Journal of Zhejiang University SCIENCE A
ISSN 1009-3095 (Print); ISSN 1862-1775 (Online)
www.zju.edu.cn/jzus; www.springerlink.com
E-mail: jzus@zju.edu.cn

RTP-based broadcast streaming of high definition H.264/AVC video: an error robustness evaluation*

HILLESTAD Odd Inge, JETLUND Ola, PERKIS Andrew
(*Centre for Quantifiable Quality of Service in Communication Systems, Norwegian University of Science and Technology, Trondheim N-7491, Norway*)
E-mail: {hillesta; jetlund; Andrew}@q2s.ntnu.no
Received Nov. 26, 2005; revision accepted Feb. 16, 2006

Abstract: In this work, we present an evaluation of the performance and error robustness of RTP-based broadcast streaming of high-quality high-definition (HD) H.264/AVC video. Using a fully controlled IP test bed (Hillestad *et al*., 2005), we broadcast high-definition video over RTP/UDP, and use an IP network emulator to introduce a varying amount of randomly distributed packet loss. A high-performance network interface monitoring card is used to capture the video packets into a trace file. Purpose-built software parses the trace file, analyzes the RTP stream and assembles the correctly received NAL units into an H.264/AVC Annex B byte stream file, which is subsequently decoded by JVT JM 10.1 reference software. The proposed measurement setup is a novel, practical and intuitive approach to perform error resilience testing of real-world H.264/AVC broadcast applications. Through a series of experiments, we evaluate some of the error resilience features of the H.264/AVC standard, and see how they perform at packet loss rates from 0.01% to 5%. The results confirmed that an appropriate slice partitioning scheme is essential to have a graceful degradation in received quality in the case of packet loss. While flexible macroblock ordering reduces the compression efficiency about 1 dB for our test material, reconstructed video quality is improved for loss rates above 0.25%.

Key words: H.264/AVC, Video streaming, Error robustness
doi:10.1631/jzus.2006.AS0019 **Document code:** A **CLC number:** TN919.8

INTRODUCTION

Streaming media services delivered over IP networks are gaining popularity and momentum, and consumers show increased interest in being able to play back, and enjoy their media wherever they are. The ever increasing deployment of high speed wired and wireless Internet access networks, and continuous development of more efficient compression schemes for audio and video, are two of many important factors that enable content and service providers to deliver higher quality IP-based multimedia services to end-users.

The goal of any multimedia communication system should be to maximize the end-user's perceived quality of the delivered service. However, to measure the end-user perception of audiovisual service quality, or similarly, to which extent they react objectively to distortions introduced by compression and packet-switched transmission, is extremely difficult and depends on factors that are not easily modelled.

Examples of such factors are: context in which the service is being used, user expectation, human diversity, preferences, and application knowledge. Still, to this end, VQEG (video quality experts group) (http://www.vqeg.org) is currently working towards standardization of quality metrics that could be employed by multimedia streaming applications. Even though the peak signal-to-noise ratio (PSNR) has been shown not to correlate very well with perceived quality (Wang *et al*., 2004), we will still use it in this work because of its simplicity and familiarity.

This paper is organized as follows. Section 2

* Project supported by the Research Council of Norway, Norwegian University of Science and Technology (NTNU), and the Norwegian Resarch Network (UNINETT)

introduces H.264/AVC and presents related work. Section 3 describes the test setup and the error resilience schemes to be compared. The subsequent section presents resulting reconstructed PSNR for different packet loss rates. Finally, Section 5 gives a short summary.

H.264/AVC AND RELATED WORK

H.264/AVC (Advanced Video Coding) is the current state-of-the-art international video coding standard (ITU-T and ISO/IEC JTC-1, 2003). Similar to previous ITU-T and MPEG video coding standards, H.264/AVC is a block-based motion-compensated hybrid video coding scheme. The standard defines two main layers: a video coding layer (VCL) and a network abstraction layer (NAL). The definition of the VCL includes tools and methods to code a set of macroblocks of a picture into a slice partition or a data partition. The slices/data partitions are conceptually transferred by the network abstraction layer as NAL units, which are the basic transport units. Essential picture header information is assembled in parameters sets, more specifically the sequence parameter sets (SPSs) and picture parameter sets (PPSs). For more information on the new tools and features of the H.264/AVC standard, see (Wiegand *et al*., 2003; Sullivan and Wiegand, 2004).

H.264/AVC specifies several error resilience tools for efficient delivery in error-prone environments. The partitioning of a picture into slices is very flexible, both in terms of slice sizes and with respect to which macroblocks are allocated to each slice. In addition to simply assigning the macroblocks of a picture to slices in raster scan order, a scheme known as flexible macroblock ordering (FMO) can be used to partition a picture into a number of slice groups using a macroblock allocation map (Ostermann *et al*., 2004). Pre-defined allocations like e.g. Interleaved and Dispersed slice groups are specified, but the standard also enables explicit allocation of macroblocks to a slice group. Other error resilience tools include arbitrary slice ordering (ASO), data partitioning (DP), and redundant slices (RS). Because the coding mode can be decided on a macroblock level, the standard supports insertion of intra-coded macroblocks into P or B pictures to stop error propagation. For more information on the error resilience features of H.264/AVC and transport in error-prone environments, see (Ostermann *et al*., 2004; Wenger, 2003a; Stockhammar *et al*., 2003).

The fidelity range extension (FRExt) amendment to H.264/AVC includes new tools and profiles for high-quality consumer and broadcast applications (Sullivan *et al*., 2004). However, no advanced error resilience tools are allowed in the FRExt profiles, making them less suitable for delivery in error-prone environments unless e.g. forward error correction codes, application level ARQ or a reliable transport protocol are used.

RFC 3984 specifies the mapping of NAL units to RTP packets and describes issues related to fragmentation and aggregation of NAL units (Wenger *et al*., 2005). An H.264 decoder is assumed to be given NAL units in decoding order. In this work, we consider only the simple packetization mode, in which no interleaving is used and NAL units are transmitted in decoding order. If the size of a slice partition, data partition or parameter set is larger than the size of the MTU on the access network, then the NAL unit has to be fragmented. The fragments are called FU-A NAL units, and if one FU-A is lost, then the entire corresponding NAL unit is corrupted, and has to be discarded.

Most of the previously published work on H.264/AVC error robustness has been done using JVT's common test conditions for IP-based transmission (Wenger, 2001) which includes a simple simulator discarding packets based on a loss pattern file. The loss patterns are obtained from Internet experiments, and with a few exceptions mainly consist of scattered packet loss (Wenger, 2003b). Wenger (2003a) gave a comprehensive overview of the error-resilience tools in the H.264/AVC standard, and presented results for six different encoder configurations comparing the use of intra macroblock updates, slice partitioning, interleaved FMO, dispersed FMO, and data partitioning for two CIF resolution sequences. FMO and data partitioning were concluded to be particularly valuable error resilience tools.

Hallbach and Olsen (2004) evaluated the use of motion-sensitive intra macroblock updates, i.e., areas with high motion are more likely to be intra coded, as a means to stop error propagation. Stockhammar *et al*.(2003) discussed the use of H.264/AVC in a wireless environment and presented an error robustness evaluation for a conversational service with and without feedback using the JVT common test conditions for RTP streaming over 3GPP (Roth *et al.*, 2001). While slice partitioning and rate-distortion optimized mode decision give good results without the use of any feedback, excellent results were reported when multiple reference frames and feedback were used to effectively stop error propagation. Calafate *et al.*(2004) studied the error resilience of H.264/AVC in an ad-hoc wireless network scenario.

While previous work to a large extent has focused on low-resolution low-quality applications, this work studies the error robustness of a high-quality high-definition H.264/AVC broadcast service over an emulated IP network that introduce randomly distributed packet loss. We will not consider the delay and delay jitter introduced in real packet-switched networks, assuming that the playout buffer can be set large enough so that all delay jitter is absorbed. Although simple, random packet loss may be a valid model for low and transient congestion if queue management policies like Random Early Detection or Early Random Drop is employed in networks routers (Lin and Morris, 1997). For typical Internet packet loss, more complex state models incorporating temporal dependencies of lost packets are often more appropriate (Yajnik *et al.*, 1999).

One difference between this work and the above-mentioned contributions is that the proposed setup could easily be used to test and compare the error resilient encoding and packetization of real commercial H.264/AVC over RTP/UDP/IP broadcast solutions. Further, using the packet flow regenerator described in (Hillestad *et al.*, 2005), the trace file from a given network emulation can be played back to a streaming media client for development, subjective testing, demonstration or other purposes. For instance, it can be used to recreate a specific network condition for subjective testing of a streaming media session that has been impaired by packet loss.

MEASUREMENT SETUP

For our measurements we used an excerpt from the mini-movie “standardized evaluation material” (StEM) under license from digital cinema initiatives (DCI) (DCI & ASC, 2004). StEM is available in 4k resolution (4096 by 1714 pixels), has a temporal resolution of 24 frames per second and contains 16605 frames in total. The excerpt from StEM, from now on referred to as STEM-A, consists of 576 frames (from frame number 1716 up to and including frame number 2292) giving a playing time of 24 s. To create a 720p version of STEM-A, cropping was first employed to remove 768 pixels on both the left and right side of each frame, and also to remove 137 pixels from both the top and bottom of each frame. The resulting 2560 by 1440 pixel frames were then converted to the target 1280 by 720 resolution by bilinear downsampling. The clip has four scene changes, cross-fades, camera panning, complex object motion and a high level of texture detail, thus providing a good challenge for the error resilience tools and concealment algorithms in the JM reference software.

STEM-A was encoded using the H.264/AVC reference software JM version 10.1 (http://iphome.hhi.de/suehring/tml/download/). An “IPPP” GOP structure was employed, with intra period of 24 frames. IDR (Instantaneous Decoding Refresh) pictures were used to try to prevent reconstruction errors to propagate across GOP boundaries. Rate-distortion optimized macroblock mode decision was turned on, while random macroblock refresh was turned off.

The eight different configurations we compared in this work are listed in Table 1 together with the quantization parameter (QP) used in the encoding and the resulting average bit rate. “SL1” denotes the case when using no slice partitioning at all, that is, the entire primary coded picture is transported as a single NAL unit. Similarly, “SL5”, “SL9” and “SL16” is encoded using 5, 9 and 16 slices per frame, respectively.

All configurations having names starting with “SLMTU” are encoded in such a way that no NAL units are larger than 1450 bytes, which is the maximum

Table 1 The eight configurations, quantization parameter (QP) used in encoding, and resulting bit rate

Name	*QP*	Bit rate (kbps)
SL1	26	2950
SL5	26	3019
SL9	26	3078
SL16	26	3178
SLMTU-CIP	27	2910
SLMTU-FMO-CIP	28	3016
SLMTU-FMO-MBLINE	28	3016
SLMTU-FMO-PYR	27	2990

transmission unit (MTU) payload on our Ethernet network, making sure that no NAL units are fragmented by the network abstraction layer. This feature, together with constrained intra prediction, which prevents intra coded macroblocks from using inter-coded macroblock for prediction, is introduced in "SLMTU-CIP". Furthermore, "SLMTU-FMO-CIP" adds FMO using two slice groups per picture and the dispersed slice group map type, also known as the "checkerboard" FMO pattern (Ostermann *et al*., 2004).

While the above configurations have all been encoded using I and P frames only, and an intra period of 24 pictures, the next configuration instead uses groups of intra coded macroblocks to refresh the prediction. In "SLMTU-FMO-MBLINE", one line of macroblocks, or rather group of macroblocks, are intra coded in every picture. Considering that there are 45 lines of macroblocks in a 720p frame, the number of regularly intra coded macroblocks would correspond to an I frame about every 2 s for our test material.

The next configuration, "SLMTU-FMO-PYR", was encoded using a four level hierarchical coding structure also referred to as a pyramid (Schwarz *et al*., 2005). There are three pyramids inside a GOP of 24 pictures, and the coding order is "I_0 P_8 B_{R4} B_{R2} B_{R6} B_1 B_3 B_5 B_7 P_{16}", where B_R denotes B pictures that are used for reference in a higher level of the pyramid, and the lower case numbers indicate the output order of the pictures.

It is well known that extensive slice partitioning decreases compression efficiency, since prediction is constrained within one slice of a picture and also because of the increased overhead associated with slice and protocol headers (Hallbach and Olsen, 2004).

Fig.1 shows the rate-distortion performance of the different configurations (without taking protocol over-head into account). The quantization parameter was fixed for all pictures in the sequences and was chosen so that the average bit rate of all configurations was as close to a target rate of 3 Mbps as possible. We observe that for our test material, the dispersed FMO scheme decreases the compression efficiency by roughly 1 dB for the target rate. This is due to the restriction FMO places on the intra-picture prediction techniques used (Wenger, 2003b).

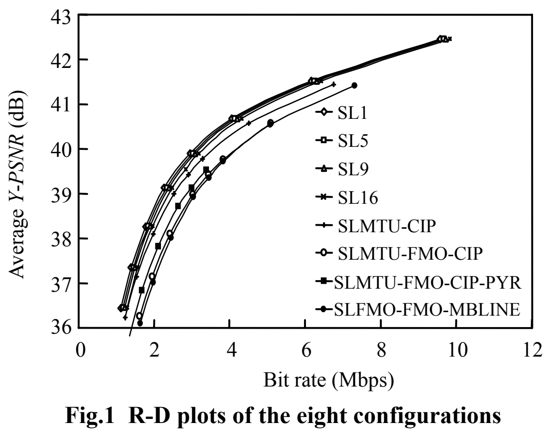

Fig.1 R-D plots of the eight configurations

The eight H.264/AVC video files were wrapped in MP4 containers and hinted with an MTU of 1450 bytes using tools from MPEG4IP (http://mpeg4ip.Sourceforge.net). The media was then broadcasted over the test network using the QTBroadcaster tool included in the open-source Darwin Streaming Server package, sent through an Empirix PacketSphere IP network emulator (http://www.empirix.com), and captured by an Endace DAG 3.5E network interface monitoring card (http://www.endace.com). The DAG card is capable of capturing packets at high speeds, and writes the entire packet with MAC/IP header and IP payload to a trace file (in the PCAP format) together with a microsecond-accuracy time stamp that reflects the packet's arrival time. The network emulator was configured to introduce uniformly distributed packet loss at different rates, namely 0.01%, 0.1%, 0.25%, 0.5%, 1%, 1.5%, 2%, 3% and 5%. Each media file was transmitted five times over the test network at each configured packet loss rate, giving a total of 45 measurements for each of the eight configurations. For more information on the network emulator and the network interface monitoring card used in these measurements, see (Hillestad *et al*., 2005).

Fig.2 summarizes the test setup. Due to the high complexity involved in real-time decoding of high-definition H.264/AVC, packet reception and decoding is done offline. A purposely-built application, "pcap2avc", was developed to implement the packet reception and stream assembly. Packets are sequentially read from the trace file, MAPI (http://mapi.uninett.no/) is used to filter out the video packet flow, and each NAL unit is assembled according to RFC 3984 (Wenger *et al*., 2005). NAL units which are correctly received are then written to an H.264/AVC Annex B byte stream file. As previously mentioned in Section 2, corrupted NAL units have to be discarded. Finally, decoding of the received bit-stream is done using H.264/AVC reference software (http://iphome.hhi.de/suehring/tml/download/).

JM version 10.1 supports the ability to conceal entire frame losses, which is particularly helpful for configuration "SL1", where the loss of a single FU-A NAL unit means that the entire frame is lost. The current limitation of the picture error concealment algorithm in JM 10.1 is that it can only conceal I and P pictures, and that it is not able to conceal the last P picture of a GOP if this is lost. This has to be taken into account when calculating the *PSNR* to prevent misalignment between frames in the original and decoded sequences. As long as at least one slice is correctly received for a picture, the JM reference software is capable of concealing lost slices using temporal and spatial error concealment algorithms (Wang *et al*., 2002).

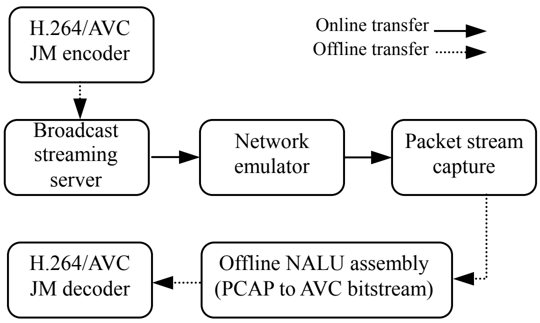

Fig.2 Measurement setup

RESULTS

In this section we present results of the reconstructed video quality for the eight different configurations described in Section 3. The peak signal-to-noise ratio (PSNR) between the uncompressed original and the decoded video is calculated based on the luminance video component only, denoted Y-PSNR. When calculating the *Y-PSNR*, no penalty was introduced for missing, non-concealed frames.

Fig.3 shows reconstructed *Y-PSNR* at packet loss

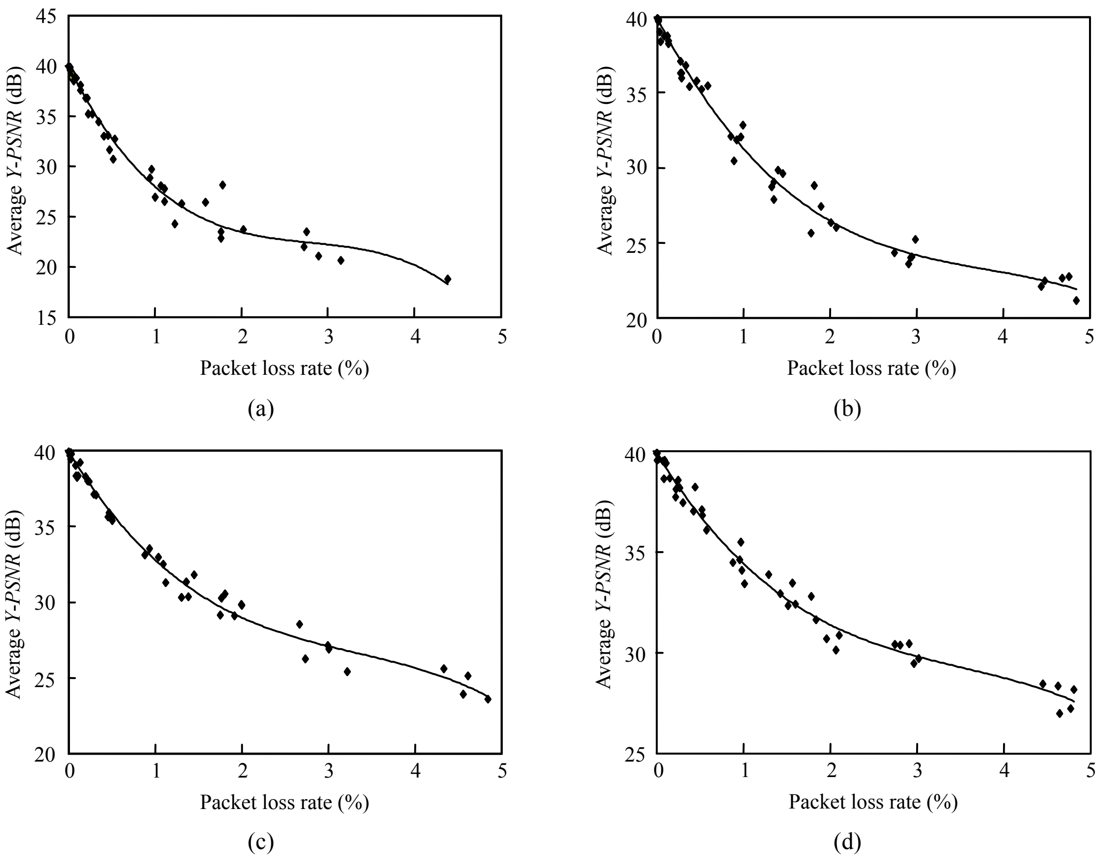

Fig.3 Average reconstructed *Y-PSNR* for the configuration (a) "SL1"; (b) "SL5"; (c) "SL9"; (d) "SL16" at different packet loss rates

rates from 0.01% to 5%. One dot in each of the scatter plots corresponds to the average *Y-PSNR* across all the frames of the reconstructed video sequence for a single run through the test bed. Curves estimated using second or third order polynomial regression (depending on which order made the best fit) are also shown in the graphs for purpose of illustration.

From the plots we can clearly see that the error resilience with respect to reconstructed *Y-PSNR* improves significantly when more advanced slice partitioning schemes are being employed.

For instance, if we compare at which packet loss rate the configurations seem to have quality reduction to 35 dB *Y-PSNR* on average, we see that "SL1" reaches this point at a loss rate of approximately 0.3%. The corresponding figures for "SL5", "SL9", and "SL16" are 0.5%, 0.6%, and 0.9%, respectively.

When the slice partitioning is matched to the MTU of the underlying network, and constrained intra prediction is employed, we see from Fig.4 that the performance is further improved. For "SLMTU-CIP", the reconstructed *Y-PSNR* stays above 35 dB for packet loss rates lower than 1.5%.

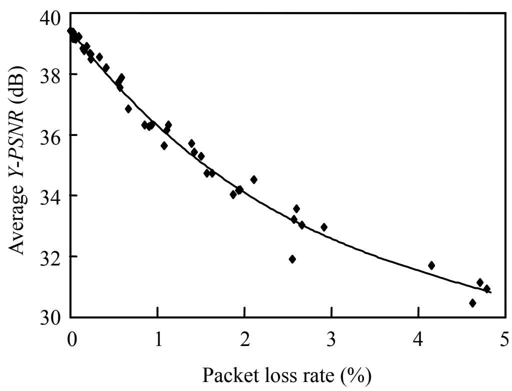

Fig.4 Average reconstructed *Y-PSNR* for the configuration "SLMTU-CIP" for different PLR

Fig.5 shows how the deficiency in proper slice partitioning affects the video delivery when packet loss occurs. The effective packet loss rate, here equal to the total number of transmitted packets divided by the sum of lost and discarded packets, is plotted as a function of packet loss for five sets of measurements. For "SL1", where no slice partitioning is employed, nearly 20% of the packets have to be discarded when 3% of the packets are lost in the network. The situation improves with increasing slice partitioning, and for "SLMTU-CIP", no packets have to be discarded, meaning that the effective PLR is equal to the PLR inflicted by the network.

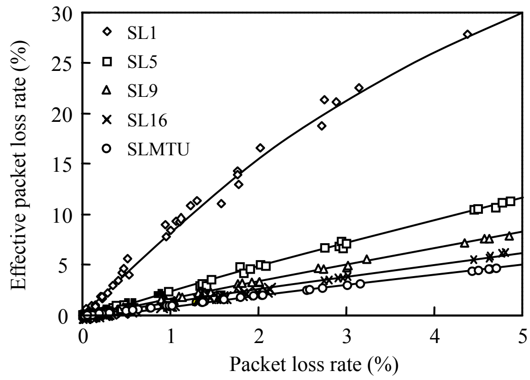

Fig.5 Effective packet loss as a function of packet loss for different slice partitioning configurations

From Fig.6a we can clearly see what can be gained by using FMO for moderate and higher loss rates. While the reconstructed *Y-PSNR* for "SLMTU-CIP" decreases below 35 dB at around 1% packet loss, "SLMTU-FMO-CIP" can maintain an average *Y-PSNR* above 35 dB for up to 3.5%~4% of random packet loss for our test material. Also note from Figs.4 and 6a that, while decreasing the compression efficiency by 0.5 dB in the error free case, the FMO-based scheme seem to outperform the pure slice partitioning based scheme at packet loss rates higher than 0.25%.

Measurements for "SLMTU-FMO-CIP-MBLINE" reveal that intra coding a single group of macroblocks in every picture gives error robustness performance comparable to that of using I pictures for random packet loss. In Fig.6b we see that the reconstructed *Y-PSNR* decreases just slightly faster than that for "SLMTU-FMO-CIP" in Fig.6a, and that it stays above the 35 dB line for packet loss rates smaller than 2.8%.

Interestingly, we see from Fig.6c that the hierarchical coding scheme using B pictures in "SLMTU-FMO-CIP-PYR" actually has reduced error robustness compared to "SLMTU-FMO-CIP", and that reconstructed *Y-PSNR* falls below 35 dB for PLR higher than 2.5%. This indicates that, without some prioritization mechanism in place to protect the all-important base layer (consisting of the I and P

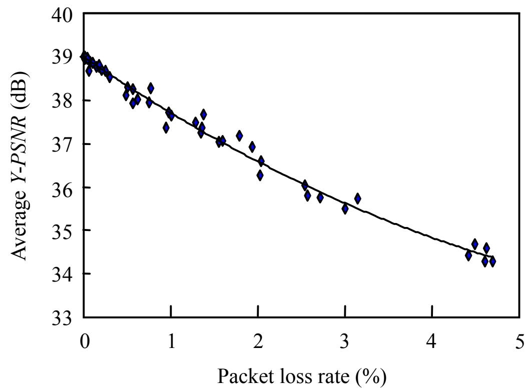

(a)

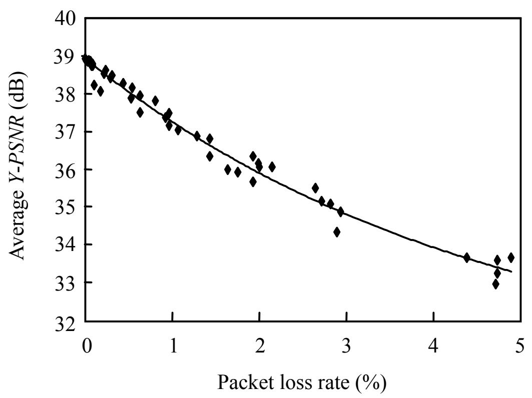

(b)

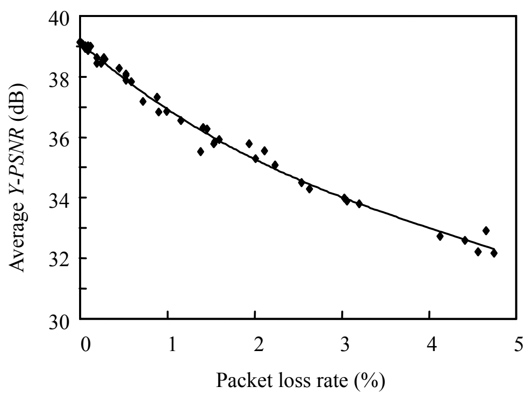

(c)

Fig.6 Average reconstructed *Y-PSNR* for the configuration (a) "SLMTU-FMO-CIP"; (b) "SLMTU-FMO-CIP-MOLIVE"; (c) "SLMTU-FMO-CIP-PYR" at different packet loss rates

pictures in the pyramid), the IPPPP coding structure of "SLMTU-FMO-CIP" is superior to the pyramid coding in "SLMTU-FMO-CIP-PYR" in terms of error robustness. In the error free case, however, pyramid coding improves the compression efficiency by 0.2 dB for our test material when FMO is used (Fig.1).

SUMMARY

In this paper, we have presented an evaluation of the error robustness of RTP-based broadcast streaming of high-quality high-definition H.264/AVC video. IP network emulation was used to introduce a varying amount of randomly distributed packet loss. A network interface monitoring card was used to capture the video packets into a PCAP trace file. Purposely-built software was developed to parse the trace file, analyze the RTP stream and assemble an H.264/AVC Annex B byte stream file, which was then decoded by JVT JM 10.1 reference software. The proposed measurement setup is a novel, practical and intuitive approach to conduct error resilience testing of real-world H.264/AVC broadcast applications.

Through a series of experiments, we evaluated some of the error resilience features of the H.264/AVC standard, and observed how they performed at packet loss rates from 0.01% to 5%. The results confirmed that an appropriate slice partitioning scheme is essential to obtain graceful degradation of application behaviour in the case of packet loss. While FMO reduces the compression efficiency by about 1 dB for our test material, reconstructed video quality is improved for loss rates of 0.25% and higher. At loss rates up to 3%~4% for our test material, flexible macroblock ordering together with slice partitioning matched to the MTU of the underlying network can give remarkable objective and subjective visual quality.

References

Calafate, C.M., Malumbres, M.P., Manconi, P., 2004. Performance of H.264 Compressed Video Streams over 802.11b Based Manets. 24th International Conference on Distributed Computing Systems Workshop, p.776-781.

DCI & ASC (Digital Cinema Initiatives and the American Society of Cinematographers), 2004. StEM Mini-movie Access Procedyre. Available at http://www.dcimovies.com/.

Hallbach, T., Olsen, S., 2004. Error Robustness Evaluation of H.264/MPEG-4 AVC. Proceedings of the International Conference on Visual Communications and Image Processing (VCIP).

Hillestad, O.I., Libak, B., Perkis, A., 2005. Performance

Evaluation of Multimedia Services over IP Networks. Proceedings of the IEEE International Conference on Multimedia and Expo (ICME). Amsterdam, the Netherlands.

ITU-T and ISO/IEC JTC-1, 2003. Advanced Video Coding for Generic Audiovisual Services. ITU-T Rec. H.264 and ISO/IEC 14496-10 (MPEG-4 part 10) AVC.

Lin, D., Morris, R., 1997. Dynamics of Early Random Detection. ACM SIGCOMM'97. Applications, Technologies, Architectures and Protocols for Computer Communication. Cannes, France.

Ostermann, J., Bormans, J., List, P., Marpe, D., Narroschke, M., Pereira, F., Stockhammar, T., Wedi, T., 2004. Video coding with H.264/AVC: Tools, performance and complexity. *IEEE Circuits and Systems Magazine*, **4**(1):7-28. [doi:10.1109/MCAS.2004.1286980]

Roth, G., Sjberg, R., Liebl, G., Stockhammar, T., Varsa, V., Karczewicz, M., 2001. Common Test Conditions for RTP/IP over 3GPP/3GPP2. ITU-T SG16, VCEG-M77.doc.

Schwarz, H., Marpe, D., Wiegand, T., 2005. Hierarchical B Pictures. Joint Video Team (JVT) of ISO/IEC MPEG and ITU-T VCEG Document JVT-PO14.

Stockhammar, T., Hannuksela, M.H., Wiegand, T., 2003. H.264/AVC in wireless environments. *IEEE Trans. on Circuits and Systems for Video Technology*, **13**(7): 657-673. [doi:10.1109/TCSVT.2003.815167]

Sullivan, G.J., Wiegand, T., 2004. Video compression—from concepts to the H.264/AVC standard. *Proceedings of the IEEE*, **86**(5):18-31.

Sullivan, G.J., Topiwala, P., Luthra, A., 2004. The H.264/AVC Advanced Video Coding Standard: Overview and Introduction to the Fidelity Range Extension. SPIE Conference on Applications of Digital Image Processing, XXVII.

Wang, Y.K., Hannuksela, M.M., Varsa, V., 2002. The Error Concealment Feature in the H.26l Test Model. Proceedings of the IEEE International Conference on Image Processing, **2**:729-732.

Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: From error visibility to structural similarity. *IEEE Trans. on Image Processing*, **13**(4):600-612. [doi:10.1109/TIP.2003.819861]

Wenger, S., 2001. Common Test Conditions for Wire-Line Low Delay IP/UDP/RTP Packet Loss Resilient Testing. ITU-T SG16, VCEG-N79r1.doc.

Wenger, S., 2003a. H.264/AVC over IP. *IEEE Trans. on Circuits and Systems for Video Technology*, **13**(7):645-656. [doi:10.1109/TCSVT.2003.814966]

Wenger, S., 2003b. Error Patterns for Internet Experiments. ITU-T VCEG, Document Q15-I16r1.

Wenger, S., Hannuksela, M.M., Stockhammar, T., Westerlund, M., Singer, D., 2005. RTP Payload Format for H.264 Video. IETF, Request for Comments, RFC 3984.

Wiegand, T., Sullivan, G.J., Bjontegaard, G., Luthra, A., 2003. Overview of the H.264/AVC video coding standard. *IEEE Trans. on Circuits and Systems for Video Technology*, **560**(576):282-300.

Yajnik, M., Moon, S., Kurose, J., Towsley, D., 1999. Measurement and Modelling of the Temporal Dependence of Packet Loss. Proceedings of INFOCOM'99. Eighteenth Annual Conference of the IEEE Computer and Communications Societies. New York, **1**:345-352.