Tomorrow's World

High Efficiency Video Coding: the Future of Video Compression?

Ken McCann, Chairman TM-AVC

Ken McCann is a director and co-founder of ZetaCast, an independent technology consultancy company specializing in digital TV and related areas.

The current generation of video compression technology (H.264/AVC) provides about a factor of two improvement in compression efficiency compared to the previous generation (MPEG-2). This ability to get the same quality of video at half the bitrate has provided an important technology enabler for the growing number of high definition (HD) television services that we see today. However, the development of the next generation video compression standard is already starting to take shape. This promises to deliver another factor of two improvement, providing a further boost to the number of HD services that can be transmitted and potentially paving the way to new ultra high definition (UHD) services.

This new video compression standard, known as High Efficiency Video Coding (HEVC), is being jointly created by ISO/IEC MPEG and ITU-T VCEG, the same two standardization bodies whose previous collaboration resulted in both MPEG-2 and H.264/AVC. A wide range of video resolutions will be covered, from low resolution right up to the UHD resolutions of 4K x 2K and 8K x 4K.

Call for Proposals

The first step towards creating HEVC was the launch of a joint Call for Proposals in January 2010. A total of 27 proposals were received – quite an impressive response, given the amount of work that was required to prepare a proposal. Each proposal was tested using 18 video sequences, 5 bitrates and 2 constraint sets. The 18 sequences covered five classes of video resolution, ranging from quarter WVGA up to 4K x 2K. The 5 bitrates were chosen to provide a range of video qualities from highly stressed to good quality encoding. For example, the 1080p 50 and 60 Hz sequences were tested at bitrates ranging from 2 to 10 Mbit/s. Given the target of a factor of two increase in coding efficiency, this roughly corresponds to the quality of a current H.264/AVC encoder operating at 4 to 20 Mbit/s. Two constraint sets were defined for the encoding - the first, typical of broadcasting applications with limits to the time for a channel change, the second, typical of conversational services with limited encode/decode delay.

Two different types of quality measurements were performed, using a standard software implementation of H.264/AVC as a lower anchor. Firstly, objective measurements were made using the Peak Signal-to-Noise Ratio (PSNR), i.e. the ratio between the energy in the encode/decode error and the energy in the original picture. PSNR is a commonly used and convenient method of giving an approximate indication of the likely video quality. However, despite many years of trying to come up with an accurate objective measurement, the only way to really determine video quality remains the rather awkward and expensive process of running subjective tests with real people.

Formal subjective testing was therefore carried out for all test cases, apart from the two UHD sequences. Around 23,000 video clips were tested; this appears to be the largest subjective quality testing effort ever carried out in the history of video compression. The test was organized in 134 sessions of approximately 20 minutes each, using a total of 850 test subjects over a one month period. This was really too much for any single laboratory to cope with, so the burden of testing is being shared between three sites: FUB in Rome, EBU in Geneva and EPFL in Lausanne.

Results of HEVC Call for Proposals

The technical details and results of the HEVC Call for Proposals were analyzed during the first meeting of the new Joint Collaborative Team on Video Coding (JCT-VC) in April 2010. The basic architecture for all of the proposed algorithms was a hybrid motion-compensated block transform, the same basic architecture used by MPEG-2 and H.264/AVC. Within this traditional architecture many innovative new tools were proposed, with larger block sizes and a

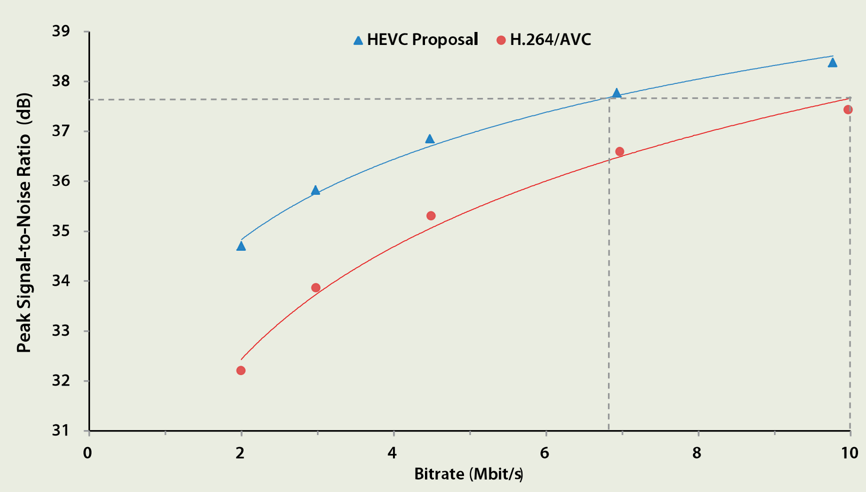

Fig. 1 - 1080p 60Hz sequence "BasketballDrive" encoded at bitrates of 2, 3, 4.5, 7 and 10 Mbit/s. Reproduced with the kind permission of Samsung Electronics Co. Ltd.

more flexible block structure being common themes.

The test results were very encouraging, both using objective and subjective measures. The graph in Figure 1 shows a typical example of an objective rate distortion curve, using results obtained from the individual proposal that gave the best overall results. This example is a 1080p 60Hz sequence "BasketballDrive" encoded at bitrates of 2, 3, 4.5, 7 and 10 Mbit/s using the encoding constraints typical of broadcasting applications. The red dots represent the results obtained using a software model of H.264/AVC, the blue triangles represent the results from the proposal. The vertical axis is PSNR, with a higher number indicating better results.

At 10 Mbit/s, the H.264/AVC software gives a PSNR value of about 37.6 dB. The same PSNR is obtained by the proposal at approximately 6.6 Mbit/s – about 34 percent less. The bitrate saving increases to about 48 percent if the same analysis is done when H.264/AVC is operating at 5 Mbit/s. This pattern is fairly typical: a significant bitrate saving for good quality video which increases further if the comparison is done at a lower quality point. Figure 2 shows the percentage bitrate savings implied by the objective measures for all five 1080p sequences, averaged across the five test points. The overall average bitrate reduction is 43 percent.

However, what really matters is the bitrate reduction at the same subjective video quality, which is shown in Figure 3 for the same five 1080p sequences. Based on these subjective results, the overall average bitrate reduction was 61 percent. As with the objective measures, the savings are largest at the lowest bitrates – if we were to focus purely on the results at bitrates that gave broadcast quality video then the average bitrate savings from the proposal would be closer to 50 percent.

Sequence	Bitrate Reduction (%)	Average (%)
Kimono	44	43
ParkScene	33	
Cactus	41	
BasketballDrive	44	
BQTerrace	54	

Fig. 2 - Bitrate Reduction implied by PSNR Measurements

Sequence	Bitrate Reduction (%)	Average (%)
Kimono	60	61
ParkScene	60	
Cactus	57	
BasketballDrive	57	
BQTerrace	70	

Fig. 3 -Bitrate Reduction implied by Subjective Results

These results relate to only one of the 27 proposals; several other proposals also performed well. At its April meeting, JCT-VC decided not to adopt a "winner takes all" approach and instead defined a "Test Model under Consideration", which combined specific elements from seven of the leading proposals.

The first formal HEVC test model, "HM1", was defined in October 2010. This was largely a selection of the better-performing tools from the "Test Model under Consideration", following a detailed analysis of their individual contributions. HM1 will be used as the basis of further "Core Experiments" and the HM can be expected to evolve through several versions before the standard is completed. Throughout this process it is anticipated that there will be further incremental improvements to the compression performance.

The plan is to complete the new HEVC standard by January 2013. It is then intended to be published by ISO/IEC as MPEG-H and by ITU-T as H.265.

Will the HEVC standard be a success in the market?

Before looking at the future, it's often instructive to consider the past. Figure 4 shows the major video compression standards that have been developed over the past 20 years. Two of these have been particularly significant: MPEG-2, the standard that launched the digital broadcasting revolution, and H.264/AVC, which has had such a large influence since its creation that it has been difficult for alternatives to grow in its shadow. It is interesting to note that, like HEVC, both of these standards were the result of collaboration between MPEG and VCEG.

So how will HEVC fare? In general, we can observe that improvements in video compression come in two forms. Firstly, there is a gradual evolution of encoders giving greater efficiency within a specification, where existing decoders can continue to be used. Secondly, there are occasional moments when there is a revolution caused by a change of algorithm, where new decoders are needed. It seems that the market is prepared to consider such a revolutionary change of algorithm roughly once a decade, provided that it can be justified by about a factor of two improvement in coding efficiency.

So HEVC appears to have it about right in terms of both timing and functionality. If history repeats itself, we might expect to see publication of HEVC in 2013, addition to DVB in 2014 and services ready to launch in 2015, perhaps with Ultra HD as a pioneer launch service. Watch this space!

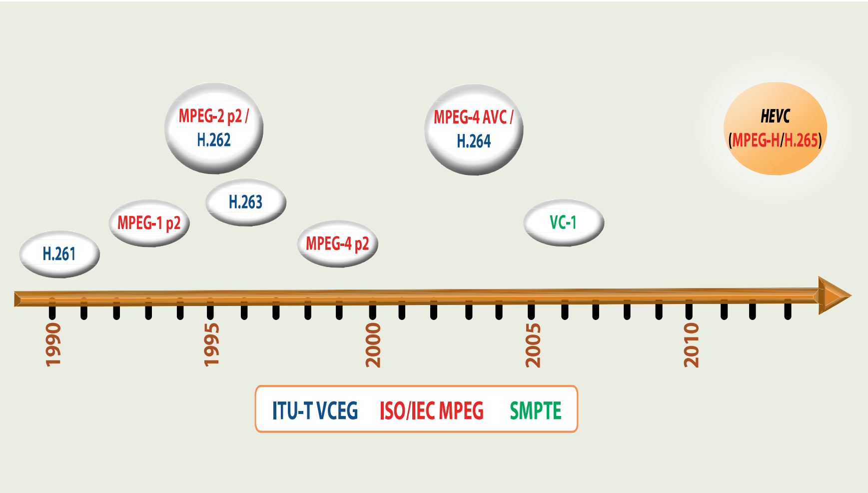

Fig. 4 - Codec Timeline