Charles Poynton
tel +1 416 535 7187
charles @ poynton.com
www.poynton.com

Poynton's Vector 18 LUTs in 1000 words

"LUTs" – that is, *lookup tables* – are widely used in digital cinema (D-cinema), professional HD, and consumer HD. LUTs are found within cameras, within boxes used on-set, within post-production equipment, in external boxes used with display systems, and within certain displays. The contents of LUTs are established using a variety of techniques. In this article, we'll explore LUTs in various applications.

We start with displays. Practical D-cinema and HD displays produce additive mixtures of red, green, and blue primary components. Many display systems exhibit near-perfect independence of the three components (*channels*); however, even with channel-independence, most displays have a nonlinear relationship between signal value input and light output from each channel. For example, studio CRTs historically produced light power proportional to the 2.4-power of signal input. For high-quality display, the goal is to match the mastering display's conversion of signal to coloured light, and so the viewer's display must incorporate the 2.4-power or something close to it.

The 2.4 "gamma" is standardized in ITU-R BT.1886. See Poynton's Vector Issue 12, *Gamma estimation*.

Some display technologies, including PDP and DLP, use pulse-duration modulation. These displays exhibit near-perfect linear behaviour in each channel. In order to approximate the BT.1886 display "gamma," a suitable power function must be implemented in the signal processing chain. The required *degamma* function is typically implemented in three independent 1-D LUTs. Each LUT contains a mapping that implements the same power function. Signal values in the three channels are processed independently; the contents of three 1-D LUTs can be scaled independently to alter white balance.

By *linear,* I refer to digital image coding where across a large range of light power levels, halving of light power is associated with halving pixel value. In *logarithmic* coding, across a large range of light power levels, halving light power is associated with decreasing pixel value by a certain increment. (For example, in 10-bit Cineon coding, halving of scene exposure decreases pixel value by 90).

A PDP or DLP display may have primaries and white point that differ from the interchange standard (BT.709/BT.1886 for HD). Owing to the near-perfect additive colour mixing behaviour of PDP and DLP displays, mapping from the BT.709 interchange primary set such to the native primaries of a particular display device can be accomplished by a 3×3 matrix multiplication operation specified by 9 numbers: Small, predetermined fractions of each channel's value are added to or subtracted from the other channels. Three independent 1-D LUTs can't do the job of mixing channels together. (If the display primaries match the interchange primaries, the matrix isn't necessary.)

LCD displays typically exhibit a small degree of unwanted "coupling" between the channels. The required correction is a nonlinear function of the three signal component values. Three 1-D LUTs alone can't do the job, because that wouldn't allow channels to combine. A matrix can't do the job, because the required compensation isn't linear. Several other approaches could potentially work, but the practical solution is to implement a lookup table that is accessed by a combination of all three signal component values, and that produces all three values required by the display. If input and output values were 8-bits each, then we could use a table having 256×256×256 (that is, about 16 million) entries, where each entry contains three bytes. Such a table occupies about 48 megabytes. For professional applications, and for still images, that size isn't necessarily out of the realm of possibility; but for HD at about 60 megapixels per second, even at the high end of consumer space, it isn't practical. The solution is to rely upon the mapping being relatively smooth, and to *interpolate* – in simple terms, average – across a few thousand entries instead of a few million. You might get away with a table 10×10×10 – that is, 10^3 or 1000 words – but LUTs used in practice have perhaps 17^3 (that is, 4913) entries, and can easily have 10, 12, or even 16 bits per component at the output. Interpolation is typically either *trilinear* or *tetrahedral.* This *3-D interpolated LUT* solution allows compensating the subtle nonlinear mixtures exhibited by LCDs, but that treatment is necessary for only the highest quality studio displays.

When I say *3-D* here, I'm not referring to stereoscopy!

Other applications of 3-D LUTs are found in cinema. Camera film exhibits nonlinear mixture among colour components; so does print film. In film, the coupling ("crosstalk") components aren't linear, and so can't be removed just with a matrix. If you seek to remove these components, the most practical way is to use a suitable 3-D LUT. On the other hand, if your application seeks the film "look," you may need to *insert* such crosstalk components. You can use a 3-D LUT to do so. But by the time such imagery is made available for consumer viewing, the material has been mastered in a perfectly additive colourspace such as BT.709, and approved in that colourspace, with or without the "film look." If the display device is additive, a 3-D LUT isn't necessary downstream. For digital cinema exhibition, material is mastered in the DCI P3 *RGB* colourspace, also perfectly additive.

We've discussed 1-D LUTs and 3-D LUTs, but you may be asking, "What about 2-D LUTs?" In principle, an acquisition technology or a display technology could have two colour components that interacted nonlinearly, but a third that was independent of the other two. In that case, a 2-D LUT could be used to combine or compensate the two interacting channels, leaving the third alone. In practice, colour devices nearly always either exhibit channel independence or they exhibit interaction among all three channels, not just two of them. In principle, you use a 2-D LUT to alter chroma components – such as x and y, or C_B and C_R – leaving the third (luminance or luma) channel untouched. However, in practice, manipulation of all three channels is usually required; so, 2-D LUTs are very rarely found.

To calibrate a display – whether on-set, in the post-production suite, in a review room, or in a home theatre – you are tasked with bringing the display system into conformance with standards such as BT.709 and BT.1886. Generally, a 3×3 matrix and a set of three 1-D LUTs will suffice. To calibrate an LCD for studio use, a 3-D LUT may be necessary. If you are brought into the cinema production or computer-generated imaging/visual effects (CGI/VFX) world and asked to remove or insert the "film look," a 3-D LUT will be needed.

This article has 1000 words, excluding words in marginal notes!