MRO Fast Tip-Tilt

Andor iXon X3 897 Custom Clocking Evaluation II

MRO-TRE-CAM-1120-0162

Bodie Seneta
bodie@mrao.cam.ac.uk
Chris Haniff
cah@mrao.cam.ac.uk

rev 1.1
19 April 2013

Cavendish Laboratory
Madingley Road
Cambridge CB3 0HE
UK

Change Record

Revision	Date	Authors	Changes
0.1	2013-03-23	EBS	First draft.
0.2	2013-03-28	EBS	Added many diagrams and corrections. Changed centroid analysis. Changed temperature performance test.
0.3	2013-04-04	EBS	Qualitative image fidelity section reworked after a discrepancy in the data was discovered. Temperature performance data and tables reworked. Many minor corrections.
1.0	2013-04-05	EBS, CAH	Many minor edits. First major release.
1.1	2013-04-19	EBS	Added document number, fixed version typo.

Objective

To compare the performance of two readout modes, a standard mode and a recently developed fast readout mode, for an Andor iXon X3 897 camera. The camera is to be used in the Magdalena Ridge Observatory Interferometer fast tip-tilt (FTT) system and so the tests focus on dark current and centroiding performance.

Reference Documents

RD1 MRO-TRE-CAM-0000-0101: Derived Requirements – rev 1.0, August 31st 2010
RD2 MRO-TRE-CAM-1120-0118: Andor iXon X3 897 Custom Clocking Evaluation – rev 1.0, 15 July 2011

Scope

This document forms part of the documentation for the fast tip-tilt system to be developed for and installed at the Magdalena Ridge Observatory Interferometer. It describes tests undertaken to evaluate the performance of a customised CCD clocking scheme devised by Andor for their iXon X3 897 camera, under conditions similar to those to be encountered in operation. This document supersedes RD2, which documents similar tests with an earlier custom clocking mode.

It is assumed that the reader is familiar with the astronomical concept of stellar magnitudes, as well as readout methods for charge coupled devices (CCDs) and electron multiplying CCDs (EMCCDs) in particular.

Contents

1 Introduction **3**

2 Conventional and Custom Clocking Schemes **4**

2.1 Conventional Clocking 5

2.2 Custom Clocking 5

3 Experimental Layout **7**

4 Assessing the Impact of Noise on Centroiding Performance **9**

5 Description of Tests **9**

6 Test Results **10**

6.1 Readout latency 10

6.2 Qualitative Image Fidelity 11

6.3 Dark Frame Quality 13

6.3.1 Effects of clocking scheme and parallel clock voltage . . . 13

6.3.2 Effect of temperature on custom clocking 19

6.3.3 Effects of integration time on custom clocking 20

6.4 Illumination Tests 21

6.4.1 Effect of using a faint strobed source during custom clocking 21

6.4.2 Effects of clocking scheme on centroiding accuracy 23

6.5 EMCCD Temperature Stability 27

7 Summary **30**

1 Introduction

The Fast Tip-Tilt (FTT) system for the Magdalena Ridge Observatory Interferometer is required to image a celestial source, calculate the position of this image, and move a compensating mirror — with minimal latency — to correct for the tilt component of the wavefront perturbations introduced by the atmosphere. Hence the image sensor component will form the detection element of a servo system. Two types of readout are required, a full frame for acquisition, and a sub-frame for the fast tip-tilt correction. Some relevant requirements for the imaging sensor (RD1) include:

- A minimum size of the rapidly read out sub-frame of 3.6×3.6 arc-seconds on the sky;
- The ability to centre this sub-frame anywhere within 10 arc-seconds of the center of a total field of view of 60×60 arc-seconds.
- The need for a closed-loop bandwidth of 40 Hz, with a goal of 50 Hz. This bandwidth requirement need not be met for the faintest specified targets (16th magnitude). For these faintest targets RD1 suggests a closed loop bandwidth of around 15 Hz. These bandwidths correspond to maximum exposure lags of 1.74, 1.39 and 4.63 ms respectively (Subsection 6.1).
- A two-axis detection noise centroiding error of less than 34 milliarcsec.

After RD1 was written, the optical design of the FTT system was frozen and a camera selected. This defined the pixel scale as 0.147 arcseconds per pixel and allowed the above requirements to be restated as:

- Minimum sub-frame size 25×25 pixels.
- Minimum detector dimension 408×408 pixels.
- A two-axis detection noise centroiding jitter of less than 0.23 pixels.

It is clear from vendor literature that the iXon X3 897 camera meets the sub-frame size requirement in principle, but the determination of its ability to meet the other requirements requires further testing. The conventional "region of interest" (ROI) readout available with the camera allows for a centrally placed sub-frame of adequate size to be read out at roughly 400Hz, but the readout latency is too long to allow our closed-loop bandwidth requirement to be met.

In 2011, Andor delivered a customised readout mode for us to evaluate, in which a 23×23 pixel region could be read out rapidly. This scheme minimised the number of serial shifts needed to get the data from the ROI out of the chip at the expense of some contamination in the images from light outside the ROI. Testing of this custom mode is documented in RD2. The mode was found to perform well, and in fact had lower latency than was required. However, there were

concerns about excessive spurious signal, due to charge in many pixels outside the region of interest getting summed into those pixels inside it, and a larger subframe size was also preferred. These concerns led us to wonder if it might be possible to deliver a new readout mode that was slightly slower, but accessed an ROI that was larger and had less contamination from spurious signals.

Late in 2012, Andor delivered a new readout mode to us which aimed to do this. This mode reads out a 32×32 pixel region in two steps, each of which is similar to the method used to read out the original 23×23 region. It was necessary to evaluate this mode in much the same way as the original mode, although some shortcuts were possible based on previous experience.

This report summarises the tests we have carried out to determine whether this custom clocking mode meets the requirements of the FTT system. Where possible we have repeated the tests using the conventional ROI readout so as to act as a comparison, and to assess whether this mode could act as a suitable fall-back for situations where a lower closed-loop bandwidth (i.e. 15 Hz) were acceptable.

Three kinds of tests were carried out:

- Latency measurements, to establish whether the mode allowed our servo bandwidth requirements to be met;
- Dark frame tests, in which no external light was allowed to reach the camera, thereby allowing assessment of the array's noise performance;
- Illuminated tests, in which an image of a point source was focused on the CCD. The source was either steady, or strobed in synchronisation with the camera frames to simulate a rapidly moving image. These tests allowed assessment of the centroiding accuracy as a function of signal level.

For all the tests a range of CCD temperature and clocking voltage parameters were explored.

Our overall conclusion is that both the custom and the conventional ROI clocking schemes produce centroid errors of similar magnitude for a 15 Hz servo at the signal levels tested. The standard ROI readout does not meet our latency requirement but the custom readout mode meets both the requirement and the goal.

The reader should note that the results in this document are not an official endorsement or detraction of Andor's products, and no guarantee is made that our interpretation of Andor's custom clocking scheme is correct.

2 Conventional and Custom Clocking Schemes

A clocking scheme is the sequence of voltages applied to a CCD in order to move the charge in the pixels to a pre-amplifier, where it can be measured and digitised. The Andor X3 897 camera uses an E2V CCD97 chip [1], which allows

considerable flexibility in how this is done. In particular, it is not necessary to clock out every pixel in the CCD if one is only interested in a subregion of the array.

For the purposes of this memorandum, two clocking schemes were investigated: conventional ROI clocking and a faster custom clocking scheme developed by Andor to meet our high frame readout requirement. The following two subsections present our current understanding of how each of these clocking strategies were implemented. Should these summary descriptions be inaccurate, we would welcome clarification.

2.1 Conventional Clocking

In conventional ROI mode our understanding is that the charge in the CCD image section is firstly parallel-shifted into the CCD storage section, which can be independently clocked out while the image section begins exposure of a new image. The charge in the region of interest is then parallel-shifted towards the readout register, with the charge from all the pixel rows prior to the first row containing the ROI being shifted into the readout register. Next, the readout register is cleared by reading it out. The CCD rows that contain the region of interest are then, one at a time, clocked into the serial readout register, read out and digitised. Although all the row values are read out, only those pixels within the ROI are subsequently made available to the user.

The following aspects of this "standard" readout mode are worth noting:

1. Because every pixel in any row containing the region of interest is read out, there is no time penalty associated with increasing the width of the ROI.

2. Conversely, this scheme is inefficient if the width of the region of interest is small compared to the full CCD width, because a large fraction of time will be spent clocking out and reading pixels outside the ROI.

3. Within the readout register, charge generated by the exposure in any pixel never gets added to the charge in any other pixel, so there is never ambiguity as to where on the CCD image section the charge that is eventually digitised came from.

In our lab tests with the Andor X3 897, we found that that a 32×32 sub-frame centred on the 512×512 pixel light-sensitive region of the chip could be read out at around 414 Hz.

2.2 Custom Clocking

The customised clocking scheme tested was based upon a proposal to Andor by our team, further refined through face-to-face and on-line discussions between

the two groups, and testing of the earlier 23×23 custom mode as discussed in RD2.

This is a scheme for rapidly reading out a 32×32 subframe. It is our understanding that in this scheme, and as before, the exposed CCD pixels are all first shifted to the storage region. Subsequently, all the rows prior to the first containing the ROI are shifted into the readout register. Next, the readout register is cleared by reading it out. Then, the first row containing data of interest is clocked into the readout register. However, instead of then reading out this whole row, the readout register is only clocked along by the width of the ROI, that is, by 32 pixels. The next row is then clocked into the readout register, which places the charge from the pixels of interest in the most recent row of the ROI adjacent to those pixels from the previous ROI row. This process is repeated 16 times so that charge from all the pixels in the lower half of the ROI are stored entirely within the readout register. The data in the serial readout register is then clocked out and digitised so as to obtain the pixel values in a sequential form.

The upper half of the readout region is then read out in similar fashion: each of the remaining 16 rows is clocked into the serial readout register, which is then shifted along by 32 pixels. After all 16 rows have been shifted into the serial register, it is read out once more to obtain the pixel values for the upper half of the 32×32 subframe.

Once again, there are a number of aspects of this custom readout mode worth noting:

1. This clocking scheme minimises the number of serial clock shifts required to get the data out of the ROI, greatly decreasing the acquisition latency. For regions of interest whose width is a small fraction of the full CCD width, the saving in time is considerable.

2. The readout rate is slower than it is for the previously tested 23×23 custom scheme (in that scheme, only one readout of the serial register was necessary, whereas the 32×32 scheme requires two). However the subframe size is larger, so a new compromise between size and readout rate has been found.

3. The primary penalty of this strategy is that as each new row is shifted into the output register, charge from pixels outside the ROI gets summed with charges from earlier rows of the ROI that have previously been dumped into the serial register. The charge from such pixels is expected to be small due to the high probability that the rest of the chip contains no other sources, but any target-independent charge, e.g. arising from thermally induced fluctuations or clock-induced charge effects, will contribute to this "spurious background"

4. The size of the region of interest is limited by the length of the readout register. We believe that in the case of the iXon X3 this is what limits the size of the custom region of interest being fixed at 32×32 pixels if a maximum of two readouts of the serial register are allowed.

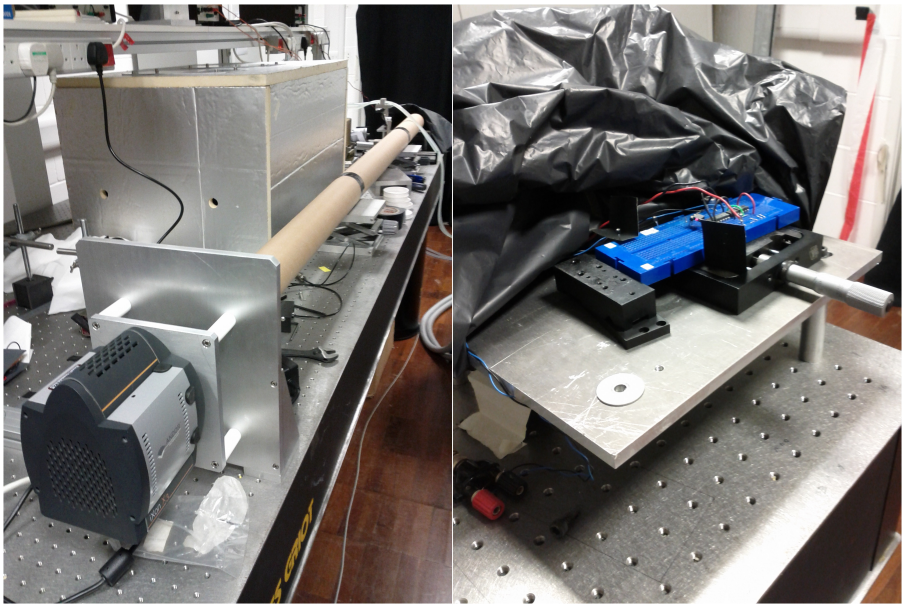

Figure 1: Left: The Andor camera in its custom mount, facing the light source at the far end of the optical table. A cardboard tube acts as a light shield between the two. Right: The LED light source (mounted on the blue breadboard). Neutral density attenuators and light shields have been removed.

Unlike the 23×23 region previously tested, no restrictions have been placed on the position of the 32×32 region on the chip. When it is centrally located, our lab tests have found that it can be read out at 1140Hz, some 2.8 times faster than a conventional readout of the same region.

3 Experimental Layout

The experimental apparatus used for our tests consisted of an iXon X3 897 camera at one end of an optical table and a light source at the other. Cardboard tubing between the two was used to shield the experiment from extraneous sources of light, and the test room was also blacked out. Apart from the cardboard tubing and the mount for the light source, the experimental hardware and arrangement was the same as for the tests in RD2.

The camera (Figure 1, left) was attached by its front plate to a custom mount similar to the mount proposed for the fast tip-tilt system. The imaging optic was a commercial telephoto lens, which was adjusted to produce an image of the light source with a Gaussian profile and FWHM of 3.75 pixels. This was broadly similar to what would be expected of a real star when focused onto the FTT sensor at one of the MROI unit telescopes. The lens was stopped down to f/32 to reduce the input light level to one comparable to those likely to be encountered in practice.

The light source (Figure 1, right) used was an Avago Technologies HLMP-6000-

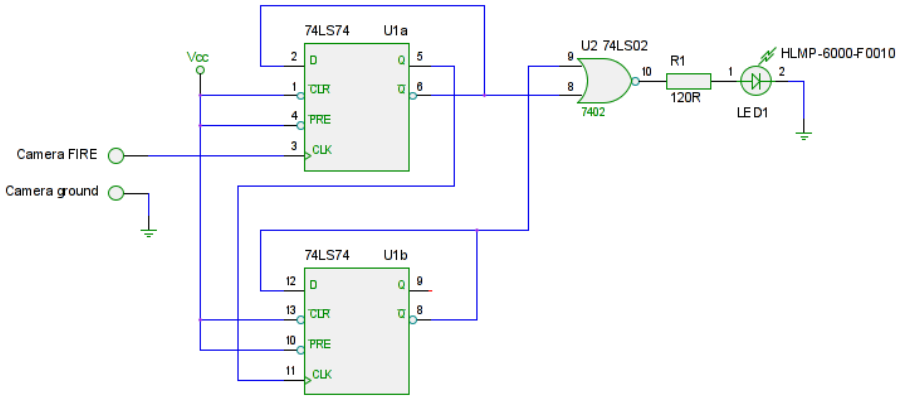

Figure 2: Circuit schematic for the strobed light source.

F0010 light emitting diode (LED), chosen for its rapid 15 ns response time, peak wavelength of 640 nm (at which the CCD has a quantum efficiency of 95%) and circular, diffused face. It could be completely turned on or off between one camera exposure and the next. The LED was wired with additional circuitry onto a breadboard that could be moved horizontally via a micrometer gauge. This allowed the image of the source on the CCD to be moved by calibrated sub-pixel distances for testing the effects of the different readout modes on the centroid error.

The circuitry driving the LED (see figure 2) was simple. The camera FIRE line signalled the beginning of an exposure [3] and this signal was input to a pair of D flip-flops wired as a two bit counter. The counter lines in turn fed a NOR gate, which only went high when the inputs were both low — that is, for one cycle in every four of the counter. The NOR gate provided sufficient power to light the LED via a current limiting resistor. This provided more than sufficient illumination for the camera: in fact a 4.0 neutral density filter was necessary to reduce the illumination to the levels we would expect from the faintest observable objects expected in astronomical use.

The camera was driven by a host computer running Windows 7 and Andor's Solis software to acquire the data reported on in this memorandum. For each acquisition, 1000 consecutive sub-frames were acquired and saved on the local hard disc as a FITS image stack. Two frame rates were tested: 300 Hz and 1000 Hz. The first of these is the rate suggested in RD1 as that likely to be used for the observation of the faintest sources, whereas the second was used to establish how well the hardware performed at the highest frame rates. Analysis of the image data was carried out with the FITS analysis program ImageJ, as well as in-house coded Python scripts.

4 Assessing the Impact of Noise on Centroiding Performance

In order to assess the impact of the two clocking schemes on the image quality delivered by the Andor camera a fixed light source was used to illuminate the CCD so as to mimic a stellar image with a FWHM of 3.75 pixels. The RMS quad-cell centroid "jitter" as measured from 1000 consecutive image frames was then used as a metric to assess the level of noise in the image data.

It is well known (see RD1 and references therein) that the RMS one-axis centroid jitter derived from a sequence of noisy images can be written as [1]:

$$\sigma_{jitter} = \frac{3\pi}{16} \frac{1}{SNR} \,\mathrm{FWHM}\,, \tag{1}$$

where SNR is the signal to noise ratio and FWHM is the full-width at half maximum intensity of the assumed diffraction-limited image. The reader should note that since the test image was not a diffraction-limited Airy pattern but closer to a Gaussian, this formula will be inaccurate at some level.

The SNR, allowing for electron multiplying gain noise, is given by:

$$SNR = \frac{N}{\sqrt{2N + n\sigma_r^2}} \tag{2}$$

where N is the number of photons detected per image frame, n is the number of pixel reads, and σ_r is the effective readout noise (referred to the electron multiplying input). The factor of 2 under the square root represents the noise introduced by the stochastic nature of the electron multiplying process.

For the test camera with a pre-amplifier gain of $5.2\times$ and a 100 ns serial shift period, the readout noise is specified as ~ 52 electrons. The faintest observing target expected for the MROI FTT system is a 16th magnitude star, which corresponds to a detected flux at the FTT sensor of roughly 11400 photons per second. Finally, the centroid method used for the tests employed a 10×10 pixel quad cell, and so the appropriate value for n was 100.

These equations and data allowed us to predict the centroid jitter expected due to readout noise, stochastic electron multiplication noise, and the Poisson fluctuations in the incident light signal. The extent to which the measured centroid jitter exceeded these predictions was used to assess whether or not additional sources of noise were present in the image data.

5 Description of Tests

A range of tests were conducted to compare the performance of the conventional and custom clocking schemes, in order to determine whether the custom scheme was suitable for the fast tip-tilt project. This set of tests comprised:

[1] The two-axis jitter will be $\sqrt{2}$ greater than this.

1. Latency measurements, to establish whether the servo bandwidth requirements could be met;
2. Qualitative image fidelity tests imaging a steady source, as a function of clocking scheme, parallel clock voltage and parallel clock rate;
3. Measurements of dark frame quality as a function of clocking scheme, parallel clock voltage and parallel clock rate;
4. Measurements of dark frame quality in the custom clocking scheme as a function of temperature;
5. Measurements of dark frame quality in the custom clocking scheme as a function of integration time;
6. Checks for residual charge in the custom clocking scheme using a faint strobed source.
7. Checks of the centroiding accuracy for faint sources, with both conventional and custom clocking.
8. Checks of the ability of the camera's Peltier cooler to maintain a set temperature during customised readout under conditions of air and water cooling.

All the tests took place between the 24th of January and the 3rd of April, 2013.

6 Test Results

6.1 Readout latency

RD1 states that the total exposure lag (half the exposure time, plus the readout time, the centroid computation time and the mirror actuation time) should not introduce a servo phase lag of more than 25°. For servo bandwidths of 50Hz, 40Hz and 15Hz (Section 1) the corresponding allowed time lags are 1.39 ms, 1.74 ms and 4.63 ms.

The centroid computation time and mirror actuation time were not measured as part of these tests, but were conservatively estimated to take 0.1 ms in total. This, together with the measured readout latency, was used to estimate the maximum exposure time and hence whether the servo bandwidth requirements could be met.

In this first test, the camera FIRE signal and the DMA transfer interrupt on the camera PCI card were connected to an oscilloscope for monitoring the readout latency. The camera FIRE signal goes from high to low at the end of an exposure, while the interrupt line goes low when the data has been transferred to computer memory and is available for processing. Hence the time between these two events can be used to measure the readout latency.

As mentioned earlier, the test sub-frame was a 32×32 region centred on the light-sensitive area of the chip. Readout times for both the conventional and custom clocking schemes were measured. The camera parallel and serial clock periods for this test were set to 500 ns and 100 ns respectively. For this sub-frame size and location, the conventional and custom readout times were measured to be 2.42 ms and 0.88 ms respectively.

We can see that the conventional ROI readout scheme will not allow closed loop bandwidths of 50 Hz or 40 Hz to be met, while the custom clocking scheme does. These tests confirm that the trial custom clocking scheme can deliver the servo bandwidth that we require, but that conventional ROI readout can only meet the needs of the lowest 15 Hz closed-loop bandwidth that is compatible with faint source sensing.

6.2 Qualitative Image Fidelity

The qualitative fidelity of an image of a faint, steady source was investigated as a function of the parallel clocking voltage for both the conventional and customised clocking schemes. In Andor's programming interface, the parallel clock voltage can be set to values of "Normal" ("0") though to "4", with "4" being the highest. A higher value improves the ability of the camera to shift charge from one row to the next, but it also causes the generation of additional clock-induced charge. Parallel clock voltages of "3" and "4" were not tested here, as experience gained in RD2 had shown these settings would produce noisier data than the others.

In our tests, the following parameters were held constant:

- CCD temperature = $-60°$C;
- EMCCD gain = 250 & Pre-amp gain = $5.2\times$;
- Exposure time: 2.15 ms;
- Serial clock period = 100 ns (the minimum available);
- Sub-frame size 32×32 pixels and sub-frame corner position (240, 240).

The pre-amp gain setting was chosen to minimise the readout noise, which should have been ~ 52 electrons per pixel for these settings. The EMCCD gain was that chosen in RD1 for observations of the faintest required targets.

To ensure that comparisons of image brightness were valid between modes, all the acquisition types were made to have the same exposure time, which was the exposure time of the slowest mode tested (conventional readout, 500ns parallel clock).

The light from the LED source was attenuated with a 4.0 neutral density filter, so that the signals and noise would be comparable in the resulting images. The parallel clock periods were restricted to 300 ns and 500 ns, the shortest available,

Clock Sequence	Clock Voltage		
	0	1	2
Conventional, 300 ns	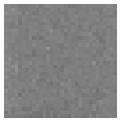	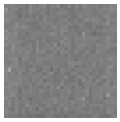	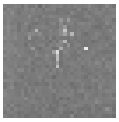
Custom, 300 ns	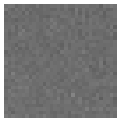	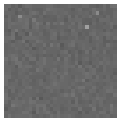	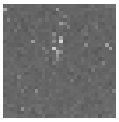
Conventional, 500 ns	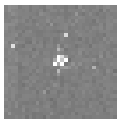	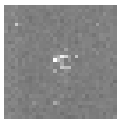	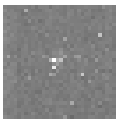
Custom, 500 ns	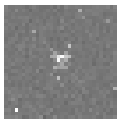	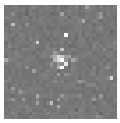	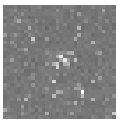

Table 1: Typical images as a function of parallel clock voltage, clocking scheme, and parallel clock period. Light from the LED has been deliberately attenuated so that its signal is comparable in strength to spurious signal sources. An identical grey scale has been applied to all images so as to allow straightforward comparison. Note that for the shortest clocking time of 300 ns, a clock voltage of at least 2 was necessary to transfer charge out of the CCD and even then the charge was smeared in the parallel clocking direction.

as our latency requirement implies the readout must be as fast as possible, but results were secured for parallel clock voltage settings of ”0” to ”2” to see which gave better image quality. Stacks of 1000 images were acquired for each setting, and the first image in each stack is displayed in Table 1.

Results for this test were qualitatively the same as those in RD2.

The most striking similarity was that no images were seen for a parallel clock period of 300 ns and a parallel clock voltage of 0 or 1, regardless of whether conventional or custom clocking was used. It appears that parallel clock voltages of less than “2” were not sufficient to shift charge from the exposure region into the serial register at this clock rate. A clock voltage of “2” produced images for both readout schemes, but in both cases the images were smeared in the parallel clocking direction.

In contrast, when a 500 ns parallel clock period was used, an image was seen for every clock voltage setting. For both conventional and custom clocking, there was more speckle in the background with increasing clock voltage, and increased speckle was also apparent in the custom case.

These results were useful for constraining the parameter space requiring further investigation. 300 ns clocking is not suitable for use with any of the clock voltages tested. It is also apparent that of the remaining options, the background speckle (which we interpret as noise) can be minimised by using the lowest clock

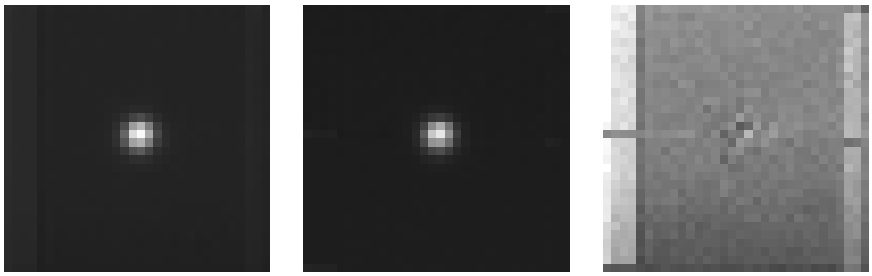

Figure 3: Comparison of averaged images acquired with conventional and custom clocking. Left, average of the stack of 1000 frames acquired with conventional clocking and a parallel clock with period 500ns and "normal" voltage. Middle, with custom clocking and the same settings. Right, the result of subtracting the middle image from the left image. The grey scale for the left and middle images is 50 to 170 counts. For the right image it is -10 to +10 counts.

voltage available.

Finally, to assess if there were significant differences between images acquired in conventional and custom modes, the two image stacks acquired with the "normal" 500ns parallel clock setting were each averaged, then the custom average was subtracted from the conventional average.

The result is shown in Figure 3. Some artifacts are visible in both average images (to be discussed further in Subsubsection 6.3.1). However, the main result is that differences between the LED profiles in both images are only of the order of 2–3 counts, less than the signal from the artefacts.

A custom readout region with a 500 ns parallel clock still exceeds the latency requirement and permits the use of lower clock voltages (and consequently less noisy images). Therefore, we do not investigate 300 ns parallel clocking any further.

6.3 Dark Frame Quality

6.3.1 Effects of clocking scheme and parallel clock voltage

In these tests, dark frame performance was explored. The following parameters were held constant:

- CCD temperature = $-80°$C;
- EMCCD gain = 250 & Pre-amp gain = $5.2\times$;
- Exposure time: 2.15 ms;
- Serial clock period = 100 ns;

Clock Sequence	Clock Voltage		
	0	1	2
Conventional, 500 ns	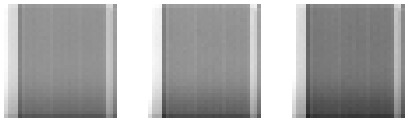		
Custom, 500 ns	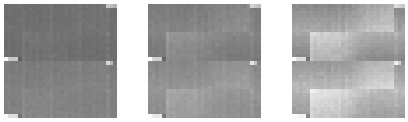		

Table 2: Dark-frame averages of 10000 frames as a function of parallel clock voltage and clocking scheme. The conventionally clocked and custom clocked images have grey scale mappings of 95 to 105 counts and 90 to 100 counts respectively. These mappings were selected based on the distribution of counts for the two modes (discussed later in this section).

- Parallel clock period = 500 ns;
- Sub-frame size 32×32 pixels and sub-frame corner position (240, 240).

A CCD temperature of -80°C was used because data for this test was taken prior to the discovery that the temperature could be set to -60°C without compromising performance (Subsubsection 6.3.2).

In each case an image stack of 10000 frames was acquired. Each stack was processed to give an average value per pixel. This emphasised any systematic artifacts introduced by either clocking scheme. Our results are shown in Table 2.

Significant systematic artifacts were visible for both the conventional and custom clocking schemes. With conventional clocking, there were vertical bands evident at the left and right of each average image. In the custom clocking scheme, the pattern was quite different, consisting of horizontal streaks a few pixels long at the top, middle and bottom of the image, and some blocks that fade into the background on alternate sides. The latter pattern became more dominant with increasing clocking voltage.

Some artifacts appear to be common to both conventional and custom clocked images (Figure 4). When a conventional subframe is clocked out, the serial shift register only contains one horizontal row from the image at a time, which is subsequently read out. In contrast, in the custom scheme an entire half of the subframe is stored in the serial shift register at a time prior to readout. As these artifacts occur at those points in the image data where digitisation of the shift register contents begins and ends, it is possible that something about the process of initiating or completing digitisation of pixels is responsible for these artifacts in both clocking schemes.

The origin of artifacts seen only when the custom clocking mode was used is less clear. We investigated one possible source by way of a numerical simulation, where the detector array was modelled in software. Patterns of charge

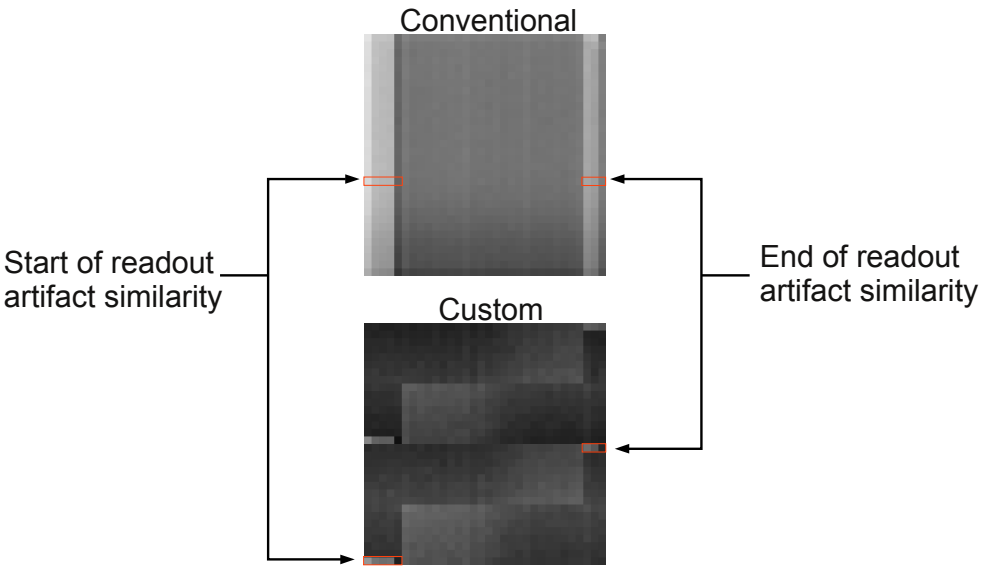

Figure 4: An example of artefacts in conventional and custom mode that appear at the start and end of a serial register read. They appear to have a similar pattern in both modes. In this sample we use the averaged images for a parallel clocking voltage of "2" from Table 2.

are placed on the array and it is clocked out according to the custom clocking scheme to derive 32×32 synthetic images. During this procedure, as noted in Subsection 2.2, charges from other nearby pixels get "added into" the pixels of interest within the serial register. Our simulation calculated the final charge values for each pixel in the output image by the time it would have been read out. A schematic of the simulation and the initial charge patterns appears in Figure 5. The resulting images, and a real averaged image for comparison, are in Figure 6.

Two kinds of initial charge pattern were tested. Firstly, every element in the array was given a charge of "1", and this was called "Model 1". This is consistent with a uniform noise source being present in all pixels, even those rendered insensitive to light by the array mask. The result from this run is shown in the middle image of Figure 6. It can be seen that the simulation correctly predicts the positions of the bright blocks within the experimental data, but does not predict how they fade on alternate sides. So the result suggests that there is some "non-target-related" contribution from every pixel in the CCD in our dark frames, but that this contribution varies according to position on the detector. Secondly, "Model 2" introduced a more complex initial charge pattern in which the charge in every pixel is "1" except close to the left and right edges of the array, where it increases towards the edges. The result of this run is shown in the right image of Figure 6 and can be seen to be a much closer match to the real image than Model 1. It should be stressed that Model 2 is purely empirical — the cause of the patterns seen in practice might have some completely different cause.

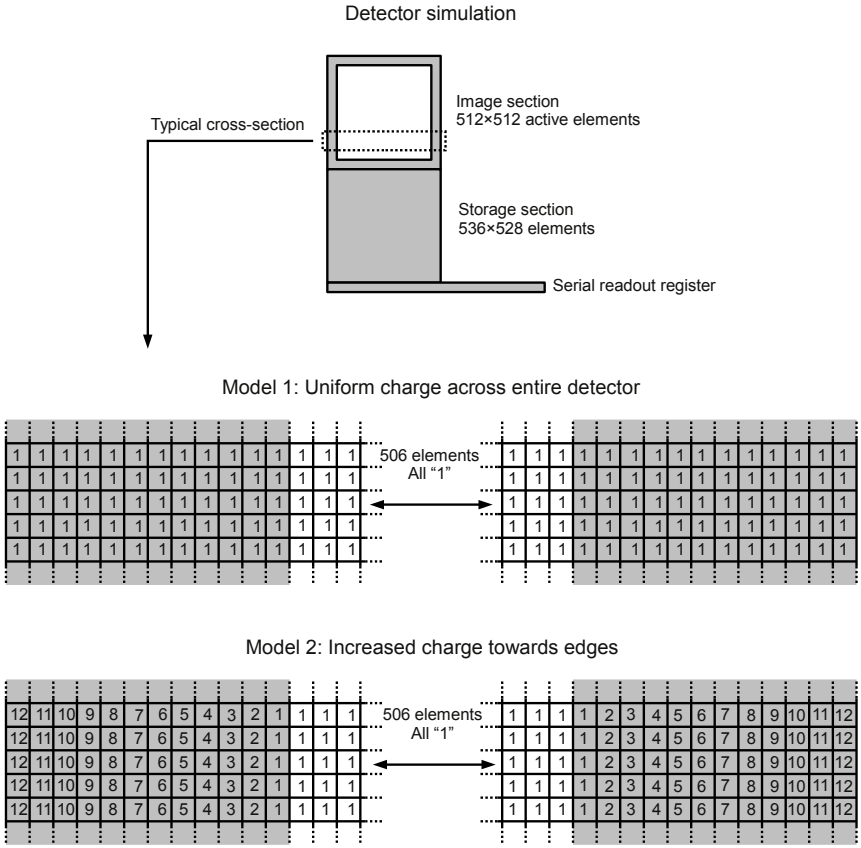

Figure 5: A numerical simulation for studying how initial patterns of charge on the EMCCD array appear when they are clocked out according to the custom clocking scheme. Top, a simplified schematic of the simulated detector showing the major regions. Middle, "Model 1": an initial condition where every pixel has a charge of "1". Bottom, "Model 2": an initial condition where every pixel has a charge of "1", except for an increase close to the left and right edges. Regions coloured grey are masked and are not sensitive to light.

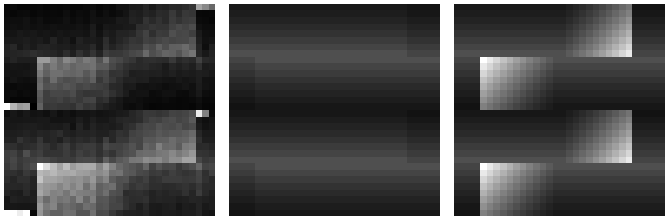

Figure 6: A comparison of our custom clocking simulation with real dark frame data. Left, a real 10000 frame dark average at "+2" clocking voltage. Middle, simulation showing the charge that gets clocked into each image pixel, assuming identical initial charge in every pixel in the detector prior to clocking out ("Model 1"). Right, simulation as before, except that the initial charge in the left and right mask pixels increases uniformly towards the edges of the array ("Model 2").

As with the 23×23 mode investigated in RD2, a possible candidate for the background is clock-induced charge. Our understanding is that clock-induced charge events occur when the shifting of charges between pixels generates one or more additional electrons. This can occur randomly in any pixel, and increases as the clock voltage increases (as occurs in Table 2). The electron multiplying register simply amplifies this additional charge, so unlike readout noise it cannot be overcome by increasing the electron multiplying gain.

As a consistency check, the clock-induced-charge specification for the Andor camera (at -85° C, with a 100 ns serial shift time, a 500 ns vertical shift time, and an EM gain of $1000\times$) is 0.005 events per pixel for a 30 ms exposure time[2]. In our case, with an electron multiplying gain of 250 and a conversion factor of 12.27 electrons per data number, a single spurious electron will generate on average 20.4 counts. If there are 0.005 of these per pixel in an average frame we would expect to see 0.1 additional counts per pixel due to clock induced charge if there was no pixel summing. In the custom clocking mode, however, each pixel is actually the sum of between 4 and 16 pixels of charge, so we would expect to see a variation of ~ 1.2 counts across the field due to clock induced charge. If we discount the hypothetical edge effects discussed above, we find the remaining variation in practice is ~ 1.7 when the parallel clocking voltage setting is "2", i.e. not inconsistent with our expectation. However, the figure is much lower in practice when the parallel clocking voltage is "normal", in fact it is not easily measureable.

As for why clock induced charge would increase at the left and right edges of the chip, we have no explanation. One possibility is that propagation, edge or reflection effects are distorting the parallel clocking signal along the chip edges so that it does not step as cleanly as elsewhere (the rise time of the transition is known to affect how much clock induced charge is generated [2]).

We have attempted to assess the impact of this quasi-deterministic non-uniform background by examining the histograms of the pixel values (in digitised counts) and RMS pixel deviation from the mean ("σ", in electrons) for each of the image stacks shown in Table 2. These results are displayed in Figure 7 and the RMS values shown in Table 3.

Several trends are apparent:

- The probability curves for custom clocking are offset from those of conventional clocking by about 5 counts.
- Each line is a combination of a smooth curve and a wing that trails off towards higher counts according to a power law. This has the effect of increasing the RMS pixel deviations, because the standard deviation calculation gives extra weight to outliers.
- The counts in the wing are higher for custom clocking, and higher for

[2]Unfortunately Andor does not specify what clocking voltage setting is used for this measurement.

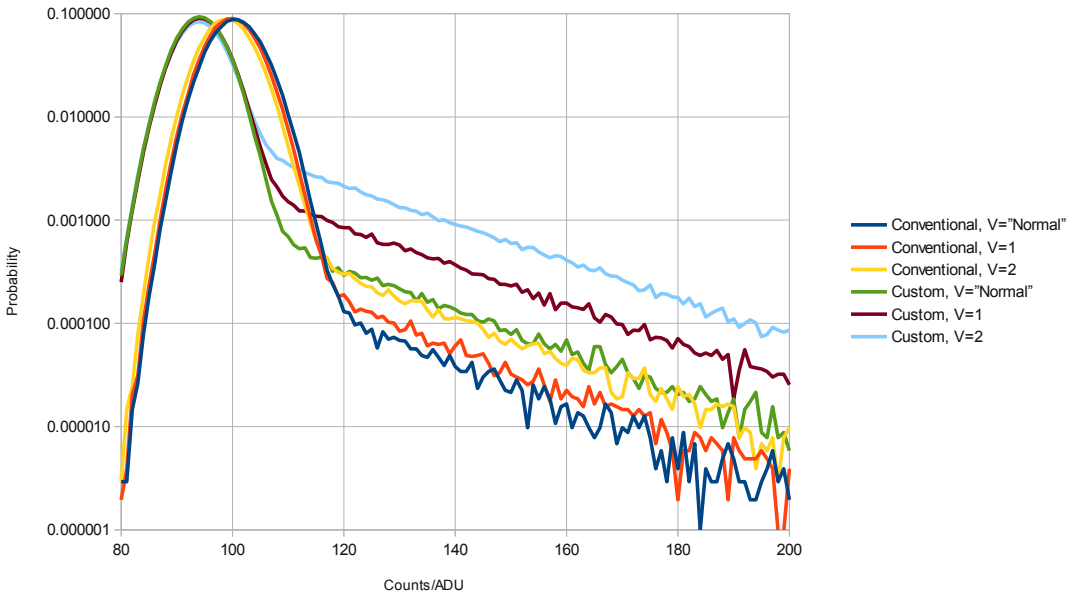

Figure 7: Pixel histograms of the image stacks shown in Table 2. Note how use of custom clocking causes a shift in the distribution to the left by about 5 counts, and how use of custom clocking and higher parallel clocking voltages lead to an undesirable wing towards higher pixel values.

Clocking scheme	Clocking voltage setting		
	V="normal"	V=1	V=2
Conventional	62/56	63/54	72/55
Custom	78/52	109/52	164/54

Table 3: RMS pixel deviation (in electrons) from the mean for the data in Figure 7, referred to the output of the gain register, for several parallel clocking voltages. Two figures for the RMS deviation are given for each dataset: the first includes all the pixels, the second excludes pixels whose counts are in the high wing in the histogram.

higher parallel clocking voltage. This is consistent with the dominant noise source being clock induced charge.

- The distribution of values (with the exception of those in the wing) is about the same for all the curves. The curves vary mostly in how many counts are found in their wings.
- When the data in the wing is excluded (Table 3) then the RMS deviations for all the curves are almost the same. So the wings are responsible for most of the variation in RMS deviations between curves.

We can see that the lowest noise option available with conventional ROI clocking is where the parallel clock time is 500 ns and a "Normal" voltage is used. In this case, this gives a dark frame RMS pixel deviation $\sigma = 62$ electrons. For the custom clocking scheme, the same parallel clock time and voltage give the lowest RMS pixel deviation ($\sigma = 78$ electrons) and the smallest wing of the custom schemes. For an electron multiplying gain of 250, the effective readout noises for these two cases are thus 0.25 and 0.36 electrons respectively.

For the faintest targets expected for the MROI FTT system, we know that the typical photon rate will be of order 50–100 detected photons per ~ 3 ms frame. At this level, the centroid error will be wholly dominated by photon noise fluctuations and we expect these marginal increases in effective readout noise to be largely irrelevant in compromising the performance of the FFT system as a whole. As a result, we opted to conduct the remaining tests with a 500 ns parallel clock and a voltage of "Normal", as being representative of the settings to be used at the MROI installation.

6.3.2 Effect of temperature on custom clocking

A key element of the system design of the MROI is that effects that have the possibility to degrade the local seeing at the observatory site be minimised. As a result, it is desirable that the power consumed by the camera should be as low as possible. Because most of the power is consumed by the camera's Peltier cooler, it follows that reducing the demand on this (by operating the CCD at a warmer temperature) should significantly reduce the camera power consumption. The trade-off is that a warmer CCD also generates greater thermal noise, which raises the question of how warm the CCD can be before this becomes significant for FTT operation.

In RD2 it was found that detector temperatures higher than -60°C caused increasing amounts of noise in the images. Hence for this test, only -80°C and -60°C were tested.

To assess the impact of running the chip at these two temperatures, tests were run where series of 1000 dark frames were acquired for these two temperatures. The parameters held constant across these tests were:

- EMCCD gain = 250 & Pre-amp gain 5.2×;

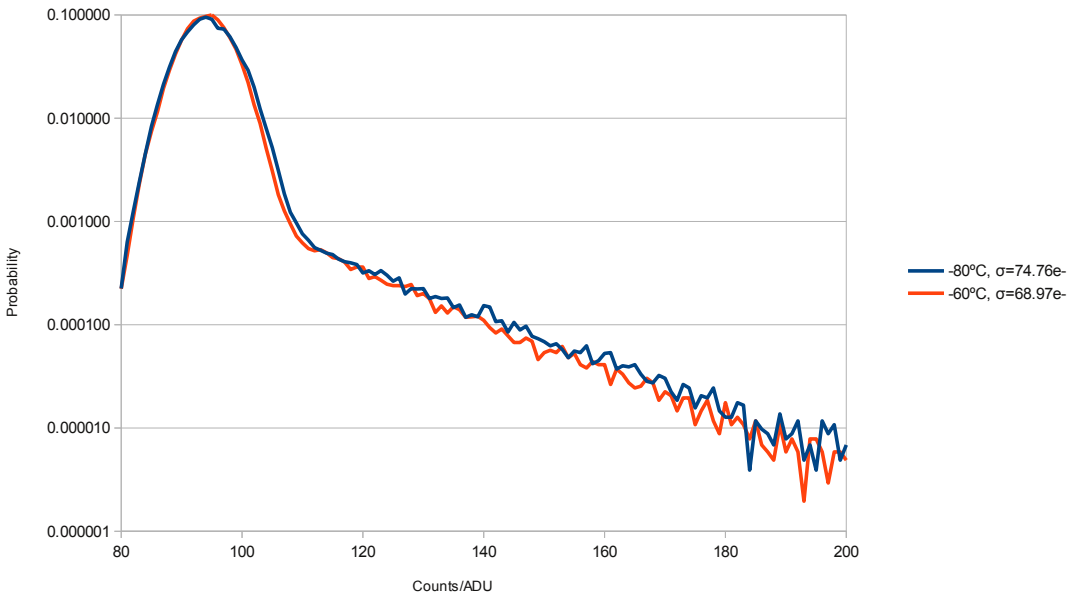

Figure 8: Comparison of dark count distribution for two image stacks of 1000 frames each, one acquired at -80°C and the other at -60°, together with the standard deviations of each in electrons.

- Parallel clock period = 500 ns;
- Parallel clock voltage = “normal”;
- Serial clock period = 100 ns;
- Frame rate = 300 Hz;
- Sub-frame size 32×32 pixels and sub-frame corner position (240, 240).

The frame rate chosen was the slowest FTT frame rate suggested in RD1, and as such was chosen to maximise the thermal noise contribution, thereby providing a “worst case” scenario. Our results are summarised in Figure 8.

Figure 8 shows almost identical count distributions at -80°C and -60°C. If anything, there are slightly fewer excess counts at -60°C. Hence it is clear that the camera could be operated at -60°C rather than -80°C with negligible CCD thermal noise penalty and significant reduction in power dissipation.

6.3.3 Effects of integration time on custom clocking

Although the FTT system is expected to normally operate with a servo bandwidth of 40–50Hz, and hence a relatively short integration time, a longer integra-

Temperature/°C	Integration Time/ms		
	1	5	10
-80	75.5	74.7	75.8
-60	67.8	70.5	71.7

Table 4: RMS electron noise per pixel per frame, as a function of temperature and integration time for the custom clocking scheme.

tion capability is potentially useful for observations when the seeing is particularly good. Hence, the dark frame performance of the camera was also checked for integration times of 1, 5 and 10 ms and CCD temperatures of -60°C and -80°C. For these tests we used the custom clocking scheme and the following "standard" settings:

- EMCCD gain = 250 & Pre-amp gain = 5.2×;
- Parallel clock period = 500 ns;
- Serial clock period = 100 ns;
- Parallel clock voltage = "normal";
- Sub-frame size 32×32 pixels and sub-frame corner position (240, 240).

The results of these tests are summarised in Table 4 and show that an increase in integration time to 10 ms at -60°C causes only a minor noise penalty, and there will likely be no benefit to running the camera colder than this when observing at the MROI.

In fact, when acquiring full 512×512 pixel frames with a 10 second integration time (as might be desirable for target acquisition) no increase in background level is seen with a CCD temperature of -60°C.

6.4 Illumination Tests

Additional tests were also carried out using the faint LED target. Firstly, the source was strobed to test for residual charge on successive frames, and secondly the source was reduced in brightness so as to test the centroiding accuracy as a function of light level.

6.4.1 Effect of using a faint strobed source during custom clocking

During observing, the source image will move across the CCD on millisecond timescales as the turbulent atmosphere refracts the source light. If the charge generated by this light is not completely removed after an exposure, then the

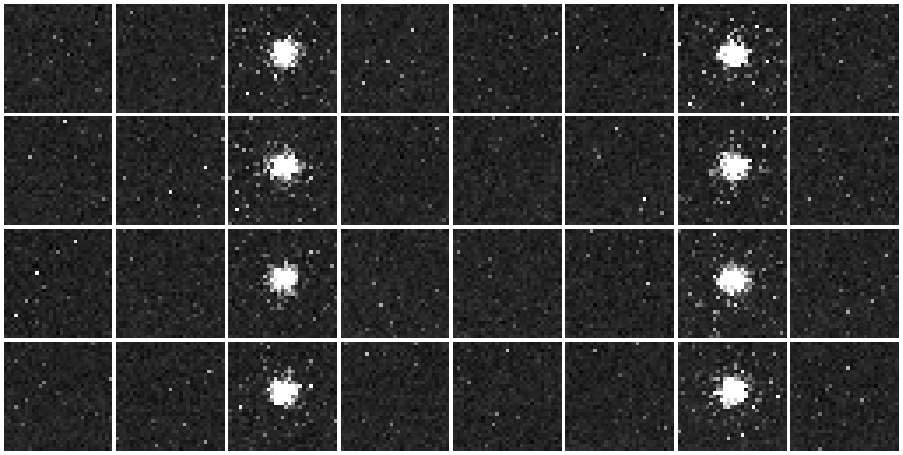

Figure 9: Image sequence in which the CCD is illuminated by the light emitting diode source for one frame in every four. The images should be read from left to right and down the page. The un-illuminated frames are inspected for residual charge. All images here have the same grey scale, which has been intentionally adjusted to make background noise visible.

remaining charge will contaminate the following exposure and cause an effective lag in the computed centroid.

To determine the extent of any residual charge contamination, a 1000 frame image sequence was acquired with the source light emitting diode on for one frame in every four and the frames examined for the presence of residual charge in the images when the LED was off. For these tests we only used the custom clocking scheme with a standard set of camera parameters:

- EMCCD gain = 250 & Pre-amp gain 5.2×;
- Parallel clock period = 500 ns;
- Serial clock period = 100 ns;
- Parallel clock voltage = “0”;
- Temperature = -80°C;
- Sub-frame size 32×32 pixels and sub-frame corner position (240, 240).

A typical sequence of frames is shown in Figure 9.

As a further check, all the 250 frames that immediately followed strobe frames were summed together to form an image (Figure 10).

Under these conditions, no evidence was found for residual charge in the un-illuminated frames or even in a summed image designed to expose such an artifact.

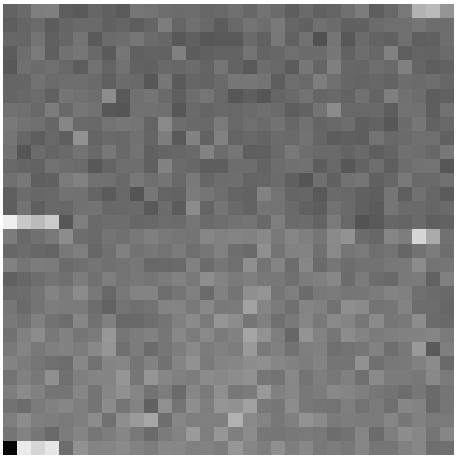

Figure 10: Image which is the sum of the images that immediately follow images in which the strobe light appears. No strobe signal is visible.

6.4.2 Effects of clocking scheme on centroiding accuracy

In order to test the impact of the custom clocking scheme on the ability of the camera to be used to centroid a faint target, a test was undertaken whereby the RMS "jitter" of the centroid of a fixed target was compared to that expected from theory (Section 4), and this was repeated for a range of light levels.

The absolute photon flux from the light source was firstly calibrated using the camera in photon counting mode, with the test room blacked out and the light path shielded by a cardboard tube. The flux of the source was adjusted by placing calibrated neutral density filters in front of it. When an N = 5.5 neutral density filter was placed in the beam the camera counted 48.9k photons over a total integration time of 20 seconds, and 32.6k counts of background over the same period with the source switched off[3]. Hence ∼16.3k photons, or ~ 816 photons per second, could be attributed to the source. Given the repeatability of the data, together with the uncertainty in the optical density of the ND filter used, we estimate that this rough calibration was accurate to approximately 20%.

It is known from RD1 that a 15.4 magnitude target is expected to produce a flux of 19800 photons per second on the detector. Hence, given any flux F we can calculate what the stellar magnitude equivalent should be using the formula

$$m = -2.5\, log_{10}(\frac{F}{19800}) + 15.4\,. \tag{3}$$

To produce a range of fluxes a range of neutral densities were used, ranging from 5.5 to 1.5 in 0.5 increments. Each reduction in neutral density by 0.5 increases the flux by $\sqrt{10}$.

[3]The camera actually took 10000×2 ms exposures in each case. This kept the maximum photon count per pixel per exposure at no more than 0.2, that is, well within the photon counting domain.

After this calibration, 1000 frame image stacks were collected in analog detection mode over the following region of the parameter space:

- Neutral densities of 1.5 to 5.5, corresponding to celestial sources of magnitudes 8.9 to 18.9;
- Conventional ROI and custom clocking modes at 300 Hz, the frame rate suggested in RD1 for the faintest sources;
- Custom clocking at a frame rate of 1 kHz to determine the practical limiting performance at that rate.

As per usual the following parameters were held constant:

- Temperature = -80°C;
- EMCCD gain = 250 & Pre-amp gain = 5.2×;
- Parallel clock period = 500 ns;
- Serial clock period = 100 ns;
- Parallel clock voltage "normal";
- Sub-frame size 32×32 pixels and sub-frame corner position (240, 240).

Examples of images for the different clocking parameters and neutral density levels are presented in Table 5. Centroids for each image stack were then calculated from the pixel data using the following quad cell algorithm:

- A 10×10 quad cell (four quadrants of 5×5 pixels) was overlaid on the data, with the centre within a pixel's distance of the known centre of the image. To achieve this during observing, one could do a first pass through the data calculating a barycentre, which would determine where to place the quad cell centre for a second pass using the algorithm listed here.
- All the pixels in the image stack that were not contained within the quad cell were averaged to determine a background level.
- This background level was subtracted from every pixel in the image stack, so that the background was approximately zero.
- For each frame of data in the resulting image stack, the position of the image in pixels was determine using a standard quad cell algorithm, i.e. with equal weights[4]

[4]The calibration of the quad-cell algorithm, so as to determine the scale-factor between normalised displacement and actual pixel values, was determined using simulated image data generated for moderately high light levels, and with the same image FWHM to pixel size ratio as our experimental conditions.

Neutral Density / Magnitude	Clock Sequence		
	Conven., 300Hz	Custom, 300Hz	Custom, 1000Hz
1.5 / 8.9	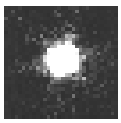	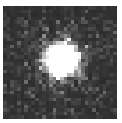	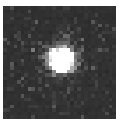
2.0 / 10.1	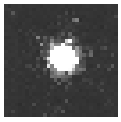	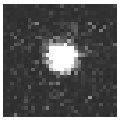	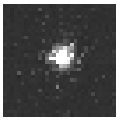
2.5 / 11.4		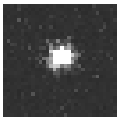	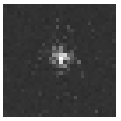
3.0 / 12.6	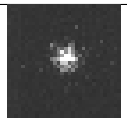	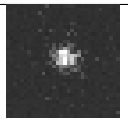	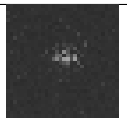
3.5 / 13.9	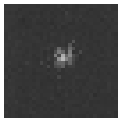	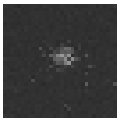	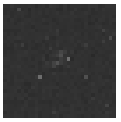
4.0 / 15.1		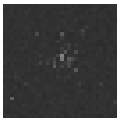	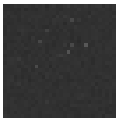
4.5 / 16.4	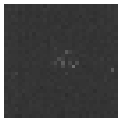	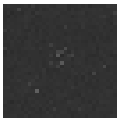	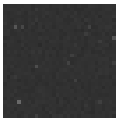
5.0 / 17.6	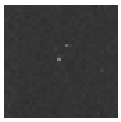		
5.5 / 18.9			

Table 5: Examples of sub-frame images as a function of clocking scheme, frame rate and neutral density (i.e. target magnitude). The grey scale is constant across all images and ranges linearly from 0 to 500 counts.

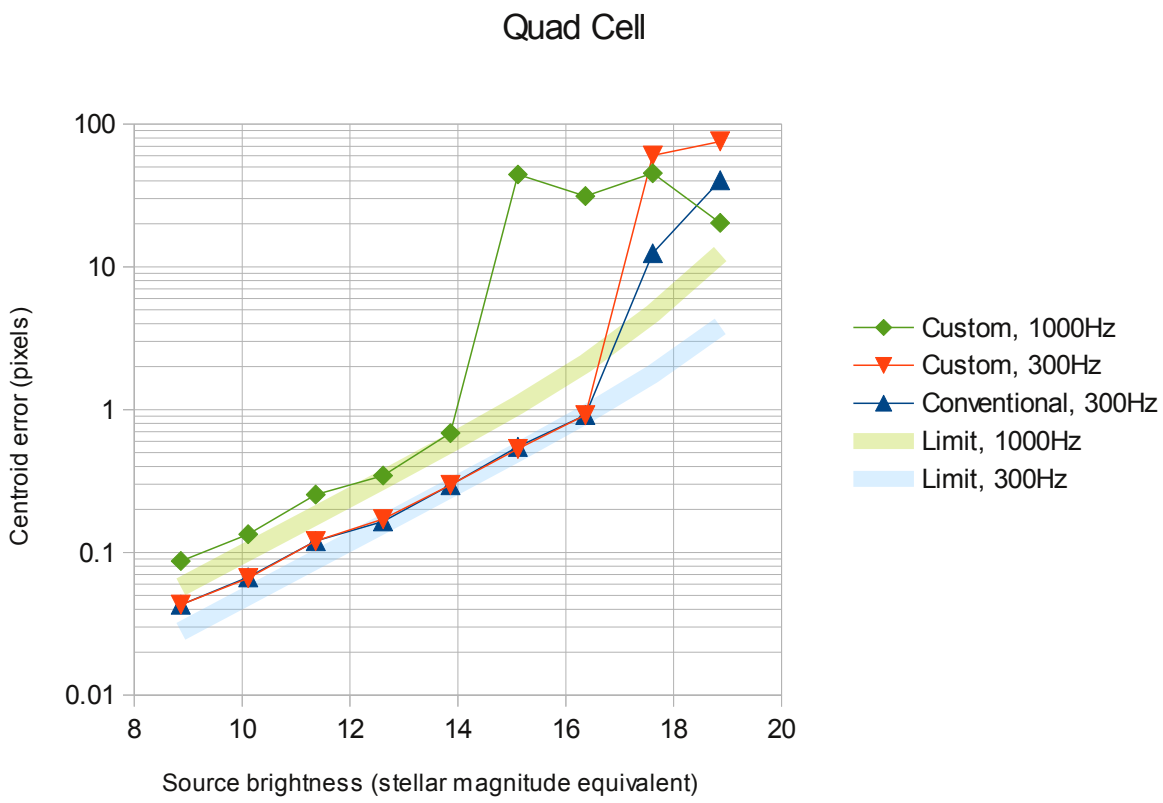

Figure 11: Two-axis error of calculated centroid, as a function of clocking scheme, source magnitude at the MROI and frame rate. The theoretical limits to centroiding performance for frames rates of 300 Hz and 1000 Hz are also shown. Due to the calibration uncertainties of our experimental set up these predictions are only accurate to the 10% level, and hence these limits are shown as shaded bands.

The measured RMS jitters of these centroid estimates were then compared with the expected values predicted by the equations of Section 4 using values of $n = 100$ pixels and σ_r=51.8/250 electrons (the manufacturer's specified readout noise divided by an electron multiplying gain of 250). Our results and these theoretical predictions are shown in Figure 11.

It can be seen that the experimental results are close to the theoretical predictions until the source brightness falls below some critical level. More specifically, the centroid error appears to degrade rapidly once the source brightness reaches magnitude 15 (1000Hz frame rate) or magnitude 17 (300Hz frame rate). These rather dramatic thresholds are actually artefacts of the quad-cell analysis that we have used, since in the low signal regime small errors in the background estimation can lead to large errors in the derived centroids.

In order to confirm this, the same data were analysed independently using a "knife edge" method to validate the increase in "centroid" error as a function of

source brightness. For this analysis, the following steps were taken:

- For each image in a sequence of 1000 frames, the difference between the counts in the left and right halves of the quad cell was calculated.
- The average RMS value of the difference, Δ, was calculated for the sequence.
- The value was normalised by the flux of the incident light, using the brightest test with neutral density of 1.5 as a reference. So if a sequence was acquired with a neutral density of 2.5, the raw Δ would be multiplied by 10.
- The test was repeated for all the datasets listed in Table 5.
- The results were plotted in Figure 12.

Figure 12 shows no unstable behaviour at the faintest magnitudes, indicating that the data are well behaved. Furthermore, the test data converges there, as would be expected if noise rather than signal is dominating.

The principal conclusions to be drawn from our tests are as follows:

- At a 300 Hz frame rate, there appears to be no substantial difference in centroid estimation between the conventional ROI and custom clocking modes, at least down to photon rates comparable to those expected from a magnitude 16 stellar target at the MROI.
- For both the conventional and custom clocking schemes, the measured centroid jitter is broadly speaking consistent with that expected for data compromised by the photon, electron-multiplication and readout noise levels expected for the camera when operating at a frame rate of 300 Hz.
- As mentioned in the introduction, there is a derived requirement that the centroid error be less than 0.23 pixels. As RD1 notes, this requirement cannot quite be met even in principle due to the scarcity of available photons. However, it can be approached: at 300Hz, objects down to 13th magnitude qualify and at 1000Hz, objects down to 11th magnitude qualify.

  Furthermore, all these results are based on a quad cell implementation because that case can be treated analytically. Other algorithms, such as those incorporating Kalmann filtering, may have better performance.

6.5 EMCCD Temperature Stability

Under observing conditions, the ROI of the CCD will be read out rapidly and continuously over periods of several minutes at a time. The charge clocking signals will generate significant amounts of heat in the detector, with the potential to alter readout parameters such as background baseline level and signal

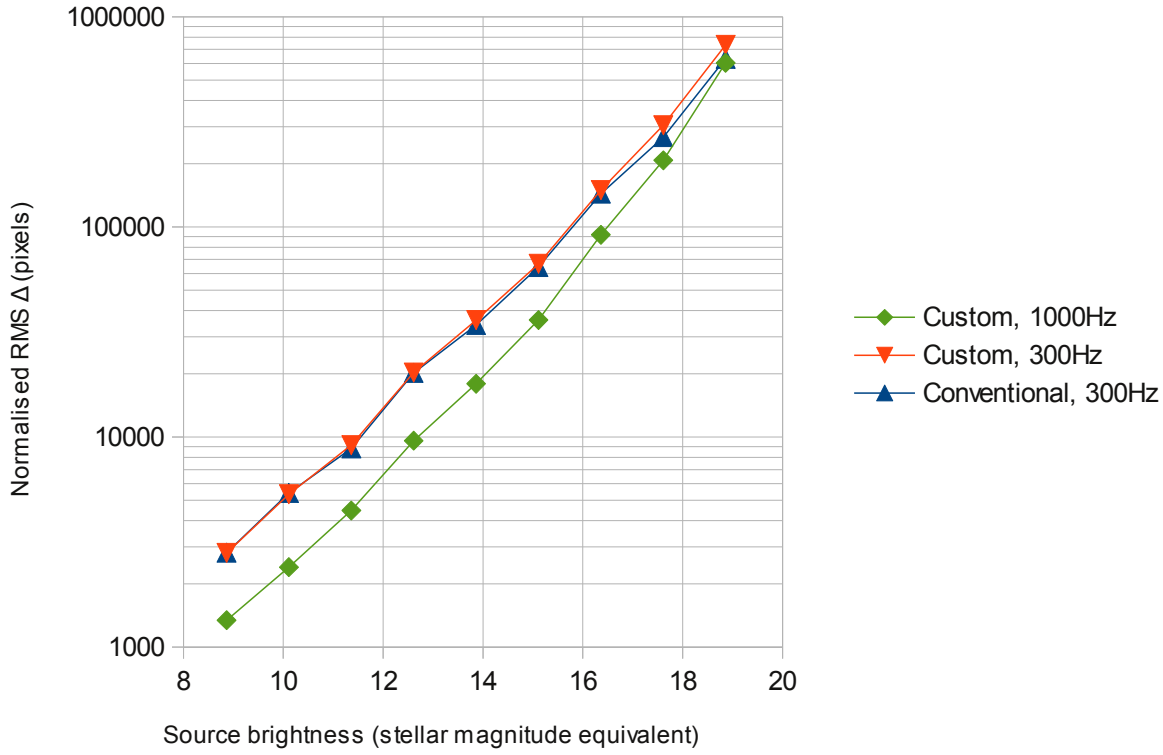

Figure 12: Knife edge tests for the data in Table 5. There is no unstable behaviour at the faintest magnitudes.

to noise. It is therefore of interest to establish the ability of the camera's Peltier cooler to maintain a set temperature under conditions of sustained readout, and "stop-start" readout, as will be encountered at the MROI when switching between targets. Water cooling is also required as that is how the camera will be cooled at the MROI.

Unfortunately it is not possible to read the temperature during a continuous image acquisition, it is necessary to stop the readout first. The Andor user interface software has additional limitations: it samples every 10 seconds and displays temperature to the nearest 5°C.

For the investigation, the following were held constant:

- EMCCD gain = 250 & Pre-amp gain = 5.2×;
- Parallel clock period = 500 ns;
- Serial clock period = 100 ns;
- Parallel clock voltage "normal";

Clocking scheme	Set Temperature/°C		
	-80	-70	-60
Conventional, 300Hz	-80	-70	-60
Conventional, 400Hz	-80	-70	-60
Custom, 300Hz	-80	-70	-60
Custom, 400Hz	-80	-70	-60
Custom, 500Hz	-80	-70	-60
Custom, 600Hz	-80	-70	-60
Custom, 700Hz	-80	-70	-60
Custom, 800Hz	-75, -85	-75	-60
Custom, 900Hz	-75, -90	-75	-60
Custom, 1000Hz	-70, -90	-75	-65

Table 6: Greatest changes in detector temperature from set point for a range of set temperatures and clocking rates after continuously reading out the ROI for five minutes. A reading of "-70, -90" means that immediately after readout was stopped, the temperature increased to -70°C and then decreased to -90°C before stabilising at the set temperature.

- Sub-frame size 32×32 pixels and sub-frame corner position (240, 240).

The ambient temperature was 21°C. The cooling water was regulated at 18°C and had a flow rate of 3.3 litres/minute. This is slightly higher than the 3 litres/minute expected at the MROI, but a check was done for this as discussed below.

The performance of the Peltier cooling servo was investigated for a subframe clocked out using the custom scheme at rates from 300Hz to 1000Hz. Conventional clocking at 300 and 400Hz was also tested. Peltier servo set points of -80°C, -70°C and -60°C were used.

While the camera was idle, the temperature would be set and allowed to stabilise at the set point. The camera was then made to continously read out the ROI for five minutes, then was stopped, and the temperature response was watched for excursions from the set point. The maximum deviations in either direction are listed in Table 6.

The temperature measured after readout stopped was always the same as the set temperature as long as the frame rate was below 800Hz.

When the set temperature was -80°C and the frame rate was 800Hz or more, the first reading taken at the end of the run was at a higher temperature, indicating that warming had occurred. The detector temperature then cooled until it was below -80°C, before stabilising at the set point. This is likely to be due to lag in the Peltier servo: stopping the readout suddenly removes the heat load on the servo but the servo responds too slowly to prevent an overshoot.

When the set temperature was -70°C, at fast frame rates the response was more

benign. There was no evidence that warming had occurred during readout, but there was still a mild overshoot at the end of it.

When the set temperature was -60°C, the behaviour was better: a mild overshoot was seen, but only for 1000Hz readout. For all other frame rates tested, the temperature remained at the set point.

This is a good result: we already know from Subsubsection 6.3.2 that the camera can be run at -60°C without sacrificing performance. Here we see that at this temperature we can operate at 1000Hz without warming the CCD, and if we want complete stability even when switching between targets we can still operate at 900Hz. At this frame rate the servo bandwidth is still 49.6Hz, just short of our goal of 50Hz and comfortably above our requirement of 40Hz. In addition, at the MROI the cooling water temperature is likely to be much cooler than 18°C, which further reduces the load on the Peltier cooler[5]. This might improve the stability when readout at 1000Hz stops.

The flow rate for this test was 3.3 litres a minute, which is slightly more than the 3 litres a minute expected at the MROI. To ensure that this did not affect the results, the flow rate was then reduced to 2 litres/minute (less than expected at the MROI, but still within the bounds of what is possible there). The 900Hz, -60°C run was repeated, with the same result: there was no temperature fluctuation observed.

7 Summary

The key conclusions from these initial tests of the custom clocking scheme developed by Andor can be summarised as follows:

- The custom clocking scheme allows our 40 Hz closed loop bandwidth requirement and our 50 Hz goal to be realised with some margin.
- At a frame rate of 300 Hz there is no evidence to suggest that the custom clocking scheme degrades the ability to centroid a stellar image, beyond that which is available using conventional ROI readout.
- The custom clocking scheme leads to a non-uniform "background" in the ROI of the order of 0.5 counts, and some peripheral artefacts of the order of 5 counts. However, there is no evidence that this is at a level which affects centroiding at the lowest light levels.
- When the camera is water cooled, the CCD temperature can be held steady at -60°C, even when readout is intermittent, as long as the frame rate does not exceed 900Hz. At 1000Hz the temperature holds steady during readout, but fluctuates by 5°C when readout stops.

[5]Colder water was not tested in this experiment due to the risk of condensation.

These conclusions are very reassuring, and suggest that we have an operating mode that meets the most critical requirements of the MROI FTT system. Hence we recommend that Andor be formally contracted to deliver the scheme.

References

[1] E2V Technologies, "CCD97-00 Back Illuminated 2-Phase IMO Series Electron Multiplying CCD Sensor", E2V Technologies Limited, UK (May 2004).

[2] E2V Technologies, "Low-Light Technical Note 4: Dark Signal and Clock-Induced Charge in L3Vision(TM) CCD Sensors", E2V Technologies Limited, UK (June 2004).

[3] Andor Technology PLC, "Hardware Guide: iXonEM+", Version 1.2, UK (August 2008).