

SOURCE: UK, NETHERLANDS, SWEDEN
TITLE: LOOP FILTERS IN THE OKUBO RM2.

Introduction

A general class of post motion compensation low pass filters are considered here. The purpose of the filter is to minimise the prediction error. Several possible methods exist by which the data can be filtered in a 'post motion compensation' sense. We examine the relative merits of filtering within the 8×8 block boundaries, filtering over the block boundaries in the motion predicted frame, filtering over the block boundaries in the previous reconstructed frame and the use of adaptive filters as opposed to fixed filters.

All the results presented in this report are generated using the Trevor/Split Screen CIF sequence at 30kbits/coded frame and 10 coded frames /second. The first 19 coded frames in the data are Split Screen the remaining 29 coded frames are Trevor.

Low Pass Filter Post Motion Compensation

The position of the low pass filter in the coding loop is illustrated in Figure 1. The current CCITT hardware specification of the CODEC states that the motion compensation is optional at the encoder, therefore any coder design may not include the motion compensation block indicated in Figure 1. For the results shown here the filter is applied adaptively.

The introduction of a low pass filter post motion compensation (MC) would appear to have the following advantages:

- A reduction of high frequency artifacts introduced by MC.
- A reduction in quantisation noise circulating in the feedback loop.
- A reduction in unwanted high frequencies circulating in the feedback loop.
- For systems with motion compensation the filter could be made adaptive as a function of motion vectors with no additional overhead cost.
- Assuming that the filtered block is coded in an INTER mode, then the high frequency components will actually be coded INTRA and the low frequency components INTER.

Three distinct forms of filter (in a spatial application sense) may be defined. These are:

- 1] The filter acts strictly on data within an 8×8 block.
- 2] The filter acts over an area greater than 8×8 in the previous predicted frame domain.
- 3] The filter acts over an area greater than 8×8 in the motion predicted frame domain.

Various forms of low pass filters can then be applied under the conditions above. The filters that have been examined are:

- a] Non-separable filter
- | | | | |
|--|---|---|---|
| | | 1 | |
| | 1 | 4 | 1 |
| | | 1 | |
- b] Separable filter
- | | | | |
|--|---|---|---|
| | 1 | 2 | 1 |
| | 2 | 4 | 2 |
| | 1 | 2 | 1 |
- c] 3 separable filters
- | | | |
|-------|---|-------|
| 1 2 1 | 1 | 1 2 1 |
| | 2 | 2 4 2 |
| | 1 | 1 2 1 |
- d] Separable filter
- | | | | |
|--|---|----|---|
| | 1 | 4 | 1 |
| | 4 | 16 | 4 |
| | 1 | 4 | 1 |
- e] 3 separable filters
- | | | |
|-------|---|--------|
| 1 4 1 | 1 | 1 4 1 |
| | 4 | 4 16 4 |
| | 1 | 1 4 1 |

Results are presented for filters 'a' - 'c' however it was found that filters 'd' and 'e' yielded similar results.

Experiments indicate that applying filter 'a' in RM2 regardless of motion vector or whether the block would have been fixed degrades the performance of the model. Thus all of the above filters were applied only when the motion vector was non zero, unless indicated otherwise.

Filter 'c' can be applied adaptively when motion compensation is used. When using filter 'c' for the data presented herein if the motion vector was found to be purely horizontal then the data was filtered with a 1 2 1 horizontal filter. If the vector was found to be purely vertical then the data was filtered with a $[1 \ 2 \ 1]^T$ vertical filter. If both horizontal and vertical motion vectors were non-zero then the 2-dimensional filter was used.

It would appear that different combinations of 1 - 3 and a - c give different hardware complexities, therefore some of these possible combinations are now examined in terms of coding performance.

Note: The results presented here are not in accordance with the recommended methods, as it is thought the presentation here illustrates the results more clearly.

Filtering within the 8 * 8 block

Applying filter 'a' or 'b' directly to an 8 * 8 block implies that unless special conditions exist for the peripheral pels in the block these pels will not be filtered. Thus 43% of the data will not be filtered with this arrangement. Experiments have shown that while this limited action filter actually gives an improvement when used in the Okubo Reference Model 2 (RM2) the improvement is small compared to other schemes illustrated below. If one were to apply filter 'c' within a block then the 3 * 3 mask could be approximated by applying the two 1-dimensional filters one after

the other. In this case only 4 pels (ie the corner pels) would remain unfiltered. For the pels not on the edge of the block the 2 pass filter approach is identical to a 1 pass approach using the full 3×3 mask, however this is not true for the peripheral block pels. This arrangement has only been investigated when no motion compensation was used in the coding loop and is discussed below. The other scenarios will not be discussed further.

The action of filter 'a' or 'b' can be significantly improved either by modifying the coefficients at the block edges or by extrapolating block data so as to generate synthetic data around the block being filtered. Both these schemes appear to be equivalent and allow the peripheral pels in each block to be filtered to some extent. The block extrapolation technique has been examined for filters 'a', 'b' and 'c'. The block extrapolation process consisted of pel repeating at the block edges.

Figure 2 illustrates the comparative results between RM2 and RM2 using filter 'a'. Figure 3 illustrates the comparative results between RM2 and RM2 using filter 'b'. Figure 4 illustrates the comparative results between RM2 and RM2 using filter 'c'.

Clearly the introduction of filter 'a', 'b' or 'c' improves the performance compared to RM2. The difference in the improvement between filters 'a' and 'c' is very small ie filter 'c' is 0.025 dB better than filter 'a'. Filter 'b' shows only a small improvement over the reference model. It is thought that this filter is too severe for the source data used here.

Filtering over a 10×10 block in the previous predicted frame domain to produce an 8×8 block prediction.

This method is outlined in [1] and represents the method that has been most examined. Figure 5 illustrates the comparative results between RM2 and RM2 using filter 'a'. It has been shown that filter 'c' produces marginally superior results compared to filter 'a'. The margin is similar to the above case. Again an improvement compared to the reference model is obtained. The results are very similar to those illustrated in Figures 2 and 4.

Instead of using motion vectors to switch the loop filter on or off an alternative switching mechanism was tried. This operated by comparing the prediction error for filtered blocks and unfiltered blocks, regardless of motion compensation. The block with the minimum prediction error was used and 1 bit was used to indicate whether the block was filtered or not. Filter 'b' was used. The objective results remain similar to those shown however the subjective results appeared to be the marginally better compared to when the decision was made as a function of the motion vector.

Filtering over a 10×10 block in the motion predicted frame domain to produce an 8×8 block prediction.

The motion predicted data is initially generated for the block to be filtered and all adjacent blocks then a filter is applied to this data such that filtering occurs over block boundaries which are contiguous in the motion predicted data. Figure 6 illustrates the comparative results between RM2 and RM2 using filter 'a'. Again an improvement compared to the reference model is obtained. The results are very similar to those illustrated in Figures 2, 4 and 5. Experiments have also indicated that filters 'b' and 'c' yield similar results.

The Application of a Loop Filter when Motion Compensation is not performed in the Encoder.

As the current CCITT specification for hardware states that motion compensation should be optional in the encoder, then one must consider how a filter can operate (if desirable) without motion compensation.

In order to allow a degree of adaptivity, attribute bits must be sent to the decoder as overhead information. Experiments have been conducted applying filter 'c' to data within the 8×8 block (1) and using external block data (2 and 3 are equivalent here). Applying the filter over the block edges yielded better results than restricting the filter within the block (no block extrapolation was used). Both these methods assumed that 2 bits of side information were sent for each block to indicate the filter type. The signal to noise ratios produced by the RM2 with motion compensation were typically 2 - 3 dB better than without motion compensation. The subjective quality of the coded sequences was also significantly better for RM2 with motion compensation. It has been found however that the subjective performance of the RM2 without motion compensation is improved significantly with the use of a loop filter. Further it is thought that the subjective performance of the RM2 without both a loop filter and motion compensation is not acceptable.

Conclusion

It is clear that the application of a filter within the coding loop as described above improves the quality of the coded pictures. So far it has not been demonstrated that filters which are made adaptive to the motion vectors (as in filter 'c') exhibit significantly superior performance to a fixed filter. However it has been shown that the filter must be capable of being switched on or off. With the exception of the RM2 without motion compensation it has been shown that the objective quality of the coded pictures remains relatively invariant of how the filter is applied (ie methods 1 - 3). It has been indicated that the subjective quality of the coded pictures also remains invariant of how the filter is applied (assuming that the filter is applied as a function of motion vectors).

The results for RM2 without motion compensation indicate that the data should be filtered over the block boundaries for best subjective performance.

Proposal

The filter should be capable of accessing a 10×10 area in order to produce a prediction of 8×8 . The decision for the filter to access either the motion predicted frame data or the previous reconstructed frame data should be made on hardware grounds as the two techniques yield very similar results both objectively and subjectively. For the case where no motion compensation is used the two situations are identical. It has not been shown that adaptivity more than being able to switch a filter on or off is required therefore a single filter is proposed which may be switched on or off. The filter can be switched using either motion vectors or some other criterion requiring additional attribute bits.

It has been shown that any improvement in performance is invariant of whether the

filter is separable or non-separable, therefore this parameter should not be considered as important for filter performance. Further it has been shown that all the filters considered yield similar performances, therefore the decision to use filter 'a' , 'b' or 'd' should be made on hardware grounds.

The actual coefficients used in the filter should be programmable and could be optimised at a later date.

References

[1]...A Hardware Aspect of multi-pel Predictors. Doc #120 CCITT SGXV June 86.





Figure 1. Outline of the coding configuration

L H11A


FIGURE 2

Filter process uses mask $\begin{matrix} 1 \\ 4 \\ 1 \end{matrix}$



The filter is applied to an extrapolated block.



Mean value of reference = 36.207180
Mean value of modified = 37.438083




Mean value of reference = 4
Mean value of modified = 3.4845175




8'441#

FIGURE 3



Filter mask is $\begin{matrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{matrix}$ and is
to an extrapolated block



Mean value of reference = 26.207180
Mean value of modified = 27.3087




Mean value of reference = 4
Mean value of modified = 3.548809




#144,
9

FIGURE 4



Filter mask is 1 2 1 and is applied adaptively as a function of the motion vector.



Mean value of reference = 36.207160
Mean value of modified = 37.482932




Mean value of reference = 4
Mean value of modified = 3.484359




#144, 10

FIGURE 5



Filter is a $\begin{matrix} 1 & & 1 \\ & 4 & \\ 1 & & 1 \end{matrix}$ mask
 applied to the previous
 predicted frame.
 The filter accesses a
 10*10 area and produces
 an 8*8 prediction



Mean value of reference = 38.207180
 Mean value of modified = 37.4592




Mean value of reference = 4
 Mean value of modified = 3.4858075




11'4411#

FIGURE 6


The filter is a $\begin{matrix} & 1 & \\ 1 & 4 & 1 \\ & 1 & \end{matrix}$ mask
 applied to the motion
 predicted frame.
 The filter accesses a 10*10
 area and produces an 8*8
 prediction.




Mean value of reference = 36.207160
 Mean value of modified = 37.591941



Mean value of reference = 4
 Mean value of modified = 3.439291



Mean value of reference = 39.367520



Mean value of reference = 11