Proceedings of 33rd Annual European Meeting on Atmospheric Studies by Optical Methods.
IRF Sci. Rep., 292, 61–66, 2008
http://www.irf.se/publications/proc33AM/mcwhirter.pdf
ISSN 0284-1703, ISBN 978-91-977255-1-4, © Author(s) 2008

Electron Multiplying CCDs - New Technology for Low Light Level Imaging

I. McWhirter

Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT

Abstract. Low light level imaging has been revolutionised in recent years with the introduction of Electron Multiplying Charge Coupled Devices (EMCCDs). The performance of conventional CCDs has always been limited by readout noise, inherent in the output amplifier of the device. In order to minimise this, the readout speed has to be relatively slow, typically taking several seconds to read the image. EMCCDs incorporate an additional gain register through which the signal passes as it is read out of the CCD array. Because the signal is amplified prior to readout, the signal-to-noise ratio is much improved and the restriction of slow readout times is also lifted. In spite of the promised improvements in performance, they have been received by the scientific community with a certain amount of caution. The advantages of EMCCDs, compared to earlier generations of imagers, are described and the specific issues relating to their operation are discussed

1 Introduction

Low light level imagers have always been an essential tool in the field of atmospheric research by optical methods. Until recently there were two main contenders, bare CCDs and intensified CCDs (ICCDs).

1.1 Bare CCDs

With these devices the light is imaged on the face of the CCD chip itself. They have the advantages of:

- Highest and broadest Quantum Efficiency curve.
- Excellent resolution.
- Very low dark signal with sufficient cooling.

However, they suffer from noise inherent in the process of reading out the signal. Even in the best devices, this is typically 2 - 4 electrons RMS with a slow readout rate (several seconds to read the entire chip) and much more if the readout amplifier bandwidth is increased to give a fast readout.

1.2 Intensified CCDs

The readout noise limitation can be overcome by boosting the light level before it reaches the CCD, using a vacuum-tube image intensifier. In this case the light is focused onto a light-sensitive photocathode and the resulting photoelectrons directed into an electron multiplier before hitting a phosphor screen to give a much brighter image than the original. This brighter image is transferred to the face of the CCD, usually with a fibre-optic coupler and the resultant signal is sufficiently high that the background noise becomes insignificant. However:

- The photocathode has lower QE than that of a bare CCD, over a narrower range of wavelengths.
- The vacuum tube intensifier has limited life and is easily damaged.
- There is additional background noise in the form of thermionic emission from the photocathode.
- The electron gain mechanism is statistical and gives rise to an additional noise factor.

2 The architecture of the Electron Multiplying CCD

The status quo had been established for some years and the above compromises accepted among the community. When, in the early part of this decade, it was announced that a bare CCD chip had been produced that had no effective readout noise, the news was received with a certain amount of scepticism. The new device was the Electron Multiplying CCD

Correspondence to: I. McWhirter (i.mcwhirter@ucl.ac.uk)

Published by the Swedish Institute of Space Physics.

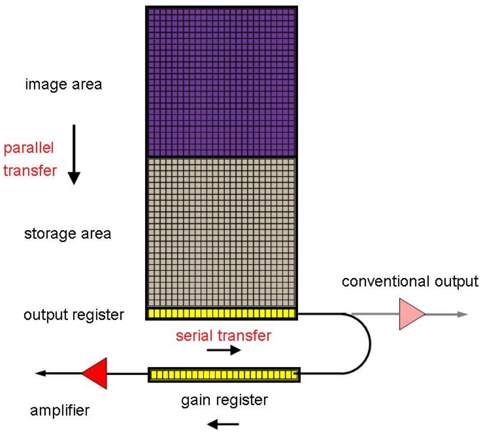

Fig. 1. The EMCCD architecture: the image is accumulated in the upper half of the CCD and then rapidly transferred to the storage area from where it is read out, one line at a time, into the gain register and output amplifier.

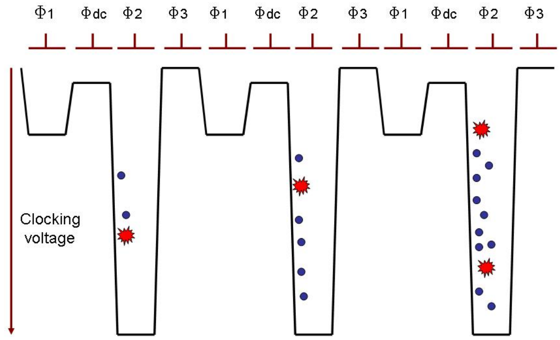

Fig. 2. The electrons are moved from one cell to the next by the clocking pulses, Φ1, Φ2 and Φ3, which are applied to the surface electrodes. The electron multiplication effect, which occurs as a result of the enhanced Φ2 pulse, is highly exaggerated in the diagram for clarity.

(EMCCD). Although it is a bare CCD, it has a gain mechanism built into the chip itself which amplifies the signal before it enters the readout amplifier. Fig. 1 shows the architecture of a frame transfer CCD, the format typically used for EMCCDs.

The charge generated by incident light is read out by shifting it with clock pulses which are applied to electrodes on the chip. The charge is passed from one cell to the next in 'bucket brigade' fashion. For a frame transfer CCD the electrons are first shifted into a storage area and then read out one line at a time through the serial register into the readout amplifier. However, in an Electron Multiplying CCD, the charge is also passed through a second register, known as the gain register. This consists of an additional row of serial readout elements which are operated at a high clock voltage to encourage the phenomenon of impact ionisation. As the electrons are passed from one element to the next, a small number of additional electrons are generated at each transfer. The statistical gain is small, in the region of 1.01 per element. However, the serial register of a typical device consisting of n elements has a gain of 1.01^n. A typical array where n is 512 or 1024 will, therefore, have a gain of several hundred for the entire row. Fig. 2 illustrates the gain mechanism. CCDs commonly use three phase clocking. There are three gates per cell which each receive different clock pulses, Φ1, Φ2 and Φ3. In EMCCDs an extra fixed-potential gate is added and a much higher voltage than usual (35 - 50v) is used to drive Φ2.

3 The signal-to-noise ratio of a CCD

In order to quantify the effect of the gain register it is necessary to examine the signal-to-noise ratio (SNR) of the CCD, which is given by the equation:

$$SNR = \frac{P.QE}{\sqrt{P.QE + N_{DARK}^2 + N_{READOUT}^2}} \quad (1)$$

The magnitude of the signal is given by $P.QE$ where P is the number of photons incident on a picture element, or pixel, of the CCD and QE is the quantum efficiency of the CCD, the fraction of incident photons that are converted into electrons. The noise has three components:

- Poisson or $\sqrt{N}$ noise. This is a consequence of the statistical nature of any signal made up of discrete elements, in this case electrons, and is equal to $\sqrt{P.QE}$.
- Dark noise - background signal created within the CCD chip due to thermal effects.
- Readout noise - electronic noise generated during the readout process.

The total noise is given by the square root of the sum of the squares. The dark noise decreases dramatically with cooling, so CCD detectors are usually cooled in order to minimise this. The $\sqrt{N}$ noise is minimised by using techniques such as back-thinning to ensure that the QE of the device is as high as possible and then simply by taking a sufficiently long exposure. The most significant source of noise is, therefore, the readout noise. In a conventional CCD, the best case scenario

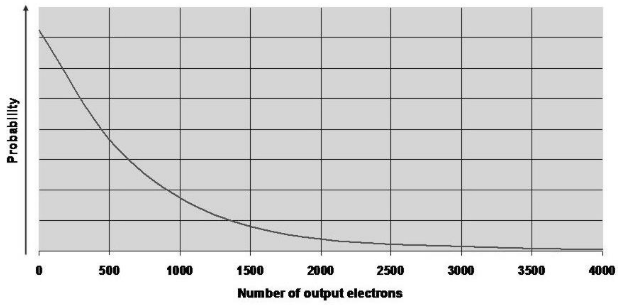

Fig. 3. This shows the trend in probability for a given number of output electrons to be generated from one input electron. The mean gain in this example is in the region of 600.

is a noise contribution of 2 - 4 electrons RMS, obtained at the expense of a slow readout time of several seconds.

4 The effect of the gain mechanism

As we have seen, the gain mechanism amplifies the signal prior to the output amplifier. It appears in the revised equation (2) as a factor G, affecting everything apart from the readout noise. With a sufficiently high gain, usually several hundred, the readout noise can be ignored, even at high readout rates, and so it is removed from the equation. This means that the restriction of slow readout is no longer a hindrance for high time resolution measurements. However, the noise factor F now also appears in the equation:

$$SNR = \frac{G.P.QE}{\sqrt{G^2.F^2(P.QE + N_{DARK}^2)}} \quad (2)$$

The noise factor is a result of the statistical gain process. Fig. 3 shows the trend in probability for an input of one electron. The curve has a form similar in shape to a negative exponential. The mean value of the gain is in the region of 600, but the stochastic nature of the mechanism means that this additional noise is introduced. According to theoretical analysis, (Mackay et al , 2001), F has a value of $\sqrt{2}$ but subsequent experimental measurements are typically lower and a value of 1.3 is now commonly stated. It should be remembered that a similar noise factor is also present in the noise equation relating to intensified CCD detectors because of the image intensifier tube. It is even higher if the microchannel plates of the intensifier have an ion barrier film to protect the photocathode. Generally accepted values for F are:

Conventional bare CCD	1.0
EMCCD	1.3
Intensified CCD (unfilmed)	1.5 - 2.0
Intensified CCD (filmed)	2.0 +

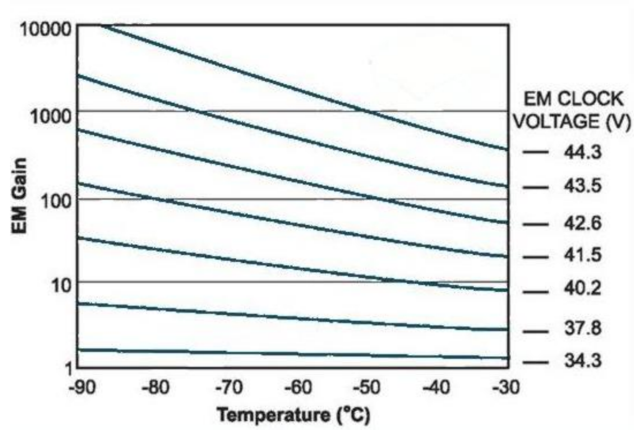

Fig. 4. The variation of electron multiplying gain with clock voltage and temperature (Source: Andor Technology Ltd).

5 Temperature dependence of the gain mechanism

One important feature of the gain mechanism is that it is highly temperature dependent. This is due to energy loss by phonon scattering, which increases with temperature. The effect is shown in Fig. 4. The enhanced cooling required to minimise dark signal does, therefore, offer a further advantage in the increased gain that it provides. The majority of applications will demand that the temperature of the device is stabilised in order to provide a consistent gain.

6 Reducing the dark noise component

Because the dark noise is amplified by the gain mechanism it now represents the dominant form of noise in the device. It consists of two components: the thermal noise and the clock-induced charge (CIC). The latter is generated by the clocking process which moves the signal from one cell to the next during readout (E2v Technologies Ltd , 2003).

6.1 Thermal noise

There are two ways to reduce the thermal component of the dark signal: cooling and a process known as ’inverted mode’ operation (IMO).

6.1.1 The effect of cooling

CCDs have traditionally been cooled to around -50°C to make the thermal noise component small when compared to the readout noise. Now that the effect of the readout noise has been practically eliminated this is no longer enough. In order to minimize the thermal noise to an insignificant level, deep cooling down to -85°C or lower is needed. In order to obtain this degree of cooling it is necessary to house the chip in a hermetically sealed, evacuated enclosure. Because the gain is also temperature dependent, it is essential to provide not just cooling but also accurate temperature stabilization. The improved engineering required for this has led to a big improvement in reliability when compared to conventional CCD cameras, which have usually been housed in o-ring sealed chambers filled with dry air. Many conventional CCD camera failures are due to corrosion of the delicate bonding wires as a result of ingress of water vapour through ineffective seals.

6.1.2 Inverted mode operation

The thermal signal is caused by generation of electron-hole pairs. It has two parts, the surface-generated component and that arising from within the bulk of the material:

$$S_{Total} = S_{Surface} + S_{Bulk} \tag{3}$$

A process known as inverted mode operation (IMO), in which the clock voltages are driven more negative, floods the surface interface with holes, inhibiting electron-hole pair generation. This eliminates the surface component. By this process, thermal noise is reduced by at least two orders of magnitude. The process is also known as Advanced Inverted Mode Operation (AIMO) or Multi-Pinned Phase (MPP) operation and is now standard for most scientific CCDs.

6.2 The clock induced charge

The CIC is caused by impact ionisation within the main body of the chip. This is exactly the same process that provides the electron multiplication in the gain register. In the latter it is encouraged by using high clocking voltages but when generated in the image and store registers it is a noise source to be minimised. In fact, it has always been present in conventional CCDs but because it is insignificant compared to the readout noise it has largely been ignored. Unfortunately, because of the higher clock voltages involved, it is significantly greater when inverted mode operation is used. CIC occurs during the parallel transfer into the readout register. It is usually expressed in electrons generated per pixel per parallel transfer. For example, 1056 parallel transfers are required to read out a 512x512 pixel chip such as the e2v CCD87. Fig. 5 shows a graph of CIC versus the parallel clock amplitude for the CCD87 chip. This shows a level of CIC of 10^4 e^-/pixel/transfer for non-inverted mode, which gives an overall value of 0.1e^-/pixel/frame. However, since these data were published, advances in manufacturing techniques have improved matters by up to a factor of 10, so the figures may now be regarded as pessimistic. It is seen from fig. 5 that the situation for non-inverted mode devices is significantly better by a factor of about 30.

There are three main ways to minimise the CIC:

- Keep the parallel clock amplitude to a minimum.
- Increase the parallel clock frequency.
- Slow down the rising edge of the parallel clock.

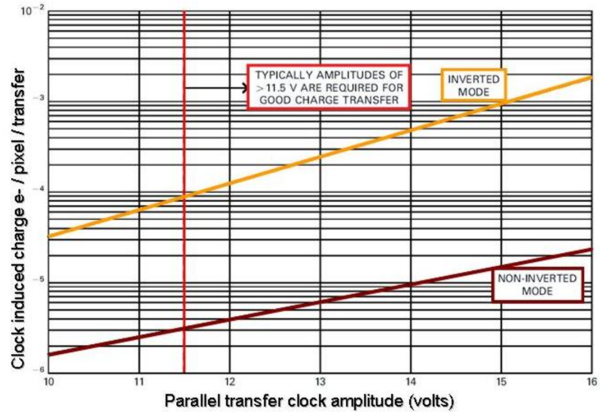

Fig. 5. Comparison of CIC levels for inverted and non-inverted mode operation for the CCD87 (Source: E2v Technologies Ltd , 2003).

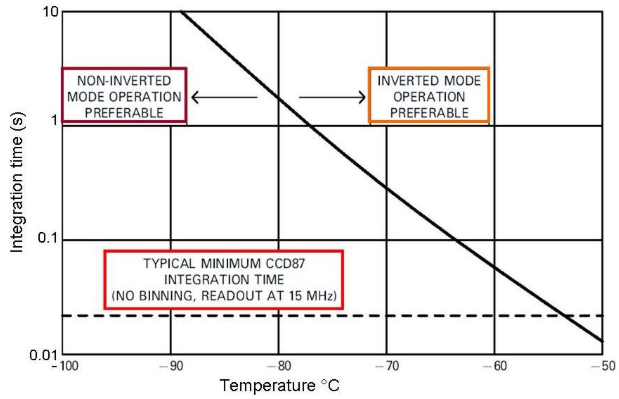

Fig. 6. At lower temperatures and shorter exposure times lower noise is obtained by operating in non-inverted mode (Source: E2v Technologies Ltd , 2003).

Unfortunately, the requirements are conflicting. For example, faster parallel transfer rates require higher clock amplitudes for efficient transfer. It is clear that if the thermal noise is low, due to good cooling or short exposure times, it may be advantageous to operate in non-inverted mode. Fig. 6 shows the preferred regions for both modes of operation.

7 Ageing effects in the gain mechanism

Ageing processes in the gain register mean that, for a fixed clock voltage, the gain falls as the device is run. There are several mechanisms involved, but generally it is thought to be caused by an accumulation of embedded charge in the insulating layer between the electrodes and the active silicon. It is a function of both gain and illumination. The gain can be recovered simply by increasing the voltage of the gain register, although that can only be done up to the point at which reliable operation ceases - generally an increase of 4 to 5 volts (E2v Technologies Ltd , 2006). The ageing mechanism has short term and long term components. The effect of the short term components, with time constants of up to several hundred hours, can be minimized for the end user by pre-ageing the chip before it is released for camera manufacture. The long term effects have a time constant of several thousands of hours and so a long device life can be expected. Simple precautions serve to minimize this effect. The user should only use the amount of gain necessary to provide the required signal-to-noise ratio and the chip should not be left running in high gain mode when data are not required. Camera manufacturers' software now includes routines for measuring the gain and re-calibrating the gain mechanism.

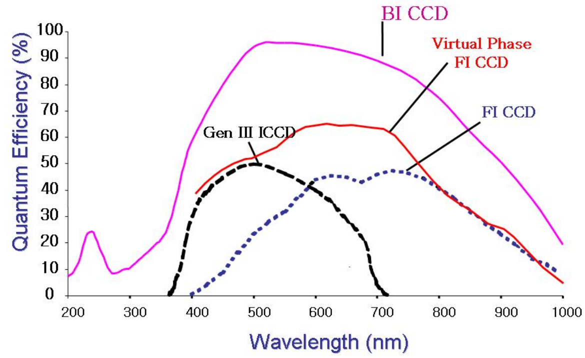

Fig. 7. The quantum efficiencies are compared for back illuminated, front illuminated and virtual phase CCDs together with the response of a GenIII image intensifier photocathode.

8 Quantum Efficiency

CCDs have far superior quantum efficiency when compared to the photocathodes of image intensifiers. However, the gate structure on the face of a conventional 'front illuminated' chip obstructs some of the active area and so reduces the QE. The peak QE of such a device is in the region of 45%. This effect can be avoided by thinning the back of the chip to the point at which it becomes transparent and then illuminating it from behind. This yields a QE of up to 95%. However, this is a mechanical process and consequently the cost of the chip is substantially increased. An alternative approach is used by Texas Instruments, with their 'Virtual Phase' chip, in which only one electrode is physically outside the silicon and a specially doped layer within the silicon performs the function of the remaining conventional electrodes. This method obscures much less of the face of the device and a QE of up to 65% can be obtained, without the cost penalty of back thinning. Fig. 7 shows typical QE curves for the different devices, together with the response of a GenIII photocathode for comparison.

9 EMCCD manufacturers

Although the chip manufactures also market their own cameras, the most sophisticated cameras come from separate companies who have specialised in making camera systems.

9.1 The chip manufacturers

There are two manufacturers of Electron Multiplying CCDs. The techniques are heavily patented and it is unlikely that there will be any more for the foreseeable future. They are e2v Technologies in the UK, who market their chip as L3Vision and Texas Instruments (TI) in the USA who produce the Impactron. The chips produced by e2v are back thinned to give the highest QE and consequently are more expensive than the virtual phase chips made by TI.

9.2 The camera manufacturers

The stringent requirements for getting the best performance from EMCCDs mean that the range of suppliers is more limited than for conventional CCDs. In Europe, Andor Technology is regarded as the most prominent manufacturer and in the USA, Roper Scientific, which includes Princeton Instruments and Photometrics. Hamamatsu of Japan has also emerged as a main contender. More companies are beginning to offer EMCCD cameras but the above companies, experienced in the technology involved in manufacturing the more sophisticated camera housings have so far dominated the market. The latest cameras offer Peltier cooling down to -100°C and support a range of chip options with pixel readout rates of up to 10MHz (L3Vision) or 35MHz (Impactron).

10 Comparisons with ICCDs

The problems that are presented by EMCCDs have meant that some experimenters have been reluctant to make the transition from the ICCDs that they know and love. Furthermore, not surprisingly, there have been strenuous efforts on the part of some ICCD manufacturers to show that EMCCDs are no better than ICCDs. However, these have not, in the opinion of the author, stood up to close scrutiny. In fact, most of the shortcomings of EMCCDs are also applicable to ICCDs:

- Ageing: Image intensifier photocathodes are also subject to ageing and, moreover, they can be destroyed by exposure to bright light to which EMCCDs are immune. The EMCCD ageing process is now well understood and can be minimised and compensated for to a great extent.

- Noise factor: ICCDs are also subject to a noise factor, which is, in fact, worse than that of EMCCDs.
- Clock-induced charge: Whilst this may be considered as the limiting factor for EMCCDs, it is substantially smaller than the total noise of a conventional CCD and smaller than the noise associated with image intensifiers. Ongoing improvements in chip design and manufacturing techniques are successfully reducing it.
- Reliability: The chips themselves are no less rugged than conventional CCDs. However, because of the more stringent cooling requirements they are usually housed in evacuated 'sealed for life' chambers which are considerably more reliable than some conventional CCD housings.

The one area in which EMCCDs cannot replace ICCDs is where fast gating of the light signal is required. The photocathode of an image intensifier can be biased off in just a few nanoseconds so that light generated throughout a very short time interval can be measured. This mode of operation is not possible for a non-intensified CCD.

11 Conclusions

The EMCCD has revolutionised low light level imaging. Its shortcomings are now well understood and ongoing work is successfully reducing the effects of clock induced charge and ageing. Its ruggedness and immunity to damage by high light levels make it an excellent choice for instruments that routinely operate unattended, such as all-sky cameras. The capability of high speed readout makes it ideal for the high time-resolution auroral studies, an area of research which is becoming increasingly important. The author has used these devices in the field since they were first commercially available and is of the opinion that the claims of their superior performance are fully substantiated.

Acknowledgements. The author gratefully acknowledges inclusion in this paper of data supplied by Andor Technology Ltd. and e2v Technologies Ltd.

The editors thank one referee for assistance in evaluating this paper.

References

Andor Technology Ltd: Technical publication - Digital Camera Fundamentals.

e2v Technologies Ltd: Low Light Technical Note 4 - Dark Signal and Clock-Induced Charge in L3Vision CCD Sensors, July 2003.

e2v Technologies Ltd: Low Light Technical Note 5 - An overview of the Ageing Characteristics of L3Vision Sensors, March 2006.

Mackay, C.D., Tubbs, R.N., Bell, R., Burt, D., Jerram, P., and Moody, I. Sub-electron Read Noise at MHz Pixel Rates SPIE 4306 Conference Proceedings, 289-298, January 2001.