Experimental Testbed Results for Broadband Residential Video Service QoS Management

R. J. Green, S. I. Woolley, N. W. Garnham*, and K. P. Jones*

The University of Birmingham, School of Electronic and Electrical Engineering, Edgbaston, Birmingham, B15 2TT, UK
*Fujitsu Telecommunications Europe Limited, Solihull Parkway, Birmingham, B37 7YU, UK

Abstract-**This paper presents residential video service Quality of Service (QoS) results obtained from an experimental testbed employing constant bitrate MPEG-2 transport streams on a real ADSL/ATM access network testbed.**

The experimental design addresses the choice of MPEG encoding parameters and the spatio-temporal characterization of video test content. The Institute Telecommunications Science video quality metric is used for objective video fidelity measurement.

Our results demonstrate the complex multilayer propagation of QoS parameters, using deterministic and exponentially distributed cell losses, and their relationship to video quality and video bitrate. The paper provides a basis for developing and comparing these results with similar video service systems; for example, using MPEG encoded video over IP. These results provide an important insight into the factors relevant to the provisioning and management of new multi-service network infrastructures.

I. INTRODUCTION

Use of the public access network for the provision of broadband services is now feasible with the availability of Asymmetric Digital Subscriber Line (ADSL) and fibre technologies. Home users can use Video-on-Demand (VoD) and the Internet by connecting to service providers via multiplexers to a new ATM backbone network.

The telecommunications industry has put considerable investment into the development of Asymmetric Digital Subscriber Line (ADSL) technology in order to deliver medium-bitrate (2-8Mbps) services to the home. These data streams require an access network infrastructure to manage services. The access network performs concentration of subscriber connections; multiplexing; signalling and routing operations. Integrated operational support and network management systems are essential features of access networks. Video services having stringent timing and error constraints, present a particular challenge. The delivery of high data rate video at satisfactory quality levels requires an understanding of the cumulative effects of QoS propagation through multiple protocol layers. In this work ATM bitrates of 3-5 Mbps have been used to meet the data rate of ADSL.

The standards for VoD and ATM connectivity are defined by a number of bodies, the most notable of which are the ATM Forum and the Digital Video Broadcasting (DVB) groups. Many years of activity in the area of video coding has established the work of the Motion Picture Coding Experts Group (MPEG) as an industry standard. The coexistence of these standards provided a platform for the adaptation of MPEG-2 streams [1] across ATM.

II. MPEG VIDEO TRANSPORT STREAMS

MPEG video streams consist of groups of pictures ((a) in figure 1) containing I-, P- and B-frames. (I)ntra frames are coded on the basis of spatial information contained within them and without reference to other frames. (P)redicted difference frames are derived from I-frames, and (B)i-directional frames are estimated from past and future I- and P-frames [2].

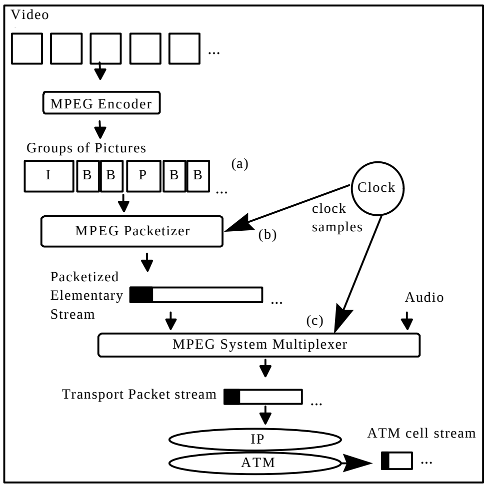

Figure 1. MPEG Transport of Video over IP/ATM

For the purpose of successful video decoding, the I- frames should be error-free and the clocks, associated with their presentation, reliable. A decoder uses periodic timing references to playback a video stream containing groups of pictures. *Decoding timestamps* indicate when to start decoding a video frame and *presentation timestamps* dictate the time to display a frame to the user ((b) in figure 1). Together, these timestamps aid long-term synchronisation and are supplied at least every 700ms. Transport packets, carrying video and audio data, are received at irregular intervals due to the accumulation of delay variations (jitter) within the broadband network. Unique Program Clock

References (PCRs), occurring at least every 100ms, are used by the decoder to reconstruct the encoding clock ((c) in figure 1). This enables synchronised data delivery at a Constant Bit Rate (CBR) to the audio and video decoding stages. To maintain jitter bounds in the CBR environment the baseline VoD model proposed by the ATM Forum delivers 2 transport packets (2 x 188 bytes) in each frame of ATM Adaptation Layer 5 (AAL5) data; where AAL5 [3] provides the means to deliver frames between sender and receiver at a given quality of service (e.g. delay, error rate and jitter). Provided that PCR-carrying packets are delivered within acceptable delay bounds the elementary streams (for video, audio and other data) can be regarded as decodable.

III. THE EXPERIMENTAL TESTBED

The experimental results were obtained from a real ATM/ADSL testbed with content encoded with an MPEG-2 hardware encoder and a commercial video server. Figure 2 shows the key components of the testbed which was configured as a residential broadband access network.

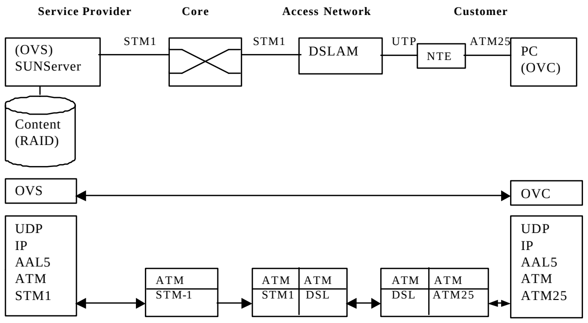

Key: OVS – Oracle Video Server; OVC – Oracle Video Client; DSLAM - Digital Subscriber Line Access Multiplexer NTE – Network Termination Equipment;

Figure 2.Logical configuration of the Testbed

The following equipment was used for the experiments:

1. ATM/ADSL Testbed Components.

Fujitsu FETEX-150 ESP Broadband ATM Switch;

Fujitsu Digital Subscriber Line Access Multiplexer (DSLAM); Network Termination Equipment (NTE) – ATM 25Mbps interface.

2. Video Content

Optibase Moviemaker 200 Basic hardware encoder was used with a Betcam SP video recorder.

3. Video Client/Server

An Oracle Video Server (OVS) [4] running on a SUN Enterprise Server with a SUN 155 Mbit (STM1) ATM interface card providing classical IP over ATM. The Oracle Video Client (OVC) was run on Windows NT PC with an Efficient Networks ATM25 interface card a providing classical IP over ATM. An ATM Permanent Virtual Circuit (PVC) was provisioned between the customer equipment and the video server. A Realmagic Netstream 2 MPEG hardware decoder was installed on the client [5]. The Optibase encoder did not support concealment techniques. Therefore, the result of transport error propagation on the decoded video is determined by the MPEG decoder-board's implementation.

4. ATM Impairment and Capture

A Hewlett-Packard BSTS E4219A Network Impairment Emulator Module (NIM) was used to generate ATM errors [6]. A Radcom Prismlite analyser was used to capture AAL5 Protocol Data Units (PDU).

Specific considerations involved in designing the experiment are now reviewed.

Video Server Parameters: Six MPEG2-TS packets were used per AAL5 Service Data Unit (SDU) (1128B) and, within OVS, four AAL5-SDUs were accumulated before being passed to the network interface. This was a strategic decision made because the ATM Forum's suggestion [1] of 2 MPEG packets per AAL5-PDU (minimising jitter) imposed too high a processing overload on the server. Each OVS data frame is segmented into 25 cells.

Cell Discard Experiment: This investigation used single cell losses. Two types of cell loss distributions were used:

1.) Deterministic Discards – This case provides a control to act as a quality differentiator by providing a fixed cell loss interval (e.g. 1 in 1000 cells); 2.) Exponentially Distributed Discards – selected to approximate buffer losses; where

$$\text{Mean(CLR)}= \int_{-\infty}^{\infty} x f(x) dx \qquad (1)$$

where ,

x = Error interval (cells)

$$f(x) = \text{Exponential pdf} = \frac{1}{m} e^{-\frac{x}{m}} \qquad (2)$$

m = Distribution Mean

Table 1 shows the 3 MPEG-TS bitrates used for the test along with their respective ATM bitrate, cell rate and the length of sequence that can be captured by the Radcom analyser. Considering consecutive losses, each PDU has 25 cells and the HP Network Impairment Emulator can generate 1-8 consecutive losses. Tests show that in practice, with entire PDUs being dropped in the event of one cell loss, the propagation of consecutive losses into the next PDU results in negligible differences in video quality.

MPEG-2 TS (Mbps)	ATM (Mbps)	ATM (cps)	Max Radcom Capture
2.379	2.883	6,800	95 seconds
3.449	4.197	9,800	65 seconds
4.516	5.469	12,900	50 seconds

Table 1. Radcom Analyser Capture buffer capacities for 2, 3, and 4 Mbps MPEG2 Transport Streams

Impairment Insertion Point: Cell losses were inserted at the ATM25 connection to the PC. Cell losses were applied only to the ATM assigned cells forming the reassembly process for the MPEG-TS. This is a logical choice because the same CLRs can be considered in the presence of other data interspersed on the same ATM channel in future work. For example, a video stream may be presented on ATM Virtual Circuit 73 and an Internet connection on ATM Virtual Circuit 32 using the same network termination equipment. Therefore data are asynchronously interspersed and errors are dispersed over both connections.

IV. TEST CONTENT

This experiment used Half D1 MPEG2 video encoding to encode TV quality video MPEG at data rates between 2-4 Mbps. A Half D1 (HD1, resolution of 352x576) encoder was used as opposed to Full D1 (FD1) (704x576) encoder because this reduces the amount of data requiring compression. Typically, HD1 operates in the 2.2-5 Mbps range giving near Betacam quality and FD1 is used in the 5-10 Mbps range giving Betacam Quality [7]. Four test clips were selected for the test sequence. The clips were chosen to have a variety of spatial and temporal information content. Figure 3(a-f) are representative scenes from the four 8-second clips forming the test sequence.

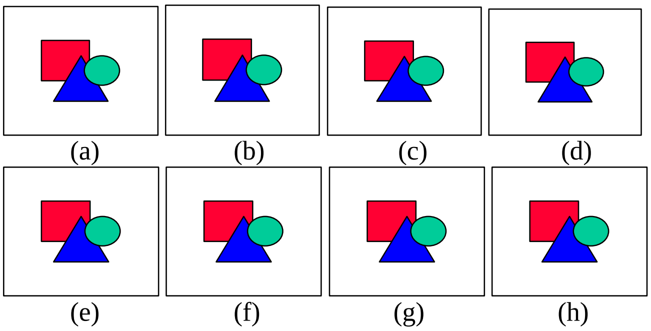

(a) (b) (c) (d)

(e) (f) (g) (h)

Clip Name	Description	Reference
Advert	Screen pan and still frames.	a, b
News	News presenter moving hands/face and background road traffic.	c, d
CITV	Morphed computer graphics against rolling background.	e, f
Cricket	Three scene cuts containing close up action and fast pans.	g, h

Figure 3 (a-h) Sample videoframes from the test sequence showing the range of spatial and temporal content

The clips were edited into a 101 second sequence in which each of the four clips are repeated three times. This video was loaded onto the video server and transferred over the testbed. Impaired AAL5 frames were captured and the decoded video captured on the Betacam video recorder. The next section details the analysis of this data.

V. VIDEO QUALITY ASSESSMENT

Decoded video may be assessed by either *subjective* or *objective* means. Subjective assessment analyses the opinions of a carefully selected group of human viewers. Observers vote on perceived quality according to a 5-point scale where 5 is excellent and 1 is very annoying. Subjective assessment uses well established techniques developed for evaluating conventional television. However, these tests are time consuming to setup and cannot be used for in-service monitoring. For consistent, real-time, video-quality measures objective assessment algorithms may be used. These may either calculate a distortion measure based on the input and output images or attempt to compute a quality rating by mimicking the Human Visual Systems (HVS) response to spatio-temporal abnormalities.

Objective methods are required for real-time, constrained bitrate, encoding to minimize visible compression artefacts. They are also required to manage errors for provisioning, system evaluation (in-service or out-of-service) and comparison of transmission channels and decoder functionality (e.g. recovery and concealment). Webster [8] presents the Institute of Telecommunications Sciences (ITS) objective assessment system that makes meaningful evaluations of quality without viewing panels. The ITS quality rating is mapped to the standard quality scale of 1-5. As the ITS metric was designed for low bitrate applications the scores are overestimated. Variation of the metric in the 3-15Mbps range is reported to be 0.2 impairment units [9].

VI. RESULTS

A visual indication of the performance of each test case is presented using the Institute for Telecommunications Science (ITS) video quality indicator. The ITS mean opinion score (ITS_{mos}) metric is used to assess differences between test cases relative to each other, having a common set of parameters and testbed. Scaled plots of the reference and degraded Temporal Information (TI) [10] (against time (frames) – 25fps) provides an index into the test sequence. TI shows the degree of motion in the sequence; e.g. the news sequence has low motion, whereas the cricket sequence contains fast motion, camera pans and quick scene cuts.

A. Categories of Observable Artefacts

This section presents a sample of the range of observable artefacts that were in the first 650 frames of the 2Mbps encoding which was subjected to exponentially distributed discards at CLR=10^{-4}. shows the first 650 frames of the test sequence. The Temporal Information (TI) of the reference sequence is shown in blue and the degraded TI shown in green. The ITS_{mos} score that ranges between 4.77 and 1 is shown as a black staircase plot where each step represents 10 frames. The time of the AAL5-PDU discards are shown by the red pluses (+). Each frame represents 40ms intervals.

Magenta line segments show the observable video errors that were observable on a frame analysis of the sequence. The long line segments represent contiguous errors.

Figure 4. ITS$_{mos}$ Time History Plot Of The 2Mbps Encoding At CLR=10^{-4} With Exponential Discards.

Figure 4 shows that there were video frame losses in the second and third clip where the ITSmos score is 1; the lowest score on the five-point scale. In particular, the third clip (Frames 500) the loss of video frames causes the alignment of the feature vectors to be lost thereby resulting in a score of 3.25. Figure 5 to Figure 7 show example error bursts that occurred in this sequence. The frame numbers presented correspond to those of Figure 4.

161 169 185

Figure 5. Selected frames from a 1 second burst of tiling errors.

481 491 494

Figure 6. Examples from frames 478-494.

650 *652* *653*

Figure 7. Frames 650-653 contain a significant glitch.

Figure 5 shows that starting on frame 161 severe tiling errors extend to frame 172. Frame 173 to 185 are affected by horizontal slice blocking. In Figure 6 blocking and tiling errors span 17 frames. In Figure 7 full screen tiling is observable with a “synchronisation” error on frame 653.

All the errors were very noticeable on playback and are representative of all other test cases considered. These artefacts characterise the observable errors seen at higher or lower cell loss rates, i.e. similar errors are seen nearly-continuously (e.g. CLR=10^{-3}) or less frequently (e.g. CLR=2×10^{-5}).

B. *Results Analysis*

Undegraded Performance: Figure 8 shows the difference between the 6Mbps, t6ok, reference and the 3Mbps encoding, t3ok.

Figure 8 ITS$_{mos}$ Time History Plot comparing the 3Mbps Undegraded Encoding with a 6Mbps Reference.

The set of ITS values for the sequence has been time collapsed into an overall mean value. The means for t2-4 are all 4.7 (to one decimal place), the maximum value of the ITS$_{mos}$ metric. Figure 8 shows small deviations from 4.7 where coding performance has dropped during changes in temporal activity. ITS values were computed for exponentially distributed frame errors. The sequences are resynchronised at each scene cut (as necessary) to compensate for the loss of play-back frames resulting from data losses.

Distribution Accuracy:

The HP network impairment emulator generated the exponentially distributed errors according to a targeted mean CLR. Analysis of these distributions used the time stamps of the arrival time of the last cell of each errored AAL5-PDU. This data are limited by the capture buffer’s memory size where 50 to 90 seconds of data can be stored. For the deterministic cases simple cell counting is used for their distributions. Table 1 shows the calculated means for the error distributions used in this experiment.

Case	Mean	Median	Min	Max	Skew	CV
t2clx2e5	6.02	5.03	0.208	12.9	0.09	0.72
t3clx2e5	6.51	4.23	0.049	17.6	0.70	0.93
t4clx2e5	3.20	2.94	0.100	8.53	0.75	0.84
t2clx1e4	1.650	1.210	0.011	6.04	1.17	0.93
t3clx1e4	0.933	0.749	0.0182	3.09	0.88	0.77
t4clx1e4	0.740	0.465	0.0155	3.71	1.62	1.02
t3fmcx1e4	0.955	0.508	0.0016	4.64	1.65	1.08
t3thclx1e4	0.874	0.641	0.0354	4.86	2.27	1.01
t2cld1e4	1.470	1.470	1.470	1.48	2.30	0.002
t3cld1e4	0.999	1.010	0.385	1.02	-5.85	0.091
t4cld1e4	0.775	0.774	0.770	0.78	-0.02	0.003
t2clx1e3	0.162	0.103	0.00167	1.06	1.77	0.978

$t\langle \text{videostream\_bitrate}\rangle cl\langle \mathbf{d}\text{eterministic} \mid \text{e}\mathbf{x}\text{ponential}\rangle\langle CLR\rangle$

Key: CV: The coefficient of variation (CV), or relative standard deviation.; th, fm – defined below.

Table 1. Summary Statistics Measuring the Distribution of AAL5 Errors used in this Experiment (seconds).

Changing the MPEG Video GOP length: The standard group of picture (GOP) structure for PAL TV was used (N=12, M=3). Two additional encodings at 3Mbps were made with different GOP lengths:

I. fm – fast motion N=10, M=3. This may improve coding performance for fast, complex, action sequences.

II. th – talking heads N=18, M=3. This increases compression efficiency to reduce bandwidth. However, this is only suitable for slowly changing sequences, for example, for a news reader.

Overall, it was concluded that without the encoder supporting adaptation of the GOP structure to motion content, the standard GOP gave the best performance for bitrates in the 2-4Mbps range.

C. *Overall Results*

The results of the experimental test cases have been shown in the preceding section. Figure 9 shows the probability distributions of ITS_{mos} scores for the test cases.

$t\langle \text{videostream\_bitrate}\rangle cl\langle \mathbf{d}\text{eterministic} \mid \text{e}\mathbf{x}\text{ponential}\rangle\langle CLR\rangle$

Figure 9. Probability distributions of ITSmos Scores for the Cell Loss and Video rate Test Cases.

Figure 9 shows that a CLR of 10^{-3} produced very poor results. This was due to frame-freezing during playback. Picture freezes are not acceptable for video services. The deterministic and exponential test cases at CLR=10^{-4} also clearly show poor performance. The most important conclusion of this work was that a CLR of $2\text{x}10^{-5}$ (i.e., 1 in 50,000 cells) is required to avoid video frame losses for video encoded at 2 and 3Mbps. At this CLR, 32 video frames (1.3 seconds) were lost from the 4Mbps encoding, having a cell rate of 12,900cps. It was noted that bursts of errors were visibly more disturbing than regular, deterministic errors. This is not evident from the histograms. Table 2 summarises the relationship between video performance at a given CLR and the ITSmos scores of Figure 9.

CLR	Bitrate (Mbps)	Observable Artefacts
$1\text{x}10^{-3}$	2	Screen freezing, jerkiness and continuous tiling.
$1\text{x}10^{-4}$	2, 3 4	Widespread tiling and jerkiness. Significant jerkiness in fast-action sequences.
$2\text{x}10^{-5}$	2, 3 4	Artefacts significantly fewer. Tiling and jerkiness are observable. Artefacts become more frequent.

Table 2. Categorisation of Artefacts for CLR and Video Bitrate.

D. *Traffic Shaping*

This section analyses the relationship between the MPEG2-TS packets, the video server and the DSL-ATM25 connection.

Approximation of Packet Inter-arrival Times (IATs): For the 2.3796Mbps stream with 6 TS packets per AAL5-SDU and 4 SDUs buffered by OVS the time between data being passed to the network interface is 14.8ms. The data are then shaped onto the 5Mbps PVC to the ADSL DSLAM. At the ADSL DSLAM the data (ATM cells) are rate adapted to the 6Mbps circuit. Table 3 shows the minimum IAT between AAL5-PDUs before and after the DSLAM.

Raw Circuit Bitrate	Cell IAT (ms)	PDU IAT (ms)	IAT of 4 PDUs
5.5Mbps	0.077	1.92	7.68ms
6.2Mbps	0.068	1.7	6.8ms

Table 3. Approximate Minimum cell and PDU Interarrival Times for the 2.38Mbps MPEG-TS with Rate Adaptation.

Figure 10 shows the histogram of AAL5-PDU IATs at the NTE ATM25 interface. Here the IAT of AAL5-PDUs is determined primarily by the video bitrate, the buffering within the video server and the buffering and rate adaptation at the DSLAM. Figure 10 clearly shows the minimum AAL-PDU IAT of 1.7ms, indicating brief buffering in the DSLAM. The cells of each PDU are not passed back-to-back and the cells on the 6Mbps circuit are not at the same spacing they are on the 5Mbps circuit (1.92ms, Table 3). The buffering within the video server can also be seen where there is 14.8ms between AAL5-PDUs.

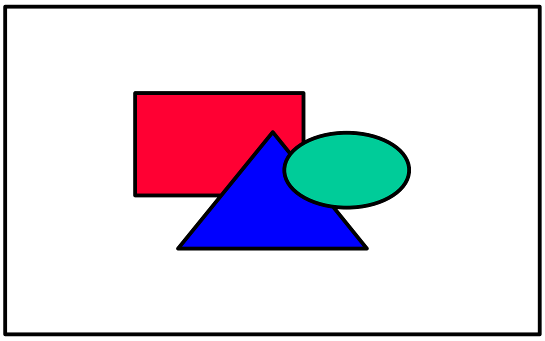

Figure 10. Histogram of AAL5-PDU IAT for the 2Mbps Encoding with 5Mbps Traffic Shaping.

Figure 11 plots PDU IAT against time (seconds) showing the near regular spacing of AAL5-PDUs.

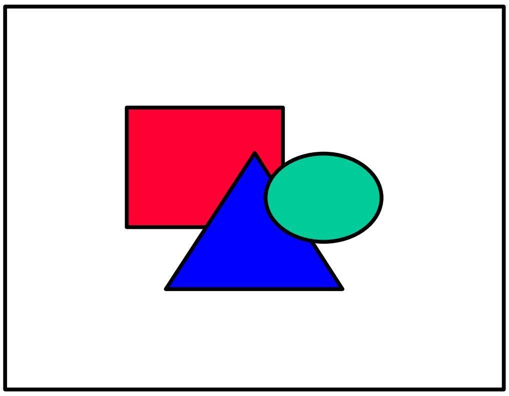

Figure 11. Graph showing the IAT of AAL5-PDUs against time (seconds).

The cell spacing, or more accurately, AAL5-PDU and IP/UDP spacing is determined by the video bitrate, the video server, IP packetization and Traffic shaping and rate adaptation performed on the ATM cells.

The traffic from the video server has been shown to be not truly constant bitrate. Primarily, this is due to the use of a 5Mbps Traffic Shaper on the server's ATM card for the video stream's PVC. To enable the video bitrate to be varied from 2-6Mbps, without changing the PVC's configuration for each experimental case, the ATM circuits to the DSLAM via the FETEX core switch were provisioned to be Unspecified Bit Rate (UBR). Provided the rate of traffic from the video server did not exceed the maximum rate provisioned for the DSL circuit no data were lost. To use CBR circuits the ATM shaper would need to operate at the video stream's ATM cell rate and meet the bounds of CDV tolerance (CDVT).

Traffic shaping within the video server has been shown to be very important to ensure that traffic meets the provisioned bitrate of each xDSL subscribers. Using UBR circuits through the core switching environment is feasible for CBR MPEG transport streams provided back-to-back cell interarrival times are within the bounds that the xDSL access multiplexor's buffer can absorb when multiplexing the stream onto the subscriber's provisioned DSL line rate. This is an issue only when DSL lines are provisioned at different speeds. Using true CBR circuits through core switches implies the video server's traffic shaping must meet the CDV tolerance of the CBR circuits.

The choice of circuits for a commercial deployment are important. This paper has shown that the use of CBR circuits for CBR video requires the video server's traffic characteristics to match the CDVT provisioned throughout the core, access and home (NTE) network.

E. *MPEG Packetizing and Clock References*

The MPEG standard specified Program Clock References (PCRs) should be sent at least every 100ms [11]. Analysis of the video streams coded using the Optibase Moviemaker encoder showed that PCRs were inserted every 22.29ms. This is approximately the interlaced field time of each PAL frame lasting 40ms.

The Optibase encoder used fixed length Packetized Elementary Stream (PES) packets. This is the layer above MPEG-TS packets. The Video Elementary System (VES) data (2-4Mbps) are split into 8,192 byte PES packets before being segmented into transport packets and multiplexed with the Audio and System table packets. There were 45 MPEG-TS packets per PES packet. This accounts for approximately 1 PES packet every 8 AAL5-PDUs or every 185 cells.

With a CLR of 5×10^{-3} (1 in 185 cells) every PES packet would be corrupted. As the encoder did not support variable length PES packets (an optional packetizing layer in the MPEG standard), their relationship to video quality could not be investigated further in this paper.

F. *ATM AAL5 and IP/UDP Encapsulation*

This research has included the novelty of an IP/UDP layer. Although an ATM card was used within a client PC, UDP/IP encapsulation was used to overcome an incompatibility between OVS and the SUN 155Mbps ATM card.

As AAL5 corrupt data delivery was not supported by OVS, when bit errors, cell loss or cell gain, are detected using the CRC, the entire PDU (carrying 1,128Bytes of MPEG data) is discarded.

Gringeri [12] showed use of a padding algorithm to align MPEG-TS packet data after loss of cells from AAL5-PDU was shown to reduce the number of observable errors by 7% compared to discarding the entire PDU. The paper used 2 MPEG-TS packets per AAL5 PDU and a CLR of 4.24^{-5} on a 6Mbps stream. Therefore, only 376 bytes of MPEG data were being lost with each discarded AAL5-PDU. Delivering AAL5-SDUs that had bit errors was shown to reduce noticeable errors more than 1 frame by 20%.

With this in mind this work draws attention to the CRC within the UDP/IP protocol. UDP, which uses IP as the underlying transfer protocol, is a connectionless, best effort protocol. Transportation over ATM ensures that data are received in order, without duplication and on-time (not too early or late). The results of this work are directly applicable to transferring UDP/IP over transport protocols other than ATM, where the UDP packet header's length and CRC will lead to discarding the 1,128Bytes of MPEG data. Discarding 6 TS packets produces poor video-quality results. It should be noted that the UDP CRC is optional, though it is usually implemented.

VII. CONCLUSIONS AND FURTHER WORK

This paper has shown the results of an investigation into the multilayer impact of QoS for high data rate video services delivered over a real access network platform by making use of commercial equipment. Ultimately, QoS is the degree of satisfaction subjectively defined by the user of the service. However, from the networking viewpoint it is the collective effect of performances that may be measured at a hierarchy of access points. Using the Fujitsu Telecommunications testbed the effects of errors have been shown to have a direct relationship on real access network deployments using ADSL/ATM/AAL5/IP/UDP/MPEG-TS/MPEG-2 video protocol layers. This research used half-D1, half TV resolution, MPEG video coding of test sequences in order to provide high quality compression in the 2-5Mbps range. The results of this work convey the key issues that need to be understood, and points the way for future research work. These issues included the coding of video test content, the testbed, video server, network interface parameters and the capture of results by video-recording and by data capture.

This paper has used the objective video quality measure known as the ITS mean opinion score metric to compare the results that were obtained by modifying the experimental variables, namely, changing the MPEG-TS data rate and the ATM cell loss ratio; i.e. to compare the performance of video relative to the impairment parameters.

The work has shown that CLR needs to be less than 1 in 50,000 cells ($2x10^{-5}$) to prevent severe video quality degradation to provide services using bitrates greater than 3Mbps. Although the hardware decoder was developed for ATM/ADSL networked video applications, the visibility of artefacts are shown to be tiling or blocking errors. The results presented here have shown that 4Mbps video is particularly intolerant to cell losses below 10^{-5}. The loss of 1,128-byte data packets is severely detrimental to playback performance. The results of this experiment have been related to the performance of MPEG over UDP/IP. The paper provides a basis for developing and comparing these results with similar systems, for example MPEG over RTP/IP/ADSL/Fibre.

Protocol layers are often thought of as transparent and the QoS definitions are hard to determine in advance and even apply to evaluate service quality. With diverse service requirements for the Internet, Voice over IP and video services the cohesion of layers needs to be much better. This paper has shown the effects of exponentially distributed errors and it is these bursts that need to be transparent to the video service viewer. The framework presented provides a basis for future system development and comparison.

ACKNOWLEDGEMENTS

This work is funded by the UK Engineering and Physical Sciences Research Council in collaboration with Fujitsu Telecommunications Europe Limited.

REFERENCES

[1] ATM Forum, "Audiovisual Multimedia Services: Video on Demand Specification 1.1", Technical Committee, af-saa-0049.001, March 1997.

[2] ITU H262/95, "Information Technology – Generic Coding of Moving Pictures and Associated Audio: Video", July 1995.

[3] ITU-T recommendation I.363 'B-ISDN ATM Adaptation Layer (AAL) Specification, March 1993.

[4] Oracle Video Server, http://www.oracle.com/itv/products/index.html

[5] Sigma Designs – Real Magic Netstream 2 MPEG decoder, http://www.sigmadesigns.com.

[6] Hewllet-Packard Broadband Series Test, www.agilent.com/comms/bsts.

[7] Optibase Moviemaker 200 Basic MPEG hardware encoder, http://www.optibase.co.uk.

[8] A A Webster, C T Jones, M H Pinson, "An Objective Video Quality System Based on Human Perception", Proc. SPIE, 1913 ISBN 0-8194-1146-9, 1993, pp 15-26.

[9] A Basso, I Dalgic, F A Tobagi, C J van den B Lambrecht, "Study of MPEG-2 Coding Performance based on a Perceptual Quality Metric", Proc. Picture Coding Symposium, vol. March 1996, pp 263-268.

[10] ANSI TI.801.03, "Digital Transport of One-Way Video Signals – Parameters for Objective Performance Assessment", ANSI Inc., February 1996.

[11] ITU-T H222.0, "Information Technology – Generic Coding of Moving Pictures and Associated Audio: Systems", July 1995.

[12] S Gringeri, B Khnasnabish, A Lewis, K Shuaib, R Egorov, B Basch, "Transmission of MPEG-2 Video Streams over ATM", IEEE Multimedia, 5(1), January-March 1998, pp58-71.