Proposal of An AudioVisual SSCS with Forward Error Correction

Ahmed Mehaoua [1,3], Raouf Boutaba [2], Jean-Pierre Claudé [3], Guy Pujolle [3]

[1] *University of Cambridge - CCSR*
10 Downing Street
CB2 3DS Cambridge, UK
a.mehaoua@ccsr.cam.ac.uk

[2] *University of Toronto - ECE Dept.*
10 King's College Road
Toronto (On.), Canada
rboutaba@comm.utoronto.ca

[3] *University of Versailles – PRiSM Lab.*
45 av. des Etats-Unis
78035 Versailles, France
{mea, jpc, gp}@prism.uvsq.fr

ABSTRACT

This paper addresses the transport of real-time multimedia traffic generated by MPEG-2 applications over ATM best effort services (ABR, UBR+). To cope with network congestion and the unreliability of the ATM Adaptation Layer type 5 (AAL5), we propose a new Service Specific Convergence Sublayer for AudioVisual applications. The proposed AV-SSCS includes adaptive Forward Data Error Correction (FEC) based on Reed-Solomon and Parity Codes. With respect to network load and video packet loss ratio measured at the destination, the source can dynamically increase or reduce the FEC efficiency (i.e. the amount of data redundancy) to provide a higher protection to the video packet or a better use of the shared bandwidth respectively.

Keywords: AAL, Packet Video, MPEG, Cell Discard, FEC

I. INTRODUCTION

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 ATM Adaptation Laye (AAL) and the associated end-to-end error detection and correction algorithms.

This AAL is responsible for making the networ 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 computer data traffic over ATM networks, which makes too simple to provide reliable connection fo multimedia applications. Additional features such as error localization and forward error recovery are required to improve AAL5 reliability for the transport of real-time MPEG-2 video data.

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 ma 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. randoml 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, cell 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 proposed AV-SSCS in this paper is associated with an intelligent switch discarding technique called Partial Video slice Discard scheme with FEC-aware (i.e. FEC-PSD).

The paper is organized as follows. Section 2 is devoted to the description of the ne AudioVisual SSCS. In section 3, we introduce the cooperative intelligent packet video discard scheme and we finally conclude in section 4.

2. AUDIOVISUAL SSCS FOR BEST EFFORT VIDEO OVER ATM (FEC-SSCS)

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 [1][2] 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 [3], 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 [4] and Parity Codes [5], and on the video packet encapsulation mechanism proposed in [6]. 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 i described bellow.

2.1 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 [7], as illustrated in figure 1.

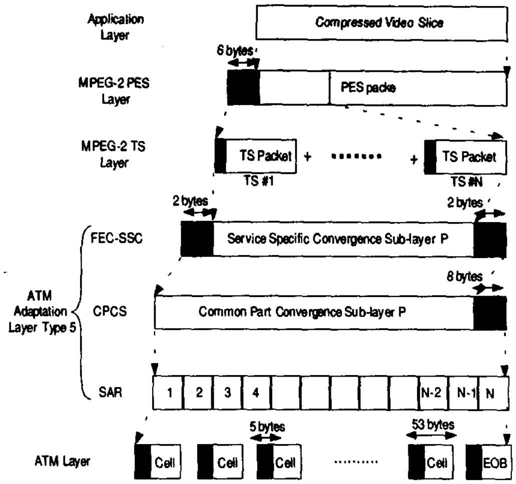

Figure 1 - the Audio-Visual AAL5 with FEC

The following primitive is used: AAL UNITDATA request(ID, M, SLP, CI). The 'Interface Data' (ID) parameter specifies the exchange of MPEG-2 T 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 [8], 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. 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.

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 1.

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 1 - 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 2. The header is composed of a 4-bit Sequence Number (SN), a 4-bit Sequence Numbe Protection (SNP), a 4-bit Payload Type (PT), and a 4-bit Control Block Length (CBL).

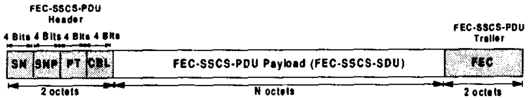

Figure 2 - The FEC-SSCS Protocol

The trailer is composed of a 2-byte Forward Error Correction field (FEC) applied only to the payload. Th 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. Th 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 contain 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 i 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 b XORing the columns at the cell basis. A single cell loss per block can be recovered or an entire SSCS PDU.

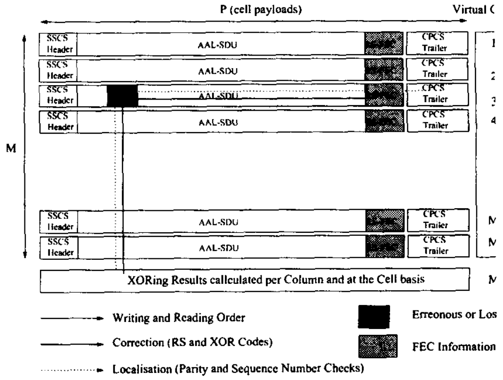

Figure 3 - Virtual Matrix of Control Blocks

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 immediatel 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) laye using the SAR-UNIDATA-Invoke primitive. Th underlaying 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 [7].

2.2 Receiver behavior

At destination, three tasks have to be performed b 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 finall (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 CPC 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, dumm 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, th SSCS PDUs are immediatelv 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.

3. A FEC-AWARE ADAPTIVE PARTIAL VIDEO SLICE DISCARD SCHEME (FEC-PSD)

3.1 The FEC- PSD Algorith

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 thei 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 [9], 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 FEC-aware Partial Slice Discard (FEC-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 FEC-PSD before considering the control block as lost.

Therefore, unlikely the simple A-PSD, FEC-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. PE encapsulation, Dex-PAS, and FEC-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 i 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 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. 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 priorit slice. The proposed approach avoids congestion increase while maintaining the mean cell transfer dela in acceptable value. This proactive strategy is performed gradually by including high priority cells if necessary. As evaluated in [9], the proposed approach can significantly improve the network performance b minimizing the transmission of non-useful video data before buffer overflow. The proposed Selective and Adaptive Partial Slice Discard algorithm is highlighted bellow.

3.2 FEC- PSD Parameters

FEC-PSD scheme runs per-VC and employs fou state variables and one counter variable to control eac 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 th 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 fo 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.

3.3 FEC- PSD Operation Modes and Fairness

FEC-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.

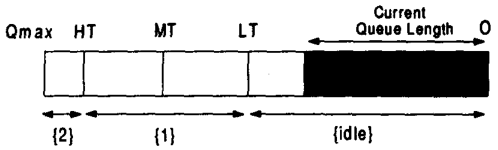

Figure 4 - FEC-PSD operation modes

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:

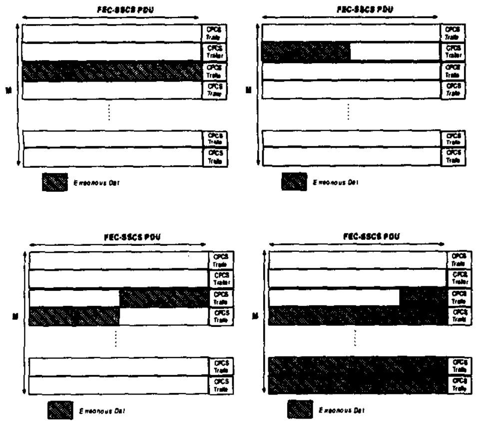

Figure 5 - Cell Loss Distribution using FEC-PSD

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 buffe 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. 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 an 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 the provide indication of the next slice. Cells with higher priority are always accepted in the buffer. This mode stops when 'QL' falls bellow the "Low Threshold".

3. *Mode 3*: This mode is activated when QL exceeds "High threshold". Incoming slices are eligible fo discarding regardless to their priority level. The last cell of a control block and slice are preserved to avoid th 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 fall 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 quickl 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 b gradually including I-frame cells, if necessary.

4. CONCLUSION

In this paper, we have surveyed complementar issues related to the transmission and control of MPEG-2 video streams over ATM best effort services. We have proposed a video slice-based service specific convergence sublayer with FEC (i.e. FEC-SSCS) which enhanced the classical ATM Adaption Layer type 5. Among the additional features 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 have also designed a novel cell discard scheme, referred to FEC-PSD, which intelligently operates according to the FEC ability of the source and destination.

Compared to classical AAL5, the proposed FEC-SSCS scheme can be easily optimized according to network congestion 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 latenc if the redundancy data ratio is carefully selected. The major drawback of our FEC-SSCS protocol is that it introduces redundant information.

REFERENCES

[1] ATM Forum. Draft proposal for specification of fec-sscs for aal type 5. af-saa-0326R2, October 1995.

[2] AK.Kanai and R.Grueter et al., Forward error correction control on aal5 Fec-sscs, ICC'96, Dallas, TX, pages 384--391, June 1996.

[3] ATM Forum SAA SWG. Audiovisual multimedia services: Video on demand. af-saa-0049.000, December 1995.

[4] S.B. Wicker and V.K.Bhargava. Reed-solomom codes and their applications. IEEE press, 1994.

[5] N. Shacham et P. McKenny, 'Packet Recovery in High Speed Networks using coding', IEEE INFOCOM '90, San Francisco, CA, Juin 1990, p.124-131.

[6] 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, Februar 1998, pp. 2021-2037.

[7] ITU-T, B-isdn adapation layer (AAL-l), November 1995, Perth

[8] ISO/IEC International standard 13818-2, "Generic Coding of Moving Pictures and Associated audio information: Part2 - Video", 1995.

[9] A.Mehaoua and R.Boutaba. Performance analysis of a partial video slice level discard scheme for mpeg connections over ubr+ service. Proceedings of IFIP ICCC'97, Cannes, France, November 1997.