Investigation into L3 CCD Performance

University of Durham Sep 2002-June 2003

Simon Tulloch

Introduction to L3 technology

An L3 chip, the CCD60, is shown schematically below. In a conventional CCD the serial register contains approximately the same number of pixels as there are columns in the image. L3 chips have an unusually serial register design. This register is extended by 520 pixels into what is referred to as the Multiplication Register. Photo-electrons generated in the imaging part of the chip are amplified as they pass along this register, by a process of avalanche multiplication. The degree of amplification can be adjusted externally from unity up to several tens of thousand. A single photoelectron entering the multiplication register exits as a substantial packet of charge that produces a signal in the output amplifier far in excess of its read noise. Single photon detection is therefore possible.

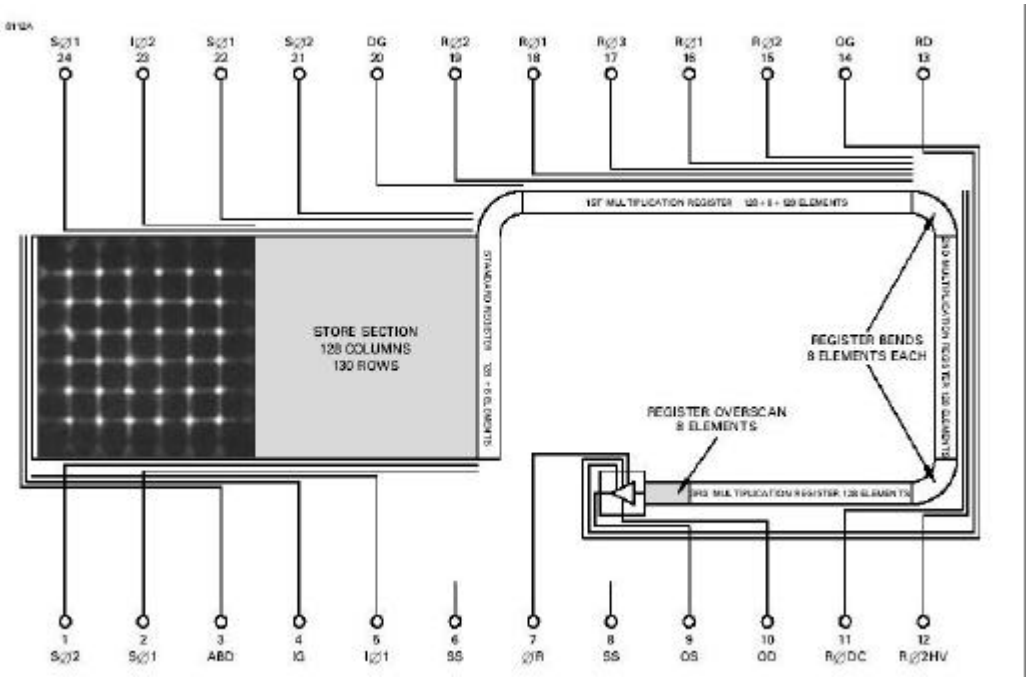

Two scientific grade (i.e. thinned, backside illuminated) L3 chips are currently available from the manufacturers, E2V technologies. The smallest, shown above, is the CCD60. The second is the CCD87, also a frame transfer device with a 512 x 512 pixel imaging area. The QE of these devices matches that of conventional E2V thinned CCDs. The L3 technology could be applied to any desired image format for those willing to pay for a custom manufacturing run.

The multiplication process offers a large potential gain in SNR at low signal levels. Unfortunately the process introduces an unavoidable extra noise source, known as Multiplication Noise that actually degrades the SNR at high signal levels. There is

thus a crossover point , only below which L3 chips offer a real advantage. The graph below shows this. The crossover point comes at the 20-30e signal level for normal slow scan levels of noise (3 to 5e)

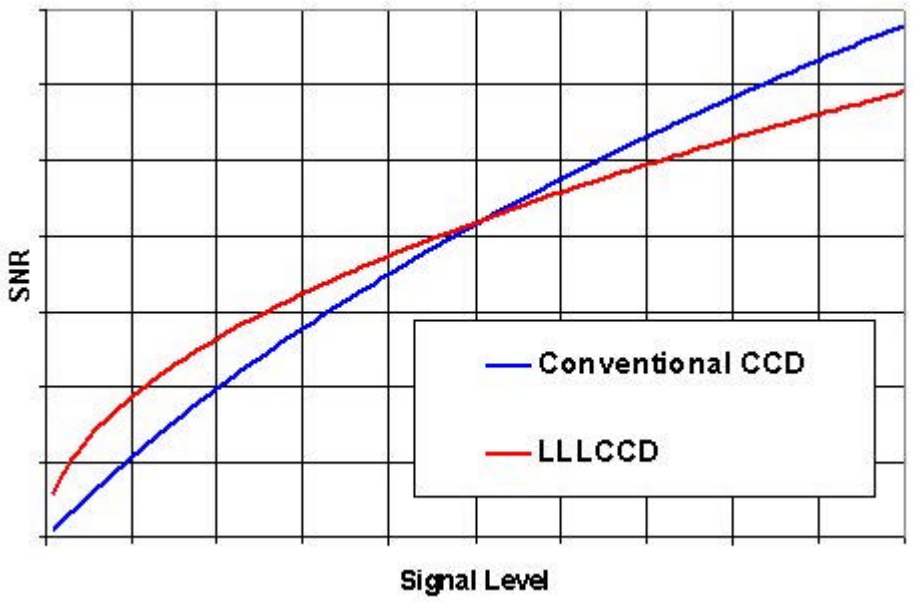

The precise effect of the multiplication noise is to degrade the SNR in the high signal regime by a factor of $2^{1/2}$.This is tantamount to a halving of quantum efficiency. An L3 chip can therefore be modelled very closely (for the purposes of limiting magnitude calculations for example) as a zero-noise CCD with a QE one half of a conventional scientific device.

Investigation program at Durham.

A CCD60 scientific grade CCD was purchased in 2002 for the study. The principal aim was to assess the chip's likely performance in NAOMI, but the photon counting performance was also investigated. The study aims were:

1) Investigate maximum frame rate possible with an SDSU controller
2) Measure SNR improvements at low signal levels
3) Measure the gain stability of the CCD60
4) Confirm the presence of multiplication noise
5) Investigate other noise sources in the chip.

Test Camera Design.

The CCD60 was incorporated into a standard ING 2.5l Oxford Instruments dewar. The equilibrium temperature once filled with LN2 was –100°C and the hold time >20hrs. The camera was connected to an SDSU II controller which supplied all the clock and bias signals to the CCD with the exception of H2ϕHV : the hi voltage clock that induces the avalanche gain mechanism. This clock, although timed from the SDSU clock board, was actually generated by an extra card bought in from E2V. In their literature this is described as an L3MHV module. Another connection was made between one of the DAC channels on the SDSU video board and the gain control pin of the E2V module so that the multiplication gain could be adjusted in software.
The E2V module was not specifically designed for the CCD60. Although it worked well, it had very unpleasant start up transients that had to be limited through the addition of a zener clamp on the output. It was physically small enough to mount on a Eurocard that was inserted into a spare slot inside the SDSU. The backplane connector of this card was used to supply power, all signal connections were made from a separate connector at the front.

Test Camera part assembled.

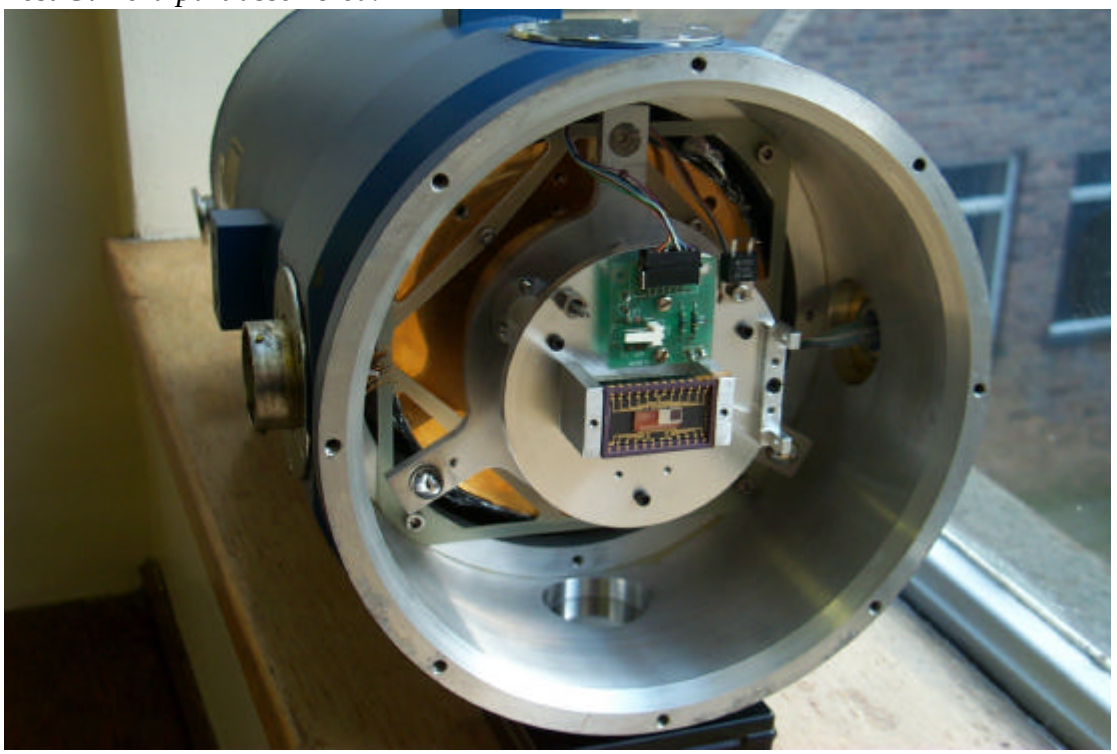

Test Camera after wiring

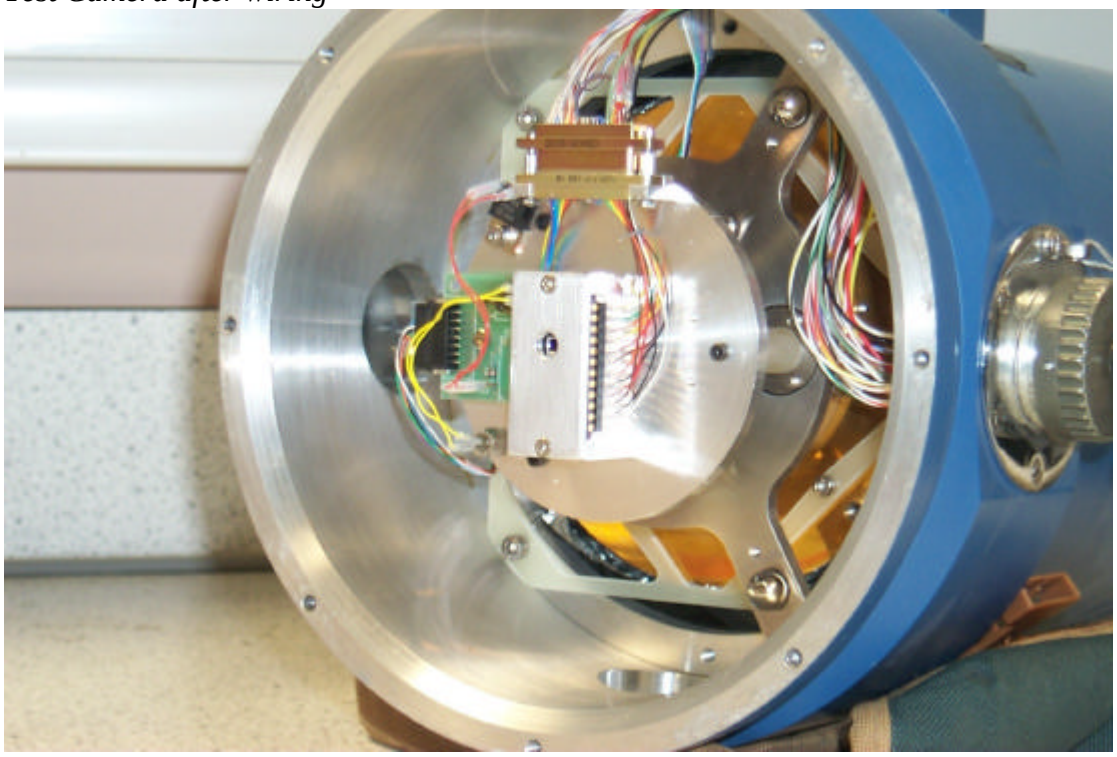

Demonstration of L3 Gain and Sub-Electron Noise

The first measurement to be done was of the multiplication effect itself. A stable light source was observed for short exposures with a variety of gain settings and at different temperatures.

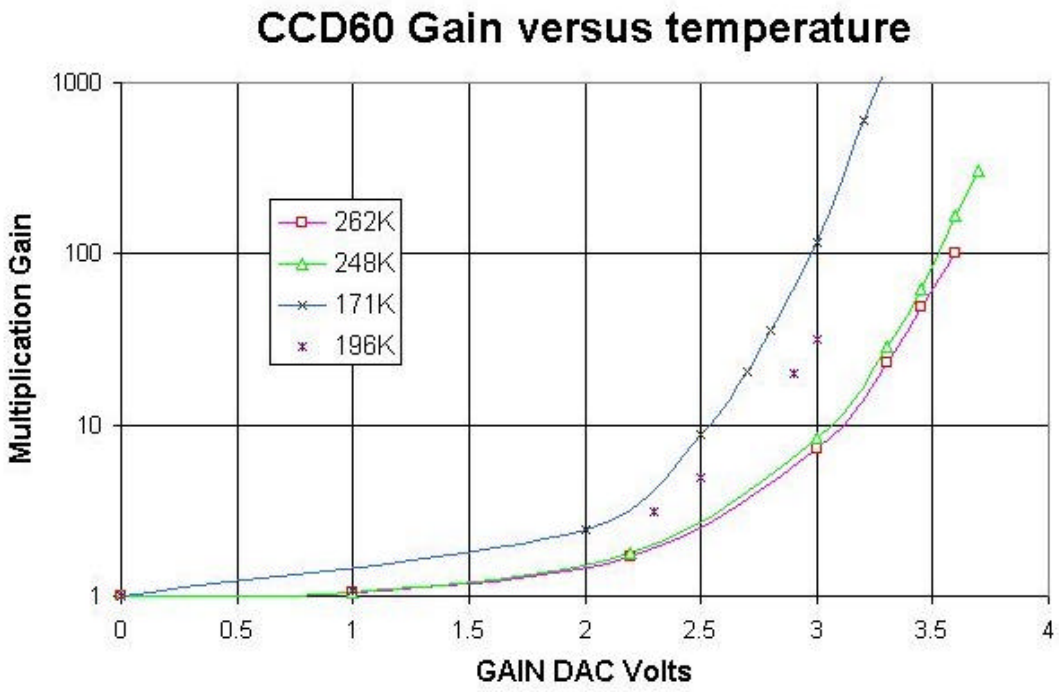

The illumination of the chip was varied by making changes to the exposure time. The CCD was of frame transfer architecture (i.e. it has integral electronic shuttering) so the exposure time could be accurately varied down to a few ms duration. It was therefore possible to calculate the gain by measuring the change in exposure time required to maintain the signal level constant. The gain DAC of the SDSU video board had a range of 0 to +5V. When set to 0V, the multiplication gain of the CCD60 was 1, DAC voltages above 4V were not used as the chip maximums were approached and in any case a protection zener on the output would start to clamp.

As can be seen from the graph the multiplication gain was strongly temperature dependant, increasing at lower temperatures.

The intrinsic amplifier noise was 50e (a consequence of very rapid video waveform sampling and the use of unshielded cables in the camera design). At multiplication gains above 50 then, the *effective* readnoise of the amplifier was sub-electron. With higher gains, and zero illumination, the noise in the bias frames (as measured in ADU) eventually started to increase and the effective readnoise bottomed out at about 0.1e RMS. This noise had a strange speckled character. It was later realised that the speckles corresponded to individual electron events. Unfortunately these were not photo-electrons but were instead produced inside the CCD itself through a process known as Spurious Charge Generation. This was explored in some depth later in the study (see the section on Photon Counting).

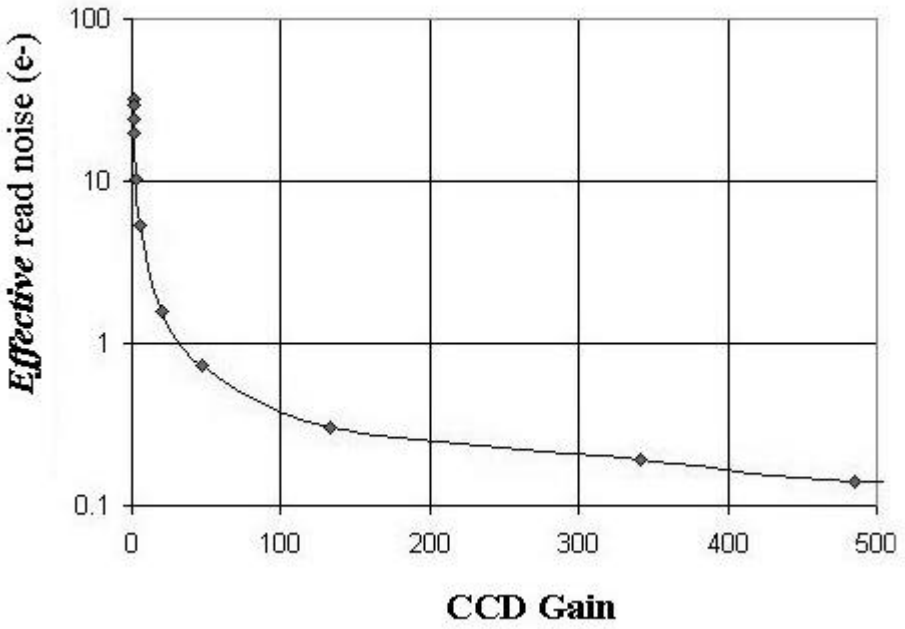

Measurements of the Multiplication Noise

Monte Carlo simulations of the multiplication process were done using GnuC and the Gnu Scientific Library 9see Appendix) that confirmed what the manufacturers earlier stated about the $2^{1/2}$ degradation in SNR at high signal levels. This was later backed up by actual measurements.

These measurements were done by uniformly illuminating the CCD with a stable light source whose intensity in photoelectrons per pixel per second was known. The exposure time was gradually increased and the relationship between mean signal level and the RMS variation in the signal level was plotted. In the case of a noise free system this would result in a straight line graph (indicated below by the bold black line) with a gradient of one i.e. Poissonian statistics. The measurements were repeated at 3 different gain settings between 11.8 and 283 to see if there was any dependence.

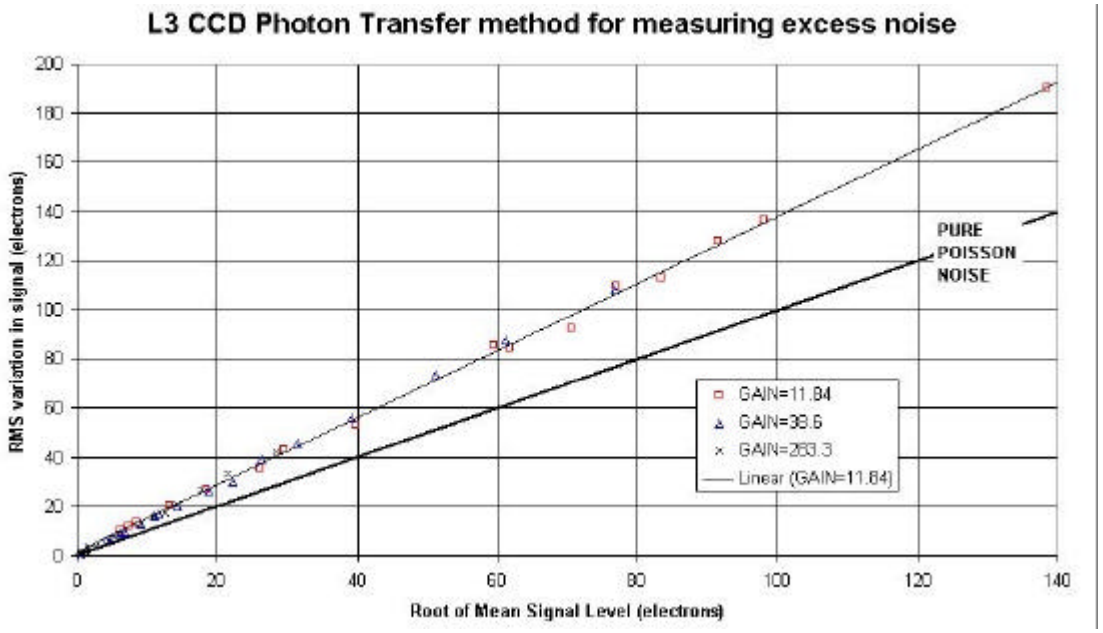

As can be seen , the data at all gain settings fits very well to a line of gradient $2^{1/2}$, the excess noise factor.

Monte Carlo Simulations of Multiplication Noise

The simulation model assumed that the gain of each multiplication stage was equal, and that the probability of an electron being multiplied was independent of the number of other electrons in the pixel. The model was very valuable as it showed up some non-intuitive results and showed that multiplication noise is an unavoidable feature of L3 CCDs. In the Monte Carlo model, a single electron was simulated entering a 520 stage L3 gain register with a total gain of 1060. The number of output electrons was recorded and the simulation run again, a total of 2 million times. A histogram of output values is shown below:

Monte Carlo Simulation. L3 Gain=1060, Input=1e, Binwidth=10e

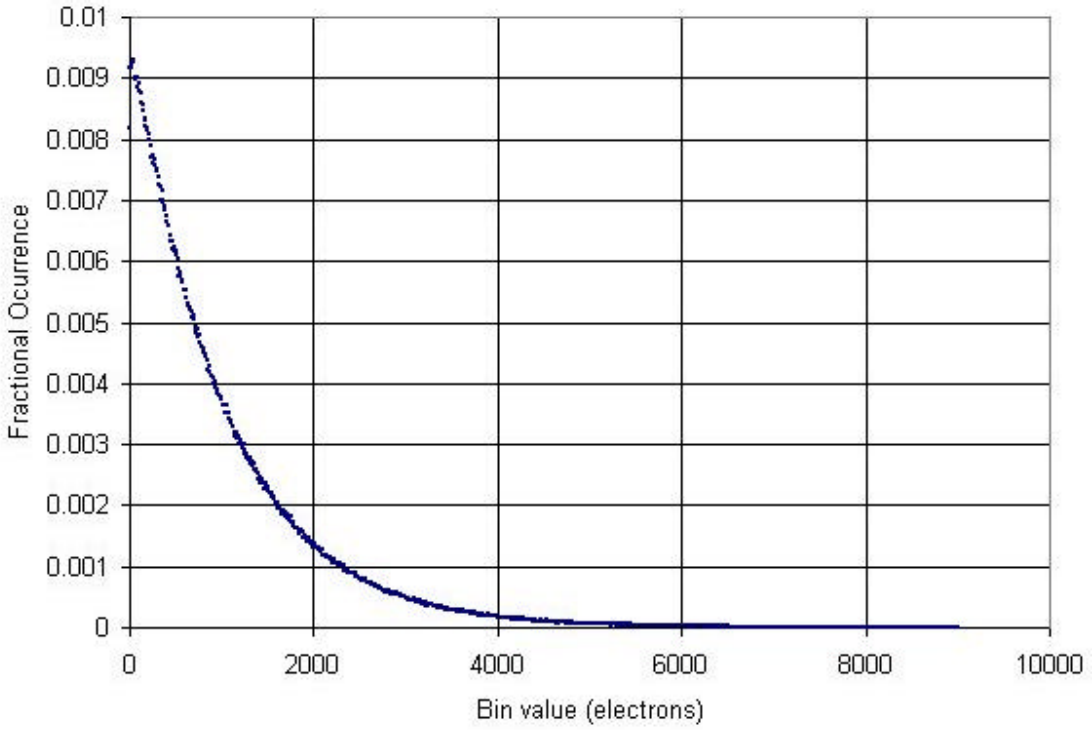

Rather than showing a peak at 1000e, the curve shows a peak in the first histogram bin. So a single electron entering the gain register can give rise to a wide range of output signals.

As the number of electrons in a pixel increases the output of the model starts to look more intuitive. The Sigma of the distribution is 1.4 times higher, with the multiplication noise included in the model, as expected.

Monte Carlo Model. 10e mean input, L3 gain 100

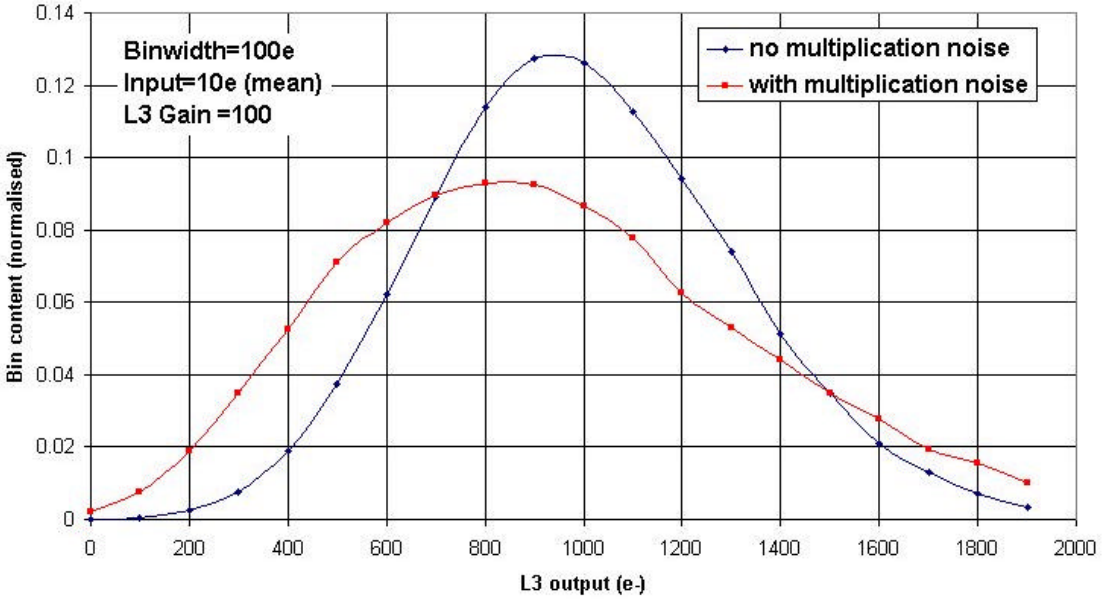

Monte Carlo Model. 30e mean input, L3 gain 100

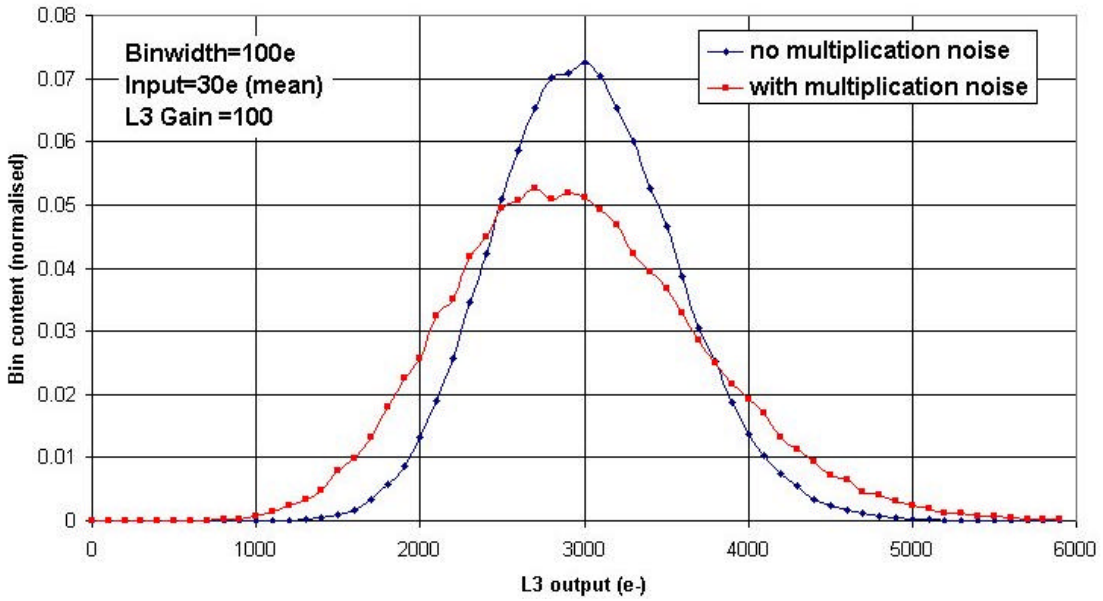

Application to Wavefront Sensing

Some preliminary simulations done by Richard Wilson indicated that a zero noise detector with multiplication noise could offer quite an advantage to the Naomi SH WFS. Richard's result is shown below.

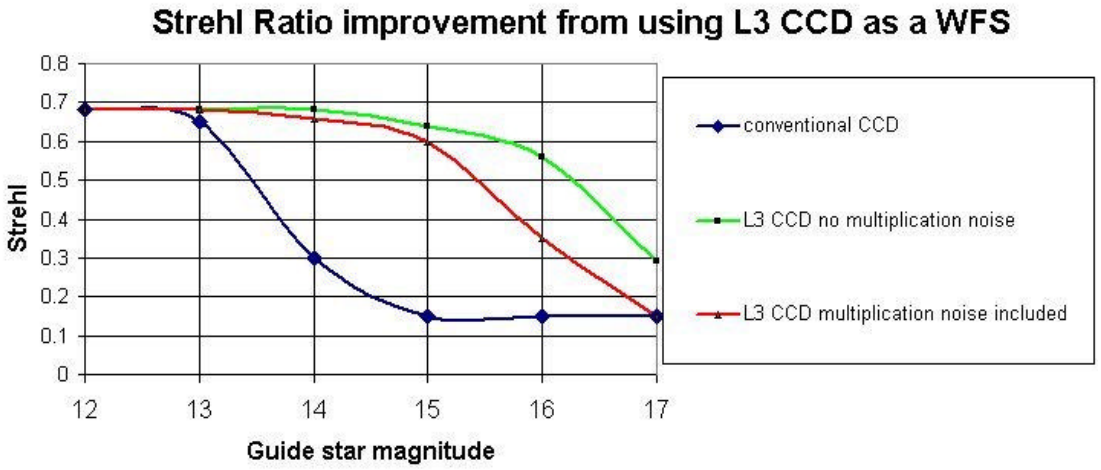

Another simulation is being done for the end of August to include more accurate predictions of the CCD60s electro-optical performance. Its suitability as a NGS tip tilt sensor will also be simulated.

Tests were done to verify that the required frame rates could be achieved. These are described in the following section.

Readout Speed through an SDSU II controller

The maximum read out speed of the CCD is quoted as 16Mpix per second i.e. 1000 frames per second when using an E2V controller. This was not available for the study so all the measurements presented here were obtained with an SDSU II controller.

If applied to the NAOMI, the CCD will need to be windowed in quite a complex fashion. The NAOMI Shack Hartman (SH) spot pattern consists of an 8 x 8 matrix of spots with 8 pixels separating the spots in both the vertical and horizontal directions. The centroiding is calculated from a 4 x 4 pixel window around the nominal position of each spot . There are thus 4 dead (or 'guard') pixels between each spot that must be skipped over by the windowing algorithm. The total number of pixels that need to be digitised in a NAOMI WFS frame is therefore 32 x 32 = 1024. In order to check that the windowing algorithm was indeed functioning correctly, a SH test target was made using pinholes in an aluminium foil target illuminated from behind by an extended diffuse source.

Various windows were set up and their timings are shown below:

A) Full CCD60 128 x 128 pixel frame

Frame rate **28Hz**

B) Windowed down to 100 sub-windows of 4 x 4 pixels

The dead space between windows was skipped over.

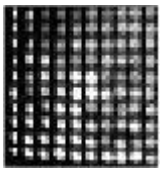

Frame rate **211Hz**

C) Windowed down to 64 sub-windows of 4 x 4 pixels

This is the format required by NAOMI.

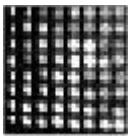

Frame Rate **320Hz**

The peak pixel rate was about 400Kpix/second, well below the 2.5Mpix/s controller limit. The bottleneck comes in the speed of operation of the SDSU clock card which allows a minimum of 40ns between consecutive clock transitions. Nevertheless 320Hz is quite acceptable for almost all NAOMI regimes.

More recently, the CCD60 has been proposed as an NGS tip tilt sensor for the Rayleigh laser system. This would require imaging of a single small window only. Here, frame rates of >500Hz should be achievable, although no specific tests have yet been done.

Positioning of the image on the CCD is critical to the readout speed. The image should be as close as possible to the side (left or right) of the chip containing the readout amplifier. The vertical position of the image is not important.

Image Quality

The SH spots all showed horizontal streaking, indicative of poor Horizontal CTE. Various changes to the timing and levels of the clocks were made but it was not possible to improve the CTE, which remained at 0.9998. The streaking was equally bad at the left and right hand sides of the image which indicates the poor CTE originates in the multiplication register rather than the serial register itself. This level of CTE corresponds to a deferred charge of about 8% (compared to less than 0.5% for a conventional scientific device). The manufacturers have confirmed that this is a known problem. A CCD60 image below shows the soft right hand edge of a flat field where the image area charge is trailed into the bias area. The exposure level was a few thousand electrons per pixel.

CCD60 Flat Field

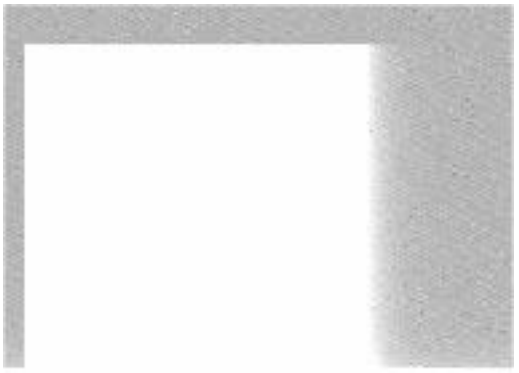

Using L3 CCDs for Photon Counting

If the multiplication gain is increased above several hundred, and the illumination kept sufficiently low, the L3 images take on an unusual appearance. A cut through such an image is shown below:

Cross section of L3 image taken with high gain

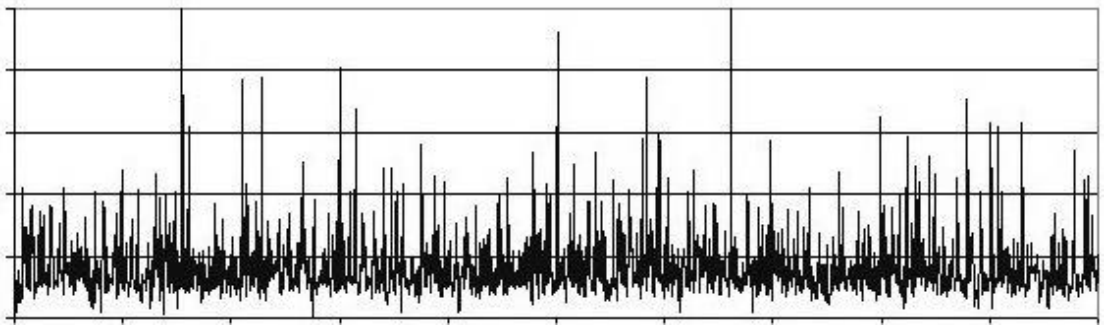

Overlying the Gaussian amplifier noise is a forest of 'spikes', which are all positive going. These spikes are in fact individual electrons entering the multiplication register. These spikes stand out clearly from the noise and could in principle be thresholded , either using a hardware comparator or by post processing of saved images, to give photon counting performance. Using photon counting mode removes the effects of SNR degradation due to multiplication noise and gets us one step closer to the ideal detector. In practice, there are a number of problems that arise. The most serious of these is 'Spurious Charge'.

Spurious charge in L3 CCDs

Even when the illumination was reduced to zero , the electron 'spikes' did not entirely disappear from the image, although a considerable reduction could be produced simply by letting the CCD settle for about 10 minutes after it was first powered up. These spurious electrons are created within the CCD by the high electric fields in the region of the transfer clock electrodes. Specifically, they are created when the CCD goes from the inverted to non-inverted states (as the clocks pass above approximately the substrate voltage). This effect is present in normal CCDs but is generally masked by the read noise except when high binning factors are used. In photon counting applications its effect is more serious.

A reduction in spurious charge can be achieved by reducing the clock swings or by using 'tri-level clocking'. In this mode, the clocks pause in their positive going transitions at an intermediate level close to the substrate voltage. The manufacturers also suggested running the CCD (an AIMO or 'Advanced Inverted Mode Operation' device) in its non-inverted state, where the default state of the vertical clocks was high rather than low. This would produce at least a 2 order magnitude increase in dark

current , but at cryogenic temperatures and short exposure times , this would not be significant.

Tri-level clocking was not tried; with the SDSU controller there would be a large speed penalty. Reducing the clock swings did not help as the charge transfer efficiency suffered very badly with only a small reduction in swing. Running non-inverted, however, gave a considerable improvement.

In order to investigate the spurious charge further, the CCD was read out with large overscan areas. It would then be easier to work out where in the chip the noise was actually being generated. The image format is shown below. The actual light sensitive 128 x 128 pixel region of the chip is delineated by the large red square. It is being weakly illuminated in this image and a large number of photon events are visible. The other two delineated regions occupy the serial overscan (to the right) and the parallel overscan (at the top). These regions are not light sensitive and the events that they show are spurious charge only.

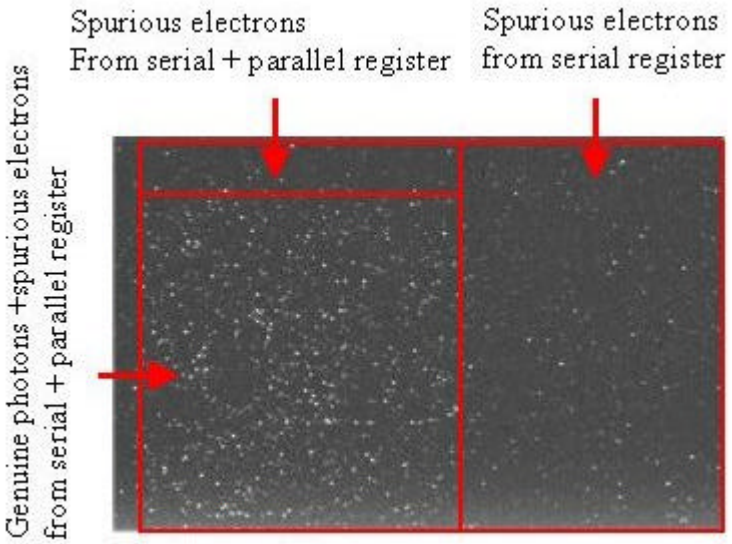

The pixels in the serial overscan have only ever transferred through the serial register so any spurious charge they contain must have originated there. On the other hand, pixels in the image area and the parallel overscan have been transferred through both parallel and serial registers and contain spurious charge generated in both places.

The density of events in the serial overscan and parallel overscan regions is identical indicating that there is now no spurious charge generation from the vertical clocks.

Once run in non-inverted mode the CCD60 only generates significant spurious charge in the serial/multiplication registers.

A large number of the frames like the one shown above were analysed statistically and a histogram of the pixel values in the image area and the serial overscan compared. This histogram is shown on the next page overlaid with the output of the Monte Carlo model. The first thing to notice is that the model agrees very closely with the actual distribution of photon event pixel values at least towards the right side of the graph.

The second thing to notice is that the spurious charge events in the serial overscan have a different distribution to the photon/spurious charge from the image area. There are clearly less of them, as this region is un-illuminated and contains no genuine photon events, but they also show a large excess of low signal events.

This low signal excess can be explained if the spurious charge in the serial overscan is generated at all positions along the multiplication register. For example, if a spurious electron is generated at the start of the L3 register then it will be amplified by the maximum amount possible. However, it is more likely that it will be generated at some intermediate point along the gain register where it will not subsequently be transferred through all stages of amplification and will therefore give rise to a low signal event.

This further complicates the question of where is the best level at which to threshold the photon events. Rising the threshold slightly will reduce the false events from spurious charge generated in the serial register quite sharply whilst only slightly reducing the number of genuine photon detections. Each photon counting system will have to be individually tuned , depending on its intrinsic amplifier noise and spurious charge levels. The intensity of the source will also be a factor.

Histogram analysis of weakly illuminated images and comparison with Monte Carlo model.

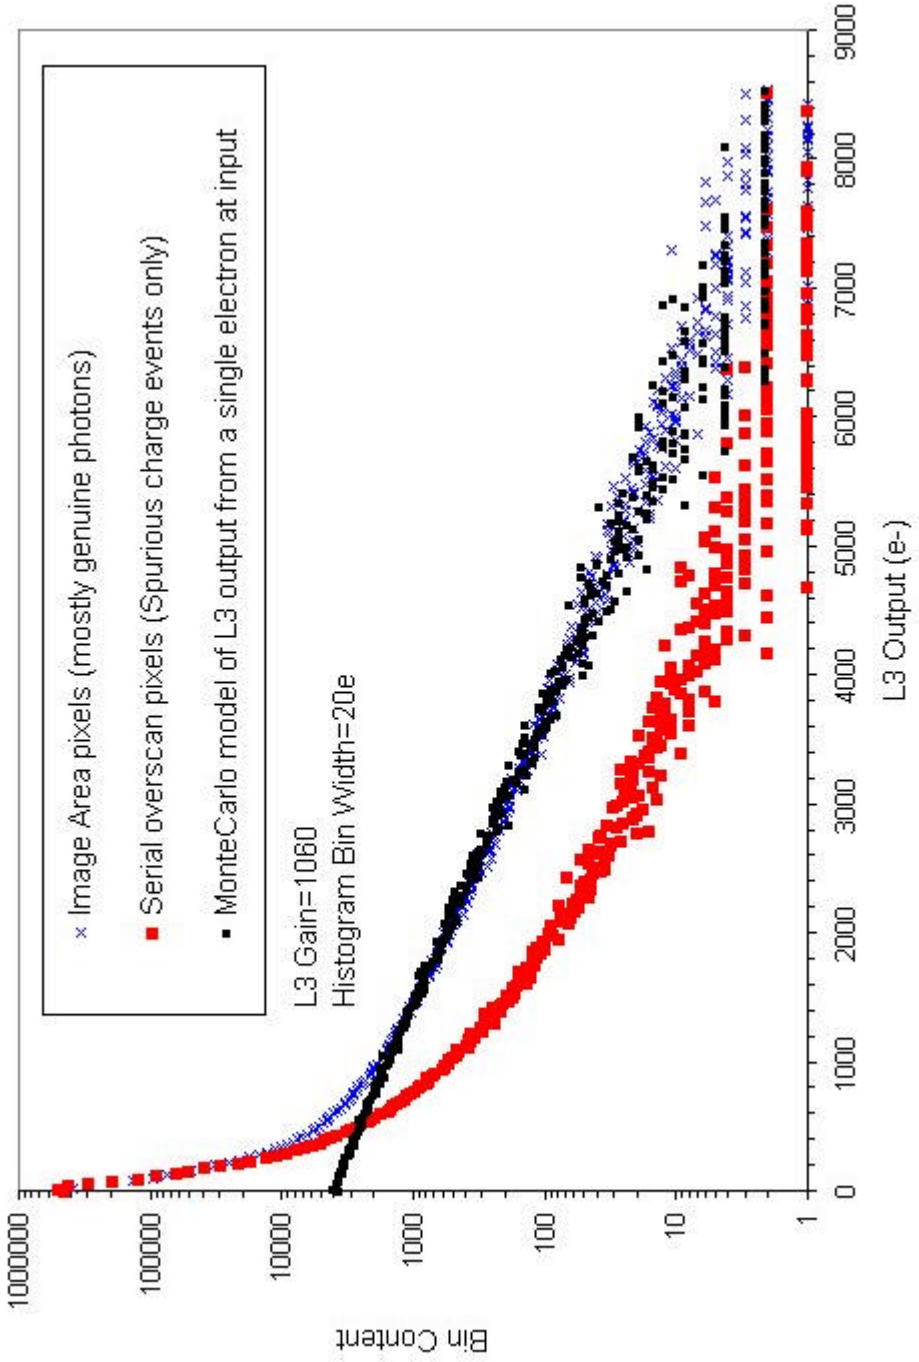

The curve below shows exactly how the detected fraction of single photon events varies with the threshold level. The curve was derived from the Monte Carlo result obtained earlier. Placing the threshold at the 0.1e level (as referenced to the input of the L3 register) should catch 90% of the photo-electrons. Placing the threshold any higher will reduce the *effective* QE of the chip unacceptably, as some of the photo-electron events will fall under the threshold.

How photon counting efficiency changes with threshold position.

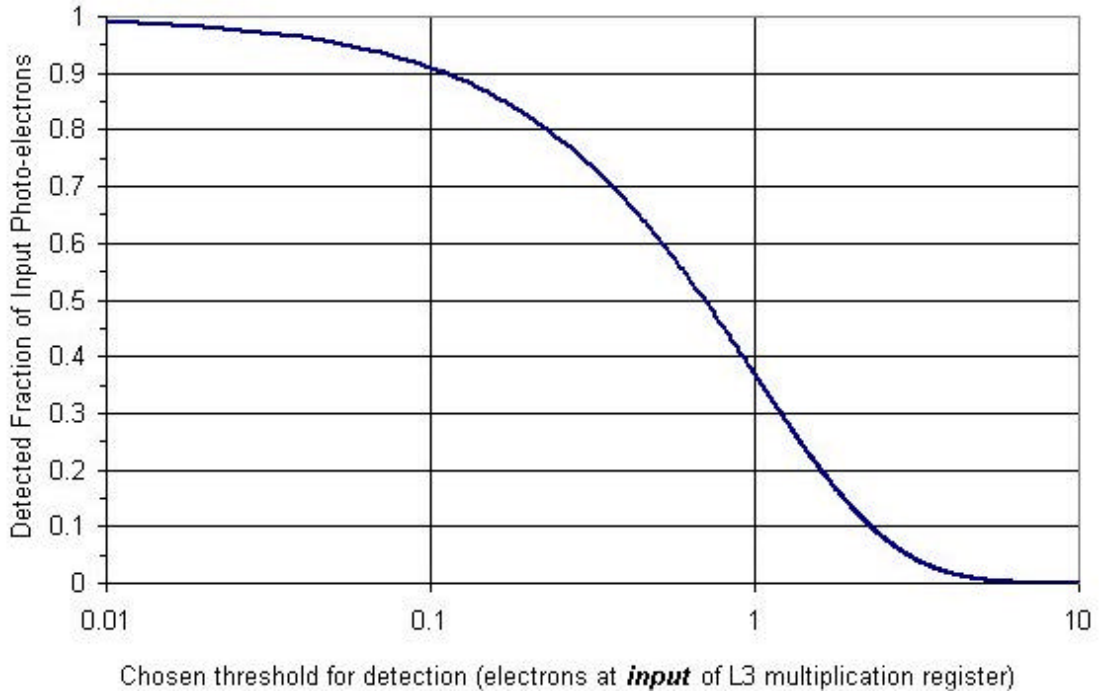

At the lower end it is important that the threshold should not lie within the intrinsic *output* referenced noise of the amplifier. The RMS noise of the CCD60 amplifier in the test camera was 50e. To avoid intrinsic noise triggering false photon events we need the threshold to be at least 5 times higher than this i.e. 250e output referenced. To ensure we detect 90% of the photon events, this 250e level at the output needs to correspond to 0.1e at the input of the L3 gain register. This places a constraint on the minimum L3 gain permissible if we are to photon count efficiently , in this case 2500. E2V do not recommend gains in excess of 1000 so a low intrinsic noise output amplifier and good shielding practice continue to be important even in this regime.

The spurious charge in the CCD60 affected about 1 in 25 pixels. Its effect on the SNR of scientific observations is the same as that of a weak background illumination. It should be noted that spurious charge will worsen in its effect as the frame rate goes up since its contribution is measured *per readout* rather than per second in the case of a normal background noise source. For best SNR, the photon counting camera frame rate will need to be tuned to the source intensity.

Actual Photon Counting Performance

To test the photon counting performance it was necessary to add an ND4 filter to the front of the test camera so as to cut down the peak illumination to about 0.1 photons per pixel per frame. The test target used was the same SH spot pattern used for the WFS windowing tests. A series of 300 images of this target were read out and stored. The individual images showed no recognisable features, merely a scattering of electron events. The images were then fed into a thresholding program (threshold=0.4e input referenced, that should catch 70% of the photons) that detected the photo-electrons and added them to an output image array which is shown below:

Photon counted image (300 frames, thresholded and summed)

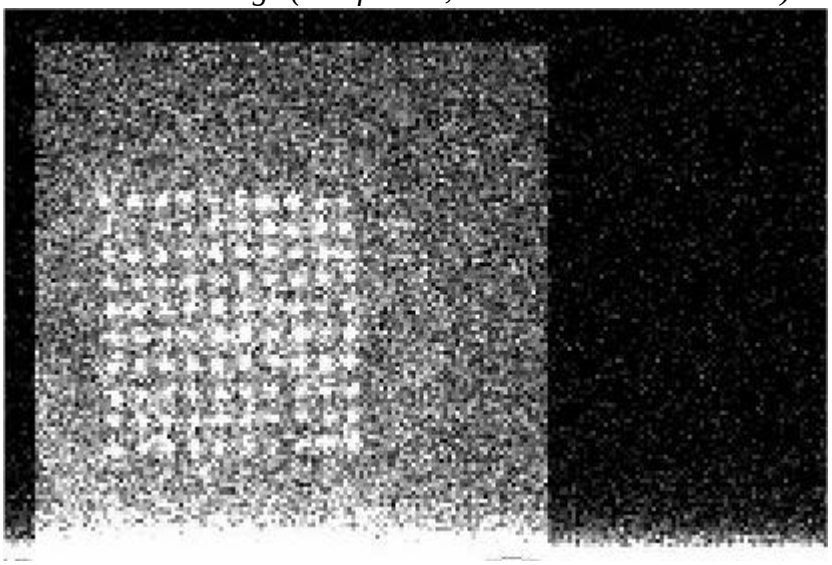

The square image area can be seen bordered on three sides by the underscan and overscan regions. A cut through a row of this image , plotted below at the same horizontal scale, shows that these non-light sensitive regions contain about 8e of signal : the spurious charge from the serial transfer clocks.

Cut through above image

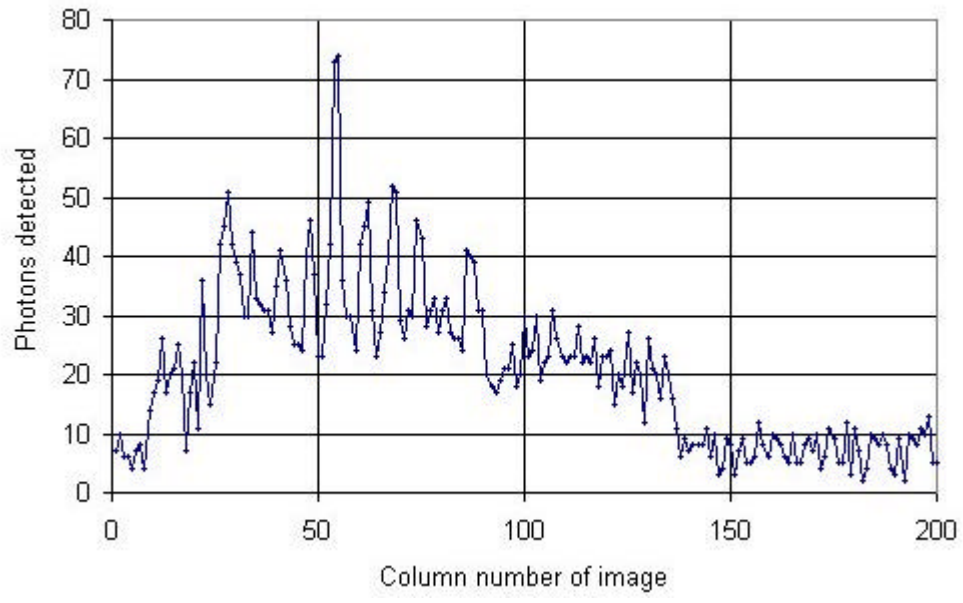

To get good dynamic range from a photon counting camera it is important that they be used at high speed. If the speed is too low or the source too bright then coincidence losses give a reduction in sensitivity.

It should be added that photon counting mode cannot be used for wavefront sensing where the required frame rates are just too high.

Conclusion

The CCD60 could be extremely useful for the NAOMI AO system, both as a SH sensor and as a NGS tip-tilt sensor. It remains to be seen whether there are any applications for photon counting operations. A Java application based around the WHT uDAS Corba grab client (being used currently for rapid simplexing) could be developed quite straightforwardly to add a photon counting capability to the WHT.

Links

Photon Counting Movie made with a CCD60:
http://www.ing.iac.es/~smt/WFS/counting/movie.htm

L3 CCD Publications:
http://www.ing.iac.es/~smt/LLLCCD/marcl3.htm

ING L3 Project Page:
http://www.ing.iac.es/~smt/WFS/WFS_index.htm

Appendix A

A program to calculate the multiplication noise using the Monte Carlo Method. A charge packet is followed through the multiplication register and its size at the output added to a histogram. The program finishes with the histogram loaded with the outcome of 50,000 such trials. The input charge packet varies around a mean entered as a parameter. The actual size of the input charge packet varies around this mean with the Poissonian distribution, as it would in a real flat field image.

```
//////////////////////
// Simon Tulloch        //
// Isaac Newton Group   //
// La Palma             //
// Aug 2003             //
// smt@ing.iac.es       //
//////////////////////

#include <math.h>
#include <stdlib.h>
#include <stdio.h>
#include <time.h>
#include <string.h>
#include <gsl/gsl_histogram.h>
#include <gsl/gsl_randist.h>
#include <gsl/gsl_rng.h>

// Program to calculate the histogram of output values from
// electrons entering the CCD60 gain register
// Expects 2 parameters: 1st is the mean input electrons, 2nd is gain.
// The mean is turned into a poissonian distribution so that each trial
// mimics a seperate pixel
// in a real flat field image of even illumination.

// gcc myhist2.c  -static  -I/usr/local/include/gsl -lgsl -o myhist2 -lm
// Gnu Scientific Library required for compilation.
// To get library, see http://www.gnu.org/software/gsl/

int main(int argc, char *argv[])
       {
       int start_electrons=atoi(argv[1]);
       int overall_gain=atoi(argv[2]);
       int stages=520;
       int trials =50000;
       float gain_per_stage=pow(10,(log10(overall_gain)/stages));
       float thresh =(RAND_MAX*(gain_per_stage-1));
       int stage_num=0;
       int electron=0;
       unsigned long total=0;
       int seed=0;
       int rn=0;
       int x;
       float randnum=0.0;
       int count;
       unsigned long inp_electrons=start_electrons;
       double mu;
       float y=0.0,check=0.0;

       // setup histogram bins so that each bin is 1 electron wide
       // (referenced to input)

       gsl_histogram *h = gsl_histogram_alloc(2*start_electrons);
       gsl_histogram_set_ranges_uniform(h,0,2*start_electrons*overall_gain);
```

```
	// start Poissonian RNG
	const gsl_rng_type * T;
	gsl_rng * r;
	gsl_rng_env_setup();
	T=gsl_rng_default;
	r=gsl_rng_alloc(T);

	printf("\nMonte Carlo Simulation of An L3 Electron Multiplication
Process\n");
	printf("running.............................\n");
	seed=time(NULL)%3600;
	for(count=1;count<=trials;count++) {
		mu=start_electrons;
		total=gsl_ran_poisson(r,mu);	// Mean into a Poiss. Dist.
		srand(seed+count*10);
		// Now follow the input electron charge packet through each
		// stage of the L3 register
		for(stage_num=1;stage_num<=stages;stage_num++) {
			inp_electrons=total;
			// Follow each electron in the packet from one gain
			// stage to the next.
			for(electron=1;electron<=inp_electrons;electron++) {
				rn=rand();
				if (rn<thresh) total++;
				}
			}
		gsl_histogram_increment(h,total);

		}
	printf("Input electrons have a POISSONIAN distribution\n");
	printf("Input Electrons(mean)=%d\n",start_electrons);
	printf("L3 gain=%d\n",overall_gain);
	printf("Sigma of Output signal=%g e-\n",gsl_histogram_sigma(h));
	printf("Output e- from  L3 register (as a fraction of the
total)\nwould be :\n\n");
	printf("Bin Value (e-)  Bin content\n");
	for (x=0;x<2*start_electrons;x++) {
		// Read out bin values and recalculate bin boundaries
		// referenced to output of L3 register
		y=(gsl_histogram_get(h,x)/trials);
		check+=y;
		printf("%d\t\t%g\n",x*overall_gain,y);
		}
//	printf("Check=%g\n",check);
	gsl_histogram_free(h);
	return(0);
}
```

//

Appendix B

The following setup voltages were used with the CCD60:

Parallel Clocks	+6/-5V
Serial Clocks	+10/0V
Dump Gate	+10/0V
RϕiDC	+3V
OD	+30V
RD	+20V
VSS	+2.5V
OG	+3V
ABD	+22V
IG	–5V

Appendix C

Test Camera to SDSU wiring schedule:

CCD60 External Wiring.

SECOND ATTEMPT W

Function	Volts	CCD	Cryostat microD	Cryostat 2041	E2V Board 9 way D	E2V Board 10 way IDC	SDSU Clock Board	SDSU Bias Board	
I1		5	30	K			13		
I2		23	11	D			15		
S1		2,22,24	27,10,12	G			19		
S2		1,21	26,9	H			17		
DG		20	8	L			18		
PhiR		7	32	B			1		
R1		16,18	5,7	E			3		
R2		15,19	3,6	F			5		
R3		17	4	C			7		
RPhiDC		11	36	J			9		
RPhiHV		12	37	A	6	1			
VSS		6,8	31,33	S				12	<R143 on video board dropped to 10R
OS		9	34	N				SMA1	
ABD		3	28	V				2	
IG		4	29	P				11	
RD		13	2	U				3	
OD		10	35	T				1	
OG		14	1	R				9	
Chassis					7		24,26,25	13	
R1-TTL					2	5	2		
R2HV TTL					3	6	4		
GAIN					4	4		10	
12V						8			
0V						7			
GND						3			
spare GND				W				25	

OS and RPhiHV are in coax cables

smt jan 03 2003