MRO Fast Tip-Tilt

Andor iXon X3 897 Custom Clocking Evaluation

MRO-TRE-CAM-1120-0118

Bodie Seneta, Alex Rea and Chris Haniff
bodie@mrao.cam.ac.uk
a.rea@mrao.cam.ac.uk
cah@mrao.cam.ac.uk

rev 1.0
15 July 2011

Cavendish Laboratory
Madingley Road
Cambridge CB3 0HE
UK

Change Record

Revision	Date	Authors	Changes
0.1	2011-03-30	EBS	First draft.
0.2	2011-06-10	EBS	Second draft, using results from purchased camera.
0.3	2011-06-26	EBS	Complete rework including fix to very bad noise calculation error and new centroid test results.
0.4	2011-07-04	EBS	Centroid analysis reworked for consistency between real data and theory. Many minor corrections.
0.5	2011-07-08	CAH	Corrections to centroid analysis and photon rates (fig 5 revised). Reorganised sections. Finalised conclusions and recommendations. Additional minor stylistic corrections.
0.6	2011-07-14	CAH/JSY	Corrections to many sections via DFBs comments. Revised recommendations. Updated figures.
1.0	2011-07-15	CAH	Final tweeks to request suggested by DFB. First release.

Objective

To assess the suitability of an Andor iXon X3 897 camera for the Magdalena Ridge Observatory Interferometer fast tip-tilt (FTT) system when the camera has been upgraded with a trial customised fast CCD readout mode written by Andor. Results using the standard readout mode are included for comparison.

Reference Documents

RD1 MRO-TRE-CAM-0000-0101: Derived Requirements – rev 1.0, August 31st 2010

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.

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

The 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. 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.

Since RD1 was written, the optical design of the fast-tip-tilt (FTT) system has been frozen, and the pixel scale of the camera has been fixed at 0.147 arc-seconds per pixel. This implies that the sub-frame size must be at least as large as 25×25 pixels, and that a detector with dimensions of at least 408×408 pixels will be needed. Furthermore, the camera readout latency has been measured which has allowed us to establish a minimum frame rate requirement (see Subsection 6.1).

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 500Hz, but the readout latency is too long to allow our closed-loop bandwidth requirement to be met.

In order to deal with this problem, Andor has written a trial customised fast clocking scheme for evaluation. This custom scheme minimises 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. This was felt to be a reasonable approach because the typical number of stars in the 60×60 arc-second field of view is expected to be no more than one or two in all cases.

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 the required minimum frame rates;

- 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 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 the signal levels tested. The standard ROI readout does not meet our latency requirement but the custom readout mode does. However, a number of features of the custom readout scheme still require modification and these are reported in Section 8

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 all the pixel rows prior to the first row containing the ROI being discarded. The CCD rows that contain the region of interest are then 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 in 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 23×23 sub-frame centred on the 512×512 pixel light-sensitive region of the chip could be read out at around 520 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.

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, read, and these values discarded. 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. 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 until charge from all the pixels from the region of interest are stored entirely within the readout register. Finally, the whole of the serial readout register is read out a single time so as to obtain the pixel values for the ROI in a sequential form.

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 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"

3. 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 it was this that led to the custom region of interest being fixed at 23×23 pixels.

The trial custom readout scheme was limited to a region of interest that was displaced 100 pixels in the serial and parallel shift directions from the origin of the light-sensitive area of the chip. In our lab tests with the Andor X3 897, we found that this 23×23 sub-frame could be read out at approximately 1500 Hz, i.e. some three times faster than with the conventional ROI readout.

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. The tests were conducted in a darkened room, and where possible cardboard sheets were used to shield the experiment from any stray ambient light.

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-F0010 light emitting diode (LED), chosen for its rapid 15 ns response time, peak

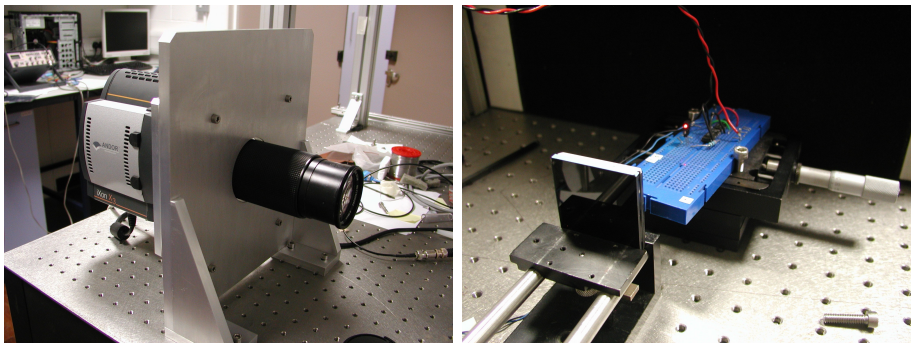

Figure 1: Left: The Andor camera in its custom mount, facing the light source at the far end of the optical table. Right: The LED light source (mounted on the blue breadboard) with its neutral density attenuators in front.

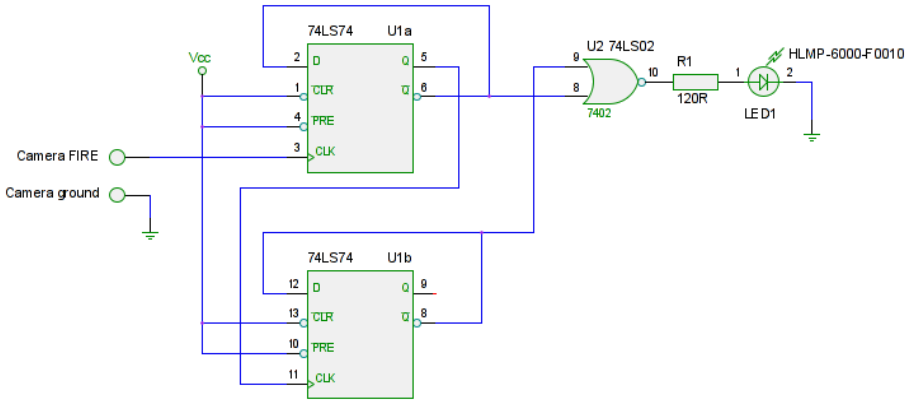

Figure 2: Circuit schematic for the strobed light source.

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 [2] 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 4.1 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 programs Spydr and 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 the required minimum frame rate;
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.

All the tests took place between the 1st and the 21st of June, 2011.

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 the minimum frame rate requirements.

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 could be used to measure the readout latency.

As mentioned earlier, the test sub-frame was a 23×23 region displaced 100 pixels in the serial and parallel shift directions from the origin of 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 1.85 ms and 0.69 ms respectively.

If we assume a minimum exposure time equal to the readout time, we can see that 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.

In our tests, the following parameters were held constant:

- CCD temperature = $-80°$C;
- EMCCD gain = 250 & Pre-amp gain = $5.2\times$;
- Frame rate = 300 Hz;
- Serial clock period = 100 ns (the minimum available);
- Sub-frame size 23×23 pixels & Sub-frame position (100, 100).

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 and frame rate were chosen to be those values suggested in RD1 for observations of the faintest required targets. The parallel clock periods were restricted to 300 ns and 500 ns, the shortest available, as our latency requirement implies the readout must be as fast as possible, but results were secured for all five parallel clock voltage settings to see which gave better image quality. Our results are displayed in Table 1.

The most striking result observed 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. Clock voltages of "2" or higher produced images for both readout schemes, although images acquired via custom clocking had slightly more background noise, and this noise increased with parallel clock voltage.

Clock Sequence	Clock Voltage				
	0	1	2	3	4
Conventional, 300 ns	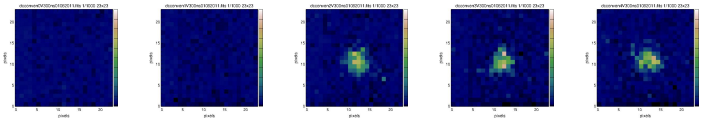				
Custom, 300 ns	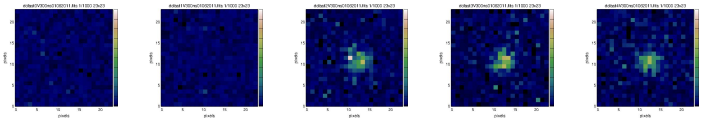				
Conventional, 500 ns	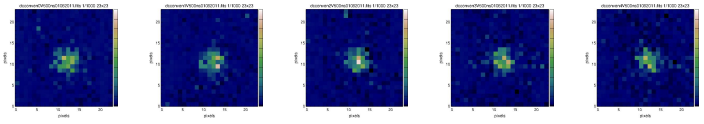				
Custom, 500 ns	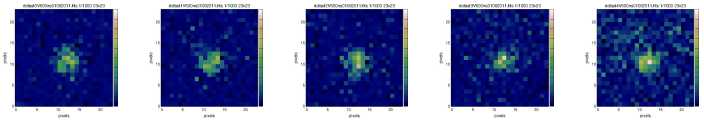				

Table 1: Typical images as a function of parallel clock voltage, clocking scheme, and parallel clock period. An identical colour mapping has been applied to all images so as to allow straightforward comparison. Note that for the shortest clocking times of 300 ns, a clock voltage of at least 2 was necessary to transfer charge out of the CCD at all.

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, the noise was seen to increase with increasing clock voltage, with more noise apparent in the custom case.

These results were useful for constraining the parameter space requiring further investigation. If 300 ns clocking is used, clock voltages of "0" or "1" cannot be used. It is also apparent that of the remaining options, the noise can be minimised by using the lowest clock voltage available.

6.3 Dark Frame Quality

6.3.1 Effects of clocking scheme and parallel clock voltage

In these tests, the parameter space outlined in Subsection 6.2 above was explored with no light falling on the image sensor. In each case an image stack of 1000 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 a darker vertical band on the left and a

Clock Sequence	Clock Voltage				
	0	1	2	3	4
Conventional, 300 ns	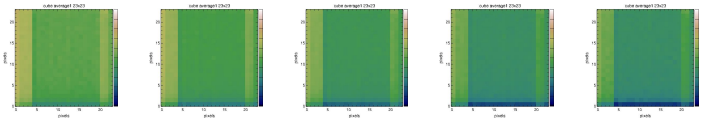				
Custom, 300 ns	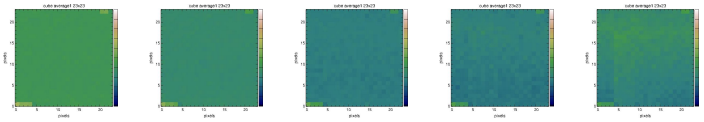				
Conventional, 500 ns	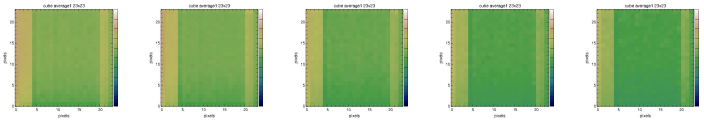				
Custom, 500 ns	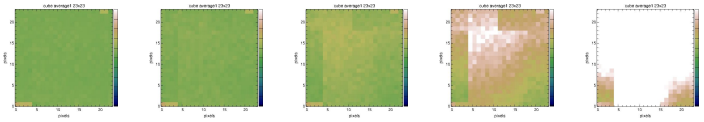				

Table 2: Dark-frame averages of 1000 frames as a function of parallel clock voltage, parallel clock period, and clocking scheme. The colour mapping is the same for all images and ranges from 85 to 110 counts.

bright horizontal region towards the top, and this pattern became more dominant with increasing clocking voltage. In both cases, we believe that the systematic offset is likely due to charge from pixels outside the region of interest getting summed with charge from pixels inside it. Even if those external pixels are unilluminated, they can still carry thermally- and clock-induced charges.

In the case of conventional clocking, it is assumed that the charges contained in any light-sensitive pixels never get summed together. However, there are masked "guard pixels" surrounding the light-sensitive CCD region, which are used to calibrate the background level, and it is possible that these are the source of the bands seen in Table 2.

The origin of artifacts seen when the custom clocking mode was used is less clear. We investigated one possible source by way of a numerical simulation in which every pixel of a "virtual" CCD was given a charge value of "1", and this chip was then clocked out according to the custom clocking scheme. 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, or equivalently, the number of pixels that had been summed into each output image pixel, by the time it would have been read out.

Figure 3 compares the result of this simulation with our real averaged dark frame data, and it is clear that there is qualitatively good agreement between the two. The surprise from the simulation is how much charge from extraneous pixels gets summed into the output image. Every image pixel in the final ROI appears to be the sum of charges from at least four pixels (at the lower left), and up to twenty-

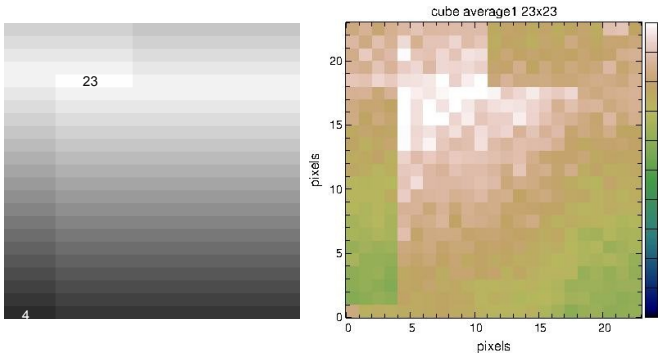

Figure 3: A comparison of our custom clocking simulation with real dark frame data. Left, simulation showing the number of CCD pixels that get summed into each image pixel. Right, a real 1000 frame dark average at "+3" clocking voltage.

three pixels towards the upper left of the ROI. In an ideal CCD at the MROI, this would not matter, as we expect the field to be mainly dark (i.e. with only a single target in the full field of view) and any additional pixels should thus contribute no charge. However, in reality there appears to be a roughly equal "non-target-related" contribution from every pixel in the CCD in our dark frames.

This contribution is unlikely to be thermal dark current, because the thermal contribution is only expected to be 0.001 electrons per pixel per second under our measurement conditions. Nor can it be readout noise, because that is only added on readout and does not depend on how the pixels are summed beforehand.

A remaining possible noise source is clock-induced charge, and we believe this may be the likely candidate. 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.

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. Hence, for a summed readout of 1000 frames, but with a shorter exposure time as in Figure 3, and assuming a conversion factor of 12.27 electrons per data number, one would expect at most roughly two data counts at the bottom left and up to ~ 9 events in the worst-case summing region towards the upper left of the ROI[2]. Our summed data, collected with an exposure time of a few milliseconds, showed an enhanced count of roughly 0.5 data numbers in the brightest region of the stacked dark frame, i.e. not inconsistent with our expectation.

[2]For an exposure time shorter than 30 ms, one would expect a decreased contribution from thermally-generated spurious events.

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 4 where the RMS values are refered to the output of the gain register.

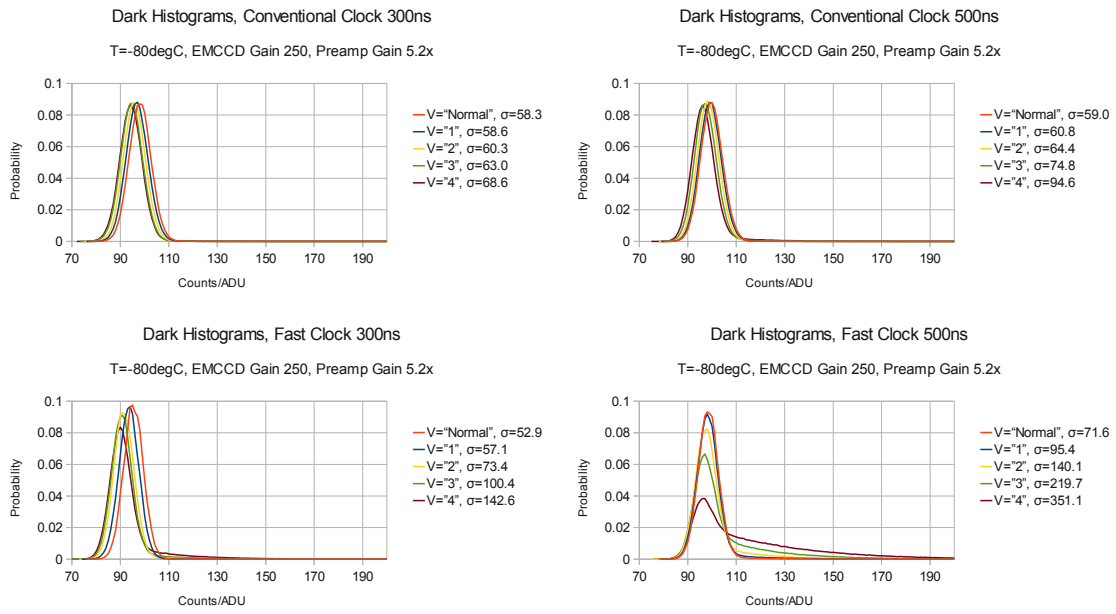

Figure 4: Pixel histograms of the image stacks shown in Table 2 and corresponding RMS pixel deviation from the mean ("σ"), in electrons, referred to the output of the gain register. Note how for the custom "fast" clocking scheme (lower row) the use of a high clocking voltage leads to an undesirable wing towards higher pixel values.

If we eliminate those regions of parameter space that we know are uninteresting, i.e. where 300 ns parallel clocking is combined with a clock voltage of "0" or "1" (Subsection 6.2), then 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 $\sigma = 59.0$ electrons. For the custom clocking scheme, the same parallel clock time and voltage give the lowest RMS pixel deviation, although in this case it has $\sigma = 71.6$ electrons. For an electron multiplying gain of 250, the effective readout noises for these two cases are thus 0.24 and 0.29 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.

To assess the impact of running the chip "warm", tests were run where series' of 1000 dark frames were acquired for a range of temperatures and parallel clock voltages. The parameters held constant across these tests were:

- EMCCD gain = 250 & Pre-amp gain 5.2×;
- Parallel clock period = 500 ns;
- Parallel clock voltage = 0;
- Serial clock period = 100 ns;
- Frame rate = 300 Hz;
- Sub-frame size 23×23 pixels & Sub-frame position (100, 100).

The frame rate chosen was the slowest acquisition 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 Table 3.

Dataset	Temperature °C				
	-80	-60	-40	-20	-10
Stack RMS image	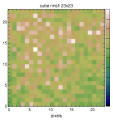	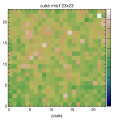	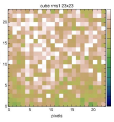	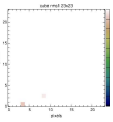	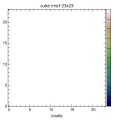
RMS noise (e^-)	71.6	66.5	86.1	152.1	249.5

Table 3: Dark-frame RMS noise in custom clocking mode, as a function of CCD temperature. The colour mapping is the same for all images and ranges from 0 to 100 electrons per pixel per frame.

Table 3 shows that there was little change in the background noise between sensor temperatures of -80°C and -60°C, but that it increased from -40°C upward. Hence

Temperature/°C	Integration Time/ms		
	1	5	10
-80	68.8	71.3	73.5
-60	64.8	66.7	70.2
-40	80.0	84.4	89.7

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

it was 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. It is noteworthy that there appears to be a "sweet spot" at -60°C where the noise contribution was lowest.

6.3.3 Effects of integration time on custom clocking

Although the FTT system is expected to normally operate with a servo bandwidth of 40Hz, and hence a relatively short integration time, a longer integration 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 -40°C, -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 = "0";
- Sub-frame size 23×23 pixels & Sub-frame position (100, 100).

The results of these tests are summarised in Table 4 and showed 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.

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.

Frame number							
1	2	3	4	5	6	7	8
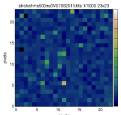	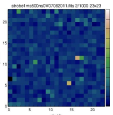	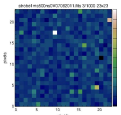	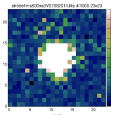	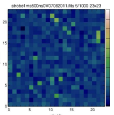	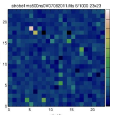	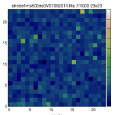	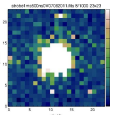

Table 5: Image sequence in which the CCD is illuminated by the light emitting diode source for one frame in every four. The un-illuminated frames are inspected for residual charge. All images here have the same colour map scale.

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 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 23×23 pixels & Sub-frame position (100, 100).

A typical sequence of frames is shown in Table 5.

Under these conditions, no evidence was found for residual charge in the un-illuminated frames.

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 additionally shielded by a cardboard enclosure. 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 30.9k photons over a total integration time of 20 seconds, and 14.3k counts of background over the same period with the source switched off[3]. With the source off and the camera shutter closed, 14.4k counts were measured. Hence $\sim$16.5k photons, or ~ 830 photons per second, could be attributed to the source.

The neutral density was then adjusted to a value of 4.1 so as to yield a measured photon rate of approximately 20800 photons per second, close to that expected for a 15.4 magnitude target (19800/second). Once this calibration point was determined, other fluxes could be tested simply by changing the neutral density filter. Given the repeatability of the data, together with the uncertainties in the optical densities of the ND filters used, we estimate that this rough calibration was accurate to approximately 20%.

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

- Neutral densities of 1.6 to 4.1, corresponding to celestial sources of magnitudes 9.1 to 15.4;
- 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 = -60°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 23×23 pixels & Sub-frame position (100, 100).

Examples of images for the different clocking parameters and neutral density levels are presented in Table 6. Centroids for each image stack were then calculated from the pixel data using the following quad cell algorithm:

[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.

Clock Sequence	Neutral Density					
	1.6	2.1	2.6	3.1	3.6	4.1
Conven., 300Hz	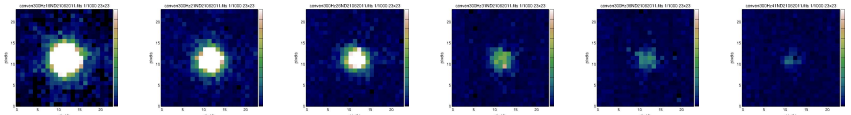					
Custom, 300Hz	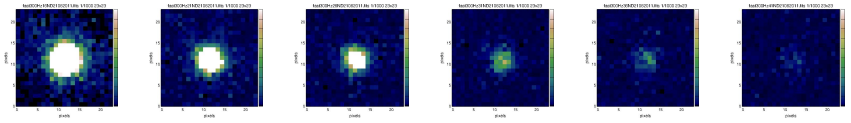					
Custom, 1000Hz	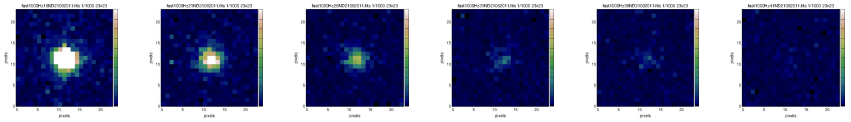					

Table 6: Examples of sub-frame images as a function of clocking scheme, frame rate and neutral density (i.e. target magnitude). ND values of 1.6 and 4.1 correspond to a target magnitudes of roughly 9.1 and 15.4 respectively at the MROI. The colour scale is constant across all images and ranges from 85 to 1000 counts.

- 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]

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 5.

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,

[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.

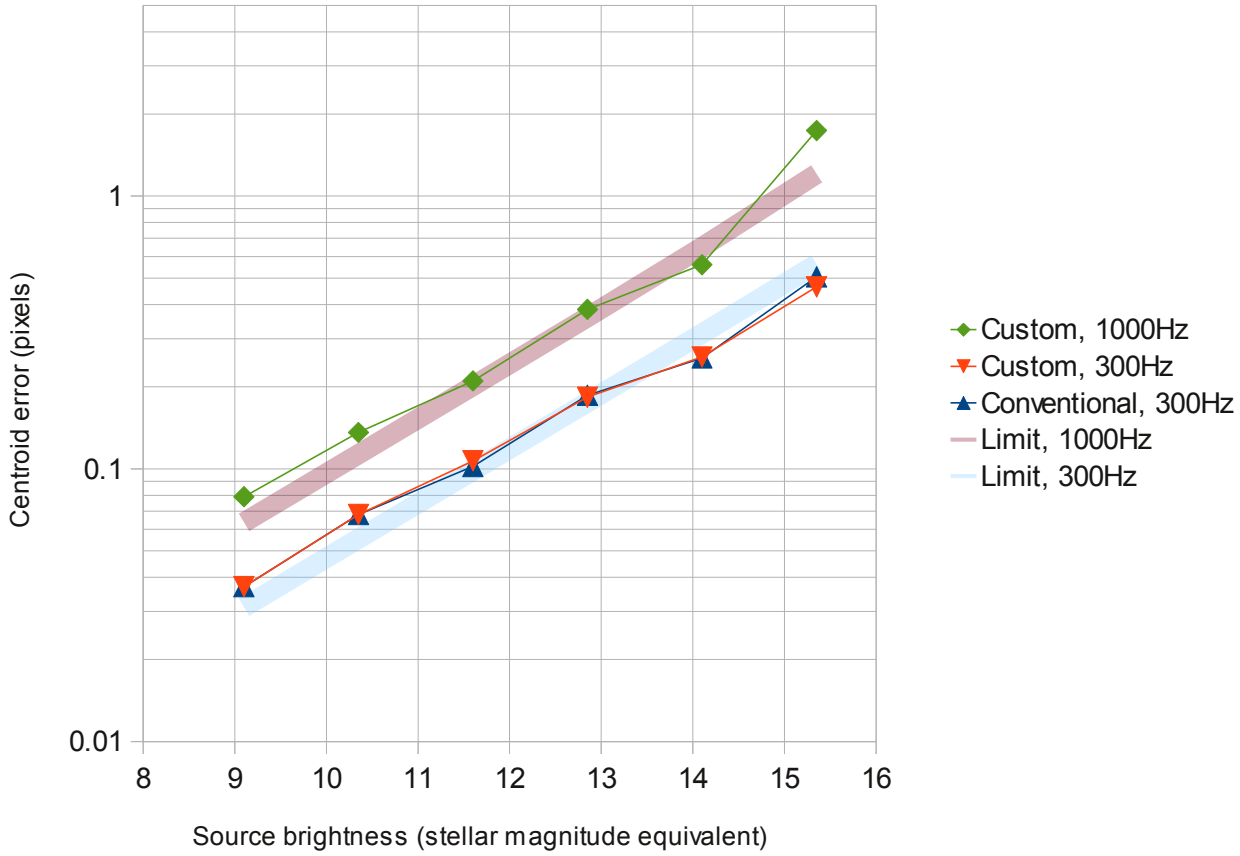

Figure 5: Two-axis error of calculated centroid, as a function of clocking scheme, source magnitude at the MROI and frame rate. The theoretical limits to centroiding performance for frames rates of 300 Hz and 1000 Hz are also shown. Due to the calibration uncertainties of our experimental set up these predictions are only accurate to the 10% level, and hence these limits are shown as shaded bands.

at least down to photon rates comparable to those expected from a magnitude 15.4 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.
- Our data suggest that at a frame rate of 1000 Hz, and using the custom clocking scheme, the RMS centroid jitter is greater than expected, indicating the possible presence of an excess noise contribution in this region of parameter space.

7 Summary

The key conclusions from these initial tests of the trial 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 convention ROI readout.
- The trial custom clocking scheme does not currently allow positioning and sizing of the sub-frame to be read out to meet our requirements. Furthermore, it leads to a non-uniform "background" in the ROI which may make precision centroiding difficult at the lowest light levels.

The first two of these conclusions are very reassuring, and suggest that we are close to demonstrating an operating mode that meets the most critical requirements of the MROI FTT system. Furthermore, the final conclusion was expected given that the trial clocking scheme had not been designed with our desired flexibility built-in.

In view of these very positive conclusions, we believe that relatively minor adjustments of the custom readout strategy will be able to realise the majority of our needs, and recommend that Andor be commissioned to deliver a modified version of their scheme, as outlined below, in the very near term.

8 Recommendations

Having undertaken an evaluation of Andor's custom clocking scheme we believe that our key technical requirements can be met with the following adjustments and enhancements of the current trial code. Our first, and most pressing, request is that the basic readout strategy be adjusted so as to limit the probability of signal buildup from regions outside the region of interest. The other requests is to allow for more convenient switching between different readout modes.

8.1 Change in basic readout strategy

The trial custom clocking scheme has been able to meet our low readout latency requirement by filling all the pixels of the ROI into the serial readout register, as though they originated from a "single row" of the CCD. Our understanding is that it is this that both allows an effective sub-frame readout rate higher than 1 kHz and constrains the ROI sub-frame size to be no greater than $23 \times 23 \quad (= 529)$ pixels in size. In addition, because multiple parallel row transfers are needed to

populate the readout register, charges from outside the ROI are added "on top of" charges already residing in this register.

As a result, we wish to request a modification to the trial custom clocking scheme such that a *fixed-size* 32×32 ROI is clocked out, but *in four steps*, each of which would correspond to a 32×8 sub-region of the ROI.

The center of the ROI would normally be located at the center of the 512×512 pixel active area, but the new clocking scheme must allow for the central location of the ROI to be varied by up to ± 68 pixels in the horizontal and vertical directions.

Our revised clocking scheme to allow readout of the ROI would comprise the following steps:

1. The exposed CCD pixels would all initially be shifted to the storage region of the chip;
2. Subsequently, all the rows prior to the first containing the 32×32 pixel ROI would be parallel clocked into the readout register;
3. The *whole* of the readout register would then be serially clocked out, with none of the data being recorded;
4. Then, the first row containing data of interest from the ROI would be parallel clocked into the readout register. The register would then be clocked along by 32 pixels, after which the next row of the ROI would be parallel clocked into the readout register. This would be repeated until the values in the first lower quarter (i.e. 32×8 pixel region) of the ROI had been transfered into the readout register;
5. Thereafter, pixels from the readout register would be read out until all the pixels from the bottom quarter of the ROI had been digitised. Note that this would correspond to fewer serial clocks than in step [3] where every pixel in the readout register is assumed to have been clocked out;
6. The next three quarters of the ROI would then be readout by repeating steps [4] and [5] of the procedure three times;

This would be slower than the current trial custom readout, but we believe it will meet all our requirements. In particular, for parallel and serial clock periods of 500 ns and 100 ns respectively, we estimate that the 32×32 ROI could still be read out and digitised in around a millisecond.

8.2 Mode switching

With the current trial implementation, in order to be able to switch between the "conventional ROI" and "custom" readout modes it is necessary to exit the program and start again, reading in the appropriate configuration file at startup. This

restarting demands that the camera's Peltier cooler switch off and on and hence it can take several minutes for the CCD temperature to stabilise.

It is essential that we be able to switch between conventional full-frame and custom ROI clocking on a scale of seconds, without a need to power off the Peltier cooler.

8.3 Proposal for way forward

Finally, in order to facilitate the expedient delivery of a suitable custom clocking scheme, we suggest that this report be forwarded to our contacts at Andor, and that a teleconference be set up as a mutually convenient time, as soon as possible.

References

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

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