Estimation of Video Queueing Performance Using Markov Chains

A. Awad, M. W. McKinnon, R. Sivakumar
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0250

***Abstract*— In this work, we present an algorithm to analytically estimate the queueing performance of MPEG-2 video using goodput as the metric; where we define goodput as the ratio between the number of cells in uncorrupted and correctly displayable frames to the total number of cells that arrive at the queue. The estimation algorithm is used to evaluate three buffer management schemes. The effect of congestion at the output link is also investigated. The algorithm produces a good approximation of the frame goodput metric and closely agrees with the corresponding simulation results.**

I. Introduction

Multimedia applications have become a main part of today's networks and many network protocols are built to support those applications. In order to provide efficient support of multimedia applications such as video, we need to understand its behavior and how to treat video when congestion and loss is encountered. For example, video frames are delay constrained and frames cannot be used if they miss the display time. Also, compressed video formats such as MPEG-2 contain dependencies between frames. Therefore, losing a frame may result in having several frames to be incorrectly displayed due to their dependencies on the lost frame.

In this work, we study and estimate the video quality loss due to congestion and buffer overflow at a network queue using Markov chains. We segment video frames into fixed size packets that we will call cells. This segmentation is used because Variable Bit Rate (VBR) video frames vary in size and so do the slices that compose a picture. Therefore, measuring losses in terms of frames or slices would not be accurate as opposed to fixed size cells as a measuring unit.

Due to the information dependency between video frames, measuring video quality using cell loss metric will neither be suitable nor accurate because cell loss does not capture frame corruption that is caused by error propagation between frames. Hence, we use the frame goodput as measurement metric for video quality because it captures error propagation. We define frame goodput as the ratio between the cells in uncorrupted and correctly displayable frames to the total number of cells in the video stream that arrive to the buffer. The goodput can be studied at the slice level rather than the frame level, but this will require complex analysis and difficult tracking of error propagation, which is not practical for deployment in a network node. Even though our definition of the goodput over-estimates the loss in video quality, it is easier and more simple to apply at network nodes than other more accurate definitions.

We apply the goodput estimation algorithm for three existing buffer management schemes and compare the results with simulation results to compare their goodput performance and to verify the accuracy and flexibility of the estimation process. These schemes are: Tail Dropping (TD), Partial Buffer sharing (PBS), and Triggered Buffer Sharing (TBS). In TD, if a cell arrives to the buffer when it is full, it is discarded along with consecutive cells from the same frame and following frames that depend on the lost frame. In PBS, if the buffer occupancy is below a threshold T, both low and high priority cells are accepted into the buffer, otherwise only the high priority cells are accepted until the buffer is full where all cells are discarded. Cells that belong to I- or P-frames (anchor frames) are marked as high priority cells while those belonging to B-frames are low priority cells. In TBS, two thresholds are defined, $T_H > T_L$. When the buffer occupancy exceeds T_H it will only accept high priority cells. The buffer goes back to accept all cells only when the buffer occupancy decreases below T_L . When considering the system goodput, discarding a cell will result in discarding the frame it belongs to and all frames that is dependent on it[1].

The analytical results are verified with results from a discrete event simulation of the same system. While a simulation based evaluation can provide the same results, analytical estimation techniques are often faster and can be applied in actual systems for call admission control, rate allocation, buffer management schemes and as a base for testing new performance improvement techniques. Estimation techniques can usually be used in predicting the system's performance under several network conditions. However, the more accurate and detailed the estimation is in modeling an actual system, the more complex it becomes. Hence, in order to make the estimation reasonably practical, approximations are made to simplify the estimation process as long as these approximations do not significantly affect the accuracy of the network model we are studying.

In summary, the main contribution of this paper is to present *an analytical estimation of frame goodput for a node buffer with MPEG-2 VBR video traffic using Markov chains*. The frame goodput is used as a performance metric to establish a better representation of the video quality than cell loss metric. A simulation for the system is used to verify the accuracy and cor-

[1]The PBS and TBS are compared in performance in [1].

rectness of the estimation algorithm. Then, the estimation algorithm is applied for three existing buffer management schemes to evaluate and compare their frame goodput performance.

This paper is organized as follows. In Section II, the system description is presented. The concept and theory behind the proposed estimation algorithm is illustrated in Section III. The simulation procedure and a discussion of the results are shown in Section IV, and the conclusion is presented in Section V.

II. System Description

To transmit a video stream in the system, frames are segmented into cells according to the network protocol applied. Figures 1(a) and (b) give the frame order and dependencies during both frame generation and display, and during transmission, respectively. The frames are reordered prior to transmission so that they are transmitted before all dependent frames. This frame reordering creates an overlap between consecutive GOP boundaries, which makes goodput analysis more difficult.

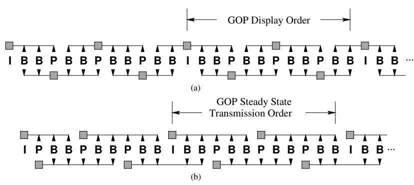

Fig. 1. Dependencies between MPEG-2 frames in the transmission order.

When the buffer overflows during a receipt of a frame, it drops ensuing cells that belong to that frame. Due to the information dependencies among video frames, dependent frames cannot be fully reconstructed when they arrive their destination and are discarded[2]. Hence, when a B-frame causes the buffer to overflow, only that B-frame is discarded. However, when an anchor frame causes the buffer to overflow, other frames in the GOP are also discarded. In our system, we will not empty the buffer of the arrived cells that belong to a discarded frame before the overflow happens; this is an approximation to the goodput in order to accommodate the definition of TD, PBS, and TBS and to simplify the theoretical analysis significantly. Emptying the buffer will also make it difficult to implement in practice. Buffered traffic is then served in a FIFO fashion by being transmitted over the output link to its destination.

III. Queueing Analysis

In order to analyze the system described above, a Markov Chain is embedded at the observation instants at the end of each slot. An arrival event occurs at the start of a cell arrival at the beginning of the slot and a departure event occurs when a cell transmission is completed and after the arrival event occurs. Therefore, if the buffer is full and both a departure and an arrival occur in a given slot, the queue overflows producing a pessimistic evaluation of the system's performance (i.e., maximal loss is incurred). Statistics are gathered at the end of the slot, after the departure completion.

The queueing system with different queue management schemes will result in slightly different models for the system. Due to space limitation, we will only describe the system characteristics using PBS scheme as an example. This queueing system is analyzed using a five dimensional Markov Chain with a state representation (t, w, z, y, x), where:

- t is the slot number within a frame period of time, such that $t \in \{0, 1, 2, \cdots, T-1\}$, which is a function of line speed and cell size in the studied system;
- w is the frame number in a GOP, such that $w \in \{0, 1, 2, \cdots, W-1\}$[3];
- x indicates if the current frame has completely arrived or not, such that $x \in \{0, 1\}$;
- z indicates if the buffer is accepting cells, overflowing on their arrival, or not accepting them at all, $z \in \{$*accept, drop frame, drop GOP, drop GOP**$\}$; the *Drop GOP** state is needed because of the frame overlap between two consecutive GOPs; and
- y is the instantaneous length of the queue, such that $y \in \{0, 1, \cdots, B\}$.

Given the state description above, we notice that the repeating structure of the GOP may be exploited to realize a p-cyclic Markov Chain representing the transition probability between slots of a GOP. The state machine's probability transition matrix has the block form given in (1). $\mathbf{P}$ is a $TW \times TW$ p-cyclic block matrix that contains a top level description of the cyclic state machine mentioned earlier and the total size of $\mathbf{P}$ is $O(WTB)$ for the three different buffer management schemes. It defines the t and w dimensions of the Markov Chain. Each block matrix represents a transition probability matrix between consecutive slots in the GOP. The dimension of the block matrix is equal to the total number of slots in the GOP whether they are occupied with cells or empty. We define the value τ, where $\tau = \tau(t, w) = wT + t$, to reference slots in the GOP when distinction between frame boundaries is not necessary.

The submatrix $\mathbf{R}(t, w)$ describes the system behavior during the transition between the end of slot τ and the end of slot $\tau \oplus 1$.[4] Cells arrive at a constant burst rate (25 or 30 frames per second) according to the video frame structure . Cell bursts are usually followed by unused slots; therefore, a slot can either be part of the burst or part of the unused period. We notice that in any slot, exactly one of two cases can occur: arrival or no arrival. We describe each case with a separate submatrix in the Markov Chain transition matrix due to their different effects on the system. These two matrices are used to define the x dimension in the Markov chain. The first matrix, denoted as $\mathbf{C}_a = \mathbf{C}(t, w, 1)$, represents the case of the system transitioning

[2]Error concealment algorithms can be applied to corrupted frames and some information can be retrieved, however, their effect is not considered here.

[3]The frames in a GOP are numbered according to their transmission order, therefore, frames 1 and 2 are the last two B-frames of the previous GOP.

[4]Note that the symbol $\oplus$ denotes modulo-(TW) addition.

$$\mathbf{P} = \begin{array}{c} \tau_{\downarrow}^{\rightarrow} \\ 0 \\ 1 \\ \vdots \\ \\ \\ TW-1 \end{array} \begin{array}{c} \begin{array}{ccccc} 0 & 1 & & \cdots & TW-1 \end{array} \\ \begin{pmatrix} 0 & \mathbf{R}(0,0) & 0 & & \cdots & 0 \\ 0 & 0 & \mathbf{R}(1,0) & 0 & \cdots & 0 \\ \vdots & \vdots & & \ddots & & \vdots \\ & & & & & 0 \\ 0 & 0 & & & \cdots & \mathbf{R}(T-2,W-1) \\ \mathbf{R}(T-1,W-1) & 0 & & & \cdots & 0 \end{pmatrix} \end{array} \quad (1)$$

from a slot τ to slot $\tau \oplus 1$ given that an arrival occurs during slot $\tau \oplus 1$. Similarly, the second matrix, denoted as $\mathbf{C}_n = \mathbf{C}(t,w,0)$, represents the case of the system transitioning from slot τ to slot $\tau \oplus 1$ given that *no* arrival occurs during slot $\tau \oplus 1$. Using these two matrices, we can write the general structure of $\mathbf{R}(t,w)$ in the following form:

$$\mathbf{R}(t,w) = \begin{array}{c} x_{\downarrow}^{\rightarrow} \\ 1 \\ 0 \end{array} \begin{array}{c} \begin{array}{cc} 1 & 0 \end{array} \\ \begin{pmatrix} \mathbf{C}_a & \mathbf{C}_n \\ 0 & \mathbf{C}_n \end{pmatrix} \end{array} \quad (2)$$

In each $\mathbf{C}(t,w,x)$ submatrix, we can define four different states to describe the buffer's behavior; we will denote them as *accept, drop frame, drop GOP*, and *drop GOP**. In the *accept* state, the buffer performs the normal operation of receiving and enqueueing cells, if adequate space exists. When the buffer occupancy exceeds the frame threshold value, the buffer starts discarding arriving cells. In PBS, the threshold for B-frames and anchor frames are T and B, respectively. The length of cell discard depends on the frame type during which overflow occurs. If the buffer occupancy exceeds T while receiving a B-frame, the buffer state transitions to the *drop frame* state to discard the cells of that frame before it returns to the *accept* state. If the buffer occupancy exceeds B while receiving an anchor frame, however, the buffer transitions to the *drop GOP* state because cells of the rest of the frames in the current GOP are to be discarded.

$$\mathbf{C} = \begin{array}{c} z_{\downarrow}^{\rightarrow} \\ accept \\ D'frame \\ D'GOP \\ D'GOP* \end{array} \begin{array}{c} \begin{array}{cccc} accept & D'GOP & D'frame & D'GOP* \end{array} \\ \begin{pmatrix} \mathbf{D}_{11} & \mathbf{D}_{12} & \mathbf{D}_{13} & \mathbf{D}_{14} \\ \mathbf{D}_{21} & \mathbf{D}_{22} & \mathbf{D}_{23} & \mathbf{D}_{24} \\ \mathbf{D}_{31} & \mathbf{D}_{32} & \mathbf{D}_{33} & \mathbf{D}_{34} \\ \mathbf{D}_{41} & \mathbf{D}_{42} & \mathbf{D}_{43} & \mathbf{D}_{44} \end{pmatrix} \end{array} \quad (3)$$

$$\mathbf{D}_{f_x} = \begin{bmatrix} 0 & 1 & 0 & \cdots & & 0 \\ 0 & \alpha & 1-\alpha & & & \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ & & 0 & \alpha & 1-\alpha & \\ 0 & \cdots & & 0 & 0 & 0 \\ 0 & \cdots & & 0 & 0 & 0 \end{bmatrix} \quad (4)$$

Recall that the GOP transmission order introduces an overlap in the boundaries between successive GOPs because the last two B-frames in a GOP are dependent on the I-frame of the next GOP, and that I-frame has to be transmitted before these B-frames. In the *drop GOP** state the buffer accepts the I-frame of the next GOP and drops the following two B-frames then transitions to *accept* state. In the *drop GOP* state, the buffer drops frames until it transitions to *drop GOP** when it receives a new I-frame. The difference between *drop GOP* and *drop GOP** is that, in the latter, the I-frame that belongs to the next GOP is skipped when dropping the current GOP while in the former the I-frame that belongs to the current GOP is dropped.

Each of the $\mathbf{C}(t,w,x)$ matrices has the general block matrix structure given in equation (3), where *drop* is abbreviated as D'. We will elaborate on $\mathbf{D}_{ij}(t,w,x,z)$ possible forms so that the specific structure of $\mathbf{C}(t,w,x)$ will become more clear.

The matrix $\mathbf{D}_{ij}(t,w,x,z)$ is a $(B+1)\times(B+1)$ matrix; it describes the instantaneous buffer length given the submatrix position in $\mathbf{P}$, which defines the y dimension in the Markov Chain. We use three types of $\mathbf{D}_{ij}(t,w,x,z)$ matrices to define $\mathbf{C}(t,w,x)$ matrices; these matrices are: $\mathbf{D}_{f_x} = \mathbf{D}(t,w,1,0)$, $\mathbf{D}_e = \mathbf{D}(t,w,x,1)$, and $\mathbf{D}_{t_x} = \mathbf{D}(t,w,1,2)$, where x is the frame type.[5] The matrix $\mathbf{D}_{f_x}$ represents the buffer operation under normal conditions when there are cell arrivals and no buffer overflow. $\mathbf{D}_{f_x}$ has the form given in equation (4), where α (in cells/slot) denotes the buffer's service rate[6]. Notice that the number of non-zero rows in $\mathbf{D}_{f_x}$ equals to the threshold value (in cells) for frame x. Since there is a cell arrival, the buffer size either increases or remains constant depending on the occurrence of a cell departure.

The matrix $\mathbf{D}_e$ describes the buffer operation when it is not accepting cells either because the buffer is discarding arriving cells due to an earlier overflow, or because there are no cell arrivals at all. Therefore, $\mathbf{D}_e$ has the form given in equation (5). Since the buffer is not accepting new cells, the buffer size either remains constant or decreases depending on the occurrence of a cell departure.

$$\mathbf{D}_e = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ \alpha & 1-\alpha & 0 & & \\ 0 & \alpha & 1-\alpha & & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \alpha & 1-\alpha \end{bmatrix} \quad (5)$$

Finally, the matrix $\mathbf{D}_{t_x}$ describes the buffer operation at the

[5] Notice that $\mathbf{C}$(t,w,0,0) and $\mathbf{C}$(t,w,0,2) cannot occur because when there is no cell arrival, the buffer neither accepts a cell nor overflows.

[6] The system is analyzed with a probabilistic (statistical) service rate α.

instant of an overflow when the buffer occupancy exceeds the frame threshold value. $\mathbf{D}_{t_x}$ has the form given in equation (6). As in $\mathbf{D}_{f_x}$, the number of *zero* rows in $\mathbf{D}_{t_x}$ equals to the frame threshold value. Since the arriving cell is discarded, the buffer size remains constant or decreases. $\mathbf{D}_{f_x}$, $\mathbf{D}_e$, and $\mathbf{D}_{t_x}$ are used to specify the entries of $\mathbf{C}$ for each slot. In TD, the B-frame threshold value in the previous analysis becomes B. In TBS, an additional state is introduced to model the operation of the two thresholds. Due to space limitations, we omit TBS analysis description.[7]

$$\mathbf{D}_{t_x} = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \cdots & \alpha & 1-\alpha & 0 \\ 0 & \cdots & 0 & \alpha & 1-\alpha \end{bmatrix} \quad (6)$$

In summary, the slot number is defined using **P** and the arrival process for each slot is defined using **R**. The system behavior depends on the occurrence of cell arrivals. If there is a cell arrival ($\mathbf{C} = \mathbf{C}_a$) in the current slot, the possibility of buffer overflow will be indicated by the buffer occupancy that is preserved in $\mathbf{D}_{ij}$. If an overflow occurs, the system's next state and the time the system spends in that state is governed by the type of the current frame (i.e. the current slot number). The system transitions back to *accept* state after frame referencing (if any) terminates. The dropping algorithm is mainly implemented by the definition of **C** for each slot because it contains the state space that governs cell loss.

IV. Simulation Results

Segmented frames arrive to the buffer where the order of frames is governed by a twelve-state machine that repeatedly produces an $IBBPBBPBBPBB$ pattern. This way the deterministic ordering between the frames is preserved within GOPs to resemble better source modeling accuracy. Upon arrival, if space permits, segments are buffered in a FIFO fashion. If a frame segment overflows the buffer, it is discarded along with all successive cells that contain information dependent on that frame. Buffered segments are served at a constant rate.

In this work, frames are sized according to lognormal distributions [3] that are bounded by a maximum frame size. The mean frame size is selected according to the average of statistics for real MPEG-2 coded movies shown in [4][8]. The I-, P-, and B-frames are set to have normalized mean frame sizes, $\mu_I : \mu_P : \mu_B$, of $1 : 0.3 : 0.13$ and relative standard deviations, $\sigma_I : \sigma_P : \sigma_B$, of $1 : 0.76 : 0.32$, respectively. The ratios between different frame statistics reflect the correlation between frames in MPEG-2 video streams. The inter-GOP correlation, however, is not implemented in this traffic model since the variation in the total GOP size is greatly affected by scene changes.

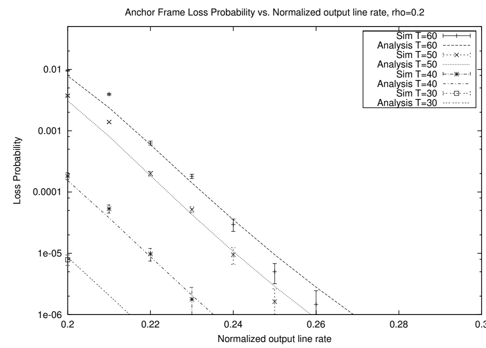

Fig. 2. PBS cell loss in anchor frames.

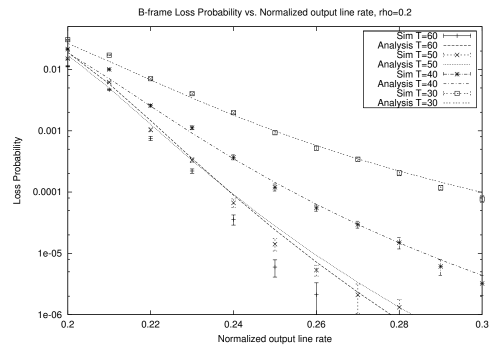

Fig. 3. PBS cell loss in B-frames.

Scenes may last for a relatively long time, which minimizes the effect of a finite buffer in improving the system performance or absorbing traffic bursts at GOP level. Maintaining the above average frame size ratios limits the system's offered load[9] to 23% because, even when μ_I is maximal, $\mu_B = 0.13\mu_I$. Even though the maximum load is low, video streams are time sensitive, which cause the input arrival rate to reach the peak rate during a cell burst (frame). This arrival pattern can cause buffer overflow when the output link is congested.

A transition matrix is built for each buffer management scheme and then used to estimate the frame goodput of the system. As previously mentioned, goodput is defined as the ratio of the cells belonging to correctly displayable frames that are transmitted by the server on the system's output link to the total number of cells . Furthermore, we define $P_{loss} = 1 - goodput$; we choose these definitions of goodput and loss probability because they are more representative measures of the integrity of the end user's video stream quality than other measures [5], [6].

We compare the results of our analysis with the corresponding simulation results. In Figures 2 – 5 we plot cell loss probability versus normalized output line rate[10] for a buffer size of 60 cells and a load of 20%. The estimation and simulation results closely agree.

The data in Figure 2 shows the loss due to dropping anchor

[7] See [2] for a detailed estimation analysis for TD, PBS, and TBS.

[8] Results were prepared by Oliver Rose of the Computer Science Institute at University of Wuerzburg

[9] Offered load = (full slots/Total slots) at the input link.

[10] The normalized output rate is the ratio between the output line rate and the input line rate. It can be looked at as a measure of link congestion.

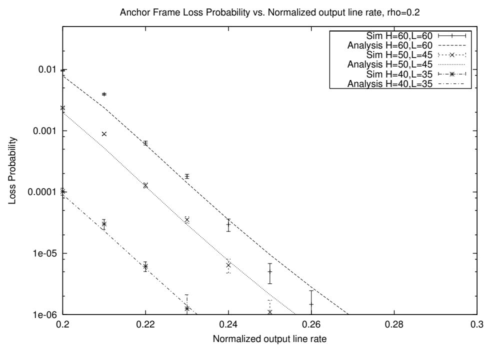

Fig. 4. TBS cell loss in anchor frames.

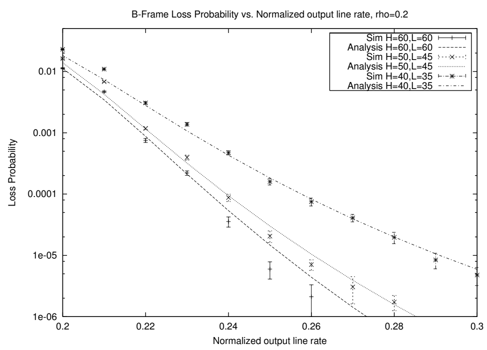

Fig. 5. TBS cell loss due to B-frames.

frames for the PBS case. We notice that threshold value change significantly affects the loss. As the threshold value decreases the loss decreases. The results are compared with the TD case which is shown to encounter the highest loss. The opposite is shown in Figure 3 where the loss due to dropping B-frames. In this case, the curve representing the TD case performs the best. Similar figures are shown for the TBS case in Figures 4 and 5.

The results for the PBS and TBS shows higher total loss probability than the TD case. This is due to the less buffer space available for B-frames. Other dropping schemes more efficient than TBS and PBS are needed for video traffic. The loss probability in anchor frames can be reduced without increasing the loss probability in the B-frames if the output bandwidth is efficiently utilized. This can be achieved by introducing multiple dynamic thresholds that are a function of the buffer size, output bandwidth and a knowledge of the incoming frame sizes.

V. CONCLUSION

The goodput of a buffer with a video traffic input source is investigated in this work. An approximate method is proposed to estimate the buffer behavior under highly correlated traffic source and this method was verified using simulation. The simulation results and the estimation analysis results closely agree.

The estimation algorithm flexibility and accuracy is shown using three cases of buffer management schemes. It is also shown that anchor frame protection using threshold-based priority buffer management schemes increases the number of total frame loss but controls anchor frame loss.

There has been some similar work on queueing analysis for VBR and priority assignment network traffic, however, these studies did not consider the video traffic dependency structure, subsequent packet dropping, and deterministic modeling of video sequences. In [7]–[8], Markov chains were used in queueing analysis with focus on priority queueing or bursty VBR traffic, however, usefulness of transmitted packets were not recognized. In [9]–[10], the authors did not study subsequent dependency dropping of cells within a frame or frames within a GOP. The authors in [11] studied the goodput of CBR video traffic and presented simple online dropping algorithms. The study provides neither a goodput analysis nor an estimation for the studied system.

This estimation algorithm can be applied in network nodes that perform call admission control, rate allocation, and video quality evaluation under different network congestion conditions. It can either be applied online and utilize the network condition feedback or can be used to frequently produce and update lookup tables that can be utilized for the the above applications. The next step in this work is to design appropriate dropping schemes to enhance the video quality. Another area of interest is studying the goodput of multiple video streams entering a network node.

REFERENCES

[1] Pedro Cuenca, Antonio Garrido, Francisco Quiles, and Luis Orozco-Barbosa. Performance Evaluation of Cell Discarding Mechanisms for the Distribution of VBR MPEG-2 Video Over ATM Networks. *IEEE Transactions on Broadcasting*, 44(2), June 1998.

[2] A. Awad, M. W. McKinnon, and R. Sivakumar. Goodput Analysis of Network Access Node Priority Buffers Carrying MPEG-2 Video Source. *Unpublished Manuscript, http://www.cnd.gatech.edu/˜ashraf/publications/Report.ps*, 2002.

[3] M. Krunz, R. Sass, and H. Hughes. Statistical Characteristics and Multiplexing of MPEG Streams. *INFOCOM'95*, 2:455–62, April 1995.

[4] M. Frey and S. Nguyen-Quang. A Gamma-Based Framework for Modeling Variable MPEG Video Sources: The GOP GBAR Model. *IEEE/ACM ToN*, 8(6):710–9, December 2000.

[5] Olivier Verscheure, Pascal Frossard, and Maher Hamdi. Joint Impact of MPEG-2 Encoding Rate and ATM Cell Losses on Video Quality. *GLOBECOM '98*, 1:71–6, Nov. 1998.

[6] P. Frossard and O. Verscheure. AMISP: A Complete Content-Based MPEG-2 Error-Resilient Scheme. *IEEE Transactions on Circuits and Systems for Video Technology*, 11(9):989–98, Sept. 2001.

[7] C.G. Kang and H.H. Tan. Queueing Analysis of Explicit Priority Assignment Partial Buffer Sharing Schemes for ATM Networks. *INFOCOM*, 2:810 –819, April 1993.

[8] K. Kawahara, K. Kitajima, T. Takine, and Y. Oie. Packet Loss Performance of Selective Cell Discard Schemes in ATM Switches. *IEEE Journal on Selected Areas in Communications*, 15(5):903–13, 1997.

[9] C. H. Tan and L. Zhang. Effects of Cell Loss on the Quality of Service for MPEG Video in ATM Environment. *IEEE Singapore International Conference on Communications and Networks*, pages 11–15, July 1995.

[10] O. Rose. Discrete-time Analysis of a Finite Buffer with VBR MPEG Video Traffic Input. *Proceedings of the 15th International Teletraffic Congress- ITC 15*, 2:413–22, 1997.

[11] J. R. Li, X. Gao, L. Qian, and V. Bharghavan. Goodput Control for Heterogeneous Data Streams. *NOSSDAV*, June 2000.