Exploring the impact of detection physics in X-ray CCD imagers and spectrometers

Thesis

Exploring the impact of detection physics in X-ray CCD imagers and spectrometers

Thesis submitted for the degree of

Doctor of Philosophy

at The Open University

by

David James Hall

e2v centre for electronic imaging

Planetary and Space Sciences Research Institute

The Open University

22nd June 2010

Exploring the impact of detection physics in X-ray CCD imagers and spectrometers

David James Hall

Abstract

This thesis is concerned with exploring the way in which the physics of the detection process affects the quality of a CCD-based X-ray detector system. The physical processes which lead to the final images and spectra achieved with a CCD-based camera system are investigated through a combination of simulations and experimental techniques with the aim of improving the detector performance and allowing future detectors to be designed with optimal characteristics. Techniques developed throughout the study and the results of the simulations have wide-ranging impacts on the areas concerned. The study is split into two main sections, the first regarding a high-resolution, high-energy, photon-counting X/γ-ray camera. In medical imaging, X-rays and gamma-rays are often used for the purposes of diagnostic imaging. In many synchrotron based research programmes, such as protein crystallography and X-ray diffraction imaging, X-rays are used, once again, for imaging purposes. In both cases, a high-resolution detector with a high frame-rate is required such that images can be taken with a spatial resolution of the order of micrometers to tens of micrometers. If one is able to distinguish the energy of the incident X-rays and gamma-rays (with energies of 20-200 keV) then these spectral capabilities add to the functionality of the detector, allowing, for example, the removal of fluorescence X-rays. Chapter 2 reviews the relevant detector physics and theory before providing a critical review of current gamma-cameras. Chapter 3 outlines the feasibility study for the scintillator-coupled EM-CCD detailing the development of a new energy discrimination methodology. Also described is the development of a full system simulation which can be used to troubleshoot problems

found when calibrating and optimising the device. Chapter 4 details the characterisation and optimisation of the detector making use of the aforementioned simulations where appropriate. Chapter 5 presents the results of the study, showing how the resolution can be dramatically improved and how energy discrimination can be implemented. The second section of the thesis regards instrument background. The use of CCDs for space borne X-ray detection in scientific satellites is wide-spread. Whilst in-orbit the CCDs are subjected to an incident flux of high energy particles. These particles may be detected, both as the primaries themselves and by means of secondaries produced in the detector shielding, and will produce a background level formed by components indistinguishable from the X-rays for which the mission was designed to detect. Chapter 6 presents an introduction to the theory behind the instrument background experienced by CCD-based detector systems in orbit. A simulation has been developed which is in very good agreement with data received from the spacecrafts, described in Chapter 7. Finally, Chapter 8 summarises the outcomes of these studies and provides insight into future work which will further aid the improvement of gamma-cameras for medical imaging and synchrotron-based research and will allow future CCD-based camera systems to be designed for increased sensitivity in-orbit.

Declaration

I hereby declare that no part of this thesis has been previously submitted to this or any other university as part of the requirement for a higher degree. The work described herein was conducted solely by the undersigned except for those colleagues and other workers acknowledged in the text.

David James Hall

22nd June 2010

Dedication

To my mother, father and Coralia.

Acknowledgements

I would like to thank many people for their help and support over the duration of my thesis work, without whom this study would not have been made possible.

Thanks to my supervisors Andrew Holland, David Smith and George Seabroke for their technical support and advice, and for guiding me along the right path.

Thanks to David Burt for his teaching and support on the CCD and for all the fascinating discussions which subsequently followed.

Thanks to Mark Robbins, Peter Pool and Bill Bruns at e2v technologies for their support on the devices and their operation.

Thanks to Martin Turner for his insightful discussion on instrument background at the outset of this study and for setting me on my way with this project.

Thanks to all the remaining members of the e2v centre for electronic imaging, past and present, for their advice, guidance, friendship and entertainment: Neil Murray, Tom Greig, Amir Intisar, Jason Gow, Billy Kinchin, Ben Dryer, James Tutt, Anthony Evagora, Richard Harriss, Richard Ingley, John Osmand, Ross Burgon, Karen Guyler, Gopal Jeyesundra and Ian Hutchinson.

I would like to gratefully acknowledge the financial support received from the Science and Technologies Facilities Council and e2v technologies during my research studentship.

Finally, a big thank you to my family and to Coralia for their continued support and enthusiasm and for putting up with my endless ramblings whilst I was thinking through the finer details of the study, and thanks to Donald for his company whilst composing this thesis.

List of acronyms

ADU	Analogue-to-Digital Unit
BI-CCD	Back-illuminated Charge-Coupled Device
CCD	Charge-Coupled Device
CDS	Correlated Double Sampling
DN	Digital Number
EM-CCD	Electron-Multiplying Charge-Coupled Device
ESF	Edge Spread Function
FI-CCD	Front-illuminated Charge-Coupled Device
FOP	Fibre-optic plate
FWHM	Full Width Half Maximum
HEO	Highly elliptical orbit
LEO	Low Earth Orbit
LSF	Line Spread Function
MTF	Modulation Transfer Function
PTC	Photon Transfer Curve
Rϕ2HV	Rϕ2 High Voltage
SRIM	Stopping and Range of Ions in Matter (Software)

Contents

Chapter 1: Introduction

1.1 Secondary event detection in CCD X-ray detectors

In many areas of space and terrestrial imaging, the Charge-Coupled Device (CCD) is the detector of choice when high resolution imaging or spectroscopy using X-rays or gamma-rays is required. The 'primary events' may be defined as those which one is aiming to detect. In this case, the detectors are designed to allow the detection of the X-rays or gamma-rays incident on the imaging devices. In most cases, however, one finds that the detector is also subjected to a number of secondary events, many of which are indistinguishable from the primary events. These secondary events may range from K-shell fluorescence X-rays to the passage of high energy protons through the detector. Although some unwanted signal may be removed through event processing techniques, many secondary events cannot be removed and have a detrimental effect on the camera resolution or sensitivity.

This thesis aims to investigate the impact of secondary events in CCD-based X-ray detectors through a study of the physical processes involved. The study of secondary events is often ignored in order to concentrate on the improvements that are possible through the investigation of the primary events. Through a better understanding of the physics behind and the impact of secondary events in a detector, however, one is able to further improve the imaging and spectral capabilities of a CCD-based camera system.

The first section of this study involves the investigation of the impact of secondary event detection in a scintillator-coupled Electron Multiplying CCD (EM-CCD). Current devices based around similar technology are capable of achieving a spatial resolution (FWHM) of approximately 80-120 µm in integrating mode. Through the use of a novel event detection technique this resolution has been improved to 60-80 µm, with some studies claiming a FWHM of 50 µm (although this occurs to the detriment of the spectral capabilities of the detector).

In theory, it should be possible to improve both the spatial and spectral capabilities of the scintillator-coupled EM-CCD over the results found in the literature - the physical limits on the system are currently far from being met. Through the use of thorough analysis using various novel techniques, this study aims to provide evidence that a high-resolution, hard-X-ray detector is possible using this technology. The study aims to initially determine the improvements which may be made over current devices through the analysis of the interaction chain in the scintillator and the impact this has on the spatial and spectral results. Secondly, the results of this study should provide an insight into the requirements for the design of an optimal detector for use in commercial applications. The designs implied from the results of this study are presented for further work following the future availability of faster drive electronics and larger area EM-CCD camera systems.

In the second case described in this study, the X-rays forming the primary events are those in space being detected in orbiting camera systems such as the XMM-Newton and Swift EPIC cameras (Sections 7.6.1 and 7.6.3). Whilst the camera systems are designed for the optimal detection of low energy X-rays, the camera design has a large impact on the background continuum found in the results. If the background is too high then the primary X-ray events may be buried within the background signal. The 'in field of view' background experienced by a CCD in-orbit consists of two main components, the Cosmic X-ray Background (CXB) and the instrument background, with the instrument background dominating the X-ray spectrum above approximately 2 keV. Taking the XMM-Newton mission as an example, pre-launch predictions of the instrument background under-estimated that experienced in-orbit by nearly two orders of magnitude (Sections 6.9 and 7.7.6). The increase in instrument background experienced in-orbit has serious implications on the sensitivity of the camera systems.

Through a better understanding of the components forming the instrument background the background may be more accurately predicted for future missions. Both the mission orbit and the detector design have a major impact on the instrument background detected by the

camera systems in-orbit and it is hoped that this study will aid in the design process for future missions. Through the better understanding of the formation of the instrument background, considerations can be made for future missions in order to improve the quality of the science obtained.

Throughout this study, the real data, whether taken in-orbit or in the laboratory, have been analysed and compared with data produced using various Monte Carlo simulation techniques. Through the analysis of real data alongside the simulated data, it is possible to gain an insight into what is occurring 'behind the scenes' in a way which is not possible experimentally.

1.2 Research goals

The work carried out for this thesis comprises of two main sections, as mentioned above, both relating to the investigation of the physical processes behind the detection of secondary events in X-ray CCD detectors. The first part of this study involves the design, testing and simulation of a high resolution, hard X-ray imager. The detection efficiency of the EM-CCD to hard X-rays is very small and for this reason a scintillator is required to convert the hard X-rays into light of a lower energy. At this lower energy the EM-CCD has a very high detection efficiency. The inclusion of the scintillator significantly complicates the detection process and previous studies have encountered many difficulties in imaging at a high resolution whilst retrieving information on the energy of the incident X-ray.

Through the careful characterisation and optimisation of the EM-CCD the basic system can be improved. Using many Monte Carlo simulations of various components of the system, this study aims to make large advances in the performance for the high-resolution, hard X-ray detector. The mathematical theory of scale-space is developed through this study in order to optimise the performance for the individual case of Gaussian-like event detection in a series of two-dimensional images. In depth study of the energy component of the detected events aims to provide a method of improving the quality of the images produced. The problems

encountered with the detection of K-shell X-ray fluorescence generated in the scintillator itself can be dramatically reduced through the study of these secondary events.

The second part of the thesis involves the study of the instrument background found in X-ray CCDs used for space missions in-orbit. The instrument background is formed from the detection of secondary events indistinguishable from the detection of primary X-rays. This work aims to use similar Monte Carlo simulation methods to study the formation of the spectral continuum of the instrument background for three individual missions. The study aims to provide a more thorough understanding of the way in which the instrument background continuum is formed and, as a consequence, allow the detector design and orbital parameters to be considered with respect to their implications on the camera sensitivity.

Throughout this work, Monte Carlo based simulations have been used to study the physics of the generation and detection of secondary events in X-ray CCDs, with many simulation methods developed in tandem between the two projects. The combination of experimental study and simulation work allows the secondary events to be explored in a much greater detail than may be found in previous literature and the improvements in the performance of the high-resolution, hard X-ray imager and in the understanding of the implications of detector design and orbit on the instrument background continuum may be seen throughout this thesis.

1.3 Thesis organisation

This thesis is organised into seven further chapters. Chapters 2-5 detail the high resolution, hard X-ray imaging achieved with a scintillator-coupled EM-CCD and the influence of secondary interactions. Chapters 6-7 describe the development of a thorough understanding of the instrument background for the XMM-Newton, Swift and Suzaku missions and the implications this background has on current and future missions through careful detector design and orbital considerations. The chapter contents are summarised in more detail below.

Chapter 2 reviews the theory behind the scintillator-coupled EM-CCD study. The chapter begins with an overview of the detector physics, working from the CCD through to the more complex EM-CCD. A basic knowledge of the CCD is assumed before following the theory behind the EM-CCD, with the reader directed to, for example, Janesick (2001) and Sze (1981) for further reading on CCDs and semiconductor physics respectively if required. The chapter continues with an explanation behind the process of scintillation and describes the properties of the scintillator of choice in this study: thallium doped caesium iodide, CsI(Tl). A review follows in which hard X-ray imagers are discussed from their early beginnings using arrays of photo-multiplier tubes through to the more recent use of CCDs. Various methods used in the literature are described, with further investigation of these methods left to Chapter 3. The chapter finishes with a brief history of the Scale-Space theory, a theory originating from the Fourier and Wavelet transforms in which event profiles in an image can be analysed.

Chapter 3 describes the preliminary study to the experimental work carried out. The chapter details Monte Carlo simulations implemented to show the requirement for a structured scintillator as opposed to a solid volume of scintillator. The theory of scale-space is discussed in more detail before a deeper investigation into the mathematics is instigated to specialise the technique for this particular study. A further Monte Carlo study, based in the Geant4 toolkit, is then described in which the passage of the lower energy photons (generated in the scintillator) to the EM-CCD is modelled in order to demonstrate the problems which must be overcome in order to improve the detector. A study of various energy discrimination techniques follow, detailing the advantages and disadvantages of the various methods found through the literature and with further simulation techniques. The chapter continues with a simulation of the operation of the EM-CCD in order to provide suitable images (with known characteristics) for future testing of the techniques. The simulation of the EM-CCD output is also used in the optimisation and characterisation of the device, allowing various problems encountered in the optimisation process (detailed in Chapter 4) to be solved. The chapter

finishes with a study of the various centroiding techniques which may be used to aid the image analysis.

Chapter 4 presents the characterisation and optimisation techniques used to achieve a detector with suitable low noise performance to be capable of photon-counting. The challenges encountered are described and a detailed explanation is given on the methods used to overcome the various problems encountered. The chapter is completed with initial results from the X-ray imaging process, verifying the origin of the events and describing the various techniques used in the preliminary image processing. The point-search function is described with reference to the optimisation process used to achieve the best results with the detector.

In Chapter 5 the imaging and spectral results are presented and discussed in turn, with particular reference to the influence of the secondary interactions in the scintillator and their indistinguishability from real X-ray events. The detailed analysis of the influence of fluorescence on both spatial and spectral results provides a new insight into the problems faced in high-resolution hard X-ray imaging and provides solutions to counter the detrimental effects. The improvements in spatial resolution and energy discrimination are presented with reference to those found in the literature in Chapter 2. A look to the future of the scintillator-coupled EM-CCD as a high resolution, hard X-ray detector is also presented. The results of the study described thus far indicate the main areas of improvements required and here is outlined a proposed detector design and test plan in order to further improve performance.

Chapter 6 begins the second part of the study into the physics of secondary event detection for space-borne CCD-based X-ray spectrometers. Instrument background is discussed from the basic theory through to the individual processes causing the detection of secondary events, with these events also classed as ‘X-ray like’ events.

Chapter 7 describes the Monte Carlo simulation developed to replicate the instrument background found in orbit. The simulation was implemented through the Geant4 toolkit, with

particular attention paid to the process of detection. The results are then analysed and discussed, providing a new and more complete understanding of instrument background formation. The implications of the detector design and orbital parameters are then clearly defined in order to allow the greater optimisation of future detectors to be used in-orbit.

The final chapter, Chapter 8, concludes the study and presents a brief analysis of the outcomes of the experimental and simulation work.

1.4 Publications

The results from the work in this thesis are partly contained within the following publications:

1. D. J. Hall, A. D. Holland

   "Space radiation environment effects on X-ray CCD background"

   Nuclear Instruments and Methods: A, **612** (2010), 320-327.

2. D. J. Hall, A. D. Holland

   "The use of automatic scale selection to improve the spatial and spectral resolution of a scintillator-coupled EMCCD"

   Nuclear Instruments and Methods: A, **604** (2009) 207-210

3. D. J. Hall, A. D. Holland, N. J. Murray

   "High-resolution X-ray and γ-ray imaging using a scintillator-coupled electron-multiplying CCD"

   Proc. SPIE **7449** (2009) 74491G.

4. D. J. Hall, A. D. Holland, D. R. Smith

   "Imaging and spectroscopy using a scintillator-coupled EMCCD"

   Proc. SPIE **7021** (2008) 70211Z.

5. D. J. Hall, A. D. Holland, M. Turner

“Simulating and reproducing instrument background for X-ray CCD spectrometers in space”

Proc. SPIE **7021** (2008) 70211Y.

6. D. J. Hall, A. D. Holland, M. J. L. Turner

“Modelling instrumental background for CCD X-ray spectrometers in space”

Proc. SPIE **6686** (2007) 66860J.

Chapter 2: Detector physics

2.1 Introduction

The Charge Coupled Device (CCD) has many uses, from the simple digital camera to the latest space telescopes. The work carried out here relates to the scientific CCD, and, for this reason, this chapter gives some details relating the operation of the CCD (specifically the Electron-Multiplying CCD), how it works and the problems which can be encountered. The majority of the experimental work detailed in this thesis is based around the EM-CCD. The charge multiplication aspect of the EM-CCD adds several potential problems to the device operation. These difficulties are discussed in theory here, with the solutions used in this study described later in Chapters 4 and 5. The basic theory behind the CCD and its operation is assumed as prior knowledge for the reader; only the most relevant theory with regards to this study is presented here. Further reading may be found in Janesick (2001) amongst many other sources if required.

The detector described in this study makes use of a scintillator in order to convert X-rays and gamma-rays into multiple photons in the visible region of the electromagnetic spectrum. Thallium doped caesium iodide is the scintillator of choice in this study and as such, special attention is paid to inorganic scintillator crystals and the scintillation processes involved.

The penultimate section of the chapter, Section 2.4, presents a review of current X-ray/gamma cameras and discusses the troubles experienced in previous attempts to produce a high-energy, high-resolution spatial imager with energy discrimination. The choice between imaging and spectral analysis presents a balance which must be carefully considered when designing the optimal detector.

When an X-ray or gamma-ray interacts in the scintillator, the visible photons produced can be imaged. The events produced in the image, each relating to X-ray and gamma-ray interactions

in the scintillator, require careful profiling for which many methodologies are available. The final section of the chapter details the history of the scale-space method from its origins in the Fourier transform and wavelet transform and gives a brief introduction to one method which may be used to profile these events. The particular scale-space formulation used in this study is described later in Chapter 3.

The penultimate two chapters of this thesis, Chapters 7 and 8, describe work carried out to reproduce the spectral continuum of the instrument background as found in orbit by three missions (XMM-Newton, Swift and Suzaku). The theory behind 'instrument background' and its formation are discussed in Chapter 7 as this is felt to lead more fluidly into the simulation methodology and results described in Chapter 8.

2.2 The Charge-Coupled Device (CCD)

The full theory behind the Charge-Coupled Device (CCD) is not described in this work and a basic level of understanding of the CCD and its operation is assumed of the reader. The theory behind the CCD and the operation of such a device is described in many literature sources, with the reader directed towards, for example, Janesick (2001) and Sze (1981) for further reading on the charge-coupled device and solid-state physics respectively if required. Several aspects of the CCD and its operation are essential to later sections of this work and as such are mentioned briefly below.

2.2.1 Dark current

Charge carriers may be thermally excited from the valence band to the conduction band, adding to both the size of the charge packet in a pixel and the noise on the charge packet. The thermal generation of charge carriers is defined by Poisson statistics such that the variance on the generated charge is equal to the number of charge carriers generated. The thermally generated charge carriers are produced even when the CCD is in complete darkness. The charge generated per unit time is defined as the 'dark current' or 'dark signal'.

The quantity of thermally generated charge added to the signal charge passing through the CCD can be minimised by operating the CCD at a faster rate (the charge is generated at a certain rate) or by reducing the temperature (reducing the rate of thermal excitation).

There are three main regions of dark current generation in a CCD: the neutral bulk material beneath the depleted region, the depleted material within the depleted region and the Si-SiO_2 interface states (Janesick, 2001). The generation of dark current is dominated by the Si-SiO_2 interface states, where a miss-match in the lattice parameters can lead to the creation of a population of states across which electrons can be promoted from the valance band to the conduction band (Figure 2.1). This dominant 'surface' component of the dark current can be reduced through the operation of the device in 'inverted mode'.

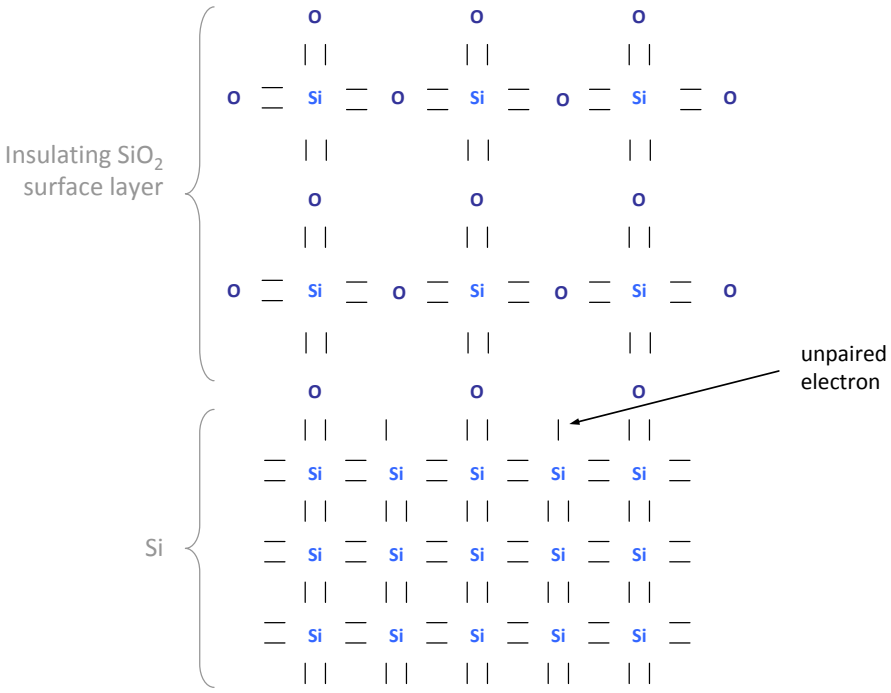

Figure 2.1. The silicon to silicon dioxide interface contains many unpaired electrons because of the lattice structure miss-match.

2.2.1.1 Inverted Mode Operation (IMO)

The substrate bias of the CCD (V_{SS}) may be increased above that used under standard operating conditions. If the bias is increased above a certain level (approximately +7-8 V for most e2v devices), holes will accumulate at the silicon surface causing an increase in surface recombination. The accumulation of holes at the surface effectively changes the surface silicon from n-type to p-type, 'inverting' the surface layer. This accumulation of holes will act to suppress the dark signal generation at the Si-SiO_2 interface (typically from ~1 nA cm^{-2} to ~10 pA cm^{-2}). The remainder of the dark signal is now due to traps in the underlying silicon at a much lower value (typically at ~1 % of the interface generation rate).

2.2.2 Clock-Induced Charge (CIC)

The onset of inverted mode operation is determined by the difference between the substrate voltage and the voltages applied to the electrodes. Taking an example V_{SS} for a standard CCD of +9 V running in inverted mode, with a clock voltage swing from 0 V to 10 V, with the clock voltages high (10 V) the electrode is taken out of inversion and the free holes will move to the less positive areas. In more technical terms, the lattice miss-match shown in Figure 2.1 leads to many active traps between the valence and conduction bands. Whilst the traps near the valence band remain full and those near the conduction band have a low probability of capture, those near the middle of the band are active. In the pinned condition there is a high concentration of holes at the surface, causing the trapped electrons to recombine with the holes. This leaves the traps in an empty state giving no excess dark signal. When the clocks change such that the silicon is placed back into normal depletion, the electrons start to refill the traps. This refilling is not instant and is subject to exponential based time constants, leading to a slow build of dark-current (hence the dependence of the CIC on the clock timings).

When the holes pass through the region of higher field at the surface between the electrodes (10 V to 0 V over a short distance, with the highest field being at this surface close to the electrodes) the process of impact ionisation may occur (see Sze, 1981, Section 1.6.4 of the 2nd

edition), with each hole being multiplied generating an electron which may then be collected as dark signal. The process of impact ionisation may be considered as simply the acceleration of an electron as it passes through a high electric field, causing the electron to effectively slam into the lattice and release further electrons. The resulting generation of electrons is called 'clock induced charge' (CIC). Although the rate of generation is relatively low, the large number of clock transfers involved in transferring the charge through the CCD can lead to a large component of CIC. The rate of generation of CIC is dependent on the clock amplitude, and as found later in this work, the significantly higher voltages used in the gain register of the EM-CCD (Section 2.3) can lead to a dominance of this component in the EM-CCD (Section 4.4).

2.2.3 Noise sources

The main noise sources encountered when imaging with a CCD are detailed below and include the shot noise on the signal (including dark current and CIC), the reset noise and the transistor noise. The readout noise is discussed briefly in preparation for the more detailed discussion of the Electron-Multiplying CCD in the following section, where the problems found with the readout noise may be dramatically reduced.

2.2.3.1 Shot noise

The detection of photo-generated electrons is a random process with the number of photon-generated carriers as the random variable. The standard deviation of this random variable is the input photon (or shot) noise. It is a property of the photon source and not of the detector. The variation in the emission of photons from the source will affect the number of electron-hole pairs produced. This may also be applied to leakage current, also given by Poisson statistics on the number of electrons generated.

The number of electron-hole pairs N_{eh} produced by an X-ray photon of energy E_x is given by:

$$N_{eh} = \frac{E_x}{\omega} \tag{2.1}$$

where ω is the energy required to create a single electron-hole pair. The value of ω is dependent on temperature, but is approximately 3.65 eV for silicon. The band gap of silicon is approximately 1.12 eV, but the energy ω required is that to make an electron hole pair and the corresponding phonons. The number of photons emitted is described by a Poisson distribution and hence has variance (N_{noise}) and noise (σ) defined by:

$$N_{noise} = N_{eh} = \sigma^2 \tag{2.2}$$

2.2.3.2 Johnson noise and KTC noise

Fluctuations in charge across a resistor due to the random thermal motion of electrons lead to a fluctuation in the voltage. This voltage fluctuation is known as 'Johnson noise' (Johnson, 1928). The voltage noise (rms.) is given by the square root of $4kTRB$, where B is the bandwidth of frequencies passing through the resistor, R is the resistance, k is Boltzmann's constant and T is the temperature in Kelvin. KTC noise (Barbe, 1975) is one specific consequence of Johnson noise.

When a capacitor (with capacitance C) is placed in series with a resistor, the -3dB bandwidth is given by $1/2\pi CR$, with the noise bandwidth given by $\pi/2$ times the standard bandwidth. When the capacitor is disconnected from the resistor, the fluctuations in the resistor are held on the capacitor with a noise V_n of $\sqrt{kT/C}$ volts rms. on the capacitor. The noise on the charge of the capacitor is given by CV_n, giving a noise on the capacitor of $\sqrt{kTC}$ coulombs rms., the KTC noise.

2.2.3.3 Reset noise

It is necessary to reset the voltage across the output in order to set a reference level for each charge packet coming from subsequent pixels. In resetting this voltage across a capacitance one is effectively sampling and holding the noise in any series resistance. This leads to

fluctuations known as 'reset noise' (Figure 2.2), a form of *kTC* noise as described above. These fluctuations in the reset level can be removed through the use of Correlated Double Sampling (CDS) where a measurement is taken of the reset level for comparison with the signal level.

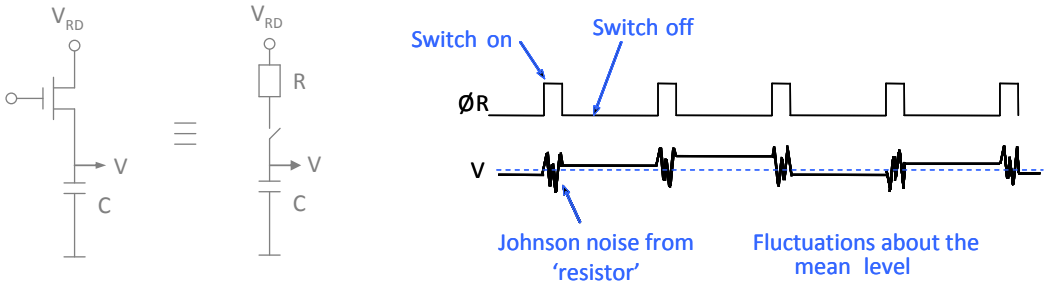

Figure 2.2. Reset noise, shown right, for the simple circuit, shown left (adapted from Burt, 2006), relating to the reset of the CCD.

2.2.3.4 Transistor noise

The thermal energy of the electrons can lead to random fluctuations of the source-drain current. These fluctuations in the source-drain current lead to fluctuations in the output known as transistor noise (Figure 2.3).

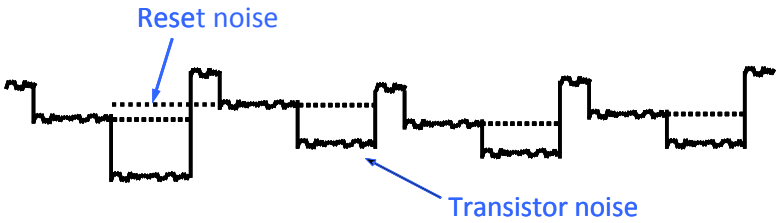

Figure 2.3. The reset noise and transistor noise are shown on a diagrammatic CCD output with this noise component added across the transistor at the output.

2.2.3.5 Readout noise

At readout, the three main components of noise experienced are the shot noise on the signal packet, the reset noise and the transistor noise. The shot noise cannot be reduced and is defined by the statistics of the signal packet. The reset noise can be removed through the use of Correlated Double Sampling (CDS). Using CDS, the signal is sampled first after the reset and again after the signal has been deposited in the output node. The difference between these two values gives a measure of the signal level without relying on low reset noise giving a constant 'zero' level. The transistor noise cannot be removed (although the effect can be reduced) and defines the overall noise performance of the device, usually quoted as noise-equivalent-signal in electrons rms. The noise in electrons rms. gives the equivalent number of electrons noise that would equate to the fluctuations in the signal level caused by the transistor noise.

Although the readout noise cannot be removed, the effect of this noise on the signal-to-noise ratio can be reduced through the use of an Electron-Multiplying CCD (EM-CCD), detailed in the following section.

2.3 The Electron-Multiplying CCD (EMCCD)

When charge moves through a region of silicon over which a high electric field is present, such as in the case of a high voltage placed over an electrode in the CCD, the process of 'impact ionisation' may occur. Through this processes, in the most basic terms, charge carriers (electrons) can lose energy by the creation of more charge carriers, thus multiplying the signal for the same noise (increasing the signal-to-noise ratio).

In the mid 1990s, e2v carried out experiments to show that this multiplication in signal could be used in CCDs, giving rise to the 'Low-Light-Level' (L3) devices, Electron-Multiplying CCDs, which are now commercially available (Jerram *et al.*, 2001). Initial problems involving clock-

induced charge and the subsequent multiplication of the dark current have been overcome and the devices may be used in conditions where the incident photon flux is low.

2.3.1 Low-light level (L3) detection

The Low-Light-Level camera uses an Electron-Multiplying CCD (EM-CCD) to provide 'all light level' imaging from bright sunlight down to shadowed overcast starlight. The variable gain allows the device to be run as a standard CCD in sunlight with unity gain down to very low light levels. The multiplication of the signal before readout effectively reduces the readout noise which can be reduced to the sub-electron level (Mackay *et al.*, 2001). The very low readout noise allows the detection of signals of only a few photons, signal levels which would otherwise be lost beneath the readout noise of a standard CCD.

EM-CCDs have many uses, from 24-hour surveillance to military imaging at night. The flexibility of the camera has led to the technology being exploited in night-time surveillance, scientific imaging and medical imaging. One region of low-light-level imaging which has not been fully exploited involves the imaging of individual hard X-ray interactions in a scintillator. This novel 'photon-counting' technique forms the basis of the first part of this thesis.

2.3.2 Impact ionisation

When an electron passes through a region of high-electric field it can be accelerated. If an electron passes through a region of high electric-field in silicon then the electron can gain sufficient kinetic energy to effectively slam into the lattice, breaking the silicon-silicon covalent bonds, the process of 'impact ionisation' (Janesick, 2001). This process occurs naturally in the silicon when an electron-hole pair is generated through thermal means or at a lattice defect. The generated electrons can in turn gain enough kinetic energy to break further bonds, creating extra electron-hole pairs and causing further impact ionisation.

Although this process can generate spurious signal through the formation of electron-hole pairs from non-signal electrons, the process can be controlled through the application of

specific voltages to a specially designed gain structure in the CCD. The probability of impact ionisation increases as the electric-field increases in magnitude. Varying the voltage applied to the multiplication electrode structure alters the electric-field and allows the gain process to be controlled.

At thermal equilibrium, in the absence of an applied electric field, carriers both emit and absorb phonons (quanta of energy relating to modes of vibration) and the net exchange of energy is zero. In the presence of an electric field the carriers gain energy from the field. With an increase in electric field, the average energy of the carriers increases and they can be said to have an effective temperature which is higher than that of the bulk lattice. The difference in temperature (or effective temperature) between the lattice and carriers determines the rate of absorption/emission of phonons by the carriers (see Sze, 1981, Section 1.6.4 of the 2nd edition, for further details).

As the temperature of the lattice (in this case, the temperature of the EM-CCD) increases, the ionisation coefficient (responsible for impact ionisation) decreases due to increased lattice scattering through the emission and absorption of phonons by the carriers and lattice. This acts to impede the ability of the carriers to attain sufficient energy to trigger impact ionisation (Marsland *et al.*, 1991).

The impact ionisation responsible for the multiplication gain in the EM-CCD gain register can be seen to decrease with increasing temperature for the CCD97 (Section 2.3.5), with measured results shown in Chapter 5.

2.3.3 Structure

The EM-CCD shares the same basic structure as a standard frame-transfer CCD (Figure 2.4). The main difference between the EM-CCD and those described previously is the addition of a gain register following on from the standard readout register. A standard CCD (Figure 2.4) contains only the register elements (as seen directly below the store Section in Figure 2.5). In-

order to save space in the design of the CCD97 (one example of an e2v EM-CCD), the gain register is placed directly below the readout register using 16 corner elements to join the two registers (Figure 2.5, bottom right).

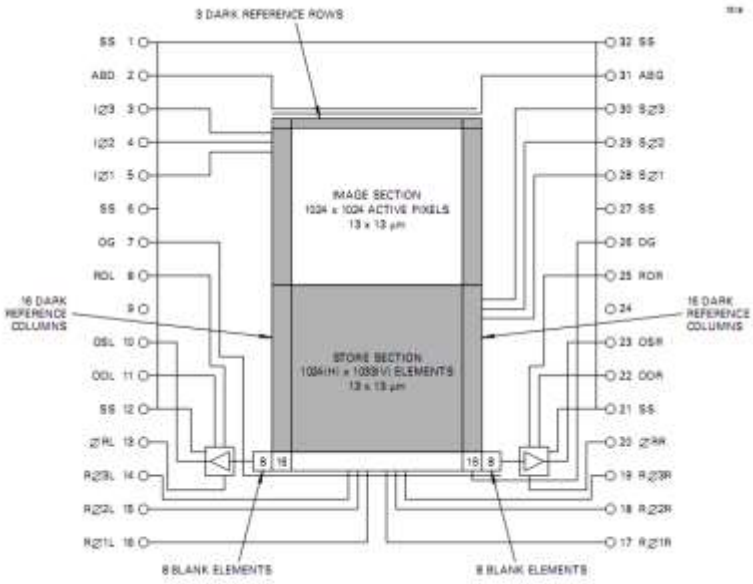

Figure 2.4. Standard frame-transfer device (CCD47-20 datasheet, 2006).

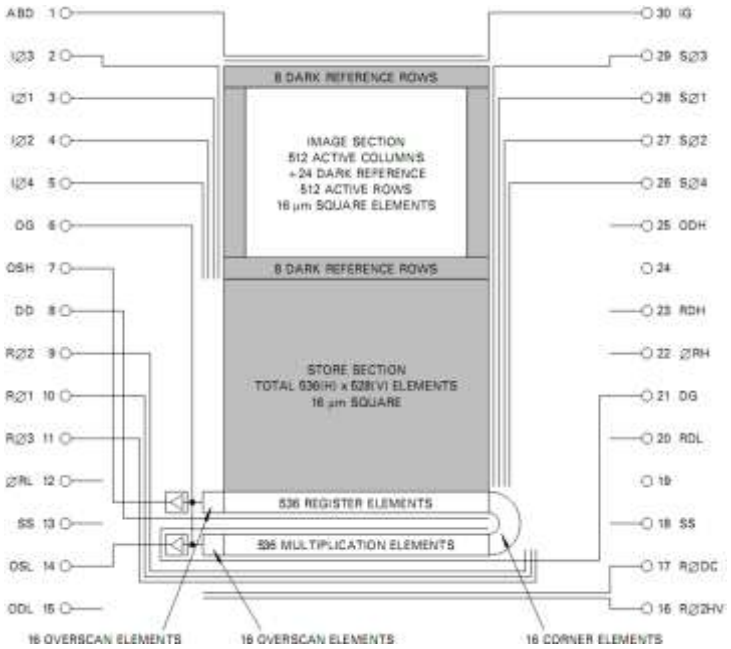

Figure 2.5. Schematic of the CCD97 low-light-level EM-CCD (CCD97 datasheet, 2004).

The standard CCD97 operates in frame-transfer mode and is available in both front-illuminated and back-illuminated forms. The back-illuminated version is used in this study due to the higher quantum efficiency at lower energies. The image section of the device consists of 512 × 512 pixels, with the addition of 24 dark reference columns and 16 dark reference rows, evenly split on each edge of the device. The dark reference regions are covered with a thin layer that is opaque to visible light and allow the image area signal to be compared to the dark current.

Although designed with four image and store electrode phases, the device can be run in two-phase mode, made possible due to the novel potential structure present in the device. Steps in potential as charge passes through the device prevent charge movement opposing the clocking direction, allowing two-phase operation.

2.3.4 The multiplication gain register

The gain register in the CCD97 operates in much the same way as the standard readout register. Phases Rϕ1 and Rϕ3 are common between the two registers. The high voltage, providing the electric field for the impact ionisation, is placed on an additional electrode structure, labelled Rϕ2HV (Rϕ2 High Voltage). Between Rϕ1 and Rϕ2HV is found an additional constant DC level. This acts as a potential barrier, separating the two phases. The avalanche multiplication process is optimised if the high voltage has been reached at Rϕ2HV before the passage of charge. To allow the high voltage to be reached before charge is transferred the DC level blocks the electrons from passing forward. Once the high voltage has settled, the first phase is taken low, allowing charge to flow over the DC barrier and giving the full impact of the voltage difference. It is worth noting that the electric-field is therefore not set by the voltage applied to Rϕ2HV, but to the difference between Rϕ2HV and the DC level (Figure 2.6). This value is referred to as the corrected Rϕ2HV value later in the study.

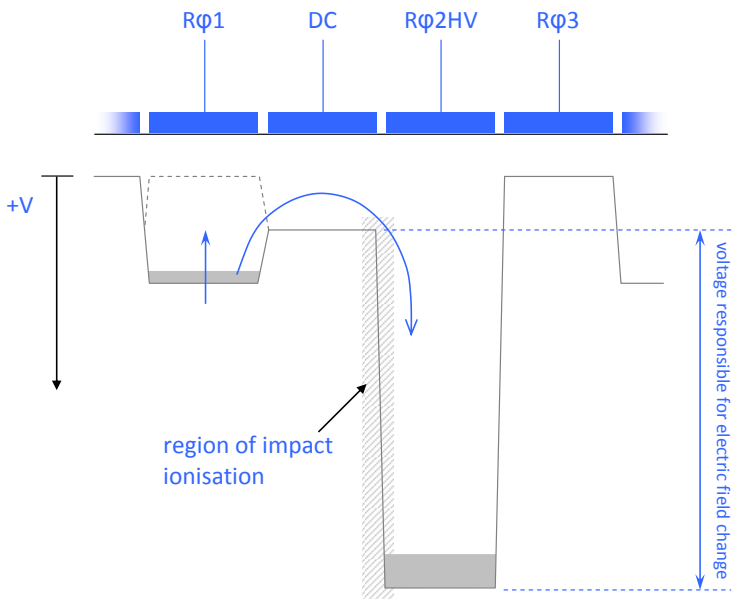

Figure 2.6. The gain register contains an extra DC level and replaces Rϕ2 with the high voltage phase Rϕ2HV. The DC level allows the high-voltage phase to be reached before transferring the charge over the barrier, giving full multiplication.

Through the variation of the high voltage applied, the gain process can be controlled. The probability of impact ionisation increases with applied voltage. For example, a gain probability of 1% per step over 512 gain elements will give an average total gain of approximately × 163 (1.01^{512}). The process of impact ionisation can only give an average multiplication across the gain registers, and as such, an extra noise factor is added to the signal (see below).

2.3.4.1 Shot noise on the gain process

The shot noise on the gain process is described here in detail as it is not encountered in conventional CCDs. Due to the stochastic nature of the multiplication gain, during the gain process each signal electron can be assumed to behave independently and may generate a

differing number of avalanche electrons. The excess noise factor, *F*, was defined in (Robbins *et al.*, 2003) as:

$$F^2 = \frac{\sigma^2_{n_{out}}}{G^2 \sigma^2_{n_{in}}} \tag{2.3}$$

where G is the total multiplication gain, σ^2_{nin} is the variance on the signal before the gain process and σ^2_{nout} is the variance of the output signal. The following results are described by Robbins *et al.* (2003) but are expanded upon here for completeness.

Defining the mean number of electrons before gain as n_{in} and the mean number of electrons after gain as n_{out}, and combining these with the mean total gain *G* and the mean variance on the gain σ^2_G, one arrives at Equations 2.4 and 2.5, assuming the mean gain is independent of the signal input.

$$n_{out} = G n_{in} \tag{2.4}$$

$$\frac{\sigma^2_{n_{out}}}{n^2_{out}} = \frac{\sigma^2_{n_{in}}}{n^2_{in}} + \frac{\sigma^2_G}{G^2} \tag{2.5}$$

Assuming that the multiplication probability is constant and that successive trials are independent, the events can be described by the Binomial distribution and the variance of the added electrons is given by Equation 2.6, where *g* is the probability of multiplication in each step.

$$\sigma^2_{added} = n_{in} g(1-g) \tag{2.6}$$

The mean variance on the gain can then be described by Equation 2.7.

$$\sigma^2_G = \frac{\sigma^2_{added}}{n^2_{in}} = \frac{g(1-g)}{n_{in}} \tag{2.7}$$

The variance on the number of electrons for each gain element can then be described through the combination of Equations 2.4, 2.5 and 2.7.

$$\sigma_{n_{out}}^2 = (1+g)^2 \sigma_{n_{in}}^2 + n_{in} g(1-g) \tag{2.8}$$

Applying Equation 2.8 to the first gain element, and noting that σ^2_{in} is equal to n_{in} by Poisson statistics, gives Equation 2.9.

$$\sigma_{n_{out}}^2 = (1+3g) n_{in} \tag{2.9}$$

Applying Equation 2.8 to the second gain element, noting that the new mean input signal is given by $n_{in}(1+g)$ and the mean variance on the new input signal is given by Equation 2.9, gives Equation 2.10.

$$\sigma_{n_{out}}^2 = (1+g)(2g^2+5g+1) n_{in} \tag{2.10}$$

Taking the mean input signal to element 3 of the gain transfer to be $n_{in}(1+g)^2$ and the mean variance on the new signal input to be given by Equation 2.10 gives Equation 2.11.

$$\sigma_{n_{out}}^2 = (1+g)^2 (2g^3+6g^2+7g+1) n_{in} \tag{2.11}$$

This series can be given for the N^{th} gain element by Equation 2.12, noting that the mean total gain, G, is given by $(1+g)^N$.

$$\sigma_{n_{out}}^2 = \frac{G n_{in}}{(1+g)}(2G+g-1) \tag{2.12}$$

With reference back to Equation 2.3 for the definition of F and inserting Equation 2.12 gives the final equation for the extra noise factor, F.

$$F^2 = \frac{(2G+g-1)}{G(1+g)} \tag{2.13}$$

When N is large then $G = (1+g)^N >> g$, leaving the approximation for the CCD97 (given in Equation 2.14), where there are 536 gain elements. The form of F^2 is shown in Figure 2.6, where F^2 is equal to unity for a total gain of 1 and approaches 2 for high gain.

$$F^2 = \frac{2G-1}{G} \tag{2.14}$$

This noise factor acts to increase the noise on any signal which undergoes multiplication gain by a factor of *F*. Although the noise on the signal and dark current are increased by a factor of *F* there is no effect on the subsequent readout noise. The impact of *F* must also be considered when taking the photon transfer curve (Section 4.4) at gains other than unity since the variance on the signal is no longer equivalent to the signal itself, but instead is equivalent to the signal multiplied by F^2.

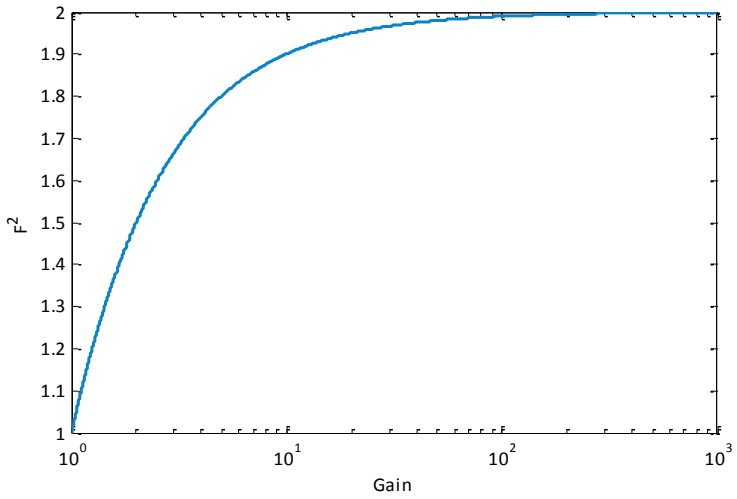

Figure 2.7. The excess noise factor, F^2, shown from Equation 2.14 for 536 multiplication elements (as in the CCD97).

2.3.5 CCD97

The CCD97 has been chosen as the detector to be used in this study. Full details about the device can be found in the device datasheet (CCD97 datasheet, 2004), but the relevant data are outlined here. The device is supplied in a ceramic package, as shown in Figure 2.8. The figure shows the image Section as the darker square (top right). The lighter regions of the device are the dark reference columns/rows and the covered store Section coated with ~1 µm of aluminium to stop the detection of visible photons in these regions.

The CCD97 has a peak in quantum efficiency between approximately 500 nm and 650 nm, Figure 2.9. This matches very well with the emitted spectrum of photons from the scintillator CsI(Tl), discussed in Section 2.4, which has a peak emission at 550 nm.

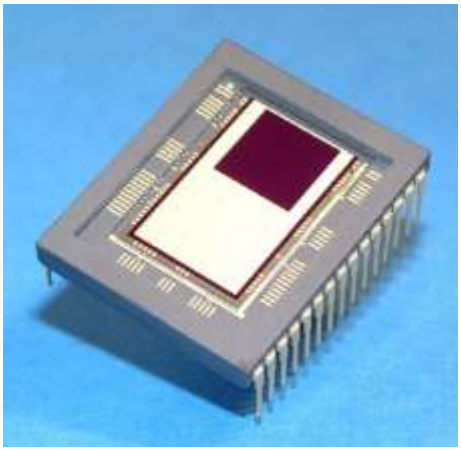

Figure 2.8. A back-illuminated CCD97 in its ceramic package. The dark square is the uncovered imaging area, whilst the lower, lighter area is the shielded store area (CCD97 datasheet, 2004).

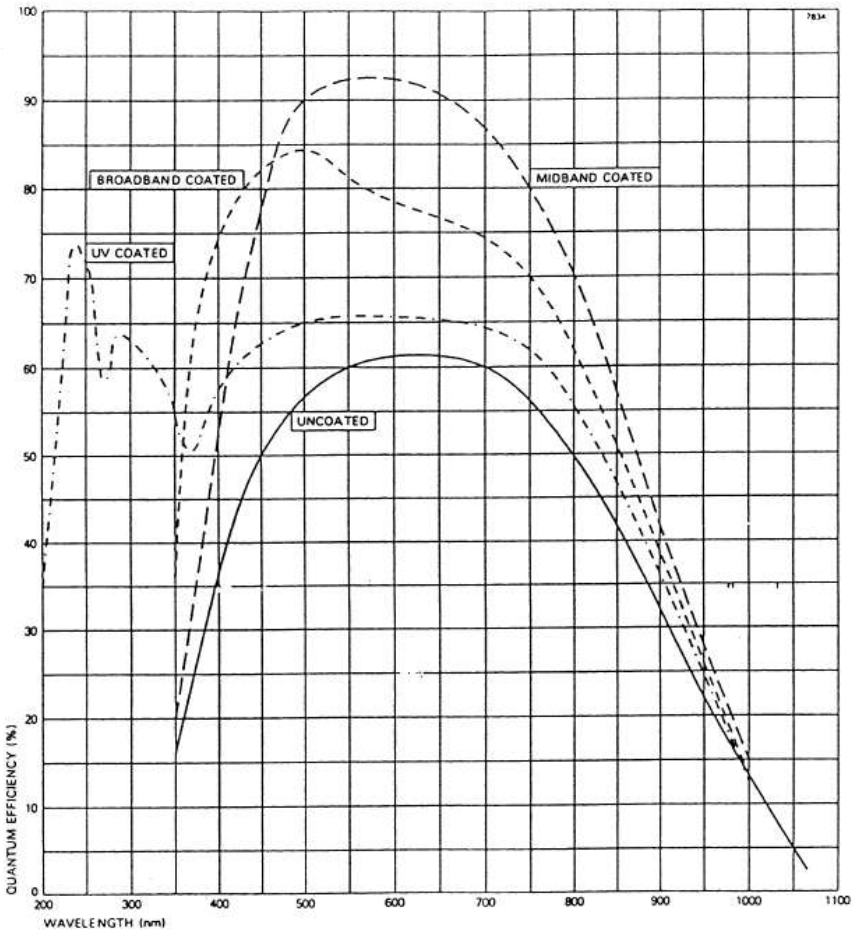

Figure 2.9. The quantum efficiency of the CCD97 at varying wavelengths. Note that the peak in quantum efficiency lies at approximately 550 nm, the peak in emission of CsI(Tl). The various coatings provide increased QE at different wavelengths. The device to be used in this study is mid-band coated to maximise the QE at 550 nm (private communication, David Burt).

The CCD97 has controllable gain through the variation of the high voltage applied to Rϕ2HV. The increase in gain with applied voltage is shown in Figure 2.10, taken from the device datasheet. The gain is seen to increase as the temperature decreases (as described in Section 2.3.2). The specific device used in this study has been characterised in Chapter 4, with more accurate measurements taken for the gain of the device under the operating conditions

used. Each CCD requires individual characterisation due to the differences between devices following production and due to the subtle variations in the operating conditions used.

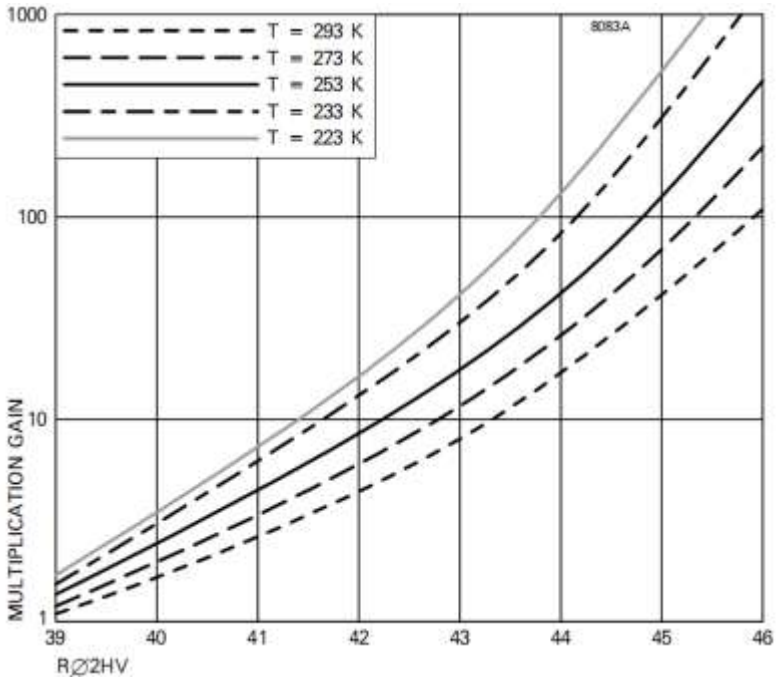

Figure 2.10. The avalanche gain can be seen to increase with applied voltage. The temperature dependence of the gain process shows an increase in gain with decreasing temperature (CCD97 datasheet, 2004).

2.4 X-ray scintillators

Scintillators have been used in the field of ionising radiation detection for over one hundred years. Solid scintillation was observed in 1903 by Elster and Geitel where individual light flashes were seen from a ZnS screen when placed near an alpha-emitting source (Elster *et al.*, 1903). Since then, the development of scintillators has provided many new uses throughout high-energy physics and astrophysics, along with the continued use in medical imaging, from the first X-ray film to the application of modern CCD dental imagers.

The scintillator converts the energy from absorbed ionising radiation into a flash of light of a longer wavelength, usually in the visible or UV region of the electromagnetic spectrum. Detectors based around scintillators are often run in a counting mode, counting the flashes of light and hence counting the incident particles and high energy photons, but this work concentrates on the use of scintillators for more detailed detection and analysis.

Current use of scintillators ranges from use in astrophysics through to Positron Emission Tomography (PET) and Single Photon Emission Computed Tomography (SPECT) in medical imaging. The use of scintillators in medical imaging is covered further in Section 2.5, with further comment on the possible use of scintillators for high-energy, high-resolution imagers in synchrotron research also covered. Here we concentrate on the principles of scintillation itself and the processes involved.

2.4.1 Inorganic scintillation process

The scintillation process acts to convert a high-energy quantum into many lower energy quanta. These lower energy quanta may then be detected with a much higher efficiency by a solid-state detector than the original high-energy quantum. The scintillation process can be described in five stages (Rodnyi, 1997):

1. Creation of electron-hole pairs through the absorption of ionising radiation.
2. Relaxation of primary e-h pairs, producing multiple secondary electrons, holes, photons, phonons and other electronic excitations.
3. Thermalisation of secondary e-h pairs through interactions with the vibrations of the environment.
4. Energy transfer to the luminescence centres.
5. Emission of energy from luminescence centres in the form of lower energy photons.

At the energies considered in this study, the photo-electric effect dominates (due to a larger interaction cross-Section than the Compton interactions, Figure 2.11). Through the photo-electric effect, a hole is created in the inner shell (K-shell). This leaves an ionised atom and a free electron with energy equal to hv minus the binding energy of the electron. The ionised atom in the lattice may relax through the emission of a photon, as another electron drops to fill the hole, or through the Auger effect, where further electrons are released. The electrons then lose energy through further scattering or the emission of photons. This process continues until ionisation is no longer possible. Electrons lose excess energy through inelastic scattering until only low energy excitations in the lattice remain.

When the energy of the excited electrons is below the ionisation threshold, the electrons begin to interact with vibrations in the environment: the process of thermalisation. The holes move to the top of the valence band, whilst the electrons move to the bottom of the conduction band, leaving many electron-hole pairs each separated by the band gap energy E_g. This stage leaves N_{eh} electron-hole pairs, Equation 2.15, where ε is the average energy required to create a single electron-hole pair ($\varepsilon \approx E_\gamma$).

$$N_{eh} = \frac{E_\gamma}{\varepsilon} \tag{2.15}$$

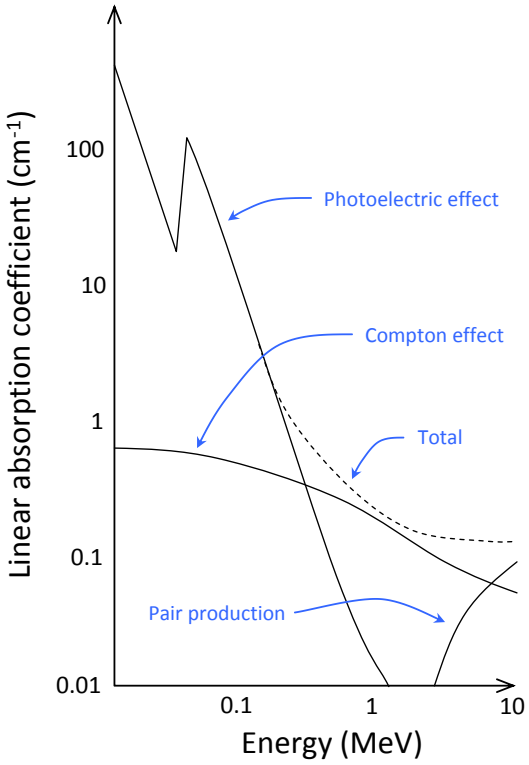

Figure 2.11. The linear absorption coefficient for CsI(Tl) against incident gamma-ray energy. The photo-electric effect dominates in the region of interest in this study (below 200 keV).

The electronic excitations then transfer energy to their final excitation state at the 'luminescence centres' (due to the potential structure), so called due to the emission of photons during relaxation from these excited states. If the electron/hole is close to a luminescence centre then the process is very efficient. If, however, the electron/hole is not close to a luminescence centre then the electron/hole can be captured by traps (caused by impurities and lattice defects). The emission of photons follows through Equations 2.16 and 2.17 (Rodnyi, 1997), where L^+ represents the luminescence centre.

$$L^+ + h \rightarrow L^{2+},\ L^{2+} + e \rightarrow h\nu \tag{2.16}$$

$$L^+ + e \rightarrow L^0,\ L^0 + h \rightarrow h\nu \tag{2.17}$$

2.4.2 Ideal properties of scintillators

When choosing a scintillator, many factors must be taken into consideration. Much research is devoted to developing scintillators which best fulfil the requirements of imaging and event counters. Many properties must be considered for each application, allowing the user to choose the most suitable material. The following list of desirable properties (shown in no particular order) is supplied by Rodnyi (1997), but developed and discussed here with reference to the specific requirements for this study.

1. *Scintillation yield*: The scintillation yield is a physical characteristic of the scintillator: the greater the light output per MeV in, the better. The efficiency is defined as the ratio of the energy emitted to the incident energy. The higher the yield, the more photons there are to be detected, giving better statistics for detection.

2. *Decay time*: A short pulse duration is preferable for high frame-rate operation. A short decay time is also preferable to avoid lag between absorption and detection. $I(t) = I(0)\exp(-t/\tau)$ describes the decay, where $I(t)$ is the intensity of the emission at time t, and so a scintillator with a small decay constant τ is desired.

3. *Level of afterglow*: Afterglow occurs when carriers are thermally released from traps (filled when the e-h migrate to luminescence centres). The afterglow has detrimental effects on the energy resolution and maximum possible frame-rate.

4. *Temperature stability*: It is required that the light yield is stable over the temperature range used to ensure the best-possible energy resolution.

5. *Index of refraction*: A refractive index close to that of glass allows the scintillator to be more easily coupled to a fibre-optic plate, minimising problems of reflection/refraction.

6. *Transmission of own emission*: The light emitted by the scintillator may be reabsorbed by the crystal (depending on the energy band structure) or absorbed by impurities and defects. A high purity decreases this absorption and reduces the loss in energy resolution. The density of impurities may vary along the length of the crystal due to growth methods, again, affecting the energy resolution.

7. *Radiation hardness*: The radiation hardness is dependent on the presence of impurities and defects, but should be negligible at the doses used in this study.

8. *Stopping power*: Higher density scintillators provide greater stopping power, which in turn increases the detection efficiency of the scintillator-coupled device.

9. *Wavelength of emission*: The peak wavelength of emission will ideally match the peak in the quantum efficiency of the detector to maximise detection.

10. *Hardness/ruggedness*: The material should have good mechanical stability to avoid physical damage during testing and use.

11. *Chemical stability*: Limited chemical interactions, such as those with water, are desirable.

12. *Ability to grow large samples*: For commercial and industrial applications, this is essential.

13. *Uniformity*: The energy resolution is affected by the distribution of impurities. Consistency throughout the crystal is desirable.

14. *Cost*: The cost at the desired purity must be considered.

15. *Proportionality of energy response*: In order to gain information about the energy of the incident quantum, one requires the scintillator response to be proportional to the incident energy.

2.4.3 Caesium Iodide[1]

Thallium doped Caesium Iodide has been chosen for this study due to its many highly desirable properties, including the high emission yield, the similarity of the refractive index to that of glass (for ease of coupling to a fibre-optic plate), low absorbance of the emitted lower energy photons and the ability to grow the scintillator in columnar form. If the scintillator is used in an unstructured form, the light spreads radially from the initial interaction point and the peak signal at the CCD is very low (explained in Section 3.2). To increase the peak number of photons reaching the CCD the scintillator can be grown in a columnar form, as seen in Figure 2.12. The scintillator is coated with a thin parylene layer in order to prevent the absorption of water which can interact with the CsI(Tl).

Figure 2.12. An SEM of the scintillator CsI(Tl). The scintillator can be seen to be formed of columns of diameter 5-6 µm (Bill Bruns, private communication).

[1] Taken from Hall *et al.* (2009), written in the original form by the thesis author.

The columnar structure acts in a similar way to a fibre-optic plate, channelling the light to the CCD. The scintillator is grown from the substrate and is not perfect in structure and does not give perfect channelling of the lower energy photons. It does, however, act to increase the peak signal at the CCD. The peak wavelength of emission from the scintillator falls at 550 nm, matching the peak in quantum efficiency of the CCD97 of over 90% (Figure 2.9).

The scintillator produces approximately 60 photons per keV of incident energy, forming a linear dependence on energy over the range of interest. Through the analysis of the number of photons emitted, it is in theory possible to relate this figure to the energy of the incident photons. It is not possible to relate the peak signal to the energy however, as the depth of interaction in the scintillator has a large impact on the peak signal (Section 3.5.4). The whole event profile of each X-ray interaction must be considered to approximate the energy of the incident photons.

2.5 X-ray and gamma cameras

The standard gamma-camera does not include the EM-CCD, but many make use of a standard CCD. The use of an EM-CCD is covered in very few papers, but the literature which is available is discussed here, with particular attention paid to the methods used alongside the results achieved. A review of the implications of the methods and analysis used in the literature follows.

2.5.1 Medical imaging

The coupling of a scintillating layer to a CCD has been developed in the past for various medical applications. Dental imaging sensors are commercially available from e2v technologies and allow more freedom in the taking of dental X-ray images (Figure 2.13). The devices consist of a standard CCD coupled directly to a thin, unstructured Gadox (Gd_2O_2S) scintillating layer. Gadox, whilst having a higher light yield, cannot be produced in a structured form (such as the columnar form in which CsI may be grown), and hence gives a lower spatial

resolution. These devices allow the operator to see the X-ray image directly on-screen. If the image is not satisfactory, such as it being in the wrong region of the mouth, another image can be taken immediately without the waiting period required to develop X-ray film. In today's digital age it is also more convenient to have an image which can be manipulated, stored and transported electronically.

Figure 2.13. The e2v technologies CCD dental imaging sensor. The sealed units can be placed inside the patient's mouth using a separate sheath.

With the use of unstructured scintillators, the spread of signal inside the scintillation layer can be considered to be of the same order as the thickness of the scintillator. Using a thinner scintillator can lead to a better spatial resolution but this comes at the expense of the quantum efficiency – the thinner the scintillating layer the fewer X-rays that will be converted to visible light. With all medical X-ray imaging it is important to minimise the dose received by the patient, so the trade-off between resolution and quantum efficiency is an important point to consider.

Another use of X-rays, or gamma-rays, in medical imaging is in Single Photon Emission Computed Tomography (SPECT), as first demonstrated by Kuhl and Edwards in 1963, (Kuhl *et*

al., 1963). A gamma camera is used, similar to those normally employed in nuclear medicine planar imaging, to provide images taken in pre-defined positions around the patient. The gamma camera images the distribution of a radioisotope inside the patient. By taking multiple images around the patient, a 3D distribution map of the isotope can be produced.

In general planar medical imaging, a gamma camera can be used to take a 2D image of the distribution of a radionuclide. In much the same way, a series of 2D images, or projections, can be reconstructed to produce a 3D map of radionuclide density. The 2D images are taken at predefined angles around the patient and reconstructed using a tomographic reconstruction algorithm to provide the 3D information. In many cases it can be more efficient to use several gamma camera heads such that a 360° rotation can be cut to a 180° rotation by two camera heads and the image acquisition time halved. General spatial resolutions achieved with current SPECT systems are of the order of 1 mm, although any improvement in resolution is highly beneficial to diagnosis and treatment (Barret *et al.*, 1990).

One method which can be used to improve the spatial resolution of an integrated X-ray image is to use a structured scintillator. Thallium doped Caesium Iodide, CsI(Tl), is another scintillator which is widely available (Section 2.4.3). Through appropriate growth techniques a structured layer of CsI(Tl) can be produced. This layer takes the form of many column like structures. These columns act to channel light through the layer parallel to the column direction, much like light in a fibre optic plate. The process is not an ideal case of total internal reflection due to the imperfect structure produced with current techniques, but the improvements are clear over an unstructured scintillation layer (Byoung-Jik *et al.*, 2007).

2.5.2 Synchrotron-based research

Although the work detailed in Chapters 4 and 5 makes use of a small solid source to supply X-rays for imaging, other alternatives are available. A synchrotron is capable of producing a

very intense beam of electromagnetic radiation from infrared, through ultraviolet, to X-ray regimes.

Electrons generated in an electron gun can be accelerated to very high speeds through a series of particle accelerators. Once inside the storage ring, the electrons are deflected round the polygonal path using an array of dipole magnets. On deflection around the ring, the electrons lose energy which is emitted in the form of electromagnetic radiation. This radiation can be channelled out of the storage ring and into experimental stations called beam-lines (www.diamond.ac.uk).

One example of synchrotron-based research in which a high-resolution gamma-ray camera could be used is that of protein crystallography (Phillips *et al.*, 1976). A synchrotron beam-line can provide many times the flux of a sealed X-ray tube, allowing the use of smaller protein crystals than would otherwise be possible. The beam-line can also provide a tuneable source over a much lower angular divergence, meaning that a much higher resolution can be achieved if one uses a detector with a suitably high spatial resolution.

2.5.2.1 Benefits of synchrotron beam-lines

Synchrotron X-ray imaging techniques have numerous uses in biomedical applications, materials science and industrial applications. The techniques are particularly well suited to studies of living processes and in-situ materials processing. In vivo imaging of small animals can advance medical practices by removing the effects of variation between individuals and can help reduce the number of animals used for medical research. In situ imaging for materials science applications allows processes such as alloy solidification and precipitation phenomena to be investigated (www.synchrotron.org.au).

The high-intensity beam of X-rays which can be produced by a synchrotron has many advantages over solid sources. Such a high intensity allows very high signal-to-noise ratios to be achieved in short spaces of time. The monochromatic nature of the beam that can be

produced allows energy dependent measurements to be obtained. The energy range over which the monochromatic beam can be produced allows testing at many energies to be undertaken within the same environment. Furthermore, one of the problems encountered in this study, caused by the geometry of the imaging system using a solid source, may be overcome due to the highly parallel beam produced.

Within current synchrotron research, there is great interest being directed towards the development of high performance photon-counting detectors capable of resolving in energy.

2.5.2.2 Medical applications

The synchrotron beam line, with an appropriately high-resolution detector, provides an excellent source for medical imaging applications. Some beam-lines, such as the Australian Synchrotron, are currently building a medical imaging facility, allowing in-vivo human imaging, scheduled for completion in 2012. The facility is hoped to allow the acquisition of high-resolution images of cells and tissues, tumours and lung tissue, such as in mammography and lymph node imaging, and to allow cell movements to be traced in real-time using specialised markers. The planned studies aim to show how radiation interacts with cancer cells and provide a comparison with healthy cells, progressing radiotherapy treatments. Other examples of work planned to be carried out in the new beam line include (www.synchrotron.org.au):

1. development of better asthma treatments through the study of lung function;
2. bone density and porosity measurements;
3. enhanced mammography techniques;
4. soft tissue imaging with high contrast to study structures inside plants in order to develop more drought-tolerant and salt-tolerant plants.

2.5.2.3 X-ray diffraction imaging

The combination of a series of X-ray diffraction patterns can allow the reconstruction of a three dimensional image of a structure to be produced. Using coherent X-rays, one can study the crystalline structure of a material without destroying the sample (such as that required for powder diffraction).

For such purposes, a high-resolution hard X-ray imager is advised. In order to gain the information required a high frame-rate is preferred whilst maintaining the image quality. A high frame-rate photon-counting camera fits these needs, with a requirement of below 100 µm spatial resolution and up to 100 frames per second. The work detailed in Chapters 4 and 5 describes the development of such a multi-purpose camera. Future work covered at the end of Chapter 5 discusses the implications of this study and investigates the future of this technology.

2.5.3 The gamma camera

Until recently, the standard gamma-camera consisted of an arrangement of Photo-Multiplier Tubes (PMTs) coupled to a scintillating layer. When an X-ray from a patient passes into a scintillator a flash of photons is produced that can be detected by the PMT array. Through the detection of many photons an image is integrated to provide a map of radionuclide distribution in the patient.

As an example, consider the Hamamatsu flat panel type multi-anode photomultiplier tube assembly H8500 (Hamamatsu Data Sheet, 2007) used by Pani *et al.* (2006). The flat panel consists of an 8 × 8 array of PMTs, giving an effective area of 49 × 49 mm. Each 'pixel' has an area of 5.8 mm × 5.8 mm. The spatial resolution achieved with the detector described in Pani *et al.* (2006) varies with the scintillation crystal thickness. The energy resolution achieved gives a FWHM of 7.5–12% at 140 keV. A similar camera, described in Loudos *et al.* (2003), was found to have a spatial resolution of 2 mm in planar and 2-3 mm in tomographic imaging.

In 2007, Pani *et al.* described an $LaBr_3$:Ce scintillation gamma camera prototype. The camera, using continuous $LaBr_3$:Ce crystals coupled to the H8500, was found to give improved energy resolutions of 6% at 140 keV and 3% FWHM at 662 keV. The improved energy resolution is attributed to the increased linearity of the light yield from $LaBr_3$:Ce over the original scintillator. The spatial resolution is quoted to be of the order of 1 mm.

These results, all based around similarly designed systems, provide spatial resolution measurements of the order of several millimetres. Although this is a minimum requirement for SPECT, improvements in spatial resolution can provide much more detailed information (Barret *et al*., 2009). Limitations of the PMT array can be overcome through the use of a CCD with a much finer pixel structure. In order to further improve the spatial resolution, photon counting can be implemented, in which many frames are taken to capture individual scintillation events. These individual events can later be reconstructed to produce an image. With the development of the internal multiplication gain register of the EMCCD, it is possible to reduce the effective readout noise to the sub-electron level, allowing the individual photon flashes to be extracted above the noise, something which is not possible with a standard CCD. Developments in scintillator-coupled EMCCD cameras are described below. Although the spatial resolution is greatly improved it can be seen that this increase comes at the cost of the energy resolution.

2.5.4 The scintillator-coupled EM-CCD

Through the use of a non-invasive gamma-ray imager one particular animal can be examined throughout the development of a disease without the animal being sacrificed en route. This allows the study of the same animal throughout all stages of disease development and removes the statistical variations which would arise if different animals were examined at each stage in the disease. The desired spatial resolution specification for Computed Tomography (CT) and Single Photon Emission Computed Tomography (SPECT) are 100 µm and 1 mm respectively.

Over the last decade, the rapid advancements in CCD technology have lead to significant developments in the field of low-light-level CCDs (EM-CCDs). With the addition of a gain register before output the signal electrons can be multiplied without increasing the noise. This has the effect of decreasing the effective readout noise of the device to the sub-electron level. With this low effective readout noise very low signal levels can be detected above the background. These advancements have allowed the development of scintillator coupled EM-CCD detectors, in which a scintillator is either directly coupled to the EM-CCD imaging surface or coupled through a fibre-optic plate.

The ability to see signal of only a few photons can be used in an alternative imaging system to the standard integrating camera. If the number of X-ray interactions in the scintillator per image frame is reduced the individual X-ray interactions can be separately resolved. Each interaction of an X-ray in the scintillation layer produces a multitude of photons. These photons pass through the scintillator radially outwards from the point of interaction with half therefore travelling towards the CCD surface. This spread of photons produces a small Gaussian-like peak signal in the image and a form of photon-counting can be achieved provided there is a low probability of event pile-up. The probability of pile-up can be reduced by either decreasing the activity of the source or by shortening the frame integration time.

In 2004, de Vree *et al.* developed an Ultra Gamma Camera (UGC) capable of taking 50 images per second. The UGC consisted of an EMCCD (the e2v CCD65) coupled to a 1mm thick collimated CsI(Tl) scintillation layer through a fibre optic taper. The camera was designed to operate below -30 °C in order to reduce the dark current to a level suitable for single X-ray photon extraction. The CCD65 uses a rectangular pixel design, where the pixels have edge lengths of 20 µm and 30 µm, leading to differences in the spatial resolution with direction.

Rather than integrating an image over a certain time period, many images are taken in quick succession in which each image only contains a few X-ray interaction events. Each event is

then separately analysed and the resulting processed data recombined to produce the final image.

To test their system, a lead slit of width 50 µm was placed over the scintillation layer and flood illuminated with ^{57}Co and ^{125}I sources (X-ray energies of energies 122 keV and 28 keV respectively). This gave a spatial resolution with FWHM (corrected for the width of the slit) of approximately 60 µm in the short pixel direction and approximately 100 µm in the long pixel direction. The variation in spatial resolution between the *x* and *y* directions is caused by the rectangular pixels in the CCD65 (20 µm × 30 µm). This resolution of approximately three times the pixel width/length shows promise for this type of detector, but still leaves considerable room for improvement.

In terms of calculating the energy resolution, these earlier studies were rather unsuccessful. The 'energy' of each event was taken as being proportional to the intensity found at the maxima of the smoothed event profiles. This is shown in Section 3.5 to provide serious limitations to the energy resolution. Although 'peaks' are seen at the two energies tested (28 keV and 122 keV), it is not possible to determine the energy of an interacting X-ray photon due to the large overlap in the spectrum at the two energies. The overlap, seen in much of work of this kind in the literature, is shown diagrammatically in Figure 2.14.

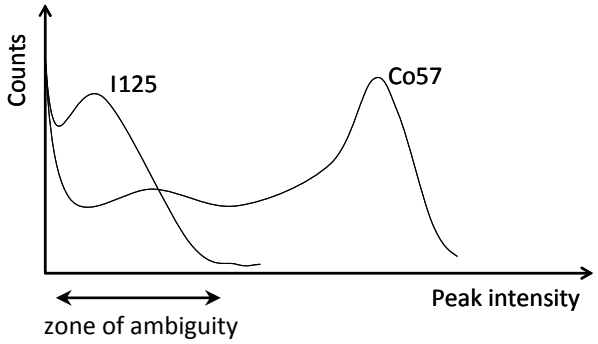

Figure 2.14. A typical energy spectra obtained from a scintillator coupled EM-CCD using the peak intensity as a measure of energy. If an event has a peak in the 'zone of ambiguity' then it is difficult to assign an energy value to the event.

In 2005, de Vree and Beekman developed this work further and attempted energy discrimination in order to remove scattered X-rays and to discriminate against 'noise' events. At 122 keV a FWHM in the peak of the spectrum of 40 keV (33 %) was calculated. The form of the spectrum was very similar to that shown for ^{57}Co in Figure 2.14 and so only the main peak was considered. The FWHM at the lower energy of 30 keV was calculated as 42 keV. Although some form of energy discrimination has been attempted the results are not ideal.

Beekman *et al.* (2005) stated that a high detector resolution combined with energy discrimination capabilities was essential for future single photon emission computed tomography (SPECT) devices. The system was altered to include a 1mm thick CsI(Tl) layer with 6 µm fibres and a 2:1 Fibre Optic Taper (FOT), now running at -50 °C. Large losses in the number of photons detected are detailed when using the FOT and this must be balanced against the increase in imaging area. It was also found that if the count rate was too high then the energy spectrum was further degraded due to the combination of multiple events in the event detection operation (pile-up). Large improvements were noted over an integrating camera system, as shown in Table 2.1.

Source	FOP (taper)	FWHM with Integration	FWHM with photon counting
^{99m}Tc (140 keV)	1:1	342 µm	64 µm
	2:1	343 µm	96µm
^{125}I (28 keV)	1:1	270µm	66µm
	2:1	307µm	110µm

Table 2.1. FWHM measurements comparing integration against photon counting. Results taken from Beekman *et al.* (2005).

The energy resolution was also calculated with results quoted as 33 keV at 140 keV and 36 keV at 28 keV using the 1:1 FOP. With the 2:1 FOT this was degraded to 76 keV at 140 keV and the lower 28 keV peak was indistinguishable from the background.

The improvement in spatial resolution through photon counting is clear from Tables 2.1 and 2.2. It is worth noting the greater than 50 % increase in spatial FWHM when using a 2:1 FOT in photon-counting mode and the severe degradation in the energy resolution. Combining this with the loss in photons when using the taper, the increase in imaging area must be carefully considered.

The processes previously described were investigated in further detail by Heemskerk *et al.* (2007). A comparison of the front and back illuminated EM-CCD devices was also implemented, comparing the e2v CCD65 (front-illuminated) and the e2v CCD97 (back-illuminated). The back illuminated (BI) device was found to be superior in all areas due to the increase in quantum efficiency over the front-illuminated device.

Device	Source energy	Spatial FWHM	Energy FWHM	E/ΔE
FI (CCD65)	28 keV	85.3 µm	not distinguishable	-
	140 keV	59.4 µm	41 keV	3.4
BI (CCD97)	28 keV	44.1 µm	22 keV	1.3
	140 keV	49.4 µm	36 keV	3.9

Table 2.2. FWHM measurements comparing the spatial and energy resolutions of the FI-CCD (CCD65) and BI-CCD (CCD97). Results taken from Heemskerk *et al*. (2007).

Several key points can be obtained from the work detailed above:

- The spatial resolution can be improved threefold (or greater) through the use of photon-counting in comparison to an integrating system.
- The BI-CCD provides better spatial and energy resolution than the FI-CCD, in which low energy events are indistinguishable from the background.
- The use of a fibre optic taper to increase the imaging area must be carefully considered as the spatial and energy resolutions are severely affected.
- Energy discrimination through the 'peak intensity of smoothed image' method is not possible due to the large spread in the energy spectrum (see Section 3.5).

Concurrent studies carried out by Miller *et al.* (2006; 2007) used a similar scintillator-coupled CCD system. Miller *et al.* (2006) describe the Bazooka SPECT system, so called due to its appearance. The system uses a lens system, allowing up to 8 × magnification, and a second generation image intensifier (incorporating a micro-channel plate) between the CsI(Tl) and CCD instead of an EM-CCD based system.

Events were isolated through background subtraction, followed by thresholding above the noise and finding contiguous regions of signal. The event positions were extracted through centroiding in an 11 × 11 pixel window around the event. Using a 25 µm tungsten slit and a ^{57}Co source (peak at ~120keV), a full field of view FWHM of 150 µm was obtained. If some of the field of view was sacrificed, allowing the system to run faster, a FWHM of 100 µm could be achieved. These values are similar to those found by de Vree *et al.* (2004). The intensification of the signal photons from the scintillator before reaching the EM-CCD allows the system to operate without cooling as the signal is significantly higher than the noise on the dark current. The noise on the dark current limited the previously described system to operation only when cooled. The energy resolution of this device is not quoted due to poor separability of the peaks, but the method attempted is still of interest. Here, the signal is summed around the event centre in a region of 9 × 9 pixels. This takes into account some of the variation in peak intensity due to depth of interaction events in the scintillator, but does not yet yield convincing results.

Developments in the 'BazookaSPECT' system have recently been presented, (Miller *et al*., 2009). The improvements shown for the BazookaSPECT II system are described below, and show the most complete alternative to the scintillator-coupled EM-CCD system. The differences in the design and performance of the BazookaSPECT II system require further discussion and comparisons with the scintillator-coupled EM-CCD show many similarities with previous performance.

In 2006, Miller *et al.* also considered the scintillator-FOP-EMCCD system. A secondary peak in the energy spectrum at approximately 30 keV was suggested to be related to the K-shell fluorescence X-ray. This was not discussed further and the implications of this peak on the spatial resolution and energy discrimination capabilities of the device were not investigated. Using the basic system an integrated image FWHM of 162 µm and photon-counting FWHM of 78 µm using centroiding was obtained. Again, these results are consistent with those

described previously. The centroiding method uses a 15 × 15 or 21 × 21 pixel window around the brightest pixel. The brightest pixel is however not always at the centre of the event due to the photon shot noise on the signal. This may be the cause of the slightly worse resolution than described by de Vree *et al.* previously, as the centroiding may be taking into account excessive noise.

A 'Most Likely Estimation (MLE) technique' was introduced where the system was calibrated to try to improve the energy resolution obtained. Through a novel calibration step, shown in Figure 2.15, preliminary data were obtained. This was used to provide the MLE with the necessary data to estimate the energy of each event detected in the CCD. The results from the use of the MLE are shown in Table 2.3 with approximate FWHM measurements taken from data presented in the literature.

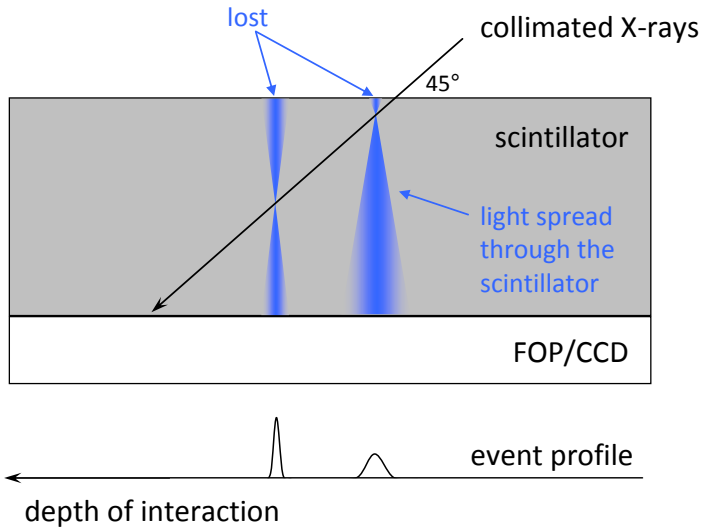

Figure 2.15. A collimated X-ray beam is incident on the scintillator at 45° such that the depth of interaction is given by the coordinate of the event along the plane of the CCD surface. The peak intensity and spread of the signal varies as one would expect with depth.

Source	^{125}I	^{241}Am	^{133}Xe	^{57}Co	^{99m}Tc
Peak energy	28 keV	60 keV	80 keV	120 keV	140 keV
FWHM	25 keV	60 keV	78 keV	100 keV	110 keV

Table 2.3. Approximate FWHM measurements taken from data in (Miller *et al*., 2006). The energy spectra were taken using a MLE method after calibrating the system.

An alternative implementation of the scintillator-coupled CCD system was described by Tipinis *et al.* (2004) where the CsI(Tl) was coupled via a fibre-optic bundle to a silicon PIN photodiode. This was designed to be used as a hand-held probe for such processes as locating tumours. This was developed in Shestakova *et al.* (2006) using the e2v CCD97 and a 140 µm CsI(Tl) layer coupled with a flexible, coherent fibre-optic cable. The spatial resolution measured was quoted to be of the order of 50 µm. The collimated CsI(Tl) reduces pixel-pixel cross-talk, thus increasing the image contrast. Running at -30 °C and a gain of 35 the total noise was brought to the sub-electron level. In this study, the 'energy' of the events was calculated by summing a 25 × 25 pixel window around the event. The resulting spectra, although peaking at differing points, had excessive overlap and energy discrimination was not possible. A diagrammatic representation is shown in Figure 2.16.

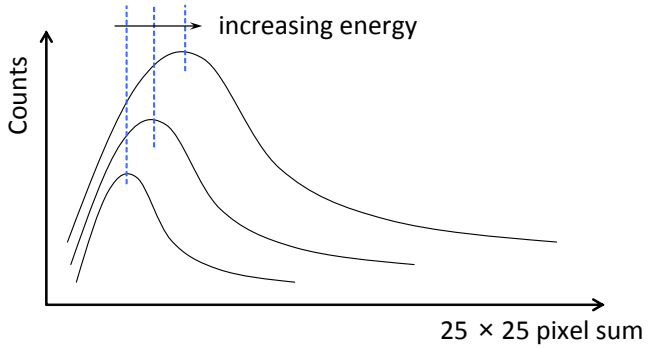

Figure 2.16. A representation of the 'energy' spectra obtained by Shestakova *et al.* (2006). Although the peak varies on the *x*-axis, energy discrimination is not possible.

Energy determination is not only useful for multi-label imaging, but can be used to discriminate between signal generated by primary X-rays and that generated by fluorescence X-rays interacting in the scintillator. 91% of interactions in the scintillator are due to the photoelectric effect, of which 79% will also produce K X-rays (Nagarkar *et al.*, 2006). In thicker screens, the re-absorption of fluorescence X-rays may lead to a blurring of the image. In the same paper, results are discussed to compare integrated images against those produced with photon counting, leading to a FWHM of 206 µm and 87 µm respectively. These results only reinforce those previously described and show that the improved spatial resolution when using photon counting methods, as opposed to the standard integrating methods, are certainly worth investigating. In medical imaging applications such as SPECT even small improvements in the spatial resolution of the system can bring great benefits, Barret *et al.* (1990).

Although several scintillator-coupled EMCCD cameras have been designed and tested with clear improvements in spatial resolution over the standard scintillator-coupled PMT gamma camera, there has been little detailed study into energy discrimination in the systems. Early attempts have shown how the peak in the 'energy spectrum' will sit at different levels related to the energy of the incident X-rays. Unfortunately these 'peaks' have a FWHM of the order of

their energy value, leaving large overlaps in the energy spectrum. One attempt to calibrate the system has lead to small improvements using a MLE technique, but the errors involved in the experimental calibration procedure minimise the benefits with respect to the extra processing time required. If depth of interaction effects are ignored, such as by taking the event peak intensities, and if smoothing is applied at a single level, then the spread of the peaks in the energy spectra cannot realistically be expected to allow energy discrimination. These results show that depth-of-interaction effects and shot noise are important factors which need to be taken into account in order to investigate a multi-label imaging system.

The EM-CCD provides many benefits over arrays of single pixel detectors, such as arrays of photomultiplier tubes, as the EM-CCD can significantly improve the spatial resolution. If such a device could be improved in the spatial sense whilst also providing sufficient energy discrimination then significant advances can be implemented.

2.5.5 BazookaSPECT II

Recent developments of the BazookaSPECT system were presented at the SPIE Optics and Photonics conference in August 2009 (Miller *et al.*, 2009). The new detector system, named 'BazookaSPECT II', is a straight-forward development of the BazookaSPECT system described previously, using a similar lens and image-intensifier system. The system can operate at 200 fps, making use of a modified PC graphics card to process the images. The system, at 1:1 magnification has an intrinsic resolution of 69 µm. Although the physical size of the system has been reduced, the camera system alone measures over 21 cm in length. The high frame rate is possible through the use of a fast-frame-rate CCD camera, as opposed to an EM-CCD, which necessitates a separate image intensifier and optical system for focussing. The BazookaSPECT II system therefore sacrifices the device size for the faster frame-rate (made possible through the use of a non-EM-CCD). In use, multiple BazookaSPECT systems can be placed around a patient to take standard SPECT images for reconstruction at a later stage.

The BazookaSPECT II system is discussed further at the end of Chapter 5 with reference to the results of this study. The system is also compared to the expected performance of a more developed scintillator-coupled EM-CCD, making use of faster drive electronics (currently at the prototype stage).

2.5.6 Commercial systems

Several commercial gamma-camera systems are available for medical imaging at the time of writing (or in the final stages of testing). Of these, three are presented here showing a selection from the range of systems available. Each system is described with reference to the advantages and disadvantages of the particular design and characteristics, ranging from the more portable ez-Scope to the larger, room-based Sentinella Suite. The final system is undergoing clinical trials and is of a similar design to that detailed in this work.

2.5.6.1 eZ-Scope

eZ-Scope is a hand-held gamma-camera with a small field of view, measuring just 7 cm × 7 cm × 22 cm, with a mass of 800 grams (www.nuclemed.be). It was designed for intra-operative imaging to replace current audible indicators for areas of interest. The camera has a field of view of 3.2 cm × 3.2 cm, with a pixelated area of 16 × 16 pixels. The detector is fabricated from a single piece of CdZnTe, pixelated to give a pitch of 2 mm. The camera system has a detection efficiency of 99 % at Tc^{99m} with an energy resolution of 7 % FWHM. Although this system provides a much higher detection efficiency and energy resolution than the detectors discussed earlier, the spatial resolution falls short of those achieved with the scintillator-coupled EM-CCD. This means that, although the eZ-Scope is a large improvement over audible detectors, it is not a natural imaging device and may only provide an indication of the location of the region of interest.

2.5.6.2 Sentinella Suite

The Sentinella Suite provides "the new standard in oncological surgery" (www.gem-imaging.com). The system is much larger than the eZ-Scope and combines a camera on a manoeuvrable arm attached to a base stand including the monitor and drive system.

The camera system is designed for sentinel node imaging for breast cancer diagnosis and treatment, and for imaging aimed towards other cancer diagnosis. The system offers a spatial resolution of 1.6 mm and a larger field of view than the eZ-Scope, but consequently is no longer hand-held.

The image resolution is limited by the detectors used in both the Sentinella Suite and the eZ-Scope, offering only 1-2 mm imaging resolutions. This is countered by a high detection efficiency and high energy resolution, but it is the imaging accuracy that is most vital in many medical imaging applications.

2.5.6.3 The Mini Gamma-Ray Camera (MGRC)

The Mini Gamma-Ray Camera (MGRC) was developed at Leicester University and is currently undergoing clinical trials (CNRS Press Release, 2007). The camera is based around a scintillator-coupled EM-CCD, using the scintillator Gadox (Gd_2O_2S:Tb). The camera was designed to image the gamma-ray emissions from Tc^{99m}-labelled tracers (140 keV) in the detection of sentinel lymph nodes in order to aid cancer treatment. The imaging is designed to facilitate a non-intrusive and non-surgical method of tumour localisation. The camera can operate in the energy range of 30 keV to 160 keV, with an intrinsic sub-millimetre spatial-resolution, quoted at approximately 0.5 mm.

The literature surrounding the MGRC makes an interesting comment on the use of smaller field of view gamma cameras. Whole-body gamma cameras come at great cost and offer spatial resolutions of the order of 10 mm (largely due to scattering), the smaller gamma-camera can provide much high spatial resolutions over smaller fields of view and provide a

natural complement to whole body cameras once the area of interest has been sought using the full body imager. The reduction in cost of smaller cameras also opens them to much wider usage throughout the field of medical imaging.

2.5.7 Summary

Current commercially available systems have very low spatial resolutions compared to devices in development (such as those discussed in Section 2.5). Although the energy resolutions are better in the commercial devices, this comes at the expense of the spatial resolution and allows the cameras to be used only as an indicator to the areas of interest.

Detectors based around Photo-Multiplier Tube (PMT) arrays coupled to scintillators provide similar results to the commercial systems. Large fields of view can be formed by combining multiple PMTs in an array, but the resolutions achieved with such arrays are only of the order of 1 mm. Early scintillator-coupled EM-CCD devices, such as scintillators coupled to the e2v CCD65, provide better results than previous PMT arrays. Resolutions of approximately 60-100 µm are quoted in the literature, providing large advances over PMT arrays, but still leaving room for further improvements. Energy discrimination using these early scintillator-coupled EM-CCD systems was found difficult due to the very broad spectra recorded. No energy discrimination is possible with an integrated image, and throughout all tests of such camera systems, the method of photon-counting was found to achieve at least three-fold improvements in spatial resolution over integrated imaging.

Using a CCD97 (the detector used in this study), a best-case resolution of 50 µm was quoted in the literature. This must be considered carefully, as no mention is made of the re-absorption of X-ray fluorescence (see Chapter 5 for further discussion on the effect of re-absorbed fluorescence and the implications of ignoring this component on overestimation of the MTF). One research group did, however, discuss the possibility of secondary detection of re-absorbed

K-shell fluorescence X-rays. Although the possibility was mentioned, no detail of any further investigation has been given, leaving the effects of the re-absorption open for further study.

If secondary events are being detected due to the re-absorption of X-ray fluorescence from the scintillator itself, the image quality may be decreased (in both spatial and spectral resolution). If energy discrimination can be achieved with greater accuracy through improvements in device design and data processing techniques then the detection of events in the scintillator can be investigated in much greater detail. It is worth noting that in imagers with much lower spatial resolutions, the majority of the re-absorbed fluorescence X-rays may be imaged in combinaton with the primary events, cloaking the issue in previous studies. Through a more thorough investigation into fluorescence re-absorption, this factor may be partially removed, improving the spatial resolution further still. The sparseness of discussion on the effects of X-ray fluorescence re-absorption in the literature leads one to assume that this factor has been ignored as anomalous, left unqualified or left hidden by the low resolutions achieved.

This study paves the way for the design of an optimised detector. Following the findings discussed here, a future system is described at the end of Chapter 5 where soon-available faster drive electronics are suggested as one method to improve performance with relation to the methods discussed in this work.

2.6 Scale-Space

In Section 2.5 the problems of ‘depth of interaction’ effects were mentioned in relation to several previous studies. The depth at which the X-ray or gamma-ray interacts inside the scintillator has a major impact on the event profile detected by the EM-CCD. The depth of the interaction can only be determined from a more detailed analysis of each event profile individually. For this reason, the event profiles must be characterised. In an idealised case a simple 2-dimensional Gaussian profile can be fitted to each event through a standard least-squares process. The least-squares fitting process is shown in Chapter 3 to be unsuitable for

low signal-to-noise cases, leaving an alternative method to be found. One such alternative, developed in this study, may be derived through the analysis of the event profiles in Scale-Space as described below.

When considering a one dimensional standing wave such as a sound wave, the Fourier Transform (FT) provides information regarding the frequencies present. From this FT, the original wave can easily be reproduced.

Through the use of the Wavelet Transform (WT) a non-standing wave can be analysed to give a three-dimensional map which not only shows the frequencies, but also visually shows the positions in time at which these frequencies are present. This provides a much clearer visual statement of the properties of the wave.

The theory of scale-space is similar in many ways to the wavelet transform and has many applications in image processing. The process detailed in this chapter is related to work carried out by Lindeberg (1993; 1994; 1996), before further developments as detailed in Chapter 3. The theory of scale-space can be considered to be an adaption of the Fourier transform and the Wavelet Transform. Although not specifically used in the scale-space analysis, the following introduction helps to place the theory in context with more commonly used techniques.

2.6.1 Fourier analysis

If presented with a one-dimensional wave such as a sound wave one might be interested as to what frequencies are present. Using the Fourier transform it is possible to transform from the variable in which the data are presented (in this example: time) to a different related variable (in this example: frequency). The Fourier transform of the wave *f(t)* in the time domain is given the notation *F(v)* in the frequency domain and is described by Equation 2.18.

$$F(v) = \int_{-\infty}^{\infty} f(t)\exp(-2\pi ivt)dt \qquad (2.18)$$

The inverse of the Fourier transform can be calculated for a wave and is given by Equation 2.19. This reverts the transformed wave back to its original form as is illustrated in Figure 2.17.

$$f(t) = \int_{-\infty}^{\infty} F(v)\exp(2\pi ivt)dv \tag{2.19}$$

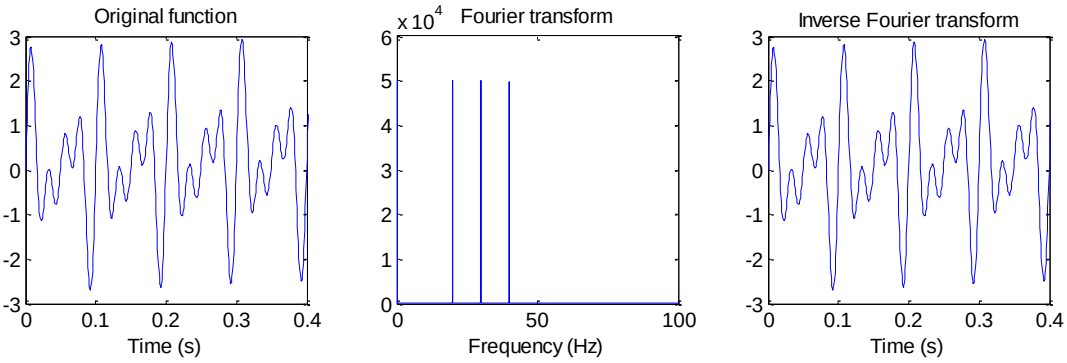

Figure 2.17. A demonstration of the Fourier transform and inverse Fourier transform applied to a standing wave.

The case of interest occurs when the wave is not a stationary wave – the case in which the frequencies vary with time. In this instance, whilst the Fourier transform will still specify the frequencies present in the wave, the location in time of each frequency is not easy to visualise in the results produced. The wavelet transform, however, provides a much clearer visual representation of the location of the various frequency components in the wave.

2.6.2 The Wavelet transform

The Fourier transform has limitations visually when encountering a non-stationary wave. What is needed is a method of extracting the frequencies present in the wave whilst visually maintaining their location in time.

The wavelet transform, in its basic form, allows the extraction of the frequencies present in a signal whilst preserving their location in time. The wavelet transform produces a 3-dimensional plot of the wave in frequency *and* time using a series of wavelets (literally 'small waves'). A wavelet, in the simple case, could be defined by a Gaussian profile.

By convoluting the original wave with a series of wavelets, each based on the same mother wavelet but with a different width, a 3-dimensional map of the wave can be built up, as shown in Figure 2.18. The width of the function is named the 'scale' of the wavelet and as such gives the translation-scale map shown. The map is not an exact time-frequency map, although the scale and frequency are strongly correlated (inversely), as are the translation and time. The two axes of translation (time) and scale (frequency) are given along with the colour scale representing the relative presence of the frequency. The smaller the separation between steps calculated, the greater the resolution in time-scale space.

Each line in the map at a constant scale (across all times) is the result of the convolution of the original waveform with the wavelet at that scale. When considered at all scales the map is completed for the discrete wavelet transform. Through this process the wave has been converted to an image, increasing the dimensionality of the data. If this process is applied to an image the resulting map is rather more complicated due to the increase again in the number of dimensions. Scale-space theory uses similar ideas to those found in the wavelet transform theory and has some very useful mathematical intricacies.

Whilst the Fourier transform can provide visually the frequencies present in a signal, the wavelet transform gives a clear representation of the position in time of the different frequency components.

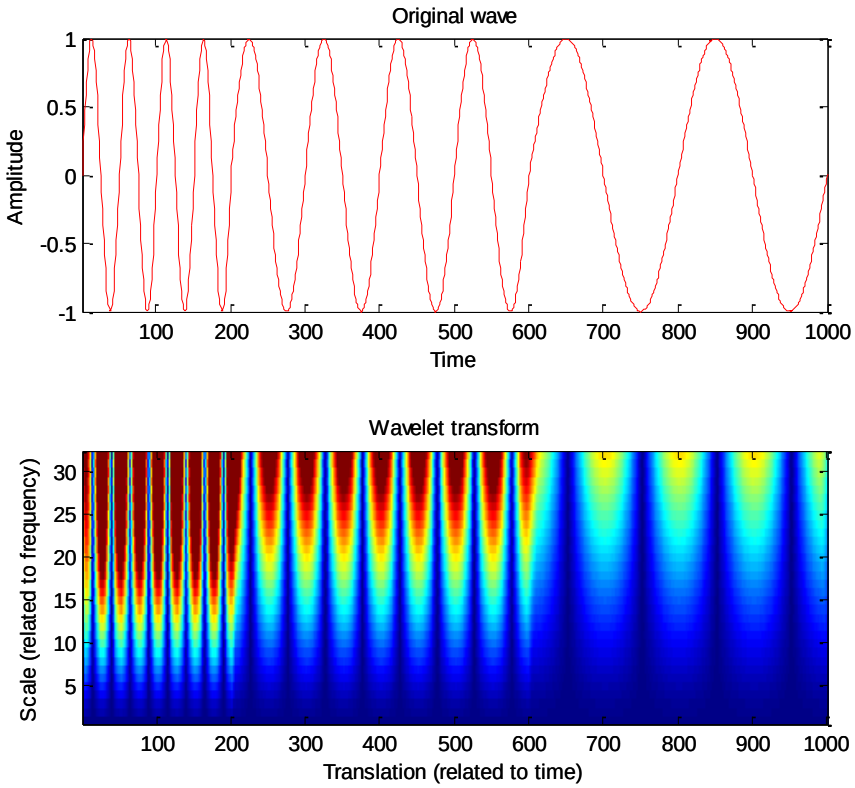

Figure 2.18. An example wavelet transform generated in MATLAB for a wave with three distinct frequencies at different times. The position in time of each frequency can be seen from the extra dimension. The 'scale' is related to the frequency, with low scales relating to high frequencies, and the 'translation' is directly proportional to the time.

An image can be considered in the same way as the previously described wave, except that the image has a further dimension. Each row is equivalent to a wave, but each row must be considered concurrently with the surrounding rows to fully represent the 2-dimensional features of the image.

Scale-space theory is similar in many ways to the wavelet transform. The image is convoluted with a set of 2-dimensional wavelets in much the same way as described above for the 1-dimensional case. The theory lends itself to the analysis of 'Gaussian-like' structures in images, such as those found when an X-ray photon interacts in a scintillator coupled to an

EMCCD. The outputs from the Scale-space analysis, for the purposes of this study, are required to aid the visualisation of the imaging capabilities of a detector and to aid the understanding of the processes involved in the imaging chain. The Scale-space analysis lends itself to the visual analysis of two-dimensional images and is explained below in principle and in greater detail in Section 3.4.

2.6.3 Scale-space

A commonly used example to represent the concept of scale is that of how we perceive a simple tree (Lindeberg, 1993; 1994; 1996): from a distance, the tree appears to be formed from a round shape and a cylindrical shape. As you get closer, the branches, and then the individual leaves, become visible. If you move closer still, the structure of the veins in the leaves can be seen. As you move closer to the tree from your initial viewpoint the tree is being seen on different scales, from the large scale at a distance to the finer scale up close.

The development of scale-space theory began back in the late 1970s with the intention of developing a method to allow computational analysis of images at different scales. Since then, the theory has been developed and provides a wider scope for computational analysis of images in many different sectors. The scale-space approach used in this study is developed from work described in the mid-1990s by Lindeberg (1993; 1994; 1996).

2.6.4 The history of scale-space

In the late 1970s, a multi-scale representation of images in terms of pyramids was developed. The aim of the multi-scale representation was to improve computational efficiency in the analyses of the image, but relating the objects at differing scales (tracing objects from coarse to fine scales) was found difficult. Edge pyramids are described in detail by Hong *et al.* (1982). In this case, an initial grey-scale pyramid is produced by averaging the original image in non-overlapping 2 × 2 pixel areas to create the second stage and repeating the process to produce the higher levels of the pyramid. An edge detecting algorithm is then applied to the image at

each level of the pyramid. The edges can then be linked to produce a tree-like structure representing all the borders found in the image. Each edge is then provided with a score in the pyramid measuring "how well the edge obeys requirements of good continuation and closure" (Hong *et al*., 1982).

The theory was developed in 1984 by Witkin using a Gaussian kernel (Witkin, 1984). He explains that the scale-space image can be "collapsed, using its qualitative structure, into a tree providing a concise but complete qualitative description covering all scales of observation". A stability criterion is also described, allowing events that persist over changes in scale to be identified.

In order to examine a signal at different scales, one must convolute the signal with a kernel varying in size. It is worth pointing out that the work described above using a gray-scale pyramid leads clearly into the more recent work, but with the original kernel only covering a 2 × 2 pixel area. As one passes through incremental scales, the image may appear differently, with new extremal points appearing and disappearing, giving a very different view of the same image. One smoothed image cannot, however, be considered more 'correct' than any other. It is important to consider the image over all scales: "the ambiguity introduced by scale is inherent and inescapable, so the goal of scale-dependent description cannot be to eliminate this ambiguity, but rather to manage it effectively" (Witkin, 1984).

The Gaussian kernel was selected for Witkin's work due its "well-behavedness": the Gaussian kernel is symmetric, it decreases with scale strictly about the mean and as such the weighting applied decreases smoothly with distance, the convolution tends to the original image for small scales and to the average value over the whole image for large scales, and finally, the Gaussian kernel is readily differentiated and integrated (Witkin, 1984). In fact, considering the 'zero-crossings' with the Gaussian kernel, and those of it's derivatives, leads to the use of the Gaussian kernel as the only kernel to satisfy all the required criteria (Babaud *et al.*, 1986). This

theory can be applied to one-dimensional signals to leave a two-dimensional 'scale-space image'. The similarities between this approach and the wavelet transform using a Gaussian wavelet cannot go unnoticed. In this instance, scale-space is but a different form of the discrete wavelet transform.

The development of scale-space to the analysis of two-dimensional images was described in Lindeberg (1993; 1994; 1996). The two-dimensional theory is based on the proviso that local extrema are not enhanced when the scale parameter is increased continuously.

In this work, the scale-space theory is described for the specific case of the Gaussian-like event profiles found with the scintillator-coupled EM-CCD. The technical details of this scale-space analysis are covered in Section 3.4.

Chapter 3: Simulating the X-ray/gamma-ray camera

3.1 Introduction

In order to better understand the individual light flashes from single X-ray interactions in a scintillator and the effect the interaction processes have on the quality of the image produced one must delve deeper into the detection process. By way of simulating the individual components of the system it is possible to gain an insight into the processes affecting the spatial and spectral resolutions achieved with the detector.

Following an initial feasibility study, the individual components of the detection process may be investigated using a variety of Monte Carlo simulations. The initial simulation work is coded using Geant4, a toolkit which allows the simulation of the passage of visible photons from a point interaction inside the columnar scintillator structure.[2] These initial results can be combined with further simulations, ranging from the readout of the EM-CCD and gain structure to the centroiding techniques available to analyse the images produced. Through the use of such simulation techniques it is possible to gain a better understanding of many aspects of the detection process, giving a greater insight into how the detector performance can be improved.

The basic Scale-Space theory detailed in Chapter 2 may be expanded upon for the case of the Gaussian-like event profiles produced when using a scintillator-coupled EM-CCD. These results are described theoretically in this chapter, with a full derivation of the results left to the appendix.

[2] The Geant4 toolkit will be used in more detail in Chapters 6 and 7 for the study of the instrument background in space-borne X-ray spectrometers.

3.2 The requirement for a structured scintillator

The two main scintillator alternatives for use in the detector are a solid block of scintillator and a more structured (columnar) form. Although the columnar structure has been selected for this detector, it is worth considering the solid-block alternative to show this requirement.

A simple analytical model can be used to give approximate values for the expected detectable signal from a 30 keV X-ray interacting in a solid block of CsI(Tl) (approximately 60 photons are produced per keV of incident X/γ-ray). A distance of 50 microns from the imager to the interaction position in the scintillator is used for these basic calculations. With no losses, assuming radial emission of photons from the point of interaction and including the 80 % QE across the emission range of the scintillator, CsI(Tl), a peak signal of approximately 10 photons could be expected at the central pixel of the event (using a 16 µm pixel). Each of the neighbouring pixels could be expected to receive 5-10 photons. The shot noise on this low signal equates to approximately 30-50 % of the signal level itself. Following multiplication gain, the noise on the signal levels (shot noise on the incident photons and the extra multiplication gain noise) is equivalent to 40-70 % of the signal levels.

The above approximations assume that all photons passing radially towards the EM-CCD are detected, but this is not the case. Losses inside the scintillator due to scattering away from the central pixel, absorption of the emitted photons (minimal) and losses at the interface between the scintillator and detector (or fibre-optic plate between) cause fewer photons to be detected than those quoted above. Lower signal levels give much higher equivalent noise levels on the signal electrons and this causes major problems in analysis and in searching for events above the background noise.

A Monte Carlo simulation has been implemented to simulate the passage of photons through a solid block of scintillator for a 30 keV X-ray photon interacting at 50 µm from the imager. The first results show approximately 10 photons peak signal and 0-5 photons surrounding the main

peak (varying with Poisson statistics), Figure 3.1. These results do not, however, take into account the noise on the multiplication gain, which, when included, causes great problems in event detection. The additional Poisson noise on the signal from the multiplication gain causes a non-uniform distribution, with large fluctuations in the signal level in each pixel (Figure 3.2). Further noise will be seen in the detected images due to readout noise (although at a low level if using high gain) and shot noise on the dark current. Noise of only a few electrons rms. will cause major problems in the detection of the event profiles at these low signal levels.

The use of a structured scintillator is designed to add a degree of collimation to the emitted light, channelling the photons towards the detector with each CsI column acting as a simple light guide. The structure does not provide perfect collimation due to the imperfect surfaces (unlike a mirror or fibre-optic plate), but does aid the increase in the peak signal at the detector (Byoung-Jik *et al.*, 2007).

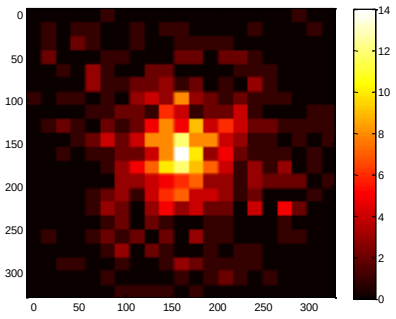

Figure 3.1. An example of a simulated profile resulting from photon passage through the scintillator. With no losses, a solid 70 µm thick block of CsI(Tl) would generate a peak signal of 14 photons for an incident 30 keV X-ray interaction 50 µm from the imaging detector, simulated here for a noiseless imaging device directly coupled to the scintillator. In reality, losses will lower this value (interface losses and QE of detector). The Poisson variance on the signal will be double that shown in this figure due to the multiplication gain.

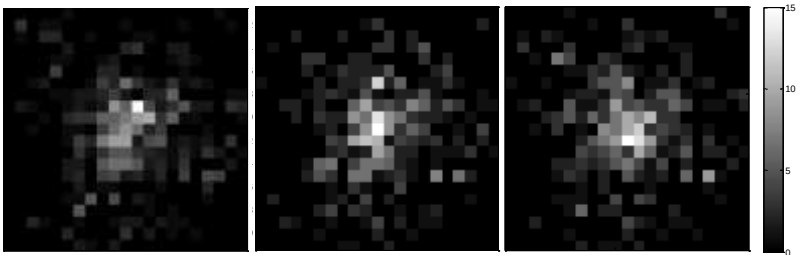

Figure 3.2. The noise on the multiplication gain has been added to the simulated profile shown in Figure 3.1. The signal is again shown with no losses or detector noise. In reality photons will be lost at the interfaces between the scintillator, fibre-optic plate and EM-CCD. The average QE across the emission range (~80-90 %) will further lower the signal detected.

3.3 Event location using a point-search function

When an X/γ-ray interacts in the scintillator a small 'blob' of light will be recorded by the EM-CCD. This 'blob' of light indicates an interaction in the scintillator at those co-ordinates in the plane of the detector. No information is available directly regarding the depth of the interaction, leading to the requirement for further analysis to be applied to the event profiles recorded. Before analysis of the event profiles, however, one must first perform an automated search for the event profiles themselves in the image.

Each image processed contains many individual events for which the location is unknown. The position of each event is clear to the human eye, especially if the images are smoothed, but an automated search process must be produced in order to automate the analysis process.

In order to locate the events, or 'points of interest', an image can be convoluted with a Gaussian profile such that the noise is sufficiently suppressed but the spread of the signal does not lead to the coming together of neighbouring events. A limit is set to one or two sigma above the average level of the noise. The noise in the image following the convolution is much lower than in the initial image, allowing low signal-to-noise images to be analysed.

In order for an 'event' to be selected, regions of 5×5 pixels with signal above the limit are found in which the central pixel must contain the highest signal of the group, Figure 3.3. The shot noise on the signal does not allow this method to be used on the original image, as it is not necessarily the central pixel which has the highest signal. The convolution with the Gaussian kernel emphasises Gaussian-like profiles in the image, effectively averaging the signal event profiles to a smoother Gaussian-like profile. This method allows for events to be detected with a separation of 3 pixels or more, Figure 3.3. Events which occur closer than this (and those around this distance) will 'join' as a result of convolution with the larger Gaussian kernels.

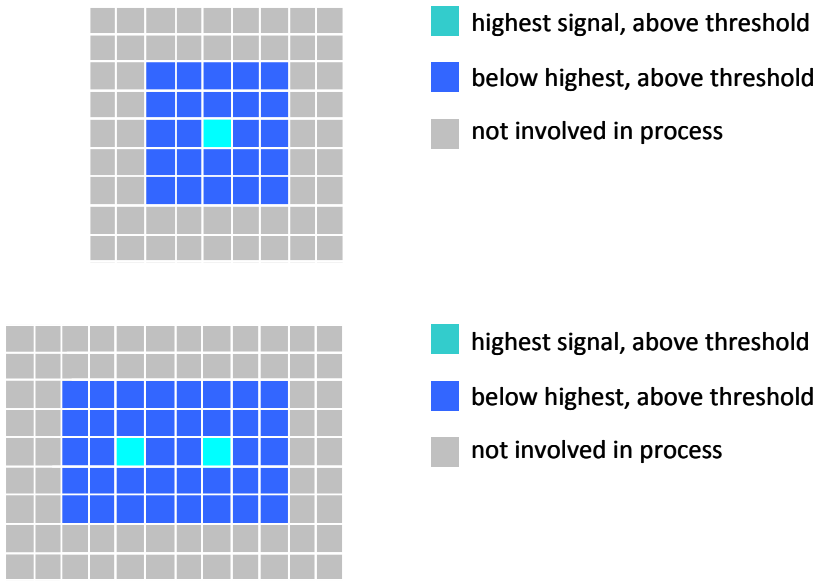

Figure 3.3. The event profile required for an event to be detected (top). Two closely interacting events (possibly a photo-electron and K-shell X-ray) which will be detected as individual events (bottom). Events with lower separation will not be individually detected.

3.3.1 Optimisation of the point-search function

When searching for points of interest in an image, two parameters define whether a position in the image is designated a point of interest or ignored. The first of these parameters is the value of sigma for the Gaussian kernel used to suppress the noise in the image. This does not directly determine which events are accepted and which are to be ignored, but must be combined with a lower limit on the event intensity (the second parameter).

The values of these two parameters, the sigma and the limit, affect the efficiency of the system in two ways. If the limit is high and the sigma large then fewer positions in the image will be assigned as 'events', but the confidence in the events (whether the event is real) increases: low numbers of counts, but a high percentage of the counts are 'real'. If the limit is low and the sigma small then the total number of detected events increases, along with an increase in the number of 'false' events: the confidence in the selection of only 'real' events decreases.

In order to optimise the parameters of point selection, one must carefully consider the application and a balance must be chosen between the number of events detected and the accuracy of the detection. Once this has been considered, it is possible to test the set-up in-situ to give the best results.

A series of X-ray images is taken with an opaque edge placed over the device. Whilst it is possible to apply this test with one image, using multiple images provides more stable results. Images are taken at a sufficiently high frame-rate such that individual events can be resolved. Several images are taken under experimental conditions with a final image taken using a longer frame time (or with an LED) to provide an image of the edge itself. This set of images can then be used to optimise the point-search parameters.

Regions of the image are designated as dark *D* and light *L*, referring to the covered region of the image and that remaining open to X-rays respectively. It may be assumed that all events detected in *D* are non-events: regions of high noise or pixels with dark current above the limit.

For any combination of the two parameters, if one assumes that the number of non-events across a number of frames is constant over the surface of the device then one can predict the total number of false events per pixel per frame. The number of real events per pixel per frame can then be calculated from the number of detected events in *L* following subtraction of the number of false events per pixel per frame (calculated from *D*). The detection efficiency (determining the accuracy of the point-search function) is then given by the ratio of true events to total events. The total number of real events per pixel per frame can be taken from these same calculations.

This process is repeated, using the same batch of images, for varying limits and sigma using automated code to provide a three dimensional surface plot. The comparison of the parameters with the requirements of the detector (the balance between detection efficiency and total number of real events detected) allows the appropriate parameters to be selected. An example of this code in use with real data may be found in Section 4.6 where the process is used to select the appropriate parameters for this study.

3.4 Using Scale-Space with Gaussian-like profiles

The case of general automatic scale selection is covered by Lindeberg (1993) and the basic theory was introduced in Chapter 2. The theory allows the scale of any feature on an image to be estimated. In the case required here, where specific features of events are to be analysed, the examination of how the convolution affects the Gaussian-like profiles in the image leads to the formulation detailed below.

The scale-space representation for a two-dimensional image is produced by convoluting the image with a two-dimensional Gaussian kernel. In order to save computational time, it is often easier to convolute the image in each direction separately with a one-dimensional Gaussian kernel, giving the same effect. This convolution produces one-level in the scale-space image

'stack'. The scale of the Gaussian kernel, σ^2, can then be increased/decreased incrementally to produce a complete stack of images ranging from small to large scales, Figures 3.4 and 3.5.

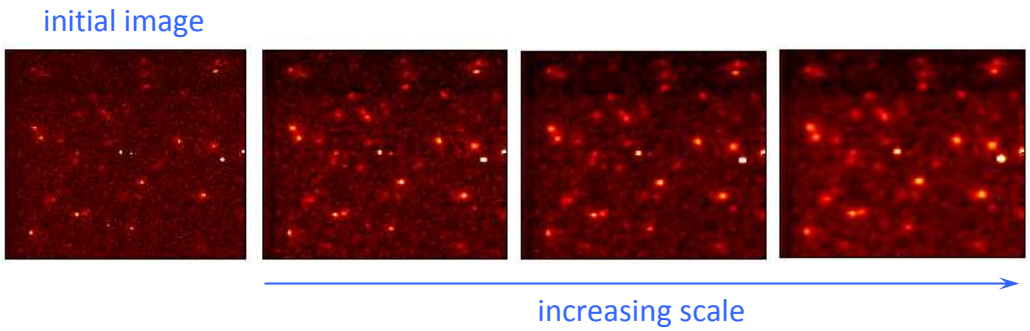

Figure 3.4. An example image (taken from experimental data) shown at four different scales, with scale increasing from left to right.

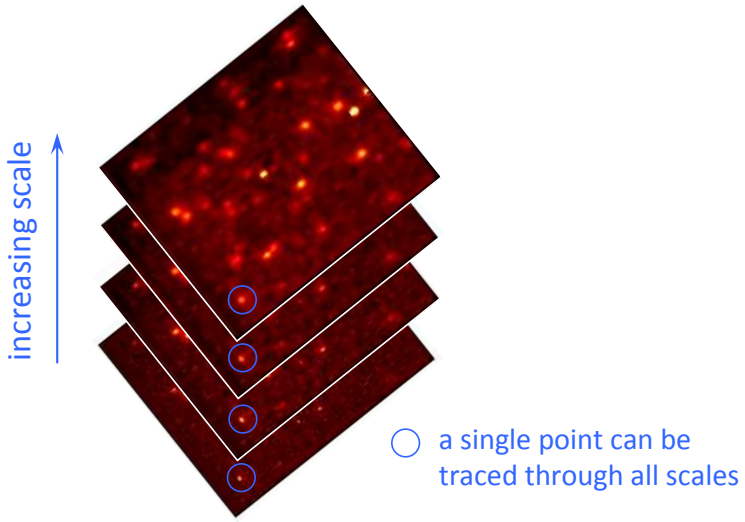

Figure 3.5. The images can be described as an 'image stack' where each pixel in the image can be traced through all scales.

In the case of the analysis of Gaussian-like profiles, the scale-space representation can be developed to provide accurate information on the event profiles, based upon the work of Lindeberg (1993; 1994; 1996). Here, a representation is described from first principles, specialising the scale-space theory for the analysis of the Gaussian-like profiles found from X-ray interactions in the scintillator-coupled EM-CCD. The following derivation is applicable to Gaussian-like profiles. The analysis holds provided the profile follows the form of Equation 3.1, where *s* is defined as the scale.

$$I_0(x,y) = \alpha \exp\left[\frac{\beta\left(x^2+y^2\right)}{s}\right] \tag{3.1}$$

If one considers the convolution of a normalised Gaussian kernel, defined by scale σ_k, and a Gaussian-like profile in the image of scale σ_i and peak intensity *A*, the resulting formula (Equation 3.2) can be used to develop the automatic scale selection method for this specific case. The mathematics leading to this formula are shown in Equations A3.1 to A3.19 in the appendix at the end of the chapter.

$$I(x,y) = \left(\frac{A\sigma_i^{\,2}}{\sigma_i^{\,2}+\sigma_k^{\,2}}\right)\exp\left[\left(-\frac{1}{2}\right)\left(\frac{1}{\sigma_i^{\,2}+\sigma_k^{\,2}}\right)\left(x^2+y^2\right)\right] \tag{3.2}$$

Equation 3.2 shows that the resulting image profile is dependent on the peak intensity and scale of the original profile and the scale of the Gaussian kernel. In the very low noise case, a two-dimensional Gaussian could be fitted to the smoothed image, where the smoothing has suppressed the noise. In the real case this is not possible as the shot noise on the signal along with the extra noise factor for the avalanche gain leads to a low signal-to-noise ratio at the low signals experienced with the scintillator-coupled EMCCD.

The detailed calculations for the following results are again shown in the appendix, with only the main formulae given here. The function *J* can be defined as the trace of the Hessian matrix on *I*, as shown in Equation 3.3 (Equation A3.25). This is applied to give an image of *J* at each

scale. The function *J* is then normalised to take account of the scale of the Gaussian kernel and its effect on the intensity of the profile, giving an image of J_s at all scales. The function J_s needs only to be considered at the centre of each event and leads to Equation 3.4 (Equation A3.30).

$$J(x,y) = \frac{\partial^2 I(x,y)}{\partial x^2} + \frac{\partial^2 I(x,y)}{\partial y^2} \tag{3.3}$$

$$J_s = 2A \frac{\sigma_k^{\,2} \sigma_i^{\,2}}{\left(\sigma_i^{\,2} + \sigma_k^{\,2}\right)^2} \tag{3.4}$$

If the value of J_s at the central point of an event is plotted through all scales then the function can be considered as a function of J_s against σ_k, both of which are known. The unknowns of the peak intensity and scale of the event profile can then be found by fitting to the 'simple' function in Equation 3.4 using a simple least-squares method.

The advantages of the scale-space analysis method for low signal-to-noise images are discussed later in the chapter. The convolution of the image at many scales allows the noise to be suppressed and bright pixels and regions of higher noise (through random statistics) to be discounted and ignored with no weighting given to a particular event profile width.

3.5 Simulating the passage of photons through the scintillator

The implications of the interaction position in the scintillator are far reaching and affect both the spatial and spectral performance of the camera. The 'depth of interaction effects' lead to the requirement for investigation into the optimal energy reclaim technique from the event profiles, covered theoretically in the previous chapter. The thickness of the scintillator affects both the absorption of the incident X-rays/gamma-rays and the range of event profiles possible, with thicker scintillating layers providing a much larger spread of event profiles.

The simulation detailed here uses the Geant4 toolkit (see below) to implement a Monte Carlo simulation of the passage of the emitted visible photons through the scintillator in order to investigate the profiles presented at the detector. The interaction cross-sections for the

incident high-energy photons are included in the simulation to give a more realistic representation of the interaction depths. The columnar structure of the scintillator is modelled as accurately as possible with the Geant4 toolkit to allow the passage of the visible photons to be tracked from their interaction position in the scintillator.

Analysis of the event profiles follows, making use of the scale-space analysis techniques. The results from the Geant4 photon-transport simulations are further evaluated later, where the simulation results are passed through the EM-CCD output simulation.

The Geant4 toolkit, described below, is also used in the instrument background section of this study, described in Chapters 6 and 7.

3.5.1 The Geant4 toolkit

The Geant4 toolkit resulted from a worldwide collaboration between physicists and software engineers. It is designed to cover as wide a range of applications as possible and as such requires a large amount of programming to be performed by the user. The toolkit offers many physical processes to be simulated including optical, hadronic and electromagnetic processes. A wide range of particles can be modelled passing through a user defined geometry, with each particle traceable as a separate entity. Event data are generated for all particle-geometry and particle-particle interactions and can then be stored and manipulated as required.

Geant4 is an object-orientated simulation toolkit implemented in the C++ programming language. The Geant4 toolkit aids the production of Monte Carlo simulation code to model the passage of particles through matter. Code is written to represent each aspect of the system such as the geometry and the logging of data, and is then compiled through the Geant4 toolkit to combine the many individual components. The compiled code can then be operated through a command line interface.[3]

[3] Agostinelli *et al.* (2003) provides a more detailed explanation of the history behind Geant4 along with its main features. A brief summary only is given here.

Geant4 also allows the visualisation of particle progress through the user-defined system, such as that shown in Figure 3.6. This shows an OpenGL visualisation taken with ten incident protons entering a simple geometry. The colours are representative of the charge of the particle, with blue being positive (protons), green being neutral (gammas, neutrons) and red being negative (electrons). In this visualisation, the upper layer is defined as aluminium and the lower layer as silicon. For consistency with the use of individual particles in the model, X-rays are referred to in a particulate manner as photons.

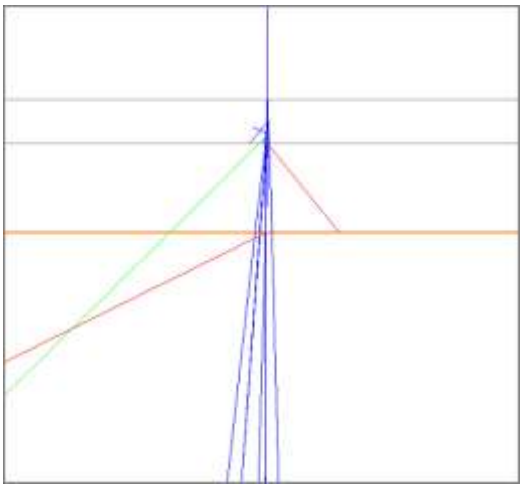

Figure 3.6. An OpenGL visualisation of a simple geometry (described in Chapter 7), shown in grey and orange. The tracks from negative particles (electrons) are shown in red, positive particles (protons) shown in blue and neutral particles (photons) shown in green.

Individual files are required to detail each part of the simulation. These files, coded in the C++ programming language, describe each part of the simulation separately, before finally being compiled using the Geant4 toolkit. Geant4 enables the user to create the particles described, to track them through their individual geometry and to log the results in a method suitable to the purpose.

The Geant4 toolkit (version 4.8.0) has been used in more depth in the instrument background work detailed in Chapters 6 and 7. A more thorough investigation using the Geant4 toolkit is described in Chapter 7 including a full description of the workings of the simulation package, with only the specifics of the photon simulation work included in this chapter.

3.5.2 Simulation structure

A full description of the file structure required for the Geant4 simulations used in this study is given in Chapter 7 where a more complex example is presented. The first stage in the simulation is to define the materials used.

A vacuum cannot be coded in Geant4, but it is possible to set the pressure of any defined gas. Air is defined simply as oxygen and nitrogen and is set with a very low pressure (effectively a vacuum) for the 'world' volume (see below).

Caesium and iodine are defined as individual elements before combining to give CsI, nominated as the scintillator. The scintillation properties are then added to include the emission spectrum (Figure 3.7), the refractive index (approximated as a constant across the emitted wavelength range), the absorption coefficients and the fast/slow decay components. Further general properties are then defined such as the scintillation yield (52 photons/keV) and the time constants.

The detector, although silicon based in reality, has been set as a vacuum. Nearly all visible photons (90%) are assumed to be absorbed in the silicon meaning that particles are logged as they pass through the detector layer. If the layer was silicon based, then further interactions in the silicon would complicate the detection logging in this case.

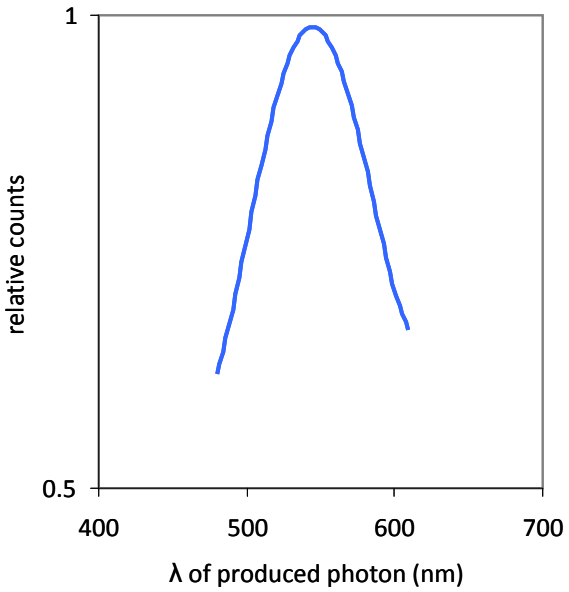

Figure 3.7. The emission spectrum for caesium iodide implemented in the simulation (supplied by the manufacturer).

The world volume defines the outer reaches of the simulation geometry. Particles are lost once they have passed outside the extremes of the world volume. The volume is set as a cube of 5 cm edge for simplicity. The volume is defined as a very low pressure gas (oxygen and nitrogen), effectively a vacuum.

Inside the world volume is placed a scintillation layer. The 70 µm thick scintillation layer is formed from 1000 × 1000 individual columns of CsI, each a cylinder of 6 µm diameter and separated by 0.1 µm. The optical properties of surfaces can be defined specifically using the toolkit and for the scintillator columns the surface is set to 'ground' which defines the surface

as a multi-faceted (rough) surface. This rough surface aims to mimic the imperfections in the columnar structure of the scintillator when grown in reality.

For the image analysis, a lead slit is implemented above the scintillator. The slit is formed from two lead edges, each 2 cm thick, and separated by 100 µm. The results from the image analysis are discussed later in this section.

In this part of the study, the Geant4 toolkit is being used to examine the event profiles produced when the visible photons (emitted by the scintillator) pass through the imperfect columnar structure of the caesium iodide. The interaction depths of the X-rays in the scintillator are generated using the appropriate interaction cross-sections using the physics lists defined in the toolkit. The interaction depth can be compared with the event profiles, allowing the analysis of event profiles from known exact interaction depths, a practice which is not possible in laboratory testing. In this case, the Geant4 toolkit is being used to analyse the photon passage through the material structure, showing an alternative use to that described in Chapter 7 where the consequences of specific interactions are of greater interest.

3.5.3 Initial data analysis

The visible photons emitted by the scintillator are detected and logged as they pass through the 'sensitive detector' described in the Geant4 code. Each photon is logged in a file quoting the pixel through which the photon passed. A marker is inserted in the output file between each incident gamma-ray so that events related to each gamma-ray photon can be analysed individually.

The list of pixel positions can be parsed using bespoke C++ code (written by the author) to produce a series of 'images'. The positions are histogrammed in Matlab to produce the raw images, an enlarged example of which can be seen in Figure 3.8. The images are restricted to approximately 200 events per frame in order to minimise the overlapping of events.

At this stage, the raw photon counts are analysed with no further image processing. At a later stage, as described later in the chapter, appropriate noise statistics are included in the images for further analysis.

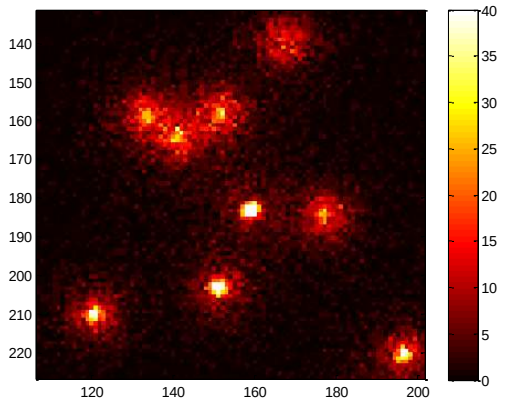

70 µm thick scintillator, CsI(Tl)
16 µm pixels
60 keV incident X-ray energy

Figure 3.8. An example of a region of a raw image generated in the simulation, containing nine event profiles at random positions in the scintillator (where the *x,y,z* position of the interaction in the scintillator is determined by the appropriate statistics). The events can be seen to vary dramatically in peak intensity and spread. This variation is explained in section 3.5.4 through depth of interaction effects.

3.5.4 Depth of interaction effects

As the visible photons generated in the scintillator pass through to the fibre-optic plate the photons will spread out radially. In a structured scintillator, the initial photon spread is radial, but when hitting a column wall, the photon may be totally internally reflected and in this way may pass down the columns to the fibre-optic plate. The 'rough' surface of the columns ensures that the process is not perfect and hence the spread is somewhat less predictable.

The depth of the interaction directly affects the spread of the visible photons on their passage to the fibre-optic plate. Those photons interacting near the upper surface of the scintillator (such as profile 1 in Figure 3.9) will spread by a large degree before reaching the fibre-optic

plate. Events resulting from interactions deeper in the scintillator (such as profile 3 in Figure 3.9) will be found with higher peak signals and smaller spreading of the signal across multiple pixels.

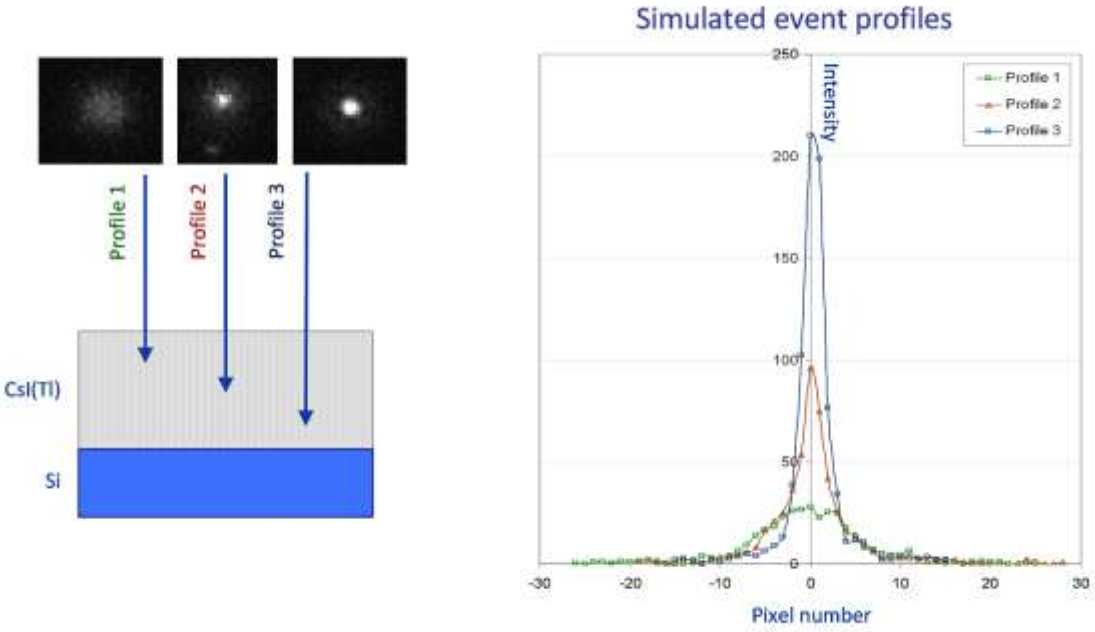

Figure 3.9. Three simulated event profiles with different interaction depths. The integrated signal remains constant (integrating in two dimensions) whilst the peak and width vary. Deep interactions provide higher peak intensities with sharper profiles.

In Figure 3.9 (right side), a slice through the centre of each of the three demonstrated profiles shows how the profiles vary with depth of interaction in the scintillator. The peak intensity and spread of each event varies dramatically and this shows the 'peak intensity' method of energy determination to be inaccurate and incomplete. The profiles show an approximately constant integrated signal (integrating over two dimensions), implying that it is the integrated signal which must be considered for any energy reclamation calculations. The variation in peak intensity of over one order of magnitude shows the implications of using basic energy

determination techniques. An approximation to the integrated signal can be calculated through the summation of signal over many pixels or through the fitting of profiles to the events, reclaiming the peak and spread of the event to allow the calculation of the total signal (Section 3.6).

3.5.5 A simulation of the imaging capabilities of the system

The raw data from the simulation can be analysed to form frames containing a set number of events (defined by the user). Taking approximately 200 events per frame, a series of images was produced. This set of images will be referred to from this point as the raw images. The images were simulated using a 100 µm lead slit and using a 70 µm thick collimated CsI scintillation layer. The incident gamma-ray energy has been set at 140 keV to simulate the upper end of the energy range to be tested in the laboratory (narrowly above the maximum gamma-ray energy emitted by ^{57}Co).

Taking the raw images as a starting point, one can create three main types of image, one based on the integrated image method and two based on centroiding techniques. The integrated image is created by summing all the events into one image. A simulated integrated image of the 100 µm slit is shown in Figure 3.10. The FWHM of the imaged slit is approximately 380 µm, showing a large degradation in image quality through the use of the integrated method. Each detected X-ray event contains a certain spread in signal and by summing the raw events, one also sums the spread of all the events. This summation of the event profiles blurs the slit over a much wider area.

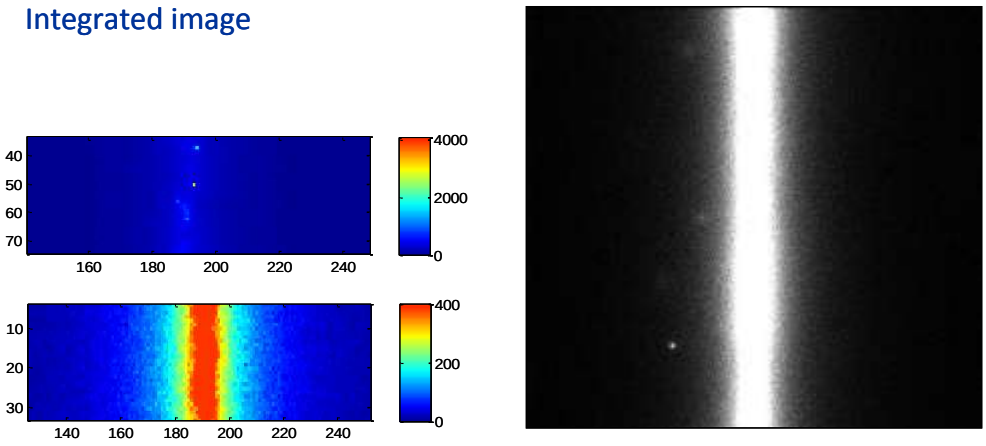

Figure 3.10. The integrated image of the 100 µm slit from the simulated data, created by summing all the events into one image. The images shown left are labelled in pixels (shown at two different intensity scales), with a pixel size set at 16 µm square for the CCD97. The rogue points seen to the left of the slit in the image shown right are X-rays which have passed through the lead shielding.

The 'maximum intensity' image (a form of 'photon-counting'), Figure 3.11, is produced by first searching each raw image for events using the point search function described earlier. Each event is then labelled by the pixel in the event profile containing the highest signal. This method is affected by the Poisson statistics on the detected photons (shot noise), meaning that some events are 'centroided' to the wrong pixel. This photon-counting method, although not the most accurate, does provide much better results than the integrated image method and the 'centres' of the events can be calculated more quickly than with the more accurate method detailed below. The FWHM of the slit is equivalent to seven pixels (112 µm), and here is only limited by the pixel width (choosing a location to the nearest pixel).

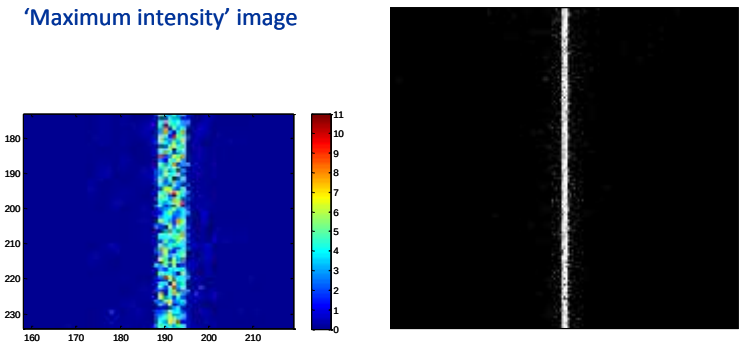

Figure 3.11. The 'maximum intensity' image of the 100 µm slit (from the same simulated data-set as displayed in Figure 3.10). The image is created by taking the pixels in the event profiles containing the highest signals. The image shown left is labelled in pixels, set at 16 µm square.

The second 'photon-counting' method examined here uses a full centroiding methodology. The centroid is taken for each event in the raw images. The images shown in Figure 3.12, for comparison with previous results, show these centroid positions histogrammed in two dimensions to 16 µm pixels. At first glance, little difference can be seen between the two centroiding methods, but on closer inspection near the edges of the slit, one can see a decrease in the number of events located outside of the slit position.

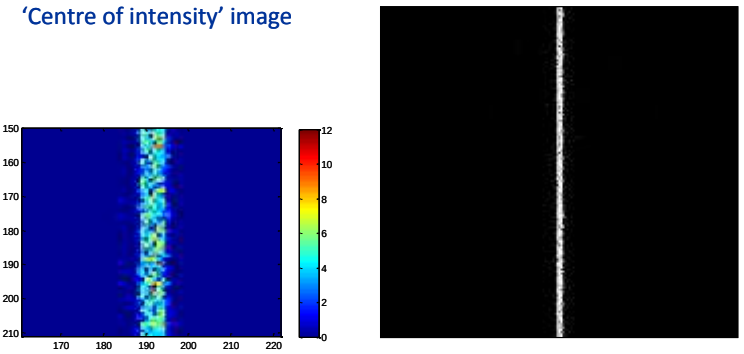

Figure 3.12. The 'centre of intensity' image of the 100 µm slit (from the same simulated data-set as displayed in Figures 3.10 and 3.11). The centroid locations of each profile are combined to form the image. The image shown left is labelled in pixels, set at 16 µm square.

The summed signal across the slit is plotted in Figure 3.13 and this shows more clearly the differences between the three methods. The photon-counting technique shows a large improvement in the FWHM of the Line Spread Function (LSF), where the LSF is taken as the sum of all the rows of signal across the slit, with the centre-of-intensity measurement providing small improvements over the maximum-intensity method. The two photon-counting methods provide similar results on larger scales. When examined in finer detail it is seen that the 'centre of intensity' method provides a cleaner edge. This can be seen specifically by looking at the region between pixels 180 and 200 where the two photon-counting methods differ at the base of the LSF. The slight blurring of the edges when using the maximum intensity method arise from the shot noise on the signal causing a fluctuation in the signal level in each pixel – the pixel with the highest signal in each event is not necessarily the central pixel.

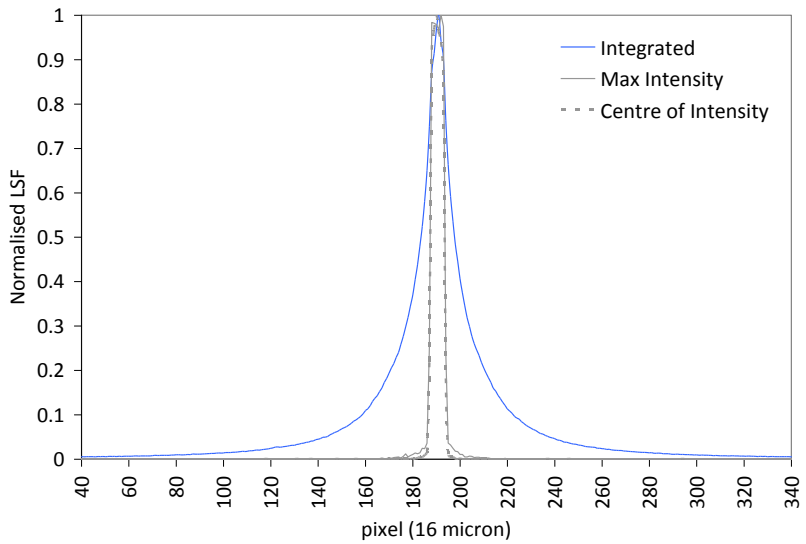

Figure 3.13. The normalised signal summed across the 100 µm lead slit for the three imaging methods simulated above. The improvements in using photon-counting over the integrated methodology are clear.

3.6 Energy discrimination through event profiling

In order to determine the energy of an incident X/γ-ray, one must choose an appropriate method of analysis. Taking the peak intensity of an event does not give an accurate indication of the energy of the event. The peak intensity is altered by many factors, namely the depth-of-interaction in the scintillator and the dual shot-noise apparent on the signal (input and gain). Events of identical energy will give very different results depending on the distance from the image plane at which the scintillation occurs. The shot noise on the event is not only present due to the Poisson statistics of the incident photons, but the variance increases by a factor of two due to the multiplication in the gain register (Section 2.3).

Several methods have been suggested in the literature to overcome this problem, including the summation of signal in pixels surrounding the event peak. These methods do not give an indication of the width of the event and hence provide only one dimension of information. Fitting to the noisy events using the scale-space methodology provides a further indication as to the origins of each interaction, providing information on the interaction position inside the scintillator.

3.6.1 Comparison of previous methods found in the literature

The use of photomultiplier tubes connected to solid scintillators was discussed in Section 2.5.3, showing poor spatial resolution but countering this with a much better energy resolution of 3-10 %. The photomultiplier tubes allow many more of the emitted photons to be 'detected' as part of the events due to their larger pixel size. The lower spatial resolution records many re-combined events, where fluorescence X-rays are combined with their initial interaction due to the inability to the separate events.

The use of scintillator-coupled EM-CCDs has, in the past, yielded poor spectral results. Whilst some studies have discounted the spectral capabilities of the scintillator-coupled EM-CCD,

either through design or following poor results, several results have been reported in the literature (as discussed in section 2.4).

Early studies, using the peak intensity of smoothed events to discount some of the degradation due to the Poisson noise, give energy FWHMs of approximately 40 keV at 30 keV. The spectra displayed show large, broad peaks over which the FWHM is hard to calculate and for which there is severe blurring of the boundaries between differing energy peaks. Although the use of the peak intensity of the smoothed peaks decreases the effect of the Poisson noise, the depth of interaction in the scintillator has been ignored, significantly decreasing the spectral performance of the device (leading to a FWHM of over 100 % of the energy value).

When calculating the energy resolution through the summation of pixels the FWHM has been reduced to closer to 30 keV but again the spectral peaks cannot be well separated and take the form of those shown from previous studies in Chapter 2. The method gives improved spectral performance, but only gives a single value for the intensity of the event. Although the depth of interaction is in some way included through the integration of the signal over several pixels, any depth effects aside from the geometrical aspects are ignored.

The presence of re-absorbed fluorescence would lead to severely degraded spectral resolution measurements if not fully taken into account. The reduction in peak energy and presence of multiple additional peaks must be fully understood to improve the spectral performance, and consequently the spatial performance, of the gamma-camera.

An attempt at applying ‘most likely estimate’ parameters to the events yielded very similar results to those found through the methods mentioned above. The method does, however, confirm that the depth of interaction effects must be taken into account when considering the energy of the events. The complicated process of calibration is currently not yielding the results expected from such a time consuming method.

This study aims to resolve these problems and allow much improved performance of the scintillator-coupled EM-CCD camera through the more thorough understanding of the processes occurring when imaging hard X-rays in a scintillator.

3.6.2 Event profile fitting methods

Through the fitting of a profile to an event, a greater amount of detail about the event can be found. This greater detail provides further insight into the interactions occurring inside the scintillator and can yield factors needing consideration. The event profiles provide depth of interaction information which is otherwise lost through peak intensity and pixel summation methods. A least squares fitting algorithm is considered here and compared to the scale-selection method (discussed in Section 3.4).

In order to compare the scale-space analysis against a standard least-squares fitting routine it was necessary to investigate pre-determined Gaussian-like profiles. An image area of 41 × 41 pixels was created, with the events centred at (20,20).

3.6.2.1 Simulating Gaussian profiles for fitting method analysis

The spread of the signal in the event profile is determined by the depth of the interaction of the initial X-ray in the scintillator and is hence variable from a single pixel out to the width where detection is limited by the noise floor. The spread of the Gaussian has been randomly generated about a mean value (chosen from initial experimental data), with $\sigma < 0.5$ ignored to avoid excessive run times. The peak has then been calculated from Equation 3.5, assuming a Gaussian formation, where γ is a predetermined value to give a constant total signal in each 'event'.

$$\text{peak} = \gamma\left(\frac{1}{2\pi\sigma^2}\right) \tag{3.5}$$

Following the peak and spread generation, the signal has been generated in the image area, centred upon (20,20), for the perfect noise-less case, Figure 3.14 (left). The initial testing of

the scale-selection process has been implemented using this noise-less data. The results from the scale-selection on noiseless data are shown in Figure 3.15. The results show an accurate depiction of the input data (straight line with a gradient of -½ in logarithmic space) and the tight fitting of the data-points confirms the accuracy of the scale-selection process with the noiseless data.

The Gaussian profiles were then adapted to include shot noise and background noise (readout noise and noise on the dark current). New signal levels are applied to each pixel, randomly generated with Poisson statistics to account for the shot noise on the incident photons. The process is then repeated on the new signal values to account for the shot noise on the gain process (Section 2.3). A background noise level has then been added to each pixel.

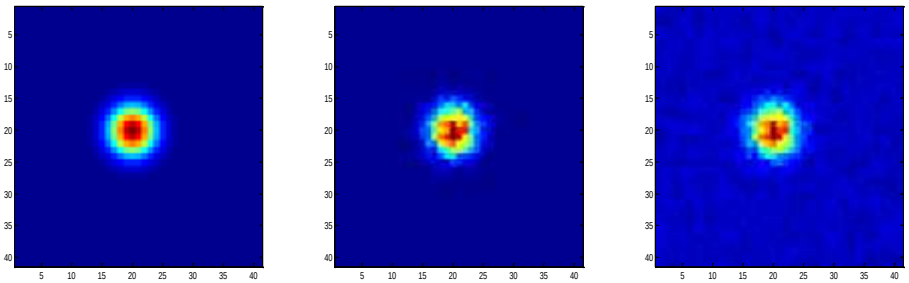

Figure 3.14. A sample Gaussian profile generated for the simulation is shown left. The central image shows the profile generated with shot noise. The right hand image shows the profile with shot noise and a background noise component.

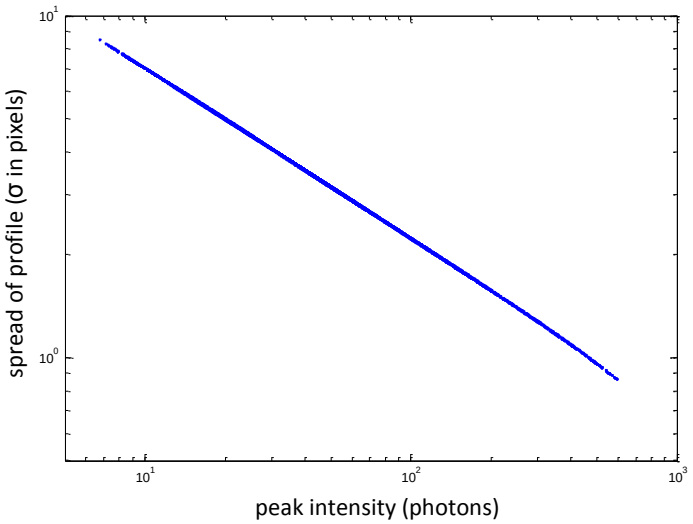

Figure 3.15. The fitted data points show a linear relationship (in log-log space) between the peak intensity and σ, with a gradient of -½. This fits with Equation 3.5, showing that the scale-selection is working as expected with the noise-less data.

Plotting the results from the scale-selection code on the noisy data gives a comparison with the noise-less data, as shown in Figure 3.16. The fitting process loses accuracy at low peak signal (high spread) as the shot noise has a much greater impact on the signal level. If a signal level of 9 is compared to a signal level of 25, the shot noise (from the incident photons and the gain process) gives noise levels of 4.2 and 7.1 respectively, leaving a noise-to-signal ratio of 0.47 and 0.28 respectively. The larger noise-to-signal ratio for lower signal levels leads to a decrease in the accuracy of the fitting method. This loss in accuracy is investigated in further detail below. Taking Equation 3.5 and rearranging for γ gives the conversion from Figure 3.16 back to a total signal (or 'energy') value.

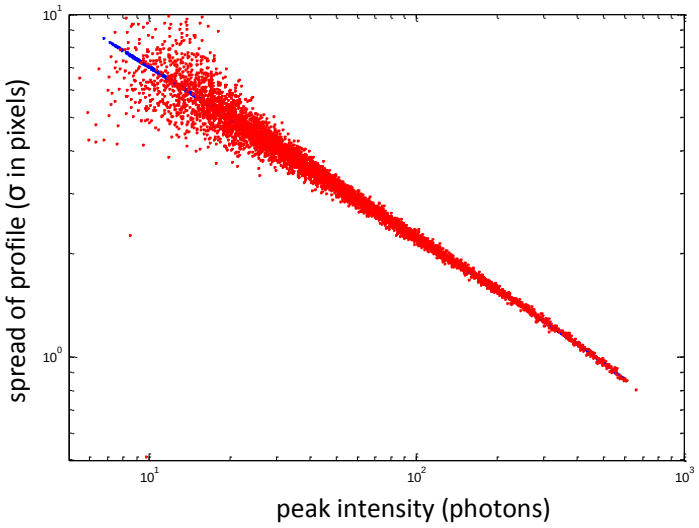

Figure 3.16. The fitted data points to the simulated data (red) on the noiseless simulated data (blue) show a linear relationship in logarithmic space between the peak intensity and σ. The red data points show the results of the scale-selection process on the noisy data, showing an increasing accuracy as the peak intensity increases (increasing the signal-to-noise ratio).

3.6.2.2 Comparison of Scale-Space with least-squares fitting methods

Taking an increasing value of sigma from 0.2 to 6 in steps of size 0.2, images were generated with appropriate shot noise on a background image. The scale-selection code was implemented for images at four peak signal levels of 5, 10, 20 and 40 photons. Alongside the tests for the scale-selection code, a 2D-Gaussian was fitted using a least-squares method to compare the two processes. The results from these tests are shown in Figures 3.17 and 3.18, with the scale-selection method shown in red and the 2D-Gaussian fitting method shown in blue. The results show much higher accuracy in predicted spread and peak for the scale-selection method compared to the least-squares fitting method.

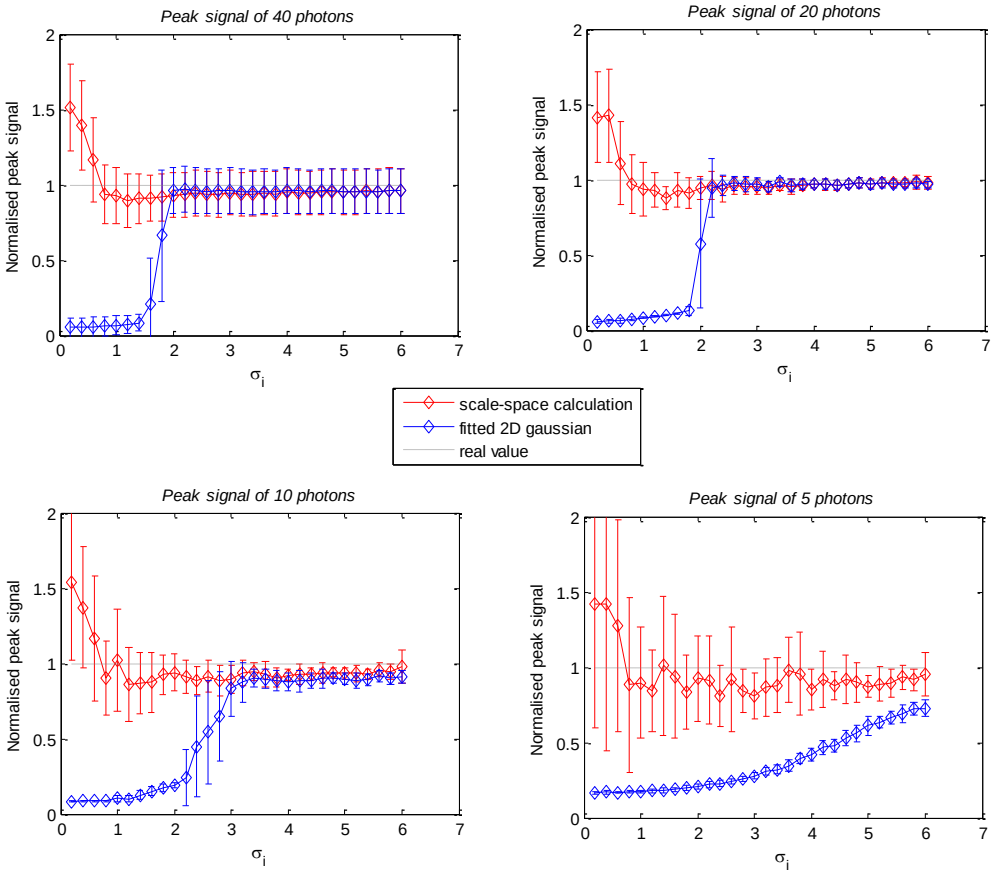

Figure 3.17. The calculated peak signal (normalised) for signal spread sigmas from 0.2 to 6 pixels over four intensities. Readout noise of 1 e^- is included on top of the shot noise and noise from the avalanche gain registers. The scale-space calculations (using the derivations shown in the appendix) provide a much more accurate value over all sigma. The error-bars relate to error analysis applied to multiple runs of the simulation.

The plots in Figure 3.18 show the most striking example of the benefits of the scale-space representation over a least-squares fitting method. The value of σ taken as the input for the simulated signal profile is shown on the *x*-axes, with that calculated by the code shown on the *y*-axes. The blue data points, corresponding to the least-squares fitting of a 2D-Gaussian profile, show large deviations from the input values. Many of the data points for the least-square fits fall over the top of the displayed plots and in many cases only the error bars can be

seen (calculated from multiple runs of the simulation). The red data points, however, follow a linear relationship with a gradient of unity and for higher signal-to-noise ratios have very low errors. The errors at lower peak signals are several orders of magnitude lower than those found for the least-squares fitting method.

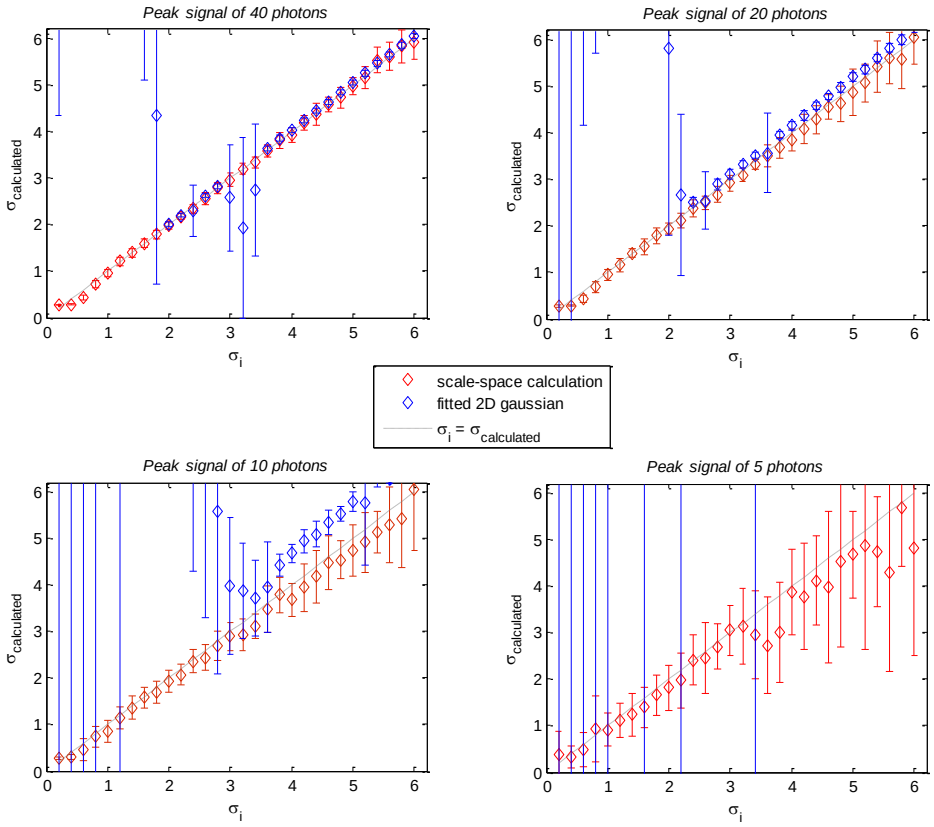

Figure 3.18. The calculated sigma (spread of the Gaussian-like event) for initial sigmas from 0.2 to 6 pixels over four peak intensities. The scale-space calculation (using the derivations shown in the appendix) provides a much more accurate value across all sigma. At low signals, the least-squares fitting algorithm provides results off the scale of the plots shown and only the error bars are visible.

The customised scale-space event profiling technique has been shown here to accurately analyse Gaussian-like profiles under conditions in which standard fitting techniques have failed. The ability of the customised scale-space analysis technique to operate effectively in the very low signal conditions encountered with the photon-counting scintillator-couple EM-CCD detector makes this method the subject for the analysis of the real data in later chapters.

3.7 Scale-space event profiling

The peak profiles detailed previously show how the energy of the incident X-rays cannot accurately be given by the peak intensity of the event profiles. The scale-space methodology can be used to analyse the event profiles delivered by the simulation.

Three incident gamma-ray energies were simulated at 20 keV, 60 keV and 120 keV. The events simulated were generated from the single interaction of each high energy photon and present the transport of the visible photons through the imperfect columnar structure. For the full generation of the appropriate energy spectrum, the specific processes involved in the scintillation mechanism must be modelled. The energy spectrum is considered in more detail in Chapter 5, whereas the scale-space methodology is the main area of interest in the results detailed here.

Following the operation of the point-search function, each event was processed through the use of scale-selection in scale-space. The process is detailed in the flow-chart shown in Figure 3.19. The output from the scale-selection code includes a data file consisting of the position of the interaction and the spread and peak of the event profile. The peak and spread of the event profile can be mapped in a ‘Peak vs. Spread’ map (PvS), as shown in Figure 3.20.

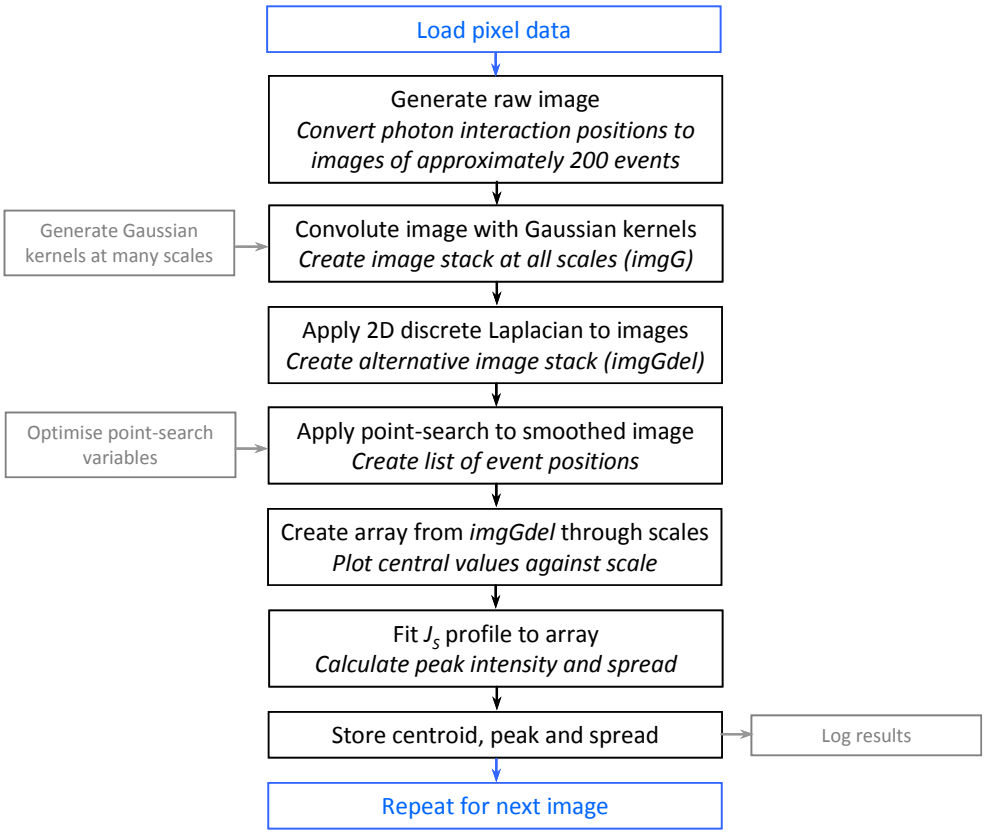

Figure 3.19 A flow-chart of the scale-selection procedure followed in the analysis code.

The PvS map shows a very different view to methods such as peak selection and signal summation where only a single value is obtained. The correlation between the peak and the spread shows a constant integrated profile for each incident gamma energy.

The PvS map also gives an indication of the depth of interaction in the scintillator. The depth of interaction of the initial gamma photon is related to the spread of the event profile. The deeper the interaction in the scintillator (the closer the interaction is to the detector) then the lower the spread on the event as the visible photons are detected earlier in their path.

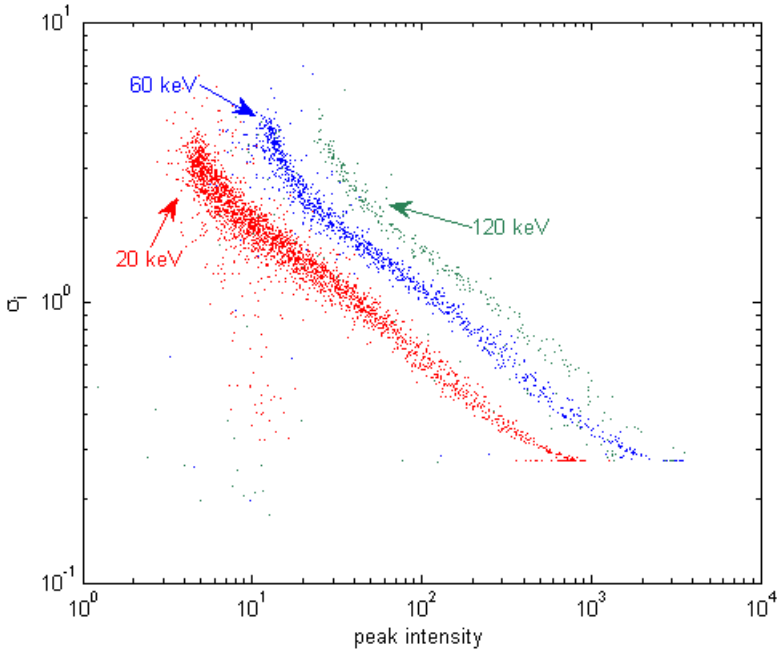

Figure 3.20. The 'Peak vs. Spread' (PvS) map for the simulated data analysed using the scale-selection method. The results are presented for three different gamma-ray energies through a 70 µm caesium iodide scintillating layer.

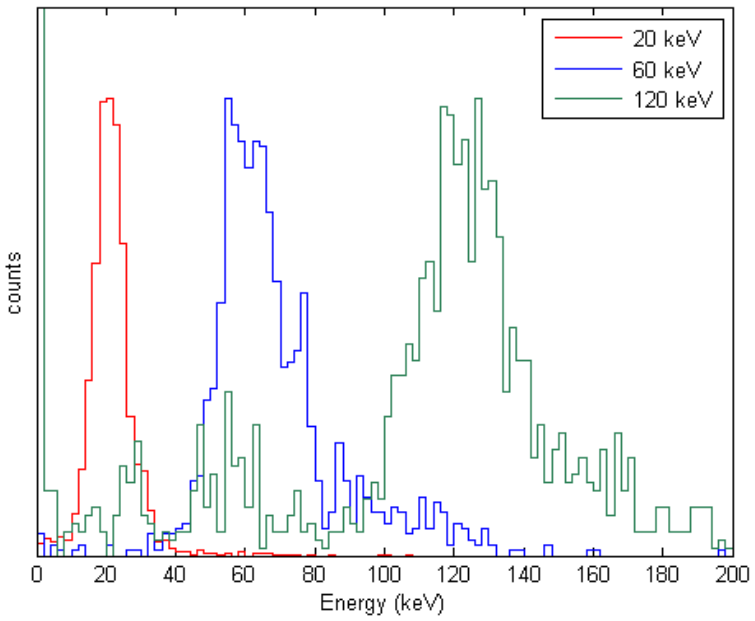

Figure 3.21. The energy spectrum generated using the scale-selection method, generated using the PvS map shown in Figure 3.20, normalised to the peak height. The signal located away from the main peak (particularly visible in the 60 keV and 120 keV cases) is due to the events in which the signal remains mostly in a single pixel (small σ_i) and hence the 'spread', and subsequently the 'peak', of the event cannot be accurately determined.

At low peak intensities, the spread appears to be slightly overestimated and this is believed to be due to the very low signal-to-noise ratio of the low signals due to the shot noise on such few photons. At higher peak intensities, the linearity of the spread determination is lost due to the low spread of the events and the quantisation due to the finite pixel size.

The total signal can be obtained by taking the result shown in Equation 3.5. Using this formula and applying to the data generated by the simulation, the energy spectrum shown in Figure 3.21 was produced. The clear separation of the energy peaks gives a good indication that the scale-selection process can produce 'energy spectra'. The linearity of the energy scale provides further reasoning to progress this method further.

3.8 Simulating the EM-CCD

Following the simulation work carried out to simulate the passage of visible photons through the scintillator, it is necessary to consider the output of an image from the EM-CCD. The results taken from the Geant4 simulation show event profiles which, although they include the shot noise on the signal, do not include the readout noise and noise due to the gain mechanism. The influences of dark current and clock-induced charge (CIC) are also not included.

The simulation detailed here processes the full output chain from the EM-CCD. Basic operation of the simulation produces a background image showing the running of the EM-CCD in the dark. The signal from the Geant4 simulation can be included in the simulation such that the output from the scintillator-coupled EM-CCD under bombardment from gamma-rays can be predicted.

The simulation has many uses, including aiding the analysis of anomalous results and solving problems arising in the characterisation process. One example of the use of the simulation in the characterisation process is given in Section 4.5.1, resolving the issues of the readout noise 'tail'. The optimal gain required by the device can also be considered to give the best performance and is discussed.

3.8.1 EM-CCD operation chain

The simulation is based around the probabilities of charge generation and multiplication with each transfer through the device. The code works systematically through device readout from integration to the output of the signal following passage through the gain register. All stages are governed in the same way: a random number is generated between zero and one and compared to the probability (zero to one) of charge generation/multiplication. If the random number is below the probability then charge is generated, otherwise the charge moves to the next pixel.

The device must be modelled according to the nature of the CCD, Figure 3.22. The image and store regions of the device are run in inverted mode with the readout and gain registers given the choice of inversion or non-inversion. The temperature of the device can be altered through the modification of the charge generation rates.

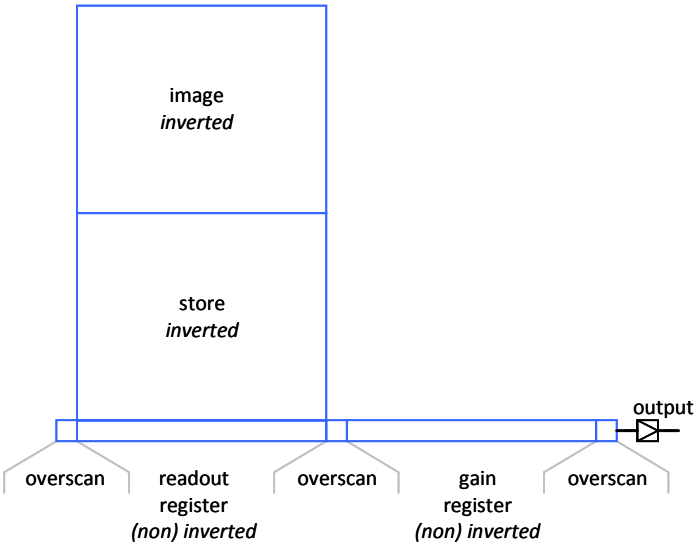

Figure 3.22. The main regions of the CCD97 to be simulated.

If the device is run in non-inverted mode then the CIC can be assumed to be negligible at approximately 0.003 electrons/pixel/frame (Low-Light technical note, 2004). Under inversion, this increases to approximately 0.1 electrons/pixel/frame, giving one electron on average from CIC per 10 pixels. In the simulation, the CIC is generated from the probability of generation per line transfer and gives a rate of 1×10^{-4} electrons/pixel/transfer in inverted mode. The CIC generation increases in the gain register due to the higher voltages across the electrodes (Section 2.3).

The dark current generation rate can be edited to simulate different readout frequencies and temperatures. The plot shown in Figure 3.23 gives the standard level of dark current under

two-phase inverted mode operation in electrons per pixel per second. The readout rate and temperature can be altered in the simulation by adjusting the dark current generation rate.

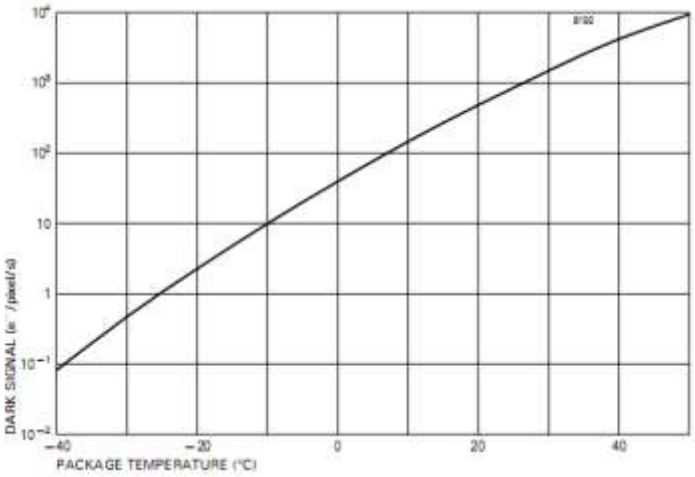

Figure 3.23. The CCD97 dark signal generation profile for varying temperature in two-phase inverted mode operation (CCD97 datasheet, 2004).

For initial testing, the data sheet values can be used. For future analysis using the EM-CCD simulation one can change the variables to suit the operating conditions required. The images produced with the simulation show the same properties as those achieved experimentally in the presence of no X/γ-ray source. Figure 3.24 shows the simulation output at multiplication gains of 1 and 100 respectively using the device parameters as specified on the device datasheet, running the device at 150 kHz pixel readout rate (as with the early experimental results). Signal can be added to the simulation of the CCD97 output and this is discussed in section 3.9.

3.9 Simulation of the full output chain

The flowchart shown in Figure 3.25 details the full output chain simulation. When including signal, an array of electron/photon values must be set. The signal array can be generated from

the Geant4 scintillation simulator output using automated code. The scintillator output has been set to a slightly smaller area than the CCD97 area due to time constraints on processing, and as such, the signal included in the simulation does not cover the entire image area (as can be seen on the images).

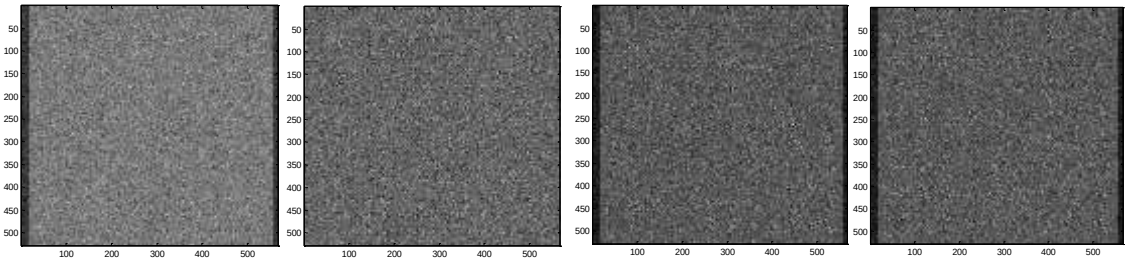

Figure 3.24. The simulation output at unity gain (images 1 and 2) and at a multiplication gain level of 100 (images 3 and 4). The left images (1 and 3) show the signal in electrons. The right images (2 and 4) show the signal in ADU. The finer structure of the device (e.g. the dark reference columns) are lost at unity gain due to the readout noise. At higher gains the device structure becomes apparent due to the reduction in the effective readout noise.

The results detailed below were simulated at 150 kHz, with a calibration factor of 75 e^-/ADU (Analogue-to-Digital Unit). The parameters used for the calculations involved are taken from the datasheet (set at -15 °C). An efficiency for the device was set at 50% to account for the loss in photons at the interfaces between the fibre-optic plate and the EM-CCD and the quantum efficiency of the device.

At unity gain, the simulated signal does not appear above the general background noise of the device (consisting of the noise on the dark current, the CIC if the gain register is inverted and, more dominantly, the readout noise). Increasing the gain of the device increases the signal before readout, hence reducing the effective readout noise (Section 2.3).

Figure 3.26 shows a selection of output images from the simulation. At unity gain, the signal (in ADU) shows no variation across the device. The readout noise masks the fine structure and the signal is lost. When the image is smoothed to suppress the noise, the signal is still not visible and a regions of higher noise (by random chance) appear as possible events.

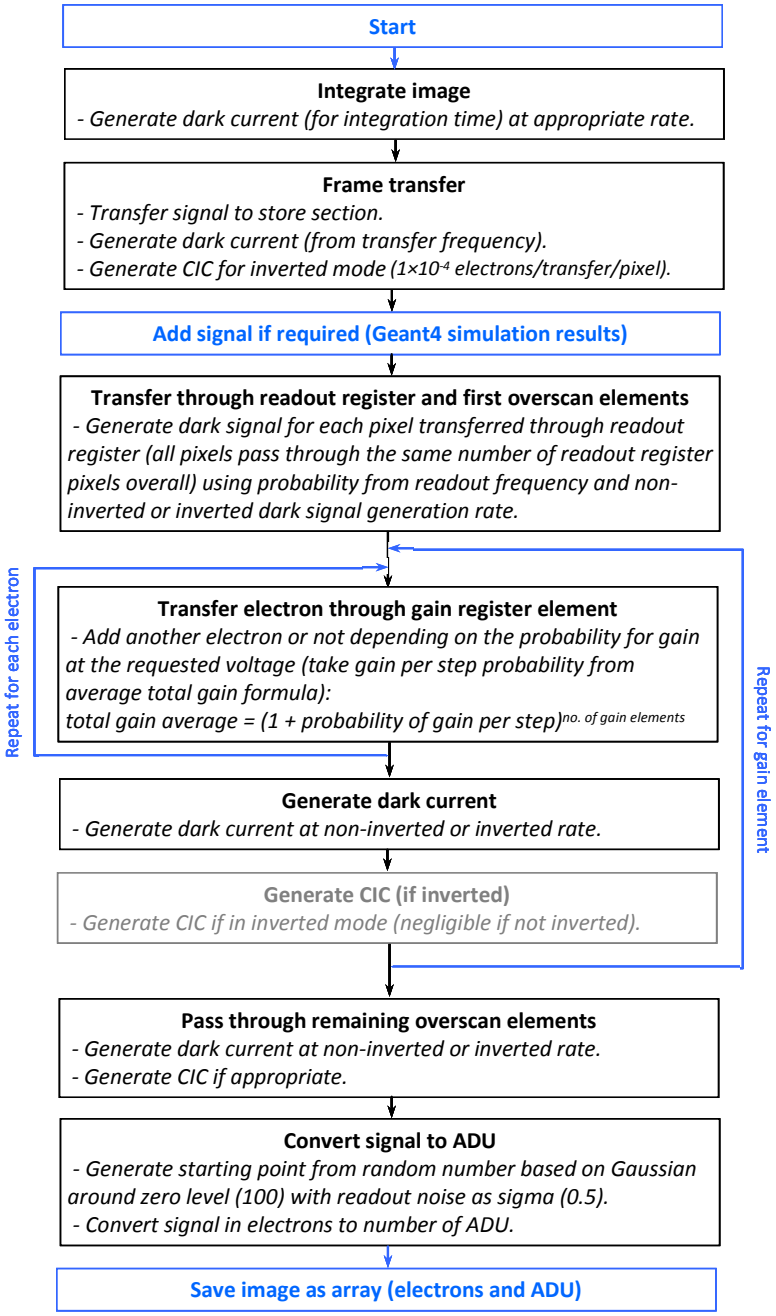

Figure 3.25. A flowchart to show the operation of the EM-CCD simulation code.

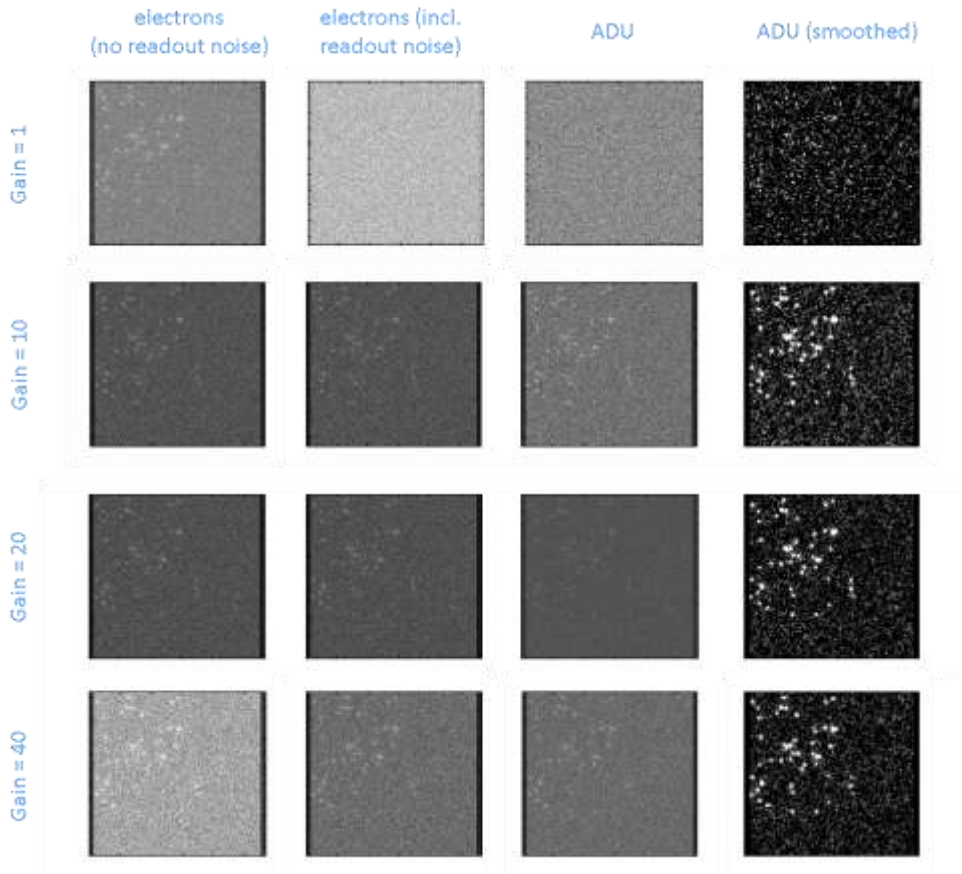

Figure 3.26. A selection of output images from the simulation of the CCD97 output under increasing multiplication gain.

At higher gain the effective readout noise is sufficiently reduced such that the signal electrons can be seen above the background variations. At higher gains the readout noise is no longer the dominant noise factor and the shot noise on the dark current and CIC take precedence. The shot noise on the number of photons detected in the CCD cannot be improved upon through increasing the gain and the noise on the dark current (and CIC) cannot be improved through the use of higher multiplication gain.

The extra noise factor (Section 2.3) increases rapidly for increasing gain below a multiplication gain level of × 10 but the square of the extra noise factor only rises from 1.9 to 2 at higher gains. As a multiplication gain of greater than × 10 is preferred to reduce the readout noise

sufficiently, the extra noise factor can be considered to have negligible impact in the gain selection process.

3.10 Event location through centroiding

For improving the resolution the photon-counting method has many advantages over the integrated image. Although the computational time required to generate a photon-counted image over an integrated image is much longer, the resolution achieved can be greatly improved. The resolution in an integrated image is limited by the spread of the visible photons from the point of interaction of the initial gamma photon in the scintillator. In photon-counting mode, each event profile can be reduced to a single point. The differing methods available for determining the location of the single point need to be carefully considered, as all are relevant for different cases. The search for the optimal centroiding technique for this study is discussed below.

The most basic form of 'centroiding' is to take the pixel with the highest intensity from the raw images. This method requires the least computational time and gives a grid size equal to the pixel spacing (determining the ultimate maximum resolution). This method is, however, subject to errors from shot noise. In all events the signal levels are subject to Poisson statistics (shot noise) such that the maximum intensity in the event profile may not be the real centre of the profile. This is particularly noticed in profiles with large spread as neighbouring pixels share similar signal levels, resulting in the shot noise having a greater anomalous effect.

In order to avoid the deviations in centroid accuracy caused by the shot noise, the image can be convoluted with an appropriately scaled Gaussian kernel. The Gaussian kernel, as described in the theory behind scale-space in Chapter 2, will act to suppress noise in the signal and emphasise the Gaussian-like nature of the event profiles. The resultant profile will, although having a larger spread, maintain the same central location and as such the position will not be affected. The pixel of highest intensity can then be regarded with much greater confidence as

the central pixel of the event. In general, a kernel of the same approximate width as the event profile will provide sufficient smoothing (without including excessive noise). This method will leave the same grid size as the previous method, with the best-case resolution determined by the pixel spacing.

To achieve sub-pixel positioning, further analysis of the event profiles is required. Centroiding techniques are widespread and vary greatly from method to method. In order to estimate the interaction position to a sub-pixel accuracy one must consider more than only the central pixel of the event. The surrounding pixels will contain lower signal levels than the central pixel and will hence suffer from a lower signal-to-noise ratio. In order to define the central pixel of an event, the most effective method uses a Gaussian kernel to smooth the image and suppress the shot noise. The same convolution can act to decrease the influence of shot noise on the outer pixel signal levels, but there are some negative consequences to this application. As shown earlier in Equation 3.2, the convolution of a Gaussian with a Gaussian-like profile increases the spread of the event, such that the new σ is equal to the sum in quadrature of the original profile σ and the Gaussian kernel σ. Although the smoothing will act to reduce the noise in the outer pixels, the increase in spread of the event leads to problems when calculating the centroid position. These 'edge effects' can be seen in the simulated centroid calculations presented below.

Taking the most extreme example the centroid 'edge effects' can be explained. If an event falls on the boundary between two pixels then each pixel will have equal signal and the peak intensity cannot be obtained. In the presence of shot noise however, one pixel will have slightly higher signal than the other and hence, in practice, a 'central pixel' may always be selected. On selection of the central pixel of an event profile, an area of pixels is selected surrounding this 'centre'. As the real centre falls between two pixels, the area over which the centroid is calculated (through a 'centre of mass' technique) will not sit symmetrically over the

event profile, causing a weighting away from the real central location. The closer to the edge of a pixel the real event centre lies, the greater the error in the centroid calculation.

This process can be countered by selecting a larger area of pixels, but some error will still remain due to the calculation favouring the nearest neighbour to the event centre. The greater the proportion of the total signal contained in the pixel area analysed, the greater the accuracy of the centroid, although with larger areas the shot noise and readout noise can further affect the accuracy of the centroid location.

The smoothing of the image using the previously discussed Gaussian kernel will act to suppress the readout noise and shot noise. The smoothing does, however, increase the spread of the event and requires a much greater area to be considered to minimise the errors in the position calculation. By increasing the area over which the centroid is calculated the computational requirement increases, as does the influence of the noise on the image. As the area is increased more outer pixels are included in the calculation, including pixels with much lower signal-to-noise ratios and the readout noise/dark current dominated signal.

The balance between computational time and noise (the area over which the calculation is taken and the size of the convolution kernel) must be carefully considered and the most appropriate solution obtained. Simulations using the Matlab software package illustrate these points and are detailed below.

With no noise on the Gaussian event profiles (with randomly generated but known central locations) the centroid calculation was implemented over a 5 × 5 pixel area (Figure 3.27). The results show that the centroid accuracy decreases as the spread of the event profile increases, as described above. For the smoothed images, where the spread of the events is increased, the results show greater deterioration. In the noiseless case, the centroid positions can be corrected to account for these problems, following a near-linear error function (Figure 3.28). This correction is not so successful however in images in which noise is present.

In reality, the images produced cannot be noise-free. This means that the above results can only be considered indicative to the real scenario. In order to analyse the centroiding technique a background level of noise must be added. In addition to the background noise level, a further shot noise component must be added to the signal.

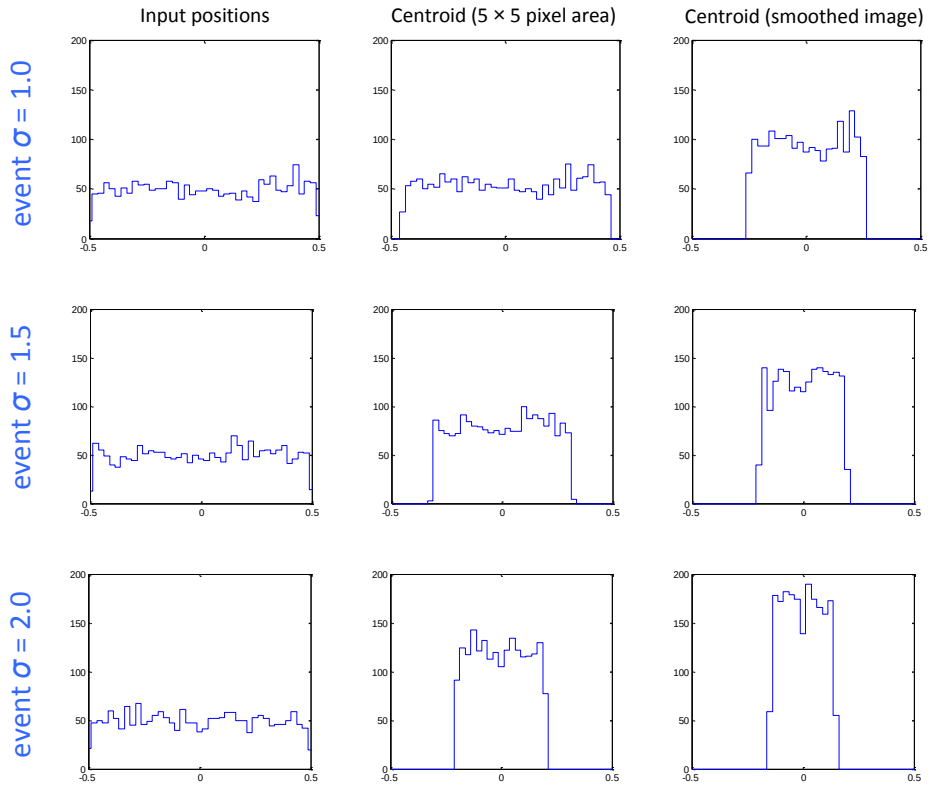

Figure 3.27. The centroiding accuracy (vertical axes) for images with no noise taken over a large number of randomly generated event profiles (event location on horizontal axes). The left column shows the real positions of the events generated across the central pixel. The middle column shows the calculated centroid position across a 5 × 5 pixel area. The right column shows the centroid positions on the smoothed event profiles.

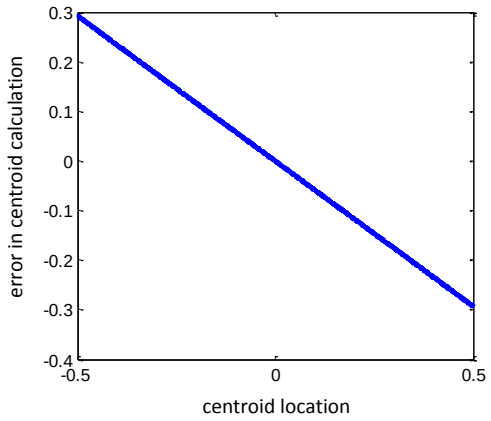

Figure 3.28. The error in the centroid calculation (shown here for the third example in Figure 3.27) is linear with the centroid location. This allows a correction to be included to accurately pin-point the event centre. However, these results are dependent on the noise-free nature of the simulated profiles.

Profiles generated with appropriate noise statistics have been analysed for varying centroid areas, ranging from 3 × 3, through 5 × 5 to 7 × 7 pixels, centred on the peak intensity in the noise-suppressed (smoothed) images. The results from these tests are shown in Figures 3.29 to 3.31. Moving clockwise from the top left, the first plots show the centroid locations of the randomly generated Gaussian profiles. In all three cases, the centroid locations generated spread evenly across the whole pixel area. The second plots show the centroid locations calculated from the raw images (no noise suppression). The fifth plots relate to the upper-middle plots to show the shift in the calculated centroid position from that originally generated. The third and fourth plots equate to the middle plots for the analysis of the noise-suppressed (smoothed) images.

The first area to be considered is a simple 3 × 3 region surrounding the central pixel. The results (Figure 3.29) show the same linearity in the calculated-to-real centroid position as found with the noise-free results. However, although the linearity holds true, the noise on the centroid location shift has increased dramatically, giving a noise on the correction of

approximately ± 0.1 pixels for the raw images due to the noise on the signal. Smoothing the images acts to increase the correction required, but the noise on the correction is reduced. For larger event profiles (larger σ), using a constant area for the centroid, the gradient of the required correction increases. Whilst maintaining the same centroid area, increasing the spread of the event includes a lower proportion of the total signal in the centroid area, hence requiring a greater shift to correct for the offset described previously.

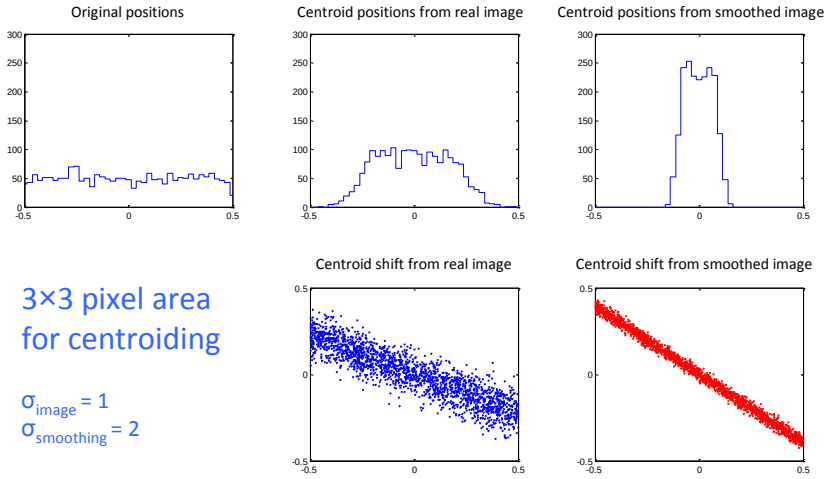

Figure 3.29. Taking a 3 × 3 pixel area for centroiding, with image and kernel parameters as specified, provides the least noise on the correction needed to be applied, but the correction is greatest of the three data sets.

Increasing the centroid calculation area to 5 × 5 pixels includes more of the total signal, requiring a smaller shift in the calculated centroid position (Figure 3.30). This is shown further still when using a 7 × 7 pixel area (Figure 3.31). The noise on the correction required increases as the centroid calculation area increases because the noise included in the centroid area is increased.

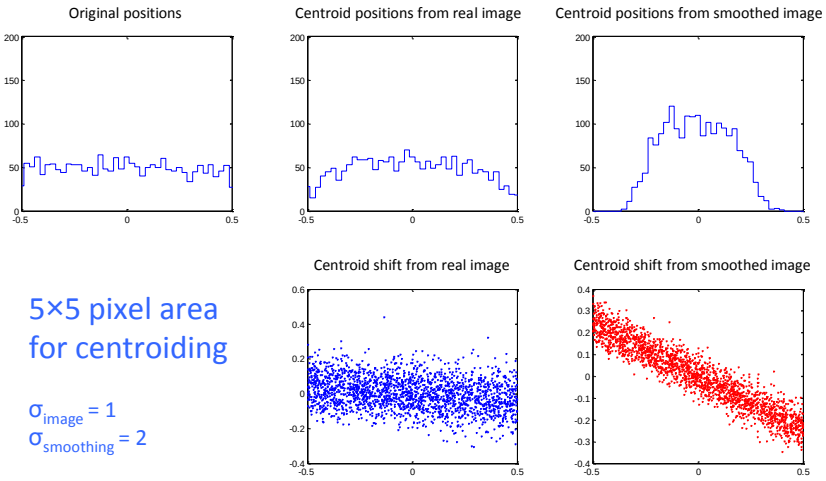

Figure 3.30. Taking a 5 × 5 pixel area for centroiding, with image and kernel parameters as specified, requires a smaller correction on the centroid position, but the noise on this correction is increased over the 3 × 3 centroid calculations.

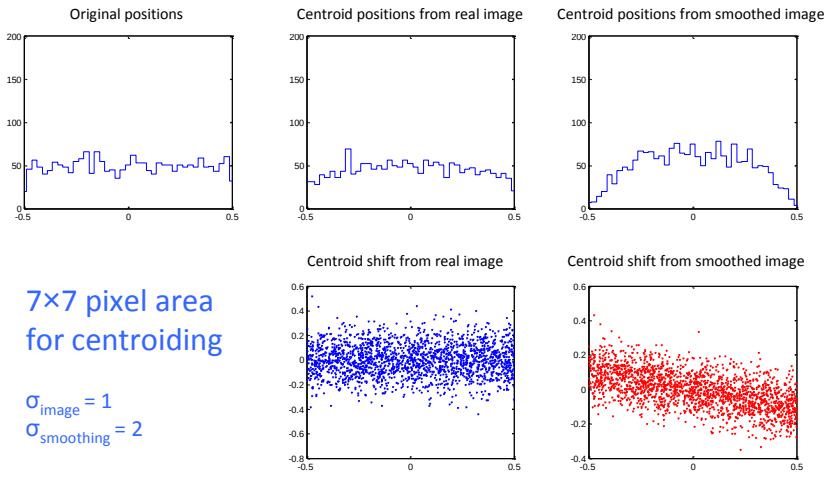

Figure 3.31. Taking a 7 × 7 pixel area for centroiding, with image and kernel parameters as specified, leads to the smallest required correction, but the error on the correction is much larger.

Rather than centroiding the event profiles as described above where the location cannot be determined accurately, a simpler method may be more appropriate. Using a less complex calculation algorithm can dramatically increase the speed of locating the event positions and may not decrease the accuracy to a great degree. A balance must be obtained between the speed of the centroiding procedure and the accuracy of the results. It is also worth considering the overall limiting factors on the resolution. If the centroiding algorithm is not the limiting factor and, for example, it is the non-uniformities in the scintillator causing the greatest detrimental effects on the resolution, then a less complex algorithm is more appropriate. The balance must also consider the accuracy of the result obtained. Producing a value for the location at the sub-micron level is possible, but if the location is only accurate within several microns, then this method is inappropriate. It is for these reasons that the method described below was designed and selected.

The method involves taking each event profile and locating the central point to the nearest quarter of a pixel. This provides sub-pixel imaging, but does not show a disrespect for the maximum accuracy achievable with the centroiding technique.

The outer pixels of an event profile are subject to much lower signal-to-noise levels. Taking only the central pixel and the four nearest neighbours minimises the effect of the shot noise on the smoothed event profile. The central point can be allocated to one corner of the pixel through two simple comparisons. The first comparison takes the pixels above and below the central pixel and compares the intensity. This decides the half of the pixel in which the central point lies in the columnar direction. The second comparison takes the pixels to the left and right of the central pixel and this determines the half of the pixel in which the centre lies in the row direction. With knowledge of the two halves, the intersecting corner is found and this provides the approximated position for the centre of the event profile.

Using this method, taking the centre to the nearest quarter of a pixel, with a CCD97 of pixel size 16 µm the central point can be determined to a precision of ± 4 µm in each direction (Figure 3.32). The use of two simple comparisons leads to much lower run times than the more complex centroiding methods previously described. The new pixel size of 8 µm provides a best-case resolution of 16 µm. Provided the measured image resolution is between approximately 16 µm and 32 µm (the resolution limit for the real pixels) this method is the optimal choice.

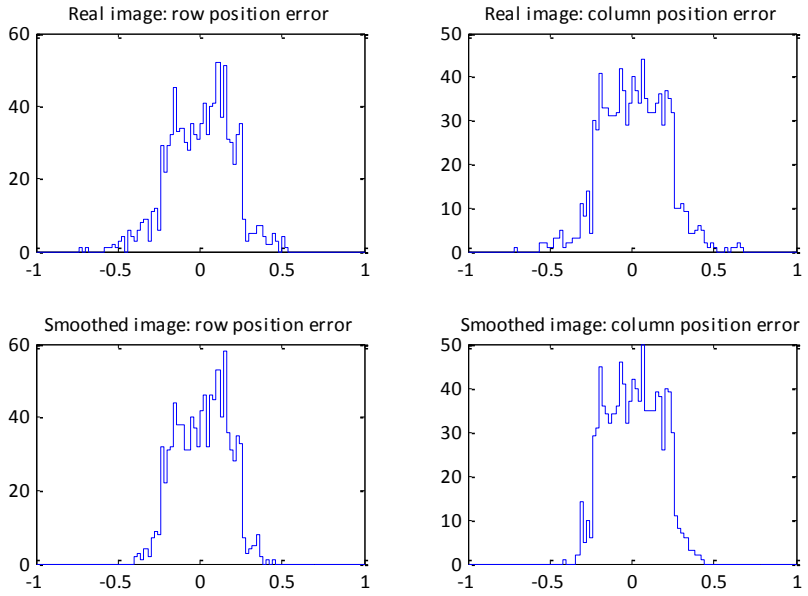

Figure 3.32. The error on the located position should lie within ± 0.25 pixels, but in the noisy conditions simulated here, whilst the code is no longer perfect the results are impressive, particularly when using the smoothed images.

3.11 Summary

The studies detailed in this chapter not only act as a feasibility study for the detector, but also aid in the understanding of the processes which lead to the formation of the images from the

initial interactions of the X-rays/gamma-rays in the scintillator. It has been shown that the number of photons forming the event profile at the EM-CCD is very low and, as such, that the device must be optimised carefully to subsequently reduce the noise in the system. The events must themselves be extracted from the raw image data and the parameters used in the point-search function must be carefully selected to match the performance specifications. The point-search function is visited again in Section 4.7, making use of real experimental data. The use of the scale-space methodology in the analysis of the Gaussian-like event profiles has been shown to be the most accurate method for characterising the parameters of each X/γ-ray interaction in the scintillator. Standard least-squares fitting methods do not allow accurate parameterisation of the events profiles and as such one cannot deliver appropriate depth of interaction information. With the use of the scale-space methods, it is possible to produce a 'peak vs. spread' map that, in future work, may be used in a more developed way to achieve a better spectral resolution.

The Geant4 toolkit has been used to simulate the passage of photons through the scintillator and these simulations show how important the depth of interaction effects are on the peak intensity of the event profiles. Previous studies on the spectral capabilities of the scintillator-coupled EM-CCD in which the peak intensity is used as an indication of the energy of the interacting photon have therefore been shown to be lacking in accuracy.

The use of simulations has again shown how effective the photon-counting technique can be on improving the spatial resolution over the standard integrated image, giving a large improvement in the FWHM in the simulated data. The design and implementation of a full simulation of the detector chain (scintillation through to the readout from the EM-CCD gain register) has given a very useful tool to be used in the next chapter, in which the device optimisation is studied in more detail. The simulation allows the physics behind the effects observed to be investigated.

Finally, studies of centroiding processes have shown that the accuracy of the centroid location from pixellated data, such as the event profiles obtained in this study, leads to significant edge effects due to the low signal across very few pixels. Although these edge effects can be corrected for in the noiseless case for a single image profile width, the error on the correction in low signal-to-noise conditions is of the order of the correction itself and the centroiding techniques break down. The added noise component from the need to know the image profile width in advance in order to calculate the required shift function also complicates the standard centroiding procedures further still. To account for this problem, a simplified approach has been suggested in which the events are located to the nearest quarter of a pixel, giving the event locations accurately within ± 4 μm, without the suggestion of any greater accuracy being achievable. When compared with the results found in Chapter 5, the quarter pixel 'centroiding' method is found to be the most appropriate method for the finding the location of events to the sub-pixel level.

3.12 Appendix: Mathematics of Scale-Space formulation

This short appendix details a more complete derivation of the scale-space theory adapted for the purposes of this study.

3.12.1 The image after convolution

If the structure of interest in the image is taken as Gaussian-like in two dimensions then the following analysis can be applied.

Let f_i be the form of the structure in the image and f_k be the kernel.

$$f_i = A'\exp\left(\frac{-x^2}{2\sigma_i^2}\right)\cdot A'\exp\left(\frac{-y^2}{2\sigma_i^2}\right) = A\exp\left(\frac{-\left(x^2+y^2\right)}{2\sigma_i^2}\right) \tag{A3.1}$$

$$f_k = \frac{1}{\sigma_k\sqrt{2\pi}}\exp\left(\frac{-x^2}{2\sigma_k^2}\right)\cdot\frac{1}{\sigma_k\sqrt{2\pi}}\exp\left(\frac{-y^2}{2\sigma_k^2}\right) = \frac{1}{2\pi\sigma_k^2}\exp\left(\frac{-\left(x^2+y^2\right)}{2\sigma_k^2}\right) \tag{A3.2}$$

$$I(x,y) = f_i * f_k = \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f_i(m,n) f_k(x-m, y-n) dmdn \tag{A3.3}$$

$$I(x,y) = \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} A \exp\left(\frac{-\left(m^2+n^2\right)}{2\sigma_i^{\,2}}\right)\frac{1}{2\pi\sigma_k^{\,2}}\exp\left(\frac{-\left((x-m)^2+(y-n)^2\right)}{2\sigma_k^{\,2}}\right) dmdn \tag{A3.4}$$

$$I(x,y) = \frac{A}{2\pi\sigma_k^{\,2}}\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \exp\left(\frac{-\left(m^2+n^2\right)}{2\sigma_i^{\,2}} + \frac{-\left((x-m)^2+(y-n)^2\right)}{2\sigma_k^{\,2}}\right) dmdn \tag{A3.5}$$

Through the use of the substitution given below, the result can be calculated more easily.

$$\text{let } -\frac{Q}{2} = \frac{-\left(m^2+n^2\right)}{2\sigma_i^{\,2}} + \frac{-\left((x-m)^2+(y-n)^2\right)}{2\sigma_k^{\,2}} \tag{A3.6}$$

$$I(x,y) = \frac{A}{2\pi\sigma_k^{\,2}}\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \exp\left(\frac{-Q}{2}\right) dmdn \tag{A3.7}$$

$$\text{let } Q = Q_m + Q_n \text{ to give } Q_m = \frac{m^2}{\sigma_i^{\,2}} + \frac{(x-m)^2}{\sigma_k^{\,2}}, Q_n = \frac{n^2}{\sigma_i^{\,2}} + \frac{(y-n)^2}{\sigma_k^{\,2}} \tag{A3.8}$$

$$Q_m = \frac{m^2}{\sigma_i^{\,2}} + \frac{x^2}{\sigma_k^{\,2}} + \frac{m^2}{\sigma_k^{\,2}} - \frac{2mx}{\sigma_k^{\,2}} \tag{A3.9}$$

$$Q_m = \left(\frac{\sigma_i^{\,2}+\sigma_k^{\,2}}{\sigma_i^{\,2}\sigma_k^{\,2}}\right)\frac{m^2}{\sigma_i^{\,2}} - \frac{2mx}{\sigma_k^{\,2}} + \frac{x^2}{\sigma_k^{\,2}} \tag{A3.10}$$

$$Q_m = \left[\left(\frac{\sqrt{\sigma_i^{\,2}+\sigma_k^{\,2}}}{\sigma_i\sigma_k}\right)m - \left(\frac{\sigma_i}{\sigma_k}\right)\frac{x}{\sqrt{\sigma_i^{\,2}+\sigma_k^{\,2}}}\right]^2 + \frac{x^2}{\sigma_k^{\,2}} - \left(\frac{\sigma_i}{\sigma_k}\right)^2\frac{x^2}{\left(\sigma_i^{\,2}+\sigma_k^{\,2}\right)} \tag{A3.11}$$

$$Q_m = \left[\left(\frac{\sqrt{\sigma_i^{\,2}+\sigma_k^{\,2}}}{\sigma_i\sigma_k}\right)m - \nu x\right]^2 + \frac{x^2}{\sigma_k^{\,2}}\left(1 - \frac{\sigma_i^{\,2}}{\left(\sigma_i^{\,2}+\sigma_k^{\,2}\right)}\right) \quad \text{where } \nu = \left(\frac{\sigma_i}{\sigma_k}\right)\frac{1}{\sqrt{\sigma_i^{\,2}+\sigma_k^{\,2}}} \tag{A3.12}$$

$$Q_m = \left[\left(\frac{\sqrt{\sigma_i^{\,2}+\sigma_k^{\,2}}}{\sigma_i\sigma_k}\right)m - \nu x\right]^2 + \frac{x^2}{\sigma_k^{\,2}}\left(\frac{\left(\sigma_i^{\,2}+\sigma_k^{\,2}\right)-\sigma_i^{\,2}}{\left(\sigma_i^{\,2}+\sigma_k^{\,2}\right)}\right) \tag{A3.13}$$

$$Q_m = \left[\left(\frac{\sqrt{\sigma_i^{\,2}+\sigma_k^{\,2}}}{\sigma_i\sigma_k}\right)m - \nu x\right]^2 + \left(\frac{x^2}{\left(\sigma_i^{\,2}+\sigma_k^{\,2}\right)}\right) \tag{A3.14}$$

$$Q_n = \left[\left(\frac{\sqrt{\sigma_i^2+\sigma_k^2}}{\sigma_i\sigma_k}\right)n-\omega x\right]^2+\left(\frac{y^2}{\left(\sigma_i^2+\sigma_k^2\right)}\right) \tag{A3.15}$$

This gives the new image function shown below (as $Q_m = Q_n$).

$$I(x,y)=\frac{A}{2\pi\sigma_k^2}\exp\left[\left(-\frac{1}{2}\right)\left(\frac{1}{\sigma_i^2+\sigma_k^2}\right)\left(x^2+y^2\right)\right]\left\{\int_{-\infty}^{+\infty}\exp\left[\left(-\frac{1}{2}\right)\left(\left(\frac{\sqrt{\sigma_i^2+\sigma_k^2}}{\sigma_i\sigma_k}\right)m-\nu x\right)^2\right]dm\right\}^2 \tag{A3.16}$$

The integrals are now just Gaussian functions shifted in position and as such can be considered with a further change in variables.

$$I(x,y)=\frac{A}{2\pi\sigma_k^2}\exp\left[\left(-\frac{1}{2}\right)\left(\frac{1}{\sigma_i^2+\sigma_k^2}\right)\left(x^2+y^2\right)\right]\left\{\int_{-\infty}^{+\infty}\exp\left[\left(-\frac{1}{2}\right)\left(\frac{\sigma_i^2+\sigma_k^2}{\sigma_i^2\sigma_k^2}\right)m'^2\right]dm'\right\}^2 \tag{A3.17}$$

$$I(x,y)=\frac{A}{2\pi\sigma_k^2}\exp\left[\left(-\frac{1}{2}\right)\left(\frac{1}{\sigma_i^2+\sigma_k^2}\right)\left(x^2+y^2\right)\right]2\pi\left(\frac{\sigma_i^2\sigma_k^2}{\sigma_i^2+\sigma_k^2}\right) \tag{A3.18}$$

This gives the image function $I(x,y)$ in terms of the Gaussian structure peak and spread and the spread of the Gaussian kernel convoluted with the image.

$$I(x,y)=\left(\frac{A\sigma_i^2}{\sigma_i^2+\sigma_k^2}\right)\exp\left[\left(-\frac{1}{2}\right)\left(\frac{1}{\sigma_i^2+\sigma_k^2}\right)\left(x^2+y^2\right)\right] \tag{A3.19}$$

3.12.2 Automatic scale selection

Now that the image function after convolution with the Gaussian kernel is known, it is possible to manipulate the image such that the peak and spread (A and σ respectively) can be found. This method is much better suited to low signal-to-noise situations than a 2D Gaussian least-squares analysis, as shown in section 3.6.2.

The image function for the region of interest post-convolution was shown above as:

$$I(x,y)=A\left(\frac{\sigma_i^2}{\sigma_i^2+\sigma_k^2}\right)\exp\left[\left(-\frac{1}{2}\right)\left(\frac{1}{\sigma_i^2+\sigma_k^2}\right)\left(x^2+y^2\right)\right] \tag{A3.20}$$

Making the substitutions shown below for ease of notation, the function can be simplified.

$$\text{let } A' = A\left(\frac{\sigma_i^2}{\sigma_i^2+\sigma_k^2}\right) \text{ and } \sigma_{eff}^2 = \sigma_i^2+\sigma_k^2 \tag{A3.21}$$

$$I(x,y) = A'\exp\left[-\frac{\left(x^2+y^2\right)}{2\sigma_{eff}^2}\right] \tag{A3.22}$$

Taking the partial differential twice in both x and y then gives $J(x,y)$ as defined below.

$$\frac{\partial I(x,y)}{\partial x} = -\frac{A'x}{\sigma_{eff}^2}\exp\left[-\frac{\left(x^2+y^2\right)}{2\sigma_{eff}^2}\right] \tag{A3.23}$$

$$\frac{\partial^2 I(x,y)}{\partial x^2} = \left(\frac{A'x^2}{\sigma_{eff}^4}-\frac{A'}{\sigma_{eff}^2}\right)\exp\left[-\frac{\left(x^2+y^2\right)}{2\sigma_{eff}^2}\right] \tag{A3.24}$$

$$J(x,y) = \frac{\partial^2 I(x,y)}{\partial x^2}+\frac{\partial^2 I(x,y)}{\partial y^2} \tag{A3.25}$$

$$J(x,y) = \left[\left(\frac{A'x^2}{\sigma_{eff}^4}-\frac{A'}{\sigma_{eff}^2}\right)+\left(\frac{A'y^2}{\sigma_{eff}^4}-\frac{A'}{\sigma_{eff}^2}\right)\right]\exp\left[-\frac{\left(x^2+y^2\right)}{2\sigma_{eff}^2}\right] \tag{A3.26}$$

$$J(x,y) = \frac{A'}{\sigma_{eff}^4}\left(x^2+y^2-2\sigma_{eff}^2\right)\exp\left[-\frac{\left(x^2+y^2\right)}{2\sigma_{eff}^2}\right] \tag{A3.27}$$

The central point of the new image J is given by $(x,y) = (0,0)$, such that J can be further simplified.

$$J(0,0) = -\frac{2A'}{\sigma_{eff}^2} = -2A\frac{\sigma_i^2}{\left(\sigma_i^2+\sigma_k^2\right)^2} \tag{A3.28}$$

The scale factor s_k is defined as σ_k^2 and J can be normalised as shown below.

$$\text{let } s_k = \sigma_k^2 \text{ and } J_s = -s_k \cdot J(0,0) \tag{A3.29}$$

$$J_s = 2A\frac{\sigma_k^2\sigma_i^2}{\left(\sigma_i^2+\sigma_k^2\right)^2} = 2A\frac{s_k\sigma_i^2}{\left(\sigma_i^2+s_k\right)^2} \tag{A3.30}$$

The turning point in J_s is of the most interest.

$$\frac{\partial J_s}{\partial s_k} = 2A\sigma_i^2\left[\left(\sigma_i^2 + s_k\right)^{-2} - 2s_k\left(\sigma_i^2 + s_k\right)^{-3}\right] \tag{A3.31}$$

Taking the differential of J_s with respect to s_k shows that the turning point occurs when $\sigma_i = \sigma_k$.

$$2A\sigma_i^2\left[\left(\sigma_i^2 + s_k\right)^{-2} - 2s_k\left(\sigma_i^2 + s_k\right)^{-3}\right] = 0 \tag{A3.32}$$

$$\left(\sigma_i^2 + s_k\right) = 2s_k \tag{A3.33}$$

$$\sigma_i^2 = s_k \tag{A3.34}$$

$$\sigma_i^2 = \sigma_k^2 \tag{A3.35}$$

Chapter 4: Detector characterisation and optimisation

4.1 Introduction

Following the preliminary investigative work detailed in Chapter 3, it is now necessary to develop the insights gained into the processes involved with the detection process with real experimental data. The feasibility of the camera system has been positively confirmed in the previous chapter, but first one must suitably prepare the detector for the collection of data.

Hard X-ray photon-counting using a scintillator requires very low detector noise levels. These low noise levels require full optimisation of the detector. The CCD97 EM-CCD is more complex than a standard frame-transfer CCD and requires further characterisation than is normally necessary. The work outlined in this chapter describes a selection of the processes in the characterisation and optimisation and details some of the main problems overcome.

The EM-CCD uses multiplication gain in order to reduce the effective readout noise. The multiplication gain process brings with it several complications over those encountered when optimising a standard CCD. The gain itself requires calibration such that one is aware of the gain characteristics of the specific device. Measurement of the gain against the applied voltage requires careful thought due to factors such as the 'excess noise factor' (Section 2.3), the possibility of increased CIC (due to the high voltages applied) and the 'selective inversion' required to run the device with sufficiently low dark current at the temperatures used in this study (Section 4.5).

Before characterising the multiplication gain, one must first optimise the general running of the device itself under unity gain. The majority of the optimisation is not detailed here due to the similarities in the optimisation of any standard CCD, with the specifics to the optimisation of the EM-CCD given more in depth discussion. Various techniques are available to aid the characterisation and optimisation of the device, with the Photon Transfer Curve (PTC) the

method of choice here due to the amount of information that can be accessed from one simple curve. The sequencer used to operate the device requires subtle changes in order to simplify the process of generating the PTC and these may be found detailed in this chapter. As with the need to compare simulation results with experimental data, it is important to ensure that the calibration data are consistent with what one would expect theoretically, and as such, measurements are compared with theoretical calculations and datasheet values where appropriate.

With the specific low-noise requirements of the scintillator-coupled EM-CCD, it is essential that all aspects of the noise in the system are accounted for and any anomalous components resolved. One such anomalous component is described in Section 4.5 with simulations again taken as a method of investigating the hypotheses drawn on the problem.

Finally, before acquiring and analysing the data presented in Chapter 5, the initial testing of the system from an imaging perspective must be discussed. It is essential to verify the origin of the signal in the images obtained from the system and to generate the necessary code to analyse the experimental data and this process is discussed at the end of the chapter.

Through the careful and thorough device optimisation, subsequent characterisation of the system and with the attention to detail in the choice of components and method of assembly of the camera system, the results discussed in Chapter 5 will be seen to bring major improvements to the field of high-resolution hard X-ray imaging.

4.2 Test equipment: the detector

The detector designed for this study is based around a scintillator-coupled EM-CCD where the EM-CCD is a modified back-illuminated CCD97 (Section 2.3). A small fibre-optic plate has been added to the image section of the device, surrounded by a thin layer of aluminium shielding.

A secondary fibre-optic plate, on which the scintillator has been grown, can be coupled to the primary fibre-optic plate. The detector design allows the easy interchange of the scintillator

without the requirement for re-characterisation and repeated optimisation. The fibre-optic plates serve dual purposes. Not only do they allow the easy interchange of the scintillator, but they also act to block any X-rays/gamma-rays that are not absorbed by the scintillator. This prevents direct detection in the EM-CCD and ensures that there will be no signal input in the shielded regions of the device (such as the store section). The arrangement of the fibre-optic plates also importantly provide a very close effective contact between the scintillator and the EM-CCD, minimising the spread of the visible photons at the interfaces.

4.2.1 Camera design

The CCD97-scintillator combination is first inserted into the device headboard, Figure 4.1. The headboard contains a small hole beneath the ceramic of the EM-CCD in order for an aluminium cold-finger to be placed against the surface of the ceramic. The cold-finger is clamped over a thermo-electric cooler, which is in turn clamped against a copper cooling block. The copper block is cooled through two external water pipes connected to a chilled water pump. The thermo-electric cooler is used to reduce the temperature of the device to approximately -30°C, with the water cooling used to remove the excess heat generated by the thermo-electric cooler.

The copper block is then placed into the front wall of the vacuum chamber. The vacuum chamber is required to ensure that no water vapour from the air is deposited on the device under cooling. The copper block is partially isolated from the stainless steel chamber using an O-ring. Vacuum feed-throughs are used to connect the appropriate electronics to the device headboard, thermo-electric cooler and to measure the temperature of the EM-CCD ceramic.

The device is operated using the commercially available XCAM drive system (XCAM website), which allows full customisation of the sequencer, voltages and timings in order to run the EM-CCD. The high-voltage required for the multiplication gain is provided by a separate power supply and the thermo-electric cooler is operated through a temperature controller.

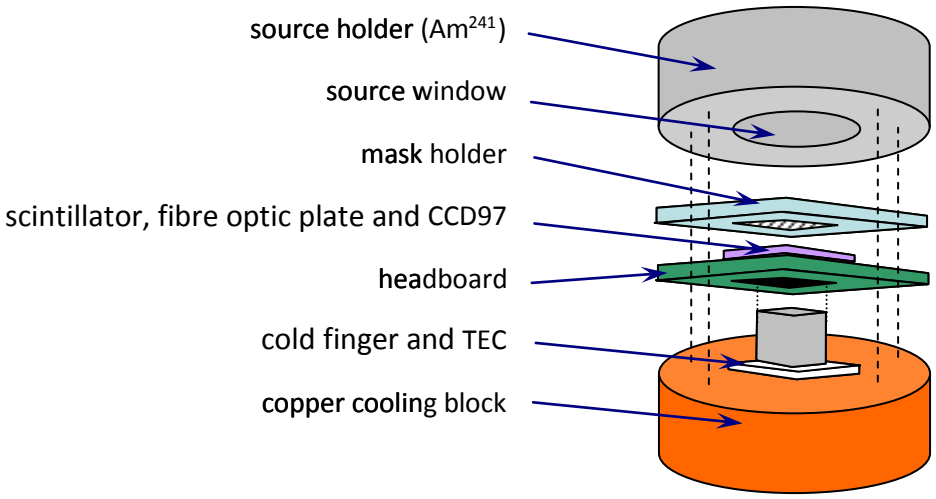

Figure 4.1. An illustration of the experimental set-up using the ^{241}Am source, with the equipment shown measuring approximately 8 cm diameter by 12 cm.

4.2.2 X-ray/gamma-ray sources

X/γ-ray imaging, as described towards the end of this chapter and in Chapter 5, is made possible in the laboratory through the use of two small solid sources (^{241}Am and ^{57}Co), each measuring approximately 1-2 cm in diameter and a few millimetres thick. The initial testing was implemented using the americium source, with further spectral response investigated using the cobalt source. The solid sources are detailed in Table 4.1.

The radiation sources require careful safety consideration. In order to handle the sources they have been placed in two custom designed holders. The holders, manufactured from stainless steel for the americium source and lead for the higher energy cobalt source, allow the safe handling of the sources from the rear side. The open side of the containers (required to allow the gamma-rays to reach the scintillator) are covered by thin aluminium windows such that the gamma-rays can pass through the window but alpha and beta particles are blocked. The entrance windows serve a secondary purpose, required for placement of the sources in the

vacuum chamber. The sources themselves are not classified as vacuum safe. The entrance windows sit over a threaded column which, with the use of an O-ring, allows the small volume containing the source to be isolated from the vacuum, thus rendering the sources (whilst in their holders) safe to be placed in the vacuum chamber.

Source	Emission			Half-life	Activity
	γ (keV)	α (MeV)	β (keV)		
^{241}Am	**59.5 (35.9%)**	5.49 (84.5%)		432.2 years	18.5 kBq
	26.3 (2.4%)	5.44 (13.0%)			
	13.9 (42%)				
^{57}Co	14 (9.54%)		136	271.8 days	85 kBq
	122 (87%)				
	136 (11%)				
	692 (0.02%)				

Table 4.1. The ^{241}Am and ^{57}Co sources used in the experimental section of this study. The activities are quoted for the start of the experimental work. The energies of emission used in this study are shown in black, with emission to be ignored through appropriate thresholds (or due to the minimal contribution to the emission) in grey.

4.2.3 The scintillator

The scintillator used in this study is thallium doped caesium iodide. The scintillator is grown with a columnar structure directly onto a fibre-optic plate and is described in more detail in Section 2.4. In using a scintillator grown on a fibre-optic plate one has a combination which

can be directly coupled to the fibre-optic plate found on the customised CCD97 used in this study.

Using the standard industrial processes at e2v technologies (at the time of acquisition) a 70 µm thick layer of columnar CsI(Tl) was produced. The absorption of X-rays/gamma-rays in the scintillator is limited at higher energies but allows the system to be evaluated. At 60 keV approximately 22 % of the X-rays/gamma-rays are absorbed in the 70 µm layer. This provides the main limit to the detection efficiency of the system. Section 4.6 describes the analysis of the efficiency of the point-search function, providing the secondary limitation on the detection efficiency (the detection efficiency of the detector is discussed further in this section).

4.2.4 Experimental set-up

The source holders can be supported in front of the scintillator such that the source is placed between 0.5 cm and 4.5 cm away from the scintillator surface. A mask can be placed against the scintillator, allowing the imaging of the mask using the scintillator-coupled CCD97 (Figure 4.2).

To achieve the best-case results from the detector, one should use parallel X-rays/gamma-rays arriving perpendicular to the scintillator surface. A parallel beam is not possible with the solid sources due to the limited activity. If the sources were placed at a sufficient distance to get best-case operation then a much larger (and impractical) vacuum chamber would be required and the sources required would pose further safety risks due to the much higher count rates required. The solid sources, placed at 4.5 cm from the detector, allow the practical study of the detector and the scintillation detection process. The distance has no impact on the spectral results.

The system is operated under vacuum (approximately 10^{-3} mbar) and is designed to be operate at temperatures between -20°C and -30°C. Water cooling is supplied at approximately 4°C to

remove excess heat generated by the thermo-electric cooler. The conditions can be assumed constant during the data acquisition unless otherwise stated.

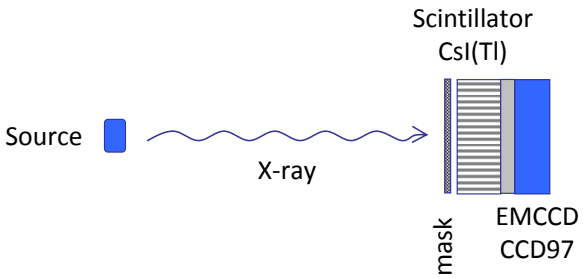

Figure 4.2. A representation of the experimental set-up. The source can be placed at approximately 0.5 cm to 4.5 cm from the mask.

4.3 Sequencer operation for imaging

The CCD is driven by the XCAM commercially available camera drive system (XCAM USBRem2) which allows the user to fully customise the operation of the device. The custom sequencers used to operate the EM-CCD were developed under the work of this thesis.

Two main sequencers were developed to allow both Time Delay Integration (TDI) operation (see below) and traditional imaging. Although the sequencers are based around the same framework (the code required to move the charge through the CCD and to readout the signal is constant), several changes are required to run in the different operating modes. The sequencers are written in assembler code using a series of loop operations to be executed by the Digital Signal Processor (DSP) in the camera drive system.

4.3.1 Imaging modes

When a standard image is required the sequencer must first allow a delay to integrate the signal charge in the image section of the CCD. Once the designated integration time has passed, the charge is passed into the store region of the CCD through the frame-transfer

operation. Whilst the charge packet is in the store section of the CCD no further signal charge can be added except that produced by CIC and dark current. This frame-transfer helps to prevent smearing of the image along the columnar direction whilst the charge is read out from the CCD.

Following the frame-transfer, the store electrodes are clocked to pass the charge row-by-row into the readout register to then be passed through the gain register and digitised.

4.3.2 TDI operation for device calibration

To run the device in TDI mode, as required in the device calibration in the following section, the image and store section of the CCD must be clocked as one continuous region. No external signal will be added to the charge packets once inside the store region of the CCD, with only dark current and CIC increasing the signal. The sequencer operation follows that set out in the flowchart shown in Figure 4.3 using a 'dual slope' CDS method.

The sequencer can, in theory, be run indefinitely. If the device contains *r* rows then the first *r* rows of the image will show a gradient in the dark current and signal profile due to the differing amount of time the charge packets will spend in transit before readout. All rows after the r^{th} row will spend an equal amount of time in each physical row of the device and should hence show a constant profile. This mode of operation is very useful for moving targets, but in this case will be used to define the device calibration value (Section 4.4) through the production of a photon transfer curve.

The TDI mode can also be used to provide 1-dimensional images at a very high effective frame-rate, providing localisation along the row direction.

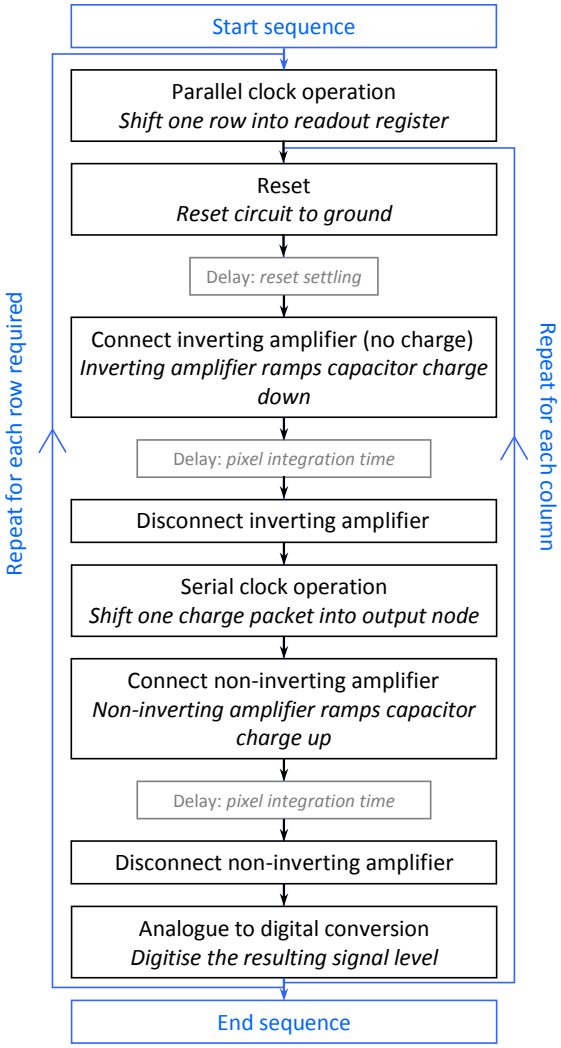

Figure 4.3. The sequencer for CCD image capture follows this flow-chart. The sequencer for TDI operation follows the chart above without the frame integration and frame transfer steps.

4.4 Photon transfer curve (PTC)

The Photon Transfer Curve (PTC) is a remarkably useful tool which can provide various details about the characteristics of a particular device.

4.4.1 The theory behind the PTC

The theory here is split into several sections to explain the various constituent components of the PTC method, starting with a brief discussion of Poisson statistics.

4.4.1.1 Derivation of the PTC from Poisson statistics

If the number of events λ over a large period of time T is known then the probability of an event occurring in a shorter time interval t is given by λ/n where n is given by T/t (Equation 4.1), provided the events are independent of each other and occur with a known average rate.

$$p = \frac{\lambda}{n} \tag{4.1}$$

The probability of no event occurring in any of the n time windows is thus given by Equation 4.2.

$$P(0) = (1-p)^n \tag{4.2}$$

Using the expansion of $e^x = (1 + x)$ in the limit where x tends to zero we arrive at Equation 4.3.

$$P(0) = e^{-np} = e^{-\lambda} \tag{4.3}$$

If there are a total of n small time intervals, where n tends to infinity, the probability of exactly N events occurring is the combination of N of the time periods containing an event and the remaining (n-N) time periods containing no event. The probability of exactly N events occurring over the total time period is thus given by Equation 4.4.

$$P(N) = {}^{n}C_{N}\, p^{N} (1-p)^{n-N} = \frac{n(n-1)(n-2)\ldots(n-N+1)}{N!} \left(\frac{p}{1-p}\right)^{N} (1-p)^{n} \tag{4.4}$$

As n tends to infinity p tends to zero, giving that the second term tends to p^N. The third term, with reference to Equation 4.2 and Equation 4.3 is equal to $e^{-\lambda}$. Finally, the first term simplifies in the limit to $n^N/N!$ giving Equation 4.5.

$$P(N) = \frac{n^N}{N!} p^N e^{-\lambda} = \frac{(np)^N}{N!} e^{-\lambda} = \frac{\lambda^N}{N!} e^{-\lambda} \tag{4.5}$$

The expected value, or mean number of events in the time window will be given by Equation 4.6.

$$E(N) = \sum_{N=0}^{\infty} NP(N) = \sum_{N=0}^{\infty} N\frac{\lambda^N}{N!}e^{-\lambda} = \lambda e^{-\lambda}\sum_{N=1}^{\infty}\frac{\lambda^{N-1}}{(N-1)!} = \lambda e^{-\lambda}\sum_{N=0}^{\infty}\frac{\lambda^N}{N!} = \lambda e^{-\lambda}e^{\lambda} = \lambda \tag{4.6}$$

Using the formula for variance Var(N) = $E(N^2)$ - $E(N)^2$, we can calculate the variance, first calculating $E(N^2)$ as in Equation 4.7.

$$E(N^2) = \sum_{N=0}^{\infty} N^2P(N) = \sum_{N=0}^{\infty} N^2\frac{\lambda^N}{N!}e^{-\lambda} = e^{-\lambda}\sum_{N=1}^{\infty}\left(N(N-1)+N\right)\frac{\lambda^N}{N!} = e^{-\lambda}\sum_{N=2}^{\infty}N(N-1)\frac{\lambda^N}{N!} + E(N)$$

$$= e^{-\lambda}\lambda^2\sum_{N=2}^{\infty}\frac{\lambda^{N-2}}{(N-2)!} + E(N) = \lambda^2e^{-\lambda}\sum_{N=0}^{\infty}\frac{\lambda^N}{(N)!} + E(N) = \lambda^2 + \lambda \tag{4.7}$$

The variance is therefore given by Equation 4.8.

$$Var(N) = \sigma^2 = E\left(N^2\right) - E(N)^2 = \left(\lambda^2 + \lambda\right) - \lambda^2 = \lambda \tag{4.8}$$

So, if a variable can be described using the Poisson distribution, it has a mean value and variance of λ (more details may be found in Francis, 1979).

The probability that exactly *n* emission events from a photon source (or of *n* electrons being thermally generated as dark current in a CCD) occurring inside a given time period *T* is described by the Poisson distribution. As shown above, the variance is equal to the mean of the distribution, and consequently the standard deviation is equal to the square root of the mean. These facts can be used with the photon transfer curve to derive a number of characteristics of the device.

4.4.1.2 Full PTC theory

The input to the camera system is in the form a number of incident photons, *P*. The output signal *S* from the camera system to the user is in the form of ADU channels or Digital Numbers (DN), notated here as S_{DN}. Each pixel has been encoded to give a 12-bit number, and as such, a conversion factor is needed to get back to the original incident signal S_{e^-} in photons or electrons.

The output signal can be given by Equation 4.9, where *P* is the number of incident photons, *QE* is the quantum efficiency of the device (interacting photons per incident photon), η is the quantum yield (electrons generated, collected and transferred per interacting photon), S_V is the sensitivity (V/e$^-$) and G_{OA}, G_{SP} and G_{ADC} are the gains of the output amplifier (V/V), the signal processor (V/V) and the ADC respectively (DN/V) (Janesick, 2001).

$$S_{DN} = P \cdot QE \cdot \eta \cdot S_V \cdot G_{OA} \cdot G_{SP} \cdot G_{ADC} \tag{4.9}$$

Here it is convenient to introduce a calibration in electrons per digital number, *C*, defined by Equation 4.10. It is also convenient for our purposes to redefine incident signal to be in terms of electrons produced in the CCD, as shown in Equation 4.11.

$$\frac{1}{C} = S_V \cdot G_{OA} \cdot G_{SP} \cdot G_{ADC} \tag{4.10}$$

$$S_{e^-} = P \cdot QE \cdot \eta \tag{4.11}$$

This gives a new simplified form as shown in Equation 4.12.

$$S_{DN} = \frac{1}{C} \cdot S_{e^-} \tag{4.12}$$

Using the propagation of errors formula, and including in quadrature the read noise from the output amplifier, σ_{DN}^R, we arrive at Equation 4.13, where the calibration is assumed constant over all signal levels.

$$\sigma_{DN}^{S\,2} = \left(\frac{\partial S_{DN}}{\partial S_{e^-}}\right)^2 \sigma_{e^-}^{S\,2} + \left(\frac{\partial S_{DN}}{\partial C}\right)^2 \sigma_C^{\,2} + \sigma_{DN}^{R\,2} = \left(\frac{1}{C}\right)^2 \sigma_{e^-}^{S\,2} + \sigma_{DN}^{R\,2} \tag{4.13}$$

Using the Poisson statistics described previously, such that the variance on the incident photons is equal to the mean number of incident photons, we arrive at Equation 4.14 and consequently Equation 4.15.

$$\sigma_{DN}^{S\,2} = \frac{1}{C} S_{DN} + \sigma_{DN}^{R\,2} \tag{4.14}$$

$$\log(S_{DN}) - \log(C) = \log\left(\sigma_{DN}^{S\,2} - \sigma_{DN}^{R\,2}\right) = \log\left(\sigma_{DN}^{'\,2}\right) \tag{4.15}$$

This means that by taking a plot of the noise on the signal against the signal in logarithmic space the calibration can be easily determined. When $\sigma_{DN}^{'2}$ is equal to one, the logarithm is equal to zero and $\log(S_{DN}) - \log(C) = 0$. The calibration is therefore given by Equation 4.16.

$$C = S_{DN} \qquad \text{at } \left(\sigma_{DN}^{S\,2} - \sigma_{DN}^{R\,2}\right) = 1 \tag{4.16}$$

This formula will only hold for very low multiplication gain. The excess noise factor has to be taken into account for higher gains (see Section 2.3).

4.4.1.3 Signal vs. Noise Technique to calculate the device calibration

By continuation from Equation 4.15 above, if the log of the signal is plotted against the log of the noise, both in DN, the curve should be defined as follows:

$$\log\sigma_{DN} = \tfrac{1}{2}\log S_{DN} - \tfrac{1}{2}\log C \tag{4.17}$$

The curve can then be used to calculate the calibration factor. This result applies to any region of signal which is dominated by the Poisson distribution.

4.4.1.4 The Photon Transfer Curve (PTC)

In general, the photon transfer curve is formed from three components. The low signal regime is dominated by the read noise. With the read noise dominating, the change in signal, and hence the increase in shot noise, is lost below the read noise. This level gives the read noise of the device.

Once the signal is equivalent to the square of the read noise, the shot noise or fixed pattern noise begins to dominate. As the shot noise is governed by Poisson statistics the noise is equal to the square root of the mean, giving a gradient of one half in logarithmic space. By taking a fit to this region of the curve, the read noise is removed and one can calculate the calibration factor *C*. The fixed pattern noise, or pixel non-uniformity noise, is caused by the differences

between individual pixels and is due to the nature of the device and the device production. The fixed pattern noise is proportional to the signal, acting like a percentage difference in each pixel between the incident signal and the measured value. This gives a gradient of one so can be clearly differentiated from the shot noise.

Once the CCD pixels reach full well capacity the charge can overflow into neighbouring pixels and this creates a smoothing effect of the signal across the pixels. The variance is rapidly reduced at this point which should be clear from the photon transfer curve. The signal at which the curve becomes non-linear gives the full well capacity of the device.

4.4.2 Methods for PTC generation

The fixed pattern noise can be removed through use of a TDI (Time Delay Integration) sequencer, where each pixel has passed through all other pixels in the column and so any fixed pattern noise is removed. By removing the integration time period from the sequencer and running the clocks in the image and store regions simultaneously, signal can be transferred from the top of the CCD to the bottom with the 'time of integration' being equal to the time it takes to transfer the row from the top to the bottom and to readout the previous rows. Each pixel in a column will then have been 'integrating' charge for the same period of time.

By taking an image of area 558 columns by 2056 rows and removing the top 1056 rows (the image and store sections of the CCD at the start of run-time) an image of 558 × 1000 pixels can be obtained (Figure 4.4, left).

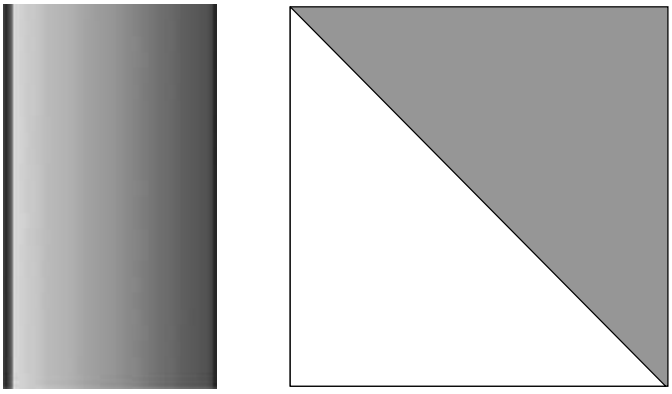

Figure 4.4. An example TDI image (left) using the triangular mask (right). The left column is fully open and the right column is fully shielded, with the columns graded in-between.

This image (Figure 4.4, left) can be used to generate a PTC if the light level is appropriately adjusted. The 11 left-most and right-most pixels show the bias regions of the device. The noise from these regions gives the approximate readout noise. The central 512 columns (from the imaging areas) are graded from left to right with each pixel in a column being subjected to the same amount of signal. The differing light level across the image is achieved through the use of a diagonal mask (Figure 4.4, right). The left-most pixels will be subjected to incident light for the full 512 vertical pixel positions as the charge packet is clocked through the CCD. The right-most pixels will not be subjected to any incident light and the signal here will be purely dark current built up over the transfer period. The central pixels will be subjected to the incident light for half the open 512 pixels and will be shielded for remaining half. This mask produces the banded image seen in Figure 4.4 (left-hand side).

Through the use of some bespoke MATLAB code (written by the author) this image can be manipulated to provide mean signal and noise readings at 512 signal levels (for the 512 columns). The average bias level is subtracted from each of the 1,000 pixel values in each column. The readout noise can be calculated from the overscan region and confirmed from the PTC. The 1,000 pixel values are then averaged in each column to produce the mean signal.

The variance of the 1,000 pixel values provides the total noise, from which the average readout noise can be subtracted in quadrature to leave the signal noise remaining.

4.4.3 PTC results for the CCD97

The results from the above method have shown great potential as a method for determining the calibration of a device when X-ray calibration is not possible or is inconvenient. An example output from this study is shown in Figure 4.5, where the calibration here is shown to be 75 ± 3 e^- per ADU channel.

As can be seen from the example in Figure 4.5, the photon transfer curve is a very powerful tool in characterising a CCD. Under appropriate light conditions a single image can be used to achieve these results.

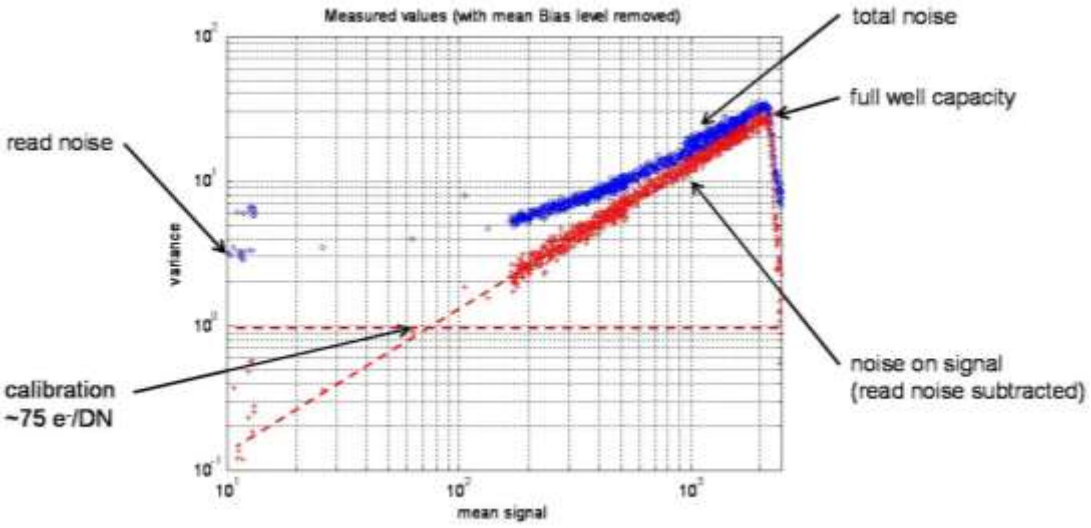

Figure 4.5. A photon transfer curve taken with the CCD97 used in this study under no avalanche gain.

4.4.4 Theoretical calibration value

The calibration for the CCD is a conversion factor used to convert between electrons in the CCD to ADC channels or Digital Numbers (DN) at the output. The calibration will therefore be described in units of electrons per digital number. This conversion factor can be calculated experimentally using the photon transfer curve, as shown in Figure 4.5, but can also be approximated theoretically.

The ADC used in the XCAM drive system is an AD976A, which has an input voltage of ± 10 V, giving a voltage range of 20 V. Although the ADC is 16-bit, this 16-bit data are divided down to give 12-bit values in the XCAM software. As all the data analysis is carried out with data received from the XCAM software, in essence, the ADC can be considered to be 12-bit only. A 12-bit ADC has a range of 2^{12} channels, or 4096 digital numbers. If these are evenly spaced over the input voltage range this gives 4,882.81 µV per ADU channel. From the data-sheet for the EM-CCD used in this study (the CCD97) a typical value of output amplifier responsivity is quoted at 1 µV per electron (CCD97 datasheet, 2004). With an XCAM gain setting of 'A', and a CDS IntT value of 1.2 µs (see the XCAM USBRem2 manual for more information), the initial conditions used for the PTC in Figure 4.5, a gain of 65.6 is achieved between the CCD and the ADC input. Combining these values (Equation 4.18) gives a value of 74 electrons per DN for gain setting A (set in the XCAM software to simply change the resistor connected to the amplifier) with no avalanche gain (Rϕ2HV = 20 V). This value is consistent with that found from the photon transfer method. This calibration value will scale inversely with the multiplication gain.

$$\frac{(4{,}882.81\ \mu\text{V per DN})}{(1\ \mu\text{V per electron}) \times (65.6\ \text{gain})} \approx 74 \text{ electrons per DN} \qquad (4.18)$$

4.4.5 Multiplication gain characterisation

In order to accurately characterise the multiplication gain of the CCD97 used in the experimental work detailed in this study, one must select the method carefully. The method chosen here was designed to provide the most accurate gain measurements.

Through the variation of the R2HV voltage the different gain values are achieved. For each gain, a separate TDI image was taken and a PTC plotted using the method detailed previously following one important modification. As explained earlier in Section 2.3, the impact ionisation process introduces an extra noise factor. Following gain the variance is no longer equivalent to the signal but is instead equivalent to F^2n, where n is the number of electrons in the signal packet. The PTC theory must therefore be adapted to take into account the excess noise. For a standard CCD with no gain the signal in DN is related to the noise in DN by Equation 4.19, where C is the calibration in electrons per DN.

$$\sigma_{DN}^{S\ 2} = \frac{1}{C} S_{DN} + \sigma_{DN}^{R\ 2} \tag{4.19}$$

This must be adapted to include the noise from the gain process to give Equation 4.20.

$$\sigma_{DN}^{S\ 2} = \frac{F^2}{C} S_{DN} + \sigma_{DN}^{R\ 2} \tag{4.20}$$

The gradient of the mean signal against variance plot is no longer equal to the inverse of the calibration but is now given by Equation 4.21.

$$\text{gradient} = \frac{F^2}{C} \tag{4.21}$$

Unfortunately, as one is trying to ascertain the gain at a particular value of R2HV, it is not possible to calculate the value of F in advance and include this in the measurements. The method employed here is detailed below.

The 'false gain', G_f, was defined as the gain that would be implied if the gain process did not introduce further noise. The false gain is related to the 'real gain', G, by Equation 4.22, which, following the substitution $F^2 = (2G\text{-}1)/G$ (Section 2.3, Equation 2.14) gives Equation 4.23.

$$G_f = F^2 G \tag{4.22}$$

$$G_f = \left(\frac{2G-1}{G}\right) G = 2G - 1 \tag{4.23}$$

The false gain profile can thus be converted to the real gain profile using Equation 4.24 at each of the measured points. This approximated conversion gives an error of below 1% up to a gain of × 200.

$$G = \frac{G_f + 1}{2} \tag{4.24}$$

The plot shown in Figure 4.6 gives the gain profile for the CCD97 used in this study using the method described above, including the correction for the noise on the gain. For a set voltage the gain achieved increases as the temperature decreases. The reasons for this increase are described in Section 2.3 and follow the expected trend.

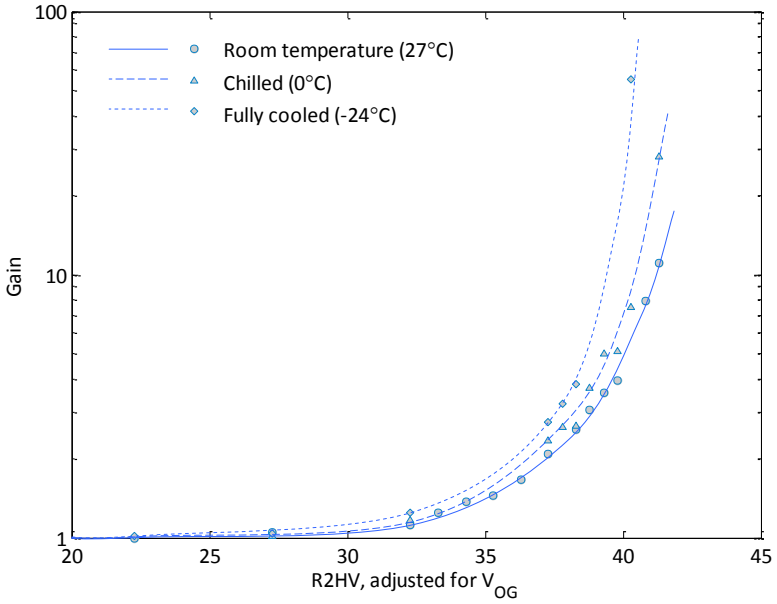

Figure 4.6. Avalanche gain variation with increasing R2HV. This plot was created through the generation of a photon transfer curve for each R2HV and applying the conversion shown in Equation 4.24 to account for the noise on the gain process. The inverse of the calibration, when related to the calibration at unity gain, gives the multiplication gain at the specified voltage.

4.5 Device characterisation and optimisation

In the running of a CCD many voltages and timings can be altered to give improved performance. In some cases there is a delicate balance, for example between the readout speed and the readout noise, whereas in other cases there is a seemingly 'better' option, such as with the substrate voltage. The optimisation of the device is briefly described here, with the effects some elements have on the overall performance discussed. Problems encountered when running the entire CCD in inverted mode are described below. The solution to these problems is also supplied with comparison to the initial results and discussion of the causes of these difficulties illustrated through simulative methods.

The PTC is a very useful tool throughout the characterisation and optimisation process. When certain voltages and timings are changed (such as the pixel integration time when using dual slope integration CDS) the calibration of the device will change. Without knowing the equivalence between electrons signal and the ADU values the readout noise cannot be compared. It is essential to optimise the device with regards to electrons and not ADU values, as the real device operates in electrons and ADU should only be viewed as the digitised signal.

4.5.1 Selective inversion and CIC

Clock Induced Charge (CIC) is found to be at a low level throughout the image and store regions of the EM-CCD. The spurious signal, generated as the electrons pass through the device and encounter variations in the electric field, cannot be removed but may be reduced.

The level of CIC generated is dependent on the transfer frequency (inversely proportional) and clock rise times (inversely proportional). It has been reported that the CIC is independent of temperature in the range -15°C to -50°C (Low-Light technical note, 2004). The greatest impact on CIC is supplied by the applied bias levels. When run in inverted mode, the level of CIC generated is one to two orders of magnitude higher.

The CCD97 image and store sections generate 1×10^{-4} electrons/pixel/transfer in inverted mode at normal clock amplitudes (Low-Light technical note, 2004). In non-inverted mode, the CIC is reduced to 3×10^{-6} electrons/pixel/transfer, but the increase in dark current of over two orders of magnitude outweighs the benefits of running the image and store regions in non-inverted mode.

Initial testing with the device at unity gain showed benefits in running the full device in inverted mode. Running the readout and gain register in inverted mode allowed clocking to be uniform across the device from a low clock voltage of 0 V. The substrate voltage can be set to 8-9 V for full inversion and the dark current is significantly reduced.

Problems were encountered, however, when running the device with higher gain. With the full device running in inverted mode, although the dark current is reduced, the CIC was found to increase dramatically at higher R2HV voltages. The charge induced as electrons pass through regions of high electric field increases as the electric field increases (cf. avalanche gain, Section 2.3). In inverted mode, holes are forced to flood the surface where the clock voltage is low. As the R2HV electrodes are placed high, the holes escape the region below the electrode and inversely, electrons pass in the opposing direction. As the electrons pass through the high electric field on either edge of the R2HV electrodes, charge is generated (CIC). In non-inverted mode the effects are minimised through the lack of inversion/non-inversion switching as the voltages change on the electrodes.

4.5.1.1 Readout noise tail

The indicator that high amounts of CIC are present may be found in the 'readout noise peak'. The overscan can generally be histogrammed to provide what is usually referred to as the 'readout noise peak'. The overscan is chosen as no signal can be integrated in the overscan elements externally and only internal charge generation can be accounted for. The dark current is assumed negligible if the full device is cooled and/or run in inverted mode. If signal is present (such as if the device is operated at room temperature, generating large amounts of dark current) then the readout may appear normal under unity gain, with the dark current and CIC lost in the readout noise itself. When gain is applied, however, the readout noise peak is seen to shift to higher signals, broadening in the process, due to signal now present in the overscan elements being multiplied.

If small amounts of signal are present, such as a signal of a few electrons, then the readout noise profile will show significant differences to the higher signal case described above. At unity gain the signal will still be lost in the readout noise peak but at higher levels of gain, however, a tail becomes apparent on the right of the readout noise peak (Figure 4.7).

Taking readings from the 'overscan' and 'underscan' elements, one would expect significant differences in the histogram if dark current was causing the tail to the readout noise. All elements pass through the gain register, but the overscan and underscan will pass through differing numbers of pixels in the readout register (hence taking different amounts of time) and one would expect a difference of 50% in the dark current. This was not seen in the results using the cooled device under full inversion and indicates that the tail is not due to dark current generated in the readout and gain registers.

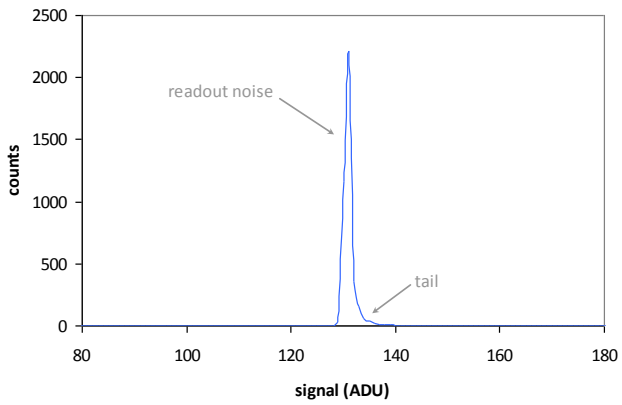

Figure 4.7. Experimental results showing the 'readout noise' peak.

The tail was found to increase as the gain was increased.

The level of CIC increases as the gain voltage is increased due to the increased voltages/electric-fields applied. As any charge present due to CIC is passed through the gain register, it is itself multiplied. This leads to a complex chain of charge generation, where an electron generated at the start of the register is multiplied by the full length of the gain register, and conversely, an electron generated near the output node is only multiplied by a small amount. The multiplication of small signals leads to broad signal profiles due to the statistical nature of the charge multiplication process. The combination of the Poisson profiles (Figure 4.8) for the single figure charge levels leads to the tail generated on the 'readout noise

peak' (Figure 4.9). The same statistics apply to optical photons detected in the CCD, which, combined with the multiplication gain, leaves difficulties in determining the true number of photons/electrons at lower levels.

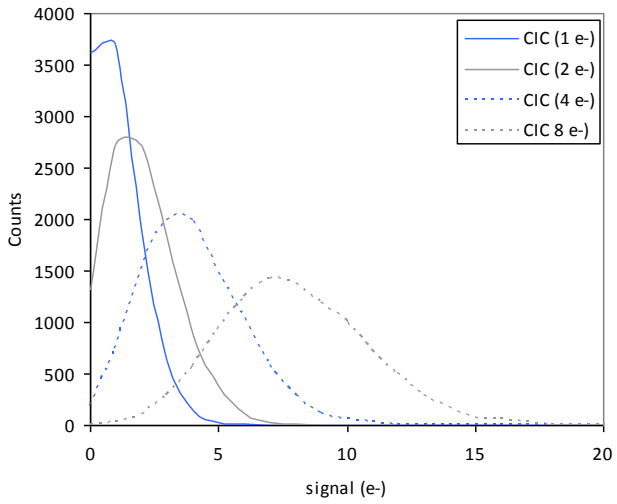

Figure 4.8. The Poisson statistics for signals of $1e^-$, $2e^-$, $4e^-$ and $8e^-$ following multiplication gain, normalised to the signal levels for the *x*-axis.

The diagrammatic representation shown in Figure 4.9 must be considered with the low signal profiles generated across all points in the 'readout noise' peak, and not only from the average zero signal level as shown for simplicity. At low gain, the CIC is very low. As the gain is increased, not only does the CIC increase, but the gain under which the signal is multiplied is also increased.

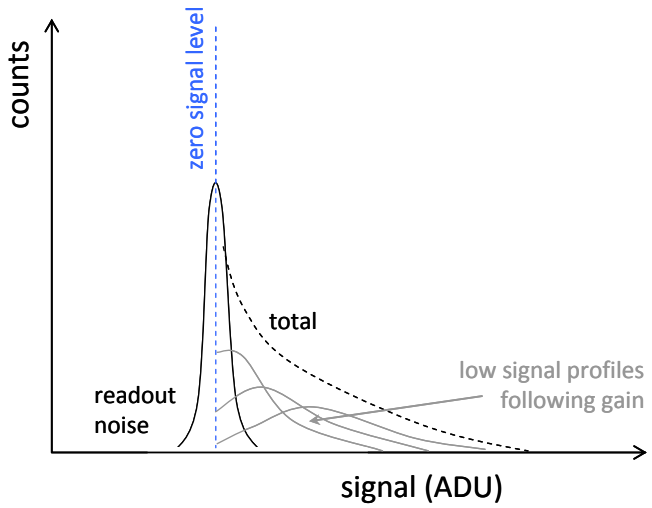

Figure 4.9. A representation of small signal profiles following gain (grey) and the effect these have on the readout noise peak (black). The one-sided nature of the tail occurs due to the impossibility of negative signal.

4.5.1.2 Simulating the 'readout noise' peak

The EM-CCD simulation detailed earlier in sections 3.7 and 3.8 also allows the analysis of results produced in the experimental data. The hypothesis that the CIC was causing problems in the gain register was verified through the use of the EM-CCD Monte Carlo simulation. By changing the inversion properties of the readout and gain register in the simulation, the effects on the overscan region (taken as an indication of the readout noise containing zero time signal) were investigated. Results from the simulation are shown in Figure 4.10, in comparison to the experimental data. The tail on the peak can only be seen at low temperatures - at room temperature the signal from the overscan elements accumulates too much dark current and the whole peak is shifted to a higher ADU value. It is only at lower temperatures that the CIC begins to be visible over the dark current and acts to increase the high-signal side of the overscan peak.

Following on from the simulations detailed above on the readout of the EM-CCD, one can simulate the influence of increased CIC on the 'readout noise' peak. The simulation results shown in Figure 4.10 account for an average generation of 2 electrons per pixel of CIC. A value of 0.5 ADU was taken for the readout noise to match that measured in early experimental results.

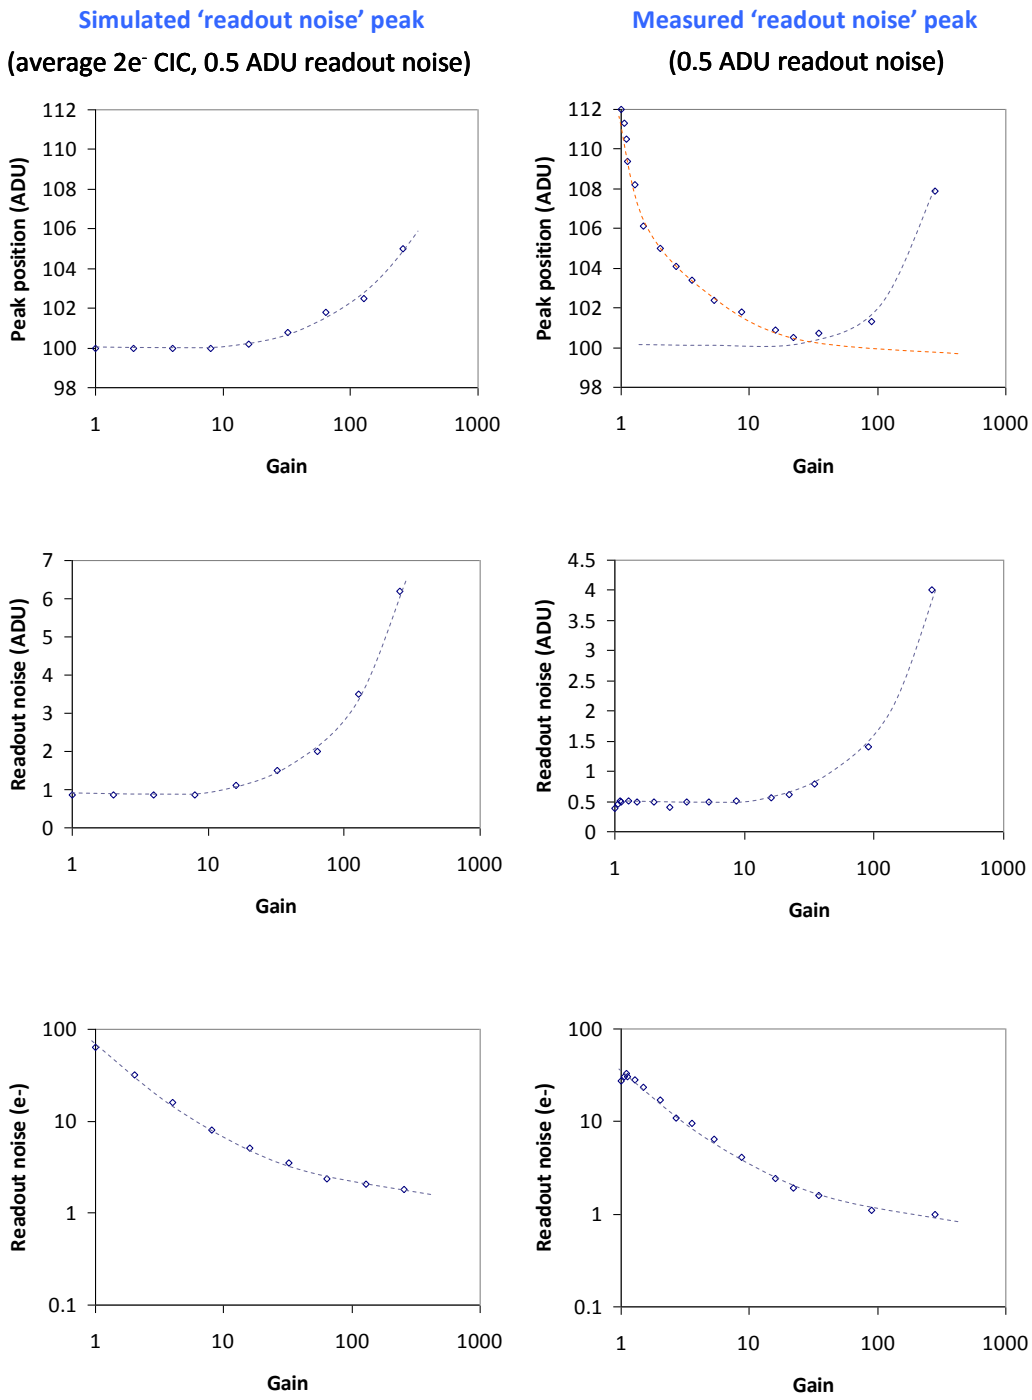

Figure 4.10. Results from the analysis of simulated 'readout noise' peaks giving apparent peak localisation and readout noise in ADU rms and electrons rms. The variation in the 'zero level' in the upper right plot , indicated by the red dashed line, is assumed to be a consequence of the drive electronics although this is not fully understood and does not affect the results.

The variation in CIC with R2HV was not simulated here and hence the readout noise values are not appropriately scaled. By setting the simulated readout noise to the level seen with the experimental 'apparent readout noise' it is seen that the simulated results predict an 'apparent readout noise' higher than that found experimentally. This indicates that it is, indeed, an external factor (i.e. CIC) increasing the noise level in the experimental data. The measurements and simulation results do, however, show trends consistent with the measured data, allowing comparisons to be made.

The generation of CIC in the gain register and the multiplication of this signal at higher gain leads to a tail in the readout noise which must be eliminated in order to see the low signal levels required in photon-counting with the scintillator-coupled EM-CCD.

4.5.2 Selective inversion between the image/store regions and readout registers

To run the device with high gain, and hence high R2HV voltages of 40-50 V, whilst reducing the dark current and CIC as far as possible the device must be run with selective inversion: the image and store sections operated in inverted mode, whilst the avalanche gain register is operated in non-inverted mode. The readout register and gain register are linked through the clocking mechanism and are therefore very difficult to separate.

In order to run the device under selective inversion whilst maintaining a constant substrate voltage across the device, the low clock voltages must be altered. If a substrate voltage of 3 V is used, a clock running from a low value of 0 V will not lead to inversion. If, however, the low clock voltage is run from -5 V then the effective substrate voltage will lie at 8 V, resulting in inverted mode operation. Through this process, the image and store regions can be run in inverted mode with low clock voltages of -5 V whilst the readout registers can be run with low clock voltages of 0 V in non-inverted mode.

With the XCAM drive system used in this study (XCAM USBRem2), it was not possible to run the device with low clock voltages other than 0 V. To enable the use of the system with the

CCD97 it was therefore necessary to insert an adapter between the drive system and the CCD headboard.

4.5.2.1 Modifying the clocking signal from the XCAM drive system

The output from the XCAM drive system can only run from a low clock voltage of 0 V (at the time of writing). To alter this, a clock can be output from 0-5 V with the timings as required by the CCD. This signal is fed into a high speed logic gate (HCPL 2630 datasheet, 2001), Figure 4.11. The ground of the logic gate can be set to -5 V using the converted -15 V output from the drive system, creating a clock waveform from -5 V to 0 V. This waveform is fed into a high speed metal-oxide-semiconductor field-effect transistor (MOSFET) driver (EL7202 datasheet, 1996). This is set with a ground of -5 V and a V+ value of +5 V, taken from the converted +15 V output of the drive system. The clock signals are then regenerated from -5 V to +5 V, with timings generated by the drive system as specified in the sequencer. Providing these signals to the image and store clocks and running with a substrate voltage of 3 V leaves the image and store regions inverted, whilst maintaining no inversion in the readout and gain registers.

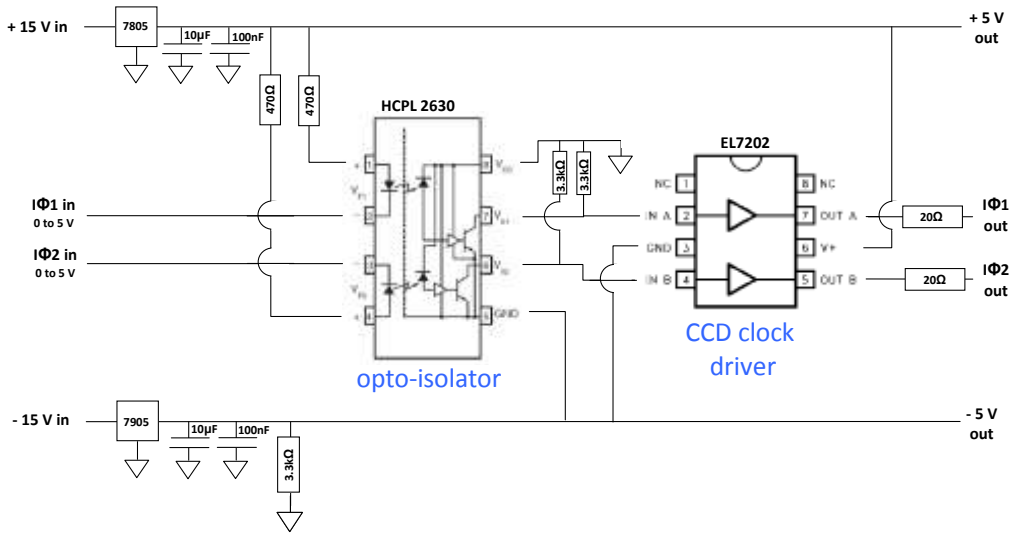

Figure 4.11. The modification to the drive system to generate clock voltages from -5 V to +5 V, allowing selective inversion of the CCD97.

4.5.2.2 Before and after the application of selective inversion

Figure 4.12 shows the performance of the device with the gain register inverted (before the clocking modification) and with the gain register not inverted (after the clocking modification). The dramatic increase in the clock-induced charge in the gain register when the gain register was operated in inverted mode is clear in all four plots.

4.5.3 Substrate bias (V_{ss})

There is a contribution to the total dark current in a CCD from three main regions. These regions consist of the depleted region, the neutral bulk silicon and the Si-SiO_2 interface states on the front and rear sides of the CCD (Janesick, 2001). The main contribution arises from the surface states and it is the dark current from this region which can be reduced through the use of a higher substrate bias voltage, V_{SS}.

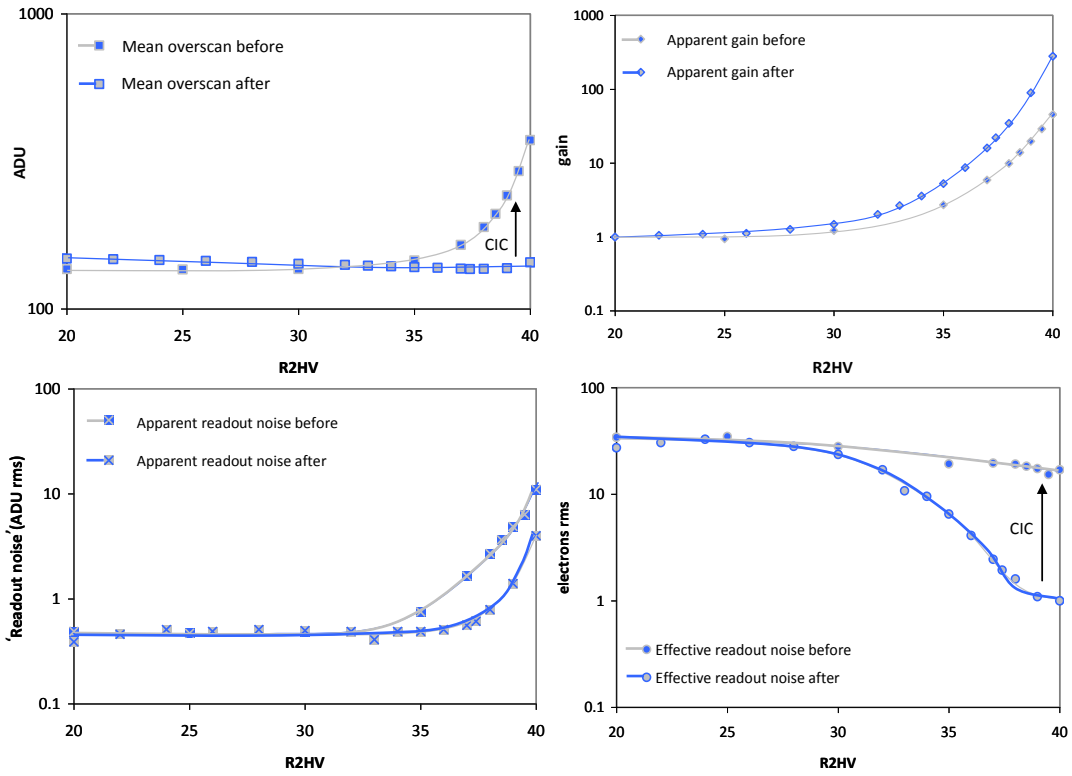

Figure 4.12 The overscan value increases with applied R2HV when the gain register is inverted due to increases in CIC with increasing R2HV (top left). Running the gain register in non-inverted mode leaves the overscan approximately constant. The 'apparent gain' (signal minus overscan) appears lower under inversion due to the increase in the signal in the overscan (top right). The 'readout noise' (in ADU rms) was calculated from the width of the histogram of overscan signal values and increases when inverted due to excessive CIC from the high R2HV voltages (and hence high fields) in the gain register (bottom left). The effective readout noise (the readout noise divided by the gain) does not decrease as expected when the gain register is inverted due to the CIC (bottom right). When not inverted, the effective readout noise decreases to the sub-electron level.

4.5.3.1 Inverted mode using the default clocking signals

When V_{SS} is raised to approximately + 9 V, holes will accumulate at the silicon surface, changing the surface from n-type to p-type, giving rise to the name 'inverted mode operation' (IMO). The potential at the surface of the silicon is then no longer dependent on the gate voltage (due to the presence of the holes) and the surface is thus referred to as being pinned. The holes which have accumulated at the surface will recombine with any thermally induced electrons (dark current), dramatically reducing the total dark current level. Additional implants can be added to allow the whole surface to be inverted. If V_{SS} is increased too far the full-well capacity and depth of depletion can be adversely affected and the larger signal packets can begin to recombine with the excess holes.

Each device may differ in the optimal V_{SS} value and so individual calibration is required. It can be seen in Figure 4.13 that the dark current is rapidly reduced with a V_{SS} of approximately 7 V and above before modification. There is a rapid decrease in the dark current until a V_{SS} of approximately 8 V is reached. Signal charge transfer is affected above approximately 8.6 V, so a value of 8.4 V has been taken as the optimal value – the highest V_{SS} possible with full signal charge transfer.

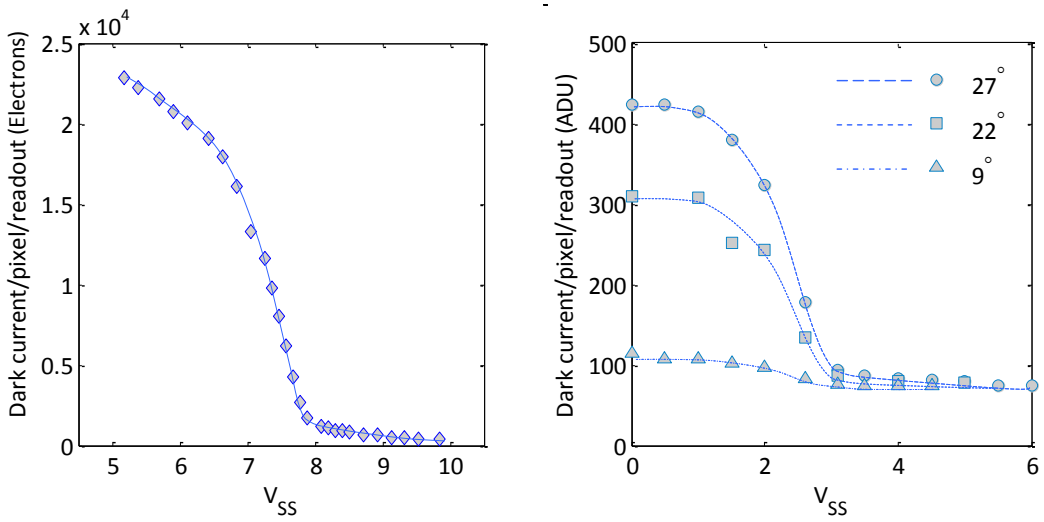

Figure 4.13. The effect on the dark current through the increase of the substrate voltage, before (left) and after (right) the clock modification. The plot on the right has a calibration factor of 75 electrons per ADU, with temperatures given in °C. The plot on the left hand side was taken at approximately 25°C. A 5 V shift can be seen in the inversion position.

4.5.3.2 Inverted mode using the modified clocking signals

Under selective inversion the image and store clocks run from -5 V to +5 V, offsetting the 'zero level' from 0 V down to -5 V. The readout and gain registers remain at a 'zero level' of 0 V as there is no change to their clocking scheme from the full inversion case. The applied substrate voltage V_{SS} must now be compared to the low clock voltage; an applied substrate voltage of 0 V is now equivalent to 5 V in the image and store regions and 0 V in the readout and gain registers. The 5 V offset to the substrate between the two regions makes it possible to invert the image and store regions whilst keeping the readout and gain registers in the non-inverted mode. This acts to dramatically reduce the dark current in the image and store regions whilst not causing problems with excessive clock-induced charge in the gain register with the presence of the high electric field.

The effects of the clock level shift can be seen in Figure 4.13. A shift by 5 V in the point of inversion can be seen when comparisons are made left to right. The shift seen in Figure 4.13 shows that the dark current can be reduced in the image section whilst the readout and gain registers remain non-inverted. Although it would be ideal to also invert the readout register, the coupling of the clock signals between the readout and gain registers means that either both or neither must be inverted. The increase in dark-current caused by leaving the readout and gain registers in normal operation mode is outweighed by the clock-induced charge if the gain register is inverted.

4.6 Initial X-ray imaging

Before embarking on the characterisation of the gamma-ray event profiles it is important to test the system at a lower level. The initial testing described below shows that the events to be processed are due to X-ray interactions in the scintillator and not due to random fluctuations in noise, bright pixels or direct detection of gamma-rays in the CCD.

The initial image processing is described here, showing the basic background removal procedure, chosen following the consideration of several methods, and the primary cleaning of the image to remove bright pixels, whether due to the device itself or the random fluctuations in dark current and noise. The results described here detail the testing of the system only and are not designed as a test of the quality of the imaging system – full results from the device testing are described in Chapter 5.

4.6.1 Image acquisition

Images are initially acquired with the XCAM system and software. The system provides an image of the ADU values which can be converted to a matrix in MATLAB using purpose-written code. From this matrix the image can be processed. Once full calibration of the timings and voltages has been achieved, the images can be captured directly in MATLAB using a code system and GUI written in this study.

4.6.2 Initial image processing

There are two stages to the initial image processing: removing the background and 'cleaning' the image. Several background removal processes were tested. The first tests were based around the removal of a 'local background', averaging the area surrounding the pixel in question and removing this as the 'background level'. A background image with no source could not be taken without opening the system to remove the source (at room temperature), and hence changing the operating conditions. The local background removal was found to produce a background image with standard deviation of only 0.6 ADU. Considering the uncertainty in the signal levels due to the shot noise on the signal and the dark current the fluctuations in the background image were considered negligible, particularly considering the processing time required to remove the local background compared to a simple image average removal. For this reason, the average signal was calculated over the image area and this value was subtracted from the signal level in each pixel.

The image was then cleaned to remove the majority of the bright pixels in the image through a dual-pass system. The initial bright-pixel removal stage removes the brightest single pixels in the image, whilst the second pass aims to remove the larger fluctuations in the background noise in the image. For each removal, the surrounding 8 pixels were averaged and the signal in the bright pixel replaced with this average signal value.

4.6.3 Basic imaging in photon counting mode

Taking one thousand images at the same multiplication gain and temperature with the 5.55 kBq ^{241}Am X-ray source gives a basic sample of X-ray interactions in the scintillator. A lead sheet edge was placed between the scintillator and X-ray source at an angle such that half of the CCD97 image region was covered.

Some 'scintillation flashes' can be seen in Figure 4.14, but one must first confirm that these are scintillation flashes from X-ray interactions in the scintillator and not bright pixels or direct X-ray interactions in the CCD.

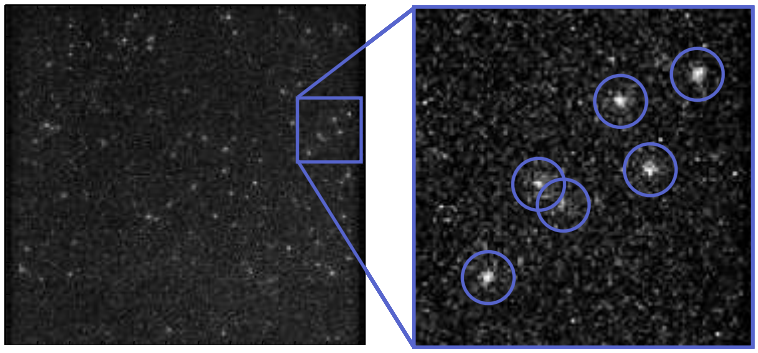

Figure 4.14. A single frame (approximately 1cm^2) showing approximately 200 events with a region containing six individual Gaussian-like profiles enlarged, each caused by a single X-ray photon interacting in the scintillator.

4.6.4 Basic integrated image

Here, the 'integrated image' referred to is not produced using the standard integration method. The rate of interactions in the scintillator has purposefully been set up to allow individual flashes to be distinguishable in each frame and so integrating in the traditional sense will only give a high noise background from the accumulated dark current. One thousand images were taken, each with an integration time of 2 seconds (integration occurs whilst the previous frame is read out at ~150 kHz). To simply add these images would lead to the same noise problems as found by taking a standard integrated image.

Here, a 'signal reject level' was set above the image offset (the base level of the image). All pixels above this reject level are added to the "integrated" image. Again, MATLAB code using matrix multiplication is used for increased efficiency. In this way, signal from scintillation flashes will dominate but bright pixels which have not been removed as single pixel events will also be present. The results of summing one thousand images in this way are shown in Figure 4.15.

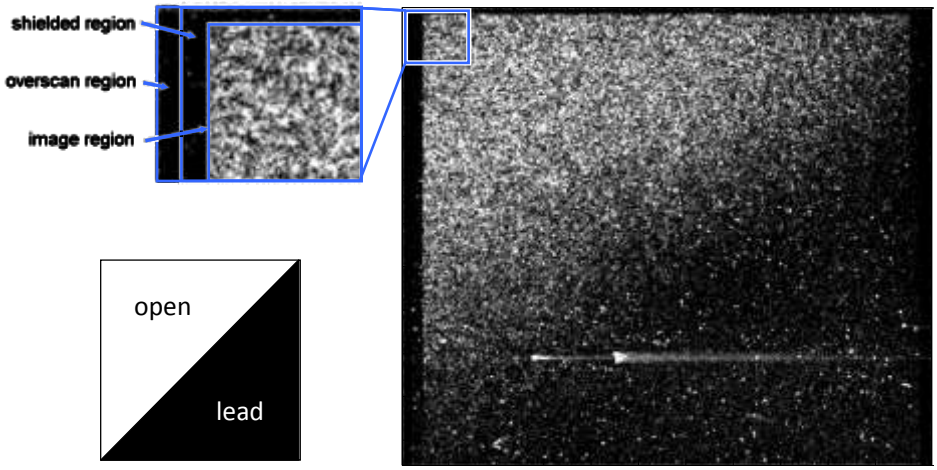

Figure 4.15. The 'integrated' image of one thousand individual frames showing the position of the lead sheet (bottom right corner), including the overscan and shielded regions. The events can be seen to be inside only the image region (not in the shielded regions). This implies that the signal is from photons passing through the face plate of the CCD and not X-rays interacting in the silicon of the CCD itself.

4.6.4.1 Verification of signal

Figure 4.15 shows the diagonal lead sheet placed across the image area of the CCD. The probability of the signal being from bright pixels and not a source external to the CCD is

negligible. This implies that the signal is not caused by bright pixels and is supplied by an external source.

The lack of definition of the edge is due to the system geometry. The source was placed against the lead edge in this case to increase the count rate and hence there is a large angular variation in the incident X-rays. This does not represent the same geometrical set-up as used in the experimental part of the study (Chapter 5).

The smearing seen in row 400 (Figure 4.15) is due to a phenomenon which is an unwanted consequence of the gain register structure. As the charge passes under the high voltage electrode it is located nearer to the electrodes leading to a higher probability of contact with the surface states. If the surface states are filled by the signal electrons then they will later be released over the subsequent time. The released electrons can be easily seen due to the gain on the signal. This effect is less obvious when the electrons reach the surface states towards the end of the gain register as the subsequently released electrons undergo lower gain. The threshold signal for this phenomenon to occur is approximately 200k electrons and the visibility of the smearing is dependent upon the position in the gain register at which this signal level is reached (D. Burt, private communication).

4.6.4.2 Gamma-ray interactions in the scintillator

The shielded region of the CCD97 will not be subjected to photons around 550 nm, the peak wavelength of light emitted from CsI(Tl). Gamma-rays may, however, pass through the shielding and interact directly with the CCD. Here, the question is whether these 'flashes' seen in the image are due to X-rays interacting in the CCD97 or in the scintillator. It is worth remembering that the system has been designed to minimise the number of X-rays interacting in the CCD through the use of a fibre optic plate (of several millimetres thickness) between the scintillator and CCD surface in order to absorb the X-rays.

In Figure 4.15 it is clear that the events recorded are purely in the image section of the CCD and do not pass into the shielded regions, verifying that the signal recorded is due to gamma-ray interactions in the scintillator.

4.7 Optimisation of the point-search function

It is important to optimise the point-search function to achieve the most desirable results possible for the required application. For the device testing described in this study the detector is being considered for several possible applications (namely medical imaging and synchrotron beam line imaging) and as such the specific conditions are not precisely defined. A high accuracy in detection is given with a high efficiency of detection (Figure 4.16), but this must be balanced with the number of 'real counts'. A high detection efficiency with a low count rate would not be desirable for medical applications (due to the requirement of higher doses) or for beam line use (time/cost).

With reference to Figure 4.17, it can be seen that lower sigma values and lower signal limits give a higher number of 'real counts'. This is unfortunately countered by a low detection efficiency (as shown in Figure 4.16). Whilst increasing the smoothing kernel width dramatically decreases the 'real-count-rate', increasing the signal limit has a lesser effect.

It can be seen in Figure 4.16 that a smoothing kernel sigma below 2 rapidly decreases the detection efficiency. Below a sigma of 2 pixels, the background noise, particularly due to any brighter single pixels, is not sufficiently suppressed, leading to a larger number of erroneous events being detected.

With the initial device set-up an efficiency of over 80% can be achieved whilst maintaining high count-rates (sigma of 2 pixels and a signal limit of 3-4). The count-rate can subsequently be increased, but only to the detriment of the detection efficiency. These results relate only to the initial set-up of the device without optimising the cooling mechanism. With further cooling the dark current is reduced, leading to a reduction in the noise and hence the signal limit can

be decreased. Optimisation of the cooling system is only possible when assembling the full system, as the thermal contact between the TEC and EM-CCD is affected with each system adjustment.

With the final set-up, taking the absorption of 60 keV X-rays/gamma-rays of 22 % in the 70 µm thick CsI(Tl) layer, a total detection efficiency of incident photons may given at approximately 16 %. This efficiency may be improved in future work (see Section 5.13) where a thicker scintillating layer may be used to increase the absorption probability of the incident X/γ-ray photons. For example, layers of thickness 120 µm, 200 µm and 500 µm will give absorption at 60 keV of 30 %, 51 % and 83 % respectively, increasing the total detection efficiency of the system by up to a factor of four.

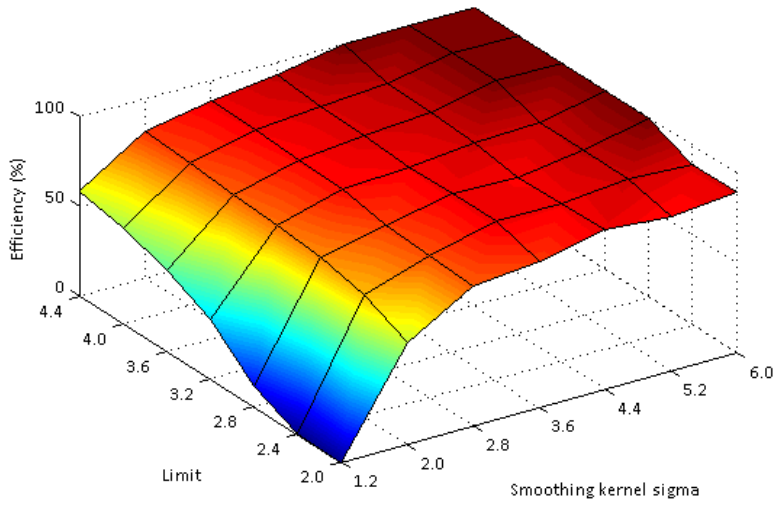

Figure 4.16. The efficiency of the detection (the number of 'real counts' as a proportion of the total counts) for varying limits and smoothing kernels. The axes for the limit and sigma are inverted with reference to Figure 4.17. The efficiency of detection increases as the limit and sigma increase, but the number of 'real counts' drops dramatically (Figure 4.17). Both the efficiency and number of real counts must be considered.

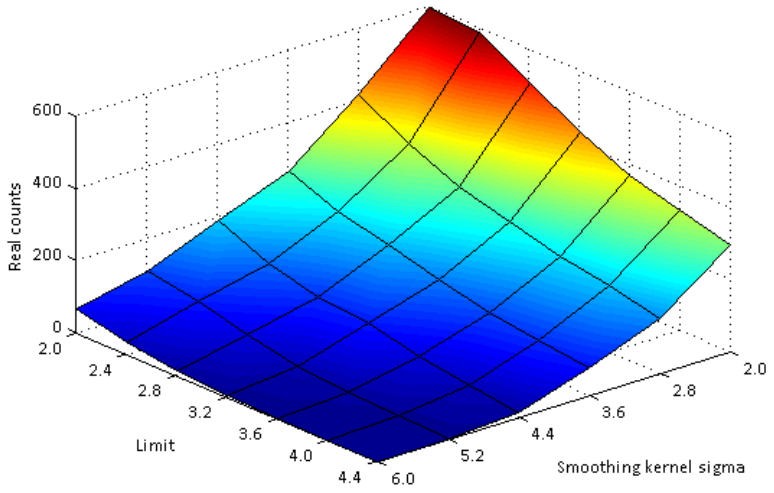

Figure 4.17. The number of 'real counts' detected with varying kernels and limits above the background. At low sigma the smoothing kernel does not depress the noise sufficiently and leads to discontinuities – a selection of conditions must be met for consideration as an 'event'. As the signal limit and sigma increase the number of detections decreases as the number of events falling below threshold increases.

4.8 Conclusions

Following the optimisation and characterisation of the device it has been shown that the device, setup as described, is capable of resolving the individual X/γ-ray interactions in the scintillator. The various problems encountered with reducing the noise to a sufficiently low level have been overcome through a variety of methods, including the use of simulations to gain a better understanding of the processes involved. The initial images shown at the end of this chapter show that the signal measured in the device is indeed due to X/γ-ray interactions in the scintillator. The characterisation and optimisation of the device, with the understanding gained from the techniques used, will be developed in Chapter 5 to show the true capabilities of the detector and the advancements that this brings to the field of high-resolution, hard X-ray imaging.

Chapter 5: Photon-counting X-ray/gamma-ray camera

5.1 Introduction

The simulations carried out in Chapter 4 show potential problems that may occur when using a photon-counting X/γ-ray camera in practice. These simulations must be verified through experimental methods. The calibration and optimisation of the device, some of which was detailed in the previous chapter, is essential in order to view the individual photons with the photon-counting methodology. The quality of the optimisation has major implications on the resulting use of the detector. Without sufficient noise reduction and sensitivity the event profiles will not be extractable from the images. Without photon-counting, an integrating mode must be used, having a detrimental effect on the resolution through the inclusion of the visible photon spread from the photon passage through the scintillator. Without the use of photon-counting there is also no way of gaining energy discrimination data, removing many of the benefits of this type of detector. This chapter concentrates on the experimental testing and the results achieved with the device. The spatial measurements are first presented, with the justification of the methods used left until after the presentation of the spectral results for better continuity. The results of a simulated energy spectrum may then be considered with respect to the experimental results and an interaction process flow, and the effect this has on the detector quality, can be discussed.

Through a combination of simulation techniques and a thorough analysis of experimental data one is able to gain a much deeper understanding of the scintillator-coupled hard X-ray imager and the challenges encountered. The development of any detector is a continuous process and new discoveries and better understanding always leads to new problems to overcome in order to design the 'perfect' device. This chapter concludes with a brief look into the future of the scintillator-coupled EM-CCD, discussing where advances in detector design and drive electronics may enable even greater strides to be made.

5.2 Spatial measurement techniques

The spatial resolution of a detector is a complicated value to define. There is large disagreement within the community as to what provides a suitable resolution measurement for the comparison of different detector systems. When testing devices with the same apparatus/technique, the measurements will be valid for comparison there only, and will not necessarily provide a fair comparison with measurements taken using a different technique. As an example, considering the Modulation Transfer Function (MTF) as described below, a value of 4 % is often quoted for optical resolution measurements, whereas some prefer a 10 % or 50 % value for comparison. Other tests may simply provide the MTF value at a specified spatial frequency. The FWHM of various peaks may be quoted, although the peaks themselves may have been created using inaccurate fitting methods to the experimental data (see Section 5.5.2).

Throughout this work comparable parameters are taken as far as possible to provide a comparison to results quoted in the literature (Chapter 2). The method of using the whole MTF as an indication of the imaging abilities of a detector is generally considered the most complete, as this provides an indication of the imaging capability across all spatial frequencies. This, however, cannot exclude results quoted in the literature in which the fitting of curves to data has provided misleading values. As shown later in this chapter, a simple Gaussian fit will not accurately map the Line Spread Function (LSF) and provides a large over-estimation in the MTF over all spatial frequencies. Ignoring the 'noise-floor' can also lead to large overestimations of the MTF, as the 'noise-floor' generally considered present with optical measurements may hide subtle variances only present with the X/γ-ray detectors detailed here.

For this reason, the quoted spatial measurements in the literature should be read with some caution. It is true, however, that an under-estimation of the MTF is very rarely quoted, and as

such, values quoted in the literature should be viewed as a 'best-case' limit to the spatial resolution as far as the method is capable.

The MTF is described below along with a description of various MTF measurement techniques. The method selected here was chosen to enable measurement with the solid gamma-ray sources available for bench-top testing. In integrating mode, the signal levels are generally high across the device and as such the statistical fluctuations in the signal level are much lower. With the photon-counting method, only a single count is provided per event profile, leading to larger statistical fluctuations in signal level across the image region for the same time-period. Modifications to the standard method (as designed in this study) are shown later in this section and the results show the benefits of this method for photon-counting analysis.

5.2.1 The Modulation Transfer Function (MTF)

The Modulation Transfer Function (MTF), or 'spatial frequency response', is a method used to characterise the components of a visual system. In Figure 5.1, the column on the left shows a sample mask which could be imaged by a series of lenses or a CCD based imaging system. The amplitude of these masks varies from black to white in distinct steps. As one moves down the column, the stripes get thinner: the spatial frequency increases. The image which is obtained through the imaging system will not be as sharp as the input. The right side of the figure shows some example outputs from a generic imaging system. The edges of the stripes are blurred, and although the blurring has little impact in the ability to distinguish the independent stripes when spread far apart, as the stripes get closer together (as the spatial frequency increases) the distinction between the stripes is diminished.

The final column of the figure shows how the amplitude of the image varies across each row. Rather than using human judgement to determine the highest spatial frequency that can be resolved, a numerical value can be assigned. Using human judgement creates ambiguity and each person may have a different opinion of when a set of stripes is resolvable. This numerical

value can be plotted against the related spatial frequency, giving the modulation transfer function, Figure 5.2 (relating to the example shown in Figure 5.1). The spatial frequency is generally defined in line pairs per millimetre (lp/mm), with the MTF normalised to one at the lower limit of spatial frequency.

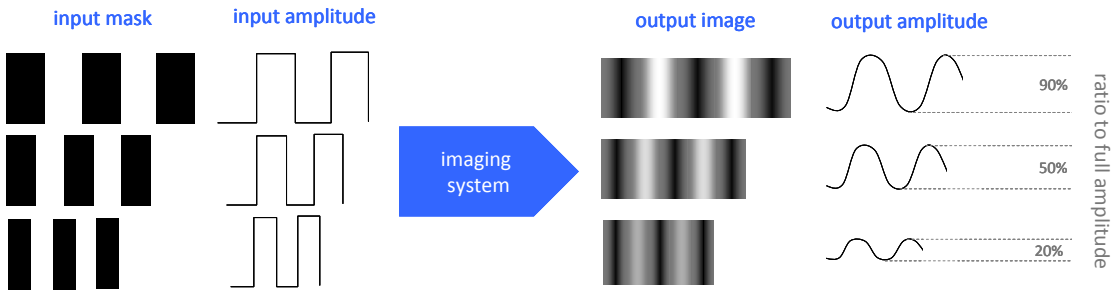

Figure 5.1. The incident amplitude is altered by the imaging system. The difference in the output amplitude gives the MTF of the device.

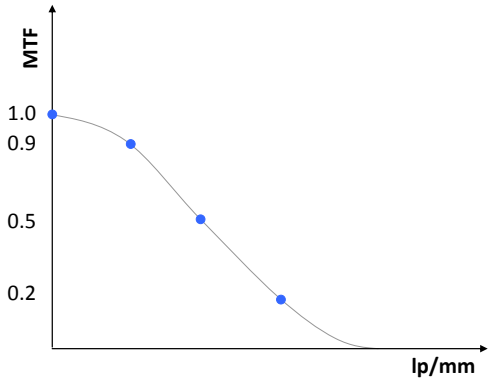

Figure 5.2. An illustration of how the MTF curve would take shape for the simple example shown in Figure 5.1.

Relating the MTF to a single value for the spatial resolution of the detector is a difficult process, and many different values may be obtained. Whilst the 'image sharpness' may be

considered to be at an MTF value of 50%, the standard resolution measurement for optical based systems falls somewhere between 2% and 10%, depending on the source. For this reason, and to avoid ambiguity, this chapter ensures the method used to generate any signal resolution values is fully specified or that, in the most complete case, the MTF is displayed in full.

5.2.2 MTF measurement techniques

Many methods are available to measure the MTF of a detector. In optical measurements using purely visible light one can use a simple patterned mask. The patterned mask allows the imaging of a collection of lines which gradually get closer and closer together (an example is shown in Figure 5.3). As the lines get closer together the spatial frequency is increasing. The contrast variation can be measured using the test pattern, and the resulting function gives the MTF. This method is not practical at 60 keV. The thickness of material required to absorb sufficient gamma-rays will not allow the imaging of high spatial frequencies due to the limits on the hole size (the minimum hole size is of a similar order to the thickness).

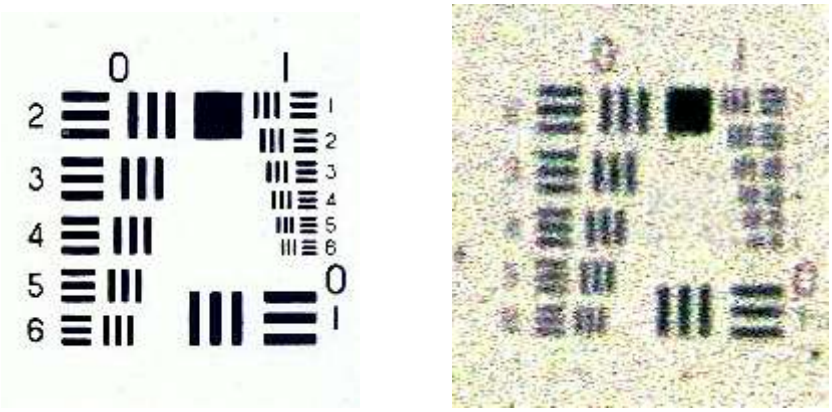

Figure 5.3. The test pattern mask, left, can be used with low energy photons to measure the MTF (USAF 1951 Test Pattern). An example output on film, right (Bob Atkins, 2008).

With high-energy imaging using sources at 60 keV or 120 keV as desired here, the standard mask techniques cannot be used and other methods must be investigated. Two methods can

be used, both of which make use of an edge in order to generate an Edge Spread Function (ESF).

The virtual knife-edge technique does not require a mask, but instead uses a very narrow beam of X-rays/gamma-rays of known spread. The beam is imaged by the detector in progressive steps across the device (approximately perpendicular to the columns of the CCD). Signal is summed, or events are counted in photon-counting mode, on one side of a pre-defined column edge (the 'virtual knife edge'). The edge of the selected column acts as the virtual knife edge. As the beam passes over the column edge, fewer signal photons will be detected as more and more fall on the dark side of the selected column edge, Figure 5.4. The sum of the signal on the selected region of the detector is plotted against the position of the beam, providing a function analogous to the ESF defined below. The ESF can be differentiated to give the LSF, from which the Fourier transform can provide the MTF. If the beam width is not minimal, the width must be removed from the LSF.

The virtual knife-edge technique requires a narrow beam of X-rays/gamma-rays of known spread. This method can give excellent results, but the beam requirements lead to severe impracticalities for lab-based testing.

A second method for obtaining the ESF uses a simple absorbing material to create an edge. In this case a 0.5 mm thick tungsten sheet was placed over the detector absorbing photons incident on one half of the detector whilst allowing the imaging of photons over the uncovered area. The edge is placed at a shallow angle to the columnar direction in order to create an oversampled image and avoid problems related to the finite pixel size.

The image of the edge must be re-aligned to the vertical through rotation or shearing to allow the summation of the signal in the columnar direction. The sum of the signal in the columnar direction (following the alignment of the edge to the vertical) gives the ESF. In conditions where the statistical fluctuations in the signal are large (due to the low numbers of counts),

the method must be adapted. A method designed in this work to overcoming the noise in the ESF is detailed below.

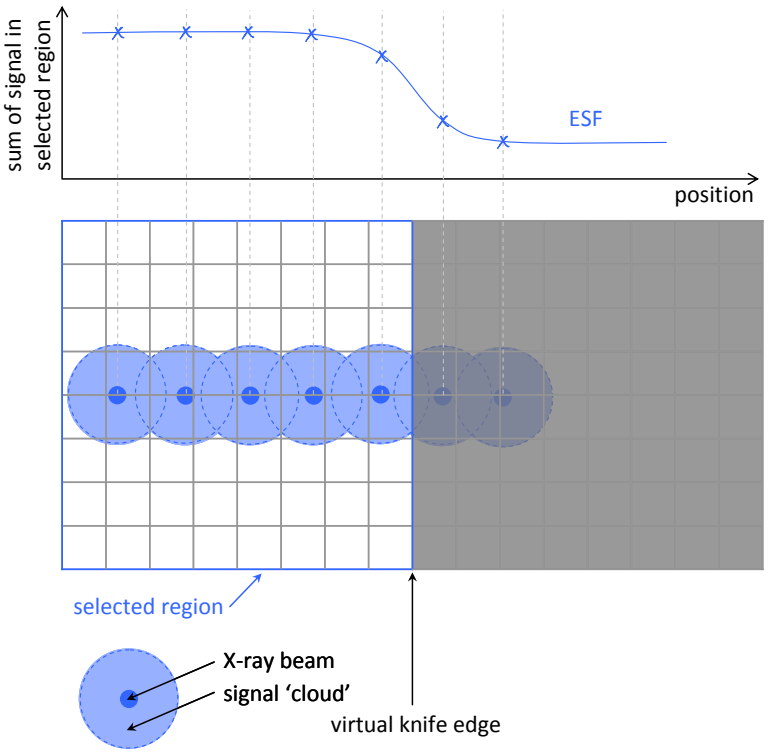

Figure 5.4. A representation of the virtual knife edge technique which can be used to determine the device ESF.

5.2.3 Oversampled MTF method for low count-rate images

The standard method used to generate the ESF from an image of an edge is shown in Figure 5.5. The centre of each pixel is passed parallel to the edge to the base of the image. The ESF is generated by taking the intensity of each pixel against the new position. This method samples each pixel once and with low count-rate images the results are very noisy, Figure 5.6 (blue).

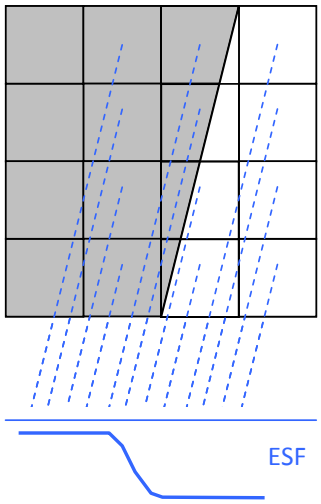

Figure 5.5. The standard method for sampling the image of an edge to generate the ESF.

The method developed in this study allows the generation of an ESF with much lower noise characteristics. Each pixel is split into a number of sub-pixels with the signal averaged over the sub-pixels, the number being determined by the angle ϑ of the edge against the pixel edges such that the number of sub-pixels per pixel is given by rounding $(\tan\vartheta)^{-1}$ to the nearest integer. The image can then be sheared in the new sub-pixel increments such that the edge is shifted to the vertical position. Binning the image into a single row provides an oversampled ESF such that the signal fluctuations found in the standard ESF are greatly reduced, Figure 5.6 (red). The errors introduced with this method are minimal. Taking the example of 2.8° rotation (as seen in the results presented in Section 5.4), the pixels are split into 20 increments (1/tan(2.8°)=20.45). This gives an error in the shift of 0.45 of a 'sub-pixel' over 20 rows, equating to an error of ± 0.09 µm per 20 rows (for those used to produce the ESF). Over the full image, this equates to a blurring of the true edge position ± 4.5 µm. Minimising the number of rows used to produce the ESF has a linear effect on the error (where the reduction in the number of rows used is either through choice or through the need to remove sections of the image due to bright regions, regions of high noise or regions of lower counts due to the inverse square law incidence of the gamma-rays from the solid sources). For the results

presented later in this chapter, a sub-section of the image has been used, reducing the errors in the subsequent calculations. It is worth noting that the errors in the edge location are strongly dependent on the angle of edge. For example, an edge at 2.9° presents half the error found for an edge at 2.8°.

Referring to Figure 5.6, the new method of ESF generation can be seen to give much smaller fluctuations in the ESF using the low-count data. The characteristics of the ESF are preserved with greatly increased visibility of the edge fall-off. Using this method, the time taken to collect a suitable data-set is greatly reduced due to the requirement of fewer frames to achieve the required signal to noise ratio in the ESF.

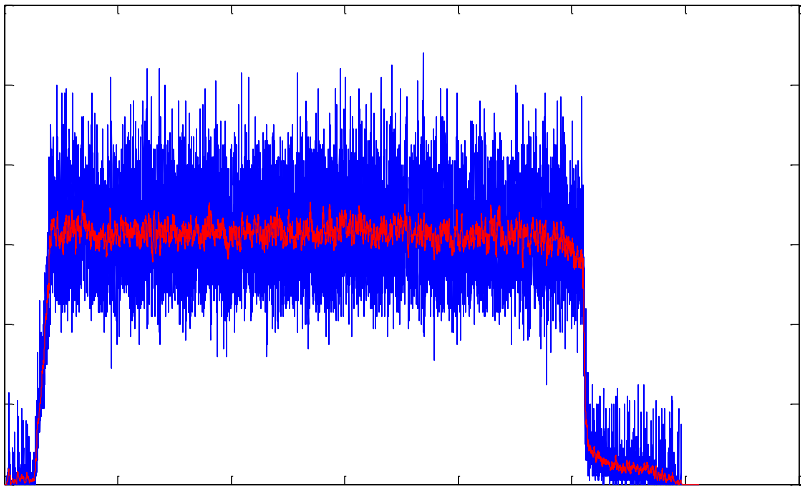

Figure 5.6. The two edge spread functions shown in this plot are generated from the same data-set produced in this study. The blue data is generated using the standard method, whereas the red data is generated using the improved method designed here.

5.2.3.1 The Line Spread Function (LSF)

The ESF must be converted to the Line Spread Function (LSF) in order to generate the MTF. The LSF is taken as the differential of the ESF to produce a single peak corresponding to the

fall-off from the edge response. The i^{th} component of the LSF function is generated using Equation 5.1 (the 'finite difference') as the function cannot be generated analytically.

$$\mathrm{LSF}(i) = \frac{\mathrm{ESF}(i+1) - \mathrm{ESF}(i-1)}{2\left[\mathrm{x}(i+1) - \mathrm{x}(i-1)\right]} \tag{5.1}$$

The LSF gives a Gaussian-like peak which can be analysed in its own right before using the function to generate the MTF. In studies detailed in the literature, the tails of the LSF peak are often ignored. The high noise at the low signal edges of the peak are often disregarded with the central peak of the LSF extrapolated below what is termed the 'noise floor', Figure 5.7. Although this may appear to give improved results, and the validity in purely optical systems is verified, with a scintillator based system the method is invalid. The validity of the extrapolation is discussed later in this chapter with relation to the experimental results.

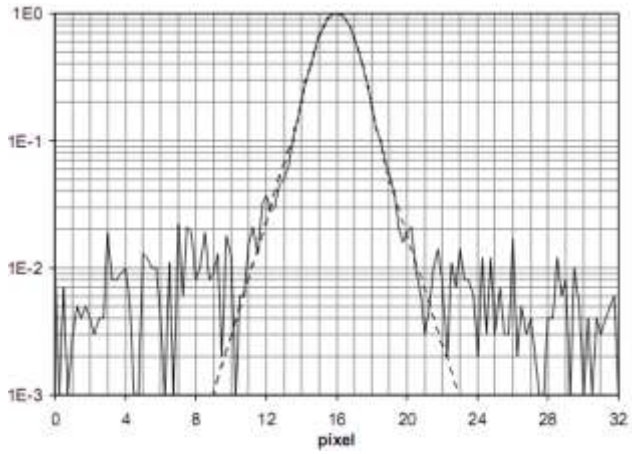

Figure 5.7. The LSF can be extrapolated below the 'noise floor' in optical systems. In this optical plot (reproduced from Graeve *et al.*, 2001) the dashed line shows the extrapolated tails of the LSF. In the scintillator based system this method is flawed, over-predicting the MTF.

5.2.3.2 MTF

The MTF is generated from the Fourier transform of the LSF. The Fast Fourier Transform (FFT) can be used when developing code to generate the MTF. The MTF should be normalised to

unity at the low spatial frequency limit (with a solid black or white input there are no visible blurring effects). This normalisation is often performed following the generation of the MTF but this, again, can cause an overestimation of the device MTF. In order to accurately normalise the MTF the LSF can be normalised to generate an MTF with no further corrections required.

Taking the FFT of noisy data will generate a noisy MTF plot and the results are often inconclusive, however, taking a fit to the ESF or LSF before the FFT removes the unwanted fluctuations in the MTF. Fitting to the ESF is a complicated process and requires several assumptions. The fitting to the ESF in examples in the literature with averaging functions may remove the finer structure of the edge profile and this generally leads to an overestimation of the MTF. Fitting to the LSF provides more accurate results and the fitted curve can be used to generate a fitted ESF (through integration of the curve analytically). Through the fitting to the LSF of the experimental data, the finer details of the imaging process can be examined. This is discussed with reference to the experimental results in Section 5.5.

5.3 Experimental layout

The experimental setup was designed to allow testing on a bench-top system. With the low count rate gamma-ray sources available for integration into bench-top apparatus the flux is limited. Due to the low flux, the source must be placed in close proximity to the detector in order to achieve suitable count-rates of tens to hundreds of events per frame.

An americium-241 source was used to investigate the imaging capabilities of the detector. The emission of the source is dominated by the 59.5 keV gamma-ray emission and this is used to provide the events in the scintillator. The source active area measures approximately 5 mm in diameter. For the initial results (Section 5.4), the source was placed against the tungsten edge. This provides the highest count-rate possible but does not provide the optimal conditions for testing the imaging capabilities of the device. The directional variation in the incident

gamma-rays severely limits any measurements taken to determine the spatial resolution of the device. The high-count rate does, however, provide data suitable for initial verification of the system and a large number of event profiles suitable for energy analysis – the distance of the source from the scintillator has no impact in the formation of the event profiles for spectral analysis.

Moving the source further from the tungsten edge improves the trajectories of the incident gamma-rays. In-order to maintain a reasonable count-rate, an edge-to-source distance of 3 cm was implemented. Again, this is not optimised for the best possible MTF measurements, but the results are indicative of the imaging capabilities of the detector.

In order to achieve a more optimal set-up to determine the true MTF of the device, a columnar beam of gamma-rays is required. This beam could be provided by a synchrotron, such as that at the Diamond Light Source (based on the Harwell Science and Innovation Campus), but this was not possible in the testing period designated for this project. Future work includes the testing of the device at a synchrotron facility and improvements in device design, generated through this study, are detailed at the end of this chapter.

A second source has been used to ensure the spectral calibration of the device is as declared with the ^{241}Am source. The cobalt-57 source was placed in close proximity to the scintillator. Due to a lower count rate, the source was not deemed suitable for imaging, but provides an energy spectrum of great interest. The higher energy of the cobalt source (approximately double that of the ^{241}Am source) provides further information on the implications of energy discrimination on the imaging capabilities of the detector.

5.4 Initial results

The results detailed here were taken during the initial testing of the system. The source was placed a few millimetres from the scintillator to maximise the count-rate. This was necessary to allow multiple data-sets to be taken whilst making small adjustments to the system without

consuming large amounts of operating time. The initial images have been used to calibrate the system and to allow the testing of the bespoke imaging and analysis code designed during this study.

Event profiles are selected from multiple frames using the methods described in Chapter 4. The code analyses the event profiles, here centroiding the events to the nearest pixel to optimise the operation time (further centroiding is not necessary due to the geometry-limitations of the system). Each event location is saved to an array of sub-pixel locations. These locations can be histogrammed in two dimensions (following event profiling) such that the image is reconstructed. The image can then be analysed using the MTF method described previously. Figure 5.8 shows this process pictorially, displaying the reconstructed image (with the edge found to be placed at an angle of 2.8° in this case), the sheared image (such that the edge falls vertically) and the signal binned to a single row displayed as the ESF.

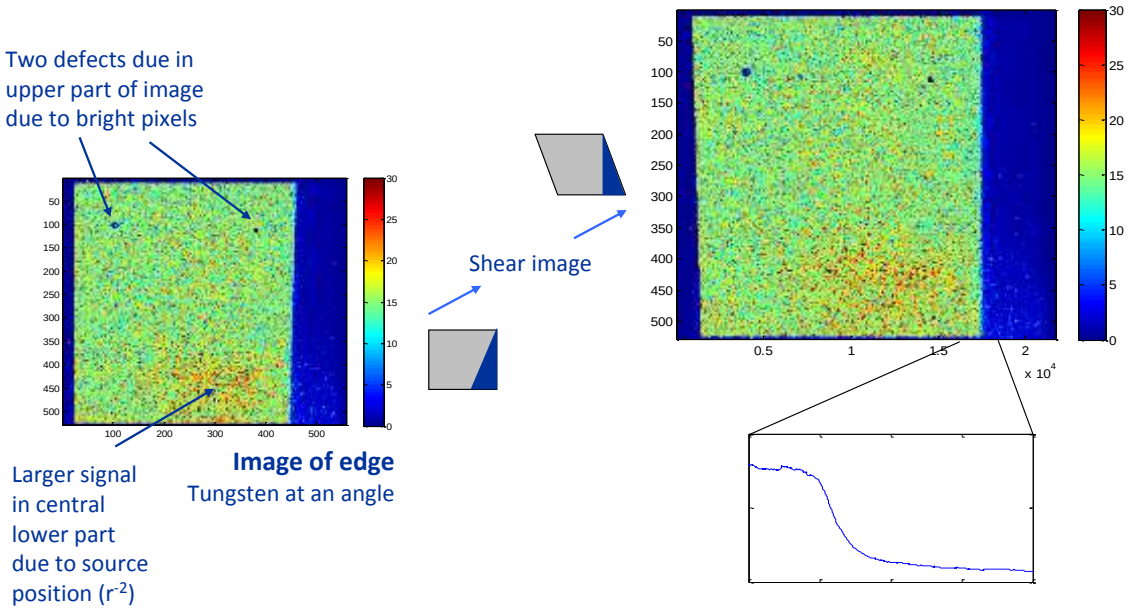

Figure 5.8. A pictorial representation of the generation of the ESF (bottom right) for the initial data from a sample image (left). The ^{241}Am is located several millimetres from the scintillator. These data were used to test the analysis code and confirm correct operation of the system

The ESF in figure 5.9 shows the transition from 'black-to-white' from left-to-right. The average count-rate on the vertical axis shows that some 'events' may be found in the non-irradiated region. These events can be attributed to both X-ray events (limited transmission through the tungsten edge) and fluorescence events, Section 5.7. The ESF falls from 90% to 10% over approximately 100 µm. This in itself shows an improvement over many detectors found in the literature, but with the severely geometry-limited (Figure 5.10) system used for this initial testing, the results can be greatly improved.

The system testing detailed above is substantially geometry-limited. The geometry of the system allows the gamma-rays from the ^{241}Am source to pass at many angles into the scintillator, with many passing under the edge due to the layout of the system. Figure 5.10 shows the limitations of the initial testing layout, which, as mentioned previously, was designed to optimise the count-rate in order to test the code used to analyse the data.

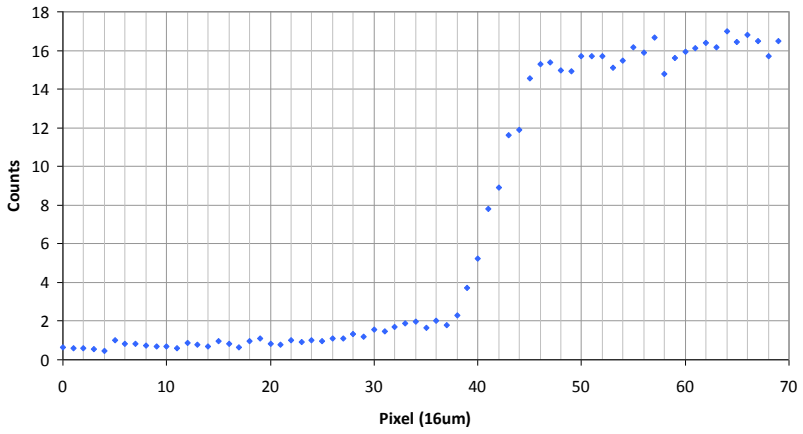

Figure 5.9. The ESF for the initial data (enlarged from figure 5.8). The resolution demonstrated here is very poor due to the divergence of the gamma-rays from the ^{241}Am.

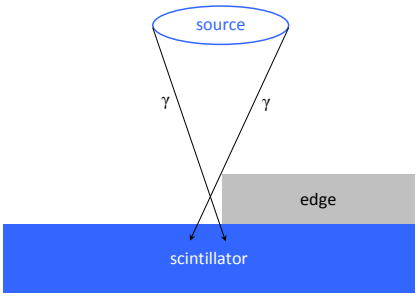

Figure 5.10. The geometry of the system limits the maximum resolution possible.

5.5 Imaging capabilities

The device has been tested in both integrating and photon-counting modes. The integrated image is taken through the accumulation of a large number of events in a single frame. This imaging technique includes the spread of each event in the final image. The photon-counted image was generated using the centroiding technique described in Section 3.10 using 'quarter-pixels' (taking the positions to the nearest quarter pixel). The system resolution is limited by the geometry, with the source placed approximately 3 cm from the mask. The centroid positions are stored in an array and used to generate an image through two-dimensional histogramming.

5.5.1 The edge profile

Images of the edge are shown in Figure 5.11 for the integrating and photon-counting techniques. The improvement in the edge sharpness can be seen, but through the processing of the image as described above one can quantify the improvements. The imaged edge in the photon-counting case is shown enlarged.

The images are sheared (as described previously, using the method for low count-rate images) and binned to a single row to give the ESF. This is shown in Figure 5.12 for the photon-counted image. The ESF shows sufficient noise such that, if processed directly, the MTF will contain

undesirable noise levels. As discussed earlier, fitting to the ESF is a complicated process, and as such, the LSF is calculated through discrete differentiation and a profile is fitted to the LSF.

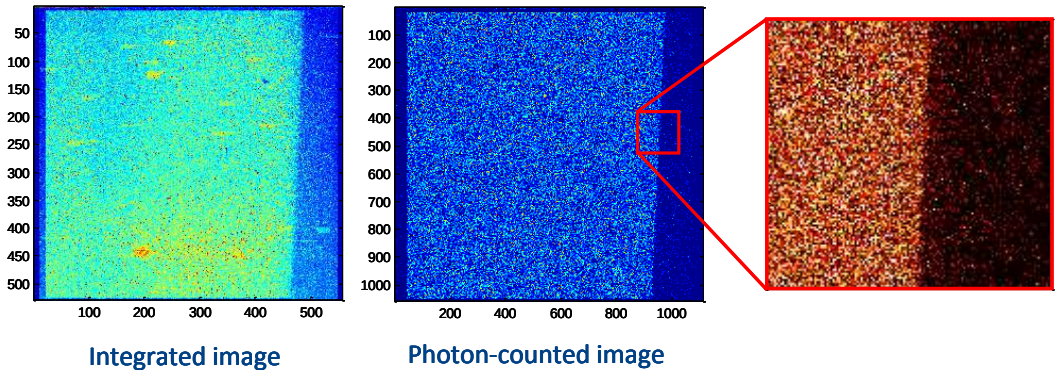

Figure 5.11. Experimental data showing images of the tungsten edge in the integrating and photon-counting modes. The photon-counting mode centroids the events to the nearest quarter pixel, hence the doubling in the number of rows and columns. An enlarged image of the edge in photon-counting mode is shown right.

5.5.2 Data analysis

Taking a best-fit to the LSF with a Gaussian profile provides a fit to the ESF through progressive integration across the function, Figure 5.13. Although the main slope of the ESF is well approximated, the upper and lower roll-overs are poorly matched. The upper and lower roll-over regions are over and under-estimated respectively. Although further analysis of this function provides a satisfyingly high MTF curve, the imaging capabilities are dramatically over-estimated.

The MTF generated from the fitted data shown in Figure 5.14, once normalised, provides a satisfying curve, as shown by the solid lines in the figure for the photon-counted and integrated images. The MTF generated directly from the experimental data can be 'falsely normalised' to coincide with the MTF from the fitted curve. At the low frequency end

(below 3 lp/mm) the experimental MTF data are found to be much 'higher' than that produced by the normalised single Gaussian fit.

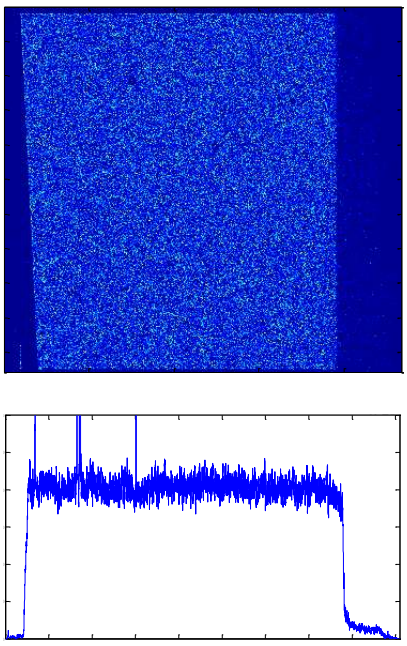

Figure 5.12. The photon-counted image is sheared to create a vertical edge. The edge is then binned to a single row, giving the profile of the ESF. The high-signal peaks are due to remaining bright-pixels.

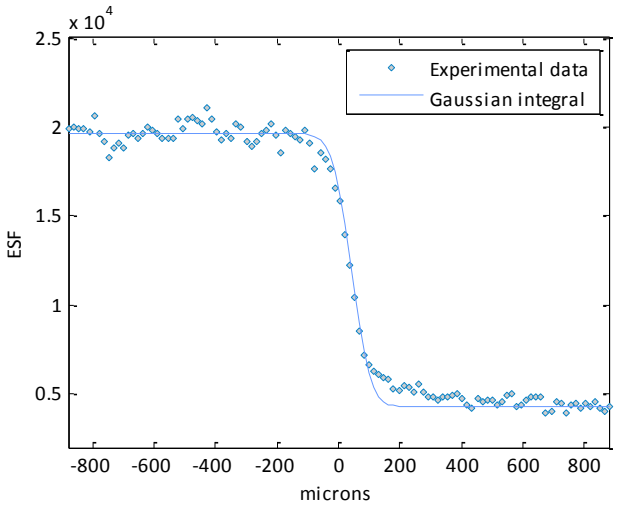

Figure 5.13. Taking the progressive integral of a Gaussian (fitted to the LSF) the ESF is not well approximated. Using a single Gaussian ignores the presence of the extended tails at the turn-overs of the function and hence these turn-overs are not well fitted. This can be compared with the more accurate fitting shown in Figure 5.15.

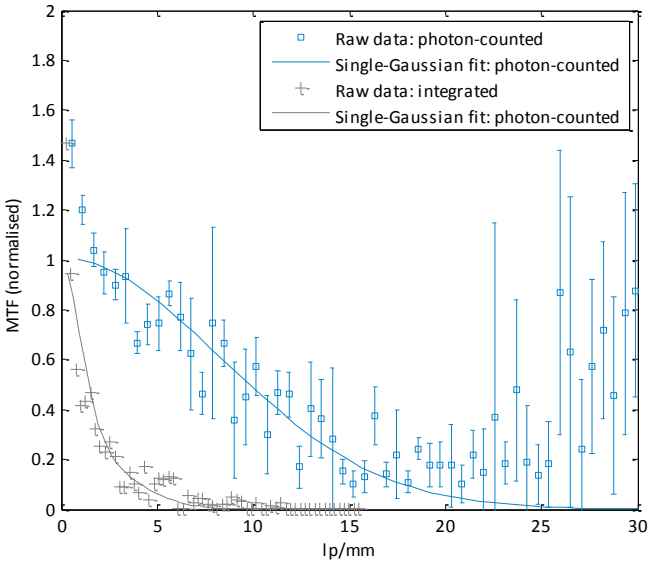

Figure 5.14. The normalised MTF curve produced from the single Gaussian fit to the LSF. Although the experimental data can be 'falsely normalised' to be consistent over the majority of the MTF curve, the lower frequency data points are seen to be much greater than the MTF calculated from the single Gaussian fit. The fluctuations at higher frequencies are due to high frequency noise in the experimental data.

The standard method of fitting a Gaussian to the LSF hides the processes which define the operation of the detector. In order to properly match the form of the ESF, further investigation is required. The slope of the ESF appears to follow two distinct forms. The sharp slope over zero microns is coupled to the slower fall-off to zero in the positive region. Both these components must be modelled to capture the full essence of the ESF.

It was found that the use of a LSF peak formed from the sum of two independent Gaussian profiles gave the best fit to the experimental ESF, Figure 5.15. The peak consists of a low amplitude wide Gaussian and a sharper peak of greater amplitude. The corresponding LSF is shown in Figure 5.16. The fitting function can be seen to be much more accurate over the full ESF.

The two-Gaussian fit is calculated through a least squares fitting routine to the LSF. The LSF is first fitted with a single Gaussian to find the central peak location in pixel space in order to minimise the number of parameters sequentially. The central position of the peak is fixed and a starting width (for the narrow peak) is taken from the initial fit. These parameters are fed into the fitting code to fit the two Gaussian peaks to the LSF (Figure 5.16). The fitted LSF is then normalised to give a 'total signal' of one. The normalised LSF can then be Fourier transformed to give the MTF, with the normalisation at the low spatial frequency limit already accounted for.

Following the manipulation of the experimental geometry, such that the source is placed 3 cm from the tungsten edge, the imaging capabilities of the detector can be further investigated. Although this placement is not optimal to test the spatial resolution of the device, the lab-bench test set-up allows much greater freedom than would be possible in a synchrotron facility. The final results will not display the full capabilities of the detector but will give a 'worst case' limit to the final device resolution. The ESF and consequent LSF produced, using the previously described method, are shown in Figure 5.17. The integrated image and photon-counted image can be compared directly, Figure 5.18, where the differences between the results are more obvious.

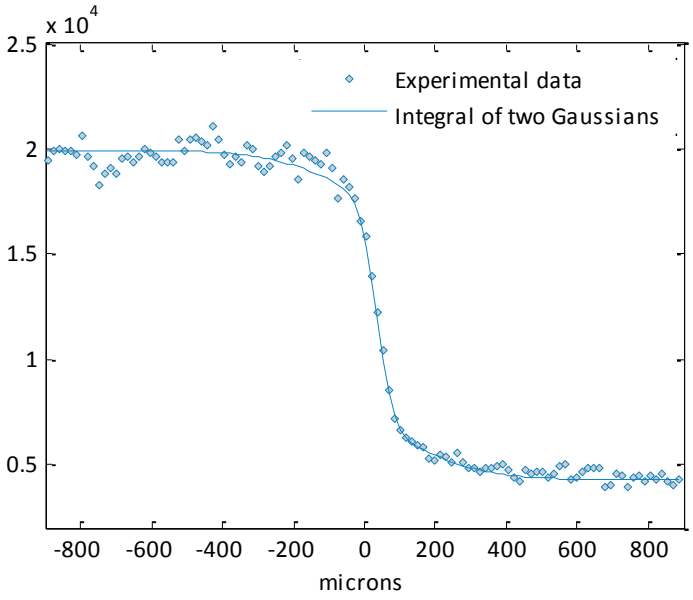

Figure 5.15. The experimental ESF data are more appropriately fitted with a profile formed from the sum of two Gaussians (shown in Figure 5.16). This fit can be compared with that of the single Gaussian function shown in Figure 5.13.

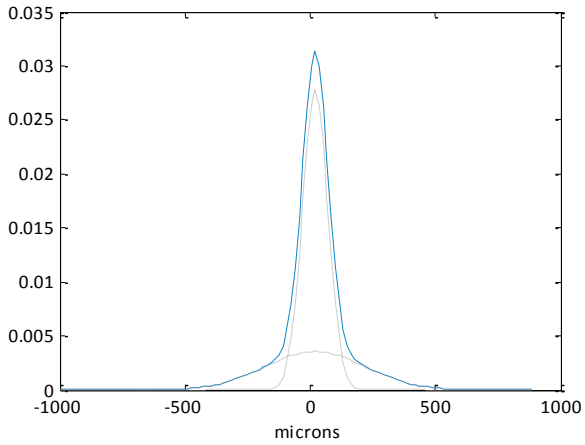

Figure 5.16. The fitted LSF corresponding to the ESF shown in Figure 5.15. The two Gaussian profiles are shown in grey (dashed lines), with the sum shown in blue.

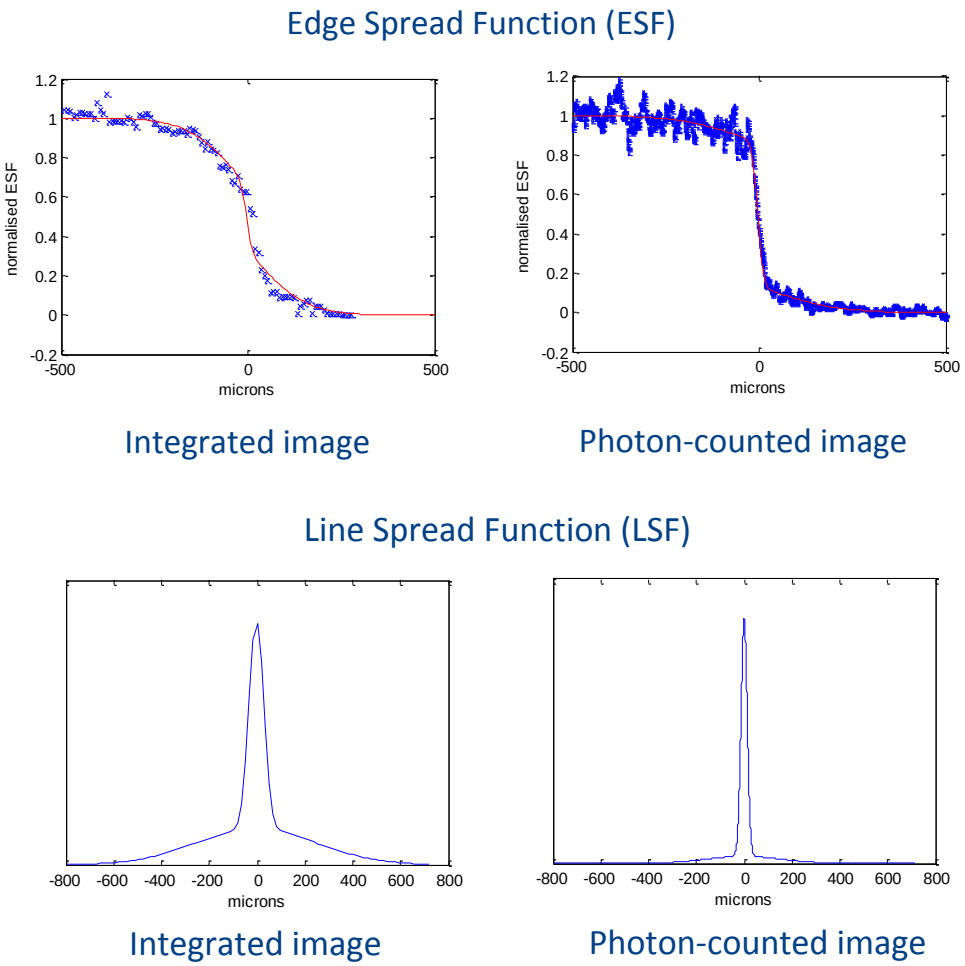

Figure 5.17. The ESF and LSF results for the integrated and photon-counted imaging. The ESF is shown with the experimental data in blue, showing the fitted curve (the progressive integral of the LSF) in red.

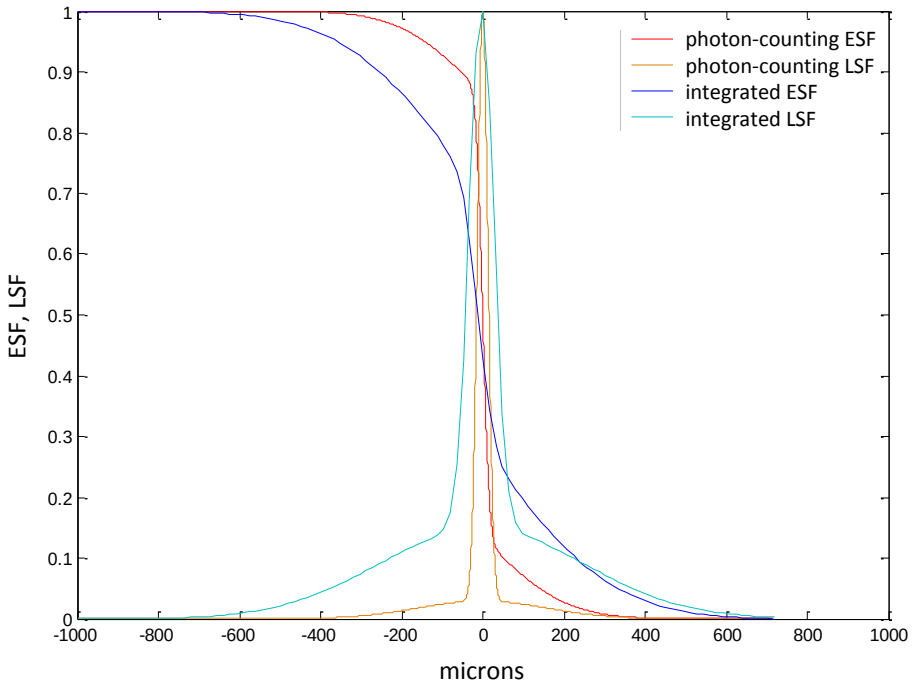

Figure 5.18. A direct comparison of the fitted ESF and LSF profiles in integrating and photon-counting modes. The photon-counting method shows large improvements in spatial resolution over the integrated mode operation.

The spatial resolution and MTF data are discussed in the following sections, but it is worth here considering the form of the functions. The summed Gaussian approach gives a much more appropriate fit to the data. The integrated image method provides a much more prominent broad component to the two Gaussian approach. If one considers two LSF peaks being brought from infinity (zero spatial frequency) to closer proximity (increasing spatial frequency), the space between the two peaks will decrease much more rapidly than in the photon-counting case. As the peaks come within approximately 1 mm of each other, the contrast difference will decrease between the light and dark portions of the peak. The decrease is much more pronounced in the integrated case than in the photon-counting case, where, for equal contrast reduction, the peaks must fall within approximately 0.1 mm.

Whilst discussing the spatial resolution of the detector, it will be taken as read that the LSF is formed using the summed Gaussian approach. The reasons for the use of the two distinct Gaussians are discussed in detail following the introduction of the energy spectrum produced by the detector in Section 5.9.

Referring back to Figure 5.14, showing the MTF generated by the experimental data using the single-Gaussian fit, the summed-Gaussian approach can be verified. A single Gaussian fit to the LSF is dominated by the narrow peak and therefore the form of the experimental data at higher spatial frequencies could be expected to coincide with the single-Gaussian fit. There are, however, problems occurring with the normalisation in the single-Gaussian fit. The experimental data can be seen in Figure 5.19 to be consistent with the summed-Gaussian fit across all frequencies (now including the low spatial frequencies), and the correct normalisation is achieved.

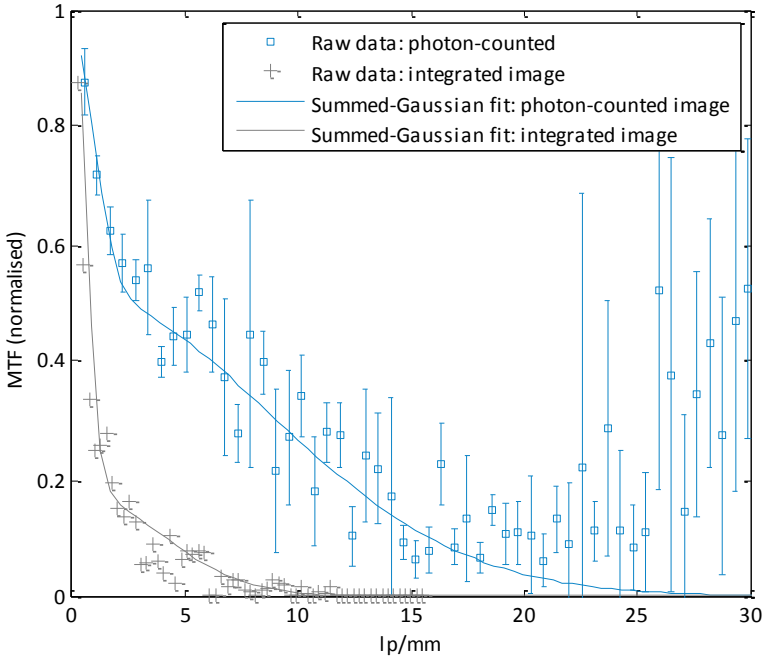

Figure 5.19. The MTF produced from the experimental data was found to be consistent with the summed-Gaussian fit. The normalisation was applied to the LSF for increased accuracy.

5.5.3 Spatial resolution

As discussed previously, the 'resolution' quoted in the literature is often misleading. There is little agreement as to a standard method for obtaining the spatial resolution of a device. For this reason, the MTF curve is often used to define the device capabilities across a range of spatial frequencies. The MTF has been calculated for this detector using the ESF and LSF from above, and the MTF data are presented in the next section.

Here, several single figures are quoted to 'define' the spatial resolution of the device, Table 5.1. The two main methods quoted use the FWHM of the LSF and the ESF at various ratios. The 10%-90% measurement from the ESF takes the distance from 10% of the maximum ESF value to 90% of the maximum ESF value. This value is somewhat misleading due to the extended tails on the ESF (due to the wider Gaussian of the LSF).

These values are quoted for imaging using all the recorded event profiles. Using the energy discrimination capabilities of the detector in photon-counting mode, the resolution can be improved. This process is discussed in Section 5.10, following the investigation of the energy discrimination capabilities of the device in Section 5.6.

Measurement	Integrated image	Photon-counted image
FWHM (LSF)	**83 µm**	**31 µm**
20%-80% (ESF)	125 µm	45 µm
15%-85% (ESF)	182 µm	53 µm
10%-90% (ESF)	250 µm	58 µm

Table 5.1. The imaging capabilities of the scintillator-coupled EM-CCD detector at 59.5 keV (using data taken from the fitted ESF and LSF). Several measurements are given due to the complex shape of the line spread profile.

5.5.4 Measured MTF

The FFT of the LSF is performed through automated code, with the scales adjusted to convert from the physical dimensions of the LSF to line pairs per millimetre (lp/mm). Figure 5.20 shows the MTF of the integrated and photon-counted images for comparison. The 'knee' seen in the MTF curves at 2-3 lp/mm relates to the double-Gaussian fit to the LSF. The ratio between the narrow and wide Gaussian profiles determine the position of the 'knee' (where the narrow or wide Gaussians dominate the detector response).

The formation of the two-stage MTF through the summation of two Gaussian profiles in the LSF is left to Section 5.9, as the energy discrimination results, detailed below, provide further insight into the effects on the detector resolution of the physical processes occurring in the scintillator.

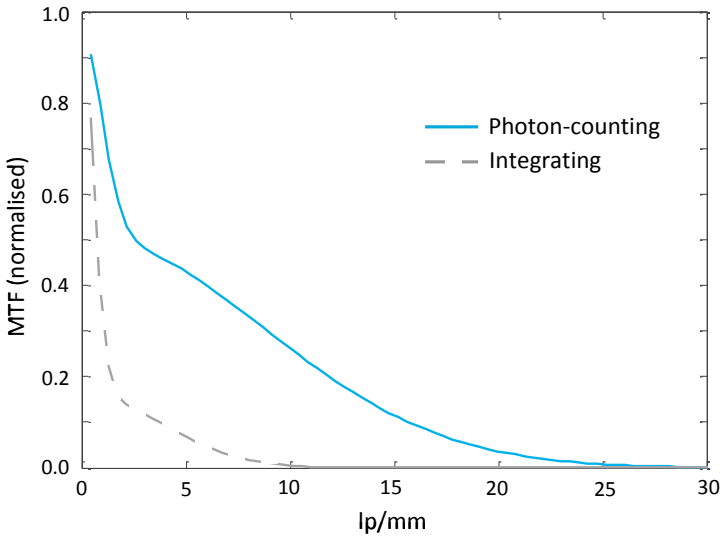

Figure 5.20. The MTF curves for the photon-counting and integrating methods. The improvements in resolution are dramatic over all spatial frequencies.

Referring back to the line spread functions shown previously, the wider Gaussian profile can be seen to be more prominent in the integrating case than in the photon-counting case. It was expected before the experimental testing that the improvement in resolution would result from the exclusion of the spread in the individual event profiles. In the integrating case the spread in the event profiles (from the photon passage through the scintillator) is included in the image and hence the image sharpness is compromised. In the photon-counting case the spread in the event profiles is removed. To confirm this hypothesis, the fitted LSF for the photon-counted case was convoluted with a standard event profile – a Gaussian with width appropriate to the measured event profiles. The results are shown in Figure 5.21 and it can be seen that the lower Gaussian and main peak are spread accordingly. The convolution does not represent the system perfectly, as each event shows a different spread, but the approximation provides sufficient consistency to the hypothesis. The removal of the individual event spread gives a large improvement in the spatial resolution of the detector. The use of photon-counting and the analysis of the individual event profiles also provide information regarding

the depth and energy of the interacting X/γ-ray, information that is lost in the integrating case. The following section describes the energy discrimination work carried out in this study and leads into a deeper investigation of the processes involved in the image formation and degradation. A method to improve the resolution of the detector using the outcomes of the study may then be considered.

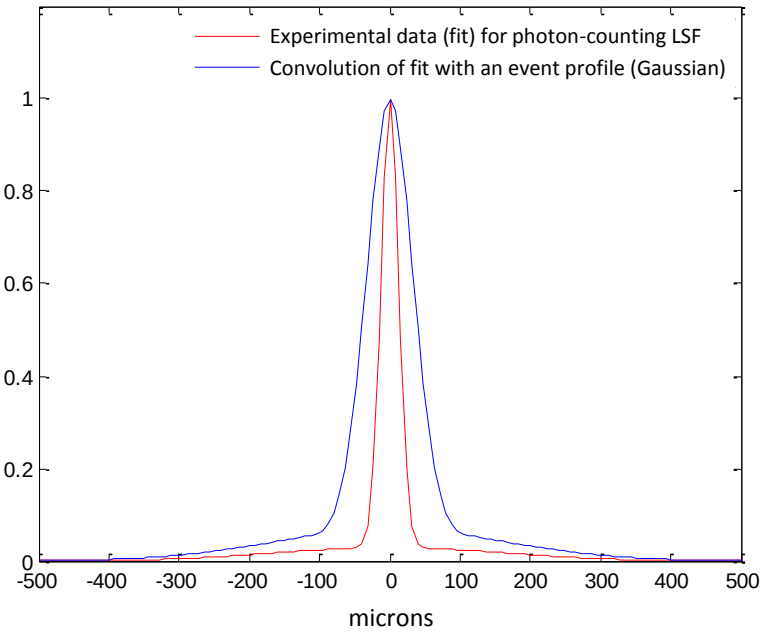

Figure 5.21. The fit to the experimental data LSF for the photon-counted data (red). The blue curve shows the result of convoluting an average event profile with the fitted LSF function.

5.6 Overview of energy discrimination methods

Various energy-discrimination techniques were introduced in Chapter 3. Here, a brief overview of the two chosen methods is presented with specific reference to the use of the methods with the experimental data obtained.

5.6.1 Scale-space

Although the scale-space analysis technique is the most time consuming of the methods presented, the results paint a more detailed picture of the event profiles. The accuracy of the event profile characterisation is determined by both the noise in the image and the number of levels in scale-space that are generated. Increasing the number of scales generated in the analysis code increases the accuracy, but also increases the time taken for processing. The scale-space method produces a 'Peak vs. Spread' (PVS) map of the event profiles in the image, Figure 5.22, providing information not only on the event profile itself, but also on the depth of the interaction and the presence on 'non-events' (such as bright pixels).

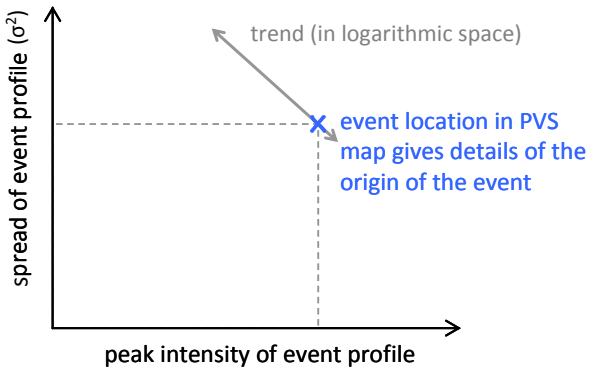

Figure 5.22. The 'peak vs spread' (PVS) map gives information about the origin of an event. The location of the event on the map details the depth of the interaction using the spread of the profile.

5.6.2 Event integration

Whilst the event integration method does not provide the same level of detail as the scale-space method, the analysis of the events can be implemented over much shorter time-scales. The event integration method provides a single energy value for each event. The measured energy results appear to outperform those found with the scale-space method, and as such,

this method was selected as the method of choice for further analysis (using a second gamma-ray source). If the run-time of the analysis is not required to be optimised, then the two methods can be considered in tandem, combining the results to get more information from the same data.

5.7 Energy discrimination for ^{241}Am

The ^{241}Am source emission is dominated by 59.5 keV gamma-rays. Each gamma-ray will generate, on average, approximately 3000 photons (of approximately 550 nm wavelength) when interacting in the scintillator. These photons spread radially from the initial interaction point, with some photons channelled vertically in the scintillator along the columnar structure of the scintillator. The number of photons detected per event by the EM-CCD should provide information as to the energy of the incident gamma-ray.

5.7.1 Scale-space

The scale-space code was implemented on the experimental data. For demonstration purposes the code can be edited to include several displays in order to show the results pictorially. Initially, a set of ten images was passed through the point-search optimisation code in order to specify the most appropriate parameters to be used in the point search function.

The scale-space code takes the events from the point-search function and convolutes the events in the images with Gaussian profiles from σ = 0.1 to 6.0 in steps of 0.1, producing a stack of images for each profile. The code applies the calculations detailed in Chapter 3, fitting a curve to the final data array. The fitted curve provides the peak intensity and the spread (or σ value) for each event profile.

The peak intensity and spread data for the profiled events can be combined to form a single plot, the 'Peak vs. Spread' (PVS) map, showing the position of the event in peak/spread space, Figure 5.23. The PVS map shows clearly the events corresponding to bright pixels, allowing these events to be selected and excluded from the energy spectrum and any future image

processing. The lower right corner of the PVS map shows events with a high peak intensity and low spread – small numbers of pixels with very large intensities (such as a single bright-pixel). Around the outside of the bulk of the events shown in the centre of the map may be found events for which the code could not properly calibrate the profile. These events can again be removed from the energy spectrum. The diagonal form of the PVS map is used to generate the energy spectrum.

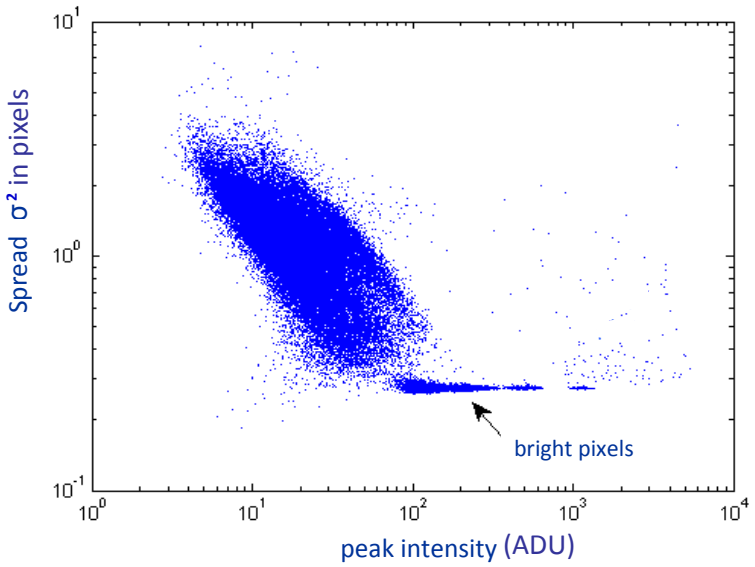

Figure 5.23. The 'Peak vs. Spread' (PVS) map provides much greater detail on the event profiles detected. The presence of bright pixels is notated, whilst the two bands of events can be seen, relating to the two peaks seen in the energy spectrum (Figure 5.24), explained further in Section 5.8. There are very few events seen around the main grouping (those which cannot be profiled accurately due to the noise).

Following the gradient of the bulk of events shown in the PVS map, the events are summed to give the energy spectrum displayed in Figure 5.24 showing two main peaks. A major peak falls

at approximately 30 keV, relating to the re-absorbed fluorescence (from the Cs and I in the scintillator) and the escape peak. The peak at 59.5 keV corresponds to the full energy peak.

A point of interest may be found between the two main peaks. The 30 keV peak is not symmetrical and shows a large 'bulge' at approximately 30-40 keV. This discrepancy was investigated and was found to be due to the scale-space code used to characterise the events. Events lying in close proximity can be merged under convolution with wider Gaussian profiles. The scale-space code is adversely affected by events in close proximity, increasing the calculated spread and intensity of the events. The combined spread and intensity data (used to calculate the energy) leads to an overestimation below approximately 5 pixels separation, Figure 5.25. As the separation is increased, the calculated energy converges to the correct value.

Using the results shown in Figure 5.24, it is not ideal to quote an 'energy resolution' as such. The events can be split into approximate 60 keV and 30 keV regions, but the merging of the peaks in the central region is a cause for concern. The method has major benefits in the ability to analyse events in more detail from the PVS map, but the 'energy resolution' using other methods has been found to be preferable.

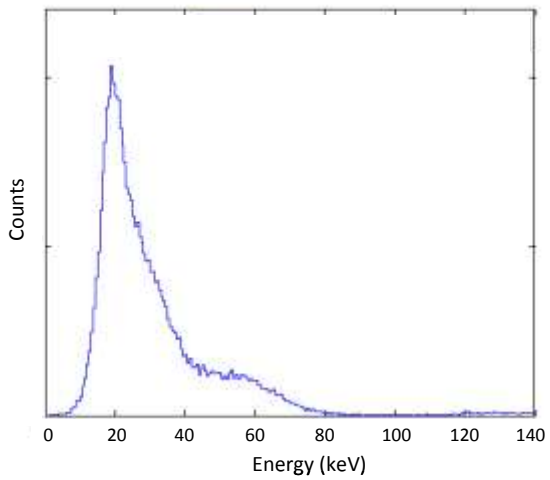

Figure 5.24. The energy spectrum generated for ^{241}Am using the scale-space method.

5.7.1.1 Depth of interaction

The spread of a registered event profile is independent of the energy of the interacting X/γ-ray. The spread of each event is dominated by the depth of interaction of the X/γ-ray in the scintillator. Through the use of the PVS map, the depth of interaction of the X/γ-ray in the scintillator can be approximated.

It may be possible, through future study, to use the depth of interaction data gained through the scale-space analysis to further improve the energy determination of the system. If an interaction occurs far from the CCD then the visible photons passing through the scintillator will suffer increased scattering and increased absorption by the scintillator. These processes act to decrease the intensity of the event profiles and lead to a broadening of the energy spectrum. Through the use of the depth data one may introduce a compensation factor to account for the losses from interactions occurring far from the CCD. This is a method which has been left for future study.

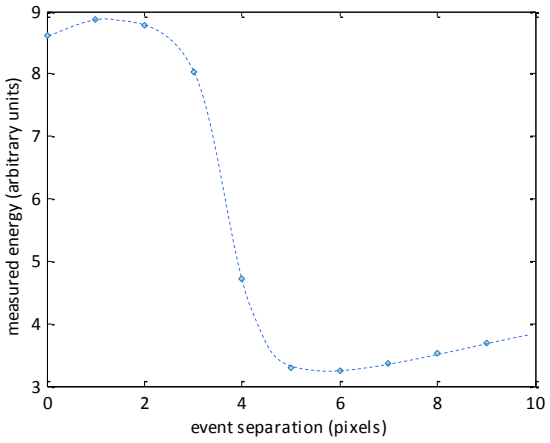

Figure 5.25. The number of pixels between two events affects the energy measured if the events are in close proximity. Below a separation of 5 pixels, the energy is overestimated due to the combined spread of the two events when smoothed. As the separation increases, the energy approaches the real value (4 in this case) for the two individual events.

5.7.2 Event integration

The second method of energy discrimination discussed previously is that of event integration. After selecting the central pixel of an event through the point-search process, an area of 11 × 11 pixels is selected, centred upon the brightest pixel. The area of 11 × 11 pixels was chosen to optimise the energy spectrum produced. Areas of 9 × 9 pixels and below, whilst producing separable spectral peaks, were found to produce a broader energy spectrum. The spectral broadening occurs due to the loss of the signal in the outer pixel region of the event profiles. For areas of 13 × 13 pixels and larger, the signal included from the outer regions of the event profiles is reduced to below the level of the noise in the image. Including pixels outside this area brings increased noise to the 'energy' values, again broadening the spectral peaks.

Using an area of 11 × 11 pixels, the spectrum for ^{241}Am was produced using the same raw data as used for the scale-selection method. The spectrum is shown in Figure 5.26, and shows three peaks. The higher energy peak at 59.5 keV corresponds to interactions in the scintillator where all the photon energy is transferred to the scintillator lattice in one instance. The two lower energy peaks, just above and below 30 keV, correspond to the escape peaks and the re-absorbed fluorescence X-rays. The fluorescence X-rays correspond to the K-shell α and β emission from the caesium and iodine atoms in the scintillator lattice. The separation of the two lower energy peaks cannot be found in the literature due to the much broader spectra produced in previous studies. The cartoon (orange) shown in Figure 5.26 over the experimental data (blue) displays the expected formulation of the energy spectrum.

The 'energy resolution' of the detector in this instance should only be loosely defined, as although the peaks can be seen individually, the detector is only truly capable of 'energy discrimination' as opposed to more detailed spectral analysis. The spectrum can, however, be used to improve the spatial aspects of the detector through the removal of unwanted events (discussed in Section 5.10).

Regarding the energy resolution of the device, approximate values can be quoted where the presence of the other peaks is ignored. At 59.5 keV, a FWHM of can be quoted by assuming a symmetrical peak and doubling the half-width using the right-hand side of the peak (unaffected by the 30 keV peaks). The FWHM of 19.7 keV (assuming a symmetrical peak) at 59.5 keV gives an energy resolution of 33 % at 59.5 keV, Figure 5.27. The resolution is harder to calculate at lower energies due to the overlapping of several peaks, but a rough approximation from the data gives a FWHM of 8.5 keV at 24 keV (the lowest energy peak), giving a resolution of 35 % at 24 keV. This demonstrates an approximately constant 'energy resolution' of 30-40 % in keeping with other scintillator/gas counters (see Section 2.5), whilst maintaining a much better spatial resolution.

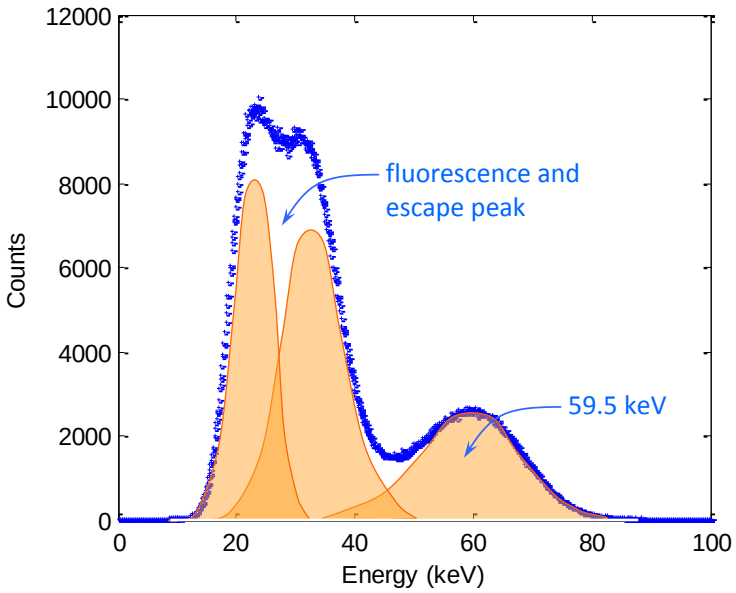

Figure 5.26. A representation of the different components forming the energy spectrum for ^{241}Am. The blue points show the experimental results achieved in this study. The orange peaks show the expected distributions. The lower energy peak is more complex than shown but is demonstrated by two peaks for simplicity.

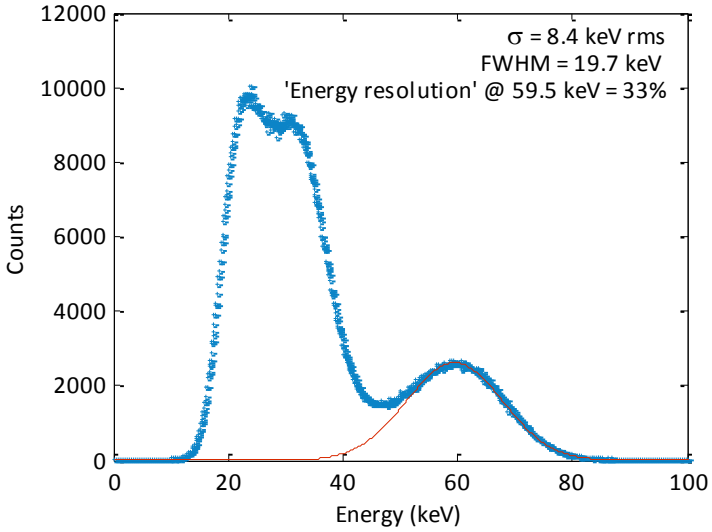

Figure 5.27. A Gaussian curve was fitted to the 59.5 keV peak, giving an 'energy resolution' (ΔE/E) of 33% at 59.5 keV.

5.8 Energy spectrum

The interactions occurring in the scintillator when subjected to gamma-ray irradiation require a more detailed analysis to ascertain the spectral and imaging capabilities of the detector. Unfortunately, the interaction process does not simply supply a series of identical scintillation flashes all occurring at the incident gamma-ray energy. Taking the example for the 59.5 keV gamma-rays emitted by an ^{241}Am source from the experimental data detailed previously, the following simulation and analysis of the spectral capabilities of the detector allow further analysis of the imaging capabilities of the device.

For the perfect scintillator-based detector one must only supply X-rays with energies below the binding energy of the atoms present. For caesium iodide based scintillators the binding energies of importance are those for caesium at 35.99 keV and iodine at 33.17 keV. Incident photons of lower energy cannot fluoresce the caesium or iodine K-shell electrons. For hard X-ray imaging, however, the addition of the K-shell fluorescence to the interaction chain complicates the analysis.

Approximately 90 % of interactions in the caesium iodide at 59.5 keV are due to the photo-electric effect. Of these interactions, the fluorescence yield is approximately 88-90%. The differences between the probability of interaction with the caesium and iodine atoms is negligible. At 59.5 keV the incident photons have sufficient energy to knock an electron from the inner shell of the atoms, releasing an electron of energy equal to the incident photon energy minus the binding energy for the atom. The ejected electron has insufficient energy to travel far from the initial interaction position, but the few microns over which the electron travels degrades the imaging capabilities to a small degree. Excitations remaining in the atom, combined with the scintillation from the ejected electron leave a flash of photons in the scintillator (Chapter 2).

The Auger process accounts for approximately 10 % of the relaxations of atoms following removal of the K-shell electron. The electrons emitted will cause scintillation surrounding the initial interaction position and will account for the 59.5 keV peak (full energy) in the energy spectrum. The influence on the image quality due to the electron motion at this energy is minimal and limits the resolution by a few microns due to the spread of the electrons from the initial interaction position.

The remaining 90 % of relaxations result in K-shell X-ray fluorescence, where an outer shell electron falls to fill the hole left by the ejected electron, releasing the difference in energy in the form of a characteristic X-ray. The Kα emission dominates over the Kβ emission by 90 % to 10 %. The fluorescence X-rays may travel through the scintillator and leave the material undetected. In this case, only the ejected electron and excitations in the atoms will produce visible photons through scintillation. This accounts for a part of the peak produced at approximately 30 keV. If, however, the fluorescence X-ray interacts with the scintillator, a secondary interaction site will be noticed relating to an energy of approximately 30 keV in this case. The re-absorption of the fluorescence X-ray leads to an interaction site away from the

initial interaction position. The position cannot be determined in advance, but is subject to a probability of interaction at any given distance.

For the case detailed here, ^{241}Am at 59.5 keV, problems are noticed due to the similarity in energy of the re-absorbed fluorescence peak and the escape peak, both being formed from components around the 30 keV mark. The interaction chain is detailed in the flowchart in Figure 5.31 at the end of this section, with the real energies calculated for each possible outcome. In order to simulate the energy spectrum produced, one must consider not only the scintillator materials and energies involved, but also the dimensions of the scintillator, as these determine the probability of the re-absorption of the fluorescence X-rays.

A Monte Carlo simulation was designed to simulate the energy spectrum produced by the ^{241}Am source. A series of gamma-rays are produced, diverging from the ^{241}Am source. The simulation uses the appropriate interaction cross-sections to determine the interaction distance of each incident gamma-ray in the scintillator. If this interaction position is found to fall outside the scintillator geometry then the gamma-ray escapes. If, however, the gamma-ray interaction position lies inside the scintillator geometry then the interaction chain proceeds. Using the appropriate interaction cross-sections the interaction chain progresses with energy deposition noted in the log file. The interaction energies are calculated and summed to form a spectrum of delta peaks. The width of the peaks (the energy resolution of the detector) can be added from this point to varying degrees of complexity.

The plot in Figure 5.28 shows the proportion of each interaction that builds up the energy spectrum produced, as shown in Figure 5.29. This correlates very well with the spectrum found experimentally. The distance between the initial interaction and that of the re-absorbed fluorescence X-ray has also been simulated. The results, shown in Figure 5.30, indicate the problems that the re-absorbed fluorescence pose to the imaging capabilities of the device. If one separates the initial interaction positions from those of the re-absorbed fluorescence

X-rays the spatial resolution would be expected to improve dramatically. This is investigated in the following sections and shows the improvements possible through energy discrimination and analysis.

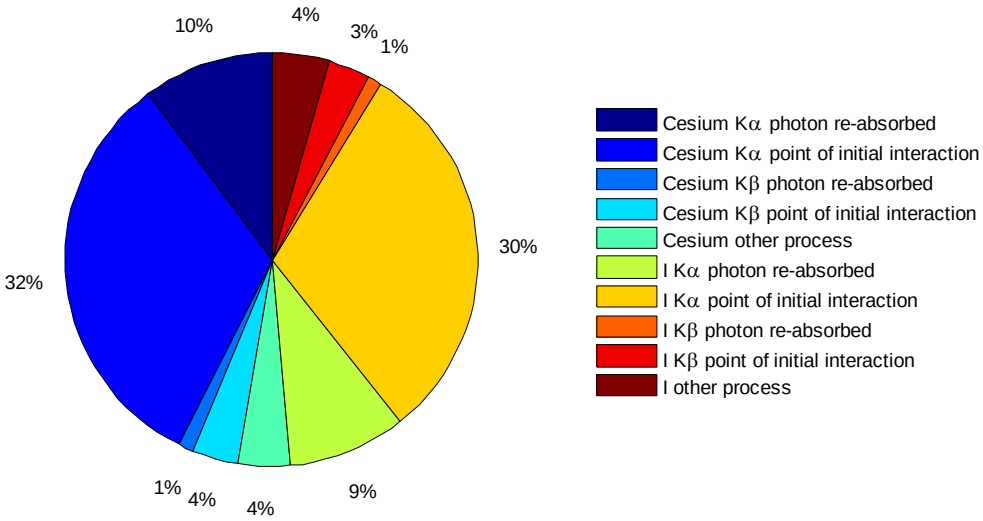

Figure 5.28. The proportions shown in this figure relate to the interaction chain in the scintillator. These show the complications in spectral distinction with the scintillator-based detector in high-resolution imaging.

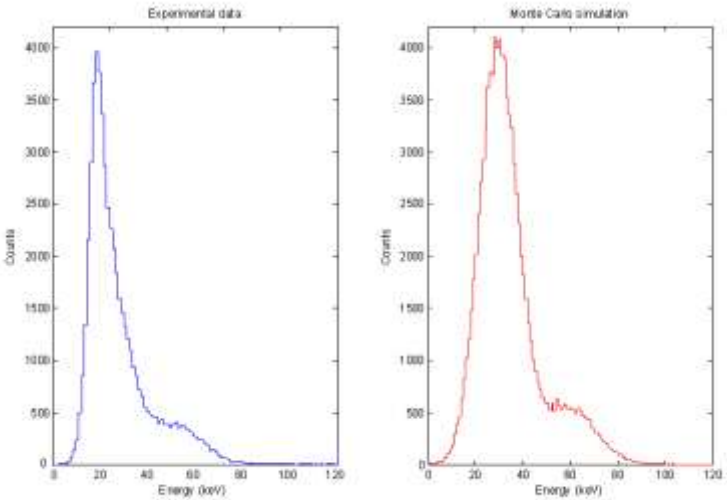

Figure 5.29. Spectra of ^{241}Am, measured and simulated. The experimental data (left) shows an energy spectrum using the scale-space method. The Monte Carlo simulation results are displayed for comparison (right) using a pre-determined energy resolution calculated from the experimental data.

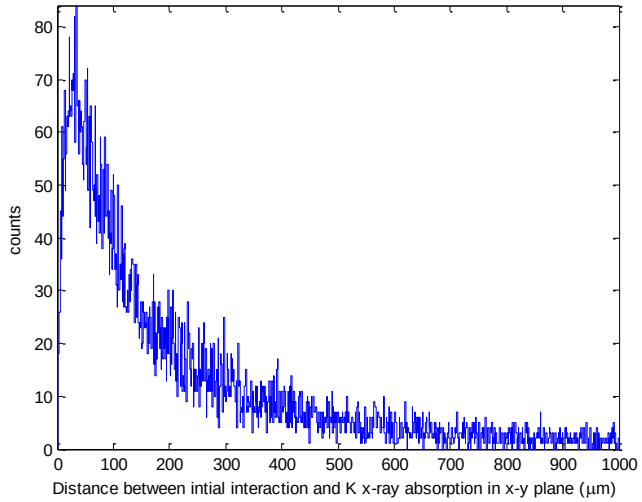

Figure 5.30. The distance between the initial interaction position and that of the re-absorbed fluorescence X-ray poses problems when high-resolution imaging is required, simulated using bespoke C++ code (written by the author) and the appropriate geometry and absorption cross-sections. If the detector spatial resolution is low then the reabsorbed fluorescence X-ray has a greater probability of being recorded as part of the initial event.

The ^{241}Am spectrum observed with the scintillator-based detector is relatively simple compared to the spectra produced by higher energy X-rays/gamma-rays. The photo-electrons and fluorescence X-rays do not have sufficient energy to knock out further electrons from the caesium and iodine atoms to produce further Auger cascades or fluorescence X-rays. At higher energies, the process is further complicated by the possibility of further fluorescence X-rays being emitted in the reaction chain.

Taking ^{57}Co as an example, the second source used in the experimental work, the 122 keV component can be examined. The chain of reactions follows a similar initial pattern as that detailed for ^{241}Am in the flowchart shown in Figure 5.31. The photo-electron produced for the caesium atom now has an energy of 86.01 keV, which can combine with the excitations following the fluorescence to give a second peak (or escape peak) in the energy spectrum at 91.03 keV. The energy will be deposited over a slightly larger area, with the electron able to travel of the order of 10 µm (Casino based calculation, see Drouin, 2007). This electron has sufficient energy to break the binding energy of the inner-shell electrons and fluorescence can occur for a second time. Following the emission of a secondary fluorescence X-ray, the energy remaining can be deposited within approximately 5 µm (Casino based calculation), leaving a secondary escape peak at 55.04 keV. The ejected electron still has sufficient energy to break the binding energy and a further peak is possible at 19.05 keV, emitting a further fluorescence X-ray.

Repeating this process with the iodine atoms leads to a complex spectrum with four main peaks at 122 keV (initial peak), ~90 keV (escape), ~60 keV (secondary escape) and ~30 keV (fluorescence re-absorption). The experimental data for the ^{57}Co source is presented in Section 5.11.

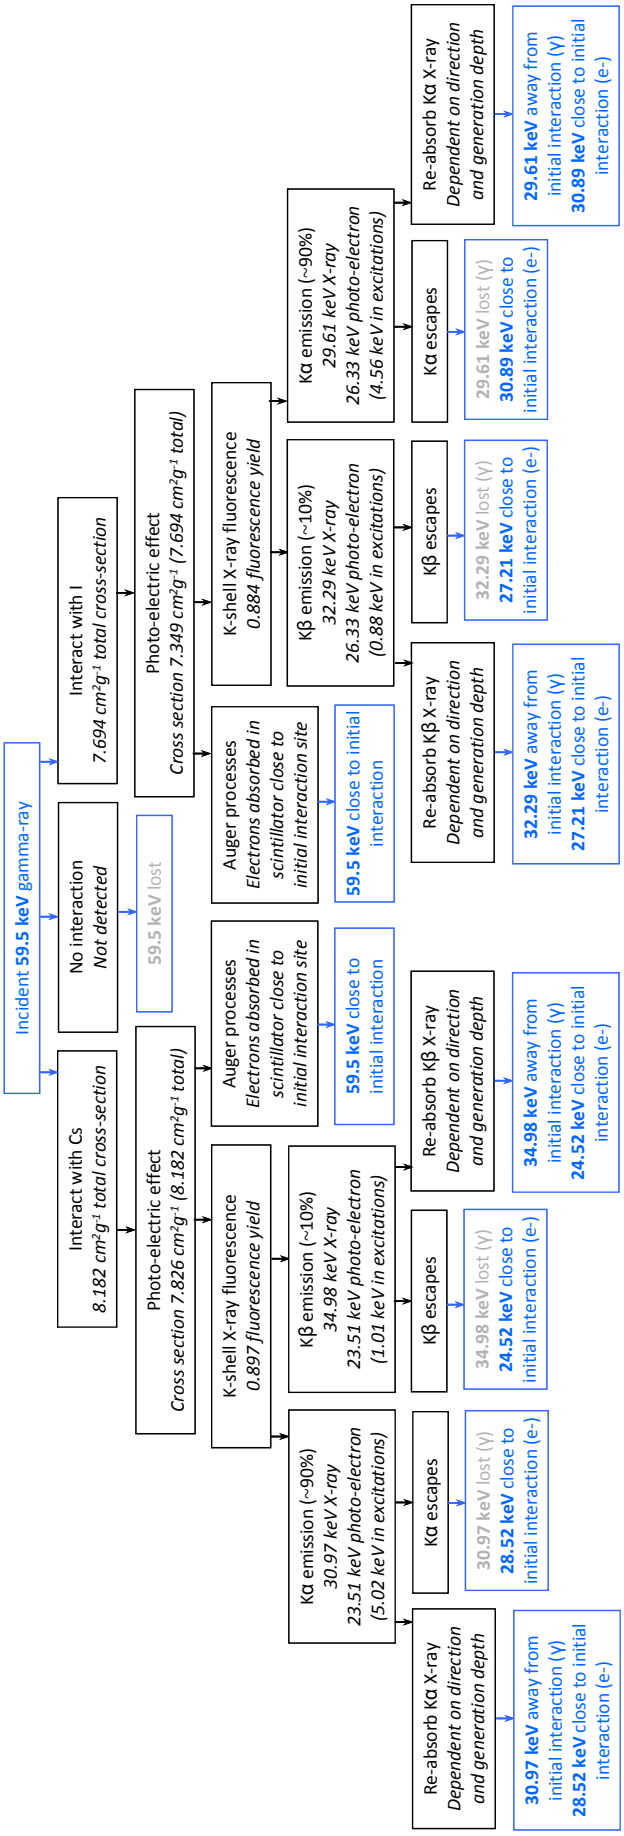

Figure 5.31. Flowchart of the interactions occurring inside caesium iodide under irradiation by 59.5 keV gamma-rays from an Am^{241} source, designed in this work to aid the simulation code.

5.9 The impact of fluorescence on the resolution

The 'two-Gaussian' approach implemented earlier closely relates to the energy spectrum produced by the detector. The energy spectrum is seen to be composed of primary interaction peaks (at 59.5 keV and the escape at approximately 30 keV) and a secondary peak at approximately 30 keV corresponding to the re-absorbed K-shell X-ray fluorescence from the caesium and iodine atoms in the scintillator lattice.

The primary interaction locations produce the sharp, central peak in the LSF. The spread in the sharper central peak is determined by the detector system and analysis technique. The wider Gaussian found in the LSF is of greater interest here. This secondary peak has been found to be produced by the re-absorption of the K-shell X-ray fluorescence generated by the caesium and iodine. The fluorescence X-rays can travel some distance through scintillator in the imaging plane before interacting and producing a secondary event profile for a single incident X/γ-ray. This spreading outwards of the signal causes a much wider peak. The width of the secondary peak is determined by the scintillator's elemental make-up and the geometry of the scintillating layer used. The elements present in the scintillator produce the K-shell fluorescence X-rays, in this case caesium and iodine. The material chosen for the scintillator (and the corresponding fluorescence yields and absorption coefficients) therefore determine the intensity and spread of the secondary LSF peak.

The re-absorption of the internal fluorescence in the scintillator is unavoidable and one can only ensure the chosen scintillator is suitable for the application. The fluorescence can, however, be partially removed using the energy discrimination capabilities of the device, although this leads to a decrease in counts per image. A brief simulation of the effect of fluorescence re-absorption is presented below to show the impact of fluorescence on a hypothetical 'perfect device'.

5.9.1 Monte Carlo simulation of a 'perfect' device

In order to demonstrate the impact that the re-absorption of the internal fluorescence can have on the imaging capabilities, a Monte Carlo simulation was performed for a 'perfect device' using code written in this study. Here, the 'perfect device' is defined as a camera system for which the MTF of the primary interactions has a value of one over all spatial frequencies. The LSF for such a device gives a delta function, and the ESF is a step function from one to zero. The device is shown to have 8 µm pixels for ease of display.

The simulation is coded in Matlab and allows customisation of the geometry of the scintillator. In this case, the scintillator is defined as a single 70 µm thick block of caesium iodide, with side-lengths of 800 µm (100 quarter-pixels) and 1600 µm (200 quarter-pixels). The appropriate interaction cross-sections and fluorescence yield data were entered into the simulation and used to generate the random interactions appropriately in the scintillator.

A series of 59.5 keV gamma-rays were simulated, entering the scintillator perpendicular to the plane of the image. The incident gamma-rays are only placed into a 800 µm × 800 µm region at the edge of the scintillator (the left-most 100 × 100 'quarter-pixels'). The interactions of the gamma-rays in the scintillator are determined by random number generation combined with the interaction cross-sections, with each gamma-ray traced until interaction with, or escape from, the scintillator geometry. If fluorescence occurs then a second X-ray is produced and the path of the secondary X-ray is traced. The interaction positions of the X-rays/gamma-rays are recorded for analysis. The results are shown in Figure 5.32, where the colour-scale runs from blue to red (low to high signal).

The image shows that the majority of the events are recorded in the region of the detector in which the incident gamma-rays were placed. The signal is not, however, constrained to the area bombarded with gamma-rays and spreads into the incident-radiation-free region of the simulated scintillator.

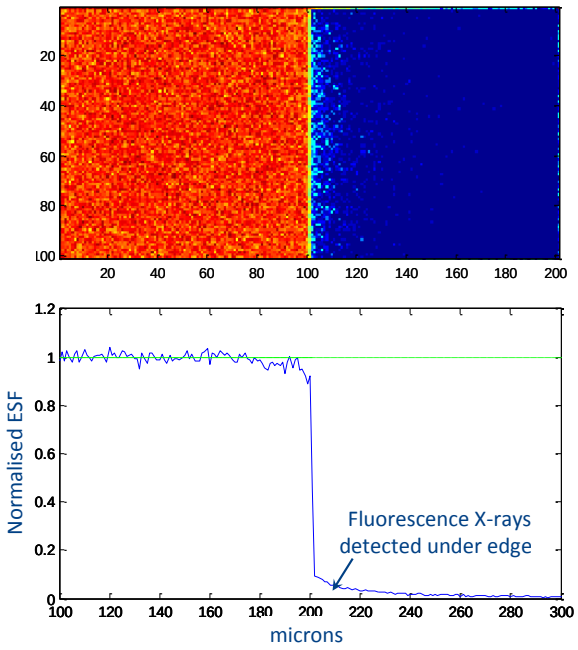

Figure 5.32. The upper panel shows the image output from the simulation of a 'perfect' detector, showing the fluorescence X-rays creeping into the un-irradiated region (right). The lower panel shows how the ESF is displaced from the 'perfect device' step function.

5.10 Improving the MTF

With a standard integrated image there is no way of regaining any energy information about the gamma-rays producing the image. Through the use of the energy discrimination methods described above, using the photon-counting imaging technique, post-radiation image techniques can use the energy information for each event to improve the image definition.

We know the different components from which an image is produced from the simulations described previously. The full-energy events are those which are most accurately located with regards to the mask – these events are found with no translational motion in the plane of the image. Escape events are generated following a small motion from the photo-electrons. These electrons can only travel a few microns, having minimal impact on the resolution, but an

impact none the less. The re-absorbed fluorescence events have the greatest impact on the image quality, with the fluorescence X-rays able to travel many tens or hundreds of microns before being re-absorbed (and hence generating a secondary event profile).

The energy spectrum for ^{241}Am will be considered here, with the two main peaks at 59.5 keV and approximately 30 keV. The lower energy peak is formed from many fluorescence related components merging into a single peak which cannot be successfully split. One can split the energy spectrum into two parts, one for lower energy events and one for higher energy events (Figure 5.33). The high energy events are, in the majority, the 59.5 keV events (total energy deposition). The tail from the low energy peaks may spread into this region, but the number of counts must be kept to a reasonable level and hence the split must optimise both the energy distinction and the count rate.

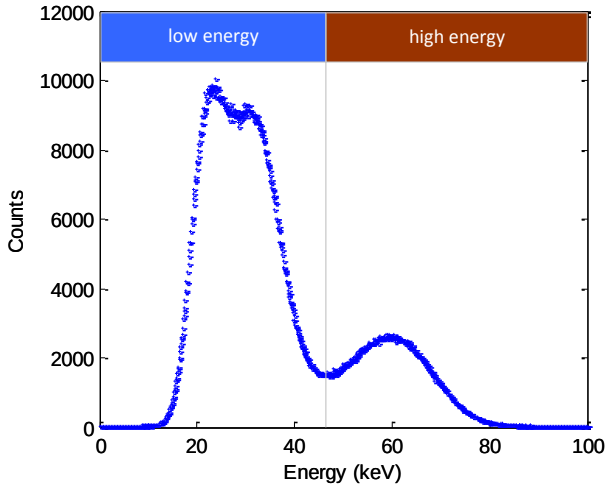

Figure 5.33. The measured energy spectrum for the 59.5 keV ^{241}Am source, showing the split in energy used to demonstrate the following analysis.

One of the many benefits of the photon-counting method is the storage of the image data in a three-column array (row, column, summed signal). The energy value was used to sort the

array into the high and low energy events and two images were generated using the energy separated data. Analysing the image using the same code as that formulated to calculate the results shown previously, the ESF, LSF and MTF were generated.

The edge spread functions for the three bands (low energy, high energy and combined data) are shown in Figure 5.34. The number of counts in each grouping must be taken into account when considering the figures shown below (each plot is normalised, although the number of counts causes subtleties in the ratios of peak heights, particular with regards to the LSF). The 'all energies' data are the same as those presented earlier in the chapter and are included for ease of comparison.

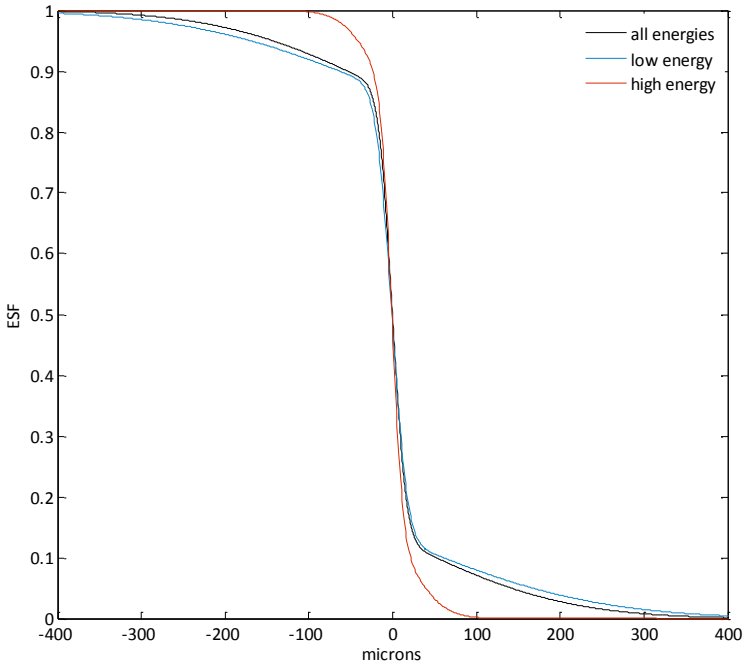

Figure 5.34. The edge spread functions for an edge image taken using all energy events, low energy events (approximately 30 keV) and high energy events (approximately 59.5 keV) using the experimental data shown previously. The impact of the fluorescence is visible in the tails of the ESF.

The low energy ESF shows a slight decrease in counts in the open section of the image (left side of Figure 5.34) and a corresponding slight increase in counts in the covered section of the image (right side of the figure) – note that the ESF has been normalised to one for direct comparison and that the number of total counts is actually decreased. This is due to the fluorescence X-rays passing away from the initial interaction positions. The 'low energy' data are formed from a combination of the reabsorbed fluorescence and the escape events.

The high energy data show a much improved ESF. The maximum gradient in the ESF (determined by the device itself) matches that found with the two alternative data sets, as would be expected. The fall off in the ESF occurs much closer to the edge transition in the high energy case as the majority of the fluorescence events have been removed. Through the removal of the fluorescence events, it can be seen that the imaging quality of the overall detector has been improved dramatically. This can be seen more clearly with the line spread functions, in Figures 5.35 and 5.36.

The line spread functions, shown in Figure 5.35, demonstrate the impact of the energy discrimination further still. The main peak of the LSF is dominated by the primary events (59.5 keV events) and the escape events. Some fluorescence events remain, but show much lower impact than in the all/low energy data (displayed more clearly in the non-normalised data in figure 5.36). The narrowing of the central peak is thought to be due to the exclusion of the bulk of the escape events from the analysed data. The escape events are spread by several microns due to the passage of the electrons through the material. This extra spread of only a few microns is consistent with the shift in peak width shown in the line spread functions.

The spatial specifications of the detector are shown in Table 5.2, giving a numerical comparison between the three data sets. Through the removal of the low energy events, leaving only the high energy events to be analysed, the FWHM of the LSF has been improved

from 31 µm to 25 µm, giving an improvement of approximately 20 % over the original data analysis.

The use of 'quarter pixels' is thus a very appropriate choice. With events taken to 'whole pixel' positions, the 'maximum full resolution' could be seen as restricted to 32 µm (two pixels). The use of 8 µm pixels fits very well with the maximum measured resolution at 25 µm, fractionally over four pixels. These results are taken with the non-optimal lab-based geometry – with a more optimal testing set-up the pixel size may require further reduction. This can be assessed following initial results with the system, where the measured resolution using quarter-pixels can be compared to the theoretical resolution limit.

Two points should be recalled here. The first is that the split in energies used in this data aims to maximise the count rate in the high energy case. If the split was taken several keV higher then the spatial resolution could be improved further, although this would lead to a loss of larger amounts of data. The second point relates to the geometrical set-up of the testing system. The system was optimised for lab-bench testing and low cost requirements. Further testing is planned for future work using a synchrotron beam-line at varying energies (Section 5.13). With a parallel beam of gamma-rays the resolution is expected to improve further still. A second generation device is proposed in Section 5.13, using the results found in this study to optimise the system and to provide the best-case performance following testing at a synchrotron beam-line.

Measurement	Low energy events	All events	High energy events
FWHM (LSF)	35 µm	31 µm	**25 µm**
20%-80% (ESF)	55 µm	45 µm	41 µm
15%-85% (ESF)	60 µm	53 µm	47 µm
10%-90% (ESF)	68 µm	58 µm	56 µm

Table 5.2. The imaging capabilities of the scintillator-coupled EM-CCD detector at 59.5 keV. The 'low energy' results relate in the majority to the escape and fluorescence events, 'high energy' results refer to the 59.5 keV events only, and 'all events' refers to the full data set. The improved FWHM relates to the removal of the escape events. A more dramatic improvement is seen when the fluorescence event removal is observed in the LSF and MTF.

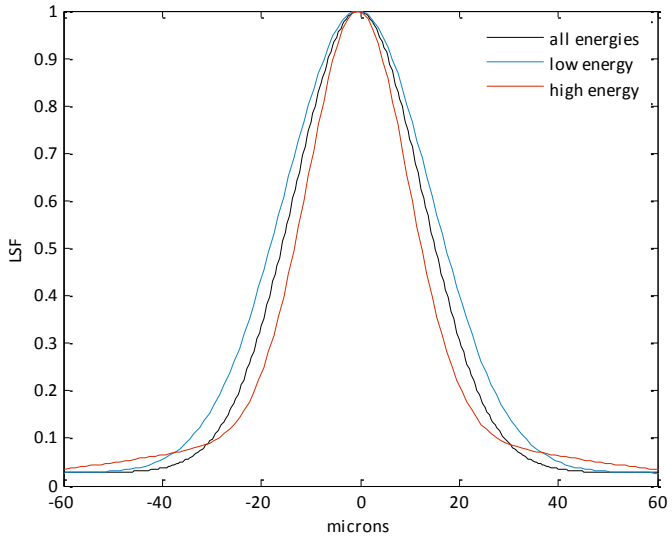

Figure 5.35. The normalised line spread functions for an edge image taken using 'all energy' events, 'low energy' events (approximately 30 keV) and 'high energy' events (approximately 59.5 keV). The results show a narrowing of the peak and hence an improvement in resolution when energy discrimination is used.

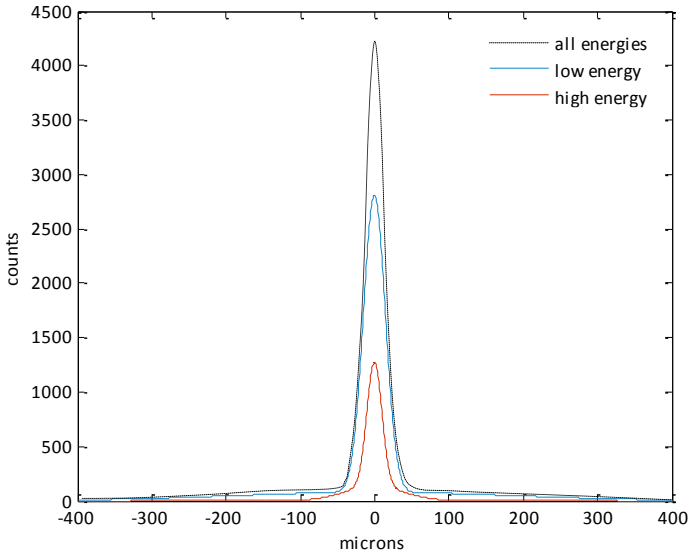

Figure 5.36. The line spread functions for the three energy categories, shown here with no normalisation. Here it can be seen that the fluorescence tails are much reduced.

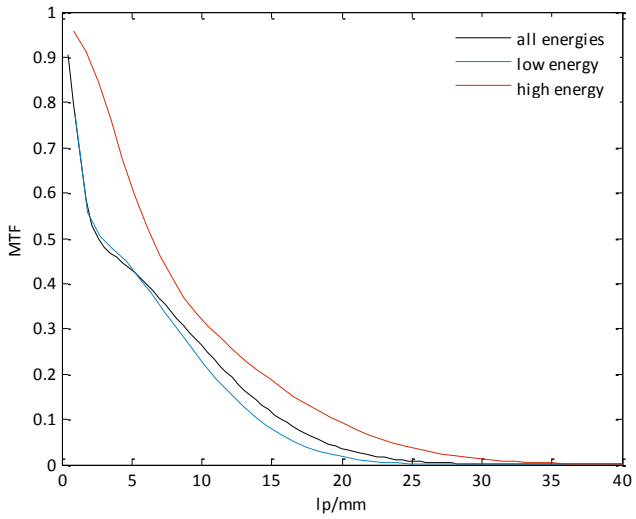

Figure 5.37. The modulation transfer functions for an edge image taken using all energy events, low energy events and high energy events. The MTF is improved over all spatial frequencies when using only the high energy events.

The MTF for the high energy data set shows improvements across all spatial frequencies, particularly at low spatial frequencies (where the MTF is dominated by the fluorescence tails on the LSF) and at high spatial frequencies, where the narrower central peak extends the MTF further. The spatial resolution from the FWHM of the LSF is improved from 31 µm (all energies) to 25 µm when using only the higher energy events. Comparisons of these results with previous studies in the literature are discussed in Section 5.12, following a secondary study of the energy discrimination of the device using a ^{57}Co source at 122 keV.

5.11 Energy discrimination for ^{57}Co

Testing was continued using a higher energy gamma-ray source. The ^{57}Co source was chosen due to the approximate doubling of the peak gamma-ray energy at 122 keV. The 122 keV peak will be broadened slightly by the presence of a higher energy peak at 136 keV, but this should

not dominate. The spectrum shown below in Figure 5.38 for the ^{57}Co was generated as for the ^{241}Am results using the event integration technique.

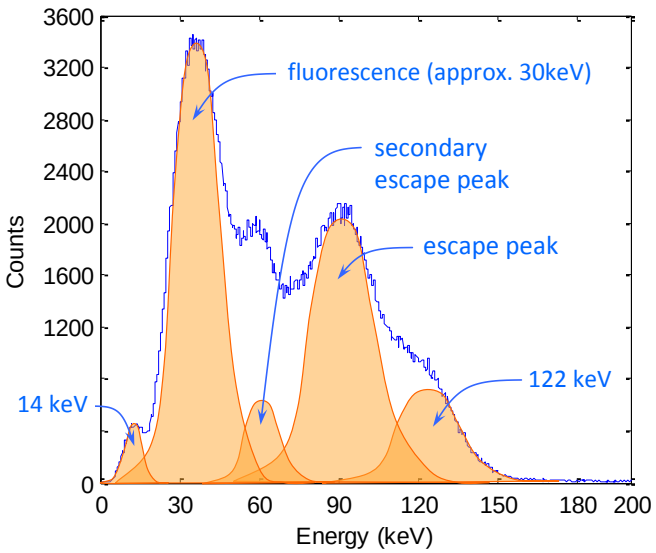

Figure 5.38. The experimental data for the ^{57}Co spectrum (shown in blue). The orange overlay shows the relative dominance (very approximately) of the peaks forming the spectrum.

The ^{57}Co energy spectrum is further complicated over the ^{241}Am spectrum due to the higher energy of the incident gamma-rays. At 59.5 keV, the energy of the photo-electron and fluorescence X-ray is not sufficient to further ionise a K-shell electron and thus the fluorescence occurs over a single stage. With the ^{57}Co source, the photo-electron has sufficient energy to further ionise the caesium and iodine in the scintillator, leading to further fluorescence X-rays and a secondary escape peak. The cartoon overlaid onto the spectrum in Figure 5.38 shows the various components found in the spectrum. The re-absorbed fluorescence at 30 keV is separated from the remaining peaks.

The correlation between the two spectra shown here gives further evidence of the impact of fluorescence in imaging applications. The energy of the incident gamma-ray has a major

impact on the components forming the event profiles which are then analysed for imaging. Further discussion of the implications of these results is given below.

5.12 Discussion

The use of photon-counting methods has been shown to enable energy discrimination when imaging with a gamma-ray source. The high-energy, high-resolution imaging made possible with this device opens the doors for multi-colour X/γ-ray imaging. Further development, however, is required to enable true multi-colour imaging due to the complications in the energy spectra arising from the caesium and iodine fluorescence from the scintillator. A scintillator in which the fluorescence yield is lower may allow increased separation in the energy peaks, allowing multi-label imaging using a user-defined colour scale. The device discussed in the following section explores the possibility of using a scintillating fibre-optic and uses the results from this study to allow a theoretical study of a future detector. Multi-colour imaging would allow the use of multi-label systems – several gamma-emitting isotopes could be used for a single imaging study, allowing the simultaneous imaging of differing systems.

It must be noted that the geometry of the system set-up was not optimal for spatial measurements and was designed for bench-top testing in the laboratory. The results show large improvements over devices discussed in the literature despite the imaging set-up, as shown in Table 5.3. Future testing with a parallel synchrotron beam-line is hoped to improve the measured resolution further still. The testing of the device on a synchrotron beam-line is dependent on a much higher readout rate, discussed further in the following section.

The 'energy resolution' of the detector also shows improvements when compared with previous studies detailed in the literature (Chapter 2). Previous studies give a main peak at the full gamma-ray energy, but leave a continuum at lower energies. The full energy peak is of a higher count rate compared to the continuum, but this is due to the lower resolution – the blurring of event profiles leads to the merging of many fluorescence and escape peaks to

create an increased number of full energy peaks but with a much broader LSF. It is only with the better event separability due to the much higher spatial resolution achieved in this work (through careful calibration, optimisation and simulative analysis), that the individual event profiles can be distinguished and the energy spectrum can be improved to include more near-independent peaks.

Study	Resolution (FWHM)
Hall (2009)	25 µm
Tipinis (2004)	50 µm
de Vree (2004)	60 µm
Miller (2007)	64 µm
Miller (2009)	69 µm
Miller (2006)	78 µm
Heemskerk (2007)	86 µm
Nagarkar (2006)	87 µm

Table 5.3. A comparison of the resolution measurement for this study with previous studies using a similar scintillator-CCD/EM-CCD system (this study shown in blue).

A final comparison is required between the opposing methods used to determine the energy of the event profiles to ensure consistency. The plots shown in Figure 5.39 give a direct comparison between the scale space method (lower right quadrant) and the pixel summation method (upper left). The comparison shows a near-linear response between the two methods (upper right quadrant), although the spread across the events shows the uncertainty of the measurement in both methods. Each method samples the noise in the image in a different

way (both the background noise and the intrinsic shot noise present in the signal and the multiplication gain noise increase), leading to the differing results. As discussed previously, each method has benefits depending on the requirements of the system and the method must be chosen carefully with reference to the processing time and the results required.

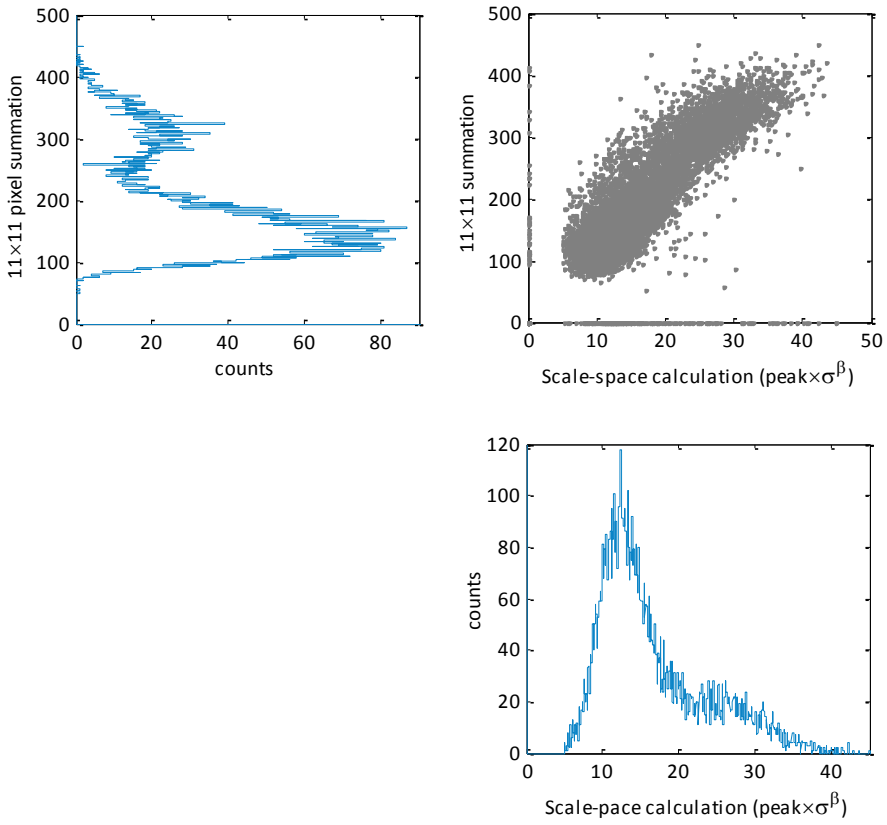

Figure 5.39. A comparison between the energy value generated through the pixel summation and scale-space method for a sub-set of the data. The spectrum for the pixel summation shows greater separation between the energy peaks. The comparison plot gives a near-linear relationship, although the grouping is widely spread.

5.13 Implications for future devices

The results from the analysis detailed in this chapter must be considered against current detectors and applications. The application, in many ways, defines the requirements for the future development of the detector. The remainder of this chapter discusses further the implications of the study and investigates the changes which are now known to be required to further improve the detector's capabilities. This introduces a possible angle for future study: the implementation of a detector designed to overcome, as far as possible, the downsides of current devices with particular regard to those detailed in this study. The results from this work already show marked improvements over previous studies detailed in the literature and, through a thorough understanding of the processes gained by attention to detail in the analysis and through the use of simulative techniques, a device with superior characteristics may be produced for future testing and implementation.

5.13.1 Impact of fluorescence

The choice of scintillator plays a major role in the imaging capabilities of the detector. It was shown previously in the chapter that the fluorescence from the elemental components of the scintillator had the most detrimental effect on the spatial resolution and spectral discrimination capabilities of the device. The re-absorption of the fluorescence X-rays generated inside the scintillator lead to the spreading of the signal across a wide area, resulting in the secondary Gaussian inferred by the LSF. The secondary Gaussian was shown to be reducible through the removal of events of a pre-selected energy, but the removal of such events decreases the overall number of counts forming the image, increasing the flux or time required to generate a suitable image. The energy of the K-shell fluorescence X-rays can also interfere with the escape peak (such as with the ^{241}Am source discussed previously), resulting in difficulties in achieving true energy discrimination.

Although the fluorescence was shown to be the dominant factor decreasing the resolution, the photo-electron interaction lengths were also found to decrease the resolution to a lesser extent. The FWHM of the LSF was found to decrease by several microns when the escape events were removed in the ^{241}Am data.

The choice of scintillator should therefore act to minimise the fluorescence yield at the energies required for imaging. A scintillator with a lower fluorescence yield would decrease the impact of the fluorescence on the spatial and spectral resolutions achievable with the detector.

The detection of the 14 keV gamma-rays in the ^{57}Co results shows that the number of visible photons required to be resolvable above the background is much lower than first assumed. The 14 keV gamma-rays produce less than a quarter of the number of photons in the scintillator than those at 60 keV and as such, a scintillator with a lower scintillation yield may be chosen if the scintillator in question is more favourable in other areas. The variable gain achievable with the EM-CCD allows low signals to be amplified, allowing much lower signals to be measured. Consequently, the background must be reduced further, but this is possible through other methods (see below).

It is possible that previous studies have not encountered the same problems regarding fluorescence due to the poorer resolutions achieved. If the resolution of the detector is lower, the fluorescence events in the image may combine with the primary events, lowering the impact of the fluorescence on the image compared to the ideal case where no fluorescence is present. At the higher resolutions achieved here, the impact of the fluorescence is noticeably greater. With a spatial resolution of 25 µm FWHM the impact of the fluorescence has been shown to be very important, and further improvements in the spatial resolution require this problem to be tackled. With an improved test-geometry, decreasing the impact of gamma-ray

divergence from the source to detector, it is predicted that the fluorescence will have a greater impact still.

5.13.2 Proposed detector

The detector proposed below is designed to take note of the results detailed in this study in order to develop an optimised detector. Through the greater understanding of the multiple components of the LSF and MTF, the detector can be optimised to provide a detector capable of a much higher resolution and greater energy separation.

5.13.2.1 CCD drive system

The drive electronics for the detector used in this study have a maximum readout rate of 2 MHz, giving a maximum full frame readout rate (558 × 528 pixels for the CCD97) of 6.7 frames per second. The readout rate can be increased dramatically through the use of an alternative drive system. A high-speed drive system is currently being developed at the Institute of Astronomy at the University of Cambridge, Figure 5.40. At the time of writing, the development team has a prototype camera capable of taking pictures of "reasonable quality" at pixel rates in excess of 20 MHz ('Lucky' website). Once the drive system is made available, the scintillator-couple EM-CCD may be tested with the faster drive electronics, allowing the development of the camera and the implementation of the testing outlined here. A frame rate of 20 MHz using a CCD97 at 558 × 528 pixels results in a full frame readout rate of over 60 frames per second.

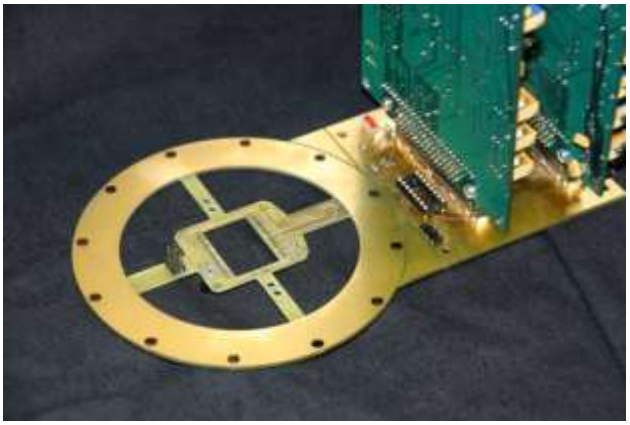

Figure 5.40. The CCD is mounted on the fibreglass support (left), located only a few centimetres from the two control boards (right).

The camera described in Chapter 4 and previously in this chapter is based around an operation at 200 kHz in order to achieve the best-case results with the current drive system. The use of a 20 MHz drive system (with a relevant increase in gain to overcome any readout noise increase) would act to decrease the dark current integration period by a factor of one hundred. Operation at 20 MHz would offer the same level of dark current (electrons/pixel/frame) at approximately 5°C compared with that at 200 kHz at -25°C (CCD97 datasheet, 2004). The increase in temperature made possible with the higher frame-rate may therefore negate the need for the vacuum system. A reduction to 5°C is theoretically possible through the use of a TEC without the requirement for water cooling or a vacuum system. Alternatively, the dark current may be reduced by a further factor of 100 (reducing the noise on the dark current by a factor of 10) through the use of the same cooling system as used previously.

5.13.2.2 Scintillator

One of the major problems in using caesium iodide is the high fluorescence yield of approximately 90 %. With a lower fluorescence yield, gadolinium oxysulphide (Gd_2O_2S), 'Gadox', a commonly used scintillator, would be an ideal scintillator for this purpose.

Unfortunately, Gadox is currently only manufacturable in single blocks – collimation has proved particularly difficult. Attempts have been made to 'force' a solution of Gadox into micro-channel plates but with little success. With only a solid sheet/block of scintillator the spread in the light signal will have detrimental effects on the imaging capabilities of the detector.

An alternative to the scintillator-coupled fibre-optics discussed thus far may be investigated. A scintillating fibre-optic, in which the fibre-optic itself contains the scintillator, is one such possibility. The scintillating fibre-optic is formed like a standard fibre optic plate but from fibres of scintillating glass surrounded by a cladding material (often borosilicate glass). Some of the emitted light will undergo total internal reflection and this light will pass to the ends of the fibres. Those photons which pass through the cladding will continue to pass through the scintillator, spreading the event profiles. Many scintillating fibre-optics therefore make use of extramural absorption (EMA) fibres which are placed at regular intervals in place of the scintillating fibres in order to absorb some of the stray photons.

The reduction in the spreading of the photons from the fibre in which the initial interaction occurs would provide a much tighter event profile than for the collimated caesium iodide scintillators where the columns are far from perfect. Reflective coatings can be applied to the upper surface of the plate in order to reflect back any photons passing 'upwards' towards the detector, increasing the number of photons detected by up to a factor of two (results from the manufacturer show an increase of approximately 60 % in the signal with a mirror on the input). The decay time of the scintillating fibre-optic tends to be longer than that found with caesium iodide, with a decay to 40 % at 2 ms typical of plates provided by Collimated Holes (R. Mead, private communication). At frame rates of 60 frames per second a standard frame integration time stands at 16 ms, minimising the impact of the longer decay time and allowing the majority of the signal to be integrated in a single frame.

The current generation of scintillating fibre-optics from Collimated Holes Inc. produce 7-8 'green photons' per keV (R. Mead, private communication). At 60 keV this equates to approximately 450 photons per interaction (almost ten times lower than the output from caesium iodide). Half of all photons produced will pass away from the detector and will be lost if no reflective coating is used. The fibre structure of the scintillator will channel a certain percentage of photons towards the CCD with these photons remaining in the fibre in which the interaction took place (the interaction fibre). The fibres are formed from columns of LKH-6 glass approximately 10 µm in diameter surrounded by a borosilicate cladding. The refractive indices for the LKH-6 glass and cladding are 1.58 and 1.49 respectively. This structure leads to approximately 12 photons in the interaction fibre at the detector-fibre interface at 60 keV. With a reflective coating on the input side, results from the manufacturer show an increase in recorded signal of approximately 60 %, relating to an interaction fibre peak signal of 19 photons at 60 keV.

The peak signal in the central pixel of each event will be slightly higher than the 12 or 19 photons quoted above due to the mismatch in pixel-fibre size. The 16 × 16 µm pixels will collect signal from the surrounding fibres to increase the signal by a small fraction. The low signal level (compared with caesium iodide) means that one must ensure that the interface between the device and scintillating fibre-optic are not causing a loss in signal. In this case, direct coupling between the EM-CCD and the scintillating fibre-optic is preferable. The perfect collimation found in the scintillating fibre-optic brings benefits over the rougher collimation in the caesium iodide, despite the lower signal produced per interaction

The 'perfect' formation of the columns in comparison to the CsI(Tl) would be expected to enable imaging at a much higher resolution, with the minimisation of surfaces at which photons can be lost/spread. The energy resolution, however, would be reduced through the decrease in the signal-to-noise caused by the reduction in the number of photons emitted. Although the shot noise on the lower signal would have a detrimental effect on the energy

discrimination, the better collimation would reduce scatter. No fluorescence yield data are available from the manufacturer at present.

5.13.2.3 Detector system

The lower emission intensity found with the scintillating fibre-optic requires that the cooling system may still be required (even though the detector will run with lower dark current due to the higher frame-rate). The fluctuations in the dark current must be minimised in order to not only see the peak signal but also to see the signal in the surrounding pixels necessary for centroiding. If, however, one was to continue with the CsI(Tl) detector tested in this study, the use of the faster drive electronics may permit the operation of the detector sufficiently above 0 °C to negate the need for the vacuum system and water cooling for the TEC. The reduction in size and operation technicalities with such a system would give great benefits to the use of the system in situ.

The two detector possibilities for the future of the scintillator-coupled EM-CCD detector revolve around the scintillating fibre-optic (with cooling and vacuum system) and the CsI(Tl) system (with no vacuum system or cooling beyond the TEC). The caesium iodide system is shown in Figure 5.41, where the size of the detector (excluding the drive system) is minimal. The scintillating fibre-optic system is shown also, where the detector requires a vacuum sealed enclosure (with pump) and a water cooling system to remove the excess heat from the TEC.

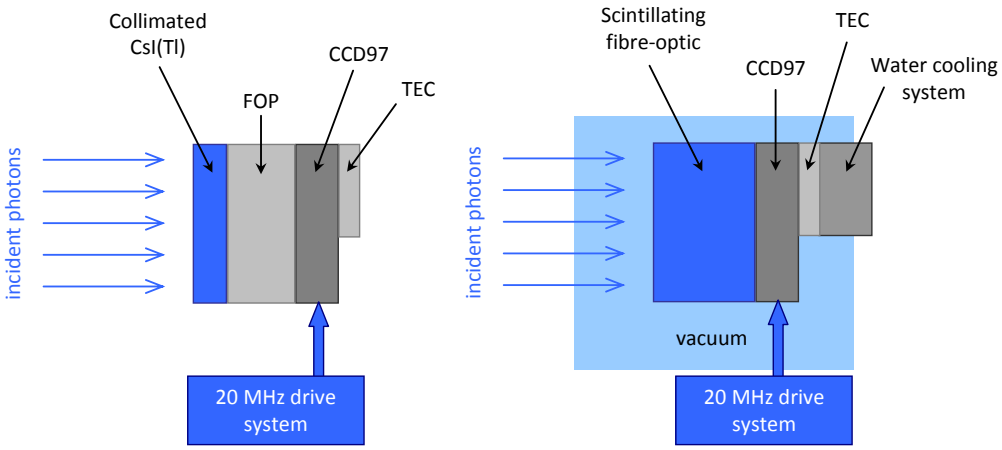

Figure 5.41. The proposed system if the collimated CsI(Tl) scintillator is used (left). The system is similar to that described previously, except for the exclusion of the water cooling and vacuum chamber which is no longer required due to the faster drive system. The proposed system to make use of the scintillating fibre-optic is also shown (right). The cooling system and vacuum chamber are still required due to the lower scintillation yield.

5.13.3 Future testing

Subject to the future availability of the high-speed drive system and time allocated at a suitable synchrotron facility, further testing will be carried out to further the development of the next generation of scintillator-coupled EM-CCD detectors.

The study thus far has concentrated on a lab-based testing system. Although this system allows easy adaption and geometry changes within limits, the system does not allow the optimal resolution to be truly recorded. A parallel beam of X-rays/gamma-rays is required to get the best-case results. Two beam types are available – the narrow pencil beam and a larger beam of uniform flux (see Section 5.2 for MTF details). This experimental work is possible at many facilities, including the Diamond synchrotron in Oxford, UK, providing easy access. A new beam line (I12) is, at the time of writing, due to start testing over the coming months and is scheduled to come into general use mid-2010.

5.13.4 Potential applications

The expected performance of the new and improved detector can be compared to the specifications required in practice. The maximum flux which the detector is capable of imaging in photon counting mode is dominated by the limits on the frame rate at which the device can be operated. The faster drive system described earlier is, at present, being optimised, but taking the prototype described in the literature 60 fps is currently possible. For photon-counting with energy discrimination, an event must be alone centrally in an area of 11 × 11 pixels (the 'event area'). Using two-dimensional Poisson statistics, the probability of two events coinciding in the same area within a single frame is given by Equation 5.2.

$$P\left[N(area)=k\right]=\frac{\left(\lambda|area|\right)^{k}\exp\left(-\lambda|area|\right)}{k!} \tag{5.2}$$

This predicts that the expected number of events, λ, must be approximately 0.12 or less per event area for a less than 10% probability that two events will occur within the same region (events separated by 11 rows and columns). Over the image area of the CCD97 (1 cm^2) with 60 fps, this gives that the maximum flux detected in the CCD97 must be lower than 140,000 events per cm^2 per second. Taking into account the absorption of 140 keV gamma-rays (primary emission from Tc^{99m}) in 5 mm of scintillating fibre-optic and the 90 % QE of the CCD97 at the emission wavelengths, the maximum incident flux is higher.

Taking the medical application of sentinel node imaging as an example, the device compares very favourably with the required operating conditions. Taking sample data from Tsuchimochi (2006), the gamma-ray flux at the detector in standard sentinel node imaging can be calculated. Taking an injected dose rate of 10 MBq of Tc^{99m} implies that 9 million 140 keV gamma-rays will be emitted per second. Surgeons suggest that only 0.15% of the injected dose is absorbed per lymph node, and that these lymph nodes can be found approximately 3 cm below the skin (effectively below 3 cm of water). With the attenuation of the water taken into account, along with the absorption of the injected dose by the lymph nodes, this

can be approximated to an effective source with a flux of 8,600 gamma-rays per second at 140 keV. This leads to approximately 250 events per second (assuming radial emission) at the entrance of the fibre-optic plate. Injected doses vary depending on the patient and the time between injection and imaging, with doses up to 100 MBq often injected. These higher doses would provide 2,500 events per second at the entrance to the fibre-optic. These values fall within the maximum capabilities of the device, allowing imaging at higher doses, imaging closer to the skin surface and imaging over shorter time periods than required in current systems.

The dominating limit on the imaging resolution of the detector in this case is the collimator. The collimator is required to avoid geometrical uncertainty in the origin of the detected gamma-rays and further developments in collimators are required to make optimal use of the high detector resolution in standard medical imaging. Imaging in synchrotron research (e.g. crystal diffraction) is not dependent on a solid source (or injected source solution) and hence does not require the same level of collimation.

High-resolution hard X-ray imaging for synchrotron research (or that for use in synchrotron beam-lines) does not require collimation and hence the resolution is dominated by the device itself. Fluorescence X-rays may be isolated from the main peak energy and as such, the original targeted gamma-rays can be imaged independently (within the limitations of the energy discrimination of the system). The high resolution and high frame rate may open the doors to new imaging applications and will act to improve the quality of results from current imaging arrangements.

A final consideration is the imaging area of the detector. The image region of the CCD97 measures only 1 cm^2. This is smaller than many current imagers and is seen as the main limitation to the detector. The compact nature of the device is not limited to the image region however, as the device as a whole is much smaller than the systems used in many hospitals.

The image region of the CCD97 can be increased through the use of a fibre-optic taper, readily available with a magnification of 3.1:1. The increase in image area comes at a cost: signal photons can be lost in the taper, and as such, this method is more suitable for the caesium iodide case where the signal levels are much higher. Detailed studies of the use of the fibre-optic taper would be required to ascertain the degradation in resolution obtained with the taper in place.

A second method available to increase the image area of the detector involves an upgrade to the EM-CCD used. The latest EM-CCD from e2v technologies, the CCD201-20 (CCD201-20 datasheet, 2005), is formed from 1024 × 1024 pixels in the image area. The total store section can be read out at approximately 20 fps using a single output node at a 20 MHz pixel readout rate. Although the number of pixels is increased by a factor of two in each direction, the imaging area of the device is not quadrupled due to the smaller pixel size. The CCD201-20 pixels measure 13 µm square, giving a total image area of 13.3 mm × 13.3 mm. This gives an increase in area by a factor of 2.64 compared to the CCD97 (8.192 mm × 8.192 mm image area). Using a 3.1:1 fibre-optic, the maximum image area achievable is therefore approximately 16 cm^2. This gives a large increase over the CCD97 imaged alone at 0.67 cm^2. Using the two individual output nodes of the CCD201-20 it may be possible to operate the device at higher frame-rates, giving a large area, fast-readout detector suitable for implementation into medical and security imaging and for synchrotron research.

5.14 Conclusions

The simulation work carried out in this study, backed up by the promising experimental results, shows the impact the X-ray fluorescence of the scintillator itself has on the imaging and spectral capabilities of the detector. The higher the resolution achievable with the basic detector, the greater the impact of the fluorescence, hence the more pronounced spectral results from this study (resolution of 25 µm) compared to previous studies of lower resolutions

(FWHM greater than 60 µm). To further increase the resolution the issues with the fluorescence must be tackled, a subject discussed briefly in this chapter with regards to the use of a scintillator with a lower effective atomic number. The imperfect collimation of the caesium iodide causes further scattering and hence the use of the scintillating fibre-optic in future devices may improve the resolution further still.

The limitations on the readout rate caused by the drive system may be overcome following the release of the faster drive electronics discussed earlier. The higher frame-rate will allow the use of higher flux sources and the lower dark current per pixel per readout will decrease the noise on the dark current and hence reduce the requirements of the cooling system.

The planned testing of the device has been outlined, and future results will aim to improve the resolution further still, with the lower noise levels improving the energy discrimination and aiding the removal of the re-absorbed fluorescence. It is hoped that the future testing of the device described above will be carried out over the coming months, subject to the availability of the faster drive electronics and time allocated at a suitable synchrotron facility, with the aim of producing a commercially available high-resolution hard X-ray imager to give dramatic improvements over currently available gamma-ray imagers.

Chapter 6: Instrument background in space borne X-ray spectrometers

6.1 Introduction

Secondary event detection in CCD cameras is not limited to terrestrial applications. Charge-Coupled Devices (CCDs) are often utilised for observing X-rays in space. The instrument background experienced in-orbit is caused by the interaction of particles with the detector and surroundings and the subsequent detection of non-X-ray events, a major component of the total background, having a large impact on the overall sensitivity of the camera system. Results from the European Space Agency (ESA) X-ray Multi Mirror (XMM-Newton) mission and the Japanese Space Agency and NASA Suzaku mission can be seen to show a variation in the instrument background with differing orbit and detector design (see the end of this chapter), yet a detailed study on the causes of this instrument background has previously not been undertaken. The instrument background for the NASA Swift mission has been studied in less detail but leads to an interesting comparison with the XMM-Newton results. Chapters 6 and 7 describe work carried out to gain a better understanding of the instrument background experienced in X-ray CCD cameras in space using computer simulations. A Monte Carlo simulation has been used, implemented with the Geant4 toolkit described previously in Chapter 3, and the methodology and data analysis is detailed. The somewhat surprising results follow, and their implications on possible future space missions such as the joint NASA, ESA and JAXA International X-ray Observatory (IXO) mission (ESA IXO website) are briefly discussed.[4]

[4] The work described in this chapter has been published in the papers Hall *et al.* (2007), Hall *et al.* (2008) and Hall *et al.* (2010) at different stages, with some figures shared between this chapter and the publications detailed (all by the thesis author).

6.2 The purpose of simulation

In an ideal world, one would be able to accurately predict the instrument background pre-launch for any X-ray observatory in space. In the past the instrument background for many missions has been found to exceed that predicted before launch. The instrument background for the XMM-Newton mission when in-orbit was found to be nearly two orders of magnitude higher than was predicted pre-launch.

By using a Monte Carlo based method it is possible to define a geometrical and physical model of the system which can then be subjected to a suitable spectrum of incident particles. The geometrical model is defined by the detectors and surrounding craft used in the mission. The physical model should be consistent throughout all missions with regards to the interactions of particles with matter. These two factors can be replicated through the use of the Geant4 toolkit and associated packages. The difficulties thus far are only those involved with the decisions on how to best approximate the system to optimise the implementation and running time of the simulation without adversely affecting the accuracy of the results.

Whereas it is simply a matter of cost to replicate the detector and physical surroundings for ground based testing, it is very difficult, if not impossible, to accurately replicate the spectrum of particles incident on the detector as would be found in space. One of the main problems encountered when attempting to simulate the instrument background experienced in-orbit is the complexity of the spectrum of the incident particles. These cosmic-ray particles vary dramatically with the orbit of the mission. The spectrum of cosmic-rays incident on any space mission can be predicted, but it has to be noted that the models available are themselves not fully accurate or well understood.

This chapter contains details on the theory behind instrument background, detailing the processes involved in the formation of instrument background and the influence the orbit and structure of the detectors has on the particles incident on the detector. A discussion follows,

describing the basic details of the mission, followed by a more in-depth look at the detectors themselves. A brief comparison of the in-orbit instrument background is given before the detailed analysis is provided in Chapter 7.

6.3 What is instrument background?

When examining a source (e.g. a star) within a small spatial region of interest (e.g. a region of sky) it is possible to remove the background to a certain degree by subtracting the 'local' background. This background is taken from a region close to the source and can give good results if used appropriately. If, however, a larger or extended source is being examined, it is not possible to remove a local background as the locality is not specifically defined. In this case it may be necessary to remove the background using a region which may not be entirely 'local'.

One source of background originates from instrumental fluorescence. This occurs when high energy particles and X-rays/gamma-rays pass through the detector giving rise to additional background from interactions in the detector and shielding structure. This X-ray fluorescence is dependent on the materials used in the surrounding structures and is generally well understood (e.g. Freyberg *et al.*, 2002).

XMM-Newton, designed for the detection of "low surface brightness emission features from extended and diffuse galactic and extra-galactic sources" (ESA XMM-Newton website), was found to have a much higher background than was predicted pre-launch. This is a scenario which it is hoped will not be repeated with future missions. The under-prediction of the instrument background for the XMM-Newton mission is discussed further in Sections 6.9 and 7.7.

The X-ray like background continuum of a CCD in space is composed of two main components: the unresolved Cosmic X-ray Background (CXB) and the instrument background. The CXB

dominates below approximately 2-3 keV, with the instrument background dominating at higher energies, as shown in Figure 6.1.

The CXB is approximately three times larger than the instrument background at 1 keV, falling to approximately 0.4 times at 5 keV (data taken from De Luca *et al.*, 2001). The simulation described here models the system from 1-12 keV, comparing the simulated data to the instrument background only, where it has been possible to remove the CXB through the use of closed observations. The CXB consists of un-resolved X-rays which have passed through the detector aperture. With a closed filter these X-rays are blocked but it is not possible to gain data from the sources of interest whilst the filter remains closed.

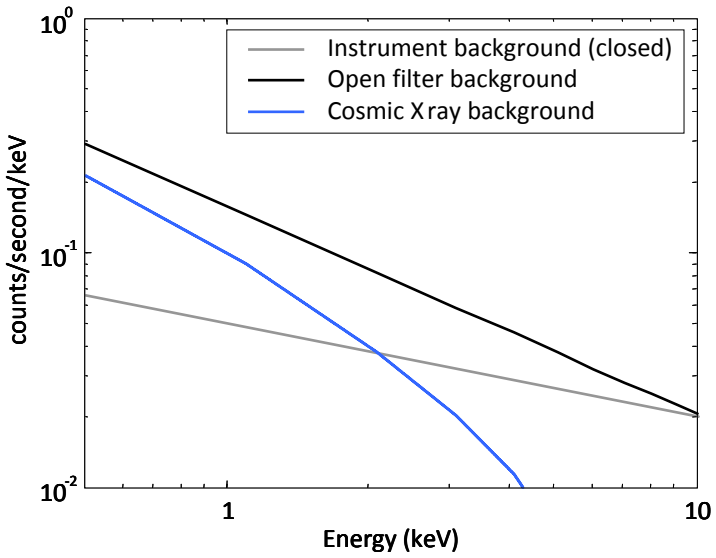

Figure 6.1. The Instrument background is shown from the closed filter background. The total background is taken as the open filter background. The Cosmic X-ray Background (CXB) dominates below 2-3 keV. At higher energies the instrument background dominates. Original data taken from De Luca *et al.* (2001) for the MOS1 X-ray CCD camera aboard XMM-Newton.

The second component, that which dominates at higher energies, is the instrument background, or the "quiescent non X-ray background" (De Luca *et al.*, 2001). This is the component of the background that is simulated here. The instrument background is particle-induced and is mostly due to the interaction of high-energy particles with the detector surroundings (particles of the order of a few MeV and higher). Some of the background is therefore due to secondaries from the incident particles and some due to un-rejected minimally ionising charged particles. The instrument background is the component of the background which is most easily reduced as it is dependant on the most easily variable factors such as the design of the surroundings to the detector and the camera shielding materials. Some discussion has been made on using multi-layered shielding (such as Pfeffermann *et al.*, 2004), yet it is believed that only with a better understanding of the real causes can one aim to reduce the instrument background.

6.4 What causes instrument background?

There are many causes of instrument background. The most dominant of these are described briefly below.

6.4.1 The photo-electric effect

When a photon hits a metallic substance, electrons can be emitted through the photo-electric effect, provided the photon is of sufficient energy. Equation 6.1 shows that the kinetic energy of the emitted electron, E_{KE}, is dependant on both the incident photon energy, hf (where h is Planck's constant and f is the frequency of the incident photon), and the work-function, W.

$$E_{KE} = hf - W \tag{6.1}$$

The work-function is the minimum energy required to release an electron from the metal, and as such, the photo-electric effect may only take place if the energy of the incident photon surpasses the work-function. At much higher energies, the Compton Effect may take place.

6.4.2 The Compton effect

When a high energy photon collides with a target, loosely bound electrons can be released (Figure 6.2). The photon loses energy related to the energy of the ejected electron (equation 6.2, where θ is the scattering angle, m_0 is the electron mass, c is the speed of light and h is Planck's constant).

$$\lambda_{initial} - \lambda_{final} = \Delta\lambda = \frac{h}{m_0 c}(1 - \cos\theta) \tag{6.2}$$

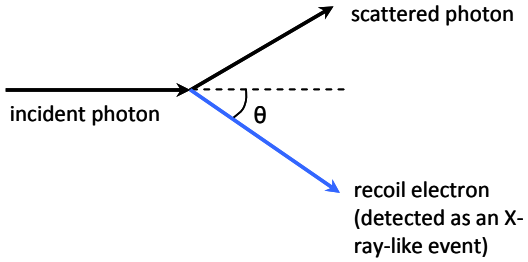

Figure 6.2. The Compton effect**:** The recoil electron can be detected in the CCD as an X-ray-like event

If the Compton effect takes place in the silicon of the detector itself then these electrons can be detected as X-ray-like events, contributing to the instrument background. This component depends on both the incoming flux of photons and the thickness of the detector, as the thicker the detector is the greater the chance there is of Compton scattering.

6.4.3 Minimally ionising charged particles

When a charged particle passes through a material, such as silicon, the particle will lose energy through interactions with atomic electrons (ionisation). The energy lost is dependent on the velocity (or energy) of the particle and the material in which it is interacting. The energy loss for the particles in question here is mostly insignificant compared to the total particle energy

over the distances that may be traversed in a CCD and as such the energy loss can be assumed to be approximately constant.

If the particle travels across pixel boundaries in such a way as to deposit signal over a pattern that does not appear X-ray-like (Section 6.7) then this signal will be rejected, but those particles which pass through the CCD such that energy is deposited in an X-ray-like arrangement of pixels may be indistinguishable from the real X-ray events.

6.4.4 Secondary particles

It is not only the interaction of primaries with the detectors which must be taken into consideration when investigating the instrument background. Primary particles may interact in the material surrounding the detectors and any secondary particles (or those of higher order) may be detected by the devices. As an example, electrons knocked by ionisation from the material surrounding the devices will be absorbed after traversing very small distances through the shielding material, but those near the surface may be able to pass from the shielding material and travel on to the devices themselves.

6.5 Influence of the orbit

The orbit at which the detectors reside has a large impact on the instrument background observed. This is as one would expect, as the orbit plays a major role in determining the flux of incident particles. At Low Earth Orbit (LEO) the spectrum of incident particles is very different to that at Highly Elliptical Orbits (HEO), and as such, one would expect this to play a major part in determining the instrument background observed. The flux of incident protons was found here to be of particular importance due to the secondary particles produced. The height of the orbit is not the only factor to consider, as the angle of incidence of the orbit plays a major role on the passage through the South Atlantic Anomaly (SAA).

6.6 The South Atlantic Anomaly (SAA)

The Van Allen radiation belts are held in place by the earth's magnetic field and consist of a torus of charged particles (Barth, 1999). These belts make their closest approach to the Earth's surface in a region over the South Atlantic due to the tilt in the rotational axis of the Earth, Figure 6.3. In this region, named the South Atlantic Anomaly, the radiation environment has the highest intensity that will be experienced. To take this into account, the instrument background spectra are taken over many orbits such that the total average flux over this time can be considered. The SPENVIS calculations used to generate the incident particle spectra for the simulation are based on the orbital parameters of the missions and are likewise averaged over many orbits (Heynderickx *et al.*, 2000).

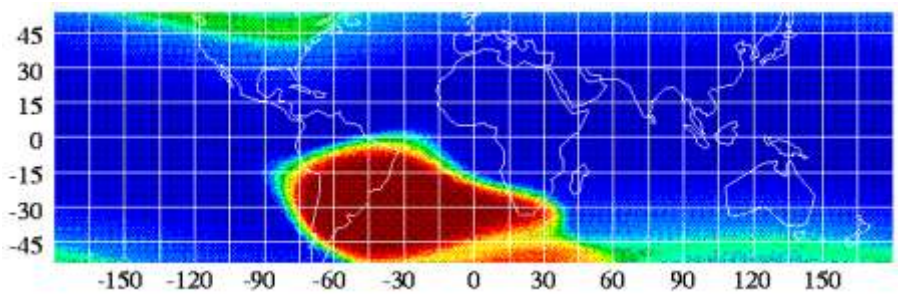

Figure 6.3. The dimensions of the South Atlantic Anomaly at approximately 560 km (NASA website, SAA).

6.7 Influence of the surrounding structure

The structure surrounding the camera is often designed to provide shielding to the detector but poor choice of material and design can have major disadvantages. X-ray fluorescence from the material surrounding the detector can lead to unnecessary and unwanted peaks in the spectrum. A good example of this can be seen with the pn-CCDs (Freyberg *et al.*, 2002) on board the XMM-Newton mission (Section 6.8.1), where large copper fluorescence peaks were

observed due to the printed circuit boards placed directly beneath the detectors (Freyberg *et al.*, 2002). The secondary particles produced in the materials surrounding the detector can lead to high instrument background levels even though the components of instrument background resulting from the primary particles may actually have been reduced. Careful consideration of these factors is important in the design of the camera system.

6.8 The missions studied

6.8.1 The ESA X-ray Multi Mirror (XMM-Newton) Mission

The ESA X-ray Multi Mirror (XMM-Newton) mission, the biggest scientific satellite ever built in Europe, was launched on the 10th December 1999. The mission has examined X-rays from accretion onto black holes, explored the properties of exploding stars, the nature of exotic matter and made many observations of gamma ray bursts (ESA XMM-Newton website).

The satellite was launched with an operational Highly Elliptical Orbit (HEO) with a perigee of approximately 7 000 km and an apogee of approximately 120 000 km (Table 6.1). With this orbit it was expected to experience “close to interplanetary galactic cosmic ray spectra at solar minimum conditions” (Lumb *et al.*, 2002). These cosmic ray spectra will be explored in more detail in Chapter 7.

Parameter	Current values	Values at launch
Apogee	108 898 km	114 000 km
Perigee	24 985 km	7 000 km
Inclination	52.608°	40°
Launch Date	Dec 10th 1999	
Orbit length (hours)	48 (approx.)	

Table 6.1. Orbital parameters for the XMM-Newton mission (ESA XMM-Newton website).

The XMM-Newton spacecraft contains, amongst other apparatus, the European Photon Imaging Camera (EPIC), consisting of two Metal Oxide Semi-conductor (MOS) CCD cameras, Figure 6.4, and one pn-CCD, Figure 6.5. The pn-CCD is subjected to an unobstructed beam of X-rays in normal operation, whilst the MOS cameras are placed behind the X-ray telescopes where the gratings of the Reflection Grating Spectrometer (RGS) divert less than 50% of the incident flux to the CCDs (Read *et al.*, 2003). This does not affect the instrument background as such, with the instrument background mainly due to particles external from the main beam interacting with the detector surroundings and the CCD itself, but does affect the overall sensitivity of the camera (Section 7.8). The Cosmic X-ray background, however, will be affected.

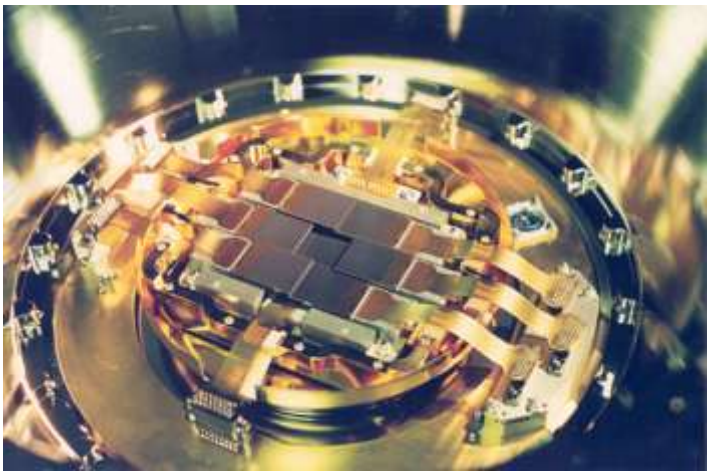

Figure 6.4. The MOS1 array onboard XMM-Newton comprising of the seven e2v technologies type 22 front-illuminated (FI) CCDs. The outer six CCDs are stepped upwards to more closely match the curvature of the focal plane (ESA XMM-Newton website). The imaging sections of the CCDs (dark grey in the picture) measure approximately 2.4 cm across.

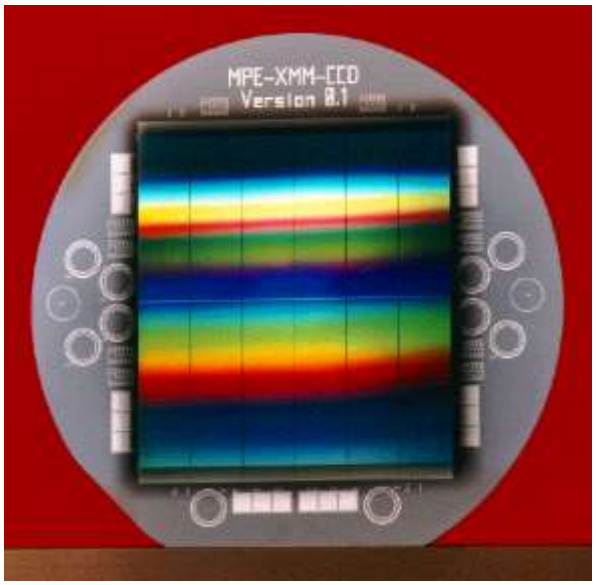

Figure 6.5. The 12 twelve pn-CCDs, fabricated on a single wafer, forming the pn-CCD camera focal plane array aboard XMM-Newton (ESA XMM-Newton website). Each of the 12 pn-CCDs measure approximately 3 cm × 1cm.

Each MOS camera consists of seven e2v technologies type 22 front-illuminated (FI) CCDs. The central CCD is set at the focal point with the outer six CCDs stepped upwards by 4.5mm in order to better follow the curvature of the focal plane. Each CCD image section consists of 600 × 600 pixels of 40 µm width. The energy range of the device (with suitably high quantum efficiency) is from approximately 0.2 keV to 10-12 keV, and this is therefore the range of energy simulated. In order to increase the quantum efficiency for X-rays below 1 keV the pixel structure is partially etched away to leave 'open' 40 % of the structure. Although the CCDs have an epitaxial layer thickness of 80 µm, the mean depleted region of the CCDs is approximately 30-40 µm (Read *et al.*, 2003).

The pn-CCD camera consists of twelve pn-CCDs on a single wafer, each containing 200 × 64 pixels of width 150 µm. The device uses a fully depleted 280 µm layer of silicon with a useful energy range of 0.15 keV to 15 keV (Read *et al.*, 2003).

To determine which events are to be accepted as X-ray-like, a grading system is used as shown in Figure 6.6. The detected events are usually separated into single pixel events or all X-ray-like events for analysis.

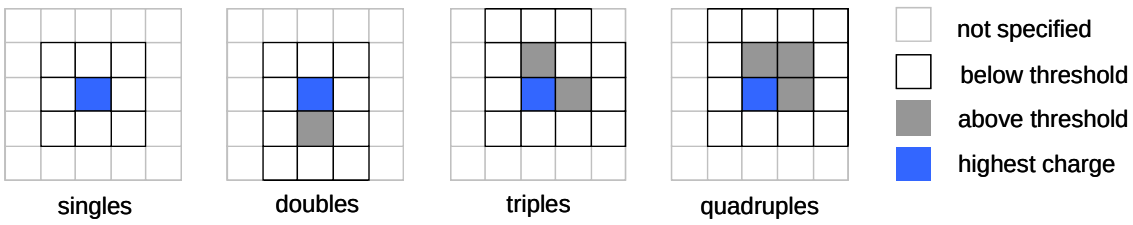

Figure 6.6. The pixel grading used for this study is based on that used for the XMM-Newton mission X-ray CCD cameras. Four-fold rotational symmetry is applied to the double, triple and quadruple pixel events.

For XMM-Newton, the out of field of view background is dominated by the instrument background as the CXB is not present. The out of field of view area of the CCDs corresponds to the overlapped regions of the seven MOS CCDs forming the focal plane and the corners outside of the circular field of view for both the MOS and pn (see Figure 6.7). The MOS in-orbit data, Figure 6.8, was acquired with thanks to Andrew Read at the University of Leicester. The in-orbit pn data, Figure 6.9, taken from Freyberg *et al.* (2004) was calculated as an average over the first three years of the mission, and as such the variation in proton flux (and the influence and direction of the South Atlantic Anomaly, SAA) with the orbital position will be averaged over this period. The MOS instrument background component over any given observation is within approximately 10 % of the mean level (Ness *et al.*, 2009). The continuum has been reproduced and normalised and is shown for ease of comparison in Figure 6.15 at the end of this chapter.

The instrument background for the pn-CCD is approximately three times higher than that for the MOS when comparing like-for-like single pixel events. This was previously thought to be due to the thicker silicon in the pn detector leading to an increased number of Compton scattered electrons (Dyer *et al.*, 1996). Although a logical argument, there has been no previous study to either prove or disprove this hypothesis.

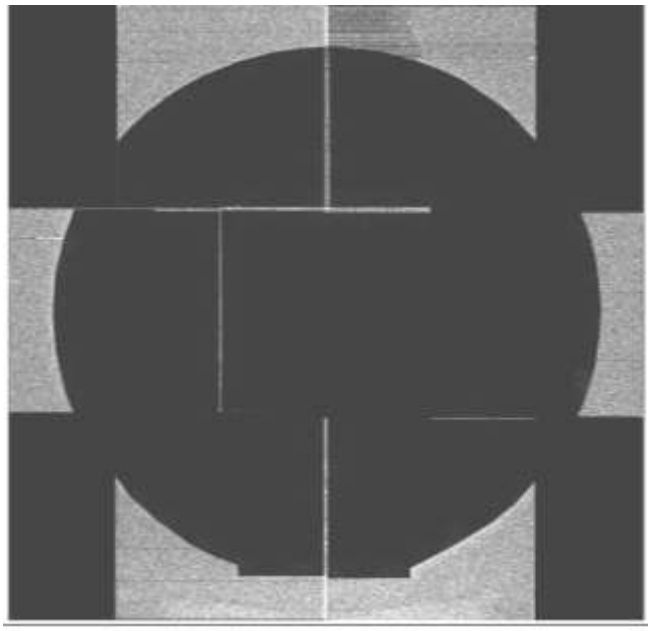

Figure 6.7. The out of field of view region of the MOS focal plane, shown by the lighter grey areas. The outer regions are outside of the circular aperture and the remaining interior light grey areas are obscured regions due to the overlapping of the CCDs. The remaining area (dark grey) at the bottom of the image has been removed due to unexplained anomalous results. Each of the seven 'squares' measures approximately 2.4 cm across.

The data shown here were taken over the first three years of the mission, leaving the variations in proton flux and any influence of the South Atlantic Anomaly (SAA) averaged over many orbital positions. The proton flux during this period of orbit was calculated using SPENVIS (Heynderickx *et al.*, 2000) using the orbital parameters for the mission, Figure 6.10.

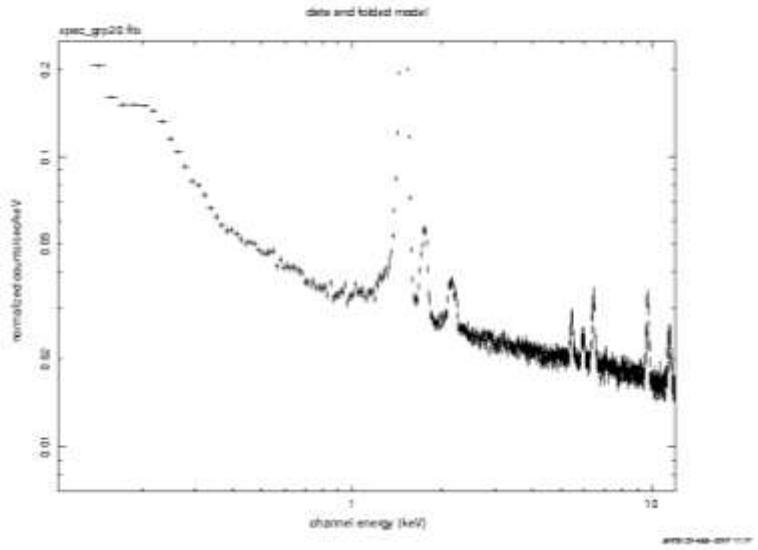

Figure 6.8. The out of field of view spectrum over the region described for all X-ray like events in the MOS camera aboard XMM-Newton. There is a clear continuum to the spectra interrupted by the X-ray fluorescence peaks (spectra provided by Andrew Read, University of Leicester).

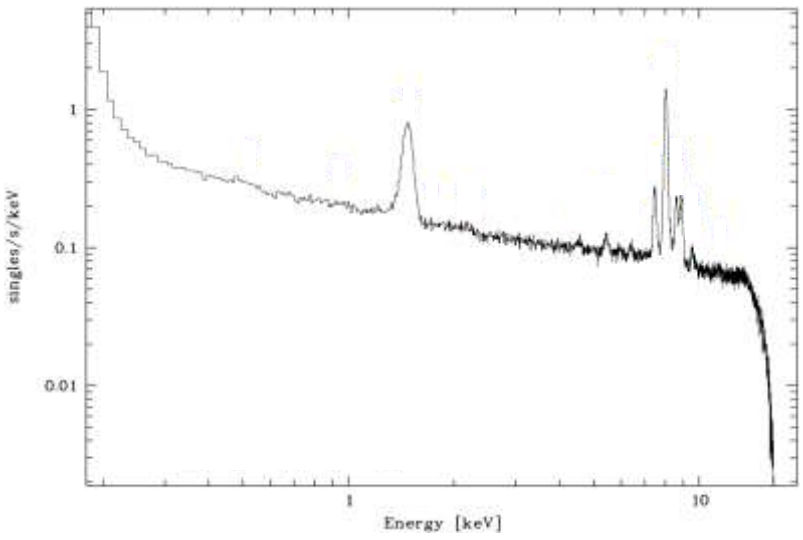

Figure 6.9. The out of field of view spectrum for the pn CCD camera aboard XMM-Newton. This spectrum was edited from Freyberg *et al.* (2004).

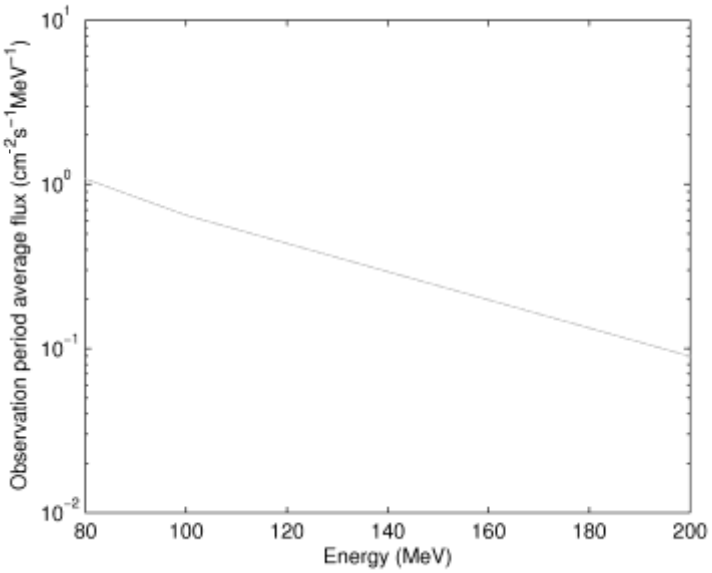

Figure 6.10. The incident proton spectrum for the XMM-Newton simulation including trapped protons and a solar component averaged over the first three years in orbit (produced in SPENVIS).

6.8.2 The Japanese Space Agency and NASA Suzaku mission

The Japanese Space Agency and NASA Astro-E2 mission was launched on the 10th July 2005. The mission was later nick-named Suzaku, meaning “red bird of the south”. The mission has an operational low Earth orbit of approximately 570 km (Table 6.2). The Suzaku satellite was designed to gain a better insight into high energy objects such as neutron stars, active and merging galaxies, black holes and supernovae (NASA Suzaku website).

Parameter	Value
Apogee	568 km
Inclination	31.9°
Launch Date	Jul 10th 2005
Orbit length	96 minutes

Table 6.2. Orbital parameters for the Suzaku mission (NASA Suzaku website).

The Suzaku X-ray Imaging Spectrometers (XISs) contain three front-illuminated (FI) X-ray CCD cameras (XIS0, XIS2 and XIS3) and one back-illuminated (BI) X-ray CCD camera (XIS1), Figure 6.11. The high energy resolution and large effective area of approximately 450 cm^2 at 1.5 keV make these detectors highly suited to this purpose. The X-ray CCD cameras have a stable background at a lower level than similar missions such as XMM-Newton. It is again the continuum component of the background that is of most interest here.

Figure 6.11. The four XIS cameras aboard Suzaku (NASA Suzaku website).

Each CCD measures approximately 2.5 cm across.

The instrument background spectra for the two detector types, the FI-CCD cameras and the BI-CCD camera, are significantly different to both each other and that of the detectors on the XMM-Newton mission. Both the FI-CCD and the BI-CCD have imaging areas of 1024 × 1024 pixels with a pixel size of 24 µm × 24 µm. The depletion layers for the FI-CCD and the BI-CCD are approximately 75 µm and 45 µm respectively.

To determine which events are to be accepted as X-ray like, the ASCA Grade system is used, splitting the events by grade. Grades 0, 2, 3, 4 and 6 are regarded as X-ray like. Figure 4 in Yamaguchi *et al.* (2006) gives a good diagrammatic definition of this grading system. For consistency, the same grading system was employed in the simulation as that used in the XMM-Newton mission, as was shown in Figure 6.6.

The in-orbit background taken with a closed filter is dominated by the instrument background. The data are shown here in Figure 6.12, with the in-orbit data shown for one of the FI-CCD cameras and for the single BI-CCD camera. It is clear from this data that the BI-CCD has a higher level of instrument background than the FI-CCDs. The spectra here are more complex than those in the XMM-Newton detectors but, being subjected to a different incident cosmic-ray spectrum, this is to be expected.

A previous study, Murakami *et al.* (2006), studied the proton flux experienced by the Suzaku XIS detectors in orbit. The incident proton flux was found to relate to 0.6 protons per cm^2 per second. The spectral form of the incident protons for this study was calculated using SPENVIS, as with the XMM-Newton mission, with the results shown in Figure 6.13 for the orbital period considered here. A comparison is also shown with the Swift mission, described below.

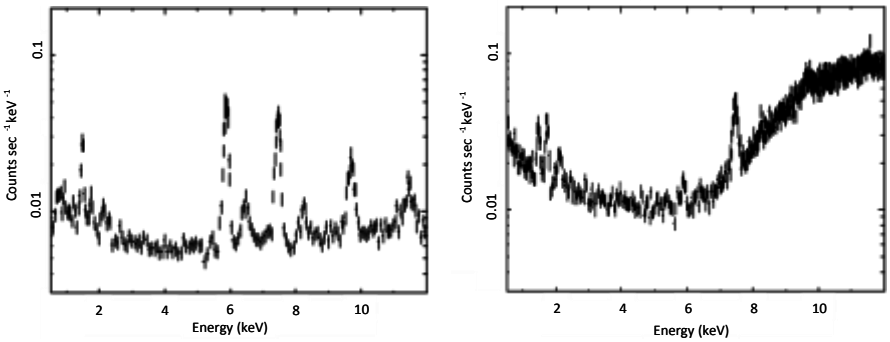

Figure 6.12. The in-orbit closed filter spectra for two of the XIS cameras aboard Suzaku. The left spectrum shows the instrument background for one of the front-illuminated CCD cameras, the right spectrum shows the instrument background for the back-illuminated CCD camera (Murakami *et al.*, 2006).

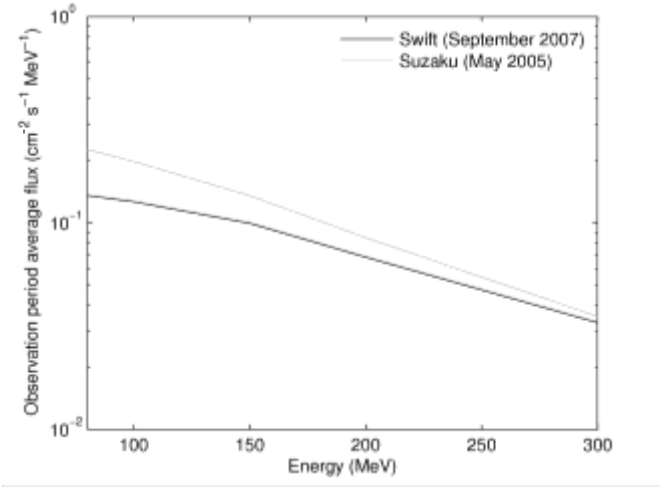

Figure 6.13. A comparison between the incident proton fluxes for the Swift and Suzaku missions, with calculations taken for the two periods over which the spectra were obtained (data obtained using SPENVIS).

6.8.3 The NASA Swift mission

The NASA Swift mission was launched on the 20th November 2004 into a low earth orbit of approximately 600 km, very similar to that of Suzaku (Table 6.3). The Swift mission was designed to investigate Gamma-Ray Bursts (GRBs) and their origins, a section of astrophysics for which there was a great deal of study but about which little was known. One of the three instruments aboard the mission is an X-Ray Telescope (XRT) which is based around the e2v CCD-22, the same device that was designed for the XMM-Newton MOS instruments.

Parameter	Value
Apogee	600 km
Inclination	20.6°
Launch Date	Nov 20th 2004
Orbit length	90 minutes

Table 6.3. Orbital parameters for the Swift mission (NASA Swift website).

Since launch, there has understandably been less work undertaken on the instrument background than for the XMM-Newton mission, launched 5 years earlier. A spectrum of the instrument background was obtained through Claudio Pagani at Penn State University in January 2008 and is shown in Figure 6.14. The dataset was collected in September 2007 when one of the three dual-gyro modules exhibited anomalous behaviour, causing the pointing accuracy to be lost. During this period of observation, the camera was pointed too close to the sun and the shutter was closed. This was beneficial for this work in isolating the instrument background.

The spectrum between 1 and 10 keV, as before, is formed from a relatively flat continuum combined with X-ray fluorescence peaks. A quick comparison between this spectrum and that shown for the MOS1 in Figure 6.15 shows a similar form.

The radiation environment has been considered over the period of observation and the incident proton spectrum is given in the previous section (Figure 6.13), showing the small variation between the Swift and Suzaku proton spectra. Calculations have been implemented through the use of SPENVIS and results show that the Swift data collected over the period in September 2007 is consistent within ± 25 % of the proton flux since the mission launch.

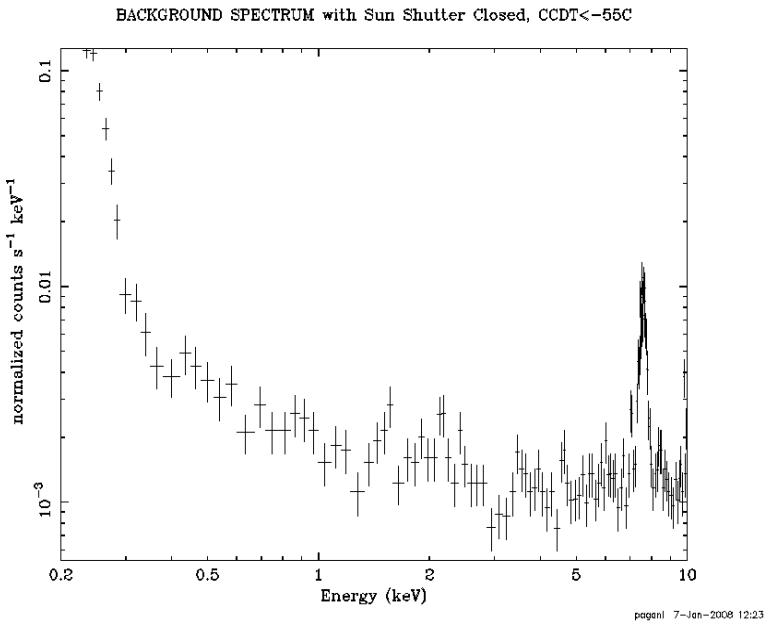

Figure 6.14. The instrument background for the Swift XRT. The dataset was collected in September 2007 when pointing accuracy was lost. With the XRT pointing too close to the Sun the XRT Sun shutter was closed. The shutter was re-opened two days later, during which time the data collected was purely instrumental background (C. Pagani, private communication).

6.8.4 Initial comparison of instrument background

In order to aid the comparison of the instrument background between the missions and the simulated data, the continua of the spectra (figures 6.8, 6.9, 6.12 and 6.14) were reproduced mathematically to allow accurate normalisation. Without normalising the area and time the continua cannot be legitimately compared. The continua for the XMM-Newton cameras were approximated to single exponentials, but the continua for the Suzaku cameras were more complex. For the FI-CCD, the spectra was split into two parts: that above and below 3.5 keV. The BI-CCD was best approximated with a sinusoidal component.

As simulating the fluorescence would lead to poor processing efficiency it is the continuum of the instrument background which has been studied in detail here. The X-ray fluorescence lines depend on the materials used and are generally much better understood in origin than the continuum (e.g. Freyberg *et al.*, 2002).

The data-set used for the Swift instrument background spectrum shown in Figure 6.14 was extracted using a circle of radius 200 pixels. Following the removal of bad columns, the effective area is 121,494 pixels, giving an effective area of 1.944 cm^2. Using this information, an approximated form of the instrument background continuum can again be produced.

A direct comparison of the in-orbit instrument background continua for the three missions and detector types can be seen in Figure 6.15. The XMM-Newton results are taken as the single pixel events whereas the Suzaku results are taken as all X-ray like events due to the availability of the data used. There are two important comparisons to consider: that of the similar devices at different orbits, and that of the differing devices at the same orbit. At the same orbit, the front-illuminated devices have the lower instrument background. For the similar devices, those at LEO have a lower instrument background.

The Swift instrument background can be seen to be approximately three to four times lower than the same detector, the XMM MOS detector, at the higher orbit. The reasons for this are investigated in Chapter 7.

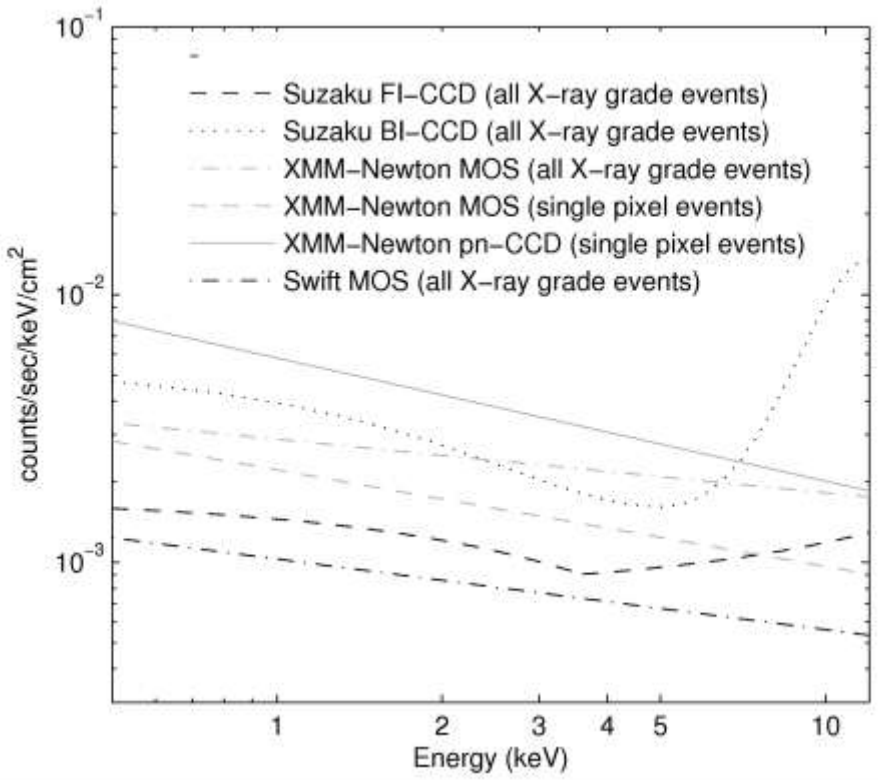

Figure 6.15. The normalised in-orbit spectral continua approximations reproduced in a simplified form for ease of comparison in this study (referenced in text).

6.9 Previous studies

In an ESA-funded study, completed before the launch of XMM-Newton, a simulation was implemented using the Integrated Radiation Transport Suite (IRTS) to predict the instrument background that would be experienced in-orbit by the EPIC MOS CCDs (Dyer *et al.*, 1996).

The pre-launch study concentrated on the Compton-interactions resulting from secondary gamma-rays produced through interactions of cosmic-rays and solar flares with the spacecraft

and detector materials. By integrating the response functions of the detectors according to the cosmic-ray and solar flare induced gamma spectrum, it was thought that the background spectra could be predicted. A full three-dimensional simulation was produced for the EPIC CCDs, calculating a total response of 3.5×10^{-5} counts keV^{-1} per incident proton.

Work by Lumb *et al.* (2002) states a rate of protons passing through the detectors for XMM-Newton of approximately 2.2 protons cm^{-2} s^{-1}, consistent with a rate of 4.4 protons cm^{-2} s^{-1} at solar minimum conditions, reduced by a factor of two due to solar modulation. This incident proton rate lead to a pre-launch prediction for the instrument background of approximately 7.7×10^{-5} counts cm^{-2} s^{-1} keV^{-1} with a flat spectrum. This prediction is almost two orders of magnitude below that observed in-orbit at 1 keV of 3×10^{-3} counts cm^{-2} s^{-1} keV^{-1}. The discrepancies between the pre-launch predictions and the background observed in-orbit are discussed in Section 7.7.

Chapter 7: Simulating the instrument background

7.1 Introduction

A Monte Carlo simulation can only be made reliable through rigorous testing and investigation. Through this chapter, a step-by-step account of the formation of the basis of the simulation is described, including details of the validation of the simulation at each stage. Individual parameters such as pixel size and the more global camera design have been investigated to confirm the source of each component of the background. The results are described with reference to the device parameters and the spectra produced, with further discussion provided on each aspect of the instrument background before finally discussing the impact of the instrument background on the overall sensitivity of the device.[5]

The simulations detailed here are based around a simple framework. The initial raw data are taken from a Monte Carlo simulation written to utilise the abilities of the Geant4 toolkit (first described in Chapter 3). These raw data are then analysed using a selection of code written specifically for the purposes of this project. The code used here allows more flexibility to the user and enables adjustments to be made to the method of analysis without the need to repeat the simulations, providing further consistency throughout. The simulation can be checked at various stages in order to verify the integrity of the results obtained.

The model is initially verified against basic physical processes, namely the specific energy loss of electrons and protons in silicon. The data output is constantly monitored to ensure both a logical and scientifically feasible output is being obtained. The model is then advanced to a more accurate standing following which the results are compared to those obtained in-orbit.

[5] As stated in the previous chapter, the work described here has been published in the papers Hall *et al.* (2007), Hall *et al.* (2008) and Hall *et al.* (2010) at different stages, with some figures shared between this chapter and the publications detailed (all by the thesis author).

The model is then adjusted to both more accurately represent the real-life detectors and more closely match the data obtained in-orbit. Using a process of continual adjustment it is possible to see the factors of the detector design which have the greatest impact on the instrument background and it enables a more complete understanding of the instrument background to be obtained through a continual, progressive process.

The data obtained from the model have been analysed and the results show a distinct set of components which form the instrument background, namely the Compton recoil electrons, the knock-on electrons, the minimally ionising charged particles and the X-ray fluorescence. The dominance of these components in the energy range of interest shows clear trends and correlates well with factors of the detector design and orbital parameters.

Finally, this chapter concludes with a brief look at the implications that this study has on the overall sensitivity of the camera systems, taking into account the differing device structures and the influence of the QE.

7.2 Initial testing and basic model verification

A simulation is only a valid method of scientific investigation if one can trust the results produced. It is never possible to have full faith in simulated data without significant verification of all components of the model. The description which follows is based on the initial testing of the simulation whilst climbing the steep learning curve involved with the Geant4 toolkit.

Data are readily available for the specific energy loss of protons and electrons in silicon. In order to check that the basics of the simulation are operating correctly the energy loss in a silicon layer was plotted for varying energies of incident protons and electrons. The output from Geant4 works on the basis of energy deposited in a step – a small movement inside the sensitive detector. All the steps are read from the Geant4 output file and summed using a separate C++ program.

7.2.1 File structure

Individual files are required to detail each part of the simulation. These files, coded in the C++ programming language, describe each part of the simulation before finally being compiled using the Geant4 toolkit. Geant4 enables the user to create the particles described, to track them through their individual geometry and to log the results in a way suitable to the purpose. It is in no way a complete software package and should be viewed as a means to move the particles through the geometry and trace their interactions with matter.

7.2.1.1 Detector Construction

The detector is defined as the geometry of the system and the materials from which it is made. It is defined in two files: the header file *include/DetectorConstruction.hh* and the source file *src/DetectorConstruction.cc*.

The world volume defines the limits of the simulation universe. Only interactions and movement inside the world volume are considered and any particle leaving the world volume is lost from the system. For this model, the world volume consists of a 1 m sided cube at near vacuum. In Geant4 it is not possible to create a perfect vacuum so air with density 1×10^{-25} g cm^{-3} was used.

The sensitive detector in this initial model is defined as a 10 µm thick silicon layer of area one square metre. This is effectively an infinite sheet in relation to the interaction distances involved. It is defined as with the world volume. All interactions in the sensitive detector can be accessed through code in the file *src/SensitiveDetector.cc* with the inclusion of the appropriate header file.

7.2.1.2 Primary Generator Action

Particles that enter the system are generated through the use of code in the file *src/PrimaryGeneratorAction.cc* and are fired from a particle gun at a user specified energy and position/angle. For the basic model, the position of the gun is fixed at half a metre from the

centre of the silicon sheet at the edge of the world volume. The particles used as an input in this example are the proton and the electron.

7.2.1.3 Physics List

The physics of the model and the particles used are defined through the file *src/PhysicsList.cc*. This file calls six more files which define individually the particles, the electromagnetic physics, the general physics, the hadronic physics, the ionic physics and the leptonic physics. These files were adapted from those included in the standard Geant4 *GammaRayTel* example file.

7.2.1.4 Sensitive Detector

In this basic model the sensitive detector is defined as the thin silicon sheet. All interactions in this layer can be accessed through code in the file *src/SensitiveDetector.cc*.

Geant4 manages the movement of particles and their interactions using a short step-by-step methodology. Processes can be called for particles at rest, at the start of a step, during a step and after a step. Steps which occur inside the sensitive detector are logged and the particle is tracked. In this model the initial and final coordinates of each step are sent to the *RunAction* along with the energy deposited in the silicon during the step. All interactions occurring outside of the silicon are not recorded.

7.2.1.5 Run Action

An 'event' occurs if energy is deposited during a step. If an event occurs then the code in the file *src/RunAction.cc* is called and the coordinates and energies are recorded as the tab delimited data file *results.dat*.

7.2.1.6 GNUmakefile

In order to compile the various files used in the simulation the *GNUmakefile* calls the file containing the main method and builds the executable in the default folder *../bin/WIN32-VC* if using *cygwin* with Windows, or *../bin/Linux-g++* if using Linux.

7.2.2 Specific energy loss

The specific energy loss is a measure of the differential energy deposited per unit length (dE/dx) of a particle's path through a material. The deposition of energy in the materials found in the simulation is an important factor in the model and the specific energy loss is therefore a useful test of the basics of the code operation.

7.2.2.1 Protons

The data for the comparison with protons was taken from Ziegler *et al.* (2008). A simulation was implemented using 100 protons at energies of 1 MeV, 10 MeV, 100 MeV, 1 GeV and 10 GeV through a 10 µm layer of silicon, Figure 7.1. The value of -d*E*/d*x* in keV per µm of silicon was approximated by averaging the energy deposited over the 10 µm. This is only an approximation as the differential should be taken as the limit as δx tends to zero, but in the case of the protons the energy loss is negligible compared to the energy of the incoming particle. These results show that, at least to this basic level, the simulation performs as expected.

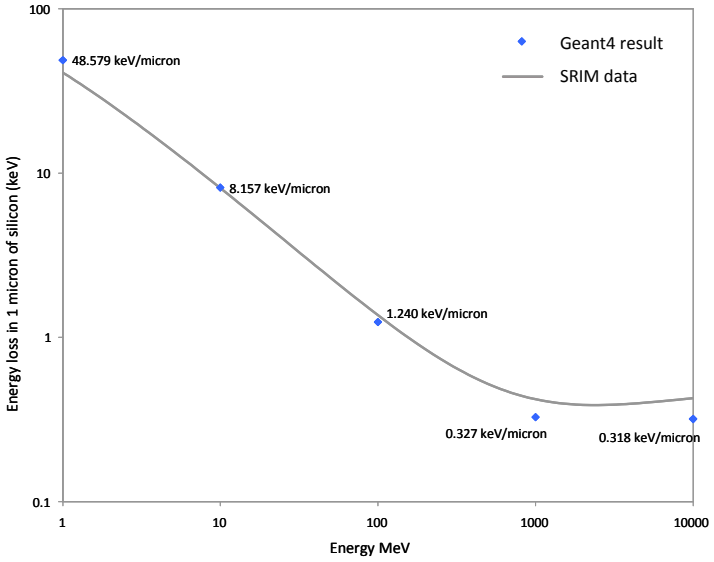

Figure 7.1. The specific energy loss of protons in silicon - a comparison of The Stopping and Range of Ions in Matter (SRIM) software data (SRIM website) and the Geant4 output.

7.2.2.2 Electrons

The specific energy loss for the electron data was found through the same method as described above for the proton. The curve of the plot for the 10 µm silicon layer case, as shown in Figure 7.2, is higher than would be expected - the model shows more energy being deposited per µm in the silicon. This apparent increase can, however, be explained as the energy loss per µm is no longer negligible compared to the energy of the incoming electrons.

The values of -d*E*/d*x* in keV per µm of silicon was calculated as the average per µm over 10 µm. After *x'* µm, following this logic, the energy of the electron will be:

$$E_{e^-} = E_{e^-}^{initial} - \left(dE_1 + dE_2 + \ldots + dE_{x'}\right) \tag{7.1}$$

where dE_x corresponds to the energy loss between the $(x'\text{-}1)^{th}$ and x'^{th} µm of silicon. It is clear from the curve in Figure 7.2 that the energy lost per µm (*-dE/dx*) increases as the energy of the electron decreases.

$$-\frac{\delta E}{\delta x} \approx -\frac{\sum_{i=1}^{x'} dE_i}{x'} > -\frac{dE}{dx} \tag{7.2}$$

This is confirmed with the use of a 1 µm layer of silicon in place of the 10 µm layer. Following this through, one can predict that if the limit is taken (Equation 7.3) then the data will be consistent. It was not possible to verify this result in more detail with the simulation due to the large run times required to gain enough data with an infinitesimally thin silicon layer.

$$\frac{dE}{dx} = \text{limit as } \delta x \text{ tends to zero of } \frac{\delta E}{\delta x} \tag{7.3}$$

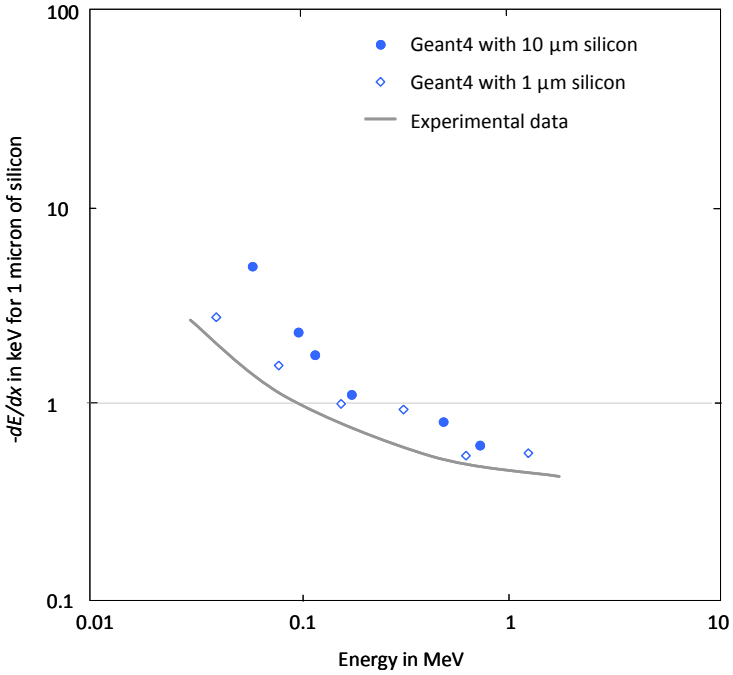

Figure 7.2. The specific energy loss of electrons in silicon - a comparison of experimental data and Geant4 output for both 1 µm and 10 µm thickness of silicon.

7.2.3 X-ray fluorescence

The need for X-ray fluorescence in the simulation is debateable. Whilst it is desirable to include X-ray fluorescence for completeness, it is the continuum of the instrument background which is of most interest. In the final model, following these results, X-ray fluorescence is ignored. The results from these tests showed that the inclusion of X-ray fluorescence would lead to much longer run times for the simulation due to the added complexity involved in modelling the fluorescence X-rays. Detail of these tests are included here for completeness. The materials tested were those of which the majority of the geometry comprises: aluminium and silicon.

7.2.3.1 Aluminium

The geometry of the model was changed to use a sensitive detector, defined as a vacuum, which is simply a volume in which the energy of all particles passing through is logged. The particle gun was set between an aluminium sheet and the detector in order to aid the detection of fluorescence from the aluminium surface.

The incident photons were taken with an energy of 3 keV. A run of one million photons yielded the results shown in Figure 7.3. The energies which are in close proximity may be grouped and differences assumed to be due to the method used to compute energy values from the Geant4 data sets. The Geant4 code produces fluorescence peaks at 1.47 keV and 1.55 keV which relate to the Kα lines at 1.487 keV and 1.486 keV and to the Kβ line at 1.553 keV respectively.

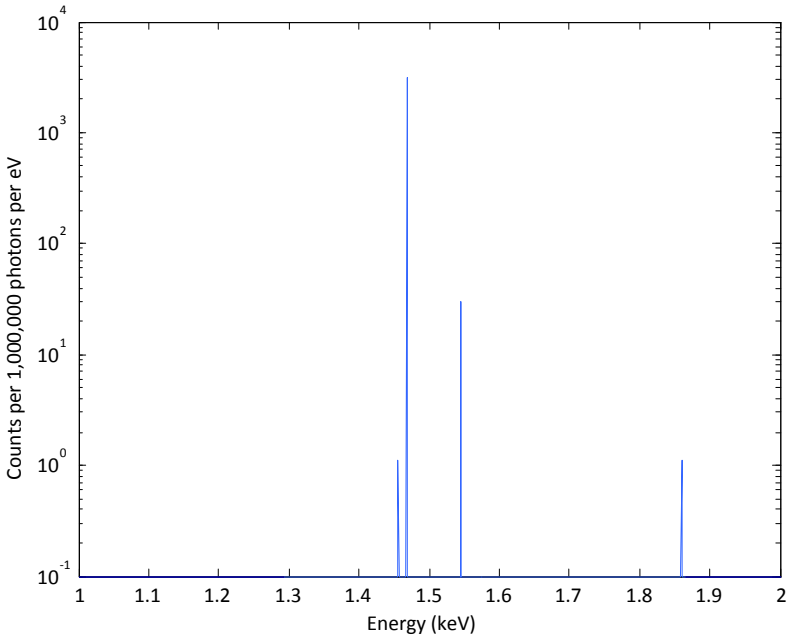

Figure 7.3. Simulated fluorescence from aluminium: counts per million photons from an aluminium sheet bombarded with photons.

7.2.3.2 Silicon

The procedure used for silicon was the same as that for aluminium with the results shown in Figure 7.4. The energies which are in close proximity may again be grouped with the differences due to the method used to compute energy values from the Geant4 data sets. The Geant4 simulation produces fluorescence peaks at 1.72 keV and 1.82 keV which are assumed to relate to the Kα lines at 1.740 keV and 1.739 keV and to the Kβ line at 1.832 keV respectively.

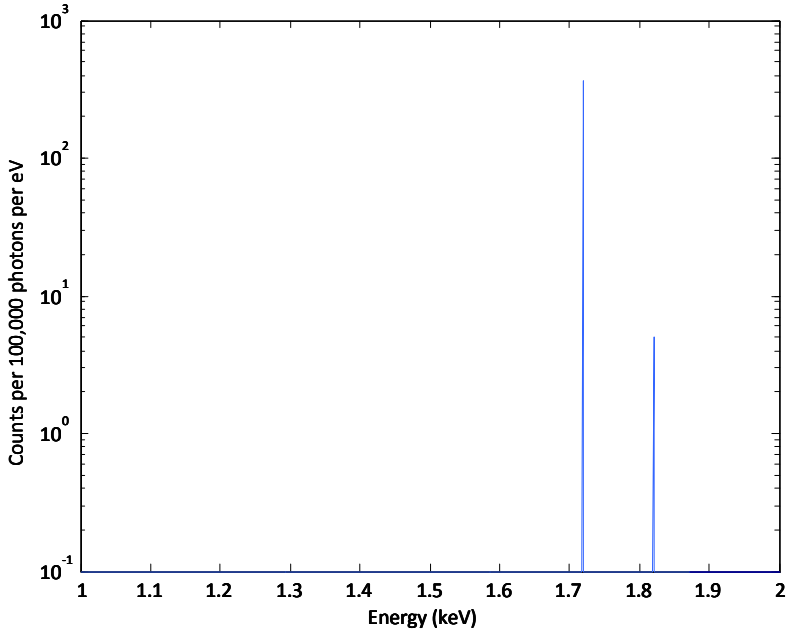

Figure 7.4. Simulated fluorescence from silicon: fluorescence from silicon sheet detected for 100,000 incident 3 keV photons.

7.3 The Geant4 instrument model

For each mission, the same building blocks for the simulation are used. There are four main components to the simulation: the detector and surroundings (the geometry), the physics to be implemented, the incident particle spectrum and the collection and analysis of the data.

Each of these components can be treated separately and coded independently prior to compilation through the Geant4 toolkit.

7.3.1 The detector and surroundings

The detector and surrounding materials are defined in a single geometry file. The code first describes the elements to be used, followed by the definition of the materials formed from these elements. Full control over the elements and materials is given to the user. There then follows the description of the individual segments of the detector. It is possible to define the detector in full detail down to the individual components and gate structures, but here this was thought to be unnecessary. The effect of approximating the detector construction was studied by Murakami *et al.* (2006) for the Suzaku mission and it was found that a simple spherical shell can be used to simulate the detector surroundings to a reasonable degree of accuracy. Any complication to the construction adds substantially to the run time of the simulation as the particles have to follow more complex paths, with the increase in the running time dependent on the complexity that is added to the system. The results detailed later in this chapter confirm that the presence of 4π of appropriate materials at the required separation is all that is required.

The detectors used in the XMM-Newton and Suzaku missions were simulated here, surrounded by spherical shells of aluminium of approximately 3 cm in thickness (Figure 7.5). This aluminium represents the material that surrounds the detectors. The individual physical interactions which occur outside of the detector itself are unimportant; it is the spectrum of particles emerging from these materials which is of greatest importance. The pixel size and detector thickness was accurately modelled on the real detector parameters. A section of the detector was nominated the 'Sensitive Detector' (SD) and it is in this section of the detector that Geant4 is set to log the raw data output from the simulation.

In initial tests of the simulation, the silicon based detector was removed from the system and replaced with a spherical shell (vacuum) nested inside the aluminium. This shell was set as the SD in order to log the spectra of particles which are able to pass through to the detector.

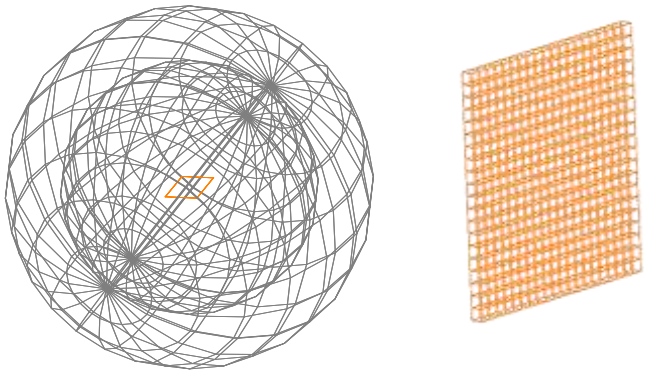

Figure 7.5. The CCD detector sits inside an aluminium shell with approximate inner radius of 4.5 cm and outer radius of 7.5 cm (shown, left, from the Geant4 simulation). A larger cube surrounding this arrangement defines the 'world' of the simulation. Particles are only tracked whilst inside this volume. An example pixelated layer of silicon is shown (right) replicating the parameters of the detector.

The SD Sections of the detectors were split into pixels, matching the dimensions of the real-life devices. For the Suzaku devices with pixel side of 24 µm the detectors were created with 250 × 250 pixels. For the XMM-Newton devices, the MOS CCD with pixel size 40 µm was created with 250 × 250 pixels and the pn-CCD with the larger 150 µm pixels was created with 80 × 80 pixels. The pn-CCD was created with fewer pixels, but covered a slightly larger area due to the larger pixel dimensions. All results were normalised to 1 cm^2 in the analysis of the data.

7.3.2 Physical processes

The physical processes to be implemented in the model are defined in a separate file. The Geant4 toolkit includes all the physics that is required in such a model, but it is still necessary

to list the physical processes which are thought necessary to the running of the simulation. The processes dominating here are included in the Geant4 Low Energy Electromagnetic package which is designed to describe the interactions of photons, electrons, hadrons and ions with matter down to 'low' energies (on a 'particle physicist' scale), ranging from approximately kiloelectron volt levels up to the order of hundreds of megaelectron volts.

7.3.3 Incident particle spectrum

The incident particle spectrum is in many ways the most difficult component of the simulation to compute. This is not due to difficulties in the coding but is due to the somewhat incomplete nature of current knowledge on cosmic-ray spectra. Approximations must be made to the current model of cosmic-rays at the appropriate orbit for each mission. In order to realistically represent in-orbit conditions it is necessary to randomise certain aspects of the particle generation process. The particles are split by type, with separate simulation runs to account for protons, electrons and photons. This not only simplifies the particle generation process but also allows instant separation of the instrument background by 'incident particle type'.

The energies of the incident particles are generated using a Monte Carlo method. The equation for the particle flux is normalised to unity at the minimum energy required to maximise efficiency. A first random number is generated between the minimum and maximum energies being considered. A second random number, f, is generated with a value between 0 and 1. The normalised flux at the randomly generated energy is calculated and compared to f. If the flux at this energy is higher than f then the particle is generated. If the flux is lower, this particle is not generated and the process repeats (Figure 7.6). This provides a simple method of generating the correct incident particle fluxes.

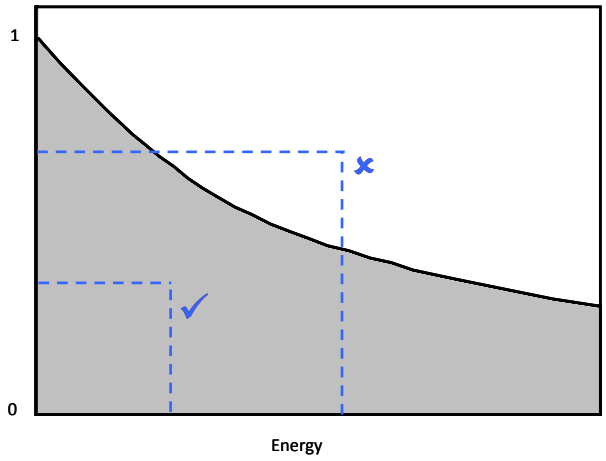

Figure 7.6. Two random numbers are generated, one for each axis. These represent the energy and *f*. In the first case (tick), *f* is lower than the flux, so this particle is generated. In the second case (cross), *f* is greater than the flux at this energy and so the particle is rejected and the process repeated until a particle is accepted.

The directions and initial points of interaction of the incident particles are randomly generated using the following process. A random point *x* is generated on the outer surface of the aluminium shell using Equations 7.4-6 (where *rndm* is a random number and all other variables use the standard notation for spherical coordinate systems). This gives a uniform distribution (in area) for the initial particle positions. A second random point is generated and the direction between these points is calculated, Equation 7.7. The particle is then generated at the first point and set in motion to follow the direction to the second point, defined by the vector *u* in Equation 7.7. This provides an isotropic distribution of incident particles, accounting for the angular independence of the incident particle flux averaged over many orbits.

$$\vartheta = 2\pi \times \mathrm{rndm}_{\vartheta} \tag{7.4}$$

$$\varphi = \arccos\left(\left(2 \times \mathrm{rndm}_{\varphi}\right) - 1\right) \tag{7.5}$$

$$\bar{x} = \left(R\cos\vartheta\sin\varphi,\quad R\sin\vartheta\sin\varphi,\quad R\cos\varphi\right) \tag{7.6}$$

$$\bar{u} = \bar{x}_1 - \bar{x}_2 \tag{7.7}$$

7.3.4 Collection of data

Geant4 operates through the use of short particle 'steps'. A 'step' is a small segment in the path of a particle. Each 'step' has a start point and end point and contains various pieces of information (such as the particle type and the energy lost in the 'step'). It is also possible in Geant4 to aggregate all the steps along a particle path to create one 'hit' per track. In this case aggregating steps to create one track, although less time-consuming, was thought to be undesirable. By leaving the raw data in a step-by-step form it is possible to extract more information and more importantly it allows the full tracking of the energy deposited by the particle through the detector independently of the pixel size and thickness of the detector. Although not implemented here for reasons explained later, it is possible to use the same raw data files with separate code to simulate the charge spreading through the CCD. By implementing the analysis of the raw data using separate code it is possible to carry out multiple forms of analysis on the same data set without the need to repeat the more time consuming Geant4 simulations.

The sensitive detector is the area in which data can be logged. Data were output on a step-by-step basis to a single data file designed to aid further analysis. The data file was based around a ten line layout per 'step', each line representing a separate piece of information. The incident particles were split into batches representing one second in real-time (normalised to the incident proton flux observed in orbit). Following each 'second' of data, a marker was placed in the file before continuing onto the next 'second'.

7.4 The simulation at run-time

The geometry of the system is defined in the *DetectorConstruction* file. A particle is generated with a random energy and is either accepted or rejected depending on the weighting system defined in *PrimaryGeneratorAction* (see Section 8.3.3). The form of the spectrum for this

particle is reproduced mathematically in the file. A second random number is generated, allowing a Monte Carlo fitting of the particles to the form of the spectrum on an accept-or-reject basis. If the particle is rejected then another is generated. If accepted, the particle is placed in the geometry at a randomly generated position on the surface of a sphere, just outside of the geometry, and set in motion to another randomly generated position on the spherical surface, producing an isotropic range of *n* incident particles.

Upon entering the aluminium shell the particle is tracked on a step-by-step basis. All physical processes included for the particle in the *PhysicsList* file are considered and the Geant4 toolkit selects the process which will first occur. At this step, the selected interaction occurs and any secondary particles will then be tracked (along with the primary if necessary) until leaving the aluminium shell and again passing into the vacuum. If the particle passes out of the world volume the particle is removed and no longer tracked. If the particle passes into the aluminium shell at another point then this process continues.

If the particle passes into the detector then the process is slightly different. Any energy which is deposited in the sensitive area of the detector is logged in the results file. The *RunAction* file and *SensitiveDetector* file define how the data are logged. In this simulation the sensitive area of the detector is split into pixels. The coordinate number of the pixel is recorded to allow data processing at a later stage along with the particle type and the energy deposited in the step. Various other pieces of information are recorded regarding the physical processes involved in the deposition of the energy and the way in which the particle was produced if it is a secondary particle.

This process continues until a pre-determined number of particles have entered the system (related to the integrated flux from the spectrum for that particle). Once the set number of particles have been generated, a marker is placed in the output file and the process restarts. In this simulation the number of particles corresponding to one second of real-time is used to

separate the results file. This enables the distinction between two separate events occurring in adjacent pixels and one event split over two pixels.

It is possible to analyse the data whilst the simulation is still running, allowing any problems to be notified early in the run. After a suitable number of runs the data logged during the simulation is analysed using a selection of bespoke C++ code (written by the author).

7.5 Data analysis procedures

Once the simulation has been completed the data generated requires further processing to enable a detailed analysis. As it stands, the data output from Geant4 is not in a suitable form for analysis. These data are of little direct use in providing an overview of the background but will be used at a later stage to gain a deeper understanding of X-ray-like event formation.

7.5.1 Normalising to the incident proton rate

When a proton passes through the silicon of a CCD a track of electrons is left behind which may then be detected. In reality these tracks are clearly visible in the images and can be counted. Work by Lumb *et al.* (2002) and Murakami *et al.* (2006) states a rate of protons passing through the detectors for the XMM-Newton and Suzaku missions respectively. These rates correspond to 2-2.5 protons cm^{-2} s^{-1} for XMM-Newton and 0.6 protons cm^{-2} s^{-1} for Suzaku. The difference in the rates is due to the differing composition of the cosmic-ray spectra at the two orbits (see Chapter 6 for more detail).

With the output from the simulation it is possible to carry out a similar calculation to determine the incident proton flux in the model. Using the particle type and step number it is possible to calculate the total number of protons from the output file. Each proton passing through the detector will run in step number from x, x+1, x+2 through to x+n, where x is the step number upon entering the detector and n is the number of steps during passage through the detector. Using this information one can sum the total number of individual protons. Combining this with the total number of runs (through summing the occurrences of the

'end of run' marker in the file as described above, marking each second in real-time) one arrives at a count rate of protons per second per unit area. Using this conversion factor, the 'second' of data as defined in the output file can be normalised to match that of the in-orbit devices. It is not desirable to continuously log data without separation into these blocks as without this separation there may be event pile-up giving a disproportionate spectral form.

7.5.2 Converting the data file

The output file from the simulation contains a lot of information in a somewhat complex form, Table 7.1. In order to convert this data to give a spectrum of counts per unit energy it is necessary to carry out further analysis using purpose-written code.

Output	Data	Explanation
247	Pixel number 247 (row)	Pixel # in *y*-direction (row)
30	Pixel number 30 (column)	Pixel # in *x*-direction (column)
0.00507723	0.00507723 MeV deposited	Energy deposited
e-	Electron	Particle type
1	1st step in material	Step number through material
hLowEIoni	Particle created by ionisation by hadron	Creator process of the particle
eIoni	Energy deposited through ionisation	Energy deposition process
(-3.762,-0.00578,4.903)	(x,y,z)	Position at start of step
(-3.763,-0.00584,4.902)	(x,y,z)	Position at end of step
1	Parent particle was a primary particle	Parent particle ID number
endOfRun	End of one 'second'	Marker

Table 7.1. The output from the simulation is explained with sample data. This corresponds to one 'step' in the SD falling at the end of one 'second'.

The data are analysed on a second-by-second basis. There are four 2D arrays initiated at the start of the program, matching the pixel arrangements of the detectors: the energy array, the proton array, the creator process array and the depth array. The meanings of these arrays will become apparent following the explanation below. The first two lines of the data file determine the element in the energy array. The third line, the energy deposited in the step, is added to the value of this element. If the particle is a proton, as declared in the fourth line of the data file, then the corresponding element in the proton array is set to true. The sixth line of the data file is written to the creator process array using a numerical code. This enables the events to be split by the physical creation interactions. The depth at which the particle initially deposits energy in the CCD is extracted from the eighth line of the data file.

After one second of data has been extracted, the energy array is sequentially tested for the X-ray-like patterns to distinguish single, double, triple and quadruple pixel events, rejecting all others as non-X-ray-like events. For both missions the same grading system was used to simplify the code. This information is output to a separate file with the pixel coordinates, total energy deposited, pixel pattern and initial depth of entry of the particle recorded. Depending on the detector type, the proton array can be used to exclude the signal caused by protons. In the case of the pn-CCD and back-illuminated CCD the signal from the protons will be detected as X-ray-like providing the pixel pattern is X-ray-like and the energy is in the selected X-ray energy band. When a proton passes through the MOS CCD or front-illuminated CCD there is a large amount of charge spreading in the silicon below the depleted region. This leads to the spreading of charge over many pixels and in reality the event would be rejected due to non-conformity to the X-ray-like grading system. In the simulation these events are removed by rejecting all events in which the proton array has been set to true. This removes the requirement for the modelling of charge spreading in the CCDs. The arrays are then returned to zero before continuing with the next ‘second’ of data. The data file produced is now in a format which enables simple conversion to both images and spectra.

7.5.3 Producing images

Using MATLAB, the data from the energy array described above can be converted into an image file. The data are initially processed through separate code, where the 'time interval' of the image can be selected by the initial and final 'seconds' of the data. The images provide a useful insight into the simulated data (an example simulated image is shown in Figure 7.7). In these images there are clear tracks of higher energy left by protons passing through the sensitive area of the detector. These tracks can be rejected on processing through the use of pixel pattern determination, where only single, double, triple and quadruple pixel events are not rejected. The image on the right shows a section taken in a smaller area with a shorter time scale where individual events can be seen in the image. When producing the initial spectra the position of pixels producing the counts was cross-checked with the images and this was found to match as expected.

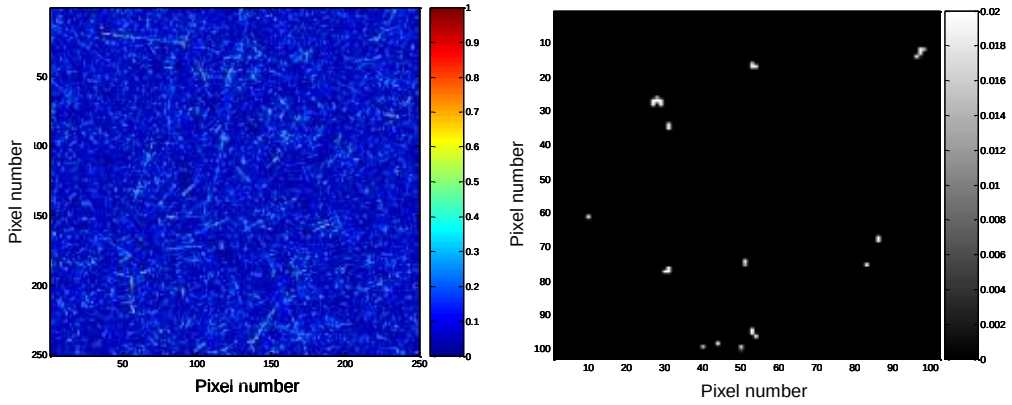

Figure 7.7. Sample images of a CCD output from the simulation. The image on the left has the energy shown on the right bar in MeV. Tracks formed by protons can be seen. In the image on the right, a small region of the detector is shown with the energy scale again shown in MeV. The events are restricted to those from 0 keV to 20 keV corresponding to possible X-ray like events. For example, single pixel (11,61) and multiple pixel (30,77) events can be clearly seen, along with events which would be rejected by the grading system (30,28).

7.5.4 Producing spectra

The data files can be easily converted to a spectral form using MATLAB's histogram functionality. These simulated results can then be compared to the in-orbit spectral continua as shown in the previous chapter. It is these spectra which are the main purpose of this study. The level of the instrument background can be clearly seen from the spectra and comparisons can be made between different orbits and detector designs. Understanding the background is the first stage in optimising the detector performance for future missions.

7.5.5 Analysing the spectra

Once the spectral continuum has been reproduced it is necessary to determine the constituent factors contributing to the instrument background. To deconstruct the spectrum it is necessary to refer back to the output data file and the converted data file. Through the running of the simulation in particle-specific batches the data are naturally split by incident particle type. From the converted data file it is possible to separate the spectra into pixel grades (single, double, triple and quadruple pixel events) and to separate by the 'particle creator process'. The results that have been split by 'particle creator process' are of great interest as they determine the interactions which produce the background. Through an understanding of the constituent components producing the background it is possible to find ways in which to reduce these effects.

By referring back to the initial output data files it is possible to track individual particles from the point at which they first deposited energy in the detector through to the point at which they leave. From such information it is possible to deduce the path of the particles before they enter the detector. A good example of this is taken from the XMM-Newton results. It was found that the majority of the signal was being detected within a few micrometers of the upper surface of the device. The particles producing the signal were electrons produced

through ionisation by hadrons, *hLowEIoni*. This shows that a proton has knocked electrons from the aluminium shell and these electrons are being detected as X-ray like events in the detector and are hence being detected at the surfaces of the device, as electrons of several keV may only travel a few micrometers through the silicon.

7.6 Process and Results

The results for the simulation are described below, following the simulation process from the initial testing of the models through to the final spectral simulations. The incident particle spectra are discussed for each mission with regards to the aspects involved in efficiently simulating an appropriate incident particle flux whilst maintaining the integrity of the results produced. The results are given initially for the XMM-Newton mission, followed by the results for the Suzaku mission, giving an overview of the major components involved. These results then allow analysis of the Swift mission, using the same detectors as those for XMM-Newton, but with a similar orbit to the Suzaku mission (Chapter 6). A full discussion of the results is left for the following section.

7.6.1 XMM-Newton

A study by Lumb *et al.* (2002) determined that, at the specified orbit, the XMM-Newton mission would be subject to "close to interplanetary galactic cosmic ray spectra at solar minimum conditions". This spectrum is dominated by protons, of which only those with energies greater than approximately 80 MeV will pass through the 3 cm aluminium shielding. The approximation to the spectrum used here runs from 80 MeV over a 600 MeV band (at which point the diminishing of the spectrum leads to a comparably negligible flux). The spectral form of the proton component was based upon the SPENVIS calculations detailed in the previous chapter using the appropriate orbital parameters and time periods. The spectra were considered over many orbital periods such that the influence and direction of the passage through the SAA could be averaged. The total flux of protons through the detector is

normalised in all cases to 2.2 protons cm^2 s^{-1} as described from in-orbit observations (Lumb *et al.*, 2002). Initial results used to verify the system were primarily based on 100 MeV protons only, moving through 200 MeV and 400 MeV components before utilising the complete spectrum, approximating the SPENVIS model (shown in Figure 6.10) using a power law over the energy range indicated above.

Preliminary spectra were examined to determine the most prominent components of the incident spectra and the features in the simulated device which were of most importance. Figure 7.8 shows the preliminary results for the XMM-Newton detectors, based on an incident particle spectrum of purely 100 MeV protons. The basic structure of the detectors was modelled, ignoring the complexities of the pixel structure and the material below the pn-CCD.

These initial results showed a low count rate at lower energies for the MOS CCD and indicated that further development was needed on the pixel structure. The major components of the initial spectra also aided the preliminary understanding of the more dominant components of the background. The initial results pointed towards the dominance of the knock-on electrons: secondary particles produced when the protons pass through the aluminium shell. It is interesting at this point to investigate the spectrum of particles leaving the inner surface of the aluminium shell (those which could eventually hit the detector). A spectrum was recorded containing the electron component leaving the inner surface of the aluminium shell, shown in Figure 7.9. The spectrum has been normalised to the MOS and pn-CCD instrument background continua to allow a comparison of the two spectra. This shows that the signal detected in the CCDs is consistent with the spectrum of electrons emerging from the inner surface of the aluminium shell.

The effect of the pixel size was investigated and the background was considered for detectors of matching pixel sizes. The pixel size was found to have negligible effect on the background continua. The pixel size has most effect on the separation of single pixel and multi-pixel

events, as would be expected – the smaller pixel size leading to fewer single pixel events. Minimally ionising charged particle tracks are more likely to pass through an accepted pixel pattern if the pixels are larger, but the knock-on electron component is not greatly affected due to the short interacting distances of the electrons in the silicon.

The initial results showed that the simulation was modelling the system to a reasonable degree, without fully simulating the more intricate parts of the structure. The MOS CCD single-pixel events initially appeared low at lower energies whilst using a basic gate structure in the simulation. The results shown in Figure 7.10 use a more accurate gate structure taking into account the 'open' nature of the electrodes. The low count rate in the initial results was due to absorption of the lower energy electrons in the gate structure, whereas the open pixel structure allows the lower energy electrons greater passage into the detector. The pn-CCD structure was updated in two ways: a more accurate gate structure was applied to the rear of the device and a thin layer of copper (representing the printed circuit board) was placed beneath the device.

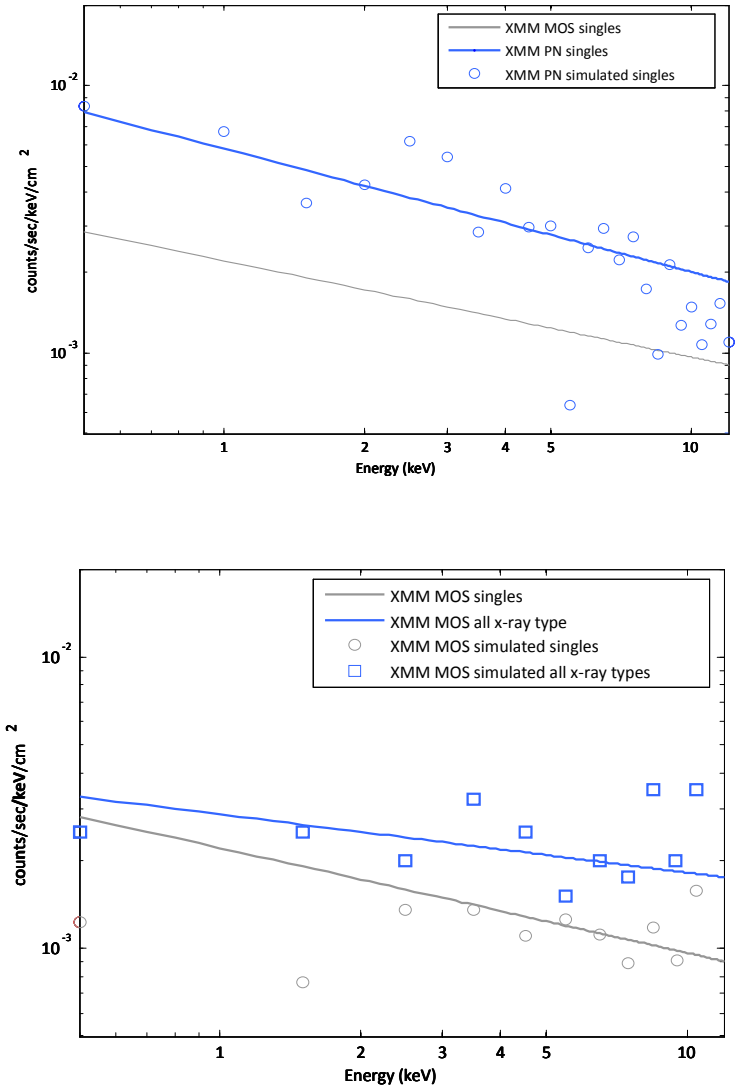

Figure 7.8. Early spectra for the basic simulation of the XMM-Newton detectors. The early runs aided the analysis of the various components of the simulation such as pixel design.

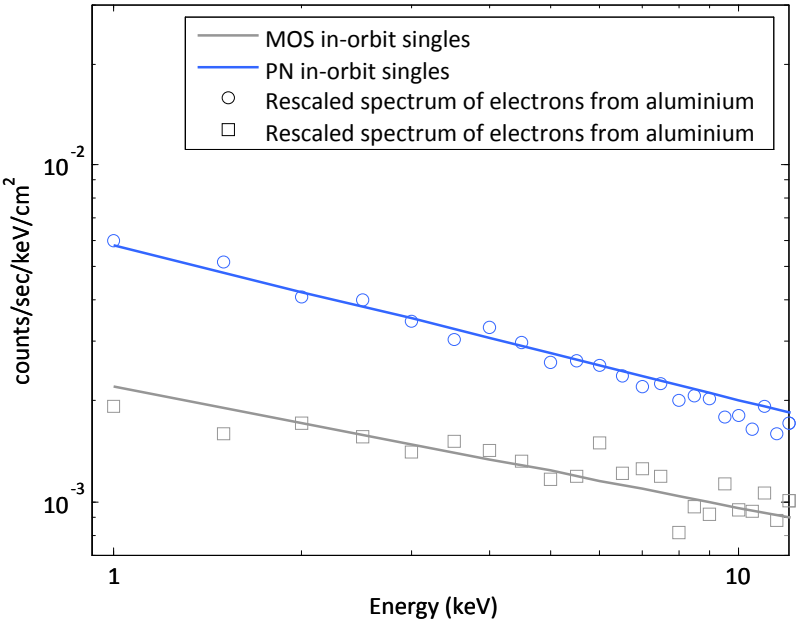

Figure 7.9. The spectrum of electrons leaving the inner surface of the aluminium shell (the knock-on electrons), normalised for comparison with the in-orbit components.

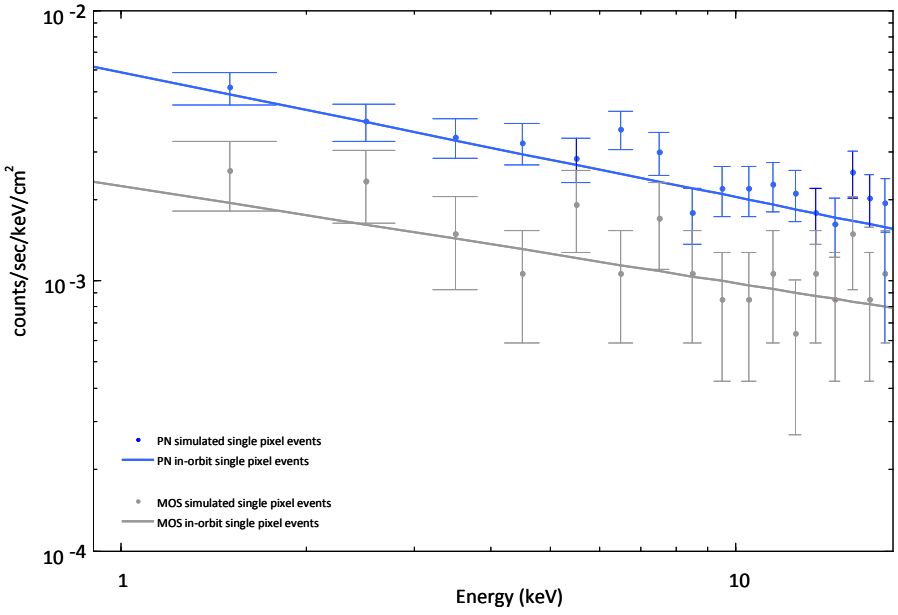

Figure 7.10. The spectra from the Geant4 simulation shown against the in-orbit continua for the more accurate model of the devices. These results allow discussion of the components forming the instrument background through further analysis of the data files.

7.6.2 Suzaku

At the Suzaku mission orbit the craft was expected to experience the standard cosmic-ray flux for this altitude. Using a combination of data from Mizuno *et al.* (2004) and Murakami *et al.* (2006) the cosmic-ray spectrum was split into three main components: protons, electrons and photons. The relative dominance of the varying energy bands was investigated for each particle type with the region used in the simulation being that which had the greatest impact on the instrument background. The lower limit was generally determined by the energy at which passage through the aluminium became probable, whilst the upper limit was taken several orders of magnitude lower in flux, taking into account the energy dependence of the individual components. This approximation to the total incident spectrum was required to save on the simulation run-time. The spectrum for each particle shown in Figure 7.11 was used for the simulation. The incident particle flux was normalised to the integrated proton flux as found by Murakami *et al.* (2006) at 0.6 protons cm^{-2} s^{-1}, normalising to the incident proton spectrum detailed in Chapter 6 from the SPENVIS calculations.

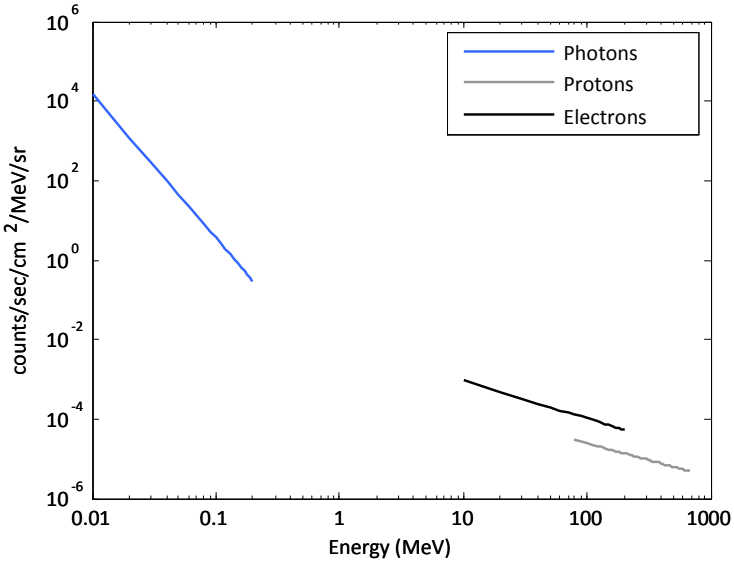

Figure 7.11. The incident particle spectrum used for the Suzaku mission simulation. The energy ranges used are shown here and were selected as the regions of most interest from the initial testing, allowing optimal accuracy in the simulation whilst not compromising excessively on the simulation run-time.

In a study by Murakami *et al.* (2006), aiming to reproduce some aspects of the instrument background for the Suzaku mission, it was found that a constant factor was required to 'shift' the simulated spectra in line with the results taken in-orbit. It is thought that this constant factor is included to account for inaccuracies in the complex incident particle spectra. The results produced by simulations detailed here using the Geant4 toolkit were also found to be a factor of 3-4 different to those found in-orbit, Figure 7.12. The spectral form is however reproduced and allows for an analysis of the components forming the instrument background. Fluctuations in the simulated spectra at lower energies (lower than 2 keV) are due to problems encountered in achieving large data sets for the Suzaku mission due to the complexities of the devices and particularly the incident particle spectrum. Individual components have been modelled based on the incident particle types and interaction processes to provide a detailed analysis of the overall form of the spectra.

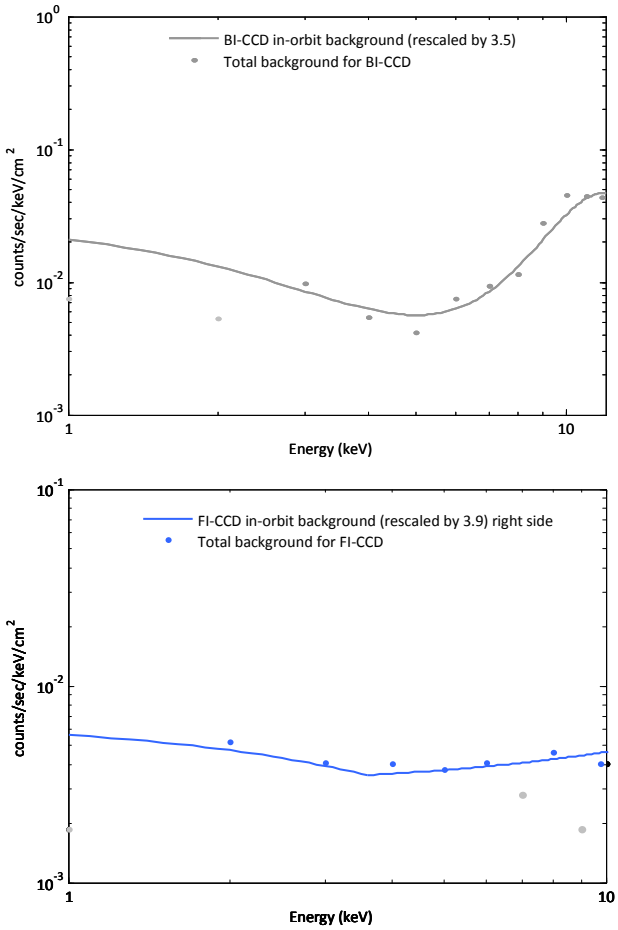

Figure 7.12. The simulated spectra for the Suzaku mission. The BI device results are shown top. The lower energy results (below 2 keV, shown in light-grey) show a similar trend as those found with the XMM-Newton mission due to the approximated surface gate structure. The lower plot shows the FI device. The spectral form of each component producing the total instrument background can be extracted from the data files.

7.6.3 Swift

The MOS CCD detectors used for the Swift mission are the same as those used on the XMM-Newton mission. At the orbit of Suzaku, approximately 0.6 protons cm^{-2} s^{-1} pass through the detector (based on in-orbit measurements by Murakami *et al.*, 2006). The incident proton spectra were determined through SPENVIS calculations for both the Swift and Suzaku orbits

and inclination angles. The variation between the particle fluxes was found to be approximately 20 % at 100 MeV and only 6 % at 300 MeV. With the orders of magnitude underestimation found in previous studies, this minor variation could be considered negligible, but for completeness here the measured in-orbit Suzaku proton flux is altered appropriately for the Swift mission. The model for the Swift mission was tightly based around the XMM-Newton model, re-normalising to the proton flux at the new, lower, orbit.

Minimally ionising charged particles have little impact in FI-CCDs, due to charge spreading in the non-depleted region of silicon. Comparisons with the in-orbit data are discussed below and show good correlation with the XMM-Newton and Suzaku results, providing further backing for the hypotheses discussed.

7.7 Discussion[6]

Analysis of the simulated data shows that the instrument background is built from three main components. The relative dominance of each component is dependent on not only the incident flux of particles, and hence the orbit, but also the design of the detector.

7.7.1 Compton scattering

When a photon passes through the silicon of the detector it is possible for this photon to be Compton scattered. The probability of a high energy photon (energy greater than 100 keV) being detected conventionally is negligible. The recoil electron is detected in the silicon as an X-ray-like event, providing it does not deposit energy over an excluded pixel pattern. Such an event is indistinguishable from an event caused by an X-ray in the energy range of interest. This component of the instrument background is more prominent in detectors with a thicker active layer due to an increased probability of interaction for the primary photon. The thicker depletion layers of the Suzaku XIS detectors lead to a greater Compton electron component

[6] The majority of this section is taken from the paper “Space radiation environment effects on X-ray CCD background”, Hall & Holland, NIM A (2009), written by the thesis author.

than for the MOS detector of the Swift mission. This can be seen in the spectra as the increase in the continuum at lower energies.

7.7.2 Minimally ionising charged particle

When a minimally ionising charged particle passes through the silicon of the detector a track of electrons is left behind. If this track crosses several pixel boundaries then the event is rejected by the grade-selection process. This does not exclude all cases, as it is possible for a particle to pass through the detector depositing energy in only those pixels of an accepted X-ray-like grade. For the back-illuminated devices (Suzaku XIS1 and the XMM-Newton pn-CCD) a peak in the spectrum will be seen at an energy related to the thickness of the silicon. For the pn-CCD the peak falls above the energy range of interest, but for the XIS1 detector, the peak falls at approximately 13 keV and the lower energy tail dominates the background over the upper regions of the energy region of interest. The front-illuminated devices are not fully depleted and the increased spreading of charge generated below the depleted region leads to the events being reduced by the event grading system. The reduction in the peak energy depends on the amount of charge collected from the undepleted silicon, but nonetheless leads to a significant reduction compared to the back-illuminated devices.

7.7.3 Knock-on electrons

The component of most interest in this study is that caused by electrons ejected from the material surrounding the detector through ionisation, predominantly from the aluminium shell: the knock-on electrons. When a proton or electron passes through the aluminium shell it liberates electrons through ionisation along its path. The knock-on electrons are absorbed within a few micrometers of the aluminium and so only those electrons from the inner surface of the aluminium are of importance, all others are absorbed by the shielding. Similarly, the electrons which are able to reach the detector will be absorbed within a few micrometers of the surface.

The energy deposited by the knock-on electrons in the simulation could be traced back to the outer surfaces of the CCDs and the electrons were found to be released in the aluminium housing around the detector by the incident protons, Figure 7.13. The surface structures of each detector define how these knock-on electrons will impact the instrument background.

The MOS CCDs and XIS FI-CCDs are insensitive to the knock-on electrons from the rear-face due to the thick layers of non-depleted silicon. Any surface electrode structure acts to reduce the effect of the knock-on electrons on the instrument background by absorbing some of the energy of the electrons. The absorption of some of the energy from the knock-on electrons acts to shift the spectrum to a lower energy, and due to the slope of this component of the spectrum this lowers the total instrument background. The open surface structure of the back-illuminated devices (such as the XIS BI-CCD and pn-CCD) increases the quantum efficiency at the expense of a greater component of instrument background from the knock-on electrons. The open structure allows all knock-on electrons to be detected and the effects of this can be seen by comparing the instrument background for the XMM-Newton MOS CCD and pn-CCD. The knock-on electrons can only be detected on the upper surface of the MOS CCD and this component is somewhat reduced, even with the open-electrode structure. For the pn-CCD, the knock-on electrons may be detected in full at the open surface and partially detected at the rear-electrode surface. The combination of the absorption of some of the energy of the knock-on electrons and the two-sided against one-sided detector produces a large difference in the total instrument background. This effect can also be seen with the two XIS detectors types. If the Compton recoil electrons were the dominant component of the instrument background then one would expect the thicker layer of depleted silicon in the FI-CCD to give rise to a higher instrument background. This is not the case. The BI-CCD has a higher level of instrument background, explainable by the knock-on electrons.

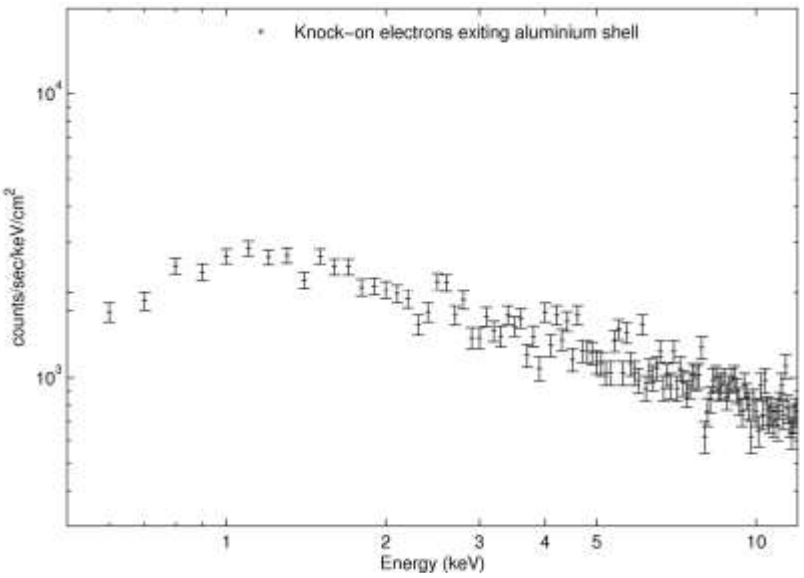

Figure 7.13. The simulated spectrum of electrons emerging from the inner surface of the aluminium shell following bombardment with protons using the same spectrum as that used for the XMM-Newton simulation.

7.7.4 Further FI-CCD vs. BI-CCD comparison

As a comparison with the front-illuminated devices, it is worth considering what would happen if the back-illuminated devices were mounted onto a thick silicon dummy layer. The dummy layer would reduce the knock-on electron component from the detector shielding, but any protons passing through the dummy layer would generate knock-on electrons of their own.

It may be possible to reduce this new source of knock-on electrons as the proton tracks will now be in close proximity to the detected knock-on electrons, whereas protons interacting in the shielding produce knock-on electrons which can arrive at the detector independently of the proton trajectory. A dummy layer would not reduce the component of the instrument background produced by the minimally ionising charged particles as charge spreading in the dummy layer would not be detected. The MOS FI-CCD does not only consist of the image region, but also contains a frame-store region (for frame-transfer). The frame-store region is not identical to the image region, with the main differences being namely the ‘integration

time', the pixel size and the structure of the electrodes. Signal will be accumulated at an average rate equal to half of that seen in the image region (ranging from zero at one end of the store to 2.4 seconds at the other). The frame-store pixels measure 12 µm × 40 µm compared to the 40 µm × 40 µm pixels in the image region. These two factors lead to a basic reduction in signal to 0.15 that in the image region. The electrodes in the frame-store region also follow the standard pixel construction without the open area. Combining these effects with the findings of the simulation leads to the conclusion that the instrument background accumulated in the frame-store region of the MOS FI-CCD is an order of magnitude lower than that in the image region.

The spectral form for the two types of X-ray CCD camera aboard the Suzaku mission is shown in Figure 7.14, giving a rough diagrammatic view of the instrument background based on results from the simulations, showing the three main components as described above.

The relative dominance of each component is dependent on the structure of the device. The presence of an electrode structure on the active surface of the device acts to absorb some energy from the knock-on electrons, shifting the continuum to a lower energy and hence lower count rate in the FI-CCD.

The absence of the field free region in the BI-CCD leads to a lower rejection efficiency for minimally ionising charged particles, producing the larger Landau peak – with the peak at a lower energy due to the thinner device.

The Compton recoil electron component is greater for the FI-CCD due to the thicker depletion layer, but this is overshadowed by the increase to the knock-on electron component in the BI-CCD.

The total instrument background for most devices in orbit can be considered as a combination of the three main components shown in Figure 7.14, allowing the detector design to be considered from several 'angles' with regards to the instrument background. A careful

balance, dependent on the orbit and surrounding material, must be considered to increase the overall sensitivity of the device (see Section 8.8).

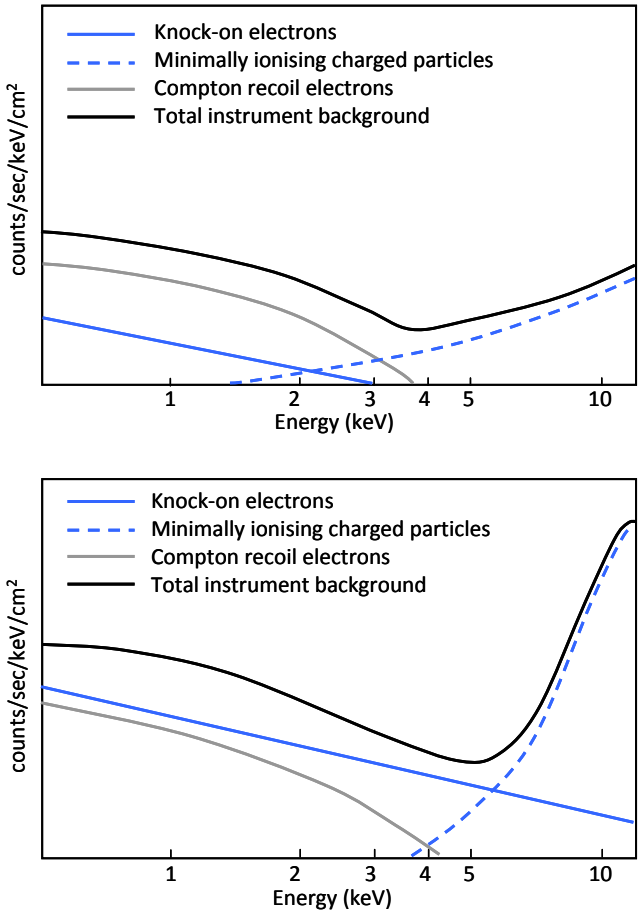

Figure 7.14. A schematic representation of the main components producing the instrument background produced from the results found in the simulations detailed previously in the chapter, with the upper plot detailing the FI-CCD and the lower plot detailing the BI-CCD. The dominance of each component is dependent on the structure of the device.

7.7.5 Orbital dependence

The MOS camera system used in both the XMM-Newton and Swift missions allows a comparison of the effect of the orbit on the instrument background and clarification of the

results previously discussed. The component of instrument background from minimally ionising charged particles is negligible as the MOS CCD is a front-illuminated device comprising of both depleted and non-depleted regions of silicon. The charge spreading in the non-depleted region causes these events to be rejected with the X-ray-like grading system.

If the instrument background is knock-on electron dominated, one would expect the level of background to scale predominantly with the flux of protons. It was stated previously that the flux of protons through the detector was 2-2.5 protons cm^{-2} s^{-1} at XMM-Newton's orbit and 0.6 protons cm^{-2} s^{-1} at the Suzaku orbit. From this ratio, taking into account the slight decrease in incident proton flux for the Swift orbit with respect to the Suzaku orbit (Figure 6.13), a knock-on dominated instrument background hypothesis for a similar solid angle of shielding would predict that the Swift MOS background would be approximately three to four times lower than the XMM-Newton MOS instrument background. This is indeed the case, as shown in Figure 7.15. The Swift MOS instrument background shares a similar form to the XMM-Newton background, but shows an increase at lower energies. At the lower orbit, with the decrease in proton flux and relative increase in photon flux, the Compton recoil electron component for the front-illuminated device is more prominent. The increase in the Swift instrument background can therefore be explained by the Compton recoil electron component.

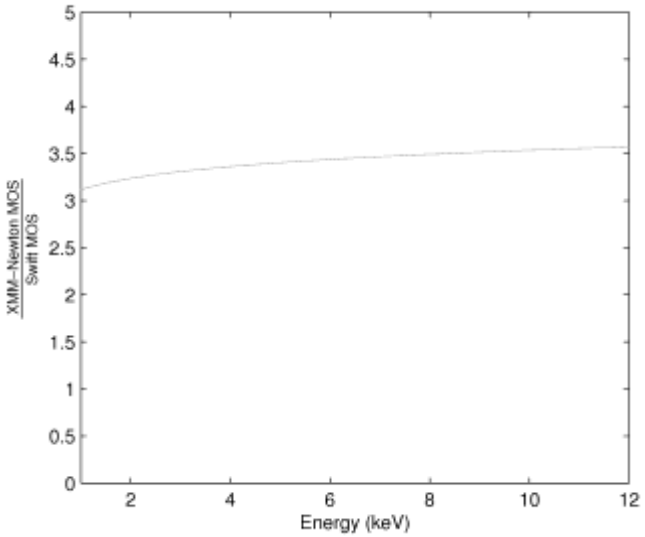

Figure 7.15. The ratio between the instrument background continua for all X-ray grade events in the XMM-Newton MOS and Swift MOS spectral continua previously detailed in Figure 6.15.

7.7.6 Comparison with pre-launch study

The pre-launch study (Dyer *et al.*, 1996), described in Section 7.8, concentrated on the influence of the Compton interactions of secondary gamma rays and not the component caused by the secondary electrons (knock-on electrons). Although a Compton dominated background would show a higher count rate for the thicker pn-CCD over the MOS CCD, the knock-on electrons were found in this study to be a dominant factor in producing the instrument background experienced in-orbit, particularly at HEO, and can also explain this difference. The flux of knock-on electrons, recorded at the CCD as simulated here, closely matches that found in-orbit and dominates over the flux of secondary gamma rays detailed in the pre-launch study by nearly two orders of magnitude. The lower incident proton rate at LEO reduces the knock-on electron component and this reduction can be seen in the Swift MOS in-orbit instrument background. The Compton component at LEO is no longer negligible due to the differing incident spectrum and contributes to increase the instrument background above the purely knock-on electron dominated continuum at lower energies.

7.8 Implications on sensitivity[7]

When considering the sensitivity of each camera, the instrument background must be taken into account. For the EPIC systems aboard XMM-Newton, the mirror (including the effects of the mirror on the amount of diffuse X-ray background which is resolved), filter, source and detector all need to be taken into account when considering the total sensitivity, and hence the minimum observable flux. The net effective area, found through the product of the X-ray telescope effective area and the quantum efficiency of the detector, must be combined with the background level to find the overall detector sensitivity. For the XMM-Newton EPIC MOS systems the net effective area is lower than would be expected relative to the pn-CCD system as only 44 % of the telescope incident flux is directed to each of the MOS cameras, the remainder being directed to the Reflection Grating Spectrometers (RGS) and lost due to the structural obstruction by the RGS grating and supports (Turner *et al.*, 2001). This must be taken into account when comparing with the unobstructed beam reaching the pn-CCD camera. The differences between the quantum efficiencies of the detectors are quite significant, Figures 7.16 and 7.17, and considering the loss in effective area of over 10 % with the two MOS detectors combined, the difference in sensitivity is greatly reduced due to the significantly lower instrument background for the MOS cameras.

Defining a Quality Figure (QF) for the cameras related to the signal-to-noise ratio, it is possible to examine this further. The energy range used depends on the science to be considered and will be taken at 1 keV here as an example. Assuming the mirrors and filters to be common to the two detector types, the signal can be taken as the product of the proportion of the beam reaching the detector, the QE of the device at 1 keV and the collection efficiency (CE), where the CE gives the integration to dead-time for the detector (assumed to be approximately 0.9 for the pn-CCD and 1 for the MOS). At the low signal limit, the noise can be approximated to

[7] This section is taken from the paper “Space radiation environment effects on X-ray CCD background”, Hall & Holland, NIM A 2009 (written by the thesis author).

the square root of the background, dominated by the instrument background and the cosmic X-ray background. We define the QF here as the ratio of the signal to noise as described above for each camera system.

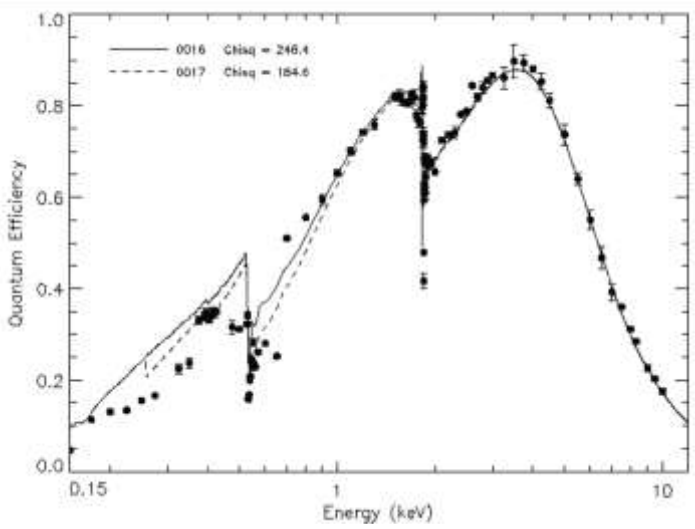

Figure 7.16. The QE for the MOS CCD (Guainazzi *et al.*, 2008), showing a QE of 0.7 at 1 keV.

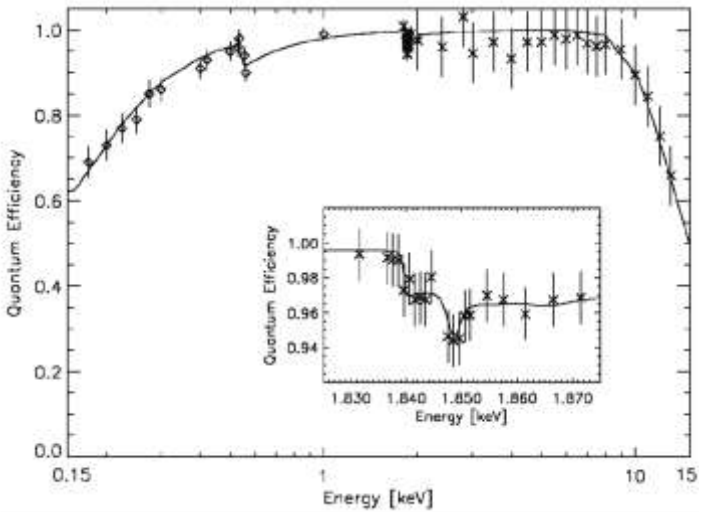

Figure 7.17. The QE for the pn-CCD (Guainazzi *et al.*, 2008), showing a QE of 0.97 at 1 keV.

Using background data from De Luca *et al.* (2001) for the MOS background and from Freyberg *et al.* (2004), the relative background count rates for the MOS and pn-CCD have been extracted and normalised to the MOS instrument background. It should be noted that although the pn-CCD has approximately three times the instrument background of the MOS at 1 keV, the total background from the above data gives a slightly lower total background count-rate ratio of 9:4 at 1 keV (normalised to the XMM-Newton MOS instrument background, MOS IB = 1). These calculations are only relevant for one aspect of the science (low signal source detection) and other aspects of the science at higher energies will generate a different QF. The results of these calculations are shown in Table 7.2.

The Quality Figures calculated in Table 7.2 show an interesting comparison between the camera systems using the front-illuminated and back-illuminated devices. If subjected to the same full beam, the difference between the Quality Figures for the two devices is greatly reduced, and despite the variations in device structure, within the errors of the approximations used here the devices offer equivalent performance.

	Beam proportion	QE	CE	'Signal'	Background (MOS IB = 1)	Quality Figure (QF)
pn-CCD	1	0.97	0.9	0.87	9	0.29
MOS (×1)	0.44	0.7	1	0.31	4	0.16
MOS (×2)	0.88	0.7	1	0.62	8	0.22
MOS (if unobstructed)	1	0.7	1	0.7	7	0.26

Table 7.2. The Quality Figures (QF) for the MOS and pn-CCD cameras. The QF is defined here in a similar way to the signal-to-noise ratio, but with constant factors between the systems ignored for simplicity. The 'signal' is taken as the product of the beam proportion (BP) hitting the detector, the QE and the Collection Efficiency (CE), given by the integration to dead-time ratio. The calculations are based on the low-signal, background-dominated case. The two MOS detectors are first examined individually (×1), then as a pair (×2). The final row shows the QF for the MOS detector if it were presented with the same unobstructed beam encountered by the pn-CCD, allowing a more direct comparison. All background levels are normalised to the XMM-Newton MOS instrument background at 1 keV.

The simulated data suggest that the rear of the front-illuminated devices offer better resilience to the knock-on electrons. The layers of non-depleted silicon at the rear of the devices absorb the knock-on electrons, and consequently they do not add to the instrument background count as in the case of the pn-CCD and back-illuminated devices. The front faces of the devices is where the balance must be achieved between the quantum efficiency and the level of the instrument background. A high QE device such as the pn-CCD is able to detect the lower energy photons more effectively than the MOS CCD. Unfortunately, this optimisation of the QE increases the background; as well as allowing the detection of lower energy X-rays, the knock-on electrons are also able to be fully detected. The balance between the QE and the

instrument background must be carefully considered at HEO where the knock-on electron component dominates. If one is solely looking at X-rays with energies greater than ~0.5 keV then a back-illuminated structure (particularly one which is able to detect knock-on electrons from both major faces) is not necessarily ideal from a sensitivity perspective, although the QE at lower energies may appear more desirable. The QE at higher energies (> 4 keV) for front-illuminated detectors (similar to the MOS) has improved dramatically over the last decade and should no longer be a deciding factor in the detector choice (Murray *et al.*, 2008).

7.9 Conclusions[8]

A simulator has been developed which is in very good agreement with data from spacecraft. Previous studies based on Compton interactions from secondary gamma rays under-predict the instrument background at HEO. The simulation carried out here shows the dominating component of the background at HEO is caused by the knock-on electrons produced by the passage of protons through the detector shielding. The influence of the knock-on electrons is greatly affected by the design of the device, with the surface structure defining how the electrons will be detected. It is worth, therefore, not only considering the effect of the structure of the device on the quantum efficiency of the system but also the influence the structure will have on the instrument background. The possibilities of using a graded-shield for the innermost camera layers could be explored in future work to assist in reducing the instrument background from knock-on electrons and may help reduce fluorescence (Pfefferman *et al.*, 2004). Any considerations to reduce the instrument background must be carefully balanced with their effect on the X-ray fluorescence peaks. Through these considerations the sensitivity of future camera systems can be optimised.

[8] This section is taken from the paper "Space radiation environment effects on X-ray CCD background", Hall & Holland, NIM A (2009), written by the thesis author.

Chapter 8: Conclusions

This chapter provides a summary of the studies carried out for this thesis, the conclusions on the outcomes of the work and indicates possible directions for future investigation.

8.1 High-energy, high-resolution imaging

Through the use of simulations and modelling techniques, various aspects of the scintillator-coupled EM-CCD have been studied. The implementation of the results of this study in an experimental device have resulted in a high-resolution, high-energy, photon-counting gamma-camera. The points outlined below summarise the main outcomes of the study.

- A customisable point-search function was designed to enable the detector output to optimised for the current purpose.
- The theory of scale-space was developed with regards to the analysis of images from a photon-counting gamma-camera. Since this work was first published at an international conference in June 2008 (Hall *et al.*, 2008), the Mini Gamma-Ray Camera (MGRC) described in Section 2.4.6.3 has made use of the theory and further publications can be found on similar methods.
- Preliminary feasibility studies have shown how a structured (columnar) scintillator was essential to provide sufficient photons for a photon-counting detection technique.
- Simulations have been used to analyse event profiles with known interaction locations in the scintillator to allow the analysis of the depth-of-interaction effects which have plagued earlier studies.
- A simulator was developed for the whole scintillator-coupled EM-CCD, from the interactions in the scintillator and light spread to the noise and EM gain. This has provided

a tool which allows the investigation of the processes involved and allowed the eradication of several problems in the detection and analysis chain.

- Various centroiding techniques were analysed in detail and a novel simplified technique was developed to give the required sub-pixel positioning (within suitable system errors) with a very quick algorithm for coding in the data processing chain.
- The experimental set-up was designed to allow ease for future testing of different scintillators and gamma-ray sources (whether solid or at a synchrotron).
- The device was carefully characterised and optimised to achieve the best performance possible with the system, with simulations used to investigate unknown sources of noise that could then be corrected for.
- The standard 'oversampled edge' technique was used to determine the device MTF, with the method then adapted and developed to give much improved performance for analysis at very low signal-to-noise ratios.
- The basic resolution (FWHM) of the detector was improved from 83 µm in the integrated image case to 31 µm in the photon-counting image. This is a factor of two to three times better than results presented in the literature, showing that attention to detail in the system design and optimisation (through simulations and modelling of each component of the detection chain) and the development of event analysis techniques specific to this purpose can give large improvements in the system performance.
- The event analysis techniques developed in this study have allowed large steps to be taken towards energy discrimination in a photon-counting gamma-camera. The energy discrimination techniques developed in this study have allowed the removal of many self-fluorescence X-ray events, improving the spatial resolution (FWHM) further still to 25 µm, and, more dramatically, improving the MTF over all spatial frequencies (particularly the higher spatial frequencies). Many cameras described in the literature do not have

sufficiently low noise, high enough resolution or suitable analysis techniques to distinguish the fluorescence components in the image which have been shown to degrade the imaging capabilities of the detector.

This study paves the way for the development of a fully fledged high-resolution, high-energy, photon-counting gamma-camera for use in medical imaging and synchrotron based research (as outlined in Section 5.13).

The energy discrimination techniques developed in this study are already being implemented in the Leicester-based Mini Gamma-Ray Camera (CNRS Press Release, 2007) which is, at the time of writing, starting medical trials.

With appropriate funding and time allocated at a suitable synchrotron facility (such as the I12 at Diamond, due to open for general use in summer 2010), the developments outlined in Section 5.13 have the potential to lead to an energy-discriminating, high-energy, high resolution gamma-camera, giving a much improved resolution compared with cameras currently in production.

8.1.1 Research outcomes

Current gamma cameras based on photo-multiplier tubes have very strong spectral performance but are let down by their severely limited spatial resolution. Using an EM-CCD it is possible to improve the spatial resolution dramatically, in part down to the smaller pixel size. The improvements in spatial resolution do, however, come at a cost. The spectral resolution is severely affected by the increase in spatial resolution: it was found in this study that the improved spatial resolution allows the camera to resolve not only the incident X/γ-rays, but also the fluorescence X-rays from the caesium and iodine in the scintillator itself. With camera systems based on photo-multiplier tubes one includes the reabsorbed fluorescence X-rays within the primary interactions because of the limited spatial resolution. The ability to resolve the fluorescence X-rays with the scintillator-coupled EM-CCD, coupled with the ability to

distinguish the differences between the incident X/γ-rays and the fluorescence X-rays through energy discrimination, allows the spatial resolution to be improved further still through the exclusion of the escape electrons and the wider ranging fluorescence X-rays. Further improvements must focus on the improvement of the spectral resolution of the camera system in order to increase the efficiency of the energy determination and hence the separation of events into primary and secondary classes.

8.2 Instrument background

A pre-launch study of the background for XMM-Newton under-predicted the level of the instrument background continuum by nearly two orders of magnitude. This under-prediction was a consequence of the assumption that the instrument background was dominated by Compton-interactions resulting from secondary gamma-rays produced in the space-craft and detector material. Following the analysis of the background found in-orbit, it was deemed necessary to understand the real causes of the higher background to enable the reduction of the instrument background in future missions.

- A simulator was developed to aid in the investigation of the true causes of the instrument background and to analyse the individual components of the background spectrum.
- The simulated spectra are in very good agreement with the data obtained in orbit and lead to the conclusions presented below.
- The instrument background was found to be strongly dependent on both the detector design and the orbit of the mission (due to the varying spectra of incident radiation).
- The background spectral continua were found to comprise of three main components:
    - Compton scattering – More prominent in detectors with a thicker active layer due to the increased probability of interaction of the incident photon.

- Minimally ionising charged particles – The trail of charge following the passage of a minimally ionising charged particle leads to a peak in the spectrum at an energy determined by the detector thickness. Thinner detectors tend to pull the peak into the region of interest at lower energies. FI-CCDs do not suffer to the same magnitude as BI-CCDs due to charge spreading in the lower regions of the CCD (leading to spreading of signal over a non-X-ray-like grade pattern).
- Knock-on electrons – These provide the most dominant component of the instrument background spectra. When a high energy charged particle passes through the detector shielding many lower energy electrons are released. Although most of these electrons are absorbed in the shielding (due to the absorption lengths of only a few micrometers), those released close to the surface of the shielding may follow a path to the CCD and be detected as X-ray-like events near the detector surfaces. In this case, the surface structure of the detectors has a major impact on the level of the instrument background continuum.

This study has shown how the impact of the detector structure on the instrument background should aid in the development of detectors for future missions. Through a consideration of the total sensitivity of the device, one finds that a detector designed to purely optimise the quantum efficiency may not provide the greatest sensitivity due to the increased detection of the knock-on electrons.

There is great potential for future work to continue this instrument background study. The following section gives a brief overview of one possible direction for future study, with initial results showing a promising outlook.

8.2.1 Research outcomes

The pre-launch study for XMM-Newton under-predicted the continuum level of the instrument background. The study assumed that the Compton Effect was the dominant production

mechanism for the background continuum. This thesis asked the question as to why the instrument background was found to be much higher in reality for the mission when in-orbit. It was found, through the use of the simulations detailed in Chapters 6 and 7, that the instrument background is in fact dominated, particularly at HEO, by the detection of knock-on electrons produced from the interactions of the incident protons with the shielding surrounding the detector. This study therefore concludes that the pre-launch studies were found to under-predict the background continuum through the use of an inaccurate assumption (that the continuum was dominated by the Compton Effect) and that the continuum can in fact be well simulated if all components able to generate the instrument background are included.

8.2.2 Future work

Any method which might reduce the instrument background without a significant reduction of the desired signal is worth investigating in more detail. One possible method for reducing the background is the use of Z-graded shielding. This method does not affect the desired signal but may significantly reduce the instrument background.

8.2.2.1 Z-Graded shielding

A study of Z-graded shielding was published in Pfeffermann *et al*. (2004). The following paragraph is given as a summary of this paper, followed by an extension to the work implemented in this study.

A 1 mm layer of aluminium oxide is intransparent to X-ray photons below approximately 10 keV, see Figure 8.1. Such a layer would significantly reduce the background from fluorescence photons from outside of this layer. The fluorescence X-ray photons from the aluminium oxide layer lie below 2 keV and would therefore be detected as X-ray like signal unless otherwise absorbed. A 1 mm layer of boron carbide is intransparent to X-ray photons below approximately 4 keV, see Figure 8.1. If placed onto the aluminium oxide layer, the

fluorescence X-ray photons from the aluminium oxide layer would be absorbed. The fluorescence yield of carbon and boron are of the order of 0.1 percent. The majority of the energy re-emitted from the boron carbide is in the form of Auger electrons of energy below 280 eV. These electrons can be absorbed in a layer of just several nanometres of Si or SiO_2 and would therefore be absorbed in the passivation layer of the CCD. However, any impurities in the materials would lead to an increase in the number of different fluorescence lines. Although a reduction in the strength of the fluorescence lines is clear from the above, the effect of the Z-graded shielding on the continuum is not covered in the study.

The geometry of the simulation for XMM-Newton was altered to include a 1 mm layer of Al_2O_3 inside the aluminium shell with a further 1 mm layer of B_4C inside this. A significant reduction by a factor of approximately 2-3 in the background continuum was seen in the initial simulated results, although further work is required to confirm this result with a full implementation of the fluorescence. Electrons ionised from the inside of the aluminium shell will be absorbed within the 1 mm layer of aluminium oxide, as from the aluminium oxide to the boron carbide, so only those electrons ionised on the inner surface of the boron carbide layer can be detected.

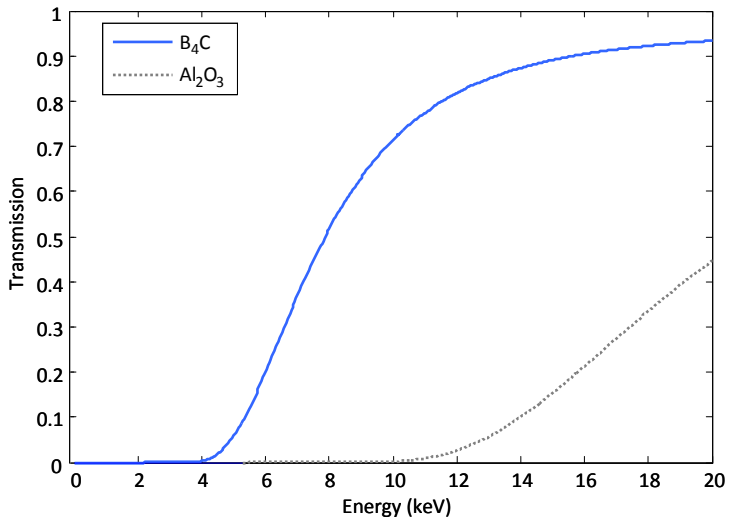

Figure 8.1. Transmission for 1 mm layers of boron carbide and aluminium oxide. The aluminium oxide is intransparent to X-rays below approximately 10 keV and the boron carbide below approximately 4 keV. The boron carbide layer is therefore opaque to the fluorescence photons from the aluminium oxide.

Although the B_4C and Al_2O_3 layers appear to reduce the level of the continuum of the background, the effect of the B_4C on the fluorescence peaks also has to be taken into consideration. B_4C and Al_2O_3 were chosen for their positive effects on the background from fluorescence, but when introducing new materials one has to be careful not to also introduce impurities to the system. Impurities in the B_4C particularly will show up in the spectrum as further fluorescence peaks and will reduce the advantage gained from any reduction in the background continuum.

For the case of the Suzaku mission, Murakami *et al.* (2006) carried out tests using multi-layer shielding consisting of beryllium and tungsten. They found a reduction in intensity of X-rays in the energy band 10-500 keV to approximately one third, but approximately double the intensity of X-rays in the energy band of 2-8 keV. They state that the total background will be reduced as long as the background is dominated by the Compton scattering. The impact of an Al_2O_3 and B_4C layer structure is not covered here for the incident particle spectrum at the Suzaku orbit.

Bibliography

1. Agostinelli, S., Allison, J., Amako, K., Apostolakis, J., Araujo, H., *et al*., "G4—a simulation toolkit", *Nuclear Instruments and Methods in Physics Research Section A,* **506**(3), (2003) 250-303.

2. Anzai eZ-SCOPE AN - Handheld Gamma Camera, Nucle Med (accessed 5th June 2010): http://www.nuclemed.be/product.php?cat=149&prod=520

3. Atkins, B., Optical Articles, Bob Atkins Photography, 2008 (accessed 5th June 2010): http://www.bobatkins.com/photography/technical/mtf/mtf1.html.

4. Australian Synchrotron website: www.synchrotron.org.au (accessed 5th June 2010).

5. Babaud, J., Witkin, A. and Duda, R., "Uniqueness of the Gaussian kernel for scale-space filtering", *IEEE Trans. Pattern Anal. Machine Intell.* **8**(1), (1986) 26–33.

6. Barbe, D. F., "Imaging devices using charge coupled devices", *Proc. IEEE* **63**(1), (1975), 38-67.

7. Barrett, H., "Objective assessment of image quality: effects of quantum noise and object variability", *J. Opt Soc Am A* **7** (1990) 1266-1278.

8. Barth, J., "The Radiation Environment", *GSFC*, (1999).

9. Beekman, F., J., and de Vree, G., A., "Photon-counting versus an integrating CCD-based gamma camera: important consequences for spatial resolution", *Phys. Med. Biol.* **50**, (2005) N109-N119.

10. Bruns, B., private communication, e2v technologies.

11. Burt, D., "Introduction to CCD Technology", presentation series (2006).

12. Burt, D., e2v technologies plc., private communication.

13. Byoung-Jik Kim, Gyuseong Cho, Bo Kyung Cha, and Bosun Kang, "An X-ray imaging detector based on pixel structured scintillator", *Radiation Measurements* **42**(8), (2007) 1415-1418.

14. CNRS Press Release, "A mini camera for the surgical treatment of breast cancer", (Paris, June 1, 2007).

15. Diamond website: www.diamond.ac.uk (accessed 5th June 2010).

16. Drouin, D., Couture, A. R., Joly, D., Tastet, X., Aimez, V., and Gauvin, R., "CASINO V2.42 - A Fast and Easy-to-use Modeling Tool for Scanning Electron Microscopy and Microanalysis Users", *Scanning (Wiley)* **29**(3), (2007) 92-101.

17. Dyer, C. S., Truscott, P. R., Evans, H. E., and Peerless, C. L., "Simulation of spacecraft secondary particle emissions and their energy deposition in CCD X-ray detectors", *IEEE Transactions on Nuclear Science* **43**(6), (1996) 2709-2714.

18. e2v CCD201-20 datasheet, A1A-100013 Issue 1 (2005).

19. e2v CCD47-20 data-sheet, A1A-100039 Issue 5 (2006).

20. e2v CCD97 data-sheet, A1A-CCD97BI_2P_IMO Issue 3 (2004).

21. EL7202 datasheet, Intersil Americas Inc., FN7282 (2006).

22. Elster, J. and Geitel, H. F., "About the scintillating phosphorence of the Sitoblende excited by radioactive emanations", *Physikalische Zeitschrift* **4**, (1903), 439.

23. ESA XMM-Newton website (accessed 5th June 2010): http://sci.esa.int/science-e/www/area/index.cfm?fareaid=23

24. ESA IXO website: http://sci.esa.int/science-e/www/area/index.cfm?fareaid=103 (accessed 5th June 2010).

25. Francis, A., “Advanced Level Statistics”, *ST(P) Ltd* (1979).

26. Freyberg, M. J., Pfefferman, E., and Briel, U. G., “The XMM-Newton EPIC pn camera: spatial distribution of the internal background fluorescence lines”, *ESA SP* **488** (2002).

27. Freyberg, M. J., Briel, U. G., Dennerl, K., Haberl, F., Hartner, G., Pfeffermann, E., Kendziorra, E., Kirsch, M., and Lumb, D., “The EPIC pn-CCD detector aboard XMM-Newton: status of the background calibration”, *SPIE* **5165**, (2004), 112-122.

28. Graeve, T., and Weckler, G.P., “High-resolution CMOS imaging detector”, *Proc. SPIE* **4320**, (2001), 68-76.

29. Guainazzi M., *et al.*, “EPIC status of calibration and data analysis”, XMM-SOC-CAL-TN-0018-2-7-1, 2008 (last update).

30. Hall, D. J., Holland, A., and Turner, M., “Modelling instrument background for CCD x-ray spectrometers in space”, *Proc. SPIE* **6686**, (2007), 66860J.

31. Hall, D. J., Holland, A., and Turner, M., “Simulating and reproducing instrument background for x-ray CCD spectrometers in space”, *Proc. SPIE* **7021**, (2008), 70211Y.

32. Hall, D. J., Holland, A. and Smith, D.R., “Imaging and spectroscopy using a scintillator-coupled EMCCD”, *SPIE* **7021**, (2008), 70211Z.

33. Hall, D. J. and Holland, A., “High resolution X-ray and γ-ray imaging using a scintillator-coupled Electron-Multiplying CCD”, *Proc. SPIE* **7449**, (2009), 74491G.

34. Hall, D. J., Holland, A. and Smith, D. R., “The use of automatic scale selection to improve the spatial and spectral resolution of a scintillator-coupled EMCCD”, *NIM A* **604**, (2009) 207-210.

35. Hall, D. J., and Holland, A., “Space radiation environment effects on X-ray CCD background”, *Nuclear Inst. and Methods in Physics Research: A* **612**, (2010), 320-327.

36. Hamamatsu data-sheet, H8500C_H8500D_TPMH1308E01, (2007)

37. HCPL 2630 datasheet, Fairchild Semiconductor Corporation, DS300202 (2001)

38. Heemskerk, J., W., T., Westra, A., H., Linotte, P., M., Ligtvoet, K., M., Zbijewski, W. and Beekman, F., J., "Frontilluminated versus back-illuminated photon-counting CCD-based gamma camera: important consequences for spatial resolution and energy resolution," *Phys. Med. Biol*. **52**, (2007), N149-N162.

39. Heynderickx, D., Quaghebeur, B., Speelman, E., and Daly, E., AIAA 2000-0371, *38th Aerospace Sciences Meeting and Exhibit, Reno, NV*, SPENVIS (2000).

40. Hong, T. H,. Shneier, M., and Rosenfeld, A., “Border extraction using linked edge pyramids”, *IEEE Trans. Systems Man Cybernet* **12**(5), (1982), 660–668.

41. Janesick, J., “Scientific Charge-Coupled Devices”, *SPIE Press*, 2001.

42. Jerram, P., Pool, P., Bell, R., Burt, D., Bowring, S., Spencer, S., Hazelwood, M., Moody, I., Catlett, N., and Heyes, P., "The LLLCCD: Low light imaging without the need for an intensifier," *Proc. SPIE* **4306** (2001) 178-186.

43. Johnson, J. B., “Thermal agitation of electricity in conductors”, *Physical Review* **32**, (1928), 97-109.

44. Kuhl, D. E., and Edwards, R. G., “RQ: Image Separation Radioisotope Scanning”, *Radiology* **80**, (1963), 653-662.

45. Lifshin, E., “X-ray characterisation of materials”, Wiley (1999).

46. Lucky website (accessed 5th June 2010): http://www.ast.cam.ac.uk/~optics/Lucky_Web_Site/LI_CCD%20Cameras.htm.

47. Lindeberg, T., “Scale-space Theory In Computer Vision”, Springer (1993).

48. Lindeberg, T., “Scale-space theory: A basic tool for analysing structures at different scales”, *J. of Applied Statistics* **21**(2), (1994), 224—270.

49. Lindeberg, T., “Scale-Space Theory in Computer Vision”, *Kluwer Academic Publishers* (1994).

50. Lindeberg, T.,: "Scale-space: A framework for handling image structures at multiple scales", *Proc. CERN School of Computing, Egmond aan Zee, The Netherlands* (1996).

51. Loudos, G., *et al.*, Comp*. Med. Imaging & Graphics* **27**, (2003), 307.

52. Low-Light Technical Note 4, “Dark Signal and Clock-Induced Charge in L3VisionTM CCD Sensors”, A1A-Low-Light TN4 Issue 2, June 2004.

53. De Luca, A., and Molendi, S., “Studying the cosmic X-ray background with XMM-Newton”, *Proc. Symposium ‘New Visions of the X-ray Universe in the XMM-Newton and Chandra Era’*, 26-30 November, (2001).

54. Lumb, D. H., Warwick, R. S., Page, M., and De Luca, A., “X-ray background measurements with XMM-Newton EPIC”, *A&A* **389**, (2002), 93-105.

55. Mackay, C., D., Tubbs, R., N., Bell, R., Burt, D., Jerram, P., Moody, I., "Sub-electron read noise at MHz pixel rates," *Proc. SPIE* **4306**, (2001), 289-98.

56. Marsland J. S., “Temperature dependence of ionisation coefficients in silicon derived from physical model”, *Electron. Lett*. **27**, (1991).

57. Mead, R. W., Collimate Holes Inc., private communication.

58. Miller, B. *et al.*, “ Progress of BazookaSPECT“, *SPIE Proc. Optics & Photonics 2009*, **7450-24** (2009)

59. Miller, B. W., Barber, H. B., Barrett, H. H., Wilson, D. W., and Liying C., “A Low-Cost Approach to High-Resolution, Single-Photon Imaging Using Columnar Scintillators and Image Intensifiers”, *IEEE Nuclear Science Symposium Conference Record* **6**, (2006), 3540-3545.

60. Miller, B. W., Barber, H. B., Barrett, H. H., Shestakova, I., Singh, B., and Nagarkar, V. V., "Single-photon spatial and energy resolution of a columnar CsI(Tl)/EMCCD gamma-camera using maximum-likelihood estimation," *SPIE* **6142**, (2006), 61421T.

61. Miller, B., W., Barber, H., B., Barrett, H., H., Chen, L. and Taylor, S., J., "Photon-counting gamma camera based on columnar CsI(Tl) optically coupled to a back-illuminated CCD", *SPIE* **6510**, (2007), 65100N.

62. Mizuno, T., Kamae, T., Godfrey, G., Handa, T., Thompson, D. J., Lauben, D., Fukazawa, Y., and Ozaki, M., “Cosmic-Ray Background Flux Model Based on a Gamma-Ray Large Area Space Telescope Balloon Flight Engineering Model”, *Astrophysical Journal* **614**, (2004), 1113–1123.

63. Murakami, H., Kitsunezuka, M., Ozaki, M., Dotani, T., and Anada, T., “Origins of the instrument background of the X-ray CCD camera in space studied with Monte Carlo simulation”, *Proc. Space Telescopes and Instrumentation II: Ultraviolet to Gamma Ray* **6266***,* (2006), 62662Y.

64. Murray, N. J., Holland, A. D., Smith, D. R., Gow, J. P., Pool, P. J., and Burt, D. J., “The X-ray quantum efficiency measurement of high resistivity CCDs”, *Nuclear Instruments and Methods: Section A* **604**(1-2), (2009), 180–182.

65. Nagarkar, V.V., Shestakova, I., Gaysinskiy, V., *et al*., “Fast X-ray/γ-ray imaging using electron multiplying CCD-based detector”, *Nuclear Instruments and Methods in Physics Research Section A* **563**(1) (2006) 45-48.

66. NASA Suzaku website (accessed 5th June 2010): http://www.nasa.gov/mission_pages/astro-e2/main/index.html

67. NASA Swift website: http://www.nasa.gov/mission_pages/swift/main/index.html (accessed 5th June 2010).

68. NASA website, SAA image (accessed 5th June 2010): http://imagine.gsfc.nasa.gov/docs/ask_astro/answers/961004.html

69. Ness, J. U., Ehle, M., González Riestra, R., Guainazzi, M., Rodríguez, P., Talavera, A., and de la Calle, I., XMM-Newton Users Handbook, Issue 2.7 (17/09/2009)

70. Pagani, C., Penn State University, private communication.

71. Pani, R., Bennati, P., Betti, M., *et al.*, "Lanthanum scintillation crystals for gamma ray imaging", *Nucl. Instr. and Meth. A* **567**(1), (2006), 294.

72. E. Pfeffermann, P. Friedrich, M. Freyberg, G. Kettenring, L. Krämer, N. Meidinger, P. Predehl and L. Strüder, "Shielding of cosmic ray induced background in CCD detectors for X-ray astronomy", *Proc. High-Energy Detectors in Astronomy* **5501**, (2004).

73. Phillips, J. C., Wlodawer, A., Yevitz, M. M. and Hodgson, K. O., "Applications of synchrotron radiation to protein crystallography: preliminary results", *PNAS* **73**(1), (1976), 128-132.

74. Read, A. M., and Ponman, T. J., "The XMM-Newton EPIC Background: Production of Background Maps and Event Files", *A&A* **409**, (2003), 395.

75. Robbins, M., S. and Hadwen, B., J., "The Noise Performance of Electron Multiplying Charge-Coupled Devices", *IEEE Transactions On Electron Devices* **50**(5), (2003), 1227-1232.

76. Rodnyi, P., "Physical processes in inorganic scintillators", *CRC Press*, (1997).

77. Sentinella Suite, Onco Vision, Gem-imaging group (accessed 5th June 2010): www.gem-imaging.com/productos/sentinella102/sentinella102.php?idm=en

78. Shestakova, I., Nagarkar, V.V., Gaysinskiy, V., Entine, G., Stack, B.C., Miller, B., *Proc. SPIE* **6319**, (2006), 63191E.1.

79. Smith, D. R., and Holland, A. D., "Proton irradiation of e2v technologies L3 Vision devices", *Nuclear Instruments and Methods: Section A* **546**, (2005), 544-52.

80. SRIM website: www.srim.org (accessed 5th June 2010).

81. Sze S.M., "Physics of semiconductor devices", Wiley (1981).

82. Tipinis, S.V., Nagarkar, V.V., Shestakova, I., Gaysinskiu, V., Entine, G., Tornai, M.P., Stack, B.C., "Feasibility of beta-gamma digital imaging probe for radioguided surgery", *IEEE Trans Nuclear Sci.* **51**, (2004), 110-116.

83. Tsuchimochi, M., "Small high resolution Gamma Camera for Sentinel node biopsy" *Clinical Case Studies in the Third Dimension* **2**, (2006), 275-282.

84. Turner, M. J. L., *et al.*, "The European Photon Imaging Camera on XMM-Newton: The MOS cameras", A&A **365**, (2001), L27-L35.

85. de Vree, G., A., van der Have, F. and Beekman, F., J., "EMCCD-based photon-counting mini gamma camera with a spatial resolution < 100 µm", *IEEE Nuclear Science Symposium Conference Record* **5**, (2004), 2724-2728.

86. Witkin, A., "Scale-space filtering: A new approach to multi-scale description", *IEEE International Conference on Acoustics, Speech, and Signal Processing* **9**, (1984), 150-153.

87. XCAM USBREM2 – Scientific CCD Camera Drive System (accessed 5th June 2010): www.xcam.co.uk/products/electronics.html

88. Yamaguchi, H., Nakajima, H., Koyama, K., Tsuru, T. G., Matsumoto, H., Tawa, N., Tsunemi, H., Hayashida, K., Torii, K., Namiki, M., Katayama, H., Dotani, T., Ozaki, M., Murakami, H., and Miller, E., "The background properties of Suzaku/XIS" *SPIE Proc. Int. Soc. Opt. Eng.* **6266**, (2006), 626642.

89. Ziegler, J. F., Bierzack, J. P., and Ziegler, M. D., "The Stopping and Range of Ions in Matter", Lulu Press, Morrisville, North Carolina (2008).