 7 shows an internetwork of LANs. Shown in Figure 7 are a repeater, local (Bridge A), remote (Bridges B1, B2), and wireless (Bridge C) bridges, and two routers connected over a wide area network (WAN) link. The wireless

Figure 8 Fields in a token-ring frame for TCP/IP

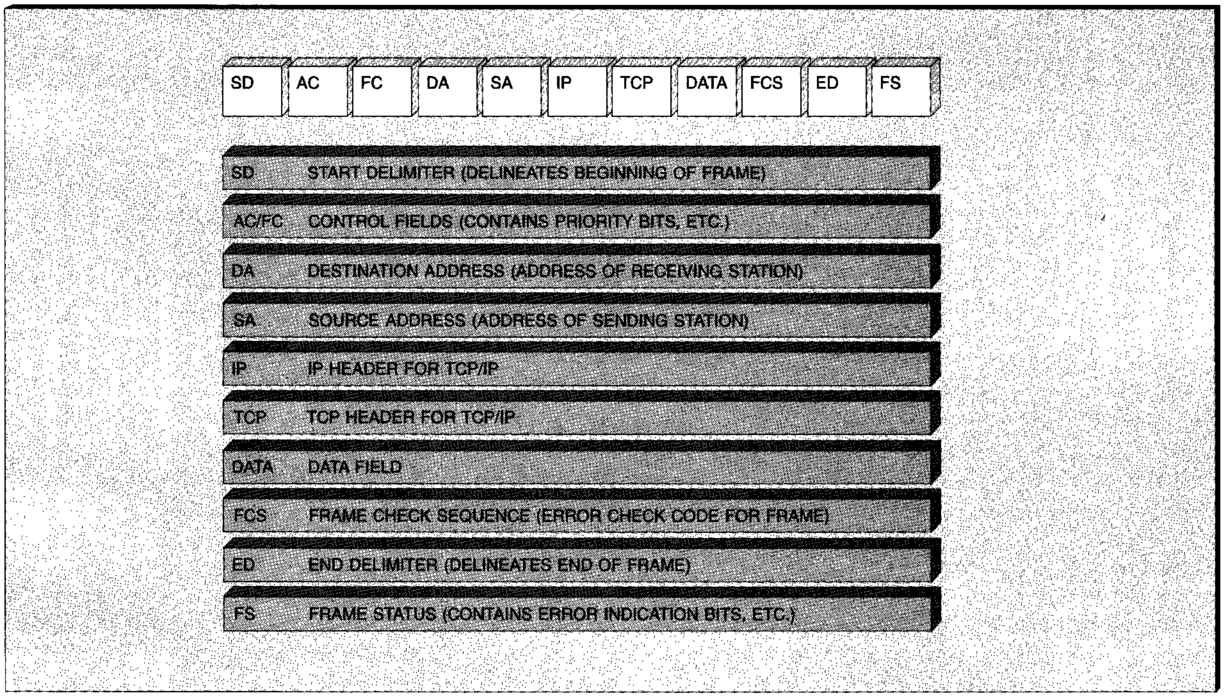

bridge in Figure 7 connects nomadic stations to LAN B via radio-frequency or infrared links.

There are two types of bridges—transparent and source routing. Transparent bridges[14] are typically associated with Ethernet and source routing bridges[14] with token ring. Bridges can interconnect geographically neighboring LANs via local bridging (Bridge A in Figure 7), or geographically distant LANs via remote bridging (Bridges B1, B2 in Figure 7). For remote bridging, the link connecting the two bridge halves will typically be a relatively low-speed WAN link. Options for the WAN link include Integrated Services Digital Network (ISDN) links at 64 to 128 Kbps and frame relay services at up to 1.544 Mbps. A WAN link with a much lower data rate than a LAN is a performance bottleneck if large amounts of traffic need to be forwarded across this link.

Transparent bridges filter traffic based on the destination address contained in each frame received. Figure 8 shows a token-ring LAN frame and its address fields. The bridge "learns" which source addresses are on which side of the bridge and stores these addresses in forwarding tables. Source addresses are unique for all stations. Learning occurs as a bridge copies all frames on a LAN segment and examines the source address of each frame. Any time a frame is received with a source address not in the forwarding tables, the source address is entered into the forwarding table. Subsequent frames received with a destination address on "this side" of the bridge are then not forwarded. For example, in Figure 7, Bridge A learns that stations A and B are on the left side of the bridge and stations C, D, E, F, and G are on the right side. Traffic between stations A and B is not forwarded across the bridge. Traffic between stations A and C is forwarded. With transparent bridges, the network must be a spanning tree configuration—that is, there can be no parallel bridges (parallel bridges connect between the same pair of LAN segments). Parallel bridges would confuse the learning process since a single station may not appear to be on both sides of a bridge. Transparent bridges execute a protocol defined by IEEE 802.1D[14] that enforces the spanning tree configuration. Parallel bridges are automatically detected and disabled from frame forwarding. The disabled bridges still participate in the IEEE 802.1D spanning tree protocol and are au-

Figure 9 LAN segmentation for improved performance

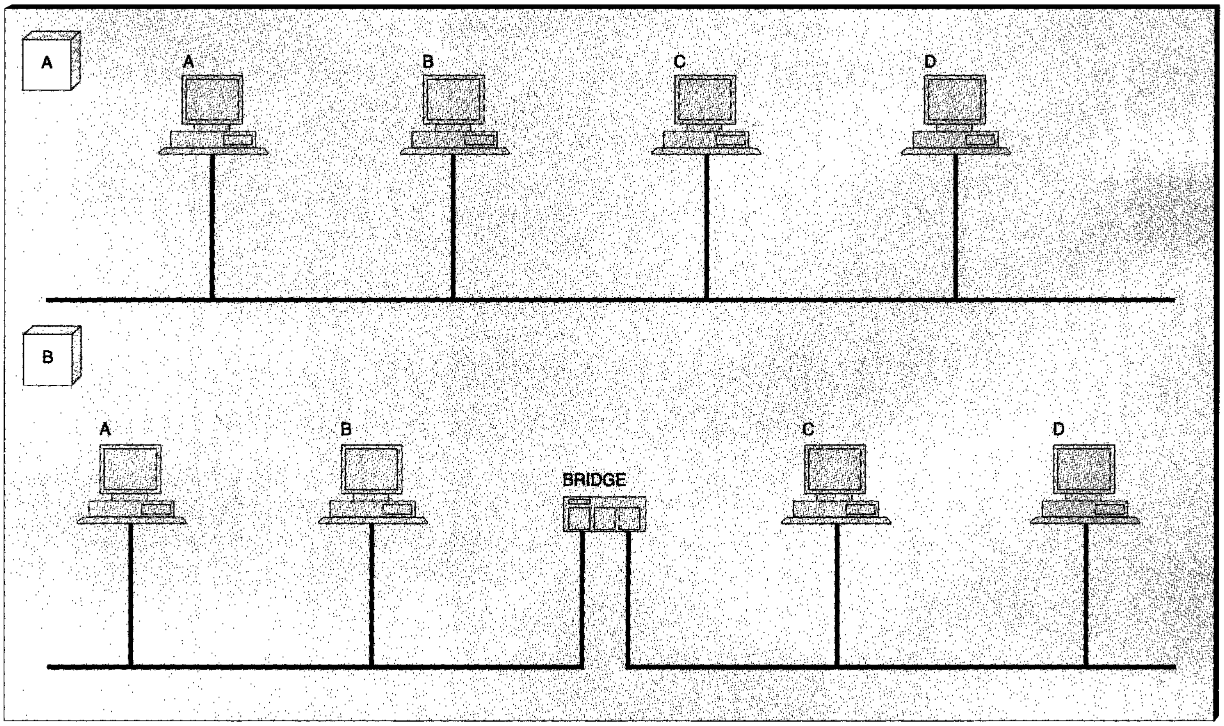

tomatically activated if a parallel frame-forwarding bridge should fail.

Source routing bridges filter traffic by a special Routing Information Field (RIF) inserted into a LAN frame by a sending station. The RIF immediately follows the source address field in a LAN frame and is defined in the IEEE 802.5 token-ring standard. The contents to be placed in the RIF are typically determined via a discovery process initiated by the sending station. In the discovery process a sending station broadcasts a special "discovery frame" destined for the intended receiving station. Each bridge that this frame crosses adds its unique two-byte route designator to the RIF field. Communicating stations then save and use the RIF field for one of possibly several routes. Reference 14 describes the discovery process for networks interconnected with source routing bridges. With source routing, active parallel bridges are possible since each bridge will contain a unique two-byte route designator.

Routers filter LAN traffic based on fields contained within protocol headers. These protocol headers occur "inside" the LAN frame after the destination and source address fields. The protocol headers contain address fields. The addresses in these fields contain subfields identifying the location, or LAN segment, on which the sending and receiving station are located. TCP/IP is a commonly routed protocol. Protocols that do not contain address fields with location-specific subfields cannot be routed (but they can be bridged). For example, NetBIOS* cannot be routed since its protocol headers contain only information on station names and not their locations. Thus, routers are designed to handle specific sets of protocols, for example Internet Protocol (IP) and Systems Network Architecture (SNA). Routers perform functions beyond those of a bridge including the determination of optimal network routes to minimize network traffic, provide congestion control, and filter broadcast frames. The filtering of broadcast traffic is very valuable when a relatively low-speed WAN link is used to interconnect two LANs (e.g., routers A and B in Figure 7). Often, routers can also function as bridges to be able to forward frames for non-routable protocols.

Segmenting LANs for improved performance. As the utilization of a LAN increases, user response

Figure 10 LAN backbone for improved performance

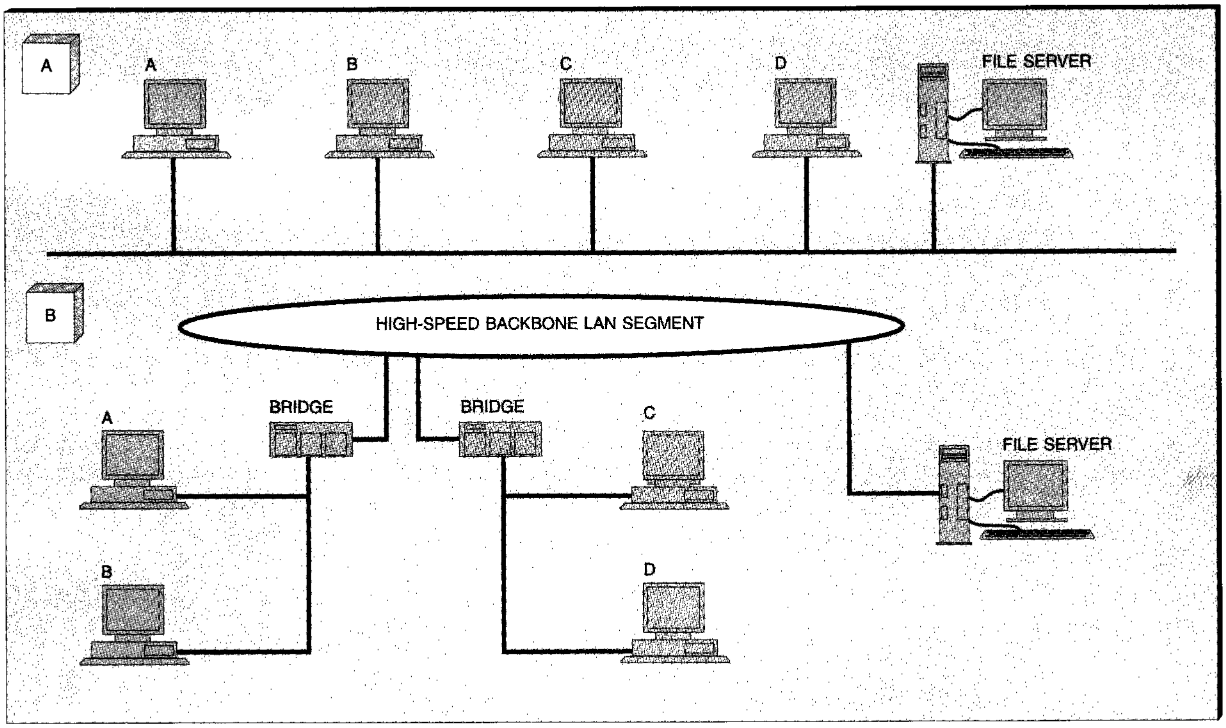

times for file transfers and interactive computing also increase. This is due to the additional delay for accessing the LAN and transmitting a frame. One way to decrease the utilization of a LAN segment is to reduce the number of stations on the segment. By localizing stations that communicate primarily among each other, overall utilization can be decreased. As a simple example, Figure 9A shows a single LAN segment with four users. In this example, assume that it has been found that users A and B communicate primarily with each other and users C and D communicate primarily with each other. By segmenting this LAN into two segments, attached by a bridge, the network utilization as seen by a workstation is approximately reduced by one half. This is shown in Figure 9B.

If, in the example of Figure 9, all four stations communicate with a common fifth single station, the solution of Figure 9B does not yield any performance improvements. In client/server computing it is typical to have many stations (the clients) communicating primarily with one station (the server). To improve performance for client workstations in a client/server network, a higher speed backbone LAN segment is needed. Figure 10A shows the four stations of Figure 9 with an added server station. Figure 10B shows a network where the combined traffic on each bridged segment is approximately half the segment traffic of the network in Figure 10A. The higher speed backbone segment can handle the bandwidth demands of multiple attached lower-speed LAN segments. Figure 10B shows what is called a distributed backbone with local LAN segments attached with bridges. These bridges must have sufficient buffering to handle "speed matching" of the traffic between the low-speed local and high-speed backbone LAN segments. The backbone LAN segment can be used to directly attach frequently accessed stations such as file servers, host gateways, and WAN attached routers. Fiber Distributed Data Interface (FDDI) is a 100-Mbps token-ring LAN technology well suited for distributed backbones.[15]

The backbone LAN for Figure 10 could also be implemented within a multiport bridge, router, or switch. This is commonly called a *collapsed back-*

Figure 11 Collapsed LAN backbone

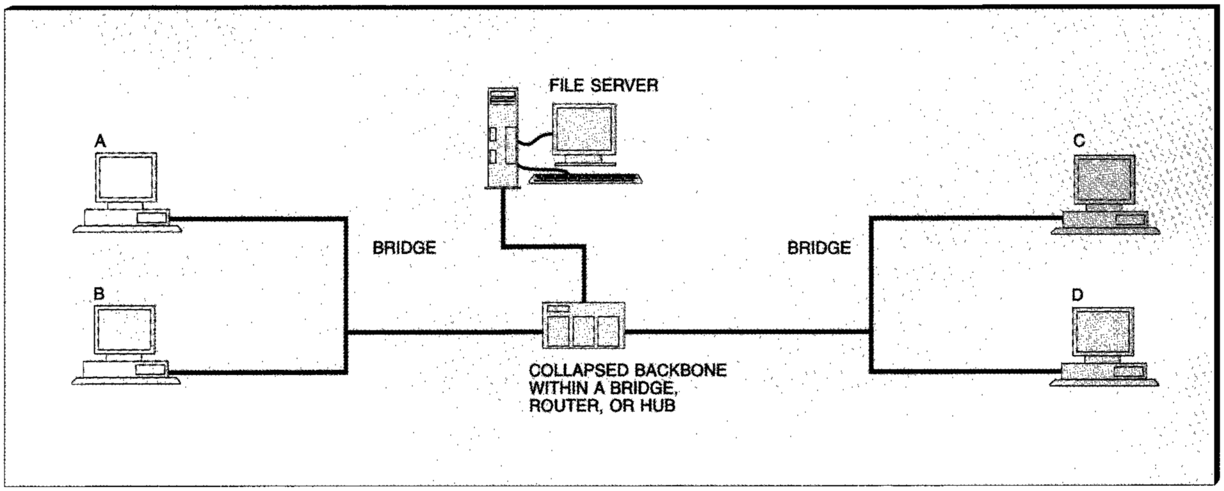

bone and is shown in Figure 11. A collapsed backbone physically brings all of the LAN segments into one wiring closet and into a single hub (as shown in Figure 6, for example). If the attachment of a file server or servers is via a higher speed LAN segment, this attachment is often referred to as a *fat pipe*. There are both advantages and disadvantages to the collapsed backbone. One key advantage is the ability to consolidate network devices in a secure, controlled, centralized location. This is especially viable for a single building or for clustered buildings. However, an industrial or university campus will probably find the distributed backbone concept to be more workable. One disadvantage of a collapsed backbone scheme is that the network is vulnerable to a single component failure.

Segmentation of LANs is often done on an ad hoc basis. When performance problems are reported, a LAN is segmented based on the best available network management data. For example, RMON data can be used to identify groups of communicating stations that could be isolated on separate segments. Networks are often configured on organizational or purchasing boundaries; for example, each department in a large organization will have its own LAN segment. A careful analysis of existing and future applications can be done to determine bandwidth requirements for LAN-attached workstations.

Motivations for the future—Bandwidth and management

Much has been written about the anticipated upcoming explosion of multimedia applications and their huge network bandwidth requirements (see, for example, Reference 16). A question that remains unanswered is whether high-bandwidth capacity fiber optic cabling should be installed to every desktop.

Requirements for increased bandwidth. This section describes bandwidth requirements for existing and future applications. The need for bandwidth is a two-dimensional question, with the main discussion points being: where and when. The "where" has two main groupings, desktop and network, with desktop encompassing all desktop machines, from Intel-based DOS systems to high-end PowerPC*AIX* workstations. The "when" is presumed to be the 1997–1998 time-frame. With LAN switching, the existing infrastructure of LAN adapters and cabling can accommodate even the most optimistic future for multimedia.

While there are many different types of LAN applications, they can be broadly categorized into three different types,[17] each with different bandwidth requirements. The three types of LAN applications are data, client/server multimedia, and interactive multimedia. Examples of data applications include file transfers and electronic mail. Client/server multimedia applications include video-on-demand for distance learning and entertainment. Interactive multimedia applications include videoconferencing and real-time voice (e.g., telephony on LAN).

The outstanding characteristic of data traffic is its extreme random bursts, which is why shared bandwidth models such as token ring and Ethernet can be so effective. If one assumes that 2 Mbps were available for data to the end user, then all but the largest file transfers would see subsecond transmission times, and most traffic would see transmission times in the order of 10 to 100 milliseconds. With such small delays, the network is significantly more responsive than all but the most lightly loaded server systems.

Multimedia traffic is typically continuous in its nature (e.g., video and audio) and thus requires continuous guaranteed bandwidth over an extended period of time, unlike the short-term bursts of data traffic. Video clips may last several minutes and a videoconference may last several hours. Throughput demands for client/server and interactive multimedia are about the same. However, an interactive multimedia application requires very high quality bandwidth. High quality bandwidth has low delay and low variance in delay (or low "delay jitter"). Why are delay and delay jitter important? For human conversations, the maximum tolerable delay in each direction is about 150 milliseconds, as determined from the experience of the long distance carriers.[16] Delay jitter is removed at the receiving station by buffering a set amount of received data. This buffer will allow a constant rate of data to be passed to the multimedia coder/decoder (CODEC) hardware despite variances in the interarrival times of frames from the network. If the delay jitter is large, a large "jitter buffer" is needed. The size of the jitter buffer directly affects end-to-end delay.

Video traffic has a wide variation of possible throughputs, from 96 Mbps for an uncompressed digitized television signal to 64 Kbps for a very slow frame rate videoconference signal. A common standard for video transmission is MPEG-2 (Moving Pictures Experts Group).[18,19] MPEG-2 produces television-quality full-screen video and requires 6 Mbps of transmission bandwidth. MPEG-2 and other video standards employ various compression methods to reduce the very high "raw" bit rate of video to much lower rates. As compression technologies continue to advance, even lower bandwidth requirements can be expected (see Reference 18 for a discussion of 64-Kbps MPEG-4), or lower cost for existing compression rates. Voice requires 64 Kbps of bandwidth for standard digitized voice as used, for example, in the integrated services digital network (ISDN). Lower compressed bandwidths of 8 Kbps are possible.

Voice is generally left out of most bandwidth discussions, since it is a known fact that channelized carrier systems carry toll quality voice at 64 Kbps or less. However small, voice bandwidth needs should not be dismissed as trivial. First of all, the human ear is remarkably sensitive to deviations in voice quality, including omissions and pauses that can be caused by excessive delays, or data loss from overflows or line errors. Delay jitter is a significant issue when transmitting voice over non-time-multiplexed systems. Delay jitter can be reduced by careful reservation of end-to-end bandwidth, by prioritization of delay-sensitive traffic, or the approximate allocation of excess bandwidth. For example, estimating 1 Mbps per voice stream can provide sufficient margins to reduce delay jitter to an acceptable amount in most networks.

The total bandwidth requirement to a typical desktop workstation is less than 10 Mbps to support simultaneous data, video, and voice; about 2 Mbps for data, 6 Mbps for video, and about 1 Mbps for voice. There may be exceptional applications that may require very high bandwidth (e.g., 100 Mbps) to the desktop. These applications include supercomputer scientific visualization and very specialized medical applications. Even the case of multiple video windows on a single screen does not require more than the nominal bandwidth. As a video window is reduced in size, the bandwidth requirement for the video stream is also reduced. Thus, one large video window or several smaller video windows will all require about the same bandwidth. Given the need for less than 10 Mbps to the desktop, any of the three "networks of choice"—10-Mbps Ethernet, 16-Mbps token ring, and 25-Mbps ATM—can be tailored to meet the needs of the end user, and all three are supported by the widest possible transmission media, including UTP categories 3, 4, and 5 as well as STP.

Where future applications will have their biggest bandwidth impact is in the network infrastructure

Table 1 Features of telecommunications transmission methods

System	Cable Type	Data Rate (Mbps)	Symbol Rate 10^6/sec	Symbol Type	Encoding: Bits per Symbols
10BaseT	UTP-3[2]	10	20	2-level (binary)	1 per 2
Token ring	UTP-3[2]	16	32	2-level	1 per 2
ATM[1]	UTP-3[2]	25.6	32	2-level	4 per 5
ATM[1]	UTP-3[2]	51	12.75	16-CAP[4]	4 per 1
100BaseT[1]	UTP-3[2] (4 pairs)	100 (33.3 Mbps/pr)	25	3-level	4 per 3
100 VG-AnyLAN[1]	UTP-3[2] (4 pairs)	100 (25 Mbps/pr)	30	2-level	5 per 6
FDDI	UTP-5[3]	100	125	3-level	1 per 1
ATM[1]	UTP-5[3]	155	155	2-level	1 per 1

Notes:
[1] Emerging standard.
[2] Operation over UTP-5 and STP is also available.
[3] Operation over STP is also available.
[4] Carrierless Amplitude Modulation/Phase Modulation (closely related to Quadrature Amplitude Modulation as described in Reference 20).

and at the video servers. Video servers, depending on the number and characteristics of attached clients, may require bandwidths of 100 Mbps and up. Introducing LAN switches in the infrastructure can support both the network and desktop bandwidth requirements for most foreseeable data and multimedia applications.

Requirements for improved management. Small LANs with less than ten stations require little management. Changes to a small network can be handled manually. With large networks of hundreds to thousands of users, management is much more difficult. Because organizations are dynamic, the ability to easily add, remove, and move users in a network is important. It is estimated that one-half of all network users move geographic locations (e.g., move between new offices) within one year.[3] New users are also added and departing users removed from a network. User adds, removes, and moves are estimated to be the most significant cost in the operation of a large network.

Much of the cost of moving a user comes from required software configuration changes within the moved workstation. In other words, a workstation software configuration that works on one portion of a network may not work in another. Workstations that use TCP/IP, for example, cannot be easily moved between LAN segments connected with routers. TCP/IP defines a LAN segment or a network of bridged LANs as a *subnetwork*. Individual subnetworks are connected with IP routers. The workstation IP address contains specific information describing the subnetwork to which a workstation is attached. This address is part of the IP header of a LAN frame. This network location-specific IP address is needed for routers to correctly forward IP frames between subnetworks.[13] Thus, moving a workstation from one subnetwork to another requires that its IP address be changed. This IP address change involves manually configuring software files in the workstation and may also require changes to be made in address tables of network routers. These changes incur labor costs—costs that are typically increasing rather than declining. Thus, there is a motivation for future LAN technology to be able to simplify user moves within a network and between protocol-defined subnetworks. The virtual LAN, described later in this paper, is one solution to simplifying network moves.

Increasing data rates. The 1990s brought higher data rates to the desktop. It was understood that to be widely accepted, new offerings intended for the majority of office applications must be able to operate over existing or planned building cabling. This meant that higher data rate LANs must operate over available, standardized (shielded or unshielded), twisted pair cabling.

The FDDI committee (American National Standards Institute, or ANSI X3T9.5, now X3T12) began a Twisted Pair Physical Media Dependent (TP-PMD) Working Group in 1990. Advanced encoding methods were proposed for 100-Mbps operation over UTP category-3 cable. However, after five years, these FDDI-over-UTP-category-3 efforts have been largely discontinued. In 1991 six major companies, working together, completed a simple, cost effective, proposal for FDDI operation over STP. The following year, this STP proposal was jointly supported by eleven companies and given the name SDDI. SDDI products became available in 1992.

The FDDI TP-PMD standard providing for operation over STP and UTP category-5 cable completed final technical review in December of 1994. This FDDI transmission method utilizes three signaling levels to reduce the amount of higher frequency energy in the cable. This was necessary to meet United States Federal Communications Commission (FCC) regulations for the amount of high frequency emissions permitted from LAN cabling.

In other telecommunication systems, different transmission methods are being employed in order to offer higher data rates over UTP category-5 cable. Table 1 shows some of the key features of each of the methods. There are two new emerging standards for 100-Mbps technologies, "Fast Ethernet" (100BaseT) in the IEEE 802.3 and 100 VG-AnyLAN in the IEEE 802.12. 100BaseT is a ten-times "speed-up" of 10-Mbps IEEE 802.3 (10BaseT); 100 VG-AnyLAN uses a new demand priority protocol that supports Ethernet or token-ring frame formats.

Emerging trends—fourth-generation LANs

The emerging fourth generation of LAN technology is characterized by the introduction of very high-speed switching elements in the LAN backbone and in the departmental LAN or workgroup (see, for example, References 21–23). A workgroup is a collection of client workstations that intercommunicate and share common devices such as servers and printers. With the introduction of switching technology, very little has to change in the infrastructure itself; the network cabling, workstation LAN adapters, and workstation software do not have to change.

Switch technology and packet (i.e., frame) switching concepts have been implemented in voice and data networks for over a decade. Significant advances in integrated circuit technology now make it practical to consider frame switching at data rates that far exceed those that were possible just ten years ago. Most private branch exchange (PBX) switching systems developed in the mid-1980s provided full-duplex switching for 64-Kbps links. Today's technology can easily switch frames at LAN speeds of 10, 16, 25, and even 100 Mbps.

A LAN switch is very similar to a LAN bridge with multiple ports. The switch must forward incoming frames to the appropriate output port. Also, the switch must ensure that frames from two or more simultaneously transmitting workstations are not lost due to contention within the switch. LAN switches usually contain both input and output port buffering. LAN switches maintain address tables, or port forwarding tables, and SNMP MIBs similar to transparent bridges. If the switch supports only single workstation attachment per port (e.g., dedicated media and dedicated bandwidth) it is a workgroup switch. This dedicated LAN segment has two stations, the user workstation and the port of the switch. For collapsed backbone switches, LAN segments with multiple stations can be attached to each port. LAN switching not only increases the bandwidth available to a workstation, it can also enable software-defined segments called virtual LANs.

Switch architectures. A number of switch architectures have been developed and evaluated in both industry and academia. Industry product developers weigh switch performance versus switch cost in determining the optimum architecture for a specific environment or application. Switch cost is typically measured as a per port cost: total switch cost divided by the number of ports. Switch architectures usually include either disjoint or common buffer memory and input/output ports that are interconnected by a switch fabric. The switch fabric can be a cross-bar, a multistage switching array, or a high-speed bus.[24,25]

In a switch with input buffering each input link to the switch is buffered so that an incoming frame can be stored as it is being received. This may be necessary to avoid frame loss when two stations on separate ports attempt to simultaneously transmit to the same target workstation. Input buffering is also necessary for some broadcast and multicast schemes. A condition known as head-of-line blocking will occur when frames waiting in a queue cannot be forwarded to the appropriate destination buffer while the frame at the head of the queue is waiting for its destination path to become free. An alternative design allows the switch control logic to look at the address of the next frame or frames in the input queue to see if the destination path for that frame is free, and thus bypass the frame at the head of the queue. Performance studies have shown that overall switch throughput can be improved and queueing delays reduced if head-of-line blocking is avoided.[24]

In normal operation a burst of frames may enter the switch from different sources and be destined

Figure 12 A four-port cross-bar switch

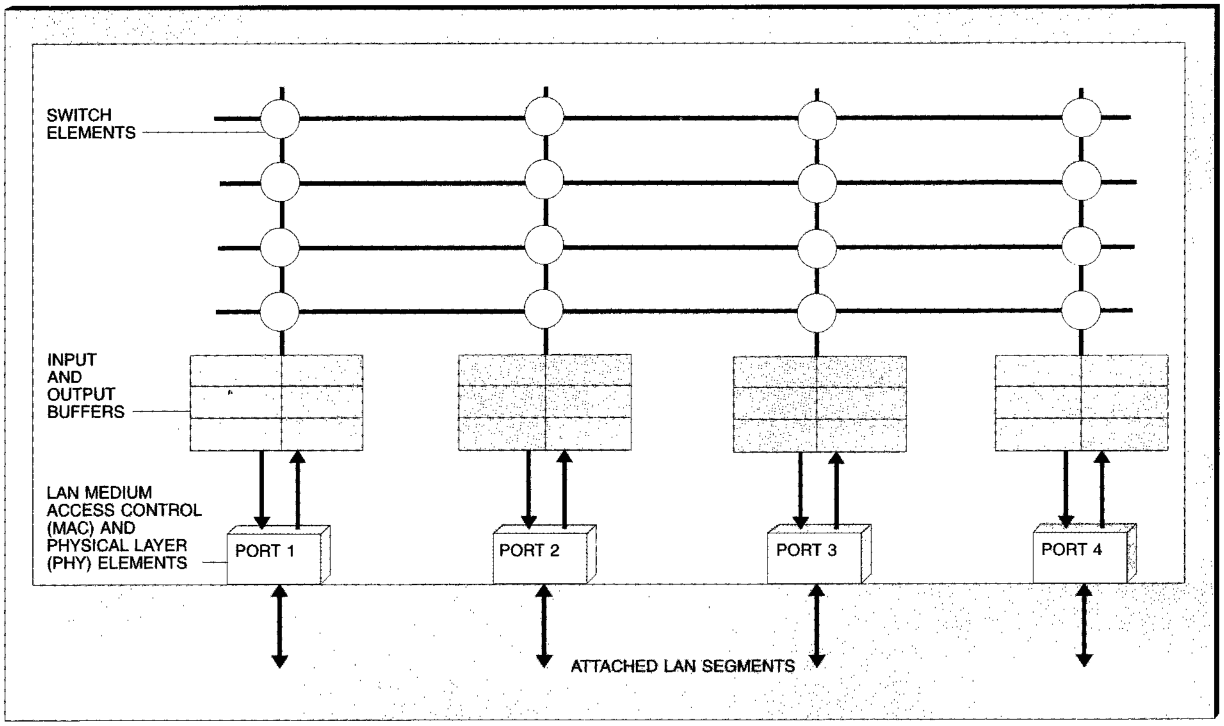

to the same output port (i.e., same target station). With output buffering, the frames can move through the switch to the appropriate output port rather than remain at the input buffer where they can block other frames from traversing the switch. Common buffer pool designs usually provide multiple output queues within a shared-memory space. High-speed, multiaccess memory is often required to permit the simultaneous transfer of multiple data streams both into and out of the shared memory. Output port buffering is valuable if a port is attached to a shared-media LAN, such as a token ring or Ethernet. Since the port must contend with other stations for access to the LAN segment, frames may be moved from the input buffer to the output buffer prior to transmission. Output port buffering is thus needed if ports operate at different speeds or if traffic loads are asymmetrical (i.e., a single port is the destination for a large amount of traffic from many other ports). A busy output port may have several frames in its output transmit queue awaiting transmission. New frames arriving at other ports will be moved to this output buffer, assuming that the internal paths are not busy.

Switch architectures can be classified as blocking or nonblocking. In a nonblocking architecture, a frame being forwarded between two ports cannot prevent, or block, the forwarding of another frame between two other ports. Figure 12 shows a four-port switch with a cross-bar switch and both input port and output port buffering. For example, in the switch of Figure 12, a frame being forwarded from port 1 to port 4 does not prevent frames from being forwarded between ports 2 and 3. Cross-bar and multistage Banyan architectures[25] are nonblocking. Switches that do not avoid internal "collisions" are blocking architectures. These architectures may rely upon internal retries in successive cycles. If a frame cannot be forwarded after a set number of retries, it is discarded. Nonblocking architectures are marginally more expensive than blocking architectures. In a nonblocking architecture the total internal switch band-

Figure 13 Switch forwarding tables and central processor

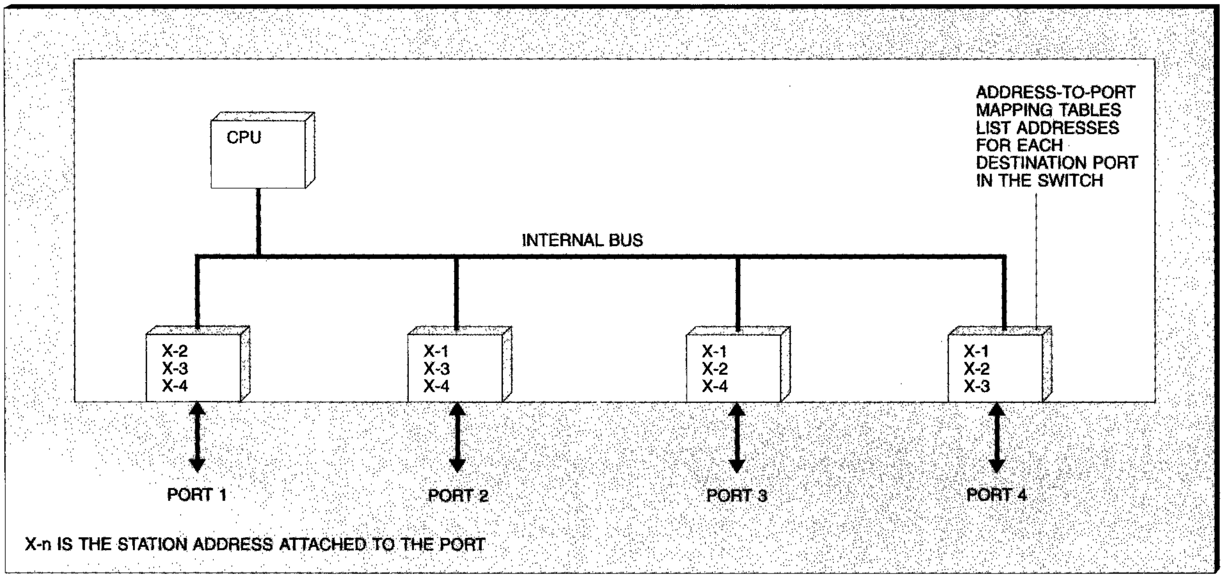

width is higher than in a blocking architecture. It is this additional bandwidth that adds to the switch cost.

All LANs today rely upon broadcast and multicast capabilities to reach multiple stations with a single frame transmission. Multicast is the ability for one frame to be received by many stations, each of which is a member of the designated multicast group. A broadcast frame is received by all stations. Thus, another function of the interconnect switch is to support broadcast and multicast frame forwarding to multiple output ports.

LAN switches learn and maintain address tables in similar fashion to transparent bridges. Each switch port contains a forwarding table that correlates switch port numbers with LAN medium access control (MAC) destination addresses. When a frame is received at the switch port, its destination address is compared to all addresses in the port forwarding table to find the correct destination port. Port forwarding tables are typically implemented with Content Addressable Memory (CAM). A CAM allows for very rapid comparison of a received address with all the addresses in a large table. Similar to transparent bridges, address learning and table updating can be handled by a central processor in the switch. The central processor can also handle switch management tasks such as SNMP MIB updates and, as described later in this section, maintain virtual LAN tables. LAN switches based on source-routing frame forwarding are also possible. Figure 13 shows a LAN switch with its internal forwarding tables and central processor.

Cut-through versus store-and-forward operation. Various buffering and internal frame forwarding schemes have been studied and reported in the literature.[24,26] Two internal frame forwarding schemes are possible, cut-through and store-and-forward. In cut-through switching, frames are forwarded through a switch "on the fly." The destination address of a frame arriving at an input port is compared against the port forwarding table as the frame is being received. Assuming that the path between the input port and destination port is available (i.e., no other frame is being moved to the destination port at this instant in time), the switch can immediately begin moving the frame to the destination port buffer. Then, assuming that the destination port is not actively transmitting a frame at this instant in time (i.e., the port buffer is empty), the frame transmission to the destination station can begin. This transmission occurs in parallel with the continued movement of the frame from

Figure 14 Cut-through switching

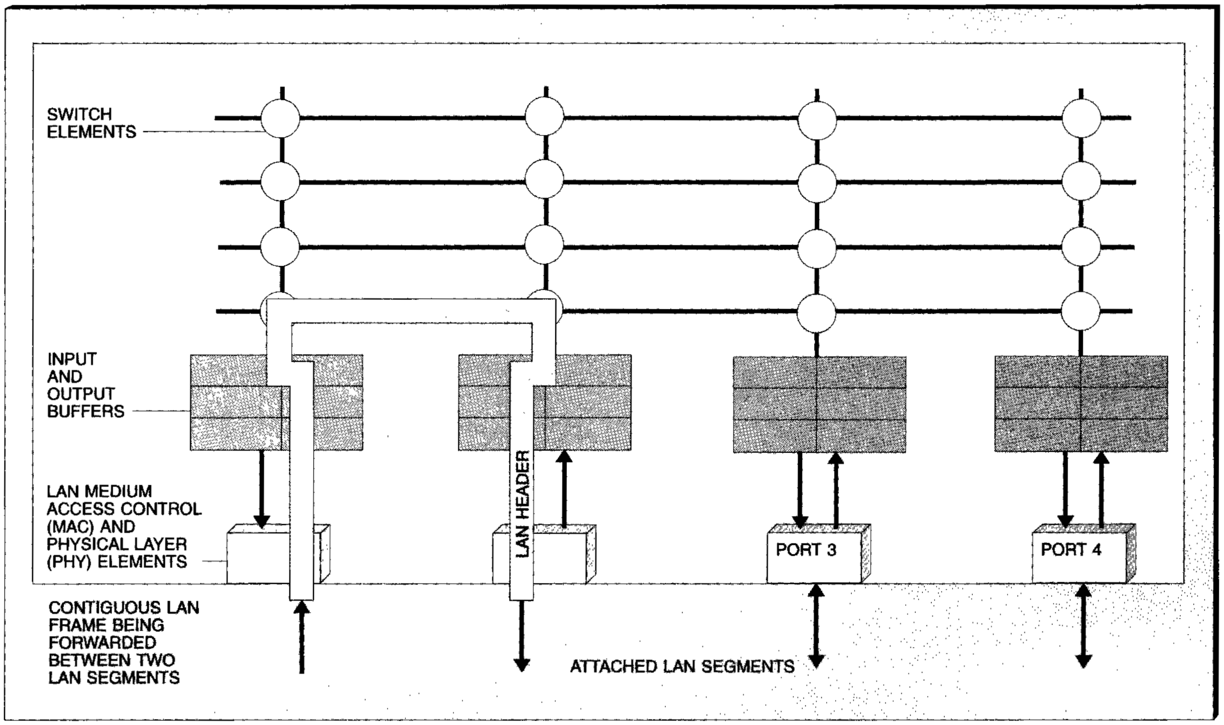

the input port. Figure 14 shows a frame being forwarded via cut-through switching in a cross-bar switch.

In store-and-forward switching the entire frame must be completely received into the switch buffer prior to the beginning of transmission at the destination port. Cut-through switching reverts to store-and-forward switching when the output port for a received frame is busy. In this case, the received frame is stored, either in the input or output port depending on the architecture, before being transmitted onto the destination LAN segment. Thus, in the case of a destination LAN with high levels of utilization, cut-through and store-and-forward switching provide equivalent delays.

Cut-through operation has much lower latency at low-to-moderate loads than does store-and-forward operation. A cut-through switch can begin transmitting a received frame onto its destination LAN segment within a latency of tens of microseconds—the time required to match an address in the port table. For store-and-forward switching, latencies of several milliseconds are incurred for large frames (e.g., 2 milliseconds for a 4-kilobyte frame at 16 Mbps). For small frames, cut-through and store-and-forward operation have essentially the same latency. Some protocols and applications exhibit poor performance in networks with large latencies. For example, interactive multimedia applications require very low latencies. One disadvantage of cut-through switching is that it forwards frames with errors. In store-and-forward switching, since frames are fully received before being forwarded, error checking can be done and erroneous frames discarded.

Full-duplex operation. LAN adapters support a MAC (medium access control) protocol to regulate access to a shared bandwidth LAN segment. In a switched LAN with a dedicated LAN segment for each station, the requirements for a MAC protocol are minimized. No longer must a station "contend" with other stations for permission to transmit onto the medium; a station can transmit whenever a frame is queued in the adapter. A station can also receive at any time. Full-duplex operation is then

Figure 15 Full-duplex operation

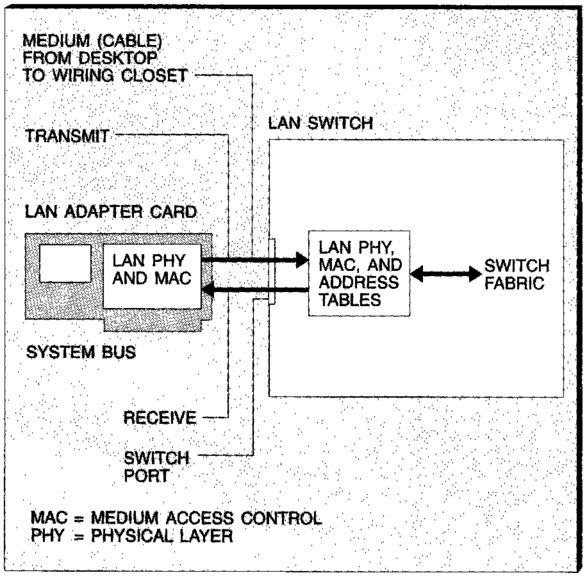

possible; a station can transmit and receive independently of other stations attached to a LAN switch. Depending upon the capabilities of the station and the interconnect switch to which it is attached, sustained media-speed transmit rates are possible. Even more significant, the adapter can also receive frames at the media rates on the receive link, for an aggregate capacity of double the normal half-duplex rate. Full-duplex operation has been proposed for Ethernet in the IEEE 802.3[27] Standard and for token ring in the IEEE 802.5[28,29] Standard. Figure 15 illustrates a full-duplex adapter and switch port.

A key aspect to full-duplex operation is that the standard LAN (e.g., token-ring and Ethernet) frame formats are preserved without change. Full-duplex operation is transparent to the user. Thus, the software applications that already exist do not have to be changed to support this enhanced mode of operation.

Throughput analysis for switched LANs. In shared bandwidth LANs the total bandwidth is shared among all attached workstations. For example, 50 stations with similar data transmission requirements on a shared 16-Mbps token ring can each expect an average transmit bandwidth of 16 Mbps / 50 = 320 Kbps. The peak bandwidth available to a workstation is the full 16 Mbps under light load conditions. Similar per-station bandwidth degradation is typical of other shared bandwidth LANs, such as 10-Mbps Ethernet and 100-Mbps FDDI.[30,31] In a shared-bandwidth LAN segment, the addition of each new station decreases the average bandwidth that any one station may acquire during high load conditions. For a dedicated bandwidth LAN, the addition of a new station does not provide less bandwidth to every station. Figure 16 is a graph showing the average transmit bandwidth available to each station on shared and dedicated-bandwidth LANs, assuming that ideal load distribution and bandwidth efficiency are achieved.

The analysis shown in Figure 16 for dedicated bandwidth is sensitive to load distribution and changes if the stations are randomly transmitting to each other. If two stations try to transmit to the same station at exactly the same time, then one of the frames will get through while the other will be temporarily buffered in the switch. There will be periods of times (statistically) when several stations try to transmit to the same station, thus resulting in multiple frames being queued. Researchers have shown that, in general, for input buffered systems, less than ideal throughput is achieved between two ports when there is contention for the output port.[9] The input queue lengths of all ports will increase if a significant portion of the switch traffic is targeted for a single port that is operating at near capacity. More efficient frame service disciplines, such as replacing a first-in first-out service scheme with an input windowing scheme, can significantly reduce delays for frames destined to uncongested ports.[9]

In client/server computing, a single station, the file server, is often the target or source for most traffic from and to a large number of client workstations. For such configurations, a higher bandwidth "fat pipe" attachment is needed for the file server. For example, a single 100-Mbps Fast Ethernet port can be used to attach a file server that is accessed by a multitude of 10-Mbps Ethernet client workstations. The client workstations attach to the switch via 10-Mbps ports. This is an example of scaled bandwidth, providing the needed amount (but not more) of bandwidth to each class of attachment.

Switched LAN topology—workgroup and backbone. LAN interconnection can be defined at the workgroup and backbone levels. LAN switches will play a key role in both of these areas.

Figure 16 Performance of shared and dedicated bandwidth LANs

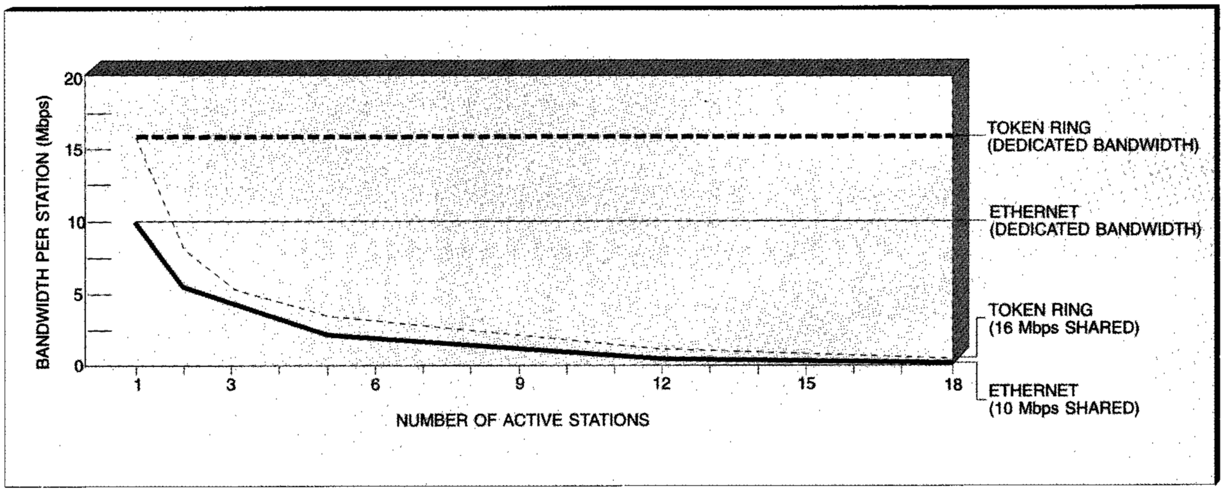

In a workgroup, the client workstations may have high-bandwidth requirements as characterized by some multimedia or computer-aided design applications. Today, a workgroup is typically constrained to the workstations attached within a single LAN segment. Virtual LANs will allow workgroups to be defined largely independent of geographical constraints. As workgroup applications expand to consume more network bandwidth, the workgroups can be further segmented. Workgroup LAN switches can replace existing concentrators, repeaters, or hubs, to provide an overall increase in the aggregate bandwidth available to the workgroup. By simply replacing the repeater of a 20-port Ethernet workgroup with a 20-port dedicated LAN switch, the effective transmit bandwidth available to each user will increase by a factor of 20. Full-duplex operation provides an additional avenue for dedicated bandwidth growth. A workgroup switch may have a high-speed port for file servers, or other frequently accessed devices. Figure 17 shows a workgroup switch with attached client workstations (at 10 Mbps or 16 Mbps dedicated) and a high-speed "fat pipe" port for a 100-Mbps full-duplex file server.

From a practical viewpoint, workgroup stations do not need sustained peak throughput on a continuous basis. Thus, a workgroup switch must only be able to support the full aggregate bandwidth for traffic in bursts, with sufficient buffering to withstand short bursts from multiple source ports to a common destination port. Collapsed backbone switch nodes can be expected to handle more network traffic than workgroup switch nodes. For this reason, the internal bandwidth of the switch, as well as its internal buffering capacity, will be greater than that of a workgroup switch. Backbone switches that rely upon traditional transparent bridging concepts for frame forwarding will require greater address storage and handling capacity than workgroup switches. Due to the importance of the network backbone for interconnectivity, critical components that are subject to failure should be replicated and be made quickly replaceable with an available spare.

The virtual LAN. A physical LAN segment is defined by the one or more concentrators or repeaters into which workstations are attached. LAN segments are interconnected by bridges or routers. LAN segments that are isolated by bridges have the following properties:

- All frames with destinations on the segment remain local (i.e., are not forwarded by the bridge). Local traffic is thus private to the LAN segment.
- All frames with destinations not on the segment are forwarded out of the segment by the bridge.

If the LAN segments are isolated by routers, then the above properties hold and, in addition, broadcast frames do not leave the segment. Broadcast frames are received and processed by all stations,

Figure 17 A high-performance LAN workgroup example

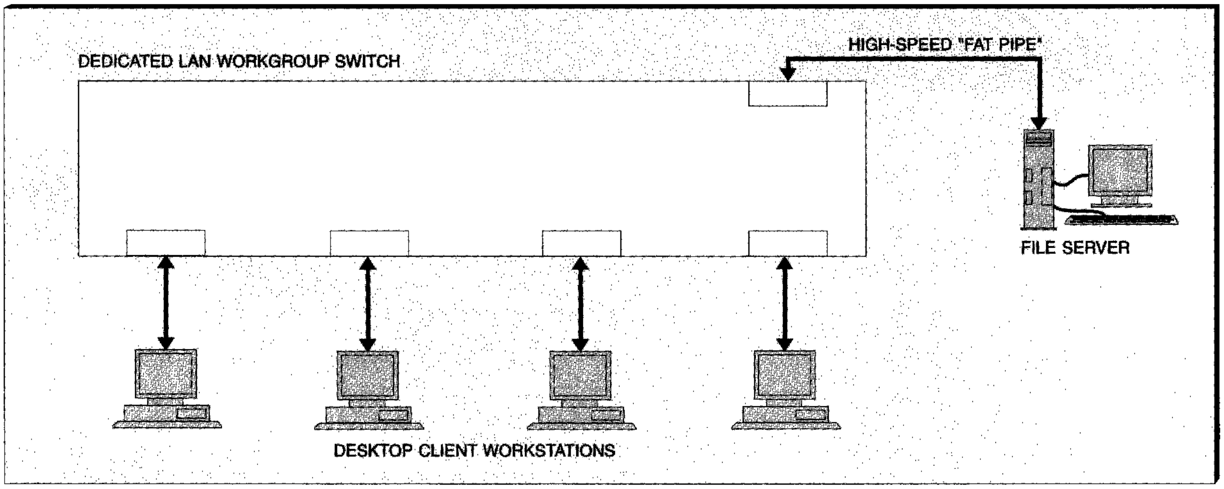

and excessive broadcast frames can cause performance problems to workstations, bridges, and routers. Thus, minimizing the impacts of broadcast traffic is one of the goals of a good network configuration.

A *virtual LAN* is a LAN segment that is defined by software and not by the underlying cable distribution hardware.[32,33] A fully virtual LAN can exist where all workstations are directly attached to a switched network. Frame-forwarding tables in the switches group stations into logical segments by the ports to which they are attached. (Stations could also be grouped based on other common attributes, such as networking protocol capability, e.g., IP subnetwork.) Only stations within a logical segment—a virtual LAN—can communicate with each other. To communicate between virtual LANs requires interconnection with a bridge or router. Interconnection of virtual LANs is no different from interconnection of physical LAN segments. In networks with multistation LAN segments attached to switch ports, a fully virtual LAN is not possible. All of the stations on a physical LAN segment must belong to the same virtual LAN. Figure 18 shows a sample network with a four-port switch. The virtual LAN membership table is shown within the switch. The virtual LAN table and frame forwarding tables in each port together determine the frame forwarding characteristics of a LAN switch. Figure 18A shows two virtual LAN segments. Stations A and B are on the sample physical LAN segment, and for this example E is added to the same virtual LAN segment. Likewise, F is added to the virtual segment containing C and D.

Since the members of a virtual LAN are defined in software, a dedicated station can be physically relocated to another dedicated port and still remain in the same virtual LAN segment. This is done by manual or automatic remapping of the new port of the relocated station to the virtual LAN of the original port. This ability to relocate and still communicate is very important for nomadic users. By their very nature, nomadic users are constantly changing location. Control of port assignments in the virtual LAN table can be done via a remote network manager. This demonstrates how virtual LANs simplify user relocation. Figure 18B shows the change to Figure 18A with station B relocated to port ID 4 and station F removed. The updated virtual LAN table is shown within the switch.

With virtual LANs it is thus possible to quickly build workgroups, independent of physical location, with shared resources and greater security. The common virtual LAN segment enables access to the same resources and limits traffic to only the designated members of the workgroup (i.e., the members of the same virtual LAN segment). Most existing and proposed virtual LAN implementations are limited to the same site; crossing a WAN boundary will be possible in the future, but is not possible today.

Figure 18 **Example of a virtual LAN with two segments**

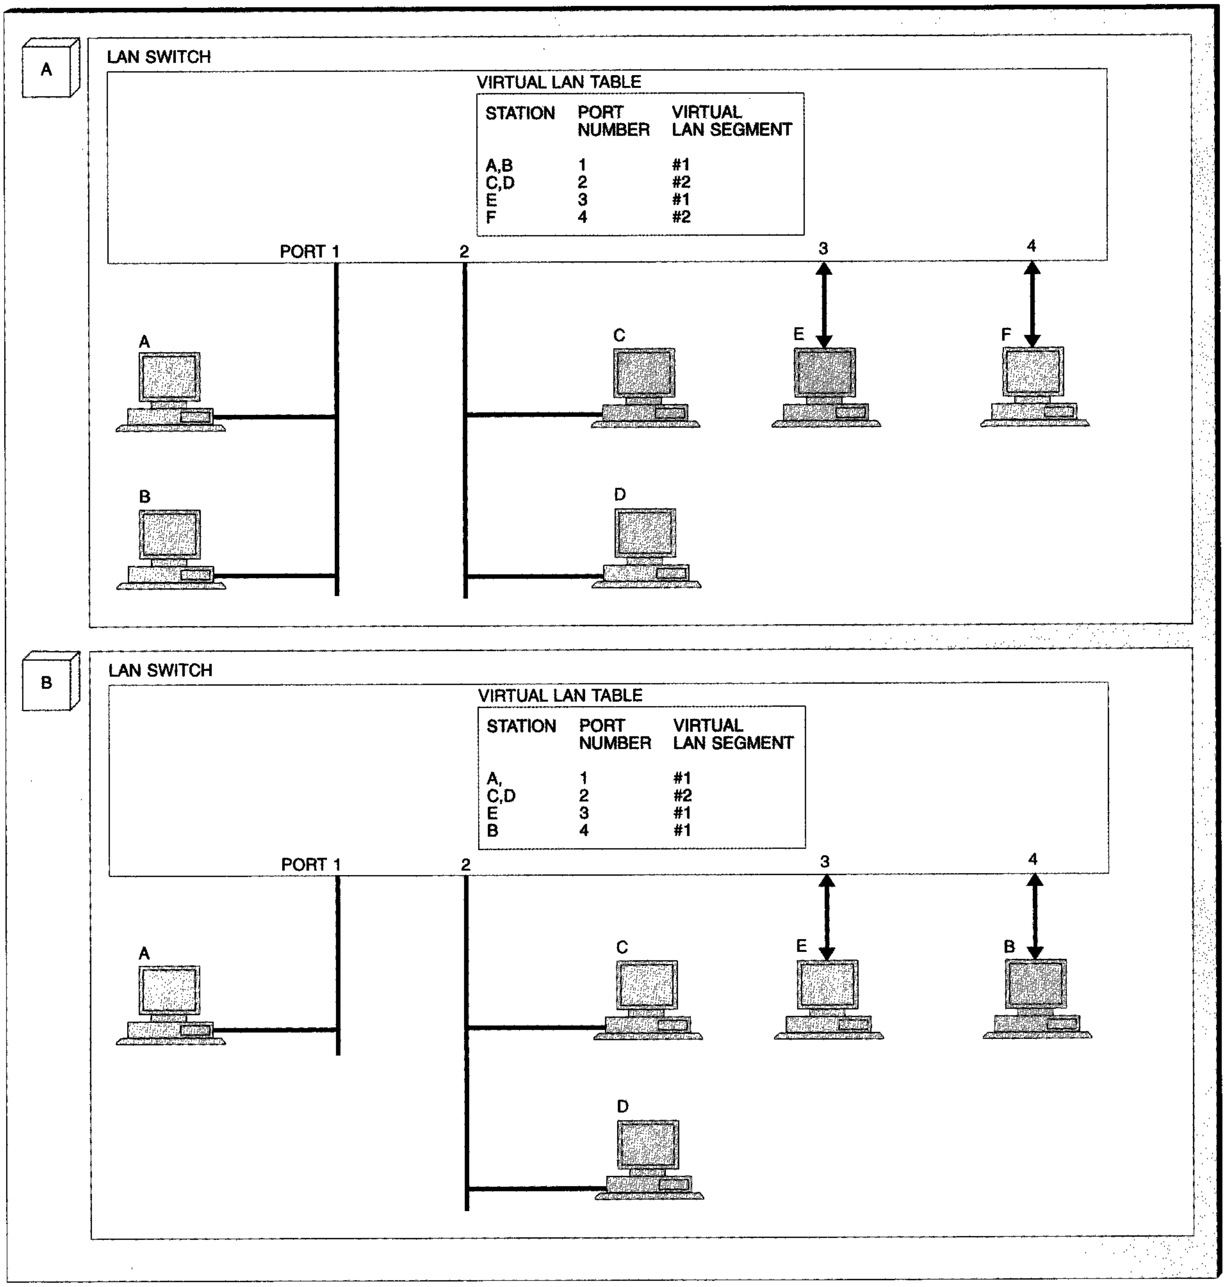

A: VIRTUAL LAN TABLE

STATION	PORT NUMBER	VIRTUAL LAN SEGMENT
A,B	1	#1
C,D	2	#2
E	3	#1
F	4	#2

B: VIRTUAL LAN TABLE

STATION	PORT NUMBER	VIRTUAL LAN SEGMENT
A,	1	#1
C,D	2	#2
E	3	#1
B	4	#1

A more advanced form of virtual LANs could allow stations to be members of more than one virtual LAN. For example, a file server could belong to several virtual LANs. Support for overlapped virtual LANs would require bridging or router functions (between virtual LANs) within the LAN switch. Virtual LANs are one of the precursors to asynchronous transfer mode (ATM) networks. With ATM and

Figure 19 Segmentation and reassembly of LAN frames

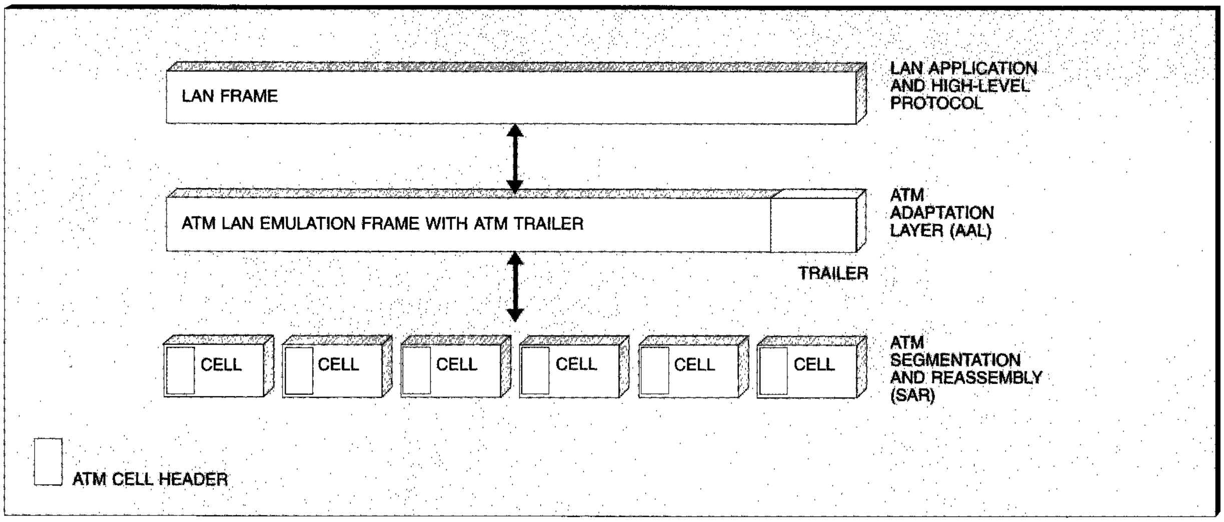

the ATM LAN emulation specification being defined in the ATM Forum,[34] it will be possible to define virtual LANs with multiple memberships and across WAN boundaries. ATM is described later in this paper.

Upgrading the existing LAN infrastructure. Switching preserves installed cabling and adapters. When upgrading a system, often what is *not* changed is of paramount importance to both the network provider and the end user. The former is interested in minimizing labor cost and the latter in minimizing the disruption to the workstation hardware and software. When moving to a switched architecture, there is only one mandatory change to the network: the installation of the switch in the wiring closet that services the desired workstations. The interesting thing to note is that this is all that must be changed. With this single change, and no associated change to the workstation hardware or software or associated cabling system, the user will be tied to a network access point (the switch) via a dedicated 16-Mbps token ring or 10-Mbps Ethernet.

Seamless LAN and WAN with ATM

The emergence of switched LANs, with dedicated bandwidth to the desktop, will complement the growth of ATM as an establishment backbone. The network user has become accustomed to the seamless LAN interconnectivity that is now possible with legacy LANs. Likewise, the user will expect this seamless operation to continue when switched LANs are integrated with ATM backbones, as network-attached workstations migrate to full ATM connectivity in the future. Multivendor interoperability will be governed by guidelines set forth by the ATM Forum. The IBM 8260 Intelligent Switching Hub[35] is an example of a high-function LAN hub that includes ATM switching functions applicable to ATM backbones.

ATM is based on high-speed switching of fixed-length 53-byte packets called *cells*. An ATM cell contains 5 bytes of header and 48 bytes of information. The small size of the ATM cell is intended to reduce delay jitter for real-time applications and simplify switch designs. As users migrate from legacy-LAN-based systems to local ATM networks, many of the applications and application interfaces will remain "frame based." This means that the LAN frames, whether in IEEE 802.3 or IEEE 802.5 format, will be subdivided into 53-byte ATM cells for transport through the ATM network (see Figure 19). The ATM adaptation layer (AAL), along with the segmentation and reassembly (SAR) functions, provide this service.

The ATM model for interconnectivity is different from that of today's frame-based LAN networks. With frame-based LAN networks, intermediate nodes can simply examine frame header contents, the 6-byte destination address or the routing information field, to determine the appropriate handling of the frame. Source routed LAN protocols are connection-based, in that the intermediate route is established at the beginning of the connection and is carried in every frame via the routing information field. Nonsource-routing LAN switches are connectionless in that the LAN switches transparently forward frames to the appropriate output links. The ATM model for cell forwarding is a connection-based protocol. However, unlike LAN source routing, an ATM cell does not carry the full routing information. Rather, the cell header contains a virtual circuit identifier and virtual path identifier (VCI/VPI) field that has meaning to the intermediate switch node that is receiving the cell. This information can be used in conjunction with stored information within the cell switch to forward the cell on the appropriate output link. Also, the switch may modify the contents of the VPI/VCI field within the cell header to facilitate the handling of the cell in the next switch node. The VPI/VCI information, along with the intermediate switch table entries, is established at call or connection setup time. ATM network standards describe a call setup procedure, called *signaling*, that is executed by the workstation prior to the establishment of each end-to-end connection. Thus, ATM cells are generally routed along a fixed, predetermined physical path within the network. The ATM cell routing model is the same for both wide area and local area network configurations. See Reference 36 for a description of ATM cell formats and virtual connections in ATM networks.

In combining frame-based switched-LAN networks and ATM in the same network, the salient aspects of both must be preserved. These include:

- The nodes at the boundaries of the frame-based and cell-based systems must provide the appropriate domain mapping, transparent to the LAN MAC level. This means that the LAN stations transmit frames with the destination MAC address of the target station. The target station may be a LAN-attached or ATM-attached station.
- The ATM transport is based on the VPI/VCI mapping that is established at the beginning of the connection. This may require that the LAN station initiate a call setup with the appropriate LAN/ATM server node for connection establishment.
- The WAN transport for this model can be existing transports such as ISDN or frame relay attached by a router to the local ATM network. It is, however, likely that the WAN transport for this model will evolve to become ATM. With ATM as the WAN transport, the LAN-to-LAN link will be entirely transparent to the end users.

Users can immediately begin to migrate to full ATM attachment as new workstations are added to the existing network under this model. Several vendors, including IBM, now offer ATM adapters and software that are designed especially for workstations and personal computers. IBM was one of the first to market low-cost (less than $400 per adapter) ATM adapters with its well-received 25-Mbps ATM technology. The 25-Mbps ATM links provide more than ample bandwidth to meet the requirements of the data, voice, and video applications that are needed, or will be needed, by desktop workstations. In addition, 25-Mbps technology supports operation on existing UTP category-3 cabling, making it lower in cost than higher speed alternatives.

There are several choices for network access speeds when selecting ATM adapters and interface hardware. This scaling of speed that is possible with ATM is one of the big benefits that the user can leverage in expanding an enterprise network. The well-informed user will only purchase the bandwidth that is necessary for a given requirement, rather than purchase excess bandwidth. For example, a client station that typically uses less than 5 Mbps aggregate bandwidth does not require a 100 Mbps ATM adapter. However, a server or mainframe that is accessed simultaneously by a large number of users could readily justify a 100 or 155 Mbps ATM attachment interface. Wide area ATM links allow variable bandwidth channels to be allocated to specific links, thus permitting the user to pay only for the bandwidth that is actually required. The IBM Transport Network Node[37] provides for wide-area connectivity with sophisticated management capabilities to optimize usage of tariffed wide-area bandwidth.

Direct ATM attachment at the workstation will enable many ATM services to be available end-to-end. One of the major advantages is the quality-of-service (QOS) bandwidth guarantees that can be provided on ATM systems. This level of service is more critical to networked real-time applications, such as interactive videoconferences. Direct ATM attachment distributes the frame-to-cell conversion

Figure 20 ATM and LAN network with LAN emulation

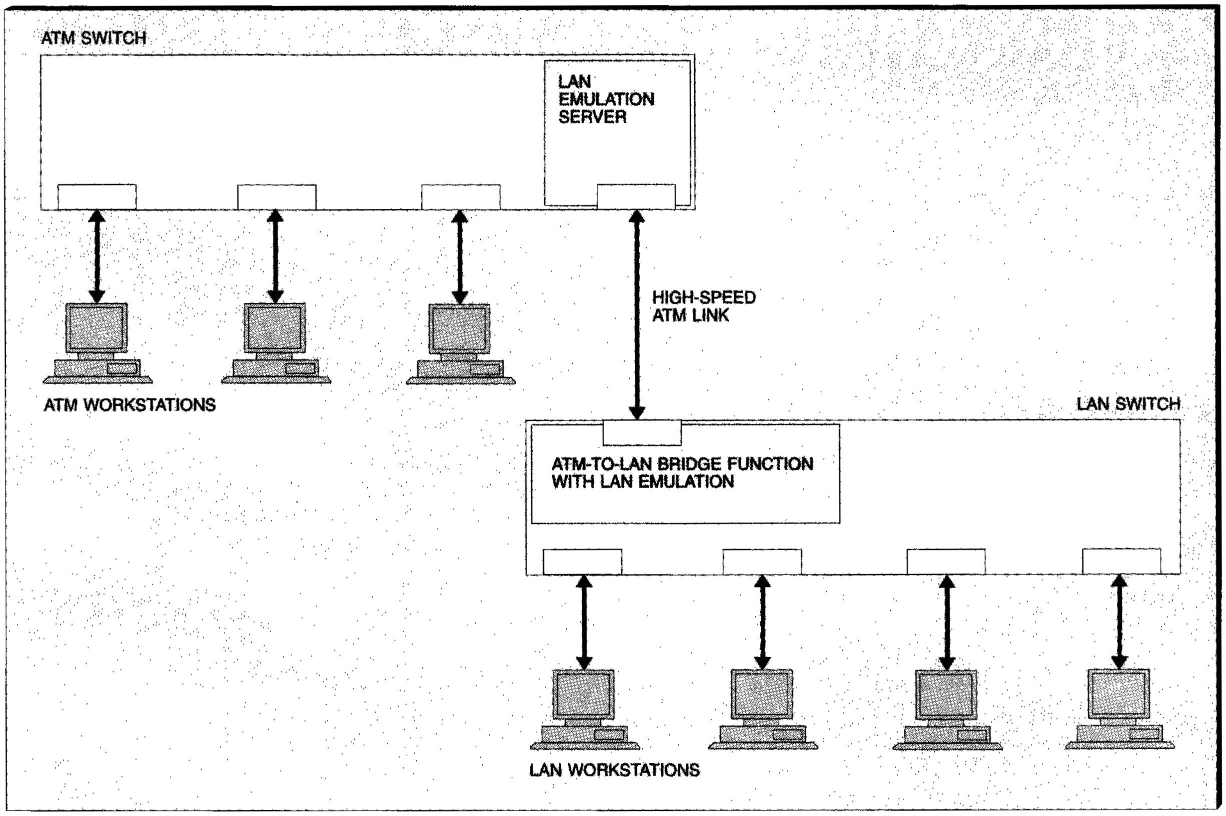

functions to the end stations, thus reducing the complexity of network interconnect nodes.

LAN emulation servers allow frame-oriented applications to operate within an ATM infrastructure. The servers provide the VCI/VPI information that is necessary in the call setup phase to establish an end-to-end connection. Once the two parties are linked, the role of the LAN emulation server is complete. However, for an ATM station to communicate with a legacy LAN-attached station, an intermediate bridge node is required. In this case, the LAN emulation server provides the link information to enable the ATM station to communicate with the bridge, where the cells are reassembled into frames for forwarding onto the legacy LAN. Figure 20 illustrates an ATM network with LAN emulation and a LAN network.

All of the stations that share a common LAN emulation server are part of the same virtual LAN. That is, broadcast information will be forwarded from one station to all other stations within only that virtual LAN. Stations that may be within another virtual LAN are accessible through an intermediate router node. Eventually, virtual LANs may be directly linked via the exchange of VPI/VCI information among the LAN emulation servers. Thus, the concept of virtual LANs can eventually be extended to remote ATM-attached workstations interconnected via an ATM WAN. This would allow users in geographically distant locations to appear to be on the same LAN segment.

Summary

This paper has described the evolution of LANs through four generations. With each generation, the amount of function in the wiring closet has increased. In the emerging fourth generation, switching technology is added to the wiring closet and in the network infrastructure. Switching technology increases the desktop and infrastructure bandwidth and does this, in many cases, without requiring changes on the desktop. Switching preserves existing investments in LAN adapters, workstation software, and cabling infrastructure.

Desktop bandwidth requirements for the foreseeable future can easily be met with full-duplex data rates of 10, 16, or 25 Mbps to the desktop. Even emerging multimedia applications do not require larger bandwidths to the desktop, but will certainly require much larger bandwidth capacities of the network infrastructure. Network management costs for user relocation can be reduced with switching technology and the software-defined virtual LANs that it enables. With LAN emulation capabilities being defined for ATM, switched LANs migrate easily into an ATM future. ATM in the infrastructure and to the desktop enables fully scalable bandwidth and virtual LANs that cross even large geographic domains.

Acknowledgments

The authors would like to acknowledge Max Brownstein for his insight into customer requirements, his ability to integrate technology trends, and his contributions to IBM's LAN-switching development strategy.

*Trademark or registered trademark of International Business Machines Corporation.

Cited references and notes

1. *YankeeWatch White Paper* **9**, No. 5, p. 1, The Yankee Group, Boston, MA (July 1994).
2. In this paper, transmission rate is referred to as "bandwidth."
3. S. Salamone, "Ethernet Switching Hubs: Sorting Through the Offerings," *Business Communications Review* **24**, No. 7, 38–41 (July 1994).
4. R. M. Metcalf and R. R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," *Communications of the ACM* **17**, No. 7, 395–404 (July 1976).
5. N. Strole, "A Local Communications Network Based on Interconnected Token-Access Rings: A Tutorial," *IBM Journal of Research and Development* **27**, No. 5, 481–496 (September 1983).
6. R-Loop was very similar to the subsequent IBM S-Loop, which was tailored to the needs of retail point-of-sale terminals and to IBM B-Loop, which was tailored to the needs of banking teller stations.
7. J. Schoch, D. Yogen, D. Redell, and R. Crane, "Evolution of the Ethernet Local Computer Network," *IEEE Computer* **15**, No. 8, 10–26 (August 1982).
8. G. Held, *Ethernet Networks*, John Wiley & Sons, Inc., NY (1994).
9. IEEE Standard 802.3, "Local Area Networks: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications," 1985.
10. P. Fortier, *Handbook of LAN Technology, Second Edition*, McGraw-Hill Series on Computer Communications, McGraw-Hill, Inc., NY (1993).
11. P. Abramson and F. Noel, "Matching the Media to Local Network Requirements," *Data Communications* **12**, No. 7, 115–123 (July 1983).
12. M. Rose, *The Simple Book—An Introduction to Management of TCP/IP-based Internets*, Prentice Hall, Englewood Cliffs, NJ (1991).
13. D. E. Comer, *Internetworking with TCP/IP, Volume 1: Principles, Protocols, and Architecture, Second Edition*, Prentice Hall, Englewood Cliffs, NJ (1991).
14. R. Perlman, *Interconnections—Bridges and Routers*, Addison-Wesley Professional Computing Series, Reading, MA (1992).
15. F. E. Ross, "An Overview of FDDI: The Fiber Distributed Data Interface," *IEEE Journal on Selected Areas of Communication* **7**, No. 7, 1043–1051 (Sept. 1989).
16. J. Russell, "Multimedia Networking Performance Requirements," *Asynchronous Transfer Mode Networks*, Y. Viniotis and R. Onvural, Editors, Plenum Press, NY (1993).
17. K. Amer, K. Christensen, and T. Toher, "Experiments with Client/Server Multimedia on 16-Mbps Token-Ring," *Proceedings of the 18th Conference on Local Computer Networks* (September 1922), pp. 2–6.
18. L. Chiariglione, "The Development of an Integrated Audiovisual Coding Standard: MPEG," *Proceedings of the IEEE* (IEEE, 345 E. 47th, New York, NY 10017) **83**, No. 2 (February 1995), pp. 151–157.
19. B. Cole, "Interactive Multimedia—The Technology Framework," *IEEE Spectrum* **30**, No. 3, 32–39 (March 1993).
20. E. A. Lee and D. G. Messerschmitt, *Digital Communications*, Kluwer Academic Publishers, Boston, MA (1988).
21. J. McQuillan, "The End of Shared-Medium LANs," *Business Communications Review* **22**, No. 10, 10–12 (October 1992).
22. J. McQuillan, "The Wheel of Reincarnation," *Business Communications Review* **21**, No. 5, 14–17 (May 1991).
23. J. McQuillan, "The Present and Future of LAN-in-a-Box," *Business Communications Review* **21**, No. 8, 12–14 (August 1991).
24. M. J. Karol, M. G. Hluchyj, and S. P. Morgan, "Input versus Output Queueing on a Space-Division Packet Switch," *IEEE Transactions on Communications* **35**, No. 12, 1347–1356 (December 1987).
25. F. A. Tobagi, "Packet Switch Architectures for Broadband Integrated Service Digital Networks," *Proceedings of the IEEE* (IEEE, 345 E. 47th, New York, NY 10017) **78**, No. 1 (January 1990), pp. 133–167.
26. P. Kermani and L. Kleinrock, "A Tradeoff Study of Switching Systems in Computer Communications Networks," *IEEE Transactions on Computers* **29**, No. 12, 1052–1059 (December 1980).
27. D. Cullerot, "Full Duplex Switched Ethernet (FDSE)," *IEEE 802.3 Plenary*, IEEE, 345 E. 47th, New York, NY 10017 (November 1992).
28. A. Guruge, "Token-Ring—The Next Generation," *Business Communications Review* **24**, No. 9, 48–52 (September 1994).
29. N. Strole, "Full-Duplex Option for 16-Mbps Token-Ring Stations," *IEEE 802.5 High Performance Architecture Study Group*, Marlborough, MA (September 1993).

30. Numerous performance studies, including those by Bux (see Reference 31), have demonstrated that the portion of the shared bandwidth available for data transmission is dependent upon the access protocol and other network parameters. The examples shown here assume ideal operation with near-optimum bandwidth efficiency.
31. W. Bux, "Local-Area Subnetworks: A Performance Comparison," *IEEE Transactions on Communications* **29**, No. 10, 1465–1473 (October 1981).
32. R. Strauss, "Plan Your 2001 Network Now," *Datamation* **40**, No. 15, 32–37 (August 1, 1994).
33. R. Strauss, "Virtual LANs Pave the Way to ATM," *Datamation* **40**, No. 16, 20–24 (August 15, 1993).
34. R. Jeffries, "ATM to the Desktop: Prospects and Probabilities," *Telecommunications* **28**, No. 4, 25–32 (April 1994).
35. R. A. Sultan and C. Basso, "ATM: Paving the Information Superhighway," *IBM Systems Journal* **34**, No. 3, 375–389 (1995, this issue).
36. J. Le Boudec, "The Asynchronous Transfer Mode: A Tutorial," *Computer Networks and ISDN Systems* **24**, No. 4, 279–309 (1992).

Accepted for publication March 13, 1995.

Kenneth J. Christensen *IBM Networking Hardware Division, P. O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: kenc@vnet.ibm.com).* Dr. Christensen is an advisory engineer in IBM's Networking Hardware Division, where he has been a member of the token-ring development team since its inception in the early 1980s. Dr. Christensen is currently a member of the LAN/WAN Advanced Development Department, focusing on future enhancements for token-ring and other LAN technologies. He received a B.S. degree from the University of Florida in 1981 and M.S. and Ph.D. degrees from North Carolina State University in 1983 and 1991, respectively. Dr. Christensen is a senior member of the IEEE.

Lee C. Haas *IBM Networking Hardware Division, P.O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: leehaas@vnet.ibm.com).* Mr. Haas is a senior engineer in IBM's Networking Hardware Division. His present work is in physical layer advanced development for local area networks and related standards activities of ATM, ANSI, EIA/TIA, and IEEE. He is the originator of IBM's proposal for the ATM 25.6 physical layer and was a key contributor to the EIA/TIA Commercial Building Cabling Standard. Mr. Haas joined IBM in 1967 and has worked in the areas of circuit design, signal processing, data communications, and modem product development. He is a registered professional engineer, holds B.S. and M.S. degrees in electrical engineering from Virginia Polytechnic Institute and State University, and San Jose State College, respectively. He is pursuing a Ph.D. in electrical engineering at North Carolina State University. He is national vice president of Theta Tau, a professional fraternity for engineers.

Franc E. Noel *IBM Networking Hardware Division, P. O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: franc@vnet.ibm.com).* Dr. Noel is a Senior Technical Staff Member in IBM's Networking Hardware Division and is manager of LAN/WAN advanced development. He joined IBM in Poughkeepsie, New York, in 1974 and has been working in the field of local area network communications since 1981. Dr. Noel developed the media specifications for the IBM Cabling System, for which he received an Outstanding Technical Achievement Award, and led the transmission system development effort to support token ring on unshielded twisted pair cabling. He is on the board of directors for the National Science Foundation Center for Advanced Communications and Computing. Dr. Noel received his M.S.E.E. from Syracuse University in 1979, and the Ph.D. in computer engineering from North Carolina State University in 1991. He is a senior member of the IEEE and a registered professional engineer.

Norman C. Strole *IBM Networking Hardware Division, P. O. Box 12195, Research Triangle Park, North Carolina 27709 (electronic mail: ncstrole@vnet.ibm.com).* Dr. Strole is a senior engineer with IBM in the Research Triangle Park, North Carolina, laboratory where he is currently involved in the design and performance analysis of high-speed switched LAN systems. He has been a member of the LAN development team for the past 14 years. Dr. Strole received a B.S. degree in electrical engineering from North Carolina State University in 1973. He received M.S. and Ph.D. degrees in electrical engineering from Duke University in 1975 and 1980, respectively. He currently holds an adjunct appointment in the School of Electrical Engineering at Duke University and is a senior member of the IEEE.

Reprint Order No. G321-5572.