New Network and ATM Adaptation Layers for Real-Time Multimedia Applications: A Performance Study Based on Psychophysics

Xavier Garcia Adanez, Olivier Verscheure and Jean-Pierre Hubaux
Telecommunications Services Group, TCOM Laboratory
Swiss Federal Institute of Technology, Lausanne
email: {garcia, verscheure, hubaux}@tcom.epfl.ch
URL: http://tcomwww.epfl.ch/{˜garcia, ˜verscheu, ˜hubaux}

Abstract

We present in this paper Network and ATM Adaptation Layers for real-time multimedia applications. These layers provide a robust transmission by applying per-cell sequence numbering combined with a selective Forward Error Correction (FEC) mechanism based on Burst Erasure codes. We compare their performance against a transmission over AAL5 by simulating the transport of an MPEG-2 sequence over an ATM network. Performance is measured in terms of Cell Loss Ratio (CLR) and user perceived quality. The proposed layers achieve an improvement on the cell loss figures obtained for AAL5 of about one order of magnitude under the same traffic conditions. To evaluate the impact of cell losses at the application level, we apply a perceptual quality measure to the decoded MPEG-2 sequences. From a perceptual point of view, the proposed AAL achieves a graceful quality degradation compared to AAL5 which shows a critical CLR value beyond which quality drops very fast. The application of a selective FEC achieves an even smoother image quality degradation with a small overhead.

1 Introduction

ATM technology is reaching a certain level of maturity that allows for its deployment in local as well as in wide area networks. Concurrently, audiovisual applications are foreseen as one of the major users of such broadband networks. However, it has already been shown in [1] that the cell and frame loss ratios might not be negligible in ATM based environments especially if the operators employ statistical multiplexing to efficiently use network resources. MPEG-2, the standard for full motion video, will be the audio and video compression tool of such multimedia applications. The semantic and syntactic structure of the data flows generated by MPEG-2 makes the applications sensitive to data loss. Errors can spread across a single image causing holes or may spread out over several images within a Group of Pictures (GOP). The adoption of AAL5 by the ATM Forum [2] for the transmission of MPEG-2 video even increases the impact of data loss due to its packet level granularity which implies packet discard in

1

case of cell loss. To reduce the impact of cell losses on video applications, using a cell level granularity reduces the CLR seen by the application compared to AAL5 based transmission [3].

Traditional error recovery methods are based on retransmission, also known as Backward Error Correction (BEC). While this technique has proven to be very efficient for data applications, it fails to protect real-time multimedia applications due to their stringent timing constraints, especially in Wide Area Networks (WAN) and point-to-multipoint configurations. Conversely, Forward Error Correction (FEC) techniques do not rely on data retransmission but instead, by adding redundancy, are able to correct the errors at the receivers end. The major drawback of FEC techniques is that they add delay and overhead. We propose in this paper a Burst Erasure [4] selective FEC technique based on the analysis of the syntactic and semantic components of the stream to be transmitted. We efficiently protect the most sensitive elements of the data by adding a small overhead since the header to raw data ratio is small. ATM Adaptation Layers (AAL) by definition are generic in the sense that they have to support any kind of applications, albeit not any ATM Transfer Capability (ATC). To get rid of any application specificity we have developed a Network Adaptation specific to MPEG-2 on top of the AAL. This layer is able to identify the headers encapsulated into the Transport Stream (TS) packets [5]. We show by simulation that the proposed AAL gives better results in terms of cell loss. As it is very difficult to map the impact of cell losses onto MPEG-2 based video applications, we use a perceptual quality metric based on psychophysics to bring to the fore the improvements obtained by our proposal from the end user perspective.

The paper is organized as follows: Section 2 describes the requirements of today's multimedia applications and the currently available mechanisms. In the next section we develop the mechanisms that we propose for a multimedia AAL and we describe the operation of the FEC in combination with a Network Adaptation. Section 4 describes the simulation setup and shows the improvements measured in terms of network and perceptual quality parameters. We derive some conclusions in Sec. 5 and we provide an outlook of how this work will be continued.

2 ATM Adaptation Layer Mechanisms for Multimedia Applications

2.1 Requirements of Multimedia Applications

The principal characteristic of multimedia data streams is that it is of a continuous nature. As such, it has very stringent constraints in terms of delay and delay jitter. This has been referred to as Timely Information [6]. This means that data arriving beyond a certain point in time is considered as lost by the application. Moreover, if the audiovisual information makes use of compression techniques such as MPEG-2, it becomes also sensitive to loss due to the lack of redundancy in the transmitted data albeit the impact of loss on the image heavily depends on the type and the location of the lost information. Data

Figure 1: *Data Loss Propagation*

loss spreads within a single picture up to the next resynchronization point (e.g. slice headers). This is referred to as spatial propagation. Due to the predictive nature of the MPEG-2 algorithm, when losses occur in a reference picture (Intra-coded, I, or predictive, P, frames) it will remain until the next intra-coded picture is received. This causes the errors to propagate across several pictures which is known as temporal propagation (see Fig 1). The impact that the loss of syntactic data may have is in general more important and difficult to recover than the loss of semantic information. So, the transport of real-time multimedia data has to be reliable and timely. ATM networks fulfill both conditions but, in some cases, they may fail to guarantee loss ratios.

2.2 Current AAL Mechanisms

Nowadays, two AALs are available for the transport of multimedia data: AAL1 and AAL5 [7]. AAL1 was basically designed to cover circuit emulation services. It therefore offers Constant Bit Rate (CBR) services to the applications. This may be a limitation if constant quality encoding or Variable Bit Rate (VBR) in general is used. The advantages of AAL1 are that it provides a cell level granularity to detect cell losses via a sequence number and also provides cell loss recovery via FEC combined with interleaving. However, using interleaving introduces delays proportional to the size of a FEC block at both the sender and the receiver. The FEC scheme applied is based on Reed-Solomon codes that are able to correct erasures (cell losses) and also random errors (impulse noise). However, the bit error ratios to be considered in fiber-based ATM networks are close to 10^{-12}. In fact, with such impulse noise error values, it becomes unnecessary to have bit or octet granularity. Therefore, and taking into account that it introduces considerable delay and overhead, an octet based interleaver is not a major need.

On the other side, AAL5 is able to cope with any ATM transfer capability. One of the reasons for this is its simplicity and low overhead. It receives

PDUs from the upper layer, appends an 8 byte trailer containing a CRC-32 parity check, a length indicator and padding information for boundary alignment, and sends the AAL5-PDU to the Segmentation and Reassembly (SAR) sublayer. This simplicity has however some drawbacks for multimedia applications. AAL5 does not provide enough protection against erasures because it was mainly designed for data transfer applications that rely on robust transport protocols for error correction. It provides an error detection mechanism based on the parity check calculated on a PDU basis and the check of the length of the received packet. Due to the lack of more sophisticated error detection functionalities, in case of cell losses, i.e. length indicator mismatch, AAL5 is unable to know the position of the cells lost inside the PDU and so, no error correction even at a higher layer can be applied. Therefore, when cell losses are detected, the packets are discarded. If no retransmission mechanism is applied, the packet discard leads to a data loss at the application level which is higher than the data loss (cell loss) at the network level. This is clearly not adequate for real-time applications for two main reasons. First, a single cell loss causes the loss of several data that could be still used by the decoder with techniques such as early resynchronization[8]. Second, AAL5 does not notify errors to the upper layers. Therefore, the decoder may not detect data corruption and may not be able to apply any of the multiple error concealment mechanisms [9].
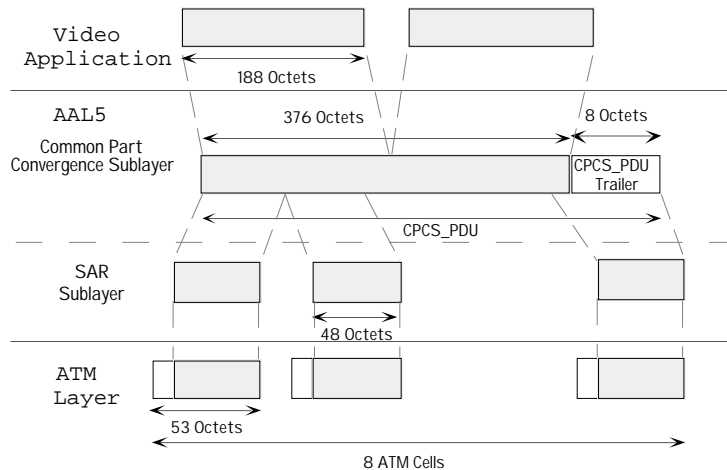

Figure 2: *AAL5 Segmentation Mechanism of MPEG-2 Video*

3 New Mechanisms for Real-Time Multimedia Applications

3.1 ATM Adaptation Layer Mechanisms

To meet the requirements of real-time multimedia applications, we propose an AAL that includes the following mechanisms: cell level granularity to improve the error detection capability and a selective FEC mechanism to selectively protect essential data. The cell level granularity needs the introduction of a per-cell
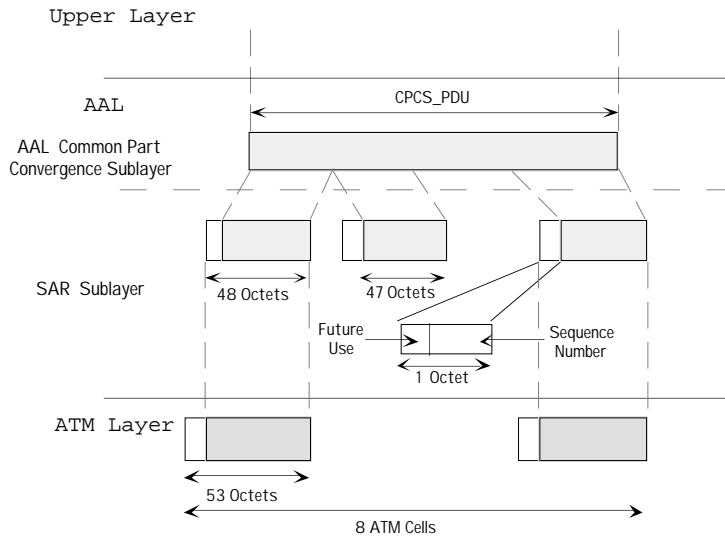
Figure 3: *Proposed AAL Segmentation Mechanism*

sequence number. To make this possible, we use 47-byte cell payloads which frees an octet per ATM cell to insert the sequence number and also control information for the FEC mechanism (see Fig 3). If we consider the transport of MPEG-2 TS packets, our packetization has the advantage of always giving an integer number of cells, since the length of the TS packets is 188 bytes which gives exactly 4×47 byte payloads. This increases the per cell overhead. However, we have reduced the PDU overhead since we do not add any header or trailer information as AAL5 does. Since our packetization process does not provide such kind of PDU delineation, our approach consists in assuming that the PDU size is negotiated at connection setup and that it will remain constant for the duration of the connection. Therefore, the receiver does not need any extra information to delineate the packets. Moreover, this scheme reduces the number of states that the protocol has to deal with in case of loss. This scheme allows to reduce the data loss seen by the receiver because it increases the resolution of the data loss detection algorithm avoiding the packet discard. The receiver uses the sequence numbers to detect the number and position of the lost cell in a packet. This information can thus be passed to the upper layers that can take necessary action to conceal data loss. In our proposal, we use a dummy cell insertion mechanism, when cell losses are detected (see Fig 4). This has two advantages: first, it guarantees packet length integrity. Second, in the case of MPEG-2, it can be used as an error message since sequences of 47 zero-octets are not allowed by the MPEG-2 standard [10]. So this mechanism can be exploited by the decoder as an error indicator. Also, if FEC is used, the insertion of dummy cells simplifies the error correction. However, since the proposed mechanisms are not specific to MPEG-2, we also consider the fact that dummy cell insertion may be useless or even harmful to other applications. Consequently, we propose this mechanism to be user selectable at connection setup.
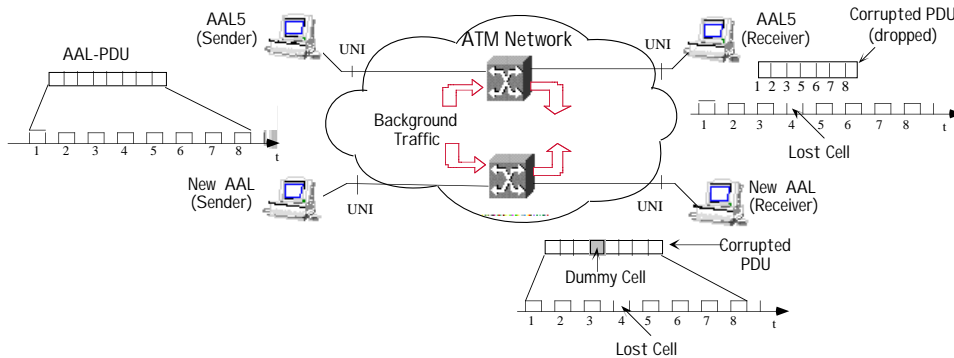
Figure 4: *Dummy Cell Insertion Mechanism*

The utilization of FEC for real-time multimedia applications has several advantages. Besides the fact that it introduces relatively low delay compared to retransmission based error correction mechanisms, it is also better suited to point-to-multipoint configurations that should be widely used by such multimedia applications. The major drawback of FEC is that it introduces overhead. Indeed, it adds redundancy packets that will be able to recover the missing data when losses are detected; moreover, FEC does not guarantee zero loss. To reduce the overhead generated by the FEC scheme, we propose to include a selective mechanism in our AAL. The advantage of such a selective method is twofold: first it reduces the overhead, and second, it adapts to highly structured information such as compressed video. Indeed, the impact of syntactic data loss (i.e. headers) is rather different than the loss of semantic data. Headers are points of synchronization. When a header is lost, all the underlying data cannot be recovered. This leads to a loss of quality and also to a waste of network resources since data correctly received could not be decoded which may either generate artifacts or loss of synchronization. Moreover, such selective mechanism can be used to protect separately audio from video since losses in audio are much more noticeable and disturbing to the user than video losses. Besides, also, timing information can be selectively protected to reduce the probability of losing synchronization.

The flexibility proposed by our mechanism cannot be achieved by using the method already developed for AAL1. The latter uses Reed-Solomon Codes with interleaving which gives a fixed matrix structure. Due to this structure it is difficult to apply a selective mechanism. We propose a mechanism based on burst erasure codes (RSE) [4]. The advantage of this method is that it takes into account the specifics of ATM. It relies on the fact that erasures are limited to fixed boundaries (cells) to correct cell losses only. Conversely, it cannot correct octet or bit errors (impulse noise). This is not a fundamental problem since we assume that the ratio of errors to erasures is very small and can be considered as negligible.

6

3.2 Network Adaptation Mechanisms

By definition, an AAL has to be generic in the sense that it does not have to provide mechanisms specific to applications. The mechanisms proposed in Sec. 3.1 albeit designed with real-time applications in mind are not specific to any audio-visual coding scheme. To provide efficient protection of data with selective FEC mechanisms, it is necessary to know the syntax of the data to be transmitted. The protocol stack depicted in Fig. 5 presents the network adaptation as the layer specific to an application. In our example, we have used MPEG-2 based applications so we have developed for our experiments a network adaptation layer specific to this standard. The network adaptation main functionalities are: delivery of fixed size PDUs (in the case of MPEG-2 it is straightforward) which includes packet segmentation/reassembly and alignment to boundaries (padding), and detection of loss sensitive data combined with the generation of FEC_request messages. The second functionality in the case of MPEG-2 applications consists of detecting headers and generating FEC_request messages accordingly. Based on these messages passed with the PDU, the AAL will or will not generate the FEC data.


Figure 5: *Protocol Stack*

The proposed Network Adaptation depicted in Fig. 6 receives the TS packets from the MPEG-2 system layer and generates a single NA-PDU. Since we assume a fixed packet size, no PDU delineation is done and therefore no overhead is generated. However, applications using other encoding systems such as JPEG or MJPEG that generate variable length packets will need to add some information to align the packet sizes to PDU boundaries.

4 Comparison of AAL5 and the Proposed AAL

4.1 Simulation Setup

The simulation setup used for our experiments is depicted in Fig. 7. The simulator is composed of four multimedia workstations and two ATM switches. Both switching stages, implemented as multiplexers with limited buffer size, are loaded with background traffic provided by several On-Off sources. This type of source model is widely used to simulate a multiplex of traffic such as the one that could be found at the entrance of an ATM switch. Moreover, two state Markov source models encompass the peak cell rate parameter which is
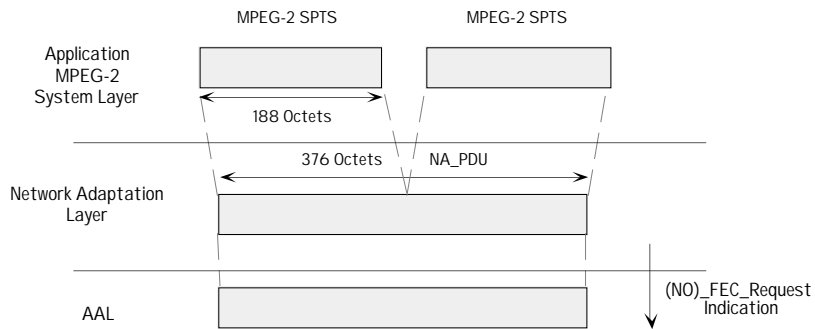
Figure 6: *Proposed Network Adaptation for MPEG-2 based applications*

currently the most important traffic contract parameter [11]. To guarantee the same CLRs to both cell streams, the background traffic is replicated and sent simultaneously to both multiplexing stages. The multimedia workstations are connected as two point-to-point communications. One of the connections uses AAL5 to transmit the MPEG-2 bit stream while the other uses the proposed AAL.

The traffic under test (TUT) consists of a ski sequence of 1000 frames (720×576) encoded with MPEG-2 at 4 Mbit/s with a structure of 12 images per GOP with two B pictures between every reference picture and a single slice per line. The encoded bitstream is encapsulated into TS packets prior to be sent to the network. The resulting traffic is sent to the switch at a constant cell rate where it is multiplexed with the background traffic. Since the switch buffer is limited in size, some of the TUT or background cells may be lost. The TUT is then routed to the receiver end system where the data is reassembled prior to decoding. The background traffic is assumed not to interfere further with the TUT and thus is directly routed to a traffic sink after leaving the switching stage.

The transmission of video over AAL5 is based on the approved ATM Forum Video on Demand specification [2]. This document describes the encapsulation of MPEG-2 TS packets into AAL5-SDUs (Service Data Units). This scheme packetizes two single program transport streams (SPTS) packets regardless of


Figure 7: *Simulation Scenario*

8

their information contents, being of audio, video or timing nature into a single AAL5-SDU. The AAL5 adds its 8 byte trailer and the resulting AAL5-PDU is segmented into 8 ATM cells without any padding (see Fig. 2). This encapsulation method may support other PDU sizes, 376 being the default. However, some padding may be necessary to align larger AAL5-PDUs to ATM cell boundaries.

The transmission of video over the proposed AAL is based on the segmentation described in Sec. 3.1. Two TS packets are passed from the Network Adaptation layer to the AAL. The AAL-PDU is then segmented into 47-byte payloads giving exactly 8 ATM cells. When FEC is applied, we add a single redundancy cell obtained by XORing the 8 data cells. A single cell loss per PDU can be recovered.

4.2 A Perceptual Quality Measure as a Performance Criterion

Several studies have shown that a correct estimation of subjective quality has to incorporate some modeling of the Human Visual System [12]. A spatio-temporal model of human vision has been developed for the assessment of video coding quality [13, 14, 15]. The model is based on the following properties of human vision:

- The responses of the neurons in the primary visual cortex are band limited. The human visual system has a collection of mechanisms or detectors (termed channels) that mediate perception. A channel is characterized by a localization in spatial frequency, spatial orientation and temporal frequency. The responses of the channels are simulated by a three-dimensional filter bank.

- In a first approximation, the channels can be considered to be independent. Perception can thus be predicted channel by channel without interaction.

- Human sensitivity to contrast is a function of frequency and orientation. The *contrast sensitivity function* (CSF), quantizes this phenomenon, by specifying the detection threshold for a stimulus as a function of frequency.

- Visual masking accounts for inter-stimuli interferences. The presence of a background stimulus modifies the perception of a foreground stimulus : masking corresponds to a modification of the detection threshold of the foreground according to the local contrast of the background.

The vision model described in [13] has been used to build a computational quality metric for moving pictures [14] which proved to behave consistently with human judgments. Basically, the metric, termed Moving Pictures Quality Metric (MPQM), first decomposes an original sequence and a distorted version of it into perceptual components. The quality measure is then computed, accounting for contrast sensitivity and masking (see Fig. 8) and scaled from 1 to 5 as described in Tbl. 1 [16].

Recently, an improved metric, termed Normalized Video Fidelity Metric (NVFM), based on a finer modeling of vision, has been introduced in [17]. This

9

Rating	Impairment	Quality
5	Imperceptible	Excellent
4	Perceptible, not annoying	Good
3	Slightly annoying	Fair
2	Annoying	Poor
1	Very annoying	Bad

Table 1: *Quality scale that is often used for subjective testing in the engineering community*

new metric adds a modeling of the saturation characteristic of the cortical cells' responses and a modeling of inter-channel masking. It is an extension of a still-picture model developed by Teo&Heeger [18].


Figure 8: *Moving Pictures Quality Metric (MPQM) block diagram*

4.3 Performance Measurements and Analysis

To study the performance of the proposed AAL, we have carried out simulations based on the setup described in Sec. 4.1 for different background loads varying between 79 and 86%. The background sources generate a balanced load. Figure 9 shows three curves with the CLRs measured at the receiver, which for AAL5 includes all the data lost due to the packet discard mechanism.

The CLR measured for the new AAL is equal to the network loss ratio since there is no extra discard of information at the AAL. For all the experiments, we obtain an improvement factor close to 8 in terms of CLR. This is basically due to the size in cells of the PDUs transmitted and suggests a low correlation in the cell loss process. It is interesting to note that using larger PDU sizes with the proposed AAL will not affect the CLR measured. This will not be the case if AAL5 is used since the impact of cell losses will be amplified by the packet discard mechanism albeit the probability of having more than one cell lost within the same packet will also increase. To perform the selective FEC experiment, we have tuned the network adaptation layer to protect three types of MPEG-2 headers, namely: sequence, picture and Packet Elementary Stream (PES) headers. We have protected a PDU with a single cell as described in Sec. 4.1. The overhead due to this mechanism is very small (4.75×10^{-3}) since there is little amount of loss sensitive data. Consequently, the results do not
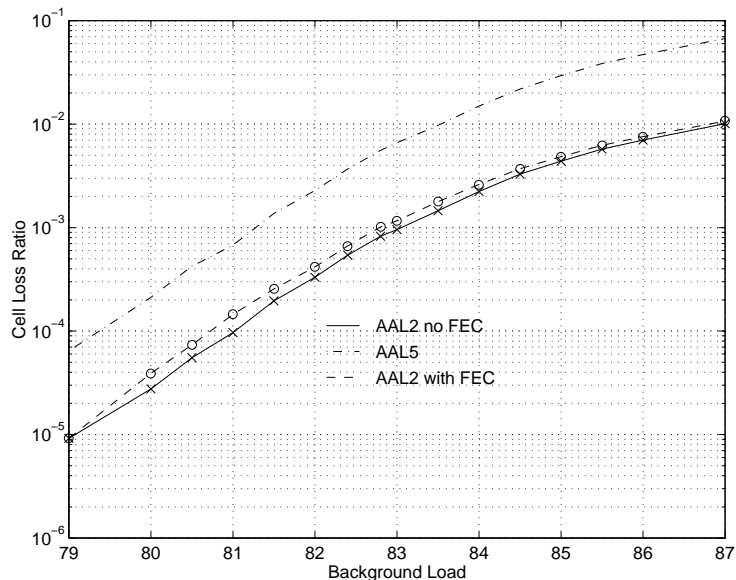
Figure 9: *Cell Loss Ratios seen by receivers*

show any particular improvement on the CLRs. They are even a little bit higher than in the no FEC case. However, the measurements depicted in Tbl. 2 show that a certain number of cells and therefore headers have been recovered. This means that syntactic information that may have been lost in the other two experiments will be available for decoding.

Load	Data Cells Recovered
79	0
80	2
81	7
82	15
83	46
84	97
85	147
86	244
87	339

Table 2: *Recovery Efficiency of the Selective FEC Mechanims*

One of the major difficulties with video applications is the mapping of data loss impact into the quality perceived by the user. Due to the different types of data carried in a multimedia stream, the impact of loss may be different. To assess the efficiency of the proposed mechanisms, we have used a perceptual quality metric described in Sec. 4.2. We have also applied error concealment techniques [8] that reduce the impact of cell losses on the perception of decoded

11

video sequences. Using these techniques allows us, for a given confidence interval, to estimate the mean value of the perceptual quality on a lower number of frames. The quality estimation has been performed on 100 frames.
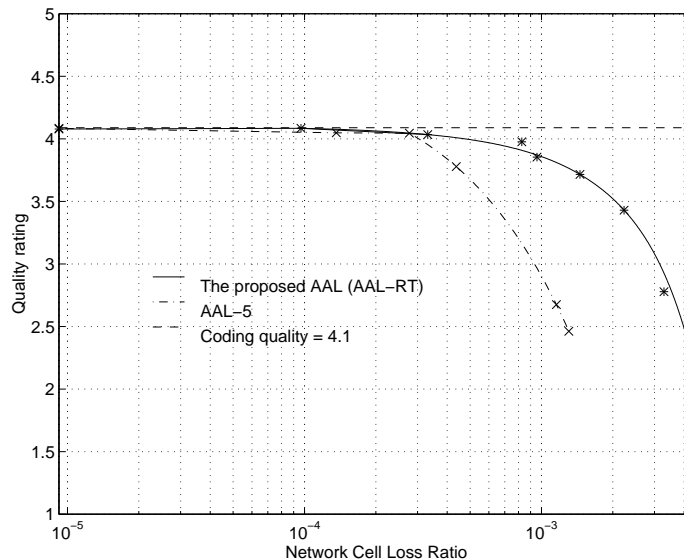

Figure 10: *Quality Rating vs. Network CLR for AAL5 and New AAL*

Figure 10 shows the MPQM quality assessment as a function of the cell loss ratio for the transmission of MPEG-2 coded streams over AAL5 and the proposed AAL. While increasing the CLR, the perceptual quality obtained with the proposed AAL remains nearly constant at the maximum value corresponding to the quality of the MPEG-2 coded sequence compared to the original not coded one (horizontal dotted line). This is due to relatively sparse cell losses easily masked by the use of error concealment techniques which therefore leads to a very low impact onto the quality. Beyond a CLR close to 3×10^{-4}, the perceptual quality drops smoothly since only a few frames have been lost, keeping efficient the error masking. Conversely, the curve obtained with AAL5 shows a different behavior. Two regions separated by a critical CLR can be distinguished: a first one with almost no quality degradation and a second one where the quality drops fast. This is mainly due to the packet discard mechanism. Considering the fact that for a single cell loss two TS packets are discarded, it is clear that the probability of losing syntactic information is higher than for the proposed AAL. Therefore more frames and PES may be lost. Also, if the data has been lost within the picture, the *holes* will be bigger and the masking less efficient than for the proposed AAL. Having less information available, the concealment mechanisms are less efficient leading to a faster quality degradation. Indeed, beyond the critical CLR, for a given loss ratio we achieve a significant gain in terms of perceived quality.

Figure 11 presents the gain in perceptual quality we have obtained by protecting syntactic information, namely, sequence, frame and PES headers with
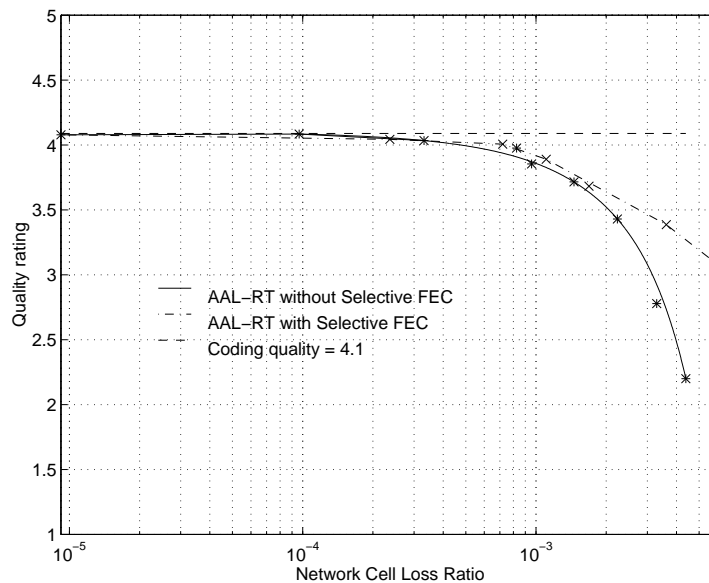
Figure 11: *Quality Rating vs. Network CLR for New AAL and New AAL with Selective FEC*

the selective FEC mechanism. The utilization of selective data protection that generates very small overhead (0.47%) reduces the slope of the curve, hence, smoothing out even more the quality degradation. Under severe cell loss conditions, the transmission with selective FEC proves to be very robust. It is worth to note that we have applied a very simple protection scheme that proves to be efficient. Using more elaborated data protection algorithms may improve even more the quality degradation figures while still having a small FEC overhead.

5 Conclusions

We have presented in this paper mechanisms at the AAL to improve the transmission of real-time multimedia applications. We increase the cell error detection resolution to the cell level. We also propose a cell-based selective FEC mechanism to recover from cell losses. A basic network adaptation layer has also been presented that covers the segmentation of packets into fixed size PDUs and a selective FEC request mechanism. Although network adaptation layer mechanisms are generic they necessarily have to be specific to the application since the request of FEC data relies on the knowledge of the information to be transmitted. Given the difficulty to map cell losses onto image quality, we have used a perceptual quality metric as a tool to evaluate network performance. We have carried out experiments with MPEG-2 video streams. The results obtained show an improvement of the cell loss figures for the simulations with and without selective FEC compared to current transmission methods as defined by the ATM Forum. From a perceptual point of view, the proposed AAL achieves a graceful quality degradation compared to AAL5 which shows a critical CLR value beyond which quality drops very fast. The application of a low overhead

13

selective FEC scheme does not show any significant CLR improvement however it smooths out even more the quality degradation observed under severe CLR ratios. The FEC mechanism applied which protects a subset of the syntactic information, even though it is very basic, gives good results. Future experiments will be done with improved data protection schemes to increase the robustness of audiovisual applications to data loss.

References

[1] Thomas Stock and Reto Gruenenfelder, "Frame Loss vs. Cell Loss in ATM Concentrators and Multiplexing Units", *in EFOC&N'94*, 1994.

[2] The ATM Forum, "Video on Demand Specification 1.0", Jan. 1996.

[3] Xavier Garcia Adanez, Andrea Basso, and Jean-Pierre Hubaux, "Study of AAL5 and a New AAL Segmentation Mechanism for MPEG-2 Video over ATM", *in 7th Workshop in Packet Video*, Brisbane, Mar. 1996, available on `http://tcomwww.epfl.ch/~garcia`.

[4] A.J. McAuley, "Reliable Broadband Communications Using a Burst Erasure Correcting Code", *in ACM SIGCOMM'90*, pp. 297–306, Philadelphia, PA, 1990. ACM, Academic Press.

[5] ISO/IEC JTC-1, editor, *Information Technology - Generic Coding of Moving Pictures and Associated Audio Information - Part 1: Systems Specification*, ISO/IEC JTC-1, 1994.

[6] H. Leopold, G. Blair, A. Campbell, G. Coulson, P. Dark, F. Garcia, D. Hutchinson, N. Singer, and N. Williams, "Distributed Multimedia Communication System Requirements", OSI 95 Consortium, May 1992.

[7] ITU-T, editor, *Recommendation I.363 B-ISDN ATM Adaptation Layer (AAL) Specification*, ITU-T, Geneva, 1993.

[8] Carlos Lopez Fernandez, Andrea Basso, and Jean-Pierre Hubaux, "Error Concealment and Early Resynchronization Techniques for MPEG-2 Video Streams Damaged by Transmission over ATM Networks", *in SPIE*, Brisbane, Mar. 1996.

[9] Mohammad Ghanbari and Vassilis Seferidis, "Cell Loss Concealment in ATM Video Codecs", *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 3, June 1993.

[10] ISO/IEC JTC-1, editor, *Information Technology - Generic Coding of Moving Pictures and Associated Audio Information - Part 2: Video*, ISO/IEC JTC-1, 1994.

[11] Rastislav Slosiar, *Performance Analysis Methods of ATM-Based Broadband Access Networks Using Stochastic Traffic Models*, PhD thesis, TCOM Laboratory, Swiss Federal Institute of Technology, CH 1015 Lausanne, Switzerland, 1995.

[12] Serge Comes, *Les traitements perceptifs d'images numerisées*, PhD thesis, Université Catholique de Louvain, 1995.

[13] Christian J. van den Branden Lambrecht, "A Working Spatio-Temporal Model of the Human Visual System for Image Restoration and Quality Assessment Applications", *in ICASSP*, pp. 2293–2296, Atlanta, GA, May 7-10 1996.

[14] Christian J. van den Branden Lambrecht and Olivier Verscheure, "Perceptual Quality Measure using a Spatio-Temporal Model of the Human Visual System", *in Proceedings of the SPIE*, vol. 2668, pp. 450–461, San Jose, CA, January 28 - February 2 1996.

[15] Christian J. van den Branden Lambrecht, *Perceptual Models and Architectures for Video Coding Applications*, PhD thesis, Swiss Federal Institute of Technology, CH 1015 Lausanne, Switzerland, 1996, available on `http://ltswww.epfl.ch/pub_files/vdb/`.

[16] M. Ardito, M. Barbero, M. Stroppiana, and M. Visca, "Compression and Quality", *in Proceedings of the International Workshop on HDTV 94*, Brisbane, Australia, October 1994.

[17] Pär Lindh and Christian J. van den Branden Lambrecht, "Efficient Spatio-Temporal Decomposition for Perceptual Processing of Video Sequences", *in ICIP*, Lausanne, Switzerland, September 16-19 1996, accepted for publication, available on `http://ltswww.epfl.ch/~vdb`.

[18] David J. Heeger and Patrick C. Teo, "A Model of Perceptual Image Fidelity", *in ICIP*, pp. 343–345, Washington, DC, October 23-26 1995.