Video Adaptation for Consumer Devices: Opportunities and Challenges Offered by New Standards

James Nightingale, Qi Wang, Christos Grecos, and Sergio Goma

ABSTRACT

Video and multimedia streaming services continue to grow in popularity and are rapidly becoming the largest consumers of network capacity in both fixed and mobile networks. In this article we discuss the latest advances in video compression technology and demonstrate their potential to improve service quality for consumers while reducing bandwidth consumption. Our study focuses on the adaptation of scalable, highly compressed video streams to meet the resource constraints of a wide range of portable consumer devices in mobile environments. Exploring SHVC, the scalable extension to the recently standardized High Efficiency Video Coding scheme, we show the bandwidth savings that can be achieved over current encoding schemes and highlight the challenges that lie ahead in realizing a deployable and user-centric system.

INTRODUCTION

Advances in portable consumer devices such as smartphones have boosted global popularity of high-quality multimedia-rich applications over WiFi, 3G, and 4G wireless networks. Cisco [1] has predicted that over half of all Internet traffic will be generated by mobile devices by 2017, with a remarkable 13-fold increase from 2012, and the vast majority of the traffic is expected to be video related. Though substantial volumes of mobile traffic are currently offloaded to wired networks via WiFi connections, the trend is toward ubiquitous content consumption over all available wireless mediums. This rapid growth, while offering myriad opportunities for content providers and advertisers, will place huge burdens on the network infrastructure and pose significant challenges for Internet Service Providers (ISPs).

The main implication associated with this unprecedented traffic growth is the demand for more advanced management of network capacity and application bandwidth requirements, while meeting consumers' expectations. Users expect to receive high-quality services by fully exploiting the capabilities of their devices. Given the massive improvements in hardware over the last few years, with High-Definition (HD) or better displays and quad-core processors becoming the norm for even modestly priced handsets, this is a daunting task. Service providers will need to deliver enormous volumes of delay-sensitive traffic, possibly over imperfect wireless channels, to a plethora of portable devices with a wide variety of processing and playback capabilities, at the same time, ensuring that each customer's experience is maximized and their loyalty retained.

In this article we highlight three separate but intricately woven challenges. First, by improving video compression techniques, bandwidth requirements can be reduced without compromising on video quality. Second, employing scalable and adaptive video streaming technologies, network resources can be conserved by matching video stream characteristics to user device capabilities. Third, video streams can be dynamically and intelligently adapted to the varying status of the transmission channels.

With the standardization of High Efficiency Video Coding (HEVC) [2] as H.265 by the ITU-T and as 23008-2 MPEG-H Part 2 by ISO/IEC in 2013, a major step forward has been taken in delivering high-quality video streams at greatly reduced bandwidths. Bandwidth reductions of up to 50 percent can be achieved using HEVC, when compared with current video compression standards, without loss of picture quality. HEVC is an effective means to address bandwidth consumption, the first of the three highlighted challenges.

SHVC [3], a more recent and as yet incomplete, scalable extension to HEVC, will offer the facility to adapt an HEVC-encoded video stream to meet end user device capabilities and network resource constraints. By offering a single encoded bitstream containing multiple representations of the same video at different spatial, temporal, or quality resolutions, SHVC can be adapted, statically or dynamically, to match a device constraint such as the display resolution. Using SHVC will address the second of the three highlighted challenges by conserving bandwidth, as only those sub-streams that the client device is capable of playing will be sent over the network.

James Nightingale, Qi Wang, and Christos Grecos are with the University of the West of Scotland.

Sergio Goma is with Qualcomm Inc.

0163-6804/14/$25.00 © 2014 IEEE

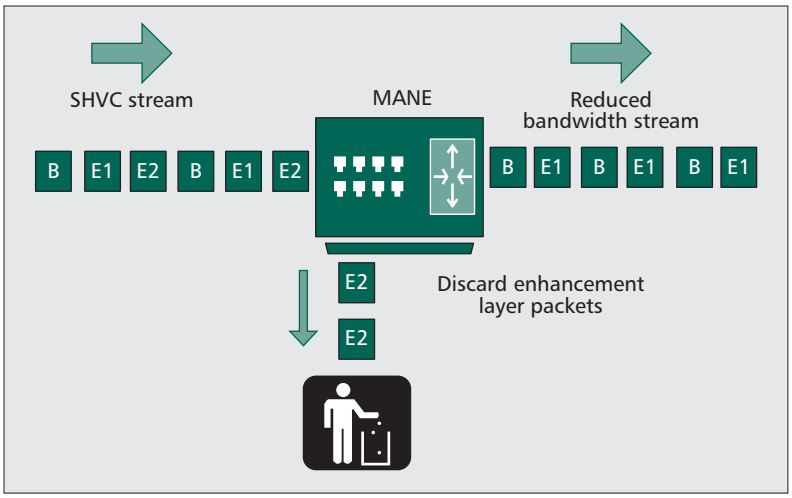


Figure 1. Adapting an SHVC scalable video stream at a MANE to reduce the bandwidth requirement.

In the third and most demanding of the challenges, SHVC adaptation solves the problem of meeting prevailing network constraints by only transmitting the number of sub-streams or layers that can be successfully delivered under the current network conditions. Adaptation of scalable video streams takes place at a Media Aware Network Entity (MANE) located within the network. The principle of scalable video stream adaptation to meet a bandwidth restriction is shown in Fig. 1. In the illustration a video stream consists of three layers (Base, Enhancement 1, and Enhancement 2). A MANE reduces the bandwidth of the stream by discarding packets from the highest enhancement layer (E2). If further reduction is required, the lower enhancement layer (E1) would also be discarded. Transmitting only the base layer will deliver a lower visual quality.

Matching video streams to network capabilities helps to prevent issues such as buffering waits or heavily distorted pictures, which would affect a user's perception of quality.

This article provides both valuable insights into video stream adaptation using the new compression technologies of HEVC and SHVC, and the results of a comparative study (between current and evolving scalable video encoding schemes) conducted on a wireless networking platform. The relative merits of competing adaptive streaming technologies are discussed primarily within the context of mobile/portable devices, and an outline model of an SHVC-based streaming system is presented. Results of an experimental implementation of this model show that bandwidth savings of between 42 percent and 58 percent can be achieved using SHVC without any significant quality penalty. The limitations of the current SHVC design are also explored by evaluating its performance under adverse network conditions.

High Efficiency Video Coding

In 2010 the Joint Collaborative Team on Video Coding (JCT-VC) was formed with a goal of designing a new generation of video encoding standards that would offer a 50 percent bandwidth saving over the H.264 Advanced Video Coding (H.264/AVC) standard. In addition to improving coding efficiency, the JCT-VC also sought to deliver a standard that would improve integration with transport system protocols, include resilience to data loss, and be suitable for implementation on parallel processing architectures. Another important consideration in the codec design was the need to address the encoding and transmission requirements of emerging Ultra-High Definition Television (UHDTV) formats such as 4K (2160p) and 8K (4320p) with spatial resolutions of up to 16 times that of current HD (1080p) displays.

Target applications for the new standard included broadcast television at HD and higher resolutions over cable, satellite, and terrestrial transmission channels, video delivery over Internet and mobile networks, video-conferencing, and storage applications such as blue-ray discs. Over a period of almost three years many encoding tools, proposed mainly by industry, were evaluated, culminating in the adoption of HEVC as a new standard for video coding. HEVC retains the same two-layer hybrid coding structure as its predecessor (H.264/AVC), in which a Video Coding Layer (VCL) handles encoding and a Network Abstraction Layer (NAL) organises the HEVC-encoded bitstream as a series of logical data units (NAL units) for encapsulation and transmission.

Adaptive Video Streaming

Adapting video streams to meet device and network constraints has been the subject of extensive research in recent years. Two distinct approaches have evolved to deliver video streams to a diverse range of devices over fluctuating network channels.

DASH

Dynamic Adaptive Streaming over HTTP (DASH) [4] has recently become popular with content providers. In DASH, several representations (typically three) of the same video, each with a different bitrate, are either separately encoded or derived by transcoding, up-sampling, or down-sampling from a single representation. Each representation, which is contained in a separate file, is divided into segments marked in relation to their temporal position within the video sequence. The client device monitors the arrival of segments from the server over short time intervals. Using its knowledge of the current delivery rate and the bandwidth requirements for each representation, it selects which representation to request segments from during the next time interval. Switching between representations can take place frequently during streaming sessions where bandwidths fluctuate or losses and delays in the network lead to the need for retransmissions.

H.264/SVC

The alternative approach is to employ a single video file containing a number of layers that can be extracted, either independently or in combination, to match a given constraint (such as

bitrate, spatial resolution, etc.). This approach has previously been adopted as the scalable extension (H.264/SVC) [5] to the H.264/AVC standard. A similar extension (SHVC) is currently being developed by the JCT-VC for HEVC. H.264/SVC permits encoding of a video stream consisting of an H.264/AVC-compliant base layer, ensuring compatibility with legacy decoders, and a number of enhancement layers.

Three possible scalable dimensions (spatial, temporal, and quality) are available in H.264/SVC. Spatial scalability offers a course-grained mechanism whereby the H.264/AVC-compliant base layer is encoded at a lower spatial resolution with a spatial enhancement layer carrying the additional encoded picture data required to increase the number of pixels by a factor of four. For example, base layer = Height (H) × Width (W), and base layer + enhancement layer = 2H × 2W.

In temporal scalability, the base layer is encoded at a low frame rate with each enhancement layer increasing the frame rate by a factor of two. For instance, base layer = 7.5 frames per second (f/s), enhancement layer 1 = 15 f/s, and enhancement layer 2 = 30 f/s. Temporal scalability can be combined with spatial scalability with each spatial layer comprising a number of temporal layers.

Quality scalability alters the Quantization Parameter (QP) used by the encoder for each quality layer in a bitstream. A high QP will produce a more condensed bitstream, encoded using fewer bits but with a higher level of error (noise) in the resultant picture. In quality scalability the highest QP is chosen when encoding the base layer. Enhancement layers increase quality by reducing the QP, resulting in higher bitrates and reduced noise in the picture.

All three scalable dimensions can be combined to produce very granular scalable bitstreams containing many layers. Examples of each type of scalability are shown in Fig. 2. H.264/SVC scalable video coding lowers coding efficiency by up to 20 percent when compared with H.264/AVC [5].

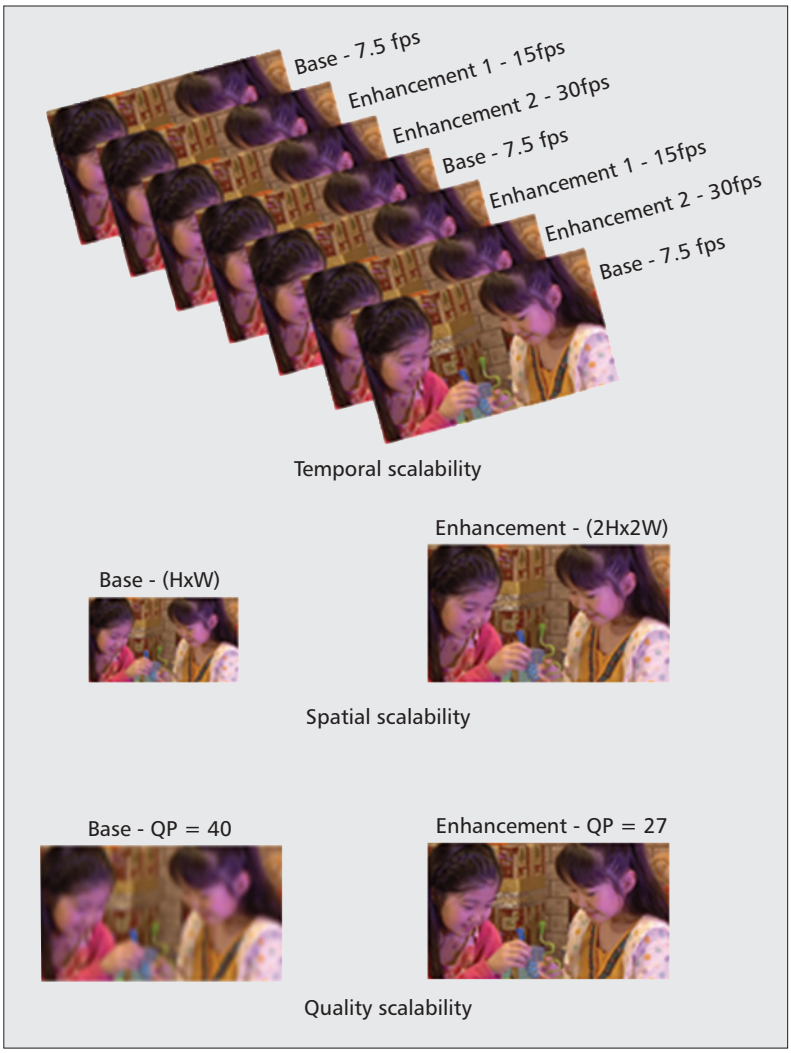


Figure 2. The three scalable modes. Temporal (frame rate) scalability is achieved by removing frames belonging to enhancement layers. In spatial (picture resolution) scalability the base layer has a lower spatial resolution and in quality scalability the base layer has a 'noisier' picture than an enhancement layer.

SHVC

SHVC has not yet reached the same level of maturity as H.264/SVC. The current working model only supports a two-layer scalable bitstream (base layer and a single enhancement layer). In addition to the three scalability dimensions employed in H.264/SVC, hybrid scalability, Region Of Interest (ROI) scalability, and color gamut scalability are all actively being considered for inclusion in the SHVC specification.

Temporal scalability has been incorporated into the HEVC design, allowing a specified number of temporal layers to be extracted as a standard-compliant bitstream. Similarly to H.264/SVC, SHVC quality scalability is achieved by varying the QP value of each layer. Spatial scalability in SHVC has two modes. The resolution of an SHVC spatial enhancement layer can be either twice the height and width of the base layer (2× mode) or one and a half times the height and width of the base layer (1.5× mode).

The first new scalability mode being considered for SHVC is hybrid scalability where the base layer complies with the H.264/AVC standard and the single enhancement layer is encoded by employing HEVC. In this hybrid mode, legacy H.264/AVC decoders will still be able to decode the base layer. Example applications for this mode may include broadcast transmission of television programs at greatly increased spatial resolutions such as UHDTV. Legacy TVs and set top boxes would be able to decode the base layer at the current HD quality, while new devices would decode the entire UHDTV stream. Some of the current proposals [6] being considered for the hybrid mode include having an HEVC-encoded temporal layer of the H.264/AVC base layer and an enhancement layer that offers combined scalability (temporal and spatial). The latter proposal would allow the enhancement layer to transmit streams with both a higher spatial resolution and an increased frame rate.

ROI scalability is another interesting proposal for SHVC, where the enhancement layer carries a higher quality representation of a selected

Although SHVC is still under development and no SHVC-specific MANE implementations have been reported, models previously proposed for H.264/SVC can be readily updated to utilize the new scalable encoding format.

region within the viewable area. There are many potential applications for ROI scalability in, for example, security surveillance and sports broadcasting. For security applications, an observer may select an area (or object) to enhance. This area would then have additional picture information transmitted in the enhancement layer to perhaps reduce the noise (quality scalability) in the selected region. This method offers the advantage of only carrying an additional bitrate burden for the chosen region of the picture in the enhancement layer, thereby saving bandwidth and reducing congestion in transmission channels.

Color gamut scalability recognizes that current HD standards (ITU-R Rec. BT 709) employ an 8-bit color encoding depth, whereas UHDTV applications will utilize a 10-bit or 12-bit color depth (ITU-R Rec. BT 2020). The aim of this scalability vector mode in HEVC is to permit a base layer that is encoded with an 8-bit color depth and an enhancement layer using 10-bit or 12-bit color depth. Legacy decoders would be able to decode the base layer only. This scalability mode is seen as a priority for the implementation of HEVC in broadcast transmission systems.

SHVC vs. DASH

Storage Requirements

Although both DASH and scalable video encoders such as H.264/SVC and SHVC address the same issues of stream adaptation to meet a set of resource constraints, they operate in significantly different ways. DASH requires that multiple representations of the same content are individually stored on the server, while scalable encoding requires only a single file. It is noted that although employing SHVC encoding reduces compression efficiency compared with HEVC in the range from 8 percent to 22 percent, dependent on encoding mode [7], SHVC layered bitstream files are actually significantly smaller than the sum of the multiple representations required by DASH. Disk space savings of 35 percent to 40 percent were observed in SHVC when compared with DASH (based on HEVC encoding) in our empirical studies. Hardware costs of content delivery networks, which typically have several copies of popular content distributed across the network, could be reduced by utilizing scalable encoding methods.

Flexibility

A second consideration when comparing DASH with scalable video encoding is the location within an end-to-end network path where video stream adaptation occurs. DASH-based streaming is controlled by the client device responding to changes in bandwidth and segment delivery. The choice of which representation (quality level) to request from the server during the next time interval is derived from delivery reports for the current time interval, whereas in SHVC and other scalable video streams, adaptation occurs at a MANE, which can conceptually be located at either the end point(s) of a transmission path or intermediate nodes such as media gateways or proxies. Network packets carrying SHVC bitstreams can be easily parsed with low complexity at intermediate nodes to determine the scalable layer identity of the packet contents. This approach is more flexible than that of DASH as it allows network operators to apply global traffic management policies to video streams traversing the network. DASH does, however, offer advantages over SHVC and H.264/SVC in that HTTP traffic can readily traverse firewalls and NAT access controls, and adaptation is performed at the client without the need for a MANE at the network side. The third, and less well researched, consideration when comparing DASH and scalable video streaming is that of a user's perceptual Quality of Experience (QoE) when delivery takes place over an imperfect transmission channel. Again SHVC and other scalable encoding schemes offer more flexibility in how they can be implemented than DASH does. SHVC-based streaming can be implemented over either reliable (e.g. TCP) or unreliable (e.g. UDP) transmission protocols to accommodate different use case scenarios. For example, it may be more appropriate to have an unreliable transport protocol with channel coding (e.g. Forward Error Correction (FEC)) for packet loss mitigation and SHVC-based adaptation in some real-time streaming scenarios where packet loss or interference in wireless media is an issue. In such a scenario reliable delivery mechanisms, including DASH, may experience significant buffering delays [8] due to large numbers of retransmissions. In contrast, SHVC-based schemes with error and loss compensation will deliver video content, although at a lower perceptual quality, while minimizing buffering waits. Objective comparisons between sent and received files using, for example, the Peak Signal to Noise Ratio (PSNR) metric may well score the reliable transport stream delivery higher than the unreliably delivered stream. However, in a subjective evaluation, viewers may find the buffering wait intolerable. Current standard methods of video quality assessment do not accurately reflect the detrimental effect of buffering on a user's QoE. The Video Quality Experts Group (VQEG) is currently considering new methods of subjective assessment for adaptive streaming mechanisms such as DASH, in which additional factors related to a user's perception of quality will be measured [9].

SHVC Pilot Implementation

Although SHVC is still under development and no SHVC-specific MANE implementations have been reported, models previously proposed for H.264/SVC can be readily updated to utilize the new scalable encoding format. The SHVC evaluation described here is implemented by building upon the H.264/SVC streaming framework proposed by Nightingale, Wang, and Grecos in [10].

The SHVC MANE is implemented as a software agent co-located and fully integrated with the streaming server. The combined streaming server and MANE consist of three principal components: network resource matching, scheduling, and dependency checking. Network resource matching ensures that only those layers that can be successfully delivered by the currently available bandwidth are transmitted. On a

Group Of Pictures (GOP) basis, the available bandwidth during the transmission period of the GOP is compared with the cumulative encoded bandwidth of all the layers in the scalable SHVC bitstream. Where the encoded bitrate exceeds the available bandwidth the enhancement layers are dropped recursively, beginning with the highest enhancement layer, until the encoded bitrate of the remaining layers is less than or equal to the available bandwidth. Our experimental evaluation employs a software agent located within the Wide Area Network (WAN) emulation nodes (Fig. 3) to inform the MANE of currently available bandwidth and end-to-end delay.

Conceptually, this monitoring and feedback mechanism could be placed at other point(s) in the network, including at the client device. Second, delays in the transmission network are accommodated by ensuring that only those packets anticipated to arrive at the client in time to be decoded and displayed without causing buffering waits are transmitted. This is achieved by permitting an initial start-up delay (buffering window) of 150ms at the client. The arrival time of each packet is estimated relative to the first packet in the stream. Only those packets estimated to arrive within the buffering window are transmitted. Finally, as with all recent video encoding standards, some frames are predicted from others and cannot be successfully decoded unless the frame(s) upon which they are dependent are available at the decoder. To conserve bandwidth, a dependency checking mechanism is employed to ensure that only packets containing video content with no unmet dependencies are transmitted. This three-pronged approach ensures that the streaming experiments adapt the SHVC streams efficiently under prevailing network conditions.

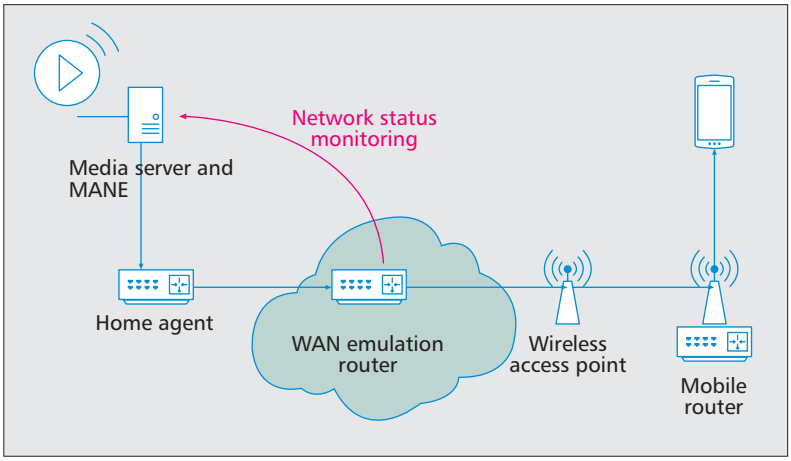


Figure 3. The wireless testbed used to compare SHVC and H.264/SVC stream adaptation. Stream adaptation takes place at the MANE, which is informed of the current network conditions by the WAN emulation router.

Sequence name	Spatial resolution	Frame rate
Traffic	2560 x 1600	30
PeopleonStreet	2560 x 1600	30
Cactus	1920 x 1080	50
BasketballDrive	1920 x 1080	50
Kimono	1920 x 1080	24

Table 1. Test sequences used for comparison.

COMPARING SHVC AND H.264/SVC

The following comparison of scalable video adaptation schemes employing H.264/SVC and SHVC highlights the bandwidth saving that could be achieved by adopting the newer SHVC encoding scheme. Video sequences drawn from the HEVC/SHVC standardization test sequences were encoded with both SHVC and H.264/SVC at two spatial resolutions chosen to be representative of popular display technologies. The full HD resolution (1920×1080) is widely available in a range of consumer devices including TVs, laptop computers, and high-end tablet/smartphone devices. A higher resolution (2560×1600) is currently popular on Apple retina displays and some other high-end portable devices.

The evaluation investigated spatial scalability and quality scalability, the two elements of SHVC that extend HEVC and have reference software implementations currently available. For spatial scalability, H.264/SVC sequences were encoded using a 2× configuration only. SHVC sequences were encoded using both 1.5× and 2× spatial scalability modes. Three QP testing points (22, 26, 30) were employed with both base layer and enhancement layer encoded at the same QP. The test sequences used are shown in Table 1.

The bandwidth requirements for each sequence encoded with 2× SHVC, 1.5× SHVC, and H.264/SVC spatial scalability are shown in Fig. 4 (for HD sequences) and Fig. 5 (2560×1600 sequences). We observed bandwidth savings in the range of 42 percent to 58 percent. On average, SHVC requires 47 percent less bandwidth than H.264/SVC. Spatial scalability in SHVC typically has a 12 percent lower bitrate than quality scalability. As with most video coding schemes, SHVC has different profiles designed for specific use cases. Although some variation in bandwidth requirement was seen between SHVC profiles, the differential with H.264/SVC remained consistent.

In both SHVC and H.264/SVC, quality scalability was evaluated at two testing points. In the first, the base layer was encoded with a QP of 32 and the enhancement layer with a QP of 30. The second testing point has a base layer QP of 22 and an enhancement layer QP of 20. Although employing a larger QP step leads to increased coding efficiency [11], this would be achieved at the expense of increased encoding time. When scalable streams are adapted at a MANE to reduce bandwidth requirements, employing spatial scalability results in a substantial reduction in bitrate of up to 63 percent, whereas using quality scalability, with a two QP differential, reduces the bitrate by a much lower margin of up to 28 percent. With respect to video quality, the objective PSNR metric was used to assess

The comparison between SHVC and H.264/SVC presented in this article has demonstrated the capabilities of an early implementation of SHVC adaptive streaming and highlighted its potential to improve streaming services in the consumer marketplace.

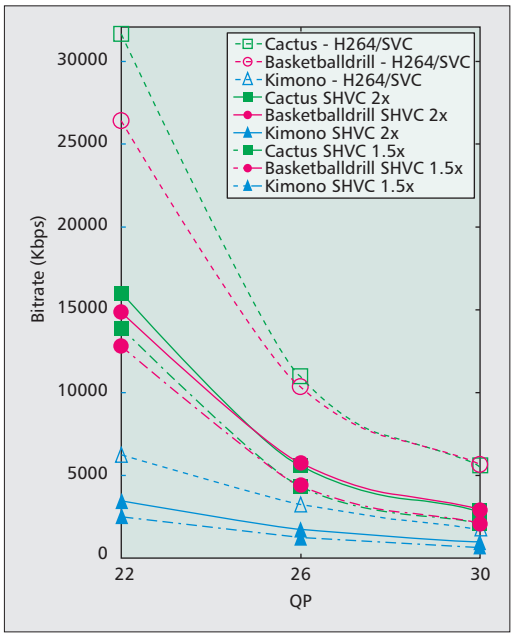


Figure 4. A comparison of bitrates for 1920 × 1080 sequences encoded using spatial scalability.

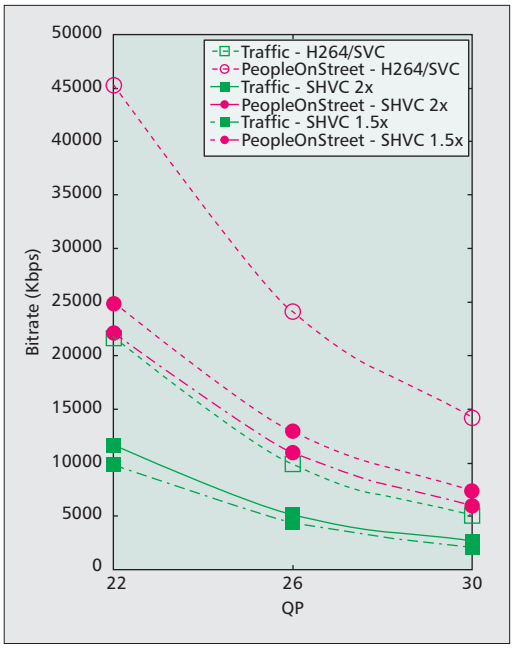


Figure 5. A comparison of bitrates for 2560×1600 sequences encoded using spatial scalability.

the impact of adaptation on scalable video streams, following the convention of JCT-VC in its comparative studies. Quality scalability produced better PSNR results with a modest degradation of 0.7dB of PSNR. Spatial scalability with a 2× configuration yielded an average PSNR reduction of 1.8 dB after adaptation. The drop in PSNR after adaption was not remarkably different between SHVC and H.264/SVC: on average H.264/SVC performed better than SHVC by a margin of 0.2 dB. However, when we applied a common bandwidth constraint of 3.3Mb/s to all of the encoded sequences, requiring that all but Kimono SHVC sequence be adapted by the MANE, the overall average quality of the H.264/SVC streams was 4.2dB lower than that of the SHVC streams. From these experiments, we can conclude that SHVC is capable of delivering high-quality video streams that need barely more than half the bandwidth of H.264/SVC with equal or better quality.

Opportunities and Challenges

The newly standardized H.265/HEVC and its proposed scalable extension will, when widely adopted, have a substantial impact on the design and development of the next-generation consumer products. The comparison between SHVC and H.264/SVC presented in this article has demonstrated the capabilities of an early implementation of SHVC adaptive streaming and highlighted its potential to improve streaming services in the consumer marketplace. Lower bandwidth requirements will mitigate the accelerating growth in video traffic and facilitate delivery of high-quality video over wireless channels. Higher spatial resolutions up to 8k will facilitate broadcast transmission of UHDTV content over terrestrial and satellite channels, and scalable video streams will enable transmission of the same content to many different consumer devices (from low-end phones to 8k televisions) over varied transmission networks. However, there are many challenges that will need to be addressed before SHVC-based streaming can be adopted as a core technology for consumer products. The first of these challenges arises from the computational complexity of HEVC and SHVC in comparison with previous encoding standards. Commercial broadcasters and content providers would need to embrace hardware-accelerated and/or distributed encoding architectures using parallel processing, perhaps cloud based, to facilitate encoding of UHDTV content and the additional representations required for scalable SHVC streaming or DASH. Real-time encoding and delivery of UHDTV content will pose particular challenges. Low-cost "on-chip" hardware decoders would be needed to drive the introduction of HEVC/ SHVC-enabled consumer products.

Although improved error resilience was one of the design objectives of the HEVC/SHVC standardization effort, considerable further work will be required to make robust, resilient delivery over impaired network channels a reality. In particular, robustness to packet loss in challenging wireless environments and novel methods of error concealment will be required.

Perhaps most interestingly of all, novel user-centric approaches to adaptive video streaming would need to be adopted. The current reliance on objective measurement of quality should be replaced by new metrics that fully consider the user's perception of not just visual quality, but his/her entire viewing experience. Current human vision system approaches to subjective or pseudo-subjective measurement are based on highly controlled laboratory experiments and largely ignore real-world factors. When considering immersive QoE, context factors (such as user

expectations of particular devices, service levels or contexts, for example, when travelling on public transport) that can detract from the user experience may also need to be explored. It may be time to consider moving to real-world 'total-experience' subjective evaluation of content delivery to everyday consumer devices in natural situations. By adopting user-centric approaches to service delivery, quality enhancements can be driven by user perception and expectation rather than the current quality of service mantras of availability, reliability, and integrity. Additionally, by employing a scalable video format it becomes practical to consider and measure the impact on a user's perception of quality when presented with different representations of the same video. The insights gleaned from such studies could be used, in conjunction with other factors such as video content type and prevailing network conditions, to develop new user-centered video streaming applications.

Acknowledgment

This work was partially funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/J014729/1: Enabler for Next-Generation Mobile Video Applications and the UWS-BUU DBasS project.

References

[1] Cisco, "Visual Networking Index: Forecast and Methodology, 20122017," white paper, 2013.

[2] G. J. Sullivan *et al*., "Overview of the High Efficiency Video Coding (HEVC) Standard," *IEEE Trans. Circuits Syst. Video Technol*., vol.22, no.12, Dec. 2012, pp. 1649–68.

[3] JCT-VC, "High Efficiency Video Coding (HEVC) Scalable Extension Draft 4," JCTVC-O1008_v3, 2013.

[4] ISO/IEC, Information Technology — Dynamic Adaptive Streaming over HTTP (DASH) — Part 1: Media Presentation Description and Segment Formats 23009-1, 2012.

[5] H. Schwarz, D. Marpe, and T. Wiegand, "Overview of the Scalable Video Coding Extension of the H.264/AVC Standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 17, no. 9, Sept. 2007, pp. 1103–20.

[6] JCT-VC, "Description of HEVC Scalable Extensions Core Experiment SCE1: Color Gamut and Bit-Depth Scalability," JCTVC-P1101_v2, 2014.

[7] J. Nightingale, Q. Wang, and C. Grecos, "SHVC Based Video Stream Adaptation in Hybrid Wired/Wireless Networks," *Proc. IEEE Int'l Symp. Personal Indoor and Mobile Radio (PIMRC) 2013*, 2013.

[8] I. Irondi, Q. Wang, and C. Grecos, "Empirical Evaluation of H.265/HEVC Based Real-Time Dynamic Adaptive Video Streaming over HTTP (HEVC-DASH)," *Proc. SPIE Photonics Europe 2014: Real-Time Image and Video Processing*, 2014.

[9] Minutes of VQEG Meeting Jan. 2014, available at http://www.its.bldrdoc.gov/vqeg/meetings/boulder,-co,-usa-january-21-25,-2014.aspx

[10] J. Nightingale, Q. Wang, and C. Grecos, "Optimized Transmission of H.264 Scalable Video Streams over Multiple Paths in Mobile Networks," *IEEE Trans. Consum. Electron*., vol. 56, no. 4, Nov. 2010, pp. 2161–69.

[11] B. Grkemli, Y. Adi, and A. M. Tekalp, "Effects of MGS Fragmentation, Slice Mode and Extraction Strategies on the Performance of SVC with Medium-Grained Scalability," *Proc. 17th IEEE Int'l Conf. Image Processing (ICIP)*, 2010, pp. 4201–04.

Biographies

James Nightingale (james.nightingale@uws.ac.uk) is a postdoctoral research fellow at the University of the West of Scotland (UWS), UK. His research interests include video streaming techniques, mobile networks, multihoming, and cloud computing. He is a member of IET and IEEE. He received his Ph.D. in video networking from UWS with the Outstanding Progression Award.

Qi Wang (qi.wang@uws.ac.uk) is a senior lecturer at the University of the West of Scotland (UWS), UK. He is the principal investigator of the UK EPSRC project "Enabler for Next-Generation Mobile Video Applications" (EP/J014729/1) and several industry-sponsored projects. He has over 70 publications on video networking, wireless/mobile networks, and cloud computing, which are his main research interests. He received his Ph.D. in mobile networking from the University of Plymouth, UK, with an ORS Award.

Christos Grecos (christos.grecos@uws.ac.uk) is a professor in visual communications standards, and the head of the School of Computing at the University of the West of Scotland (UWS), UK. His main areas of expertise are image and video compression algorithms and standards. He has over 150 publications, and has obtained significant research funding. His is a senior member of IEEE. He received his Ph.D. in image/video coding algorithms from the University of Glamorgan, UK.

Sergio Goma (sgoma@qti.qualcomm.com) is a senior director at Qualcomm Inc. in San Diego, USA, where he leads the multimedia R&D standardization group for imaging, driving Qualcomm's vision on future imaging technologies. His research interests include computational photography and programmable hardware architectures. Before joining Qualcomm he developed the image processing solution present in AMD's Imageon series of chips. He received a Ph.D. in reliability and fault tolerance of computers, and he holds several US patents on image processing algorithms and architectures.

By adopting user-centric approaches to service delivery, quality enhancements can be driven by user perception and expectation rather than the current quality of service mantras of availability, reliability and integrity.