Whitepaper:

Video encoding in live broadcast production

Sept 2024

Contents

Introduction

Need for compression

Live production requires many video, audio and other signals to be transported within and between production locations. Within a campus setting, the network is typically designed to carry all these signals uncompressed. However, once Wide Area Networks (WAN) are involved – such as those used for remote production and contribution off-campus connectivity – capacity can be much more constrained. Inevitably, there is a compromise between the number of signals required for the production, and the bandwidth available to transport all the signals. Using video and audio compression allows more signals to be packed onto the same bandwidth, but compression brings its own set of compromises, between compression ratio, quality loss and latency. This is particularly relevant for video, which typically represents the bulk of the bandwidth usage.

Compression in distribution vs live production

It is easy to forget that the compression (or encoding) requirements of live production are fundamentally different to those of distribution.

Latency

Latency is important for live production, as even small delays can lead to synchronization issues, and in a multi-camera production environment it is important to avoid discontinuous cuts. It is also important for coordination within the production crew that there is no significant delay between captured content and what remote operators can see and interact with. For example, in interview situations involving multiple locations it is essential to keep latency to a minimum to avoid impacting the storytelling experience.

Several studies show that the human eye can detect latency as low as 13ms, but generally speaking around 50ms latency is considered very good for production purposes. However, in remote production across larger distances it may not be possible to achieve 50ms latency, since network latency alone also contributes significantly to overall end-to-end latency. For instance, typical one-way latency within Europe is less than 15ms, within North America 30ms and trans-Atlantic approx. 70ms. For this reason, it may be necessary to adjust the workflow to operate with higher latency in some situations.

For the distribution chain, latency is not the main concern and generally speaking 2-3s or greater latency is considered acceptable. This means that the encoding process can focus on more optimal compression with longer latency. Simply speaking, bandwidth reduction and quality of experience are more important than latency in distribution. Also, distribution decoding is done on the consumer side, so it is important that this process is less complex, and that it can be supported by both hardware and software implementations.

Asymmetricity

Distribution typically involves the one-way transport of media from a single point to multiple locations (transmitters or even end-devices) through a path with potentially very limited bandwidth. In some cases, e.g. distribution over the Internet, the decoding will be performed by devices that may have limited computing capabilities such as smart phones.

As a result, the focus for compression in distribution is on achieving a high degree of compression and on minimizing the cost of decoding – both in terms of computation power and expense. This favors asymmetric compression methods, such as those of the MPEG family including H.264/AVC and H.265/HEVC, where a lot of effort is put into getting a high compression ratio and less effort is required to decode.

In contrast, live production requires very fast transport between locations, usually in both directions (e.g. for both feeds and return feeds). As a result, the biggest concern is latency: while a delay of multiple seconds is totally acceptable for distribution, delays must be measured in milliseconds or a few video frames in real-time production. At the same time, there is not such a compelling need for the compression methods to be asymmetric, as the ratio of encoding devices to decoding devices is not skewed towards decoding as it is in distribution. Furthermore, the availability and affordability of IP bandwidth is increasing rapidly meaning that very high levels of compression are not necessarily needed.

Varying bitrates

Video compression quality and efficiency depend on the bitrate mode. In video transport, different approaches are taken to bitrates, commonly referred to as Constant BitRate (CBR) and Variable BitRate (VBR).

Constant BitRate encoding, gives a consistent and predictable bitrate, which is good for streaming over limited or fixed bandwidth networks, but the drawback is that you may use more bandwidth than required for simple scenes or lose quality for complex scenes. However, its predictability makes it the preferred choice for live production.

Variable BitRate encoding, uses a variable bitrate depending on the scene complexity or motion. VBR encoding gives a more consistent quality of the output video but can cause unpredictable and fluctuating bitrates.

A separate approach, called Adaptive BitRate (ABR), is a way to vary bitrates dynamically depending on network conditions. It is widely used in OTT (over-the-top) internet delivery to ensure that video streams are not interrupted when the available bandwidth becomes squeezed.

Encoding in distribution

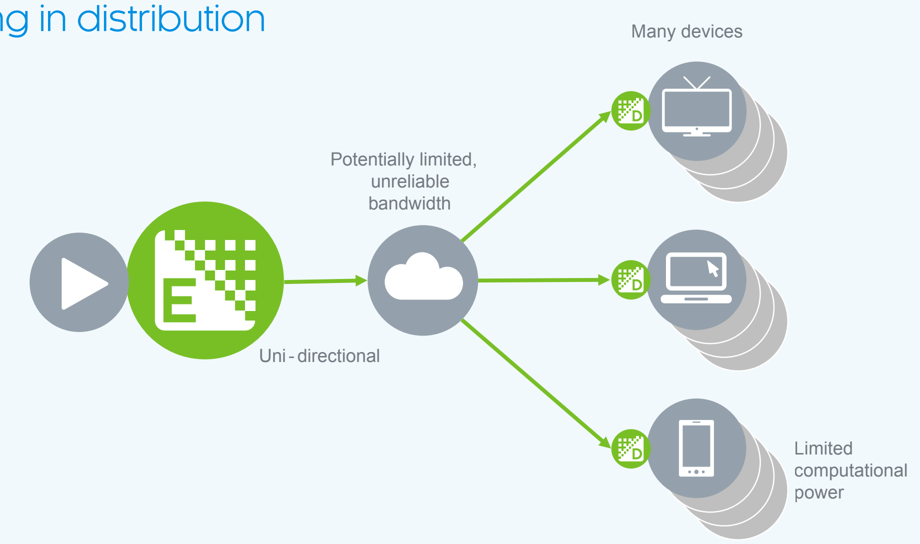

Encoding in live production

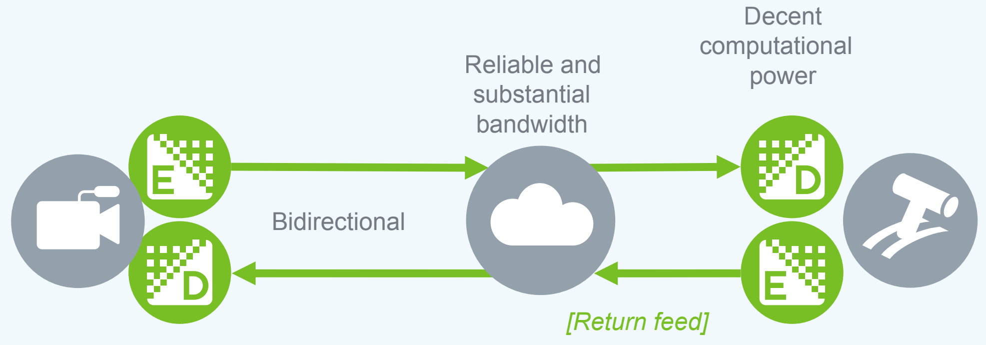

Handling higher resolutions

The move to higher resolutions, higher frame rates (HFR) and higher dynamic range (HDR) in high-end content production and distribution, substantially increases bandwidth and processing requirements.

The SDI infrastructure of most facilities was not designed to handle resolutions beyond 3G, which makes handling resolutions like UHD/4K somewhat complicated.

A commonly used approach is to split UHD/4K images into 4 identically sized quadrants, which allows the video to be transported as 4 synchronized 3G signals. Sometimes, compression techniques are used to squeeze those 4 signals onto a single 3G link. This obviously requires careful synchronization and monitoring but has been shown to work successfully.

As the infrastructure in the facilities and beyond moves to IP, the quad-split approach makes much less sense, and as a result 'full raster' or 'full frame' compression and transport with single UHD interfacing are becoming the norm.

It is worth noting that, while the industry is heading towards solutions that offer full-raster UHD/4K compression with SMPTE ST 2110 interfaces, many media transport products have limitations that mean they require higher resolutions to be handled through 4 streams of SMPTE ST 2110-20.

Quad-split

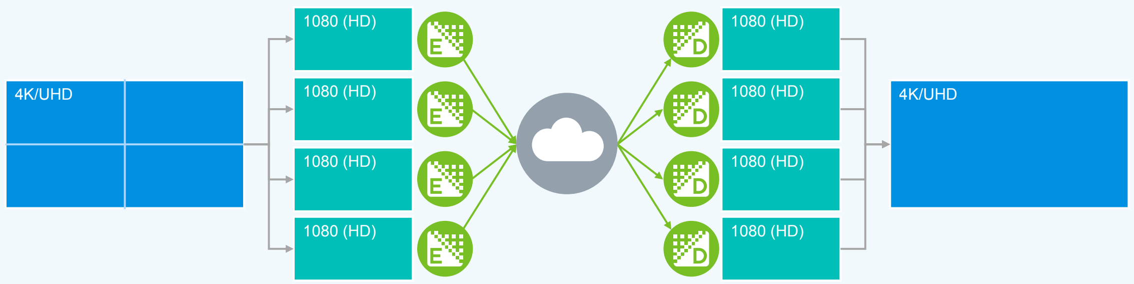

The SDI infrastructure of most facilities was not designed to handle resolutions beyond 3G.

Commonly used **compression**

When it comes to video compression, there are multiple codec (encode-decode) options.

In live production, the landscape is fairly stable, meaning the adoption of new codecs is relatively infrequent, primarily because of the need for compatibility with existing investments. However, innovation within those compression methods is gathering pace, especially in terms of reducing latency.

Comparing compression methods

Multiple factors can be used typically to compare these codecs, most commonly:

- **Compression ratio:** reducing bandwidth requirement is generally desirable.
- **Latency:** the lower, the better in live production, as delays complicate the interactions between those involved in production on site and remotely.
- **Quality:** while all codecs used in live production maintain a decent image quality over the course of a compress/decompress cycle, the main issue is how well quality is maintained over multiple cycles (also known as concatenation, or multi-generational compression or encoding).

In addition, from an implementation perspective, the following factors are sometimes considered:

- **Hardware and software support:** with live production shifting from dedicated encoding hardware to more generic platforms, like optimized software-defined platforms, COTS (commercial-off-the-shelf) hardware (i.e. IT servers) and even the cloud, the consistent support for those various hardware and software platforms is important.
- **Complexity**: This factor is largely related to the one above, i.e. if a codec is highly complex to implement, this reduces the potential for it to be supported more widely; complexity also affects the computing and electrical power needed to execute.
- **Cost:** the cost of a codec, including licensing, hardware and power has been a significant factor in the past, and continues to be in some cases.

As no codec achieves the best performance across the board, these factors need to be weighed against each other when selecting the most appropriate compression method.

Uncompressed (linear)

Today, broadcasters can transport uncompressed 3G/HD and UHD signals over IP networks using SMPTE ST 2110-20 or SMPTE ST 2022-6 (limited to 3G/HD) standards. SMPTE ST 2110-20 uses approximately 2.5 Gbps for 1080p and 10 Gbps for UHD.

Key advantages of uncompressed video include the lowest option for latency, pristine video quality and wide support – all of which are highly desirable for remote production. The major disadvantage is of course the bandwidth required to transport uncompressed video.

For that last reason alone, compression quickly becomes an attractive or even essential option – especially as the quantity of signals to be transported increases and higher resolutions are used.

Spatial (intra-frame) compression

Spatial compression methods involve encoding one frame (or each field in the interlaced world) at a time. These methods achieve a decent compression ratio with relatively low latency, which make them very suitable for real-time production.

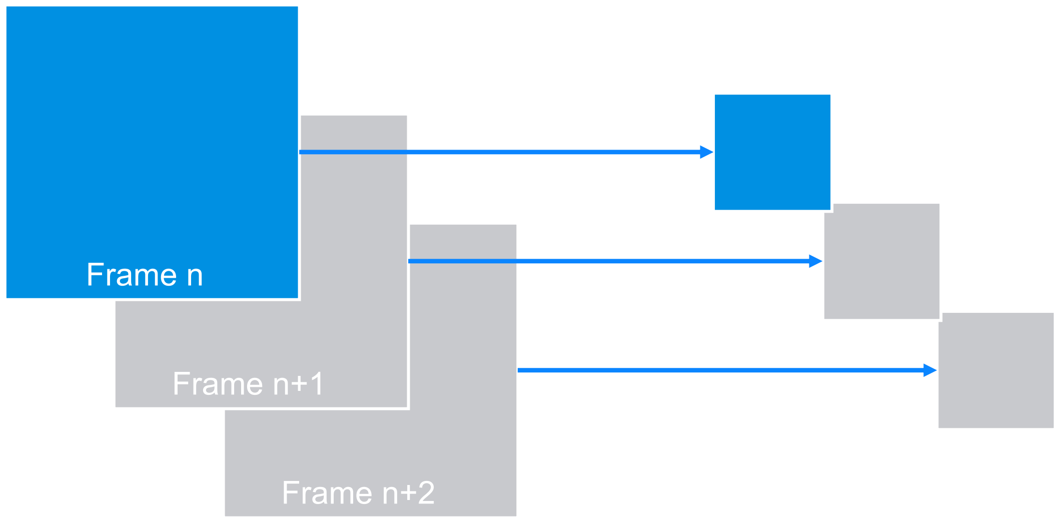

JPEG 2000 (J2K)

Developed between 1997 and 2000, JPEG 2000 (or J2K) quickly established itself as the compression of choice for low latency high-quality contribution and is still widely in use – though the more recently released JPEG XS technology is challenging that dominant position.

JPEG 2000 offers a visually lossless compression at around 1/8th to 1/10th of the uncompressed bandwidth. In practice, using JPEG 2000 compared to uncompressed means fitting 8-10 times more feeds on the same network. Nevion equipment achieves latencies of less than 60ms for HD 1080p50, i.e. fewer than 3 frames.

JPEG 2000 is used to transport UHD/4K video signals onto 4 x 3G links, using a quad-split technique.

The more recent versions of the VSF TR-01 JPEG 2000 interoperability document covers the use of ultra-low latency (ULL), which is achieved by compressing horizontal stripes of each image. Each stripe can be transmitted as soon as that stripe has been received and compressed. This technique, referred to as JPEG 2000 ULL achieves latency reduction to less than one video frame end-to-end (plus transit delay).

JPEG XS

The most recent lightweight compression standard to emerge is JPEG XS (ISO/IEC 21122). It is a wavelet-based intra-frame spatial compression (i.e. like JPEG 2000 ULL) that compresses horizontal stripes of each image which it transmits immediately.

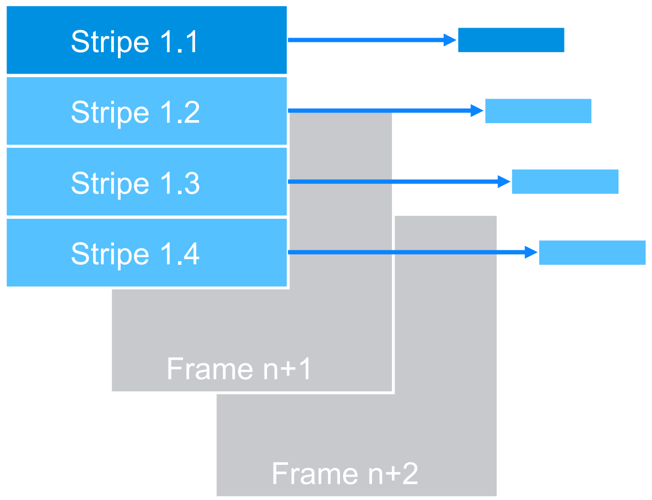

Spatial (sub-intra-frame) compression

This standard, which can also handle 4K and 8K full-raster, has been rapidly adopted in the industry because it combines a visually lossless compression rate close to that of JPEG 2000 (up to 1/8) with an ultra-low encode/decode latency (32 video lines or just a few hundred microseconds). That latency is negligible compared with that of other processes in the media network like transport (especially in a WAN environment) and buffering. It also maintains a pristine visually lossless output, both for one-time and multiple compression/ decompression cycles.

From an implementation point of view, JPEG XS is a lower complexity technology than its predecessors, which means it has low processing power requirements, making it very suitable for inclusion in cameras for example. It has been optimized for use on FPGA, CPU and GPU, which makes its use in COTS server technology highly viable. It is also a great cross-over technology for lightweight compression interfacing to and from cloud-based processing systems.

JPEG XS's quality, performance and versatility means that it has now made very substantial in-roads in production, not only for WAN applications (e.g. contribution or remote production) but also in a LAN environment (e.g. in the facilities). In effect, it has become the de facto lightweight compression – not only for UHD/4K and higher resolutions, but also for HD.

TICO, VC-2 (Dirac-Pro) and LLVC

When UHD/4K started to make serious in-roads into production and distribution, the industry began to seek fast and more efficient ways to squeeze the higher resolution signals onto existing SDI infrastructure.

The first codecs to emerge were the TICO, SMPTE VC-2 (aka Dirac-Pro) and Sony's LLVC (Low Latency Video Codec). All three are wavelet-based and provide 4:1 compression with minimal loss, and latency of less than 1 ms.

In many ways, these compression methods are largely surpassed by JPEG XS, which offers greater flexibility and higher compression ratios (see above). However, they can still play a useful role where the requirements are simply to transport UHD/4K over existing 3G-SDI infrastructure.

Temporal compression

Some encoding methods combine spatial (complete frame or field) compression with temporal compression (encoding just the change between reference or key frames/fields).

Such methods achieve much higher compression ratios than spatial-only compression but have historically done so at the cost of significantly higher latency. However, recent technical developments have reduced that latency substantially.

While temporal compression can achieve an excellent quality in a single compression/decompression cycle, that quality tends to fade after repeated cycles.

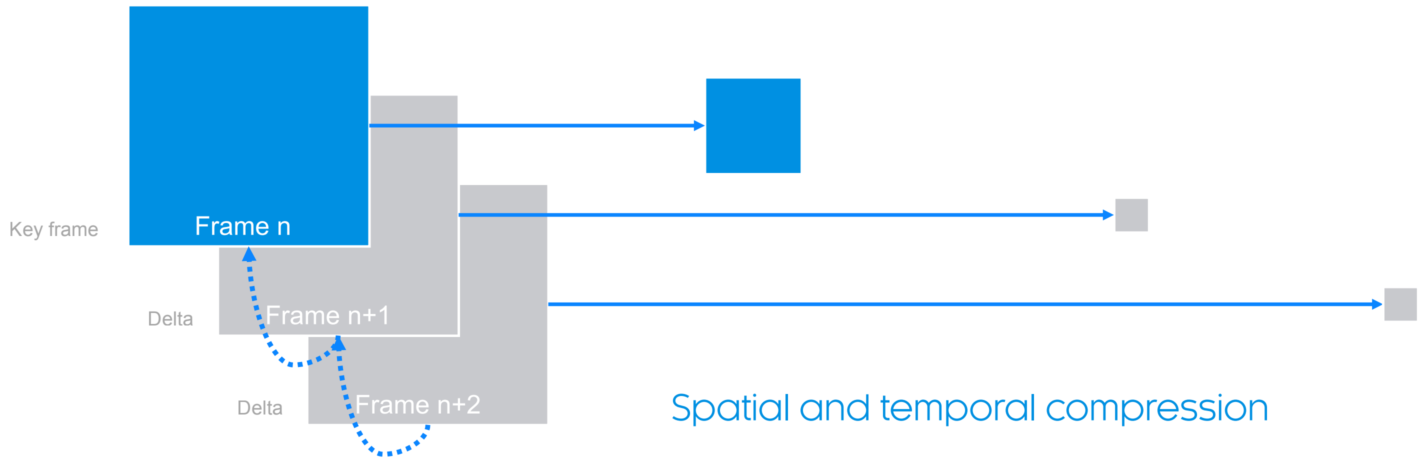

Spatial and temporal compression

I/P/B-frames

In temporal compression, reference frames (or more accurately key frames) which are encoded and can be decoded independently of other frames, are referred to as I-frames (Intra-coded picture).

The relative frames (or more accurately predicted frames) that are encoded relative to other frames, are known as P-frames (Predicted picture) if they are compared to the previous frames only, and B-frames (Bidirectional predicted picture) if they are compared to both previous and following frames.

H.264/AVC and AVC-I

In situations where bandwidth is at a premium, Advanced Video Coding (H.264/AVC) 4:2:2 10-bit encoding offers significant improvements in compression efficiency, while still providing pristine video quality.

H.264/AVC can squeeze video into about 1/100th of the bandwidth. Latencies of approximately 250ms for HD 1080i50 can be achieved, which is acceptable for most contribution needs.

For production H.264/AVC normally uses 4:2:2 10-bit encoding, while in distribution most content is encoded as 4:2:0 8-bit. The latter is supported by leading CPU vendors and therefore suitable for hardware accelerated decoding by consumer devices. Supporting 4:2:2 10-bit in software typically needs to be supported by GPU to achieve sufficient performance.

AVC-I is a variant of H.264/AVC where only I-frames are used, whereas the codec normally also uses P-frames and B-frames – effectively making a spatial compression. AVC-I has traditionally been used in some contribution applications as an alternative to spatial compression like JPEG 2000 – with similar characteristics, but less tolerance to multi-generational encoding.

H.265/HEVC

High Efficiency Video Coding (HEVC), also known as H.265 offers up to double the compression ratio of H.264/AVC, i.e. reducing the bandwidth by 50% compared to it. This has typically been achieved at a cost of a much higher complexity and latency.

Aside from high latency, one of the biggest obstacles to the widescale adoption of HEVC has been that of the licensing of the intellectual property within the standard, which has made it costly to implement. However, these obstacles have been much reduced in recent years, and as a result H.265/HEVC has been challenging the dominance of H.264/AVC.

Ultra-low latency (ULL) versions of the HEVC codec, based on slicing technology, have now begun to emerge – for example the technology developed by Sony, which achieves latencies as low as two frames. The cost of implementing the standard has also come down.

All in all, this means that HEVC is currently a very attractive compression for production workflows and strikes a good balance between quality, latency and bandwidth usage.

That said, like H.264/AVC the quality can be significantly impacted by multiple encode/decode cycles. This is typically addressed by using higher bitrates (i.e. lower compression ratios) in each cycle.

GOPs

In temporal compressions such as H.264/AVC and H.265/HEVC, a Group of Pictures (GOP) is a collection of consecutive frames within a video stream that are encoded together.

A GOP typically starts with an I-frame, followed by P-frames and B-frames (see I/P/B-frames). The length and structure of a GOP can vary depending on the encoding settings and the content of the video.

In live production, the GOP normally needs to be short (e.g. IPPB) to minimize the encode/decode latency, while in distribution applications long GOPs are typically used to achieve higher bandwidth reductions with the same picture quality.

It is also possible to further decrease latency by various techniques to slice the frames horizontally and vertically.

NDI

NDI is a video connectivity system originally developed by NewTek, widely used in streaming and cloud-based productions but far less commonly in high-end broadcast applications. It supports multiple codecs, including the proprietary SpeedHQ (NDI High Bandwidth) and the more common AVC/H.264 and HEVC/H.265 (NDI HX formats).

SpeedHQ, based on MPEG-2 intra-frame encoding, is known for low latency (less than 100ms end-to-end) and configurable bandwidth, typically 150 Mb/s for HD signals. While it supports 4:2:0 8-bit (and 10/12-bit HDR from NDI version 6), SpeedHQ has issues with multigeneration encoding (like other MPEG-2 codecs), limiting its use in high-quality productions.

NDI HX formats, using AVC or HEVC compression, offer similar video quality at lower bit rates, making them suitable for limited bandwidth environments but with increased latency. NDI HX support is commonly found in hardware devices like PTZ cameras and mobile phones, and can also be used in software applications.

H.266/VVC

Versatile Video Coding (VVC) also known as H.266 further increases the compression ratio compared with HEVC. However, this comes at the cost of much higher encoding (and decoding) complexity. Currently, there is no widespread adoption of VVC in the production space, but with advances in hardware/software and algorithmic efficiency, it should be expected that VVC will also play a role in this space in the future.

Key codecs compared

The figure shows how different compression codecs (including uncompressed) compare in terms of end-to-end latency, typical compression ratio and ability to support multiple codec cycles without artifacts. As can be seen different codecs have different sweet spots and the choice depends what is most important for specific production requirements.

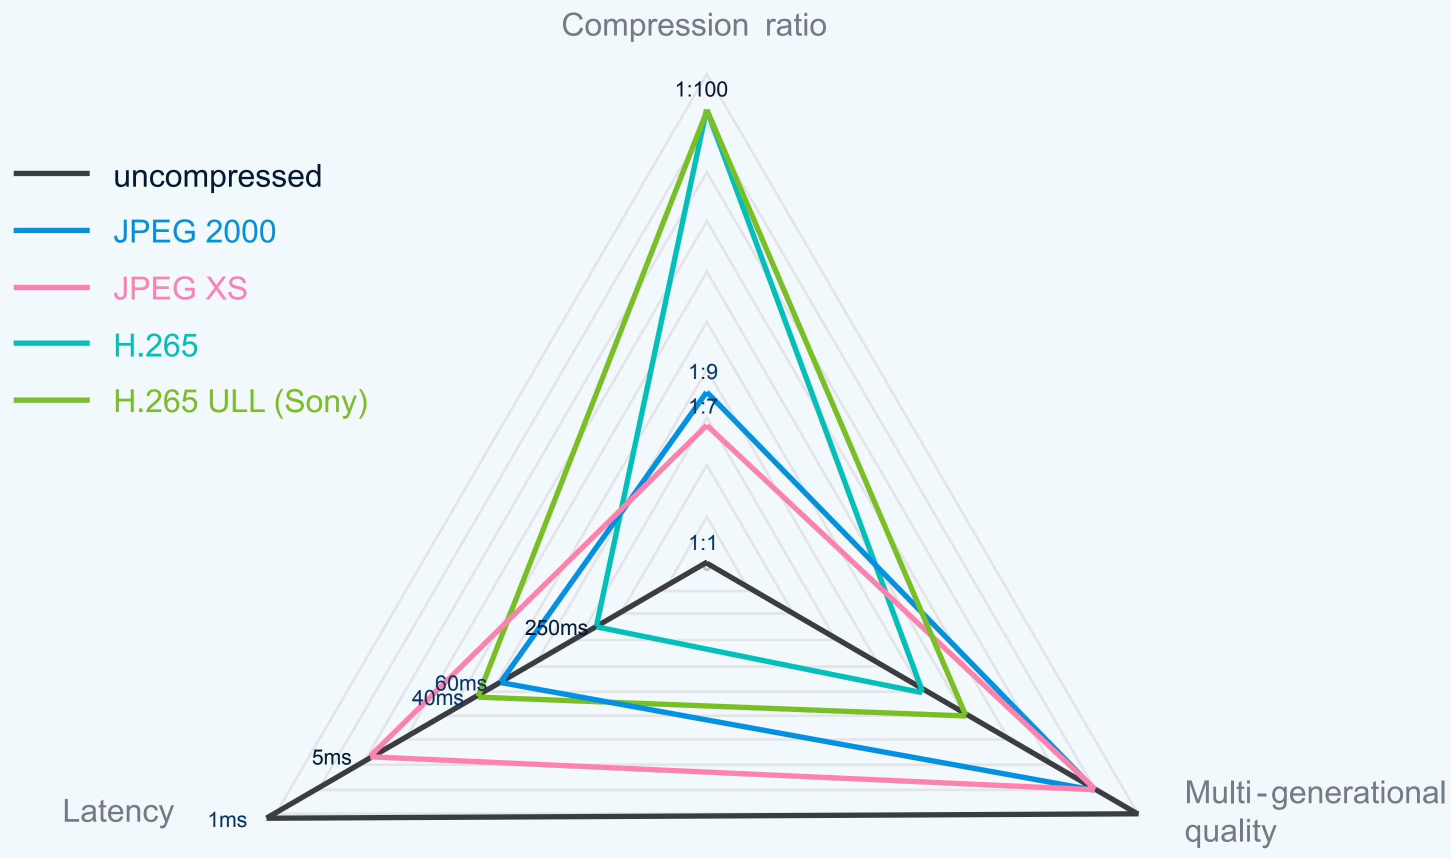

NB: The latency and compression ratio scales are logarithmic

Picking the right compression

Facilities

Historically live production facilities have used uncompressed video throughout the workflow.

One of the reasons for this has been that, in an SDI world, network capacity has been constrained by cabling and the availability of ports, not the bandwidth occupied by video. More importantly, using compression in the facilities would require repeated encoding and decoding cycles that increase latency and potentially reduce the quality of the video.

In an IP environment however, bandwidth can become an issue, especially as more productions take place simultaneously on the premises (not least because of the trend for remote production) – each involving more video signals and higher quality images (UHD/4K, HDR, HFR, etc.). Adding capacity can be comparatively cheap if the IP network has been architected properly and is controlled flexibly. Nonetheless, a case for video compression may also be made.

In such situations, JPEG XS is currently the only choice, as it offers the best combination of ultra-low latency, maintenance of image quality (after multiple encoding and decoding cycles) and cost.

Many suppliers of broadcast equipment are now building in JPEG XS capabilities into their products, facilitating the adoption of the standard in the facilities.

Contribution and remote production

When it comes to WAN applications such as contribution or remote production, the choice of compression is wide.

Typically, the decision of which compression to use will be affected by several factors.

Bandwidth availability

A key consideration is obviously the amount of bandwidth available in the network in relation to the bandwidth required for the signals to be transported.

If bandwidth is plentiful (and cheap), compression might not even be required.

At the opposite end, if bandwidth is very constrained and expensive, high compression (aka low-bit rate compression) such as H.264/AVC or H.265/HEVC are likely to be the preferred options. These do come at a cost though (see some of the other factors below) so for many mainstream applications JPEG XS (or sometimes JPEG 2000) may be the preferred choice.

Network reliability

Unmanaged networks, such as the Internet or 5G can provide cheap, ubiquitous and very convenient ways of transporting video in contribution and remote production. On the other hand they can offer unpredictably fluctuating bandwidth, which is an issue for live production.

Selecting the right compression technology in such an environment is challenging, but generally it is best to try to reduce the bandwidth requirements as much as possible. For this reason, HEVC is a good choice for unmanaged networks in combination with network healing (transport protection) technologies like SRT or RIST.

Latency

Keeping latency as low as possible throughout the production workflow has always been a primary goal for broadcasters. However, compromises are often made if the benefits are deemed to be sufficiently great or important. One example of this is the use of AR (Augmented Reality), which adds considerable latency, but is seen as a valuable tool in production.

Video compression is another source of latency, which needs to be outweighed against the benefit of using less bandwidth.

Until relatively recently, the decision was quite simple: spatial compression methods, especially those that use a slice approach (e.g. JPEG XS) offered considerably lower latency than temporal methods (e.g. H.265/HEVC). So if low latency was more important than compression ratios, JPEG XS or even JPEG 2000 have been the codecs of choice.

However, as mentioned before, low and even ultra-low latency versions of temporal codecs have now been developed. These are still slower than codecs such as JPEG XS, but the difference is not as significant as before, meaning that HEVC can now also be considered in situations where latency is the critical factor.

Type of production

The type of event being produced, and the purpose of the production could have a significant impact on the choice of codec.

High quality and low latency are always crucial, but they become even more critical when creating high-value feeds for major international events to distribute to rights holders in various countries, who in turn will process them further in their own productions. In such cases, uncompressed or lightweight (spatial) codecs may be the preferred options. Conversely, for comparatively low value productions, bandwidth usage may be the priority, favoring temporal compressions.

Another consideration is the type of events being produced. For interactive applications, like two-way interview links, quality may not be as critical as for high profile events, but latency is. This means that codecs like JPEG 2000 and JPEG XS will be preferred.

Fast-paced sports, regardless of the value of the event, require high-quality pictures and have always been a challenge for codecs like MPEG-2, H.264/AVC, H.265/HEVC, as they tend to introduce compression artifacts in areas with high motion. For example, in ice hockey the puck could disappear with such codecs. This makes JPEG 2000 and JPEG XS much better choices – these codecs also being very efficient with images with a predominant color, such as the white of the ice-rink. In contrast, “talking heads” (e.g. political debates) are prime examples where temporal codecs are very efficient.

Image quality

The intended image quality will also influence the choice of codec. For example, if the event is an HDR production, the codec would have to support at least 10-bit color. This could rule out some codecs that do not support 10-bit (such as NDI prior to version 6).

If the format is 4K (or higher), this might rule out JPEG 2000 and JPEG XS because the required bandwidth (even with compression) is much higher than what is available at the venue. This is in fact directly related to the bandwidth availability factor above.

Production hops

A key consideration in picking a codec is whether the video is likely to be compressed and decompressed multiple times as it progresses through the production workflow.

Spatial codecs are more suited than temporal codes to maintain image quality through multiple cycles. If multi-generational quality is a key requirement, then JPEG XS, JPEG 2000 and TICO are the best options.

Multi-platform interoperability

Nowadays, video compression is required in many aspects of the production workflow, for transport between locations, but also to and from the cloud, and even within facilities. Furthermore, the encoding or decoding might be performed by a variety of devices from different vendors, including gateways, broadcast equipment and COTS servers (on the ground and in the cloud). For multi-platform interoperability, the best codec is without doubt JPEG XS, as it's relatively simple in computational terms so lends itself to widespread implementation on multiple platforms. This includes both hardware and software-based implementations without the need for additional hardware acceleration.

Cost

Cost is a multifaceted consideration. The cost of using compression needs to be outweighed against the cost of bandwidth, in the first place.

When it comes to the cost of the codecs themselves, H.264/AVC and H.265/HEVC used to be substantially more expensive than standards like JPEG 2000. But that is now changing and those high compression rate (or low bitrate) compressions are now at a price point that is much more comparable.

Another point that is worth consideration is that H.264/AVC and H.265/HEVC still require hardware acceleration to achieve sufficient density and performance for production purposes. This also has cost implications as it means that dedicated hardware is typically required.

Transport options

Most codecs are compatible with MPEG-2 TSoIP (SMPTE ST 2022-2) and this is important for industry inter-operability. However, when it comes to SMPTE ST 2110, only JPEG XS (TR-08) and VC-2 are currently formally supported through SMPTE ST 2110-22. Typically, JPEG XS is preferred as it performs better than VC-2, e.g. it offers a high compression ratio.

TSoIP is attractive for WAN transport as video, audio and data are combined in a single stream. This means that sync is maintained during transport, but packets are small, which requires additional bandwidth compared with SMPTE ST 2110.

SMPTE ST 2110-22 is attractive for production workflows as video, audio and data are sent as separate essence streams, which means that the workflow is the same as for uncompressed SMPTE ST 2110. This is a better choice if compression is used in a facility, for instance to reduce bandwidth for UHD feeds.

Unmanaged networks can have variable characteristics and need additional protection to recover from bandwidth contention and other intermittent network issues. Compression is typically combined with SRT and RIST technologies across such networks.

Summary

In summary, for transport applications in the facilities, JPEG XS is the only alternative to uncompressed video.

For WAN transport (including ground-to-cloud-cloud-to-ground), if latency and multiplatform support are more important than bandwidth usage, JPEG XS is again the best option. However, JPEG 2000 can also be a good alternative for some applications because it is long-established, and TICO can be useful for 4K/UHD over legacy (3G-SDI) infrastructure.

For WAN transport where bandwidth is a critical factor, H.264/AVC and H.265/HEVC provide the best options, as they offer the highest compression ratios. The traditional disadvantage of those temporal codecs, higher latency, is rapidly disappearing as ULL versions come on the market. Broad hardware and software support and multi-generational robustness remain weaknesses though.

End-to-end latency, compression ratio and multi-generational quality are key parameters for selecting the right codec.

Encoding **flexibility**

There isn't a universal answer to the question of which compression technology should be used in live production: it very much depends on the circumstances, the needs and the priorities of the production.

Sometimes, these factors are fairly constant for specific types of productions. For example, bandwidth is always constrained when transmitting from a particular location, so H.265/HEVC should always be used.

In other cases, the factors will vary: 10 camera feeds might be needed for one event, necessitating a light-weight compression like JPEG 2000 or JPEG XS; then the following production from the same location requires 20 camera feeds, which demands a higher compression such as H.265/HEVC to fit down the same link.

This in turn will influence the encoding/decoding devices needed – whether single codec or multi-codec.

Sony and Nevion offering

As part of their Networked Live ecosystem, Sony and Nevion offer several video encoding / decoding solutions including:

Sony CBK-RPU7

The CBK-RPU7 is an H.265/HEVC 4K/HD remote production unit with ultra-low latency, low bit rate and high picture quality. Acting as a portable 5G video transmission device attached to a camera, the CBK-RPU7 is designed to support wireless and remote production, which enables flexible and robust content creation. The CBK-RPU7 is perfect for low-bandwidth 5G connectivity.

For more information: pro.sony/cbk-rpu7

Sony NXL-ME80

Acting as a gateway between LAN and WAN, the NXL-ME80's advanced H.265/HEVC technology achieves ultra-low latency, high picture quality and low bitrate, resulting in an efficient use of network bandwidth for remote production and video contribution. The NXL-ME80 is ideal ultra-low latency point-to-point transport over for low-bandwidth networks.

For more information: pro.sony/NXL-ME80

Nevion Virtuoso

Virtuoso is a software-defined media node that offers media transport, processing and monitoring capabilities. It includes support for the reliable transport of uncompressed and JPEG XS, JPEG 2000 (including ULL mode), TICO and H.265/HEVC encoded video signals. Virtuoso is the ultimate solution where choice of codecs (and other media transport and processing functionality) is needed.

For more information: nevion.com/virtuoso/

Nevion is a Sony Group Company

pro.sony/
networked-live

nevion.com