ISO/IEC JTC1/SC29/WG11
MPEG 92/
July 1992

CCITT SGXV
Working Party XV/1
Experts Group for ATM Video Coding

Document AVC-297
July 1992

SOURCE: Australia

TITLE: The ATM Adaptation Layer for video services in the B-ISDN

PURPOSE: Information

Abstract

In this document the ATM Adaptation Layer (AAL) for the transport of video services on the broadband ISDN (B-ISDN) is discussed. An introduction to the B-ISDN and the AAL for video services is given. Methods of multiplexing user information on the B-ISDN are described. Example call scenarios are used to illustrate the advantages of cell based multiplexing. AAL functions for the transport of video services are summarised, and an example AAL is given.

1. Introduction

Asynchronous Transfer Mode performs multiplexing and routing functions in the broadband ISDN (B-ISDN). It is a variable rate, source independent transport mechanism, in which user information is transported in fixed length cells. An ATM cell consists of a 5 byte header and a 48 byte user payload. Routing information in the cell header identifies a virtual channel link. Virtual channel links are concatenated to form a virtual channel connection. The B-ISDN guarantees cell sequence integrity on a virtual channel connection. ATM errors are such that some cells will be missing, and some will be incorrectly inserted. In addition user data in some cells may contain bit errors.

A virtual channel connection and its relation to the B-ISDN Protocol Reference Model (PRM) is shown in Figure 1. The ATM layer does no processing of user data. Service specific functions reside above the solid line at the network edge.

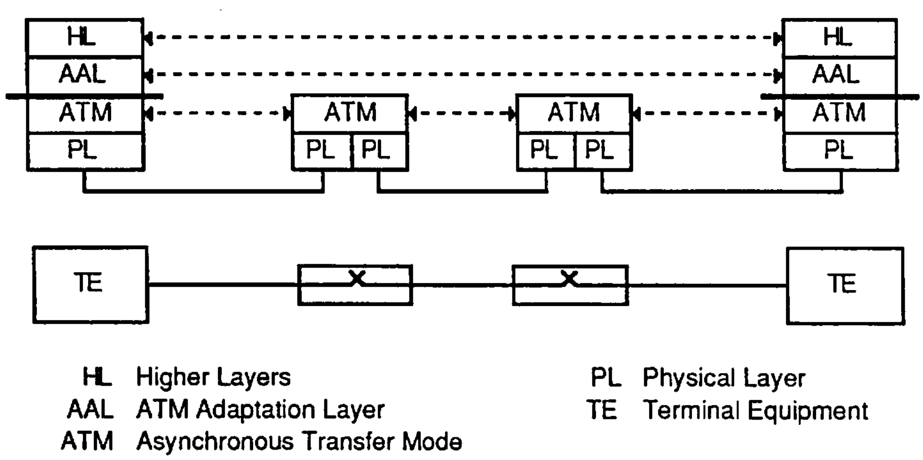

Figure 1. B-ISDN Protocol Reference Model showing a virtual channel connection.

B-ISDN call admission control sets up a new call of the required Quality of Service (QOS), if enough network resources are available, such that the quality of existing calls is maintained. Traffic descriptors such as peak rate, average rate, and burst length may be used to express the bandwidth requirement of a call. During the call, rate control at the terminal prevents violation of the agreed traffic parameters, while usage parameter control acts upon violating cells at the network edge.

The ATM Adaptation Layer (AAL) enhances the service provided by the ATM layer on an end to end basis, and in a service dependent manner. The AAL is subdivided into a Convergence Sublayer (CS) and a Segmentation And Reassembly (SAR) sublayer.

Figure 2 illustrates the B-ISDN PRM and data unit terminology. Data is passed across protocol layer boundaries at Service Access Points (SAPs), as Service Data Units (SDUs). SDUs are transported to peer protocol layer SAPs in one or more Protocol Data Units (PDUs). No SAP is defined at the CS/SAR boundary.

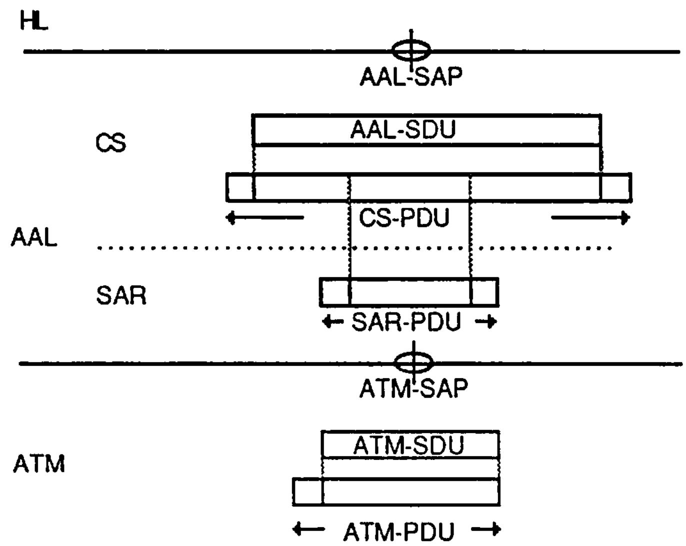

HL	Higher Layers	SDU	Service Data Unit
AAL	ATM Adaptation Layer	PDU	Protocol Data Unit
CS	Convergence Sublayer		
SAR	Segmentation And Reassembly		
ATM	Asynchronous Transfer Mode		

Figure 2. B-ISDN Protocol Reference Model and data unit terminology.

To minimise the number of AAL protocols, four service classes, based upon end to end timing requirement, constant or variable bit rate, and connection mode, have been recognised. Class B covers real time, variable bit rate services. AAL type 2 is mapped to Class B, and is appropriate for variable bit rate coded video. CCITT SGXVIII has not yet agreed on AAL type 2 functional requirements.

2. Multiplexing

B-ISDN will support calls of different media type. Examples of media include audio, video, text and graphics. Bitstreams for some media may be segmented into components eg layered video. Different media, and components, may have different QOS requirements. This impacts upon the method of media multiplexing.

2.1. Methods of multiplexing

Four options are available.

- Higher Layer multiplexing - different media are assembled into a continuous bit stream within the Higher Layer. An example is the systems specification of ISO/IEC CD 11172 (MPEG1). Only one virtual channel is used, which must have a QOS matching that of the most demanding media. Different media in general occupy the same cell payload.
- CS multiplexing - AAL-SDUs are segmented in turn such that all the cells carrying segments from one AAL-SDU appear in the ATM stream before cells carrying segments from the next AAL-SDU. A multiplex identifier is attached to each AAL-SDU. One

virtual channel is used, which must have a QOS matching that of the most demanding media. Each cell, however, carries information from only one media. The CS is able to support the Higher Layers, with functions such as error control, in a media dependent manner.

- SAR multiplexing - a multiplex identifier is attached to each payload segment. Cells carrying segments from different AAL-SDUs may thus be interspersed in the ATM layer. Both the SAR and CS sublayers may support the Higher Layers in a media dependent manner.
- ATM multiplexing - the inherent multiplexing capability of the ATM layer is employed. A separate virtual channel is used for each media. The QOS can be adapted to the requirements of each media. No Higher Layer, or AAL multiplexing overhead is incurred. Media may be routed independently. Network interworking may be simplified.

  The use of separate virtual channels raises the question of media synchronization. In addition, the cost in using multiple virtual channels should not be prohibitive if the benefits of ATM layer multiplexing are to be realized.

Only ATM layer multiplexing fully exploits the service management features provided by B-ISDN.

2.2. *Features of ATM layer multiplexing*

The ability to route media independently within the network, as provided by ATM layer multiplexing, offers advantages in terms of efficient use of network bandwidth. Some interworking scenarios illustrate this:

- interworking between audiovisual terminals and audio only terminals in a multipoint video conference. Video information should not be sent to audio only terminals.
- interworking between video terminals of different capabilities. Layered coding might be used in such a case. A low resolution video terminal should not receive higher resolution enhancement layers that it is not capable of using.
- compatibility between MPEG1 and future MPEG2 video standards. The use of layered coding and separate virtual channels allow MPEG1 compatible video to be placed in one virtual channel, and an MPEG2 enhancement layer to be placed in a separate virtual channel. The layers are routed only to where they are required. This causes no redundancy of existing MPEG1 receivers.

The use of separate virtual channels for multimedia, multipoint calls provides efficient use of network capacity. QOS requirements can be optimised for each media. Introduction of new services is simplified. Placing all interworking components in the same virtual channel may complicate charging, since it may not be possible to determine which interworking components are being used in any particular call.

3. AAL type 2 functions

In determining functions required of the AAL type 2, it may be convenient to classify the video decoder according to its ability to accept errored slices. Non delivery of errored slices may be an acceptable approach in a layered coding scheme. Alternatively errored slices might be delivered. Where a high quality virtual channel is used, error frequency might be low enough such that errors can be displayed. Where error location is also known, that part of a slice preceding the first error location can be used. Loss of variable length coding word synchronization may require the remainder of the slice to be ignored. Error concealment actions are then required eg display of previous frame information, or interpolation of baseband layer in a layered coding scheme.

3.1. *Framing of data*

Framing of video data assists in minimising error propagation. By coding and segmenting the information from a specific area of screen pixels independently, loss of one cell effects only one part of the image.

The MPEG1 coding standard uses such a strategy with the concept of a slice. The standard allows for the number of macroblocks in a slice to be selected according to the bit error rate of the storage/transmission medium. To be successful in the ATM environment this framing needs also to be carried into the AAL. The start of each slice should begin a new AAL-SDU.

The number of pixels to be coded at a time is chosen with respect to efficiency and error propagation. Study is required as to whether the number of pixels should be fixed, or selected according to the error rate of the connection, at connection setup. Since the last segment of each packet may be only partially filled, the efficiency of framing, particularly at low bit rates, needs to be studied.

Framing of user data within the Segmentation And Reassembly (SAR) sublayer may be by indication of segment type. The position of the segment from within the CS-PDU is indicated with signals of Beginning of Message (BOM), Continuation of Message (COM), End of Message (EOM), and Single Segment Message (SSM). Alternatively one signal to indicate the first segment of the CS-PDU, and one to indicate following segments, may be adequate.

3.2. *Length indication*

Partially filled cell payloads suggests the requirement for cell payload length indication. VLC coding produces slices which are in general not multiples of eight bits. However, bit stuffing may precede start codes in existing coding standards, allowing slices to be filled to the next byte or segment boundary.

Indication of the number of used bits in a cell payload may be an expensive requirement. If AAL-SDUs are restricted to be multiples of eight bits in length, then length indication can be reduced to indication of the number of bytes used in the SAR-PDU payload. Alternatively if AAL-SDUs are restricted to be an integer number of SAR-PDU payloads in length, minus CS-PDU header and trailer field lengths, SAR-PDU payload length indication is unnecessary.

3.3. *Sequence Number*

Sequence numbers applied to the SAR header can be used to detect cell loss and cell insertion. Indication only of missing or additional cell segments could be done using a CS-PDU length field. This cannot, however, determine error location.

Sequence numbers use a modulo N count. Loss of up to N-1 cells can be reliably detected. The value of N, and hence the field length required, must be selected such that the probability of N cells lost in succession is made small. Knowledge of B-ISDN cell loss characteristics is required to reliably select N.

Cell loss in a statistically multiplexed environment occurs when queue buffers at ATM nodes overflow. Consecutive cell loss is likely.

Block error correcting coding might be used in the CS sublayer, and requires reliable cell loss detection performance.

3.4. *Multiplexing*

In addition to multiplexing required for service interworking of multimedia and layered video calls, layering within interworking layers, to cope with ATM cell loss characteristics, may also be desirable. It may be appropriate to place video control data and coded DCT frequency coefficients into separate cell streams. Further layering of luminance and chrominance coefficients may also be of value. Such layering allows interworking layer components to be treated differently by the network eg different cell loss rates, or differently within the AAL eg different forward error correction techniques. The AAL may need a multiplex function to support this additional layering. SAR based multiplexing may be more able to satisfy particular timing constraints compared to CS multiplexing. It is however, more expensive in terms of bits used, as identification is required for each SAR-PDU, rather than for each CS-PDU in the latter case.

3.5. Forward Error Correction

The cell payload is likely to require protection from random bit errors. Where the Synchronous Digital Hierarchy based interface is used for the Physical Layer, the payload scrambler causes error multiplication of two. The current AAL type 3 and type 4 SAR trailer CRC may also be suitable for video transport.

Block forward error correction schemes are also proposed to deal with cell loss. In the CS sublayer redundancy cells may be periodically inserted. Interleaving of cell data may also be appropriate. The delay introduced by such schemes may make these techniques applicable only to higher rate services.

3.6. Timing and synchronisation

The AAL type 2 may have to provide end to end timing and synchronization functions. The receiver local clock may need to be synchronised to the transmit terminal clock. Cells come from variable rate sources and are subject to network jitter introduced by queueing buffers at ATM nodes.

Time stamping of cells is an appropriate method. At the receiver, time stamps drive a phase locked loop, whose clock reads out received cells from a local buffer. Time stamps may be sent regularly as a sequence number, or irregularly as a timer value. The phase locked loop time constant must be large enough to remove network jitter.

Where separate virtual channels are used, delay jitter between interworking components should be bounded. Appropriate methods of synchronization must be studied.

4. Example AAL type 2

An example AAL type 2 is shown here to illustrate AAL functionality. This AAL has been used in simulations of the B-ISDN for the transport of signals from layered video codecs based on CCITT H.261 and MPEG1. An AAL-SDU holds one MPEG1 slice. The decoder is assumed to be able to use errored slices up to, but not past, the first error location. The AAL thus passes to the decoder only that part of an errored slice up to the first lost segment. A signal indicates to the Higher Layers if the AAL-SDU is complete, or if trailing segments are missing.

4.1. SAR sublayer

The SAR sublayer of the example AAL adds a one byte SAR-PDU header and a two byte SAR-PDU trailer as shown in Figure 3. SAR fields should be kept short to maximise efficiency. The SAR-PDU header fields are:

- a 2 bit segment type field to indicate one of BOM, COM, EOM or SSM.
- a 4 bit sequence number field allowing detection of loss of up to 15 cells in sequence.
- a 2 bit SAR multiplex identifier field to identify a specific component of the SAR multiplex.

The SAR-PDU trailer fields are:

- a 6 bit reserved field. The reserved field might be used for SAR-PDU header error correction.
- a 10 bit cyclic redundancy check field. The CRC code is applied to the whole SAR-PDU, excluding its own field. The CRC code shown in Figure 3 is that used in AAL type 3 and 4, and is able to detect and correct one bit error, and two correlated bit errors.

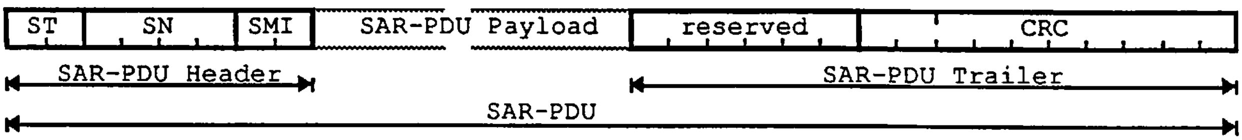

ST - Segment Type BOM 10
COM 00
EOM 01
SSM 11

SN - Sequence Number

SMI - SAR Multiplex Identifier

CRC - Cyclic Redundancy Code

$G(x) = 1 + x + x^4 + x^5 + x^9 + x^{10}$

Figure 3. SAR fields for example AAL.

4.2. *CS sublayer*

The CS sublayer of the example AAL adds a one byte CS-PDU header, and no CS-PDU trailer, as shown in Figure 5. The header fields are:

- a 4 bit reserved field. This field might later be used for transport of end to end timing information.
- a 4 bit CS multiplex identifier field to identify a component of the CS multiplex. This field enhances the multiplexing capability provided by the SAR sublayer.

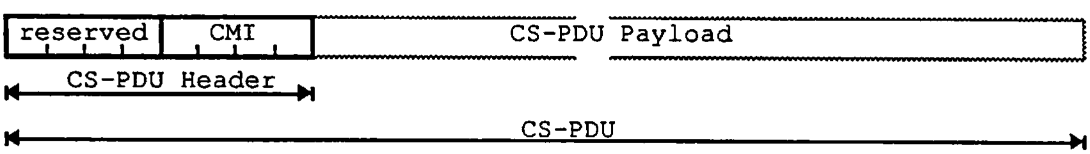

CMI - CS Multiplex Identifier

Figure 5. CS fields for example AAL.

5. Conclusion

This document has presented an overview of the ATM Adaptation Layer for variable bit rate video services. The advantages of cell based multiplexing for multimedia and video service interworking have been stated. Likely functions required for the AAL type 2 have been described. An example AAL for video transport has been presented. Future work is required to determine the suitability of the example AAL fields.

It should be noted that, since cell loss and error recovery performance is highly dependent on AAL functions, cell loss experiments must make some assumptions regarding AAL. The example AAL shown here could be the basis for definition of a suitable AAL for such experiments.