A Robust Codec for Transmission of Very Low Bit-rate Video over Channels with Bursty Errors

W. S. Lee, M. R. Frater, M. R. Pickering and J. F. Arnold

***Abstract*— We describe a robust codec for the transmission of very low bit-rate video over channels with a variety of errors including random and bursty bit errors and packet loss. The codec exploits adaptivity to give good performance with a low overhead. By only protecting macroblocks which would otherwise be poorly concealed by the decoder, the codec allows adaptive selection of the parts of video to protect. For protection, it uses multiple description codes which indirectly provide frequency based adaptivity by protecting the more significant DCT coefficients. Simulations show significant improvements in the performance of the codec when compared to codecs which use intra macroblock updating (raster scan and random) at the same overhead. The codec is efficient in its use of bits and has good error resilience properties both objectively and subjectively over a wide range of conditions. Further, transcoding of the received bit-stream to the standard H.263 syntax is relatively easy.**

I. Introduction

A realistic video transmission system can be represented by the model shown in Figure 1.

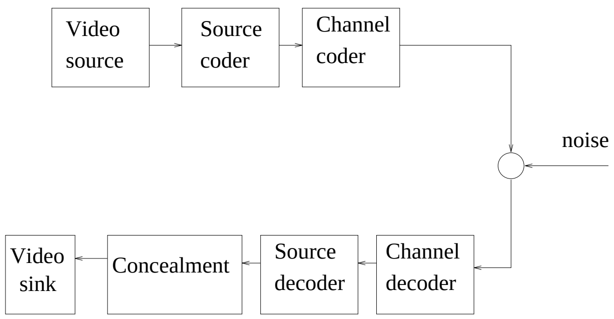

Fig. 1. Model of a video transmission system.

The source coder provides compression (usually lossy) of the video while the channel coder introduces redundancy in order to combat errors introduced by a noisy channel. The concealment stage is a post-processing stage (usually found only in lossy compression systems such as video) which is useful for reducing the effects of residual channel errors. In the concealment stage, operations such as spatial or temporal filtering are carried out to improve the quality of corrupted video (see e.g. [9]).

In this paper, we describe a codec which is suitable for use over channels with bursty errors. Such channels include wireless channels which suffer from fading and ATM networks which suffer from cell loss. We have concentrated on designing a codec for very low bit-rate channels which suffer from bursty bit errors.

Our approach is to do adaptive joint source-channel coding that takes into account the concealment strategy used by the decoder. Joint source-channel coding is achieved by protecting parts of the bit-stream which will result in heavy distortion if corrupted by errors and leaving the rest of the bit-stream unprotected. Such an approach is more efficient than uniformly protecting the whole bit-stream as is usually the case for channel coding.

Our video codec is modified from the H.263 video coding standard [3]. It uses block-based motion compensation for prediction and the Discrete Cosine Transform (DCT) to code the prediction error. One major difference between our codec and the H.263 codec is that it assumes that the encoder has knowledge of the concealment strategy used by the decoder. As such, it is able to make use of this knowledge to protect only those parts (macroblocks in H.263 terms) of the video which the decoder is unable to conceal effectively.

Various other schemes for unequal error protection have been proposed in the literature. The MPEG-2 international standard [2] provides various methods for producing bit-streams with dual priority (levels of protection) including data partitioning, signal-to-noise ratio (SNR) scalability and spatial scalability. A comparison of these methods can be found in [1].

For error protection, our codec uses multiple description codes [7; 8]. In multiple description codes, it is assumed that two independent channels (channel 1 and channel 2) connecting the sender and the receiver are available. Each channel may be in a working or non-working state. This information is available to the decoder (possibly through error detection) but not to the encoder. The encoder must code the information in such a way that a good quality signal is obtained by the receiver if both channels are working and an acceptable quality signal is obtained if either of the channels is not working.

Multiple description codes have been used in image coding and some preliminary work was also done for video coding [8]. Our implementation shows that multiple description codes can be profitably used for very low bit-rate video coding with relatively little modification to the current international standard.

The multiple description scalar quantizer which we use

This work was supported by the Australian Research Council. The authors are with the School of Electrical Engineering, University College, The University of New South Wales, Australian Defence Force Academy, Canberra ACT 2600, Australia.

The material in this paper was presented in part at the International Conference on Image Processing, Santa Barbara, October 26-29, 1997.

indirectly protects only the larger DCT coefficients, giving some frequency based adaptivity which again helps the efficiency of the codec. In our codec, we approximate the two independent channels by separating the information temporally i.e. all the information (in one frame) for one channel is sent first followed by all the information for the other channel. Provided the burst length of the error is less than the number of bits (in one frame) of one channel, this approximation should be reasonably good.

In our simulations, we compare the performance of our codec against codecs which use intra macroblock updates, for a wide range of error conditions including bursty bit errors, random cell loss and random errors. Intra macroblock updates is a method (also described in H.263) commonly used to reduce the effect of temporal error propagation (propagation of errors due to motion compensated prediction). We find that for fairly low overheads (less than 15%) our codec significantly outperforms codecs which use intra macroblock updates at the same overhead.

The main contributions of this work are (i) showing that by protecting only poorly concealed macroblocks, significant gains in efficiency can be obtained for protection of very low bit-rate video sequences and (ii) showing that a combination of unequal error protection and multiple description codes can outperform a standard technique (intra macroblock updates) for error resilience of very low bit-rate video coding for codecs similar to that used in an existing international standard (H.263).

In Section II, we give a detailed description of the concealment based unequal error protection scheme and multiple description codes. The experimental procedures are described in Section III. In Section IV, we give simulation results comparing the new scheme with intra-coded macroblock updating for a wide range of error conditions. The simulations show significant gains in performance of our codec compared to intra macroblock updating schemes. We give conclusions in Section V

II. Detailed System Description

Our codec is based on the H.263 standard [3] for low bit rate video coding. Each frame is divided into 16-pixel by 16-pixel macroblocks, each of which is further divided into four 8 by 8 blocks. The macroblocks are grouped into group of blocks (GOBs). Instead of having fixed sized GOBs as in H.263, we use an adaptive GOB size, where a GOB consist of all macroblocks for which coding is started within 256 bits of the start of the GOB. The GOB header consists of a synchronization marker and has the same information as in H.263 except for an additional 2 bits which describes whether the GOB comes from channel 1 or channel 2 and whether it is an intra or inter frame. This extra information helps the decoder to place the GOB correctly in the video sequence and use the appropriate variable length code table for decoding the GOB. The synchronization marker used is also different from the one used in H.263. It consists of parity bits which are used for both synchronization purposes and for detecting errors [5]. The remaining syntax is the same as that used in H.263 except where indicated. Fixed quantization parameters (variable bit-rate) are used throughout to simplify our experiments.

A. Concealment-based Unequal Error Protection

For each frame, the information is sent in two separate (virtual) channels. Information from all the macroblocks is sent in channel 1. Most of the macroblocks on channel 1 are the usual H.263 macroblocks with the exception of the protected macroblocks. Information from protected (poorly concealed) macroblocks is sent on both channel 1 and channel 2 using multiple description codes. Information from both channels are combined before the frame is used for prediction of the next frame.

Multiple Description Codes

Macroblocks which would be poorly concealed if lost are protected using multiple description codes. For the motion vectors in inter mode and the DC coefficients in intra mode, we simply repeat the information in both channels. This means that twice the amount of information needs to be sent but the exact same information is recovered even if one channel is unavailable. The rest of the DCT coefficients are quantized using a multiple description scalar quantizer similar to those used in [7; 8]. This allows less than twice the amount of information to be sent. However, if one of the channels is unavailable a poorer reconstruction results. The quantization table for the multiple description scalar quantizer which we used is shown in Table 1.

	·	·	**-3**	**-2**	**-1**	**0**	**1**	**2**	**3**	·	·
·	·	·									
·		·	·								
-3			-6	-5							
-2				-4	-3						
-1					-2	-1					
0						0					
1						1	2				
2							3	4			
3								5	6		
·									·	·	
·										·	·

Table 1. Table for multiple description scalar quantizer. The index values that are sent are represented in bold with those on the rows sent on channel 1 and those on the columns sent on channel 2.

To use the quantization table, the index on the row containing the quantized value is sent through channel 1 and the index on the column containing the quantized value is sent through channel 2. If both values are received, the decoder can reconstruct the original quantized value by looking up the appropriate row and column of the table. If only the row index is received, the average value over the indexed row is used as the reconstructed value. Similarly, if only the column index is received, the average value over the indexed column is used as the reconstructed value.

The quantization table shown in Table 1 is slightly different from those in [7; 8] in that it is symmetric about zero.

This introduces a slight imbalance that results in channel 1 being slightly more accurate and costly to send. However, the resulting probability distributions of both channels are symmetric about zero. This is useful for our simulations because we can use the variable length codes used in H.263 which assumes a symmetric distribution about zero. (It is possible to redesign the variable length codes but it makes comparison with a H.263 like scheme, which is designed to operate over a wide range of quantization parameters, difficult. The resulting bit-rates are also more difficult to interpret since the codes are not trained on the same sequences and quantization parameter values as H.263.) Since the H.263 codec is designed to work over a wide range of quantization parameters, we expect its variable length codes to perform reasonably for our system. Of course, optimizing the variable length codes can bring further improvements to the codec.

H.263 is designed so that the quantization bin containing zero is approximately twice as large as the other quantization bins. From Table 1, we can see that this means that the outputs for channel 2 are approximately what would be output by H.263 at twice the quantization parameter. The output for channel 1 is similar except that the quantization bin containing zero is the same size as all the other quantization bins (since the original H.263 bin is approximately twice the other bins). In the absence of one of the channels, the resulting decoded macroblock codes are similar to that produced by a coder with twice the quantization parameter. This effect is similar to that produced using SNR scalability when the enhancement information is lost.

Deciding Protected Macroblocks

A macroblock is protected only if the result of concealment would be significantly poorer than the result of protection using multiple description codes, assuming that channel 1 is lost.

To decide whether to protect a macroblock, a motion compensated prediction for the macroblock is first determined from previous macroblocks in channel 2 using the same prediction scheme as that used in H.263 (median of the previous motion vector, the above motion vector and the above right motion vector). Note that the previous macroblocks are all previous macroblocks from channel 2, hence are independent of those used in channel 1. Call this predicted macroblock M. This has the effect of simulating the concealment action of the decoder for the case where the macroblock has been lost. We would like to compare this macroblock against a macroblock which has been protected using multiple description codes, with the aim of protecting the macroblock only if the concealment performs significantly poorer than protecting the macroblock.

For the comparison, the macroblock is coded in protected mode as described in Section II.A and then decoded assuming that channel 1 has been lost. Call this macroblock N. If the mean squared error (compared to the uncoded original) of the macroblock M is more than k times the mean squared error of the macroblock N, we mark the macroblock to be protected. If the macroblock is not protected, it is sent on channel 1 as a normal H.263 macroblock. We also protected all macroblocks which are coded in intra mode, where the decision to code in intra mode is made according to the heuristic algorithm used in video coding test model TMN5 [4].

To approximate the independence between the two channels by using a multiplexed single channel, information from each channel is sent contiguously before information from the other channel is sent.

Protected macroblocks need to be decoded using multiple description codes. For channel 1, the protected macroblocks are indicated for the benefit of the decoder using the macroblock type. In this work, we do not use the advanced prediction mode of H.263 and use this mode to indicate a protected macroblock instead. Intra coded macroblocks are also known to the decoder to be protected. For very low bit-rate video coding, the protected macroblocks tend to be very sparsely distributed. To improve the efficiency of sending the macroblocks of channel 2, a run-length code is used to indicate the presence of the macroblock to be coded. It is possible to design efficient run length codes if the statistics of the video sequences are known. However, to avoid training the codes on the sequences used for simulation, we use a simple scheme for run-length coding the presence of coded macroblocks.

In this scheme each uncoded macroblock is indicated using one bit (zero) if the total number of bits used this way (plus one bit to indicate the end of uncoded macroblocks) is smaller than the number of bits needed to code the run-length of uncoded macroblocks (approximately the log of the remaining macroblocks in the frame). Otherwise, the length of the skipped macroblocks is coded. One extra bit is sent before sending the run-length to inform the decoder which of the two methods is used.

At the decoder, information from both channels is used if both channels are available. If only one channel is available, the available channel is used. If both channels are not available, motion compensated concealment is done using the same prediction scheme as that used in H.263 (median of the previous motion vector, the above motion vector and the above right motion vector) where missing motion vectors are substituted with estimated motion vectors. Error detection is used to indicate the availability of the channels. This is further described in Section III.B.

III. Experimental Procedure

We compare the concealment-based unequal error protection codec with codecs which use intra macroblock updates, where the rate of intra macroblock updates is adjusted to give the same number of bits as the concealment-based unequal error protection technique for each video sequence.

A. Codecs with intra macroblock update

For comparison with the concealment-based unequal error protection codec, we use two codecs which use intra updating but are otherwise identical to our codec. One codec uses a raster scan intra macroblock update while the other codec uses a random intra macroblock update. The two codecs are identical except for the way that they do macroblock updating. The raster scan codec does intra macroblock updating in raster scan order. For the random updating codec, a random permutation of the ordering of the macroblocks is first generated. The intra macroblock updates are then generated using this permutation of the macroblock ordering for all video sequences. The random intra macroblock update scheme is used in practical systems to avoid the visually displeasing 'windscreen wiper' effect present in the raster scan updating.

The syntax used is very similar to H.263 except that it uses adaptive GOB size, the parity resynchronization scheme (described in section III.B) and it has an additional bit in the GOB header describing whether the frame is an intra or inter frame. For concealment, the codecs use motion compensated concealment using the same prediction scheme as that used in H.263 where missing motion vectors are substituted with estimated motion vectors.

B. Error Detection

Our video codec, like H.263, uses variable length codes. Hence errors in the bit-stream can cause loss of synchronization because the decoder is unable to locate the start of the next codeword. Like H.263, our coded bit-stream contains resynchronization markers which allow the decoder to regain synchronization at the appropriate place in the bit-stream. A crucial part of a robust video codec is the action of the decoder when an error in the bit-stream is detected. A simple method would be to abandon decoding at that point and restart at the next resynchronization marker. Because errors in the bit-stream are usually not detected immediately, this often results in highly perceivable artifacts appearing in the video which results in poor objective and subjective quality [5]. We take the approach of assuming that all the macroblocks between the two resynchronization points are corrupted when an error is detected between the two resynchronization points. This approaches eliminates virtually all the obnoxious artifacts caused by undetected errors at the expense of losing some useful information. The resulting bitstream is similar to that produced by an erasure channel (e.g. cell loss in ATM networks).

Methods we used for error detection include detecting invalid variable length codes, detecting incorrect macroblock numbers at resynchronization points and using parity bits as a resynchronization marker [5]. For the concealment-based codec, an additional method of error detection (which we did not implement) is possible by checking that motion vectors and DC coefficients which are transmitted in both channels agree.

Resynchronization

The same resynchronization strategy is used by all our codecs (concealment-based unequal error protection and intra update codecs). After an error is detected, the decoder moves through the bit-stream and searches for a resynchronization marker at every possible resynchronization point (i.e. at every byte boundary). The use of parity bits as a resynchronization marker (or the presence of errors in normal markers) means that it is possible that a false resynchronization marker will be found. To mitigate the effect of false resynchronization, we save the state of the decoder at the resynchronization point before we continue decoding. If no error is detected up to the next resynchronization point, we assume that the resynchronization was correct and continue decoding, otherwise we restore the state of the decoder and continue searching for a resynchronization marker.

IV. Simulation Results

For simulation, we use the following four well known test sequence at QCIF resolution: Mother and Daughter, Hall Monitor, Container and Akiyo. The sequences have low spatial detail and low amounts of movement. They are typical of sequences coded at very low bit rates. The sequences are coded at 10 frames per second with a fixed quantization step-size giving 100 frames of coded video for each sequence. The first frames for all the sequences are coded intra without any protection (we do not corrupt the first frame with errors in our simulations). For the concealment-based protection scheme we set the value k defined in Section II.A to be 2 (recall that a macroblock is protected only if the resulting mean square error of the protected macroblock assuming channel 1 is lost is more than k times that of an unprotected concealed macroblock). The encoding results showing number of bits used and the PSNRs (without errors) are shown in Table 2.

The overhead used for increasing the robustness can be used for improving the quality of the video if robustness is not required. Since we are coding using fixed quantization levels, it is not possible show the exact PSNR improvements by matching the number of bits used in the robust sequences. Instead we show the PSNR of the unprotected sequences at the next quantization parameter level in Table 3 to give an indication of the possible improvements in the PSNR if the overhead were used to code the video instead of improving robustness.

Sequence	QP	Total Bits	PSNR
Mother Daughter	9	243400	34.57
Hall Monitor	11	227504	32.62
Container	10	254792	32.32
Akiyo	7	242488	36.60

Table 3. Total bits and PSNR (luminance) of the unprotected sequences at the next level of the quantization parameter. This gives an indication of the potential gain in PSNR if the overhead were used for improving the quality instead of for error robustness.

In most video communications systems, a packet-based multiplex is used to permit the audio and video data to be carried on a single transmission channel. Consequently any errors introduced in the transmission channel pass through the demultiplexer before entering the video decoder. Three

Sequence	Coding Scheme	Total Bits	Ch. 1 Bits	Overhead	PSNR (dB)
Mother Daughter	unprotected	203720		0%	33.96
	concealment-based	232016	201571	13.9%	33.96
	intra raster	231984		13.9%	34.02
	intra random	232168		14.0%	34.00
	all protected	258560	197085	26.9%	33.96
Hall Monitor	unprotected	200592		0%	32.02
	concealment-based	231648	187582	15.5%	32.02
	intra raster	231656		15.5%	32.10
	intra random	232120		15.7%	32.08
	all protected	247072	181126	23.2%	32.02
Container	unprotected	230416		0%	31.84
	concealment-based	238792	231204	3.6%	31.84
	intra raster	239168		3.8%	31.91
	intra random	239416		3.9%	31.89
	all protected	260800	225581	13.2%	31.84
Akiyo	unprotected	196768		0%	35.63
	concealment-based	213480	194050	8.5%	35.63
	intra raster	213992		8.7%	35.65
	intra random	214550		9.0%	35.68
	all protected	228808	189787	16.3%	35.63

Table 2. Total bits and PSNR (luminance) without errors for coded sequences. The PSNR should be almost the same for the different coding methods since each sequence because the same quantization parameter is used (the difference being caused by the intra coded macroblocks which has slightly better PSNR).

classes of error pattern have been chosen as typical of the types of errors passed to the video decoder: bursty bit errors, random (independent) bit errors, and packet loss. These same patterns are currently being used for evaluating error resilience techniques in the MPEG 4 standard.

A. Gilbert Channel

To test the codec, we performed experiments with a two state Gilbert channel with no error in one state and bit error rate of 0.2 in the other state. The average burst length was set at 240 bits and the average bit error rate set at 1×10^{-3}, 3×10^{-3}, 5×10^{-3}, 7×10^{-3} and 1×10^{-2}. The first frames of all the sequences are left uncorrupted by errors to help concealment. The plots of the PSNR vs bit error rates averaged over 20 runs for the concealment-based unequal error protected codec and the intra codecs are shown in Figure 2.

We also conducted hypothesis testing at each of the bit error rates. Let μ_0 be the average PSNR of the intra updating scheme (either random or raster scan, whichever has higher average PSNR) at the particular bit error rate being tested. Let μ_1 be the average PSNR of the concealment-based protection scheme at the particular bit error rate. Then the null hypothesis used is $H_0 : \mu_0 = \mu_1$ and the alternative hypothesis is $H_1 : \mu_1 > \mu_0$. Since the PSNRs of each run for the different schemes are not independent (the same seed is used for the random number generator for each run), we used a t-test for paired observations with 19 degrees of freedom (for 20 runs) at 0.05 level significance. Rejection of the null hypothesis means that the sample evidence refutes the hypothesis that the average PSNR of the new scheme is no different (not better) than the other schemes. The results are also shown in Figure 2 where circles on the solid curves indicate that for these points we did not manage to reject the null hypothesis.

The plots show significantly better performances for the concealment-based unequal error protection scheme compared to both the intra updating schemes for all the sequences except the container sequence where the performance is comparable. This result for the container sequence is mainly because (i) concealment is quite effective for the container sequence and (ii) quantizing the DCT coefficients at twice the resolution does not leave many nonzero coefficients. This is reflected in the low overhead (3.6%) used by the concealment-based protection scheme for this sequence.

We also ran experiments comparing the performance of the codecs with the average burst length of errors set at 24, 120, 240 and 1200 and the bit error rate fixed at 5×10^{-3}. With the bit error rate fixed, the frequency of bursts increases as the average burst length decreases. The PSNR results averaged over 20 runs are shown in Figure 3. The same hypothesis tests are also done.

The results show that the PSNR decreases as the average burst length decreases. This is because the higher frequency of bursts cause a much higher rate of information loss as at least a whole packet of information (between two resynchronization points) has to be discarded for each burst. Comparing Figures 2 and 3, it would appear that the main factor determining the performance of our two codecs is the frequency of error bursts.

The performance of the codecs are closer at lower burst lengths. This is mainly because undetected errors or errors caused by both channels of the concealment-based unequal

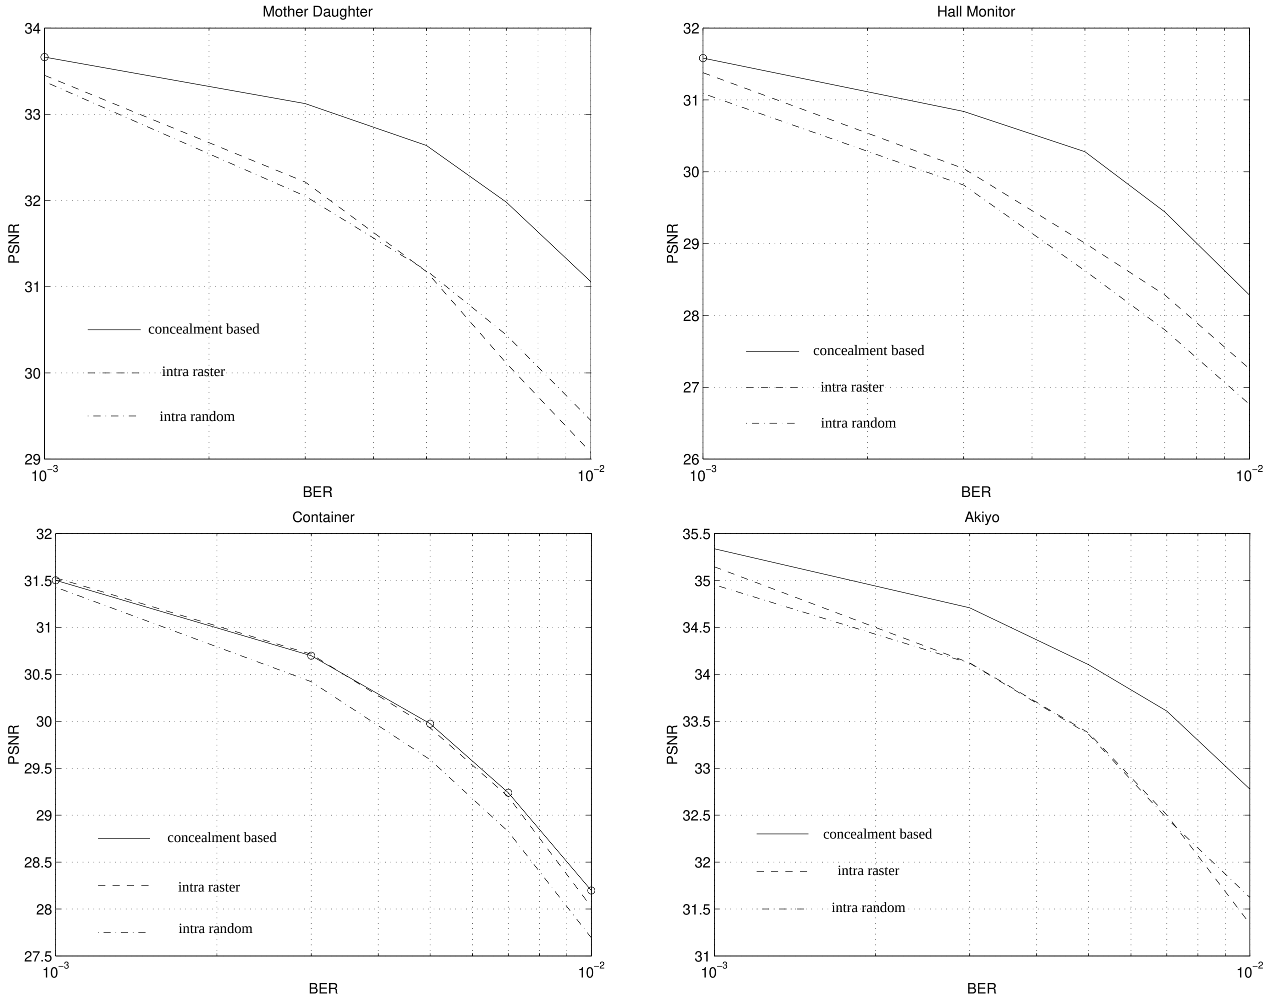

Fig. 2. Plots of PSNR vs BER for Gilbert channel. Circles on the solid curves indicate that for these points, we did not manage to reject the null hypothesis, H_0: the average PSNR of the intra macroblock updating scheme (raster scan or random whichever has higher average PSNR at that point) is the same as the average PSNR of the concealment-based protection scheme, at 0.05 level significance.

error protection scheme being unavailable tend to remain unless the corrupted macroblocks are intra updated. The more frequent error bursts cause more of these errors, decreasing the performance of the concealment-based protection scheme. At very long burst lengths, the performance of the codecs are also very similar, possibly because (i) the independence assumption between the two channels of the concealment-based codec is starting to fail and (ii) with very few bursts, there is little information loss, and hence little to test the performance difference of the codecs.

Subjectively, in the concealment-based unequal error protection codec, the errors usually take the form of mild blurring while in the intra updating codecs, the errors are usually macroblocks which are not properly updated. Examples of errors when using the intra updating schemes can be seen in Figure 4. These errors are visually more annoying than slight blurring and also propagate more temporally.

B. Packet Loss

We conducted a number of experiments to observe the performance of the codecs in a packet loss situation. For simplicity, we assume that when a packet is lost, the network simply replaces the lost bits with zeros. The decoders still have the job of detecting the packet loss. We ran simulations with random packet losses with a packet size of 256 bits at packet loss rates 0.001, 0.005, 0.01, 0.05 and 0.1. The results for the Mother and Daughter and Hall Monitor sequences, averaged over 20 runs for the concealment-based protection scheme and the raster scan intra update scheme are shown in Figure 5. These results show performance improvement for the new scheme similar to the Gilbert channel case. The results for the other two sequences are also similar to the Gilbert channel case showing that the concealment-based unequal error protection scheme outperforms the intra macroblock update scheme for all sequences except the container sequence.

C. Random errors

We ran some simulations on the codecs for the binary symmetric channels to observe the performance of the codecs without additional forward error correction codes. The values of p are set to 1×10^{-4}, 3×10^{-4}, 5×10^{-4}, 7×10^{-4} and 1×10^{-3}. Using the sequences with concealment-based unequal error protection and raster scan intra updates from

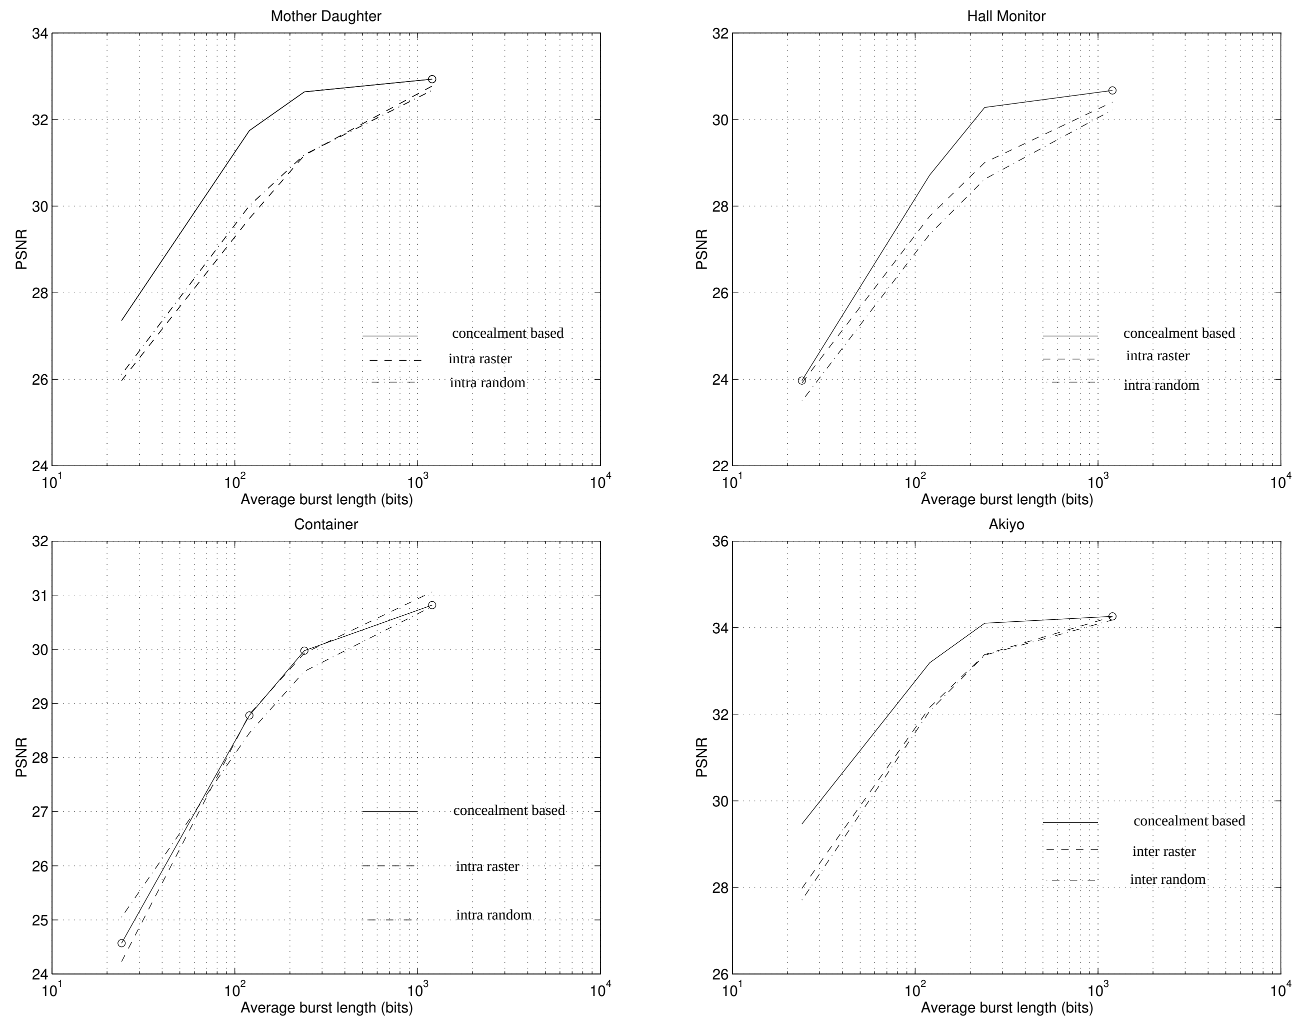

Fig. 3. Plots of PSNR vs average burst length at BER of 5×10^{-3} for Gilbert channel.

Table 2, the PSNR's averaged over 20 sequences are shown for the Mother and Daughter and Container sequences in Figure 6. The results shows that the concealment-based protection scheme maintains its edge over the raster scan intra updating scheme in the random error case. The results for the other two sequences are similar to the Gilbert channel case.

D. Scaling behaviour

In this section, we investigate the effect of reducing the overhead used for protecting the video. For the concealment-based unequal error protection scheme, two possible ways to reduce the overhead are (i) to use a different multiple description code which has poorer quality when one of the channels is lost and (ii) increasing the value of k.

We investigate the second option of increasing k and use a k value of 3. Again, we compare the performance of the scheme with that of intra macroblock updating (raster scan). The encoding results are shown in Table 4. In these sequences, the intra updating rate is chosen so that the number of bits approximately matches that of the concealment-based unequal error protection scheme.

We show the simulation results for the Gilbert channel with no error in the good state and bit error rate of 0.2 in the bad state. The average burst length was set at 240 bits and the average bit error rate set at 1×10^{-3}, 3×10^{-3}, 5×10^{-3}, 7×10^{-3} and 1×10^{-2}. The plots of the PSNR vs bit error rates for the Mother and Daughter and the Hall Monitor sequences averaged over 20 runs are shown in Figure 7 together with the PSNR of the sequences with higher overhead. The results for random error and packet loss channels are similar to the results for the Gilbert channel.

As expected, for the concealment-based unequal error protection scheme, the sequences with the higher overhead perform better than the sequences with the lower overhead. The performance of the concealment-based unequal error protection scheme at lower overhead is also better than the performance of the raster scan intra updating scheme at higher overhead for all sequences except the container sequence. Surprisingly, for the raster scan intra update scheme, increasing the overhead does not improve the performance at some bit-rates for some of the sequences. One likely explanation for this is that the effects of the additional errors suffered by these sequences (since a higher number of bits for a given frame implies a higher num-

Fig. 4. Typical errors of the intra updating scheme. The frames on the left are the ones produced by the concealment-based codec while the frames on the right are the ones produced by the raster scan intra updating codec. Note the errors in the intra updated frames. For the Mother Daughter frame, part of the hand was not updated correctly. For the Hall Monitor frame, part of the person was not updated correctly.

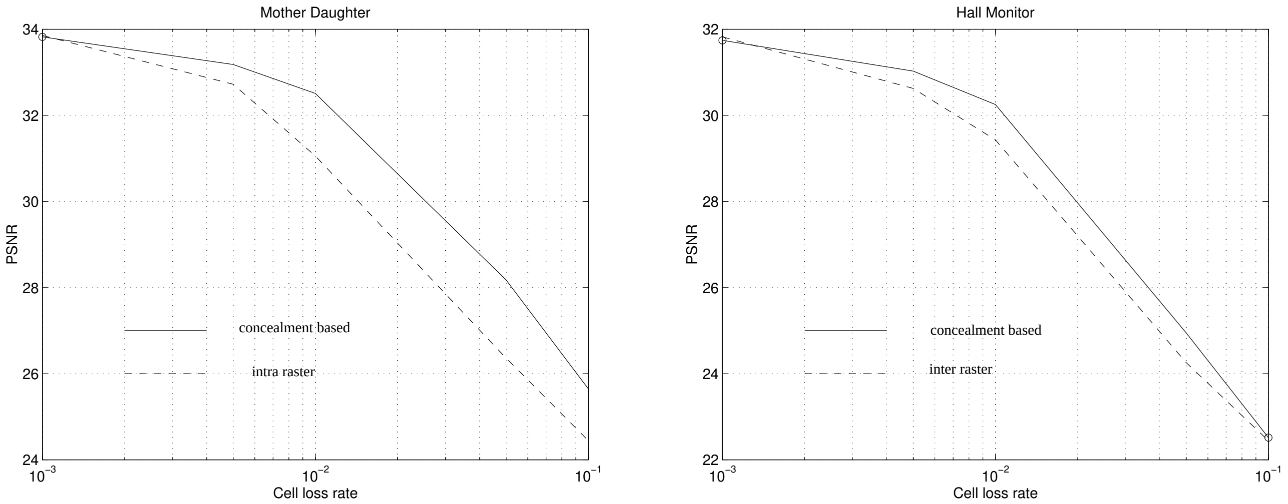

Fig. 5. Plots of PSNR vs packet loss rate for packet loss.

ber of errors) outweighs the benefits of the additional intra macroblocks.

E. Discussion

Compared to the intra updating codecs, the concealment-based unequal error protection codec is not as robust against catastrophic errors such as losing several frames at the same time. The codec will also not correct undetected

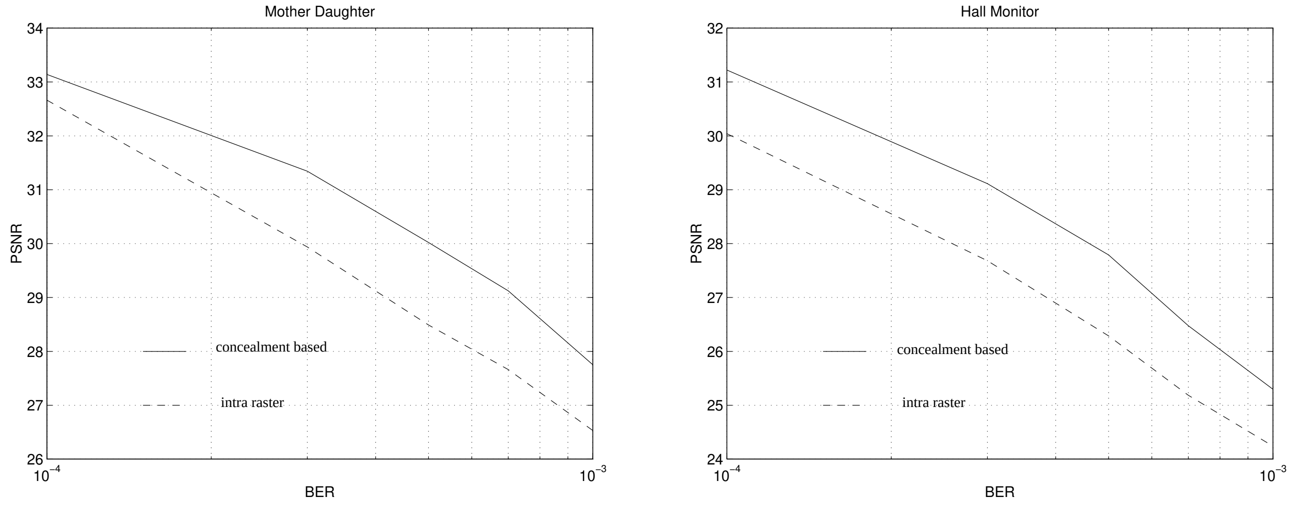

Fig. 6. Plots of PSNR vs BER for random error channel.

Sequence	Coding Scheme	Total Bits	Overhead	PSNR (dB)
Mother Daughter	concealment-based	222704	9.3%	33.96
	intra raster	222704	9.3%	34.00
Hall Monitor	concealment-based	223000	11.2%	32.02
	intra raster	222640	11.0%	32.07
Container	concealment-based	234168	1.6%	31.84
	intra raster	235032	2.0%	31.87
Akiyo	concealment-based	207560	5.5%	35.63
	intra raster	207632	5.5%	35.64

Table 4. Total bits and PSNR (luminance) without errors for coded sequences with lower overheads.

errors or errors caused by losing both channels simultaneously. Loss of quality due to the loss of one channel is also not corrected. However, the overhead required for the protection is relatively low. It should therefore be possible to combine the protection scheme with intra macroblock updating in useful ways.

We have not compared the scheme against more intelligent macroblock updating schemes such as those found in [6]. While these schemes perform better than simple raster scan or random updating, we expect that a lot of the annoying artifacts from the raster scan and random macroblock updating schemes will remain. This is because macroblock updating schemes will only correct errors after it happens unlike the concealment-based scheme that we use. The scaling experiments also suggest an inherent flaw in intra updates - the fact that a macroblock transmitted in intra mode is more likely to be lost due to errors.

The proposed codec is computationally slightly more demanding than a standard H.263 codec. In the encoder, an additional inverse DCT and two additional calculations of the mean squared errors of each block is needed compared to a standard H.263 encoder. The complexity of the decoder is similar to that of a standard H.263 decoder.

One advantage of the codec as described is that transcoding of its bit-stream into a standard H.263 bit-stream can be done easily (without doing the inverse DCT). All that needs to be done is to combine the multiple description codes in the two channels for the protected macroblocks and then recode the macroblock in the H.263 syntax.

V. Conclusions

We have designed a robust codec for low bit-rate channels with bursty bit errors. The codec exploits adaptivity to give good performance with a low overhead. It achieves this by protecting only macroblocks which would be poorly concealed by the decoder. To achieve good performance over channels with bursty errors, it uses multiple description codes. Simulation results show that the codec gives significant improvements in performance compared to intra macroblock updating (raster scan and random) over a wide range of conditions (bursty bit errors, packet losses and random errors). Transcoding of a received bit-stream produced by our codec to the standard H.263 syntax can also be achieved with relative ease and without the need for full decoding.

References

[1] R. Aravind, M. R. Civanlar, and A. R. Reibman. Packet loss resilience of MPEG-2 scalable video coding algorithms. *IEEE Trans. Circuit and Systems for Video Technology*, 6(5):426–435, October 1996.

[2] ISO/IEC International Standard 13818-2. *Generic Coding of moving pictures and associated audio information: Video.* 1995.

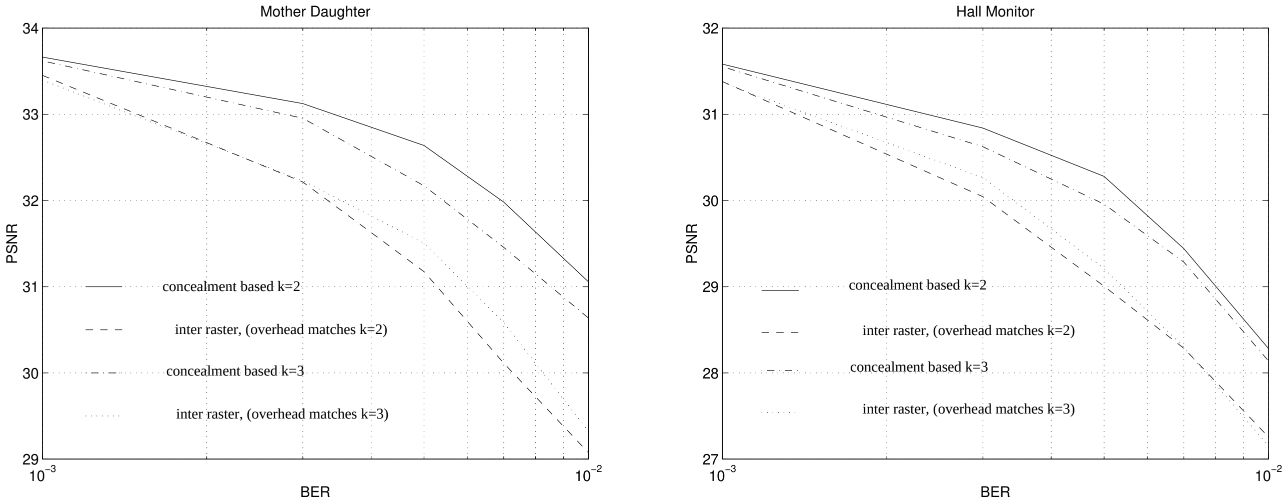

Fig. 7. Plots of PSNR vs BER for Gilbert channel showing effects of reducing protection.

[3] ITU-T Recommendation H.263. *Video Coding for Low Bitrate Communication.*

[4] ITU Telecommunication Standardization Sector LBC-95. Video codec test model, TMN5.

[5] W. S. Lee, M. R. Pickering, M. R. Frater, and J. F. Arnold. Error resilience in video and multiplexing layers for very low bit-rate video coding systems. *IEEE Journal on Selected Areas in Communications*, 1997. Accepted for publication.

[6] J. Y. Liao and J. D. Villasenor. Adaptive intra update for video coding over noisy channel. In *Proc. Int. Conf. on Image Processing*, pages 763–766, 1996.

[7] V. A. Vaishampayan. Design of multiple description scalar quantizers. *IEEE Trans. on Information Theory*, 39(3):821–834, May 1993.

[8] V. A. Vaishampayan. Application of multiple description codes to images and video transmission over lossy networks. In *7th International Workshop on Packet Video*, pages 55–60, 1996.

[9] Q. F. Zhu, Y. Wang, and L. Shaw. Coding and cell-loss recovery in DCT-based packet video. *IEEE Trans. Circuit and Systems for Video Technology*, 3(3):248–258, June 1993.

Wee Sun Lee (S'91-M'96) received the B.E. degree in Electrical Engineering from the University of Queensland, Australia, in 1992 and the Ph.D. degree in Engineering from the Australian National University in 1996. He is currently a Research Fellow in the School of Electrical Engineering of the Australian Defence Force Academy. His research interests include error resilience for video and image coding, image and video compression and machine learning.

Michael R. Frater (S'89-M'92) was born in Sydney, Australia, in 1965. He received the B.Sc. degree in mathematics and physics and the B.E. (Hons.) degree in electrical engineering from the University of Sydney, Australia, in 1986 and 1988 respectively. He received the Ph.D. degree from the Department of Systems Engineering at the Australian National University, Australia, in 1991.

He is a Senior Lecturer in Electrical Engineering at the University College, the University of New South Wales, Australian Defence Force Academy, Australia. His research interests include aspects of the coding and transmission of video services, teletraffic and queueing theory, stochastic processes, and signal processing.

Mark R. Pickering (S'91-M'95) received a B.E. degree in 1988 from Capricornia Institute of Advanced Education, Rockhampton, Australia, and M.E. and Ph.D. degrees in 1992 and 1995 respectively, from The University of New South Wales, Sydney, Australia. He is currently a Lecturer in the School of Electrical Engineering at the Australian Defence Force Academy in Canberra. His research interests include the application of visual perception to image sequence coding and techniques for rate control, error concealment and motion estimation in video coding.

John F. Arnold (S'77-M'85-SM'96) received the B.E. and M.Eng.Sc. degrees from the University of Melbourne, Australia in 1976 and 1979, respectively, and the Ph.D. degree from the University of New South Wales, Australia, in 1984.

Since 1978, he has been with the Department of Electrical Engineering, University of New South Wales, initially at the Royal Military College, Duntroon and more recently at the Australian Defence Force Academy, Canberra where he currently holds the position of Professor of Electrical Engineering. His research interests lie in the fields of video coding, asychronous transmission mode (ATM) networks, and coding of remotely sensed data.

List of Figures