

Systematic Lossy Error Protection based on H.264/AVC Redundant Slices and Flexible Macroblock Ordering

Pierpaolo Baccichet*, Shantanu Rane, and Bernd Girod†

Abstract

In this paper we propose a scheme for Systematic Lossy Error Protection (SLEP) of an H.264/AVC compressed video bit-stream, using standard compatible features such as redundant slices, and flexible macroblock ordering. In this scheme, the systematic transmission consists of a conventional H.264/AVC bit-stream. For error resilience, an additional bit-stream known as the Wyner-Ziv bit-stream is also transmitted. The Wyner-Ziv bit-stream allows the decoding of a coarsely quantized description of the original video signal, and is efficiently generated by using H.264/AVC redundant slices in conjunction with Reed-Solomon coding. The Wyner-Ziv bit-stream is decoded in order to recover the redundant video descriptions, which are used in lieu of portions lost from the original video signal due to channel errors. SLEP allows the video quality to degrade gracefully with worsening channel conditions, and provides a flexible trade-off between the achieved error resilience and the coarseness of the redundant description. The performance can be improved, especially for low motion video sequences, by applying SLEP to a region-of-interest in the video frame, using Flexible Macroblock Ordering (FMO). We provide experimental results for two video transmission scenarios, which demonstrate the advantages of SLEP over FEC as an error resilience scheme.

Keywords: Wyner-Ziv coding, distributed video coding, side information, systematic source-channel coding, redundant slices, flexible macroblock ordering.

1. Introduction

In typical video transmission systems, a video signal is compressed, and the resulting bit-stream is transmitted over an error-prone channel. The errors may consist of symbol errors caused by fading, as observed for wireless channels, or packet erasures caused by congestion, as observed in the Internet. If a received video packet contains errors, then the portion of the video signal contained in the packet is lost and must be concealed. Error concealment schemes alone cannot guarantee acceptable video quality at the error probabilities encountered in video transmission systems. Hence, a combination of application layer forward error correction (FEC) and feedback (if permissible for the given application) is used to protect the video packets, in exchange for an increase in

latency and a small expansion of the transmitted bit-rate. For e.g., MPEG-2 transport uses a (204,188) Reed-Solomon (RS) code to protect the packets in the broadcast stream. A conventional (n, k) RS code can be used to provide FEC as long as the number of symbol erasure is less than $n - k$. However when the error probability increases, the number of erroneous or erased packets overwhelms the error correction code and this results in rapid degradation of the picture quality, a phenomenon commonly referred to as the “cliff” effect.

In our recent work [1–3], we have proposed an error-resilient scheme called Systematic Lossy Error Protection (SLEP), which uses Wyner-Ziv coding ideas to achieve error resilience. This scheme achieves a graceful trade-off between the decoded video quality and the resilience to channel errors, effectively mitigating the cliff effect of FEC. In this work, we present an implementation of the SLEP scheme for H.264/AVC video transmission. In particular, we leverage H.264/AVC features such as redundant slices and flexible macroblock ordering in order to construct the

*P. Baccichet is with I.E.I.I.T. - Consiglio Nazionale delle Ricerche - Italy and is partially supported by MIUR (Italian Ministry of Education and Research) under Research Project PRIMO - “Reconfigurable platforms for wideband wireless communications”.

†This work has been supported by NSF Grant No: CCR-0310376.

Wyner-Ziv bit-stream, thus motivating a standard compatible implementation of an error resilience scheme for common video transmission application. We provide a detailed experimental comparison between SLEP and FEC for two scenarios, viz., transmission of video in IP packets over a wireless link where the application sees symbol errors as a result of wireless fading, and transmission of video over the Internet, where the application sees packet erasures due to congestion.

The SLEP system is based on the systematic lossy source-channel coding framework in which, a source X is transmitted over an analog channel A without coding. A second encoded version of X is sent over a digital channel D as enhancement information. The noisy version Y of the original serves as side information to decode the output of channel D and produce the enhanced version Z . Thus, source coding with decoder side information, i.e., Wyner-Ziv coding [4], is an integral part of the lossy source-channel coding configuration. The term systematic coding has been introduced as an extension of systematic error-correcting channel codes to refer to a partially uncoded transmission. Shamai, Verdú, and Zamir established information theoretic bounds and conditions for optimality of such a configuration in [5].

The systematic coding scheme for error resilience differs from other recently proposed schemes for distributed video coding [6–8]. The difference is that, in these schemes, the Wyner-Ziv codec is an integral part of the video encoding and is necessary for both source coding efficiency and resilience to channel errors. In contrast, the systematic source-channel coding scheme uses the Wyner-Ziv codec solely for error resilience and is, in principle, independent of the video compression scheme employed for the systematic transmission.

The remainder of this paper is organized as follows. Section 1 outlines the principle of Systematic Lossy Error Protection. In Section 3, the H.264/AVC standard support for redundant slices and flexible macroblock ordering is discussed, followed by an implementation of the SLEP system using these tools. Section 4 contains results of experimental simulation of two video transmission scenarios, in which we compare the error resilience of FEC with that of SLEP. Concluding remarks are presented in Section 5.

2. Systematic Lossy Error Protection

We now describe the concept of systematic lossy error protection, using H.264/AVC video transmission as an example. As shown in Fig. 1, consider a compressed video bit-stream which is transmitted across an error-prone channel without channel coding, comprising the systematic part of the transmission. For error resilience, a second independent bit-stream is generated using Wyner-Ziv encoding of the video signal. The Wyner-Ziv bit-stream allows the decoding of a coarsely quantized description of the original video signal. Since the H.264/AVC bit-stream is generated without consideration of the error resilience provided by the Wyner-Ziv coder, we refer to the overall scheme as systematic source-channel coding. The receiver decodes the compressed video signal, and conceals erroneous or lost portions. Even after concealment, some portions of the recovered video signal may contain unacceptably large errors. These errors are corrected up to a certain residual distortion by the Wyner-Ziv decoder.

The Wyner-Ziv bits allow error-free reconstruction of the coarser second description, employing the decoded video signal S' as side information. The coarser second description and side information S' are then combined to yield an improved decoded video signal S^* . In portions where the waveform S' is not affected by transmission errors, S^* is essentially identical to S' . However, in portions of the waveform where S' is substantially degraded by transmission errors, the second coarser representation transmitted at very low bit-rate in the Wyner-Ziv bit-stream limits the maximum degradation that can occur. Instead of the error-concealed decoded signal S' , the signal S^* at the output of the Wyner-Ziv encoder is fed back to the H.264/AVC decoder to serve as a more accurate reference frame for decoding of the future frames.

3. SLEP based on H.264/AVC Redundant Slices

3.1. Redundant Slices and FMO standard support

The standard specification for H.264/AVC [9] provides several error resilience tools [10]. In our scheme, we exploit *Flexible Macroblock Ordering* (FMO) and *Redundant Slices*, both available in the “Baseline” and “Extended” profiles. Using the *slice group* concept, every coded picture may

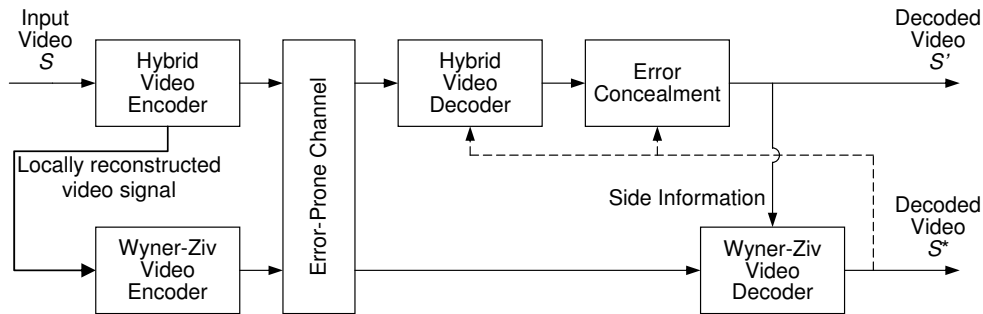


Fig. 1. A Wyner-Ziv decoder uses the decoded error-concealed video waveform as side information in a systematic lossy source-channel setup.

be divided into several macroblock partitions, defined by means of a “macroblock to slice group map” function. Then, a slice is defined as a sequence of macroblocks from one slice group, taken in raster scan order. The parameters that control the mapping functions are included in a *Picture Parameter Set* (PPS) which is sent to the receiver. Seven predefined mapping functions have been provided in the standard.

In our proposal, we exploit the FMO type 2 mapping, also called *Foreground with Left-Over*, that allows to define a maximum of seven slice groups for the foreground, plus one for the background. To activate these slice groups, 7 sets of macroblock coordinates are included in the PPS, signaling the **up-left** and **bottom-right** corners of each rectangle. Macroblocks belonging to overlapping slice groups are assigned to the one with lowest identifier. This FMO mapping allows to distinguish the foreground region of the image from the background, thus enabling the underlying transmission layers to adopt different protection policies for the two regions.

In order to construct the Wyner-Ziv description, as explained in Section 1, we exploit redundant slices. The H.264/AVC standard introduces support for redundant slices by defining the *Access Unit* concept. Practically, an access unit is one *Primary Coded Picture* (PCP) with some optional additional information, that may include one or more redundant pictures or Sequence Enhancement Information. Of course, the specification poses some constraints about the content of these redundant slices. In practice, (1) the redundant representations must follow the corresponding PCP in the same access unit, (2) the redundant picture cannot present a different sampling structure (frame/field) than the primary one and

(3) the reference picture list for the redundant picture must contain the same set of images as that of the PCP.

On the other hand, much freedom is left to the encoder about the actual content of the redundant slices. For example, an encoder can use different coding modes and quantization parameters than the ones used for the corresponding PCP. Redundant slices can have different shapes from those in the PCP. They may contain a different number of macroblocks and, more relevant to this work, decoded slices within the same redundant picture need not cover the entire picture area. Furthermore, redundant pictures are allowed to make use of different Picture Parameter Sets, and different FMO mapping functions, provided that the Picture Order Count (i.e., the timing information necessary for ordering the list of reference pictures) of the primary coded picture is preserved.

3.2. Wyner-Ziv encoding

The implementation of the SLEP scheme is shown in Fig. 2. The following operations are performed on the encoder side:

1. *ROI determination*: The image is analyzed to check for the existence of a Region of Interest (ROI). In practice, we determine the portions that do not need protection because decoder-based error concealment would reconstruct them with an acceptable distortion. This process is described in detail in Section 3.4 and may result in the generation of a PPS that specifies the FMO mapping for encoding the redundant slices.
2. *Generation of Redundant Slices*: Each macroblock belonging to the redundant description is encoded with the same coding mode,

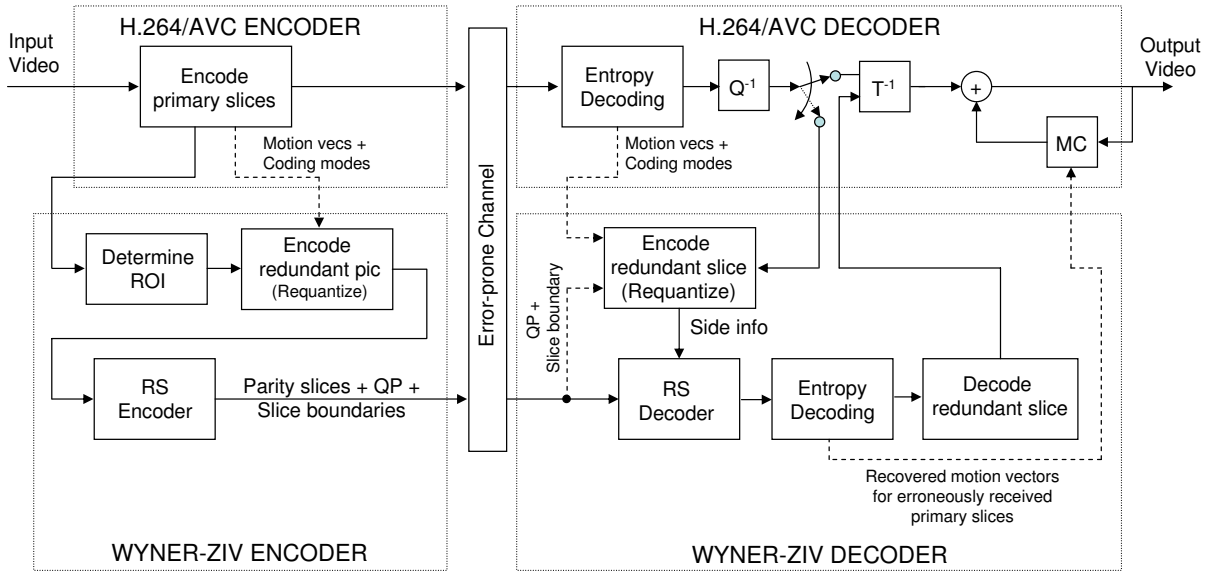


Fig. 2. Implementation of a SLEP system using H.264/AVC redundant slices and FMO-based region-of-interest determination. Reed Solomon codes applied across the redundant slices play the role of Slepian-Wolf codes in distributed source coding. At the receiver, the Wyner-Ziv decoder obtains the correct redundant slices using the error-prone primary coded slices as side information. The redundant description is used in lieu of the lost portions of the primary (systematic) signal.

motion vectors and reference pictures of the corresponding primary coded macroblock.

3. *Reed-Solomon encoding:* Reed-Solomon (RS) codes perform the role of Slepian-Wolf coding in this system. A byte-long Reed-Solomon code is applied across the redundant slices, to generate parity slices, as shown in Fig. 3. The number of parity slices generated per frame depends upon the allowable error resilience bit-rate, and can vary slightly from frame to frame. The redundant slices are then *discarded* and *only the parity slices* are included for transmission in the Wyner-Ziv bit-stream.
4. *Wyner-Ziv bit-stream generation:* In addition to the parity slices resulting from the previous step, we encode for each slice (1) the number of the first macroblock, (2) the number of macroblocks in a redundant slice and (3) the QP difference. This latter parameter could be also specified per macroblock instead of per slice, if the rate control is enabled to change it on a macroblock-by-macroblock basis.

3.3. Wyner-Ziv decoding

The Wyner-Ziv decoding process is activated only when transmission errors result in the loss of one or more slices from the bit-stream of the primary coded picture. Wyner-Ziv decoding consists of the following operations:

1. *Requantization to obtain Redundant Slices:* This step involves the requantization of the prediction residual signal of the primary coded picture, followed by entropy coding. This generates the redundant slices used as side information for the Wyner-Ziv decoder. Note that redundant slices can only be generated only for those portions of the frame, where the primary bit-stream has not experienced channel errors. The redundant slices corresponding to the error-prone portions are treated as erasures. Since the coding modes for the redundant macroblocks are identical to those in the primary bit-stream, the requantization procedure is straightforward. Note that this simplification is a slight departure from the conceptual SLEP system of Section 1, which required a full re-encoding of the error-concealed primary video signal. This would have a large complexity cost because motion estimation




Fig. 3. During Wyner-Ziv encoding, RS codes are applied across the redundant slices and only the parity slices are transmitted. At the decoder, the redundant slices can be recovered, possibly with a few erasures. During Wyner-Ziv decoding, the received parity slices are used to correct the erasures and recover redundant slices corresponding to the lost portions of the original video signal.

would have to be re-performed at the decoder. The requantization process sacrifices a small amount of coding efficiency, but requires very low complexity in addition to being more robust to Wyner-Ziv decoder failure.

2. *Reed-Solomon Slepian-Wolf decoding:* The parity slices received in the Wyner-Ziv bit-stream are now combined with the redundant slices, and erasure decoding is performed to recover the slices which were erased from the redundant bit-stream, as shown in Fig. 3. In the language of distributed video coding, the RS decoder functions as a Slepian-Wolf decoder, and recovers the correct redundant bit-stream using the error-prone redundant bit-stream as side information.
3. *Concealment of lost primary slices:* If Wyner-Ziv decoding succeeds, the lost portions of the prediction residual from the primary (systematic) signal, are replaced by the quantized redundant prediction error signal. The H.264/AVC decoder then performs motion compensation in the conventional manner, using the redundant prediction error signal and the motion vectors recovered from Wyner-Ziv decoding. The coarse fall-back operation results in a quantization mismatch which propagates to the future frames, but a drastic reduction in picture quality is avoided.

3.4. ROI signalling and coding

SLEP is essentially based on a tradeoff between error robustness and distortion introduced by the Wyner-Ziv decoded signal used to replace the original video. For error resilience at high error rates, it is usually desirable to generate redundant descriptions with a very coarse quantization step. In these cases, the distortion introduced by the Wyner-Ziv description can be large, especially for intra-coded macroblocks. We try to limit this effect by avoiding redundant encoding of “unimportant” regions of the image. This allows better management of the bit-rate budget by encoding only the macroblocks belonging to the ROI with a finer quantizer.

In order to determine the ROI, we follow the process reported in Fig. 4. First, we evaluate the impact of the loss of each slice, simulating the error concealment process at the encoder. The error signal, obtained as the difference between the encoded frame and the concealed frame, provides a measure of the expected distortion in case of losses. For each macroblock, we compute the “Mean Absolute Error” (MAE), thus producing a *Significance Map* for the whole image. This solution is a simple but effective compromise between the extra computational complexity and the effectiveness in determining the perceptual significance of the region-of-interest.

After constructing a matrix that essentially captures the perceptual significance of each macroblock, the ROI is obtained by thresholding the MAE. Then, it is necessary to cover the ROI with up to seven rectangles (slice groups) that could be

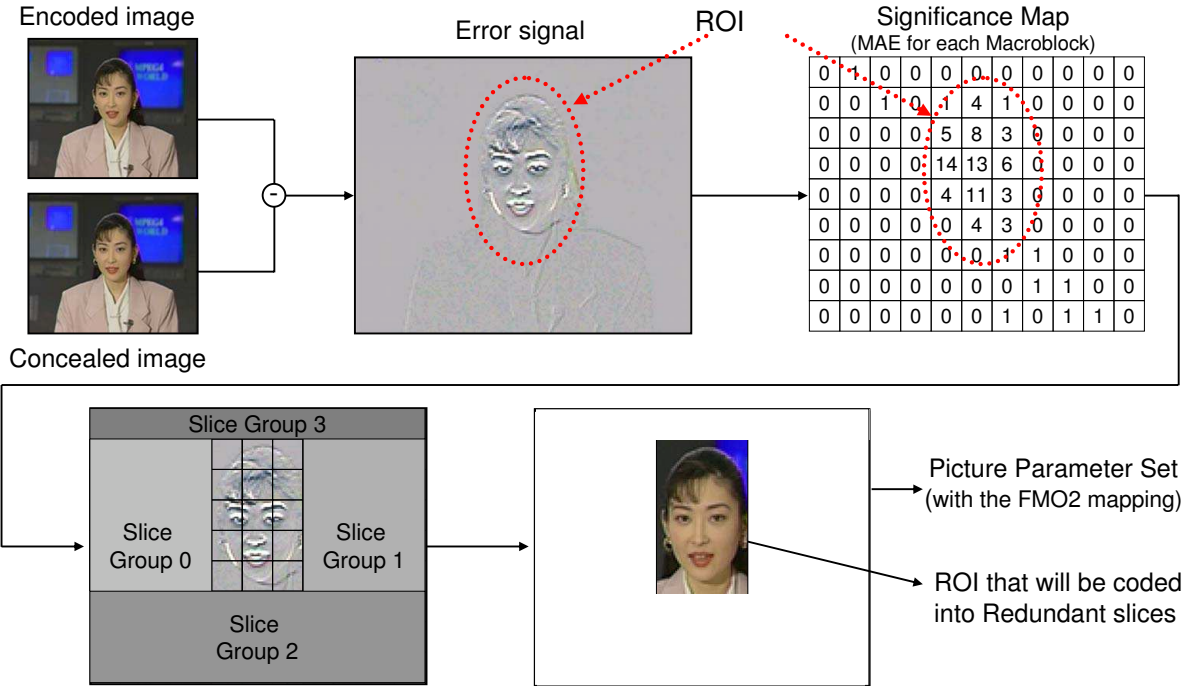


Fig. 4. A Region Of Interest (ROI) is determined at the encoder by finding the mean absolute error between the current image and its locally error concealed version. The error image is then thresholded to obtain the ROI.

signalled by means of FMO type 2. In practice, FMO 2 specifies the area *not covered* by any of the seven rectangles (usually the background) by placing it into an eighth slice group. Therefore, we attempt to cover the greatest possible number of “unimportant” macroblocks with the first seven rectangles, leaving the relevant macroblocks in the background. This is performed by means of a dichotomic search in the space of possible rectangles. A list of rectangles is constructed, and the largest one is sub-divided into two. Then the MAE of the contained macroblocks is compared to a threshold, to determine whether the rectangle contains a portion of the ROI. This procedure is iterated until 7 slice groups covering the greatest possible number of unimportant macroblocks are found. The procedure for determining the ROI is detailed in [11].

4. Experimental Results

In this section, we describe simulation results demonstrating the error resilience of the proposed SLEP scheme. We investigate two scenarios, viz., video transmission over a wireless link, and video transmission over the Internet. The following settings are common for both experiments:

- We use the JVT version JM 10.1 [12] of the H.264/AVC video codec for our simulations. The experiments are carried out on CIF resolution sequences and all the graphs in this section contain results averaged over 30 channel realizations. The GOP structure used is *I-P-P-P...* and the sequences are encoded at 15 frames per second.
- We use two high-motion sequences, *Foreman* encoded at 1 Mbps, and *Bus* encoded at 1.5 Mbps, and one low-motion sequence, *Akiyo* encoded at 400 kbps.
- We use redundant descriptions encoded at 100%, 50%, 25% and 10% of the bit-rate of the primary stream. In the following sections, schemes using these redundant descriptions will be referred to as FEC, SLEP50, SLEP25 and SLEP10 respectively. Note that these numbers refer to the bit-rate of the redundant description and *not* to the transmitted Wyner-Ziv bit-rate. The transmitted Wyner-Ziv bit-rate will be specified and will be identical for *all* the schemes in a given experiment.




Fig. 6. Improvement in the error-resilience of SLEP when the Wyner-Ziv bit-rate is increased from 10% to 20% of the bit-rate of the systematic transmission, which consists of the *Bus* sequence encoded at 1.5 Mbps.




Fig. 5. SLEP provides a graceful fall-off in decoded picture quality as the symbol error probability increases. In addition, it allows for a flexible trade-off between the error-resilience provided by a redundant description and the residual distortion that can be tolerated after Wyner-Ziv decoding. The plot is for the *Bus* sequence with the Wyner-Ziv bit-rate being 150 kbps, which is 10% of the rate of the systematic transmission.

- To avoid the delay incurred for buffering an I-frame, and to ensure that the bit-rate profile does not fluctuate drastically over time, we use Intra macroblock line-refresh in all our simulations, i.e., one row in each frame is encoded using the intra macroblocks, thus generating a complete frame-refresh every 18 frames for a CIF resolution sequence.
- The rate control method specified in the standard codec [9] is used at the encoder to determine the modulation of the quantization parameters while encoding the macroblocks in the video sequence. Once the encoding rates for the primary and redun-

dant slices has been decided, the rate control algorithm for the primary picture proceeds independently from that for the redundant picture.

- The slices of the primary picture, i.e., the systematic transmission are constrained to be 400 bytes long. To this, we add the RTP, UDP and IP headers which occupy a total of $12+8+20=40$ bytes per slice. Since each video slice travels inside an IP packet, the words “slice” and “packet” will be used interchangeably. Thus, unlike our previous experiments with the MPEG-2 codec [2,13], the slices are no longer constrained to contain the same number of macroblocks. The packetization for the slices in the Wyner-Ziv bit-stream will differ according to the experiment and will be described below.
- At very high error probabilities, the number of erased redundant slices is too large and Wyner-Ziv decoding fails for some or all video frames. In this case, the non-normative error concealment scheme [14] included in the reference JVT codec, is used to conceal the lost portions of the systematic transmission.

4.1. Video Transmission over a Wireless link

We now discuss the simulation results for a communication scenario in which a video sequence is transmitted over a wireless link. In this case, the transmission experiences symbol errors. For this scenario, we constrain each redundant slice to contain the same number of macroblocks as the corresponding primary slice. The rationale for this is as follows: For a given symbol error rate p_s ,



Fig. 7. Decoded frames from the *Foreman* sequence at a symbol error probability of 10^{-4} . 100 kbps of FEC fails to correct the channel errors, reducing the picture quality. At the same bit-rate, SLEP replaces the lost portions by their Wyner-Ziv decoded redundant counterparts and ensures acceptable picture quality. The 2 dB difference between the error-free picture and the picture decoded by SLEP is due to the quantization mismatch from Wyner-Ziv decoding.

the probability that the packet arrives in error, is proportional to its length l , since $1 - (1 - p_s)^l \approx lp_s$ for small p_s . Further, for the same number of macroblocks, a redundant slice is smaller than a primary coded slice, since the redundant slice in SLEP uses coarser quantization. Thus a redundant slice containing the same number of macroblocks as the primary slice, is less likely to be lost. Of course, if there are many small redundant slices, then the bit-rate overhead corresponding to the slice headers is quite large. However, since the redundant slices themselves are not transmitted to the decoder in SLEP, this overhead is restricted only to the parity slices in the Wyner-Ziv stream, and is not very large. Note that applying RS codes across the redundant slices dictates that the length of the Wyner-Ziv slices in a frame is equal to the length of the largest redundant slice in that frame.

Fig. 5 shows a plot of average video quality versus symbol error probability for the *Bus* sequence. The error resilience bit-rate is chosen to be 10% of the systematic bit-rate, i.e., 10% of 1.5 Mbps = 150 kbps. The cases compared are SLEP with redundant description bit-rates of 1.5 Mbps (FEC), 750 kbps (SLEP50), 375 Kbps (SLEP25), and 150 Kbps (SLEP10). It is clear that the error resilience of FEC is poor at high symbol error probabilities, while that of SLEP increases as the redundant picture quality becomes coarser. Thus, SLEP10 has the highest error resilience at high symbol error rates. The curves also highlight the trade-off between the error resilience of a SLEP scheme and the picture quality after Wyner-Ziv

decoding. Specifically, SLEP10 experiences the highest quantization mismatch after Wyner-Ziv decoding, and hence at low symbol error rates, it pays the price of slightly reduced picture quality. The performance of SLEP50 and SLEP25 is in between that of FEC and SLEP10, in keeping with this trade-off. As expected, the error resilience of each scheme improves when the Wyner-Ziv bit-rate is increased from 10% (150 kbps) to 20% (300 kbps) of the source bit-rate of the primary picture, as shown in Fig. 6. Interestingly, at high error rates, SLEP25 with only 150 kbps Wyner-Ziv bit-rate outperforms the scheme with 300 kbps FEC.

Fig. 7 shows the improvement in visual quality for a frame of the *Foreman* sequence, when the error concealment artifacts resulting from the failure of FEC are avoided in the SLEP system.

Fig. 8 compares FEC with the SLEP10 scheme for the *Akiyo* sequence, with 10 % error resilience bit-rate, with and without ROI estimation. Since *Akiyo* is a low motion sequence, it is beneficial to use ROI estimation before applying Wyner-Ziv coding, and this is borne out by the experiments which show about 1-2 dB improvement in average picture quality when ROI estimation is used before encoding the redundant slices. The reason for the improvement is that, while generating the redundant description in SLEP10, the entire redundant description bit-rate (10% of 400 kbps = 40 kbps), is reserved for encoding the ROI instead of the whole frame, resulting in better picture quality for the ROI. The improvement can be better appreciated by observing the instantaneous




Fig. 9. A video sequence (*Akiyo*) decoded at an error probability of 4×10^{-4} . The Wyner-Ziv bit-rate is 10% of the source coding bit-rate of 400 kbps. Both the SLEP10 plots have the same bit-rate allocated to the redundant slices, however, superior quality is observed when the redundant encoding is restricted to the ROI. The dotted lines represent the average PSNR for the video sequence.

variation of video quality over time, as shown in Fig. 9. Besides providing a higher average video quality, redundant encoding of the ROI also lowers the fluctuation in video quality, since the distortion in the ROI is now closer to that in the original picture.

4.2. Video Transmission over the Internet

For the Internet video transmission experiment, we note that the losses are only due to congestion at the network routers. Thus the receiver sees an erasure channel, in which video packets are lost with a certain probability that does not directly depend on the length of the packet. It would appear then, that it is beneficial to use

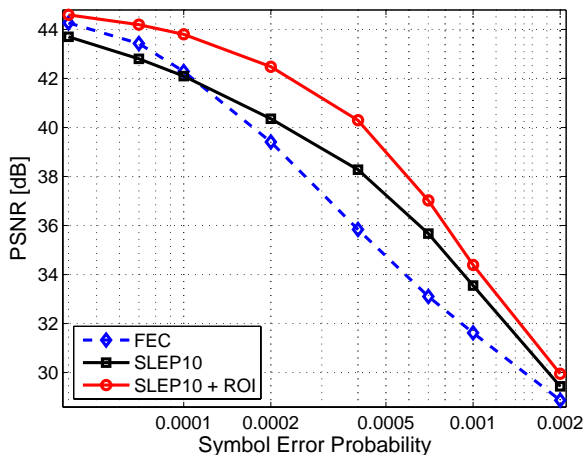


Fig. 8. For a sequence with low-motion (*Akiyo*), SLEP is only performed on the region of interest. This significantly reduces the residual distortion resulting from Wyner-Ziv decoding, and improves the overall picture quality.

the same packet length for the primary and the Wyner-Ziv slices. This would require that the redundant slices also be 400 bytes long. Now, since the redundant description is coarsely quantized, a redundant slice will contain more macroblocks than a primary slice of the same length. Since redundant and primary pictures are encoded independently, it can happen that the loss of one primary slice results in the loss of two redundant slices. This is clearly undesirable and hampers the performance of SLEP, as shown in Fig. 10. To prevent this, we again impose the restriction that primary and redundant slices contain the same number of macroblocks, as in the wireless scenario. Thus, the length of the transmitted Wyner-Ziv slices will again be equal to the length of the largest redundant slice in a given frame. Unlike the wireless scenario however, since losses are due to congestion, and not symbol errors, the error probability experienced by all the packets is kept the same irrespective of their length. By preventing Wyner-Ziv decoding failure due to loss of synchronization between the primary and redundant slices, the performance of SLEP improves as shown in Fig. 10. Note that for FEC, the performance in both cases is identical, since primary and redundant slices are identical by construction.

5. Conclusions and Ongoing Work

The Systematic Lossy Error Protection scheme uses of a Wyner-Ziv bit-stream to provide error resilience in a systematic lossy source-channel




Fig. 10. Effect of the method of packetization on the performance of SLEP for the *Foreman* sequence. Using the same packet length as the primary slices results in poor performance due to loss of synchronization between the primary and redundant slices. Using the same number of macroblocks in the primary and redundant slices dramatically improves the performance of SLEP, and trends similar to the wireless case are observed even though all packets experience the same loss rate.

coding configuration. The Wyner-Ziv bit-stream is generated efficiently by using H.264/AVC redundant slices in conjunction with Reed-Solomon coding. Wyner-Ziv protection can be enhanced by restricting the encoding of redundant slices to the high-motion regions of the video sequence by specifying a region-of-interest to the SLEP scheme, using standard-compliant flexible macroblock ordering (FMO 2). Experimental simulations show that SLEP provides graceful degradation of video quality when the channel error probability increases, without using scalable video representations. Further, the quality of the redundant description can be used to trade off the amount of error resilience desired with the amount of degradation tolerable after Wyner-Ziv decoding.

The bit-rate/quality of the redundant descriptions used in SLEP results in an extra degree of freedom in the optimization of a SLEP system, compared to the optimization of traditional FEC-based systems. Parallel to the present work, a model has been developed [15] to study the end-to-end behavior of the SLEP system and to find the bit-rates for encoding the primary and redundant pictures, as well as the Wyner-Ziv bit-rate that optimizes the received picture quality given the channel capacity, channel model and the error probability.

References

- [1] B. Girod, A. Aaron, S. Rane, D. Rebollo-Monedero, Distributed video coding, Proc. IEEE, Special Issue on Advances in Video Coding and Delivery 93 (1) (2005) 71–83.
- [2] S. Rane, A. Aaron, B. Girod, Systematic lossy forward error protection for error resilient digital video broadcasting - A Wyner-Ziv coding approach, in: Proc. IEEE International Conference on Image Processing, Singapore, 2004.
- [3] S. Rane, B. Girod, Analysis of Error-Resilient Video Transmission based on Systematic Source-Channel Coding, in: Picture Coding Symposium (PCS 2004), San Francisco, CA, 2004.
- [4] A. D. Wyner, J. Ziv, The rate-distortion function for source coding with side information at the decoder, IEEE Transactions on Information Theory IT-22 (1) (1976) 1–10.
- [5] S. Shamai, S. Verdú, R. Zamir, Systematic lossy source/channel coding, IEEE Transactions on Information Theory 44 (2) (1998) 564–579.
- [6] A. Aaron, S. Rane, B. Girod, Wyner-Ziv video coding with hash-based motion compensation at the receiver, in: Proc. IEEE International Conference on Image Processing, Singapore, 2004, to appear.
- [7] R. Puri, K. Ramchandran, PRISM: A new robust video coding architecture based on distributed compression principles, in: Proc. Allerton Conference on Communication, Control, and Computing, Allerton, IL, 2002.
- [8] A. Jagmohan, A. Sehgal, N. Ahuja, Predictive encoding using coset codes, in: Proc. IEEE International Conference on Image Processing, Vol. 2, Rochester, NY, 2002, pp. 29–32.
- [9] J. V. T. J. of ISO/IEC MPEG & ITU-T VCEG, Draft ITU-T recommendation and Final Draft International Standard of Joint Video Specification (ITU-T Rec. H.264 — ISO/IEC 14496/10 AVC - JVT-G050r1.doc), ISO/IEC MPEG & ITU-T VCEG, 2003.
- [10] T. Stockhammer, M. Hannuksela, T. Wiegand, H.264/AVC in wireless environments, IEEE Trans. Circuits and Systems for Video Tech. 13 (7) (2003) 657–673, special Issue on the H.264/AVC Video Coding Standard.
- [11] P. Baccichet, A. Chimienti, Forward Selective Protection exploiting Redundant Slices and FMO in H.264/AVC, in: Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2006), Toulouse, France, 2005, submitted.
- [12] K. Suehring, JVT JM reference software home page (online), <http://iphome.hhi.de/suehring/tml/>.
- [13] S. Rane, A. Aaron, B. Girod, Error Resilient Video Transmission using Multiple Embedded Wyner-Ziv Descriptions, in: Proc. IEEE International Conference on Image Processing, Genoa, Italy, 2005.
- [14] Y. Wang, M. Hannuksela, V. Varsa, A. Hourunranta, M. Gabbouj, The Error Concealment Feature In The H.26L Test Model, in: Proc. IEEE ICIP, 2002.
- [15] S. Rane, P. Baccichet, B. Girod, Modeling and Optimization of a Systematic Lossy Error Protection System, in: Picture Coding Symposium (PCS 2006), Beijing, China, 2005, submitted.