Log Jam

… a bit more than you ever wanted to know about shooting log

GTC member DoP/lighting cameraman **Ben Turley** will be well known to users of the GTC Forum as he is rapidly becoming the Guild's 'go-to' expert on the often bemusing subject of 'looks' and LUTs. Based on the very thorough and valuable work and research undertaken in the process of creating his LUTCalc app, Ben explains some of the whys and wherefores of 'shooting log'.

It used to be so simple. Once upon a time, if you were a video shooter, you'd try to keep everything interesting below 90%, maybe have 65–70% zebras for skin tones and probably some grads and a pola on hand to help hold things in range. You might play with master black and detail from time to time and have dabbled with knee and matrix settings, either by experiment or via whispered suggestions from friends, broadcasters or the internet, but as someone who assisted dozens of excellent camera people, and subsequently did my own shoots through the Digibeta nineties, once happy with a set-up, that was generally it.

For film people, things were arguably even simpler. At any given time you might work with two or three stocks (particularly as the likes of Agfa, Ilford and eventually Fuji fell away). You'd get to know how they responded, whether you liked to rerate them a little, where your printer lights would tend to end up for thick or thin negs. You'd meter to mid gray and, whether you were a zone-system obsessive or just preferred to light to what pleased your eye, there was the safety net of film's wide latitude and a colour-timer or telecineist keeping things together in post.

But now it's all about options. Anything is possible; you can make adjustments before, during and after filming to radically alter the picture (or even find you have someone further down the line change it all anyway) and – going by fashion and marketing – if it's not raw, then it has to be log. Flexibility, complexity, lots of opinions – all in all, a bit of a mess!

Different ends of the scale

My three-year-old daughter and I recently found ourselves in a toyshop with £10 of Christmas money and an exciting array of plastic animals to choose from. She decided upon a shark and a unicorn, but unfortunately these came to £11. So, in an effort to teach her about money, I'm afraid I made her choose just one.* To her, £1 had made all the difference in the world. But my wife works in the NHS… and I don't expect to hear the phrase 'savings of one pound' in the Chancellor's budget statements any time soon!

As with pocket money and the national economy, so it goes with computers and digital cameras. In simple terms, a video sensor spits out values proportional to the number of photons hitting it. This is known as a linear response (or scene linear). The thing is that our eyes, brains, light meters, camera scales and filters are all much more attuned to ratios (or f-stops). Each stop equates to a doubling of photons.

The simplest way to store a linear value would be as an integer, with 0 equating to 0% black and 65,535 (the largest number a positive 16-bit integer can hold) equating to the camera's clip. The Sony F5 and F55 have roughly 6 stops range from mid gray to clip, so using that as a guide and treating 0 as being 0% black, then a correctly exposed 18% gray card would have a 16-bit integer value of about 860; a 90% white would be at about 4300; one stop below clip would be around 32,800; and clip itself 65,535.

> If not handled well, log can lead to many headaches and, by its very flexible nature, the final image can easily end up utterly disconnected from your artistic intention but with your name still on the credits.

That brightest stop alone, in the most overexposed of highlights would be covered by a range of nearly 33,000 possible values – nearly eight times as many as from pitch black to 90% white! Echoing the gulf between my daughter and the Chancellor's purses, a linear difference of one makes a huge difference in the shadows five or six stops below

mid gray (particularly if you are going to push the ISO upwards) but is effectively meaningless in the upper reaches of the highlights.

So, clearly, this is not a good way to store a picture.

A bit of maths (sorry) that fixes everything(ish)

A standard way to even up the spread breaks up a 16-bit number into three pieces. First, one bit is used to signify if the number is positive or negative; the next 10 bits (1024 possible values) are used to store a fractional value between 1 and 2; and then the big trick is that the remaining five bits are used to store an 'exponent' value. The exponent is the 5-bit value minus 15, so the range goes from –14 to +15. The actual number stored** is then of the form:

(sign +/–) fraction x (2 raised to the power exponent)

This is known as a floating point number (in the case of 16-bit a half float) and is the basis of how computers handle all non-integer numbers. If you have heard of the cinema encoding system ACES, it is also the storage encoding for that.

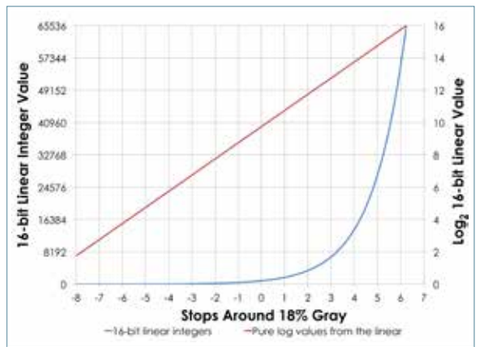

Fig 1: The raw data from a video sensor is linear, but calculating a log of that value gives a much more even spread of picture information from stop to stop.

There are still only 65,535 possible different values, but the encoding spreads them more evenly from stop to stop. The ACES spec covers a 33-stop range – far more than any cameras likely in the near future – and most video codecs are integer based with the two most common 'depths' (8-bit (0–255) and 10-bit (0–1023)) too small to make worthwhile floats. Instead, the principle of the even spread of values can be taken and scaled to fit the actual dynamic range of a camera within the integer ranges available. This is what is commonly known as a log curve.

Sony recently published (and then at the time of writing this article removed) an F5/F55/FS7 comparison chart which mentioned under raw recording: "16-bit offers seamless interoperability with half float workflow tools…", suggesting that linear raw uses half floats. If so, even linear is really log!

Maths over. Phew.

A bit about log flavours

Manufacturers often refer to their particular log curves as "Cineon based". Going back to the days of film scanners, Cineon was the system developed by Kodak for storing the wide dynamic range of scanned film data within a 10-bit digital format. It was a log encoding tuned for images and based on the response of the film itself.

> Setting exposure to a LUTted log recording gives you a degree of control over the final image rather than ceding it all to the colourist.

At this point I'd like to introduce some software I have been developing, both as a practical tool and also in order to help me to understand all of this. It's called **LUTCalc** and can be found at *www.lutcalc.net*. You can either download it as an offline web app or click where it says 'LUTCalc online' and use it like a webpage. Using either of these you can change recorded and output gammas (strictly 'transfer functions', but 'gamma' is what most people seem to understand and it's a bit pithier) on the left and see how they compare in the charts on the right. You can also see the effect on high and low contrast test images by clicking 'Preview'.

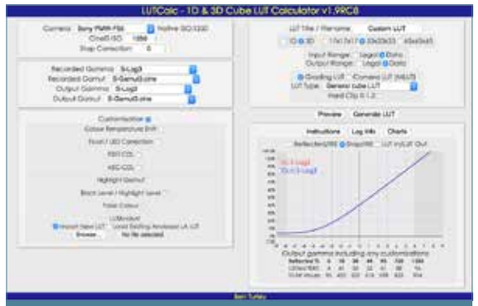

Fig 2: LUTCalc – changing the gamma settings in the second box down on the left will show their response curves in the chart on the bottom right. The 'correct' IRE values for the chosen output gamma are given in the table below the chart.

Cineon, Arri LogC, Sony S-Log3 and Panasonic V-Log are all very evenly spread and look flat and washed out with low contrast and thin blacks. Basically horrible. This brings up the crucial point that they are purely designed to convey information about an image as efficiently as possible, not to display the image itself. In an ideal world, no one would ever feel the need to look directly at a log recording.

They are not even good for setting exposure as the extreme low contrast means that a relatively large change in aperture makes for a small shift in the waveform or zebras. Log curves are meant to be viewed graded; whether fully by hand, using automated functions or lookup tables (LUTs), or more generally by a combination of both.

Setting exposure to a LUTted log recording (provided you know the characteristics of the LUT) is easier and more accurate than with the base log. In addition, it gives the camera person a degree of control over the final image, rather than ceding it all to the colourist.

Having said that, it's not an ideal world; software may not know how to do the 'prettifying' corrections automatically, grading may well be done by someone who isn't a colourist or cut back altogether through budget pressures, or the preferred codec may not be really suitable for log (i.e. 8-bit).

For those reasons there are other less data-centric log curves, namely S-Log, S-Log2, Canon C-Log and before them Panalog. These are all still log curves but the parameters chosen keep the shadows closer to conventional gammas with more of the even spread in the highlights above mid gray. This gives an image that tends to look a bit dark, fairly 'normal' in the shadows and midtones, and only milky in the high highlights.

Fig 3: Three common log flavours on low contrast and high contrast images. From top: Arri LogC, Sony S-Log3 and Sony S-Log2.

In the case of Canon it also allows a fairly wide dynamic range to be crammed into only 8 bits, though without the exposure shift flexibility of other formats. Gain is baked in with C-Log on a C300, with maximum dynamic range only at native ISO. The 'purist' approach to log recording is Sony's – the log recording captures the camera's entire dynamic range and the chosen ISO is just a piece of metadata information; post software is then expected to make appropriate exposure corrections. Great when it works but pictures that bounce up and down in brightness with ISO changes when it doesn't.

ARRI takes a third approach, adjusting the log parameters with ISO. This leads to a gentle knee in high ISOs to keep the full dynamic range without clipping.

All flavours of log are far better suited to 10 bits or deeper.

A bit about bits

The two most common bit depths in video codecs are 8-bit and 10-bit per channel; 8 bits can represent 256 values from 0 to 255 and 10 bits 1024 values from 0 to 1023. In video, the legal range from 0% to 100% IRE is represented as 16–235 in 8-bit and 64–940 in 10-bit. While YUV is a common colour space for codec storage, transfer functions and LUTs tend to operate on RGB data, which is also rather more intuitive to think in.

Fig 4: Not just a gray box!

With all that in mind, a difference of one between two colours in an 8-bit image will be invisible, which is why it has become so popular in both photography (jpegs, where it is referred to as 24-bit (8-bit by three colours)) and in video for capture, streaming and broadcast. Figure 4 is an 8-bit image featuring two different shades of gray one value apart. If you can spot the hidden message then very, very, very well done; if not, download it from *http://cameramanben.github.io/LUTCalc/images/fig4.png*, load it into Photoshop and have a play with the levels.

Once you start to manipulate an 8-bit image – pushing contrast, brightening some areas and darkening others – it is relatively easy to take things to a point where artifacts such as banding or a lack of detail can become apparent. Going from a very low contrast log capture to a (generally) much higher contrast finished Rec709 grade is really pushing 8-bit to the limit. 10-bit, with four times the range, has far more leeway in being pushed and pulled around before artifacts become an issue. 8-bit CP Lock on the C300 introduced high dynamic range and log to many people, but 10-bit is the point at which it becomes a truly flexible tool.

Flat pictures and noise

The flat contrast and lifted shadows of S-Log3 or LogC cause shadow noise to be shifted upwards, closer to the midrange that our eyes are attuned to looking at. Also, where the midtones and highlights are much flatter than in a conventional image, the contrast in the shadows is actually increased. This leads to a common belief that S-Log3 is noisier than S-Log2. The reality is that the underlying sensor noise is identical between the two. Corrected, the noise of S-Log3 will disappear back into the shadows, just as with S-Log2. Log is not supposed to be viewed untouched.

I have heard a number of people suggest that, while they generally shoot S-Log3, they switch to S-Log2 for high ISO, as it looks less noisy. Shooting on a Sony in CineEI without baking anything in, this logic is really the wrong way around.

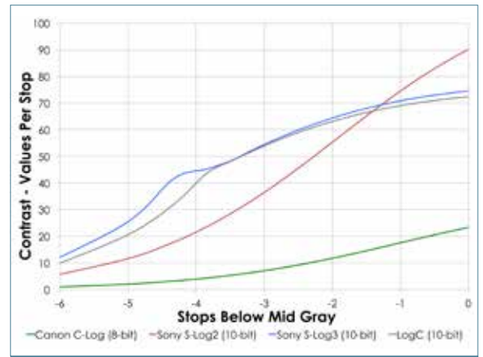

Fig 5: Contrast in the shadows – a higher line means more detail stored, but noise will be more obvious before correction. The Canon line is low because the C300 records 8-bit; the 10-bit curves exceed Canon's stop zero (18% gray) value deep in the shadows.

In normal CineEI the entire dynamic range is captured by whichever log curve is chosen, with the camera operating at a nominal ISO of 1250 for an F55 or 2000 for an F5 or FS7. Whatever CineEI ISO has been set does not affect the recording itself, but sets an item of metadata which is supposed to be picked up by the post software and only then used to make exposure adjustments. For S-Log2 in

10-bit this means that a two-stop push will be working from a rather smaller range of values than for the equivalent shot in S-Log3.

In reality, within a two-stop push, the difference between the S-Logs is not particularly troublesome, but the point is that if the general set-up is S-Log3, there should be no merit in switching to S-Log2 for high ISO sequences.

In 8-bit MPEG2, log is not a terribly good idea anyway, but S-Log2 in CineEI at high ISO is a very bad idea. One way to improve things is to bake in an S-Log2 monitor LUT (MLUT), but that will prevent further MLUTs and reduce the dynamic range as the ISO is moved up.

Does it really need to be log?

By now I hope it is a becoming more clear why log curves exist and how they work to squeeze a wide dynamic range into limited codecs. Between clever marketing, the widespread success of the C300 and articles such as this one, all the cool kids know they need to be shooting log or raw, or both, to be getting the most out of these cameras, but is this really true?

My personal feeling is that no colourist is likely to complain about being handed beautiful, consistent conventionally shot rushes to buff to a perfect shine. You are also unlikely to get complaints about richly coloured greenscreen shots which fill up a narrower dynamic range curve and key cleanly with minimal artifacts rather than muted, milky log material which barely dips into the dynamic range available and risks compression artifacts – particularly if shooting an 8-bit codec, which I suspect the majority of us still do for established workflows.

Log is great if you know that production have the time and resources to handle it well and if you as the cameraperson are able to have some kind of relationship with whoever is going to be turning that log picture into something beautiful. The flip side is that, if not handled well, log can lead to many headaches and, by its very flexible nature, the final image can easily end up utterly disconnected from your artistic intention but with your name still on the credits.

As a first step I would suggest developing your own, preferably small, set of 'looks' generated as LUTs. You can do this in grading software such as Resolve Lite, Sony Catalyst Browse, Amira Color Tool or my own LUTCalc. Some people like to roll their own from scratch but I think it's much easier to develop from established looks.

Fig 6: Rec709 variants: top – Sony's Alexa-alike LC709A; bottom – Rec709(800%) with Rec709 colour.

> Burning in removes the danger that someone who can cause you trouble will see unprocessed or poorly handled log rushes. Nobody ever got sacked for good dailies.

Once you have a look, you can generate versions suitable for use as MLUTs and, if need be, larger, more tight-fitting LUTs for use in grading software. Sony cameras work with the widely supported cube format. Sony's LC709 and LC709A look profiles are just 33x33x33 3D LUTs, which shows that larger LUTs are not always worthwhile unless the corrections performed are complex and very specific.

With the AMIRA, depending on your licence, you may be able to use either LUTs, the ASC-CDL (American Society of Cinematographers Color Decision List) controls, or a combination of the two. The ASC-CDL provides lift-gamma-gain controls (they call them offset, power and slope) along with saturation, and the camera both generates a MLUT and stores the details in metadata for grading software to automatically apply in post... hopefully.

Once you have a look or two that you are happy with, you can shoot log with your MLUT and then supply a copy with the rushes for post to use or ignore, or you can bake in your LUT. Both Sony and ARRI allow for this second option.

The proof is in the baking

Baking in irreversibly changes the recording. It is no longer a log recording. What is the point then? Why not use a conventional video approach of knee, black and matrix controls?

The answer is that there is no reason not to use the conventional approach if you have a set-up you are happy with. It is simple, well established and, as in Sony's custom mode, may include additional processing such as sharpening, which may not be available in log mode.

On the other hand, LUTs open up the flexibility of modern grading software to the camera itself. Obviously, you can't have local effects such as power windows or effects that mix together pixels such as blurs, but in colour terms the only limitation is the resolution of the 3D LUT and, at the reasonably standard 33x33x33 size, this affords a great deal of flexibility. Baking in also fixes ISO shifts, avoiding issues of post software failing to read and apply metadata.

As a cameraman, I can think of two other big advantages to burning in a look. First is that I can be confident that the edit will at least start out looking how I intend things to be. Second, it removes the danger that someone who can cause me trouble will see unprocessed or poorly handled log rushes. Nobody ever got sacked for good dailies.

Care tips for home colourists

I have a Sony F55 and, on the whole, I personally keep two looks on the go: one based on the Sony LC709A look profile (which is itself based on the ALEXA Rec709 look) and

> Is log recording a useful tool? Definitely. Is it a useful tool for every job? Definitely not.

a second which is a rescaling of the first to record all the way up to 109% IRE in extended range. Both capture the full dynamic range of the camera but with the contrast rolling off into the highlights rather than evenly reaching clip as with the underlying log.

Fig 7: Top: Sony LC709A; Bottom: LC709A rescaled to clip at 109%. Less headroom in legal range, but brighter 'correctly' exposed skin tones and less work in post to make a useable image.

The standard Sony LC curves record between 2% or 4% IRE and 100% IRE. As such, mid gray and skin tones sit rather lower than they do in 'conventional' Rec709. This is in order to squeeze in six and a bit stops above mid gray without making the knee overly harsh. By comparison, Rec709 (800%) and Hypergammas 7 and 8 cover a stop less highlight range and record to 109%, so that their legal range highlight headroom is more like 3 to 3½ stops. By rescaling my LC LUTs to 109%, skin tones sit closer to their conventional levels, I have around four stops of headroom in legal range and the remaining two and a bit stops are still there in extended range for a colourist to play with.

When shooting in a contrasty, high dynamic range environment I can use the standard version, but where all that headroom is excessive I can switch to the extended range version, knowing that whoever tweaks the pictures will have a bit less to do to even things up.

As an added bonus, when baking in, the extended range version can be pulled down in ISO (in the case of my F55, to ISO 400) and still reach 100% IRE. I use low ISO both for reducing noise and as 1/3 stop NDs for aperture control.

I also like the interpretation of Rec709 colours in LC709A; it is rather more complex than a simple matrix could achieve, and reasonably forgiving of the greeny compact fluorescents that are ubiquitous now. I do switch between more saturated versions depending on the job. The 'ASC-CDL' control in LUTCalc allows you to dial colour saturation up or down.

Wrapping things up

I hope that all this has made reasonable sense and covered some of the conflicting opinions about log that the internet turns into facts and confusion. There may well be a correction or two to be made, but the factual basis has come from a great deal of reading along with experimentation, both in the camera and on the computer. Getting a bit of software to work is a great, if frequently frustrating, way to confirm understanding.

If nothing else, you should now have plenty of useful material for ending conversations at parties.

Is log recording a useful tool? Definitely. Is it a useful tool for every job? Definitely not.

In a workflow where the cameraperson has a relationship with the colourist, and production has the time, budget and inclination to take on the added complexity, log is great. Where the rushes are being handed over to a post-production black hole and the reason for using it is because log is 'better', it is not a recipe for showreel material.

Within a couple of years, things should be a lot more bulletproof but today – particularly considering the number of editors and directors wedded to FCP7, and production companies and edit houses not ready or able to upgrade large systems – log support is spotty and reliant upon the skills of those working with it in post.

On the subject of FCP7, Apple recently and unexpectedly released an update to the Pro Codecs for Quicktime, so FCP7 can now work directly with 10-bit XAVC files from Sony's F cameras. If you have the option of 8- or 10-bit and intend to shoot log, 10-bit is the only worthwhile way to go.

Are conventional approaches or burning in your looks lesser alternatives? Well, if Gordon Willis had shot *The Godfather* in LogC on an ALEXA, the producers could have quietly fixed all those shadowed eyes, put the twinkle into Brando and we'd all possibly be leaving more lights on today. Hmmm.

**After a suitable period of time (two days as I recall), Daddy caved in and got the other toy as well. Valuable lesson learned though!*

***This is a simplified description of a half float – there are special arrangements for zeros, +/– infinity and division by zero / not a number (NaN).*

Fact File

Ben Turley first lugged boxes in the 80s for his father, documentary cameraman Patrick Turley. He spent much of the 90s lugging boxes for a wide array of camera people in documentaries, commercials, dramas, features and corporates, on film and video. This millennium he has been lugging boxes into and out of locations where he also turns on lights and points cameras.
Personal website – *www.turley.tv*

LUTCalc (*www.lutcalc.net*): offline (web app), online (webpage) and Mac OSX native versions available with instructions.The page also features a list of links to numerous standards papers relating to log, gamma and colour spaces.

Apple Pro Video Formats v2:
https://support.apple.com/kb/dl1396