TUTORIAL

KEYWORDS COLOR MANAGEMENT // COLOR WORKFLOWS // VIRTUAL PRODUCTION // COLOR PIPELINES // LED PANELS

Color Management Principles for LED Panels in On-Set Virtual Production

By David Long, Rod Bogart, Spencer Stephens, Catherine Meininger, Joseph Goldstone, Tim Kang, Kathryn Brillhart, and Jim Geduldick

Abstract

The goal of color management is simple: a displayed image should look the same wherever it is viewed: on-set, on production monitors, in visual effects (VFX) on an artist's desktop monitor, in the Digital Intermediate (DI) suite on a reference monitor, on a cinema's screen, and so on. If the image is displayed correctly and consistently, color management maintains the filmmaker's creative intent. If a filmmaker approves an image color from a display that is not calibrated to a known standard, there is a potential that the colors seen on that display cannot be communicated correctly to image consumers downstream. Correcting these image color errors may be challenging, time-consuming, and expensive. These errors may accumulate, with the image colors deviating further from those originally established as the image proceeds through the uncalibrated workflow. Color management is not a novel concept in creative industries. Numerous tools and best practices have been developed for decades, with modern technologies requiring new techniques and workflows to maintain creative intent. Virtual Production (VP) techniques are some of the latest additions to a filmmaker's toolbox, surfacing new questions and needing new recommendations to maximize effectiveness. This paper from SMPTE Rapid Industry Solutions (RIS) provides introductory material to support future efforts that focus on specific color workflow applications using LED panels for On-Set Virtual Production (OSVP), such as performing color calibration for in-camera visual effects (ICVFX) and image-based lighting. This paper covers color engineering principles that are well-known in traditional color workflows but expands upon their new roles within OSVP color workflows and their challenges.

The goal of color management is simple: a displayed image should look the same wherever it is viewed: on the set, in VFX, in the DI suite, in a theater, and so on. If the image is displayed correctly and consistently, color management maintains the filmmaker's creative intent. If a filmmaker approves an image color from a display that is not calibrated to a known standard, colors seen on that display may not be communicated correctly to downstream artists and final image consumers. Correcting these image color errors may be challenging, time-consuming, and expensive. The errors may accumulate, with the image deviating further from the originally established colors as it flows through post-production.

In a poorly managed color workflow, the delivered color of digital backgrounds, set extensions, skin tones, wardrobe, and so on will lose the harmonious appearance created on-set. Which version of the changing image is correct? This leads to disputes over where the colors deviated from those creatively approved.

On the other hand, a properly managed pipeline preserves intent from creative selection to final delivery—the colors remain as chosen.

Color management is not a novel concept in creative industries. Numerous tools and best practices have been developed over many decades, and modern technologies require new techniques and workflows to achieve the same goal of maintaining creative intent. Virtual Production (VP) techniques are some of the latest additions to a filmmaker's toolbox, bringing along new questions and the need for new recommendations that help to maximize effectiveness.

SMPTE formed the On-Set Virtual Production (OSVP) Rapid Industry Solutions (RIS) initiative in response to these new challenges. "On-set" describes VP techniques involving real-time interaction of traditional production tools with virtual and augmented reality (VR/AR) tools, computer-generated

1545-0279/25©2025SMPTE

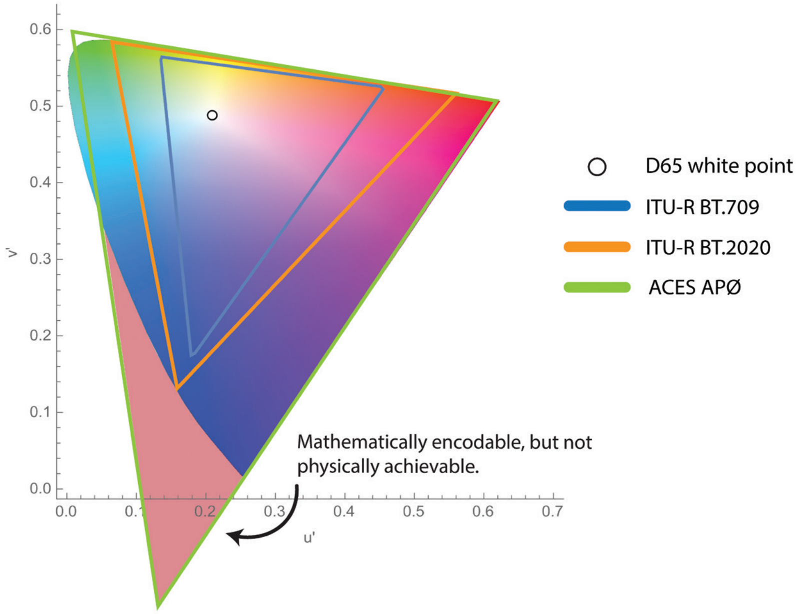

Mathematically encodable, but not physically achievable.

FIGURE 1. Chromaticity coverage of two ITU-standardized color gamuts for broadcast applications and one SMPTE-standardized gamut used for color management and image interchange.

imagery (CGI), and game-engine technologies. These interactions on-set significantly impact color-managed workflows, as critical color control and evaluation common in post-production now occurs earlier and more frequently in the production process.

In virtual production, traditional display technologies (i.e., LED panels) do not present an image for human viewing but rather be captured directly as an in-camera visual effect (ICVFX) or used indirectly as image-based lighting fixtures to illuminate actors and objects to complete the immersion of a virtual scene. Digital assets are not rendered for the human eye, as they are in post-production, but rather for the camera imager. Camera imagers do not respond to light the same way as the human eye, requiring a different and unique image from the display panels. This exposes an intermediate stage of color management to creatives on-set that are not typically viewed, which can negatively impact color decisions if misunderstood or used incorrectly. Characterizing these interactions, understanding the characterization, and accommodating the panel-to-imager interaction as part of the conservation of creative intent is fundamental to color-managed virtual production.

Industry professionals experienced with color management concepts and the challenges of adapting them for OSVP have joined together to form a specialized working group within SMPTE RIS. Their combined experiences inspire authorship with goals to highlight the importance of color management, educate readers on color engineering principles, and suggest best practices for applying proper color management within an OSVP workflow. While the scientific principles of color management don't fundamentally change for OSVP workflows, there is a need for education regarding their particular application in OSVP and guidance for engineers and practitioners using the tools available.

This paper provides introductory material for color workflow components specific to using LED panels in OSVP, such as color calibration for ICVFX and image-based lighting. LED panels are not the only tool for OSVP, and most of the concepts in this paper apply to other technologies. However, the language and scenarios explored in these articles will be specific to LED technology.

Those new to color management or virtual production should use the references in the bibliography to understand the core concepts and terminology of digital color management.

Introduction

Those experienced with camera and display equipment traditionally used in production will appreciate the complexities of maintaining a color-managed workflow. Designing the camera's encoding color space and optoelectronic transfer function (OETF) to complement the display's color space and electro-optic transfer function (EOTF) is critical to delivering appropriate scene-to-screen color rendering to the viewer. Intermediate post-production encodings allow artists to manipulate the images efficiently and intentionally to produce the ultimately rendered deliverable. Production and display standards for color are commonly paired with native manufacturer-dependent transforms to give users extended control over gamut, dynamic range, and ultimate 'look' development in a comprehensive ecosystem. These imaging pipelines are well-engineered and can provide everything the creative team needs to deliver aesthetic intent to the viewer when they are well-understood.

OSVP color management extends standardized color-managed workflows to accommodate color and vision dynamics unique to photographing a computer-generated image directly from an LED-based VP screen. Some attributes of the OSVP color pipeline are directly manipulatable. Some are locked into a theoretical standards mentality unless operators manually make adjustments. Some attributes are the constrained capability of the equipment selected and cannot be fully adjusted.

There are several major color rendering paradigms in OSVP that productions must be intentional about. Within the Virtual Production Department, the virtual art department (VAD) is responsible for authoring colorimetrically-referenced content as a facsimile for, or at minimum, complement to, a real lit scene. While real scenes comprise actual dynamic light sources with definable spectral and geometric signatures interacting with scene objects that transfer perturbation of spectra and lighting geometry down the optical axis of the camera, VAD assets are a (usually) three-channel encoding of comparable physics in a defined color volume standard (demonstrated in **Fig. 1**, such as ITU. Rec. 709[1] or 2020,[2] or proprietary, such as manufacturer wide gamut or SMPTE ST 2065-1 [ACES][3]). And that encoding may or may not be linearly proportional to radiometry. The first premise of OSVP color management is VAD assets that are rendered to the LED volume are the scene physics to be metamerically and radiometrically captured and encoded faithfully by the camera. The camera-encoded values should equal the master VAD color values so long as color space and tone encoding are set equivalent. But, of course, the LED does not offer a perfect spectral analog of the real scene. Three-channel metamerism realized in the camera to match VAD encoding is possible, but only when careful characterization of the spectral emission of the LED is accounted for relative to the unique spectral sensitivity function of the camera. Since each LED product and camera are independently engineered, the transforms must also be applied to ensure the camera aligns with the VAD. The VAD ultimately creates original scenes using physics engines and rendering approaches that mimic real light-matter interaction. They also enjoy the benefit of defining as large a color volume or dynamic range as they wish, considering that asset authors can preview the colors correctly on a mastering display with faithful color volume reproduction. A computer-generated VAD asset is anything the production team envisions. That is their ultimate creative control. But they must understand and respect that the colors they've created could lie outside the gamut and dynamic range of the subsequent image encoding(s), LED volume, and/or camera the OSVP team uses. If these boundaries are well understood and respected, VAD and OSVP teams can absolutely employ sound color management, outlined next, to address spectral metamerism and radiometric linearity in the most common workflows (**Fig. 2**).

Before we proceed, however, we must acknowledge elements of color physics that may prove obstacles to delivering the best OSVP color outcomes. Some of these, too, might be out of the direct control of the creative team. First, no motion picture camera today replicates the response of the human visual system. Camera sensor spectral sensitivity functions are not linear transformations of human color matching functions; furthermore, there is acknowledged diversity amongst human observers not currently well accounted for in any mainstream motion picture workflow.[4-6] Many cameras have perturbed color rendering to yield more preferred results versus more accurate results, not to mention encoding spaces that don't always cover the entire human visual system gamut. In practice, camera color-encoding is usually engineered as a balance to get the most important scene spectra rendered colorimetrically right, at the expense of "best fit" for all conceivable spectra. There are color science-based approaches to match the colorimetry of practical foreground scene elements to that of the VAD renders on the

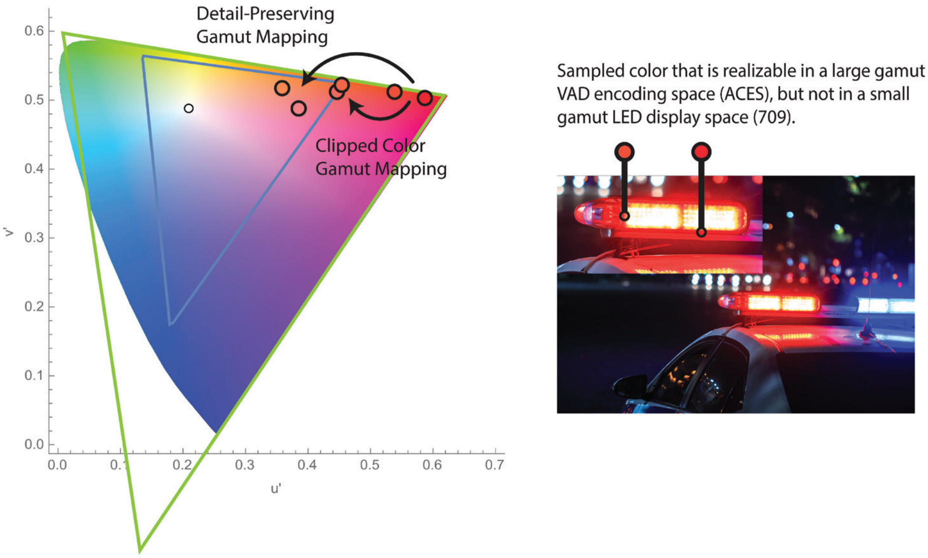

Sampled color that is realizable in a large gamut VAD encoding space (ACES), but not in a small gamut LED display space (709).

FIGURE 2. Example gamut mapping considerations that must be understood when displaying and capturing the content of a highly saturated color stimulus. Understanding the desired appearance of the color and reproducing it throughout each color workflow stage is critical to maintaining creative intent.

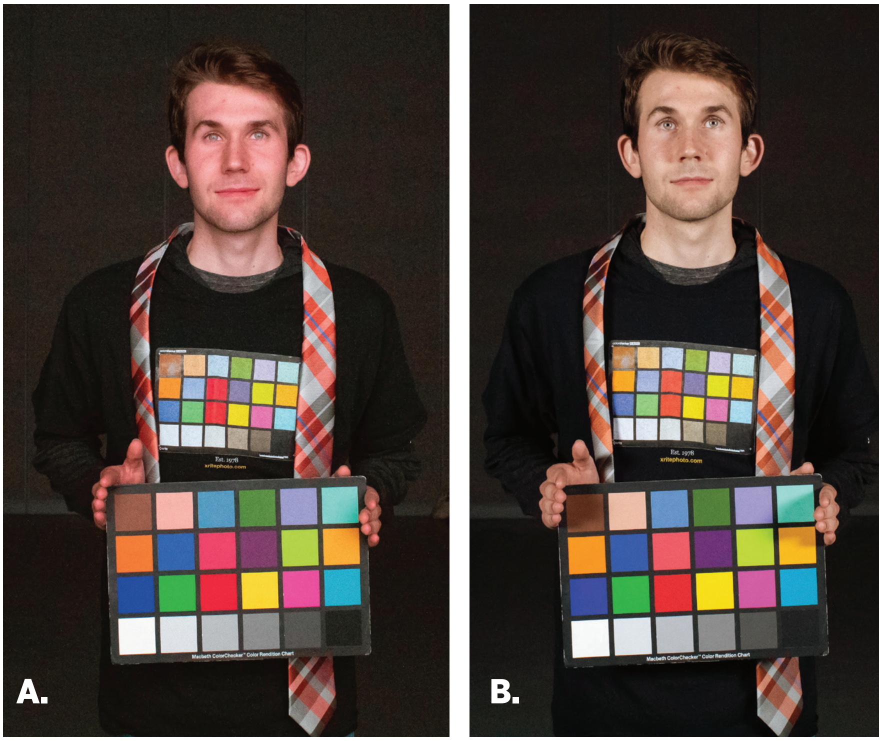

FIGURE 3. Demonstration of illuminant metamerism from two different lighting sources typically found on an LED volume stage. (a) Illumination from LED walls often results in more saturated color rendition of objects due to the narrowband spectral characteristics of the LED emitters. This is especially seen in objects that reflect significant amounts of long-wavelength energy from red LED components. (b) Light sources with a more broadband spectral characteristic do not exhibit this behavior.

LED volume, as captured in the camera. Such color management becomes important for applications requiring highly accurate color matching, such as set extensions or product placements. This topic is beyond the scope of this paper but is under discussion within the RIS-OSVP working group.

Second, one particularly unique aspect of OSVP versus more traditional production is the use of the LED emitters to directly or indirectly illuminate foreground actors and set objects. As part of the overall OSVP lighting design, three-channel LEDs can be adjusted to mimic the colorimetry produced by a modern cinema light—with anywhere from four to six different types of LED emitters—but not the spectrum. As such, the workflow suffers from illuminant metamerism failure (**Fig. 3**), a phenomenon where objects convey inaccurate color based on spectral peculiarities of their interaction with light sources of different spectral compositions but the same integrated white point or color temperature.[7] This topic, too, will not be covered in this paper.

Ultimately, the goal is to make scenes that "match." This is easy to say, but specificity is required to define the criteria for matching. For OSVP, the preference is for the final camera capture, when rendered for viewing on a reference monitor, to match the look of the VAD assets as seen in the real-time engine. To do this, the specific processing that happens to the final camera capture must be considered, as well as the methods used to view the VAD content while it is being created. Additionally, a great deal of color processing happens in the various pieces of hardware and software between the initial creation of the VAD assets and their final display on the LED wall. The following section will highlight the individual conversions in a typical OSVP color pipeline and discuss the parameters that must be set to ensure proper matching.

OSVP Color Pipelines for LED Panels

From a high level, there are three major sections to the OSVP color pipeline for LED panels (**Fig. 4**). The first is the creation of linear light values by a real-time engine. The goal of this step is to create pixel values that represent a virtual scene as if it were a real scene (hence the commonly used term "scene-referred colorimetry"[8]). The second step is the display of those linear values on the LED wall. The goal of this step is to produce linear light from the screen that presents that virtual scene to a camera system. The third step is to capture that virtual scene with the camera system. The goal of this step is to capture the virtual scene on the wall and the real objects in front of the wall in a way that best represents the entire scene. Caveats and details for the above steps are discussed below. This discussion will be focused on colorimetry; for overviews of the general OSVP color pipelines, please refer to the references, e.g., the "VES Handbook of Virtual Production."

Step 1: Creation of Virtual Scene Linear Values

Color calculations are performed in a specific linear RGB color space in a real-time engine. For much of the industry, the primaries and white point of that color space come from ITU-R BT.709 (also referred to as '709', the primaries and white

1545-0279/25©2025SMPTE

point of which are equivalent to those of IEC-61966-2-1[9] or 'sRGB'). In more VFX-heavy workflows, another common linear color space with greater gamut capacity is ACEScg. The lighting and shading calculations are performed using linear values, which may be mapped directly to real-world equivalent colorimetric units such as CIE 1931 2-degree standard observer tristimulus values or 'XYZ.' The choice of one space over another may impose gamut limitations on the creative team. A traditional RGB color space comprised of physically realizable primaries necessarily excludes regions of color space that are achievable with real stimuli. Color scientists such as Dr. Michael R. Pointer have executed several statistical studies to characterize the extent of real-world material colorimetry, and traditional capture systems have emphasized accurate encoding of these reference stimuli.[10]

In OSVP, VAD artists may wish to creatively push the boundaries of scene colorimetry beyond the traditional Pointer gamut. These excursions from classic scene behavior must be considered when specifying real-time color calculations and when considering capture encoding in camera. It is not a requirement that the working color space of the rendering exactly matches the color space of the output devices to be used. However, appropriate signal conversion must be applied to achieve the ultimate goal of a match between the displayed camera output and VAD assets. Gamut mapping compromises must be expected should any VAD color lie outside the physical capability of either the LED wall or the camera encoding, as demonstrated in **Fig. 2**.

The purpose of the linear rendering is to create a virtual scene that can be used as a dynamic background on an LED wall or volume. During the creation of this scene, the VAD will typically view the scene on a color-calibrated desktop monitor. They may additionally use a "look" transform that emulates the eventual look that will be applied to the camera output for the final presentation of the finished show. In a completely color-managed workflow, it is preferable to have advanced knowledge of the eventual "look" transform that will be used at the end. This treatment is demonstrated by including a "show look" transform in the beginning steps of **Fig. 5**. However, this is often not available during asset creation, so the VAD artists use the default look of the real-time system instead. Since the goal of color management is to provide a workflow that allows matching from end-to-end, using different looks for the VAD and the final footage will cause different interpretations of the scene, which will cause an otherwise managed pipeline not to match the final creative intent.

To alleviate this issue, the match can instead be confirmed on the linear output of the VAD with the linearized camera output. To do this, a specific color space must be chosen for the comparison. While it would be best to use the real-time engine's working color space, an alternative color space can be used as long as it is also linear and a known conversion exists to a traditional colorimetric connection space such as CIE XYZ. If the comparison color space has a smaller gamut than the working color space of the real-time engine, then care must be taken to observe any clamping or otherwise nonlinear mapping that may occur due to the conversion to the smaller gamut.

The two images that will be compared were created in quite different ways. The first is the linear version of the virtual scene, which exists during the lighting and shading in the real-time engine. The second image is the result of the camera capturing that scene rendered to the LED wall.

The first image, from the real-time engine, is not the scene with the aforementioned "look" applied; it is the input to that "look" step. It has the properties that it is linear, represents real-world objects, and has a known working color space. It also has an overall exposure that positions the image data at a reasonable level for the wall and the final output. These specific characteristics are differentiated in the first section of **Fig. 5**.

The second image, from the camera, is typically stored in a vendor-specific format that involves a known encoding color space and OETF. Once that encoding space and encoding curve are converted to the chosen linear comparison space

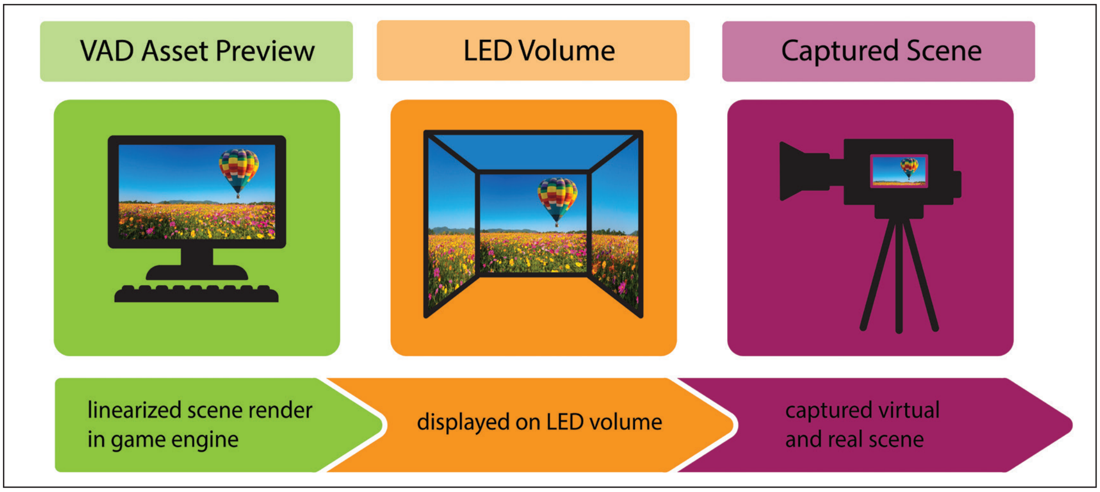

FIGURE 4. High-level overview of the three major sections of an OSVP color workflow for LED panels. A more detailed breakdown of specific operations can be seen in Fig. 5 and described in more depth throughout this paper.

through published methods, the image will be in the same color space as the first image, and comparisons are possible. This comparison is demonstrated in the overall workflow diagram provided in **Fig 5**. Because the image is the result of photography, some artistic choices may be made that alter the default or nominal behavior of the camera, such as exposure rating, the addition of internal or external ND filters, or possible changes to the white balance. These are on top of the spectral response of the imager and the processing of demosaiced photosite data into the camera's native color encoding that makes perfect colorimetric capture impossible (due to gamut restrictions or nonlinear spectral integration limitations). Some optical phenomena, such as lens flare, will transform the spectral power distribution produced by a well-calibrated LED wall into a different spectral power distribution impinging on the sensor. VAD blacks may be lifted during capture simply by the physics of photographic flare.

The following sections will delve deeper into the color processing that occurs in creating each of these images. For the real-time engine, the color science is fairly simple because all calculations are done in a single linear working color space. Texture files may be provided in a different color space (for example, using sRGB textures in an ACEScg project), necessitating performing the transform between the two known spaces by linearizing the texture and then converting to the linear working space. The final step in real-time processing is converting the linear working space image into a signal that can be sent to the LED processors.

Step 2: Displaying the Virtual Scene on LED Walls

The creation of the final camera-captured image involves many more steps. The first is the conversion from the engine's linear working color space to the color space configured on the LED processor, and the second is a nonlinear encoding of the signal that again matches the configuration on the LED processor. Typically, the color space of the LED processor—the RGB primaries, white point, and EOTF—is a user interface parameter setting; therefore, it is one of the first color management options that must be appropriately set.

The role of the LED processor is to convert the signal from the engine into values that can be used by the LED panel drivers to create specific amounts and colors of light, ostensibly matching the colorimetric values specified by the VAD mastering space, but which may also require adjustment for cameras incapable of perfect colorimetric metamerism upon capture of the LED stimuli. System operators provide information about the input signal, including the desired colorimetry of the desired display, and the processor is responsible for the proper conversions to the actual panel's native physics. Internally, there are subtleties to the color math that must be performed, but largely, it involves a conversion from the data in the signal (based on the provided settings that describe that input) to the panel-specific values that are needed to produce the desired color and brightnesses of the light. Factors at play in this dynamic include the native primary spectra and colorimetry of the panel, the balanced white point of the multi-channel array, the explicit EOTF of the display hardware, and the physical dynamic range limitations of the panel. This color processing is reasonably well understood as it is fundamentally similar to the processing expected in any calibrated display for accurate reproduction of display-rendered color signals.

The conversion method between the visual signal values and the control of the distinctive LEDs is called pulse width modulation (PWM). The resolution of this conversion (the number of bits) is important for a good visual result. It may appear that a 16-bit driver, therefore, has sixteen stops of information, but the bottom few stops are represented by only a minimal number of values. This results in significant quantization at the dark levels if no additional improvements are attempted. To mitigate this, LED processor vendors use time multiplexing or spatial dithering to reduce perceived quantization artifacts. Rather than calculate a linear PWM value in a sixteen-bit controller working space, they use more integer bits (in the 20-22 range) and use the additional bits to effectively produce a higher bit depth result. The perception is that there are more brightness levels in the dark regions, but with temporal or spatial noise rather than individual quantized values.[11]

The LED processor knows the panel's native color gamut. This is part of the initial wall setup and calibration process and usually does not involve user adjustments. The user indicates the color space for the input signal to the processor, and the processor drives the panels to reproduce that colorimetry via the actual display physics and processor capabilities. OSVP has no "best" choice for the input color space. If, for example, you know that all of your content has been defined and created within a particular color space (i.e., Rec. 709), it is ok to choose that color space as the one to send to the processor and thus use as the input of the processor. Typically, the native color space of the panels will be larger than Rec. 709, so this does mean that some producible display colors will not occur because they are not required by the content. Inversely, if you choose to send ACEScg data to the processors, a portion of those values are potentially out of gamut for the panels, so color clipping or more sophisticated gamut mapping will need to occur in the processor (**Fig. 2**). Gamut-mapping and gamut-clipping algorithms may vary widely amongst manufacturers and could represent a feature the operator has no manual control over. Some users choose to set the panel's native color space as the input color space, which results in no color space conversion or clipping in the processor. The onus is then on the users to ensure that they have properly produced that particular native color space in the real-time engine, where the conversion (and potentially clipping) will occur. The user would further be responsible for confirming the LED panels are calibrated in conjunction with the processors to faithfully deliver the declared colorimetry. This is the first point in the workflow where the user can directly measure the produced light to verify that the hardware is performing as represented by the manufacturer.

In addition to understanding the chromaticities of the primaries of the incoming signal, the LED processor needs to understand what EOTF was used in encoding the input signal from the VAD real-time engine. The processor will perfectly invert this encoding to produce the linear values that

1545-0279/25©2025SMPTE

eventually drive the panels. Camera OETF and display EOTF curves are explicitly defined in standardized video ecosystems. For example, scene-to-screen standards for broadcast video have widely adopted the Rec. 709 OETF and Rec. 1886 EOTF[12] for the respective hardware. However, it is extremely important to note that concatenating these two transfer functions does not yield a perfect linear map of scene light to screen light. In fact, the net opto-optic transfer function (OOTF) produced via the cascade of these two curves yields a significantly boosted net contrast. This boost is necessary to yield pleasing images for human observers who are subjected to different adaptation influences from reference white point, white brightness, surround adaptation, and optical flare variations in scene versus screen environments. For OSVP workflows, these phenomena do not require the traditional OOTF contrast boost, as the camera is meant to interpret the VAD signal as literal scene radiometry. As such, tone transfer encodings used by the engine and the LED processor and wall should be perfect inverses of one another and not conform to more classic video standards. One potential exception could be to allow for the correction of lens flare in the camera should that be well characterized, though this is not currently a common practice. The two most common tone transfer encodings for the OSVP signal workflow are perceptual quantizer (PQ) and "gamma."

PQ is defined in SMPTE ST 2084.[13,14] The design of PQ involved careful modeling of human vision perception of luminance differences in various spatial patterns across a wide dynamic range of stimuli. PQ may be considered a direct EOTF, indicating a produced brightness output for each quantization value input, or the direct inverse, representing quantized encoding of the display-referred brightness. The goal of the encoding curve is to most closely conform to human threshold perception as defined by the Barten model for modulation detectivity (also known as contrast sensitivity). When a luminance signal is PQ encoded at 12 bits (the nominal standard), each integer digital count step in a displayed pattern should not yield a detectable lightness change. In 10 bits, there may be a perceived difference per step, but the PQ curve helps ensure that the difference is the same across the entire range from dark to bright.

Additionally, the PQ curve is an absolute specification of luminance. Each unique signal value defines a unique luminance that should be produced by the display device. For example, a 10-bit value of 520 defines a luminance of 100.23 nits (candelas per meter squared); a value of 769 defines a luminance of 998.93 nits. The maximum 10-bit value of 1023 produces 10000 nits, which is the limit of the PQ encoding system. Most LED panels produce brightnesses in the 1200 to 1600 nit range, which means that PQ encoding exceeds the range of the device. At the low end, the smallest PQ values represent extremely dark values. The first twenty-two values of a 10-bit PQ encoding represent brightnesses below 0.01 nits.

The other option for LED processor input encoding is "gamma." This is simply a mathematical power function where the exponent is declared using the γ symbol. Many users choose a 2.4 gamma as this value is standardized in BT.1886, but it is reasonable to select values from 2.2 to 2.6. These values are historically connected to the explicit light rendering physics of early cathode-ray tubes and the camera tone transfer curves that paired with them to make pleasing images. However, they persist in modern signal processing as a reasonable accommodation of human perception phenomena in encoding. The gamma method differs from PQ when considering the brightness of a particular signal value. The gamma function is a relative value rather than an absolute, so a maximum 10-bit value of 1023 in a gamma-encoded input signal will produce the peak brightness of the panel. On most LED processors, the peak brightness of the panel is selected separately from the gamma to facilitate this particular function. Therefore, for an example system that uses panels with a peak brightness of 1500 nits, the signal would use the entire range of values to illuminate those panels. However, if a neighbor wall (or ceiling) used panels with a different peak brightness, then the signal to that section would need to be created differently to ensure that it is possible to see both walls at the same brightness for objects that require it.

The choice of gamma versus PQ for the signal to the LED processor is open to the needs of the production. Neither option is inherently better; however, whichever one is chosen, the final step in the real-time engine must involve encoding the selected option. This ensures invertibility and maintains a completely linear relationship between the VAD color space luminance and the luminance of the displayed LEDs. When the option is PQ, the virtual scene displayed should have known values that can subsequently be measured on the wall. For gamma, some exposure gain (or multiplier) should be present to adjust the scene values to the percentage of peak brightness required when the LED wall is used in this relative decoding paradigm.

Step 3: Photographing the LED Walls

Once the LED processor has converted and displayed the signal, the light from the panels must be considered, as the capture camera will see it. In LED workflows, the camera will capture not only the LED wall but also the actors, sets, and props in front of the wall. The light fixtures illuminating the actors may have differing spectral emissions. All visible light is part of the spectrum, specifically the wavelengths in the 380 to 780 nanometer (nm) range. In a traditional outdoor setting, the light from the sun and sky will provide all of the frequencies in that range in varying proportions. In an indoor setting, the lights can vary widely, from tungsten to fluorescents to LEDs. Most stage light instruments have broad spectrum emissions, but the LED walls typically do not. There is work by manufacturers to produce panels with additional emitters to produce more wavelengths, but the most common LED panels that are used in OSVP have spectral emissions that are quite narrow. This means that for each reproduced red, green, and blue signal sent to a pixel, the emitted spectral power distributions of the individual channel emitters do not overlap. If the LED walls were the only light source, many illuminated objects in the scene would look unusual. One unfortunate property of using only LED walls as light sources is metameric failure. This is where two

1545-0279/25©2025SMPTE

objects that normally convey different color appearances under natural lighting look the same under LED panels because a large portion of the spectrum is not available to illuminate them (**Fig 3**). For example, the spectral region between green and red is where the subtle reflectance distinctions of yellow and orange objects occur. But with the narrow emission of the LED panels, these object colors may look quite similar under that lighting.

To correct this potential lighting issue on an LED volume set, foreground set elements are often illuminated by an array of traditional cinema lighting instruments. Still, errors can occur for objects subject to mixed lighting if they are under the heavy influence of both the panels and the augmenting stage lighting. More important is the dilemma caused by the direct capture of the narrow emission spectra of the LED panels. The stunning background VAD elements that make OSVP so compelling are extremely precarious spectral stimuli to be photographed. In fact, variability in colorimetric accuracy of the capture is exacerbated specifically by the narrow emission spectra of the LEDs themselves.

In a digital cinema camera, the color is determined by filtering the light that travels through the lens into a red, green, and blue image. In most cameras, this happens on a single sensor with an array of filters, one per photosite. To reconstruct a full image, the demosaicking process creates blended pixel values for locations with only one filter. Each camera manufacturer has their own method of doing this demosaicking and also has different types of filters on the photo sites. This means that different cameras will see the world differently via different spectral integrations and will particularly see LED panels differently from one another because the spectral responses of the camera filter overlap with the narrow LED spectral emission responses in a mathematically variable absorption profile.

More importantly, the cameras detect incoming light energy differently than the human visual system. As described previously, the cameras have explicit spectral responses. The human visual system also has explicit spectral responses called cone responses, referencing the three major cone classes on the human retina, sensitive to long, mid, and short wavelength regimes of the visible spectrum. If a camera existed with the same responses as our cones, it would be colorimetrically similar to our vision system. Unfortunately, virtually no cameras have this property, and no camera has the same response as our eyes, even with a linear combination of the cone responses. If they did, the camera would meet the Luther-Ives condition, allowing it to distinguish the same color differences as human vision.[15] In actuality, most cameras are different from human vision, so some compromise must be made when doing calculations that relate color using human vision metrics. This makes the white balancing process critical during the LED wall calibration process.

In the color space conversions described earlier, a standard connection space based on the CIE 1931 2-degree Standard Observer was inferred. Linear algebra allows traditional RGB renditions of color, where the three primaries and white points are all explicitly defined in terms of CIE chromaticity coordinates, to be transformed with perfect mathematical preservation of color representation. Two different RGB color space renditions of a stimulus may yield an equivalent color appearance to the standard observer so long as the conversions of each RGB to the standard XYZ tristimulus representation match. It should be noted that color scientists currently accept multiple different colorimetric observer standards and that linear algebra can only guarantee a perfect match for the specified observer. There are also the aforementioned concerns when one RGB color space gamut is smaller than another, so stimuli from the larger space may suggest matching stimuli in the smaller space that are physically unrealizable (even if mathematically plausible).

In the end, some account must be taken for the camera's response relative to the LEDs' narrow responses. This is the linear camera color calibration matrix, which is needed to ensure end-to-end matching. The matrix should be placed in the conversion from linear scene to encoded signal sent to the LED processor and is an adjustment to the matrix that is already present to convert from the real-time engine working color space to the LED input color space (a matrix which may be a simple identity function if the two spaces are already equivalent). The rest of this section will assume that the matrix was correctly made to continue the discussion of the remaining color management steps that help ensure end-to-end matching. Validation of the performance of the calibration matrix will be discussed in future articles.

The cinema camera will capture the light on the panels using spectral responses, which will be demosaiced to produce red, green, and blue channel data. It is important to recognize that the resulting sensor RGB values, the result of applying a demosaicing algorithm to the integrated interaction of scene illuminant, object reflectance, lens transmission, and sensor spectral sensitivity, are not in any well-defined color space. At that point in the camera's processing, the values are "camera native" RGB. Each camera vendor has its own conversion method from its "camera native" space to an RGB camera encoding color space.

Most camera encoding spaces are large enough to minimize required gamut mapping for typical captured colors and have color primaries well outside the spectral locus (a mathematically plausible treatment of color even though said primaries are physically unrealizable). For such spaces, this means that the camera should never, for example, produce an encoded value of (1,0,0) because to do so would imply that it saw a color that is not physically possible. Instead, the vendors choose those primaries that help them comfortably encode all the possible scene values that they will photograph. The physically realizable colors captured should (after conversion to a Yxy representation) be close to the Yxy values that could be measured on location.

Vendors strive to minimize residual error rather than eliminate it, as they know the latter is impossible (i.e., the Luther-Ives condition is not achieved). Instead, each vendor defines a conversion from camera native RGB to their chosen encoding space. Factors differentiating one encoding space from another include quantization efficiency, scene-colorimetry coverage, and the mathematical alignment of the chromaticity of selected primaries with those found in tra-

1545-0279/25©2025SMPTE

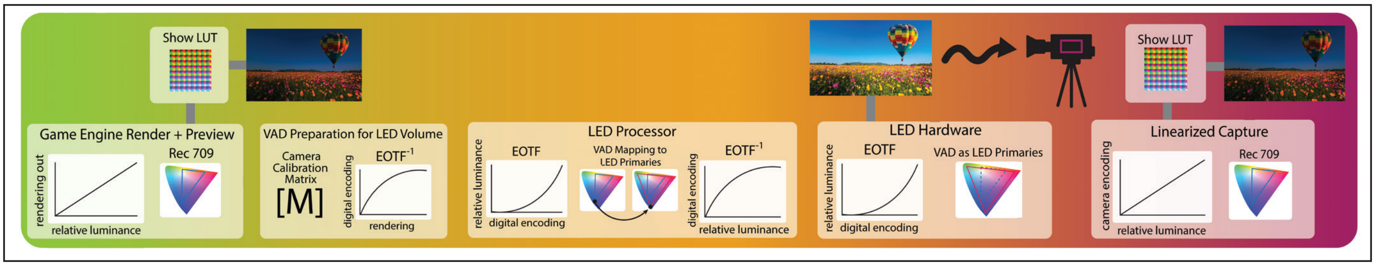

FIGURE 5. Detailed color operations that occur within the three major sections of an OSVP color workflow described in this paper. Special attention should be made to when and where encoding transforms, color space transforms, and camera calibration matrices should be applied, as well as which parts of the workflow an image should "visually" match (demonstrated through viewing the image via a designated "show look").

ditional post-processing operations such as color correction.

It should be noted that some manufacturers employ preferred perturbations of captured colorimetry that intentionally yield inaccurate color rendition of the scene to provide an image that is, in some definition, more pleasing to the average user. Studies of photographic rendition have proven accurate color representation is seldom a general audience's prime driver of image quality versus "remembered" color.[16] Not every camera vendor is guaranteed to strive for Luther-Ives processing, and a vendor might offer multiple encoding possibilities, some more accurate than others. When accurate color capture is not a primary goal, and a camera deviates far from the basic goals of metameric capture encoding, it should be expected that significantly more complex transforms beyond a simple matrix will be needed for color management of VAD assets. Further, if the camera encoding gamut is smaller than the VAD content or the LED wall gamut, in-camera gamut-mapping or gamut-clipping may make matches for more chromatic stimuli especially difficult.

Although the sensor in a cinema camera is largely linear, the image undergoes processing inside the camera, typically resulting in an OETF-encoded output. This output may be in an encoding space that differs from the real-time engine's linear working space. To return to a linear version of the scene for comparison, the effect of the OETF curve must be undone by inversion, and the resulting linear RGB values must be converted from the camera encoding color space to the linear working space of the real-time engine. If all is working well, this should match within an overall exposure change (or a linear multiplier value). This assumes the camera calibration matrix was designed to ensure an end-to-end match. If not applied, the difference between the scene and the linearized camera data should correspond solely to the missing matrix. These calculations and their appropriate locations within a color workflow are described in **Fig. 5**.

Overall, the OSVP color pipeline consists of many primarily linear devices (at least, compared to highly nonlinear behavior in older photochemical workflows). This makes color management reasonably straightforward and fairly obvious when one of the steps has an incorrect non-linear conversion. For example, a simple mistake would be sending tone-mapped data from the engine to the wall. This would be similar to displaying a DVD signal on the wall; the image already has a look applied, so the light on the wall is not truly linear, while the light elsewhere (actors, props, sets) is linear. If the tone curve is properly disabled in the real-time engine, the next likely mistake would be a mismatch between the encoding tone transfer function applied during signal creation and the input signal settings for the input signal (e.g., full range signaling vs. narrow range). When this is incorrect, brightness relationships between objects in the scene are not maintained when the overall scene exposure is changed in the engine. For example, a box that is twice as bright as another box could be found to exhibit that behavior on the wall, but it should continue to have that double relationship if the scene is brought down a few stops. If this does not occur, the settings on the LED processor should be checked against the settings on the signal output of the engine.

Understanding the core principles behind color transforms, where they occur within a color pipeline, and the limitations of the hardware ecosystem can significantly speed up the troubleshooting process when things go wrong on set, both in terms of identifying the location of the error and resolving the issue in a timely manner. The information presented in this article aims to provide that education for engineers and practitioners alike.

Conclusion

At this point, the article has stepped through the basics of each stage of an LED-based OSVP color pipeline, from VAD asset and scene creation through capture of that scene with a camera, along with a description of appropriate places to apply colorimetric transforms to determine if the final captured result was a "match" to the VAD original intent. With so many new roles and locations to apply color management in the OSVP color workflow, this article aimed to review these color principles and introduce how they are used to maintain creative intent, or "match," throughout the production process. The article also presented the technical challenges within this color workflow that frequently cause these images not to "match," such as metamerism failures introduced by the use of LED walls, both in the illumination of real-world objects and their photographic capture with different cameras. Preferred perturbation of images by different camera manufacturers also introduces complexities into the color workflow.

The RIS-OSVP team will explore the topics and challenges introduced in this paper in more depth in the future. The team hopes these articles will help engineers and practitioners better understand how to use the tools available and inspire readers to investigate improvements to existing practices to fully maximize the benefits that OSVP offers today's filmmakers.

Acknowledgment

Thanks to Evelyn Sutkus, student in the BS motion picture science program at RIT, for support in preparing illustrations.

References

1. International Telecommunication Union—Radiocommunication (ITU-R), Recommendation BT.709-6, "Parameter values for the HDTV standards for production and international programme exchange," Jun. 2015.
2. International Telecommunication Union—Radiocommunication (ITU-R), Recommendation BT.2020-2, "Parameter values for ultra-high definition television systems for production and international programme exchange," Oct. 2015.
3. SMPTE, ST 2065-1:2021, "Academy Color Encoding Specification (ACES)."
4. M. Leonhardt, H. Brendel, "Critical spectra in the color reproduction process of digital motion picture cameras," *Proc. Color and Imag. Conf.*, 23: 167-170, Nov. 2015.
5. D. Long, M. D. Fairchild, "Modeling Observer Variability and Metamerism Failure in Electronic Color Displays," *J. Imag. Sci. and Technol.*, 58(3): 30406-1–30406-9, Nov. 2014.
6. D. K. Prasad, L. Wenhe, "Metrics and statistics of frequency of occurrence of metamerism in consumer cameras for natural scenes," *J. Opt. Soc. Am. A*, 32(7): 1390-1402, Jun. 2015.
7. X. Guo, K.W. Houser, "A review of colour rendering indices and their application to commercial light sources," *Lighting Research & Technology*, 36(3): 183-197, Sep. 2004.
8. International Organization for Standardization (ISO), ISO 22028-1:2016(en), "Photography and graphic technology—Extended colour encodings for digital image storage, manipulation, and interchange—Part 1: Architecture and requirements," Mar. 2016.
9. International Electrotechnical Commission (IEC), IEC 61966-2-1:1999, "Multimedia systems and equipment—Colour measurement and management - Part 2-1: Colour management—Default RGB colour space - sRGB," Oct. 1999.
10. M.R. Pointer, "The Gamut of Real Surface Colors," *Color Res. Appl.*, 5(3): 145-155, 1980.
11. S. Daley, X. Feng, "Bit-depth extension: Overcoming LCD-driver limitations by using models of the equivalent input noise of the visual system," *J. Soc. Inf. Display*, 13: 51-66, Jun. 2012.
12. International Telecommunication Union—Radiocommunication (ITU-R), Recommendation BT. 1886, "Reference electro-optical transfer function for flat panel displays used in HDTV studio production," Mar. 2011.
13. SMPTE, ST 2084:2014, "High Dynamic Range Electro-Optical Transfer Function of Mastering Reference Displays."
14. 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 2015.
15. R. Luther, "Aus dem Gebiet der Farbreizmetrik (On color stimulus metrics)" *Z. Tech. Phys. 8: 540-558*, 1927.
16. D. L. MacAdam, "Quality of Color Reproduction," *Proc. IRE, 39(5): 468-485*, May 1951.

Bibliography

Goldstone J., "Color Management," *The VES Handbook of Visual Effects: Industry Standard VFX Practices and Procedures*, edited by J. A. Okun, S. Zwerman, C. McKittrick, L. Sepp-Wilson, New York: Routledge, 2021.

Madden, T. E., Giogianni, E. J., *Digital Color Management: Encoding Solutions*, 2nd ed., edited by M. Kriss, West Sussex: John Wiley & Sons, Ltd., 2008.

Zwerman S., Okun, J. A., *The VES Handbook of Virtual Production*, 1st ed., New York: Routledge, 2024.

About the Authors

David Long is Director of the Center for Media, Arts, Games, Interaction & Creativity (MAGIC) and associate professor of motion picture science at Rochester Institute of Technology (RIT).

Rod Bogart is the Principal Color Scientist in the Virtual Production Tools team at Epic Games, concentrating on the color pipeline of Unreal Engine. His initials are RGB.

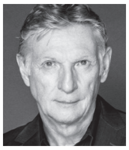

Spencer Stephens is SVP of Production Technology and Security at MovieLabs. He has worked at three Hollywood studios in technology and production roles.

Catherine Meininger is Sr. Director of Color Science at Portrait Displays, where she leads the design and development of color calibration algorithms. She earned her bachelor's degree in Motion Picture Science from RIT.

Joseph Goldstone consults on color management for visual effects, continuing themes from his time in the Image Science group at ARRI Cine Technik and Industrial Light & Magic. He is a SMPTE Fellow.

Tim S. Kang, a cinematographer, serves as Principal Engineer, Imaging Applications, for Aputure. He chairs the Lighting Committee in ASC's Motion Imaging Technology Council.

Kathryn Brillhart is a cinematographer and virtual production supervisor who leverages real-time technology, volumetric capture, visualization techniques, and supervising visual effects to enhance projects.

Jim Geduldick is chair of SMPTE RIS and a renowned VFX/virtual production supervisor and cinematographer whose work spans film, visual effects, virtual production, and artificial intelligence.

A contribution received from the authors and SMPTE Rapid Industry Solutions (RIS) Group, August 2024. Copyright © 2025 by SMPTE.

DOI: 10.5594/JMI.2025/XEJK1138
Date of publication: 1 April 2025