HDTV Subjective Quality of H.264 vs. MPEG-2, with and without Packet Loss

Margaret H. Pinson, Stephen Wolf, and Gregory Cermak

***Abstract*— The intent of H.264 (MPEG-4 Part 10) was to achieve equivalent quality to previous standards (e.g., MPEG-2) at no more than half the bit-rate. H.264 is commonly felt to have achieved this objective. This document presents results of an HDTV subjective experiment that compared the perceptual quality of H.264 to MPEG-2. The study included both the coding-only impairment case and a coding plus packet loss case, where the packet loss was representative of a well managed network (0.02% random packet loss rate). Subjective testing results partially uphold the commonly held claim that H.264 provides quality similar to MPEG-2 at no more than half the bit rate for the coding-only case. However, the advantage of H.264 diminishes with increasing bit rate and all but disappears when one reaches about 18 Mbps. For the packet loss case, results from the study indicate that H.264 suffers a large decrease in quality whereas MPEG-2 undergoes a much smaller decrease.**

***Index Terms*—H.264, HDTV, MPEG-2, packet loss, quality, subjective testing, transmission errors.**

I. INTRODUCTION

THE bit-rate reductions claimed for comparable video quality of H.264 (MPEG-4 Part 10) and MPEG-2 invite scientific investigation: Are these claims valid for HDTV resolution video streams? A few months after the H.264 standard was approved by the ITU-T, the chairmen of the JVT/MPEG Ad Hoc Group on AVC Verification Test compared H.264 and MPEG-2 coding-only impairments using subjective data. Their HDTV experiment examined bit rates between 6 and 10 Mbps and showed that AVC/H.264 coding efficiency was increased by a factor of 1.7 or more over MPEG-2 in 7 of 9 cases [1]. Other researchers have used objective metrics to compare H.264 and MPEG-2 coding-only impairments for smaller image resolutions, and also demonstrated that H.264 has an improved coding efficiency [2], [3] and [4].

In addition to examining the perceptual quality of coding-only impairments, there are other questions to consider. How does H.264 compare to MPEG-2 when packet loss is present? Is there any potential disadvantage in the improved compression efficiency of H.264 in terms of robustness to dropped IP packets? To answer these questions, a controlled subjective experiment was designed and implemented utilizing HDTV scenes, video encoders, an Internet Protocol (IP) network impairment emulator, and a commercial set top box (STB). This paper describes the subjective experiment and presents comparative results for H.264 and MPEG-2. While some of these results confirm commonly held beliefs regarding H.264's compression efficiency, they show that there is a price to be paid.

II. SUBJECTIVE TEST DESIGN

The test was designed to have equal numbers of observations for each of the settings of the following variables: codec type, bit-rate, and packet loss rate (PLR). For MPEG-2, the bit-rates chosen were: 6 Mbps, 8½ Mbps, 12½ Mbps, and 18 Mbps. For H.264, the bit-rates chosen were: 2 Mbps, 3.5 Mbps, 6 Mbps, and 10 Mbps. The MPEG-2 bit-rates were chosen to evenly span the quality range from 6 Mbps to 18 Mbps, where 6 Mbps was chosen as the lowest bit rate that produced usable video from the STB and 18 Mbps produced excellent quality video that could easily fit within the allowed HDTV broadcast bandwidth of 19.4 Mbps. The H.264 bit rates were chosen to test the claim that H.264 required only one-half to one third the bit rate of MPEG-2 for comparable quality.

The MPEG-2 and H.264 encoders were commercial grade software encoders. The group of pictures (GOP) settings were the default values recommended by the encoders for encoding HDTV. For MPEG-2, this was an I-frame distance constrained between 8 and 15 frames, and a P-frame distance of 3 frames. For H.264, this was an I-frame distance of 33 frames, and a P-frame distance of 3 frames.

Streaming software was used to perform live streaming of the coded video through a network impairment emulator to a STB. Each combination of codec-type and bit-rate was streamed through two IP network conditions: no errors, and 0.02% random PLR. A 0.02% PLR was chosen to emulate a well managed, dedicated network. Higher levels of packet loss caused the STB's H.264 decoder to freeze for long periods of time (e.g., 20 seconds).

Each of the above codec / bit-rate pairs was matched with eight HDTV source sequences (to be described later). Additionally, H.264 at 17 Mbps was included for coding-only

Manuscript received June 1, 2009. This work was supported by the U.S. Department of Commerce and Verizon Labs.

Margaret H. Pinson and Stephen Wolf are with the National Telecommunications and Information Administration, Institute for Telecommunication Sciences (NTIA/ITS), Boulder, CO 80305 USA. (303-497-3579; fax: 303-497-3680; e-mail: mpinson@its.bldrdoc.gov, swolf@its.bldrdoc.gov).

Gregory Cermak is with Verizon Labs, Waltham, MA 02451 USA. (e-mail: gregory.cermak@verizonwireless.com).

using four video sequences, which allowed a coding-only comparison of quality with high bit-rate MPEG-2 coding at 18 Mbps. Operation of the H.264 codec at bit rates higher than 17 Mbps (e.g., 18 Mbps) did not produce usable output video from the STB. The MPEG-2 and H.264 encoding was performed using standard Main Profile/High Level settings for coding 1080i 30fps video (i.e., 1920 pixels by 1080 lines interlaced, 30 frames-per-second). High quality two-pass encoding was utilized. One hardware STB was used for all decoding.

This subjective test utilized twelve HDTV source video scenes that were shot with broadcast quality cameras.[1] Three different cameras were used to shoot the source sequences. One animated sequence was created digitally, and was thus entirely uncompressed. Eight of the sequences were shot using a broadcast quality camera that recorded at 100 Mbps. Three of the sequences were shot using a broadcast quality camera that recorded at 50 Mbps. One sequence was shot using an HDV camera that recorded at 25 Mbps. Using more video scenes reduces viewer boredom but increases the potential number of scene by HRC pairs (an HRC, or Hypothetical Reference Circuit, is a fixed combination of a video encoder operating at a given bit-rate, a network condition, and a video decoder or STB). The twelve scenes are summarized in Table I, ordered by increasing coding difficulty. Subjective data from this experiment were used to estimate coding difficulty.[2]

To reduce the number of scene by HRC combinations, an intelligent block design (see Table II) was utilized that would provide the quality comparison results to answer the questions posed in the introduction of this paper. Eight sequences were split into two pools of four scenes each with approximately similar characteristics (Pool A and Pool B in Table I). For example, each of these scene pools had one sequence containing vertically scrolling text on the left side of the screen (scenes WestWindEasy and Hope). Four video sequences were paired with all HRCs (i.e., the scenes included in both Pool A and Pool B). These were chosen to be sequences that appeared to evenly span nearly the full range of quality after coding. In addition, these four scenes had similar quality for MPEG-2 and H.264 (judged visually by the experiment designers). Pool A was matched with H.264 at 3½ & 10 Mbps, and MPEG-2 at 8½ & 18 Mbps. Pool B was matched with H.264 at 2 and 6 Mbps, and MPEG-2 at 6 and 12½ Mbps. These pairings of scenes and HRCs enabled 8 scenes to be used for comparing the quality of MPEG-2 with H.264 (at one half to one third the bit rate) and 4 scenes for general quality comparisons across all bit rates (see Table II). All of the video sequences were edited to 15-seconds duration after HRC generation.

TABLE I
SCENE POOL DESCRIPTIONS

Name	Description	Pool
Hope	Scrolling text on left (white on black) with flowers blowing in the wind and a goose on the right.	B
PowerDig	Close-up shot of digging equipment; irregular motion.	A, B
SnowMnt	Slow pan of snow covered mountains; high detail, low motion	B
Aspen	Aspen trees in fall color consisting of a variety of shots with leaves blowing in the wind; rapid scene cuts.	A
Kickoff	Wide-angle (distant) shot of college football kickoff, small figures and a pan.	B
Speedbag	Close and medium shots of a boxer striking a speed bag and gesturing to explain technique; fast motion.	A
FoxAndBird	Animation in style of old cartoon without color fading and scratches depicting a fox and bird; rapid scene cuts.	A
ControlledBurn	A controlled burn of a house showing firemen, flames, and smoke; rapid scene cuts and unpredictable flame motion	A, B
TouchdownDay	College football touchdown following the scoring player while running; blurred background, high motion.	A, B
RedKayak	Red kayak in a stream consisting of a variety of shots showing motion of paddle and water; difficult coding complexity and rapid scene cuts.	A, B
WestWindEasy	Scrolling text on left (white on black) and grasses moving in the wind on right; blurred background.	A
RushFieldCuts	Football crowd rushing onto a field with a variety of shots from wide to moderate zoom; unpredictable crowd motion.	B

TABLE II
EXPERIMENT DESIGN

	MPEG-2 6 Mbps & H.264 2 Mbps	MEPG-2 8½ Mbps & H.264 3½ Mbps	MPEG-2 12½ Mbps & H.264 6 Mbps	MPEG-2 18 Mbps & H.264 10Mbps	H.264 17 Mbps
Hope	x		x		
PowerDig	x	x	x	x	x
SnowMnt	x		x		
Aspen		x		x	
Kickoff	x		x		
Speedbag		x		x	
FoxAndBird		x		x	
ControlledBurn	x	x	x	x	x
TouchdownDay	x	x	x	x	x
RedKayak	x	x	x	x	x
WestWindEasy		x		x	
RushFieldCuts	x		x		

---

[1] These source sequences will be made available free of charge for research purposes at the Consumer Digital Video Library [5].

[2] The coding difficulty for a given scene was estimated as follows. For each of the eight scenes that belong to Pool B, the subjective ratings at 6 Mbps MPEG-2 and 2 Mbps H.264 HRCs (coding-only) were averaged. For each of the eight scenes that belong to Pool A, the subjective ratings at 8½ Mbps MPEG-2 and 3½ Mbps H.264 (coding-only) were averaged. The four scenes that were in both Pool A and Pool B thus had values for both the higher bitrates and the lower bitrates, The mean difference in subjective ratings for those four scenes was used to shift the Pool A ratings down (i.e., extrapolate the missing subjective ratings for 6 Mbps MPEG-2 and 2 Mbps H.264 HRCs, coding-only). These final subjective ratings were rank-sorted from best quality to worst quality, as given in Table I and Table II. This process produced an approximate ranking, where changes in magnitude of the subjective rating from one scene to the next are not reflected.

Fig.1 depicts one representative video frame from each source sequence.

Hope

PowerDig

SnowMnt

Aspen

Kickoff

Speedbag

FoxAndBird

ControlledBurn

TouchdownDay

RedKayak

WestWindEasy

RushFieldCuts

Fig. 1. One representative frame from each of the 12 source video scenes.

In summary, the overall experiment was designed using the following rules:

- Source videos were included once in the overall experiment to anchor the high end quality expectations.
- Each HRC (except H.264 at 17 Mbps) was matched with 8 sources, such that the MPEG-2 HRCs could be directly compared with corresponding H.264 HRCs operating at approximately one-half to one-third the bit rate.
- Four source videos that spanned nearly the full range of coding difficulty were matched with all HRCs to allow general comparisons across bit-rates.
- Each combination (scene, bit-rate, coding algorithm) shown in Table II was tested for both coding-only impairments, and with a random PLR of 0.02% (except H.264 at 17 Mbps, which was tested only for coding-only impairments).
- One commercial-grade STB was used for all HRCs.

The subjective test was performed using the single stimulus Absolute Category Rating (ACR) methodology as described in [6], where viewers rated each sequence on a scale of: excellent, good, fair, poor, and bad. These words are typically mapped to the numbers 5, 4, 3, 2, and 1 respectively. Subjective ratings were gathered for 24 naïve viewers. The subjective test contained a total of 144 video sequences: 12 source and 132 processed.

III. DATA ANALYSIS

A. *Coding-only*

The source video sequences' average mean opinion score (MOS) was 4.46, which lies approximately halfway between "excellent" and "good." Viewers in most subjective tests tend to be reluctant to score any video sequence as "excellent." Thus, some drop in quality from the theoretical maximum of 5.0 is expected. The source sequences' MOS ranged between 4.04 and 4.68. The quality of the camera recording (from 25 Mbps to 100 Mbps) did not appear to be related to the source sequence MOS. For example, the ControlledBurn sequence recorded at 25 Mbps had a MOS = 4.50. Variation of MOS for the source scenes may be due to viewer preferences for certain types of scene content.

Fig. 2 and Fig. 3 plot the change in MOS distribution for H.264 and MPEG-2 as a function of bit rate for the coding-only case. H.264 is plotted in blue with solid lines, MPEG-2 is plotted in red with dotted lines, and the source video sequences are plotted in black. Fig. 2 presents the mean MOS computed from the four sequences that were paired with all the HRCs, while Fig. 3 presents box plots that show the data distribution for all the sequences. The bottom and top of the box indicate the 25th and 75th percentile, respectively. The dark bar in the middle of the box identifies the mean MOS for that HRC, averaged across all scenes ($\mathrm{MOS}_{\mathrm{HRC}}$). The range spanned by the minimum and maximum MOS are drawn as a bar extending below and above the box, respectively. In some cases, the minimum or maximum fall very close to the box, and is thus not visible on the plot (e.g., the 6 Mbps MPEG-2 maximum and 75% points are approximately equal). Most of the data points in Fig. 3 represent eight sequences passed through the HRC at the bit-rate shown on the X-axis. There are two exceptions: 17 Mbps H.264 includes only four scenes, and the source video includes all twelve scenes. Note that while the source is plotted near 30 Mbps for visual convenience, the actual bit-rate of most sequences is 50 to 100 Mbps.

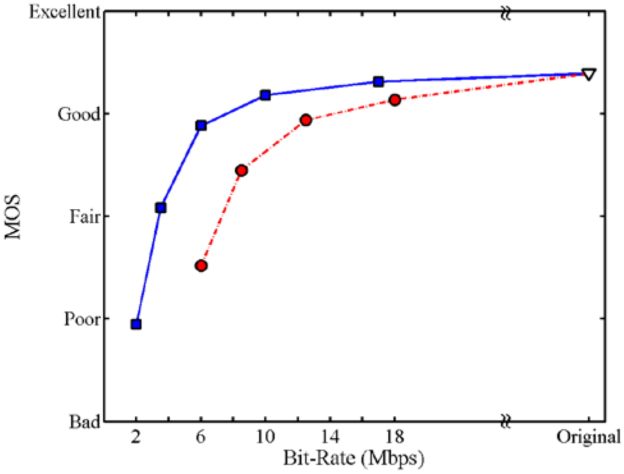

Fig. 2. For coding-only impairments, quality comparison of H.264 (solid blue line with squares) and MPEG-2 (dotted red line with circles) as a function of bit rate, using the same four scenes.

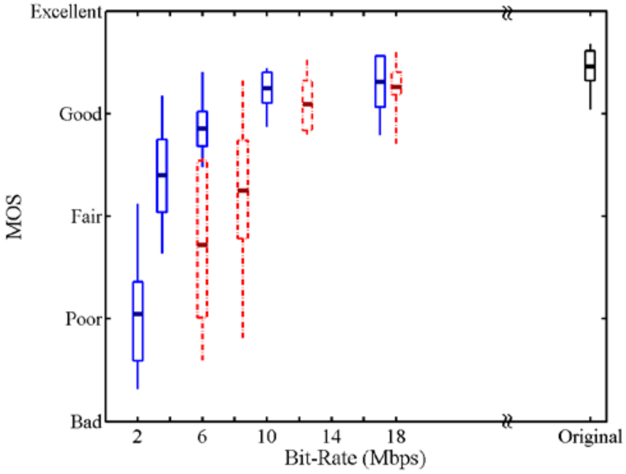

Fig. 3. For coding-only impairments, quality comparison of H.264 (solid blue) and MPEG-2 (dotted red) as a function of bit rate. Box-plot identifies minimum, 25%, mean, 75%, and maximum MOS, using all available scenes.

The student T-test indicates that H.264 at 3.5 and 10 Mbps are statistically equivalent to MPEG-2 at 8½ and 18 Mbps respectively. However, H.264 at 2 and 6 Mbps are significantly worse than MPEG-2 at 6 and 12½ Mbps respectively. This partially supports the general rule of thumb that H.264 produces a quality equivalent to MPEG-2 while using approximately one-half the bit-rate.

B. Packet Loss

Fig. 4 compares H.264 with coding-only impairments (solid green line) to H.264 with a 0.02% random PLR (solid blue line with squares) and also MPEG-2 with coding-only impairments (dotted purple line) to MPEG-2 with a 0.02% random PLR (dotted red line with circles). This figure uses only the four source video sequences that were paired with all HRCs, so that trends between the measured bit rates may be inferred. Fig. 5 compares H.264 with 0.02% random PLR (solid blue) to MPEG-2 with 0.02% random PLR (dotted red) using all eight scenes for each HRC. The box plots are defined as for Fig. 3. These two figures show that H.264 is more sensitive to dropped packets than is MPEG-2. For a PLR of only 0.02%, the quality advantages enjoyed by H.264 disappeared at bit rates above 8-10 Mbps. A similar observation was noted in a prior subjective experiment that utilized standard definition television sequences and a different set of H.264/MPEG-2 codecs [7]. The improved compression efficiency of H.264 versus MPEG-2 comes at a price, namely, greater vulnerability of visual quality to dropped packets.

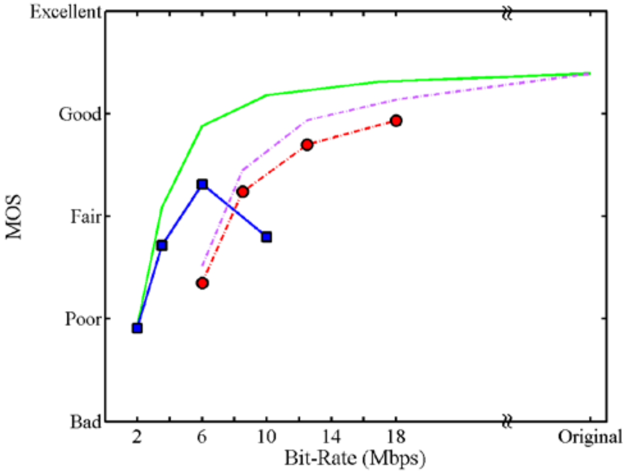

Fig. 4. Quality comparison of H.264 with 0.02% random PLR (solid blue line with squares) with H.264 coding-only (solid green line); and MPEG-2 with 0.02% random PLR (dotted red line with circles) with MPEG-2 coding-only (dotted purple line), using the same four scenes.

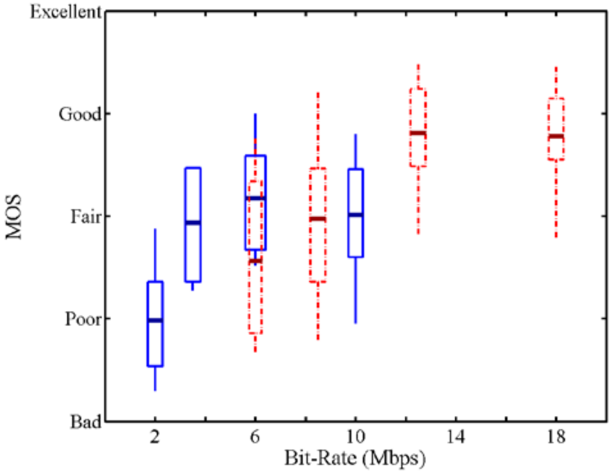

Fig. 5. Quality comparison of H.264 (solid blue) and MPEG-2 (dotted red) for a PLR of 0.02%, using all available scenes.

C. Visual Appearance and Propagation

Although all processed video sequences (PVSs) used the same STB, H.264 and MPEG-2 produced visually different responses to dropped packets. To characterize the STB’s response to dropped packets, one of the experiment designers watched each PVS twice, in real time, and counted the number of impairments that appeared to be caused by dropped packets in each PVS. The PVSs were then loaded into a video editing program that was capable of single frame stepping. This allowed the duration, type, and severity of each observed impairment to be tabulated. More rigorous experiments of this nature are presented in [8], however those experiments did not compare H.264 with MPEG-2.

The visual duration of dropped packets was more than twice as long for the H.264 decoder as for the MPEG-2 decoder. The H.264 decoder displayed artifacts from 0.033 to 1.10 seconds (1 to 33 frames), with an average duration of 0.63 seconds. The MPEG-2 decoder displayed artifacts from 0.033 to 0.60 seconds (1 to 18 frames), with an average duration of 0.25 seconds. These differences in duration are likely caused by the GOP structure[3]. The MPEG-2 coder placed I-frames at intervals of 3 to 15 frames (i.e., 0.5 seconds maximum), while the H.264 coder placed key-frames 33 frames apart (i.e., 1.10 seconds).

The visual impact of dropped packets was also different. The H.264 decoder typically responded by frame freezing the video stream (e.g., 0.6 second pause followed by playing with some loss of video). Because the H.264 decoder was slow to recover from packet loss, these freezes seemed to have a major impact on the viewing experience. The H.264 decoder once responded to a dropped packet by displaying a magenta overlay on most of the picture for a full 1 second.

The MPEG-2 decoder had a variety of visual responses that combined two strategies: (1) brief frame freezes, and (2) distorting part of the image, while continuing to play the rest of the image without packet loss visibility. Most of the time (approximately 80%), the artifact displayed by the MPEG-2 decoder in response to packet losses had a minor impact visually. These minor responses included 1-frame and 2-frame pauses, artifacts that impacted less than 5% of the screen, and artifacts that were partially masked by scene cuts, coding artifacts, or scene content. The MPEG-2 decoder's frame freezes ranged in duration from 0.033 to 0.27 seconds, with an average pause duration of 0.066 seconds.

Fig. 6 shows the average number of visual errors observed in each PVS, as a function of algorithm and bit rate. Two important trends can be observed. The first trend is that the number of visual errors increases with increasing bit rate. This happens because the packet size and packet loss rate are held constant. Thus, a packet loss is more likely to occur in each second as the bit rate increases. This phenomenon helps to explain why a fixed PLR has a greater effect on quality at higher bit rates and may even cause higher bit rates to have lower perceived quality levels than lower bit rates (for example, in Fig. 4 compare the average quality level of H.264 at 10 Mbps with 6 Mbps, both with a PLR of 0.02%). The logical conclusion is that under packet loss scenarios, there is a point where transmission bits can be better utilized to increase robustness to packet loss rather than encoding accuracy.

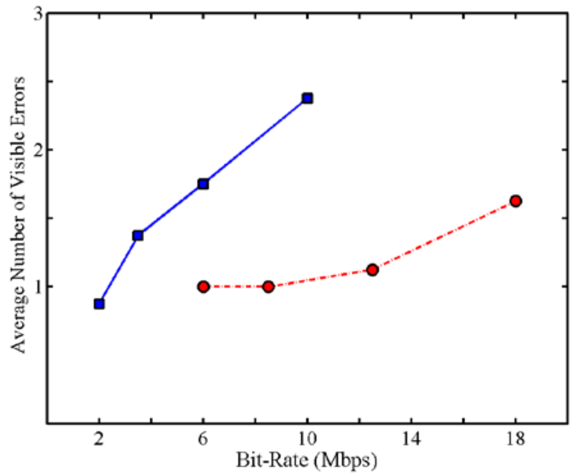

Fig. 6. A rough count of the per-clip frequency of visual H.264 impairments (solid blue line with squares) and MPEG-2 impairments (dotted red line with circles) for a PLR of 0.02%.

The second trend shown in Fig. 6 is that the H.264 decoder was more likely to display a visible artifact in response to packet loss than the MPEG-2 decoder. H.264 is a more efficiently coded information stream than MPEG-2 – that is, each packet represents more video information.[4] So perhaps the decoder is less able to compensate for a missed packet. On the other hand, H.264 is a relatively new codec compared to MPEG-2, so improvements in its response to dropped packets may be forthcoming. Either way, forward error correction and/or error concealment strategies will be important features for H.264 decoders if they are to be deployed anywhere other than very well managed networks.

D. Diminishing Returns with Increased Bit Rate

Fig. 7 and Fig. 8 show the MOS and 95% confidence intervals for two video sequences that were matched with all HRCs: PowerDig and RedKayak. By examining these plots, we can discern scene dependent responses of both MPEG-2 and H.264 and observe a "diminishing returns" relationship between judged video quality and bit rate: the improvement in quality becomes smaller as bit rate increases and the point where quality improvements are negligible depends upon the scene coding difficulty (e.g., approximately 8½ Mbps for PowerDig, 12½ Mbps for TouchdownDay, and above 18 Mbps for RedKayak and ControlledBurn).

---

[3] The GOP begins with an I-frame for MPEG-2 and a key-frame for H.264/MPEG-4 AVC. I-frames and key-frames stop the propagation of decoding errors, because the entire video frame is decoded without reference to previous data.

[4] This same trend might occur when examining one coder and either increasing packet size or decreasing bit rate. In both cases, more information is transmitted in each packet.

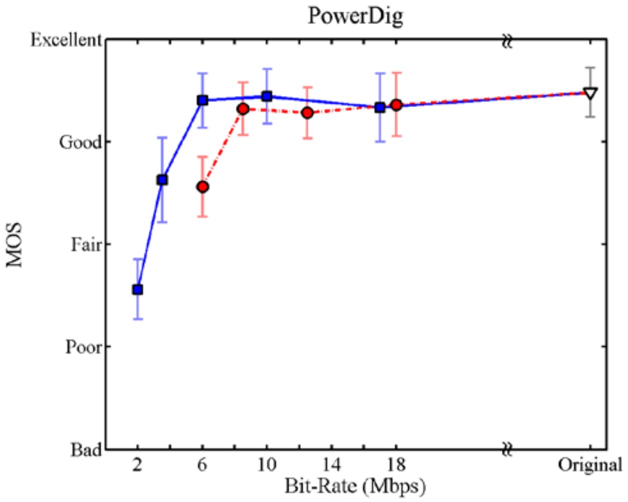

Fig. 7. Quality comparison of video sequence PowerDig coded with H.264 (solid blue line with squares) and MPEG-2 (dotted red line with circles).

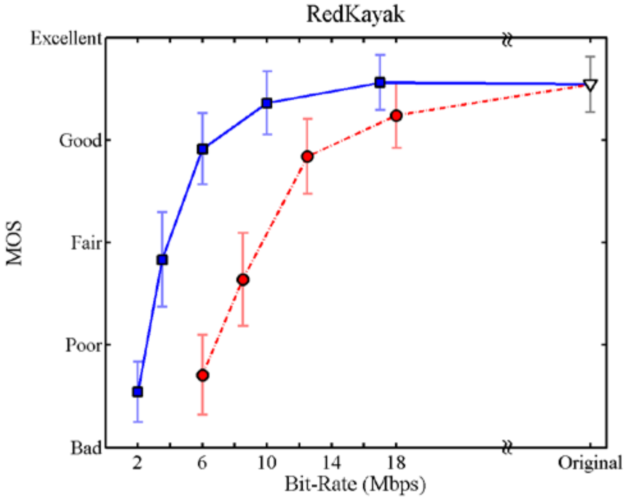

Fig. 8. Quality comparison of video sequence RedKayak coded with H.264 (solid blue line with squares) and MPEG-2 (dotted red line with circles).

Fig. 9 plots the average quality for H.264 (solid blue with squares) and MPEG-2 (dotted red with circles) as a function of bit rate, with the clips separated into two sets of six scenes each: easy-to-code, and hard-to-code. Here, video clips were separated by the MOS received at the lowest two bit rates of each coder type (see footnote 2). Thus, the easy-to-code set contains the first six scenes listed in Table I, and the hard-to-code set contains the last six scenes. These two plots show that the quality of MPEG-2 becomes increasingly dependent on characteristics of the source sequence as bit rate falls. H.264 is less impacted by characteristics of the source sequence. These same phenomena can be observed in Fig. 3 by observing the quality range spanned by the 25th to 75th percentile as bit rate drops. So, as bit rate drops, one would expect to see less quality variation in H.264 than in MPEG-2. Thus, H.264 may have fewer opportunities for statistical multiplexing than MPEG-2.[5]

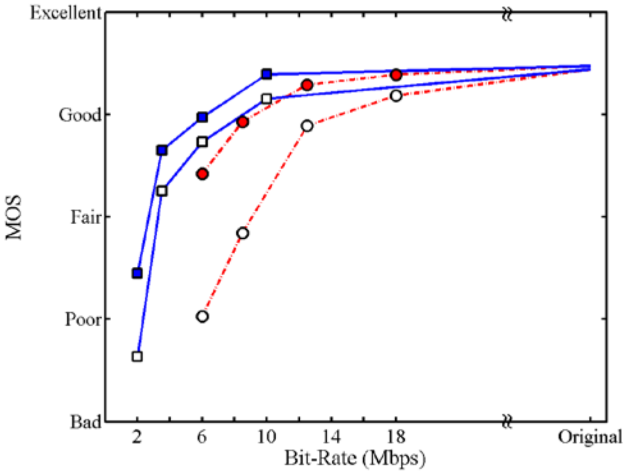

Fig. 9. For coding-only impairments, comparing quality of easy-to-code scenes (colored circles & squares) with hard-to-code scenes (white circles & squares). H.264 is plotted in solid blue lines with squares and MPEG-2 is plotted in dotted red lines with circles.

Fig. 10 shows the average drop in quality when the coded bit streams have a random PLR of 0.02%, comparing the same groupings of easy-to-code and hard-to-code scenes presented in Fig. 9. For MPEG-2, the quality impact of dropped packets is similar for easy-to-code and hard-to-code scenes. This impact increases slightly as bit rate increases. For H.264, the quality impact of dropped packets is more pronounced as bit rate increases. The easy-to-code and hard-to-code scenes also seem to respond differently to dropped packets, yet no definitive pattern is evident.

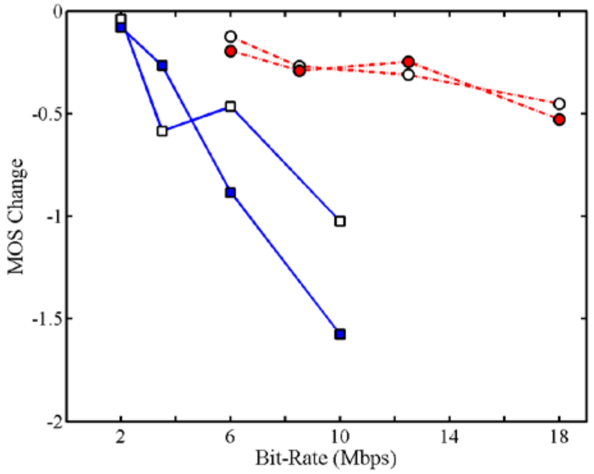

Fig. 10. Comparing drop in quality for easy-to-code scenes (colored circles & squares) with hard-to-code scenes (white circles & squares) when dropped packets are introduced. H.264 is plotted in solid blue lines with squares and MPEG-2 is plotted in dotted red lines with circles.

---

[5] As used here, statistical multiplexing refers to dynamically allocating bits across many video channels in such a manner so as to achieve an aggregate reduction in transmission bits while maintaining the video quality of the individual channels.

IV. CONCLUSIONS

The subjective test results are consistent with the commonly held belief that H.264 coding provides quality similar to MPEG-2 coding at approximately half the bit rate. However, the advantage of H.264 diminishes as bit rate increases, such that at sufficiently high bit rates (above 18 Mbps), there is very little difference between MPEG-2 and H.264.

However, the results of this experiment also demonstrate that H.264 quality drops steeply for even low packet loss rates (0.02%), while MPEG-2 quality drops by much less. In an environment where dropped packets are expected to occur, the bit-rate advantage of H.264 over MPEG-2 disappears above about 8-10 Mbps. This sensitivity appears to be partially caused by H.264's long GOP structure (chosen to diminish bandwidth), which causes the visual impact of transmission errors to propagate longer. The results of this paper suggest that in packet loss environments, the transmission bits would be better used in H.264 systems to provide forward error correction or another robust response to packet loss rather than improved coding accuracy. Alternately, better error concealment techniques could be deployed for H.264. This is an area for future study, when H.264 codecs become available that implement improved error concealment and/or forward error correction.

REFERENCES

[1] T. Oelbaum, V. Baroncini, T. K. Tan, and C. Fenimore, "Subjective quality assessment of the emerging AVC/H.264 video coding standard," in *Proc. International Broadcasting Convention (IBC)*, Amsterdam, Netherlands, 2004.

[2] N. Kamaci and Y. Altunbasak, "Performance comparison of the emerging H.264 video coding standard with the existing standards," in *Proc. 2003 International Conference on Multimedia and Expo.*, vol. 1, July 6-9, 2003.

[3] O. Nemcic, M. Vranjes, and S. Rimac-Drlje, "Comparison of H.264/AVC and MPEG-4 Part 2 coded video," in *Proc. Electronics in Marine (ELMAR), 2007*, Sep. 12-14, 2007.

[4] P. Begovic, N. Behlilovic, and A. Mastilovic, "Parameters comparison of stored audio/video contents, transcoded using different compression algorithms," in *ELMAR, 2008. 50th International Symposium*, Sep. 10-12, 2008, vol. 1.

[5] The Consumer Digital Video Library (CDVL) will be available at www.cdvl.org. This web site is scheduled to be deployed by the fourth quarter of 2009.

[6] ITU-T Recommendation P.910, "Subjective Video Quality Assessment Methods for Multimedia Applications," Geneva, 2007 (available at www.itu.int).

[7] M. H. Pinson, S. Wolf and R. B. Stafford, "Video performance requirements for tactical security video applications," in *Proc. 2007 IEEE Conference on Technologies for Homeland Security*, Woburn, MA, May 16-17, 2007. Available: http://www.its.bldrdoc.gov/pub/n3/video/ieee_07.pdf

[8] A. Reibman and D. Pool, "Predicting packet-loss visibility using scene characteristics," *in Proc. Packet Video 2007*, Nov. 12-13, 2007.