WHITE PAPER

HIGH DYNAMIC RANGE VIDEO

Implementing Support for HDR using Software-Defined Video Solutions

TABLE OF CONTENTS

TABLE OF FIGURES

INTRODUCTION

Video continuously moves forward. Black and white video ultimately gave way to standard definition color. In turn, the limited SD palette gave way to high definition with its higher resolution and more numerous and richer colors. Technology continued to march on, and ultra-high definition with even better resolution and deeper colors is upon us. The next step in the evolution of video is high dynamic range (HDR), an emerging standard aimed at increasing the number of colors and luminance available with video to more closely approximate what is visible to the human eye in the natural world.

There are three techniques by which to improve video experiences that are likely to be implemented in the near term: increase the video frame rate, increase the number of pixels in each frame, and/or improve the pixels themselves. This paper focuses on the third approach and makes the case that the video ecosystem is best served by implementing support for HDR in software given the uncertainty of the current landscape with regard to standardization.

HDR: A WORK IN PROGRESS

Much of the real excitement in the video field today focuses on improving pixels and enabling those improvements to be accurately disseminated through pay TV, Internet, satellite and other video workflows to produce lustrous images on the 4K monitors that are becoming more commonly available in the market.

The standards supporting HDR are very much a work in progress. For both technical and business reasons, the industry has not reached a consensus as to what standards will look like. More specifically, fundamental questions remain around the way in which HDR metadata will be encoded, transmitted and decoded. Subsectors in the industry – for instance, pay TV providers and the content owners that support them – will thrive or be challenged, depending upon the set of specifications that are ultimately adopted.

A LONG HISTORY

Efforts to translate electronic signals into colors that approximate the real world are as old as broadcasting. Color research was underway decades before color television was introduced, although study was interrupted by World War II. In 1931, the International Commission on Illumination (CIE) created the CIE 1931 RGB color space and the CIE 1931 XYZ color space. Six decades later – in 1990 – the high definition television standard was set under ITU-R Rec. 709. It limited the color gamut to a comparatively narrow region of the CIE color space requirements. In 2012, ITU Rec. 2020 – ultra high definition – was published. High dynamic range is an effort to implement the expanded color space available with Rec. 2020.

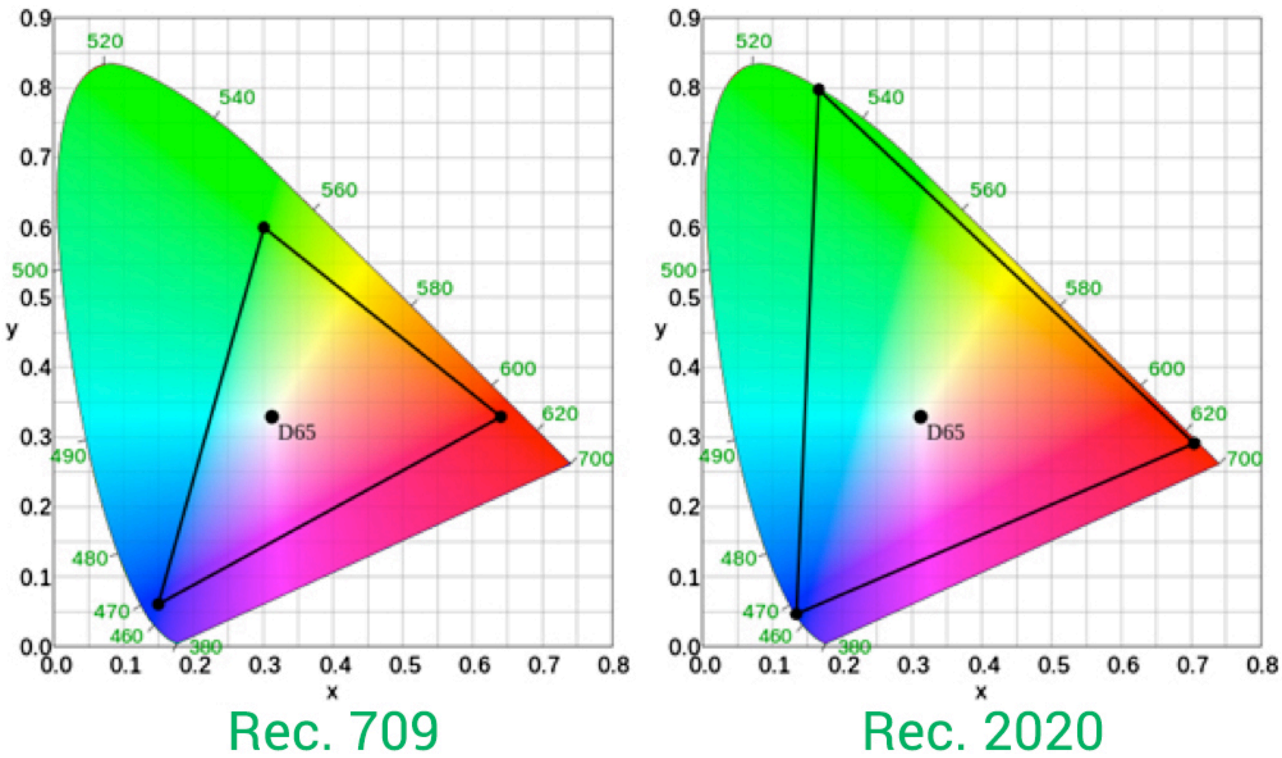

Figure 1 - Rec. 2020 offers expanded color space for High Dynamic Range video

EVERY BIT HELPS

The number of colors that can be displayed is known as color depth or bit depth. Color depth is either the number of bits used to indicate the color of a single pixel, in a bitmapped image or video frame buffer, or the number of bits used for each color component of a single pixel. There is a color depth of 256 when eight-bit pixels are used. Color depth jumps to 1,024 with pixels composed of ten bits. The quadrupling of color depth not only adds shades and gradations, but enables smoother visual transitions within color families. This can be critical: In the eight-bit environment, a stepping artifact can occur in which incremental moves from one gradient of a color to another are visible. The bottom line is that ten-bit pixels offer a range of colors closer to the limits of what human beings can perceive; and this occurs smoothly and without visible artifacts.

Figure 2 - 10-bit color depth portrays a smooth visual transition within color families

Increasing the number of bits is a good first step to improve pixels. It means little, however, without other equally significant changes. An evolution closely related to the addition of bits to pixels is the improvement in overall display technology. Cathode ray tubes (CRTs), which for decades were the only option for television displays, use electron guns and phosphorescent screens to render images. Similarly, some liquid crystal displays (LCDs), are backlit in only two places. If a bright event occurs on screen, the resulting image is comparatively over saturated. The image does not “pop” as the content creator almost certainly intended. Such images don’t do the eye justice.

This problem has been alleviated with modern LEDs and OLEDs that have regional or uniformly dispersed backlighting capabilities. With these displays, the transitions between light and dark – between the flames of a bonfire on a beach at night set against the darkness of the ocean beyond – are uniformly more extreme, dramatic and lifelike.

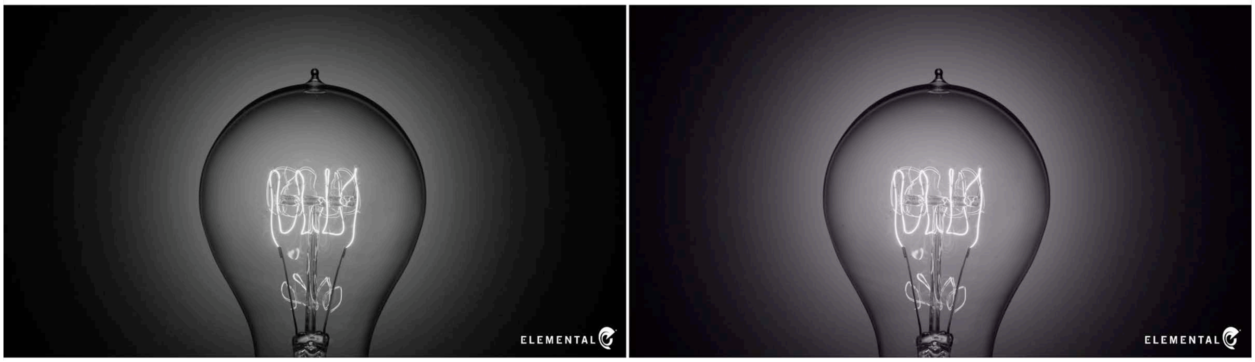

Figure 3 - Increased bit depth reduces visible banding to create a smooth gradient

A related advance that underlies the move to HDR is the improved ability of infrastructure to understand and mathematically represent the nature of encoded images and to use that information in the decoding process. Historically, a technique called a gamma curve was used to carry instructions from the encoder (and artistic director) to the display. The gamma curve, however, is another relic of the CRT era. It is no longer relevant to today’s modern displays.

HDR provides a means by which to describe and protect the content creator’s intentions via metadata. It contains in essence a language used by the content creator to instruct the decoder. HDR provides metadata about how content was created to a display device in an organized fashion such that the display can maximize its own capabilities. As displays evolve, HDR will allow existing devices to always make a best effort in rendering images rather than running up against unworkable limitations.

A formula called the electro-optical transfer function (EOTF) has been introduced to replace the CRT’s gamma curve. Some engineers refer to EOTF more simply as perceptual quality, or PQ. Whatever the name, it offers a far more granular way of presenting the luminance mapping according to the directions given by the content creator. EOTF is a part of the High Efficiency Video Coding (HEVC) standard.

The use of EOTF in HDR is particularly important because the end-to-end workflow – from the camera, through the encoder and distribution path to the decoder and ultimately the television screen – requires content to go through many transformations. Different standards – some relatively new, some old and set – are used at different points in the distribution path. The EOTF is a way to make sure that quality doesn’t degrade through these stages and provides a vital, high-level means by which to assess whether the artistic integrity of video content is being maintained.

FROM SCIENCE TO PRACTICE

The next large challenge facing the video industry is translating the science behind HDR into a system or systems that can actually perform the required tasks of making HDR a reality for consumers and provide a return on investment for providers. This adds complexity by bringing the laboratory into the marketplace.

Figure 4 - Standard Dynamic Range vs. High Dynamic Range video

It is important to understand the landscape and that HDR is in its infancy. A relatively large group of HDR solutions exist, but with significant technical differences between them. It is a facet of the industry that is only just emerging, but, at the same time, content creators and vendors want to quickly take advantage of the great images marketers have demonstrated with HDR. In addition, TV manufacturers are promoting HDR directly to consumers, with a desire to enhance 4K television offerings and achieve higher price points.

It will be some time before a stable environment based on standardized approaches to HDR coalesces for both traditional and OTT delivery. Today, end-to-end solutions are just beginning to emerge. Moreover, it is impossible to say which approach will be favored by different industries. In the absence of such basic foundational information, pay TV companies (cable operators, telcos, etc.) will likely demand different solutions based on local factors and the politics of their corporate ownership and alliances. In the absence of industry consensus, a scenario could arise in which content providers and owners must support multiple silos and workflows. This situation, which played out in the adaptive bitrate streaming sector, would be highly inefficient.

There are no lack of entities putting forward solutions. There are candidate approaches from Philips, Dolby, Technicolor and BBC/NHK as well as offerings from less well known entities. Other proposed solutions are likely to emerge. The most notable standards-based approach to date is HDR-10 (also known often as vanilla HDR), which is a combination of specifications from the Society of Motion Picture & Television Engineers. Other groups, such as the European Telecommunications Standards Institute, are working on full or partial systems. The Motion Pictures Experts Group (MPEG) is adding supplemental enhancement information (SEI) messages to its HEVC standard to support HDR. And, of course, these contender approaches are pushing behind the scenes to have their intellectual property become the standards-based approach.

The details of the differences between approaches are significant, but at a high level, there are two major considerations. The first is if a particular system is backwards compatible (i.e., whether or not it uses existing encoding/decoding equipment). The second is how the backwards compatible metadata for HDR and SDR is transmitted. Neither are trivial, to say the least.

BACKWARDS COMPATIBILITY

Settling on an approach that doesn't require the replacement of encoders/decoders is very important to some pay TV companies, distributors and device manufacturers. Dolby Vision, Technicolor, Philips and BBC/NHK are all backwards compatible. In a backwards compatible approach when an SDR television receives the video signal, the HDR metadata simply is ignored by the set-top boxes or display.

However, not all prospective customers of HDR care about backwards compatibility. Backwards compatibility is less of an issue in some distribution ecosystems, such as over-the-top (OTT) and Blu-ray2. A situation could arise, therefore, in which supporting HDR would require the change out of encoders, decoders and cable set-top boxes. Non backwards-compatible approaches no doubt will lead to significant expense to various members of the ecosystem because maintaining two sets of content (SDR and HDR) may become necessary.

ONE OR TWO LAYERS

The other structural difference between the approaches is how metadata is carried through the workflow. In a dual layer approach, SDR and HDR video streams are carried separately through the workflow. In a single layer approach, the extra metadata – what must be added to SDR to make it HDR – is integrated into a single layer plus side-car metadata.

The challenge of the dual layer approach is that legacy systems (encoders, televisions, set-top boxes, etc.) expect a single video stream. The introduction of a second video stream makes vital ancillary operations such as presentation of emergency messaging, on-screen displays and ad splicing and insertion more difficult. A dual layer approach provides the highest quality viewing experience without constraint with regard to data size. A single layer approach is better suited to fit within an existing video workflow.

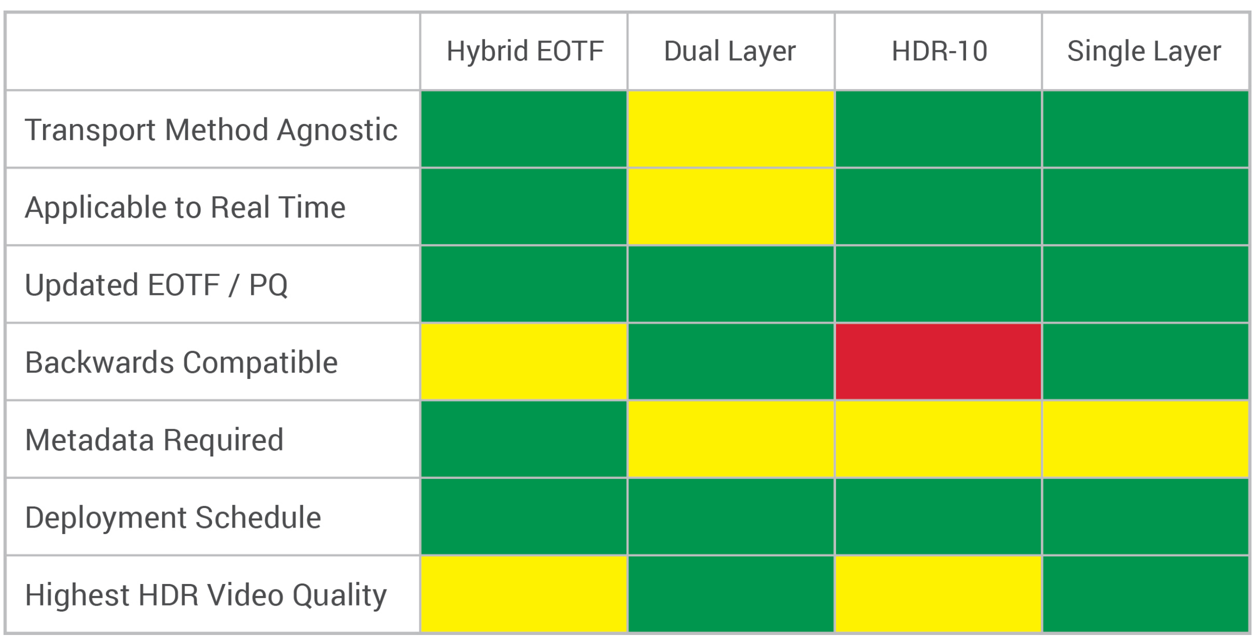

Figure 5 - HDR solution comparisons

A final concern is that the industry needs to settle on approaches to make use of live HDR possible in the near term. This is particularly important because live sports will be a key showcase and driver for HDR. If, for instance, a soccer match is played on a pitch in which one area is brilliantly lit and the other is darker (due to shadows caused by the stands, for example), HDR allows all areas of the field to be seen equally well. In contrast, the use of SDR forces a decision about which area of the pitch to make visible to ensure a satisfactory viewing experience.

CONCLUSION: SOFTWARE IS KEY

The benefits of HDR are great, and the industry wants to offer them to customers sooner rather than later. Powerful companies are introducing viable techniques for implementing HDR, while others are urging open approaches designed to keep the industry from descending into divided camps that could end up in a VHS versus Betamax or Blu-ray versus HDDVD scenario.

It will take some time for the issues to clarify and work themselves out. There is no guarantee that a single approach to HDR will be agreed upon. This is especially important for content owners. For example, a content owner might be asked to deliver HDR content in two different formats for two different distributors. This could be an issue for the content creator when it tries to market its library for reuse and rebroadcast.

There is an elegant answer to this seemly intractable quandary, however. It is a simple one that will enable content producers and the companies charged with encoding the content they create to keep the business side of the house happy and future-proof themselves. By creating a software-based HDR infrastructure using video solutions from Elemental, new and vibrant color and luminance standards can be implemented today without forcing expensive decisions or the maintenance of running parallel hardware-based HDR silos.

HDR is a key element in the overall evolution of video and part of a bigger picture in which many things change and evolve on an ongoing basis. In such an environment, keeping as many components in software as possible is an advantage: It reduces both capital and operational expenses. Elemental approaches the HDR challenge with software-defined video solutions. Flexible and scalable software from Elemental is designed to take advantage of the huge increases in general purpose processing power and the ascendency of the cloud to maximize video processing and delivery tasks that can be performed without significant hardware investments.