An Extended Priority Data Partition Scheme for MPEG Video Connections over ATM

Ahmed Mehaoua[1,2], Raouf Boutaba[1] and Guy Pujolle[2]

[1] Computer Research Institute of Montréal,
1801, McGill College, Montréal (Qc), Canada
Email: {amehaoua, rboutaba}@crim.ca

[2] University of Versailles,
45 Av. des Etats-Unis, 78000 Versailles, France
Email: Guy.Pujolle@prism.uvsq.fr

Abstract

Transmission of compressed video over ATM networks requires efficient data priority partition techniques. In association with intelligent cell discard schemes, these techniques aim to minimize loss probability of critical information in the situation of congestion. In this paper, we propose a new video-oriented priority data partition mechanism named Extended Priority Assignation Scheme (ExPAS). This mechanism better uses the cell header and allows the definition of up to three service classes per connection. To evaluate its performance, we have also designed an adaptive cell dropping scheme which takes benefits of the new features. In comparison with previous priority partition techniques based on CLP mechanism, ExPAS with the Adaptive Selective Cell Discard (A-SDC) scheme shows better results in minimizing cell losses of Intra-coded frames.

Keywords : ATM, VBR MPEG, Priority, QoS.

1. Introduction

Next generation video services, such as high definition TV broadcasting and video-on-demand, will work on the basis of ATM networks and will widely use MPEG compression standards to save network resources.

MPEG video coding relies on two basic techniques : block-based Motion Compensation (MC) for reduction in temporal redundancy and Discrete Cosine Transform (DCT) for spatial redundancy. Video data are organized in a hierarchical format in order of increasing spatial size [1]. Namely, 8x8 pixels Block, 16x16 pixels Macroblock, Slice, Frame, Group of Pictures and Sequence.

Frame or picture is the basic unit of display. Three types of frame exist in the standard which differ from the used coding method : Intra-coded (I) picture, Predictive-coded (P) picture and Bi-directionally predictive-coded (B) picture. For three main reasons, I-pictures are essential and have to be preserved from corruption during transmission.

Firstly, interactive video applications requiring random access into the video sequence, like fast-forward or fast-reverse playback, will frequently use I-pictures.

Secondly, since I-frames are encoded using information only from themselves, they intend to assist the decoding system for resynchronization in situation of long error bursts.

Finally, a non-available reference picture (I- or P-frame) leads on perceptible picture degradation. With error propagation, I-frame impairments will affect all the subsequent frames on the same group of picture (GOP). Similarly, the impairment of P-frames will affect the following P- and B-frames until the next I-frame. Only B-frame impairments have no adverse effects on other frames.

To protect efficiently essential frames from loss during overload periods, we propose a new accurate cell discrimination strategy and an associated adaptive cell discard scheme. Both mechanisms take into account the specific data structure of MPEG encoded video.

The proposed Extended Priority Assignation Scheme (ExPAS) allows the discrimination of up to three prioritized cell flows within every single video connection. These cell flows are associated with MPEG frame types (e.g. I, P and B).

Since these frames have different impact on decoding and display process, we propose to eliminate B-frame cells (e.g. B-cells) first, while preserving I and P-cell during light congestion. If the congestion worsens, the other frames are progressively candidate for discarding.

This paper is organized as follows. In section 2, we review the main priority data partition techniques employed with MPEG video. The third section is devoted to the multi-level priority assignation mechanism and the associated discarding scheme. In section 4, we present the simulation model and discuss the results in section 5. Finally, concluding remarks are given in section 6.

2. Video Data Partition Techniques

2.1 Block Level Partition

Since human perception is less sensitive to low frequency components of a video signal, subsequent blocks are transformed into the frequency domain using DCT. Each transformed block may then be partitioned into an essential layer (comprising the lowest frequency DC coefficient), and an enhancement layer (consisting of the set of high frequency AC coefficients). The information contained in the essential layer is packetized and transmitted at high priority, which ensures a guaranteed quality of service. Information in the enhancement layer is transmitted at a low priority, which provides only a best effort service. The cell loss priority (CLP) bit in ATM headers is used to provide a two-level cell priority mechanism within single channel.

In [2], this approach is applied at the macroblock layer. Similarly to the block-based approach, the DC value for each of the 8x8 block are assigned to the high priority (HP) stream. The macroblock header, and the motion vector in case of P-frames are also included in the HP stream. For the remaining 63 DCT coefficient of each block, the authors define a parameter β which specifies the number of AC coefficients that are to be placed in the HP stream. The remaining ($63-\beta$) coefficients are transmitted in a low priority (LP) stream. To permit the regeneration of the original bit stream by the destination, the macroblock address is joined to the LP information.

In [3], a connection-level prioritization approach is evaluated to transport a 2-layer MPEG-2 video sequence over different ATM service classes. The scheme uses static data partitioning between two connections by means of a Load Balancing factor (LBF). The virtual connections are associated with guaranteed service class (i.e VBR-rt) and best effort service class (e.g. ABR) to respectively carry the base and the enhancement layers.

The drawbacks of these techniques are the added complexity and the special devices required at the destination end points to synchronize and recover the original video stream.

2.2 Frame Level Partition

Data partition with priority assignation can simply be implemented at the frame layer. The cells belonging to following frames are set to different priorities. For instance I-frame cells may have the highest priority over P and B-frame cells. In [4], two static priority partition strategies are proposed and use the cell loss priority (CLP) mechanism :

- Static I/PB priority partition : in this method, I-frame cells are considered with a high priority and have the CLP bit set to '0', while P- and B-frame cells are assigned a lower priority with CLP flag set to '1'. If a congestion occurs, cells from P- and B-frames are discarded first.

- Static IP/B priority partition : in this variant, I- and P-frames cells are assigned a high priority and better protected from elimination. Here, only B-frames cells are considered with a lower priority.

The main drawback of these methods is that they can not dynamically adapt to different network loads.

3. Proposal of an Extended Priority Assignation Scheme (ExPAS)

All the cited techniques use the cell loss priority (CLP) scheme to discriminate between video data. Using this two-state priority mechanism, they are not able to efficiently capture MPEG data structure complexity.

Therefore, a new scheme is proposed to extend ATM prioritization capability. The scheme is called Extended Priority Assignation Scheme (ExPAS) and offers three cell priority levels. To achieve this performance, a new field is defined referenced as Extended CLP (ExCLP). This field is located in the cell header and comprises the classical CLP bit and the adjacent PTI ATM-user-to-ATM-user bit (AUU) [5]. Table 1 presents the mapping of MPEG data frames into the ExCLP field. The PTI AUU bit is currently used by cell group discard schemes, such as Tail Drop and Early Packet Discard [6], to determine ends of upper layer packets. To allow these mechanisms to properly work with ExPAS, a similar AUU flag is reserved in the ExCLP field. This new definition of the CLP and the PTI-AUU bits permits a better utilization of the cell header. Using this new ExCLP field, three priority services are now available in a single channel. Whereas the traditional approach restricts the number of priority to two and under utilize the cell header capabilities.

Cell Type	CLP	PTI-AUU	Priority Level
Intra-frame	0	0	High
Predictive	0	1	Medium
Bi-directional	1	0	Low
End of Message	1	1	n/a

Table 1 - ExCLP field Mapping

4. Proposal of an Adaptive Selective Cell Discard Scheme (A-SCD)

One of the simplest switch buffer scheduling algorithm is to serve cells in first-in first-out (FIFO) order. If buffer congestion occurs, the incoming cells are dropped regardless to their importance. This random discard (RD) strategy is not suitable for video transmission.

In this section we propose a variant of the selective cell discard (SDC) scheme [7], which provides better performance for carrying video streams over lossy environment. The proposed adaptive-SDC scheme is associated with ExPAS to form a quality of picture (QoP) control framework. The aim of this framework is to ensure graceful picture degradation during overload periods. It allows accurate video cell discrimination and progressive drop by adjusting dynamically A-SCD mode in respect with cell payload types and switch buffer occupancy.

During light congestion, we propose to drop a lower priority cell first rather than delayed it and give its buffer space to a higher priority cell. This approach ovoid congestion worsening and remain the mean cell transfer delay in respectable level. This proactive strategy is performed gradually by including medium and high priority cells .

As illustrated in Figure 1, four buffer thresholds are used : Stop_Threshold (ST), Low_Threshold (LT), Medium_Threshold (MT) and High_Threshold (HT). These thresholds define up to four discarding modes depending on the used priority assignation scheme (e.g. CLP-based or ExCLP). The utilization of four thresholds instead of two reduces the speed of oscillation and have shown better performance.

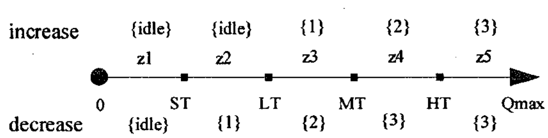

Figure 1 - Buffer thresholds

- Mode {idle) : If the buffer queue length (Ql) is lower than Low_Threshold (e.g. in regions z1 and z2), no cells are discarded.
- Mode {1} : When the queue length exceeds Low_Threshold but is still below Medium_Threshold, a light congestion is detected and only the low priority cells are eligible for elimination. This mode stops when Ql falls down to Stop_Threshold.
- Mode {2} : When the queue length exceeds Medium_Threshold but is still below High_Threshold, a medium congestion is detected. With CLP-based priority techniques, the scheme behaves like in mode {1}, while with ExPAS the low and medium priority cells are both eligible for elimination. This mode stops when Ql falls down to Low_Threshold.
- Mode {3} : When Ql exceeds High_Threshold, all the incoming cells are eliminated until the queue length drops below Medium_Threshold.

Using this adaptive discarding technique low-priority cells are firstly dropped to quickly reduce buffer occupancy during light congestion, while higher priority cells are preserved from elimination. If the congestion worsens, all the cells are progressively candidate to discarding.

This approach allows a graceful and controllable picture degradation with low operation complexity.

5. Simulation Model

In this paper, we consider a simulation model composed with an ATM switch (*SW*), an OC-1 bottleneck link (*L*) and three VBR MPEG connections. As depicted in Figure 2, the distances between the sources and the switch are equal and set to 0.125 miles (e.g. 0.2 km). 'L' is initialized to 310 miles (e.g. 500 km).

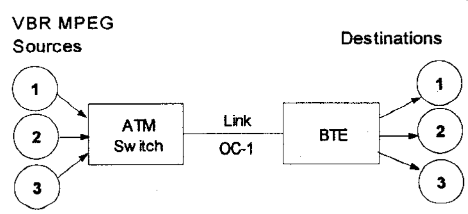

Figure 2 - Network Model

The VBR video connections are generated using three MPEG-1 frame traces : 'Star-Wars', 'Tennis' and 'Soccer' [8]. The main statistics of the video sequences are summarized in Table 2.

	Star Wars	Tennis	Soccer
MCR (Mbps)	0.36	0.55	0.63
PCR (Mbps)	4.24	1.58	2.29
Peak/Mean ratio	11.7	2.87	3.63

Table 2 - Statistics of the video sources

Every connection starts transmitting at different times in the range [0, 83.3 ms.]. Indeed, we have noticed that the deterministic flow properties of the compressed video have an important impact on switch buffer occupancy, which confirms the conclusions in [9]. Therefore at simulation start, the first frames (e.g. I-frames) of the three connections are desynchronized. We also assume shared output FIFO buffers and constant decoder output rates during transmission of a frame. The level of congestion is measured by means of thresholds.

Let us define the aggregate cell loss ratio, as the number of lost cells of the three video connections vs. the total transmitted cells.

Let us define the I-frame cell loss ratio, as the number of lost cells belonging to I-frame from the three connections vs. the total number of transmitted I-cells. The same metric is applied on P- and B-frames.

Let us also define the end-to-end cell transfer delay (*CTD*) as the time between the departure of cell K from the source node (t_{iK}) and its arrival at the destination node (t_{0K}) : $D_K = t_{0K} - t_{iK}$

For all the proposed scenarios, the processing delay at the ATM layer is not explicitly modeled. We assume that its contribution to the end-to-end delay experienced by each cell is relatively constant, and thus it can be omitted.

Emphasis is on the variation of the mean-CTD for the aggregate stream and for the video cell sub-flows associated with I, P, and B frames.

6. Results Analysis

For this study, we compare the performance of Extended PAS with the three following techniques.

- No priority cell assignation
- Static I/PB frame partition using CLP bit.
- Static IP/B frame partition using CLP bit.

Adaptive SCD is used with all these techniques. Exception is on 'No priority cell assignation' technique where random discard is performed when HT is crossed.

Seven switch buffer configurations are investigated. For each of them, the same method is applied to determine the values of the four thresholds. ST, LT, MT and HT are respectively set to 1.0, 0.9, 0.8 and 0.7 fraction of the maximum buffer size (Qmax). Qmax is varying in the range of 40 to 165 Kbits.

Figures 3, 4, 5, and 6 show the cell loss ratio for respectively the aggregate, I-, P- and B- cell flows. From Figure 3, no significant difference appears regarding to the cell loss probability. Since the same discard mechanism (e.g. A-SCD) is used with every priority assignation technique, we get approximately the same loss ratios regardless to the cell types.

As illustrated in Figure 4 and 6, ExPAS concentrates the loss within the B-frames and protects efficiently the referenced I-frames. With 'No priority' and 'Static CLP I/PB' approaches, the I-cell and P-cell loss are much higher than the two other priority partition techniques. This is due to the GOP structure of the sequence (N=12 and M=2) and the sizes of the Intra-frames. Indeed, when an I-frame is transmitted, the buffer queue length rapidly increases to accommodate the burst. The two schemes start to drop the I-cells immediately and stop when the queue length decreases below ST. The same phenomena is repeated with the following P-frame. Indeed, 'Soccer' sequence has several scene changes during the simulation period (e.g. 1.43 min.). This yields to code some of the frames as intra-coded picture rather than predictive coded. To illustrate this characteristic, the mean size of the I and P-frames are respectively 65.9 Kbits and 38.2 Kbits for 'Soccer', and 57.4 Kbits and only 16.3 Kbits for 'Star-wars'.

When the buffer size is small (e.g. Qmax lower than 120 Kbits), ExPAS performs like the other schemes. This can be explained by the fact that the buffer can not simultaneously accommodate cells from different connections. Thus when the lower and medium thresholds are crossed due to the arrivals of an I-frame burst, few B- or P-cells are available for discard. The queue keeps on rising and the A-SCD scheme switches to the mode {3}.

B-frames are the most concern by loss and contribute largely to the overall loss ratio. Afterwards, I- and P-cells are, in this order, the most subject of drop. This is explained by the frequency of B-frames in video sequences, as well as the drop policy of A-SCD. Indeed, in our MPEG video samples, the proportion of I, P and B data are respectively to 53, 24 and 23 % of the aggregate stream. Due to the GOP pattern and the multiplexing process B-frames occurs more often and are more likely to be discarded, e.g. 48 B-frame occurrences per second, for only 2 and 6 for I and P-frames.

As depicted in Figure 7, the mean-cell transfer delay increase in order of magnitude of the buffer size. We may notice that ExPAS has not much effect on the variation of the mean-CTD. With limited buffer size it performs similarly to 'No priority' and CLP-based schemes. When 'Qmax' rises it even shows better results than 'static CLP IP/B' partition scheme. This is because 'ExPAS' starts to drop P-cells earlier than 'Static CLP IP/B' and thus reduce the buffer occupancy much faster. When Queue length exceeds the medium threshold (MT), both B- and P-cells are eligible for elimination with ExPAS, which represent about fifty percent of the whole stream. With 'Static CLP IP/B' scheme, the buffer must be filled until High threshold to start P-cells elimination. The impact of this partition approach is a higher mean buffer occupancy and thus a greater mean CTD.

7. Conclusion

Since loss is a major concern for compressed video data, we have proposed in this paper a new priority data partition technique which extends ATM prioritization capability. With a better use of the ATM cell header, the Extended Priority Assignation scheme (ExPAS) provides three different service classes per connection. In association with an intelligent video-oriented cell discard scheme, it overcomes the problem of picture quality degradation caused by random drop.

We have compared ExPAS with former approaches based on the two-state CLP mechanism. The results have shown a better protection of Intra-coded frames with moderate increases of the mean cell transfer delay. By applying a more accurate cell discrimination strategy and progressive cell dropping the proposed quality of picture framework is found to be suitable for the transmission of hierarchical MPEG video over ATM networks.

8. References

[1] ISO/IEC International standard 13818-2, "Generic Coding of Moving Pictures and Associated audio information : Video", 1995.

[2] P. Pancha and M. El Zarki, 'MPEG coding for variable bit rate video transmission', IEEE Communication magazine, May 1994, pp.54-66.

[3] A. Mehaoua and R. Boutaba, 'Layered transmission of MPEG over VBR and ABR connections', Fifth International conference on Telecommunication systems, Modeling and Analysis, Nashville, TN, March 1997.

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

[5] A. Romanov and S. Floyd, 'Dynamics of TCP Traffic over ATM networks', ACM SIGCOMM'94, Sept. 1994, pp. 79-88.

[6] ITU-T I.361 specification of ATM layer, March 1993.

[7] ITU-TI.371, 'Traffic Management and Congestion Control in B-ISDN', Perth, November 1995.

[8] O. Rose, 'Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems', TR-101, University of Würzburg, Feb. 1995.

[9] O. Rose and M. R. Frater, 'Impact of MPEG Video Traffic on an ATM multiplexer', IFIP High Performance Networking '95, Palma, Sept. 1995, pp. 157-168.

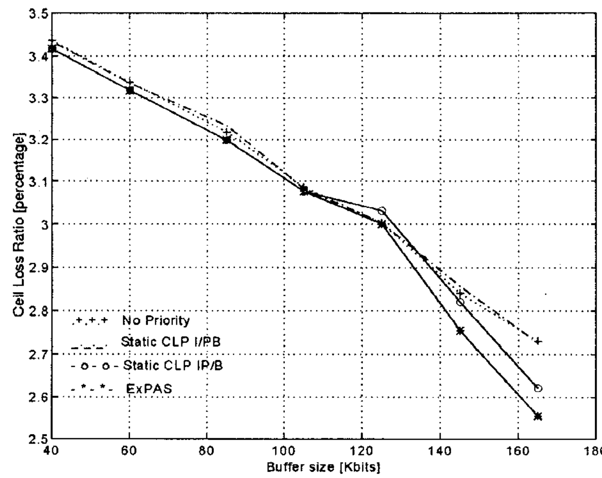

Figure 3 - Aggregate Cell Loss Ratio

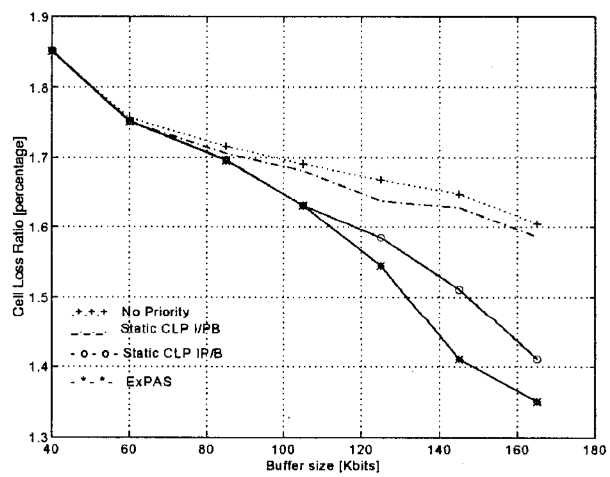

Figure 4 - (I)ntra-frame Cell Loss Ratio

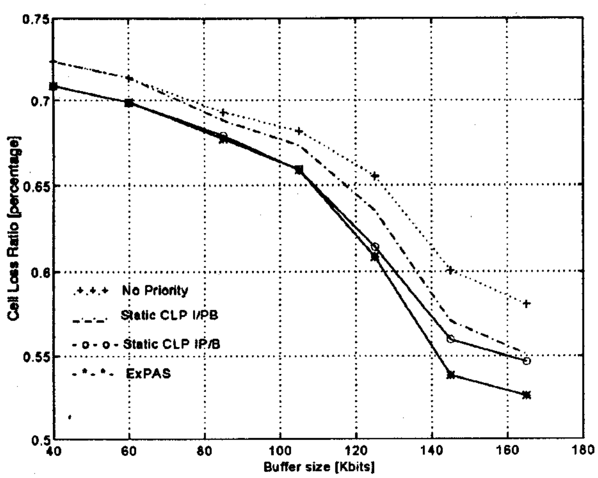

Figure 5 - (P)redictive-frame Cell Loss Ratio

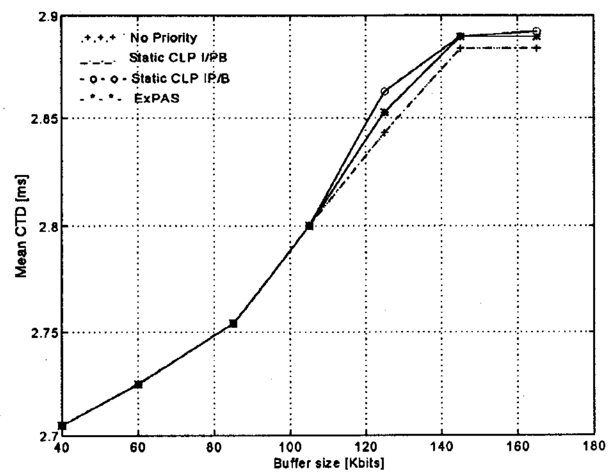

Figure 7 - Mean-CTD for the aggregate stream

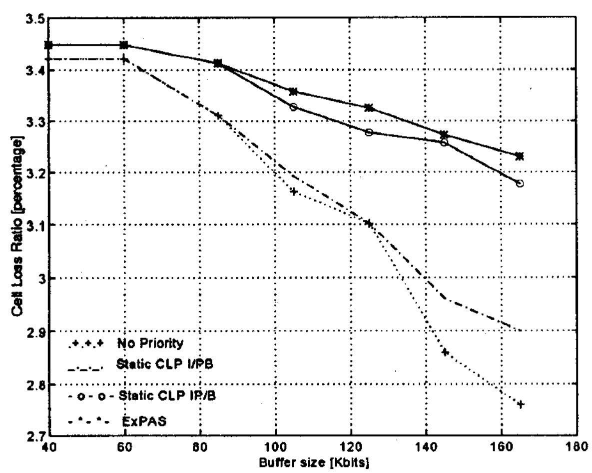

Figure 6 - (B)idirectional Predictive-frame Cell Loss Ratio