NOISE REDUCTION OF IMAGE SEQUENCES AS PREPROCESSING FOR MPEG2 ENCODING

P.M.B. van Roosmalen[1] A.C. Kokaram[2] J. Biemond[1]

Information and Communication Theory Group[1]
Faculty of Information Technology and Systems
Delft University of Technology, The Netherlands
bruce@it.et.tudelft.nl

Electrical Engineering Dept.[2]
Trinity College, Dublin 2
Republic of Ireland
anil.kokaram@tcd.ie

ABSTRACT

This paper investigates the influence of three classes of noise reduction filters as preprocessors to MPEG2 coding of noisy image sequences. From each class we select a representative noise filter that performs well and we investigate the effects of these on the coding efficiency. We also investigate the use of an adjusted MPEG2 compression scheme for simultaneous noise reduction and compression. The quality of the filtered sequences are evaluated objectively and subjectively after MPEG2 coding at bitrates varying from 1.0 to 6.5 Mbit/s.

1. INTRODUCTION

With the advent of digital compression techniques such as MPEG, the home viewer will be able to choose from a great number of channels in the near future. These channels require programming, and, due to the high cost of creating new programs, many old films will be re-used. This will only be viable if the quality (both visual and audio) of those films meets the standards that the modern viewer is accustomed to. Many authors have proposed methods for correcting artifacts common to old film sequences such as noise [1-5], line scratches [6], and blotches [7]. Removing artifacts, and noise in particular, can lead to significant gains in image quality at identical bitrate, or, conversely, to a much lower bitrate at identical quality. This is because less bits need to be spent on irrelevant information. Several authors recognize a relationship between noise reduction and coding efficiency [8, 9].

This paper investigates the influence of three classes of noise reduction filters as preprocessors to MPEG2 coding of noisy image sequences. The goal is to determine whether the class of preprocessing filter influences the coding efficiency and what the effect of filtering is on the subjective quality of the coded sequence. From each class we select a representative noise filter that performs well. We also investigate the use of the MPEG2 compression scheme itself as a noise reducer. Comparisons are made using both objective and subjective evaluations.

The first filter we evaluate is the *3D Wiener* filter as described by *Kokaram* [1]. The Wiener filter is optimal in the class of linear filters in the sense that the mean squared error between filtered noisy image and original noise free image is minimal. The next filter we investigate is a wavelet based filter based on the *Algorithm a Trous* [3]. Wavelets have shown to be very powerful in many fields of image processing, including noise reduction. Wavelet noise filters form a class of non-linear filters that operate in a (discrete wavelet) transform domain. The transform coefficients are *cored* (or *thresholded)* according to some non-linear coring function. The third filter we investigate is an Order Statistics (OS) filter. OS filters are non-linear filters that operate in the spatial domain in which data is ordered before a weighted averaging is applied. Because of the ordering operation, spatial and temporal information is ignored in favor of magnitude information. A well known OS filter is the median filter. We apply the *3D Range Test* filter described in [4]. The MPEG2 compression scheme itself can also be used as a noise reducer: coring the DCT coefficients prior to quantization may lead to considerable noise reduction.

In Section 2 we describe the details on how the MPEG2 coding scheme can be applied to achieve noise reduction and compression simultaneously. Then, in Section 3, we describe the experiments and the results. We conclude this paper in Section 4.

2. MPEG2 FOR NOISE REDUCTION AND COMPRESSION

The MPEG2 system is based on *I*-frames and predicted *P* and *B* frames. The *I*-frames are coded efficiently by dividing the frame in 8x8 blocks, applying the *discrete cosine transform* (DCT) to these blocks and quantizing the DCT coefficients. Because for natural images most of the signal energy is concentrated in relatively few DCT-coefficients, most of the quantized coefficients are zero and need not be transmitted. In the presence of additive white gaussian noise, the noise is spread evenly over all the transform coefficients and less quantized coefficients are zero. To maintain efficient coding in the presence of noise, one must either apply coarser quantization or one must selectively, depending on the expected signal to noise ratio, reduce (*core)* the magnitude of each of the transform coefficients prior to quantization. The former removes much of the noise energy but unfortunately also strongly reduces the signal quality. The latter is a better solution and can be achieved following [10] where the procedures for estimating the signal pdf in the presence of white gaussian noise and for deriving the optimal coring coefficients are described.

The *B* and *P* frames are predicted from frames coded previously. The frame differences between predicted and current frame are coded like the *I*-frames, i.e. by using DCTs and quantization. Finding the ideal coring coefficients is much more difficult now because the signal and noise distributions of the frame differences are not known (these depend on the non-linear coring and quantization of frames coded earlier and on the quality of the motion estimation and compensation). However, the DCTs of the differences are the same as the differences of the DCTs of the respective frames. Assuming that the predicted frame is already noise reduced (since it is derived from previously processed frames) the only remaining requirement in coring the DCTs for noise reduction is to core the DCT coefficients of the original current frame. This can be done in the same manner as for the *I* frames. Thus we avoided the need to infer the signal and noise statistics of the frame differences between current frame and predicted frame.

Extra gains in noise reduction can be achieved for the *B* and *P* frames by realizing that the differences of the DCTs are in fact a high-pass signal in the temporal direction that contain both signal and (residual) noise. Applying, for example, soft-thresholding to the DCT coefficients thus further reduces the amount of noise.

3. EXPERIMENTS AND RESULTS

The goal of the experiments is to examine whether the type of prefiltering influences the coding efficiency. The experiments are set up as follows (Fig. 1). First we generate noisy test sequences with varying amounts of noise. These sequences are filtered using the filters described before. The filtered results are then compressed. After decompression, the final results are compared to the original sequences, using the PSNR as an objective measure or by subjective evaluation.

Section 3.1 describes the objective experiments. Section 3.2 describes the subjective experiments.

Filter	$\overline{\text{PSNR}}$ (dB)	$\overline{\text{PSNR}}$ (dB)	$\overline{\text{PSNR}}$ (dB)
None	33.0	27.0	23.5
3D Range Test	35.2	31.3	29.3
3D Wiener	35.3	31.4	28.6
3D Wavelet	35.4	32.4	30.1
MPEG2 (adjusted)	34.6	31.6	29.7

Table 1. Average PSNRs of noisy and noise reduced test sequences using various filters and using the adjusted MPEG2 coder. The latter was set to a bitrate of 15 Mbit/s.

3.1 Objective experiments

Three noisy test sequences were generated consisting of the *Mobile* sequence with low, moderate and high amounts of additive white gaussian noise with variances of respectively 25, 100 and 225. The intensities of all sequences were clipped to lie in a range between 16 and 240. The noisy sequences were filtered using each of the filters described before. The noise reduced results were then compressed using the TM5 MPEG2 encoder at bitrates ranging from 1.0 Mbit/s to 6.5 Mbit/s in steps of 0.5 Mbit/s. The noisy test sequences were also coded using our modified MPEG2 encoder (no prefiltering was applied here).

Table 1 shows the average PSNRs of the noisy and noise reduced sequences before they are MPEG encoded. We note that the three selected filters considerably increase the PSNRs. Table 1 also shows the results for our adjusted MPEG2 encoder when set to a high bitrate (this to measure its merit as a noise reducer more than its merit as a combination of a compressor and noise reducer) It too brings considerable noise reduction.

We note that for the Wiener filter the performance drops more than for the other filters for the third test sequence. The reason for this is that we assumed the noise spectrum to be flat. This assumption is not entirely correct due to the clipping operation mentioned before.

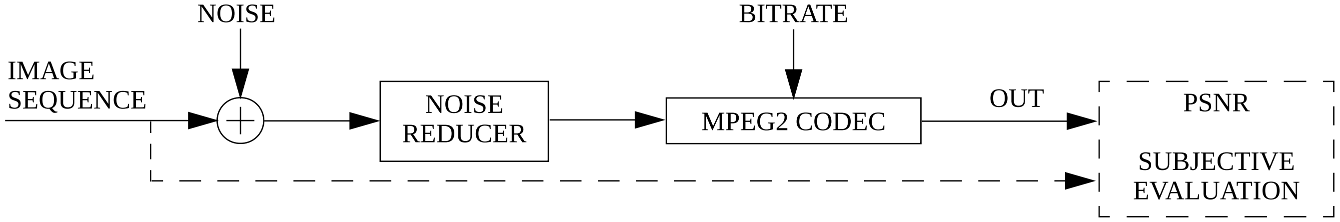

Figure 1. Schematic overview of experimental setup.

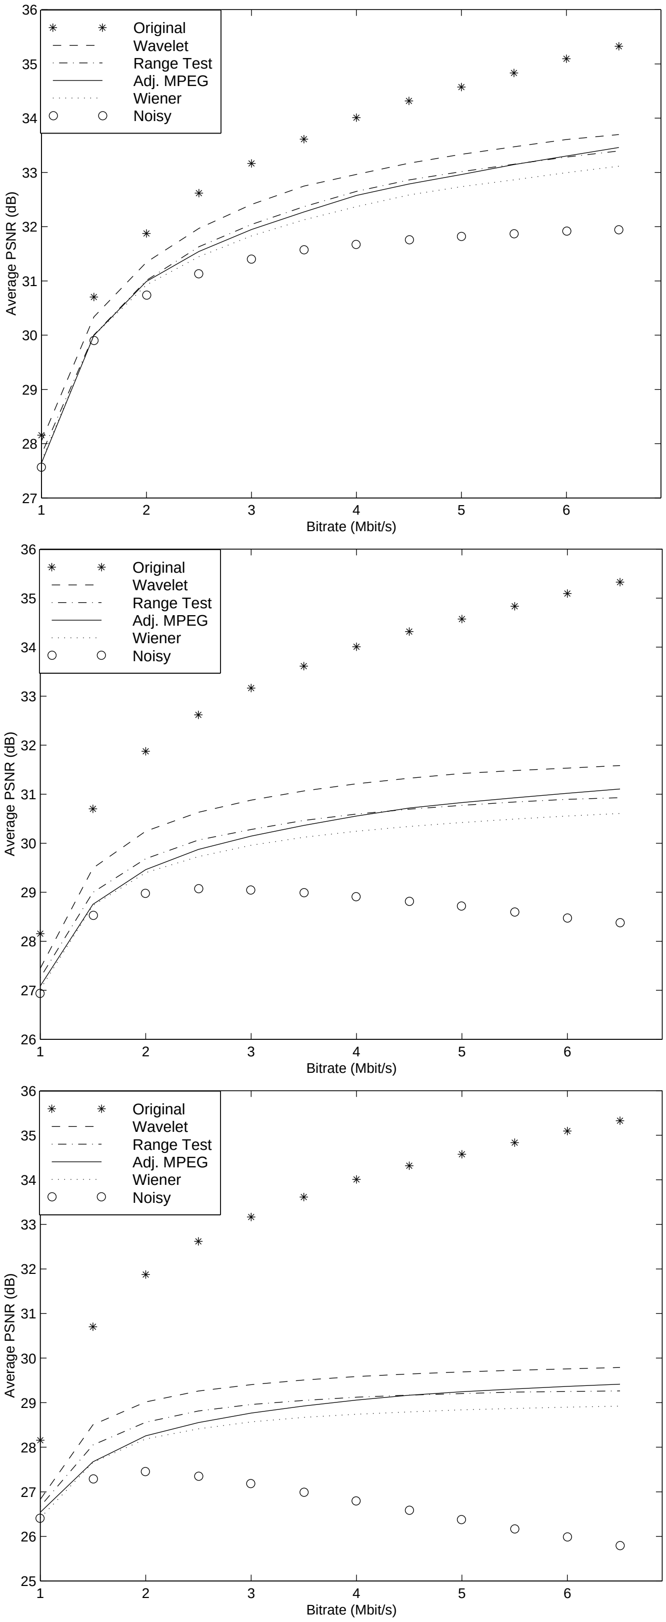

Figure 2. PSNRs of preprocessed sequences after MPEG2 encoding at various bitrates. Also shown are the PSNRs for the original clean sequence, the noisy sequences after coding with TM5. The results of our adjusted MPEG2 coder are shown as well. The noise variance in the test sequences are from top to bottom 25, 100 and 225 respectively.

Figure 2 shows the resulting PSNRs for various bitrates and filters after MPEG2 encoding/decoding of the filtered test sequences using the TM5 coder. Also shown are the PSNRs for various bitrates obtained from our modified MPEG2 for the noisy test sequence without prefiltering. The results for the original and noisy sequence coded by TM5 are shown as well.

It is interesting to note the following from figure 2. First, in the case of sequences containing moderate and high amounts of noise, the TM5 coder itself performs the function of a noise reducer. This can be deduced from the fact that for the two test sequences containing the most noise, the PSNR increases when using the standard TM 5 coder at low bitrates. At higher bitrates the PSNR decreases for these sequences. This is because at low bitrates the coder applies coarse quantization and relatively much noise energy is removed. At higher bitrates more bits are available for coding the noise more accurately, giving rise to lower PSNRs.

Second, we note that in all the combination of the wavelet denoising and standard coder outperforms the other filters in terms of PSNR. This is consistent with the results in Table 1. However, the PSNRs of the sequences coded by the modified MPEG2 coder are higher than that of the coded Wiener filter results. This is surprising because from Table 1 we see that for the first noisy test sequence (noise variance 25) the Wiener results before compression are better than those resulting from the modified coder. We see something similar for the Range Test filter compared to the adjusted coder. At low bitrates the coded results from the Range Test filter show higher PSNRs than the adjusted coder in all cases. At high bitrates, the adjusted coder has the advantage. Finally, we see that the Range Test filter outperforms the Wiener filter in all cases after coding even though we see from Table 1 that these filters show similar performance in the presence of low and moderate amounts of noise. We do not offer an explanation for these findings.

Third, we observe that, in cases of moderate and high amounts of noise, the PSNR of the filtered sequences initially rises sharply with increasing bitrate. However, at a certain point the slope reduces and there is little increase in PSNR with increasing bitrate. For example, in the bottom graph at 6.0 Mbit/s the PSNR for the wavelet filter is 29.8 dB whereas at 4.0 Mbit/s it is 29.6 dB, thus over 2 Mbit/s bandwidth there is only minor difference in PSNR. The explanation for this behaviour is that at low bitrates relatively many bits are spent on coding the original signal. As the bitrate increases more bits are spent on the coding of residual noise in the filtered sequences and on the coding artifacts introduced by the filter. The decrease in quality that would result from these is compensated by the fact that the original signal is also coded more accurately.

From the last observation one can conclude that 33% bandwith may be saved with neglible loss of quality! In the next section we verify this statement from a subjective point of view.

	A: filtered (5 Mbit/s) B: filtered (4 Mbit/s)	A: filtered (6 Mbit/s) B: filtered (4 Mbit/s)	A: filtered (6 Mbit/s) B: filtered (5 Mbit/s)	A: filtered (6 Mbit/s) B: original (6 Mbit/s)
Votes for A	56.2 %	43.8 %	50.0 %	2.5 %
Votes for B	43.8 %	56.2 %	50.0 %	97.5 %

Table 2. Results of subjective evalution using 2AFC testing. See text for explanation.

3.2 Subjective evaluation

In the previous section we noted that, in the case of sequences containing moderate and high amounts of noise, from a certain point the PSNR of prefiltered and MPEG2 coded sequences increases only by small amounts with increasing bitrate. In this section we investigate whether differences in quality are visible between filtered sequences coded at various bitrates.

The experiments are set up as follows. We selected the test sequence containing additive white gaussian with noise variance 100 and we filtered this sequence using the wavelet filter (PSNR: 32.4 dB, see table 1). The filtered sequence was then coded at 4.0, 5.0 and 6.0 Mbit/s (PSNR: 31.2, 31.4 and 31.5 dB). We also coded the original (noise free) sequence at 6.0 Mbit/s (PSNR: 35.1 dB) as a reference.

Next we formed four pairs of coded sequences:
1. (A) filtered at 5.0 Mbit/s - (B) filtered at 4.0 Mbit/s,
2. (A) filtered at 6.0 Mbit/s - (B) filtered at 4.0 Mbit/s,
3. (A) filtered at 6.0 Mbit/s - (B) filtered at 5.0 Mbit/s,
4. (A) filtered at 6.0 Mbit/s - (B) original at 6.0 Mbit/s.

Each pair of sequences was presented to a testpanel of eight persons. The panel was asked to indicate which sequence of each pair, A or B, had the best visual quality. The responses were obtained using the *two alternative forced choice* (2AFC) method [11]. Table 2 shows the voting results.

From table 2 we conclude that, from a subjective point of view, there is no clear preference for the bitrate at which filtered sequences should be coded. The fact that (at 6.0 Mbit/s) in 97.5 % of the cases the coded original sequence was preferred over the coded prefiltered noisy sequence shows that significant differences in quality would have been visible, i.e. that a PSNR of 31-32 dB is not so high that no differences would have been noticed no matter what bitrate had been chosen.

4. CONCLUSION

Prefiltering of image sequences leads to increased coding efficiency due to the fact the less bits have to be spent on coding irrelevant information (noise). We found the wavelet filter to be the most succesful preprocessor, the wiener filter was least succesful. The maximal difference in PSNR of the coded sequences was 1.2 dB (at 6.5 Mbit/s) for these filters.

Objective and subjective evaluations have shown, for sequences containing moderate and high amounts of noise, that coding at 4 Mbit/s is sufficient when using the TM5 coder in combination with prefiltering. More advanced coders may well allow even lower bitrates without loss of visual quality.

References

1 A.C. Kokaram, "3D Wiener Filtering for Noise Suppression in Motion Picture Sequences using Overlapped Processing", Signal Processing VII, Vol3, pp 1780-1783, 1994.
2 D.L. Donoho, "De-noising by soft-thresholding", IEEE Trans. on Information Theory, Vol 41, No. 3, pp. 613-627, 1995.
3 M. Holschneider, R. Kronland-Martinet, and J. Morlet, "A real-time algorithm for signal analysis with the help of the wavelet transform", in *Wavelets, Time-Frequency Methods and Phase Space*, Berlin: springer, IPTI, pp. 286-297, 1989.
4 R.P. Kleihorst, R.L. Lagendijk, and J. Biemond, "An Adaptive Order-Statistic Noise Filter for Gamma-Corrected Image Sequences", IEEE Trans. on Image Processing, Vol. 6, No. 10, 1997.
5 P.M.B. van Roosmalen, S.J.P. Westen, R.L. Lagendijk and J. Biemond, "Noise Reduction for Image Sequences using an Oriented Pyramid Thresholding technique", Proc. of ICIP-96, Vol. I, pp. 375-378, Lausanne, Switzerland, IEEE, 1996.
6 A.C. Kokaram, "Detection and Removal of Line Scratches in degraded Motion Picture Sequences (Signal Processing VIII, Vol 1, pp 5-8, 1996.
7 A.C. Kokaram, R.D. Morris, W.J. Fitzgerald, P.J.W. Rayner, "Interpolation of Missing Data in Image Sequences", IEEE Trans. on Image Processing, pp. 1509-1519, Vol. 4, No. 11, 1995.
8 B.K. Natarajan, "Filtering Random Noise from deterministic Signals via Data Compression", IEEE Trans. on Signal Processing, Vol. 43, No. 11, 1995.
9 D.L. Donoho, and I.M. Johnstone, "Ideal Denoising in an orthonormal basis chosen from a library of bases", Technical Report 461, 1994.
10 E.P. Simoncelli, and E.H. Adelson, "Noise Removal via Bayesian Wavelet Coring", Proc. of ICIP-96, Vol. I, pp. 379-382, Lausanne, Switzerland, IEEE, 1996.
11 J. Allnat. "Transmitted-picture Assessment", John Wiley & Sons, -Chichester, 1983