White Paper

November 5, 1995

Video Over ATM and Existing Networks

ATM and the Future of Digital Video

As video applications developers, content providers, and operating system vendors develop higher-quality methods of communication using digital video as a medium,[1] chip vendors and platform designers are creating lower-cost technologies for doing the same.[2] The rapid convergence of these two trends has made digital video a highly promising medium with many implications for the networks that will carry it. The simultaneous emergence of Asynchronous Transfer Mode (ATM) networking technology in the last several years has made it a favorite contender as the network of choice for video transmission.

One clear advantage of ATM is its ability to support video, voice, and data simultaneously.[3] As ATM deployment begins, however, ATM and existing networks—Ethernet, Token Ring, and circuit-switched WANs such as T1 and Integrated Services Digital Network (ISDN)—will coexist for many years. To fully Take advantage of ATM and existing networks, video applications must therefore interoperate over these network types, and these networks must interoperate with each other.

There are numerous considerations and tradeoffs to consider when developing an infrastructure for video applications, and no two networks will be identical. However, in the emerging milieu of hybrid shared LAN, switched LAN, and WAN networks, many of which will include ATM and other technologies, there are guidelines to follow in the creation and deployment of these networks. Moreover, the addition of ATM as a common video, voice, and data transport mechanism that transparently spans the LAN and the WAN creates additional options.

The following questions, therefore, are answered in this paper along with the respective rationales and technical discussions:

- What video applications might be deployed?
- What type of video codecs will be used?
- What network infrastructure will support these applications?
- What networking products create these infrastructures?

From the perspective of the network, video applications fall into one of three areas:

- *Packetized video*, which runs over traditional LANs at the MAC (for example, Ethernet) or network (for example, IP) layers, and that will therefore be transported into ATM networks through LAN emulation, layer 3 encapsulation, or multiprotocol over ATM (MPOA). These applications will use LAN switches, routers, and ATM switches, such as the LightStream® ATM switches, Cisco routers, and Catalyst™ and EtherSwitch® LAN switches.

---

1. Microsoft's Windows '95 will support some type of video decompression, and Apple Computer has long supported video through its QuickTime extension to the Macintosh OS.
2. Intel, for example, plans to support Multimedia extensions technology in its chips, while numerous vendors have already delivered video add-on boards for various platforms.
3. As a cell-switching technology based on small, fixed-length cells, ATM has the capability to support any traffic type—video, voice, and data—that has been adapted to ATM with the appropriate ATM Adaptation Layer (AAL). All of these cells of various traffic types are switched along virtual connections through an ATM network.

- *Constant-bit-rate video*, which runs over traditional 64-kbps or multiple 64-kbps (or ISDN) lines, and will be transported over ATM using circuit emulation. These applications will use ATM service multiplexers, ATM switches, or existing leased-line facilities.
- *Packetized video*, which will run natively over ATM using an MPEG2-to-ATM convergence layer. Standards for doing such video transmission are still in the nascent stages of development and probably will not be widespread for several years. These applications will make use of specialized MPEG2 video to ATM codecs,[4] still in early stages of development.

As video applications mature and find additional native support in end stations—whether they are PC hardware or software, videoconferencing codecs, video servers, or the like—network infrastructures will lead the way in paving the support for widespread deployment of video. Thus, corporations and network providers faced with an array of video applications can remain confident in the underlying infrastructure, recognizing that they can continue to leverage their existing infrastructures while moving to the world of cheaper and higher-quality video applications that run over networks of increasingly higher bandwidth using ATM.

This paper is structured as follows:

- Executive Summary
- Packetized Video vs. CBR Video
- Network Protocols to Transport Video
- Networks and Network Equipment to Support Video
    - Packetized Video (Motion-JPEG, MPEG, MPEG2, Cinepak, etc.)
    - Constant Bit Rate Video
- Technical Requirements for Video Transmission
    - Dimensions of Video Variability
    - Technical Requirements for Delivery
    - Video Compression
    - Compression Standards
    - Display and Transmission Formats
    - Bandwidth Requirements
        - Bandwidth Requirements for Moving Pictures
        - Bandwidth Requirements for Still Images
    - Latency Requirements
- ATM Internetworking and Video Support
    - LAN Emulation
    - Native Layer 3 over ATM—Multiprotocol over ATM (MPOA)
- Performance Guarantees in Existing Networks
    - Resource Reservation Protocol (RSVP)
    - Protocol-Independent Multicast (PIM)
- Native ATM Support for Video Applications
    - MPEG2 convergence to ATM
    - Real-time VBR for Video Applications
    - ATM Adaptation Layers for Video Traffic
    - Multicasting in ATM
    - ATM Service Multiplexers
- Network Capabilities for Video Applications
    - Bandwidth Requirements
    - Performance Requirements
- Video Application Types and Network Design
    - Packetized Video
    - Distance Video
    - Video on Demand (Cable TV / Telco TV)
- Summary
- Appendix A: MPEG2
- Appendix B: Performance Issues for High-End Video over ATM

4. A codec is a device that provides compression and decompression of a digital video signal, although it is commonly used to refer to a product that also digitizes an analog video signal, such as the NTSC, PAL, or SECAM signal from a television or VCR, and then applies compression.

Table 1.

Application Type	LAN-Based video	Distance Video	Video on Demand (Cable TV / Telco TV)
Video Applications	• Video courseware/training • Desktop videoconferencing • Application sharing • Graphic visualization • Video kiosks	• Remote classroom/distance learning • Videoconferencing • Telecommuting • Telemedicine • Telejustice	• Video on demand • Near video on demand • Interactive video games
Video codecs and Servers	• PC-based codecs, hardware, and software • Video servers	• Standalone codecs • PC-based hardware/software codecs • Video servers	• Standalone codecs (Other) • Set-top box • Video servers
Video Codec Formats	• MPEG, MPEG2, H.320 • Proprietary—motion JPEG, AVI, Indeo, Cinepak, others	• MPEG, MPEG2 • MPEG4 (future) • H.320 / H.261	• MPEG2 • Existing analog protocols (QAM RF modulation)
Network Infrastructure: Protocol and Format Perspective	• Packetized video running over layer 2 or layer 3 • Layer 2 or layer 3 internetworking with ATM	• CBR video running over circuit emulation • Packetized video over layer 2 or layer 3; internetworking with ATM	• Packetized video running natively over ATM (future) and Coax or ADSL • Analog video using RF modulation (today)
Network Infrastructure: Configuration	• Highly segmented LANs with one or few users per segment • ATM backbone requirement depending on number of videos • ATM to desktop	• Dedicated or on-demand WAN lines (leased lines, ISDN, etc.) • Minimum 64 kbps for H.320 protocols • Minimum 1.5 Mbps for MPEG protocols	• ATM fiber networks to head-end and coaxial cable to home (hybrid fiber coax) • Fiber to the curb or home (FTTC, FTTH) • ADSL
Network Infrastructure: Performance Guarantees	• ATM Quality of Service • ATM switched multicast circuits • RSVP (Resource Reservation Protocol)	• ATM Quality of Service • ATM switched or permanent circuits • RSVP (Resource Reservation Protocol)	• ATM Quality of Service • ATM switched or permanent virtual circuits
Network Infrastructure: Products	• ATM switches • LAN switches • Routers with LAN switching and ATM ports	• Enterprise switches that support CBR, ATM, and LAN connections	• Enterprise switches • RF modulators for coaxial connections
Recommended Cisco and Partner Products	• LightStream 100 ATM switch • Catalyst LAN switches with ATM cards where needed • Cisco routers with ATM Interface Processor (AIP) where needed • Cisco NIC cards • Partner NIC cards	• LightStream 2020 ATM switch • LightStream 100 ATM switch • ATM Multiplexer	• LightStream 2020 ATM switch • LightStream 100 ATM switch

Packetized Video vs. Constant Bit Rate (CBR) Video

Video applications are either packetized, which means that they are *bursty or variable bit rate*, or they are *constant bit rate*. The packetized video types were designed to run over traditional LANs, and make use of some type of compression algorithm whose output is dropped into traditional packets, whether they are IP packets (or the equivalent) or Ethernet frames (or the equivalent). See the section entitled “Video Compression” for more information.

Because packetized video runs over LAN infrastructures, it can perform even better when the LAN infrastructure has protocol support and enhanced features to guarantee quality of service. Cisco has been very active in developing the standards and software to do just that by developing Protocol-Independent Multicast (PIM) and the Resource Reservation Protocol (RSVP), for example. See “Performance Guarantees in Existing Networks” for an in-depth discussion of these protocols and how they enhance video applications.

In addition, when LAN-based video will run over ATM, there must be some translation from ATM to the legacy LAN, whether by LAN Emulation or by Multiprotocol over ATM. See ATM Internetworking and Video Support (section 6.0) for a discussion of these ATM internetworking standards; for an in-depth discussion, refer to a white paper entitled “ATM Internetworking” by Anthony Alles, Cisco Systems. In addition, there is an effort under way to provide translation between RSVP and ATM Qualities of Service.

When packetized video runs natively over ATM, there must be a convergence layer between the video stream and the ATM Adaptation Layer (AAL)—most likely AAL5, according to the most recent work of the ATM Forum. Such work is in the early stages, and the development of standards-based video compression algorithms for two-way, low-bandwidth, high-quality video in the form of MPEG4 is also still under way. The ATM Forum has not even begun to address MPEG4 over ATM, and instead is tackling MPEG2 over ATM, which is well suited for applications such as video broadcasting in LANs or by cable TV. MPEG2 is a ratified open standard.

Compressed video can be bursty and lends itself well to framing or packetizing. In the case where video must run over p x 64-kbps lines for wide-area transmission, the compressed bursty video bitstream must be transformed into a CBR bitstream to meet the requirements of transmission over these CBR digital lines. Traditional videoconferencing equipment, such as that made by PictureTel, VTel, and Compression Labs, uses some variant of the H.320 / H.261 protocol suite. This protocol suite uses buffers to ensure that bits are always sent at every “clock tick,” regardless of the original structure or traffic shape of the bursty video stream.

Such CBR video, when transmitted over an ATM network, requires circuit emulation, where the ATM transport is emulating the traditional p x 64-kbps or ISDN circuits that the H.320 equipment is expecting as its underlying network.

Network Protocols to Transport Video

Video per se includes numerous types of packets and bitstreams, depending on the compression and convergence protocols. Because manufacturers of video codecs (with the possible exception of WAN videoconferencing H.320 manufacturers such as PictureTel, VTel, and Compression Labs) do not assume anything about the characteristics of the underlying network, video packets can theoretically ride on top of numerous network and data-link layers. Even H.320 / H.261 protocols, however, can be made to interoperate with LAN architectures.

Because the value of having network layers is to hide higher-layer transport and applications from the specifics of the underlying data-link layers, a video application that will interoperate across a heterogeneous network will make use of network layers such as IP, or at minimum, a MAC layer such as Ethernet, although using MAC layers without network layers carries all of the traditional problems of bridging—unroutability and flat network design.

Certain video applications will ride directly over ATM, such as MPEG2 program or transport streams for video broadcasting in the cable TV industry, but these will require modulation to put on analog cable.

Video can be transported over different layers in the protocol stack. For interoperability, Figure 1 shows that different codecs will employ different methods for networking. The clear implication is that the choice of application will determine the type of network and vice-versa.

Figure 1. Video and the Protocol Stack

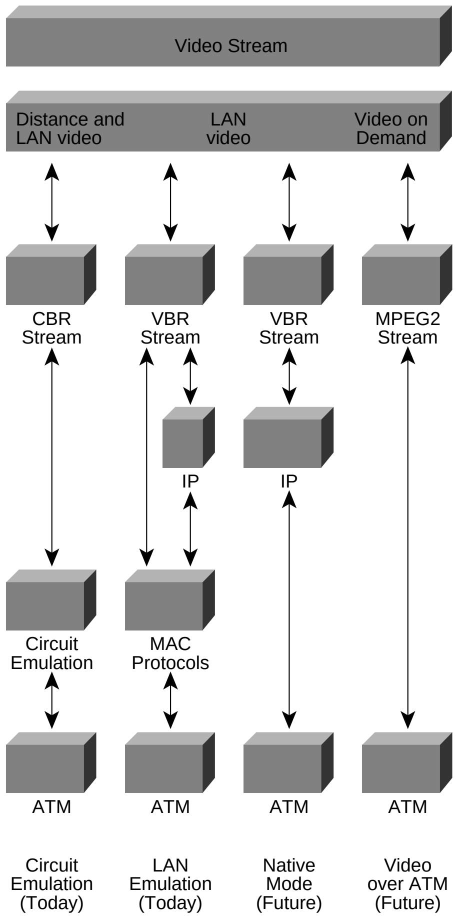

If the codec includes an IP layer, it can take advantage of IP multicasting and routing as well as multiprotocol over ATM internetworking. If it lacks an IP layer, but instead puts video into Ethernet or Token Ring frames, it will not be routable, but it still can make use of Ethernet data-link multicast on a LAN or LAN Emulation services over ATM.

Similarly, H.320 video that is transported over circuits can make use of circuit emulation over an ATM network but will not be inherently routable. Such video connections will typically terminate on another H.320-compliant device for two-way videoconferencing.

Networks and Network Equipment to Support Video

Packetized Video (Motion-JPEG, MPEG, MPEG2, QuickTime, AVI, Cinepak, and Numerous Others)

For packetized, LAN-based video, high-performance LAN switches and routers make an excellent infrastructure for deployment. For bandwidth reasons, the more highly segmented the LANs (and hence the fewer collisions), the better performance a video application will display.

Hence, LAN switches such as the Catalyst or EtherSwitch family of switches make an excellent choice for video-to-the-desktop deployment. In addition, Cisco routers with high-performance switching engines and support for the PIM and RSVP protocols allow for bandwidth efficiency and delivery of the quality-of-service guarantees needed by the video application. Quality of service requirements are discussed further in the section “Performance Guarantees in Existing Networks.”

As networks are required to carry more video streams, more bandwidth to the server and in the backbone is required. ATM comes into play for high-end video to the desktop or where multiple video streams create high backbone bandwidth demands. Therefore, for effective deployment, the use of an ATM uplink from the LAN switches and routers, such as Cisco’s AIP is necessary, along with the LightStream switches for ATM connection.

For effective ATM connections, today the LightStream 100 is an excellent choice, because it has switched virtual circuit (SVC) and multicast SVC support as well as independent queuing for high-priority traffic (such as video) and for multicast traffic (for example, video broadcasts). These features are required for LAN Emulation and MPOA, as well as for bandwidth efficiencies. In the future, other members of the LightStream family will also support SVCs.

A hybrid network for video deployment might appear as shown in Figure 2, using lower-cost LAN switching to the desktops, and deploying ATM in the backbone and to very-high-resolution equipment.

Figure 2. Hybrid Corporate Video and Data Network

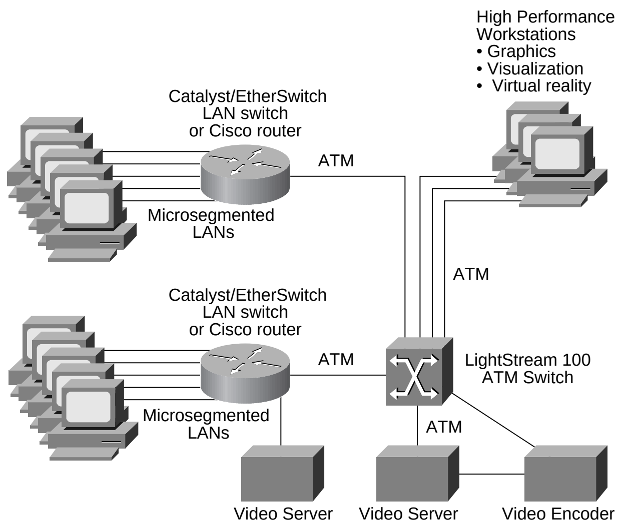

In the future, it is expected that the service guarantees of RSVP will map to ATM QoS (Qualities of Service). Once such work is complete, the service guarantees will interoperate across a hybrid network.

In the interim, using LAN switching and ATM Forum standards-based LAN Emulation over ATM, the network depicted in Figure 2 can be deployed with each end user given a dedicated Ethernet connection to guarantee delivery from the server to the desktop.

As more and more desktop videoconferencing equipment is deployed, the vast majority will be directly attached to LANs, especially as switched LANs become more the norm. One analyst group has projected that eventually, 75 percent of videoconferencing equipment will be packetized and attached to a LAN. Note that these numbers do not account for the rapidly growing small office and home office market (see Figure 3).

Figure 3. Growth of LAN-Attached Videoconferencing

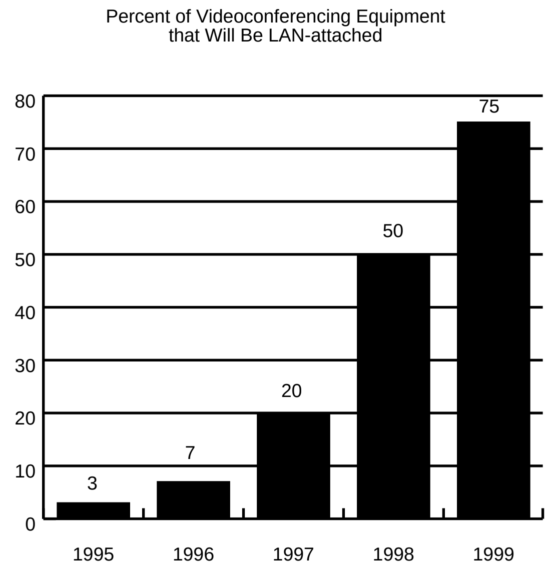

Source: Gartner Group, 1994

Constant Bit Rate Video

For constant bit rate (CBR) video connections, a service multiplexer is needed for ATM switches without CBR cards, such as the LightStream 100, because the multiplexer has inputs for CBR traffic at T1/E1 or T3/E3 speeds and can adapt those streams to ATM. An ATM multiplexer and provides ATM adaptation with an OC-3 (155 Mbps) ATM port for connection to ATM switches.

The LightStream 2020 ATM switch has built-in CBR line cards, so a service multiplexer is not required when the video codec is colocated with the switch; the codec can feed directly into the CBR card. For remote sites where there are multiple video feeds coming into a LightStream 100 or 2020 switch, multiplexers play the role of providing higher connection density per card or port for fan-in or fan-out, and distance connectivity where low-cost remote connections to the enterprise switch are desired. See Figure 4 for an example of such a network.

Figure 4. Service Multiplexers for CBR Video

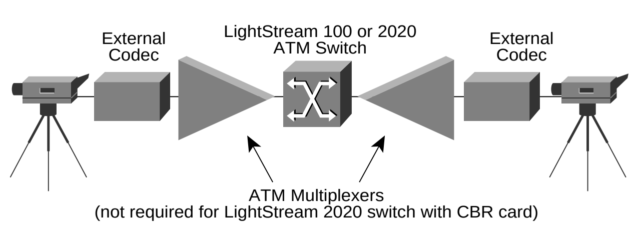

Technical Requirements for Video Transmission

Dimensions of Video Variability

These applications highlight distinct points on a spectrum of video applications. Video applications can be classified along a number of dimensions, such as real-time versus stored data, point-to-point or multicast, bandwidth required, and delay (latency) and jitter tolerance.

The simplest division is between real-time and stored data streams. Stored data can be transmitted in bursts, because it is up to the end user to reconstruct the timing information by playing back the application. Real-time applications, however, are those where the latency between source and destination must be low to ensure a continuous, smooth conversation.

Video applications also span point-to-point to multicast or broadcast. Point-to-point applications include video telephony; multicast can include distance learning where one instructor is transmitted live to several locations, LAN TV for corporate broadcasts, and video transmission by cable companies.

These applications can be mapped along the dimensions shown in Table 2.

Table 2. Video Application Types

Application	Stored Data Streams	Real-Time Interactive
Point to Point	Multimedia mail, multimedia notes (which include images or movies)	Video telephony Videoconferencing
Multipoint	LAN TV (stored information) Corporate training Financial broadcasts	Distance learning (where students can respond to instructor) Kiosks Videoconferencing Live broadcast

Stored data streams that are locally played in bulk can be treated as data (albeit large amounts of data), because bandwidth does not need to be reserved—the information can be downloaded as bandwidth is available. A CD-ROM is an example of this capability—the data is stored and then downloaded directly into the PC.

The challenge in adding video to a network is the integration of real-time information into an existing network. Some vendors propose creating parallel networks to carry video, but the enormous equipment and administrative cost to the user makes this solution impractical in most cases, except for those few critical locations where applications such as videoconferencing justify the existence of a separate network. This scenario exists today in corporations that make use of H.320 videoconferencing equipment.

The advantage of video over data networks is that users can use their existing data infrastructures—whether they are traditional LANs or ATM networks—to carry video alongside traditional data.

Just as today's voice, data, and video networks are separate, the future promises convergence to a common infrastructure. It is not likely that voice networks will be the first integrated, because of the high efficiency of voice and TDM (time-division multiplexing) networks today. It is probable that video applications, most of which are packetized, will share a common infrastructure with data—with future integration of voice as ATM standards and technology mature.

As networks begin to integrate ATM into their infrastructures, certain video applications become even more cost-effective, even while they are enabled today with a LAN-switched- and-routed network (in the case of corporate networks) or with existing analog wire or coax (in the case of telcos and cable companies).

Technical Requirements for Delivery

Real-time and near-real-time video may require several attributes from the network in order to be of high enough quality to please users with the service or application offered.

The main requirements for delivery of video are as follows:

- *Sufficient bandwidth*—depending on the amount of video information transmitted and the compression method
- *Low latency*—necessary for "live" sessions such as long-distance teleconferencing when users at either end must respond to each other without undue delays
- *Low jitter*—required even in one-way applications to avoid frame slips or poor synchronization between audio and video
- *Efficient multicast*—so that the entire video stream is not replicated n times for n users, but is replicated only at the last possible branch point to each user

In general, using standards-based compression schemes such as MPEG2 (to be discussed later), video applications can be supported with the bandwidth and latency requirements shown in Figure 5.

Figure 5. Latency and Bandwidth Requirements for Video Applications

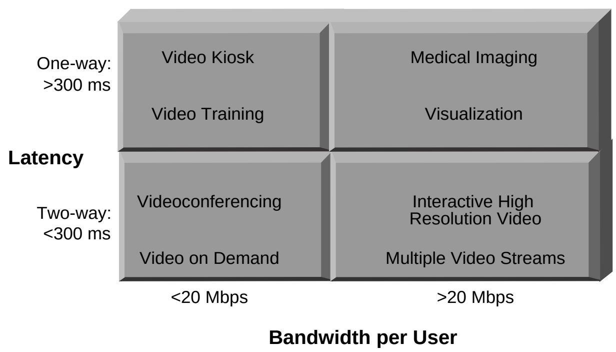

Note that the bandwidth cutoff is more accurately portrayed as a continuum; this chart highlights the ability of switched Ethernet or Token Ring to support most video applications. When multiple streams are shared on one backbone, however, bandwidth requirements may be such that ATM is called for.

Video Compression

There are numerous video compression formats, thanks largely to the emergence of video from the arenas of the motion picture, television, and video game industries. Many of these formats are proprietary, such as Intel's Indeo, but there are a number of open standards as well.

Protocols include methods for:

- *Digitizing*—translation from analog signal to digital signal and vice versa
- *Compression/Decompression*—includes compression within a frame (spatial compression) and compression between frames by some interpolation scheme (temporal compression)
- *Convergence* to lower protocol layers
- *Multiplexing* of different video and audio sources into one flow of information or transport stream, then demultiplexing them on the other end

For example, in the H.320 protocol suite used for videoconferencing, the H.221 protocol provides for convergence and multiplexing, while H.261 is the protocol for video compression and decompression.

Compression Standards

Predominant compression standards today include the following:

- *H.320 / H.261 (circuit-switched) protocol suite*—developed by the International Telecommunications Union (ITU), formerly the CCITT, and used for videoconference device connection to p x 64-kbps lines or ISDN lines.
- *MPEG 1 and 2*—the ISO's Motion Pictures Expert Group has promulgated these standards for video compression. MPEG 2 is fast becoming the most common version of these standards. MPEG provides for compression by selecting a certain number of frames that are transmitted in full and interpolating or reconstructing other frames that appear between the frames fully sent. This is known as temporal compression, because the compression is applied along the time dimension, where every few frames (called intra-frames, or I-frames) are transmitted in full, while others in between are interpolated from the I-frames.

See Appendix A for a brief discussion of MPEG2.

- *MPEG4*—a standard still under development. Its key application is for low-bandwidth videoconferencing where image quality is important.
- *JPEG*—the ISO's Joint Photographer's Expert Group is a standard for still picture (frame) compression, or spatial compression.
- *Motion JPEG*—a number of vendors use motion JPEG, which is proprietary, where all frames are transmitted but each individual frame is a compressed JPEG image. Motion JPEG (M-JPEG), can be used for moving video, but requires more bandwidth than MPEG protocols, because it only provides compression within each frame but not between frames. Bandwidth requirements may vary from several Mbps to 40 Mbps, depending on the image resolution and frames per second.
- *Indeo, Cinepak*—proprietary standards developed by vendors for desktop-based video applications that are designed to run over LANs.

Display and Transmission Formats

- *Common Interface Format (CIF)*—an ITU standard for low-resolution screen format used for desktops and conference rooms. QCIF, or quarter CIF, gives half the linear resolution of CIF, giving the image one-quarter the overall resolution. The CIF standard allows for a continuum of frame rates up to 30 per second and a continuum of compression ratios.

- *CGA, EGA, VGA, SVGA*—standards for screen display used by computer monitors. See Figure 7 for screen resolution information.
- *NTSC, PAL, and SECAM*—the standards used by televisions and VCRs for analog video signals in the U.S., Europe, and France, respectively.

Bandwidth Requirements

Bandwidth Requirements for Moving Pictures

Uncompressed video is usable where bandwidth is not an issue. Different standards require varying levels of bandwidth because of both the picture size and resolution and the number of frames transmitted per second. Table 3 shows bandwidth requirements for various full-screen display standards, for both computer monitors (CGA through SVGA), and for television screens (NTSC through SECAM).[5]

5. Hodge, Winston *Interactive Television* 1995: McGraw-Hill, Inc., San Francisco, CA, p. 21

Table 3. Bandwidth Requirements without Compression

Format	Pixels per line	Line per frame	Pixels per frame	Frames per second	Million pixels per second	Bits per pixel	Megabits per second
CGA	640	200	128,000	60	7.7	4	30.7
EGA	640	350	224,000	60	13.4	6	80.6
VGA	640	480	307,200	60	18.4	6	110.6
SVGA	800	600	480,000	72	34.6	8	276.5
NTSC	600	485	291,000	30	8.7	24	209.5
PAL	580	575	333,500	50	16.7	24	400.2
SECAM	580	575	333,500	50	16.7	24	400.2

Compression dramatically reduces bandwidth requirements. The image size on the screen, resolution, image depth, and transmission rate all affect bandwidth. By using compression at different resolutions and frame rates, the raw bandwidth requirement is reduced substantially. The bandwidth requirements for compressed video images is as shown in Table 4.[6]

Table 4. Bandwidth Requirements with Compression

Standard/ Format	Approximate Bandwidth	Compression Ratio*
Motion JPEG	10-20 Mbps	7-27:1
MPEG-1	1.2-2.0 Mbps	100:1
H.261	64 kbps-2 Mbps	24:1
DVI	1.2-1.5 Mbps	160:1
CDI	1.2-1.5 Mbps	100:1
MPEG2	4-60 Mbps	30-100:1
CCIR 723	32-45 Mbps	3-5:1
CCIR 601 / D-1	140-270 Mbps	reference
U.S. commercial systems using “mild compression”	45 Mbps	3-5:1
Vendor methods (e.g., PictureTel SG3)	0.1-1.5 Mbps	100:1
Software compression (small windows)	Approximately 2 Mbps	6:1

With MPEG2, the level of compression can be controlled to deliver varying quality levels. To make the network requirements of the compression more clear, the estimates for bandwidth in Table 5 are given for varying MPEG2 quality levels.[7]

6. Keinath, R. and D. Minoli, *Distributed Multimedia Through Broadband Communications*, Boston: Artech House, 1994.
7. Hodge, p. 33

Table 5. MPEG2 Quality and Bandwidth

MPEG2 Quality Level	Bandwidth Required
VHS	1.5 Mbps
Broadcast	5.0 Mbps
Studio	7.0 Mbps

By now, it should be readily apparent that very high quality images often require less than the 10 Mbps that a dedicated Ethernet or Token Ring LAN can provide. Bandwidth requirements become higher when compression is not possible, such as in video production environments;[8] when lossy compression[9] algorithms cannot be used, such as in certain medical environments; or when many video streams must be transported over the same infrastructure, requiring higher aggregate bandwidth.

Bandwidth Requirements for Still Images

When still images are transmitted, bandwidth can be an issue when multiple merges must be transmitted in rapid succession. In typical applications, still images can receive bandwidth as shown in Table 6.

8. In some video post-production environments, compression of any type is not used because artifacts might be introduced and degrade the final movie. This situation is common in movie production, studio, and special effects houses.
9. Compression can be lossy or lossless. Lossy compression allows for interpolation, whether spatial or temporal, and all the MPEG and JPEG compression formats are lossy. Lossless compression includes methods for removing redundant information (such as run-length encoding) but information is not interpolated and hence no information is “lost”; it is merely transmitted in nonredundant form.

Table 6. Still Image Bandwidth Requirements

Application	Pixels	Bits per pixel	Compression Ratio	Mbits	Transfer Time (sec)	Mbps
Document image scan 8.5x11 @ 200 dpi, monochrome	1700 x 2200	1	15:1	0.25	4	0.06
Document image scan 8.5x11 @ 200 dpi, gray scale	1700 x 2200	8	5:1	6.0	4	1.5
X-ray digitizer 14x17 @ 140 dpi	1960 x 2380	12	1:1	56.0	15	3.7
Ultrasound digitizer	512 x 512	8	1:1	2.1	5	0.42

Latency Requirements

Latency is important in the case of two-way (or many-way) conversations, because the delay in response is noticeable. This perceived problem is common with current H.320 videoconferencing applications, where users must sometimes wait until the other speaker is completely finished before responding due to delay. This latency is more a function of the compression and decompression rather than a function of the switching, routing, or underlying network. For applications such as remote classrooms or improved videoconferencing, latency must be minimized.

Because the choice of codec is the most significant determinant of latency, in two-way videoconferencing applications the most commonly used codecs are H.320 and motion-JPEG, which provide lower latency than MPEG or MPEG2 codecs. However, MPEG or MPEG2 codecs using less than full I, B, and P frames can also reduce latency (see Appendix A). MPEG4 codecs promise to provide low latency along with lower transmission bandwidth requirements.

Low latency is important in videoconferencing and telephony applications, because delays above 300 milliseconds make conversations difficult. In addition, in videoconferencing applications, synchronization between audio and video becomes important, so the low latencies apply to both audio and video components.

As higher-end video applications are developed, such as those using MPEG2, HDTV or Super High Definition (SHD) television, latency and other performance specifications become more critical. See Appendix B for an in-depth treatment of high-end video over ATM.

The wide-area network and the codecs are much more likely sources of latency introduction that the network equipment itself. Figure 6 illustrates the equipment latency as compared to the latency tolerance of video applications.

Figure 6. Application Latency Tolerance and Network Product Latencies

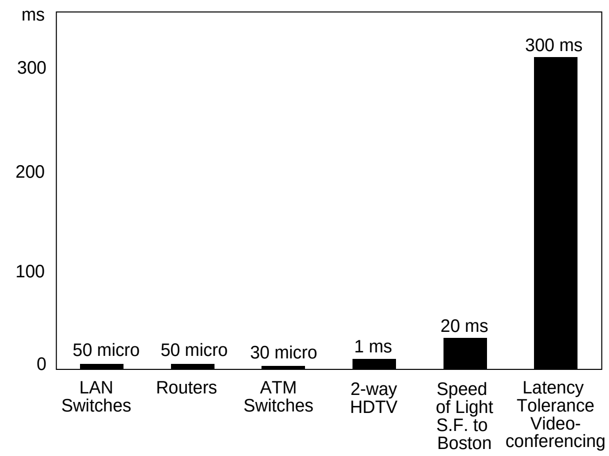

Different compression algorithms are also designed with different uses in mind. Some, like the H.320 suite (using H.261 compression), were designed for bidirectional communications, so that compression and decompression equipment costs roughly the same because it would be required at both ends. Others, like MPEG, were designed with broadcast (unidirectional) traffic in mind, so that the compression equipment (needed once) could be more expensive, while the decompression equipment (required at all receiving points) would be less expensive. These costs, however, are falling rapidly.

ATM Internetworking and Video Support

Note: much of the information in this section is adapted from the "ATM Internetworking"[10] white paper by Anthony Alles of Cisco Systems.

LAN Emulation

Given the vast installed base of LANs and WANs today and the network and link-layer protocols operating on these networks, clearly a key to the success of ATM will be the ability to allow for interoperability between these technologies and ATM. Few users will tolerate the presence of "islands" of ATM without any connectivity to the rest of the enterprise network. The key to such connectivity is the use of the same network-layer protocols, such as IP and IPX, on both existing networks and on ATM, because it is the function of the network layer to provide a uniform network view to higher-level protocols and applications.

There are, however, two fundamentally different ways of running network-layer protocols across an (overlay mode) ATM network. In one method, known as native mode operation, address resolution mechanisms are used to map network-layer addresses directly into ATM addresses, and the network-layer packets are then carried across the ATM network. The alternate method of carrying network-layer packets across an ATM network is known as LAN Emulation (LANE). The ATM Forum has recently completed a Phase 1 LAN Emulation specification.

As the name suggests, the function of the LANE protocol is to emulate a local-area network on top of an ATM network. Specifically, the LANE protocol defines mechanisms for emulating either an IEEE 802.3 Ethernet or an 802.5 Token Ring LAN. What this means is the LANE protocol defines a service interface for higher-layer (that is, network-layer) protocols that is identical to that of existing LANs, and that data sent across the ATM network is encapsulated in the appropriate LAN MAC format.

The rationale for doing this is that it requires no modifications to higher-layer protocols to enable their operation over an ATM network. Since the LANE service presents to network layer drivers the same service interface given by existing MAC protocols—for example, an NDIS or ODI like driver interface—no changes are required in those drivers.

10. Alles, Anthony *ATM Internetworking* Cisco Systems: 1995

It is envisaged that the LANE protocol will be deployed in two types of ATM-attached equipment:

- *ATM Network Interface Cards (NICs):* ATM NICs will implement the LANE protocol and interface to the ATM network, but will present the current LAN service interface to the attached end system. Hence, as far as the network-layer protocols on the end system are concerned, end systems will believe that they continue to communicate on a known LAN, using known procedures; they will, however, be able to enjoy the vastly greater bandwidth of ATM networks.

  This is the type of emulation that may well be used by video servers that are ATM-attached. In particular, these servers may be large servers such as Silicon Graphics, Hewlett-Packard, or other large-scale servers, they may be workstation-based, or they may be high-end PCs. In any case, the NIC card drivers for these servers will have LANE client drivers because they will be clients of the LAN Emulation services offered by the network. They may also incorporate the LANE Server (LES) and LANE Broadcast and Unknown Server (LEBUS) to eliminate bottlenecks in video multicast service using LANE.

  In some cases where LANE drivers are not available for these servers, the network manager must manually configure a PVC between the server and any devices, such as an ATM-attached LAN switch or router, or use Application Program Interface (API) in the server that makes calls to the SVC driver software.

- *Internetworking and LAN switching equipment:* The second class of network gear that will implement LANE is ATM-attached LAN switches and routers. Together with directly attached ATM hosts equipped with ATM NICs, these devices are used to provide a "virtual LAN" (VLAN) service, where ports on the LAN switches are assigned to particular LANs independent of physical location. LAN Emulation is a particularly good fit to the first generation of LAN switches that effectively act as fast multiport bridges; LANE is effectively a protocol for bridging across ATM. Internetworking equipment such as routers also implement LANE to allow for VLAN internetworking.

  The basic function of the LANE protocol is to resolve MAC addresses into ATM addresses. By doing so, it actually implements a protocol for MAC bridging on ATM; hence the close fit with current LAN switches. The goal of LANE is to perform address mappings so that LANE end systems can set up direct connections between themselves and forward data.

With LAN emulation, devices internetworking over an ATM network gain the advantages of ATM bandwidth, which is useful for multiple video streams. In an application where a server or set of servers are transmitting numerous video streams—such as 100 independent streams at 1.5 Mbps each—an ATM backbone of 155 Mbps is required.

The disadvantage of LANE is that it hides the qualities of service from higher-layer protocols that can make use of qualities of service, such as IPv6 and CO-IPX (connection-oriented IPX).

Native Layer 3 over ATM— Multiprotocol over ATM (MPOA)

The main rationale for using a native mode protocol as opposed to LANE was implied in the conclusion of the previous section. LANE deliberately hides ATM, so any network-layer protocol operating over ATM cannot gain access to the QoS properties of ATM and must, therefore, use Available Bit Rate (ABR) or Unspecified Bit Rate (UBR) connections only. At the moment, this is not a major restriction. This is because all current network protocols were developed for use over existing LAN and WAN technologies, none of which can deliver a guaranteed QoS. Consequently, no existing protocol has any ability to request a specific QoS from the network or to deliver such to a higher-layer protocol or application. In turn, most network applications today do not expect to receive any guaranteed QoS from the underlying network protocol.

Therefore, at best, all current network-layer protocols today both expect and deliver only a “best effort” service, which is precisely the service that the ABR service was designed to offer. Similar to the way LANE adapts ATM’s connection-oriented nature to offer the same type of connectionless service expected by network-layer protocols, ABR “hides” the guaranteed QoS features of ATM to offer the best-effort service expected by these protocols. As such, ABR and LANE perfectly complement each other.

In a video application that was designed to run over LANs, including high-bandwidth LANs such as 100BaseT or ATM, this is not a severe limitation. The network infrastructure itself can be designed to provide sufficient bandwidth for video applications so that in many cases, QoS guarantees are not required because there are minimal collisions. For example, a 1.5-Mbps video stream running along a dedicated Ethernet connection to a workstation will not suffer degradation if that end station is not simultaneously engaged in file transfers. For applications such as video server training “tapes” or “LAN-TV,” the user is not likely to engage in other intensive applications at the same time. Moreover, only 1.5 Mbps of the available 10 Mbps are being used for the video transmission.

In the future, however, this situation is unlikely to endure. In the first instance, as ATM networks proliferate, it is likely that demand will grow to be able to use their QoS benefits, because this is one of ATM’s major selling points. Independent of ATM, moreover, considerable work is being done on building a networking infrastructure capable of supporting a wholly new class of multimedia applications that combine voice, video, image, and intensive data traffic simultaneously, without need for precise control by the user or network manager. In order to support such applications, QoS guarantees are required from the network (for example, to minimize jitter and latency for interactive video or voice applications). One way in which such applications could be built is by running the applications or transport protocols directly across ATM. The ATM Forum is working on developing models for an API for direct ATM access within operating systems, and more specifically, for MPEG2 video transport streams over ATM.

However, one of the principal functions of network-layer protocols is to offer universal connectivity and a uniform service interface to higher-layer protocols; in particular, to transport-layer protocols, independent of the nature of the underlying physical network. Correspondingly, the function of transport-layer protocols is to provide session control services (such as reliability) to applications, so that these can be built without being tied to a particular network type. Unless applications run over common network and transport protocol, interoperability between two applications running on two different networks (for example, ATM and a conventional network), would be difficult, if not impossible.

Hence, other than for a small class of applications that can only run on ATM (because they require more bandwidth than available from any other technology—for instance, studio-quality video processing that will be transported only to other ATM-attached devices), most multimedia applications will continue to be built upon enhancements of current network-layer protocols and will be deployed on a wide variety of high-speed networking technologies.

Performance Guarantees in Existing Networks

Resource Reservation Protocol (RSVP)

RSVP is a control protocol that will be used by applications within IP end systems to indicate to transmitting nodes the nature (such as bandwidth, jitter, or maximum burstiness) of the packet streams that they want to receive. Intermediate systems along the path from the source to the destination IP end systems will also interpret RSVP control packets in order to perform admission control (analogous to ATM call admission control) and allocate the resources required to support the requested traffic flows.

Such systems will maintain "soft-state" about such traffic flows, much as ATM switches maintain connection state, and will perform packet-level traffic policing, shaping, and scheduling the same way that ATM switches groom cell streams to provide the guaranteed QoS. RSVP can hence be thought of as providing very much the same traffic contract specification functions with respect to packet-level traffic flows that ATM UNI and NNI signaling play with respect to cell flows as illustrated in Figure 7.

Figure 7. Resource Reservation Protocol (RSVP)

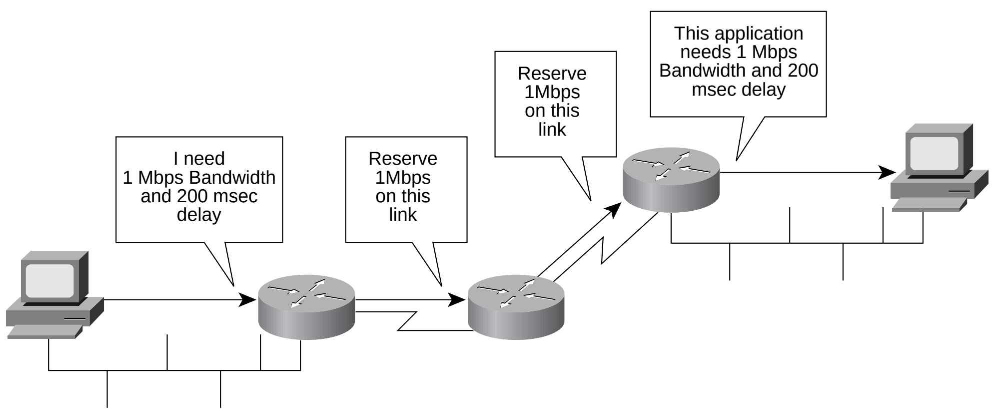

Protocol Independent Multicast (PIM)

RSVP is fundamentally built upon a multicast paradigm and routes traffic flows along source-routed point-to-multipoint paths (with unicast handled as a special case of multicast). New multicast protocols like Protocol-Independent Multicast (PIM), illustrated in Figure 8, and their associated packet routing protocols are closely coupled with RSVP, much as VC routing protocols are closely coupled with UNI and NNI signaling.

Figure 8. Protocol-Independent Multicast in routed networks

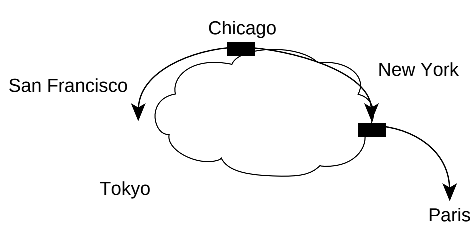

Such protocols rely upon the use of a flow specification that characterizes the expected traffic patterns for a stream of IP packets between two nodes and that the network can employ, using packet-level policing, shaping, and scheduling mechanisms, to deliver a requested QoS. In other words, a flow can be thought of as a layer 3 connection, because it identifies and characterizes a stream of packets between two or more nodes even though the protocol remains ostensibly connectionless.

The IP Next Generation (IPng) protocol, or IPv6, which the IETF is currently developing, incorporates support for a flow ID within the packet header that the network can use to identify flows, much as VPI/VCI identifies streams of ATM cells. Protocols like RSVP will be used to associate with each flow a flowspec that characterizes the traffic parameters of the flow, much as the ATM traffic contract is associated with an ATM connection. IPng will incorporate full support for integrated services through the use of those mechanisms and the definition of protocols like RSVP. This type of support might also be extended back to the current IP v4 protocol. It is likely that IPng and the other protocol components of the Integrated Service Internet will be fully standardized by the end of 1995; however, components may be deployed even earlier. As and when such protocols are widely deployed and applications are developed to use them, there will be a demand to run these protocols in native mode over ATM. It would be pointless to obtain QoS support from the network layer, only to have LANE preclude that support from being mapped to their equivalents in the ATM network. There is a very clear and natural mapping between the concepts and mechanisms of the Integrated Services Internet and ATM (flowIDs and flowspecs to ATM connections and traffic contracts, respectively).

Thus the latter can be thought of as eventually providing the packet-level control infrastructure for the physical network infrastructure of ATM where the former provides application services and the latter realizes the requested QoS guarantees. Thus, the true value of ATM can be exploited while preserving a network-independent service infrastructure for application portability. In order to realize this vision, however, there must be native-mode protocol support over ATM, not to be confused with native MPEG2 over ATM.

Native ATM Support for Video Applications

MPEG2 Convergence to ATM

In the cable TV industry, where potentially many hundreds of video streams require transport, and in studio video processing and post-production that require bandwidth unavailable elsewhere, internetworking with non-ATM devices is not a consideration. However, with the emerging development of set-top boxes that understand MPEG2 transport streams, a ATM infrastructure can be used for the MPEG2 transport stream. Set-top boxes by companies such as General Instrument, Scientific Atlanta, and others will attach to so-called “hybrid fiber-coax” networks, whose backbone between headends consists of fiber carrying MPEG2 over ATM, and whose connections to private homes are traditional coaxial cable carrying an RF signal.

In those cases, video riding directly over ATM is possible. Moreover, where bandwidth efficiencies are essential—as in the transport of hundreds of streams across wide areas through fiber—minimizing protocol overhead is critical.

For such applications, the ATM Forum’s Service Aspects and Applications (SAA) subworking group has voted on a mechanism for supporting MPEG 2 over AAL5. This mechanism is enabled via an Audio-Visual Service-Specific Convergence Sublayer (AVSSCS) in its nascent development stage. The sublayer will also be designed to support other video compression and coding protocols.

The ATM Forum’s SAA group voted on AAL5 as the adaptation layer for the AVSSCS for two reasons. First, AAL1 provides a timestamp—the Synchronous Residual Time Stamp (SRTS)—and as such is used for CBR connections. While a timestamp is necessary for recovering timing information, MPEG2 has its own Program Clock Reference (PCR) in its transport stream. It may be redundant to provide timing information in the packet stream at the MPEG2 layer as well as the ATM layer in the cell stream.

In addition, the existence of SRTS and other additional overhead in the AAL1 adaptation layer makes AAL5 a better choice. It has lower overhead and the full 48 bytes per cell are usable, so multiples of these cells coincide neatly with the size of MPEG2 packets. Moreover, with widespread vendor chipset support for AAL5, deployment can proceed at a lower cost.

Real-Time VBR for Video Applications

In addition, the creation of a real-time VBR traffic class is under development for UNI 4.0 in recognition of bursty video traffic that requires a real-time component—control over latency and jitter.

While it was initially assumed that a particular traffic type, such as CBR or ABR, would be tied to a particular set of Quality of Service (QoS) guarantees, the latest work of the ATM Forum enables different traffic types to request a range of QoS guarantees according to Table 7.[11]

ATM supports both latency and jitter control by explicitly allowing the end user to negotiate a contract with the network that guarantees the level of latency and jitter. In addition, the allowable cell loss and cell error are negotiable, bringing to the user an acceptable level of cell loss and error.

11. ATM Forum document 94-0730R2 *QoS Baseline Document* 1994. Note the UNI 4.0 will also allow for leaf-initiated joins—the ability for an end station to request to be added to a multipoint connection. This feature is useful where video broadcasts or conferences are in progress and an additional user wants to participate in the video conversation. As with any application using any feature of ATM, a mechanism will be required, whether through MPOA or an ATM API, for the application to use the ATM UNI 4.0 leaf-initiated join.

Table 7. ATM Quality of Service Parameters

Feature	Constant Bit Rate (CBR)	Real-time Variable Bit Rate (VBR)	Non-Real-Time VBR	Available Bit Rate (ABR)	Unspecified Bit Rate (UBR)
Cell Delay Variation (Jitter)	Specifiable	Specifiable	Not specifiable	Not specifiable	Not specifiable
Max Cell Transit Delay (Latency)	Specifiable	Specifiable	For further study	Not specifiable	Not specifiable
Cell Loss Ratio (% dropped)	Specifiable	Specifiable	Specifiable	Specifiable	Not specifiable
Cell Error Ratio (% erred)	Specifiable	Specifiable	Specifiable	Specifiable	Not specifiable

Generally, video traffic is more tolerant of cell loss and cell errors because viewers do not notice the loss of a few bits; the video decompression can often compensate for such errors. In extreme cases, lost picture frames simply drop out until the next frame arrives one-thirtieth of a second later, in the case of 30-frame-per-second transmission. The implication is that unlike data packets that are lost or erred, video is never retransmitted to avoid mixing up frame sequences.

ATM Adaptation Layers for Video Traffic

AAL5 will be used for VBR and ABR traffic to deliver bursty data as LANs do today. AAL1 will be used for CBR traffic where a timestamp is needed for precise transmission clockings.

Figure 9. Video over ATM Protocol Stack

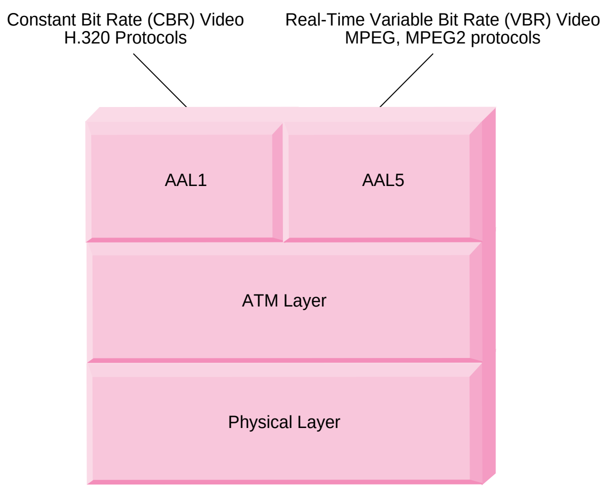

The ATM adaptation for compressed video and for isochronous services such as CBR video or voice appears as shown in Figure 9. As discussed earlier, packetized video will likely ride on top of data-link or network layers on top of ATM to allow internetworking with legacy LANs and protocols, and those will be adapted to ATM with AAL5.

Multicasting in ATM

ATM networks inherently support multicasting, where a single source can set up connections to multiple users. This multicasting is designed to be able to copy cells down the multicast tree only at the last possible branch point, so that multiple individual connections are not set up through the network to each user. This conserves network bandwidth.

ATM also supports near video on demand (NVOD) applications where users can decide whether to view one of a number of video streams in progress or to wait until the next viewing arrives (such as in five or fifteen minutes). The ATM Forum is working on mechanisms to allow leaf-initiated joins—the process by which an end node joins into a multicast of its own volition. Currently, a leaf (end user or viewer) becomes part of a multicast tree only at the time the multicast circuit is set up. With leaf-initiated joins, a viewer who did not join a multicast tree during its setup can choose to join after that multicast circuit has been set up.

ATM Service Multiplexers

Because video streams can be compressed, they do not require the full dedicated bandwidth given to an ATM 155-Mbps (OC-3) port, even though MPEG or MPEG2 streams can require at least 1.5 Mbps each. To maximize the investment payoff in ATM switches, multiple input and output feeds can be multiplexed up to full ATM DS3 / E3 or even OC-3 speed, to be fed into the ATM switch. Such an approach allows fanout of the video ports, giving high port density per ATM connection without tying up switch ports for individual video connections. The benefits of ATM, however, are presented through the service multiplexer.

For proper support of various video applications, encoding in either AAL1 or AAL5 is important, in order to match the application—AAL1 for H.320-compliant videoconferencing equipment and AAL5 for most other compressed video streams.

Network Capabilities for Video Applications

Bandwidth Requirements

By now it should be readily apparent that ATM is not imperative for providing bandwidth requirements except in specific cases of high-performance visualization or where many video streams are shared on one link or backbone, such that the aggregate bandwidth required exceeds the capacity of traditional LANs. Figure 10 illustrates the bandwidth requirements for various applications, along with a classic user population curve that exists in many networks today.

Clearly, the bandwidth needs of many of these users can be met with existing infrastructure, which can be augmented by ATM where needed.

Performance Requirements

Table 8 summarizes how different networks support performance requirements for video. In some cases, such as internetworking between the Resource Reservation Protocol and ATM Qualities of Service, work is still in the earliest phases.

Table 8. Video Requirements and Network Support

Video Performance Requirement	Support in ATM	Support without ATM
Bandwidth	Scalable to many hundreds of Mbps	Up to 100 Mbps
Latency and Jitter Control	Quality of Service Requests	Resource Reservation Protocol
Constant or Variable Bit Rates (Packetized or CBR)	ATM Adaptation layers 1-5; LAN emulation, MPOA, circuit emulation	Packetized network separate but internetworking with n x 64 kbps network
Multicasting	Multicast switched or permanent virtual circuits	Protocol Independent Multicast

Figure 10. Bandwidth requirements and network capability

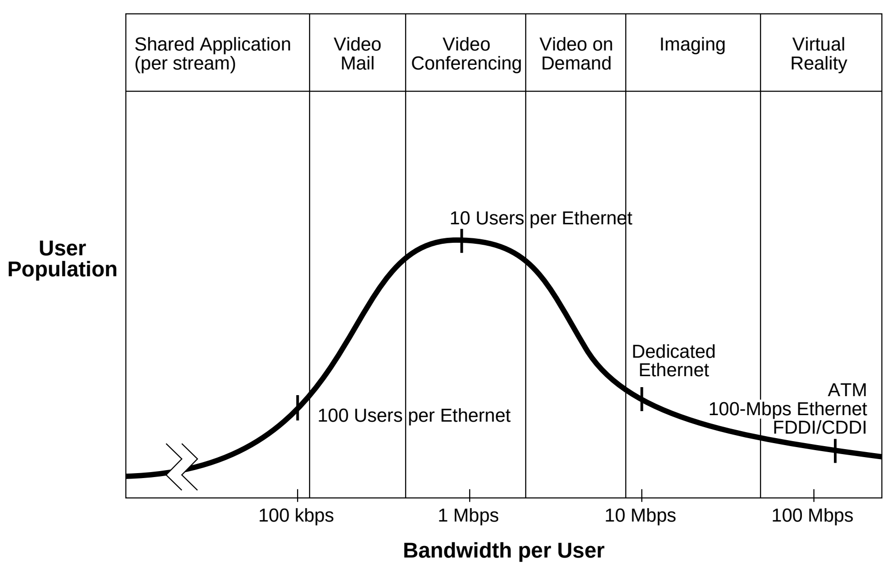

Video Application Types and Network Design

There are numerous video applications, and they range from more simple packetized, LAN-based video applications to Video on Demand (VOD) for thousands of users.

Packetized Video

Typical packetized video applications include inexpensive PC-based videoconferencing, as well as whiteboard-sharing, in which users simultaneously view and "mark up" an image. They also include corporate training, which can run over a LAN to multiple users who can view training videotapes that are digitally stored on a video server (see Figure 11). These applications usually run on LANs but could be transmitted over a WAN as needed.

Figure 11. Corporate Training via Networked Video

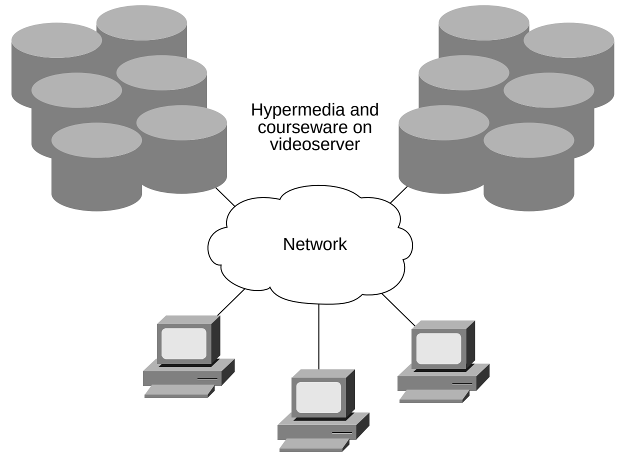

Typical packetized video applications are summarized in Table 9.

Table 9. Summary: Packetized Video

Video Codecs	Network	Network Protocols	Network Equipment
Numerous, including PC-based open and proprietary standards such as JPEG, MPEG, MPEG2, AVI, Indeo, etc.	• ATM • Switched LANs • Shared LANs	• Numerous layer 2 and layer 3 protocols • LAN Emulation over ATM • Multiprotocol over ATM	• LightStream 100 and 2020 ATM switches • Catalyst/ EtherSwitch LAN switches • Cisco routers

Distance Video

Videoconferencing is the traditional WAN, or distance, application that corporations employ to improve collaboration and communication while reducing travel costs. It provides the ability for several users to communicate visually; they see each other and share images such as drawings or presentations as illustrated in Figure 12. The application can include point-to-point connections or multiuser connections, and can span a campus LAN or a cross-country WAN.

Figure 12. Videoconferencing

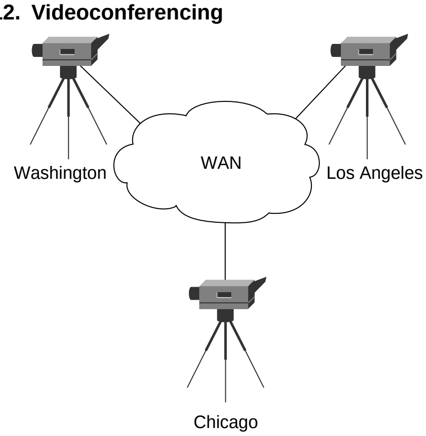

Hence, there are many levels of quality and cost for videoconferencing equipment, as well as transmission costs depending on the bandwidth used. While traditional videoconferencing equipment uses CBR video (the H.320 protocol suite), packetized video is increasingly viewed as a strong contender for such applications because of its high quality levels and potentially better WAN transport mechanisms via ATM when WAN ATM services are deployed.

Other wide-area video applications include distance learning or remote classroom applications, where an instructor in one location can teach classrooms of students in remote locations. Typically, these classrooms are equipped with two-way communication equipment so that the lesson is interactive (see Figure 13).

Figure 13. Distance Learning / Remote Classroom

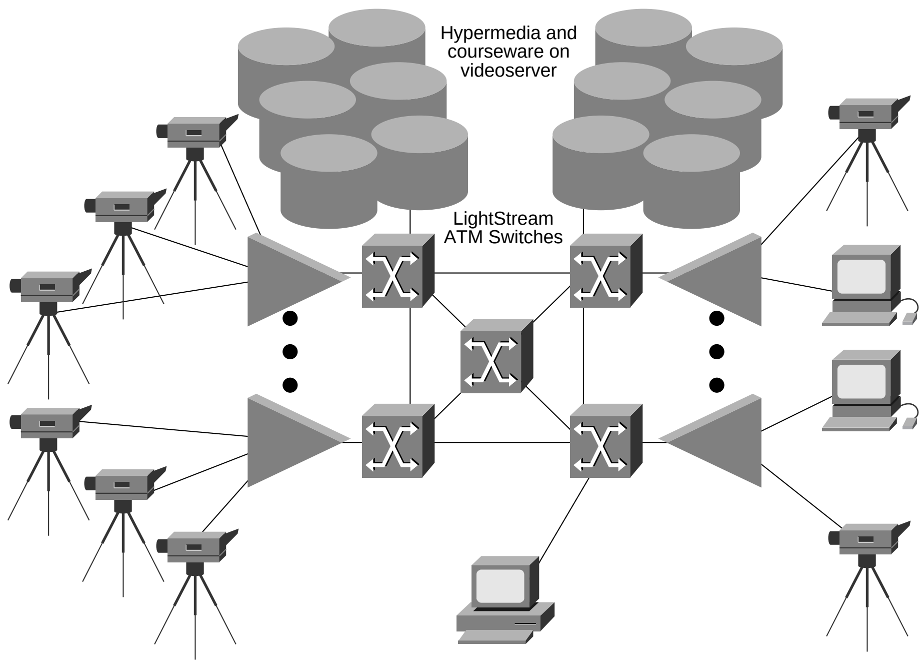

Table 10. Summary: Distance Video

Video Codecs	Network	Network Protocols	Network Equipment
Mostly MPEG, MPEG2 (packetize and H.320 (CBR)	• ATM • Minimal LAN use at edges • ISDN, n x 64 for legacy applications	• Circuit Emulation over ATM if CBR • MPEG2 over ATM if packetized • ISDN, leased lines for legacy CBR	• LightStream 100 and 2020 ATM switches • ATM multiplexers

Video on Demand (Cable TV/Telco TV)

Video on demand (VOD) is the video transmission typically provided by cable companies today, but it is of strong interest to the telephone companies (telcos), some of whom are beginning to enter this market. It requires a high degree of image and audio quality, because movie consumers expect to get high-quality service. In addition, VOD requires the ability to multiplex different streams of information, because different customers will want different movies at different times (see Figure 14)

In order to meet transmission cost-effectiveness requirements, VOD may be deployed as Near Video On Demand (NVOD) instead of True Video On Demand (TVOD). TVOD allows a user to request any movie at any time, while NVOD transmits the same movie at predefined but frequent intervals (say, every 15 minutes) so that home viewers can choose times that are comfortably close to the times they actually prefer to begin viewing the movies.

Figure 14. Video on Demand

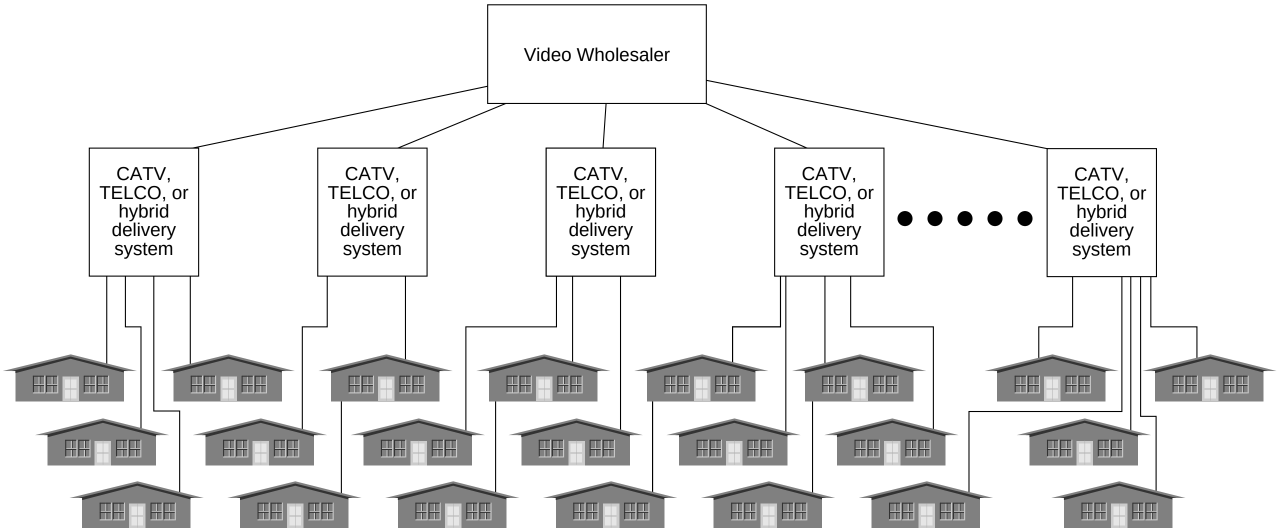

Table 11. Summary: Video on Demand

Video Codecs	Network	Network Protocols	Network Equipment
• Mostly MPEG, MPEG2, and MPEG4 (future)	• ATM	• MPEG2 over ATM	• LightStream 100 and 2020 ATM switches • Litton-FiberCom ATM multiplexers

Summary

The growth of videoconferencing, video on demand, and other applications is just beginning to move out of the early-adopter phase. As these applications and newer ones, such as visual simulation, virtual reality, and HDTV or SHD[12] transmission are developed, the demands on networks will accelerate.

Clearly there are numerous video applications, but regardless of the application, networks must provide for:

- Sufficient bandwidth
- Latency and jitter control
- Provision for VBR or CBR traffic, as needed
- Efficient multicasting

Since these networks will probably include ATM and other infrastructures like switched LANs and WAN circuit switching, these infrastructures must all interoperate. Companies that face numerous choices of video applications can evolve their network infrastructures to ATM, recognizing that they can continue to leverage their existing infrastructures while moving to the world of cheaper and higher-quality video applications running over networks of increasingly higher bandwidth using ATM.

As new applications are developed, the network infrastructure must and will evolve to ensure interoperability and cost-effective deployment of these applications that promise to provide cost savings and increased differentiation for the companies using them.

12. Super High Definition television, which includes HDTV as a subset

Appendix A: MPEG2

MPEG2 is rapidly becoming an industry standard for image transmission. MPEG2 is an open standard and is governed by an ISO subcommittee. It can be used by numerous devices, including encoders that are card or PC based and inexpensive decoders such as PCs, cards, or set-top boxes for TV.

MPEG2 compresses images spatially and temporally. Spatial compression is an algorithm by which a particular frame is compressed in size through a discrete cosine transform or DCT (see Figure 15).

Figure 15. Spatial Compression

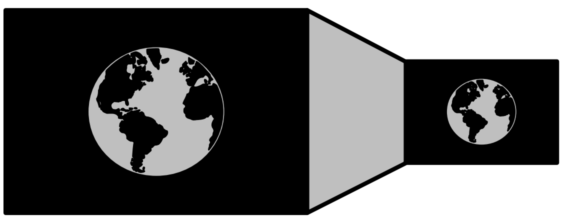

Essentially, a DCT reduces the number of bits needed to describe a still image or frame by eliminating redundant information from the frame.

Other algorithms, such as JPEG for still images, also use this same basic algorithm for spatial compression. However, JPEG is used only for still frames, while motion-JPEG is a proprietary method for applying JPEG to all the frames inside a movie clip.

MPEG and MPEG2 are the open standards designed for motion pictures. They compress the still frames and then eliminate redundancy between frames using a temporal (time) compression algorithm that consists of three types of frames:

- I-frames, or intra-frames, which are complete (spatially compressed) frames
- P-frames, or predicted frames, which are interpolated as the difference between the current frame and the last I-frame
- B-frames, or bidirectional frames, which are interpolated between these other frames.

To apply the full temporal compression, an MPEG2 coder buffers a set of a few frames at a time while it looks forward and backward in the frame set to do this interframe interpolation (see Figure 16). This allows real-time compression, but with a small degree of latency to allow for buffering and compression.

Figure 16. MPEG2 temporal compression

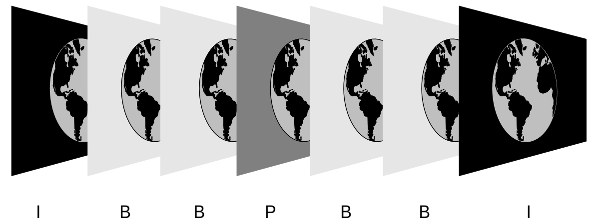

MPEG2 creates streams of compressed video with timestamps. This timestamp is provided inside an MPEG or MPEG2 stream, so that the timing and ordering of the I, B, and P frames is correct. Note that because the MPEG layer itself has a timestamp (Program Clock Reference, or PCR), it does not depend on the network layer to provide a timestamp for it. That is one reason why the ATM Forum decided that AAL5 would be used for carrying MPEG2 streams rather than AAL1, which has a synchronous residual timestamp (SRTS). AAL5 is a lower-overhead adaptation layer and can better transport MPEG2 packets.

Latency, Bandwidth, and Quality Control with MPEG2

MPEG2 requires lengthier encoding time than decoding time because of the bidirectional search and compression in a given frame sequence. The decompression is faster, because the decoder only reconstructs information found in the compressed I, P, and B frames. Note that much of the compression time on the encoder is given to the B frames—the bidirectional frames that require interpolation both forward and backward in time.

Hence it is possible to reduce the encode and decode time by eliminating the use of B frames. Consequently, the bandwidth required to transmit a similar-quality image will be higher, but the ability to control whether that MPEG2 stream will transmit I frames only, I and P frames, or I, P, and B frames allows some degree of control over the encode/decode latency (useful in two-way communication scenarios) and the bandwidth and image quality.

Appendix B: Performance Issues for High-End Video over ATM

ATM has been defined as a networking technology to support data, voice, and video transmission. In particular, ATM is being developed to support numerous high-end applications such as high-end video, which includes the emerging MPEG and MPEG2 video standards.[13] While the predominant standard for videoconferencing today is the circuit-switched H.320 protocol suite,[14] one key reason for the adoption of ATM is its ability to scale to higher bandwidth and to provide control for latency and jitter in the network and enable numerous emerging video applications.

The requirements for high-end video applications include bandwidth and delay control. Delay control, including both end-to-end latency and jitter control, becomes important for high-quality video transmission. It is important to recognize, however, precisely how much latency and jitter is tolerable for high-end video.

Carrier ATM Latency and Jitter Requirements to Support Video

A significant level of research has been conducted in this field to determine these levels of latency and jitter. Based upon the extensive research in this field, carriers are beginning to define the requirements for latency and jitter on a per-switch level to support high-end video. These requirements are as follows to support two-way high-end video:[15]

- A cell-loss ratio across the network of less than 1.7 x 10-9
- A cell-transfer delay (99 percentile) across the network of 4 milliseconds across the network and 150 microseconds per switch[16]
- A cell-delay variation (99 percentile) across the network (not per switch) of 500 microseconds

Note that for the next several years of MPEG and MPEG2 equipment evolution, the vast preponderance of delay will introduced in the video compression equipment, which is on the order of hundreds of milliseconds. Therefore, it is not likely that full I, B, and P frame MPEG or MPEG2 formats will be used for two-way communications in the near-term; rather, they will be used for store-and-forward applications such as video broadcasting in the wide area or LAN-video such as corporate training. In such applications, the delay can be much greater than the specifications noted above.

ATM switches which meet or exceed these latency and jitter requirements are capable of supporting high-end video applications in the wide area. Given that local-area (campus) transmission introduces lower transmission delay and allows for greater control over both video encoding and decoding equipment as well as control over the network itself, it is clear that local-area ATM switches can even exceed the above requirements and allow for high-end video applications.[17]

General Latency and Jitter Requirements for High-End Video

In particular, two-way communications requires tighter control of latency than one-way transmission, which can be store-and-forward.[18] Such requirements stem from human tolerance for delays and pauses in communications. Even for two-way high-end video, however, average latency and jitter requirements are in the range of milliseconds, and not microseconds, while the maximum tolerance of delay can be double that of the average tolerance (for example,

13. Defined by ISO/IEC JTC1/SC2/WG11, also known as the ISO Motion Pictures Experts Group
14. The specific standard for video coding is ITU-T/CCITT H.261, which defines coding of moving pictures for p x 64 kbps lines.
15. Keinath, R. and D. Minoli, *Distributed Multimedia Through Broadband Communications*, Boston: Artech House, 1994. Note that these specifications are for public-carrier switches, which are deployed, by definition, for wide-area services, and hence have stricter requirements than for local-area switches. Such required specifications do not exist as such for local-area switches.
16. Cell-transfer delay is defined as the elapsed time from which the first bit of a cell traverses the ingress point of a switch, to the time when the last bit of the cell passes the egress point of the switch. The term “99 percentile” means that 99 percent of the cells will experience the stated delay or less.
17. It should be noted that Cisco’s LightStream ATM switches provide for latencies in the range of 30 microseconds and thus fall well within the range of switch requirements for supporting high-end video.
18. Keinath and Minoli, ibid.

MPEG can tolerate latency of 11 milliseconds for two-way sessions).[19] Table 12 illustrates the tolerance of high-end video for delay and jitter.[20]

Table 12. Delay and Jitter for Two-Way Video and Voice Applications

Application Type	Application	Average Delay Tolerance (msec)	Average Jitter Tolerance (msec)
Low-end	64-kbps video conferencing	300	130
	16-kbps compressed voice	30	130
High-end	1.5 Mbps MPEG NTSC video	5	6.5
	256-kbps MPEG voice	7	9.1
Extremely high-end	20-Mbps HDTV video	0.8	1

Since end-to-end delay of a few hundred milliseconds is acceptable in traditional two-way real-time video, ATM networks can easily meet the requirements.[21] For high-end, two-way video, the high transmission speeds of ATM networks also ensures that absolute increases in delay time are not critical. In fact, the encoding operation delay (compression delay) will itself be more significant than the delay of the ATM network and its component switches.[22]

Additional research regarding network requirements for high-performance multimedia indicate that jitter of milliseconds is also tolerable for these applications[23] (see Figure 17).

19. Russell, J. "Multimedia Networking Performance Requirements," *ATM Networks*, I. Viniotis and R. Onvural (editors) New York: Plenum Publications, 1993
20. Onvural, Raif "Asynchronous Transfer Mode Networks: Performance Issues" Boston: Artech House, 1994, p. 81
21. Lee, S. and Wu, L. "Variable Rate Video Transport in Broadband Transport Networks," *Proceedings of SPIE Conference on Visual Communications and Image Processing*, 1988, p. 955
22. Ohta, Naohisa *Packet Video: Modeling and Signal Processing* Boston: Artech House, 1994, p. 27

Figure 17. Jitter and Cell Loss Requirements for Voice and Video

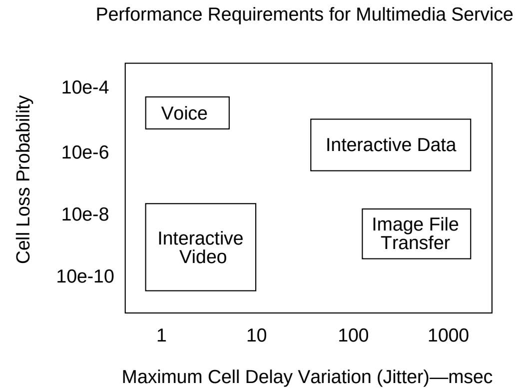

Above the ATM Layer: How High-End Video Deals with Delay & Jitter

Recognizing that there will be cases where delay and jitter are introduced into ATM networks, the higher layers—MPEG and MPEG2—are designed to deal with the introduction of delay and jitter. In particular, codecs have internal buffers that can dejitter the signal in case of small jitter introduction. Moreover, the existence of a clock signal in MPEG and MPEG2 streams enables recovery of timing in the event of network jitter.[24] Such clock recovery is controlled with a phase-locked loop (PLL) at the decoder side. Note, however, that in high-jitter scenarios, the PLL can lose phase lock.

Since *extreme* jitter and delay can effectively result in MPEG packet loss,[25] layers higher than the network layer must effectively ensure that the video stream recovers from such loss. Such protection and recovery techniques include structured packing of encoded video, where MPEG

23. Woodruff, G. and R. Kositpaiboon "Multimedia Traffic Management Principles for Guaranteed ATM Network Performance," IEEE Journal on Selected Areas of Communications, Vol 8, 1990, pp. 437-445
24. Such clock signals include the STC, or System Time Clock, generated at a 90 kHz reference time, which is inserted into a stream with a value ranging between 0 and 2^{33}-1. Such values in the stream are called PTSs, or Presentation Time Stamps. Given the range of values for the PTSs, unique identification of the information stream is possible over a 24-hour range.
25. Ohta, p. 27

macroblocks[26] are assigned addresses.[27] If something prevents a packet from proper reception by the decoder, synchronization can still be maintained by discarding all data until the next macroblock address is recognized.[28]

There are numerous higher-layer mechanisms for protection and recovery of video streams. These include error correction, packet priority protection, structured packing, leaky prediction encoding, insertion, cyclic refresh, lapped orthogonal transform, concealment, and coordinated operation of coder and decoder. For a more complete treatment of these mechanisms, see Naohisa Ohta's text on Packet Video.[29]

While some of these mechanisms already exist in MPEG and MPEG2 specifications, additional mechanisms may well become part of new video standards now evolving, such as digital HDTV. Implementation layers of protection and recovery mechanisms are illustrated with respect to the ATM protocol stack as shown in Figure 18.[30]

Figure 18. Future Methods for Video Application Performance Controls over ATM

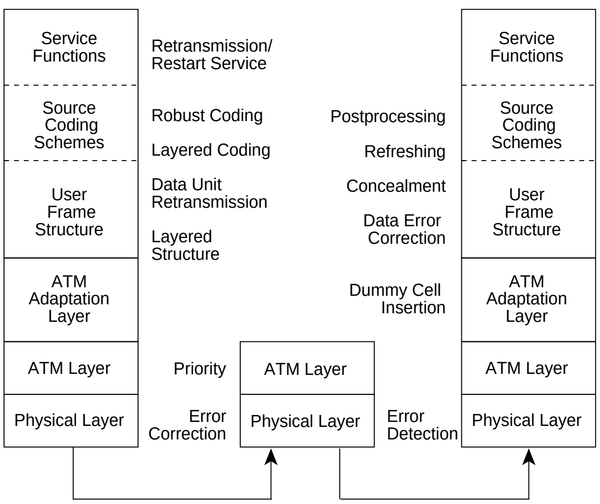

Higher-Layer Video Application Performance Control

It should be clear that the performance issues of delay and jitter in ATM networks are both understood and are being designed into specifications for ATM switches and networks. Moreover, as higher-quality video applications emerge such as SHD, (Super High Definition television or images)[31] emerging methods and algorithms for compensating for undue delay and jitter will likely be introduced into emerging video standards that require higher performance.

Bandwidth and Burstiness

One reason why the ATM Forum created the real-time VBR specification[32] is to enable transmission of variable bit rate video such as MPEG and MPEG2. The bandwidth requirements for compressed video images is shown in Table 13.[33]

Table 13. Bandwidth Requirements for Moving Pictures

Standard/ Format	Bandwidth	Compression Ratio[1]
Motion JPEG	10-20 Mbps	7-27:1
MPEG-1	1.2-2.0 Mbps[2]	100:1
H.261	64 kbps-2 Mbps	24:1
DVI	1.2-1.5 Mbps[2]	160:1
CDI	1.2-1.5 Mbps	100:1
MPEG2	4-60 Mbps[3]	30-100:1
CCIR 723	32-45 Mbps	3-5:1
CCIR 601 / D-1	140-270 Mbps	Reference
U.S. commercial systems using "mild compression"	45 Mbps	3-5:1
Vendor methods (e.g., PictureTel SG3)	0.1-1.5 Mbps	100:1
Software compression (small windows)	Approx. 2 Mbps	6:1

1. *Compared to broadcast quality*
2. *Baseline standard; other rates also possible*
3. *Image quality becomes asymptotic above 8 to 10 Mbps*

26. A macroblock is the basic unit for motion compensation, and consists of 16 x 16 pixels.
27. LeGall, D. "MPEG: A Video Compression Standard for Multimedia Applications," *Communications of the ACM*, Vol. 34, No. 4, April 1, 1991
28. Ohta, p. 150
29. Ibid., Chapter 6
30. Ibid, p. 142
31. This emerging video format includes digital HDTV, or definition High-Definition TV, as a subset.
32. ATM Forum document 94-0730R2 *QoS Baseline Document* 1994
33. Keinath, p. 188

The variable-bit-rate results from the encoding scheme—as MPEG or H.261 methods for compression are used, the bit rate varies, depending on the level of motion between frames. The level of burstiness for moving pictures has been characterized for several video types as shown in Table 14.[34]

Table 14. Burstiness of Video Applications

Video Type	Bandwidth Peak-to-Average Ratio
Studio-quality video	1.9
Broadcast-quality TV	2.7
Videoconference	3.1
Video telephone	4.4

For different levels of video and data priority, bandwidth allocation and call admission control allows provisioning for different traffic types. Constant bit rate (CBR) traffic, for example, will require highest priority levels, while latency-insensitive data can be assigned lower priority. Real-time VBR traffic may be assigned higher or lower levels of priority, depending on the needs of the network users.

Given that there will likely be numerous services deployed by an ATM network, including data, video, and voice, it should be apparent by now that to ensure performance guarantees for different services, call admission control will be necessary in ATM networks.

Usage Parameter Controls

In particular, it is not possible to speak of large-scale ATM networks deployed for high-end video and ignore the requirement for usage parameter controls. The usage parameter control (UPC), or traffic policing function, ensures that incoming traffic does not exceed the contract negotiated between the user and the network. Without such a function, already-established connections may incur performance degradation from users who violate their contract.

34. Verbiest, W. and L. Pinno, "A Variable-Rate Video Codec for Asynchronous Transfer Mod-Networks," *IEEE Journal on Selected Areas in Com- munications*, Vol. 7, No. 5, June 1989, pp. 761-770.

The use of UPCs, therefore, is required for ATM networks that support video traffic.[35] The UPCs ensures that all traffic types conform to their contract and receive their requested level of performance from the network.

In addition, the use of codecs with buffers that provide some traffic shaping can reduce peak rates and enable more robust conformance to the UPC.[36] Allocating the correct amount of bandwidth for the type of video transmission ensures proper allocation of network facilities. Using both bandwidth allocation and UPCs[37] has been confirmed in the research literature to be an effective method for controlling VBR video traffic.[38]

Conclusion

Clearly, then, performance measures for ATM switches that provide transport for high-end video are well-understood and defined. As new video formats such as MPEG and MPEG2 become widely accepted in transmission networks, ATM will be ready to provide the guarantees for performance expected of applications using these new video formats. Moreover, additional control methods above the ATM layer will be used for optimal MPEG and MPEG2 service. Finally, if and when yet fully undefined super-high-end video such as SHD is standardized, additional controls above the ATM layer will be developed to provide for services built upon SHD.

35. Ohta, p. 189
36. Kawashima, M. and Tominaga, H. "A Study on VBR Video Transmission under the Usage Parameter Control" *5th International Workshop on Packet Video*, F3, March 1993, Berlin
37. In this study, parameters were chosen using Markov-chain processes. As the state of the art advances, one may assume that video traffic behavior will be better modeled and understood and network managers will be able to more easily apply bandwidth allocation and UPC parameter control for their traffic.
38. Heeke, H. "A Traffic Control Algorithm for ATM Networks" *IEEE Transactions on Circuits and systems for Video Technology*, Vol. 3, No. 3, June 1993

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
World Wide Web URL:
`http://www.cisco.com`
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Europe
European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic - Batiment L2
16 avenue du Quebec
BP 706 Villebon
91961 Courtaboeuf Cedex
France
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Austria
Cisco Systems Austria GmbH
World Trade Center
A-1300 Vienna Airport
Austria
Tel: 43 1 71110 6233
Fax: 43 1 71110 6017

Belgium
Cisco Systems Bruxelles
Complex Antares
71 avenue des Pleiades
1200 Brussels
Belgium
Tel: 32 2 778 42 00
Fax: 32 2 778 43 00

Denmark
Cisco Systems
Larsbjoernsstraede 3
Dk-1454 Copenhagen K
Denmark
Tel: 45 33 37 71 57
Fax: 45 33 37 71 53

Germany
Cisco Systems GmbH
Max-Planck-Strasse 7, 3rd Floor
85716 Unterschleissheim
Germany
Tel: 49 89 32 15070
Fax: 49 89 32 150710

Italy
Cisco Systems Italy Srl
Centro Direzionale Milano Oltre
Palazzo Raffaello Scala B 4P
Via Cassanese 224 20090
Segrate (MI)
Italy
Tel: 39 2 269 290 85
Fax: 39 2 269 290 06

The Netherlands
Cisco Systems
Stephensonweg 8
4207 HB Gorinchem
The Netherlands
Tel: 31 18 30 22988
Fax: 31 18 30 22404

Norway
Cisco Systems
Holmens Gate 4
N-0250 Oslo
Norway
Tel: 47 22 83 06 31
Fax: 47 22 83 22 12

South Africa
Cisco Systems South Africa
Prestige Business Center
Sloane Park 90 Grayston Drive
2152 Sandton
South Africa
Tel: 27 11 784 0414
Fax: 27 11 784 0519

Spain
Cisco Systems Spain
Avenida de Burgos
28036 Madrid
Spain
Tel: 34 1 383 2178
Fax: 34 1 383 8008

Sweden
Cisco Systems AB
Arstaangsvagen 13
11760 Stockholm
Sweden
Tel: 46 8 681 41 60
Fax: 46 8 19 04 24

Switzerland
Cisco Systems Switzerland
Grossrietstrasse 7
CH-8606 Naenikon/ZH
Switzerland
Tel: 41 1 905 20 50
Fax: 41 1 941 50 60

United Arab Emirates
Cisco Systems (Middle East)
Sheik Zayed Road
P.O. Box 26095
Dubai, UAE
Tel: 971 4 310 433
Fax: 971 4 313 681

United Kingdom
Cisco Systems Ltd.
4 New Square
Bedfont Lakes
Feltham, Middlesex TW14 8HA
UK
Tel: 44 1 81 818 1400
Fax: 44 1 81 893 2824

Intercontinental and
Latin American Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
Tel: 408 526-7660
Fax: 408 526-4646

Asia
Cisco Systems (HK) Ltd
Suite 1009, Great Eagle Centre
23 Harbour Road
Wanchai, Hong Kong
Tel: 852 2583 9110
Fax: 852 2824 9528

Cisco Systems (HK) Ltd
Beijing Office
Room 821/822,
Jing Guang Center
Hu Jia Lou, Chao Yang Qu
Beijing 100020 PRC
Tel: 86 10 501 8888 x821
Fax: 86 10 501 4530

Cisco Systems (HK) Ltd
New Delhi Liaison Office
Suite 1/9, Hyatt Regency Delhi
Bhikaiji Cama Place, Ring Road
New Delhi 110 066, India
Tel: 91 11 688 1234 x119
Fax: 91 11 611 7688

Cisco Systems, Indonesia
Level 12, Wisma Bank
Dharmala
Jl Jenderal Sudirman Kav. 28
Jakarta Selatan 12910, Indonesia
Tel: 62 21 523 9132
Fax: 62 21 523 9191

Cisco Systems Korea
Samik Rabidol Building
5th floor
720-2 Yuksam-2-dong,
Gangnam-ku
Seoul, 135-082, Korea
Tel: 82 2 3453 0850
Fax: 82 2 3453 0851

Cisco Systems (HK) Ltd
Kuala Lumpur Office
Level 5, Wisma Goldhill
67 Jalan Raja Chulan
50200 Kuala Lumpur, Malaysia
Tel: 60 3 236 5147
Fax: 60 3 236 5146

Cisco Systems (HK) Ltd
Singapore Office
Shell Tower, Level 37
50 Raffles Place
Singapore 0104
Tel: 65 320 8398
Fax: 65 320 8307

Cisco Systems (HK) Ltd
Taipei Office
4F, 25 Tunhua South Road,
Section 1
Taipei, Taiwan
Tel: 88 62 577 4352
Fax: 88 62 577 0248

Cisco Systems (HK) Ltd
Bangkok Office
23rd Floor, CP Tower
313 Silom Road
Bangkok 10500, Thailand
Tel: 66 2 231 8300
Fax: 66 2 231 8121

Argentina
Cisco Systems Argentina
Cerrito 1054, Piso 7
(1010) Buenos Aires, Argentina
Tel: 54 1 811 7526
Fax: 54 1 811 7495

Australia
Cisco Systems Australia Pty Ltd
Level 17
99 Walker Street
North Sydney NSW 2060
Australia
Tel: 61 2 9935 4200
Fax: 61 2 9957 4077

Brazil
Cisco Systems Do Brasil
Rua Helena 218, 10th Floor
Cj 1004-1005 Vila Olimpia
Sao Paulo, SP CEP 04552-050,
Brazil
Tel/Fax: 55 11 822-6095
Tel/Fax: 55 11 822-6396

Mexico
Cisco Systems de México,
S.A. de C.V.
Ave. Ejecito Nacional No. 926
3er Piso
Col. Polanco C.P. 11560
Mexico D.F.
Tel: 52 5 328 7600
Fax: 52 5 328 7699

New Zealand
Cisco Systems New Zealand
Level 16, ASB Bank Centre
135 Albert Street
P.O. Box 6624
Auckland, New Zealand
Tel: 64 9 358 3776
Fax: 64 9 358 4442

Japanese Headquarters
Nihon Cisco Systems K.K.
Seito Kaikan 4F
5, Sanbancho, Chiyoda-ku
Tokyo 102, Japan
Tel: 81 3 5211 2800
Fax: 81 3 5211 2810

North America
Canada
Cisco Systems Canada Limited
150 King Street West
Suite 1707
Toronto, Ontario M5H 1J9
Canada
Tel: 416 217-8000
Fax: 416 217-8099

United States
Central Operations
5800 Lombardo Center
Suite 152
Cleveland, OH 44131
Tel: 216 520-1720
Fax: 216 328-2102

Eastern Operations
1160 West Swedesford Road
Suite 100
Berwyn, PA 19312
Tel: 610 695-6000
Fax: 610 695-6006

Federal Operations
1875 Campus Commons Drive
Suite 305
Reston, VA 22091
Tel: 703 715-4000
Fax: 703 715-4004

Northeastern Operations
One Penn Plaza
Suite 3501
New York, NY 10119
Tel: 212 330-8500
Fax: 212 330-8505

Northern Operations
8009 34th Avenue South
Suite 1452
Bloomington, MN 55425
Tel: 612 851-8300
Fax: 612 851-8311

Service Provider Operations
(Telecommunications)
111 Deerwood Road
Suite 200
San Ramon, CA 94583
Tel: 510 855-4800
Fax: 510 855-4896

Southwestern Operations
14160 Dallas Parkway
Suite 400
Dallas, TX 75248
Tel: 214 774-3300
Fax: 214 774-3333

Western Operations
2755 Campus Drive
Suite 205
San Mateo, CA 94403
Tel: 415 377-5600
Fax: 415 377-5699

Cisco Systems has over 120 sales offices worldwide. To contact your local account representative, call the company's corporate headquarters (California, USA) at 408 526-4000 or in North America call 800 553-NETS (6387).

11/95