Software vs Hardware Compression: Striking the Perfect Balance

Abstract

The compression of high-quality video has always been a compute-intensive operation requiring dedicated hardware. But the steady progression in computational power of commercial-off-the-shelf (COTS) equipment in the past few years has enabled the gradual transition of some video processing operations to software-based environments. A software-based approach can offer significant flexibility and agility improvements over purpose-built hardware, but at a potential cost to video quality. As the computational power of generic computing resources continues to grow, media companies must regularly evaluate the relative flexibility, cost, and quality attributes of hardware- and software-based compression engines. Additional factors influencing compression strategies include shifting video consumption patterns that require media companies to expand productivity and to improve operational efficiencies to meet growing consumer demand for anyplace, anytime, and any device access to video content.

This document provides video processing professionals with the information needed to construct a comprehensive strategy that delivers the optimal balance of flexibility, quality, and cost to meet not only today's but also tomorrow's compression requirements.

Introduction

Compression technology has long played an instrumental role in the distribution of video content. Consistent improvements in the ability to compress a series of images into a reduced payload allow media companies to steadily improve bandwidth efficiency while maintaining the required quality. Advances in compression technology, especially in the digital realm, have enabled media companies to improve productivity and competitive positioning by increasing the number of services within existing bandwidth, improving the quality of those services, or a combination of both.

All modern video delivery systems rely on some form of compression, whether it's a traditional linear TV service to a set-top box, or a streaming service that terminates on a gaming console, smartphone, or tablet device. Though several compression types have preceded it, MPEG is one of the most widely used compression standards. This compression type has predominantly required hardware encoding, but high-quality software-based solutions have become more prevalent due to the steady increase of computational capabilities in generic computing platforms.

These incremental gains can, in part, be attributed to Moore's Law, a long-standing axiom that projects that the number of transistors in dense integrated circuits will double roughly every two years, which loosely contributes to

chip performance doubling roughly every 18 months. Moore's Law and the associated performance increases in off-the-shelf computing platforms play a pivotal role in determining the cut-over point between hardware-assisted and software-based compression.

Another influential factor in determining when to utilize hardware-assisted or software-based compression is recent shifts in the way video content is distributed and consumed. Media companies are now confronted by the need to produce more content than ever before and to distribute it to an ever-growing diversity of receiving devices. These demands require media companies to support a growing assortment of protocols, formats, and resolutions, which can require frequent upgrades and code revisions. Completely replacing large amounts of existing equipment with each upgrade is costly and inefficient. As an alternative, software-based and hybrid compression approaches offer flexibility and versatility that enable media companies to leverage existing content and equipment for future technology upgrades.

Despite the flexibility and cost-efficiency benefits of moving compression operations to software, media companies must consider an assortment of variables. Any shift from hardware-based encoding to software-based encoding will depend upon a number of factors, including cost, existing infrastructure, distribution requirements and technical factors, such as video quality and latency.

The remainder of this document will supply media companies with the information needed to construct an optimized compression strategy. It will provide insight into multiple issues confronting media professionals, including the potential to move all compression operations to a software-based environment, and how best to match specific applications and operations to compression resources

Compression Options

Hardware Compression

Dedicated hardware encoding has, in general, been regarded as the better option for high video quality compression platforms due to the fact that these encoders use processors that are specifically designed for the job at hand — they only process the algorithms required.

Hardware compression uses ASICs (Application Specific Integrated Circuits) for encoding processing and FPGAs (Field Programmable Gate Arrays) to support any additional Transport Stream (TS) processing or video analysis. But the higher processing performance and power efficiency comes at a cost. A significant disadvantage of using dedicated ASICs is the lack of flexibility. ASICs cannot be re-programmed. Every time a new codec is introduced, new hardware is needed. An FPGA-based encoder allows for more flexible programming, but adds additional cost and consumes more power than an ASIC-only implementation.

Additional benefits of hardware encoding include lower latency. In hardware compression, latency depends on the encoding type and profile. Figure 1 lists the contribution encoding latencies of several hardware encoding types. Distribution encoding latencies are typically higher, but deliver better bitrate efficiency.

Hardware Encoding Type	**Latency (end-to-end)**
MPEG-2	**250ms, 750ms (Profile dependent)**
H.264	**250ms, 300ms, 750ms (Profile dependent)**
JPEG2000	**40ms, sub-frame (licensable)**

Figure 1: Hardware Encoding Latency

Hardware encoders are often used in more static, point-to-point systems that are reconfigured rarely or not at all during their operation. Advances in compression technologies, such as H.264 and HEVC (High-Efficiency Video Coding), have enabled service providers to deliver more and higher-quality services over their existing infrastructures, such as terrestrial, satellite or cable. (H.264 delivers the same quality as MPEG-2 at approximately half the bandwidth. HEVC is roughly twice as efficient as H.264).

Hardware-based encoding is often initially more expensive than a software-based approach. However, the lower running costs (including maintenance and power consumption) of hardware encoding make it comparable to software encoding over time.

Software Compression

The main difference between hardware and software encoding is that software encoding uses standard, widely available, COTS server platforms that can simultaneously process a multitude of computations. Software-based encoders are regarded as not delivering the same quality as in hardware encoders. But those deficiencies are disappearing. As processors become more and more efficient, powerful, and cost effective, media professionals are able to utilize them in denser and more compact solutions. Video quality is now offered at all ranges as codec engineers focus on software implementations.

Running compression operations on generic server platforms offers significant flexibility and versatility. It is possible, for example, to change and grow the system while it's in operation. It's also possible to easily redirect computational power where it is most needed. This allows software encoding to be more flexible and responsive for commissioning or on live in-use systems. Software-based encoding is also well tailored to cloud-based and OTT applications, which require the transcoding of multiple ABR (Adaptive Bit Rate) streams of the same content for differing receiving devices.

Latency in software encoding can depend greatly upon the type of codec (MPEG-2, H.264 and HEVC). Independent of codec latency, there is also system latency due to the applications typically associated with software encoding, and is largely determined by the streaming protocol used. For H.264/AVC compressed video over RTMP (Real-Time Messaging Protocol), for example, latency can be in the region of 2 seconds from the input of the encoder to the user device. Latency associated with HLS/HSS can be in the 30 second range end-to-end, depending on settings. These figures will also depend upon the chunk duration, segments, and network before playback. As with hardware encoding, the way that the encoder is set up can affect latency, as well as video quality.

Hybrid Software and Hardware Compression

Combining software and hardware compression solutions in a hybrid approach often provides media professionals with the ability to customize a compression strategy that most efficiently balances video quality, flexibility and costs. A hybrid approach, for example, may rely on ASICs-based hardware to manage the complexities of encoding and software to manage other requirements, including TS (transport stream) processing or video analysis, removing the need for FPGAs.

A hybrid solution often gives media companies the ability to provide denser solutions while at the same time, simplifying the process of adding new features — all while retaining video quality due to the advanced computational power of the hardware used. The hybrid approach allows for fast application of the right tools for a given requirement at any time.

Evolving Landscape

This section will examine factors that are influencing the direction of the compression market.

Moore's Law

As previously noted, Moore's Law posits that the number of transistors in a dense integrated circuit will approximately double every two years. Though there is speculation that the computational capabilities of semiconductors will eventual reach a saturation point, cumulative advances in generic computing continue to improve the power and efficiency of software-based compression.

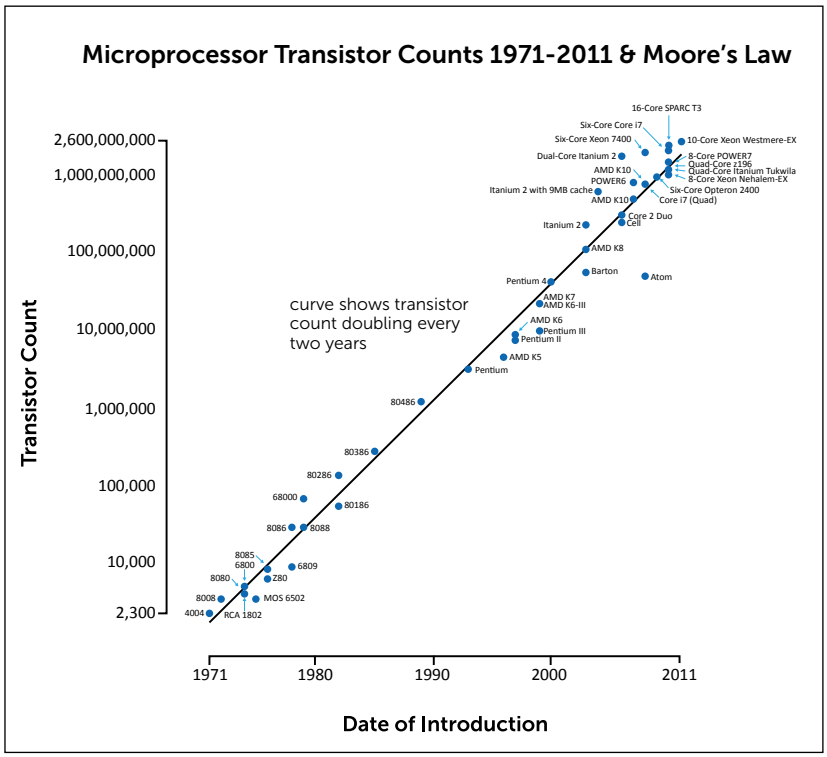

Figure 2: Moore's Law Timeline

COTS Equipment

The cumulative impact of Moore's Law has led to the availability of COTS platforms that provide media companies with flexible and versatile alternatives to purpose-built hardware. The agility and versatility benefits of moving operations to virtualized environment, as well as the need to future-proof infrastructures to seamlessly incorporate the latest technology advances, are prompting many of the largest media companies in the world to migrate operations, including production, playout and distribution, to generic computing and networking platforms.

Shifting Consumption Patterns

As video consumption habits evolve, so too must the way media companies distribute their content. Steady advances in Internet-connected portable devices allow consumers to watch content on the move and whenever they want, rather than having to view content in a static environment at a fixed time. Broadcasters and content creators and distributors will require platforms offering increased agility to keep pace with the rapidly evolving media and entertainment industry.

Transition to IP and Virtualized Environment

While moving to software-based systems on standardized platforms allows for a more flexible and cost-effective approach for the distribution of content, the cost and operational disruption of a wholesale move to an all-IP environment is not feasible for most media companies. The move to an IP-based infrastructure, with rare exception, will be gradual.

Similarly, the transition from purpose-built hardware to software-based, or virtualized, environments will not be done overnight. A survey conducted by Imagine Communications at the end of 2015 found that only about 15% of media companies surveyed had converted more than half of their operations to software-only environments.

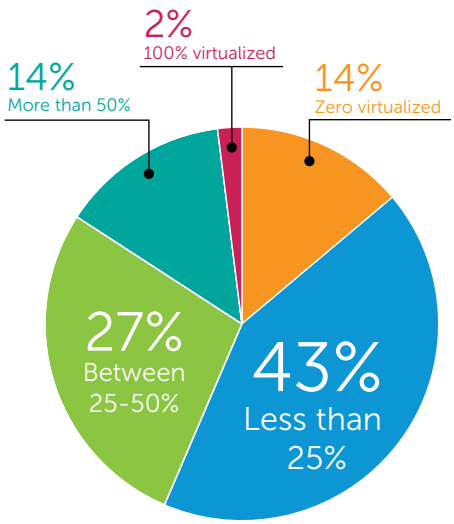

Figure 3: Hardware-to-Software Transition
Nearly 15% of the more-than-700 media professionals surveyed said their companies had transitioned more than 50% of all media operations to a software-based, virtualized environment.
Source: 2016 Focus Forward Technology Trends Report

The Impact of Compression Enhancements

As Ultra-High Definition (UHD) production becomes more prevalent, the need to provide a compression technique that delivers the best quality available in an efficient manner increases. A proposed successor to H.264, HEVC (also known as H.265), which is based on improved algorithms over H.264 and MPEG-2, delivers substantially improved video quality at the same bit rate as H.264. Currently, most HEVC codecs are software based, which promotes flexibility but requires huge amounts of computational power. A hybrid approach to HEVC compression would enable some media companies to preserve flexibility and quality while lowering computational requirements. As codec implementations mature, latency associated with HEVC is expected to decrease.

Software vs Hardware: Putting it All Together

Achieving a compression strategy that strikes the perfect balance between hardware and software depends on a number of variables, which will be unique to each media company. Broadcasters, content distributors and others will need to prioritize the importance of flexibility, density, video quality, and price-performance on a service-by-service basis. In addition, media companies will need to closely monitor additional factors and market drivers, including codec maturity, recent advances in the computational power of generic platforms and shifts in video consumption practices, all of which are constantly evolving. The latency associated with software-based compression, for example, is likely to improve over time with advances in generic compute equipment and as codecs become more efficient.

Figure 4 provides a matrix for comparing the benefits and capabilities of software- and hardware-based compression strategies, as well as examples of applications where each is likely to be used today.

	Hardware	Software
Latency	Superior for low latency contribution requirements	Optimized for distribution applications, where latency is less of a concern
Cost	Initially more expensive, reducing over time due to lower running costs	Lower cost for high-density solutions at balanced video quality, also lower costs for COTS hardware
Flexibility	Rigid in that new functionality and capabilities require new hardware	Extensive, easy to add new functionality and codecs
Quality	Superior for applications that require high video quality	Better for balanced video quality, high-density applications
Application sweet spot	Contribution, Primary Distribution, Secondary Distribution	Distribution for both OTT/ABR and linear

Figure 4: Compression Comparison

Summary

In today's dynamic market, a well-tuned compression strategy is a competitive imperative. The ability to flexibly and cost-efficiently repurpose existing content to OTT and ABR platforms increases revenue and market share for content owners and distributors. Achieving a balanced compression strategy depends on a number of factors, including technical and business requirements, as well as all-important budgetary considerations. Whether to provide high-quality linear service over cable and satellite or to provide hundreds of thousands of services to multiple devices, a comprehensive and optimized compression strategy has never been more important.

The Imagine Communications Compression and Distribution Portfolio

For high-quality and efficient linear hardware compression, Selenio™ MCP is a comprehensive,modular solution that provides reliable encoding and processing in a converged baseband/ broadband environment. Available in 3RU (14 application slots) and 1RU (3 application slots) form factors, Selenio™ MCP works with a broad selection of processing modules to support a simplified infrastructure from baseband to IP.

The SelenioFlex™ Live platform, powered by the extensible Zenium™ software framework, provides superior software-only encoding. It is available in a 1RU appliance, with single-, dual-, or quad-channel SDI input models, as well as a software-only offering for virtualized and cloud deployments. SelenioFlex™ Live supports OTT, web, and mobile live streaming formats, with the ability to offer linear and ABR profiles for input SDI or IP transport stream sources.

SelenioNext™, available in a variety of form factors, is a high-density hybrid compression solution, which offers up to 10 times the density and power efficiency of competing solutions.

A uniquely agile, software-defined platform, Selenio™ One enables service providers to efficiently process and deliver current and future compressed streams using a single platform — while also ensuring the integrity and quality of their high-value video content. The initial Selenio™ One offering is a high video quality transcoder with supporting PCIe acceleration. Selenio™ One will support software-only and hybrid configurations in the future.

Each solution can provide best-in-class-scalability and redundancy profiles, along with user-friendly management and operations systems to provide efficient and streamlined systems economically priced and optimize for a broad spectrum of requirements.