A comparison of the noise characteristics of EMCCDs and CMOS image sensors for astronomical lucky imaging

Harry Thomas Fox

A thesis presented for the degree of

Doctor of Philosophy

Centre for Electronic Imaging

School of Physical Sciences

The Open University

August 2022

Abstract

Over the last few decades, the improvements to CMOS image sensor technology have made it possible for them to be considered for high-precision astronomy applications. GravityCam, a new concept for a ground-based imaging instrument, is proposing an upcoming CMOS technology to achieve significantly higher image quality over a wide field of view compared to previous instruments using EMCCD cameras. Observing faint signals, such as lunar masses via gravitational lensing, requires the ability to measure extremely small changes in signal, and therefore controlling image sensor noise is extremely important. For this reason, investigations into clock-induced charge for EMCCDs and the readout noise in CMOS image sensors are completed in the thesis to see the impact of these noise sources on low signal observations. A simulation-based approach is taken to investigate how CMOS image sensor noise impacts the limiting magnitude of the instrument and what the expected star loss is for each given magnitude versus mean readout noise.

Declaration

I hereby declare that no part of this thesis has been previously submitted to this or any other university as part of any other degree or professional qualification. This thesis has been wholly written by the undersigned, except for colleagues and others acknowledged in the text.

Harry Fox

2022

Acknowledgements

Firstly, I would like to thank my supervisors, Jesper Skottfelt and Konstantin Stefanov, for all their hard work helping me throughout the PhD, teaching me and especially for the many, many hours helping writing this thesis over weekends and late at night.

A big thank you to Oliver Hetherington and Nathan Bush for the many hours helping me solve various issues during the experimental parts of this PhD.

Thank you to all of the remaining members of the CEI, current and past, for being excellent colleagues and friends, including James, Daniel-Dee, Alice, Tom, Daniel, Matthew, Dave, Chiaki, Julian, Steven, George, Ben, and especially the weekly DnD group, Zoe, Chris, Anton, Saad, Lawrence and Neil for keeping me sane and social during and after the lockdown period.

Thanks to my friends and family back home for the support and care throughout the course, especially during the covid lockdown.

Finally, the biggest thank you to my partner Ellie, for looking after me and for the tremendous love and support over the last four and half years of this project.

Table of Contents

Table of Tables 13

Table of Figures 15

Chapter 1: Introduction 21

Chapter 2: Charge-Coupled Device (CCD) 23

- 2.1: History of Charge-Coupled Devices 23
- 2.2: CCD Architecture 25
  - 2.2.1: Charge Generation 25
- 2.3: Charge Transfer Method 30
  - 2.3.1: Charge Transfer Efficiency 31
  - 2.3.2: Charge Transfer Mechanisms 32
  - 2.3.3: Traps 33
- 2.4: Charge Output Circuit 34
- 2.5: Amplifier Noise Sources 35
  - 2.5.1: White Noise 36
  - 2.5.2: Reset Noise 36
  - 2.5.3: Flicker Noise 37
  - 2.5.4: Shot Noise 38
- 2.6: CCD Noise Sources 39
  - 2.6.1: Photon Shot Noise 39
  - 2.6.2: Dark Current 40
  - 2.6.3: Clock-Induced Charge 40
- 2.7: Electron-Multiplying CCD (EMCCD) 41
  - 2.7.1: Impact Ionisation 42
  - 2.7.2: Multiplication Register 44
  - 2.7.3: Excess Noise Factor 45

2.7.4: Clock-Induced Charge in EMCCDs......48
2.7.5: EMCCD Ageing......51
Chapter 3: Complementary Metal-Oxide Semiconductor (CMOS) Image Sensors...54
3.1: Active Pixel Sensors (APS)......54
3.2: Correlated Double Sampling (CDS)......55
3.3: 4T Pixels......57
3.4: 5T Pixels......58
3.5: CIS Readout Modes......58
3.5.1: Rolling Shutter......58
3.6: CIS Characteristics......59
3.6.1: Random Telegraph Signal Noise......60
3.7: CCD vs CIS......61
3.7.1: Charge Transfer Inefficiency (CTI)......61
3.7.2: Radiation Hardness......62
3.8: Scientific CMOS (sCMOS)......63
Chapter 4: Ground-based Observations......66
4.1: Atmospheric Seeing......66
4.1.1: Measuring Seeing Conditions......67
4.1.2: Aberrations and Zernike Polynomials......69
4.2: Adaptive Optics......71
4.2.1: Problems with Adaptive Optics......72
4.3: Lucky Imaging......73
4.3.1: Fried Probability......73
4.3.2: Lucky Exposures......75
4.3.3: Lucky Imaging Issues......77
Chapter 5: GravityCam......81
5.1: Device Requirements......82

5.2: Science Cases 83

5.2.1: Microlensing Surveys 84
5.2.2: Occultation Surveys 85
5.2.3: Dark Matter Surveys 86

5.3: Gravitational Lensing 86

5.4: GravityCam Summary 87

Chapter 6: CCD Tests and Results 89

6.1: Spurious Charge Measurements 89

6.2: Spurious Charge vs Gain 95

6.3: Column and Row Noise 99

6.4: Potential Tail Impact Values 109

Chapter 7: CMOS Image Sensor Tests and Results 111

7.1: Introduction to the CIS115 Sensor 111

7.2: Why the CIS115 and the aims of this chapter 111

7.3: CIS115 Tests 112

7.3.1: Dark Current 112
7.3.2: Stability 115
7.3.3: Readout Noise 118
7.3.4: Noisy Pixel Locations 119
7.3.5: Noisy Pixel Characteristics 126

7.4: Discussion 129

Chapter 8: The effect of CIS-like noise on stellar detection limit 133

8.1: Atmospheric Seeing 133

8.2: CIS-like Noise 136

8.3: ODIN Reduction Software 141

8.3.1: DanDIA IDL Starfit 142
8.4: Danish Telescope Specific Case 143

8.5: Comparing Magnitudes 144

8.5.1: Sub-electron Mean Distributions 145
8.5.2: One Electron and above Mean Noise 149
8.6: Summarising Star loss vs Mean Noise 153

Chapter 9: Conclusion and Future Work 159

9.1: Future Work 161

References 162

Table of Tables

Table 1 - A comparison of two Teledyne e2v EMCCD devices, CCD 220 and CCD351, and two sCMOS devices, Fairchild CIS1910 and CIS2521 64

Table 2 - The low-order aberrations for the expansion of the Zernike polynomial sequence. 71

Table 3 - Comparison between Red and Visual camera. 94

Table 4 – Number of tail pixels as a function of EM gain. 97

Table 5 – High-value pixel threshold with multiplication gain. 99

Table 6 - Mean and standard deviations for Figure 7-4 and Figure 7-5. 117

Table 7 - The three correlation tests were conducted for the adjacent pixels. 128

Table of Figures

Figure 2-1 Shows the clocking mechanism in a three-electrode pixel. 23

Figure 2-2 - An example of how a combination of n-type and p-type material creates a depletion zone. 27

Figure 2-3 - The layout for CCD. 31

Figure 2-4 – A single-stage output for a CCD. 35

Figure 2-5 – Taken from (Castro, 2017), Spectral Noise Density versus Frequency. 38

Figure 2-6 - The layout of a LLLCCD device, as shown by (Jerram, 2001). 42

Figure 2-7 – A four electrode clocking scheme that is used in EMCCD to multiply the charge packet. 44

Figure 2-8 - Excess noise factor as a function of the multiplication gain. Taken from (Robbins, 2003). Note that the equation (8) in the Figure refers to equation (21) in this text. 47

Figure 2-9 – A comparison between 1000 frames taken on the CCD97 at Gain = 1 and Gain = 300. The X-axis shows the signal in DN, and the Y-axis is the number of pixels for a given output value. 50

Figure 3-1 - Basic three transistor (3T) CMOS pixel layout. 55

Figure 3-2 - A simple layout of a 4T CIS pixel. 57

Figure 3-3- 5T pixel schematic. 58

Figure 3-4 - A simple layout for how dual amplifiers and dual ADCs allow for both high and low gain operation. Based on Figure 4 in (Chang, 2018). 63

Figure 4-1 - A diagram, showing plane waves becoming perturbed due to the turbulent atmosphere. 66

Figure 4-2 - shows the relationship between $\boldsymbol{DR0}$ and the probability of getting a near-diffraction limited exposure. 75

Figure 6-1 - A histogram of the dark signal from 1000 frames at EM gain 1 (i.e., no EM gain) for CCD97 and a Gaussian fit. 90

Figure 6-2 - A histogram of the dark signal in a CCD97 at gain 300 for 1000 frames and a Gaussian fit. 91

Figure 6-3 - A cumulative plot shows that pixel values below 138 DN contain 90% (red line) of all pixels. 92

Figure 6-4- The signal distribution for the visual and the red cameras at a gain of 300, obtained from 1000 frames. 93

Figure 6-5 – Number of tail pixels per frame against increasing EM gain. 95

Figure 6-6 - A comparison between the red and visual cameras. This includes all the pixels within the tail. 98

Figure 6-7 - Average tail pixel counts per pixel, per frame versus column number, for the red camera at gain 300. Taken over the full 512x512 image area. This is averaged over 1000 frames. 100

Figure 6-8 - The bias pattern of the Danish telescope's CCD97 red camera, taken from (Harpsøe, 2012). 102

Figure 6-9 – Average tail pixel quantity per frame versus row number for the red camera at gain 1. 1000 frames are being investigated. This has been windowed to 450x450 to remove any contribution from edge effects. 103

Figure 6-10 – Similarly to Figure 6-9, the average amount of tail pixels per pixel in each row, per frame is plotted but instead of gain1, gain 300 is used. 104

Figure 6-11 – Average tail pixel counts per pixel, per frame versus column number, for the red camera at gain 1. 105

Figure 6-12 – Similarly to both Figure 6-10 and Figure 6-9, this figure effectively a gain300 version of Figure 6-11, to show the impact of increasing the device's gain. . 106

Figure 6-13 - The device architecture of the backside-illuminated CCD97, taken from (Teledyne e2v, 2021). 107

Figure 6-14 - Average tail pixel counts per pixel, per frame versus column number, for the visual camera at gain 300. 108

Figure 6-15- Right shows average high-value tail pixel counts per pixel, per frame versus column number for the red camera at gain 300, left shows the same for the visual camera. 109

Figure 7-1 - Dark signal with the Transfer Gate (TG) off and on. The dark signal has been removed by keeping the TG low. 113

Figure 7-2 - The pixel timings for a standard readout. 114

Figure 7-3 - Modified pixel timings for effectively removing dark current from the measurements. This timing is used for the readout noise measurements and Figure 7-1(a). 114

Figure 7-4 - Mean, DSNU-corrected signal from a stack of 1000 images (top); spread of the standard deviations across the stack (bottom). 116

Figure 7-5 - The same plot as Figure 7-4, except the dark current, has been removed with the previously mentioned clocking modification. 116

Figure 7-6 - Readout noise histogram of CIS115. 118

Figure 7-7 - Cumulative noise distribution in CIS115. 119

Figure 7-8 shows the search area around a noisy pixel in the centre (in black) and 4 noisy pixels. 120

Figure 7-9 – A cumulative bar plot shows how many additional noisy pixels appear within a certain radius of a noisy central pixel. The y-axis is plotted as the percentage of noisy pixels, which comprise approximately 1% of the entire device, as defined by Figure 7-7. 121

Figure 7-10 - Top plot shows the measured data's noise distribution, as shown in Figure 6. The bottom plot shows the simulated array distribution based on the actual data. .. 123

Figure 7-11 – A box plot shows a comparison of the experimental and simulated datasets for quantity of detections within a three-pixel radius. 124

Figure 7-12 - Shows three pixels that are adjacent to each other, positioned at (3,105), (3,106) and (3,107)........ 126

Figure 7-13 – Distribution of the mean values of pixels, with standard deviation about the peak of Figure 7-6, per frame. 130

Figure 8-1 - An example of a point-like object becoming blurred by the telescope simulator. 134

Figure 8-2 – Globular Cluster NGC6528, high-resolution image taken using the Danish Telescope. 136

Figure 8-3 – Simulated CIS-like noise, pixel-quantity versus pixel value (e^-). 138

Figure 8-4 – Similar to but with a PDF multiplier of 5x, showing a notably longer high value pixel tail........ 139

Figure 8-5 – Simulated best resolution image of NGC6528 with a mean CIS-like read signal of 0.01 e^- and standard deviation of 0.2 e^- rms. 140

Figure 8-6 - Apparent magnitude versus quantity of stars. 145

Figure 8-7 – Two CIS-like noise distributions with a mean value of 0.01 e-. The right plot has 3x the pixel readout noise........ 146

Figure 8-8 – Similar to Figure 8-7 with both distributions having a mean value of 0.1 e- instead. 146

Figure 8-9 – Same format as previous with a mean value of 0.5 e-. 146

Figure 8-10- Group of magnitude plots with different mean noises in electrons......... 148

Figure 8-11 – Same pixel value distribution format as Figure 8-7 with a noise distribution with mean 1 e- and std 5.9 e- rms (left) and 9.0 e- rms (right) 150

Figure 8-12 - A noise distribution with mean 5 e- and std 8.4 e- rms (left) and 19.9 e-rms (right). 150

Figure 8-13 - A noise distribution with mean 10 e- and std 8.4 e- rms (left) and 26.8 e-rms (right).......... 150

Figure 8-14 - Magnitude plots for mean distribution values 1e- to 10e-. 152

Figure 8-15 - Star loss versus magnitude for distributions with means ranging from 0.01 e- to 1 -e. 154

Figure 8-16 - Star loss versus magnitude for distributions with means ranging from 1 e- to 10 -e.......... 154

Figure 8-17 - Difference between two distributions at mean five e-, one with an increased 3x standard readout noise.......... 156

Figure 8-18- Difference between two distributions at mean ten e-, one with an increased 3x standard readout noise.......... 156

Chapter 1: Introduction

In recent years, the development of Complementary Metal-Oxide Semiconductor (CMOS) Image Sensors (CIS) has reached a point where the technology is being considered a serious alternative to Charge-Coupled Devices (CCDs) in the field of observational astronomy. CCDs have long been the primary choice for astronomical instruments owing to their high sensitivity to a wide range of wavelengths, Teledyne e2v's CCD97 datasheet states the range of sensitivity as 193nm to 1000nm (Teledyne e2v, 2021). Other beneficial CCD characteristics are near 100% quantum efficiency and low device noise.

The CIS technology has been improving in these areas and has reached a point where they can be considered a viable alternative to, and maybe even surpass the performance of, CCDs. The current advantages of CIS devices are; higher framerates, radiation hardness, and low power consumption.

One of the key differences between the two technologies is the method of charge readout. CCDs transfer charge through the device's image and storage areas, passing charge packets through a series of electrodes before reaching the readout circuitry. This transfer approach means that all charge packets are read out through the same circuitry and therefore has the same readout noise applied to it. In comparison, CIS devices do not transfer charge across the device; each pixel has its own readout circuit within the pixel's boundaries. As a result, each pixel in a CIS device has its own readout noise distribution. This multiple readout mechanism significantly impacts low-signal observations, where controlling the contribution from noise sources is critical.

A variant of the CCD design, the Electron-Multiplying CCD (EMCCD), introduces an additional step in the signal readout phase of the device. This new step involves multiplying or gaining the charge packet using Avalanche Multiplication before readout, thus effectively making the readout noise negligible. Whilst this is advantageous for high framerate imaging, there is a drawback in that noise sources prior to the readout circuit are also multiplied. One notable source is Clock-Induced Charge (CIC), also called Spurious Charge, which contributes to charge packets in varying amounts, from small to large quantities of additional charge carriers. However, for a non-multiplying CCD, this contribution is negligible compared to readout noise, but for EMCCDs the multiplication register elevates CIC to notable levels.

This thesis aims to explore the impact of EMCCD CIC on high-speed low-signal observations, primarily frames used for high framerate imaging, and similarly, the impact of CIS pixel readout on high framerate Imaging, and whether there is any spatial dependency of CMOS noisy pixels. Finally, the relationship between different of CIS-like noise the quantity of stars observed is investigated via a simulated approach. This analysis on EMCCD and CIS noise will be used to help inform the choice of detectors for a proposed ground-based instrument called GravityCam.

Chapter 2: Charge-Coupled Device (CCD)

2.1: History of Charge-Coupled Devices

The Charge Coupled Device (CCD) concept was first invented in 1969 by Willard Boyle and George Smith (Smith W. S., 1970) as a novel approach to storing information electronically. The base idea is that the device can store charge in potential wells at the surface of the device and then transfer the charge across the device. This transfer method was initially demonstrated using a three-phase voltage-driven electrodes (also called gates) to create potential wells. This method is shown in Figure 2-1.

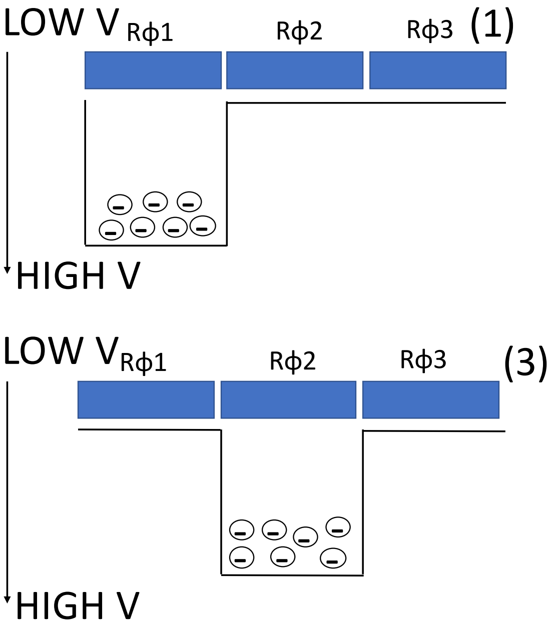

Figure 2-1 Shows the clocking mechanism in a three-electrode pixel.

Step 1 shows the initial phase, with the charge held under the first electrode, RΦ1. Next, the potential well is expanded when the second electrode, RΦ2, is set at the same voltage as RΦ1; this causes the charge to occupy this area, and then finally, the voltage of the RΦ 1 is brought back to the off-level, and all the charge is stored within the potential well of RΦ2. These steps are repeated until the charge is moved to the target location. This three-phase approach would form the basis of charge transfer still used in modern devices. The first devices had numerous issues, such as poor charge transfer efficiency (CTE) due to surface level charge transfer losses, much of the charge would be lost as the packet moved from well to well, and these issues were solved during the subsequent developments, such as the innovation of buried channel CCDs, which make of a doped silicon layer to create a potential well below the surface level to reduce surface level charge transfer losses.

It was discovered that combining the photosensitive properties of certain materials, for example silicon, the charge storage and transfer method of the CCD could be used to capture electrons generated by the photoelectric effect to create an electrically stored image of the incident photon source. This led to the development of CCDs as an imaging instrument, both for commercial and scientific purposes.

The main competitor of the imaging CCD was a type of photographic film (Janesick J. , 2001). The film had been used for an extended period and was a well-established and comfortable technology that astronomers at the time preferred. However, NASA needed an electronic solid-state detector device for the upcoming Large Space Telescope mission, later renamed the Hubble Space Telescope (HST). This requirement prompted much development for the CCD imaging sensor technology.

While the HST mission had candidates from other imaging technologies, both photographic film and image tube detectors, these had issues that stalled the HST team from committing to either of these technologies. In 1972 the Jet Propulsion Laboratory in California launched a concentrated effort to get the CCD to a point where it would be acceptable for space-based astronomy missions. One of the companies involved in this effort, Fairchild Semiconductor, created a commercially available buried channel, 100x100 pixel CCD in 1974. The Fairchild device showed greater CTE and reduced noise. More innovations would be made over the next few years, adding new device properties such as back-side illumination to increase quantum efficiency and polysilicon gates, which were much more resistant to shorting. As CCDs improved, many cameras were developed and used across observatories worldwide. These cameras produced incredible results that impressed the astronomical community, and the demand for this technology soon rose (Janesick J. , 2001). Eventually, the HST team used CCDs as part of the Wide-Field Planetary Camera. CCD proved to have better quality exposures over more extended observational periods, better sensitivity in visible and infra-red wavelengths, and much better linearity over imaging tubes and photographic films.

The success and discoveries made by the HST showed that the CCD technology was highly worthwhile investing in and continues to be today, with space missions such as Cassini (Lebreton, 1992), Kepler (Koch, 2010), and GAIA (Prusti, 2016) having utilised the CCD technology successfully.

2.2: CCD Architecture

The first thing to explore in CCDs is the structure of the semiconductors and how photons become converted into charge carriers.

2.2.1: Charge Generation

Two types of charge carriers are encountered within CCDs; firstly, there are electrons. These are negatively charged sub-atomic particles and are the type of charge carriers most referred to when discussing CCDs and charge manipulation. The second type of charge carrier is holes. These are not physical particles and are simplified to refer to the absence of an electron in atomic states. Holes are treated as a positively charged virtual analogue to electrons. CCDs are typically constructed from silicon because of silicon’s excellent purity and the high-quality native oxide that can be formed on its surface. Incident photons can free electrons via the internal photoelectric effect only if they have high enough energy. This energy requirement is called a bandgap, the difference between the valence and the conduction bands in a material. This property allows for greater control over the state of charges present within the material than in fully conductive or insulating materials; Silicon has a bandgap of approximately 1.12 eV at 300 K (Sze, 1969).

Additionally, adding or doping the silicon material with other elements that either add additional charge states near the conduction band or near the valance band results in a type of silicon called N-type or P-type, respectively. P-type silicon contains an excess of hole charge carriers (an analogue for a positive electron charge carrier), whereas N-type silicon contains an excess of electrons. Some of the materials used for doping are boron and gallium for P-type silicon and phosphorus for N-type silicon. This process works by added materials that either form holes in the valance band of the silicon lattice, creating a net positive charge (P-type) or introduce additional valance electrons, creating a net negative charge across the lattice (N-type). The introduction of these materials allows for a particular interaction between P-type and N-type silicon to create the P-N junction. The original theory behind the P-N junction can be found in (Shockley, 1949).

A P-N junction is created by doping one side of the semiconductor to p-type silicon and then the other side to n-type silicon. As the free charges diffuse across the border, a fixed charge is created in this region, called the space charge region. The fixed ions create an electric field that prevents further diffusion of free charge carriers. This effect results in a potential difference, diffusion potential, calculated in (1).

$$V_{\mathrm{diff}} = \frac{kT}{q} \ln\left(\frac{N_A\, N_D}{n_i^2}\right). \tag{1}$$

Here V_{diff} is the diffusion potential, q is the electron charge, T is the ambient temperature, k is the Boltzmann constant, N_A is the concentration of acceptor ions, N_D is the concentration of donor ions and n_i is the intrinsic carrier concentration. This potential repels any other free charge carriers, effectively creating a region of space in the material with none; this region is called the depletion region or the space charge region, as seen in Figure 2-2.

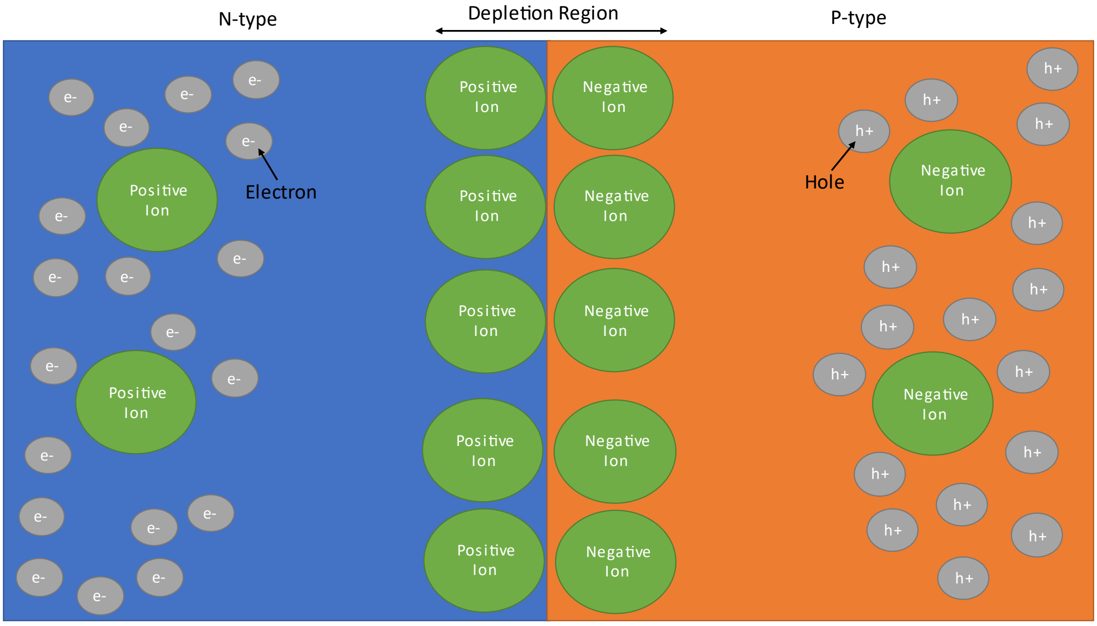

Figure 2-2 - An example of how a combination of n-type and p-type material creates a depletion zone.

This depletion zone's size depends on the material, the concentration of both ion types, and the diffusion potential. It is shown in (2) (Sze, 1969).

$$d = \sqrt{\frac{2\varepsilon_m \varepsilon_0 V_{\text{diff}}}{q}\left(\frac{N_A + N_D}{N_A\, N_D}\right)} \tag{2}$$

Here ε_m is the relative permittivity of the material, ε_0 is the permittivity of free space, and d is the width of the depletion zone. Note here that the width of the depletion zone changes when voltages are applied to the p and n regions. Forward biasing is when a positive voltage is applied to the p-zone and a negative voltage to the n-zone. This biasing causes the depletion to shrink and allows charges to flow between the two zones. Conversely, reversing the voltages creates the opposite effect where free charges are drawn deeper into their respective zones, thus causing the depletion zone to widen.

2.2.1.1: Charge Storage

Understanding how p-n junctions can be manipulated creates one of the basic building blocks of the pixel, the Metal-Oxide-Semiconductor (MOS) Capacitor. Combining the p-n junction with a polysilicon gate allows the voltage biasing of the p-n junction. By creating a potential difference across the gate structure, the depletion zone can be widened, as discussed, and this creates a high potential well that attracts electrons, specifically those generated from incident light.

2.2.1.2: Photoelectric Effect

As described initially by Einstein, the internal photoelectric effect happens when an incident photon has enough energy to move a valence electron to the conduction band; Photon energy, $E,$ is described as in (3).

$$E = \frac{hc}{\lambda} \tag{3}$$

Here h is the Planck constant, c is the speed of light in a vacuum, and λ is the wavelength of the incident photon. This equation means that a wavelength limit can be described for a photosensitive material by setting E equal to the bandgap, E_{band}. Moreover, rearranging for the wavelength.

$$\lambda_{\text{limit}} = \frac{hc}{E_{\text{band}}} \tag{4}$$

Using silicon as an example, at a temperature of 0 K, the bandgap energy is approximately 1.18 eV; this yields a wavelength limit of approximately 1050nm. This limit shows that silicon is photosensitive in the near-UV, optical and near-infrared wavelengths, from about 200 to 1050nm. Wavelengths longer than these do not carry enough energy to promote a silicon electron to the conduction band, but those shorter do. Once an electron is moved from the valance band, the atom's net charge increases positively by one electron charge. This increase in positive charge is characterised by generating an electron-hole pair. The space previously occupied by the electron is referred to as a hole. Holes are analogous to a positively charged electron. So, when the photon interacts with the silicon lattice, it generates electron-hole pairs. These electrons are then subject to the nearby potential well within the pixel and stored there during the collection phase of the CCD's operation. The holes generated now occupy the valance band positions that the electrons previously occupied.

2.2.1.3: Quantum Efficiency

Not all incident photons are converted into charge, the rate at which this happens is called Quantum Efficiency (QE), and in an ideal device, this would be 100%. When CCDs were first invented, QE was one of the device characteristics that warranted improvement. One of the critical improvements was using back-side illuminated devices instead of the original front-side illumination.

The main issue with front-side illumination is that incident photons can also interact with the gate structures and other non-photosensitive material on this side of the device. The electrode structure can absorb and reflect the incoming photon flux. This effect will reduce the number of photons that the device can record. This problem is avoided with back-side illumination, with the significantly more exposed photosensitive area. Back-side illumination is usually achieved by thinning the back-side substrate layer, allowing wavelengths with less energy and therefore less penetrating power to reach the epitaxial layer, where electron-hole pairs may be detected. More photosensitive area means that more of the incident light can be captured and recorded, which increases the overall QE of the device, up to values of high 90% QE. Anti-reflection (AR) coatings are also used in conjunction with this approach to further aid the QE. Backside thinning results in a surface with high specular reflectivity and is problematic for UV and other high-energy wavelengths (Lesser, 1994). AR coatings can be optimised for various wavelengths and combined into multi-layer coatings that increase sensitivity to multiple wavelengths.

Another issue created by the backside thinning approach is a generation of a considerable backside potential well. This potential well happens because a silicon-oxide film grows on the back-thinned surface when exposed to air. This silicon-oxide layer contains impurities, resulting in a positive charge that generates a potential well, attracting free electrons within the silicon interface within the device. Therefore, some electrons generated near this silicon-oxide potential well will be taken away from the pixel generated potential wells and therefore do not contribute towards the recorded signal.

This problem can be counteracted by applying layers of negative charge to the back-thinned surface, typically accomplished by the controlled growth of an oxide layer (Janesick J. C., 1987).

Back-thinned CCDs offer much greater sensitivity to most wavelengths, especially those at higher frequencies, making it a desirable modification.

2.3: Charge Transfer Method

Once the photon collection phase of the device’s operation is over, the charge is then transferred towards the readout circuit. The charge in the image area and subsequent storage area for CCDs is moved towards the serial register, between each pixel and the areas between each pixel, known as channel stops. Channel stops effectively isolate charge packets in a column from other packets on the device using methods such as doping.

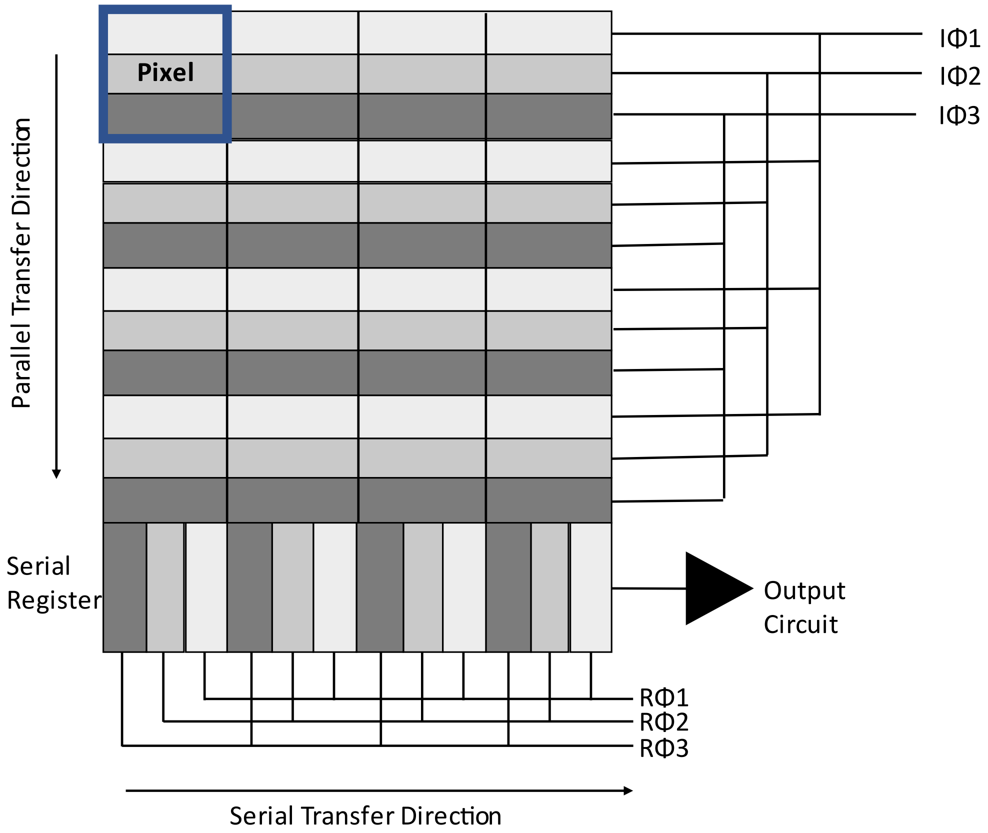
Pixel
Parallel Transfer Direction
IΦ1
IΦ2
IΦ3
Serial Register
Output Circuit
RΦ1
RΦ2
RΦ3
Serial Transfer Direction

Figure 2-3 - The layout for CCD.

In Figure 2-3**,** the charge packets in the image area are moved towards the serial register, and each parallel clock, labelled IΦ, is responsible for moving the charge packets through each electrode gate. Once the charge packet reaches the serial register, the serial (or register) clocks, labelled RΦ, move the charge packet toward the output circuit, as shown in Figure 2-4.

2.3.1: Charge Transfer Efficiency

Charge Transfer Efficiency (CTE) measures how much of the charge packet remains intact as it is moved through the device. Sometimes manufacturers will refer to Charge Transfer Inefficiency (CTI), defined as 1 – CTE = CTI. Typically for modern devices, CTE is above 99.99% per pixel.

2.3.2: Charge Transfer Mechanisms

(Janesick J. , 2001) refers to three charge transport mechanisms: transfer of thermal diffusion, self-induced drift and the fringing field effect (FFE). Thermal diffusion and FFE are dominant motion mechanisms for small charge packets. For more significant charge packets, the electrostatic repulsion from individual charge carriers within the packet is the dominant source of motion and is called self-induced drift.

If no additional external electric field influences the charge packet, the charge packet will slowly diffuse and disperse due to thermal diffusion. The random Brownian motion of electrons causes this dispersion due to the thermal kinetic energy these electrons receive from collisions with the silicon lattice. Self-induced drift is similar to thermal diffusion because electrons repel due to their negative charge. A sufficiently dense population of electrons will, over time, become less dense as electrons push each other away from the central position of the population. The dispersion caused by these effects has the impact of increasing CTI because as the charge packet becomes more spread out, the likelihood for charge packets to not be completely transferred increases. (Janesick J. , 2001) gives two equations, one where thermal diffusion applies –

$$\mathrm{CTI}_D = e^{-t/\tau_{th}} \tag{5}$$

where CTI_D refers to the charge remaining after transfer time, t. The diffusion drift time constant τ_{th} is given as –

$$\tau_{th} = \frac{L^2}{2.5D_N} \tag{6}$$

where L is the gate length and diffusion coefficient is given as –

$$D_N = \frac{kT}{q}\mu_{\mathrm{SI}} \tag{7}$$

here k is the Boltzmann's constant, T is the temperature, q is the electron charge and μ_{SI} is the electron mobility. Although electron mobility improves with lower operating temperature, due to less scattering collisions, the overall diffusion time increases because of the kT term.

The equation for self-induced drift applies –

$$\mathrm{CTI}_{SID} = \left(1 + \frac{t}{\tau_{SID}}\right) \tag{8}$$

where τ_{SID} given as –

$$\tau_{SID} = \frac{2L^2 C_{EFF}}{\pi \mu_{\mathrm{SI}} q Q} \tag{9}$$

here, C_{EFF} is the effective capacitance (F/cm^2) and Q is the number of electrons per unit area (e-/cm^2) for the charge packet at $t = 0$. This CTI contribution decreases with time because as the charge is transferred the field strength decreases and so the repulsion force causing the charge packet to spread out decreases.

2.3.3: Traps

Another contribution towards CTI is defects in the silicon lattice, known as traps. These can occur for two main reasons: defects introduced during manufacturing and radiation damage. The traps can retain a fraction of the charge packet as it is transferred through a region where they are present. The charge in the trap is released later, typically joining another charge packet. Hence traps can cause some pixels to report a lower-than-expected signal and others higher. This impact can be particularly problematic if the trap, or traps, are within the serial register or near the output circuit, as this region of the device would see the vast majority of all the charge packets, thus allowing the trap to influence nearly every pixel on the device. However, since these traps have a higher probability of being filled by the passing charge packets, the serial CTI is typically much less than the parallel CTI contribution. Modern manufacturing processes are much better than when CCD was invented, but modern devices are still subject to defects introduced during the manufacturing process, with values typically ranging between 10^{-5} and 10^{-6} electrons per pixel.

2.4: Charge Output Circuit

Once a charge packet reaches the end of the serial register, it enters the output circuit region of the device.

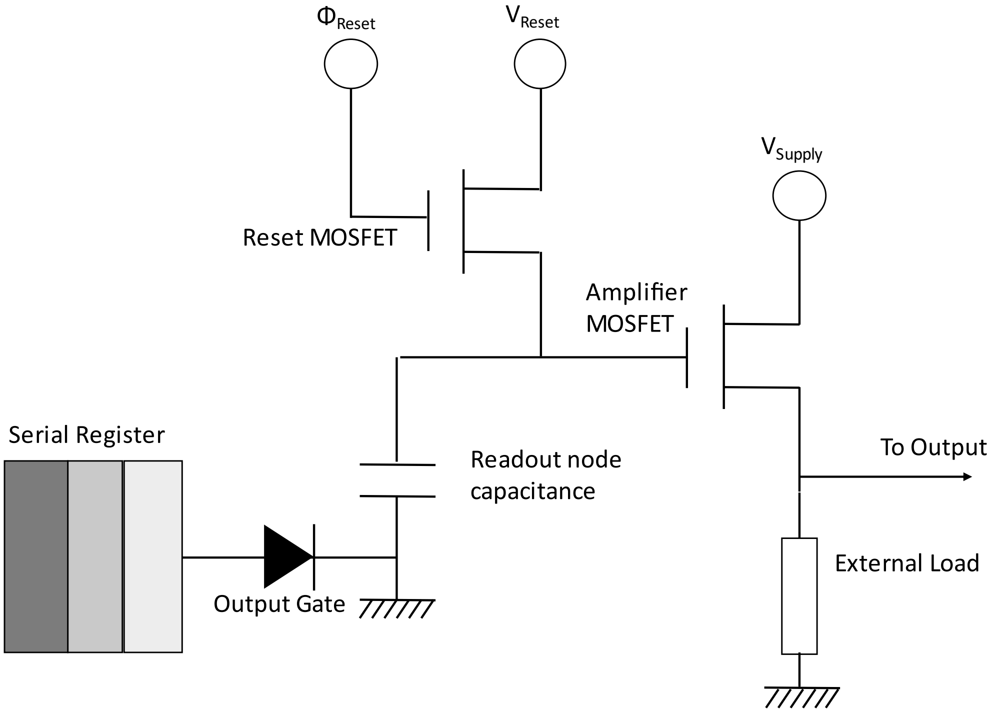
Φ_{Reset}
V_{Reset}
V_{Supply}
Reset MOSFET
Amplifier MOSFET
Serial Register
Readout node capacitance
To Output
External Load
Output Gate

Figure 2-4 – A single-stage output for a CCD.

In Figure 2-4, the charge is moved from the potential well via the final output gate onto the readout node capacitor. The presence of the charge packet on the readout node increases the voltage across the amplifier MOSFET. This change in voltage changes the current between V_{Supply} and the output paths, resulting in an approximately equal change in potential across the external load and output circuit. This potential is recorded as the digital measurement of the charge packet. The reset clock Φ_R then pulses the reset MOSFET, which both resets the voltage and current across the output circuit and clears out the output node in preparation for the next charge packet.

2.5: Amplifier Noise Sources

The source follower and the signal readout using a reset are subject to various noise sources, as outlined in (Janesick J. , 2001).

2.5.1: White Noise

First recognised by and discovered experimentally by Johnson in (Johnson, 1927), White noise, or Johnson noise is the noise associated with the thermal excitation of charge carriers, which creates a potential difference between the two ends of conductor material. Johnson provided evidence of this phenomenon using a vacuum tube amplifier and thermocouple. Very soon after, Nyquist published a paper (Nyquist, 1928) that theoretically agreed with Johnson's experimental data, showing the noise described in (10).

$$V_n = \sqrt{4kTBR} \tag{10}$$

Here, V_n is the rms noise voltage, k is the Boltzmann's constant, T is the absolute temperature, R is the resistance of the conducting path, and B is the noise power bandwidth (Hz)

We can see that White Noise is dependent on the device temperature, scaling with $\sqrt{T}$ and $\sqrt{R}$. So ideally, the device should be as cooled as possible since this will reduce both the temperature dependence and resistance, given there will be fewer thermal oscillations of the atomic structure. This effect also happens during the reset phase of the readout circuit.

2.5.2: Reset Noise

Reset noise is generated when the reset MOSFET resets the sense node. Because there is thermal noise generated by the channel resistance of the reset transistor, the reset level exhibits thermal fluctuations described by (11)

$$N_R = \sqrt{4kTBR_R} \tag{11}$$

N_R is the noise voltage across the sense node (rms V) and R_R is the effective channel resistance across the reset transistor (Ω). This noise manifests as a deviation in the reset (reference) voltage level from pixel to pixel.

2.5.3: Flicker Noise

Flicker noise, or $1/f$ noise, is generated by surface interface states trapping electrons and later releasing them with a wide range of release times. The reduction of charge carriers from the channel current is given by the ratio r_e (12).

$$r_e = \frac{\tau_e}{t_r} \tag{12}$$

Where τ_e is the emission time constant of the given interface state and t_r is the transit time for charge carriers that are in the channel. A $\frac{1}{f}$ relationship occurs because of the long-time constant states contributing charge with greater, significant variance to the output signal compared to those surface states with faster release times because of the ratio mentioned above. This means that typically, flicker noise is most impactful at low frequencies, where these long-time constant states are more visible in the spectral noise density. For higher frequencies, typically, white noise dominates over flicker noise.

An equation for approximating the total noise of the output amplifier is given by (Janesick J. , 2001). Here we assume that the two dominant noise sources are white noise and flicker noise.

$$N_{CCD}(f)^2 = W_{CCD}(f)^2 \left(1 + \left({}^{f_c}\!/_{f}\right)^m\right) \tag{13}$$

In equation (13), $N_{CCD}(f)$ is the amplifier MOSFET noise power density (V Hz$^{-1/2}$), $W_{CCD}(f)$ is the white noise voltage (V Hz$^{-1/2}$) and f_c is the flicker noise corner frequency

(Hz), defined as the frequency f (Hz) where the flicker noise power equals the white noise power. The constant m varies between 1 and 2 depending on the drain bias. Figure 2-5 shows this relation between white noise, flicker noise and the noise power density.

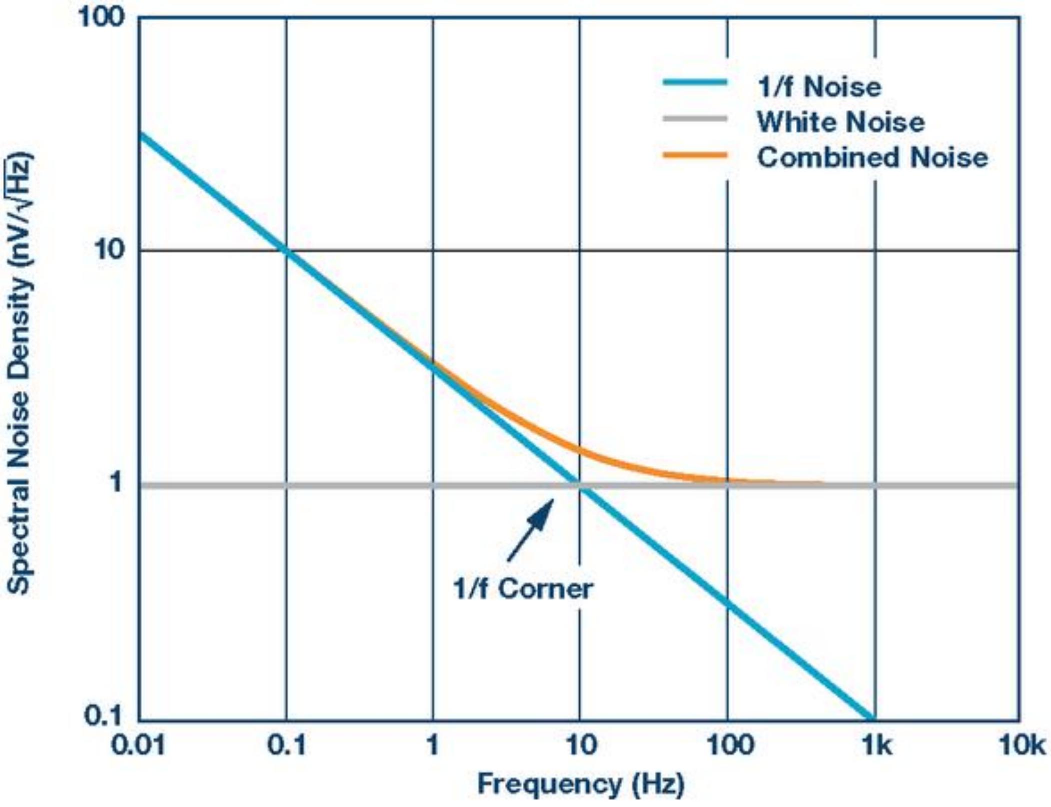

Figure 2-5 – Taken from (Castro, 2017), Spectral Noise Density versus Frequency.

2.5.4: Shot Noise

Shot noise is caused by the sudden appearance and disappearance of charge carriers, i.e. spontaneous generation of charge carriers results in a spike of shot noise. It is described in (14).

$$I_{SN} = \sqrt{2qI_D B} \tag{14}$$

Here I_{SN} is the full Shot noise (rms A), I_D is the drain current (A). We can see that it is wholly dependent on current and not temperature. This dependency means it is crucial to keep the drain current as optimised as possible; shot noise will also be unnecessarily high. However, the contribution of transistor shot noise is minimal, and as a result, it is only a visible impact during sub-threshold operation and is entirely negligible if the charge packet is large enough.

2.6: CCD Noise Sources

Next, several other noise sources that impact the charge packet are described. These noise sources do not occur within the output amplifier or the reset MOSFETs.

2.6.1: Photon Shot Noise

Similar to the electronic shot noise discussed in 2.5.4:, a noise component is associated with the non-uniform arrival rate of incident photons. The standard deviation of this arrival rate is termed photon shot noise. A Poisson distribution (15) describes the noise as.

$$P(k) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{15}$$

The mean signal level is λ and k is the observed signal level. Owed to the properties of Poisson statistics, the variance of the signal level is equal to the mean signal level, so the standard deviation, and therefore the noise, is given as $\sqrt{\lambda}$.

Incident photons that generate many additional electron-hole pairs do not show the same noise level; as discussed in (Fano, 1947), the ratio between the variance and mean of the observed signal is called the Fano factor (16).

$$F = \frac{\sigma^2}{\mu} \tag{16}$$

For a process without excessive additional electron-hole creation, the Fano factor becomes 1.

2.6.2: Dark Current

Dark current is the charge generated through thermal excitations naturally occurring within semiconductor materials. This signal can be detected even when there is no photons incident on the device. Dark current effectively limits how long a CCD can collect photo-electrons before the collection potential wells become full. The dark current generation has two sources: the interface between the silicon and silicon dioxide layers and the bulk silicon.

The surface generation occurs due to many traps at the Si-SiO_2 interface. The traps here have a distribution of energy levels that allow for easy promotion of charges from the valence band into the conduction band. Three common approaches to reducing the impact of these surface states are high-temperature annealing with hydrogen. This annealing allows the traps to form silicon-hydrogen bonds, which are much more stable and inactive under normal conditions, reducing the total density of traps across the surface interface.

Another technique for reducing surface interface trap charge contribution is Inverted Mode Operation (IMO). IMO requires the substrate voltage to be raised above the voltage at the interface, whilst pixel clocks are all at the same low voltage level. Holes migrate from channel stop regions to the surface when this happens, and they saturate the traps, reducing the electron generation rate by keeping the traps occupied.

Finally, the most common method of reducing dark current is to cool the device. Doing so reduces the system's total amount of thermal energy, thus reducing the thermal energy available for promoting electrons into the conduction band. This cooling additionally reduces the contribution from all traps, and the diffusion dark current, which is generated without traps.

2.6.3: Clock-Induced Charge

This charge is a form of noise generated when charge packets are moved through the device during the device operation. The voltage clocks that create the potential wells to move charge packets through the device create a high potential electric field. This field can accelerate the free charge carriers, i.e. holes for n-type material. This acceleration can provide enough kinetic energy for the electrons to generate additional electron-hole pairs via impact ionisation. Impact ionisation occurs when a moving, energetic electron or hole collides with the silicon lattice and breaks Si-Si covalent bonds, resulting the generation of new electron-hole pairs. This collision into the lattice transfers some of the incident particle's kinetic energy and a new electron-hole pair is created if the energy transferred is sufficiently above the silicon bandgap energy, previously discussed in 2.2.1.2:.

The generated electrons are added to the nearby charge packet and contribute to the output signal as noise. This noise type is called Clock-Induced Charge (CIC). However, this contribution is typically negligible for CCDs compared to other noise sources, such as the readout noise, given that the likelihood of electrons becoming frequently energised enough to create a substantial amount of additional electron-hole pairs is extremely small, so in almost all circumstances the contribution from CIC is minimal.

2.7: Electron-Multiplying CCD (EMCCD)

Electron-Multiplying CCDs, or EMCCDs, are variants of the conventional CCD design. The principle of charge multiplication was discovered in the 1980s (Madan, 1983). Teledyne E2V developed the first commercially available devices utilising this effect in the early 2000s. These devices were initially called Low Light Level CCD, LLLCCD, and one of the first papers discussing this technology is (Jerram, 2001). They are called LLLCCDs because these devices are capable of operating within dim lighting, low flux environments where a traditional CCD would be unable to reach the required S/N within the limiting integration time. Below is the explanation behind how LLLCCDs operate. This paper describes the initial layout design for the LLLCCD technology and explains the reasoning behind this variant's development.

The two primary noise sources on CCDs are amplifier noise and dark signal. While dark signal can be reduced via external cooling, amplifier noise, or readout noise, it remains at any pixel rate. The readout noise is especially problematic for imaging applications that rely on high pixel rates, so a new device architecture was developed to overcome this problem.

The general architecture of the device is essentially the same as a conventional CCD. In addition, an extended segment of the serial register, the multiplicative or gain register, is added, as seen in Figure 2-6.

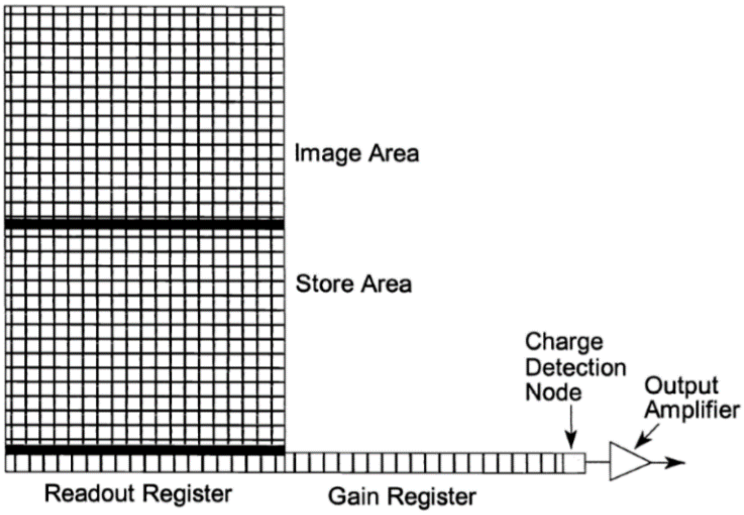

Figure 2-6 - The layout of a LLLCCD device, as shown by (Jerram, 2001).

2.7.1: Impact Ionisation

The gain register uses avalanche multiplication to increase the signal-to-noise ratio (SNR). To understand how avalanche multiplication works, first, we must understand how the signal is multiplied. This multiplication is due to a process called impact ionisation. Impact ionisation occurs when an electron has enough energy that when it collides with another electron held in orbit with an atom, the impacting electron frees the other electron. As a result, the atom becomes ionised because it has lost an electron, and the total amount of free electrons has increased. The freeing of this bound electron creates a hole.

A common circumstance for impact ionisation is when a strong electric field is present, typically generated by a high potential difference. Due to this acceleration, the charge carrier has high kinetic energy and can exceed the threshold energy required to ionise an atom, creating new free charge carriers. For decades, it has been widely accepted that impact ionisation rates scale with electric field strength; papers like (Grant., 1973) discuss how electron and hole ionisation rates vary in different materials at high electric field strength. For example, (Grant., 1973), equations (17) and (18) show how the electron and hole ionisation rates vary for epitaxial silicon. Here we define the ionisation rate as the rate of electron-hole pairs per cm^3.

$$\alpha \approx 6.2 * 10^5 \exp\left(-\frac{1.05 * 10^6 + 1.3 * 10^3\mathrm{T}}{\mathrm{E}}\right) \tag{17}$$

$$\beta \approx 2.0 * 10^6 \exp\left(-\frac{1.95 * 10^6 + 1.1 * 10^3\mathrm{T}}{\mathrm{E}}\right) \tag{18}$$

Here α refers to the electron ionisation rate, β is the hole ionisation rate, E is the electric field strength measured in V/cm, and T is the local temperature in K. Typically, we see significant variance over ranges of 10^2 K and 10^3 kV/cm. These equations both vary exponentially with electric field strength and temperature, so as the electric field strength increases, so do the ionisation rates. Conversely, as ambient temperatures decrease, the ionisation rates also increase; as thermal energy decreases for the particles, there is less random particle motion. This reduced motion means that the mean free path becomes longer for the moving charges as they are less likely to be scattered by random particle motion.

The hole ionisation rate is about a factor of 10 lower than the electron ionisation rate, as seen by the numerical weights in both equations. This difference is significant because holes have a larger effective mass, yet the saturation velocity is roughly the same, at about 10^7 cm s^{-1}. This means that because the electron is lighter, the electric field strength does not have to be as strong because it is easier to accelerate the electrons to the required speed. Because of this, impact ionisation for holes is typically ignored. However, it cannot be ignored for CIC because hole ionisation is often attributed as a source of CIC. Holes moving in and out channel stops during the readout phase generate additional spurious charges.

2.7.2: Multiplication Register

To make use of this effect, each multiplication register element uses two specific electrodes to create the conditions for impact ionisation to occur in a controlled manner. This mechanism can be seen in Figure 2-7.

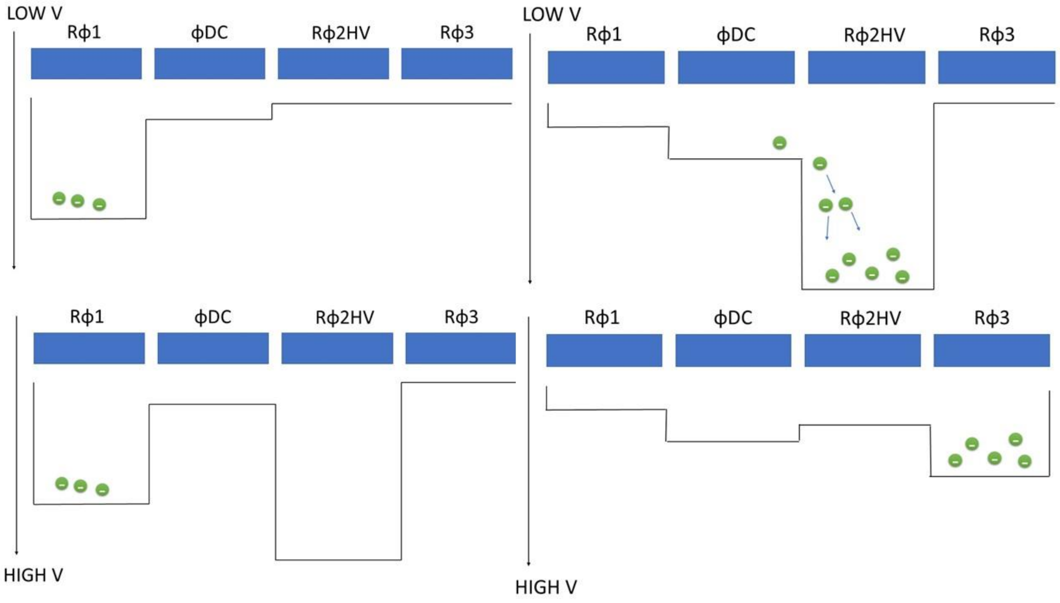

Figure 2-7 – A four electrode clocking scheme that is used in EMCCD to multiply the charge packet.

The initial phase, charge carriers are held under the RΦ1 electrode. A high voltage pulse is applied to RΦ2HV, creating a high potential difference between ΦDC and RΦ2HV. RΦ1 is reduced below ΦDC's voltage, and this causes the charge carriers to move from RΦ1 to ΦDC, which are then accelerated into the RΦ2HV well. This acceleration gives the charge carriers enough energy to cause impact ionisation. (d) The RΦ2HV returns to low voltage, lower than RΦ3 which moves the charge packet into RΦ3. This cycle continues n times, equal to the number of gain elements in the multiplication register.

The mean gain per multiplication element (stages) R is very low; typically about 1%. The device's gain G is the mean gain over the whole register and can be much larger, as calculated in (19).

$$G = (1 + R)^n \tag{19}$$

Given the small size of R, to reach reasonable levels of device gain, n must be in the order of several hundred. An example of how a typical gain register works with n = 625 and R = 0.01 (1% mean gain), the total gain calculated from (19) would be 502.18. Using a mean gain per stage of 0.01, it is expected that the multiplication register should be 625 stages long to achieve an average total gain of 500 electrons per signal electron. However, given that the number of electrons gained during the process relies on probability, the output varies, and the total amount of electrons gained during this process is not the same for every charge packet. For N input electrons, we expect to see some variance in output electrons expressed as $\sqrt{N}$.

2.7.3: Excess Noise Factor

The amount of input electrons is also uncertain due to effects such as shot noise. Photon shot noise occurs when observing a photon source and is an inherent property of light. The rate at which photons are incident onto the device can be modelled as Poisson distribution. Since the standard deviation of shot noise is given as the square root of the average amount of input signal, the input signal to noise ratio (SNR) is (20).

$$SNR = \frac{N}{\sqrt{N}} = \sqrt{N} \tag{20}$$

For EMCCD devices, the SNR does change at high gain levels. This change is due to the signal dependent uncertainty added by the gain mechanism. This mechanism is described by both (Robbins, 2003) and (Hynecek J. N., 2003) as an excess noise factor F(21):

$$F^2 = \frac{\sigma_{out}^2}{G^2 \sigma_{in}^2} \tag{21}$$

Here σ_{out}^2 is the variance of the outgoing signal from the multiplication register. This variance increases as the gain G increases. σ_{in}^2 is the variance of the incoming signal into the multiplication register; this is independent of the gain but is subject to shot noise. Both papers above show experimentally and analytically that F^2 tends toward two at high gain. The SNR for this noise component is given in (22).

$$SNR = \frac{S}{N} = \frac{N}{\sqrt{2}} \tag{22}$$

Typically, this is combined, in quadrature, with the Poisson shot noise, (15), yielding this overall SNR.

$$SNR = \frac{S}{N} = \frac{N}{\sqrt{2N}} \tag{23}$$

This SNR is only valid for EM devices at high gain. At low gain, the device reverts to (20). This can be seen in Figure 2-8.

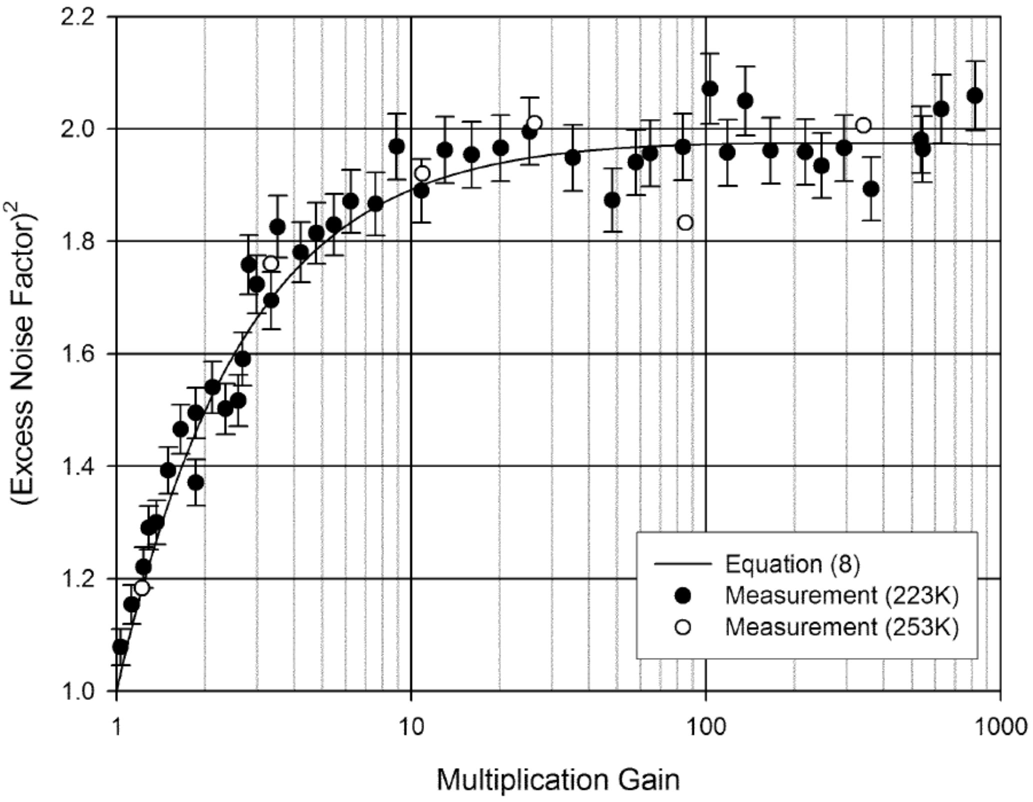

Figure 2-8 - Excess noise factor as a function of the multiplication gain. Taken from (Robbins, 2003). Note that the equation (8) in the Figure refers to equation (21) in this text.

(Jerram, 2001) gives us a total input-dependant noise equivalent signal (NES), in darkness, by considering the amplifier (readout) noise in electrons RMS N_a and the EM gain G.

$$NES = \sqrt{F^2 N_d + \frac{N_a^2}{G^2}} \tag{24}$$

N_d is the mean dark current signal in electrons per pixel. NES is typically quoted in electrons rms. However, (24) shows us that the EM gain should be greater than the readout noise to achieve a higher SNR; the readout noise and dark current should be kept as small as possible via the reduction methods mentioned previously. This reduction will keep NES as low as possible, which is needed for high SNR.

At high gain $F^2 \approx 2$ and equation (24) becomes (25):

$$NES = \sqrt{2N_d + \frac{N_a^2}{G^2}} \tag{25}$$

A method of measuring NES is discussed by (Robbins, 2003). The method plots the output signal variance divided by the system gain against the total output signal. The gradient of the graph shows the $F^2$of that gain value. We can see that these controlled impact ionisation events can allow EMCCDs to multiply the input signal, causing any noise added after the multiplication register to become effectively suppressed at the cost of some additional excess noise.

(Stefanov & Dunford, 2018) explain that the unique clocking scheme, seen in Figure 2-7, must be used; otherwise, little to no gain would be achieved. This requirement is because the electrons transfer between gates at timescales of a few nanoseconds, whereas the time it takes to raise the electric field of the high voltage gate is much greater. This reason is why the charges are held under the first gate, and the use of a DC gate is required to maintain a potential barrier.

The structure design of the gates themselves can better facilitate impact ionisation. (Stefanov & Dunford, 2018) show that bespoke gate structures either increase the electron charge density, allow for more focused electron impact ionisation and higher rates, or increase the effective length between the two gates, causing the electrons to spend more time travelling through the high potential allowing more impact ionisation events to occur. Both methods allow for more impact ionisation without increasing the electric field's strength or voltage over the HV gate.

2.7.4: Clock-Induced Charge in EMCCDs

While increasing the voltage does increase the impact ionisation rates, there are problems with this approach due to a stronger electric field. The primary issue is an increase in spuriously generated clock-induced charge. During charge transfer, CIC appears in both the multiplication register and, in the image/storage area. The primary source of CIC are holes that undergo impact ionisation.

The electric fields that manipulate the electrons also cause holes to move in the opposite direction. As they are accelerated, they can undergo impact ionisation, causing the generation of new electron-hole pairs. The additional electrons can then be collected into the same potential wells that hold charge packets, thus contributing to the total noise. For EMCCDs, CIC is generated in the store and image areas like standard CCDs and within the multiplication register under similar circumstances. CIC Generation before the multiplication register means that the electrons added by CIC are subjected to EM gain. So even a tiny contribution from CIC can be significantly magnified at higher EM gains to a point where its contribution is easily noticed. This effect can be seen in Figure 2-9.

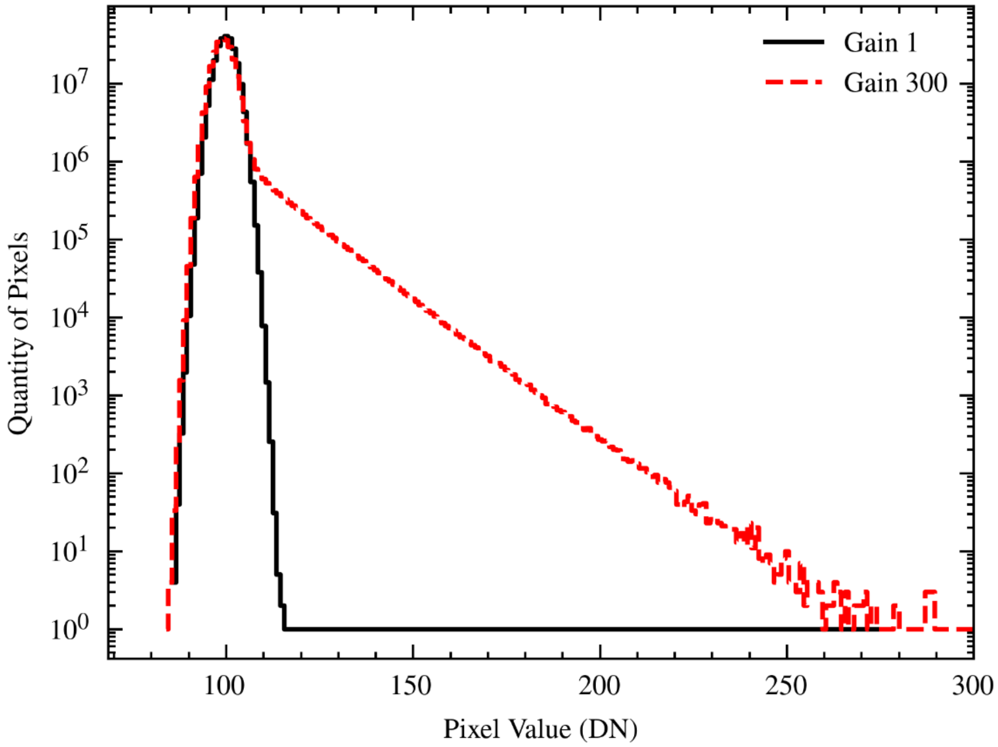

Figure 2-9 – A comparison between 1000 frames taken on the CCD97 at Gain = 1 and Gain = 300. The X-axis shows the signal in DN, and the Y-axis is the number of pixels for a given output value.

Rather than showing a signal distribution that follows a Gaussian distribution, as would be the case for Gain = 1 (no EM gain), there is instead a side tail that contains pixels with much higher signals. It is important to note that the Y-axis is plotted logarithmically, the peak value for Gain 1 is about $0.5 * 10^7$pixels greater than Gain 300, these pixels are instead distributed into the tail of Gain 300 distribution.

Figure 2-9 shows that high EM causes CIC to become a prominent noise source. Although not all of the noisy tail pixels can be attributed to CIC, a significant amount of tail pixels spread over an extensive range of values. There are device optimisation techniques that can be applied to reduce the amount of CIC at high EM gain. (Janesick J. , 2001) discusses that the serial and parallel voltage rising clock swings gradient notably impacts spurious charge generation, with a sharper gradient resulting in the more spurious charge being generated.

Device temperature also has a notable impact on the generation of spurious charge. The colder the device, the less energy is available for thermal vibrations from atoms. This reduction means that the free mean path for free electrons becomes longer due to the smaller likelihood of colliding, not impacting, with atoms, which leads to more free electrons and holes being able to undergo impact ionisation, so the rates increase. Unfortunately, lower operating temperatures are also used for reducing the dark current, the benefits of which outweigh the decreasing CIC. As a result, it is far better to control the clock voltage shapes than increase the operating temperature.

Papers like (Daigle O. Q., 2010) and (Bush, 2021) show approaches to minimising the impact of the CIC, such as the generation of finely tuned voltage clock waveforms. Daigle (Daigle O. Q., 2010) shows that changing the clock waveform shape can have an impact on the quantity of CIC generated, with triangular and sinusoidal waveform shapes generating noticeablely less CIC than square waveforms.

2.7.5: EMCCD Ageing

As discussed in (Ingley R., 2009) this effect, also called ageing, shows a monotonically decreasing EM gain with the device’s operational time. This decrease is permanent and to achieve a constant total gain, the high-voltage electrode must be operated at increasingly higher voltage levels. The gain decrease appears logarithmic, with two distinct timescales. The first timescale occurs over the device’s first few hours of operation. During this time the gain drops by as much as 10x. Manufacturers refer to this period and loss as conditioning. EMCCDs that are sold commercially are put through intense conditions, i.e., operation at high gain values, so that the conditioning phase of the device’s lifetime is complete.

Once conditioned, the multiplication gain for an EMCCD will decrease much more gradually over the lifetime of the device. Although this does depend on the operation conditions of the device, as operating the device at high gain values, or high voltages, will contribute more loss of EM gain.

It is thought that the cause of ageing is a result of trapped electrons within the Si – SiO_2 layer beneath the high-voltage electrode. These trapped electrons contribute to reducing the electric potential beneath the high-voltage electrode, meaning that to achieve the same total gain, more voltage must be applied to this electrode to maintain the same level of electric potential.

Chapter 3: Complementary Metal-Oxide Semiconductor (CMOS) Image Sensors

CCD and EMCCDs are the primary devices for many ground-based and space-based cameras. However, there is a competitor architecture that has seen an incredible amount of development over the last few decades called CMOS image sensors (CIS). This chapter aims to introduce this device architecture, how it differs from CCDs, and its advantages and disadvantages.

3.1: Active Pixel Sensors (APS)

One of the critical differences in the layout of the image area of CCDs and CISs is the use of APS. APS is when individual pixels contain their readout circuit rather than having one readout circuit for the entire frame.

Simple CIS pixels contain a photo-sensitive area that collects photogenerated electrons (a photodiode), a reset transistor, a row select transistor, a source follower, and a common column bus in Figure 3-1.

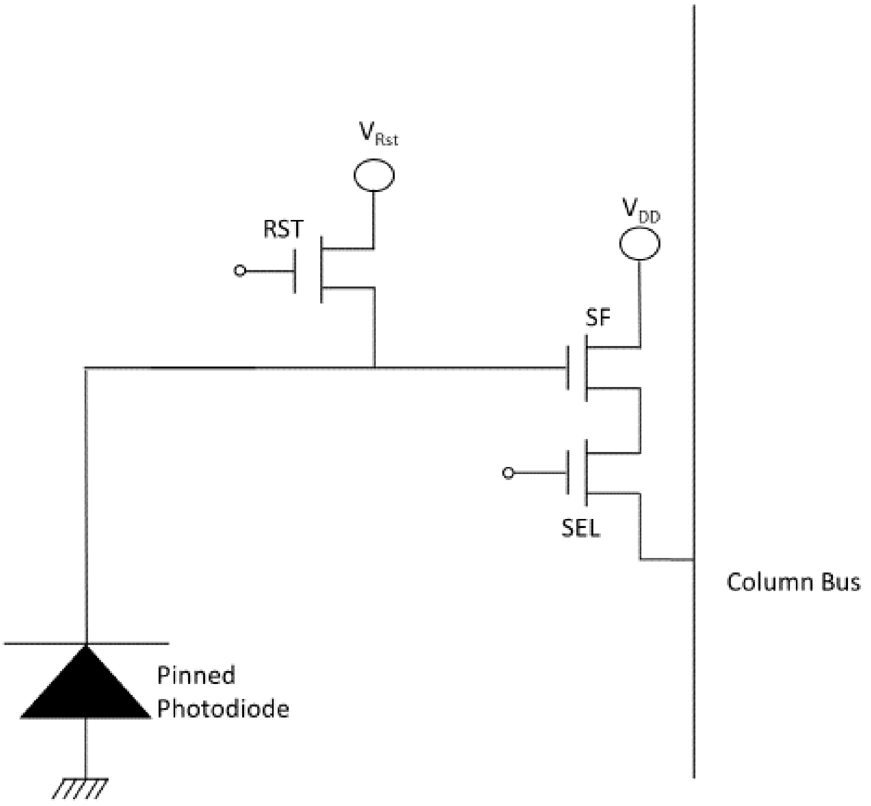

Figure 3-1 - Basic three transistor (3T) CMOS pixel layout.

The initial step is the signal collection phase, where the photodiode captures any electrons generated within the photosensitive material. As the charge within the photodiode increases, the voltage across the photodiode decreases. The change in voltage is buffered by the source follower MOSFET, which, when enabled by the row select MOSFET, passes the voltage change onto the column bus. Next, the column bus connects to the output for the sensor. Finally, the reset MOSFET is pulsed and clears all the charge on the photodiode, resetting it to the reference voltage V_{RD} to begin a new measurement.

3.2: Correlated Double Sampling (CDS)

Correlated Double Sampling is a technique used during the pixel readout process to reduce the impact of fixed pattern noise (FPN). FPN describes an offset of the pixel signal and gain due to non-uniformities in device manufacture. This offset can be split into two parts: photosensitive and non-photosensitive. The photosensitive component is usually referred to as the Photo-Response Non-Uniformity (PRNU) and appears as a variance in output signal when the sensor is exposed to a uniform light source. This variance is measured across the device for different illumination levels, and then PRNU is expressed as a percentage of the average signal.

The non-photosensitive component is called the Dark Signal Non-Uniformity (DSNU). All pixels have an inherent offset level added to prevent pixels from reading out negative values caused by the noise added by the readout circuit. This offset value is entirely independent of input illumination and ideally would be the same for each pixel. However, this offset value is not uniform across all pixels, and the fluctuation in this offset level, caused by threshold voltage non-uniformities between each pixel, is called DSNU.

DSNU can be measured by taking exposures in darkness without incident light on the device and creating a mean signal value for each pixel over as many frames as possible. Once this mean signal is collected, this can be subtracted from each frame, approximately correcting for each pixel’s DSNU.

CDS happens during the charge collection phase, the transfer gate is biased at a low voltage, usually that of the substrate, and acts as a potential barrier to the collecting charge packet, confining it to the pinned photodiode. Whilst the photodiode collects charge, the voltage level within the empty floating diffusion collection zone is measured using the row select MOSFET. This measured voltage is stored on a capacitor within the CDS circuit called the Sample Hold Reset (SHR).

Once the collection phase of the pixel has ended, the transfer gate voltage is changed so that the potential barrier is low enough to allow the charge packet within the photodiode to flow into the floating diffusion. Once the packet is fully transferred, the transfer gate returns to its previous voltage, re-establishing the potential barrier.

The voltage at the floating diffusion is measured by the row select MOSFET, and this value is now stored on another capacitor within the CDS circuit called the Sample Hold

Signal (SHS). The voltage value held by the SHS capacitor has the SHR’s value subtracted from it. This removes any DC off-set at the floating diffusion prior to the charge packet’s storage and removes the reset noise.

3.3: 4T Pixels

4T pixels are the next step in CMOS pixel design, containing an additional transistor called the transfer gate. This transfer gate is a barrier between the photodiode and the floating diffusion storage area. The benefits of this addition allow for CDS to be used in each pixel, allowing the removal of the reset noise as described in chapter 3.2:

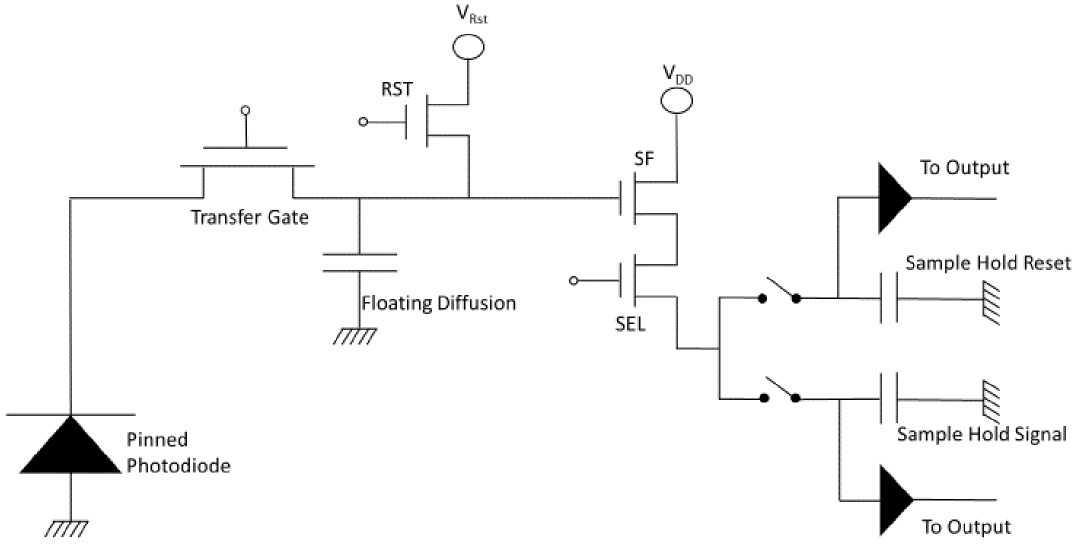
V_{Rst}
RST
V_{DD}
SF
To Output
Transfer Gate
Sample Hold Reset
Floating Diffusion
SEL
Sample Hold Signal
Pinned Photodiode
To Output

Figure 3-2 - A simple layout of a 4T CIS pixel

3.4: 5T Pixels

5T pixels include the use of a global reset transistor, this allows the device to reset the photodiode within all pixels at the same time. The benefit of this is that it eliminates any rolling shutter lag across the frame.

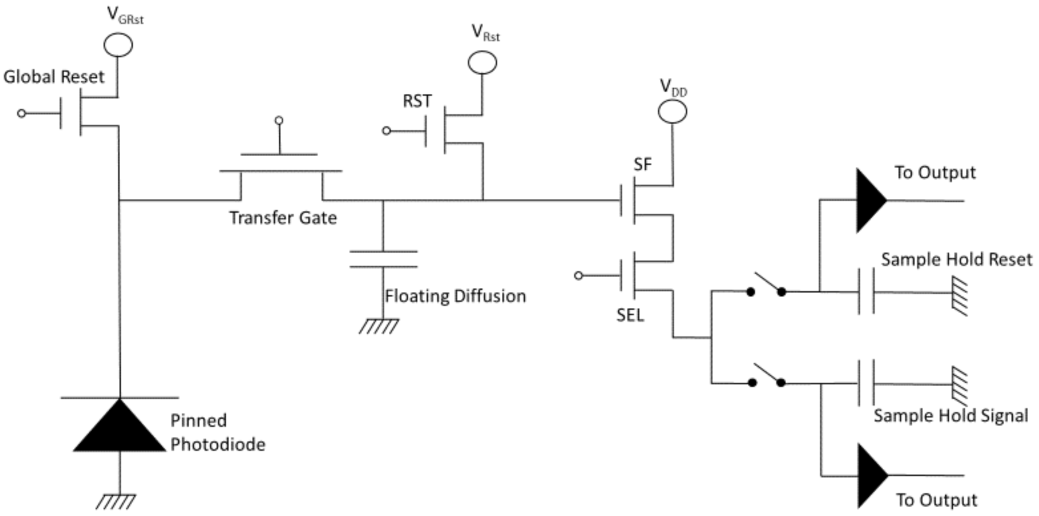
V_{GRst}
Global Reset
V_{Rst}
RST
V_{DD}
SF
Transfer Gate
Floating Diffusion
SEL
To Output
Sample Hold Reset
Sample Hold Signal
To Output
Pinned Photodiode

Figure 3-3- 5T pixel schematic.

3.5: CIS Readout Modes

3.5.1: Rolling Shutter

Rolling shutter is a capture method that acquires the scene sequentially across the device, either vertically or horizontally, but most importantly, not all simultaneously. This capture method effectively means that the rows are not read out simultaneously. Capturing the scene across the device at the same time is called Global shutter. The benefit of a rolling shutter is that the framerate for this operation mode is typically higher than the global shutter.

This comes with some trade-offs, however. For example, operating in the rolling shutter mode can introduce imaging artefacts, mainly in the form of image distortion. A typical example of this is when imaging an object that is moving or rotating very quickly, at a rate close to or greater than the readout speed, then the object will appear smeared across the frame as the object moves across the frame.

The global reset transistor enables a global shutter, as readouts between pixels can be synchronised. However, there is an issue with this. If the global reset transistor is used, correlated double sampling cannot be used. This is because CDS requires sampling the reset level, but the global reset transistor does not apply a reset pulse here; it resets the photodiode directly.

To use CDS, the device must be operated in rolling shutter mode, which effectively reverts the pixel to a 4T pixel, given that the operation will be the same.

3.6: CIS Characteristics

One of the critical differences between these two sensor architectures is how the readout circuit is structured. In CIS each pixel has its own internal readout circuit, as discussed. CCDs, however, have a single readout circuit for the whole array of pixels. This has several impacts that make the treatment of noise in CIS devices different to CCDs.

Firstly, the noise from the readout circuit is not the same for each pixel. In CCDs, the whole device is served by the same readout circuit, so some noise sources are the same for each pixel. For example, the spread of reset transistor noise values is the same for every CCD pixel in the frame. This is not the case for CIS devices. One pixel could have a low spread of reset transistor values, and another could have a large spread of reset transistor values. Other noise sources in the reset circuit, like white and flicker, will always vary between readouts in CCD and CIS devices.

This will also apply to other noise sources previously discussed in the readout circuitry, such as flicker and white noise. However, another noise source is much more impactful in CIS devices than in CCDs, called Random Telegraph Signal (RTS).

3.6.1: Random Telegraph Signal Noise

Physically similar to flicker noise, RTS is generated from traps existing within the readout circuity. Single trap sources typically generate RTS compared to flicker noise which is generated by a large number of traps. RTS noise is more prominent in CIS than CCDs because CCDs have larger readout circuits, allowing for statistically more significant quantities of traps, whereas, CIS devices have much smaller readout circuits, and as a result, are more impacted by RTS noise. The capture τ_c and emission τ_e trap time constants are given in (Janesick J. A., 2006) as –

$$\tau_c = \frac{1}{\sigma_t v_{th} n} \tag{26}$$

$$\tau_e = \frac{\exp\left(\frac{\Delta E_c}{\mathrm{kT}}\right)}{\sigma_t v_{th} N_c} \tag{27}$$

Here σ_t is the average capture cross-section in cm^2, v_{th} is the average thermal electron velocity in cm s^{-1}, n is inversion layer electron density in carriers cm^{-3}, N_c is the effective density of electron states in the conduction band in states cm^{-3}, ΔE_c is the energy of a trap below the conduction band edge in eV, k is the Boltzmann constant and T is the absolute lattice temperature in K.

Both time constants depend inversely on the average capture cross-section and the average thermal electron velocity. The capture time constant is also inversely dependent on the density of carriers, so a more significant charge density leads to a higher capture rate. Similarly, the emission time constant is inversely dependent on electron-state density, so a greater state density leads to quicker emission times. We can see that cooling the device will increase the emission time constant and lower emission rates.

Ultimately, RTS noise is the limit for CIS if the reset noise is reduced dramatically. There are some methods available that do reduce the impact of RTS. Device cooling can reduce the emission rate of traps and also help with dark current reduction but does not reduce the capture rate of traps.

Correlated double sampling does help to some extent. Much of this theory is covered in detail (Janesick J. A., 2006) and (Janesick J. E., 2015).

Another reduction in RTS noise is found by utilising P-type MOSFET (PMOS) transistors. Experimental data shown in (Janesick J. E., 2015) and (Janesick J. E., 2017) shows that PMOS CIS yield less RTS noise when compared to NMOS. In PMOS, most free charge carriers are holes, and holes are less likely to be captured by the oxide traps because the potential barrier for holes is more significant to overcome than electrons.

3.7: CCD vs CIS

There are several benefits to the CIS technology for imaging compared to the CCD.

3.7.1: Charge Transfer Inefficiency (CTI)

Physically moving charge packets across long distances in CCDs introduces additional noise sources, firstly, charge transfer inefficiency. For an average device, charge packets are subjected to thousands of transfers. The CTI must be as small as possible, often below 10-5 per pixel in modern CCDs. However, for a typical CCD with 1024×1024 pixels, a CTI as low as 10^{-5} leads to a charge loss of 2%. (Janesick J. , 2001) defines the amount of charge lost for an N-stage CCD register as –

$$S_{N_{p+n}} = \frac{S_i N_p!}{\left(N_p - n\right)!\, n!} (1 - \mathrm{CTI})^n \mathrm{CTI}^{(N_p - n)} \tag{28}$$

S_i is the initial charge held within the start pixel (in electrons). N_p is the number of pixel transfers, and n is the number of subsequent pixels after the starting pixel. $S_{N_{p+n}}$ is the charge held within the N_{p+n} pixel. This only gets worse for larger devices as the number of transfers increases. None of this is an issue for CIS-type devices, given that they do not have any charge transfer requirements. CIS devices do, however, suffer from pixel lag which has similar effects to CTI. Sufficiently bright pixels can impact subsequent pixel readouts in the form of the residual signal. This residual signal is caused by an incomplete reset, leaving behind some signal from the initial charge packet. This leftover signal is then readout in the next readout phase, contributing to the next charge packet. A hard reset can somewhat mitigate this where the reset gate voltage exceeds the transistor threshold voltage and the source follower transistor bias (Holst, 2011). This leads to the next issue that CCDs face.

3.7.2: Radiation Hardness

Traps in a CCD's image and storage area can have a huge impact on the CTI. This is because traps have the potential to capture charge from a charge packet and then release it into the following charge packet.

Radiation-induced traps are created by damage caused by highly energetic particles displacing atoms within the silicon lattice. This damage cause is primarily an issue for devices that operate in high-radiation environments, such as space.

Space has a variety of sources that can cause radiation damage; solar radiation and activity, galactic cosmic rays, and trapped particles within Earth’s magnetosphere. All of these sources have different particle energy and flux densities.

Given that CIS either do not need to move the charge or perform only one charge transfer, the impact of traps and radiation damage is much less significant. In addition, any damaged pixel only impacts its readout values, reducing the individual impact a trap could have on the whole array.

3.8: Scientific CMOS (sCMOS)

In recent years, many companies have begun to produce scientific-grade CIS , offering very low noise and making them competitive for many scientific and astronomical observational instruments.

(Chang, 2018) An example of a sCMOS chip that uses a dual column amplifier and dual analogue-to-digital converter (ADC). This additional circuit outside the CMOS 5T pixel allows for a combination of low noise (via the high gain path) and high full well capacity (via the low gain path) as seen in Figure 3-4.

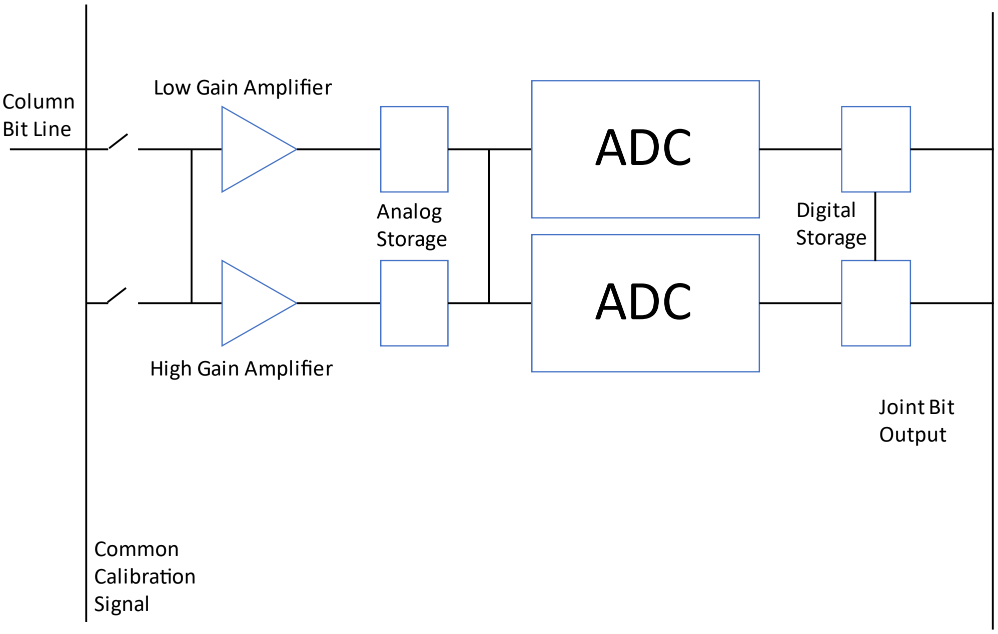
Column
Bit Line
Low Gain Amplifier
ADC
Analog
Storage
Digital
Storage
ADC
High Gain Amplifier
Joint Bit
Output
Common
Calibration
Signal

Figure 3-4 - A simple layout for how dual amplifiers and dual ADCs allow for both high and low gain operation. Based on Figure 4 in (Chang, 2018).

The dual gain readout allows for a strong comparison between EMCCDs and sCMOS devices. A table in (Chang, 2018) demonstrates this comparison.

Sensor	CIS1910F	CIS2521	CCD220	CCD351
Active Pixels	1920 x 1080	2560 x 2160	240 x 240	1024 x 1024
Pixel Size (µm)	6.5 x 6.5	6.5 x 6.5	24 x 24	10 x 10
Frame Rates (fps)	100	100	1300	30
Read Noise (e- rms)	< 1.2	< 2	< 1	<< 1
Quantum Efficiency Max	> 52%	> 55%	> 90%	> 90%
Dark Current (e- pixel^{-1} s^{-1})	< 30	< 35	10	100
Power Consumption (mW)	< 800	< 2000	< 2500	-

Table 1 - A comparison of two Teledyne e2v EMCCD devices, CCD 220 and CCD351, and two sCMOS devices, Fairchild CIS1910 and CIS2521.

Table 1 shows that sCMOS devices can have higher framerates whilst having a larger device size. This faster readout can be partially attributed to the parallel pixel readout that CIS devices utilise rather than the serial readout of CCDs. Both sCMOS devices have readout noise below two e- rms, demonstrating that an sCMOS device with readout noise equal to or better than what EMCCD can achieve is attainable.

One issue for the above sCMOS devices is that the quantum efficiency for both is much lower than the EMCCD, and as stated before, the sCMOS must capture twice as much light to achieve the same SNR as the EMCCD. The CCDs also have larger pixel sizes allowing for more light capture per pixel than the sCMOS devices. The advantage of having a larger pixel size is that the total area of photosensitive material is greater, and assuming equal quantum efficiency, the amount of time needed to achieve low-light sensitivity is lower. This is because the larger photosensitive area means that the device can capture more incoming signal, compared to a device with smaller pixels, which has less photosensitive area due to more non-sensitive pixel architecture.

Although there is no record for power consumption for the CCD351, we can see that power consumption for sCMOS is typically lower than that of EMCCDs, especially the CIS1910F being much lower. This is primarily a result of power supply requirements for EMCCDs needing high voltage clocks to clock and multiply charge packets during the readout phase. However, this difference is not that big considering that the units used in Table 1 is mW, so the difference between the CIS1910F and CCD220 is about 2W.

Chapter 4: Ground-based Observations

4.1: Atmospheric Seeing

When observing stars through the atmosphere, they appear blurred rather than point-like objects. This blurring is referred to as atmospheric seeing. Atmospheric seeing is due to the turbulent airflows that naturally occur within the atmosphere, which distort the image of the target object, typically by blurring the object. An example of this can be seen in Figure 4-1.

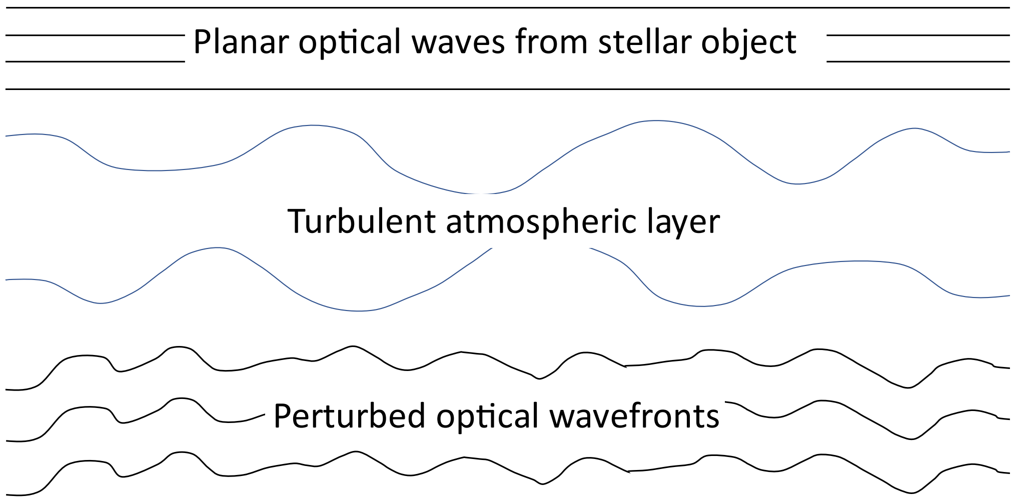
Planar optical waves from stellar object
Turbulent atmospheric layer
Perturbed optical wavefronts

Figure 4-1 - A diagram, showing plane waves becoming perturbed due to the turbulent atmosphere.

Due to different temperatures and densities in the turbulent atmospheric layers, the optical refractive index varies slightly. This variation in the index leads to the plane optical wavefronts becoming perturbed and distorted.

4.1.1: Measuring Seeing Conditions

Different methods can be used to measure the seeing conditions at the observer’s location, and most of these use the Point Spread Function (PSF) quantity of the observer’s lens, as defined (Guenther, 2018).

4.1.1.1: Point Spread Function

The PSF is a way of quantifying the performance of the optical system. It defines how the object's intensity is translated onto the image, i.e. how the light from a star is spread across the telescope's focal plane. The amplitude of the PSF defines the 2D spatial variation of the amplitude of a single point-like object. So, if there is no variation, i.e. the image and object are identical, then the PSF amplitude is zero, and there is no spread of the object’s intensity. However, this can never be the case due to the wave-like nature of light causing the signal to appear as a diffraction pattern when the light enters the optical system. The physically best image obtainable is an Airy disc. First described by (Airy, 1835), an Airy disc is the best possible image that a perfect lens can create, one without optical imperfections, with a circular aperture. Since telescopes typically use circular apertures, we assume the Airy disc to be the best possible image for these types of telescopes. In an ideal case, the PSF of a point-like source is distributed like an Airy disc, with the central peak containing most of the incident light, which then spreads away across the plane. These limiting effects that impact the degree of spatial variation will be discussed in further detail in subsequent chapters. More details on PSF can be found in (Corle, 1996).

4.1.1.2: Strehl Ratio

One method to quantify the seeing condition is the Strehl ratio. Named after and proposed by Karl Strehl (Strehl, 1895), the Strehl ratio (S) is the ratio of the peak intensity of the real image (I_{real}), which includes any optical errors introduced by the real optical system, and peak intensity of the ideal image taken by a perfect optical system (I_{ideal}), as shown by (29).

$$S = \frac{I_{\mathrm{real}}}{I_{\mathrm{ideal}}} \tag{29}$$

A high Strehl ratio means that the seeing conditions are near that of the conditions required for an ideal image, thus implying high image quality. A low Strehl ratio infers the opposite, poor image quality because the light is spread out over a larger region, blurring the image and having a lower peak intensity.

4.1.1.3: Full-Width at Half-Maximum

The Full-Width at Half-Maximum (FWHM) is another approach to measuring the quality of an image. As the name suggests, it is a measure of the width of the PSF of a stellar object at half the maximum intensity. A good quality image should therefore have a small FWHM of the PSF, meaning that the light from the star is concentrated across very few pixels. However, for images where the point-like source is blurred, i.e. a poor quality image, the FWHM of the PSF is larger because the light from the object is spread over a larger area, so the spatial variance is more significant.

4.1.1.4: Diffraction and Turbulence limit

As previously mentioned, I_{ideal} is defined as the peak intensity of the ideal PSF taken by a perfect optical system. While this assumes no atmospheric turbulence impacting the observation, the ideal PSF is still impacted by diffraction. Given the wave-like nature of light, we know that light propagates around edges, resulting in the wavefront no longer being planar but curved instead. Diffraction around the telescope aperture impacts the resulting image, leading to blurring and image degradation. The diffraction limit of a telescope is a function of the size of the telescope aperture, the focal length and the wavelength of the incident light, as shown in (Riedl, 2001) and in (30). This describes how the diameter of the central airy changes with optical wavelengths and the F-number, which is the ratio of focal length to aperture diameter.

$$B_{\mathrm{diff}} = 2.44\lambda F_{\#} \tag{30}$$

In an ideal scenario, the quality of all images taken by an optical system would only be diffraction limited, but in practice, it will also be impacted by optical aberrations in the imaging system and atmospheric seeing. In addition, atmospheric seeing can be limited by atmospheric turbulence. This limiting factor is referred to as the turbulence limit. Finally, as discussed in 4.1:, the quality of the image is further degraded if the optical wavefront is heavily distorted before entering the telescope's aperture.

4.1.2: Aberrations and Zernike Polynomials

It is possible to quantify how the observed PSF varies from an ideal diffraction-limited PSF. Any deviation from the ideal diffraction-limited performance is called an Aberration, both optical and atmospheric.

Zernike polynomials can describe these aberrations. First derived by Fritz Zernike in 1934 and well described by (Lakshminarayanan, 2011), these are a sequence of continuous and orthogonal polynomials used to describe the shape of a waveform over the pupil of an optical system. In addition, Zernike polynomials help define the magnitude and characteristic differences between the image and the original object. These polynomials use a set of polar coordinates r, θ.

Zernike polynomials use a double indexing notation and are described in (31).

$$W(r,\theta) = \sum_{n,m} C_n^m Z_n^m(r,\theta) \tag{31}$$

Here the polar coordinates are defined as $W(r,\theta)$, and can be expanded into a sequence of polynomials Z, that are orthonormal over the optical pupil's surface. C is the amplitude of the polynomials. n is the radial order, typically showing the degree of complexity of the aberration, n must be a positive integer with increasing complexity being a larger value. m is the angular frequency of the aberration and is an integer value ranging from $-n$ to n.

Aberration Type	n	m	RMS	Polar form
Piston/Bias	0	0	$\sqrt{2}$	1
Tip	1	-1	2	$r\ \sin\theta$
Tilt	1	1	2	$r\ \cos\theta$
Astigmatism	2	-2, 2	$\sqrt{6}$	$r^2\sin 2\theta, r^2\cos 2\theta$
Defocus	2	0	$\sqrt{3}$	$2r^2 - 1$

Table 2 - The low-order aberrations for the expansion of the Zernike polynomial sequence.

Whilst the sequence continues, Table 2 shows the first six aberration types; Piston/Bias, Tip, Tilt, negative Astigmatism, positive Astigmatism and Defocus. Calculating and removing these different aberration types for imaging systems allows a much more accurate representation of the source image. Two methods that use these principles to improve image quality, adaptive optics and Lucky Imaging, are described in the following.

4.2: Adaptive Optics

Adaptive Optics is a widely used technique to improve image quality and resolution. The general idea is to use a deformable mirror to correct any image distortion caused by aberrations. A bright reference star, either a suitable bright star near the target or a powerful laser pointed into the upper layers of the atmosphere to create an artificial guide star, is constantly observed to measure the shape of the wavefront so that corrections can be made using the deformable mirror. Incoming light is split into a wavefront camera and the scientific camera. The wavefront camera is a high-speed camera that usually uses a Shack-Hartmann lens array to measure the differences between the observed reference object and the known properties of the reference object. Shack-Hartmann wavefront sensors are composed of an array of lenses, typically microlenses, which focuses some of the incident light into points onto a detector array beneath the lens array. Initially calibrated with some reference wavefront, any observed incident wavefront subject to distortion causes the focusing points on the wavefront sensor to move off the reference position. This deviation in position can be used to create a correction for the deformable mirror, which would correct the wavefront distortions so that focused points return to match the reference positions. More information on the operation of these arrays can be found in (Aftab, 2018).

Then, typically using methods like the Zernike Polynomials, corrections are calculated and applied to the image to reduce the difference between the observed and expected reference images. These corrections improve the observed science image by reducing the effects of the aberrations., with image resolutions reaching those of the Hubble space telescope, which is about 0.03 arcseconds. More complex architectures have been developed to further aid in better measurements of atmospheric turbulence and correction precision. A good overview of current and past adaptive optics technologies can be found in (Hippler, 2019).

One critical upcoming adaptive optics system is the Multi-adaptive optics imaging camera for profound observations (MICADO) system (Davies, 2021), which is designed for the Extremely Large Telescope (ELT). It aims to capture images at a resolution of 10 microarcseconds over its full field-of-view of 50.5 x 50.5 arcseconds^2, making it comparable to the James Webb Space Telescope.

4.2.1: Problems with Adaptive Optics

Whilst the improvements provided by using adaptive optics are generally worth the drawbacks, these can be problematic for smaller telescopes and low budgets.

- Field-of-view

The field-of-view (FOV) for telescopes using adaptive optics is generally restricted to areas with suitably bright reference stars or being able to use and maintain a laser guide star. Typically, these guide stars need to be close to the target star, which causes the FOV to be constrained. This, in turn, leads to more time needed to observe multiple targets over a large area. Additionally, the FOV for AO telescopes changes depending on how many Zernike polynomial terms are corrected for. As the amount of terms increases, the area for which this correction is valid decreases. Effectively, AO can perform many high-order corrections, which would be valid for a smaller FOV, or do fewer term corrections which would be valid over a larger FOV.

- Expense

Adaptive optics systems cost substantial money to build and maintain. The specialised mirrors, machinery and computing requirements are costly to construct and maintain. This is typically why most high-end adaptive optics systems are only found in observatories where this budget is available and the cost-benefit is ideal. There are some cheaper, alternative adaptive optics systems available commercially, but for smaller observatories, adaptive optics is an expensive option.

4.3: Lucky Imaging

Another technique used to obtain higher-quality images is Lucky Imaging (LI). It is known that the negative impact on image quality that atmospheric turbulence has varies, randomly, over the duration of the exposure. Therefore, a probability can be described for when the amount of atmospheric distortion is minimal, leading to images that are near diffraction limited.

4.3.1: Fried Probability

The probability of getting a diffraction-limited exposure for a given coherence length, R_0, and telescope diameter D, is given by (Fried, 1978). Recall that image quality, or exposure resolution, is either turbulence-limited, due to atmospheric distortion effects on the optical wavefront or diffraction-limited, given the size of the telescopic aperture compared to the scale of the optical wavefront.

(Fried, 1978) refers to diffraction-limited resolution R_D as defined in (32).

$$R_D = \frac{\lambda}{D} \tag{32}$$

Here the optical wavelength is λ. This means that the diffraction-limited resolution becomes smaller for larger telescopes. Turbulence-limited resolution R_T is defined as (33).

$$R_T = \frac{\lambda}{l} \tag{33}$$

Here, l is the path length of the incoming wavefront, determined by the strength of the turbulence S_r over the propagation path r from source to the observer as seen in (34).

$$l = \frac{S_r}{r} \tag{34}$$

This shows that as the strength of the turbulence increases, the wavefront's path length increases. We know that the strength of the turbulence is random; this means that very rarely, S_r becomes very small, leading to the path length becoming much more minor. A smaller path length means that the image is less impacted by turbulence, so it is possible to acquire near diffraction-limited images.

Fried shows that the probability of getting a good image, which is defined as the probability that at any moment in time, the wavefront distortion and any change in high-order aberration is effectively negligible, is given as a negative exponential function of the ratio of telescope diameter D to the turbulence-limited coherence diameter, or the Fried parameter, R_0. The Fried parameter is a measure of the length, typically cm, at which the optical wavefront can be considered planar, given the atmospheric turbulence. The larger the Fried parameter, the better the atmospheric conditions are. (Littlefair, 2022) states that for a site with good seeing, the Fried parameter is typically about 10cm for an optical wavelength of 500nm. The probability of a good image is given as (35).

$$\text{Probability} \approx 5.6\exp\left[-0.1557\left(\frac{D}{R_0}\right)^2\right] \tag{35}$$

This equation is only valid for the limit $D > 3.5R_0$ as these limits are a result of the reduction of the Karhunen-Loève integral equation solved in (Fried, 1978).

The probability is suitable for telescopes with low $\frac{D}{R_0}$ values, but decreases exponentially as this value increases, as seen in Figure 4-2.

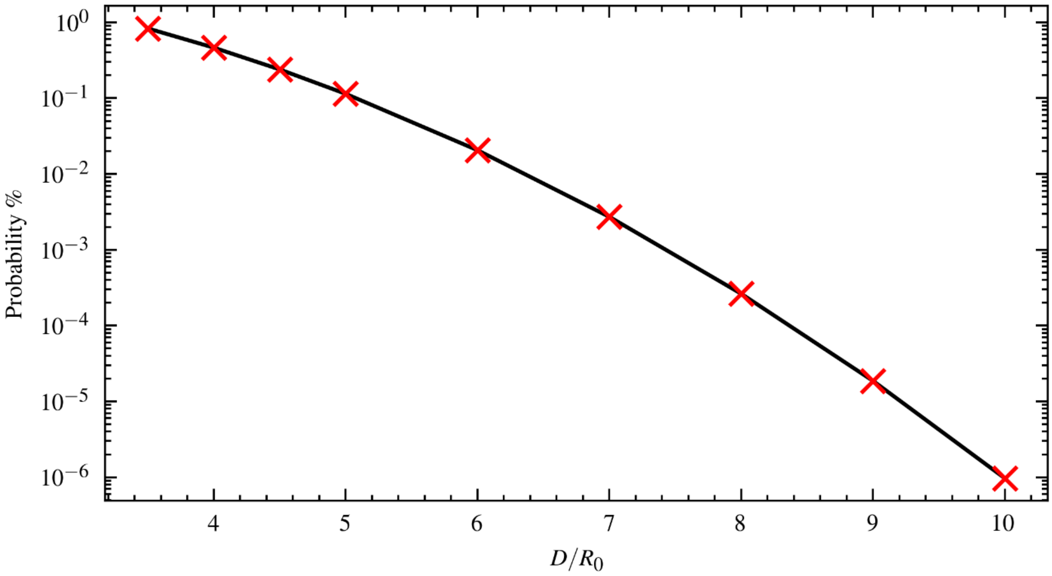

Figure 4-2 - shows the relationship between $\frac{D}{R_0}$ and the probability of getting a near-diffraction limited exposure.

This is what makes Lucky Imaging a viable approach for relatively small telescope diameters.

4.3.2: Lucky Exposures

The lucky imaging technique uses this relation between telescope diameter and probability in equation (35). Using telescopes with relatively low $\frac{D}{R_0}$ values, the likelihood of getting a near-diffraction-limited exposure is still high enough that it is feasible to take large quantities of short exposures, with a timescale low enough that the change in atmospheric impact across the frame is negligible, intending to get enough diffraction-limited exposures.

The first step for lucky imaging is taking a large number of exposures. Naturally, these images will vary due to the variables mentioned above that can impact image quality. Applying image calibration techniques, such as bias frame subtraction and flat-field subtraction can help improve the signal-to-noise ratio of the images.

Given that readout noise can be positive or negative, to prevent pixels from reporting a negative value, each pixel has a zero level, or bias level. This bias level is a digital number that represents the pixel at zero electrons, for example a pixel with zero electrons reports a digital number (DN) value of 100. A whole bias frame is subtracted from the exposure because each pixel will report a different bias level. However, when the bias frame is measured it is subject to readout noise. To reduce the impact of readout noise on the bias frame, is it necessary to create a master bias frame which is average of many bias frames. While the bias frame should be quite constant, it is good practice to repeat this measurement at regular intervals as external conditions can impact the bias level, such as device wear over time.

The flat-field frame is used to correct any imbalances between the light sensitivity of each pixel and removes any vignetting on the frame. Vignetting describes when the edges of frame appear darker than the central image. For example, many flat frames are acquired by taking images of a uniformly illuminated target (e.g. a screen in the dome lit up by a lamp or an area of the sky during twilight without any bright stars). A master flat frame is obtained by taking the mean over these frames to decrease noise and normalising. This frame can then be used to calibrate the scientific exposures to reduce the impact of each pixel's non-uniformly distributed light sensitivity.

Next, the image quality should be quantified for each image, using either the Strehl ratio or FWHM techniques, and then the images can be arranged into best to worst quality. Finally, a cut-off can be defined so that only images above a user-defined quality threshold are used for the subsequent shift-and-add.

Several factors could be considered when choosing a quality threshold. Firstly, there could be some fixed signal-to-noise level that images must be above. Some observations can be obscured easily by poor signal-to-noise, so only choosing images above a certain threshold would be appropriate. Alternatively, the width of the combined PSF should be below some width, or the Strehl ratio should be above a chosen value. The chosen threshold should be evaluated depending on the type of observation and the observational targets.

Shift-and-add technique is used to align and stack the images. Image stacking is a popular and well-used technique in astronomy as it is an easy way to produce images with an improved signal-to-noise ratio and a higher dynamic range. The brightest object in the frame aligns the images, and then the pixel values of the frames are summed to a single image. This reduces the average random background noise variance, and both faint and bright targets increase their values, making identifying fainter targets easier.

Finally, this produces an image of better spatial resolution, signal-to-noise and dynamic range than a single exposure.

4.3.3: Lucky Imaging Issues

Lucky imaging has some issues, making it less ideal for larger telescopes.

- Low probability at larger telescope diameters

As discussed in 4.3.1:, the probability of getting a lucky exposure is related to the telescope's diameter and the Fried parameter. Assuming the Fried parameter is fixed, the probability of getting a near-diffraction limited exposure decreases exponentially as telescope diameter increases. Assuming that the Fried parameter is fixed at 10cm, which is considered good seeing conditions, (Fried, 1978) tells us that having a telescope diameter seven times greater than this length, so 70cm in this example, means that the probability of getting a lucky exposure is approximately 1 in 934579 exposures. The European Southern Observatory’s (ESO) Very Large Telescope (VLT) in Cerro Paranal, Chile, has an aperture diameter of 820cm, so large telescopes of this scale are not suited for the lucky imaging technique.

- High framerate requirement

The camera needs to capture exposures on a timescale where the change in turbulence power, or how the turbulence impacts the wavefront, is approximately uniform. Imaging at a framerate lower than the atmospheric timescale makes it much more complicated to try and remove the atmospheric aberrations without additional equipment, such as Adaptive Optics.

It is generally accepted that lucky imaging requires the camera to operate at 30 frames per second or higher (Hippler. S. Bergfors, 2009). The timescale, or speckle coherence time τ_e required is related to the Fried parameter and the average wind speed velocity v, as defined in (36) by (Tubbs, 2003).

$$\tau_e \approx 0.56 \frac{R_0}{\Delta v} \tag{36}$$

- Data loss

The remaining data is discarded when choosing only to use some percentage of the higher quality images. This method means that choosing the top 1% of best frames means discarding 99% of the data taken. This data loss can lead to a reduction in visibility for faint targets. For these targets, the observer would have to take many more frames to compensate for the signal loss when discarding large percentages or accepting poorer spatial resolution in exchange for more signal.

Chapter 5: GravityCam

GravityCam (Mackay, 2019) is a proposed ground-based instrument using the lucky imaging technique to; improve the detection rate of exoplanet surveys via gravitational microlensing; improve the quality of weak shear dark matter observations in distant galaxy clusters, and provide a large amount of information for asteroseismology surveys. Additionally, there is an objective to observe objects within the Kuiper belt, and potentially the Oort cloud, to better understand the characteristics of objects within these regions.

The telescope GravityCam is planned to be designed for, is the 3.6m New Technology Telescope (NTT) at the European Southern Observatory (ESO) site at La Silla in Chile. The seeing conditions at La Silla are excellent, with a median seeing of ~ 0.75 arcseconds and a median turbulent scale r_0 of ~ 0.24m. This should allow for very high-resolution images to be acquired.

GravityCam's current sensor of choice is an array of CIS devices. EMCCDs are still very good and potentially a viable alternative due to their ability to operate at high framerates, however, the main drawback of these devices is the low fill factor of the focal plane. EMCCDs contain more non-photosensitive material compared to CIS devices.

Major developments towards large array-sized detectors and greater quantum efficiency have been made for CIS devices over the last decade. CIS readout noise is also getting consistently lower, with some ultra-low noise CIS devices reaching near or less than one electron RMS readout noise at high framerates. These characteristics mean that CIS devices could be the ideal device for the GravityCam instrument.

5.1: Device Requirements

The requirements for such a device are listed as follows:

- It must be a three-edge buttable device, preferably of large device size.

This format has a better fill factor of the field-of-view. Having a better fill factor means that the instrument can cover the field of view with fewer observations, saving on observation time. By having large devices, fewer would need to be used to achieve a good fill factor, which keeps data processing complexity, and costs down.

- A pixel scale of 85.76 milliarcseconds per pixel to match the spatial resolution at the telescope's focal plane.

85.76 milliarcseconds per pixel is the correct pixel scale for the focal plane of the NTT. Having the correct pixel scale means that the instrument does not under or oversample the observed wavelengths. Either under or oversampling will lead to the loss of information and additional aberrations that will reduce data quality. This can be corrected with additional lenses; however, this will introduce extra system complexity. This can be all be avoided by getting the correct pixel scale.

- A frame rate of at least 25 Hz.

As previously discussed, one of the critical requirements of the Lucky imaging technique is that the camera can capture frames at a rate equal to or greater than the speckle coherence time, τ_e. With a frame rate above 25 Hz, this should sufficiently meet the requirements for τ_e for best seeing conditions. Having a high framerate also improves the dynamical range of the instrument because bright objects can be observed without saturating the pixels.

- Low readout noise.

One of the key advantages of using EMCCDs for Lucky imaging is that the EM gain mechanism effectively reduces the readout noise to below one electron RMS. For a CIS to compete, it should also have a low readout noise level. This will be further explored in Chapter 8:.

- Real-time data processing and digitising on the chip.

Data management and processing optimisation are vital for this instrument. The aim to have a large array of devices to achieve a wide field of view means that vast quantities of data will be generated with each observation. The current estimate with 36 9-Mpixel devices with a framerate of 25 frames per second would generate around 25Gbs per second. Over an entire night, this would yield about 400TB. Since being able to store this amount of data is not feasible, on-chip digitisation and real time data processing is key to managing data volumes of this scale.

- High red sensitivity and deep-depletion technology.

GravityCam aims to observe many targets are red stars, i.e. much brighter in the red and NIR wavelengths than the blue ones. It is critical that the instrument is appropriately sensitive to this range (700 – 900nm), especially for fainter targets where the incoming photon flux is low. This can be achieved with thicker devices.

- Very flat detector array mounting.

To achieve the best possible image resolution, the array of detectors needs to be mounted as flatly as possible. With 16μm pixels and an incident photon-flux beam at f/8, the overall flatness variation must be no more than 15 to 20 microns.

5.2: Science Cases

GravityCam has three prominent science cases: exoplanet and lunar mass surveys, asteroseismology surveys in the Kuiper belt, and weak shear studies of dark matter distributions. These cases will massively benefit from the wide field of view and high spatial resolution that GravityCam aims to achieve.

5.2.1: Microlensing Surveys

GravityCam aims to detect exoplanets and potentially lunar masses using the gravitational microlensing method. Gravitational microlensing can occur when two stellar bodies, a pair of stars for example, are in a line of sight with the observer. The mass closest to the observer is known as the lensing mass, and the object further away is the target mass. It is well-known from general relativity that gravity can distort the path of light via the bending of space-time.

For microlensing to occur, the masses must be aligned so that the lensing mass bends and focuses the light from the target mass onto the observer’s location, like an optical lens. The result of this lensing is that the observed brightness of the target mass increases during the lensing, peaking when the two masses are aligned.

There is a particular case for a planetary body orbiting the lensing mass. As the planet orbits the lensing star, its mass contributes to the bending of space-time, and therefore the degree of change in observed brightness for the target mass. Two key elements determine the impact of that the planetary body has on the degree of lensing. Firstly, the mass of the planetary body. More massive planets will have a more notable impact on the change of observed brightness—secondly, the position and orientation of the planet’s orbit. Maximum impact occurs when the planet is aligned with the observer, lensing star, and target star. Minimal impact occurs when the planet is at its furthest from this alignment. Additionally, depending on how far from the lensing star the planet orbits, it could be aligned very briefly or for a much more extended time, depending on how long the planet takes to complete an orbit. How the planet contributes to the brightness increase, gives astronomers information about the size of the planetary body and its orbital period and radius.

GravityCam aims to detect earth-sized planetary masses and even potentially lunar masses. Regardless of orbit, the apparent brightness change due to these masses is minimal and additionally, these events can be short, so the camera must be able to detect these minor changes in signal and have a high enough time-resolution to be able to resolve them. These targets will be extremely difficult to detect if the signal-to-noise is too great. This leads to the requirement for a device with low readout noise and high framerate to enable the lucky imaging technique to provide images with high enough resolution to capture these events.

As stated in (Mackay, 2019) , the main problem with using gravitational microlensing to detect exoplanets, specifically earth-sized exoplanets, is that these events are infrequent. To counteract this problem, the main requirement becomes to be able to observe many, many stars at once. GravityCam will therefore follow the standard microlensing survey approach of observing in the dense stellar region towards the galactic bulge, known as the Bades Window, which is a dust-free region towards the galactic centre.

5.2.2: Occultation Surveys

GravityCam intends to observe and study small solar system bodies within the Kuiper belt and potentially the Oort Cloud via occultation from background stars. Occultation is simply the act of an object moving into view, blocking, or partially blocking, line-of-sight with a background object. A typical example is a solar eclipse, where the moon moves into view, blocking our line of sight with the sun, resulting in an eclipse.

For asteroids and other smaller system bodies, occultation is used to infer size and shape measurements for the occulting object. Like microlensing, occultation will benefit from GravityCam’s high signal-to-noise ratio, framerate and wide field-of-view. In addition, the large field-of-view will allow more stellar populations to be observed at once, allowing for more potential occultations.

The high framerate allows GravityCam to detect minor changes in signal as the occulting body rotates and moves, giving a more precise image of the body's shape and rotation and any surrounding material such as faint orbital rings.

The better signal-to-noise ratio would improve the quality of data collected, similar to the higher framerate benefits, and allow fainter stars to be observed, increasing the number of potential occultations that could be observed.

5.2.3: Dark Matter Surveys

Similar to using microlensing to infer the mass of exoplanetary bodies, galaxies and galaxy clusters can act as gravitational lenses. If the alignment between the lensing and target mass is good, then several highly magnified images can be observed, an effect called Strong gravitational lensing. Conversely, if the alignment is poor, weak gravitational lensing occurs instead, where a single, less distorted image is produced.

5.3: Gravitational Lensing

High image resolution is essential for both weak and strong lensing. Weak lensing requires the measurement of tiny distortions in the lensed object. Whereas for strong lensing, the details of distorted images are much easier to observe with higher image resolution. The example given by (Mackay, 2019) suggests that weak lensing distortions can be as small as 0.1 arcseconds. For an average ground-based observation this is difficult to achieve and means that the only targets which are available to these observations are nearby bright galaxies where the impact of weak lensing is more visible due to less obscuring material in-between the observer and target. Because GravityCam aims to achieve better seeing, of up to 0.3 arcseconds, the amount of observable galaxies increases.

5.4: GravityCam Summary

Overall, it has been shown that GravityCam relies on having a good signal-to-noise ratio and high framerates. The ideal CIS detector should reach 25 frames per second, ideally without windowing, to achieve as much field-of-view fill factor as possible. Also, the readout noise must compete with the current sub-electron noise of EMCCDs, to have a good signal-to-noise ratio. Therefore, the CIS device might also need sub-electron readout noise to be a viable detector choice. Many upcoming sCMOS cameras report sub-electron noise, but it is unclear what circumstances the device may be operated under to achieve these values. For example, whether the device is operated with a rolling or global shutter and whether correlated doubled sampling is used.

Chapter 6: CCD Tests and Results

Currently, the cameras used to perform these tests are part of the Two-Colour Instrument (TCI) on the Danish Telescope at La Silla, Chile. The instrument uses two Andor iXon 897 cameras to observe visible and near-infrared light simultaneously. Each cameras uses a single Te2v CCD97 EMCCD with 512x512 active pixels in the image and storage areas, 12 dark reference pixel columns either side of the image and storage areas, leading to a total of 24 additional columns, and with a multiplication register length of 536 elements. At a chosen pixel scale of 0.09 arcseconds per pixel, the field-of-view is about 45 arcseconds^2. They have a maximum pixel readout rate of 10MHz, which enables them to provide 35 full frames per second. Using a thermoelectric cooling system, the EMCCDs are cooled to a temperature of approximately -85 °C. Additional information on the instrument and its setup can be found in (Skottfelt, 2015). The data collected during this chapter only uses the 512x512 image areas, not the full 512x536 array. Since this is the area that would be used during real observations.

6.1: Spurious Charge Measurements

This characterisation aims to discover how often spurious events occur and if there is any spatial dependence. By finding out the frequency of spurious events and where on the device they form, the next step is to estimate the likelihood of a spurious event impacting a lucky image, i.e., how likely it is that a pixel could be influenced by enough spurious charge to impact observation. For example, this could happen if spurious events frequently occur in a single pixel or a small area.

The tests collect 1000 dark images, called an image stack, for each gain setting.

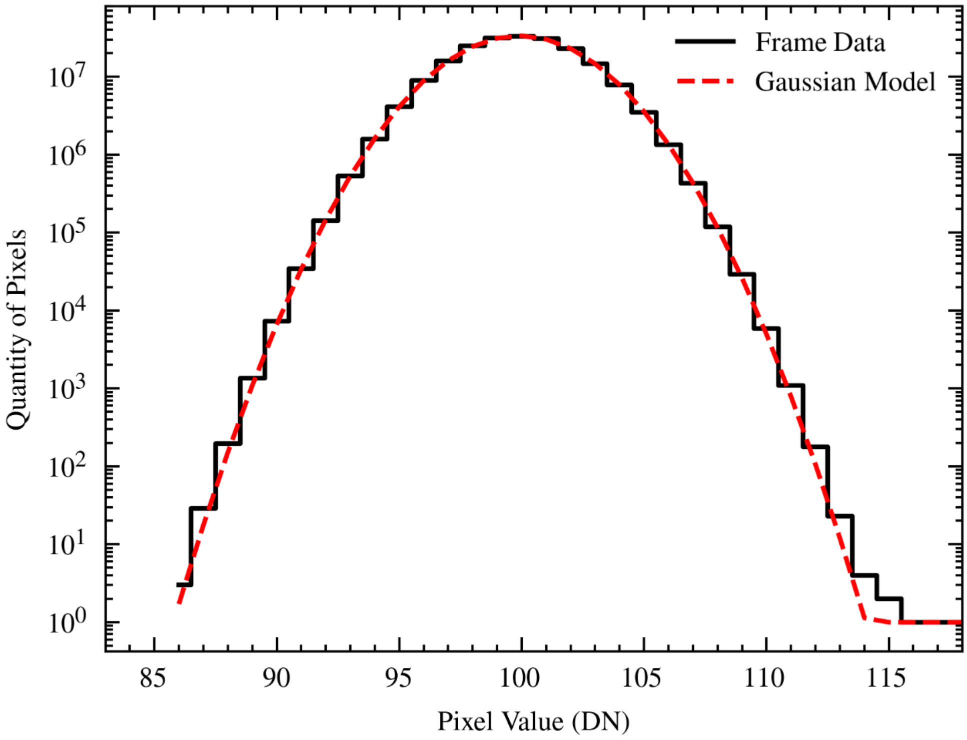

Figure 6-1 - A histogram of the dark signal from 1000 frames at EM gain 1 (i.e., no EM gain) for CCD97 and a Gaussian fit.

Figure 6-1 shows a histogram of the dark signal without EM gain and its Gaussian fit. The literature has widely recorded that readout noise on CCDs is distributed as a Gaussian. The dark signal in Figure 6-1 also appears to have a Gaussian distribution, so it can be assumed that the readout noise is the dominant contributing factor to this dark signal. There is little contribution from dark current, about 0.001 e-/pixel/sec (Skottfelt J. B., 2015), because it is suppressed at the low operating temperature. As expected, the noise distribution of CCD97 at no EM gain does not contain any notable CIC because multiplication events yield enough additional signal carriers very infrequently. Therefore, the noise distribution is similar to a traditional CCD sensor.

The noise distribution acquires a “tail” outside the Gaussian peak because the CIC contribution is amplified at higher gains. Pixels that contain significantly more CIC than other pixels now become distinguishable, as can be seen in Figure 6-2.

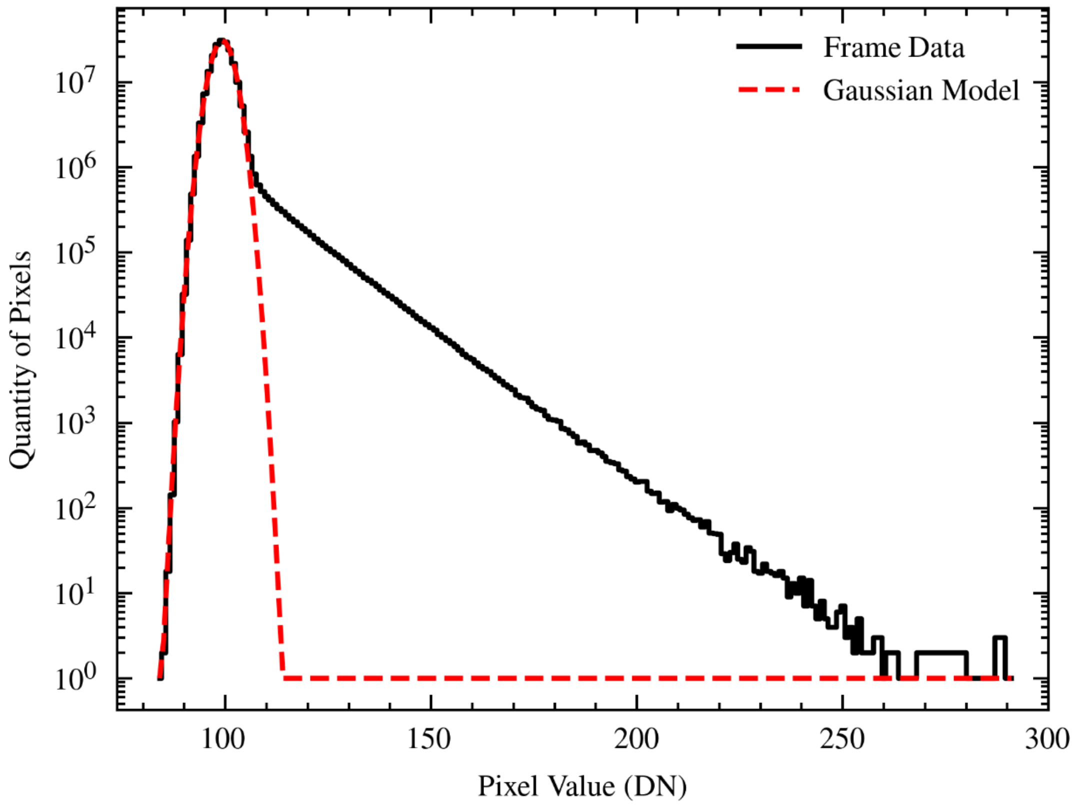

Figure 6-2 - A histogram of the dark signal in a CCD97 at gain 300 for 1000 frames and a Gaussian fit.

Because the y-axis of this plot is plotted on a logarithmic scale, it appears that the tail contains a significant signal, but it contains about 5% of the total amount of pixels, the rest being within the Gaussian peak. Here we define the tail as any pixels outside the Gaussian model. As previously stated, these tail pixels receive a notable signal contribution from multiplied CIC noise, which is why they appear outside the Gaussian fit. Not all pixels are impacted equally by CIC, and a small amount is heavily impacted. These heavily impacted pixels shall be referred to as high-value pixels. Isolating the pixels in the tail portion of Figure 6-2, the high-value pixels are chosen to be defined as the highest 10% of pixel values in this tail. This choice should allow the observation of the pixels that CIC most heavily impacts and determine where these pixels appear on the frame to find any spatial preference and how frequently the same pixel is flagged as a high-value pixel across the frame stack.

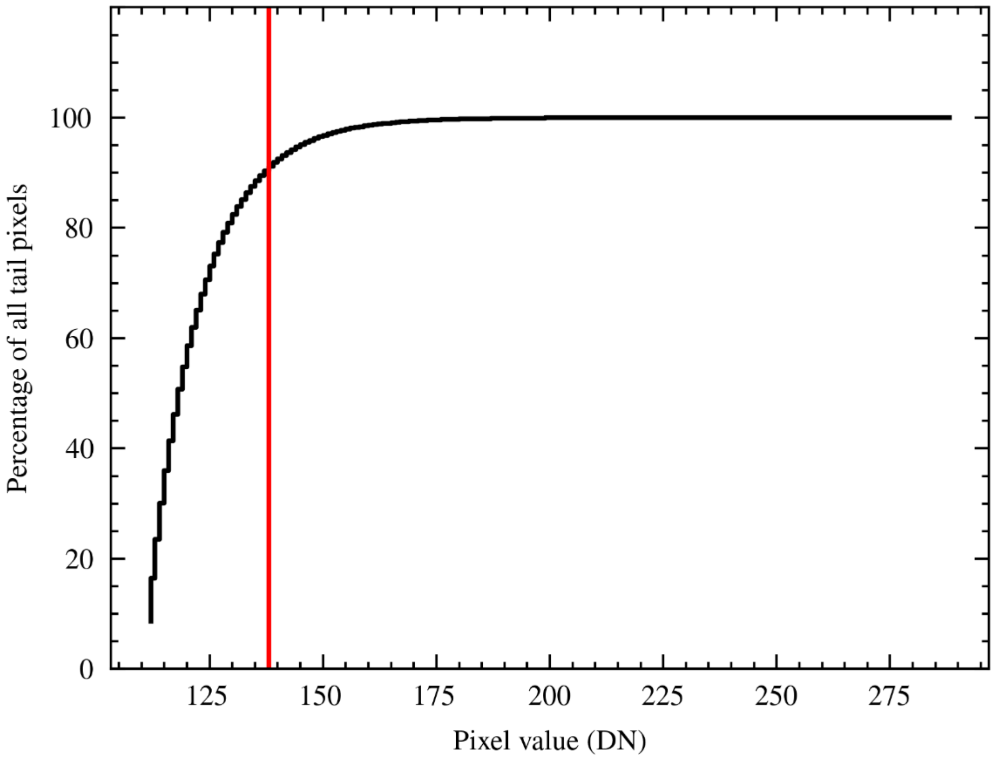

Figure 6-3 - A cumulative plot shows that pixel values below 138 DN contain 90% (red line) of all pixels.

Using Figure 6-3, the top 10% of these tail pixels are above the threshold of 138 DN, so any pixel in any frame above this value is flagged as a high-value tail pixel.

To see if there are any device-specific effects, i.e. a noise source caused by some defect on the device, two CCD97s were tested, and the results were compared. To distinguish them, the device with data in the prior figures, Figure 6-1, Figure 6-2 and Figure 6-3, will be referred to as the red camera, owing to the fact it is used to observe red wavelengths, and the second CCD97 will be referred to as the visual camera, as it is used to observe visual wavelengths.

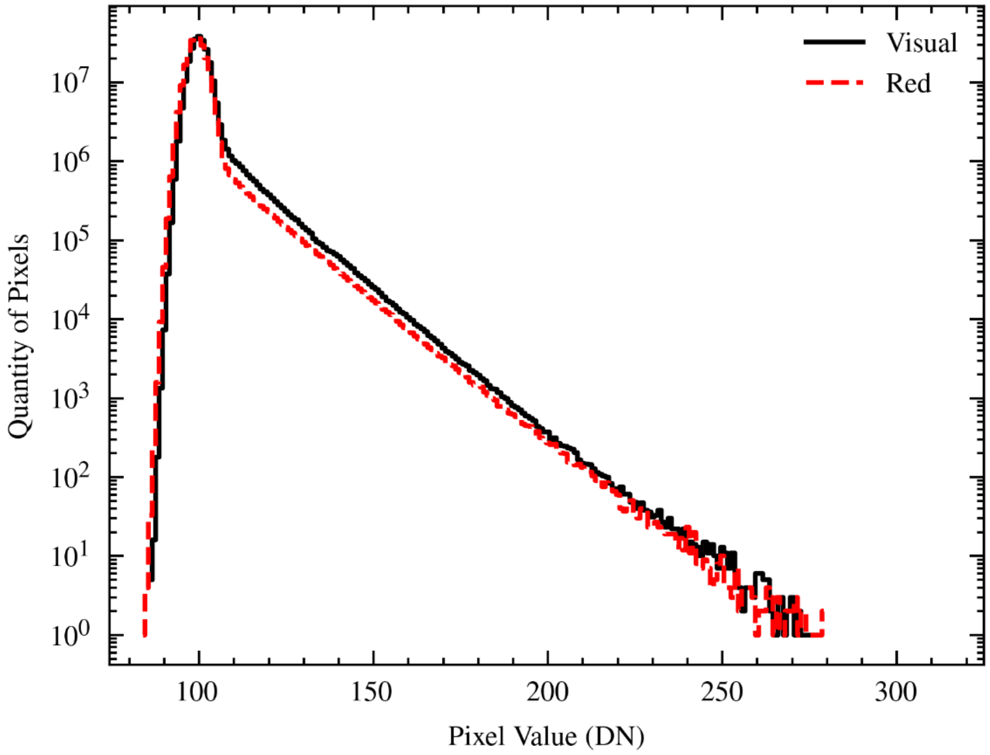

Figure 6-4- The signal distribution for the visual and the red cameras at a gain of 300, obtained from 1000 frames.

Figure 6-4 shows that both signal distributions have a similar shape. Table 3 shows a more detailed comparison between the two distributions.

(DN)	Red	Visual
Gaussian peak value	99	100
Max pixel value	291	277
Standard deviation	4.75	4.54

Table 3 - Comparison between Red and Visual camera.

The following section looks at how the quantity of tail pixels changes as the EM gain setting increases. Both devices will be examined for any differences or similarities.

6.2: Spurious Charge vs Gain

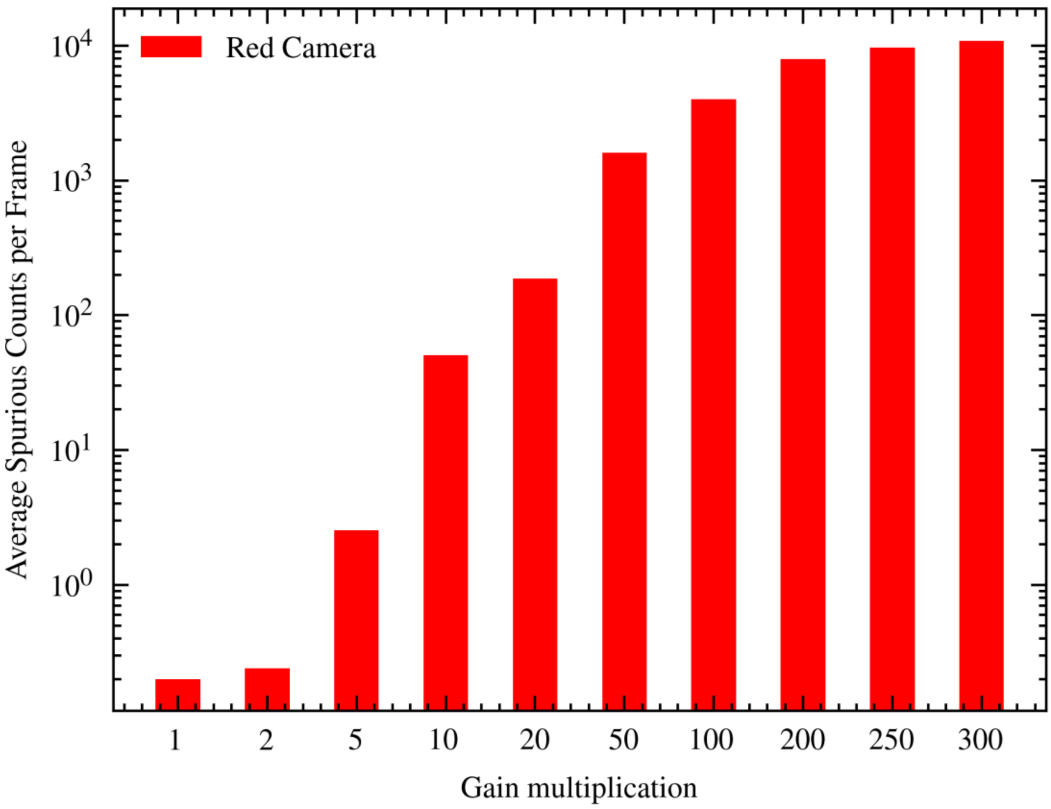

Figure 6-5 – Number of tail pixels per frame against increasing EM gain.

At gains below 20, we can see in Figure 6-5 that there are below one tail pixel, or 1/(512*512) % of the total 512x512, per frame, which indicates that spurious charge and other possible pre-multiplication noise sources do not contribute enough charge to have any significant impact on the overall signal at these gain levels, meaning that the distributions are well modelled with a Gaussian distribution. Above this, the fraction of tail pixels increases dramatically, up to 300 with a tail pixel count of 10923, or 4.16% of the total. The distribution's growth does not appear linear; instead, the increase in spurious counts slows down as the gain increases above 200. This slowing rate could suggest that there could be some degree of saturation for tail pixel counts; however, confirming this would require access to a much higher range of gain values. The issue with operating the devices at extreme gain values, 1000 or greater, increases the chance that the device would become damaged due to an EM gain ageing effect, which is previously mentioned in chapter 2.7.5: EMCCD Ageing.

Gain multiplication	Average tail pixels per frame	Increase in average tail pixels per frame quantity.	Percentage of frame taken up by tail pixels
1	0.01	--------	$3.8\text{x}10^{-6}$
2	0.01	0	$3.8\text{x}10^{-6}$
5	0.02	0.01	$7.6\text{x}10^{-6}$
10	0.04	0.02	$1.5\text{x}10^{-5}$
20	0.28	0.26	$1.0\text{x}10^{-4}$
50	45.9	45.62	$1.7\text{x}10^{-2}$
100	552	506.1	0.21
200	2336	1784	0.89
300	4226	1890	1.61

Table 4 – Number of tail pixels as a function of EM gain.

In Table 4, we can see that the increase rate of tail pixels rapidly increases at gains above 50. Finally, the number of tail pixels becomes substantially higher at an EM gain of 100, reaching more than four thousand pixels.

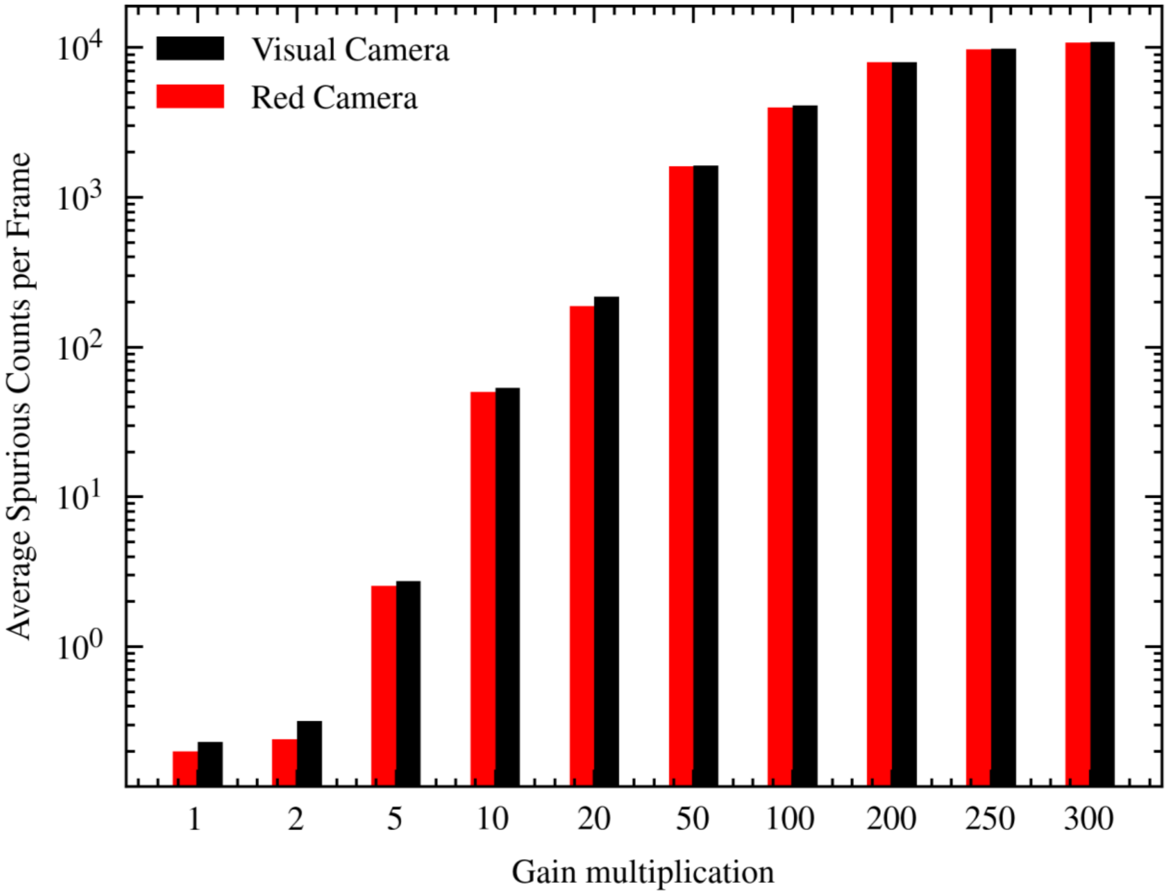

Figure 6-6 - A comparison between the red and visual cameras. This includes all the pixels within the tail.

In Figure 6-6, we can see that the cameras appear well matched, which suggests that both cameras do not have any significant anomalies. Both cameras have the same number of rows and columns polled to make the comparison fair. At EM gains of 100, 200, 250 and 300, the average tail pixels per frame account for as much as up to 1.5% of the entirety of the frame.

For tail pixels in the top 10%, the threshold value for a tail pixel being flagged as a high-value tail pixel increases as the gain multiplier, showing that these high-value tail pixels become more prominent as gain increases. This can be seen in Table 5, where for gains between 1 and 20, there is no change in the high-value pixel threshold. This lack of increase is likely due to these gains not significantly impacting the Gaussian distribution of the pixel values. Because these distributions lack any notable tail, there are no tails pixels and no actual high-value pixels.

Gain multiplication	High-value pixel threshold (DN)
1	113
2	113
5	113
10	113
20	113
50	116
100	120
200	129
300	138

Table 5 – High-value pixel threshold with multiplication gain.

Once at gain 50, there are enough tail pixels that the high-value pixel threshold deviates from the prior, lower gains and starts to increase, reaching 138 DN for a gain of 300. This increase shows that the pixel value of these high-value pixels is more significant than at lower gains, indicating that the overall pixel-value range has increased.

6.3: Column and Row Noise

Spurious events are expected to occur randomly and without any spatial preference. Several factors influence the likelihood of a spurious event happening in a pixel. One of the most significant factors is voltage clock characteristics, i.e., the voltage swing and the rise and the fall time. Looking at the spurious counts for each column and row, we can see whether there are any unusual patterns in the spatial distribution of the spurious charge generation. To investigate this, the following plots contain the amount of tail pixels counted in each row/column, the x axis variable, per pixel in that row/column, the y axis variable. This is to show whether there is a preference for tail pixels to appear in a particular row or column on the device.

For these devices, taking measurements at the full 512x512 active imaging area, there appears to be edge effects that can be seen to impact the column data, as seen in Figure 6-7.

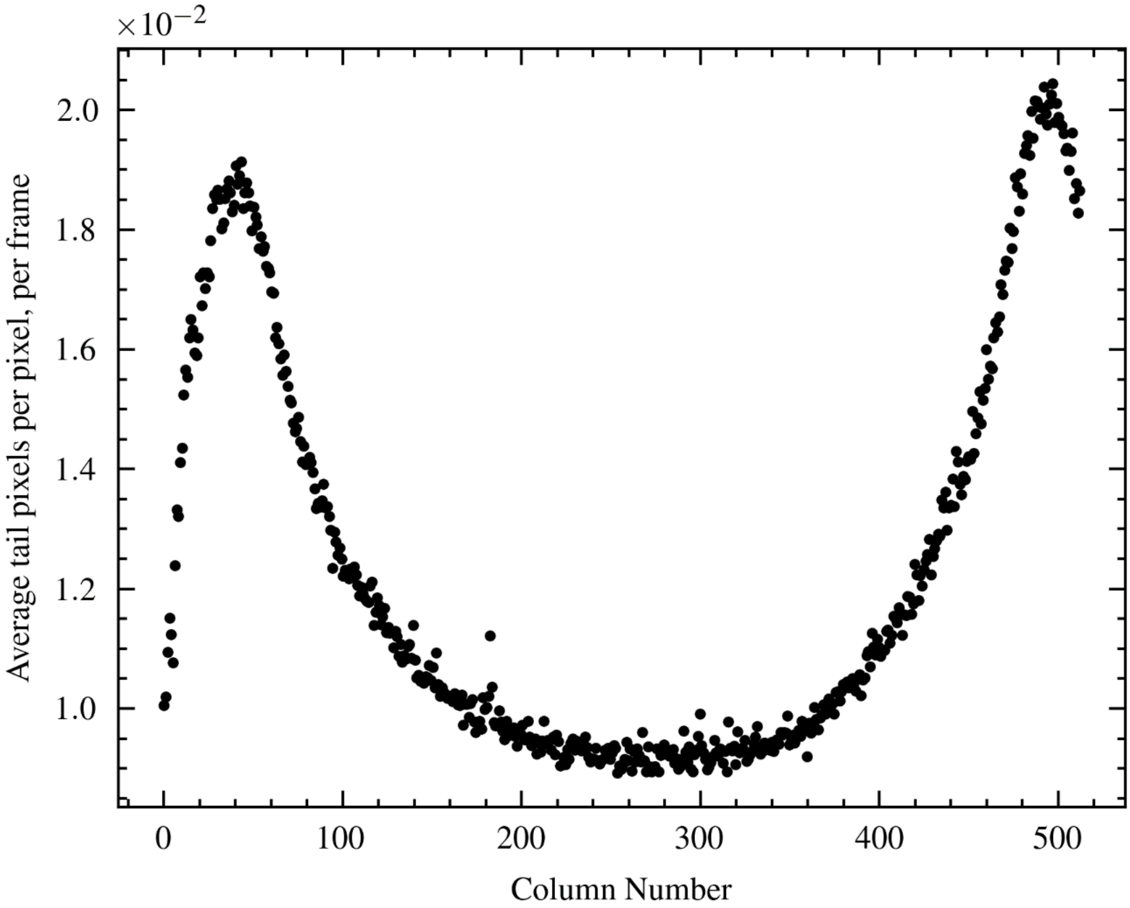

Figure 6-7 - Average tail pixel counts per pixel, per frame versus column number, for the red camera at gain 300. Taken over the full 512x512 image area. This is averaged over 1000 frames.

Both left and right side of the device show a notable decrease in tail pixels, this could be due to an edge effect reduces the pixel values, which would prevent them from appearing within the tail's distribution, preventing them being flagged as tail pixels. The causing mechanism for this is currently unknown, as vignetting, which can cause similar effects is reduced by taking flat frames, as discussed in 4.3.2: Unfortunately, due to the inability to access to the CCD itself, and check the clock voltage levels on these columns, determining the cause of this is not possible. The lower column numbers, show a much more dramatic decrease in tail pixel quantity, this can be attributed to a lower bias level across this part of the device. (Harpsøe, 2012) shows that in Figure 6-8 the bias level of these columns and the bottom rows, have a much lower bias level. This would also achieve the result of reducing the pixel values in these regions, moving them from the tail distribution, and therefore these pixels would not be flagged as tail pixels.

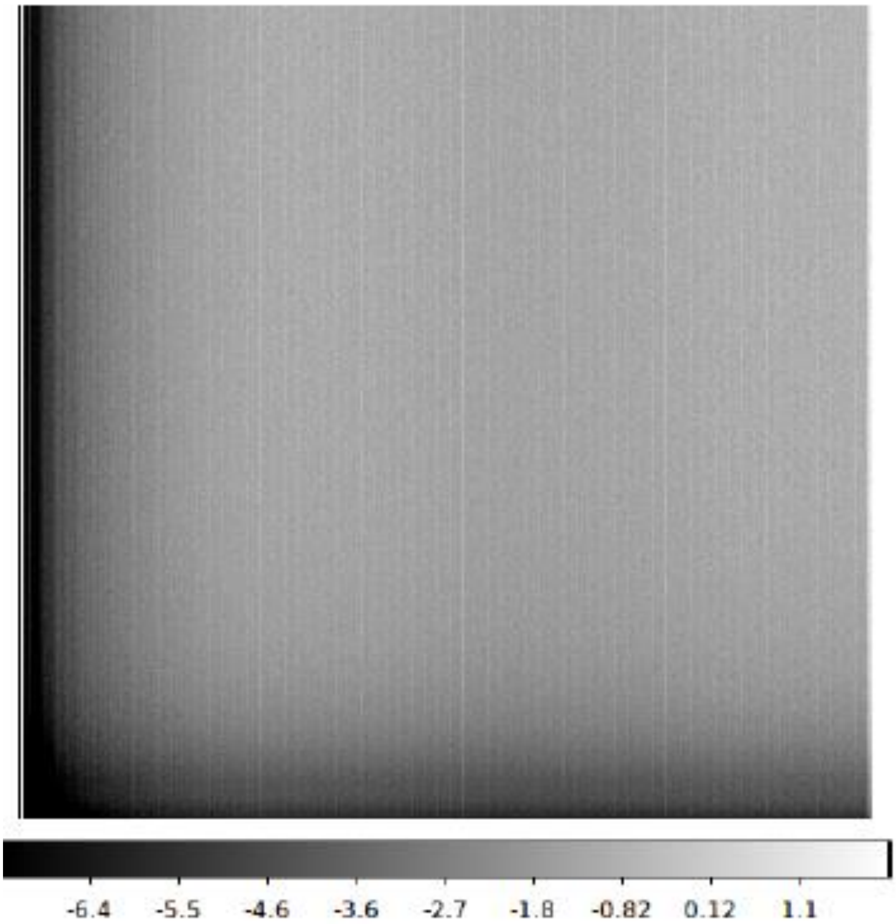

(a) Bias pattern

Figure 6-8 - The bias pattern of the Danish telescope's CCD97 red camera, taken from (Harpsøe, 2012).

To prevent these regions of the device from skewing the analysis of spatial preference for these devices, these left and right columns will be removed the analysis. The subsequent distributions will be drawn from column 50 and row 50, till column 500 and row 500, a 450x450 area of the device. This will allow a suitable study of the middle area of the device, where any spatial deviation will not be impacted by the bias level and any erroneous edge effects.

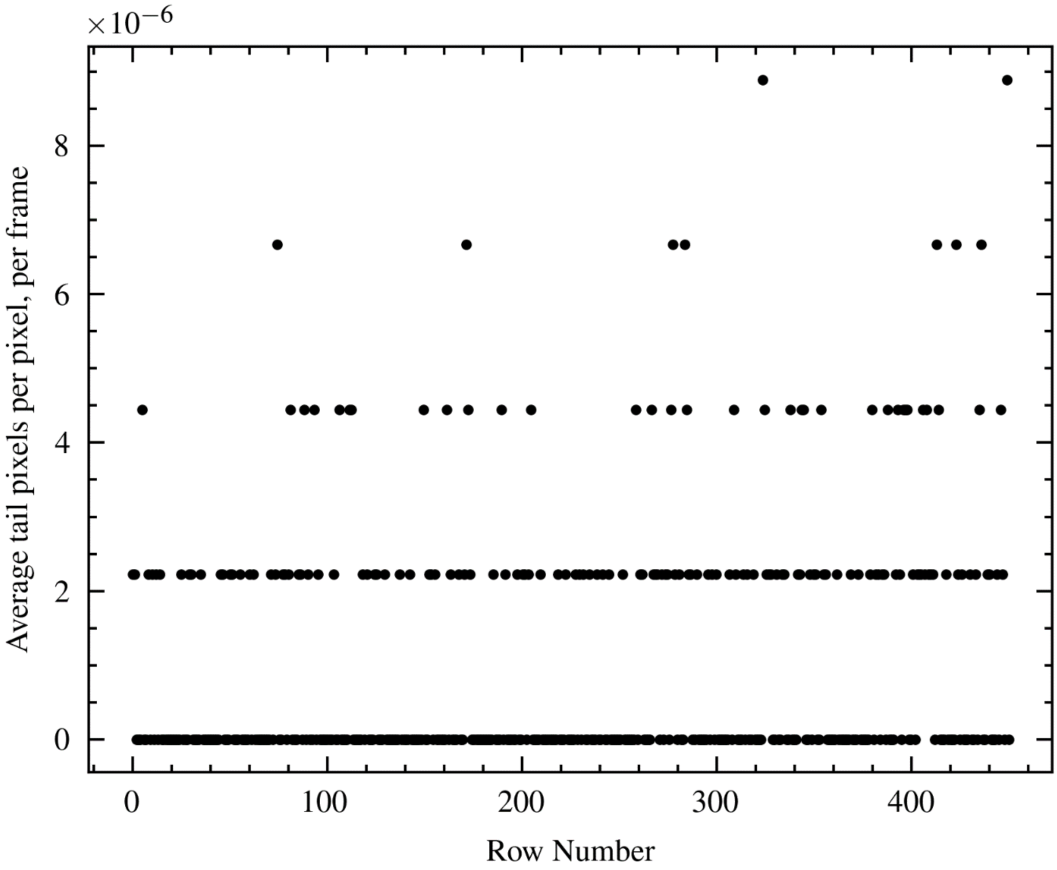

Figure 6-9 – Average tail pixel quantity per frame versus row number for the red camera at gain 1. 1000 frames are being investigated. This has been windowed to 450x450 to remove any contribution from edge effects.

Figure 6-9 shows that most rows contain an average of zero tail pixels per frame throughout the 1000-frame stack. The discreet values are resultant from these rows containing either zero tail pixels, or one to four tail pixels over the entire 1000 frame stack. The mean row contains 1.1x10^{-6} tail pixels per pixel per frame. The standard deviation is 2.1x10^{-6} tail pixels per pixel, per frame rms. At 3σ, which is about 7.4 x10^{-6} tail pixels per frame, we expect to see at most 1.35 rows. The two outliers at 8.8 x10^{-6} are more likely to be noise than the result of some external factor, as they are between 3σ and 4σ thresholds.

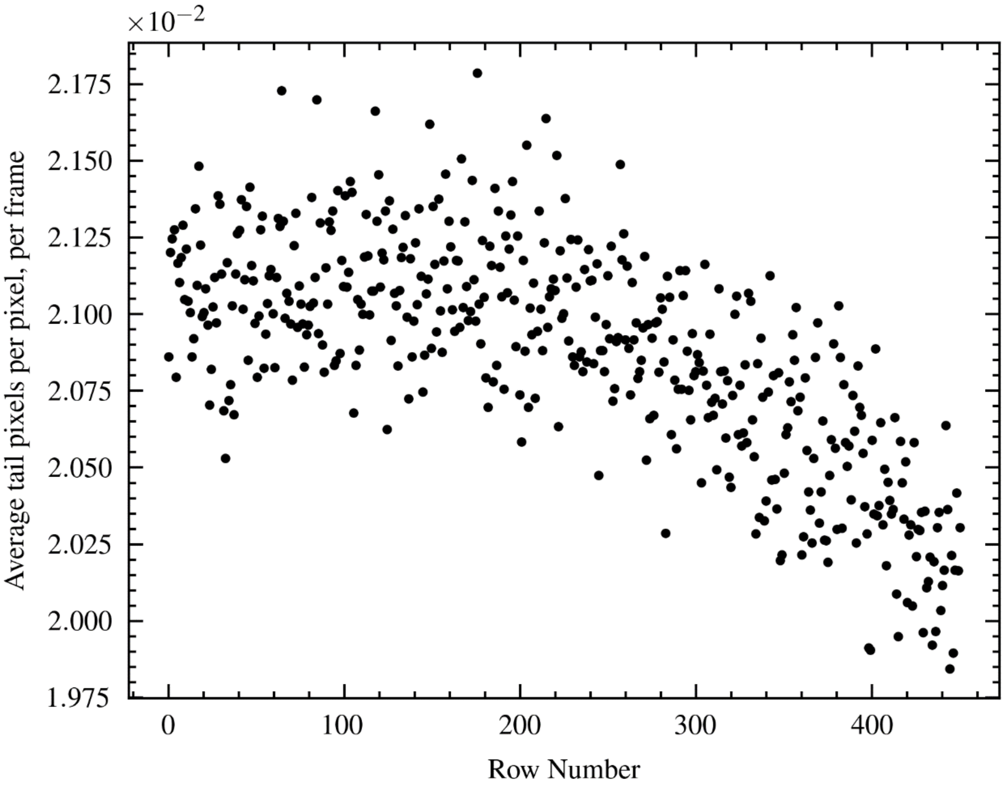

Figure 6-10 – Similarly to Figure 6-9, the average amount of tail pixels per pixel in each row, per frame is plotted but instead of gain1, gain 300 is used.

Figure 6-10 shows a very slight preference for more tail pixels detected towards row zero, which in this case is the row furthest from the serial register. This is because charge packets originating from the serial register must travel longer distances through the device than other charge packets. This means that these charge packets have more opportunities to accrue CIC and are more likely to be tail pixels.

The mean row value in Figure 6-10 is 2.08 x10^{-2} tail pixels per pixel per frame and a standard deviation of 3.63 x10^{-4} tail pixels rms per pixel per frame. The 3σ threshold is 2.18 x10^{-2} tail pixels per pixel per frame, which no row goes above, and 1.97 x10^{-2} tail pixels per pixel per frame, which no row goes below.

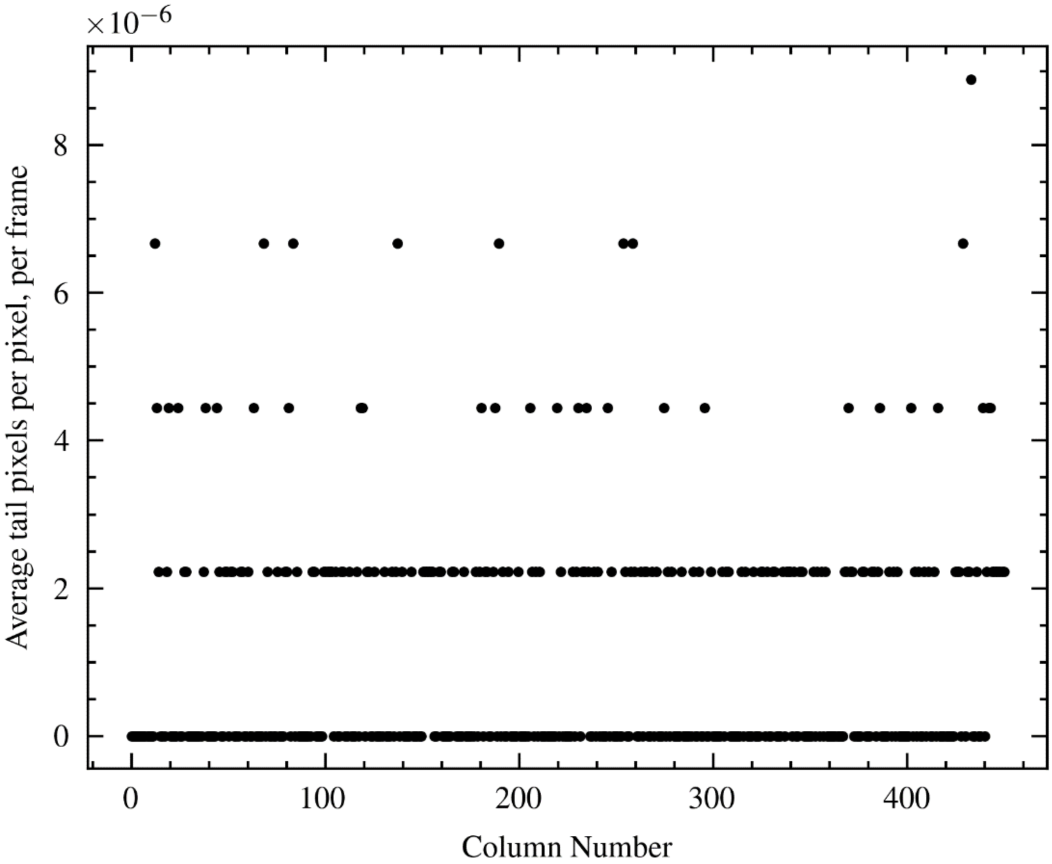

Figure 6-11 – Average tail pixel counts per pixel, per frame versus column number, for the red camera at gain 1.

Figure 6-11 shows us the column average tail pixel distribution, many columns see an average of zero tail pixels per frame, and the mean frame has 1.05x10^{-6} tail pixels per pixel per frame. Similarly, to Figure 6-9, the discreet values are resultant from these columns containing either zero tail pixels, or one to four tail pixels over the entire 1000 frame stack. The standard deviation is 2.0x10^{-6}. There is a single outlier column, which has an average of 9x10^{-6} tail pixels per frame, the only column above the 3σ threshold of 7.05x10^{-6} tail pixels per frame, and therefore is more likely to be a statistical anomaly than the result of external influence.

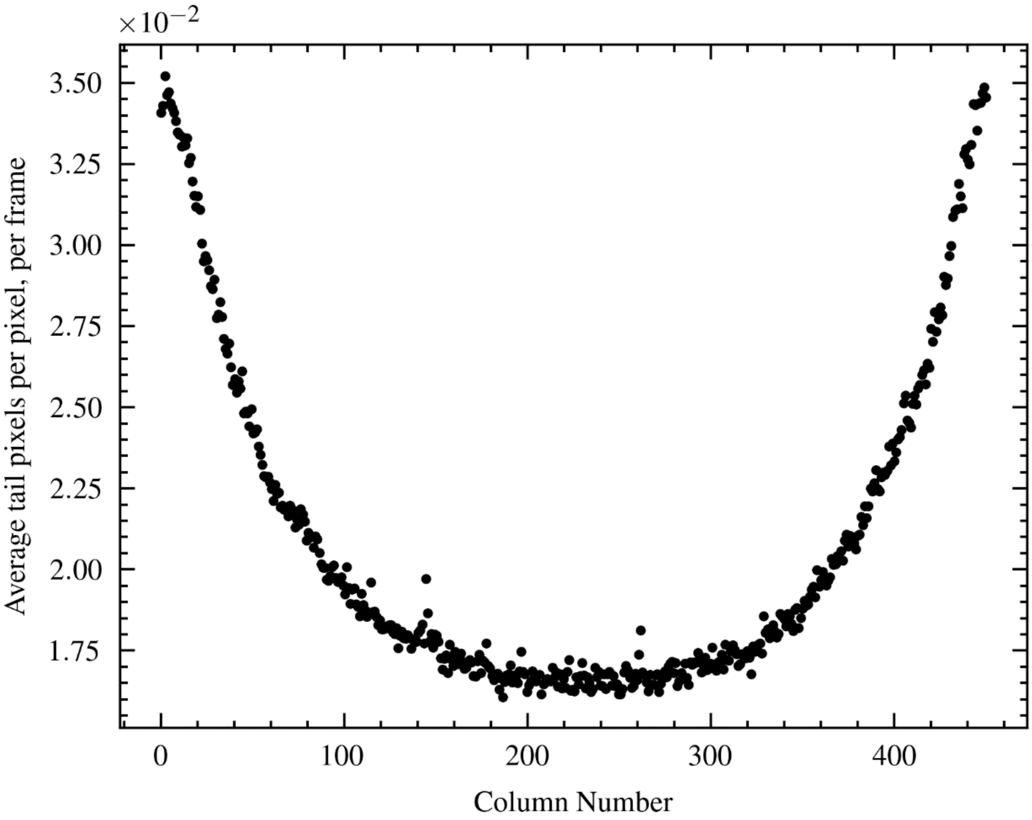

Figure 6-12 – Similarly to both Figure 6-10 and Figure 6-9, this figure effectively a gain300 version of Figure 6-11, to show the impact of increasing the device’s gain.

In Figure 6-12, there is an apparent disparity between the left and right sides of the device and the central columns, columns closer to the sides of the device show twice as many average tail pixels per frame compared to columns in the middle of the device. Figure 6-13 shows us the device architecture.

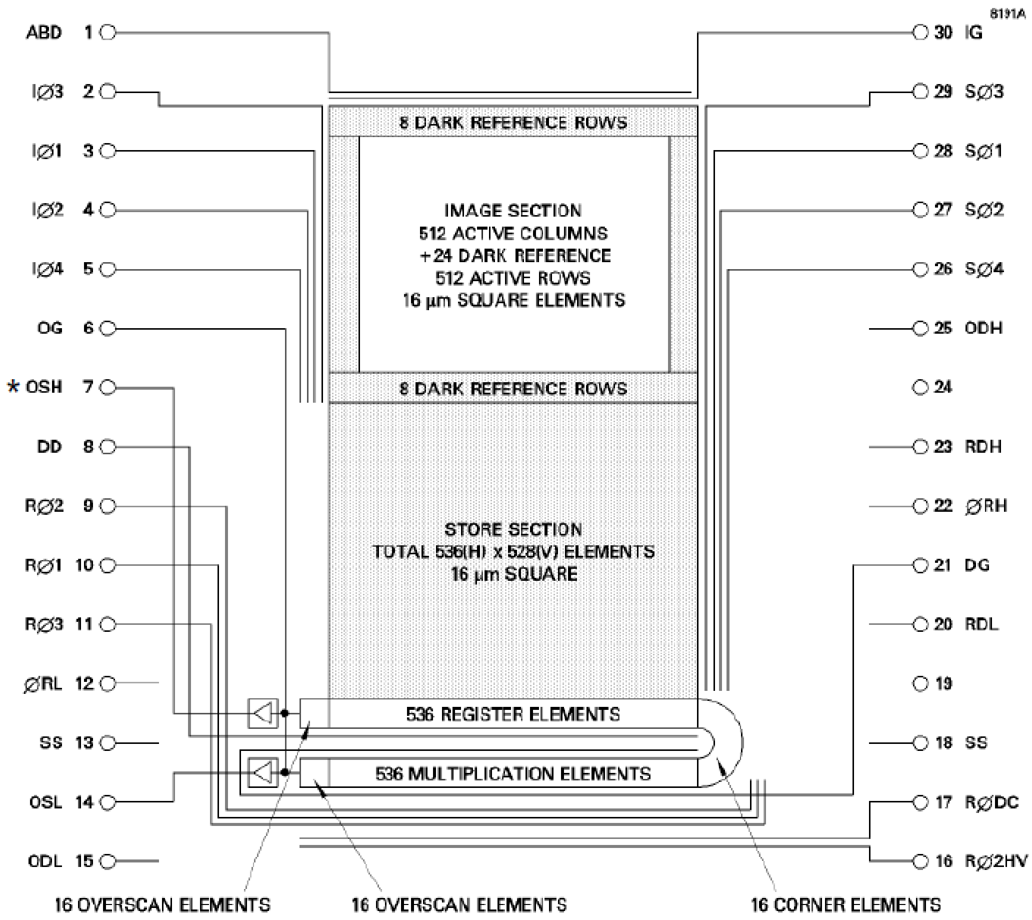
8191A
ABD 1
IØ3 2
IØ1 3
IØ2 4
IØ4 5
OG 6
* OSH 7
DD 8
RØ2 9
RØ1 10
RØ3 11
ØRL 12
SS 13
OSL 14
ODL 15
30 IG
29 SØ3
28 SØ1
27 SØ2
26 SØ4
25 ODH
24
23 RDH
22 ØRH
21 DG
20 RDL
19
18 SS
17 RØDC
16 RØ2HV
8 DARK REFERENCE ROWS
IMAGE SECTION
512 ACTIVE COLUMNS
+24 DARK REFERENCE
512 ACTIVE ROWS
16 µm SQUARE ELEMENTS
8 DARK REFERENCE ROWS
STORE SECTION
TOTAL 536(H) x 528(V) ELEMENTS
16 µm SQUARE
536 REGISTER ELEMENTS
536 MULTIPLICATION ELEMENTS
16 OVERSCAN ELEMENTS
16 OVERSCAN ELEMENTS
16 CORNER ELEMENTS

Figure 6-13 - The device architecture of the backside-illuminated CCD97, taken from (Teledyne e2v, 2021).

We can see that the positions for the image section clocks are on pins 2 to 5 and notably on the left side of the device. Similarly, the store section clocks are on pins 26 to 29 and positioned on the device s right side. Recalling that (Janesick J. , 2001) states that rising clock edges are a prominent source of spurious charge generators, the phenomena in Figure 6-12 could be attributed to this. As the CCD97 moves charge packets through the image and store sections, the clock voltage waveform starts with a sharp rising edge, producing a high spurious charge. This sharp rising edge flattens out further through the device due to the resistance along the polysilicon gates and as the rise time decreases, the less and less spurious charge is generated, creating a sloping spurious charge population. This means that the spurious charge generation is highest at the column positions nearest to the pins where the voltage waveform enters, and varies greatly due to the rising edge, and reduces towards the lowest values on the other side of the device, where the clock waveform is much more stable. Because during readout, the charge packets move through both the image and storage area, this initially high rate of spurious charge generation and then decreasing across the device effectively happens twice - once beginning on the left-hand side of the device and then once again on the right-hand side. This pattern has been noted in other papers, such as (Harpsøe, 2012). The same effect can be seen in the visual camera in Figure 6-14.

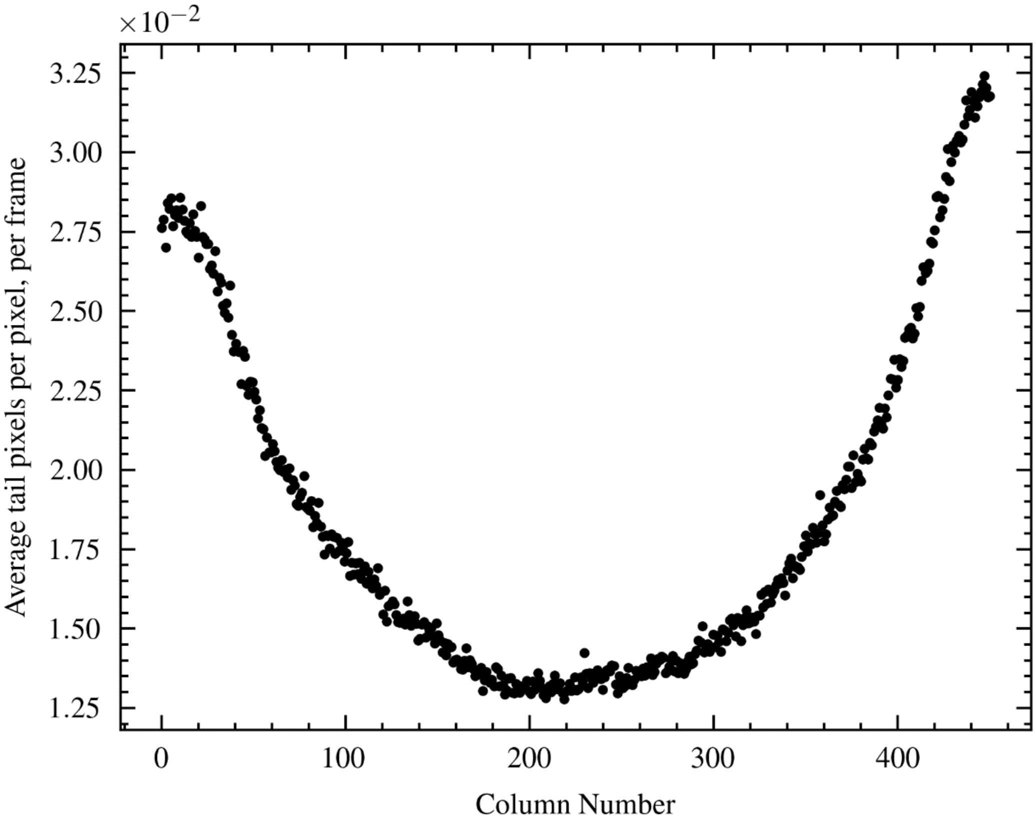

Figure 6-14 - Average tail pixel counts per pixel, per frame versus column number, for the visual camera at gain 300.

Additionally, if we look only at high-value tail pixels, this column preference is still observable in both cameras in Figure 6-15.

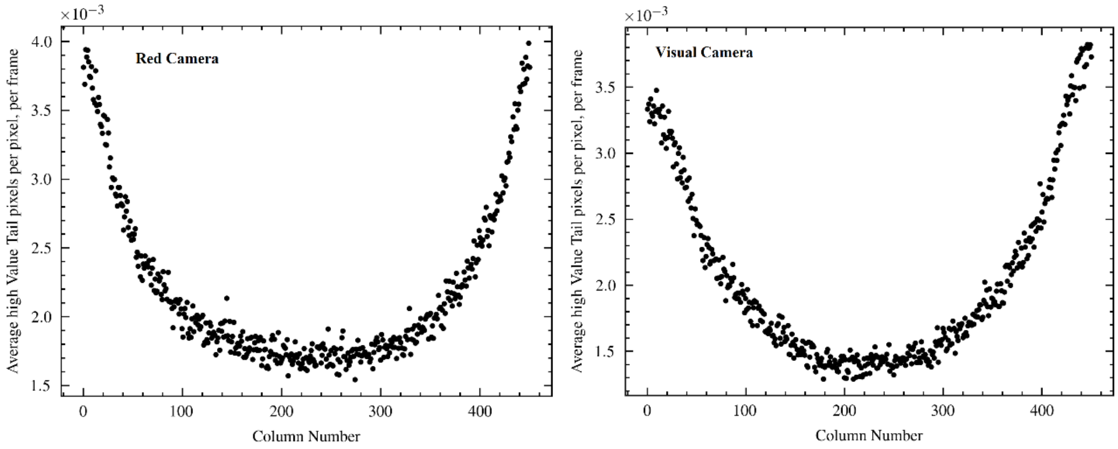

Figure 6-15- Right shows average high-value tail pixel counts per pixel, per frame versus column number for the red camera at gain 300, left shows the same for the visual camera.

Whilst there is no preference for tail pixels, or high-value tail pixels, to appear in any row at any gain, there is a clear spatial column preference at a high gain in both cameras. As previously mentioned, this is suspected to be a result of the image and storage voltage clocks used to drive each pixel’s gates. Therefore, to reduce the CIC, voltage clock optimisation of the rising and the falling edges can be done, as discussed and shown in both (Daigle O. T., 2018) and (Daigle O. Q., 2010).

6.4: Potential Tail Impact Values

Given that it has been shown that multiplication gain and image and storage clocks impact CIC and, therefore, whether a pixel is a tail pixel or not, the potential value of an average tail pixel can be estimated. Using this estimation, a pixel frequently flagged as a tail pixel can be compared to the mean pixel value, and a pixel infrequently flagged.

Assuming that the red camera device is operating at a high gain of 300, the worst column in Figure 6-12 is column 2, having the highest amount of $3.53\text{x}10^{-2}$ tail pixels per pixel per frame. Looking in this column, we can choose the pixel flagged most over the 1000 frames. This ‘worst pixel’ has a flagged amount of 53 times. Assuming that we can use the median tail value as a reasonable estimate for the average tail pixel value and that any time that the pixel is not flagged as a tail pixel, it has the average value of a pixel within the gaussian limits, we can estimate the impact in average pixel signal over the 1000 frames. By multiplying the number of times, a pixel is flagged as a tail pixel, T_F, by the median tail value for the device, M, then adding the times the pixel was not flagged, T_N, multiplied by the average gaussian pixel value, A, gives the total amount of signal, S_T, accumulated by the pixel over all frames. Dividing this value by the total amount of frames, F, we can find the average pixel value for the worst pixel, the most flagged pixel, and the best pixel, the least flagged pixel, as shown in (37).

$$S_T = [(T_F * M) + (T_N * A)] \tag{37}$$

The worst pixel has an average value of 103.41 DN per frame; conversely, in this case, the best pixel has an average value of 101.63 DN per frame. So, for a gain of 300, on the red camera, it appears that the maximum average impact from CIC is 2.22 DN or 0.244 e- per frame.

Whilst this is not a significant signal difference, as discussed in Chapter 8:, CMOS noise distributions with mean values as low as 0.1 e- impacts these faint objects.

Chapter 7: CMOS Image Sensor Tests and Results

7.1: Introduction to the CIS115 Sensor

The Sirius CIS115 device was developed by Teledyne e2v (Te2v) for the European Space Agency's Jupiter Icy Moon Explorer (JUICE) mission, currently planned to launch in 2023. The sensor is part of the optical camera scientific payload to image and map Jupiter, the Jovian moons, and icy rings. The device's characteristics are detailed in the device datasheet available on Te2v's web page (Teledyne e2v, 2021) and other sources (Soman, 2014) and (Lofthouse-Smith, 2018).

The device format is 2000 rows by 1504 columns of 7 μm square pixels. The total imaging area is 14 x 10.528 mm^2. The pixel array is split into four identical blocks with 376 columns. Frame readout times depend on how many rows are read out, so higher frame rates are achievable if a smaller region of interest is addressed.

At a temperature of 293 K, the device specifications are:

- Mean readout noise = 5 e$^-$ rms.
- Dark current = 20 e$^-$/pixel/s.
- The maximum frame rate for the whole device readout = 7.5 Hz.

7.2: Why the CIS115 and the aims of this chapter

Being one of the latest devices available from Te2v, the CIS115 is an appropriate choice of an image sensor to establish how suitable CMOS imaging devices are for high-speed imaging techniques, specifically Lucky Imaging. Due to device availability, the tests were performed with a front-side illuminated CIS115 device.

The tests presented in this chapter aim to show the dominant noise sources, the magnitude of the noise, and frame stability. Given that lucky imaging requires the device to take large quantities of images at a high frame rate, it is crucial to evaluate how well the CIS115 performs under these circumstances.

Some of the main requirements for lucky imaging are to take images at a high cadence, usually above 10 Hz, as stated in (Mackay, 2018). It is also essential that noise sources, i.e. readout noise, do not significantly reduce the signal-to-noise ratio (SNR). Readout noise can obscure faint astronomical targets, reducing useable data or ultimately preventing any measurement. For example, having too much noise from the sensor could prevent GravityCam from detecting low-mass objects via gravitational microlensing. The following tests show how the device performs at high frame rates, what dominant noise sources are present and how stable the frames' noise is.

7.3: CIS115 Tests

7.3.1: Dark Current

A dark current is a signal that can be measured when there is no light on the device. Thermal excitations cause the generation of electron-hole pairs within the device's imaging region. The output circuit receives the electrons and reads them off as a signal. The dark current decreases exponentially with temperature; therefore, cooling the device can reduce the dark current significantly.

The CIS115 device tested in this thesis was operated at room temperature, as there was no available method for cooling. First, the dark current is measured by taking sequential dark images, i.e. images with no incident light, and increasing the integration time linearly. Then, the measured signal is plotted against the integration time, giving the relation between the measured signal in electrons and the time used for their collection. This relation is the dark current at a given temperature, usually given in units of electrons per second (e^{-}/s). Figure 7-1 shows an example of such a measurement.

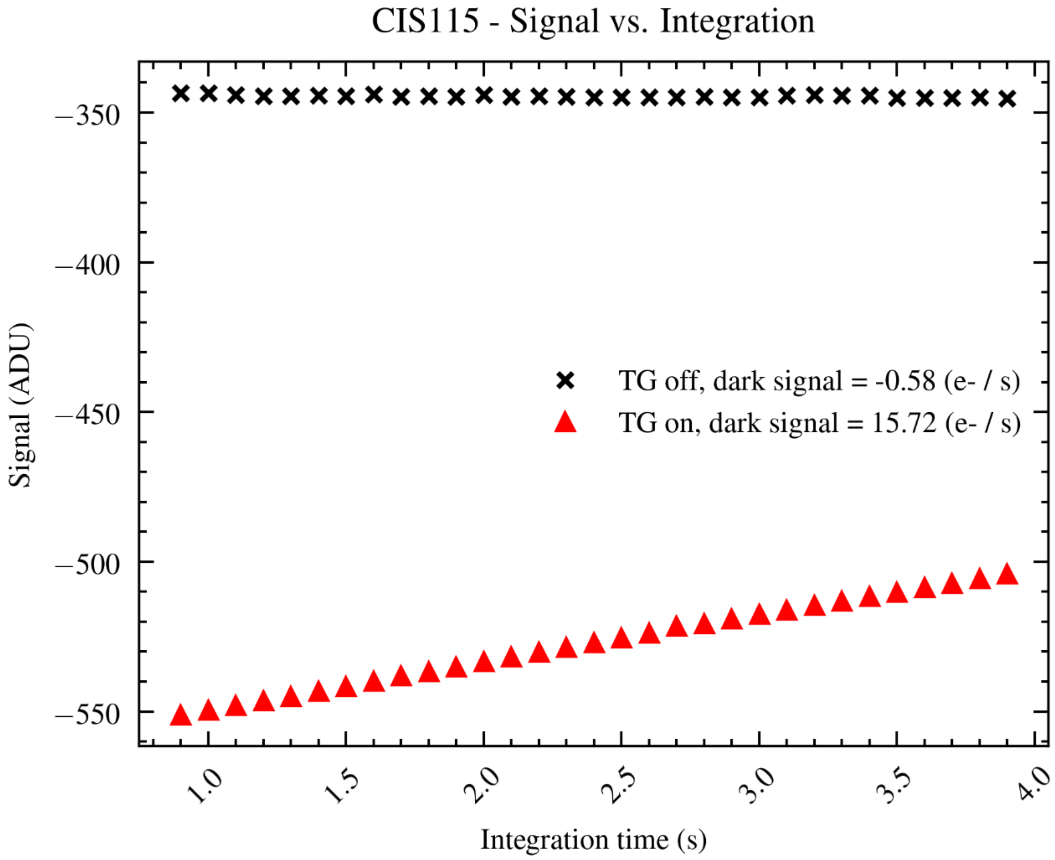

Figure 7-1 - Dark signal with the Transfer Gate (TG) off and on. The dark signal has been removed by keeping the TG low.

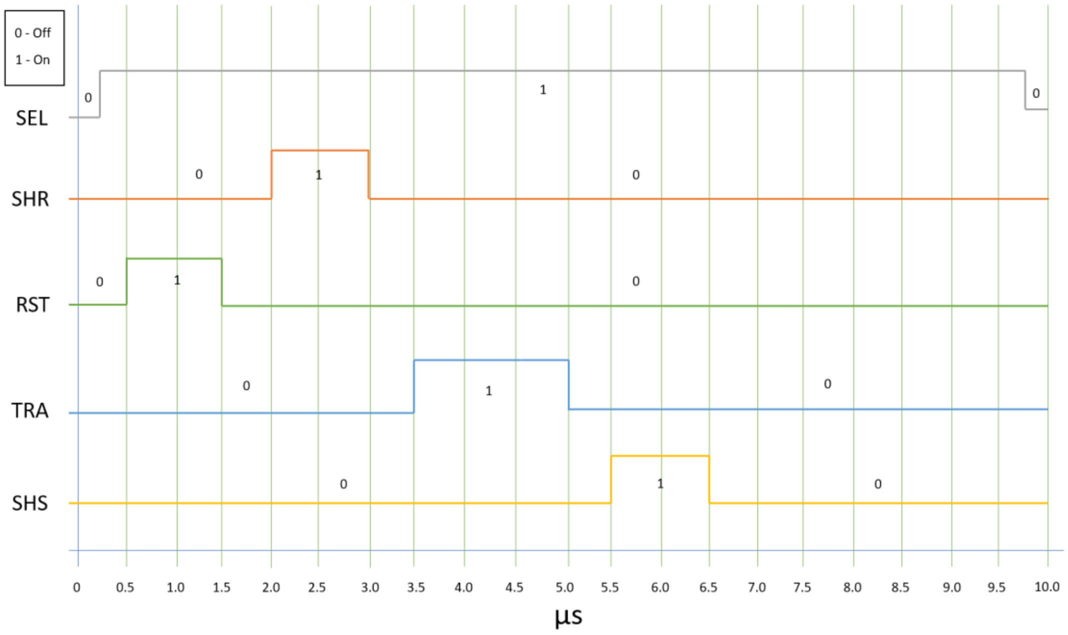
0 - Off
1 - On
SEL
SHR
RST
TRA
SHS
0
1
0
1
0
1
0
0
1
0
0
1
0
0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
4.0
4.5
5.0
5.5
6.0
6.5
7.0
7.5
8.0
8.5
9.0
9.5
10.0
μs

Figure 7-2 - The pixel timings for a standard readout.

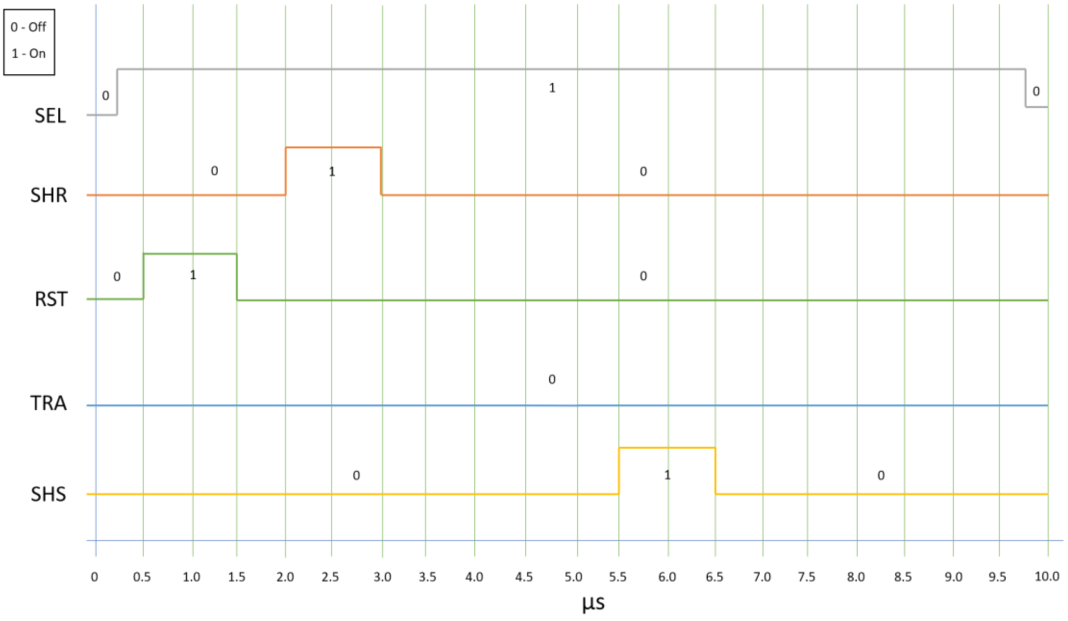
0 - Off
1 - On
SEL
SHR
RST
TRA
SHS
0
1
0
1
0
1
0
0
0
1
0
0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
4.0
4.5
5.0
5.5
6.0
6.5
7.0
7.5
8.0
8.5
9.0
9.5
10.0
μs

Figure 7-3 - Modified pixel timings for effectively removing dark current from the measurements. This timing is used for the readout noise measurements and Figure 7-1(a).

Figure 7-2 shows the device's timing for standard imaging. This timing scheme will result in a standard device readout, including a dark current in the signal output. The dark current of this device is measured to be 15.72 e^{-}/s at room temperature. Therefore, the device could be cooled to reduce the dark current, which could be required if the integration time is extended, i.e. greater than 1 s. However, the dark signal could be reduced by operating the camera at a short integration time; this will also reduce the dark signal for Lucky Imaging.

Figure 7-3 shows the modified pixel timing where TRA is kept off, preventing charge transfer to the sense node. Thus, the output signal contains only readout noise because there is no charge from the PPD. This technique makes studying the device's readout noise considerably easier without worrying about any signal coming from a dark current.

7.3.2: Stability

Knowing that Lucky Imaging requires the device to be operated at short integration times below 100 ms, it is crucial to understand how the device behaves when operating at this speed and if there is any notable impact on the signal that occurs frame to frame.

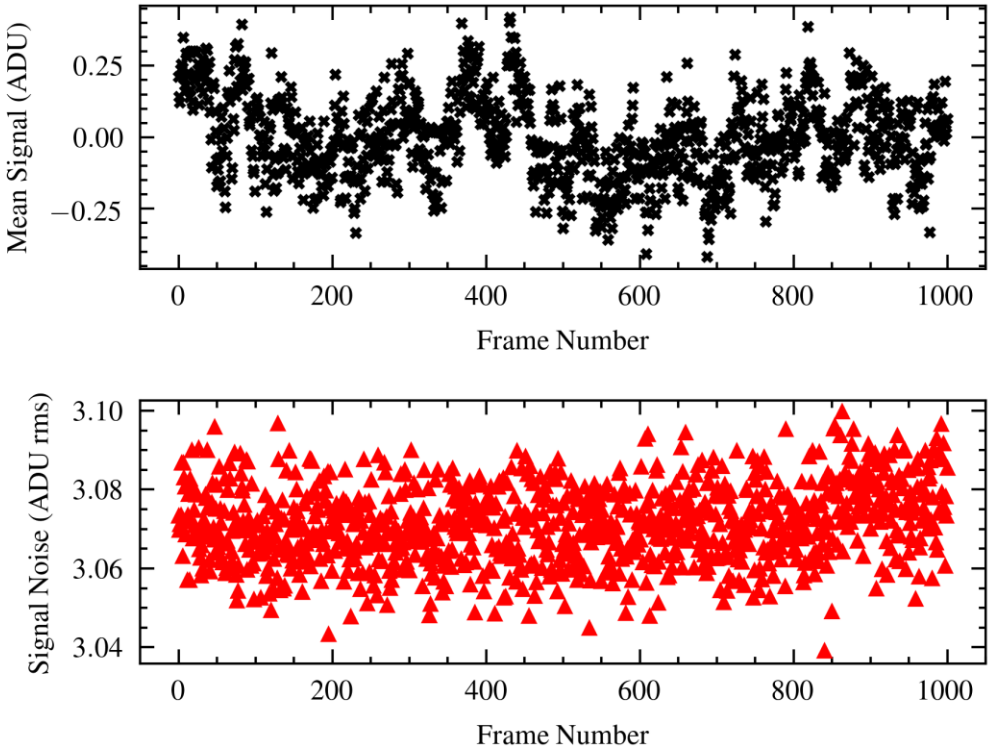

Figure 7-4 - Mean, DSNU-corrected signal from a stack of 1000 images (top); spread of the standard deviations across the stack (bottom).

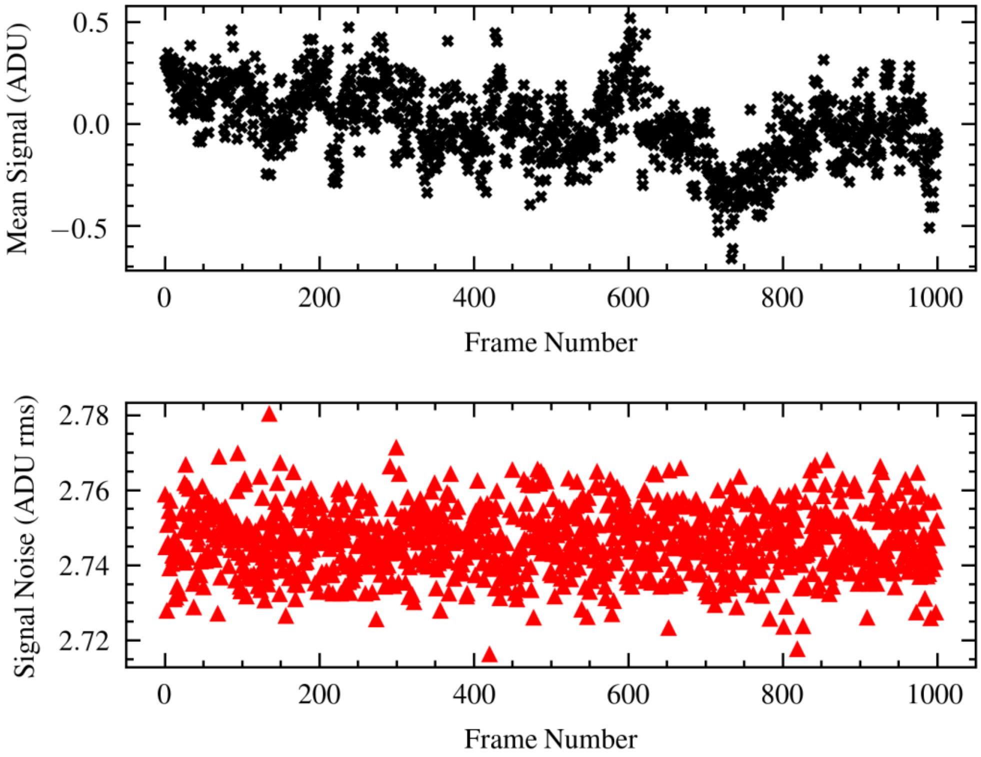

Figure 7-5 - The same plot as Figure 7-4, except the dark current, has been removed with the previously mentioned clocking modification.

Figure 7-4 and Figure 7-5 show the mean and standard deviations for 1000 frames (each with 202x376 pixels) taken at 82 ms (>10 Hz) integration time. This pixel window size allowed the device to run at this speed. Each frame has had a mean image subtracted from them to remove the fixed pattern noise. This subtraction frame was acquired by averaging the values of each pixel across the 1000 frame stack. Figure 7-4 shows this with a stack of 1000 frames with a dark current, and Figure 7-5 is without a dark current.

	Figure 7-5	**Figure 7-4**
Mean of the mean signal (ADU)	≈ 0	≈ 0
Standard deviation of the mean signal (ADU rms)	0.14	0.18
Mean of the signal noise (ADU rms)	3.07	2.74
Standard deviation of the signal noise (ADU rms)	9.75x10^{-3}	8.67x10^{-3}

Table 6 - Mean and standard deviations for Figure 7-4 and Figure 7-5.

Table 6 shows that the mean of the mean signal for Figure 7-5 is extremely close to zero, with a standard deviation of the mean signal being 0.14 ADU rms. Therefore, with a standard deviation of 0.14 ADU rms, we expect to see less than 1% of frames to be beyond a mean value of ±0.42 ADU, a 3σ limit.

The most significant difference between the two datasets is that the dark current dataset has a higher frame signal noise by 0.33 ADU rms, so there appears to be more significant variance in frame output with the dark current than without dark current. However, the standard deviation of frame signal noise values appears similar for both datasets.

Overall, the main impact of dark current on high framerate dark images is an increase in the variance of the pixel output, but the mean frame value appears only minorly impacted.

7.3.3: Readout Noise

Readout noise is added to the signal during the charge to voltage conversion and amplification. Unlike CCDs, each pixel's readout noise differs from other pixels because they are separate readout circuits. As a result, individual pixel distributions impact the device's global read noise distribution, which appears as a log-normal distribution instead of the normal distribution seen in traditional CCDs. This difference is primarily owed to a small number of pixels with much higher read noise than the majority, extending the distribution with a “tail”. This log-normal distribution can be seen in Figure 7-6.

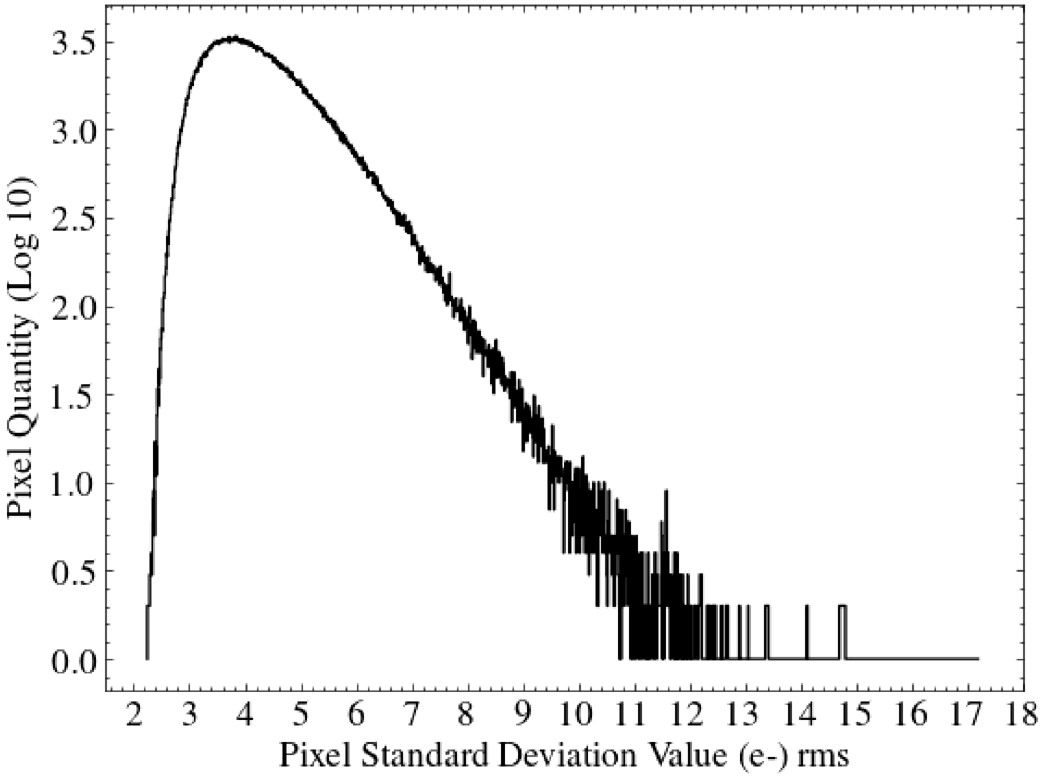

Figure 7-6 - Readout noise histogram of CIS115.

A distribution like this shows that some pixels have a much higher standard deviation than the average. These pixels are referred to as noisy pixels. Though they make up less than 10% of the pixels on the image area, it is essential to know where they are and how they behave with time to mitigate their impact on low-signal observations, i.e. high speed in low light levels.

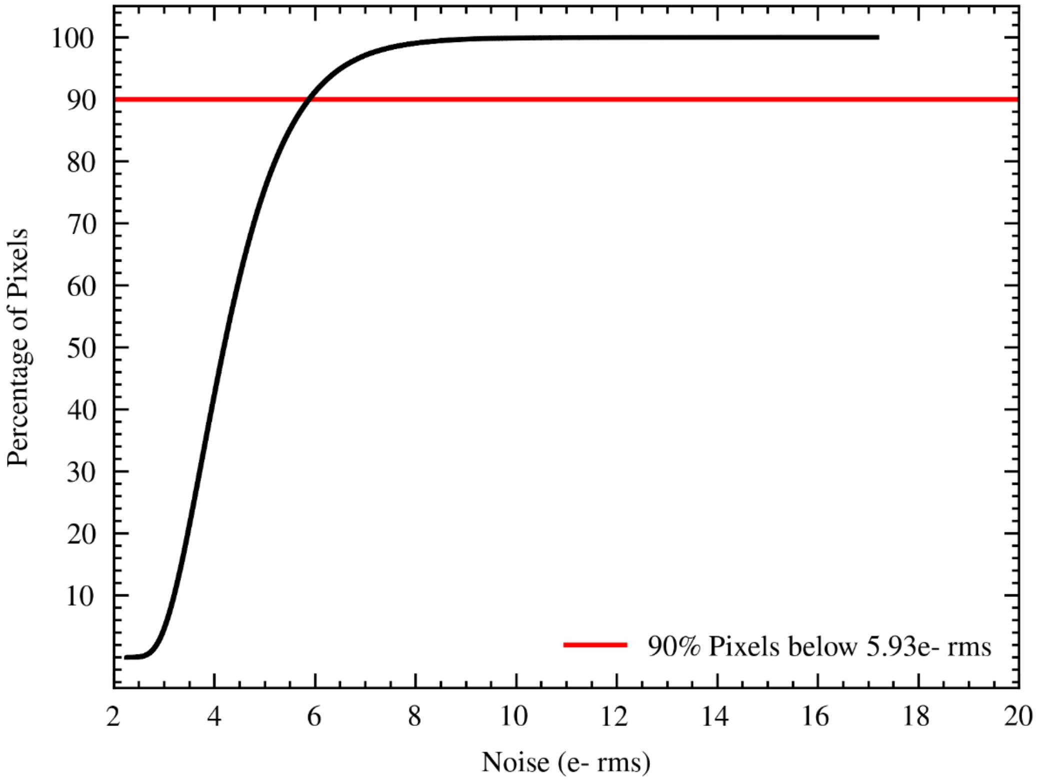

Figure 7-7 - Cumulative noise distribution in CIS115.

Figure 7-7 shows that 90% of the pixels are below 5.93 e- rms, with less than 1% of pixels below 3 e-rms. Here we can define the top 1% of the distribution, pixels with noise above 8.56 e- rms as “noisy” pixels, and the following sections will look at these pixels in further detail.

7.3.4: Noisy Pixel Locations

When determining the impact of noisy pixels on the image, it is interesting to know if there is any spatial preference for where these pixels appear. For example, do noisy pixels appear in clusters or pairs, or is there an observable pattern to their position?

The first step is calculating how many noisy pixels appear on the frame. Again, the same dataset is used, and now the threshold is the highest 1% of standard deviations across the frame. All pixel values are compared to this threshold (8.56 e- rms), with pixels above this value being flagged as noisy pixels. Next, the locations of these pixels are determined, and the area around each noisy pixel is scanned for additional noisy pixels. Figure 7-8 illustrates how this is done.

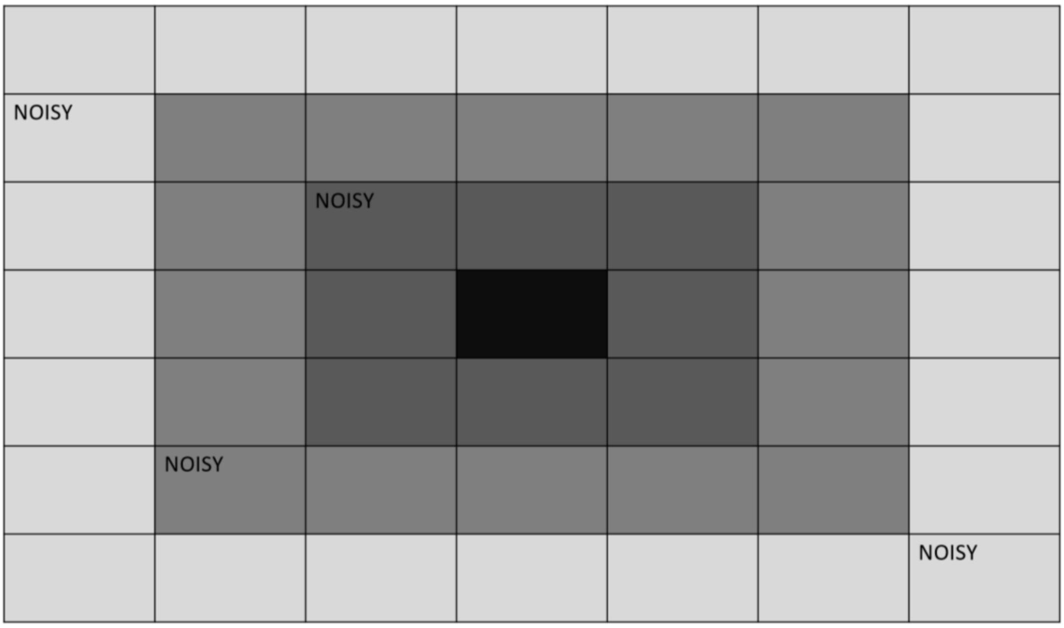

Figure 7-8 shows the search area around a noisy pixel in the centre (in black) and 4 noisy pixels.

Figure 7-8 shows the defined 3-pixel search radius around the central, flagged noisy pixel—the first step to search the entire area comprising 48 pixels. Again, the central pixel is not counted to avoid double counting. Next, the four pixels in the search area that exceed the standard deviation threshold are flagged as noisy pixels. For example, Figure 7-8 would contribute one entry to the third column at the Radius-3 mark since the noisy central pixel has three or more noisy pixels within a 3-pixel radius.

Next, the search area is limited to a 2-pixel radius. In Figure 7-8, this is the grey and a dark grey area which comprises 24 pixels, again not counting the central black pixel. In the above case, there are two noisy pixels within this range. This would contribute to the second column at the Radius-2 mark since the noisy central pixel has exactly two nearby noisy pixels within the 2-pixel radius.

Finally, the search is limited to the adjacent 8 pixels to the noisy central pixel, the dark grey area in Figure 7-8. Here, there is only one other noisy pixel, so this would contribute an entry to the first column on the Radius-1 mark since there is precisely one other noisy pixel within a radius of one pixel.

This process is repeated for each noisy pixel, determining how close these pixels are. For example, if many noisy pixels have many other pixels within their respective search areas, this would indicate that the noisy pixels may be clustered together. Conversely, a small number of noisy pixels with other nearby pixels would suggest that noisy pixels do not cluster.

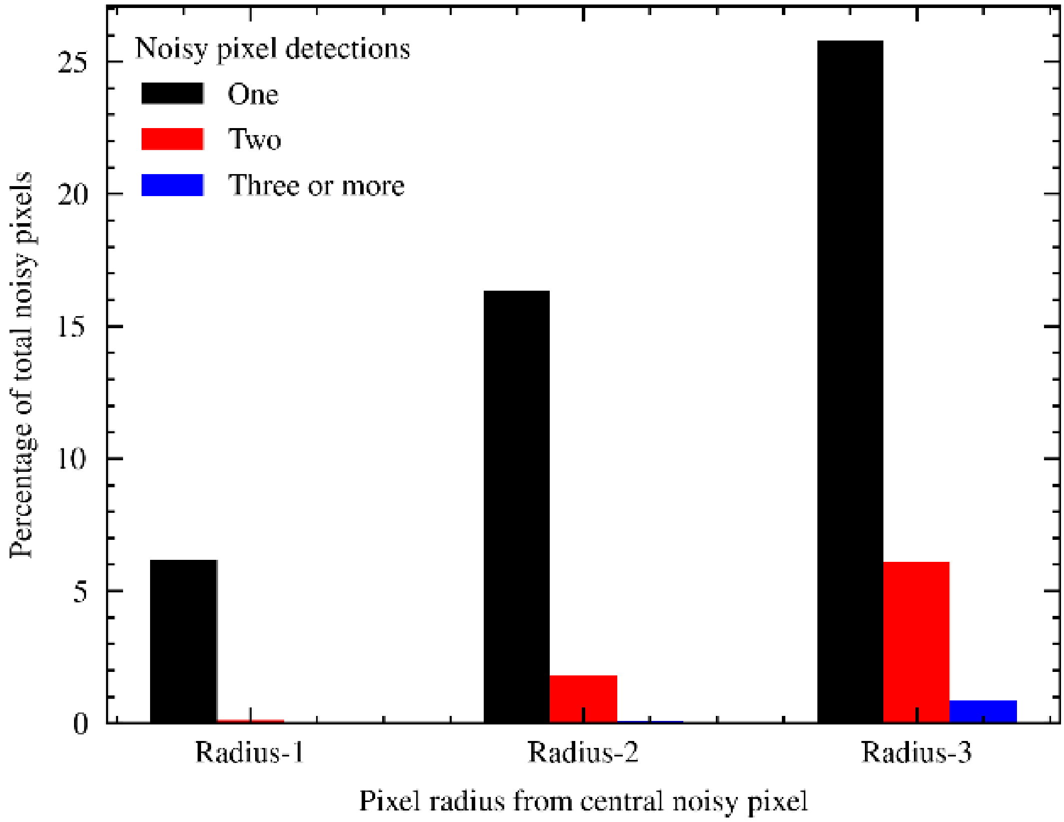

Figure 7-9 – A cumulative bar plot shows how many additional noisy pixels appear within a certain radius of a noisy central pixel. The y-axis is plotted as the percentage of noisy pixels, which comprise approximately 1% of the entire device, as defined by Figure 7-7.

Figure 7-9 shows an increase in noisy pixel detections as the search radius increases. Radius-1 shows almost only single noisy pixel detections. About 6.15% of all 460 noisy pixels detected appear to have one other noisy pixel within a radial range of one pixel; this is the dark grey area in Figure 7-8. For the same radial range, only 10 pixels had two other noisy pixels nearby. There are no pixels found that have three or more other noisy pixels. If we expected significant clustering of noisy pixels, this is not supported by the data from this device.

For a two-pixel search range (the grey area in Figure 7-8), 1221, or 16.3% of the noisy pixels have precisely one other noisy pixel within the area; 137 pixels, or 1.8%, had precisely two other pixels and just 6 pixels had three or more. At the largest 3-pixel radius, 1928 pixels had precisely one other pixel, 457 pixels had exactly two other pixels, and 63 had three or more pixels.

Whilst all three counts increase as the radial search area increases, this is expected because more pixels are searched. With only 0.84% of pixels ultimately having three or more pixels within a three-pixel radius, there does not appear to be a clustering of noisy pixels. However, 100 simulated images of the same size (2000x376 pixels) were used to see if these measurements were typical. This simulation was used to investigate whether the clustering of noisy pixels is typical for a random spatial distribution. Each image was generated via Python’s *random.choices* method. This method allows the user to create random distributions based on a specified population. In this case, the population used is the experimentally measured noise distribution with weights derived from the frequency of appearance. For example, a pixel value of 100 may appear 100 times, but a value of 120 appears only 40 times. This difference in frequency means that while 100 and 120 appear in the population list, the value 100 has a higher weighting due to it appearing more frequently in the original dataset. Figure 7-10 shows an example of a generated distribution vs a measured distribution.

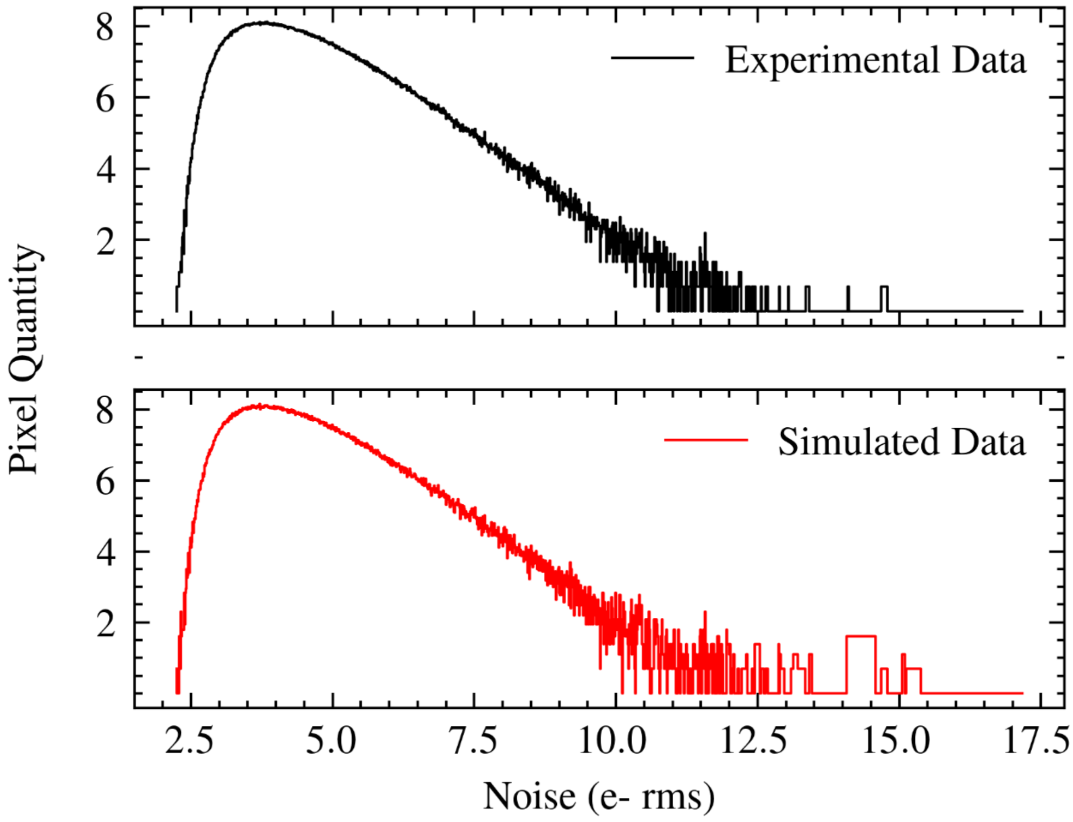

Figure 7-10 - Top plot shows the measured data's noise distribution, as shown in Figure 6. The bottom plot shows the simulated array distribution based on the actual data.

Figure 7-10 shows that the method can generate random distributions similar to the measured dataset. The experimental dataset has a mean noise value of 7.6 e- rms and a standard deviation of 3.2 e- rms, whilst the simulated dataset has values of 7.3 e- rms and 3.0 e- rms, respectively. Both sets' mean and standard deviation values are close in value, showing that these distributions are similar. The simulated datasets are based solely on pixel bin value and the number of pixels in each pixel bin; given this is the case, there is no dependence on spatial positioning of pixel values. Because there is no relation to any pixel positioning in the experimental dataset, if there is any spatial preference for noisy pixel generation in the experimental dataset, the simulated datasets should not be impacted by this. This also aids in making the comparison fair and ideal to find any abnormalities within the experimental dataset.

In this approach, 100 datasets are generated and compared to the experimental dataset below.

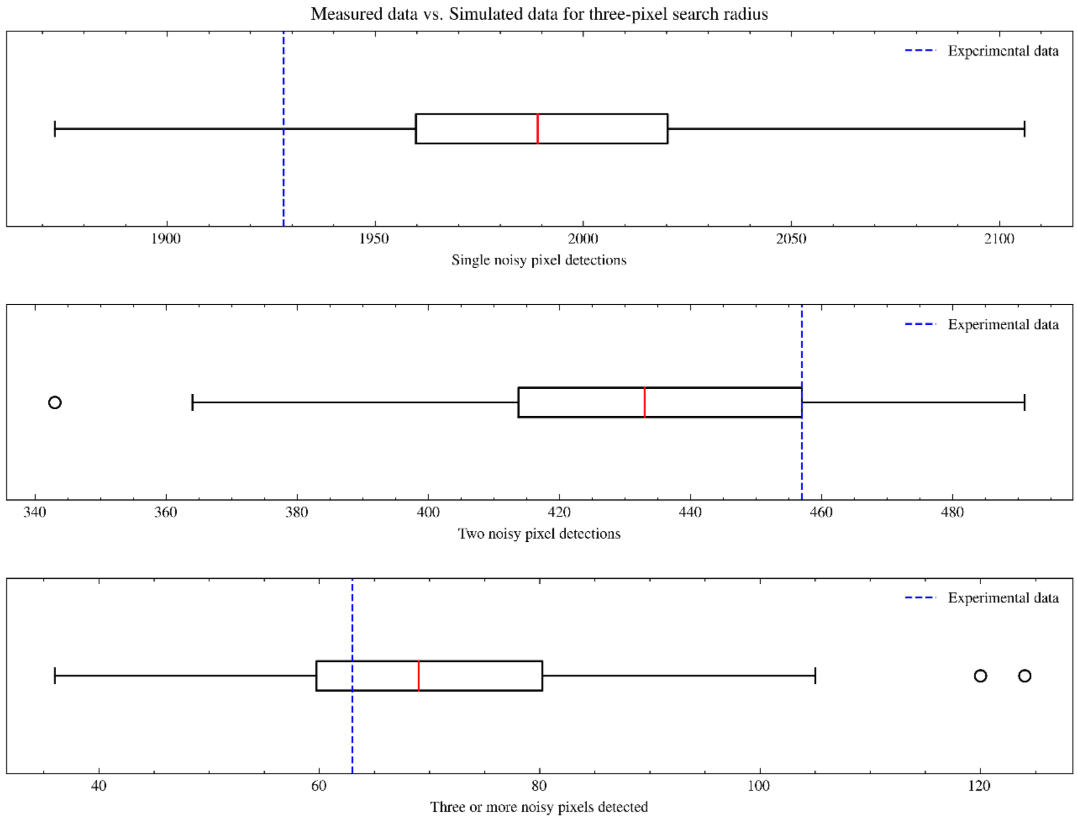

Figure 7-11 – A box plot shows a comparison of the experimental and simulated datasets for quantity of detections within a three-pixel radius.

(Top) Experimental data measured 1928 noisy pixels with exactly one nearby noisy pixel within the 3-pixel radial range. This value appears below the first quartile (Q1), about 1961, but above the minimum value, about 1871.

(Middle) Here, the experimental data value measures 457 pixels with two other noisy pixels within the three-pixel search radius. This value is slightly above the third quartile value, 456 pixels, for the simulated distribution and below the maximum value. There is a single outlier much below the minimum value at 342 pixels.

These outliers are defined as outcomes that are notably higher or lower than the rest of the dataset. The standard limits defined for this are:

$$\text{outlier} < \mathrm{Q1} - 1.5 * \mathrm{IQR} \tag{50}$$

$$\text{outlier} > \mathrm{Q3} + 1.5 * \mathrm{IQR} \tag{51}$$

Here Q3 is the third quartile value, and IQR is the interquartile range, equal to the difference between Q1 and Q3.

(Bottom) The experimental value is 63 pixels with three or more other pixels within the three-pixel search radius, within the interquartile range of the distribution of the simulated sets.

Other ranges, the two-pixel search and one-pixel search radii, show similar results, with the experimental data being within the bounds set by the simulated datasets. This suggests that the experimental dataset is not significantly different from the simulated datasets. Given that the simulated datasets do not have a spatial preference, given that the method of generating them is inherently random, we can assume that the experimental dataset also shows no spatial preference for noisy pixel generation. If there were some spatial preferences, it should show significant clustering of noisy pixels.

A correlation test should be conducted between the pixels to determine whether noisy pixels impact others nearby properly.

7.3.5: Noisy Pixel Characteristics

This section aims to discover whether noisy pixels have any impact on the signal of nearby noisy pixels throughout 250 frames.

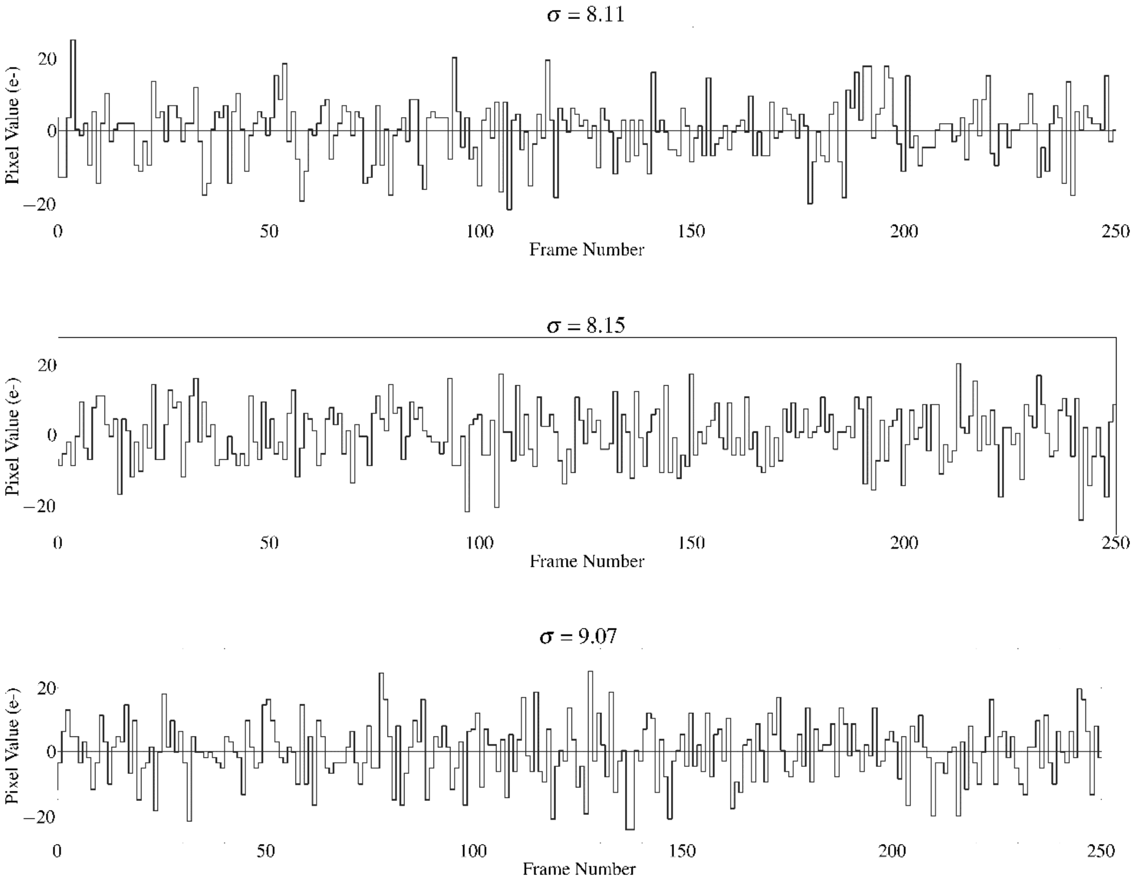

Figure 7-12 - Shows three pixels that are adjacent to each other, positioned at (3,105), (3,106) and (3,107).

Figure 7-12 shows three adjacent pixels, a, b and c, in the same row. When deciding what test should be used for these datasets, it should first be established whether they are linear or monotonic to establish what type of correlation test can be used. The Pearson test is used to investigate whether there is any correlation between two linear datasets. A dataset is linear if the dataset values increase linearly over time. In this case, if the pixel values were to increase throughout the frames, our pixel datasets would be linear. A quantifiable method would be to apply a linear regression model to the dataset. Linearly distributed data should have high accuracy, whereas low accuracy would suggest that the dataset is non-linear. In this case, the linear regression model shows inferior accuracy, so none of the pixels' frame-to-frame signal outputs is linear and therefore, the Pearson test is not appropriate for looking for pixel-to-pixel correlation.

The Spearman's rank correlation test, which does work with non-linear datasets, would be a candidate test; however, it is not designed for non-monotonic datasets. A monotonic dataset contains a function that increases or decreases in value over the dataset's length, i.e. a linearly increasing dataset is also a monotonic dataset. Again we can see in Figure 7-12 that our pixel datasets are not monotonic because they continuously vary between high and low values without any particular pattern. To show that the datasets are not monotonic, Spearman's rank correlation can be fitted to each dataset, the pixel values and frames, and if there is a good correlation between these values, then the dataset can be described by a monotone function. However, in the case of these pixels, the Spearman rank correlation between pixel signal and frame number is poor, giving evidence that these pixel datasets are non-monotonic.

Because of this constraint, a test capable of handling non-linear and non-monotonic datasets is required. One such test is Hoeffding's D independence test. The details of how this test works can be found in (Hoeffding, 1948) and (Hollander, 1999). This test's results can be interpreted similarly to Pearson and Spearman; values that are close to 0 are considered not correlated. Positive values tending towards one are positively correlated, and negative values tending towards -1 are negatively correlated.

	Pearson R test	Spearman R test	Hoeffding D test
Pixel 51a & 51b	-0.0307	-0.0132	-0.0001
Pixel 51a & 51c	-0.0231	-0.0154	-0.0002
Pixel 51b & 51c	0.0518	0.0598	0.0007

Table 7 – Results for the Correlation Tests.

Table 7 clearly shows that the noise signals from the pixels in Figure 7-12 are not correlated. This non-correlation is expected if each pixel's readout noise is independent of nearby pixels.

This analysis is continued with additional noisy pixel configurations, for example, a pair of noisy pixels that appear in the column. Overall, there are no groups of noisy pixels that show any correlation with other nearby noisy pixels. Overall, it appears that readout noises are not correlated; to an extent, this is expected. The readout circuits of each pixel are separate from each other, so there should be no interference. Readout column noise could occur, but this would impact all the pixels in that column, not just pixels with high readout noise. In addition, this noise is typically easier to detect because the entire column of pixels would exhibit higher noise than other columns on the device.

7.4: Discussion

These tests aim to understand what noise sources contribute to device noise for the CIS115 and their impact on a Lucky Imager's performance.

Figure 7-4 and Figure 7-5 demonstrate that the contribution from dark current is close to zero electrons per frame after bias subtraction but does have a slight impact on the noise, as shown in Table 6, when operating the device at readout rates above 10 Hz, as per Lucky Imaging's requirement. Additionally, although the device tested in this chapter was not cooled, external cooling will further reduce the impact of the dark current.

Figure 7-6 and Figure 7-7 show that the noise spread between pixels is 2 to 16 e^- rms, and 99% of the pixels have noise below 8.56 e^- rms. The majority of pixels have noise ranging from 3 to 5 e^- rms. Unfortunately, this readout noise is already relatively high, which makes the CIS115 not suitable for photon counting applications unless other noise reduction methods are available.

GravityCam's intention to effectively use the microlensing technique to observe smaller planetary and even lunar masses would require a high signal-to-noise ratio in order to be able to confidently determine whether an electron is likely to be a photoelectron and not a noise electron. Because of this requirement, the device noise needs to be close to EMCCD's sub-electron effective read noise.

Section 7.3.4: suggests that noisy pixels have no spatial or clustering preference across the device, and Figure 7-9 shows that there is no significant degree of clustering. 7.3.5: shows that there are no pixels with any significant noise correlation, so any noisy pixels that appear near each other do so purely by chance. No evidence suggests that noisy pixels that appear clustered have localised impacts on other nearby pixels.

Assuming that noisy pixels can appear in any location on the device, then any pixel has a 1% chance of being defined as a noisy pixel, given that the definition of a noisy pixel, in this case, is being in the top 1% of pixel readout noise across the frame. The difference in the average noisy pixel value compared to the average pixel value can be seen in Figure 7-13. It shows that for over 1000 frames, the mean value of all noisy pixels varies between -1.3 and 1.7 ADU, with a mean value of approximately 0 ADU per frame, which should be expected for a device operating in darkness with bias subtraction. Comparing these values to pixels with a standard deviation between 3.5 and 5 ADU rms, denoted as average pixels, given that these occupy the highest quantity bins in Figure 7-6. These average pixels have a greater concentration of about a mean value of 0 ADU per frame, with a minimum value of -0.8 ADU and a maximum value of 0.7 ADU.

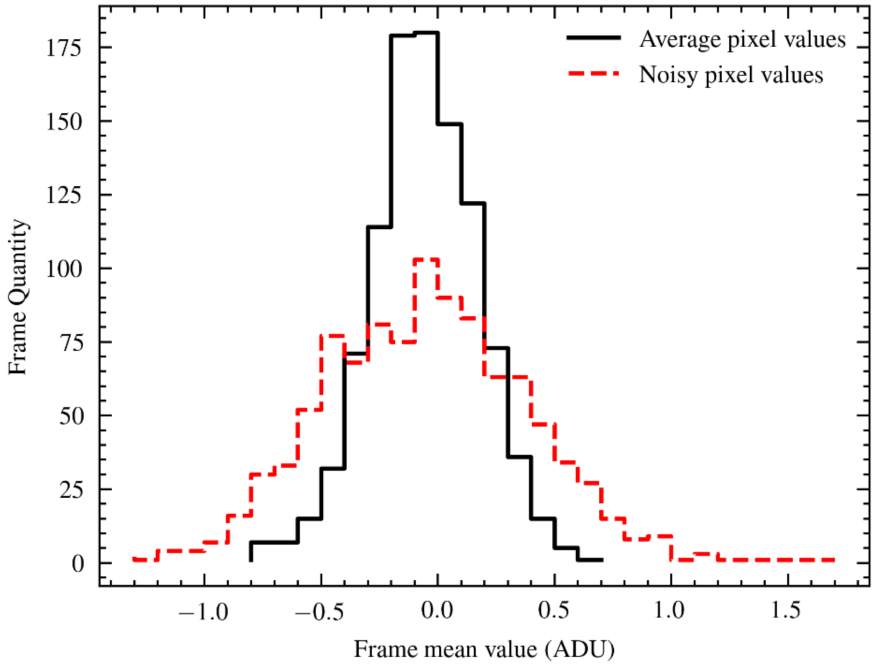

Figure 7-13 – Distribution of the mean values of pixels, with standard deviation about the peak of Figure 7-6, per frame.

The critical difference between the noisy and average pixel values is the spread. The noisy pixels show a more significant standard deviation of frame means and a large difference between the minimum and maximum value, and given that we know noisy pixels show a more significant standard deviation, this makes sense. However, most noisy pixels (60%) show a frame mean value of about 0 ADU per frame, which shows that for an average frame, many noisy pixels will yield a value similar to average pixels. There are a few frames, less than 10, where the frame mean value for the noisy pixels is notably higher than the average pixels, so a small number of frames can be impacted by these pixels. However, given that Lucky Imaging requires taking numerous frames, the contribution from these frames will be averaged. Therefore, it is unlikely that these noisy pixels will significantly impact the Lucky Imaging reference frame and, therefore, the overall image quality analysis.

Chapter 8: The effect of CIS-like noise on stellar detection limit

This simulation aims to find a relation between the faintest detectable stellar magnitude and the mean of CIS-like readout noise on the frame. Ideally, this will allow us to understand better how the Lucky Imaging reduction technique is impacted by CIS-like noise, what noise levels correspond to the loss of targets at different stellar magnitudes and how the presence of noisy pixels influences star detection. Furthermore, knowing at what noise level certain magnitudes become challenging to observe will help demonstrate the importance of noise reduction for lucky imaging.

The method of constructing this simulation is split into three different parts. Firstly, the atmospheric seeing is simulated using a python library called MegaScreen, detailed in section 8.1:. The second step is to simulate the CIS-like noise. This is generated using the same method as the noise simulations described in section 7.3.4:, with some changes discussed in section 8.2:. The final step is to generate a FITS cube using the simulated dataset and process it via the same IDL pipeline used at the Danish telescope, discussed in section 8.3:.

8.1: Atmospheric Seeing

The atmospheric seeing is simulated using the python library MegaScreen, which implements an algorithm described in (Buscher, 2016). The Buscher algorithm uses a Fast Fourier Transformation (FFT) approach to generating phase screens, which randomly perturbs the optical wavefront’s phase based on input parameters. Firstly, a telescope is defined, in terms of aperture diameter, focal length and focal ratio. Using these characteristics, the diffraction-limited resolution can be found as defined in equation (32).

Secondly, the atmospheric turbulence is simulated using Kolmogorov phase screens. First defined in (Kolmogorov, 1991), this method splits the atmosphere into discrete layers describing how the phase of optical propagation changes at each layer. Increasing the resolution of each screen will improve the ability to simulate smaller and smaller turbulent aberrations and increase the processing time and power required. More information on these techniques can be found in (McAulay, 2000).

A speckle pattern can be generated using phase screens to simulate atmospheric seeing. Assuming that the entire frame is subject to the same blurring, each object in the frame can be folded with the same speckle pattern or PSF stamp. An array is generated with a single, central high-value pixel. Passing this single pixel array through the telescope simulator discussed above will create a PSF stamp of this point-like object, as seen in Figure 8-1. The stamp can then be folded with the reference flux of a star to create a blurred star.

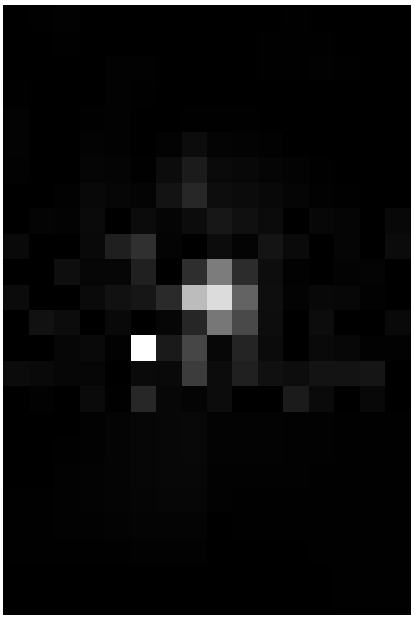

Figure 8-1 - An example of a point-like object becoming blurred by the telescope simulator.

This should be repeated for the number of stars desired, and then each star can be positioned to match the original star's position in the scene chosen for the simulation.

Completing this, yields a simulated approximation of the original scene, with any prior device noise substantially removed from the background.

The observational data chosen for this chapter is the globular cluster NGC6528 seen in Figure 8-2. The data were obtained at the Danish telescope using the EMCCD instrument as detailed in (Skottfelt J. e., 2014). The reason for choosing a globular cluster is that they are densely packed stellar scenes, allowing for numerous stars of differing magnitudes over a relatively small field of view, which is the case for the NGC6528 cluster. Additionally, the dataset obtained from the cluster is easily accessible. The magnitude list of all the detected stars gives both position of each star on the frame and its reference flux. Using this list, a simulated version of NGC6528 can be easily constructed.

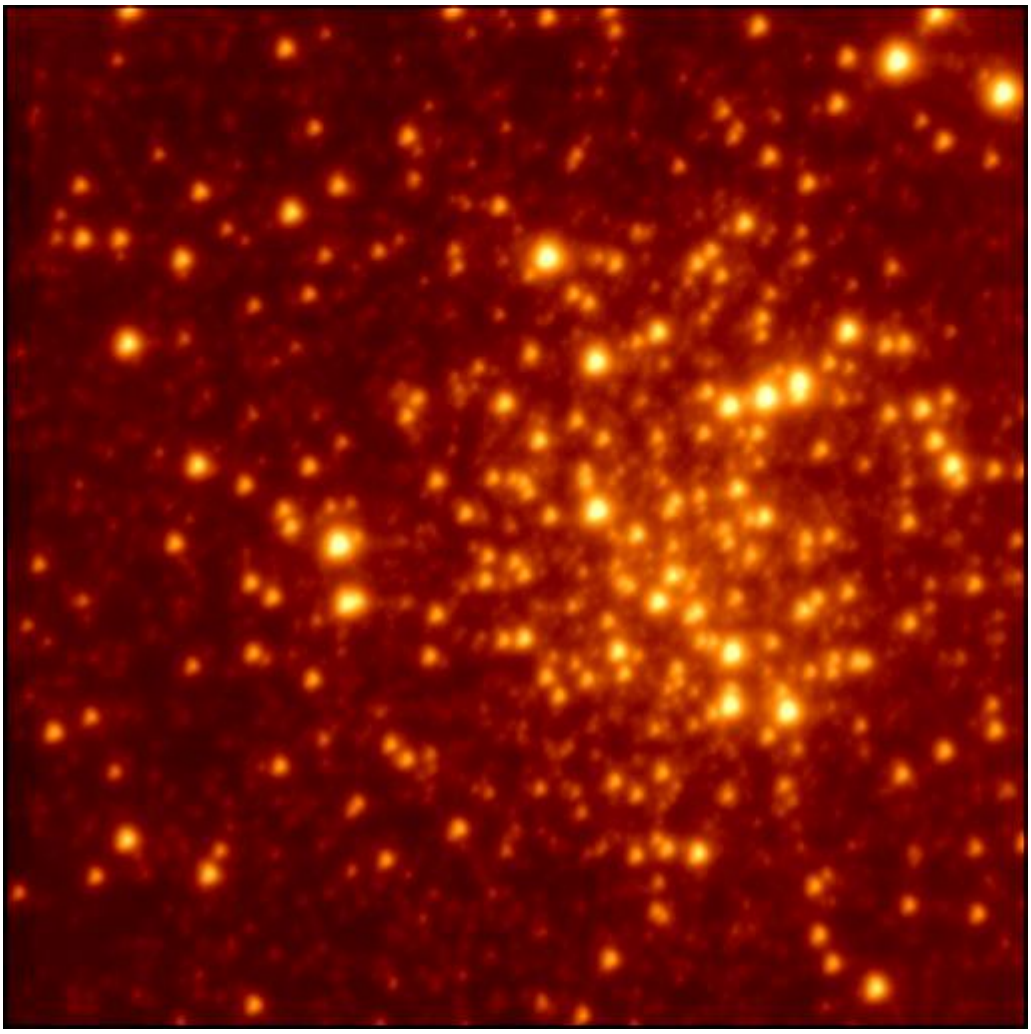

Figure 8-2 – Globular Cluster NGC6528, high-resolution image taken using the Danish Telescope.

8.2: CIS-like Noise

The next step is to add CIS-like noise to the image array. As discussed in Chapter 7, CIS noise is described well via a log-normal distribution. This distribution shows that the vast majority of pixels have a readout noise close to the mean readout value of the entire frame; however, there is a long tail of higher value pixels.

The first step draws a list of each pixel's readout noise taken from the experimental data in Chapter 7:. This list generates a new array using the *scipy.random.choices* function.

This function allows the user to generate a probability density function (PDF) using an input list of values. This means a new array of pixel readout noises can be generated based on the experimental data. This ensures that the simulated dataset correctly models a CIS device.

Once the simulated readout noise array is generated, these readout noise values can be used to generate an array of output signals. The output distribution of a CIS pixel can be modelled using a normal distribution, with a mean value chosen by the user and a standard deviation given by the pixel’s readout noise value. This generates a new array of CIS-like noise with the chosen mean. PDF can be modified by increasing or decreasing the values of the input list. This will respectively increase or decrease the pixel readout noise array values. An increase should effectively increase the number of noisy pixels by increasing the standard deviation of pixels within the simulated array whilst decreasing, having the opposite effect, making the simulated array more uniform.

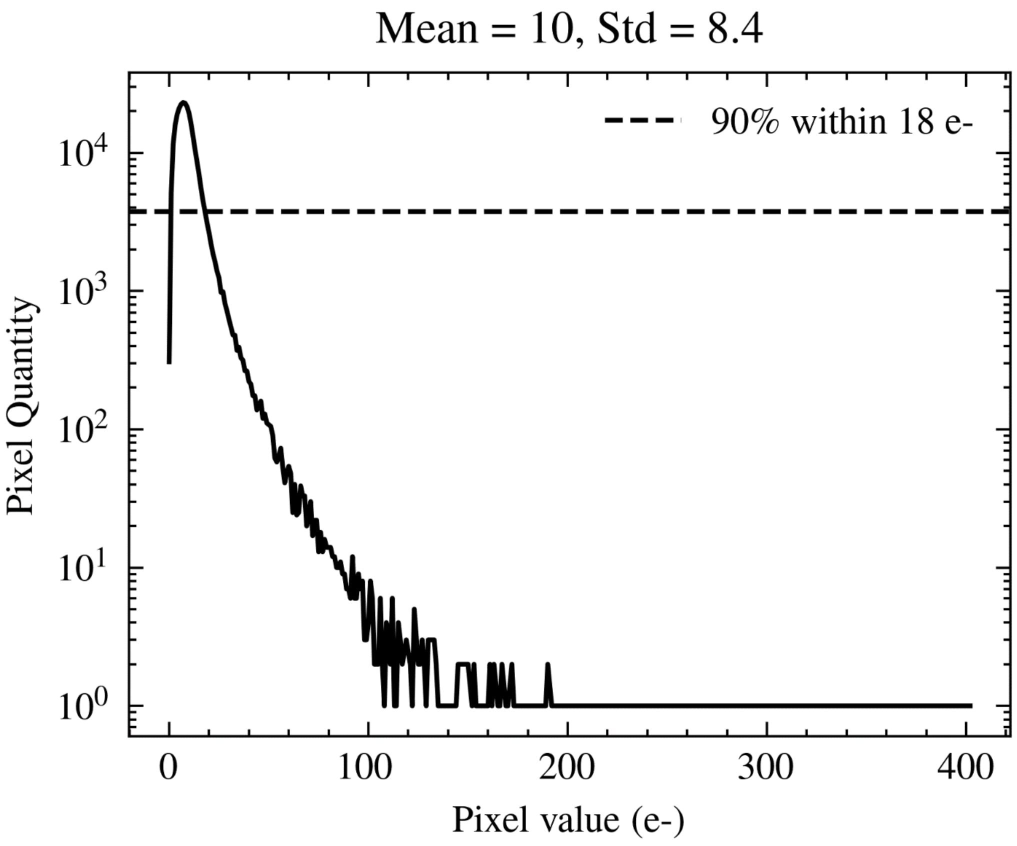

Figure 8-3 – Simulated CIS-like noise, pixel-quantity versus pixel value (e^-).

Figure 8-3 shows that for a mean signal of 10 electrons and standard deviation of 8.4 e^- rms, 90% of pixels exhibit a signal less than 18 electrons. Although there is a large tail of pixels tending towards a signal of 200 electrons, these bins contain much less than 10% of the total amount of pixels and therefore have a small but the non-negligible probability of impacting the frame’s quality for star detection.

In order to explore how different mean and standard deviation values impact star detection, more distributions with mean values ranging between 0.01 and 20 electrons are generated. The readout noise is also increased in alternate arrays by 3x to simulate an increase in noisy pixels.

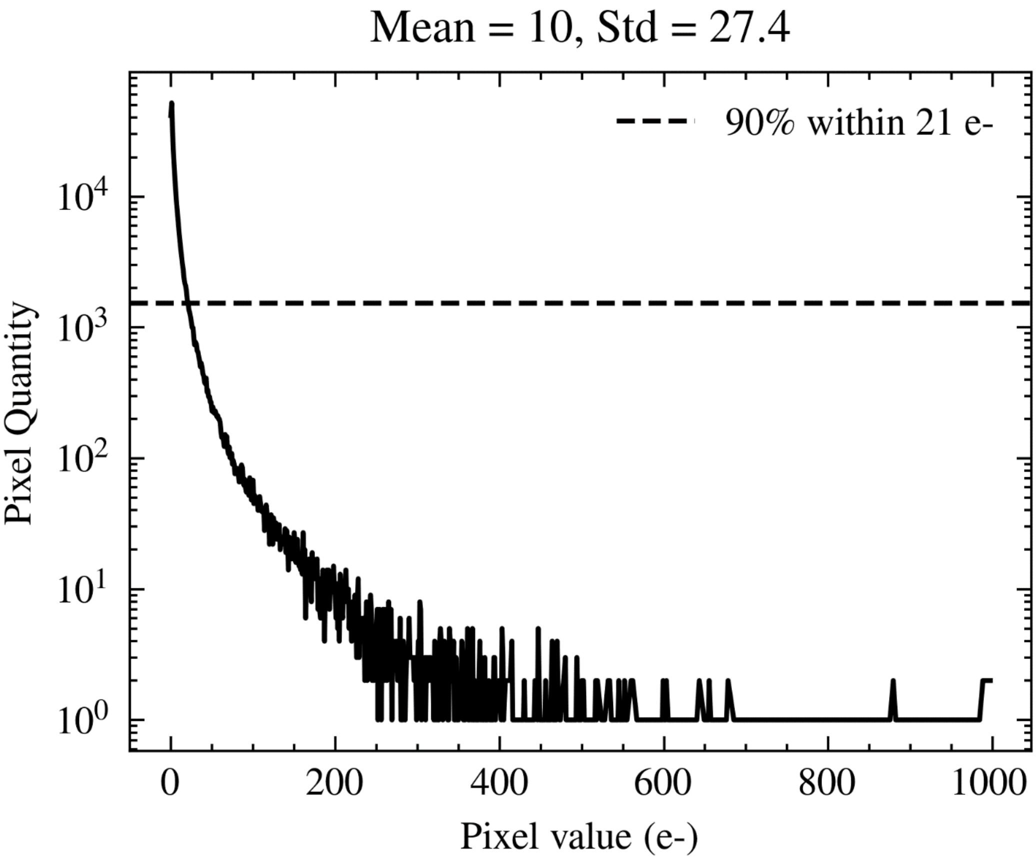

Figure 8-4 – Similar to but with a PDF multiplier of 5x, showing a notably longer high value pixel tail.

Figure 8-4 shows a much greater standard deviation value at 27.4 e^- rms, resulting in a much longer high-value pixel tail tending towards 600 electrons. Although this tail still contains only 10% of the total amount of pixels, the higher value of these pixels could lead to a more significant impact on the quality of the frame for star detection.

Once the noise is generated, it is then added to the star map's array creating an image of the globular cluster with CIS-like noise.

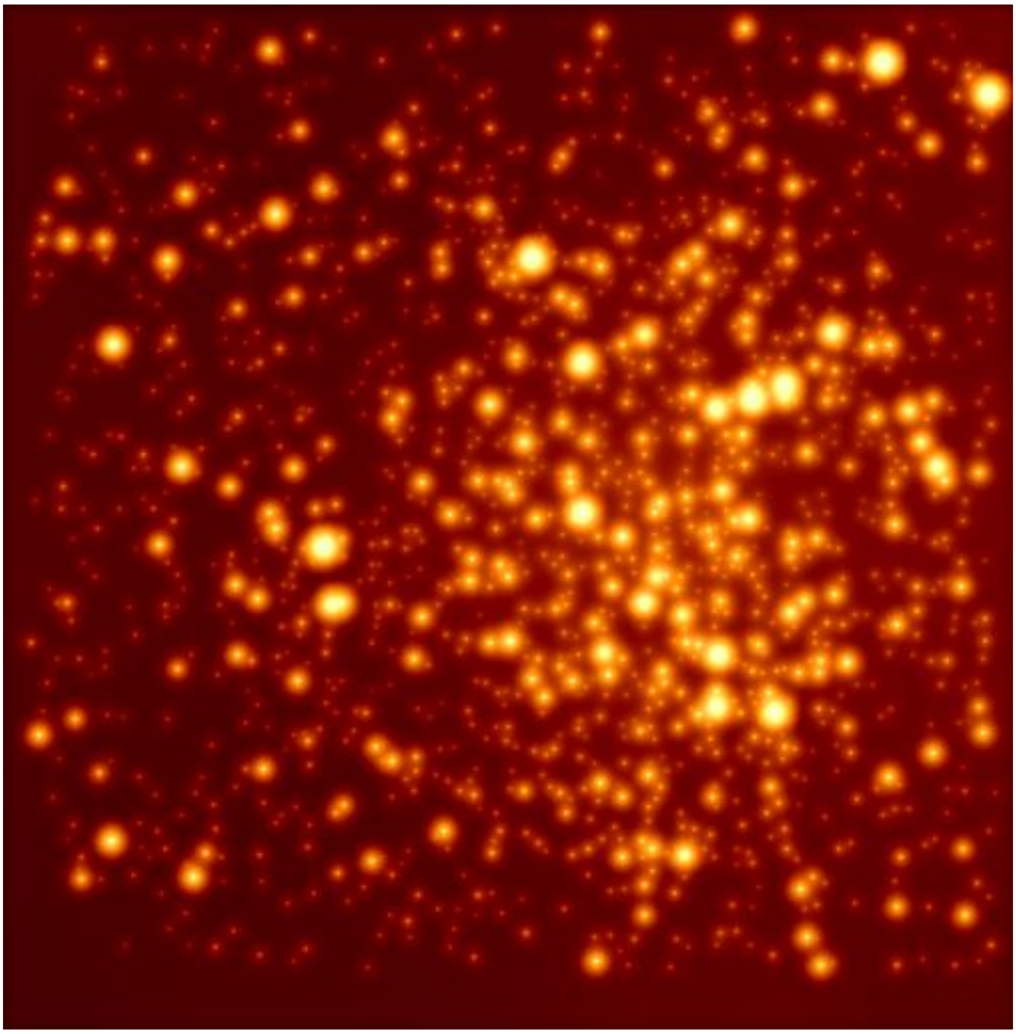

Figure 8-5 – Simulated best resolution image of NGC6528 with a mean CIS-like read signal of 0.01 e^- and standard deviation of 0.2 e^- rms.

Figure 8-5 shows the best quality image produced by the simulation, which is the closest to the reference dataset, as shown in Figure 8-2. The next step is for this simulated dataset to be analysed by the ODIN lucky reduction software used at the Danish telescope. Whilst there are other reduction methods, this is the most familiar to use and potentially the easiest to modify for this simulation. Additionally, this is the same software used for the NGC6528 cluster.

8.3: ODIN Reduction Software

Each of the simulated frames is stored in an HDF5 library, (HDF Group, 2022), which is a file format that can store many different types of n-dimensional datasets and other types of complex objects. This format is used as observational data from the EMCCD imagers needs to contain different types of data besides the array of science exposures, such as a master bias frame, a master flat frame, and meta data about the observation. The same file can then be used to store the calibrated data, or just the parameters needed to produce the calibrated data.

For a normal observation, image correction starts with image calibration, however, for simulated data these steps are not necessary.

Next step is correcting the tip and tilt of the image. These are types of aberration discussed in 4.1.2:. Tip and tilt are corrected by constructing a reference image and using this reference image to shift and align the image exposures properly. The reference image is made by taking the mean of all the images after master bias and flat correction. More detail on how this is accomplished can be found in (Harpsøe, 2012) and (Arazia, 2008).

Finally, the image quality of each image can be calculated by cross-correlating the PSF peaks of each image with the reference frame. It is assumed that the reference image is of good quality; therefore, images with a high correlation with the reference frame are also of good quality. Using this approach, each image is assigned a quality factor, and the highest quality images are then re-used to construct a new master reference frame. Then using this new master reference frame, quality factors are finalised for each image. Calculating the master reference frame in this way should reduce the impact of bad quality images on the reference frame and, therefore, improve the calculation of the quality factor of each image.

Once all quality factors are finalised, this output is typically a 10-layer FITS file, where each layer is a composite image generated from shifting and adding frames with a quality factor above a given threshold. In this case, the thresholds, or cuts, are defined from highest quality to lowest quality in percent (1, 1, 3, 5, 10, 30, 40, 8, 1, 1). The first cut thus contains the best 1% of quality factor images. The second contains the second best 1%. The third contains the next best 3% until finally reaching the end which contains the worst 1%. This allows the user to view the best images first, looking at the highest spatial resolution before looking at the cuts that contain far more images but at a lower spatial resolution, effectively seeing if any important targets are lost in the highest quality images.

8.3.1: DanDIA IDL Starfit

Once the FITS cube has been generated, it can be passed through the DanDIA IDL pipeline (Bramich, 2008), used at the Danish Telescope. This method works by utilising the high-quality reference image mentioned above to perform a technique called difference image analysis (DIA). DIA effectively tries to compare images from the stack to the reference image. It does this by trying to model each image I_{ij} using the reference frame R_{ij} and a pixel kernel. The pixel kernel K_{lm} is some user defined area of l columns and m rows, which contains a multiplier changing the reference image pixel values (i, j) to closely model M_{ij} that of the target image pixel values. The differential background is defined as some unknown constant B_0.

$$M_{ij} = \sum_{lm} K_{lm} R_{(i+l)(j+m)} + B_0 \tag{38}$$

(38) shows this modelling technique, where K_{lm} and B_0 are solved for using a least-squares method, trying to achieve the best possible model of the targe image I_{ij}. Finally

a difference image is constructed by subtracting the target image I_{ij} and model image M_{ij}. This difference image will primarily consist of Poisson noise, mostly shot noise, and locations where stellar objects have varied in brightness or size. This can then identify areas where constant stellar objects are located on the frame and where highly varying stellar objects exist. Ultimately, this produces a list of possible stars, their location on the frame and a reference flux (F).

8.4: Danish Telescope Specific Case

The above subsections detail a general case of how the simulation is carried out. This subsection describes the specific parameters used for mimicking the original dataset taken using the 1.54m Danish telescope.

To best compare the simulated CIS dataset with the original EMCCD dataset, many of the original parameters were copied for the simulation. The framerate is 10Hz, the seeing conditions are copied from the site, La Silla. The telescope diameter is 1.54m and focal length about 13m. Using these parameters, the PSF stamps are generated. Given that the original dataset uses thousands of frames for the Lucky Imaging process, the simulated dataset should be of the same scale. For this reason, 5000 PSF stamps are generated which represent the initial exposures which are currently only impacted by atmospheric turbulence.

The simulated CIS device size is copied from the CCD97 device that it is being compared, the device used at the Danish telescope, which is a 512x512 array with 24 overscan columns, leading to a total array size of 512x536. Using the CMOS data taken in Chapter 7: the CIS noise distribution is used to generate many distributions as seen in Figure 8-3. A selection of different means are chosen, ranging from 0.01 e^- to 10 e^-. This is done to show the impact of increasing CIS noise on stellar detection limits and to be able to compare with the original EMCCD dataset. Overall, each CIS mean dataset contains the same quantity of noise frames as there are PSF stamps. This creates a set of 5000 simulated exposures, each with different stamps and a different noise distribution to best approximate a non-simulated observation dataset.

The final step is to run the ODIN reduction software, which takes the 5000 exposure stack and generates a HDF5 file. Some modifications were made to this phase of the data processing to accommodate for the simulation, such as there being no bias in a simulated dataset or any need for flat-field reduction. This ultimately results in a FITS cube as described above and run through the DanDIA IDL software, generating a starlist which is then used below to compare how the simulated CIS datasets compare to the original, observational, EMCCD dataset.

8.5: Comparing Magnitudes

Once a starlist has been generated, the list details how many stars are found and what the reference flux (ADU) is for each object. Using a zero point (M_0) taken from (Skottfelt J. e., 2014), a relationship between reference flux (F) and magnitude (M) can be defined –

$$\mathrm{M} = \mathrm{M}_0 - 2.5\log_{10}(F) \tag{39}$$

Using (39), the list of reference fluxes can be converted into magnitudes and then binned to see how many stars of each magnitude are detected by the ODIN software.

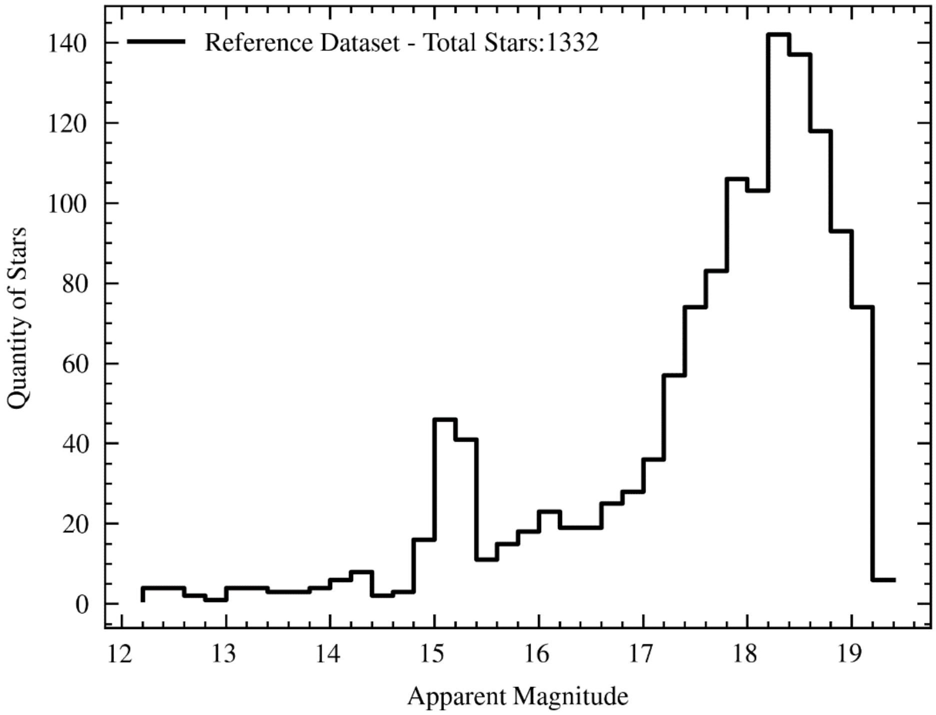

Figure 8-6 - Apparent magnitude versus quantity of stars.

Figure 8-6 shows that for the reference dataset that 80% of stars are within the magnitude bins 17.0 and 19.4. The stars are very faint at these magnitudes and could easily become obscured at higher noise levels, as seen in Figure 8-14.

8.5.1: Sub-electron Mean Distributions

At the lowest noise mean value setting of 0.01 electrons, it is expected that the number of stars should be the same as the reference dataset due to a lack of noise obscuring any stars.

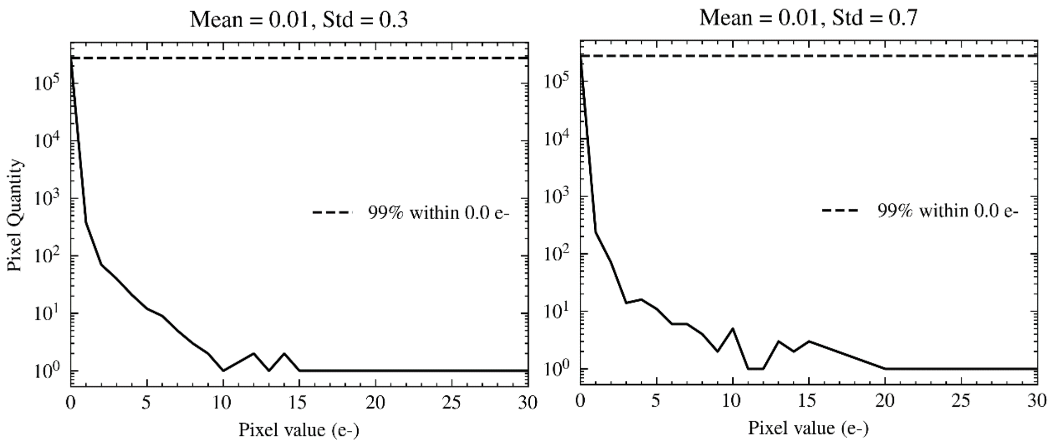

Figure 8-7 – Two CIS-like noise distributions with a mean value of 0.01 e-. The right plot has 3x the pixel readout noise.

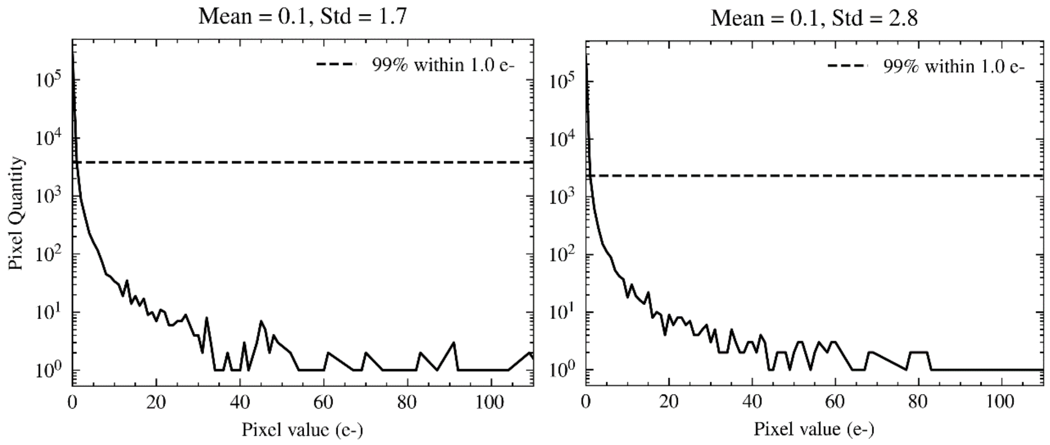

Figure 8-8 – Similar to Figure 8-7 with both distributions having a mean value of 0.1 e- instead.

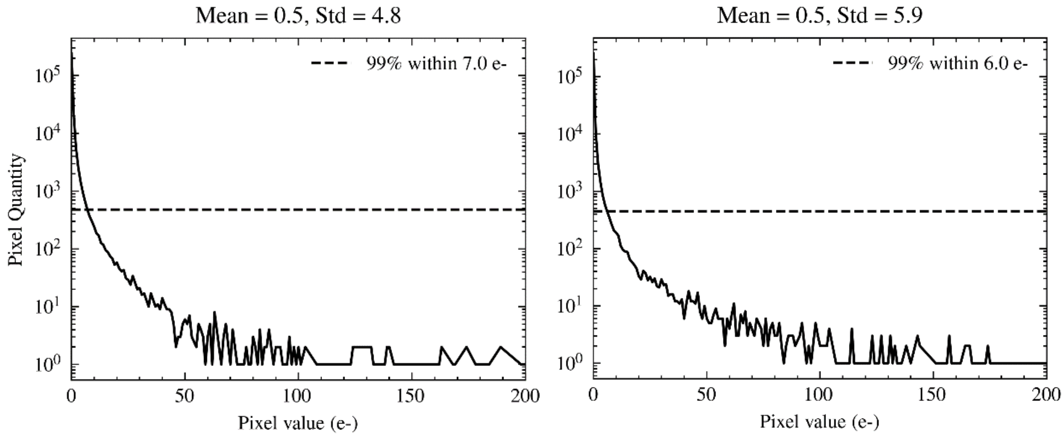

Figure 8-9 – Same format as previous with a mean value of 0.5 e-.

Figure 8-7, Figure 8-8 and Figure 8-9 show plots with increasing mean values for a CIS-like noise distribution. Right-hand plots have 3x on the readout noise. However, on all three plots, the increased readout noise has little impact beyond increasing the maximum pixel value observed because the impacted pixel values at these means are less than 1% of the total amount of pixels.

Figure 8-7 shows a noise distribution with a mean value of 0.01 electrons, 99% of the pixels in the noise array do not contain any noise, with 1% containing mostly 1-3 electrons and extremely few beyond up to 20 electrons.

Figure 8-8 shows that for a mean of 0.1 electrons, 99% of the pixels have either no noise contribution or a contribution of 1 electron. 1% of pixels tend decreasingly toward 80 electrons. There are so few pixels with noise contribution above 10 electrons that magnitudes that require noise signals above 10 electrons to be obscured should always be visible from this noise distribution.

Figure 8-8 right-hand plot shows an increase in the quantity of pixels in higher value bins of 20 electrons and upwards, but 99% of the pixels still are within the 1 electron value range. Because 99% of the pixels are still below this limit and 99% of the signal is identical, the expected change between these two distributions for star detection should be negligible.

Figure 8-9 clearly shows that the amount of signal contained within 99% of the pixels has increased from 1 e- to 7 e-. This means that any magnitude obscured by noise at 7 -e or below will have a good chance of being impacted by this noise.

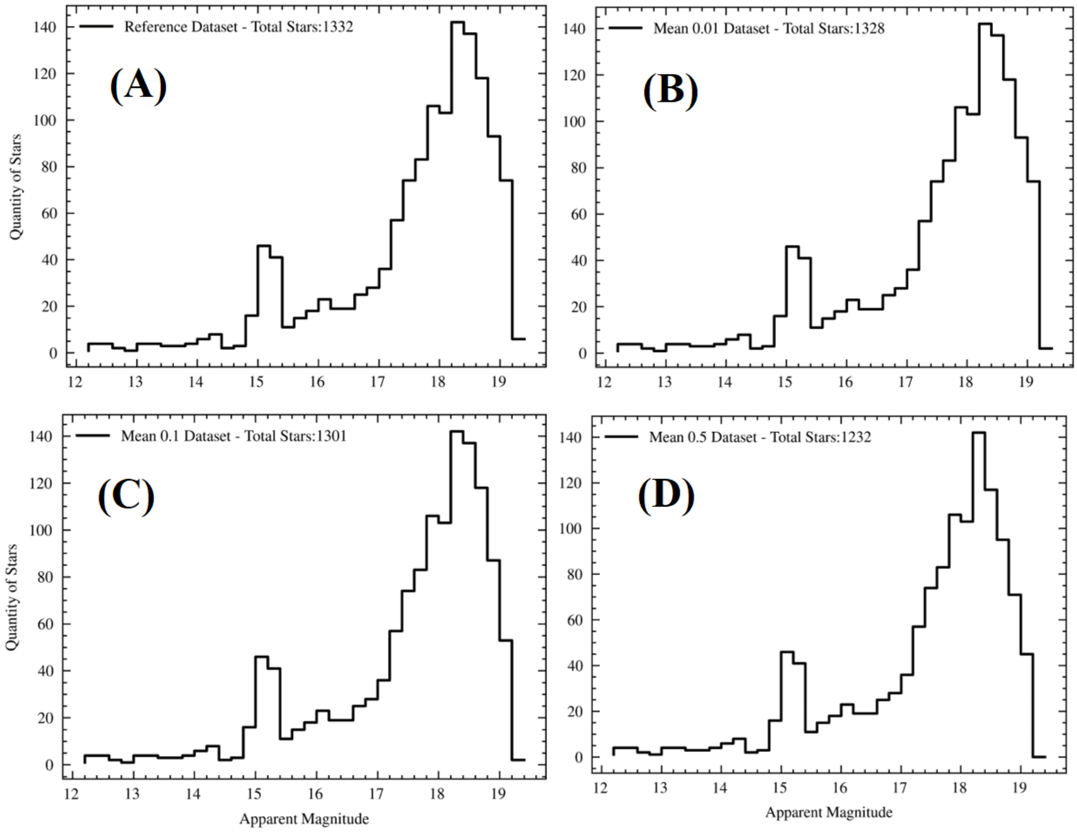

Figure 8-10- Group of magnitude plots with different mean noises in electrons.

Figure 8-10 (a) shows the reference magnitude dataset. This is the result dataset taken from the EMCCD instrument used at the Danish Telescope, yielding a star list containing 1332 stars. All subsequent magnitude datasets are subject to the previously shown CIS-like noise and will be compared to this dataset.

Figure 8-10 (b) shows that only four stars are lost compared to the reference dataset, resulting from losing the faintest detectable objects at magnitude 19.4. The rest of the star list shows that shallow noise levels are best at preventing star loss, which is expected given that a minimal amount of noise has little power to obscure observation. This suggests that the simulation is working as expected. The next level of noise explored is a mean of 0.1 e- in (c)

The only difference between (c) and (b) is the loss of stars at magnitudes 19.2 and up. This suggests that the increase of the distribution mean value from 0.01 electrons to 0.1

electrons can obscure objects in this range, although not all of the 19.2 magnitude stars are lost, only approximately 1/3rd.

Compared to the other three figures, (d) shows a greater loss of stars in 18.8 and 19.0 magnitudes due to the increase in mean distribution value of 0.1 to 0.5 electrons. Furthermore, given that the 99% threshold increases from within 1 electron to within 7 electrons, the quantity of pixels that have a value more excellent than 1 electron is much greater; therefore, magnitudes, 18.8 and 19.0, are further obscured by this noise.

8.5.2: One Electron and above Mean Noise

Sub-electron mean noise has been shown to have an impact limited to magnitudes of 18.6 and above. So, if the observational targets have a magnitude at or lower than this range, then sub-electron mean noise distributions are suitable to reduce target loss. Next, the step is to investigate the limiting magnitudes for distributions with mean values of 1 electron and greater.

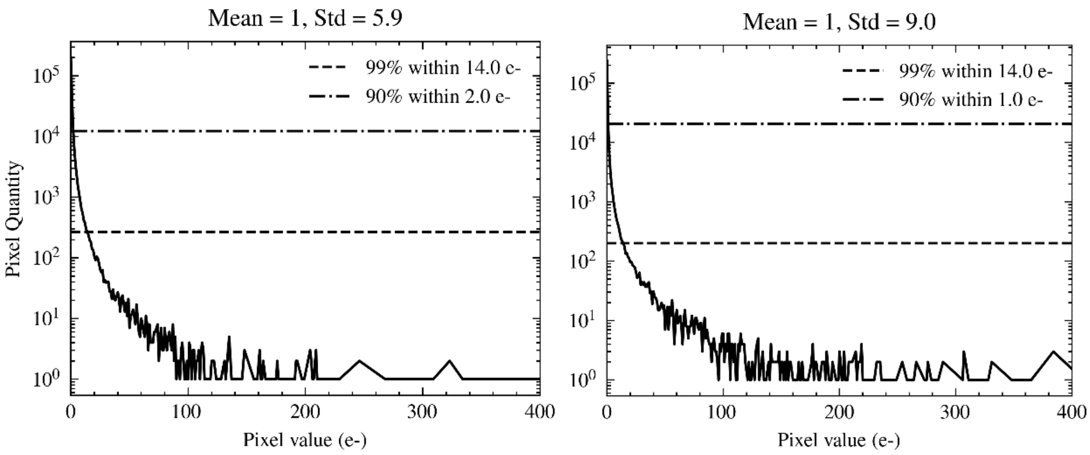

Figure 8-11 – Same pixel value distribution format as Figure 8-7 with a noise distribution with mean 1 e- and std 5.9 e- rms (left) and 9.0 e- rms (right) .

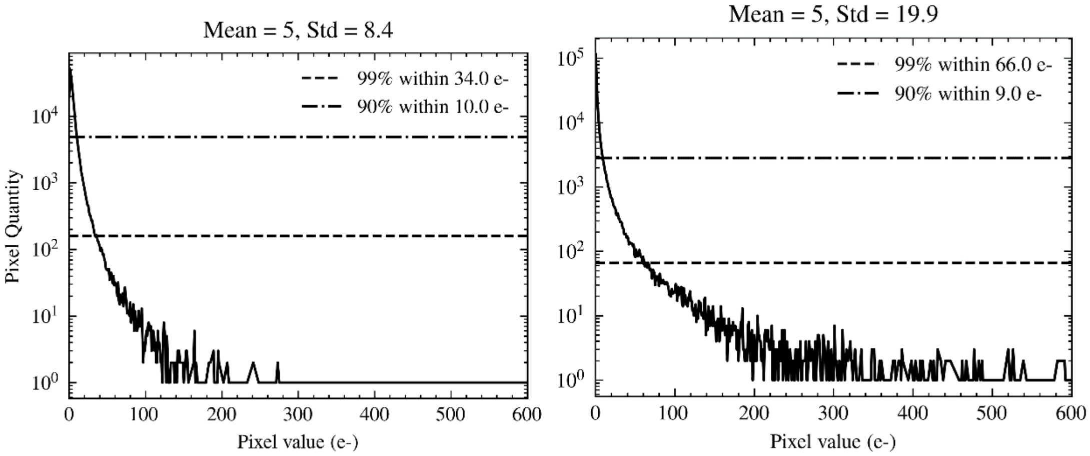

Figure 8-12 - A noise distribution with mean 5 e- and std 8.4 e- rms (left) and 19.9 e-rms (right).

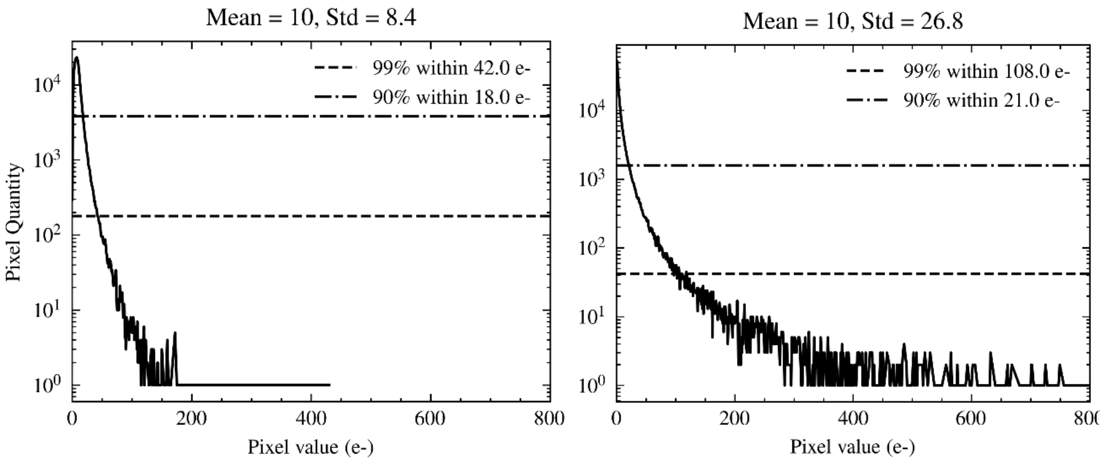

Figure 8-13 - A noise distribution with mean 10 e- and std 8.4 e- rms (left) and 26.8 e-rms (right).

Figure 8-11 shows that increasing the mean value from sub-electron to one electron has increased the signal within the 99% threshold from 6/7 electrons to 14 electrons. The increase in the standard deviation generation array shows some increase in the overall standard deviation of the noise distribution produced. However, the 99% threshold remains in the exact location.

Figure 8-12 shows, with a mean noise of 5 electrons, a significant increase of noise at the 99% to 34 electrons. Additionally, 90% of all the noise lies between 0 and 10 electrons, with a significant peak of about 3. With many pixels exhibiting noise between 1 and 10 electrons, it is expected that magnitudes impacted by these noise levels will become more heavily obscured.

There is a noticeably longer tail in Figure 8-12 right-hand plot due to the higher standard deviation. This is because the 99% signal threshold has increased from 34 to 66 electrons, while the 90% threshold has remained unchanged. This lack of increase in the 90% threshold means that the number of pixels contributing to bin values above 9 electrons has changed very little, but the number of bins occupied has increased. So effectively, more pixels exhibit higher noise than in the left-hand distribution.

Compared to Figure 8-12’s mean and std values, Figure 8-13 shows a notable rise in signal at both 90% and 99% thresholds. Also, there is a prominent peak about the 10 electron measurement. This, combined with the previous noise plots, shows that increasing the mean noise will prevent observing higher magnitude stars. In contrast to previous plots, the right-hand plot shows a substantial increase in the 99% threshold compared to the left-hand plot, increasing by 66 electrons. Whereas 90% increases by three electrons. The tail length increases from a maximum value of 160 electrons up to 780 electrons, increasing the potential impact of these tail pixels on much brighter objects.

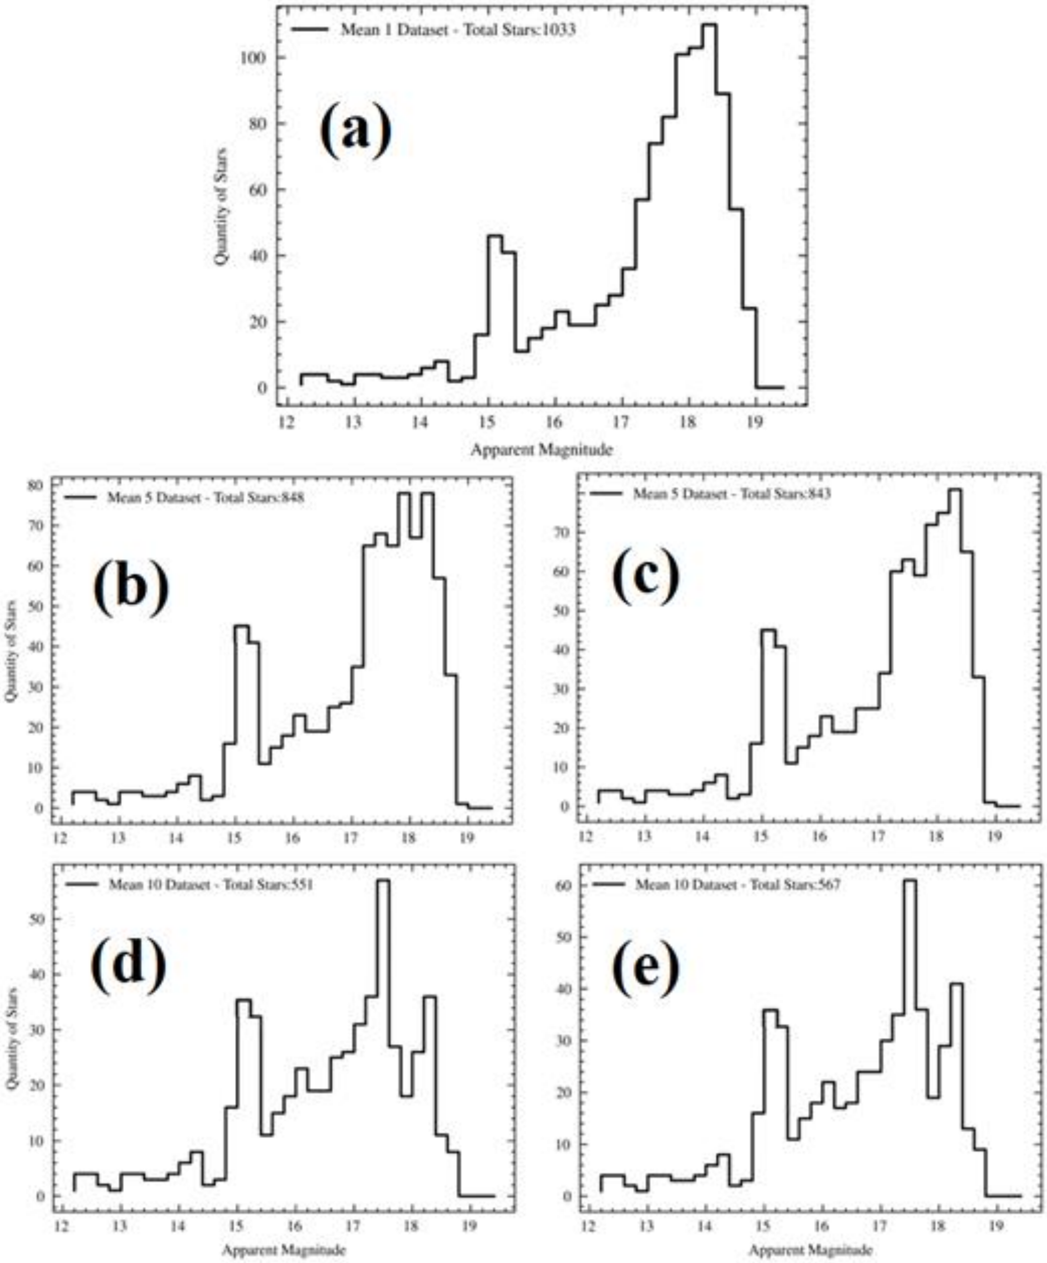

Figure 8-14 - Magnitude plots for mean distribution values 1e- to 10e-.

Figure 8-14 shows that as the mean electron noise increases, there is a decrease in the number of stars for a more extensive range of magnitudes. This is especially notable if Figure 8-14(a) is compared to Figure 8-14(d), which sees substantial loss at magnitudes 17.4 and upwards; the peak about magnitude 18.0 in Figure 8-14(a) is lost, showing that there is considerable impact on these higher magnitudes if the mean electron noise increases from one electron to ten electrons. A direct star comparison between each mean electron noise can be seen in Figure 8-15 and Figure 8-16.

Increasing the readout noise, as is done in Figure 8-14(c) and Figure 8-14(e), shows a decrease in star quantities at lower value magnitude bins, although this decrease is not substantial. A direct comparison can be seen in Figure 8-17 and Figure 8-18.

8.6: Summarising Star loss vs Mean Noise

Figure 8-7 to Figure 8-14 show that as mean noise increases, the number of stars lost increases. The results from these two figures have been summarized in Figure 8-15 and Figure 8-16, respectively.

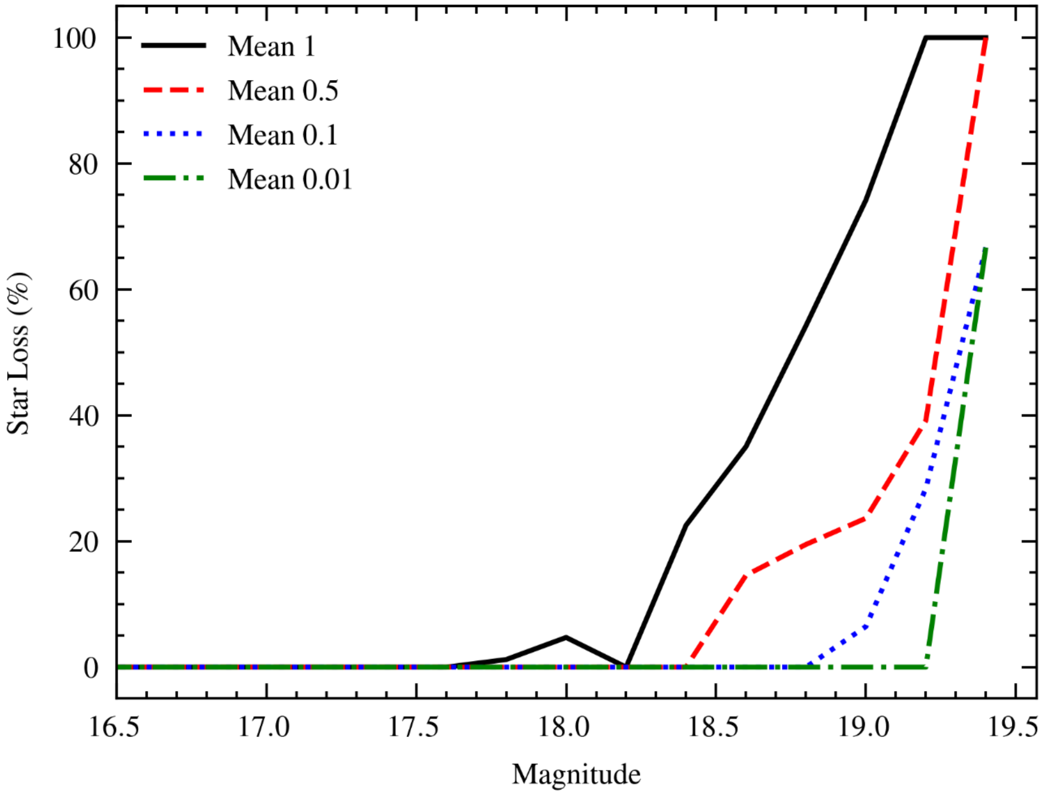

Figure 8-15 - Star loss versus magnitude for distributions with means ranging from 0.01 e- to 1 -e.

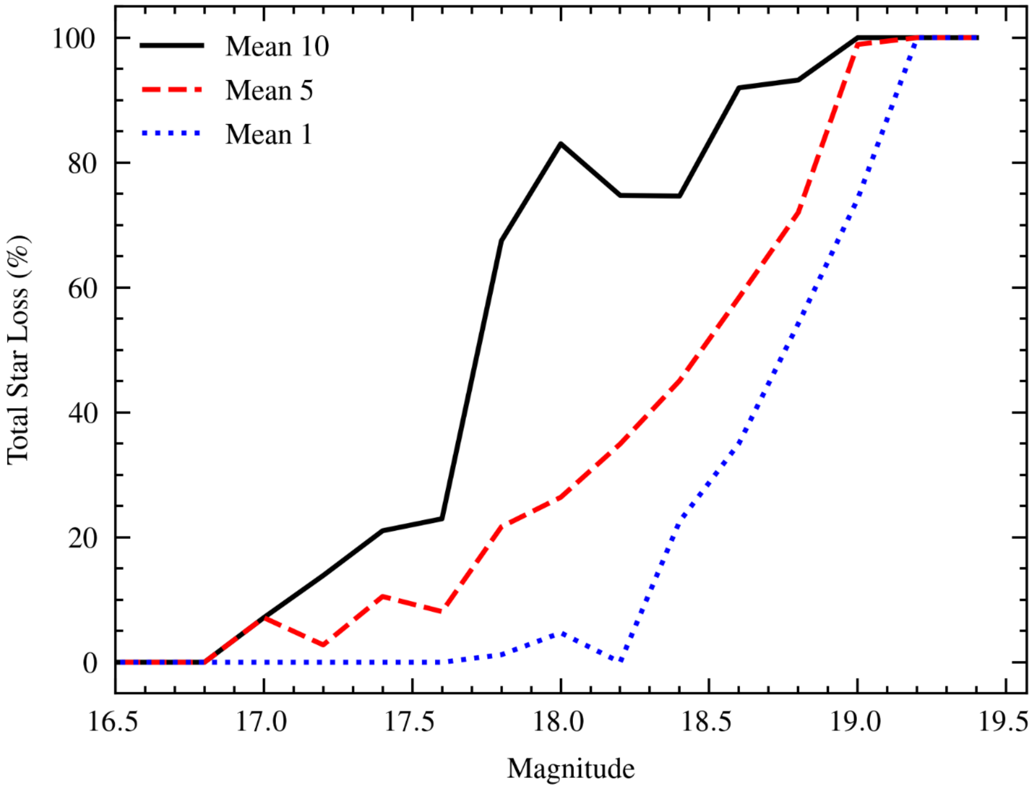

Figure 8-16 - Star loss versus magnitude for distributions with means ranging from 1 e- to 10 -e.

Both Figure 8-15 and Figure 8-16 indicate that an increase in the noise distribution's mean does cause an increasing loss of stars and stars of a lower magnitude becoming obscured by the noise. To prevent the loss of objects at these magnitudes, the mean noise of the CIS device must be as low as possible. To compete with EMCCDs, the noise ideally needs to be below 1 electron mean as these mean noises show only a small to tiny loss of stars in the faintest magnitudes of 18.4 and upwards, especially mean noises of 0.1 and 0.01.

Increasing the standard deviation of each distribution does not appear to have notable impacts until the noise is above 1 electron.

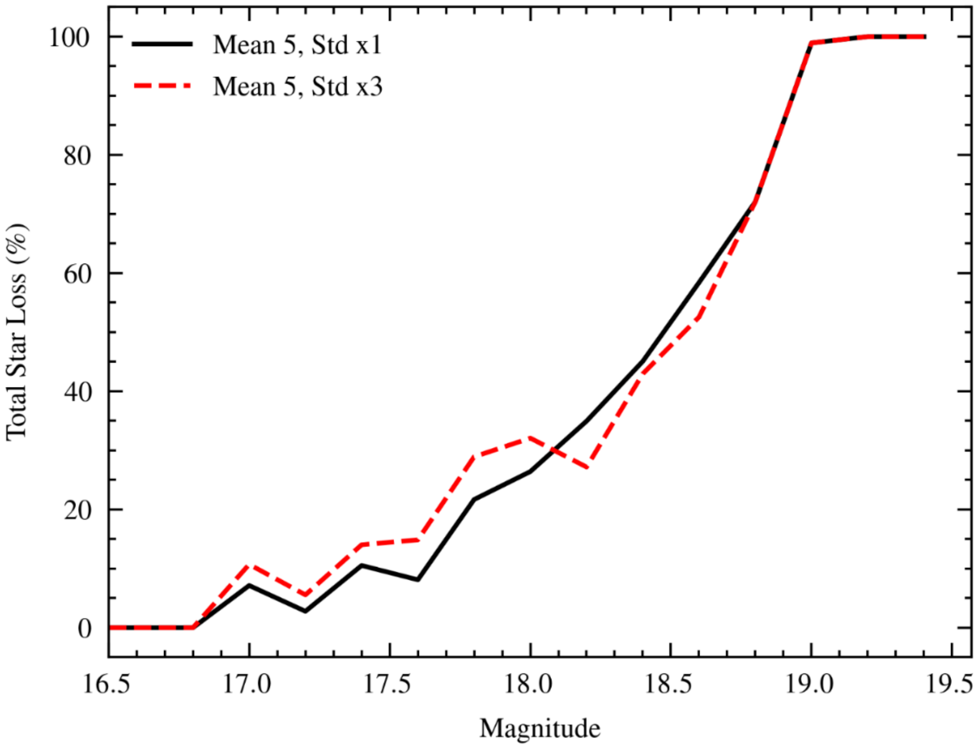

Figure 8-17 - Difference between two distributions at mean five e-, one with an increased 3x standard readout noise.

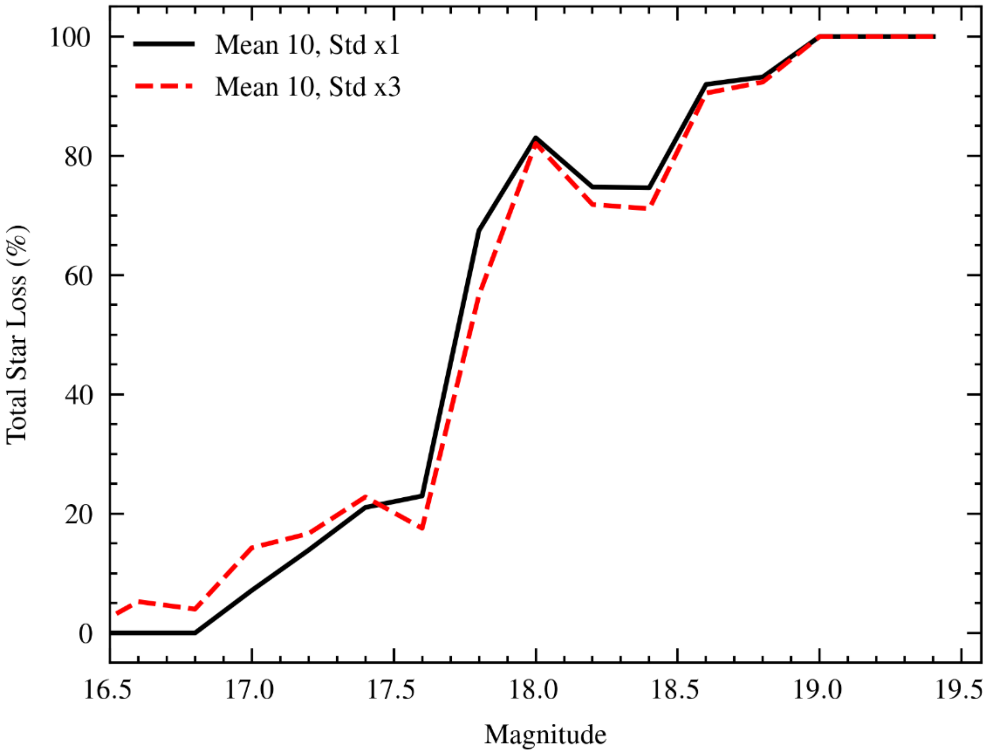

Figure 8-18- Difference between two distributions at mean ten e-, one with an increased 3x standard readout noise.

Both Figure 8-17 and Figure 8-18 show a slight increase in the star loss at lower magnitudes, around the 17.4 magnitude mark for Figure 8-17 and 17.0 for Figure 8-18. This additional loss results from an increase in noise output from a small population of pixels.

The additional star loss is unideal; however, it only appears to manifest at higher mean noise thresholds. Increasing the standard deviation source array at more minor mean noises has little or no notable impact on the star loss.

Figure 8-18 shows a much more significant impact on the total amount of star loss, with minor increases shown by changing the standard deviation. This shows that controlling the mean noise of the device is vital to improving the quantity of star identification and reaching fainter, higher magnitude targets. Ideally, to compete with EMCCDs, the mean total noise should be sub-electron. If CIS devices can achieve this signal to noise ratio, this would allow for new instruments to be able to benefit from CIS characteristics that EMCCDs lack, such as on-chip digital readout.

Both Figure 8-17 and Figure 8-18 also show that the CIS-like tail does not appear to have a large impact on the total amount of star loss, which is good because this suggests that controlling the individual pixel readouts is not as important as originally assumed in order to achieve good observational results.

Chapter 9: Conclusion and Future Work

The purpose of this thesis is to investigate the impact of noise from different detector technologies on the scientific output of GravityCam, for which the objectives and requirements have been discussed in Chapter 5. To make a high-precision microlensing survey possible from a ground-based instrument, high-frame observations and a large field-of-view is necessary. This means that upcoming sCMOS devices are becoming preferential compared to the more established EMCCDs. The thesis investigates how and where spurious charge is generated in EMCCDs, where noisy pixels on CMOS devices appear and the impacts of CIS-like noise on a simulated astronomical scene. Looking at these areas should aid in comparing EMCCDs to CIS devices, thus helping to inform the choice of detectors for GravityCam.

In Chapter 6: it is shown that spurious charge in the CCD97, does have a preference for generating near to the edges of the image and storage areas, as seen in Figure 6-12, and the impact of this is an increase in noise of 2.22 DN or 0.24 e^- as shown in Chapter 6.4:. If we compare this noise quantity in electrons to the CMOS mean noise added in Chapter 8, which is also measured in electrons, we can show the impact of the CIC measured in Chapter 6: on stellar detection rates. Although this seems a small contribution, we do see in Figure 8-13, which has a CMOS-like noise mean of 0.1 e^-, that some stars at the faintest magnitude, 19 and upwards, are lost, so there is evidence to suggest that CIC contribution at this level can have a potential impact on very faint targets.

There are some methods for reducing the CIC generated during the device operation, as shown in work by (Bush, 2021) and (Daigle O. Q., 2010), so there is potential for EMCCDs to reach even lower mean noise levels.

Chapter 7: shows that there is no evidence for pixels with high standard deviation, i.e. noisy pixels, to have a spatial preference in CIS, and they do not show any correlation to other local noisy pixels. The main impact of these noisy pixels appears to be the same as the average pixel noise contribution. However, in about 5% of the frames, the mean impact of these noisy pixels is notably different to the average pixel, as seen in Figure 7-13.

Given that the contribution from these noisy pixels appears to be spatially random across the image area, frame stacking should be an effective method for reducing the impact of these noisy pixels, because 5% of frames where the mean impact of noisy pixels is notably higher will be reduced by averaging with the remaining 95% of frames where the noisy pixel contribution is not notably greater.

Image stacking will be less effective if there is some spatial preference for noisy pixel generation. Although there is no evidence that any such preference exists for the CIS115 test device used, it would be informative for a future study to investigate other CIS devices to see if noisy pixel generation does or does not have any spatial preference.

Chapter 8: demonstrates the potential impact of CIS-like noise on counting stars within a dense stellar region. From the test data from the 1.5m Danish telescope that is used as input for the simulation, the impacts are restricted to magnitudes 18.4 and above at sub-electron noise levels. It should be noted that these magnitude limits are specific to the Danish telescope. Other telescope sizes will have different magnitude limits due to differences such as larger aperture size allowing for greater light collection, which leads to being able to view fainter magnitudes. However, compared to mean value of one electron rms and above, there is a notable impact up to magnitude 16.8. Furthermore, many faint background stars are wholly lost at these magnitudes for this particular globular cluster used in this test data, as seen in Figure 8-16. This shows that to be a viable alternative to current EMCCD detectors, any potential CIS device must have a readout noise below 1 electron for optimal results, especially for observing faint targets like lunar mass signals in microlensing surveys.

Overall, it appears that the main goal for any CIS to become a viable choice for the GravityCam instrument is to achieve sub-electron readout noise to prevent the loss of magnitudes above 19.0. As a secondary concern, the contribution from CIC for EMCCDs and noisy pixels for CIS could have impacts on targets fainter than 19.0; however, this should be further investigated using real detectors collecting actual on-sky data for more accurate measurements, which was not achievable for this project due to external circumstances.

9.1: Future Work

For future studies, there are several avenues. First, it would be very informative to test several different CMOS imagers with on-sky data to confirm the findings in chapter 8.5:, for a more varied set of device parameters.

Additionally, the reduction software used in Chapter 8: is not designed to be used this way and is designed explicitly for EMCCDs. So similar software should be designed with CIS devices in mind, i.e. not needing to reduce CIC contribution.

There are upcoming CIS that report having sub-electron readout noise. Once cameras using such devices become commercially available, characterising the noise, frame rate and QE performance for an application in the GravityCam instrument would be extremely useful.

The findings in Chapter 7 should also be compared with other CIS devices to see if the random nature of noisy pixels is consistent across device platforms, using the same methods outlined in this chapter.

References

Aftab, M. C. (2018). Adaptive Shack-Hartmann wavefront sensor accommodating large wavefront variations. *Optics Express, Vol 26. I. 26.* , 34428 - 34441.

Airy, G. B. (1835). On the Diffraction of an Object-glass with Circular Aperture. *Transactions of the Cambridge Philosophical Society, Vol 5, Part III.*, 283-291.

Andor. (2022, 08 01). *sCMOS Cameras for Physical Science & Astronomy*. Retrieved from Andor Oxford Instruments: https://andor.oxinst.com/products/scmos-for-physical-science-and-astronomy

Arazia, R. (2008). Image Referencing Using Fast Fourier Transform. *International Journal of Intelligent Technologies and Applied Statistics*, 127-158.

Bramich, D. M. (2008). A new algorithm for difference image analysis. *MNRAS, 386, 77*.

Buscher, D. (2016). Simulating large atmospheric phase screens using a woofer-tweeter algorithm. *Optics Express, Vol. 24. Issue 20. pp. 23566-23571*.

Bush, N. H. (2021). Measurement and optimisation of clock-induced charge in electron multiplying charge-coupled devices. *Journal of Astronomical Telescopes, Instruments and Systems*.

Castro, G. (2017, Feb). *This Noise Will Keep You Up at Night.* Retrieved from Analog Dialogue: https://www.analog.com/en/analog-dialogue/raqs/raq-issue-138.html

Chang, W. D. (2018). The challenge of sCMOS image sensor technology to EMCCD. *Proc. SPIE 10697, Fourth Seminar on Novel Optoelectronic Detection Technology and Application*.

Corle, T. R. (1996). Confocal Scanning Optical Microscopy and Related Imaging Systems. *Academic Press*, 1-66.

Cropper, M. e. (2013). Defining a weak lensing experiment in space. *Monthly Notices of the Royal Astronomical Society, Volume 431, Issue 4*, 3103–3126.

Daigle, O. Q. (2010). The darkest EMCCD ever. *Proc. SPIE 7742, High Energy, Optical, and Infrared Detectors for Astronomy IV.*

Daigle, O. T. (2018). Preliminary characterisation results of a large format 4k x 4k EMCCD. *SPIE 10709, High Energy, Optical, and Infrared Detectors for Astronomy VIII.*

Davies, R. H.-B. (2021). MICADO: The Multi-Adaptive Optics Camera for Deep Observations . *The Messenger, vol. 182*, 17-21.

Fano, U. (1947). Ionization Yield of Radiations. II. The Fluctuations of the Number of Ions. *Phys. Rev. 72.*, 26-29.

Fried, D. (1978). Probability of getting a lucky short-exposure image through turbulence. *Journal of The Optical Society of America*.

Grant., W. N. (1973). Electron and Hole Ionization Rates in Epitaxial Silicon at High Electric Fields. *Solid-State Electrons, Vol 16.*

Guenther, B. S. (2018). *Encyclopedia of Modern Optics, 2nd Edition.* Academic Press.

Harpsøe, K. B. (2012). High Frame-rate Imaging Based Photometry, Photometric Reduction of Data from Electron-multiplying Charge Coupled. *Astronomy & Astrophysics*, 4.

HDF Group. (2022, 8 26). *The HDF5 library & file format.* Retrieved from The HDF Group: https://www.hdfgroup.org/solutions/hdf5/

Hippler, S. (2019). Adaptive Optics for Extremely Large Telescopes. *Journal of Astronomical Instrumentation, Vol. 8, No. 2*.

Hippler. S. Bergfors, C. B. (2009). The AstraLux Sur Lucky Imaging Instrument at the NTT. *The Messenger, 137.*, 14-17.

Holst, G. C. (2011). *CMOS/CCD Sensors and Camera Systems, Second Edition.* SPIE Press.

Hynecek, J. (1992). CCM - A New Low-Noise Charge Carrier Multiplier Suitable for Detection of Charge in Small Pixel CCD Image Sensors. *IEEE Transactions On Electron Devices, Vol. 39 No. 8.*, 1972 - 1975.

Hynecek, J. N. (2003). Excess Noise and Other Important Characteristics of Low Light Level Imaging Using Charge Multiplying CCDs. *IEEE TRANSACTIONS ON ELECTRON DEVICES, VOL. 50, NO. 1*.

Ingley R., S. D. (2009). Life testing of EMCCD gain characteristics. *Nuc. Inst. Meth. Vol A600.*, 460-465.

Janesick, J. (2001). *Scientific Charge-Coupled Devices.* Bellingham, Washington: SPIE—The International Society for Optical Engineering.

Janesick, J. A. (2006). Fundamental performance differences between CMOS and CCD imagers: Part 1. *Proc. SPIE 6276, High Energy, Optical, and Infrared Detectors for Astronomy II*.

Janesick, J. C. (1987). Flash Technology For Charge-Coupled-Device Imaging In The Ultraviolet. *Optical Engineering*.

Janesick, J. E. (2015). Fundamental performance differences of CMOS and CCD imagers: part VI. *Proc. SPIE 9591, Target Diagnostics Physics and Engineering for Inertial Confinement Fusion IV*.

Janesick, J. E. (2017). Performance and design differences between PMOS and NMOS CMOS imagers. *Proc. SPIE 10209, Image Sensing Technologies: Materials, Devices, Systems, and Applications IV*.

Jerram, P. P. (2001). The LLCCD: low-light imaging without the need for an intensifier. *Proc. SPIE 4306, Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications II*.

Johnson, J. (1927). Thermal Agitation of Electricity in Conductors. *Nature*, 50–51.

Koch, D. G. (2010). Kepler Mission Design, Realized Photometric Performance, and Early Science. *ASTROPHYS J LETT, 713 (2)*, 79-86.

Kolmogorov, A. (1991). The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Proc. R. Soc. Lond. A434.*, 9-13.

Lakshminarayanan, V. a. (2011). Zernike polynomials: a guide. *Journal of Modern Optics Vol. 58, No. 7.*, 545-561.

Lebreton, J. M. (1992). An overview of the Cassini mission. *https://doi.org/10.1007/BF02506708*, 1137–1147.

Lesser, M. P. (1994). Improving CCD quantum efficency. *Proc. SPIE 2198, Instrumentation in Astronomy VIII*.

Littlefair, S. (2022, June 27). *L10: Adapative Optics*. Retrieved from Observational Techniques for Astronomers: http://slittlefair.staff.shef.ac.uk/teaching/phy217/lectures/telescopes/L10/index.html

Low, J. K. (2008). Band Gap Energy in Silicon. *American Journal of Undergraduate Research*.

Mackay, C. D.-R. (2019). GravityCam: Wide-field Imaging Surveys in the Visible from the Ground. *arXiv:1911.07544*.

Madan, S. B. (1983). Experimental Observation of Avalanche Multiplication in Charge-Coupled Devices. *IEEE*.

Mao, S. P. (1991). Gravitational Microlensing by Double Stars and Planetary Systems . *Astrophysical Jopurnal Letters v.374*.

Massey, R. e. (2013). Origins of weak lensing systematics, and requirements on future instrumentation (or knowledge of instrumentation). *Monthly Notices of the Royal Astronomical Society, Volume 429, Issue 1*, 661–678.

McAulay, A. (2000). Generating Kolmogorov phase screens for modeling optical turbulence. *Proc. SPIE 4034, Laser Weapons Technology*.

Nyquist, H. (1928). Thermal Agitation of Electric Charge in Conductors. *Phys. Rev 32.*, 110-113.

Paczynski, B. (1986). Gravitational Microlensing by the Galactic Halo . *Astrophysical Journal v.304*.

Prusti, T. e. (2016). The Gaia mission. *Astronomy & Astrophysics, Volume 595*.

Riedl, M. (2001). *Optical Design Fundamentals for Infrared Systems, 2nd Edition.* Bellingham, WA: SPIE Press.

Robbins, M. S. (2003). The Noise Performance of Electron Multiplying Charge-Coupled Devices. *IEEE TRANSACTIONS ON ELECTRON DEVICES, VOL. 50, NO. 5*.

Shockley, W. (1949). The Theory of p-n Junctions in Semiconductors and p-n Junction Transistors. *Bell System Technical Journal, 28: 3,* 435-489.

Skottfelt, J. B. (2015). The Two-Colour EMCCD Instrument for the Danish 1.54m Telescope and SONG. *A&A*.

Skottfelt, J. e. (2014). Searching for variable stars in the cores of five metal-rich globular clusters using EMCCD observations. *Astronomy & Astrophysics, Vol. 573, A103*.

Smith, B. A. (1981). Encounter with Saturn: Voyager 1 Imaging Science Results. *Science (212.4491.163),* 163-191.

Smith, W. S. (1970). Charge coupled semiconductor devices. *The Bell System Technical Journal, Vol. 49*, 587-593.

Stefanov, K. D., & Dunford, A. a. (2018). Electron Multiplying Low-Voltage CCD WIth Increased Gain. *IEEE Transactions on Electron Devices, Vol 65 No. 7.*, 2990-2996.

Strehl, K. (1895). Aplanatische und fehlerhafte Abbildung im Fernrohr. *Zeitschrift für Instrumentenkunde*, 362-370.

Sze, S. M. (1969). *Physics of Semiconductor Devices.* New York: John Wiley & Sons.

Teledyne e2v. (2021, May 12). *Sensors Overview CCD97.* Retrieved from Teledyne Imaging: https://www.teledyneimaging.com/en/aerospace-and-defense/products/sensors-overview/ccd/ccd97/

Tubbs, R. (2003). Lucky Exposures: Diffraction Limited Astronomical Imaging Through the Atmosphere. doi:https://doi.org/10.17863/CAM.15991

Wambsganss, J. (2006). Gravitational Microlensing. *ArXiv: astro-ph/0604278*.