

EBU


OPERATING EUROVISION AND EURORADIO

TECH 3355 s1

ASSESSMENT OF THE COLORIMETRIC PROPERTIES OF LUMINAIRES (TLCI-2012)

SUPPLEMENT 1: AN INTRODUCTION TO SPECTRORADIOMETRY

Geneva
October 2013



Contents

1.	Introduction	5
2.	Types of spectroradiometer	6
3.	Problems in spectroradiometry.....	6
4.	Cures for problems in spectroradiometry	7
4.1	Calibration	7
4.2	Sensor and electronics noise.....	8
4.3	Dark current	8
4.4	Internal flare	9
4.5	High-order spectra	10
4.6	Directionality	11
4.7	Linearity.....	11
5.	Recommended procedure for making measurements.....	12
6.	Further Reading	14

Introduction to Practical Spectroradiometry for users of the Television Lighting Consistency Index

<i>EBU Committee</i>	<i>First Issued</i>	<i>Revised</i>	<i>Re-issued</i>
TC	2013		

Keywords: Spectroradiometry, spectrometry, TLCI-2012

1. Introduction

Spectroradiometry (also known by the abbreviation 'spectrometry') is a measurement science, a metrology, the measurement of the power-distribution of radiated energy. It is not restricted to visible light; it is used in the analysis of all radiated energy, from X-rays 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, 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.




Figure 1: Prism

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 is typically about 600/mm; the pitch 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 500 nm it maximises the output over the range from about 330 - 1000 nm, 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. Thus efficiency comes at the expense of bandwidth. The rulings on a blazed grating resemble a sawtooth pattern.

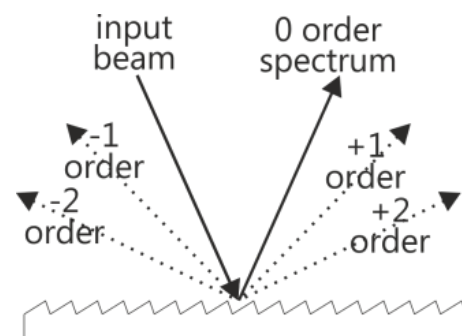


Figure 2: Ruled Grating

Also, 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) in wavelengths, since they are occur at exactly octave spacing. While blazing can increase the efficiency of the wanted output, it reduces the efficiency of the other orders which helps immensely.

Gratings don't have to be physical, a holographic image of a grating can also be used, in fact anything which reflects or transmits light and involves some form of light interference will do the job.

2. Types of spectroradiometer

There are two basic types of grating spectroradiometer, but both do the same job, spreading 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. 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.




Figure 3: Sequential monochromator

A *snapshot spectroradiometer* measures all wavelengths at the same time. 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 photo-cells which can be read out separately, thus giving the power distribution directly. Thus operation is very rapid, but the sensitivity is usually quite low.

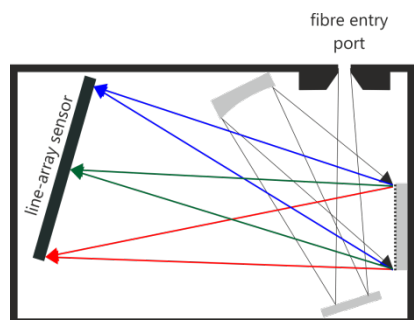


Figure 4: Snapshot spectroradiometer

3. Problems in spectroradiometry

Both types suffer from the same types of problem:

Calibration. It is not enough to have the data output of the spectroradiometer to correctly indicate the wavelengths, it also has to correctly indicate the light level at each wavelength. If the grating has been 'blazed' to improve its efficiency, then it will have a humped wavelength response, peaking typically at around 500 nm and falling steeply below 400 nm and above 700 nm. Failure to correct for this leads to highly misleading results. The calibration must include any component in the light path which could possibly affect the 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. Thus charge is read into an amplifier when the voltage on the diode is reversed. Since the amounts of charge are spectacularly low (10^{-12} to 10^{-15} Amps, 1 pico-amp to 1 femto-amp) special electronics are needed, which will always generate some noise.

Dark current. The silicon sensors are diodes; therefore there will always be some leakage current

even when they are properly reverse-biased. 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 dependant 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 a problem only when the range of wavelength content in the input light exceeds an octave. It cannot be cured in a simple snapshot spectroradiometer. It can be cured at a price, by the addition of extra filters, which makes the device 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 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 which is the problem. In theory, if you change the exposure (either by changing the actual light level, or by changing the exposure duration) then the curve shape should not change. It is unusual for this to be a problem, but not unknown.

4. Cures for problems in spectroradiometry

4.1 Calibration

For **wavelength calibration** either a set of known emitters is needed (e.g. lasers) or a single emitter exhibiting peaks at known wavelengths. Probably the simplest solution is to use an ordinary fluorescent tube, which uses mercury vapour to generate radiation which is partially intercepted by a phosphor coating to emit visible light. The spikes of the mercury emission are always clearly visible, and can be used to confirm the accuracy of wavelength calibration, or to provide the data from which a calibration process can be made.

The known peaks of emission for mercury vapour occur at 184.5 nm, 253.652 nm, 296.728 nm, 302.150 nm, 313.155 nm, 334.148 nm, 365.015 nm (the I line), 404.656 nm (the H line), 407.783 nm (the G line), 497.604 nm, 546.074 nm, 576.960 nm and 579.066 nm.

The peaks at 545, 576 and 579 nm are particularly visible since they are powerful and green, and are the basis of the 'green spike' which characterises fluorescent lamps.




Figure 5: Fluorescent spectrum

For **amplitude calibration**, a light source is needed whose spectral output has been measured and tabulated by an organisation which can trace it's processing back to a standards body, such as NIST or NPL.

This known, standard, light source has a spectral power distribution S_λ , where λ denotes wavelength in nm. It will be supplied by the calibrating authority as a text listing or as a 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 , taking all the normal precautions. Then 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 TLCI-2012 mathematics makes calculations over the visible range from 380 - 760 nm in 5 nm steps, it follows that the calibration must cover at least that range and in 5 nm steps.

Typically, the responsivity is far from flat, and the resulting compensation to obtain a flat response can emphasise noise, dark current variations and higher-order spectral content.




Figure 6: Spectroradiometer wavelength responsivity

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

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 2 measurements will reduce the noise level by 3dB. Averaging 4 reduces the noise by 6dB. Averaging 8 reduces by 9dB, 16 by 12dB and so on. The ASEQ LR1 allows for up to 64 measurements to be averaged.

4.3 Dark current

This cannot be cured, but can be compensated for. Each photo-cell 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, and then to subtract that measurement when making subsequent measurements. Since this process will increase the noise level, this should ideally be done using averaged measurements.

Here, the two uppermost 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 the large variations must be due to dark current, while the 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.

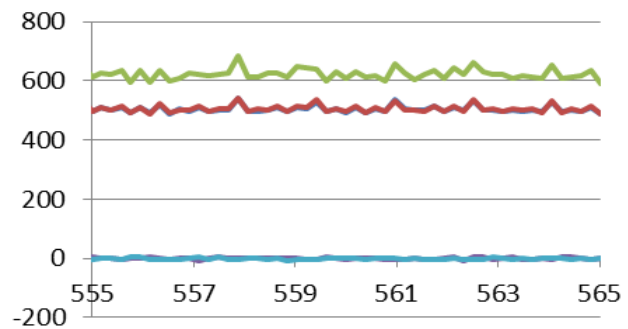


Figure 7: Background signal and noise

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 uppermost line is a later dark-current measurement of the spectroradiometer, capped, when the temperature of the unit had risen after about 30 minutes.

Clearly, 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.

4.4 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, it is shiny. Black paint is not good enough since 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. Here, a layer of self-adhesive black velour material has been stuck to the black anodised metal interior.



Figure 8: ASEQ LR1 internals

Note the difference from the shiny metal, indicating that light is being absorbed rather than scattered. Even with only this amount of treatment, the level of flare is halved in this spectroradiometer, and much better can be done with more full treatment.

Special optical paints are available for the treatment of cameras and telescopes, but even these do not eliminate flare (the manufacturers discuss the suitability of flour and poppy seeds for mixing to improve the performance:

<http://www.gerdneumann.net/english/instrument-building-parts-teile-fuer-den-fernrohrbau/totmatte-schwarze-optikfarbe-deep-black-optical-paint.html>).

4.5 High-order spectra

This cannot be solved cheaply in a snapshot spectroradiometer.

Here, a 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 450 nm. Much of the blue light is absorbed by a layer of phosphor and re-emitted as yellow/amber, centred at about 545 nm.

However, there is a higher-order spectral component just visible, a peak at 900 nm and a lower, broader peak at 1090 nm (i.e. largely off the 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 it should not be a problem for TLCI use because we need data only between 380 and 760 nm, and the level is quite low at about 780 nm.

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.

The effect can be cured by inserting a filter in the light path between the grating and the sensor.

This will have a 'long-pass' characteristic, of ever increasing wavelength. This can be simplified to a set of filters, perhaps 6, with successively increasing wavelength-pass features. If this is correctly placed, it will pass only the desired subset of wavelengths to each part of the sensor, blocking the higher orders.




Figure 9: LED emissivity, uncalibrated




Figure 10: LED emissivity, calibrated

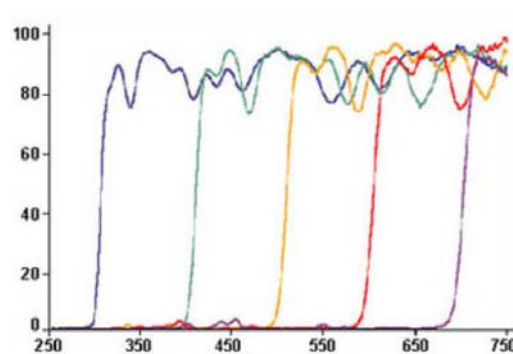


Figure 11: Order-sorting filter

4.6 Directionality

There are two ways to ensure that light enters the entry port in the right 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 (Ulbricht sphere) is always the best approach, since it uses total internal reflection within a sphere to evenly spread light in all directions, but it requires a lot of light for the same reason, and spheres can be large and expensive. If the sphere is large enough, it can totally contain the luminaire being measured, and then it is possible to measure to total light output.

In the 1980s, the National Physical Laboratories in Teddington, UK, had a 5-metre diameter sphere, which is perhaps slightly overkill. Much smaller spheres can be used as diffusers, with an entry port and an exit port to the spectroradiometer, arranged such that neither port is visible from the other within it.



Figure 12: Integrating sphere

The example given in Figure 12 is made from a table-tennis ball, which is wholly inadequate, but it illustrates the principle; light enters via the large hole on the left, is scattered internally, and exits via the 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 such is used to route light to it.




Figure 13: Diffuser polar diagrams

The problem with diffusers is that they do not have uniform properties at all wavelengths; PTFE diffuses well in blue, but in the far red it is almost transparent, and thus separate calibration would be needed. Not only that, but since this means that 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.

4.7 Linearity

There is no guaranteed cure for electronic non-linearity, unless any such non-linearity can be measured and a correction applied. Should non-linearity be 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: exposure 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.

Figure 14 overleaf shows the results of taking 10 exposures of a tungsten light source, without

moving anything, so the light level reaching the sensor is constant. On the left, it is clear that the curve flattens as exposure and signal level increases, therefore there is an electronic non-linearity somewhere. On the right, the curves are all normalised to the results for 40 ms exposure.

Clearly, the signals from the lower exposure levels are all linear, but for exposures higher than about 50 ms, non-linearity is an increasing problem.

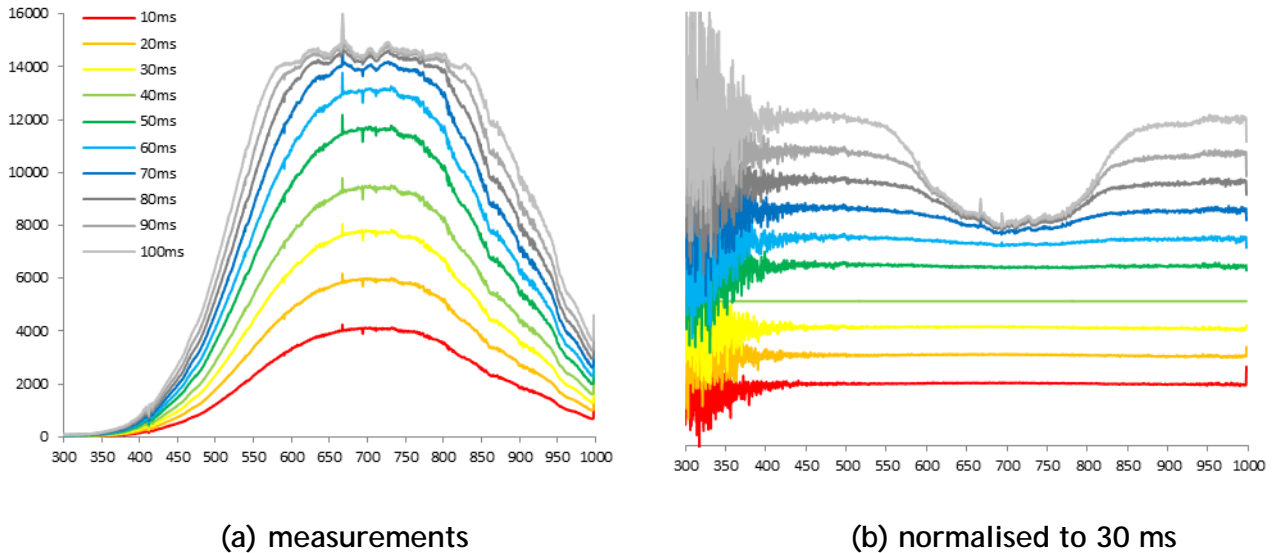


Figure 14: Linearity, 10 exposures

The signal level at 50 ms is just starting to show non-flatness, therefore it is reasonable to say that the linear limit for this device is the signal level achieved between 40 and 50 ms. Note that this is not a limit on the exposure duration, but on the signal that it results in. Therefore we can refer back to the left-hand diagram and estimate that the linear limit is between 10000 and 12000, of the 16384 possible levels from the 14-bit ADC (analogue-to-digital converter). In this case, the maximum advisable signal level should be taken as 11000.

5. Recommended procedure for making measurements

For this instruction, the ASEQ LR1 spectroradiometer will be used as an illustration, but the processes are the same for all spectroradiometers.

1. Connect the spectroradiometer and allow it at least 15 minutes for the internal temperature to stabilise. Also power the light source to be 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' dialogue, and set the Y Maximum value, 11000 in this case.

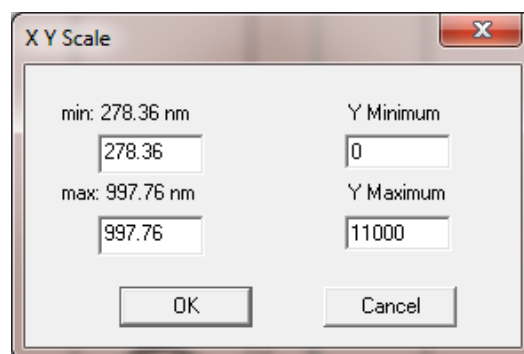


Figure 15: X Y Scale dialogue

- Uncap the input port, and allow the light source to be measured to illuminate it, via a diffuser if that is appropriate. Check the 'Run' box to allow measurement. Set 'Boxcar' to a suitable number: 0 means that you get fine resolution (a recorded value for every sensor photo-site, very useful for checking wavelength accuracy or for identifying the peaks of spectral spikes), 10 means that values are averaged from 10 photo-sites before to 10 after (which gives a bandwidth of approximately 5 nm which is all that is needed for the TLCI, and produces lower noise levels).

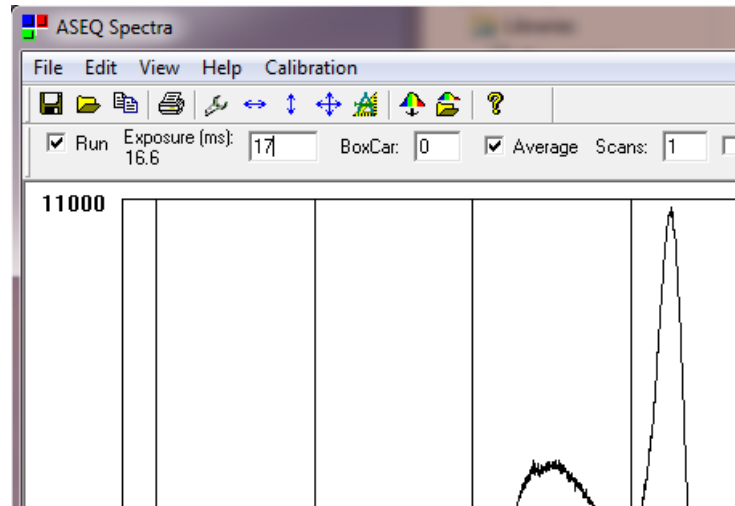


Figure 16: Run and set exposure time

- 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.
- 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.

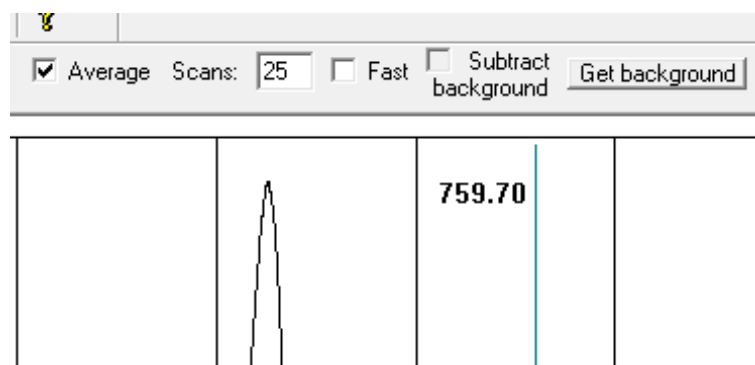


Figure 17: Average 25 measurements

- Check 'Subtract background', and check that the background signal level falls to near-zero.
- 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.




Figure 18: Measurement, background subtracted

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 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 also is 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.

6. 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 entry http://wikipedia.org/wiki/Blazed_grating |
| Technical Note 11 | Determination of the blaze wavelength. Richardson Gratings,
http://gratings.newport.com/library/technotes/technote11.asp . |