FEC_PSD: A FEC-aware Video Packet Drop Scheme

Ahmed Mehaoua [1], *Sijing Zhang* [2], *Raouf Boutaba* [3]

[1] University of Versailles - PRiSM Lab.
45 Av. des Etats-Unis
Versailles, France
mea@prism.uvsq.fr

[2] University of Cambridge - CCSR
10 Downing Street
Cambridge, United Kingdom
sz207@ccsr.cam.ac.uk

[3] University of Toronto - ECE Dept.
10 King's College Road
Toronto (On.), Canada
rboutaba@comm.utoronto.ca

ABSTRACT

While delivering real-time video streams across a resource constrained network, loss of frame may be unavoidable. Under such circumstances, it is desirable to design smart video packet drop policies thattake into account the ability of the end terminals to recover totally or partially the corrupted data using forward error correction and error concealment feature. To address this issue, in this paper we introduce the concept of "Cell Drop Tolerance" at the switch node. This concept is supported by a new video slice drop mechanism named FEC-PSD, that adaptively and selectively adjusts cell drop level to switch buffer occupancy, video cell payload type and error correction/concealment ability of the destination. The aim of this proposal is twofold. First, minimizing loss for critical video data, and second, reducing the bad throughput crossing the network. This algorithm is evaluated using MPEG-2 video traces.

Keywords: Packet Video, MPEG, Cell Discard, FEC.

I. INTRODUCTION

In a broadband ATM network where resources such as the network bandwidth and buffering capacity are limited, it is a major challenge to design an efficient video delivery service that can achieve high resource utilization while maximizing user's perceived quality of service (QoS).

Various ATM services have been proposed and evaluated for transporting compressed video: classical Constant-Bit-Rate (CBR) [1], Renegotiated CBR [2], Renegotiate Deterministic Variable-Bit-Rate (RED-VBR) [3], and Available-Bit-Rate (ABR) [4]. Nevertheless, Unspecified-Bit-Rate (UBR) is the true and simplest ATM best effort service available. Since today it is widely and is based on the excess bandwidth in the network with lower usage cost, it is predictable that it will also support a non-negligible part of the multimedia traffic. Therefore, this paper particularly focuses on unidirectional delay-tolerant video applications that can efficiently make use of such simple and low-cost transport service.

In this paper we present a new video packet drop scheme, named "FEC-aware Partial Video Slice Discard" (FEC-PSD), that intelligently cooperates with end terminals (i.e. AAL-5) for improving picture quality and network resource utilization.

The reminder of this paper is as fellows: Section 2 introduces the concept of "Cell Drop Tolerance" and "FEC SSCS Control Block". We describe the new Packet video mechanism in Section 3, and give some performance simulation results in Section 4. Finally, we conclude in section 5.

II. AUDIOVISUAL SSCS CONTROL BLOCK AND CELL DROP TOLERANCE

The ATM adaptation layer is responsible for making the network behavior transparent to the application. AAL5 is currently the most commonly used adaptation layer in industry and can support VBR MPEG2 traffic. However, AAL5 was initially designed simple to carry non-real-time data and lacks to provide reliable connection for multimedia applications. Additional features such as error detection, localization and recovery are required.

In [5], a new Audiovisual Service Specific convergence Sublayer (AV SSCS) for AAL5 that satisfies the above requirements has been proposed. This SSCS is based on Reed-Solomon and Parity Codes [6], and rely on the video packet encapsulation and priority assignation mechanism called DexPAS introduced in [7]. Compared to other mechanisms based on only RS codes and byte interleaving, this approach allows the use of flexible virtual data matrix, called control blocks (CB) that are built at the Service Specific Convergence Sublayer (SSCS) with no modification to the AAL5 Common Part (CP).

0-7803-5796-5/99/$10.00 © 1999 IEEE

At the sender side, a Control Block (CB) is defined as a two-dimensional matrix of P cells column x M rows into which consecutive fixed length Service Specific SSCS PDUs are written row by row (see Figure 1.). An extra single redundancy row is appended at the tail of the matrix that is obtained by XORing the columns at the cell basis. Therefore, a single cell loss per column can be recovered and a "Drop Tolerance can be associated with every CB. This parameter depends on the SSCS PDU length (i.e. number of encapsulated MPEG2 Transport Stream (TS) video packets per SSCS PDU) and may be negotiated at the call set up with reference to the protection level desired by the connection. The lower its value is and higher the recovery power of the AV- SSCS mechanism. The drawback is a proportional increase to the control information overhead. Since the FEC information is obtained using line-by-line XORing method, the data matrix is only an abstract structure and no buffering is required at the sender. The destination process is also pipelined and the correct SSCS PDUs are immediately transmitted to the upper layer without latency.

For more information on AV- SSCS and the Dynamic and Extended Priority Assignation Scheme (DexPAS) see [5] and [7] respectively.

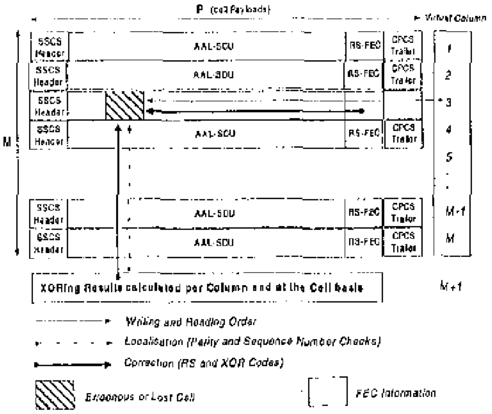

Figure 1 - A Virtual Control Block Matrix

III. THE FEC - AWARE PARTIAL VIDEO SLICE DISCARD

Various video packet discarding techniques have been proposed to minimize user perceived quality degradation during network congestion [8].

In this section, we propose enhancement to the Tail Drop mechanism [9] to take into account Forward Error Correction feature. The new scheme, is named FEC-aware Partial video Slice Discard (FEC_PSD) and performs at video slice level. Our approach is to reduce the number of corrupted slices by assuming that a number 'T' of cells per SSCS Control Block can be recovered by the destination using FEC codes. This parameter 'T' is defined as the "Cell Drop Tolerance" (DT) and corresponds to the maximum number of cells per control block that may be discarded by FEC_PSD before considering the Control Block as definitely lost.

Therefore, unlikely to the simple Selective Cell Discard (SCD) that stops cell discarding cells immediately when the congestion ends, or the Tail Drop (TD) scheme that automatically keeps discard the video packet, FEC_PSD always evaluates the possibility of the destination to recover the partially discarded video slice. It makes its own decision to stop or discard the whole video slice when congestion ends, according to the "Cell Drop Tolerance" parameter and the number of previously eliminated cells.

In terms of performance, SCD under utilizes network resource by forwarding corrupted video slices to the destination, but can offer better picture quality when combined with Forward Error Correction and Error concealment techniques.

In the other hand, Tail Drop better optimizes network resource by avoiding non usable tail packet to cross the network, but reduce the probability of the destination to recover partially lost video data at destination.

FEC_PSD aims to combine the advantageous of both mechanisms, by monitoring the level of cell discards per video slices and compares it to the associated "Cell Drop Tolerance" parameter.

Using this approach, the proposed scheme acts at a finer data granularity, e.g., "SSCS Control Block" rather than video slice, and better preserves entire slices from elimination. Let is remind that a video slice can be divided into one or more "Control Blocks" depending on the spatial resolution of the video sequence. The flexibility proposed by our mechanism can not be achieved without the use of the "Dynamic and Extended Priority Assignation Scheme" (DexPAS) which allows the detection of both video slice and SSCS Control Block boundaries at the cell level.

Let us define a low (resp. high) priority slice as a slice belonging to a low (resp. high) priority video frame. During light congestion, we propose to drop a lower priority slice first rather than delaying it. Then we could assign the buffer space of the dropped slice to a higher priority slice. The proposed approach avoids congestion increase while maintaining the mean cell transfer delay in acceptable value. This proactive strategy is performed gradually by including high priority cells if necessary. As evaluated in [8], the proposed approach can significantly improve the network performance by minimizing the transmission of non-useful video data before buffer overflow. The proposed selective and adaptive partial video slice discard algorithm is highlighted bellow.

FEC_PSD Operation Modes and Fairness

FEC_PSD runs per Virtual Channel and employs four state variables and one counter to control each video connection. Two of them are associated with the slice level and the remaining ones with the control block level (*S_Priority*, *S_discarding*, *CB_dropped*, *CB_discarding*, and *CB_EFCI_marking*).

The use of both *CB_discarding* and *CB_EFCI_marking* indicators allow us to manage losses occurring at subsequent switches and belonging to a control block more efficiently. Indeed, when a block is partially discarded by a switch node, the following switches are not capable to take into account these cell losses to update the associated Drop Tolerance. As a consequence the switches handle erroneous cell Drop Tolerance with adverse effect on algorithm performance. At the Control Block level, the Drop Tolerance can be seen as a cell loss credit distributed and shared among the switches along the communication path.

To address this issue and make implementation easy, we propose to entirely consume the loss credit as soon as a cell loss occurs at a switch. *CB_discarding* is used to ensure that, for every control block, losses are concentrated in a single switch. If cells from a CB tail arrive at a congested node, the use of EFCI bit permits the detection of an early discard and to consider the CB as non-recoverable blocks since a previous switch has used the entire drop credit. In such situation, we propose to commit to the slice level control by entirely dropping the remaining video slice.

FEC_PSD makes use of three buffer thresholds as depicted in figure 2: Low Threshold (LT); Medium Threshold (MT); and High Threshold (HT).

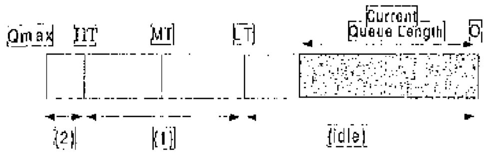

Figure 24 - Buffer Thresholds and FEC_PSD operation modes

The thresholds define three operation modes which in turn limit the distribution of the cell loss within the stream (i.e. Control Block) to four as in figure 3:

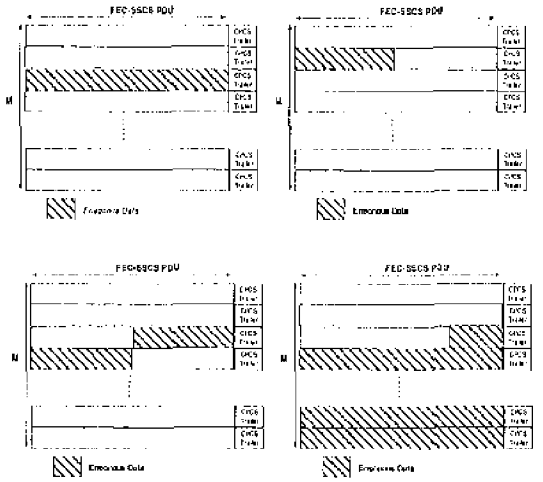

Figure 3 - Cell Loss Distribution per Control Block using FEC_PSD

1. *Idle Mode* : If the buffer queue length (*QL*) is lower than "Low Threshold", for every connection, the cells are accepted and may have EFCI marked, whenever CB_EFCI_marking is activated.

2. *Mode 1*: If the total number of cells in the buffer exceeds 'Low Threshold' but is still below 'High Threshold', for every video connection currently emitting a low priority video slice (i.e. belonging to B- or P-frames), FEC_PSD starts to discard their incoming cells. The discarding is done in respect to the "*Cell Drop Tolerance*" parameters associated with each connection. We propose a fair distribution of the elimination among the targeted connections using round robin scheduling. If the light congestion is subsisting, the algorithm switches to the slice level, and starts to eliminate any incoming low priority cell until receiving an "*End Of Slice*" (EOS) cell. The procedure is done in a round robin fashion in order to guarantee fairness among connections. The last cells (EOS) are always preserved from elimination since they provide indication of the next slice. Cells with higher priority are always accepted in the buffer. This mode stops when '*QL*' falls bellow the "Low Threshold".

3. *Mode 2*: This mode is activated when *QL* exceeds "High threshold". Incoming slices are eligible for discarding regardless to their priority level. The last cell of a Control Block (*EOB*) and video slice (*EOS*) are preserved to avoid error propagation. This is feasible, since usually 10% of switch buffer has been set aside to accommodate the system control and management. This mode behaves like Mode 1 to spread the losses over connections with respect to their "*Drop Tolerance*". It stops when queue length falls below "HT".

The I/PB Resource Management (RM) cells, used to change DexPAS cell priority assignation, are transmitted to all the video sources when Medium Threshold (MT) is exceeded, while IP/B RM cells are send only when *QL* drops below Low Threshold (LT). At the reception of feedback signals, the sources immediately change their operation mode. Consequently, some P- frames may transmit cells with different priority.

Using this adaptive strategy, B- slices are quickly dropped first to reduce buffer occupancy during light congestion, while P- and I- slices are preserved from elimination. If the congestion becomes worse, B- and P- slices are both candidates for elimination, followed by gradually including I-frame cells, if necessary.

IV. PERFORMANCE EVALUATION

Simulation Configuration

In this paper, the proposed FEC_PSD mechanism is integrated with the UBR+ ATM service and its performance is evaluated using simulation. The network topology is shown in figure 4 and composed of two ATM switches, and ten (10) MPEG2 connections crossing the bottleneck link of a capacity of 155 Mbps (OC-3). We evaluate the framework with a LAN configuration, by setting the backbone link to 1 km. All the other link distances, between the source/destination and the switch nodes, are constant and set to 0.2km. The ATM switches are implemented to be non-blocking, output-buffered with finite amount of buffering. Switch buffers size varies from 80,000 to 220,000 cells for both SWITCH_1 and SWITCH_2 in the simulation experiment.

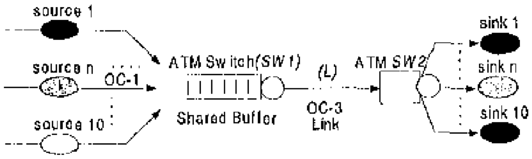

Figure 4 - Network Simulation Model

The video sources generate VBR MPEG2 data at a mean and peak rate of 5.07 Mbps and 20.03 Mbps respectively. A trace file, obtained from Michael R. Izquierdo, IBM Corporation is used as input for the sources. A detailed description of this file could be found in [10].

The level of congestion is monitored through the occupancy of the switch buffers. We assume shared output FIFO buffer, with three congestion thresholds: Low Threshold (LT); Middle Threshold (MT); and High Threshold (HT).

As to the transfer delay, we have the following:

- Propagation delay between the sender and the receiver is set to 0.005 ms, which corresponds to the propagation distance of about 1 km.
- Queuing delay varies from 0 to a maximum value of 0.6 sec, which corresponds to the maximum buffer size of 220,000 cells, when transmitted using 155 Mbps link (OC3).
- The process delay for the sender can be assumed as negligible, due to pipelined data transmission and encoding of the appended data. At the receiver, following FEC processing time for error recovery in AV_SSCS layer is assumed:
  - AV_SSCS (with error): 0.46 ms/slice. (12 x M cell transmission time for 55 Mbps link, where "M" = 5 in most of our cases)
  - AV_SSCS (without error): 0.092 ms/slice. (12 cell transmission time for 55 Mbps link)

The additional processing delay generated at the other layers (e.g. SAR and ATM) is not explicitly modeled. We assume that their contribution to the end to end cell transfer delay (CTD) is relatively constant, and thus can be omitted.

We carried out our simulation with seven switch buffer configurations. For each of them, the same method is applied to determine the values of the three thresholds. *HT*, *MT* and *LT* are respectively set to 0.9, 0.8 and 0.7 of the maximum queue size (*Qmax*), where *Qmax* is set to one of the following values: 80,000, 100,000, 120,000, 140,000, 160,000, 180,000, 200,000 and 220,000 cells.

Data loss are monitored at both ATM cell and MPEG video slice layers. Cell Loss Ratio (CLR) and Slice Loss Ratio (SLR) are evaluated for the aggregate video streams and individually for each cell sub-flows belonging to a certain type of frame (i.e; I, P and B). Regarding, to the application-oriented SLR, it takes into account decoding, e.g., cell loss, and delay propagation, e.g., late cells, constraints. In addition, it also takes into consideration FEC capacity to decide if a slice is usable or not at the destination. In this paper, error concealment at the destination is not considered.

Results Analysis

The forward error correction capability of the video packet delivery service is based on the redundancy data ratio.

In our algorithm, the redundancy ratio is determined by the "M" parameter (i.e. the number of rows per SSCS Control Block). The smaller the "M" value, the larger amount of FEC coding data and higher "Cell Drop Tolerance", and thus the stronger error correction capability. How many redundancy data per FEC frame is needed in order to obtain sufficient performance mainly depends on the cell loss pattern. Larger number of redundancy is required for a strongly correlated cell loss. As mentioned before, the proposed FEC scheme can be easily optimized with respect to amount of redundancy data per FEC frame, and the actual frame size even during a session. As a result, this

FEC scheme is adaptive and will be able to achieve sufficient throughput and latency performance with reasonable transmission overhead.

Figure 5 and Figure 6 show the Cell Loss Ratio (CLR) for the aggregate video stream and for the individual video sub-flows respectively, when different values of "M" are used. Figure 7 and Figure 8 present the same measurements at the video Slice layer (i.e. Slice Loss Ratio), and finally, Figure 9 depicts the Cell Transfer Delay for the aggregate flow.

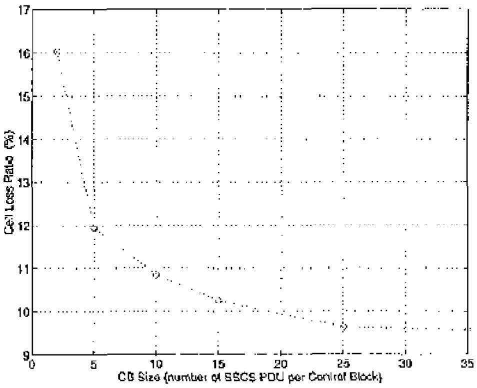

Figure 5 - Cell Loss Ratio vs Data Redundancy Ratio (Aggregate Stream)

For CLR, it is obvious that the loss ratio decreases as redundancy data decreases since fewer overhead is transmitted to the network. In SLR, it is interesting to notice that above some value of "M", the SLR decreases with redundancy. When "M" goes down, the SLR goes up gradually.

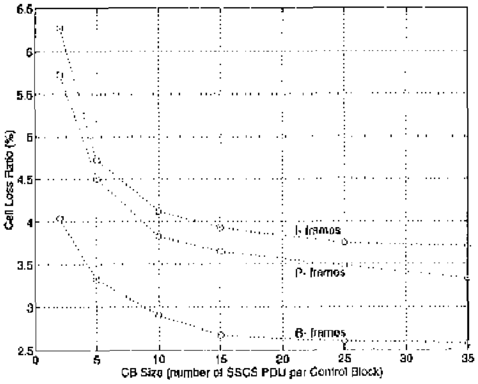

Figure 6 - Cell Loss Ratio vs Data Redundancy Ratio (IPB Sub-streams)

The two effects introduced by appending redundancy data could explain this result. If we append more redundant information, we get more powerful error recovery capability. However, on the other hand, the appended data require more data to be transmitted in the network, hence, consume more bandwidth. This could make the congestion in network worse. As a result, the cell loss ratio increases and so does the slice loss ratio.

Consequently, when "M" is low, the overhead effect dominates. Consequently, SLR decreases with overhead decrease (M increase). When "M" exceeds a certain value, the error correction power provided by redundant data becomes insufficient. Thus, the SLR increases. Therefore, one has to select carefully the redundancy data ratio per FEC frame in order to achieve the best end-to-end performance.

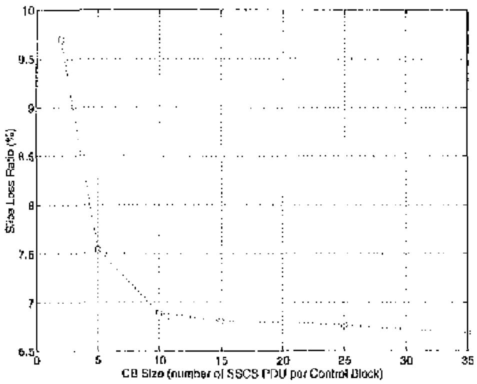

Figure 7 - Slice Loss Ratio vs Data Redundancy Ratio (Aggregate Stream)

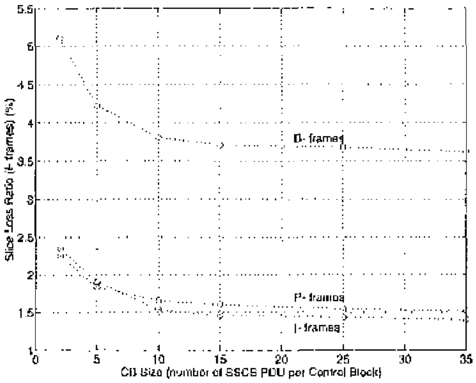

Figure 8 - Slice Loss Ratio vs Data Redundancy Ratio (IPB Sub-streams)

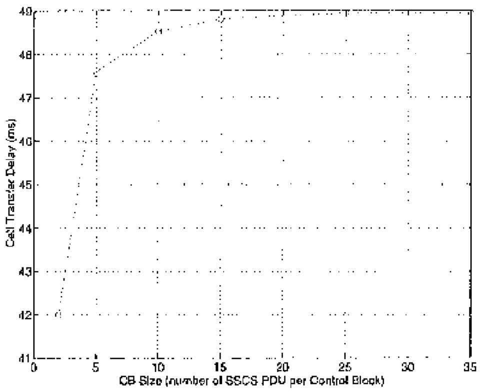

Figure 8 - Cell Transfer Delay vs Data Redundancy Ratio (Aggregate stream)

V. CONCLUSION

Asynchronous transfer of video requires careful integration between the network and the video end systems. A number of issues must be addressed in order to tackle the problem on an end-to-end basis. Among these issues is the design of an ATM adaptation layer (AAL) that provides forward error detection/correction, and a smart cell discard mechanism for managing network congestion. [13].

To cope with this problem, we have proposed and evaluated a novel smart video packet drop scheme called FEC_PSD, that takes into account the specific encoding and stochastic properties of MPEG2 video sources. This mechanism intelligently operates not only according to the network congestion level, but also in respect to the forward error correction ability of the source and destination to recover lost data.

The performance results shown that the proposed FEC_PSD can concentrate the data loss within the B-frames and thus better protect critical video data (e.g. referenced I-and P-frames). Despite of an increase of cell losses, the proposed mechanism demonstrated an improved result at the slice level, i.e. a significant increase of the number of non corrupted slices arriving at destination. A slight increase of the mean cell transfer delay for the aggregate video stream is also experienced because of the overhead introduced by the Forward Error Correction mechanism.

Finally, the enhanced UBR+ service can be easily optimized in terms of both number of redundancy data per video frame and data recovery power (i.e. cell drop tolerance). Consequently, it can achieve sufficient performance with respect to the throughput and latency if the redundancy data ratio is carefully selected.

Acknowledgments

We are grateful to Mickael Izquierdo from IBM Corp., Triangle Research Park, NC, USA, for providing us with the MPEG2 video slices traces.

REFERENCES

[1] ATM Forum SAA SWG, "Audiovisual multimedia services: Video on demand", AF-SAA0049.000, December 1995.

[2] S. Keshav M. Grossglauser and D. Tse, "RCBR: A simple and effective service for multiple timescale traffic", in ACM Computer Communication Review, vol. 25, n.4, pp. 219-230, October 1995.

[3] H. Zhang and E.W. Knightly, "REDVBR: A renegotiation based approach to support delay sensitive VBR video", in ACM/SpringerVerlag Multimedia Systems Journal, Vol.5 n.3, pp. 164-176, May 1997.

[4] T.V.Lakshman, P.P. Mishra, and K.K.Ramakrishnan, "Transporting compressed video over ATM networks with explicit rate feedback control", in IEEE INFOCOM 97, Kobe, Japan, March 1997.

[5] ATM Forum, "Draft proposal for specification of FEC-SSCS for AAL type 5", AF-SAA 0326R2, October 1995.

[6] T. Han and L. Orozco-Barbosa, "Performance requirements for the transport of MPEG video streams over ATM networks", in IEEE ICC '95, Washington, pp. 221-225, June 1995.

[7] A. Mehaoua , R. Boutaba, Song PU, Y. Rasheed and A. Leon-Garcia, "Towards An Efficient Best Effort Video Delivery Service", in IEEE ICC '99, Vancouver, Canada, June 1999.

[8] N. Shacham et P. McKenny, "Packet Recovery in High Speed Networks using coding", in IEEE INFOCOM '90, San Francisco, CA, pp.124-131, June 1990.

[9] A. Mehaoua, R. Boutaba and al., Proposal of An Audiovisual SSCS with Forward Error Correction, In IEEE Symposium on Computers and Communications (ISCC '99), Red Sea, Egypt, July 6 - 8, 1999

[10] A. Mehaoua and R. Boutaba, "Performance Analysis of Cell Discarding Techniques for Best Effort Video Communications over ATM Networks", in Computer Networks and ISDN Systems, Vol.29, issue 17-18, Elsevier Science, pp. 2021-2037, February 1998,.

[11] M.R. Izquierdo and D.S. Reeves, "MPEG VBR slice layer model using linear predictive coding and generalized periodic Markov chains", in IEEE International Performance, Computing and Communications Conference (IPCCC '97), Scottsdale, AZ, February 1997.

[12] A. Romanow, and S. Floyd, "Dynamics of TCP Traffic over ATM Networks", in IEEE JSAC, Vol. 13 No. 4, pp. 633-641, May 1995.

[13] T. V. Lakshman, A. Ortega, and A. R. Reibman, "Variable bit-rate (VBR) video: Tradeoffs and Potentials", in Proceedings of the IEEE, vol 86, No 5, pp. 952-973, May 1998.