

ING Technical Note 132

Simon Tulloch

Department of Physics & Astronomy
The University of Sheffield

October 25th 2010

Contents

1	Introduction	4
2	Key EMCCD concepts	4
2.1	EMCCD structure	4
2.2	Multiplication noise	5
2.3	Gain	5
2.4	Clock-induced charge	9
2.5	Non-inverted mode operation	10
3	Modes of EMCCD operation	11
3.1	Conventional mode	11
3.2	Linear mode	11
3.3	Photon counting mode	12
3.4	Optimum choice of mode	13
4	Modelling photon-counting performance	13
4.1	What was modelled	13
4.2	Modelling of CICIR	15
4.3	Modelling of realistic EMCCD images	15
4.4	Testing of the model	16
4.5	Optimum PC threshold	16
4.6	Simplification of SNR equation in PC mode	18
4.7	SNR predictions from model	18
5	Recipe for using QUCAM2 and QUCAM3 on ISIS	20
5.1	Worked example of the recipe	24
5.2	Binning	24
5.3	Phase folding	25
A	Symbols used	26

B Fractional charge of CICIR	26
C SNR of an ideal photon counter	27
D SNR of a photon-counting EMCCD	28
E Sky backgrounds	29

1 Introduction

Conventional CCD detectors have two major disadvantages: they are slow to read out and they suffer from read noise. These problems combine to make high-speed spectroscopy of faint targets the most demanding of astronomical observations. It is possible to overcome these weaknesses by using electron-multiplying CCDs (EMCCDs). EMCCDs are conventional frame-transfer CCDs, but with an extended serial register containing high-voltage electrodes. An avalanche of secondary electrons is produced as the photon-generated electrons are clocked through this register, resulting in signal amplification that renders the read noise negligible. Two such cameras : QUCAM2 and QUCAM3 are now available for use on WHT/ISIS. Their performance is almost identical with QUCAM3 displaying slightly lower noise due to CIC (see Section 2.4) Using a combination of laboratory measurements with the QUCAM2 EMCCD camera and Monte Carlo modelling, we show that it is possible to significantly increase the signal-to-noise ratio of an observation using these cameras, but only if they are utilised correctly. We also show that even greater gains are possible through the use of photon counting. We present a recipe for astronomers to follow when setting up QUCAM2 and QUCAM3 on ISIS which ensures that maximum signal-to-noise ratio is obtained.

2 Key EMCCD concepts

Although the principle of operation of an EMCCD is very similar to that of a conventional CCD, there are some additional features that need to be considered.

2.1 EMCCD structure

The structure of an EMCCD (Figure 1) has already been described in some depth by Mackay et al. (2001), Tulloch (2004), Marsh (2008) and Ives et al. (2008). Photo-electrons are transferred into a conventional CCD serial register, but before reaching the output amplifier they pass through an additional multi-stage register (known as the *electron-multiplication* or *EM register*) where a high-voltage (HV) clock of $> 40\text{V}$ produces a multiplication of the photo-electrons through a process known as impact ionisation – see Figure 2. The EM output amplifier is similar to that found in a conventional CCD but is generally faster and hence suffers from increased read noise. Nevertheless a single photo-electron entering the EM register will be amplified to such an extent that the read noise is rendered insignificant and single photons become clearly visible. Most EMCCDs also contain a conventional low-noise secondary amplifier at the opposite end of the serial register; use of this output transforms the EMCCD into a normal CCD.

Most EMCCDs are of frame-transfer design – see Figure 1. Here, half the chip is covered with an opaque light shield that defines a storage area. The charge in this storage area can be transferred independently of that in the image area. This allows an image in the storage area to be read out concurrently with the integration of the next image, with just a few tens of millisecond dead time between exposures. Incorporating frame-transfer architecture into any CCD will greatly improve observing efficiency in high frame-rate applications where the readout time is comparable to the required temporal resolution (Dhillon et al. 2007). In the case of an EMCCD the use of frame-transfer architecture is essential otherwise the SNR gains will be nullified by dead time.

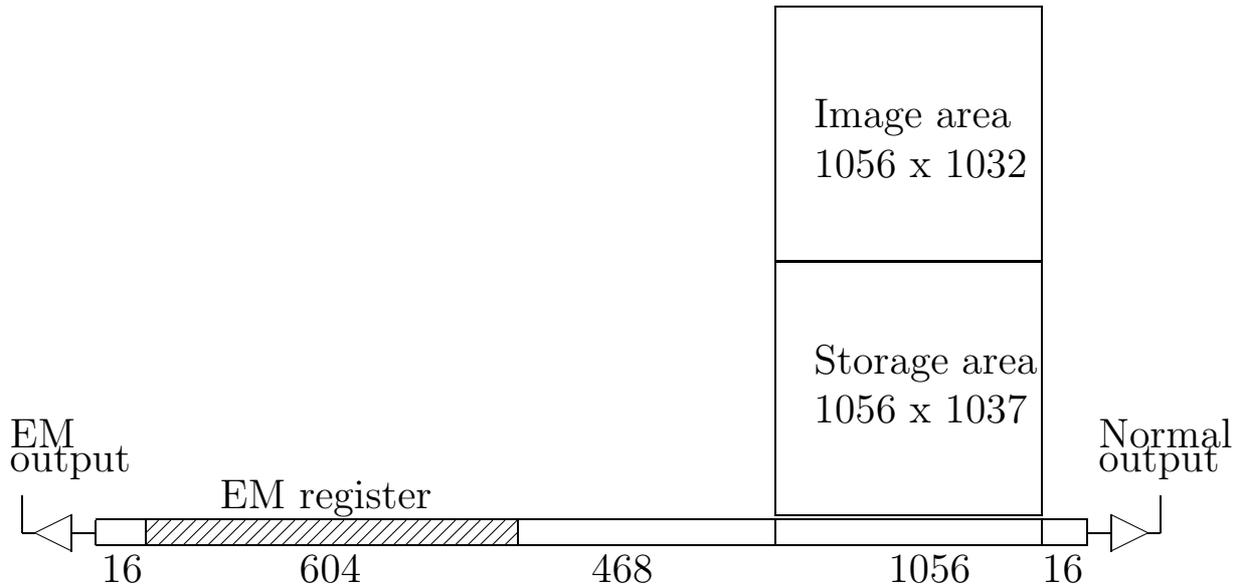


Figure 1: Schematic structure of the EMCCDs used in QUCAM2 and QUCAM3. Photo-electrons produced in the image area are vertically clocked downwards, first into the storage area, and then into the 1056-pixel serial register. For EMCCD operation, the charge is then horizontally clocked leftwards, through the 468-pixel extended serial register and into the 604-pixel EM register, before being measured and digitised at the EM output. For conventional CCD operation, the charge in the serial register is horizontally clocked rightwards to the normal output.

2.2 Multiplication noise

A single photo-electron entering the EM register can give rise to a wide range of output signals. This statistical spread constitutes an additional noise source termed multiplication noise (Hollenhorst 1990). Basden et al. (2003) derive the following equation describing the probability $p(x)$ of an output x from the EM register in response to an input of n (an integer) photo-electrons:

$$p(x) = \frac{x^{n-1} \exp(-x/g_A)}{g_A^n (n-1)!}. \quad (1)$$

This is evaluated for several values of n and with g_A the EM gain (see Section 2.3) equal to 100, in Figure 3, which shows that for an output signal of $300 e^-$, the input signal could have been either 3 or $4e^-$ with almost equal probability. The overall effect is to double the variance of the signal, which is statistically equivalent to halving the QE of the camera (see Section 3.2). In the photon-noise dominated regime this means that conventional CCDs will actually give a higher performance. It is in the read-noise dominated regime that EMCCDs come into their own, where their lack of read noise more than compensates for the effects of multiplication noise. Note also that for signal levels where there is a low probability of a pixel containing more than one photo-electron it is possible to use a photon-counting analysis of the image to remove the effect of multiplication noise (see Section 3.3).

2.3 Gain

Astronomers typically refer to the gain (or strictly speaking system gain, g_S) of a CCD camera as the number of photo-electrons represented by 1 analogue-to-digital unit (ADU) in the raw

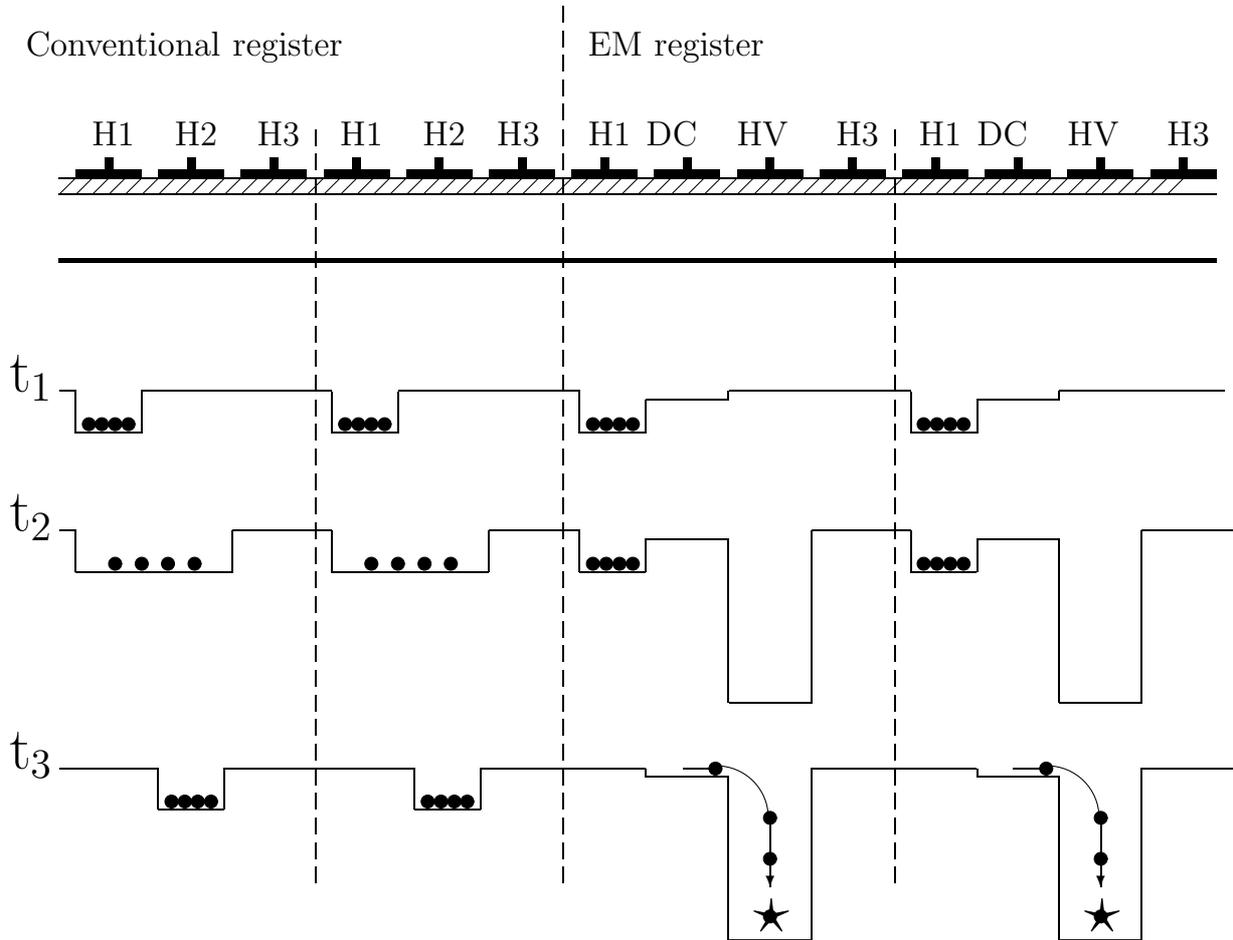


Figure 2: The geometry of the serial and EM registers, showing how multiplication occurs. The top part of the diagram shows a cross-section through the EMCCD structure with the electrode phases lying at the surface. Below this are three snapshots showing the potential wells and the charge packets they contain at key moments (t_1, t_2, t_3) in the clocking process. At t_3 the photo-electrons undergo avalanche multiplication as they fall into the potential well below the HV clock phase. Note that this diagram does not show a complete pixel cycle.

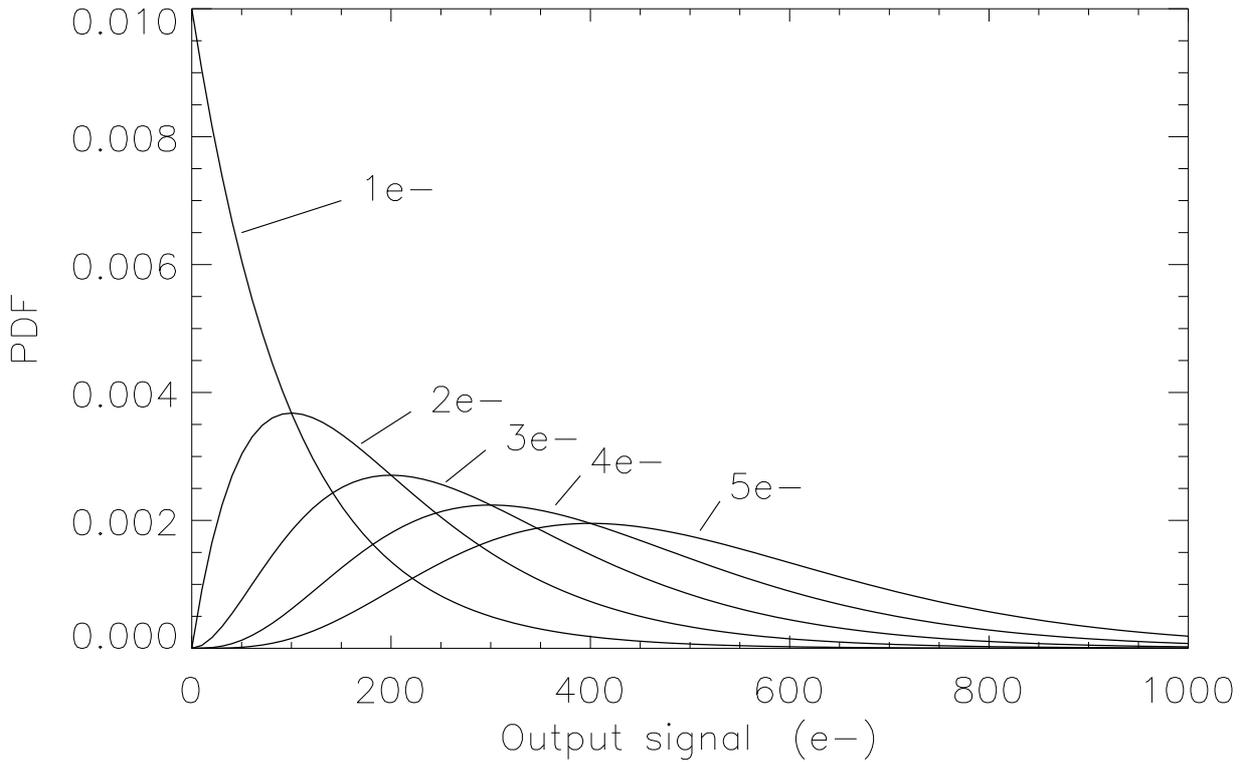


Figure 3: Output of an EM register with $g_A = 100$ in response to a range of inputs from 1 to $5e^-$. The y -axis shows the probability density function (PDF) of the output signal, i.e. the fraction of pixels lying within a histogram bin.

image, i.e. it has units of e^-/ADU . An EMCCD camera has another gain parameter that we need to describe: the avalanche multiplication gain g_A (hereafter referred to as the *EM gain*), and there is a risk of confusion here with g_S . EM gain is simply a unitless multiplication factor equal to the mean number of electrons that exit the EM register in response to a single electron input. It is hence related to g_S by the relation $g_A = g_{S0}/g_S$, where g_{S0} is the system gain (in units of e^-/ADU) measured with the EM gain set to unity.

To measure the various EMCCD gain parameters, we need to first turn off the EM gain by reducing the HV clock amplitude to 20V. At this level the EM register will then behave as a conventional serial register, i.e. 1 electron in, 1 electron out. We can now measure g_{S0} , just as we would with a conventional CCD (there are various methods, for example the photon transfer curve, Janesick 2001). To measure g_S , we weakly ($< 0.1e^-/\text{pix}^{-1}$) illuminate an EMCCD with a flat field so as to avoid a significant number of pixels containing more than a single electron. A histogram of such an image, with a vertical log scale, is shown in Figure 4. In this histogram, the pixels containing photo-electrons lie along a curve that is linear except at low values where the effects of read noise become dominant. Figure 4 also shows a least-squares straight-line fit to the linear part of the histogram: the gradient of the line is equal to $-g_S$ (Tulloch 2004). In this particular case the camera had a system gain $g_S = 0.005 e^-/\text{ADU}$, i.e. a single photo-electron entering the EM register would produce a mean signal of 200 ADU in the output image. When discussing noise levels in an EMCCD it is more convenient to express this in units of input-referenced photo-electrons (e^-_{pe}). So in the above example, if the read noise is 5 ADU this would be quoted as $5 \times 0.005 = 0.025e^-_{pe}$.

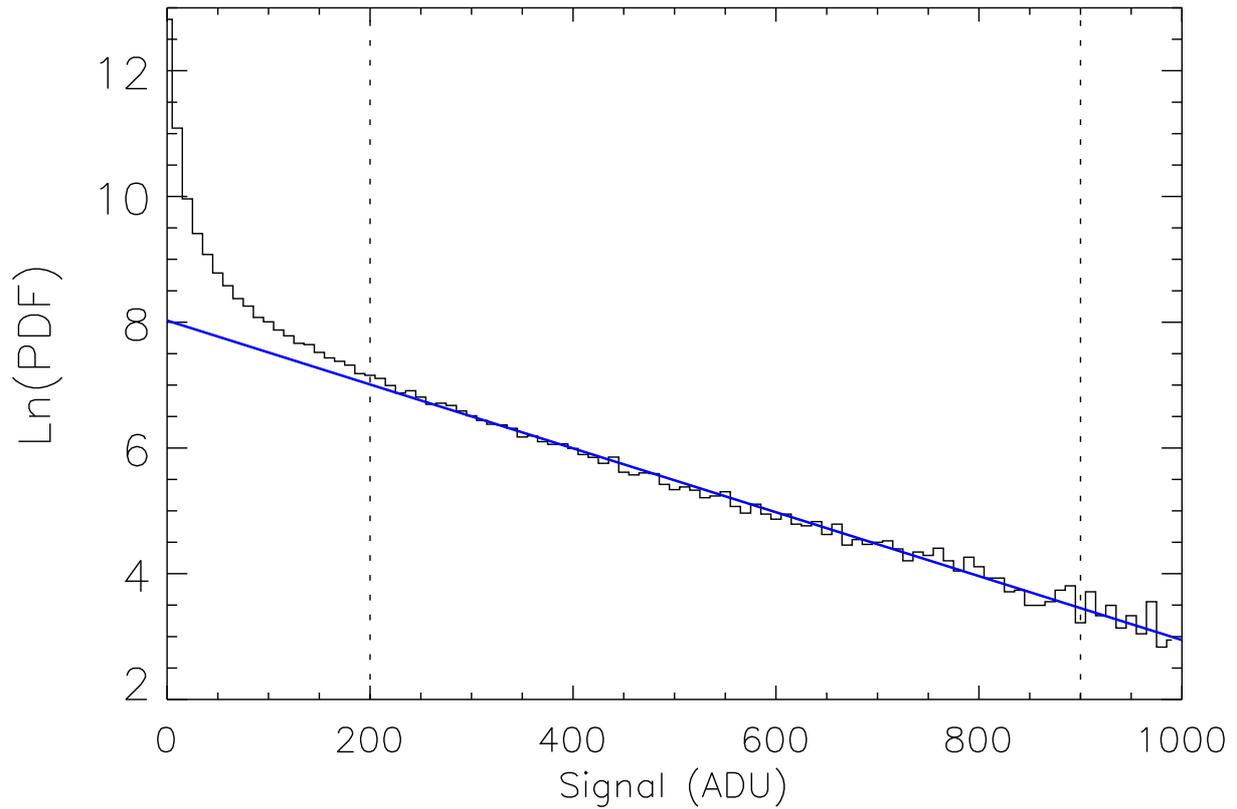


Figure 4: The histogram, plotted on a \log_e vertical scale, of pixels in a weakly-illuminated EMCCD image. The solid line is a least-squares linear fit to the photo-electron events lying between the two vertical dotted lines. The slope of this fitted line can be used to calculate the system gain g_S of the camera in e^-/ADU .

2.4 Clock-induced charge

Clock-induced charge (CIC) is an important source of noise in EMCCDs. Its contribution needs to be minimised. It consists of internally-generated electrons produced by clock transitions during the readout process. CIC is visible in EMCCD bias frames as a scattering of single electron events which at first sight are indistinguishable from photo-electrons. It is only when a histogram is made of the image that they appear different.

CIC is dependent on a number of factors. The amplitude of the clock swings is relevant, as is the temperature (Janesick 2001). At first sight, one might assume that, since CIC is proportional to the total number of clock transitions a pixel experiences during the readout process, a pixel lying far from the readout amplifier should experience a higher level of CIC. One would then expect CIC gradients in both the horizontal and vertical axes of the image. This would be true if the CCD is entirely cleared of charge prior to each readout. However, this is never the case. One must consider that prior to each readout the chip has either been flushed in a clear operation or read out in a previous exposure. These operations leave a ‘history’ of CIC events in the CCD pixels prior to our subsequent measurement readout. The distribution of these historical events will be higher the closer we get to the output amplifier since the CIC charge residing in these pixels will have accumulated through a larger number of clock transitions than for pixels more distant from the amplifier. When these historical events are added to the events created in the most recent measurement readout the overall effect is that each pixel of the image will have experienced the same number of clock transitions regardless of its position, and the resulting CIC distribution will be flat.

CIC is produced by both vertical and horizontal clocks, as well as the clocks within the EM register. Vertical CIC can be virtually eliminated through the use of non-inverted mode clocking (where the clock phases never fall more than about 7 volts below the substrate of the CCD, the exact value depending on the CCD type). Serial-clock CIC can be reduced by using lower clock amplitudes. CIC generated within the EM register (described from here on as CICIR) is harder to remove since any changes to the EM clock amplitudes produces large changes in EM gain. A well optimised EMCCD will have its performance limited only by CICIR. Other noise sources such as amplifier read noise, dark current, image-area and serial-register CIC should all have been reduced to an insignificant level with respect to CICIR. This optimisation process for QUCAM2 is described in Tulloch (2010). As an example of this, the histogram of a bias image from the QUCAM2 EMCCD camera is compared with the histograms of two other images generated using a Monte Carlo model (see Section 4) in Figure 5. The first of these models consists of CIC originating prior to the EM register, the second consists of CIC originating at random positions within the EM register. The latter will on average experience less multiplication than the former since it will pass through fewer stages of the multiplication register, producing a histogram that shows an excess of low value pixels (as shown in Figure 5). Various other models were created with mixes of the two noise sources. The best fit was found to correspond to 85% CICIR and 15% pre-EM-register CIC, demonstrating that QUCAM2 has been well optimised (at least as far as its CIC performance is concerned with some other parameters such as CTE remaining non-optimal.) and is dominated by CICIR.

Pre-EM register CIC electrons should not be confused with dark current generated during the readout. These can easily outnumber CIC events if the operational temperature is too high, if the CCD controller has recently been powered on or if the CCD has recently been saturated to beyond full-well capacity. The recovery time for these last two cases is approximately 2 hours, and an accurate measurement of CIC should not be attempted until after such a period.

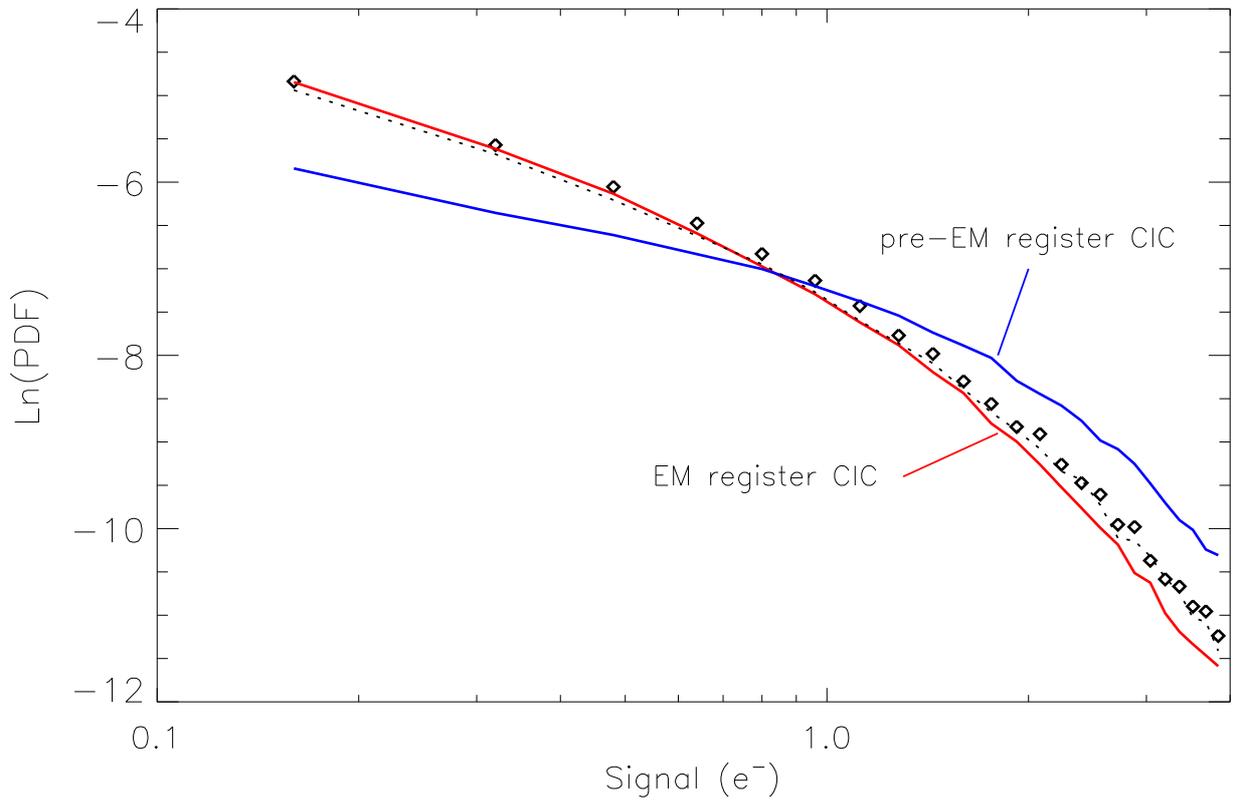


Figure 5: Histogram (diamonds) of a QUCAM2 bias image compared with two models: pure in-register CIC and pure pre-EM register CIC. A fit to the data, consisting of 85% CIC generated within the EM register and 15% generated prior to the EM register is shown as a dotted curve.

CIC electrons generated within the EM register do not contribute as much charge to an image as a CIC electron generated prior to the register: CICIR has a fractional charge when expressed in units of input-referenced photo-electrons. If we assume that CICIR is generated randomly throughout the EM register then we can calculate the average charge ν_C that it will contribute to an image pixel. The calculation, derived in Appendix B, shows that:

$$\nu_C \approx \frac{B_C}{\ln(g_A)}, \quad (2)$$

where B_C is the mean number of CICIR events experienced by a pixel during its transit through the EM register.

2.5 Non-inverted mode operation

If the vertical clock phases are held more than about 7V below substrate then the surface potential of the silicon underlying the phases becomes pinned at the substrate voltage (so-called Inverted-mode operation or IMO). This has important consequences for both dark current and CIC. If during the read-out of the image the clock phases never become inverted (so called Non-inverted mode operation or NIMO) then the CIC is greatly reduced (in the case of QUCAM2 it fell from $0.2e^- \text{ pixel}^{-1}$ to a level that was unmeasurable even after 50 bias frames had been summed). At higher operational temperatures NIMO will cause an approximate 100-fold increase in dark current: something that will negate any gains from lowered CIC. The trade-off between CIC and dark current does not, however, hold at lower temperatures. For QUCAM2 at

an operational temperature of 178K, the NIMO dark current was $1.5e^- \text{pixel}^{-1} \text{hour}^{-1}$, a value that was found to be constant for exposure times of up to 1000s. Switching to IMO reduced this to $\sim 0.2e^- \text{pixel}^{-1} \text{hour}^{-1}$: a value somewhat difficult to measure since it is ~ 4 times below the current delivered by cosmic ray events. This demonstrates that operation at cryogenic temperatures permits NIMO without a loss of performance from high dark current. It should be noted that at higher temperatures, such as those experienced by Peltier-cooled CCDs, the situation becomes more complex since dark current is no longer constant with exposure time. This effect was seen during the optimisation of QUCAM2 when operated experimentally at 193K. Using NIMO, the dark current for 60s exposures was measured at $12e^- \text{pixel}^{-1} \text{hour}^{-1}$ whereas for 600s exposures it was $40e^- \text{pixel}^{-1} \text{hour}^{-1}$.

3 Modes of EMCCD operation

EMCCDs can be utilised in three separate modes, each offering optimum SNR in certain observational regimes. In this section the equations describing the SNR in these three modes are shown.

3.1 Conventional mode

The SNR obtained through the conventional low-noise amplifier is given by:

$$\text{SNR}_C = \frac{M}{\sqrt{M + \nu_C + D + K + \sigma_N^2}}, \quad (3)$$

where M is the mean signal per pixel from the source, σ_N is the read noise from the conventional amplifier, ν_C the mean CICIR per pixel, D the dark charge and K the charge received from sky photons.

3.2 Linear mode

In linear mode, the digitised signal from the EM output is interpreted as having a linear relationship with the photo-electrons, as is usual for a CCD. The SNR obtained through the EM output is then given by:

$$\text{SNR}_{lin} = \frac{M}{\sqrt{2 \cdot (M + \nu_C + D + K) + (\sigma_{EM}/g_A)^2}}. \quad (4)$$

The factor of 2 in the denominator accounts for the multiplication noise (see section 2.2). The derivation of this factor can be found in Marsh (2008) and Tubbs (2003). The read noise in the EM amplifier σ_{EM} will typically be tens of electrons due to its higher bandwidth but its contribution to the denominator is rendered negligible by the use of high EM gain, g_A . High speed means high read noise so high frame rate cameras will need higher gains than the more leisurely QUCAMs (1.6s read-out time in EM mode).

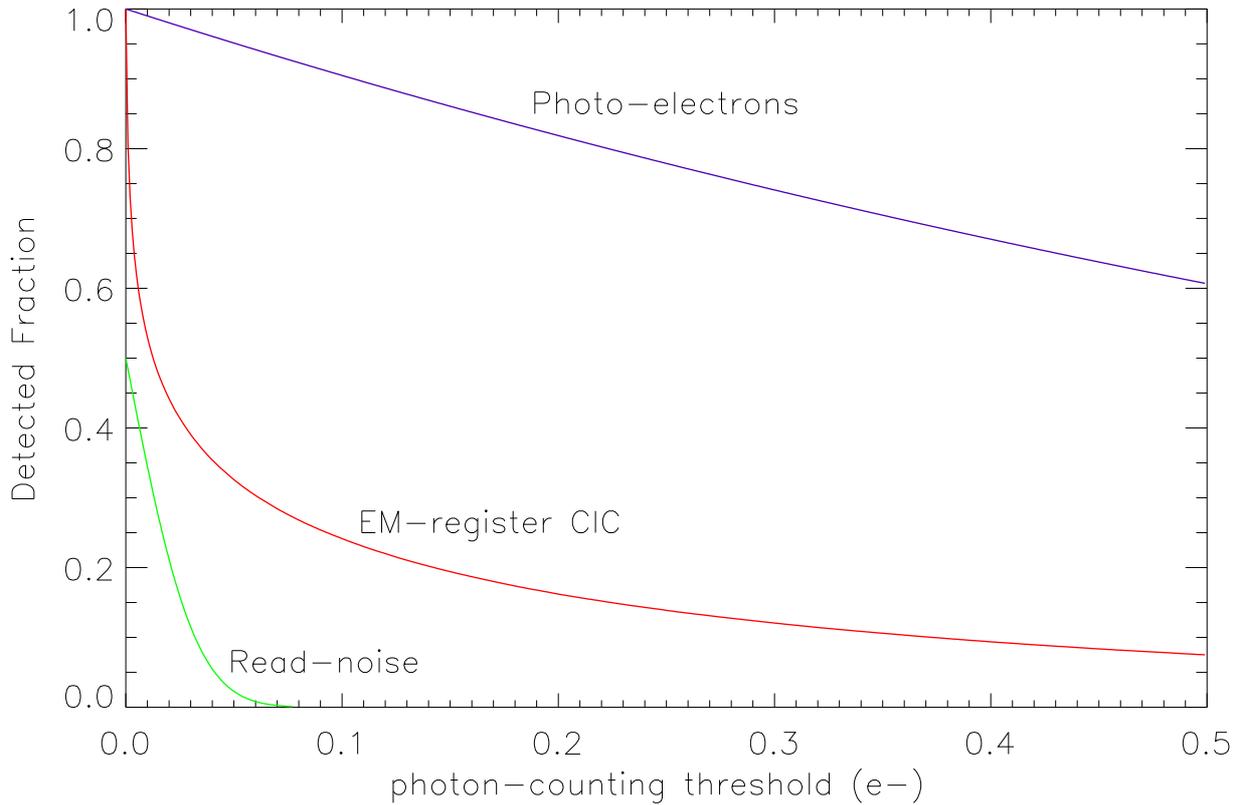


Figure 6: The effect of the PC threshold value on the detected fraction of the signal, the CIC and the read noise in an EMCCD image. In this case the read noise $\sigma_{EM} = 0.025e_{pe}^-$ and $g_A = 2000$.

3.3 Photon counting mode

In PC mode we apply a threshold to the image and interpret any pixels above it as containing a single photo-electron. This leads to coincidence losses at higher signal levels where there is a significant probability of a pixel receiving two or more photo-electrons, but at weak signal-levels where this probability is low, PC operation offers a means of eliminating multiplication noise (Plakhotnik et al. 2006) and obtaining an SNR very close to that of an ideal detector. When photon counting we must aim to maximise the fraction of genuine photo-electrons that are detected whilst at the same time minimising the number of detected CICIR and read noise pixels. Figure 6 shows how this can be done. The graph shows us that if we set a photon-counting threshold of, say, $0.1e_{pe}^-$ (a value that was later found to be optimum, see Section 4.5) we will detect 90% of photo-electrons but only 23% of the CICIR. False counts from the read noise will be negligible.

The SNR of an ideal photon-counting detector, including the effects of coincidence losses, is given by:

$$\text{SNR}_{pc} = \frac{M}{\sqrt{e^M - 1}}. \quad (5)$$

This equation (derived in Appendix C) needs to be modified to accurately describe a photon-counting EMCCD since it makes no allowance for the complex effects of CICIR and choice of threshold level. Since the distribution of CICIR and photo-electron events are different, the SNR can vary greatly depending on the precise choice of threshold. These complexities have been explored in detail by modelling (see Section 4) but, in short, the following SNR relation

(derived in Appendix D) is found to hold, assuming a PC threshold of $0.1e_{pe}^-$ (close to optimum, see Section 4.5):

$$\frac{0.9M}{\sqrt{\delta}\sqrt{\exp[(0.9(M+D+K)/\delta) + 0.23\ln(g_A)\nu_C] - 1}}. \quad (6)$$

It should be noted that the maximum possible SNR in a single PC frame is ≈ 0.8 (see Figure 11) and it is then necessary to average many frames to arrive at a usable image (i.e. one with $\text{SNR} > 3$). The number of frames that would need to be ‘blocked’ together in this fashion is given the symbol δ in Equation 6.

3.4 Optimum choice of mode

Considered from the point of view of maximising the SNR per pixel, the choice of readout mode is quite simple. Figure 7 shows the range of per-pixel illuminations over which each mode offers the best SNR, based on the equations presented earlier in this section. A single pixel is, however, not usually the same thing as a single wavelength element in a reduced spectrum. Many additional factors affect the choice of mode, such as plate scale, seeing and sky background, and this is explored in greater depth in Section 5.

4 Modelling photon-counting performance

Modelling was required for two reasons. First there was no equation available describing the output signal distribution of pixels affected by CICIR and second to check that the assumptions underlying the derivation of Equation 6 (see Appendix C) are valid.

4.1 What was modelled

Equation 1 describes the output of an EM register for any integer number n of photo-electrons. Using the Poisson distribution it is possible to calculate the proportion of pixels that contain $n = 1, 2, 3, \dots$ photo-electrons as a function of the mean illumination M . This result can be combined with Equation 1 to yield the output distribution of the EM register for any *mean* input signal level. No equivalent relation describing the output distribution of CICIR could be found and it is here that Monte Carlo modelling is required (see Section 4.2).

The final output of the model is a pair of 3D vectors. The first of these shows the output pixel value distributions (i.e. histograms) for a wide range of signal and CICIR levels. The second is the associated cumulative distribution function (CDF) derived from the histograms in the first 3D vector. The CDF is extremely useful since it indicates the mean photon counts per-pixel that we can expect for any combination of signal and CICIR for any given PC threshold. These 3D vectors can be visualised as two cubes of normalised histogram values and CDF values. The x -axis of the cubes are labelled with a logarithmic-spaced range of CIC values and the y -axes with a logarithmic-spaced range of signal values (extending up to a maximum of $2.5e^-$). The z -axes are labelled with pixel values extending up to a maximum of $10e_{pe}^-$. In the case of the CDF cube, the z -axis units can also be interpreted as the PC threshold setting and the data values as being the mean per pixel-PC signal at that given threshold.

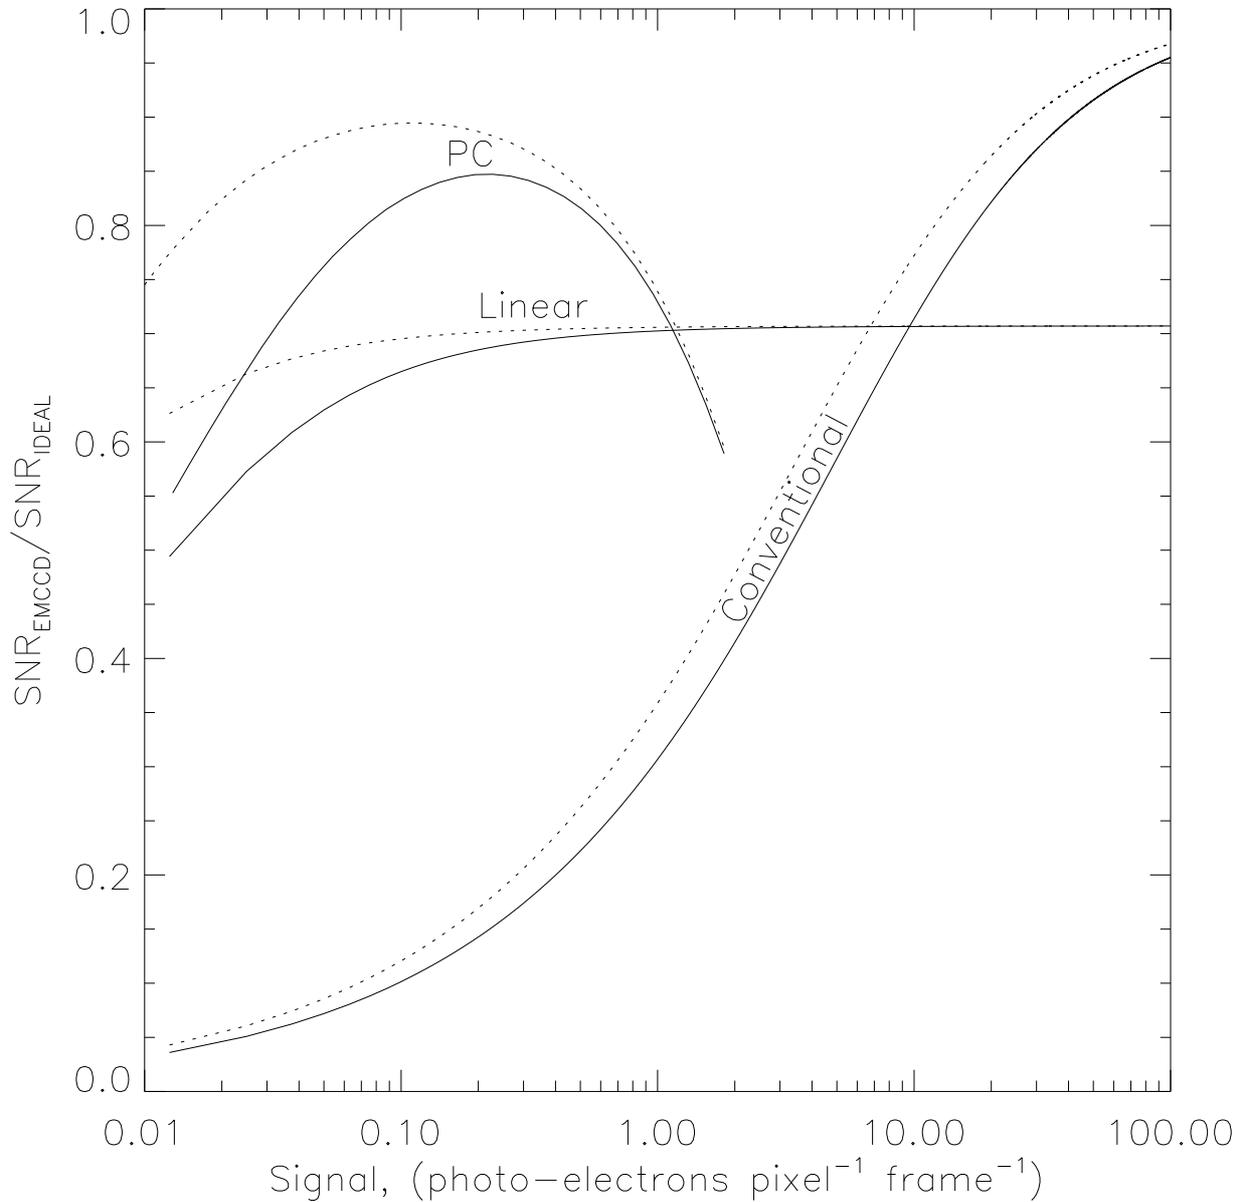


Figure 7: The relative SNR (compared to an ideal noise-free detector of the same QE) of each EMCCD mode (conventional, linear and PC) is shown over a wide range of illuminations. The solid curves show the performance of a detector similar to the QUCAMs ($\nu_C=0.013e^-$, $\sigma_N=3.1e^-$), the dotted lines show the performance that would be expected from a more highly-optimised detector ($\nu_C=0.003e^-$, $\sigma_N=2.6e^-$).

4.2 Modelling of CICIR

Since no analytical formula describing the distribution of CICIR events could be found it was necessary to do a Monte Carlo model of the EM register. Binomial statistics describe processes whose final outcome depends on a series of decisions each of which has two possible outcomes. It therefore applies to the creation of CIC as a pixel is clocked along the EM register. Synthetic images were generated containing between 1 and 6 CIC events per pixel. It was not really necessary to go any higher than 6 CIC events since the binomial distribution shows that there is an insignificant probability of more than 6 events being generated per pixel for mean CIC event levels of up to 0.3 per pixel, well beyond the useful operational range of an EMCCD (indeed QUCAM2 gave ~ 0.08 CIC events per pixel). Histograms of these 6 images were then calculated to yield a set of output CIC distributions for integer input, analogous to the distributions for photo-electron events given by Equation 1. It was then necessary to combine these histograms using the binomial distribution formula to yield the output distribution of the EM register for any *mean* CIC event level.

The model was implemented by simulating the transfer of charge through an EM register with 604 elements, one pixel at a time. At each pixel transfer a dice was thrown for each electron in the pixel to decide if a multiplication event occurred. An overall EM gain of $g_A=2000$ was used. Six thousand lines were read out in this way to get a good statistical sample of pixel values. At the start of the readout of each simulated image row, the EM register was charged with a single electron per element. This simply amounted to initialising the array representing the EM register with each element equal to 1. This was then read out, simulating the effect of charge amplification, to yield an image with width equal to the length of the register. The resulting image was then scrambled (i.e. the pixels were reordered in a random fashion) and added to its original self to yield an image containing an average of 2 CIC events per pixel. This scrambling was necessary since the raw images contained pixels with values that were approximately proportional to their column coordinate, with the pixels in the higher column numbers having higher values. Further scramble-plus-addition operations were performed to yield images with 3, 4, 5 and 6 events per pixel.

4.3 Modelling of realistic EMCCD images

The distributions of the CICIR (calculated in Section 4.2) and the distribution of the photo-electron events were then combined, together with read noise, through the use of intermediate model images. Photo-electron events were first generated in the image using a random number generator that was weighted by the distribution of the EM register output. CICIR events were then added to the image in the same manner as for the photoelectrons. Finally, read noise of $\sigma_{EM} = 0.025e_{pe}^-$ was added to every pixel in the image. These model images contained a bias region from which photo-electrons were excluded. Histograms of the image and bias regions were then calculated to yield the distributions and their CDFs. These CDFs effectively gave the mean photon-counting signal from the image and bias areas as a function of t the PC threshold. The SNR could then be determined using the following equation (derived in Appendix C):

$$\text{SNR}_{pc}(t) = \frac{-\ln[1 - CDF_I(t)] + \ln[1 - CDF_B(t)]}{\sqrt{[1 - CDF_I(t)]^{-1} - 1}}, \quad (7)$$

where CDF_I and CDF_B are the image and bias area CDFs, respectively. Since the CDFs were calculated over a wide range of threshold values it was possible to find the optimum threshold value or alternatively just calculate the SNR for any given threshold.

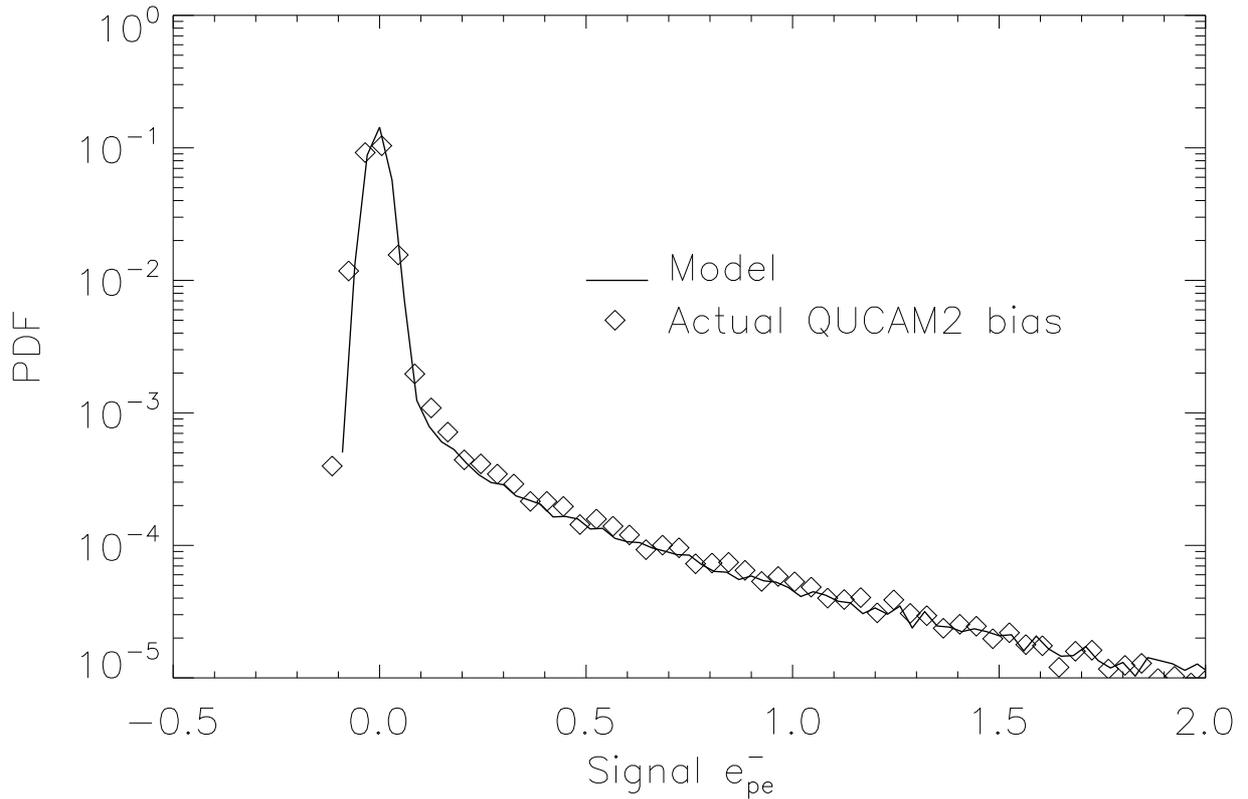


Figure 8: The plot compares the histogram of pixels in a genuine QUCAM2 EMCCD image (diamonds) with that of a model image (solid line).

4.4 Testing of the model

The model output was tested against a stack of 45 QUCAM2 bias frames of known CIC level, EM gain and system gain. The comparison between model and data is shown in Figure 8. The agreement is good, although QUCAM2 shows a slight excess of low value events compared to that predicted. This can be explained by the imperfect charge transfer in the EM register which boosts the relative number of low-value events, an effect that was not included in the model.

4.5 Optimum PC threshold

It is the read noise that sets the lower limit on the photon-counting threshold. Gaussian statistics predict that a threshold set $\sim 3\sigma$ above this noise gives 1 false count per 1000 pixels, falling to 1 pixel in 32000 if we choose a threshold of 4σ . Other groups (Daigle et al. 2006, Ives et al. 2008) have chosen quite high thresholds (5 - 5.5σ). This is a good choice as long as it is combined with a high EM gain, high enough to ensure that the mean level of a photo-electron is at least 10 times the threshold value. This high ratio between mean photo-electron level and read noise permits a threshold to be set low enough to include a majority of photo-electrons. There is a limit to how high the EM gain should be pushed, though, since it can risk damage to the chip if gain is applied during overexposure for long periods.

Figure 9 shows that the PC threshold needs to be tuned depending on the signal level, so that the maximum number of genuine photo-electrons are counted and the maximum number

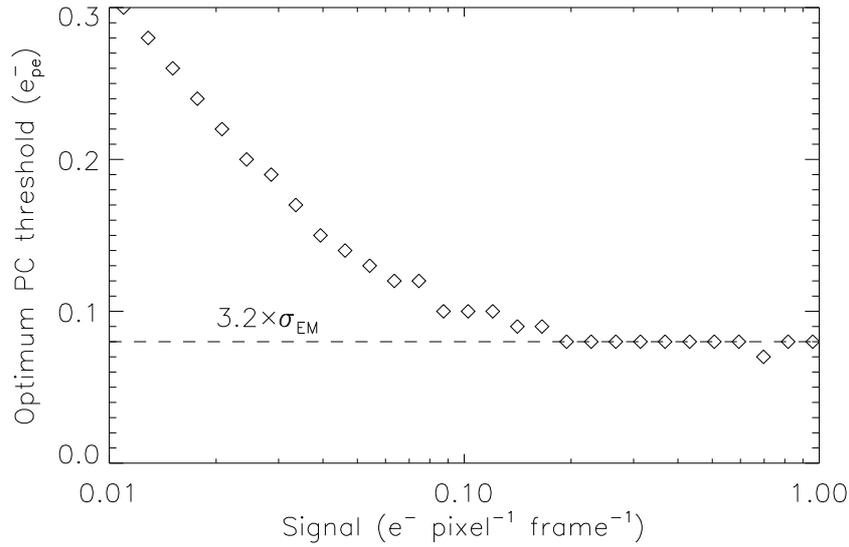


Figure 9: The PC threshold that gives maximum SNR is plotted as a function of signal level. Model images were used with parameters close to those of the QUCAMs. The read noise σ_{EM} was $0.025e^-_{pe}$ and the multiplication gain g_A was 2000.

of CICIR are rejected. Figure 9 also shows that in the case of a detector with read noise $\sigma_{EM} = 0.025e^-_{pe}$, the optimum threshold falls as low as $3.2\sigma_{EM}$. From a data-reduction point of view, using a variable threshold adds complexity but may be necessary to extract maximum SNR, particularly at low signal levels. One example of this would be the measurement of a faint emission line, the peak of which would be placed at an optimum signal level through a suitable choice of frame rate. Here, the threshold would be set low, in the region of $0.1e^-_{pe}$ according to Figure 9. The wings of this same line, which may be an order of magnitude fainter would then benefit from an increased threshold, say $\sim 0.25e^-_{pe}$. The use of an adaptive threshold would create many side effects, such as noise artifacts (the amount of background signal from CIC, sky and dark current would be modified depending on threshold setting), so would require extra data reduction effort. Note also that non-Gaussian pattern noise is a particular problem in high-speed detectors in an observatory environment which may require the threshold to be pushed higher than would otherwise be optimum. The models described so far have assumed a read noise of $0.025e^-_{pe}$ (equal to that of QUCAM2). A simulation was performed of the effect of higher ($0.05e^-_{pe}$) and lower ($0.012e^-_{pe}$) read noise on the SNR performance of an EMCCD. Whilst the higher noise definitely impinges on the photon-counting performance by forcing the threshold higher and giving a lower detected fraction of photo-electrons, the lower noise gives very little additional benefits. This can be explained by the fact that the bulk of the CICIR has a distribution that falls between $0-0.05e^-_{pe}$ and it will dominate any read noise lying within the same range.

In the specific case of QUCAM2 there is an additional reason (apart from staying above the read noise) why the threshold must be kept slightly elevated. This is the effect of poor charge transfer efficiency (CTE) in the EM register. Figure 10 shows the autocorrelation of a low-level flat-field which demonstrates the problem. The autocorrelation was performed along an axis parallel to the serial register. The slight broadening of the autocorrelation peak is indicative of less-than-perfect CTE. Using a threshold much below $0.1e^-_{pe}$ would cause complex effects from multiple counting of single electron events due to the slight tail on each event being above the threshold.

Note that Basden et al. (2003) have modelled a multiple threshold technique that can ex-

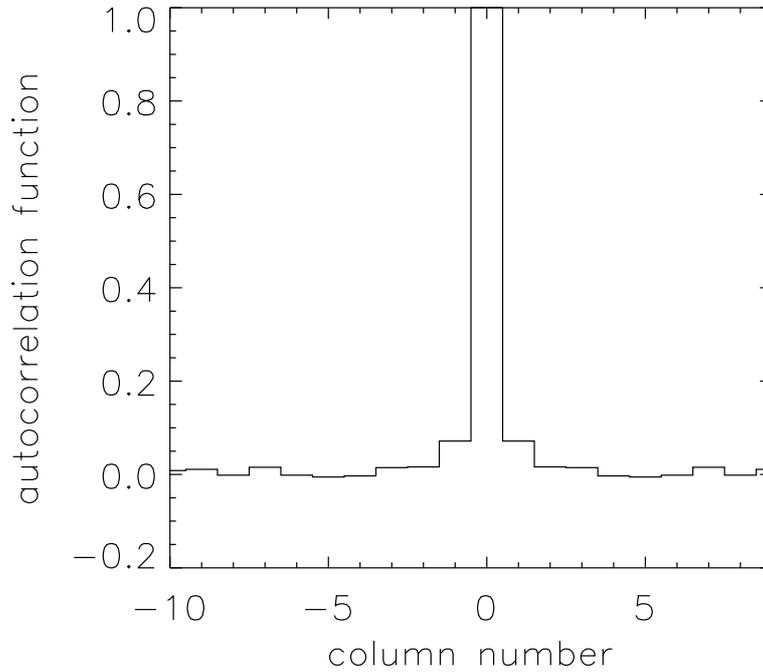


Figure 10: The autocorrelation of a weakly illuminated QUCAM2 image showing the elongation of the single electron events through CTE degradation in the EM register.

tend photon-counting operation well into the coincidence-loss dominated signal regime (see Section 3.3) whilst maintaining high performance.

4.6 Simplification of SNR equation in PC mode

The method used to calculate SNR in PC mode (Equation 7) is rather complex since it requires the analysis of both image and bias areas in a large model image. A simpler SNR formula that can be applied more generally was therefore sought. One simplification would be to consider CIC, read noise and photo-electrons separately, i.e. assume that they only interact in the digital domain after thresholding. This is, of course, an approximation and in reality there is a complex interplay between CIC events and photo-electron events in the analogue domain. For example, a photo-electron event could be ‘helped over’ the PC threshold by an accompanying CIC event, or a CIC event could be lost by occurring within an illuminated pixel. A second simplification would be to assume that, as the signal or CIC level increases, the detected fraction of these events does not change. It is thought reasonable to make these approximations since it will only seriously fail in the high signal regime where there is a high probability of coincidence losses and consequently low SNR compared to an ideal detector. The resulting simplified SNR equation (Equation 6) is derived in Appendix D. This approximation was tested against the earlier more comprehensive model (i.e. that which used Equation 7) for a whole range of signals and at two CIC levels. The comparison is shown in Figure 11. As can be seen, there is excellent agreement, justifying our simplifications.

4.7 SNR predictions from model

The model was used to evaluate the photon-counting SNR (SNR_{pc}) over a range of CIC and signal levels. The results are shown in Figure 12, expressed both as a fraction of the SNR of

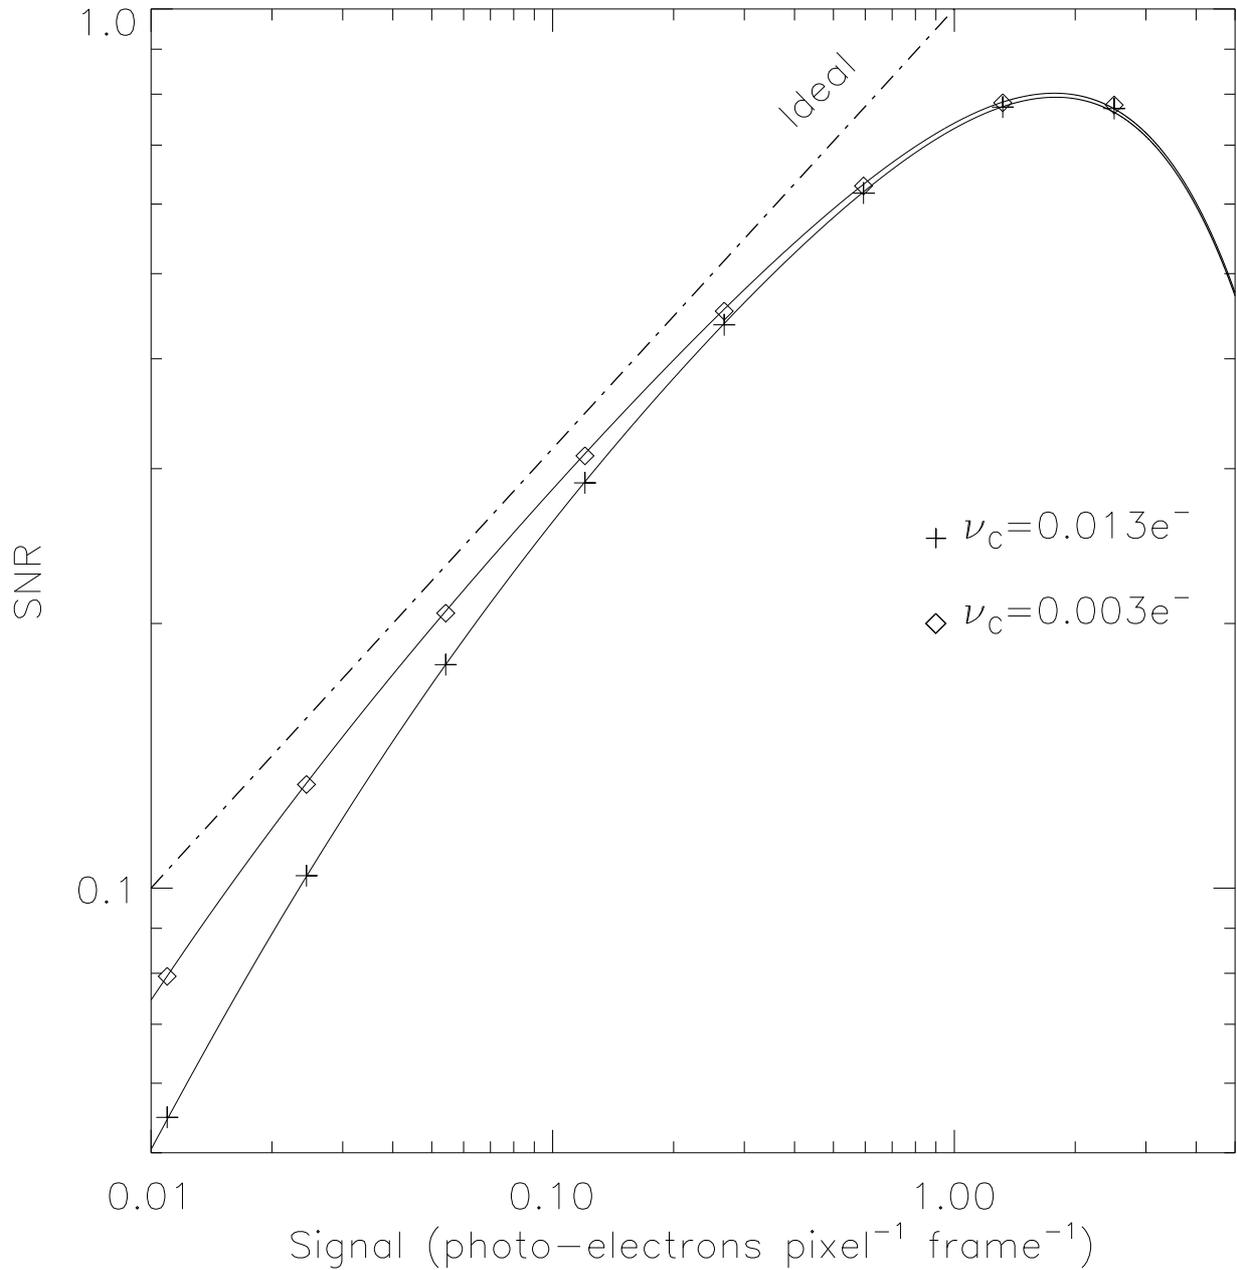


Figure 11: The SNR of two hypothetical EMCCDs in photon-counting mode, one with CICIR ν_C similar to that of the QUCAMs, another with ν_C equal to that which might be obtained in a more optimised detector. g_A in both cases is 2000 and the read noise is $0.025e_{pe}^-$. The crosses and stars show the predictions of the Monte Carlo model (Equation 7), the solid lines show the approximation (Equation 6). The SNR of an ideal detector (one where the $SNR = \sqrt{\text{Signal}}$) is plotted as a dot-dash line for comparison.

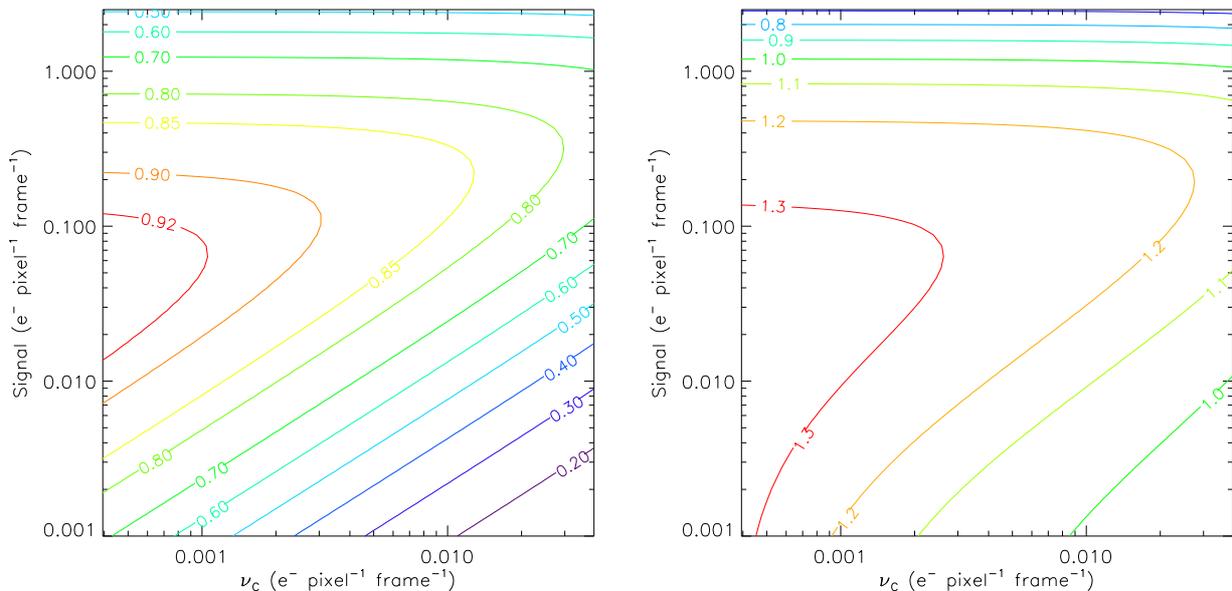


Figure 12: Relative SNR contours of a photon-counting EMCCD similar to the QUCAMs, compared in the left panel to an ideal detector and in the right panel to a linear-mode EMCCD. It has been assumed that all the CIC is generated in the EM register, the read noise $\sigma_{EM}=0.025e_{pe}^-$, the threshold is fixed at $0.1e_{pe}^-$ and the EM gain, $g_A=2000$.

an ideal detector SNR_{ideal} and the SNR of an EMCCD operated in linear mode SNR_{lin} (as described in Equation 4). The ideal detector is assumed to have the same QE as the EMCCD but does not suffer from read noise, coincidence or threshold losses. The figures reveal the presence of a photon counting ‘sweet-spot’, where an SNR in excess of 90% of ideal is possible if the CIC can be sufficiently reduced (to around $\nu_C = 0.002 e^- \text{ pix}^{-1}$). The sweet-spot is quite narrow and extends between signals of ≈ 0.07 and $0.2 e^- \text{ pix}^{-1}$. The plots also show that for signals of less than $1.2e^- \text{ pix}^{-1}$, photon counting is superior to linear-mode operation, and this is fairly independent of CIC level.

Figure 12 shows that the sweet-spot of an EMCCD actually covers a very small range of exposure levels, however, we can effectively slide the sweet-spot along the per-temporal-bin exposure scale to quite high signal levels through the use of blocking, i.e. summing together a number of frames whose total exposure time equals our required temporal resolution. For the QUCAMs, if we use the fairly generous definition of the sweet-spot as occupying the exposure range over which $\text{SNR}_{pc} > \text{SNR}_{lin}$ then this dynamic range is about 30:1. This would be like using a normal science CCD camera with a 7-bit (16-bit being more usual) analogue to digital converter (ADC) and could cause problems if the spectrum we wish to observe has a set of line intensities that exceeds this range.

5 Recipe for using QUCAM2 and QUCAM3 on ISIS

QUCAM2 and QUCAM3 are relatively slow cameras giving a minimum frame time in EM mode of 1.6s. This then dictates the highest temporal resolution that is available. The small size of the CCD means that when used on the ISIS spectrograph it measures only 3.3 arc-minutes in the spatial direction. Full frame readout is then generally needed in order to locate a suitable comparison star for slit-loss correction and this frame rate will therefore be hard to improve

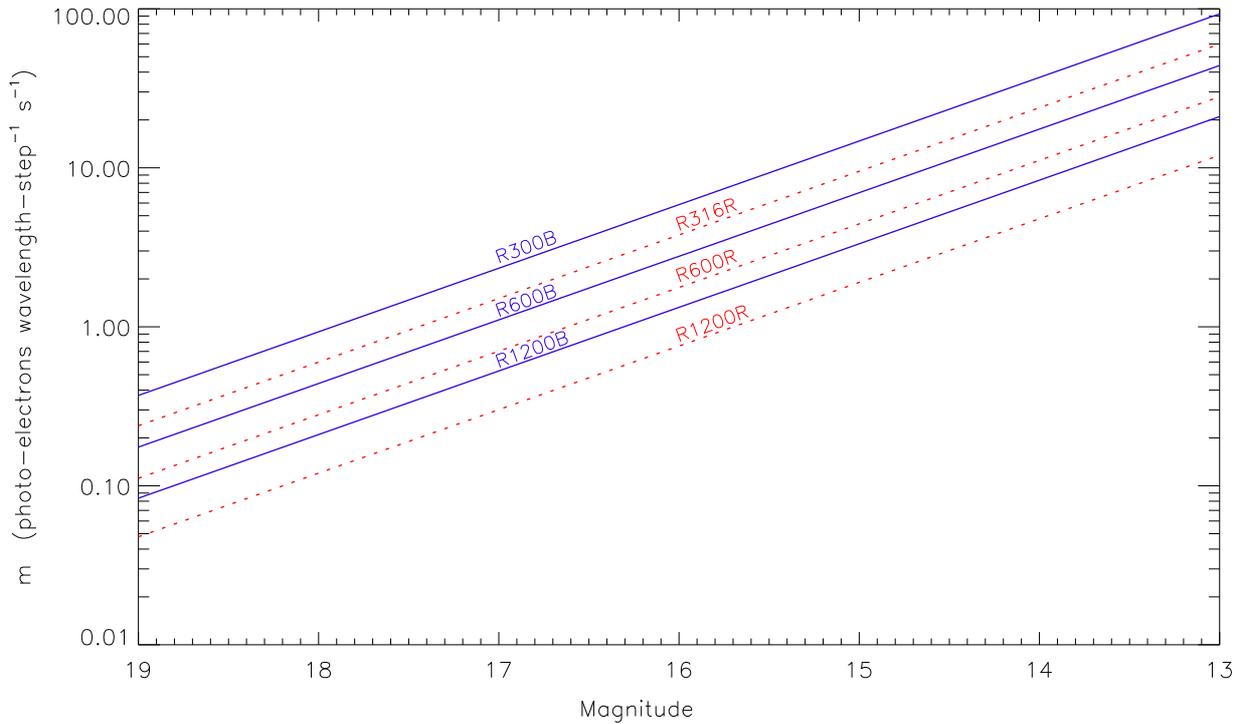


Figure 13: The flux m , that can be expected from each grating on ISIS when used with the QUCAMs as a function of source magnitude. R -band magnitudes as shown for the R -gratings and B -band magnitudes for the B -gratings. Calculated for airmass=1.

upon through the use of windowing. The linear EM mode should be considered as the default mode about which the observations are planned. The reason for this is that it gives an SNR that is a constant fraction of an ideal detector for almost all signal levels (see Figure 7). The observer will not go far wrong by selecting this mode. It may be possible to coax extra SNR (as much as 40%) from the observations by switching to one of the other modes but if the observation is not prepared carefully the data could prove useless. With linear mode the observer is guaranteed a practically noise-free detector without the dangers of potential coincidence losses and worse SNR than with a conventional CCD.

The first stage in planning the EMCCD observation is to refer to Figure 13. This allows us to calculate m , the flux per pixel step in wavelength that we can expect from the object at our chosen spectral resolution. Using this datum we then need to refer to Figure 14 to see how many seconds of observation ($T_{\text{SNR}=1}$) would be required on our object, using an EM detector in linear mode, to reach an SNR=1 in the final extracted spectrum. Figure 14 shows the calculated times for both dark and bright-sky conditions (new and full moon) for both arms of ISIS. Seeing of $0.7''$ and a slit width of $1''$ has been assumed. Given that the spatial plate-scale of the QUCAMs on ISIS is $0.2'' \text{ pix}^{-1}$, this implies that each wavelength element in the final extracted spectrum contains ~ 5 pixels-worth of sky (assuming the spectrum is extracted across $\sim 1.5 \times \text{FWHM}$ pixels in the spatial direction). Once we know the time it takes to reach an SNR=1, it is then straightforward to calculate the time needed to reach any arbitrary SNR, since with linear mode, the SNR is proportional to the square-root of the observation time.

The observer can now either play safe and use linear mode or explore the possibility of up to 40% higher performance from either PC or conventional modes. This will depend on the mean per-pixel signal (from all sources including the sky) that we can expect during an exposure of

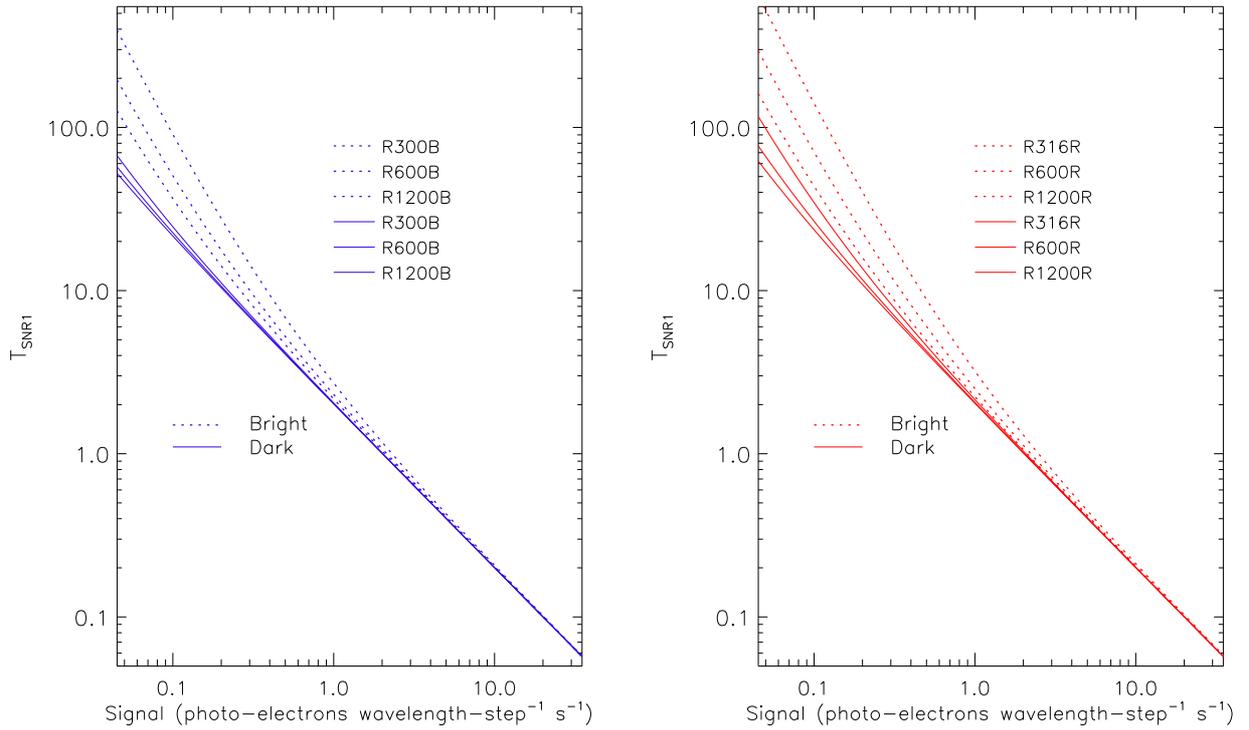


Figure 14: T_{SNR1} , the observation time on the WHT required to reach an SNR of 1 in the final extracted spectrum with the ISIS blue (left panel) and red (right panel) gratings as a function of source brightness. The detector is QUCAM2 operated in linear mode. A slit-width of $1''$, seeing of $0.7''$ and a spectral extraction over 5 pixels in the spatial direction are assumed. Observations are in the B-band (left) and R -band (right).

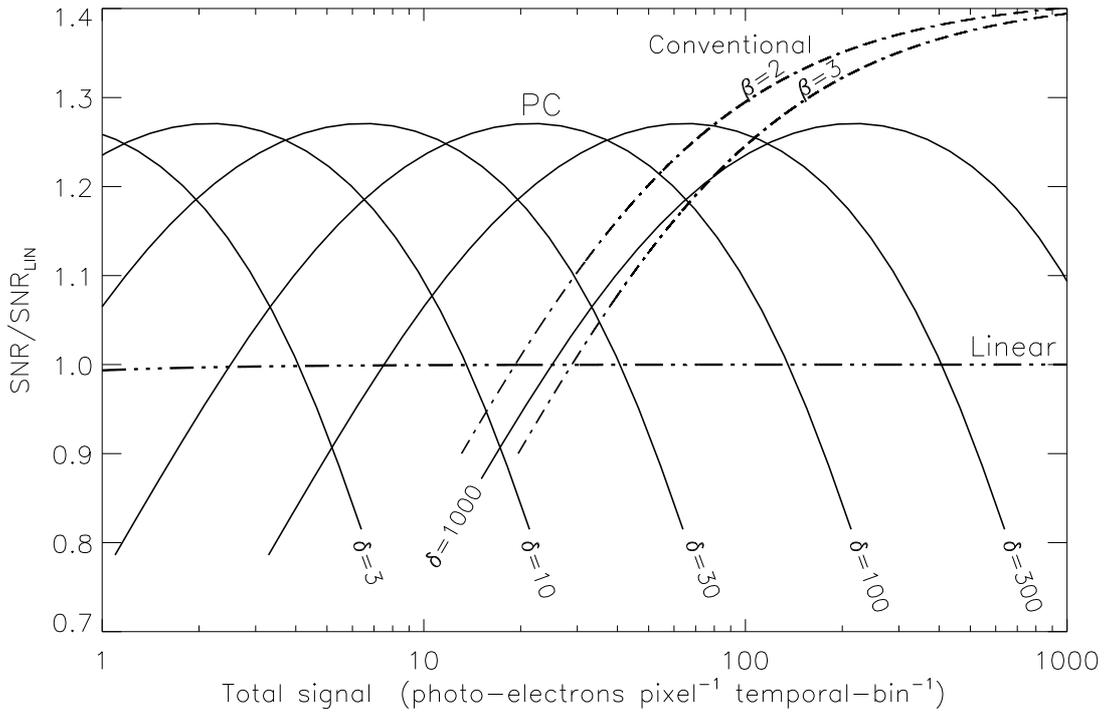


Figure 15: The relative SNR achievable using each mode as a function of the total (i.e. source+sky) per-pixel signal. β is the off-chip binning factor. δ is the photon counting blocking factor (the number of thresholded PC frames that are summed within each temporal bin). The read noise and CIC levels experienced with QUCAM2 are assumed.

duration equal to our required temporal resolution, τ . This is calculated as follows:

$$\text{signal pix}^{-1} = \tau \cdot \left[\text{sky} + \frac{m}{d} \right], \quad (8)$$

where d is the seeing-induced FWHM of the spectrum along the spatial axis of the CCD frame, measured in pixels. This dictates what spatial-binning factor we later need to use when extracting the spectrum. Note that in conventional and linear mode, τ is defined by the individual frame time, whereas in PC mode we divide τ into δ separate frames that are later photon counted and summed (in order to observe brighter objects without incurring coincidence losses). Now that we have an estimation of the per-pixel signal we can use Figure 15 to find which mode will give the optimum SNR on a per-pixel basis. If we find we can use photon counting then, according to Figure 15, the SNR gain will be $\sim 25\%$, which we can translate into less telescope time. The SNR increases as the square root of the observation time in PC and linear mode, since the detector is dominated by noise sources with variances that increase linearly with signal. If instead we find that we have enough photo-electrons to permit conventional mode, the reduction in the observation time is harder to estimate since the relatively high read noise gives a non-linear relation between the square root of the exposure time and the SNR. The saving in telescope time from use of conventional mode could, however, be as much as 50% (relative to linear mode) as the per-pixel signal tends to higher and higher values and the read noise becomes insignificant relative to the Poissonian noise in the sky and target.

5.1 Worked example of the recipe

We wish to observe an $R=18.5$ eclipsing binary star with the red arm of ISIS at the highest possible spectral resolution under dark-sky conditions. The emission line is $4\times$ brighter than the underlying continuum. The star undergoes an eclipse that lasts 7 minutes that we wish to resolve spectroscopically. What SNR can we achieve? Which mode should we use?

To begin this problem, we use Figure 13 to establish the signal that we can expect from the emission line. If we choose the R1200R grating (resolution ~ 7000) we will receive $0.07 \times 4 = 0.28$ photo-electrons per wavelength step per second. If we observe the object with a temporal resolution of 30s we will be able to easily resolve the eclipse. Referring to Figure 14 we can then see that this combination of signal and spectral resolution will require ~ 7 s of observation to give an SNR=1 in the final extracted spectrum if we use linear mode. Since we can actually observe for 30s, the SNR will be equal to $\sqrt{30/7} = 2.1$.

We now turn to the choice of observing mode. Assuming $0.7''$ seeing, a slit-width of $1''$ and that the spectrum will be extracted over 5 pixels in the spatial direction, we can calculate that the peak signal per temporal bin per pixel will be $0.28 \times 30/5 = 1.7 e^-$. To this we must add the sky signal tabulated in Table 1. Since we are observing in dark time and at high resolution this will be a negligible $0.003e^-s^{-1}$. Referring to Figure 15 we can then immediately rule out conventional mode as the SNR would collapse at such a low signal level. The best mode would then be PC with δ somewhere between 3 and 10. Interpolating the figure we can estimate that $\delta \approx 7$ would be optimum. This is feasible since the minimum read-time for 7 PC frames is 11.2s with the QUCAMs, well below the required temporal resolution of 30s. In conclusion, we would obtain the best SNR by observing at a frame time of 4.3s, photon counting the raw images and then averaging them into groups of 7 to obtain our required time resolution. Tuning the exposure time in this way to keep the spectral line on the PC sweet-spot affords us a $\sim 25\%$ SNR improvement over linear-mode operation (see Figure 15).

Note that there is an upper-limit to the useful PC blocking factor δ , set by the read noise of the conventional amplifier. Large blocking factors imply low-temporal resolution and there comes a point where it becomes favorable to use the conventional mode (see Figure 15). The degree of off-chip binning β used is critical since this adds noise to conventional mode but not to photon counting mode.

For the QUCAMs it can be demonstrated (using the SNR equations in Section 3) that the maximum useful blocking factor is $\sim 20\sigma_N^2/\beta$. For larger values of δ the SNR that can be obtained with conventional mode then exceeds the maximum obtainable with PC operation.

5.2 Binning

CCDs can be binned on-chip in a noiseless fashion in both axes. Since a spectrum will always be spread by seeing in the spatial direction, some degree of binning is often required. Some of this can be done on-chip but it is usual to do some of it off-chip (i.e. post-readout) during extraction of the spectrum. An EMCCD will suffer much less from off-chip binning than a conventional detector, which has an effective read noise multiplied by the square-root of the binning factor. As the off-chip binning factor β increases, the balance tips ever more in favour of the EMCCD.

5.3 Phase folding

For observations of objects that vary on a regular period, such as short-period binary stars, we can also consider extending our observations over many orbits and then ‘phase folding’ the data. This is an equivalent form of off-chip binning. It increases our signal by the folding factor F and permits us to observe fainter sources or, alternatively, to observe the same source at higher time resolution. If the SNR we require per wavelength element in our final extracted spectrum is given by SNR_{req} then the number of phase folds F we need to use (i.e. the number of orbits we must observe) in linear mode is given by:

$$F \approx \frac{T_{\text{SNR1}} \times \text{SNR}_{req}^2}{\tau}. \quad (9)$$

References

- Basden A. G., Haniff C. A., Mackay C. D., 2003, *MNRAS*, 345, 985
- Daigle O., Carignan C., Blais-Ouellette S., 2006, in Dorn D. A., Holland A. D., eds, *SPIE Conference Series Vol. 6276*. p. 42
- Dhillon V. S., Marsh T. R., Stevenson M. J., Atkinson D. C., Kerry P., Peacocke P. T., Vick A. J. A., Beard S. M., Ives D. J., Lunney D. W., McLay S. A., Tierney C. J., Kelly J., Littlefair S. P., Nicholson R., Pashley R., Harlaftis E. T., O’Brien K., 2007, *MNRAS*, 378, 825
- Hollenhorst J. N., 1990, *IEEE Transactions on Electron Devices*, 37, 781
- Ives D., Bezawada N., Dhillon V. S., Marsh T. R., 2008, in Dorn D. A., Holland A. D., eds, *SPIE Conference Series Vol. 7021*. p. 10
- Janesick J. R., 2001, *Scientific Charge-Coupled Devices*. SPIE Press, Bellingham
- Mackay C. D., Tubbs R. N., Bell R., Burt D. J., Jerram P., Moody I., 2001, in Blouke M. M., Canosa J., Sampat N., eds, *SPIE Conference Series Vol. 4306*. p. 289
- Marsh T. R., 2008, in Phelan D., Ryan O., Shearer A., eds, *American Institute of Physics Conference Series Vol. 351*. p. 75
- Plakhotnik T., Chennu A., Zvyagin A. V., 2006, *IEEE Transactions on Electron Devices*, 53, 618
- Tubbs R. N., 2003, PhD thesis, University of Cambridge
- Tulloch S. M., 2004, in Moorwood A. F. M., Iye M., eds, *SPIE Conference Series Vol. 5492*. p. 604
- Tulloch S. M., 2010, PhD thesis, University of Sheffield

A Symbols used

e_{pe}^- , input referenced photo-electrons.
 ν_C , mean charge in a pixel from CICIR.
 B_C , mean number of CICIR events in a pixel.
 g_A , avalanche, multiplication or EM gain (unitless).
 g_S , gain or system gain (e^-/ADU).
 g_{S0} , system gain of EM amplifier with EM gain=1.
 p_C , per transfer CIC generation probability.
 S , number of stages in the EM register.
 σ_N , conventional output read noise.
 σ_{EM} , EM output read noise.
 SNR_{pc} , Signal to noise ratio in PC mode.
 SNR_{lin} , SNR in linear mode.
 SNR_{ideal} , SNR of a noise free detector.
 SNR_{req} , required SNR in extracted spectrum.
 τ , required time resolution of observation.
 T_{SNR1} , exposure time to reach $SNR_{lin} = 1$ in extracted spectrum.
 M , mean signal per pixel.
 K , sky signal within one temporal bin.
 D , dark charge within one temporal bin.
 m , signal per-pixel-wavelength-step s^{-1} .
 t , photon-counting threshold (units e_{pe}^-).
 n , mean counts in PC mode (counts pix^{-1}).
 N , mean illumination (photo-electrons pix^{-1}).
 d , spatial extent of spectrum (due to seeing), pixels.
 δ , photon-counting blocking factor.
 β , off-chip binning factor.
 F , phase-folding factor.

B Fractional charge of CICIR

A CIC electron originating within the EM register will effectively have a fractional charge whose value is equal to the average charge \bar{q}_O of a single photo-electron charge packet during its transit through the EM register. The instantaneous charge q_O of a pixel within the EM register that originated as a single photo-electron at the register input is:

$$q_O = (1 + p)^x, \quad (10)$$

where x is the position within EM register and p the per-transfer multiplication probability. The mean value, \bar{q}_O , of this pixel 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 S . If we reference this charge to an equivalent signal at the input to the register, such that $q_I = q_O/g_A$, we get:

$$\bar{q}_I = \frac{1}{Sg_A} \int_{x=1}^S (1 + p)^x dx. \quad (11)$$

This integral by the standard solution:

$$\bar{q}_I = \frac{1}{Sg_A \ln(1+p)} [(1+p)^x]_{x=1}^S, \quad (12)$$

which gives the result

$$\bar{q}_I = \frac{(1 - g_A^{-1})}{\ln g_A} \approx \frac{1}{\ln g_A}. \quad (13)$$

Equation 13 now allows us to calculate the per transfer probability of CIC p_C in the EM register from a knowledge of the gain g_A and the mean CIC charge in the bias ν_C . The generation of this charge is a binomial process so B_C the total number of CIC events in a pixel exiting the EM register is given by

$$B_C = p_C S. \quad (14)$$

Since each of these CIC events will contain an average charge of \bar{q}_I , the mean per pixel CIC charge ν_C is given by

$$\nu_C = \bar{q}_I B_C, \quad (15)$$

substituting \bar{q}_I from Equation 14 and rearranging, we get

$$p_C \approx \frac{\ln(g_A)\nu_C}{S}. \quad (16)$$

We can then show how the mean CIC charge in a pixel ν_C (relevant for linear-mode operation) and the mean number of CIC events in a pixel B_C (relevant for PC operation) are related as follows:

$$\frac{B_C}{\nu_C} \approx \ln(g_A). \quad (17)$$

C SNR of an ideal photon counter

We present a full derivation of SNR_{pc} , the SNR in a photon-counting detector. This differs from that of an ideal detector due to the effects of coincidence losses.

Let N be the mean illumination in photo-electrons per pixel and n the mean photon-counted signal per pixel, i.e. the fraction of pixels that contain one or more photo-electrons. Poisson statistics tells us that

$$n = 1 - e^{-N}, \quad (18)$$

therefore:

$$N = -\ln(1 - n). \quad (19)$$

The noise in a photon counted frame can be derived straightforwardly by considering that only two pixel values are possible: 0 and 1. Pixels containing 0 will have a variance of N , those containing 1 will have a variance of $N - 1$. Knowing the fraction of pixels containing each of these two values then allows us to combine the variances in quadrature to yield σ_{pc} , the rms noise:

$$\sigma_{pc} = \sqrt{[e^{-N}N^2 + (1 - e^{-N})(N - 1)^2]}, \quad (20)$$

$$= \sqrt{(e^{-N} - e^{-2N})}. \quad (21)$$

The photon-counted images must then be processed to remove the effects of coincidence losses. This is done after the component frames within each temporal bin have been averaged to yield a mean value for n for each pixel. The original mean signal N prior to coincidence losses is

then recovered by using Equation 19. Although coincidence loss tends to produce a saturation and a smoothing of the image structure, the overall effect is to add a great deal of noise to the observation and for this reason we must avoid a photon-counting detector entering the coincidence loss regime. The amount of extra noise generated can be calculated by considering the change dN in N produced by a small change dn in n . The noise in the final coincidence-corrected pixel will then be equal to that in the unprocessed average pixel multiplied by dN/dn . From Equation 19 we get

$$N + dN = -\ln[1 - (n + dn)]. \quad (22)$$

This standard differential is then solved to yield

$$\frac{dN}{dn} = (1 - n)^{-1}. \quad (23)$$

Substituting N for n using Equation 18 we get

$$\frac{dN}{dn} = e^N. \quad (24)$$

We then multiply the uncorrected noise given in Equation 21 by this factor to yield the noise in the final coincidence-loss-corrected PC image. SNR_{pc} is then given by:

$$\text{SNR}_{pc} = \frac{N}{\sqrt{e^N - 1}}. \quad (25)$$

This can be expressed in units of n , using Equation 19, which is more useful since it is n that we actually measure from our images:

$$\text{SNR}_{pc} = \frac{-\ln(1 - n)}{\sqrt{(1 - n)^{-1} - 1}}. \quad (26)$$

D SNR of a photon-counting EMCCD

We have already shown in Appendix C that the SNR of an ideal photon counter is:

$$\text{SNR}_{pc} = \frac{N}{\sqrt{e^N - 1}}, \quad (27)$$

with N representing the signal per pixel. This basic equation is now altered to a more realistic form to include the noise sources found in an EMCCD. Certain approximations are made during this process. The validity of these approximations have been verified by the Monte Carlo modelling.

So to begin with, we replace N in the numerator with the detected fraction of photo-electrons at our chosen threshold and we replace N in the denominator with the detected signal plus the detected CIC. In Figure 9 we have already shown that a threshold of $0.1e_{pe}^-$ is close to optimum and that at this level 90% of photo-electrons and 23% of CICIR will be detected. We then get:

$$\text{SNR}_{pc} = \frac{0.9M}{\sqrt{\exp(0.9M + 0.23B_C) - 1}}, \quad (28)$$

where M is the signal per temporal bin and B_C the number of CICIR events per pixel. We now need to consider that any photon counting observation will require the blocking of δ separate

images if we are to achieve a usable SNR (the SNR of a single PC image with the exposure level lying within the sweet-spot is ≈ 0.4 , see Figure 11). We then get:

$$\text{SNR}_{pc} = \frac{0.9M}{\sqrt{\delta} \sqrt{\exp(0.9M/\delta + 0.23B_C) - 1}}. \quad (29)$$

Next we need to include other noise sources such as sky K and dark charge D (units of $e^- \text{pix}^{-1}$). Since these are indistinguishable from photo-electrons they will have the same detected fraction for a given PC threshold. We also need to express the CICIR in terms of ν_C (see Equation 17), i.e. the mean per-pixel charge that CICIR contributes to the image. This is a parameter that can be measured directly from the bias frames of an EMCCD camera. So we get:

$$\text{SNR}_{pc} = \frac{0.9M}{\sqrt{\delta} \sqrt{\exp [(0.9(M + D + K)/\delta) + 0.23 \ln(g_A)\nu_C] - 1}}. \quad (30)$$

Note that ν_C the CICIR is multiplied by a factor of $\ln(g_A)$ in the denominator of Equation 30. This is a consequence of the fractional charge of a CICIR electron (see Appendix B).

Note also that the read noise has been entirely ignored in this simplified description. This is fair since the threshold was set well above the noise, resulting in very few false counts. The read noise has a secondary influence, however, since it causes an effective blurring of the threshold level. This ‘fuzzy-threshold’ must add some noise to the images since it can make all the difference as to whether an event lying close to the threshold is counted or not. The close fit between the approximation and the comprehensive model (see Figure 11) would, however, indicate that this noise source is not significant.

E Sky backgrounds

Table 1: WHT+ISIS sky backgrounds with QUCAM2. Units: photo-electrons $\text{pix}^{-1}\text{s}^{-1}$. Slit width= 1". Obtained using the SIGNAL program from the ING web pages.

Grating	Bright	Dark
R316R	0.12	0.014
R600R	0.05	0.006
R1200R	0.023	0.003
R300B	0.07	0.004
R600B	0.03	0.002
R1200B	0.016	0.001