WHITE PAPER

Gear up for a brighter future with HDR

Contents

Introduction

The development of television technology - in the studio, on location, along the transmission chain and in viewers' homes - has been a continuing search for greater picture resolution and better looking images. From 405 line black and white - described as 'high definition' for the BBC's broadcasts from 1930 onwards - to 480i and 576i line colour, the medium made the significant jump into digital transmission in the 1990s, paving the way for HD broadcasting.

These transitions were relatively gradual over a period of 50-plus years. The rate of technological innovation and research has accelerated considerably during the 21st century, with the result that HD is now being challenged by Ultra High Definition (UHD) and its 4k higher resolution. The pace has picked up even more over the last five years as OTT and streaming services have changed the broadcasting landscape.

The new breed of entertainment distribution services, notably Netflix and Amazon Prime, have made UHD a primary specification for programming and are delivering it to homes as standard rather than being a special feature or test channel.

Satellite broadcasters, BSKyB and BT Sport in particular, are also keen to embrace UHD as with greater visual depth and clarity it delivers a more immersive viewing experience for drama and particularly for live sport such as football. However, as with digital terrestrial television (DTT), the constraints of the production and distribution infrastructure make adoption a more difficult proposition.

Why HDR & WCG?

UHD was developed by NHK Science and Technology Research Laboratories to deliver a screen resolution that was at least 4 times greater than HD, operating with 2160 lines progressive. The main thinking behind 4k/UHD is to present a cinematic experience in the home. This is not only because of the higher resolution but also a greater depth to the images that creates an immersive, near three-dimensional sensation without the need for special glasses.

While UHD itself is still in the early stages of rolling out for both production and transmission, broadcasters and particularly the new breed of entertainment distribution services, notably Netflix and Amazon Prime, are looking at techniques to further enhance and improve the viewing experience.

UHD has already brought us more pixels. The next logical step is not any additional pixels but better pixels that give a picture with more definition and, most important of all, improved colour reproduction and a higher dynamic range. Which is why the initials HDR (high dynamic range) and WCG (wide colour gamut) have been challenging those of UHD for prominence on booths at trade shows and in new TV sets advertisements.

HDR is able to produce a dynamic range of 200,000:1 (or 17.6 stops in camera terms) when shown on a 2,000 cd/m2 display with a bit depth of 10-bits per sample. This compares to the 64:1/approximately 6 stops from standard dynamic range (SDR) on a conventional gamma curve with a bit depth of 8-bits per sample.

By extending the dynamic range, more information can be accommodated in an image, bringing with it more dynamics in luminance to an image. In a similar way WCG allows more vibrant colours to be displayed because it can store a wider range of colour values than established RGB (Red Green Blue) colour spaces.

Both HDR and WCG are being discussed in professional and consumer electronics circles but HDR is the more recognisable of the two terms. It has become almost a catch-all for improved image quality, greater definition - which improves the immersive feel of 4k and creates a sense of reality - and lifelike colour. Because of these many broadcasters and consumers are saying they will move to HDR but by that do they mean just that format on its own or HDR and WCG? The broadcast industry in particular needs to be sure this is clear. It is very likely it will mean both. HDR means a blacker black and a whiter white. If we go to WCG as well the reds will be more red, the greens greener and the blues more blue, with everything else in between. The whole (image) space is bigger and more beautiful.

The broadcast industry needs to do this right because we already have a serious competitor in the form of new services such as Netflix, which can present pictures in excellent, extreme quality. Which means broadcasters and broadcast equipment manufacturers need to provide viewers with an absolutely stunning picture because the competition will only become stronger.

So broadcast has to move into more pixels, more colours and higher dynamic range. We need to provide better pixels and do it right first time. We have one chance and if we get it wrong the audience will switch off.

What is the new challenge?

If broadcasters implement both HDR and WCG it has to be in best possible way. This is necessary so we can preserve our precious industry. To do this properly, and be able to compete on a level playing field with competition from new media platforms, there can be no compromise.

The technical challenge behind this is the move from the HDTV colour space (BT.709) to the colour space for UHD plus HDR and WCG (BT.2020). In doing this we are going from a relatively small area (Figure 1: BT709 colorspace compared to BT.2020 colorspace) with a maximum brightness/luminance of 100 nits, all represented in the well-known 10-bit YUV definition that is used in broadast. Compared to BT.709, BT.2020 has a much larger colour space to work in, with luminescence measuring 1000/10,000 nits. (??), but as broadcasters want/need to be able to continue to use their existing infrastructure, we still need to code this in the same 10-bit YUV infrastructure.

In effect it may seem we are trying to put a square peg in a round hole. This is normally impossible without using brute force but with the correct forms of conversion and compression we can make it fit.

The whole issue becomes even more complicated if WCG is part of the equation. Higher dynamic range with some sort of backwards compatibility is only possible when everything is working remains in the same colour space. But because the broadcast market has to be part of the real world, services will have to include WCG, which will mean the whole issue of backwards compatibility in practice is down the drain.

To fully understand these issues we should be familiar with some of the history of colour technology, its terminology and the various formats being used during these formative stages of HDR and WCG.

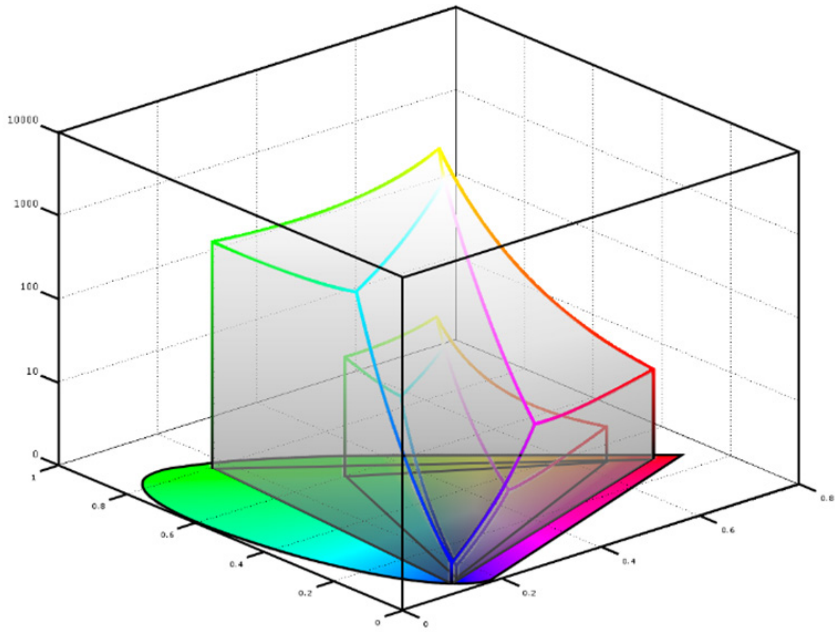

Figure 1: BT.709 colorspace compared to BT.2020 colorspace

History

Worldwide standards for colour and light are drawn up by the CIE, the Commission Internationale de l'eclairage or International Commission on Illumination. Founded in 1913 and headquartered in Vienna, the CIE defined the quantitative links between physical pure colours - otherwise called wavelengths - in the electromagnetic visible spectrum in 1931. The CIE 1931 RGB and CIE 1931 XYZ colour spaces laid down the basis for colour technologies that followed and still have an influence today.

The colour space chromaticity diagram that forms part of these spaces is quite an impressive curve (Figure 2: CIE 1931 colour space). Shaped like a horseshoe it encompasses everything the human eye can see, starting at 460 nanometers, just before ultraviolet, all the way up to 630 nanometers. It has a narrow spectrum from blue to green to red, with other colours created in the middle by mixing the three primaries. Right in the middle is D6500 (a colour temperature of 6500 Kelvin), which corresponds to natural daylight at midday in northern/western Europe.

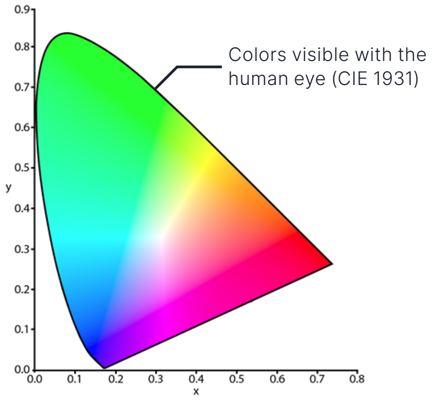

Figure 2: CIE 1931 colour space

Terminology

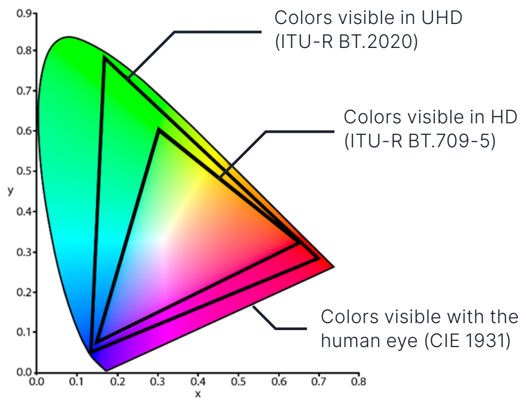

Figure 3: BT.709 and BT.2020 curves

Since the publication of CIE 1931 RGB and CIE 1931 XYZ, a succession of colour spaces and curves has appeared, each tailored to either the requirements of broadcast/film technology at the time or a specific application.

In 1982 the predecessor of the Radiocommunications section of the ITU (International Telecommunications Union) issued Recommendation BT.601 (Rec.601 or BT.601). This set out the colour space for standard definition images in 525-line 60Hz and 625-line 50Hz broadcasting, using YCbCr 4:2:2 coding.

BT.601 was superseded by BT.709, which was approved for HD, widescreen (16:9 aspect ratio) TV in 1990. This was based on 1980-lines and produced more colours than 601 but was nothing to get excited about. It has a fairly small colour palette in terms of what the human eye can see.

The big jump came with BT.2020, which appeared in 2012 and laid out parameters for UHD. This covers two resolutions: 3840x2160, otherwise known as 4k, and 7680x4320, or 8k, which forms the basis of NHK's Super Hi-Vision TV system. It also specifies bit-depths of either 10-bits per sample or 12-bits per sample. BT.2020 much closer to the capability of the human eye but takes up only a relatively small proportion of the CIE 1931 colour space.

As we move from BT.709 to BT.2020, the volume increases logarithmically and exponentially from 100 nits to the 1000s and even 10,000s. As stated earlier, the problem is that these larger ranges in colour and luminance amounts still have to fit into 10-bits. In a linear environment we would need much more than those 10-bits. If you go from black to maximum white plus all the edges of the colours you would need approximately 40-bits. In photographic terms that's the equivalent of 40-stops on a camera lens. But 40-bits as an infrastructure is too expensive and definitely not backwards compatible. So we need to make fit everything into 10-bits.

This is where gamma curves come in. We need some sort of curve that will use the 10-bits representation of YUV more efficiently. It is a method of squeezing data into a 10-bit environment. And you need to do this in a way the human eye will not notice artefacts. This can be done in different ways. A fundamental aspect is that more bits are necessary in the darker portions of the picture (where the human eye is very sensitive for slight variations) and fewer in the light (where the human eye is much less sensitive).

The volume of the BT.2020 colour space is much larger than that of BT.709. What is also obvious and makes moving from one colour space to another a bit messy is that the area taken up by BT.2020 does not expand over the centre of the available space. If the colour space had been a perfect circle with white in the middle then compatibility would be easier because moving from one colour space to another could be achieved by changing the gain. But because the BT.2020 space is a triangle that does not expand over the centre, there are a number of inconsistencies and differences in the amounts of the primary colours that could be used. On the BT.2020 curve green has increased considerably but is not well represented in the BT.709 curve. Consequently it is mathematically almost impossible to go from one to the next and you need something else to make the transition.

The gamma curves we are using to make it fit 10-bit, and the fact there are two colour spaces that are not expanding over centre points, cause a problem with interoperability between signals. If you record material in a BT.2020 space and want to show it in a system that expects BT.709, the process goes wrong because material recorded into one curve does not look right on a display that is expecting a different curve.

Three different ways to represent HDR

To accommodate what are becoming necessary additions to the broadcast vision production and transmission processes, three new curves have been developed, each with its own approach to improving dynamic range.

Perceptual Quantisation (PQ) is a curve that has been optimised for the human eye through extensive research carried out by Dolby Laboratories. It is part of Dolby Vision, the 4K video display technology that also incorporates HDR and WCG. Dolby Vision has been adopted by both TV set/ receiver manufacturers - including LG and VIZIO - and streaming video services such as Netflix, Amazon Video and VUDU.

Although developed by Dolby, PQ is an open curve for HDR that has been standardised by SMPTE (Society of Motion Pictures and Television Engineers) as ST 2084. It is the result of testing 2000 humans throughout the world, because, depending on where you come from, you perceive light and colours slightly differently. There is also a difference between the perception abilities of men and women.

Without making such accommodations and adjustments in the curve there is the chance that banding artefacts would be produced. These can be seen in badly compressed YouTube movies. For example, if there were a dark underwater scene you would see strange artefacts in the dark edges. Which is why more grades of darkness are necessary combined with a higher number of bits.

The PQ curve was created to achieve this and goes all the way up to 10,000 nits. This is probably overkill for the average household TV but you certainly need more than 1000 nits, maybe around 4000, to produce the best viewing experience.

Despite its positive features, PQ is still an odd curve. It hits only 10 percent of maximum light output at 75 percent of the 10-bit space used. This is because the human eye will see errors in the dark much sooner than in the light.

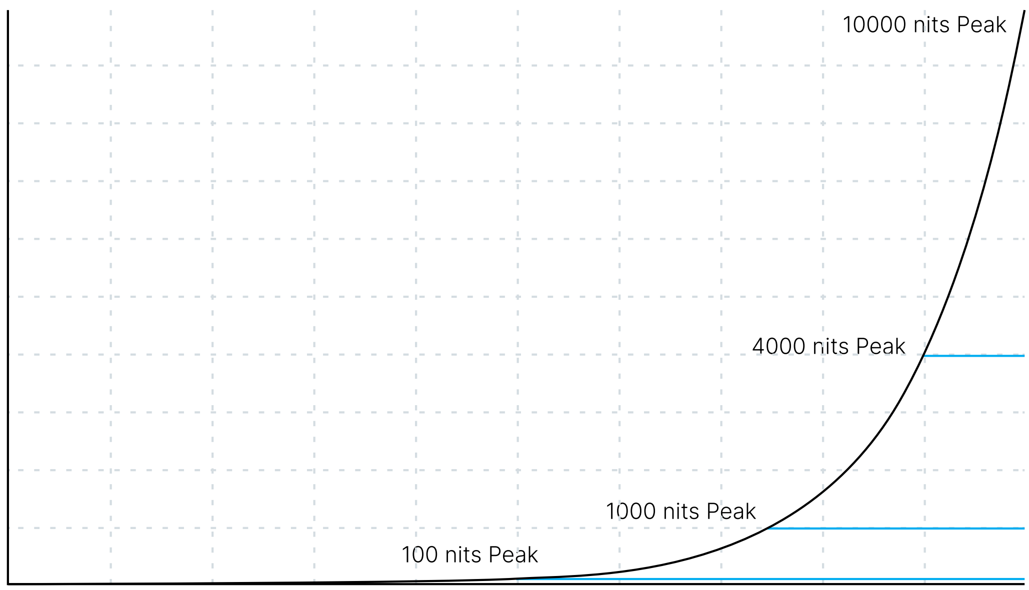

Figure 4: PQ curve

The second curve is **HLG (Hybrid Log Gamma)**, which was developed jointly by the BBC and NHK. It has some backwards compatibility, although only in BT.709, and increases the dynamic range of the video more than conventional gamma curves. This is done by using a logarithmic curve for the upper half of the signal values. HLG is royalty-free and has been approved as ARIB (Association of Radio Industries and Businesses) STD-B67. It is a non-linear transfer function in which the lower signal values, which make up half the total amount, use a gamma curve and the upper half of the signal values use a logarithmic curve.

The third curve, **S Log 3**, is purely a production/acquisition tool and features on Sony digital cameras. The Sony F65, F55 and F5 use two new grading spaces, G-Gamut3.Cine/S Log3 and SGamut3/S-Log3, which are intended for Log Base Grading. There has been some production in Slog3 but it involves converting to another output, going from an OB truck, for example, to a PQ or HLG curve.

In effect, there are different curves available depending on how much peak brightness you want to achieve.

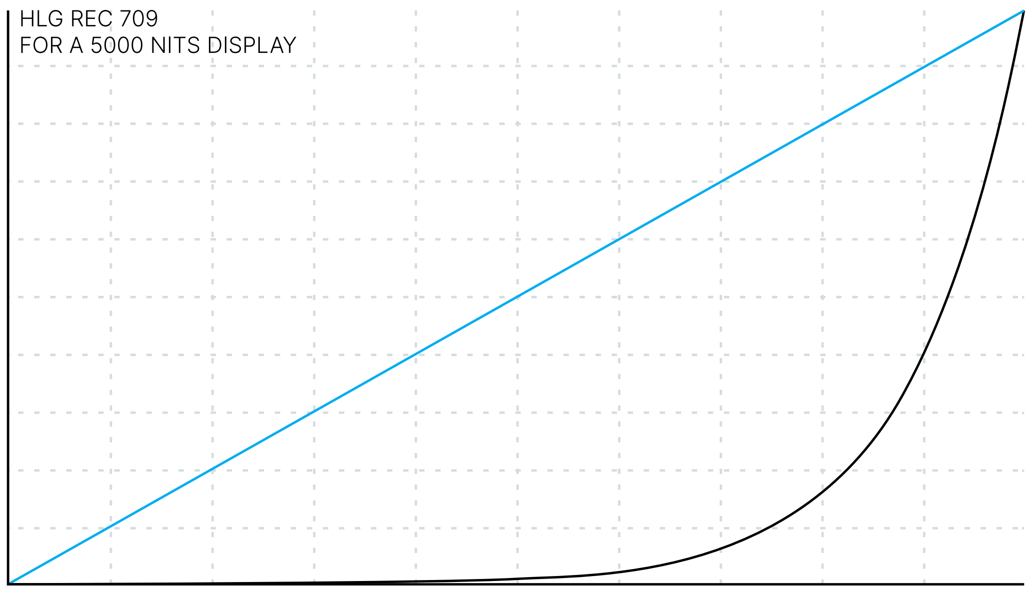

Figure 5: HLG curve

Making the changes

Broadcast image technology is in a state of flux right now, with several formats and systems - both old and new - likely to co-exist for some time to come. Going from SDR to HDR and back is difficult. The broadcast industry is still learning what all this means and what needs to be done. Live event productions such as motor racing coverage are a case in point. While HDR/WCG production cameras are available now, the technologies do not yet extend to the miniature cameras that are used for on-board shots. Converting seamlessly from one to the other will be difficult, just it was when switching between HD and SD in the early days of high definition TV.

LUT based conversion

Going from one colour space to another and dealing with different dynamic range systems is also challenging. Fortunately there is a long established technique for this involving look-up tables (LUTs). A LUT is a list of parameters - a row of numbers - with an input and an output. It tells you that if you go in with a number, you have to come out with a different number.

A LUT has to be made for the different colours. So a single LUT would consist of LUTs for R, G and B. The ultimate goal is to maintain equal and realistic skin colours. If skin colours start to shift any problems with the colour become immediately very apparent. The human eye is very perceptive to skin colour changes.

A 1D LUT is quite easy and basic, with brightness and tones in conveyed in RGB (Figure 6: Example of a 1D LUT). Any particular value of red becomes another value; for example 60 in becomes 80. If you put in RGB 60-60-60, the LUT comes out as 80 R, 85 G and 75 B. 70 equals 90 R, 95 G and 85 B and so on. It is a fairly straightforward method of transferring a value to another.

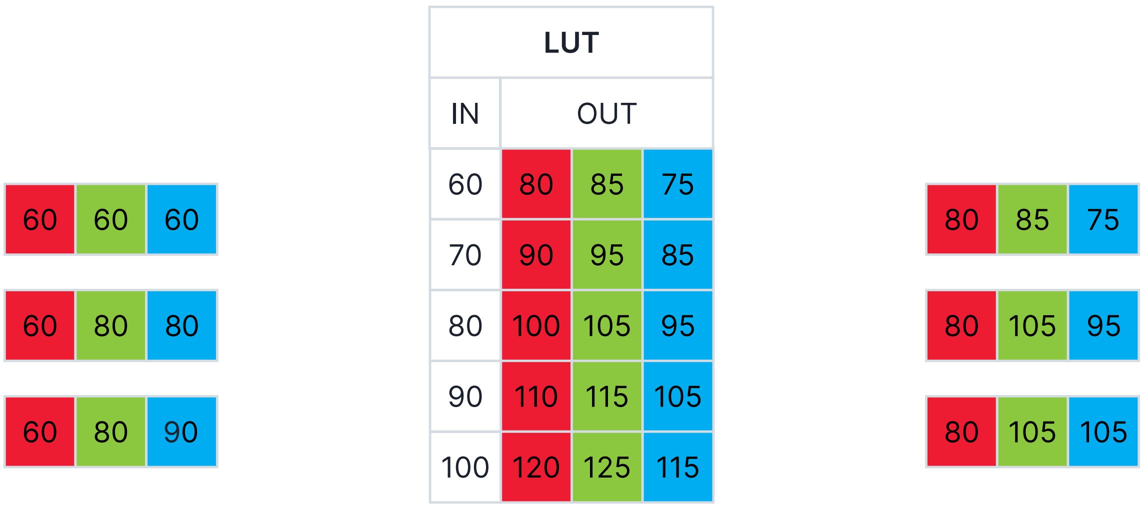

Figure 6: Example of a 1D LUT table

A 3D LUT is more complicated, as can be seen in the figure below. Instead of applying independent modification to each R, G, B code values, 3D LUT applies conversion based on the RGB triplet.

A further complication is that LUTs are handled in a RGB domain. This means a matrix is needed to convert from YCrCB to RGB. By taking the colour corrector in one of our cards, the incoming YCrCB signal is converted into RGB, colour corrected and then outputted as YCrCB using an inverse matrix. The problem is that the matrix going from YCrCB to RGB is different in BT.709 and in BT.2020, calling for two conversion matrices. As explained, the LUT file is just a simple conversion table to be applied for a specific color adaptation. Because of this the colour space and HDR profiles of the source and the destination have to be clearly identified to match with the LUT purpose. It is expected that very soon we will have a standard embedded in the signal just as happened with widescreen signaling between 4:3 to 16:9. We need colour space signaling that tells you what the source is because you have to adjust your matrix according to what comes in.

Adaptive HDR conversion

Another method of HDR conversion is by means of a frame by frame adaptive algorithm. This technique is much easier to setup since it doesn't require any customized LUTs or parameters, other than choosing whether you're converting SDR to HDR, HDR to SDR or HDR to HDR (changing curves, for instance from S-LOG 3 to HLG) and choosing your curve. It also doesn't require using any metadata.

The algorithm ensures that you have an optimal conversion, regardless of the video content. The conversion adapts to the content automatically in real-time (with ultra low latency) without the need for any manual adjustments. This is of course easier in use, especially in case of ad-hoc conversions, but you cannot change the color grading of the end result.

Adaptive HDR conversion does however give a perfect roundtrip result, meaning that if you convert SDR to HDR and then back to SDR, the SDR end-result looks exactly the same as the original picture, no matter if you're working with HD or UHD or whether your working with HLG, S-LOG3 or PQ curves.

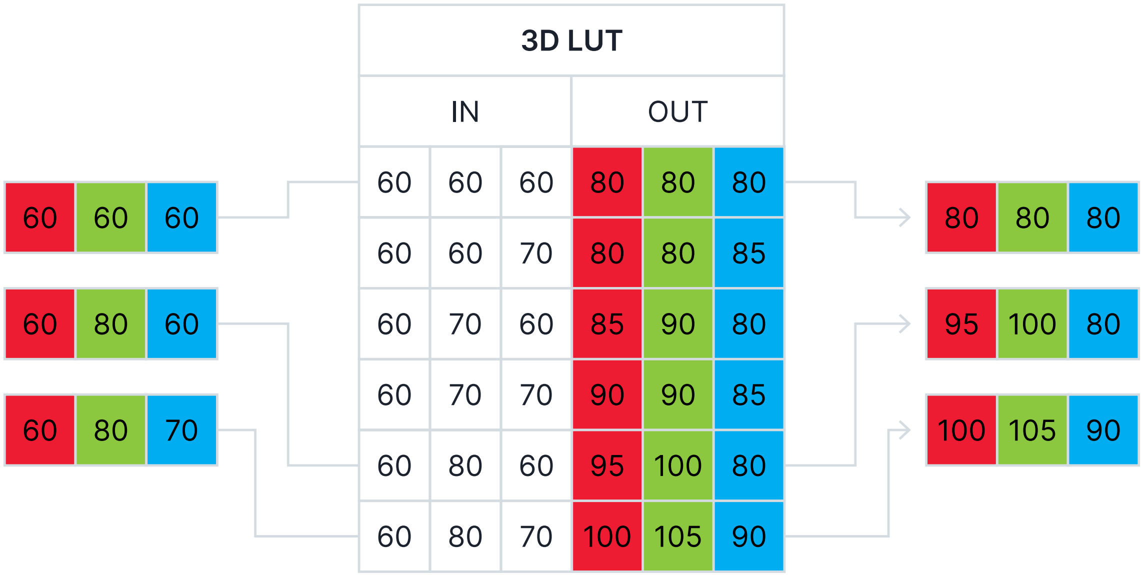

3D LUT					
IN			OUT		
60	60	60	80	80	80
60	60	70	80	80	85
60	70	60	85	90	80
60	70	70	90	90	85
60	80	60	95	100	80
60	80	70	100	105	90

Figure 7: Example of a 3D LUT table

The solutions from EVS

Real time signal conversion solutions

EVS produces a number of **Synapse** modules that are able to convert the various colour spaces and LUTs while at the same time working with HDR and WCG. The UXU500-510 is designed to convert from HD (SDR709) to UHD (HDR2020) and back for SDI I/O. A standard set of (BBC) LUTs is included, along with a piggyback PCB for storing the LUT presets. The unit is compatible wit 3D LUTs and dual (20 bit) 1D LUTs for very precise LUT based conversion with minimal interpolation artefacts. A selection of presets is available, allowing many different LUTs to be held on the card. These can be switched on the fly, with the ability to use split screen mode to compare and see what is going in and out.

The **Synapse** U4T200-240 is designed as a UHD 4 wire toolbox. It features 4 framesyncs to allign the 4 wire interface two de-embedders and two embedders (in Q1 and Q2) It comes with a 4 quadrant LUT based HDR<>SDR converter identical to the UXU500 mentioned above. A Dolby E decoder and encoder is included in the U4T240.

For Live IP or hybrid SDI+IP environments there is an HDR conversion option available in our **Neuron** product line. This option offers an adaptive/dynamic HDR to SDR (and back) conversion as well as HDR to HDR conversion. These converters offer real-time frame-by-frame HDR <-> SDR and HDR <-> HDR conversion techniques, without any need for manual adjustment, correcting scene based adaptation that replaces the additional shader workflow needed in many environments. **Neuron** will also offer real-time LUT-based conversion.

As part of the **XS-VIA** and **XT-VIA** server offering, the embedded multiviewer also plays a role in real-time conversion from HDR to SDR. Having an HDR to SDR converter built in, the Total Cost of Ownership can be reduced by using traditional SDR monitors instead of having to use more expensive HDR monitors. The internal mutliviewer will apply a conversion to SDR so that the inputs and outputs can be viewed as close to the real picture as possible, for the replay or highlights operator, or for any other position that is leveraging the EVS Multiviewer.

File-based conversion solutions

The EVS solution for transforming audio and video files from and to HDR or SDR is **VIA XSquare**. VIA XSquare is behind various tools which EVS offers to production staff.

For simple and centralized restore and archive to the server, **XFile3** leverages VIA XSquare. With the adaptive/dynamic conversion feature it makes it easy for the operator to restore and archive directly to the EVS servers. The BBC LUTs are included in case a more advanced 3D LUT based conversion is needed. The operator can see what type of HDR or SDR flavour is used at any time during the production and can decide on-the-fly if conversion is needed or not.

In setups like broadcast centers our asset management system, called **IPDirector**, can show all metadata information and share this metadata with NLE systems. Here, VIA XSquare also runs in the background and facilitates the file movement and conversion in the background. Multiple transcoding engines can be scaled and orchestrated to make sure that delivery between EVS solutions, but also towards and from 3rd party systems, is not delayed and processing is done in the most efficient way.

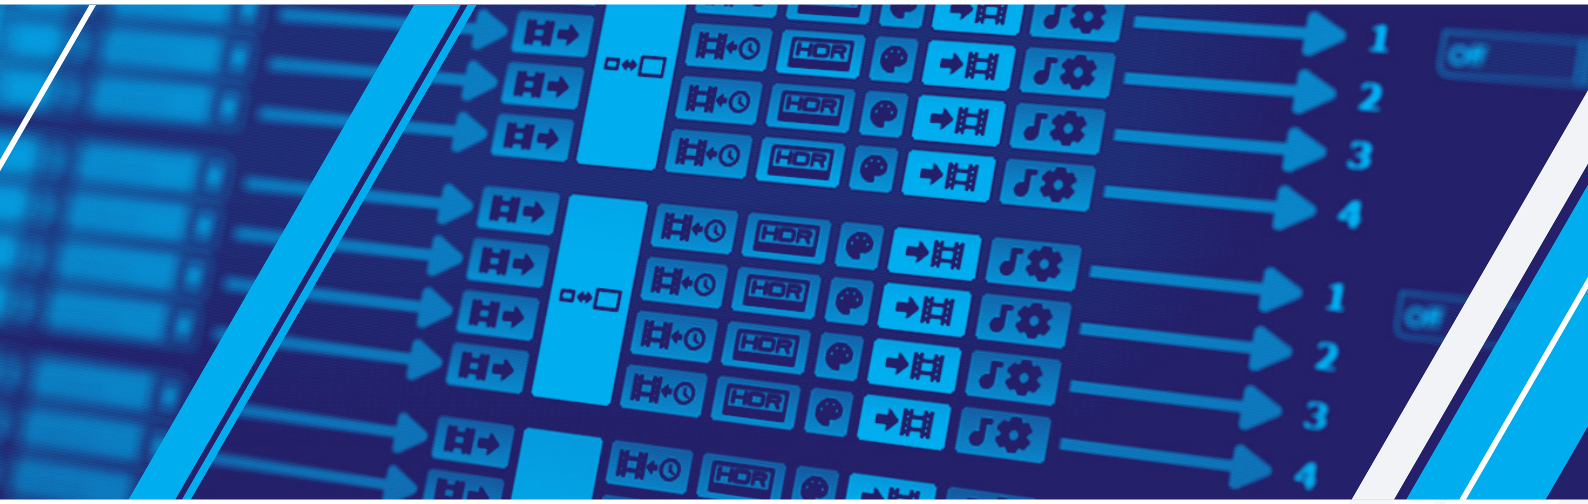

Conclusion

Unlike other technologies such as 3D, HDR is a real game-changer. It provides a transformative immersive experience, but one which requires minimum effort from the viewer and it's in easy reach of the mass market. More and more broadcasters are launching HDR services, and it is in live sports that HDR and WCG truly come into their own, delivering vivid life-like pictures with wider contrast and greater depth – banishing shadows that often plague viewers' enjoyment of the action.

Thanks to accessible technologies, using various curves, the much higher range of colors and luminance in the BT.2020 colorspace still fits in a 10-bit YUV definition. The PQ, HLG and SLOG-3 curves ensure that the information which would usually require 20 bits, can still be broadcast using the same formats and bandwidth as a non-HDR transmission.

However, when you are working in HDR productions, you will run into the challenge of having to convert SDR to HDR, HDR to SDR or converting from one HDR curve to another. This is a challenge for real-time baseband signals, as well as for file-based video.

To address this issue, two types of conversion can be used, based on static Look Up Tables (LUTs) or based on adaptive/dynamic algorithms. Using LUTS has the advantage that you can choose between various LUTS with different outcomes. You can even fine-tune the resulting colors and luminance with the right color grading expertise. You can use various LUTS for various events, or create certain effects. The conversion is static though, so fast changes in brightness are not compensated.

Using Dynamic conversion is much easier compared to LUT based conversions. There are less settings to configure and you don't need a color grading expert to choose or create a LUT for you. Dynamic conversion compensates the color and luminance output when the brightness between scenes changes. On top of that, it ensures the perfect round-trip conversion. This immediately highlights the down-side of this type of conversion: you do not have any control of the output. What you see is what you get, without the possibility to make adjustments.

The choice of which conversion method is best, differs per use case. EVS' solutions integrate both methods to ensure the highest quality SDR and HDR outputs in every situation, whether you are working with baseband or file-based video.