Subjective Quality of Service Performance of Video-on-Demand under Extreme ATM Impairment Conditions

Javier Zamora, Dimitris Anastassiou, Shih-Fu Chang, and Koichi Shibata

Department of Electrical Engineering
Columbia University, New York, N.Y. 10027, U.S.A.

ABSTRACT: *The Columbia VoD testbed is used to study the impact of extreme conditions of impairments (delay variation, loss and error) on the Quality of Service (QoS) for VoD. The goal of our experiments is to assess the subjective QoS performance perceived by the end-user under different testing scenarios and extreme ATM impairment conditions. These scenarios comprise different combinations of video server mode of transmission, impairments and type of video client used. The results of our experiments allow us to get valuable information for the design of video servers and video clients.*

1. INTRODUCTION

The new service of Video-on-Demand (VoD) will transform the television into a smart, interactive device capable of searching and providing interactive communications. Content will include first run releases, old classics, and other collections of movies and television shows not necessarily found in video stores [1]. However, the widespread acceptance of VoD will depend strongly on the ability to provide a user friendly interface and a refined quality of audio and video to the end-user.

A generic VoD system consists of three components: video server, network and video client. Concepts of semantic transparency and time transparency [2], [3] characterize the QoS performance of video services over networks. Such services require specific constraints regarding the delay, specifically the delay variation or jitter, experienced across the connection, as well as constraints on the rate of errors from video server to video client.

The Columbia VoD testbed [4] is used to study the impact of extreme conditions of impairments (delay variation, loss and error) on the Quality of Service (QoS) for VoD. In previous work, we studied the impact of the video server architecture and of a long distance and multi-switch connection on the end-to-end QoS in VoD services.

Two different scenarios can be considered when a VoD system is analyzed. The first is when the video server and the video client are connected to the same ATM-LAN. In this case, the main source of QoS degradation is the video server since the low utilization of such networks prevents the switches from building up significant congestion. That is the case for the experimental results obtained in the Columbia VoD testbed [5]. The second scenario is when the video server and the video client are interconnected through an ATM-WAN. In this case, the switches along the connection introduce QoS degradation since the utilization of the network is higher than in the ATM-LAN case [6].

In this paper, we concentrate on the subjective QoS viewed on the video client under extreme ATM impairment conditions. The goal of our experiments is to assess the subjective QoS performance perceived by the end-user under different testing scenarios. These scenarios comprise different combinations of video server mode of transmission, impairments and type of video client used. Overall, more than 300 combinations have been tested and taped (about 11 hours of video material). The video material is assessed according to ITU methodology [7].

The results of our experiments allow us to get valuable information for the design of error protection and concealment systems in video server and video client respectively, synchronization recovery mechanisms in the presence of jitter, as well as what kind of traffic contract we have to establish with the network in order to provide a good QoS under cost constraint.

The organization of this paper is as follows. The VoD testbed is presented in Section 2. In Section 3 the video streams used, the impairment scenarios and the methodology for subjective QoS evaluation are presented. Section 4 presents the performance results of our experiments and their implications on video clients. In Section 5 some concluding remarks are given.

2. VIDEO-ON-DEMAND TESTBED

In this section we describe the VoD testbed used in our experiments. Figure 1 shows the portion of the Columbia VoD testbed used in our experiments.

Correspondence to: javier@ee.columbia.edu
K. Shibata is a Visiting Scholar at Columbia University.

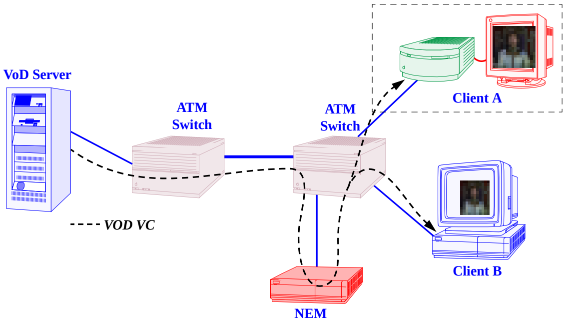

Fig. 1. Columbia VoD Testbed - Network Impairment Emulation.

2.1 Video Server

Our video server is implemented in software using a general purpose UNIX workstation. Specifically, it includes a Silicon Graphics ONYX 6-processor system running IRIX 5.3. The video content is stored on an array of disks in MPEG-2 Transport Streams (TS) [8] format, that is to say, a consecutive number of 188-byte MPEG-2 TS packets. The video server reads these packets and constructs AAL-5 Protocol Data Units (PDU) with a Service Data Unit (SDU) size of N MPEG-2 TS packets ($N = 2, 3, \ldots$). These PDUs are pumped over the ATM network either in Constant Bit Rate (CBR) or Variable Bit Rate (VBR) mode compliant to the ATM Forum and DAVIC recommendations on VoD [9], [10], [11]. Figure 2 shows the AAL-5 adaptation of an MPEG-2 stream into ATM cells.

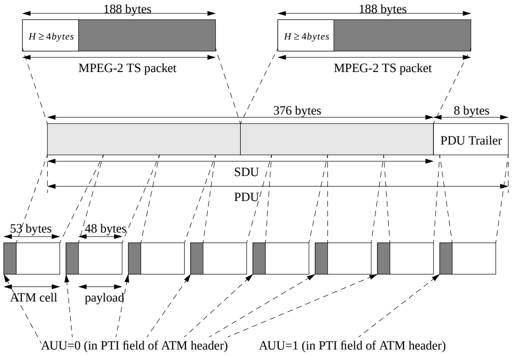

Fig. 2. AAL-5 Adaptation of an MPEG-2 Transport Stream into ATM Cells.

For the network interface, a *ForeRunner* SBA-200 ATM VMEBus Adapter is used. This adapter allows us to control how equispaced cells sent over a PDU period are, in the case of CBR transmission.

2.2 ATM Network and NEM

The Virtual Circuit (VC) established between the video server and the video client traverses two ATM switches and is routed through an ATM Network Impairment Emulator (NEM) [12]. This NEM impairs cells of a specific VC by injecting cell delay variation, cell errors or cell losses. The impairments are generated according to different distributions and statistical values associated with them.

2.3 Video Clients

On the client side, we use two types of video clients. The first type, *Client A*, uses a digital set-top-box with direct ATM connection (DS3 line). The output of the decoder is connected to an NTSC TV monitor. The second type, called *Client B*, uses a personal computer with an MPEG-2 TS decoder and ATM Network Interface (OC3 line) boards.

Client A is an all hardware solution which can decode up to 15 Mbps MPEG-2 TS Streams (CBR and VBR). It only permits the reception of small PDUs ($N = 2$ or SDU size of 376 bytes). This condition is imposed by memory constraints on the video client in order to minimize its cost. This type of video client discards incoming corrupted AAL-5 PDUs (i.e., with wrong CRC field in the PDU trailer).

Client B allows the reception of bigger PDUs ($N < 21$ or SDU size of 3,948 bytes) by expanding the kernel memory of the operating system (Windows NT). Its ATM driver can be programmed to deliver corrupted PDUs to the video decoder. Table I summarizes the main features of each type of client.

TABLE I
MAIN FEATURES OF *Client A* AND *Client B*

Feature	*Client A*	*Client B*
Type	Set-Top-Box	PC
CRC	Check	Check/ No Check
Mode	CBR/VBR	CBR
SDU	$N = 2$	$N = 2 \ldots 21$

3. TESTING SCENARIOS

Once the VoD testbed is described we define the testing scenarios in which the subjective QoS assessment takes place.

3.1 Testing Video Streams

We use two types of MPEG-2 TSs in our testing (Table II). The first type of video streams, encoded from the movie *Robin Hood*, has an average bit rate of 2.9 Mbps which corresponds to VCR quality. The second type of video streams, called *Ice Skating*, has an average bit rate of 5.4 Mbps which corresponds to digital TV quality. For each type of video stream, *Robin Hood* and *Ice Skating*, we have one CBR stream and two VBR streams with 10 and 15 Program Clock References (PCR) per second, respectively. The PCR frequency controls the variability of the rate for MPEG-2 VBR transmission, since the bit rate is constant between a pair of consecutive PCRs [10]. The minimum duration for all streams is about 100 seconds, the streams are looped to obtain an arbitrary video stream length.

TABLE II

MPEG-2 TS TESTING SEQUENCES

Video	Mode	PCR (sec^{-1})	Avg. (kbps)	Peak (kbps)	length (sec)
robin	CBR	10	2,902	2,902	106
ice	CBR	10	5,346	5,346	112
robin	VBR	10	2,987	3,981	104
ice	VBR	10	5,500	10,547	115
robin	VBR	15	3,361	5,227	93
ice	VBR	15	5,607	20,423	113

3.2 Impairment Scenarios

Table III shows the different impairment scenarios used in our testing. All scenarios with the exception of scenario *VBR* use CBR streams. *Client A* and *Client B* are compared in all scenarios except in scenario *PDU SIZE*, which uses *Client B*, and scenario *VBR*, which uses *Client A*. The default SDU size is 376 bytes (8 ATM cells) for all scenarios except for scenario *PDU size*. Each type of impairment generated by the NEM emulates different extreme conditions in the network.

TABLE III

IMPAIRMENT SCENARIOS

CELL ERROR			
Test	Parameter	Test	Parameter
RATE	*robin*	DISTR.	Uniform
	ice		Exponential
ERROR	{1,2,3,4}		Normal
PDU LOSS		**PDU SIZE**	
Test	Parameter	Test	Parameter
DISTR.	Deterministic	CRC	Check
	Exponential		No check
	Normal	BURST	{1,2,4,8}
LOCATION	Last	CDV	Equispaced
	Any		Burst
	Not last	SDU	376,752
BURST	{1,2,4,8}		2632
CDV		**VBR**	
Test	Parameter	Test	Parameter
CELLS	Burst	ERROR	{1,2,3,4}
	Equispaced	LOSS	{1,2,4,8}
DISTR.	Binomial	CDV	Binomial
	Geometric	PCR	10,15

3.2.1 CELL ERROR

Errors in cell payloads emulate bit errors (single or consecutive) caused by noise in transmission media. First, we test the impact of the video stream rate on the video clients under different cell error probabilities. We use a deterministic distribution with error rates ranging from 10^{-7} to 10^{-2}. Second, we study the video client behavior when cell errors are generated according to the uniform, exponential and normal distribution. In the case of the normal distribution, we consider three coefficient of variations (CV): 0.5, 1.0, and 2.0. Finally, we consider the generation of 2, 3 or 4 consecutive bit errors in the cell payload instead of isolated bit errors.

3.2.2 PDU LOSS

Cell losses emulate PDU overflow in the video client and errors in the cell headers, since they provoke the lost of the entire PDU. First, we test different cell loss distributions, such as deterministic, exponential and normal ($CV = 1$). Second, we focus on how the PDU loss is produced. The PDU loss can be generated in the NEM by losing the last cell of the PDU, by losing any cell of the PDU or by losing any cell but the last one of the PDU. Each of these policies have different implications on the clients as we will see in Section 4. We consider 10^{-5}, 10^{-6}, and 10^{-3} cell loss rates of the last cell of the PDU. To obtain equivalent PDU loss rates, for the other two cases, we divided those cell rates respectively by 7 and 8. Finally, we consider the case when the losses are generated in bursts of 2, 4, and 8 PDUs.

3.2.3 CDV

Cell Delay Variation (CDV) emulates cell buffering, network congestion, multiplexing delays and QoS degradation of software video servers [5]. First, we study the impact of how cells are sent over a PDU period (in a burst or equispaced) on the video clients under CDV conditions. Second, we test CDV generation under geometric and binomial distributions.

3.2.4 PDU SIZE

In this scenario, we study the impact of the SDU size (*Client B*) under different impairments (i.e., CDV, losses in bursts). We also test the implications of discarding or not discarding corrupted PDUs by enabling or disabling the CRC checking in the client.

3.2.5 VBR

In this scenario, we reproduce the previous scenarios, but using VBR streams instead of CBR ones.

3.3 Subjective QoS

In order to asses the subjective QoS perceived by the end user, users watch the video signal decoded on the video client in all traffic scenarios. The video material is assessed according to ITU methodology [7]. Under this methodology, the user grades the video material in a 5-value impairment scale, from impairment "Imperceptible" to impairment "Very annoying" (Table IV). This kind of subjective measurements more directly anticipate the reactions of those who might use the VoD system tested.

TABLE IV
Five-grade impairment scale (ITU-R BT.500.7).

Grade	Impairment
5	Imperceptible
4	Perceptible, but not annoying
3	Slightly annoying
2	Annoying
1	Very annoying

4. SUBJECTIVE QoS PERFORMANCE

In this section, we present the subjective QoS performance of the scenarios described in Section 3 in the video client. These scenarios have different effects on the display video, such as frame loss, lip synchronization, scrambled picture, chopped audio, frozen frame, frame repetition, jerkiness, block effect, etc.

4.1 Performance of Scenario *CELL ERROR*

A bit error in an ATM cell payload causes an incorrect CRC checking in the PDU containing the MPEG-2 TS packets (see Fig. 2). A corrupted PDU is discarded in *Client A*, after checking the CRC field. Therefore, a single bit error causes the loss of 8 ATM cells (for SDU size of 376 bytes). On the contrary, *Client B* passes the corrupted PDU to the decoder. The effect on the final QoS for the latter will depend on the information carried by those errored bits (i.e., header information or payload in the Elementary Stream).

When we apply a deterministic distribution with a cell rate of 10^{-6}, a good picture quality is obtained, for both streams and both clients. For a cell error rate of 10^{-3}, it is still possible to recognize the audio but with an intense and annoying effect on the video. We observe a slight improvement on the quality in the low rate stream *Robin Hood* over high rate stream *Ice Skating* for the same cell error rate, because the number of errors per frame is higher for the last stream.

The worst case scenario is given by the deterministic distribution, because of the periodic nature of the error pattern. When we test other distributions (i.e., uniform, exponential, and normal), we observe that good quality is achieved for a cell error rate of 10^{-5}. In this case, we observe a quality improvement on *Client B* over *Client A*. Figure 3 shows two frames of the video stream *Robin Hood* using CBR mode, SDU size of 376 bytes, single cell error with mean 10^{-4} and exponential distribution. Figure 3-a corresponds to *Client A*, while Fig. 3-b corresponds to *Client B*. For this particular example, the video streams were graded as "Slightly Annoying" and "Perceptible, but not annoying", respectively.

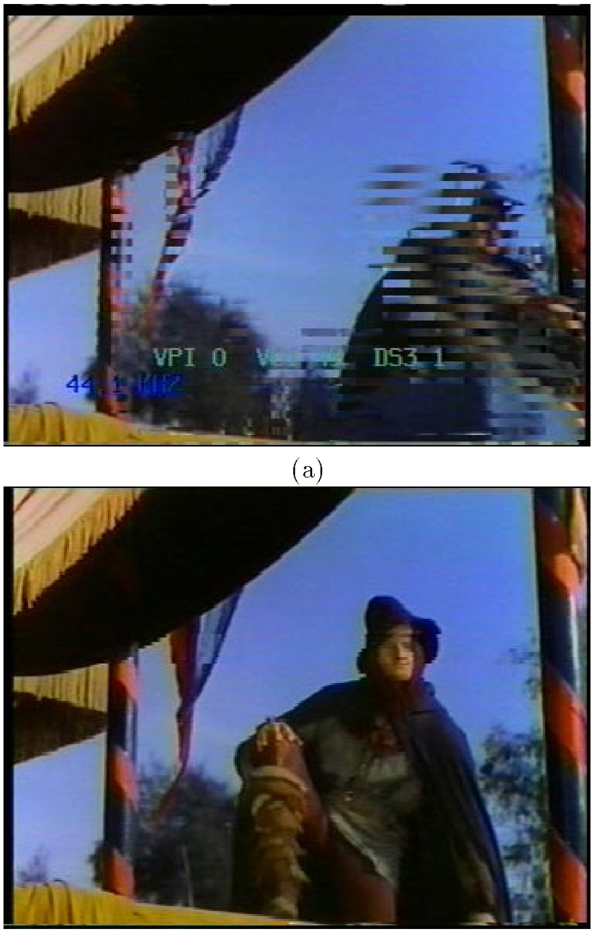

(a)

(b)

Fig. 3. Subjective QoS assessment. *Robin Hood* frame. CBR mode, 376-byte SDU size, single cell error mean of 10^{-4} and exponential distribution. (a) *Client A* "Slightly Annoying", (b) *Client B* "Perceptible, but not annoying".

Different *CV* (i.e., 0.5, 1.0, and 2.0) do not have an effect on the perceived QoS. In addition, consecutive errors in the cell payload (i.e., 2, 3, and 4) cause more visible block effect on the *Client B*, but less frequent than in the single error case. From the above results, we conclude that, in order to guarantee a good picture quality we need at least a bit error rate of 10^{-5} in the physical layer, regardless of how these bit errors are distributed. Moreover, a quality improvement is achieved when corrupted PDUs are passed to the video decoder instead of discarding them.

4.2 Performance of Scenario *PDU LOSS*

Client A and *Client B* detect the boundary of a PDU by checking the third bit (AAU) of the payload type field (see Fig. 2) in the cell header (i.e., "1" for the last cell of the PDU and "0" otherwise).

The NEM can generate a PDU loss in three different ways. First, we can lose the last cell of the PDU, that is to say, the cell with AAU=1 in its header. For instance, for a SDU of 376 bytes (8 ATM cells), *Client A* will detect a PDU of 15 ATM cells but with a wrong CRC field. Therefore the loss of the last cell of the PDU implies discarding two consecutive PDUs. On the other hand, *Client B* will identify and pass a larger wrong PDU to the video client. Second, we can lose any cell of the PDU but not the last one. In this case, the loss is confined to one PDU, since the cell with AAU=1 is preserved. Finally, we can lose any cell of the PDU. This case is a combination of the previous ones.

A PDU loss probability of 10^{-3} provides a "Slightly annoying" picture (i.e., block effect and blanks) for both clients, for all the tested distributions and regardless how the PDU loss is generated. In contrast, a PDU loss probability of 10^{-5} gives a good quality for both clients. However, for a PDU loss probability of 10^{-4}, the performance of *Client B* is much better than *Client A* when the PDU loss is not caused by the loss of the last cell of the PDU. Therefore, the last cell of the PDU has a great impact on the final QoS. Finally, both clients are sensitive to bursts (i.e., congestion situations) of PDU losses. The bigger the burst, the more noticeable the degradation.

4.3 Performance of Scenario *CDV*

When cells are sent in a burst, we have more tolerance to jitter than when cells are sent equispaced. Obviously the implications in the traffic contract with the network are different in each case. *Client B* behaves better than *Client A* when cells are sent equispaced. However, *Client B* has more problems to receive the cells in burst since it uses a software ATM driver.

TABLE V

SCENARIO *CDV* WITH A BINOMIAL DISTRIBUTION, *Client A* AND STREAM *Robin Hood*

Equispaced			Burst		
μ	σ	A	μ	σ	A
0.01	0.0049	5	0.2	0.02	4
0.1	0.010	4	0.25	0.01	5
0.1	0.011	3	0.25	0.02	1
0.1	0.0125	2	0.3	0.005	1
0.1	0.015	1	0.3	0.01	1

For equispaced cells, we obtain good quality when the mean of the geometric distribution is less than 10 μsec. If the CDV is increased to 30μsec, the QoS degrades rapidly.

If cells are sent in a burst at the beginning of the PDU period, we have an "Imperceptible" impairment for a CDV of 100μsec, a "Slightly annoying" impairment for a CDV of 175μsec, and a "Very annoying" for 300μsec. The performance with a binomial CDV distribution is much better than the geometric case. The binomial distribution is a more realistic assumption about the CDV generated through the switches in the connection [13]. Table V shows the evaluation of *Client A* and stream *Robin Hood* for different values of the mean (μ) and standard deviation (σ) of the binomial distribution expressed in milliseconds. The greater the standard deviation, the higher the degradation.

4.4 Performance of Scenario *PDU SIZE*

Some digital set-top-boxes only admit small PDU size ($N = 2$) for memory cost constraints. This is the case for *Client A*. The use of small PDUs has an adverse effect on the performance of the video server [5]. *Client B* permits the reception of larger PDUs by expanding the kernel memory of the operating system. The bigger the PDU size is, the fewer the number of interruptions that occur both in the video server and video client. This fact allows the support of greater number of simultaneous video streams and higher bit rates. In this scenario, we test for $N = 4$ and for $N = 14$, which is the next PDU size, after $N = 2$, without wasted bytes in the mapping process of MPEG-2 TS packets over AAL-5 [5].

For a normal PDU loss distribution with mean 10^{-3} and $CV = 1$, we observe a greater degradation when we use $N = 14$ instead of $N = 4$. This is caused because more TS packets are lost in each PDU. However, for bursts of PDU losses, in other words, for congestion situations in the network, we observe a severe degradation regardless of the PDU size. We also observe a more robust behavior in the presence of CDV, for $N = 14$. For instance, we observe a "Slightly annoying" quality with a binomial CDV distribution with $\mu = 0.1$ msec, $\sigma = 0.015$ msec and stream *Robin Hood* in the case of $N = 4$. In contrast, for $N = 14$, we can achieve similar quality for $\mu = 0.2$ msec and $\sigma = 0.0125$ msec.

4.5 Performance of Scenario *VBR*

In MPEG-2, VBR is constrained to be piece-wise CBR, that is, the bit rate is constant between a pair of PCRs [10]. VBR reception is only supported by *Client A*, since it can handle the reception of PDU of variable size (i.e., the last PDU before a rate change can be shorter than the rest of PDUs). The frequency of PCR controls the smoothness of the traffic pattern. The higher the frequency, the higher the variability of the traffic.

VBR mode is more robust to a PDU loss distribution than CBR mode. In the VBR mode, we obtain an "Imperceptible" impairment for a normal ($CV = 1$) PDU loss distribution with mean 10^{-4} and stream *Robin Hood*. Under the same conditions, we obtain a "Perceptible, but not annoying" impairment for the CBR case. The explanation of this behavior is that in VBR the PDU loss is concentrated on the high bit rates area, where more redundant information exists.

With regard to CDV, VBR mode is more sensitive to CDV than CBR mode. For instance, we need to apply a CDV binomial distribution with $\mu \leq 0.08$ msec and $\sigma \leq 0.01$ msec in order to obtain an "Imperceptible" impairment for VBR *Robin Hood* and 10 PCR/sec. In the case of CBR, we can increase these values up to $\mu \leq 0.25$ msec and $\sigma \leq 0.01$ msec. When we test the VBR *Robin Hood* with 15 PCR/sec, these CDV values are reduced to $\mu \leq 0.06$ msec and $\sigma \leq 0.01$ msec, because of the higher variability of the traffic pattern.

5. CONCLUSIONS

From experimental measurements, we have presented the impact of extreme ATM impairment conditions on the subjective QoS performance of VoD. This paper complements our previous results on video client architecture and long-distance VoD connection. All of them are in the general framework of the Columbia VoD testbed.

We have described the ATM VoD testbed used in this paper, which includes a video server, two ATM switches, a NEM module to generate impairments and two different types of video clients. Five impairment scenarios have been considered, covering cell errors, PDU losses, and CDV. Each scenario represents an extreme condition in the VoD system, such as network congestion, noise in transmission media, buffer overflow, etc. In addition, we have considered both CBR and VBR modes of transmission. The subjective QoS assessment is done following the five-grade impairment scale of ITU-R BT.500.7.

We have shown that 10^{-5} cell error rate in the physical layer is enough to guarantee a good quality video. This quality can be enhanced when corrupted PDUs are passed to the video decoder instead of discarding them.

Preserving the last cell of the PDU (i.e., assigning high priority) is essential to the provision of a good QoS in the presence of PDU losses. Video clients are very sensitive to congestion situations which cause bursts of PDU losses.

Insuring that the CDV experienced by the PDUs remains within certain bounds is critical for the operation of MPEG-2 decoders. The performance degrades rapidly beyond the CDV value that causes dropping of cells in the video client. The subjective QoS is also affected by the way ATM cells are sent over the PDU period.

Furthermore, we have studied the impact of the PDU size on PDU losses and robustness to CDV. Finally, we have shown that VBR is more robust to PDU losses, but it is also more sensitive to the presence of CDV.

The results of our experiments allow us to get valuable information for the design of error protection and concealment systems in video server and video client respectively, synchronization recovery mechanisms in the presence of jitter, as well as what kind of traffic contract we have to establish with the network in order to provide good QoS under cost constraint.

ACKNOWLEDGMENTS

The authors would like to thank the members of the Columbia VoD Testbed and COMET Group for their contribution and suggestions. Columbia University also acknowledges the help of the following companies for contributing equipment to the VoD testbed: C-Cube Microsystems, GTE, Hitachi, HP-IDACOM, NEC and Philips.

REFERENCES

[1] The ATM Forum - IDC Consulting, *ATM Residential Broadband Market Requirements*, Mar. 1996.

[2] M. de Prycker, *Aynchronous Transfer Mode. Solution for Broadband ISDN*, Ellis Horwood Series in Computer Commmunications and Networking. Ellis Horwood Limited, New York, NY, second edition, 1993.

[3] J. Zamora, "Issues of videoservices over ATM," Tech. Rep. 405-95-11, Center for Telecommunications Research, Columbia University, New York, NY, May 1995.

[4] S.-F. Chang, A. Eletheriadis, and D. Anastassiou, "Development of Columbia's video on demand testbed," *Image Communication Journal*, vol. 8, no. 3, pp. 191–207, Apr. 1996, Special Issue on Video on Demand.

[5] J. Zamora, S. Jacobs, A. Eleftheriadis, S.-F. Chang, and D. Anastassiou, "A practical methodology for guaranteeing quality of service for video-on-demand," Tech. Rep. 447-96-13, Center for Telecommunications Research, Columbia University, New York, NY, Apr. 1996, submitted to the *IEEE Trans. on Circuits and Systems for Video Technology*.

[6] J. Zamora, D. Anastassiou, S.-F. Chang, and L. Ulbricht, "Quality of service performance measurement and analysis of video-on-demand in ATM-WAN," Tech. Rep. 467-97-01, Center for Telecommunications Research, Columbia University, New York, NY, Apr. 1997, submitted to *INFOCOM'98*.

[7] Recommendation ITU-R BT.500-7, *Methodology for the Subjective Assesment of the Quality of Television Pictures*, Oct. 1995.

[8] ISO/IEC IS 13818-1 and ITU-T Recommendation H.222.0, *Information Technology - Generic Coding of Moving Pictures and Associated Audio - Part 1: Systems*, June 1994.

[9] The ATM Forum, Technical Commitee, *Audiovisual Multimedia Services: Video on Demand. Specification 1.0*, Jan. 1996, af-saa-0049.

[10] The ATM Forum, Technical Commitee, *Audiovisual Multimedia Services: VBR MPEG-2*, Apr. 1997, saa-ams-vbrmpeg2-02.01.

[11] Digital Audio Visual Council, *DAVIC 1.0 Specification*, Dec. 1995.

[12] Hewlett Packard, *ATM Network Impairment Emulator User's Guide*, Oct. 1995, HP Broadband Series Test System- E4219-92000.

[13] S. A. Wright, D. W. Jarret, D. Kataria, and Y. Viniotis, *Accumulation Algorithms for CDV*, The ATM Forum, Technical Working Group, May 1995, ATM Forum/95-0556.