The University of Birmingham

The Development of Positron Imaging Systems for Applications in Industrial Process Tomography

Thomas William Leadbeater

This thesis submitted for the degree of DOCTOR OF PHILOSOPHY

Positron Imaging Centre,
School of Physics & Astronomy,
The University of Birmingham

January 2009

"He used radioactive tracers a good deal, to follow individual termites around. He'd been tracking one with a Geiger counter when I first met him, in fact"

Arthur C. Clarke, The Next Tenants (1956 Renown publishing Co.), Taken from 'Tales from the White Hart', Sidgwick & Jackson ltd. 1980.

Abstract

The diagnostic medical imaging technique of Positron Emission Tomography (PET) has been adapted at the University of Birmingham for use in imaging industrial processes. A particularly powerful technique, called Positron Emission Particle Tracking (PEPT) is used to study granular materials and fluid flows in a number of different applications.

The PEPT technique is used to follow the motion of a single particle labelled with a positron emitting radioactive isotope. The decay of the isotope and subsequent annihilation of the positron results in two back-to-back gamma rays, the detection of which defines a line along which the tracer particle is located. Triangulation of a small number of successive events allows the particle to be accurately located in three dimensions on a short timescale. The dynamic behaviour of the material can be studied in detail using this method.

This thesis describes the most recent development of the detector systems which are used for PEPT studies at the Positron Imaging Centre. Primarily, the work has focussed upon developing a modular positron camera which can be arranged in custom geometries around the system under study. This camera is transportable and has been used to investigate a number of applications *in situ*, opening a wide range of potential applications which had previously proved impossible to study using the PEPT technique.

The modular positron camera resulting from this work offers a novel instrument with potential to deliver new information on industrial processes. As well as proven successful operation on a number of applications the performance and limitations of the camera have been investigated and are described here.

The detector systems are derived from clinical imaging systems and have thus required significant modifications to allow their use for the study of industrial systems; these modifications are discussed in detail. One of the major limiting factors has been the original data acquisition system provided with the camera systems. This was designed to produce medical images over a long timescale and has proved unsuitable for the high speed storage of information with the time resolution required for a dynamic PEPT study.

A new data acquisition system has therefore been designed to better enable the storage and analysis of the raw data produced by the detectors. This device is based around a single desktop computer: it has significantly reduced the complexity of the original system and offers much in the way of flexibility for data acquisition and analysis. The system has proven to operate reliably and with superior performance compared to that of the original device; for example the data storage rate has been increased by a factor of around 40, allowing for more frequent and precise PEPT locations to be made.

The modular camera and the new data acquisition system are supported by a software suite allowing for PEPT analysis and the display of both raw and PEPT location data. Real time analysis and real time display of the tracer location is made possible with software which is described here.

The modified devices have been used for a number of PEPT studies on different systems in various locations. Results and brief discussion are given here to illustrate the power of the techniques and of the new systems for deriving relevant data regarding the processes studied. Applications include improving process efficiency and gaining understanding of the flow dynamics within industrial applications such as fluidised beds and polymer extrusion plant.

Louise,

Thank you for sharing those moments of frustration, many sleepless nights and the great times when everything finally works. This is for you.

Acknowledgments

There are many people who deserve my thanks for their continuing support during the three years it has taken to produce this work. Firstly I am indebted to my supervisor Professor David Parker for initially accepting my application and then for providing a relaxed and excellent working environment with much encouragement. I have thoroughly enjoyed working with David and the members of the Positron Imaging Centre, they also deserve a mention here:

Thanks go to Andy Ingram for many discussions and good company, Xianfang Fan for tracer production and always a cheery word and to Monjur Ahasan and Marc Hausard. Thanks to the late Michael Hawkesworth for funding, which I hope, was money well spent.

I have had much support from Paul Jagpal (nuclear physics & electronics) and Chaw & Ian in physics stores during this period, they all deserve my extended gratitude. I have found the physics store an invaluable resource without which would have resulted in many frustrating delays. Thanks also to Tony and John in the physics workshop where I have spent many enjoyable hours.

Thanks and love go to my father Chris, mother Anne-Marie and sister Sarah. You have constantly encouraged and supported me and pushed me in what appears to have been the right direction. This thesis is indirectly a product of your efforts and is therefore also written for you.

To my extended family Honda & Vicki, Louise & Rob, Ben & Ali, Pete, (Gid)Iain, Richard Limbs, Ian Bone Egg, Gareth, Nigel, Penny and Laura: a long time ago I charged you all with 'keeping me sane', time will tell if you have managed this challenge, but I am more then happy with the result so far! I could not wish for a better bunch of friends. Thank you all.

Other people who deserve a mention are the inhabitants of E320 both past and present: Dave, Ant, Leon, Tom, Pete, Ramone and Louise; the nuclear physics staff and postdocs: Martin, Garry, Peter, Neil, Nick, Lee, Dave and Carl for discussions as far ranging as physics, climbing and the attributes of oompa loompas. You have all made my time in the Nuclear physics group both productive and enjoyable.

Finally shouts go out to Bong Ra & End user, Sir Roni Size & Sir Bryan G, Mickey Finn, LTJ Bukem & MC Conrad, Fabio & Grooverider, Orbital, The Prodigy, Primal Scream, The Marley family, Michael Eavis, everybody at the Goat Lab, Bristol, Ed Cox and the pandemonium circus, UNKLE, Birth Control recordings inc., The D & G brewery, Arthur C. Clarke, Isaac Asimov, Robert Heinlein, Larry Niven, Terry Pratchett and Arthur Ransome, without who life would have been a very different and dull experience.

Contents

List of Figures

Chapter 1

Introduction

The characterisation of industrial processes and the modelling of granular and fluid flows are still not fully understood. The process industry is increasingly using solid particles, bubbles and liquid drops within their manufacturing processes: understanding of the behaviour and dynamics of such systems is essential for further improvement.

Industrial process tomography has been an area of active research and interest worldwide for the past two decades. The fields covered are wide ranging and advances have been made in both experimental and modelling techniques such as Computational Fluid Dynamics (CFD) and Discrete Element Modelling (DEM). The use of experiment and theory provide validation for both methods, and can offer new insights into physical processes and suggest further research techniques.

At the University of Birmingham Positron Imaging Centre the diagnostic Nuclear Medicine imaging technique of Positron Emission Tomography (PET), and a complementary technique called Positron Emission Particle Tracking (PEPT), have been developed for the study of particulate and multi-phase flow dynamics in industrial processes.

In both of these methods a radioactive tracer is inserted into the system of study and its emissions are recorded by a dedicated set of detectors placed around the device. The tracer isotope decays by emitting a positron (β^+) which subsequently annihilates with an electron after travelling a short distance.

Due to energy and momentum conservation this annihilation results in two 511 keV gamma photons which are emitted back-to-back (i.e. 180° apart). The coincident detection of both of these gamma photons then defines a Line Of Response (LOR) joining the two detector elements, along which the annihilation is assumed to have occurred.

Figure 1.1 shows a number of LORs arising from a point source located in the field of view of two detectors. Within the spatial resolution of the detectors, uncorrupted LORs intersect at the position of the source. The PEPT method uses a small number of LORs to frequently locate a single tracer particle in the field of view by triangulation.

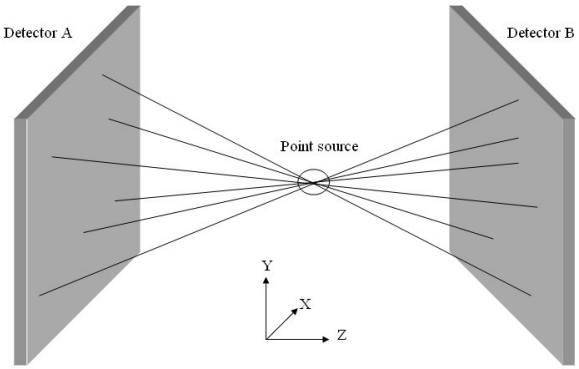

Figure 1.1: Detection of a number of LORs from a point source

Positron imaging is a widespread technique that is primarily used for medical diagnosis and clinical studies. The first report of the use of a positron camera applied to flow tracing and measurement in engineering and industrial subjects was given by Hawkesworth *et. al.* in 1986 [1], working at the University of Birmingham. Since then a number of position sensitive gamma ray detector systems have been developed and used for the imaging of industrial processes in the laboratory. This includes both custom built systems and commercially available devices originally designed for clinical studies.

Since 2003 the Positron Imaging Centre has been acquiring redundant clinical systems from hospitals in the UK and Europe; after some initial modification these scanners can be used for PEPT studies. As well as performing tomographic experiments on a wide range of chemical engineering and physical systems, a research program has been initiated in developing these devices to better enable the study of engineering processes.

Recent investigations have seen the components of these systems utilised to provide custom made cameras for the imaging of specific applications. This development has involved changes to system electronics and software, along with considerations in physics and chemistry. It offers a rapidly progressing and exciting field with a wealth of potential applications and possible future developments.

Presented in this thesis is the culmination of three years research and development of the detector systems used in the Positron Imaging Centre at the University of Birmingham, UK. This introductory chapter provides the reader with the necessary background to positron imaging techniques, with specific focus upon the methods employed to image industrial applications.

The systems used to record event data have all been modified to some extent to allow the study of engineering processes: thus in the following chapter a description of the systems used and the modifications that have been performed is made. This introduces some of the original work done by the author.

The main focus of the thesis is split between the continuing development of a modular positron camera, described in chapter 6, and the design and application of a new data acquisition system supporting the detector devices, described in chapter 4.

A description of the many supporting software applications that have been developed to operate the detector systems and analyse and display the experimental data is given in chapter 5.4.

Finally a summary review, conclusions and suggestions for further work in this field are described, before a list of submitted publications.

1.1 PET and PEPT

In a PET system a quantitative map of tracer concentration is built up over time by the collection and processing of a large number of LORs, typical acquisition times are minutes to hours. These LORs are binned into histograms (called sinograms, section 1.4) representing the total tracer activity along each measured line of response. Further analysis by standard techniques such as filtered backprojection or iterative reconstruction lead to tomographic images used for medical diagnosis, these techniques are described in the book edited by Zaidi [2].

Unlike other imaging techniques (X-ray, CT, MRI) which produce detail about the physical structure of the system, PET imaging reveals metabolic information from the relative concentration of the radiotracer in the system.

The ability of the different techniques to distinguish physical structure and metabolic information is extremely useful for clinical diagnosis; in the image given below it is possible to correlate metabolic information derived from PET (e.g. brain energy usage) with the physical structure derived from a MRI study (or X-ray CT). In the medical discipline the combination of the two images can be used for example to accurately locate the position of a cancer tumor given its activity measured by PET, and the physical structure measured by MRI / CT.

The image shown in figure 1.2, adapted from a Siemens press release [3], shows a Magnetic Resonance (MR) image (left) giving detail about the physical structure, and a PET image (right) giving detail about the physiological function, of the human brain.

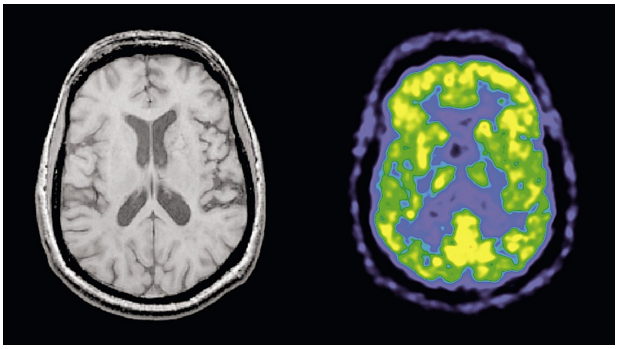

Figure 1.2: MR and PET images describing both anatomy and physiology

First suggested in 1986 by Bemrose *et. al.* at the University of Birmingham [4] in 1988, Positron Emission Particle Tracking (PEPT) has since been developed [5] into a powerful method for the study of industrial processes. In this method a small number of annihilation photon trajectories are used to locate the position of a single tracer particle many times per second by triangulation.

Disregarding corrupt events due to scattering, a sequential set of LORs should cross each other at the location of the tracer particle within the spatial resolution of the camera. Typically a fast moving tracer particle can be located with a precision of 1 mm 500 times per second, using 50 to 250 trajectories per location.

The power of the PEPT technique stems from the use of a tracer particle (section 2.4) that has similar properties to the bulk material used in an industrial system. The overall behaviour of the bulk material can be deduced from analysis of the tracer particle motion. This is provided that the experiment covers a long enough time scale such that the tracer particle exhibits the full range of motion and visits all possible locations in the experimental setup.

Example behaviour can be the dynamic motion of the entire bulk over a long time, short timescale positional data, spatial and velocity distributions, and tracer residence time distributions. This information is used to understand the studied process and to further develop the process techniques. Examples include increasing efficiency, lowering of running costs and offering new techniques and methods.

The data shown in figure 1.3 are results from a PEPT study of a spouted fluidised bed with 7 mm thick steel walls. They show the tracer occupancy plot (left); which over a long period of time represents the density of particles at each point. In the centre of the figure time averaged velocity vectors for the tracer particle are displayed. Finally the plot on the right shows the tracer trajectory for 5 seconds.

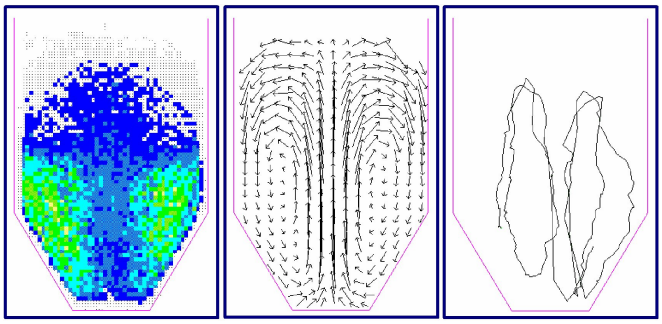

Figure 1.3: Examples of occupancy, time averaged velocity and trajectory plots

1.2 Background

The industrial process measurement community first gathered from 1992 until 1999 under the European Concerted Action on Process Tomography (ECAPT) banner. Since then there have been biennial meetings known as the World Congress on Industrial Process Tomography (WCIPT) where industrial imaging techniques and methods are discussed. The community is steadily growing, with many meetings (e.g. WCIPT) and a large internet community provided by the Virtual Centre for Industrial Process Tomography (VCIPT) playing a central role.

There are a large number of groups and facilities performing studies on chemical processes, fluid dynamics and industrial processes across the world. Many techniques using a wide range of physical data have been developed for imaging these processes; examples range from radiation systems such as infrared, microwave, optics, X-ray computed tomography (CT), gamma ray imaging and neutron devices, to electrical resistance, inductance and capacitance measuring devices and ultrasound and acoustical methods.

These all have their respective advantages and disadvantages, for example optical methods cannot be used to study opaque apparatus, X-ray penetration power is too low to study dense systems and electrical devices inserted into the process system can interfere with its operation (e.g. disturb the fluid flow) and can only be used on vessels which do not conduct if the electrodes are to be placed externally.

Since the invention of the scintillation gamma camera by Anger in 1957 [6] and the further development also by Anger [7] of the first positron scintillation camera, the use of radioactive tracers for medical diagnosis has become a widespread and significant research field. Positron Emission Tomography (PET) allows detailed information on metabolism to be obtained by mapping the concentration of a radioactively-labelled species as a function of time after introduction to the body. In parallel with developments in the medical sector the University of Birmingham Positron Imaging Centre has led an active role in the development of positron imaging techniques for applications outside of the medical field.

Anger himself suggests 'The motion of tracers in plants and animals, as well as industrial processes, could be studied [in] this way.' in his 1957 paper [6]. It is therefore surprising that from the initial development as a proof-of-concept method for mapping the concentration of gamma ray emitters within a volume, coupled with the huge amount of investment in R&D for these devices in the medical sector, that the use of positron imaging techniques has not found more widespread use for imaging other systems. Particularly the study of multi-phase flow and process technology where gamma ray imaging offers great potential for such studies. The University of Birmingham has been one of very few facilities worldwide successfully producing results in these fields for over 20 years.

Due to the large amount of energy released in nuclear decay, positron imaging is sensitive to very small amounts of material and therefore offers a sensitive measurement probe. As the penetrating power of the resultant gamma rays is high enough to penetrate a considerable thickness of material (approximately 50 % are transmitted through 11 mm of steel), the method readily lends itself to the study of industrial systems.

In PET imaging, short timescale dynamic studies are difficult to perform due to the time required to generate a statistically significant volume of data to provide an image. The PEPT technique allows the study of dynamic processes because it uses a small amount of data to provide accurate and frequent locations of the tracer particle in three dimensions. The number of events used in each location is of the order of a few hundred (typically 100 events) and data are generated rapidly (depending on the source activity and detector system sensitivity).

The spatial and temporal resolution of PET is not as high as that can be achieved using X-ray methods (CT), but for the study of multi-phase flows in particular positron imaging techniques can easily distinguish between the different phases. One of the biggest disadvantages of positron imaging is that radioactive tracers are used, these have to be made (at the University of Birmingham in-house cyclotron), transported, and then inserted into the system of study. On campus this is generally easy to arrange, however for studies off-campus the use of radioactive substances can cause many administrative delays.

It is possible that the high cost of the detector systems and facilities required to make the radioactive tracer has limited the use of positron imaging for industrial purposes.

However the department of Chemical Engineering at the University of Groningen, The Netherlands, has reported a number of trials where they have used the existing facilities in a local hospital to circumvent these financial considerations. Reports of these trial PEPT studies are given by Hoffmann *et. al.* [8] and Dechsiri *et. al.* [9] where the PEPT method is used to study a fluidised bed system using a modern ECAT EXACT HR+ camera manufactured by Siemens. Unfortunately the hospital prioritises clinical studies using these systems and modification is not possible.

The number of studies done by this group is consequentially low as they do not have access to the imaging systems, unlike in the Positron Imaging Centre at Birmingham where dedicated systems are used full time for the study of industrial processes.

The Delft University of Technology, Delft, Netherlands, have reported the use of two collinear non-collimated gamma cameras donated to the department by Siemens for PEPT work.

Stellema *et. al.* [10] describes the development of new tracer particles using neutron activation and gives an account of using absorbers placed in front of the scintillator screens to remove the Compton scattered component from the gamma ray signal, thus improving the signal to noise ratio. They report increased count rate performance when compared to the Multi-wire Proportional Chamber used at the time in Birmingham (section 1.2.1).

Despite the fact that they have their own dedicated imaging system this group has not provided many publications on their particle tracking studies, again unlike the situation in Birmingham where the PEPT method is used on a frequent basis to study a wide range of applications.

Several groups have developed their own method, commonly referred to as Computer Automated Radioactive Particle Tracking (CARPT) or Radioactive Particle Tracking (RPT) [11] [12] [13] [14]. This technique locates a single tracer particle by comparing the response to single gamma emissions received in an array of detectors surrounding the device under study. Comparing the relative singles count rate in each detector allows for accurate location of the particle by consideration of the attenuation and intensity of the gamma ray signal.

Unfortunately this method requires significant calibration (by either experiment using sources fixed in known positions, or by simulation) and therefore it is difficult and time consuming to study unique applications.

1.2.1 Positron Imaging at Birmingham

Positron imaging at Birmingham has been an active research interest since the acquisition of the RAL MkII PET system designed and built at the Rutherford Appleton laboratory in 1986 [1]. This system was one of a pair of devices, the other of which was used for medical research at the Royal Marsdon Hospital, Sutton, UK [15]. The prototype device is described by Bateman *et. al.* in [16]. Both devices consisted of a pair of large area multi-wire proportional chambers (MWPC) operated in coincidence; the system under study was placed between the two active faces of the detectors.

The initial goal was to observe the lubricant distribution in operating engines and gearboxes, for this conventional PET imaging was performed using oil labelled with a positron emitting radiotracer. As the technique proved effective, many other industrial and engineering applications were studied in a similar manner.

The MWPC camera had a maximum practical count rate of around 5 k s^{-1} due to deadtime in the readout electronics. It also suffered from a relatively low efficiency for detecting 511 keV photons; this meant that only apparatus in the steady state could be imaged effectively due to the time taken (1 hour for 3D, 1 minute for 2D [17]) to acquire enough data to reconstruct the image. It was soon realised that to study dynamic processes the labelling and tracking of a single particle in the bulk material could be used to provide equivalent data. This has since led to the development of the PEPT method, which has proved to be a particularly powerful technique for the study of industrial systems.

At Birmingham the MWPC camera was replaced in 1999 by a much more powerful positron camera. This device, the Forte model made by ADAC Laboratories California, was originally designed for clinical imaging but has been modified to enable PEPT experiments to be performed; the characteristics of this system are described by Parker *et. al.* [18].

The Forte camera consists of two opposing digital gamma camera heads operating in coincidence, with the system under study placed between the two. Each camera head contains a single sodium iodide scintillator crystal lightly doped with thallium (NaI(Tl)). The active area of each device covers 590 $\times$ 470 mm^2 and is 16 mm thick. The crystal is optically coupled to an array of 55 photomultiplier tubes (PMTs) each with a separate analogue to digital convertor controlled by an on-board computer.

The centroid position of a scintillation occurring in the crystal is then determined by software comparing the relative light intensity in each PMT with a spatial resolution of around 8 mm. Having separate digital channels allows for high speed data acquisition as different regions in the crystal can be processed in parallel.

The quantum efficiency for each head at a gamma ray energy of 511 keV is 23% resulting in a useful coincidence rate of approximately 100 k s^{-1} for most practical applications. Compared to the MWPC camera the spatial resolution has been found to be at least 25% improved, with a factor of 10 increase in sensitivity and resolving time reduced by around 60% enabling data rates of factors 20 to 40 times higher [18].

Using the Forte camera for a PEPT study on a slow moving particle locations accurate to 100 μm (in 3D) can be made 50 times per second, and a tracer moving at 1 m s^{-1} can be located to within 0.5 mm 250 times per second.

The MWPC has now been decommissioned and replaced by the Forte system which is used on a daily basis for PEPT studies of a wide range of applications in the laboratory. Other detectors, described below, are used for research and development into detector methods and now increasingly for PEPT studies.

In recent years many medical PET centres have replaced their existing scanners with new systems combining both PET and X-ray Computed Tomography (CT) imaging capabilities allowing simultaneous imaging of physical structure and physiology. It has been possible for the Positron Imaging Centre to acquire many of the older PET devices for research and development purposes as they are being replaced.

These systems consist of many small detector elements arranged in a cylindrical geometry providing a uniform field of view where the object of interest is placed. They are designed for the imaging of biochemical processes and metabolism across the whole body, and offer a field of view which uniformly covers a diameter of 76 cm with an approximate depth of 10.4 cm [19].

The diagram in figure 1.4 shows the two fixed geometry systems used at Birmingham. The left figure shows the geometry of the MWPC system and of the ADAC Forte, the separation between the detector heads is adjustable but rotation of the system has to be performed in order to sample a full 180° set of projection data (described in section 1.3). The figure on the right shows the fixed ring system for the other cameras, in these a ring of individual detecting elements allows full angular sampling to occur simultaneously.

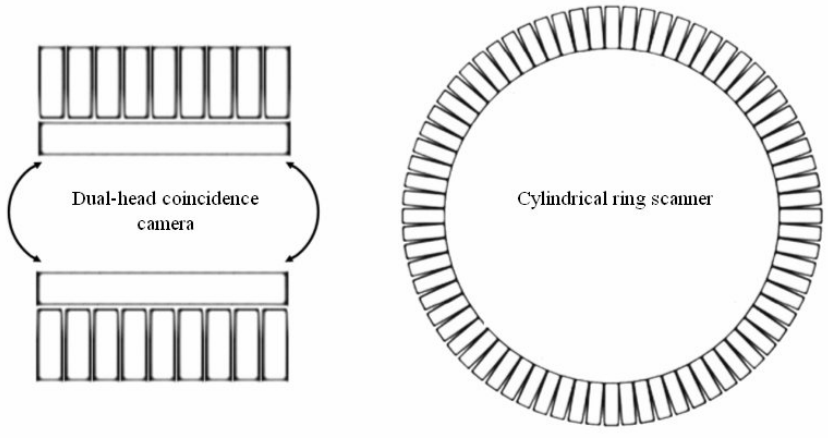

Figure 1.4: The two fixed geometries of PET systems used at Birmingham

In addition to the ADAC Forte the Centre now has four complete scanners of various generations made by Siemens systems / CTI (Knoxville, TN), plus components from two further devices.

One of these scanners has been left intact and performs PET imaging studies on small scale systems, the others are being used in research and development roles where unique geometries and applications are being tested. A major breakthrough in recent years (since 2003) has been a modular camera; components from these scanners have been adapted to provide a number of modular free-standing detectors. This system can be used for PEPT studies over a range of scales, from the bench top to full scale industrial plant operating *in situ*.

1.3 Positron Imaging Techniques

Positron imaging relies upon the detection of the radiation emitted by a radioactive isotope introduced to the system of study. Unlike Single Photon Emission Computed Tomography (SPECT) which relies upon the detection of single gamma photons emitted from the tracer substance, the technique of positron imaging is to detect the pair of back-to-back gamma rays arising from the decay of the tracer isotope. As PET uses a highly correlated source of photons, the spatial resolution and signal to noise ratio are improved and the sensitivity higher and more uniform compared to SPECT.

To provide an image the detector response to radiation emitted from an unknown position within the detector volume must give enough information in order to resolve the spatial distribution of the isotope throughout the system. As only the position (and energy and time) of the radiation interactions within the detector volume is known, it is necessary to provide some form of collimation (either physical as in SPECT, or electronic as in PET) in order to define the vector along which the photon travelled prior to detection. This is opposed to transmission tomography such as X-ray CT where the positions of both the radiation source and the interaction end point are known.

Single photon imaging records the spatial distribution of emissions from the body, this is generally done with a gamma camera based on the original design by Anger [6].

The Anger scintillation camera generally consists of a large scintillator coupled to a parallel-hole collimator which has the function of limiting the acceptance angle and therefore defining the spatial distribution of radiation entering the crystal.

A number of photo-sensitive devices register the scintillations occurring in the crystal and analogue 'Anger logic' is used to generate a signal comprising the (x, y) spatial position and energy of each detected event. Over time a position dependent histogram is used to map the spatial distribution of the gamma ray events recorded.

Figure 1.5 depicts a typical early Anger camera viewing an object which has been loaded with a gamma emitting isotope. Gamma rays with trajectories which are perpendicular to the collimator pass through to the scintillator crystal and produce a localised scintillation. The flash of scintillation light is received by the photomultiplier tubes (PMTs) and the signal is passed on to the Anger logic processor. Here the (x, y) coordinates of the scintillation are calculated along with the energy (Z) of the gamma excitation.

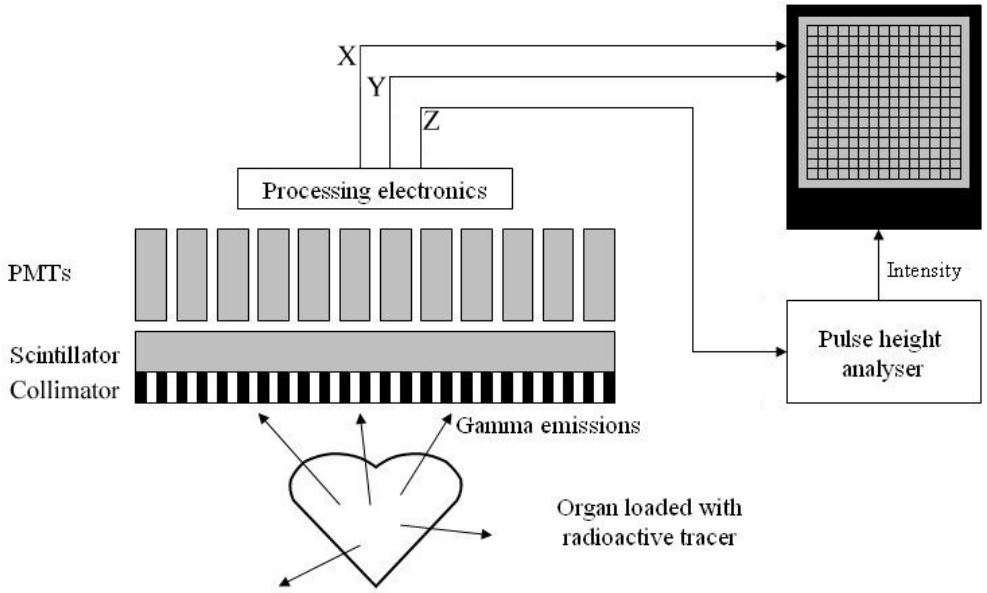

Figure 1.5: Typical gamma camera

The energy signal (Z) is passed through a discriminator which provides a pulse if the measured energy is within the allowed limits, thus removing the Compton scattered background from the signal. The (x, y) coordinates are passed onto the respective inputs of an oscilloscope while the discriminator pulse is passed onto the readout intensity control. If the pulse is accepted by the discriminator the plot intensity is turned on and a point appears on the oscilloscope screen at the relative (x, y) coordinates.

Initially the oscilloscope screen was photographed using a long exposure allowing many events to be recorded, areas of high activity would appear brighter than those regions with low activity. In this manner a projection image would be constructed. Modern techniques record this data electronically for post processing.

The image resulting from a study using a gamma camera consists of a two-dimensional representation of projections of three-dimensional distributions of activity in the detector field of view [20]. Images of activity deep within the system of study are sometimes difficult to resolve due to the added volume component. Computed Tomography (CT) studies can be done by sampling the distribution over a wide range of angles to remove this superposition.

This is illustrated in figure 1.6. In this diagram three spheres with increasing radii are placed in front of a gamma camera head. In the first frame the spheres are lined up perpendicular to the camera face, the resultant image shows the three spheres overlayed on top of each other. The second and third frames show rotation of the camera head (or object) and the resulting projection images.

As the camera head rotates, allowing for large angular sampling, the resultant image shows the activity profile in a number of projections giving detail about the profile in the volume. Provided the acquired data for each projection angle is stored, the data can be reconstructed using mathematical and computational techniques to display a three dimensional image.

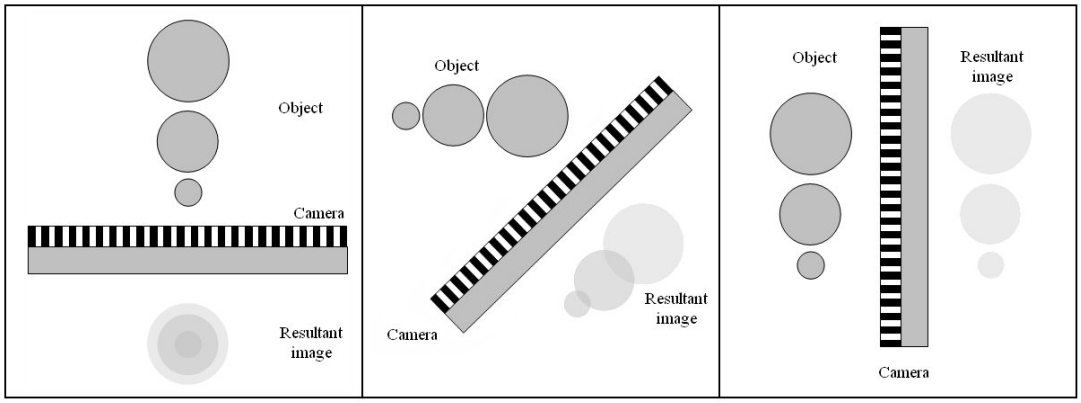

Figure 1.6: Projection imaging leading to an image

PET imaging improves upon SPECT with the use of electronic collimation [2]. Detection of both gamma rays gives more information than single gamma imaging; this can be used to resolve the position of the source without the need for physical collimation as each pair of detectors is only sensitive to the line of response coupling them (strictly this covers a volume and is therefore a tube of response). As the line of response for each event is defined at both ends, the vector describing the photon trajectory is more rigidly defined, and thus the geometric response is more uniform for PET.

Considerations regarding photon attenuation in the material studied also lead to improved image quality. In single photon imaging the probability of a photon impinging upon the detector is proportional to $exp(-\mu x)$, where μ is the linear attenuation coefficient and x is the depth of the gamma emission in the object. This factor is difficult to measure as the attenuation can vary across the material and the depth of the source is unknown.

In coincidence imaging, the probability that photons will be detected in coincidence is given by:

$$P \propto exp(-\mu x) \times exp[-\mu(T - x)] = exp(-\mu T) \tag{1.1}$$

Where μ is again the linear attenuation coefficient, x is the depth of emission within the object and T is the thickness of the object. This factor is easily determined by imaging a suitable source with and without the object present; the ratio of the two count rates giving an attenuation correction factor for each possible LOR. It is therefore much more straightforward to correct for the effect of photon attenuation with PET then in SPECT.

PET devices can therefore be constructed without the heavy collimator needed for single gamma imaging, allowing for a number of improvements in image quality and sensitivity (as genuine events are not lost in the dead material of the collimator). The spatial resolution of a PET system is then dependent upon the ability to distinguish between gamma events occurring close together in space at the detector.

As PET events are correlated, the signal to noise (uncorrelated) ratio is greatly improved over single gamma techniques, despite the open geometry allowing for many single gamma events. Geometric effects are seen in single photon imaging which degrade the spatial resolution and sensitivity as functions of location and depth. Again due to the correlated source, PET has more uniform sensitivity and resolution across the entire field of view [21].

1.3.1 Coincidence Detection

Following the emission and annihilation of a positron within the system under study the resultant gamma rays are both assumed to interact within the detector apparatus. Ideally both pulses representing the interaction in each detector element occur simultaneously and are registered as a coincident event by the processing unit.

A coincidence unit in its simplest form consists of two identical logic inputs. A digital output is produced by the coincidence unit if pulses arrive within the resolving time at each input, following a pulse arriving in either input. The order of arrival is not significant. Figure 1.7 illustrates a simple detector system for measuring coincident events.

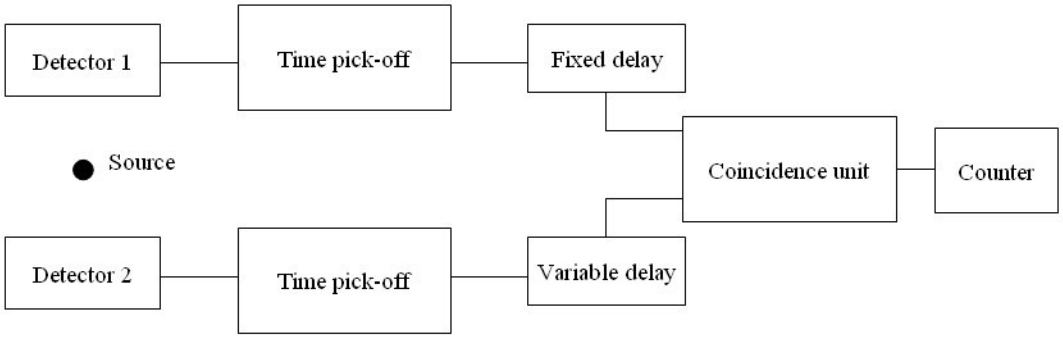

Figure 1.7: Schematic of a simple coincidence detector system

For the systems used in Birmingham a coincidence is then defined as two events occurring within a 12 ns time window between opposing detectors (see section 3.3). The length of this window (denoted 2τ) is dependent upon fluctuations in the speed of the system in responding to each event and is empirically determined. It allows all true events to be detected but it also allows for random events (not associated with the same annihilation) to be recorded as genuine data: consequences of this are described in section 1.3.2.

The diagram in figure 1.8 illustrates the coincidence gate method. In the figure labelled A a coincidence gate is opened by the first pulse to arrive at the processing unit (from detector A): following this trigger pulse a coincidence gate is opened, illustrated by the dashed grey lines. As no pulse is present in the data arriving from detector B within the gate width there is no coincidence response.

In B, a coincident gate is opened by a pulse (again in detector A) and a short time later another pulse is seen in the second detector. This corresponds to a genuine coincidence, thus a digital pulse representing the event is seen at the output. The third diagram, labelled C, shows the same case as in B; but it is the other detector which receives the initial trigger event.

The final illustration, D, shows that each received event opens a new coincidence gate; in this the initial coincidence is detected as in part B. The final pulse in detector B also opens a gate in which coincidences can be seen. This double trigger response is the reason that the coincidence gate width is commonly denoted 2τ as the total gate width is potentially twice the width of the trigger initiating the event.

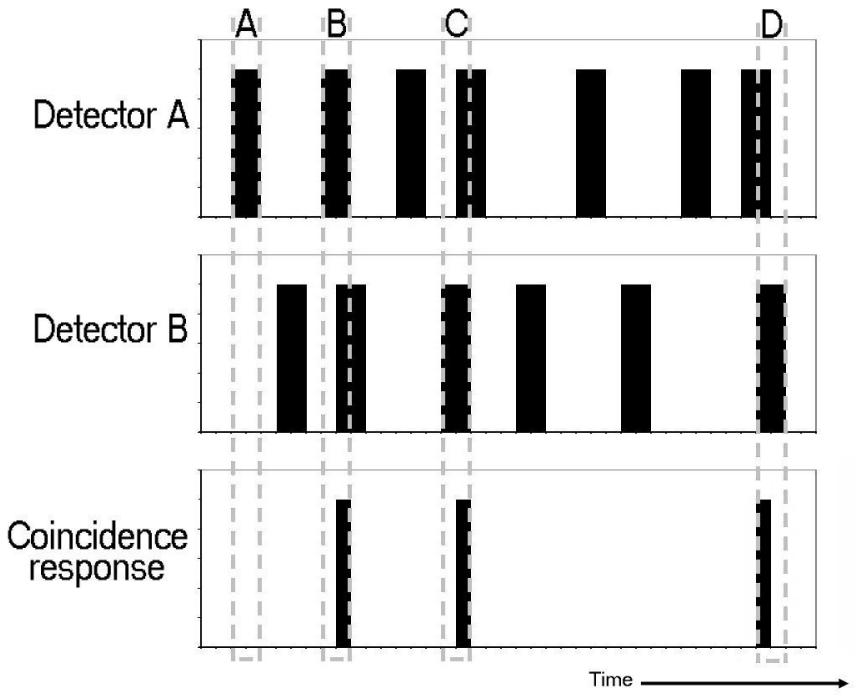

Figure 1.8: Illustration of coincidence gate method

In newer systems utilising fast scintillators (lutetium oxyorthosilicate; LSO) it is possible to measure the difference in arrival time for the two photons comprising an annihilation event. It is then possible to localise the source even further by using time-of-flight measurements. As light travels at a speed of 30 cm ns^{-1} these fast systems can localise the annihilation centre to around 5 - 10 cm which is useful to reduce statistical noise at the expense of added complexity.

1.3.2 Types of Event

There are a number of different types of event possible in a PET system, these are illustrated in figure 1.9 where the solid lines show actual photon trajectories, and the dotted lines show the reconstructed line of response. The data signal is made up from the coincident detection of genuine correlated pairs which have not undergone a scatter process before detection. These are referred to as Reals or Trues. Other events which are detected then contribute to the background noise in the signal.

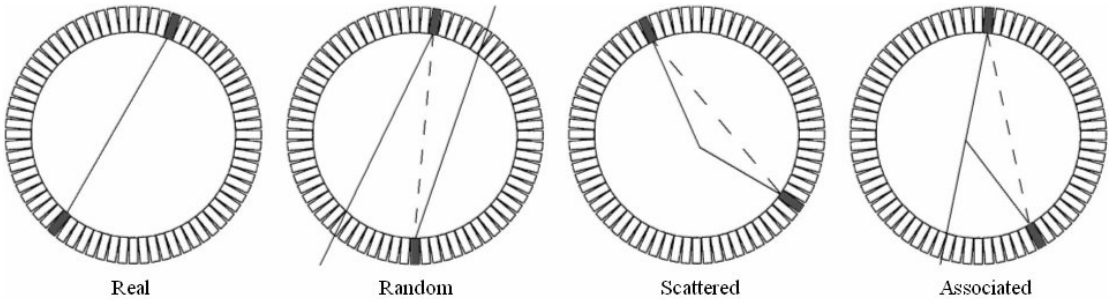

Figure 1.9: Schematic of possible events

Random (and Multiple) events occur when two (or more) uncorrelated gamma rays are detected within the coincidence window. Events arise from separate annihilations occurring close enough together in time that photons from the separate events arrive within the same coincidence gate. This is a major source of background noise, especially when a high activity source is used; the emission rate is then high enough that photons from different interactions can be received within the same timing window.

If the count rates in the two opposing detectors are R_1 and R_2 respectively, with a coincident gate width of 2τ, the corresponding random rate is given by:

$$R = 2\tau R_1 R_2 \tag{1.2}$$

If both count rates are approximately the same ($R_1 = R_2$) the randoms rate is then proportional to the square of the coincidence count rate, C.

$$R \propto 2\tau C^2 \tag{1.3}$$

This random signal can be a source of significant background, and therefore provides an upper limit to the useful activity allowed in the field of view of the detectors.

There is no physical method of distinguishing between the Real signal and the Random component. To deal with this signal statistical techniques can be used to determine the fraction of random events occurring in the detector system. In this approach another coincidence timing gate is used, this gate has the same length as the original gate but is delayed relative to the initial trigger event. This delay is long enough that events which are found to be in coincidence with the initial trigger cannot possibly be genuine and thus the rate found by the delayed gate gives a measure of the amount of background caused by random coincidences.

In sinogram formation as used in the medical technique (section 1.4), the delayed signal is processed in the same manner but is subtracted from the genuine data signal when the sinogram is formed. This is possible in a medical study because the data represents a time averaged response to the radiation emitted from the body. Thus a statistical correction is made which accounts for the distribution of randoms in the image, however this is at the cost of additional statistical uncertainty. In PEPT work the algorithm variables are dependent upon the amount of corrupt events (and thus the delayed rate) but the delayed signal is not used to perform a statistical correction. The algorithm has proved effective at rejecting corrupt events during the process of locating a single particle within the field of view of the camera (refer to chapter 2 for further discussion).

Genuine events which are corrupted by Compton scattering within the system under study or the detectors are termed Scatters, these can be rejected to some extent with the use of energy discrimination as the scattered gamma photons have lower energy than genuine photons. As the angle of scatter in a Compton event has an energy dependence, significant scattering (greater then 30°) results in a measurable decrease in incident energy. Low angle scatters do not lose enough energy in the interaction to be rejected as there is little energy loss (low energy loss scatter is forward biassed), for these scatters and the scintillator energy resolution is too poor to discriminate against such events. Ideally only events with characteristic energy of 511 keV would be accepted. In reality, an energy range is chosen to increase efficiency as events have a natural Gaussian energy spread due to the statistical nature of the interactions.

Associated events arise from the use of tracers which decay by emitting a positron and an associated gamma ray. An example of this is Sodium-22 which decays emitting a positron and leaves the daughter nucleus in an excited state. The daughter (^{22}Ne) returns to the ground state by the emission of a gamma ray with energy of 1274 keV. The signal from this gamma ray can be rejected using energy discrimination, however it can still be a cause of background if this high energy photon suffers a Compton interaction or undergoes a pair production process resulting in gamma photons with energy within the discriminator window.

1.4 Medical Techniques

A brief summary is given on the techniques used in nuclear medicine for completeness. Medical systems are used to image the physiological function of the human body or brain and necessarily have a uniform field of view covering an aperture suitable for the patient.

In medicine a PET image is built up over a long period of time and shows a static image of the concentration of the radio-labelled tracer fluid. At a typical data rate of 1 MHz a PET acquisition lasting for an hour would require 14.4 GBytes of storage space as a raw (uncompressed) data stream. The Sinogram is then used instead to store this data in a compressed manner.

The sinogram is a histogram representing the total possible phase space of the LORs allowed in the camera system. The abscissa of the sinogram gives the perpendicular displacement of the LOR from the centre of the field of view, the ordinate gives the angle the LOR makes with the horizontal axis of the field of view.

Figure 1.10 shows the Angle / Displacement convention, a number of LORs for a point in the field of view and the resulting sinogram construct. As there is a finite number of possible LORs the sinogram is broken into a number of pixels which correspond to the particular tube of response between two detectors, tallies are made in each pixel when corresponding events are detected. Data are acquired for a period of time such that the Poisson error in each histogram bin is reduced to an acceptable level.

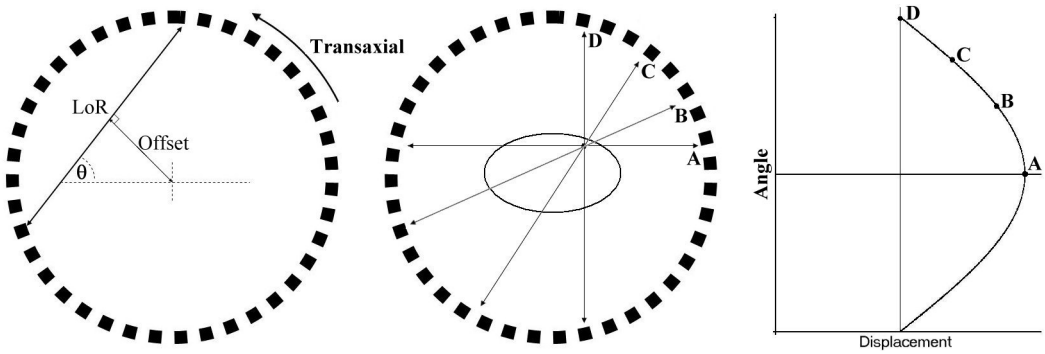

Figure 1.10: Sinogram convention

The sinogram represents the integrated data for each image. This consists of a set of two dimensional projections seen through the subject material. A number of techniques are used to process these data into an image, the analytical technique of filtered back-projection is a simple form of many advanced analytical techniques. Recently, due to increased computer speeds, iterative techniques are becoming favoured solutions.

Figure 1.11 shows the sinogram and resulting back-projected image for a PET scan of a human torso. Note how the arms are clearly visible as sine curves at the extremities of the sinogram. A typical medical PET scan consists of a number of two dimensional projections at different axial positions, these can be combined to form a true three dimensional image.

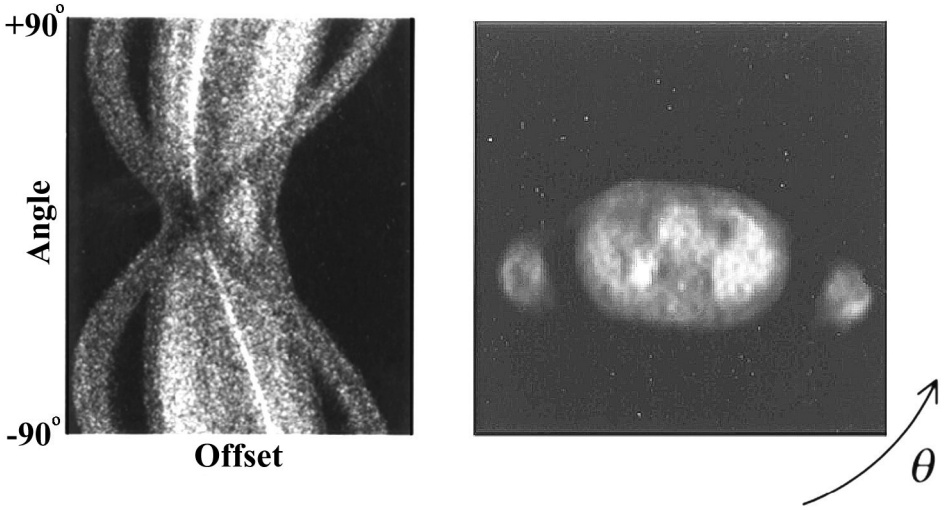

Figure 1.11: Sinogram and projection image for cross section of human torso

Medical PET can give dynamic information about a patient's metabolism if the acquired data are segregated into different cases. A common example of this would be cardiac gating where data are acquired into a different sinogram for each phase of the heartbeat. Over a long enough time sufficient data can be acquired for each of the three stages of the cycle, resulting in images describing each part. In this manner dynamic information over a small timescale can be extracted from averaged data over a long time.

This is done using a gating system monitoring the patients heartbeat or respiration. The system sends timing information to the scanner support electronics allowing the separate images to be made. This method can also be extended to separate the images by a fixed time (e.g. once per second) for imaging a long timescale dynamic process or for viewing portions of a repetitive cycle.

1.5 Modular Camera

Whilst the cylindrical ring systems offer increased data rates (due to higher efficiency) and improved spatial resolution compared to the Forte camera one of the major disadvantages is the fixed ring geometry of the detector gantry. Unlike the fairly open geometries of the other cameras where the separation of the two detectors can be adjusted and the entire camera rotated, the cylindrical ring detectors are fixed in place.

The field of view of these scanners is uniform over a diameter of 76 cm by 10.4 cm depth with the allowable aperture being 102 cm diameter (with reduced sensitivity at the edges). The systems that can be imaged then have to be small scale or arranged in such a manner that only a portion of the device is within the field of view, as the ring is continuous it is very difficult to position equipment within the camera volume.

It was determined that due to the modular construction of the new scanner: with the detectors arranged together in a number of identical small sub-systems, it would be possible to remove the detectors from the gantry frame and arrange them in more suitable geometries.

Sadrmomtaz [22] reported initial trials with this idea in 2005, where a number of detector elements were removed from the gantry and organised in a cuboidal geometry. The data acquisition system was modified to allow list mode data capture (required for PEPT) and correction for the new geometry. These trials were successful and the modular camera has since undergone further development and has been used for a number of relevant experiments.

Currently there are sixteen separate detector modules each containing 256 scintillator crystals (totaling 4096 individual detecting elements) with the supporting electronics housed in sturdy boxes for use in hostile environments and off-campus. The method has been extended to newer detector systems and similar modules have been made up from these components. For PEPT studies custom support gantries have been constructed to hold the modular detector elements around the apparatus under study.

The modular camera is particularly powerful as it can be arranged into custom geometries which are appropriate for the device being studied. For example, the imaging of pipelines has been considered by using a number of the modular devices distributed along the length of the pipe.

The device has enabled the potential for PEPT studies to be greatly extended as the detector system can now be transported to the experimental site and the flexible geometry allows for the study of a wide range of differently shaped systems over a considerable volume scale.

1.6 New Data Acquisition System

Commercial research and development in positron imaging (mainly in the medical discipline) has centred around the improvement of the detector systems. Advances include higher spatial resolution, improvement in scintillator response time: allowing for time-of-flight data to be used, and multi-modality devices used for simultaneous imaging of physical structure and physiological function (combined PET/CT or PET/MRI).

Despite these advancements there has been relatively little progress made in developing the data acquisition hardware. This can be explained with two arguments: the data describing events in the new systems has remained in a similar format (32-bit word), and the acquisition system, cheaply built from standard components, has remained sufficient to handle the data.

For example one of the major limiting factors in the acquisition of list mode data is the maximum disk write speed. When acquiring list mode data events are written sequentially to disk; this can be done in a buffered manner where events are first stored in system memory, however there is always a delay whilst data describing each event is transferred to the disk. At high data rates this can result in loss of data as the disk write speed is lower then the event acquisition rate.

In a clinical environment the data are formatted and processed into a compressed form (a sinogram, section 1.4) describing the detector response and is constructed in system memory as this compressed form of data requires little storage space. The sinogram data is then written to disk at intervals. Furthermore, in a clinical study where the patient dose is limited by radiation protection law, the total activity in the field of view, and therefore the data rate, is low enough to be handled by these old systems without need for further improvement.

In an industrial environment where a PEPT experiment is to be performed, list mode data acquisition is essential and the tracer activity can be made significantly higher then in a clinical study: resulting in much higher event rates. The disk write speed is then a serious limitation to data acquisition. This has motivated the development of a dedicated device that can be used to replace the older data acquisition systems and which offers high speed list mode data storage capabilities.

A high speed PC-based data acquisition system for use with the positron imaging systems at the Centre has been developed and described by Leadbeater *et. al.* [23]. This system replaces old dedicated hardware with a compact, flexible device with the same functionality and superior performance compared to the original systems.

Control of the entire detector system is provided by high-level software running on the same computer.

Chapter 2

Positron Emission Particle Tracking

The University of Birmingham Positron Imaging Centre has been actively developing and using positron imaging techniques for 25 years to study applications in physics and engineering. During this time the technique of Positron Emission Particle Tracking (PEPT) has proved to be a particularly powerful technique for the study of these processes. This chapter outlines the development of the PEPT technique, the methods used and the computer algorithm in full detail. Following this a section describing the analysis of PEPT data and the development of tracer particles is given, before a discussion of the fundamental limits in positron imaging.

Unlike PET where a static image representing the tracer concentration over a volume is generated over a long time (minutes), dynamic PEPT data relies on the knowledge that only a single point-like source is present in the field of view. Then, using just a small number of LORs acquired over a short timescale (ms) the spatial position of the tracer can be determined by triangulation. To perform PEPT each event must contain the data describing the position of the end points of each LOR, along with data representing the time that the event occurred. The data acquisition systems have thus been modified (see chapters 3 and 4) allowing the storage of list-mode positional data with the required temporal information.

The original Birmingham positron camera, described by Hawkesworth *et. al.* [1], was conceived to perform PET studies of industrial processes. The fact that the two gamma photons emitted from positron annihilation have relatively large energy motivated the development of PET technology over the use of single photon imaging using lower energy emissions (i.e. SPECT) for the study of industrial systems. This meant that it was possible to accurately image large metallic structures due to the penetrating power of the photons.

The Birmingham group had an interest in single particle tracking from early on, however the first mention in literature regarding particle tracking was by Bemrose *et. al.* in 1988 [4]. Here the possibility of tracking a single radio-labelled particle in the field of view of the camera, rather then a volume of labelled fluid, is considered and trials are performed. The method was fully developed by 1993 and used in a trial study where a labelled particle was accurately tracked in a rotating drum of powder, see Parker *et. al.* [5].

2.1 Philosophy

Fundamental to PEPT is the ability to label realistic tracer particles (section 2.4). In many cases it is possible to label a particle taken directly from the bulk material in the system under study, in this manner tracer motion is assumed to be identical to that of the unlabelled bulk. In other cases tracer particles similar in nature (e.g. volume, density) can be used, tracer motion is then assumed to be representative of the bulk. The tracking of a tracer particle for an extended period of time in a closed, circulating system builds up an integrated picture of particle behaviour at each point in space; provided that the experiment is allowed to run for long enough that the tracer particle explores all possible types of motion within the system.

Figure 2.1 shows a snapshot of software written by the author (described fully in chapter 5) to display raw data from the camera system. In this image 100 LORs arising from a 4 mm bead loaded with ^{22}Na placed in the field of view of a small detector setup are displayed. The view shows the three orthogonal planes in Cartesian coordinates for a cubic system consisting of 2 stacks of 4 detector modules. The position of the source can be clearly identified, alongside a small number of corrupt or random LORs.

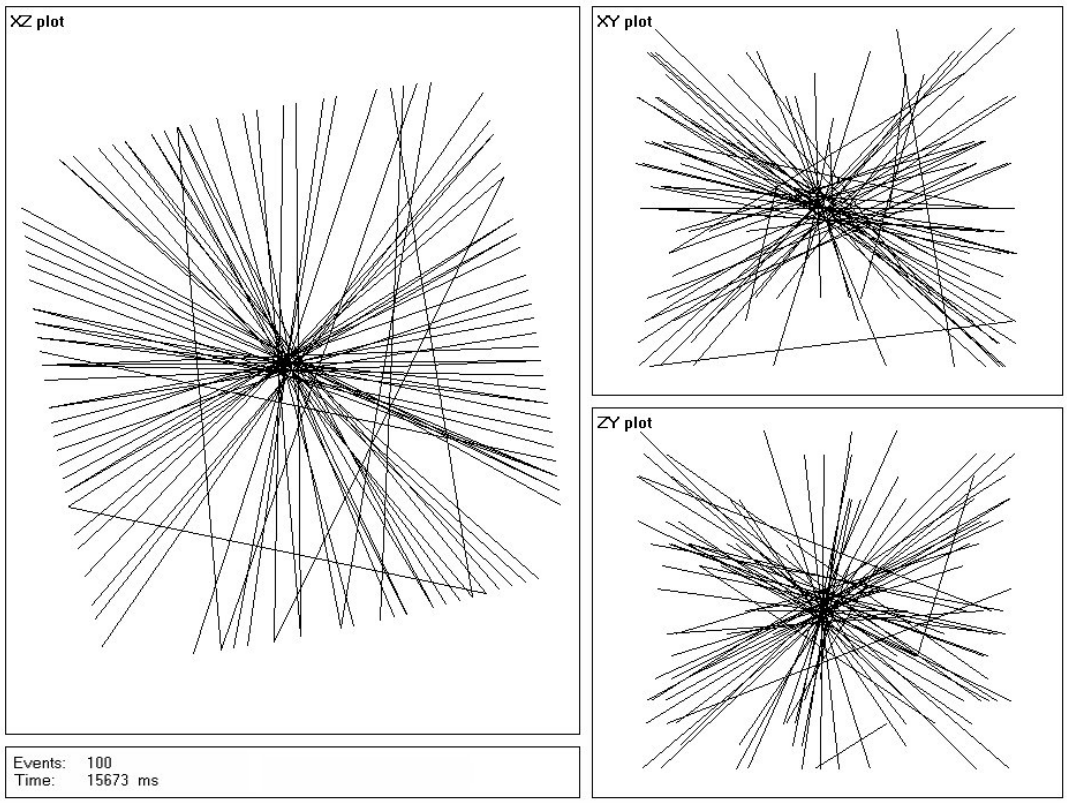

Figure 2.1: Location of the tracer particle in PEPT

For a given set of events such as shown in the figure, the point in the field of view closest to where all the trajectories pass is known as the minimum distance point: this can be assumed to be the particle location. Figure 2.2 illustrates this concept for three LORs in the field of view of the camera. Minimising the perpendicular distance between each line leads to the determination of the minimum distance point.

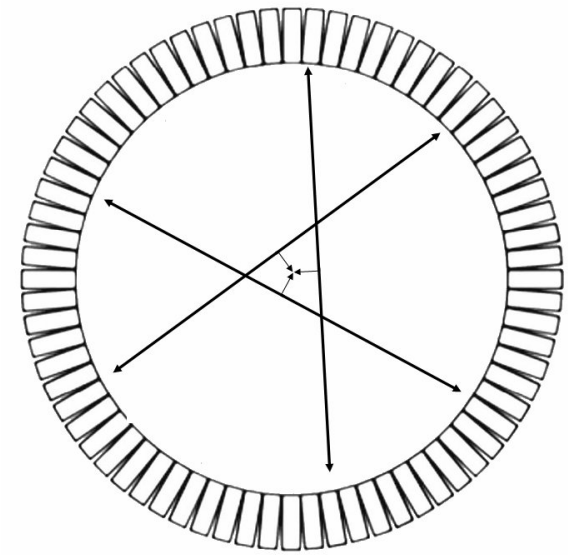

Figure 2.2: Minimum distance point for 3 LORs

The location algorithm uses an iterative approach where the minimum distance point for a given set of trajectories is calculated, the trajectories passing furthest away are discarded, and the minimum distance point recalculated using the remaining set. This approach is computationally intensive but has been shown to perform well for the tracking of single [18], and now multiple [24], particles within the field of view.

Provided that the rate of genuine event data is higher than that from random and scattered events the tracer particle can be localised in this manner. As the background events are generally broadcast randomly in space, the algorithm efficiently discards these events as they lie away from the concentrated position of the tracer particle.

Recent developments of this philosophy are used to track multiple particles. A modified algorithm is used where a particle is located, the LORs used to locate the original particle are removed from the data set, and the second particle is found using the reduced set. By using particles of differing activities particle identification can be performed as generally the most active particle is located first.

2.2 Algorithm

In list mode the data are recorded as a list of sequential trajectories; L_1, L_2,..., from which a set, S, of trajectories is selected for processing: L_1... L_N. The sum of the perpendicular distance from any point (x, y, z) is given by:

$$D_s(x, y, z) = \sum_S \delta_i(x, y, z), \tag{2.1}$$

where $\delta_i(x, y, z)$ is the perpendicular distance of the ith trajectory from the point (x, y, z). to calculate the minimum distance point, m_s, the solution is given by equation 2.2:

$$\nabla D_s(x, y, z) = 0 \tag{2.2}$$

For a given number of events $N(S)$, in the set S, the mean deviation of these trajectories from the minimum distance point m_s is:

$$d(S) = \frac{D_s(x, y, z)}{N(S)} \tag{2.3}$$

The order of events followed to locate the particle is then:

1. Start by calculating m_s for the given set of trajectories
2. Discard trajectories for which $\delta_i(x, y, z)$ is larger then $kd(S)$ where k is a fixed parameter related to the number of events discarded
3. This leaves a remaining subset of trajectories S_1 with a smaller mean deviation and therefore an improved particle location, m_{s1}
4. The original set of trajectories is re-examined using m_{s1} in case the new value is close to previously rejected trajectories
5. Proceed by iteration over further subsets S_2, S_3, ...
6. Ensure convergence by the condition that if $N(S_n) = N(S_{n-1})$ the trajectory with the largest value of $\delta_i(x, y, z)$ is removed
7. Stop iteration when the subset is reduced to a certain pre-defined fraction of events, f.

The size of the subset N determines how many events are originally used to resolve the particle location; generally this number lies between 50 and 500 depending on tracer activity, speed and experimental conditions.

The parameter k determines how many of the original events are rejected per iteration, Parker *et. al.* [5] show that the optimum value for k lies between 1.0 and 1.5.

The value of f is dependent upon the amount of corrupted data within the set, and therefore the amount of scattering material present in the field-of-view.

The geometry of the camera system also has an effect upon the value of f, so a stationary tracer particle within the field-of-view is usually studied to determine the optimum value. Typical values of f lie between 5% and 40% of the number of events in the set N.

A measurement of the mean deviations between the trajectories used in the final location, $d(S)$, is used to quantify an error value for each PEPT location. This value is given the nomenclature d and is referred to as the PEPT error.

2.3 Techniques

For completeness the following section outlines some of the common principles employed for the imaging of industrial systems using the PEPT technique.

Fan *et. al.* [25] show that for a set of N detected events, with a fraction f used to locate the tracer, the precision Δ of locating a stationary particle is given approximately by:

$$\Delta \approx \frac{w}{\sqrt{fN}} \tag{2.4}$$

where w is the spatial resolution of the camera. It can be seen that the value of Δ can be made arbitrarily small by making N sufficiently large. Over a long enough time scale (increasing N) a stationary particle can then be located arbitrarily well.

For example by acquiring event data at a rate of 100 k s^{-1}, with f equal to 10%, a central tracer can be located within about 1 mm every 10 ms using the Forte camera (w of around 10 mm). During this time however, a tracer particle moving at 2 m s^{-1} will move 20 mm with the LORs from this spread across its trajectory. To deal with fast moving tracers it is therefore necessary to locate more frequently with the corresponding loss of precision.

To first order the data rate achieved is proportional to the activity of the source although experimental conditions (such as the amount of scattering material present) and the motion of the source (if it moves into a region of low sensitivity) also contribute. The activity loaded on the source then has to be high enough to provide an adequate data rate, but low enough that the randoms signal does not interfere excessively. Generally optimum tracer activity lies in the range of 10 to 40 MBq.

Positron cameras with high efficiency and data logging rates are then suitable for the tracking of fast particles and particles with low activity. Low activity particles occur in experiments that run long enough so that the activity appreciably decays (commonly used ^{18}F has a half life of 110 minutes) or in experiments with small tracer particles that cannot be loaded with large amounts of activity.

As illustrated in figure 1.3 (above) the data derived from a PEPT experiment can be processed and represented in a variety of different ways. Perhaps the most simple method of displaying the data is to describe the particle track in three dimensions.

Provided the experiment covers a long enough period of time to give statistically significant location data at each position in space, an integrated picture can be built up of particle behaviour at each location. The spatial distribution of the experimental setup is divided into voxels (5 mm^3 is typical) and the fraction of experimental time during which the particle is located in each voxel is calculated. If the tracer motion is representative of the bulk, and the particles making up the system are allowed to mix freely, this integration gives the distribution of particle density across the system. This method can be taken further by using differently sized tracer particles mixed in with the bulk; here segregation effects can be investigated.

The instantaneous velocity of the tracer particle can be calculated given a set of particle locations with time. A time averaged velocity field can be constructed to demonstrate the overall dynamic behaviour of the system.

When studying mixing processes it is useful to select an initial volume and study how material disperses from the selected volume by following the tracer trajectory each time it leaves the volume. This technique can be used to determine a mixing index as a function of time. Recently conventional PET has been used to study mixing in detail by the labelling of a small amount of powder and performing PET imaging after a small number of revolutions in a pharmaceutical blender [26].

It is also possible to divide a system into individual compartments and study residence time or flow dynamics in each section of the apparatus. Most recently this has been done to analyse bulk rotation within a large scale fluidised bed; here the bed was divided into four compartments and tallies were made for particles moving in clockwise and anti-clockwise directions from compartment to compartment [27].

Described in section 1.4 the use of a remote trigger is useful in medicine to study dynamic motion of various cyclic processes. This type of information is also useful in an industrial experiment. For example when studying a stirred vessel it is interesting to correlate the position of the stirring impeller with the positional data derived from PEPT. The existing ability of the camera to perform this type of gated study has been modified to insert data describing this remote trigger into the data stream.

There are a number of different versions of the PEPT algorithm offering different performance. Usually a version which takes a sequential set of events is used to determine the location of the particle. However, it should be possible to use part of the previous set in the next location, especially if the tracer has a high speed. The algorithm variables are optimised by the study of a stationary tracer within the experimental apparatus. Minimising the standard deviation of the locations over a range of the algorithm variable f gives the most accurate results. It has been determined that the optimum value depends on spatial position when the field of view is non-uniform. It might be possible to use the delayed component of the data to better determine the algorithm variables, in this case it also should be possible to use it for dynamically finding the ideal values as the tracer moves through the system.

2.4 Tracers

Fundamental to Nuclear Medicine / PET and PEPT techniques is the ability to label materials with positron emitting radionuclides and introduce them into the system under study. Due to the high energy signal arising from nuclear decay, positron imaging is sensitive to picomolar masses (nanograms) of radio-labelled material allowing for near non-invasive studies to be made.

A common technique in nuclear medicine is to use chemical analogies of natural substrates such as ^{18}F fluorodeoxyglucose where a deoxyglucose molecule contains the radioactive isotope of fluorine. For PEPT a lot of effort has been made to develop the tracer chemistry and physics to provide adequate tracer particles with a range of sizes, for many materials sufficient activity can now be attached to tracer particles as small as 100 μm [25].

Many of the same isotopes used in the medical discipline are used with equal success for PEPT experiments. Isotope production methods are well known due to the vast experience with medical applications, although methods for labelling single tracer particles have had to be developed. Isotopes with short half lives are preferred for PEPT as it is sometimes difficult to separate the active material from the bulk for recovery after an experiment. Due to its ease of production and the fact that its half life (110 min) is ideal for most experiments; being long enough to acquire a decent statistical level of data but short enough that the activity has decayed to background level shortly after the experiment has finished, ^{18}F is the most popular isotope used for PEPT tracers.

The School of Physics and Astronomy at the University of Birmingham has a long history of using and developing cyclotrons for experiments in nuclear physics and development of radioisotope and medical techniques. An interesting account of the history of cyclotrons operated in Birmingham is given by Clarke [28]. The radioisotopes used in the Positron Imaging Centre have all been produced using the in-house cyclotrons, most recently a modern device: the Scanditronix MC40, has been installed. It is a flexible research cyclotron, delivering variable energy beams of hydrogen and helium ions with maximum energies of 40 MeV (protons or alphas), 20 MeV (deuterons) and 53 MeV (^{3}He).

The cyclotron is used to produce the isotopes used for PEPT, for nuclear physics research, radiation damage studies and detector testing. This cyclotron also provides a resource for students and produces ^{81}Rb daily for use in hospitals across the UK.

In labelling tracer particles for PEPT work a number of physical and chemical methods have been developed, these are direct activation, ion exchange and surface modification, described by Fan *et. al* in [29].

Tracers with a size range from 1 mm to 10 mm can be made by direct activation provided that they have a high oxygen content (e.g. silica or alumina). Direct activation allows the actual particles of interest to be used as tracers, rather then dissimilar materials of unknown behaviour. In this technique the particle is directly bombarded by a 33 MeV ^{3}He beam from the cyclotron. Some of the oxygen atoms in the particle are converted into the ^{18}F isotope by the reactions:

$$^{16}0(^3_2He, p)^{18}F \qquad \text{and} \qquad ^{16}0(^3_2He, n)^{18}Ne \rightarrow^{18} F \tag{2.5}$$

Using direct activation, the radioactivity A achieved in a single particle is determined by the irradiation time, t (min), the electrical current on the target, I (μA), and the cross sectional area of the particle, S (mm^2). This is shown in equation 2.6:

$$A = f_0(S, I, t) = f(d^2, I, t) \tag{2.6}$$

Where d is the diameter of the particle. The half life of the produced isotope limits the useful irradiation time (essentially this fixes an upper limit), thus for a fixed beam current (ions per unit area per second) the activity achieved is proportional to the square of the particle diameter. It is therefore difficult to generate adequate activity in particles with diameters less then 1 mm and the heating effects of the beam impose further limitations.

Short-lived radioisotopes are produced by activation of other structural elements in the irradiated particles, it is necessary to provide a 20 min 'cooling' period to allow these isotopes to decay. Following this over 90% of the remaining radioactivity is from ^{18}F and the tracer particle can be used for PEPT experiments.

Advantageously, the activated component of the particle produced in this manner exists as a structural element within a layer approximately 0.3 mm deep. This means that that activity is firmly fixed in place and effects of exchange with ions in the experimental system and surface degradation are negligible.

An ion exchange method is used when a smaller particle with size ranging from around 60 μm to 1 mm is required. ^{18}F is produced by the irradiation of water following the same nuclear reactions as in direct activation. Particles made from a strong base anion exchange resin are then immersed in the radioactive water solution, ions residing on the resin are replaced by ions in the solution and the particle gains a radioactive uptake. Disadvantageously the activity can leach off the particle, especially when used in an aqueous environment. Some success has been made with coating the resin particle to seal the activity in.

Many materials have a poor capacity to naturally take up fluoride via the ion exchange process. In order to increase the number of materials that can be effectively labelled a surface modification technique has been developed [29]. Metallic ions (e.g. Fe^{3+}) are introduced to the particle, these act as bridges by which the fluoride ions (e.g. $^{18}F^-$) bind. Using this technique a number of different materials have been labelled. Examples of these are polymers, metals, MCC and gamma alumina, coal, minerals, plant seeds and catalysts.

2.5 Limits

Many factors influence the precision of a PEPT measurement. Generally these factors can be grouped into two subsections: event related and background related. The former discusses the problems with resolving a real event, the latter with data that is recorded, but not in fact associated with a genuine event.

The spatial resolution of the system is limited due to the finite size of the detecting element. In a gamma interaction, the gamma ray can interact at any position in the detecting crystal element. The data that is passed onto the algorithm gives the centroid of the crystal element as the interaction position.

This is illustrated in figure 2.3 where the acceptance of a pair of opposing detectors is shown in grey, with the line of response describing the event as the arrow joining the centres of the pair. LORs occurring at any point within the tube of response coupling the detectors are all represented by the same coordinates; those of the central position of each crystal $(x_1, y_1, z_1; x_2, y_2, y_3; t)$.

Figure 2.3: Incongruity of gamma events and location

The PEPT algorithm assumes a uniform and continuous field of view with uniform angular sampling. The description of gamma interactions is not continuous across the detector face and if there are gaps between the detectors potentially some of the volume is not sampled, regions of over-sampling can occur if the sensitivity of the detector system varies across the field of view.

It is suspected that this lack of continuity degrades the positional data given by the PEPT algorithm. Investigations into these effects are reported in sections 6.4.2 and 6.4.3.

In an energetic decay where the positron has high kinetic energy, the positron can travel a short distance before coming to rest and annihilating with a local electron. The annihilation site is not necessarily at the location of the tracer isotope. The stopping power for positrons varies as a function of energy and the surrounding material and is therefore dependent upon experimental conditions.

Positron range varies as a function of the energy of the radioactive decay. Different isotopes have different decay energies and therefore some are more suited to high spatial resolution studies then others.

Table 2.1 shows the energies, ranges in tissue and the intrinsic spatial resolution for an ideal system for common PET/PEPT isotopes.

Isotope	Q-Value (keV)	β^+ energy (keV)	Range in tissue (mm)	Intrinsic resolution FWHM for ideal detectors 50 cm apart (mm)
^{11}C	1982.2	385.6	0.28	1.3
^{18}F	1655.50	249.8	0.22	1.2
^{66}Ga	5175	1740	1.35	2.0

Table 2.1: Positron range and spatial resolution [20], [30]

The tracer particle used in PEPT has finite dimensions with the isotope spread evenly across its volume. The positron emissions and subsequent annihilations do not all originate from the same point in space and therefore there is a distribution of events describing the location of the tracer. For an evenly loaded tracer particle, with isotropic positron decay, by using a number of genuine events particle location can be localised to a small region.

The precision of the measured location is given by the value labelled d: in this the standard deviation of the perpendicular distances from the minimum distance point to each LOR used in the final set of trajectories is calculated. A low value of d corresponds to an accurate location, in an ideal system the minimum value of d would be dependent upon the size of the tracer particle and the positron range in the surrounding material.

Finally the momentum of the centre of mass of the electron-positron pair at annihilation contributes to noise in the data signal. Ideally the annihilation should occur with zero CM momentum; resulting in the gamma ray trajectories being 180° apart. However, whilst the cross section for positron annihilation increases with decreasing momentum in the CM frame, the momentum of the lab frame (from the electron motion) can be significant, especially in a thermal system.

Figure 2.4 illustrates this case. It shows a genuine response in light grey and the resultant corrupted response due to the acollinearity of the emitted photons. The momentum of the electron is represented as the arrow labelled P_e, and is generally the cause of the angle of deflection, denoted θ.

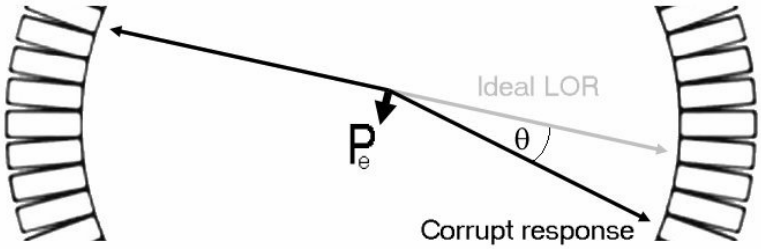

Figure 2.4: Diagram of non-collinear two gamma event

These effects can be reduced statistically by the recording and processing of a large number of events. It is therefore important to have a high activity tracer particle, good efficiency for detecting genuine events and as large as possible solid angle coverage in order to maximise the number of genuine events received per unit time.

As discussed above (section 1.3.2), the random coincidence signal is the major source of background in a positron imaging experiment. Other background comes from corrupt events such as scattered and associated gamma rays, cosmic rays and interference from other sources both radiative and electrical.

Cosmic rays form a very small component of the background signal because a cosmic event would have to cause a correlated capture signal that falls within the energy discriminator as well as being between two detectors operating in coincidence (section 3.3.2).

Radiative interference occurs when other positron or gamma sources are present when an experiment is being performed. The effect of these can be minimised by removing non-essential sources and by proper shielding.

Again event data from this type of interference has to meet the requirements for a genuine event. Furthermore there is a small but finite probability that the electron-positron pair can annihilate to produce more then two gamma rays, in this case coincidences between two of the resulting interactions can lead to corrupt LORs. The effects of these can be reduced by having a large sample of genuine data.

Finally, electrical interference provides a significant source of background signal, especially when the camera is operated in an electrically noisy location such as in the study of large engineering plant.

All the sub-systems used are adequately shielded from electrical interference and communication between the various devices is done over twisted pair cabling allowing for common mode rejection of noise.

In some applications, when not all detector modules are used, spare inputs can contribute to this noise. By rejecting data that arrives from non-existent detectors these corrupt data can be identified and rejected.

Chapter 3

Camera Systems

The Positron Imaging Centre has a wide range of positron camera systems available for performing experiments and for detector development applications. As discussed in section 1.2.1, the multi-wire proportional camera: designed for studying engineering applications, was the first such device used at the Centre. This system has now been de-commissioned. The second device: the ADAC-Forte, was originally designed for medical studies but has been modified to allow engineering applications to be studied. This system currently performs the majority of the PEPT research done at the Centre.

The most recent positron camera devices acquired by the Centre are the ECAT series of scanners developed for clinical PET imaging by CTI / Siemens. The development of dual modality devices offering combined X-ray CT (Computerised Tomography) and PET images has led to the replacement of the majority of PET systems in hospitals across the UK in recent years making the older machines readily available for research.

It has thus been possible to obtain these now redundant medical scanners and develop these systems further for the study of engineering applications. The Centre has acquired a large number of components for both ECAT-931 and ECAT-951 series scanners and an ECAT-EXACT from hospitals in the UK and Europe.

The ECAT series of scanner is described in detail in the following chapter. Components from both the 931 and 951 designs are used for developing novel systems for PEPT and PET. The major advantage of these scanners for development work is their modular nature; the scanners consist of a number of small detector blocks arranged in a ring geometry. From the two complete ECAT 951 systems that have been acquired, one has been stripped down to its component pieces and the extraneous parts (patient bed, support gantry etc.) disposed of. The other scanner has remained intact and has been used to perform PET studies on small systems [26]. With the arrival of the ECAT EXACT scanner the older system will probably be stripped to its component parts with the ECAT EXACT remaining intact for PET studies.

3.1 ECAT Scanners

The ECAT 931 and 951 series of PET scanners are high resolution positron emission computed tomography systems for imaging and quantification of biochemical processes and metabolism across the whole body. These scanners provide three dimensional acquisition in static, dynamic or triggered modes. They offer high spatial resolution in both the transaxial and axial dimensions and with a large axial field of view can be used to display organs in three dimensions. The block detector technology allows for high speed acquisitions to occur simultaneously from many detectors with excellent temporal resolution. The system provides simultaneous data acquisition and image reconstruction (performed in system hardware) allowing a relatively high patient throughput.

Both scanners have relatively large gantry diameters and were originally designed for imaging both the human head and the whole body. Due to the large aperture the sensitivity (cps μCi^{-1} cm^{-3}) of the system is lower than scanners designed for imaging the head only, which have smaller diameter gantries, however the uniformity of the spatial resolution is superior.

The photograph in figure 3.1 shows the physical support structure for the ECAT 951 scanner in the Positron Imaging Centre. The patient port is surrounded by the supports which hold the individual detector elements and electronics in a set of two rings around the port. The scanner consists of 2 rings of detector blocks (section 3.1.1), with 64 blocks per ring. Each block is subdivided into an array of 8 $\times$ 8 crystal elements. This results in 16 rings (axially) of 512 scintillator crystal elements around the ring (transaxially). Refer to section 3.3.1 for further discussion.

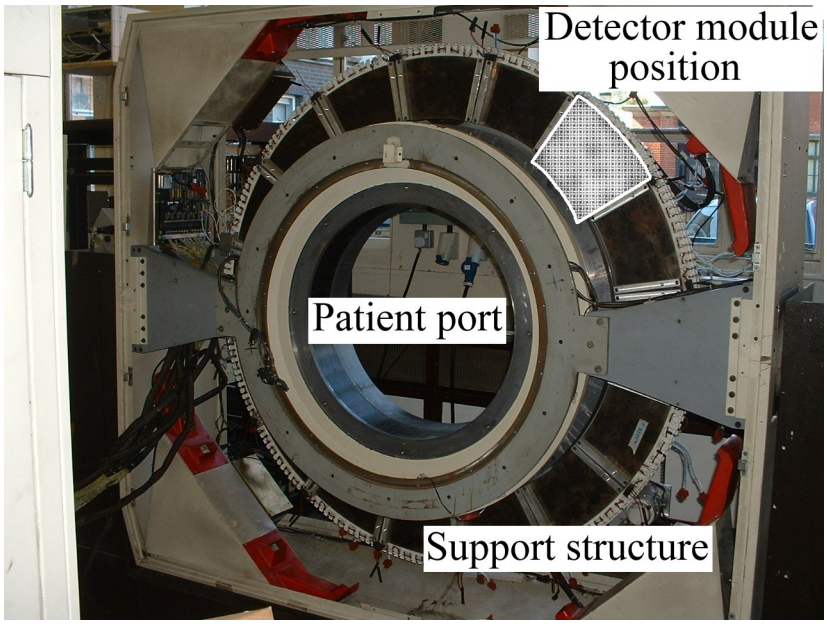

Figure 3.1: ECAT-951 scanner gantry

Table 3.1 shows some of the properties of the ECAT 951 scanner. The primary difference between this model and the 931 series is the number of scintillator crystals in each detector block. The 931 detector blocks each contain 8×4 crystals, whilst the 951 blocks contain 8×8 crystals, these detector blocks are described in section 3.1.1.

Spatial resolution	Axial $<$ 5mm, Transaxial $<$ 6mm
Energy resolution	9.4% (FWHM 511 keV photopeak)
Data acquisition rate	2 Mcps coincident events
Scintillator	Bismuth Germanate ($Bi_4Ge_3O_13$)
Crystal size	30 mm depth, 6.25×6.75 mm^2 area
Number of crystals	8192
Number of PMTs	512 (4 per block)
Number of blocks	32 (16 per ring)
Axial field of view	10.8 cm
Resolving time	4 ns
Dead time	1.2μs per block

Table 3.1: Properties of the ECAT 951 scanner [19]

Table 3.2 gives details about the original data processing and storage electronics used in both the 931 and 951 models.

Data Processing & Storage electronics	68020 host CPU
	Expandable 32Mb standard High Speed VME Memory
	VME Array Processor & Backprojector
	Graphics based User Interface Software
	SCSI Controller
	Fixed Disk & Tape Drive interface across the VME Bus
	670 Mb Disk Drive
	SunOS (UNIX) NFS access to VME Disk Resources
	Operator Console SUN 4/65 SPARC

Table 3.2: Data processing & storage properties [19]

The following sections (3.1 – 3.4) describe the subsystems of the scanner architecture: this includes the front end data acquisition components, data analysis and storage, and the system controller.

These subsystems communicate with each other using three separate data buses. Control signals and calibration data (e.g. histogram data) are carried by a serial bus operating via the master clock. Event data are carried by a separate serial bus from each detector to the Image Plane Coincidence Processor (IPCP); this device analyses the singles data from each detector and determines the coincidence response. The IPCP transmits coincidence data words in parallel to a hardware processing unit called the Real Time Sorter (RTS), where sinograms are constructed and image analysis is performed. The IPCP and RTS modules communicate in parallel over a high speed VME bus [31] which gives access to system memory and disk storage.

Conventionally the data are compiled into a sinogram (section 1.4) in VME memory and eventually transferred to disk storage. Image reconstruction is performed using the RTS hardware operating on sinograms stored in memory. For PEPT work where a list mode data stream is required, the RTS system has been modified to allow this type of data handling.

In summary this modification allows list mode data to be acquired into a VME memory buffer which is transferred to disk each time it is filled. This limits the maximum data storage rate for the system due to the relatively slow physical disk write speed. The maximum rate for each buffer depends upon the memory access and write speeds: this is significantly higher than the disk write speed.

The memory is configured into double buffered mode where event data are written to one buffer whilst data in the other buffer is transferred to disk. Data can be lost if the transfer is not quick enough and there are short gaps between each buffer.

3.1.1 Detector Block

This is the fundamental modular component of the ECAT scanner series containing the scintillator crystal viewed by four photomultiplier tubes as shown in figure 3.2. In the 931 series the crystal is divided into an 8×4 matrix separated by a slotted light guide designed to illuminate the PMTs in such a manner that the individual crystal segment can be determined. In the newer 951 series this matrix has been increased to 8×8.

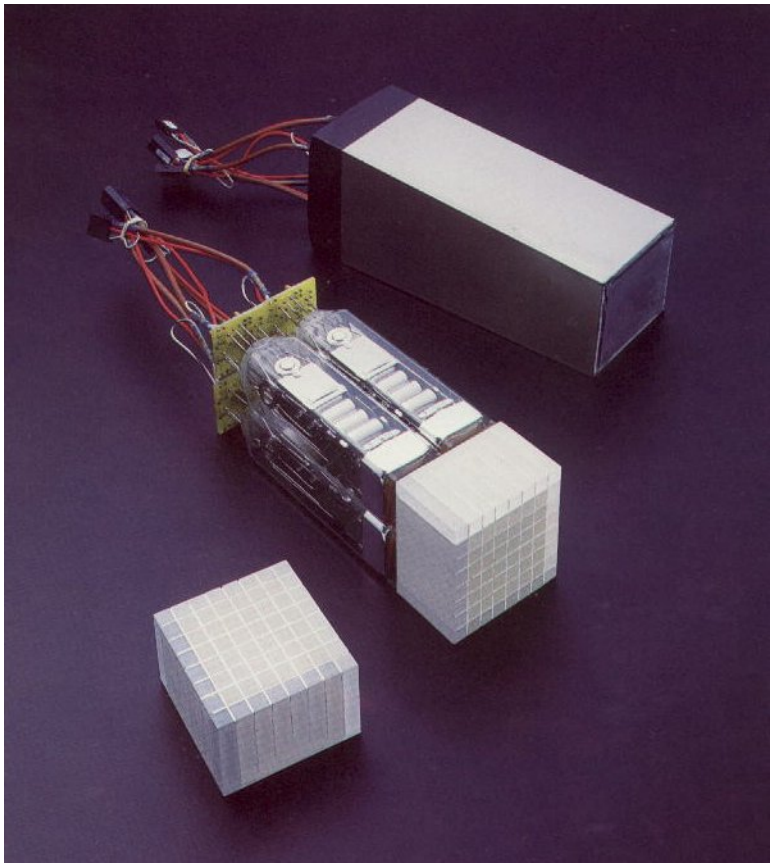

Figure 3.2: ECAT-951 detector block, showing segmented crystal element and PMTs

Considerations in geometry and scatter show that excitation of each of the crystal segments produces a different scintillation light intensity (i.e. measured energy) into each of the four PMTs. Casey and Nutt [32] ahve shown that the ratio of the number of photons collected in one PMT to the number of photons collected in another PMT depends only upon the probability that a photon will be collected, and is constant for each crystal segment. A decision circuit is used to compare the relative light intensities in each PMT leading to an explicit determination of the excited crystal element.

The measured energy from each PMT is compared to a lookup table that has been determined by measurement in known conditions and preprogrammed into the control system placed on the detector controller, as described in section 3.1.2.

3.1.2 Detector Bucket

The detector bucket holds four separate detector blocks and the associated electronics for these blocks. Each block has an analog signal processor, event position and energy processor, and time digitiser. Included in the analogue processor each PMT has its own preamplifier and flash analogue to digital converter.

Once the signal from each PMT is converted into a binary number there is further on-board processing provided by the rest of the bucket controller. A schematic for this is shown in figure 3.3, adapted from [19].

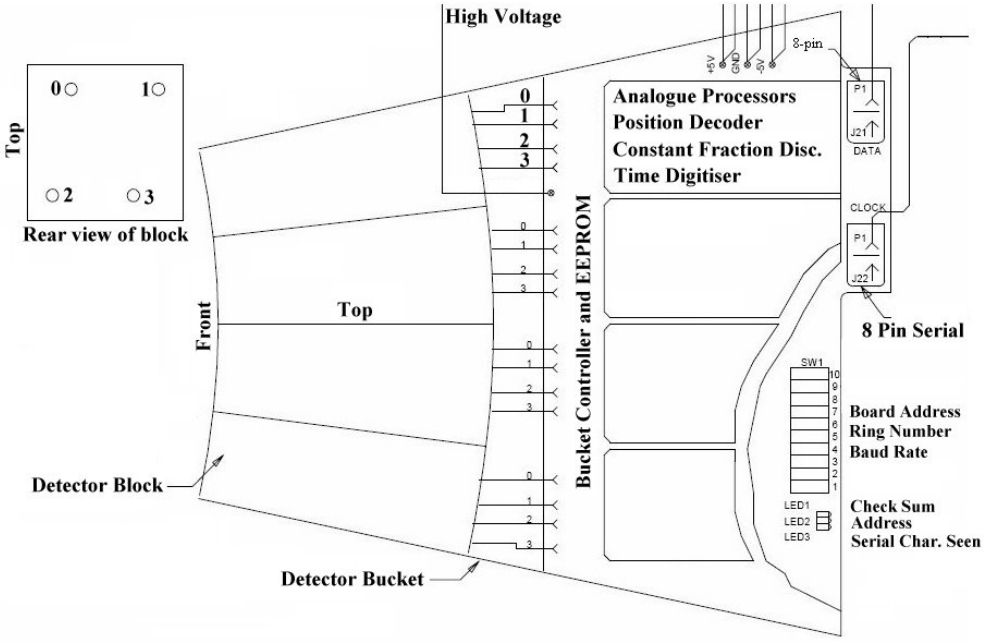

Figure 3.3: Schematic of a detector bucket

The signal from each of the 4 PMTs in a block is sent directly to the analogue signal processor where it is applied to the inputs of four variable gain amplifiers. The gains are set during bucket calibration using the bucket controller operating over the serial clock bus.

The output from each amplifier has a fast rise time from 0 mV to -150 mV, with a gradual decay. A following amplifier stage provides gain of a further factor of 10 and some smoothing of the waveform. The amplifier outputs are connected to summing RC networks which provide composite signals describing:

Total energy	PMT 0 + PMT 1 + PMT 2 + PMT 3
Transaxial energy	PMT 0 + PMT 1
Axial energy	PMT 0 + PMT 2

The composite signal describing the total energy is passed onto a constant fraction discriminator which is used to provide a statistically constant timing window for the calculation of the event energy using the area under the curve. The output from each summing network is sent to the input of an integrating amplifier which is triggered by the constant fraction discriminator; these integrators measure the transaxial, axial and total energy of the event.

The energy signals are then converted from an analogue pulse into a 6-bit digital word using Flash conversion. This process is controlled by the master clock, where the voltage is sampled on the first clock pulse, and presented on the digital output lines following the second clock pulse.

The position energy processor consists of a RAM containing the lookup table used to describe the crystal position of the event within the block. The two 6-bit words describing X (transaxial) and Y (axial) energy are applied to the address lines of the lookup RAM: the resulting memory address contains a 6-bit number describing the relative crystal position (section 3.1.1).

The total energy signal is then tested to determine if the event is valid; i.e that the total energy lies within the discriminator levels and is therefore not a Compton scattered or associated gamma event.

Finally the event time is digitised into another 6-bit word using the 256 ns period master clock as reference, this allows the time resolution for each event to be +/- 2 ns. Time data are presented as a time offset (from 0 to 256 ns) from the master clock signal.

A functional block diagram is given in figure 3.4 showing the flow of event data from the analogue signal received from the detector block, through energy and time processing to output as a digital serial word. For clarity this diagram shows one of the four block processors contained on each bucket; the other three block processors are identical to this.

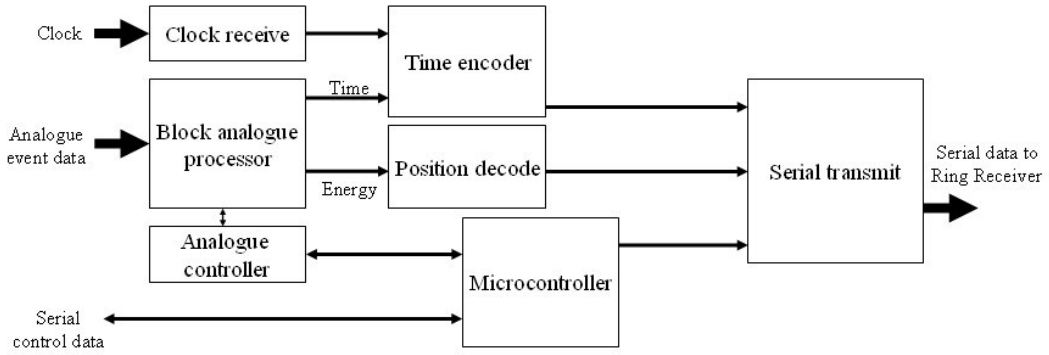

Figure 3.4: Block diagram for a detector bucket

The bucket controller is responsible for the overall orchestration of the event processing, transferring event data through the position energy processor and time digitisation, and presenting events into a readout buffer ready for transmission to the central coincidence processor described in section 3.3.

A number of different operating modes are supported: Run mode provides normal acquisition of events, position histogram mode is used to define the crystal boundaries, and energy histogram mode is used to calibrate the energy discriminators.

The firmware run by the bucket microprocessor upon startup is stored in a user programmable EEPROM where discriminator levels and amplifier gains are stored. It is simple to access and adjust these values using the serial communications link allowing for ease of calibration and hardware updates.

The bucket controller performs the housekeeping for the module. This includes a self-test on startup, loading the PMT gains and discriminator levels from the onboard EEPROM and loading the crystal position table into the lookup RAM.

It can be used to sample the count rate in each crystal segment, report corrected and uncorrected singles rates, set energy discriminator levels and perform detector calibration. The controller oversees resets, mode selection, programming of onboard EEPROMs and communications. Control is provided via the serial clock bus.

3.1.3 Front End Data

The bucket controller transmits data to a central processing unit in the form of a 16-bit serial data word every 256 ns. The format of this word for the 931 series buckets is:

- 6-bits time (T_0 - T_5) (0 - 127)
- Positional data
  - 2-bits block in bucket (0 - 3)
  - 3-bits transaxial segment in block (0 - 7)
  - 2-bits axial segment in block (0 - 3)
- 3-bits describing hit, multiple and scatter

Where event time is relative to the 256 ns period master clock: each bit then represents a time period of 4 ns. The positional data are given as a block within a bucket and then a crystal segment within the given block.

The hit bit describes data that has passed the energy discrimination test and describes a single event that the bucket has thus deemed valid. The hit bit is not set if the bucket has not had an event within the last 256 ns. The scatter bit describes events which have not passed the energy discrimination test and are therefore regarded as scattered events.

The multiple bit is a feature of the 931 series buckets only; it is set when more then one event was seen by the bucket within the 256 ns time period between generating data words.

In the later models (951) the multiple bit is not used for this purpose because the activity used in a clinical study is always too low to result in multiple events being seen in a single bucket over this time period. In the later models one extra bit is required to describe the extra axial crystal segments. The format of the data word remains similar to that of the 931 series, however the redundant multiples bit has been used to describe this extra positional data. The format of the data word for the 951 series buckets is then:

- 6-bits time (T_0 - T_5) (0 - 127)
- Positional data
  - 2-bits block in bucket (0 - 3)
  - 3-bits transaxial segment in block (0 - 7)
  - 3-bits axial segment in block (0 - 7)
- 2-bits describing hit and scatter

This data is transmitted in serial over 4 twisted pair cables in differential Emitter Coupled Logic (ECL) format.

3.2 Ring Receiver

A Ring Receiver consists of up to 16 identical input devices, driven by the data signals from the detector buckets. For a scanner with 32 detector buckets, two ring receivers are required. Each input consists of 4 twisted pair cables receiving serial data in differential Emitter Coupled Logic (ECL) format. The ring receiver converts the signal from ECL to Transistor Transistor Logic (TTL) and then performs a conversion from serial data into parallel data under the control of the master clock. The 4 twisted pairs carry data describing:

- Master clock signal 31.25 MHz (125 MHz / 4)
- Strobe clock signal 4 MHz (31.25 MHz / 8)
- High order byte 8-bits
- Low order byte 8-bits

The master clock signal moves the two data signals into two serial to parallel shift registers with depth of 8-bits. On the strobe signal (following 8 clock cycles) the data from the shift registers is latched into two 8-bit parallel registers. The following clock pulse places this 16-bit word onto the backplane of the VME bus where it is read by the coincidence processor. Thus every 8 clock cycles (256 ns) a new data word describing an event in each bucket is presented to the coincidence processor.

The coincidence processor master clock signal originates in one of the ring receivers. This clock signal is generated by the strobe signal with a 50% duty cycle. This results in a 4 MHz clock which is used to drive the coincidence processor. Figure 3.5 shows a block diagram for the ring receiver.

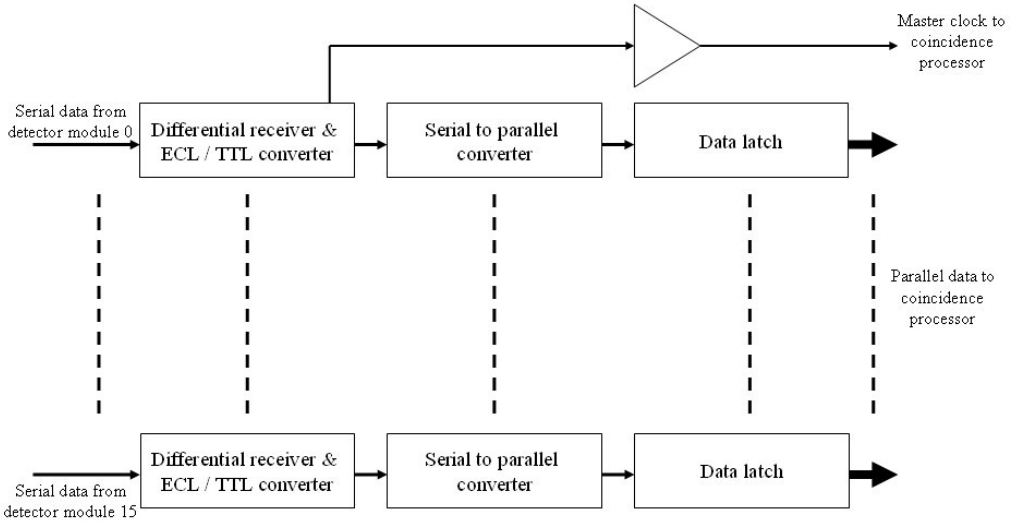

Figure 3.5: Block diagram of a ring receiver

3.3 Image Plane Coincidence Processor

The Image Plane Coincidence Processor (IPCP) receives data from each of the 16 bucket controllers in parallel via the ring receivers. This data describes the excited crystal element within a given block in each of the buckets, along with the time the event was detected. The bucket identification number is not directly encoded but is assumed from the ring receiver input number.

The coincidence processor is a pipelined parallel processing board which is hard wired to look for coincidence events between 56 possible pairs of detectors, described in section 3.3.1. The block diagram describing its operation is given in figure 3.6.

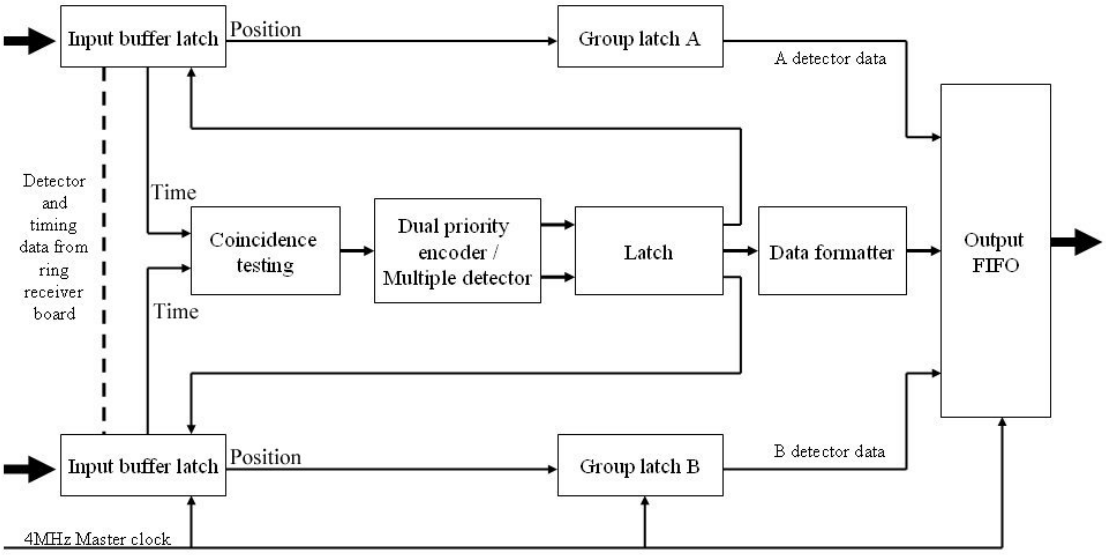

Figure 3.6: Block diagram of a coincidence processor

The incoming data is latched into a buffer on the first pulse from the 4 MHz master clock. The timing data and positional data are separated from the data word into different memory buffers where they are stored for the duration of the processing period.

The processor pipeline is split into 2 sections. The first part of the pipeline uses the time information to determine if any coincidences have occurred during the 256 ns time period using a 12 ns coincidence window (2τ) programmed into a set of dedicated PROMs. The second part of the pipeline chooses a single valid event, checks the positional data, and constructs the output word describing the processed LOR.

When one event is detected in a bucket, an event in the opposing buckets is sought within a time interval from 4 ns before the primary event to 4 ns after it. This yields the total coincidence time window of 12 ns (2τ), illustrated in figure 3.7.

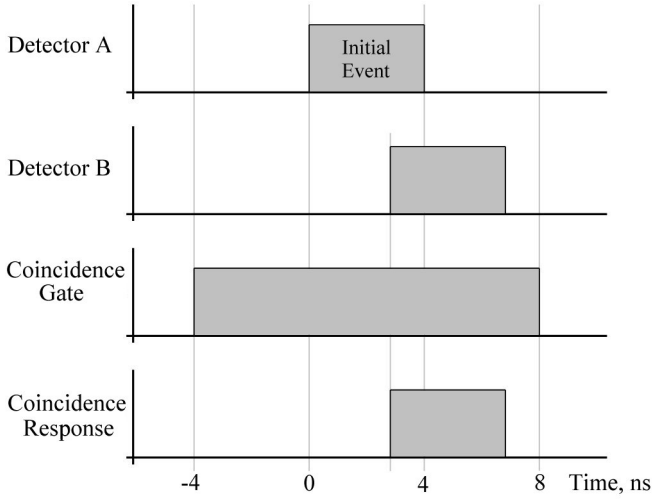

Figure 3.7: Prompt coincidence gate for ECAT IPCP

The processing of delayed (random) events occurs at the same time as the prompt coincidence test, this is handled by a further set of PROMs programmed to have a 12 ns coincidence window (2τ) which is delayed by 48 ns compared to the trigger pulse, described by Spinks *et. al.* in [33]. This is illustrated in figure 3.8 (note the scales are different on both of these diagrams), refer to section 1.3.1 for further discussion.

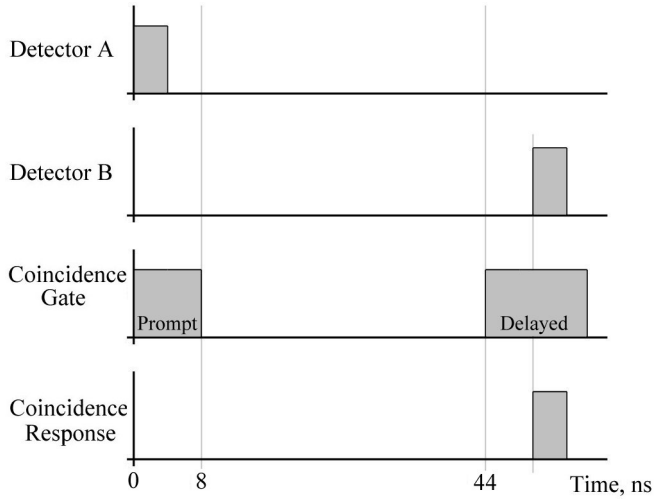

Figure 3.8: Delayed coincidence gate for ECAT IPCP

Events which are found to be in coincidence are flagged and the coincidence processor checks to see if there have been any multiple coincidence events. A multiple is defined as one bucket being in coincidence with two or more buckets, a prompt and a delay are not considered to be multiple events.

A dual priority encoder then chooses one and only one of the flagged coincident events for further processing. The dual priority encoder consists of two priority encoders with the 56 possible coincidence pairs listed in priority from 1 to 56. The second priority encoder mirrors the first, with the priority order reversed. Upon each new event (256 ns given by 4 MHz strobe clock) the priority encoder used is toggled. There is a dual priority encoder for each detector, A and B.

The second part of the pipeline receives the two tagged events which have been determined to be a valid coincidence. It refers to the first latch and recovers the positional data for the two tagged events. The data are tested for coincidence again and whether the event is a prompt or delayed coincidence is determined. The positional data is processed to determine that the events occurred between a valid pair of detectors (as 2D or 3D imaging limits the axial plane numbers). Finally the module pair number describing the bucket combination (section 3.3.3) is determined.

The data describing the coincident event is then latched into a 32-bit wide output FIFO under the control of a handshake with the Real Time Sorter (RTS) system, described in section 3.4. Events are produced at the FIFO output at a maximum rate of 4 MHz for each of four coincidence processors. The first processes ring 0 data, the fourth processes ring 1 data and the second and third process data occurring across both rings. The format of the output data word is:

- Transaxial position in bucket (crystal in bucket)

  5-bits crystal A (0 - 31)

  5-bits crystal B (0 - 31)

- Axial position in bucket

  2-bits (931) / 3-bits (951) plane A

  2-bits (931) / 3-bits (951) plane B

- 6-bits module pair number (giving bucket identification 3.3.3) (0 - 63)

- 2-bits identify the coincidence board (giving ring identification) (0 - 3)

- 3-bits multiple, delayed and scattered

This data word is transmitted in differential TTL format to the RTS system (section 3.4) following a successful handshake process. This is described in further detail in section 4.1.1. Note that there is no time information included in the final output describing coincidence words: time of arrival gives the approximate (to within 256 ns) time of the event.

3.3.1 Coincidence Data

It is necessary to describe the manner in which event data is represented and the format used in storage. A LOR is typically described as a vector with components along three orthogonal axes, to describe this vector the (x, y, z) locations for each of the endpoints are recorded. For PEPT data, time information also has to be saved. The scanner system contains a gantry holding the 32 bucket controllers with the electronics used for the running of the scanner and processing and storage of event data. Overall the gantry holds 2 separate rings of 16 buckets in each, arranged in such a manner that there are 8 (931) or 16 (951) rings of crystal segments axially and 512 crystal segments transaxially around each ring.

In the original scanner geometry, the coordinate system can be represented in a similar manner to cylindrical polar coordinates. In this the data is broken down into 2 components; the angle given in crystal numbers around the ring and axial data representing the depth into the ring. This is illustrated in figure 3.9.

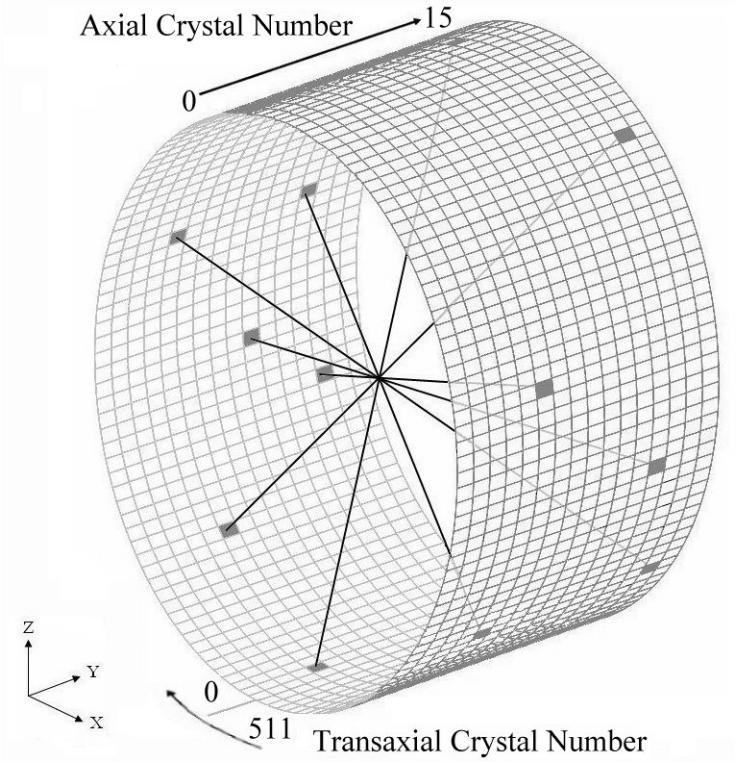

Figure 3.9: Crystal identification convention for 951 scanner

For PEPT analysis these data are converted into a Cartesian vector describing the end points of each LOR and assigned a time given by the time of arrival of the event into the data storage system: $(x_1, y_1, z_1; x_2, y_2, y_3; t)$.

Buckets that are directly opposite each other in the transaxial plane operate in coincidence along with 1, 2 or 3 neighbours on either side: a total of 7 buckets. This is illustrated in figure 3.10. The coincidental region swept out by bucket zero is highlighted, with a uniform field of view occurring where the regions from each bucket overlap.

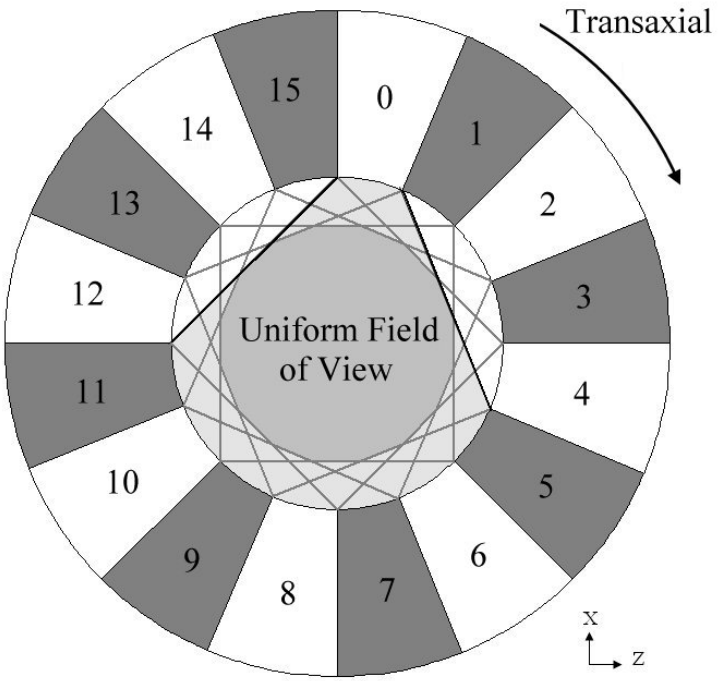

Figure 3.10: Possible coincidental planes in the transaxial orientation

The transaxial data represents the end points of the LOR in two dimensions around the scanner ring. In order to keep the data word describing the LOR below 32-bits, a compression technique is used. Each possible pair of detectors buckets is assigned a unique module pair number:

In a system with 16 buckets per ring with each bucket in coincidence with the 7 opposing buckets (the maximum allowed) there are a total of 56 possible transaxial bucket combinations, this is given by :

$$\frac{7 \times 16}{2} = 56 \tag{3.1}$$

The factor of one half is included because there is degeneracy across the scanner ring. For example bucket A being in coincidence with bucket B is indistinguishable from bucket B being in coincidence with bucket A. The module pair number which describes both of the buckets receiving the event can then be represented as a 6 bit number (0 - 63).

In the original scanner geometry an individual segment can be located as described in figure 3.11. The process is as follows:

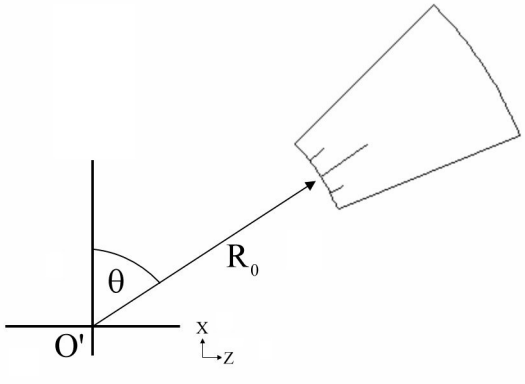

Figure 3.11: Segment identification process

For the origin, O', placed at the centre of the scanner ring, an arbitrary bucket is positioned at angle θ relative to the vertical, at distance R_0 from the origin. The distance R_0 is the original radius (510 mm) of the scanner gantry. As there are 16 buckets in a full transaxial ring each bucket subtends an angle of $2\pi/16$ radians. With bucket numbers increasing in the clockwise direction, with bucket zero centred on the vertical axis, the angle describing the centre of each bucket is given by the bucket number (0 - 15) multiplied by this factor.

An arbitrary segment within the given bucket is then given by the addition of a further angle describing the segment within the bucket, this is described in a similar manner. In a full ring there are 512 segments, therefore each segment subtends the angle $2\pi/512$ radians. Each bucket contains 32 segments numbered 0 - 31. The centre of the zeroth crystal clockwise around the bucket is then positioned at -15.5 multiplied by this factor.

The centre of an arbitrary segment is then located at an angle $(Segment - 15.5)2\pi/512$ in the bucket.

Segment location across the whole gantry then includes the bucket location:

$$\theta = Bucket.\frac{2\pi}{16} + (Segment - 15.5).\frac{2\pi}{512} \tag{3.2}$$

In Cartesian coordinates the X and Z locations are then given relative to the centre of the scanner gantry, O', by the relations:

$$\begin{aligned} X &= R_0 \cos\theta \\ Z &= R_0 \sin\theta \end{aligned}$$

In the axial direction direct coincidences are defined when events occur in the same ring; giving 8 direct planes. Cross coincidences occur with events in adjacent rings; leading to a further 7 coincidental planes; these modes are used to perform 2 dimensional imaging of a particular subject. The format is configured in the coincidence processor and events are rejected if they do not lie between accepted crystal rings.

For 3 dimensional imaging the axial sampling is increased to allow coincidences between all rings of the system. As the coincidence processor only relies on the time of each event to determine its response this means all events are then passed on for further data processing. Figure 3.12 illustrates the allowed axial combinations for the 931 scanner: all planes are illustrated as just one crystal operating in coincidence with the opposing 8 crystals for clarity.

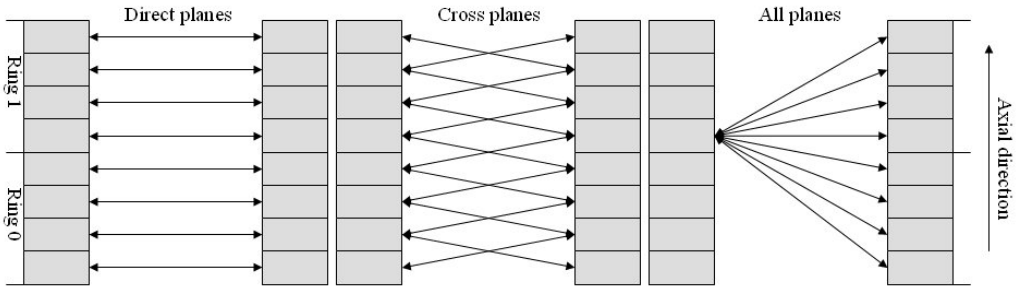

Figure 3.12: Possible image planes in the axial direction (931)

With each bucket operating in coincidence with the maximum of 7 opposing modules, with 32 transaxial crystals per bucket there are 224 possible transaxial combinations per crystal. In a full system with 16 buckets per ring there are 114688 total combinations, with degeneracy this results in 57344 possible LORs.

A typical sinogram of 65536 (256×256) pixels is then sufficient to store the data for one transaxial slice. Axial data is combined into 15 (for the 931) separate slices describing the activity at 15 positions along the axial axis. A data set describing all of the possible LORs in the camera then contains 15 sinograms of 256×256 pixels.

If the data storage type used for each pixel element is of integer type (4 bytes) then the storage space required for one sinogram is 262 kBytes, and for the entire data set is 3.93 MBytes. Using integer type storage allows a maximum count above 4×10^9 events per pixel resulting in Poisson error of 0.0015 percent.

The equivalent amount of memory can be used to store around 100k events (4 bytes each) in list mode. The standard 32 MByte buffer can then be used to store up to 7 complete sets of sinograms at any one time, or 8 million events (of 32 bits each) in list mode.

At the design maximum data rate of 16 MHz (total from all 4 IPCPs) the 32 MByte buffer would be sufficient to store list mode data for 0.5 seconds, but is sufficient to store many weeks of data as a sinogram.

A separate set of sinograms is constructed to hold the data from random and multiple events. After acquisition these data are subtracted from the prompt data set in order to give a statistical correction for random events that have occurred during the acquisition. This subtraction can be configured to occur during the acquisition, thus saving the extra memory required for the random event data.

3.3.2 Coincident Pair Representation

The four IPCP boards process coincident events for the different ring combinations allowed in the scanner gantry. Each IPCP is functionally identical, with 16 separate inputs configured to handle coincidences between 56 combinations of these inputs. The event data presented to each coincidence processor follows a different route from the ring receivers, presenting data from different detectors to each processor. Each IPCP board processes events across all axial ring combinations (4 for the 931, 8 for the 951) within the physical ring of the bucket.

The first IPCP board; given the serial address 64, processes events that occur between buckets located in the first ring. Hence data from buckets 0 - 15 is transmitted to inputs 0 - 15 respectively.

The last IPCP board; with the serial address 67, processes events occurring in the second ring. Hence data from buckets 16 - 31 is transmitted to inputs 0 - 15 respectively.

Events with one end of the LOR in the first ring and the other end of the LOR in the second ring are processed by IPCP board 65, and vice-versa for IPCP board 66. In the original configuration bucket 0 sits on top of bucket 16, therefore the two buckets are equivalent in the transaxial direction. The set of buckets in the second ring (16 - 31) are then equivalent to those in the first (0 - 15).

Each IPCP board then reads the input from the ring receiver as originating in buckets zero to fifteen. Thus an LOR between buckets 0 & 8 is equivalent to buckets 16 & 8, buckets 0 & 24 and buckets 16 & 24. All these combinations are represented as being between buckets 0 & 8 and are designated the module pair number 4 (see section 3.3.3). The distinction between the 4 combinations is then given by a further two bits of data representing the address of the IPCP board.

This is illustrated in figure 3.13 where the two rings are shown with their respective bucket numbers, the tube of response describing events between buckets 0 & 8, 16 & 24, 0 & 24 and 16 & 24 is shown in green.

In the transaxial direction the different ring combinations cannot be distinguished, however in the axial direction the IPCP board number which processed the event gives these combinations.

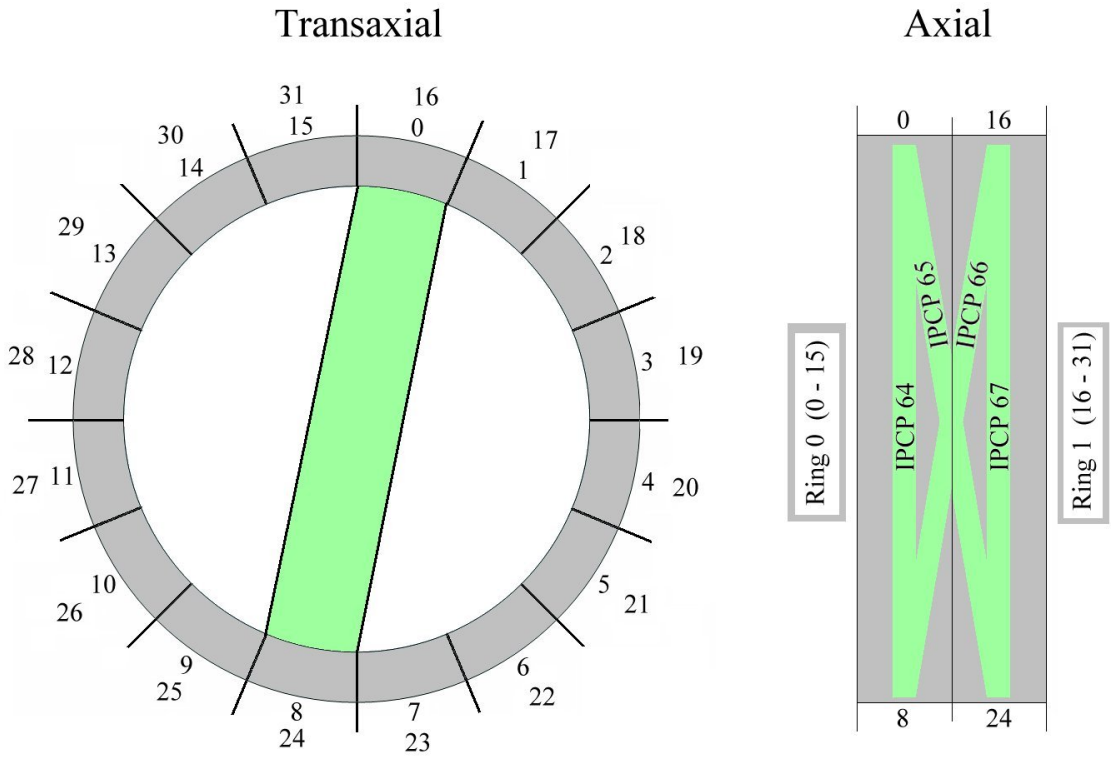

Figure 3.13: Coincident pairing across detector rings

Given for reference are the coincident combinations for all of the possible coincidence pairs in a system with the maximum amount of buckets (32). Starting with the serial address of the IPCP board concerned with the acquisition of each pair, the table displays the bucket number with its coincident pairs in each column.

IPCP 64:

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
5	6	7	8	9	10	11	12	13	14	15	0	1	2	3	4
6	7	8	9	10	11	12	13	14	15	0	1	2	3	4	5
7	8	9	10	11	12	13	14	15	0	1	2	3	4	5	6
8	9	10	11	12	13	14	15	0	1	2	3	4	5	6	7
9	10	11	12	13	14	15	0	1	2	3	4	5	6	7	8
10	11	12	13	14	15	0	1	2	3	4	5	6	7	8	9
11	12	13	14	15	0	1	2	3	4	5	6	7	8	9	10

IPCP 65:

0	1	2	3	4	5	6	7	8	9	10	21	22	23	24	25	26	27	28	29	30	31
21	22	23	24	25	26	27	28	29	30	31							0	1	2	3	4
22	23	24	25	26	27	28	29	30	31							0	1	2	3	4	5
23	24	25	26	27	28	29	30	31							0	1	2	3	4	5	6
24	25	26	27	28	29	30	31							0	1	2	3	4	5	6	7
25	26	27	28	29	30	31							0	1	2	3	4	5	6	7	8
26	27	28	29	30	31							0	1	2	3	4	5	6	7	8	9
27	28	29	30	31							0	1	2	3	4	5	6	7	8	9	10

IPCP 66:

5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26
						16	17	18	19	20	5	6	7	8	9	10	11	12	13	14	15
					16	17	18	19	20	21	6	7	8	9	10	11	12	13	14	15	
				16	17	18	19	20	21	22	7	8	9	10	11	12	13	14	15		
			16	17	18	19	20	21	22	23	8	9	10	11	12	13	14	15			
		16	17	18	19	20	21	22	23	24	9	10	11	12	13	14	15				
	16	17	18	19	20	21	22	23	24	25	10	11	12	13	14	15					
16	17	18	19	20	21	22	23	24	25	26	11	12	13	14	15						

IPCP 67:

16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
21	22	23	24	25	26	27	28	29	30	31	16	17	18	19	20
22	23	24	25	26	27	28	29	30	31	16	17	18	19	20	21
23	24	25	26	27	28	29	30	31	16	17	18	19	20	21	22
24	25	26	27	28	29	30	31	16	17	18	19	20	21	22	23
25	26	27	28	29	30	31	16	17	18	19	20	21	22	23	24
26	27	28	29	30	31	16	17	18	19	20	21	22	23	24	25
27	28	29	30	31	16	17	18	19	20	21	22	23	24	25	26

3.3.3 Module Pair Numbers

As discussed, each coincident pair has been assigned a 'module pair number' which identifies the two buckets providing the end points of each LOR.

Each coincidence processor handles events occurring in one of 56 different bucket combinations; the storage convention for LORs occurring in each bucket can be reduced to a six bit number (0-63). In actual fact only 56 different states are required as seen in section 3.3.1.

Given in table 3.3.3 are the module pair numbers which represent each of the possible 56 bucket combinations. Previously knowledge of the module pair numbers was unnecessary as the data was formatted into a known form by subsequent processes in the RTS system. This data was determined by a series of experiments by acquiring data using only the module pair in question.

Given in the table are the low order bucket numbers and their corresponding module pair numbers. It is often necessary to determine the same data for the second ring of buckets and cross ring combinations so the high order bucket numbers are given in brackets.

Any combination of the bucket numbers describing Bucket A with any of the bucket numbers describing Bucket B then have the unique module pair combination given (but the IPCP board identifier will be different).

Bucket A		Bucket B		Module Pair Number
0	(16)	5	(21)	1
0	(16)	6	(22)	2
0	(16)	7	(23)	3
0	(16)	8	(24)	4
0	(16)	9	(25)	5
0	(16)	10	(26)	6
0	(16)	11	(17)	7
1	(17)	6	(22)	8
1	(17)	7	(23)	9
1	(17)	8	(24)	10
1	(17)	9	(25)	11
1	(17)	10	(26)	12
1	(17)	11	(17)	13
1	(17)	12	(18)	14
2	(18)	7	(23)	15
2	(18)	8	(24)	16
2	(18)	9	(25)	17
2	(18)	10	(26)	18
2	(18)	11	(17)	19
2	(18)	12	(18)	20
2	(18)	13	(19)	21
3	(19)	8	(24)	22
3	(19)	9	(25)	23
3	(19)	10	(26)	24
3	(19)	11	(17)	25
3	(19)	12	(18)	26
3	(19)	13	(19)	27
3	(19)	14	(20)	28
4	(20)	9	(25)	29
4	(20)	10	(26)	30
4	(20)	11	(17)	31
4	(20)	12	(18)	32
4	(20)	13	(19)	33
4	(20)	14	(20)	34
4	(20)	15	(21)	35
5	(21)	10	(26)	36
5	(21)	11	(17)	37
5	(21)	12	(18)	38
5	(21)	13	(19)	39
5	(21)	14	(20)	40
5	(21)	15	(21)	41
6	(22)	11	(17)	42
6	(22)	12	(18)	43
6	(22)	13	(19)	44
6	(22)	14	(20)	45
6	(22)	15	(21)	46
7	(23)	12	(18)	47
7	(23)	13	(19)	48
7	(23)	14	(20)	49
7	(23)	15	(21)	50
8	(24)	13	(19)	51
8	(24)	14	(20)	52
8	(24)	15	(21)	53
9	(25)	14	(20)	54
9	(25)	15	(21)	55
10	(26)	15	(21)	56

Table 3.3.3. Module pair numbers for all coincident combinations

3.4 Real Time Sorter

In a conventional system the coincidence data word is transmitted from the image plane coincidence processors to the Real Time Sorter (RTS) system held in a second VME rack. This transmission is done over a ribbon cable (the format of which is discussed in section 3.5) in differential TTL format following a handshake process which is initiated by the IPCP.

The RTS system formats the data word into a form suitable for sinogram formation and image processing and places the data in a local memory buffer. When constructing a sinogram this formatted data are placed directly into sinogram memory where image processing and randoms correction can take place in real time. In list mode the data are stored sequentially into a double buffer with occasional bursts written to disk.

The system has a fixed physical disk drive and is allowed both read and write access via suitable VME controllers. The write access allows the event data (either as a set of sinograms or as list mode) to be written to the disk for permanent storage. Read access allows data which has been previously stored on the disk to be loaded into the VME memory for image analysis and processing post acquisition [31].

In the original system the 68020 host CPU was used for control of the entire system. Byars and Sadrmomtaz upgraded this system in 2003 [34] and replaced the Mizar processor with a SPARC CPU (Scalable Processor ARChitecture Central Processing Unit). This oversees the running of the scanner and communicates between the operator and the bucket controllers. Its tasks include:

- Interaction with the operator
- Routing jobs to various ECAT subsystems
- Communication with bucket controllers over clock serial link
- Gantry control system
- Monitors data and image processing systems

The upgrade allows list mode data storage to be performed and inserts timing data into the list mode stream as data are acquired: this is essential for PEPT work. The time information is inserted using a flagged data word every 1 ms. The value of the timing tag word gives an absolute value of time since the acquisition began.

The overall system architecture is shown in figure 3.14. It shows the front end data system with the two serial buses, clock control and data, for one bucket. The ring receivers and image plane coincidence processor share a VME rack with 64 bit wide data paths. The data from each ring receiver is sent along this bus to each of the coincidence processors in parallel, each coincidence processor receives data from a different combination of buckets in the system.

The coincidence processor transmits coincidence words to the RTS unit which is situated in a VME rack with a 32 bit wide data path. This bus is indirectly controlled by the SPARC CPU, it allows the dedicated data collection and image processing boards to access VME memory and hard disk.

All of these systems are controlled by the operator's console which communicates with the RTS CPU over a standard ethernet link where remote procedure calls can be initiated and memory and hard disks accessed.

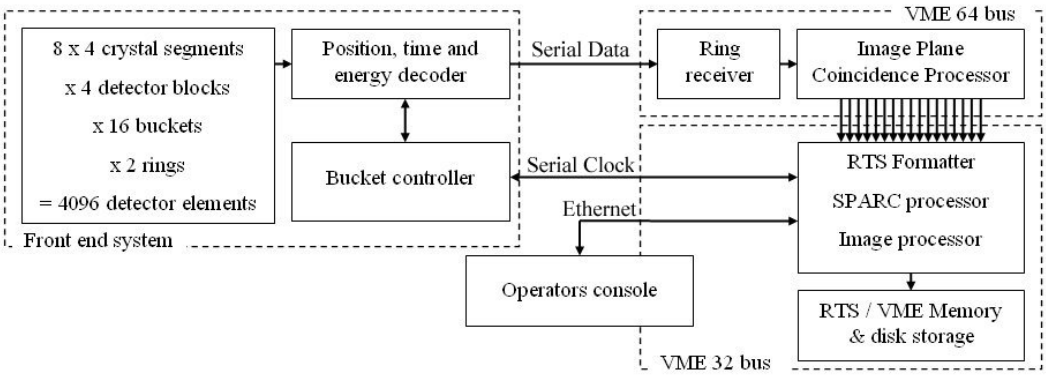

Figure 3.14: Architecture of the ECAT scanner system

3.5 Data Storage Conventions

Analysis of the circuit diagrams [35] and supporting documents [19] is not sufficient to describe the data word formats in full. For this information to be determined a number of studies have had to be performed using collimated point and line sources to excite specific crystals in the system and thus determine the data word representing activation of the particular crystal.

Using a broadcast source to excite all of the crystals in a single pair of buckets, with the rest of the buckets in the system disabled, gives the most significant bits of the data word, those describing the bucket pair that was used. The same can be done for determining the data describing individual blocks in the system: by disabling all but the block in question the block number can be determined. Rows and columns of crystal segments were then excited using a collimated line source: by exciting all rows at once, the row identification could be performed, and similarly for column data.

It is the purpose of this section to outline the storage conventions and data word formats used in moving data from the IPCP to eventual disk storage. Chapter 5 describes the software used to perform these tasks. Presented here are the word formats for each of the ECAT scanner systems at the Positron Imaging Centre.

3.5.1 931 Scanner

The original 931 system recorded event data after some initial processing via the RTS system. In this, the data are formatted into a 32-bit word describing the Angle and Offset of each LOR (see section 1.4). The system has been modified by Sadrmomtaz *et. al.* [22] to write this formatted event word to disk in list mode, where events are written sequentially to disk, rather then conventional storage as a sinogram.

Data which is stored to disk using the original 931 RTS system then consists of a 32-bit word describing the LOR for each event in terms of Angle and Offset. For PEPT further software is used to convert this Angle / Offset data back into data describing the individual crystal segments that were excited in each event. Figure 3.15 shows the format of the data word as written to disk.

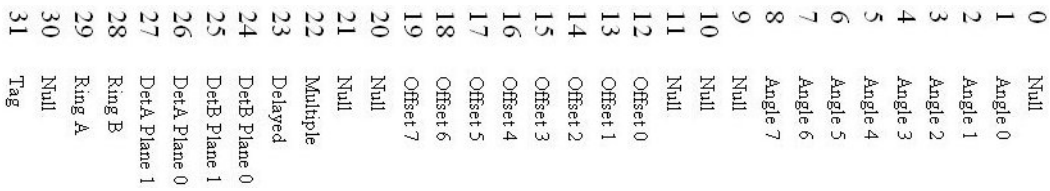

Figure 3.15: 32-bit data word from 931 RTS formatter

As this system uses a UNIX operating system for the operator's console and SPARC CPU the binary word is stored in Big-Endian format; this has to be converted into Little-Endian if analysis using a computer with the Windows operating system is required.

The event datum is converted from the Angle / Offset storage convention into the absolute crystal numbers describing the LOR and then into Cartesian coordinates describing the crystal position. First the Endian conversion is performed, then using bit operators the axial positional data and the angle and offset data are determined directly from the data word. A function is then used to determine the crystal number in the ring from the sinogram data, the conversion into a crystal number in a detector bucket and finally the spatial coordinates of the crystal element.

Development of the 931 system has led to a need for the format of the data word describing event data from each detector bucket, and the event data word describing each coincident LOR to be known prior to formatting by the RTS. Consideration of the circuit diagrams can be used to partially derive this information, studies using collimated sources were used to finalise and confirm this derivation.

In figure 3.16 the data word describing a coincident LOR at the output FIFO of the 931 IPCP is described. This word is presented at the output FIFO in parallel and transmitted over a 60 conductor ribbon cable. This cable is setup with the even numbered conductors carrying a reference ground and odd numbered conductors carrying signal data. Included in this are 4 data cables which carry the handshake control lines.

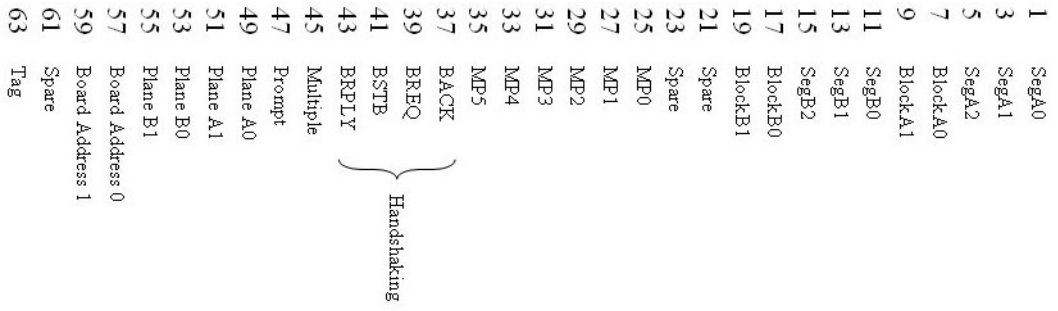

Figure 3.16: Data word describing ECAT 931 LOR at the output FIFO

3.5.2 951 Scanner

The 951 system is similar in operation to the 931 system, however the positional event data requires an extra 2 bits to describe the extra axial crystal segments. The format of the coincidence processor output word has changed to allow this extra data.

Once again data is transmitted in differential TTL format, but in the 951 system this is over an 80 conductor ribbon cable. The extra data lines allow the transmitted data word (with extra 2 bits), the four handshaking control lines, and the serial communications data to be transmitted together. Figure 3.17 shows the format of the data word for the 951 coincidence processor.

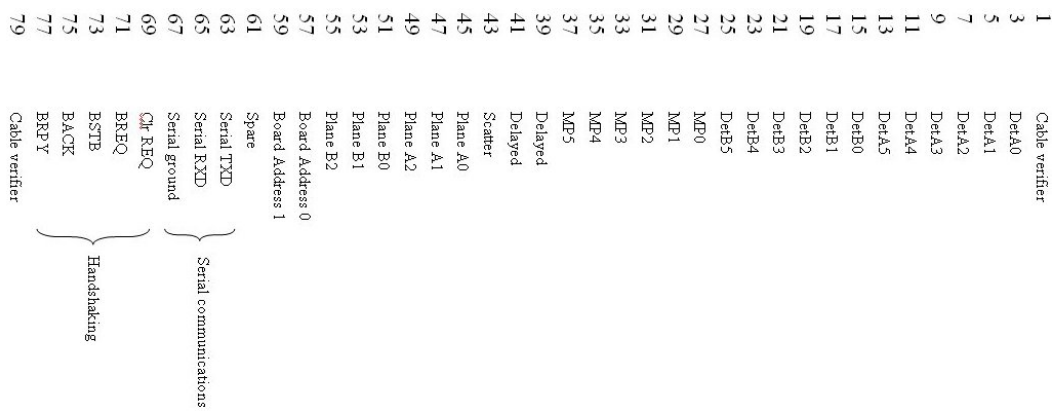

Figure 3.17: Data word describing ECAT 951 LOR at the output FIFO

3.6 Combining the Systems

Due to the age of the components and the availability of spare parts it has been beneficial to determine whether components from the two different scanner systems can be combined. Further advantages of this are that the available systems are then very flexible in their relative abilities.

The first test determined whether it was possible to use the detector blocks from the 951 system (8×8 crystal segments) with the older 931 bucket controllers (expecting 8×4 crystal segments). This was mainly motivated by the lack of spare 931 blocks and the fact that many have broken crystals or PMTs. A study was performed in a similar manner to that described in section 3.5, using a collimated source to excite individual crystal segments in each block. As expected the increased positional resolution from the extra crystals was lost as the data was processed to be a 8×4 device. The transaxial data was seen to remain unchanged from the original system, as expected because both systems have eight transaxial crystals per block.

In the axial direction the bucket controller was seen to process events occurring in the extra crystals into the same format as the old system; effectively reducing the axial resolution to four large crystal segments, as illustrated in figure 3.18.

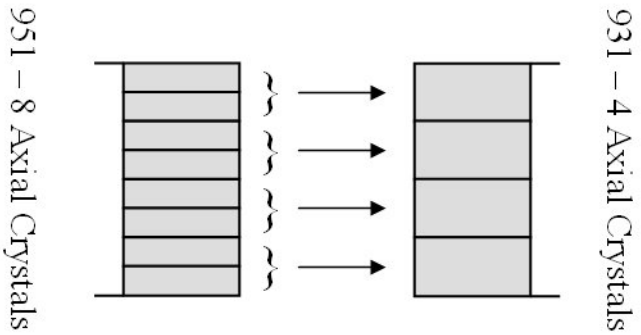

Figure 3.18: 951 block as seen by 931 bucket controller

It did not seem necessary to test how the 951 bucket controller saw the older 931 blocks, this was because of the lack of spare 931 blocks and the fact that the newer controllers are designed for use with the higher resolution newer blocks.

Analysis of the circuit diagrams for both models of coincidence processor led to the conclusion that it would be possible to use the newer 951 buckets with the older 931 coincidence processor system. The ring receivers for both systems are functionally and schematically identical, and the coincidence processors for both models are fairly similar.

After testing the 931 coincidence processor with the inputs driven by the newer 951 buckets, coincidence events were seen to be written to disk. Unfortunately the extra positional data (describing the extra four axial rings of crystals) was lost as it passed through the formatter board in the RTS unit. This motivated the development of a new data acquisition system to replace the RTS for both systems, described in chapter 4. It was also seen that the older bucket controllers operated correctly with the newer coincidence processors; however this combination has little potential benefit for PET or PEPT studies.

The discussion given in section 3.1.3 regarding the multiple bit shows that the incoming data contains the relevant positional data in a different form to that which is expected by the IPCP. However, as only the time information is used (and this remains in the same format) to determine a coincidence the 931 coincidence processors can still be used with 951 buckets. The difference results in the multiple bit on the output word from the IPCP being set whenever an event occurs in the extra crystals (as this bit now carries the extra positional data).

The positional data streams through the buffer system unaffected, and the multiple / axial crystal position data can also be recovered from the system; i.e. it appears as 2 data bits where in the previous system these bits were left unconnected (spare). Unfortunately this extra data is lost as it enters the formatter board of the RTS unit. The format of the data word output by a 931 series coincidence processor, operating with 951 series detector buckets is given in figure 3.19.

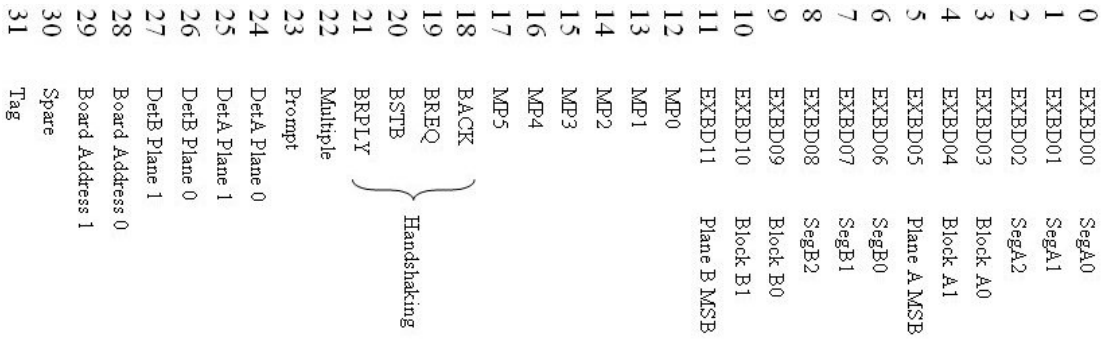

Figure 3.19: Data word describing ECAT 931 LOR

Chapter 4

Data Acquisition Systems

The original ECAT data acquisition system (the Real Time Sorter (RTS), described in section 3.4) was custom designed using off-the-shelf components in a compromise between speed and cost. The design described by Jones *et. al.* [31] allowed a maximum data rate of 1 MHz transmission of data words from the Image Plane Coincidence Processor (IPCP) to the data storage and processing system (RTS). Data storage rates for the original systems could be improved by using two separate RTS acquisition systems operated in parallel, however this came at the expense of added cost and complexity.

Primarily the speed limitations were due to memory access times in the RTS, as in the medical application the data are formatted into a sinogram which is constructed in system memory, with occasional transfer to permanent disk storage. However storage of the data stream directly to disk in list-mode was significantly slower due to large disk access times; this is seen in the performance data for the small ring camera (section 7.1) where the maximum data storage rate achieved is approximately 700 kHz.

Compare this to the maximum theoretical data rate at the output FIFO of the coincidence processor under the control of the master clock: data can be presented at the output FIFO every 256 ns or at a rate of 4 MHz. This is further complicated by the use of multiple coincidence processors. In a full system 4 would be used, giving an output rate of 16 MHz. As each output word is of size 4 bytes (32-bits in length) each coincidence processor streams data at a rate of 16 MBytes per second. The whole system operating at maximum speed would then produce up to 64 MBytes of event data per second for processing and disk storage.

In a clinical application it is doubtful that the activity required to generate this amount of data would ever be injected into a patient due to radiological protection arguments. Thus limitations to the original design are not seen to be factors affecting the operation of a clinical study, and subsequent developments for medical application have seen improvements in efficiency and spatial resolution instead of high speed data acquisition. For PEPT experiments the activity is not limited by radiological protection considerations (aside from exposure of the experimentalists) and therefore a much higher activity source can be used, thereby providing frequent and precise PEPT locations (section 2.3).

The demand for high speed list mode data storage required for PEPT experiments has therefore motivated the development of an improved data acquisition system to replace the original RTS device. This system is based around the PCI-7300A Digital I/O Card manufactured by Adlink Technologies inc. It features 80 MBytes per second, 32-bit data transfer via direct memory access (DMA) along the PCI bus, with a further 8-bits auxiliary I/O and simple handshaking/trigger process capability.

In order to interface with the existing coincidence processor system a small number of dedicated support circuits have been designed and built. These use TTL technology and deal with the handshaking requirements, timing tag word insertion and support for remote trigger devices.

This chapter describes the design of the data acquisition system and associated electronics in detail. Given as an introduction figure 4.1 shows a block diagram for the camera system as a whole.

The detectors provide data to the central Image Plane Coincidence Processor where the coincidence response is determined. The support electronics oversee communication between the IPCP and the data acquisition system built into a desktop computer, and are used to insert further data (time and remote signal) into the data stream. The computer also provides control signals (described in section 5.1) to the detectors and IPCP electronics.

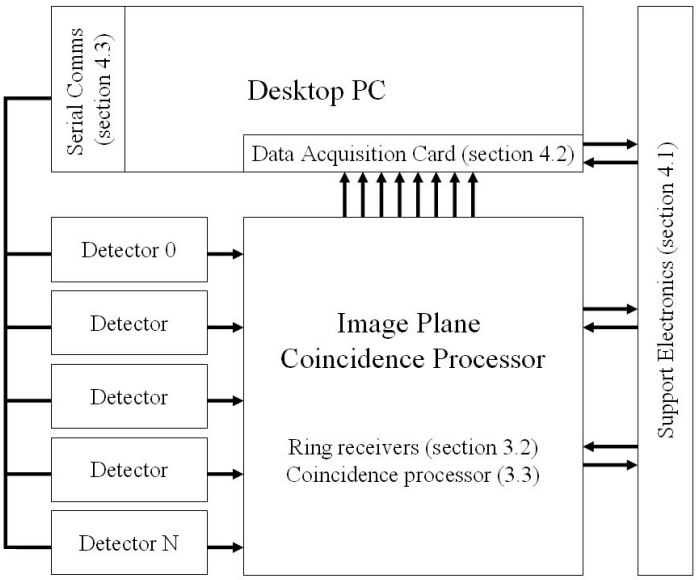

Figure 4.1: Top level block diagram for the camera system

A high level block diagram for the components of the data acquisition system is given in figure 4.2. Here the three subsystems are given in more detail with the connecting signals between modules illustrating the data flow.

The support electronics are broken into four sections dealing with signal communication (handshaking), the inclusion of time and remote trigger data, and a trigger system used to capture data.

The coincidence processor output data lines are connected to a FIFO built into the PC data acquisition card, further data from the support electronics are included here. The handshake driver is used to activate the output of the coincidence processor when event data are ready to be captured, and controls the data capture onto the PC.

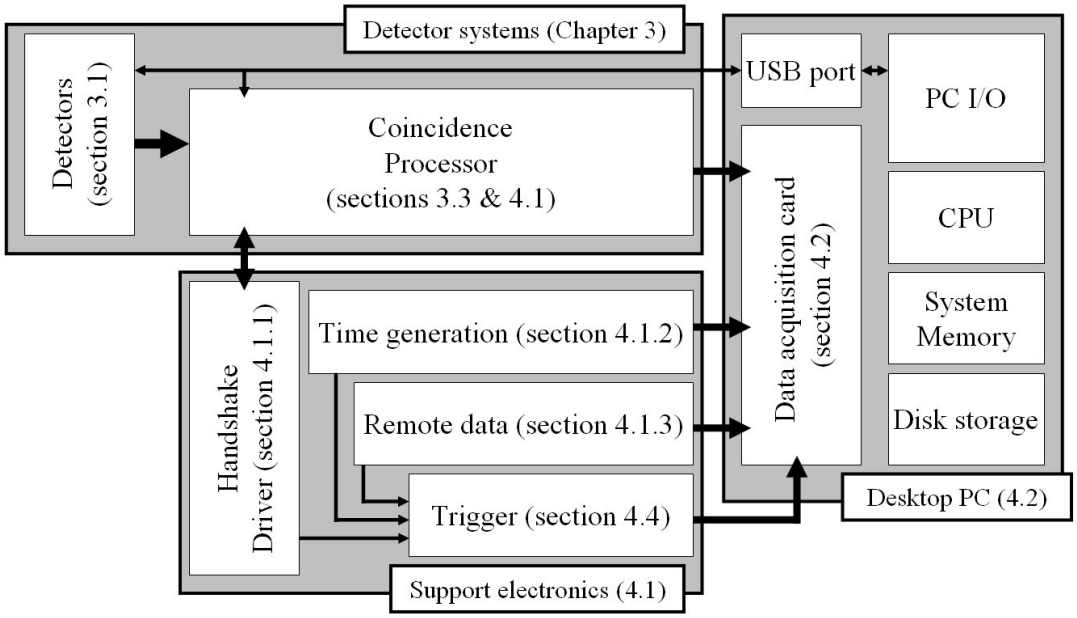

Figure 4.2: Block diagram for the data acquisition system

Finally figure 4.3 illustrates the connecting signals between subsystems. The notation given is used throughout the remainder of this chapter, given with reference for each section.

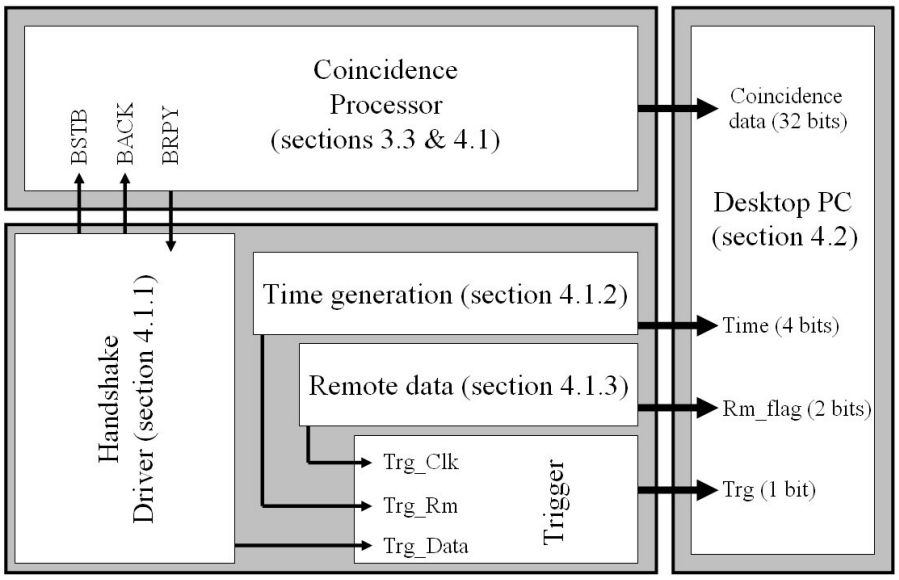

Figure 4.3: Block diagram showing data signals

The Adlink PC card and support circuits make a modern PC based, high speed, data acquisition system with superior performance compared to the original RTS system. Advances in computer technology (specifically speed and cost) have thus made it possible to replace the dedicated hardware comprising the RTS unit with a single desktop PC and off-the-shelf components at minimal cost.

This system is advantageous because, as will be shown, it can operate at a rate competitive with (or superior to) the original system, has far less complexity and offers much in the way of flexibility for both data storage and processing.

4.1 Acquisition Electronics

A block diagram for the output FIFO and control of the IPCP is shown in figure 4.4; the corresponding circuit diagram is given in figure 4.5. The output FIFO on the IPCP consists of eight parallel asynchronous FIFOs with 16-bit deep, 4-bit wide memories (giving 16 by 32 bit data words). These Integrated Circuits (ICs) of type 74ALS232; offer up to 40 MHz data transfer rate, independent input and output clocking and have 3-state outputs.

The 3-state output allows the ICs to be directly connected to the output bus which is shared by all four coincidence processors. In the 3rd output state the output lines are allowed to float between logic levels at high impedance, this is when the output enable flag (OE) is held low, during this time other devices (the rest of the coincidence processors) can drive logic levels on the bus. If the OE flag is high, the output is enabled and the device drives the bus with logic 0 or 1 depending upon the output word.

A handshake process is used to control which of the devices sharing the data bus is allowed control (i.e. allowed to operate its FIFOs with true logic states). The signals for the handshake are transmitted serially through each of the IPCPs in turn, starting at IPCP 64, then through IPCP 67, to IPCP 65 and finally IPCP 66.

If the current coincidence processor has data, the handshake signal is used to activate the FIFO outputs on that particular board; during this time the handshake signals are prevented from reaching the remaining devices. If the board does not have data ready for collection it passes the handshake signals onto the next board down the line.

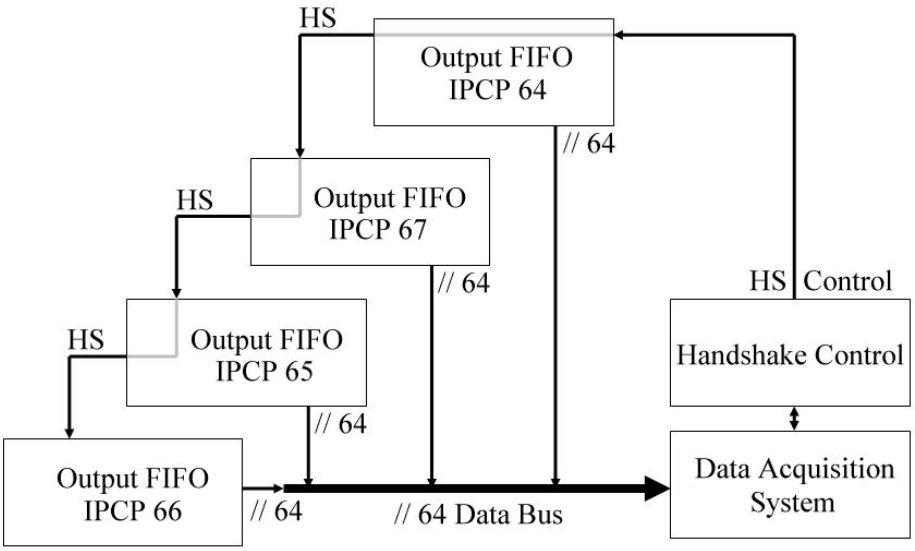

Figure 4.4: Coincidence processor output FIFO and control block diagram

Data arrives from the coincidence processor at random intervals (dependent upon data rate) to each of the 4 inputs on each FIFO IC. A following clock pulse on the LDCK (load clock) pin inserts each word into the FIFO memory provided that the memory is not full. If the memory is full the LDCK signal has no effect and the data is lost. It is therefore important to be able to empty the FIFO at a higher rate then data arrives, slower operation results in the loss of data and limits the coincidence count rates.

Once event data is loaded into the FIFO, the OR (output ready) flag is driven high, indicating that there is data ready to be read from the FIFO. The OR flags from each of the ICs making up the full 32-bit output word are combined: provided all ICs have data ready this sets a flag telling the handshake driver that data is ready. The independent clocking for both input and output allows the data to be temporarily stored until the readout system is ready to capture it.

The handshake process is then initiated (described in section 4.1.1), if this is successful the UNCK signal is sent from the handshake driver to the FIFO ICs which then assume control of the data bus and present the data state of the stored memory word at their output. This state is held until the data acquisition system has captured the data and released the handshake signal. The FIFO ICs then revert to the 3rd output state allowing the other devices to drive the output bus lines.

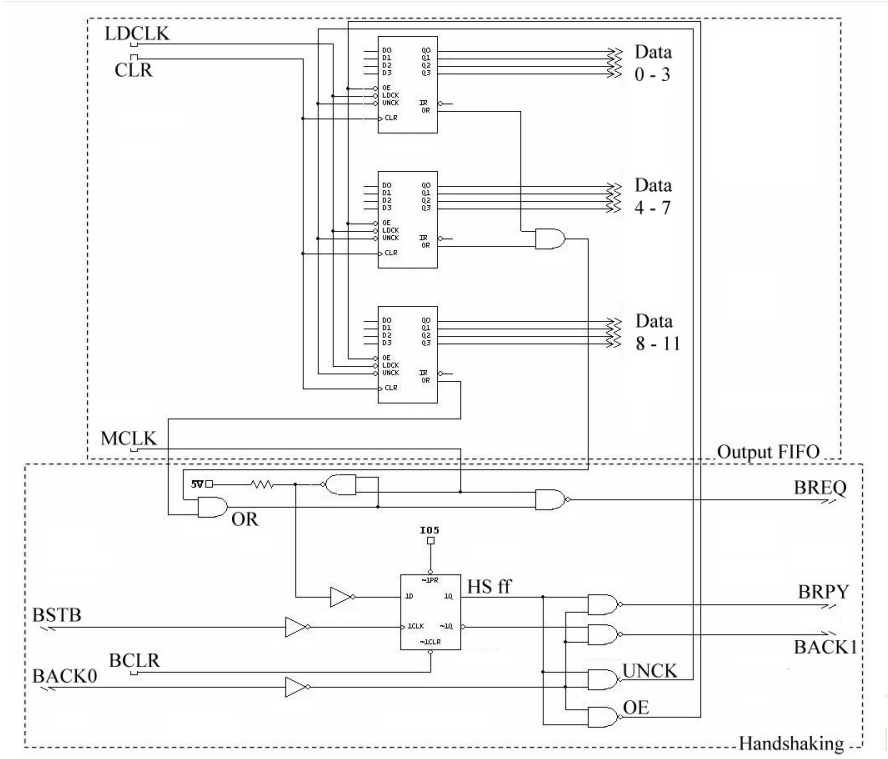

Figure 4.5: Coincidence processor output FIFO and control circuit diagram

4.1.1 Handshaking

The output control for each coincidence processor FIFO is handled by a handshake with a registered PAL (Programmable Logic Array) in the original RTS system. After a successful handshake between the RTS and a single coincidence processor within the IPCP cabinet, a coincidence data word is presented in parallel on the data bus. The data can then be captured and the handshake reset allowing further events.

There are four handshake control lines: BSTB and BACK are provided by the RTS system, BREQ and BRPY are generated by the coincidence processor. The signals are transmitted by 3-state devices as they share a number of control lines, the signals are therefore active low. The traditional handshake sequence is as follows:

- A coincidence board has data ready (output ready OR is set on each IC making up the FIFO) and requests service by holding BREQ (request) low. This is done with a 3-state device allowing communication with all coincidence processors.

- The RTS PAL sees the request and transmits a responding high to low pulse on the BSTB line.

- The BSTB pulse is seen simultaneously by all the devices on the bus, this pulse is used to trigger a flip-flop. The input to the flip flop contains a flag which is set if the FIFO has data ready to be read, and is not set if there is no data waiting. The output from the flip flop drives a decision circuit which waits for a further handshake signal to choose the bus controller.

- The RTS then holds the BACK signal at logic level low. This signal is passed through the coincidence processors along the serial control line. Depending on whether the coincidence processor has data (and therefore the result of the decision circuit) this pulse is either passed onto the next device along the chain (no data), or is used to complete the handshake (data). If the device does have data the BACK signal is prevented from transmission further down the line.

- The device that originally requested service (as it has data) sees this BACK signal. This is sent through the decision circuit (triggered by the data ready and BSTB signals) with the result of clocking the UNCLK pin on the FIFO ICs; this presents the data onto the data bus. At the same time the UNCLK pulse is sent along the handshake lines as BRPY (reply).

- The RTS sees the reply pulse on BRPY, delays a short time to allow for propagation delay in the system electronics and then latches the transmitted data in a local FIFO. The control signal BACK is then released by the RTS.

- The released BACK signal is seen by the device controlling the data bus, this disables the output enable (OE) flag on the FIFOs and returns them to the high impedance 3rd logic state.

- The system is then reset and ready to initiate another handshake sequence. The coincidence board that responded is forced to delay a short time before it is allowed to request service again, allowing the other devices higher priority for data bus control.

The BREQ pulse is placed on the handshake lines by all devices in parallel, therefore any device can request service at any time, however this causes some problems when two devices make simultaneous requests as only one pulse will be seen, the BREQ pulse is a single pulse transmitted when the device has valid data ready for readout.

The BSTB pulse is sent in parallel to all devices on the bus, this simultaneously clocks devices which have data enabling them to respond on the next clock cycle.

The BACK signal is seen sequentially along the serial link between each of the devices and enables the device to assume control of the data bus, the first device which has data along this line gains bus control, and stops the serial transmission of the BACK pulse.

The BRPY pulse is generated following the device gaining control of the bus, this is seen in parallel from each device on the bus as only the bus master can assert this line.

As the BACK pulse is held low by the RTS until the BRPY response is seen and data captured the bus remains in a fixed output state until the released BACK signal is seen. Due to the propagation delay through the serial link each device captures the data bus for different periods of time (during which other data cannot be captured), with the longest being the final device on the serial link.

Figure 4.6 shows a block diagram for the handshake process and illustrates the manner in which each device receives the handshake control lines.

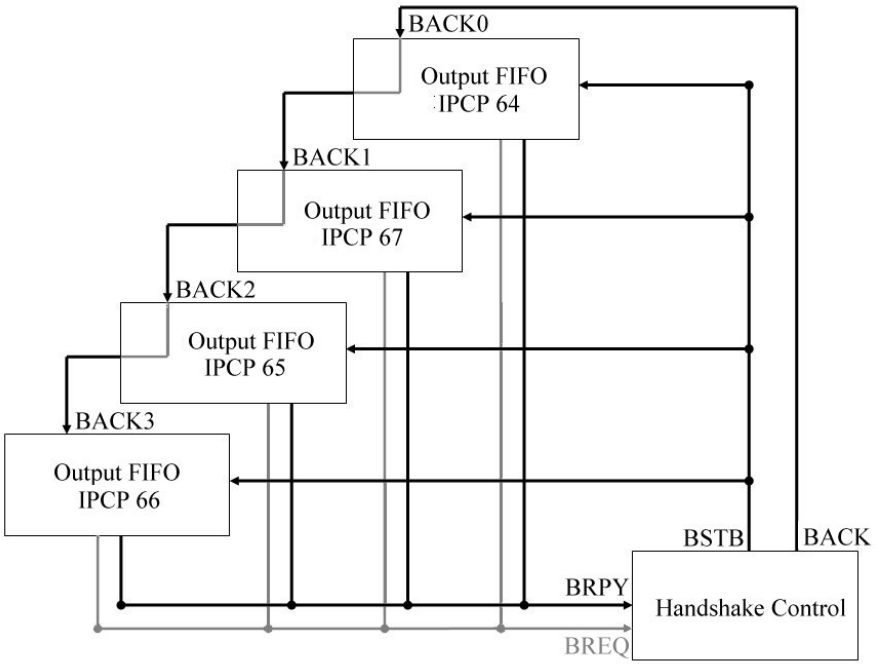

Figure 4.6: Block diagram for handshake process

In order to replace the RTS unit with an improved data acquisition system the handshaking requirements have to be met by the new system in order to have the coincidence processors placing data onto the data bus.

Initially a circuit which performed an identical function to the PAL in the RTS system was designed and built. It was shown to operate in the same manner as the RTS handshake and confirmed that it would be possible to interface with the coincidence processors without the need for the RTS system. The device itself operated relatively slowly due to propagation delay in the electronics and in the serial link (The RTS suffers from similar problems).

This device (and the RTS) had problems operating at high count rates because whilst a handshake is in progress other request pulses are ignored. If an event is ignored it is not lost, the FIFOs on each coincidence processor are 16 words deep, therefore the data are stored until a request can be granted. The request pulse is not given a second time however and this occasionally locks the handshake system up if no other board has a request. The RTS handshake has a similar problem, to this end the designers have placed a 50 ms watchdog timer on the handshake lines to reset the process if this occurs.

To operate the data transfer at a higher rate than possible with the RTS handshake unit the complexity of the handshake process can be reduced. The handshake system exists in order that the individual devices on the data bus can all obtain control of the data bus at some point during the operating cycle. As the RTS system cannot process or store events at a higher rate than about 1 MHz the coincidence processors initiate the data transfer process by effectively enquiring if the RTS is ready to accept data by the request (BREQ) pulse.

The new system is designed to have a higher processing speed than the coincidence processors, it can then be used to initiate the handshake sequence itself, rather than waiting for the request pulse from the coincidence processors. This can be thought of in terms of the relative speed of both devices. The digital I/O card which replaces the RTS has a maximum operating speed of 80 MHz, thus the effective dead time in the data capture process is negligible when operating at the speed of the coincidence processor (max 16 MHz combined from all devices). The data capture device can then be thought of as constantly available to accept data, and the BREQ pulse is then a formality which effectively limits the handshake speed and can be disregarded.

The new handshaking philosophy then sends a request to the coincidence processors to determine if any have data (this is effectively the BSTB pulse), if no device has data ready there is no change in the logic state of the decision circuit in the handshake controllers on the coincidence processors. The BACK pulse is then sent along the serial line, but as no device has data this is not responded to.

If a device does have data, the BSTB request pulse clocks the device flip flop and changes the logic state of the decision circuit. The BACK pulse is sent down the serial line, reaches the device with data and generates a response which is used to initiate data capture.

The new handshake sequence is as follows:

- BSTB pulse sent to IPCP to inquire if any device has data
- BSTB received and if any device has data the onboard flip-flop is set
- BACK signal is sent low and held
- BACK seen sequentially along serial bus. If a device has data the FIFO output is enabled and the BACK pulse aborted from further transmission. BRPLY response pulse transmitted back to handshake driver.
- BRPLY seen and a short delay is used before data is latched into a local FIFO on the data acquisition PC card.
- BACK pulse released allowing for the next event.

As the handshake is controlled by the data acquisition system, rather then the IPCP as before, the desired speed can be set and devices which have data but have lost the BREQ pulse can still be accessed, rather then having to wait for the handshake reset by the watchdog timer.

A simple circuit was constructed to provide these control signals, this uses a crystal oscillator with a fixed resonant frequency to generate a timing clock pulse. By careful considerations of the propagation delay in each of the circuit elements a square wave pulse is used to generate the required signals. As this is periodic, the entire handshake must take place during a single cycle of this clock, the following cycle then repeats the process.

The BSTB pulse clocks the decision circuit flip-flop (74ALS74A) following reception in a Schmitt invertor (74LS14). The invertor provides two functions, firstly it discriminates against electrical noise picked up in the transmission lines, then inverts the signal from active low (used for bus transmission) to active high. The flip-flop is clocked by a low-to-high transition on its clock signal and is therefore triggered by this rising pulse edge.

Following the propagation delay in the flip-flop (74ALS74A: 3 ns minimum, 14 ns maximum) from receiving a clock pulse to producing data at the Q output, the device is then ready to switch states upon the reception of the BACK pulse. However as the signal from the flip-flop and the BACK pulse are ANDed together, the BACK pulse is allowed to be received first as the state will only change when both signals are seen at the input of the AND gate.

The BACK pulse has to be held low for long enough that the signal propagates through the handshake circuit and triggers the FIFO (74ALS232) output enable, this requires a low-to-high transition with a typical delay of 5 ns (min 1 ns, max 14 ns) from trigger to the data being presented at the output. The output is disabled by a high-to-low transition on the same pin, with similar timing characteristics.

Whilst the BACK pulse is held low, and conditions are right that the device is the bus controller (i.e. it has data), the output enable is held high by the output from the decision circuit. The transition of BACK from high to low gives the OE state transition and the BRPY response pulse.

The data are presented on the output FIFO with a typical delay of 5 ns (min 1 ns, max 14 ns) from the transition of the BACK pulse within the device under consideration. The BACK signal is transmitted serially between devices and therefore this transition takes longer relative to the master clock for devices further down the bus. The FIFO output is enabled at the same time as transmission of the BRPY pulse, this is fixed to be a maximum of 44 ns per coincidence processor following the BACK transition and is set by the propagation delay in the handshake electronics.

Data capture is performed after a short delay following the reception of the BRPY pulse, conventionally this is designed to be for the worst case scenario, in this a trigger circuit has been designed to provide a 14 ns delay from reception to activation (see the following section).

The handshake driving circuit consists of a square wave which provides a falling edge on the BSTB line to trigger the flip-flop. The signal is also sent to the BACK line where the logic level is held low for half of the clock cycle. It was decided that there was no need to delay the BACK pulse following the BSTB pulse as the logic levels in the coincidence processor only change when both signals are received. The following high logic level of the remainder of the clock pulse then serves to reset the current handshake for the next event. This is illustrated in figure 4.7.

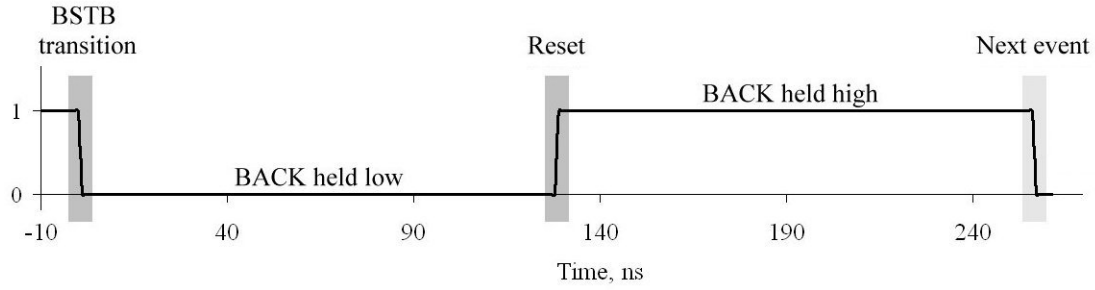

Figure 4.7: Handshake signals for one clock cycle

Figure 4.8 shows the handshake driver clock signal (4 MHz) with the corresponding responses in handshake flip-flop output, FIFO output enable and data ready for the first device along the serial control bus. This device is assumed to have data, therefore the output of the decision circuit (here labelled HS flip flop) is activated following the propagation delay in the handshake electronics. This is the delay from a single Schmitt invertor (74ALS14: 15 ns typical, 22 ns maximum), plus the delay for the flip flop to change state (74ALS74: 3 ns minimum, 14 ns maximum) following the transition on the BSTB line.

The FIFO output enable (FIFO OE) is then set and the data clocked from the device (UNCK) following the delay in the transmission of the BACK signal, ANDed with the output from the decision circuit. The BACK signal reaches the input to the AND gate before the decision circuit signal in this case, thus total delay is from the decision circuit output pulse through the AND gate (74S38: 4 ns typical, 22 ns maximum).

Data are presented at the output FIFO following the propagation delay from the FIFO receiving the OE and UNCK signals, to data presented at the output (74ALS232: 1 ns minimum, 5 ns typical, 14 ns maximum). These data are held on the data bus until the BACK signal is released and seen at the OE flag on the FIFO, a delay of a Schmitt invertor plus the AND gate plus the transition in the FIFO (total of 24 ns typical, 58 ns maximum). The diagram shows the worst case (i.e. maximum propagation delay) scenario.

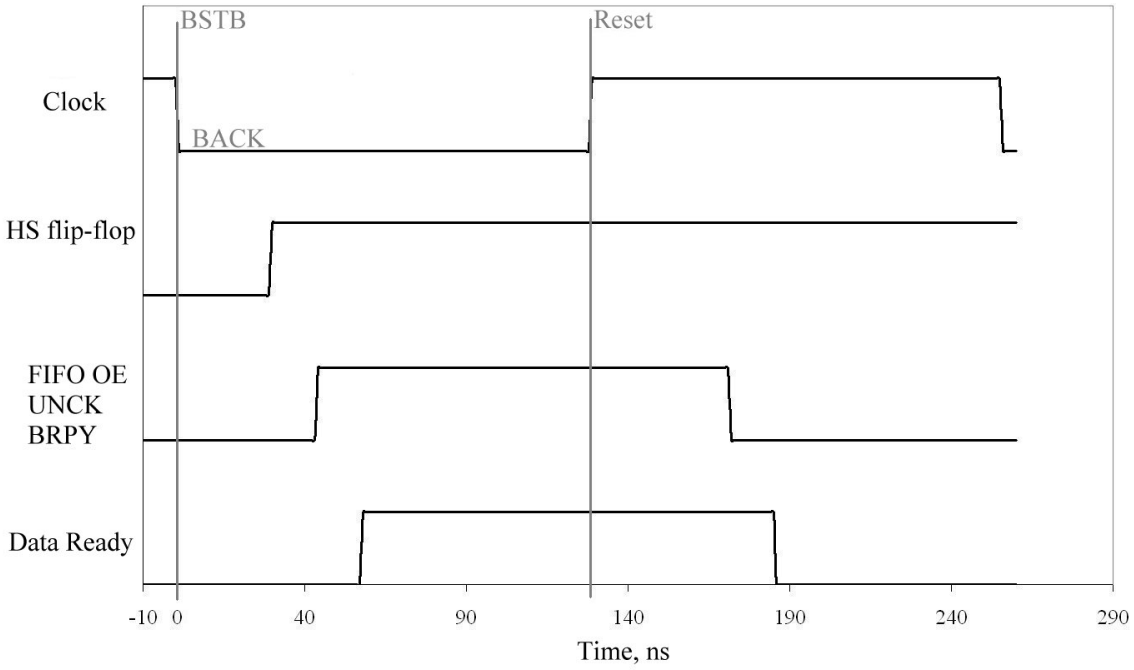

Figure 4.8: Control signals for 1st device on bus

The BACK pulse has to travel through each device in series to reach the final device on the bus and give it control of the data bus. For all four coincidence processors the maximum delay from the transmitted BACK signal, to a response (in the last device) is then the sum of the delays from the Schmitt invertors receiving the signal on each board (74ALS14: 15 ns typical, 22 ns maximum), and the delay in the output AND gates (74S38: 4 ns typical, 22 ns maximum). Each coincidence processor on the serial link then contributes a 44 ns (maximum), 19 ns (typical), delay in the BACK signal reaching the next device on the bus.

The total delay for a bus containing four coincidence processors is then 176 ns (maximum), 76 ns (typical) for the final device to receive and process the BACK signal. A further delay is given by the output FIFO switching (74ALS232: 1 ns minimum, 5 ns typical, 14 ns maximum) until the data is presented at the output, and a short time is required to capture the data (almost instantaneous). Thus the total time taken for data to be captured from the final coincidence processor is 190 ns (maximum) or 90 ns (typical).

It might be thought that signal transmission times would limit the time that the BACK signal has to be held low for a minimum of 190 ns to ensure that the last device has time to reply, this is the case for the RTS system, which waits for the response before freeing the data bus. This would limit the maximum clock frequency to be 2.6 MHz, which is a faster data transfer then the original RTS system allows.

However, once the signal has reached the first device there is a synchronous transfer delay which allows the BACK signal to be reset, whilst the pulse continues to propagate down the serial line. Provided the BSTB pulse is not triggered within this time the logic still operates correctly, this means the BACK signal can be released before the reply is seen (allowing a square wave pulse with equal mark/space ratio to be used to run the handshake).

The minimum time between handshake pulses can then be reduced to 190 ns as a worst case design, although the typical switching time is faster at 90 ns. The maximum frequency that can then be used to gather data is then around 5.2 MHz, a factor of 5 higher then the maximum design rate for the RTS system.

The clock frequency was chosen to be 4 MHz (with period 256 ns, giving control signals every 128 ns (activation followed by reset)) due to the availability of crystal oscillators with frequencies in the high end of the available range. Higher clock frequencies (up to 5.2 MHz) are possible, provided the correct oscillator can be chosen. Further issues with timing are discussed shortly.

The handshake sequence is illustrated in figure 4.9 where the handshake driver clock signal (4 MHz) and the corresponding responses in handshake flip-flop output, FIFO output enable and data ready are shown as before. In this figure though, the timing diagram is shown for the fourth device on the data bus. The data are presented at the output FIFO towards the end of the handshake clock cycle, but is still active for long enough to be captured by the data acquisition system.

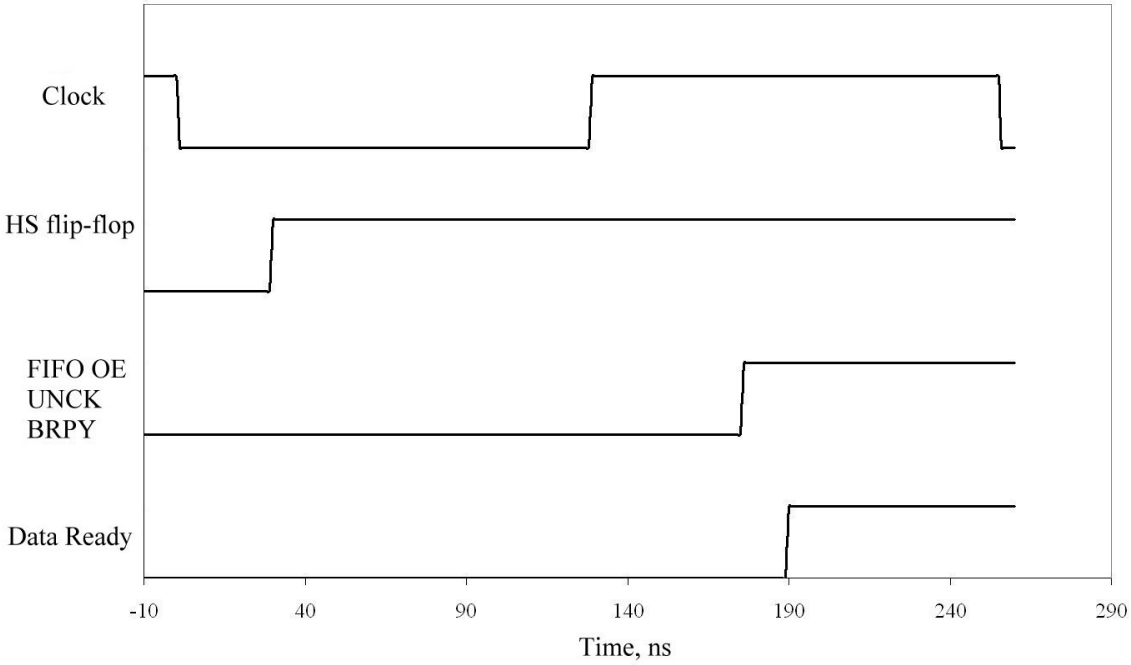

Figure 4.9: Control signals for the 4th device on the bus

The reset follows the BACK signal with the same time delay. Control of the data bus is relinquished by the fourth device after the 256 ns clock period. This is fast enough however, that control is released before the propagation delay of the next handshake signal is seen in the first coincidence processor.

The circuit diagram for the handshake device is shown in figure 4.10, it contains two separate parts, the upper is the crystal oscillator (4 MHz) driving a voltage controlled oscillator IC (74LS124) which provides a square wave output pulse with equal mark / space ratio at the crystal frequency. The device can provide stable pulses from 1 Hz up to 20 MHz with correct choice of resonant component. The square wave output from this device is tied directly to the transmission line for the BSTB and BACK signals as described above.

The lower part of the circuit diagram deals with the BRPY response pulse generated when data are presented on the data bus. Due to the propagation delay in the FIFO electronics (74ALS232: 1 ns minimum, 5 ns typical, 14 ns maximum) a 14 ns delay is required between receiving the BRPY pulse and the capture of data from the bus. As the signal is active low by using a Schmitt invertor (74ALS14: 15 ns typical, 22 ns maximum) the pulse can be sufficiently delayed, and converted into an active high pulse. The resulting signal is then used to trigger the PC data capture.

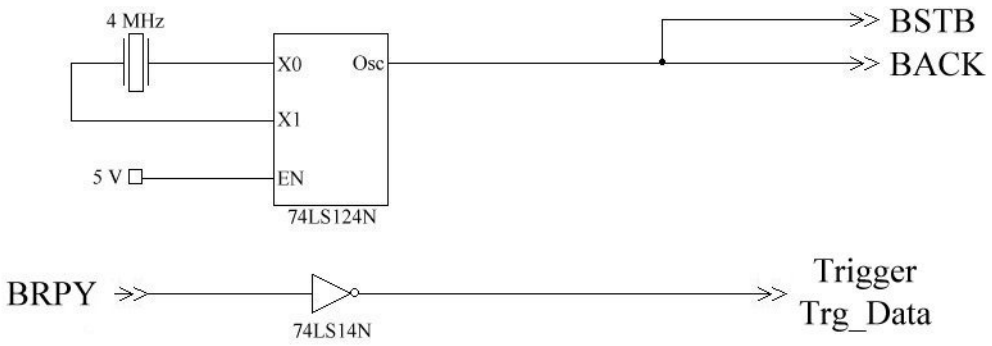

Figure 4.10: Circuit diagram for the handshake driver

Finally a word about timing. The clock has a minimum period (calculated above to be 190 ns) for successful handshaking with 4 coincidence processors. This time can be decreased with fewer devices on the data bus or by using a handshake pulse with an uneven mark-space ratio allowing the BACK pulse to be held low for the majority of a clock cycle. A 256 ns clock (4 MHz) is then sufficient to deal with the handshake requirements, this would be the maximum data rate from a single coincidence processor operating at full speed. For higher speeds (up to 16 MHz as from all four devices) the serial nature of the handshaking is then the limiting factor, although clever design can reduce these effects.

Unfortunately this is not the primary limiting factor for the maximum speed. The handshake signals and the data signals are transmitted in TTL format over a metre length of cable to the data acquisition system. This cable has a certain capacitance associated with it and at TTL voltage levels this capacitance limits the useful operating frequency. The handshake and data signals are received correctly at 4 MHz, however increasing this frequency to 8 MHz results in a capacitative deformation of the signal and transmission errors degrade both the handshake and data signals beyond recognition.

As the maximum operating speed of the RTS is only 1 MHz, the new system offers a factor of four increase in performance for the camera in this design. In the future if higher data rates are needed, active line drivers (with transmission in high speed ECL) and shorter cable lengths are required to overcome these problems.

The IPCP output FIFO is 16 bits deep allowing for deadtime in the original RTS data acquisition system. This means that each IPCP can continue to produce up to 16 coincidence words at 4 MHz; whilst the RTS system is busy processing the previous data the following data is buffered. The RTS can then catch up with the data stream by accessing and emptying this buffer. This cannot continue indefinitely however, if the data rate is high enough to keep the buffer constantly filled (as each RTS access removes the current event), events are dropped if the buffer reaches its maximum capacity (16 words).

4.1.2 Timing Data Insertion

Once the data can be effectively captured by the PC data acquisition card following the handshake process, event time information needs to be included. Conventional PET imaging does not require high resolution temporal information to be stored, however this data is essential for PEPT experiments as locations are made on a frequent basis.

Time information for each event is determined by the time of arrival of the event in the processing unit. Conventionally the RTS processor is used to insert a timing tag word into the data stream once per millisecond as the data are streamed to disk. Unfortunately the PC based system does not give access to the memory buffers used in data capture and this approach, whilst possible, is highly complex.

Instead of a software approach, a dedicated board has been constructed which adds a 4-bit time stamp to individual events with a temporal resolution of 0.5 ms. Every 16 ms a time tag word is inserted into the data stream to deal with events that occur with period of less than 16 ms and for error checking. The time resolution is sufficiently high for PEPT as the location time is calculated from the average time of the trajectories used in the final solution. This can cover a number of milliseconds, for example an event rate of 1 MHz using 1000 events per location (N of 1000) provides locations once per millisecond with the time of the location given as 0.5 ms offset from the absolute time.

With the increased data rate expected using the new system it might be beneficial to increase the frequency of time data if many locations are made within the timing period. This is easily done by changing the resonant frequency of the circuit, to this end two variable resistors are included providing coarse and fine frequency control.

The block diagram for this circuit is given in figure 4.11 and the circuit diagram in figure 4.12. A 555 timer is setup to give a square wave of frequency 2 kHz, this provides the clocking signal for a binary counter. Every 0.5 ms the counter is incremented, the 4-bit data word is sent from the counter to the high nibble of the 32-bit data word describing the event. The format of the word is described in detail in section 4.2. Any events which arrive during the 16 ms described by the counter then have a time stamp included in their description.

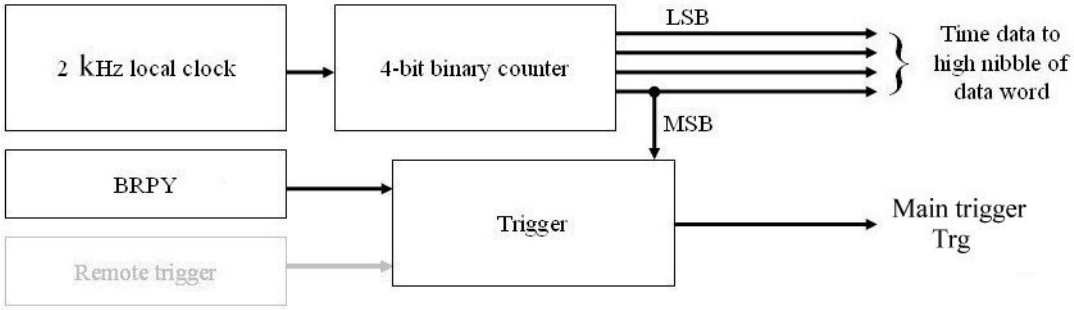

Figure 4.11: Block diagram for time generation

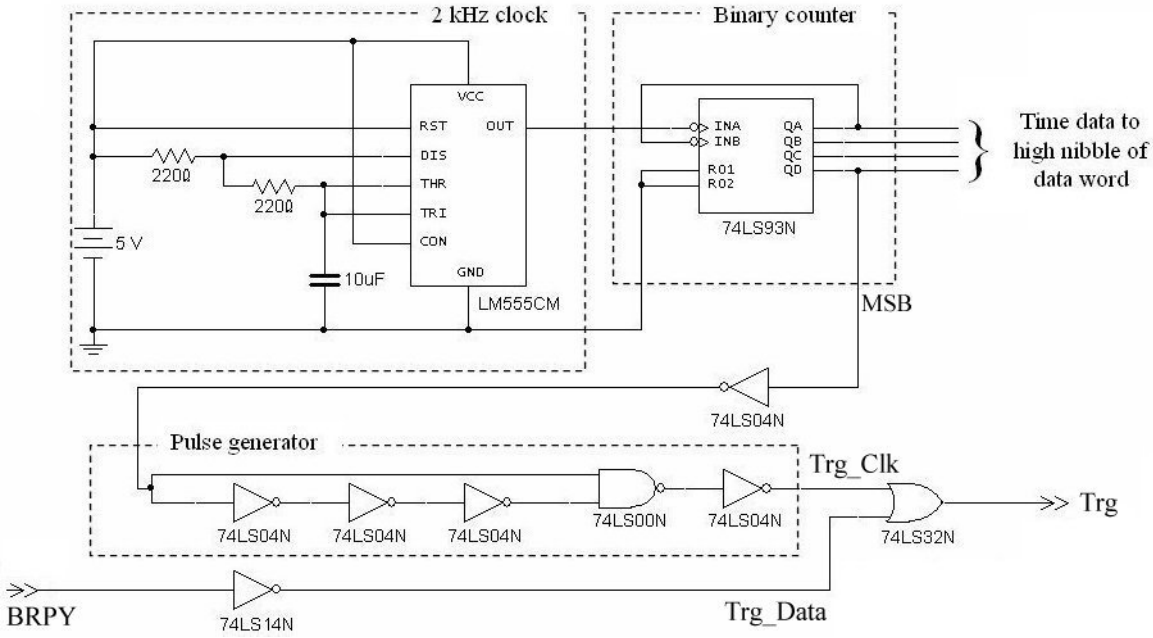

Figure 4.12: Circuit diagram for time generation

If the event rate is low, the counter still needs to provide accurate timing data. To this end the most significant bit (MSB) from the counter is used as a trigger to insert a data word every 16 ms, describing the time data when the counter is reset. This is much the same as extending the counter to include a carry operation. On the transition of the counter MSB from high to low (corresponding to binary 1111 to binary 0000) a pulse is generated to trigger the data capture and insert the timing tag word.

This is done using a falling edge pulse generator. The falling pulse (from the MSB transition high to low) is separated into 2 signals, one signal is delayed by a set of 3 inversions (74LS04: 9 ns minimum, 15 ns maximum) totalling 27 - 45 ns. This signal is NANDed with the original falling pulse which generates a logic spike of 27 - 45 ns in length.

This signal is then ORed with the BRPY response pulse set when the output FIFO has data and finally sent to the PC acquisition system trigger. The OR (74LS32: 3 ns minimum, 15 ns maximum) gate includes a further propagation delay to the trigger pulse, this is short enough that the trigger occurs during the time data is still available on the IPCP FIFO output.

The pulse edge generator is used to minimise the deadtime (0.25 ms) during which the MSB from the timing signal would hold the trigger high; this would result in the data acquisition system being unable to accept events during this period. That said, there is still deadtime during this trigger pulse, any data (at high rates) arriving during this time will not be captured, however the handshake will allow this data to be collected on the next cycle as the data word is stored in the 16-bit deep output FIFO on the coincidence processor.

It is not a concern if the timing trigger is missed due to acquisition system deadtime during high speed data capture. If this occurs the captured data has the new time stamp included which will be processed in the algorithm in the same manner. This time stamp lasts for 0.5 ms (at 2 kHz clock rate) thus any data arriving during this time will include the time data (as this would only be lost at high data rates).

4.1.3 Remote Trigger

To increase the functionality of the data acquisition system the facility to include trigger signals from up to 2 extra devices has been included. In PEPT it is sometimes useful to measure other properties of a system simultaneously with the tracer motion, for example in a stirred system it is often useful to correlate tracer motion with the position of the stirring impeller. A tachometer can be used to measure the rotation speed of the impeller shaft.

In a clinical application this extra data could take the form of a cardiogram or ECG to divide the data stream into cardiac gated cycles.

The trigger arrives from the measurement instrument (e.g. tachometer) as a TTL voltage pulse with a width of a few ms. This pulse is passed through a rising edge pulse generator to provide a voltage spike that can be used for triggering. The spike is ORed with the other trigger pulses that activate the data capture and therefore provides its own independent trigger pulse.

The TTL signal from the device is also placed onto the data bus on one of the spare inputs (2 bits are spare once coincidence data and time data is included). This flags the trigger pulse as being generated by the remote trigger, rather then the coincidence data stream or time trigger. Analysis of the data file can then be done to resolve the individual event. As the time data are included in the data word the remote trigger data are provided with time resolution equal to that of the timing signal.

Figure 4.13 shows the circuit diagram for the remote trigger circuit.

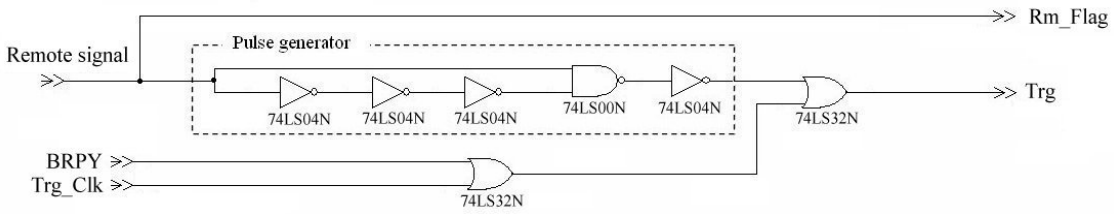

Figure 4.13: Circuit diagram for remote trigger

It should be pointed out that this circuit only deals with a single bit of data from the remote device. This is acceptable when applications are involved with simple counting, but devices presenting data greater then the available 2 bits input cannot be used. In the future a multiplexed system could be used to insert up to 32-bits of data from a remote device provided a trigger signal can be generated.

Figure 4.14 shows data taken with the remote trigger for a tachometer measuring the rotation period for a variable speed electric drill. The device inserts a remote trigger word each time a label attached to the drill shaft was at a certain position during the rotation.

Three different drill speeds are shown, these are labelled high, medium and low speed and were measured to be 2800 rpm, 1747 rpm and 746 rpm respectively. The manufacturers guidelines state that when new the drill has a variable speed ranging from around 1000 rpm up to 3000 rpm.

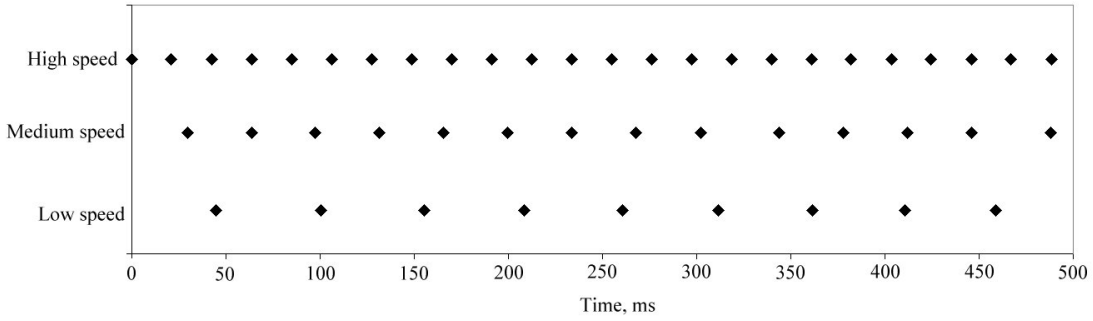

Figure 4.14: Results for tachometer attached to an electric drill

4.2 Computer Hardware

The coincidence data and BRPY trigger, time information and the remote trigger information is sent, along with the final trigger pulse, to a modern PC based data acquisition card manufactured by Adlink Technologies inc. This card is an ultra high speed digital I/O card consisting of 32 programmable digital I/O channels operating at differential TTL voltage levels, with simple trigger / handshake capabilities. The card features a further 8 I/O channels in an auxiliary buffer and active cable termination for high speed and long distance data transfer.

For data acquisition the card is programmed to have all 32 channels operating as digital inputs, the maximum data acquisition rate from the input channel to the host memory is 80 MBytes per second; this is given by the 32-bit bus width multiplied by a 20 MHz trigger clock. 160 MBytes per second can be reached in short bursts when the acquired data length is less then the size of the onboard receiver FIFO (16 k $\times$ 32-bit samples).

The acquisition is triggered by a handshake signal generated by the support circuits, this trigger provides a low to high transition to initiate data capture.

Data transfer from the onboard FIFO is performed by bus mastering direct memory access (DMA) over the host computers PCI bus. The PCI bus is 32-bits wide and operates at a maximum throughput of 132 MBytes per second. The data acquisition card takes control of the PCI bus as the bus master, transfers FIFO data at the PCI burst speed, and then releases the bus control. This DMA is processor free, therefore the processor is available to perform real time analysis and display data etc.

The data are transferred directly from the onboard FIFO into system host memory in a set of circular buffers. The data does not have to be stored as a single continuous block of physical memory as the card uses scatter / gather technology. Each buffer contains an address and length register and provides a linked list allowing data to be stored in addresses scattered through the DMA address space.

The circular buffer consists 8 buffers of 32 kbits (10k array of 32-bits) wide. Data are transferred asynchronously from system memory to storage on a high speed large capacity hard drive as each buffer is filled.

This is a potential bottleneck as disk write speed is again the slowest data transfer rate in the system. Buffers which are dropped due to slow disk access are flagged. In tests to date this has not limited the data transfer performance, however if data rates increase above the maximum already achieved this could be a potential problem.

The onboard system memory for the PC is 2 Gbytes, this could be used as a temporary storage buffer and data written to disk after the acquisition. This would limit the maximum run time at high speed as system memory would be quickly filled.

Alternatively the system memory could be used to buffer data during high event rates, and transfer data to disk during periods with lower event rates (i.e. when the computer resources are free to handle the data transfer).

Figure 4.15 shows the data word format for the system when connected to a 931 coincidence processor operating with 951 detector buckets (section 3.5).

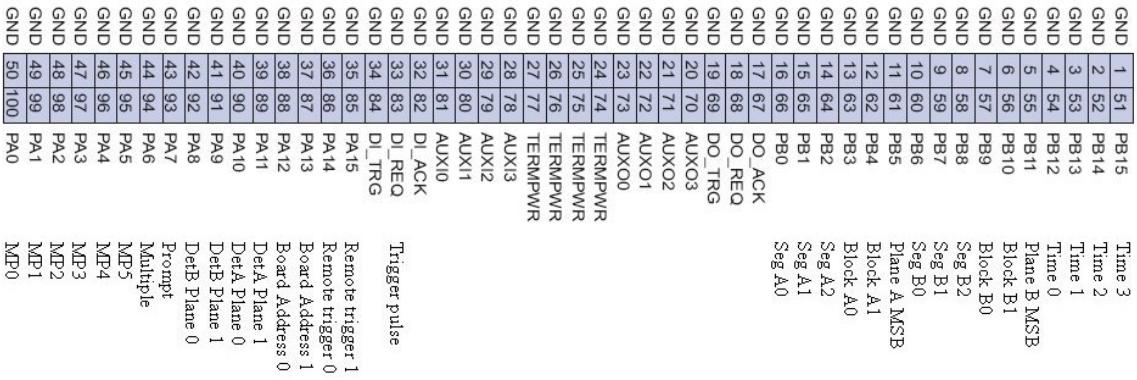

Figure 4.15: Data word format for the new system

The format of the word depends only upon the connections between the coincidence output FIFO and the PC card input FIFO, and therefore is relatively arbitrary. For other coincidence processors the format will remain unchanged although if 931 detector buckets are used the MSB describing the extra 8 axial crystals will not be active.

Currently the auxiliary I/O port has not been used, this remains for a future extension if it is necessary to input or output extra data. The card offers handshake triggers for both input and output, as the handshake is dealt with in dedicated hardware the trigger pulse is provided to the appropriate handshake line. The remainder are left unconnected.

4.3 Control

Conventionally control of the entire system is provided by a RS-232 serial link connecting the Sparc CPU to the IPCP and detector buckets via a device called the current loop convertor. This transmits serial data along the clock bus lines, all devices listen to the bus so a serial address is sent to activate one device and allow reply data.

Commands are sent from the SUN operator's console via an ethernet link to the Sparc CPU, these commands initiate a remote procedure call running on the Sparc which sends ASCII characters to the addressed device. The characters are interpreted by the destination microprocessor into control commands. These could be in the form of initialising the master clock signal, inquiring the coincidence count rate within a coincidence processor or calibration of a detector block.

In the new system a new control link has had to be designed to replace the complex RTS / Sparc / SUN control system. The Universal Serial Bus (USB) for the data acquisition PC is used to generate the appropriate serial commands via software running on the PC. Data are transferred from the USB port and patched into the current loop convertor using the standard RS-232 convention. High level software similar to the original operators console and written in C is used to operate the serial link and send the appropriate control signals.

Data is received from the serial bus in the same manner. Thus detector setup, calibration and single photon information can be obtained from the detectors, coincidence data and control commands can be accessed in the coincidence processors and the clock system can be initiated. The motion control for the original scanner gantry is also handled in the same manner, however the software support has not been included as for PEPT the patient couch etc. has been removed. The software is described in detail in chapter 5.

Figure 4.16 shows the cable connections for the serial link from the USB port to the current loop convertor, the serial data format is setup to have a baud rate of 9600 transitions per second with 8 data bits, 1 stop bit and with no parity checking.

The hardware control lines are not used in communication with the current loop convertor so hardware flow control needs to be disabled. In the cable built for this device the hardware checking is looped back to itself in the same manner as a null modem cable allowing the system to communicate without handshaking flow control.

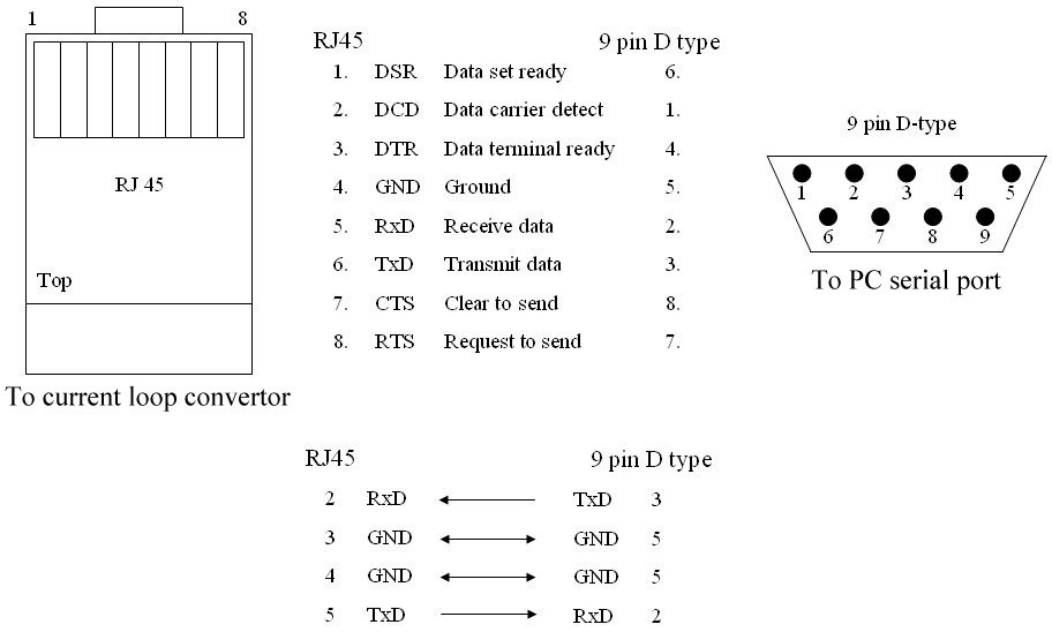

Figure 4.16: RS-232 cable connections

4.4 Complete System

As the acquired data and control all operate from a single PC it is far simpler to acquire, analyse and store the data. There is no need to transfer acquired data to a separate system for analysis (as before) and analysis code and a display of data can be provided for the PC.

Figure 4.17 shows a block diagram for the acquisition system electronics. The coincidence processor is driven by the handshake driver and event data words are placed on the output FIFO following a successful handshake and trigger. Time information and remote trigger information is inserted into the FIFO by triggering the data capture as events occur.

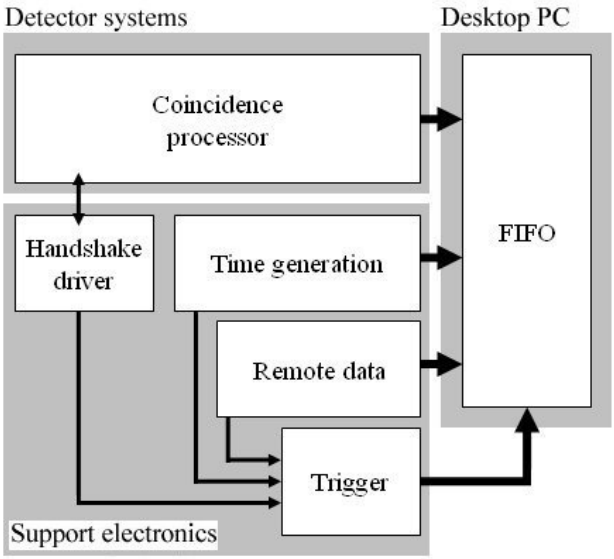

Figure 4.17: Acquisition system electronics

The new data acquisition hardware has been designed to interface with the ECAT 931 coincidence processor made by Siemens / CTI (section 3.3). Recently it has been used to capture and process data from ECAT 951 series detectors and is able to process data from similar systems with little or no modification, provided the correct interface cables can be constructed.

Control of the coincidence processor and individual detectors is made possible by high level software running on the same computer. In this a serial RS-232 link is used to interface with the systems, the software converts user level commands into machine code used to communicate with the camera. In this manner detector calibration is performed and system checks and event data can be accessed.

The entire system is illustrated in the block diagram given in figure 4.18. Here the detectors and coincidence processor are controlled by serial commands issued from the computer's USB port.

Data from the individual detector modules undergoes coincidence processing and the resultant coincidence data word is captured following successful handshaking.

Data are transferred as a block from the local FIFO to a circular buffer sitting in system memory via Direct Memory Access (DMA) along the PCI and system busses. The PC resources can be used for processing the data, performing data analysis, and displaying real-time LOR data.

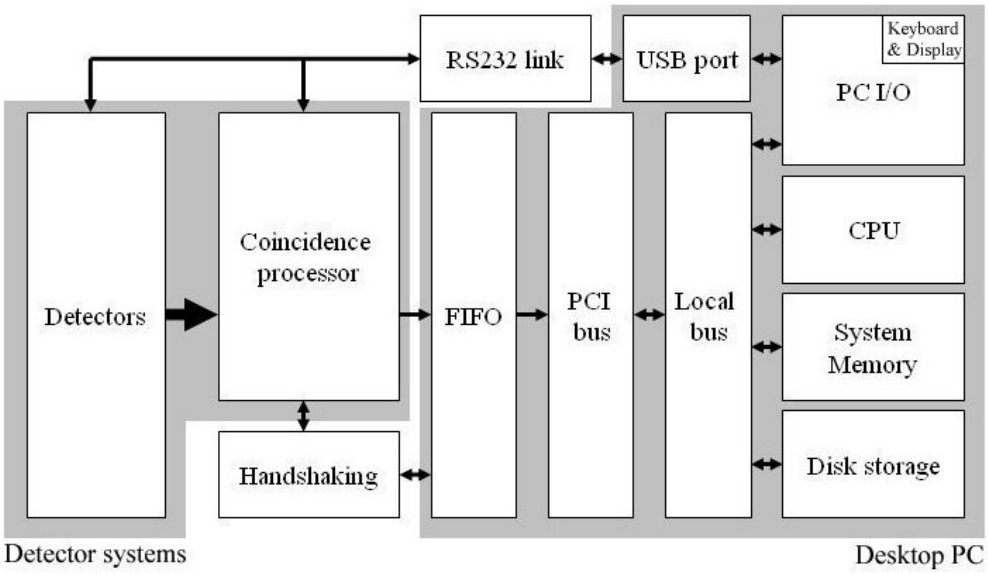

Figure 4.18: System block diagram

Raw event data can be saved to disk or removable storage media (CD/DVD) in binary format. These raw data files can then be processed using the PEPT algorithm to determine the particle locations. This allows different processing techniques to be used to optimise the location data.

The output from the PEPT algorithm is an ASCII file containing the average time of the location, three Cartesian coordinates describing the particle position, and a value which determines the quality of the location by calculating the perpendicular standard deviations of the trajectories used to locate each event (d).

4.5 Results

The new data acquisition system has been used with the small modular camera setup (section 7.5) consisting of 8 detector buckets. These have all been configured to operate in coincidence and the data rate is consequentially high.

Leadbeater and Parker [36] report upon trial PEPT runs on a plastic extrusion process using this system. Sustained data acquisition rates of approximately 120 kHz are demonstrated with short bursts of above 1.7 MHz seen during periods with the tracer particle located in areas of high sensitivity.

More recently additional trials have been carried out, these have used highly active tracers fixed within the field of view of the small ring. Continuous data rates of up to 3.8 MHz have been demonstrated with this system: approaching the maximum data rate that is attainable with the 4 MHz crystal oscillator driving the handshake control lines.

This corresponds to writing up to 16 MBytes per second to disk; if this rate is constant 3 GBytes will be written to disk in 188 s meaning that different storage methods have to be developed for long experimental runs (e.g. live data compression).

Furthermore the raw data file will have to be broken into pieces during acquisition so that it does not exceed the maximum file size and that it remains portable (4.7 GBytes is the maximum file size that can be written to DVD).

Full analysis of a data file this size using the PEPT algorithm will take much longer time than the acquisition (of the order of 10s of minutes) even for a modern PC. There is a possibility of using parallel processing techniques to improve the analysis speed in the future.

The increased data rate has allowed more trajectories (increased value of N) to be used for each PEPT location without the corresponding loss of precision due to tracer motion. In these experiments accurate locations were made every 11.3 ms on average using 2000 events per location (N).

Shown in table 4.1 are the PEPT location data for three different event rates, a low value of 25 kHz corresponds to the older systems in the positron imaging centre, whilst the higher rates of 100 kHz and 1 MHz correspond to operation with the new data acquisition system. In all of these cases the tracer particle has to be sufficiently active and remain in the field of view to generate the required data rate.

Results were analysed with an f value of 10% for each data set. The particle displacement is measured as the distance travelled by a tracer particle moving at 2 ms^{-1} during the time taken for each location. It can be seen from this data that the increased data rate results in higher location rates and increased measurement precision.

Event rate	N	Location precision	Location period	Particle displacement ($2\ \mathrm{m\ s^{-1}}$)
kHz		mm	ms	mm
25	250	1.00	10	20
100	250	1.00	2.5	5
1000	250	1.00	0.25	0.5
25	1000	0.50	40	80
100	1000	0.50	10	20
1000	1000	0.50	1	2

Table 4.1: Location data for different count rates

For this application a laser tachometer was used to correlate tracer motion with that of the process system, the device inserts data using the remote trigger. The new system has been seen to operate with performance exceeding the maximum performance of the initial system and offers much in the way of flexibility for acquiring and processing / displaying data.

Further details on this experiment are given in section 7.5. The high speed data transfer has allowed accurate measurements to be performed frequently, this device has opened up a fantastic opportunity to acquire high quality PEPT data.

Chapter 5

Software & Visualisation

A discussion of the various software applications that have been developed is given in this chapter. These applications are used for interfacing with the camera hardware and data acquisition systems and offer data handling utilities, data analysis and visualisation techniques. They can be broadly divided into two sections: software used to interface with the camera hardware and software used to analyse the retrieved data.

The camera hardware communicates with the user via an RS-232 serial link where simple ASCII character data are transmitted and received. There are a number of different methods used to communicate these commands and data; in the original system commands are transmitted over an ethernet link to the host CPU contained in the RTS rack of components. Here commands are translated into RS-232 standard and transmitted to the gantry system via the current loop convertor. In the new systems the RS-232 link is driven directly via the serial port of the operator's computer.

Data describing coincident events is transmitted in parallel to the RTS system or the new data acquisition system where it is stored to disk and eventually processed. Software used to acquire the data to system memory and to disk is described in the following sections. A number of software utilities have been written to analyse this data stream both for examining the raw data and for performing PEPT locations.

The majority of the software is written in the programming language C and compiled using MINGW, DJGPP or Microsoft visual C (MSVC) compilers. These different compilers offer many unique functions, for instance many of the older functions used to drive the serial ports and VME bus drivers have required the use of legacy code that is not necessarily supported by newer compilers. Conversely the visual elements written to run on modern systems (Windows XP) have required up-to-date software libraries and functions.

5.1 Control Software

The entire camera system is controlled using relatively simple ASCII commands sent over the serial clock bus. In the original system commands are sent from the operator's console over an ethernet link to the SPARC CPU situated in the RTS VME rack. These commands initiate a remote procedure call (RPC) on the SPARC CPU. This allows executable code to be run on the SPARC via remote commands from the SUN machine connected over the ethernet link.

The RPC activates the serial port on the SPARC system and commands are sent to the current loop convertor where they are converted from RS-232 voltage levels into the high current drive necessary to accommodate the distance and number of devices on the serial bus. Signals are then transmitted to the clock board where bus connection is made.

The clock transmits the commands serially on two of the four twisted pair cables (one for the high byte and one for the low byte of the data word). The other two twisted pairs carry the master clock signal and a strobe signal (master clock frequency divided by 8) in order that the receiving device can parse the data word.

Each device on the serial bus is assigned a serial address by setting a DIP switch on the particular hardware board. Detector buckets are given the serial addresses from 0 to 31, the coincidence processors are assigned 64 to 67, and the master clock has the address 255. The motion control systems for moving the patient bed etc. are accessed in the same manner (with different addresses) but are not used in the Positron Imaging Centre.

The data word that is transmitted then contains the serial address for the device being communicated with, followed by command characters. All the devices on the serial bus listen to the data stream, but only the addressed device responds to the commands and takes control of the bus in order to send information back along the same path.

The diagram in figure 5.1 shows the conventional manner in which the serial data is transferred around the ECAT systems.

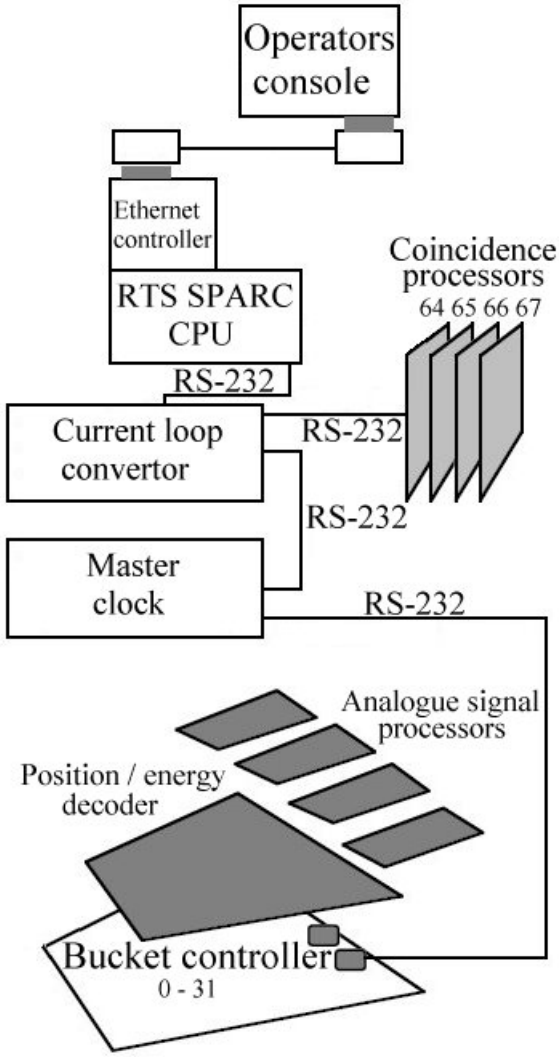

Figure 5.1: Conventional serial communication strategy

With the new system, where the RTS cabinet has been replaced with a single desktop PC, the serial communications giving control of the camera system have been replaced. Instead of the complex system of ethernet control, RPC and serial communications, the serial port of the new PC is connected directly to the current loop convertor using a standard RS-232 interface (described in section 4.3).

The new system architecture is shown in figure 5.2. The front end data remains identical to the original system, but the host device controlling the camera is replaced with the PC. Data from the coincidence processors is transmitted directly to the PC data acquisition card and then to system memory and disk.

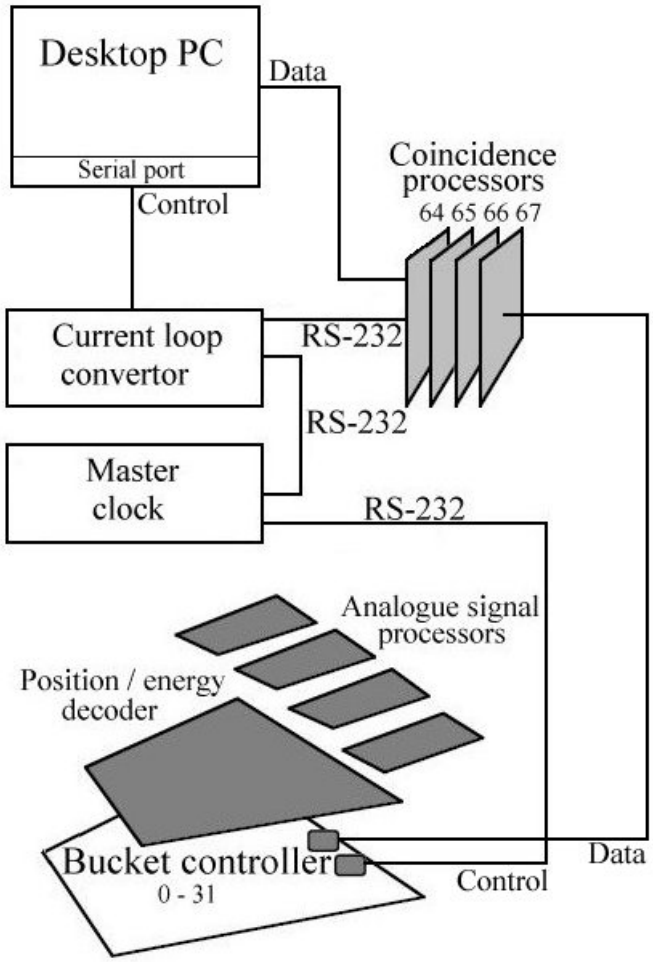

Figure 5.2: New serial communication strategy

Software running on the PC is then used to initialise the serial port and setup a communications link between the PC and the camera system. The serial port is setup to have a baud rate of 9600 transitions per second with 8 data bits, 1 stop bit and with no parity checking. The data that is transmitted describes characters in the extended ASCII character set. Commands are given in [19], pages 184 - 191.

Typical examples of command words are given in table 5.1 where # represents the serial address of the device queried.

Command characters	Function
255R	Start clock
#C	Report countrates in coincidence processor
#Jx	x = 1 disable, x = 0 enable, 50 ms FIFO reset
#A	Enable all coincidence plane combinations
#S	Select bucket
#Sx	Select # bucket, block x (0 - 3)
#H1	Select position histogram
#Hx	Start histogram for x seconds
#X6	Report position histogram

Table 5.1: Commonly used serial commands

Responses are generated by the addressed device and transmitted serially back to the current loop convertor where it is converted into RS-232 signals and sent to the PC (or SPARC CPU). Responses range from a simple response indicating correct operation or confirmation that the serial commands have been seen and interpreted properly (typical response is the characters N 0), to count rates or histogram data.

Figure 5.3 shows the command and a typical response for selecting bucket 0, block 0, and a query into the coincident count rate in the coincidence processor given the serial address 64. The response from the coincidence processor returns three numbers separated by commas. The first number is the prompt coincidence count rate, the second is the measured randoms rate and the third shows the combined multiples and scatters rates.

```
Command  > 0S0
Response > N 0

Command  > 64C
Response > 33179,9981,893
```

Figure 5.3: Commands and responses for serial communications

The character level commands can be sent with any standard communications software such as Microsoft Terminal or Putty, where the serial bus can be activated and used to send data. The process of acquiring a histogram involves sending many commands and a detector calibration involves many such histogram stages which would be tiresome to perform by hand. Software has therefore been written to interface with the serial port (with extensive diagnostic utilities) and send character commands following higher order commands from the user.

For example the commands that need to be sent to retrieve a detector block histogram are shown in table 5.2. With the high level user software this chain of events can be initiated by using the simple command **histo bucket, block**. More complex commands can be issued with simple language, for example the entire system can be calibrated with individual PMT gains and energy discriminator levels correctly set using the command **setupbucket all**.

Command characters	Comments
#S4	Select bucket
#I4	Disable all blocks in bucket
#Sx	Select block x within bucket
#H1	Enable histogram
#H0	Confirm histogram status is ready
#Ht	Begin histogram acquisition for t seconds
#Sx	Select block x within bucket
#H0	Check histogram status, repeat until histogram is ready
#X6	Read out finished histogram
#S4	Enable all blocks in bucket

Table 5.2: Commands to acquire histogram in bucket #

A flow chart for the program is shown in figure 5.4. The first function of the program initialises the serial port with the correct settings to interface with the camera systems. The program then accepts commands from the keyboard or from a predefined file (where setup data can be saved and reloaded).

The program (called **931Comm** or **951Comm** depending on the camera system) then has to parse the user command into a set of functions that operate the camera. Provided the command is recognised the program runs code that sets up the problem (i.e. the commands needed to initiate a histogram) then uses the function **comms(cmdbuf, response)** to place data on the serial port.

This function takes pointers to two character strings as arguments. The first argument (cmdbuf) points to the string giving the ASCII command (e.g. 0S0). The second argument points to a large character array (response) where the response from the camera (e.g. histogram data) is stored.

The received data are then parsed from the response buffer into either a text response displayed on screen (e.g. N 0), or into integers / floats describing histogram data etc. The parsing makes it possible to operate on the numbers received and calculate further information from the data (e.g. average count rate in a detector block).

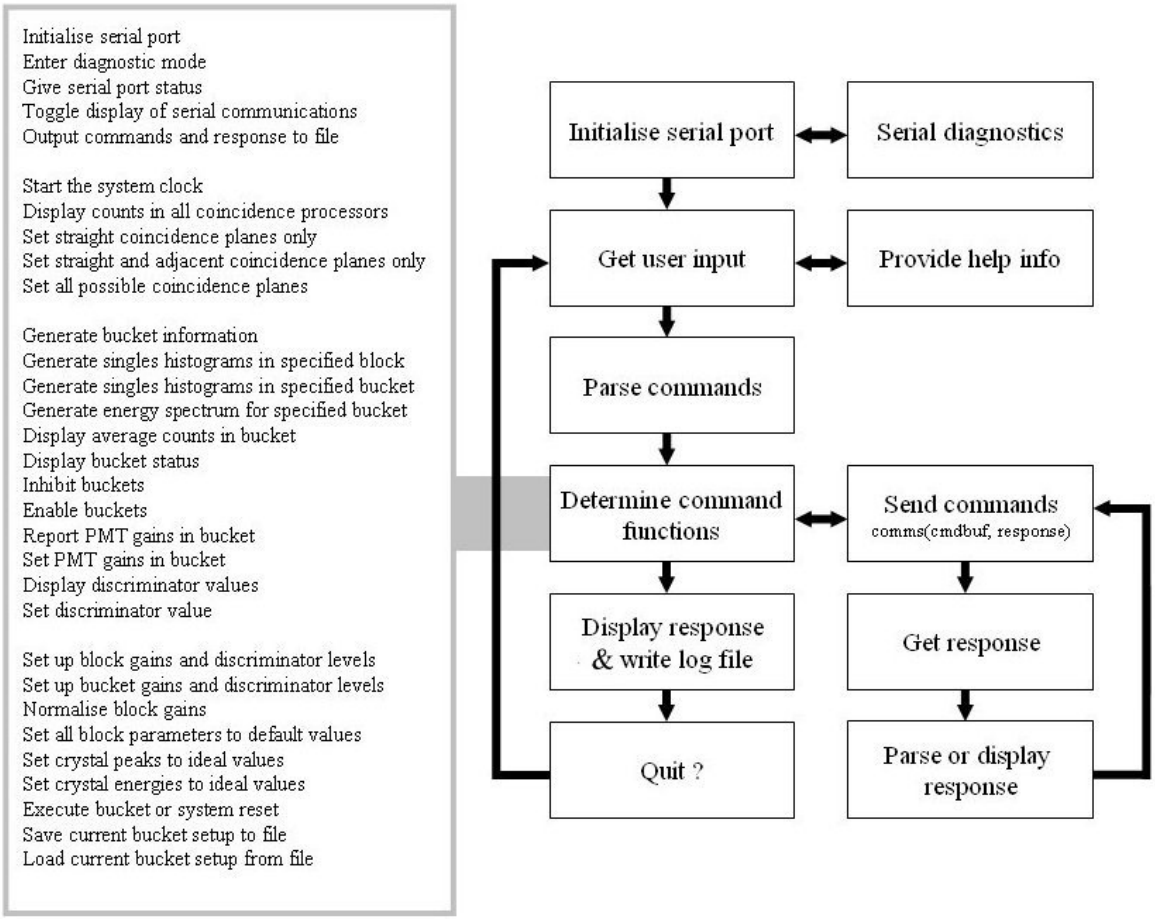

Figure 5.4: Flowchart for serial communications program

Figure 5.5 shows a typical position histogram for bucket 8, block 1. The 8 × 8 matrix of crystal elements in the detector block (951 series) is shown with a histogram of the number of events occurring in each crystal during a 10 second time period.

Histogram for bucket 8, block 1:

475	678	636	606	528	738	585	561
471	613	571	617	645	542	613	619
484	520	592	632	556	658	556	549
454	594	570	614	551	537	642	582
473	658	556	621	647	634	707	597
504	544	611	565	641	733	549	618
494	553	532	634	493	549	628	541
520	516	587	571	549	551	612	637

Mean = 580 (+/- 24) Stdev = 11.31 Center = (0.08, 0.03)

Figure 5.5: Histogram for detector block

This data was acquired using a broadcast source resulting in each crystal element seeing an approximately uniform radiation field. Each crystal element should then detect an equal number of gamma events (within the Poisson counting error).

By comparing the average and the standard deviation across the whole crystal it can be seen that the standard deviation is within the counting error, and that this detector block is functioning correctly.

Detector problems can easily be identified using the histogram function. Figure 5.6 shows the histogram for an 8 × 4 matrix of crystal elements (931 series). The top left corner of crystal elements have not had any events during the 10 s period that the histogram was acquired for.

This is characteristic of one of the four photomultiplier tubes viewing the crystal not operating correctly, this could be caused by mechanical damage or a bad connection between the PMT and the analogue signal processor.

Histogram for bucket 16, block 0, time 10:

2	0	0	0	3	2	20	1325
2	0	0	0	10	121	376	708
3	0	8	139	231	281	281	499
1292	602	628	352	466	297	365	994

Mean = 281 (+/- 17) Stdev = 67.16 Center = (-4.16, 1.91)

Figure 5.6: Histogram for broken detector block

A useful function when performing PEPT experiments is **Rate**. This function polls each coincidence processor in turn to determine its instantaneous count rate for prompts, delayed and multiple events. During a PEPT acquisition this is useful to determine when the tracer particle is within the field of view of the camera.

For example figure 5.7 shows the output upon giving a **Rate** command during data acquisition for one of the PEPT runs in the University Interdisciplinary Research Centre (IRC) described in section 7.5. Tracer motion is a 'once through' event where the tracer starts at one side of the camera, passes through the field of view, and ends up on the other side. The tracer particle is then recovered and reintroduced to the system upstream of the camera.

```
Command  > rate
  0,1,0                 7,15,0                1,0,0                 2,0,0                 Wed Aug 20 17:02:24 2008
  0,1,0                 3,15,0                1,0,0                 0,0,0                 Wed Aug 20 17:02:26 2008
  0,0,0                 3,15,0                0,0,0                 0,0,0                 Wed Aug 20 17:02:28 2008
  6,0,0                 2,15,0                1,0,0                 3,0,0                 Wed Aug 20 17:02:29 2008
  6,0,0                 2,15,0                1,0,0                 3,0,0                 Wed Aug 20 17:02:31 2008
  9,8,0                 0,7,0                 10,0,0                4,6,0                 Wed Aug 20 17:02:33 2008
  30,8,0                11,7,0                10,3,0                4,6,0                 Wed Aug 20 17:02:35 2008
  30,45,0               11,20,0               33,72,0               26,51,0               Wed Aug 20 17:02:36 2008
  414,45,0              296,20,0              33,72,0               419,51,35             Wed Aug 20 17:02:38 2008
  414,623,6             296,559,0             448,775,0             419,586,99            Wed Aug 20 17:02:40 2008
  1411,1775,6           1237,1246,0           2600,775,14           1457,586,99           Wed Aug 20 17:02:42 2008
  1411,1775,32          4319,1246,30          2600,1584,53          1457,1442,226         Wed Aug 20 17:02:43 2008
  5203,3402,32          4319,2547,30          5428,1584,53          2875,2555,226         Wed Aug 20 17:02:45 2008
  33179,3402,212        17642,2547,159        5428,4082,375         9070,2555,599         Wed Aug 20 17:02:47 2008
  33179,9981,893        17642,6934,585        22809,7121,375        9070,5503,599         Wed Aug 20 17:02:49 2008
  68283,9981,893        26203,7864,585        30259,7121,656        14550,5503,1044       Wed Aug 20 17:02:50 2008
  68283,11534,1440      26203,7864,1004       30259,9058,910        14550,7718,1472       Wed Aug 20 17:02:52 2008
  87372,11676,1440      46006,9572,1004       52836,9058,910        22004,8817,1472       Wed Aug 20 17:02:54 2008
  66852,11676,1232      70524,9572,780        52836,9431,1713       22004,8817,2009       Wed Aug 20 17:02:56 2008
  66852,11467,1232      70524,10495,1085      65616,9399,1713       26253,8835,2009       Wed Aug 20 17:02:58 2008
  51380,11467,1349      60861,10495,1085      66045,9399,1151       26934,8835,1616       Wed Aug 20 17:02:59 2008
  51380,10160,948       60861,9581,1519       66045,9623,1151       26934,9319,2574       Wed Aug 20 17:03:01 2008
  36071,8706,948        78934,10444,1519      65671,9623,1366       57864,9319,2574       Wed Aug 20 17:03:03 2008
  36071,8706,590        40600,10444,1095      65671,8566,1099       57864,11320,3397      Wed Aug 20 17:03:05 2008
  24047,7240,590        40600,8861,1095       28055,5437,1099       69278,8388,3397       Wed Aug 20 17:03:06 2008
  8001,7240,240         15031,8861,389        28055,5437,236        7802,8388,1015        Wed Aug 20 17:03:08 2008
  8001,3256,52          15031,3034,32         5475,1963,236         7802,1983,178         Wed Aug 20 17:03:10 2008
  1464,695,52           1958,775,32           992,1963,9            2041,1983,178         Wed Aug 20 17:03:12 2008
  1464,695,2            1958,775,0            992,255,2             2041,415,14           Wed Aug 20 17:03:13 2008
  162,26,2              178,43,0              78,255,2              131,93,14             Wed Aug 20 17:03:15 2008
  50,26,2               44,43,0               78,21,0               101,93,5              Wed Aug 20 17:03:17 2008
  50,15,2               44,30,0               80,13,0               101,29,5              Wed Aug 20 17:03:19 2008
  57,15,2               45,18,0               84,13,0               71,29,5               Wed Aug 20 17:03:20 2008
```

Figure 5.7: Output display for Rate command

The tracer moves through the field of view of the camera starting with detectors connected to IPCP 64, then passes through the middle of the system where IPCP 65 and IPCP 66 receive the most data as detectors in both rings are seeing a response. Finally the source passes through the second ring and IPCP 67 sees the most data. The total coincidence response rapidly reaches a high number of events and remains roughly constant (as the field of view is uniform) until the source leaves the field of view on the other side of the camera.

The serial transmission is slow and **Rate** data is produced on average every 1.75 s from all four coincidence processors at a time.

These data are plotted in figure 5.8 along with the PEPT locations for this experimental run.

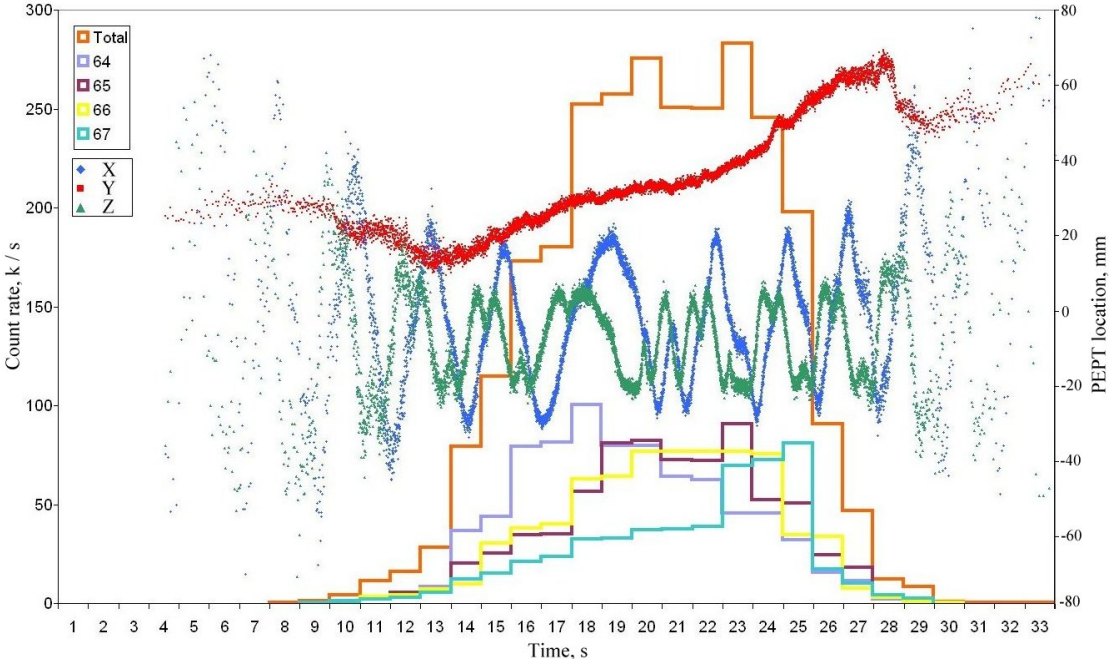

Figure 5.8: Coincidence count rates and PEPT locations for an IRC run

5.2 Data Acquisition Software

Software has been written to interface with the data acquisition card and to store data to disk. The program **DAC** is used for this, a flow chart describing its operation is shown in figure 5.9.

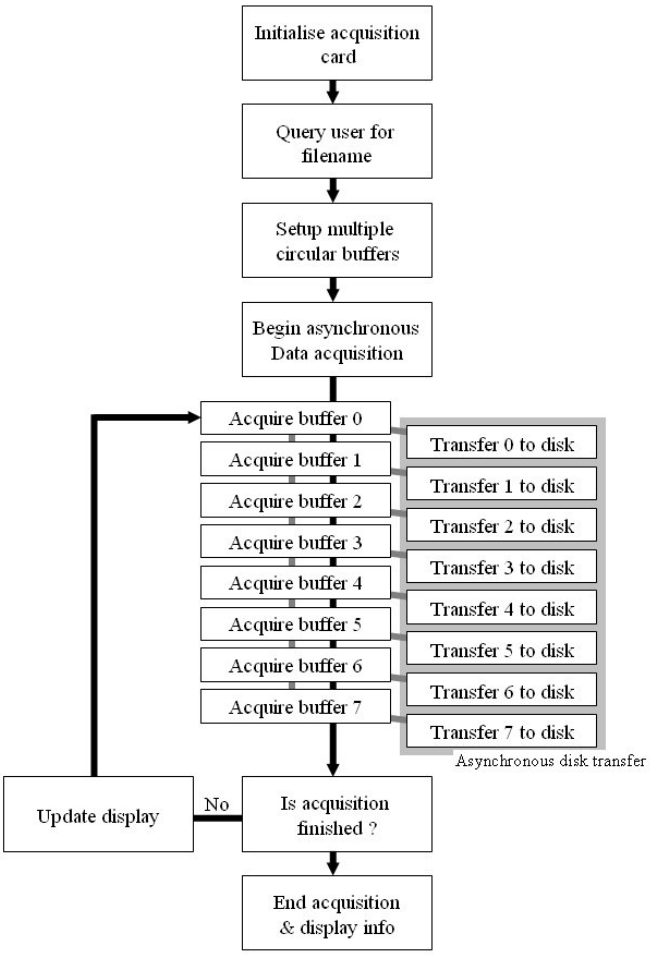

Figure 5.9: Flow chart for data acquisition software

The first step taken is to initialise the PCI card, this is done using C functions that are provided in a software library by the manufacturer of the hardware card. Up to 16 identical cards can be used in a single computer (provided the PCI bus has enough empty ports), with each device sharing control of the PCI and system busses. In this application only one card is used and this is designated the device identification number 0.

The card is setup to have a 32 bit wide input (using both 16 bit ports A & B) with active termination, the input sampling is provided by the trigger signal on the DI_REQ pin on the card.

The software produces a list of the files in the local directory where the program has been initiated, and queries the user for a filename to save the data into. Error checking is provided to prevent existing data files being overwritten. A header containing information about the experimental conditions and the modular camera geometry used can be written to the output file at this stage. Eight buffers are then constructed in system memory as a set of multiple circular buffers for temporary storage of the data before it is streamed to disk.

The function **DI_ContMultiBufferStart(card, port, samplerate)** is then run. This function takes the arguments card (device identification number equal to 0), port (32 bit wide input port, designated the value 0) and samplerate; as the sampling is triggered by the external trigger pulse the sample rate is set to the integer value of 1 (designating external sampling).

The function begins the acquisition in hardware. This is processor free and gives the acquisition card control of the PCI bus where it begins streaming data into the allocated memory space of the first circular buffer. As this occurs asynchronously the function returns a value reporting success (0) or an error code (integers less then 0). The processor is then free to run the remaining code.

The program then enters a while loop which is terminated by the pressing of any key, signifying the end of the acquisition process. The first step of the while loop is a further while loop which checks to see if the first buffer has been filled and terminates when data acquisition into the first buffer is completed. The asynchronous data transfer then acquires data into the next circular buffer.

The acquired buffer is then transferred to disk by appending the data file with the data from the completed circular buffer. During this time the incoming data are streamed into the next circular buffer via DMA without interruption.

The program updates the display with information regarding the progress of the acquisition, this involves displaying the number of the current circular buffer and the total number of buffers acquired. It then checks the error flags to ensure that buffer overflow has not occurred and that the transfer of data to disk has not delayed the program long enough to miss the following buffer. The end of the while loop is then met and the code again waits for the next buffer to acquire.

Once the acquisition is complete the while loop is terminated and the memory space for the circular buffer is released and returned to available system memory. The output file is properly closed and the acquisition hardware card is turned off. Acquisition information is displayed (filename, start and end times and the total number of buffers acquired) and any error messages generated during the acquisition are reported.

5.3 Data Stream Analysis

A number of utilities have been written to analyse the data stream in various manners. The first program described here, called **Readbin**, reads in the data file and displays the binary data words. The flowchart for this program is given in figure 5.10. The program opens the data file, reads in the data words and converts them into a string of thirty two 0's and 1's representing the binary number. Each bit is tallied in a histogram and the string is displayed on screen and written to an output file.

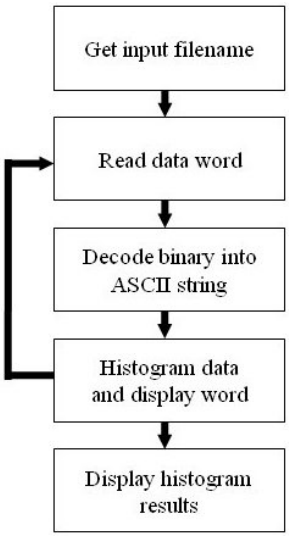

Figure 5.10: Flow chart for Readbin

Figure 5.11 shows a section of a typical output file and table 5.3 shows the histogram data for a PEPT run.

Time 0 - 3	PlaneB 2 (MSB)	BlockB 1	BlockB 0	SegB 2	SegB 1	SegB 0	PlaneA 2 (MSB)	BlockA 1	BlockA 0	SegA 2	SegA 1	SegA 0	Remote Trigger 1	Remote Trigger 0	Board Address 1	Board Address 0	DetA Plane 1	DetA Plane 0	DetB Plane 1	DetB Plane 0	Multiple Delayed	Module Pair 0 - 5	Hexadecimal value	Absolute time, ms
1101	1	10	001	1	10	110	00	10	00	00	00											100111	c762027	6.5
1101	0	01	000	1	01	100	00	10	01	10	00											100111	22c2627	6.5
1101	1	11	001	0	11	100	00	10	10	10	00											100111	e5c2a27	6.5
1101	1	10	100	0	10	010	00	10	00	11	00											000100	d122304	6.5
1101	0	01	101	1	10	010	00	10	10	10	00											100111	3722a27	6.5
1110	1	11	001	0	10	101	00	10	11	10	00											000100	e552e04	7.0
1110	1	11	010	0	11	100	00	10	00	11	00											100111	e9c2327	7.0
1110	1	10	101	0	10	010	00	01	10	11	00											000100	d521b04	7.0
1110	0	10	110	1	11	000	00	01	11	01	00											100111	5b81d27	7.0
1110	1	01	010	0	01	110	00	01	01	11	00											100111	a8e1727	7.0

Figure 5.11: Output from Readbin

Bit	Histogram
Time 3	3109347
Time 2	3109850
Time 1	3107561
Time 0	3109575
BlockB 1	2493498
BlockB 0	2822784
SegB 2	3044843
SegB 1	3056832
SegB 0	3080462
BlockA1	3801219
BlockA 0	3326898
SegA 2	3181182
SegA 1	3137549
SegA 0	3129459
Remote Trigger 1	0
Remote Trigger 0	1422
Board Address 1	3776616
Board Address 0	3196536
DetA Plane 2	3217258
DetA Plane 1	3222525
DetA Plane 0	3120952
DetB Plane 2	3182188
DetB Plane 1	3219356
DetB Plane 0	3139239
Delayed	176629
Multiple	32471
MP5	2921958
MP4	440093
MP3	6498
MP2	4817597
MP1	2921952
MP0	3890781

Table 5.3: Histogram data for typical file from Readbin

The program **Readword** is similar to **Readbin**, however it analyses the data further. **Readword** reads in the data file and processes the data to determine the type of event (prompt, random, multiple), the time of the event and the identification numbers of the two crystals which detected the LOR. The flowchart for **Readword** is shown in figure 5.12 and a typical output is shown in figure 5.13.

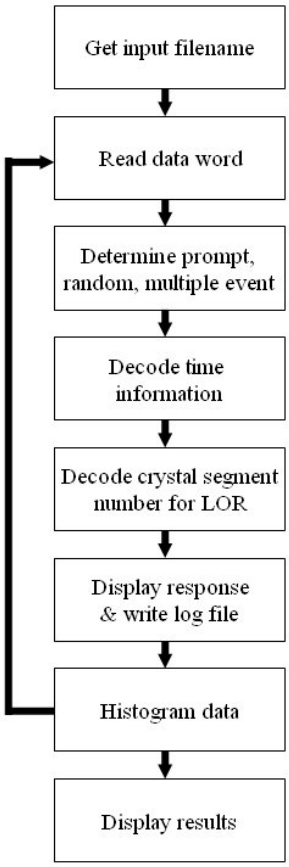

Figure 5.12: Flow chart for Readword

The detector bucket number is determined from the module pair number encoded in the list mode data. The module pair numbers are hard-coded into the array **MPnum[112]** which has 112 elements holding both bucket numbers describing the module pair (56).

An array with 56 elements is constructed of the structure **ModulePairs**; defined to contain two integers named DetectorA and DetectorB which is populated with the numbers for both detectors in a module pair. The index of the array is then the module pair number, with the returned values being the numbers of each detector.

The axial and transaxial crystal numbers are described in the conventional manner outlined in section 3.3 and can be derived directly from the data word. Time information is decoded and used to calculate the average event rate for the data file, and any remote trigger events are tallied and reported.

A typical output from **Readword** is shown in figure 5.13.

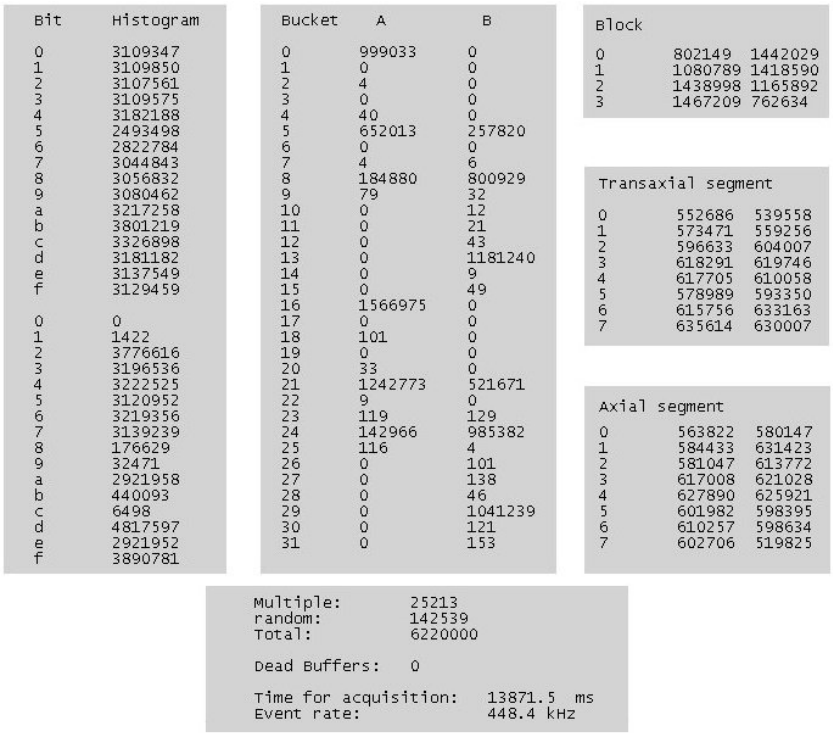

Bit	Histogram
0	3109347
1	3109850
2	3107561
3	3109575
4	3182188
5	2493498
6	2822784
7	3044843
8	3056832
9	3080462
a	3217258
b	3801219
c	3326898
d	3181182
e	3137549
f	3129459
0	0
1	1422
2	3776616
3	3196536
4	3222525
5	3120952
6	3219356
7	3139239
8	176629
9	32471
a	2921958
b	440093
c	6498
d	4817597
e	2921952
f	3890781

Bucket	A	B
0	999033	0
1	0	0
2	4	0
3	0	0
4	40	0
5	652013	257820
6	0	0
7	4	6
8	184880	800929
9	79	32
10	0	12
11	0	21
12	0	43
13	0	1181240
14	0	9
15	0	49
16	1566975	0
17	0	0
18	101	0
19	0	0
20	33	0
21	1242773	521671
22	9	0
23	119	129
24	142966	985382
25	116	4
26	0	101
27	0	138
28	0	46
29	0	1041239
30	0	121
31	0	153

Block		
0	802149	1442029
1	1080789	1418590
2	1438998	1165892
3	1467209	762634

Transaxial segment		
0	552686	539558
1	573471	559256
2	596633	604007
3	618291	619746
4	617705	610058
5	578989	593350
6	615756	633163
7	635614	630007

Axial segment		
0	563822	580147
1	584433	631423
2	581047	613772
3	617008	621028
4	627890	625921
5	601982	598395
6	610257	598634
7	602706	519825

Multiple:	25213
random:	142539
Total:	6220000
Dead Buffers:	0
Time for acquisition:	13871.5 ms
Event rate:	448.4 kHz

Figure 5.13: Output from Readword

It is sometimes useful to generate sinogram data from the list mode data stream. This sinogram could then be analysed in the normal fashion and used to produce an image of the distribution of radioactivity within the detector volume or used for fault finding. This would be especially important if medical imaging (PET) was to be performed using the new data acquisition system.

A utility called **Sino** has been written in order to compute and display a sinogram from the raw data file. The flowchart for the program is shown in figure 5.14.

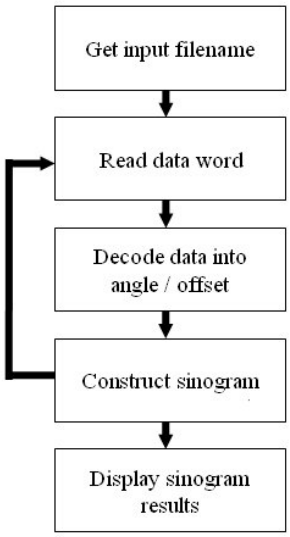

Figure 5.14: Flow chart for program Sino

Figure 5.15 shows the output of **Sino** for a large cylindrical source positioned in the centre of the small ring camera described in section 6.3. As the camera radius was reduced by removing the odd numbered detector buckets in the ring there are gaps in the sinogram corresponding to the missing detectors.

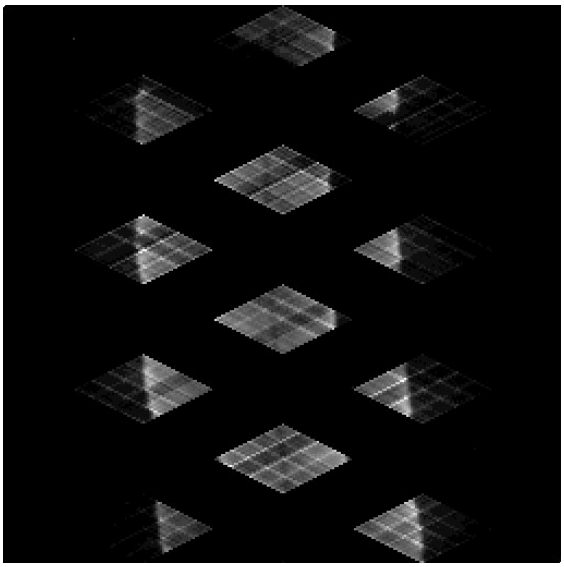

Figure 5.15: Output sinogram from program Sino

5.4 Visualisation

A program has been written using the Windows Graphical Device Interface (GDI) to display the raw data saved in the file. This has proved extremely useful as the geometry can be easily changed in the software code and the graphical display used to confirm the software modelling of the camera geometry.

The flow chart for this program, called **Displor_win**, is shown in figure 5.16.

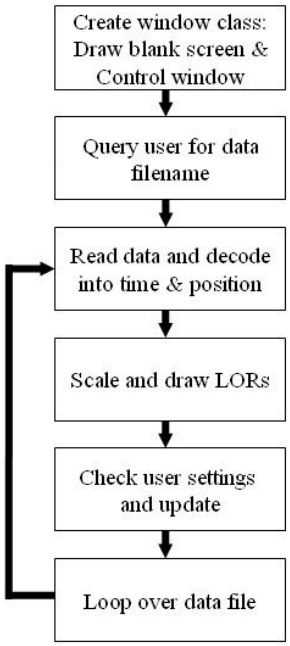

Figure 5.16: Flowchart for program Displor_Win

Upon execution the program creates a window class within the Windows operating system. This involves creating a blank window with maximise, minimise, resize and close functions and a control window which is used to interact with the user and change display settings. The blank window is then filled with the graphical background and data file information. The user is asked to provide a data file using a standard windows open file class.

Depending upon the user display settings the program reads between 5 to 900 events from the data file into system memory. These events are processed to calculate the end points of the described line of response and event type and time information. A line is then drawn on screen between the coordinates of the LOR end points (scaled by a factor to fit in the window) and coloured according to the event type (prompt, delayed, multiple).

The program can be instructed to display events in real time by either delaying the display process between events (low event rate) or by dropping events to keep time (high event rate). Colour coding is also used to display data from different detectors or extra modules (via copied data, section 6.5) as well as colour coding the data from each coincidence processor.

Prompt and random events can be separately displayed and three orthogonal projection views of the camera system are shown. Figure 5.17 shows a typical display from the **Displor_Win** program, for data acquired using the small camera described in section 7.5.

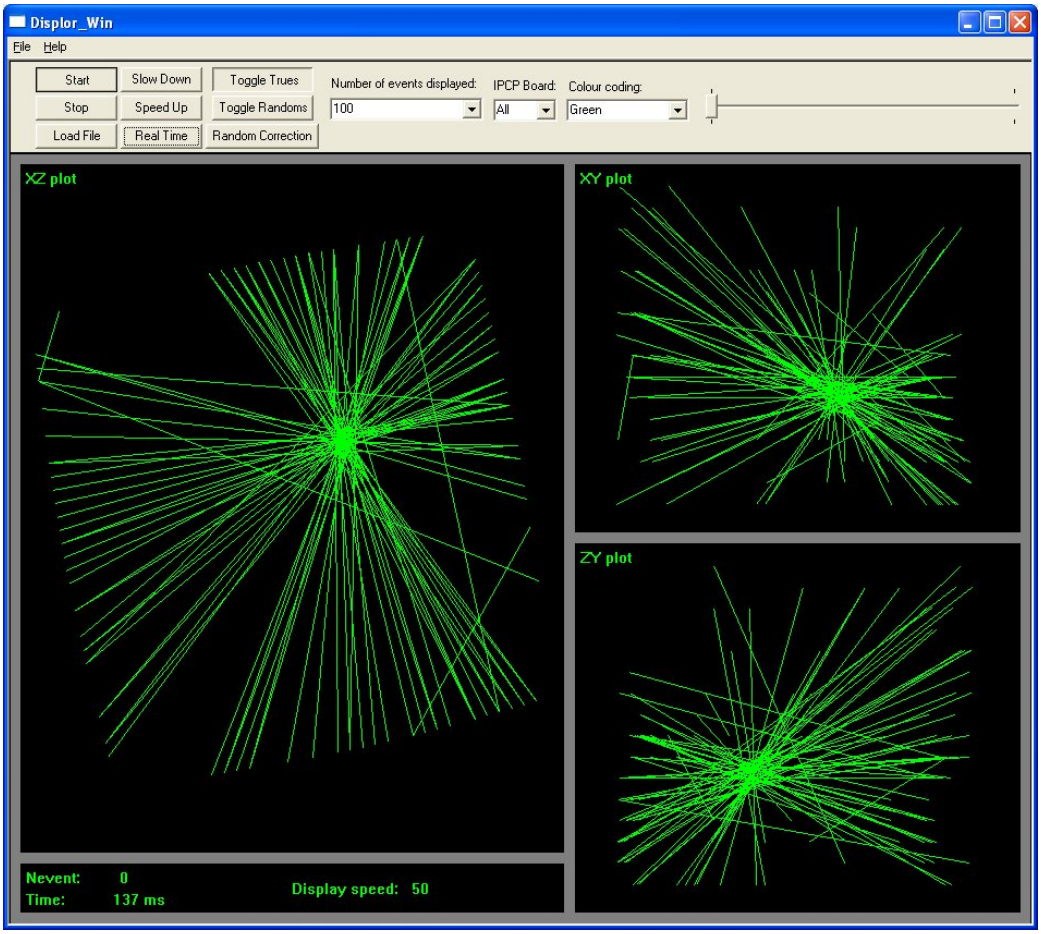

Figure 5.17: Displor_Win output display

The Windows GDI operations are relatively slow as the program uses the CPU to calculate the display as a bitmap in memory, and then transfers this bitmap to the screen at the refresh rate of the computer display. Game programming libraries offer a wealth of high speed graphic facilities utilising the graphics processor (GPU) built into the computer graphics card.

The graphics card is designed for high speed display of visual data and performs many graphics calculations in hardware. Example functions include the display of triangles and lines and the ability to correct for depth perception and lighting of a scene. Other useful features include object scaling and rotation in a three dimensional space.

Using the open source graphics library OpenGL it has been possible to write software which utilises the graphics processor to display data from the camera. The program **Displor_GL** is used to display data in a similar manner to the **Displor_Win** program described above. The flow chart given in figure 5.16 adequately describes the program structure.

In the **Displor_GL** program the lines of response are displayed in three dimensions using the perspective control functions. The field of view can be scaled, rotated and translated allowing for different views to be seen for the data. Again processor time is wasted if the data occurs at a low rate and frames are dropped if the data rate is higher then the display rate.

Figure 5.18 shows the display output of **Displor_GL** for the same data file shown for the **Displor_win** program in figure 5.17 above.

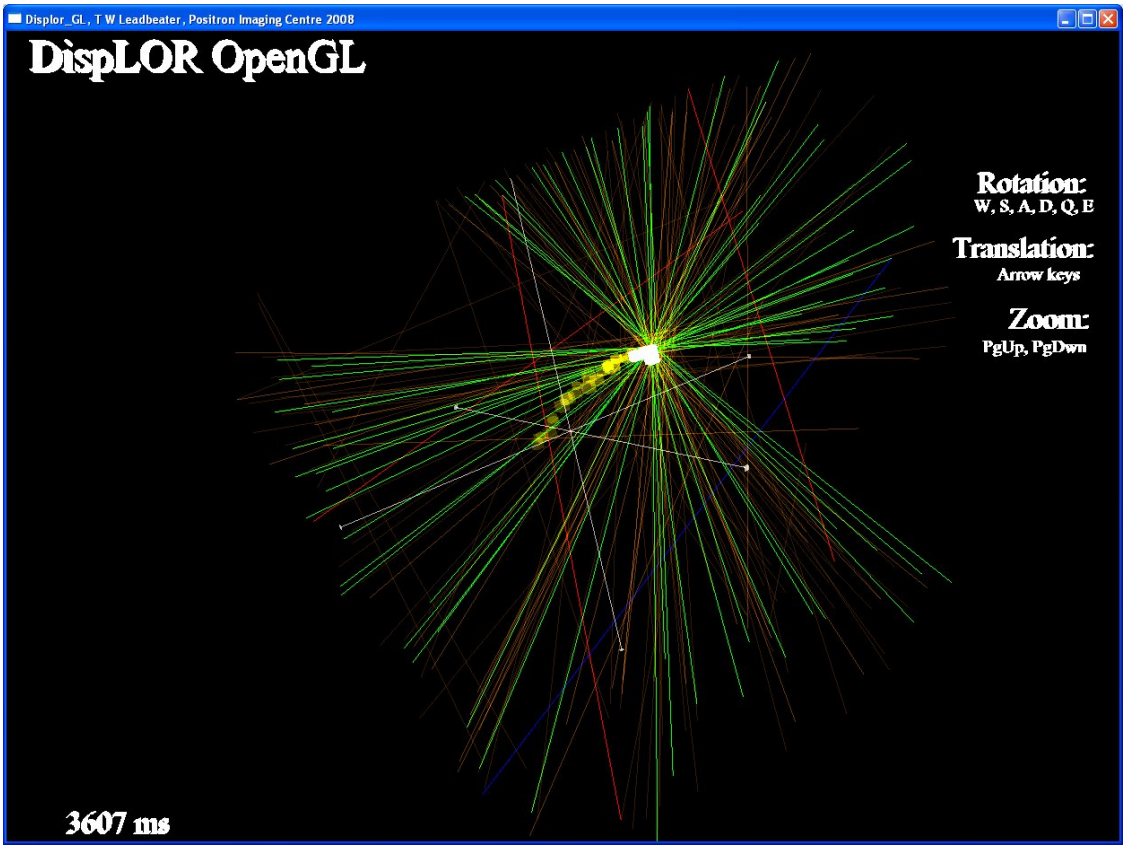

Figure 5.18: Displor_GL output display

The PEPT location is calculated on the fly and displayed as a cube with dimensions given by the PEPT error (d) overlaid upon the LOR data, algorithm parameters are currently fixed due to speed considerations: f is set to be 10 % with N equal to 100.

A version of **Displor_GL** has been written in order to display the raw LOR data and a fast PEPT location for real time data. This program combines the data acquisition software **DAC** with the code for the **Displor_GL** display application.

This software can be used for demonstration and display purposes where the display is projected onto the experimental rig in real time using a LCD projector. The source would then be tracked using the PEPT technique with the visual element of the software overlaid onto its true position. This could prove to be a very powerful technique for visualisation and for demonstration of PEPT experiments.

5.5 PEPT Algorithm

Finally a discussion of the PEPT algorithm is given. The flow chart describing the PEPT program is given in figure 5.19.

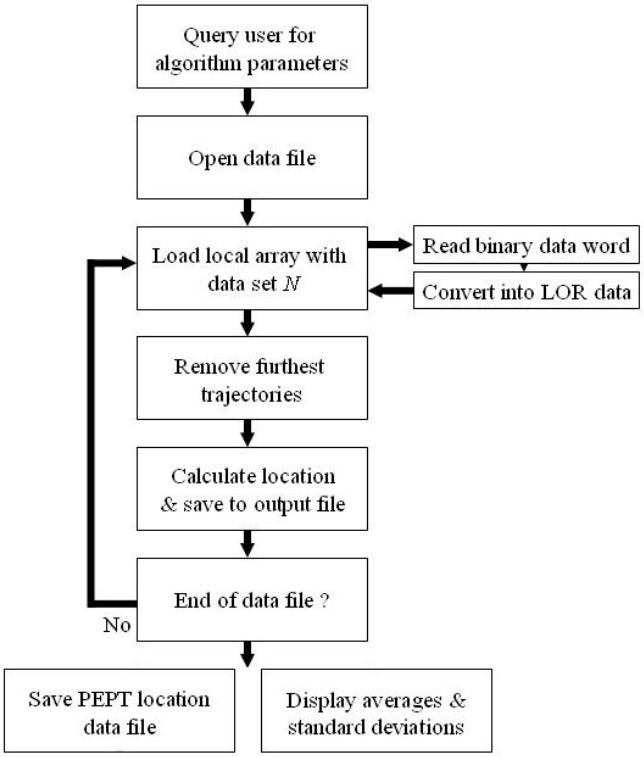

Figure 5.19: Flow chart for the PEPT algorithm

The user is asked to provide a filename where the raw data are to be saved and to select the algorithm variables (N, f). The data are read into a local array on an event by event basis. Events which are genuine (prompt events) are converted into a set of two Cartesian coordinates describing the end points of the line of response.

The other events (delayed, multiple) are tallied and rejected, time information and remote trigger data words are handled at this point.

The data word is decoded to give the bucket numbers, block numbers and crystal elements within the block for each of the end points of the LOR. When the relationship between crystal identification number and the position of the crystal in space is known, the coordinates describing the LOR can be determined and processed by the PEPT location algorithm.

The detector geometry is modelled as the structure type **BucketLocations** which contains three variables of type double. These variables describe the (X, Y, Z) locations of the bottom left corner of the detector module. The variable **BucketPos**, of type BucketLocations is declared to be an array 32 items in length, each describing the position of the detector bucket relative to a defined origin (usually the position of detector module 0).

Simple relationships between the crystal element within a bucket are then used to describe the position of the LOR event. Once the detector number and the location of the bottom left of the detector is known, the position of the excited crystal element can be determined.

For example in the Y direction (axial) the centre of the bottom crystal is positioned half a crystal height above the origin describing the detector position. Crystals are labelled from 0 to 7 increasing with height, therefore the crystal number multiplied by the height of a crystal gives the offset of the crystal from the origin, half a crystal height is added in order to locate the event in the centre of the crystal. Similar arguments can be used for the other two axes.

Once the positions of the end points of each LOR are known this data can be passed onto the conventional PEPT algorithm. This code has existed for many years in various guises, modifications are made as necessary to include extra data or if further calculations are required. The basis of the PEPT algorithm is described in section 2.2.

The data set N is loaded into an array and processed to determine the minimum distance point. The program iterates over the data set and removes the trajectories which are furthest away from the minimum distance point until a pre-defined fraction f is left. The minimum distance point is calculated for the last time and the error d is determined. The final location, error and the average time of the trajectories used in the final location are then written to a file describing the PEPT location response for the data.

The program loops over the entire data file and provides PEPT locations for the particular data file which can be analysed using standard techniques (e.g MS Excel, Mathcad etc.). The final output displays the average location along each axis, its standard deviation and the average and deviation in the PEPT error.

Chapter 6

Modular Camera

The advantage of the ECAT scanner series is its high efficiency for detecting 511 keV gamma photons; compared to sodium iodide scintillator crystals the BGO scintillators used in the ECAT scanners have a photopeak efficiency a factor of approximately 3 higher. The high spatial resolution (from using many small crystal elements covering a large angle) and its corresponding high data rate as many parallel channels allow simultaneous event acquisition are further benefits. Unfortunately the construction of the scanner gantry (figure 3.1) is highly restrictive to the study of engineering applications because the tightly packed detector elements do not offer a wide field of view, nor do they allow large equipment to be positioned within the field of view due to physical restrictions in size.

Once the data storage conventions are known, and data describing an event in an individual crystal is explicitly defined, the possibility of using the modular detector blocks free from the constraint of the support gantry can be entertained. Provided that the detectors are arranged around the system under study in such a manner that the allowed coincidence combinations are satisfied (section 3.3.1), the detector system will provide genuine coincidence LOR data. It is then only necessary to know the relationship between crystal number and the relative spatial position of the crystal to define the coordinate system, and thus give the locations of both end points of each received LOR.

This chapter discusses the physical changes that have been made to the original scanner system such that the detectors can be optimised for engineering studies. Primarily this has involved the removal of the detectors from the scanner gantry; resulting in a flexible system allowing the study of a wide range of devices on many different volume scales. The flexible geometry allows customisation of the field of view for the study of a particular piece of apparatus, it also makes a transportable system that can be used when equipment cannot be brought into the Positron Imaging Centre.

6.1 Motivation

The acquisition of the ECAT scanners by the Positron Imaging Centre has been seen as a great advancement. This allows for continuing PEPT studies on many systems using the ADAC Forte camera already commissioned in the lab, but also allows the development of the imaging systems used for these studies.

The new systems offer increased data rates and improved spatial resolution, but in order to fully realise the potential of these systems for PEPT experiments they have required significant modification. As mentioned above, the format of the data word is required to be known in order to represent the crystal elements describing each LOR in a coordinate system that can be used for the PEPT algorithm. List mode data storage and the inclusion of time information in the data stream is also an essential modification for PEPT work.

The ADAC Forte camera is highly flexible in that it can accommodate relatively large pieces of equipment due to its open geometry. The camera has two opposing gamma camera heads, both of which can be rotated and translated in both horizontal and vertical directions. The separation of the camera heads can be adjusted allowing the optimisation of the field of view for the particular system under study.

Shown in figure 6.1 and 6.2 are photographs of the ADAC camera used to study a fluidised bed and a horizontal-axis scraped surface heat exchanger respectively.

Figure 6.1: ADAC camera used to study a fluidised bed

It is clear from the photographs that these systems would not fit within the tightly confined space of the patient port of the ECAT scanners (76 cm diameter, 10.6 cm depth, shown in figure 3.1).

Figure 6.2: ADAC camera used to study a scraped surface heat exchanger

These considerations have motivated the development of the modular camera, an imaging system that makes use of the modular construction of the ECAT scanners where the detector blocks are independently placed around the scanner support gantry.

6.2 Application

As discussed in sections 3.1.1 and 3.1.2 the ECAT scanners consist of a large number of identical detector blocks, each with its own support electronics. Detector buckets each hold 4 blocks mounted on a sturdy metal plate which holds the blocks in the arc of the scanner gantry, and holds the bucket controller and associated electronics. It proved to be relatively easy to remove the buckets (and their support structure) from the scanner gantry and experiment with these devices outside of the fixed ring geometry.

Sadrmomtaz [34] performed these initial trials in a proof-of-concept experiment. Fourteen buckets were removed from the gantry of a 931 scanner and arranged in a cuboidal geometry, shown in figure 6.4. In this diagram each bucket module is illustrated as a block as shown in figure 6.3, note that the bucket remains mounted on its original mounting plate and thus the blocks are arranged in a curve with radius 510 mm as depicted.

Two orthogonal views are shown of the cuboidal setup in figure 6.4; a side view and the plan view. Some of the allowed lines of response are illustrated between the modules making up the system.

The sensitivity varies as a function of position as in some places there are more allowed LORs then in others; this is illustrated as the density of the lines and is discussed further in section 6.4. Note that these diagrams are not drawn to scale.

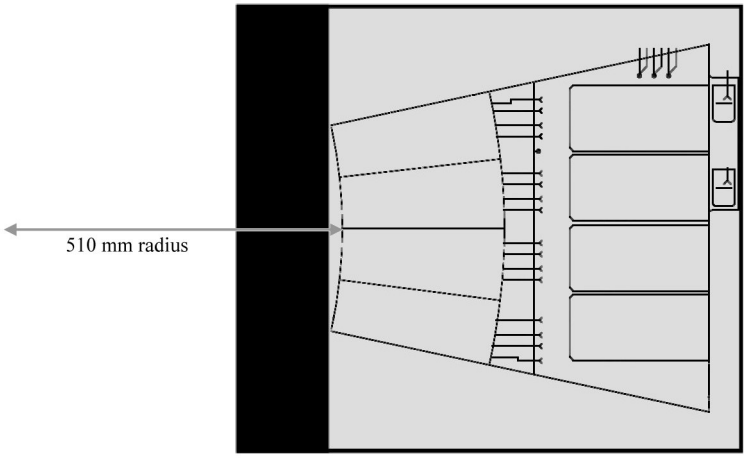

Figure 6.3: Depiction of detector bucket as a module

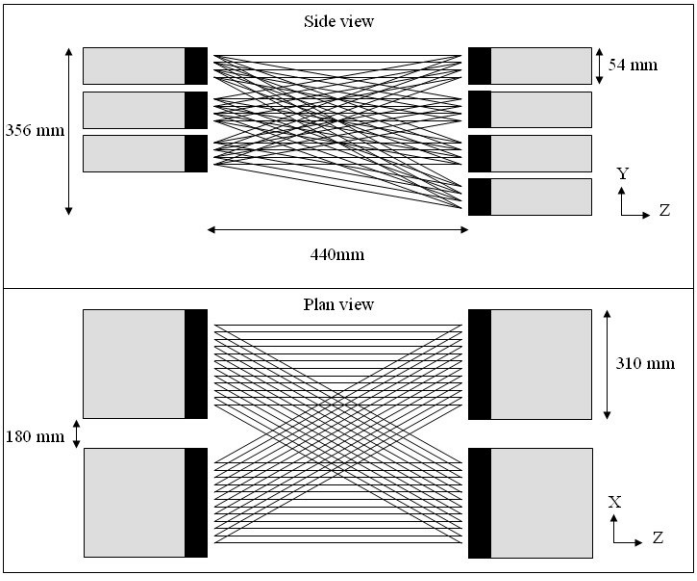

Figure 6.4: Cuboidal proof-of-concept geometry

This system was used to prove that the detector modules would operate correctly in a newly configured geometry. Careful measurement of the relative position of each bucket gave a conversion from the crystal number data to a (x, y, z) location for each event. This took the form of measuring the relative position of the bottom left crystal in each bucket, once this was known knowledge about the bucket geometry then gives the position of each crystal.

For example in the axial direction (now the vertical) the bottom crystal is located at zero, each axial (vertical) crystal is then 6.75 mm apart (given by the crystal pitch in the 951 series blocks). Similarly in the transaxial (now horizontal) direction, crystal numbers increase from the zeroth crystal in block zero, to the seventh crystal in block three. The pitch of the crystals in this direction is 6.25 mm, this is illustrated in figure 6.5. At this stage the arc of the detector blocks was not taken into consideration.

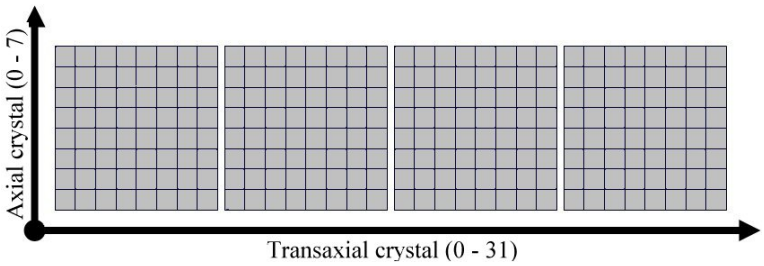

Figure 6.5: Crystal location within a 951 series bucket

It was confirmed that the sensitivity varies strongly across the field of view for this geometry of detector modules by a series of experiments using a point source placed at different locations within the field of view [22].

This can be demonstrated from considerations in geometric efficiency as illustrated in figure 6.6. In A, a source is central to the detectors and offers a wide cone of allowed response. In B, the source has moved off-centre and the allowed response is therefore reduced. In C, the source has moved to the edge of the field of view and only one LOR can link two detectors whilst simultaneously passing through the source; therefore the sensitivity is reduced to zero at the edge of the field of view.

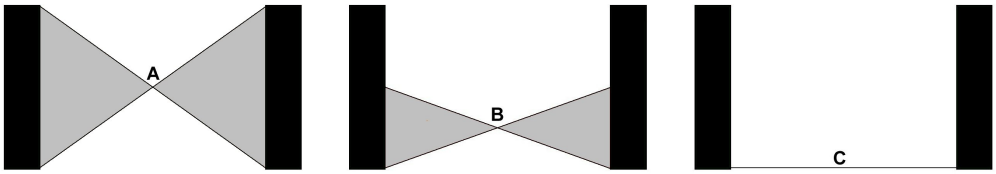

Figure 6.6: Variations in sensitivity as a function of source position

A uniform field of view is important when performing PET imaging because each image depends upon the concentration of the radioisotope in each volume element. In PEPT experiments however, the uniformity of the field of view is not such an issue. This is because each location occurs on a short timescale (order of ms) and depends only on the number and position of the LORs used in the particle location. Variations in sensitivity across the field of view then result in more or less frequent locations with the same precision, or a fixed location rate with the precision varying as a function of position.

This camera was used for a number of trial studies in particle tracking [22]. These used point like sources either positioned at various locations within the field of view, or sources subject to some well defined motion, such as a turntable or pendulum. These studies showed that it was feasible to use this system to generate meaningful PEPT data.

It was decided that it would be beneficial to continue to develop the modular camera concept for PEPT studies as the flexible geometry (and the fact that the camera would then in principle be transportable) opens up a wide range of potential applications on apparatus ranging in size scales. To this end the 16 available 931 detector buckets (used in the proof-of concept experiment) were removed from their support plates and mounted into protective boxes containing 4 detector blocks and the bucket controller electronics. The boxes are fan cooled and provide connections for power supplies, clock serial bus, data serial bus and high voltage.

Given as reference in figure 6.7 is a diagram of a module, with the cable format given in figure 6.8.

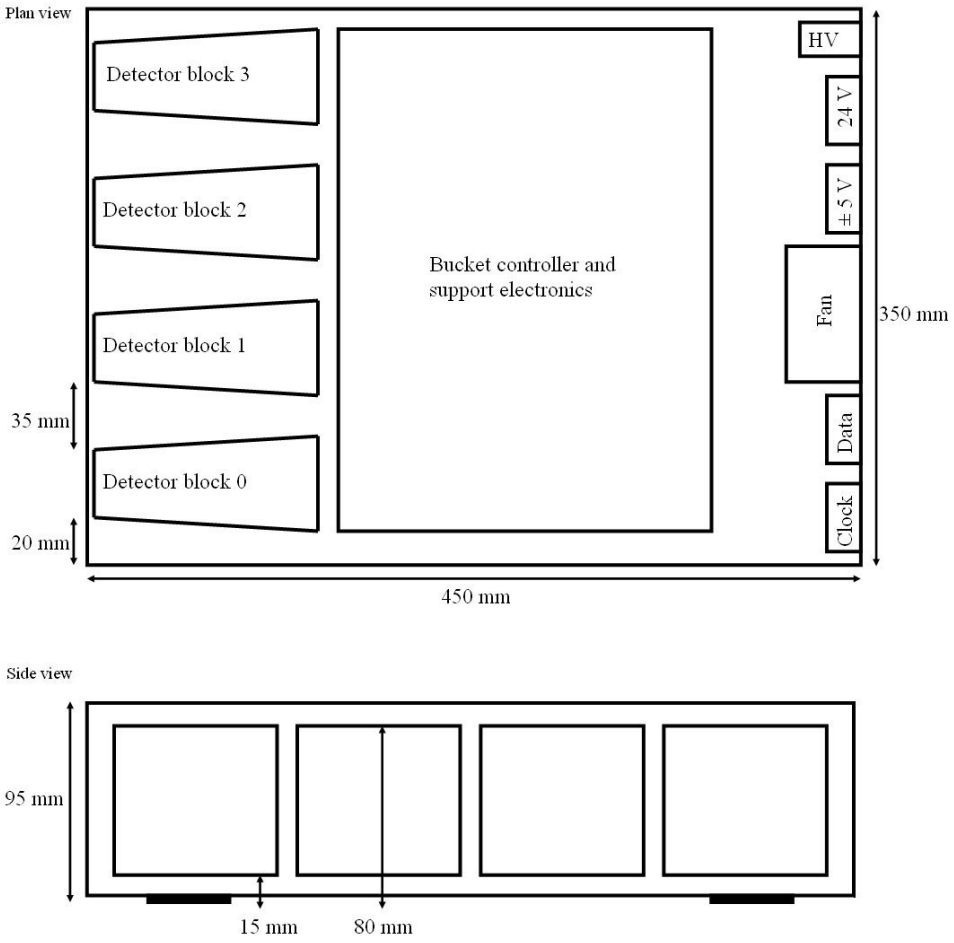

Figure 6.7: Detector module

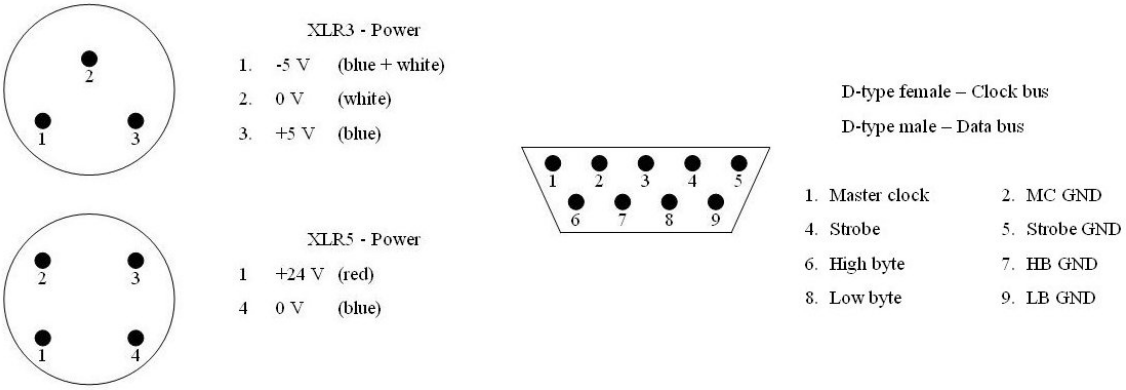

Figure 6.8: Module cable pin-outs

The photograph in figure 6.9 shows a typical detector block (right) and a bucket (left) with the four blocks mounted in an arc as in the original scanner system. The detector bucket also contains the bucket controller from a 951 series scanner.

The central image shows a detector module with the blocks separated and placed in an orthogonal geometry, the bucket controller is placed behind within the protective casing.

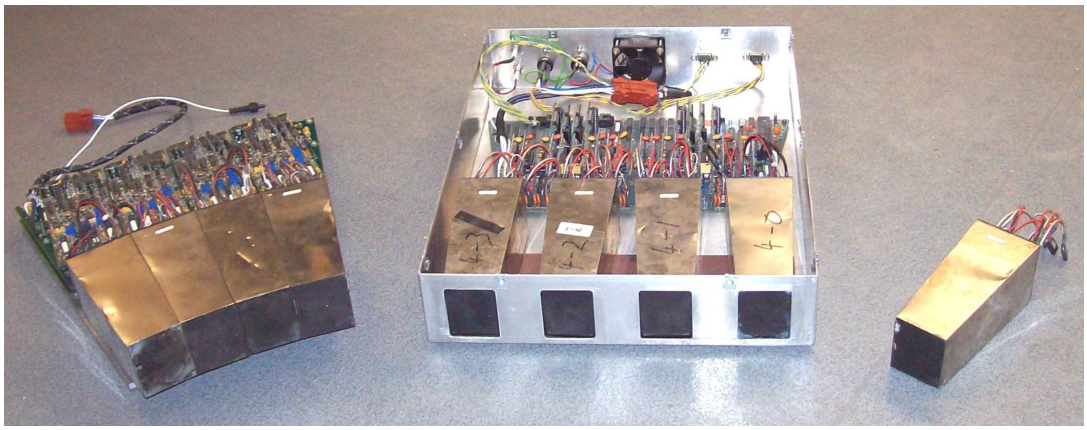

Figure 6.9: Bucket, detector module, and block

Figure 6.10 shows all of the allowed lines of response for two of these modules separated by approximately 300 mm. Note that due to the separation of the blocks there are regions of high and low sensitivity, shown by the density of the lines. In reality the lines joining the two detector elements are tubes and therefore the granularity of this image is reduced, this is discussed further in section 6.4.2.

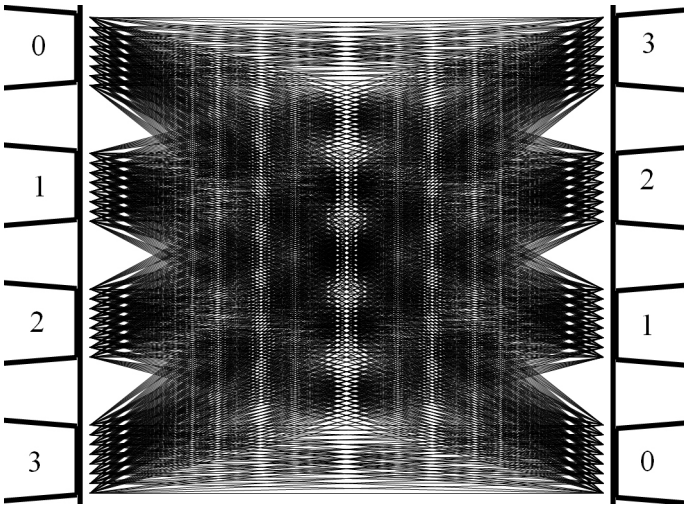

Figure 6.10: Transaxial LORs for two directly opposing modules

6.3 Changing the Geometry

As discussed in section 3.3.1 each detector bucket only operates in coincidence with the opposing seven buckets within the gantry ring. It is then important to arrange the detector modules in an arrangement so that the system under study is placed between valid module pairs. The field of view is no longer uniform using these devices, however the geometry that is covered by the field of view is flexible and adaptable. By correct choice of geometry the regions of interest within the experimental apparatus can be covered with a sensitive, and relatively uniform, field of view.

The first trial of the modified detector system saw the modules arranged in the same cuboidal geometry used for the initial trials. The software was modified in order to take account of the separation of the detector blocks in the new modules. Advantageously the arc correction required for the previous system was not needed as the detectors were then arranged in a cubic manner. This system was used to study a tracer placed on a turntable rotating in the XZ plane within the field of view of the camera and was demonstrated to operate as expected following the proof-of-concept trials.

Following the success of these trials the modular camera was transported to the Chemical Engineering Department on the University campus. This was to perform PEPT experiments on a high pressure fluidised bed. These experiments used a geometry similar to the initial trial geometry.

Details are given in figure 6.11 with a photograph of the detector modules during commisioning. The fluidised bed had a diameter of 150 mm with active height above 330 mm. The modular detectors were arranged such that the bed was in the centre of the XZ plane and the field of view covered the lower portion of the active region.

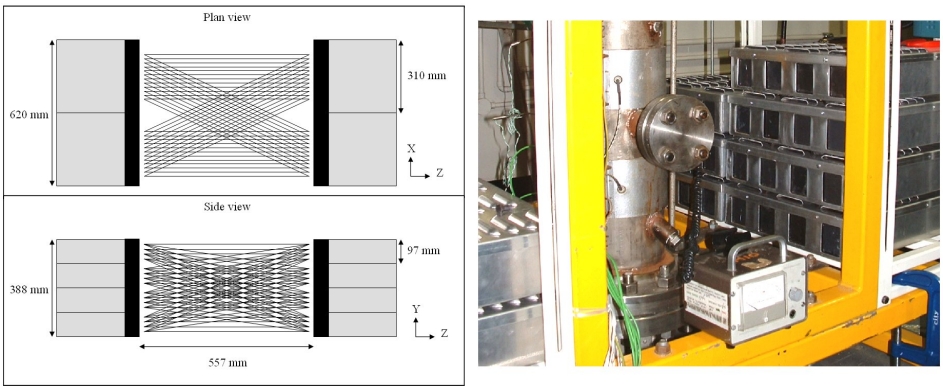

Figure 6.11: Modular camera geometry for fluidised bed study

The results from this study, summarised in section 7.2, show that the modular camera could be used to obtain meaningful and useful data for a system that it would not otherwise be possible to study. It demonstrated the first use of a modern positron camera for engineering studies outside of the Positron Imaging Centre since the MWPC, and demonstrated its superior properties.

At the loss of some spatial resolution (due to decreasing the angular sampling) the geometry was modified in order to study motion along the full height of the active region of the bed. In this geometry the modules were stacked on top of each other with only one module on either side. This allowed the height to be doubled, but resulted in loss of precision as there were only 4 blocks along the X-axis. The geometry for this is shown in figure 6.12.

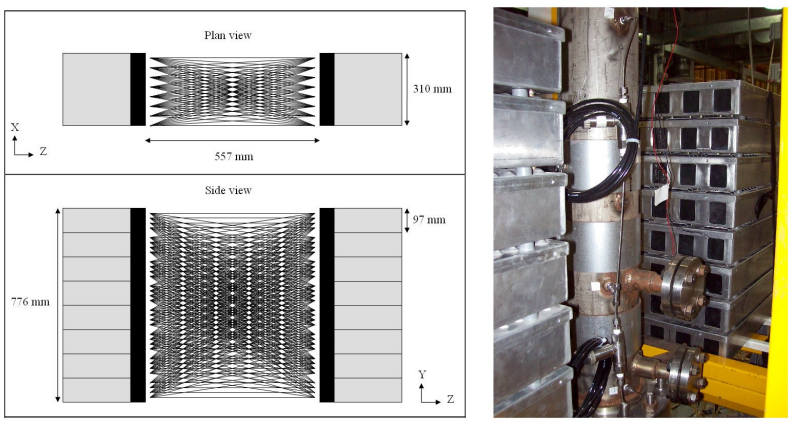

Figure 6.12: Second modular camera geometry for fluidised bed study

After the encouraging success of the experiments in the Chemical Engineering Department, the camera was used for a study of real industrial plant operating in situ. Figure 6.13 shows the large scale geometry that was used for experiments performed between February and June 2006 at the BP chemicals Ltd. (BP) site at Saltend in Hull, UK [27].

A 750 mm diameter pressurised fluidised bed pilot plant reactor was studied using the modular camera. In this experiment it was not necessary to have a field of view covering the whole apparatus; however it was important that the tracer was located when it passed over the top, or underneath, a dividing baffle placed inside the reactor.

The bed circulation time was measured for different flow regimes and bed turn over motion observed qualitatively. Further details are given in section 7.3.

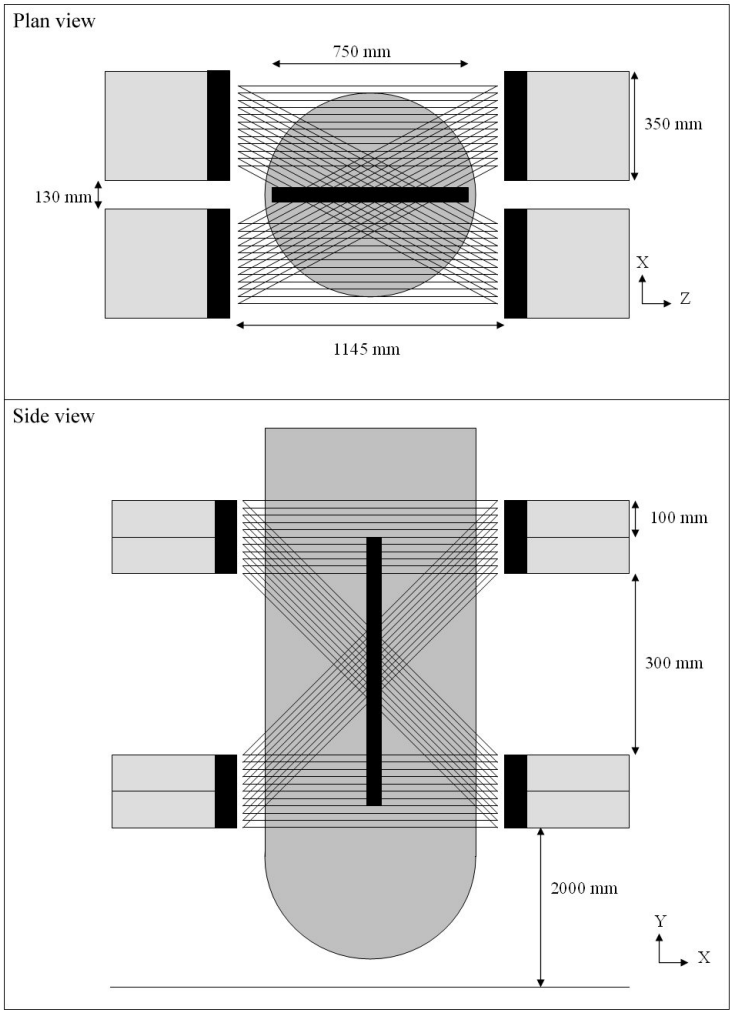

Figure 6.13: Large scale geometry

In figure 6.14 the cylindrical geometry of the original design was kept, resulting in a uniform field of view within the detector ring. The diameter of the ring was reduced from the original by a factor of one half, the new radius being 255 mm: a reduction from 510 mm.

This was done by removing the odd numbered detector blocks from the original ring, thereby reducing the number of detectors in the scanner system by one half. As each detector was setup to operate in coincidence with the seven opposite detectors in the original ring, the new system had the number of coincident pairs reduced to three detectors.

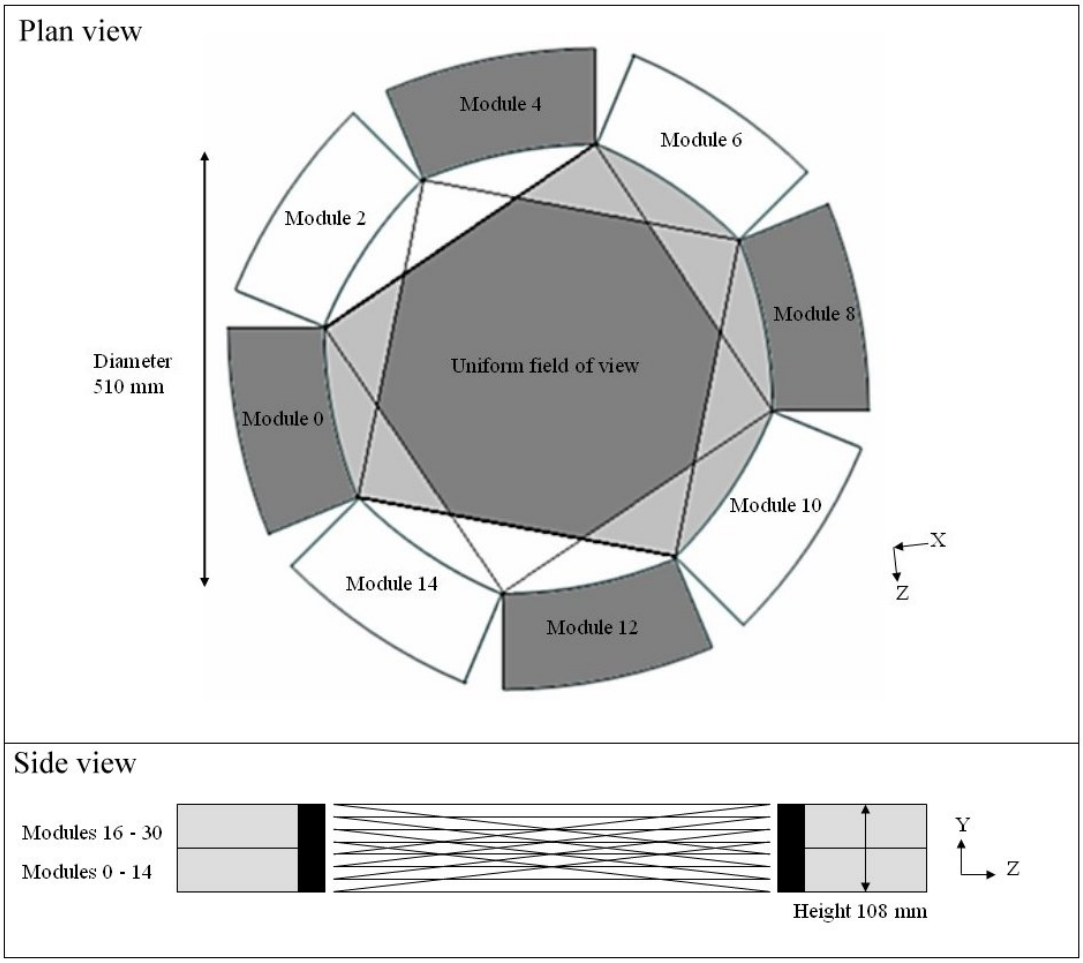

Figure 6.14: Small Scale Ring Geometry

As the radius of the ring is reduced the segment location as given for the original scanner (section 3.3.1) is no longer valid. Furthermore due to the arc of the detector buckets a correction needs to be derived to accurately locate each crystal segment. The diagram in figure 6.15 illustrates the new segment location concept.

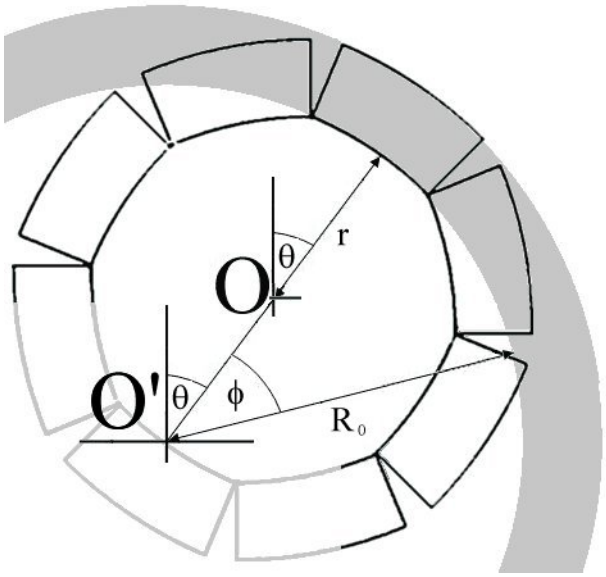

Figure 6.15: Segment location for new ring setup

The original arc of the bucket is 'focussed' upon the origin O'. We consider the location of the new origin, O, in Cartesian coordinates. Once again the bucket numbers give the angle around the scanner ring. For an original gantry radius of R_0 (510 mm), and a new radius r (255 mm), the centre of the reduced ring is given as:

$$O = [-(R_0 - r)\cos\theta, -(R_0 - r)\sin\theta] \tag{6.1}$$

As before the angle θ is given by the bucket number multiplied by the factor 2π / 16, and the angle ϕ is given by (Segment - 15.5).2π / 512. The position of the crystal segment is then given in terms of the original radius, R_0 and an angle $(\theta + \phi)$. This gives the position relative to the original origin, O' in Cartesian coordinates as:

$$X = R_0\cos(\theta + \phi) \tag{6.2}$$

$$Z = R_0\sin(\theta + \phi) \tag{6.3}$$

If we then take account of the new origin, O, the new coordinates of an arbitrary crystal segment is given by the Cartesian relations:

$$X = R_0\cos(\theta + \phi) - (R_0 - r)\cos\theta \tag{6.4}$$

$$Z = R_0\sin(\theta + \phi) - (R_0 - r)\sin\theta \tag{6.5}$$

Figure 6.16 shows the geometry that was used to image a PEPT tracer introduced into a metal cast in the metallurgy department at the University in 2007.

In previous experiments the location precision has been dominated by errors in the Z direction because of the low number of detection elements along the axis. In this geometry however, for the first time the detectors were arranged with equal numbers of detectors along each orthogonal axis, resulting in the errors in location along each axis being of a similar order of magnitude.

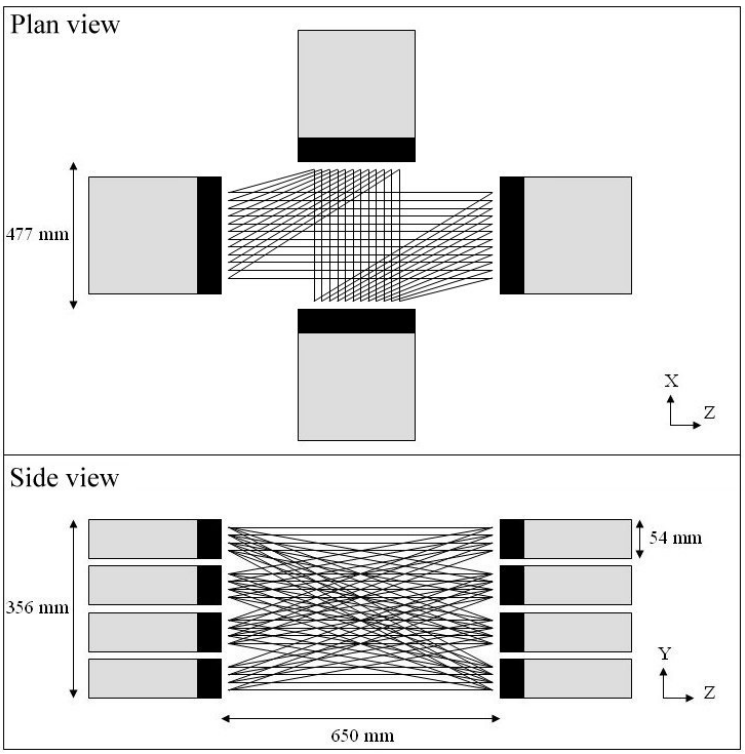

Figure 6.16: Orthogonal Geometry

As can be seen from the lines of response in the diagram, due to the restricted module pair combinations not all of the modules operate in coincidence. This means that there are some areas of reduced sensitivity, however the experiment was arranged to minimise this effect and produced data summarised in section 7.4.

The problem with the restricted module pair combinations is discussed in detail in section 6.5

Having a small geometry is advantageous because the sensitivity is greatly increased with smaller systems due to solid angle considerations. The amount of scattered gamma rays is also reduced with a small system because there is less scattering material present in the field of view.

The geometry shown in figure 6.17 was used for the study of a horizontal twin screw extruder operating in the Interdisciplinary Research Centre (IRC) on the University campus in 2008. In this 8 buckets from the 951 scanner were arranged in a tightly packed geometry around the equipment. As this is essentially another reduction in radius from the original system the field of view is relatively uniform in the XZ plane and reaches a maximum in the centre along the Y axis.

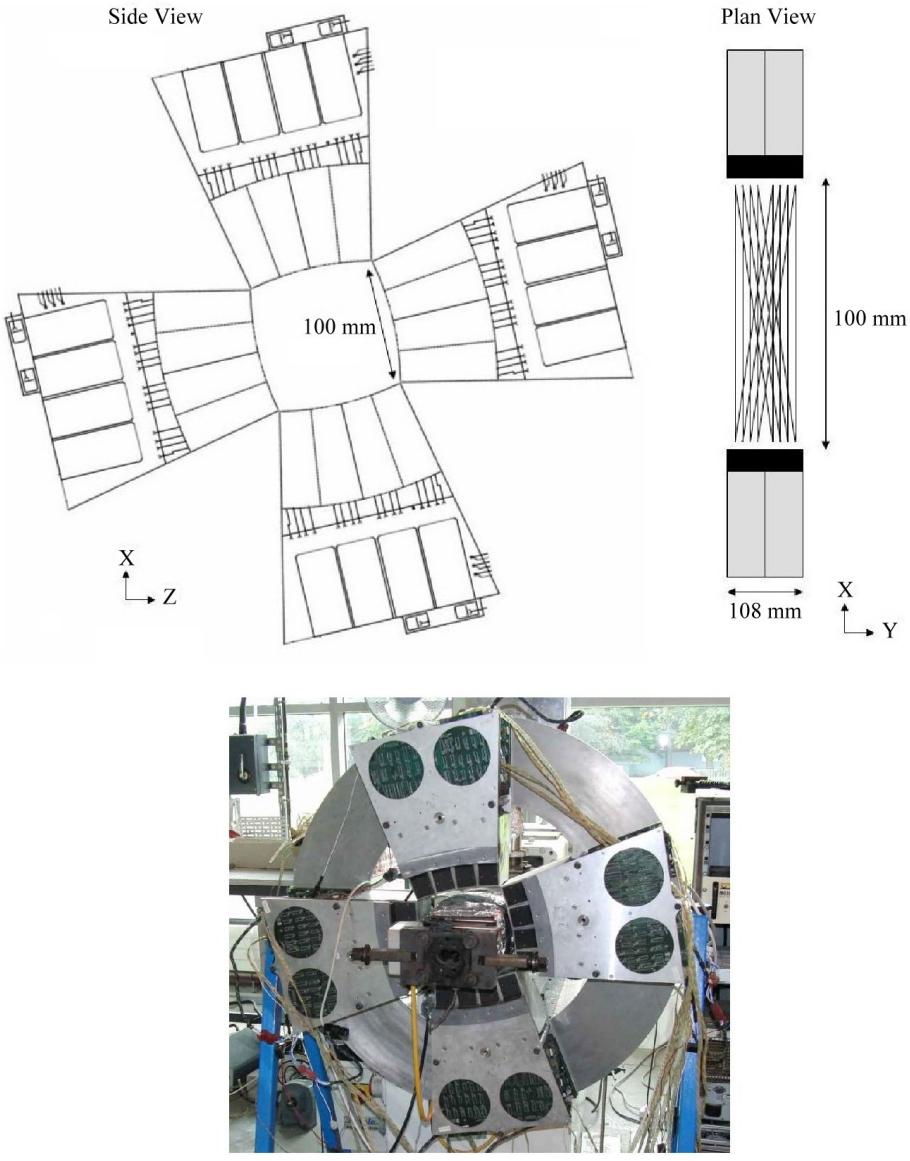

Figure 6.17: Small Scale Ring Geometry

This geometry is modelled in software by initially considering it as a cubic geometry. Knowing the original radius of the gantry gives an expression for the curvature of the arc of the detector blocks. The arc correction is applied by subtraction or addition of the arc correction (sagitta) to the detector module in question. Finally the whole coordinate system is rotated using a standard rotation matrix to correct for the angle the block faces make with the horizontal: this angle is designed to be 12.5 °.

This is illustrated in figure 6.18, the actual detector geometry is shown in grey, with the cubic model overlaid in black. The perpendicular distance from the cubic line to the arc of the detector block face is known as the sagitta, this is denoted S. The maximum value of the sagitta is found in the centre of the bucket, this can be found by considerations in the original ring circumference, divided into 16 equal segments (one for each bucket). The maximum value of the sagitta is then calculated to be 9.79 mm for the 951 series buckets. A sine function is then used to describe the path along the arc depending upon the crystal location.

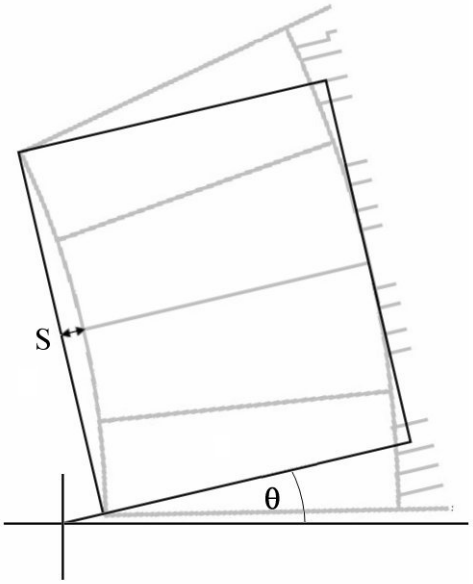

Figure 6.18: Modelled geometry

Finally the model is rotated around the Y-axis using the rotation matrix, R:

$$R = \begin{bmatrix} \cos\theta & -\sin\theta \\ \cos\theta & \sin\theta \end{bmatrix} \tag{6.6}$$

This results in expressions giving the X and Z transforms: X' and Z' for original coordinates X and Z respectively:

$$X' = X\cos(\theta) - Z\sin(\theta) \tag{6.7}$$

$$Z' = X\sin(\theta) + Z\cos(\theta) \tag{6.8}$$

6.4 Challenges

6.4.1 Module Alignment

One of the biggest practical problems with the modular camera is accurately positioning the detector modules around the system under study. Measuring the relative positions of the modules is non-trivial, especially when the apparatus to be studied is large and bulky. Errors in positioning and measurement of the detector locations can be highly detrimental to the precision of PEPT locations.

This was studied in detail for the camera geometry used for the extruder experiments outlined above, and described in section 7.5, during the calibration before the experimental runs. One of the modules (bucket 0) was offset from its measured position and approximately 12k locations were performed on a stationary source placed in the centre of the field of view.

Figure 6.19 shows how the PEPT location deviates from the true location (measured using a correctly defined geometry) as the detector module is offset from its measured location along the X axis. Algorithm variables N of 500 and the measured optimum f of 30 % were used in this study.

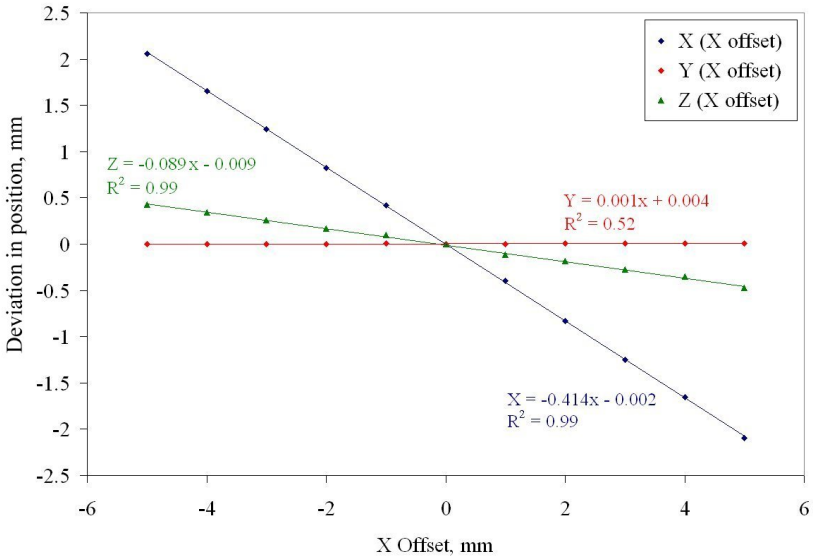

Figure 6.19: Deviation in position for a stationary source for module offset along X

As expected, misplacement along a single axis affects the locations along that axis the most, however this does affect the location accuracy along the other axes to a lesser extent. Linear best fit lines are drawn between the data points to illustrate that all location values are affected, and to what extent.

Similar data are shown in figure 6.20 for the detector module moved along the Y axis; it did not seem necessary to include Z offset data as this is equivalent to the X axis in this geometry.

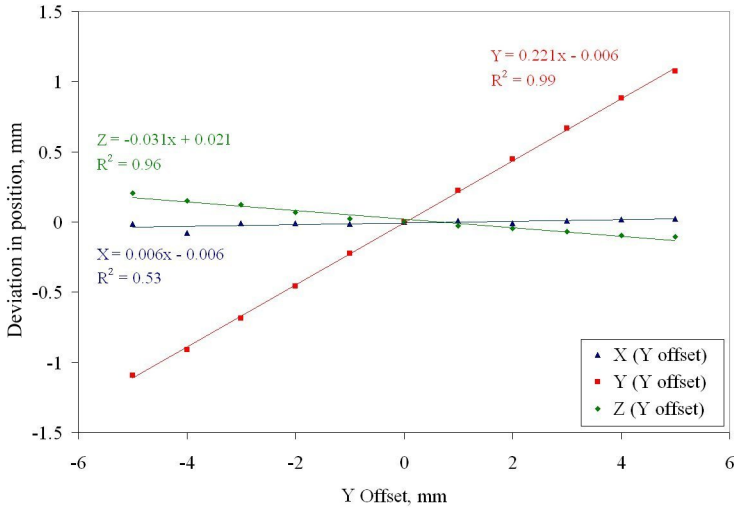

Figure 6.20: Deviation in position for a stationary source for module offset along Y

The average PEPT error for the locations made on the stationary source is shown in figure 6.21 for the same two cases as module 0 is offset along X and Y axes.

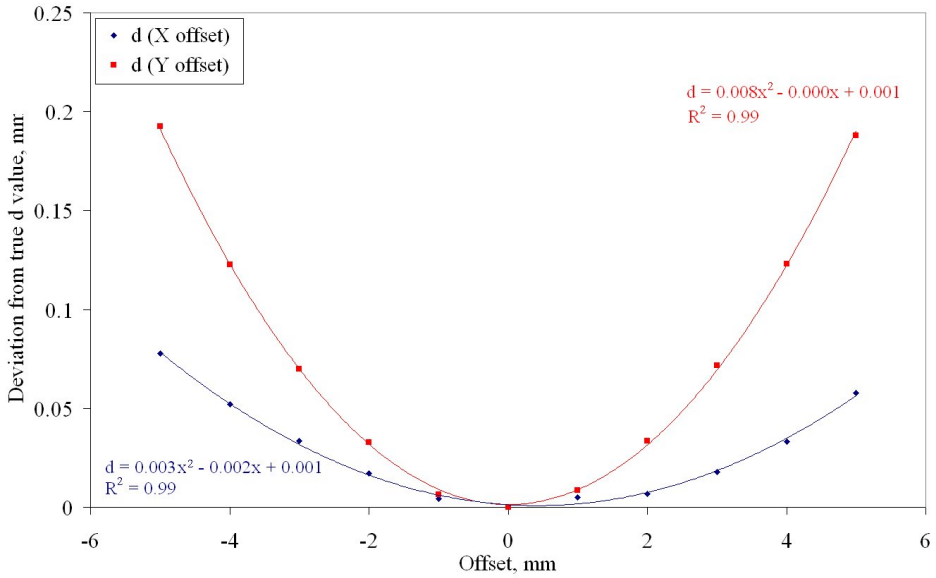

Figure 6.21: Error for stationary source as module is offset from true position

Polynomial fit lines of second order are shown between the measured data points to illustrate the extent of the effect for mis-location along each axis. That offsets along the Y axis are seen have a greater effect then the same offsets along the X axis can be explained due to the amount of sampling along each axis.

In the Y direction there are 2 detector modules resulting in a total of 16 crystal elements along the axis. Of these 16, when bucket 0 is moved, this offsets 8 crystal elements from their true locations. As the source is central and stationary, the response of all of the crystals is approximately uniform so the affect of this linear offset is large.

In the X direction each detector module presents 32 crystal elements along the axis, but the X and Z locations for a central source are determined by the coincidence response from all modules. Thus a smaller proportion of the data describing the particle location in X and Z is perturbed by the offset of the module.

Surprisingly the standard deviations in the measured position for the stationary source remain relatively constant when the detector module is offset from its true position. This is plotted in figure 6.22. The detector module is offset along the X axis, with negative values placing the detectors closer together and positive values separating the detectors.

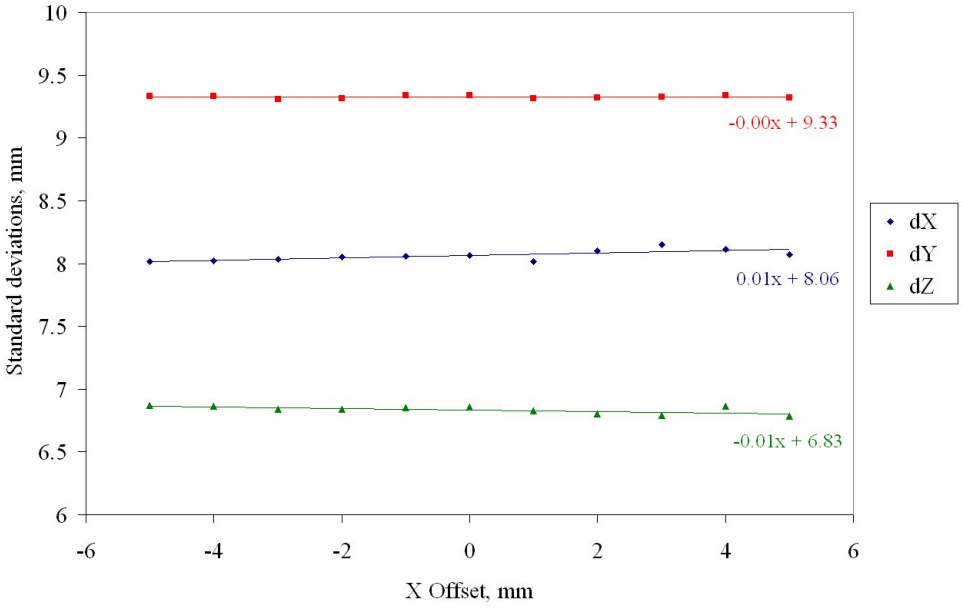

Figure 6.22: Standard deviations for stationary source as module is offset

The standard deviation is seen to increase slightly as the detector offset is increased, this is due to the detector separation increasing, meaning the angles each LOR makes with the normal to the plane of the detector are effectively increased, and therefore resulting in a larger error in the triangulation.

In the Y direction the standard deviation remains at a constant value within the quoted precision; this shows that the offset in the X direction has negligible effect on the precision in the Y direction. In the Z direction however the standard deviation is seen to decrease slightly as the module offset is increased.

Extrapolation results in the standard deviation tending to zero for locations in the Z direction as the detector module is offset from its location. This effect, whilst linear in this region, breaks down at large values of detector module offset as the algorithm eventually rejects all of the data from the offset detector module; compared to data arriving from correctly positioned modules this data are regarded as highly corrupt.

The effect was also tested with a dynamic system: in this a source was placed on a turntable rotating with components in all three orthogonal axes. The measured trajectory for the source is shown in figure 6.23.

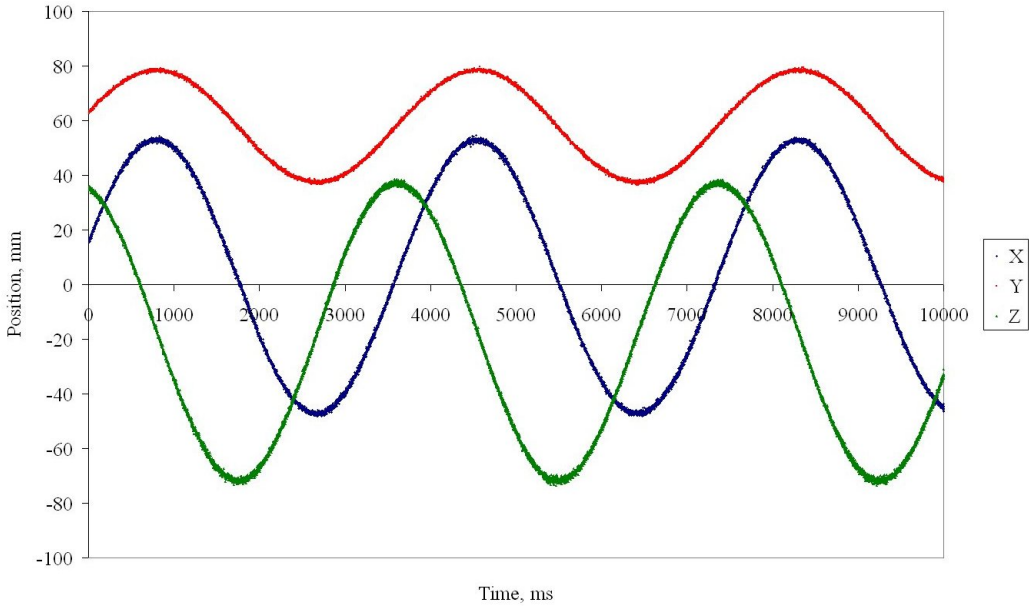

Figure 6.23: Trajectory for test source in field of view

140k events were recorded for the turntable setup and analysed using the PEPT algorithm with variables f of 30 % and N of 500 events per location. A single module was positioned in software with different offsets along the X and Y axes and the data reanalysed.

The PEPT algorithm calculates the average perpendicular offset of all the trajectories used in each location from the minimum distance point. It is then possible to quantify the relative precision of the location using this figure (denoted d, described in section 2.2).

The detector module is offset along the X axis, with negative values placing the detectors closer together and positive values separating the detectors. The average PEPT error ($\bar{d}$) for all the locations produced in this experiment is plotted in figure 6.24, showing that the error is minimised when the module is correctly positioned, and increases as the module location is further offset.

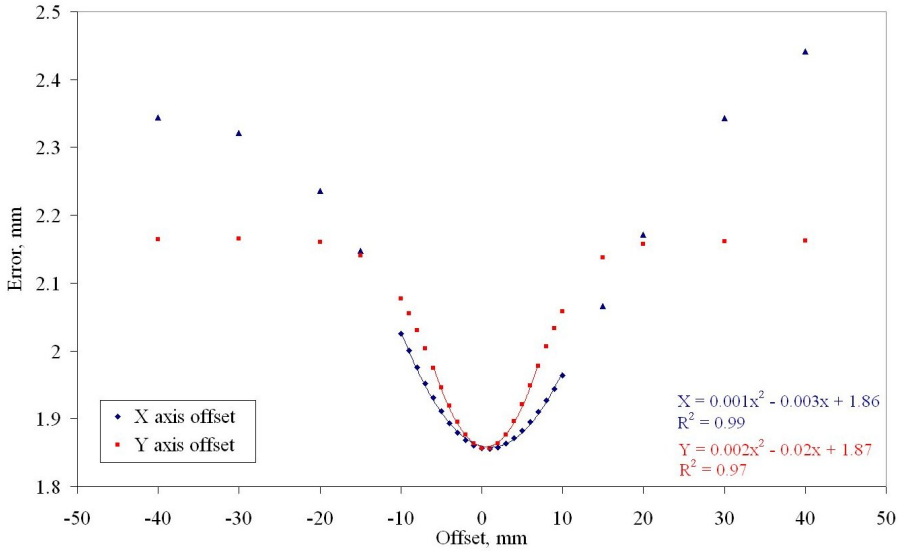

Figure 6.24: Average PEPT error as modules are offset from their true positions

As the module offset is increased this error increases; at this stage the PEPT algorithm is rejecting data from the mal-aligned module but some data remains in each set, increasing the error. As the module offset increases still further, the data from the module is rejected leaving a smaller set of events with which to describe the real location; thus the statistical error is increased and the location error tends to a constant value.

It is encouraging to see that the PEPT error is minimised when the detector module is positioned correctly; this implies that the camera geometry is modelled correctly. The use of sources fixed in known positions within the field of view is suggested as a method to locate the detectors. This idea would use a comparison between the recorded data and the known source location in an iterative approach to locate the detector. In this manner the detector system could be used to effectively locate itself.

This work needs much more development to become viable; it has been seen that the PEPT locations are very sensitive to bucket positioning and with the large number of degrees of freedom for the positioning of each detector block the problem is highly complex.

For example each detector module can be translated and rotated along 3 axes giving 6 degrees of freedom in location. For a detector system using 32 detector modules the number of degrees of freedom is then 296. If detector blocks are positioned independently of each other this is increased by a further factor of 4.

6.4.2 Gaps Between Detector Modules

As discussed, there is no requirement for a uniform field of view when performing a PEPT experiment. This allows the modular camera geometry to be highly flexible for PEPT studies but results in regions of higher and lower sensitivity. It has been seen however, that having gaps between the detector blocks can cause a significant decrease in performance despite the given arguments allowing a non uniform field of view.

In a PET system these gaps can be accounted for as a reduction in the angular sampling. However in the PEPT algorithm where a uniform field of view is not required this should not cause significant effect to the location accuracy.

If we consider a purely random source of radiation exciting the detectors equally across the system, there will be an equal density of LORs throughout the field of view. When the PEPT algorithm is used to define the minimum distance point for this set of LORs, the result is a location in the centre of the volume as this is the average midpoint for all of the LORs used in each location.

If we apply this argument to a field of view where there are regions of higher and lower sensitivity (as the detector blocks are spaced apart), we see a similar effect. There are regions of higher and lower sampling across the volume which tends to bias the location towards the region of heavier sampling.

6.4.3 Discrete Crystal Elements

Even without gaps between detector blocks there is still a problem with the discrete nature of the spatial sampling. Each crystal element has a defined spatial distribution (951 detector crystals are 6.25×6.75 mm^2 area by 30 mm depth).

An interaction within one of these crystal elements can occur anywhere in the crystal, but is always described as being in the centre of the crystal by the event data.

There is heavy sampling at a discrete location, surrounded by regions of zero sampling. This is acceptable in a PET image as the sinogram represents all possible combinations of these discrete points, however in PEPT this seems to cause a similar problem as with gaps between the detector modules albeit on a much smaller scale.

Figure 6.25 shows an example of this: here 900 LORs from a point like source within the XZ plane are drawn between a number of different blocks. In the left figure the LORs are seen to all originate from the same position within the crystal, resulting in regions with no sampling between these lines. The figure on the right shows the same case with each LOR arising from a random position within the crystal. This results in much more uniform sampling across the whole volume.

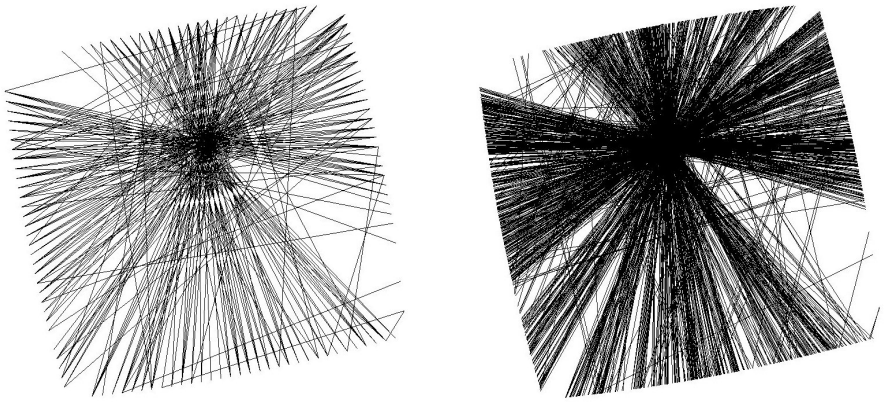

Figure 6.25: LORs from discrete and continuous detection elements

This is illustrated further in figure 6.26 which shows a zoomed in version of the LOR combinations between two opposing detector modules as shown in figure 6.10. With each crystal positioned at a discrete point, the field of view varies on a microscopic scale due to the discrete nature of the LOR description. Physically however this is not the case as each line of response is strictly a tube, thus the field of view is continuous and constant across each crystal.

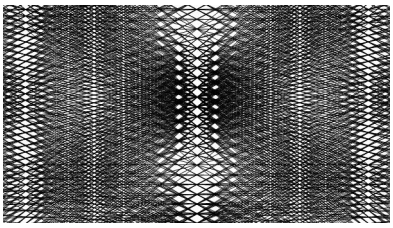

Figure 6.26: Magnified image of LOR combinations between two opposing modules

The approach of smearing the distribution across each crystal using a random number has been investigated for use in the PEPT algorithm in order to make the spatial sampling more continuous.

A one dimensional argument can be used to justify the investigation: in moving the interaction site to a random position in the crystal the assumed position can be either closer to the actual interaction site, or further away. Over a large sample on the average half of the events are moved closer (improving the data set) and the other half are moved further away (degrading the data set).

Due to the spatial sampling arguments given above it was thought that the PEPT algorithm would perform more efficiently with continuous sampling, provided by this random correction. Unfortunately this has been proven to have a negative effect upon the precision in making a PEPT location. In figure 6.27 the average PEPT error for a stationary and a rotating source is shown with varying orders of random correction. The geometry used was the IRC geometry (section 7.5) and a source giving a prompt coincidence rate of around 500 kHz was used.

In this, and the following plots, the random correction is varied from zero (each event located in the centre of the crystal) up to a continuous distribution extending to 150% of the crystal size. At 100% crystal size the random distribution covers the entire crystal, with 50% of event locations spread below the midpoint of the crystal and 50% spread above the midpoint.

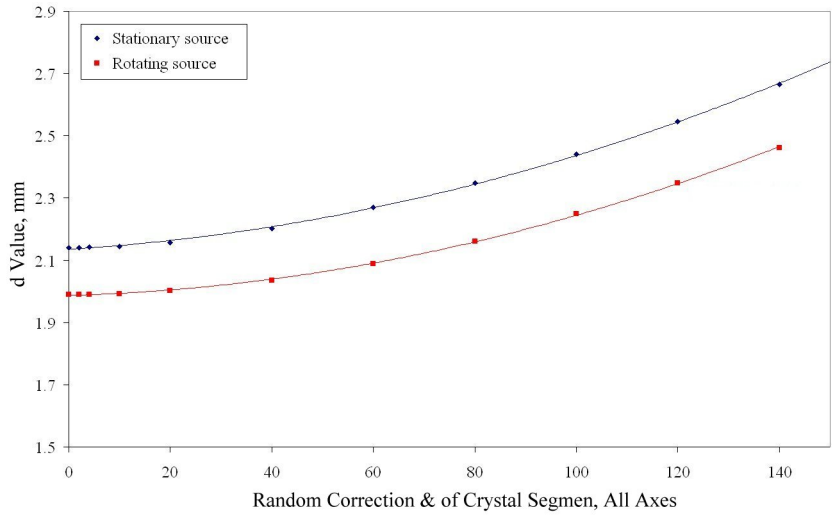

Figure 6.27: Average PEPT error for random crystal correction

Both sources exhibit similar behaviour due to the random correction, with the PEPT error increasing with increasing random smearing. The deviation between the two curves is due to the relative size of the sources used.

The stationary source used was a 4 mm diameter bead containing ^{22}Na, the rotating source was a 100 μm diameter particle loaded with ^{18}F, count rates for the rotating source were around 700 kHz, and for the stationary source 350 kHz.

The measured optimum value of 30 % was used for the algorithm variable f, with an N value of 2000 events per location, data was acquired using the new acquisition system.

Figure 6.28 shows the deviations of the average location calculated for the stationary source with different random distributions from the calculated value without any correction. Again the random distribution is increased smearing the events across the crystal element, in this case the smearing only occurred along the Y axis. This shows that the correction affects the value of the calculated location but assumes the zero correction value is the true location. From the PEPT error results above this is a valid assumption.

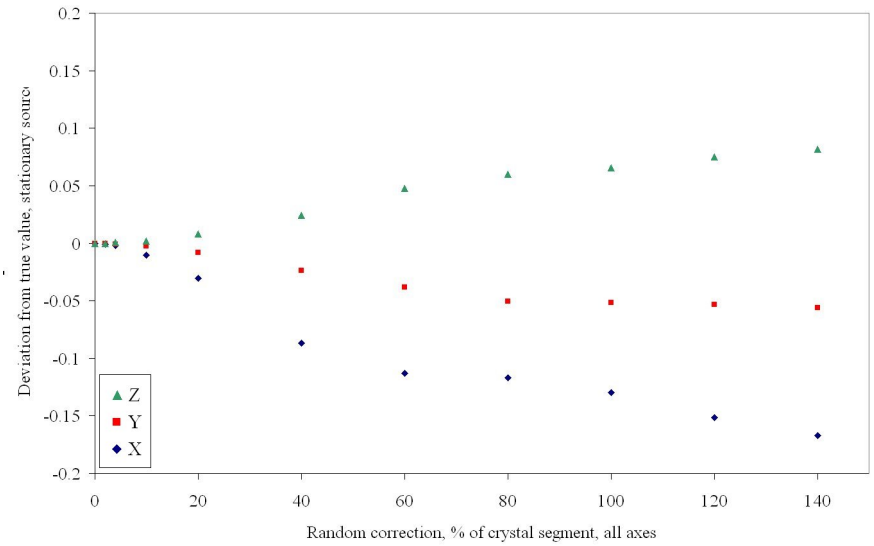

Figure 6.28: Deviation for stationary source for random distribution

The same experiment was performed for the rotating source with similar results, these are shown in figure 6.29.

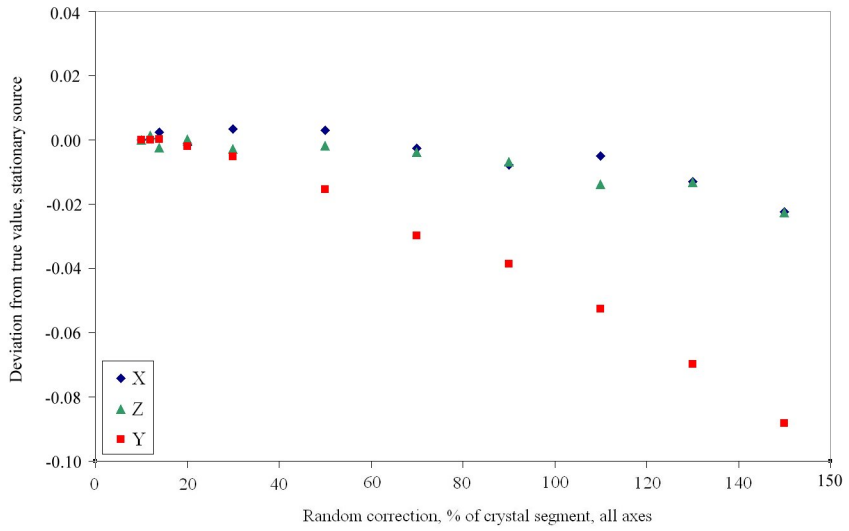

Figure 6.29: Deviation for rotating source for random distribution

6.4.4 Source Outside the Field of View

Due to the background of purely random coincidences (from random and scattered / associated events) the camera system will still produce event data when the tracer particle is outside the field of view (FOV). The data that is produced is randomly orientated to the first order (this will be explained further in the following section) and is processed by the PEPT algorithm to provide locations over the number of events selected. The solution generally has a large deviation between the trajectories and therefore a large PEPT error (d) due to this random nature.

With a purely random source the locations provided by the PEPT algorithm tend towards the average of the spatial distribution (i.e. the centre of the FOV). It is considered that this is the case when the source is outside of the FOV of the detectors, then the PEPT algorithm outputs these average locations with a large standard deviation between the LORs used in the location. It is therefore possible to identify locations which are calculated for a source outside the FOV as these locations have a larger associated error.

As the tracer approaches the field of view from a large distance, the exposure of the detectors is no longer purely random, there seems to be some form of correlation between locations provided by the PEPT algorithm and the tracer position even when outside of the FOV. This is shown in figure 6.30, where a tracer is rotating in the XZ plane with radius 10 mm, centred on the origin, and approaches the detectors along the Y axis from -1000 mm through the detector system to +25 mm.

The detectors are setup in the IRC geometry detailed in figure 6.17. In this the X and Z axes range from -104 mm to 104 mm, the Y axis is two detector blocks high so ranges from 0 mm to 108 mm. The middle of the field of view then has the coordinates (0, 0, 54) in the X, Y and Z directions respectively.

The plot shows the PEPT locations for both the Y and Z directions and a rolling average of these data points plotted as a line to aid clarity. This data can be divided into three separate regions.

In the first, on the left (section A), the source is far enough away from the detectors that the exposure is approximately uniform, and the PEPT algorithm produces locations near the centre of the field of view. Average Y for this section is 43.4 mm which is off-centre towards the position of the tracer (centre is 54 mm), this is explained as the detectors nearest the tracers receive a higher exposure then the detectors further away due to geometric and absorption arguments. The average Z for this section is -1.2 mm, the actual centre of the rig is around 0 mm.

In the second section (middle, section B) the source is close enough to the detectors to bias the PEPT locations towards the detectors closest to the source. The PEPT locations in the Y direction are now heavily biassed towards the source with average location being 25.6 mm, this is approximately the mid point in the first set of detectors (27 mm).

As the source is outside the field of view, the data are still composed entirely of random coincidences although some correlation between tracer motion and detector response is evident in the Z direction. The PEPT locations in the Z direction exhibit a sinusoid with period equal to 1200 ms, which is approximately equal to the actual period of rotation. The average position along the Z axis is 0.81 mm, showing that the rotation is approximately centred on the origin. The amplitude of the sinusoid fluctuates but is generally 1.5 to 2 times the amplitude of the real system.

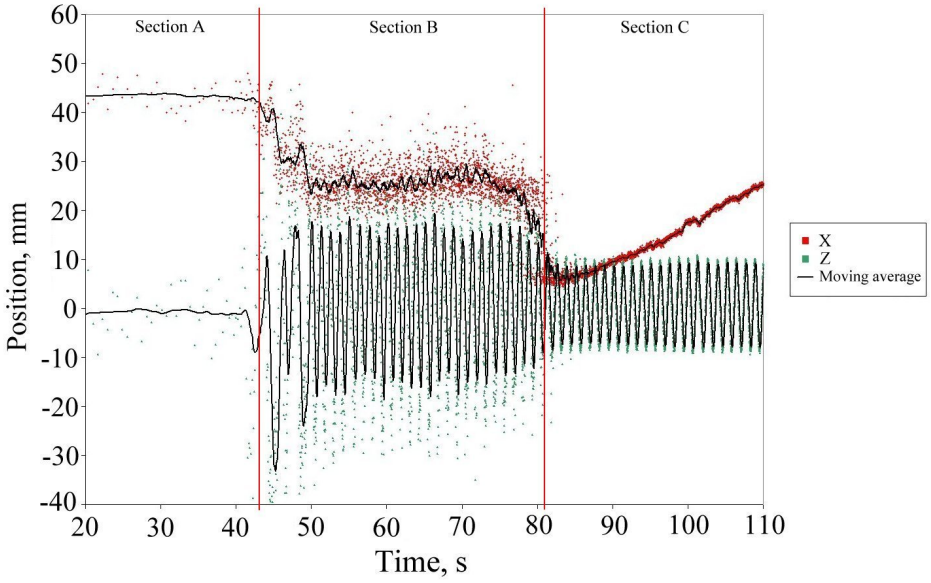

Figure 6.30: PEPT location data as tracer is moved into the field of view

The third part of the plot (right, section C) sees the source entering the field of view. In the Y direction the PEPT locations approach zero as the tracer enters the field of view. However they never reach zero as the sensitivity of the camera for a tracer particle positioned at the origin is zero (section 6.2). As soon as the tracer enters a region where there is some sensitivity the genuine data rate is much higher than the random rate and therefore the PEPT location produces accurate results, the actual trajectory of the tracer particle is described in this section.

This data, along with the countrate data, is given in figure 5.8.

The three regions can be clearly identified and corrupt PEPT data rejected by consideration of the relative error for each location. Refer to section 2.2 for a discussion about how this error is calculated. The PEPT error for each location, given the nomenclature d, measures the standard deviation in the perpendicular distance from each LOR used in calculating the location to the measured position. For well defined locations this value is small, for LORs with a large spatial distribution this value increases. Plotted in figure 6.31 is the value of d for the data set given above.

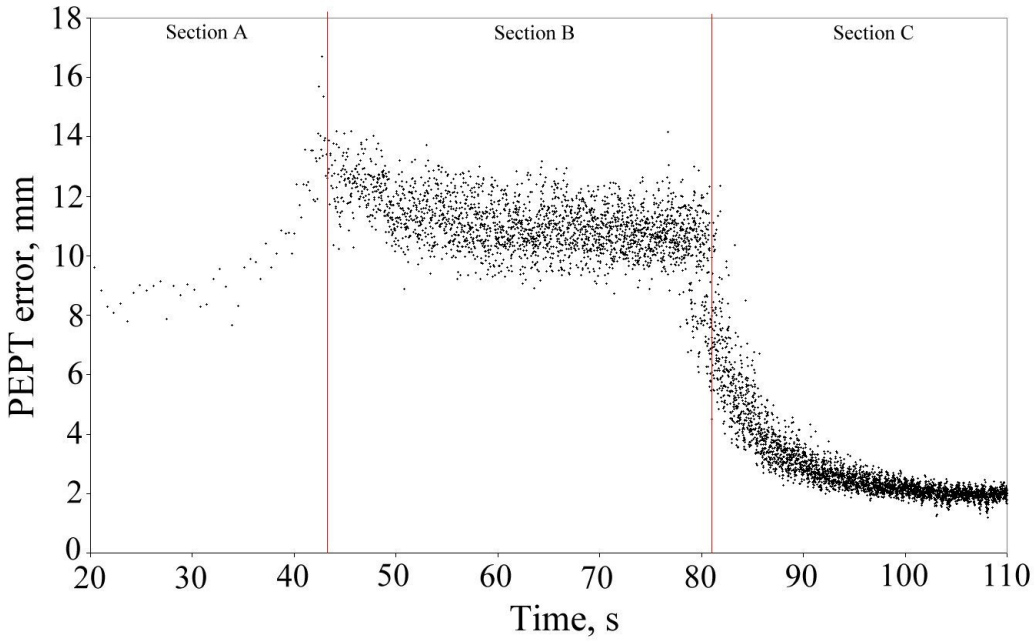

Figure 6.31: Location error as the tracer is moved into the field of view

Whilst the tracer is a large distance away from the detectors the d value remains around 9 mm on average. This is probably due to the fact that there will not be many coincidence data and the furthest of these from the centre are rejected.

As the tracer gets closer the d value actually rises, this is due to the higher data rate and the fact that the movement of the source causes some correlation with the PEPT locations; the source is then measured to be off centre, but the random lines of response still have a high component of scatter around the centre.

As the source enters the field of view the value of d drops drastically and remains approximately constant whilst the tracer is within the field of view. By rejecting data which has d value larger then a fixed constant (in this case around 3 mm) genuine PEPT data can be identified.

Figure 6.32 shows the total count rate data and the Y position for the same data set shown in figure 6.30. This data shows that the count rate increases dramatically as the tracer enters the field of view of the camera. It is encouraging to see that the count rate due to random coincidences (whilst the tracer is outside the FOV) is relatively low whilst the tracer is outside the active region of the detector; this implies that the coincidence window is correctly setup.

At this point the coincidence rate increases at approximately the same rate as the position along the Y axis. This is because the sensitivity is increasing linearly with the tracer position as the tracer enters the field of view.

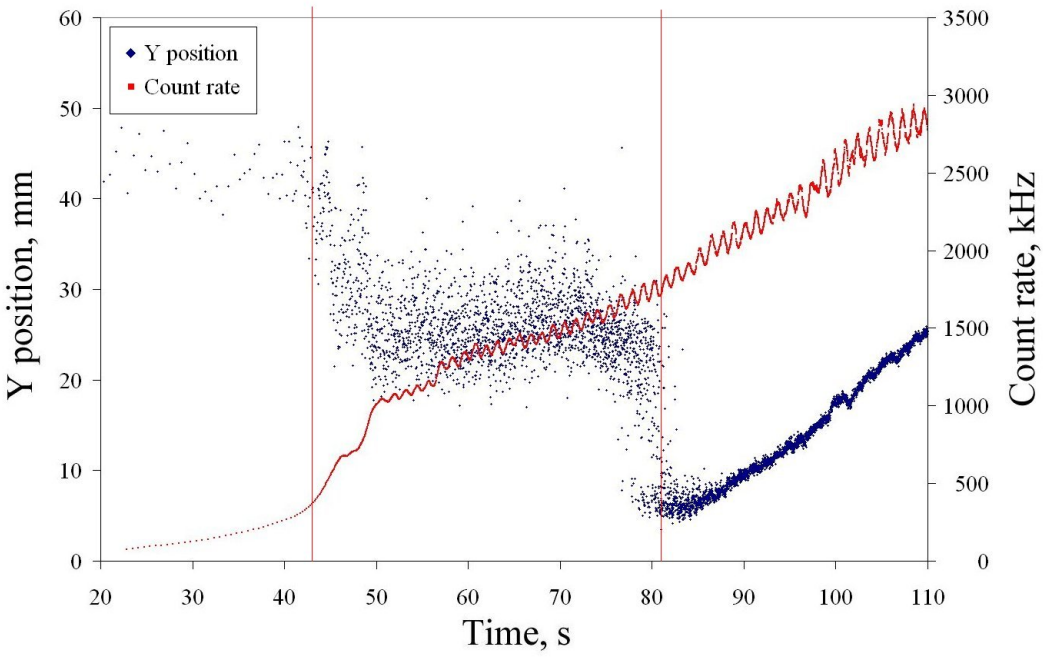

Figure 6.32: Coincident countrate as tracer is moved into the field of view

6.5 Coincidence Combinations

The need to keep the coincident combinations opposite each other can cause some problems when changing the geometry for a specific application. Note in figure 6.16 that not all diagonal combinations are possible; resulting in a varying sensitivity across the field of view.

This effect becomes more apparent as the number of buckets (and hence the size of the field of view) is reduced. With such a small system it is not possible to arrange the module pairs in a manner such that each detector operates in coincidence with every other detector in the system, without changing the coincidence logic.

For example, in the IRC geometry (figure 6.17 and 6.33) bucket zero is in coincidence with buckets five and eight, but not with bucket thirteen. Conversely if the bucket numbers are arranged such that coincidence combinations for bucket zero (in this example) cover the whole system: possibly bucket zero with buckets five, eight, and eleven, other coincident combinations in the ring are lost.

In this example bucket eight is then only in coincidence with bucket zero. This would result in the field of view being heavily non-uniform with much higher sensitivity around bucket zero, and low sensitivity especially off-centre near bucket eight.

With the modular camera the ability to use more then the maximum 56 bucket combinations would be highly beneficial. Consider for example the setup given in figure 6.33 showing the LOR combinations for the original buckets on the left and same diagram on the right if further bucket combinations were allowed.

In this example, where such a small number of detectors are used, it is impossible to arrange the coincident combinations in such a manner that all of the detectors in the system are operating in coincidence (figure 6.33 left).

It is clear that a system with all of the coincident combinations allowed offers a more uniform field of view and offers greater sensitivity provided the tracer particle remains within the detector volume (as shown in figure 6.33 right).

The example shown illustrates a similar case to the setup used for the IRC extruder experiments outlined above (and in section 7.5). Direct coincidences are not shown in the diagram to aid clarity.

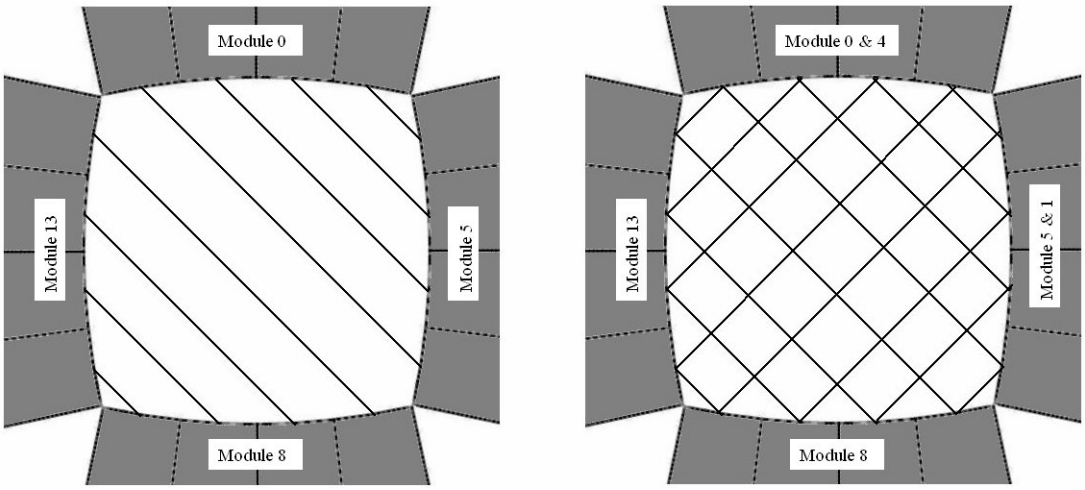

Figure 6.33: LORs for different bucket combinations

Without the limitation on the maximum number of bucket pairs there are a possible 120 bucket combinations (((16 - 1) * 16) / 2) in a 16 bucket system.

In order to process the data from the extra modules more coincidence processors would be needed and a wider output bandwidth would be required. These further combinations would not present any useful data in the original medical configuration as the field of view for this pairing is outside the patient port, in fact these events would contribute to the random signal due to this geometrical argument.

It was originally proposed that the connections on the VME 64 backplane which contains the IPCP boards be modified in order that extra coincidence pairs can be accepted. Consider the circuit philosophy for the Image Plane Coincidence Processor described in section 3.3.

The coincidence processor receives events simultaneously and in parallel from each of up to 16 detectors. The detector bucket numbers are assumed from the input channel on the Ring Receivers and events are directly routed to the appropriate processing module. Each coincidence processor is configured to determine events between 56 possible bucket combinations (i.e. opposing bucket elements within the original gantry).

It became obvious that modifying the VME backplane to the IPCP would be a complex task and had the risk of making the VME crate unusable. Instead it was more practical to copy the event data arriving from each bucket to different ring receiver inputs.

This allows the new bucket combinations to be identified in the data stream as they would have the module pair numbers associated with the additional inputs. This does require however that there are spare inputs on the ring receiver which offer beneficial pairing. In the case where all the ring receiver inputs are used, a second coincidence processor could be used to accept the copied data.

In the eight bucket system described above (the IRC device) two ring receivers and all four coincidence processors were used. This allowed combinations between opposing buckets and between diagonal buckets as illustrated in figure 6.33. These combinations cover both rings (i.e. IPCP 64 covers combinations in the zeroth ring, IPCP 67 covers combinations in the first ring, and the further two IPCP boards cover coincidence combinations between both rings). Using 8 buckets meant that there were 24 spare inputs available for receiving detector data.

The incoming data are received in serial format over 4 twisted pair cables (section 3.2) and is controlled by the master clock; this master clock signal is provided on one of these twisted pairs in order to parse the data word. This means that both the event data and the timing signals required to parse the data correctly are all easily accessible.

In copying the data it was possible to piggy-back a second data cable from the original input, onto a connector going to the secondary input. This parallel connection allowed the same data to be presented at both the original and the secondary ring receiver inputs thereby allowing one detector to present data to multiple inputs.

Fortunately the fanout drive of the output devices on the buckets is of high enough current rating to drive both inputs on the ring receiver (separated by a significant length of cable) simultaneously. However attempts to drive a number of greater than two inputs failed; this is due to the high current needed and the added capacitance of the extra cabling. In the future, if a larger fanout of data is required, active line drivers will be needed to compensate for these effects.

Problems deriving the optimum bucket pairing are greatly increased with this system as further combinations need to be considered. It is vitally important that a bucket is not setup in such a manner that it can operate in coincidence with itself (via copied data).

It has been seen that if the copied data replicates a coincidence pair that already exists this does not adversely effect the data stream. For example bucket zero is in coincidence with bucket eight; if the data from bucket zero is copied to the input for bucket one (which is also in coincidence with bucket eight) the coincidence processor sees two valid coincidence events occurring at exactly the same time.

The coincidence processor flags both events as being a multiple, and chooses one of the events at random for further processing, eventually outputting a coincidence word describing an LOR between buckets eight and either bucket zero or bucket one. As buckets zero and one are actually the same device the output data represents the actual LOR. The multiple bit is set by the IPCP as it has indeed processed a multiple event, this can be accounted for when processing the data using the PEPT algorithm.

The first trial of this method was attempted for the small ring system used for the IRC experiments. It was determined that data from two buckets out of the four in each ring required copying in order to achieve all of the coincident pairs in the ring. Previously the number of coincident pairs was four; this can be increased to the maximum of six.

This was done by copying the data from bucket zero to the ring receiver input for bucket four, thus giving the pair four and thirteen (zero and thirteen). Bucket five data was copied to the ring receiver input for bucket one, thus giving the pair one and eight (five and eight).

The module numbers given in figure 6.33 (right) show this setup. These combinations gave only the extra pairing required and no further combinations in the system, thus the argument that multiple events would remain unchanged was not needed for this arrangement.

As each coincidence processor is functionally identical the same method could be used for buckets in the second ring. Thus, bucket sixteen was copied to input twenty (giving pairing between sixteen and twenty nine) and twenty one was copied to seventeen (giving pairing between twenty one and twenty four). This is equivalent to the lower order bucket numbers given above.

The symmetry of the system allowed the cross ring combinations for the copied data to be accepted, thus all possible combinations throughout the system were accepted. This was two extra from IPCP 64, two extra from IPCP 67 and four extra from the cross planes: two from IPCP 65 and two from IPCP 66.

The PEPT algorithm was modified to include data from the copied buckets; in this the physical location of the new buckets was given to be the same as their copied pair, data arriving from either of the pairs then gave the position of one end of the LOR. Further information and results obtained from this study is given in section 7.5.

Chapter 7

Experiments

The purpose of this chapter is to illustrate some of the applications and the variety of data that the modular camera has been used to acquire during the past 3 years. All of these examples have had significant input from the author, however these results are not deemed to be the main subject of this thesis. Where possible more detailed analysis is presented in the references given.

The first section describes the physical performance of a small camera system designed for a high sensitivity, high resolution PEPT study. Before the experiment was performed the camera system was characterised, this is described here.

The remaining chapter describes the applications which the modular camera has been used for, challenges and issues with certain experiments and results obtained from the study of various applications. The various modular camera geometries have been described in chapter 6 and are not repeated here.

7.1 Small Ring Camera Performance

As noted in sections 6.2 and 6.4, having gaps between the detector blocks causes a number of problems as the field of view is no longer uniform and therefore the angular sampling is irregular. In order to perform a high resolution experiment it is beneficial to keep the ring geometry of the original system, with its uniform field of view.

The geometry described in figure 6.14 was used for a high resolution study. By reducing the size of the field of view the sensitivity can be greatly increased. Before the system was used for any PEPT experiments its performance was measured for a number of different properties.

The precision in locating a tracer particle is proportional to the tracer activity (section 2.3), therefore the use of high activity tracer particles is preferred. However, at high activities the rate of random and corrupt data increases faster then the rate of true coincidences.

For a central point source of activity A, with a camera system subtending Ω solid angle with efficiency ϵ, the prompt singles rate in the camera is given by:

$$R_s = \frac{2A\Omega\epsilon}{4\pi} \tag{7.1}$$

The true coincidence rate for the system then depends on the singles rate in each detector and the probability that both gamma rays will be detected, this is given by:

$$R_c = \frac{2A\Omega\epsilon^2}{4\pi} = \epsilon R_s \tag{7.2}$$

The rate of random coincidences is a function of the count rates in individual detectors (R_1, R_2) operating in coincidence and the length of the coincidence gate (τ) used to resolve the event:

$$R_r = 2\tau R_1 R_2 \tag{7.3}$$

For a central point source the singles rates in each detector are approximately equal and can then be given the nomenclature R_s, the random count rate for a central point source is then given by:

$$R_r = 2\tau R_s^2 = \frac{2\tau}{\epsilon^2} R_c^2 \tag{7.4}$$

From these expressions it can be seen that the random rate increases with the square of the source activity, whilst the true coincidence rate increases linearly with activity.

It is then possible for the number of random events to completely overwhelm the true coincidence events with high source activity as the data rate from the coincidence processor has a fixed upper limit (theoretically 4 MHz per IPCP, practically 700 kHz (1 MHz theoretical) for the RTS unit).

It can also be seen from the above expressions that the factor $2\tau/\epsilon^2$ must be minimised to reduce the effects of random coincidences. This can be done by reducing the gate width or by increasing the efficiency of the detectors. This makes the BGO detector superior to NaI(Tl) scintillators for the same timing properties.

A carbon 11 source with an initial activity of 200 MBq was used to study the count rate performance of the camera system. This isotope has a half life of 1224 s (20 min) and the experiment was run for 4.5 hours after which the tracer activity had decayed to approximately 2 kBq.

In figure 7.1 the count rates for prompt, delayed and multiple coincidences are shown along with the total and true coincident count rates for the small ring camera setup using the RTS data acquisition system [37].

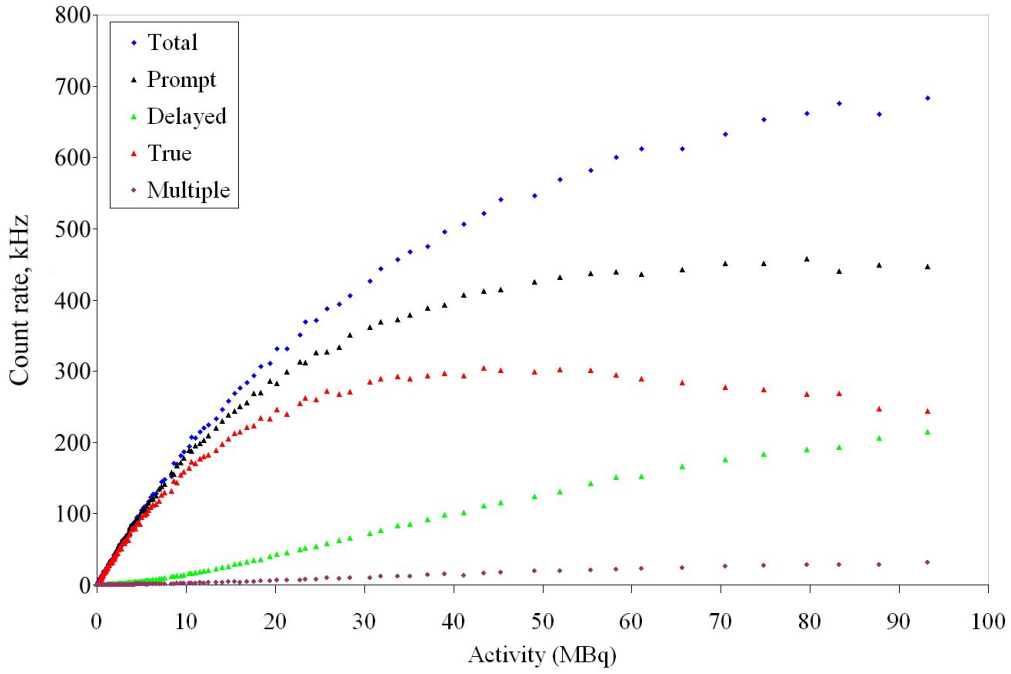

Figure 7.1: Count rates for small system

The prompt signal gives the total number of events occurring within the coincidence window following a trigger pulse; this includes genuine events and a background of random coincidences.

The delayed signal gives the number of events that occur in coincidence with the trigger event following a long time delay. This delay is long enough that the events cannot be associated with the trigger event; thus, the delayed count gives a measure of the number of random events occurring in the system.

The true event rate is given by the prompt count rate minus the delayed count rate. These concepts are discussed further in section 1.3.1.

It can be seen from the total count plot that at high activities: > 80 MBq, the detector system is saturated and cannot accept any further events, this occurs for a maximum event rate of 720 k s^{-1}. As the activity decreases the detector is no longer saturated, but the level of random events is high enough to interfere with the genuine data signal. As the activity decreases further the number of random events occurring also decreases since the randoms rate is proportional to the square of the singles rate. This allows the genuine data rate to increase as the coincidence processor is no longer overwhelmed with random events.

The maximum prompt coincidence rate was found to be 460 k s^{-1}; the activity of the tracer at this point was 53 MBq. The maximum true coincidence rate occurs with a tracer activity of 22 MBq; the data rate at this point is 305 k s^{-1}.

Figure 7.2 shows these countrate data, with some simple modelling used to describe the countrate performance of the system.

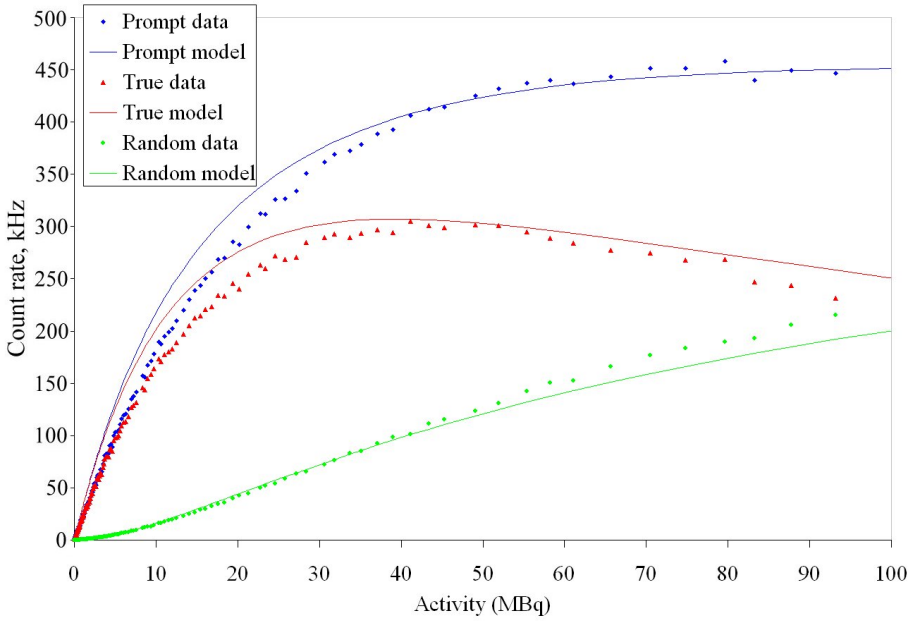

Figure 7.2: Count rates for small system with modelled lines

The model is calculated using geometrical arguments as outlined above for a central point source with two deadtime parameters included [38]. One of these parameters describes the deadtime in each detector bucket (δ_s), whilst the other describes the deadtime in the data acquisition system (the RTS in this case) and effectively limits the maximum data transfer rate.

The singles rate in each bucket, S, is given by:

$$S = \frac{R_s}{1 + \delta_s R_s} \tag{7.5}$$

Where the unperturbed singles rate, R_s, is given by $2A\Omega\epsilon/4\pi$ as above. The efficiency of the detector block, ϵ was taken as 50% as determined by an independent study. The value of the deadtime parameter, δ_s was measured to be 0.6 μs by comparing the corrected and uncorrected count rates provided by the bucket controllers over a wide range of event rates.

The true coincidence rate detected by the system is then:

$$T_0 = \frac{A\epsilon^2\Omega}{4\pi(1 + \delta_s R_s)^2} \tag{7.6}$$

And the randoms rate can be calculated in a similar manner:

$$R = 2\tau N_c R_s^2 \tag{7.7}$$

Where N_c is the number of possible coincidence combinations (48 in this case) and the resolving time (τ) is given as 6 ns.

The total event rate at the output of the coincidence processor is then calculated as the sum of the prompt coincidences and the delayed signal: this is a rate of $T_0 + 2R$. The parameter describing the deadtime in the data acquisition system (RTS) is now included as a reduction in this event rate by the factor:

$$[1 + \delta(T_0 + 2R)]^{-1} \tag{7.8}$$

Where δ is the deadtime parameter describing the data acquisition system. The fitted lines in figure 7.2 were calculated using a value of τ equal to 1 μs. This is an oversimplification of the complexity of the RTS system but the fitted lines reproduce the measured data to first order. The value of 1 μs effectively limits the maximum data transfer rate to 1 MHz (as the design allows), and justifies the development of the new data acquisition system.

Data from a moderate activity (8 MBq) point source were processed using the PEPT algorithm to determine the optimum location parameters. Figure 7.3 shows the standard deviation in resulting location (in three dimensions) as a function of the PEPT algorithm variable f for several values of N using a stationary tracer.

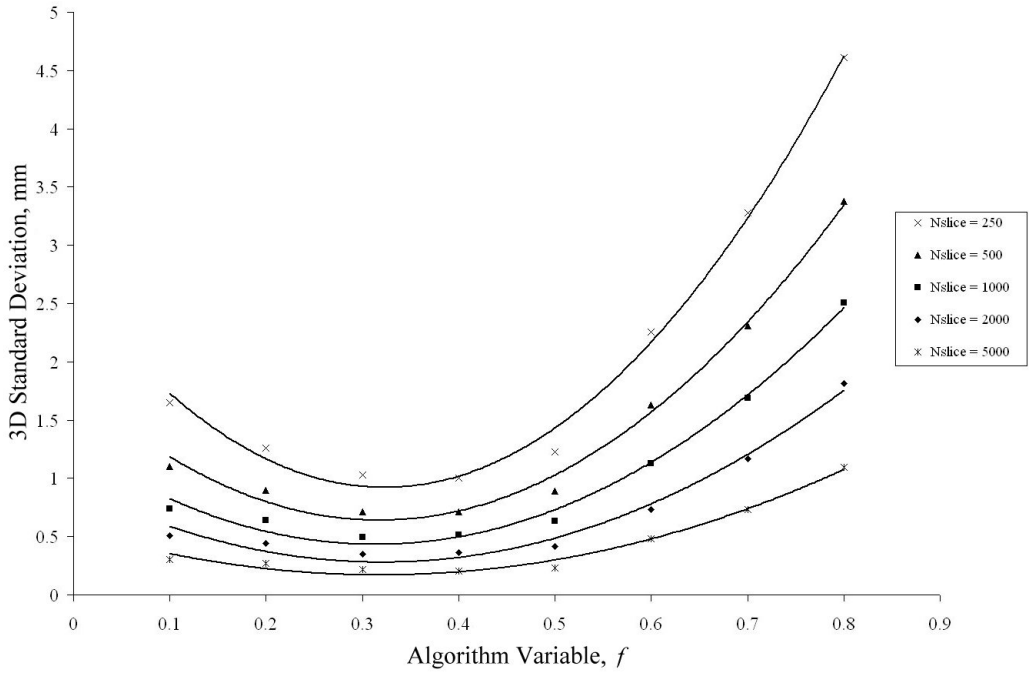

Figure 7.3: Variation in 3D std for PEPT algorithm variable f

Minimizing the standard deviation showed the optimum value of f to be approximately 32% for a 'bare' tracer in this camera configuration.

The variation in standard deviation remains small for a wide range in f. At the extremities of the plot the standard deviation is seen to increase; this occurs at high values of f because corrupt events are being retained and at low values of f because too many events are rejected to give good statistics.

Using the optimum value of f (32%), figure 7.4 shows how the same standard deviation varies as a function of the variable N for the stationary tracer. As expected, the standard deviation reduces as $N^{-0.5}$.

For a N value of 250 and an f value of 32% the standard deviations for this stationary tracer were found to be 0.52 mm, 0.51 mm and 0.72 mm along the X, Y and Z axes respectively.

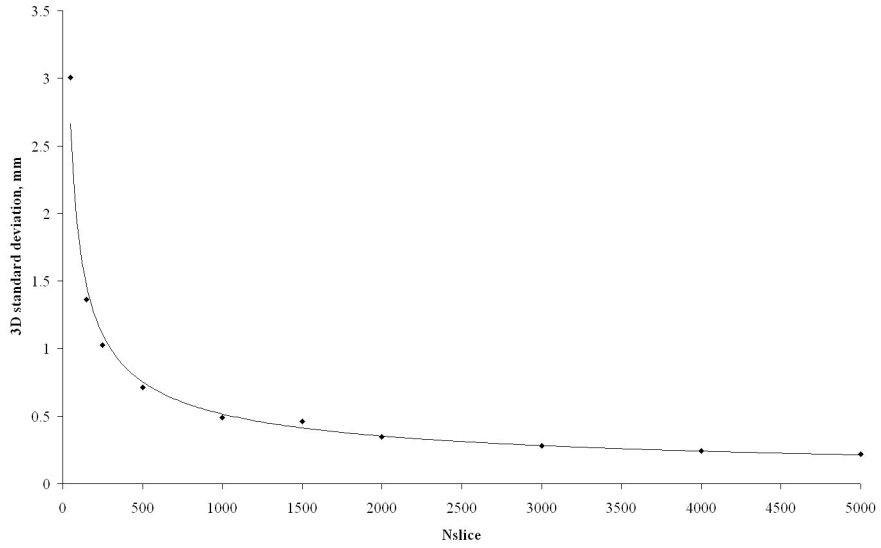

Figure 7.4: Variation in 3D std for PEPT algorithm variable N

This data was analysed by the PEPT algorithm in order to determine the particle location rate as a function of the variable N. Ideally the value of f will reduce with activity, but for simplicity a fixed value of 30% was chosen. The maximum location rate was found to be 7 k s^{-1} using an initial set, N equal to 50. Further data are shown in the following table:

N	Location rate, s^{-1}	Location precision (d), mm
50	7016	3.04
75	4650	2.10
125	2773	1.53
250	1379	1.03
500	701	0.72

A turntable of radius 85 mm was placed approximately 15 mm off-centre in the the X and Z (horizontal) and 10 mm below the midplane in the Y (vertical) directions. A ^{18}F source of activity 7.6 MBq was placed on the turntable which was rotating at approximately 2.2 revolutions per second, corresponding to a tracer velocity of 1.2 m s^{-1}.

The recorded event rate for this source was around 105k s^{-1}, these data were processed using algorithm parameters $N = 250$ and $f = 32\%$ resulting in approximately 410 locations per second. Figure 7.5 shows the data for this experiment.

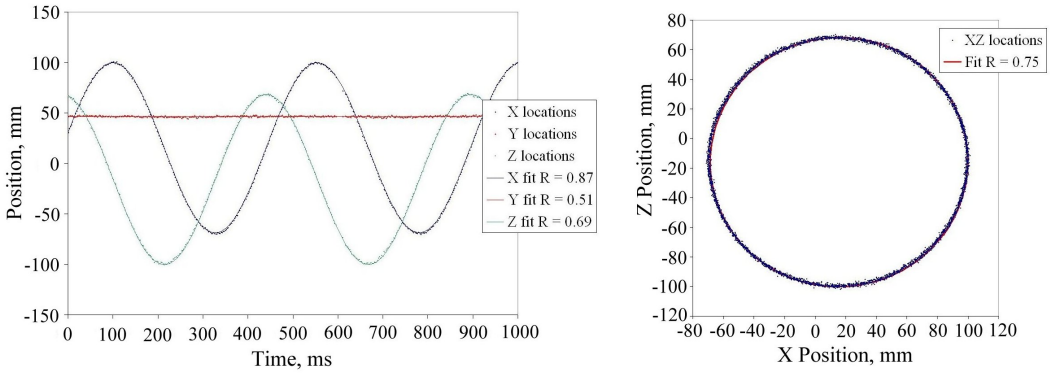

Figure 7.5: Particle locations and fitted lines for turntable data

Fitting sinusoidal functions to the set of locations gives a measure of the rms deviations from the path of the tracer.

This was found to be 0.7 mm in X, 0.6 mm in Y and 0.8 mm in Z, so that overall the standard deviation of the locations relative to the true positions of the tracer is 1.2 mm. This is limited by the accuracy in positioning the detector blocks and can certainly be improved.

Considering the regression on the Y axis it can be seen that the detector system or the turntable was not quite levelly positioned, resulting in a relatively low R value for this fit. This can be seen in more detail in figure 7.6 where the residuals for the fit are plotted for each data point, there is clearly some sinusoidal structure here as a result of errors in positioning which confirms the errors are due to the position of the system. As expected for the best fit line, the sum of these residuals over the entire data set is zero.

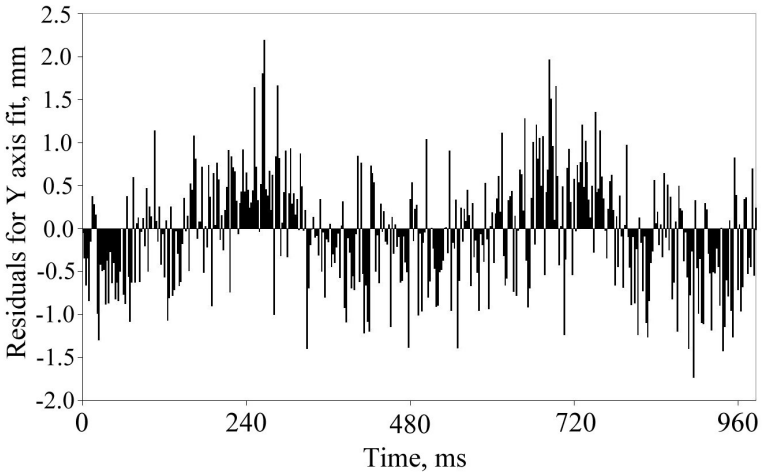

Figure 7.6: Residuals for horizontal fit

7.2 Fluidised Bed

The first experimental use of the modular camera (rather than testing of the camera itself) was conducted in the Chemical Engineering Department at the University of Birmingham. This was to perform PEPT experiments on a high pressure fluidised bed described by Seville *et. al.* in [39], it has 7 mm thick steel walls allowing operation at pressures up to 25 bar.

The fluidised bed was filled with polyethylene beads and ethylene gas was used to provide excitation. Catalyst particles were introduced to the bed to induce the chemical reactions, as the reaction occurs the catalyst particles expand which consequentially affects their motion through the bed. The PEPT experiment tracked labelled catalyst particles to determine their behaviour as the reaction continued. Lam Cheun *et. al.* [40] describes this work in further detail, Stein *et. al* [41] describes similar PEPT work and Rowe *et. al.* [42] describes the general concepts that are being studied by this experiment.

The modular camera geometries used for this experiment were the stacked cuboidal geometries described in figures 6.11 and 6.12. In these a cuboidal geometry is used made up of the square detector modules, with gaps between each detector block. The following three figures show some example data from these experiments.

Figure 7.7 shows the horizontal (X and Z) and vertical (Y) positions for the tracer particle over a period of 30 s in the bed whilst it was operating.

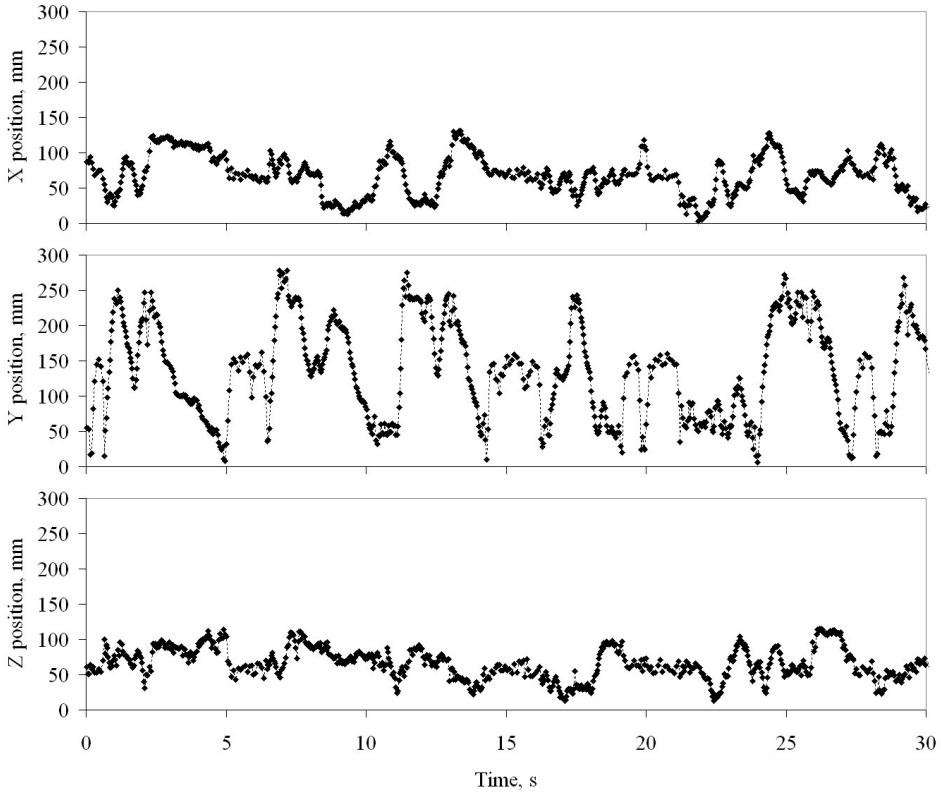

Figure 7.7: Positional data for high pressure fluidised bed

The fraction of the experimental run time spent by the tracer in each location is shown in figure 7.8. Unfortunately due to the short experimental timescale the statistics for this figure are rather low, however the bounding of the bed walls is clearly demonstrated.

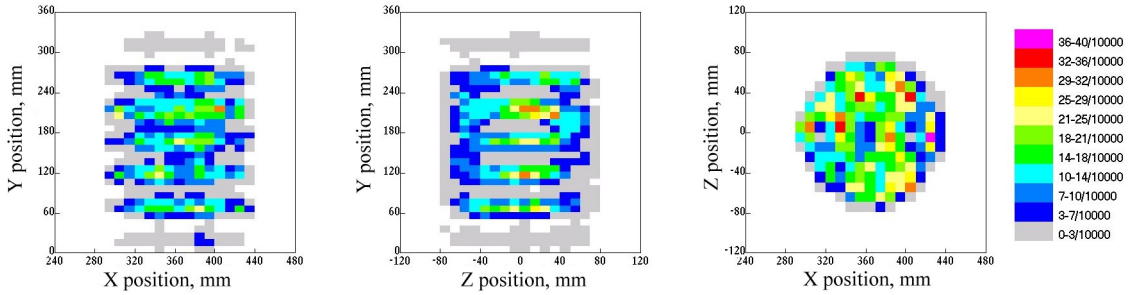

Figure 7.8: Occupancy data for high pressure fluidised bed

Figure 7.9 shows histograms of the velocities in each direction and the total speed of the particle for the same data. Over a much longer run enough statistics can be generated to provide additional velocity data. This is shown in the lower figure, where the average velocities at each point in the bed are shown. This can give information regarding the motion of the bed particles at different positions within the active volume.

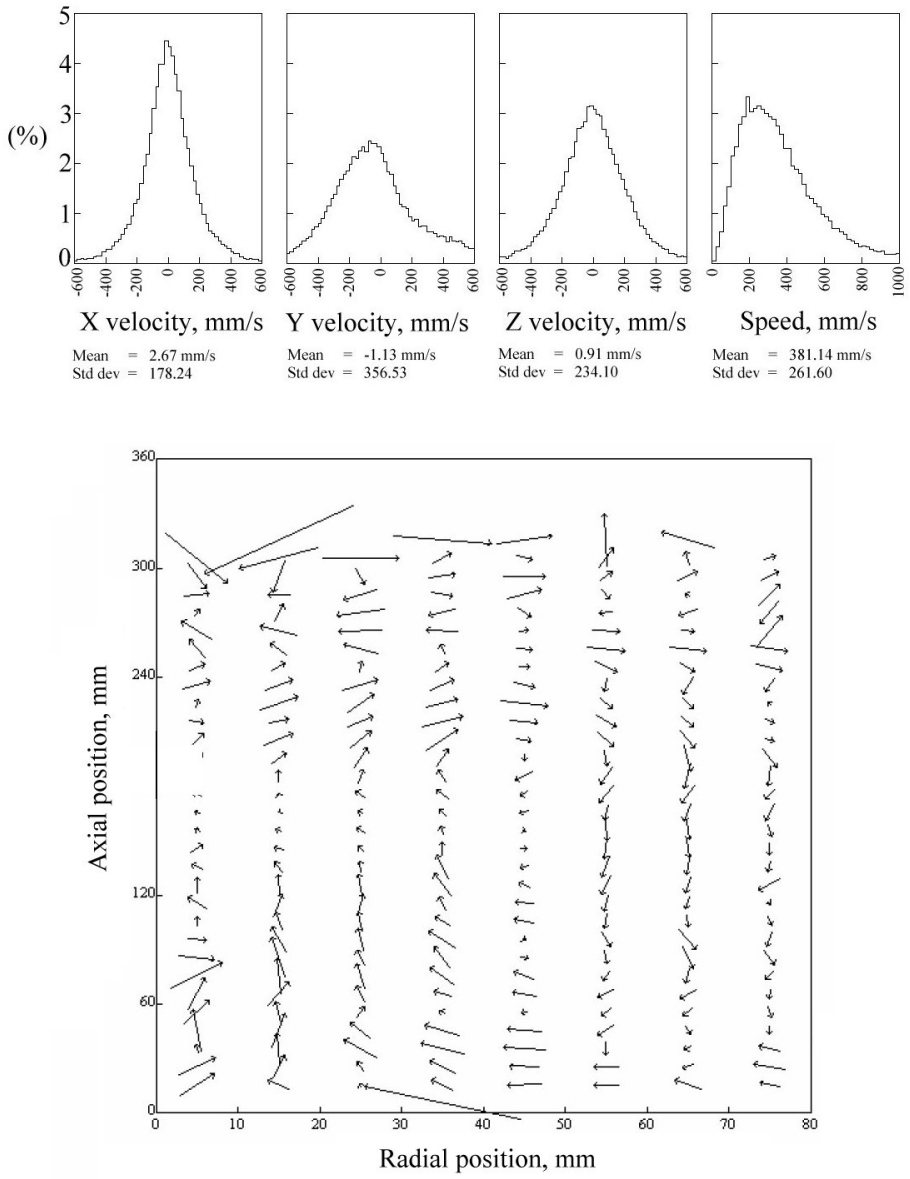

Figure 7.9: Velocity data for high pressure fluidised bed

7.3 Large Scale Fluidised Bed

Following the success of the experiments in the Chemical Engineering department it was decided to trial the modular camera on a pilot scale industrial fluid bed at the BP chemicals Ltd. (BP) site at Saltend in Hull, UK. As this was to be a study on an industrial site there were many concerns in safety and in radiological protection. Once these concerns had been adequately solved the camera was transported to the site and commissioned in the geometry detailed in figure 6.13.

The bed is divided into two semi-cylindrical sections with a vertical baffle along the X axis. Due to the scale of the bed (750 mm diameter by over 4 m height) it was not possible to cover the entire bed with the field of view of the modular camera. The data required from the experiment was concerned with measurement of the circulation time of particles within the bed (i.e. average time taken for a particle to complete a trip around the bed). It was therefore decided to have a field of view covering the top of the baffle (in order to see particle motion when it is not confined to a sector) and near the bottom of the bed. Using the cross coincidences a section in the centre of the bed was also covered by the field of view.

The fluidising gas can be injected with equal flow rates on both sides of the bed, or with asymmetric flow rates in each side to promote particle circulation. The aim of the experiment was to determine the effect differing flow rates had on particle circulation. The bed was operated at high temperature and pressure as it would be operated in industry.

The tracer particles were made in Birmingham and then transported by car to the site in Hull. Typically the transportation time was around 3 hours which resulted in the tracer particle (loaded with ^{18}F) losing over half of the loaded activity. This meant that the location rates were not as high as could be achievable due to the lower tracer activity. For experiments such these, tracers using isotopes with longer half lives: such as ^{66}Ga, were suggested. Production methods are outlined by Lewis [43].

For some of the BP experiments these new tracers were tested, however the maximum activity attainable on a tracer particle using ^{66}Ga was far lower then the highly developed methods used for ^{18}F. Thus the initial activity for both tracers upon arrival at the plant were approximately equal.

Figure 7.10 shows the PEPT locations of the tracer particle for 500 s of an experimental run. The Y axis shows the vertical position of the tracer particle within the bed. The red boxes (Y axis) illustrate the field of view arising between opposing sets of detector modules. Locations are also made between these regions due to coincidences between buckets that are not directly opposite; this has not been illustrated on the diagram.

The X axis shows which of the semi-cylindrical sections the particle is in, the red line illustrates the position of the dividing baffle. This data shows that particles move up and down on each side of the central baffle and occasionally move between the sections by transferring either under the baffle, or moving over the top of the baffle.

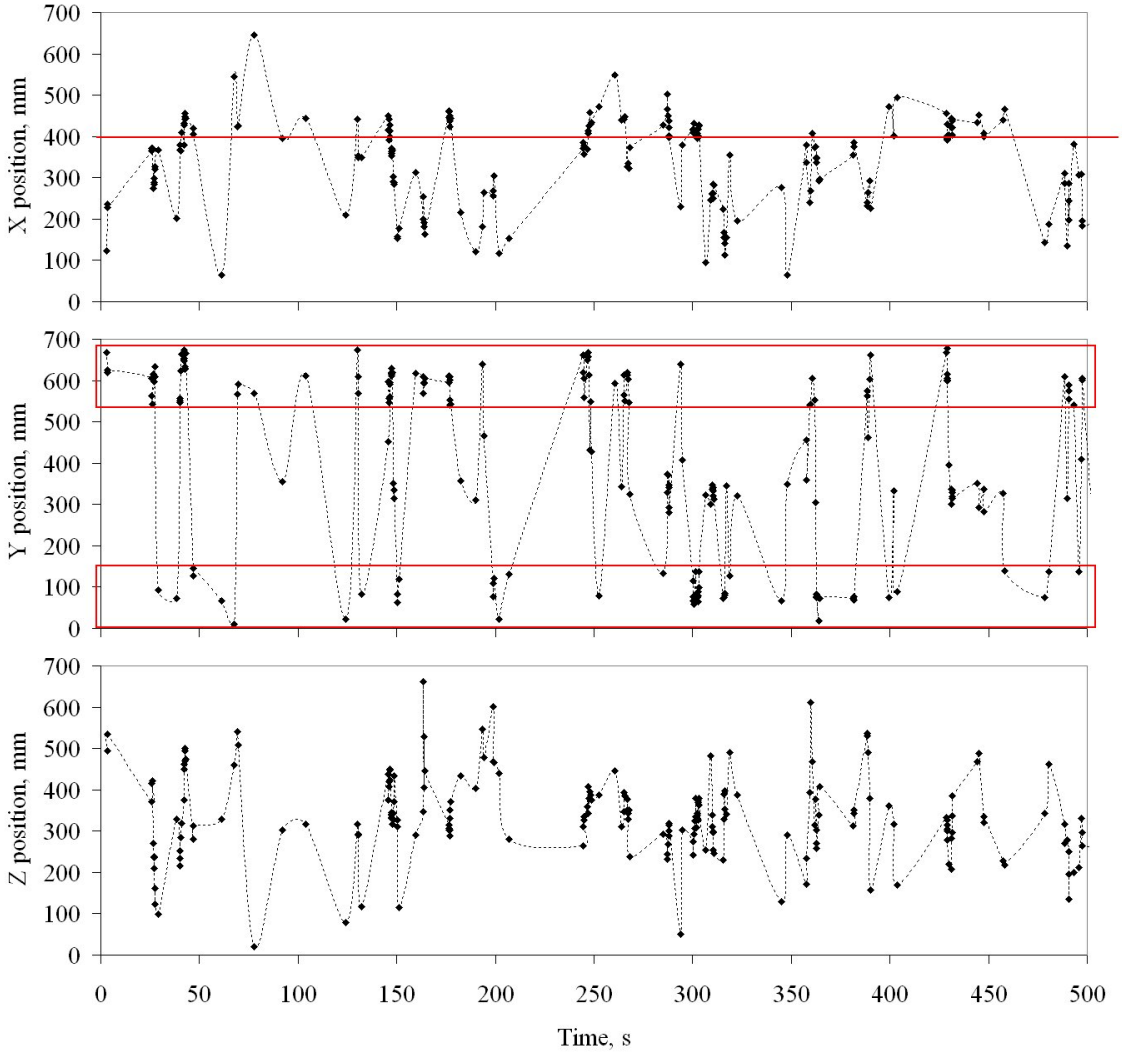

Figure 7.10: X, Y and Z location data for a large scale fluidised bed

In order to analyse the circulation time of the bed the volume was divided into four regions: top left, top right, lower right, lower left. Following this sequence then gives a single clockwise rotation. The number of transitions between consecutive regions is then plotted as a histogram. In figure 7.11 the bed is operated with a symmetrical gas flow, this results in the particle undergoing both clockwise and anticlockwise rotations in approximately equal amounts; therefore the cumulative effect is zero.

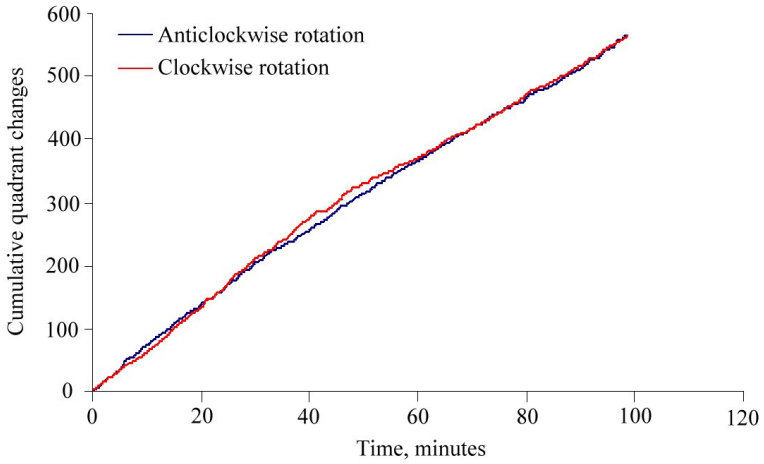

Figure 7.11: Clockwise and anticlockwise rotation for symmetric gas flow

In figure 7.12 the gas flow is anti-symmetric resulting in the anticlockwise rate of motion being approximately twice that of the clockwise rate. Further analysis of these data by Ingram *et. al.* [27] results in the net effect being an anticlockwise solids circulation of around 1.7 rpm for this data.

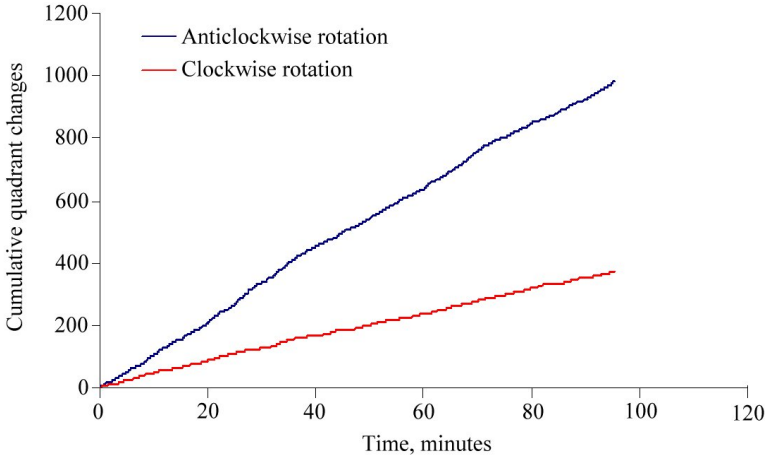

Figure 7.12: Clockwise and anticlockwise rotation for anti-symmetric gas flow

Perhaps of more interest for this thesis is the following data. As discussed in section 6.2 the field of view for many applications of the modular camera is non-uniform. For PEPT experiments this is not regarded as an issue as regions of high sensitivity will result in many accurate locations, whilst regions of low sensitivity result in few locations.

For a PET image however this would seriously degrade the data, as a PET image corresponds to the distribution of the radioactive tracer fluid over the volume. Regions of high sensitivity would therefore receive more events then regions of low sensitivity for a fixed exposure.

This effect can be clearly seen in figure 7.13, this shows three orthogonal projections through the fluidised bed with the total number of location events tallied for small pixels dividing the area. The sensitivity of the camera is then shown by this image (compare this to figure 6.13), as locations are made more frequently in regions of high sensitivity.

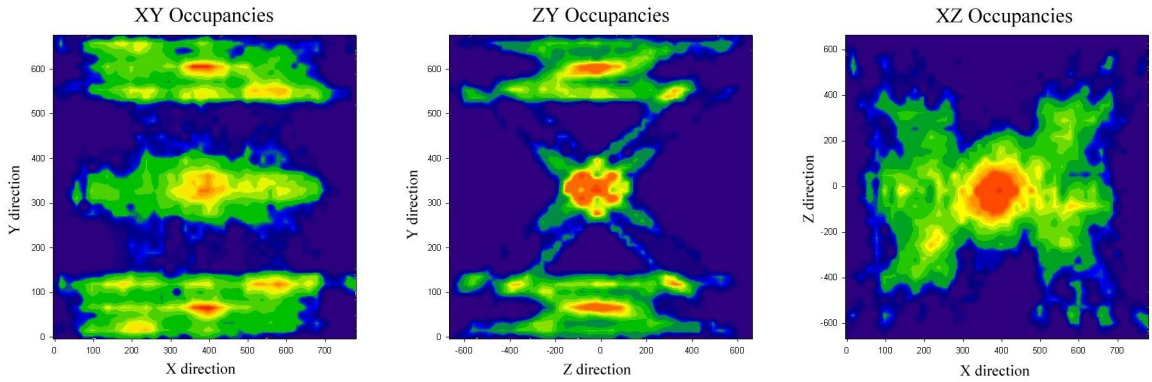

Figure 7.13: Sensitivity of the large scale modular camera as a function of position

If this image is corrected by taking into account the variations in sensitivity the image in figure 7.14 is obtained. This is done by considering the relative amount of the experimental time the tracer particle spends in each volume element, rather then the total number of times the element was 'hit'.

The resulting image is equivalent to a PET image given that the experiment is run for a long enough period of time to get high statistics and provided the tracer exhibits all possible forms of motion within the field of view.

The corrected image then shows the tracer particle spending most of its time near the bottom of the bed, occasionally rising and falling on either side of the baffle. The baffle location can be inferred from the lack of data in the centre of the XY plot.

The XZ plot shows how the tracer spends approximately equal time on both sides of the bed for this data. Note that the scales on the axes of these plots differ, this is due to the size of the field of view along each axis.

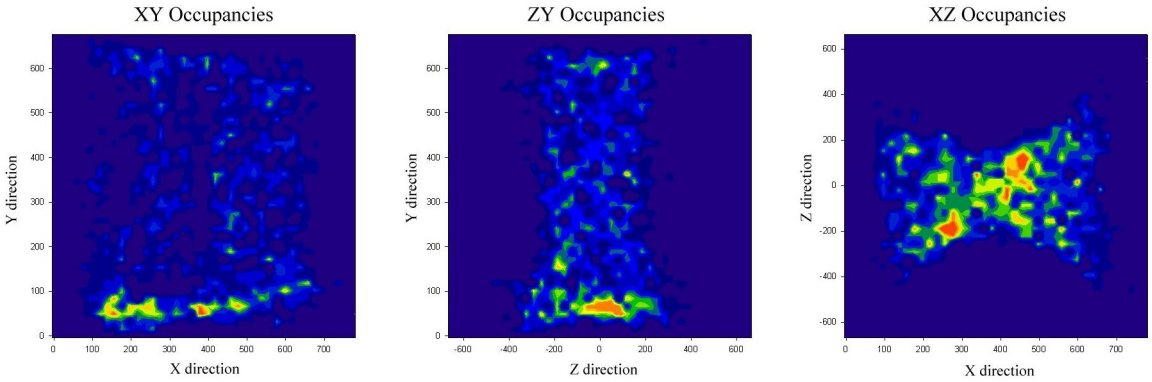

Figure 7.14: Corrected occupancy data for BP experiment

7.4 Metal Casting

The geometry described in figure 6.16 was used to study the motion of small inclusions during pouring of liquid aluminium into a casting mould. In these experiments a tracer particle was added to the liquid aluminium and then poured into the mould, the tracer particle then represents an inclusion in the metal melt. The figure in 7.15 shows the resultant tracer trajectory overlaid on a photo of one of the aluminium casts as the metal was poured into the mould and allowed to solidify.

The typical event rates for a small alumina tracer of activity around 20 MBq were measured to be between 10 and 3 kHz. Data was recorded using the RTS data acquisition system.

Using algorithm variables f equal to 10 %, and the number of events used per location N equal to 250, the tracer particle could be located once every 25 ms with a theoretical precision of around 0.6 mm. In practice this value was slightly worse (precision of around 1 mm) because of errors in positioning the detector modules and due to the lack of coincidence combinations in one of the cross planes (section 6.5).

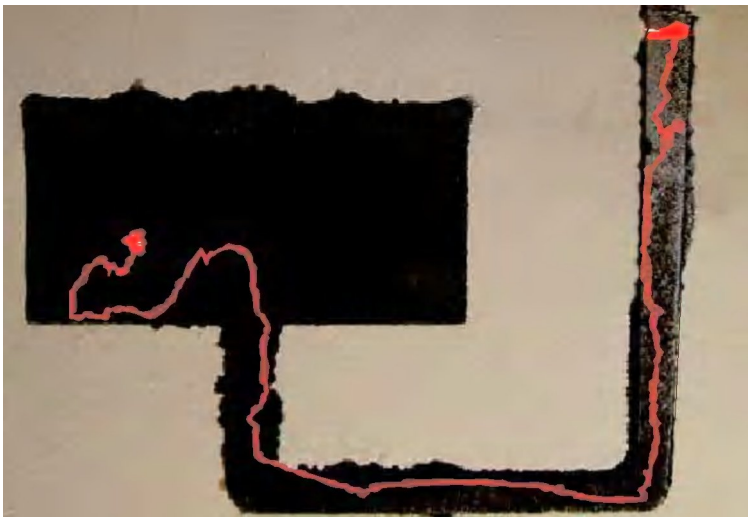

Figure 7.15: Trajectory of a particle introduced into liquid metal during casting

7.5 Plastic Extrusion

The final example given in this section used the tightly packed geometry given in figure 6.17 for PEPT experiments on an industrial twin screw extruder in the University Interdisciplinary Research Centre (IRC). This was the first experimental campaign to use the new data acquisition system discussed in chapter 4 with the added benefits of the new system.

The data presented below (figure 7.16) shows the tracer trajectory as it passes through the extruder system. The tracer was inserted into the device at time 0 s and takes approximately 40 s to reach the field of view of the camera. It then takes another 40 s to pass through the field of view (along the horizontal Y axis) where it is eventually collected and the experiment repeated. The Z axis shows the vertical motion.

This data file was chosen as a representative example of many data runs as it clearly shows two different flow regimes; the first part sees the tracer confined to motion around one of the extrusion shafts whilst in the second part the tracer is swapping between the two shafts upon each revolution.

Simplistic analysis of the gradient of the Y position for each regime gives the average tracer velocity for each regime. This shows that when the tracer particle swaps between the shafts (second regime) the velocity through the system is approximately 5.4 times higher than when the tracer is confined to motion around one shaft.

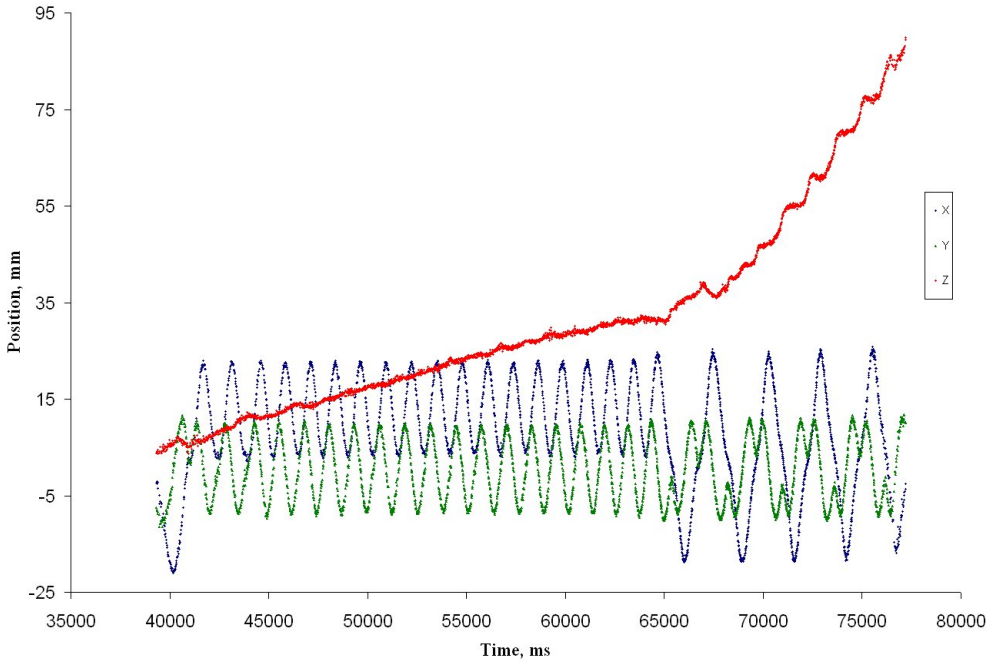

Figure 7.16: Trajectory data for extrusion experiment

Figure 7.17 shows the corrected occupancies for each orthogonal view in this experimental run. The shaft the tracer was confined to at the start of the run is clearly identified from these plots, along with the manner that the particle swaps between shafts.

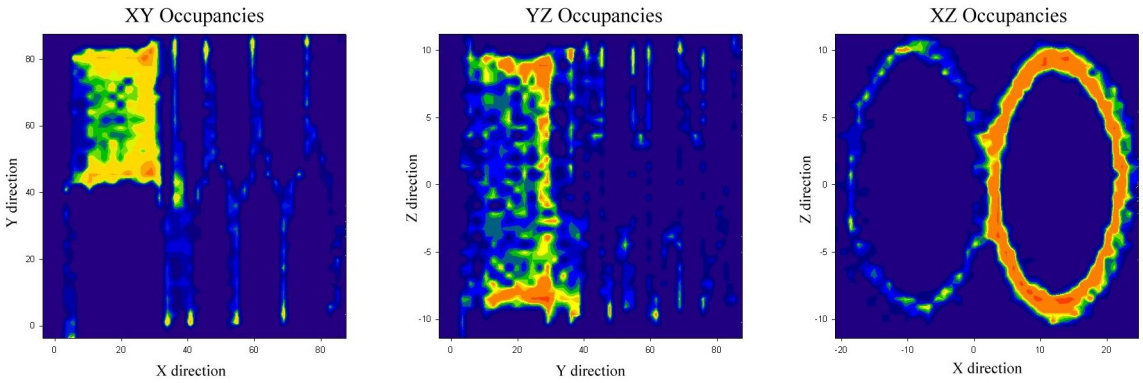

Figure 7.17: Occupancy data for extrusion experiment

These data were recorded at a sustained data rate of 120 kHz, with short bursts reaching 1.2 MHz. This high data rate allows for PEPT analysis with an increased value of N resulting in frequent and accurate locations. For this data N was set to 2000 events per location, which corresponded to an average location rate of once per 11.3 ms, with the corresponding precision from using such a large set of events. The measured optimum value of f 32% was used for this analysis.

Chapter 8

Conclusions

Over the past three years the systems used to perform Positron Emission Particle Tracking (PEPT) on industrial applications have seen some very interesting and useful developments. Primarily the development of the modular positron camera, with all that it has to offer in efficiency, spatial resolution and event data rate, has been a great advancement for the Positron Imaging Centre. The development of this system has resulted in a novel instrument with a great potential for delivering new information related to granular and fluid flows used in industrial and physical processes.

The flexible geometry offered by the camera is useful for two reasons: firstly the field of view of the camera can be custom made for the application in question, allowing study of large or bulky apparatus, secondly the camera system is transportable and it has been possible for the first time to study systems fixed *in situ*.

The modular camera has seen a number of different developments during this time as it started as a proof-of-concept system proving that the modular camera idea was feasible for PEPT work. Soon after this the camera was placed in rugged boxes and the first PEPT work, both on a large scale and outside of the University campus, was performed.

Valuable lessons about the modular camera were learned from these initial trials, examples of these are that having gaps between detector modules can degrade the data and that the original data acquisition systems provided with the camera were overly complex and slow for PEPT experiments. The performance and limitations of the modular camera and the data acquisition system have been characterised, allowing the use of the device to be optimised for each particular study.

The original data acquisition system was used with a small ring camera developed for high resolution studies. This camera was used to perform PEPT work, but also allowed for investigation into its operating characteristics. The count rate data and modelling showed that the maximum data acquisition speed was limited by the relatively slow data acquisition hardware and motivated the development of a new, high speed system.

Improvements in computing technology are consistently producing smaller features on silicon integrated circuits, this has the effect of making complex devices both faster and cheaper. It proved possible to purchase a data acquisition hardware card for a desktop computer, with superior performance to the original data acquisition hardware for minimal cost.

A number of dedicated support circuits have been designed and built to interface the front end detector systems with the computer allowing high speed data acquisition for PEPT (or PET).

The new systems have been used for a number of successful experiments and have produced data at rates up to 4 MHz, a factor of 40 improvement in the data rates previously seen in the Positron Imaging Centre. This high data rate has allowed the PEPT location algorithm to be used to locate the tracer particle more frequently and with increased precision then previously possible (provided the tracer remains in a sensitive region of the field of view of the camera).

A number of software applications have been written to run the new hardware and display both the raw data arriving from the camera and display the PEPT location data. These software applications have seen considerable use during the development of the camera systems, running of PEPT experiments and for displaying the and analysing PEPT data post acquisition.

8.1 Further Work

Suggestions for further work have been included at various points in this thesis. Some of the immediate considerations for further development are using the modular camera in a range of different geometries and for a number of different applications. This work would extend the potential of PEPT into new regimes, however it does not describe future work in the development of the imaging systems for industrial process tomography.

Provided that connections can be made between the data buses it seems reasonably possible to connect extra coincidence processors to the new data acquisition system allowing coincidence data to be recorded for more then the currently available 32 detector modules. This could be used to make a large sensitive system or to study, for example pipelines.

There are a total of 120 possible coincidence combinations between the 16 detectors providing data to each coincidence processor. It is therefore proposed to design a new backplane (similar to the VME 64 standard) with which to connect the ring receivers and coincidence processors. With a custom design it would be possible to add extra ports for additional ring receivers and coincidence processors, the current aim is to accept data from 48 detector modules by using 2 extra ring receivers and a further 4 coincidence processors. It would be possible to provide the connections with fewer components but the system is to be designed to be more flexible.

Development of the software would benefit from further work. With the new data acquisition system the raw data files have the potential to be extremely large. Already discussed are techniques for keeping these data files in a manageable size to they are easily transported and analysable.

However the processing of such large data files using the PEPT algorithm takes a lot of computer time, and whilst the program can be left running overnight or during weekends, the processed data file is not necessarily immediately accessible.

Techniques to speed up the PEPT algorithm need to be investigated. It should be possible in principle to use parallel processing techniques to process different parts of the data file simultaneously. Modern PCs now contain multiple processor cores, so this technique is now possible using a modern desktop PC such as the one containing the data acquisition card and software. Other options include the use of the Nuclear Physics group processor farm, or the University's new computing cluster Blue Bear, both of which offer huge computing power spread across a multitude of processor elements.

Other software advances that are worth investigating are applications which can be used to display both the raw data file and the computed PEPT algorithms. As discussed in section 5.4, some software tools have already been developed to display these data using both the Windows GDI and graphics processor libraries.

These programs have proven to be extremely useful for visualising how the PEPT method and camera systems work, and provide a useful way of calibrating the detector geometry for the modular camera. Commercial software such as MathCad (by Parametric Technology Corporation) can be used to generate rendered images from data files. The process is slow, however once the calculation has been performed the data image is easy to manipulate; these techniques merit further investigation.

Bibliography

[1] M.R. Hawkesworth, M.A. O'Dwyer, J. Walker, P. Fowles, J. Heritage, P.A.E. Stewart, R.C. Witcomb, J.E. Bateman, J.F. Connolly, and R. Stephenson. A positron camera for industrial application. *Nucl. Inst. Meth.*, A253:145–157, 1986.

[2] H. Zaidi, editor. *Quantative analysis in nuclear medical imaging*. Wiley & Sons Inc., 2000.

[3] Siemens press release, 2007.

[4] C.R. Bemrose, P. Fowles, M.R. Hawkesworth, and M.A. O'Dwyer. Application of positron emission tomography to particulate flow measurement in chemical engineering processes. *Nucl. Inst. Meth.*, A273:874–880, 1988.

[5] D.J. Parker, C.J. Broadbent, P. Fowles, M.R. Hawkesworth, and P. McNeil. Positron Emission Particle Tracking - A technique for studying flow within engineering equipment. *Nucl. Inst. Meth. A*, 19:9, 1993.

[6] H. Anger. A new instrument for mapping gamma ray emitters. *Berkeley, CA: University of California Radiation Laboratory*, 3653, 1957.

[7] H. Anger. Positron scintillation camera. *Nucl Med.*, 4:326–330, 1963.

[8] Hoffmann A C, Dechsiri C, van de Wiel F, and Dehling H G. Pet investigation of a fluidized particle: spatial and temporal resolution and short term motion. *Measurement Science and Technol.*, 16:851–858, 2005.

[9] Dechsiri C, van der Zwan E A, Dehling H G, and Hoffmann A C. Dispersion of particle pulses in fluidized beds measured by positron emission tomography. *AIChE Journal*, 51:791–801, 2005.

[10] C.S. Stellema, J. Vlek, R.F. Mudde, J.J.M. Goeij, and C.M. van den Bleek. Development of an improved positron emission particle tracking system. *Nucl. Inst. Meth. A*, A404:334–348, 1998.

[11] M.S. Fraguio, M.C. Cassanello, M.A. Cardona, D. Hojman, and H. Somacai. Tracer fast ascending velocities in a bubble column inferred from CARPT calibration experiments. 5th world conference on industrial process tomography, 2007.

[12] D. Mosiemian, N. Devanathan, and M.P. Dudukovic. Radioactive particle tracking technique for investigation of phase recirculation and turbulence in multiphase systems. *Rev. Sci. Instrum.*, 63(10), 1992.

[13] S. Bhusarapu, M.H. Al-Dahhan, and M.P. Dudukovic'. Solids flow mapping in a gassolid riser: Mean holdup and velocity fields. *Powder Technology*, 163:98–123, 2006.

[14] Y.B. Yang, N. Devanathan, and M. P. Dudukovic. Liquid backmixing in bubble columns via computer-automated radioactive particle tracking (carpt). *Experiments in Fluids*, 16:1–9, 2004.

[15] P.K. Marsden, R.J. Ott, J.E. Bateman, S.R. Cherry, M.A. Flower, and S. Webb. The performance of a multiwire proportional chamber positron camera for clinical use. *Phys. Med. Biol.*, 34:1043–1062, 1989.

[16] J.E. Bateman, J.F. Connolly, R. Stephenson, G.J. Tappern, and A.C. Flesher. The Rutherford Appleton Laboratory's mark I multiwire proportional counter positron camera. *Nucl. Inst. Meth.*, 225:209, 1984.

[17] D.J. Parker, M.R. Hawkesworth, C.J. Broadbent, P. Fowles, T.D. Fryer, and P.A. McNeil. Industrial positron-based imaging: princliples and applications. *Nucl. Inst. Meth. A*, A348:583–592, 1994.

[18] D.J. Parker, R.N. Forster, P. Fowles, and P.S. Takhar. Positron Emission Particle Tracking using the new Birmingham positron camera. *Nucl. Inst. Meth. A*, A477:540–545, 2002.

[19] Siemens Medical Systems. *ECAT aquisition systems maintenence manual*, 55-55-693 edition.

[20] S. Webb, editor. *The Physics of Medical Imaging.* Institute of Physics, 1988.

[21] C. Bomford and I. Kunkler. *Textbook of radiotherapy.* Walter and Miller, 2002.

[22] A. Sadrmomtaz, D.J. Parker, and L.G. Byars. Modification of a medical PET scanner for PEPT studies. *Nuc. Inst. Meth. A*, A573:91–94, 2007.

[23] T.W. Leadbeater and D.J. Parker. A high speed pc-based data acquisition & control system for positron imaging. 8th International Position Sensitive Detectors Conference, 2008.

[24] Z. Yang, D.J. Parker, P.J. Fryer, S. Bakalis, and X. Fan. Multiple-particle tracking - an improvement for positron particle tracking. *Nucl. Inst. Meth.*, A564:332–338, 2006.

[25] X. Fan, D.J. Parker, and M.D. Smith. Labelling a single particle for Positron Emission Particle Tracking using direct activation and ion-exchange techniques. *Nucl. Inst. Meth.*, A562:345–350, 2006.

[26] D.J. Parker, T.W. Leadbeater, X. Fan, M.N. Hausard, A. Ingram, and Z. Yang. Positron imaging techniques for process engineering: recent developments at birmingham. *Meas. Sci. Tech*, A326:592–607, 2008.

[27] A. Ingram, M. Hausard, X. Fan, D.J. Parker, J.P.K. Seville, N. Finn, and M. Evans. Portable positron emission particle tracking (PEPT) for industrial use. The 12th international conference on fluidization - New horizons in fluidization engineering, 2007.

[28] N.M. Clarke. Fifty-one years of the nuffield cyclotron, 2001.

[29] X. Fan, D.J. Parker, and M.D. Smith. Enhancing ^{18}F uptake in a single particle for positron emission particle tracking through modification of solid surface chemistry. *Nucl. Inst. Meth.*, A558:542, 2006.

[30] G. Audi, O. Bersillon, J. Blachot, and A.H. Wapstra. The NUBASE evaluation of nuclear and decay properties. *Nuclear physics A*, 729:3–128, 2003.

[31] W.F. Jones, M.E. Casey, L.G. Byars, and S.G. Burgiss. A VME based, real time sorter design for positron emission tomography. *Nucl. Inst. Meth.*, A326:592–607, 1993.

[32] M.E. Casey and R. Nutt. A multicrystal two dimensional BGO detector system for positron emission tomography. *IEEE Nucl. Sci. Symposium*, 1985.

[33] T.J. Spinks, T. Jones, M.C. Gilardi, and J.D. Heather. Physical performance of the latest generation of commercial positron scanner. *IEEE trans. Nucl. Sci.*, 35:1:721–725, 1988.

[34] Alireza Sadrmomtaz. *Modification of a medical PET scanner for PEPT studies.* PhD thesis, Dept. Physics and Astronomy, the University of Birmingham, 2005.

[35] CTI. Circuit and systems diagrams for ECAT scanners.

[36] T.W. Leadbeater and D.J. Parker. A high speed PC based data acquisition and control system for positron imaging. *Accepted for publication in Nucl. Inst. Meth. A*, December, 2008.

[37] T.W. Leadbeater and D.J. Parker. A positron camera with flexible geometry for the study of industrial processes. 5th world conference on industrial process tomography, 2007.

[38] D.J. Parker, T.W. Leadbeater, X. Fan, M. Hausard, A. Ingram, and Z. Yang. A high speed PC based data acquisition and control system for positron imaging. *Accepted for publication in Nucl. Inst. Meth. A*, December, 2008.

[39] J.P.K. Seville, A.M. Salleh, A. Ingram, A. McCormack, R.W. Greenwood, and V. Reiling.

[40] Y.V. Lam Cheun, A. Ingram, and J. Seville. Solids motion in fluidised beds of fine particles comparing particle and bubble velocities using pept. Particulate Systems Analysis, 2008.

[41] M. Stein, Y.L. Ding, J.P.K. Seville, and D.J. Parker. Solids motion in bubbling gas fluidised beds. *Chem. Eng. Sci.*, 55:5291–5300, 2000.

[42] P.N. Rowe, B.A. Partridge, A.G. Cheney, G.A. Henwood, and E. Lyall. The mechanisms of solids mixing in fluidised beds. *Trans. Inst. Chem. Engrs.*, 43:271–286, 1965.

[43] M.R. Lewis et al. Production and purification of gallium-66 for preparation of tumor-targeting radioparmaceuticals. *Nucl. Med. Bio.*, 29:701–706, 2002.