EAVC4 "The Ultimate H.264 Encoder"

Providing a Major Leap Forward in Video Encoding Technology

By Pierre Larbier and Jim Hunt

Introduction

Driven by ever-escalating consumer demands for rich video content, as well as the proliferation of streaming connectivity methods and user-side display devices, the broadcast and broadband industries require advanced video encoding solutions that can deliver an optimal blend of processing speed, video quality, multi-screen capabilities and resource efficiency.

At the heart of every video processing system is the encoding "library" and programming API that provides the foundation for all of the higher level functionality. While this piece of software is not directly accessed by the end user, it represents a critical technology choice that will make or break the overall performance for the system as a whole. During the video encoding step, these core software processes can consume 90 percent or more of the total CPU power. Therefore, the efficiency and robustness of the encoding library directly determines the ultimate performance and effectiveness of the entire application.

In addition to providing the speed needed to support large-scale, continuous video encoding operations, today's video encoding systems must also offer a high degree of flexibility for outputting different levels of video quality, frame-rates, bit-rates and screen resolutions. With video content now being streamed and/or downloaded to a widening range of consumer devices, broadcasters and service providers need to be able to optimize the end-users' viewing experience by tailoring the video content for specific screen sizes and available bandwidths.

This white paper provides an overview of EAVC4, the major new 4th-generation release in ATEME's proven family of industry-leading H.264 video encoding technologies. EAVC4 represents the ultimate MPEG-4 AVC / H.264 Encoder, and is positioned to fulfill all of the industry's ongoing requirements for advanced standards-based MPEG-4 encoding over the coming years, while also laying the technological foundation for the introduction of next-generation High Efficiency Video Coding (HEVC).

In addition to providing major improvements in speed and video quality, EAVC4 also uses an innovative new approach to processing video for multi-screen output. This patented technique, called "Multi-Screen by Design," makes use of intelligent parallelization to perform common processes only one time for all bit-rates and screen formats, yielding dramatic speed improvements and optimizing video fidelity for multiple outputs from a single content stream.

Leveraging ATEME's Proven Foundation of EAVC Codecs

Building on 20 years of technology leadership in creating video encoding solutions— highlighted by three major releases of H.264 encoders over the past decade—ATEME is widely renowned by video-processing system builders as well as professional video content creators and service providers.

EAVC4 benefits from both a strong lineage of previous ATEME EAVC encoders, and significant innovations that are unique to this newest release. Highlights of the previous three EAVC generations are as follows:

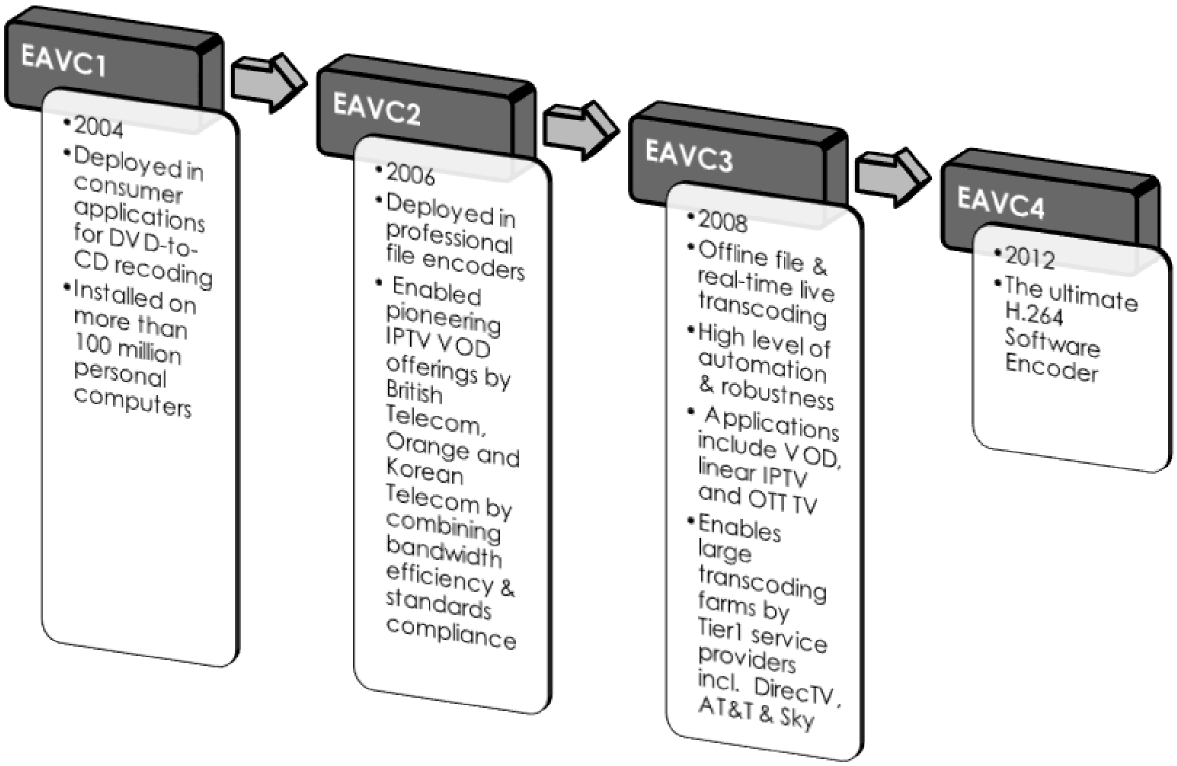

EAVC4: Meeting Advanced Performance Needs Now…
With an Eye Toward the Future…

In response to the broadcast and video management industries' dynamically evolving needs, this major release of EAVC4 encoding software addresses today's most pressing requirements for improvements to both performance and flexibility.

Few companies can build and release a complete MPEG-4 encoder from scratch, and ATEME's introduction of EAVC4 comes at a critical time for the industry, as most technology providers have already shifted their entire development focus and resources toward meeting the new HEVC standard.

As a long-time leader in the video encoding space, ATEME also is well underway with plans for HEVC product releases. However, it is clear that the need for higher performance, encoding flexibility and improved multi-screen capabilities is not something that users can simply defer until the full testing and adoption of future HEVC solutions.

ATEME's in-depth experience as a developer of core video encoding technologies and designer of complete video processing systems provides a unique perspective on the critical interrelationships between low-level and high-level architectural requirements. The decision to invest in the development and release of one final improvement to MPEG-4 AVC encoding capabilities was driven by real world market requirements that cannot wait for HEVC.

It is also important to keep in mind that the usage of MPEG-4 is continuing to grow and expand, even as much of the future-focused discussion turns to HEVC. Depending on the specific market segments and industry requirements, some video processing operations may continue to be addressed by the well-established MPEG-4 standard while other segments move to HEVC with a varied adoption rates.

By providing this new release of EAVC4 in Q2 2012, ATEME has assured users with continued availability of the enhanced MPEG-4 AVC capabilities needed to support the full range of user requirements throughout at least the next three to four years. EAVC4 will meet emerging requirements, support continuing MPEG growth demands, and bridge the gap while HEVC matures.

The following sections provide a more detailed discussion of the three major focus areas that have driven the creation of EAVC4:

- Delivering the highest quality video
- Providing the best quality and encoding performance at any specified speed
- Optimizing for video processing for multi-screen output

Delivering the Highest Video Quality

With the entire EAVC4 encoder redesigned from top to bottom, this newest release significantly raises the level of available video quality in a number of different aspects.
Some of the key video quality improvement areas include:

- Enhanced objective and subjective picture quality
    - Up to 20 percent bandwidth gains for progressive and interlaced content based on objective measures, and even higher gains based on visual perception
- Reducing bandwidth for use cases with high and/or emerging market demand
    - iPad high-resolution format (1080p)
    - High quality on smartphones, even at 64 kbps
    - Highly efficient software-based CBR and VBR bit-rate regulation, delivering constant quality over time and designed for multi-threading operation
- Expanding the range of available use cases
    - 4K and 8K image resolution, 4:2:2 and 4:4:4 chroma, up to 14-bit coding precision, AVC-I intra only coding
    - Mezzanine encoding formats in addition to multi-screen delivery formats
    - Laying an early foundation for high-resolution use case scenarios that will ultimately be required in HEVC

Providing the Best Video Quality for Any Specific Encoding Speed

EAVC4 offers a major leap forward in performance, with dramatic speed improvements for all video quality (VQ) settings. This provides users with the capability to either encode faster at their current VQ level or to significantly increase VQ while maintaining the current encoding speed.

The all-new EAVC4 video encoding architecture draws its efficiency from:

- Multi-processor and multi-core 64-bit architecture
- Multi-threading from slice level up through macroblock level
- Algorithmic optimization at multiple levels: video sequence, GOP, frame and macroblock
- Assembly code optimization to leverage the newest Intel SandyBridge capabilities and its specific instruction sets
- Software architecture optimization (patented approach, free of any IP challenges)

The combination of these technology enhancements yields major performance improvements for customers in both the file transcoding and live transcoding environments.

- File transcoding workflows, such as VOD processing, can benefit from up to a three times speed improvement, enabling jobs to be completed in one third of the previous time
- Live transcoding workflows, such as IPTV, can benefit by delivering three times as many linear channels per processor blade

Maximizing Performance Tradeoffs and the Range of Choices

Most software encoders can be configured to provide various tradeoffs between video quality and speed. EAVC4 is designed to offer maximum flexibility, which allows users to tailor both the performance and speed to better achieve their specific application requirements.

EAVC4 offers eight different VQ/speed encoding modes with a 50x difference between the slowest and fastest modes, as compared to a 10x range of difference in EAVC3.

The fastest mode in EAVC4 is designed not only to be significantly faster that EAVC3's fastest mode, but to provide video quality that could be used in broadcast applications.

By taking maximum advantage of both multi-thread software architectural improvements and the newest multi-core hardware capabilities, EAVC4 assures the most efficient usage of available hardware resources. This enables users to more effectively tailor their capital equipment deployments (e.g. number of blades, CPUs, etc.) in order to achieve optimal tradeoff between CapEx investments and the targeted performance level.

This more efficient use of hardware resources gives users the option to achieve the same amount of video encoding throughput with a lower equipment investment or to achieve significantly higher throughput from existing equipment.

Performance Comparisons for EAVC4 vs. the Alternatives

As previously mentioned, most software encoders can be configured to make tradeoffs between the encoding speed and video quality, therefore it can be quite challenging to accurately evaluate the actual performance vs. perceived performance differences.

In order to provide an empirical basis for comparison, the following sections focus on measuring the raw performance of the codec at specified quality and bit-rate levels, without any adaptive bit-rate optimizations. This provides a better understanding of the relative coding efficiency and speed for EAVC4 as compared to previous ATEME encoders, as well as other encoding solutions currently on the market.

For these comparison tests, video quality has been measured using two well-established industry standard metrics:

- Peak-signal-to-noise-ratio (PSNR) vs. bit-rate: to determine the level of errors introduced during the encoding process
- Bjøntegaard-delta (BD) vs. bit-rate: to compare the relative percentage savings in bit-rate for each encoder using normalized PSNR levels (ref: VCEG-M33, April 2001)

Evaluation of Coding Efficiency

In the following tests, all of the encoders are set to their maximum video quality settings, which also equates with the slowest encoding speeds. The first set of graphs (see Figures 1 and 2) show PSNR measurements of the various encoders on two common test sequences encoded at constant bitrates ranging from 2.5Mbps to 20Mbps.

As can be seen, EAVC4 consistently provides better quality PSNR results across all bit-rates when compared with the other commercial or open-source encoders.

Visual evaluation confirms these objective results as, for instance, the large tree behind the runners in the *CrowdRun* sequence encoded with EAVC4 does not exhibit intra-frame flicker while this annoying defect is visible with all other encoders.

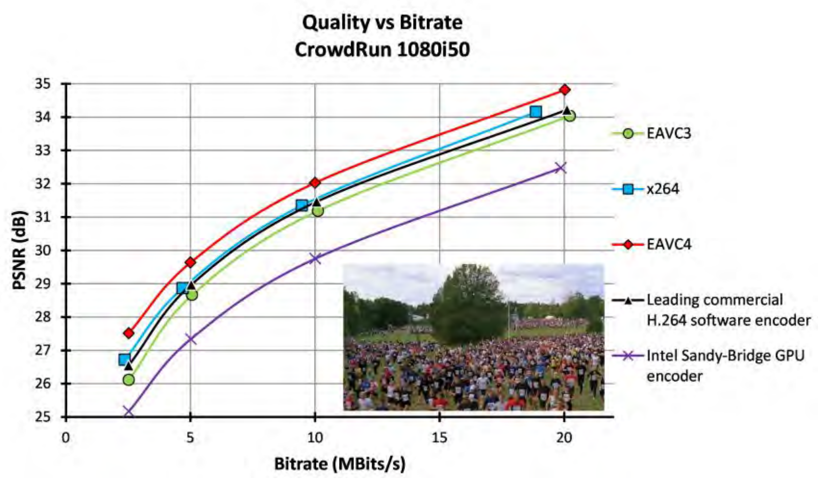

Figure 1 – CrowdRun Rate-Distorsion Measurement

Figure 2 – ParkScene Rate-Distorsion Measurement

The next series of graphs show results of the Bjøntegaard-delta (BD) vs. bit-rate analysis, comparing the same set of commercial and open-source encoders with regard to their percentage of bit-rate gain when normalized at the same PSNR level as compared with the baseline reference encoder (GPU).

This analysis shows that EAVC4 provides an average bit-rate gain that is 20 percent greater than EAVC3 and 10 percent greater than x264 across all of the tested ranges (see Figures 3 and 4).

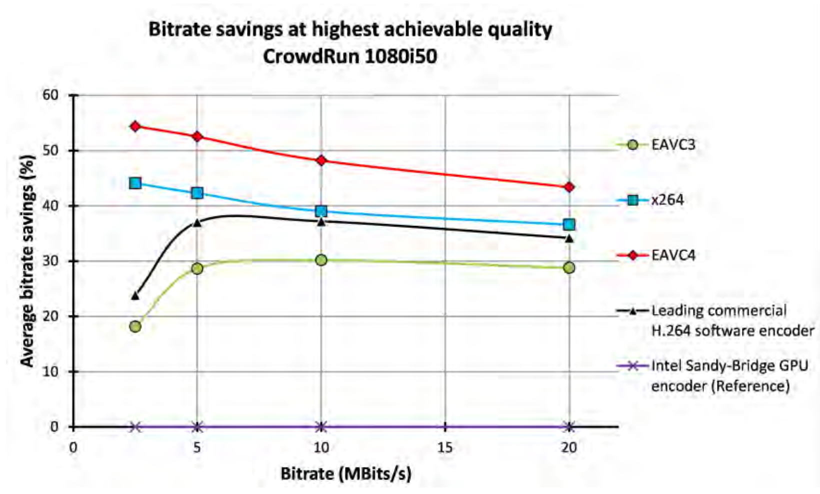

Figure 3 – CrowdRun Bjøntegaard-delta Analysis

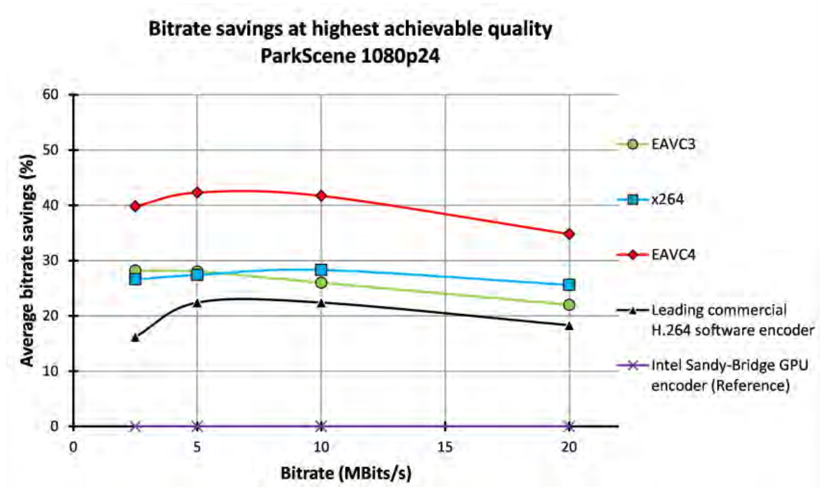

Figure 4 – ParkScene Bjøntegaard-delta Analysis

Evaluation of Quality vs. Speed Tradeoffs

In summary, these test results show that EAVC4 consistently provides the best performance results on both industry standard benchmarks across the full range of bit-rate use cases.

The following analysis provides an aggregated view of the overall advantages that EAVC4 offers with regard to optimizing tradeoffs between video quality and encoding speed. The graphs in Figures 5 and 6 show a summary figure-of-merit for the fastest encoders from the previous section (EAVC4 and x264) along with ATEME's previous generation EAVC3 encoder as a baseline reference for the relative measurements.

For these tests, all three encoders have been normalized to operate at a fixed bit-rate and the graphs show each encoder's relative positioning with regard to both the encoding speed and BD bit-rate gain. For this summary analysis, the optimal positioning is toward the top-right portion of the graph. As can be seen, EAVC4 offers the best figure-of-merit for combining quality and speed across all of the use cases.

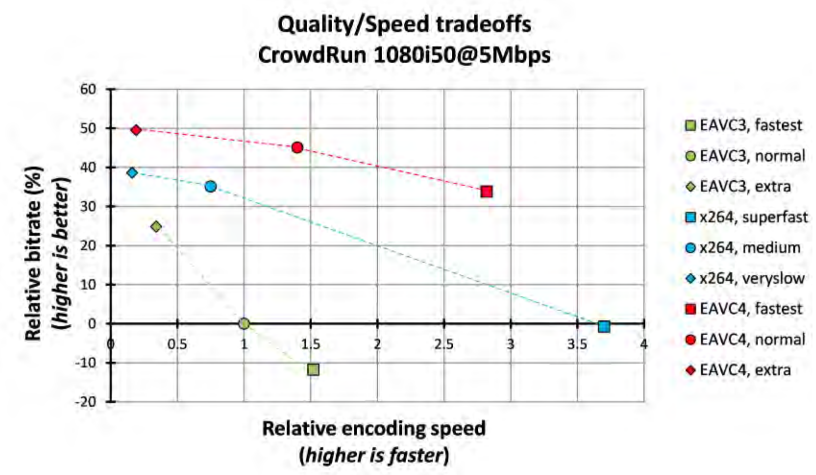

Figure 5 – CrowdRun Relative Bitrate vs. Relative Speed

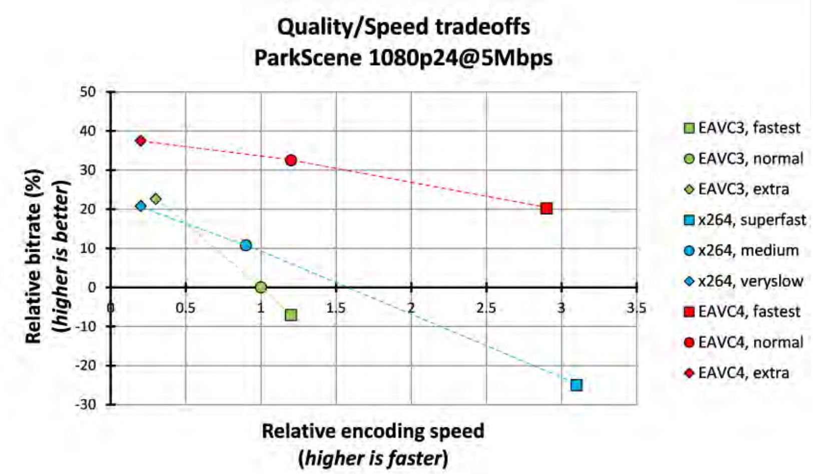

Figure 6 – ParkScene Relative Bitrate vs. Relative Speed

Optimizing for “Multi-Screen by Design”

One of the most significant breakthroughs in EAVC4 is the new patented approach to intelligent parallelization of the video encoding pipeline called “Multi-Screen by Design.” The EAVC4 encoding core is specially designed to provide high efficiency output of multiple formats as required for Adaptive Bit Rate delivery.

As illustrated in Figure 7, EAVC4 uses an optimized architecture that performs all of the common processes only once for multiple formats and bit-rates that are to be output from a single content stream.

The key steps shown in the figure are as follows:

- DEC: decoding process
- PP: pre-processing to clean up the input video before encoding (noise reduction, etc.)
- DS: downscale process to change video resolution
- ENC1: H.264 encoding portion that is not dependent on the output bit-rate
- ENC2: H.264 encoding portion that is dependent on the output bit-rate

As can be seen, in the non-optimized design used by all existing encoders, the input video is decoded

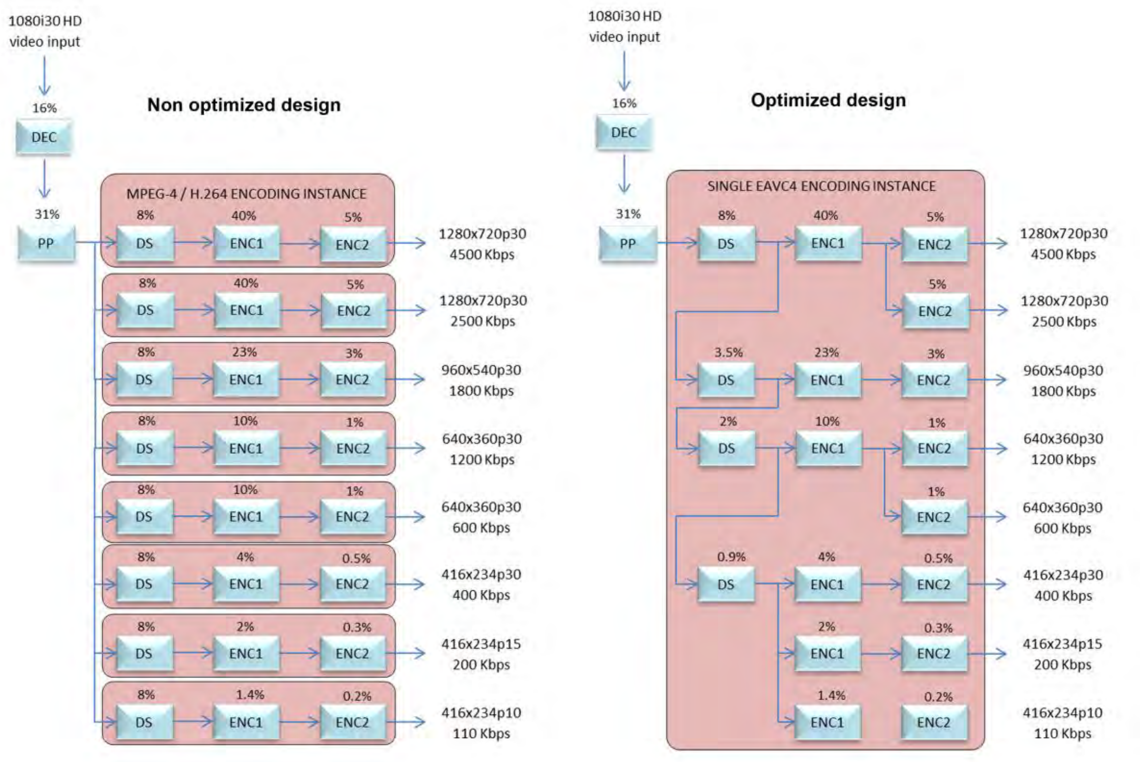

Figure 7 – Video Encoding Pipeline Optimization with “Multi-screen by Design”

and cleaned up once but then the content stream is repeatedly downscaled and encoded separately for each output format. In contrast, the EAVC4 optimized design performs the input decoding, clean-up, downscaling and the non-output-dependent encoding portion (ENC1) only once for the entire content stream.

This innovative approach yields very significant speed improvements, especially for multiple instances of the same output resolution and frame-rate but using different bit-rates. In these cases, all of the processing steps for that common screen resolution and frame rate, except for the last one (ENC2) are performed only once for all required output bit-rates. Similarly, for instances that share a common resolution but have different frame-rates and bit-rates, all steps except for the last two (ENC1 and ENC2) are performed only once.

Depending on the mix of output resolutions, frame-rates and bit-rates required by any specific customer's video processing operations, the aggregate benefits from using EAVC4 can range from quite significant to very compelling.

For example, content streams that are intended for multiple devices that share a 1280x720 resolution at 30fps frame-rate but need different output bit-rates (such as AppleTV, iPad, and iPhone 4), the EAVC4 approach can cut the processing load by 90 percent or more for every bit-rate instance after the first one.

The charts in Figure 8 illustrate a representative sampling of the total processing load for a variety of multi-screen output scenarios, which compare the non-optimized traditional encoder approach with the EAVC4 "Multi-Screen by Design" optimized approach. (In all instances, the input is HD 1080i30 video and output best practices are as defined in Apple use case referenced at https://developer.apple.com/library/ios/#technotes/tn2224/_index.html .)

Conventional Transcoders - - Non-Optimized Design

	DEC	PP	DS	ENC1	ENC2	TOTAL
1280x720p30	16%	31%	8%	40%	5%	**100%**
1280x720p30			8%	40%	5%	**53%**
960x540p30			8%	23%	3%	**33%**
640x360p30			8%	10%	1%	**19%**
640x360p30			8%	10%	1%	**19%**
416x234p30			8%	4%	1%	**13%**
416x234p15			8%	2%	0,3%	**10%**
416x234p10			8%	1,4%	0,2%	**9%**
TOTAL			***63%***	***131%***	***16%***	***257%***

EAVC4 - - Optimized Design

	DEC	PP	DS	ENC1	ENC2	TOTAL
1280x720p30	16%	31%	8%	40%	5%	**100%**
1280x720p30					5%	**5%**
960x540p30			3,5%	23%	3%	**29%**
640x360p30			2%	10%	1%	**13%**
640x360p30					1%	**1%**
416x234p30			0,9%	4%	0,5%	**6%**
416x234p15				2%	0,3%	**2%**
416x234p10				1,4%	0,2%	**2%**
TOTAL			***14%***	***81%***	***16%***	**158%**

Figure 8 - Comparison of Total Processing Loads for Non-Optimized vs. Optimized Designs

To compare performance, a reference load of 100 percent has been assigned to the most demanding output format (1280x720p30) and a comparative percentage is shown for each of the other formats.

For instance, in the non-optimized design a second instance of 1280x720p30 using a different bit-rate would require 53 percent as much load as the first instance. By comparison, by using the EAVC4's optimized design, the second instance of 1280x720p30 at a different bit-rate would only require 5 percent as much load as the first instance!

Looking further down the list of samples, one can see that in the non-optimized conventional design, each 640x360p30 output requires 19 percent of the resources needed for the baseline 1280x720p30 output (regardless of how many bit-rate variations are output at 640x360p30 they all consume the same 19percent). In contrast, with the EAVC4 optimized design, the first instance of 640x360p30 only needs 13 percent and subsequent instances of 640x360p30 at different bit-rates can be encoded using just 1 percent of the baseline resources.

While every video processing operation will have its own unique mix of multi-screen output requirements, it is clear that the optimized design can readily provide overall savings whenever input video content streams need to be output for a range of different screen resolutions, frame-rates and bit-rates.

Conclusion

As the last major release that the market is likely to see for this compression standard, EAVC4 represents the ultimate H.264 codec; not only in terms of its timing but also in terms of the large leap forward in speed, video quality, flexibility and multi-screen output capabilities.

ATEME has carefully designed EAVC4 from the ground up to provide the additional performance and functionality needed to meet the escalating requirements of the broadcast, broadband and content creation segments over the next three to four years; this will ensure unbroken, leading-edge support for the huge number of companies and users who depend on MPEG-4 encoding both now and in the future. At the same time, the unique innovations pioneered in EAVC4, such as Multi-Screen by Design, signal a major leap forward in encoder architecture that will also benefit the next generation of HEVC encoders from ATEME.

EAVC4: *Beyond the industry's requirements for high fidelity MPEG-4 encoding*

About the authors

Pierre Larbier is ATEME's Chief Technology Officer and a renowned expert in the field of algorithmic design and MPEG / H.264 video compression. He has specifically dealt with video encoder architecture and designs for a wide range of platforms and applications: from low-power single DSP to ultra-high performance FPGA and multi-core software.

Pierre can be contacted at p.larbier@ateme.com.

For the past 20 years, **Jim Hunt** has written extensively about a wide range of software and hardware technologies that are revolutionizing the ways that people work, live and interact with information. Among the key focus areas of Jim's research and writing projects are advanced content creation and optimization methodologies as well as communications and delivery infrastructures.

Jim can be contacted at jim@techstoryteller.com.