REAL TIME CROSS-MAPPING OF HIGH DYNAMIC RANGE IMAGES

J. A. Pytlarz, K. D. Thurston, D. Brooks, P. Boon, R. Atkins

Dolby Laboratories, USA

ABSTRACT

The ability to deliver high dynamic range (HDR) and wide colour gamut (WCG) imagery is crucial to next generation broadcast. It is a key feature of both DVB UHD-1 Phase 2 and the latest ITU-R recommendation: BT.2100.

While this is an important step towards the creation of broadcast HDR-WCG systems, if HDR-WCG production is to be deployed commercially, it is necessary to use a mix of both conventional standard dynamic range (SDR) and HDR cameras in a single HDR-WCG production. It is also necessary to derive a high quality conventional ITU-R BT.709 (SDR with gamma non-linearity) programme for regular contribution and transmission. Additionally, it is necessary to cross-map SDR programmes, interstitials and adverts into an HDR-WCG service for transmission.

This paper describes the techniques that have been developed to perform these transforms to meet broadcast production standards in real-time. These techniques are built on the experience gained in the creation of the first fifty HDR theatrical releases, as well as trials with HDR broadcast productions. Finally, the operational practices to ensure consistency in HDR-WCG production, high quality programme interchange, and a pleasing viewer experience are examined.

INTRODUCTION

Broadcast production today utilises a single colour volume workflow, as majority of footage is captured in one format: SDR (gamma non-linear curve and ITU-R BT.709 (1) colour primaries). With this single colour volume workflow, graphic overlays, adverts, and television programmes are routinely intermixed during production.

In the past year, availability of HDR and WCG televisions has grown substantially as has the demand for HDR-WCG imagery. To meet this demand, the broadcast production workflow must incorporate a multi-format pathway. It is likely that a mix of both SDR and HDR-WCG cameras will be in live production for several years. Additionally, the incorporation of archive SDR material in HDR-WCG production will be required for the foreseeable future.

The colour primaries (red, green, and blue) used in a WCG system are the ITU-R BT.2100 (2) colour primaries. The difference between BT.709 and BT.2100 colour primaries is shown in a chromaticity plot in Figure 1.

As specified in the ITU-R BT.2100 recommendation, HDR will replace the BT.709 gamma curve with an HDR non-linear curve: PQ or HLG. Unlike the traditional BT.709 gamma curve, PQ is an absolute encoding, meaning each code value relates directly to an expected luminance output on a reference display, in a reference environment. In that reference environment, whether a display can produce 500 cd/m^2 or 3000 cd/m^2, skin tones will remain constant around 26 cd/m^2.

The difference between SDR and HDR-WCG reproduction can be better visualised as a colour volume (the colour gamut at all reproducible luminance levels). All colours that a display can replicate are contained within this volume. Figure 2 is a comparison between the BT.709 and the BT.2100 colour volumes. The larger HDR-WCG colour volume not only increases headroom for highlight and shadow detail, it also allows for increased brightness of highly saturated colours. BT.709 primary blue has a maximum brightness of just 7.3 cd/m^2 in a 100 cd/m^2 peak white system. Increasing the maximum luminance of the display increases the possible maximum luminance of highly saturated colours, allowing blue to reach 593 cd/m^2 in a 10000 cd/m^2 peak white system. It should be noted that the HLG system proposed in BT.2100 will not utilize the full BT.2100 colour volume due to the OOTF application on luminance in a display (limited to 103 cd/m^2 with a 10000 cd/m^2 peak white).

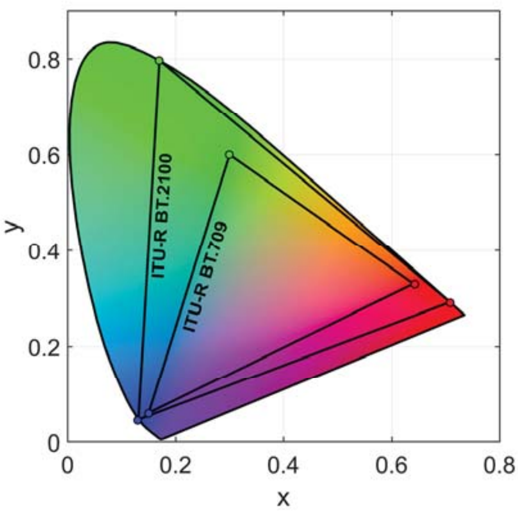

Figure 1 – Colour gamut comparison

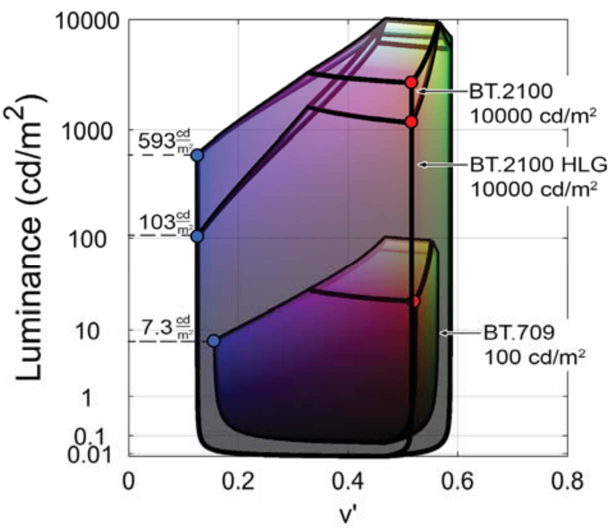

Figure 2 – Blue primary comparison

During production, HDR-WCG and SDR imagery will be intermixed, derived, and converted from one source to another. Colour volume mapping is the process used to convert between SDR and HDR-WCG with the goal of preserving hue, saturation, and luminance. ITU-R BT.2100 (2) recommends two colour representations, IC_TC_P (3) and $Y'C_BC_R$, for efficient international programme exchange. To avoid costly colour conversion in a real-time system, one of these colour representations may be used to perform all operations, including colour volume mapping. Efficient and high quality mapping (including round trip) is essential to successful real-time cross-mapping of SDR and HDR-WCG imagery.

MERGING SDR AND HDR-WCG IMAGERY

In a traditional SDR workflow, all content has the same colour primaries and non-linearity. Eventually an HDR-WCG workflow will have this same property. However, even as the availability of HDR-WCG imagery increases, SDR imagery will remain a common input source for years to come. Therefore, in the transition to an entirely HDR-WCG workflow, a hybrid system must be implemented to merge these two formats into a single container solution from which a high quality SDR programme, that either matches or exceeds current

SDR programme practice, can be derived. Once inside this common container, standard broadcast switchers and other tools will work equally well as with traditional SDR sources. In addition, post-production tools may be employed as necessary (colour correction/LUTS, blends, and graphic overlays).

What should this common container be? For perfect compatibility with SDR displays, an SDR container would be the logical choice. However, by limiting the container size to SDR, the HDR-WCG experience will never reach the home and many of the advantages of capturing HDR-WCG content is lost. In order to maintain the highest quality imagery, the common container should be large enough (wide enough primaries and high enough dynamic range) to hold all incoming formats. This way, no information is prematurely lost during the merge. The largest container available should be chosen: IC_TC_P.

Figure 3 is an example merge of SDR and HDR-WCG content inside the IC_TC_P container. The HDR-WCG content already has an HDR tone curve and wide colour primaries, so a 1:1 transfer into the IC_TC_P colour representation is appropriate. The SDR content has significantly less dynamic range and narrower colour primaries than the IC_TC_P container allows, therefore the SDR content may be mapped into the IC_TC_P container in a variety of ways. In this example, it is mapped with a 1:1 ratio (assuming 100 cd/m^2 peak white). Tone mapping maximum SDR to 100 cd/m^2 maintains the SDR standard (4) and maximises backward compatibility. As shown in the example, colour gamut mapping can also be a straight forward 1:1 map, placing the BT.709 content inside the IC_TC_P container, ensuring the content information remains unchanged. In practice, the broadcaster may choose to alter saturation and luminance placement to best fit his needs.

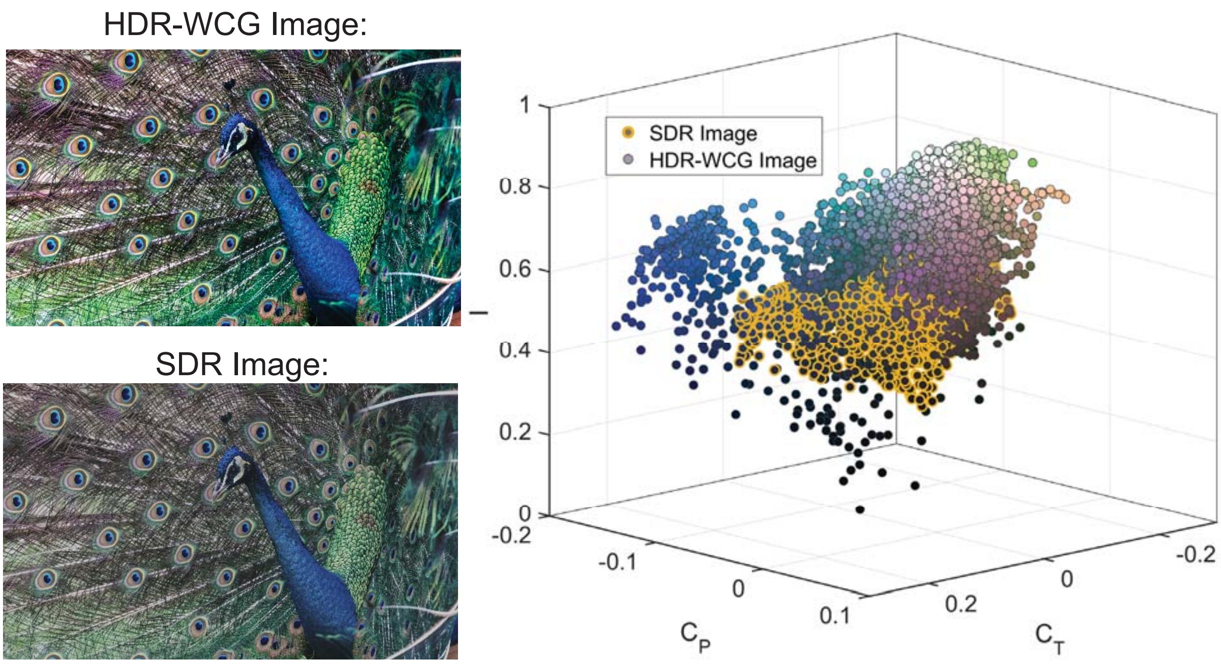

Figure 3 – SDR vs HDR-WCG image

While making the tone and gamut mapping decisions, the cross-mapping capabilities of the workflow should be considered. SDR input/output (and HDR-WCG input/output) should yield near visually identical results. Therefore, when making choices of where to place and if to alter content, be aware of the assumptions made during the inverse down-map and take this cross-mapping result into consideration.

From trials thus far (in a reference viewing environment), maintaining the 100 cd/m^2 reference level works well during live HDR-WCG broadcast. In addition, straight conversion from the BT.709 colour primaries into the IC_TC_P container (no up-mapping/stretching) provides anticipated results and allows systems to predict the images, and recover the original SDR signals. If a custom brightness level or colour re-mapping algorithm is applied, additional information about the SDR to HDR-WCG mapping may be required for inversion.

DERIVING SDR FROM AN HDR-WCG SIGNAL

Even after the transition to an entirely HDR-WCG workflow, SDR displays will remain in the market for a number of years to come. In order maintain perceived hue, saturation, and brightness as much as possible through the SDR mapping, all content should first be placed in a container that is uniformly spaced and has constant hue and luminance. To minimise processing power and time, this container should also be an efficient encoding space. This way the incoming signal will not require colour transformation before or after colour volume mapping for the purposes of efficient signal interchange. The IC_TC_P colour representation has all of the colour attributes necessary for high quality colour volume mapping and is also an efficient signal interchange format (5). Therefore, as has been shown in practice, the IC_TC_P colour representation may be the most appropriate container to enable real-time cross-mapping of SDR and HDR-WCG imagery.

Colour volume mapping must take two attributes into consideration: tone mapping and gamut mapping. The large dynamic range and wide colour primaries of HDR-WCG content must be reduced to fit inside SDR display capabilities.

Tone Mapping

A tone curve addresses the change in luminance from the HDR source to the SDR display. It is most commonly an "S"-shaped curve that rolls off the shadows and highlights to minimise detail loss. Figure 4 shows one example of HDR to SDR tone mapping. This curve is applied to the luma channel of a luma/chroma separated system such as IC_TC_P. Assuming this system has no crosstalk between chroma and luminance, adjusting luma (I of IC_TC_P for example) alters the brightness of the image without compromising hue or saturation. Care must be taken when tone mapping inside systems that have crosstalk between chroma and luminance, such as $Y'C_BC_R$, as colour shifts may occur.

The design of the tone mapping curve (choice between highlight/shadow roll-off and contrast) may be modified after analysing image content. If all content lies within the SDR range, tone mapping may not be required. On the other hand, if a large portion of the HDR-WCG container is being utilised, a form of slow roll-off (depending on image content and artistic choice) should be applied.

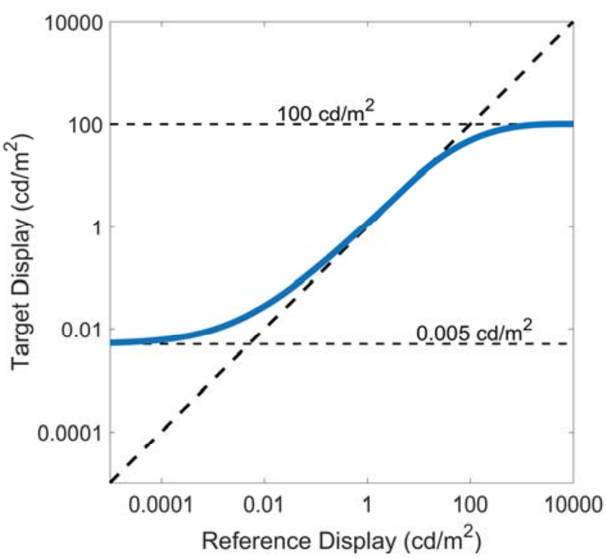

Figure 4 – Example Tone Mapping

Gamut Mapping

Once the tone mapping is complete, colours may still be more saturated than the BT.709 container allows. Various methods exist for mapping out-of-gamut colours into the BT.709 container including clipping out-of-bounds colours to the gamut boundary. Although clipping may be computationally cheap and fast, it may also introduce artefacts such as colour shifting and detail loss.

An alternative method of gamut mapping is applying a saturation knee function near the boundaries of the BT.709 colour gamut (6). As seen in Figure 5, this method squeezes colours near the gamut boundary to make room for those out-of-gamut colours. In this way, saturated colour detail is better preserved. If the image content does not fully utilise the BT.2100 gamut initially however, and image content is not analysed, unnecessary gamut mapping for scenarios where content already lies within the BT.709 gamut boundary will occur.

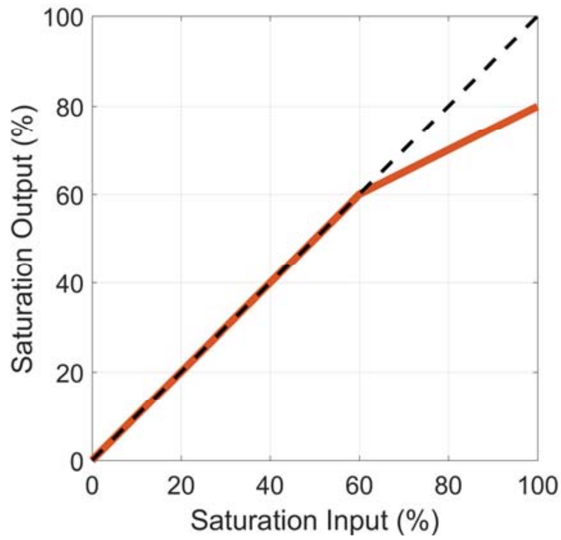

Figure 5 – Example gamut mapping

Implementation

As demonstrated by Grass Valley and Dolby Laboratories at NAB 2016, these generalised colour volume mapping techniques (Figure 6), utilising the IC_TC_P container, are efficient and accurate enough to be computed live without any special data being passed through standard broadcast equipment and infrastructure. Because of this, the SDR rendered content can be displayed live in the production facility for quality assurance, while retaining an uncompromised HDR-WCG signal through the broadcast chain.

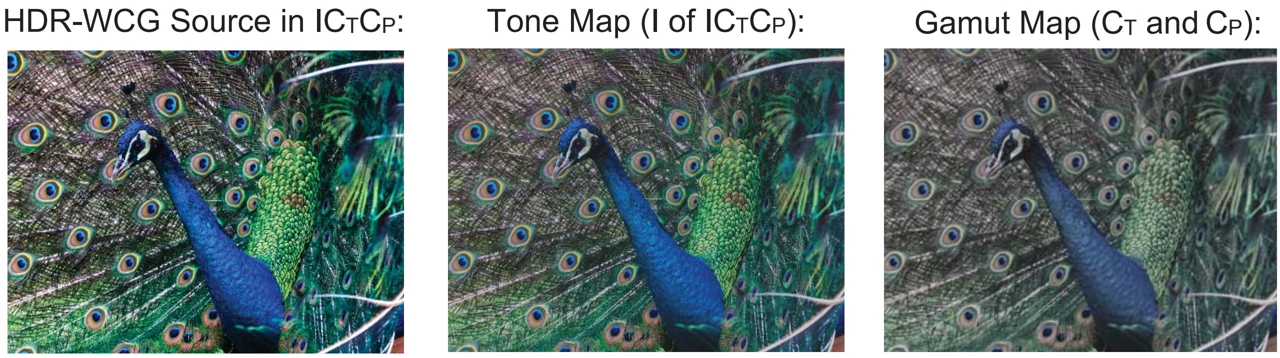

Figure 6 – Example colour volume mapping from HDR-WCG to SDR

Due to the speed and convenience of producing a live map of HDR-WCG to SDR, minimal workflow adjustments are required to produce HDR-WCG content. The rendered SDR signal may be monitored on traditional multi-view monitors, scopes, and other tools. Tests by broadcasters indicate that if the rendered SDR signal is in range, the more flexible HDR-WCG signal most likely is as well. Therefore, initially only a few critical viewing HDR-WCG monitors need be deployed. Furthermore, HDR-WCG training for camera operators is reduced because traditional SDR scopes, monitoring tools, and controls remain unchanged.

IC_TC_P COLOUR REPRESENTATION

The colour representation (container) chosen to perform the tone and gamut mapping inside plays a vital role in the quality of the SDR render. For example, tone mapping the BT.2100 red, green, and blue channels independently may result in significant hue shifts due to the native crosstalk between hue and saturation. Ideally, the chosen colour representation has entirely de-correlated hue, saturation, and luminance to allow control of each channel independently.

An additional consideration is the ability to reduce colour space conversions during production and transmission to minimise processing power and time. Preferably no colour space conversions are necessary. To accomplish this, mapping may be applied in either $Y'C_BC_R$ or IC_TC_P as those are the two format recommendations put forward by the ITU-R for programme or signal interchange of HDR-WCG content (7).

The IC_TC_P colour representation is both an efficient transmission space and a quality colour volume mapping space (8). It holds similar properties as $Y'C_BC_R$ (chroma/luma separation for chroma subsampling and compression performance), yet has the added benefits of hue linearity and de-correlation of luminance and chroma. To analyse performance of colour volume mapping in $Y'C_BC_R$ or IC_TC_P, decorrelation of hue, saturation, and luminance along with uniformity of code word distribution is examined.

Chroma and Luminance Decorrelation

During tone mapping, the luma component (I of IC_TC_P or Y' of $Y'C_BC_R$) is mapped to the appropriate display range. During this process, chroma should remain constant. Likewise, during gamut mapping, the hue and saturation (chroma) components are adjusted and luminance should remain constant. For this to happen, chroma and luminance must be decorrelated (no crosstalk).

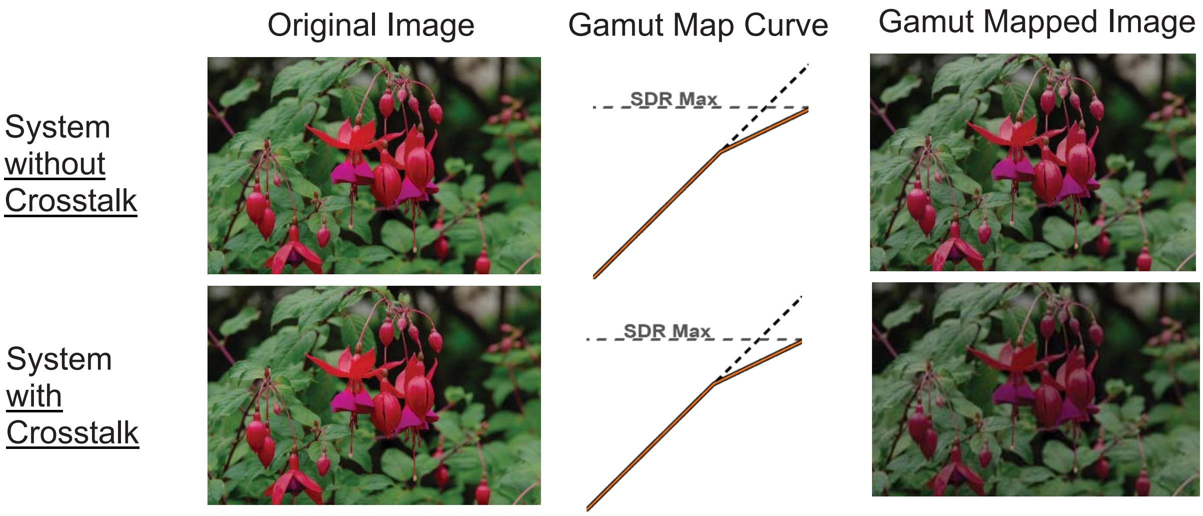

Figure 7 – Tone mapping and crosstalk effects

Figure 7 demonstrates the consequence of gamut mapping in a colour representation that has chroma/luminance crosstalk. Gamut mapping in a system with crosstalk (bottom right) causes luminance changes when chroma is adjusted (flower gets darker). A system without crosstalk (top right) maintains the luminance of the original scene.

To have the least amount of chroma/luminance crosstalk, the mapping space should be near constant luminance. To analyse the constant luminance attributes of each colour representation, a cube of BT.2100 colours from 0-10,000 cd/m^2 is converted to $Y'C_BC_R$ and IC_TC_P. The luma components are compared to PQ encoded luminance. A perfectly constant luminance space has a 1:1 ratio.

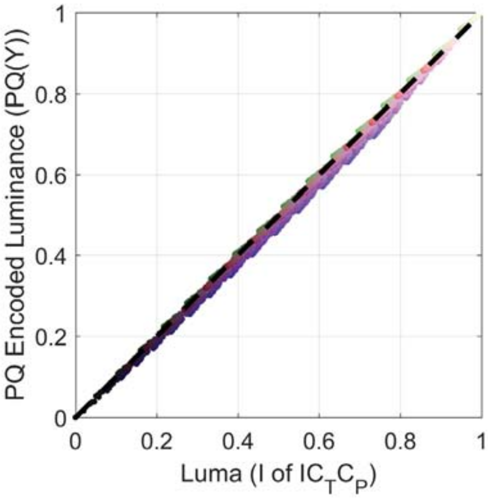

Figure 8 - IC_TC_P luma vs. luminance

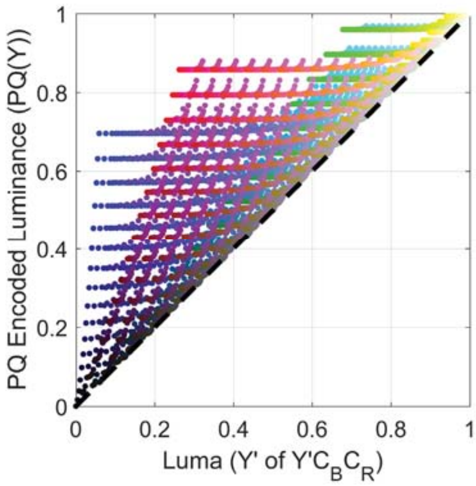

Figure 9 - $Y'C_BC_R$ luma vs. luminance

As seen in Figure 8, IC_TC_P is nearly constant luminance. Therefore, during tone mapping, saturation and hue are not affected as luma (I) is changed. In addition, during gamut mapping, luminance is not affected as saturation and hue are changed. As seen in Figure 9, $Y'C_BC_R$ has crosstalk between chroma and luminance. Therefore, with luma changes during tone mapping, the chroma may also be affected. In addition, with chroma changes during gamut mapping, the luminance may be affected. Therefore, more distortion between hue, saturation, and luminance will occur when colour volume mapping inside $Y'C_BC_R$.

Saturation and Hue Decorrelation

Gamut mapping applies adjustments to the saturation component. As desaturation is applied, hue should remain constant. For this to be the case, the colour representation must be hue linear (decorrelated saturation and hue). To analyse the hue linearity of $Y'C_BC_R$ and IC_TC_P, the Hung and Berns (9) constant hue data set is used. In Figure 10 and Figure 11, the solid colour points are the reference points used in the matching experiment.

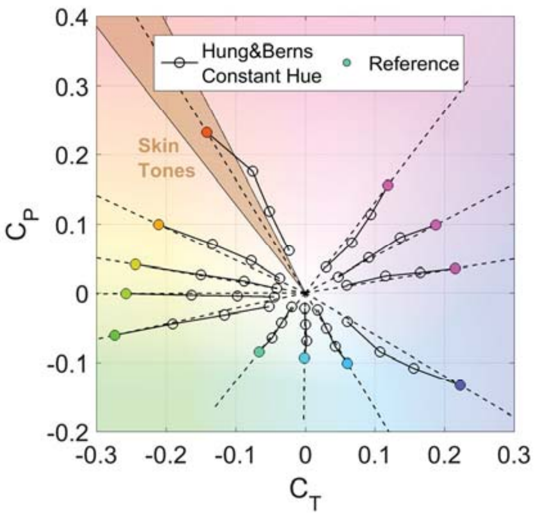

Figure 10 - IC_TC_P hue linearity

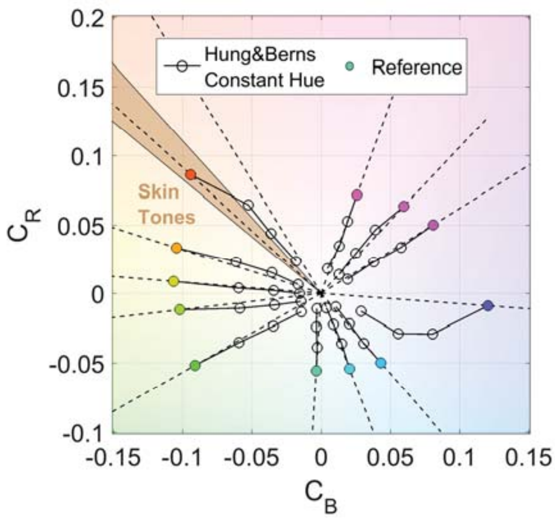

Figure 11 - $Y'C_BC_R$ hue linearity

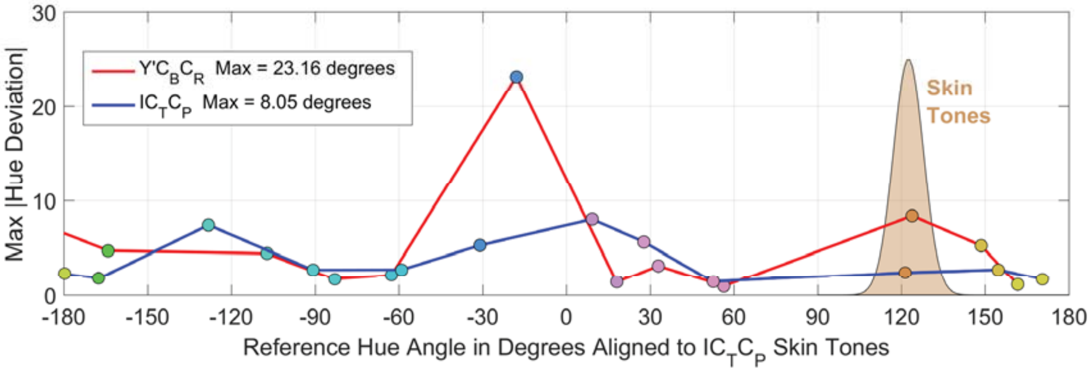

Figure 12 – Hue linearity analysis

Compared to these reference points, three test patches at each hue (desaturated in increments of 25%) were matched by an observer. The observer was asked to adjust the hue to match the reference. The experiment results are the solid lines. The dotted lines are the predicted hue in each representation. Where the solid and dotted lines differ, is where hue shifts may occur during gamut mapping.

As seen in Figure 12, the hue angle deviation in IC_TC_P is significantly less than in $Y'C_BC_R$, especially in the blue region. Compared with $Y'C_BC_R$, IC_TC_P reduces the maximum hue angle deviation by nearly three times. Therefore, using the IC_TC_P colour representation for gamut mapping will result in less average hue deviation than using $Y'C_BC_R$.

Code Word Uniformity

During gamut mapping, it is expected that every code word has the same perceptual difference. Changing saturation in red and blue by 50% should result in an equal perceptual decrease in saturation. This may not be the case, however, if more code words are allocated to certain areas of the colour representation.

To analyse code word uniformity, MacAdam ellipses (10) were plotted (scaled 10x) inside IC_TC_P and $Y'C_BC_R$. MacAdam ellipses were generated from an experiment to find the just-noticeable-difference (JND) around a reference point. A perceptually uniform colour representation would render these ellipses as circles, equal in size.

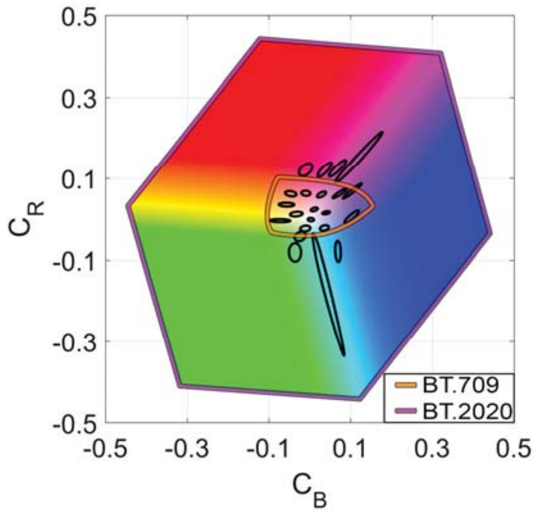

Figure 13 – $Y'C_BC_R$ uniformity

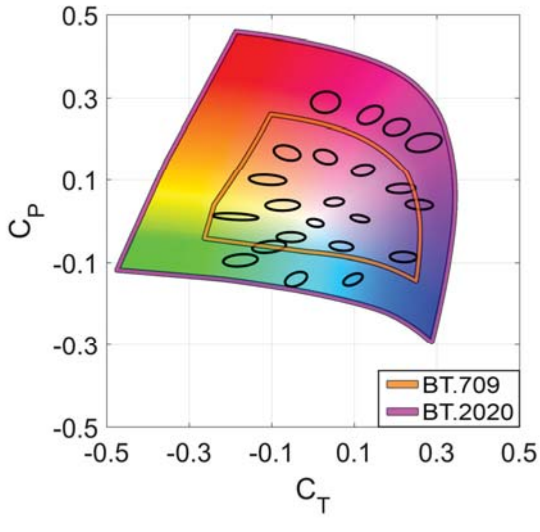

Figure 14 – IC_TC_P uniformity

As seen in Figure 13, $Y'C_BC_R$ renders the MacAdam ellipses very different in size and shape with the worst case being the cyan and magenta regions. Due to the long skinny nature of the ellipses in these regions, saturation can be decreased significantly without any visible change. The BT.709 region of $Y'C_BC_R$ does not include these long ellipses, which is why SDR is not largely affected by this characteristic. IC_TC_P, as seen in Figure 14, renders near circular ellipses very similar in size. This indicates that every step in IC_TC_P is nearly perceptually equivalent.

Image cross-fades illustrate this effect as seen in Figure 15. When blended 50%-50% in $Y'C_BC_R$, the magenta colour from the flowers and the cyan colour from the peacock overpower the crossfade. This is due to the large perceptual allocation of code words in saturated regions. Blending in uniform IC_TC_P results in perceptually equal weighting of saturated and desaturated colours. Gamut mapping inside the IC_TC_P colour representation therefore, will result in expected and even perceptual changes.

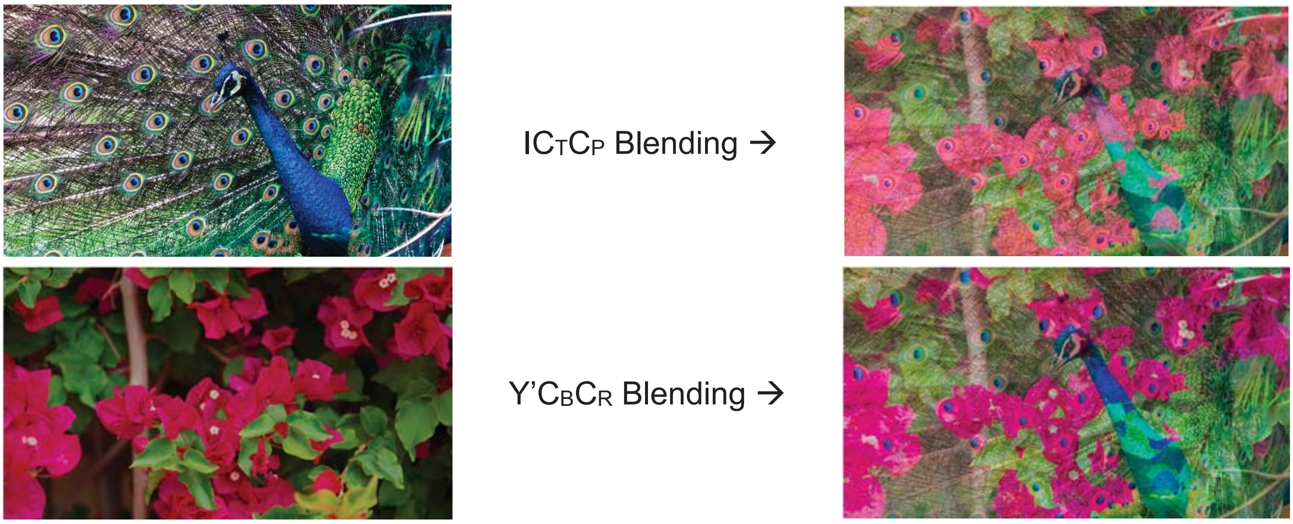

Figure 15 – Uniformity and blending

Practical Application and Benefits

With the introduction and adoption of BT.2100, a new colour representation must be implemented: either $Y'C_BC_R$ or IC_TC_P. For efficiency, not only will this colour representation be used for bandwidth reduction, but also for merging SDR and HDR-WCG content. Mapping from HDR-WCG to SDR involves both luminance and colour reduction. Using IC_TC_P for these operations will better preserve hue, saturation, and brightness than $Y'C_BC_R$ will. Gamut mapping in $Y'C_BC_R$ will result in luminance distortion. Likewise, tone mapping in $Y'C_BC_R$ will cause hue distortions. IC_TC_P has the same mathematical operations as $Y'C_BC_R$, making it simple to implement, and has the benefit of significantly reducing distortions caused by $Y'C_BC_R$ crosstalk and non-uniformity.

CONCLUSION

With the rising number of HDR-WCG displays in the market, availability of HDR-WCG content is rapidly increasing. The need to merge SDR and HDR-WCG content is inevitable. To avoid the massive cost of a complete system overhaul, as HDR-WCG cameras, adverts, and interstitials enter broadcast production, a single-stream, merged path of HDR-WCG and SDR content is economically necessary. Given the real-time algorithms outlined, from this single HDR-WCG stream, uncompromised BT.2100 HDR-WCG may be delivered, and where necessary, high quality BT.709 SDR may be derived. This enables a gradual workflow transition as few changes are immediately required in a broadcast pipeline.

Broadcast companies have a corporate identity and brand invested in the way their imagery appears, therefore, the image quality of the HDR-WCG single-stream should be preserved as much as possible during the colour volume mapping to BT.709 SDR. The decorrelated hue, saturation, and luminance as well as the uniform code word distribution of the IC_TC_P colour representation allows for predictable change plus fine-tuning. Using the IC_TC_P colour representation for colour volume mapping enables high-quality mapping that best preserves the HDR-WCG source.

REFERENCES

1. International Telecommunication Union (ITU), 2015. Parameter values for the HDTV standards for production and international programme exchange, Recommendation ITU-R BT.709-6, 2015.

2. International Telecommunication Union (ITU), 2016. Image parameter values for high dynamic range television for use in production and international programme exchange, Recommendation ITU-R BT.2100-0.

3. Dolby Laboratories, 2016. What is ICtCp?- White Paper, Dolby Laboratories < http://www.dolby.com/us/en/technologies/dolby-vision/ICtCp-white-paper.pdf>.

4. International Telecommunication Union (ITU), 2011. Reference electro-optical transfer function for flat panel displays used in HDTV studio production, Recommendation ITU-R BT.1886. March 2011, pp. 3.

5. Lu et al., 2016. ITP Colour Space and Its Compression Performance for High Dynamic Range and Wide Colour Gamut Video Distribution, ZTE Communications February 2016, Vol. 14 No. 1, pp. 32-38.

6. Fairchild, M. D. and Braun, G. J., 1999. General-Purpose Gamut-Mapping Algorithms: Evaluation of Contrast Preserving Rescaling Functions for Color Gamut Mapping, IS&T/SID Seventh Color Imaging Conference: Color Science, Systems and Applications (1999), pp. 167-172.

7. International Telecommunication Union (ITU), 2016. High dynamic range television for production and international programme exchange, Recommendation ITU-R BT.2390-0.

8. Pytlarz, J. A., 2016. Versatile Color Space for High Dynamic Range and Wide Gamut Signal, National Association of Broadcasters Show 2016 Conference Proceedings.

9. Hung, P. and Berns, R., 1995. Determination of Constant Hue Loci for a CRT Gamut and Their Predictions Using Color Appearance Space, COLOR research and application. Volume 20, Number 5, October 1995.

10. MacAdam, D. L.,1942. Visual Sensitivities to Color Difference in Daylight, Journal of the Optical Society of America, 32, May 1942, pp. 247-274.

ACKNOWLEDGMENTS

Thanks to Timo Kunkel for providing all of the HDR-WCG source imagery used in this paper.