

TECHNICAL NOTE 1: INTRODUCTION TO EMCCDs

Simon Tulloch, QUCAM

Calle Miranda 85

Breña Alta, Tenerife 38700

España

June 2009

1 History of CCD development

Up until the 1970s astronomical images were recorded on photographic plates. In the absence of alternatives astronomers had to content themselves with throwing away up to 95% of the photons collected by their telescopes; the maximum quantum efficiency (QE) of a photographic plate being in the region of 5%. The detector revolution began in 1970 in Bell labs where W. Boyle and G. Smith invented the charge coupled device (CCD). It was initially conceived of as a memory storage device but its potential as an imager was quickly recognised. The technology has matured incredibly quickly, driven mainly by the needs of the military and consumer camera market. Now, less than 40 years since their invention, CCDs are approaching the theoretical limits of their performance. QEs of greater than 90% are now easily obtained, meaning that a 1m diameter telescope with a CCD can perform more or less the same science as a 4m telescope equipped with only a chemical-emulsion camera. In terms of detection threshold a CCD camera is limited by the electronic noise in its read-out amplifier. For a conventional CCD this has now reached approximately $2e^-$ rms, not quite enough to resolve a single photon

but still impressively and surprisingly low. CCDs have also grown steadily larger during their development. The first astronomical CCD measured 100x100 pixels. The current record is now held by DALSA Semiconductor who have manufactured a 111 Megapixel device, the worlds largest integrated circuit. In 2001 a new type of detector was announced: the electron multiplying CCD or EMCCD. This was first described by E2V Technologies [7] and Texas Instruments [5]. These devices incorporate an avalanche gain mechanism that renders the electronic noise in their read-out amplifiers negligible and permits the detection of single photo-electrons. Whilst this has been possible for some time with image tube detectors such as the IPCS, it has never been available with the pixel count, convenience of use, relatively low cost and high quantum efficiency of a CCD. The EMCCD thus combines many useful properties that are sought by astronomers working in photon-starved regimes and whilst currently somewhat under-exploited they have the promise to make a large contribution to astronomy in the near future.

2 Key CCD concepts

2.1 Photoelectric effect

This effect lies at the heart of CCD operation. To understand it one has to consider the energy bands in the Silicon from which the CCD is constructed. Atoms normally have narrow discrete energy levels at which electrons are stable. In crystalline solids such as Silicon these levels are smeared out into wider bands due to the perturbing influence of each atom on its neighbours. Most electrons are to be found in the valence band where they are bound to a specific Silicon atom and unable to move around. A small fraction of electrons are to be found, however, in the conduction band situated above the valence band. These electrons have absorbed sufficient energy, either from an incident photon or from thermal excitation, to reach the conduction band from where they are then free to move throughout the Silicon and contribute to its conductivity. The band-gap in Silicon is approximately 1.26eV so any photon with an energy more than this will be able to excite a valence electron

into the conduction band. Below this energy, which corresponds to that of a $1\mu\text{m}$ infra-red photon, the Silicon becomes transparent since the photons are unable to interact with the electrons through the photo-electric effect. It is a lucky coincidence that Silicon, so readily usable for the manufacture of micro-electronic devices, has a band-gap so nicely matched to the energy of photons in the visible part of the spectrum.




Figure 1: The photoelectric effect. The band-gap of the Silicon is 1.26eV , approximately equal to the energy of a near-infrared photon. An electron excited into the conduction band leaves a hole in the valence band. Both are then free to move throughout the Silicon and contribute to its conductivity.

2.2 Structure of a CCD

There is a useful analogy (see figure 2) that can be used to describe how a CCD works. One can imagine a series of parallel conveyor belts arranged in columns on which are placed a series of buckets. With the conveyor belts stationary the buckets are exposed to a shower of rain. At the end of the exposure the conveyor belts then start up transferring the collected rain towards a single measuring cylinder where the contents of each bucket is recorded. Each bucket corresponds to a pixel in this model, the measuring cylinder is the output amplifier and the rain corresponds to photons (or rather the photo-electrons that they generate). So collected charge packets are transported physically through a CCD device during readout. Figure 3 shows how the earlier geometry, introduced using the analogy of a conveyor belt is actually implemented in the Silicon. This figure considers only the imaging area of the CCD, also known as the parallel area, where the photo-charge is collected.




Figure 2: A CCD Analogy. Here the rain drops are analogous to photons, the buckets can be considered as pixels and the measuring cylinder corresponds to the readout amplifier. Once the rain stops (i.e. the shutter closes) the conveyor belts starts up and transports the collected charge to the output amplifier one pixel at a time.

In cross-section the CCD consists of two layers of Silicon one of n-type the other of p-type. This constitutes a p-n diode and therefore has an intrinsic electric field perpendicular to the junction between the layers. On top of the Silicon is deposited an insulating layer of SiO_2 and above this are deposited a series of transparent electrodes formed of Polysilicon. These electrodes are generally grouped into three alternating sets known as 'phases'. During integration of the image, one of these phases is held to a low voltage whereas the other two are held high. This induces an electric field in the underlying Silicon that defines a series of potential wells into which photo-electrons are attracted. Each potential well then constitutes a pixel. Note that the electrodes only define the extent of each pixel in the axis parallel to the pixel columns. Pixels have to be defined in the orthogonal axis by an additional structure in the Silicon known as a channel stop.

The above is only an approximate description and there are several variants on this

theme. For example, some CCDs have 2 or 4 parallel phases whereas others function with all phases held low during integration.




Figure 3: Structure of a CCD. A small part of the imaging area is shown. Pixels are defined horizontally by the channel stops and vertically by the potential wells induced by the voltages on the electrodes.

2.3 Quantum efficiency

The quantum efficiency or ‘QE’ is one of the most important parameters that describe a CCD. It shows what fraction of the incident photons produce a measured photo-electron in the device. QEs in excess of 95% are possible. Silicon has a refractive index of almost 4, which requires the use of an anti-reflective (AR) coating to avoid most photons simply being lost through reflection at the surface of the CCD. Since AR coatings have to be tuned to a certain wavelength it is usual to optimise them either for the red or blue ends of the spectrum. Broader-band coatings are available but they tend to have a lower peak efficiency. Other losses in QE can occur through ‘recombination’. When a photo-electron is generated it leaves behind a hole in the valence band. If this electron-hole pair are not rapidly separated, for example by an electric field, they will tend to recombine and be lost. It is therefore important to ensure that there are no field-free regions within the imaging area that might permit this to happen. This is normally done by thinning the CCD during manufacture using an etching process.

2.4 Charge transport

The same structures that are used to collect the photo-charge during integration are also used to transport that charge to the output amplifier at the start of read-out. The electrons will always seek the region of lowest potential energy (highest +ve electrical potential). They are therefore attracted to the region beneath the most +ve electrode. By alternating the voltages on the electrode phases in a precise sequence we can then push the photo-charge packets through the CCD. By decreasing the voltage on the phase under which the photo-charge is stored and at the same time (with a small overlap) increasing the voltage on the adjacent phase we move the charge by one step. This can be repeated as many times as necessary until the charge reaches the output amplifier where it is measured. During the transfer of charge from one phase to its neighbour the voltage on the third phase must be held low so as to maintain the potential barrier between adjacent pixels. A serial transfer, in an EMCCD, is illustrated in figure 11.

Charge transport is typically very efficient with 0.999999 of the photo-charge being transferred at each stage. This fraction is known as the ‘charge transfer efficiency’ or CTE and is another important CCD parameter. The amount of charge not transferred (known as the ‘deferred charge’) can increase, however, if the clock voltages are not properly optimised. This can give rise to trailed images where stars appear elongated in the direction of charge transfer.

2.5 Full-well

There is a limit to how much photo-charge a pixel can contain. If this limit, known as the ‘full-well’, is exceeded then any additional charge leaks into neighbouring pixels. Since the channel stops provide a very effective isolation between pixels in adjacent columns, this charge leaks out preferentially along the columns. The potential barrier between pixels is lower along this axis. The result is that overexposed stars appear to have bright streaks extending above and below them. The general term for this effect is ‘blooming’ and it is illustrated in figure 4.



Figure 4: CCD blooming. When a pixel becomes saturated, additional photo-charge will then spill out into neighbouring pixels. Image courtesy of Nik Szymanek.

2.6 Frame-transfer (FT) CCDs

A frame-transfer CCD has its parallel-clock area sub-divided into two equal parts that can be independently clocked. One of these areas has an opaque Aluminium layer deposited on its surface that renders it insensitive to light. This insensitive area can then be used as a temporary frame store and is termed the ‘store area’. This is shown schematically in figure 5. Sacrificing 50% of the area of the CCD in this way may seem somewhat wasteful but it permits much higher observing efficiencies in fast frame-rate applications. Normal astronomical CCDs are used in conjunction




Figure 5: Geometry of a frame transfer CCD. The image and store areas have independent clocks.

with a mechanical shutter that typically operates in a few 10s of ms. To avoid image smear the shutter must be closed prior to the commencement of readout. Whilst shut, the CCD is clearly blind and photons are wasted. Another drawback is that mechanical shutters are not fully reliable, especially if operated repeatedly at high frequency. CCDs intended for TV rate applications (25Hz) will be of FT design. At the end of one exposure the image is rapidly transferred under the opaque shield of the store area. This can be an extremely rapid process: just a few ms. Once in the store, the image can then be read out more slowly concurrently with the integration of the subsequent frame. The CCD is therefore almost continuously exposing and gains are made not just in efficiency but also in reliability. Almost all EMCCDs are of this design, one of which (the E2V CCD201-20, [4]) is shown in figure 6.




Figure 6: A Frame Transfer CCD. The image area is visible as a dark square offset from the centre of the device. The store area is immediately adjacent and of equal size but hidden under the Aluminium light shield.

2.7 Dark current

Electron-hole pairs are generated in a CCD when a photon is absorbed. The photon provides the energy required to excite an electron from the valence into the conduction band where it has a good chance of being collected in the potential well of a pixel. This energy can, however, also be provided by the thermal motion of the atoms in the CCD material. At room temperature these thermally generated electrons, which constitute a ‘dark current’ can saturate a CCD pixel in a second or

so. For optimum noise performance it is therefore necessary to cool the CCD. Many EMCCD cameras are used with Peltier coolers at temperatures of around -40°C . In this higher temperature regime special measures have to be taken to ensure low dark current. One of these is to design the device for ‘inverted mode’ operation. Here, the parallel clock phases (those present in the image and store areas) are all maintained in their low states during image integration. This is in contrast to the non-inverted mode, used by most cryogenically (i.e. liquid Nitrogen) cooled detectors, where one of the parallel phases is maintained in the high state so as to provide a potential barrier between adjacent pixels. Figure 7 shows a cross section through the image

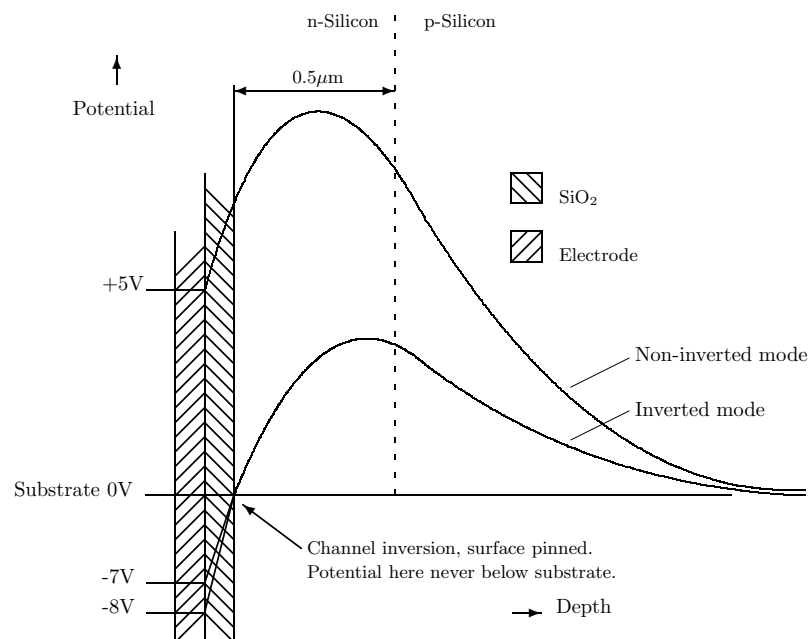


Figure 7: Potential gradients within the CCD when both inverted and non-inverted.

area of a CCD and indicates the electrical potential through the device in each of these modes. In both cases the electric fields generated by the externally applied electrode potential and the intrinsic field generated by carrier diffusion across the p-n junction maintain a potential maximum within which the photo-charge accumulates. As the externally applied electrode voltage is reduced the potential at the surface of the CCD, i.e. the interface between the Silicon and the SiO₂ insulating layer that underlies the electrodes, also falls and the device approaches ever more closely its inversion point. At inversion this surface potential becomes equal to that of the substrate and an interesting phenomenon occurs. Holes flow out from struc-

tures in the CCD known as channel stops (which define the columns of the CCD) and populate the surface region. Any further reduction of electrode potential then has no further effect on the potential of this interface: it becomes locked, or ‘pinned’ to the substrate potential. These holes have an important effect on the dark current of the device. The major dark current contribution comes from mid-band interface states at the surface and these holes effectively mop-up this charge greatly reducing the dark current. For this reason inverted mode operation is preferred for Peltier cooled devices with their intermediate operational temperature. Figure 8 shows how the dark current in a CCD201 experienced a rapid drop as the device became inverted.

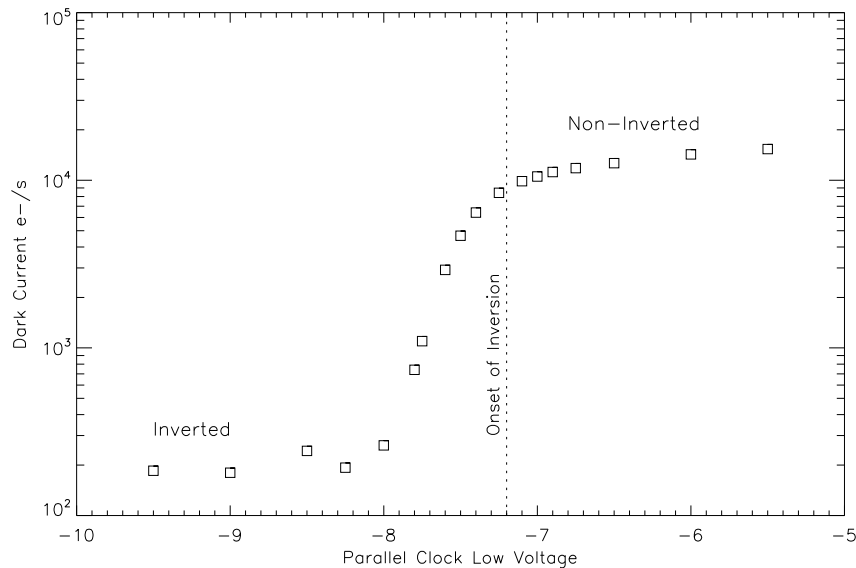


Figure 8: Dark Current as a function of parallel clock-low voltage (with respect to substrate).

2.8 Dark current suppression using Dither

Dither is a term used by E2V to describe the effect of vertical clock modulation on the dark current rate [8]. It is an important effect since it allows a CCD to be used in non-inverted mode, so as to gain the advantages of low CIC, whilst at the same time avoiding the effects of increased dark current.

During the clear and readout operations of a CCD, the vertical clocks are modulated in order to effect the vertical transfer of the charge. This modulation, even if the

clocks remain non-inverted at all times, can inject holes into the mid-band surface states where they neutralise the surface dark current. The hole lifetime is strongly temperature dependent and below 193K can exceed 1 hour. Even in the Peltier regime, at 230K the lifetime is of the order of 10s [1]. So, at whatever is our chosen operating temperature, as long as we keep the exposure time below the hole lifetime, we should see a suppression of surface dark current. This was borne out by experiment. A series of dark current frames was taken at three different operating temperatures whilst operating an EMCCD in non-inverted mode. The exposure times were varied between 25s and 1000s. The results are shown in Figure 9. Note




Figure 9: Non-Inverted mode dark current as a function of exposure time and temperature

that at lower temperatures the dark current stays fairly constant with exposure time, as one would expect. The hole lifetime is several hours and the surface component of the dark current is entirely suppressed. At higher temperatures, closer to the Peltier cooler regime, the dark current actually increases as a function of exposure time. Here the hole lifetime is much lower and as they gradually disappear the surface dark current slowly increases.

One must be very careful when measuring dark current with a Peltier cooled CCD. The dark current in a long exposure may be more than a magnitude greater than that present in an exposure lasting tens of milliseconds.

2.9 Read-out noise

This is what limits the performance of a conventional CCD at the low-signal end of its dynamic range. The on-chip output amplifier of a CCD consists of at least two MOSFET transistors. One of these acts as a buffer amplifier transforming the pixel charge, stored on a capacitive measuring node, into a voltage. Once the measurement is performed this node is then reset through the second transistor. Both transistors contribute noise but through clever processing of the video signal (so called ‘correlated double sampling’ or CDS) it is possible to remove almost entirely that produced by the reset transistor. The noise on the buffer amplifier, however, is much more problematic. It consists of Johnson noise which originates in thermal motion of the electrons in the channel of the transistor. It has a white (i.e. flat) noise spectrum and its effects can only be removed by reducing the measurement bandwidth or by extreme cooling. Reducing the measurement bandwidth means increasing the time taken to read-out the CCD image so there is an inevitable trade-off to be made. Fast readout inevitably means high read-noise. For example, operating a CCD at 10MPixs^{-1} may incur a read-noise of $50e^-$ rms, whereas to reach $2\text{-}3e^-$ read-noise may require read-out at around 50kPixs^{-1} . One further noise source caused by thermally induced variations in the carrier concentration in the FET channel, known as ‘flicker noise’ imposes an ultimate noise floor. The spectrum of flicker noise has a $1/\text{frequency}$ distribution so further reductions in read-out speed can, whilst reducing the Johnson component, actually increase the total read-noise.

2.10 Photon noise

The arrival of photons in a detector is described by Poissonian statistics. The Poisson distribution describes the probability $P(M, n)$ that a pixel experiencing a mean illumination of M photo-electrons will receive a signal of n electrons. This is defined as follows:

$$P(M, n) = \frac{M^n \exp(-M)}{n!}. \quad (1)$$

The Poissonian distribution has the property that its variance is equal to its mean. What this means is that the number of observed photons fluctuates about its mean M with a standard deviation $= \sqrt{M}$. So, if we detect 100 photons from a faint astronomical object, the noise in our photometry will be 10 photons; we are unable to decide on its true brightness to any better than 10% accuracy. This is one consequence of the discrete nature of photons.

3 Key EMCCD concepts

EMCCDs have an effective read-noise that is so low as to be negligible ($0.025e^-$ would be typical). The output MOSFET amplifier of an EMCCD is, however, actually intrinsically quite noisy: 2 or 3 times that of a conventional CCD. The reason that they manage to achieve such low noise is that the photo-charge of each pixel is hugely amplified prior to being measured by this amplifier. A single photo-electron may have been amplified to $1000e^-$ prior to being measured. In that case the *effective* noise of the amplifier is reduced by a factor of 1000.

3.1 EMCCD gain structure

The structure of an EMCCD is actually very similar to that of a conventional CCD in terms of how the photo-charge is collected and transported. It is when the photo-charge arrives at the serial register that the main difference appears. In an EMCCD the serial register is greatly extended to include extra stages in which the multiplication gain occurs. These extra stages are known as the ‘multiplication register’. This register is connected to the same serial clock phases $\phi_{1,3}$ as drive the conventional part of the serial register. The multiplication register, however, shown as a shaded section in Figure 10, contains in place of ϕ_2 , an extra serial phase known as ϕ_{2HV} . Additionally there is an extra DC phase, called ϕ_{2DC} inserted between ϕ_1 and ϕ_{2HV} . This is shown in Figure 11. In the course of the readout, photo-charge transfers from ϕ_1 , passing through the small potential step produced by ϕ_{2DC} and then falls into the deep potential well beneath ϕ_{2HV} . In doing so it gains considerable energy and has a small probability of generating additional

electrons through impact ionisation. This probability rises steeply with the voltage on $\phi 2HV$.

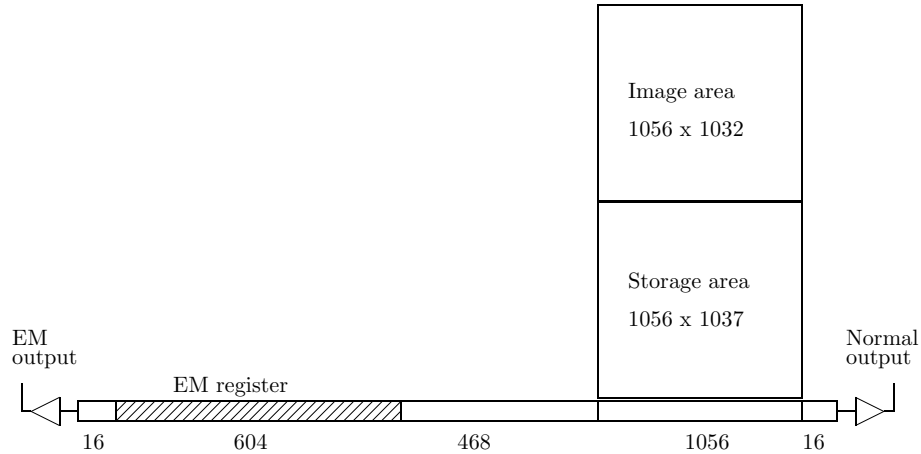


Figure 10: CCD201-20 gain structure. In the case of this EMCCD, the gain structure takes the form of an extended serial register which contains a 604 stage multiplication (or ‘EM’) register.

3.2 Output signal distribution from EM register

This has been studied in some depth by Tubbs [11]. He obtained the following result, showing that the probability of an output signal of n electrons resulting from a single input electron traversing the EM gain register with overall gain g was :

$$P(n) \begin{cases} = \left(\exp\left(\frac{1}{g-1/2}\right) - 1 \right) \exp\left(-\frac{n}{g-1/2}\right) & \text{if } n \geq 1 \\ = 0 & \text{otherwise,} \end{cases} \quad (2)$$

As long as the $g \gg 1$, then for the case of a single electron input we can approximate the output signal distribution by :

$$P(n) \begin{cases} = \frac{1}{g} \exp\left(-\frac{n}{g}\right) & \text{if } n \geq 1 \\ = 0 & \text{otherwise,} \end{cases} \quad (3)$$

If we plot a histogram of pixel values from a faintly illuminated EMCCD it should therefore have this form. Pixel values are normally quoted in ‘ADUs’, one ADU representing a single digital interval. Conveniently, if we plot log base e of the number of pixels in each histogram bin along the ordinate, the histogram (see figure 12.)




Figure 11: Serial-clock geometry of an EMCCD. Four consecutive register elements are shown straddling the join between the multiplication register and conventional part of the serial register. The graphs indicate the potential wells lying beneath the electrodes in which the photo-charge sits. Three consecutive clock states are displayed (t_1, t_2, t_3). Multiplication occurs when the voltage on ϕ_1 goes low and the charge falls into the deep potential well under ϕ_2HV .

will then have a large linear region whose gradient is equal to $-1 \times$ the system gain in e^-/ADU . This useful property can be exploited as a simple method of measuring the system gain.

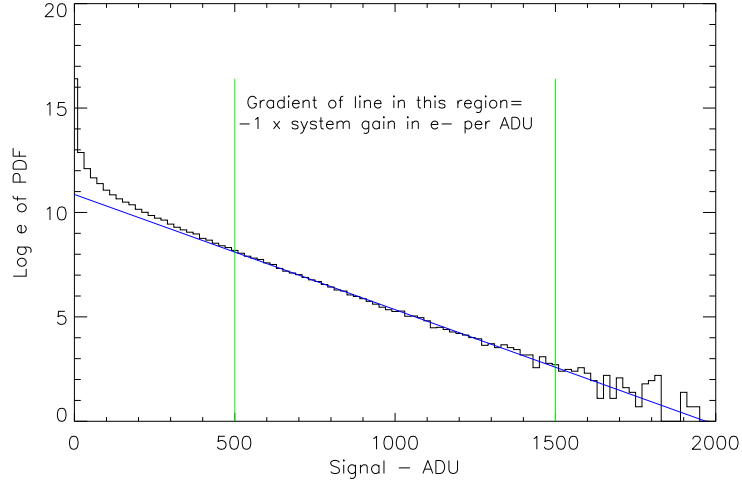


Figure 12: EMCCD pixel histogram used for gain calculation. The gradient of the histogram between the vertical blue lines is equal to $-1 \times$ the system gain in e^-/ADU .

Tubbs then extends his equation to a more general form to show the probability of an output signal n arising from an input of m electrons to the EM register :

$$P(n, m) \begin{cases} = \frac{(n - m + 1)^{m-1}}{(m - 1)! \left(g - 1 + \frac{1}{m}\right)^m} \exp\left(-\frac{n - m + 1}{g - 1 + \frac{1}{m}}\right) & \text{if } n \geq m \\ = 0 & \text{otherwise.} \end{cases} \quad (4)$$

This relation is illustrated graphically in figure 13 for inputs of between 1 and $5e^-$ to the multiplication register.

3.3 Clock-induced charge

In an inverted CCD the holes present in the surface layer clearly have an important role to play in reducing the dark current. Unfortunately, during the readout phase the presence of these holes has a drawback. During readout the parallel electrodes make fast transitions to their high states. This repels the holes that have previously populated the surface downwards into the bulk of the Silicon at high velocity where they generate stray electrons through impact ionisation. These electrons constitute




Figure 13: Output of multiplication register in response to a range of inputs from 1 to $5e^-$.

‘clock induced charge’ (CIC). In a conventional CCD the CIC is barely noticed since it is swamped by the generally much higher read-noise. In an EMCCD, the CIC can become the dominant noise source due to the effective absence of read noise. Figure 14 shows the appearance of CIC at three different levels. The CCD in question had a very high gain and the individual CIC electrons were clearly visible as discrete events. Note that these events have a large spread in height. The reason for this spread is ‘multiplication noise’ which is explained in the next section.




Figure 14: Appearance of CIC in cross-sections through the image. Three profiles are shown from 3 cameras experiencing a range of CIC levels.

3.4 In-register CIC

A well set-up EMCCD should be limited only by the CIC that is generated within the multiplication register. CIC events that appear close to the register's input end will receive almost the full multiplication gain. Those that appear close to the output will in contrast receive almost no EM gain. On average an electron generated within the EM register will produce a lower signal than a genuine photo-electron. This has important consequences if we use the CCD for photon counting (see section 3.8). The effective size of an 'in-register' CIC event can be calculated as follows.

The mean size, q_O , of a pixel charge packet containing 1 photo-electron of signal as it transits the EM register is given by:

$$q_O = (1 + p)^x, \quad (5)$$

Where x = position within EM register and p = per-transfer multiplication probability. The mean value, \bar{q}_O of this charge packet during its EM register transit can then be obtained by integrating this function over the length of the register and then dividing by the number of stages. If we reference this charge to an equivalent signal at the input to the register, such that $q_I = q_O/g$, where g = overall gain, then we get:

$$\bar{q}_I = \frac{1}{Ng} \int_{x=1}^N (1 + p)^x dx, \quad (6)$$

Where N =total number of stages within EM register. This is given by the standard integral:

$$\bar{q}_I = \frac{1}{Ng \ln(1 + p)} [(1 + p)^x]_{x=1}^N, \quad (7)$$

Which gives the result:

$$\bar{q}_I = \frac{(1 - 1/g)}{\ln g} \approx \frac{1}{\ln g}. \quad (8)$$

This shows that an electron that originates *within* the EM register effectively has a mean charge, when referenced to the input of the register (i.e. in units of photo-electrons), of $\approx 1/\ln g$ e^- . The total noise-charge that in-register CIC contributes to the output image (when expressed in equivalent photo-electrons) can therefore

be reduced by the use of high gain.

3.5 Multiplication noise

The multiplication process in an EMCCD has a low probability of actually producing gain. In moving from one element in the EM register to the next, an electron has approximately a 1% chance of creating an additional electron. The high gains we see are only achievable if we use multiplication registers containing several hundred stages. This distribution of gain throughout many stages has the effect of producing a large variance in the output signals resulting from single electrons entering the EM register. This uncertainty in the gain experienced by a photo-electron constitutes an additional noise source in an EMCCD and is known as multiplication noise.

This noise source has been studied in some depth by Robbins et al.[10]. They introduce a noise factor F . This is related to the standard deviation of the input signal σ_{in} , the output signal σ_{out} and the multiplication gain g as follows:

$$F = \frac{\sigma_{out}}{\sigma_{in}g}, \quad (9)$$

They then go on to derive :

$$F^2 = \frac{1}{g} + 2(g-1)g^{-\left(\frac{N+1}{N}\right)}. \quad (10)$$

where N =number of stages within EM register. It can be shown that as N and g become very large (as they do in a practical camera system) that F tends to $\sqrt{2}$. This has the effect of multiplying the noise in the input image by $g\sqrt{2}$. Since the signal will only be multiplied by g the overall effect is to reduce the signal to noise ratio in the image by a factor of $\sqrt{2}$. Statistically this is equivalent to a halving of the quantum efficiency of the camera. This important result was tested in the laboratory by using an EMCCD camera to take a series of flat field frames of increasing intensity. At each signal level the mean and the variance of the flat fields was measured. The sequences were repeated at a total of three different EM gain levels. The results are shown in figure 15. As can be seen the variance in the EM

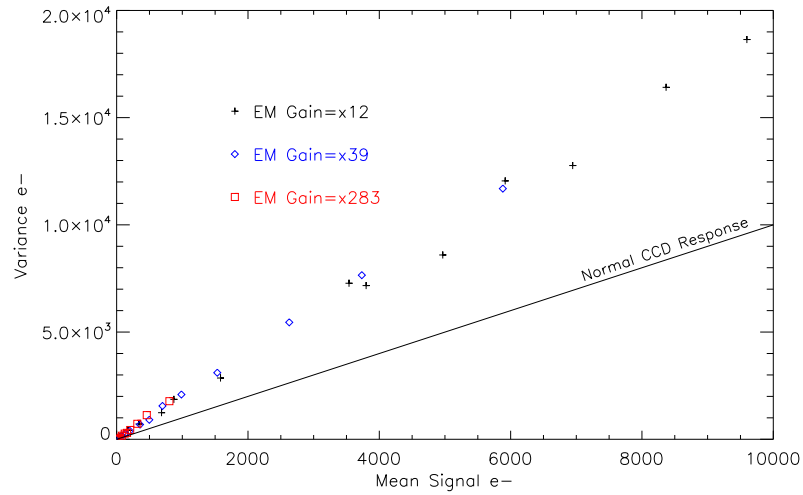


Figure 15: Variance as a function of mean signal level for an EMCCD as opposed to a conventional CCD. Three distinct EM gains were investigated. The variance in the EM frames is higher by a factor of 2 due to the effects of multiplication noise.

frames was equal to $2\times$ their mean, thus confirming the earlier theoretical result. The effects of multiplication noise have in fact already been shown earlier in figure 13. Note that for an output signal of 300 e^- and an EM gain of $\times 100$ the input signal could have been either 3 or 4 e^- with almost equal probability.

To help illustrate multiplication noise further a series of synthetic images, shown in figure 16, was generated. The top row of images was obtained using a conventional CCD model, incorporating the effects of Poissonian noise and read noise. The second row was obtained using an EMCCD model incorporating Poissonian noise, multiplication noise and clock induced charge. The bright areas of these synthetic images had mean illuminations of 0.1 , 1 , 5 and 25 e^- . The read noise of the conventional CCD was set to 5 e^- and the CIC level of the EMCCD was 0.03 e^- per pixel. At very low illuminations the conventional CCD sees nothing whereas at the higher end the effects of multiplication noise mean that the conventional detector actually does better than its electron multiplying counterpart.




Figure 16: Comparison of an EMCCD and a normal CCD of $5e^-$ read-noise. Signal increases to the right. At low signal levels the EMCCD is greatly superior. At higher signal levels the conventional detector produces less-noisy images.

3.6 Normal or Conventional mode

If the CCD has a conventional output (i.e. a normal non-EM gain amplifier is used to measure the photo-charge) in addition to the EM output, as is the case with the CCD201-20 [4], then we can use the device with great flexibility. If the signal is high, and therefore not read-noise dominated then we can use this output to improve our SNR. The detector then behaves as a normal scientific CCD. As there is no EM gain, neither is there any multiplication noise. The standard equation for calculating the SNR is shown in equation 11 :

$$\text{SNR}_N = \frac{M}{\sqrt{(M + \nu_C) + \sigma_G^2}}. \quad (11)$$

Where M =mean photo-charge per-pixel, ν_C =mean CIC-charge per-pixel and σ_G =read-noise in the output amplifier expressed in e^- RMS.

3.7 Proportional or Linear mode

If we then switch to the EM output of the CCD we add multiplication noise and at the same time reduce the read-noise by the EM multiplication factor g . The achievable SNR is shown in equation 12. The factor of 2 in the denominator accounts

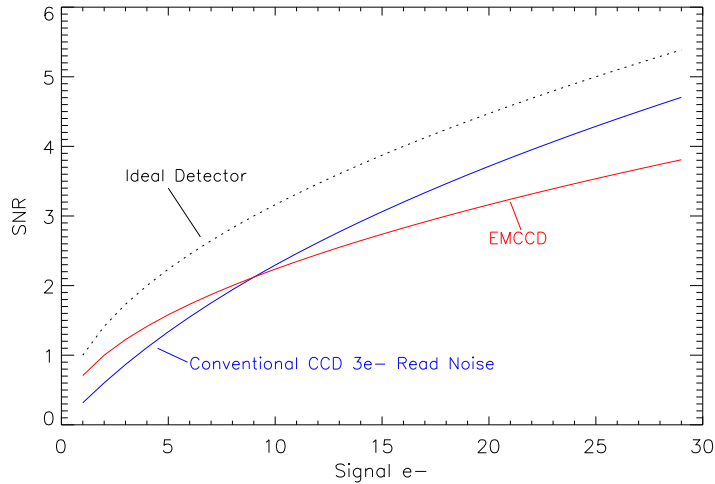


Figure 17: Signal to noise ratio as a function of signal level. The SNR of an EMCCD is compared with a conventional low noise CCD and an ideal detector. The amplifier read noise is $3e^-$.

for the multiplication noise (see section 3.5).

$$\text{SNR}_L = \frac{M}{\sqrt{2(M + \nu_C) + (\sigma_G/g)^2}}. \quad (12)$$

There is therefore a cross-over point above which the SNR is higher with a conventional detector. This is shown in figure 17.

3.8 Photon counting mode

This involves setting a threshold and interpreting all pixel values that lie above it as representing 1 photo-electron. All pixels below the threshold are then interpreted as having no signal. Since a pixel is either equal to one or zero photo-electrons, the effects of multiplication noise are eliminated. The relevant SNR equation is shown in simplified form in equation 13. It is assumed that all the events are counted and that the amplifier noise is negligible. Note that a CIC term is included but its units are different to those used in equation 12. Now the CIC, B_C , needs to be quoted in units of mean CIC counts per pixel.

$$\text{SNR}_{PC} = \frac{M}{\sqrt{(M + B_C)}}. \quad (13)$$

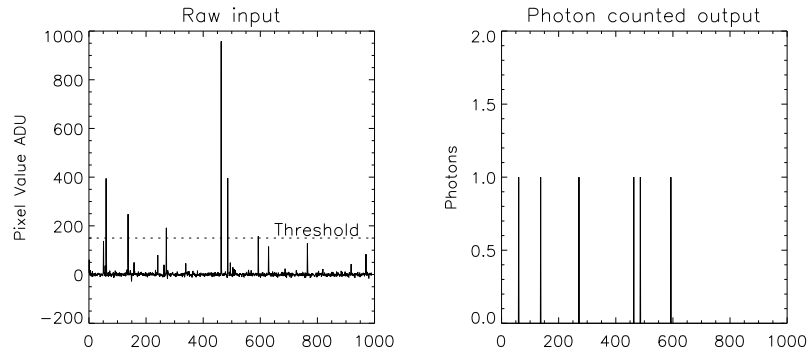


Figure 18: Photon counting method using a pixel value threshold.

Given the fact that single photo-electron events have a distribution heavily skewed to lower values it is necessary to set the threshold as low as possible whilst at the same time ensuring that it does not go so low as to produce false triggers from pixels containing just the Gaussian read-out noise. As the threshold rises the proportion of ‘lost’ photo-electrons steadily rises and the effective QE of the photon-counting system falls. We can calculate this effect by integrating equation 3. The probability of a photo-electron being counted $P_d(T)$, where T =threshold setting in electrons at the output of the EM register, is then given by:

$$P_d(T) = \frac{1}{g} \int_{n=T}^{\infty} \exp\left(-\frac{n}{g}\right) dn = \exp\left(-\frac{T}{g}\right), \quad (14)$$

It is perhaps easier if we express the threshold as a fraction of the mean size of a photo-electron :

$$P_d(t) = \exp(-t) \text{ where } t = T/g. \quad (15)$$

So if a threshold of $0.1e^-$ is chosen then the probability of a photo-electron event being counted is 0.905.

Photon-counting will only work in the low signal regime. If there is a significant probability of a pixel receiving two photons within a single exposure then only one of these will be photon-counted. This effective loss of sensitivity is called coincidence loss and is described by the following equation :

$$F(M) = \frac{1 - \exp(-M)}{M}. \quad (16)$$




Figure 19: Demonstration of the wave/particle nature of photons. The panel on the left shows a single input frame. Shown on the right is the sum of 50 thresholded input frames. The Airy rings, a wave phenomena, are shown composed of discrete photon events.

Here $F(M)$ is the detected fraction of photons and M the mean number of photons per pixel expected during the integration. So for $M=0.1$ the equation shows us that we will count at most 95% of the photons.

Earlier, in equation 8, it was shown that a CIC electron generated within the EM register has a lower significance than a genuine photo-electron, at least in terms of the noise charge it contributes to the final image. In photon-counting mode, however, it will have exactly the same significance. On average an in-register CIC event will be of lower amplitude than a genuine photo-electron but as long as it remains above the photon-counting threshold it will still be counted. The overall effect of this is to decrease the effectiveness of photon counting as opposed to linear mode operation at least in the case of cameras limited by in-register CIC. Note that in the earlier SNR equations (equations 12,13) that the CIC was expressed in two separate units : B_C for photon-counting and ν_C for linear mode. These are related as follows, where g =EM multiplication factor :

$$B_C \approx \nu_C \ln(g). \quad (17)$$

An interesting example of photon counting is shown in Figure 19.¹ Here an EM camera was set up in the laboratory with a $50\mu\text{m}$ pinhole placed a small distance in

¹A movie of this is available at <http://www.ing.iac.es/~smt/WFS/counting/movie.htm>

front of the CCD. This pinhole was illuminated by a very faint LED approximately 50cm distant. No lens was required in the setup. Diffraction through the small pinhole then projected an Airy diffraction pattern onto the CCD. The LED intensity was varied, along with the exposure time, so that the brightness of the Airy rings was suitable for a photon counting analysis of the images i.e. mean intensity smaller than $0.1e^-$ so as to avoid significant coincidence losses. A series of 50 frames were then taken and their photon counted sum combined into a single image. This is shown in the right hand panel of the figure. On the left is one of the 50 raw input frames. This gives a nice demonstration of the dual nature of the photons: wave-like when passing through the pinhole but particle-like when detected by the CCD.

3.9 System gain and sensitivity

Astronomers typically refer to the gain (or system gain) of a CCD camera as meaning the number of photo-electrons represented by 1 ADU in the raw image and therefore it has units of e^-/ADU . This can be somewhat confusing since if the numerical gain of the video amplifier in the camera controller is increased, the e^-/ADU in the output image actually *decreases*. The use of these units is so ingrained (even amongst CCD engineers who should know better!) that it was thought best to stay with this convention in this technical note. System gain has to be well chosen so that at the low signal end of the camera's dynamic range the size of a single ADU is small compared to the read-noise. Typically the gain is adjusted such that the read-noise is around 3 ADU. This prevents the quantisation noise of the analogue to digital converter (ADC) from degrading the SNR of the system. In a well designed camera the upper end of the dynamic range, ultimately limited by the range of the ADC, should then match well to the range of illumination over which the CCD will be linear before blooming sets in. A typical value for the gain would be $1e^-/ADU$ for a CCD with $3e^-$ read-out noise. With a 16bit ADC, which is the standard, the upper end of the dynamic range would then be $65Ke^-$. This is rather less than the blooming limit of most scientific CCDs. In order to be able to extend the camera operation to the high-signal regime most systems therefore allow switchable gain:

low gain for faint objects and high gain for brighter objects. Most video processors contain an integrator. Conveniently, these give more e^-/ADU if they are operated at higher speeds. They also give more noise. The gain switch mechanism therefore normally just involves changing the integrator speed : the faster the integration the higher the gain in e^-/ADU and the higher the read-noise. These overall effect is that the image noise measured in ADU stays reasonably constant. Since the higher gains will only be used for measuring bright signals, the fact that they are noisier is of no consequence: the images will be photon-noise dominated anyway.

An EMCCD camera has yet another gain parameter that we need to describe, that of the multiplication gain factor g and there is a risk of confusion here with the e^-/ADU gain of the complete camera system. If we measure the overall gain of an EM camera then the value we arrive at will be a combination of voltage gain in the video processor and the multiplication factor of the EM register. Our measurement of e^-/ADU gain will not tell us directly what the EM multiplication factor is. The most convenient way to measure this important parameter is to temporarily turn off the EM multiplication gain and measure the camera e^-/ADU gain. This result is then described as the video processor sensitivity S_V . For the conventional output of the CCD, the system gain is the same as S_V for that amplifier. In the case of the EM output the system gain = S_V/g .

4 Astronomical applications

When it was first announced it was clear that EMCCD technology was a solution looking for a problem. They were first intended for military and security imaging applications where moonlight and even starlight could provide sufficient illumination to yield a useful image. At an astronomical observatory, astronomers frequently complain that the read-noise of their detectors is too high. A noiseless detector should surely be of interest even with the presence of multiplication noise. In the 1980s the IPCS (Image Photon Counting System) remained in demand long after CCD detectors with their vastly superior QE were made available. The fact that IPCS had zero read-noise (due to an internal avalanche gain mechanism) more than

made up for its 5% QE (it was based on vacuum tube technology) in certain applications. Read-noise dominated regimes are typically encountered when faint sources are observed at high spectroscopic resolution in the absence of moonlight. Here, an already faint signal is spread out by the spectrograph over many pixels on the detector and the per-pixel signal can easily be just a few photons or less. EMCCDs are greatly superior to the IPCS but the situation is the same in that certain astronomical observations are completely dominated by read-noise and a 50% drop in effective QE due to multiplication noise in exchange for \approx zero read-noise is a price that many astronomers are willing to pay.

The first people to start exploiting EMCCDs in astronomy were Craig Mackay's group at the University of Cambridge. They used a CCD97 on the Nordic Optical Telescope for high-speed imaging. Their technique, called 'lucky imaging' [9], used the high speed capabilities of EMCCDs to produce high resolution images (in some cases diffraction limited) during periods of low atmospheric turbulence. These occur in a random fashion and for short periods only, interspersed with longer periods of poorer atmospheric conditions. By taking many images at high speed (25Hz) and then adding together only those images that had the highest resolution, they were able to offer a way of removing to a large extent the blurring effects of atmospheric turbulence. A few years later the European Southern Observatory (ESO) decided to use EMCCD technology for their future high-speed wavefront sensor cameras [3]. These cameras will be used to measure at high speed ($>1\text{KHz}$) the wavefront distortions induced by the atmosphere on the light of a faint star. These measurements are then used to control, in real time, the shape of a deformable mirror that compensates for these distortions. The combination of low signal and high frame rate makes EMCCD technology the logical choice in this application. An EMCCD based wavefront sensor has been in use since 2005 at the Isaac Newton Group (ING) in La Palma [12]. The ING also offers two additional EMCCD cameras dedicated to spectroscopy. Observations of short-period binaries would appear to be a valuable niche for these cameras [13]. At least two other groups have also built EMCCD cameras intended for astronomy: Olivier Daigle's at the Université de Montréal [2]

and Derek Ives at the UKATC [6]. Several commercial cameras are also available, aimed at the bio-analysis market. Of these Andor is the best known.




Figure 20: An EMCCD at the ING designed for rapid spectroscopy. The 1k x 1k detector is shown mounted in a liquid Nitrogen cryostat.

References

- [1] B. E. Burke and S. A. Gajar. Dynamic suppression of interface-state dark current in buried-channel CCDs. *IEEE Transactions on Electron Devices*, 38:285–290, February 1991.
- [2] O. Daigle, J.-L. Gach, C. Guillaume, S. Lessard, C. Carignan, and S. Blais-Ouellette. CCCP: a CCD controller for counting photons. In *SPIE*, volume 7014 of *Conference Series*, August 2008.
- [3] M. Downing, G. Finger, D. Baade, N. Hubin, O. Iwert, and J. Kolb. Detectors for AO wavefront sensing. In *SPIE*, volume 7015 of *Conference Series*, July 2008.
- [4] E2V. Ccd201-20 backside illuminated 2-phase imo series emccd. Data sheet, E2V Technologies, 2000.

- [5] J. Hyneczek. Impactron—a new solid state image intensifier. *IEEE Transactions on Electron Devices*, 48:2238–2241, October 2001.
- [6] D. Ives, N. Bezawada, V. Dhillon, and T. Marsh. ULTRASPEC: an electron multiplication CCD camera for very low light level high speed astronomical spectrometry. In *SPIE*, volume 7021 of *Conference Series*, August 2008.
- [7] P. Jerram, P. J. Pool, R. Bell, D. J. Burt, S. Bowring, S. Spencer, M. Hazelwood, I. Moody, N. Catlett, and P. S. Heyes. The LLCCD: low-light imaging without the need for an intensifier. In M. M. Blouke, J. Canosa, and N. Sampat, editors, *SPIE*, volume 4306 of *Conference Series*, pages 178–186, May 2001.
- [8] P. R. Jordan, P. Pool, and S. M. Tulloch. Secrets of E2V Technologies CCDs. In P. Amico, J. W. Beletic, and J. E. Beletic, editors, *Scientific Detectors for Astronomy, The Beginning of a New Era*, volume 300 of *Astrophysics and Space Science Library*, pages 115–122, 2004.
- [9] N. M. Law, C. D. Mackay, and J. E. Baldwin. Lucky imaging: high angular resolution imaging in the visible from the ground. *AAP*, 446:739–745, February 2006.
- [10] M. S. Robbins and B. J. Hadwen. The noise performance of electron multiplying charge-coupled devices. *IEEE Transactions on Electron Devices*, 50:1227–1232, May 2003.
- [11] R.N. Tubbs. PhD thesis, University of Cambridge, 2003.
- [12] S. Tulloch. L3 CCD Wavefront Sensor Developments at the ING. In J. E. Beletic, J. W. Beletic, and P. Amico, editors, *Scientific Detectors for Astronomy 2005*, pages 303–+, March 2006.
- [13] S. M. Tulloch, P. Rodriguez-Gil, and V. S. Dhillon. Radial-velocity study of the post-period minimum cataclysmic variable SDSSJ143317.78+101123.3 with an electron-multiplying CCD. *ArXiv e-prints*, May 2009.