



R&D White Paper

WHP 074

September 2003

The development of ATM network technology for live production infrastructure

C.J. Chambers

**The development of ATM network technology
for live production infrastructure**

C.J. Chambers

Abstract

The broadcasting infrastructure provided within production centres is placing greater reliance on the use of standard IT structured cabling to provide a cost effective and flexible method of delivering broadband data connectivity throughout the building. However, the provision of live broadcast and production facilities poses particular problems when attempting to distribute those in a similar way. This is primarily due to the requirements of these high bandwidth streams to pass between internal areas with extremely low latency and stability if such structures are to be considered suitable when replacing more traditional methods of connecting live audio and video signals. This paper explores the development of a suitable open standard network structure to meet the needs of live broadcast production systems.

This document was originally published in the Conference Publication of the International Broadcasting Convention (IBC 2003) Amsterdam 11th-15th September 2003

White Papers are distributed freely on request.
Authorisation of the Chief Scientist is required for
publication.

© BBC 2003. All rights reserved. Except as provided below, no part of this document may be reproduced in any material form (including photocopying or storing it in any medium by electronic means) without the prior written permission of BBC Research & Development except in accordance with the provisions of the (UK) Copyright, Designs and Patents Act 1988.

The BBC grants permission to individuals and organisations to make copies of the entire document (including this copyright notice) for their own internal use. No copies of this document may be published, distributed or made available to third parties whether by paper, electronic or other means without the BBC's prior written permission. Where necessary, third parties should be directed to the relevant page on BBC's website at <http://www.bbc.co.uk/rd/pubs/whp> for a copy of this document.

THE DEVELOPMENT OF ATM NETWORK TECHNOLOGY FOR LIVE PRODUCTION INFRASTRUCTURE

C. J. Chambers

BBC R&D, UK

ABSTRACT

The broadcasting infrastructure provided within production centres is placing greater reliance on the use of standard IT structured cabling to provide a cost effective and flexible method of delivering broadband data connectivity throughout the building. However, the provision of live broadcast and production facilities poses particular problems when attempting to distribute those in a similar way. This is primarily due to the requirements of these high bandwidth streams to pass between internal areas with extremely low latency and stability if such structures are to be considered suitable when replacing more traditional methods of connecting live audio and video signals. This paper explores the development of a suitable open standard network structure to meet the needs of live broadcast production systems.

INTRODUCTION

Asynchronous Transfer Mode (ATM) takes its roots from the evolution and development of switching and data transmission technology over that last 30 years. As high bandwidth transmission paths were digitised, one of the main driving forces in the effective utilisation of these paths was the ability to provide a well-managed mix of voice, video and increasingly data over these connections. ATM was developed to meet these needs by switching and multiplexing fixed length packets known as cells that carry a wide range of low to high bandwidth services simultaneously. As these services may well be carrying a mixture of voice, video and data it was clearly essential for ATM to be able to support the different transmission quality requirements that are important to ensure their successful passage and use through a common digital infrastructure.

The evolution of ATM has produced a thorough and extremely well thought through set of standards published by the International Telecommunication Union (ITU) which are recognised and used worldwide. It is these standards (see references 1 to 7) that can form the bases of open, high performance live programme infrastructure within production centres and the larger broadcasting organisations working along side other network protocols all running over common structured cabling and telecommunication interconnections. This enables significant opportunity to provide both live and file based production and broadcast structure with unprecedented scalability and flexibility while reducing the cost traditionally associated with such abilities.

1. THE BASICS OF ATM STRUCTURE

1.1 The ATM cell

ATM data is divided into a 53 byte cell structure of which 48 bytes are assigned to carry the payload data being transported across the network. The term cell is defined in the ITU-T Recommendation I.113 as “A cell is a block of fixed length. It is identified by a label at the ATM layer of the B-ISDN PRM”.





Figure 1: UNI Cell format

The ATM cell has two main formats, the User-to-Network Interface (UNI) format shown in figure 1 and the Network-to-Network Interface (NNI) format sometimes also referred to as the Network Node Interface shown in figure 2.





Figure 2: NNI Cell format

These two formats are slightly different being designed for use in differing parts of the structure. The UNI cell is designed to be used between the user or equipment interface and the first network switch. It contains a five-byte header consisting of a 4-bit field for Generic Flow Control, an 8-bit field for the Virtual Path Identifier, a 16-bit field for the Virtual Channel Identifier. This is followed by a 3-bit field containing the Payload Type Identifier, a single bit flag known as the Cell Loss Priority bit and finally a 8 bit field known as the Header Error Check providing error checking and in some cases correction for the header of the cell. The remainder of the cell contains 48 bytes of payload data being passed by this cell.

The NNI format is designed for use between ATM network switches and appears either as Public or Private versions (PNNI) depending on the public or private address support requirements of the path the cells are passing on the network.

The key difference between a UNI cell and an NNI cell is the removal of the GFC field and the use of these bits to increase the VPI field to 12 bits. This has the effect of greatly increasing the number of possible virtual paths on physical routes between ATM switches.

The important issue to remember with ATM cells is to think of their structure in a similar way to an athletic relay race. The runner is equivalent to the header and the baton is equivalent to the payload. As an ATM cell passes a switch the payload is handed to a new header just as the baton is handed to a new runner at points around the track. To continue the analogy, the athletic track can contain many runners all running at different speeds in a similar manner to the different services being carried on the network. In this way the payload is provided with appropriate header information and uses the appropriate header types as it passes through the ATM structure. The importance of the cell size will become clear later in this paper as the issue of passing services running at different speeds down a single physical path is explored.

1.2 ATM transport network

The ATM transport structure is shown in figure 3 and basically split into two layers.

In simplistic terms, the physical layer deals with:

- The transmission and framing of data over copper or fibre using standardised physical interfaces.
- Header Error Check head sequence generation and verification.
- Re-acquiring cell boundaries and header structure. (One example is by the use of Physical Layer Convergence Protocol [PLCP] frames as specified in IEEE 802.6 and a second more commonly use method is set out in the ANSI SONET and the similar but incompatible ITU-T SDH framing standards. Most modern ATM switches will support all of these standards based on the physical interfaces that are available for individual products as set out by the ATM Forum.)
- Cell-rate de-coupling and bit timing. (Also defined in the SONET and SDH standards.)

It should be remembered that it is also the task of the SONET or SDH based structures to add and remove “stuffing” bits and idle cells to maintain the required framing structure while passing the asynchronous cell streams as efficiently and as quickly as possible. Many processes in ATM switch fabric are carried out in hardware to maximise and maintain consistent performance under high load conditions.

Put simplistically, the ATM layer deals with:

- The packing and sequencing of ATM cells through switches and interfaces. (Cell multiplexing and de-multiplexing)
- Cell Virtual Path and Virtual Circuit Identifier (VPI and VCI) translation to provide route addressing.

ATM supported services	Higher layers	Segmentation and Reassemble. Convergence	
	Application Adaptation Layer		
ATM transport network	ATM layer	VC level	
		VP level	
	Physical layer	Transmission path level	
		Digital section level	
Regenerator section level			

Figure 3: ATM transport network

- Cell header generation and extraction: at each point the payload data needs to be passed on through the structure.
- Generic flow control. A concept set out in ITU-T Recommendation I.361 which allows UNI attached network interfaces to control data flow capacities through shared data trunk paths.

At this point in the description of ATM, the transport structure can be shown as a physical data path containing virtual routes or paths each containing one or more virtual circuits providing connections for the routed data. These are set up as service routes across the structure with the appropriate circuit performance being specified at the beginning of each circuit session.

Figure 4 shows this approach in pictorial form and shows the obvious resemblance to traditional physical cabling long employed in the broadcast and production world. This resemblance is not accidental as it demonstrates one of the fundamental aims of ATM in providing true data connections via a routed structure on shared network topology.

However, ATM goes further as it also uses the fact that it runs on a network structure using common data bandwidth also to provide circuit distributions as well as point-to-point routes. Certain details covering how this is achieved will be referred to in section 1.4 however; a full description of the ATM version of the multicast process falls outside the scope of this paper. Full details of the signalling and multicast process can be studied in Chapter 5 of the book ATM Telecommunications, Principles and Implementation by P.S. Neelakanta. (Ref. 8)

1.3 ATM supported services

An important layer for the implementation of practical ATM services is the ATM Adaptation Layer (AAL) which sits above the ATM transport structure in Figure 3. This is divided into two sub layers known as the Segmentation and Reassemble sub layer (SAR) and the Convergence sub layer (CS). The basic purpose of these layers is to standardise the way large Protocol Data Units (PDU) are structured to suit particular higher layer services, i.e. the way they are segmented for transmission across the ATM transport structure and subsequently reassembled. PDUs consist of control information and user data being exchanged between peer layer processes. These adapt the ATM transport layers to suit the services being carried such as high quality video streams, voice telephony and packet based data for example.

The standardised AAL developed by the ATM Forum (Ref. 9) are:

AAL-1: Concerned with constant bit rate (CBR) traffic such as audio and video. It supports applications sensitive to both cell loss and cell delay variation (CDV). It effectively emulates Time Division Multiplex (TDM) circuits.

AAL-2: Supports connection oriented services with variable rates that have timing transfer and limited delay variation requirements. This is used with time sensitive, variable bit rate traffic such as “packetised” voice or video.




Figure 4: A pictorial representation of the ATM transport structure.

AAL-3/4: Intended for both connectionless and connection oriented variable rate services. Originally this standard defined two distinct protocols but later merged them into a single AAL. This was primarily adopted to support switched multi-megabit data services (SMDS).

AAL-5: Used for connection oriented and connectionless variable bit rate data services without any particular timing transfer requirements. An example of this would be the connection of bursty data traffic typically found in IP based LANs.

Most existing services use the most appropriate AAL listed above to transport them across ATM. This is important to ensure basic interoperability with products not requiring any additional standardisation from an ATM point of view.

However, an extremely important key issue in the development of broadcast and production services on ATM is the realisation that to produce particularly efficient interfaces best suited to Audio and Video live production, none of the existing application layers are completely satisfactory. It is important to note that a particular interface can use any AAL it requires to set up a connection for a particular service. As an example, a video connection may be set up using AAL-1 and a data management connection may be set up using AAL-5 concurrently from the same interface. It will become clear later in this paper that, as additional standardisation has been necessary to transport AES3 digital audio and SDI video over ATM, a specialist version of the application layer can also be built in to these new standards. This, from ATM perspective produces a user defined AAL generally known as AAL-0. The interoperability issues of such a development is dealt with by additionally specifying the standard defining the application layer being used, over those already defining ATM as a whole when interfacing with the network.

1.4 ATM Signalling and Circuit types.

Up until this point this paper has explored the approach used for ATM to assemble and move the cells through the network structure as well as providing a route addressing structure using VPI and VCI. The next stage in passing a practical service over the network is in using standardised signalling structures to set up and tear down service routes as required. As already stated in previous sections, exploring the details of this subject is beyond the scope of this paper however the key elements will be set out here and those wishing to go into greater depth should explore the book referred to earlier by P.S. Neelakanta (Ref. 8)

The basic ATM signalling standards are defined in ITU-T Recommendation Q.2931 and are further extended by the ATM Forum UNI specifications.

To summarise the key elements;

- Fundamental to ATM networks is the use of a network address to identify each interface device attached to the edge of the network. Along side this, each ATM switch has a unique “prefix” address which is used within the signalling structure to set and identify different routes through the ATM structure. There are four network addresses formats specified for ATM use, each format being identified within the first field of the address. These are;
 - Data Country Code (DCC) format for private use. Internally identifies with which country and registration authority the address is registered.
 - International Code Designator (ICD) format for private use. Internally identifies which international organisation owns the address registration. The British Standards Institute administers the organisation codes.
 - E.164 ATM format for private use. The fields in this format are specified by the ITU-T in recommendation E.164 originally planned for use with ISDN networks.

- E.164 ATM format for public use. Identical with the previous format. However, it is administered and assigned publicly as set out in ITU-T recommendation E.164.
- Permanent Virtual Circuits (PVC) in effect use the VPI and VCI structure already shown in figure 4 to manually route services through ATM switches. This is achieved by entering the switch management of each ATM switch on the chosen route and setting up the service. This involves;
 - Setting up a service contract particularly if Constant Bit Rate (CBR) or Variable Bit Rate (VBR) is required. For example, A typical service contract will require the entry of the service category – CBR or VBR, the switch policing scheme, the peak cell rate, the sustained cell rate, and the maximum cell delay variation allowable on this route.
 - Setting up the service path. This will require the setting up of the reverse path as well if bi-directional routes are required. A typical example would require the entry of the service input port, the input port VPI and VCI, the service output port, the output port VPI and VCI, the service category (i.e. CBR, VBR) and the service contract required within this switch.

As can be seen from above, the sole use of PVC's could involve a very labour intensive operation to keep pace with operational requirements within production centres and would be difficult to manage.

- Switched Virtual Circuits (SVC) use signalling protocols described in ITU-T recommendation Q.2931 and extended by the ATM Forum. This can be enabled on any network port to automatically set up integrated signalling routes between "end devices" or network interfaces and between switches. This could also include the use of Integrated Local Management Interface (ILMI) which is a Simple Network Management Protocol (SNMP) based structure defined by the ATM Forum to provide interface, status and configuration information across an ATM structure. By using this SVC signalling structure, it becomes possible to set up and clear down routes from devices attached to the network. In effect it becomes possible to think of the ATM switch structure much as we would think of a telephone exchange structure. That is to say, that as users we tend not to think of it at all, as our interaction with the telephone switches is automatic. As we dial the number of the telephone we want to call, the exchange switches works out the service route and set up the call from the information entered as the dialled number. Figure 5 illustrates a simple ATM structure where the end devices could be interfaces in production areas capable of making professional audio or video calls across the network to other interfaces similarly equipped. Using the multicast provisions within the ATM standards it would also become possible to set up calls between a single source and a number of destinations that all have the same quality of service as the original connection.

- Permanent Virtual Paths (PVP) connections enable paths to be set up through which both signalling and service connections may be routed. This has the effect of allowing the bundling of a number of connections through the same ATM route. This can also provide a useful method of controlling the routes taken by SVC over a wider area ATM structure.




Figure 5: Simple ATM structure.

- There are a number of circuit types available within ATM. In order of priority these are:
 - Constant Bit Rate (CBR) which reserves a specified amount of guaranteed performance bandwidth between two or more ports on an ATM structure.
 - Variable Bit Rate (VBR) allows the reservation of a sustained amount of bandwidth that is allowed to expand to a specified peak value as the service requires. If at any time this service is not using its peak bandwidth, services with lower priority may use the spare capacity.
 - Available Bit Rate (ABR) allows the reservation of a basic sustained amount of bandwidth that can be allowed to expand to a specified peak if that bandwidth is available across the service route.
 - Unspecified Bit Rate (UBR) allows the use of any free bandwidth capacity across the service route not already reserved for any other service including signalling. It can have a specified peak cell rate.

2. QUALITY OF SERVICE

A key requirement of any network carrying live professional audio or video suitable for use in broadcast and production centres is the need for this to meet the industry standards. A network-based solution for the routing and distribution will have to match these standards in order to allow integration with current systems and working practices. In the case of production audio systems, the Audio Engineering Society (AES) publish a number of standards aimed at this area for digital audio. One of the aims with live production audio structures is to be able to route and mix audio streams with a fair degree of freedom within a production centre and not allow the technology to prevent reasonable programme making ideas from being attempted. Networking the audio structure may well reduce the cost of providing new live programme making if the quality of service requirements can be achieved. However, the author has found that the understanding of what is meant by the phrase Quality of Service (QoS) varies immensely across the networking and broadcast world.

In a pilot project which will be discussed in this paper, a port to port audio latency of 3 milliseconds or better with a latency stability or jitter of 5 microseconds or better is the target.

2.1 Basic packet serialisation

There is a continual discussion in the networking world on the addition of features to improve the quality of service. It is certainly true that huge improvements have been made to switch designs with far more powerful architectures being applied to one of the key areas of packet handling, this being the switch fabric. The pros and cons of different switch designs could easily form the subject of a separate paper.

However, if the true distributed nature of networking is to be exploited to its fullest, the quality of service of the overall structure has to be taken into consideration. This has to include the interconnection of switches and the interface equipment. In the author's view it would be a backward step in both design flexibility and cost to restrict high QoS network topographies to "star" cabled single box designs. If true distributed network structures are to replace more traditional production systems then the far more basic issues raised by serialisation of packets down ports and its effect on system QoS for media streams has to be understood. This is fundamental to service performance and applies at a basic level despite other high level schemes such as VPN and packet tagging schemes that may be applied.




Figure 6: Packet / Cell serialisation

Figure 6 sets out a representation of a network port with the concept of multiple queues feeding port buffers before the data is framed and serialised for transmission. As can quickly be seen if the port is serialising a range of packet sizes from 64 octets to 9,000 octets the likelihood of the next packet having to wait in the buffer for anything up to 76 microseconds (including the inter-packet gap) could be quite high. This is assuming that all the ports are running at 1 Gigabit per second and that no other queue priority mechanism is in use. It can be argued that the majority of packets in this kind of network are likely to be around 1,518 octets in size but even then the wait could be around 16 microseconds and this figure is per port used in the path at 1 Gigabit. There are also the significant additional delays coping with collisions and retransmission as well as the important issue of the use of 64 octet packets in moderately loaded 1 Gigabit network paths being lost due to the collision avoidance mechanism not working effectively with packets of this size. It would improve the situation if such a network were only used to deliver a single service and nothing else. This however would mean a proliferation of stand-alone networks as well as concealing the problems as the system expands, as the QoS is not managed, it is just delivered as an offshoot of extremely low utilisation.

By just taking this serialisation issue alone, it can be seen that delivering a system wide QoS for professional audio with latencies of 3 milliseconds or less with a 5 microsecond or better jitter performance will be virtually impossible in any standardised way unless ATM is used. Jitter can be removed by the use of buffers at the receiving end of the stream but unless

these can be kept very small in size (as they can when using ATM) they rapidly add up to unacceptable amounts of system latency within production centres. This would be unacceptable in the context of **live** broadcast services.

ATM can provide each of its service paths by setting up reserved end-to-end bandwidth connections. Port capacity on these routes is therefore also reserved for each service by using types of stuffing bits and marked idle cells within the port framing structure. This, along with the small cell size delivers very consistent port latency with much improved jitter for each service being delivered through the port. If all the capacity of a port is already in use or reserved, a diverse ATM structure using SVC call management will automatically use an alternate route between the source and destination if one is available. If additional "on line" capacity is required through any route, ATM simply allows for additional parallel connections to be added as required without any complex switch management.

There are a number of other QoS issues within network design but this single issue is used to illustrate the problems faced by networks in delivering high quality audio and video streams with low latency and jitter performance.

3. THE DEVELOPMENT OF A NEW STANDARD FOR AUDIO IN LIVE PRODUCTION SYSTEMS

3.1 Background

The requirements to route and distribute live audio over standard structured cabling within production centres were discussed in detail during 1998. Ideas were developed in conjunction with a few interested manufactures the following year. Around the same time, a paper was published by the AES (Ref.10) and prototype equipment was constructed to enable early tests to be carried out. At least one manufacture had already delivered a hybrid MADI / ATM product. However, no standardised way of packaging AES-3 frames directly into ATM cells had been established, so the industry was in danger of developing a number of proprietary methods of delivering live audio over ATM. In early 2000 a second paper (Ref. 11) was presented and published during an AES UK conference on a theme of networking audio.

It was clear that the BBC would only consider such solutions if equipment could be specified from a number of manufactures and the interface specifications on both the audio and network side of the interface was to an open and published standard. It was also very clear that the audio transport stream specification carried over the structure had to be similar to an open standard. This would then support audio stream transitions between legacy and the new structure with greater ease.

3.2 Setting the requirement

The assumption made was that the core of this new structure would be digital and the interconnection of legacy analogue streams would be serviced by installing converters where required. This structure would also need to be at least as simple to use as current local and centrally controlled audio routers while extending the flexibility and cost effectiveness to cope with the rapidly changing requirements of production in today's environment.

There was little doubt that the AES-3 standard (Ref. 12) should be specified for the audio interface as well as the audio stream structure, this being in wide use in professional audio today. This just left the question of which network interface standard and network fabric would support the requirements of live production audio and was currently an established standard.

It was found by discussion with production staff and by simple measurement of the existing

audio structure within two production centres, that certain performance criteria would have to be met if the techniques used in live audio production were to remain unchanged and continue to be supported on future data networked systems. In particular, digital audio routes would have to exhibit port-to-port latencies of around 3 milliseconds or better. They would also have to exhibit connected path delay variations certainly better than 125 microseconds for general use and better than 5 microseconds for production audio paths, such as those forming a set of “common audio scene” sources in a live multi stream mixing environment. In addition to the digital port-to-port audio route latency of better than 3 milliseconds, the analogue to digital and digital to analogue conversion delays need to be considered in the overall performance to support traditional audio equipment.

Despite the fact that new methods of production could be introduced in many areas, it was clear that the new audio structure had to support any of the existing production equipment and programme making methods. The route performance figures already mentioned would need to be supported for existing as well as new live production systems. It was important to ensure that both local and wide area routing of the audio streams could be supported without the use of large numbers of specialised interfaces within it. For this to be achieved, the network structure would have to be a current ITU approved international standard, capable of delivering low latency switched services with the required quality of service.

The new structure must also be capable of being specified to be highly stable and resilient, providing consistent route performance regardless of system loading. Ideally, the new structure should be able to provide point-to-point routing, audio distribution and audio monitoring as well as being able to provide managed and flexible combinations of local and central control. Finally, the new structure had to be capable of supporting all of the standard AES-3 word lengths and sample rates simultaneously as well as being capable of utilising any standard building structured IT copper and fibre cable infrastructure. Larger production centres may well be supporting a number of output services and there are no guarantees that they will all use common sample rates across all the production process chain. The system must also be able to have new services added over time, which may also have different sample rate requirements from those already in place.

With these performance requirements in mind, and following a number of tests with different network structures, ATM was chosen to provide the network structure, meeting the requirements for this task.

3.3 The work within the AES

As already discussed, back in 1999 at least one manufacture had already produced a hybrid MADI / ATM product and there was no pre-existing standard which could be used to bridge AES3 and ATM. So this was at a point where a number of proprietary ATM solutions could have evolved. As it was also clear that the BBC would require some structures to meet the requirement described in section 3.2, a task group was set up in the appropriate standards area within the AES to work on this problem.

Following due AES process, AES47-2002 (Ref. 13) and AES-R4-2002 (Ref. 14) were published in May 2002.

AES47 is the standard itself and AES-R4 contains guidelines to aid those wishing to implement AES47.

In summary, AES47 covers the following issues.

- The format of the audio samples and the packing of these samples into ATM cells.
- Data protection.
- The grouping of AES3 channels and multi channel formats.

- Defining the ATM Adaptation Layer for use with AES47. (AAL-0 User defined)
- ATM Addressing and the use of Switch Virtual Circuits to set up audio connections.
- AES3 Clock recovery.
- The use of Permanent Virtual Circuits.
- Management and signalling. Setting up, clearing down and joining existing audio routes.

These two documents can be accessed via the AES web site in the “Standards in print” area.

4. THE RADIO 4 ATM PILOT PROJECT

As a result of demonstrations with equipment developed from the prototype used in the standards work, the BBC decided to proceed with a pilot project using AES47. This will transfer the Radio 4 current live production structure from its traditional systems and place this onto ATM running over structured cabling. This is a particularly well timed project as not only has it continued and enriched the work already done, it has triggered the formation the ATM Audio and Video Alliance (ATM AVA) in which the manufactures who have designed and supplied the equipment for this pilot are members. This will prove a useful forum to ensure interoperability and to help push new standards work forward from the manufactures perspective. You will see its trademark on member’s equipment.

4.1 Overview

There are a number of particular features this pilot will be utilising. These are:

- High quality audio for the “on air” programme chain will be carried together with associated services (cue audio, DC signalling etc) on an ATM network, using Switched Virtual Circuits, to the AES47 standard - with low latency.
- Digital audio signals will be carried at their native sampling rates (e.g. 32/44.1/48/96/192kS/s) and bit depths (16 to 24), minimising the number of sample rate conversions and truncations needed.
- All terminal equipment will be connected using structured cabling (CAT5 or better) at 155Mb/s - with each CAT5 cable capable of carrying all eight stereo audio circuits in each direction at 192kS/s 24 bit depth (the terminal equipment sizing chosen for this project). Two unused pairs in each CAT5 cable will carry an AES3 stereo digital audio feed back from each studio to a central area for “emergency” use and a multi-rate reference signal allowing all areas to operate synchronously at any standard sample rate. No central or zone “broadcast” cabling distribution frames etc are needed. Fibre links are used for backbone connections (at 622Mb/s).
- ATM Audio Interface Units will be used in studios and apparatus rooms to interface to conventional analogue and AES3 digital audio equipment. At least two units will be used in on-air areas to provide resilient paths, with each cable connected to a different ATM Switch (hub).
- Transmission output streams (e.g. Radio 4 FM) will each be processed through Sample Rate and Format Converter Units (SRFC) allowing individual control of output sample rate, headroom level, “time signal” insertion timing, etc per transmission stream.
- Control and monitoring of all equipment will be over the ATM network, allowing users to self-select audio feeds etc, and for “control room areas” to override selections, monitor activity, etc.
- Users will only be able to control and monitor facilities appropriate to their location or job function, by “soft” segregation of the available facilities (e.g. a “control room” will have

access to all facilities, but a workshop studio will only be able to access particular sources).

4.2 The studio connections

From the overview above it can be seen that each studio can be connected via an ATM interface unit. This unit is connected to the ATM structure via a single CAT5 or higher cable. In transmission areas, two of these units are installed to cater for reserve connections into the structure for robustness.

Although AES47 will regenerate the AES3 sampling frequency across the ATM structure, this will not allow a phase accurate clock to be used across all the ports within the studio. To overcome this problem a simple synchronisation signal is sent over a spare pair on the CAT5 cable, which is used to provide a digitally locked reference at the required sample rate in the studio. One of the earlier ideas produced by members of the ATM Audio and Video alliance (ATM AVA) was the calculations which showed that a frequency of 153.6 kHz could be used with an industry standard crystal to produce any of the AES3 sample rates required to be locked within the system. This standard frequency is generated in a stable oscillator that is locked to a GPS receiver. In this way it is possible to set up wide area AES47 links and, by using GPS at each site, be able to lock all of the audio ports to the same stable reference.




Figure 7: Connections to the studio

By repeating the block of equipment shown in Figure 7 it becomes possible to build up as large an infrastructure as required without the need for any major central plant. This is distributed around the production centre being connected on standard structured cabling giving complete flexibility.

4.3 The transmission chains

The other key building blocks to complete the Radio 4 installation are the devices that make up the transmission chain. As with the studios, these are connected via CAT5 or better to the ATM switch infrastructure. All the control requirements for each distribution are contained within a 1 U unit connected with a CAT5 cable. These are the Sample Rate and Format Converter (SRFC) units shown in Figure 8. These units also contain stored audio messages




Figure 8: The transmission chains

and facilities to insert time signals. Many of the functions available within this unit are controlled from the studio feeding the transmission chain while others are only available to supervisory or control room staff. These units may also be distributed between different ATM switches and it should be apparent just how easy it is to expand the services delivered from the structure.

Other devices that are required for the transmission process are connected via an ATM interface unit identical to that used within the studios. These include the loudness processors, off air monitoring and the feed to the transmitters themselves.

4.4 Control and Management

In section 3.3 it can be seen that AES47 includes a very simple management structure for setting up, clearing down and joining audio routes. While this provides a very important function within the structure, it clearly does not have the detailed facilities to manage a complete production chain. These have been developed as the Radio 4 pilot has progressed and now form a new specification called "Common Control Interface Specification for ATM Audio Products". All the equipment produced by the manufactures in the ATM AVA conforms to the relevant parts of this document. This specification utilises Simple Network Management Protocol (SNMP) and defines a number of Management Information Base (MIB) objects that are passed between devices using SNMP commands. This forms the bases of the command structure for the ATM attached equipment. Touch screen terminals are used where control of any part of the ATM structure is required. These control points also conform to AES47 and the Common Control Interface Specification.

This additional specification forms the bases of a new standard working group set up within the IEC which will be expanded to include control functionality for video products as well.

5. CONTINUING WORK

Standardising work is continuing in the IEC on AES47 and the Common Control Interface for digital audio and video products on asynchronous transfer mode networks. The object of this area of work is to establish a universally accepted control structure for all ATM attached production equipment.

The next phase of our work at BBC R&D is to work on similar SVC based ATM structures for 270 Mbit SDI video and MPEG 2 video equipment attached to ATM networks.

The BBC's Radio 4 Pilot Project installation and commissioning is continuing, and we hope to be able to report progress and findings at the IBC 2003 presentation.

While standards are in place for MPEG 2 over ATM, there is no common control structure between products. As for 270 Mbit SDI video, there is a standard in draft currently being worked on within a SMPTE working group. However, this does not address SVC control on the SDI routes. We would also like to connect the IEC work on common control with both these video areas. This would lead to a greatly simplified approach to the management of ATM attached live production structures and almost certainly allow the manufacture of a number of management tools for such structures.

6. REFERENCES

1. ITU-T Recommendation I.150 - Broadband ISDN asynchronous transfer mode functional characteristics.
2. ITU-T Recommendation I.361 - B-ISDN ATM layer specification.
3. ITU-T Recommendation I.363.5 - B-ISDN ATM Adaptation Layer specification: Type 5 AAL.
4. ITU-T Recommendation Q.2931 - Digital Subscriber Signalling System No. 2 (DSS 2) – User-Network Interface (UNI) layer 3 specification for basic call/connection control.
5. ITU-T Recommendation Q.2961 - Digital subscriber signalling system No. 2 – Additional traffic parameters.
6. ITU-T Recommendation Q.2971 - Broadband integrated services digital network (B-ISDN) – Digital subscriber signalling system No. 2 (DSS 2) – User-network interface layer 3 specification for point-to-multipoint call/connection control.
7. ITU-T Recommendation E.164, The international public telecommunication numbering plan. Geneva CH. International Telecommunication Union.
8. A Textbook on ATM Telecommunications, Principles and Implementation. P.S. Neelakanta. CRC Press. ISBN 0-8493-1805-X
9. ATM Forum specifications link. <http://www.atmforum.com/>
10. Distributed Network-Based Broadcasting Structures. C J Chambers, BBC. 107th Convention of the Audio Engineering Society, New York, 1999 September 24-27, Pre-print Ref. 5043.
11. Network design for Audio Production Systems. C J Chambers, BBC. Audio Engineering Society, UK meeting: Moving Audio: pro-audio networking and transfer. 8-9th May 2000, Church House, Westminster, London.
12. AES3-1992 (r1997) AES Recommended practice for digital audio engineering -- Serial transmission format for two-channel linearly represented digital audio data.
13. AES47-2002 AES Standard for digital audio — Digital input-output interfacing — Transmission of digital audio over asynchronous transfer mode (ATM) networks.
14. AES-R4-2002 AES Standards project report – Guidelines for AES Recommended practice for digital audio engineering – Transmission of digital audio over asynchronous transfer mode (ATM) networks.