CMOS-testing and configuration

Abstract

The observation and recording of long sequences of high-quality images required for Exoclock measurements of exoplanet transits put major demands on the observers and therefore on the instrumentation used. The availability of high-quality and low-cost CMOS (or sCMOS) cameras has opened wide fields of astronomical imaging to many amateur astronomers. However, it was recognised that the complex features now available with these cameras presented a veritable zoo of possibilities, some of which might well interfere with the objective of high-quality photometric observations, as demanded by the Exoclock project and other fields, including variable star measurements and DSO observations. The top-end CMOS and CCD cameras currently being used by amateur and professional observers within Exoclock have thus been reviewed. In particular, user-controlled features such as gain / contrast etc. which can interfere with high-quality photometric imaging and also be confusing to the amateur have been assessed and explained. It is hoped to obtain four top-end cameras from QHY in the near future that will be tested in some detail under both laboratory conditions and on various telescopes of 20 – 40 cm aperture that are readily available to the Exoclock observers. The Report is primarily aimed at the Exoclock observing community. However, it is expected that the calibration procedures discussed and presented, as well as the general conclusions, will be of use to the broader amateur astronomy community.

Introduction

This document is intended primarily to aid the testing and configuration of modern CMOS cameras. Much, if not all, of the testing can also be applied to CCD sensors but some of the attributes and uses of CMOS sensors discussed at the end may not be easily applicable to CCD sensors.

Whilst also relevant to colour CMOS cameras (DSLR and the like) intended for general photography, it is aimed primarily at monochromatic, astronomy cameras using either CMOS or sCMOS sensors. For the purposes of this document the name CMOS will be used to indicate both CMOS and sCMOS sensors, unless otherwise indicated. For readers who are unclear as to the main differences between CMOS and CCD technology it is suggested that they consult the web, which contains many references on this.

The document is not intended to promote CMOS over CCD sensors since CCDs are more than suitable for Exoclock-type work but, since CMOS is a relatively new technology for more complex astronomical use, it was felt necessary to look at this new technology in more detail with a particular emphasis on photometry.

Glossary of terms

1	Well	The collection area of the silicon which converts photons to electrons.
2	Full Well	The maximum number of electrons that can be contained in the Well.
3	Gain	Amplification applied by the electronics to the analogue signal contained in the Well
4	ADC	Analogue to digital conversion of the signal into machine-readable units
5	ADU	Analogue (to) digital units. This expresses the value of the well per single digital unit when the analogue signal is digitised in an analogue to digital converter (ADC). This value depends on the digitisation used-normally 12, 14 or 16 bits (4096, 16384 or 65536 units).
6	eGain	The number of electrons corresponding to a single ADU. This is Gain dependent.
7	Gain Index	A number through which the manufacturer expresses Gain in dimensionless units (e.g. as a percentage or a number).
8	Read Noise	the noise added by the electronics when the analogue signal is read
9	Bias current	A current applied by the electronics to each pixel to bring the value of each signal slightly above zero
10	Dynamic range	general term denoting the ratio between a minimum and maximum signal
11	Noise	Part of the signal which represents the uncertainty within measuring a signal
12	Offset/Black level	This is a charge applied to the sensor by the electronics to ensure that no pixels are at or below zero. This ensures that all incoming photons are converted to electrons and are thus measured by the electronics

Overview and General Assumptions

Testing methods and testing-related theory are reviewed, with emphasis on several key areas related to CMOS technology. However, the document is not intended to discuss basic sensor operation, except to the extent that it is related to actual testing or the specific configuration of a sensor during test.

The analysis provided applies only to sensors which are linear. That is, the output increases linearly with the number of incoming photons. These analyses are not, for example, generally applicable to intensified or to electron-multiplying cameras (EM-CCD).

The analyses also assume that only the dark current of the device is temperature dependent.

Several essential work areas are considered:

- Equipment required for Testing.
- Standard Calibration Techniques.
- Setting of Gain.
- Setting of Offset/Black Level in relation to Gain Indices (GI).
- Testing for Linearity.
- Measurement of Electronic Gain (eGain, eG), Read Noise (RN) and Full Well (FW).

- Effects of eG, RN and FW on Dynamic Range (DR).
- Practical implications of CMOS properties.

The cornerstone of the approach discussed is that it is entirely possible to characterise the basic performance of a camera using three simple tests-flat fields, bias frames and flat fields at varying light intensity or exposure time.

Please note that the test results discussed here are all based on the use of a Finger Lakes Instrumentation Kepler 400 TVISB. It should not be assumed that all of the results outlined here will apply to all CMOS cameras. The only way to understand the performance of other sensors/cameras is to conduct tests. The main tests are discussed here but, given the pace at which CMOS development is taking place and the diversity of CMOS sensors, it is entirely possible that other sensors may respond differently.

A series of rather more diverse tests is planned on QHY sensors being loaned for test and evaluation.

Equipment Required for Testing

Most tests can be conducted either on a test bench or with the camera on a telescope.

The ability to obtain accurate Flat Fields is also essential and this will be further discussed later on in the document.

Testing some of the practical implications will require that the camera be installed on a telescope to image and measure star fields.

It is entirely possible to conduct all of this work with the camera on a telescope, although more accurate and reproduceable results would most likely be obtained on a test bench. The methodology discussed assumes that the tester has an appropriate way of obtaining good dark fields, flat fields and bias frames on the telescope. This is discussed below in slightly more detail.

Standard Calibration Techniques

Dark Frames

Dark frames are frames which measure the signal level which arises purely from the electronics. Some signals contained within the dark frames are not time dependent (mainly bias signal and read noise) but some are time dependent (thermal signal). Calibration is conducted to remove as much of these signals as possible, leaving “behind” only those portions of the signal which are statistically intrinsic to the uncertainty of the signal (noise).

Flat Fields

Flat fields are exposures taken of a flat, evenly illuminated area. These flats are then used to reduce the pixel by pixel variation due to response differences between individual pixels as well as dust or other light obscuring elements which affect the response of pixels.

Bias Frames

Bias frames are frames which measure the response of a sensor at, or very close to, zero exposure time. They are part of a dark frame and, apart from using them to measure Read Noise, they are not used here since it is suggested that full dark frames be taken at every time and temperature used.

Testing requires the use of calibration and measurement techniques which are typical of what is required in these circumstances:

- Dark frames.

Dark frames should only be taken under the same conditions under which final “science” exposures will be taken. This implies principally that no scaling should be conducted-either with time or with temperature. This is to avoid possible non-linear variations in single pixel response, even if average values taken over many pixels may show linearity. Indeed, these single pixel variations are largely taken into account by measuring values over many pixels (e.g. over a 500x500 matrix).

- Bias frames.

Bias Frames are used only to measure Read Noise. It is not suggested that separate bias frames be subtracted from uncalibrated images. Rather, full dark frames (which also contain the bias signal) should be subtracted as above.

- Flat Frames

Flat Frames are perhaps the most difficult to produce accurately and require a flat, evenly-illuminated light source or the use of the sky just after sunset or just before sunrise. If artificial illumination is used, it is probably best to use continuous-spectrum, incandescent sources (such as tungsten halogen) with some form of current stabilisation and allow these sources to reach operating temperature before beginning to use them. Broad-spectrum LED sources can also be used but, since these still do not produce fully continuous spectra, some testing may be required to determine their suitability.

The use of standard fluorescent tubes or compact fluorescent tubes (CFL) is not recommended.

Setting of Gain Index/eGain

Most CMOS cameras have the ability to apply different amplification to the charge contained in the well, prior to digitalisation. The way this is implemented depends on the approach taken by the manufacturer but, in general terms, it is possible to apply several Gain Indices which will result in differing eGains.

Digitalisation varies depending on the ADC used but is typically 12, 14 or 16 bits.

This implies that ADU values can be 0-4095, 0-16383 or 0-65535.

Gain Indices are largely related to nomenclature used by the camera manufacturer but all relate a specific number (GI), which can be expressed as a percentage or some other dimensionless number, to a specific eGain which is expressed in e-/ADU.

For the purposes of this discussion, an increase in GI indicates an increase in amplification and this implies a reduction of the eG value, when expressed in e-/ADU.

For the CMOS sensor tested, several properties emerge which are relevant to this discussion:

- As amplification increases (eG in e-/ADU decreases), *RN decreases*.
- As amplification increases (eG in e-/ADU decreases), *pixel to pixel variation increases*.
- As amplification increases (eG in e-/ADU decreases), *Full Well decrease*s.

For some cameras, eG (via GI) can be selected from a continuous spectrum whilst for other cameras only discrete eG values can be selected. It is suggested that eG values be selected which cover the spectrum of amplification available for the camera. It is also suggested that not more than 5 values be selected, purely to enable testing within reasonable time frames. Subsequently, should the user desire it and the camera allow it, other eG values can be considered.

Offset/Black Level Settings (OS/BL)

For each of the eG values being investigated, an appropriate OS/BL needs to be determined. This value should be established at typical operating conditions (predominantly temperature). For this work, only bias frames need be used.

The essential principle here is that no bias frame should show zero values since a zero value would affect the linearity of the sensor's response when the bias frame is subtracted. Since a bias frame is taken with a very short exposure, it will contain very little thermal current, thus dark noise can be effectively ignored.

At each of the eG settings, OS/BL needs to be adjusted so that no pixel has a zero value.

Take a bias frame with varying OS/BL values until no pixel shows a zero value within the entire frame. The histogram should be a Gaussian looking curve just to the right of zero. Bear in mind, however, that excessive OS/BL levels will increase the bias current and will take up space in the CMOS well.

Note these values for every eG being evaluated. From then on, all frames should be taken using these values under the same operating parameters.

Testing for Linearity

Not all eG values show the same linearity, that is a linear response with photon absorption. For many astronomical applications, such as photometry, this is an essential requirement and it is strongly suggested that linearity tests be conducted even if the manufacturer supplies data.

For the purposes of this work, it is sufficient to take a number of dark-calibrated flat fields with varying exposure times to cover the lowest ADU levels to the highest.

However, there are several important requirements:

- The light source should be stable; hence some on-time must be allowed to ensure that this is so. Typically, broad spectrum LED sources can be used since the primary aspect being considered is the relative response difference between frames. Ideally, the illuminating light source would be a tungsten halogen lamp which is allowed to stabilise. Even more stable light sources for scientific uses can obviously be employed, if available.
- Since the light intensity needs to remain constant over time, it is essentially not possible to use sky flats taken at dawn or sunset, since the light varies quickly over time.
- At every eG (and OS/BL) setting, dark frames will need to be subtracted. Dark frames should not be scaled either with temperature or time so every exposure time will require its own calibration frames. Use a full dark frame (a dark frame which also contains a bias signal) to avoid possible non linearities. In those instances where the software used allows the subtraction of just a single dark frame, this is probably sufficient. Slightly better results would be obtained by combining several dark frames but this can most likely be avoided at first and may not be necessary.
- A range of exposure times needs to be used which allows at least 15 data points, going from low to high ADU levels. The actual ADU range will vary depending on digitisation (12, 14 or 16 bit). In most instances it will be possible to use the same set of exposure times for all eG values.
- The exposure times can be scaled so that at the lowest amplification factor (highest e-/ADU) the frames go to saturation at the longest exposures. As eG/amplification increases (e-/ADU decreases) the curves will simply reach saturation at lower exposure times. However, it is important that the response be adequately sampled so some trial and error may be necessary.

- To minimise thermal signal, it is suggested that exposure times be kept relatively short, going perhaps from a few tenths of a second to a few seconds. This also has the desirable effect of shortening overall testing time.
- It is best to measure a central portion of the frame to ensure that amplification noise (and other edge effects) are minimised. A 500x500 central pixel area would be sufficient.

At every combination of eG and exposure times the following parameters need to be noted/measured on the dark calibrated frame using a suitable software measuring tool:

- Exposure time.
- ADU maximum, minimum, mean, median and standard deviation.

Note that, as eG/amplification increases, pixel to pixel variation may also increase resulting in higher standard deviations.

For every eG, mean ADU level can then be graphically plotted versus exposure time.

The resulting graph should look something like the graph below.

From the graph, it should be possible to determine the eG levels which show acceptable linearity. By way of example from the plot, which is structured for 16 bit images (0-65535 ADU), it can be seen that the two lowest gain indices taper off well before the maximum ADU values for 16 bit images.

At the same time, two of the eGains (12,69 and 17,36 e-/ADU) show an impressive linearity almost to the saturation limit, while the rest show lower but acceptable linearity.

The eGain values used in the above graph are taken from data supplied by the manufacturer but eGain and RN can also be determined experimentally. This is the subject of the following section.

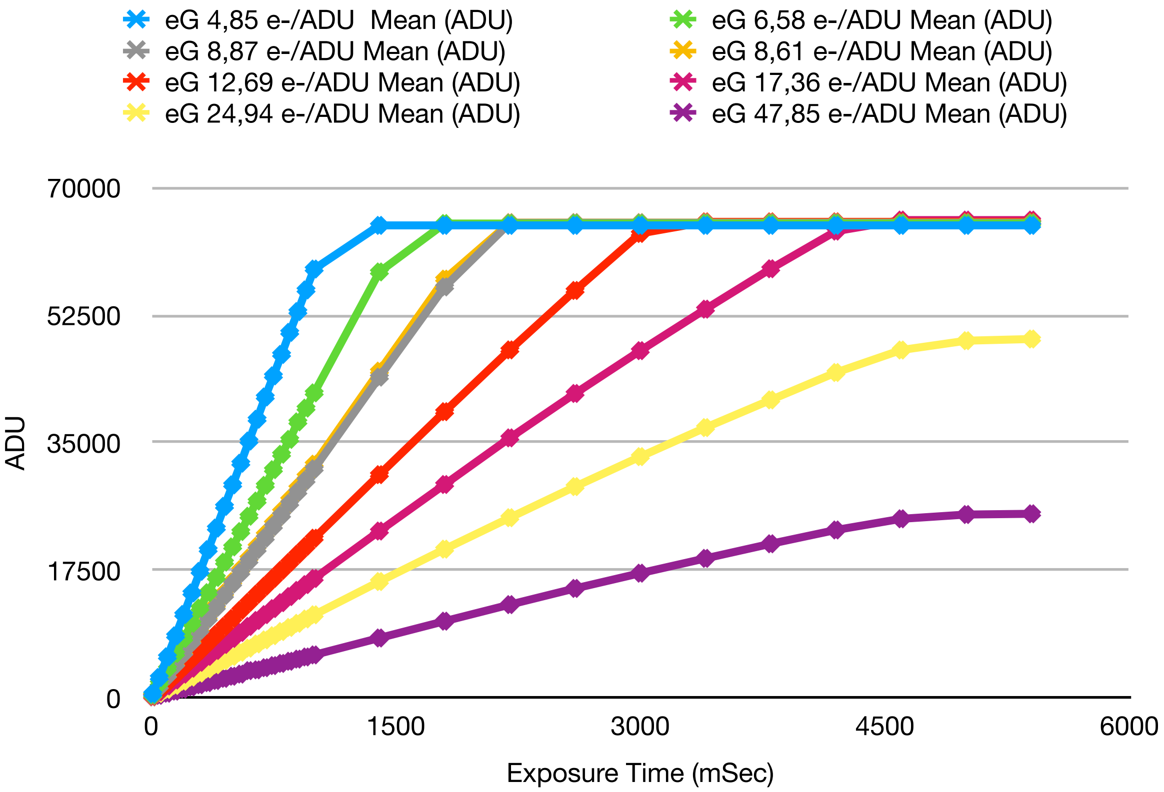

Measurement of Electronic Gain (eGain, eG), Read Noise (RN) and Full Well (FW).

Some manufacturers of CMOS cameras supply data on eG (versus a Gain Index, such as a number or percentage), RN (at a particular GI or eG) and Full Well.

However, it is possible to determine these values experimentally either to confirm the data supplied by the manufacturer or to establish a data set, in the event that these data are not supplied.

The data can then be used to determine optimum operating parameters, in relation to the properties/performance being sought.

Testing

It is suggested that, following the above tests for linearity, one should select an ample range of eGains to cover all the expected operating conditions under which the camera will be used. It is also suggested that eGains be selected which will allow acceptable linearity (as a very rough rule of thumb, so that the sensor is linear out to 85% of its saturation limit but *testing is strongly suggested*). It is possible to use other eGains with lower linearity for specific purposes but, generally, this should only be done when one is aware of these limitations.

Based on the typical properties of commercial CMOS sensor Gain , it is likely that a maximum of 5 eG (e-/ADU) values will be sufficient to cover the majority of requirements.

Testing is based on the Janesick method, who is considered to be the foremost expert on this topic since the 90s. His book, “Photon Transfer”, is available from SPIE. Current price in the USA is $66.

The test procedure is based on documents produced by Arne Henden in relation to the testing of the QHY600 camera. For those interested in doing so, a search on “AAVSO QHY600” should give access to what is essentially an exchange of communications between various AAVSO members. In any event, the link follows: https://www.aavso.org/qhy600-tests

Note that the method of altering exposure in the AAVSO document is different to that suggested for the test in this document. However, for the purposes of the accuracy required here, the method suggested in this document seems to have proven sufficiently accurate.)

Some use has also been made of the latest European Machine Vision Standard 1288 on the characterisation of image sensors and cameras. The website is:
https://www.emva.org

For those wishing to look at the functioning of sensors and cameras in much greater detail, it is suggested that the EMVA publication be studied. Note that there are two publications, one of which is applicable to sensors with a linear response (the case being considered here) and those with a non-linear response.

For the particular eG being investigated, use an exposure time which will allow a flat field exposure to achieve values (in ADU terms) of about half of the maximum value corresponding to the digitisation being used. This is purely a rule of thumb which will apply to most sensors. Other exposures can be used if the user is certain that the exposure is in the linear portion of the response curve. Normally operating at 50% of maximum ADU will be in the linear part of the curve. For example, at 16 bits (0-65535) using a flat exposure time giving about 30000 ADU will normally be in the linear portion of the response curve.

Take 2 flat exposures at this level (flat1 and flat2) and two bias exposures (zero1 and zero2).

The formulae used to calculate the desired parameters are:

1. flatdif=flat1-flat2

2. zerodif=zero1-zero2

3. eG=(((mean(flat1)+mean(flat2))-(mean(zero1)+mean(zero2)))/(sigma(flatdif))^2-(sigma(zerodif))^2)

4. RN= eG *sigma(zerodif)/sqrt(2)

Firstly, flat1 and flat2 are subtracted and the standard deviation (sigma(flatdif)) of this difference frame is measured. Note that if the difference results in zero or negative numbers a constant may need to be added to properly characterise the curve of which the standard deviation is measured. Most measuring software allows this to be an option.

Similarly, zero1 and zero2 are subtracted and the standard deviation (sigma(zerodif)) of this difference frame is measured. Again, note that if the difference results in zero or negative numbers a constant may need to be added to properly characterise the curve of which the standard deviation is measured.

Following is a table of some measured values for the Kepler400 operating at 6 different eGains, along with a test conducted on an SBIG ST-8 CCD:

Camera/Operating Mode	eGain (e-/ADU)	Read Noise (e-)	Effective Full Well (e-)	Dynamic Range (dB)
Kepler400 LG 1,85	**14,94**	41,20	61164	63
Kepler400 LG 2,49	**11,10**	37,52	45451	62
Kepler400 LG 3,70	**7,56**	33,43	30974	59
Kepler400 HG 1,85	**2,34**	4,37	9579	67
Kepler400 HG 2,49	**1,59**	3,51	6507	65
Kepler400 HG 3,70	**1,08**	3,05	4413	63
SBIG ST-8 Bin 1x1*	**2,86**	22,28	100000*	73
SBIG ST-8 Bin 2x2**	**2,69**	26,08	100000**	72

Notes:

* Nominally, the ST-8 is a 16-bit camera but if one applies the FW formula, the calculated FW would be in excess of 180000 e-. The FW specified by SBIG is 100000 e-, suggesting that the true digitisation is probably 15 bits.

** In binning 2x2, eG remains very similar to 1x1 suggesting that the FW remains essentially the same. Indeed, the DR is basically unvaried. In some CCD cameras, the eG would increase meaning an effective increase in FW. This does not appear to be the case for the ST-8.

Effects of eG, RN and FW on Dynamic Range (DR)

To take a practical example from the above table, for the Kepler400 operating at a measured eG of 2,34 e-/ADU (compared to the manufacturer supplied value of 2,57 e-/ADU), the RN was measured to be 4,37 e-. Compare this with CCD sensors which typically have a read noise of 3 or 4 times this value. At a measured eG of 1,08 e-/ADU (manufacturer supplied value 1,10 e-/ADU), RN was measured to be 3,05 e-.

From these numbers, at a given digitisation, it is possible to calculate both the FW and the DR. In the case of the Kepler, 12 bit digitisation is applied (0-4095) so at a gain of 2,34 e-/ADU, the FW is calculated to be:

1. FW=Max ADU*eG
2. FW=4095*2,34
3. FW=9579 e-

As eG decreases so does FW, but the effect on DR is interesting, bearing in mind that DR=20log(FW/RN).

Even though FW decreases, DR can increase as eG decreases. This is because RN decreases. The effect can be significant. Indeed, in this set of data, DR is maximised at an eG of 2,34 e-/ADU despite having a small FW of 9579 e-.

However, the interesting effect comes from stacking (summing) frames, for example at an eG of 2,34 e-ADU with a stacking of 16 frames.

Camera/Operating Mode		Read Noise (e-)	Effective Full Well (e-)	Dynamic Range (dB)
Kepler400 HG 1,85	**2,34**	**17,48***	**153264****	**79**

Notes:

* Assuming a RN of 4,37 e- for a single frame, since RN adds in quadrature, the RN of 16 frames is: 4,37*sqrt(16)=17,48 e-.

** Since the FW of a single frame is 9579 e-, the effective FW (assuming summation) of 16 frames is: 9579*16=153264.

This results in a significantly higher effective DR of 79 dB. So, if DR is an important consideration, stacking exposures is a good way of increasing DR to levels which are comparable with or better than CCD cameras. Indeed, some software allows this to happen essentially transparently by dividing a nominal integration into n number of subframes.

Some of the results of doing this will be discussed in the next section.

Some practical aspects and uses

CMOS sensors have some characteristics which can add extra functionality, compared to CCD sensors. A few are discussed here.

Stacking subframes

Following on from the above theoretical calculation of the effects of stacking on DR, below are 2 examples of the effect in practical terms.

Managing Saturation

Two exposures were taken of a star field (NGC1708), both of the same total duration (32 secs). One exposure was a single exposure of 32 secs whilst the second was divided into 16x2 sec subframes, added on the fly (as the exposure was progressing) by software. The single exposure showed saturation in the star being measured while the multiple subframe exposure shows the same star without saturation. Nominally, integration time remains the same.

Since saturation introduces a non-linearity, clearly, for photometry, it is important that both target and comparison stars not be saturated and so this represents one possible way of utilising exposures long enough to average out scintillation effects whilst maintaining the target and comparison stars unsaturated.

It should be borne in mind that the single exposure invariably will show dimmer stars but the useful DR (measured as a ratio between the dimmest star in the field with an acceptable SNR (say, greater than 10) and the brightest unsaturated star in the field) will be higher in the sub-exposure frame.

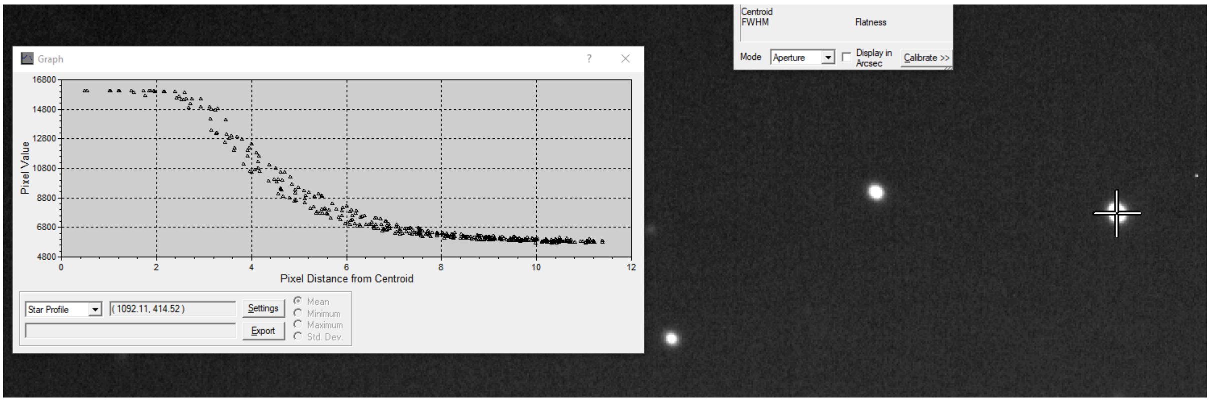
Star Profile
(1092.11, 414.52)
Mode: Aperture

32 sec single exposure, saturated star

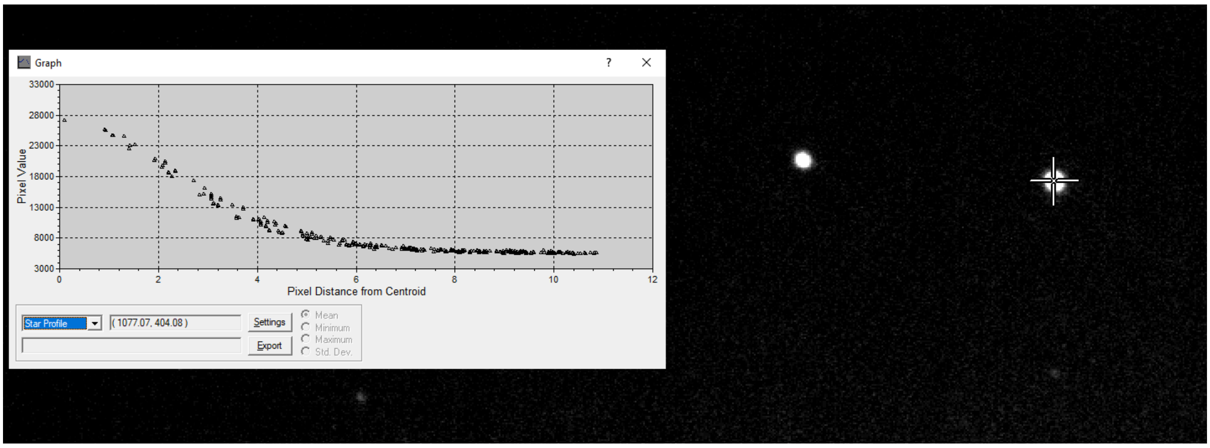
Star Profile
(1077.07, 404.08)

16x2 sec exposures stacked-on-the-fly, star not saturated

Of course, saturation can also be managed by reducing exposure time and defocusing. Perhaps more importantly, the effect on useful DR is discussed more fully in the next section.

Improving Useful DR

The same star field discussed above was analysed in the two different exposure modes as outlined above-a single 32 sec exposure and a stacked-on-the-fly image consisting of 16x2 sec exposures.

The dimmest star in the 2 fields with an SNR a little over 10 was measured in terms of ADU intensity (sum of all pixels in the measuring aperture) and the brightest, unsaturated star was also measured in the same way.

Note that the stars measured were not the same in the two fields since selection was made based on SNR for the dimmest stars and highest brightness just short of saturation for the brightest stars. The measuring aperture used was 2xFWHM.

It should be borne in mind that these numbers are purely representative of the specific camera and testing conditions used. Other cameras and different testing conditions (e.g. the use of filters, binning, differing sky background and star elevation) may well show different results.

Exposure length (sec)	Intensity (ADU)	SNR	Dynamic range (ratio)	Dynamic range (dB)
32	**3495**	**13,8**		
32	**127441**	**487**	**36,5**	**31**
16x2	**6118**	**13,6**		
16x2	**2639905**	**3492**	**431**	**52,7**

Nonetheless, it is felt that these results will be broadly typical of CMOS sensors which can use stack-on-the-fly methods, which have low read noise and fast file read-out. Obviously, if sufficient storage capacity exists, stacking post capture of very short exposures into a suitable sub-frame, will also work

The same approach can be used to establish the magnitude range which can be captured in a single exposure and which maintain linearity. If done as preparation for a photometric session, it will ensure that target and comparison stars remain in linear parts of the response curve. A magnitude range can be established for the various operating modes (eGain, black level) the user would normally employ.

By varying the operating mode, exposure time and the number (if any) of sub-exposures, the user can accomodate a number of situations, ranging from very bright stars to dim.

Lucky Imaging and the Effect on FWHM

1366 images of NGC 1708 were taken at an eG 0,54 e-/ADU and a RN of 2,54 e- and at an exposure time of 0,05 sec. These images were then classified by average FWHM over the field.

The fast read-out allows taking a large number of images in a relatively short time. However, this depends on the type of camera and operating software and the ability of the system to capture images (capturing a stream of images depends on many factors such as processor speed and type, type of connection used (e.g. USB2 or USB3), type of media used (e.g. SSD or normal hard disk), frame rate and frame size). The Kepler 400 allows capturing full frames at about 24 FPS in one of the operating modes. In this case the effective capture rate was about 2 FPS, due to some software issues but the principles remain the same.

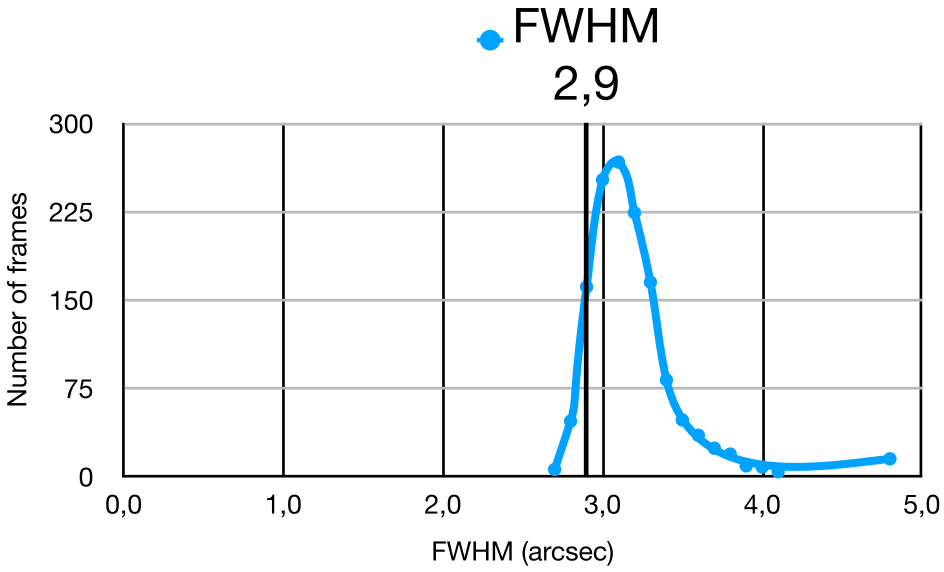

Nearly 16% of the frames showed average FWHM values better than or equal to 2,9 arcsec. The average value over the rest was about 3,4 arcsec so an effective improvement of 0,5 arcsec was seen for this limited number of frames.

Several points should be made:

- The more significant effects of Lucky Imaging are normally seen over fields of only a few arc minutes. The field of these images was about 26 arc minutes square and an average over this field was taken. Over smaller fields it would be reasonable to expect better results. The test conducted here was purely to look quickly at a broad effect and further work should be conducted.
- Using only 16% of the exposure series obviously means that a large number of exposures are not used. This implies a high overhead. Hence, this approach will be of use where such a high overhead can be accepted.
- Clearly, the very short exposure times used imply that there will be a limit to how faint potential targets can be. In this particular case, the faintest detectable stars (SNR greater than 10) were at about GMag 15,5. The above exposures were taken with a 37 cm RC telescope.

Conclusions

This document is not exhaustive but hopefully it will provide enough stimulus for a user to conduct her or his own testing and experiment with the characteristics of the camera being used. However, it is hoped that it will facilitate a degree of uniformity within the Exoclock group, in particular amongst CMOS users.

Do look at the listing of cameras being used by the ExoClock group and do add your own camera if it's not present: https://docs.google.com/spreadsheets/d/1-J0yXYvPMAilVUu5v4hXsPFr4HIZPuZe8kmo3kxh-S0s/edit?usp=sharing

This should be seen as a live document so comments, critiques and suggestions for further work are welcome.

To facilitate receiving comments please use either the CMOS-WG Slack channel or write to the following:

Roland Casali	mrcas62@gmail.com;
David Rees	walnut1@easynet.co.uk;
Leon Bewersdorf	bewersdorff@pm.me