American Society of Cinematographers Motion Imaging Technology Council Progress Report 2020

Advanced Imaging Committee

Chair: Gary Demos Vice-Chair: Joe Kane Vice-Chair: Bill Mandel Vice-Chair: Jim Fancher Secretary: David Reisner

HDR and WCG Background

In the current configuration of advanced television systems, there are two prominent architectures defined for use distributing High Dynamic Range (HDR) and Wide Color Gamut (WCG) signals. Ignoring optional metadata, both architectures are defined in BT.2100-2. The color primaries are specified for both architectures in BT.2100-2 table 2. The non-linear transfer functions are specified in BT.2100-2 table 4 (PQ) and BT.2100-2 table 5 (HLG). PQ is defined using a hypothetical display capable of extreme deep black, as well as representing pixels up to 10,000cd/m² maximum brightness. It is appropriate to consider PQ as being display-referred, although the specified display does not exist. There is also no specification of the ambient surround level associated with the on-screen PQ brightness level. HLG is defined as a transform from scene-referred light for a variety of display brightnesses and ambient surround levels. The range of scene-referred light in HLG is limited to two or three stops above diffuse scene white, and is limited in dark distinction through the use of gamma 2.0 (note that PQ allocates many more code values to dark distinction than does HLG).

BT.2100-2 table 3 provides some specification of reference HDR viewing, with the key parameters being a maximum brightness of at least 1000cd/m², a black level at or below 0.005cd/m², and an ambient surround luminance of 5cd/m² at D65. Although there is no variability in the ambient surround definition, there is a large range for bright and dark reference display. EBU Tech3320-4.1 sections 2.4.1.1 and 2.4.1.2 similarly specify at least 1000cd/m² for Grade 1A and 1B monitors. For example, a 3,000cd/m² display is allowed as well as a 1,000cd/m² display. For dark range, a .000001cd/m² black is allowed as well as .005cd/m² according to BT.2100-2 table 3. EBU Tech3320-4.1 section 2.4.2 specifies black level as capable of adjustment down to .005 cd/m2 for Grade 1A and 1B monitors.

If grading to these disparate display levels, the resulting graded HDR is likely to be inconsistent in the HDR image pixel values that it specifies, depending on the values of these display parameters within the particular reference display being used.

BT.2100-2 table 3 in note 3c states that full screen maximum luminance can be less than small area maximum luminance. There is no specification for the reference behavior of this, other than that this behavior is allowed. This adds further potential variability to reference viewing.

BT.2100-2 table 3 does not specify a "minimum gamut". If such a minimum gamut were to be specified, then there would be variation in reference viewing between that

minimum gamut and the maximum gamut as specified in BT.2100-2 table 2. Given that the minimum gamut is not specified in BT.2100-2 table 3, there can be variation between any reference monitor's gamut and the maximum gamut as specified BT.2100-2 table 2. For example, a reference HDR display could meet the requirements of BT.2100-2 table 3, and yet could be limited to the gamut of Rec709 primaries. EBU Tech3320-4.1 section 2.4.3 specifies a minimum color gamut as a range of chromaticities for each of red, green, and blue. Note that the specified green primary and its tolerance are entirely outside of P3 green. Note 1 in 2.4.3 indicates that the minimum gamut specified can be obtained using 90% of the xy chromaticity area vs. BT.2100-2 Table 2 (as also specified in BT.2020) for Grade 1A monitors and 60% for Grade 1B and Grade 2 monitors. Note 2 points out that the green primary and its tolerance are outside currently existing reference display gamuts.

Standard Dynamic Range

The situation is further complicated when attempting to also use a common single master for creation of a standard dynamic range (SDR) result, possibly using a standard dynamic range reference projector (e.g. 48cd/m² with a .024cd/m² black corresponding to a 2,000:1 dynamic range with a P3 minimum gamut). If graded in scene-referred light, the result is sometimes being called an "archival master", despite the invisibility of pixel values outside of the reference projector's dynamic range and gamut. While scene-referred pixels contain information that supports HDR and WCG, these pixels are not visible during SDR reference viewing, thus casting doubt on what is being mastered.

With historic SDR in television (e.g. Rec709 HDTV, PAL, NTSC), there was little intent that the reference match the distributed viewing displays. This has been problematic since the dawn of color television. However, the dynamic range of SDR televisions, historically using cathode ray tubes (CRT's), was limited both in bright range and dark range. Phosphor emission spectra were also constrained in variability. However, with HDR and WCG, as well as a wide variety of display technologies (and emission spectra), the variation has been greatly increased. Any given end-user viewing HDR is likely to see an image that is substantially different than the reference viewing appearance. The problem is not new, but has been greatly increased by HDR and WCG.

Multiple Versions

One approach to all of this variation is to have multiple versions. These can be individually graded, lightly adjusted from other gradings, fully automatic, or a combination. Implicit in using graded or adjusted versions is that each version is likely to express a different visual creative intent. Multiple versions will usually not be the preferred architecture.

Most challenging is the disparity between the SDR version(s) and one or more HDR versions.

Single Creative Intent

The concept of a single master is based upon a) having a single creative intent over the entire range of presentation displays, and b) having technical mechanisms to adjust appearance for each such presentation display to consistently achieve that creative intent.

The choice of a single creative intent will often be desirable for aesthetics and/or efficiency. Having the ability to achieve a single creative intent does not preclude multiple versions. However, if the technical mechanisms for consistent appearance are not available, then a single master is correspondingly also not available, even if a single master were to be the preferred choice.

Perhaps separate consideration is appropriate for SDR vs HDR, since it is challenging to span such a large range. If SDR is its own grade, then HDR is also its own grade, and there are correspondingly dual creative intents. SDR may also be distinguished between 48cd/m² DCinema but maybe having a second SDR version for video at 100cd/m², going perhaps up to 350cd/m² with appropriate appearance compensation. Having a single SDR master achieving consistent creative intent covering the range between 48cd/m² and 350cd/m² is likely to be useful, even if HDR is a separate version. It could be postulated that there might be an SDR master from 48cd/m² up to 350cd/m², and a second HDR master from 350cd/m² up to some appropriately high maximum (e.g. 3000cd/m²). On the other hand, perhaps it is possible to achieve a single master covering the entire range of SDR through HDR. Several functioning systems have been demonstrated at SMPTE, HPA, and elsewhere to support this (see Figure 19, 20).



Figure 19: Gary Demos presenting to Don Eklund at HPA Feb 2020, showing a single master being displayed over the range 150cd/m² to 2000cd/m².



Figure 20: Single Master demonstration architecture shown at HPA Feb 2020.

Color Appearance Models

Candidate color appearance models exist such as CIECAM02, hdr-CIELAB, and hdr-IPT (see "Color Appearance Models", Third Edition, Mark D. Fairchild, 2013 John Wiley & Sons).

Each of these candidates also relies upon the CIE 1931 Color Matching Functions (CMFs). In practice CMFs vary with subtended angle, age, and vary with individuals.

Each of these candidates uses the Michaelis-Menten "hyperbolic" function, which is asymptotic (goes flat) at some maximum value. The flat asymptote will not typically be invertible. However, a quasi-logarithmic representation such as "intQL" provides an invertible representation suitable for HDR (see Figure 21).



Figure 21. Invertible Quasi-Logarithmic intQL HDR function based upon scaled ACEScct (ACES S-2016-001). Note that the range extends to 55.725, about six stops above scene diffuse white.

There are also other problematic appearance attributes used by color appearance models. Hue angle and "hue quadrature" are unstable for dark colors. Neither hue nor hue quadrature could be considered reliably invertible.

The more elaborate color appearance models utilize achromatic and chromatic adaptation. Adaptation is an adjustment in our perception of the colors we are seeing (e.g. perceiving 3200K and D65 both as being white in appropriate contexts). Adaptation adjustment is likely partial (e.g. seeing 3200K as slightly yellow, or D65 as slightly blue) with large screens, and likely varies based upon the specific on-screen (and moving) images. Varying partial adaptation is therefore likely the norm, and likely is impractical to quantify.

Also, some appearance affects apply only to "related" colors or "unrelated" colors (related colors are part of chromatic adaptation). It further does not seem practical to determine the "related" aspects within HDR scenes between colors, although this is an area of active research.

Use Case

A key feature of the single master concept is that the master embodies the creative intent for the appearance of every scene. However, the mastering display itself has a specific brightness range, gamut, and ambient surround. Further, there may be pixels within the master that are outside the brightness limit, dark limit, and gamut limit, that are not accurately visible on the mastering display.

Also, the final distributed presentation dark range, bright range, gamut, and ambient surround will typically differ substantially from this mastering configuration. An approach might be to run a color appearance model forward to obtain all relevant "appearance correlates" for the master display configuration, and then run the model backward using a different display configuration appropriate for that given display. In other words, use a round-trip color appearance model using each presentation display's specifics for the reverse trip. In practice, this has inversion problems using existing color appearance models.

Also, achromatic and chromatic adaptation is affected by screen size. If the presentation is on a 75" screen, for a show mastered on a 30" screen, the degree of adaptation will likely differ. The degree of adaptation is also likely affected by the moving scenes and what they contain at various locations on screen. This will likely be different for differing screen sizes as well as differing on-screen images.

Architecturally it should be noted that many issues affecting appearance (such as brightness range) are only known at the presentation display after distribution, and are not known during creation of the master using the reference HDR display.

These are a substantial set of issues that arise when attempting to adapt or modify color appearance models to meet the needs of a single master embodying a single creative intent.

There are, however, numerous principles of color appearance that can be applied. This is the basis of the currently functioning single master systems.

Hypothetical Viewer

One approach that might be promising is to construct a hypothetical reference viewer. This is a concept similar to the CIE 1931 "Standard Colorimetric Observer" that attempts to be universal so that it can linearly represent brightnesses well beyond the capability of any human. In other words, rather than constrain appearance to a literal model of our vision (with our cones saturating), we can architect a vision model that is constructed specifically to meet the needs for an HDR-capable and WCG-capable single master architecture.

The first and foremost need is invertibility. Implicit in invertibility is the ability to accurately represent a full HDR and WCG range and wide gamut in a number of ambient surrounds and screen sizes.

This makes basic sense in that our goal for a single master is to embody creative intent over a wide range of display configurations.

Summary

Inherent limitations in color appearance models make them unlikely to meet the needs of a single HDR master. This is partly due to limitations in modelling the visual system, that functions only for a specific display configuration at any given time.

As an approach, a hypothetical viewer representation might be considered which sees the entire range of HDR and WCG display configurations in its appearance model. A

consequence of this approach is that the master itself will embody HDR creative intent for all presentation displays rather than for one mastering display being viewed during mastering.

Inquiries regarding the ASC Motion Imaging Technology Council should be sent to Alex Lopez: alex (at) theasc (dot) com