

An Integrated Framework for Efficient Transport of Real-Time MPEG Video over ATM Best Effort Service

This paper describes a framework for the transport of real-time multimedia traffic generated by MPEG-2 applications over ATM networks using an enhanced UBR best effort service (UBR+). Based on the factors affecting the picture quality during transmission, we propose an efficient and cost-effective ATM best effort delivery service. The proposed framework integrates three major components: 1) a dynamic frame-level priority assignment mechanism based on MPEG data structure and feedback from the network (DexPAS), 2) a novel audio-visual AAL5 SSCS with FEC (AV-SSCS), and 3) an intelligent packet video discard scheme named FEC-PSD, which adaptively and selectively adjusts cell drop levels to switch buffer occupancy, video cell payload type and forward error correction capability of the destination. The overall best-effort video delivery framework is evaluated using ATM network simulation and MPEG2 video traces. The ultimate goal of this framework is twofold: First, minimizing loss of critical video data with bounded end-to-end delay for arriving cells and second, reducing the bad throughput crossing the network during congestion. Compared to previous approaches, performance evaluation shows a good protection of Predictive coded (P-) and Bi-directional Predictive coded (B-) frames at the MPEG video slice layer.

© 2001 Academic Press

Ahmed Mehaoua¹, Raouf Boutaba², Yasser Rasheed² and Alberto Leon-Garcia²

¹*University of Cambridge, CCSR, 10 Downing Street, Cambridge, United Kingdom
E-mail a.mehaoua@ccsr.cam.ac.uk*

²*University of Toronto, ECE Dept., 10 King's College Road, Toronto (On.), Canada
E-mail: {rboutaba, yrasheed, alg}@comm.utoronto.ca*

Introduction

MPEG-2 and ATM have been adopted as key technologies for the deployment of broadcast and interactive video services. The confluence of these two international standards aims to provide all the advantages of transmitting variable bit rate video over packet networks, i.e. better video quality, less delay, more simultaneous connections, and lower cost.

Asynchronous transfer of video requires careful integration between the network and the video end systems. A number of issues must be addressed in order to tackle the problem on an end-to-end basis. Among these issues are the selection of: the ATM bearer capacity, the ATM adaptation layer, the mechanism of encapsulation of MPEG-2 packets into AAL, the scheduling algorithms in the ATM network for control of delay and jitter, and the error control and correction schemes.

Various proposals have been made for selecting the type of service under which MPEG-2 video streams are to be transported over ATM [1–5]. Unspecified Bit Rate is one of the simplest ATM best effort services available.

It is expected that this service will be widely available in the future since it is based on the excess bandwidth in the network with lower usage cost. It is also predictable that it will support a non-negligible part of the multimedia traffic. Unfortunately, UBR as initially defined in [6], is not appropriate for carrying such demanding traffic. This paper particularly focuses on unidirectional delay tolerant video applications and attempts to enhance the UBR transport service to efficiently support them.

In order to ensure optimal end-to-end quality, each component along the transmission path must be designed to provide the desired level of QoS. Therefore, optimizing only specific components in the path may not be sufficient for ensuring the QoS desired by the application. For example, designing a good Forward Error Recovery (FEC) scheme for the adaptation layer while using a poor cell discarding algorithm (e.g. randomly discarding) for the switch might not be sufficient to maintain the end-to-end performance of video application at the receiver. Consequently, the adaptation layer, encapsulation scheme, scheduling discipline in the ATM switches and error recovery mechanisms at the receiver must all be cooperatively designed and harmonized to provide the desired level of quality at the receiver (i.e., end-to-end). The framework proposed in this paper integrates the three following schemes: An AAL5 Service Specific Sub-layer with FEC control capability, an intelligent video data encapsulation and prioritization mechanism located at the source, and an efficient cell scheduling policy with adaptive video slice drop at the switch.

The paper is organized as follows. In the next section, we present a brief summary of the state of the art in video transport over packet networks. Following this, we describe the different components of the proposed best effort video delivery service with particular focus on the scheme for dynamic video cell priority assignment, and the intelligent packet video discard scheme. Next, we evaluate the performance of the framework using simulation, and discuss the obtained results. Finally, we conclude the paper.

Related Work

Various data protection and recovery techniques have been proposed to cope with the problem of transmitting compressed video streams over lossy networks. These schemes attempt to minimize picture quality degradation induced by data loss during network congestion.

Data protection and recovery techniques are usually implemented at the adaptation layer or above such as layered video encoding with data prioritization, which is one of the most popular approaches [7, 8]. Forward Error detection and Correction (FEC) techniques associated with byte interleaving and error concealment mechanisms at the destination have also been proposed to address this issue [5, 9, 10].

At the network level, smoothing algorithms attempt to reduce the burstiness and the peak bandwidth of video streams prior to transmission by applying complex shaping and buffering techniques at the source [11, 12]. This should minimize switch buffer oscillation, ease cell scheduling and thereby reduce cell loss probability.

To provide differentiated classes of service to the connections and ease the cell scheduling within the network, data priority assignment at the link or ATM layer is a powerful and cost-effective strategy. In the context of MPEG video communication, several implementations have been proposed.

Human perception is less sensitive to low frequency components of a video signal. Therefore in [7], the 8×8 DCT transformed video blocks are partitioned into an essential or base layer (comprising the lowest frequency DC coefficients), and an enhancement layer (consisting of the set of high frequency AC coefficients). The information contained in the base layer is then packetized and transmitted at high priority (HP), while information in the enhancement layer is transmitted at a low priority (LP) with a best effort delivery service. The cell loss priority (CLP) bit in the ATM header is used to provide a two-level block-based cell priority mechanism within a single channel.

In [13], the authors proposed to adapt the previous approach to the macroblock layer. The DC value is still assigned to the HP stream, and the macroblock header, and the motion vectors for the predictive frames (i.e. P- and B-) are also included. The remaining 63 DCT coefficients are split into two sub-streams according to a predefined parameter β . β specifies the number of AC

coefficients that are to be placed in the HP stream. The remaining $(63-\beta)$ coefficients are transmitted in the LP stream. To allow the regeneration of the original bit stream by the destination, the macroblock address is joined to the LP information.

These two techniques send the HP and LP video data onto the same virtual channel using the same ATM service class. In [14], a novel approach is evaluated performing a connection-level prioritization. The base layer and the enhancement layer of a hierarchically encoded MPEG-2 video are transmitted over distinct channels, i.e.; a VBR-rt guaranteed VC and an ABR best effort VC respectively.

The drawbacks of these priority techniques are the high complexity (i.e. bit stream parsing) and the special devices required at the destination to synchronize and recover the original video stream. Thus, simpler approaches have been proposed in [15] and [16], where video data partition and priority assignment are implemented at the video frame layer. The cells belonging to MPEG frames are set to different priority level according to the current frame coding mode. I-frame cells are assigned high priority over P- and B-frame cells. In [15], a static priority partition is proposed, while in [16], a dynamic and adaptive priority assignment is preferred with reference to the network congestion level.

Additionally, in order to cope with the problem of packet fragmentation and poor transmission performance of traditional packet services (IP, Frame-relay,) over ATM, some mechanisms have been designed to preserve packet integrity and achieve higher good throughputs. *Packet Tail Discard* or *Partial Packet discard* (PPD) has been proposed first to address this problem [17]. If a cell is dropped by a switch, the subsequent cells of the packet are also dropped. Romanov *et al.* have shown that PPD improves network performance to a certain degree, but it is still not optimal [17]. Consequently, they proposed a new mechanism called *Early Packet Discard* (EPD) that achieves better throughput performance but does not guarantee fairness among the connections [9]. When the switch buffer queue reaches a threshold, entire newly arriving packets (for example, AAL5 PDU) are preventively discarded. To improve its fairness, selective packet drop based on per-Virtual Circuit accounting have been introduced by Heinanen and Kilkki and the scheme is referred to as *Fair Buffer Allocation* (FBA).

Since MPEG video is a packet-oriented application, it may be interesting to apply similar approaches to the video streams.

As video 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, adaptation of PPD and EPD to the slice layer are proposed in [18]. The enhanced schemes, referred to as Adaptive Partial video Slice Discard (A-PSD) and Adaptive Early Video Slice Discard (A-ESD), are evaluated in their ability to gracefully degrade picture quality during network congestion and optimize network resource utilization without introducing noticeable visual artefacts.

Nevertheless, none of the previously mentioned smart data packet discard schemes are considering Forward Error Correction and Error Concealment capabilities of the end terminals. Therefore, we are proposing in the following section an enhancement to PPD in order to intelligently stop video cell discard as early as congestion stops and the forward error correction mechanism can effectively recover the missing data.

A Best Effort Video Packet Transport Service

A dynamic extended priority assignment scheme

Since the ATM cell header only includes one Cell Loss Priority (CLP) bit to discriminate between video data, it can not capture the full range of MPEG data structures. Thus, we propose a video data formatting and prioritization scheme based on the Extended CLP (ExCLP) field [8] and the Dynamic-Priority Assignment Scheme [16]. The new mechanism is referred to as the "Dynamic and Extended Priority Assignment Scheme" (DexPAS). The mechanism is sufficiently generic to be performed at any MPEG data layer (e.g. frame, slice, macroblock or block).

In this paper, the emphasis is on the video slice and frame layers. The data encapsulation is made at the video slice layer and the priority assignment is performed at the frame level. According to ITU-T and ATM Forum specifications, there are only two loss priorities in a single ATM Virtual Channel Connection. This is explained by the fact that the classical CLP bit and the adjacent PTI ATM-user-to-ATM-user bit (AUU) are interpreted separately by the network nodes. This operation underutilizes the ATM cell header and

restricts the number of cell priority to two although three priorities can be supported with the PTI-AUU.

DexPAS uses the ExCLP field to dynamically assign cell priorities according to the current MPEG frame type, e.g., (I)ntra (P)redictive or (B)I-directional predictive, and the reception of backward congestion signals from the network (see Figure 1).

Table 1 presents the mapping of MPEG data frames into the ExCLP field. Cells belonging to Intra-coded frames (I-cells) are assigned a high priority while B-frame cells (B-cells) have the lowest priority. P-cells are alternatively assigned a high or a low priority depending on the network load. At the beginning of the transmission, P-cells are initialized with a high priority. When the switch buffer Queue Length (QL) exceeds an upper threshold, an early congestion is detected and the switch sends a feedback signal to the source, which in turn adjusts P-cells priority level to low. When QL decreases below a lower threshold, P-cells priority are switched back to a high priority with an IP/B Resource Management cell. The cells having their ExCLP field set to “10”, are referenced as “End of control Block” (EOB) and delimits a group of video cells under FEC control. The PTI AUU bit is commonly employed to indicate whether it is the last cell of an upper message (e.g. TCP packet).

We propose to define a similar flag to distinguish between successive video slices. The cell having its ExCLP flag set to “11” is referred to as the End of video Slice (EOS) cell. Both EOB and EOS cells will be treated as “*very high*” priority cells in our implementation, that is, they are preserved with the most effort. As a result, DexPAS takes advantage of both static I/PB and static IP/B priority partitioning techniques [15]. Moreover, it extends ATM capabilities to provide up to four priority levels whereas the traditional approach restricts the number of possible cell types to only two.




Figure 1. DexPAS Operation with transmission of RM cells.

Table 1. New ExCLP field mapping for DexPAS

Cell Type	CLP	PTI-AUU	Priority
I-/P-frame	0	0	High
P-/B-frame	0	1	Low
End of CB	1	0	Very high
End of slice	1	1	Very high

As evaluated in [19], this dynamic priority assignment strategy minimizes loss of critical video frames and provides better performance than static CLP-based techniques. The main drawback of the scheme is that its efficiency is stringently dependent on the round trip time delay, and consequently on both network topology and link distances.

An audio visual SSCS for AAL5 with FEC

Classical AAL5 only provides error detection by means of CPCS packet length integrity and CRC-32 checks. It is not possible to locate which cell was dropped or which cell includes bit errors. Therefore, the task of the proposed Video Service Specific Convergence Sublayer is to implement a robust Forward Error Correction (FEC) mechanism targeted to MPEG video transmission.

The proposed FEC-SSCS protocol is based on both Reed-Solomon [20] and Parity Codes.

Four grouping modes are defined at the Service Specific Convergence Sub-layer that ensure an integer number of 48-byte cell payloads at the SAR layer and thus, no byte stuffing. These modes group a number “*N*” of MPEG-2 TS packets to build a SSCS-SDU. The parameter “*N*” may take the values: 3, 15, 27 and 39. After appending the CPCS-trailer information, we respectively obtain exactly 12, 59, 106 and 153 times 48-byte ATM cell payloads, which define in turn the number of ATM SDU per CPCS-PDU (i.e. parameter “*P*”) and the associated Cell Drop Tolerance “*T*”.

At the SSCS, a two-byte header and a two-byte trailer information are appended to every SSCS SDU as in Figure 2(b). The header is composed of a 4-bit Sequence Number (SN), a four-bit Sequence Number Protection (SNP), a four-bit Payload Type (PT), and a four-bit Control Block Length (CBL).

The trailer is composed of a two-byte Forward Error Correction field (FEC) applied only to the payload. The

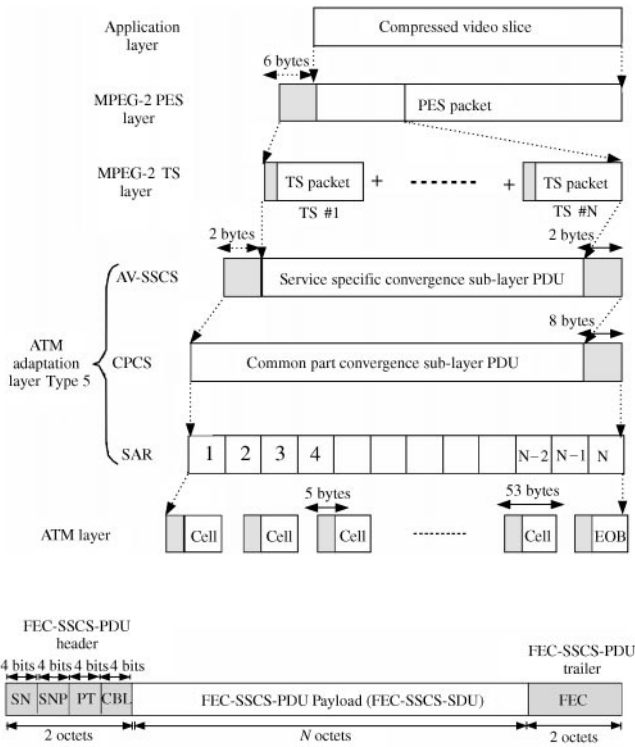


Figure 2(a). The Enhanced Audio-Visual AAL5 with FEC.
 Figure 2(b). The AV-SSCS Protocol.

FEC scheme uses a Reed-Solomon (RS) code, which enables the correction of up to two erroneous bytes in each block of 564 bytes (for example 3×188). So, it is only used for recovering of cell errors due to electrical or physical problems along the communication path. The sequence number (SN) of four bits enables the receiver entity to detect and locate up to 15 consecutive SSCS PDU losses. The SNP contains a three-bit CRC, and the result is protected by an even parity check bit. The PT field specifies the type of embedded information for discrimination purpose (I-frame, P-frame, B-frame, Audio, Data, Headers, FEC information, etc.).

Let us define a Control Block (CB) as a two dimensional matrix of P cells column \times M rows (see Figure 3). A single redundancy row is appended at the tail of the matrix, and is obtained by XORing the columns at the byte basis.

The parameter “ M ” is referenced as “Control Block Length” (CBL), and may be negotiated at the call set up with reference to the protection level desired by the connection.

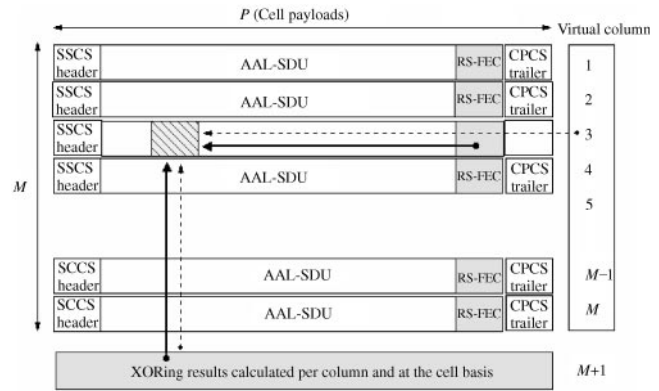


Figure 3. An adaptive FEC scheme using SSCS Control Block structure. \rightarrow Writing and reading order; $\leftarrow \blacktriangleright$ Localization (parity and sequence number checks); $\bullet \rightarrow$ Correction (RS and XOR codes). \square Erroneous or lost cell; \blacksquare FEC information.

At destination, three tasks have to be performed by the FEC SSCS receiver entity: (1) detecting error or loss in the incoming stream, (2) localize the missing cells or the position of the erroneous bytes, and finally (3) recovering the initial data.

- (1) Both SSCS and CPCS protocols assure the detection of erroneous SSCS PDUs. CPCS layer is able to identify received corrupted AAL PDUs by CRC-32 and missing cells by length mismatch. In the extreme situation of missing entire PDUs, the sequence number will permit the detection to up to 15 consecutive packet losses.
- (2) The association of the parity FEC XOR check sequence allows the FEC-SSCS layer to locate the erroneous bytes by determining simultaneously the line and the column numbers, as shown in Figure 3. Moreover, taking benefit of the fixed length of both MPEG-2 TS packets and ATM cells, the SSCS layer is capable of easily locating the missing cell.
- (3) After localization, both errors and losses can be corrected using Reed-Solomon and/or XORed FEC check codes.

A FEC-aware partial video slice discard scheme.

Due to the hierarchical nature of MPEG syntax, Random cell Discard (RD) during congestion may not be suitable for video transmission. An improvement is to take into consideration the cell’s priority when discarding, i.e., a cell with low priority is dropped first; if congestion persists, this approach gradually begins to

drop the high priority cells. This is called Selective Cell Discard (SCD). However, the useless cells, in our case the tail of corrupted slice, may still be transmitted and congest upstream switches. In [19], a scheme called Adaptive Partial Slice Discard (A-PSD) has been proposed to cope with this problem. The proposed approach consists to select the packet (i.e. slice) to be dropped with respect to MPEG data structure and congestion level.

In [21], we have proposed enhancements to the Adaptive Partial Slice Discard (A-PSD) to support Forward Error Correction feature. The new scheme, named FEC Adaptive Partial Slice Discard (FEC-PSD), is performed at both control block (CB) and video slice levels. Our approach is to reduce the number of corrupted slices, knowing that a number T of cells per control block can be recovered by the destination SSCS using FEC based on both Reed-Solomon and Parity codes. Let us define the parameter T as the Drop Tolerance (DT) which corresponds to the maximum number of cells per CB that may be discarded by A-PSD before considering the CB as definitely lost.

Therefore, unlike the simple A-PSD, FEC-PSD stops discarding cells when the congestion decreases and the number of previously dropped cells in every CB is below DT. Using this approach, the proposed scheme acts at a finer data granularity, e.g., Control Block, and better preserves entire slices from elimination. The flexibility proposed by this mechanism can not be achieved without the use of DexPAS which allows the detection of both slice and control block boundaries at the cell level.

The integration of the two mechanisms (e.g. DexPAS, FEC-PSD) with the enhanced AAL5 AV-SSCS provides us with an efficient and intelligent video delivery service with quality of picture (QoP) control optimization. The aim of this scheme is to ensure graceful picture degradation during overload periods as well as increase of network performance, e.g., effective throughput. It allows accurate video cell discrimination and progressive drop by dynamically adjusting the FEC-PSD mode in respect to cell payload types, switch buffer occupancy, and Drop Tolerance.

Let us define a low (high) priority slice as a slice belonging to a low (high) priority frame, as indicated in Table 1. During light congestion, we propose to drop a lower priority slice first rather than delaying it. Then we could assign the buffer space of the dropped slice to a

higher priority slice. The proposed approach avoids congestion increase while maintaining the mean cell transfer delay at an acceptable value. This proactive strategy is performed gradually by including high priority cells if necessary. As evaluated in [19], the proposed approach can significantly improve the network performance by minimizing the transmission of non-useful video data before buffer overflow. The proposed FEC-aware Partial Video Slice Discard (FEC-PSD) algorithm is highlighted below.

FEC-PSD parameters

The Selective and Adaptive FEC-aware Partial video Slice Discard (FEC-PSD) scheme is implemented per-VC in the switches and employs four state variables and one counter variable to control each video connection. Two of these parameters are associated with the slice level and the remaining ones are associated with the control block level.

1. **S_PRIORITY** indicates the priority level of the current slice. The indicator is modified at the reception of the first cell of this slice in respect to its priority field (the two ExCLP bits, in our case). This indicates that the switch is currently handling a high ($S_priority=0$), or a low ($S_priority=1$) priority slice.
2. **S_DISCARDING** indicates whether the switch is currently discarding ($S_discarding=1$) this slice, e.g., the tail, or not ($S_discarding=0$). Only the last cell of a slice (EOS) can change this indicator from *discarding* to *not discarding*. Other cells will only change the flag from *not discarding* to *discarding*, when necessary.
3. **CB_DROPPED** is a counter that indicates that for the current control block the number of cells discarded by the switch. It is initialized to zero at the reception of a new control block. This is needed so that we can check whether a control block is still recoverable or not.
4. **CB_DISCARDING** indicates whether the switch is currently discarding ($CB_discarding=1$) the current control block or not ($CB_discarding=0$). In contrast to the slice level control, the indicator changes from *discarding* to *not discarding* in two situations: the **CB_DROPPED** counter reaches the Drop Tolerance T ; a new block is received. Other events, e.g., cell arrivals, will only change the flag from *not discarding* to *discarding*, when necessary.
5. **CB_EFCI_MARKING** indicates whether the switch is tagging ($CB_EFCI_MARKING=1$) or not

tagging ($CB_EFCI_MARKING = 0$) the EFCI bit of the cell for the current control block. Only the last cell of a block (EOB) can change this indicator from *marking* to *not marking*. In addition, only one event may provoke the modification of the state from *not marking* to *marking*. This occurs when the arrival of a cell is concurrent with the $CB_DISCARDING$ indicator is in a *no discarding* state, and $CB_DROPPED$ equals the tolerance T .

The use of both $CB_DISCARDING$ and $CB_EFCI_MARKING$ indicators allows us to manage losses occurring at subsequent switches and belonging to a control block more efficiently. Indeed, when a block is partially discarded by a switch node, the following switches are not capable of taking into account these cell losses in order to update the associated drop tolerance. As a consequence, the switches handle erroneous cell drop tolerance with an adverse effect on the performance. At the control block level, the drop tolerance can be seen as a loss credit shared by the crossed switches.

In order to simplify the implementation, we propose to entirely consume the loss credit as soon as a cell loss occurs. $CB_DISCARDING$ is used to ensure that, for every control block, losses are concentrated in a single switch. If cells from a block tail arrive in a congested node, the use of EFCI bit allows the detection of non-recoverable blocks since a previous switch has used the entire drop credit. In such a situation, we propose to commit to the slice level control by entirely dropping the remaining slice.

FEC-PSD operation modes and fairness

FEC-PSD uses three buffer thresholds as shown in Figure 4: Low Threshold (LT), Medium Threshold (MT), and High Threshold (HT). The utilization of three thresholds, instead of two, reduces the speed of oscillation for the transmission of Dex-PAS RM cells and exhibits better performance.




Figure 4. FEC-PSD buffer architecture and operation mode.

The thresholds define three operation modes (see Figure 4.), which in turn restrict the distribution of the cell loss within a Control Block at the receiver to four when FEC-PSD has been activated in the network (see Figure 5). The dashed lines present the four possible locations of discarded cells in a Control Block.

1. *Mode Idle*: If the buffer queue length (QL) is lower than “Low Threshold”, for every connection, the cells are accepted and may have EFCI marked, whenever $CB_EFCI_MARKING$ is activated.
2. *Mode 1*: If the total number of cells in the buffer exceeds “Low Threshold” but is still below “High Threshold”, for every video connection currently emitting a low priority slice, FEC-PSD starts to discard incoming cells. The discarding is done in respect to the drop tolerance associated with each connection. If the light congestion is subsisting, the algorithm switches to the slice level, and starts to eliminate any incoming low priority cell until receiving an EOS cell. The last cells (EOS) are always preserved from elimination since they provide indication of the next slice. Cells with higher priority are always accepted in the buffer. This mode stops when “QL” falls below the “Low Threshold”.
3. *Mode 2*: This mode is activated when QL exceeds “High Threshold”. Incoming slices are discarded regardless of their priority level. The last cell of a control block and slice are preserved to avoid the error propagation. This is feasible, since usually 10% of switch buffer has been set aside to accommodate the system control and management messages as well as other important cells. This mode stops when queue length falls below “HT”.

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

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

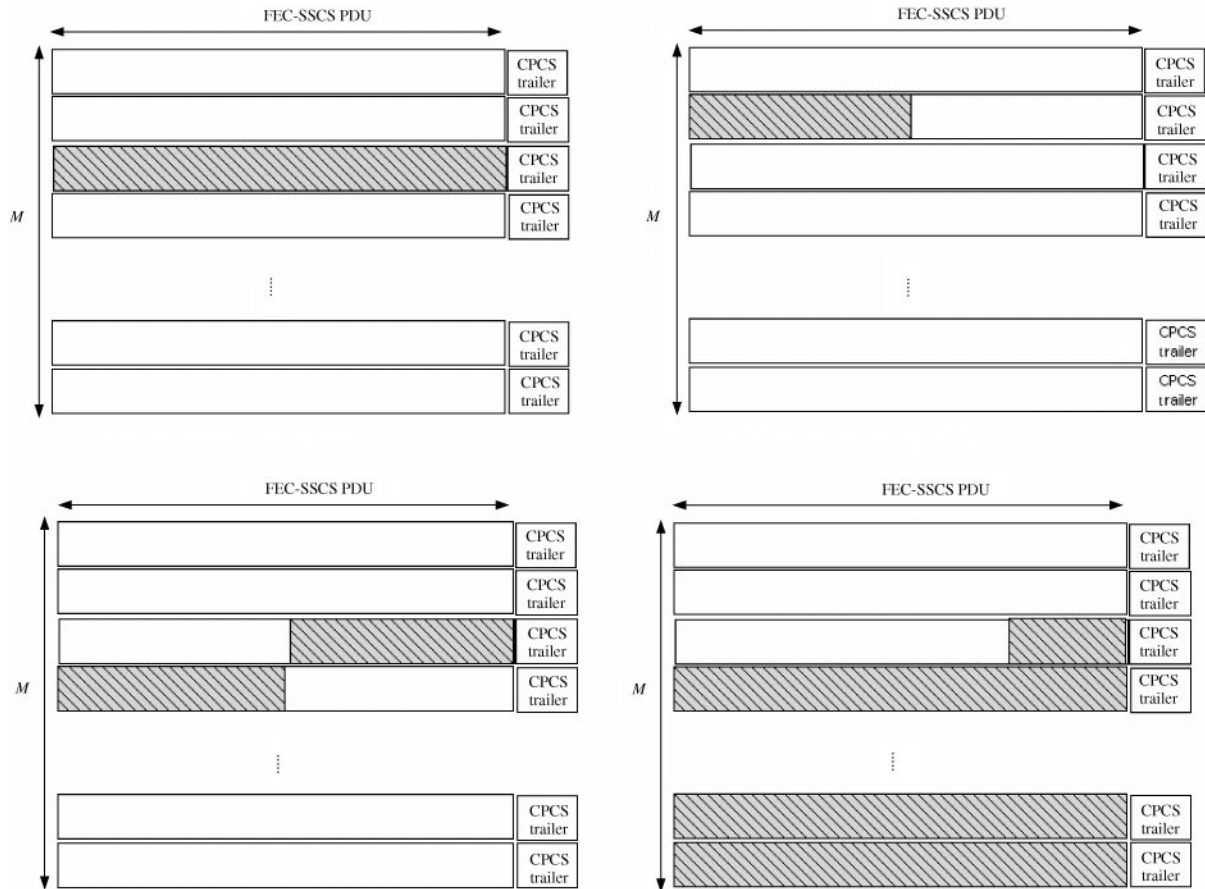


Figure 5. Cell Loss Distribution in a destination Control Block when FEC-PSD is activated; ▨ Erreous data.

Performance Evaluation

Network simulation model

The simulation network topology is depicted in Figure 6. It consists of two ATM switches, and 10 MPEG2 video connections crossing the backbone (bottleneck) link with a capacity of 155 Mbps (OC-3). We evaluate the framework in a LAN configuration by setting the




Figure 6. Network simulation model

physical backbone link distance to 1 km (one-way). Link distances between the source/destination and the access switch nodes are constant and set to 0.2km (one-way). The ATM switches are simulated as non-blocking, shared finite output-buffered switches. Switch buffer size varies from 80,000 to 220,000 cells for both SWITCH-1 and SWITCH-2 in the simulation configuration.

The video sources generate MPEG2 data at a rate specified in a trace file obtained from Michael R. Izquierdo, IBM Corporation. The video sequence shows a flower garden located in the bottom half of the screen and a row of houses in the background towards the top of the scene. The camera tracks this scenery from left to right. A detailed description of this file can be found in [22]. The video sequences uses SIF format and were encoded at a resolution of 352×240 pixels per frame, a frame rate of 30 frames/s, and 15 slices/frame. The

Peak and Mean Cell Rates are 20 and 5 Mbps respectively.

Figure 7 shows the number of ATM cells per slice for the first 20 frames. We notice distinctive pulses occurring at deterministic time intervals. The pulse period is determined by the GOP pattern, that is, every 45 slices. There are also alternating pulses caused by I and P-frames. The spacing between pulses are B-frames.

We use the same file for all of the senders. Since each sequence has the same I/P/B frame pattern, I-frames will always overlap for the duration of playback if the source sends video streams at the same time. For this reason, we shift the send time so that I and P-frames from one sequence overlap with B-frames from another. Figure 8 shows the results of multiplexing the shifted MPEG-2 traffic. No distinctive peaks and valleys are shown in contrast to the single sequences.

The level of congestion is monitored through the occupancy of the switch buffers and three congestion thresholds (LT, MT and HT). We carried out our simulation with seven switch buffer configurations. For each of them, the same method is applied to determine the values of the three thresholds. HT, MT and LT are respectively set to 0.9, 0.8 and 0.7 of the maximum queue size (Q_{max}), where Q_{max} is set to one of




Figure 7. Number of ATM cells per slice for the first 20 frames of the MPEG-2 video sequence.




Figure 8. Number of cells per slice slot time after multiplexing of all sources with time shift.

the following values: 80,000; 100,000; 120,000; 140,000; 160,000; 180,000; 200,000 and 220,000 cells.

The transfer delay is measured as follows:

- Propagation delay between the sender and the receiver is 0.005 ms which corresponds to the propagation distance of about 1 km.
- Queuing delay varies from 0 to a maximum value of 0.6 s, which corresponds to the maximum buffer size of 220,000 cells, when transmitted using 155 Mbps link (OC-3).
- The process delay for the sender can be assumed as negligible, due to pipelined data transmission and encoding of the appended data. At the receiver, following FEC processing time for error recovery in SSCS layer is assumed:
 - SSCS-FEC (with error): 0.46 ms/slice. ($12 \times M$ cell transmission time for 55 Mbps link, where $M = 5$ in most of our cases)
 - SSCS-FEC (without error): 0.092 ms/slice. (12 cell transmission time for 55 Mbps link)

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

An AV-SSCS-PDU contains three MPEG2 Transport Streams (12 cells) and a Control Block is built with 15 AV-SSCS-PDUs (except for the last CB). Consequently the operation parameters are :

- $N=3$
- $M=15$
- $P=T=12$

The video Slice Loss Ratio (SLR) is measured at the MPEG2 application layer and takes into account the effect of cell loss, and propagation delay. In addition, it also takes into consideration FEC capacity to decide if a slice is recoverable or not at destination.

Table 2, and 3 summarize the possible states of a cell crossing the network and the investigated performance parameters respectively.

The Slice Loss Ratio (SLR) is measured at the application layer and take into account decoding, e.g., cell loss, and propagation, e.g., late cells, constraints. In addition, it also takes into consideration FEC capacity to decide if a slice is usable or not.

We compare the performance of the proposed framework DexPAS associated with FEC_PSD (i.e. Dex_FEC_PSD) at the video slice level with the three following schemes associated with the classical AAL5:

- Random Discarding with no Priority Assignment Scheme (No_RD).
- Selective Cell Discarding with Extend Priority Assignment Scheme (Ex_SCD [8]).

Table 2. Data unit definitions

Data units	Definition
Lost cell	A cell dropped by discarding 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 $1/N$ sec., where N is the frame rate of the video sequence. In this paper, N is equal to 30
Correct cell	Neither a lost, dead or late cell
Correct slice	A slice received with only correct cells

Table 3. Performance parameters

Performance Parameters	Definition
I-frame cell loss ratio (CLR_i), P-frame cell loss ratio (CLR_p), B-frame cell loss ratio (CLR_b),	Number of lost and late cells belonging to I-frames from the three connections vs. the total number of transmitted I-cells. The same metric is applied for P- and B-frames
Cell bad throughput (CB)	Number of dead cells vs. the total transmitted cells. It is a performance parameter evaluated at the ATM layer
I-frame video slice loss ratio (SLR_i) SLR_p and SLR_b	Number of corrupted I-frame slices vs. number of transmitted slices. The same metric is applied for P- and B-frames
Mean cell transfer delay (mean CTD)	Time between the departure of cell K from the source node (t_{iK}) and its arrival at the destination node (t_{oK}): $DK=t_{oK}-t_{iK}$

- Partial Slice Discarding with Extend Priority Assignment Scheme (Ex_PSD [18]).

Results analysis at the video slice layer

From Figure 9, we observe that the mean cell transfer delay (CTD) increases in proportion to the buffer size.




Figure 9. Mean cell transfer delay. — Dex_PSD; ○ ○ ○ ○ Ex_PSD; + + + + Ex_SCD; - - - - No_RD.

As expected, No_RD has the largest mean CTD, since the No_RD scheme attempts to accommodate every cell in its switch buffer until it overflows, meanwhile increasing the queue delay. We also notice that Dex_PAS + FEC_PSD (Dex_PSD) has a longer mean-CTD than the other two schemes even though it tries to drop low priority cells at the light congestion stage in order to leave space for the high priority ones. This is mainly due to its overhead, which results in larger switch buffer occupancy. On one hand, it preventively discards low priority cells at light congestion and switches to slice level to discard the whole slice as in Ex_PSD, which reduces the average queue length. On the other hand, it introduces 15% overhead due to stuffing bits and FEC redundancy codes which in turn increases the average queue length.

Figures 10, 11, 12 and 13 show the variation of the buffer queue length in bytes for the switch_1 versus the simulation time when running the different cell drop schemes (RD, SCD, PSD). The horizontal axis represents the time period in ms while the vertical axis represents the queue size in bytes.

With No_RD (Figure 10), the buffer must be filled until the High Threshold is reached before it starts the elimination, which leads to a greater mean CTD.

Ex_SCD (see Figure 11) and Ex_PSD (see Figure 12) start to drop B-cells when light congestion occurs and




Figure 10. Buffer occupancy with No_RD.




Figure 11. Buffer occupancy with Ex_SCD.

thus reduce the buffer occupancy while minimizing the transfer delay of the high priority cells.

The difference of CTDs between different schemes increases with larger buffer size. For instance, with limited buffer size, the difference between random drop




Figure 12. Buffer occupancy with Ex_PSD.




Figure 13. Buffer occupancy with Dex_FEC_PSD.

and preventive discarding schemes is small, whereas when Q_{max} increases, it becomes larger. This is explained by the fact that the preventive B-frame cells elimination approach works better when more buffer space is available. With limited buffer size, the space saved by dropping B-cells is limited and therefore it performs similarly to No_RD.

Intuitively, it is expected that Dex_PAS + FEC_PSD (Dex_FEC_PSD in the figure) has a better performance at slice level. This is exhibited by Figures 14, 15, 16, and 17. The proposed framework significantly improves the percentage of arrivals at the destination of non-corrupted video slices. Indeed, the aggregated Slice Loss Ratio (SLR) is reduced to achieve an upper bound of 6.8% of the total number of transmitted video slice. In comparison, No_RD, Ex_SCD and Ex_PSD reach 16.6%, 12.2% and 8.9% respectively.

Finally, the SLR per sub-flow is analysed for the four approaches as follows. We observe that Ex_PSD and Dex_FEC_PSD outperform the other approaches by better protecting I-frames, though for aggregate SLR our new scheme has the best performance. This is consistent with the results obtained at cell level. There is a trade-off between fair distribution of cell discarding among the connections (i.e. VCs) and the speed of reactions to congestion. With B-frames,




Figure 14. Slice loss ratio (aggregate stream). — Dex_FEC_PSD; ○ ○ Ex_PSD; + + + + Ex_SCD; - - - - No_RD.

Dex_PAS and FEC_PSD (Dex_FEC_PSD) demonstrate the best SLR value, and perform correctly with P-frames. This further indicates the capability provided to protect data at the slice level by the FEC mechanism based on Parity [23] and Reed-Solomon [20] correction codes.




Figure 15. Slice loss ratio (I-frame). — Dex_FEC_PSD; ○ ○ Ex_PSD; + + + + Ex_SCD; - - - - No_RD.




Figure 16. Slice loss ratio (P-frame). Dex_FEC_PSD; ○○○ Ex_PSD; ++++ Ex_SCD; ---- No_RD.

Conclusion

In this paper, we proposed and evaluated an enhanced best-effort video delivery service based on UBR+, that takes into account the specific encoding and stochastic properties of MPEG2 video sources. This service is composed of three components: a new video data encapsulation and priority assignment technique called Dynamic Extended Priority Assignment Scheme (Dex_PAS), an intelligent packet video drop policy




Figure 17. Slice loss ratio (B-frame). Dex_FEC_PSD; ○○○ Ex_PSD; ++++ Ex_SCD; ---- No_RD.

named FEC_PSD and an audio-visual AAL5 SSCS with FEC.

By providing three different priority classes per connection and the detection of video slice and FEC Control Block boundaries at the cell level, Dex_PAS permits accurate cell discrimination and progressive cell group discard and error recovery at the slice level.

Since we are assuming that a better decoding/display result can be obtained by minimizing error propagation within a Group Of Pictures by preserving referenced video frames from corruption while concentrating the data loss in B-frames, the proposed best effort delivery service shows better performance in the transmission of non-corrupted high-priority video slices to the destination, which leads to graceful picture quality degradation and a higher throughput during network congestion.

References

1. Keshav, S., Grossglauser, M. & Tse, D. (1995) RCBP: A simple and effective service for multiple time-scale traffic. in *ACM Computer Communication Review* **25**: 219–230.
2. Kanakia, H., Mishra, P. & Reibman, A. Packet video transport in ATM networks with single-bit feedback. In: *Proceedings of the Sixth International Workshop on Packet Video*, Portland, Oregon, September 1994.
3. Lakshman, T.V., Mishra, P.P. & Ramakrishnan, K.K. Transporting compressed video over ATM networks with explicit rate feedback control. In: *Proceedings of IEEE INFOCOM 97*, Kobe, Japan, March 1997.
4. Zhang, H. & Knightly, E. W. (1997) RED-VBR: A renegotiation-based approach to support delay-sensitive VBR video. *ACM/Springer-Verlag Multimedia Systems Journal* **5**: 164–176.
5. Rasheed, Y. & Leon-Garcia, A. AAL1 with FEC for the Transport of CBR MPEG2 Video over ATM networks. In: *Proceedings of IEEE INFOCOM 96*, San Francisco CA, March 1996.
6. ATM Forum TM-SWG (1996) *Traffic Management Specification 4.0, at-tm-0056.000*, April 1996.
7. DeCleen, B., Pancha, P. & El Zarki, M. Comparison of priority partition methods for VBR MPEG. In: *Proceedings of IEEE INFOCOM'94*, Toronto, Canada, June 1994, pp. 689–696.
8. Mehaoua, A., Boutaba, R. & Pujolle, G. An Extended Priority Data Partition Scheme for MPEG Video Connections over ATM. In: *Proceedings of the IEEE Symposium on Computers and Communications 97*, Alexandria, Egypt, July 1997, pp. 62–67.
9. Wang, Y. & Zhu, Q.-F. Error Control and Concealment for Video Communication: A review. In: *Proceedings of the IEEE, Vol. 86, N.5*, May 1998, pp. 974–997.
10. Ramamurthy, G. & Raychaudhuri, D. Performance of packet video with combined error recovery and

- concealment. In: *Proceedings of IEEE INFOCOM'95*, Boston, U.S.A., April 1995, pp. 753–761.
11. Feng, W. & Sechrest, S. (1995) Smoothing and buffering for delivery of prerecorded compressed video. *ACM Computer Communications Review*, 709–717.
 12. Salehi, J.D., Zhang, Z.-L. Kurose, J.F. & Towsley, D. Supporting stored video: Reducing rate variability and end-to-end resource requirements through optimal smoothing. In: *Proceedings of ACM SIGMETRICS '96*, May 1996, pp. 222–231.
 13. Pancha, P. & El Zarki, M. (1994) MPEG coding for variable bit rate video transmission. *IEEE Communication magazine*, pp. 54–66.
 14. Mehaoua A. & Boutaba, R. (1997) A hybrid VBR/ABR Service for Layered Video Communications: A Simulation-based Analysis. *IEEE Workshop on Broadband Switching System '97 (BSS'97)*, Taipei, Taiwan, pp. 88–96.
 15. Han, T. & Orozco-Barbosa, L. (1995) Performance requirements for the transport of MPEG video streams over ATM networks. In: *Proceedings of ICC'95*, Washington, June 1995. pp. 221–225.
 16. Mehaoua, A., Boutaba, R. & Pujolle, G. An Adaptive and Selective Cell Drop Policy with Dynamic Data Partitioning for Best Effort Video over ATM. In: *Proceedings of the 22nd IEEE International Conference on Computer Networks (LCN'97)*, Minneapolis, E-U, 2–5 November 1997, pp. 519–528.
 17. Armitage, G. & Adams, K. (1993) Packet Reassembly during Cell Loss. *IEEE Network Magazine* 7: 26–34.
 18. Mehaoua, A., Boutaba, R. & Iraqi, Y. Partial versus Early Packet Video Discard. In: *IEEE GLOBECOM '98*, Sydney, Australia, November 1998.
 19. Mehaoua A. & Boutaba, R. (1998) Performance Analysis of Cell Discarding Techniques for Best Effort Video Communications over ATM Networks. *Journal of Computer Networks and ISDN Systems* 29: 2021–2037.
 20. Wicker, S.B. & Bhargava, V.K. (1994) *Reed-Solomon codes and their applications*. IEEE press.
 21. Mehaoua, A., Boutaba, R., Pu, S., Rasheed, Y. & Leon-Garcia, A. Towards An Efficient ATM Best Effort Video Delivery Service. In *Proceedings of IEEE ICC '99*, Vancouver, Canada, June 1999.
 22. Izquierdo, M.R. & Reeves, D.S. MPEG VBR slice layer model using linear predictive coding and generalized periodic markov chains. In: *Proceedings of International Performance, Computing and Communications Conference*, Scottsdale, Arizona, U.S.A., February 1997.
 23. Shacham, N. & McKenny, P. Packet Recovery in High Speed Networks using coding. In: *Proceedings of IEEE INFOCOM '90*, San Francisco, CA, June 1990, pp. 124–131.