Open Research Online

Imaging and spectroscopy using a scintillator-coupled EMCCD

Conference or Workshop Item

oro.open.ac.uk

Imaging and spectroscopy using a scintillator-coupled EMCCD

David J. Hall[*a], Andrew Holland[b] and David R. Smith[a]

[a] Imaging for Space and Terrestrial Applications, School of Engineering and Design,
Brunel University, Uxbridge, Middlesex, UB8 3PH, UK

[b] Planetary and Space Sciences Research Institute, Open University,
Walton Hall, Milton Keynes MK7 6AA, UK

ABSTRACT

The CCD97 is a low light level (L3) device from e2v technologies range of electron multiplying CCDs (EMCCDs). The device uses e2v's patented extended gain register and through the use of appropriately designed electrodes can be used to maximise the signal whilst keeping the impact of the noise to a minimum. The nature of this device makes it ideal for use with a scintillator in order to see individual flashes of light from single X-ray photons. Through the examination of individual X-ray events, it is possible to analyse each interaction in the scintillator to determine the sub-pixel position of the interaction. Using the modelling capabilities of the Geant4 toolkit it is possible to simulate X-ray events and thus examine interactions with known energy and point of interaction. Through bringing together the experimental and simulated results, the spectral capabilities of such a device are discussed.

1. INTRODUCTION

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, "impact ionisation" or "avalanche multiplication" can occur. Through this process, charge carriers (electrons) can lose energy by the creation of more charge carriers, thus multiplying the signal before the readout amplifier, leading to a better signal to noise ratio. In the early 1990s, e2v carried out experiments to show that this multiplication in the signal could be used in CCDs, spawning the "low-light-level" (L3) devices, Electron Multiplying CCDs (EMCCDs), which are now available [3,7]. The devices are ideal for use in conditions where the incident photon flux is low [11,13].

With the EMCCD's ability to work in low-light conditions, it seems to be the perfect candidate for use with a scintillator. When X-rays interact in the scintillator, approximately fifty-four thousand 550nm photons are produced per MeV of incident X-ray. Some of these photons will travel towards the CCD surface and a small 2D Gaussian-like signal (a photon flash) will be detected. With a standard CCD, this peak can be lost under the noise, but with the improved signal to noise ratio of the EMCCD device it is possible to detect these peaks.

In order to reduce the readout noise, the device requires precise optimisation. With appropriate consideration the effective noise can be lowered to the sub-electron level. Following this, the device requires calibration such that one can determine the readout noise and the calibration factor for the device (electrons per digital number). The device optimisation and calibration techniques are described briefly and the results discussed.

The incident X-rays may interact at any depth inside the scintillator and as such there are several unknown variables involved in examining the individual photon flashes, namely the change in spread and peak intensity of the flash due to depth of interaction effects and fluctuations in the gain inherent in the avalanche gain process. These variables have on the whole been ignored in the past, with basic methods such as image smoothing on a single level sometimes used to correct for the errors introduced [1,2,10,14]. Investigating these variables experimentally is a complicated process and results achieved in this way show only small improvements on the more basic methods [9].

Further author information: send correspondence to David Hall, david.hall@brunel.ac.uk +44 (0)1895 266 521
Address: School of Engineering and Design, Brunel University, Uxbridge, Middlesex, UB8 3PH, United Kingdom

This system crosses several boundaries between different areas in physics. When a raw image is presented, the data hidden inside the often random looking signal must be extracted. This, in many ways, is not dissimilar to point-source extraction from extragalactic point source searching. The method introduced here builds upon work discussed by Tony Lindeberg in various publications [4-6] based on automatic scale selection in scale-space.

Here, a more detailed correction for the depth of interaction effects is discussed. Through the use of the Geant4 toolkit it is possible to study optical interactions with matter and a scintillator can be simulated. The specific parameters for the scintillator are coded by the user in this case and so any scintillator can be modelled. The structure of the materials involved can also be freely designed using more custom C++ code. The toolkit allows the user to both investigate X-ray interactions with a scintillator and to investigate the photon flash initial conditions inside the scintillation layer. This freedom coupled with the precise nature in which the analysis can be designed gives the user the ability to investigate the system to a more detailed level then can be achieved experimentally.

Such a system has the potential for use in a variety of applications, namely in small animal Single Photon Emission Computed Tomography (SPECT) [10] in which, through the use of different isotopes, a multi-label imaging system could be developed. The main aims of this study are to highlight methods by which the imaging and spectrographic qualities of a scintillator-coupled EMCCD can potentially be improved through the use of novel image analysis techniques and the use of Geant4 as an investigative tool.

2. THE ELECTRON MULTIPLYING CCD

The EMCCD is similar to a standard frame transfer CCD except for the addition of a 'gain' register. Signal is built up in the image section by the conversion of photons to electrons. These electrons (the signal) are then quickly transferred to a shielded store section. The signal is then transferred line by line to the readout register to be transferred to the CCD output node. The EMCCD contains a standard linear readout register followed by an additional linear 'gain' register. In this gain register, the signal undergoes impact ionisation or avalanche multiplication.

2.1 Impact ionisation

Avalanche multiplication occurs when a charge moves through a region of silicon in which a high electric field is present. One phase in each gain register pixel is held at approximately 40 V. This is variable to give a gain from 1 to approximately 1000. The large change in voltage over this short distance accelerates the electrons to a much greater kinetic energy. These higher energy electrons can then impart enough energy when colliding with the lattice to free additional electrons from the silicon atoms. The electrons released in this process may then also undergo avalanche multiplication, leading to a progressive increase in the signal over the extent of the gain register. Whilst the probability, P, of multiplication over each element of the readout register is quite small, the number of gain register elements, n, is generally quite large ($n = 536$ with the CCD97) and so the multiplication gain can be of the order of thousands, Equation 1. This multiplication gain is applied to the signal before the output circuitry of the CCD and so any electronic or processing noise will only be added to the multiplied signal, greatly improving the signal to noise.

$$\text{total gain} = (1 + P)^n \tag{1}$$

2.2 Limitations for photon-counting

The total gain quoted above is only a mean value of the gain achieved. In reality, the number of electrons generated through the impact ionisation process shows considerable fluctuation. It was shown in [12] that the noise in the signal was increased by a factor of $\sqrt{2}$ for high gain. This means that it is not possible to distinguish between very small signals of a few electrons and true 'photon-counting' cannot be achieved.

In previous studies of this nature it is often the case that the peak signal (or sum of 3 × 3 pixels) of a photon flash is assumed to be proportional to the energy of the incident X-ray [1,2,10]. Ignoring the depth of interaction effects (discussed later) this would be the case for unity gain. Unfortunately, at unity gain, the noise is too high to distinguish these photon flashes which have an intensity of only a few photons per pixel. At the higher gain required, the shot noise and multiplication noise factor mean that there is a theoretical limit on the energy resolution (noting that this is still ignoring depth of interaction effects and any losses in the system). This energy resolution limit in electrons is given in Equation 2.

$$\frac{\text{noise}}{\text{signal}} = \frac{\sqrt{2 \times \text{signal}(e^-)}}{\text{signal}(e^-)} = \sqrt{\frac{2}{\text{signal}(e^-)}} \qquad (2)$$

If the energy of the incident X-ray increases, the number of photons generated in the scintillator increases, leading to an increase in the electrons produced in the CCD and a better energy resolution. At a level of 10 to 100 photons per pixel as expected in the experimental arrangement described below, this leads to a fractional error (as defined by Equation 2) of between 0.44 and 0.14 respectively. These values cannot be improved upon and depth of interaction effects increase these errors substantially. For these reasons, using the peak intensity from the raw image will not give desirable spectroscopic results. If, however, these errors in pixel signal values can be averaged out, the overall shape of the original photon flash can be retained. Using curves fitted to data may dramatically decrease the errors introduced through shot noise and the fluctuations in the avalanche multiplication. In this way, photon counting for the purposes described here becomes a possibility.

For imaging, the central point of the interaction is used to represent the incident X-ray in the final image. Although the scaling in intensity of the points requires deeper analysis, the positioning in a single plane of these points is not greatly effected by the noise in the signal output.

2.3 Experimental system

The CCD97 has an image area consisting of 512 × 512 pixels, each with a square side of 16 μm. The CCD97 also contains 24 dark reference columns along with 16 additional dark reference rows. During readout, a further 16 overscan columns are added on each side of the image due to the overscan elements in the readout register. This gives a final image size of 558 × 528 pixels. Although the image area is still only 512 × 512 pixels, the dark reference and overscan pixels do provide some useful information to aid in the validation of results (see Sections 5.4 and 5.5).

A fibre-optic plate (FOP) is bonded to the surface of the CCD, as shown in Figure 1. A thin metal plate is then placed around the FOP to protect the surface of the chip and the bond wires. The scintillator is grown as columns on a second fibre-optic plate, Figure 1, which can then be the clamped against the CCD-FOP combination. This allows the scintillator to be changed without the requirement for a new CCD.

The inclusion of the fibre-optic plate is dual purpose. Firstly, it allows a simple method by which to change the scintillator. The second purpose of the FOP is to prevent direct detection of X-rays in the CCD. Without the FOP, any X-rays which pass through the scintillator could be detected in the CCD and could cause complications in the image analysis. The FOP blocks the X-rays and, at a few mm thick, has a transmission of a fraction of a percent at 60 keV.

The device is run in inverted mode to dramatically reduce the dark current. With the CCD97 it is essential that the readout registers are not inverted as this causes excess noise in the multiplication process. In order to further reduce the dark current the device is cooled to between -10°C and -15°C using a Peltier cooler.

The X-ray source used to produce the results shown in this paper is an americium-241 source from Amersham International plc with an activity of 5.55 kBq. The Am-241 source has two X-ray emissions at 14 keV and 60 keV. These emissions will produce approximately 750 and 3250 photons respectively at around 550 nm (see Section 4.1 below). Half of these photons will immediately be lost by travelling away from the FOP. The number of photons per pixel from a 60 keV interaction peaks in the order of tens of photons.


Figure 1. The CCD, fibre-optic plate (FOP) and scintillator arrangement (left). The scintillator structure can be seen as columns of width a few μm (right). Photograph supplied by Bill Bruns, e2v technologies.

3. DEVICE OPTIMISATION AND CALIBRATION

The main reason for using an EMCCD rather than a conventional CCD is to reduce the effective readout noise to the sub-electron level in order to maximise the signal-to-noise ratio. The noise on the dark current can be minimised by simply cooling the device to reduce the dark current and by running in inverted mode.

3.1 Effective noise

The readout noise in the device is primarily defined by the noise of the on-chip amplifier and noise added to the signal through the processing electronics. This is specific to the CCD and system used but, fortunately here, this noise is added to the signal at the point of leaving the gain register (unlike the shot noise on the dark current which is added to the signal before and during readout). If the signal is multiplied during readout by avalanche multiplication then noise is added to the already multiplied signal. The signal level can be taken to be effectively the same as that leaving the image region and so the effective readout noise has been reduced, Equation 3, where $\sigma*$ is the effective noise (e$^-$ rms) and cal is the calibration factor in electrons per digital number. If the noise is known with no avalanche gain then the gain required for the signal levels being investigated can be determined using Equation 3. Minimising the noise is a complicated process which is not covered here. The process is system and device specific and is documented in the appropriate device literature.

$$\sigma*\left(e^-\ rms\right)=\frac{\sigma(DN)}{cal\left(e^-/DN\right)} \qquad (3)$$

3.2 Device calibration

The Photon Transfer Curve (PTC) is a remarkably useful tool which can provide various details on the characteristics of a particular device. The probability that exactly n emission events from a photon source occur inside a given time period T is given by the Poisson distribution function. The variance of the Poisson distribution is equal to the mean, and consequently the standard deviation is equal to the square root of the mean. These facts can be used with the PTC to derive some characteristics of the device. The PTC shown in Figure 2 is shown on a logarithmic scale such that the intersection of a line with a gradient of one fitted to the data with a line of variance one provides the calibration of the device. The error on this calibration value can then be determined from the spread.


Figure 2. The photon transfer curve from the initial testing of the CCD97 [15] obtained using the XCAM ltd. drive system [16] with an avalanche gain of unity.

The benefits of lowering the noise become apparent when considering a simple drop from 3 electrons rms to 1 electron rms. For a signal-to-noise ratio of 5:1, which may be desirable to distinguish the signal from the background, the lower background allows signals to be visible from 5 electrons, whereas the higher noise level only allows signals above 15 electrons to be detected. Such a difference can have a large impact on the detected signal at the photon-counting level.

The relationship between the calibration of the device and the gain applied is shown in Equation 4, where *cal*(G)* is the effective calibration at a gain of *G*. By increasing the avalanche gain the electron signal per DN can be reduced, thus increasing the photonic resolution of the device.

$$cal*(G) = \frac{cal(G=1)}{G}$$ (4)

4. THE SCINTILLATOR

When high energy ionising electromagnetic radiation interacts with a scintillator the incident energy can be absorbed and subsequently fluoresced in the form of photons with longer wavelengths. In this way, scintillators can be used in many areas of physics to allow the detection of electromagnetic waves or particles through the use of photomultiplier tubes or CCDs. In this study, initial results are based around the use of structured thallium doped cesium iodide (CsI:Tl).

4.1 Cesium Iodide (Thallium doped)

Cesium iodide is a relatively soft ionic compound which can be used to convert X-ray photons into lower energy photons in the visible spectrum. The material is slightly hygroscopic and as such, a thin protective film is used to prevent contact with water. CsI:Tl is one of the brightest of the more commonly used scintillators. The peak of the broad emission falls at 550nm. This wavelength is very well suited to the CCD97 which has a quantum efficiency of over 90% at this wavelength (at -20 °C). In fact, the CCD97 has a quantum efficiency of over 80% across the majority of the broader spectrum of photon emission, making CsI:Tl a prime candidate for coupling with such a device.

The decay times of CsI:Tl, although a relatively slow scintillator, are short in comparison to the image integration times of relevance in this study. The scintillator output consists of several components, ranging from 0.6 μs to 3.5 μs. With integration times considered here of the order of tens of milliseconds, the decay times should not cause difficulties.

The output of CsI:Tl is dependent on the energy of the incident electromagnetic radiation of the particle, with a light yield of approximately 54 photons per keV. The scintillator output exhibits some temperature dependence. The relative light output of the scintillator varies approximately linearly in the region of 0 °C to -30 °C over a range of approximately 15 %. A variation in detector temperature within approximately 5 °C should, however, cause little deviation in output.

4.2 Structured scintillator

Light spread from the point of interaction within a solid block of scintillator material will be large (on a similar scale to the distance between the interaction and the fibre optic plate surface). In order to achieve a reasonable efficiency at higher X-ray energies (the ratio of X-rays interacting in the scintillator compared with those passing straight through being acceptable) a thicker scintillating layer is required. The spread of light from the interaction point increases as the scintillation thickness increases and hence the light detected per pixel in the CCD is rapidly reduced. In order to combat this problem, the CsI:Tl is grown in a specific way to create a collimated structure, Figure 1. The CsI:Tl produced consists of tightly packed columnar structures with diameters of the order a few microns. These columns act in a similar way to a fibre-optic plate, with light effectively being channelled from the point of interaction to the upper and lower surfaces of the scintillator. The process is by no means perfect, but does act to decrease the spread of the emitted photons and hence increase the peak signal detected at the CCD.

The quantity of photons emitted is dependent on the energy of the interacting radiation or particle. Due to the spreading of the light in the scintillator structure (although minimised through the structuring of the scintillator) the vertical position of the interaction has an effect on the spread and peak intensity of the photon flash detected within the CCD (as found in the simulations detailed in Section 8 of this paper and in [9]). The aim of this work is to use these depth effects to trace back from the intensity and spread of the individual photon flashes to gain a fuller picture of the interaction, allowing the position, and potentially the energy, of the incident X-ray to be determined.

5. INITIAL IMAGE PROCESSING

There are two stages to the initial image processing: removing the background and 'cleaning' the image, Figure 3. Before carrying out the background subtraction a region of four very bright pixels is removed from each image and replaced with an average value from three surrounding pixels.

```
┌─────────────┐   ┌─────────────┐   ┌─────────────┐   ┌─────────────────────┐
│  Raw image  │──▶│ Remove flat  ╲  │Remove isolated╲ │Image ready for analysis│
│             │   │ background    ╱ │ bright pixels  ╱│                     │
└─────────────┘   └─────────────┘   └─────────────┘   └─────────────────────┘
```

Figure 3. The processes involved in the initial image processing.

5.1 Background removal

The source used to provide the X-rays is fixed in position. In order to remove the source to take a standard background image the system must be opened, thus changing the conditions inside the chamber, such as the temperature, leading to an increase in the errors in the background removal.

One method used to remove the background involves summing surrounding pixels and subtracting the mean value from all pixels in the image. By taking the average pixel value in a 40×40 pixel area surrounding the pixel of note, the impact of an event local to the pixel is minimised. Using this method of background subtraction, a comparatively large amount of processing time is consumed in order to produce a background image to be subtracted which only has a standard deviation of 0.6124 ADU. The shot noise in any events of even the lowest nature puts this into perspective and the variation over this background can be regarded as negligible. For this reason, a simple average over the entire image is taken (assuming the density of photon flashes is few enough to avoid any noticeable change to the mean) and subtracted from the image to set the zero signal level.

5.2 Cleaning the image

Initial cleaning by removal of consistently bright pixels will clean the image to a lesser degree, but from image to image there will be cases where a pixel may have a larger 'signal' value due to the random nature of the noise. In the images taken from the CCD there are often bright pixels which may be confused with signal when using algorithms to analyse the signal (depending on the complexity of the algorithm). Single bright pixels are therefore removed using a single pixel scrub. This assumes that signal from photon flashes will be spread over several pixels, which is the case in reality.

The pixels are run through sequentially from top left corner to bottom right corner along each row. If a pixel has a signal level above the designated reject level whilst the eight surrounding pixels are below the reject level then this pixel will be replaced with an average of the surrounding eight pixels. Two passes are used, firstly to remove single pixel events with a high reject level, and secondly to remove those closer to the background level.

5.3 Smoothing

One method of analysis which is often used in the literature [1,2,14] is to 'smooth' the image. This generally involves convoluting the image with a Gaussian kernel. Although, to the eye, this emphasises Gaussian-like signal in the images, the shape and intensity of any photon flashes are altered through the convolution. Unfortunately, the variation in the depth of the X-ray interaction in the scintillator leads to a different spreading of the photon flash and hence any simple smoothing operation will affect each photon flash in a different way. This then has the effect of sampling the results in a pre-determined way and does not give an accurate representation of the system. In some cases, smoothing the data with a Gaussian kernel can lead to erroneous results.

When the image has been smoothed, the differing noise from pixel to pixel is averaged out, so the smoothing method is worth investigating in more detail. A method is required to actively smooth the data with a variable Gaussian kernel to account for the differences between each photon flash. One such method involves the concept of Scale Space, discussed in Section 6.

5.4 Verification of results

Here, the integrated image referred to is not a standard 'sit and wait' time period integration. A lead sheet has been placed at an angle across the scintillator. 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 background from accumulated dark current. One thousand images are taken, each with an integration time of 1/100th of a second. To simply add these images would lead to the same problem as taking a standard integrated image.

Here, a reject level is set above the image offset (the base level of the image). All pixels above this reject level are added to the "integrated" image. In this way, signal from scintillation flashes will dominate but bright pixels which are not removed as single pixel events will also be present. A result of summing one thousand images in this way is shown in Figure 4. A lead sheet was placed diagonally across the scintillator covering the lower right half of the device.

Figure 4 shows the diagonal lead sheet placed across the image area of the CCD with the much greater signal in the upper left half of the image. The probability of the signal seen here being from bright pixels and not a source external to the CCD is negligible, as a similar signal level would be expected across the whole image area. This implies that the signal is not bright pixels and is from an external source.


Figure 4. The 'integrated' image of one thousand individual images showing the position of the lead sheet and the lack of signal in the shielded columns either side of the image region and at the top and bottom of the CCD image region. The events can be seen to be purely inside the image region and not in either shielded region. 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.

5.5 X-ray interactions in the scintillator

The dark reference pixels in the CCD97 are obtained through the use of a small shielded region around the outer pixels of the image area. The shielded region of the CCD will not be subject to photons around 550 nm (the peak wavelength of light emitted from CsI:Tl). X-rays will, however, pass through the shielding and interact directly with the CCD. Here, the question is whether these 'flashes' seen in the CCD are due to X-rays interacting in the CCD 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 the enlarged region of Figure 4 it can be seen that the events recorded here are purely in the image section of the CCD and do not pass into the shielded regions. This is a good indication that the events are due to X-ray interactions in the scintillator being detected as photon flashes in the CCD.

6. SCALE-SPACE

When considering a one dimensional standing wave such as a sound wave, the Fourier Transform (FT) gives useful information regarding the frequencies present. From this FT, the original wave can easily be reproduced. If, however, a wave is considered in which the frequency changes over time, although the FT gives information regarding the frequencies which are present, the original wave cannot be reproduced.

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

The theory of scale-space is similar in many ways to the wavelet transform and has many applications to image processing. An image convoluted with a Gaussian kernel provides a single level in the scale-space representation. If the convolution is repeated with kernels of incrementing scale then a stack of images can be produced. This forms the basis of the scale-space analysis. When an image is transformed to scale-space, a stacked set of images is produced, each at a

progressively larger scale. Here, the scale parameter is equivalent to the sigma squared of the Gaussian kernel. Work by Tony Lindeberg [4-6] discusses scale-space in great detail.

Over many scales, a feature in the image does not move in location. In this case, the photon flash will not move in the plane of the image, it will purely change in intensity and spread as the image is smoothed to different scales. This allows a feature in the image to be traced over scales, leading to some very useful results.

6.1 Scale selection

From the perspective of this work, the most useful implication of scale-space occurs when a feature is traced over scales. It was found that "maxima over scales of normalised derivatives reflect the scales over which spatial variations take place in the signal" [5]. In other words, the spread of a photon flash in an image can be determined.

The process used to determine the scale of the photon flash involves convoluting the image with a discrete Gaussian kernel over a selected range of scales. Once this scale-space representation of the image has been produced, the photon flash can be traced through the scales. In order to find the corresponding scale of the photon flash, the trace of the normalised Hessian matrix H_{norm} (Equation 5) is taken for each image and the value of these new images at the photon flash point can be plotted against the scale of the corresponding image. The values from these new images are usually squared to avoid sensitivity to the sign. The maximum in this corresponding curve occurs at a scale which is proportional to the scale of the photon flash.

$$H_{norm} = scale \times \begin{bmatrix} \dfrac{\partial^2}{\partial x^2} & \dfrac{\partial^2}{\partial x \partial y} \\ \dfrac{\partial^2}{\partial y \partial x} & \dfrac{\partial^2}{\partial y^2} \end{bmatrix} \tag{5}$$

Through the convolution of the original image with discrete Gaussian kernels at progressive scales and with the use of the trace of the normalised Hessian matrix the scale of any feature in the image can be determined. This method can be used to determine the spread of the individual photon flashes. Using this method, the errors in the signal discussed previously will be averaged out and a form of photon-counting can be achieved.

7. AUTOMATIC SCALE SELECTION

Following the initial image processing on the raw images, further processing is required in order to extract the relevant signal data. This process begins with point-source extraction, in which the location of the photon flashes is determined. Following the point-source extraction, the scale of the photon flash can be calculated through the use of the scale-space technique described above.

7.1 Point-source extraction

In order to differentiate between the photon flashes and the background noise, a simple point-source extraction technique is used. Due to the pixel to pixel variation in shot noise and avalanche gain noise, the image must first be lightly smoothed such that the photon-flashes take a more natural Gaussian-like appearance.

The image is scanned row-by-row from the upper left to lower right corner. At each pixel location the surrounding 8 pixels (to create a 3 × 3 sub-array) are examined and if not all are found to be above a pre-set limit then the pixel is skipped and the next examined. If all the pixels in the 3 × 3 sub-array are found to be above the limit then the surrounding 24 pixels are analysed (creating a 5 × 5 sub-array). The requirement for this larger sub-array is that all signal values must be lower than the central pixel signal value. In this way, a table of central positions for individual photon flashes observed in an image is produced.

7.2 Automatic scale selection

Using the new table of photon-flash locations, the scale of each event can be found. The image is transformed to scale-space through convolution with appropriate discrete Gaussian kernels, following which the trace of the normalised Hessian matrix is found at each point. With this image-stack the point-source locations can be traced through scales. As described above, the maxima through scales of the trace of the normalised Hessian matrix is related to the scale of the point-source itself. This method of scale selection has been initially tested on pre-defined discrete Gaussian point-

sources of similar scale to those expected for the photon-flashes in reality. The results of the testing look very promising and the scale of the 'faked' photon flashes can be calculated. This process has been repeated on data taken with the experimental system. Preliminary results for the automatic scale selection of photon-flashes from real images are shown in Figure 5.


Figure 5. Five points of interest selected from a section of a real image from experimental data. The points have been automatically ringed with circles of radius related to the scale of the photon flash determined by the automatic scale selection process.

In Figure 5, a small section of the raw image is shown. Five points of interest have been selected and the automatic scale selection code applied. The radius of the circles shown over each photon-flash are automatically provided through the scale-selection process. The circles can clearly be seen to follow the size of the photon-flashes from the upper brighter small scale flash (interaction depth of X-ray in scintillator presumed close to the FOP) to the larger less bright flashes towards the lower part of the image (interaction depth of X-ray in scintillator presumed further from the FOP).

8. SIMULATING PHOTON FLASHES

Geant4 is a Monte-Carlo simulation toolkit which allows code to be written to trace the passage of particles through matter. Here, the particles of interest are X-ray photons and photons in the visible spectrum. The scintillator can be modelled with the appropriate properties coded by the user. Optical effects both inside the material and at surface interfaces can be included. Each individual part of the simulation requires the production of bespoke C++ code, allowing full customisation of the model.

8.1 The simulated experiment

The scintillator is defined in the simulation, as with all materials, with the basic properties (such as the atomic make-up and the density). Where the scintillator differs from the standard material definition is in the inclusion of the specific scintillation properties, the most important of these properties being the scintillation yield, output spectrum and decay constants over the energy ranges required and the optical properties of the material. All edges in the scintillator are defined as 'rough' surfaces to more accurately describe the real-life material. The detector must then be modelled. Light passage through the FOP is assumed perfect and the fibre width is assumed to have negligible effects compared to the pixel size of the device. As such, a pixelated detector is inserted below the scintillator. Any visible photons which pass through the detector are logged and can be summed in the pixel structure to provide an image similar to that seen in the experimental results.

Following the definition of the physics and materials present, the experimental set-up must be supplied. In order to test the basic operation of the simulation, a simple experimental set-up has been coded. A lead sheet of thickness 2 cm is placed over the scintillator in two arrangements: firstly to create an edge and secondly a slit of width 100 µm. The

normally incident X-rays are randomly generated above the scintillator surface and pass through the scintillator. Scintillation is then dependent on the properties defined in the code. A full readout of the pixelated structure is taken for each incident particle. This data is then processed using bespoke C++ code to extract the photon-flashes and further analysis can be undertaken to produce images such as those shown in Figure 6.

Three individually simulated photon-flashes are shown in Figure 7, along with a cross-section of the pixel intensities through the flash. The variation in the peak shape with the depth of interaction in the scintillator is clearly visible, Figure 8. It is this variation in peak shape which adds to the poor spectral resolution of such a device if no correction is applied. The results of this simulation are intended to be compared with the experimental data to provide a better understanding of how the energy of the incident X-ray and the position of the interaction in the scintillator affect the profile of the corresponding photon-flash.


Figure 6. Images showing the simulated output from 100 keV X-rays incident on the scintillator. The first picture shows an integrated image over several seconds on a 100 μm lead slit, with the second image showing only the central points of each photon flash. The third image shows the integrated image of a lead edge, with the fourth image showing only the central points of each flash. The increase in resolution over the integrated images is clear.
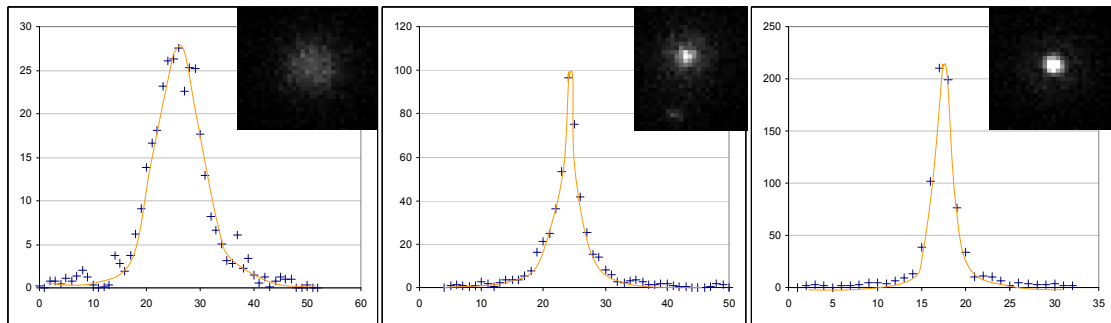

Figure 7. Three examples of simulated photon flashes with different interaction depths inside the scintillator. The vertical axes show the intensity in arbitrary units. The horizontal axes show the pixel numbers. The curves are fitted by eye to aid visualisation of the results.
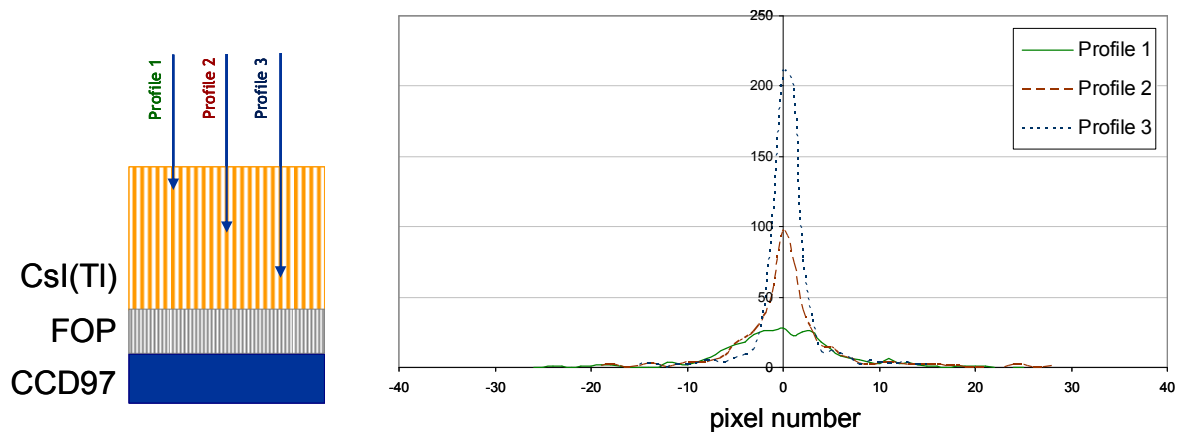
Figure 8. The three example photon flashes from Figure 8 on a comparative scale. The spread and peak intensity (shown in arbitrary units) scale directly and inversely respectively with the distance of interaction from the fibre optic plate.

9. CONCLUSIONS

Through the use of a combination of simulation and automatic scale selection the previous problems with spectral resolution in a scintillator-coupled EMCCD have been investigated. It has been shown that there are fundamental limits on the spectral resolution if raw data is used to approximate the energy of a photon flash. By fitting to the raw data through the use of scale-space techniques in future work, some of these problems can be resolved and a better energy resolution can be achieved.

Individual photon flashes caused by single X-ray photon interactions in the scintillator have been seen in the laboratory for the low energy 60 keV Am-241 source at a relatively high temperature of -10°C to -15°C (compared to temperatures as low as -50°C in the literature) [1,2]. These individual photon flashes have been shown to be due to X-ray interactions in the scintillator as opposed to bright pixels, direct interactions in the EMCCD or randomly positioned increases in noise.

The imaging capabilities of such a device have been previously discussed [1,2,8-10,14] and attempts to improve upon the spatial resolution through fitting to the raw data will be made. Automatic scale selection has been shown to fit the purpose required here with simple discrete Gaussian images. The use of automatic scale selection with the photon flashes achieved experimentally has shown promising results, with a clear correlation seen between the selected scale and the real scale of the photon flash.

The initial results from the simulation show the depth of interaction effects which have previously plagued attempts to spectrally resolve the photon flashes. The simulations also show the benefits to imaging in analysing individual photon flashes as opposed to integrating an image over an extended period of time. The theoretical increase in spatial resolution can be seen in the simulated results and similar results are hoped to be achieved experimentally in the near future.

The device detailed here may be open to many uses, depending on the final spectral resolutions achieved. One such application is in the field of small animal Single Photon Emission Computed Tomography (SPECT) [10]. Through the resolving of the energies of different isotopes a multi-label imaging system may be produced with the number of levels allowed dependent on the ability to resolve the differing energies.

REFERENCES

1. 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**, N109-N119 (2005)

2. Heemskerk, J., W., T., Westra, A., H., Linotte, P., M., Ligtvoet, K., M., Zbijewski, W. and Beekman, F., J., "Front-illuminated versus back-illuminated photon-counting CCD-based gamma camera: important consequences for spatial resolution and energy resolution," *Phys. Med. Biol.* **52**, N149-N162 (2007)

3. Jerram, P., Pool, P., Bell, R., Burt, D., Bowring, S., Spencer, S., Hazelwood, M., Moody, I., Catlett, N., Heyes, P., "The LLLCCD: Low light imaging without the need for an intensifier," *Proc. SPIE* **4306**, 178-186 (2001)

4. Lindeberg, T., "Discrete Scale-Space Theory and the Scale-Space Primal Sketch," PhD thesis, ISRN KTH NA/P--91/8-SE, all, (1991)

5. Lindeberg, T., "Feature detection with automatic scale selection," *International Journal of Computer Vision* **30**, no. 2, 77-116 (1998)

6. Lindeberg, T., "Principles for automatic scale selection", *Handbook on Computer Vision and Applications* **2**, 239-274, (1999)

7. 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**, 289-98 (2001)

8. 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**, 3540-3545 (2006)

9. Miller, B., W., Barber, H., B., Barrett, H., H., Shestakova, I., Singh, B., Nagarkar, V., V., "Single-photon spatial and energy resolution of a columnar CsI(Tl)/EMCCD gamma-camera using maximum-likelihood estimation," *SPIE* **6142**, 61421T (2006)

10. 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**, 65100N (2007)

11. Pool, P., J., Morris, D., G., Burt, D., J., Bell, R., T., Holland, A., D., Smith, D., R., "Application of electron multiplying CCD technology in space instrumentation", *Proc. SPIE* **5902**, 67-72 (2005)

12. Robbins, M., S. and Hadwen, B., J., "The Noise Performance of Electron Multiplying Charge-Coupled Devices", *IEEE Transactions On Electron Devices* **50**(5), 1227-1232 (2003)

13. Smith, D., R., Walton, D., Ingley, R., Holland, A., D., Cropper, M., Pool, P., "EMCCDs for Space Applications", *Proc. SPIE* **6276**, 62760K (2006)

14. 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**, 2724-2728 (2004)

15. CCD97 Back Illuminated Datasheet, e2v technologies public document, A1A-CCD97BI_2P_IMO Issue 3, May 2004

16. Operating Manual, Scientific CCD Camera System CCD USB2REM, XCM-OP-133/1, February 2007