EMERGING TRENDS - FULL-DUPLEX AND THE SWITCHED LAN

Ken Christensen, Franc Noel, and Norm Strole

IBM Corporation
Networking Hardware Division
Research Triangle Park, NC 27709

ABSTRACT

Ethernet, Token-Ring, and FDDI are well established protocols for regulating the access to a common transmission medium among a large number of stations. All of the attached stations share this common medium and its bandwidth. Sharing of a common medium reduces the bandwidth that is available to an individual station. To increase bandwidth to the end user, emerging LAN topologies are departing from shared-media, shared-bandwidth methods in favor of dedicated-media and dedicated-bandwidth methods. Dedicated-bandwidth switched LANs can take advantage of full-duplex operation of attached stations. This is in contrast to the normal half-duplex operation of LAN stations on a shared-bandwidth LAN. This paper describes the evolution of shared-media, shared-bandwidth LANs into dedicated-media, dedicated-bandwidth switched LANs with full-duplex operation of the attached stations.

1.0 INTRODUCTION

Early LANs were typically small - both in network span and in the number of stations. In these early LANs the attached stations were much less powerful than today′s Personal Computers and workstations. As a result, there were few concerns over bandwidth requirements. These early LANs were typified by Ethernet 10Base5 and 10Base2 and can be categorized as having shared media and shared bandwidth. With bandwidth being shared on a common medium, the LAN adapters were half-duplex in operation. Figure 1 shows an Ethernet 10Base5 LAN segment where multiple stations share a common transmission medium and its bandwidth.

During the 1980′s, LAN data rates increased from 2 Mbps for IBM′s PC-Net LAN up to 100 Mbps for FDDI. The most popular LANs during this time were 10-Mbps Ethernet and 4- and 16-Mbps Token-Ring. Cable types evolved from coax (for example, Ethernet 10Base5 and 3270 display systems) towards shielded and then unshielded twisted-pair cables. Optical fiber cabling was also installed in some locations to allow for future very high-speed technologies. As the strategic importance of LANs grew they became larger. This increased importance, along with growing station count and increased geographic coverage, encouraged the growth of systematic cabling schemes. The IBM Cabling System and the EIA/TIA-568 Commercial Building Telecommunications Cabling Standard (see [1]) describe systematic star-wired cabling systems for multiproduct and multivendor environments. Ethernet 10BaseT and Token-Ring are examples of LANs that often employ a wiring closet for concentrating each star-wired LAN segment. Star-wiring simplifies problem determination on a LAN. A wiring closet provides one place where all station attachments can be accessed. Thus, a defective station or cable, once identified, can easily be removed and repaired without the need to literally ″walk the wire.″ Figure 2 (on the next page) shows a star-wired Token-Ring LAN segment with a central wiring closet.

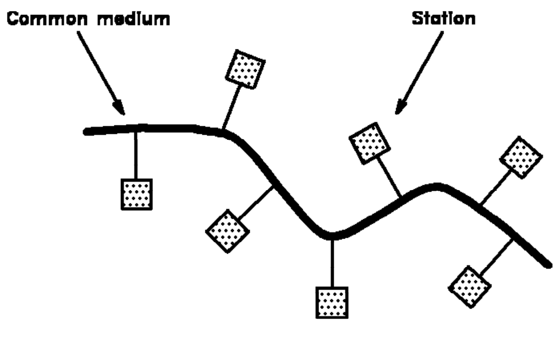

Figure 1 - Shared-media, shared-bandwidth Ethernet

In this paper, section 2 describes the evolution from shared media to dedicated media LANs. Section 3 describes the next step - the evolution to dedicated bandwidth LANs. Section 4 describes full-duplex operation of a LAN adapter and reviews proposals made to IEEE 802.5 for full-duplex operation of Token-Ring stations.

Section 5 describes performance characteristics of full-duplex dedicated-bandwidth LANs, discusses motivating applications, and briefly discusses migration and management. Finally, section 6 is a summary.

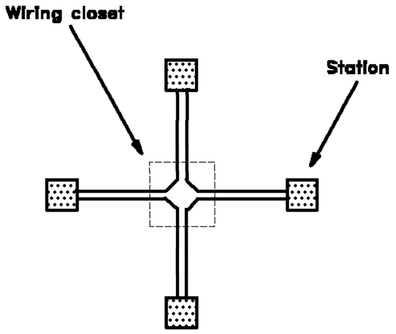

Figure 2 - A star-wired Token-Ring LAN segment

2.0 SHARED MEDIA TO DEDICATED MEDIA

Initially, wiring closets contained little intelligence - simple repeaters for Ethernet or unpowered concentrators containing insertion relays for Token-Ring. On these shared-media and shared-bandwidth LANs, any cable or station failure would likely affect all stations on the same LAN segment. For example, forgetting to install terminators on both ends of an Ethernet 10Base2 segment would cause the entire segment to be inoperable.

With the continued growth and business importance of LANs and the decreasing cost of electronics, wiring closets started to become ″intelligent″. Intelligent repeaters for Ethernet and intelligent concentrators for Token-Ring are capable of automatically isolating and removing failed stations. For example, an intelligent Token-Ring concentrator can prevent a 4-Mbps station from inserting into a 16-Mbps LAN segment. This minimizes outages and increases the reliability of a LAN. LANs built with intelligent wiring closets can be classified as dedicated media and shared bandwidth. Bandwidth is still shared among all users, however each station′s media is ″dedicated.″ Because of the star-wired cabling topology, station failures are most often isolated to a single station and its media. Intelligent repeaters and concentrators can also contain additional features such as remote network monitoring and reporting mechanisms, connection between different media types, data security facilities, and LAN bridging. Figure 3 shows a five-port Token-Ring intelligent concentrator. The active logic shown in the figure automatically removes failing stations from the main ring so, for example, a station attempting to insert at a data rate different from the main ring data rate will not be permitted to join the network.

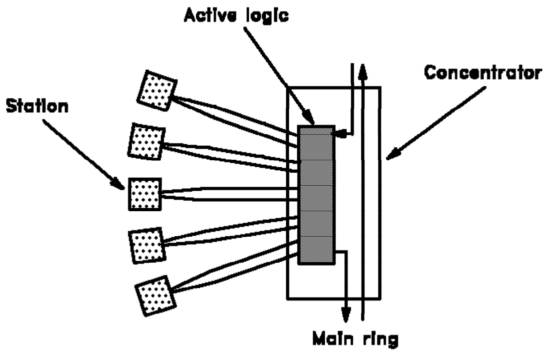

Figure 3 - A Token-Ring intelligent concentrator

3.0 THE NEXT STEP - DEDICATED BANDWIDTH

Shared bandwidth with dedicated media increases the manageability and reliability of LAN segments, but does not increase the amount of bandwidth available per station. In a shared-bandwidth LAN all transmitted frames pass by all stations, even those frames not destined for a particular station. LAN bridges or port-switching hubs can be used to partition large LAN segments into smaller interconnected segments. A port-switching hub contains multiple backplanes, each corresponding to a LAN segment. Compared to a single large LAN segment, average station bandwidth may be increased by localizing stations that primarily communicate between each other on the same (smaller) LAN segment.

By using parallel switching techniques in an intelligent concentrator, the amount of bandwidth per station is increased. For example, a cross-bar switch can allow dedicated frame-by-frame connections between switch ports. When a frame is received at a port on a switch, it is automatically switched ″on the fly″ to the appropriate port for output based on the destination address contained in the frame. A table at each switch port determines to which port a given destination address should be switched. Frames addressed to the all-stations address (that is, broadcast frames) are sent out on all ports. Thus, each LAN segment attached to a switch port receives only the frames with destination addresses on that LAN segment. A switch port can also contain frame buffering to temporarily ″hold″ frames that cannot be immediately forwarded. For example, in an Ethernet LAN if the destination Ethernet segment is currently busy, a switched frame would be temporarily stored in the switch port buffers. Figure 4 shows a four-port cross-bar LAN switch. Each switch port contains a Media Access Control (MAC) unit to support attachment to a LAN segment. Attached to each switch port can be a dedicated-media single station or shared-media LAN segment with multiple stations.

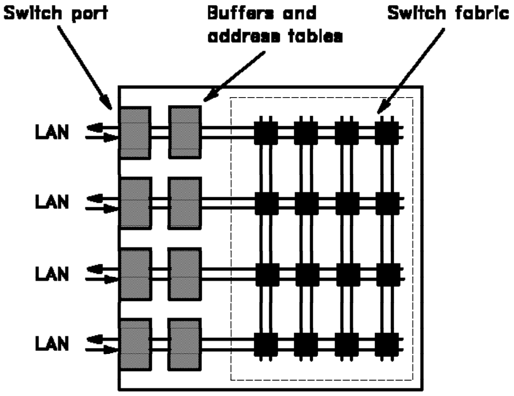

Figure 4 - A four-port cross-bar switch

Detailed descriptions of LAN switch architectures are beyond the intended scope of this paper. Manufacturer's sales literature from references [4] and [6] can be consulted. In [3] the evolution of 10-Mbps "private Ethernet", then 100-Mbps dedicated networking, and finally 155-Mbps Asynchronous Transfer Mode (ATM) is described.

4.0 FULL-DUPLEX OPERATION

All LAN adapters support a MAC protocol. The purpose of the protocol is to regulate access to a common media (for example, via tokens as in Token-Ring). In a switched LAN with a single link per station, the requirements for a MAC protocol are greatly minimized. No longer must a station "contend" with other stations for permission to transmit onto the media. Thus, in a switched LAN, stations can transmit whenever a frame is queued in the adapter. A station can also receive at any time. Full-duplex operation is now possible - a station can transmit and receive independently of other stations on the switched LAN.

Full-duplex operation has been proposed for Ethernet in IEEE 802.3 (see [2]) and for Token-Ring in IEEE 802.5 (see [5]). Full-duplex offers the following benefits over half-duplex operation:

- Increases the bandwidth of switch-attached single station LAN segments.
- Provides migration to higher bandwidths as existing and emerging applications require additional bandwidth.
- Capitalizes on emerging LAN switching technology.
- Positions the desktop Personal Computer or workstation for the future high-speed enterprise Asynchronous Transfer Mode (ATM) backbone.
- Preserves the investment in existing LAN adapters and infrastructure (for example, cabling, software applications, and network management platforms and protocols).

4.1 Full-Duplex Token-Ring

The Token-Ring MAC protocol is usually executed within a Token-Ring adapter. When operating in standard Token-Ring mode, an adapter normally transmits a frame or receives a frame, but not both simultaneously. Figure 5 shows the normal mode of operation of a Token-Ring adapter. The figure shows the repeat path for all frames, whether received or not, within an adapter. The repeat path is required in half-duplex Token-Ring operation to allow each transmitted frame to circulate around the entire ring.

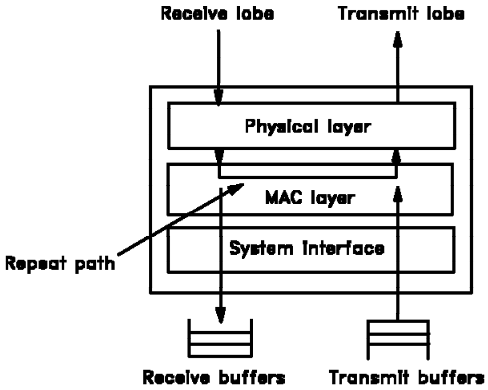

Figure 5 - Normal operation of a Token-Ring adapter

If the repeat path is removed and other changes made that are discussed later, a Token-Ring adapter can transmit a frame on the outgoing link while also receiving a frame on the inbound link, see Figure 6 (on the next page). This is full-duplex operation. In this full-duplex mode, frames are transmitted on the outgoing link immediately upon being queued for transmission at the MAC interface. Depending upon the capabilities of the adapter, its station, and the switch to which it is attached, sustained transmit rates of 16 Mbps are possible. Even more significant, the adapter can also simultaneously receive frames at the 16 Mbps rate on the receive link, for an aggregate capacity of 32 Mbps. Typically, LAN applications (for example, client/server) will generate traffic in bursts rather than in continuous streams. Thus, the achievable data rate of an aggregate 32 Mbps will be required in bursts, rather than in a continuous mode of operation.

When in full-duplex mode, a station does not employ token protocols, including:

- Token access and token recovery protocols
- Frame repeating, frame stripping, and end-of-frame status update

A key aspect of full-duplex operation is that the *Token-Ring frame format is preserved without change.* The Frame Control field, priority bits, destination and source address fields, and source routing subfields are all maintained. Thus, existing IEEE802.5 end-user applications software and network infrastructure (for example, bridges) are preserved making the full-duplex mode of operation transparent to the user.

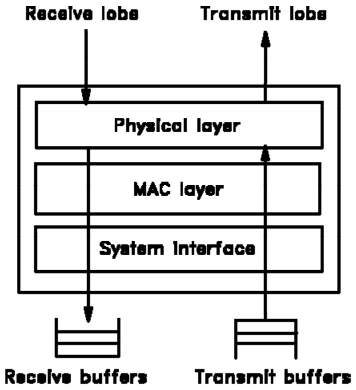

Figure 6 - Full-duplex Token-Ring adapter

A single Token-Ring adapter may support both normal (half-duplex) Token-Ring and full-duplex operation. There are two logical connectivity options that an adapter should satisfy:

1. Activate in full-duplex mode when attached to a compatible full-duplex switch port.
2. Otherwise, activate in Token-Ring mode when attached to a standard Token-Ring concentrator or switch port that does not support full-duplex operation.

A station cannot physically determine whether it is attached to a half-duplex Token-Ring concentrator or to a full-duplex switch. This determination must be accomplished via a logical protocol during the activation procedure. Figure 7 shows an example activation sequence for an adapter to a switch or concentrator port. In state S2 the normal Token-Ring insertion procedure of internal diagnostics and lobe media test is executed. State S3 advertises, for example via a MAC frame, the capability of the adapter to operate in full-duplex mode. If the port is capable of full-duplex operation, transition is made to state S6 (full-duplex operation), otherwise transition is made to state S5 (normal Token-Ring operation).

Following activation for full-duplex mode of operation, the token access protocols can be suspended, while other functions, such as those used to detect and report errors, continue to be active. When in full-duplex mode, each station or switch port must transmit using the local crystal within that station since there is no longer an active monitor providing clocking. The existing differential Manchester encoding scheme is maintained. Stations receive incoming signal via the acquired clock of the incoming link. When there are no frames to transmit, a station transmits (and expects to receive) a continuous idle stream (for example, all zeros). This idle stream provides a continuous signal to the receiving port for clock synchronization and serves as a means of detecting and/or reporting a link failure or disconnection.

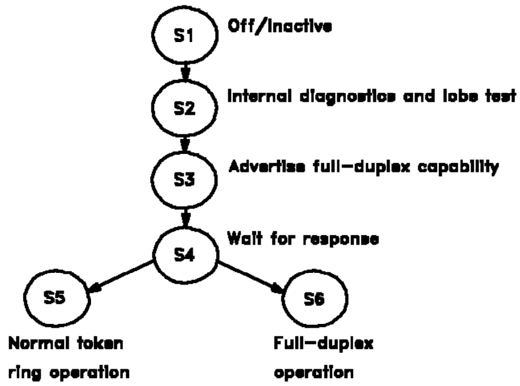

Figure 7 - Example activation sequence

All frames can be transmitted immediately upon proper setup to the MAC interface without having to wait for token access. A maximum frame length of 4500 bytes is recommended for full-duplex operation, primarily to both reduce buffer requirements in the LAN switch and end-to-end frame transit delays due to blocking or store-and-forward queueing delays. A station may maintain multiple transmit queues in order to prioritize local access to the full-duplex link. However, first-in first-out frame order within a given priority must be maintained within a station or LAN switch. The low-order three bits of the Token-Ring Frame Control field have been designated as the priority bits of a given frame. These bits may be used for priority transmit queueing. A full-duplex station will receive all frames that are properly formatted and contain the appropriate destination address.

When in full-duplex mode, the majority of the normal Token-Ring MAC level protocols are no longer required and can thus be suspended. Since many of these MAC protocols depend upon specific MAC frames, stations must not transmit these frames, nor should LAN switches forward these frames, when in full-duplex mode. The Active Monitor and Standby Monitor Present MAC frames could be used to provide a periodic "heartbeat" between a station and a switch port on a full-duplex link to indicate a logical presence to each other. When in full-duplex mode, a switch port can learn the address of its attaching adapter during the activation protocol. This is particularly applicable to mapping a station address to a physical port on an interconnect switch. Beacon MAC frames are still applicable in full-duplex mode to indicate temporary or permanent hard errors. Report Soft Error MAC frames can also still be used, but some of the error conditions that are reportable in half-duplex Token-Ring mode do not apply to full-duplex mode.

Full-duplex operation should be less sensitive to many of the disruptions that can occur in normal Token-Ring operation. For example, all errors associated with maintaining token flow are no longer a concern. Also, link errors no longer affect data on the entire ring, but only on an isolated segment. However, there can still be situations where the first indication of a problem is the appearance of soft errors in data frames. Errors in received frames may be indicative of a poor transmission link or of external electromagnetic interference. These errors will normally appear as Frame Check Sequence (FCS) errors. Improperly formatted or incomplete frames must also be handled. A frame that exceeds the maximum allowed frame size for a full-duplex link could be handled as an improperly formatted frame. In general, the receiving station is responsible for detecting frame errors. Frames that exceed the maximum frame length can be detected via the activation of a valid-frame timer (for example, 2.25 milliseconds for a maximum length 4500 byte frame at 16 Mbps) at the beginning of frame reception. If a valid End Delimiter is not detected before the timer expires, the frame reception is terminated. The detection of an Abort Delimiter by a LAN switch during the forwarding of a frame can also immediately terminate frame forwarding.

5.0 PERFORMANCE CHARACTERISTICS

This section describes the performance characteristics of a full-duplex switched LAN. In existing LANs the total bandwidth of a common medium is shared between all attached stations. For example, on a 16-Mbps Token-Ring with 50 stations, each station can expect an average bandwidth of 16 Mbps / 50 = 320 Kbps. Of course, the peak bandwidth available to a station is the full 16 Mbps assuming that no other stations are transmitting data onto the LAN. Similar per-station bandwidth degradation is typical of other LANs, such as 10-Mbps Ethernet and 100-Mbps FDDI. If a single station is attached to a switch port, then that single station has the full, or dedicated, bandwidth of the switch port. In a shared-bandwidth LAN segment, the addition of each new station decreases the average bandwidth available for every station. For a dedicated bandwidth LAN, the addition of a new station does not imply less bandwidth for every station. Figure 8 is a graph showing the average bandwidth available to stations on shared and dedicated-bandwidth LANs. This graph shows a uniform 100% load for shared (half-duplex) and full-duplex Token-Ring. This graph assumes a maximum achievable shared (half-duplex) Ethernet utilization of 80%.

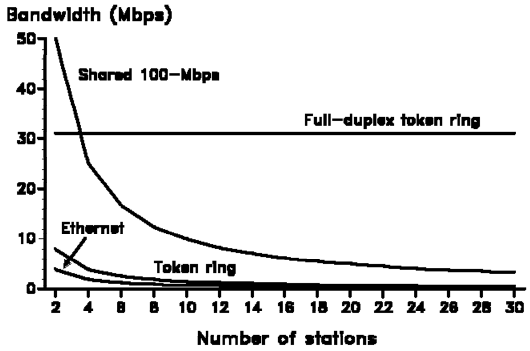

Figure 8 - Shared and dedicated bandwidth performance

Access delay is the length of time a station must wait for a queued frame to be transmitted. On a shared-bandwidth LAN the access delay varies according to how much traffic is on the LAN. For a Token-Ring LAN the access delay is the amount of time a station waits for a token. This token wait time varies according to how many other upstream stations also have data ready to transmit onto the LAN. For a Token-Ring LAN segment with very low utilization levels, token wait time is negligible. For a Token-Ring LAN segment with higher utilization levels, token wait times would rarely exceed several milliseconds. With a full-duplex link to a switch, a station has dedicated transmit bandwidth. A full-duplex link yields a full 16 Mbps for transmit and a simultaneous 16 Mbps for receive. Thus, the access delay for transmitting a frame is effectively zero since there is no token wait delay. Some frame queueing delays may, however, be experienced within the LAN switch due to the bursty nature of the traffic or the load distribution among the various stations.

Table 1 shows the bandwidth characteristics of a 16-Mbps shared-bandwidth LAN and a 16-Mbps full-duplex dedicated-bandwidth switched LAN. Table 1 also shows access delay where the value 2.25 milliseconds represents the delay caused by the transmission of a 4500 byte frame at 16 Mbps. Each LAN is assumed to have N stations attached. The minimum bandwidth is the worst case bandwidth that a single station can expect to obtain over a period of time. The peak bandwidth is the maximum instantaneous bandwidth a station can achieve. Aggregate bandwidth is the total bandwidth available in the LAN. Finally, access delay is the time a station must wait before it can transmit a frame. For maximum full-duplex performance, the aggregate bandwidth of the switch must be able to handle peak load situations. For example, a peak load situation occurs when half of the switch ports are sending frames to the other half while also receiving frames from the other half. In other words, a peak load situation occurs when all switch ports are busy in simultaneous full-duplex transmit and receive of frames.

	Token-Ring	Full-Duplex
Minimum bandwidth (Mbps)	$\frac{16}{N}$	16 Xmit and 16 Recv
Peak bandwidth (Mbps)	16	16 Xmit and 16 Recv
Aggregate bandwidth (Mbps)	16	$\frac{1}{2} \times (32 \times N)$
Access delay (milliseconds)	0 to (N x 2.25)	0

Table 1 - Performance of normal and full-duplex

5.1 Applications Motivating Switched LANs

Two applications that can benefit from a full-duplex switched LAN are client/server computing and multimedia. In client/server computing a powerful dedicated server is used to store common data and programs accessed by multiple client stations. Multimedia encompasses a large number of applications that typically combine text, images, audio, and video. Multimedia applications can operate as either client/server or client/client. Delivery on demand of video clips is an example of a client/server multimedia application. Video conferencing is an example of a client/client multimedia application.

For client/server computing the server can be a performance bottleneck. That is, multiple clients accessing a single server can overwhelm the processing capabilities or network bandwidth of the server. If network bandwidth is the bottleneck, then putting the server on a dedicated full-duplex link can improve overall performance. Figure 9 shows a switched LAN with two full-duplex file servers. The client stations are on shared-bandwidth LAN ″micro-segments″ attached to the same switch. A micro-segment is a LAN segment with very few stations (for example, five to six stations) as compared to a typical LAN segment of thirty to fifty stations. For this configuration, the two file servers have up to four times the available bandwidth than if they and all the client stations were attached to a single shared-bandwidth LAN. With full-duplex and dedicated bandwidth it is possible to improve performance for client stations by changing only the location of the file server(s) via direct attachment to a LAN switch. On a shared-bandwidth LAN, to improve the data throughput of a single station (for example, a file server), all stations have to be upgraded, at a high cost, to a faster LAN adapter. Cabling infrastructure may also have to be changed. An example of this type of upgrade would be changing from a 10-Mbps Ethernet LAN to a 100-Mbps FDDI LAN.

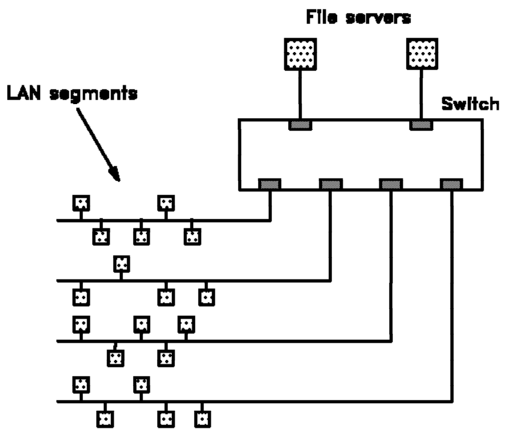

Figure 9 - Switched LAN with full-duplex file servers

Multimedia applications that contain real-time audio and video require guaranteed bandwidth and, in some cases, non-varying access delay. Data traffic does not, in general, require guaranteed bandwidth. Response times for file transfers, for example, simply take longer when less bandwidth is available. However, without guaranteed bandwidth, real-time traffic such as video will suffer in quality. This reduced quality will be observed by a user as pauses in both video and audio. Typically, full-motion video requires several megabits per second of bandwidth for each video stream. If many users on a LAN desire to view independent video streams, a large amount of aggregate bandwidth is needed. Full-duplex and switching can provide the needed bandwidth to enable many emerging multimedia applications on a LAN.

5.2 Migration and Management of Switched LANs

The evolution to LAN switching allows existing LAN adapters and infrastructure to be migrated towards higher bandwidths. The first step in this migration is to isolate the few high-bandwidth stations (for example, file servers) to dedicated full-duplex links. Stations with lower bandwidth requirements can remain on larger LAN segments. As client bandwidth demands increase, the first step towards dedicated-bandwidth can be micro-segmentation. In this case, each LAN switch port has a LAN segment consisting of more than one station attached to it. The final step is individually attaching client stations directly to a LAN switch.

By keeping existing LAN adapters, the existing network infrastructure can also remain largely intact. This includes existing cabling layouts, network management platforms and protocols, and software applications. This then lays the groundwork for attaching to high-speed backbones such as ATM.

6.0 SUMMARY

This paper has described the evolution of small shared-media LANs into large LANs with dedicated media. This first evolutionary step had significant benefits in reliability. The second and emerging step is towards dedicated media and dedicated bandwidth. Migration to dedicated-bandwidth switched LANs is motivated by demands for increased bandwidth and improved management. Switched LANs enable full-duplex operation of LAN stations. With switched LANs the existing install base of LAN adapters and infrastructure can be incrementally migrated (from shared-bandwidth to dedicated-bandwidth) for increased client/server traffic demands and many future multimedia applications.

Acknowledgments

The authors would like to acknowledge Mike Siegel, Tom Toher, Ray Zeisz (all of Networking Hardware Division, IBM Corporation), and many others for their contributions to full-duplex Token-Ring and this paper.

REFERENCES

[1] Commercial Building Telecommunications Wiring Standard, *EIA/TIA 568*, Electronics Industries Association, July 1991.

[2] Cullerot, D. ″Full Duplex Switched Ethernet (FDSE).″ *IEEE 802.3 Plenary*, November 1992.

[3] McQuillan, J. ″The End of Shared-Medium LANs.″ *Business Communications Review*, October 1992.

[4] ″MultiNet - The Future is Here (Video-Ready Switching LAN Hubs).″ Sales literature from LANNET Data Communications, Inc., Huntington Beach, California, 1993.

[5] Strole, N. ″Full-Duplex Option for 16-Mbps Token-Ring Stations.″ *IEEE 802.5 High Performance Architecture Study Group*, September 1993.

[6] ″The Full Duplex Switching Advantage.″ Sales literature from Kalpana, Sunnyvale, California, 1993.