Jumping on the right HEVC train –

Pragmatic migration strategies to the new compression standard

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Yaron Raz
Director, Video Solutions Marketing
Harmonic
4300 North First Street San Jose, CA 95134
Yaron.raz@harmonicinc.com

Overview

High-Efficiency Video Coding (HEVC) is the newest video coding standard of the ITU-T Video Coding Experts Group (VCEG) and the ISO/IEC Moving Picture Experts Group (MPEG). The main goal of the HEVC standardization effort is to enable significantly improved compression performance relative to existing standards – in the range of 50% bit rate reduction for equal perceptual video quality. A complete overview of the technical specification is detailed in (1). The HEVC specification reached the Final Draft International Specification in January 2013 and we now have to examine how this technology can be used in commercial applications. (2) provides a detailed analysis of where the HEVC standard can be used. In addition to the codec layer, in April 2013 MPEG also defined the DASH adaptation layer to incorporate HEVC into a DASH protocol (3), (4) and it therefore provides a complete standard solution for HEVC IP delivery.

For cable operators in North America, most of whom still relying heavily today on MPEG-2, this poses a significant migration dilemma: What is the best migration strategy that takes advantage of the great bandwidth savings promised by HEVC, while also considering the large installed base of MPEG2-only or MPEG-2/AVC STBs deployed today?

This paper looks at the different applications that could take advantage of HEVC and proposes a migration path for MSOs.

Second Screen ABR Video

Second screen (aka Multiscreen) video delivery refers to the delivery of video services to a breadth of consumer electronic devices including PCs, MACs, gaming consoles, tablets, smart phones, connected TVs, as well as devices such as Roku, Apple TV etc. Many of these devices utilize high performance CPUs, and while their AVC decoding is done mostly in HW today, some of these devices enable the downloading of a software-based HEVC decoder.

Back at CES in January 2013, Samsung demonstrated an HD HEVC decoder on its F8000 line of connected TVs, running a quad A15 ARM core at 1.5 GHz. Qualcomm was also showing the Snapdragon 600, which can decode HD HEVC on a quad-core ARM. This processor is deployed inside the Samsung Galaxy S4, which is the first mobile device with native HEVC HD support.
At the NAB show in April 2013, Harmonic demonstrated a player on Nexus 10, playing HD (720p60), in software using SQUID technology. Battery life at 720p30 was 7 hours, and reduced to 4 hours at a resolution of 720p60. It should be noted this is purely software-based: no GPU optimization is applied. This is of particular interest when considering porting of HEVC to legacy devices. Multimedia Research Group predicted that HEVC can be ported to 1B devices in 2012, and the Harmonic demonstrations at NAB provided real proof that this claim is valid (5).
The advantages of using HEVC for Multiscreen ABR are the following:

- Enables subscribers to consume a higher profile (higher video quality) over a sustained time within the home
- Enables delivering a 720p60 profile at a reasonable 3Mb bit-rate, which is required for devices feeding sports events to a large screen
- 50% reduction in the CDN bandwidth and storage required for ABR services
- 50% reduction in the DOCSIS capacity required to deliver Multiscreen services to the home
- 50% reduction in the storage capacity required for nDVR. This is of special importance for operators who are pursuing a copy-per-subscriber model for nDVR (set by the Cablevision precedence) in order to avoid content-rights infringement. nDVR utilizing a copy-per-subscriber model requires a huge amount of storage, and the ability to cut that cost by 50% is quite appealing.
- Since ABR technology is Unicast by nature, it allows sending HEVC video to HEVC-capable devices, while sending AVC video to devices that do not support HEVC. No overnight transition is required.

On the other hand, Multiscreen ABR session concurrency today is still below 5% of viewership, and given the typical bitrates of ~3Mb per session, it does not consume a significant portion of the DOCSIS bandwidth at this point.
For example: For a DOCSIS group of 8 QAMs serving 500 homes (250 subscribers), and assuming that 5% of these subscribers are running a Multiscreen session of 3Mb at

peak time, the overall bandwidth taken by Multiscreen is 37.5M or about one QAM of the 8 allocated for DOCSIS. Obviously, this bandwidth will grow as more eyeballs shift to consume video on second screen and connected TVs.
Some cable operators have already stated that strategically, they'd like to reduce their dependency on a managed STB in the subscriber's premises, and rely more on CE devices for delivery of video to the main TV. In such cases, the concurrency of ABR sessions to the home is likely to grow faster, and hence the saving of DOCSIS capacity, CDN bandwidth and storage will be more significant.

IP STBs

There are multiple drivers for introducing IP STBs into cable. These include the ability to offer next generation navigation capabilities that include personalized guides, recommendations and graphic-rich content selection and navigation, all based on web-based technologies such as HTML5. In addition, by migrating to IP delivery, operators can consolidate their delivery network between CE Multiscreen devices and managed STBs, as well as offer more live content while avoiding having to broadcast all live channels (basically get the benefits of Switched Digital Video that are inherent to IP).
However, the IP STBs deployed today do not have the HW capabilities nor the SW capabilities to decode HEVC. First HEVC decode ASICs are going to be available at the second half of 2013, and first HEVC-capable STBs will likely be available towards the end of 2013.
However, unlike the lengthy introduction of AVC, since HEVC over ABR is based on unicast sessions, the introduction of HEVC STBs can be done in a way that allows sending HEVC feeds to HEVC capable boxes, while sending AVC ABR feeds to boxes that do not support HEVC, similar to the transition described for CE devices.
So unlike AVC, where MSOs have been deploying AVC-capable boxes for the past 5 years, yet many have not taken any advantage of AVC and are delivering 100% of their traffic in MPEG-2 due to the large installed base of MPEG-2 only boxes, the introduction of HEVC in conjunction with ABR capabilities in IP STBs allows immediate utilization of the HEVC capability for these boxes.
This is true for delivery of IP over ABR, but for MSOs considering the delivery of IP multicast services, migration to HEVC will be as challenging as the migration from MPEG-2 to AVC.

For many cable operators a natural question that comes to mind is: "Should I leap-frog AVC altogether and go directly from MPEG-2 to HEVC?"
The answer lies in the penetration ratio of AVC capable and HEVC capable STBs over time. Some operators have already reached 40%-50% penetration of AVC in their HD STBs, and will likely get to more than 60% penetration before they deploy their first HEVC STB.

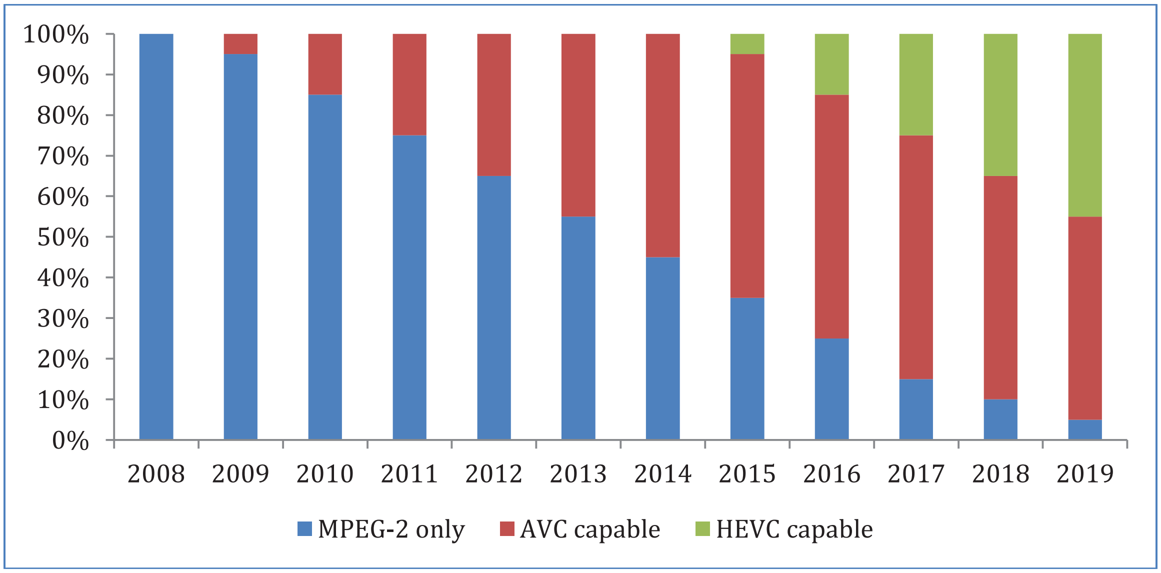

Figure 1 - Typical Breakdown of STB capability

In addition, migration to HEVC will likely be done on using unicast based on device capability. It does not preclude migrating to AVC for specific channels that are delivered only to AVC-capable boxes, utilizing SDV for AVC vs. MPEG-2 delivery. Each operator will need to consider his penetration ratio, number of boxes, and spectrum constraints in order to make that decision. But generically, the migration to HEVC seems to be independent of the introduction of AVC.

QAM STBs

Given that IP STBs are already deployed today, and HEVC STBs are yet to be available, it is highly likely that any HEVC capable STB will also be IP capable. Therefore the discussion here is not whether an MSO should deploy a pure QAM STB that is HEVC capable, but more related to running HEVC over QAM for the consumption by hybrid QAM & IP STBs. As stated before, the transition to a new codec is so more challenging in a broadcast based environment, and therefore it is much more likely that any HEVC content will be offered over IP ABR rather than over QAM. It will only be when HEVC penetration of hybrid STBs is quite high, and multicast gains are still high for the popular channels, for such an approach to be attractive. This is not likely in the next few years.

Ultra HD Video

For our purposes 4K is defined as 4096x2160x30Hz at 8 bits per pixel in a 4:2:0 representation. Ultra HD has been defined by ITU and is currently being specified by EBU and SMPTE for broadcast usage. We anticipate that the first commercial implementation of the Ultra HD specification will be 3860x2160x60Hz at 10 bits per pixel in a 4:2:0 representation, aligned with the MPEG HEVC Main 10 profile. Production profiles will increase the bit depth, and the frame rate.

Those 4K TVs currently shipping only support 30Hz, because HDMI 1.4 only supports a maximum frame rate of 30fps. HMDI 2.0 is addressing this and will support up to 120 Hz. We expect the first silicon solutions to be released towards the end of 2013, for demonstration at CES 2014.

One of the debates surrounding Ultra HD is which type of services will be deployable first, and which will reach an economically viable production/consumption point; from a content-availability and workflow perspective, VOD is the low-hanging fruit. 4K cinema files are already available and there would be very little impact on production workflows – as opposed to a live service, for which a new Ultra HD workflow supporting 10 bit 4:2:0 would have to be built; the experience gained from the rollout of HD over the last decade, shows us that this can take many years, and if there is no supporting regulatory influence, as was the case for HD, then it could take longer still. Drawing a parallel with 3D, which was a great VOD technology that was pushed by the entire industry as a live service, highlights all the roadblocks that the entire industry had to negotiate; we would therefore suggest that the first implementation of Ultra HD is likely to be VOD, with live services deployed in a second phase. We have seen a lot of interest in live broadcast of premium content (e.g. sports or special events), but we believe that this will be only for specific situations.

It is envisioned that a true broadcast Ultra HD, 3840x2160x60 Hz, will require encoding of around 14Mb/s for video using HEVC when the encoding technology becomes fully mature, but that 20Mb/s is likely to be required initially. Using AVC the required bitrate is likely going to be in the range of 28M, which makes it very challenging from a spectrum management perspective.
It is also assumed that any 4K screen will require either a new dedicated STB or more likely consume the video directly over IP.
As a result, it is quite likely that any 4K delivery will be based on HEVC from day one.

Video outside the home

While the initial rollouts of Multiscreen video have been confined to the subscriber's home, due to content rights issues, it is clear that operators would like to offer their subscribers the option of consuming these services anywhere. In fact, several MSOs have already started offering a subset of the Multiscreen lineup outside the home.
In this scenario, MSOs cannot rely on their high capacity DOCSIS connectivity to the subscriber's home, and therefore HEVC plays a major role in enabling higher video quality at lower sustained bit-rates.
However, we envision this use case as an extension of the Multiscreen use case. MSOs will likely couple their infrastructure for Multiscreen delivery within the home and outside the home, and therefore and transition to HEVC will occur for both use-cases at the same time.

Infrastructure migration to HEVC

As discussed above in this paper, we believe that pretty much any HEVC migration will be based on ABR technology. The figure below depicts a typical ABR ecosystem and identifies the components that will be affected by a migration to HEVC.

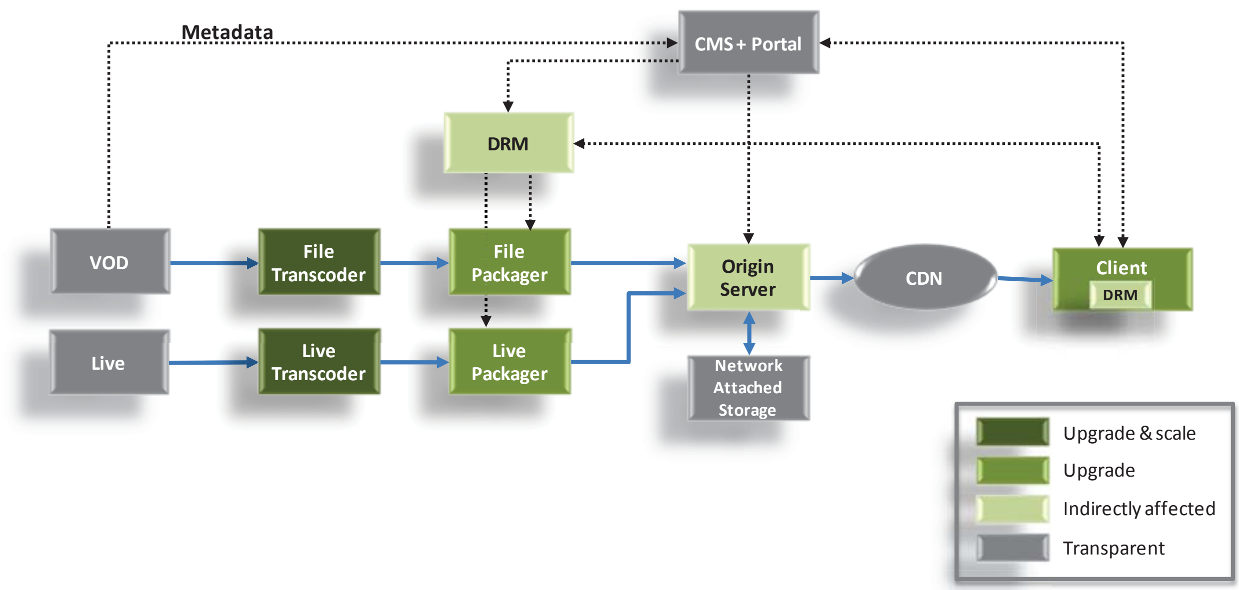

Figure 2 - Multiscreen system components affected by HEVC migration

- File Transcoder – File based transcoders are mostly based on SW technologies and can be upgraded / replaced to support HEVC encoding. It should be noted that HEVC is estimated to require about 10 times more computing capacity than AVC for the same profile. Therefore, supporting HEVC will require a larger file-based transcoding farm, and the HEVC profiles are likely going to be added to existing AVC profiles and offered to capable devices

- Live transcoder – Live transcoders for ABR today are based on software, hardware or a combination of software with some hardware / GPU assist. The hardware based ASICs available today, for the most part, do not support HEVC and will require some sort of a HW upgrade in order to support HEVC. We anticipate that next generation silicon supporting HEVC will become available towards end of 2013. The software-based transcoders have the potential of being upgraded to support HEVC. However, as stated before, since HEVC requires significantly more compute power than AVC, the channel density provided by these transcoders will significantly reduce, and therefore more live transcoder instances will be required in order to support the lineup.
- Both live and file-based packagers will need to be updated with code that supports the proper packaging of HEVC in the relevant ABR protocol(s). In addition, the publishing of manifest files will need to address both HEVC-capable clients as well as AVC-only clients, so that each client receives a manifest file that addresses its capabilities.
- Clients will obviously need to be upgraded to support HEVC. As discussed earlier in this paper, while most clients support AVC decoding natively, we believe that many of the devices deployed today, such as PCs, Macs, tablets and smart-phones can be loaded with clients that include HEVC decoders in software.
- Migration to HEVC should be transparent to the Origin Server, NAS and DRM system. However, since in many cases the HEVC profiles created will be incremental to the existing AVC set of profiles, these may require additional storage on the Origin / NAS / CDN.
- While HEVC could potentially be transported in any of the existing ABR protocols, so far only DASH has defined how to package HEVC content. As a result, most NA cable operators, which have been deploying HLS and HSS protocols for Multiscreen will need to add support for DASH in their packagers, Origin Servers, as well as their DRM. This will not be required if HLS and Smooth define a standard way for carrying HEVC content and clients support it.

Projected migration schedule

The timeline below depicts our estimation for the phases and timeframes for migration to HEVC by MSOs. Obviously these will vary on an operator-by-operator basis, but we believe the high level guidelines and order will be similar.

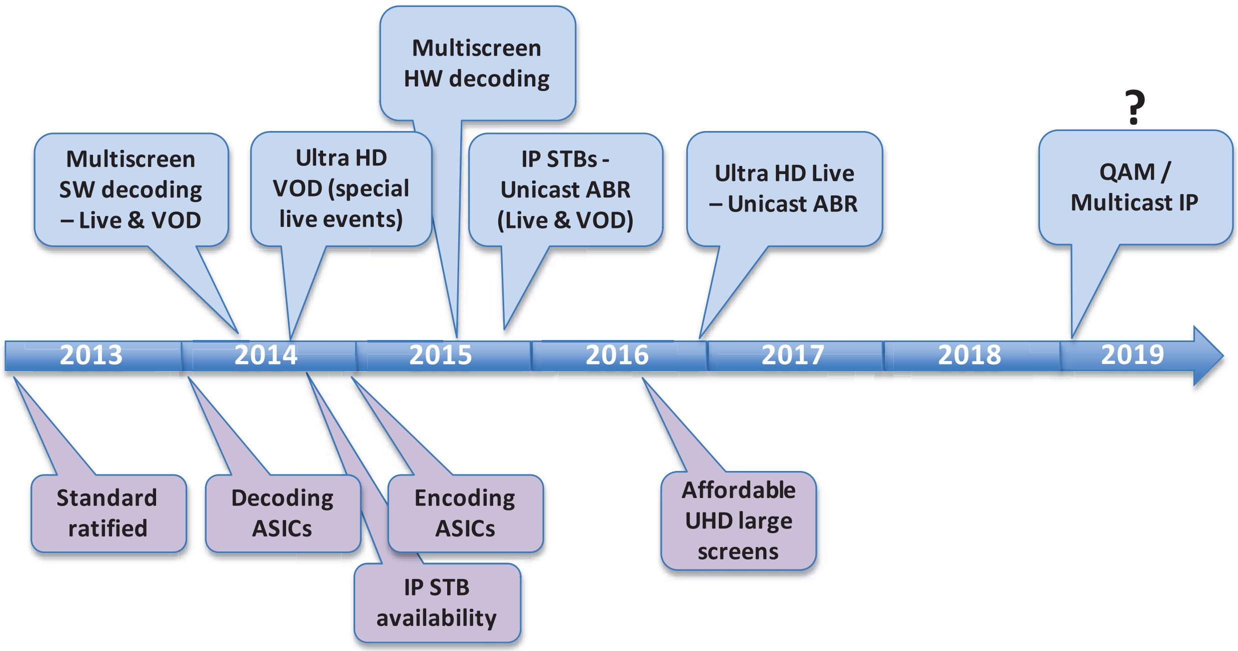

Figure 3 - Projected migration schedule to HEVC

Harmonic projects that first HEVC deployments will start in 2014 for Multiscreen devices that can be loaded with SW-based decoders such as PCs, MACs and tablets. In addition, some Ultra HD VOD content will be available for subscribers that buy 4K TVs. There may be some unique events such as the World Cup in 2014 that will provide opportunities for live 4K.

As more devices support HEVC in HW, operators that have not chosen to deploy a SW-based HEVC decoder on their client will start producing HEVC profiles in order to take advantage of the bandwidth and storage savings of HEVC.

We believe that managed IP STBs with HEVC capabilities will start deployment in 2015 and will utilize ABR unicast for both live and VOD.

Ultra HD live service will become more prevalent as large 4K TV sets become more affordable and widespread, probably around 2016.

Finally broadcasting / multicasting HEVC live content will not happen before there is a significant installed base of HEVC consumption devices and MPEG-2 STBs is small enough to start removing MPEG-2 from the live lineup. This phase might also never occur.

Summary

HEVC promises great benefits for cable operators. The 50% bit rate savings compared to AVC appears to be real, and this translates to savings in spectrum, transport network bandwidth, CDN costs, as well as storage costs.
There are two fundamental differences between the transition to AVC and the coming transition to HEVC. The first is that many of the devices deployed in the field today such as PCs, MACs, tablets etc. have the capability of decoding HEVC in software – a capability that was not present when AVC was first introduced. The second is the proliferation of ABR unicast technology for both unmanaged and managed devices, and therefore ability to avoid the simulcast vs. complete deployed base swap-out dilemma. HEVC can be introduced gradually, taking advantage of the HEVC-capable devices, while not interfering with any devices that only support AVC.
We believe that first cable deployments of HEVC will start as soon as 2014, and will focus initially on Ultra HD VOD (and potentially some unique live events such as the World Cup) as well as on Multiscreen delivery to devices that can be loaded with a SW-based decoder.
This will then be followed by delivery to managed HEVC devices late in 2015-2016 as HEVC STBs become more prevalent, while still utilizing IP ABR for delivery.
Ultra HD live will take more time as the whole live production chain will need to move to 4K, and the TV sets will need to become more affordable to justify the investment.
Finally, HEVC over QAM or IP multicast will take much more time, or may not happen altogether.
Cable operators deploying ABR ecosystems today, should make sure their equipment is HEVC future-proof. Today this means SW-based ABR transcoders, however, migration to HEVC will require much more processing power, and therefore will likely be gradual (select channels / VOD assets to begin with). Transcoding infrastructure that runs on COTF IT HW has the advantage of being more amenable to enabling HW swap out as IT technology improves, while having minimal disruption to the transcoding operation, and keeping it as transparent as possible to the operator.

Bibliography

1. Sullivan, Gary J., Ohm, Jens-Rainer , Han, Woo-Jin, and Wiegand , Thomas ," Overview of the High Efficiency Video Codec ", IEEE Trans on Circuits and Systems for Video Technology, December, 2012.

2. Thierry Fautier: “Use cases for the new HEVC standard”, Harmonic.Engineering Broadcasting Conference, NAB 2013. April, 2013.

3. Thomas Schierl, Miska M. Hannuksela, Ye-Kui Wang, and Stephan Wenger: System Layer Integration of HEVC, IEEE Transactions on Circuits and Systems for Video Technology, IEEE TCSVT, Special Issue on Emerging Research and Standards in Next Generation Video Coding, vol. 22, issue 12, pp. 1871-1884, 2012.

4. X. Wang ; Y. Xu ; C. Ai ; P. Di ; X. Liu ; S. Zhang ; L. Zhou ; J. Zhou, “DASH(HEVC)/LTE: QoE-based dynamic adaptive streaming of HEVC content over wireless networks”, 2012 IEEE Visual Communications and Image Processing (VCIP), 27-30 Nov. 2012.

5. MRG: http://www.mrgco.com/reports/hevc-decoding-in-consumer-devices/

Abbreviations and Acronyms

ABR	Adaptive Bit Rate
HLS	HTTP Live Streaming
HSS	HTTP Smooth Streaming
DASH	Dynamic Adaptive Streaming over HTTP
STB	Set Top Box
HEVC	High Efficiency Video Codec
AVC	Advanced Video Codec
VCEG	Video Coding Experts Group
MSO	Multi System Operator
CES	Consumer Electronic Show
NAB	National Association of Broadcasters
CDN	Content Distribution Network
DOCSIS	Data over Cable Service Interface Specification
nDVR	Network Digital Video Recorder
HTML5	Hyper Text Markup Language version 5
QAM	Quadruple Amplitude Modulation
NAS	Network Attached Storage
DRM	Digital Rights Management