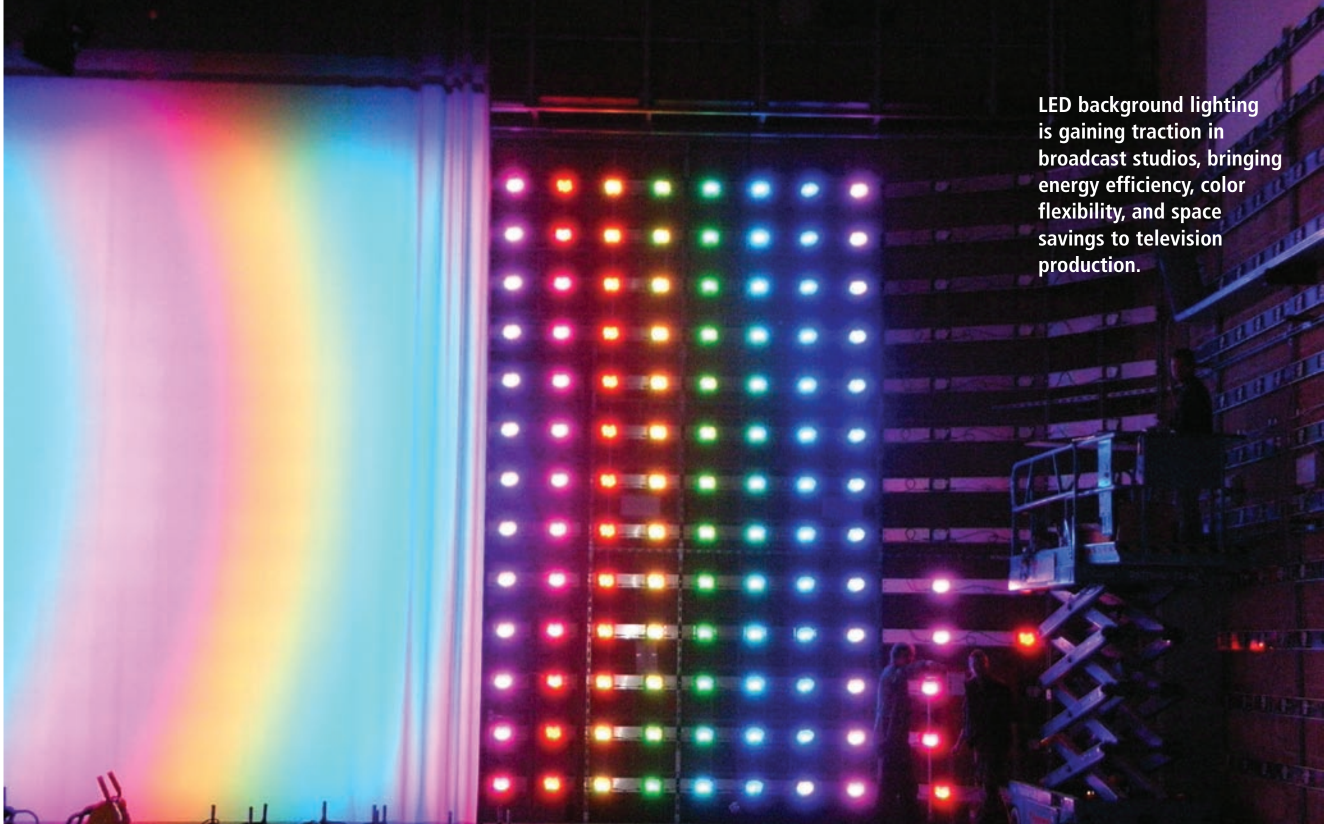

LED background lighting is gaining traction in broadcast studios, bringing energy efficiency, color flexibility, and space savings to television production.

Digital imaging with LEDs: More than meets the eye

BY **RYAN FLETCHER**

I VIVIDLY REMEMBER MY FIRST TIME lighting a dance production captured on digital video. The cameras showed up on the second night, the performance progressed smoothly, and the video crew successfully captured their shots. Thankfully the skin tones and costumes rendered well on camera, and the low-light scenes had sufficient definition.

The surprise was the appearance of the backdrop. I lit the cyclorama with RGB LED strips. The LED colors captured on camera were far different from those seen by the audience's eyes, and also looked quite dissimilar as the video cut between shots from the various camera angles around the stage. The scenes designed with a red-lit backdrop looked orange, and the bright blue sky for the birds-in-flight piece took on an under-the-sea aqua shade, much to the dismay of the choreographer.

I knew there must be a technical reason behind this difference. Little did I realize that this would be the first of many experiences where I questioned why digital cameras see LED lighting so differently than our eyes do. As both technologies rapidly push us into the digital age, we are presented with a number of new issues that challenge our traditional assumptions about the interaction between camera and light.

Capturing color

Digital image capture is fundamentally modeled after human vision. Our retinas rely upon three types of color sensors called cones, each of which is sensitive to photons of light in either the red, green, or blue areas of the spectrum. Taken together, they provide us with the information necessary to visualize a color, which we perceive when these signals are integrated within the eye, transferred to the brain, and processed. Color does not actually exist in the real world; it is a creation of our perception, assigned by our mind according to visible wavelengths of light.

The sensor in a digital camera contains millions of pixels, which, like the cones in our eyes, are designed to collect photons of light (**Figure 1**).

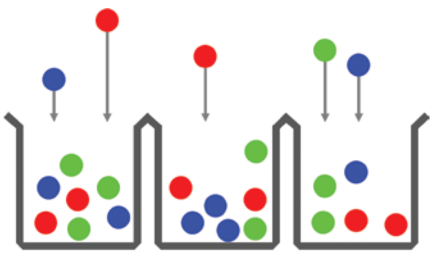

Figure 1 – Digital sensors collect light in tiny cavities, called pixels.

These photons are detected by photodiodes within each pixel, amplified, and read by the camera as an intensity value. Pixels have no built-in way to identify the wavelengths of the light they collect; only the total sum of light is determined. This is sufficient to render a grayscale image, but for color, more details are required about which wavelengths of light are received at each pixel. Several methods for discerning this information have been explored over the years, but two distinct systems have evolved to become the most widely used in modern digital cameras.

Common in the broadcast industry is the three-sensor camera. In this design a prism splits the light entering the camera into red, green, and blue components, directing each part towards its own unique sensor. Each of the three sensors contains the same number of pixels, and the intensity of light at each pixel is recorded, yielding red, green, and blue information at each point. This information is combined in order to form a color image.

The other prevalent system, and the one we will focus on in this article, is the single-sensor camera. As the name describes, only one sensor is used, but with the addition of a segmented color filter on top. The filter covers each pixel cavity, and only light with wavelengths to which the filter is transparent can pass through.

The most common filter pattern is the Bayer Array, in which red, green, and blue filters are arranged in a tile-like fashion (**Figure 2**).

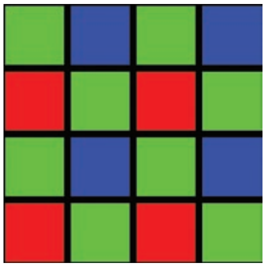

Figure 2 – Bayer Array filter pattern

Notice that the Bayer system uses twice as many green tiles as red or blue tiles. Since our eyes are most sensitive to green light, camera manufacturers are able to gain an improvement in detail and image noise by doubling the number of green tiles.

An obvious limitation of this single-sensor approach is that it reduces the resolution of information that can be captured, since each pixel only receives light within a specific color region (**Figure 3**).

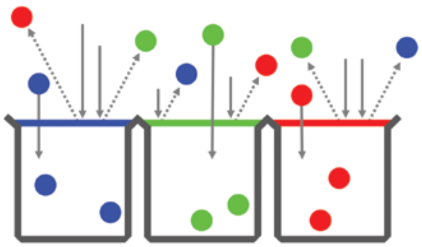

Figure 3 – The Bayer filter allows only specific color regions to pass into each pixel.

At least three pixels—one each of blue, green, and red—are required in order to fully resolve the color of light in a given area. To compensate for this decrease in resolution, an interpolation process takes place in the software of the camera to estimate the color at each pixel by taking its neighbors into account. Pixels are grouped into sets of overlapping arrays, which are then compared and averaged to extract more information from the image. The ultimate goal is to have enough color information to correctly determine the real-world colors in the scene, as they would be perceived by the human eye. Although most feature films and television productions prefer to adjust color, often unnaturally, to suit the director's intent, it is useful to start with as accurate an image as possible.

No matter how close we can come to the eye's response, there is one critical difference: adaptability. The human visual system is highly adaptive to our surrounding environment. Some of the adjustment occurs optically, when our retinas enlarge and shrink according to the intensity of light that hits them, and other adaption occurs mentally, when our brain helps to color-correct the images, adjusting the color to fit what we think "looks right." Digital sensors are quite the opposite, however. They record exactly what they see, leaving manipulation of the image to the camera software and postproduction process.

Focusing on the spectrum

The three types of cones in our eyes cover a range of wavelengths, from roughly 360 nm to 830 nm, but with greatly diminished sensitivity at either end. Although camera manufacturers would very much like to construct digital sensors with sensitivities matching those of the human visual system, they are restricted by the limits of Bayer Array filter material and manufacturing methods.

The sensitivity of a single-sensor digital camera varies widely from manufacturer-to-manufacturer and camera model-to-camera model *(sometimes even from camera-to-camera)*. Yet most generally follow a format of three spectral peaks, corresponding to the blue, green, and red transparency of the filters in the Bayer Array.

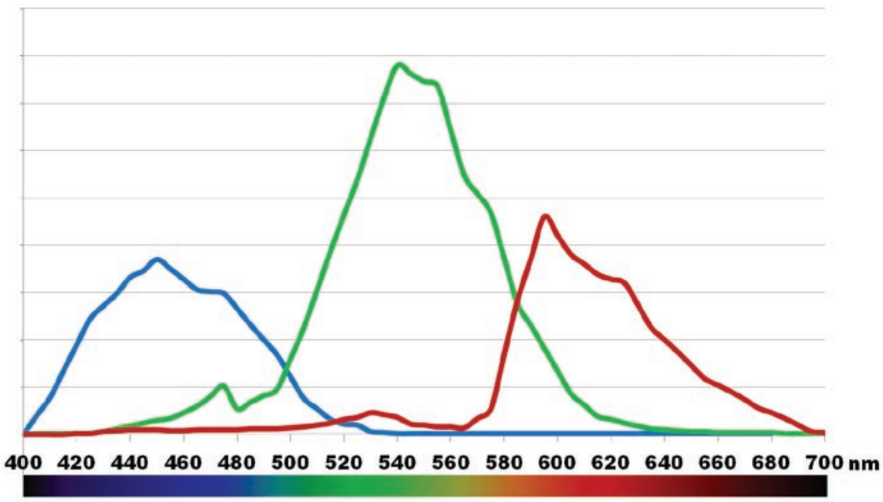

Figure 4 – Example of spectral sensitivity in a single-sensor digital camera.

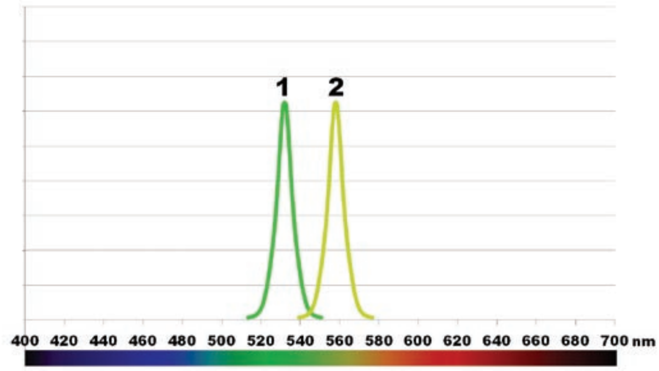

Figure 5 – Example emission from green and yellow LEDs.

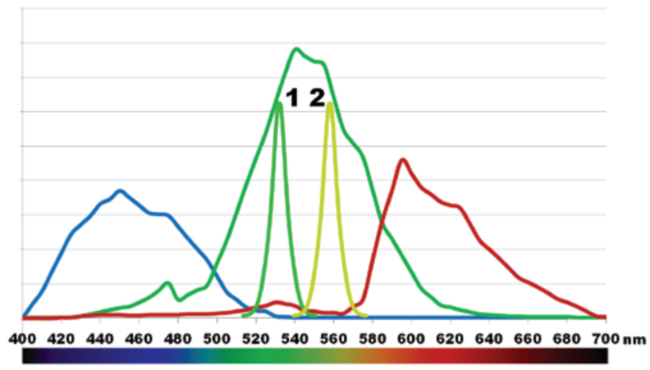

Figure 6 – Green and yellow LEDs overlaid with camera sensitivity.

When lighting a subject with a broad-spectrum source, such as the sun or a tungsten-halogen lamp, the valleys in-between the sensitivity peaks are of relatively little concern. Modern cameras use sophisticated processing algorithms to correct for color abnormalities, and camera manufacturers are resolute in ensuring that particular colors, such as skin tones, are authentically captured under broad-spectrum illumination.

Yet when the same subjects are lit with highly discontinuous sources, such as LEDs, all bets are off. LEDs are narrow-band emitters, concentrating most of their energy in very specific regions of light. Further, there are many combinations by which white and colored light can be created with LED sources. For camera manufacturers, ensuring accurate color reproduction in this new, diverse landscape becomes a far trickier challenge.

Differentiating color

Let's start with an example of two LEDs that are near each other in the visible spectrum: a primary green and a yellow (**Figure 5**).

If these two sources were to illuminate a neutral surface *(one whose degree of reflectance is the same regardless of wavelength)*, most of us could clearly see the difference in color between green and yellow, thanks to the good overlap of the green- and red-sensitive cones in our eyes. Now let's overlay the sensitivity of the sample digital sensor we referenced earlier (**Figure 6**).

Notice that the spectrum of each LED lies completely within the spectral range of the green filter material. Therefore the emitted photons are only able to pass through into the pixels with green filters; the light goes undetected by the blue and red pixels. [1]

Remember that a pixel is only able to report the magnitude of light that reaches it. There is no way to determine the color or specific wavelengths of light within a single pixel; the camera knows if the pixel has a red, green, or blue filter over it, but that is all. Since both LEDs in this example are detected only by the green channel, the camera is unable to distinguish between the chromaticity of each. To the camera they appear to have the same hue and saturation, most likely a form of primary green.

Now let's turn on a third LED. LED #3 has a peak wavelength just 10 nm to the left of LED #1 (**Figure 7**).

Those of us who come into this discussion from a lighting perspective have likely heard about color binning, which is a system for categorizing LEDs after the manufacturing process. Since LED manufacturing is not yet a perfect science, there are unintentional deviations in the color of light from one chip to the next. Therefore LEDs are grouped into bins according to color tolerances defined by each manufacturer, and the 10 nm distance between LEDs 1 and 3 on the graph could very well be within the hue tolerance of some bins.

LED 3, like the two others, is almost exclusively within the camera's green sensitivity region. The camera therefore assigns it the same chromaticity as the other two LEDs. At first glance this may appear to work in our favor, helping to even out differences within color bins. However we now have a 40 nm region of the spectrum in which the camera is virtually blind to the chromaticity of narrow-band light sources! 40 nm is far beyond all but the most generous bin-to-bin LED tolerances, and can seriously limit our color latitude when using this camera with colored LEDs.

As if this color issue wasn't enough to contend with, we must also face the fact that the sensor does not detect as much light from LED 3 as it does from LED 1 or 2—even though the camera incorrectly believes all three to have the same chromaticity. Notice that the sensor's green sensitivity level at LED 3 is lower than that at LED 1 or 2. The green Bayer filter material has a lower transmission for light at the wavelengths of LED 3, and therefore some amount of this light is unable to pass through to the underlying pixels. In this case, the luminance reduction works in favor of reproducing an accurate image, since our eyes are also a bit less sensitive to light with wavelengths around LED 3, but most of us would certainly be able to discern the color difference.

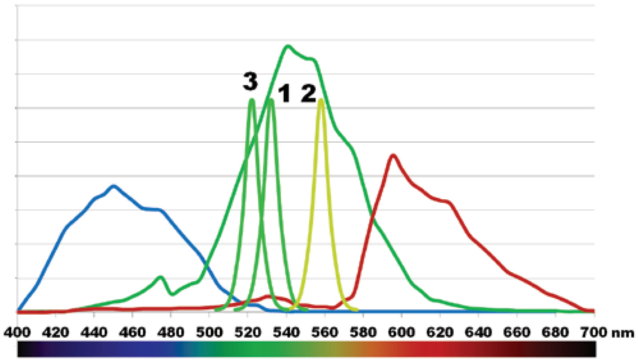

Figure 7 – LED #3 peaks just to the left of LED #1.

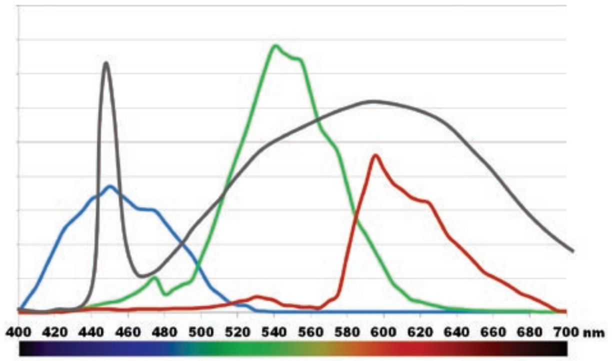

Figure 8 – White LED and camera sensitivity.

A balancing act

A frequent postproduction step in video and still photography is color balancing, a process by which the relative intensities of red, green, and blue channels in an image or sections of an image are manipulated to achieve a desired color look. Placing a filter over a tungsten-halogen or similar near-continuous-spectrum light source can create broad-spectrum colored light, allowing greater color balancing flexibility. Although the filter creates a specific color by allowing only certain wavelengths of light to pass through, the range of wavelengths is large. These images can then be easily color balanced. Yet the precise, narrow spectrum from colored LEDs can quickly oversaturate the red, green, or blue color channels in a sensor without registering any light in the other two. This makes color balancing difficult, since light was only captured in one of the three color channels.

Some suggestions have been made to circumvent this problem, such as intentionally polluting the light in a scene lit with only blue LEDs, by turning on some low levels of red and green or white light. Yet there is something to be said for ignoring a technical problem to achieve a creative effect. Colored LEDs give a pure and direct look to a scene, and can be quite effective in the right shot.

What is white?

So now that we have covered some of the issues with colored LEDs and digital capture, let's take a look at white LEDs for a moment. White LEDs are easier for camera sensors to handle, since they more closely approach the broad-spectrum light sources for which digital cameras are biased. (**Figure 8**)

The term "white LED" is slightly deceiving, since it is actually a blue (or, less often, a UV) LED with a phosphor covering, either directly on the chip itself or placed a slight distance away in the optical system. The phosphor redistributes some of the emitted energy toward the green and red regions of the spectrum. A white LED usually has as a sharp spectral peak in the blue region, takes a dip near cyan, and, thanks to the phosphor, makes a broad leap into green, yellow, orange, and red. Although there is significant variety among white LED spectrums, most follow this general pattern.

The broadcast and digital production fields often gauge three characteristics when lighting with white LEDs: correlated color temperature, plus–minus green, and color rendering.

A change in correlated color temperature is mainly effected by altering the amount of red and blue emitted from the LED and detected by the camera, with higher color temperatures having less red and more blue content. Plus–minus green *(or green-magenta shift, as some call it)*, the amount of green or magenta tint in the light relative to a neutral white, is strongly influenced by the amount of green content at a given color temperature.

Broadcast and digital cinema production is obsessively careful about the amount of green in white light, since digital sensors generally seem to detect more green from white LEDs than our eyes do. What appears to the eye as perfectly neutral white may take on a surprising green tint when viewed through a camera.

Even a small amount of green can lend a very unflattering cast to human skin (**Figure 9**). Fortunately the current era of digital production

Figure 9 – A subject lit with two different LED sources. The source on the right has a shift towards green, which is more readily apparent on skin than on the color chart.

offers us the ability to attach a calibrated monitor to a camera and scrutinize a preview—a bit of *What You See Is What You Get*—so that any issues with green can be caught and fixed before that one-time-only shot is taken. The uncertainty is much greater when shooting on film due to the lack of a truly accurate real-time preview.

Good white-light color rendition is highly product-specific, depending not just on the spectrum of the white LED, but also on the sensitivity of the camera. Both of the more commonly referenced color rendition metrics, CRI (Color Rendering Index) and CQS (Color Quality Scale), are based upon the response of the human eye. The more the spectral response of a camera differs from our eyes' response, the less useful these metrics become.

A vibrant bowl of fruit illuminated by a white LED fixture may appear vivid and natural on one manufacturer's camera but dull and unappetizing on another's, depending on the sensor design and in-camera color processing. In this case, the broader spectral coverage from a tungsten-halogen source would yield more consistent results. Although the problem of unappealing color rendition can be somewhat mitigated in postproduction, the cost and time required usually outweighs the initial effort of using a fuller-spectrum source to begin with.

Lighting for the moment

Reflecting back on my first experience with LEDs and digital video, I pondered if I could have made a different choice to ensure that my backdrop color selections looked good on camera. I saw two main options: either light without LEDs, or ignore the audience and light for the camera.

The first option, light without LEDs, was not attractive. New illumination technology should not be ignored when it can meaningfully expand the creative possibilities available to the designer.

The second option, lighting for the camera, would have sacrificed the view of the live audience in favor of a better LED-background appearance on DVD. What worked for the eye certainly did not work for these cameras, and what would have looked good on camera probably would not have worked for the eye.

As different as lighting for the camera and lighting for the eye can be, lighting designers in both disciplines have been working with the same basic set of light sources for decades. Yet now, with the advancement of LED technology finally making its way into more flexible fixtures, the light source can be built, tuned, and adjusted specifically for the eye or for the camera. However what seems to be a great new capability has also become a new risk, particularly so for lighting teams working with LEDs in broadcast and feature production. Suddenly the light seen with our eyes looks different through a camera. Very different. This leaves many of us wondering which view to trust more.

Thankfully, rapid progress is occurring on both sides of the problem: camera manufacturers continue to research new ways to more closely approximate the color sensitivity of the human eye, while LED color shows ongoing improvement through broader-spectrum phosphors and tighter binning.

Amidst all the differences that we experience when using these new, transformative technologies together, it is easy to miss how similar digital sensors and LEDs actually are. After all, photodiodes and LED chips are actually constructed from the same basic materials, and an LED can easily be modified to detect light rather than emit it. Someday they might even be one in the same—now isn't that a colorful idea! ■

Footnote 1: *There is a small amount of light from LED #1 that passes through to pixels covered by red filter material. The resulting signal is likely to be ignored in software, since it is quite small relative to the green content in this example. If it were not ignored, the results would be even more troublesome: the camera would identify the chromaticity of LED #1 to be closer to the red portion of the spectrum than LED #2!*

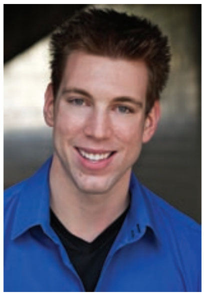

Ryan Fletcher is the LED Lighting Product Manager at ARRI. His experience spans both the technical and creative sides of light, as a former opto-mechanical engineer for the US Department of Energy and a part-time lighting designer. Ryan is a graduate of Stanford University with a degree in mechanical engineering and a minor in drama, concentration in lighting design.