Implementation of an FPGA-based DCDS video processor for CCD imaging

Simon Tulloch*

European Southern Observatory, Karl Schwarzschild Straße 2, Garching, Munich 85748, Germany

ABSTRACT

Noise modeling of an E2V CCD231 suggested that a weighted double correlated sampler (DCDS) processor could offer small noise improvements at low pixel rates. The model was used to produce synthetic video waveforms that were then processed at various ADC frequencies and analogue bandwidths to identify the best weighting strategy and preamplifier design. An FPGA-based DCDS controller was then built, first to measure the actual CCD noise spectrum and then to verify the earlier theoretical results.

Keywords: DCDS, CCD, FPGA, Video Processor

1. INTRODUCTION

The video processor in a CCD controller has traditionally been implemented as an analogue circuit. Its role is to extract the photo-induced signal from the video waveform of the CCD. This waveform (shown in Figure 1) approximates to a square-wave whose amplitude is proportional to the current pixel´s photo-charge. The amplitude of the square wave is measured by first sampling and storing the voltage of the "reference pedestal" then repeating the measurement during the "signal pedestal" a short time later. The processor then performs a subtraction of the two sampled pedestals and outputs a voltage proportional to the pixel charge. This is then converted into units of ADU using an analogue to digital converter (ADC). Finally the output node of the CCD is reset to a known value ready for the next pixel. There are certain regions of the video waveform that are extremely noisy since they correspond to periods when the pixel clocking is active and there is a great deal of noise feed-through. Most of this takes place during the first few 100ns of each pixel. It is important to avoid this region when sampling the reset pedestal. A further clock transition occurs at charge dump so a small delay needs to be added just prior to sampling the signal pedestal.
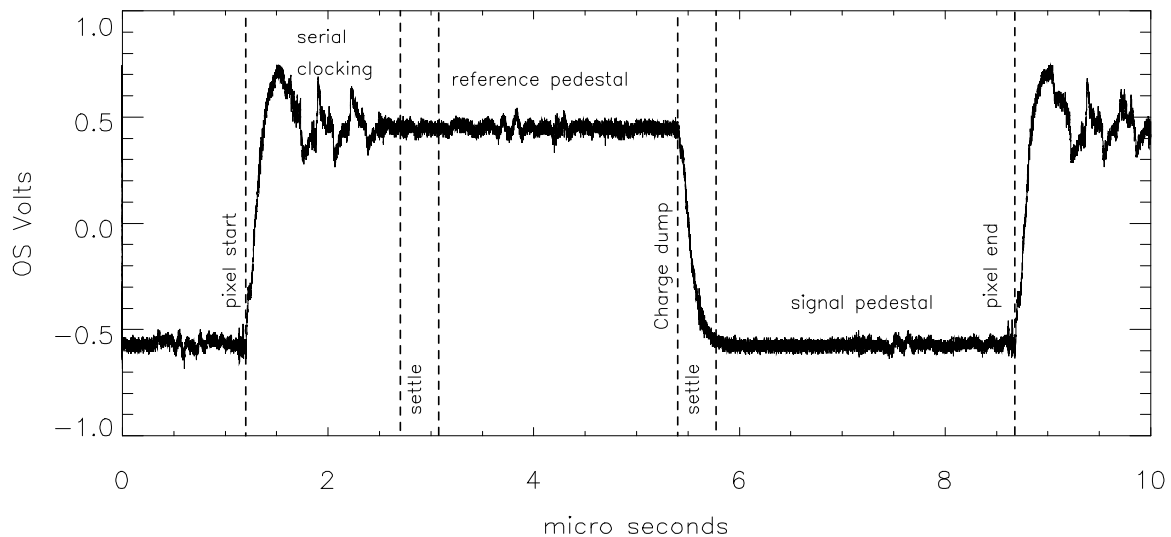

Figure 1. The video waveform from a CCD231 covering a period of just over one pixel. The reset event, where the waveform shows a positive step, separates one pixel from the next. The depth of the negative going transition is proportional to the photo-charge present in the pixel and it is this step-size that we must measure with our CDS.

*stulloch@eso.org; phone Tel. +49-89-32006170   Fax. +49-89-3200683

At first look one may wonder why this two-sample technique is required if the video waveform always returns to a known value when reset at the end of each pixel. If high speed is of utmost importance then indeed a single measure of the video waveform after the charge dump will suffice, however, it will give a noisy result due to the effects of "reset noise". Each time the output measurement node is reset it will return to a slightly different value due to thermal noise in the reset transistor. This uncertainty can amount to several tens of electrons which can be disastrous in scientific CCD cameras where we would typically hope to achieve <3e⁻ noise. The dual measurement technique completely removes the reset noise and is known as Correlated Double Sampling (CDS).

Building an analogue video processor requires a lot of circuitry and is a complex design task. Newer high-performance ADCs can now offer the chance to process the video signal almost entirely in the digital domain (a low noise analogue preamplifier is of course still needed). Such processors are known as Digital Correlated Double Samplers (DCDS).

With a DCDS we can directly digitize the preamplifier output to produce a stream of digital values. The conversion rate is high enough so that we can take multiple samples in the pedestal regions, average them together in software and then calculate their difference to yield a pixel value. With the freedom that the software implementation gives us we can also experiment with novel schemes whereby we give different weighting to the individual ADC samples in effect producing a digital filter that can be tuned to optimize our signal to noise ratio.

The theoretical noise output of any CDS or DCDS scheme can be calculated if its transfer function H(f) and the noise spectrum N(f) of the CCD are known. H(f) is essentially a filter function describing the gain of the CDS system as a function of frequency. If we multiply H(f) and N(f) and integrate over the bandwidth of our system, as shown in Equation 1, we arrive at the RMS noise output of our video processor, assuming of course that the processor does not introduce any extra noise of its own.

$$\text{Noise} = \left[\int_0^{fmax} [H(f).N(f)]^2 df \right]^{0.5} \text{Volts RMS} \quad (1)$$

The process is illustrated in Figure 2. which shows a typical CCD noise spectrum with its characteristic low-frequency excess due to "flicker noise" and H(f) for a differential averager CDS scheme. For an interesting discussion of flicker noise see Janesick[1].
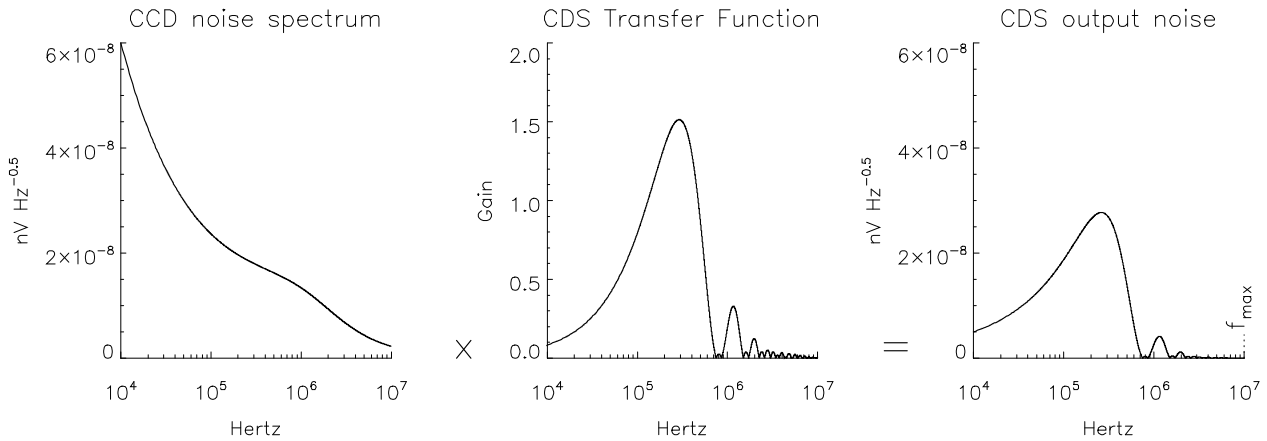
Figure 2. Calculating the read noise in the frequency domain. The noise spectrum N(f) of the CCD is multiplied by the transfer function H(f) of the CDS. The product is then integrated (see Equation 1) over the bandwidth of the system to yield the RMS read noise. An alternative time domain method is described in Section 2.

So, if we wish to model the performance of potential DCDS schemes we need to know the noise spectrum of our CCD and we also need to know our filter function H(f). Derivations of H(f) can be found in the literature[2,3] for the usual

schemes of "differential averager" (DA) and "clamp and sample" (CS). With DCDS we can now propose other more exotic sampling schemes for which H(f) cannot easily be derived. It was decided instead to take a different tack and do performance modeling in the time rather than the frequency domain. This required the generation of a synthetic CCD noise waveform N(t) using a Fourier technique that is described in Section 2.1. This waveform could then be multiplied by any CDS weighting function H(t) and the result averaged over a pixel time to yield a pixel voltage. By generating tens of thousands of "synthetic pixels" in this manner, the RMS spread in values could be calculated giving the RMS CCD read noise in Volts. This time domain technique was very attractive since once the initial problem of generating the synthetic noise waveform was solved it was then straightforward to experiment with any number of DCDS weighting schemes without first having to derive their frequency domain transfer functions.

The purpose of this study was then to first measure the noise spectrum of an E2V CCD231 and then model how low the read noise could be pushed using DCDS. Since the hardware required to measure the spectrum amounted in effect to a small CCD controller, the final stage of the study was to implement a DCDS algorithm in VHDL-based hardware and compare the predictions of the modeling with some actual real-world results.

2. THE TIME-DOMAIN NOISE MODEL

2.1 Generating a synthetic CCD noise waveform

In order to get good statistics it was necessary to generate noise waveforms over intervals of tens of thousands of pixels. The spectrum of this waveform once passed through an FFT analysis should then correspond exactly to the noise spectrum of our CCD. Since the FFT works in both directions a Fourier technique was suggested for the generation of the waveform. The technique used was as follows. An array of complex numbers was defined in IDL to describe a function in the frequency domain. Each array element had a corresponding amplitude and phase. The amplitude as a function of frequency (effectively the array index) was given by the CCD noise spectrum shown in Equation 3. The phase of each array element was simply a randomly generated number between 0 and 2π. The complex array was then passed through an FFT and appropriately scaled to yield a noise waveform. Input parameters to this IDL routine were the sample frequency, the number of samples required and the analogue bandwidth of the system (upper and lower 3dB points). An example output is shown in Figure 3. Note that the central panel shows only sparsely plotted phase values for the sake of clarity.
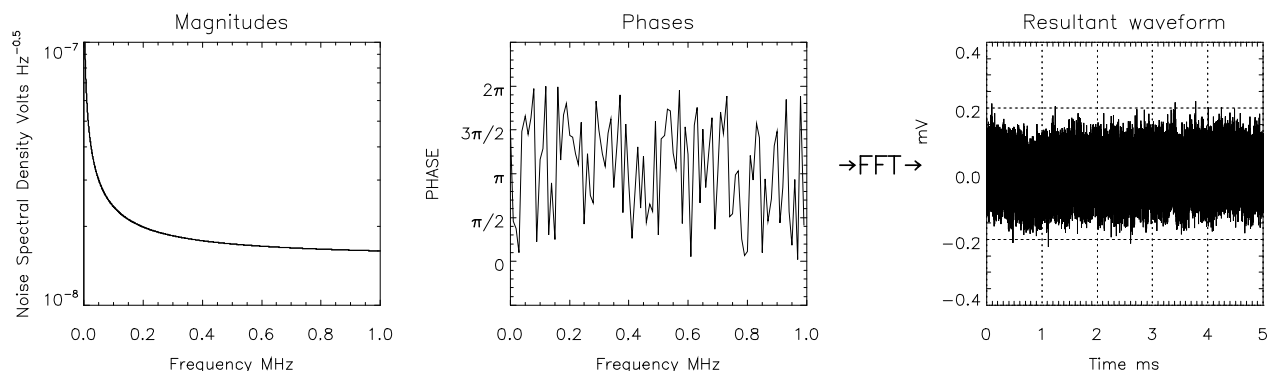

Figure 3. The Fourier technique used to generate synthetic CCD noise waveforms.

2.2 Predicting noise performance with the model

To synthesise a faithful CCD video waveform we also need to add a photo-signal to the noise waveform we have generated. This photo-signal takes the form of a constant amplitude square wave. We can then simultaneously keep track of how the DCDS gain as well as its noise varies as we change various readout parameters. We of course need to know the gain if we want to reference the calculated noise back to the input node where we can express it in units of photo-electrons RMS. To measure the read noise we must now take our noise waveform N(t) and multiply it by our CDS

function H(t). The product is then averaged over the pixel time to yield a pixel value in Volts. This is shown in Figures 4 and 5 for the two commonest CDS schemes. In both cases the non-useful parts of the waveform (labeled "settle" and "serial clocking" in Figure 1) were set to 5% and 20% of the pixel period respectively. This made the noise calculation results more directly comparable with the CCD manufacturers' data sheet since this assumed that the time between charge dump and the end of the pixel was equal to 40% of the pixel period. The CDS function is set to zero in these non-useful parts. The function is then stepped along the noise waveform, pixel at a time, building up a results array of pixel values. The RMS value of this array then gives the read noise referenced to the output of the CDS.
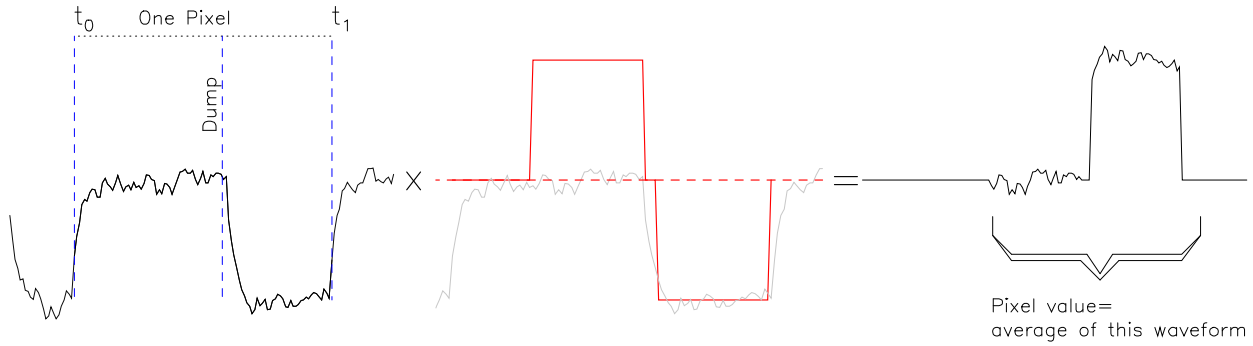

Figure 4. Calculating noise in the time domain. This was the preferred method in this study and is an alternative to the frequency domain method shown in Figure 2. The synthetic CCD waveform is multiplied by the differential averager weighting profile. The resultant product is then averaged over the pixel time to yield the pixel value. By repeating the process for many pixels we can then find the RMS read noise.
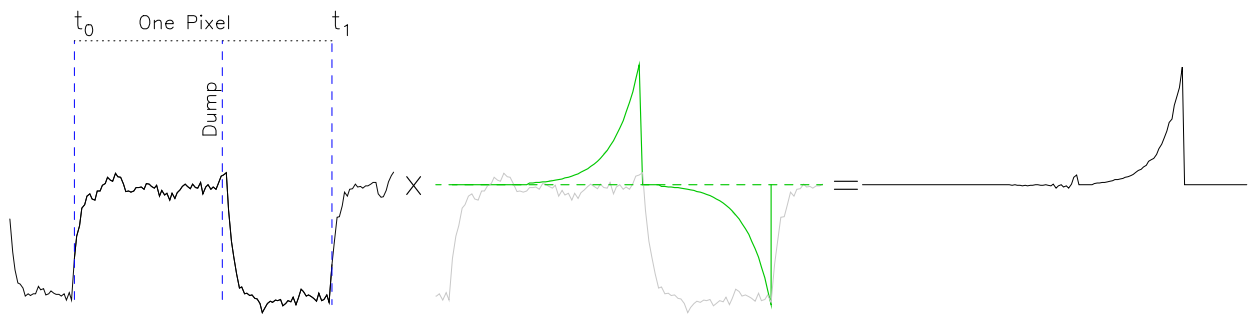

Figure 5. The same process as that shown in Figure 4. but for a Clamp and Sample CDS scheme (which essentially uses one measure of the reference and one measure of the signal pedestal, although in combination with a low-pass prefilter). Pre-filter bandwidth has been set equal to twice the pixel frequency.

2.3 Comparison with frequency domain analytic result

An equation describing H(f) the transfer function of a differential averager was available in the literature[2,3] so it was possible to perform a direct comparison of the time and frequency domain results to check their agreement. Figure 6 shows the result for a wide range of pixel frequencies using the noise spectrum shown in Equation 4. The plot also shows the read noise performance as a function of pixel rate taken from the E2V product data sheet. Their quoted measurement

conditions were that clamp and sample processing was used, the analogue bandwidth set to twice the pixel frequency and the time between charge dump and end of pixel equal to 40% of the pixel period. It is interesting that the CCD231 should be capable of lower noise than the data sheet suggests. Some margin has probably been included to account for additional noise originating in the CCD controller itself.
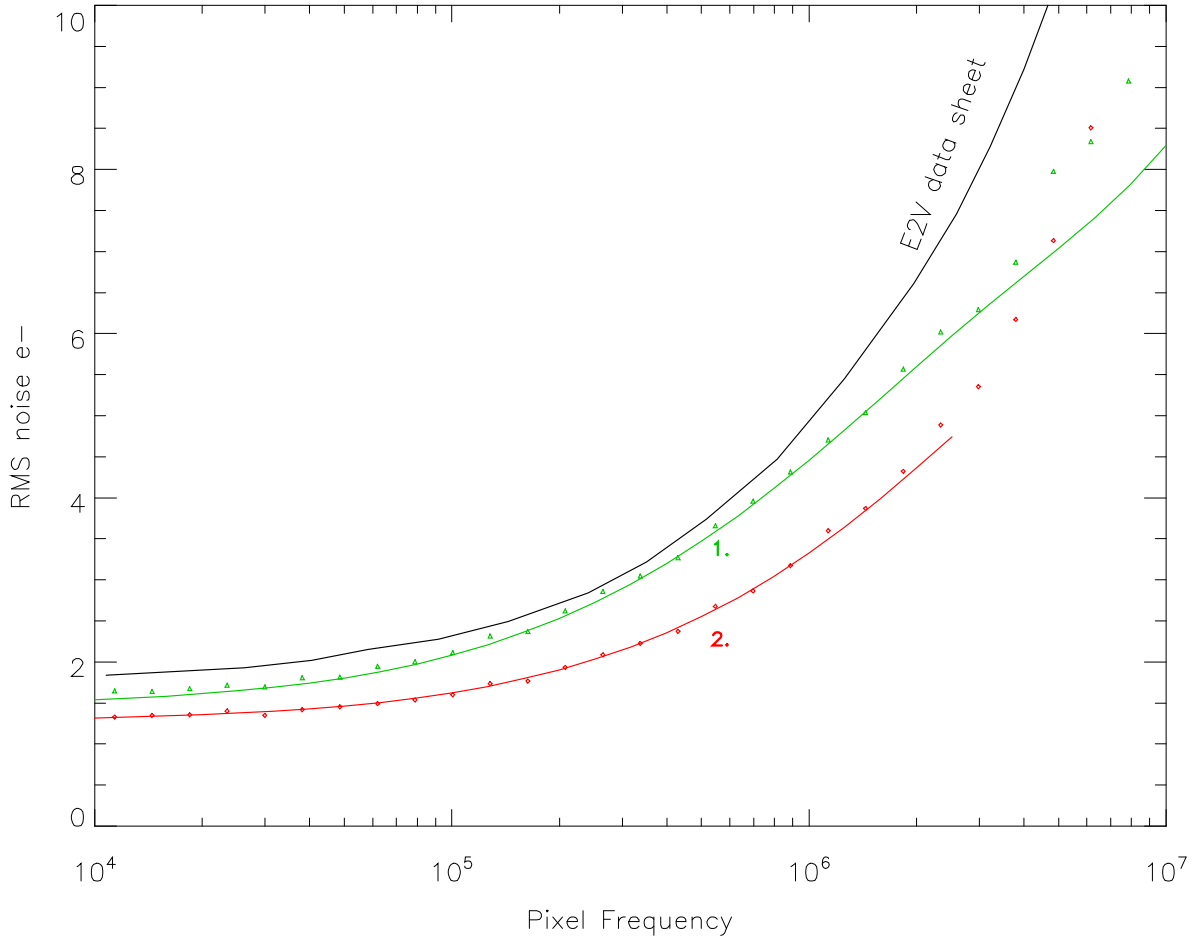

Figure 6. A check that the time domain technique result (plotted points) agreed with the frequency domain analytic solution (solid lines). Plot 1 shows the result for the Clamp and Sample scheme. Plot 2 shows the Differential Averager scheme. The E2V CCD231 data sheet values are also plotted.

3.   ANALOGUE BANDWIDTH REQUIREMENT

A DCDS system still requires an analogue preamplifier and its performance is crucial. Clearly it must not contribute additional noise but there are other things to consider. The video bandwidth is a critical parameter. The CCD output transistor will itself limit the bandwidth to about 8MHz in the case of the CCD231 but more generally we are limited by the bandwidth imposed by the cabling between camera and controller. The exception to this will be if preamplifiers are included within the camera cryostat. The video waveform captured in Figure 1 was obtained with a cable run of 50cm between CCD and controller. Examination of the waveform flanks indicates the bandwidth has been reduced to 2MHz. An insufficiently high upper-3dB point can affect both noise and image resolution. If the video waveform has insufficient time to settle between one pixel and the next then some of the signal in a pixel can *leak* downstream. The leakage has a negative sense which can cause a bright pixel to be followed by a dark pixel whose value may actually be

below bias. This can be very apparent in Electron Multiplying CCD images where individual photo-electron events appear as delta functions and easily reveal any shortcomings in the high-frequency response. The video waveform will also have its amplitude suppressed. This will reduce both the signal and the noise by the same factor, however, it will make the system more susceptible to extraneous noise or noise generated in the controller itself. A general investigation of pixel leakage was done to see just how high the upper-3dB point of the video bandwidth needs to be set to ensure that, in the case where we have a saturated pixel (65kADU), that the next pixel in the sequence does not receive more than 1ADU of leaked signal. The results are shown in Figure 7. What this shows is that we need to set the video bandwidth upper-3dB point to 6 x f_{pix} if we use differential averager (DA) CDS or 2.6 x f_{pix} if we use clamp and sample (CS). Note that DA is much more susceptible to leakage given that it places full weighting on signal samples occurring close to the start of the pixel period. The CS method experiences less leakage since it only starts to sample the waveform (with the clamp event) more than halfway through the pixel period, after which the perturbations from the earlier bright signal have had plenty of time to settle.
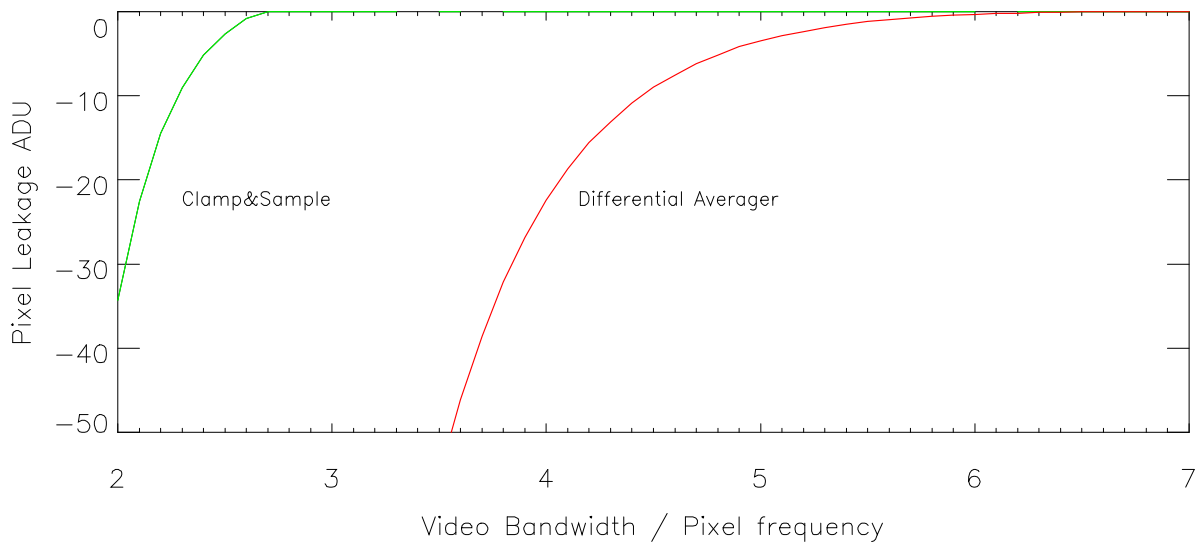

Figure 7. The analogue bandwidth required to give good image quality. The graph shows downstream signal leakage from a saturated pixel into its neighbour as a function of analogue bandwidth. Shown for the two most common CDS schemes. Note that this result applies equally to both DCDS and analogue CDS processors.

We can also do a simpler hand-waving argument using the equation describing an RC filter and ask what bandwidth is required to get a square-wave signal to settle, to say 10% within 5% of its period. The result is similar: the bandwidth needs to be >7.3x the square wave period.

4. ADC SPEED REQUIREMENTS

One fundamental parameter of any DCDS scheme is the number of ADC samples that need to be taken in each of the pedestals. This then predicts how fast an ADC we need to use for our desired pixel rate. The time domain model allows us to explore this easily. Figure 8 shows three pixel waveforms overlaid with various differential averager schemes. The first shows the transfer function H(t) of an analogue processor, the other two show the comb-like appearance of equivalent digital schemes with 15 and 8 sample-pairs per pixel. The ADC frequency should be at least twice the analogue bandwidth of a DCDS system if the Nyquist sampling criterion is to be met. One would expect the higher ADC frequencies to improve the performance but at some point there must be diminishing returns from adding extra samples to the measurement. It has already been shown that the analogue bandwidth should be at least 6- 7 x f_{pix} so this would suggest that our ADC needs to be converting at about 12-14 x f_{pix}. The question now is if going even faster offers any advantages. This was investigated using the time domain model. The modeled pixel streams were first converted from

floating point to integer using a conversion factor of $1e^-=1ADU$. A single synthetic pixel contained as many as a thousand samples in the model. To investigate various ADC frequencies it was necessary to skip over some of these values when performing the CDS processing: the higher the skip factor the lower the effective ADC frequency. The read noise was then determined as a function of the pixel rate and the coarseness of these combs (i.e. as a function of the ADC frequency relative to the pixel frequency). This was done with the analogue bandwidth of the pre-filter equal to 10x the pixel frequency, comfortably above the minimum requirement. The results are shown in Figure 9. As one would expect the noise did go up if the ADC frequency was less than 20x the pixel rate since the Nyquist criterion was not satisfied. For higher ADC rates there was a small advantage but above 50x the pixel rate (i.e. 5x the Nyquist frequency) no further gains in performance were apparent.
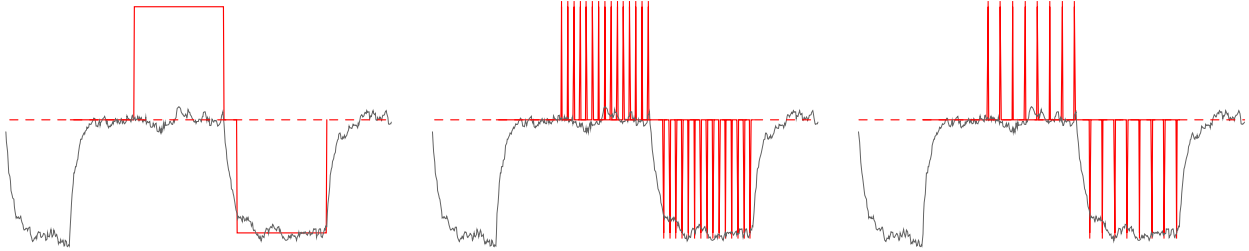

Figure 8. Synthetic pixel waveforms overlaid with first the transfer function H(t) of an analogue DA processor and then by the transfer functions of two equivalent DCDS schemes at two ADC frequencies.
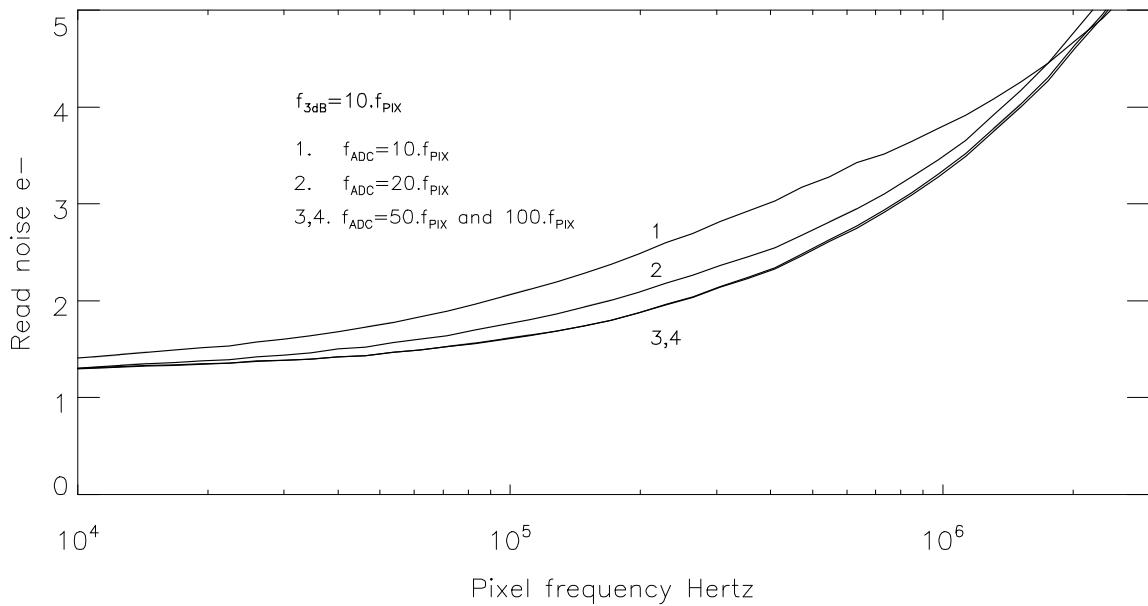

Figure 9. Effect of ADC frequency on the noise performance of a differential averager DCDS scheme. Analogue bandwidth equal to 10x the pixel frequency. Results shown for 4 different ADC frequencies expressed as multiples of the pixel rate.

5. MODELING USE OF WEIGHTED ADC SAMPLES

Now that the model had been verified in its accuracy it could then be applied to various novel weighting schemes for which analytic forms of their frequency response H(f) could not be found in the literature. One promising scheme, first described by Gach[4], involves giving extra weight to samples close to the charge dump event. This would seem sensible since the scheme is less sensitive to low frequencies in the noise waveform (where the majority of the flicker noise is found) whilst at the same time faithfully measuring the height of the charge-dump step.

5.1 Mirrored exponential weighting scheme

This scheme uses weights with an exponential profile mirrored about the time of the charge dump. It is therefore referred to here as the "mirrored exponential" scheme. It is described graphically in Figure 10. The parameter Z defines the sharpness of the profile where the weights in the reference pedestal are given by:

$$H(t) = e^{\left[\frac{2.\pi.Z\left(t-t_{dump}\right)}{T_{pix}}\right]} \quad , \quad (2)$$

and where T_{pix} is the pixel time. The weights in the signal pedestal are identical but reflected along the time axis and of opposite sign. A similar scheme is described by Clapp[5].
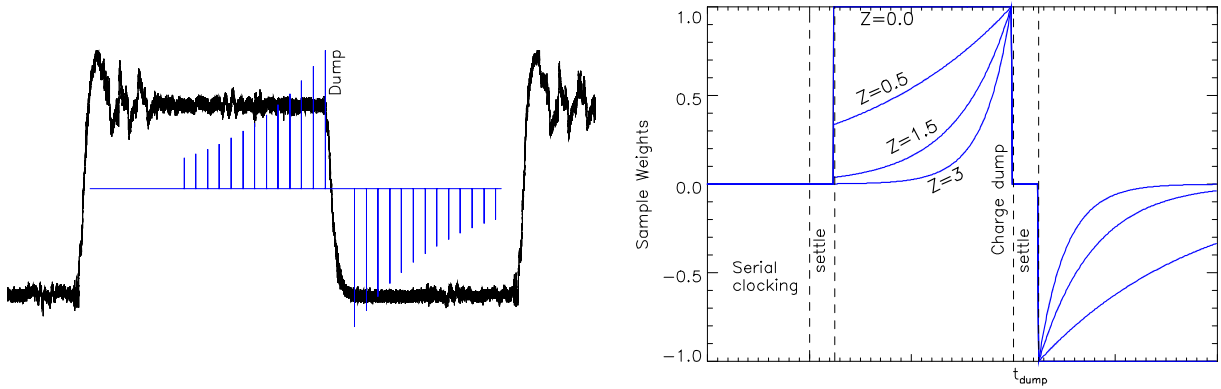

Figure 10. The appearance of the mirrored exponential DCDS weights overlaid on the pixel waveform. The narrowness of the weight profile is determined by the parameter Z (see equation 2). As Z increases, the samples closer to the charge dump are given increased relative weighting. As Z approaches 0, the scheme effectively becomes a simple differential averager.

The result is shown in Figure 11. It did offer a small reduction in read noise at very low pixel rates although not in any significant way. A reduction of 5% compared to the differential averager was found. At the higher pixel rates typically used in scientific CCD cameras it actually performed worse. Clapp has found a similar result.
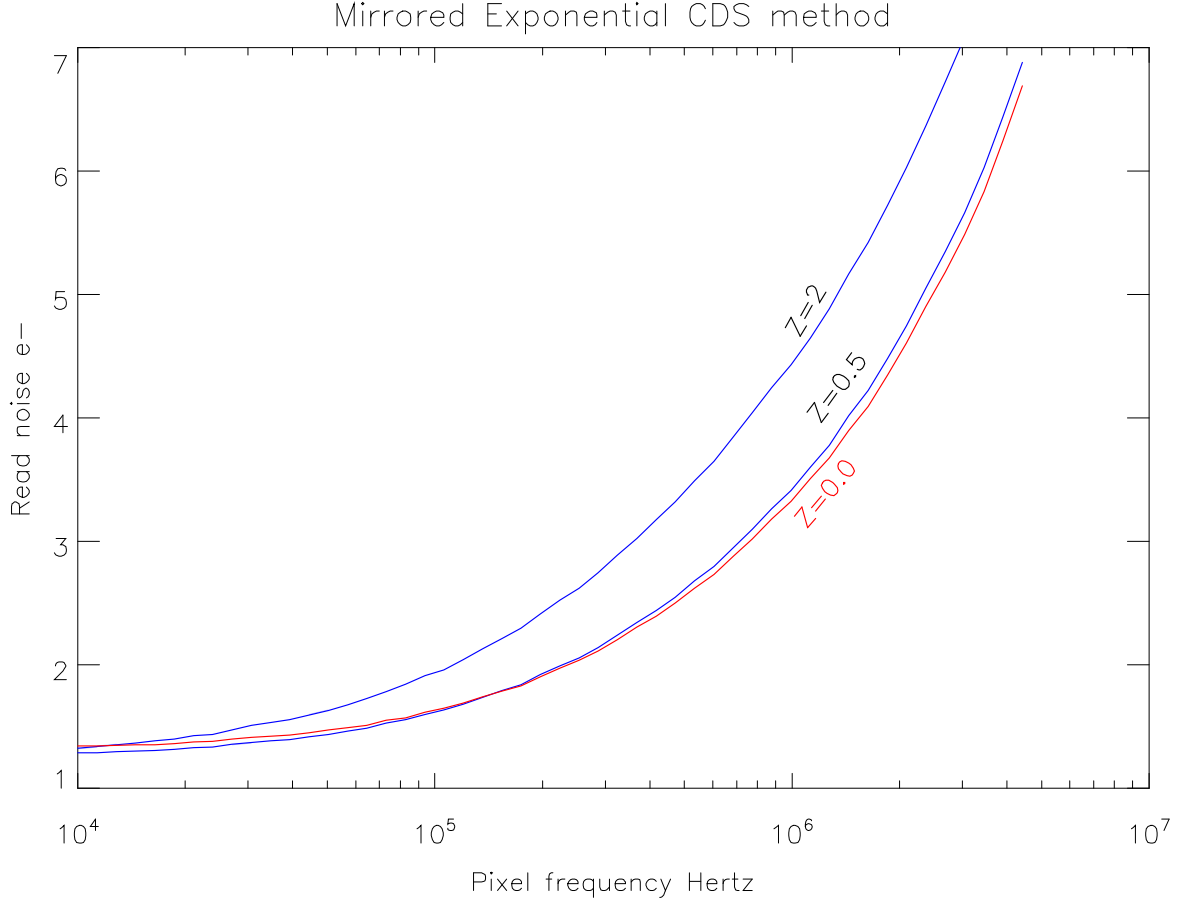
Figure 11. The noise performance of the "Mirrored Exponential" DCDS scheme with various values of the Z parameter. The case where Z=0 corresponds to the differential averager scheme. With Z=0.5 a 5% reduction in read noise was obtained at the lowest pixel rates.

5.2 Flicker Noise suppression filters

Stefanov[6] and Alessandri[7] have proposed optimised filters that maximise the signal whilst minimising the noise at each frequency. Their models show small performance improvements at very low pixel rates where flicker noise dominates. An experimental filter function $H(t)_{FS}$ (FS="flicker suppression") was developed to try and reproduce their results using the time domain modeling technique. The starting point for the calculation was the differential averager frequency response $H(f)_{DA}$ from the literature[1,2] and the CCD231 noise spectrum $N(f)$ shown in Equation 4. The normalised reciprocal of the CCD noise spectrum was then calculated as follows:

$$M(f) = \frac{1}{\left(N(f).\max\left(\frac{1}{N(f)}\right)\right)} \quad (3).$$

This function was then used to modify the differential averager frequency response so as to suppress its gain in the flicker noise region. The resultant function $M(f) \times H(f)_{DA}$ was then Fourier transformed to yield the new DCDS weighting $H(t)_{FS}$. One problem with this method was that it gave small but non-zero weights both within the serial clocking regions of the pixel and also in the neighbouring pixels. These weights were simply set to zero. The result is shown in Figure 12 and is superficially very similar to the weighting function used by Stefanov and Alessandri, although

not derived with the same mathematical rigour. Once run through the time domain model, a 5% reduction in read noise was obtained but only at pixel rates below 50kpix s^{-1}.
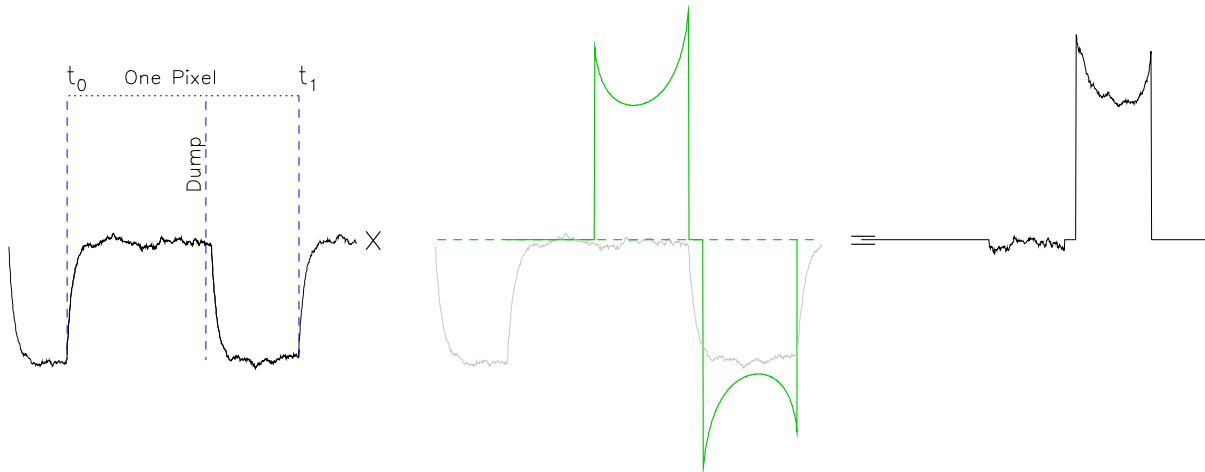

Figure 12. Experimental flicker-noise suppression weighting scheme. Shown for a pixel rate of 50kHz.

5.3 Critique of more exotic noise suppression schemes

Other researchers[8] have proposed DCDS schemes that can remove entirely the 1/f noise and achieve read noise well below 1e$^-$. Their technique involves digitizing long sequences of pixels and then disentangling the underlying low frequency 1/f noise from the photo-induced part of the video signal. The calculated 1/f noise is then subtracted off and conventional DCDS applied to the resultant waveform. In effect each calculated pixel value requires measurements from adjacent pixels to be included in the calculation. A technique that removes 1/f noise in any measurement system would be quite revolutionary but the evidence they offer for its efficacy is quite slim given the magnitude of their claim. One basic objection to this technique is that any underlying 1/f noise will have its frequency spectrum folded about the pixel frequency, effectively scrambling any knowledge of its original form. Also, any scheme that calculates the current pixel value based on data from others is bound to degrade the point spread function of the system. Any reduction in noise is then likely to be due to effective box-car averaging of a pixel´s read noise with that of its near neighbours.

6. ACTUAL MEASUREMENT OF CCD231 NOISE SPECTRUM

Dedicated spectrum analyser hardware was built to measure the noise of a thinned CCD231 that was kindly provided by Paul Jorden at E2V. The measurement required supplying all the usual bias voltages (OD, RD, SS, OG etc..) to the CCD and also the manipulation of the CCD clock phases: in effect most of the functions of a simple CCD controller. To minimize the effect of extraneous noise it was decided to mount the CCD directly onto the spectrum analyser PCB. It was then possible to place the video preamplifier right next to the video output (OS) pin. The main disadvantage of this scheme was that the device had to be operated warm although cooling to about 260K was possible by sealing the circuit board in a die-cast aluminium box and placing the whole assembly in a domestic freezer.

6.1 Spectrum analyser hardware

This was implemented on a double-sided PCB measuring 10x10cm and is shown in Figure 13. All of the CCD clocks are driven using EL7457 pin-driver ICs. One quadrant of the CCD is read out through a video pre-amplifier and digitised by an ADS807 12-bit 53MHz ADC. The charge from the other quadrants is either clocked out through the corresponding reset transistor or, in the case of the upper half of the chip, transferred directly into the Dump Drain. This is to avoid spill-over of dark current into the active video channel under measurement. The video amplifier is stable up to gains of at least x266. This was required to accurately sample the very low amplitude noise waveform of the CCD, although for later imaging tests the gain was reduced by resistor value changes. The video amplifier consists of an AD829 front-end followed by a chain of two ADA4932 differential video amps in LFCSP packages with 0603 passives. Ground plane

windows and the small component sizes minimize parasitic capacitances and ensure stability and flat frequency response across the pass-band. The PCB also contains switch mode regulator ICs to generate the various bus voltages required by the system and a mezzanine socket for an FPGA development board on which the sequencer logic was implemented. This consists of an Opal Kelly XEM3005 Spartan-3E-1200 FPGA module. CCD video output waveforms were captured by the ADC and transmitted, via the FPGA over a USB2 link, at up to 12.5MByte/s, to a PC where the spectral analysis then took place. Switch mode noise was an issue, however, it was generally extremely narrow band and at frequencies around 1.5 and 2.2 MHz (and their harmonics) which was considerably above the pixel frequency. One switcher IC used to generate 1.2V for the FPGA had to be replaced by a linear version since the combination of high step-down ratio and low demand current caused it to skip pulses in a random fashion generating broad-spectrum low-frequency noise. The load to the OS pin is link selectable between a 5k1 resistor and a 2.2mA constant-current sink. This option was included to see if the constant current load gave any reduction in read noise. It did not, although it did cause a 10% increase in amplifier sensitivity.
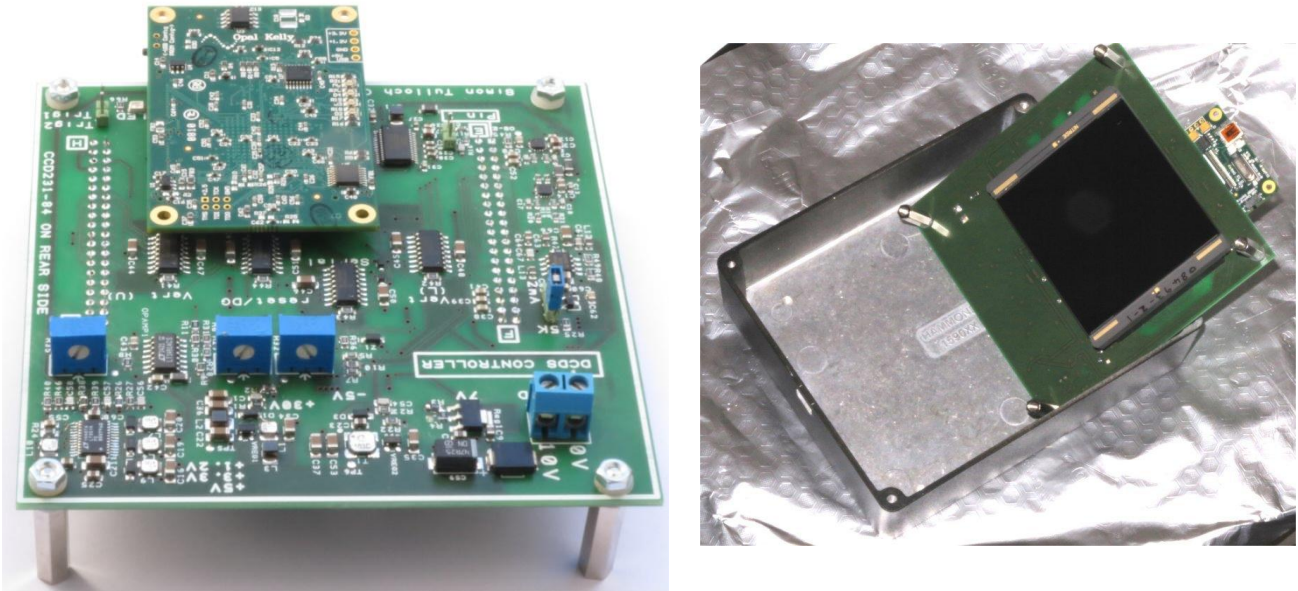

Figure 13. The spectrum analyser hardware showing the CCD231 socketed directly onto its lower surface. The FPGA development module is shown mounted mezzanine-style. The board contains a power supply, clock drivers, bias supplies, a 53MHz ADC and a 3.8MHz single-channel video preamplifier stable up to gains of x266.

There was some concern that aliased switcher noise could be contaminating the spectra. As a precaution the sampling frequency was doubled to 25MSPS when measuring the CCD231 spectrum. This exceeded the USB2 bandwidth and it was then necessary to buffer the data in a 32kB FIFO prior to transmission.

6.2 Calibration of gain, frequency response and noise spectral density

The video amplifier was characterized in terms of frequency response and base noise. Sine waves from a calibrated waveform generator were injected at the input and the RMS output measured using an oscilloscope. RMS measurements were used rather than peak-peak since it was less susceptible to high frequency noise in the system and gave a more accurate measure of the gain. A further end-to-end test of the system gain at a single frequency was done by using the on-board ADC in place of the oscilloscope. The ADC output was then analysed on a PC and the ADC sensitivity in terms of µV/ADU confirmed to be equal to the data sheet value. The high frequency roll-off of the preamp was confirmed to be 3rd order with a 3dB point at 3.8MHz.

The spectral analysis software running on the PC was implemented using the Python/Numpy FFT function. This software took a digitised noise voltage waveform at its input and output a spectrum calibrated in units of Volts Hertz$^{-0.5}$. There was some ambiguity in the documentation regarding the y-axis scaling of an FFT so the software was tested using

synthetic input waveforms consisting of bandwidth-limited white noise of known RMS amplitude (and therefore of known voltage spectral density). A final end-to-end check of the x-axis scaling was done by capturing a real square wave and confirming that the harmonics in the spectrum analyser result were at the correct positions.

The base noise of the video amplifier was tested by connecting its input to ground via a 390Ω resistor to simulate the source impedance of the CCD. This noise was found to be fairly white with a spectral density of 4nV $Hz^{-0.5}$.

6.3 CCD231 noise spectrum result

Equation 4. shows a noise model for the CCD231 provided by E2V. It consists of a white noise component with a voltage spectral density N_G of 15nV $Hz^{-0.5}$ and a "1/f" (alternative name for flicker noise) component. These are equal at f_c the corner frequency, which in this case is 150kHz.

$$N(f)_1 = 15nV . \left(1 + \frac{150kHz}{f}\right)^{0.5} VHz^{-0.5} \quad (4)$$

Measuring the actual noise spectrum required careful manipulation of the CCD clocks. The CCD231 output amplifier consists of two AC-coupled source follower MOSFET amplifiers. The bias point of the second amplifier needs to be regularly reset in order to compensate for the droop on its input coupling capacitor. This is done by pulsing the V4 phase of the CCD high every few tens of milliseconds. The noise measurement cycle therefore consisted of first resetting the output node of the CCD by pulsing ΦR, pulsing V4 high to restore the bias point, waiting a few ms to settle and then digitizing a burst of noise-waveform samples for a period of 1.3 ms at 12.5MHz. The digitised waveform was then transferred to the PC where the spectral analysis was performed. The process was repeated for up to half an hour and the captured spectra averaged to beat down the measurement noise. Between each cycle the CCD was flushed of dark current for 2 seconds. The result is shown in Figure 14. The form was quite different to that communicated by E2V and is approximated by Equation 5 below.

$$N(f)_2 = 15nV . \left(1 + \frac{600kHz}{f}\right)^{0.45} VHz^{-0.5} \quad (5)$$

The reason for the difference in corner frequency (600kHz instead of 150kHz) is not clear but could be due to the elevated temperatures during measurement (this CCD is designed to work cryogenically). It could also be that the CCD bias voltages OD and RD need additional optimisation although the fact that the output sensitivity was later measured at 7.4 µV/e- using a 2.2mA constant current load (close to the data sheet value of 8 µV/e-) suggests that it was already well optimised.

7. IMPLEMENTATION OF A DCDS PROCESSOR

The next stage of the study involved obtaining some real-world results. The FPGA used to sequence the clocks and control the ADC had plenty of spare capacity in which to implement a DCDS processor.

7.1 A DCDS processor in VHDL

Clock sequencing was done using a PicoBlaze soft processor running at 100MHz. Additional dedicated VHDL-defined hardware was then designed in the form of peripherals to this processor to capture and process the ADC samples and calculate pixel values. Figure 15 shows the internal schematics of the FPGA.
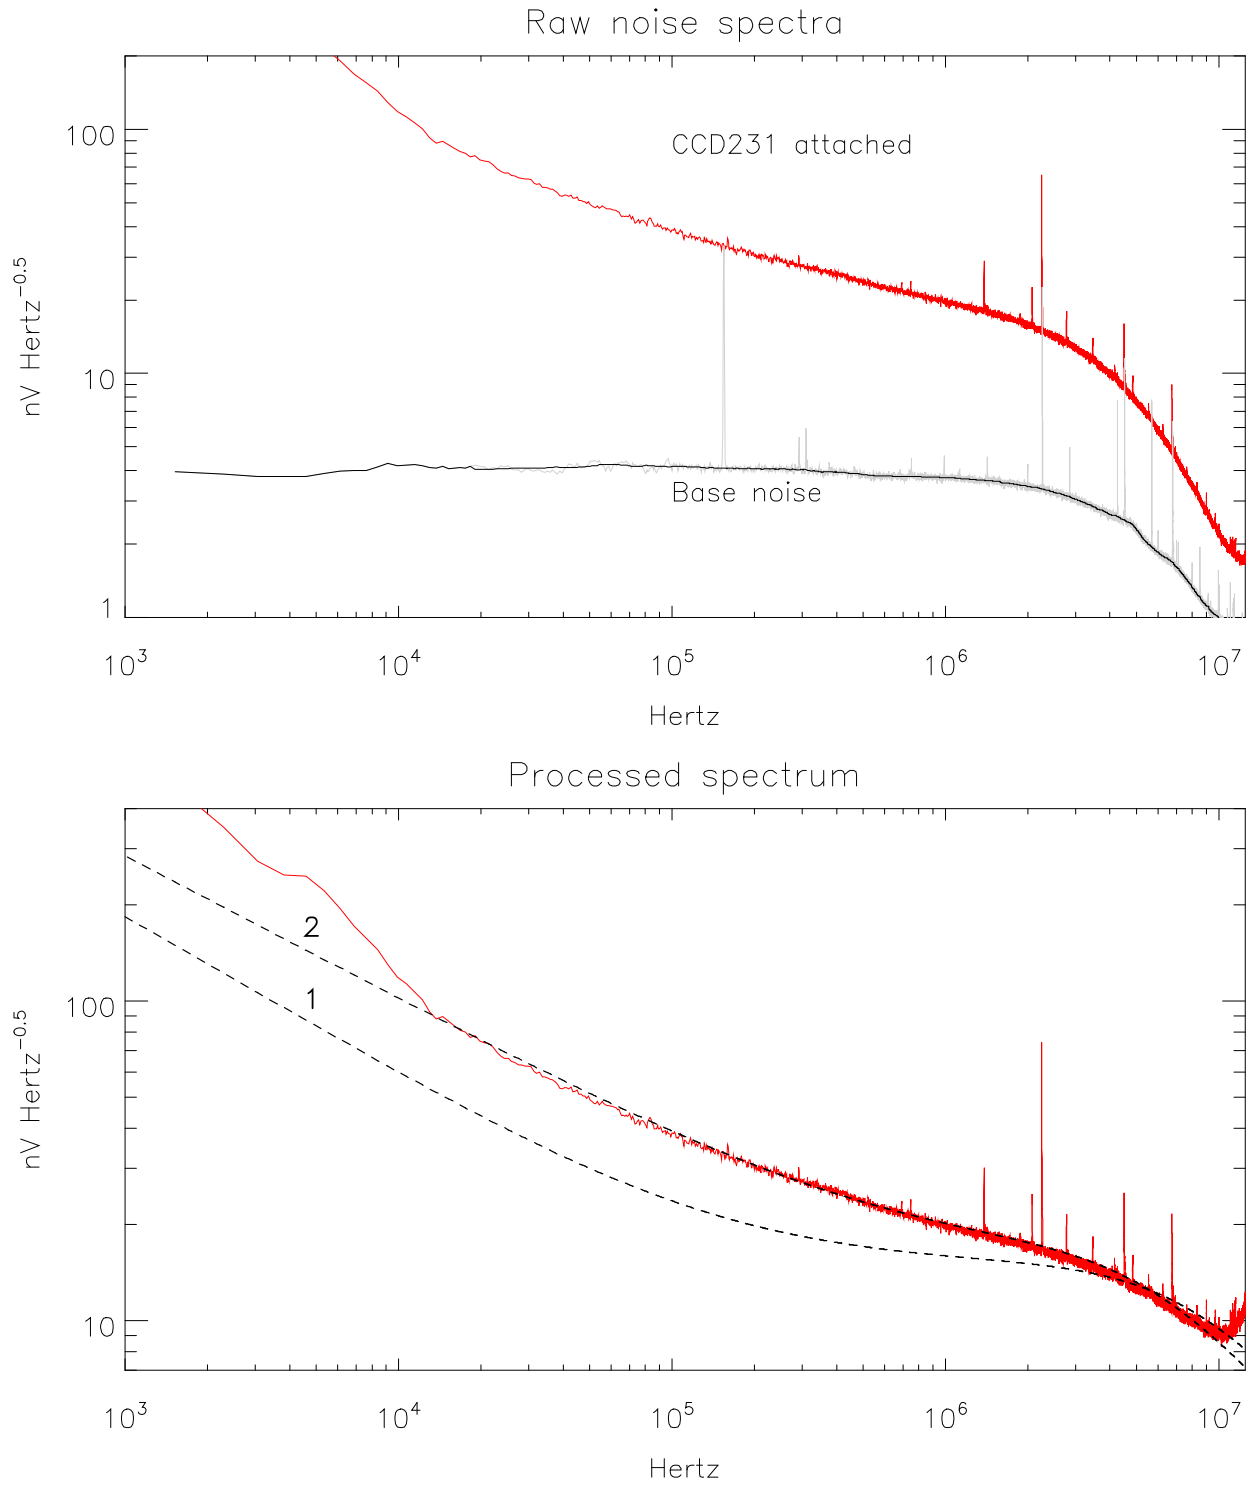
Figure 14. The CCD Noise spectrum. Top panel shows the raw output noise with CCD attached and then with the CCD replaced with a 390Ω resistor (to show underlying video processor noise).The bottom panel shows the CCD noise spectrum with video processor noise and spectral response removed. Trace 2. shows a manual fit to the spectrum (see Equation 5) and Trace 1. shows the noise spectrum quoted by E2V. Narrow-band switcher noise is apparent above 1MHz.
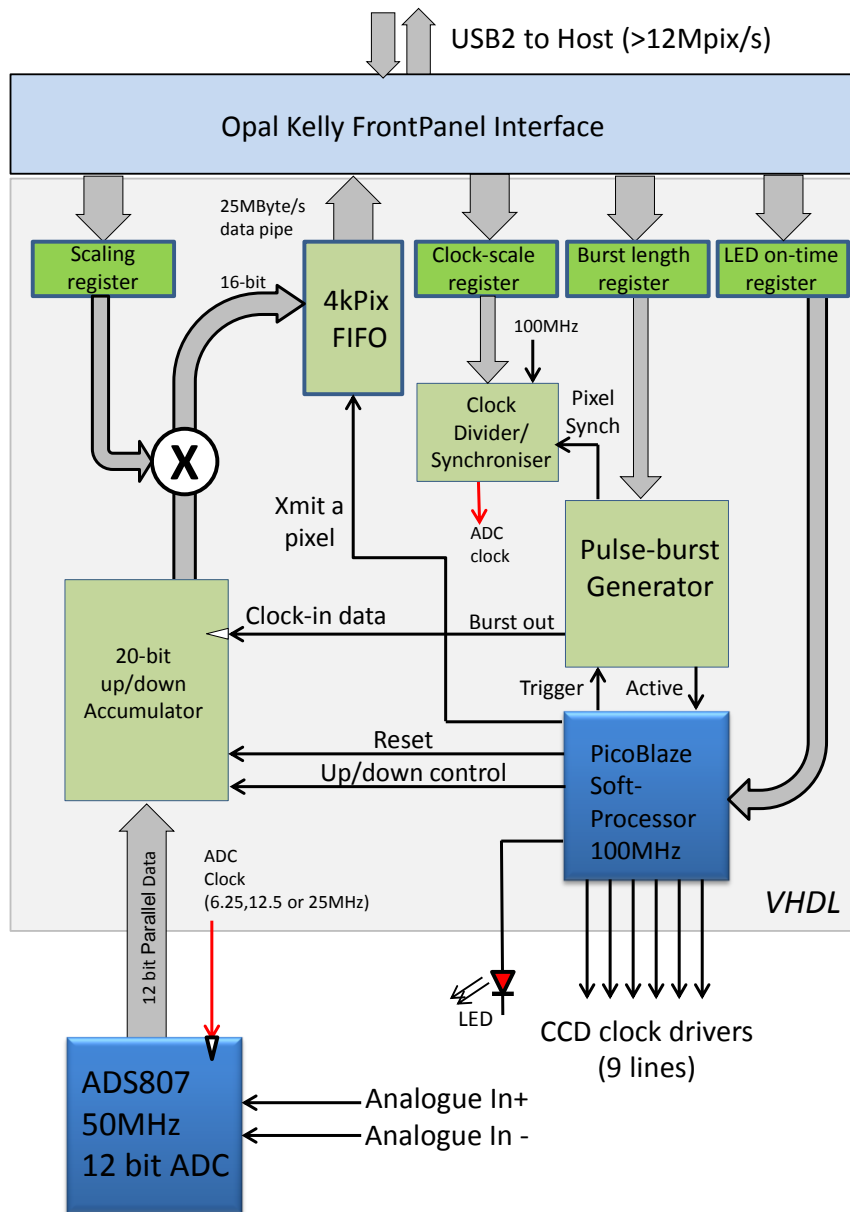
Figure 15. Main components of a DCDS video processor implemented in a Spartan-3E-1200 FPGA. Approximately 25% of the device fabric was used.

Implementing the differential averager scheme was surprisingly straightforward. All that is required is an up/down accumulator and one is provided free in the Xilinx IP Generator library. ADC samples taken during the reference pedestal are accumulated with positive sign, those taken during the signal pedestal are then accumulated with negative

sign. The final value in the accumulator is the pixel value that is transmitted via a FIFO to the PC. There was insufficient time to experiment with the weighted schemes. These would have required multiplier functions and a look-up table of weighting factors to scale each sample prior to it being added to the up/down accumulator. The principle problems encountered in the implementation were dealing with the 6-sample ADC internal pipeline and jitter between the ADC conversion clock and the start of each pixel. The PicoBlaze was only an 8-bit device so required nested loops to read out the 4k x 4k pixel CCD. This caused a small delay/advance in the point at which ADC conversions started within each pixel which was reflected in the bias images as subtle level shifts every 256 pixels. This was overcome by stretching the ADC clock to bring it into synchrony at the start of each pixel.

7.2 Imaging performance

The spectrum analyser system gave a result for the CCD noise in units of Volts RMS. What was really required was the value in units of e^- RMS which is the usual convention for quoting read noise. For this it was necessary to measure the conversion gain of the CCD and this was done using the standard Photon Transfer technique. Only the first hundred or so rows of the CCD could be read out before reaching saturation due to the large dark current and high preamplifier gain but this was enough to get good photon transfer statistics. The result is shown in Figure 16. The photon transfer graph of mean versus variance shows that the video processor is well behaved, is in its linear region and has a gain of approximately $1.6e^-$/ADU per ADC sample pair. This conversion value was optimum and ensured that quantisation noise in the ADC did not contribute significantly to the overall system noise. One complication of operating the CCD at such elevated temperatures was that the serial register was contaminated with a lot of dark current. The x-underscan region of the image did not therefore accurately represent the true bias level of the CCD (which we need to know for the photon transfer analysis). This was cured by reverse clocking the serial register by 20 pixels prior to each line transfer. This ensured that the first 20 columns of each row were free of dark current accumulated during the previous line read.
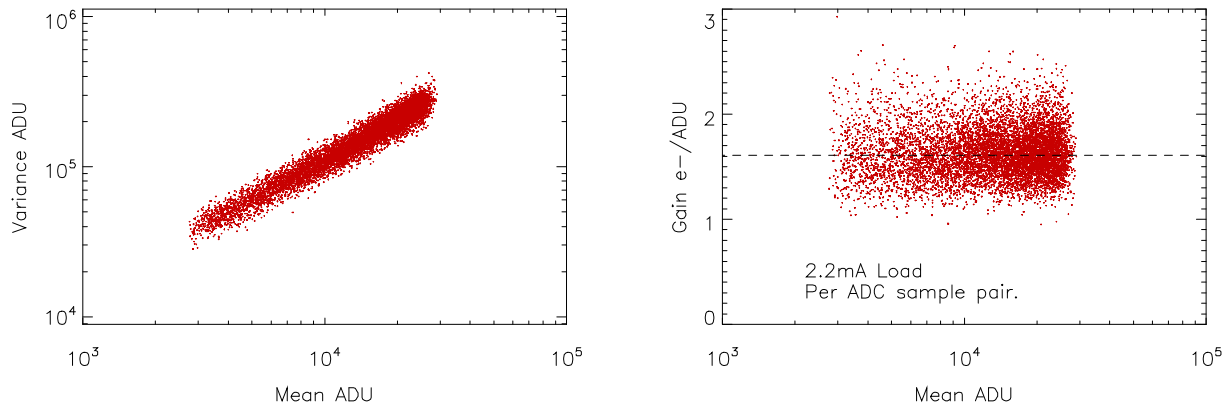

Figure 16. Photon transfer data used to calculate system gain in units of e-/ADU and also to demonstrate that the video amplifier is linear.

Figure 17. One of the frames used to calculate the photon transfer statistics. Serial underscan and overscan strips are visible to left and right. The vertical gradient in the image area due to dark current is apparent.

7.3 Actual noise performance

The ADC frequency was then varied from 6.25 to 25MHz and the number of sample pairs (one in the reference pedestal, one in the signal pedestal) varied from 1 to 255 per pixel. The samples were processed in the FPGA using the differential averager scheme. Due to the dominant dark current, the serial register was clocked entirely in reverse during the test to ensure clean bias images clear of any signal. The RMS noise in these images was measured and converted to units of e⁻ RMS. To get an idea of how well the DCDS hardware was working the theoretical noise was calculated using the frequency domain method described in Section 1, the known CCD noise spectrum N(f) shown in Equation 5 and the transfer function H(f) from the literature[2,3].The result, shown in Figure 18, shows that theoretically the system should have achieved a minimum read noise of 3.4e- RMS with pedestal widths of about 9us, corresponding to a pixel rate of close to 37kpix s^{-1}. What was actually measured was a minimum read noise of 4.2e- at pedestal widths of around 5μs. The higher ADC speeds definitely gave a better result but moving from 12.5 to 25 Msps gave only a small incremental improvement. One thing to notice in Figure 18 is that at pedestal widths below a couple of μs the read noise shoots up almost vertically. This strongly suggests a problem with the ADC pipeline and it could be that a few garbage ADC values are being included in the DCDS calculation. It could also be that the ADC clock stretching required to synchronise the ADC to the pixel disturbs its internal pipeline in some subtle fashion.
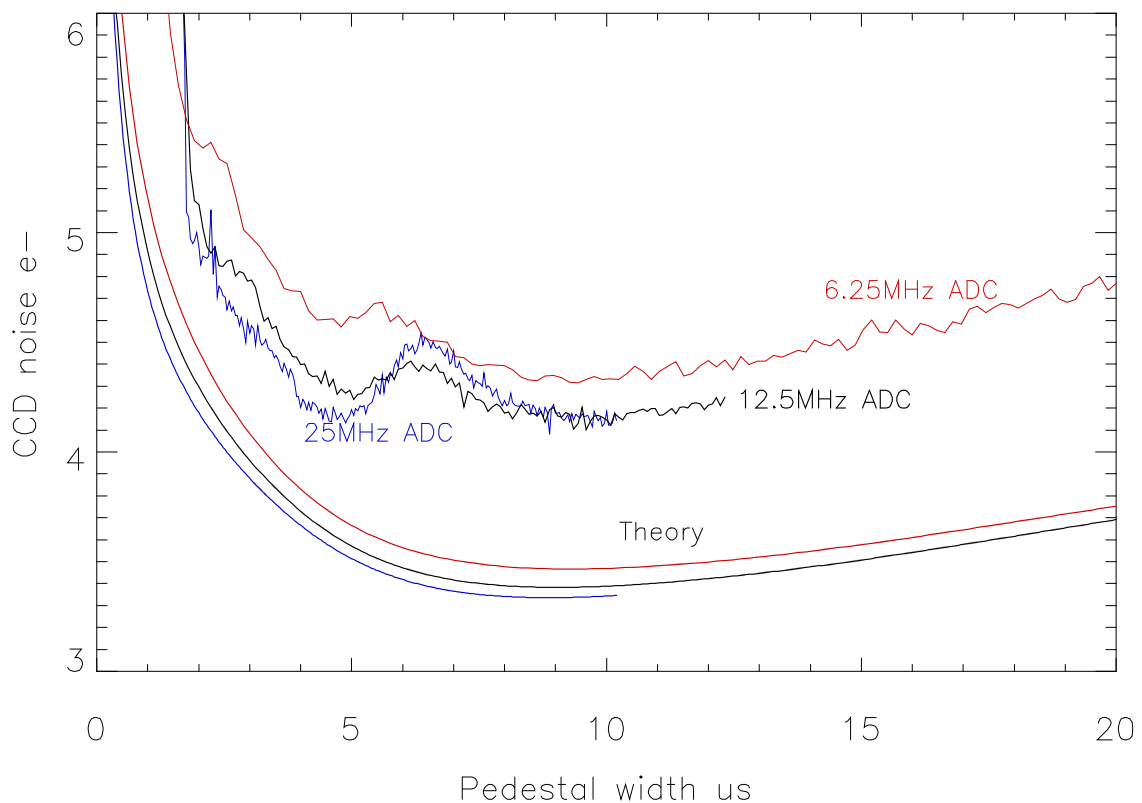

Figure 18. Actual measured noise performance at three different ADC frequencies for a whole range of pedestal widths. Plotted theoretical values were obtained by applying Equation 1 to the actual CCD231 noise spectrum shown in Figure 14. The pixel period will be about 3x the pedestal width, so a pedestal width of 10μs corresponds to 33kpix s^{-1}.

7.4 Shortcomings in implementation

The CCD231 was perhaps not the best CCD to use for this study. It would have been better to use a smaller device with an AIMO low-dark current design. It would also have been good to use a branched video preamplifier chain with low and high-gain outputs, each with its own ADC. The high-gain output could then be used to produce the noise spectrum and the low-gain output for the imaging tests without having to resort to component changes. Running the whole system from batteries and placing it in a Faraday cage would remove switcher and EMI problems and allow a better measurement of the underlying CCD noise.

8. CONCLUSION

The CCD231 should be capable of much better noise performance than its data sheet suggests. Controller noise must make up a significant part of the data sheet value. It is not clear why the measured corner frequency of the CCD231 was so high. It cannot be typical since at ESO this CCD regularly gives a noise of 2.0e- RMS which it would be incapable of reaching with a corner frequency of 600kHz. E2V have responded to this and suggest that ΦR should be held high during noise measurement but this was not tested. DCDS is very attractive for many reasons, it simplifies the electronic design and dispenses with the complex switched capacitor filter that is needed for the analogue-implemented differential averager scheme. It is less certain whether DCDS offers any significant noise advantage. The ability to weight the DCDS samples and implement novel digital-filter schemes does offer a small noise improvement. However, this advantage is really very small and is only realised at very low pixel rates where 1/f noise starts to dominate. Most scientific applications require pixel rates of several hundred kPix s^{-1} if the image is to be read out within a reasonable time. In this regime the best scheme is the classic differential averager (referred to as Dual Slope Integrator in some papers) with constant ADC weights. If a 5-10% noise reduction is sought then the designer may opt for weighted sample solutions but their time may be better spent concentrating on optimising the preamplifier and power supply performance. It seems that DCDS does not offer any "killer app" weighting schemes and we must wait for innovations from the manufacturers in order to further reduce read noise. The time domain model shows that the differential averager is the best all-round CDS scheme. This can be implemented both with analogue and digital CDS schemes and works well across a wide range of pixel rates. The model also predicts that a CCD preamplifier should have an analogue bandwidth $f_{3dB} \geq 6. f_{pix}$. If DCDS is used then the ADC should have a frequency of between 12 and 30 $f_{pix.}$ which corresponds to between 4 and 10 samples per pedestal. The DCDS system that was implemented here gave a noise minimum of 4.2e- which was 0.8e- worse than the expected value calculated by applying Equation 1. to the measured noise spectrum shown in Figure 14. In this test the pixel rate was 67kpix s^{-1}, the analogue bandwidth 3.8MHz and the ADC frequency 25MHz. This corresponds to over 100 samples per pedestal, way in excess of the minimum suggested by the model. Problems with the ADC pipeline were suspected. A Spartan 3E-1200 FPGA contains sufficient resources to comfortably implement a single channel DCDS using the differential averager scheme with ADC frequencies up to at least 25MHz. Any DCDS system would be greatly simplified by using ADCs without pipeline delays.

REFERENCES

[1] Janesick. Fundamental Performance differences between CMOS and CCD imagers: Part IV. *SPIE Conf. Series 9591,* 2015.

[2] D. J. Hegyi, A. Burrows. Optimal sampling of charge-coupled devices. *Astrononomical Journal*, 85:1421–1424, October 1980.

[3] G. Hopkinson ,D. Lumb. Noise reduction techniques for CCD image sensors. *J. Phys. E: Sci. Instrum.*, 15:1214–1222, 1982.

[4] J-L. Gach. A New Digital CCD Readout Technique for Ultra–Low-Noise CCDs, *PASP* 115, 2005.

[5] M.J.Clapp. Development of a test system for the characterisation of DCDS CCD readout techniques. *SPIE Conf.Series*, 8453, July 2012.

[6] K.D. Stefanov. Digital CDS for image sensors with dominant white and 1/f noise. *IOP J.Inst. 10*. 2015.

[7] C.Alessandri. Theoretical Framework and Simulation Results for Implementing Weighted Multiple Sampling in Scientific CCDs. *MNRAS 455*. 2016.

[8] J.Estrada, G.Cancelo. Beating the 1/f noise limit on Charge Coupled Devices, *SPIE Conf. Series*, 8453, 2012.

Software used for this study can be found at http://www.qucam.com/software/VirtualOscilloscope.html