Astronomical Imaging with L3CCDs: Detector Performance And High-speed Controller Design

Craig Mackay[*a], Alistair Basden[b], Mick Bridgeland[a]
[a]Institute of Astronomy and the [b]Cavendish Laboratory, University of Cambridge, Madingley Road, Cambridge UK, CB3 0HA

Abstract

L3CCDs represent a major step in CCD performance with great potential for astronomical applications because of their ability to work at very high pixel rates with negligible readout noise. This paper describes the results of tests on some of the L3CCDs now available and discusses how the operating conditions may be optimised for a variety of different applications. In particular, at high gain they can be used for photon counting work at photon rates well in excess of one photon per pixel per second. Readout rates which can be as high as 35MHz are entirely practical for a number of astronomical applications. This paper describes some of the compromises and trade-offs that have to be made in designing high-speed controllers to work effectively with these devices. The importance of integrating high-speed controllers for astronomy with significant amount of real-time processing power is also discussed.

Keywords: CCD controllers, signal processing, L3CCDs

1.Introduction

The first CCDs to be commercially available with an effective internal gain mechanism were produced by E2V Technologies Ltd (Chelmsford, England)[1,2]. By extending the output register and allowing it to be clocked with pulses in the 30-40 volt range rather than the more normal 10 volt range, electrons are accelerated on passing through each cell and have a small probability (up to 2%) of generating a second electron by avalanche multiplication. After passing through several hundred multiplication stages the overall gain may be several hundred or several thousand. At the end of this extended register the readout noise caused by the output amplifier is effectively divided by this overall gain giving net readout noises which are a tiny fraction of one electron. For practical purposes such a CCD can be run in a mode where it has essentially no readout noise even at high pixel rates. E2V Technologies have now produced a range of these L3CCDs (low-light-level CCDs) with sizes from 128 x 128 pixels up to 1024 x 1024 pixels of 16 microns square with other devices under development. They have available L3CCDs which are both front illuminated and back illuminated (thinned). A second manufacturer, Texas Instruments (Dallas, USA) has also produced CCDs with very similar electron multiplication technology under the trade name Impactron. They have a variety of devices available up to 1024 x 1024 pixels of 8 microns square. They also have other devices in development. At present it seems unlikely that there will be further manufacturers entering this area because of patent issues. For E2V and Texas Instruments, however there is no reason why any existing conventional full frame, frame transfer or interline transfer CCD should not have an electron multiplying output register added to the device if this level of performance was required for a particular application.

By eliminating readout noise the electron multiplying CCDs (EMCCDs) represent a major step in CCD performance and have considerable potential for certain astronomical applications. EMCCDs have all the advantages of conventional CCDs in terms of quantum efficiency, charge transfer efficiency, dark current, cosmetic quality etc. Because of that the technology effectively stands on the shoulders of all the developments already made with standard CCDs. The only significant disadvantage of EMCCDs is a consequence of the multiplication process. An individual electron passing into the multiplication register comes out with a much higher signal because of the gain process but there is a substantial

[*] e-mail:cdm@ast.cam.ac.uk; phone +44-(0)1223 337543; fax +44-(0)1223 330804

variance in this signal level because of the multiplication process. At the lower signal levels and working with high gain, individual photons may be recognised with good signal-to-noise. These allow the full quantum efficiency to be obtained in this photon counting regime. At higher signal levels the effect of the variance arising from the stochastic multiplication process is to increase the readout noise by a factor of up to $\sqrt{2}$. This is equivalent to halving the quantum efficiency of the detector in terms of the effective output signal-to-noise that is generated. Work in Cambridge[3] has shown that the transition between the photon counting regime and the more conventional integration regime where the multiplication has degraded the signal-to-noise occurs gradually with an input signal of between one and three electrons per pixel per frame . It could be argued that when one has higher signal levels, the requirement for the highest quantum efficiency is less severe. However it is always the case that astronomers are extremely reluctant to accept any diminution in quantum efficiency unless it is completely unavoidable. The application areas that are likely to be of particular importance to astronomers will include high time resolution photometry, very low light level spectroscopy (particularly of diffuse objects where the integrated photon rate per pixel is extremely low and normally not accessible because of the read noise from conventional CCDs), wavefront sensing (where high time resolution is critical if an adaptive optic system is to be able to determine wavefront errors quickly enough so that they are still relevant by the time the compensation mirror is corrected) as well as other applications such as Lucky Imaging. The ability to readout CCDs very rapidly without any noise penalty also allows a different approach to data taking. It is always possible to run at a faster frame rates and/or with a higher spatial resolution and then to average or bin temporally or spatially after the event to give whatever actual resolution is required. There is no penalty, therefore, in running fast provided there is a mechanism for coping with the data thus generated.

We have used L3CCDs from E2V Technologies extensively on the Nordic Optical Telescope on la Palma for a programme of Lucky Imaging[2,4,5,6]. This is a new method of minimising the loss in resolution usually caused by atmospheric turbulence when working with ground-based telescopes. Images are taken very rapidly (10-100 frames per second). A reference star is checked in each frame for image quality. Only the best images are kept. These are then shifted and added to create the summed output image. The overall resolution of the images obtained in this way is substantially better than is normally achievable with a ground based telescope. Resolutions of 0.1-0.15 arc seconds are routinely possible under conditions of relatively good seeing. Even under conditions of rather poor seeing an improvement of resolution of a factor between two and three is achieved routinely. The Lucky Imaging technique is described in more detail in another paper[7]. The performance of L3CCDs has been described in considerable detail by Simon Tulloch in another paper in this conference[8] and will not be repeated here. Some of his results are based on data taken with our camera in Cambridge.

2. L3CCD Operational Issues

The characteristics of L3CCDs have been described elsewhere[8,9]. The main operational issues that are important to bear in mind when building systems based on this technology include dark current, amplifier glow and clock induced charge. These devices also require a completely different approach when it comes to designing and building controllers that are going to make the best use of the capabilities of EMCCDs in general. This will be discussed in the following section.

Because it is possible to work at extremely low signal levels with an L3CCD then it is important to make sure that all sources of signal which are not wanted are dealt with rather thoroughly. Dark current can be a problem at a level that is not normally encountered with normal CCD operation. The CCDs from E2V Technologies can be cooled very considerably, and we normally operate our detectors at temperatures of approximately 130 Kelvin. If the lowest light levels are to be reached then the dark current must be well below the minimum signal level to be detected over the full exposure, which may arise from many thousands of individual frames. The dark current must be extremely low to achieve this in practice.

The presence of the output amplifier on the CCD means that very low light levels may be generated by the amplifier itself. With the CCD87 and CCD97 from E2V Technologies we see this as slightly increased dark current in the corner of the image nearest to the output amplifier (and because the extended register is folded back on itself this corner is next to the last pixel out of the first row[8]).The glow from the output amplifier can be minimised by making sure that it is turned off during integration. However with fast frame readout it is probably going to be operating essentially all the time. In this case it is possible to minimise the glow by reducing the output drain voltage (and therefore the reset drain

voltage as well). This may mean that the output amplifier is not working at its ideal point but this can be compensated for somewhat by simply increasing the overall multiplication register gain.

Clock induced charge is also important. Clock induced charge is seen as an increased dark level and arises because electrons are excited by the high electric fields that occur during clocking and in particular as a result of the clock transitions. The clocking strategy therefore is to keep the parallel transfer speed as fast as practical while at the same time making sure that the clock edge speeds are kept as slow as possible. E2V Technologies recommend (Low-light Technical Note 4) that edge speeds of 200-300 ns are ideal for parallel transfers. We find, however, that it is serial clock induced charge that is a bigger problem which again can be minimised by careful attention to the clock waveforms. Keeping transitions as slow as possible consistent with good serial transfer is critical. It is easy to produce very fast clock edges with modern drivers that produce astonishingly fast rise and fall times but it is essential to keep these under control if clock induced charge is to be managed properly. It is also important to remember that by working normally with very high gain then the maximum signal that can be tolerated in the image and store regions of the CCD will be very small indeed. This allows a substantial reduction in the amplitude of the clock swings to a level just good enough to give good charge transfer efficiency. The output registers of these devices generally have a much higher charge capacity than normal CCDs (to allow the use of some of this gain without affecting the dynamic range excessively).

3. EMCCD Controller Design

Most designers of CCD controllers for astronomical application will be familiar with approaches to relatively slow scan high precision system design. Although such experience is valuable it is very important to appreciate that it is not practical to take slow scan technologies and design strategies and then simply speed them up. You cannot expect to get the same performance at 35MHz pixel rate (the pixel rate at which the Texas instruments TC285 1024 x 1024 EMCCD is specified) that you presently get with a 100 KHz 16-bit slow scan camera system. Compromises have to be made. New design technologies need to be learned including high-speed analog electronic design, the use of high-speed simulation software packages on those designs, new approaches to high frequency printed circuit board layout strategies, the use of the latest field programmable gate arrays as well as techniques of management and suppression of high frequency noise.

However help is at hand because there is a considerable commercial interest in running CCDs at high pixel rates and achieving excellent performance. Modern handheld digital cameras use CCD and CMOS detectors and have to perform to an extremely high standard. The detectors they use can have up to 10 million pixels that have to be read out at a rate of several frames per second if they are going to be commercially acceptable. The CCDs in these cameras have pixels that are extremely small and therefore the signal levels they work with are also rather low. This means that the performance of the electronics associated with such cameras has to be exceptionally good. They have to work over a wide range of ambient temperatures and under a wide range of different operating conditions. As result there is a good selection of extremely sophisticated integrated circuits and design tools which make excellent performance much easier to achieve.

It is important therefore to be aware of the components that are available today and to see how the specification of the controller which is being designed can be modified to make the best use of what is available. Because the management of noise and transition times is so critical to the success of a design working at high pixel rates it is extremely difficult to create designs that work over a significant range of pixel rate. The fast clocks (serial clocks, high-voltage multiplication clock and reset clock) all need to be carefully filtered to minimise transients. The same is true for the analog signal coming out of the CCD. It needs to have very careful bandpass limiting filters incorporated to minimise the readout noise. The filtering has to be designed for a particular operational frequency and therefore changing frequency becomes much harder. With EMCCDs, however, there is no particular need to change the readout rate. The system can be designed to give extremely good readout noise at a specific pixel rate and left to run at that, even if the exposure times used are very much greater than the readout time. Choosing a single pixel rate also makes the selection of other signal processing components, for example, much easier since they are also usually designed also for a single processing rate.

4. EMCCD Controller Components

There are a number of areas where the approach to designing a high-speed controller is really very different from that appropriate for slow scan systems. The pixel rates at which the performance of the CCDs are specified vary from 11MHz (for the CCD 87/97 from E2V Technologies) up to 35MHz (for the TC 285 from Texas Instruments). However it is likely that the CCD 87/97 can be operated up to the 15MHz pixel rate given as the maximum rate in the datasheet and possibly as high as 20MHz. Similarly the TC 285 can probably be operated in excess of 50MHz and possibly as high as 65MHz. Please note that the author has no evidence that these pixel rates can in fact be achieved! What this means, however, is that it is important for the controller designer to bear in mind that there are likely to be pressures to push the pixel read rate up as high as possible since it is for high-speed applications that these devices are going to be particularly attractive.

The basic structure of an EMCCD controller will be very similar in outline to that of a slow scan system. A timing core generates clock waveforms and controls a double correlated sampling signal processing chain. An analog to digital converter interfaces to an output system that transfers the data back to the host computer. The timing of the various operations within each pixel has to be carefully controlled and adjusted to give the best noise performance and the clock waveforms have to be carefully tailored to give the optimum imaging performance. However operating at 35MHz gives a 30 ns pixel time and operating at 65MHz gives a 17 ns pixel time. Everything has to be done within this pixel time. Immediately we find that there is really no such thing as a digital circuit. All these circuits are analog, some of which are able to switch between two states quite quickly. What you see on your oscilloscope may or may not be what is actually happening to the CCD. Transmission line effects and reflections together with inadequate bandwidth of test equipment can mask real effects as well as appear to create problems which do not really exist.

The generation of the basic clock patterns is something that can be done in principle with field programmable gate arrays or other programmable sequencer. The big problem here is to develop the ability to adjust timings to considerably less than 1 ns in some kind of repeatable way. Digital delay lines are possible and work reasonably well although they are rather expensive (Dallas Semiconductor make a variety of programmable 8-bit delay lines. They have models with steps which range from 0.15 - 2 nanoseconds with a total of 256 steps that allow timing delays to be adjusted simply by selecting a serial or parallel bit pattern). The latest FPGAs such as the Virtex-II family from Xilinx Inc., include programmable clock generators that allow very small levels of skew to be introduced into output signals. There are also components which are integrated with CCD signal processors that include relatively complicated precision timing generators. Analog Devices (Norwood, USA) make the AD9895 integrated CCD signal processor which includes a precision timing core with 700 psec resolution at 30MHz. Even if if the performance of the the 12 bit 30MHz double correlated sampling front-end that comes with that device is not what is wanted, it may be that the timing generator part of it is worth it on its own.

The approach of slow scan system designers to clock drivers is often simply to use operational amplifiers driven by FET switches between two levels. This will simply be far too slow for this kind of work. However there are clock drivers circuits available from companies which manufacture CCDs for the high-speed commercial market (though you should certainly check availability before designing in a component: some of them are only available for minimum order quantities of several thousands). There are also high-speed drivers for automated test equipment known as Pin Drivers which can be used for CCD clocks. Elantec (Milpitas, California) make a range of these devices of which the EL7156C is probably one of the best. It is designed to provide clocking speeds up to 40MHz driving a 2000 pF load. When driving the much smaller load that we typically have with the serial clock electrode of perhaps 150 pf then the rise time can be a short as 2 ns. Alternative designs of clock drivers are provided sometimes by the CCD manufacturers. The datasheet for the TC 285 from Texas Instruments shows how a high-speed clock driver can be built using two enhancement mode FETs, although they also state that the EL7156C is a satisfactory alternative.

The driver for the high-voltage electrode is significantly more demanding. The high-voltage clock electrode of the Texas Instruments TC 285 only requires a voltage swing of about 25 volts to give a gain of a few thousand when cold whereas the E2V Technologies CCD 97 requires a voltage swing of nearly 45 volts to achieve the same gain, also when cold. E2V Technologies suggest that the high-voltage clock can be driven with a sine wave rather than a square wave that is used for the other clocks. This can lead to a much lower driver power by using a controlled oscillator

arrangement but it does require that the high-voltage clock runs continually because the oscillator will not start running instantly. The schematic given in the TC 285 datasheet mentioned above will also drive the E2V device and will give a 40 volt clock at 20MHz. Note, however, that one of the transistors specified only has a drain-to-source breakdown voltage of only 40 volts. Other enhancement mode FETs with higher breakdown voltage are available from Supertex (Sunnyvale, California) as well as from Vishay-Siliconix.

The design of clock circuits running at these speeds is distinctly nontrivial. The basic clock driver has to be able to produce a fast enough edge but it then has to be filtered to stop ringing compromising charge transfer efficiency and/or generating clock induced charge. The further complication is that running these little tiny surface mount devices at high-speed causes an surprising amount of power dissipation. The datasheet for the EL7156C discusses this problem. Careful thermal design and protection is necessary.

The signal processing chain is another area where things are quite demanding. The first thing to note is that the two varieties of CCDs are very different in terms of their output gain. The CCD 87/97 family has a charge conversion gain of approximately 1.1 microvolts per electron. By contrast the TC 285 has a charge conversion gain of 14 microvolts per electron. This makes the Texas device much easier to drive and work with because of this high sensitivity. Although the TC 285 device has a significantly smaller maximum signal capacity in the multiplication register (42,000 electrons with no gain compared with the 700,000 electron capacity of the CCD 87/97), it can be operated with much lower gain because it has intrinsically a much lower readout noise at the maximum operating frequency. The higher charge conversion gain will also make the design of the first stage of pre-amplification much easier.

Double correlated sampling is something which works just as well at these frequencies as it does at low frequencies. It is often been claimed that it is not possible to get double correlated sampling working at high speeds even though commercial controllers such as the AstroCam Capella 4100 have been on the market for 10 years with double correlated sampling working at 5MHz. These days it is possible to buy double correlated sampling components such as the AD9823 from Analog Devices which allows double correlated sampling operation at 40MHz. There is also a wide range of analogue to digital converters available that provide very high standards of operation. Examples include the AD5500 from Texas Instruments (14 bits at 125MHz) and the AD10678 from Analog Devices (16 bits at 80MHz). Before using such components one has to look very carefully at the datasheet to make sure that the guaranteed performance is adequate. Remember that if you have gain with your CCD then you do not actually need 16-bit performance. If the peak signal level is only a few hundred or a few thousand electrons then only 10 or 12 bit performance is needed although a higher resolution unit may be needed in order to give the required performance. What is difficult to divine from these data sheets is the overall noise level in data numbers produced by the integrated double correlated sampling processor. In order to get the best performance it may be necessary to set the internal gain to a rather low level and then provide much higher quality pre-amplification external to the converter.

The same considerations apply if one uses a complete analogue signal processor containing double correlated sampling components, programmable gain amplifiers and analogue to digital converters all in one package. Examples include the AD9824 from Analog Devices (a complete 14 bit 30MHz CCD signal processor) and the WM8214 from Wolfson Microelectronics (Edinburgh, Scotland) which is a complete 16-bit 40MHz CCD signal processor. For designers who have sweated for years to produce a much slower double correlated sampling CCD front-end it is mildly depressing to be able to buy one of these chips that does everything the slower system did and a hundred times the speed for just a few tens of dollars . These devices are generally well worth looking at because they provide so much functionality in a very compact space. This compactness means that many of the problems one encounters when trying to use discrete components inevitably extending over many square centimetres of board are greatly reduced. The designers of these components have made sure that they are complete and self-contained and self consistent in a way that is almost impossible for a single design engineer to achieve with the relatively modest resources usually available even in quite large astronomical organisations.

One final area that must be considered as an integral part of the design of any CCD controller operating at these speeds is the management of the data they produce. A CCD operating at 35MHz and producing 14 bit data will have the capacity to generate 70MB of data per second indefinitely. This data volume is quite close to the data input bandwidth limit of modern computers just as a simple data throughput operation unless one spends a great deal of money on more sophisticated systems. In practice it is essential to identify why such a high data volume is being used and then making

sure that there is the right kind of data management and processing hardware and software available between the output of the controller and the input to the computer. Only by dramatically reducing the data volume that needs to be stored will there be the slightest prospect of being able to use the system at a telescope.

It is worth saying in concluding this section that these modern components are extremely well documented and are very easy-to-use considering their complexity and performance. Provided that the engineer invests enough time in learning about the new technologies that have to be absorbed to make these systems work the results can be very satisfactory.

5. The Cambridge High-speed EMCCD Camera System

Many of the comments made in the previous sections arise from our experience of developing a high-speed camera system dedicated to the operation of EMCCDs at pixel rates up to 35MHz with 14 bit conversion. The camera is intended to operate L3CCDs from E2V Technologies initially but with a potential to work with the Texas Instruments TC 285 in the future.

We chose to construct a system that used an Analog Devices TigerSHARC DSP development board to generate the high-speed clock waveforms, using the 83MHz 32-bit data bus to generate a 16-bit control bus operating at 166MHz. The control waveforms are transferred to the two camera boards with LVDS high-speed drivers and receivers. The standard clock waveforms are generated with Elantec EL7156C pin drivers and are capable of working up to 35MHz although our initial system is intended only to work at up to 15MHz. The high-voltage clock driver uses a circuit that is similar to that described in the TC 285 datasheet and which has been used successfully in the camera used for our Lucky Imaging experiments over the last three years. The timing is controlled and adjusted to within 0.25 ns steps using programmable digital delay lines. The analogue signal processing chain consists principally of the Analog Devices AD9824 complete 14 bit CCD signal processor. The output data are transmitted back to the host computer again with high-speed LVDS drivers and receivers. The data will also be sent to a Bittware DSP array processor card for real-time Lucky imaging.

6. Conclusion

Electron multiplying CCDs have the potential to open up entirely new ways of operating detectors at telescopes. Although there are significant issues that need to be dealt with in using these devices and making them work there is a great deal of information and a great variety of electronic components that are now available to make this work much more tractable and much more enjoyable.

7. Acknowledgements

We should like to acknowledge useful discussions with Simon Tulloch, ING, La Palma. We are also grateful for the help and support of the staff and director of the Nordic Optical Telescope on La Palma.

8. References

1. P. Jerram, P. Pool, R. Bell, D. Burt, S. Bowring, S. Spencer, M. Hazlewood, I. Moody, N. Catlett, P. Heyes, "the LLLCCD: Low Light Imaging without the need for an intensifier, SPIE vol. 4306, 2001.
2. Lucky Imaging web site: http://www.ast.cam.ac.uk/~optics/Lucky_Web_Site/index.htm
3. A.G. Basden, C. A. Haniff and C. D. Mackay (2003) “Photon Counting Strategies with low light level CCDs”, Mon. Not R. astro. Soc, vol. 345, 985-991.
4. Baldwin, J.E, Tubbs, R.N., Mackay, C.D., et al, (2001), Diffraction-limited 800nm imaging with the 2.56m Nordic Optical Telescope, Ast & Astrophys, vol. 368, L1-4.
5. Tubbs, R. N, Mackay, C. D., et al., (2002), "Diffraction Limited CCD Imaging with faint reference stars", Ast & Astrophys, vol. 387, L21-24.
6. Tubbs. R.N., Mackay. C.D., Baldwin, J. E., (2002), “Diffraction-limited I band imaging with faint reference stars”, SPIE vol. 4839, page 1093. 22-26 August 2002, Kona, Hawaii.

7. C.D. Mackay, J. E. Baldwin, N.M. Law, P.J.Warner, (2004),High resolution imaging in the visible from the ground without adaptive optics: new techniques and results, 2004, SPIE 5492-13
8. S Tulloch, Photon Counting and Fast Photometry with L3CCDs
9. Mackay, C D, et al., (2001), "Sub-Electron Read Noise at MHz Pixel Rates", SPIE vol. 4306A, San Jose, January 2001,p289-298.