

Ultra HD Blu-ray Format Video Characteristics

By Michael D. Smith

Abstract

The Blu-ray Disc Association recently finished a new format extension called Ultra HD Blu-ray. The new format supports next-generation video technologies and several optional enhancements. This paper provides an in-depth description of the video technologies offered by Ultra HD Blu-ray and an overview of new optional features. The Ultra HD Blu-ray format allows improvement to the video representation by increasing the pixel code value quantization to 10 bits, enabling representation of high dynamic range imagery using SMPTE ST 2084 electrooptical transfer function, enabling representation of wide color gamut imagery using International Telecommunications Union Radiocommunication Sector (ITU-R) BT.2020 color primaries, increasing spatial resolution to 3840×2160 , and allowing increased frame rates up to 60p. Carriage of mastering display color volume metadata described by SMPTE ST 2086 and new content metadata parameters called maximum content light level (MaxCLL) and maximum frame average light level (MaxFALL) are also enabled by the new format. Additionally, the Ultra HD Blu-ray format improves compression efficiency by using the new HEVC/H.265 video compression standard.

Keywords

Distribution technology, HDR, HDR10, UHD, ultra HD Blu-ray disc

Introduction

The Blu-ray Disc Association (BDA) completed its work on the Ultra HD Blu-ray format specification in May 2015¹ and began licensing in August 2015.² The Ultra HD Blu-ray format specification represents the work of many companies over a multiyear period. Much of the time spent developing the new format was devoted to finding a consensus on the requirements of the new format that met the objectives of the various member companies. After the requirements for the format extension were finalized, the technical

expert groups worked efficiently to create a specification that met the agreed-upon requirements. The BDA also created authoring guidelines and compliance testing procedures designed to help ensure a consistent high-quality experience for Ultra HD Blu-ray users. This paper describes features of the Ultra HD Blu-ray format extension with a focus on the video-related characteristics.

Disclaimer: The author is knowledgeable about the Ultra HD Blu-ray format specification, has participated in its development, and believes that the information in this paper is

accurate. The BDA did not prepare this paper; therefore, it is recommended that readers refer to the specification³ and whitepaper⁴ released by the BDA for official description of the format.

This paper provides an in-depth description of the video technologies offered by Ultra HD Blu-ray and an overview of new optional features.

Mandatory, Optional, and Out-of-Scope Features

The Ultra HD Blu-ray format extension specification^{3,4} defines what can be authored to an Ultra HD Blu-ray disc and how an Ultra HD Blu-ray player plays back a disc. Ultra HD Blu-ray features can be segmented into the following categories:

1. Mandatory for discs and Mandatory for players
2. Optional for discs and Mandatory for players
3. Optional for discs and Optional for players
4. Mandatory for discs
5. Mandatory for players
6. Out of scope of the Ultra HD Blu-ray format

The features in the category “Mandatory for discs and Mandatory for players” are often referred to as the “mandatory features” of the format and will generally be the most commonly used features since they will be available on every disc and every player. The use of features in the “Optional for discs and Mandatory for players” and “Optional for discs and Optional for players” categories will generally be determined by the marketplace, since optional features for discs and players allow product differentiation. Features in the “Mandatory for discs” or “Mandatory for players” categories are usually related to compatibility.

Features that are out of scope of the format are related to everything else that may be available with a disc or a

player but are not related to a Blu-ray format itself. For example, many Blu-ray players available in the marketplace have provided other features such as over-the-top video streaming service applications, display of Joint Photographic Experts Group (JPEG) photos from storage devices, DVD-Video disc playback, Audio-CD disc playback, and multimedia file playback via removable memory or network sources that are not defined in the Blu-ray format. Other examples of out-of-scope features that have been associated with Blu-ray discs are the bundling of DVD-Video discs, Digital-Copy discs/codes, or UltraViolet redemption codes within a Blu-ray disc package that are independent from the Blu-ray format.

Mandatory for Discs and Mandatory for Players

Mandatory video characteristics describe the properties of the video signal that must be authored to a disc and presented by a player. Several combinations of properties of the video signal are allowed by the format, and the content author has the freedom to choose the most appropriate allowed combinations of properties for a particular content. The content author should select the properties that represent the native format of the video essence or the video essence should be converted into an allowed video essence format before disc authoring. A Legacy SDRⁱ high definition (HD)-resolution format, which is not a focus of this paper, using AVC/H.264 compression, is also one of the allowed formats on Ultra HD Blu-ray, and its use is expected mainly to allow for repurposing existing HD-resolution bonus material that has already been encoded for release on Legacy Blu-rayⁱⁱ format. The description of the video properties in the rest of this paper is focused on the new standard dynamic range (SDR) and high dynamic range (HDR) video formats that represent next-generation consumer video formats.

Video Resolution

The video signal must use progressive frame sampling with a resolution of either 1920×1080 or 3840×2160 . Interlaced video field sampling is not permitted. The picture aspect ratio must be 16:9. If the content author wishes to preserve the original aspect ratio of the content that is different from 16:9, then letterbox or pillarbox mattes are typically used to center the original aspect ratio image inside the 16:9 container surrounded by black pixels that fill the remainder of the 1920×1080 or 3840×2160 pixel array.

ⁱ“Legacy SDR” is a term used in the paper to describe a video format that is allowed on the Ultra HD Blu-ray format but does not represent a next-generation video format. Legacy SDR is also allowed on Legacy Blu-ray format.

ⁱⁱ“Legacy Blu-ray” is a term used in this paper to describe the original Blu-ray format that was released in 2006 and has been updated with additional profiles since its original release. Legacy Blu-ray has been mainly associated with HD-resolution video, although SD resolution was also allowed on Legacy Blu-ray.

Video Frame Rates

The video signal must use one of the following progressive frame rates: 23.976p, 24.0p, 25.0p, 50.0p, 59.94p, or 60.0p. The 25.0p and 50.0p frame rates must only be decoded by a player designed to interface with 50 Hz television systems. The format supports both “fractional” frame rates, that is, 23.976p and 59.94p, and “integer” frame rates, that is, 24.0p and 60.0p, so that the format will be frame rate compatible as the industry considers the impact of moving away from the usage of fractional frame rates that were associated with the old National Television System Committee (NTSC) analog broadcasting system and have been widely used for standard-definition (SD) and HD content production and distribution over the last several decades.

Video Bit Depth

The video signal samples must use code value quantization with 10 bits of precision.

Video Dynamic Range

Both SDR and HDR video signals are allowed in the format. SDR video signals are typically created for use with the International Telecommunications Union Radiocommunication Sector (ITU-R) BT.1886 electrooptical transfer function (EOTF).⁵ HDR video signals must be created for use with the SMPTE ST 2084 EOTF.⁶

An EOTF defines the input/output relationship between the nonlinear (electrical) video signal that is input to a display and the linear (optical) light values that are produced by a display. Cathode ray tube (CRT) display technology was used to display video signals to viewers since the introduction of video itself in the middle of the 20th century until about the last decade, during which CRT technology was largely phased out. In 2015, CRT technology was no longer used to build displays. The BT.1886 EOTF based on a gamma (power function exponent) value of 2.4 is an approximation of the CRT EOTF and continues to be used for the display of SDR video signals using modern display technology such as liquid crystal display and organic light-emitting diode. This ensures that the video content that was originally prepared for viewing on a CRT will look correct on a modern display operating in a similar viewing environment, although the actual display technology is radically different.

A typical reference viewing environment for SDR video signals conforms to Recommendation ITU-R BT.2035,⁷ which describes the configuration of the room where video mastering takes place and the configuration of the reference display. A reference display configured according to BT.2035 is set up such that the reference white level, typically corresponding to 8-bit code value 235 (10-bit code value 940) in Y’CbCr video format or 8-bit code value 255 (10-bit code value 1023) in full-range R’G’B’ video, is configured to output a luminance equal to 100 cd/sq. m (often called “nits”) and the reference black level, typically

TABLE 1. Code value examples of different formats.

Target Displayed Luminance (nits)	Codevalue to Achieve Target Displayed Luminance					
	100 nits BT.1886	100 nits BT.1886	1000 nits BT.1886	10000 nits BT.1886	10000 nits ST 2084	10000 nits ST 2084
	8 bit	10 bit	10 bit	10 bit	10 bit	12 bit
0.1	28	113	83	71	119	474
0.5	40	160	101	78	167	668
1	48	193	113	83	195	781
5	79	315	160	101	281	1124
10	100	400	193	113	327	1306
50	180	720	315	160	450	1799
100	235	940	400	193	509	2036
500	N/A	N/A	720	315	657	2627
1000	N/A	N/A	940	400	723	2890
5000	N/A	N/A	N/A	720	876	3503
10000	N/A	N/A	N/A	940	940	3760

corresponding to 8-bit code value 16 (10-bit code value 64) in Y'CbCr format or 0 code value in full-range R'G'B' video format, is configured to output a luminance equal to 0.01 nits or lower. Major Hollywood studios typically master SDR video under conditions meeting the BT.2035 specification. SDR video content produced by smaller production companies may not always reflect mastering under such well-controlled conditions.

The video signal is represented digitally by quantizing each video sample with some precision (10 bits in the case of Ultra HD Blu-ray) to a digital code value, and each code value represents a particular target displayed luminance. Each 10-bit code value has different nonlinear relationships to the previous and next code values due to the EOTF. For example, with the SDR format (100 nits BT.1886 EOTF) the 10-bit code value 400 corresponds to $(100 \text{ nits} * ((400 - 64)/876)^{2.4}) = 10.0278 \text{ nits}$, the previous code value 399 corresponds to $(100 \text{ nits} * ((399 - 64)/876)^{2.4}) = 9.9563 \text{ nits}$, and the next code value 401 corresponds to $(100 \text{ nits} * ((401 - 64)/876)^{2.4}) = 10.0995 \text{ nits}$.ⁱⁱⁱ The displayed luminance differences between these codes are $(10.0278 \text{ nits} - 9.9563 \text{ nits}) = 0.0715 \text{ nits}$ and $(10.0995 \text{ nits} - 10.0278 \text{ nits}) = 0.0717 \text{ nits}$, which represent relative percentages of $(0.0715 \text{ nits}/10.0278 \text{ nits}) = 0.71\%$ and $(0.0717 \text{ nits}/10.0278 \text{ nits}) = 0.72\%$. This means that the 10-bit SDR format can represent a minimum displayed luminance change of 0.71% or larger (by using more than 1 code value) at 10 nits. At other luminance values, the minimum relative change that can be repre-

sented can be relatively larger or relatively smaller; for example, at code value 315 (about 5 nits), the minimum relative luminance change is 0.95%, whereas at code value 720 (about 50 nits), the minimum relative luminance change is 0.37%.

Different combinations of quantization precision and EOTF result in different distributions of code values across the target displayed luminance range, which leads to different limits of quantization performance at different target displayed luminance values. **Table 1** shows the code values that would be used by various formats that use different precisions (8 bits/10 bits/12 bits) and different EOTFs (100 nits BT.1886, 1000 nits BT.1886, 10,000 nits BT.1886, and 10,000 nits ST 2084).

The Barten model of the human visual system predicts the human visual system's ability to detect a change in luminance. The Barten threshold is the percentage luminance change at which a viewer can determine a just-noticeable difference (JND).⁸ The Barten thresholds for luminance targets of 5, 10, and 50 nits displayed under reference viewing conditions are approximately 0.61%, 0.55%, and 0.46%, respectively, and the 10-bit SDR format at those luminance targets has minimum relative luminance changes of 0.95%, 0.71%, and 0.37%, respectively. Therefore, the 10-bit SDR format has performance above the Barten threshold at 5 and 10 nits ($0.95\% > 0.61\%$ and $0.71\% > 0.55\%$, respectively) and performance below the Barten threshold at 50 nits ($0.37\% < 0.46\%$). This means that, if noise-free 10-bit SDR imagery has a 1 code value change centered around 5 or 10 nits, then this change may be visible to a human viewer as a contouring or a banding visual artifact, whereas a 1 code value change centered at 50 nits would not be visible, and the viewer would not see a visual artifact.

Display technology has advanced so that advanced HDR displays can show enhanced shadow details,

ⁱⁱⁱThe calculations used in this section with the BT.1886 EOTF are performed with a theoretically ideal but impractical reference black level of $L_b = 0 \text{ nits}$. Practical SDR displays implementing the BT.1886 EOTF would likely have a nonzero black level. The example values are centered on 10 nits, which is well above the lower range of the video signal that is most affected by a nonzero L_b .

brighter specular highlight details, and brighter colors and more saturated darks. A naive approach to supporting HDR displays would be to simply use the BT.1886 EOTF with a larger luminance range, but this approach would likely lead to significant visual artifacts, such as contouring and banding, across a large portion of the increased luminance range. The combination of quantization and EOTF can be used to predict the minimum relative luminance change allowed by the format, and that minimum relative luminance change can be compared to the Barten threshold to determine the luminance quantization performance of the format. If the minimum relative luminance change of a format using a particular quantization and EOTF is below the Barten threshold, then the format will not introduce a visual artifact related to luminance quantization.

Table 2 summarizes the luminance quantization performance of the Legacy SDR format, the SDR format, the HDR format (labeled “HDR10” in the table and throughout the rest of this paper), and the Dolby Vision format (which is optional for discs and optional for players) that are available on the Ultra HD Blu-ray format. **Table 2** also provides a summary of the luminance quantization performance of two hypothetical naive systems that support HDR using 1000 nits BT.1886 EOTF and 10,000 nits BT.1886 EOTF that are not supported on the Ultra HD Blu-ray format.

A color-coding scheme is used in **Table 2**, where the green entries are less than or equal to the Barten threshold, the yellow entries are 2.5 times the Barten threshold, and the red entries are five times the Barten threshold.

The 10-bit SDR format has much better quantization performance (lower percentages) compared to the 8-bit Legacy SDR format used by Legacy Blu-ray and for most

other forms of consumer digital video distribution that have occurred over the last two decades. The HDR10 format using 10 bits with SMPTE ST 2084 EOTF also has better quantization performance (lower percentages) compared to the 8-bit Legacy SDR format.

The HDR10 format has more uniform performance across the entire luminance range compared to the hypothetical naive systems that would support HDR using 1000 nits BT.1886 EOTF and 10,000 nits BT.1886 EOTF that are not supported by the Ultra HD Blu-ray format. The Dolby Vision 12-bit ST 2084 format that is optional for discs and optional for players results in quantization performance that is below the Barten threshold across the entire luminance range, which means that luminance quantization-related visual artifacts will not be introduced by the quantization performance of the Dolby Vision format.

Entertainment imagery often contains some amount of noise, and the presence of noise in an imaging system allows the system to operate farther from the Barten threshold while not being affected by luminance quantization visual artifacts (banding or contouring). This means that quantization performance greater than the 10-bit SDR and HDR formats used by the Ultra HD Blu-ray format is not necessary to faithfully represent most entertainment content that contains some noise such that the result is free from luminance quantization visual artifacts (see, e.g., visual test results summarized in Table II in the work of Miller et al.⁹). Other factors such as color-related (not luminance-related) quantization artifacts can become significant for wide color gamut encodings of HDR,¹⁰ but the 10-bit HDR format should not be limited by substantial amounts of visually significant color-related quantization artifacts for entertainment imagery that often contains some amount of noise.

TABLE 2. Barten threshold and minimum relative percentage luminance change of different formats.

Target Displayed Luminance (nits)		Minimum Relative Luminance Change Percentage Due to Codevalue Quantization and EOTF at Target Displayed Luminance					
		Legacy SDR	SDR	Not a Format	Not a Format	HDR10	Dolby Vision
Barten Threshold		100 nits BT.1886 8 bit	100 nits BT.1886 10 bit	1000 nits BT.1886 10 bit	10000 nits BT.1886 10 bit	10000 nits ST 2084 10 bit	10000 nits ST 2084 12 bit
0.1	1.75%	18.85%	4.83%	12.17%	30.92%	4.25%	1.08%
0.5	1.03%	9.71%	2.48%	6.36%	16.29%	2.64%	0.66%
1	0.85%	7.34%	1.85%	4.83%	12.17%	2.22%	0.56%
5	0.61%	3.77%	0.95%	2.48%	6.36%	1.60%	0.40%
10	0.55%	2.83%	0.71%	1.85%	4.83%	1.44%	0.36%
50	0.46%	1.46%	0.37%	0.95%	2.48%	1.19%	0.30%
100	0.44%	1.09%	0.27%	0.71%	1.85%	1.13%	0.28%
500	0.41%	N/A	N/A	0.37%	0.95%	1.05%	0.26%
1000	0.40%	N/A	N/A	0.27%	0.71%	1.04%	0.26%
5000	0.40%	N/A	N/A	N/A	0.37%	1.06%	0.27%
10000	0.40%	N/A	N/A	N/A	0.27%	1.08%	0.27%

The HDR10 and Dolby Vision formats are both based on the ST 2084 EOTF that has a fairly uniform quantization performance across the entire 10,000-nits luminance range relative to the Barten threshold across the same luminance range. This was, in fact, one of the design goals of the “perceptual quantizer curve” (often called “PQ”)⁹ that makes up the core of the ST 2084 standard document. Dolby’s PQ proposal was standardized in SMPTE in 2014, with the original proponents of the standardization work being Dolby, MovieLabs, Pixar, and the six major Hollywood studios (Disney, Fox, Warner Bros., Universal Pictures, Paramount, and Sony Pictures Entertainment).¹¹ Its standardization in SMPTE as ST 2084 was an important step before its inclusion as a normative reference in other standards documents, including H.265/HEVC version 2,¹² CEA-861.3,¹³ DECE CFF v2.1,¹⁴ and the Ultra HD Blu-ray specification.

Video Color Encoding Primaries

An SDR video signal must use encoding primaries specified by either Recommendation ITU-R BT.2020¹⁵ or Recommendation ITU-R BT.709.¹⁶ An HDR video signal must use ITU-R BT.2020 encoding primaries.

Video Chroma Subsampling

The video signal must use chroma subsampling in the nonconstant luminance Y’CbCr 4:2:0 format. The luma and chroma equations are different when using BT.709 and BT.2020 color encoding primaries.

The luma and chroma equations used for BT.709 are according to the ITU-R BT.709 specification and are shown as follows:

$$Y' = 0.2126 * R' + 0.7152 * G' + 0.0722 * B'$$

$$Cb = \frac{(B' - Y')}{1.8556}$$

$$Cr = \frac{(R' - Y')}{1.5748}$$

The luma and chroma equations used for BT.2020 are according to the ITU-R BT.2020 specification and are shown as follows:

$$Y' = 0.2627 * R' + 0.6780 * G' + 0.0593 * B'$$

$$Cb = \frac{(B' - Y')}{1.8814}$$

$$Cr = \frac{(R' - Y')}{1.4746}$$

The two different chroma sample locations that are allowed in the format, namely, chroma_loc = 0 and chroma_loc = 2, are illustrated in **Fig. 1**. For chroma_loc = 0, the first chroma sample location is

horizontally cosited with the first luma sample and vertically interstitial between the first luma samples of the first and second lines. For chroma_loc = 2, the first chroma sample location is horizontally cosited with the first luma sample and vertically cosited with the first line’s luma sample. The ITU-R BT.709 specification does not define a 4:2:0 format, whereas the ITU-R BT.2020 specification defines a 4:2:0 format consistent with chroma_loc = 2.

If the video signal uses BT.2020 encoding primaries, the chroma location must use chroma_loc = 2. If the video signal uses BT.709 encoding primaries, the chroma location must use either chroma_loc = 0 or chroma_loc = 2. If the chroma sample location is unspecified within the H.265/HEVC bitstream (chroma_loc_info_present_flag = 0), the decoder will assume a value of chroma_loc = 0. The term “chroma_loc” used in this paper is an abbreviation for the H.265/HEVC video usability information fields called chroma_sample_loc_type_top_field and chroma_sample_loc_type_bottom_field. A high-quality process that creates 4:2:0 signals from 4:4:4 signals or reconstructs 4:4:4 signals from 4:2:0 signals would likely use filters with different phases to compensate for the different chroma sample locations indicated by chroma_loc = 0 and chroma_loc = 2.

Video Compression

SDR and HDR formats must use H.265/HEVC video compression using Main10 Profile with High Tier and Level 5.1. The maximum allowed peak bit rate of the video elementary stream is 100 Mbits/sec, as constrained by the relevant ISO/IEC HRD conformance and by the MPEG-TS T-STD decoder buffer input rate. The maximum duration of a group of pictures is 1.001 sec. The minimum and maximum numbers of slice segments are 1 and 34, respectively.

Additional constraints on the H.265/HEVC bitstream are described in the format’s whitepaper⁴ and specification.³

Allowed Combinations of Video Characteristics

Not all combinations of the video characteristics are permitted in the format. The allowed combinations are summarized in **Table 3**.

Subtitles

Subtitles must be carried in a BD-ROM Presentation Graphics (PG) stream. Subtitles use an 8-bit indexed (palletized) format. For SDR BT.709 and SDR BT.2020 videos, a color lookup table (CLUT) is used to convert each pixel in the subtitle image from an 8-bit index value into 8-bit BT.709 Y’CbCr + alpha values. A BT.709 to BT.2020 color space conversion must be applied by the player to convert the BT.709 subtitle pixel values to BT.2020 before compositing with SDR BT.2020 10-bit video. For HDR BT.2020 video, a CLUT is used to convert each pixel in the subtitle from an 8-bit index value

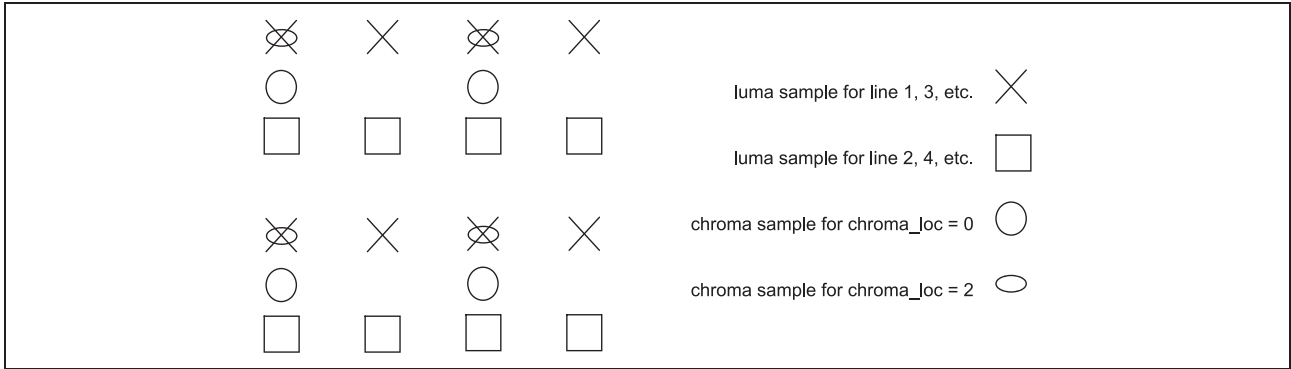


FIGURE 1. 4:2:0 luma and chroma sample locations.

into 8-bit BT.2020 ST 2084 Y'CbCr + alpha 8-bit values. The 8-bit pixel values from the CLUT are multiplied by 4 before compositing with 10-bit video pixels. The subtitle pixels and associated CLUT entry values can be changed by the content author with frame accuracy.

Most authoring tools are designed to accept subtitle input image files in a palletized 8-bit Portable Network Graphics (PNG) file format and to generate the corresponding CLUT entries and subtitle PG streams authored to the disc. Different workflows for creating HDR subtitles are being evaluated by the industry, mainly focused on either creating new HDR subtitles directly in an 8-bit ST 2084 BT.2020 format or on converting subtitles prepared for Legacy SDR video to ST 2084 BT.2020 HDR format, perhaps with some extra luminance gain to increase visibility when composited over HDR video that may contain bright pixels.

Graphics

Interactive menu graphics can use either High-Definition Movie (HDMV) or Blu-ray Disc Java (BD-J) graphics capabilities. HDMV graphics use the same basic format as is used for subtitles (8-bit indexed values with

associated CLUT) and are carried in a BD-ROM Interactive Graphics (IG) stream. BD-J graphics typically use “full-color” 32-bit sRGB graphics (8 bits for each R', G', B' + alpha sample) and typically are stored on the disc in PNG or JPEG format. BD-J graphics are used frequently on major Hollywood studio releases using the Legacy Blu-ray format.

Graphics are limited to HD resolution, and upscaling of the graphics plane from 1920 × 1080 resolution to 3840 × 2160 resolution must be performed by a player when compositing graphics with 3840 × 2160 video. Players must have the capability to convert BD-J graphics in sRGB format to BT.2020 SDR format, so that sRGB graphics can be appropriately composited with BT.2020 SDR video format while preserving their original sRGB displayed color.

The Ultra HD Blu-ray specification went through a minor specification update from V3.0 to V3.1 in December 2015, after this paper was originally authored. All Ultra HD Blu-ray discs and players that are currently commercially available support the V3.1 specification. The V3.1 specification update made some minor changes to the compositing of BD-J graphics with HDR video. The

TABLE 3. Allowed combinations of video parameters on Ultra HD Blu-ray format.

Video type	HDR	SDR	Legacy SDR
Spatial Resolution	3840 x 2160, 1920 x 1080	3840 x 2160, 1920 x 1080	1920 x 1080
Video Compression	HEVC	HEVC	AVC
Picture Format Aspect Ratio	16:9	16:9	16:9
Bit Depth	10	10	8
Color Encoding Primaries	BT.2020	BT.709, BT.2020	BT.709
Chroma Sub Sampling	4:2:0	4:2:0	4:2:0
Frame Rate	23.976p, 24p, 25p, 50p, 59.94p, 60p	23.976p, 24p, 25p, 50p, 59.94p, 60p	23.976p, 24p
Peak Video Bitrate	100Mbits/sec	100Mbits/sec	40Mbits/sec
EOTF	SMPTE ST 2084	BT.1886	BT.1886
Static Metadata	SMPTE ST 2086, MaxCLL, MaxFALL	SMPTE ST 2086	N/A

author expects the BDA to make additional technical information about V3.1 specification publicly available via its website²⁵ by the summer of 2016.

Optional for Discs and Mandatory for Players

Static Metadata

Static metadata is metadata that does not change over the duration of the playback of the associated video stream. The SMPTE ST 2086, maximum content light level (MaxCLL), and maximum frame average light level (MaxFALL) static metadata items are optional for discs and mandatory for players. Static metadata is optional for discs, so that the content owner need not set the metadata values if they are unknown to the content owner during disc authoring. Some past systems have required metadata to be transmitted on a video interface, without allowing unknown values, which can lead to situations where the metadata that is transmitted is unreliable for the receiving device.

In systems that require mandatory metadata transmission, if the metadata data values are unknown, there may be a tendency to populate the mandatory metadata fields with default values to fulfill the metadata transmission requirement, but this can result in unreliable metadata because a device receiving the default metadata value cannot tell whether the default value is in fact a valid value or the default value was set because the valid value was unknown or not available.

The downstream display devices connected to Ultra HD Blu-ray players may be designed to modify the displayed image, based on the transmitted metadata values, and therefore, the metadata items that are delivered by a player should be reliable. Because the format allows the static metadata field values on the disc to be marked as unknown, the presence of a metadata value that is not marked as unknown indicates that the content owner took it to be valid during content authoring. Allowing valid (known) metadata to be optional on disc but required to be transmitted on an interface that supports it is a good way to ensure that the metadata that is received by a connected display is reliable metadata, assuming that the content authors respect the practice of marking unknown metadata values as unknown.

The player must forward the ST 2086, MaxCLL, and MaxFALL metadata information that is present on a disc to a connected display, if the display interface between the source (player) and the sink (display) supports transfer of the SMPTE ST 2086, MaxCLL, and MaxFALL metadata. The format allows static metadata items to be associated with the video streams referenced by a playlist. The static metadata items are allowed to change when the playlist is changed.

The video interface specification CEA-861-F was recently updated by specification CEA-861.3 to support HDR signaling via the extended display identification data mechanism and static metadata transmission from source

to sink via the auxiliary video information InfoFrame mechanism, including the static metadata items ST 2086, MaxCLL, and MaxFALL. The High-Definition Multimedia Interface (HDMI) standard was recently updated to reference CEA-861.3 to support HDR signaling and static metadata transmission in the HDMI 2.0a specification.

SMPTE ST 2086

The ST 2086 metadata is optional for discs for both SDR and HDR video formats. The ST 2086 standard¹⁷ defines the metadata fields used to represent the following six important pieces of information about the mastering display that was used to perform the creative work (including adjustments and approval) for the corresponding video signal:

- CIE xy chromaticity coordinate of the red display primary
- CIE xy chromaticity coordinate of the green display primary
- CIE xy chromaticity coordinate of the blue display primary
- CIE xy chromaticity coordinate of the white point of the display
- Maximum luminance of the display as configured for the mastering process
- Minimum luminance of the display as configured for the mastering process

ST 2086 metadata, often referred to as “Mastering Metadata,” is likely to be consistent across content titles mastered using the same facility, equipment, and mastering process. The ST 2086 metadata values, if present, are stored in the playlist file on the disc. If the ST 2086 metadata values are also present within the HEVC/H.265 bitstream stored in supplemental enhancement information (SEI) messages according to the syntax specified in section D.2.27 of the HEVC/H.265 specification, then the values being represented by the SEI messages should match the values stored in the playlist on the disc, except that the maximum and minimum luminance fields are represented with different semantics in the SEI message compared to the playlist representation (in units of 0.0001 nits utilizing 32 bits for each value in the SEI message compared to units of 1 and 0.0001 nits utilizing 16 bits for each value in the playlist representation, respectively). The representation used by CEA-861.3 is the same as the playlist representation.

MaxCLL

The MaxCLL metadata value is optional for discs and is allowed for the HDR video format. MaxCLL is not permitted for the SDR video format. The MaxCLL value is expected to vary from title to title, as different mastering setups used by different content creators may impose limits on the maximum light level mastered into

the content. Additionally, the maximum light levels that appear in HDR content may be limited for creative or technical reasons.

To compute the MaxCLL for a particular content, the max(R,G,B) operator is applied to all pixels in all frames of the content to determine the maximum value (MaxCLL) for that particular content. If a simple tone mapping approach is used by the consumer display to replace pixel values in the content that are not representable on the consumer display with pixel values that are representable, then the MaxCLL could be used to define the upper bound of the value of the pixels that will be encountered in that particular content. More information about the techniques that can be used to compute the MaxCLL value appears in Annex A of CEA-861.3 and also in a recent Joint Collaborative Team on Video Coding (JCT-VC) document.¹⁸ The unit of MaxCLL is candelas per square meter (nits).

MaxFALL

The MaxFALL metadata value is optional for discs and is allowed for the HDR video format. MaxFALL is not permitted for the SDR video format. MaxFALL is expected to vary from title to title.

To compute the MaxFALL for a particular content, the average of the max(R,G,B) operator is applied to all pixels in each frame to determine the frame-average maxRGB for each frame. The MaxFALL value is set to the maximum value of the frame-average maxRGB of all frames in the content. The computation of the MaxFALL values only considers the active image areas of the frame, which is relevant when, for example, a 2.40:1 aspect ratio content is stored in a 16:9 aspect ratio pixel array with letterbox mattes for distribution to the home; hence, the MaxFALL value can remain valid if cropping or zooming is applied to the image. More information about the techniques that can be used to compute the MaxFALL value appears in Annex A of CEA-861.3 and in a recent JCT-VC document.¹⁸ The unit of MaxFALL is candelas per square meter (nits).

Playback of Dual Stream

A dual stream is a single transport stream that contains separate SDR and HDR video-elementary streams and subtitle assets but only one set of audio-elementary streams that are commonly used. The use of a dual stream allows the content author to provide both SDR and HDR versions of the video and subtitles, which permits playback to all connected display types while not relying on the player's HDR to SDR conversion implementation and while simultaneously saving disc space by not duplicating audio-elementary streams. With the dual-stream option, both the SDR and HDR versions of the video need to be multiplexed in one transport stream subject to the maximum transport stream rate defined in the specification. An alternative to the use of a dual stream is to provide separate transport streams for HDR and SDR

video versions, which requires duplicate audio assets in both the HDR and SDR transport streams, which increases disc capacity usage.

Optional for Discs and Optional for Players

BD-J Graphics With SDR to HDR Conversion

The BD-J graphics subsystem of the Ultra HD Blu-ray format was designed to reuse the BD-J graphics subsystem from the Legacy Blu-ray format while adding some small changes to support overlaying graphics on SDR BT.2020 and HDR BT.2020 videos. If BD-J graphics are used with HDR video, the player may implement an optional SDR to HDR conversion with a gain factor to convert SDR sRGB values stored in the graphics plane to BT.2020 ST 2084 HDR format. A gain factor associated with each playlist can be used by the optional SDR to HDR conversion to increase the luminance of the BD-J graphics so that they do not appear dim when overlaid on top of HDR video that may contain bright pixel areas. The BD-J application can be designed to select alternative graphic asset files (such as PNG or JPEG files) if the player does not support the optional SDR to HDR conversion.

The steps outlined in **Figs. 2** and **3** are two different reference SDR to HDR conversion algorithms with a gain factor to increase the luminance so there can be an improved match of graphics luminance with the brighter pixels that may appear in HDR video.

An SDR image that has been processed by one of these SDR to HDR conversion algorithms will not possess anymore dynamic range than it started with. If a gain value greater than 1 is used as part of the conversion, then the converted image will appear brighter than it would if the SDR image were viewed with BT.1886 EOTF using 100 nits for reference white. Conceptually, this conversion is just an EOTF format conversion with gain compared to more advanced SDR to HDR conversion algorithms¹⁹ that may attempt to actually increase the dynamic range of the displayed image, which are often called “inverse tone mapping” algorithms. Player manufacturers may choose to implement one of the algorithms described in **Figs. 2** and **3**, or they may choose to implement other SDR to HDR conversion algorithms of their own design if they choose to support the optional conversion feature.

Graphics blending in consumer devices is often performed using the Porter-Duff SRC_OVER compositing operator, which results in a linear combination of the foreground graphic plane pixel values and the background video plane pixel values using the foreground graphic plane pixel's transparency, which is called alpha, as the multiplicative factor according to the following equation:

$$\begin{aligned} \text{blending result} &= \text{alpha} * \text{foreground} \\ &+ (1 - \text{alpha}) * \text{background}. \end{aligned}$$

SDR to HDR conversion with gain and sRGB to BT.2020 color space conversion

1. Inverse quantize full range 8bit sRGB codevalues (CV_1)

$$N_1 = CV_1 / 255$$

2. Apply SDR EOTF

$$L_1 = (N_1)^{2.4}$$

3. Multiply by reference SDR luminance (100 nits) and desired gain

$$A = L_1 * 100 * \text{gain}$$

4. Conversion from sRGB color space to BT.2020 color space, using $R_{709} = A_R$, $G_{709} = A_G$, $B_{709} = A_B$

$$R_{2020} = 0.62740389593470 * R_{709} + 0.32928303837789 * G_{709} + 0.04331306568741 * B_{709}$$

$$G_{2020} = 0.06909728935823 * R_{709} + 0.91954039507545 * G_{709} + 0.01136231556630 * B_{709}$$

$$B_{2020} = 0.01639143887515 * R_{709} + 0.08801330787723 * G_{709} + 0.89559525324763 * B_{709}$$

5. Divide by SMPTE ST2084 HDR "container" luminance (10000 nits), using $A_R = R_{2020}$, $A_G = G_{2020}$, $A_B = B_{2020}$

$$L_2 = A / 10000$$

6. Apply SMPTE ST2084 inverse EOTF using $L = L_2$

$$N = ((c1 + c2 * L^{m1}) / (1 + c3 * L^{m1}))^{m2}$$

$$\text{Where: } m1 = 0.1593017578125 \quad m2 = 78.84375, \quad c1 = 0.8359375, \quad c2 = 18.8515625, \quad c3 = 18.6875$$

7. Quantize to 10bit codevalue (choose either a, b, or c depending on the desired output format)

a. legal range R'G'B'

$$CV_2 = \text{floor}(876 * N + 64 + 0.5)$$

b. full range R'G'B'

$$CV_2 = \text{floor}(1023 * N + 0.5)$$

c. legal range Y'CbCr, using $R'_{2020} = N_R$, $G'_{2020} = N_G$, $B'_{2020} = N_B$

$$Y' = 0.2627 * R'_{2020} + 0.6780 * G'_{2020} + 0.0593 * B'_{2020}$$

$$Cb = (B'_{2020} - Y') / 1.8814$$

$$Cr = (R'_{2020} - Y') / 1.4746$$

$$CV_Y = \text{floor}(876 * Y' + 64 + 0.5)$$

$$CV_Cb = \text{floor}(896 * Cb + 512 + 0.5)$$

$$CV_Cr = \text{floor}(896 * Cr + 512 + 0.5)$$

FIGURE 2. SDR to HDR conversion with gain and sRGB to BT.2020 color space conversion.

The blending result can appear differently when the blending calculation is operating on SDR pixel values versus operating on HDR pixel values. The differences are most significant where there is a large luminance difference between the foreground graphic and the background video and the graphic is semitransparent. The appearance of the blending result using fully opaque or fully transparent graphics is not expected to be dependent on the format of the pixel values. A content author that expects the blending result to be the same, if the graphic is converted by the optional SDR to HDR conversion, should carefully examine the visual result involving semitransparent graphics.

Dolby Vision

A disc that contains HDR10 primary video may optionally include an additional Dolby Vision enhancement layer within the transport stream associated with that HDR primary video. The Dolby Vision enhancement layer is a proprietary format that includes additional video data encoded using HEVC/H.265 video compression and dynamic metadata. A player may optionally support the Dolby Vision feature. A connected display may optionally support receiving Dolby Vision dynamic metadata from a player.

A specification describing the dynamic metadata associated with the Dolby Vision feature has been submitted for

SDR to HDR conversion with gain

1. Convert 8bit sRGB codevalues to 10bit SDR BT.2020 full range RGB codevalues using players required conversion.
2. Inverse quantize full range 10bit SDR BT.2020 RGB codevalues (CV_1)

$$N_1 = CV_1 / 1023$$

3. Apply SDR EOTF

$$L_1 = (N_1)^{2.4}$$

4. Multiply by reference SDR luminance (100 nits) and desired gain

$$A = L_1 * 100 * \text{gain}$$

5. Divide by SMPTE ST2084 HDR "container" luminance (10000 nits)

$$L_2 = A / 10000$$

6. Apply SMPTE ST2084 inverse EOTF using $L = L_2$

$$N = ((c1 + c2 * L^{m1}) / (1 + c3 * L^{m1}))^{m2}$$

Where: $m1 = 0.1593017578125$ $m2 = 78.84375$, $c1 = 0.8359375$, $c2 = 18.8515625$, $c3 = 18.6875$

7. Quantize to 10bit codevalue (choose either a, b, or c depending on the desired output format)

- a. legal range R'G'B'

$$CV = \text{floor}(876 * N + 64 + 0.5)$$

- b. full range R'G'B'

$$CV = \text{floor}(1023 * N + 0.5)$$

- c. legal range Y'CbCr, using $R'_{2020} = N_R$, $G'_{2020} = N_G$, $B'_{2020} = N_B$

$$Y' = 0.2627 * R'_{2020} + 0.6780 * G'_{2020} + 0.0593 * B'_{2020}$$

$$Cb = (B'_{2020} - Y') / 1.8814$$

$$Cr = (R'_{2020} - Y') / 1.4746$$

$$CV_Y = \text{floor}(876 * Y' + 64 + 0.5)$$

$$CV_Cb = \text{floor}(896 * Cb + 512 + 0.5)$$

$$CV_Cr = \text{floor}(896 * Cr + 512 + 0.5)$$

FIGURE 3. SDR to HDR conversion with gain.

standardization in the SMPTE, as part of the ST 2094 document set, but has not yet completed the standardization process as of September 2015 (the time of the submission of this paper). Dynamic metadata for the Dolby Vision feature is information that is created during mastering that is associated with the relationship between HDR and SDR versions of a particular video frame. The dynamic metadata used for the Dolby Vision feature in Ultra HD Blu-ray is stored in the Dolby Vision enhancement layer.

The Dolby Vision codec architecture²⁰ shown in Fig. 4 uses an SDR base layer that is not allowed on Ultra HD Blu-ray and an HDR enhancement layer; hence, the base layer provides simple backwards compatibility with SDR devices and SDR displays. The Dolby Vision codec architecture for Ultra HD Blu-ray disc uses an HDR10 base layer and an HDR enhancement layer, as shown in Fig. 5, and the base layer provides compatibility with HDR10 devices and displays,

but this architecture is optional for discs and optional for players in the Ultra HD Blu-ray format. The Dolby Vision codec architecture for Ultra HD Blu-ray provides compatibility with SDR devices and SDR displays using HDR to SDR conversion in the player that is controlled with dynamic metadata.

In the use case in which a disc, a player, and a connected display all support the Dolby Vision feature, the HDR10 primary video stream and the special video-elementary stream representing the Dolby Vision enhancement layer are decoded. Following decoding, the Dolby Vision enhancement layer video signal is combined with the HDR10 video signal to make a resulting 12-bit ST 2084 BT.2020 HDR video signal, and then, the 12-bit HDR video signal and the dynamic metadata are sent to the Dolby Vision capable display. The luminance quantization performance of Dolby Vision 12-bit ST 2084 format is below the Barten threshold for the entire luminance

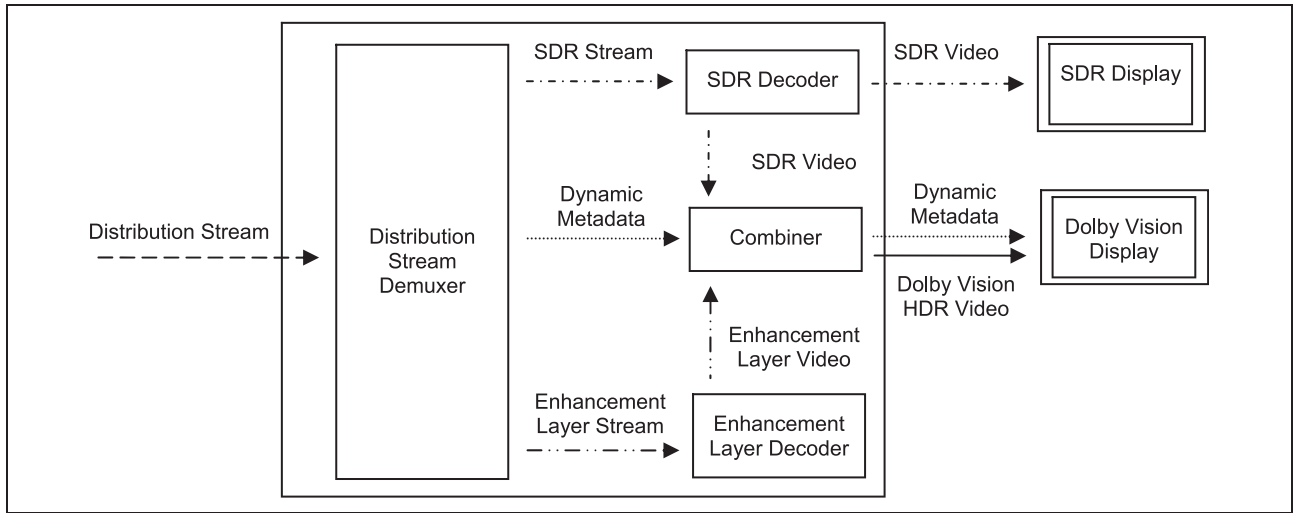


FIGURE 4. The Dolby Vision codec architecture with SDR base layer and HDR enhancement layer.

range, as shown in the rightmost column in **Table 2** and in Fig. 5(a) in the work of Miller et al.,⁹ whereas the color quantization performance of 12-bit ST 2084 BT.2020 Y'CbCr format is below 1.5 DeltaE2000 JNDs for the entire luminance range, as shown in figure 8 in the work of Nezamabadi et al.¹⁰ The resulting Dolby Vision 12-bit ST 2084 BT.2020 video signal format provided to a Dolby Vision capable display should not cause visual artifacts related to luminance quantization or color quantization for any type of imagery, including noise-free images.

In the use case in which a disc and a player support the Dolby Vision feature but the connected display is SDR and does not support Dolby Vision or the HDR10

format, the HDR10 primary video stream and the special video-elementary stream representing the Dolby Vision enhancement layer are decoded and combined to make a resulting 12-bit ST 2084 BT.2020 HDR video signal, which is then converted to SDR by the player utilizing the Dolby Vision dynamic metadata, and the SDR video signal is sent to the connected SDR display.

In the use case in which a disc and a player support the Dolby Vision feature but the connected display does not support the Dolby Vision feature and only supports the HDR10 format, only the HDR10 primary video stream is decoded by the player, and the HDR10 video signal is sent to the connected HDR10 display. In the use

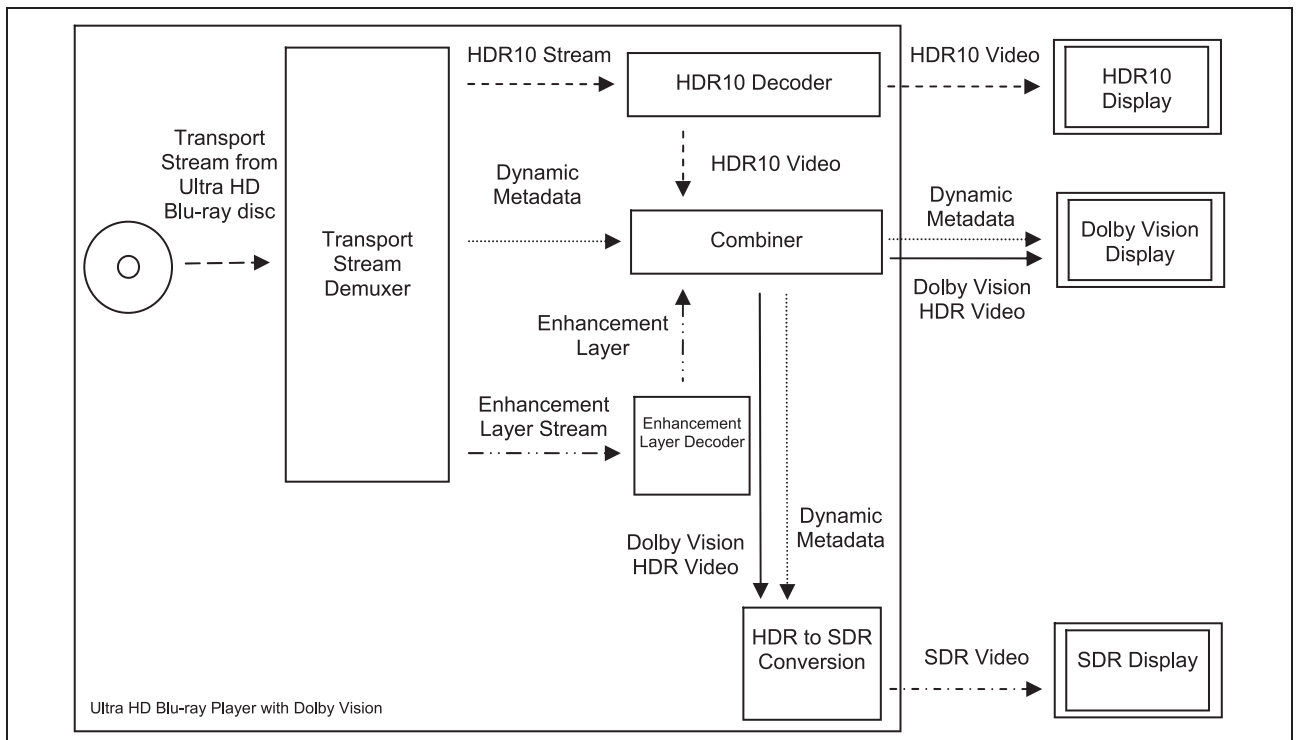


FIGURE 5. The Dolby Vision codec architecture for Ultra HD Blu-ray uses HDR10 base layer.

case in which a disc supports the Dolby Vision feature, the player does not support the Dolby Vision feature, and the connected display supports the connected HDR10 format, only the HDR10 primary video stream is decoded by the player, and the HDR10 video signal is sent to the connected HDR10 display. In the use case in which a disc supports the Dolby Vision feature, the player does not support the Dolby Vision feature, and the connected display is SDR and does not support the HDR10 format, only the HDR10 primary video stream is decoded by the player, and the player converts the HDR10 video signal into an SDR signal and sends the SDR signal to the connected SDR TV.

The use cases described in this section, involving combinations of Dolby Vision and HDR10 technology and their use in Ultra HD Blu-ray applications, are summarized in **Table 4**.

Philips HDR

A disc that contains HDR primary video may optionally include the Philips HDR feature. A player may optionally support the Philips HDR feature. A disc that supports the Philips HDR feature will contain an HDR10 primary video stream that stores Philips HDR dynamic metadata in the HEVC/H.265 bitstream’s private SEI messages. Philips HDR dynamic metadata is information that is created during mastering that is associated with the relationship between HDR and SDR versions of a particular video frame. A player that supports the Philips HDR feature can support two different use cases illustrated in **Fig. 6**.

In the first use case, if the disc and the player support the Philips HDR feature, the player will extract the Philips HDR dynamic metadata from the SEI messages that may be present in the video stream and send the extracted information along with the decoded HDR10 video signal to a connected display if the display is capable of receiving Philips HDR dynamic metadata. The connected display will use the Philips HDR dynamic metadata to map the received signal to its specific capabilities (e.g., peak luminance and black level).

In the second use case, if the disc and the player support the Philips HDR feature, the player will extract the Philips HDR dynamic metadata from the SEI messages that may be present in the video stream and use the information to perform HDR to SDR conversion in the player before sending SDR video to an SDR display.

The additional bit rate associated with the Philips HDR dynamic metadata is approximately 10 kbits/sec.²¹ A specification describing the dynamic metadata associated with the Philips HDR feature has been submitted for standardization in the SMPTE, as part of the ST 2094 document set, but has not yet completed the standardization process as of September 2015 (the time of the submission of this paper).

Color Remapping Information

Color remapping information (CRI) is dynamic metadata developed by Technicolor.²²⁻²⁴ The syntax and semantics of the CRI SEI are defined in sections D.2.32 and D.3.32 of HEVC/H.265 version 2.¹² CRI is optional for discs and players and may be used as an HDR to SDR conversion feature. CRI dynamic

TABLE 4. Use case summary for combinations of Dolby Vision and HDR10 for Ultra HD Blu-ray.

Disc Supports		Player Supports		Display Supports		Result for this Use Case
HDR10	Dolby Vision	Dolby Vision	Dolby Vision	HDR10	HDR10	
yes	yes	yes	yes	N/A		Dolby Vision video signal with dynamic metadata sent to display
yes	yes	yes	no	yes		HDR10 video signal sent to display
yes	yes	yes	no	no		Player does HDR to SDR conversion using Dolby Vision dynamic metadata, SDR video signal sent to display
yes	yes	no	N/A	yes		HDR10 video signal sent to display
yes	yes	no	N/A	no		Player performs HDR to SDR conversion without using Dolby Vision dynamic metadata, SDR video signal sent to display
yes	no	N/A	N/A	yes		HDR10 video signal sent to display
yes	no	N/A	N/A	no		Player performs HDR to SDR conversion, SDR video signal sent to display
no	no	N/A	N/A	N/A		Player decodes SDR video on disc, SDR video signal sent to display

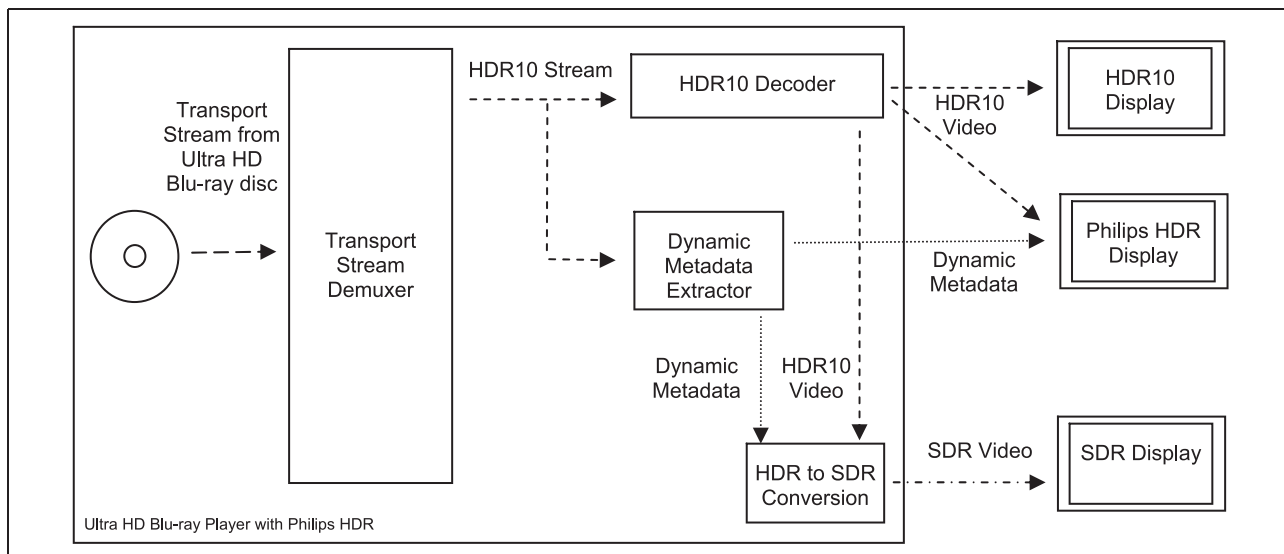


FIGURE 6. Philips HDR for Ultra HD Blu-ray.

metadata information can be created during disc authoring by analyzing both creatively approved HDR and SDR video signals after the signals are finalized during mastering. If a disc and a player both support the CRI feature and the player is connected to an SDR display, then the CRI dynamic metadata information is extracted from the CRI SEI messages that would be present in the HDR10 video stream and the information is used to control an HDR to SDR conversion that is performed in the player before sending the SDR video to the SDR display.

A specification describing the dynamic metadata associated with the Technicolor CRI feature has been submitted for standardization in SMPTE, as part of the ST 2094 document set, but has not yet completed the standardization process as of September 2015 (the time of the submission of this paper).

Mandatory for Discs

Discs must pass a verification test in order to be released as an Ultra HD Blu-ray disc.

Mandatory for Players

BT.709 to BT.2020 Conversion for HDMV Subtitles and Graphics

Players must have the capability to convert PG subtitles and IG graphics pixel values from SDR BT.709 Y'CbCr to SDR BT.2020 Y'CbCr values before compositing with SDR BT.2020 video.

sRGB to BT.2020 Conversion for BD-J Graphics

Players must have the capability to convert BD-J graphics pixel values from sRGB to BT.2020 before compositing with BT.2020 video.

BT.2020 to BT.709 Conversion

Players must have the capability to convert the output video signal from BT.2020 to BT.709 to support legacy display connection.

HDR to SDR Conversion

Players must have the capability to convert the output video signal from HDR to SDR to support legacy display connection.

Backwards Compatibility

Players must have the capability to playback Legacy Blu-ray discs. Ultra HD Blu-ray players may optionally support playback of Blu-ray 3D discs.

Disc Capacity and Bit rates

Legacy HD Blu-ray BD-ROM format supports either one or two 25-Gbyte disc layers, leading to 25 and 50-Gbyte capacities, respectively. Ultra HD Blu-ray BD-ROM format supports two 25-Gbyte disc layers, two 33.4-Gbyte disc layers, or three 33.4-Gbyte disc layers, leading to 50, 66, and 100-Gbyte disc capacities, respectively. The 50-Gbyte disc supports a maximum transport stream bit rate of 81.7 Mbits/sec. The 66 and 100-Gbyte discs support a maximum transport stream bit rate of 109 Mbits/sec. The 66 and 100-Gbyte discs also support a two-zoned disc layout with an inner low-transfer-rate zone and an outer high-transfer-rate (HTR) zone, with the HTR zone supporting a maximum transport stream bit rate of 127.9 Mbits/sec.

Conclusion

Next-generation consumer video will be enhanced by increased quantization performance, higher frame rates, higher resolution, wider color gamut, higher dynamic range, and improved compression efficiency, allowing a video representation with greater fidelity than was

previously available. This paper provided a summary of how those next-generation video enhancements are supported by the Ultra HD Blu-ray format.

Acknowledgments

The author thanks Warner Bros. for supporting him as the co-chair of the BDA UHD-TF Video Subgroup. The author acknowledges the following individuals for their diligent efforts and productive collaboration related to working on the new format: Lewis Ostrover, Michael Zink, Brian Link, David Brooks, Joop Talstra, Leon van de Kerkhof, Michael DeValue, Stephan Luka, Dominic Glynn, Arjun Ramamurthy, Joe McCrossan, Sally Hattori, Kuniaki Takahashi, Kouichi Uchimura, Yoshikazu Takashima, Masayuki Kozuka, Takahiro Nishi, Tadamasu Toma, Hiroshi Yahata, Jim Helman, and Chad Fogg.

About the Author



Michael D. Smith has worked as a digital imaging, signal processing, and intellectual property consultant since 2003. He is currently with Wavelet Consulting LLC. He has recently served, on behalf of Warner Bros., as a co-chair of Blu-ray Disc Association's UHD-TF Video Subgroup, which defined the video-

related requirements for the Ultra HD Blu-ray disc format. He has worked on technical projects for many organizations during his career, including Warner Bros., Sony Pictures Entertainment, DCI, NOAA, Dolby, and DTS. He has worked on 35 intellectual property matters related to infringement and validity analysis of patents. He has worked on patent cases for Skyhook, resulting in a \$90 million settlement, and on patent cases for TiVo, resulting in total payments of \$1.6 billion. He was editor of the book "3D Cinema and Television Technology: The First 100 Years," published by SMPTE in 2011. He received a BS and an MS in electrical engineering from the University of California at Los Angeles in 2001 and 2004, respectively. He can be reached at miksmith@miksmith.com.

References

1. Blu-ray Disc Association, "Blu-ray Disc Association Completes Ultra HD Blu-ray Specification and Releases New Logo," <http://www.businesswire.com/news/home/20150512005300/en/Blu-ray-Disc-Association-Completes-Ultra-HD-Blu-ray>.
2. Blu-ray Disc Association, "Blu-ray Disc Association to Commence Licensing of Ultra HD Blu-ray," <http://www.businesswire.com/news/home/20150805005853/en/Blu-ray-Disc-Association-Commence-Licensing-Ultra-HD>.
3. Blu-ray Disc Association, "System Description Blu-ray Disc Read-Only Format Part 3: Audio Visual Basic Specifications Version 3.0," 2015.
4. Blu-ray Disc Association, "White Paper Blu-ray Disc Read-Only Format (Ultra HD Blu-ray): Audio Visual Basic Specifications

- Version 3.0," July 2015, <http://www.blu-raydisc.com/en/Technical/TechnicalWhitePapers/BDR0M.aspx>.
5. ITU-R Recommendation BT.1886-0, "Reference Electro-Optical Transfer Function for Flat Panel Displays Used in HDTV Studio Production," International Telecommunication Union, Geneva, 2011.
6. SMPTE ST 2084:2014, "High Dynamic Range Electro-Optical Transfer Function of Mastering Reference Displays," Society of Motion Picture and Television Engineers, 2014.
7. ITU-R Recommendation BT.2035-0, "A Reference Viewing Environment for Evaluation of HDTV Program Material or Completed Programs," International Telecommunication Union, Geneva, 2013.
8. ITU-R Report BT.2246-2, "The Present State of Ultra-High Definition Television," International Telecommunication Union, Geneva, Nov. 2012.
9. S. Miller, M. Nezamabadi, and S. Daly, "Perceptual signal coding for more efficient usage of bit codes," *SMPTE Mot. Imag. J.*, 122(4):52-59, May/June 2013.
10. M. Nezamabadi, S. Miller, S. Daly, and R. Atkins, "Color signal encoding for high dynamic range and wide color gamut based on human perception," *Proc. SPIE-IS&T Electron. Imag.*, vol. 9015, pp. 1-12, Feb. 2014.
11. SMPTE, "Perceptually-Based EOTF for High Luminance and Wide Color Gamut Images," https://kws.smppte.org/kws/projects/project/details?project_id=197.
12. ITU-T Recommendation H.265, "High Efficiency Video Coding," International Telecommunication Union, Geneva, Jan. 2015.
13. CEA Standard CEA-861.3, "HDR Static Metadata Extensions," Consumer Technology Association, Arlington, VA, Jan. 2015.
14. DECE, "Common File Formats & Media Formats Specification v.2.1," Feb. 2015, <http://www.uvcentral.com/sites/default/files/files/PublicSpecs/CFFMediaFormat-2.1.pdf>.
15. ITU-R Recommendation BT.2020-1, "Parameter Values for Ultra-High Definition Television Systems for Production and International Programme Exchange," International Telecommunication Union, Geneva, 2014.
16. ITU-R Recommendation BT.709-6, "Parameter Values for the HDTV Standards for Production and International Programme Exchange," International Telecommunication Union, Geneva, 2015.
17. SMPTE ST 2086:2014, "Mastering Display Color Volume Metadata Supporting High Luminance and Wide Color Gamut Images," Society of Motion Picture and Television Engineers, 2014.
18. C. Fogg, J. Helman, M. Smith, and M. Zink, "Content light level information SEI," *JCT-VC-T0101, Joint Collaborative Team on Video Coding 20th Meeting*, Geneva, Feb. 10-18, 2015.
19. L. Meylan, S. Daly, and S. Susstrunk, "Tone mapping for high dynamic range displays," *Proc. 12 Human Vis. Electron. Imag.*, vol. 6492, pp. 1-12, 2007.
20. D. G. Brooks, "The art of better pixels," *SMPTE Mot. Imag. J.*, 124(4):42-48, May/June 2015.
21. R. Goris, R. Brondijk, and R. van der Vleuten, "Parameter Based Compatible HDR Proposal," MPEG document m35067 for Strasbourg October 2014 meeting.
22. P. Andrivon, P. Bordes, and E. François, "SEI Message for Colour Mapping Information," JCT-VC document JCTVC-Q0074 for Valencia March 2014 meeting.
23. P. Bordes, P. Andrivon, P. Lopez, and F. Hiron, "Color Mapping SEI Message," JCT-VC document JCTVC-N0180 for Vienna July 2013 meeting.
24. P. Andrivon, E. François, and P. Bordes, "AVC Update With Colour Remapping Information SEI Message," MPEG document M35665 for Geneva February 2015 meeting.
25. Blu-ray Disc Association, BDR0M Technical White Paper Website Area, <http://www.blu-raydisc.com/en/Technical/TechnicalWhitePapers/BDR0M.aspx>.

Presented at the SMPTE 2015 Annual Technical Conference & Exhibition, Hollywood CA, 26-29 October 2015. Copyright © 2016 by SMPTE.