What is 4K, UHD, SLog3, Rec 2020

And other really boring things.

Compiled By Peter Morrone

The process of capturing and reproducing images requires a collaboration of camera sensors, file formats, rendering technologies, and display or printer technologies. All of these have different ways and different capabilities of representing color and intensity. In addition, they are all different from how our eyes work which further complicates things. As a result, over the years, several standards and processes have been implemented to accomplish this. They all involve some aspects of how to capture and store colors, what range of colors can be dealt with and how to adjust intensity to best reproduce the real world. To understand the new 4k technologies, including SLOG3, HDR, Rec 2020 etc, an understanding of

- Gamut
- Bit Depth
- Gamma
- Gamma Correction
- Color spaces

are first needed.

What is a color Gamut

A color gamut is basically the “Gamut” or complete set of colors that can be captured or shown by a given system. For example, LCD monitors have the ability to represent a subset of colors, printers have a different subset of colors they can represent, Cameras can capture a different set of colors. Each of these devices had its own gamut of colors.
This chart shows the color gamut of the human eye (Visible Spectrum) as well as two other standards, RGB and CMYK. In conversation, you would say something like, “the CMYK gamut does not include all of the colors of the RGB gamut. Some colors are out of gamut.”

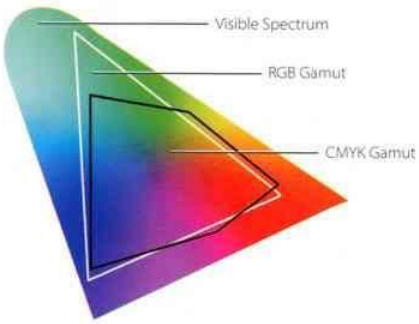

What is Bit Depth

Bit depth quantifies how many unique colors are available in an image's color palette in terms of the number of 0's and 1's, or "bits," which are used to specify each color. This does not mean that the image necessarily uses all of these colors, but that it can instead specify colors with that level of precision. For a grayscale image, the bit depth quantifies how many unique shades are available. Images with higher bit depths can encode more shades or colors since there are more combinations of 0's and 1's available.

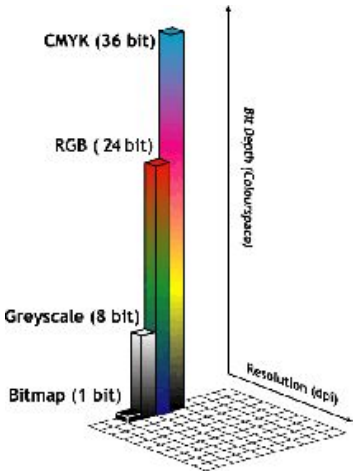

TERMINOLOGY

Every color pixel in a digital image is created through some combination of the three primary colors: red, green, and blue. Each primary color is often referred to as a "color channel" and can have any range of intensity values specified by its bit depth. The bit depth for each primary color is termed the "bits per channel." The "bits per

pixel" (bpp) refers to the sum of the bits in all three color channels and represents the total colors available at each pixel. Confusion arises frequently with color images because it may be unclear whether a posted number refers to the bits per pixel or bits per channel. Using "bpp" as a suffix helps distinguish these two terms.

EXAMPLE

Most color images from digital cameras have 8-bits per channel and so they can use a total of eight 0's and 1's. This allows for 256 different combinations—translating into 256 different intensity values for each primary color. When all three primary colors are combined at each pixel, this allows for as many as 28*3 or 16,777,216 different colors, or "true color." This is referred to as 24 bits per pixel since each pixel is composed of three 8-bit color channels.

ALPHA

Transparency of an image is can also be specified by adding a 4th “color channel” to each pixel. This channel can be used to store how transparent (or the inverse, how opaque) the pixel should be. The alpha channel, if looked at by itself, looks like a grayscale image where the white part is 100 percent opaque and the black is 100 percent transparent. All values in between are varying levels of transparency.

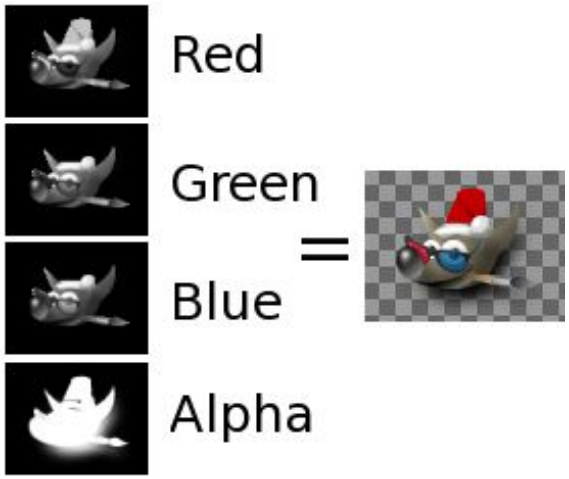

This graphic shows each color channel as an independent grayscale image representing the amount of that color used, as well as alpha. This would be an image that contains 32 bits per pixel (8xRGBA). In the resultant image shown, the checkerboard pattern represents areas of the image that are transparent.

Not every image format can contain 32bit color data. JPEG for example cannot as it was designed to store digital photos in the most efficient fashion and cameras know nothing about alpha.

A 24 bit RGB color representation (8 bits per color channel) can effectively represent all the colors the eye can see. There are applications for higher bit depth such as 10 or 12 bits per color channel. These have not be used much historically as it adds significant overhead to processing. However, they are becoming more and more common. A higher bit depth does not mean more luminance, just more color precision.

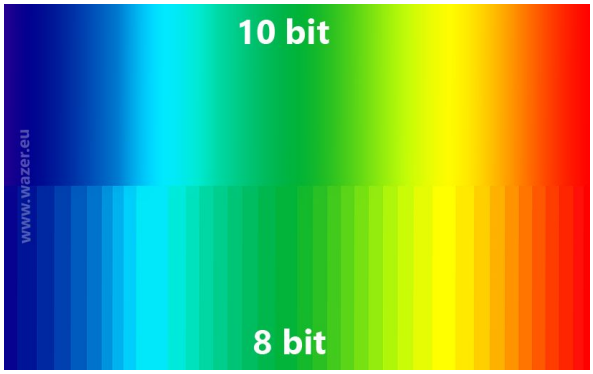

This image shows very clearly how a 10 bit color depth can represent more color gradations but not more colors.

What is Gamma and Gamma Correction?

When we talk about a computer display, an image is made up by a rectangular grid of pixels, and a pixel has 3 sets of values, for Red, Green, Blue, together they specify a color for the pixel, and all the pixels form the image. Each pixel has a brightness level, which is the average of {red, green, blue} values, and this is called its luminance. In order to reproduce an image from capture to display, the luminance needs to be accurately reproduced. Since sensors and displays can have different luminance characteristics, there needs to be a mapping or relationship between a pixel's numerical values and the actual luminance…this relationship is called the Gamma.

This relationship between data and luminance is used in various ways.

Display Gamma

A good example of this is Display Gamma. Display gamma is used to take image data (RGB data) and ensure that it is correctly shown with proper luminance on a display.

This is complicated due to the fact that the human perception of the intensities of color or brightness, does not have a linear relationship, with respect to the values of intensity of light waves in physics. That is, we are more sensitive to some color or brightness ranges than others.

Our eyes do not perceive light the way cameras do. With a digital camera, when twice the number of photons hit the sensor, it receives twice the signal (a "linear" relationship). That's not how our eyes work. Instead, we perceive twice the light as being only a fraction brighter — and increasingly so for higher light intensities (a "nonlinear" relationship).

In addition, the brightness of image display devices, such as CRTs and even LCD screens to some degree, do not always have a linear relation of its input to its output. To compensate for this we use Gamma Correction.

Gamma correction started historically due to the natural behavior of CRT tubes. This graph shows how a camera (linear) and our eyes (nonlinear) perceive luminance
So if the linear data from a camera was displayed on a device with a linear gamma, the image would look wrong to us because we don't perceive luminance linearly.

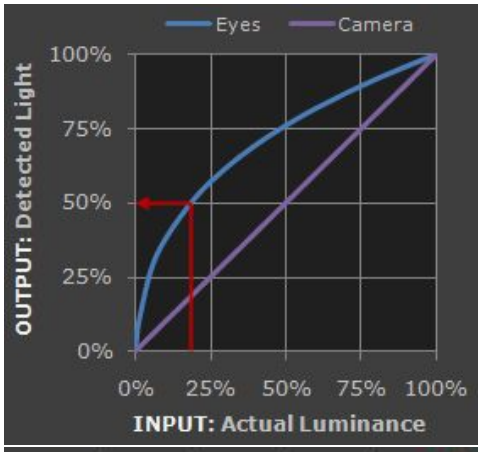

Due to an odd bit of engineering luck, CRT tubes were also nonlinear, and had an almost perfectly opposite response curve compared with that of our eyes. So a linear camera source shown on a CRT was very close to representing what our eyes expected to see.

The intensity of light reproduced by a CRT display is described, technically, as being ‘proportional to the input voltage raised to the power gamma’, or ‘the exponent of the response to power input’. This gamma value, resulting from an effect caused by the electrostatic charge in the tube’s electron gun, is typically around 2.5.

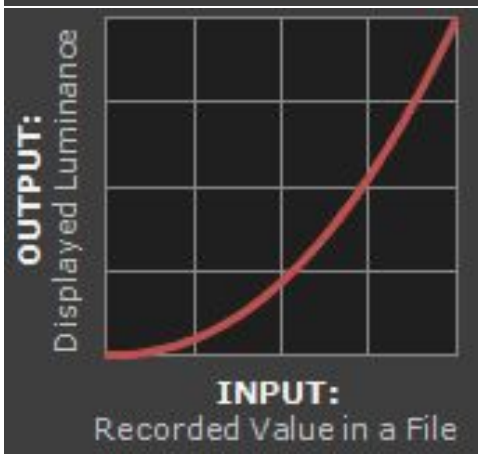

The 2.5 gamma number is derived from the fact that input brightness level of 50% (midway along the gamma curve) is presented as a light intensity of just 18%. The formula for working out the effect of the gamma on brightness values is simply input raised to the power of gamma = output, so 50% ^ 2.5 = 18%.).

The CRT Gamma was a bit too corrective and industry standards dictated that the appropriate Gamma for reproduction is 2.2 To counter this over correction, when television standards were established, an in-camera gamma correction of 0.45 was arrived at (and has been used ever since) to bring the effective result down to 2.2.

Fast forward to flat screen displays and what do we find? Plasmas and LCD displays don’t behave like CRT screens – they are relatively linear devices – but the source material is all being corrected for the old CRT displays so there is a big problem. What we now need is to make the flat panel HDTVs behave as if they were CRT devices …so they do gamma correction to compensate for that as well as slight non linear behaviors they have.

Storage Gamma

Another use for adjusting gamma is to store material more efficiently. The human eye's sensitivity to light is finer at low intensities than high intensities. That's to say, if you make a list of all the intensities you can distinguish, there are more dark ones than light ones. To put it another way, you can tell dark shades of gray apart better than you can with light shades of gray. In particular, if you're using 8 bits to represent your intensity, and you do this in a linear color-space, you'll end up with too many light shades, and not enough dark shades. You get banding in your dark areas, while in your light areas, you're wasting bits on different shades of near-white that the user can't tell apart.

If this is our source image,

If we encode it linearly, it would be represented digitally like this.

We get very little detail in the dark areas where our eyes are the most sensitive and to much in the light areas where our eyes are less sensitive.

If we apply a gamma correction or curve to the data as we encode it, and use a curve that 1/2.2 (the inverse of the 2.2 gamma curve already defined) we get a lumma distribution like this…

This is exactly what the ColorSpace sRGB does. It not only defines the set of colors but also a Gamma Curve to store them in the most efficient way.

What is a color space?

Color space is very generic term. In its purest form, basically a standard model for representing a set of colors in a way that they can be reproduced reliably on various devices.

There are simple colors spaces that define how to store a set of pixels for a still image like RGB, and “Meta” color spaces (like 709) that extend its definition to include things like

- Encoding
- Frame Rate
- Resolution
- Gamma

Encoding Colors….

There are several ways commonly used to represent colors with numbers. These are often referred to as color spaces but in reality are really encoding schemes.

RGB

RGB is a color model that uses the three primary (red, green, blue) additive colors, which can be mixed to make all other colors. RGB builds its model on different colors of light added together, where the mixture of all three colors produces white light, as shown in Figure 1. Digital cameras produce RGB images, and monitors display RGB images.

CMYK

CMYK color space is a color model based on subtracting or absorbing light – the cyan, magenta, yellow and black inks used in most commercial color printing (books, magazines, etc.) Inks absorb colored light, which is why the model is called a subtractive one. CMYK is commonly referred to as process color, and there are many individual color spaces that use the CMYK color model.

Luminance-Chrominance Encoding (YUV, YIQ, YCbCr)

The term YUV refers to a family of color spaces or encoding scheme, all of which encode brightness information separately from color information. YUV is often used as a general term encompassing

- YUV-Analog PAL Encoding,
- YIQ-Analog NTSC Encoding,
- YCbCr-Digital medium (MPEG, FireWire, SDI)

Original B+W TV was broadcast by transmitting an analog signal representing just the luminance of the image. When color was introduced into television, the black and white luma system couldn’t just be thrown away. Chrominance or Color information (designated as U and V) was added separately via a sub-carrier so that a black-and-white receiver would still be able to receive and display a color picture transmission in the receiver’s native black-and-white format.

U and V provide color information and are "color difference" signals of blue minus luma (B-Y) and red minus luma (R-Y). This graphic shows how an image is decomposed as luminance and chrominance.

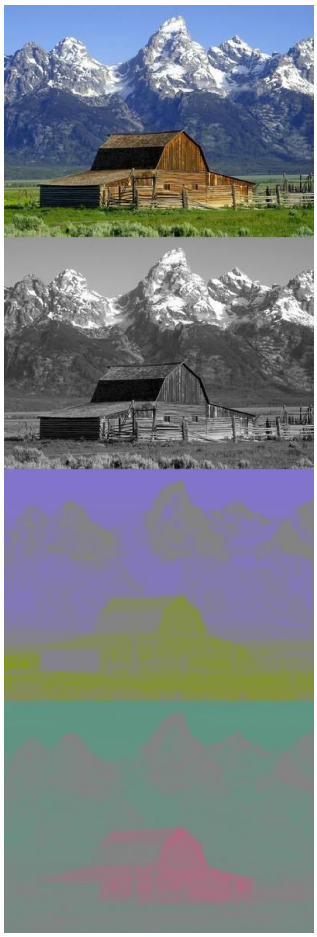

YUV has another advantage that is more relevant today. Luckily, the YUV model is similar to the human eye, which is less sensitive to changes in hue than changes in brightness. You can see that most of the detail of the image to the right is represented in the luminance. As a result, an image can have less chroma information than luma information without sacrificing the perceived quality of the image.

Since the human eye is more responsive to brightness than it is to color, many lossy image compression formats throw away half or more of the samples in the chroma channels to reduce the amount of data to deal with, without severely destroying the image quality.

If the video uses YUV color space, the most important data is in the Y channel. You can throw away a lot of the color information, and the average viewer can't tell that it's gone. The JPEG file format makes use of this color space to throw away unimportant information.

Video engineers have used this principle for years is to toss away a lot of the color information. Basically, they can toss away the color values on every other pixel, and it's not very noticeable. In some cases, they throw away even more color information. This is called Color Subsampling, and it's a big part of a lot of modern HD formats for video.

When looking at color subsampling, you use a ratio to express what the color subsampling is. Most of us are familiar with these numbers: 4:4:4, or 4:2:2, or 4:2:0. These relationships between YUV properties and how they are subsampled. A great easy to understand reference on this is here….

http://blogs.adobe.com/VideoRoad/2010/06/color_subsampling_or_what_is_4.html

Other Common Color spaces

NOTE: Color spaces are usually defined by a subset of the total spectrum in the shape of a triangle. The coordinates of the three points of the triangles are called the Primaries … you will read at times things like color space A has the same primaries as colorspace b.

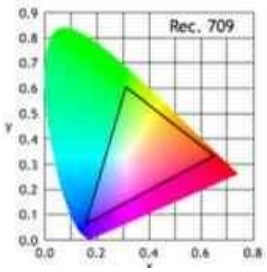

The following is a list of common color spaces and how they originated.

CIE 1931-THE STANDARD

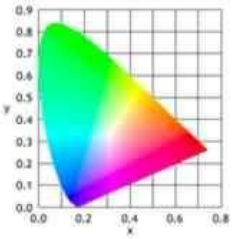

The CIE 1931 chromaticity diagram was invented in 1931 by the International Commission on Illumination (CIE). The graph represents all of the colors that are visible to a human eye.

REC. 709

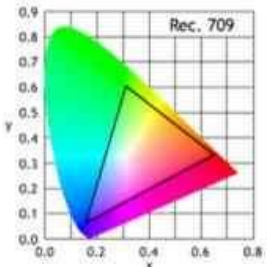

The advent of high-definition required yet another new diagram for color space, which came to be known as Rec. 709. Hollywood films that are reproduced on DVD or Blu-ray discs all use the Rec. 709 gamut. If you notice, the green areas outside the rec 709 triangle look relatively consistent in Hue as does the blue and red areas. Keep in mind that the monitor you are viewing it on is basically a rec709 monitor and even if the graphic was accurate, the monitor would not be able to display those shade and would clip them at the primary values.

SRGB

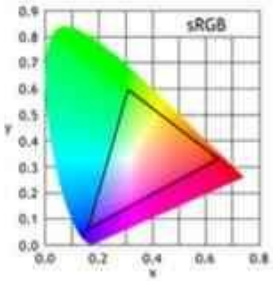

In 1996, the sRGB color space was developed as a standard for display technologies at that time. The black triangle in the graph below shows the sRGB color space. sRGB is very similar to older monitor spaces and, in fact, it's common for unmanaged computers to assume that an image is in the monitor color space. This makes sRGB a good choice to send to unknown users. sRGB uses the ITU-R BT.709 primaries, the same as are used in studio monitors and HDTV,[1] and a 2.2 transfer function (gamma curve) typical of CRTs. This specification allowed sRGB to be directly displayed on typical CRT monitors of the time, a factor which greatly aided its acceptance.

ADOBE RGB

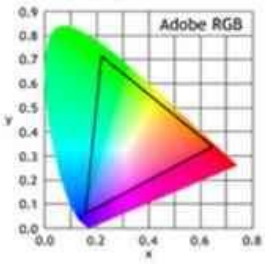

For creative artists, Adobe created the Adobe RGB standard in 1998. While the blue and red corners of the triangle are identical to sRGB, this color space accommodates more vibrant

greens. While artists do their work in Adobe RGB, they must use expensive and costly wide-gamut displays to accurately view their work.

DCI-P3

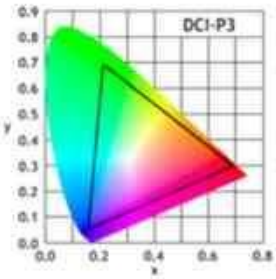

DCI-P3 refers to the color space used in digital movie theaters, and encompasses much more than sRGB. Hollywood films are provided to theatre in DCI-P3 in order to match the gamut of the digital projection equipment that is used. Currently, DCI-P3 content is limited to digital movie theaters and is not available to consumers.

NTSC

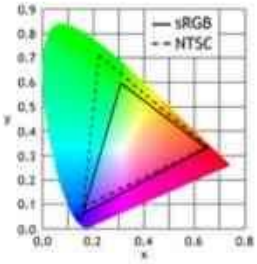

The NTSC (dashed line) color space encompasses a wider gamut than sRGB (solid line). The NTSC color space was created in 1953 at the advent of color television.

The gamut of a display is commonly listed as a percent of NTSC. sRGB is approximately 72% of the NTSC color space. Many laptop displays are listed as 40% to 60% of NTSC. Within the last few years, desktop computer monitors have been increasingly listed as 72% of NTSC or greater.
Though not always the case, in order for a display to be eligible for the Technicolor Color Certified program it must be have a native gamut of at least 72% of NTSC.

As you can see, color spaces are defined for different applications, some relate to how data is transferred from system to system (like Rec 709) and others for Storage of image data (Like sRGB)

File formats

Most image file formats like Tiff, Jpeg, png etc can store images using different color spaces. A majority of the image formats assume an sRGB color space. The details of the color space used in a stored image file is usually stored as metadata with the image. This way when a

computer program opens the file it can tell how to process it. In the case of an image file stored with sRGB, the application would have to compensate for the fact that the data is stored using a 2.2 gamma curve.

The colorspace used is usually determined based on the target application, printing, computer display etc. Each color space has different tradeoffs, here is a good example:

> *JPEG files are 8-bit color, which means you get 256 red, 256 blue, and 256 green levels whether you use Adobe '98 or sRGB. However, Adobe '98 is* broader*, meaning it spreads its bits across a broader range of colors by making the jump between each color more coarse. You get finer increments of skin tone by using sRGB, for example. But the pure cyan in HP's original logo can only be accurately represented by Adobe '98, whereas in sRGB, you'd have to pick a substitute color*

How does all this relate to what we do?

In most of our products, rendering is performed in the GPU using OpenGL. For historical reasons, OpenGL performs its rendering operations in 8 bit Linear RGBA. In a linear color-space, the relationship between the numbers you store and the intensities they represent is linear. Practically, this means that if you double the number, you double the intensity (the lightness of the gray). If you want to add two intensities together (because you're computing an intensity based on the contributions of two light sources, or because you're adding a transparent object on top of an opaque object), you can do this by just adding the two numbers together. When doing any kind of 2D blending or 3D shading, or almost any image processing, then you want your intensities in a linear color-space, so you can just add, subtract, multiply, and divide numbers to have the same effect on the intensities. THis makes processing very efficient. (*NOTE: OpenGL 3 or later can operate in native SRGBA mode*).

So if we look at a typical day in the life of a graphics rendering application….

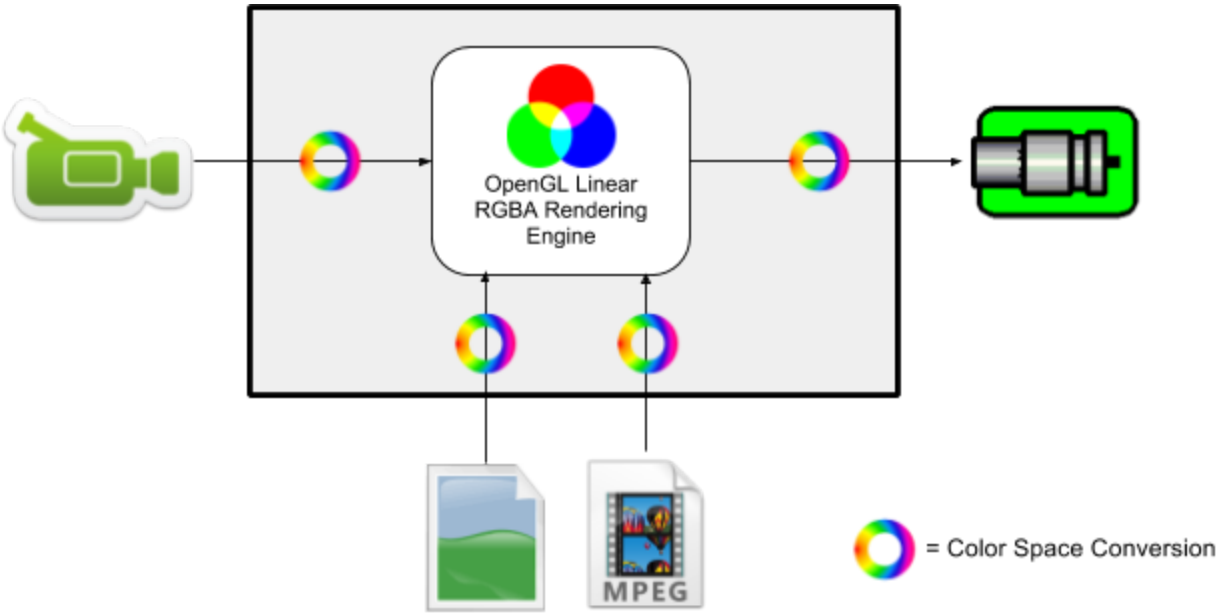

image data can come from various sources, a live camera brought in with 709 colorspace, an image file possibly in sRGB, a movie stored as YCbCr etc. Each one of these sources must be converted to linear RGBA to be operated on by the render engine. These colorspace conversions ( ) can happen in different places, the live video might be converted in the frame buffer card, the images in a software library or in OpenGL. Once converted to Linear RGBA, all this image data can be composited, textured etc along with our rendered graphics, The result, is sent out through the Frame buffer card, with yet another colorspace conversion, as 709.

What about 4K and what is HDR

4K TVs have been out for a few years. The early ones were basically HD sets with higher resolution. The display technology was basically the same and supported the standard Rec709 color space. Some sets sold in 2013 and 14 were limited to 25fps.

Most modern 4K camera sensors can capture far wider color gamut and far more dynamic range than today's monitors can display so most of that information is lost since it ends up being displayed on a relatively poor output device (rec709 displays).

The real advantage of 4K technology is only starting to be realized. THere are a new generation of sets already on the market that ultimately will be branded with the nomenclature Ultra HD Premium. The Premium means a few very important things…

- **4K / 60p** -High speed display in 60 frames per second of 4K video (3,840 x 2,160 pixels - Ultra HD), which has four times the resolution of Full Hi-vision, for highly detailed videos with extremely smooth movement.
- **10-bit Color Depth** - Previous Blu-ray Discs displayed the colour signals (Y, Cb, Cr) in 8-bit gradation each (256 gradations). By expanding this to 10-bit gradation each (1,024 gradations), even minute signals can be faithfully reproduced to realize richly textured video.
- **High Dynamic Range -**A technology that drastically expands the brightness peak from the previous 100 nit to 1,000-10,000 nit, marking a significant leap in the dynamic range of the picture. A nit is approx. the brightness of a candle. Bright light sources (e.g. lights or rays of the sun) and reflected light (from metal or water) that up to now were difficult to display can now be shown in rich textures.
- **BT/REC2020-** wide colour gamut Compliant with the ITU-R BT.2020 wide colour gamut signal formulated for 4K/8K broadcasting. Enables vividly rich coloration not previously possible on Blu-ray discs (BT.709 standard).
- **HEVC (H.265) / 100Mbps** -Support for the highest 100Mbps video signal using the latest high-efficiency video compression technology. Compression efficiency and high bit rate far beyond previous Blu-ray discs (MPEG-4/AVC (H.264), maximum 40Mbps) enabling outstanding playback of high quality video with 4K/60p/10bit, High Dynamic Range, BT.2020, etc.
- **HDMI 2.0a** - to do all this it must support HDMI 2.0a interface standard.

A TV that can achieve these characteristics ultimately can licence and display the UHD Premium logo. Some newer sets already can do this but are not yet licensed to use the logo.

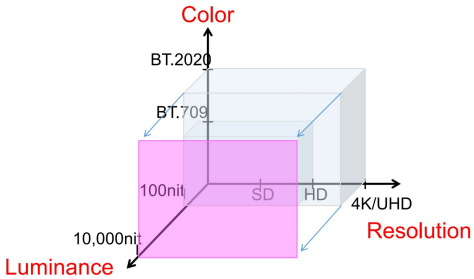

A proper 4K UHD Premium TV will provide 3 key advantages over standard HD...Color, Resolution and Dynamic range.

All this new capability will more closely match what the Cameras can already capture and allow the wider color and higher dynamic range to be brought to the home.

BR/Req 2020- A new Color Space

Most 4k sets have been basically showing Rec709 imagery as most of the source material is created that

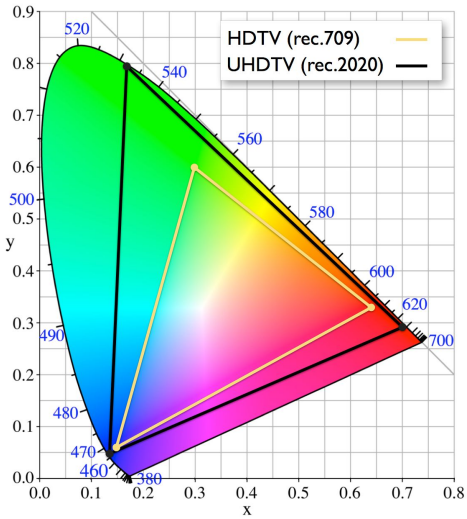

UHDTV vs HDTV Color Gamuts

way. However, just as 709 was developed for HD tv, 2020 is intended for 4K and 8K content and devices.

Basically it defines

- **Resolution/Rate** - 3840 × 2160 ("4K") and 7680 × 4320 ("8K") operating at up to 120 hz
- **Color Gamut**-A much wider color gamut. See the comparison between Rec709 and 2020 to the right.
- **Encoding** -RGB and YCbCr signal formats with 4:4:4, 4:2:2, and 4:2:0 chroma subsampling. Rec. 2020 also defines a linear encoded version of YCbCr called YcCbcCrc.[1]YcCbcCrc may be used when the top priority is the most accurate retention of luminance information
- **Color Depth**- 10 or 12 bits. These are needed as the gamut is more expansive and 8 bits would cause too much banding.
- **Gamma** -10 Bit uses the same gamma curves as Rec709. 12 bits uses a slightly modified version to make better use of the extra bit depth.

HDR-10

There is also a new Media Profile defined by the Consumer Electronics association that states that and HDR set is a set that supports (at a min) something called the HDR-10 Media Profile.

THe HDR10 Media Profile is defined as...

- EOTF: SMPTE ST 2084
- Color Sub-sampling: 4:2:0 (for compressed video sources)
- Bit Depth: 10 bit
- Color Space: ITU-R BT.2020
- Metadata: SMPTE ST 2086, MaxFALL, MaxCLL

DOLBY VISION VS HDR 10 – WHAT'S THE DIFFERENCE?

The main specification differences are that masters of movies are done at 12-bit, rather than HDR 10's 10 bits. Also, peak brightness can go – in theory, at least – right up to 10,000 lumens. In reality, most Dolby Vision masters seem to be targeting 4,000 nits – which remains a very big step up from the 1,000 nits that HDR 10 masters work to.

But most significant is how Dolby Vision uses frame-by-frame metadata to manage HDR performance. This helps to deliver the best results as it adapts the source material to the performance of your TV.

This is why Dolby Vision requires extra hardware, and why HDR 10 / Ultra HD Premium will be more prevalent, but the Dolby argument is this hardware will deliver better picture quality.

What is HDR

It stands for High Dynamic Range and basical is a general term that can be applied to anything. It basically means more and better. HDR for photography is not the same as for video…

Photo HDR: Combining multiple images with different exposures to create a single image that mimics a greater dynamic range.

TV HDR: Expanding the TV's contrast ratio and color palette to offer a more realistic, natural image than what's possible with today's HDTVs. An HDR photo isn't "high dynamic range" in this sense. The image doesn't have the dynamic range possible in true HDR. It's still a standard dynamic range image, it just has some additional info in it due to the additional exposures. HDR is the way that most of us see the world everyday. Our eyes are high dynamic range and wide color gamut sensitive equipment. On a bright sunny day, they can see cloud highlight details while still being able to see into the shadows. In dark environments, we can start to see detail in extremely low to practically no light. Some cameras are at least getting closer to being able to detect some of the same things on those sunny days. But they have also gotten better in the dark, not as good as our eyes, but certainly better as the sensor technology improves. On the other hand, the television sets or other types of displays that we typically have used to watch content has a much lower dynamic range than our eyes or even the cameras that are being used to capture the content. So TV viewing, up to this point, has been what we now call, Standard Dynamic Range (SDR) .

Overall, the dynamic range of a scene can be described as the ratio of the maximum light intensity to the minimum light intensity [1]. In digital cameras, the most commonly used unit for measuring dynamic range is in terms of f-stop, which describes total light range by powers of 2. The current ad hoc use of the term f-stop, refers to the following dynamic ranges:

- 10 f-stops = a difference of 210 = 1024: 1 contrast ratio.
- 14 f-stops = a difference of 214 = 16,384: 1 contrast ratio.
- 16 f-stops = a difference of 216 = 65,536: 1 contrast ratio.

100,000:1 is normally regarded as approximately the range that the eye can see in a scene with no adaptation.

- 20 f-stops = a difference of 220 = 1,048,576: 1 contrast ratio.

1,000,000:1 is normally regarded as approximately the range that the eye can see in a scene with minimal (no noticeable) adaptation.
In the categorization of dynamic ranges, the following definitions are typical and will be used in the present document:

- Standard Dynamic Range (SDR) is ≤ 10 f-stops
- Enhanced Dynamic Range (EDR) is > 10 f-stops and ≤ 16 f-stops
- High Dynamic Range (HDR) is > 16 f-stops

But now with these higher dynamic range displays, we're able to really start to give people more picture details in order to provide a more immersive experience. As I mentioned, getting better color saturation goes along with this whole wider color gamut piece that is part of HDR. So HDR is not just about contrast, or about resolution; or even about color; it's about being able to combine all three to represent images more accurately on consumer displays.

By leveraging the higher NIT values and the wider color space, images can have far brighter whites and more accurate color. In addition, using new gamma curves on the cameras, more detail in the darker regions can be leveraged.

For reference, these images show what some NIT values of the real world are…

A standard 4k or HD TV can display about 100 NITs so images will never really resemble what the eye can see. THe UHDP TVs use different display technology and can display up to about 1000 nits. Still no where near real life but significantly closer. The new UHDP specifications are allowing provisions for up to 10000 NITS in anticipation of newer sets. In the same way, not all UHDP sets can display the full 2020 gamut of colors. Many will advertize what percentage they can cover.

So having more NIT does not mean that everything will be brighter, in reality, the normal intensities of images should remain about the same as with a HD set, but regions on the image that should have been brighter than 100 nits can now be reproduced more accurately.

So HDR would improve this type of image where in real life, the sun and sun spots on the ground would be much brighter.

In addition to simply being able to make bright things brighter,, this new dynamic range allows new gamma curves, both at camera capture and playback, to help add more information, especially in dark areas. This is a shot from a Sony Camera And Samsung display. This first image is from a

standard dynamic range capture and display.

This is the same image, captured with a camera set to capture and store HDR content and displayed on a Samsung monitor adhering to the UHDP standards. Notice the apparent brighter sun, better colors and more depth of detail in the darks of the water. Keep in mind, you are viewing this on a Rec709 compliant display. THis image is only a representative image.

The image below is another great example of what can be achieved using HDR in the production process (capture to display)

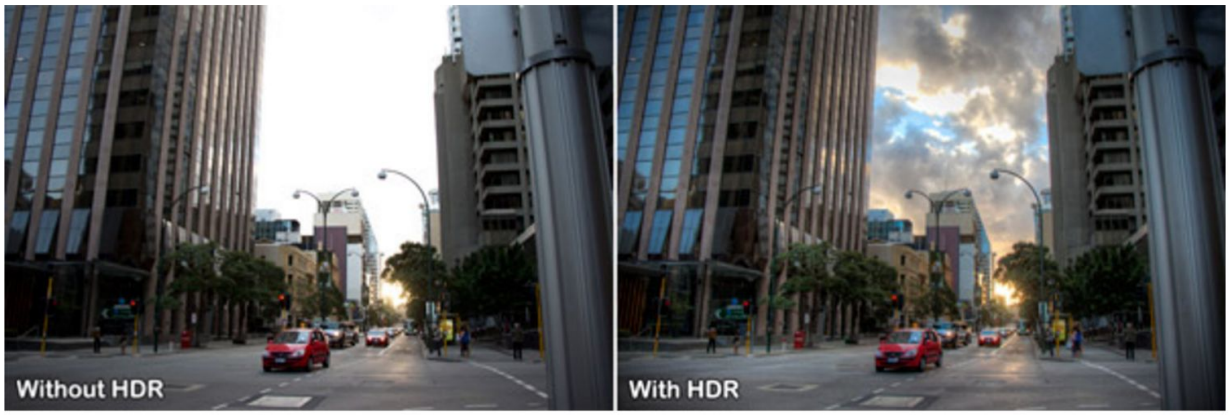

This is a great article on real world use of HDR in film making.
https://pro.sony.com/bbsccms/assets/files/cat/mondisp/articles/HDR_X300_explained.pdf

How is HDR video encoded.

There is nothing in the Rec2020 specification that allows for Dynamic range to be encoded. We still have only 10 or 12 bits of color channel to represent as much as 10x the dynamic range. In order to accomplish this, new gamma curve conversions are being used. More accurately, new Electro Optical Transfer Function (EOTF) and Inverse-EOTF (OETF) are being used.

EOTF and OETF

TVs contain lookup tables (LUTs) that describe an electro-optical transfer function (EOTF) which defines at what electrical input level the display should be illuminated, and how strongly. The EOTFs for Rec. 601 (SD) and Rec. 709 (HD) represent a gamma function with a gamma of 2.2. This describes the characteristics of phosphor in legacy cathode rate tube (CRT) TVs and is therefore known as standard gamma. On the input (camera) side, an inverse function (opto-electrical transfer function, or OETF) is used on the data after reading the camera sensor. This provides a linear relationship between the electrical signal and the luminance.

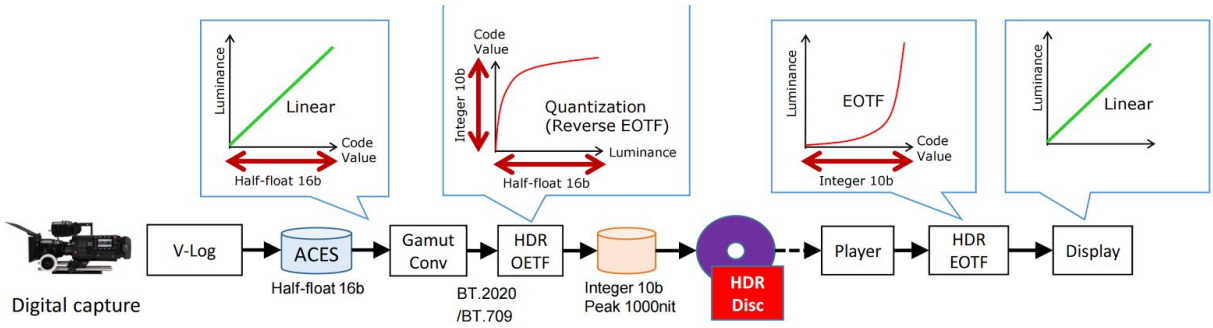

In the “modern” world we are no longer trying to model CRT display behavior and a new SMPTE standard has been generated to define how to fit 10000 nits of dynamic range into a 10 bit number. THis is SMPTE standard 2084. No TV set yet will do 10000 nit but this is intended to be a forward looking spec. This “Gamma” curve is designed to use the available bits in a way that most accurately represents human visual response, not crt behavior. THis is often call the PQ curve (Perceptual Quantizer).

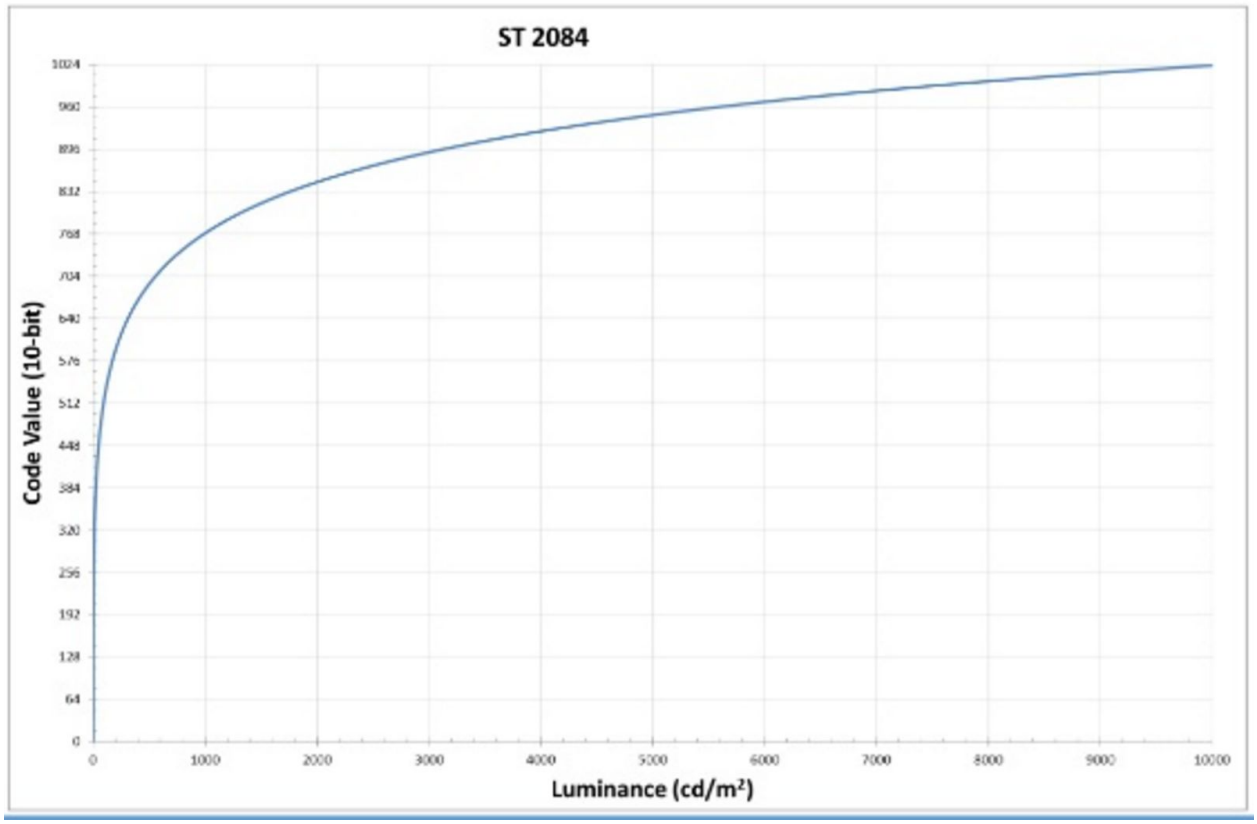

In addition to new curves, there are also approaches where metadata is also added to the signal with dynamic range information. Dolby has been a pioneer on PQ and has product (encoders and decoders of sorts) that will be used for distribution and integrated into televisions.

As 10000 nits is aspirational, other curves are being experimented with by camera manufactures that more closely match the sensor behavior and capabilities of the cameras to use the available bits more effectively, If a camera can only capture 1000 nits, there would be lots of detail lost using a EOTF designed for 10000 nits...or the behavior of a sensor may require a slightly different curve.

Sony has EOTFs called slog that do just that…

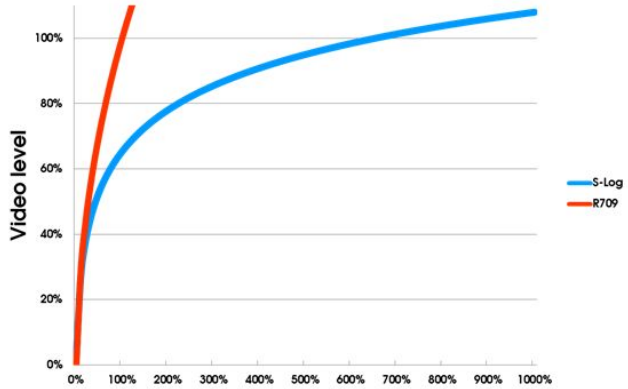

This shows a sony S-Log curve relative to rec 709 curve. Of course, if a video is encoded using S-LOG, the display device will need to know about this and use the reverse EOTF or OETF to properly display it. If you look closely, you will see that S-LOG models Rec 709 up to about 30% and then starts to deviate from it. If a S-Log signal is looked at on an TV that does not

understand S-LOG, it will look rather washed out. THis is a real problem as there are millions of 4k TV already in the market.

This means two things…

1-have two production workflows...SDR and HDR. Some cameras even output both signals.
2-use a format that works better on and SDR TV.

The jury is still out on all this however, there is a standard called Hybrid Log Gamma. This is a curve that is a Hybrid of Rec709 and then deviates to a HDR curve. Similar to SLOG except it deviates from 709 at closer to 60% instead of 30%. HLG is a clever compromise that accommodates SDR as well as HDR displays but, unfortunately, compromised both SDR and HDR viewing. SDR is not as good and viewers would be used to and HDR is limited in its effectiveness. The Hybrid Log Gamma name refers to the fact that the OETF is a hybrid that applies a conventional gamma curve for low-light signals and a logarithmic curve for the high tones.

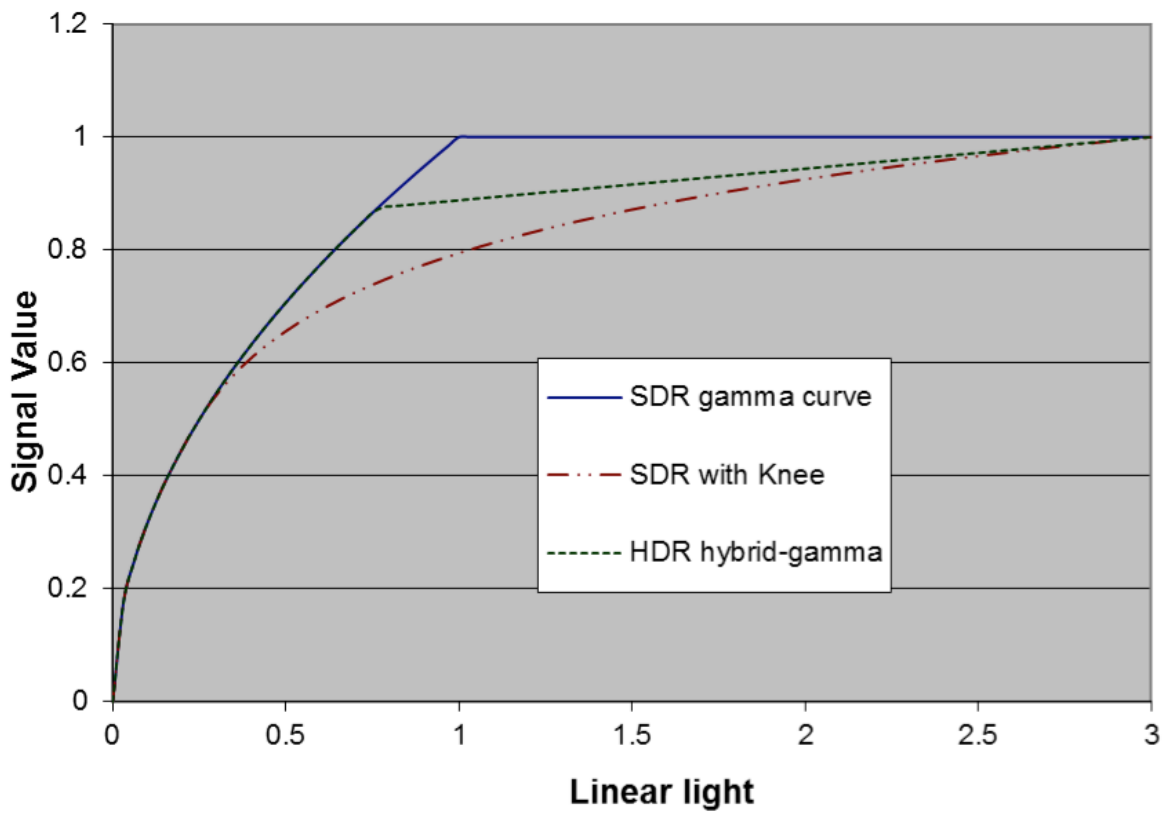

While not perfect, HLG is gaining acceptance as it allows for one production workflow and one distribution.

As an illustrative example, Figure B.1 compares an 8-bit SDR system capable of representing 0.1 to 100 cd/m2 with a BT.709 style transfer function (green curve) to a 10-bit HDR system capable of representing 0.005 to 10,000 cd/m2 with another transfer function (SMPTE ST 2084). The plots in this figure are approximate plots. They do not capture the exact form of the curve and are shown for illustration purposes only. In the figure, the integer code levels are along the horizontal axis and linear light values (scaled to log10) are along the vertical axis. This illustrative mapping includes traditional code level range proportions to accommodate both foot-room ("negative" samples below the [0.0, 1.0] real-valued range) and head-room (samples above real-value 1.0). Due to design properties, the 10-bit HDR transfer function shown here assigns approximately twice as many code levels [119 to 509] as the traditional 8-bit SDR transfer function assigns [16 to 235] in the SDR range, while providing a similar number of new code levels [510 to 940] to extend brightness. New code levels [64 to 118] are assigned for darker intensities below 0.01 cd/m2.

In a sense, the 10-bit HDR system illustrated here distributes the extra 2-bits over traditional consumer 8-bit "SDR" video by assigning approximately 1 extra bit of precision within the traditional SDR intensity range, while applying the other extra bit to extend the curve to intensities greater than 100 cd/m2. For comparison, the 10-bit SDR transfer function is also plotted (red dash curve).

Figure B.1. Mapping of linear light values for SDR and HDR representations
http://dci-forum.com/index.php?topic=362.0

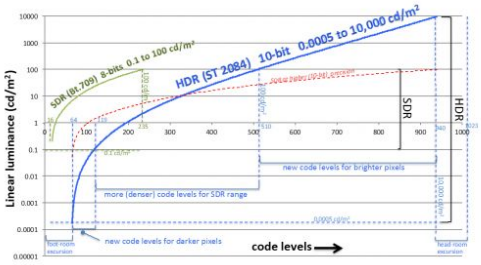

Figure B.1.

Converting Color spaces

Converting legacy material in HDTV/SDTV color spaces to the much wider UHDTV color space also requires some thought. By merely mapping the legacy image to the limits of the wider gamut (upscaling), the colors will look oversaturated. If the legacy images are mapped to their legacy equivalent in the UHDTV color space, then they will look about the same as on a legacy workflow, but then the colors in the legacy material will look muted when mixed in with the much more brilliant UHDTV compliant images.

Three methods have been identified:

1) colorimetric conversion to the UHDTV color space - content is mapped using a simple matrix, the HDTV gamut remains the same but is now placed within the UHDTV color space container (reversible if metadata was provided so that the HDTV color space is known)

2) conversion with gamut extension to the UHDTV color space – the content is automatically stretched along the lines of constant hue to increase the brightness and saturation without changing the color (reversible if the stretching algorithm is known and invertible).

3) artistic conversion to the UHDTV color space – the content is modified by a colorist to best match native UHDTV content (reversible only if the tool's output metadata which describes the colorist's actions)

Where will all this new content come from

As with HD, sports and live broadcasts will lead the way as far as content creation. 4K OTT broadcasts are already starting through Netflix, Amazon and others. H.265 supports HDR.

Ultra HD Blue Ray players are about to be released, these will feature support for Rec2020 and HDR content.

https://www.smpte.org/sites/default/files/Study%20Group%20On%20High-Dynamic-Range-HDR-Ecosystem.pdf

4K Resolution and SDI Data Rates.

For reference, the information below show how 4K, ultra HD etc relate to each other as well as how 4K is currently send over SDI.

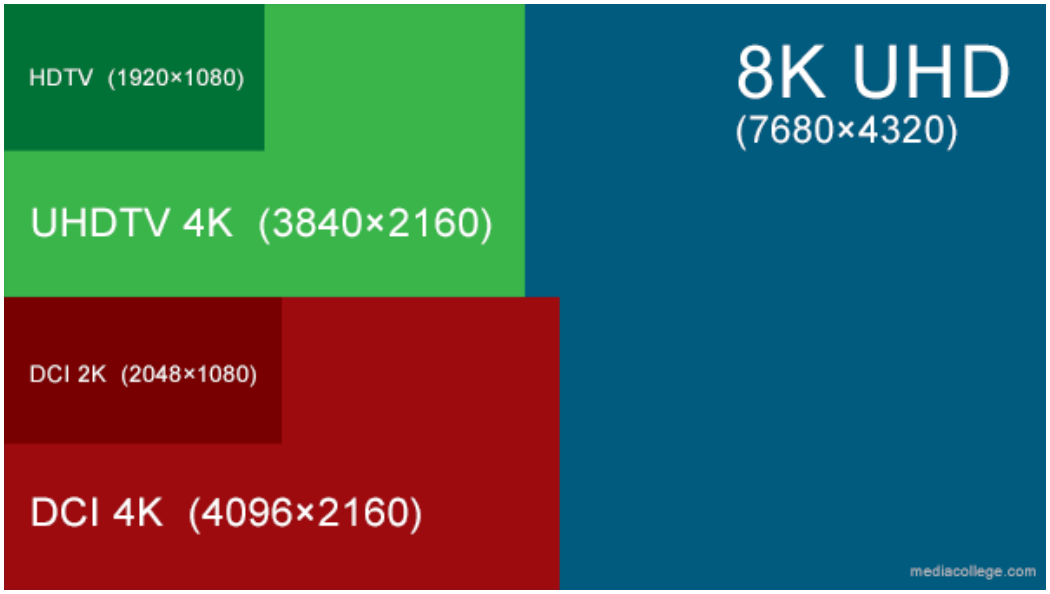

Data Rates

HD SDI	Up to 1.5 Gb/s. HD 4:2:2 
3G 1080p	3 Gb/s-going from interlaced to progressive doubles the data. ● Initially there was a standard called Dual Link which just used 2 HD SDI signals  ● Then 3G standard came out-3Gb/s -this does Dual Link (1080p on one wire)

<table>
<tr>
<td></td>
<td>


Note the Level A and B designation...this is important... The Level A format is 1920 x 1080 <u>progressive</u> 50 or 60 Hz while the Level B format is a multiplexing scheme where two streams of 1920 x 1080 <u>interlaced</u> 50 / 60 Hz video is transported over a 3G SDI link. Level A is overwhelmingly predominant. Sony uses level B. Level B is a bit more complex than described. For a detailed description go here...https://tech.ebu.ch/docs/events/webinar036-3g-sdi/presentations/ebu_3g-sdi_webinar036.pdf . The important thing to remember is there are different 3G standards.

● Dual Stream 3G is a specific variant of the 3G signal which combines two completely separate 4:2:2 image streams into a single 3G signal. This can be used to minimize confusion in stereoscopic production by keeping left and right eye signals together
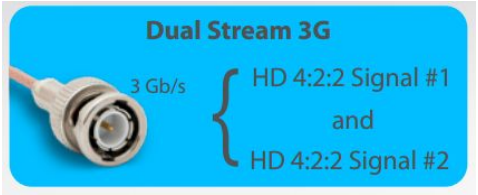

For a great review of 3G go here
http://www.aja.com/pdf/qa/Dual_Link_and_3G_Overview.pdf
</td>
</tr>
<tr>
<td>4kP60</td>
<td>
Technically fill 4Kp60 needs 12Gb/s. Currently the only way to get that is with 4 x 3G signals (3Gb/s x4). Since 4k is based on 3g...go back and read the 3G description you probably just glossed over.

12 Gig SDI is starting to emerge now.
</td>
</tr>
</table>