

INTERNATIONAL ORGANISATION FOR STANDARDISATION
ORGANISATION INTERNATIONALE DE NORMALISATION
ISO-IEC/JTC1/SC29/WG11
CODED REPRESENTATION OF PICTURE AND AUDIO INFORMATION

ISO/IEC JTC1/SC29 WG11
MPEG 93/
January 1993

SOURCE: Matthew Leditschke, on behalf on the Australian UVC Consortium
(Telecom Australia & Monash University)
TITLE: Spatial localisation for cell loss resilience
PURPOSE: Information

1 Introduction

This document presents the results of a comparison between four different AAL packetisation schemes. The four schemes consist of streaming mode and message mode, with and without macroblock resynchronisation. Macroblock resynchronisation is the process used in the structured packing approach which enables decoding to resume with the first cell received after a loss. For this study it has also been applied in conjunction with message mode.

The work presented in this document is relevant to the discussion concerning the structured packing core experiment. While the results of this work are not directly related to the core experiment because of the alternative performance measure used (see below), they provide compelling evidence of the need to resynchronise quickly after a cell loss has occurred, whether by structured packing, small slices or some other technique.

2 Experimental Conditions

Table 1 summarises the conditions under which the experiments are performed. The 30 frame test sequence is coded once, and then the resulting bitstream has cells randomly discarded (burst length of 1, using the standard unix random number generator). The errored bitstream is then examined to determine the effect of the loss (see below for more details).

The process of discarding cells and examining the bitstream is repeated using different seeds for the random number generator. Using this technique, the 30 frame encoded sequence is decoded with a different pattern of lost cells a total of 10 times, giving an effective decoded sequence length of 300 frames.

The following sections (2.1 through to 2.4) give a summary of the four AAL packetisation schemes studied in this report. Streaming mode consists of continuously filling up cells with bits generated by the encoder, regardless of what information the bits represent. Message mode [1]

Coding Parameter	Values used
Sequences (720 x 576 x 25Hz)	Mobile & Calendar Tempete
Fixed bit rates	4 Mbits/second 9 Mbits/second
Type of coder	TM2 single layer frame based
Frame types	N = 12, M = 4 (I B B B P B B B P B B B I ...) N = large, M = 1 (I P P P P P P ...)
Slice sizes (in number of macroblocks)	22, 44, 66 and 88
Number of frames used	30
Random number seeds used	10
Cell loss ratios (applied uniformly)	0.1, 0.05, 0.01, 0.005, 0.001, 0.0005, 0.0001, 0.00005

Table 1: A summary of the parameters used for the experiments.

places the bits generated from coding a slice into a *message*, comprising a whole number of cells. As the size of a slice doesn't match an integral number of cells, there will be some unused space in the last cell of a message.

Macroblock resynchronisation [2, 3] enables the decoder to resynchronise decoding at the start of every cell, enabling quick recovery from lost cells. Implementing this requires extra information to be added to every cell, which indicates the start of the first macroblock within the cell, and where on the screen it belongs. In addition the encoder must code this first macroblock in a cell absolutely (motion vectors and DC coefficients), which would require a change to the MPEG video bitstream syntax .

2.1 Streaming mode without macroblock resynchronisation

AAL Header field(s)	Sequence number (4 bits)
Cell payload size	380 bits
Recovery from errors	Cells are discarded until a cell containing the start of a slice start code is found (provided that the following cell is also received if the slice start code is split across a cell boundary).

2.2 Streaming mode with macroblock resynchronisation

AAL header field(s)	Sequence number (4 bits) Macroblock start position (9 bits) Absolute block address (11 bits)
Cell payload size	360 bits
Recovery from errors	Decoding resumes with the first complete macroblock received after the lost cell. Only macroblocks contained (or partly contained) in the lost cell are considered as lost.
Additional requirements	The first macroblock in a cell must have any relative information coded absolutely. This applies to the DC coefficient in intra coded macroblocks and to motion vectors in inter coded macroblocks.

2.3 Message mode without macroblock resynchronisation

AAL header field(s)	Sequence number (4 bits) Segment type (2 bits)
Cell payload size	378 bits
Recovery from errors	Cells are discarded until the first cell containing the start of a message is found (which will also contain the start of a slice).
Additional requirements	The first 24 bits of all start codes (always 23 zeros followed by a 1) are discarded because a start code will always be found at the beginning of a message, which is detected by the AAL.

2.4 Message mode with macroblock resynchronisation

AAL header field(s)	Sequence number (4 bits) Segment type (2 bits) Macroblock start position (9 bits) Absolute block address (11 bits)
Cell payload size	358 bits
Recovery from errors	Decoding resumes with the first complete macroblock received after the lost cell. Only macroblocks contained (or partly contained) in the lost cell are considered as lost.
Additional requirements	The first macroblock in a cell must have any relative information coded absolutely. This applies to the DC coefficient in intra coded macroblocks and to motion vectors in inter coded macroblocks. The first 24 bits of all start codes are discarded because a start code will always be found at the beginning of a message, which is detected by the AAL.

3 Results

The relative performances of the different schemes are compared in terms of the average number of macroblocks lost per cell loss. This measure is used in place of a sequence quality measure such as the signal to noise ratio (SNR), because using the SNR will require reconstruction of the sequence at the decoder, the quality of which depends highly on the error concealment techniques employed in the decoder. As this document is concerned with the AAL packetisation process rather than error concealment, the average number of macroblocks lost per cell loss gives a better indication of the ability of a packetisation scheme to transport the bitstream in an error prone environment.

By removing the need to conceal lost macroblocks the number of variables in the study has been decreased, but the results presented require interpretation in a different manner than if SNR were used. Although the eventual sequence quality depends on the concealment technique used, the lower the number of macroblock lost per cell loss (i.e. fewer macroblocks lost per frame), the less concealment that needs to be performed, and the better the resulting sequence quality. Thus there is some correlation between the number of macroblocks lost per cell loss and the quality (as measured by SNR) of the sequence.

3.1 A comparison of overhead information

The graph in Figure 1 shows the percentage of the total payload which comprises AAL imposed overhead information, for each of the four AAL packetisation schemes. This overhead information includes the header fields described in the previous section, and unused cell capacity which results from using message mode.

When coding at a fixed bit rate, the variation in the amount of overhead information between the schemes also varies the bandwidth available for coding the video. The graph in Figure 2 shows how the coded sequence quality varies with changes in the slice size, for each of the four packetisation schemes.

The following results can be ascertained from the graph in Figure 2:

- The curves representing message mode show the effects of the high overhead which results from using small slices. The lower quality for small slices corresponds with the high amount of overhead information, as shown in Figure 1.
- The difference in quality between the four packetisation schemes is larger for a slice size of 22 macroblocks (the difference in SNR values is $\sim 0.7\text{dB}$) than it is for slices containing 88 macroblocks (the difference in SNR values is $\sim 0.4\text{dB}$). Either way, the overall quality difference is quite small, even though the variation in the amount of overhead information is much larger.

3.2 Macroblocks lost per cell loss: The effect of slice size

The graphs presented in Figure 3 show the average number of macroblocks lost per cell loss as a function of the slice size and average cell loss ratio, for each of the four packetisation schemes. Because of the large number of parameters, the results are presented in four 3 dimensional graphs. The two graphs on the left show the results of both message mode and streaming mode with macroblock resynchronisation, while the two graphs on the right are for the two schemes without macroblock resynchronisation.




Figure 1: A comparison between the overhead imposed by the four different AAL packetisation schemes, for varying slice sizes. (Sequence: *Mobile & Calendar*, rate = 4 Mbits/s, N = large, M = 1)




Figure 2: A comparison between the possible quality of the different AAL packetisation schemes, taking into account the overhead information. (Sequence: *Mobile & Calendar*, N = large, M = 1)




Figure 3: The performance of the four packetisation schemes, as functions of the cell loss ratio and the slice size. (Sequence: *Mobile & Calendar*, rate = 4 Mbits/s, $N = \text{large}$, $M = 1$)

The graphs show the expected results that when macroblock resynchronisation is used, changes in the slice size have no effect on the time taken to resume decoding (except for some small fluctuations at low cell loss rates). The results presented without macroblock resynchronisation also confirm that the average number of macroblocks discarded before decoding resumes is half the slice size. The small drop at high cell loss rates is due to counting cell losses which occur during error recovery.

3.3 Macroblocks lost per cell loss: The effect of coding rate

In this section the effect of the coding rate (fixed bit rate coding) is investigated. The graph in Figure 4 shows the average number of macroblocks lost per cell loss as a function of coded bit rate, for each of the four packetisation schemes. For each of the packetisation schemes, the number of macroblocks lost per cell loss decreases as the coding rate increases.




Figure 4: The performance of the four packetisation schemes, as functions of the coded bit rate. (Sequence: *Mobile & Calendar*, macroblocks per slice = 44, cell loss ratio = 0.005, N = large, M = 1)

When coding at higher rates, the average number of bits produced by a coded macroblock increases, resulting in fewer coded macroblocks being transported per cell. Thus for the two packetisation schemes which use macroblock resynchronisation, fewer macroblocks will be lost when a cell is lost.

Without macroblock resynchronisation, the number of macroblocks lost is proportional to the size of the slice (see Figure 3 on page 6), and only coded macroblocks are counted as lost (as uncoded ones aren't sent and so can't be lost). An explanation for the decrease in the average number of macroblocks lost as the rate increases is that the higher quality obtained at the higher rate increases the performance of the prediction algorithm, producing more uncoded macroblocks.

3.4 Macroblocks lost per cell loss: The effect of frame types

The examples used in this document so far have used a coding pattern with the first frame coded as intra and all subsequent frames coded as inter: N (the intra frame spacing) = *large*

and M (the prediction frame spacing) = 1. The graph in Figure 5 compares this coding pattern with an alternative, namely $N = 12$ and $M = 4$.




Figure 5: The performance of the four packetisation schemes, as functions of the frame types used. (Sequence: *Mobile & Calendar*, macroblocks per slice = 44, cell loss ratio = 0.005, rate = 4 Mbit/s)

The graph in Figure 5 shows that there is a only a slight dependence upon the coding pattern used. The two schemes with macroblock resynchronisation show no noticeable change in the average number of macroblocks lost per cell loss for $N=12$, $M=4$ coding, whereas the schemes without macroblock resynchronisation show a slight decrease.

This observable decrease is possibly due to an increase in the number of uncoded macroblocks, resulting from the better prediction and coarser quantisation of B-frames. As with the previous section, an increase in the number of uncoded macroblocks leads to a decrease in the average number of macroblocks lost per cell loss, when macroblock resynchronisation is not used.

4 Conclusions

This document describes a simulation based comparison of four AAL packetisation schemes. The comparisons are made in terms of the average number of macroblocks lost per cell loss, rather than using the decoded sequence quality (measured by the signal to noise ratio). This measure is used to separate the packetisation and transportation from the decoding and concealment techniques, which can be applied in conjunction with any of the packetisation schemes.

The two schemes which used macroblock resynchronisation give much better results (between 3 to 10 times better, depending on the slice size used). The drawback of these schemes is quite a significant one, namely the requirement to have the first complete macroblock in a cell coded absolutely (i.e. motion vectors and intra DC coefficients are coded absolutely, rather than relatively). This requires a change in the syntax, and also implies that the bitstream generation process needs to be very closely linked into the packetisation process. This close linking is required so that the necessary macroblocks can be coded absolutely, and the cell header field values written.

The performance of message mode and streaming mode without macroblock resynchronisation is quite comparable, in that the average number of macroblocks lost per cell loss is approximately proportional to the size of the slice.

If slices smaller than 22 macroblocks were used [4], then the performance of streaming mode would approach that of the two schemes using macroblock resynchronisation (extrapolating the results in Figure 3). Message mode without macroblock resynchronisation for these small slice sizes would be inefficient due to the large overhead of unused cell capacity.

References

- [1] Australian contribution to CCITT SGXV Experts group for ATM video coding. *The ATM Adaptation Layer for video services in the B-ISDN*, July 1992. Document AVC-297.
- [2] Japanese contribution to CCITT SGXV Experts group for ATM video coding. *Cell-loss compensation scheme*, March 1992. Document AVC-235, MPEG92/184.
- [3] Japanese contribution to CCITT SGXV Experts group for ATM video coding. *Structured packing*, October 1992. Document AVC-360, MPEG92/579.
- [4] Japanese contribution to CCITT SGXV Experts group for ATM video coding. *Re-synchronisation by slice size reduction*, October 1992. Document AVC-368, MPEG92/587.