ELSEVIER

Computer Networks and ISDN Systems 29 (1998) 2021–2037

COMPUTER NETWORKS and ISDN SYSTEMS

Performance analysis of cell discarding techniques for best effort video communications over ATM networks

Ahmed Mehaoua [a,b,*], Raouf Boutaba [a]

[a] *Computer Science Research Institute of Montreal, 1801 McGill College, Montreal, Que., Canada H3H 2P2*
[b] *University of Versailles, 45 Av. des Etats-Unis, 78035 Versailles, France*

Abstract

With increasing interest in the transmission of audio–visual applications over ATM best effort services, efficient video-oriented control mechanisms for improving the video quality in the presence of loss have to be designed. In this paper, we propose and evaluate two new slice-based discard schemes for use with available bit rate and guaranteed frame rate services (e.g. formerly UBR +). The schemes adaptively and selectively adjust the discard level to switch buffer occupancy and video cell payload types. To improve their performance, we also introduce a dynamic frame-level priority data partition technique based on MPEG data structure and feedback from the network. To support these mechanisms, enhancements to the ATM adaptation layer 5 and a new MPEG-2 encapsulation strategy are also proposed. The presented quality of picture (QoP) control framework is evaluated using simulation and actual MPEG video data. The overall aim of the framework is double. First, ensuring a graceful picture quality degradation by minimizing cell loss probability for critical video data, and second optimizing the network effective throughput by reducing transmission of non useful data. In comparison to previous approaches, the performance evaluation have shown a significant reduction of the bad throughput and minimization of losses of intra- and predictive-coded frames at both cell and slice layers. © 1998 Elsevier Science B.V.

Keywords: ATM; AAL5; UBR + ; VBR MPEG-2; QoS; Loss; Delay; Priority

1. Introduction

The transport of VBR MPEG-2 applications over ATM introduces several issues that must be addressed in order to ensure a high end-to-end quality of service. These include the choice of the adaptation layer, method of encapsulation of MPEG-2 packets in ATM adaptation layer (AAL) packets, choice of scheduling algorithms in the network for control of delay and loss and the choice of the class of service with associated traffic management policies for congestion avoidance.

Best effort services, such as available bit rate (ABR) and unspecified bit rate (UBR) [1], are expected to be the primary services for carrying traditional data over ATM networks. Since they will be widely available in the future and are based on the excess bandwidth with a lower usage cost, it is not anticipated that they will also support a non-negligible part of the multimedia traffic. Even though they are targeted for the delivery of non-real time applications, the inclusion of a minimum cell rate (MCR) and intelligent video-oriented congestion control

* Corresponding author. Tel.: + 1-33-39254091; fax: + 1-33-39254057; e-mail: Ahmed.Mehaoua@prism.uvsq.fr.

0169-7552/98/$19.00 © 1998 Elsevier Science B.V. All rights reserved.
PII S0169-7552(97)00121-9

schemes makes them attractive and economical candidates for the transport of specific video applications (i.e. video playback and live video recording). ATM Forum is currently addressing this topic by the development of a new AAL5 audio visual service specific convergence sublayer (AV-SSCS) and by the introduction of real-time extensions to the ABR service [2].

These video applications will widely use motion picture expert group (MPEG) compression standards to save network resources. MPEG coding algorithm structures video data in a hierarchical format ordered by increasing spatial size: 8×8 pixels block, 16×16 pixels macroblock, slice, frame, group of pictures and sequence [3,4]. Two of them, slice and frame, have significant impact on decompression and displaying process.

Slice is the main coding processing unit in MPEG. Coding and decoding of blocks and macroblocks are feasible only when all the pixels of a slice are available. Besides, coding of a slice is done independently from its adjacent slices, making it the smallest autonomous unit. Consequently, in case of long error bursts due to transmission problems slices serve as resynchronization decoding units.

Frame or picture is the basic unit of display. Three picture types may be present in MPEG streams. They differ by the coding method used: intra-coded (I) picture, predictive-coded (P) picture and bidirectionally predictive-coded (B) picture. Intra- and predictive-coded pictures are essential and have to be preserved from corruption during transmission. Indeed, due to error propagation a corrupted or non-available reference picture (e.g. I- or P-frame) leads on perceptible picture degradation. 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.

From the presented video properties, three obvious remarks stand out. First, the smallest transmission data unit is rather a slice than cell, AAL or system packet (e.g. transport stream or program stream). Secondly, in a situation of congestion, dropping video cells indiscriminately can cause serious degradation in the picture quality. I- and P frames have to be better protected during transmission. Finally, unlikely to delay insensitive applications, error recovery techniques based on retransmission are useless to recover corrupted data. Therefore, best effort video delivery services should be evaluated in their ability to provide an efficient message-based service. In this paper, the term message is interpreted specifically in reference to MPEG video. A message may be an encoded block, macro-block, slice or an entire picture.

To address these requirements, we propose and evaluate new video-oriented discarding techniques to be used with ATM best effort services; namely, adaptive and selective cell discard (adaptive-SCD) and adaptive and partial slice discard (adaptive-PSD). To improve their performance, we have designed a dynamic data partition and cell priority assignation mechanism referenced as dynamic priority assignation scheme (dynamic-PAS). Using feedback information from the network, it dynamically adjusts cell priority in respect to the MPEG data structure and network load. To support this QoS control framework we also introduce an MPEG-2 video stream encapsulation strategy into a new AAL5 service specific convergence sublayer (SSCS).

The balance of this paper is organized as follows. In Section 2, we briefly review principal traffic and QoS control approaches for video networking and for traditional packet-oriented applications over ATM best effort services. Section 3 is devoted to the description of the different components of the proposed best effort video delivery framework. In Section 4, we introduce the network model, the investigated performance parameters and we discuss the results. Finally, we, conclude and present directions for future works in Section 5.

2. Best effort video networking

To address the problem of transmission of compressed VBR video over lossy networks, several protection and recovery techniques have been proposed to minimize video quality degradation due to cell loss. Layered coding with prioritization is one of the most popular approach [5–7]. Forward error recovery schemes with destination concealment have also been designed to cope with the problem [8–11]. Another approach consists of reducing the burstiness and the peak bandwidth requirements of video by applying complex smoothing and buffering techniques at the source [12,13].

With classical packet-oriented services (TCP, LANs) over ATM best effort services, various techniques have also been developed to preserve packet integrity and achieve a high effective throughput during overloads. These mechanisms may be classified according to three ATM services, ABT, ABR and UBR.

ATM block transfer (ABT) with fast reservation protocols (FRP) is shown to maintain high link utilization during network overloads with small buffer requirements [14,15].

In 1995, the ATM Forum standardized the ABR rate-based flow control [1] to manage congestion of unpredictable bursty traffic. However, like adaptive windowing mechanisms, the rate-based mechanism requires one or more network round-trip times before it effectively reacts to congestion [16].

Nevertheless, unspecified bit rate (UBR) is the true 'best effort' service [1]. With no flow control and no loss guarantees, it provides the least expensive service for the transport of packet-based applications. However, because of its simplicity, plain UBR with inadequate buffer sizes performs poorly in a congested network. Packet tail discard or partial packet discard (PPD) has been proposed to address this problem [17]. If a cell is dropped from a switch, the subsequent cells of the higher layer protocol data unit are discarded. Romanov and al. [18] have shown that PPD improves network performance to a certain degree, but it is still not optimal. Therefore, they proposed a new mechanism called early packet discard [18] that achieves a better throughput performance but does not guarantee fairness among the connections [19]. When the switch buffer queues reach a threshold level, the entire higher level data units (e.g. AAL5 PDU) are preventively dropped. To improve its fairness, selective packet drops based on per-VC accounting have been introduced by Heinanen and Kilkki and referenced as fair buffer allocation (FBA) [20].

However, none of the mentioned congestion control and QoS management schemes are focusing on the transmission of specific MPEG-encoded video streams over ATM best-effort services. Thus, we propose in Section 3, a new efficient quality of service control framework which ensures graceful picture quality degradation during the congestion period.

3. A QoS control framework for best effort video networking

3.1. Proposal of a slice-based dynamic priority assignation scheme (dynamic-PAS)

Since human perception is less sensitive to low frequency components of a video signal, subsequent blocks are transformed into the frequency domain using the discrete cosine transform (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) mechanism is usually used to provide a two-level priority service within a single ATM channel.

In [5,21], this previous technique is applied at the macroblock layer. Similarly to the block-based approach, the dc value for each of the 8×8 blocks are assigned to the high priority (HP) stream. The macroblock header and the motion vector in the case of P-frames are also included in the HP stream. For the remaining 63 DCT coefficients 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 allow the regeneration of the original bit stream by the destination, the macroblock address is joined to the LP information.

To overcome the limited priority capability of the CLP mechanism, a connection-level prioritization approach is evaluated in [22] to transport a layered MPEG-2 video application over different ATM service classes. The scheme uses static data partition between two connections by means of a load balancing factor (LBF). The virtual connections are associated with a guaranteed service class (i.e VBR-rt) and a best effort service class (e.g. ABR) to, respectively, carry the base layer and the enhancement layer.

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

Consequently, 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 [23], two static priority partition strategies are proposed which also use the CLP-bit:

- Static I/PB priority partition: In this method, I-frame cells are considered with a high priority and have their CLP bit set to '0', while P- and B-frame cells are assigned a lower priority with CLP flag set to '0'. 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 both considered with high priority and are better preserved from elimination. Only B-frames cells are assigned a low priority.

The main drawback of these methods is that they cannot dynamically adapt to network load changes.

Therefore in this section, we propose a new video partition and prioritization scheme, named dynamic priority assignation scheme (dynamic-PAS). The scheme is simple to implement and sufficiently generic to be performed at any MPEG data hierarchy level. In this paper, the emphasis is on the slice layer. The scheme uses the classical CLP bit and dynamically assigns cell priorities according to the current MPEG slice type (I, P or B) and the reception of backward congestion signals from the network.

Cells belonging to (I)ntra-coded slices have the highest priority and their CLP-bit is set to '0'. (B)idirectional-encoded slices have the lowest priority and the associated cells have a CLP value of '1'. Regarding (P)redictive encoded slices, they are alternatively assigned high and low priority depending on the network load. At the beginning of the transmission, P-cells are initialized with a high priority (e.g. IP/B partitioning mode).

As illustrated in Fig. 1, when the buffer queue length (QL) exceeds an upper threshold, an early

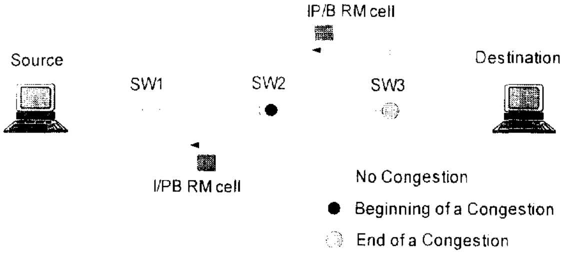

Fig. 1. Transmission of I/PB RM cells.

congestion is detected and the ATM switch emits a signal to the source which in turn, adjusts P-cells priority level to low (i.e. I/PB partitioning mode). In our implementation we use forward resource management (RM) cells, with congestion indication flag (CI) tagged, to notify the destination and afterward the source. These cells are called I/PB-RM cells. We choose RM cells rather than the EFCI mechanism to notify the sources for reliability consideration. Indeed, if the switch is highly congested, all cells are dropped and then no one is available for carrying the indication signal to the source.

When the queue length decreases below a lower threshold, a new signal (i.e. using IP/B RM cells) is transmitted to the source to modify its data partitioning mode. P-cells priority are switched back to a high priority (e.g. IP/B mode).

The PTI ATM-user-to-ATM-user bit (AUU flag) is employed to indicate whether it is the last cell of an upper message. We propose to use this flag to distinguish between successive slices. In this paper, the cell having its PTI-AUU flag set to '1' is termed the end of slice (EOS) cell. Using dynamic-PAS, each CLP and PTI bit singly assumes its original meaning according to the standard [24]. This dynamic priority assignation mechanism takes the advantages of both static I/PB and static IP/B priority partition. The main drawback of the scheme is that it is stringently dependent on the round trip time delay and thus on the network topology and link distances.

However, to support this slice-based mechanism, enhancements to AAL5 and a new MPEG-2 video stream encapsulation strategy are required. Moreover, to allow the proposed cell discarding schemes to properly work with the highest efficiency, video slice boundaries have to be correctly managed. Three

Table 1
Video data encapsulation requirements

Protocol layer	Requirements
Transport layer	every MPEG-2 transport stream packet is composed with data from a single video slice
ATM adaptation layer	an ATM adaptation layer packet is constituted with data from only one video slice
ATM layer	an ATM cell embeds data from only one video slice

data structuring conditions have to be respected by the source before transmission. They are summarized in Table 1.

3.2. *Proposal of a slice-based MPEG2 video stream encapsulation strategy*

The key factor that controls the end-to-end performance is the ATM adaptation layer. AAL 5 is currently the most commonly used. It is used to encapsulate UNI 4.0 signaling messages [25] and, in most cases for transporting non-real time data traffic [26]. In 1995, the ATM Forum has recommended the carrying of constant bit rate MPEG-2 streams over AAL5 with a null convergence sub-layer [27]. However, AAL5 is inadequate for the transmission of variable bit rate video and required extended features. Forward error correction, jitter removal, multiplexing and multi-priority support are some of them. AAL2 is supposed to address these requirements but is not yet standardized. Therefore, in [28] we propose a new AAL5 service specific convergence sublayer (SSCS) with an efficient encapsulation strategy to cope with MPEG-2 networking constraints.

As depicted in Fig. 2, uncompressed video frames, called presentation units, are individually encoded according to the MPEG standard and are referenced as access units. The stream produced by these access units is then named an elementary stream. The next step is its packetization. The resulting stream is called a packetized elementary stream (PES) [29].

There are no specific requirements for encapsulating encoded data in a packet elementary stream (PES). This means that an access unit may start at any point within a PES packet. In addition, more than one access unit may be present in one PES packet. Nevertheless, the way this packetization is done can significantly affect the performance of the decoding process and the quality of the service provided by the network. For instance, if each PES

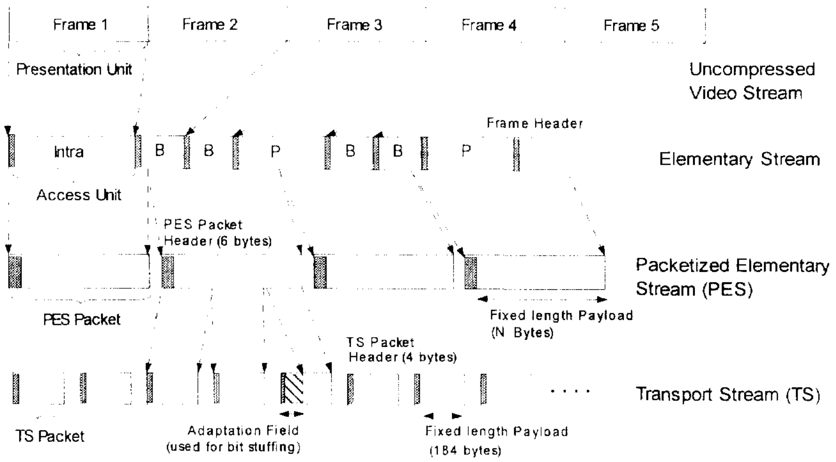

Fig. 2. PES encapsulation using fixed length packet.

packet contains exactly one MPEG data message (e.g. frame, slice, macroblock or bloc) the decoder can easily determine the start and end of the message. Similarly, network transport and control policies can take benefit of this structure to offer a guaranteed packet-oriented service. However, this approach requires the use of a variable size packet, which induces a slight added complexity for the encoding process.

Instead of encapsulating MPEG video data at the macroblock level, as in the method by Ghanbari and Hugues [30], we propose that each PES is built from a single encoded video slice (see Fig. 3). Indeed, slice is the main coding processing unit and the smallest autonomous unit. Coding and decoding of blocks and macroblocks are feasible only when all the data of a slice are available [3].

At the next step, we propose to segment the PES packet into a number of 188-byte fixed length transport stream (TS) packets. In respect to the standard [29], every TS packet embeds data from only one PES packet. In our case, the last transport packet may not be completely full since it is unlikely that a variable PES packet will fit exactly into an integer number of transport packets. Thus, stuffing bytes are placed in the adaptation field to complete the payload.

In the worst case, e.g. padding 183 bytes, we evaluate the average overhead to 0.2% with a single slice per frame, to 3.5% with 15 slices per frame. We assume a NTSC TV broadcast quality of 512 × 480 picture resolution [31]. The introduced overhead is highly dependent on the distribution and the number of slice per frame. However, it can be avoided using program stream (PS) packets.

3.3. Proposal of an AAL5 service specific convergence sublayer

At the AAL service access point (SAP), the transport layer passes the TS packets to the SSCS using message mode service with blocking/deblocking internal function [26]. The following AAL-UNIDATA-REQUEST (ID, M, SLP, CI) primitive is used. The 'interface data' parameter specifies the exchange TS packet. The 'more' parameter indicates if it is the last AAL SDU of the upper message (e.g. end of the slice). The 'submitted loss priority' parameter gives the priority level of the TS packet.

In this paper, we propose to extend SLP range from two to three possible values and to initialize it in respect to the 'picture_coding_type' field located in the MPEG frame header [3]. This field specifies for each frame, the coding mode used (e.g. intra, predictive or bi-directional predictive). Consequently, three different types of SSCS-PDUs are newly defined: high-priority (I-frames), medium-pri-

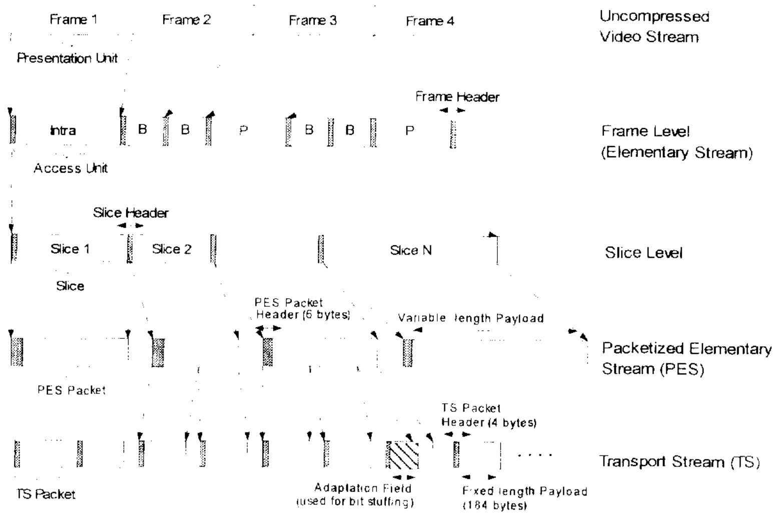

Fig. 3. Slice-based PES encapsulation using variable length packet.

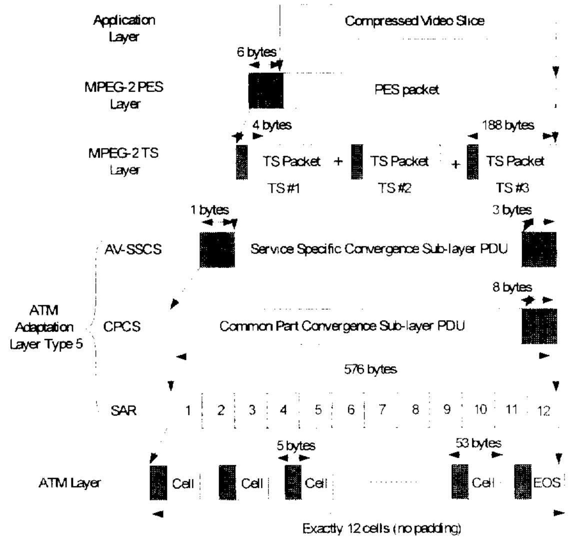

Fig. 4. New mapping of MPEG video slices on AAL5.

ority (P-frames) and low priority (B-frame). This parameter also indicates how the 'SLP' parameter of the CPCS_UNIDATA_invoke primitive shall be set. Finally, the 'congestion indication' indicates how the 'CI' parameter of the CPCS_UNIDATA_invoke primitive shall also be set.

As illustrated in Fig. 4, SSCS groups every three (3) TS packets and adds a header and a trailer information. The header is composed of a 4-bit sequence number (SN), and a 4-bit SN protection (SNP). The trailer consists of a 3-byte forward error correction (FEC).

The FEC scheme uses a Reed–Solomon (RS) code [32,33], which enables the correction of up to 2 loss bytes in each block of 564 bytes (e.g. 3×188). The addition of a sequence number modulo-16 of 4 bits enables the AAL5 receiver to detect and locate up to 15 consecutive SSCS PDU losses. When losses are detected, dummy bytes are inserted in order to preserve the bit count integrity at the receiver. The SNP contains a 3-bit CRC generated using the generator polynomial $g(X) = X^3 + X + 1$ and the result is protected by an even parity check bit. The SNP field is then capable of correcting single bit errors and detecting multiple bit errors.

The overhead introduced by the SSCS is 0.7% and the delay is three TS packets (about 12 cells) at the transmitter and the receiver.

The SSCS-PDU are then transmitted to the common part convergence sublayer (CPCS) using the CPCS_UNIDATA_Invoke primitive. The 8-byte CPCS trailer information is appended to the CPCS SDU and no byte padding is required. The resulting CS PDU is passed to the segmentation and reassembly (SAR) layer using the SAR_UNIDATA_Invoke primitive.

The under-laying SAR protocol will subsequently segment the CS-PDU into exactly twelve (12) 48-byte ATM SDUs. The ATM layer will then marked the CLP field of every cell using the 'AUU' and the 'SCLP' parameters of the AAL_DATA_Request [26].

3.4. Proposal of an adaptive and selective cell discard scheme (adaptive-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 [14], which provides better performance for carrying video streams over lossy environment. The proposed adaptive-SDC scheme [34] is associated with dynamic-PAS and the extended AAL5 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 adaptive-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

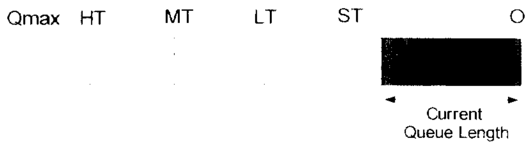

Fig. 5. Switch buffer thresholds.

Table 2
Adaptive-SCD operation modes

Queue Length	Increase direction	Decrease direction
QL ≤ ST	idle	idle
LT ≤ QL < ST	idle	{1}
MT ≤ QL < LT	{1}	{1}
HT ≤ QL < MT	{1}	{2}
QL ≥ HT	{2}	{2}

approach avoid congestion worsening and maintain the mean cell transfer delay in acceptable value [35]. This proactive strategy is performed gradually by including medium and high priority cells.

As depicted in Fig. 5, four buffer thresholds are used: Stop_Threshold (ST), Low_Threshold (LT), Medium_Threshold (MT) and High_Threshold (HT). These thresholds define three operation modes (see Table 2). The utilization of four thresholds instead of two reduces the speed of oscillation for the dynamic-PAS resource management mechanism and have shown better performance.

- *Mode {Idle}*: If the buffer queue length (QL) is lower than Low_Threshold, no cells are discarded.
- *Mode {1}*: If the total number of cells in the buffer exceeds Low_Threshold but is still below High_Threshold, a medium 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 QL exceeds High_Threshold, all the incoming cells are eliminated until the queue length drops below Medium_Threshold.

The I/PB RM cells are transmitted to all the video sources when MT is exceeded, while IP/B RM cells are send only when QL drops below LT. At the reception of feedback signals, the sources immediately change their operation mode. Consequently, a single P-frame may be transmitted with cells of different priority.

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.

3.5. Proposal of an adaptive and partial slice discard scheme (adaptive-PSD)

The drawback of adaptive-SDC is that useless cells (i.e. tail of corrupted slices) are transmitted to the destination and may congest upstream switches. An improvement is to drop all subsequent cells from a slice as soon as once cell has been dropped. This strategy is also referred as partial packet discard (PPD) or tail drop and is shown to significantly improve network performance with TCP connections.

In this section, we propose to enhance PPD to better suit MPEG video connections. The proposed adaptive and partial slice discard algorithm (adaptive-PSD) is shown to provide a better performance for carrying hierarchical video streams over ATM best effort services. The proposed scheme runs per virtual circuit and is highlighted in Fig. 6 and Section 3.5.1. It employs two state variables:

(1) S_DISCARDING indicates whether the buffer is currently discarding (S_discarding = 1) this message (e.g. slice) or not (S_discarding = 0).

(2) S_PRIORITY indicates the priority level of this slice. The indicator is modified at the reception of the first cell of the slice using its priority CLP flag. The switch is currently handling a high (S_priority = 0) or a low (S_priority = 1) priority slice.

Similarly to adaptive-SCD, adaptive-PSD uses three buffer thresholds, which also defines four operation modes.

- *Mode {Idle}*: If the current buffer queue length (QL) is lower than Low_Threshold, no cells are discarded.
- *Mode {1}*: If the total number of cells in the buffer exceeds Low_Threshold but is still below

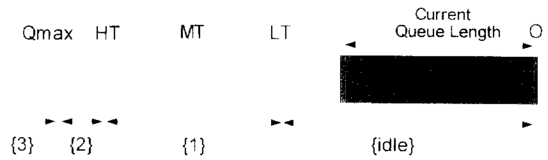

Fig. 6. Switch buffer thresholds and operation modes.

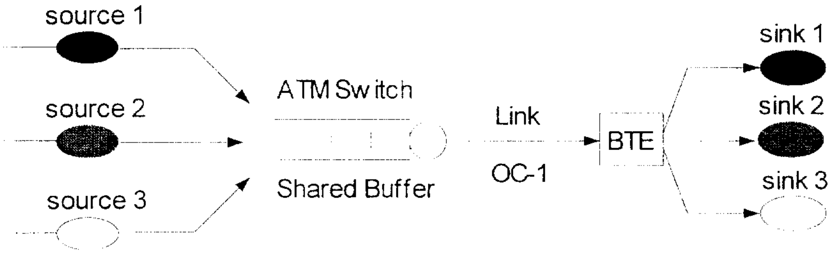

Fig. 7. Network model.

High_Threshold, for every video connection currently emitting a low priority slice, adaptive-PSD starts to discard their incoming cells until the reception of an EOS cell. The last cells are always preserved from elimination since they provides indication of the next slice. The cells with higher priority are accepted in the buffer. This mode stops when 'QL' falls down to Low_Threshold.

- *Mode {2}*: This mode is activated when QL exceeds High_Threshold, all the incoming slices are eliminated regardless to their priority level until the queue length drops below HT.
- *Mode {3}*: If the buffer is full, e.g. the maximum queue size (Qmax) is reached, all cells are dropped until the queue length falls down to High_Threshold.

Except in mode {3}, the end of slice (EOS) cells are always preserved from elimination.

3.5.1. Adaptive and partial slice discard scheme pseudo code

```
S_discarding = 0 /* Initialization */
IF (EOS) /* the last cell */
S_discarding = 0
IF QL < Qmax
accept the cell
ELSE
discard the cell
ELSE /* not the last cell */
IF S_discarding = 1
discard the cell
ELSE /* S_discard = 0 */
S_priority: = cell_priority
IF QL < Low_Threshold
accept the cell
ELSE
IF QL < High_Threshold and S_priority = high
accept the cell
ELSE
discard the cell
S_discard: = 1
```

4. Performance evaluation

4.1. Network model

We consider a network simulation model composed of an ATM switch with a finite capacity buffer, an 55.1 Mbps OC-1 bottleneck link (L) and three VBR MPEG connections (see Fig. 7). The distance between the sources and the switch is set to 0.2 km (e.g. 0.125 miles). 'L' is initialized to 2.5 km (e.g. 1.56 miles) to emulate a local area network (LAN). We assume a link propagation speed of 2.5×10^8 m/s and a propagation delay between the switch and the broadband terminal equipment (BTE) of 10×10^{-6} ms. The round trip time (RTT) between the sources and destinations is then 23.2×10^{-6} ms.

The VBR video connections are generated using three MPEG-1 frame traces: 'Star-Wars', 'Tennis' and 'Soccer' [36]. The main statistics of the sequences are summarized in Tables 3 and 4.

Table 3
Coding parameters

	Star wars	Tennis	Soccer
Compression rate (X:1)	130	121	106
Quantizer scale	I: 10	P: 14	B: 18
GOP pattern ($N = 12$, $M = 2$)	IBBPBBPBBPBB		

Table 4
Statistics of the video sources

	Star wars	Tennis	Soccer
Mean cell rate (Mbps)	0.36	0.55	0.63
Peak cell rate (Mbps)	4.24	1.58	2.29
Peak/mean ratio	11.7	2.87	3.63

Every connection starts transmitting at different times in the range (0, 41.6 and 83.3 ms). Therefore for the three connections, the I-frame beginning the first GOP are desynchronized. Indeed, we have noticed that the deterministic GOP pattern of the video sequences have an important impact on switch buffer occupancy and confirms the conclusions in [37]. Thus, the I-frame starting the first GOP of the three video connections are desynchronized.

In this paper, we assume 15 uniformly distributed slices per video frame and a simulation duration of 1.43 min. We also assume a shared output FIFO buffer and constant decoder output rates during the transmission of a frame. The level of congestion is measured by means of thresholds.

Seven switch buffer configurations are investigated. 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 1.0, 0.9 and 0.8 fraction of the maximum buffer size (Qmax). Qmax varies in the range 40 to 165 Kbits.

4.2. Performance parameters

Let us define in Tables 5 and 6, respectively, the possible states of a cell crossing the network and the investigated performance parameters.

Video slice loss ratio (SLR) is measured at the application layer and take into account decoding (e.g. lost and dead cells) and propagation (e.g. late cells) constraints.

During simulations, the processing delays at the ATM and AAL layer are not explicitly modeled. We assume that their contribution to the end-to-end delay experienced by the cell is relatively constant, and thus it can be omitted. Emphasis is on the variation of the mean-CTD for the aggregate stream composed of the three video connections.

4.3. Simulation results

Two performance evaluations are investigated. First, we compare the performance of dynamic-PAS with the three following static CLP-based partition techniques:

- Static IPB/-frame partition
- Static IP/B frame partition
- Static I/PB frame partition

For all these prioritization schemes, the switch runs the adaptive-selective cell discard (adaptive-SCD) scheme.

A second set of simulations focus on the comparison of the proposed quality of picture framework (e.g. adaptive-PSD plus dynamic-PAS) with different priority partition and discarding configurations.

All the previous static data partition schemes are associated with adaptive-SCD. An exception is on 'static IPB/-' where random discard (RD) is performed when high threshold is exceeded. This hypothesis allows us to simulate a non-selective discard mechanism.

Table 5
Data unit definitions

Data units	Definition
Lost cell	a cell dropped by adaptive-ESD scheme
Dead cell	a cell received at the destination but belonging to a partially discarded slice
Late cell	a cell arriving at destination after an ended time-out. This time-out is triggered at the reception of every first cell of a picture. Its value is set to $1/N$ s, where 'N' is the frame rate of the video sequence. In this paper, 'N' is equal to 24
Correct cell	neither a lost, dead or late cell
Correct slice	a slice received with only correct cells

Table 6
Performance parameters definitions

Performance parameters	Definition
Agregate cell loss ratio (CLR_agg)	the number of lost cells from the three video connections vs. the total transmitted cells
I-frame cell loss ratio (CLR_i), P-frame cell loss ratio (CLR_p), B-frame Cell loss ratio (CLR_b),	the number of lost and late cells belonging to I-frames from the three connections versus the total number of transmitted I-cells. The same metric is applied for P- and B-frames
Cell bad throughput (CB)	the number of dead cells versus the total transmitted cells. It is a performance parameter evaluated at the ATM layer
Video slice loss ratio (SLR)	the number of corrupted slices versus number of transmitted slices
Mean cell transfer delay (mean CTD)	time between the departure of cell K from the source node (t_{tK}) and its arrival at the destination node (t_{0K}): $D_K = t_{0K} - t_{tK}$ [38]

4.3.1. Priority data assignation strategies evaluation

Figs. 8–11 show the cell loss ratio for, respectively, the aggregate, I-, P- and B- cell flows. From Fig. 5, we can notice no significant difference between the aggregate cell loss curves for the different priority assignation techniques. Since the same dropping mechanism is applied, e.g. adaptive-SCD, we obtain approximately the same loss ratio. This ratio decreases from about 3.43 to 2.6% while the buffer size increases by a factor of four.

As illustrated in Figs. 9–11, dynamic-PAS concentrates the loss within the B-frames and protects efficiently the reference I- and P-frames. Indeed, using automatic switching between two operation modes, it takes advantages of both static partition methods. For the protection of I-frames, it performs as well as 'static CLP I/PB' and much better than the other two. For the protection of P-frames, 'static CLP IP/B' scheme gives the best results. This is due to the fact that it maintains the highest priority level to P-frames during the whole communication. However it and 'static CLP IPB/-' experience the worst I-cell loss ratio and they are not recommended to carry encoded video streams. This is due to the GOP structure of the sequence and the great sizes of the Intra-frames. Indeed, when an I-frame is trans-

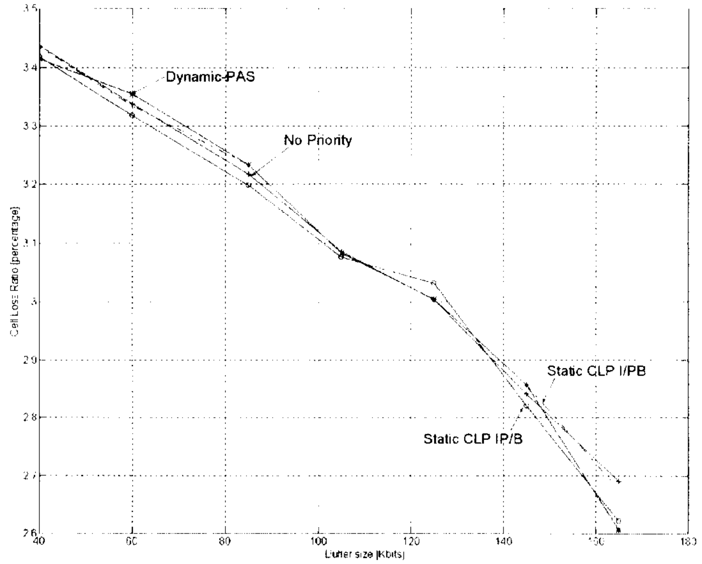

Fig. 8. Aggregate cell loss ratio.

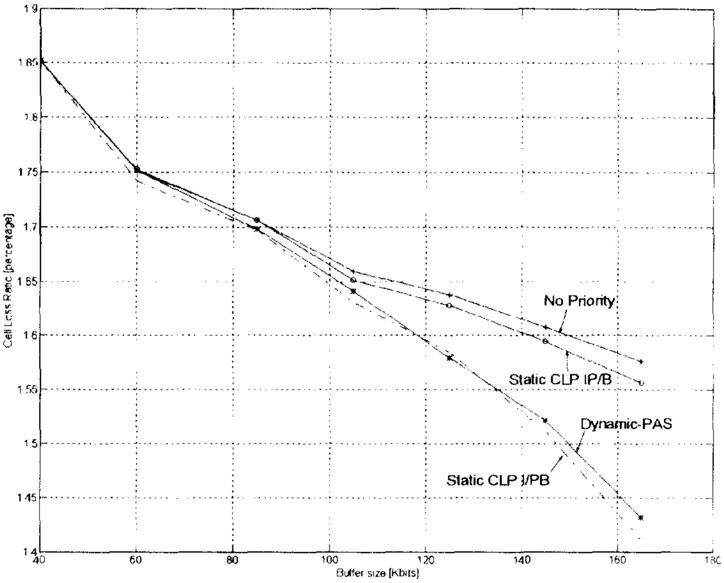

Fig. 9. (I)ntra-frame cell loss ratio.

mitted, 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. This yields to code some of the frames as intra-coded picture rather than

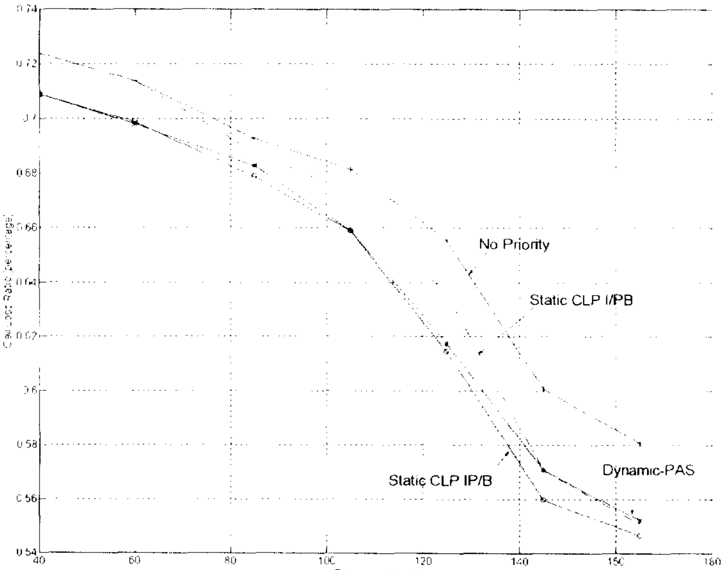

Fig. 10. (P)redictive-frame cell loss ratio.

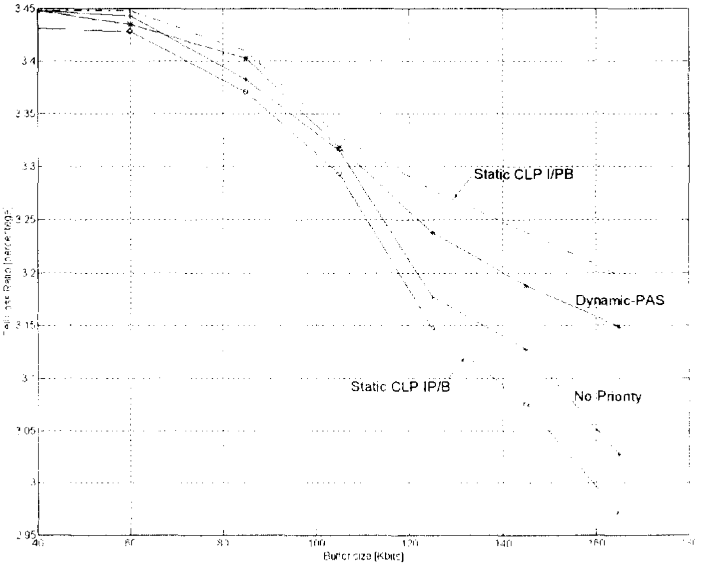

Fig. 11. (B)idirectional predictive-frame cell loss ratio.

predictive coded. To illustrate this characteristic, the mean size of the I and P-frames are, respectively, 65.9 and 38.2 Kbits for 'soccer' and 57.4 and only 16.3 Kbits for 'Star-wars'.

When the buffer size is small (e.g. Qmax lower than 120 Kbits), Dynamic-PAS performs like the other schemes. This can be explained by the fact that the buffer can not simultaneously accommodate cells

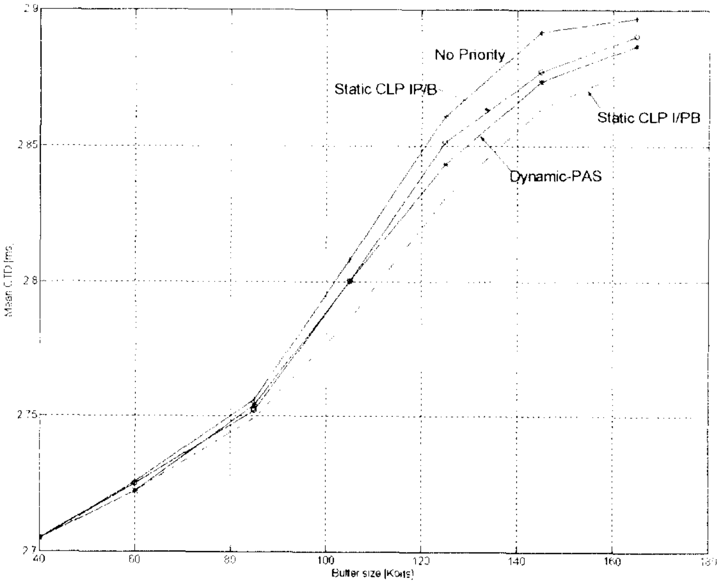

Fig. 12. Mean-CTD for the aggregate stream.

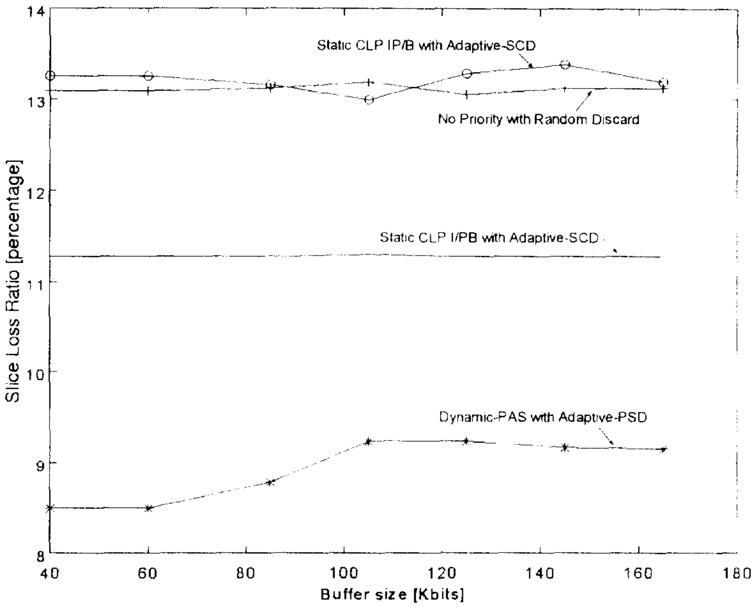

Fig. 13. Slice loss ratio for the aggregate stream.

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 adaptive-SCD switches to Mode {2}.

It is interesting to notice that, even though 'static CLP IPB/-' assigns a high priority to I-frame cells, it performs less efficiently than 'static CLP I/PB' for preserving I-cells from loss. This is due to the random discard effect when high threshold is reached.

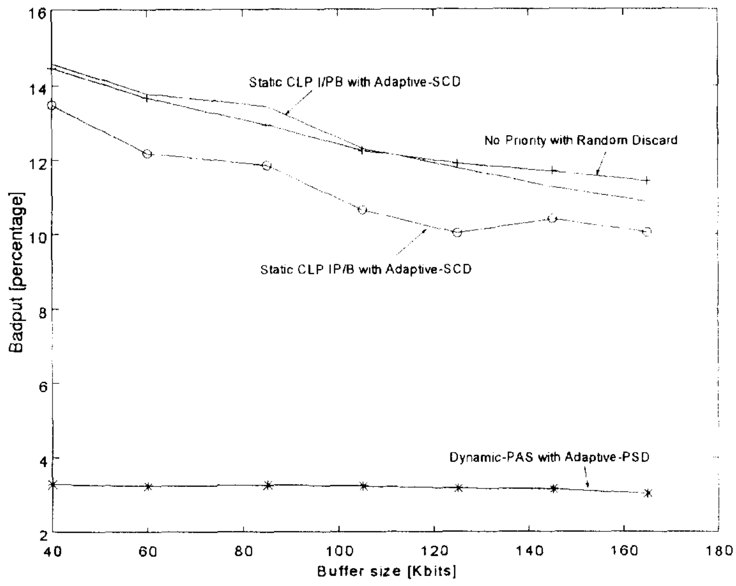

Fig. 14. Bad throughput.

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 adaptive-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 Fig. 12, the mean-cell transfer delay increase in order of magnitude of the buffer size. We may notice that dynamic-PAS has not much effect on the variation of the mean-CTD. With limited buffer size it performs similarly to the Static CLP-based schemes. A mean ratio varying from 2.7 and 2.9 ms. He even shows better results than 'static CLP IP/B' and 'static CLP IPB/-' partition schemes while 'Qmax' rises. This is because dynamic-PAS starts to drop P-cells earlier than these methods 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 dynamic-PAS, which represent approximately fifty percent of the whole stream. With 'static CLP IP/B' and 'static CLP IPB/-', the buffer must be filled until high threshold to start P-cells elimination. The impact of this strategy on buffer occupancy lead on a greater mean CTD.

4.3.2. Cell discarding policies evaluation

In this section we compare between the efficiency of the proposed quality of picture (QoP) control framework and that of other techniques at the upper application layer (i.e. MPEG slice level). As intuitively expected, the performance of adaptive-PSD is higher and it is well demonstrated by Figs. 13–15.

The slice-based delivery framework significantly improves the percentage of arrivals of non-corrupted slices at destination. Indeed, the slice loss ratio for the aggregate flow is reduced to achieve an upper bound of 9.1% of the total number of transmitted video slices. In comparison, static IPB/-, static IP/B and static I/PB reach, respectively, 11.4, 10.8 and 10.3%. Moreover, the bad throughput experienced by the network is dramatically minimized by a factor of four, e.g. from approximately 12 to 3% for the aggregate traffic.

Therefore, we observe that dynamic-PAS with adaptive-PSD outperforms the other approaches by better protecting both I- and P-frames. The differences among the static schemes are minimal with a

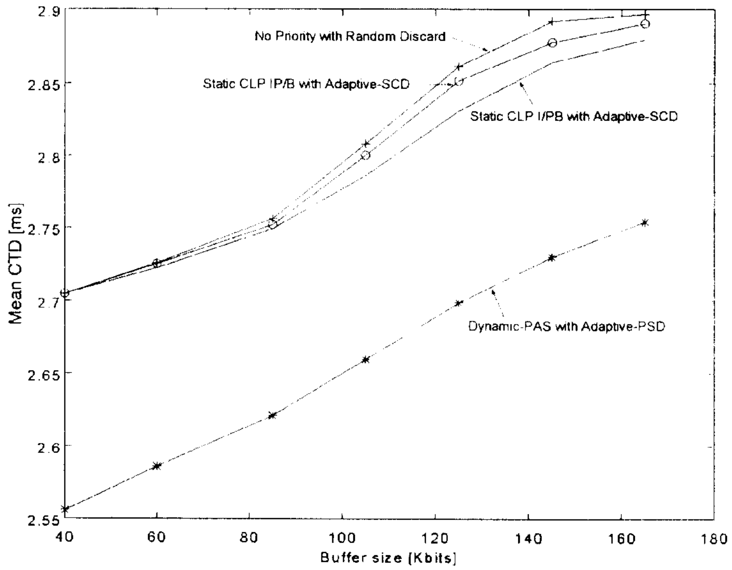

Fig. 15. Mean CTD for the aggregate stream.

slight advantage to 'static CLP IP/B'. This means that even though some prioritization mechanisms are efficient at the cell layer, they are not able to preserve their advantage at upper network layers without the addition of intelligent packet-oriented discard techniques.

Finally, Fig. 15 depicts the mean cell transfer delay gain obtained with dynamic-PAS + adaptive-PSD. The gap between the curves is constant and estimated to 0.15 ms in favor to dynamic-PAS associated with adaptive-PSD.

5. Conclusion

In this paper we have shown that a dynamic priority assignation strategy can significantly improve the quality of service provided to video encoded applications. By automatically adjusting its partition mode to the network load, the proposed dynamic priority assignation scheme (dynamic-PAS) efficiently prevent critical video data (e.g. I- and P-frames) from loss.

In association with an intelligent video-oriented discard scheme and an enhanced ATM adaptation layer 5, dynamic-PAS better exploits the structure of the MPEG video traffic and overcomes the difficulty imposed by random cell discarding. Together, they allow accurate cell discrimination and progressive cell group discard at the slice level, which lead to graceful picture quality degradation during overload periods.

The proposed slice-based discard scheme, named adaptive partial slice discard (adaptive-PSD), selects the slice to be dropped with respect to MPEG data hierarchy and the switch queue length. Simulation results show that in case of congestion the number of non-corrupted slice arriving at destination significantly increases, as well as the network effective throughput. Moreover, due to the batch elimination approach, a slight reduction of the mean cell transfer delay for the aggregate video stream is perceived.

Future work will involve studying performance evaluation of dynamic-PAS in wide area network configurations and the adaptation of early packet discard [18,38,39] mechanism to better suit with MPEG video transmission over ATM best effort services.

References

[1] ATM Forum, TM-SWG, Traffic Management Specification 4.0, at-tm-0056.000, April 1996.

[2] ATM Forum, TM-SWG, Real-time ABR: Proposal for a New Work Item, af-tm-96-1760, December 1996.

[3] ISO/IEC International Standard 13818-2, Generic Coding of Moving Pictures and Associated Audio Information: Video, 1995.

[4] D. LeGall, MPEG: A video compression standard for multimedia application, Commun. ACM 34 (1991) 47–58.

[5] P. Pancha, M. El Zarki, MPEG coding for variable bit rate video transmission, IEEE Commun. Mag. (May 1994) 54–66.

[6] B. DeCleen, P. Pancha, M. El Zarki, Comparison of priority partition methods for VBR MPEG, IEEE INFOCOM '94, pp. 689–696.

[7] A. Mehaoua, R. Boutaba, G. Pujolle, An extended priority data partition scheme for MPEG video connections over ATM, IEEE Symp. on Computers and Communications '97, Alexandria, Egypt, July 1997, pp. 62–67.

[8] W. Luo, M. El Zarki, Analysis of error concealment schemes for MPEG-2 video transmission over ATM based network, SPIE '95 2501 (1995) 1358–1368.

[9] G. Ramamurthy, D. Raychaudhuri, Performance of packet video with combined error recovery and concealment, IEEE INFOCOM '95, Boston, April 1995, pp. 753–761.

[10] S. Lee, Cell loss and error recovery in variable rate video, J. Visual Commun. Image Representation (March 1993) 39–45.

[11] A. Mehaoua, R. Boutaba, G. Pujolle, A picture quality control framework for MPEG video over ATM, IFIP/IEEE Workshop on Protocols for High Speed Networks '96, Sophia-Antipolis, France, October 1996.

[12] W. Feng, S. Sechrest, Smoothing and buffering for delivery of prerecorded compressed video, Comput. Commun. (October 1995) 709–717.

[13] J.D Salehi, Z.-L Zhang, J.F Kurose, D. Towsley, Supporting stored video: Reducing rate variability and end-to-end ressource requirements through optimal smoothing, ACM SIGMETRICS, May 1996, pp. 222–231.

[14] ITU-T, Recommendation I.371, Traffic Management and Congestion Control in B-ISDN, Perth, November 1995.

[15] J. Turner, Managing bandwidth in ATM networks with bursty traffic, IEEE Network 6 (5) (1992) 50–58.

[16] F. Bonomi, K.W Fendick, The rate-based flow control framework for the available bit rate ATM service, IEEE Network (March 1995) 25–39.

[17] G. Armitage, K. Adams, Packet reassembly during cell loss, IEEE Network Mag. 7 (5) (1993) 26–34.

[18] A. Romanov, S. Floyd, Dynamics of TCP traffic over ATM networks, ACM SIGCOMM'94, September 1994, pp. 79–88.

[19] R. Jain et al., Buffer requirements for TCP over UBR, ATM Forum 96-0518, April 1996.

[20] J. Heinanen, K. Kilkki, A fair buffer allocation scheme, unpublished manuscript.

[21] P. Pancha, M. El Zarki, Prioritized transmission of variable bit rate MPEG video, GLOBECOM '92, 1992, pp. 1135–1139.

[22] A. Mehaoua, R. Boutaba, Layered transmission of MPEG over VBR and ABR connections, 5th Int. Conf. on Telecommunication Systems, Modeling and Analysis, Nashville, TN, March 1997, pp. 425–430.
[23] T. Han, L. Orozco-Barbosa, Performance requirements for the transport of MPEG video streams over ATM networks, ICC '95, Washington, June 1995, pp. 221–225.
[24] ITU-T Recommendation I.361, Specification of ATM Layer, March 1993.
[25] ATM Forum, ATM User-Network Interface (UNI) Signaling Specification 4.0, af-sig-0061.000, July 1996, pp. 35–37.
[26] ITU-T I.363, B-ISDN Adaptation Layer (AAL), Perth, November 1995.
[27] ATM Forum SAA SWG, Audiovisual Multimedia Services: Video on Demand, Specification 1.0, ATM_FORUM/af-saa-0049.000, December 1995.
[28] A. Mehaoua, R. Boutaba, G. Pujolle, Proposal of an extended AAL5 for VBR MPEG-2 over ATM, IEEE Int. Conf. on Multimedia Computing and Systems '97, Ottawa, Canada, June 3–6, 1997.
[29] ISO/IEC International Standard 13818-1, Generic Coding of Moving Pictures and Associated Audio Information: Systems, 1995.
[30] M. Ghanbari, C.J. Hughes, Packing coded video signals into ATM cells, IEEE/ACM Trans. Networking 1 (5) (1993) 505–509.
[31] ITU-R 601, Radiocommunication Recommendation.
[32] E.W. Biersack, Performance evaluation of forward error correction in an ATM environment, IEEE J. Select. Areas Commun. 11 (4) (1993) .
[33] A.J. McAuley, Reliable broadband communications using burst erasure correcting code, ACM SIGCOMM '90, Philadelphia, PA, September 1990, pp. 287–306.
[34] A. Mehaoua, R. Boutaba, G. Pujolle, An adaptive and selective cell drop policy with dynamic data partitioning for best effort video over ATM, 22nd IEEE Int. Conf. on Computer Networks (LCN'97), Minneapolis, MN, November 1997.
[35] G. Hebuterne, A. Gravey, A space priority queuing mechanism for multiplexing ATM channels, ITC Specialist Seminar, 1989.
[36] O. Rose, Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems, TCOST 242, Cambridge, January 1995.
[37] O. Rose, M.R. Frater, Impact of MPEG video traffic on an ATM multiplexer, IFIP High Performance Networking '95, Palma, September 1995, pp. 157–168.
[38] H. Li, K.-Y. Siu, H.-Y. Tzeng, TCP over ATM with ABR service versus UBR + EPD service, Atm-forum-95-0718, June 1995.
[39] H.Li, K.Y. Siu et al., TCP performance over ABR and UBR services in ATM, IPCCC'96, March 1996.

Ahmed Mehaoua received the M.S. and Ph.D. degrees in Computer Science in 1994 and 1997 from resp. University of Paris V and University of Versailles, France. Since February 1996, he is a member of the research staff in telecommunication and distributed systems at the Computer Science Research Institute of Montreal, Canada. He was involved in several research projects concerning ATM network management and digital video networking. His primary research interests are traffic management and congestion control issues in broadband networks, and quality of service management of multimedia applications over ATM and Internet networks.

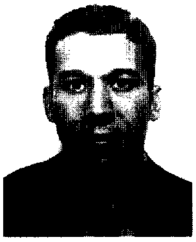

Raouf Boutaba has created and leaded the telecommunication and distributed systems unit at the Computer Science Research Institute of Montreal since 1995. He is also an Adjunct Professor at the University of Montréal, Québec, Canada. He received an "Engineer" degree in Computer Science and Engineering from the University of Annaba, Algeria, in 1988, and M.S. and Ph.D. degrees from the University of Paris VI, in 1990 and 1994, resp. He was Assistant Professor at Evry University, France and technical leader in national and international projects including ESPRIT and RACE. His research interests include integrated network and system management, wired and wireless high-speed networks, and distributed and multimedia systems and applications. He is currently directing several research projects sponsored by the computer and telecommunication industry in these areas. He has served as general chair and co-chair of several international conferences, and as a guest editor of international journals.