Ericsson Review
The communications technology journal since 1924
2013 • 6

Next generation video compression

April 24, 2013

ERICSSON

Next generation video compression

MPEG and ITU have recently approved a new video-compression standard known as High Efficiency Video Coding (HEVC), or H.265, that is set to provide double the capacity of today's leading standards[1].

PER FRÖJDH, ANDREY NORKIN AND RICKARD SJÖBERG

Requiring only half the bitrate of its predecessor, the new standard will significantly reduce the need for bandwidth and expensive, limited spectrum. HEVC will enable new video services to be launched, and the first applications that are likely to appear will be for mobile devices and OTT applications, followed by TV – and in particular ultra HD television (UHDTV).

State-of-the-art video compression can reduce the size of raw video by a factor of about 100 without any noticeable reduction in visual quality. Estimates indicate that compressed real-time video accounts for more than 50 percent of current network traffic[2], and this figure is set to rise to 90 percent within a few years[3]. New services, devices and changing viewing patterns are among the factors contributing to this growth, as is increased viewing of traditional TV and video-streaming services, such as Netflix, YouTube and Hulu, on a range of devices – from phones and tablets to PCs and home-entertainment systems. As HD shifts from luxury to commodity, it will soon be challenged by UHD, which offers resolutions up to 16 times greater.

Making standards

Most video viewed by subscribers today has been digitized and reduced in size through the application of a compression standard. The more popular include the H.26x series from ITU and the MPEG-x series from ISO/IEC. First published in 1994, the MPEG-2 standard, also known as H.262, played a crucial role in the launch of digital-TV services as it enabled the compression of TV streams to fit the spectrum available. This is also the standard used to compress movies onto a DVD. The H 264 standard (also known as MPEG-4 AVC), published in 2003, has provided the best compression efficiency to date, and is currently the most widely used video-compression codec. It has been successfully incorporated into most mobile devices, and is the best way to reduce the size of video carried over the internet. It is the preferred format for Blu-ray discs, telepresence streams and, most notably, HDTV.

Now imagine a codec that is twice as efficient as H.264. This was the target set by MPEG and ITU in 2010, when they embarked on a joint standardization effort that three years later delivered HEVC/H.265[4,5].

The new codec offers a much more efficient level of compression than its predecessor H.264, and is particularly suited to higher-resolution video streams, where bandwidth savings with HEVC are around 50 percent. In simple terms, HEVC enables a network to deliver twice the number of TV channels. Compared with MPEG-2, HEVC can provide up to four times the capacity on the same network.

Like most standards, the MPEG and ITU video codecs have been developed in a collaborative fashion involving many stakeholders – manufacturers, operators, broadcasters, vendors and academics. Ericsson has been an active participant in video standardization for more than 15 years, and was closely involved in HEVC.

Throughout the development of the standard, Ericsson has led several of the core experiments, chaired ad-hoc working groups and contributed significantly to the development of the technology behind the codec. Our greatest expertise lies in the areas of the deblocking filter[6] and in reference picture management[7].

BOX A Terms and abbreviations

AVC	advanced video coding
CABAC	context-adaptive binary arithmetic coder
CTU	coding-tree unit
CU	coding unit
fps	frames per second
HD	high definition; often refers to 1280 x 720 or 1920 x 1080 pixels
HEVC	High Efficiency Video Coding
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization
ITU	International Telecommunication Union
MPEG	Moving Picture Experts Group
OTT	over-the-top
SAO	sample adaptive offset
UHD	ultra high definition: often refers to 3840 x 2160 (4K) or 7680 x 4320 (8K) pixels
WPP	wavefront parallel processing

Concepts that create efficiency

One of the primary target areas for HEVC compression is high-resolution video, like HD and UHD. The statistical characteristics of such video streams tend to be different from lower-resolution content: frame sizes are larger, and frame rates and perceived quality are higher – imposing tough requirements on compression efficiency, as well as on the computational complexity of the encoding and decoding processes.

As smartphone and tablet architectures go multi-core, the ability to take advantage of parallel processing is key when it comes to the efficient compression of high-resolution content. All of these points have been taken into consideration during the development of the new standard.

The hybrid block-based coding used by the new codec is the same as the one used in earlier video-coding standards. To encode content, video frames are divided into blocks that are coded individually by applying prediction – based either on neighboring blocks in the same picture (intra prediction) or from previously coded pictures (motion estimation/compensation). The difference between the predicted result and original video data is subsequently coded by applying block transforms and quantization. In this way, a block can be represented by just a few non-zero coefficients. Quantized transform coefficients, motion vectors, prediction directions, block modes and other types of information are encoded with lossless entropy coding. Hybrid block-based coding is illustrated in Figure 1.

To ensure the highest level of compression efficiency, and support for parallel processing, some parts of HEVC have been significantly modified compared with previous generations of hybrid block-based codecs. For most of the previous MPEG-x and H.26x codecs, the largest entity that could be independently encoded was a macroblock (16 × 16 pixels). For HEVC, the picture is split into coding-tree units (CTUs) with a maximum size of 64 × 64 pixels. Every CTU is the root of a quadtree, which can be further divided into leaf-level coding units (CUs), as illustrated in Figure 2. The CTUs are coded in raster scan order, and each unit can itself »

FIGURE 1 Simplified HEVC encoder diagram

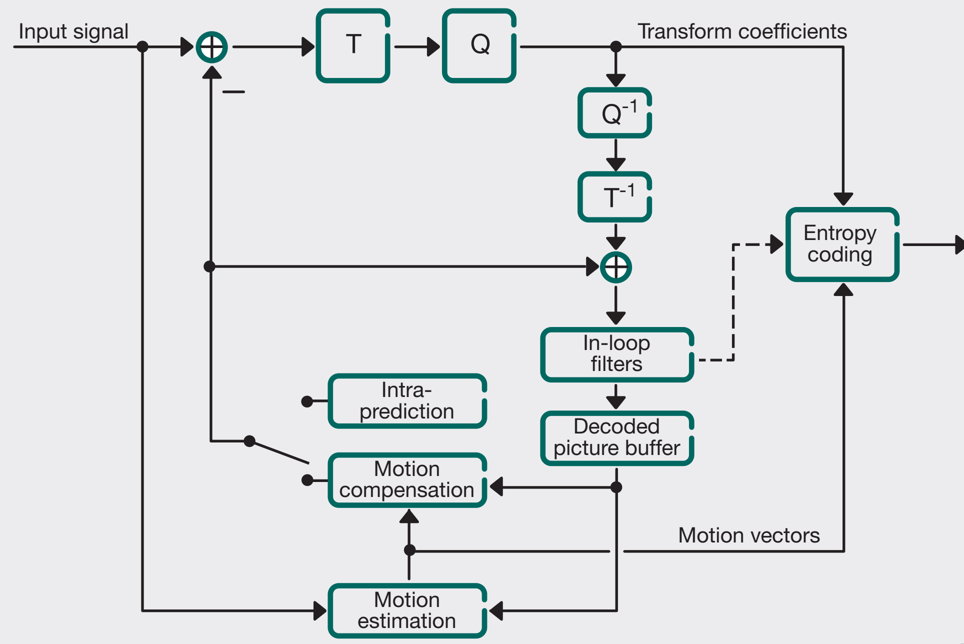

FIGURE 2 Example of the coding-tree unit structure in HEVC

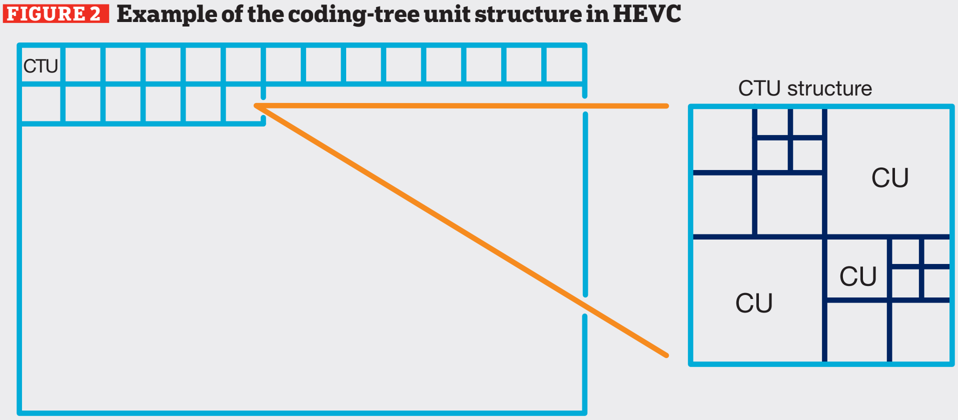

FIGURE 3 **Possible motion prediction partitions in HEVC**

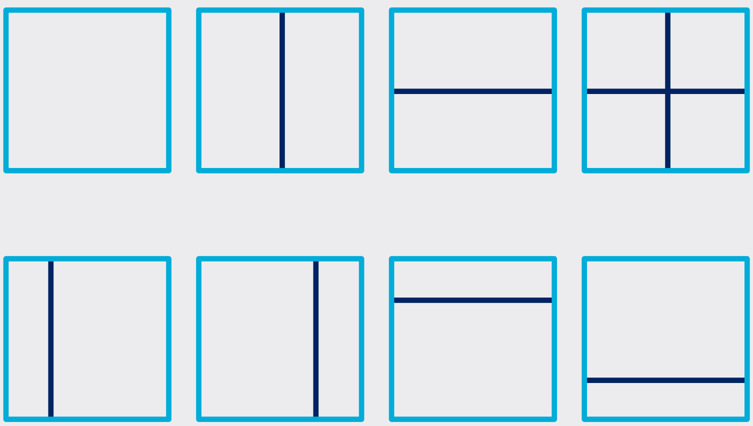

Asymmetric motion partitions are shown in the bottom row. Only square partitions are allowed for intra prediction

contain a quadtree structure. Each CU contains one or more prediction partitions that are predicted independently of each other. A CU is also associated with a transform quadtree that compresses the prediction residual and has a structure similar to that of a CTU – as shown in Figure 2.

Partitions for motion prediction can form square or rectangular shapes, which is also the case with earlier standards. HEVC also supports something called asymmetric motion partitioning, which can split the CU into prediction units of unequal width or height, as illustrated in **Figure 3**.

The size of the prediction blocks in HEVC can therefore vary from 4 × 4 samples up to 64 × 64, while transform sizes vary from 4 × 4 to 32 × 32 samples. Large prediction blocks and transform sizes are the most efficient way to encode large smooth areas, whereas smaller prediction blocks and transforms can be used to achieve precision in areas that contain finer detail.

The HEVC specification covers more intra-prediction modes than H.264, including a planar mode to approximate a surface from neighboring pixels, a flat mode and 33 angular prediction modes. Motion-compensated prediction for luma transform blocks is performed with up to quarter-pixel precision, whereas motion compensation for color components is performed with one-eighth-of-a-pixel precision. Interpolation for fractional pixel positions uses 8-tap filters for luma blocks and 4-tap filters for color.

In HEVC there is a single entropy coder for low-level data. This is the context-adaptive binary arithmetic coder (CABAC), which is similar to the one used in H.264, but modified to facilitate parallel processing. Higher-level information, such as sequence parameters, is encoded with variable-length or fixed-length encoding.

HEVC defines two in-loop filters: a deblocking filter and a sample adaptive offset (SAO) filter. The latter is applied to the output of the deblocking filter, and increases the quality of reference pictures by applying transmitted offsets to samples that fulfill certain criteria. In-loop filters improve the subjective quality of reconstructed video as well as compression efficiency. Deblocking filtering in HEVC is less complex than that of H.264, as it is constrained to an 8 × 8 block grid. This constraint, together with filtering decisions and operations that are non-overlapping between two boundaries, simplifies multi-core processing.

Parallel processing

To make the most of the increasingly widespread use of multi-core processors, plus the ever-growing number of cores used in consumer-class processors, significant attention was paid to the parallelization characteristics of video encoding and decoding when designing HEVC. As it is computationally more complex than its predecessor, maximizing parallelization has been a key factor in making HEVC an efficient real-time encoding and decoding solution.

Several HEVC tools have been designed for easy parallelization. The deblocking filter can be applied to 8 × 8 pixel blocks separately, and transform-coefficient-coding contexts for several coefficient positions can be processed

FIGURE 4 **Multi-thread decoding with wavefronts. Gray areas indicate CTUs that have already been decoded**

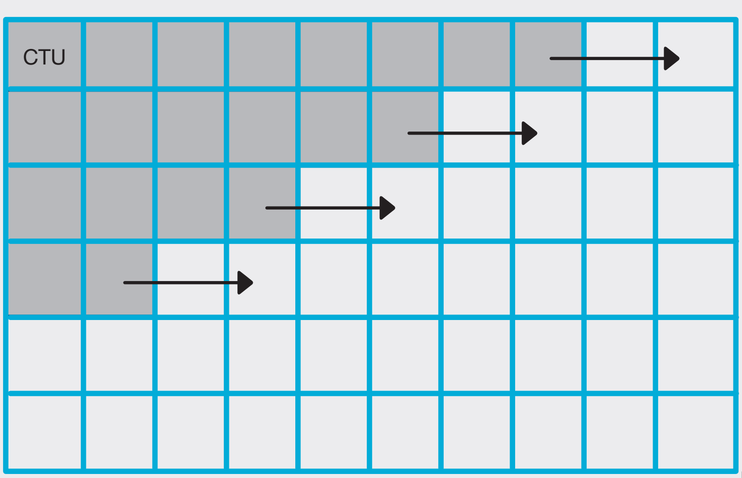

in parallel. Tiles and wavefront parallel processing (WPP) are among several HEVC tools that can provide high-level parallelism.

The concept behind WPP is to re-initialize CABAC at the beginning of each line of CTUs. To facilitate CABAC adaptation to the content of the video frame, the coder is initialized once the statistics from the decoding of the second CTU in the previous row are available. Re-initialization of the coder at the start of each row makes it possible to begin decoding a row before the processing of the preceding row has been completed. Thus, as shown in the example in Figure 4, several rows can be decoded in parallel in several threads with a delay of two CTUs between two consecutive rows.

The Tiles tool can be used for parallel encoding and decoding, and works by dividing a picture into rectangular areas (tiles) – as shown in Figure 5 – where each tile consists of an integer number of CTUs. The CTUs are processed in a raster scan order within each tile, and the tiles themselves are processed in the same way. Prediction based on neighboring tiles is disabled, and so the processing of each tile is independent. In-loop filters, however, can operate over tile boundaries. And as deblocking and SAO can be parallelized, filtering can be performed independently inside each tile, and tile boundaries can be processed by in-loop filters in a final pass.

The HEVC standard therefore enables both high- and low-level parallelization, which can provide significant benefits for multi-thread encoding and decoding of video such as 4K and 8K that has a higher resolution than HD.

FIGURE 5 Example of the way an image can be divided into tiles

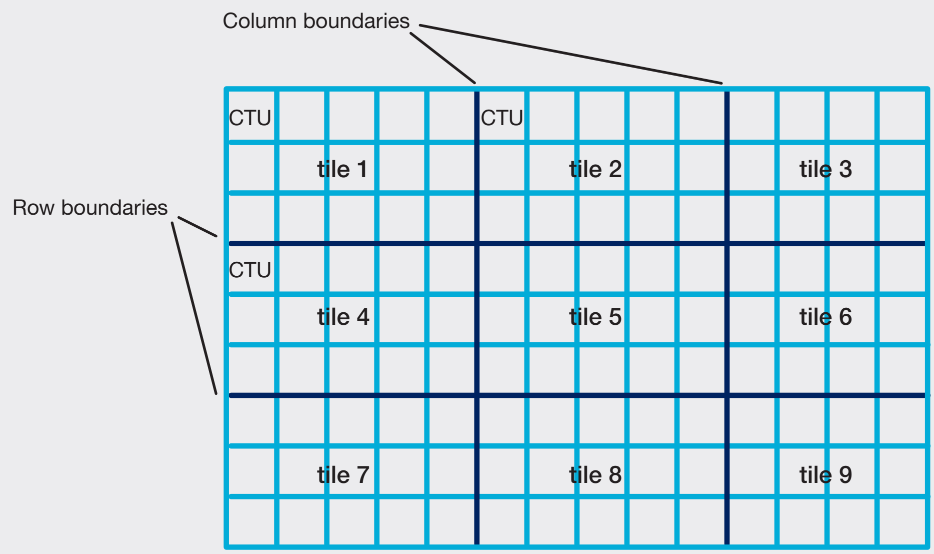

Performance and complexity

The improved coding efficiency of HEVC does however come with a price tag: increased computational complexity. Compared with its predecessor, HEVC is 50 -100 percent more complex for decoding and up to 400 percent more complex when it comes to encoding. While these comparisons are based on preliminary tests, they do give an indication of the new codec's computational complexity.

Real-time implementations of HEVC demonstrate that decoding of full HD (1080p) at 50 or 60fps is possible on fast desktop and laptop computers, running on a single core. Performance increases with multiple core implementations (hardware acceleration), so that a modern smartphone is capable of 1080p decoding at 25 or 30fps[8].

Applications

The new standard is a general one suitable for the compression of all kinds of video. The focus for the first version is consumer applications and for this, three profiles have been defined: Main, Main 10 and Main Still Picture.

Main is an all-purpose profile with a depth of 8 bits per pixel, supporting 4:2:0 – the most common uncompressed video format used by consumer devices from mobile phones to HDTVs. Main 10 extends the bit depth to 10 bits per pixel, which is well suited to consumer applications, such as UHDTV, where very high quality is critical. The increased bit depth can compress wide dynamic range video without creating banding artifacts, which sometimes occurs with 8 bits. The third profile, Main Still Picture, used for still images, is a subset of Main and carries a single still picture at a depth of 8 bits per pixel.

The initial deployments of HEVC released in 2013 will be for mobiles and OTT applications. Software implementations capable of decoding HEVC without hardware acceleration can easily be downloaded to smartphones, tablets and PCs, enabling mobile TV, streaming and download services on existing devices. To this end, in August 2012, Ericsson announced SVP 5500[9], the world's first HEVC real-time video encoder for live-TV delivery to mobile devices. However, as it is better to perform encoding on hardware and as HEVC is computationally more demanding than previous standards, it may be some time before video telephony based on this standard enters mobile platforms, whereas encoding on PCs is already feasible.

Set-top boxes with new decoders will become available soon, enabling content broadcast via satellite, cable or terrestrially to take advantage of HEVC. The new standard plays a key role in the provision of UHDTV, and as prices drop and displays become affordable, the number of services utilizing such high resolutions is expected to rise within a few years. Flat-panel displays for HDTV have been on the market for almost a decade, so this may be a good time for consumers to start upgrading to UHDTV.

What's coming

The finalized version of HEVC »

targets most consumer devices and services. However, for more specialized applications, such as 3D, content production or heterogeneous devices and networks, some additions to HEVC may prove useful. With this in mind, MPEG and ITU are working together on a number of ideas, including support for stereo and multi-view (glasses-free) 3D video, an extension that encodes multiple views by rearranging picture buffers and reference picture lists. A first drop is expected in January 2014, with a more advanced version that will support joint encoding of texture and depth information coming in the early part of 2015.

Scalability is a key attribute of any codec, as it enables trimming of video streams to suit different network conditions and receiver capabilities; scalable extensions to HEVC are planned for July 2014. Range extensions, which support several color formats as well as increased bit depths, are another area currently under development.

In addition to these extensions, further improvements are expected to take place inside the current HEVC framework, such as more efficient encoding and decoding (both software and hardware). It is likely that the full potential of HEVC will take some time to unfold, as encoding algorithms develop and the challenge posed by the optimization of encoders and decoders in multi-core architectures is overcome.

In short, HEVC or H.265 is twice as efficient as its 10-year-old predecessor, H.264. The improved efficiency that this codec brings will help to ease traffic load in networks and enable the creation of new and advanced video-based services.

The codec supports parallel processing and even though it is more complex from a decoding perspective, tests have shown that it is suitable for adoption in mobile services. Compression of mobile video streams and OTT content are the most likely initial candidates for application of the codec, and within a few years it will undoubtedly bring UHDTV into our homes.

Rickard Sjöberg

is a senior specialist in video coding in the Multimedia Technologies department at Ericsson Research, Kista, Sweden. With an M. Sc. in computer science from the KTH Royal Institute of Technology, Stockholm, Sweden, he has been working with Ericsson since 1997 and has worked in various areas related to video coding, in both research and product development. In parallel, he has been an active contributor in the video-coding standardization community, with more than 100 proposals relating to the H.264 and HEVC video-coding standards. He is currently working as the technical leader of Ericsson's 2D video-coding research, including HEVC and its scalable extensions. His research interests include video compression and real-time multimedia processing and coding.

Andrey Norkin

is a senior researcher at Ericsson Research, Kista, Sweden. He holds an M.Sc. in computer engineering from Ural State Technical University, Yekaterinburg, Russia and a Ph.D. in signal processing from the Tampere University of Technology, in Finland. He has worked at Ericsson Research since 2008, contributing to HEVC standardization through technical proposals and activities, including the coordination of a core experiment on deblocking filtering, chairing break-out groups and subjective quality tests for the Joint Collaborative Team on Video Coding (JCT-VC). He has also been active in the 3D video standardization for JCT-3V. He is currently the project manager of the 3D VISION project at Ericsson Research, working on 3D video systems, and algorithms, as well as on parts of the standardization.

Per Fröjdh

is director of media standardization at Ericsson and former head of visual technology at Ericsson Research. He holds an M.Sc. in engineering physics and a Ph.D. in theoretical physics from Chalmers University in Gothenburg, Sweden. Part of his Ph.D. work was carried out on scholarship at Imperial College London, UK. Following postdoctoral appointments in the US, and Denmark, he held the position of professor of theoretical physics at Stockholm University, Sweden. He joined Ericsson in 2000 as manager of video research and standardization. He has contributed to MPEG and ITU work on H.264 and HEVC, served on the advisory committee for the W3C, and has been the editor of 15 standards on streaming, file formats, and multimedia telephony in MPEG, ITU, 3GPP and IETF.

References

1. ITU, January 2013, press release, New video codec to ease pressure on global networks, available at: http://www.itu.int/net/pressoffice/press_releases/2013/01.aspx#.UWKhxBnLfGc
2. Ericsson, November 2012, Mobility Report, available at: http://www.ericsson.com/ericsson-mobility-report
3. Fierce Broadband Wireless, 2013, Ericsson CEO: 90% of network traffic will be video, available at: http://www.fiercebroadbandwireless.com/story/ericsson-ceo-90-network-traffic-will-be-video/2013-02-25
4. ITU-T Recommendation H.265 | ISO/IEC 23008-2: High Efficiency Video Coding, available at: http://www.itu.int/ITU-T/recommendations/rec.aspx?rec=11885
5. IEEE, December 2012, Overview of the High Efficiency Video Coding (HEVC) Standard, available at: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6316136
6. IEEE, December 2012, HEVC Deblocking Filter, available at: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6324414
7. IEEE, December 2012, Overview of HEVC High-Level Syntax and Reference Picture Management, available at: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6324417
8. IEEE, DOCOMO Innovations, December 2012, HEVC Complexity and Implementation Analysis, available at: http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6317152
9. Ericsson, 2012, Ericsson announces world's first HEVC encoder for live TV delivery to mobile devices, available at: http://www.ericsson.com/news/120822_ericsson_announces_worlds_first_hevc_encoder_for_live_tv_delivery_to_mobile_devices_244159018_c

Telefonaktiebolaget LM Ericsson
SE-164 83 Stockholm, Sweden
Phone: + 46 10 719 0000
Fax: +46 8 522 915 99

284 23-3198 | Uen
ISSN 0014-0171