Objective and Subjective Quality of Service Performance of Video-on-Demand in ATM-WAN

Javier Zamora, Dimitris Anastassiou, and Shih-Fu Chang

Manuscrit received July 25, 1997. Parts of the subjective QoS evaluation under extreme ATM impairments conditions will be presented at the AVSPN'97 workshop, Aberdeen, UK, September 14-15, 1997. Parts of the ATM-WAN performance has been submitted for publication to INFOCOM'98.

The authors are with the Department of Electrical Engineering, Columbia University, New York, NY 10027-6699 USA (e-mail: javier@ee.columbia.edu).

Abstract

The Columbia VoD testbed is used to study the impact of a long distance and multi-switch connection and extreme conditions of impairments (delay variation, loss and error) on the objective and subjective Quality of Service (QoS) for VoD. This paper extends our preliminary results of the Columbia VoD testbed in an ATM-LAN. We compare the QoS performance (i.e. jitter, losses, and subjective parameters) in both LAN and WAN scenarios, and study the implications that these have on video servers and clients. Extensive measurements from real-time tests of VoD services over the ATM-WAN, between Columbia University in New York and GTE Laboratories in Massachusetts, and ATM-NEM connections have been collected. The goal of our experiments is to assess the objective and 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, traffic, 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.

Keywords

Video-on-Demand, End-to-End QoS, ATM, Admission Control, Human factors, Performance Measurements.

I. Introduction

According to recent studies on ATM residential broadband market requirements [1], entertainment-on-demand will constitute 40% of the market opportunity for ATM-enabled applications. Among the possible entertainment applications, 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. 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 Quality of Service (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 an ATM-WAN connection as well as the impact of extreme conditions of impairments (delay variation, loss and error) on the objective and subjective QoS performance for VoD. In our previous work [5], we presented a novel and simple approach for mapping end-to-end QoS in VoD services. Our video server model allows us to consider the video server as the first switch of the connection. This has the advantage of simplifying admission control since the behavior of the video server can be characterized as an additional switch of the connection.

We defined the QoS metrics of the video server and applied them to several experimental results obtained from the Columbia video server in an ATM-LAN. We then identified a comprehensive QoS parameter for the video server (Generic Cell Rate Algorithm), and mapped it from the video server into an equivalent QoS parameter at the video client. We established a relationship between the subjective QoS and the Protocol Data Unit (PDU) loss rate at the video client. Then, we linked the Generic Cell Rate Algorithm (GCRA) performance at the video server with the PDU loss probability (i.e., overflow in the network interface) at the video client. From this relationship we defined the concept of a schedulable region for a video server which models the capacity of the video server, or how many video streams the video server can pump simultaneously, under video client QoS constraints. This model can be connected to network QoS admission control models to have a unified approach for admission control.

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], [4]. In this scenario, the QoS metrics at the output of the video server and at the input of the video client are very similar since the network acts as transparent pipeline.

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.

This paper presents the experimental results of VoD in ATM-WAN scenario.

A multi-switch ATM connection is used to interconnect Columbia University in New York and GTE Laboratories in Massachusetts, covering a distance of 250 miles. The present work emphasizes the QoS degradation originated by the network. In an ATM-WAN scenario, the QoS metrics at the input of the video client can be quite different from the QoS metrics at the output of the video server due to buffering in the network. This scenario reflects the situation of users accessing remote video servers located in other cities or states. Our analysis is based both on objective QoS parameters extracted from real-time capture video traces as well as on subjective QoS viewed on the video client. These results are the first ones published in technical literature that study the impact of such ATM-WAN connections on VoD system design.

In this paper, we also concentrate on the subjective QoS viewed on the video client under extreme ATM impairment conditions, which are emulated by a Network Impairment Emulator (NEM). 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 [6].

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 II. In Section III the video streams used, the traffic and impairment scenarios and the methodology for objective and subjective QoS evaluation are presented. Section IV presents the performance results of our experiments and their implications on video clients. In Section V some concluding remarks are given.

II. VoD Testbed

In this section, we describe the VoD testbed used in our experiments. First, we describe the Columbia VoD testbed architecture. Second, we describe the configuration used for the ATM-WAN connection. Finally, we describe the configuration used for the ATM-NEM scenario.

A. VoD Testbed Architecture

Fig. 1 shows the Columbia VoD testbed architecture, where we identify the three basic components: video server, network and video client.

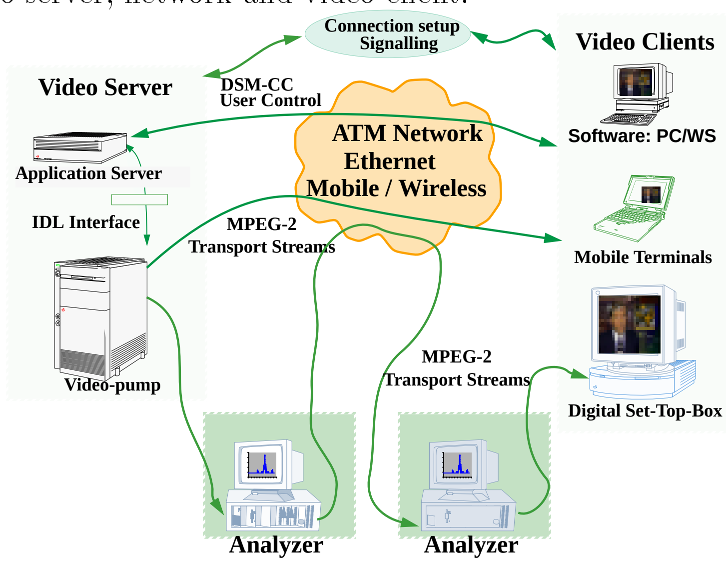

Fig. 1. Columbia VoD Testbed Architecture.

A.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. All the video material is stored in an array of disks in compressed form. We use an MPEG-2 software encoder or real-time hardware encoder to compress the video sequences and prepare both the video elementary streams and MPEG-2 Transport Streams (TS) [7]. A video stream is stored on the video server as a consecutive number of 188-byte video packets (i.e, MPEG-2 TS packets). The video server reads these packets and constructs AAL-5 PDUs 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 [8], [9], [10]. Figure 2 shows the AAL-5 adaptation of an MPEG-2 TS into ATM cells.

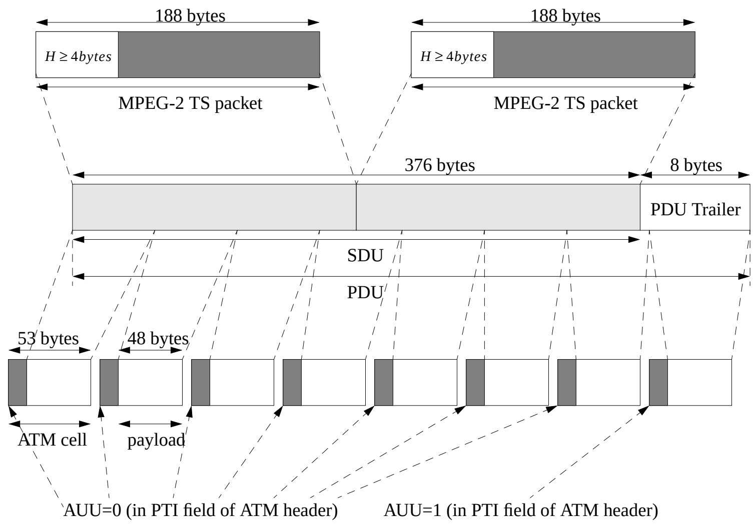

Fig. 2. AAL-5 Adaptation of an MPEG-2 Transport Stream into ATM Cells.

The video sever uses an absolute timer scheme to send out the PDUs according to a specific interarrival pattern. At time t_0 the interarrival time is set to a period, T, between two consecutive PDUs. The first PDU is sent and the pump waits for the timer to expire. The timer, t_e, will only expire when it is equal to $t_0 + T + \Delta t$, where Δt is the time by which the timer overshoots the expiration time. At this point one PDU is sent out and the timer is set to $t_e - \Delta t + T$. From the perspective of the video client, absolute timers enforce a constant average bit rate. This reduce the potential for underflow in the decoder. If one PDU is sent out late with a delay of Δt, the absolute timer will compensate for this delay by trying to send out the next PDU with a period $T - \Delta t$. However, this will also introduce substantial jitter and cell clumping [11].

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, T, are, in the case of CBR transmission.

B. Network

The transport network infrastructure of the Columbia VoD testbed includes ATM and Internet as the core technologies. Users on the campus access the video servers directly through the native ATM network (using AAL5) or via Internet (i.e. over Ethernet or ATM or wireless). An important feature is how the user can search for and select a specific video stream and control the stream. For this purpose we have an application server that communicates with the client through a CORBA-based interface using DSM-CC [12], [13] User-to-User primitives.

A broadband analyzer (*Hewlett-Packard BSTS E4210B*) is used to measure the performance parameters such as the delay jitter at different points in the network. Moreover, the broadband analyzer has a NEM board [14], which emulates network impairment conditions (i.e, cell losses, cell errors and jitter) on real video services in the testbed.

B.1 Video Clients

On the video client side, we have three different types of video clients. The first type, *Client A*, uses digital set-top-box with direct ATM connections (DS3 lines). The output of the decoder is connected to an NTSC TV monitor. The second type, *Client B*, uses a personal computer with an MPEG-2 TS decoder and ATM Network Interface (OC3 line) boards. Finally, we have mobile terminals, using a hardware decoder (*IBM Thinkpad 760CD* laptop) and mobile IP protocol.

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.

C. ATM-WAN Configuration

Figure 3 shows the Columbia-GTE ATM-WAN connection used in our experiments. The public section of the ATM-WAN connection comprises a NYNEX T3 line interconnecting an ATM-DS3 port at Columbia University, located in upper Manhattan in New York

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	$N = 2$	$N = 2 \ldots 21$

City, with an ATM-DS3 port of the MCI ATM Network, located eight miles away in lower Manhattan. The interstate connection, between New York City and GTE Laboratories in Waltham, Massachusetts, covering 250 miles, is done using 2 ATM switches from the USA-wide MCI ATM network. The traffic contract with MCI defines two Virtual Paths (VPs) supporting VBR service, one from Waltham to New York and the other in the reverse direction. Each VP have a Peak Cell Rate (PCR) of 45 Mbps and a Sustainable Cell Rate (SCR) of 16 Mbps [15].

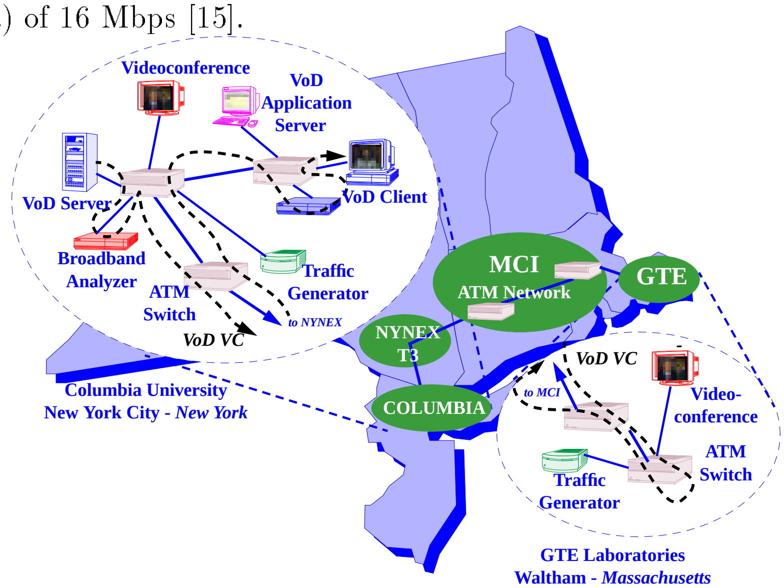

Fig. 3. Columbia-GTE ATM-WAN connection. The VoD virtual circuit is indicated in dotted line.

Since our interest resides in studying the impact of ATM-WAN on QoS performance, we create a loop in the connection to double the number of ATM switches that the MPEG-2 TS stream has to traverse. The total number of ATM switches in one direction is seven: three ATM switches at Columbia University, two in the MCI ATM network and another two ATM switches at GTE Laboratories. We create several Virtual Circuits (VCs) inside both VPs. Each of these VCs transports different type of services (i.e., MPEG-2 video, videoconference, IP traffic, background traffic, etc.).

The VC which transports the MPEG-2 TS streams between video server and video client is called VC_{VoD}. This VC is shown with a dotted line in Fig. 3.

The Columbia video server is connected through an optical 100 Mbps (TAXI line) to a port of a *FORE ASX-100* ATM switch. The VC is routed to the input of a Line Interface (LIF) of the broadband analyzer that monitors the QoS performance at the output of the video server. The output of the LIF is routed to a *FORE ASX-200* ATM switch which interconnects with the NYNEX T3 line. In GTE Laboratories the VC_{VoD} from the MCI ATM network is routed through two ATM switches where it is looped back to the MCI ATM network. Again at Columbia, the VC_{VoD} is routed through the *FORE ASX-200* and *FORE ASX-100* ATM switches and from this latter one to a *NEC M5* ATM switch. The VC_{VoD} is routed to the input of a second LIF of the broadband analyzer that monitors the QoS performance at the input of the video client. Finally the output of this LIF is routed to a DS3 port where the video client (i.e., *Client A*) is connected. *Client A* has no error concealment capabilities since we are interested in studying the impact of QoS degradation (i.e. both in video server and network) on the video displayed by the video client.

Using independent VCs but over the same VP, a videoconference system (*MultiMedia eXplorer*) interconnects both sides. The bit rate of the system can range from 1 Mbps up to 20 Mbps, by adjusting the quantization factor of the JPEG encoder. The videoconference system not only provides a communication channel during the testing, but also an additional cross-traffic source using the same network as the VoD system. Additional traffic generators are located on both sides to generate high utilization traffic conditions in the network. Finally an application server running on a PC/Windows NT is interconnected through IP over ATM, using a VC parallel to VC_{VoD}, to control the VoD system using DSM-CC.

In addition to the described configuration used for the QoS performance measurement and analysis, the ATM-WAN connection was also successfully used for VoD interoperability experiments involving the cross-connection of GTE's video server and Columbia video client and vice-versa [16].

D. ATM-NEM Configuration

Figure 4 shows the portion of the Columbia VoD testbed used in the ATM-NEM configuration. In this case, the VC_{VoD} established between the video server and video client traverses two ATM switches and is routed through the NEM board. The VC_{VoD} is terminated in both *Client A* and *Client B*, since we are interested in comparing the performance of such clients under different impairment scenarios.

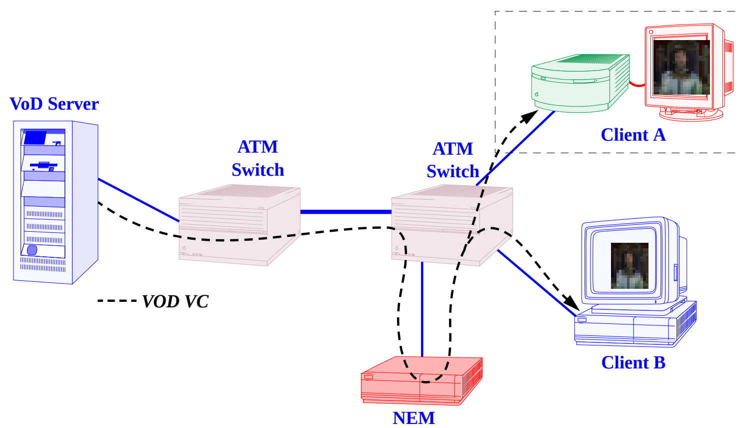

Fig. 4. Columbia VoD ATM-NEM connection. The VoD virtual circuit is indicated in dotted line.

III. QoS Parameters and Scenarios

Now that we have described the ATM-WAN and ATM-NEM configurations, we define the testing video streams used and the traffic and impairment scenarios in which the QoS measurements take place. We then define the objective and subjective QoS parameters, which we analyze from the captured video traces and from the video displayed in the video clients.

A. 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 [9]. 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 Hood	CBR	10	2,902	2,902	106
Ice Skating	CBR	10	5,346	5,346	112
Robin Hood	VBR	10	2,987	3,981	104
Ice Skating	VBR	10	5,500	10,547	115
Robin Hood	VBR	15	3,361	5,227	93
Ice Skating	VBR	15	5,607	20,423	113

B. ATM-WAN Scenarios

In this scenario, we use CBR video streams and the video server pumps SDUs of size 376 bytes (i.e., $N = 2$) since we use *Client A* as video client. Although we transmit CBR video streams, commercial ATM networks adapters do not always send equispaced cells over a PDU period. Due to this fact, the cell rate peak rate can be higher than the PDU peak rate. For this reason, the VP traffic contract is expressed in terms of SCR and PCR.

We define several traffic scenarios to test the VoD system under different network congestion situations, in order to evaluate the resources needed from with the network to provide an acceptable QoS. We have full control over the 5 ATM switches in the GTE and Columbia ATM-WAN sections. The traffic conditions in the MCI ATM-WAN section are not under our control, but reflects the real conditions a user could find in an ATM-WAN public network. These traffic scenarios reflect different level of utilization of resources in the VP. We are interested in testing traffic scenarios where we are not compliant with the VP contract to scenarios where we are compliant but with different levels of utilization. Table III shows for each CBR video stream tested and traffic scenario the utilization in terms of the SCR (16 Mbps) and PCR (45 Mbps) parameters defined in the VP contract. We called these utilizations ρ_{SCR} and ρ_{PCR} respectively. As we can see from Table III, scenario *"High"* is the only one not compliant with the VP contract. However, as we mentioned earlier some violations of the contract can occur in scenarios *"Low"* and *"Medium"* depending how bursty the cells are sent inside a PDU.

TABLE III

VP utilization (ρ_{SCR} and ρ_{PCR}) for each traffic scenario and CBR MPEG-2 TS testing sequence.

Stream	Scenario	ρ_{SCR}	ρ_{PCR}
Robin Hood	*"Low"*	0.12	0.07
Ice Skating	*"Low"*	0.38	0.14
Robin Hood	*"Medium"*	0.57	0.20
Ice Skating	*"Medium"*	0.76	0.27
Robin Hood	*"High"*	1.14	0.72
Ice Skating	*"High"*	2.19	0.78

The traffic scenarios are defined as follows:

"Low" Only MPEG-2 TS from the VoD system is present, no other sources of traffic are generated, neither at Columbia nor GTE sites.

"Medium" Besides the VoD MPEG-2 TS a continuous bi-directional videoconference session, with an average bit rate of 5 Mbps in each direction, is active between Columbia and GTE sites.

"High" The load of the VP violates the traffic contract with the network. In this case, the following sources are present: a VoD MPEG-2 TS, a bi-directional videoconference

session with an average rate of 15 Mbps, Poisson traffic of average bit rate of 5 Mbps generated at Columbia site, and a 9 Mbps CBR generated at GTE site.

"IP Cross-traffic" In this scenario, we study the impact of bursty traffic (i.e., ON-OFF sources) over a VoD session. For this purpose, we send a VoD MPEG-2 TS and at the same time we browse a server containing an image database through a PC which uses IP over ATM that follows the same route as the VC_{VoD} (Fig. 5). When the user retrieves an image, the server sends it at a maximum rate C (45 Mbps); the rest of the time it does not generate traffic.

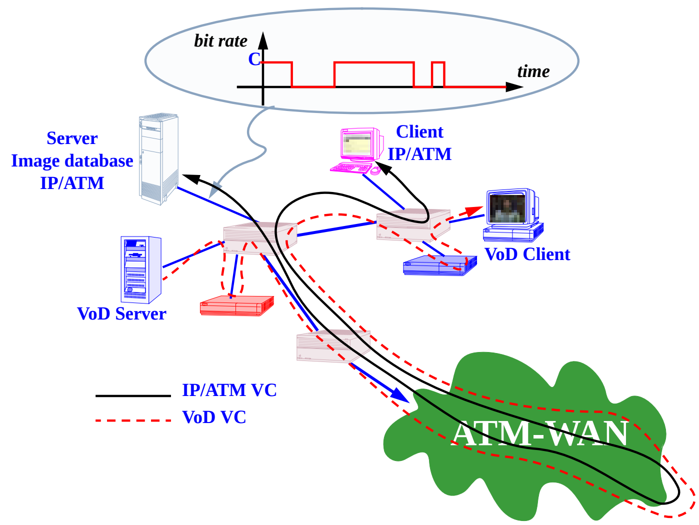

Fig. 5. "IP Cross-traffic" Scenario.

In all the scenarios the video traces are simultaneously captured both at the output of the video server and at the input of the video client. The two LIFs of the broadband analyzer use a common clock in the capture process. We test *Robin Hood* and *Ice Skating* video streams for all traffic scenarios.

C. Impairment Scenarios

Table IV 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 IV

IMPAIRMENT SCENARIOS

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

"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.

"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 IV-B. 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.

"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.

"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.

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

D. *QoS Parameters*

We divided the QoS parameter into the ones derived from the captured real-time video traces, objective QoS, and the ones derived from the evaluation of the video displayed in video clients, subjective QoS.

D.1 Objective QoS

The most important feature in evaluating the performance of a VoD system is to study the traffic pattern both at the output of the video server and at the input of the video client. The deviation from the theoretical video traffic pattern (i.e, CBR) will indicate the QoS degradation. For this reason, we are interested in studying the delay distribution of the PDU interarrival process associated with a video stream. We divide the PDU metrics into the ones associated with delay distribution moments and others associated with traffic control.

From the PDU interarrival distribution, we can calculate three basic parameters associated with the first three moments of the distribution: mean, variance and skewness. The simplest measures are the mean, μ, and the variance, σ^2, which together form the coefficient of variation, $\mathrm{CV}^2 = \sigma^2/\mu^2$, which provides information about the burstiness of the traffic. These measures are based on the first and second moments; for the third moment we use the skewness defined as $\nu = E[(X-\mu)^3]/\sigma^3$, which is a measure of symmetry of the delay distribution and indicates the uniformity of the delay variation.

There are two performance parameters associated with delay variation [15]: 1-point Cell Delay Variation (1-point CDV) and the 2-point Cell Delay Variation (2-point CDV). The 1-point CDV is too sensitive to the momentary fluctuations of the video server as well as the average rate. It also does not provide a stationary measure of the performance of the video server over a long period of time and therefore does not accurately characterize the performance of a video client [5].

The 2-point CDV for PDU k between two measurements points (video server and video client), v_k is defined as

$$v_k = x_k - d_{1,2} \tag{1}$$

where x_k is the absolute PDU transfer delay of PDU k between video server and video client and $d_{1,2}$ is the reference PDU transfer delay between video server and video client. In our case $d_{1,2}$ is the PDU transfer delay of a testing PDU sent over the ATM-WAN network under low network utilization (i.e., *"Low"* scenario).

Another group of metrics is related to traffic control. In this case, we are more interested in studying the compliance of the VoD traffic with a network contract. Example of this type of measure is the GCRA which is equivalent to the continuous leaky bucket algorithm. The GCRA(T,τ) defines a leaky bucket running at a rate of $1/T$ with a tolerance of $100\tau/T$ and is a second-order statistic that measures the burst tolerance. If the tolerance is 0%, we only admit the PDUs that are not violating the network contract, i.e. PDUs whose interarrival time is greater than or equal to the nominal period T. By increasing the value of τ we can admit more bursty PDU traffic patterns. A less bursty source will have more

PDUs admitted since it requires less tolerance in the leaky bucket to accommodate all the PDUs.

It is important to note that two video streams with the same average bit rate can have entirely different GCRA performance depending on how their PDUs are distributed over time. The GCRA captures the second-order statistics of the video stream or how the PDUs are distributed in bursts over time. It provides a measure of the deviation from an ideal VoD system (i.e., no QoS degradation in video server and network) over connection time. For this reasons, the GCRA is the most informative measurement for end-to-end QoS at the video client and a good indicator of both the video client and network performance [5].

D.2 Subjective QoS

In order to assess the subjective QoS perceived by the end user, users watch the video signal decoded on the video client in all the traffic scenarios. The video material is assessed according to ITU methodology [6]. In this methodology, the user grades the video material in a 5-value impairment scale, from impairment "Imperceptible" to impairment "Very Annoying" (Table V). This kind of subjective measurements more directly anticipate the reactions of those who might use the VoD system tested and allow a mapping between subjective QoS and the implications on objective QoS parameters in the video server and network (i.e., GCRA performance, 2-point CDV, etc.)

TABLE V

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

IV. Performance Results

In this section, we apply the QoS metrics defined in Section III-D to the experimental results obtained in the ATM-WAN and ATM-NEM VoD testbed configurations for the different traffic and scenarios described in Section III-B and Section III-C, respectively.

A. ATM-WAN Performance

We first present the performance associated with the moments of the PDU interarrival distribution. Second, we analyze the performance results associated with transfer delay and delay variation. Third, we show the performance related to traffic control (i.e., GCRA). Finally we do a subjective QoS assessment of the video material displayed at the video client.

A.1 Performance of the PDU Interarrival Moments

Table VII shows the first, second and third PDU interarrival distribution moments for video stream *Ice Skating* and for traffic scenarios *"Low"*, *"Medium"*, and *"High"*. The results are constructed with a 90% Confidence Interval (c.i.). The upper part of the table shows the performance at the video server, while the bottom part of the table show the performance at the video client. We are interested in studying the relative degradation between video server and video client, rather than in studying the absolute values at each side which depend on local conditions of the video server. The average PDU interarrival time, μ, is maintained for all scenarios except for scenario *"High"* at the video client. In this case, PDU losses occur inside the network causing a higher average PDU interarrival time. However, additional PDU losses can occur in any of the traffic scenarios. These additional PDU losses are caused by overflow in the video client buffer and depend on the delay variation introduced by the network.

We observe that the variance, σ^2, increases at the video client for all scenarios. A higher utilization causes a higher increase in the video client variance. The use of a software video server causes differences in the variance at the video server side for different scenarios, since the background CPU load can vary from one experiment to another. We observe that for scenarios *"Low"* and *"Medium"*, the skewness, ν, does not change significantly between video server and video client. This means that the shape of their PDU interarrival time

TABLE VI

First, second and third PDU interarrival distribution moments for a CBR video stream *Ice Skating*.

scenario	location	μ (μs)	90% c.i.	σ^2 (μs^2)	CV^2	ν
"Low"	server	557.02	±0.24	3427.44	0.008	88.85
"Medium"	server	557.02	±0.23	2648.70	0.009	99.31
"High"	server	557.02	±0.46	10252.96	0.033	288.06
"Low"	client	557.02	±0.24	3647.34	0.018	88.20
"Medium"	client	557.02	±0.25	3145.94	0.010	77.28
"High"	client	673.24	±1.57	105369.22	0.232	10.23

distribution is maintained; the delay variation introduced by the network has an average around zero. However, in the case of the *"High"* scenario, the PDU losses and the delay introduced by the network congestion provoke that the PDU distribution becomes heavier on the side of long delay.

Fig. 6 shows the first 5000 PDU interarrival times for the video stream *Robin Hood* under scenario *"Medium"*. We observe that the traffic pattern is preserved at the video client but with network jitter component around the mean value of 1037 μs. This jitter component has a mean around zero as can be seen in the probability mass functions (Fig. 7), where the distribution shape is preserved between video server and video client. The side lobes at the video server (Fig. 7-a) correspond to the delay introduced by the inter process communication of the video server [5].

In Fig. 8, the PDU distribution at the video client is shown for the video stream *Ice Skating* under scenario *"High"*. We observe not only a peak around the average rate of the stream, 557 μs, but also subpeaks at multiple values of the average rate, $n \times 557$ μs with $n = 2, \ldots, 5$, which correspond to the event of $n-1$ consecutive PDU losses in the network. The height of the subpeaks provide the value of the probability of $n-1$ consecutive PDU losses (i.e. 10^{-3} for 2 PDU losses). As we can see from Table VII, a sensitive variation of the skewness value is an indicator that PDU losses are occurring in the network.

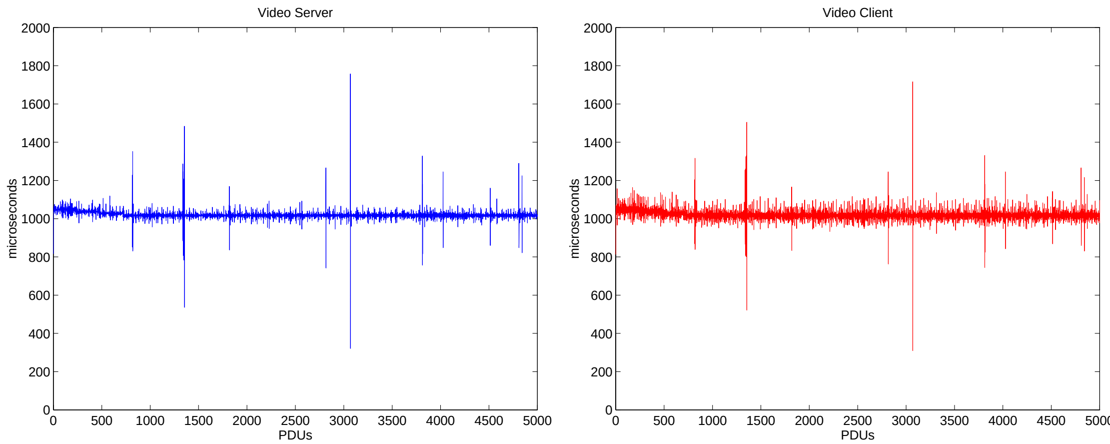

(a) (b)

Fig. 6. PDU interarrival times for the *Robin Hood* video stream under scenario *"Medium"*. (a) Video Server, (b) Video Client.

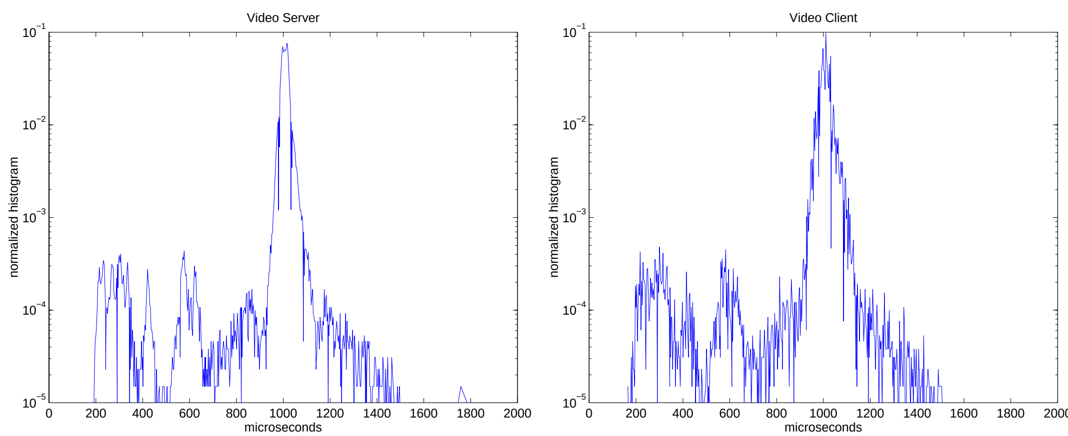

(a) (b)

Fig. 7. Probability Mass Function for PDU interarrival times for the *Robin Hood* video stream under scenario *"Medium"*. (a) Video Server, (b) Video Client.

A.2 Transfer Delay Performance

Once we have analyzed the impact of the network on PDU interarrival distribution at the video client, we study the transfer delay characteristics introduced by the network. Table VIII shows the average, $E[x_k]$, variance, ${\sigma^2}_{x_k}$, and skewness, ν_{x_k} of the PDU transfer delay as well as the average 2-point CDV, $E[v_k]$. No transfer delay data can be calculated for video stream *Ice Skating* and scenario *"High"* because of the large number of PDUs

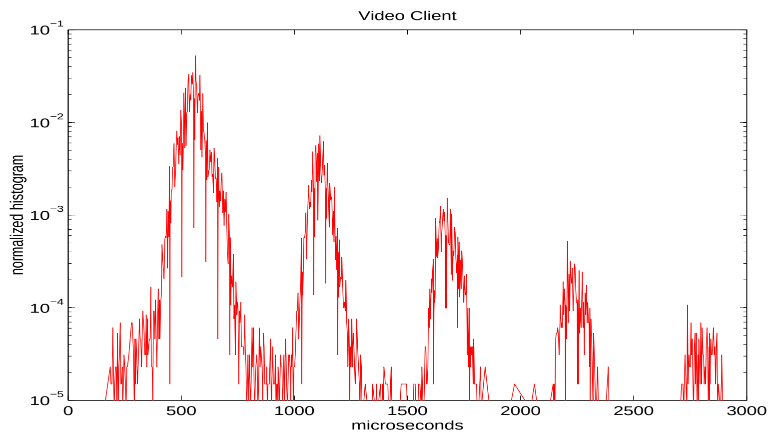

Fig. 8. Probability Mass Function for PDU interarrival times at the video client for the *Ice Skating* video stream under scenario *"High"*.

TABLE VII

FIRST, SECOND AND THIRD PDU INTERARRIVAL DISTRIBUTION MOMENTS FOR A CBR VIDEO STREAM *Ice Skating*.

scenario	location	μ (μs)	90% c.i.	σ^2 (μs^2)	CV^2	ν
"Low"	server	557.02	±0.24	3427.44	0.008	88.85
"Medium"	server	557.02	±0.23	2648.70	0.009	99.31
"High"	server	557.02	±0.46	10252.96	0.033	288.06
"Low"	client	557.02	±0.24	3647.34	0.018	88.20
"Medium"	client	557.02	±0.25	3145.94	0.010	77.28
"High"	client	673.24	±1.57	105369.22	0.232	10.23

losses caused by network congestion. The PDU transfer delay distribution is practically invariant for *"Low"* and *"Medium"* scenarios and for both video streams (Fig. 9). However, the average transfer delay as well as the variance increase for video stream *Robin Hood* and scenario *"High"*. This effect is caused by network congestion that delays the delivery of PDUs as well some isolated PDU losses.

The average 2-point CDV, E$[v_k]$, shows the same trend as the transfer delay results. For scenarios *"Low"* and *"Medium"* the average 2-point CDV is under 50 μs for both video streams. The reference transfer delay, $d_{1,2}$, is equal to 8914 μs, and reflects the

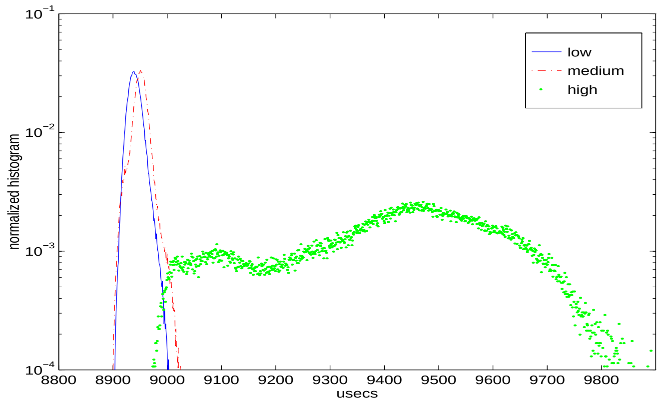

Fig. 9. Probability Mass Function of Transfer Delay for scenarios *"Low"*, *"Medium"*, and *"High"*.

TABLE VIII

PDU Transfer Delay Statistical Moments and Average 2-point CDV.

stream	scenario	E[x_k] (μsec)	$\sigma^2_{x_k}$ (μsec^2)	ν_{x_k}	E[v_k] (μsec)	90% c.i.
Robin Hood	*"Low"*	8928.24	114481.85	-26.28	13.67	±1.54
Robin Hood	*"Medium"*	8937.92	114800.15	-26.25	23.35	±1.54
Robin Hood	*"High"*	9402.13	172074.00	-17.50	487.58	±1.88
Ice Skating	*"Low"*	8914.57	114040.85	-26.31	0.02	±1.53
Ice Skating	*"Medium"*	8948.37	115066.82	-26.25	33.88	±1.54
Ice Skating	*"High"*	n/a	n/a	n/a	n/a	n/a

PDU transfer delay for low utilization in the network (i.e., 14 ATM switches in the loop). Fig. 10 shows the evolution of the 2-point CDV for the first 1000 PDUs sent for video stream *Robin Hood* and scenario *"Medium"*.

Fig. 11 shows the PDU transfer delay for video stream *Robin Hood* and scenario *"IP Cross-traffic"*. The sudden increments in the transfer delay corresponds to retrieval of images from the WWW server. At these times the server sends the image to the client at the link rate, C, blocking the PDUs from the VoD session.

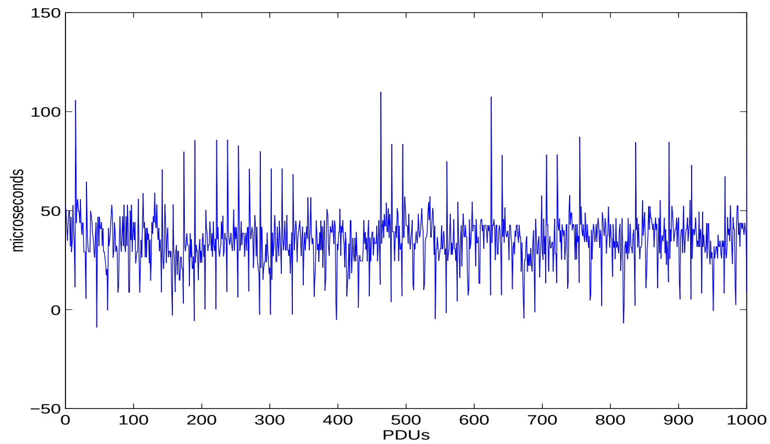

Fig. 10. 2-point CDV for PDU interarrival times for video stream *Robin Hood* and scenario *"Medium"*.

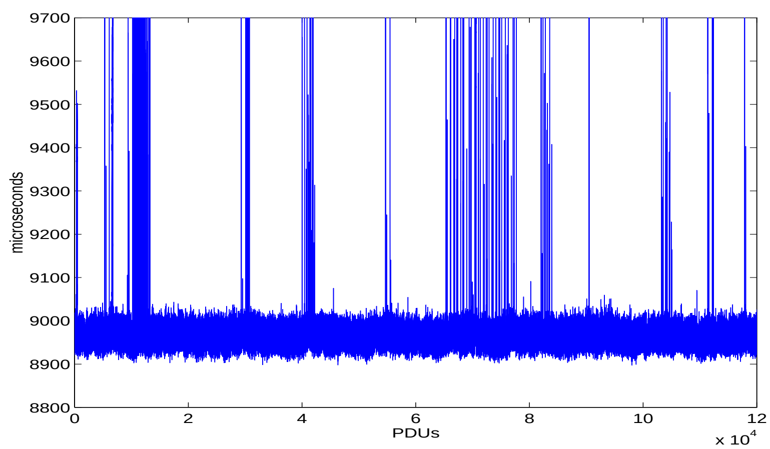

Fig. 11. Transfer Delay for PDU interarrival times for video stream *Robin Hood* and scenario *"IP Cross-traffic"*.

A.3 Traffic Control Performance

Fig. 12 is a graph of the GCRA curves for the video server and scenarios *"Low"*, *"Medium"*, and *"High"*. As the network utilization increases, so does burstiness in the PDUs arriving at the client. Clearly, the scenario *"Low"* has a higher percentage of admitted PDUs than the scenario *"High"*. For instance, for a tolerance of 10%, 97% PDUs are compliant in scenario *"Low"*, while only 89% PDUs are compliant in scenario *"High"*. The burstiness introduced by the network impacts the network interface performance at the client [5]. Fig. 13 shows the probability of PDU overflow at the video client network interface for video stream *Robin Hood* and scenario *"High"*. We observe that the jitter introduced by the network implies more memory needed in the video client for a given PDU probability of overflow.

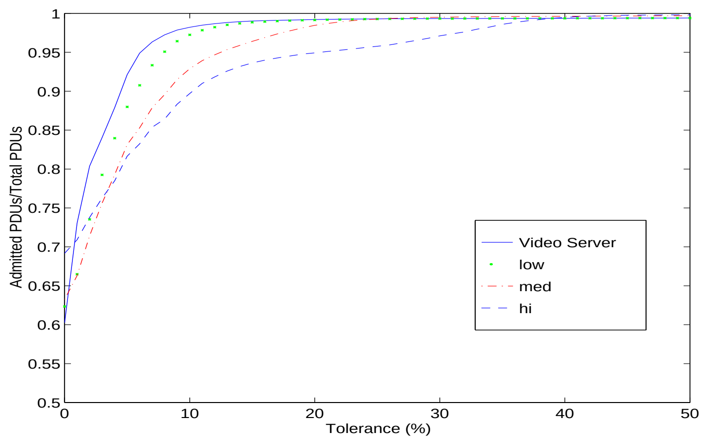

Fig. 12. GCRA for video stream *Ice Skating* for scenarios *"Low"*, *"Medium"*, and *"High"*.

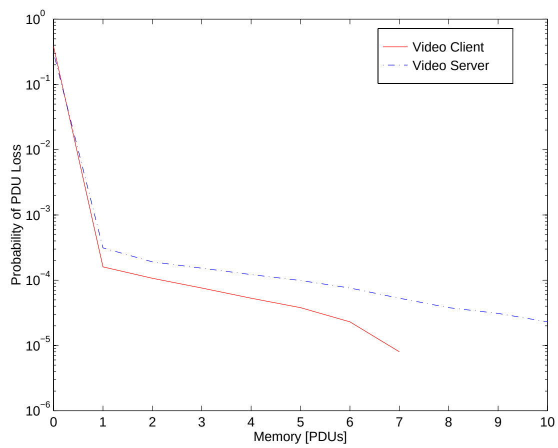

Fig. 13. Probability of PDU overflow at the video client network interface for video stream *Robin Hood* under scenario *"High"*.

A.4 Subjective QoS assessment

Table IX shows the grading of the different traffic scenarios according to the ITU subjective QoS scale [6]. This grading was done by members of our laboratory. The possible QoS degradation observed on the displayed video includes single frame losses, frozen frames, moderate block effect, lip-sync problems, intense block effect, scrambling effect and chopped audio. The impairment effect is uniformly distributed over the VoD session

TABLE IX

Subjective QoS assessment.

scenario	*Robin Hood*	*Ice Skating*
"Low"	Imperceptible	Imperceptible
"Medium"	Slightly Annoying	Slightly Annoying
"High"	Annoying	Very Annoying
"IP Cross-traffic"	Perceptible, but not Annoying	Slightly Annoying

for scenarios *"Low"*, *"Medium"*, and *"High"*. For scenario *"IP Cross-traffic"*, the QoS degradation is only present during browsing periods. In general, a non-continuous QoS is more disturbing to the user than the equivalent continuous average QoS degradation.

Fig. 14 shows two frames from video stream *Robin Hood*. Fig. 14-a corresponds to scenario *"Low"*, while Fig. 14-b corresponds to scenario *"High"* where the QoS degradation (i.e., block effect) is evident.

B. ATM-NEM Performance

In this subsection, we present the subjective QoS performance of the scenarios described in Section III-C in the video client. These impairment 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.

B.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

(a)

(b)

Fig. 14. Subjective QoS assessment. *Robin Hood* frame (a) scenario *"Low"*, (b) scenario *"High"*.

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 15 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 15-a corresponds to *Client*

A, while Fig. 15-b corresponds to *Client B*. For this particular example, the video streams were graded as "Slightly Annoying" and "Perceptible, but not Annoying", respectively.

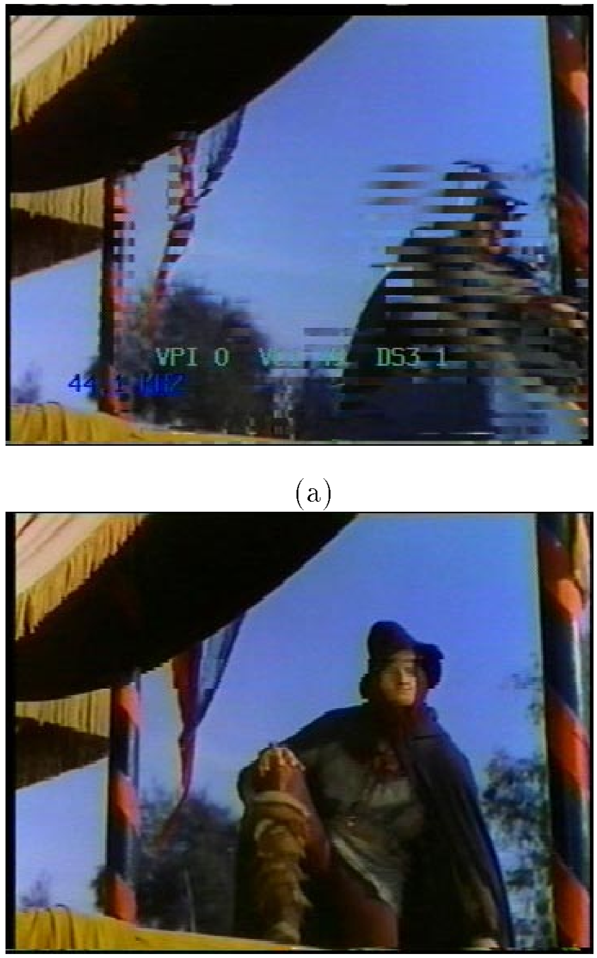

(a)

(b)

Fig. 15. 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.

B.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.

B.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.

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.

TABLE X

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

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 [17]. Table X 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.

B.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.

B.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 [9]. 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.

V. Conclusions

From experimental and novel measurements, we have presented the impact of an ATM-WAN and ATM-NEM connection on the objective and subjective QoS for VoD services.

We have compared these new results with previous results from the Columbia VOD testbed in an ATM-LAN. From that comparison, we have derived the implications of a long distance connection and extreme impairment conditions on the design of video servers and video clients. Objective QoS performance of the real-time video traces, as well as subjective QoS assessment have been considered in our work. We have done this to investigate the interaction among various types of traffic conditions in the network, and the impact of ATM cell loss, error, and delay on the final QoS perceived by the user.

We have described the ATM-WAN VoD testbed used in this paper, which includes a video server connected to a video client through a VP of 14 ATM switches covering a distance of 500 miles. Four traffic scenarios have been considered, covering a range from low utilization in the network to high utilization. We have also considered the effect of intermittent IP cross-traffic. For each scenario, we have captured video traces at the video server and video client. Two different MPEG-2 TSs have been used that correspond to VCR and digital TV quality respectively. We have used QoS metrics associated with the PDU interarrival time distribution, PDU transfer delay, traffic control parameters and subjective QoS assessment grading scale.

We have shown that in scenarios *"Low"* and *"Medium"* the PDU distribution shape is preserved from video server to video client, since the jitter component from the network has zero mean. However, scenario *"High"* for high bit rate video streams introduces significant PDUs losses at the video client. The PDU transfer delay characteristics are quite stable for scenarios *"Low"* and *"Medium"* and both types of video streams, but not with scenario *"High"*. From the traffic control performance, we have concluded that the burstiness introduced in scenario *"High"* has a direct impact on the increase of memory resources at the video client to maintain a given PDU probability of overflow in the network interface.

Furthermore, subjective QoS assessment has shown that low and moderate utilization of the network does not have a significant impact on the final QoS perceived by the user; in other words, we can consider the network as a transparent pipeline. However, high utilization of the network provokes a drastic decrease in the QoS perceived by the user. From these results, we can derive how much we can share the resources in the network without a significant QoS degradation and how much additional resources have to be added in the video client for using an ATM-WAN scenario instead of an ATM-LAN.

We have also described the ATM-NEM 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.

Finally, there is a trade-off between the video server design and the traffic contract established with the network provider. Reducing the QoS degradation level at the video server (i.e., hardware design) can be combined with a less demanding traffic contract (i.e., CDV tolerance) to get the same QoS at the video client. The choice of one option or the other depends both on the cost of network bandwidth and video server equipment.

Acknowledgments

The authors would like to thank the members of the Columbia VoD Testbed Group for their contribution and suggestions. We would also acknowledge the active partnership with GTE Laboratories for their valuable collaboration in the development of the VoD testbed. We would also like to thank the COMET group at Columbia University for the use of their Broadband Analyzer. Columbia University also acknowledges the help of the following companies for contributing equipment to the VoD testbed: C-Cube Microsystems (VideoRISC MPEG-2 Encoder), Hitachi (ATM network interface), HP-IDACOM (MPEG-2 support software and the Network Impairment Emulator board for the broadband analyzer), NEC (ATM switch) and Philips (digital set-to-boxes).

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

[7] 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.

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

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

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

[11] S. Jacobs and A. Eleftheriadis, "Video pump design for interoperability with set top units: The case against small pdus," in *Proceedings, 6th International Workshop on Network and Operating System Support for Digital Audio and Video*, Tokyo, Japan, Apr. 1996.

[12] Object Management Group, *The Common Object Request Broker: Architecture and Specification. Revision 1.2*, Dec. 1993.

[13] ISO/IEC DIS 13818-6, *Information Technology - Generic Coding of Meoving Pictures and Associated Audio - Part 6: MPEG-2 Digital Storage Media - Command and Control (DSM–CC)*, June 1994.

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

[15] The ATM Forum, *User-Network Interface (UNI) Specification. Version 3.1*, Sept. 1994.

[16] Columbia University, "DAVIC New York interoperability experiments: Report and results," Tech. Rep. DAVIC/TC/SYS/96/09/008, DAVIC, Geneva, Switzerland, 1996.

[17] 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.