Low readout noise CCDs in optical interferometry

Alastair G. Basden[a], David F. Buscher[a], Chris A. Haniff[a], Craig D. Mackay[b]
[a]Astrophysics Group, Cavendish Laboratory, University of Cambridge, UK;
[b]Institute of Astronomy, University of Cambridge, UK

ABSTRACT

Most current CCDs cannot be used as optical interferometric sensors because the high readout noise disguises the small signal. However, new low light level charge coupled devices (L3CCD) have a large on-chip gain which can allow a signal to be detected above the noisy readout amplifier. This gain has a statistical nature, meaning that the photon input cannot be predicted exactly. We investigate several techniques for photon prediction at different light levels, and demonstrate how this affects the noise on the signal. Accurate signal estimation can be achieved with very faint signals, up to about one photon per pixel per read. Above this, accuracy gradually decreases, though our signal-to-noise ratio is never worse than $\sqrt{(2n)}$. Optical interferometry requires detection of very faint signals, and the use of an L3CCD is found to allow reproduction of interferometric visibilities to high precision. Custom instrumentation used for control is also detailed.

Keywords: L3CCD, low readout noise, optical interferometry, detector, COAST

1. INTRODUCTION

Over the past twenty years, CCDs have been developed to have a very low readout noise. However, at best this is about 2-3 electrons, meaning that signals fainter than this cannot be detected. As the readout rate increases, so does the readout noise. At readout rates of about 10MHz, the lowest present readout noise corresponds to about ten photons. Many astronomical applications require measurement of a signal smaller than this. Low Light Level CCDs (L3CCDs) overcome this as they can have a high internal gain. However, the gain process is statistical, achieved by avalanche multiplication. Therefore the CCD output cannot be used to calculate the exact photon input, rather, it can only infer the most likely photon input. The accuracy of this prediction determines how useful the L3CCD is as a photon-counting device.

Optical interferometry, such as at the Cambridge Optical Aperture Synthesis Telescope (COAST), involves the detection of interference fringes from two or more light beams. The temporal variation of these fringes must be measured as the relative delay between the light paths is altered, requiring sampling at up to 10kHz. High speed detection and faint sources can result in a very low signal, often lower than one photon per pixel. For spectroscopic interferometry, we would like frame sizes of at least 100 pixels in one dimension. This requires MHz pixel rates to read out the array, and when coupled with the need for sub-electron read noise, a high performance, low noise detector is needed. The use of L3CCDs as spectroscopic detectors is investigated, with particular regard for how accurately the interferometric visibilities can be reproduced.

2. MULTIPLICATION STATISTICS

An L3CCD is able to provide a large internal gain by passing electrons through a multiplication register, an extension of the serial register. As electrons pass through this register, a high voltage on one of the pins can cause avalanche multiplication. The probability of creating an extra electron at each stage is small (1-2%), but many such stages can lead to a large overall gain, which can typically be adjusted between 1-10000. There are typically several hundred stages (591 for the E2V CCD65 giving an overall gain of $(1+p/100)^{591}$,e.g. 6629 if p=1.5). The final signal is then much greater than the readout noise. We find that the probability distribution for the electron output (x) is given approximately by:

$$P(x) = g^{-n} x^{n-1} e^{-x/g} / (n-1)! \tag{1}$$

where the mean gain, $g = (1+m/100)^r$, m is the multiplication factor, r is the multiplication register length (usually large) and n is the number of electrons entering the multiplication register. This distribution is shown in figure 1.

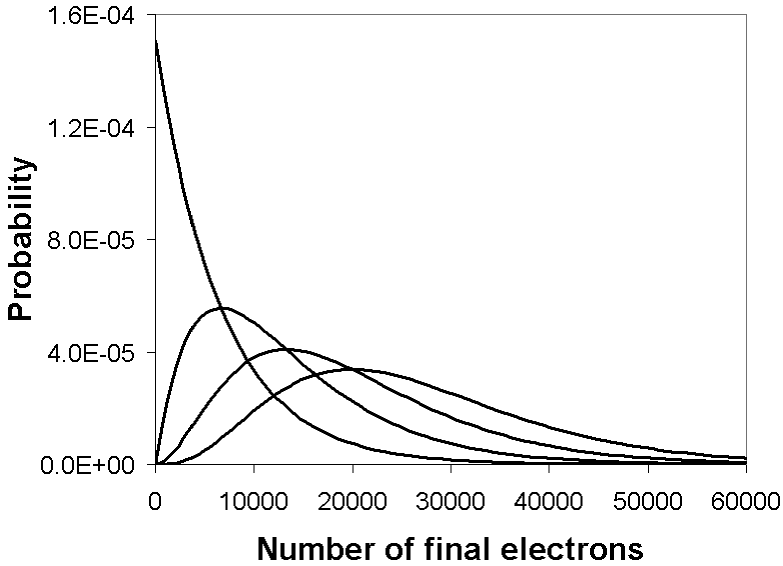

Figure 1. Output probability distribution for one (leftmost curve, n=1) to four (rightmost peak, n=4) input photons, with mean gain, g=6629. For n=1, the most likely output is one electron, and the mean output is equal to the mean gain. For n=4, the most likely output is about 20000 electrons, three times the mean gain. The mean output is about 26000, four times the mean gain.

The probability distribution can be used to estimate the number of received photons from a given output. The L3CCD can generate an output event with a high signal-to-noise ratio (SNR) from a single electron detected by the CCD imaging array. This allows processing of the output signal in a variety of ways. There is no spatial or temporal dispersion between pixels due to the amplification, with each pixel remaining independent from others. This gives advantages over some intensifier methods, allowing much higher mean photon arrival rates to be handled properly. However, the statistical nature of multiplication can add dispersion into the output signal for a given input, resulting in an extra source of noise. We are unable to determine the exact input, though as we show, we can estimate it very well.

3. READOUT MODES

There are several different ways that we can interpret the output from the L3CCD, each of which has advantages in different situations. We can use our knowledge of the mean gain to improve on our input prediction, improving the signal-to-noise ratio.

3.1 Analogue mode

The output is treated as with normal CCDs, taken to be the input multiplied by the mean gain. Due to the Poisson nature of the multiplication, if we take a large number of readings at a given light level, the output noise scales as $\sqrt{(2n)}$, n being the number of photons. This is $\sqrt{2}$ greater than we would expect from photon shot noise and effectively halves the quantum efficiency of the device. The problem we have is that there is an overlap in the output probability distributions for neighbouring input signals, as shown in figure 2. This mode is best used at higher light levels, where the increase in noise is not so important.

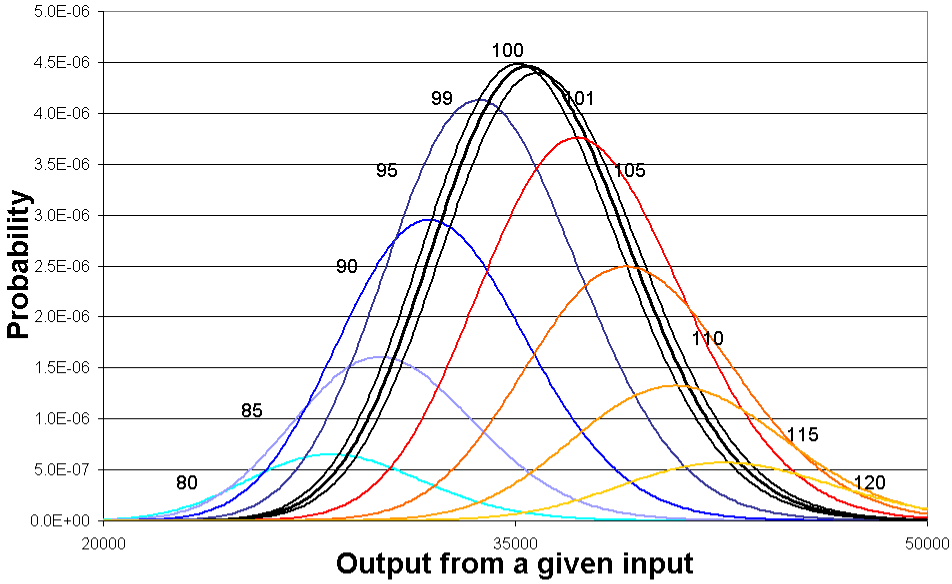

Figure 2. Output distribution for Poisson light with a mean light level of 100 photons and mean gain 358. The curves show the probability of a given output when the label gives the input number of photons. This shows that for example, an input of 90 photons could easily be mistaken to be 100 or 110 photons due to the dispersion caused by the multiplication process, as the overlap of probabilities shows.

3.2 Single threshold mode

If the input signal is low, with an average of less than about one photon per pixel per frame, then we can assume that any detected event above a threshold level corresponds to one photon. This eliminates the dispersion caused by the multiplication process, and adding many readings gives a noise of $\sqrt{n}$, returning us to the noise level due to the discrete Poisson nature of photons.

At higher light levels, there will be significant coincidence losses as shown in figure 3, and we will underestimate the true input. Where these losses are modest, they are easily estimated statistically and corrections can be introduced to improve linearity. This method only works up to a mean of about one event per pixel per readout, and so should only be used at very low light levels.

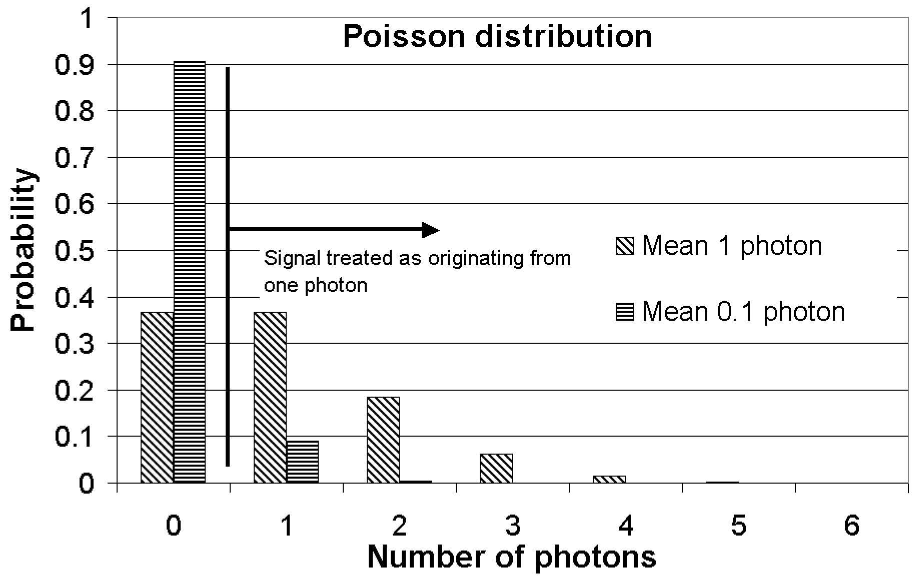

Figure 3. Coincidence losses when using a single thresholding readout mode at 0.1 mean photons per pixel (negligible) and 1 mean photon per pixel (high). In this readout mode, every signal above a background noise level is treated as originating from one photon. When the mean light level is much less than one photon per pixel, there is little chance of two photons actually being detected, so we lose very little in our assumption. However, when the light level is higher (e.g. one photon per pixel), there is a significant chance that two photons will be detected, though we would interpret this as one, thereby reducing our input prediction due to this coincidence loss.

3.3 Multiple threshold mode

In this mode, we use our knowledge of the mean gain to improve our input prediction. For a given output signal and known gain, we can estimate the most likely value of the input number of photons. Errors will be made but at low signal levels these are small, minimising the additional noise factor. In addition, knowledge of the mean light level will allow an even greater improvement in our input prediction.

By thresholding the output using multiple thresholds, are able to remove some of the dispersion caused by the multiplication process. We will not always get our thresholds in the right place for every output signal, meaning that sometimes our prediction will be wrong. However, we can choose our thresholds such that on average we will get the right prediction of the input. At low light levels we will be correct most of the time, since we would use a large threshold step size.

At higher signal levels we lose accuracy progressively. This leads to noise scaling as $\sqrt{n}$ at low light levels (standard Poisson photon shot noise statistics) where a single threshold is effectively used, and tends towards $\sqrt{(2n)}$ as the light level increases. In choosing the thresholds for selecting the input, a noise level (typically 6σ, where σ is the mean readout noise) is first subtracted from the signal. Then the remaining signal is placed into equal sized bins (the threshold step size), representing the predicted number of photons, as shown in figure 4.

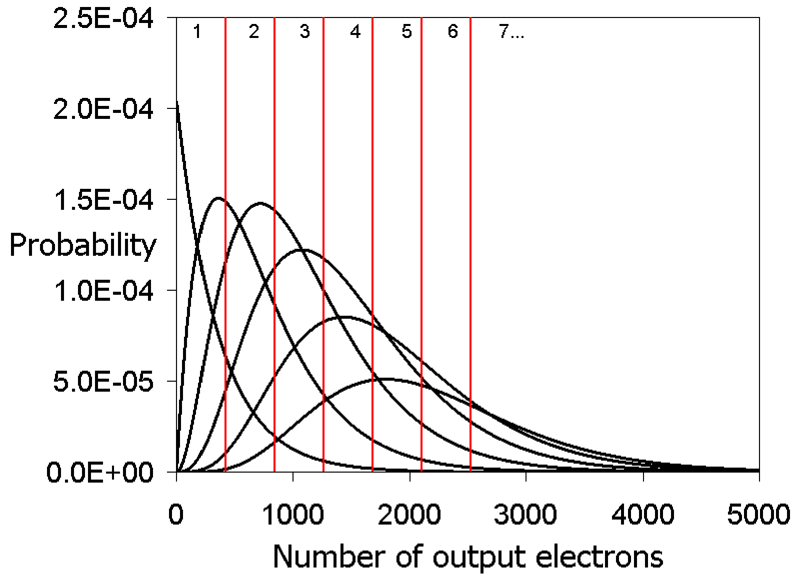

Figure 4. Multiple thresholds shown for a mean gain of 358 and mean intensity four photons per pixel. Increasing photon input shown from left (1) to right (6). Threshold step size is 421, as given by equation (2). If a signal between 1-421 is measured (after a noise level has been removed), we interpret is as corresponding to one photon. Likewise, if we measure between 843-1263 electrons, we interpret is as corresponding to three input photons.

This gives the benefits of both the single threshold mode and the analogue mode, allowing good input signal estimation at low and higher light levels.

The threshold step size that gives the best estimate for the photon input is dependent on the gain and light level. Errors in the threshold step size become less important at lower light levels where they are larger. Threshold step size tends to infinity as light tends to zero, giving rise to the single threshold mode.

At high light levels, the ideal step size tends to the expected gain and gives a noise scaling (and so quantum efficiency (QE) reduction) similar to the analogue case, as shown in figure 5.

We have investigated the ideal threshold step size, and found that it can be chosen to minimise prediction errors. This uses knowledge of the average light level (which will be known in many situations) and is given by:

$$\text{Threshold Step Size} = g(1+0.5\mu^{-3/4}) \qquad (2)$$

where g is the gain, and μ is the average light level (in photons per pixel). This gives accurate photon prediction over a wide range of gain and light levels, giving us the best of both the analogue and single threshold readout modes. If the gain is low, of similar size to the readout noise, a larger threshold step size will be required. If this is the case, an increase in gain should be considered.

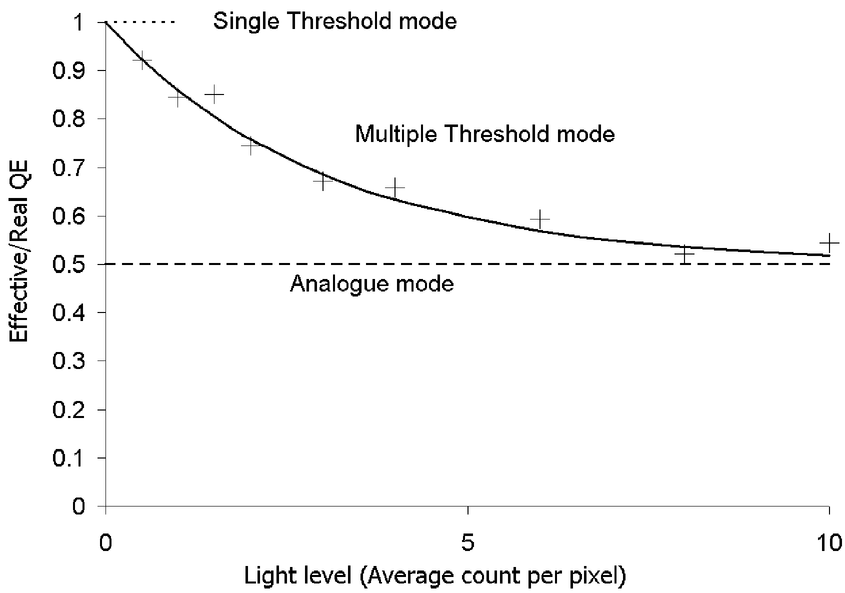

Figure 5. Effective quantum efficiency (QE) of an L3CCD for different readout modes. Shown for a gain of 6629. The effective QE is calculated from the signal-to-noise ratio of many trials, using the variance of the input prediction.

4. INTERFEROMETRIC DETECTION

Optical interferometers require the measurement of the complex visibility of a fringe. A fringe is the interference pattern generated from the combination of two or more coherent beams of light from the object of interest. The fringe is usually sampled temporally as the relative delay between the beams is altered, producing a time varying signal on the detector. Atmospheric fluctuations mean that sampling must be done quickly, typically at 5kHz. This short integration time means that the signal is very low, requiring a low readout noise detector. Spectroscopic detection is also desirable for an interferometer, and this reduces the flux available to each pixel due to dispersion of the fringe over the detector. This requires a multi-pixel detector, such as a CCD. An L3CCD fits these requirements, the low effective readout noise making it an ideal detector.

Typical light levels measured by an interferometer are very low and the mean signal level can easily be less than one photon per pixel.

At the lowest light levels it is appropriate to use a single threshold readout mode, as this minimises the extra noise introduced by the multiplication process. This also has the advantage that uncertainties in the mean gain do not affect fringe parameter estimation, provided the gain remains well above the readout noise. Visibility parameter estimation is unbiased. A multiple threshold readout mode will lead to the same results, since at low light levels, there is little difference between the two modes.

At intermediate and high light levels, a single threshold mode of operation is not appropriate due to high coincidence losses. Either an analogue readout mode, or a multiple threshold readout mode should be used. Analogue mode will produce unbiased results, while multiple threshold readout mode requires a small bias correction to be applied to the visibility amplitude. This bias is due to measurement of the fringe requiring detection at intensities both above and below the mean intensity, whilst the threshold step size will be chosen to suit the mean intensity. We find that the measurement of visibility phase of a fringe remains unbiased when using any of the L3CCD output modes discussed. Therefore, an analogue readout mode should usually be used, unless the higher signal-to-noise ratio of the multiple thresholded output gives a significant advantage.

Figure 6 shows simulation results for visibility amplitude reproduction at a mean light level of one photon per pixel per readout. We see that the visibility reproduction is good and the multiplication process does not affect the parameter estimation. It is possible to estimate low visibilities (down to about 1%) at low and high light levels.

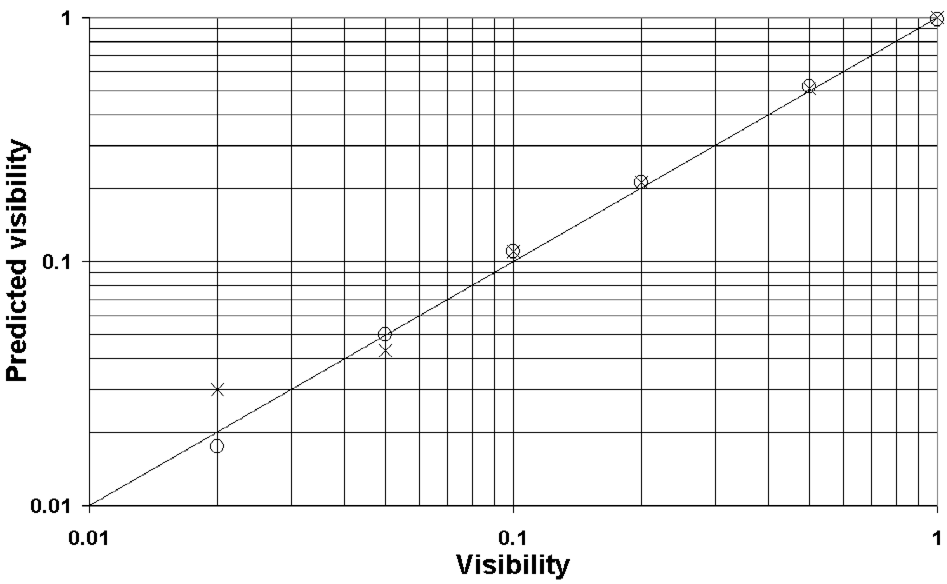

Figure 6. Visibility amplitude reproduction at a mean light level of one photon per pixel and gain 6629. Crosses show raw output (analogue mode) and circles show multiple threshold mode results corrected for bias. This shows that we are able to reproduce the true visibilities well when using L3CCDs, even at very low light levels. 250000 photons have been detected on the array.

4.1 Uncertainties in light level or gain

If the mean light level or gain is not known accurately, the threshold step sizes may be chosen wrongly. The mean gain can be controlled to about 1%[1], and the average light level can be estimated from the average raw counts in a sweep (or several sweeps) divided by the mean gain. For a large number of counts, the percentage error in light level is then only slightly greater than the error in gain. We find that the fractional error in visibility amplitude is typically much less than the error in gain or light level. This is due to the nature of visibility being composed of both a difference and a ratio ($V=(I_{max}-I_{min})/(I_{max}+I_{min})$). For this reason it is possible to estimate visibility using both analogue and thresholding modes, with little difference.

Visibility phase estimation is not affected since a wrongly chosen threshold size will only alter fringe peak height, not peak position.

5. SYSTEM DESIGN

We are currently constructing a controller for L3CCDs in our lab, capable of operating at high pixel rates, in any of the mentioned readout modes. A Digital Signal Processor (DSP, Analog Devices TigerSHARC) is used to generate L3CCD control waveforms (clock and signal process) to support a pixel rate of up to 15MHz (with capability for higher pixel rates). An embedded web server, allowing machine independent operation, controls the DSP. Custom circuit boards with a field programmable gate array will convert the DSP output into CCD waveforms, and the CCD output will be digitised to 14 bits. A host PC and frame grabber (Matrox Meteor II/dig) will then be used to capture the data and provide real time analysis of the visibilities and processing of the data for telescope control feedback. Light from an interferometer, the Cambridge Optical Aperture Synthesis Telescope (COAST) will be dispersed onto the CCD, allowing for spectroscopic operation. We hope to have the controller working by first quarter 2003, and integrated with the COAST by third quarter 2003.

6. CONCLUSIONS

The gain process of L3CCDs has been studied, and several methods of interpreting the output have been developed, depending on light level. L3CCDs have been considered as interferometric fringe detectors, and found to be ideal for this application.

- Single thresholding of the L3CCD output can be applied with accurate input prediction for photon rates up to about 0.5 photons per pixel per readout.
- Multiple thresholding of the output can be applied at any light level giving an advantage over the single threshold mode.
- With multiple thresholding the signal-to-noise ratio scales as $\sqrt{n}$ at low light levels, gradually decreasing to $\sqrt{(n/2)}$ at high light levels, as with the analogue mode.
- L3CCDs can be used for visibility parameter estimation, the multiplication process does not affect the results.
- Uncertainties in the light level or mean gain will have little effect on visibility parameter estimation, and are usually known to about 1%.
- We hope to have a system being tested at the COAST in nine months.

REFERENCES

1. Craig D. Mackay, Personal communication, 2002.
2. Jerram, P., Pool, P., Bell, R., Burt, D., Bowring, S., Spencer, S., Hazelwood, M., Moody, I., Catlett, N., Heyes, P.; "The L3CCD: Low Light Imaging without the need for an Intensifier." Marconi Applied Technologies, 2001.
3. C.D. Mackay, R.N. Tubbs, R. Bell, D. Burt, I. Moody, "Sub electron read noise at MHz pixel rates" in Sensor and Camera Systems for Scientific, Industrial and Digital Photography Applications, Proc SPIE 4306, 289-298 (2001).
4. T.R. Scott, "The Implications of atmospheric seeing for optical interferometry", Ph.D thesis, University of Cambridge, 1997.
5. J.F. Walkup, J.W. Goodman, "Limitations of fringe-parameter estimation at low light levels", J. Opt. Soc. Am, 63(4), 399-407 (1973).