divideon
digital video compression

XVC
a world-class video codec with indemnification

Introducing xvc

A novel video compression format with a revolutionary licensing model

Abstract

The digital video compression market has seen several successful video codecs during the last 25 years, in particular the ones produced by MPEG and VCEG, such as MPEG-2 and AVC/H.264. However, despite the fact that video consumption and video distribution is constantly increasing the requirements for efficient compression, an uncertain and expensive licensing situation has significantly held back deployment of the most recent video codec from MPEG/VCEG called HEVC. This paper presents the novel video codec xvc which is based on MPEG/VCEG technology and which comes with a revolutionary licensing model.

Introduction

Digital video compression has been a key component in the transformation and digitalization of the TV and media industry. Video coding standards such as MPEG-2 (Part 2) and AVC (MPEG-4 Part 10 / H.264) have enabled a wide variety of services for digital media consumption ranging from for example broadcast TV and Video-on-Demand streaming to user generated live-streaming applications. In January 2013, MPEG (formally known as ISO/IEC JTC1/SC29/WG11) and VCEG (formally known as ITU-T Q.6/SG16) completed the video coding standard HEVC [1], a new codec capable of delivering around 50% bit rate reduction compared to AVC/H.264 [2] with maintained video quality.

In this paper, background information of the video codec licensing situation is presented as well as an analysis of the components that contribute to making the video codec licensing situation so difficult. The last part of this paper introduces the novel video codec xvc, presenting the concept, the technology, and the revolutionary licensing model.

Video codec licensing background

The processes for compressing (encoding) and decompressing (decoding) digital video is based on a large number of computationally advanced steps. The overview paper of HEVC [3] shows that HEVC consists of a multitude of elements and the HEVC performance paper [4] shows that these elements combined makes it possible to encode uncompressed video to a fraction of its original size without significant loss of video quality. Most of these elements and processing steps are the result of a vast amount of research and many of the elements have been patented by the organization that first developed them.

The ability to patent, and thereby protect, video coding technology makes it possible for technology providers to disclose their innovative methods (e.g. as contributions to MPEG/VCEG standardization) knowing that they are able to get something in return for their contribution (e.g. a part of the royalty revenue).

Standards developing organizations (SDOs) such as ISO/IEC and ITU-T (the parent bodies of MPEG and VCEG, respectively) typically require contributions to be accompanied with a promise that the technology is made available under Fair, Reasonable And Non-Discriminatory (FRAND) terms. FRAND declarations are meant to ensure that no organization that has contributed to the standardization work can make unreasonable claims when the standard is completed. Although FRAND is a sound and in many ways fundamental principle, it has turned out to be very difficult to implement in practice, especially when it comes to determining a reasonable price level for using individual patented methods included in a standard [5].

In addition to the difficulties associated with the FRAND concept itself, there are, at least within ISO/IEC (the parent body of MPEG), additional difficulties associated with how patented technology is being handled during standardization of video codecs. One of the most striking problems is that patent declarations in ISO/IEC does not need to identify which patent(s) the patent declaration relates to [6].

Patent hold-up and royalty stacking

In many cases it is to the advantage of a patent holder to wait and see how popular a standard becomes before announcing the terms under which that patent holder's patents may be exercised. Once a standard has become deployed, the patent holders are in a very good position to dictate their terms since the users of the standard will be very reluctant to stop using the standard because the alternatives (removing or replacing the standard with a different standard) is typically very costly and in many cases not at all practical. This problem is known as the patent hold-up problem.

Another, related problem, is the royalty stacking problem. A standard typically contains hundreds or even thousands of patented features and even though each individual royalty claim may seem reasonable in isolation, the aggregation of all royalties may reach a completely unreasonable level. For example, when multiple patent pools are formed for the same standard, each of them might have a sensible cap for how much a licensee at most would have to pay to that pool. But if these caps are added together the amount might become unreasonable.

The problems of patent hold-up and royalty stacking are by no means unique to video codec standards. The same problems can be found for example related to mobile radio technology standards and wireless networking standards [5]. However, for video compression, the combination of highly capable execution platforms (e.g. powerful Central Processing Units – CPUs) and upgradable connected devices gives rise to the ground-breaking licensing model which is applied for xvc and which is described in more detail below.

Suggested solutions

The problems of patent hold-up and royalty stacking are very severe and are continuously threatening to delay and hold back not only deployment of standards but also innovation and technology evolution in general. When the licensing situation is unclear (because of patent hold-up) and expensive (because of royalty stacking), organizations are reluctant to start deploying a new standard – even if the new standard comes which large benefits such as reduced costs and/or new services. The result is that the uptake of the standard is low and slow and the return of investment for the developers of the standard becomes low as well – i.e. everybody loses.

The following text provides more in-depth information relating to suggestions and ideas that have been brought up as options that could potentially improve the patent licensing situation. All of these options constitute relevant background information for xvc, but readers that are more interested in the actual xvc codec and licensing scheme may jump directly to the Introducing XVC section.

There have been several suggestions for how to approach the problems of patent hold-up and royalty stacking, both in academia, within standardization organizations and in the industry. Some of the most common suggested solutions can be grouped into four categories as follows.

The first category of suggested solutions look to the completely royalty-free path. The idea is to create standards that do not include any technology which is encumbered with patents for which the patent holder would require royalty. One attempt to realize this is demonstrated in the MPEG standard for Internet Video Coding (IVC) [7] which is primarily built on technology which is old enough to ensure that possible patents on the technology would no longer be in force. Another attempt is the VP8, VP9 and the upcoming AV1 codecs [8], which are not produced by an SDO. Instead, one or more companies have made agreements that they will not be charging royalties for their patented technologies that are included in the codec.

The largest problem with the royalty-free approaches is that they do not reach the same level of performance as their royalty-bearing counterparts [9][10]. The codecs become suboptimal because of the constraint that they can only use a subset of the known technologies (i.e. those that are available on a royalty-free basis). Another problem is the uncertainty of whether the codec infringes third party patents or not, i.e. whether the codec in the end is royalty free or not. A third problem with royalty-free codecs is that they typically comes with requirements of reciprocity or grant-back, meaning that technology providers that have patented technology in the codec will have to make those available under royalty-free terms otherwise they will not be allowed to use the codec. This creates a segmentation between those that *would like to get royalty for their technology investments and those that are prepared to give away their technology for free. The risk is that this hinders deployment and/or results in long and costly litigations. In the long run, the royalty-free approach can have negative effect on the evolution of new technology and new standards since the incentive to develop and contribute new technology is low and technology providers may stop contributing their technology or stop performing research in the field altogether.*

A second category of suggested solutions look towards a hybrid- or semi-royalty free model. This is for example brought up by Tom Vaughan, VP and GM, Video at MulticoreWare [11]. The suggestion is to make certain classes of usage or certain codec instances royalty free. In particular, the suggestion is to make software instances royalty-free while hardware instances would still be subject to royalties. Another type of hybrid/semi royalty free is the idea of creating a subset (e.g. a profile) within a standard which is royalty free, while other parts of the standard would be royalty encumbered.

While hybrid- or semi-royalty free models might help adoption and might make the licensing situation better for certain cases, they do not really address the most central parts of the licensing problem. In particular, they do not give guidance on how to set the royalty rates for the cases that are still royalty bearing and they do not provide answers to how patent holders should be encouraged (or forced) to declare their licensing terms in reasonable time. With hybrid- and semi-royalty free models, new problems arise, such as which parts of the codec should be made royalty free under which conditions, and for the parts that are supposed to be royalty free, all the problems from the first category applies.

A third category of suggested solutions look towards ex ante disclosure of licensing terms. This suggestion is brought up in [12] and is already being deployed in the SDO VITA. The core idea with ex ante disclosure is that patent holders declare licensing terms before the standard is published and thereby making it possible for implementers of the standard to understand the cost of the licenses upfront.

Ex ante declarations do provide a good starting point for avoiding both patent hold-up and royalty stacking since the royalty terms will be known from day one when the standard is published. Knowing the licensing terms during the development of the standard makes it possible, in theory, to remove technology for which the licensing terms are deemed unreasonable or undesirable. However, by itself, an ex ante declaration does not solve the difficult licensing situation completely. One of the major difficulties relates to how to determine if declared terms are in fact reasonable or not. Making this type of analysis within a technical working group of an SDO may not be suitable. Another problem is that the ex ante declarations will only be made by organizations that actively contribute during the standardization phase. It is very well possible that the standard will include technology from other organizations especially in the field of video compression where a large *number of organizations hold patents and only a subset of them contribute to the standardization work.*

A fourth category of suggested solutions look at standards that are made more flexible and which may be tuned or revised based on the licensing situation (or potential legal actions taken by patent holders). A concrete suggestion along these lines is brought up in JCTVC-X0034 [13]. The suggestion in JCTVC-X0034 is to make it possible to disable each individual feature of a video codec by control information included in the compressed bitstream itself. It is argued that if a user of a codec is approached by a "troll" requiring unreasonable royalty rates for using their patent, then the user can turn off that feature and stop using it, and thereby no longer infringe the patented technology. It is asserted that this can be performed in "walled-garden" systems and that different standards groups and consortia may agree on "dialing-out" such features.

While the redesigning approach provides a good foundation for a standard that is flexible enough to avoid unreasonable royalty claims, the suggestion in JCTVC-X0034 does not provide much guidance when it comes to the process for structuring and aligning decisions of "dialing-out" features. If different consortia and organizations make different decisions on which features to dial-out, the result would be incompatible "sub-profiles" which in essence would constitute different incompatible codecs.

In the next section a novel licensing scheme is presented, which borrows primarily from the third and fourth category of solutions and which fills in the gaps which are missing in previously suggested solutions.

Introducing XVC

The concept

The xvc codec is a novel video codec which is brought to market in its first version by the software video compression company Divideon. The codec is built on technology that has been investigated thoroughly in MPEG and VCEG. A large number of the features included in xvc can also be found in very similar form in HEVC, but xvc does include a few features that provides improved performance compared to HEVC as well as some simplifications which reduces overhead in the codec compared to HEVC.

The xvc codec is not a standard (yet) and it is not defined by any specification document. Instead the codec is defined by the xvc reference software using the following definition: "A bitstream is a conforming xvc bitstream if and only if the current version of the reference xvc decoder successfully decodes the bitstream and returns *Conformance verified*."

A key element in the conformance definition is the word "current" which highlights the fact that conformance is limited to the current point in time (the current version) and no guarantees are provided for future or past versions of the reference software. By using this definition the codec can continuously evolve, and it is possible to add functionality and coding tools and, perhaps even more importantly, remove coding tools if they cannot be licensed under reasonable terms.

The version control, and addition and removal of individual tools is handled by the team responsible for the reference software (which initially consists of the software development group at Divideon) in collaboration with users of the codec (i.e. licensees) and technology providers (i.e. licensors / patent holders).

The license

The xvc codec is being made available with a “one-stop shop” license offered by Divideon. The license covers both the copyright of the software and all patent rights exercised by the software. Patent holders that have registered their relevant xvc patents with Divideon are entitled a share of the licensing revenue proportional to their patent footprint in xvc.

A key difference with the xvc license and licenses offered for other video codecs is that the xvc license comes with a guarantee that the licensee does not need to acquire any other licenses or pay any third-party patent holders in order to use the codec. How is this possible? It relates to the fact that the codec definition is non-static and that it is possible to remove and/or replace features if deemed necessary.

The xvc license comes with an indemnification clause saying that if a licensee is approached by any patent-holding entity asserting that the licensee makes infringement of the patent holders patent by using xvc, then Divideon guarantees that one of the following actions will be performed within 60 days from the date when Divideon is informed about the infringement assertion:

a) The patent in question will be added to the list of patents covered by the xvc license (e.g. by making the patent holder become an xvc licensor or by making an xvc licensor acquire the patent in question).

b) The infringement assertion is deemed inaccurate and Divideon accepts to pay all costs (including costs for litigation and damage) if a court finds the alleged infringement to be valid.

c) The patented technology is removed from the xvc codec and a new version of the codec is issued, with a new version of the reference software being made available within 60 days from the date when Divideon is informed about the infringement assertion.

An infringement assertion needs to be specific both in terms of which claim(s) of which patent is infringed and in terms of which function(s) of the reference software infringes the patent.

The technology

The xvc codec is a block-based hybrid (inter/intra) video coding format built using world-class compression technologies. The codec operates on raw YUV pictures and compresses them to a NAL (network abstraction layer) unit structured bitstream. Each picture in a video sequence is divided into rectangular blocks of samples, which are predicted from samples in the same picture (intra prediction) or samples in previously coded pictures (inter prediction). Residuals are transformed using non-square DCT-like transforms and the coded symbols are compressed using a context-adaptive binary arithmetic coder. Block boundaries are filtered using a deblocking filter.

Each of these parts of the codec consists of several individual features, and in the xvc codec each of these individual features can be turned off via control information in the bitstream. First of all, this makes it possible to adjust and tune the encoding operations so that an encoder can be tuned towards specific use cases or towards specific speed levels. Secondly, this makes it extremely easy to disable features that cannot be used due to licensing problems (fully aligned with the suggestion from JCTVC-X0034).

The xvc codec has been designed with streaming services as the primary use-case, partly because this is an area which is growing rapidly and which requires very efficient compression, but also because this is an area in which the clients are network connected and generally has support for being updated remotely. Streaming service providers, such as Netflix have emphasized the importance of more efficient compression, not only for high fidelity and high resolution video, but perhaps even more urgently for low bit rate use cases. When Anne Aaron from Netflix lists the biggest challenges for video compression at NAB 2017, 100 kbps streaming comes first in the list [14].

The xvc codec comes with a slim High-Level Syntax layer and with novel methods for efficiently handling transitions between representations in adaptive bit rate (ABR) applications.

VIDEO FORMAT SUPPORT

The xvc codec is extremely flexible in terms of supported video formats. Picture sizes can range from for example 256x144 and 320x180 to 4K, 8K and even 16K video. The encoded video can even be higher than the screen resolution of the target device - xvc has inbuilt support for resampling decoded pictures to desired resolution. The same goes for bitdepth and chroma format. The xvc codec has native support for various combinations of chroma formats and bitdepths (as shown in Table 1), but the output format of the decoder can be controlled via decoder configuration settings.

	Bitdepth		
Chroma format	8	10	12
Monochrome	Yes	Yes	Yes
4:2:0	Yes	Yes	Yes
4:2:2	Yes	Yes	Yes
4:4:4	Yes	Yes	Yes

Table 1. Supported video formats in xvc.

THE WEBSITE AND THE SOURCE CODE

The source code for the xvc reference software is publicly available on GitHub [15]. The source code is open for everyone to try out and evaluate, but commercial usage of xvc is constrained as described by the xvc license. With the source code publicly available it is easy for implementers to understand the code and integrate xvc into different platforms. It is also easy for patent holders to check if their methods are exercised by the xvc reference software. The xvc website (xvc.io) [16] provides links to the software, along with additional information related to xvc usage and licenses.

XVC INSTANCES

The license for xvc covers all the methods that are included in the reference software, both encoding methods and decoding methods. But the license is not limited to the specific implementation constituted by the xvc reference software. Any implementation which operates on conforming xvc bitstreams is covered by the xvc license to the extent that it uses methods included in the xvc reference software. For obvious reasons, it is unfortunately not possible to to make provisions for methods beyond the reference software.

It should be noted that implementations of xvc may be completely software-based, completely hardware-based, or a hybrid of the two. Implementers of xvc should be aware of the flexible nature of the xvc codec, and are encouraged to design implementations in a way which makes modifications and upgrades possible.

Continuous evolution is of course much easier in software implementations, and the xvc reference software itself constitutes a good starting point for commercial deployment. The xvc reference software can already do real time decoding of HD video at 30 fps on mobile devices.

Software implementations of video decoders typically consume more power compared to hardware implementations, which on mobile devices results in reduced battery life. How large this effect is depends on the platform as well as the resolution and bit rate of the video. Divideon has created an xvc demo which for example decodes and displays HD video consecutively for 5 hours on an Android 6 mobile phone. For really low bit rate streaming (such as 100 kbps video [14]) the power consumption for decoding becomes negligible.

divideon
digital video compression

XVC
a world-class video codec
with indemnification

Performance

The single most important aspect of a video codec is its ability to compress video sequences to as small size as possible with as high visual quality as possible. The xvc codec provides similar or slightly better quality than HEVC, as reported below. By extrapolating the results from [2] and [4] it can be concluded that xvc can cut the bit rate in half compared to AVC/H.264 without reduction of visual quality. Correspondingly, when operating at the same bit rate, xvc offers significantly better quality compared to AVC/H.264. The example in Figure 1 shows a visual comparison of xvc compared to the popular AVC/H.264 encoder x264. The screenshots are cropped out from video with a resolution of 1280x512 encoded at 300 kbps.

Figure 1. Visual comparison of x264 (top) and xvc (bottom) at 300 kbps. Video quality can also be measured using various objective quality metrics and even though objective metrics don't always correlate perfectly with

subjective quality, they provide the means for measuring small and incremental improvements of video codecs. Such improvements can vary between different sequences and different bit rates, but one well accepted approach for reporting an average number is to use the Bjontegaard-Delta method [17] which interpolates PSNR scores over several rate points. The xvc reference software has been tested against the HEVC reference software, HM-16.15 using the test conditions described in JCTVC-Z1100 [18]. The results can be seen in Table 1 with an average bit rate saving of 4%. The table also reports average differences in encoding time and decoding time for the platform which the software is run on. The reported numbers are from an Intel Xeon E5-2660 CPU.

	Random Access Configuration		
	Y	U	V
Class A	-1.6%	-1.0%	-2.4%
Class B	-2.3%	2.0%	0.0%
Class C	-6.2%	2.0%	2.6%
Class D	-6.3%	0.0%	-0.4%
Overall	-4.0%	0.8%	0.0%
Enc Time[%]		4,00%	
Dec Time[%]		-49,00%	

Table 1. Performance results of xvc compared to HM (HEVC) using the test conditions of JCTVC-Z1100 [**18**].

The best way to determine the performance of xvc is of course to try it out, applying the constraints and settings that are valid for the use case that is being evaluated. This is one of the reasons why the xvc reference software is publicly available and free of charge for evaluation purposes.

How to contribute

The xvc codec is a continuously evolving codec and anyone is welcome to suggest improvements and additional functionality. For the time being, Divideon is responsible for the development of the codec and the software, but the hope is that the xvc codec (and its flexible concept) could be included in a future standard.

Conclusion

The xvc codec is a novel video compression format with a revolutionary licensing model. It is the only modern video codec that is available with a "one stop shop" license which includes an indemnification clause to protect against third party infringement assertions. The xvc codec is already the best performing codec available to date, and with a codec structure that is flexible and continuously evolving, xvc is set out to become even more efficient in the future.

References

1. Recommendation ITU-T H.265 | ISO/IEC 23008-2 "High Efficiency Video Coding".
2. T.-K. Tan, M. Mrak, V. Baroncini, N. Ramzan, "Report on HEVC compression performance verification testing", JCTVC-Q1011, May 2014.
3. G. J. Sullivan, J. Ohm, W.-J. Han, and T. Wiegand, "Overview of the High Efficiency Video Coding (HEVC) standard", IEEE Transactions on Circuits and Systems for Video Technology, vol. 22, no. 12, pp. 1649–1668, Dec. 2012.
4. J.-R. Ohm, G. J. Sullivan, H. Schwarz, T. K. Tan, and T. Wiegand, "Comparison of the coding efficiency of video coding standards— Including High Efficiency Video Coding (HEVC)", IEEE Transactions on Circuits and Systems for Video Technology, vol. 22, no. 12, pp. 1668–1683, Dec. 2012.
5. M. Lemley and C. Shapiro, "Patent Holdup and Royalty Stacking", Texas Law Review. Vol. 85:1991, 2007.
6. L. Chiariglione, "Standards for the present and the future", Blog post accessed on 2017-08-15, http://blog.chiariglione.org/2017/07/
Jul. 2017
7. ISO/IEC 14496-33 "Internet Video Coding"
8. Alliance for Open Media, http://aomedia.org/
9. R. Wang; Z. Wang; K. Fan; T. Huang; W. Wang; G. Li; W. Gao, "MPEG Internet Video Coding Standard and its Performance Evaluation" IEEE Transactions on Circuits and Systems for Video Technology , 2016.
10. J. Ozer. "Netflix Finds x265 20% More Efficient than VP9" Web article accessed on 2017-08-15, http://www.streamingmedia.com/Articles/Editorial/Featured-Articles/Netflix-Finds-x265-20-More-Efficient-than-VP9-113346.aspx
Sep. 2016
11. T. Vaughan, "A Proposal to Accelerate HEVC Adoption" Blog post accessed on 2017-08-15, http://x265.org/proposal-accelerate-hevc-adoption/
Aug. 2016.
12. J. L. Contreras, "Fixing FRAND: A Pseudo-Pool Approach to Standards-Based Patent Licensing." Antitrust Law Journal, vol. 79, no. 1, pp. 47–97, 2013.
13. S. Wenger, "On profiles and per-feature Flags" JCTVC-X0034, Geneva, May 2016.
14. A. Aaron, Blog post on twitter accessed on 2017-08-15: https://twitter.com/AnneMargotAaron/status/857033832763072512
Apr. 2017.
15. GIT repository: https://github.com/divideon/xvc
16. xvc website: https://xvc.io/
17. G. Bjøntegaard, "Calculation of Average PSNR Differences Between RD Curves", VCEG-M33, Austin, Apr. 2001.
18. K. Sharman, K. Suehring, "Common test conditions" JCTVC-Z1100, Geneva, Jan. 2017.