Advanced HEVC encoding considerations

Ultilizing the full HEVC toolkit for greater efficiency and image quality

"H.265 was developed in response to the growing need for higher compression of moving pictures for various applications such as Internet streaming, communication, videoconferencing, digital media storage and television broadcasting...it is designed to enable the use of the coded video representation in a flexible manner for a wide variety of network environments"

Contents

Introduction

In recent years, consumer electronics manufacturers have released a flood of video consumption devices, ranging from large, high resolution televisions to powerful second screen viewing devices which can be found in every person's palm or pocket. Modern video consumption is ultimately the product not only of these spectacular viewing devices, but of a continuous drive for ever more efficient file formats that can deliver diverse and high-quality video services for those devices to consume.

Thus MPEG-1 yielded to MPEG-2, which became the foundation for digital television delivered to millions of cable, DBS and terrestrial viewers. During the last decade, MPEG-2 has had to share the stage with MPEG-4, also known as H.264 or AVC, which enabled greater diversity of services to be delivered to the home and provided the efficiency and quality needed to light up tablets, phones and PCs.

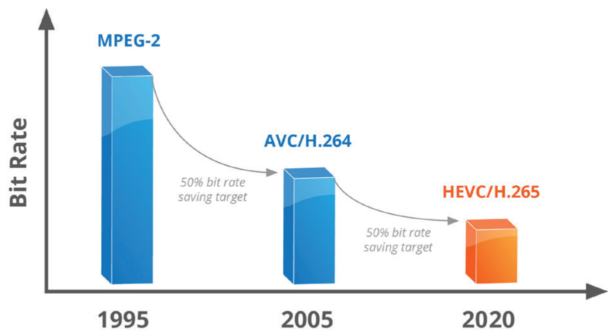

Figure 1: The evolution of professional codecs continues - each generation lasts about 10 years, and each generation promises a 50% reduction in bit rate. HEVC is the latest advance in video coding.

H.264 has been an enabling technology that has allowed for the birth and development of new viewing habits. It has supported the emergence of high definition television and has provided the flexibility and efficiency required for current distribution models. It has, however, largely reached its maximum coding efficiency limits. New technology is required for today's market.

The most recent ratified standard which offers improvements in both bandwidth usage (i.e. bit rates) and picture quality is High Efficiency Video Coding: HEVC (often referred to as H.265), which promises to deliver significantly higher quality video when compared to H.264 at the same bit rate, or to deliver the same video quality at 50% of the bit rate of H.264.

Obviously, the ability to meet these two major goals signifies a dramatic improvement when compared to current standards. HEVC offers the potential to reduce CDN and mobile carrier costs, improve quality, and make the delivery of 4K Ultra HD (3840 x 2160) possible in the real world. It also offers the potential to address the two most common technical frustrations expressed by consumers about watching video over the Internet: issues with the download stream – excessive "buffering" - and inferior video quality.

It should be noted that HEVC is still a relatively new technology, and like other compression standards, will continue to advance technically, moving it closer and closer to its maximum efficiency. At present, users of HEVC can expect to achieve around a 30% improvement in coding efficiency over H.264. As stated earlier, the ultimate target is a 50% improvement.

"HEVC is the first codec wholly designed in the modern video environment where file movement and delivery to the home is becoming the standard."

HEVC was developed in response to the growing need for higher compression of moving pictures for various applications such as Internet streaming, communication, videoconferencing, digital media storage and television broadcasting. In addition to the goals listed above, HEVC is also designed to enable the use of the coded video representation in a flexible manner for a wide variety of network environments. The HEVC standard was accepted by ITU-T in April 2013.

Benefits for the end-user

In addition to the technical requirements, the design of HEVC also accommodates (and enables) changes in the behavior and expectations of the end user. These changes include:

1. Changes in viewing habits. While AVC allowed us to get to today's changing landscape, it wasn't built from the ground up for the purpose of supporting the distribution environment we currently face and the additional services we envision for the near future.
2. The rise in prepared content as a percentage of total viewed content. As viewing has become increasingly nonlinear, opportunistic and multi-platform dependent, providers need to deliver valuable content without spending more for costly bandwidth infrastructure. That means taking full advantage of existing bandwidth.
3. 4k, and even 8k TV sets are becoming the norm. Whilst there has been little traditional programming created in these resolutions other than specialty sporting events, we are seeing the first signs of 4k resolution video services, and video distributors need a solution for the delivery of these high bitrate formats within their existing bandwidth constraints.

Benefits for the media enterprise

Although increased revenue from consumers is the most widely publicized goal, there are equally and potentially more immediately addressable applications of HEVC in the file-based operations on which media companies rely. HEVC is not simply a benefit for last-mile delivery–it has important and pressing advantages in the back-end movement of video.

"While the industry works toward the last mile distribution, there are important opportunities on the operational side of the equation that can deliver significant cost-savings and optimization"

One example can be seen in the flow of moving video material between an organization's facilities. In a recent report, the average cost of a satellite transponder (worldwide) was stated to be $1.62 million per year for 36 megahertz of capacity. At that price, it is incumbent on operators to make the most of their existing bandwidth. The ability to get more from a single transponder translates into a significant operational cost savings. So, when a network production center in Los Angeles needs to send material over a satellite to a distribution center in New York, HEVC provides the advantages of either moving better quality video (in the same transponder bandwidth), more video, or simply the same video at a lower bitrate.

Additional savings may be found in storage capacity for VOD and web distribution. HEVC provides the benefits of storing content in less space than current formats and/or storing higher quality content in storage capacity that is equivalent to what is used today. And while the benefits of monetized VOD services are predicated on consumer level devices being able to take that HEVC stream and decode it, such devices are already coming to market.

Using HEVC for delivery of prepared content

While there are potential uses for live HEVC encoding – and it has certainly been adopted for that purpose – there are constraints on a live encoder that prohibit making use of the entire HEVC toolkit to best effect. We will expand on this topic in more detail later, but there are a number of tools in the HEVC encoding toolkit that are less applicable to live encoders because they are restricted by the latency constraint that is mandatory in live world. This latency constraint, by definition, limits the number of frames which are available at any point in time for the encoder to use in its predictive activities.

Fortunately, an increasing amount of the material viewed by end users today is pre-prepared content. In truth, with the exception of live sports events or news, traditional broadcast relies on large quantities of pre-prepared material - and VOD and "catch up TV" operations are basically 100% pre-prepared.

When repurposed for re-use on websites, some live use cases also transform into pre-prepared content. For example, a news broadcast may be cut up into its constituent stories before posting to the web site and sporting events may well be reduced to more of a "highlight" package.

Freed of the constraints of live encoding, these use cases can take advantage of the advanced features of HEVC, use the full toolkit, and extract the maximum image quality (and revenue) from the source material.

The benefits of HEVC's enhanced compression capabilities can be significant in Multiscreen and VOD applications in areas other than inter-facility transfer and media storage. Consider the following real-world example: Delivery of a single 120-minute movie at Wi-Fi level bit rates (8.5Mb/s video, 64Kb/s audio) via a well-known CDN using AVC would cost approximately $0.441 per delivery. In this case, delivery to one million households would cost $441,000! HEVC offers the means to significantly reduce those delivery costs. A 50% reduction in bit rate, applied over many different titles for many different markets, soon adds up to a substantial reduction in cost of delivery.

Using less bandwidth is not the only opportunity, though. A more efficiently delivered video stream is less susceptible to stutters, stalls and buffering issues, which have all been demonstrated to improve the overall viewer experience, with clearly documented improvements in consumer retention rates - which again translates into increased profitability for the media enterprise.

It is easy to fall into the trap of oversimplifying the benefits of HEVC. HEVC is not only a solution for higher quality or greater efficiency - rather HEVC provides the flexibility to strike an optimal balance between the two, finding a "sweet spot" that makes the most sense for a particular business model, expanding the variety of services, enabling the introduction of new services and managing capital and operating expenses.

The benefits of HEVC are not binary and an operator is not forced to choose greater efficiency or higher quality. It is overly simplistic to reduce the benefits to the extremes when in reality it is a sliding scale. Operators can adjust the encoding to produce the video quality and bit rate that they believe fulfills their business needs.

How is HEVC Better?

All encoders work by removing repetitious (redundant) bits of data. There are two approaches to this.

The first is lossless compression, in which only the truly redundant information is removed (and is therefore generally limited to compression ratios in the 2:1-3:1 range). The second approach (and the one most commonly used) is lossy compression, in which not only is truly redundant information removed, but some fine detail is also removed –the idea being that the human visual system will largely be unable to tell the difference between the original and the modified versions.

This approach can yield much larger compression ratios. Determining the material which can be removed requires considerable computational power, and it is in this analysis that HEVC produces its performance gain. However, in a lossy scenario, as the codec approaches its maximum compression ratio, it is inevitable that greater loss of picture information must be accepted. Managing and masking the appearance of the resulting artifacts is the arena where encoding manufacturers differentiate themselves, and this is where the secret tools of pre-processing enable those errors to be hidden from view.

Like its predecessors, HEVC is a DCT-based compressions scheme, which means that it divides images into a number of blocks and then processes those blocks in turn. As a newer standard, however, HEVC offers several technical improvements over H.264. First, it can encode video in larger block sizes. HEVC uses "coding tree blocks" that have a maximum size of 64 x 64 pixels, whereas H.264 uses "macroblocks" with a maximum pixel size of 16 x 16. Don't be concerned about the terms "coding tree blocks" and "macroblocks" – they mean substantially the same thing. HEVC's larger blocks allow more efficient encoding, particularly for higher resolutions such as 4K. HEVC also offers more options for combining the blocks than H.264, which again means that it can offer substantial improvements in efficiency over its predecessor.

Some examples of the processor-intensive enhancements over H.264 that are available in the HEVC toolkit are:

- An increase in the number of directions for intra-frame encoding (within a frame). The intra encoding portion of HEVC looks for groups of pixels in a frame which are repeated elsewhere in that same frame (see Fig. 2)
- The use of 16-bit motion vectors, which allow for larger search areas in inter-frame encoding (across frames)
- The availability of larger coding tree units (macro-blocks) when compared to H.264, which can improve prediction efficiency
- Support for wider color gamut (broader representation of visible colors). This requires not only more processing, but also more memory
- Higher frame rates. HEVC allows for frame rates of up to 300fps. The results of an increased frame rate are quite dramatic. This is borne out by several recent studies by the BBC, Fox and others. This again requires more processing and memory

Figure 2: HEVC has significantly more search directions for intra-frame encoding. More directions means better prediction, and better prediction means lower bit rates and higher quality

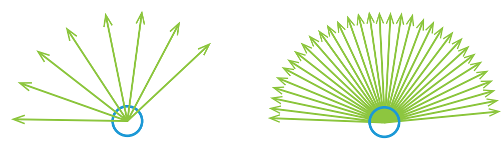

More directions means better prediction
Better prediction means **lower bit rates and higher quality**

Additional technical improvements include Adaptive Motion Vector Prediction, which increases compression efficiency by locating more inter-frame redundancies ("Inter-frame" means looking for redundancies between frames in the video, rather than within a single frame), and Wavefront parallel processing, which allows parallel processing to be efficiently harnessed in multi-core HPC or GPU-based environments.

HEVC currently offers a number of encoding profiles, of which the most important are:

- Main: 8-bits per color with 4:2:0 chroma format
- Main 10: 10-bits per color with 4:2:0 chroma format, which significantly reduces banding over 8-bit
- Main 4:4:4: This profile offers bit depths of up to 16bits per component, and up to 4:4:4 sampling, which makes HEVC the perfect vehicle for the encoding and transmission of HDR

What Markets Benefit from HEVC?

Any business that uses video can expect to benefit from HEVC – particularly those in the video distribution vertical, such as:

- Broadcasters
- IPTV companies
- Satellite and cable companies
- OTT program suppliers like Netflix
- New Ultra HD program networks
- Wireless companies for TV on mobile devices
- Teleconferencing services

As the standard evolves, everyone who watches videos or posts them online is likely to benefit from HEVC. Those who pay for bandwidth may stand to benefit the most due to the reduced costs incurred for the delivery of each movie or program.

The Benefits of looking into the future

One approach that has been used to maximize image quality from today's encoders is through "look ahead – further in time down the video stream to determine how often you can reuse the same data. This enables the encoder to effectively determine if there is an opportunity to reuse image information, which can then be turned into bit rate savings by not repeatedly resending that data. The efficiency of this approach is directly proportional to the number of "future frames" available to the encoder. In live applications, the number of "future" frames is limited by the number of frames available in the encoder's video buffer, and the buffer size is limited by the latency acceptable to the live application. In the encoding of prepared material, however, that limitation is removed–we have all of the frames that make up the entire clip (or section thereof), and can theoretically analyze the actual image content of all of them – although in practice there will still be limits to the total size of the lookahead.

This approach was used extensively in the past in Telestream's implementation of the x264 codec. This advanced H.264 codec used a concept called a "macro-block tree" to significantly improve efficiency and image quality (see Fig. 3). The macroblock tree is effectively a map for each macroblock in the current frame, which indicates where that macroblock can be used again. This may be in the next frame, or it may not be required for a second, but then be repeatedly used in the next two seconds worth of frames.

By analyzing the reuse of the macroblocks and ranking them in reuse, the encoder can then apply more of the available bits to the encoding of that macroblock and less to macroblocks that may only be reused once (or potentially not at all). This will result in a proportional increase in the video quality of the encoded picture, and an associated reduction in error propagation.

While this technology was developed for the x264 codec, the same technology and techniques are being carried forward into the x265 encoder detailed below (although to be accurate, the term "macroblock" is replaced by the term "coding tree unit" in x265). Looking seconds ahead on the video unlocks the full capability of the encoder 's capacity for efficiency and image quality.

It is important to note that these benefits are not processor dependent. These benefits accrue regardless of the hardware because the only barrier is the availability of subsequent frames of video for macro-block evaluation. The method's success is not a function of processing but simply one of data availability.
The non-transient nature of prepared material also allows for multi-pass encoding. The details of this technique are well known and beyond the scope of this paper, but in a nutshell, multi-pass encoding involves performing a first pass at encoding the material, then going back and examining the results of that encode - some frames or groups of frames are very difficult to encode and may show some artifacts from the applied compression. Others are much easier to encode and may produce quality measurements far in excess of the target value. By analyzing this, it is possible to "borrow" some of the bit budget from the easy sections, and give them to the difficult sections, which you can now re-encode with a lower compression ratio, improving the quality of the overall clip.

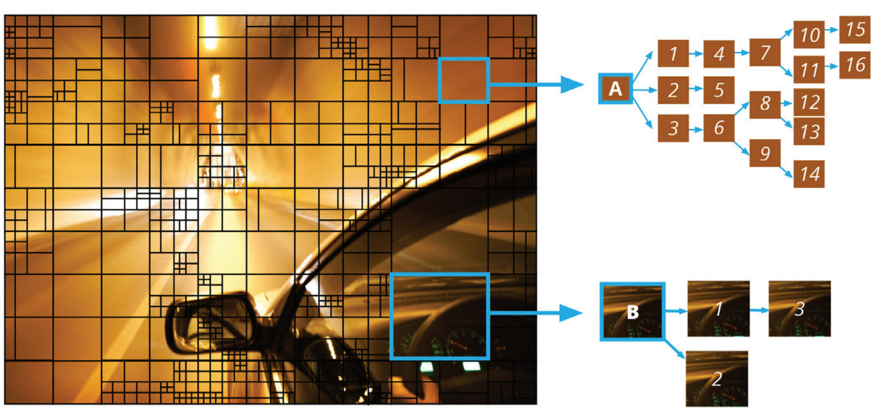

Figure 3: Lookahead can significantly improve performance. Here, we see a pair of macroblock trees, looking at 2 of the macroblocks in a frame. Macroblock A can be reused in 16 subsequent frames, Macroblock B only in 3. Clearly, budgeting more bits for Macroblock A will have greater impact on the overall clip

Collaboration breeds success

Developing an encoder is a very complex and costly activity. As a result, three distinct methodologies have emerged amongst encoder manufacturers:

1. Many manufacturers outsource development of their codecs to a third party that specializes in that technology. In fact, the majority of encoders rely on codecs created by a single vendor. While this approach can be cost-effective, it necessarily entails ceding control over development to a third party and accepting that the codec will be fundamentally similar to many competing products.

2. A second approach is for the developer to design the codec entirely in-house. While this approach gives the manufacturer full control, it is costly and limits the capabilities to the expertise of the staff on hand.

3. A third and highly compelling approach is to adopt the open source model. This is the approach that has been successfully followed by Telestream. Instead of limiting itself to an R&D department that fits into one floor of an office building, open source development utilizes the expertise of a global community of experts that span the world. Leading experts in universities and private labs bring incredible creativity and knowledge to bear on the challenge. In the case of HEVC, this approach has given rise to the x265 organization, of which Telestream is privileged to be a founding, funding member.

This approach makes it possible to build on the tremendous power of the best of the x264 technology, also developed by an Open Source team exceeding 110 in number, and apply these techniques to the x265 encoder design.

Adoption: What Issues Are Yet to Be Resolved?

Already there are over 1 billion smart phones and tablets capable of playing HEVC video simply by upgrading their software-based video codecs. Some more entrenched technologies, such as set-top cable and satellite boxes, will require new hardware. This is not at all unusual, however the same is true for any Ultra-HD services that emerge.

It is for those reasons that the earliest adopters of HEVC are likely to be wireless providers for the delivery of HDTV.

Not only are there HEVC-capable devices already in users' hands, but these devices turn over much more rapidly than TVs and other residence-based video platforms are replaced (this is especially true of cable set top boxes), and each of these replacement cycles bring greater and greater compute power. Additionally, wireless operators are facing increasing bandwidth challenges as video becomes more ubiquitous on their LTE networks, so HEVC can provide a real solution for those use cases.

For broadcasters, use of HEVC is expected to grow with the release of ATSC 3.0 – to date there has been little clear need for them to adopt the new format, but the ATSC 3.0 standard requires HEVC Main 10 profile.

The rate of adoption may still be affected by business factors, such as whether HEVC is included in upgrades to popular Internet media players such as the Flash Player. The major barrier for many companies remains the issues of royalty and licensing structures – many of which still remain to be resolved.

Conclusion

HEVC is a highly advanced encoding standard, offering the possibility of dramatic improvements in compression efficiency and deliverable image quality, and providing significant benefits to consumers and media enterprises alike. HEVC has an extensive set of tools in its arsenal, but some aspects require that the encoder have access to significant numbers of frames in the asset in order to make the best use of the available bits, and that requirement means that live encoders may not be able to make full use of all of the tools available to them.

An analysis of the media consumption models shows that most material being consumed has a high proportion of prepared material, however, and this material does allow the encoding process to make use of the advanced processing options available in both x264 and x265. The algorithms developed as part of the open source x264 project are being brought to bear in the x265 implementation, and Telestream offers this best of breed implementation in all of its HEVC-capable products, both CPU and GPU based.

To Learn More

Visit www.telestream.net or http://x265.org, or call Telestream at 1.530.470.1300

More information on the x265 project can be found at x265.org.