

# Multimedia Applications and 100VG-AnyLAN

Networks must guarantee bandwidth for multimedia traffic and must control end-to-end delay and delay jitter (fluctuation in the arrival time of packets). The new campus network, 100VG-AnyLAN, can meet these requirements in many circumstances through the basic operation of the protocol. More flexibility can be obtained through the use of bandwidth allocators and the target transmission time protocol. Until either the Broadband Integrated Services Digital Network (B-ISDN) or reliable Internet protocols become available, the use of dial-up remote bridges with existing WANs can accommodate multimedia traffic in the near term.

by **John R. Grinham** and **Michael P. Spratt**

The term *multimedia* covers a wide range of applications, ranging from simple image-in-text word processors through video conferencing. The qualifier "real-time" confines the applications to those that impose timing constraints on the underlying delivery system, such as:

- Conferencing systems using synchronized video, audio, and shared computer data
- Remote training and information databases containing digital video and audio data
- Remote control or manipulation of computer-generated models and robotic machines.

As desktop computers grow in power and decrease in price, and networks expand both in area and in data transfer rates, such applications are likely to have a major impact on all desktop computer systems and workgroup LANs in the mid to late 1990s. The factors that indicate a takeoff in multimedia can be divided into two groups: user demand and the development of appropriate technology.

## User Demand for Multimedia Applications

A major demand is for humans to interact with remote systems. This started in the fields of education and training. Computer-based training, which used to use analog video and audio, has become a mature market, and is moving rapidly towards using digital media. In the future, this data may need to be held centrally so that it can be updated easily. Training applications could be networked to this central source. This requires that control data be transmitted to a server with low delay, and that the server have the ability to transmit multimedia data back to the desktop, also with low delay.

Applications are growing in the scientific, engineering, and medical domains in which, for example, humans interact in real time with remote models and image data. These applications also place significant timing constraints on the underlying network.

Perhaps the most exciting growth area is in the support of real-time human conversation. Teams separated geographically need such systems for activities traditionally performed on a single site, such as meetings, training, marketing, and support. As an example, HP has already released a set of collaborative workgroup tools that includes imaging, digital audio, and shared data.

Systems such as these are a natural development from video conferencing systems. It is remarkable that over half of all Fortune 1000 companies use boardroom-based conferencing facilities, given that until recently they cost over \$20,000 per node and were relatively crude in operation. The possibility of low-cost desktop-based conferencing systems that integrate with existing computer applications has to be seen as extremely attractive to companies and teams distributed over several sites.

## Multimedia Technology

Demand for desktop multimedia applications connected over networks can only be satisfied if the technology exists to support it. Three areas are critical:

- The development of standards for video and audio coding that allow low-cost implementation and fast transmission
- The development of low-cost wide-area networks (WANs)
- The development of consistent interfaces to multimedia technology, both from WANs to appropriate LANs, such as 100VG-AnyLAN, and from LANs to applications.

**Video and Audio Coding Standards.** The ISO and the ITU-T (formerly CCITT) have produced several standards for audio and video coding algorithms, such as MPEG and H.261 for video conferencing and video telephony and JPEG for the compression of still images (now being extended to full-motion video). Since all these algorithms are based on discrete cosine transformations, low-cost chipsets that support all three are being developed by several manufacturers.




Fig. 1. Developments in different portions of a wide area network.

Audio CD-quality systems are already becoming standard features of workstations and PCs, and many manufacturers have released plans to incorporate hardware support for full-motion video. The increasing power of workstation and PC microprocessors is also making software compression and decompression of full-motion video feasible. HP's current workstations can already perform decompression of full-motion video. Real-time compression is likely to follow in a few years.

**Wide Area Networks (WANs).** The cost of WAN services has fallen dramatically in the last few years. For example, a long-distance video conferencing link that was around U.S.\$1500 per hour in 1984 now costs in the range of U.S. \$12 to \$30 per hour. In addition, there is increasingly widespread availability of a variety of dial-up WAN services.<sup>1</sup>

There is also a significant effort going into the development of wide area broadband services, perhaps leading to a Broadband Integrated Services Digital Network (B-ISDN) by the end of the decade.

**Consistent Interfaces.** For multimedia to work effectively over wide area networks and a variety of campus and local-area networks, there is a need for consistent interfaces between networks. 100VG-AnyLAN only covers a portion of a wide area networked system.

Various bodies are looking at other areas of network systems. Fig. 1 illustrates the situation schematically.

The Multimedia Communications Forum (MMCF), the Winsock 2.0 forum, and the Integrated Services Group of the Internet Engineering Task Force (IETF) are looking at network interface issues, while the IMA has been considering end-system services. Most major operating system and hardware manufacturers, including HP, have been involved in addressing these areas, and are developing standards to ensure consistency of interfaces.

### Demands of Multimedia on Networks

There are three areas in which real-time multimedia applications make demands on networks:

- Workgroup LAN delay must be minimized.
- It must be possible to determine delay bounds for all elements in the transfer path from source to destination so that it can be guaranteed that data will arrive at the destination within a specified time.
- Delay jitter (fluctuations in the arrival time of packets) must be controlled.

The following sections describe these three considerations in more detail and indicate how 100VG-AnyLAN addresses them.

**Control of End-to-End Delay.** The most important new requirement of real-time multimedia applications is for control of end-to-end delay.<sup>2,3,4</sup> The ITU-T's target is a maximum of

150 ms for low-quality voice services. Real-time 3D visualization transmissions are more critical. Even one-way delays of more than 45 ms have been found to cause unpleasant side effects. We decided to set a target of about 100 ms for WAN applications and 30 ms for local applications.

Table I illustrates the contribution to delay of various elements in the data path of a remote collaboration system. The system is shown in Fig. 2.

Table I  
Delays in a Remote Collaboration System

System Component	Typical Delay Value
Video Coder Delay	2 to 150 ms, depending on the CODEC used
Video Decoder Delay	2 to 150 ms, depending on the CODEC used
Packetization Delay	8 ms, using 1500-byte packets and a 1.5-Mbit/s stream
Workgroup LAN Delay	0.1 to over 200 ms, depending on the protocol used
Backbone Network Delay	10 ms typical for each backbone traversed
Router/Bridge Delay	5 to 10 ms, if priorities are implemented
Wide Area Network Delay	5 ms per 1000 km, or about 25 ms for the U.S.A., coast to coast

Of these delays, video coding and decoding have the greatest ranges. CODEC hardware is evolving extremely rapidly. Those using a JPEG processor such as the C-Cubed Microsystems CL550 have a delay of less than 10 scan lines (a few milliseconds). Currently H.261 and MPEG CODECs can impose a considerable delay. These should reduce as the technology progresses, but even so, it is possible for the




Fig. 2. Example of a remote collaboration network.

CODEC, packetization, and wide area network delays to account for all of the 100-ms delay budget. It follows then, that for effective transmission of real-time multimedia, the workgroup LAN delay must be minimized. To achieve our targets for wide area and local transmissions, we set a workgroup LAN delay target of about 10 ms.

In many cases the basic operation of the 100VG-AnyLAN protocol keeps packet delay across the LAN bounded below 10 ms. As network size and complexity increase, the high priority level can be used for multimedia traffic, possibly in conjunction with resource allocation schemes as described on page 37.

**Bandwidth Guarantees.** Application vendors need to be confident that their applications will work no matter what other traffic is on the network. The need for guaranteed high-bit-rate services to a single node or a few nodes extends across all elements within a network complex, from the LAN to bridges and routers and wide area networks. One simple method of connecting LANs to a WAN is through the use of a dial-up remote bridge (see "Remote Bridge Example" at right).

Within a LAN, 100VG-AnyLAN can guarantee bandwidth to multimedia applications either by exploiting the basic operation of the protocol or, if more flexibility is required, through the use of resource allocation schemes as described on page 37.

**Interarrival-Time Jitter.** Packet switching has been put forward as an effective technology for integrating stream-type data, such as video and audio, with conventional computer data. Packet switching can allow flexible allocation of bandwidth to individual calls, interfaces that are well-matched to computer systems, and efficient multiplexing of sources.

However, packet switching introduces a problem. Packets will be subject to variations in arrival time (known as interarrival-time jitter). How then can reconstruction of the continuous audio or video signal from a set of packets sent over a network be ensured?

Jitter is produced when a source transmits data packets at regular intervals but each packet is subject to a different delay before it arrives. Delay is caused partly by fixed elements, such as the propagation time and packetization delays, and partly by variable delays from queuing and waiting for access to the network. It is these last two that result in jitter.

Provided that any jitter is kept below a maximum value, it is possible to calculate a buffer size at the receiver that is sufficient to smooth out any variation in interarrival times. The




Fig. 3. Recovering jittered data.

## Remote Bridge Example

To illustrate the operation of a dial-up remote bridge, Fig. 1 shows two sites, each with a backbone FDDI network linking 100VG-AnyLAN hubs.

Users connected to end nodes 1 and 2 want to communicate by video and audio between the two sites. The backbone networks on each site are connected to the wide area network (WAN) through dial-up remote bridges specifically designed to carry multimedia traffic. We assume that site A and site B can also communicate through existing routes, such as some kind of corporate internet.

When the call between node 1 and node 2 is initialized, an isochronous link with the appropriate bit rate is set up between the two dial-up bridges. Bridge A is told the 48-bit MAC address of node 2. Any subsequent packets that appear on the backbone FDDI bearing that destination address will be forwarded across the WAN. Bridge B puts them on the backbone FDDI at site B, from which they reach node 2.

Any normal-priority data packets destined for node 2 will not be picked up by bridge A because they will carry the destination address of the router connecting site A to the corporate internet, not that of node 2.

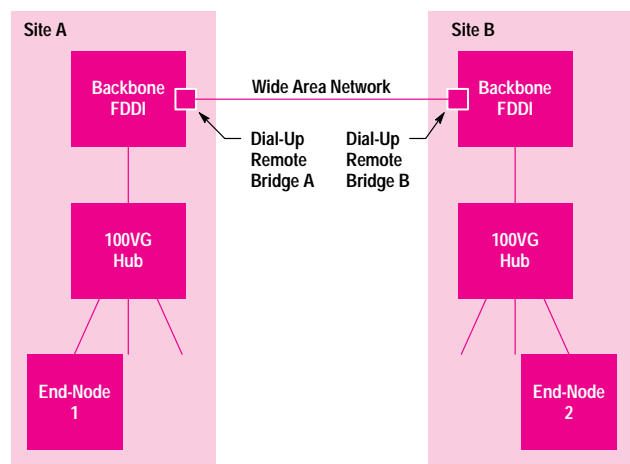


Fig. 1. Remote sites connected by dial-up remote bridges.

data can then be retrieved from the buffer in a steady stream for decoding and display. Fig. 3 shows how this is achieved when an arrival window can be determined.

The arrival window must be small, or the buffering delay will cause the total end-to-end delay of the transmitted data to exceed the 100 ms target.

## Introducing New Applications

Because interface standards and technology are still being developed for the whole wide area network structure, it is

---

---

## Higher-Level Protocols

The introduction of new LAN technologies (such as 100VG-AnyLAN) is only one element in the development of multimedia networking. Higher-level protocols are also needed that can control and transfer stream data over multiple network hops. The Internet Engineering Task Force (IETF) and the International Telecommunications Union (ITU-T) have been working in this area.

### IETF

The IETF has been developing protocols for a multiservice, packet-based network spanning the world (see Fig. 1), such as the Internet Stream Protocol (ST-II). This is a network-layer protocol roughly equivalent to Internet Protocol (IP), but specifically developed to support stream-based traffic. ST-II is being used in the Multimedia Teleservices section of the BERKOM-II program (see "Related Projects" on page 38).

The IETF has a number of groups looking at new protocols for multimedia that can be introduced into the worldwide Internet.




Fig. 1. IETF involvement in worldwide networks.

### ITU-T Narrowband ISDN Standards

The ITU-T works towards developing standards for data transmission over an enhanced telephone network, such as ISDN or other digital circuits with fixed bit rates up to 2 Mbits/s (Fig. 2). 100VG-AnyLAN could provide desktop connections to these services. ITU-T Recommendation H.320 includes the following:

- H.261: Video CODEC for audiovisual services at  $p \times 64$  kbits/s
- H.242: A system for establishing communication between audiovisual terminals
- H.221: A frame structure for a 64-to-1920-kbit/s channel in audiovisual teleservices
- Q.931/2: A D-channel signaling protocol.




Fig. 2. ITU-T involvement in worldwide networks.

---

---

important that there be an easy growth path for the introduction of new multimedia applications. This also satisfies the need for users with existing networks in place to be able to build them up to cope with advanced multimedia traffic in stages, using the existing infrastructure and avoiding a large initial outlay. Unlike other high-speed campus LANs, it is possible with 100VG-AnyLAN to provide such a growth path for the introduction of multimedia applications.

The growth path has three phases:

- Phase 1. Run the application using a single priority level throughout the network.

- Phase 2. Move some nodes to high priority as multimedia traffic increases, with simple mechanisms for monitoring and controlling bandwidth.
- Phase 3. Introduce automatic bandwidth allocation mechanisms for all high-priority traffic.

### Phase 1: Multimedia Using a Single-Priority 100VG-AnyLAN.

In the first phase, since no new API for applications has been developed, the application is simply put on normal priority all the time. Two of the properties of the 100VG-AnyLAN demand priority protocol enable it to be used with multimedia applications even in this phase: bounded delay and bandwidth sharing.

**Bounded Delay.** The access delay for a data packet is guaranteed to be less than or equal to the maximum packet duration times the number of active ports on the hub. If Ethernet-length packets are used, the maximum packet length is 12,000 bits, giving a duration of 0.12 ms. With a 32-port hub, a queued packet will have to wait at most  $0.12 \text{ ms} \times 32 = 3.84 \text{ ms}$ . This will increase to 7.84 ms if all traffic is multicast and the system is using bundles.

These figures compare with  $32 \times \text{TTRT}$  (target token rotation time) for asynchronous FDDI (giving 320 ms for a 10-ms TTRT), and unbounded for CSMA/CD.

**Bandwidth Sharing.** Every active node has an equal share of transmission opportunities. In the worst case, when every node is actively transmitting full-length packets of data at the same priority, this will give a maximum bandwidth of about 3 Mbits/s per node on a 32-port hub. This reduces to 1.5 Mbits/s if all data is multicast in bundles. This guarantee is more than sufficient for many applications such as simple conferencing.

It follows that a 32-port hub can support simple applications using real-time data at 1.5 Mbits/s or less, and application vendors can be confident that the system will work no matter what other data is on the LAN.

### Phase 2: Use of High Priority without Allocation

If the second priority is used, the two properties listed in Phase 1 apply to high-priority traffic. If 32 nodes run at high priority in a large network, say of 100 nodes, the guarantees would be the same as in the previous example.

A malfunctioning high-priority node can be prevented from saturating the network by network monitoring (the hub can refuse access to a high-priority transmission), and by configuring the network so that normal-priority traffic is raised to high-priority after waiting a certain period of time.

This provides a simple way of introducing multimedia applications, but in the longer term, a mechanism for allocating bandwidth may be required. This will control the use of bandwidth by high-priority traffic automatically, and will ensure that delay is kept within an upper bound on large networks.




Fig. 4. 100VG-AnyLAN end node stacks.

### Phase 3: Resource Allocation in 100VG-AnyLAN

A third property of 100VG-AnyLAN facilitates the allocation of bandwidth on networks: spare bandwidth utilization. Any bandwidth not used by high-priority traffic is immediately available for normal-priority data.

This property is particularly relevant if high-priority traffic is bursty. It also ensures that normal-priority traffic does not need to participate in the allocation process.

Fig. 4 illustrates how the protocol stacks for high and normal priority are independent. This is important because resource allocation protocols need only be introduced on end systems that actually run multimedia applications. End systems that only use normal priority do not need any new software.

Note that the use of a bandwidth allocator requires higher-level interfaces and services to be defined, such as those on which the IETF is working (see "Higher-Level Protocols" on page 36).

**Example of Bandwidth Allocators.** Consider the following simple scenario illustrating the use of a bandwidth allocator. Several 100VG-AnyLAN hubs are connected by an FDDI ring. The allocator ensures that the sum of the high-priority synchronous multimedia connections across the FDDI does not exceed 80 Mbits/s, and that the sum of the high-priority multimedia connections across any individual 100VG-AnyLAN network does not exceed 80 Mbits/s. This ensures that the 100VG-AnyLAN networks and the FDDI backbone are not overloaded with high-priority traffic.

Suppose that all the 100VG-AnyLAN hubs are 32-port or smaller. The protocol guarantees that in the event of all nodes simultaneously wishing to transmit normal-priority data, they will each have a bandwidth of a little over 3 Mbits/s. The allocator then only has to ensure that the sum of the high-priority synchronous multimedia traffic across the FDDI does not exceed 80 Mbits/s. A system such as this can typically support up to 80 video connections each with a peak bandwidth of 1 Mbits/s, or up to 200 364-kbit/s connections.

In 100VG-AnyLAN, the proposed mechanism for allocating bandwidth and setting delay bounds is the target transmission time.

**Example of Target Transmission Time Allocation.** A nominal time value, the target transmission time (TTT) is set, and nodes request permission from the allocator to transmit a maximum amount of data within a TTT period. Provided that the sum of all such requests is less than a preset maximum, say 80% of the network capacity, the allocator grants the request. The end node is then guaranteed that its data will be transmitted during the TTT period.

Fig. 5 shows a simple example of how this could work. In the example, there are four stations operating at high priority. Nodes 1 and 2 are each generating 30 Mbits/s, node 3 is generating 7.5 Mbits/s, and node 4 1.5 Mbits/s.

Node 4 generates one 1500-byte packet every 8 ms and requires this to be transmitted onto the network within 8 ms. If this is the lowest access delay, the TTT would be set at 8 ms. In this time period, the other nodes will generate 30kB (20 1500-byte packets), 30 kbytes, and 7.5 kbytes (5 1500-byte packets) of data, respectively.

Fig. 5 shows how the data in each of the end nodes' transmission queues is transmitted on the network.

**Bandwidth Monitoring.** High-priority traffic can be monitored centrally at the hub or through a separate device attached to a promiscuous port. From packet lengths and interpacket gap times, the monitor can easily calculate the aggregate

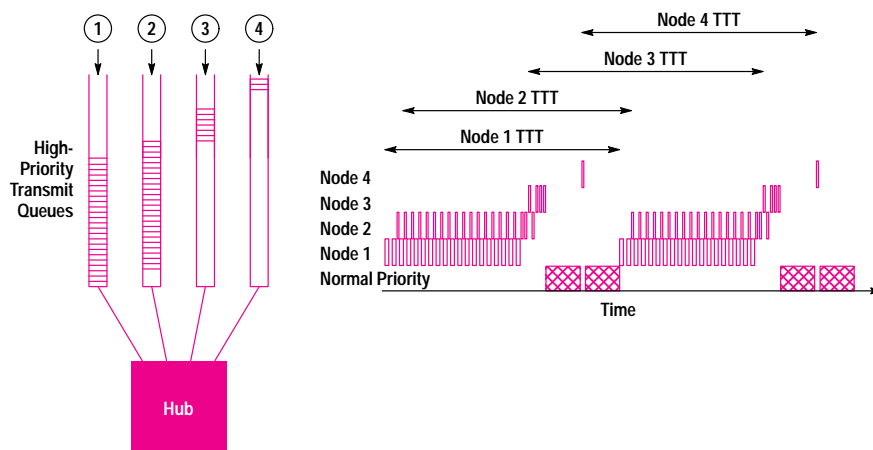


Fig. 5. Example of the use of the target transmission time (TTT).

---

## Related Projects

Two projects are under way that have a bearing on the development of networks to carry multimedia traffic.

### Teleservices Project

The Teleservices project is part of the Deutsches Bundespost (DB) Telekom BER-KOM technology and demonstrator program. The program covers wide area network and campus broadband communication studies, using broadband switches and local optical-fiber links installed in Berlin to form a large test bed. It now forms part of the German Broadband Integrated Services Digital Network (B-ISDN) asynchronous transfer mode (ATM) pilot network, with switching centers in Cologne and Hamburg.

The Teleservices Project is investigating the provision of multimedia services over wide area, broadband networks. HP is involved in two areas: multimedia collaboration (MMC), supporting conferencing and shared applications, and multimedia transport (MMT), developing high-speed networking protocols for multicast operation.

HP joined in April 1993, with collaboration from HP Germany, HP Laboratories in Bristol, and TUB (Technical University of Berlin, as an HP subcontractor). The main HP/TUB task is the development of high-speed network protocols and the overall implementation and integration of the complete collaboration application under the multimedia collaboration task for the demonstration of interoperability. Five HP 9000 Model 755 computers have been installed at the TUB for the multimedia implementation, together with a range of audiovisual equipment, communications interfaces, and associated software packages. Prototype HP ATM switches have also been installed.

The program will produce a limited HP multimedia communications and management product for use on HP 9000 Model 7x5 workstations in the German broadband trials. It will also provide measurement and performance evaluation results, including:

- Observations of the interactions between video and audio compression schemes, transport protocols, the different network technologies used, and their impact on perceived performance

---

throughput of the hub for each TTT. If the aggregate is higher than a certain threshold value, further action can be taken on individual streams. For example, the network manager could disable high-priority traffic at a particular port, or this could be done automatically through the network monitor. A technique such as this provides a high level of security to other nodes, even from malicious attempts to disrupt the high-priority service.

- Identification of particular issues and problems associated with multimedia applications related to HP products and commercial interests
- Performance measurement and behavior of a transport protocol designed for multimedia applications running over different network technologies.

A demonstration was given in August 1993 using the first version of the conference application with software video compression and relatively simple control options. In March 1994, an enhanced version was demonstrated at the CeBIT exhibition in Hannover. This had a range of management and control options, an audio server, and full-motion hardware compression (M-JPEG). It also included multivendor terminal interworking at the show over the ATM demonstration network and conferencing over ATM link connections back to a terminal in Berlin.

A demonstration of the Teleservices package between HP Laboratories in Bristol and the TUB, using ISDN at about 300 kbits/s showed that a minimum transfer rate of 1 Mbit/s is required for successful operation. Wideband operation over ISDN connections is also being evaluated.

### Distance Learning Seminar

In October 1994, a pan-European seminar on distance learning was held on a network linking three European universities (the Royal Institute of Technology, Stockholm, the Technical University of Berlin, and University College London) and various other European sites. The seminar was sponsored by HP and aimed to promote the use of distance learning and to explore the future of the technology and the role of multimedia.

The seminar was a practical demonstration of multimedia over a mix of transmission technologies. German Telekom provided B-ISDN (broadband ISDN ATM) links between Berlin and London and Berlin and Stockholm. Other links used inverse multiplexers at each site to provide contiguous bandwidth over ISDN of 384 kbits/s. This provided adequate bandwidth for 25-frames-per-second video using hardware compression based on the H.261 standard, plus good-quality audio.

## References

1. D. Blankenhorn, "MCI Announces Switched T1 and T3 Services," *Newsbytes*, 1991.
2. "General Characteristics of International Telephone Connections and Circuits, Recommendations G.100-G.181," *ITU-T Blue Book*, Vol. III, Fascicle III.1, 1988.
3. A. Shah, D. Staddon, I. Rubin, and A. Ratkovic, "Multimedia over FDDI," *LCN Conference Proceedings*, 1992, p. 110.
4. I. Corbett, "Videophony," *Audiovisual Telecommunications*, BT Telecomms Series, Chapman & Hall, 1992.