

How ATM Networks Meet Professional Broadcast Demands

By Jean Chatel, David Mouen Makoua, and Laurent Thebault

Through the combination of digital TV and digital networking technologies, the transport of professional broadcast services is undergoing a tremendous evolution. After a move from analog to digital, specialized leased lines are being replaced progressively by powerful general purpose asynchronous transfer mode (ATM)-based networks. The most critical issue for the adoption of this new technology by the broadcast industry is the transport of live realtime broadcast services, for which the most serious constraints appear. The level of performance that can be expected for this application is a key point for the end-user broadcaster. After reviewing various elements of the ATM technology, this paper focuses on performance aspects: the mechanisms put in place, leading to commitments by both the end-user broadcaster and the network operator to a given level of performance are explained; practical performance results are given; and examples of large-scale deployments are described.

The transport of professional broadcast services essentially includes contribution and primary distribution links, such as studio-to-studio, stadium-to-continuity suite, electronic newsgathering (ENG), national-to-local broadcast centers, studio-to-transmitter or up-link or head-end, etc. With the advent of digital, transport technology is undergoing tremendous change. Although some of those links continue to use analog signals at their interfaces, the vast majority of transmission networks now make use of digital technology. Digital transmissions offer better performance and better efficiency than analog, and a flexibility that analog transmissions cannot match.

Networking technology is also undergoing important change. Specialized leased lines are being replaced by powerful general purpose asynchronous transfer mode (ATM)-based networks. With optical fiber

cables and very high switching capacity being deployed in a number of regions of the world, multi-terabit transmission capacities are being offered, and a significant decrease in transmission costs is foreseen. Through permanent virtual circuits (PVCs), ATM technology offers the equivalent of conventional leased lines where appropriate, but also allows the flexible implementation of non-permanent links by means of switched virtual circuits (SVCs). Both types of service can coexist on the same ATM network, together with a wide variety of other services of different nature and requirements. From another point of view, such a network can be maintained as a whole entity, independent of the service, also contributing to improved efficiency and cost effectiveness.

Although non-realtime services can be transported through the same network, various strategies can be applied in each case to match user requirements. The key issue for the end-user broadcaster is definitely the transport of live realtime broadcast services, for which the most serious constraints appear. The level of performance that can be expected is a key point for the adoption of this technology. Realtime

broadcast services are most often transported using some form of compression like MPEG-2 or DV. Transit delay, jitter performance, robustness against errors or loss of information, and synchronization between receiver and emitter are typical parameters qualifying the related constraints.

ATM Technology and Broadcasters' Requirements

ATM networking technology offers attractive features for professional broadcasting applications, among others. The four major ones are:

- It is based on well-established standards developed within the ITU-T and the ATM forum. These widely adopted standards are the result of international cooperation and consensus.
- It has been designed to simultaneously transport audio, data, and video in the same condition as on separate networks.
- It adapts to a wide range of bit rates and works independently of the bit rate and the physical medium.
- It can be used the same way on local and wide area networks, enabling continuity of service between public and private networks.

ATM-based networks are now widespread and rapidly increasing. A large base of network service is becoming available, and the cost of transport for a video program should be comparable to that of videotape. In addition to availability, ATM networks offer the flexibility to broadcast applications: they can transport various types of applications in many forms and offer different classes of network service. For example, video files or video streams, intercom facilities, audio programs, A/V programs or multiplex of A/V programs, and other data can simultaneously share the same network access. Depending on user needs it will be possible to estab-

Presented at the 142nd SMPTE Technical Conference (paper no. 142-2), in Pasadena, CA, October 18-21, 2000. Jean Chatel, David Mouen Makoua, and Laurent Thebault are with Thomson multimedia Broadcast Solutions, F-35517 Cesson Sevigne Cedex, France. An unedited version of this paper appears in *Content Acquisition, Processing, and Distribution in the Digital Millennium*, SMPTE, 2000. Copyright © 2001 by SMPTE.

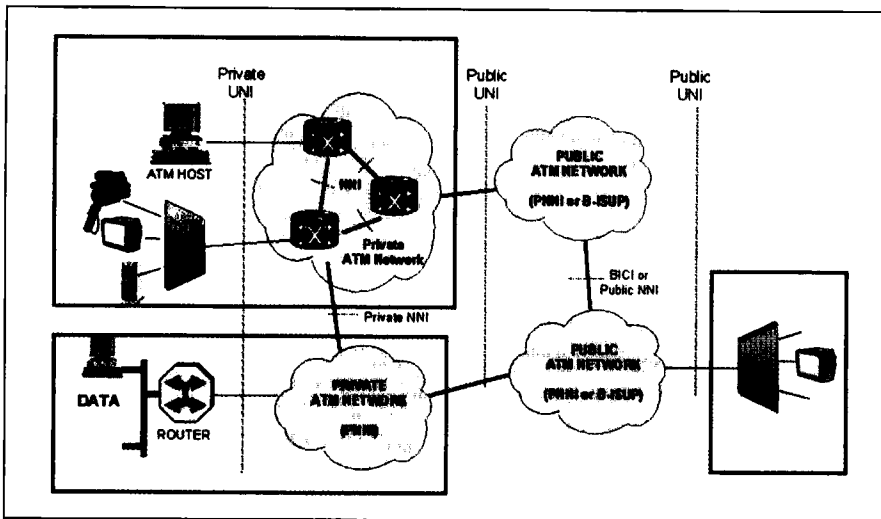


Figure 1. Basic ATM network architecture.

lish single-direction, full bi-directional, or asymmetrical connections.

Two forms of transport can be used: PVC for primary distribution towards broadcast headends; SVC for other type of communications, for example, occasional contribution links to relay sport events. Both forms of transport will co-exist on the same network and eventually share the same network access. For each type of connection, a predefined class of service will be applied. A standardized procedure is used at the connection setup to fix the traffic and quality of service (QoS) parameters—including bit rate and level of priority.

A key issue for broadcasters adopting this technology is the level of reliability and performance that can be expected from it. For the end-user a network must be transparent, in other words, it must be robust against errors and at least maintain the timing accuracy of input signals. Clock recovery must include solutions for removing jitter and wander at network output in order to match the requirements of video signal specifications, in particular, analog composite signals. In addition, the ATM network must offer the best safety and security conditions; these are easily achieved by the connection-oriented nature of the network.

Transit delay is another feature of the network. It depends on the transmission distance of course, but also on the bit rate. To be added to network transit delay, is compression delay, which also depends on the bit rate and

compression mode in use. Different trade-offs are possible to obtain a solution that best matches the requirement of an application.

Some Basics on the ATM Technology

ATM Network Architecture

An ATM backbone network is composed of interconnected ATM switches (Fig. 1). Various network topologies can be implemented. The more meshed the network, the more robust it is against blocking and congestion.

- Interfaces between the switches are called NNI (Network Node Interfaces).

- Interfaces between the switches composing the backbone and the network adaptors providing user access to the network are called UNI (User to Network Interfaces).

- ATM networks can be public or private. The interface between a private and a public ATM network is a Public UNI; it is based on ITU-T standards. An interface between public ATM networks is called B-ICI (B-ISDN Inter-Carrier Interface).

- Access to the network is performed by gateways called Network Adaptors (NA). Their role is to adapt a wide variety of user data flows to the ATM transport standard format.

The user-layered model of the backbone access network is quite simple (Fig. 2). An ATM network can work on a wide range of physical layers, which transport ATM cells between adjacent ATM layers. Bit rates on the physical layer range from less than 1 Mbit/sec to 2.5 Gbits/sec. Optical, electrical, and wireless media are used; PDH E1/T1, E3, E4, STM1-to-STM16 / OC3-to-OC48 are among the more frequent.

ATM Cells

On an ATM network, data are transported in the form of 53-byte ATM cells (Fig. 3). Each cell delivered to or from the physical layer includes a 48-byte payload and a 5-byte cell header. Header error correction (HEC) protects the header and is used for cell delineation. The most important part of the header is the 8-bit virtual path identifier (VPI) and the 16-bit VCI virtual channel identifier (VCI). These words are used to identify a connection in the network. In order to simplify setup, connections can be grouped in a VPI, in which case the network ignores the VCI for the switching operations. The VPI/VCI of a connection can be translated at each




Figure 2. Layer model of access and network backbone.




Figure 3. ATM cells and virtual connection structures.

crossed node. Since they are short and fixed-length, ATM cells are easily handled in network switches by high-speed dedicated hardware, a key element in the design of powerful switches and high-capacity networks.

ATM Adaptation Layer (AAL)

The flow of data to be transported over the network has to be adapted. This is the goal of the ATM adaptation layer (AAL) standardized by ITU-T. The most common adaptation layers in use in broadcast environments are known as AAL1 and AAL5.

AAL1 is used for the transport of time-critical application flows with constant source bit rate and delivery at the same bit rate. Each AAL1 cell payload is composed of 47 bytes of transported data, plus one cell to carry a sequence number. The sequence number is used to detect loss of cell and to maintain binary integrity. Additional timing information between source and destination can be added when the synchronous residual time stamp (SRTS) mechanism is used to synchronize source and destination. This method relies upon the network synchronization performance, which is not available on all networks.

A more generic synchronization mechanism consists of tuning the output clock frequency on the buffer occupancy level—the higher the number of cells, the higher the clock frequency. This adaptive clock recovery mechanism is the preferred solution as its performance is independent of network characteristics and network behavior. This mechanism allows performing real inter-networking.

With AAL1 a Reed-Solomon forward error correction (FEC) combined with interleaving of application bits provides robust protection against transmission errors, albeit at the price

of additional delay.

AAL5 has been targeted to transfer non-realtime data on high-quality links with a minimum of hardware/software processing. It can handle payloads of up to 64 Kbytes. Currently, AAL5 is mainly dedicated to the transport of IP data. It is also used with some MPEG-2 computer-based servers not capable of performing AAL1 encapsulation.

ATM Connections

ATM connections can involve a single source and a single destination (point-to-point) or a single source and multiple destinations (point-to-multipoint). This basic feature of ATM networks simplifies the implementation of distribution applications. ATM connections can be unidirectional, the most common for broadcast applications or bi-directional, in which case the bandwidth allocation can be symmetric or asymmetric.

There are two approaches to set up a connection. The first, called permanent virtual circuit (PVC), is a static route built into the network by the network manager. The NA must also be configured to send and/or receive cells from matching VPI/VCI and must respect an

implicit traffic contract. The second, called switched virtual circuit (SVC), is based on the use of ATM signaling and routing protocols. The ATM Forum UNI protocol is used in private networks, while the Q.2931 protocol specified by ITU-T is used in the public network. It works on a dial-up model: the calling NA sends a signaling message to its access switch; the message identifies the called NA by means of an ATM address and specifies the traffic contract parameters. This ATM access switch has a global knowledge of the network topology, of the available resources on each link, and is able to compute a path across the network that matches the request.

ATM Traffic Management and Quality-of-Service Setup

The goal of the Traffic Management toolset is to optimize use of network resources, match QoS requirements specific to each application, and guarantee that networks and users achieve the requested network performance.

Traffic Classes

ATM defines five prioritized traffic categories (Fig. 4):

1. CBR (constant bit rate) is suitable for traffic with high realtime constraints like broadcast audio or video, which requires a fixed bandwidth and supports small cell jitter.
2. rt-VBR (realtime variable bit rate) is dedicated to realtime traffic that can accommodate bandwidth variation and more jitter. Typical use is a combination, on the same connection, of voice traffic with silence suppression.




Figure 4. ATM traffic categories.

3. nrt-VBR (non-realtime variable bit rate) is for applications requiring fewer constraints on delay and jitter, but that are more sensitive to cell errors and cell losses. It is typically dedicated to ATM signaling flows or QoS IP data transport.

4. ABR (available bit rate) relies on a sophisticated mechanism that provides a fair share of the bandwidth left by the three previous categories among a set of active connections. Typical use may be the transport of IP data.

5. UBR (unspecified bandwidth) or "best effort" use the remaining bandwidth. Typical use is transport of IP data in a LAN environment.

Traffic Contract

A key element of the Traffic Management toolset is the mechanism to establish a traffic contract between the user terminal and the network when setting up a connection (Fig. 5). The terms of this contract have been standardized by the ITU-T and the ATM Forum. For each traffic category they define the network commitment, a QoS (delay, jitter, error rate), and the user commitment, a traffic profile (min/max bandwidth, jitter, burst size). Both must be fulfilled to meet a given level of performance.

Broadcast applications in a WAN environment typically use CBR and nrt-VBR. For CBR the network commitment is the cell loss ratio, cell delay transfer, and cell delay variation (cell jitter). The user commitment is

on the traffic injected in the network. It specifies the peak cell rate and the cell delay variation. For nrt-VBR the network commitment is the cell loss ratio and a mean cell transfer delay. The user commitment is the sustainable cell rate, peak cell rate and jitter, and maximum size of bursts at peak cell rate.

An important feature of ATM is that, whatever the traffic category, a connection request has to be processed and accepted by the network before any transfer of data. If the connection is established via an SVC, the contract is negotiated directly between the user and the network; if the connection is a PVC, the contract is indirectly negotiated by the network manager. This negotiation is handled by a dedicated function of the connecting switch (connection admission control [CAC]). This CAC checks all the parameters of the connection request and computes a path across the switch (PVC) or the network (SVC or S-PVC), taking into account the topology and available resources. If the requested resources are available, then the connection is agreed upon, and all along the connection path switches are configured to take the new connection into account. This includes reservation of bandwidth, cell queues, cell buffers, and scheduling capacity. Acceptance of the connection by the network is a guarantee that the QoS will be met.

To be sure that the user respects this contract, the network activates the usage parameter control (UPC), which

monitors and controls the traffic at user access for each connection. The UPC is also a standardized ATM process that is able to detect violation of negotiated parameters and capable of two types of action. The first is to drop the traffic in excess; the second is to tag the traffic in excess, but still accept it in the network, the tagged cells being the first dropped in case of congestion inside the network. Having this policy at ingress of the network ensures that the QoS of other already established connections is preserved.

Traffic Shaping

Traffic shaping is a way for the user to respect its contract. A shaper smoothes the traffic injected on a connection according to the profile requested in the contract. Shapers implement a simple or dual leaky bucket algorithm, the same standard algorithm as in the UPC network function. Traffic shaping regulates traffic and optimizes statistical multiplexing inside the network, enabling effective use of close to 100% of the bandwidth.

Congestion occurs when cells from numerous connections are competing, either at the network interfaces of NAs or at the output of ATM switches. Cell buffers and cell queues are required to solve this problem. Realtime flows do not support high delays. The buffering for related connections needs to be only few tens of cells. On the other hand, low QoS IP data are not sensitive to delays but request small cell loss. For this category a high amount of buffering is required, up to some 10k cells. In order to avoid blocking of realtime flows at the head of a line, a different queuing strategy for each category is required, as shown in Fig. 6. Having multiple queues requires a scheduler to extract a cell from a queue each time a time slot is available on the physical link. Many different scheduling algorithms exist at this point.

Main Functions of an ATM Network Adaptor

The ATM network adaptor (NA) is the gateway between a broadcast studio and a network that conforms the network performance to the end-user performance requirements. The following is an indication of the func-

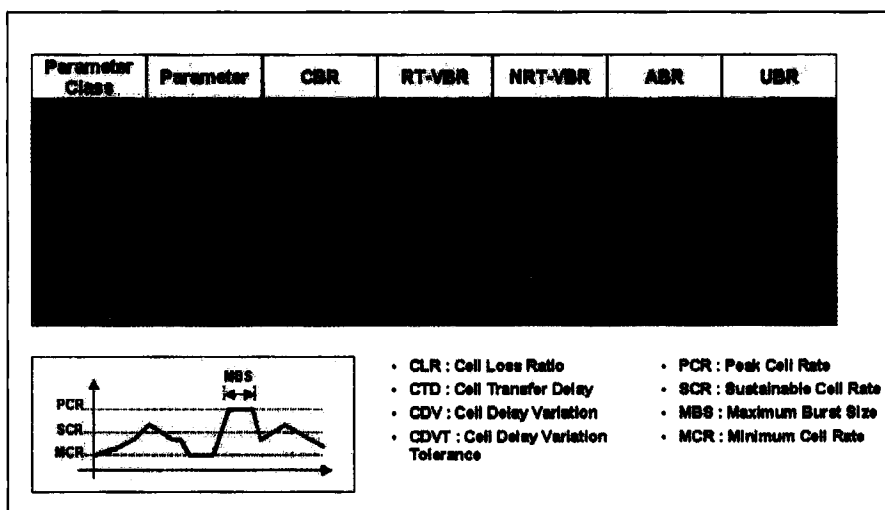


Figure 5. Parameters of a traffic contract.




Figure 6. Queuing mechanism.

tions to be offered by such a device to meet end-user expectations.

In order to solve a wide range of signal transport requests, the ATM NA must be an open platform that permits easy evolution and integration in various heterogeneous environments. It must support a variety of interfaces for the transport of multimedia data over ATM networks, including DVB-ASI and other emerging industry standards, such as SDTI. More traditional services like PBX or IP LAN interconnection should also be supported, to provide a global answer to the communication issues.

A wide pallet of network interfaces is available in order to cope with any

operational situation. The NA provides a modular architecture procuring maximum flexibility and easy integration of diverse extension capabilities as indicated in Table 1.

MPEG-2 Processing


For professional broadcast applications, an NA supports a set of independent high-quality MPEG-2 encoders and decoders, each connecting one video and several audio channels. It permits selection of a number of compression parameters to achieve the best compromise between the allocated bit rate and the rendered quality, and thus optimize telecom costs. These modules support both

MPEG-2 MP@ML, and 422P@ML profiles. The bit rate is user-selectable continuously between 1 and 50 Mbits/sec. All main MPEG-2 audio compression modes are supported. "Uncompressed AES/EBU" transmission mode, where 16 and 20-bit AES/EBU digital signal inputs to the encoder, may be passed-through inside MPEG-2 and restored on the decoder output in digital form.

The transparent of VBI line data, in accordance with relevant MPEG-2 standards, is also a key feature. This relieves the end user from the burden of supporting and interpreting a variety of VBI signals, which differ among countries and broadcasters. Different end-to-end delay modes should also be available, such as ultra-low delay (down to 150 ms) with a smallest impact on quality achieved at the price of high bit rate.

Uncompressed Audio and Video Transport

These modes are required when no signal modification is allowed or when end-to-end delay must be strictly minimized. This works at the cost of a higher bandwidth occupation.



AAL Processing

A good NA supports AAL1 with FEC for end-to-end error-free transport over quasi-error-free networks and AAL1 without FEC for delay sensitive application. For IP data transport and interoperability with some MPEG-2 servers, the NA also has to support AAL5.

Clock recovery is a very sensitive function that is absolutely required for the NA. Adaptive recovery is the preferred mechanism, as its performance is independent of the network.

ATM Multiplexing

In most situations, the NA handles multiple flows. This requires an efficient internal ATM multiplexer with per flow queuing and a sophisticated scheduler to guarantee each requested QoS for a mix of various classes of services.

Traffic Shaping

This is an important NA function performed by the scheduler. It makes sure that each individual flow shall be compliant with its traffic contract and that the addition of all flows respects a global rate limitation (Fig. 7).

Support of SVC (ATM Signaling)

This is a must for the dial-up approach of transmission over ATM. It is mainly used for contribution applications. Point-to-point and point-to-multipoint SVCs have to be supported. PVCs are required for distribution applications.

In-band Management

For remote control of the NA and to avoid an additional network to support it, the NA is able to exchange management flows via its ATM network interface. The connections dedicated to in-band management consume a very low bandwidth but are supported by a good QoS on the network.

Network Adaptor Performance

The main performance criterion for the transport of broadcast signals over an ATM network is transparency. As a key function the network adaptor must compensate for the few errors and timing impairments introduced by the digital transport. It must be empha-




Figure 7. Principle of traffic shaping.

sized that this function is common to any type of digital transport (non-ATM-leased line), with some specifics due to the asynchronous nature of the ATM transport mechanism.

Error Correction

Standard error correction of AAL1 will enable correction of four cell losses or two erroneous bytes every 128 cells. In practice this will offer “error-free” service while the network only offers “quasi-error-free” service. This mechanism is fully adapted to burst of errors because it makes use of a large interleaving, and has a demonstrated error-free service at bit error rate (BER) up to 10^{-6} , with geometric distribution of errors. This figure is similar to a worst-case BER occurring over multiple international hops and can also be compared to actual BER over typical national networks, which is commonly better than 10^{-8} .

The maximum burst of lost cells that can be corrected is 4 cells in every 128. Although this may appear to be few, recall that due to the strategies implemented by ATM switches, very high priority is given to realtime traffic, a position reinforced by a rate reservation. As a consequence, cell losses will mainly be caused by the binary errors for this type of traffic. This makes the transport of realtime services over ATM network appear to be extremely robust against errors appearing in practical situations when using AAL1.

Timing Accuracy

The key transparency requirement is, therefore, timing accuracy. The network high-frequency jitter can easily be removed. The critical timing parameter is the low-frequency wander.

At its output the network adaptor must deliver a very low drift composite video signal that meets PAL timing requirements, a set of specifications not even met by a number of video sources. Although this does not seem to be compatible with the asynchronous nature of ATM, state-of-the-art digital clock recovery devices enable this requirement to be met.

With a cell delay variation introduced by the network up to 1 ms—typical maximum value for a national network—a video signal with nominal +20 parts per billion/sec maximum drift rate and a frequency accuracy of +0.2 parts per million can be delivered. The typical time for the video signal to converge to these nominal values is less than 30 sec. While offering such performances in a composite PAL environment, the adaptor can also handle signals from inaccurate sources. The source frequency accuracy can be as poor as 30 parts per million and a large frequency drift is accepted.

Transit Delay

Due to the amount of buffering required for error correction and clock recovery, the end-to-end delay introduced by the ATM network essentially lies in the AAL1 processing layer. It

can be very large at low bit rate (0.1s at 1 Mbit/sec). Above 25 Mbits/sec, this delay is less than 15 ms (Fig. 8). For delay-sensitive applications, the AAL1 standard has stipulated another FEC principle with very low delay, and the technology to provide it is available. In the network itself the delay is relatively low, typically less than 5 to 10 ms for a connection over 1000 km.

Where applicable, the encoding and decoding delay has to be added to these network transport and processing delays. Delay introduced by MPEG-2 encoders depends very much on the compression rate. It can vary from a couple of video frames for intraframe encoding (typically 50 Mbit/sec application) up to above 600 ms when encoding with long GOPs.

Example of Large-Scale Deployment

The generic example shown in Fig. 9 is taken from actual situations. It illustrates the combination of possible, reliable, flexible contribution and distribution applications over the same ATM backbone and access network. In addition to broadcast applications, IP LAN and PBX interconnections are provided. A typical ATM network infrastructure used for broadcast applications is composed of a few tens of ATM switches, each of them connecting a few tens of network adaptors, depending on the customer distribution. These hundreds of network adaptors are spread over a region or a country.




Figure 9. Typical broadcast applications over an ATM network.

ments, each of these NAs handles 1 to 10 broadcast audio and video plus IP and circuit emulation interfaces.

Contribution

Thanks to ATM signaling over SVCs and MPEG-2 flexibility, users can, on demand, set up multiple point-to-point contribution connections for video/audio and intercom with any distant NA, on the basis of its address book content. For the operator, establishing a connection can be as simple as a click on an icon hiding more than 100 parameters concerning MPEG-2 encoding and ATM dial-up, QoS, and encapsulation. For specific occasions, experienced users themselves can tune each of the connection parameters.

The flexibility of ATM and MPEG-2 allows mixing rush transfers down to a few minutes, interview, or event

transmission and selecting the bit rate with regards to the expected picture and sound quality. Asymmetric bi-directional connections can be used, for example, for interviews providing high picture quality on the main path and a lower quality on the return path. If required, uncompressed audio and video transport can be transported, although at the expense of the higher bandwidth.

Additional side applications are also supported. They can include intercommunication to exchange orders with teams on-the-field, transparent transport of RS-232/422 signals for VTR remote control, and export of GPI for "on-air" request or return.

Distribution

Distribution of television bouquets and digital radio requires 24 hr/day, 7 day/week connections from a national head-end to regional/local cable head-ends and/or digital terrestrial transmitters.

For these applications the NA feeds permanent virtual circuit (PVC) ATM connections. In the backbone, "smart PVCs" provide the equivalent of conventional leased links, with additional capabilities such as multicasting and rerouting in case of link failure. As for contribution, the bit rate is chosen with regard to the expected service picture quality.

In order to comply with the network requirement to provide a 24 hr/day service, redundancy is implemented in the most sensible sites. Redundant interface boards at NA level, and




Figure 8. End-to-end delay introduced by AAL1 layers (input + output).

redundant NAs at site level, are used for that purpose.

Auxiliary Applications

IP LANs are interconnected via:

- Fractional E1 interfaces if an N x 64 kbit/sec link is sufficient (e.g. for management purpose).
- Ethernet interfaces when a few megabits per second and more flexibility are required.
- ATM tributary interface if fewer than 10 Mbits/sec and much more tuning are required, as for traffic classification for IP QoS.

PBX for telephony application are interconnected via fractional E1 interfaces.

Management

Each site, including one or a few NAs, is managed by a dedicated workstation. This local application is restricted to connection setup and monitoring. Configuring detailed hardware equipment and service and monitoring are globally performed in a network management center. This center is connected to all NAs via dedicated low-bandwidth ATM connections (in-band management).

Conclusion

The elements developed in this paper show that ATM networking technology, complemented by appropriate network adaptors, fulfills the need for, and performance requirements of, broadcast applications. It efficiently supports unicast and multicast applications, making it well suited for both contribution and distribution applications. Being switched-based and connection-oriented, it offers secure connections and the benefit of having to access only one network to carry all types of traffic.


Other key elements for the success of ATM networking for broadcast applications are the wide support of standards and industry forums to guarantee interoperability between various network operators and equipment vendors, and the flexibility and scalability of the technology, which supports both current and future cable plants and matches the requirements of current and future high bandwidth applications

and protocols. These conditions guarantee the long life of ATM network architectures.

Deployment of ATM networks, along with broadcast transport applications, is proceeding rapidly. It began a few years ago with PVCs being used for distribution applications, followed shortly thereafter by contribution applications. The current trend is large-scale utilization of SVCs for contribution applications. Networks covering large

regions or countries, typically including 300 to 400 broadcast-application-oriented network accesses, are currently starting operation. Most applications include a mix of contribution and distributions links on the same networks. The trend for the future seems to be toward the use of SVCs for any application. This mode of transfer, by becoming universal, would make network maintenance more rational and uniform.


THE AUTHORS



Jean Chatel



David Mouen Makoua



Laurent Thebault

Jean Chatel received an engineering degree from ESEO (Ecole Supérieure d'Electronique de l'Ouest) and a doctorate from the University of Rennes, France. He joined CCETT in 1974 to work on the early developments of digital television for studio applications.

For 15 years Chatel has been with Thomson, particularly in charge of development of realtime digital television and HDTV signals processing systems, including studio systems, motion compensated standards conversion, and compression systems. His current position is marketing manager, advanced technology, dealing with compression and networking technology. Chatel is a SMPTE Fellow.

David Mouen Makoua received a degree in electrical engineering from INPG/ENSERG France in 1984.

He spent ten years in Thomson-CSF as senior engineer and was involved in military networking

issues, mainly ATM equipment architecture.

Mouen joined Thomson Broadcast Systems in 1999 and is presently product manager for the network adaptor family, which addresses the delivery of high-quality audio/video over broadband telecom networks. He serves also as internal consultant for networking architecture issues.

Laurent Thebault received an engineering degree from Supelec (Ecole Supérieure d'Electricité, France) in 1989. He joined the Thomson group in 1990, where he currently manages the development of network adaptor and service multiplexer.

While at Thomson, Laurent has worked on the design of high-rate serial interfaces and switches for audio and video, and ATM-based multiplexing and transport systems. Currently he is involved in quality-of-service issues for the transport of high-quality video over emerging high-bandwidth IP networks.