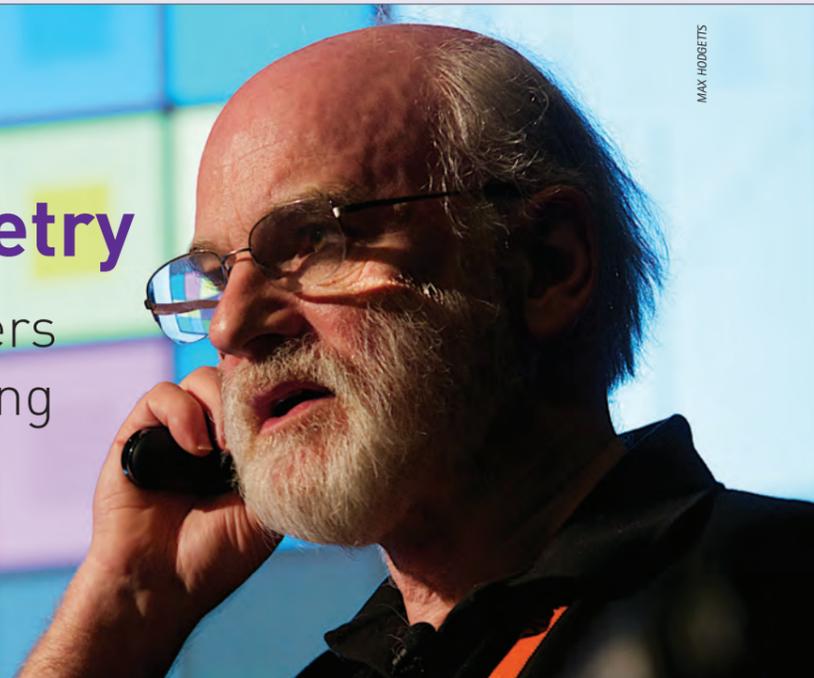


Practical Spectroradiometry

An introduction for users of the Television Lighting Consistency Index

by Alan Roberts



MAX HODGETTS

The ability of luminaires to reproduce colours faithfully compared to a natural light source has for many years been measured with the colour rendering index (CRI). In this system, a tungsten lamp is considered to have a CRI of 100% but other lamps fall short, showing an incomplete spectral output and failing to reproduce colours correctly (see also pages 66–68). Even with highly rated CRI LED lighting, problems can occur when the response of the film emulsion or digital sensor does not match the expected wavelengths of the lights, and correcting for this in post can be very expensive and time-consuming. So, an alternative to the CRI has been needed which assesses not only the output of the lamp but also takes into account the response of the TV system.

To address this problem, former BBC engineer and colour sciences expert, Alan Roberts, has worked with the EBU (European Broadcasting Union) to develop the Television Lighting Consistency Index (TLCI-2012), which measures the spectral power distribution of a test luminaire with a spectroradiometer, plus a software analysis program.

At a recent GTC workshop, held in association with the Society of Television Lighting and Design and the EBU, Alan Roberts explained how the new index works and some of the pitfalls of achieving accurate results using a spectroradiometer. For those of you not able to attend the workshop, here is the paper Alan delivered on the day where he explains in detail the processes involved and how to use the new TLCI software.

Introduction

Spectroradiometry (also known as 'spectrometry') is the science of the measurement of the power distribution of radiated energy. It is not restricted to visible light and is used in the analysis of all radiated energy, from X-rays through to heat.

The earliest experiments in spectroradiometry used the refractive properties of glass, in the form of a prism. Isaac Newton proved that light is comprised of a range of different colours by shining a beam of white sunlight through a prism (Fig 1), whereupon a rainbow of colours became visible. This spreading of the spectrum is called dispersion. His proof involved the shining of this rainbow through another prism, showing that no further splitting took place. The problem with using this refractive property of glass is that the bending angle is not linearly related to the wavelength of light, and so it can be hard to make precise measurements.

It was only when advancements in machine-tool design made it possible to make finely ruled diffraction gratings that spectroradiometry became a reliable science. Gratings have the interesting property that the angle through which light bends by reflection is proportional to wavelength. Thus it is easy to identify wavelengths accurately, using only geometry. The pitch of the engraved lines controls the spread of the output spectrum: the finer the pitch the greater the spread. The precise shape of the grooves can be modified to optimise the spectral output over a selected wavelength range; this is called 'blazing', and if done at about 500nm maximises the output over the range from about 330 to 1000nm, which is typical for the measurement of visible light. However, this blazing results in a 'peaky' response, falling rapidly at the extremes of the wavelength range. So, efficiency comes at the expense of bandwidth. The rulings on a blazed grating resemble a sawtooth pattern (Fig 2).

Gratings reflect not just one spectrum, but many. The main output (order 0) is usually fairly strong compared with the others, but the higher order outputs can overlap the main one, causing confusion if the spectral content covers more than one octave (2:1)

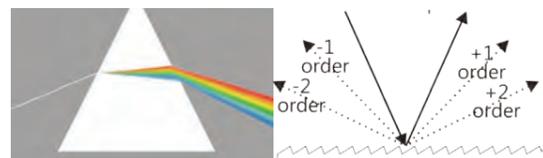


Fig 1: Prism showing dispersion; Fig 2: Diffraction grating

in wavelengths, since they occur at octave spacing. While blazing can increase the efficiency of the 0 order main output, it helps immensely by reducing the efficiency of the higher orders. Gratings don't have to be physical. A holographic image of a grating can also be used – in fact, anything that reflects or transmits light and involves some form of light interference will do the job.

Types of spectroradiometer

There are two main types of grating spectroradiometer, but both are used to spread the spectral power of the incoming light to be measured.

The heart of a sequential spectroradiometer is a monochromator, which has a grating mounted on a rotating pivot (Fig 3). Light from the entry slit is formed into a parallel beam by a focusing concave mirror, to fully illuminate the grating. The spectrally spread output is focused onto the exit slit. A sensor is placed at the output slit, or connected via relay optics, and receives only narrow-band energy. The wavelength directed to the output is changed by rotating the grating on its vertical axis. Measurement can be slow because wavelengths are measured successively, but the sensitivity can be

very high because only one detector is needed.

A snapshot spectroradiometer measures all wavelengths at the same time (Fig 4). Light from the entry slit is, again, focused onto the grating, but the spread spectral output is directed onto a line-array sensor instead of through a single output slit. The sensor has many light-sensitive photocells, which can be read out separately, thus giving the power distribution directly. This operation is very rapid, but the sensitivity is usually quite low.

Problems in spectroradiometry

Both types of spectroradiometer suffer from similar problems:

Calibration. It is not enough for the data output of the spectroradiometer to correctly indicate the wavelengths; the light level at each wavelength must also be accurately indicated. If the grating has been 'blazed' to improve efficiency, then it will have a humped wavelength response, peaking typically at around 500nm and falling steeply below 400nm and above 700nm. Failure to correct for this leads to very misleading results. The calibration must include any component in the light path that might affect performance.

Sensor and electronics noise. The line-array sensor is, essentially, a video or stills camera, operating in only one dimension. It comprises photodiodes which are normally reverse-biased to prevent current flow. In this state, each diode can intercept photons to change the state of electrons, effectively converting light into charge. As the amounts of charge are spectacularly low, special electronics is needed, which will always generate some noise.

Dark current. The silicon sensors are diodes, therefore there will always be some leakage current. This appears as a lower limit to the output signal when it is read. Since the current is proportional to temperature (absolute, K), the current increases as the temperature rises, giving a degree of unpredictability with temperature as well as unpredictability from diode cell to diode cell. This appears as a fixed pattern of variation, totally independent of the light levels, but dependent on the exposure or integration time.

Internal flare. Ideally, the light follows the paths shown and only those paths. But there is always some light which falls onto the internal structure of the device, or is not fully focused. Ideally, all such wasted light should be fully absorbed before it reaches the detector/sensor, but some is inevitably scattered and arrives at the output slit or at the sensor, causing light pollution.

High-order spectra. This is only a problem when the range of wavelength content in the input light exceeds an octave. It cannot be cured in a snapshot spectroradiometer, but can be dealt with in a motorised device by the addition of extra filters, making the device much more complex and expensive.

Directionality. The entry port to the spectroradiometer is highly directional; it is most sensitive to light arriving in exactly the direction of the onward path of the light within it. Thus, there must be some means of intercepting the light to be measured before it enters the spectroradiometer and ensuring that it is fed to the input port in exactly the right direction. Failure to do this will give strange and unrepeatable results.

Linearity. Strictly, it is potential non-linearity that is the problem. In theory, if you change the exposure (by changing either the actual light level or the exposure duration) then the curve shape should not change. It is unusual for this to be a problem, but not unknown.

Cures for problems in spectroradiometry

Calibration. For wavelength calibration, either a set of known emitters (e.g. lasers) or a single emitter exhibiting peaks at known wavelengths is needed. The simplest solution is to take an ordinary fluorescent tube, which uses mercury vapour to generate radiation that is partially intercepted by a phosphor coating to emit visible light. The spikes of the mercury emission are clearly visible, and can be used to confirm the accuracy of wavelength calibration or provide data for a calibration process (Fig 5). The known peaks of emission for mercury vapour occur at 184.5nm, 253.652nm, 296.728nm, 302.150nm, 313.155nm,

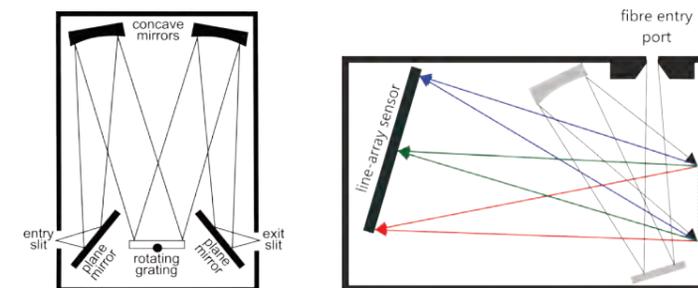


Fig 3: Sequential monochromator; Fig 4: Snapshot spectroradiometer

334.148nm, 365.015nm (the I line), 404.656nm (the H line), 407.783nm (the G line), 497.604nm, 546.074nm, 576.960nm and 579.066nm. The peaks at 545, 576 and 579nm are particularly visible since these are powerful and green, and are the basis of the 'green spike' which characterises fluorescent lamps.

For amplitude calibration, a light source whose spectral output has been measured and tabulated by an organisation that can trace its processing back to a standards body, such as NIST (National Institute of Standards and Technology) or NPL (National Physics Laboratory) is required. This known, standard, light source has a spectral power distribution S_λ , where λ denotes wavelength in nm which will be supplied by the calibrating authority as a text listing or text file, giving the light power at a series of wavelengths covering at least the range of visible light. A measurement of this source is taken, C_λ , with all the normal precautions. The responsivity R_λ of the spectroradiometer must be the ratio of C to S, over the wanted range of wavelengths:

$$R_\lambda = \frac{C_\lambda}{S_\lambda}$$

Any subsequent measurements M_λ of light sources can be normalised to deliver the actual spectral power distribution P_λ :

$$P_\lambda = \frac{M_\lambda}{R_\lambda} = M_\lambda \frac{S_\lambda}{C_\lambda}$$

Since the mathematics of the TLCI-2012 makes calculations over the visible range from 380nm to 760nm in 5nm steps, it follows that the calibration must cover at least that range in 5nm steps. Typically, the responsivity is not flat, and the resulting compensation to obtain a flat response can emphasise noise, dark-current variations and higher-order spectral content (Fig 6).

Sensor and electronics noise. There are two ways to minimise the effects of dynamic noise: spectral filtering and temporal filtering.

Spectral filtering means increasing the spectral bandwidth of the data values in the spectral power distribution. For example, the ASEQ LR1 spectrometer has 3500 light-sensitive cells in its sensor, each cell delivering the light level over only about 1/7nm, even though the actual spectral bandwidth of the image is about 2.5nm. Thus, the sensor and electronic noise can be reduced by adding together many adjacent cell values to form the 5nm values required by the TLCI. This is done automatically in the conversion from raw data to SPD and LUM files used in TLCI calculations. The user has no control over this.

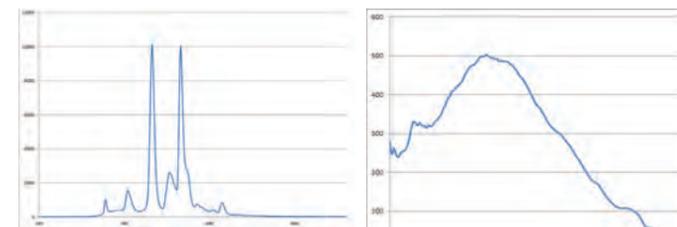


Fig 5: Fluorescent spectrum; Fig 6: Spectroradiometer wavelength responsivity

Temporal filtering is the only available process over which the user has any control. Since the exposure time of the measurement must be set to give an acceptable signal level, reasonably filling the coding range of the spectroradiometer, the only additional control is to make several measurements and average them. Thus, for a 'quick and dirty' measurement a single scan can be used, but for serious measurements it makes sense to average as many measurements as possible. Taking an average of two measurements will reduce the noise level by 3dB. Averaging four reduces the noise by 6dB; eight reduces



Fig 10: An integrating sphere

it by 9dB, and so on. The ASEQ LR1 allows for up to 64 measurements to be averaged.

Dark current. This cannot be cured, but can be compensated for. Each photocell leaks a little current, even in the absence of light, and that amount is relatively consistent with time but varies as a function of temperature. The solution is to measure the output of the spectroradiometer with the input port closed, or capped, then subtract that measurement when making subsequent measurements.

Since this process will increase the noise level, it should ideally be done using averaged measurements

(Fig 7). Here, the green and brown lines centred at around 500 are plots of measurements of dark current with the spectroradiometer input capped to eliminate light. Clearly, the lines are very similar and so large variations must be due to dark current, while small differences are due to noise. This is confirmed by the lines centred at zero, which are measurements of the capped input of the spectroradiometer, with dark-current compensation. Note that there is considerably lower amplitude in both these signals, indicating that the dark-current variations have been removed and only electronic noise is left. The green line is a later dark-current measurement of the spectroradiometer, capped, when the temperature of the unit had risen after about 30 minutes. Measurement of the dark current must be taken at the same time as the measurements, and with the same exposure time: later measurement will not do, even though the distribution is similar.

Internal flare. All internal surfaces, except those in the wanted light path, must fully absorb light. This is impossible. Black anodising does not work well as it is shiny. Black paint is not good enough as it reflects as much as 90% of long-wave energy, infra-red. Matt black paint is rather better, but can be deceptively shiny at long wavelengths. Black velour cloth absorbs reasonably well except in the far red. Special optical paints are available for the treatment of cameras and telescopes, but even these do not eliminate flare.

High-order spectra. This cannot be solved in a snapshot spectroradiometer. The only thing that can be done is to know how to recognise it and avoid being confused by it. In Figures 8 and 9, measurement has been made of a LED luminaire, with dark-current suppression but not normalised. The LED emits a narrow band of light centred at 450nm. Much of the blue light is absorbed by a layer of phosphor and re-emitted as yellow/amber, centred at about 545nm. However, there is a higher-order spectral component just visible, a peak at 900nm and a lower, broader peak at 1090nm (i.e. largely off the

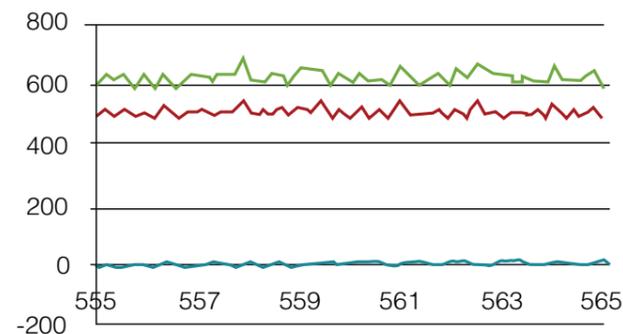


Fig 7: Background signal and noise

edge of the plot). The level of this harmonic content appears to be low, and could be thought not to be a problem, but this plot is not normalised and so does not properly describe the power distribution.

When the normalisation process is completed, the effect of the high-order content is much more dramatic, although this should not be a problem for TLCI use because we need data only between 380nm and 760nm, and the level is quite low at about 780nm.

Since this effect is ever-present, it must also be occurring in the calibration process. This means that a grating spectroradiometer cannot give reliable results if the light source under test produces significant power output over more than an octave (2:1 in wavelength). For this reason, it is useful to have a prior understanding of the type of spectral power distribution to be expected, so that unusual effects can be explained and possibly ignored.

Directionality. There are two ways to ensure that light enters the entry port in the right

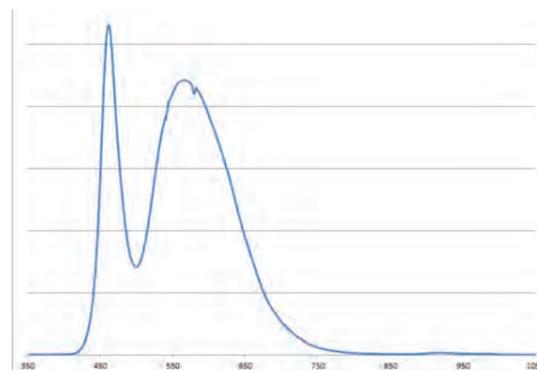


Fig 8: LED emissivity, uncalibrated

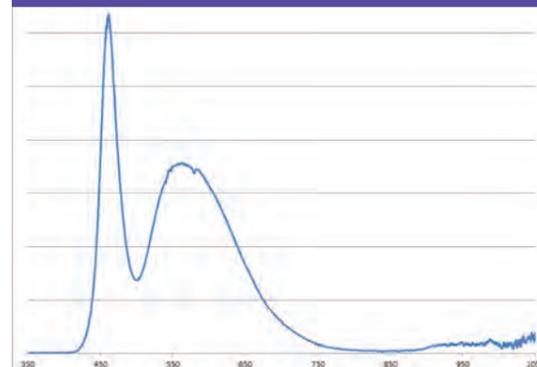


Fig 9: LED emissivity, calibrated

direction: using an integrating sphere and using a 'cosine corrector' diffuser. Both have the same effects: of scrambling any polarisation which might be in the input beam and providing an evenly lit surface at the entry port.

A reflective integrating sphere (Fig 10) (Ulbricht sphere) is always the best approach since it uses total internal reflection to evenly spread light in all directions, but for the same reason it requires a lot of light and spheres can be large and expensive. If the sphere is big enough, it can totally contain the luminaire being measured and then it is possible to measure total light output. When I was learning about light and measurement in the 1980s, the NPL in Teddington, UK, had a 5-metre diameter sphere, which is perhaps slight overkill. Much smaller spheres can be used as diffusers, with an entry port and exit port to the spectroradiometer, arranged such that neither port is visible from the other within it. I made an example for the workshop from a table-tennis ball (which is wholly inadequate but illustrated the principle): light enters via a large hole on the left, is scattered internally, and exits via a fibre at the bottom.

A simpler and cheaper alternative is to use a transmissive diffuser to intercept the light before it reaches the spectroradiometer input port. This can be much smaller, and needs only to cover the port, or the end of a glass fibre if this is used to route the light. The problem with diffusers is that they do not have uniform properties at all wavelengths: PTFE (polytetrafluoroethylene) diffuses well in blue, but in the far red is almost transparent so separate calibration is needed. Since this means the directionality is wide at blue but narrow at red, the diffuser must be used only in the direction at which it is calibrated. There is no perfect transmissive diffuser (Fig 11).

Linearity. There is no guaranteed cure for electronic non-linearity, unless any such non-linearity can be measured and correction applied. If non-linearity is suspected, then a set of tests can be used to establish the limit of linearity and used as a limiting value for measurements. The test is quite simple: expose the spectroradiometer to a known, stable, light source and make a series of measurements at ever increasing, or reducing, exposure durations. Then, if there is any change in the curve shape, an estimate can be made of the maximum usable amplitude range.



L SERIES

CUT. SHAPE. FOCUS. TUNE.

ARRI presents two additions to the popular L-Series of LED Fresnels: the L7-TT is a tuneable tungsten model that is 20% brighter than the L7-C, while the active cooling option for all models reduces size and weight for location shoots, portable kits, events and studios.

ARRI L-SERIES. TRULY CINEMATIC.



www.arri.com/qr/zm/l-series



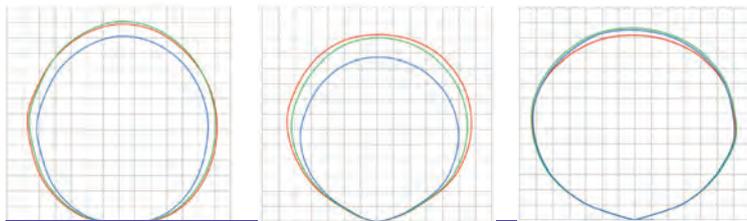


Fig 11: Diffuser polar diagrams (a) PTFE; (b) ASEQ LR1; (c) Ocean Optics CC-3

Recommended procedure for making measurements

For this section, the ASEQ LR1 spectroradiometer connected to a laptop loaded with the TLCI software, will be used as an illustration, but the processes are the same for all spectroradiometers.

1. Connect the spectroradiometer and allow at least 15 minutes for the internal temperature to stabilise. Also power the light source being measured so that it stabilises.
2. If there is a linearity limit for the device, click in the outer area of the display (e.g. on any of the horizontal or vertical axis numbers) to open the 'X Y Scale' dialog, and set the Y maximum value (11,000 in this case) (Fig 12).
3. Uncap the input port, and allow the light source being measured to illuminate it, via a diffuser if appropriate. Check the 'Run' box to allow measurement. Set 'Boxcar' to a suitable number: 0 means you get fine resolution (a recorded value for every sensor photosite, very useful for checking wavelength accuracy or for identifying the peaks of spectral spikes), 10 means that values are averaged from 10 photosites before to 10 after (which gives a bandwidth of approximately 5nm, which is all that is needed for the TLCI, and produces lower noise levels) (Fig 13).
4. Adjust the 'Exposure (ms)' (integration time) such that the measured spectral data approximately fills the reliable amplitude range (this will reduce the electronic noise level; measuring with a lower exposure duration will result in the noise being more significant). Note that this will affect the background level, which must be subtracted.
5. Cap the input port of the spectroradiometer to block all light input. Set a number of measurements to be taken (e.g. 25) and check the box to average them. Now, click 'Get background' to make a measurement of the background signal and wait for it to complete. There is no indication of completion, but you should see the spectral curve change slightly when it completes (Fig 14).
6. Check 'Subtract background' and that the background signal level falls to near zero.
7. Uncap the input port, and allow the light source to be measured to illuminate it, via a diffuser if that is appropriate. Do not change the exposure duration, since this will change the background level. The signal level will be lower than initially, because the background has been subtracted electronically (Fig 15).
8. Uncheck 'Run' and save the measurement as a file (if you do not uncheck this box, the wrong data will be saved). Use this file in the 'ASEQ LR1 conversion' utility program.

This routine should be used whenever making measurements, whether for the purposes of calibration or for actual measurements. Consistency of method is vital if meaningful results are to be obtained. Some simplifications are possible, but undesirable, and are available in the 'Engineering' version of my 'ASEQ LR1 conversion' utility



Fig 12: X Y Scale dialog

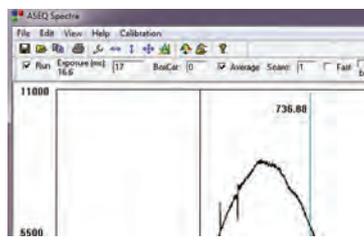


Fig13: Run, and set exposure time

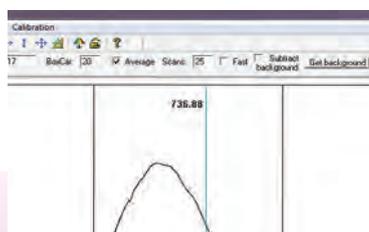


Fig 14: Average 25 measurements

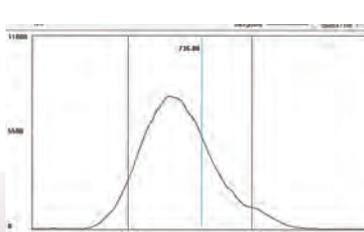


Fig15: Background subtracted

Further Reading

Spectroradiometry Methods. Application Note (A14), A guide to photometry and visible spectroradiometry, W.E.Schneider and R. Young, Optronics Labs Inc, January 1998 <http://biology.duke.edu/johnsenlab/pdfs/tech/spectmethods.pdf>

Diffraction grating, Wikipedia entry http://en.wikipedia.org/wiki/Diffraction_grating

Blazed grating, Wikipedia http://wikipedia.org/wiki/Blazed_grating

Determination of the blaze wavelength. Richardson Gratings, <http://gratings.newport.com/library/technotes/technote11.asp>

software; these tricks are not available in the 'User' version.

The background signal can be ignored completely, and compensation made in the conversion software. There is also a facility for estimating a single value to represent the dark current, which will, obviously, not correct for the dark-current variability, nor for changes in the integration time. And there is a facility for subtraction of a saved dark-current measurement, although this will not take into account any temperature variations. If flare is a problem, it can be estimated and compensated for. These methods can be effective, but are not infallible, and depend on intelligent guesswork on the part of the operator.

If flare is significant, i.e. above about 0.025%, then the measurements will not be wholly reliable. This is because each measurement will be effectively polluted with an equi-energy illuminant which will desaturate the test source. This will also happen during the calibration process, so the basic accuracy of the whole process will suffer. If the flare level reaches 0.1%, then the results can never be wholly trusted. The 'ASEQ LR1 conversion' utility software provides a means to estimate the flare level, and then apply a correction for it.

GTC Workshops Organiser Clive North adds:

Alan's workshop 'LED Lighting: Setting the Standard' on LED luminaires and his TLCI measurement techniques was an unusual event for the GTC. Organised by the GTC in conjunction with The Society of Television Lighting and Design and the EBU it was a highly technical event which, on the face of it, might well not have appealed to many cameramen. As a result I was delighted by the excellent GTC turnout and impressed by the interest those members showed in what was a highly technical presentation of a subject not immediately of obvious relevance to all cameramen.

In addition to Alan's presentation, DoP Jonathan Harrison shared his feelings on LED lighting versus tungsten and fluorescent sources (see pages 66-68) and Robert Yeo of Pro-Lite Technology thoroughly explained photoradiometry. Various luminaires were put through their paces and evaluated using Alan's technique and the results closely scrutinised. Products from ARRI, Gekko, deSisti, ETC, Dedo and Cineco were included in the tests.

Fact File

Contact Alan Roberts on: roberts.mugswell@btinternet.com

See more about the TLCI-2012 on the GTC website: www.gtc.org.uk/members-area/technical-resources/tlci-results.aspx