

# Towards An Efficient Best Effort Video Delivery Service

Ahmed Mehaoua<sup>1</sup>, Raouf Boutaba<sup>2</sup>, Song Pu<sup>3</sup>, Yasser Rasheed<sup>2</sup>, Alberto Leon-Garcia<sup>2</sup>

<sup>1</sup> University of Cambridge - CCSR  
10 Downing Street  
Cambridge, United Kingdom  
am296@ccsr.cam.ac.uk

<sup>2</sup> University of Toronto - ECE Dept.  
10 King's College Road  
Toronto (On.), Canada  
{rboutaba, yrasheed,  
alg}@comm.utoronto.ca

<sup>3</sup> Newbridge Networks Corporation  
600 March Road  
Kanata (On.), Canada  
songpu@newbridge.com

## ABSTRACT

*This paper addresses the transport of real-time multimedia traffic generated by MPEG-2 applications over ATM networks using an enhanced UBR best effort service (UBR+). It analyses the factors affecting the picture quality during transmission. Based on this analysis, we propose an efficient and cost-effective best effort delivery service. The proposed service integrates three components: An intelligent packet video discard scheme, which adaptively and selectively adjusts cell drop level to switch buffer occupancy, video cell payload type and forward error correction ability of the destination; a dynamic frame-level priority assignment mechanism based on MPEG data structure and feedback from the network; and an Audiovisual Service Specific Convergence Sublayer for AAL5. The overall best-effort video delivery framework is evaluated using ATM network simulation and MPEG2 video traces. The ultimate aim of this framework is twofold. First, minimizing loss for critical video data with bounded end-to-end delay for arriving cells. Second, reducing the bad throughput crossing the network. Compared to previous approaches, performance evaluation shows a better utilization of network resources and a minimization of data losses of referenced Intra- and Predictive-coded frames at the video slice layer.*

**Keywords:** AAL, Packet Video, MPEG2, Cell Discard, FEC.

## I. INTRODUCTION

MPEG2 and ATM have been adopted as the 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 connections, and lower cost.

However asynchronous transfer of video requires careful integration between the network and the video systems. A number of issues must be addressed in order to tackle the problem on an end-to-end basis. Among these issues is the selection of: The service type, the adaptation layer, the method of encapsulation

of MPEG-2 packets in AAL packets, the scheduling algorithms in the ATM network for control of delay and jitters, and the error control scheme.

The adaptation layer is responsible for making the network behavior transparent to the application. AAL5 is currently the most commonly used adaptation layer in industry and can support VBR MPEG-2 traffic. However, AAL5 was initially designed to carry data traffic over ATM networks, which makes too simple to provide reliable connection for multimedia applications. Additional features such as error localization and recovery are required to improve AAL5 reliability for the transport of real-time MPEG-2 video data.

Different proposals have also been made for selecting the type of service under which MPEG-2 is to be transported over ATM [1][2][3][4]. Unspecified Bit Rate is the true and simplest ATM Best effort service available. Since it is expected that this service will be widely available in the future and is based on the excess bandwidth in the network with lower usage cost, it is predictable that it will also support a non-negligible part of the multimedia traffic. This paper particularly focuses on unidirectional delay-tolerant video applications that can efficiently make use of such simple and low-cost transport service.

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 scheme for the adaptation layer while using a poor cell discarding algorithm (e.g. randomly discarding) for the switch will 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). Therefore, the framework proposed in this paper integrates the three following schemes: An AAL sub-layer with FEC control capability, an intelligent video data partition

and prioritization mechanism located at the source, and an efficient switch scheduling strategy with adaptive discarding technique.

The paper is organized as follows. We describe in Section 2 the different components of the proposed best effort video delivery service including a scheme for dynamic video cell priority assignment, an Audiovisual SSCS with FEC support, and an intelligent packet video discard scheme. In section 3, we evaluate the performance of the framework using simulations, and discuss the obtained results. Finally, we conclude the paper in section 4.

## II. A BEST EFFORT VIDEO TRANSPORT SERVICE

### 2.1 A Dynamic Extended Priority Assignment Scheme

Transmission of compressed video over ATM networks requires efficient data priority partition techniques. In association with intelligent cell discarding schemes, these techniques aim to minimize loss probability of critical information in the situation of congestion. Since the ATM cell header only contains one bit (CLR) to discriminate between video data, they are not able to efficiently capture MPEG data structure complexity. To better cope with the hierarchical MPEG video transmission requirements, we propose a new video data formatting and prioritization scheme named Dynamic Extended Priority Assignment Scheme (Dex-PAS). 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 slice and frame layers.

Indeed, the data partition is made at the slice layer and the priority assignment is performed at the frame level. In [5], a new cell header field located in the ATM call header is defined and referenced as Extended CLP (ExCLP). This field comprises the classical CLP bit and the adjacent PTI ATM-user-to-ATM-user bit (AUU) [6].

Used individually these two single bits define only three distinctive cells: high priority cell, low priority cell and End of Message (EOM) cell. Using them together permits a better utilization of the cell header with the definition of up to four available cell types within a single channel. We propose to modify the Extended Priority Assignment Scheme (Ex-PAS) introduced in [7], to support multi-layer forward error correction (FEC) described in the following section.

Dex-PAS uses Ex-CLP 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).




Figure 1 - DexPAS Operation with Transmission of RM cells

Table 1 presents the mapping of MPEG data frames into the Ex-CLP field. Whereas the traditional approach restricts the number of priority to two and under utilizes ATM capabilities. Cells belonging to Intra-coded frames have a high priority and their ExCLP flag is set to '00'. B-frames have the lowest priority and the associated cells have an ExCLP value of '01'. As to P frames, they 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 buffer queue length (QL) exceeds an upper threshold, a congestion is detected early and the ATM 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. In our implementation we use forward resource management (RM) cells with congestion indication flag (CI) marked to notify the destination and afterward the source. The '10' value is used to allow the design of a two-levels video-oriented cell discard scheme located at every switch along the connection path. The cell having its ExCLP field set to '10', is referenced as 'End of control Block' (EOB) and delimits a group of video cells under FEC control. The PTI ATM-user-to-ATM-user bit is employed to indicate whether it is the last cell of an upper message (e.g. TCP packet).

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

Table 1: New ExCLP Field Mapping

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 cell will be treated as of a very high priority in our implementation,

that is, they are preserved with the most effort. As a result, DexPAS takes the advantages of both static I/PB and static IP/B priority partition techniques [8]. Moreover, it extends ATM capabilities to provide up to four priority levels whereas the traditional approach restricts the number of possible cell types to three and thus under utilize ATM header.

As evaluated in [9], this dynamic priority assignation 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 of the round trip time delay, and thus of the network topology and link length.

To support this slice-based mechanism, enhancements to AAL-5 and efficient MPEG-2 video stream encapsulation strategy are required. These requirements are addressed in the following.

## 2.2 An AudioVisual SSCS for AAL5 with FEC

Classical AAL type 5 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 hierarchical MPEG encoded video transmission. Requirements of such a FEC-SSCS are enlightened in [10][11] and may be summarized as follows.

1. Compatibility with the specification of the existing AAL-5, e.g., compatibility with the current CPCS/SAR layers.
2. No modifications are required for the upper layer, e.g., in our case MPEG-2 Transport Stream or MPEG Program Stream);
3. Support of variable size data (e.g., slice or frame).
4. The amount of redundant data should be minimized.
5. Similarly to the ATM Forum's Video on Demand over ATM specification [12], byte padding should be avoided.
6. It would be interesting to adjust and negotiate FEC-SSCS parameters at the connection setup phase as well as during the session.
7. FEC-SSCS should be able to detect errors, localize them and finally correct them.
8. In order to avoid an increase of latency, SSCS SDU should be transferred in pipelining at the sender side. This way, no buffering is required and the processing cost is minimized.

9. At the peer destination, if no errors are detected, the packet should be forwarded to the upper layer with no delay (e.g., no buffering). The processing speed at the receiver entity should be as fast as with classical AAL5.

10. In order to recover a corrupted packet, buffering of previous packets should be avoided.

11. In order to avoid errors' propagation, slice boundaries have to be respected during cell filling.

12. A similar requirement should be applied to the frame boundaries.

The proposed FEC-SSCS protocol satisfies all the above requirements. It is based on both Reed-Solomon [13][14] and Parity Codes [15], and on the video packet encapsulation mechanism proposed in [16]. Compared to other mechanisms based only on Reed-Solomon codes with byte interleaving, our approach allows using of flexible matrix structure and correction granularity at the byte and the cell levels. Moreover, it better takes into account the fixed structures of MPEG-2 TS packet and ATM cell to avoid bit padding at the lower AAL-5 Common Part Convergence sublayer. Our mechanism can also be used selectively to protect separately audio, video and syntactic data (e.g., headers) and thus to minimize data control overhead. The algorithm is described bellow.

### Sender behavior

First, the TS packets are passed to the Specific Service Convergence Sub-layer by the MPEG-2 System Layer using message mode service with blocking/de-blocking internal functions [17], as illustrated in figure 2a. The following primitive is used: AAL UNITDATA request(ID, M, SLP, CI). The 'Interface Data' (ID) parameter specifies the exchange of MPEG-2 TS packet. The 'More' (M) parameter indicates if it is the last AAL SDU of the upper message (e.g., end of the current video slice). The 'Submitted Loss Priority' (SLR) parameter gives the priority level of the TS packet and is initialized according to the 'PICTURE-CODING-TYPE' field located in the MPEG frame header [18], the latter field specifies the used coding mode for each frame (e.g. Intra, Predictive or Bi-directional Predictive). This parameter also indicates how the 'SLP' parameter of the ATM-DATA-request primitive shall be set for cell header initialization. As described in section 2.1, we propose to extend its range from two to four possible values to allow identification of MPEG frame types and system information. Finally, the last parameter 'Congestion Indication' (CI) defines how the 'CI' parameter of the ATM-DATA-request primitive should be set to notify a congestion state both to the network nodes and destination.




Figure 2a - the Audio-Visual AAL5 with FEC

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 consist to group a number 'N' of MPEG-2 TS packets to build a SSCS-SDU. The parameter, 'N' may have the following 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 as illustrated in Table 2.

SSCS Grouping Mode (SGM)	SSCS Group Size (SGS)	FEC-SSCS SDU Size (Bytes)	FEC-SSCS PDU Size (Bytes)	SDU ATM (48 bytes)
SGM_3	3	564	568	12
SGM_15	15	2820	2824	59
SGM_27	27	5076	5080	106
SGM_39	39	7332	7336	153

Table 2 - The Four pre-defined grouping Modes (SGM)

For every connection, the grouping mode is negotiated between the source and destination at connection establishment phase, according to the required QoS. During user data transfer, this mode can be dynamically adjusted in respect to the on-line measures of the end-to-end QoS parameters, e.g., Cell Loss Ratio, Cell Transfer Delay. However, we notice that for the large group value, the overhead (i.e., stuffing bytes) to complete a SSCS-SDU increases drastically. So they could be used only for the MPEG-2 video with large slice size, such as HDTV.

At the SSCS, a two-byte header and a two-byte trailer information are appended to every SSCS SDU as in Figure 2b. The header is composed of a 4-bit Sequence Number (SN), a 4-bit Sequence Number Protection

(SNP), a 4-bit Payload Type (PT), and a 4-bit Control Block Length (CBL).




Figure 2b - The AV-SSCS Protocol

The trailer is composed of a 2-byte Forward Error Correction field (FEC) applied only to the payload. The FEC scheme uses a Reed-Solomon (RS) code, which enables the correction of up to 4 erroneous bytes in each block of 564 bytes (e.g. 3 x 188). So, it is only used for recovering of cell errors due to electrical or physical problems along the communication path. The addition of a sequence number (SN) of 4 bits enables the receiver entity to detect and locate up to 15 consecutive SSCS PDU losses. Whenever 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 resulting 7-bit codeword 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 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 x M rows into which consecutive fixed length SSCS PDUs are written row by row (see figure 3). The corresponding CS trailer is then appended. A single redundancy row is appended at the tail of the matrix which is obtained by XORing the columns at the cell basis. A single cell loss per block can be recovered or an entire SSCS PDU.




Figure 3 - Virtual Matrix of Control Blocks

The parameter 'M' is referenced as Control Block Length (CBL) which determines the ratio of data and redundancy. It is negotiated at the call set up with

reference to the protection level desired by the connection. The lower its value is and higher the recovery power of the FEC-SSCS mechanism. The drawback is a proportional increase to the control information overhead. Since the FEC information is obtained using XORing method, the data matrix is only an abstract structure and no buffering is required at the sender. The destination checking process is also pipelined and the correct SSCS PDUs are immediately transmitted to the upper layer without latency.

The virtual matrix is read row by row and the trailer is created by calculating the FEC (RS) check bytes first. The Payload Type, the Control Block Length and the Sequence Number are subsequently set. Finally the Sequence Number Protection fields is calculated and appended to the block. Since we are dealing with variable length encoded video slices, it is unlikely to have an exact number of SSCS SDUs to fill up the last virtual matrix of every slice.

Therefore, we propose to indicate, in the SSCS trailer, the length of the Control Block (CBL) that they belong to. This approach allows an easier and more reliable delimitation of the end of the block as well as a better protection of slices from error propagation.

The SSCS-PDU are then transmitted to the common part convergence sub-layer (CPCS) using the CPCS-UNIDATA-Invoke primitive. The 8-byte CPCS trailer, in formation, 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 underlying SAR protocol will subsequently segment the CS-PDU into exactly twelve 48-byte ATM SPDU. The ATM layer will then marked the CLP field of every cell using the 'AUU' and the 'SLP' parameters of the AAL-UNIDATA-Request [17].

### *Receiver behavior*

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.

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. Rather than discarding a corrupted packet, we propose to forward it to the upper SSCS together with an error indication (e.g.Reception Status (RS) parameter of the CPCS UNIDATA signal primitive).

Unfortunately, in the extreme situation of missing entire PDUs, the previous checking mechanisms are not capable to detect the problem. Therefore, the introduction of a sequence number (SN) at the SSCS

layer will permit the detection to up to 15 consecutive packet losses. When packet losses are detected, dummy bytes are inserted in order to preserve the bit count integrity at the receiver.

The association of the reported indication and 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 to the fixed length of both MPEG-2 TS packet and ATM cell, the SSCS layer is capable to easily locate the missing cell.

After localization, both errors and losses can be corrected by respectively using Reed-Solomon and XORed FEC check codes. If no error is detected, the SSCS PDUs are immediately passed to the upper layer after Sequence Numbering check and trailer moving, when only the last SSCS-PDU (redundancy part) is erroneous, no action is performed.

### ***Adaptive Partial video Slice Discard with FEC support***

#### *The SA-PSD Algorithm*

One of the simplest switch buffer-scheduling algorithms is to serve cells in first-in first-out (FIFO) order. Whenever buffer congestion occurs, the incoming cells are dropped regardless to their importance. This random discard (RD) strategy is not suitable for video transmission. A modification is to take into consideration the cell's priority when discarding, i.e., a cell with low priority is dropped first; if congestion persist, this approach gradually begin 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 hierarchy and congestion level (e.g. switch queue length).

In here, we propose enhancement to the Adaptive Partial Slice Discard (A-PSD) to support Forward Error Correction feature. The new scheme, named Selective and Adaptive Partial Slice Discard (SA-PSD), is performed at both control group and video slice levels. Our approach is to reduce the number of corrupted slices by assuming that a number 'T' of cells per control block can be recovered by the destination SSCS using FEC techniques. Let us define the parameter 'T' as the drop tolerance (DT) which corresponds to the maximum number of cells per control block that may be discarded by SA-PSD before considering the control block as lost.

Therefore, unlikely the simple A-PSD, SA-PSD stops discarding cells when the congestion decreases, and the number of previously dropped cells in every Control Block is below the drop tolerance 'T'. 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 our 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 three mechanisms (e.g. PES encapsulation, Dex-PAS, and SA-PSD) with the enhanced AAL-5 provides us 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 adjusting dynamically SA-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. During light congestion, we propose to drop a lower priority slice first rather than delaying it. Then we could assign the buffer space of the dropped slice to a higher priority slice. The proposed approach avoids congestion increase while maintaining the mean cell transfer delay in acceptable value. This proactive strategy is performed gradually by including high priority cells if necessary. As evaluated in [19], the proposed approach can significantly improve the network performance by minimizing the transmission of non-useful video data before buffer overflow. The proposed Selective and Adaptive Partial Slice Discard algorithm is highlighted below.

### *SA-PSD Parameters*

SA-PSD scheme runs per-VC and employs four state variables and one counter variable to control each video connection. Two of them are associated with the slice level and the remaining ones 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 indicate 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.

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.

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 occur when the arrival of a cell is concurrent with the **CB\_DISCARDING** indicator is in "no discarding" state, and **CB\_DROPPED** equals the tolerance 'T'.

The use of both **CB\_DISCARDING** and **CB\_EFCI\_MARKING** indicators allow us to manage losses occurring at subsequent switches and belonging to a control block more efficiently. Indeed, when a block is partially discarded by a switch node, the following switches are not capable to take into account these cell losses to update the associated drop tolerance. As a consequence the switches handle erroneous cell drop tolerance with adverse effect on algorithm performance. At the control block level, the Drop Tolerance can be seen as a loss credit shared by the crossed switches.

To make implementation easy, 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 situation, we propose to commit to the slice level control by entirely dropping the remaining slice.

### *SA-PSD Operation Modes and Fairness*

SA-PSD uses three buffer thresholds as 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 has exhibit better performance.




Figure 4 - SA-PSD operation modes

The thresholds define three operation modes which in turn limit the distribution of the cell loss within the stream to four as in figure 5:




Figure 5 - Cell Loss Distribution using SA-PSD

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, SA-PSD starts to discard incoming cells. The discarding is done in respect to the drop tolerance associated with each connection. We propose a fair distribution of the elimination among the targeted connections using round robin scheduling. If the light congestion is subsisting, the algorithm switches to the slice level, and starts to eliminate any incoming low priority cell until receiving an EOS cell. The procedure is done in a round robin fashion in order to guarantee fairness among connections. The last cells (EOS) are always preserved from elimination since they provide indication of the next slice. Cells with higher priority are always accepted in the buffer. This mode stops when ‘QL’ falls below the “Low Threshold”.

3. *Mode 3*: This mode is activated when QL exceeds “High threshold”. Incoming slices are eligible for discarding regardless to their priority level. The last cell of a control block 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 behaves like Mode 2 to spread the losses over connections with respect to their drop tolerance. It 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 send only when QL drops below Low Threshold (LT). At the reception of feedback signals, the sources immediately change their operation mode. Consequently, some P-frames may transmit cells with different priority.

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

### III. PERFORMANCE EVALUATION

#### Network Simulation Model

The performance evaluation of the proposed Best effort video service is done using simulation experiments.

The network topology is shown in figure 6. It consists of two ATM switches, and ten MPEG2 video connections crossing the bottleneck link with a capacity of 155 Mbps (OC-3). We evaluate the framework in both LAN and WAN configurations, by setting the backbone link to 1 km and 1000 km respectively. All the other link distances, between the source/destination and the switch nodes, are constant and set to 0.2km. The ATM switches are implemented to be non-blocking, output- buffered with finite amount of buffering. Switch buffers size varies from 80,000 to 220,000 cells for both SWITCH-1 and SWITCH-2 in the simulation experiment.




Figure 6 - Network Simulation Model

The video sources generate MPEG2 data at a rate specified in a trace file. The file was obtained from Michael R. Izquierdo, IBM Corporation. A detailed description of this file could be found in [20]. 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.

The video sequences uses SIF format and were encoded at a resolution of 352x240 pixels per frame, a

frame rate of 30 frames/sec, and 15 slices/frame. The sequence is 150 frames (5 sec.) long. In order to run the experiment for a sufficiently long period, we repeated transmission by playing the sequence in a cyclical way. The encoding parameters M and N are set to 3 and 6 respectively, that is, the distance between two I frame is 6 and that between I and P frame is 3 (i.e. the GOP pattern is IBBPBI...). A slice consisted of one macroblock row of 352x16 pixels. Table 3, shows the cells/slice statistics for the video sequence.

File Size (Bytes)	2,819,836
Total Pictures	150
Compression Ratio	6,741
Peak Cell Rate	105
Mean Cell Rate	26.608
Peak Rate (Mbps)	20.034
Mean Rate (Mbps)	5.077
Peak / Mean	3.9462

Table 3 - MPEG2 trace file statistics data

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




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

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 send video streams at the same time. For this reason, we shift the send time so that I and P frames from one sequence would overlap B frames from another source. 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.




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

### Network and Simulation Configuration

Before discussing the experiment results, let us make first some assumptions.

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

As to the transfer delay, we have the following:

- Propagation delay between the sender and the receiver varies from 0.005 ms to 5.0 ms. 0.005 corresponds to the propagation distance of about 1 km, while 5.0 corresponds to 1000 km.
- Queuing delay varies from 0 to a maximum value of 0.6 sec, which corresponds to the maximum buffer size of 220,000 cells, when transmitted using 155 Mbps link (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 x 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.

We carried out our simulation with seven switch buffer configurations. For each of them, the same method is applied to determine the values of the three thresholds. HT, MT and LT are respectively set to 0.9, 0.8 and 0.7 of the maximum queue size (Qmax), where Qmax is

set to one of the following values: 80,000, 100,000, 120,000, 140,000, 160,000, 180,000, 200,000 and 220,000 cells.

Table 4, and 5 summarized the possible states of a cell crossing the network and the investigated performance parameters respectively.

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 4 - Data unit definitions

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.

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 ( $t_{oK}$ ): $D_K = t_{oK} - t_{iK}$

Table 5 - Performance Parameters Definitions

We compare the performance of the proposed framework (DexPAS + SA-PSD + FEC-SSCS) with the three other schemes:

- Random Discarding with no Priority Assignment Scheme (No-RD)
- Selective Cell Discarding with Extend Priority Assignment Scheme (Ex-SCD [5]).

- Adaptive Partial Slice Discarding with Extend Priority Assignment Scheme (Ex-PSD [19]).

### Performance Evaluation at Cell level

Figure 9.1, figure 9.2, figure 9.3 and figure 9.4 show the cell loss ratio for the aggregate, I- P- and B- cell flows respectively. Figure 9.1 indicates that there exists slight difference between the aggregate cell loss curves for the different schemes. The No-RD has the minimum value. The result can be interpreted as follows: In No-RD, a switch accommodates every cell unless the buffer overflows in which case it starts discarding cells blindly.

As a result, it achieves the most usage of buffer and has the least aggregate cell loss ratio. Unlikely, the other three schemes take a preventive strategy to drop cells before the switch buffer is full in order to protect the important, i.e., high priority, information from discarding. Consequently, the switch buffer utilization is not as good as in random discarding.




Figure 9.1 - Cell Loss Ratio (Aggregate video stream)

Ex-SCD has the second least aggregate cell loss ratio because it discards data at the cell-level. In other words, it stops dropping as soon as the switch buffer size decrease to certain low level. Since both Ex-PSD and Dex-SA-PSD stop eliminatin only at the reception of the end of video slice cell (EOS), they experience higher cell loss ratio regardless to the picture types. As never Dex-SA-PSD behaves better than Ex-PSD, because we applied a round robin strategy to distribute the loss fairly among the connections. When the switch get congested, it tries to discard cells, one VC after the other. Thus, cells are preserved better than that in Ex-PSD, where all connections are subject to discarding cells if congestion happens. The cell loss ratio decrease by 13.8%, 20.4%, 24.0% and 25.6% (Ex-SCD, Ex-PSD, No-RD and Dex-PSD respectively) while the buffer size increases from 80,000 to 220,000 cells.

As illustrated in figure 9.2, figure 9.3, and figure 9.4, all of the three preventive schemes (Ex-SCD Ex-PSD and Dex-SA-PSD) concentrate the loss within the

B-frames and protect the reference I- and P- frames. To the extreme, Ex-SCD and Ex-PSD have no I-cell ratio loss at all, Dex-SA-PSD has non-zero I-cell CLR but much better than that of No-RD. The reason is still due to the round robin fashion we applied to ensure fairness, since only one connection is subject to discarding at one time. So, the reaction to protect I-frames is not as promptly as in Ex-SCD and Ex-PSD.




Figure 9.2 - Cell Loss Ratio (I Frame)

As a result, when an I-frame is transmitted running Dex-SA-PSD, the buffer queue length rapidly increases to accommodate that burst. The queue length exceeds the “High Threshold” more often due to the indiscriminately discarding, leading to increased I-frame cell loss compared to the Ex-SCD and Ex-PSD. No-RD has the highest I-cell loss ratio of all schemes, since it drops cells blindly. The data flow in Dex-SA-PSD contains redundancy for forward error recovery. As a result, it consumes more bandwidth, and thus introduces higher cell loss.

Regarding the P-cell CLR, No-RD suffers the highest loss ratio, due to the lack of protection. On another hand, Ex-SCD exhibits the minimum value for similar reasons as in I-cell CLR. In addition, it protects high priority cells, which results in a finer granulate.




Figure 9.3 - Cell Loss Ratio (P Frame)

As demonstrated previously, all predictive schemes, concentrate the loss within the B-frame. However, the

Dex-SA-PSD tries to distribute the loss fairly among VCs, at the expense of slower reaction to congestion. Therefore, it has lower B-cell loss ratio than Ex-SCD and Ex-PSD.

B-frame cells are the most affected by loss. As a result, they largely contribute to the overall loss ratio. P-frame and I-frame cells are the next in order subject for loss. This is due to the drop policy of Ex-SCD, Ex-PSD and Dex-SA-PSD, and the frequency of B-frames in video sequences. Indeed, in our MPEG-2 video samples, the proportion of I, P and B data are 53%, 24% and 23% respectively 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. For example, in figure 16, B-frame occurrences per second, while only 2 and 6 for I and P- frames.




Figure 9.4 - Cell Loss Ratio (B Frame)

From figure 9.5, we observe that the mean cell transfer delay (CTD) increases proportional to the buffer size. As expected, No-RD has the largest mean CTD. An explanation of that is as follows: No-RD accommodates every cell in its switch buffer until it becomes overflow. Thus it increases the queue delay. We also notice that Dex-SA-PSD has longer mean-CTD than the other two schemes even though it tries to drop low priority cells at the light congest stage in order to leave space for the high priority ones. This is mainly due to its overhead, which results to larger switch 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 PSD, which reduces the average queue length. On the other hand, it introduces 15% percent overhead due to stuffing bits and FEC redundancy, which dramatically increases the average queue length. Ex-SCD and Ex-PSD start to drop B-cells when light congestion occurs and thus reduce the buffer occupancy while minimizing the transfer delay of the high priority cells. This can be also shown by the buffer occupancy status, as in figure 9.6, figure 9.7, figure 9.8 and figure 9.9.




Figure 9.5 - Mean Cell Transfer Delay

With No-RD, the buffer must be filled until High Threshold is reached before to start elimination which leads to a greater mean CTD.

The difference of CTDs between different schemes increases with larger buffer size. For instance, with limited buffer size, the difference between random drop 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 are available. With limited buffer size, the space saved by dropping B-cells is limited and therefore it performs similarly to No-RD.




Figure 9.8 - Buffer Occupancy with Ex-PSD




Figure 9.9 - Buffer Occupancy with Dex-SA-PSD




Figure 9.6 - Buffer Occupancy with No-RD




Figure 9.7 - Buffer Occupancy with Ex-SCD

*Performance Evaluation at Slice Level:*

In this section, we compare the efficiency of the proposed picture quality control framework (Dex-SA-PSD-FEC-SSCS) as compared to other techniques with respect to their ability to provide a message-based service.

Video Slice Loss Ratio (SLR) is measured at the application layer. Decoding, e.g. lost and dead cells, and propagation, e.g. late cells, constraints are also taken into account. In addition, the forward error correction capability is considered in the Dex-SA-PSD framework. In order to validate our proposed algorithm, we complete the performance of Dex-SA-PSD with the same three discarding policies as in the previous section.

Intuitively it is expected that Dex-SA-PSD has better performance at slice level. This is exhibited by figure 9.10, figure 9.11, figure 9.12, and figure 9.13. The proposed framework significantly improves the percentage of arrivals at the destination of non-corrupted 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-PD reach 16.6%, 12.2% and 8.9% respectively.




Figure 9.10 - Slice Loss Ratio (Aggregate Stream)




Figure 9.13 - Slice Loss Ratio (B Frame)




Figure 9.11 - Slice Loss Ratio (I Frame)

Finally, the SLR per sub-flow is analyzed for the four approaches as follows. We observe that, Ex-PSD and Ex-SCD outperforms the other approaches by protecting I-frames, though for aggregate SLR, our new scheme has the best performance. This is consistent with the results obtained at cell level. As discussed in the previous section, there is a trade-off between fair distribution of cell discarding among the VCs and the speed of reactions to congestion. P-frame and B-frame, Dex-SA-PSD demonstrates the best SLR value. This further indicates the capability provided to protect data at the slice level by the FEC mechanism.




Figure 9.12 - Slice Loss Ratio (P Frame)

### Distance Effect

As mentioned in the previous section, the novel Priority Assignment Scheme (Dynamic Extend PAS) uses the classical CLP bit, and dynamically assigns cell priorities accordingly to the current MPEG-2 frame type (I, P or B) as well as network state. Therefore, it depends on the reception of backward congestion signal from the network.

The main drawback of this scheme is that its efficiency is stringently dependent to the round trip time delay. To validate this, we compare the two set of results for the proposed scheme, one with backbone link distance of one kilometers, simulating an ATM LAN, and the other of one thousand kilometers, simulating an ATM WAN.

Figure 9.14, figure 9.15, figure 9.16, and figure 9.17 show the results, in the WAN situation. The aggregated cell loss ratio shows only slight difference to that of LAN.

However, the I-frame protection becomes worse. This is expected since the PAS scheme assign the priority of P-frame to low or high dynamically depending on the feedback of the networks. In LAN configuration, this can be achieved promptly to reflect the current congestion state and thus, P-frame could be dropped to alleviate the congestion as soon as queue length (QL) exceed MT.




Figure 9.14 - Cell Loss Ratio (Aggregate) vs link distance

In WAN, due to the longer round trip time (RTT), the feedback comes slower and thus we expect that the priority assignation level could not be adjusted as fast and thus the protect of I-frame when congestion get worse is not as good. Since the cost of I-frame protection is to sacrifice some P-frame cells, the aggregate CLR remains almost the same. The same is verified at the slice level, as shown in figure 9.18, figure 9.19, figure 9.20 and figure 9.21.




Figure 9.15 - Cell Loss Ratio (I Frame) vs link distance




Figure 9.16 - Cell Loss Ratio (P Frame) vs link distance




Figure 9.17 - Cell Loss Ratio (B Frame) vs link distance

### Redundancy vs Data ratio

The forward error correction capability of the transport service is based on the redundancy data ratio.

In our algorithm, the redundancy is determined by the M value (the number of rows in one Control Block). The smaller the M value, the larger amount of FEC coding data, and thus the stronger error correction capability. How many redundancy data per FEC frame is needed in order to obtain sufficient performance mainly depends on the cell loss pattern. Larger number of redundancy is required for a strongly corrected cell loss. As mentioned before, the proposed FEC scheme can be easily optimized with respect to amount of redundancy data per FEC frame, and the actual frame size even during a session. As a result, this FEC scheme will be able to achieve sufficient throughput and latency performance with reasonable transmission overhead.

Figure 9.22, figure 9.23, figure 9.24, figure 9.25, figure 9.26, figure 9.27, figure 9.28 and figure 9.29 show the result of CLR and SLR respectively when different values of M are used.




Figure 9.22 - Cell Loss Ratio vs Data Redundancy Ratio (Aggregate Stream)

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

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




Figure 9.23 - Cell Loss Ratio vs Data Redundancy Ratio (I Frame)




Figure 9.24 - Cell Loss Ratio vs Data Redundancy Ratio (P Frame)




Figure 9.25 - Cell Loss Ratio vs Data Redundancy Ratio (B Frame)

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

## V. CONCLUSION

In this paper, we have surveyed complementary issues related to the transmission and control of MPEG-2

video streams over ATM best effort services. We have proposed, evaluated, and enhanced a best-effort service based on UBR+, which takes into account the specific encoding and stochastic properties of MPEG2 video sources. This service is composed of three components:

- A new priority data partition technique which extends ATM prioritization capability. With a better usage of the ATM cell header, the Dynamic Extended Priority Assignment Scheme (Dex-PAS) provides three different service classes per-connection and the detection of video slice boundaries at the cell level. In association with an intelligent video-oriented discard scheme, Dex-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.
- A slice-based service specific convergence sublayer (i.e. FEC-SSCS) which enhanced the classical ATM Adaption Layer type 5. Among the additional features supported by the new extended AAL5 are the ability to distinguish video frame types, as well as the detection and the forward correction of loss and errors at both byte and cell basis.
- Finally, we designed a novel cell discard scheme, referred to SA-PSD, which intelligently operates according to the FEC ability of the source and destination.

Compared to classical UBR, our performance results shown that the proposed best effort transport service can improve the QoS provided to video encoded applications. The integrated service can concentrate the data loss within the B-frame and thus better protect critical video data (e.g. I- and P-frames). Despite of an increase of cell losses, the proposed service demonstrated an improved result at the slice level – a significant increase of the number of non-corrupted slices arriving at destination. A slight reduction of the mean cell transfer delay for the aggregate video stream is experienced because of the overhead introduced by the Forward Error Correction mechanism. In a WAN configuration, the proposed framework endures slight performance degradation in protecting critical information, since it depends on the reception of feedback from the network. The efficiency depends on the round trip time delay. Finally, the proposed FEC scheme can be easily optimized in terms of both the number of redundancy data per FEC frame and the actual frame size. Consequently, it can achieve sufficient performance with respect to the throughput and latency if the redundancy data ratio is carefully selected.

The major drawback of our FEC-SSCS protocol is that it introduces redundant information. Further work need to be done to minimize this kind of overhead. As an example, the stuffing bits could be avoided (i.e., not transmitted) when both sender and receiver entities use the same stuffing pattern. However, this leads to SSCS-PDUs of variable length and, hence, requires an additional field in the PDU header (or tail) to indicate the actual PDU length. This has been actually implemented in our simulator and we did observe a slightly improvement in terms of performance.

## Acknowledgments

We are grateful to Mickael Izquierdo from IBM Corp., Triangle Research Park, North Carolina for providing us with the MPEG2 statistic files.

## REFERENCES

- [1] S. Keshav M. Grossglauser and D. Tse. Rrcbr: a simple and effective service for multiple time-scale traffic. Proceedings of ACM SIGCOMM'95, 25:219--230, 1995.
- [2] H. Kanakia, P. Mishra, and A. Reibman. Packet video transport in ATM networks with single-bit feedback. Proceedings of the sixth International Workshop on Packet Video, Portland, Oregon, September 1994.
- [3] T.V.Lakshman, P.P. Mishra, and K.K.Ramakrishnan. Transporting compressed video over ATM networks with explicit rate feedback control. IEEE INFOCOM 97, Kobe, Japan, March 1997.
- [4] H. Zhang and E.W.Knightly. Red-vbr: A renegotiation-based approach to support delay-sensitive vbr video. to appear in ACM/Sringer-Verlag Multimedia Systems Journal, 1996.
- [5] A. Mehaoua, R. Boutaba, and G. Pujolle. An extended priority data partition scheme for mpeg video connections over atm. IEEE Symposium on Computers and Communications'97, Alexandria, Egypt, July 1997.
- [6] ITU-T. Itu-t i.361 specification of atm layer. March 1993.
- [7] A.Mehaoua, R.Boutaba, and G. Pujolle. Proposal of an extended aal5 for vbr mpeg-2 over atm. IEEE International Conference on Multimedia Computing and Systems'97, Ottawa, Canada, June 1997.
- [8] T. Han and L. Orozco-Barbosa. Performance requirements for the transport of mpeg video streams over atm networks. ICC'95, Washington, pages 221--225, June 1995.
- [9] A.Mehaoua and R.Boutaba. Performance evaluation of slice-based discarding schemes for best effort video networking over atm. the 22th IEEE Annual Conference on Computer Networks, Minneapolis, Minnesota, November 1997.
- [10] ATM Forum. Draft proposal for specification of fec-sscs for aal type 5. af-saa-0326R2, October 1995.
- [11] AK.Kanai and R.Grueter et al., Forward error correction control on aal5: Fec-sscs, ICC'96, Dallas, TX, pages 384--391, June 1996.
- [12] ATM Forum SAA SWG. Audiovisual multimedia services: Video on demand. af-saa-0049.000, December 1995.
- [13] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. SIAM Journal of Applied Mathematics, 8:300--304, 1960.
- [14] S.B. Wicker and V.K.Bhargava. Reed-solomom codes and their applications. IEEE press, 1994.
- [15] N. Shacham et P. McKenny, 'Packet Recovery in High Speed Networks using coding', IEEE INFOCOM '90, San Francisco, CA, Juin 1990, p.124-131.
- [16] A. Mehaoua and R. Boutaba, 'Performance Analysis of Cell Discarding Techniques for Best Effort Video Communications over ATM Networks', in Computer Networks and ISDN Systems Journal, Vol.29, issue 17-18, Elsevier Science, North Holland, February 1998, pp. 2021-2037.
- [17] ITU-T, B-isdn adaption layer (AAL1), November 1995, Perth
- [18] ISO/IEC International standard 13818-2, "Generic Coding of Moving Pictures and Associated audio information: Part2 - Video", 1995.
- [19] A.Mehaoua and R.Boutaba. Performance analysis of a partial video slice level discard scheme for mpeg connections over ubr+ service. Proceedings of IFIP ICC'97, Cannes, France, November 1997.
- [20] M.R. Izquierdo and D.S. Reeves. Mpeg vbr slice layer model using linear predictive coding and generalized periodic markov chains. International Performance, Computing and Communications Conference, Scottsdale, Arizona, February 1997.




Figure 9.18 - Slice Loss Ratio (Aggregate) vs link distance




Figure 9.26 - Slice Loss Ratio vs Data Redundancy Ratio (Aggregate)




Figure 9.19 - Slice Loss Ratio (I Frame) vs link distance




Figure 9.27 - Slice Loss Ratio vs Data Redundancy Ratio (I Frame)




Figure 9.20 - Slice Loss Ratio (P Frame) vs link distance




Figure 9.28 - Slice Loss Ratio vs Data Redundancy Ratio (P Frame)




Figure 9.21 - Slice Loss Ratio (B Frame) vs link distance




Figure 9.29 - Slice Loss Ratio vs Data Redundancy Ratio (B Frame)