EMBRACING 4K HEVC: RETHINK, BE REVOLUTIONARY

Dr K. W. Moore

Mediaware, Australia

ABSTRACT

Next generation 4K and Ultra HD (4K/UHD) television formats present broadcasters with an opportunity to intrinsically transform their operations. By embracing 4K/UHD broadcasters have the opportunity to jettison legacy technology whilst rethinking and reshaping their broadcasting platform with the future in mind.

Designing the Broadcast platform around software-based compressed video processing components is now a commercially proven option that warrants careful consideration.

Latest generation software operates directly on the high efficiency video coding (HEVC) video bit stream, in its compressed form, without fully decoding and re-encoding every frame. The amount of processing required is therefore significantly reduced and video quality is preserved.

With software-based compressed video processing components, 4K/UHD HEVC services can be economically delivered and monetised on non-proprietary IP-based infrastructure today. This provides a platform enabling Broadcasters to easily scale up the number of services delivered and migrate to higher resolutions, frame rates and pixel depths, without the need to replace infrastructure.

INTRODUCTION

The television industry is going through rapid change. The adoption rate of new formats and technologies is accelerating. Facilities no sooner get up to date with one set of standards (e.g. SMPTE 292M HD-SDI) then they have to think about replacing them. Technology innovation, coupled with business advantage, is a trigger for change.

Next generation 4K/UHD television formats present broadcasters with an opportunity to intrinsically transform their operations. Along with the advantages that come with differentiating content offerings and being able to influence the rate at which new formats are adopted, embracing 4K/UHD provides broadcasters with the opportunity to jettison legacy technology and to rethink and reshape their broadcasting platform with the future in mind.

Delivery to the home of 4K/UHD television services is enabled by the new HEVC video compression format. However, since encoding 4K/UHD HEVC requires significantly more computation than encoding high definition (HD) Advanced Video Coding (AVC) services, first generation encoders are highly complex. Given the inherent lag that a new resolution introduces, waiting for market-ready 4K/UHD baseband platforms means the generational opportunity to transform operations could be lost.

Replacing Serial Digital Interface (SDI) workflows with Internet Protocol (IP)-based ones is now a widely accepted method for revamping a broadcast platform. However, the bandwidths needed to carry uncompressed 4K/UHD necessitates the deployment of cutting edge network infrastructure. 10GbE IP networks cannot carry uncompressed 4K/UHD TV formats and deploying multi-link 10GbE, 40GbE or 100GbE networks requires significant investment. On the other hand, the compressed 4K HEVC video profiles used to deliver content to the home, can easily be carried on standard IP networks.

Another developing trend is the deployment of software-based capabilities, running on commodity off-the-shelf hardware, to fulfil standard broadcast functions. The operational and commercial benefits of using standard IT hardware are obvious. The sheer processing demands of processing raw 4K/UHD video data means the current generation of hardware is not always adequate.

Redesigning the Broadcast platform around software-based compressed video processing components is one way to avoid additional infrastructure investments. Latest generation software is able to operate directly on HEVC video, in its compressed form, without fully decoding and re-encoding every frame. This style of processing greatly reduces the amount of bandwidth and processing power required, preserves quality and saves processing time. With currently available software, 4K/UHD HEVC services can be delivered and monetised today on non-proprietary IP-based infrastructure.

In live environments, content insertion such as ad replacement and regionalisation can be achieved without the expense of additional encoding hardware. File based processing and media distribution also benefits from reduced time for delivering content to consumers; whereas generating regional copies of 4K/UHD HEVC movies can take days with offline software encoding. By utilising compressed video processing, which avoids the need to re-code every frame of video, generational losses are avoided and processing times significantly reduced.

The timely combination of consumer demand for higher quality content (4K and Ultra HD) and a new compression standard (HEVC) provides a rare window of opportunity for broadcasters to challenge the status quo and revolutionise delivery platform design.

This paper looks at how broadcast platforms can be transformed using software-based compressed video processing components and how platforms can be deployed today which provide flexibility to migrate to future formats and resolutions without replacing existing infrastructure.

4K AND IP WORKFLOWS

SDI-based workflows are rapidly being replaced by IP-based workflows. IP provides significantly greater flexibility over SDI, enabling dynamically controlled software defined network configurations that allows operations to adapt to changing requirements more quickly.

Much of the hype around IP-based workflows is framed in the context of replacing point-to-point cabling with a more flexible IP network topology. However, the video moving over these links is still baseband (uncompressed), or using costly point-to-point light compression, and the style of video processing remains largely the same; IP networks are just used as a smarter way to move baseband video around.

While these new architectures make economic sense for standard SD and HD services, the benefits are less clear when looking at 4K/UHD resolutions and frame rates on current IP network infrastructure. 10GbE IP networks cannot carry a single uncompressed 4Kp60 (2160p59.94) video channel. A 40GbE or 100GbE IP network is required unless video compression is employed.

By comparison, using HEVC compression, the same video channel will be encoded at 20 to 25 Mbps for delivery to the home. At 25 Mbps, many video channels can be streamed over a standard 1GbE IP local area network (LAN).

4K VIDEO PROCESSING

The higher the video resolution, frame rate and pixel depth, the greater the computing power and I/O required to support processing. A 4K/UHD (2160p60) video has eight times as many pixels per second than HD (1080p30). While Moore's law is reducing the cost of computing power, in the near term delivering a 4K/UHD channel may require an up to eight times increase in computing infrastructure using the same type of platform used for HD video delivery.

The simple, but not always obvious, approach to reducing computing requirements is to employ video processing pipelines that minimize the sheer volume of data (frames, pixels, bytes) that are actually processed. Conceptually if not all program content needs processing then the content can remain in its compressed form (bypassing a decode-process-encode step) and then be seamlessly merged with the sections of processed content.

MINIMIZING FORMAT CONVERSION

HEVC video encoding is typically reported as requiring 10x more processing power than for the equivalent AVC encoding.

Therefore, if comparing the possible computational cost of encoding a HD AVC video stream of 1080p30 with a 4K/UHD HEVC video stream of 2160p60, there is the order of an eighty times (80x) increase in processing power requirements (4x more pixels per video frame, 2x the frame rate, 10x to encode).

Due to these significant processing requirements, first generation 4K/UHD encoders will be highly complex. Coupled with the generation loss that occurs with each decode/re-encode step, a major goal in developing an economically viable TV delivery chain is to reduce the number of times a 4K/UHD channel is decoded and re-encoded during distribution and delivery.

A typical TV delivery chain is show in Figure 1. In the diagram, the video above the line is uncompressed while the video below the line is compressed. We use the term above-the-line processing to mean any processing that is performed on the uncompressed (baseband) video signal. We use the term below-the-line processing to refer to any processing that is performed directly on the compressed (e.g. MPEG-2, AVC or HEVC) video.

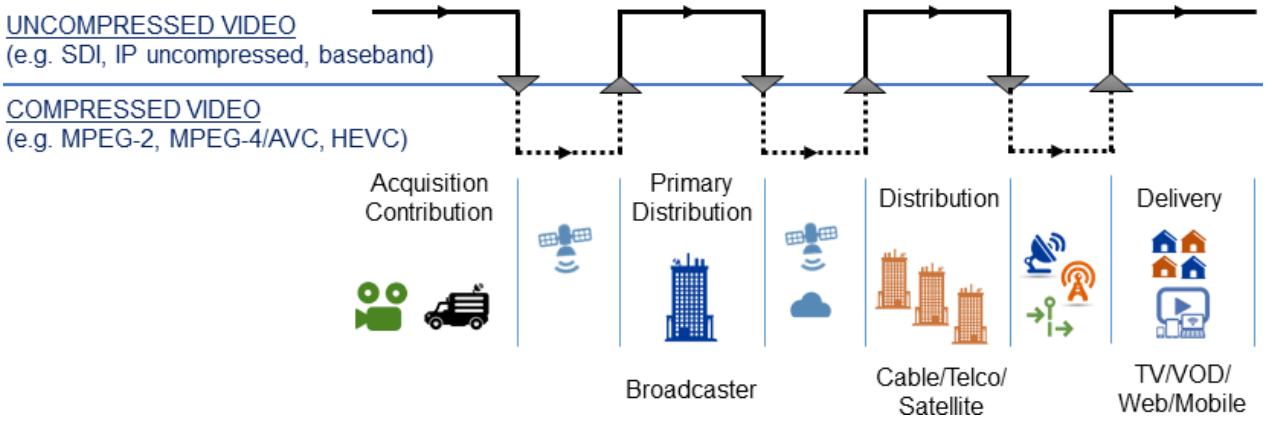

Figure 1 - TV Delivery Chain: Traditional workflow

As the video signal passes through the stages from acquisition, contribution, to distribution and delivery, it is typically encoded to a compressed format for transmission between facilities and decoded to baseband for processing at each facility.

One method to increase the efficiency and economic viability of this delivery chain for 4K/UHD and future formats is to reduce the number of these format conversions.

While increasing above-the-line baseband processing is conceptually simple, bandwidth limitations between facilities means that shifting more processing below-the-line, on the compressed video, is a compelling strategy to explore.

COMPRESSED VIDEO PROCESSING

Compressed video processing is defined as the method of processing compressed video bit streams without the need to fully decode and re-encode every video frame. A transport stream multiplexor is an example of a compressed video processing function, as it works at the bit level without needing to decode the video data back to baseband.

A more complex example is the frame-accurate switching between two live compressed 4K HEVC video streams. In this function, the majority of the compressed video stream remains untouched by the process. Only the video data in the immediate vicinity of the switch point are affected by the switch. In the simplest method, if one second worth of video were to be re-encoded either side of the switch point, the total computation and resources required to achieve this would be a fraction of what would be required for full decoding, baseband switching followed by re-encoding back to 4K HEVC.

Leading commercial compressed domain switching products employ sophisticated processing to ensure only the minimum number of video frames are touched. They also re-use much of the original compression information to further reduce computational requirements and to preserve the video quality. Proper management of buffer levels and compression structures is essential to guarantee that splice points are seamless and indiscernible to downstream decoders.

SOFTWARE-BASED COMPRESSED VIDEO PROCESSING WORKFLOWS

Software-based high-quality compressed video processing functions, specific to broadcasters operating on MPEG-2 and AVC video bit streams, have been in operation for over 7 years. The economics of these systems have to-date best suited regional

broadcasters, who localize and monetize services for many different regions. The “encode once, localize many times” model makes commercial sense when there are many geographic regions.

Given the need to upgrade or replace infrastructure to support 4K/UHD HEVC services, this approach now warrants consideration by all broadcasters, including major playout centres and content aggregators, looking to economically rethink their delivery platform to meet future needs.

Latest generation commercial software is able to work directly on compressed 4K HEVC video. Already most of the key functions required to support below-the-line regionalization and monetization of 4K HEVC services are available, including clean seamless switching and splicing, ad and promo insertion and time delay. Therefore, it is possible to start developing new distribution and delivery platforms from these capabilities today.

Broadcasters looking to showcase or roll out 4K/UHD services can now deliver and monetize them on non-proprietary standard IP-based infrastructure. Figure 2 shows a generic baseband workflow.

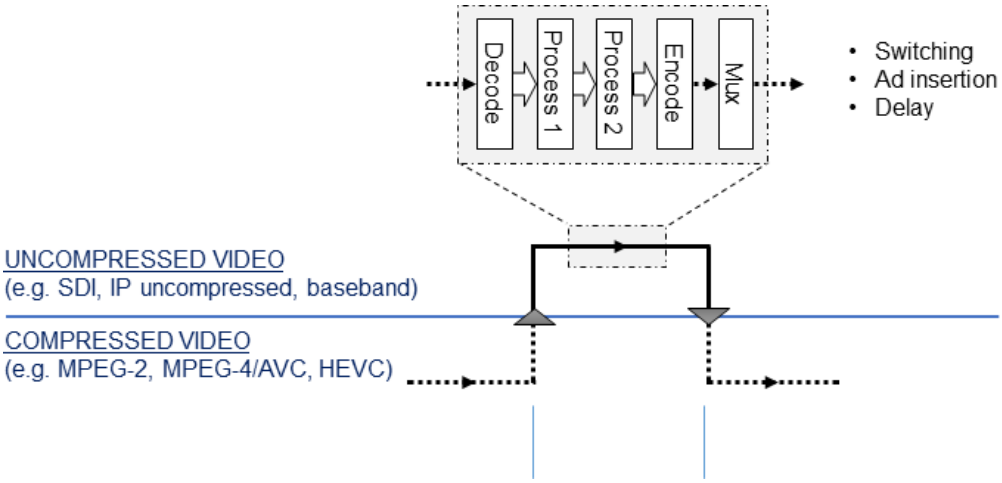

Figure 2 – Uncompressed video processing workflow

Figure 3 shows this capability shifted from an above-the-line baseband architecture to a below-the-line compressed video processing equivalent.

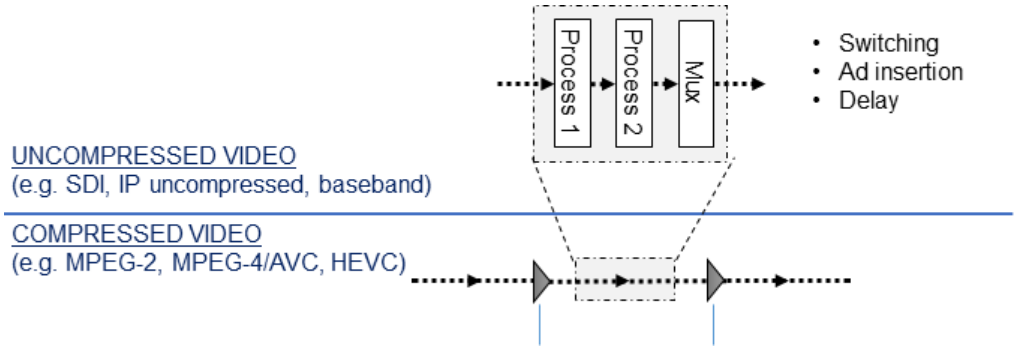

Figure 3 – Compressed video processing workflow

Figure 3 clearly shows how compressed video processing functions reduce the requirement to deploy decoders and encoders at every facility in the delivery chain.

4K AD-INSERTION EXAMPLE

The following example demonstrates how a 4K/UHD HEVC service could be distributed and delivered using software-based compressed video processing functions. The goal of the broadcaster is to take a linear video TV channel and make it compelling to the local audience and gain an additional revenue stream from selling advertisements.

In this example, a single premium 4K HEVC network feed is being localized and delivered to four different regions. Avails or breaks in the Network feed can be used to insert local advertisements and news breaks for each region as shown in Figure 4.

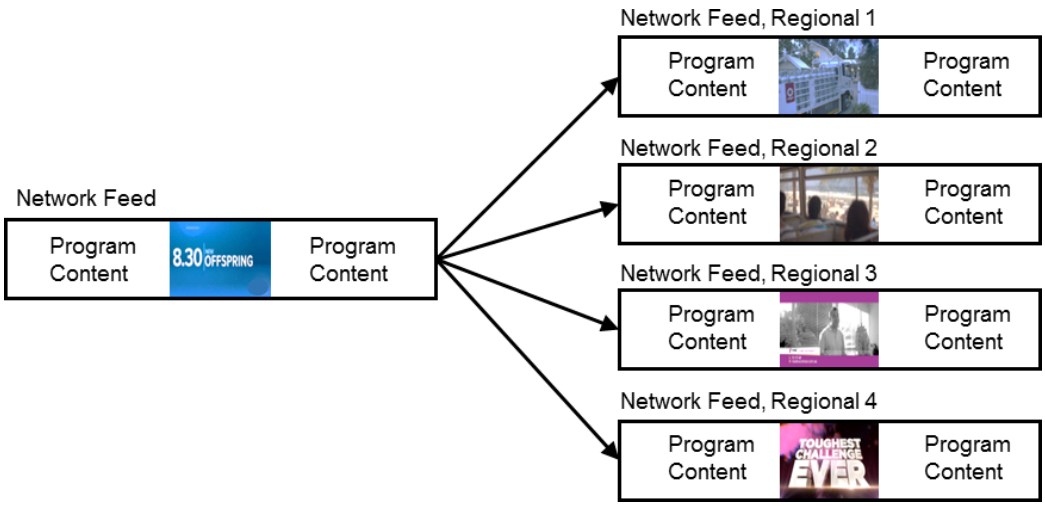

Figure 4 – Service regionalization

The Network feed originates from the Primary head-end which performs some baseband processing before encoding and transmitting it in Long GOP HEVC video 2160p60 at 20Mbps.

If deploying a traditional TV delivery platform, as shown in Figure 5, each regional head-end decodes the service back to baseband (carried over multiple 3G-SDI or 40GbE IP), switches between the network feed and the local content served from a video server, controlled by automation, followed by re-encoding the service back to HEVC for delivery to the home.

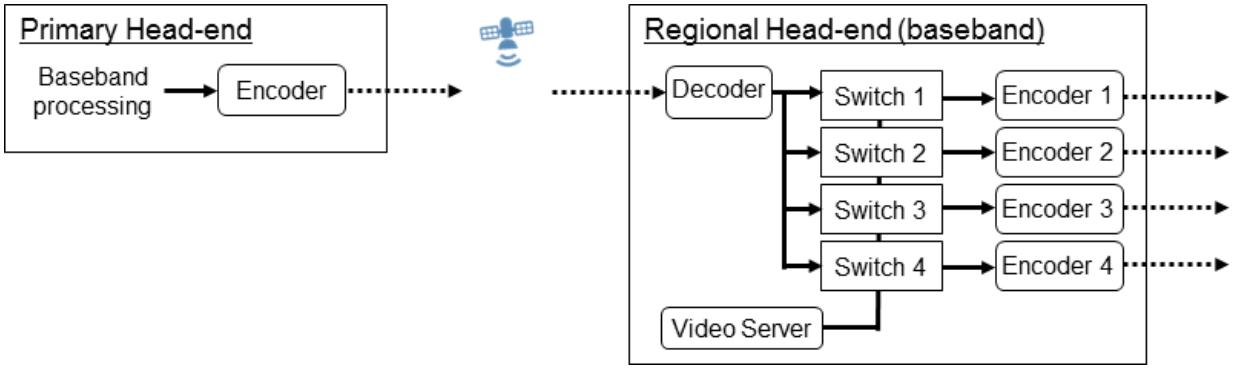

Figure 5 – Baseband regionalization

Replacing the baseband uncompressed workflow with a software-based compressed video workflow, shown in Figure 6, the regional head-end no longer needs the input decoder or the encoder per delivered regional service. Files on a network drive or local storage are spliced frame-accurately and cleanly into the compressed video bit stream prior to delivery.

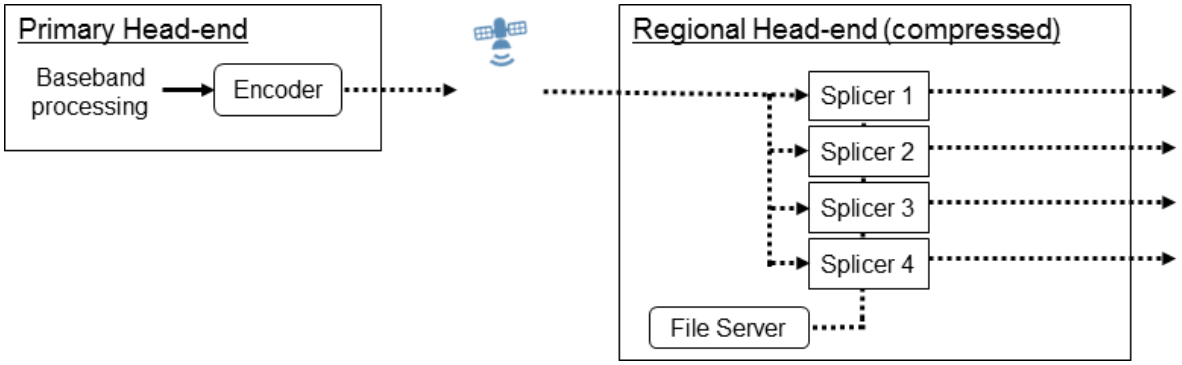

Figure 6 – Regionalization via compressed video processing

For the delivery of a 4K HEVC service, the compressed video workflow described requires less equipment that supports 4K/UHD formats. The example shows one less decoder, four less encoders and removes the need for video server and 4K/UHD signal distribution and routing equipment. The encoded video at the primary head-end is configured to match the profile for delivery to the home, saving on contribution bandwidth costs.

In addition, on-line catch up services can be created on this platform using a similar type of workflow in reverse. Playout scheduling data is used to segment on-the-fly program and commercial content into file based packages, shortening time to delivery for the consumer.

CONCLUSION

With consumer demand increasing for higher quality content (4K and Ultra HD) and a new compression standard (HEVC), broadcasters are presented with an opportunity to rethink and reshape their operations.

Next generation 4K/UHD television formats challenge existing infrastructure and workflows. Simply scaling up SD and HD operations commits the broadcaster to an expensive delivery configuration moving forward.

Developing broadcast platforms around software-based compressed video processing components is one way to minimize the additional infrastructure investments required to support new 4K/UHD channels. Shifting video processing below-the-line, to operate directly on the compressed video, reduces the number of decode/encode steps required in the TV delivery chain. This reduces computing and network infrastructure requirements and preserves video quality.

Broadcasters looking to showcase or roll out 4K/UHD services can now deploy the commercial solutions required to regionalization and monetization these services on non-proprietary standard IP-based infrastructure.

REFERENCES

1. Moore, K. W. 2009. Frame Accurate Compressed Domain Splicing, Proceedings of 2009 SMPTE, Sydney.
2. ISO/IEC 23008-2 MPEG-H Part 2 and ITU-T H.265
3. ISO/IEC 14496-10:2005, *Information technology – Coding of audio-visual objects – Part 10: Advanced Video Coding* ITU-T Recommendation H.264 *Advanced video coding for generic audiovisual services*, November 2007