

UNIVERSITY OF CAPE TOWN

MASTERS THESIS

**Sub-Millimetre Positron-Emission Particle
Tracking Using a CdZnTe Semiconductor
Array**

Author:
Nicholas HYSLOP

Supervisors:
Dr. Steve PETERSON
& Dr. Tom LEADBEATER

*A thesis submitted in fulfillment of the requirements
for the degree of Master of Science*

in the

MeASURE Unit
Department of Physics

August 1, 2021

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

UNIVERSITY OF CAPE TOWN

Abstract

Faculty of Science
Department of Physics

Master of Science

Sub-Millimetre Positron-Emission Particle Tracking Using a CdZnTe Semiconductor Array

by Nicholas HYSLOP

The Positron Emission Particle Tracking (PEPT) technique has been in development in Cape Town since 2009, and typically allows the tracking of a 1 mm diameter positron-emitting source travelling at 1 m/s with precision around 1 mm at kHz rates. Conventionally, PEPT utilises large, high efficiency scintillator detectors and relies on high data-rates obtained from high activity tracer particles to achieve the quoted performance.

However, as we look to apply the technique to more challenging liquid and mixed phase systems, smaller tracer particles are required (of the order $10\ \mu\text{m}$) which are necessarily limited in attainable activity and therefore applicability in current systems. This dissertation investigated the use of pixelated room-temperature semiconductor CdZnTe detectors with high pixel resolution to implement PEPT. Using a PolarisJ CZT detector module array built by H3D Inc. (Ann Arbor, MI, USA) with a field-of-view of $40 \times 60 \times 20\ \text{mm}^3$, a positron-emitting ^{22}Na source (1 mm diameter) with an activity of 20 kBq was located in three-dimensional space with a radial location uncertainty of $< 600\ \mu\text{m}$ (68% coverage) and location rate of $\approx 0.3\ \text{Hz}$. A ^{68}Ga source (160 μm diameter) with an activity of 1 MBq was located with a radial position uncertainty of $< 800\ \mu\text{m}$ (68% coverage) and location rate of $\approx 10\ \text{Hz}$. A 37 kBq ^{22}Na tracer was tracked with sub-millimetre precision undergoing circular motion at radii between 0.5 - 20 mm and speeds between 0.25 and 15 mm/s. We have showed sub-millimetre tracking of tracer particles with tracer activities orders of magnitude lower than currently possible using conventional PEPT systems.

A potential application of this work is in providing on-line beam diagnostics in proton radiotherapy. The PEPT technique was therefore applied to Geant4 simulation data of a 55 MeV and 200 MeV proton beam passing through an organic target to locate the beam-line with uncertainties within the incoming beam width. Such proton beams produce positron emitters at activities too low to measure with conventional scintillators. With the current CZT arrays already in use for dose verification imaging, incorporating such beam-tracking to improve image quality can be easily implemented in future.

Acknowledgements

No worthwhile science is done in a vacuum and this dissertation is no exception. I have received a great deal of support throughout my M.Sc. and would like to acknowledge all those who helped me over the past two years.

Firstly, I would like to thank my supervisors Steve Peterson and Tom Leadbeater. Your countless hours of knowledge, insight, expertise, guidance and wisdom have fundamentally changed how I think about performing great research, and this work would not exist without you. Thank you for always pushing me to improve and sharpen my work.

I would like to thank the Harry Crossley Foundation, who awarded me the fellowship that funded my fees and living expenses, allowing me to devote my whole attention to this work. I am proud to call myself a fellow and I hope you are proud to have me as one.

I would like to thank the Physics Department at the University of Cape Town, headed up by Andy Buffler, for providing the support structure that makes a M.Sc. like this possible. From the physical labs, offices and resources to the wisdom of the staff, I am grateful for making all my work possible. I would also like to thank Jill Patel for the endless administration work she has done for me over the years.

I would like to thank Willy Kaye of H3D Inc. and Zong He of the Nuclear Engineering and Radiological Sciences group at the University of Michigan for allowing me to visit them personally in Ann Arbor, Michigan to perform invaluable experiments on their equipment and discuss my work with them. I would like to thank Steven Brown and Michael Streicher for supporting me professionally throughout my stay at H3D Inc. and Daniel Shy for accommodating me at the University of Michigan. Special thanks to Michael Streicher for allowing me to live at his home throughout my stay.

I would like to thank all my colleagues at the Physics Department who made it so much fun to come to the office everyday, of whom there are too many to name. Thank you for all the advice, laughs, insights, support and friendship you have provided to me over the years. Being able to interact with people like you makes research all that much more worthwhile.

Finally, I would like to thank my friends and family who have supported my scientific journey, not just over the last two years but over my whole lifetime. Thank you to my parents, Lee-Anne and Mike Hyslop, and my siblings Chris and Robs, for their love, support, counsel and encouragement. You will always be there for me and I am eternally grateful for that. Thank you to my friends Anja, Joe, Jon, Daniel, Andrew, Scott and many others for your wonderful friendship. And thank you to my girlfriend, Bianca, for all the love and support you've given me.

Contents

Abstract	i
Acknowledgements	ii
1 Introduction	1
1.1 Dissertation Outline	1
1.2 Positron Emission Particle Tracking (PEPT)	2
1.2.1 Theory	2
1.2.2 PEPT: Birmingham (1988) to Cape Town (2020)	7
1.2.3 Tracer Production	9
1.3 Deadtime Effects on Detected Singles Event Rates	11
1.3.1 Non-Paralysable Deadtime	11
1.3.2 Paralysable Deadtime	12
1.3.3 Measured Rate from a Single Detector detector module	12
1.3.4 Measuring the Deadtime Parameters	13
1.4 Deadtime Effects on Coincidence Counting	14
1.4.1 True Coincidence Rate Model	16
1.4.2 Random Coincidence Rate Model	17
1.4.3 Prompt Coincidence Rate Model	17
1.4.4 Model Predictions	18
1.5 Extending the Positron Emission Particle Tracking (PEPT) Technique	20
1.6 Prompt Gamma-ray Imaging (PGI)	21
1.6.1 Proton Radiation Therapy	21
1.6.2 Prompt Gamma-Ray Imaging using Compton Cameras	22
1.6.3 Using PEPT to Inform DCA Filtering	24
2 Equipment	26
2.1 Cadmium Zinc Telluride	26
2.2 The H3D Inc. PolarisJ CdZnTe Detector	26
2.2.1 CZT Crystals and Signal Processing Electronics	27
2.2.2 Data Collection and Event Information Processing	27
2.3 The UCT PolarisJ Coincidence Detector Array	30
2.4 H3D Inc.'s Small Animal PET System	31
2.5 The Anet A8 Positioning Stage	33
2.5.1 Anet A8 Design	33
2.5.2 Anet A8 Modifications and Positioning Performance	33
2.6 β^+ Radiation Sources	35
2.6.1 ^{22}Na Sources	35
2.6.2 ^{68}Ga Sources	35

3	Experiment 1 - Proof of Concept	37
3.1	Locating a Central Stationary Point Source	37
3.1.1	Experimental Setup	37
3.1.2	Energy and Coincidence Timing Windows	37
3.1.3	Lines-of-Response and F-Optimisation	40
3.2	PEPT Location Uncertainty	44
3.3	Field-of-View Sensitivity	46
3.3.1	Experimental Setup	46
3.3.2	Results	47
4	Experiment 2 - Deadtime Behaviour	50
4.1	Detector Deadtime Response to Short-Lived ^{68}Ga Source	50
4.1.1	Experimental Setup	50
4.1.2	Single Module Response	51
4.1.3	Deadtime Parameter Fit	52
4.1.4	Coincidence Rates	54
4.2	Locating an Optimal Activity Stationary Source	57
4.2.1	F-Optimisation	57
4.2.2	Timing and Energy Windows	58
4.3	Improving Data Rates	63
4.3.1	The Limits of High Resolution	63
4.3.2	Effects on F-Optimisation	64
4.4	Conclusions	65
4.4.1	Further Work - Multiple Interaction Events	66
5	Experiment 3 - Moving Tracer	68
5.1	Experimental Setup	69
5.2	Stationary Source Performance	69
5.2.1	Low Activity ^{22}Na Source	69
5.2.2	High Activity ^{22}Na Source	71
5.3	Circular Motion Tracking Performance	73
5.3.1	Low Activity ^{22}Na Source	73
5.3.2	High Activity ^{22}Na Source	78
5.4	Conclusions	81
6	Proton Beam Tracking using PEPT	82
6.1	Theory	82
6.2	Experimental Setup	83
6.2.1	55 MeV Proton Beam	84
6.2.2	200 MeV Proton Beam	84
6.3	Results	85
6.3.1	55 MeV Proton Beam	85
6.3.2	200 MeV Proton Beam	86
6.4	Feasibility Estimates	92
6.4.1	Real-Time Imaging	92
6.4.2	Post-Treatment Imaging	93
6.5	Conclusions and Further Investigation	93

7	Conclusions and Further Investigation	95
7.1	PolarisPEPT	95
7.1.1	Locating a Stationary Point Source	96
7.1.2	Deadtime Behaviour	97
7.1.3	Tracking a Moving Tracer	99
7.2	PolarisPEPT Further Work	100
7.2.1	Immediate Follow Up Work	100
7.2.2	Hybrid CZT-BGO Systems	100
7.2.3	Small-Scale Systems	101
7.3	Beam Tracking using PolarisPEPT	101
7.4	Further Work	102
A	PolarisJ Module Specifications Sheet	103
B	Anet A8 Specifications Sheet	106
C	NRW-100 Ion-Exchange Resin Specifications Sheet	107
	Bibliography	109

List of Abbreviations

BGO	Bismuth Germanium Oxide
CZT	Cadmium Zinc Telluride
FOV	Field Of View
FWHM	Full Width at Half Maximum
FWTM	Full Width at Tenth Maximum
LOR	Line-Of-Response
MWPC	Multi-Wire Proportional Chamber
PEPT	Positron Emission Particle Tracking
PET	Positron Emission Tomography
PGI	Prompt Gamma Imaging
PRT	Proton Radiation Therapy
UCT	University of Cape Town

Dedicated to my Father, Mike Hyslop, who gave me every opportunity I could ever ask for. I love you and you will be sorely missed.

Chapter 1

Introduction

This thesis represents a collaboration on work done by Tom Leadbeater and Steve Peterson (the co-supervisors of this work) in the Physics Department at the University of Cape Town (UCT). The aim of this thesis is two-fold. The first is to investigate, characterise and improve the Positron Emission Particle Tracking (PEPT) technique using semiconductor detectors (traditionally PEPT is done with scintillators), building on Tom Leadbeater's research. The semiconductor detectors used were the Cadmium Zinc Telluride (CdZnTe or CZT) PolarisJ room-temperature semiconductor detectors build and sold by H3D Inc. (Ann Arbor, MI), purchased by Steve Peterson for his work in medical imaging. The second is to track a proton beam moving through an organic target using ideas and techniques from PEPT to improve the Distance-of-Closest Approach (DCA) filtering technique used in Prompt Gamma-ray Imaging (PGI), building on Steve Peterson's research as well as research done by the author in 2018 [1].

1.1 Dissertation Outline

This first Chapter provides both the necessary background motivating why this work was undertaken and the theory needed to understand the rest of the investigation and results presented in the following Chapters. It begins with the mathematical and practical principles needed to successfully implement the PEPT technique, after which it outlines the history of PEPT research and its applications. It then presents several models that characterise the effects of detector deadtime on coincidence processing. The Chapter ends with a description of Compton Camera imaging and the Distance of Closest Approach (DCA) filtering technique, outlining how PEPT can be used to track a proton beam incident on an organic target and how this is useful in improving the DCA algorithm. Chapter 2 follows up Chapter 1 by providing a technical description of all of the equipment used throughout this work, along with the procedures used to operate that equipment.

Chapters 3-5 present the full characterisation of a practical implementation of the PEPT technique in the CZT PolarisJ detector systems developed by H3D Inc. Chapter 3 starts with a proof-of-concept experiment to locate a ^{22}Na source. It presents the performance results and shows how the results were used to fine-tune the parameters required by the PEPT algorithm. It then ends with a measurement of the field-of-view sensitivity of the detector system. Chapter 4 explores semiconductor PEPT for a higher activity ^{68}Ga source, investigating the deadtime effects on the system and how deadtime influences location performance. From these results, an optimum source activity is decided on and the performance at this activity is presented. Chapter 4 ends with a discussion on how the performance could be further

improved.

Chapter 5 looks at tracking a moving tracer particle - the ultimate goal of PEPT. It presents the performance of a PEPT array that is larger than the UCT PolarisJ system: the H3D small-animal PET system. This system is representative of a complete CZT PEPT system that would be used for experiments beyond proof-of-concept, and Chapter 5 demonstrates the potential performance of such a system.

Finally, Chapter 6 presents an application of semiconductor PEPT: tracking a proton beam through an organic target to improve PGI performance by improving the DCA filter. Unfortunately much of the experimental work planned for this Chapter was interrupted by the COVID-19 pandemic, and so the results presented are primarily from Monte-Carlo simulations. However the preliminary results show promise and motivate for a more thorough investigation.

This thesis finishes with Chapter 7, presenting the final conclusions of this dissertation and the potential further work that could come from it.

1.2 Positron Emission Particle Tracking (PEPT)

1.2.1 Theory

All positron imaging techniques utilise β^+ radiation - the emission of a positron when a particular nucleus (e.g. ^{22}Na , ^{68}Ga , etc.) undergoes nuclear decay - and PEPT is no exception. If this decay occurs in a region of matter, this positron will slow down through electrostatic interactions and eventually annihilate with an electron. The rest mass of this electron-positron pair is then converted into two gamma-rays of 511 keV each, which are emitted back-to-back to conserve momentum. These two gamma-rays can then be detected in coincidence, and a line can be drawn between the two points of detection - known as a line-of-response or LOR. Therefore if the original position of the decay is not known, it is now constrained to lie on that line. Given a second decay in the same location with the back-to-back gamma-rays oriented in a different direction but in the same plane, the LOR generated by that pair will intersect with the first line, and the location of both decays is now constrained to sit on that intersection, and the previously unknown decay location has been determined. Extending this idea into three-dimensions, while the two LORs are highly unlikely to intersect, there will be a distance-of-closest-approach line that connects the two and both decays are now constrained to lie on this line. This process is shown in Figure 1.1. Thus, given a "point" β^+ radiation source that lies within the field-of-view (FOV) of a detector array detecting the 511 keV gamma-rays in coincidence, the 3D position of that source can be tracked as a function of time. This is the theoretical basis for PEPT. Note that in practice there will always be corrupted LORs, and so a statistical ensemble is required to accurately reproduce the location of the tracer.

A typical low-activity β^+ calibration source found in most nuclear labs has an activity on the order of 10 kBq, and so is giving off 10 000 of these annihilation gamma-ray pairs every second. Given sufficient detectors capable of detecting in coincidence, arranged in an array surrounding the source, it is in theory possible to locate the tracer source thousands of times per second. The frequency and accuracy with which the tracer is located is therefore limited by the activity of the tracer and




FIGURE 1.1: Demonstration of the positron annihilation process used for positron imaging. **Top:** A nucleus decays, emitting a positron which slows down and encounters an electron. **Middle:** The electron-positron pair annihilate, producing two back-to-back 511 keV gamma-rays, which are picked up by detector elements in coincidence on either side of the source. **Bottom:** These coincidences are collected and lines-of-response (LORs) are drawn in space, which converge on the source of the positrons, allowing for the tracking of the source through space. Note: this schematic is not to scale and in reality the annihilation occurs much closer to the source.

the capabilities of the detector system, which includes but isn't limited to: detector size, energy resolution, timing resolution and detector pixel size. Additionally, coincidence events can be corrupted through scatter (either with a part of the external system before entering the detector or with the detector material itself) or obscured by associated gamma-rays (e.g. the 1.3 MeV gamma-ray emitted by ^{22}Na), and so both the detector system and the processing algorithms need to account for this.

Radiation Detection in Coincidence

There are two different common techniques to detect coincidences. The first involves performing coincidence processing at the detector electronics level with the output being a list of coincidence events, which is typically how coincidence counting is implemented for dedicated coincidence detection systems such as those used in PET imaging. A full description of this process can be found at [2]. The second is to perform coincidence processing post-experiment with the raw list-mode data, which is particularly useful when developing a coincidence detection system as it allows for the re-processing of raw data with multiple processing parameters to find the best combination. This is how all the coincidence processing for this thesis was

performed. Coincidence processing is particular to the detector system being used, although the basic concepts are the same. The following is a description of how coincidence processing was implemented for the work done in this thesis using the PolarisJ detector array.

As with all coincidence processing, it is important that all the detectors in the array have their internal timing synchronised before the experiment proceeds. In the PolarisJ system one detector module is dedicated as the master and every two seconds it sends a synchronisation pulse to the other detector module(s) via coaxial cable. These sync pulses are recorded in each detector module's list-mode data-stream and used to combine the events in post-processing. In addition to a timestamp, the detector modules also record the event's energy deposition and the (x,y,z) position of the event in the detector's coordinate system. Before coincidence processing, the data from all the detector modules are combined into a single file, with all the events in order of the time they were detected in the lab time-frame. The events are then energy filtered (to see how these filters were chosen, see Section 3.1.2), after which the coordinates of each event are transformed from their detector module's coordinate system to the lab-frame - referred to as the isocentric coordinate system. Typically this is chosen with the origin at the centre of the FOV of the detector array. Once combined into a single event list, the events are processed for coincidences. Looping through each event, the algorithm searches the events before and after the current event to check if any events fall within a chosen time window from $-\tau$ to $+\tau$, which is known in the literature as 2τ . To see how this window was chosen, see Section 3.1.2. If no such event exists, the next event is visited. If there is an event in the time window, both events are stored together as a coincidence with the timestamp of the current event and removed from the event list. If more than one event is found, the event closest in time is chosen as the coincident event. This process continues until the end of the event list is reached. The output is a coincidence event file, which contains a timestamp of each coincidence and the (x,y,z) coordinate of both events (in the isocentric coordinate system).

The PEPT Location Algorithm

In practice, since no radiation source (often referred to as a tracer particle) is a true "point" source, and since typical tracer particles (^{22}Na , ^{68}Ga , ^{18}F , etc.) emit $\beta+$ radiation isotropically in three-dimensional space, a given sample of LORs almost never lie in the same plane. To solve this, first a number N_0 of these LORs are detected. Then a location in 3D space d_0 is found that minimises the sum of the squares of the perpendicular distances between the LORs and d_0 . From this sum D_0 , a mean distance is calculated: $m_0 = D_0/N_0$. Next, all the LORs whose distance to d_0 is further than some fixed multiple k of m_0 (i.e. $d_i > k \cdot m_0$) are discarded, leaving a subset of N_1 LORs. A new location d_1 is found from the remaining N_1 LORs from which a new (smaller) mean distance m_1 is calculated. Using the original set of N_0 LORs another subset of lines, N_2 , is found using d_1 . This process continues to iterate until a subset N_n is found containing a specified fraction f of the original number of LORs, i.e. $N_n = f \cdot N_0$. This f is referred to as the F-parameter, and is often used as a proxy for the signal-to-noise ratio. The k parameter determines the convergence rate, as it determines the rate at which the trajectories are discarded. A value of $k = 1$ means rapid convergence, but it discards many LORs at a time and so may overshoot the desired $f \cdot N_0$ number of LORs. For $k = 1.5$ only one or two LORs may be discarded at each iteration and so convergence is much slower. A value of

k somewhere between 1 and 1.5 is optimal. The final position d_n is then declared to be the best estimate of the position of the particle. A full description of the PEPT algorithm can be found in [3].

F-Optimisation and Location Uncertainty

Given a set N_0 of LORs used to locate a tracer, it is expected that a certain fraction of those lines are corrupt, e.g. a random coincidence, a gamma-ray that scattered before entering the detector, background radiation, etc. The remaining fraction are uncorrupted lines - e.g. true coincidence events that enter the detector directly from the tracer, so pass closest to the true location of the source. From the description above, the F-parameter represents the fraction of the LORs that are nearest to the best estimate of the location of the tracer. Therefore, in theory, there exists an F-parameter that eliminates all the corrupt lines while retaining all the uncorrupted lines. See Figure 1.2 for a visual depiction of this concept.

To find this optimal F-parameter, an F-optimisation curve needs to be produced. First, the tracer particle to be used in the experiment (at the activity to be used in the experiment) is placed inside the experimental system at the centre of the FOV of the detector array. A stationary measurement is taken for long enough to gather enough coincidences to produce a large number of tracer positions (preferably in the 1000s), keeping the number of lines per position N_0 fixed. Next, a range of F-parameters between 0.0 and 1.0 is created (usually in 0.05 steps) and for each F-parameter all the (x,y,z) tracer locations are determined using the PEPT algorithm. From this the standard deviation $\sigma_{x,y,z}$ is calculated. The standard uncertainty of the tracer location is defined as $u(x,y,z) = \sigma_{x,y,z} / \sqrt{N_f}$, where $N_f = f \cdot N_0$ is the number of lines used for the final PEPT location. Finally, this uncertainty is plotted as function of the F-parameter. Typically this looks like a U-curve with a F-parameter corresponding to a minimum uncertainty, often referred to as f_{opt} - the optimal F-parameter to minimise location uncertainty for the given system. This is because for F-parameters more than f_{opt} , corrupted lines are included in the location calculation, increasing $\sigma_{x,y,z}$ significantly. Once all the noisy lines have been eliminated by reducing the F-parameter, $\sigma_{x,y,z}$ doesn't decrease significantly for lower F-parameters. However, lowering the F-parameter further decreases $\sqrt{N_f}$, which increases $\sigma_{x,y,z} / \sqrt{N_f}$. Therefore f_{opt} is the best trade-off between removing all the noisy lines while retaining good statistics. See Figure 1.3 for an example F-optimisation curve. Naturally, increasing the starting number of LORs N_0 leads to better statistics, but at the cost of the location rate. Additionally, while the F-optimisation is performed on a stationary source, ultimately the tracer will be moving and so taking more LORs per position may decrease location uncertainty. For this reason, researchers are constantly looking to improve data rates. It is important to run an F-optimisation for every new experimental system and tracer particle used since the rate of random and corrupted coincidences is highly dependent on the experimental conditions. Once f_{opt} has been determined for a given setup it is then used for all subsequent runs.

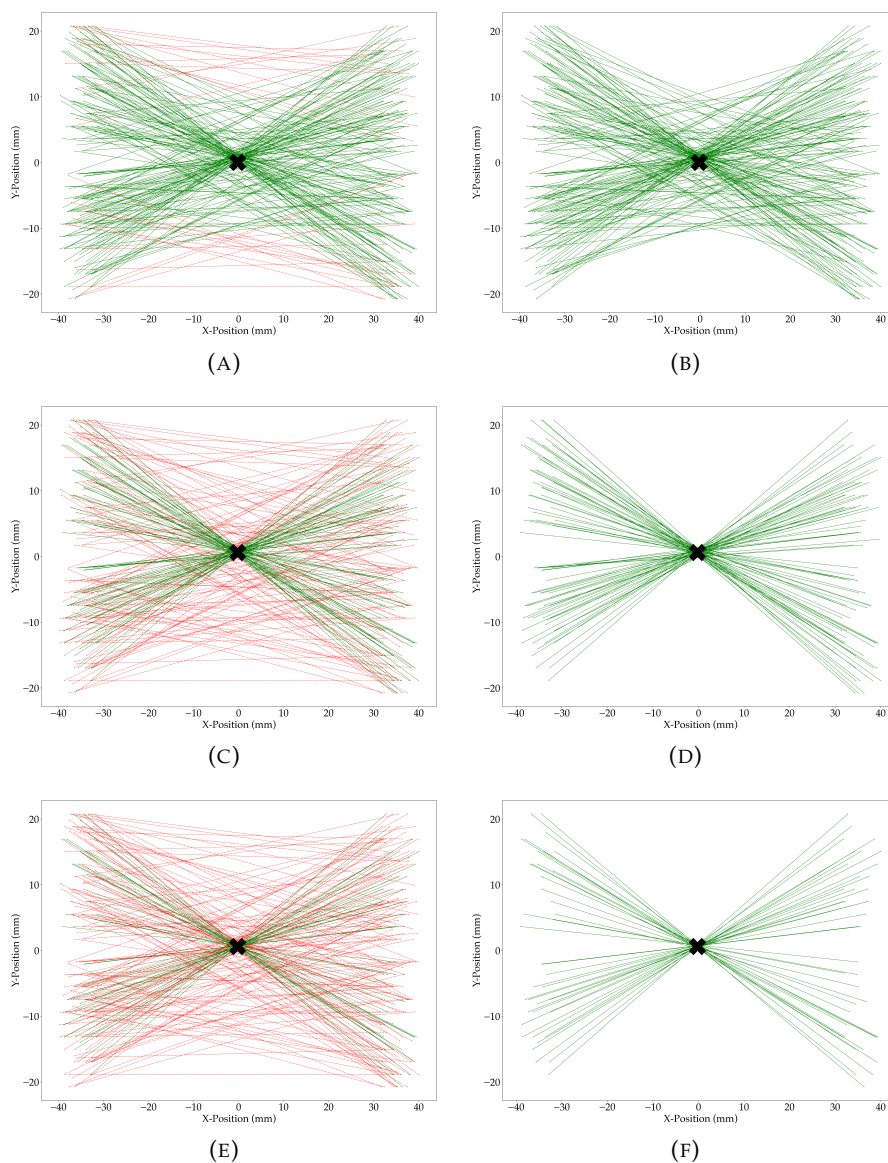


FIGURE 1.2: Demonstration of how the F-parameter operates. **(A)** 200 randomly selected LORs of a 1.5-1.8 MBq ^{68}Ga source in a CZT detector array. A large number of the LORs clearly converge onto the source (true coincidences), while many LORs are seemingly randomly distributed (noise). Shown as a black cross is the position that minimised the sum of the square of the perpendicular distances to each LOR. Using an F-parameter of 0.85, the 15% of LORs furthest from this position are highlighted in red. **(B)** The plot from (A) with the red LORs removed. Some of the noisy LORs have been removed but there are still a significant number left. **(C)** The same analysis from (A), but using an F-parameter of 0.38. **(D)** The plot from (C) with the red LORs removed. All the noisy LORs have been removed while the converging LORs remain. The F-optimisation process aims to find this optimal F-parameter that ensures all the signal is kept while retaining the most statistics as possible. **(E)** The same analysis from (C) but using an F-parameter of 0.2. This demonstrates how picking an F-parameter that is too low reduces the useful statistics. **(F)** The plot from (E) with the red lines removed. The LOR distribution is close to that in (D), but there are noticeably fewer LORs. **Note:** in practice this process is three-dimensional, and so lines that appear to converge in this 2D projection may in reality be very far from convergence.




FIGURE 1.3: An F-optimisation curve produced from a 14.7 kBq ^{22}Na source sitting in a CZT detector array, using $N_0 = 200$. The datapoints represent the radial uncertainty of the source position at a particular F-parameter. Fitted over the data is a polynomial to indicate the trend. The minimum sits at 0.8, indicating that F-parameters less than 0.8 reduce the total number of lines and therefore increase statistical uncertainty, and F-parameters greater than 0.8 include corrupted or random events and therefore increase distribution uncertainty.

1.2.2 PEPT: Birmingham (1988) to Cape Town (2020)

PEPT grew initially alongside Positron Emission Tomography (PET); researchers at the University of Birmingham were using PET as a non-invasive imaging technique to investigate optically opaque chemical and industrial systems, but were finding that the data rates were insufficient to investigate highly dynamical systems. Primarily this was because the tracer fluid used in PET was distributed and required many LORs to reconstruct an image. To solve this they switched to tracking a single "point" tracer, first proposed by Bemrose *et al.* in 1988 [4]. The PEPT algorithm described in the above section (as well as the term "PEPT") was then developed and first described by Parker *et al.* in 1993 [3]. In the same paper, Parker *et al.* showed the first application of PEPT. The detector array they used was a pair of (600×300) mm² multi-wire proportional chambers (MWPC) filled with isobutane and Freon gas, initially developed for PET imaging, with a maximum useful detection rate of 5 kHz. They were able to track slowly moving 4 MBq ^{22}Na and ^{18}F tracers to within 1 mm, several times per second, and tracers moving at 1 m/s to within 5 mm, 50 times per second. Their optimal F-parameter was between 0.15 and 0.35. Finally, they showed they were able to track a particle moving in a rotating drum and extract velocity fields, average axial, tangential and radial velocities and other useful properties of the system non-invasively.

Research at Birmingham continued, using the PEPT technique to investigate a wide variety of industrial processes, but was fundamentally limited by the capabilities of their detector. The next breakthrough came in 1999 when the University of Birmingham acquired a new positron camera made by ADAC Laboratories [5], the ADAC Forte. It consisted of two detector heads mounted on a motorised gantry with each head housing a solid NaI crystal ($590 \times 470 \times 16$) mm³ and an array of 55 photomultiplier tubes. The big advantage of this system was that each photomultiplier was connected to separate ADCs, allowing for fast, parallel event detection. With a resolving time of 7.5 ns this pushed the useful event rate up by a factor of 20-40, allowing for an order of magnitude improvement on position precision (0.6 mm) or the tracking of particles travelling an order of magnitude faster (over 10 m/s). Additionally, the optimal F-parameter increased to between 0.2 and 0.35, primarily due to the increased efficiency of NaI (25% intrinsic versus 5% for the MWPC) but also the improved timing and energy resolutions of NaI.

The next development in PEPT came when in 2002 the University of Birmingham received the ECAT931 series made by CTI Inc., originally designed for medical PET applications from Hammersmith Hospital. This system contained 2 rings of 16 detector "buckets" with each bucket containing 4 detector blocks and associated electronics. These ($50 \times 50 \times 30$) mm³ mm detector blocks were made of Bismuth Germanium Oxide (BGO) and subdivided into an 8×4 array, with the full system containing 128 detector blocks totalling 4096 BGO elements. Birmingham researchers then broke down the rings into their component parts (blocks), which gave them the advantage of modularity and allowed the investigators to arrange the blocks in the optimal geometry for the system being investigated, as described by Sadrumont *et al.* in 2007 [6]. With this system, they were able to track tracers to an accuracy of 1.2 mm at rates of 410 Hz, allowing for tracking of particles moving at speeds up to 10 m/s [7], with an optimal F-parameter of 0.4. In the period between 2002 and 2008, Birmingham acquired four complete PET scanners, some of which were broken down and some of which were kept as intact rings, allowing for a very wide variety of applications. In addition to the detector development, Birmingham researchers were also improving positron emitting tracer particles [8], further improving the capabilities of PEPT and the types of systems they could investigate. Between 2002 and 2012 researches at Birmingham used PEPT to do invaluable work in industrial process analysis [9, 7, 10, 11]. By 2012, the Birmingham BGO detector arrays were able to track small (≈ 1 mm) tracers in 3D with an accuracy of 0.5 mm, locating the particle every millisecond [12].

Up until 2009 Birmingham was the only PEPT facility in the world - the requirement for large pixelated PET-type detector arrays along with a facility to produce the tracer isotopes made it an expensive investment to build a PEPT laboratory. However, in 2009, Imperial College London and Hammersmith Hospital donated an ECAT 'EXACT3D' (Model: CTI/Siemens 966) HR++ [13] ring-geometry PET camera to iThemba Labs, a cyclotron research facility in Cape Town, South Africa [14]. The HR++ consisted of 48 rings of BGO detector crystals, each ($4.39 \times 4.05 \times 30$) mm³, arranged in 8×8 blocks. The ring diameter was 82 cm, the axial field-of-view was 23.4 cm and the acquisition system was capable of a coincidence rate of 4 MHz - this compared to the state-of-the-art at Birmingham with a coincidence rate on the order of 100 kHz. iThemba Labs was an ideal location for the second operational PEPT facility worldwide as it was already a fully licensed isotope production facility, and

one of the problems facing PEPT at this time was producing short-lived tracer particles (^{18}F , ^{68}Ga , etc.) that were small enough ($< 20 \mu\text{m}$) to investigate liquid-phase systems, such as floatation cells, but also active enough ($> 10 \text{ MBq}$) for accurate tracking. Thus PEPT Cape Town was established and has been performing PEPT research since mid-2009 to the present day. The primary isotope used at PEPT Cape Town is ^{68}Ga and there has been active research in making smaller and more active tracers, with tracers as small as $50 \mu\text{m}$ being tracked in the HR++ [15]. The data acquisition system used in the HR++ has also seen great improvements, spearheaded by Tom Leadbeater [16] who originally worked in the Birmingham lab and moved to the University of Cape Town to head up PEPT Cape Town. At the time of writing, the HR++'s coincidence rate of 4 MHz enables it to routinely track a 1 mm tracer particle moving at 1 m/s to within a 1 mm uncertainty at a location rate of 1 kHz, and has regularly performed even better. The HR++ also uses optimal F-parameters as high as 0.7 [17], showing a doubling of the signal-to-noise ratio compared to the modular systems used at Birmingham. Research continues to try and reduce the tracer size and push the boundaries of the application of PEPT to new processes and systems.

1.2.3 Tracer Production

As with all positron imaging, it is useful to be able to produce short-lived (half life ≈ 1 hour) isotope tracer particles when running a PEPT measurement. It avoids the need to recover the tracer from the system afterwards and is particularly useful for very small ($\approx 50 \mu\text{m}$) tracers in multiphase systems where it may not be possible to recover the tracer. It also allows for the production of tracers that are customised to the system being studied. Ideally, tracer production should allow the investigator to select the desired activity and physical and chemical properties (making it hydrophilic for example) for each tracer used, which is then discarded safely once it has decayed down to background levels. Additionally, because PEPT primarily investigates industrial processes, having a high activity source in the system of interest does not appreciable damage that system (compared to medical imaging for example) and therefore very high activity tracers are used to maximise the statistics. Production yield of a particular isotope is proportional to the inverse-square of the half-life, so short-lived isotopes can be produced with much higher activities than long-lived ones, and are more easily licensed. Therefore a consistent and reliable isotope production facility that can produce high-activity sources are critical to the success of any PEPT laboratory.

Since the PEPT technique developed in parallel with PET imaging (Section 1.2), so too did the development of producing PEPT tracers grow in parallel with that of PET tracer fluid. The tracers used in PET consist of distributed fluid, typically made of substances intended to mimic compounds already present in the body so that the tracer fluid will follow the fluid flow of a particular bodily process. A popular isotope used is ^{18}F ; it has a short half-life of 109 minutes, it is a pure 511 keV emitter (it only produces 511 keV gamma-rays, unlike isotopes such as ^{22}Na) and it is easily substituted directly into biomolecules [18]. ^{18}F for medical PET is typically produced through the bombardment of protons on a water target enriched with ^{18}O through the $^{18}\text{O}(\text{p},\text{n})^{18}\text{F}$ reaction, and is used by the Birmingham PEPT lab. Additionally, Birmingham produce tracers through the reactions $^{16}\text{O}(\text{}^3\text{He}, \text{p})^{18}\text{F}$ and $^{16}\text{O}(\text{}^3\text{He}, \text{p})^{18}\text{N} \rightarrow ^{18}\text{F}$ under direct bombardment from a cyclotron beam to activate solid tracers. Initially they used the 33 MeV ^3He Birmingham Radial Ridge

Cyclotron, but this was decommissioned in the 1990s and replaced in 2003 by the Scanditronix MC40 cyclotron. For tracers larger than 1 mm in size, either direct activation of the tracer or producing ^{18}F in water and adsorbing it into the tracer [8] is used. To produce tracers smaller than 1 mm, researchers turned to using an ion exchange resin to produce a tracer. Using a strong-base anion exchange resin for uptake of ^{18}F from a water target, they were able to produce tracers as small as $6\ \mu\text{m}$ with activities of 1.1 MBq [8].

One of the disadvantages of using ^{18}F as a tracer is the requirement of a cyclotron to produce them, and so an alternative is ^{68}Ga produced using a $^{68}\text{Ge}/^{68}\text{Ga}$ generator. Like ^{18}F , ^{68}Ga has a short (68 minutes) half-life and it is a pure 511 keV emitter. Tracers are produced by one of two ways: for larger tracer particles ($> 3\ \text{mm}$), a hole is drilled into the particle, the ^{68}Ga is placed inside and the hole is sealed - the so-called "drill-and-fill" method. For smaller particles an ion-exchange resin is used to uptake the ^{68}Ga , which can then be coated if required to use as a tracer particle. This tracer production technique has been adopted at the PEPT Cape Town facility at iThemba labs, which has eliminated the need for beam-time on the iThemba Labs cyclotron whenever a tracer is required. iThemba labs is also a producer of SnO_2 $^{68}\text{Ge}/^{68}\text{Ga}$ generators, making this a convenient production method. However, ^{68}Ga has its downsides and so beam-activated PEPT tracers are being actively researched. PEPT Cape Town has also become a lead researcher in producing smaller and more versatile PEPT tracers. This has been spearheaded by Mike van Heerden, who has done extensive work in manipulating the chemical properties of the tracer resin to be more hydrophobic/hydrophilic [19], in producing tracer particles as small as $50\ \mu\text{m}$ with activities between 0.05-0.5 MBq to be tracked in UCT's HR++ [15] and in removing impurities for greater uptake of ^{68}Ga in the exchange resin [20].

1.3 Deadtime Effects on Detected Singles Event Rates

Like all nuclear detector systems the PolarisJ detector modules are susceptible to deadtime effects, which have a significant influence on the maximum detectable coincidence rates and therefore the PEPT performance of the system. It is therefore important to be able to characterise these effects and so a mathematical model is presented here to describe the mechanisms that contribute to this effect [21].

Deadtime in a detector is the time, τ , after an event has interacted with a detector during which another event cannot be detected; this is caused by many different facets of the detector system. In a semiconductor detector, an event is detected when a radiation particle (gamma-ray, alpha/beta particle, etc.) interacts with an electron in the detector, transferring its energy (or some of its energy) to that electron, which then creates a shower of electron-hole pairs that have been excited into the conduction band of the semiconductor detection medium. An applied voltage then separates these electron-hole pairs and collects the electrons at an anode and the holes at a cathode. This induces a voltage across the anode and cathode proportional to the energy deposited by the radiation particle, which is then processed by the detector electronics to provide and store the event information. This whole process takes a finite amount of time, during which the detector may or may not be able to process another interaction event. To describe what happens if another event does interact with the detector during this time, there are two models of deadtime: non-paralysable and paralysable.

1.3.1 Non-Paralysable Deadtime

Suppose an event interacts with a detector and that detector goes dead for time τ_{np} . In the non-paralysable deadtime model, the detector cannot detect any events during τ_{np} , but will be available to detect another event the moment τ_{np} has elapsed, regardless of whether or not another event arrives during τ_{np} . It is also assumed that another event arriving during τ_{np} won't have any affect on the detector output; it is simply lost. Suppose that the true interaction rate with the detector is n , and the measured event rate is m . The rate of lost events is given by $n - m$. Additionally, if we know τ_{np} then the fraction of time the detector is dead is given by $m\tau_{np}$, and so the rate of lost events is $nm\tau_{np}$. Equating these expressions:

$$n - m = nm\tau_{np} \quad (1.1)$$

$$n = \frac{m}{1 - m\tau_{np}} \quad (1.2)$$

Expression (1.2) therefore gives the true interaction rate given the measured rate and τ_{np} . However, to model deadtime effects one starts with a known source activity and experimental geometry meaning that n is known, and so (1.2) is rearranged to give:

$$m = \frac{n}{1 + n\tau_{np}} \quad (1.3)$$

This provides the expected measured rate m for a particular n and τ_{np} , allowing one to model the deadtime behaviour.

1.3.2 Paralyisable Deadtime

As before, suppose an event interacts with a detector and that detector goes dead for time τ_p . In the paralyisable deadtime model there are two possible cases. In the first case there are no additional event interactions during τ_p , in which case the detector behaves as in the non-paralyisable case: it records the event information and is available to detect another event after τ_p . In the second case, a second event interacts with the detector during τ_p , let's say at time $\tau_p/2$. In this case, this causes the detector to remain dead for a further time τ_p from the moment the second event arrives and is not able to record the event information of the second event. So instead of going dead for a time τ_p , it now goes dead for a time $1.5\tau_p$, unable to record any events during this time. If a detector experiences paralyisable deadtime, there will be some maximum recordable event rate at some incoming radiation flux, but if the radiation flux increases the recorded event rate will decrease as deadtime effects compound on each other. In this case there would exist an incoming radiation flux high enough to cause the recorded event rate to be very close to zero, since events would be arriving so often the detector would never get a chance to recover.

As before, suppose the measured rate in the detector is m . This rate is not necessarily equal to the rate of true events that arrive with a time interval more than τ_p apart. As radioactive decay is a Poisson process, the distribution of time intervals between events is given by

$$P_1(T)dT = ne^{-nT}dT \quad (1.4)$$

where n is the true event rate as before and $P_1(T)dT$ is the probability of observing a time interval between events within dT about time T . The probability of observing a time interval between events greater than τ_p is thus

$$P_2(\tau_p) = \int_{\tau_p}^{\infty} P_1(T)dT = e^{-n\tau_p} \quad (1.5)$$

The measured rate is therefore this probability multiplied by the true interaction rate n

$$m = ne^{-n\tau_p} \quad (1.6)$$

1.3.3 Measured Rate from a Single Detector detector module

In practice all detectors experience a combination of paralyisable and non-paralyisable deadtime, typically with different time parameters τ_{np} and τ_p . Non-paralyisable deadtime typically arises from the signal-processing electronics; during the time it takes to process an event signal the electronics may discard any additional signals, but will be available immediately after the event has been processed. Paralyisable deadtime typically arises from the physical interactions with the detection medium; if multiple events excite electron-hole pairs within a semiconductor crystal before the pairs have been collected at the anode and cathode, the induced voltage will be some unknown combination of all the events and therefore the recorded data will be useless until all the electron-hole pairs have been collected.

Shown in Figure 1.4 is the typical detector response curve for measured event rate m as a function of true rate n (or activity of the source given a fixed detector

geometry). With zero deadtime the response is linear, with $m = n$. In the non-paralysable case the measured rate will asymptotically approach a maximum event rate of $1/\tau_{np}$ where events are, on average, arriving straight after the deadtime period has elapsed. In the paralysable case the measured rate increases to a maximum value of $1/\tau_p e$, whereafter it begins to decrease as the deadtime effects begin to pile-up on each other.




FIGURE 1.4: Typical detector response curve for measured event rate m as a function of true rate n . With zero deadtime the response is linear, with $m = n$. For non-paralysable deadtime the measured rate will asymptotically approach a maximum event rate of $1/\tau_{np}$. For paralysable deadtime the measured rate increases to a maximum value of $1/\tau_p e$, whereafter it begins to decrease.

1.3.4 Measuring the Deadtime Parameters

The best way to measure the deadtime parameter of a detector module is perform an experiment where the measured rate is recorded as a function of the true rate. A short-lived radioisotope is placed in front of the detector to provide a true rate well above the $1/\tau_{np}$ or $1/\tau_p$ values and is then observed over several half-lives until the detector response is in the linear response region - naturally this takes some trial-and-error. To extract the τ_{np} or τ_p parameters, equations (1.3) and (1.6) need to be linearised.

Non-Paralysable

The true rate n can be expressed in terms of the experimental setup as

$$n = A\gamma\epsilon\Omega \quad (1.7)$$

Where A is the activity of the source, γ is the branching ratio of the gamma-rays emitted by the source, ϵ is the intrinsic detector efficiency and Ω is the fractional solid angle subtended by the detector. From (1.2) we know n in terms of m . Equating the two

$$A\gamma\epsilon\Omega = \frac{m}{1 - m\tau_{np}} \quad (1.8)$$

$$\frac{m}{A} = (-\tau_{np}\gamma\epsilon\Omega)m + \gamma\epsilon\Omega \quad (1.9)$$

The quantity m/A is known as the Activity Function, so plotting this as a function of m gives a linear plot with gradient $-\tau_{np}\gamma\epsilon\Omega$ and y-intercept $\gamma\epsilon\Omega$. Performing a least-squares linear fit to get the slope and intercept, τ_{np} is thus $\tau_{np} = -\text{gradient}/\text{intercept}$. Additionally, any of the quantities γ , ϵ or Ω can be determined or double-checked using this fit.

Paralysable

From (1.6), we have

$$m = ne^{-n\tau_p} \quad (1.10)$$

Note that n is a function of time

$$n = n_0e^{-\lambda t} \quad (1.11)$$

where n_0 is the initial true rate and λ is the nuclear decay constant of the isotope in question. Substituting into the above expression

$$m = n_0e^{-\lambda t - n_0\tau_p e^{-\lambda t}} \quad (1.12)$$

$$\ln(m) = \ln(n_0) - \lambda t - n_0\tau_p e^{-\lambda t} \quad (1.13)$$

$$(\ln(m) + \lambda t) = (-n_0\tau_p)e^{-\lambda t} + \ln(n_0) \quad (1.14)$$

Here we have $(\ln(m) + \lambda t)$ as the Activity Function, so plotting this as a function of $e^{-\lambda t}$ provides a linear function with gradient $-n_0\tau_p$ and y-intercept $\ln(n_0)$. Finally, the deadtime parameter is given by

$$\tau_p = \text{gradient} e^{-\text{intercept}} \quad (1.15)$$

1.4 Deadtime Effects on Coincidence Counting

As described in Section 1.2.1, coincidence are found in the PolarisJ datastream using pulse-train analysis where filtered events from opposing detector modules are analysed to see if they're detected with a 2τ window of each other. Section 3.1.2 outlines how this 2τ window is selected, which for the UCT PolarisJ system is 600 ns. It is useful to have a model based on the known parameters of an experiment to be able to predict the expected measured coincidence rates to sanity-check the actual measured rates in a given system, particularly with regard to the deadtime effects on coincidence counting.

There are several different types of detected coincidence events which require definition:

True Coincidences

True coincidence pairs are two photons that correspond to a genuine annihilation event from the source. These coincidences are the useful events required for positron imaging and are the "signal" LORs in PEPT. True coincidences are proportional to the activity of the source and the fraction of emitted annihilation photons that are detected in true coincidence is a function of the system deadtime.

Random Coincidences

Random coincidences are coincidence events that don't correspond to any genuine annihilation event in the source. These are caused by the detection of multiple annihilation events from the source within the chosen timing window and therefore become increasingly significant with increasing source activity. It is not possible to know *a priori* whether a detected coincidence event is a True or Random (or even general background radiation) - this is the aim of the F-parameter in the PEPT algorithm.

Prompt Coincidences

These are all the events measured in opposing detectors within the chosen energy and timing windows, and so are the sum of the True and Random coincidences. After running coincidence processing on the data from a given experiment, the resulting coincidence rate is the Prompt coincidence rate.

Delayed Coincidences

Delayed coincidences are used to estimate the random coincidence rate. Events from one detector module in an opposing pair have a fixed time-delay added to all the event timestamps, after which the data are searched for coincidences. Delayed coincidences are the pairs of events detected within the chosen energy and timing windows, but do not correspond to any genuine annihilation event in the source. They thus represent a statistical sampling of the Random coincidences. To estimate the number of True coincidences for a given experiment, both the Prompt and Delayed coincidences are measured and the True coincidence rate is quoted as the difference (True = Prompt - Delayed).

Multiple Coincidences

Multiple coincidences are when more than two events are detected in opposing detector modules within the chosen energy and timing windows. There are therefore more than one possible pair of events that are in coincidence. This becomes significant when source activity and detector sensitivity is high enough such that multiple annihilations are detected within the timing window. One of these possible pairs is likely to be a True coincidence and the rest are Randoms, but it isn't possible to determine the True pair during coincidence processing. Typically one of these pairs is chosen at random and the rest are discarded, but with list-mode and the PEPT algorithm an alternative approach is to include all possible permutations and rely on the PEPT algorithm to remove the erroneous lines.

Total Coincidences

This is simply the total number of coincidence events after coincidence processing has been applied to a given dataset. It is typically made up of the Prompt (the True plus Random) and Multiple coincidences.

With these definitions in place, the coincidence models can be developed.

1.4.1 True Coincidence Rate Model

Since the True coincidences correspond to genuine pairs of annihilation photons from the source, it is a function of the activity, branching ratio, fractional solid angle and detector efficiency for a given experiment. The absolute rate of emitted True coincidences is

$$r_T = \frac{A\gamma}{2} \quad (1.16)$$

where A is the source activity, γ is the branching ratio for annihilation photons and the factor of two is the approximate value of the positron multiplicity distribution function (a measure of the number of photons produced per annihilation, which is slightly higher than two but rounded off for this analysis). The True measured coincidence rate (assuming no deadtime), n_T , is therefore the product of the emitted rate and the probability that both annihilation photons are detected

$$n_T = A\gamma(\epsilon_1\Omega_1)(\epsilon_2\Omega_2)W(\theta) \quad (1.17)$$

where $\epsilon_{1/2}$ are the intrinsic efficiencies of each detector module and $\Omega_{1/2}$ are their respective fractional solid angles. $W(\theta)$ is the angular correlation factor between photons and is often defined based on the detector module positions. For example, suppose the first annihilation photon is emitted in the direction of one of the detector modules. $W(0^\circ)$ is the relative yield of the second photon per unit solid angle about 0° relative to the first photon's direction. To simplify this expression, we assume that two identical detector modules are placed directly opposite each other, and the source is placed perfectly central in the FOV of the detector modules. Then $\epsilon_1 = \epsilon_2 = \epsilon$ and $\Omega_2 = W(\theta) = 1$ since annihilation photons are emitted approximately back-to-back, therefore any photon pair that strikes the first detector module must necessarily strike the opposing one. n_T then collapses to

$$n_T = A\gamma(\epsilon\Omega_1)(\epsilon) \quad (1.18)$$

Note that $A\gamma(\epsilon\Omega_1) = n_S$ is the true singles interaction rate in one of the detector modules. Since this value is often known, the True coincidence rate can be written in terms of the singles rate in one detector module

$$n_T = n_S\epsilon \quad (1.19)$$

If the detector system is experiencing deadtime losses, this can be modelled as a reduction in the probability that both photons are detected. Therefore we introduce m_T , the rate at which True coincidences are measured while detectors experience deadtime. Deadtime introduces additional factors into (1.17), which becomes

$$m_T = A\gamma(\epsilon_1\Omega_1f_{DT1})(\epsilon_2\Omega_2f_{DT2})W(\theta) \quad (1.20)$$

where $f_{DT1/DT2}$ are the respective deadtime factors for each detector module. Given the constraints of identical opposing detector modules with a central source, this becomes

$$m_T = A\gamma(\epsilon\Omega_1 f_{DT})(\epsilon f_{DT}) \quad (1.21)$$

which can be written in terms of the measured deadtime-limited singles event rate in one of the detector modules

$$m_T = m_S(\epsilon f_{DT}) \quad (1.22)$$

1.4.2 Random Coincidence Rate Model

The rate of uncorrelated, random coincidences in a detector module array is the rate of events in one detector module times the probability of observing another event in an opposing detector module within the chosen energy and timing windows. As discussed in Section 1.3.2, the probability of observing a time-interval of dT about T between events in a detector module is given by $re^{-rT}dT$ where r is the singles rate in the detector module. Given two opposing detector modules with rates r_1 and r_2 , the rate of observing an event in detector module 2 an interval dT about T after seeing an event in detector module 1 is the rate in detector module 1 multiplied by the probability: $r_1r_2e^{-r_2T}dT$. In a typical PEPT experiment the event rates are several orders of magnitude smaller than the reciprocal of the timing window, so $e^{-r_2T} \approx 1$ and the rate is approximately constant: r_1r_2dT . Therefore the rate of observing two uncorrelated events in opposing detector modules in an interval 2τ , otherwise known as the random coincidence rate n_R , is

$$n_R = n_{S1}n_{S2}2\tau \quad (1.23)$$

Assuming as before identical detectors with a central source, the two detector modules will have the same event rates and this becomes

$$n_R = 2\tau n_S^2 \quad (1.24)$$

In the case of deadtime-limited detector modules we simply put in the deadtime-limited singles event rates, giving

$$m_R = 2\tau m_S^2 = 2\tau(n_S f_{DT})^2 \quad (1.25)$$

where $f_{DT} = 1/(1 + n_S\tau_{np})$ for the non-paralysable case (1.3) and $f_{DT} = e^{-n_S\tau_p}$ for the paralysable case (1.6).

1.4.3 Prompt Coincidence Rate Model

The prompt coincidences is simply the sum of the Trues and Randoms ¹

$$n_P = n_T + n_R \quad (1.26)$$

In reality there will be deadtime-limiting and so the deadtime-limited version is more often used.

¹There will also be multiple coincidences, but modelling them is significantly more complex and the event rates seen in this thesis are low enough for them to form a very small contribution to the Prompt rate and are assumed to be negligible.

$$m_P = m_T + m_R \quad (1.27)$$

This Prompt rate m_P is what is measured when coincidence counting is performed on a given dataset. As discussed in Section 1.4, the Delayed coincidences are a statistical sampling of the Random rate, i.e. $m_R = m_D$. Since the data from the PolarisJ is recorded in list-mode, the Delayed rate can be measured after the experiment, and so the underlying measured True rate is quoted as the difference: $m_T = m_P - m_D$.

It is worth noting that for dedicated positron imaging systems where only the coincidence data are recorded, the acquisition of the Prompt and Random rates contribute an additional data acquisition deadtime. For list-mode acquisition for most of the rates being measured in this dissertation, this is not an issue. However, for the higher activity measurements (and certainly for the rates measured by the HR++) this does contribute significantly. Modelling this contribution for PolarisJ is beyond the scope of this work.

1.4.4 Model Predictions

Assuming a pair of identical detector modules facing perfectly opposite each other with a central source we can predict the system response for the various different coincidence rates. First, assume that there is zero deadtime in either detector module, i.e. the singles rate scales linearly with source activity. The coincidence rate curve predictions as a function of source activity are shown in Figure 1.5.

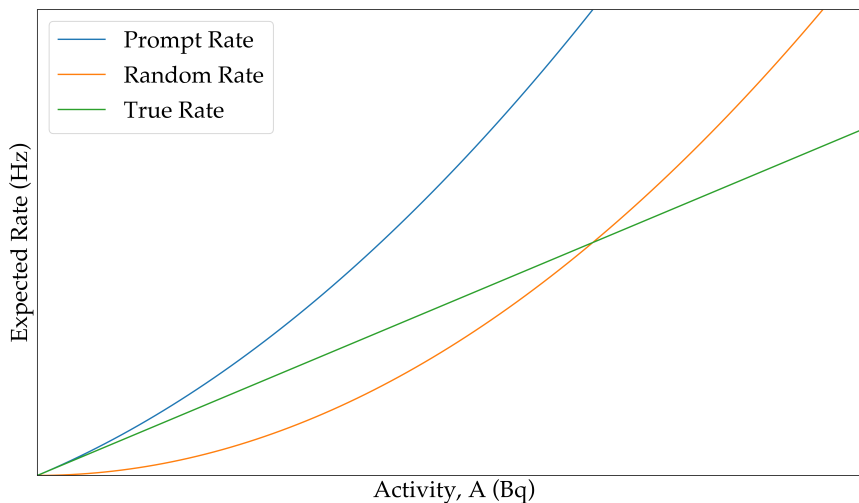


FIGURE 1.5: The Prompt, Random and True coincidence rates as predicted by their respective coincidence models with zero deadtime as a function of source activity.

Important to note here is that the Random rate scales as n_S^2 , while the True rate scales as n_S . Therefore as the source activity increases, the signal-to-noise ratio will deteriorate even if there is no deadtime-limiting.

Now to include deadtime effects; instead of assuming that the singles scale linearly with activity, it is assumed that the singles rate follows the non-paralysable curve presented in Figure 1.4. Now the coincidence curves are predicted to look like Figure 1.6.

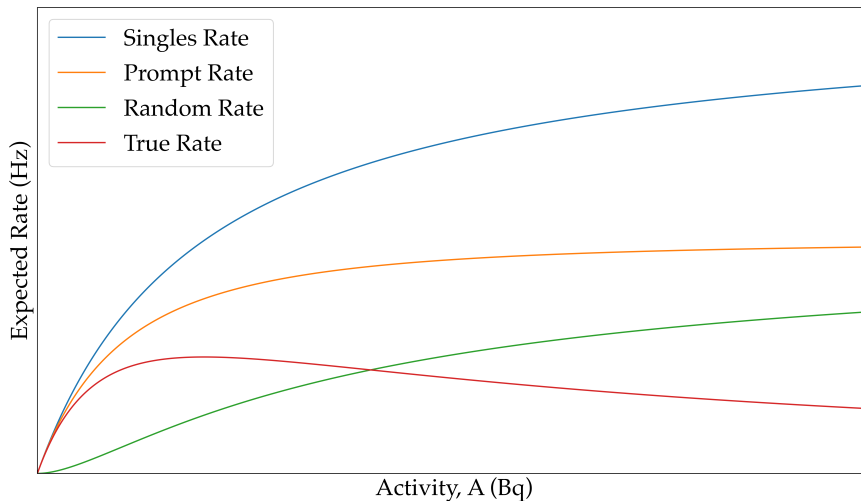


FIGURE 1.6: The Prompt, Random and True coincidence rates as predicted by their respective coincidence models including a non-paralysable deadtime model on the singles rate, as a function of source activity.

Unsurprisingly, the Prompt rate follows the general shape of the singles rate as source activity increases, asymptotically approaching some maximum rate. As with the no-deadtime case, the Random rate increases faster than the True rate and decreases the signal-to-noise ratio with increasing activity. The Random rate will also start to directly affect the True rate as deadtime-limiting will prevent the recording of genuine coincidence pairs that would have otherwise been detected. If the system then begins to experience significant paralysable deadtime, the singles rate will begin to decrease with higher activity and so too will the coincidence rates. There is therefore a theoretical optimal source activity for a given experimental geometry which maximises the True coincidence rate. Traditionally, this is measured experimentally by placing a short-lived radioisotope central to the complete detector system and observed over several half-lives to produce the curves shown in Figure 1.6. For the case of two or more pairs of detectors, all the rates are multiplied by the number of detectors in the system.

1.5 Extending the Positron Emission Particle Tracking (PEPT) Technique

Throughout the development of PEPT the primary detectors used were scintillators, namely NaI and BGO. Scintillators are large area detectors with high efficiency and fast timing resolution, with coincidence timing resolution of 7 ns for the ADAC Forte and 4-6 ns for the HR++ that allow for coincidence rates in the MHz and relatively low random coincidence rates. Additionally, towards the end of the 20th century PET-CT scanners were coming to market, offering much better imaging properties than straight PET scanners prompting hospitals around the world to replace their PET scanners with PET-CT ones. This meant the procurement of large, working, second-hand positron cameras became possible for PEPT research instead of building them from scratch, which greatly accelerated the development of the technique. However, these scintillators have relatively poor energy resolution (BGO has a FWHM of 170 keV at 511 keV) meaning that corrupted LORs are not eliminated effectively from the coincidence data. Secondly, the pixel size of the BGO blocks is limited to (4×4) mm² and doesn't have depth-of-interaction information, which geometrically limits the accuracy at which the LORs can be drawn. Traditionally these limitations have been very successfully addressed with statistics, focussing primarily on increasing the efficiency for 511s but also increasing the data-acquisition rates and tracer activities to reduce the position uncertainty.

However, as tracers are required to be smaller and smaller to investigate new processes (particularly liquid-phase) the maximum tracer activity starts to become limited. This pushes activities to far lower than optimal for scintillator systems, lowering data rates to almost unusable levels. This is particularly true for fixed-ring geometries where the solid angle cannot be increased by reducing the distance between detector elements, so while intrinsic efficiency is high, absolute efficiency is relatively low. Accurate tracking has proven difficult with very small ($< 100 \mu\text{m}$) particles, although active research is being done to push the boundaries of activity at this scale. Nonetheless, the historical approach of improving efficiency to maximum the number of useful lines per unit time was therefore not proving sufficient to track the ever-smaller tracers. This has led to the question of finding a new approach to tracking small particles.

Semiconductor detectors offer a possible solution. There has been increased development into room-temperature semiconductors such as Cadmium Zinc Telluride (CZT), which offer both high energy resolution (6 keV at 511 keV) and small pixel size ((1.8×1.8) mm²) with depth of interaction information and have been successfully used in Compton Camera applications (see Section 1.6). Theoretically, this high pixel resolution implies that accurate positions can be measured using PEPT with lower data rate requirements and therefore smaller tracer particles. However, these detectors have a lower intrinsic efficiency, detector size and timing resolution compared to scintillators and so lower data rates are expected from them. Whether the trade-off between resolution and data rates will prove effective is one of the central tenants to this thesis. Part one of this thesis aims to characterise the performance of PEPT using a CZT camera developed for Compton Camera imaging by H3D Inc. (Ann Arbor, MI).

Characterising PEPT using a new detector technology requires a deep dive into

the capabilities of that detector. Starting with a single detector module, the following needs to be described: energy resolution, timing resolution, detector pixel resolution, detector intrinsic efficiency at 511 keV, deadtime effects and how they impact the maximum possible event rate, and data-flow (i.e. how events are recorded and extracted from the detector for analysis). Next, two detectors need to be arranged in an array to be able to detect coincidences, measurements of stationary sources need to be taken and the coincidence performance needs to be quantified - namely total and random coincidence rates. Deadtime impact on these rates and response to different source isotopes also needs to be described. These data then needs to be fed into the PEPT algorithm and positions measured and compared to known source positions. The positional uncertainties need to be quantified and F-optimisation curves need to be produced to optimise PEPT performance. Moving sources can then be observed and PEPT performance can be further analysed, setting limits on tracer speeds at various tracer activities. Finally, all these processes need to be optimised to maximise the performance PEPT capabilities. The first part of this thesis lays out the full process of doing this analysis, and presents results for the PolarisJ CZT detector modules, comparing them to the performance of traditional PEPT using scintillators.

1.6 Prompt Gamma-ray Imaging (PGI)

1.6.1 Proton Radiation Therapy

Perhaps one of the hardest problems in the treatment of cancer is that of recurring tumours and metastasis. It is not uncommon after treatment and an initial remission of a tumour to see the cancer regrow and spread throughout the body. Worse, often the relapse is more aggressive and therapy resistant, so even harder to treat. The reason for this seems to be that even if only 1% of cancer cells survive treatment, by the Darwinian principle they will be the most therapy resistant cells. There is also evidence to suggest that they suffer cytotoxic stress and thereby enter a stem-cell-like state - the seeds for recurrence [22]. Therefore for any cancer therapy, it is imperative that 100% of the tumour is destroyed - 99% just isn't good enough.

Since its invention in the late 1800s, X-ray radiation therapy has been an integral part of cancer treatment and in part earned its discoverer, Wilhelm Conrad Röntgen, the very first Nobel prize in 1901 [23]. Since then the technology has progressed greatly, particularly with the advent of better imaging techniques like CT and PET scanning along with the development of intensity-modulated radiation therapy (IMRT), which allows for finer control of the beam. However, it is also well-known that prolonged exposure to radiation can cause cancer [24] and so it is important that radiation therapy minimises the dose to healthy tissue. X-rays are attractive because of their ability to penetrate through dense material and hence allows for treatment of cancers within the body without surgically revealing the tumour. However, the energy deposition of uncharged radiation into human tissue follows an exponential decay, which leads to a significant dose into the healthy tissue before the tumour as well as an exit dose (albeit not as severe). In 1946, R.R. Wilson proposed the use of high-energy charged radiation, protons, to solve this problem [25]. From the well-known Bethe-Bloch equation, protons lose energy as they move through a material like the inverse square of their velocity. As the proton enters a material it initially gives off only small fractions of its energy, and due to its large mass compared to the electron (the primary particle of interaction when

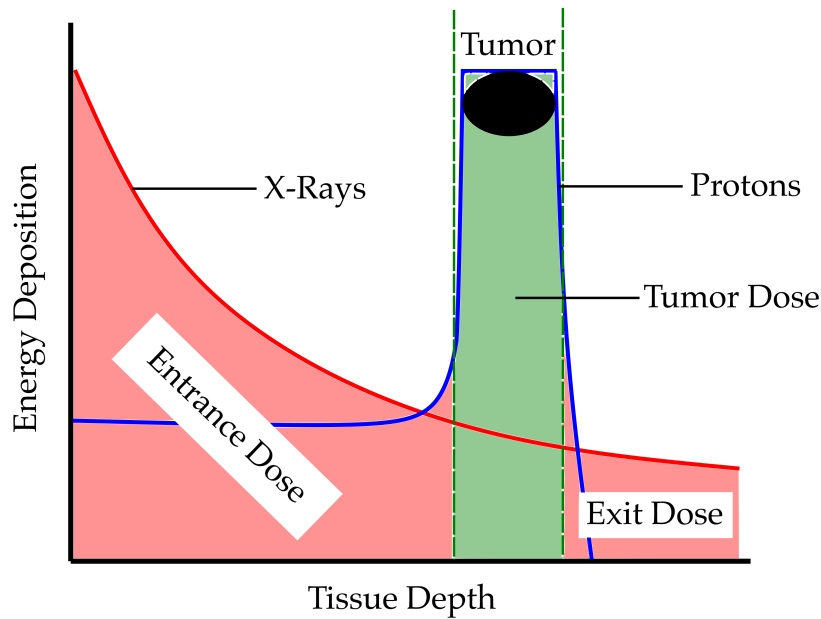


FIGURE 1.7: Use of the Bragg peak in Proton Radiation Therapy (PRT) qualitatively comparing the entrance and exit dose of x-rays to protons. The width of the peak is achieved by varying the proton incoming energy to cover the volume of the tumour.

moving through a material) it also doesn't deviate much from its initial trajectory. However, as it gradually slows down it gives off increasingly more energy, cascading into depositing almost all of its energy in a small region of space - after which it stops. This is known as the Bragg peak and its precise location in the trajectory depends on the incoming energy of the proton and the material they are travelling through. Proton Radiation Therapy (PRT) utilises this phenomenon to minimise the entrance and exit dose into healthy tissue. This contrast is summarised in Figure 1.7.

This approach brings its own challenges; the primary concern is the exact location of the Bragg peak within the patient. There have been extensive investigations into characterising these beams using Monte Carlo simulations [26, 27, 28, 29] but factors such as tumour placement within the body, patient restlessness during treatment, patient physique, moving fluids, etc. limit the precision to which the placement of the Bragg peak can be estimated. This has led to the investigation of imaging the beam and energy-deposition during treatment so that real-time adjustments can be made to ensure that the tumour is being precisely targeted.

1.6.2 Prompt Gamma-Ray Imaging using Compton Cameras

When a high energy proton makes its way through tissue it predominantly interacts with the electrons, but will occasionally interact directly with atomic nuclei (typically the abundant oxygen, hydrogen and carbon found in organic matter). One such interaction is the inelastic scatter, exciting the nuclei into higher energy states which promptly decay to the ground state within nanoseconds, giving off the energy as a gamma ray - known as a prompt gamma (PG). The cross-section for these inelastic scatters is relatively constant for proton energies between 100-200 MeV, and then increases proportionally as the energy drops below 100 MeV [30]. This implies

that emission of these prompt gammas will be proportional to the energy deposition curve, i.e. it follows the Bragg peak. These prompt-gammas can also be imaged, known as Prompt Gamma Imaging (PGI). For imaging this has several advantages: the gamma spectrum of these elements is well understood and so easy to identify and characterise, the immediacy of the gamma emission eliminates issues like biological washout and allows for real-time imaging, and the equipment needed to image prompt-gammas is relatively unobtrusive. However, while the detector systems can determine where in space a PG is detected, it also needs to be able to determine the direction that PG came from. For example, in SPECT the gammas are physically collimated, in PET the coincidence information provides a correlation and therefore emission direction and in CT the x-rays are emitted along a known trajectory and detected using a position-sensitive screen. In PGI, the incoming intensity is too low for collimation, and since a PG is emitted from inside the patient there is no spatial correlation to exploit. One solution to this problem is through the use of Compton Camera (CC) imaging.

A Compton Camera, as the name implies, uses Compton scattering of gamma-rays within or between detectors to infer the incoming direction of those gammas, and by aggregating many such events is able to reconstruct the initial emission distribution. They are particularly useful over other imaging techniques - such as PET - when the incoming gamma-ray energies cannot be selected, i.e. known isotopes cannot be inserted into the region of interest. Compton Cameras are also small compared to other gamma-ray imaging apparatus (PET, SPECT, etc.) since they're able to reconstruct images without full angular coverage of the source distribution and require relatively less data and so can be easily installed into existing PRT gantry systems. This makes them a promising option for *in-vivo* imaging.

Investigations into using Compton Camera systems for PGI have revealed a challenging filtration problem due to the large amount of noise produced by electron bremsstrahlung, Doppler broadening effects, erroneous background gammas that scatter into the detector, prompt gammas that scatter within the patient before reaching the detector and the assumption that a double or triple scatter event recorded in the detector ended with a photoabsorption when it in fact scattered out of the detector. Some ongoing research² has shown that even the neutron flux produced by PRT beams is a significant contributor to noise. Since the energy deposited by a PG scattering within a detector covers a wide range (particularly for high energy gammas), simple energy filtering has proven insufficient to resolve the true Compton scatter events from the background noise, prompting investigation into additional filtering techniques. One such technique is the distance-of-closest-approach (DCA) filter [31]: since PGs are emitted very close to the proton-nuclei inelastic scatter sites and most of the incoming protons interact approximately along the beam vector, almost all the true PGs must be emitted from somewhere along the beamline. Therefore the known incoming proton beam vector can be used to remove any reconstructed cones that don't appear to originate from the expected beamline. However, this does necessitate prior knowledge of the source distribution. See Figure 1.8 for a diagrammatic depiction of the DCA filter. A full description can be found at [31, 1].

²At the time of writing, this a new avenue of research and as such has not produced any published works.

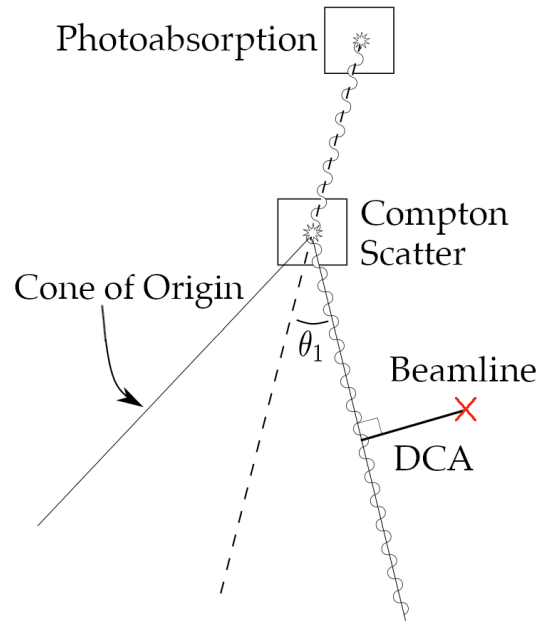


FIGURE 1.8: A visual depiction of how the DCA is determined for a particular Compton scatter event. A gamma-ray arrives in the detector, scatters inside the CZT and gets absorbed. From this, the cone-of-origin is reconstructed, and the closest distance from the surface of this cone and the assumed beamline is calculated. The DCA filter then either selects or rejects this event based on either some threshold DCA value, or on some threshold percentage of all the detected events over some period of time.

1.6.3 Using PEPT to Inform DCA Filtering

Besides scattering, proton-nuclei interactions can also lead to capture, creating positron-emitting radioactive isotopes such as ^{15}O and ^{11}C , which can be imaged using PET techniques [32, 33]. This form of delayed emission imaging, however, requires lengthy imaging times in-between treatments and so cannot produce real-time images. Additionally, in the case of ^{11}C , it has a half-life of 20 minutes and so is likely to move due to the body's natural processes before it decays. This biological washout increases uncertainty in the true origin of the initial proton interaction. ^{15}O , by contrast, has a half-life of 2 minutes so is less susceptible to washout. It is also proportionally the most produced isotope in PRT. However, most of the activated ^{15}O isotopes decay within 10 minutes of activation [34] which is too fast for producing a useful PET image, particularly if the patient needs to be moved to a different room for the scan. All this limits the usefulness of PET imaging.

One of the advantages of the PEPT technique is its ability to find a tracer location with just a few hundred coincidence lines without requiring full angular coverage (see Section 1.2). It is also possible to extend the technique to pick out a beam-line of activated isotopes within the field-of-view of the detector array. Therefore if the detector array being used to do Compton Camera imaging is set up to do coincidence processing around the treatment area, it is possible to pick out the beam-line from the short-lived ^{15}O isotopes *in vivo* and then use that to inform the DCA filter, all with the same detector setup.

In Chapter 6, the PolarisJ detector modules are used to implement PEPT techniques to perform beam-tracking on a 55 MeV and 200 MeV proton beam incident on a PMMA target. Geant4 Monte Carlo simulations are performed to simulate the experiment setup, and results from the measured rates of emitted positrons by a proton beam incident on a PMMA target done by Dendooven *et al.* [35], along with results from Chapters 3-5, are used to estimate the feasibility of such an application in a real-world scenario.

Chapter 2

Equipment

2.1 Cadmium Zinc Telluride

Cadmium Zinc Telluride (CdZnTe or CZT) is a grown semiconductor metal alloy with a band gap large enough to operate as an X-ray and gamma-ray detector at room temperature. It has an energy resolution of $\approx 1\%$ at 511 keV, which makes it useful for performing high resolution gamma-spectroscopy without the need for cryogenic cooling required in other semiconductor detectors like hyper-pure germanium (HPGe). The manufacturers of CZT do not publish the exact composition of their crystals, but the crystals used in this thesis were approximately 48% Cadmium, 2% Zinc and 50% Tellurium. All of the work was done using the PolarisJ CZT detector modules manufactured by H3D Inc. (Ann Arbor, MI, USA).

2.2 The H3D Inc. PolarisJ CdZnTe Detector

The PolarisJ is a self-contained detector module (all the processing electronics are contained in the housing), capable of hosting up to four ($20 \times 20 \times 10$) mm³ CZT crystals in a 2×2 grid. To remain consistent with the terminology used in the literature, in this thesis the term "module" will refer to a single housing (which includes the crystals and electronics inside), and the term "crystal" or "detector" will refer to individual CZT crystals within the module. The module has four external ports: an input for a 12V power supply, a USB port for storing data, an Ethernet port to stream data off the module, send/receive commands to the modules, download data, etc. and a coaxial port for connecting modules to each other for time synchronisation. See Figure 2.1. The detector array used in this work consisted of two PolarisJ modules, each with two CZT crystals. It is owned by the Physics department at the University of Cape Town and was commissioned by Steve Peterson for Compton Camera research in medical imaging.

H3D Inc. primarily focuses on building CZT Compton Camera that are portable, durable and easy to use. They produce units with an integrated 360° optical camera that provides a live view of the surrounding environment with the Compton Camera image overlaid. This allows for the identification of radioactive hotspots in the environment, and has seen many successful applications in nuclear power plants and border security. This thesis describes the first known time a CZT detector has been used to implement the PEPT technique.




FIGURE 2.1: **a)** Photo of the front of the PolarisJ module. **b)** Photo of the side of the PolarisJ module with the external ports labelled.

2.2.1 CZT Crystals and Signal Processing Electronics

The crystals used in the PolarisJ are $(20 \times 20 \times 10)$ mm³ in size with an 11×11 pixelated anode array fabricated onto the back giving a pixel size of 1.8 mm. Depth-of-interaction information is obtained from the ratio between cathode and anode signals. The resolution for the depth-of-interaction (z-coordinate) is 0.5 mm, much more precise than the xy-coordinates [36, 37]. The crystals are arranged in a 2×2 grid inside the detector module, separated by 2 mm - see Figure 2.2. All the signal-processing electronics are contained within the module housing so no external electronics are required to operate the detectors besides a computer to interface with the module.

The modules can be operated in a sub-pixel mode where electronic processing is used to determine the site of the interaction within an individual pixel. This can increase the pixel resolution from 11 mm to 0.23 mm at photon energy of 662 keV [37]. However the cost of the additional processing required effectively triples the time taken to analyse individual events, thus lowering the absolute rate at which data can be acquired and increasing deadtime effects. The advertised timing resolution is <100ns (in non-sub-pixel mode; see Appendix A) and a detailed investigation into detector deadtime effects can be found in Chapter 4. The UCT PolarisJ modules have an energy resolution of 6.66 keV FWHM at 511 keV (1.3%) (see Figure 3.6).

This is necessarily a surface-level description of the PolarisJ CZT detector modules. For more detailed information, please refer to the references and the [H3D Inc. website](#). For a thorough description of the processing electronics, see [38].

2.2.2 Data Collection and Event Information Processing

The PolarisJ outputs list-mode event information in a proprietary binary format with two options - either streamed out of the Ethernet port using the TCP internet protocol or saved onto the detector module's USB storage. H3D Inc. provides software to

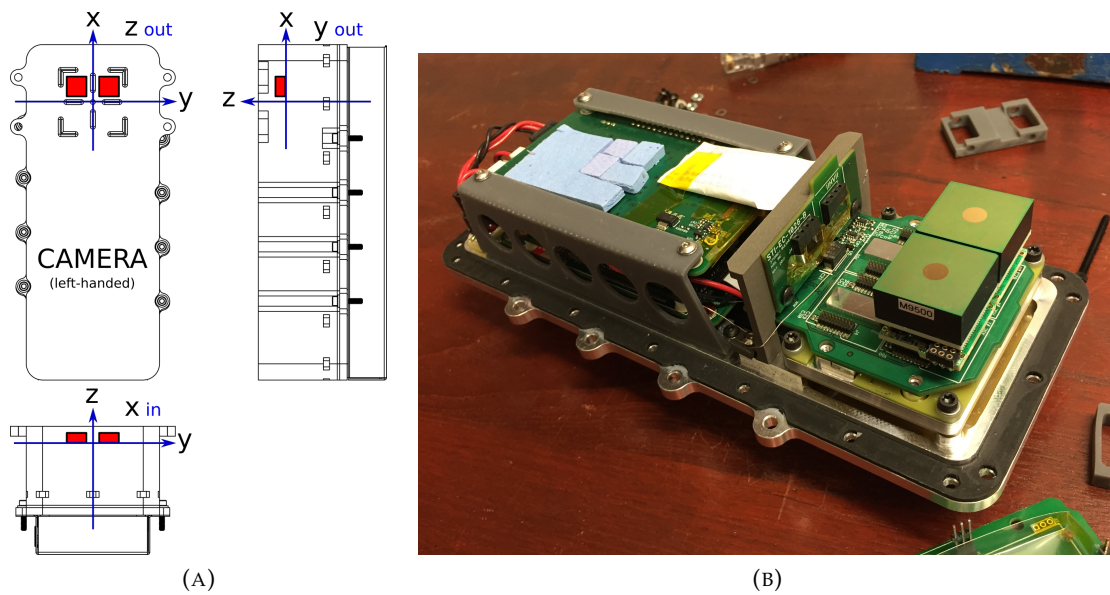


FIGURE 2.2: **(A)** Front-on, side-on and top-down schematics of the PolarisJ module provided by H3D Inc., indicating the internal coordinate system and the position of the crystals (red squares). **(B)** Photo of the internals of the PolarisJ module showing the two CZT crystals (green squares on the right).

convert the binary data into a text file for processing, but for this thesis all the conversions were done using a custom-written python script. Controlling the detector was done via the Ethernet connection; a TCP connection was established through which commands to start, stop, reset, etc. were sent. Controlling and monitoring the modules was done using a custom Python GUI written by Haijian Chen from the University of Maryland, Baltimore. The final output was a text file containing a list of all the events. An example of a recorded event with description of the recorded information can be found in Table 2.1.

Important when doing coincidence processing is to ensure that all the modules are time-synchronised. To synchronise the PolarisJ modules, first a module is designated as the master. It is then connected to the other modules via a coaxial cable, along which it periodically sends a synchronisation pulse every two seconds with a time-stamp and an ID number, which is inserted into the other modules' data-streams. These sync pulses are then used to correct all the event times recorded in the additional modules to that of the master, so that the final event file has all the detected events in order of detection in the lab's time-frame. Since each module outputs its data to a separate binary file, the sync pulses are written out in the data-stream. When these files are read and merged into a single output file, these pulses are used for synchronisation and removed from the final text file. For this work, this was either done with the software provided by H3D or the custom Python script provided by Haijian Chen. This method of synchronisation is flexible and allows for ease of modularity, i.e. allows for multiple modules to be added or removed from an array easily. However, it limits the timing resolution when counting coincidences to the precision of the sending and recording of sync-pulses. If the timing of the sync-pulses are only accurate to microseconds, then the timing window used to find coincidences must be on the order of microseconds, which could add large

TABLE 2.1: Event information provided by the H3D PolarisJ detector modules. A numbered example of a typical double scatter event is shown below the table, with the description of each datum shown in the table. Multiple scatter event information is stored in sequential lines with each line containing the information for each crystal interaction site.

#	Datum	Description
1	Number of crystal interactions	Number of interactions recorded in the module within a 10 ns window following the first event trigger. The number "2" indicates a double scatter and hence the next line contains information about the second interaction.
2	Module identifier	Indicates which module the event was recorded in. In this case it was in module 1.
3	Event energy	The energy deposited in the crystal in keV. The first scatter deposited 197.0 keV and the second 309.8 keV, giving a total of 506.8 keV, indicating this was likely a 511 keV gamma that scattered once and was photoabsorbed.
4	X-coordinate	The event location's x-coordinate in the detector's coordinate system in millimetres.
5	Y-coordinate	The event location's y-coordinate in the detector's coordinate system in millimetres.
6	Z-coordinate	The event location's z-coordinate in the detector's coordinate system in millimetres.
7	Timestamp	Event time of detection in tens of nanoseconds. Both lines have the same timestamp as it was a double scatter event.

2	1	197.0	4.73	16.96	5.5	476798566290
2	1	309.8	3.64	20.76	4.63	476798566290
1	2	3	4	5	6	7

amounts of noise to the coincidence files. This was an issue encountered during this work - see Section 4.1.4 - and occurred because PolarisJ was not designed for high resolution coincidence counting. Future detector systems built by H3D Inc. may implement better electronics to address this issue, particularly with their small-animal PET systems - see Section 2.4 below.

Full Detector Setup and Operating Procedure

Running the detector modules for a typical experiment goes as follows: first, all the modules are arranged in the array to be used for the experiment. They are then connected to the lab computer via Ethernet cables and a network switch, and to each other via coaxial cables. Next, the modules are powered on and they perform an internal booting sequence; the master begins sending synchronisation pulses and all modules start observing the environment, but do not save any event data. The lab computer then establishes a TCP connection via the Python GUI and the modules are ready to begin collecting data. The GUI also monitors the detectors, reporting event rates, internal temperature, humidity and other monitoring information. Once the detectors are set to record, the event data is streamed out via the TCP connection to the lab computer in a list-mode binary format, which is saved on the lab computer's storage for post-experiment processing.

Processing the data begins by reading in the binary data from all the modules, combining all the events into a single array and then transforming the event coordinates from the module's coordinate system into a predefined lab coordinate system, called the isocentric coordinate system - see Figure 3.1 in Chapter 3 for an example of an isocentric coordinate system. This transformation is done by reading in all the event positions into 3D vector format, determining the matrix that represents the translation and rotation operations needed to transform from the detector coordinate system into the isocentric coordinate system and applying the inverse of this matrix to each position vector.

The resulting array is written to a text file in a human-readable text format. This is useful when diagnosing strange behaviour during processing as it is often clear from this file if the data were corrupted in any way. Ultimately a more efficient system would be to store the file in a binary format, but this was not necessary for this work. Depending on the experiment, further processing is then done, including energy filtering, coincidence processing, image reconstruction, etc.

2.3 The UCT PolarisJ Coincidence Detector Array

The primary proof-of-concept and deadtime characterisation experiments were performed on the UCT PolarisJ detector array. This array consisted of two PolarisJ modules, each installed with two side-by-side CZT crystals - see Figure 2.2. For most of the experiments performed, the two detectors were bolted upright in plastic stands, which were then placed inside a 3D-printed slide to position them directly opposite each other. The front faces of the crystals were typically ≈ 60 mm apart, forming a coincidence field-of-view (FOV) of $(60 \times 42 \times 20)$ mm³, i.e. the FOV wherein coincidences could be detected from a positron-emitting source. See Figure 3.1 in

Chapter 3 for a detailed schematic and photographs of a typical setup. The UCT PolarisJ detectors were controlled using a dedicated Dell laptop which interfaced with the modules using an Ethernet cable and network switch, which was also used to stream data off the modules - see Section 2.2.2 above.

2.4 H3D Inc.'s Small Animal PET System

One of the drawbacks to the UCT PolarisJ detector array was its limited size. There were only four crystals, limiting the coincidence FOV to (20×40) mm² in the yz-plane. In theory it was possible to create an arbitrarily large FOV by choosing a large detector separation, but geometrically this increased the uncertainty in the x-direction significantly and so the larger FOV was not useful. However, in mid-2019 H3D Inc. built an experimental small-animal PET imager using four PolarisJ modules, each with four CZT crystals, arranged to form an $(80 \times 80 \times 40)$ mm³ FOV - see Figures 2.3b) and 2.3c). This array was used to perform PEPT experiments in November 2019 during a collaborative visit. Not only did this allow for a much larger FOV, but due to the fact that there were two pairs of opposing modules, the larger uncertainty in the x-axis of the UCT array due to the geometry of a single pair was also eliminated.

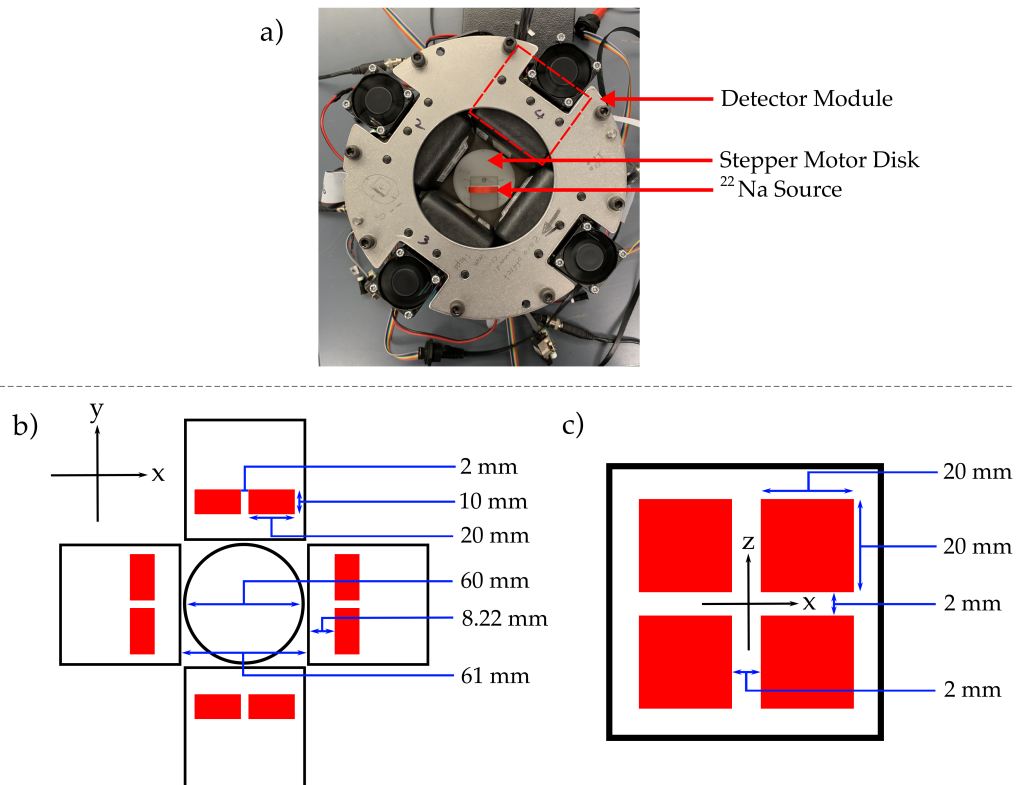


FIGURE 2.3: **a)** Top-down photograph of the H3D Inc. Small Animal PET system. There are four Polaris modules arranged in a square, each with four ($20 \times 20 \times 10$) mm^3 CZT crystals, forming an ($80 \times 80 \times 40$) mm^3 FOV. In the centre is a stepper motor with a 60 mm diameter plastic disk and source stand placing the centre of the source in the centre of the FOV. **b)** Top-down diagram of the PET system, labelled with the dimensions. The CZT crystals are shown in red. The isocentric coordinate system orientation is shown in the top-left corner. The actual origin is in the centre of the FOV. **c)** Face-on diagram of a single Polaris module, labelled with the dimensions, with the isocentric coordinate system shown in the centre.

The processing electronics in this system were the same as described in Section 2.2.1, as well as the data readout and post-processing. At the time of writing this was still an experimental piece of apparatus and so not commercially available. The one drawback to this system was its lack of modularity - it wasn't possible to rearrange the modules into a more optimal geometry for PEPT.

Built into the centre of the system was a stepper motor with a plastic disk with a 60 mm diameter attached directly onto the shaft. Cut into this disk was a groove running from the centre to the edge with small teeth spaced 1mm apart. A plastic source stand was designed to fit into these teeth to allow for the positioning of the source from directly over the centre right out to the edge of the disk. The stepper was controlled by an Arduino Uno and was capable of 700 steps per revolution. See Figure 2.3a).

2.5 The Anet A8 Positioning Stage

In order to characterise the accuracy of tracking when performing PEPT, it was important to be able to repeatedly and precisely move a tracer along a known trajectory to be able to quantify the output produced by PEPT. Since in this particular investigation we were hoping to track particles with sub-millimetre precision, we needed positioning equipment with sub-millimetre precision.

3D printers have been increasing in popularity since their invention in the late 1980s, but particularly in the last 20 years as increases in mass production made cheap, accurate printers available to the average consumer. These cheaper consumer printers are capable of precisely and repeatedly move a printer head around a print bed with sub-millimetre precision, and thus offered a solution to the positioning problem. The Physics department at UCT also owns a high-quality Ultimaker 2+ desktop 3D printer which was available to post-graduates. With this, it was straightforward to buy an existing 3D printer and design and print modifications to build a 3D positioning system.

The printer chosen was the Anet A8, built by Shenzhen Anet Technology Co., Ltd. (Shenzhen, China) and sold by Netram Technologies (Cape Town, South Africa). It was primarily chosen for its low price, simple design that allowed for easy modifications and form factor to accommodate the two PolarisJ modules.

2.5.1 Anet A8 Design

The printer was housed in a $(500 \times 400 \times 450)$ mm³ acrylic frame, which contained the $(220 \times 220 \times 240)$ mm³ print volume, electronics and LCD interface. The print head was moved in the x-direction with a rubber belt driven by a stepper motor, and in the z-direction by threaded bolts, also driven by steppers. The y-axis motion was achieved by moving the print bed with a rubber belt driven by a stepper motor. To interface with the printer, either a micro-SD card with a .gcode print file could be inserted into the printer and selected via the built-in buttons and LCD screen, or it could be connected to a computer via USB and the print file uploaded and run using the Octoprint network interface; the latter was used for this investigation. To control the print head, the printer required a .gcode file which contained standard G-code [39] commands to move the print head to coordinates in the printer's coordinate system. A custom Python script was written to produce .gcode files to move the print head along trajectories traced out by 3D parametric equations. Examples include circles, ellipses, figure-8s, etc.

2.5.2 Anet A8 Modifications and Positioning Performance

Several structural modifications were made to the Anet to attempt to improve the positioning accuracy of the printer. The acrylic material used for the frame was quite flexible and so allowed the printer to sway slightly when moving, particularly in the x-direction. Therefore several structural supports were 3D printed (using the Physics department's 3D printer) to brace it and prevent the swaying. The belts used for moving the print heads were made of rubber and could stretch slightly, and the included mechanism to tighten them didn't allow for very tight belt tension, so new tensioning mechanisms were printed and installed.

TABLE 2.2: List of 3D Printed structural upgrades for the Anet A8, along with links to the print files.

Part	Link
"Hulk" Frame Brace by Leo_N	www.thingiverse.com/thing:2189694
Front Frame Brace by Leo_N	www.thingiverse.com/thing:1857991
Rear Frame Brace by Leo_N	www.thingiverse.com/thing:1852358
Y-Belt tensioner by Photograaf16	www.thingiverse.com/thing:1959208

After the structural modifications, the printer was modified to operate as a positioning stage. The print head was removed and several tracer stands were printed to be installed into the print head's linear bearings. Several stands were required for the various form-factors of the tracer particles. To hold the detector modules, a positioning slide was printed to fit on the print bed and module stands were printed to hold the detector modules in the positioning slide. This positioned the modules directly opposite each other, and allowed for movement along the y-axis. See Figure 2.4 for a photograph and Table 2.2 for the 3D print files of all of the modifications.




FIGURE 2.4: Photograph showing the final modified Anet A8 3D printer used for sub-millimetre source positioning within the UCT PolarisJ PEPT array. The 3D printed modifications are the light blue components seen in the photograph.

The specifications sheet for the Anet A8 is shown in Appendix B. The print resolution is quoted as $\pm 200 \mu\text{m}$ and so this was the uncertainty used on all known tracer trajectories when using the Anet positioning stage. It is possible that this resolution included uncertainties regarding the precision of the laying of print material, meaning the positioning resolution was potentially smaller than $200 \mu\text{m}$. However, the

COVID-19 pandemic interrupted the experiments planned for high-precision tracking and so it wasn't critical to determine the exact precision of the stage.

2.6 β^+ Radiation Sources

2.6.1 ^{22}Na Sources

Two ^{22}Na sources were used for this investigation. The first was a calibration "point" source owned by the University of Cape Town, manufactured by Spectrum Techniques, LLC (Oak Ridge, TN, USA) as part of a standard gamma-source calibration kit, which had a quoted calibrated activity of $37.0 \pm 7.4\text{ kBq}$ in July 2016. The radioactive material was chemically plated on to the surface of a silver disk that was mounted in the centre of a plastic disk. The disk had a diameter of 25mm and thickness of 3mm. Measurements using this source were taken from January 2019 until March 2020, and so the source activity ranged from 20 kBq to 14 kBq during that period.

The second ^{22}Na source was also a lab calibration source made by Spectrum Techniques, owned by the University of Michigan. It was used for a single experiment and had an activity of $1.20 \pm 0.06\text{ MBq}$.

2.6.2 ^{68}Ga Sources

In order to characterise the deadtime behaviour of the PolarisJ detectors and how it impacted the coincidence processing, a short-lived 511 keV emitter with an activity high enough to strongly deadtime-limit the system was needed (see Chapter 4). For this experiment, two ^{68}Ga sources were produced at the PEPT Cape Town facility in iThemba Labs by Mike van Heerden (see Section 1.2.3) using a SnO_2 $^{68}\text{Ge}/^{68}\text{Ga}$ generator and NRW-100 ion-exchange resin manufacturer by Purolite (King of Prussian, PA, USA) - see datasheet in Appendix C. The first was a $\approx 180\ \mu\text{m}$ particle with activity $3.00 \pm 0.15\text{ MBq}$, and the second was $\approx 160\ \mu\text{m}$ particle with activity $5.44 \pm 0.27\text{ MBq}$. Both tracers were glued to the end of a hypodermic needle for ease of use and positioning within the FOV of the detector systems. Photos of these tracers are shown in Figure 2.5. The tracers were not produced simultaneously; the $3.00 \pm 0.15\text{ MBq}$ tracer was produced first and observed over four half lives with a final activity of 0.16 kBq. The second was then produced and observed over one half-life with a final activity of 2.02 MBq.




FIGURE 2.5: Photographs of the two ^{68}Ga sources used for the dead-time measurements as taken through a microscope. The ^{68}Ga was infused into a 160-180 μm NRW-100 ion-exchange resin (orange beads shown in the photos) and glued to the tip of a hypodermic needle before being placed in the centre of the FOV of the UCT PolarisJ array as shown in Figure 4.1. **Left:** Photo of the $3.00 \pm 0.15 \text{ MBq}$ ^{68}Ga source. **Right:** Photo of the $5.44 \pm 0.27 \text{ MBq}$ ^{68}Ga source.

Chapter 3

Experiment 1 - Proof of Concept

The first step to implementing PEPT is being able to locate a stationary tracer particle in three-dimensional space. To do this, the detectors need to be arranged opposite each other and a β^+ point-like source needs to be placed in the centre of the FOV. The detector array then needs to be able to accurately detect the annihilation photons in coincidence and locate these detection events in space. These coincidence events then need to be run through the PEPT algorithm to produce a location and uncertainty that encompasses the known location of the source. Once it is established that PEPT is possible, the performance can be investigated and characterised. This includes producing F-optimisation curves, investigating the FOV sensitivity, investigating different source activities and associated location rates and optimising the coincidence processing parameters. This Chapter shows the initial proof-of-concept experiment using a 17.8 kBq ^{22}Na calibration source (see Section 2.6.1) sitting in the centre of the FOV with the two PolarisJ modules 62 mm apart (the detailed setup is described in Section 3.1.1 below). F-optimisation and FOV sensitivity are also addressed. Chapter 4 investigates higher activity sources and coincidence parameter optimisation.

3.1 Locating a Central Stationary Point Source

3.1.1 Experimental Setup

The two Polaris modules were bolted upright in custom-machined stands, which were then placed inside a 3D-printed slide to position them directly opposite each other. The front faces of the crystals were $62 \pm 1\text{mm}$ apart, forming a field-of-view of $(62 \times 42 \times 20)\text{mm}^3$. The ^{22}Na calibration source (activity $17.8 \pm 3.6\text{ kBq}$) was seated in a 3D-printed stand on top of a linear positioning stage, which was positioned underneath the modules to place the source in the centre of the FOV. See Figure 3.1. The positioning stage was used to move the source and a Vernier caliper was used to verify that the source was equidistant from the edges of the modules. As this was all done by eye, the position was only accurate to ≈ 2 millimetres. For these measurements, however, the exact position wasn't crucial to the success of the experiment provided it was approximately central. In the isocentric coordinate system, this was the position $(0, 0, 11) \pm (2, 2, 2)\text{ mm}$. The source and modules were not moved for the duration of the experiment. A ten minute measurement was taken.

3.1.2 Energy and Coincidence Timing Windows

In order to run coincidence counting on the modules' output, the energy and coincidence timing windows needed to be selected. The energy window refers to a

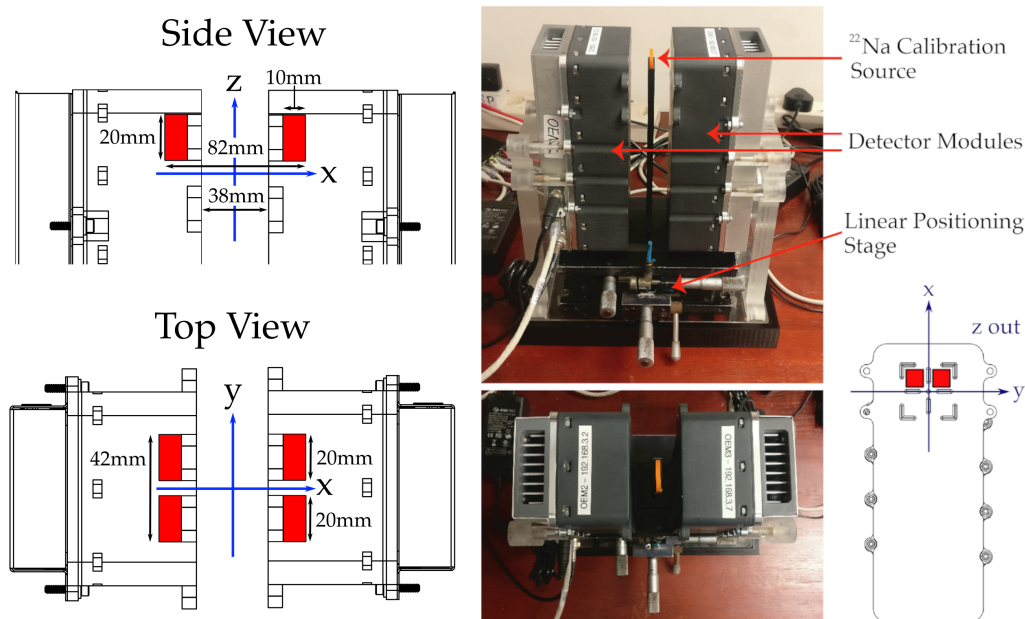


FIGURE 3.1: Experimental setup to run the low-activity stationary measurement described in Section 3.1.1. The isocentric coordinate system is shown in the left-hand diagram; the ^{22}Na source was placed at the position $(0, 0, 11) \pm (2, 2, 2)\text{mm}$. **Left:** Side-on and Top-down schematic of the PolarisJ modules with relevant distances and proportions labelled. The CZT crystals are shown in red and the isocentric coordinate system is shown in blue. **Right:** Two photographs of the experimental apparatus and a schematic of the PolarisJ detector provided by H3D Inc. showing the position of the CZT crystals within the module and the module's (left-handed) internal coordinate system.

typical energy filter: if an event's recorded energy fell within this window it was accepted for coincidence counting, and if not it was discarded. The coincidence timing window is the time around a given event's recorded time to search for a coincident event. See Section 1.2.1 for a full description.

Energy Window

The energy window was selected using the differential energy spectrum from a single hour-long measurement. The spectrum was plotted, a Normal distribution was fitted to the 511 keV peak and the distribution width σ was determined to be 2.506 ± 0.077 keV. The energy spectrum showed that there wasn't significant noise around the photopeak and as this was a low-activity source, inherently there was a low Random coincidence rate. Therefore to maximise measured coincidence rate a wide energy window of 15 keV (6σ) was selected to fully encompass the photopeak. Incidentally, a more thorough investigation into energy windows revealed that having no energy window was preferable - see Section 4.2.2. See Figure 3.2 for the energy spectrum and a visual depiction of the energy window.

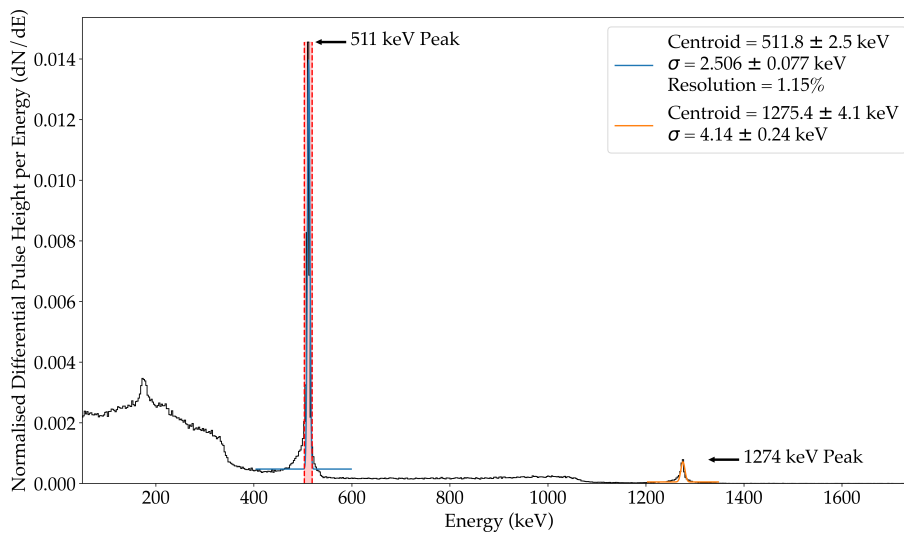


FIGURE 3.2: The normalised energy spectrum of a 17.8 ± 3.6 kBq source observed using the PolarisJ detectors. Two Normal distribution least-squares fits were performed and are shown on the 511 keV and 1274 keV photopeaks characteristic of ^{22}Na with the fit results are shown in the Figure legend. Shaded in red is the 15 keV energy window used to filter events for coincidence processing.

Timing Window

The coincidence timing window was selected using the timing resolution curve. To generate this curve, first the filtered events from module one were separated from those of module two. Next, a time-delay was added to all the events from module two that was sufficiently large to ensure that any detected coincidences were Delayed coincidences (see Section 1.4). This time-delay was then decreased in fixed intervals down to zero, and then down to the same time-delay in the negative direction, performing coincidence counting at each interval and recording the resulting coincidence rate. As the delay decreased to zero, the coincidence rate increased to a maximum as more True coincidences were being counted, and then decreased as the delay became more negative until it reached the Delayed coincidence rate again. Plotting these rates as a function of time-delay produced a curve whose width represented the coincidence timing resolution. Since the quoted timing resolution of PolarisJ is <100 ns (see Appendix A), a coincidence window of 300 ns was chosen for all this coincidence counting; provided it was within the correct order of magnitude, the size of the window should have only affected the height of the curve, not the width.

The curve was produced from the ten minute measurement of the ^{22}Na calibration source and is shown in Figure 3.3, along with a fitted Normal distribution. The centre of the peak did not fall on zero but instead was shifted 71 ns to the right. This was due to uncertainties in the timing electronics between the two detectors that cause the timing pulses to drift over time - this shift was the average of that time-wandering. Over the course of all of the experiments in this thesis, this shift remained within 80 ns from zero and investigations into correcting this shift showed

negligible differences in the final coincidence rates. The shape of the curve was not perfectly Gaussian; it had a "flat top", with a lower peak and wider width than a true Gaussian distribution. This came about from selecting a coincidence window that was larger than the system's intrinsic timing resolution. The area underneath this "flat top" represented a number of coincidence peaks that fully encompassed the True coincidences and were all Gaussian in shape. Similarly to the energy window, the final coincidence window was chosen to be 600 ns (5σ) to encompass the entire curve, and is also shown in Figure 3.3.

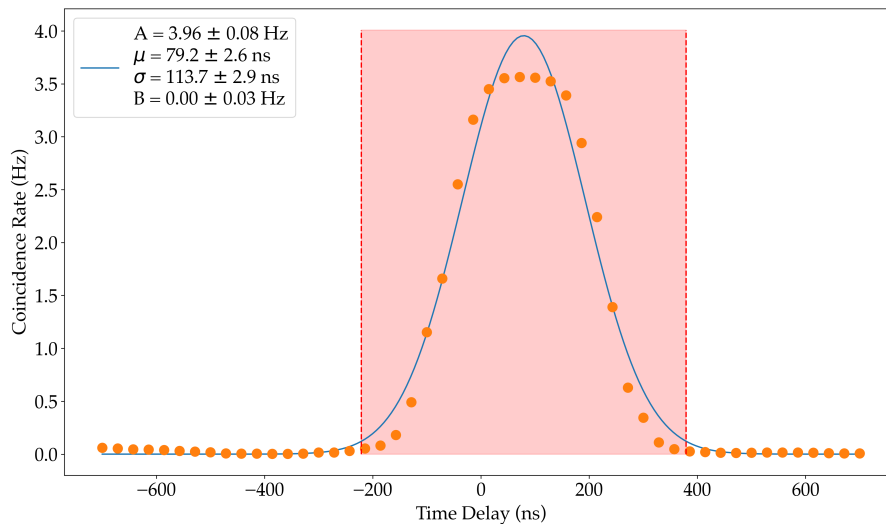


FIGURE 3.3: The coincidence rate (orange scatter points) as measured by the PolarisJ detector array as a function of delay on one of the detector modules to determine the coincidence timing resolution. A least-squares Normal distribution was fitted to the peak using the equation $R = A \exp \left\{ -\frac{(d-\mu)^2}{2\sigma^2} \right\} + B$, where R was the coincidence rate, A was the peak coincidence rate, d was the added module delay, μ was the mean, σ was the standard deviation and B was the random coincidence rate. The fitted parameters are shown in the figure legend. Shaded in red is the final coincidence window of 600 ns (5σ) around the centre of the curve, chosen to encompass the entire curve.

3.1.3 Lines-of-Response and F-Optimisation

After running energy filtering and coincidence processing, there was an average unfiltered single-interaction event rate in both modules of 1500 Hz and an average coincidence rate of 4 Hz. Next the lines of response from the first 1000 coincidences were plotted to check for convergence. The resulting plot is shown in Figure 3.4 showing that the lines mostly converged on the location of the source; the divergent LORs were most likely due to random coincidence. This was a sanity-check to confirm that the coincidence processing was operating as intended and we were able to see a convergence of the LORs.

Next in the PEPT pipeline was to generate an F-optimisation curve (see Section 1.2.1) from these data for a series of lines-per-position (N). Since the coincidence rate




FIGURE 3.4: A sample of the first 1000 coincidence LORs from a ^{22}Na source of activity 17.8 ± 3.6 kBq, plotted as straight lines between the detected positions of the annihilation gamma-rays. All positions are as measured in the isocentric coordinate system. Plotted to confirm that the coincidence processing was running as intended by showing the convergence on the source position.

was on the order of 1 Hz and this was a low-activity source, N values between 10 and 200 were chosen. The results are shown in Figure 3.5. The minima of the curves were found to be between 0.85 and 0.95. Using $f_{opt} = 0.95$ and $N = 15$, the final position determined by the PEPT algorithm was $(-0.60, 0.90, 11.00) \pm (0.50, 0.16, 0.15)$ mm. Using $f_{opt} = 0.90$ and $N = 100$, the final position was $(-0.700, 0.900, 11.000) \pm (0.132, 0.066, 0.056)$ mm. Both of these agree with each other and the known location of $(0, 0, 11) \pm (2, 2, 2)$ mm.

The quality of the results from this initial proof-of-concept experiment was the catalyst for the work done in this thesis. The first interesting result was the F-optimisation curve, specifically the f_{opt} parameter at the minimum of the curves in Figure 3.5 which was around 0.95 for $N = 15$ and between 0.9 and 0.95 for $N = 30$ to $N = 200$. As discussed in Section 1.2.2, this minimum started between 0.15 - 0.35 for the isobutane gas detector at the Birmingham lab in the 1990s. It then stepped up to 0.35 when Birmingham acquired the NaI ADAC system, and took another step up to 0.6 - 0.7 for the BGO modular detectors and the HR++ as researchers improved the detectors and processing electronics to maximise the efficiency for 511 keV gamma-rays. The f_{opt} parameter of 0.85 - 0.95 for CZT represented another step-change improvement in the signal-to-noise ratio; almost all the detected coincidences at this activity were true coincidences. This result was consistent with Figure 3.4, which




FIGURE 3.5: The ^{22}Na source position uncertainty plotted as a function of the PEPT F-parameter to determine the optimal F-parameter to minimise the position uncertainty for this particular experimental setup. Several of these F-optimisation curves have been plotted for multiple different lines-per-position (N-value). A polynomial of degree four has been fit to each curve for clarity. These N-values are shown in the figure legend. **Inset plot:** The source position uncertainty plotted as a function of the N-value using an F-parameter of 0.95. Plotted to show the $1/\sqrt{N}$ dependence of the uncertainty on the N-value.

showed visually that a very small fraction of the detected LORs did not converge on the location of the source. This improvement in signal-to-noise ratio was due to two features of the experiment. Firstly, the high energy resolution of 1.15% in the semiconductor detectors was able to exclude most of the background noise (in-scatter events, Compton-scatter events from the 1274 keV line, cosmic radiation, etc.) at the energy filter. This energy resolution is an order of magnitude better than BGO and a direct comparison is shown in Figure 3.6. Secondly, because the source activity was low (17.8 ± 3.6 kBq), the detected event rate in the modules was also low, which ensured that the Random coincidence rate was very low. The measured Delayed rate was < 1 Hz.

The second interesting result was the minimum value for the position uncertainty, which for the lowest $N = 15$ was $546 \mu\text{m}$, and for $N = 200$ was $101 \mu\text{m}$, showing that a radial position accuracy of just over 540 microns was possible with as few as 15 lines-per-position. This positional accuracy was mostly due to the millimetre pixel size and depth-of-interaction information offered by the PolarisJ modules, which meant that the primary source of uncertainty in the location of the tracer was due to source distribution, positron range and acollinearity of the annihilation photons. However, these sources of uncertainty cannot be further reduced by improved




FIGURE 3.6: The energy spectrum of a ^{22}Na source as measured by the PolarisJ detectors (in red) compared to the BGO detectors in the UCT HR++ detector array in blue. A Normal distribution was fit to each photopeak and the fit results are shown in the Figure legend. Both spectra were normalised to the area under the curves and the BGO spectrum has been scaled up by a factor of ten for clarity.

pixel resolution; to further reduce the location uncertainty it is necessary to improve the statistics by maximising the coincidence rate to increase N - see Section 3.2 below.

The final result to note was the Prompt coincidence rate of 4 Hz, which gave a maximum location rate of the tracer of 0.27 Hz using $N = 15$ (for a stationary location uncertainty of $546 \mu\text{m}$). For this activity source, with this detector array in the described geometry, to track a moving tracer would require a very slowly moving tracer or very smooth tracer motion - a highly limiting concern unless the location rate could be improved. There are two primary methods used to increase the location rate in PEPT for a given geometry: either improve the rate at which the detector array can detect coincidences (either at the electronics or data-processing level) or increase the activity of the source. There is a trade-off, however, between increasing location rate and decreasing the location uncertainty. To understand this trade-off requires an in-depth discussion about the location uncertainty in PEPT.

3.2 PEPT Location Uncertainty

With any positron imaging technique, researchers are constantly striving to improve the spatial resolution of their imagers. In PEPT this means maximising the location rate (i.e. the rate at which the particle is located in space) while minimising the location uncertainty. The location rate is maximised by maximising the coincidence rate, typically by adding more detectors or increasing the data-rates from each detector. To minimise the location uncertainty there are two aspects to consider. As described in Section 1.2.1, the location uncertainty of a stationary source location is described as $u(x, y, z) = \sigma_{x,y,z} / \sqrt{N_f}$, where $\sigma_{x,y,z}$ is the width of the distribution of the measured (x, y, z) coordinates of the stationary source, and N_f is the number of lines-per-location used. Therefore $u(x, y, z)$ can be reduced either by decreasing $\sigma_{x,y,z}$, or by increasing N_f .

Let us consider $\sigma_{x,y,z}$. At the tracer level, this is influenced by source distribution, positron range and acollinearity of the annihilation photons - the greater these factors are, the greater $\sigma_{x,y,z}$ will be. They can be reduced by reducing the size of the tracer, selecting a tracer source with lower β^+ kinetic energy or encasing the tracer in a dense material. However it is very important to match nuclear, physical and chemical properties of the tracer particle to the system being investigated. This, along with the cost of isotope production, heavily restricts the choice of tracer isotope and so these sources of uncertainty are usually treated as a constant. At the detector level, $\sigma_{x,y,z}$ is influenced by the detector's energy resolution and pixel resolution. If the system being investigated has an abundance of water in it (not unusual for a PEPT experiment) there will be a significant number of in-scattered events. If the energy resolution of the detector isn't fine enough it will include these scattered events as signal (particularly for events that scattered with small angles), which increases $\sigma_{x,y,z}$ - although the PEPT algorithm should mitigate this to a large extent. As for pixel resolution, the size of the pixels within the detector have a large effect on the size of $\sigma_{x,y,z}$ given a fixed N_f , particularly if N_f is small - see Figure 3.7 for a demonstration of how this works. Therefore $\sigma_{x,y,z}$ can be reduced by increasing detector pixel resolution, but only while the pixel uncertainty is greater than the positron range and acollinearity uncertainties; once these uncertainties dominate, reducing the location uncertainty is a matter of increasing N_f .

Now consider increasing N_f . Note that there is a nuance here: N_f is a parameter that is set by the experimenter when running the PEPT algorithm and is not directly dependent on any of the experimental apparatus. It is possible to set N_f arbitrarily high, and on a stationary tracer this will decrease the location uncertainty arbitrarily lower. However, PEPT is interested in tracking moving particles and when tracking a moving particle it is necessary to assume that for each source location found the source was "stationary". In practice this means that for each particle location calculated, the particle moved very close to linearly and moved a short distance compared to the overall movement of the tracer during the time taken to collect the LORs needed to obtain that location. Therefore the experimenter needs to consider the coincidence rate of the experiment and expected movement of the tracer and set N_f low enough to make this a reasonable assumption. Therefore to increase N_f it is necessary to increase the coincidence rate, but there is a further nuance. The coincidence rate is made up of three different rates: the True rate, which is the rate at which both 511 keV photons from a single annihilation event are detected uncorrupted and in coincidence, the Random coincidence rate, which is the rate at which

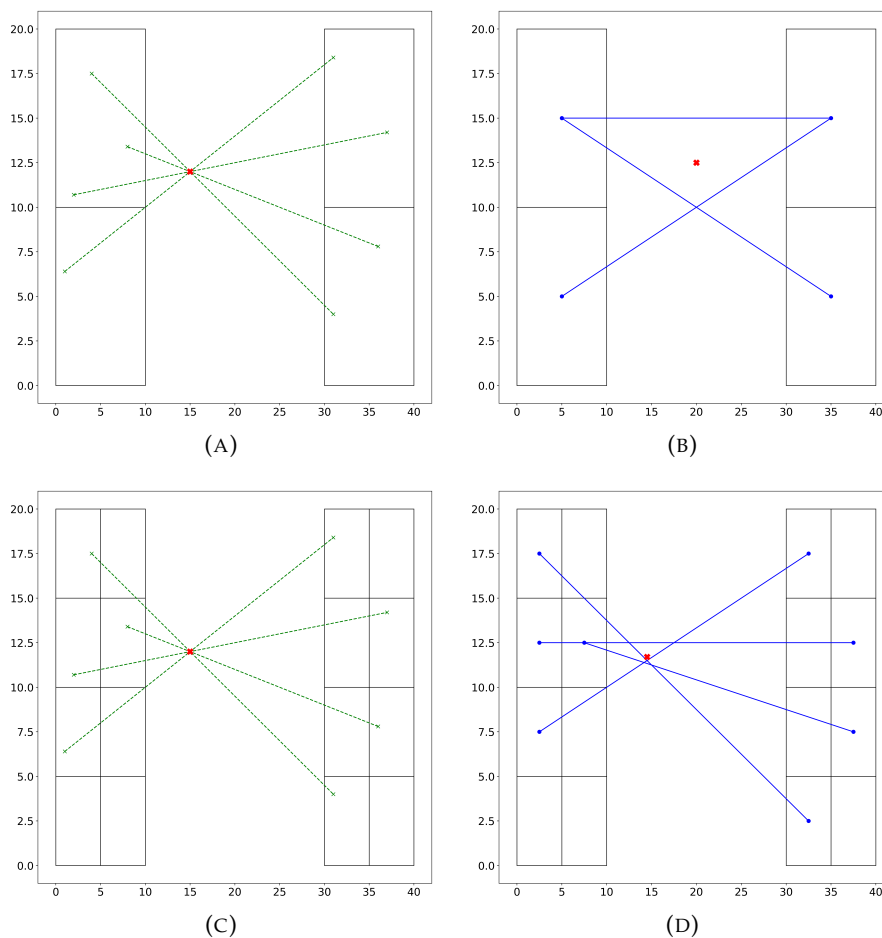


FIGURE 3.7: Demonstration of how the detector pixel size influences the tracer location uncertainty for a given number of lines-per-position (N_f). Consider a point β^+ source sitting in the FOV of two detectors, each with a pixel resolution of 1×2 , and suppose the detectors record 4 LORs from the tracer. **(A)** shows the true location of the source, along with the true interaction locations of the annihilation photons within each detector. Since the detector is pixelated, the recorded interaction locations will be at the centre of the pixels, and the recorded LORs and location will be as in **(B)**. Regardless of the orientation of the LORs, provided they trigger the same pixel pairs the recorded location and minimum distances from the final location to the LORs will be the same. Now suppose the detectors have a pixel resolution of 2×4 , as shown in **(C)**. The recorded LORs and tracer location will be as shown in **(D)** - note how much closer the recorded location is to the true location, and smaller the minimum distances are. This is how increasing the pixel resolution reduces the minimum location uncertainty on the recorded tracer location for a fixed N_f . Note that this effect can be reduced by adding more detectors and increasing N_f , as the proportion of the number of lines along each possible pair will converge to place the minimum distance point closer to the true location.

two photons within the energy window(s) are detected in coincidence by chance, and the corrupted coincidence rate which is when one of the annihilation events scattered before being detected. Therefore it is not good enough to just increase the total coincidence rate, it is important to increase the True coincidence rate. Since the PEPT algorithm is capable of discarding divergent coincidences, increasing the True rate can be done at the expense of also increasing the Random and corrupted rates, but it is important that when increasing the total rate the experimenter ensures that the True rate is also increasing.

At the tracer level, the coincidence rate can be increased by increasing the tracer activity, but deadtime effects and detector timing resolution place an upper limit on the effectiveness of this approach. This is addressed in Chapter 4. It is also possible that the tracer size sets an upper limit on the activity - typically this becomes an issue for tracer sizes below 100 μm . On the detector level the coincidence rate can be increased by improving the absolute efficiency of detecting 511 keV gamma-rays. Either the detector array is arranged to maximise the solid angle subtended around the system being investigated or the intrinsic efficiency of the detectors is improved. This has been the primary approach at improving PEPT performance with scintillator detector arrays. The development of the modular detectors at Birmingham (Section 1.2.2) allowed researchers to place the detectors in optimal geometries for each experiment, and constant improvement to the processing electronics has improved the intrinsic efficiency at 511 keV over the past few decades. This approach has been very successful, and the HR++ at PEPT Cape Town is now capable of a Prompt coincidence rate of 4 MHz, allowing for location rates in the 10s of kHz and giving a stationary location uncertainty of 1 mm [17].

The main aim of this thesis is to investigate the alternative approach to improving location uncertainty: to reduce $\sigma_{x,y,z}$ using a detector with a smaller detector pixel resolution (an order of magnitude smaller than BGO) and depth-of-interaction information, and the results in Section 3.1.3 show that this is a promising approach. However, a location rate of 1 Hz significantly limits the scope of application and so efforts needed to be made to maximise the location rate as much as possible otherwise the added benefit of reducing $\sigma_{x,y,z}$ would be lost due to N not being large enough. This required a detailed look into tracer activity, detector geometry, detector processing electronics and coincidence processing parameters. The results of that investigation is presented in Chapter 4.

3.3 Field-of-View Sensitivity

Since PEPT aims to track a particle moving throughout the FOV of the detector array and there may be very little *a priori* information about the nature of its movement, it is important to establish that there is no bias on the located position of the tracer due to its location in the detector array FOV - caused, for example, by position-dependent sensitivity within the detector crystal to 511 keV gamma-rays.

3.3.1 Experimental Setup

The PolarisJ modules were bolted upright in their stands, which were placed inside a custom 3D-printed slide. This slide was then placed on the print-bed of the modified Anet A8 3D-printer so that the crystal faces were 62 ± 1 mm apart, measured with

a Vernier Caliper. A ^{22}Na calibration source (activity 17.8 ± 3.6 kBq) was placed in a custom 3D-printed stand which was screwed into the Anet to position the source in the detector-array FOV. See Figure 3.8 for a photograph of the 3D printer setup, and Figure 3.1 for the FOV dimensions and isocentric coordinate system. For each coordinate axis, the Anet was programmed to place the source at 2 mm increments along the axis, starting and ending outside the FOV, thus spanning across the FOV in each x , y and z coordinate axis. At each of these positions, a stationary measurement of one hour was taken. These positions were plotted as a function of the known positions of the source and a weighted linear least-squares fit was performed to verify that the plot had a slope of one (i.e. the known and measured positions agreed within uncertainty). From this fit, the residuals were then calculated and plotted to check for any bias within the detector-array FOV.

Finally, the sensitivity to coincidences was determined using the data from the y -axis traverse. The coincidence rate was measured for each position along the axis and the Prompt rate was plotted as a function of y -coordinate. This was to determine how the Prompt rate responded to the location of the particle within the FOV.




FIGURE 3.8: Photograph showing the experimental setup used to produce the FOV sensitivity plots in Section 3.3. See Figure 3.1 for the FOV dimensions and isocentric coordinate system.

3.3.2 Results

Measured versus Known Locations

The results are shown in Figure 3.9. As shown in the Figure legends, the slope for the y and z axes were unitary within the fit uncertainty, while for the x -axis the slope was unitary within three multiples of the fit uncertainty. Looking at the fitted y -intercept parameters, while we expected them to be zero, the x and y axes were

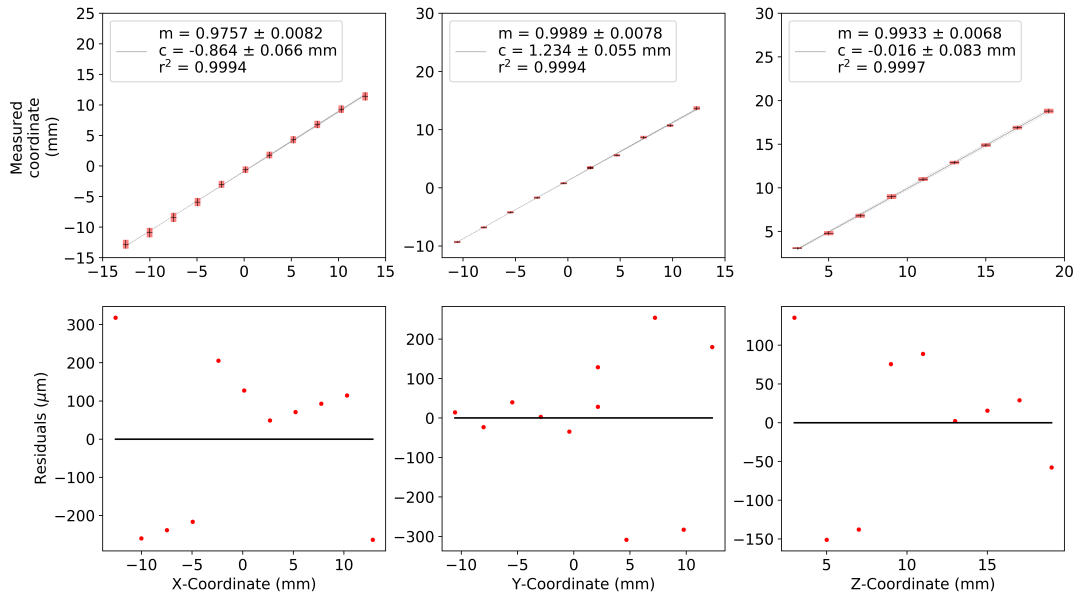


FIGURE 3.9: **Top Row:** The position of a stationary ^{22}Na source determined by the PEPT algorithm as a function of the known position of the particle. Plotted to ensure there is no clear positional bias within the FOV of the detector array. The uncertainties on the positions have been plotted as red error-boxes, with the PEPT-measured position (y-axis) error-bars plotted using three standard uncertainties for visibility. A least-squares linear fit has been applied and plotted over the data, with the fitted parameters shown in the figure legend.

Bottom Row: The residuals between the data-points and the linear fit from the data in the figures above each plot, plotted to reveal any bias within the FOV of the detector array. Note that the residuals are plotted in units of μm , not mm as in the Top Row plots.

not. This was because while the uncertainty of the known position between positions was small, the initial placement of the source had to be done by eye, which had errors on the order of millimetres. Looking at the positional uncertainties, the PEPT-measured errors for the x-axis were 2-3 times larger than for the y and z axes, due to the geometric effect of the x-axis being between the module faces and resulting in a low angular range for the photons. This could be improved by moving the modules closer together, but at the cost of a limited FOV.

The residuals plot showed that the residuals were on the order of $300 \mu\text{m}$, which was within the positioning uncertainty using an N -value of 30, and tended to be larger towards the edges of the FOV. This was unsurprising as the useful area of the detectors for detecting coincidences diminished towards the edges, and so fewer and fewer pixels could be used to determine the position. There was no clear bias in any of the coordinate axes.

Prompt Rate Sensitivity

The results are shown in Figure 3.10. The measured Prompt rates have been normalised to the maximum recorded Prompt rate. The sensitivity peaked at the centre of the FOV and dropped off to $\approx 30\%$ of the maximum on the edges. From the edges




FIGURE 3.10: Normalised Prompt coincidence rate (normalised to the maximum measured Prompt rate) as a function of the y-coordinate for a ^{22}Na source of activity 17.8 ± 3.6 kBq at stationary positions along the y-axis. Plotted to show the sensitivity for detecting coincidences within the FOV of the PolarisJ detector array.

to the centre the Prompt rate increased approximately linearly. It was expected that the sensitivity would drop slightly in the centre of the FOV as there was a 2 mm gap between the detector crystals in both modules. However, the Prompt rate still peaked at the centre, indicating that the increase in solid angle for coincidence detection across both detector crystals compensated for the loss of coincidences between the crystals.

Chapter 4

Experiment 2 - Deadtime Behaviour

The choice of source particle when performing PEPT is vital to the success of the technique - careful thought needs to be put into deciding the physical, chemical and nuclear properties of the tracer for each system being investigated. Of the nuclear properties, ideally the tracer should be short-lived to avoid complicated disposal procedures after the experiment, it should be a "pure" β^+ emitter (it only emits 511 keV annihilation photons) to avoid wasting detector resources on errant gamma-rays and it should have an optimal activity to maximise the True coincidence rate. Given a perfect detector, it is expected that the detector would respond linearly as source activity increased and so increasing useful coincidence rate would just require higher activity tracers. However, real detectors experience deadtime effects as the incoming radiation flux increases, which both limits the maximum detectable coincidence rate and may even degrade the measured event and coincidence rates. Additionally, the Random coincidence rate increases with higher activity, which adds noise to the data and may obscure the True events. This places a maximum on the tracer activity for each system being studied, giving a maximum possible True coincidence rate. This chapter investigates these limits for the PolarisJ detector array. The details of the models used in this chapter are described in Sections 1.3 and 1.4.

4.1 Detector Deadtime Response to Short-Lived ^{68}Ga Source

4.1.1 Experimental Setup

The two Polaris modules were bolted upright in their stands and placed inside a custom 3D-printed slide to position them directly opposite each other to form a FOV of dimensions $(62 \times 42 \times 20) \text{ mm}^3$. A short-lived ^{68}Ga (half-life 67.71 minutes¹) tracer with initial activity $4.72 \pm 0.24 \text{ MBq}$ was glued to the tip of a hypodermic needle and positioned in the centre of the FOV. See Figure 4.1. The initial tracer activity was chosen to strongly deadtime-limit the modules and was determined through trial-and-error. The source was then observed over five half-lives until it decayed to an activity where the detectors responded linearly to the decreasing activity (160 kBq). To ensure the source was as close to the centre of the FOV as possible, a Vernier caliper was used to place the source equidistant from the edges of both modules and the recorded event rates in each module were checked to verify that they were approximated equal. The positions of the source and modules were not adjusted for the duration of the experiment.

¹E.A. Mccutchan, Nuclear Data Sheets 113, 1735 (2012)




FIGURE 4.1: Photograph of the experimental setup used for Section 4.1 - for a cross-sectional view of the detector modules, see Figure 3.1. The source was placed at the origin of the isocentric coordinate system.

4.1.2 Single Module Response

Shown in Figure 4.2 is raw singles rate in each module as measured over the duration of the experiment. Plotted as a black dotted line is the expected linear response assuming no deadtime. This was determined by performing a linear fit on the detector data for low activity (160 - 220 kBq) and extrapolating over the full activity range; this is the linear response expected given the detector's intrinsic efficiency and subtended solid angle, assuming no deadtime effects. Plotted over the measured data-points is a non-paralysable deadtime model (the τ_{np} parameter fit is shown in Section 4.1.3 below). The measured singles rates started at 25 kHz for the initial activity of 4.7 MBq, decreasing steadily to 24 kHz at 3.5 MBq; these rates were well below the expected linear response or non-paralysable deadtime model (plotted as a solid line in Figure 4.2), indicating that the detectors were experiencing strong deadtime effects and were likely completely saturated due to the small decrease in measured rate compared to the decrease in activity. At 3.5 MBq, the

measured singles rate exhibited a sudden step-change up to 27 kHz - this was likely due to buffer-overflow effects in the module electronics, but the exact cause requires further investigation. Between 3.5 MBq and 2.5 MBq the measured rates were closer to the deadtime model, but were still lower than expected. Module 1 in particular showed significant fluctuations in measured rate. This was likely a combination of the buffer-overflow effects and paralyzable deadtime effects. From 2.5 MBq the rates then follow the non-paralyzable deadtime model closely down to 220 kBq, with the singles rates decreasing from 24 kHz to 6 kHz. The detectors then showed a linear response to source activity between 220 kBq and 160 kBq, decreasing from 6 kHz to 4 kHz. Once the modules were no longer deadtime-saturated, module 1 consistently measured singles rates higher than module 0; this was caused by uncertainties in the experimental setup which placed the source slightly closer to module 1, resulting in a slightly higher incoming radiation flux for module 1 for the duration of the experiment.

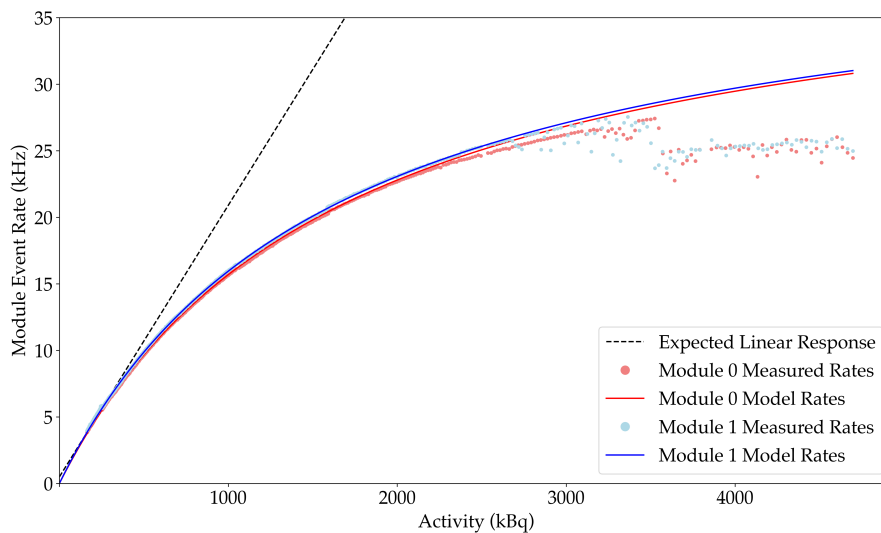


FIGURE 4.2: Measured singles event rates for both modules as a function of activity for a ^{68}Ga source with an initial activity of 4.72 ± 0.24 MBq. The black dotted line shows the expected linear response assuming zero deadtime, determined by performing a linear fit on the detector data for low activity (160 - 220 kBq) and extrapolating over the full activity range; this is the linear response expected given the detector's intrinsic efficiency and subtended solid angle, assuming no deadtime effects. The fitted lines show the expected response from a non-paralyzable deadtime model, with the τ_{np} parameter fit shown in Section 4.1.3.

4.1.3 Deadtime Parameter Fit

The raw singles data in Figure 4.2 show that for a source activity above 2.5 MBq the modules (module 1 in particular) started to experience a form of deadtime not captured by the non-paralyzable or paralyzable deadtime models discussed in Section 1.3. For this reason, the deadtime parameter fit was applied to the data for activities lower than 2.5 MBq. For a detailed explanation of how this fit was performed, see

Section 1.3.4. The resulting fit for the non-paralysable case is shown in Figure 4.3. The final average non-paralysable deadtime parameter was $\tau_{np} = 24.177 \pm 0.096 \mu\text{s}$. This was the parameter used for the fit in Figure 4.2.




FIGURE 4.3: Activity function of the non-paralysable deadtime model described in Section 1.3.1 as a function of measured singles event rate in each PolarisJ module, plotted to determine the deadtime parameter τ_{np} . The activity function of module 1 is slightly higher than that of module 0 because the source was slightly closer to module 1. The fitted deadtime parameters are shown in the figure legend, giving a final average deadtime parameter of $\tau_{np} = 24.177 \pm 0.096 \mu\text{s}$.

A paralysable deadtime fit was also applied and is shown in Figure 4.4 for module 0 (module 1 showed similar behaviour but is not shown for figure clarity). However, a paralysable deadtime model predicts that when the detector begins to saturate at high source activity, the measured rate will start to decrease as source activity increases. This compared to the non-paralysable model which predicts an asymptotic approach towards a constant value. Figure 4.2 shows that for high activity, the measured rates showed asymptotic rather than decreasing behaviour, indicating that a non-paralysable model was the more appropriate one to use for the activity range investigated here.




FIGURE 4.4: Activity function of the paralyzable deadtime model described in Section 1.3.2. Plotted over the data is a linear fit to measure the deadtime parameter, with the fit parameters shown in the figure legend. The deadtime parameter was $\tau_p = 15.83 \pm 0.15 \mu\text{s}$.

4.1.4 Coincidence Rates

Contrary to the coincidence processing applied in Section 3.1, the coincidence rates were measured from the singles data with no energy filtering and using a coincidence window of $10 \mu\text{s}$. The reason there was no energy filter was because the raw data were recorded and stored by the PolarisJ detectors in list-mode, with all coincidence processing done after the measurement. This is opposed to BGO PEPT where the energy filtering is implemented in hardware and events outside the energy window are rejected before they enter the data acquisition system, which is the dominant source of deadtime. Therefore for PolarisJ every event interaction with the detectors contributes to deadtime across the whole system. If energy filtering was applied to the data before fitting the deadtime models, this would add additional complexity that the models do not account for.

The reason a $10 \mu\text{s}$ coincidence window was used is revealed in the timing resolution curve, shown in Figure 4.5. As expected, there is a central peak near zero time-delay with a 1σ width of 130 ns. However, the Figure shows two additional peaks with similar widths where more coincidence events were being found - this is likely a buffering effect within the processing electronics, but further investigation is required to isolate the precise cause. Important to note here is the difference between Figure 4.5 and Figure 3.3 shown in Chapter 3. Both curves were generated with the same method and processing parameters (explained in Section 3.1.2), except that Figure 3.3 shows time-delays between -700 ns and 700 ns and Figure 4.5 shows time-delays between -5000 ns and 5000 ns. When the data from Figure 3.3 were reprocessed using time-delays between -5000 ns and 5000 ns, the resulting plot revealed additional peaks similar to Figure 4.5. However, since the results in Chapter 3 were a proof-of-concept and maximising coincidence rates was not the primary aim, these additional peaks were excluded through the use of a 600 ns coincidence

window. In the case of the results of this deadtime experiment, the coincidence models being used (described in Section 1.4) assume that all genuine coincidence pairs that are recorded by the detectors are counted in the measured True coincidence rate. Therefore a coincidence window of $10 \mu\text{s}$ was chosen to encompass all three peaks shown in Figure 4.5 to ensure that no genuine measured coincidences were missed during processing. As the Random rate model accounts for coincidence window, the additional Random coincidences were not a concern.




FIGURE 4.5: Timing resolution curve of a central 220 kBq ^{68}Ga source as measured by the PolarisJ module array described in Section 4.1.1. Coincidence processing was performed on the raw, unfiltered event data using a coincidence window of 600 ns. The coincidence window was chosen to reveal the additional coincidence peaks on either side of the central peak.

The measured Prompt, Random and True coincidence rates are shown in Figure 4.6, with the fitted coincidence rate models plotted over the data. Figure 4.7 shows only the True rate as a function of activity. Due to the buffer overflow effects above 3.5 MBq, the data were truncated at 3.5 MBq for this analysis, where the Prompt rate reached a maximum of ≈ 7 kHz. As expected for such a wide coincidence window, the Random rate quickly dominated the Prompt rate, reaching a maximum of 6.5 kHz. From Figure 4.7, the True rate reached a maximum rate of 600 Hz at a source activity of 0.8 - 1.5 MBq and then dropped off with increasing activity. This corresponded to a raw measured singles rate of between 15 and 20 kHz. Therefore for any given experimental setup, the tracer activity should be chosen so that the average event rate in each module is 15-20 kHz. Typically at PEPT Cape Town the source activity is chosen to be slightly higher than ideal so that it decays into the ideal region, giving 2-3 hours of useful experimental time for each tracer. The measured data fit quite closely with the coincidence models as described in Section 1.4, until around 1 MBq where the True (and therefore Prompt) rate deviated from the model, measuring lower than expected.




FIGURE 4.6: The Prompt (blue circle), Random (green triangle) and True (red plus) coincidence rates of a ^{68}Ga tracer with initial activity 4.72 ± 0.24 MBq as a function of the tracer activity. Coincidence processing was performed on the raw, unfiltered event data using a coincidence window of $10 \mu\text{s}$. Plotted over the data are the non-paralysable deadtime model fits. Plotted to show the behaviour of the different coincidence rates for high levels of deadtime in the modules.




FIGURE 4.7: The True coincidence rate of a ^{68}Ga tracer with initial activity 4.72 ± 0.24 MBq as a function of the tracer activity. Coincidence processing was performed on the raw, unfiltered event data using a coincidence window of $10 \mu\text{s}$. Plotted over the data is the expected true rate using a non-paralysable deadtime model fit on the measured singles rates. Plotted to show the True rate more clearly and thus determine the optimal tracer activity which maximises the True coincidence rate.

4.2 Locating an Optimal Activity Stationary Source

For the particular experimental geometry in this experiment, Figure 4.7 showed that a maximum True coincidence rate of 600 Hz was attainable for a source activity of between 0.8 MBq and 1.5 MBq. From this, the PEPT performance corresponding to this optimal tracer activity was investigated, focussing on the F-optimisation curve, the source positioning uncertainty and the corresponding location rates. For the following analysis, a twenty-minute window was chosen with the tracer starting at an activity 1.3 MBq and decaying to activity 1.1 MBq; from Figure 4.7 this corresponded to a True coincidence rate of consistently ≈ 600 Hz. The average singles rate for each module over the twenty-minute interval was 20 kHz. The F-optimisation results were then compared with the results discussed in Section 3.1.3 in Chapter 3. Both experiments had the same detector array setup and therefore solid angle, with the difference being the source type and activity. The experiment in Chapter 3 had a ^{22}Na calibration source with activity 17.8 ± 3.6 kBq and will be referred to as the "low-activity" ^{22}Na source. The experiment used in this section had a ^{68}Ga source with activity 0.8 - 1.5 MBq and will be referred to as the "high activity" or "optimal activity" ^{68}Ga source.

4.2.1 F-Optimisation

The coincidence processing for this section was performed with a coincidence window of $5 \mu\text{s}$. Since the coincidence model is not relevant for F-optimisation, missing a few True coincidences was acceptable and the narrower timing window decreased the Random rate without significantly reducing the True rate, which increased the signal-to-noise ratio. From Figure 4.5, a $5 \mu\text{s}$ window still encompassed all three peaks. With no energy filtering and a timing window of $5 \mu\text{s}$, the average Prompt rate was 2200 Hz, with a Random rate of 1600 Hz and True rate of 600 Hz. An F-optimisation curve was then produced for five N-values (lines-per-position) of 100, 150, 200, 500 and 1000. The resulting F-optimisation curves are shown in Figure 4.8.

The first feature to note was the much lower f_{opt} value for the five curves - all of them were between 0.10 and 0.12 which was much lower than the f_{opt} values of 0.90 - 0.95 values shown in Figure 3.5 for the low-activity ^{22}Na source. However, Figure 3.5 was produced using an order-of-magnitude lower coincidence window of 600 ns - significantly reducing the Random coincidence noise - and an energy window on the photopeak, which lowered the noise caused by in-scattered coincidence events. Additionally, the source activity was 100 times lower, further reducing the Random coincidence rate to almost zero. The second feature to note was the source positioning uncertainty. For $N = 100$ the radial uncertainty was ≈ 2 mm and it only dropped to sub-millimetre for $N = 200$, where it was $\approx 800 \mu\text{m}$. This compared to a radial uncertainty of the ^{22}Na source which was $\approx 550 \mu\text{m}$ using $N = 15$ and $\approx 100 \mu\text{m}$ for $N = 200$.

This apparent low-performance was redeemed by the coincidence and location rates of the ^{68}Ga source. The Prompt rate was 2.2 kHz for the ^{68}Ga source, so using $N = 200$ the location rate was 11 Hz at a location uncertainty of $\approx 800 \mu\text{m}$. For the ^{22}Na source the Prompt rate was 4 Hz, so using $N = 15$ the location rate was 0.27 Hz at a location uncertainty of just under $550 \mu\text{m}$. The higher activity ^{68}Ga source therefore produced two orders of magnitude improvement in location rate




FIGURE 4.8: The ^{68}Ga source radial position uncertainty plotted as a function of the PEPT F-parameter to determine the optimal F-parameter to minimise the position uncertainty for this particular experimental setup. F-optimisation curves have been plotted for multiple different lines-per-position (N-value), indicated in the figure legend. A polynomial of degree four has been fit to each curve to help differentiate between them. **Inset plot:** The source position uncertainty plotted as a function of the N-value using an F-parameter of 0.1. Plotted to show the $1/\sqrt{N}$ dependence of the uncertainty on the N-value.

for comparable location uncertainty. Additionally, it is not entirely fair to directly compare the two location uncertainties. The ^{22}Na source was approximately 1 mm in diameter and encased in a plastic disc. The ^{68}Ga particle was 160 μm and simply attached to the end of needle, meaning at least 50% of the annihilation photons were being expelled into air. The average kinetic energy of the majority of the emitted β^+ particles in ^{22}Na was $\approx 200 \text{ keV}^2$, and in ^{68}Ga is $\approx 800 \text{ keV}^3$. It was expected that the positron range was therefore higher in ^{68}Ga than in ^{22}Na , resulting in higher location uncertainties for similar N-values and geometry. For a ^{68}Ga source encased in a tracer particle that is denser than air or moving inside a solid or liquid system, the positron range is expected to be lower.

4.2.2 Timing and Energy Windows

The results from the previous Section demonstrated that sub-millimetre PEPT performance was maintained with no energy filtering and a much wider coincidence

²M. Shamsuzzoha Basunia, Nuclear Data Sheets 127, 69(2015)

³E.A. McCutchan, Nuclear Data Sheets 113, 1735 (2012)

window while showing a significant increase in location rate compared to the results in Section 3.1.3. Therefore the impact of these processing parameters on location performance was further investigated by re-processing the same dataset using an energy window of 15 keV on the photopeak and a timing window of 600 ns on the main peak in Figure 4.5. In this analysis, for each combination of coincidence processing parameters, an N-value was found such that the final location uncertainty was between 700 - 800 μm . The corresponding coincidence rates, f_{opt} values and location rates were then compared using a bar chart, which is shown in Figure 4.9. The parameters determined are also presented in Table 4.1.

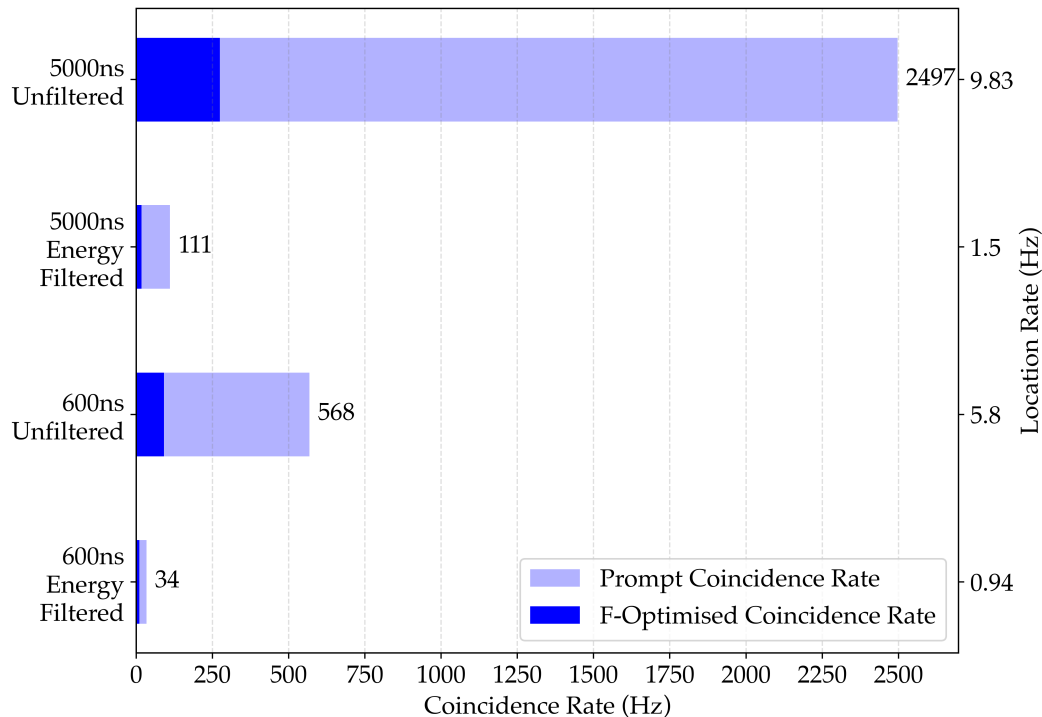


FIGURE 4.9: Bar chart comparing the Prompt coincidence rate (x-axis) and location rate (right-hand y-axis) for four different combinations of coincidence processing parameters (left-hand y-axis) that corresponded to a location uncertainty of between 700 - 800 μm . The translucent blue bars represent the total coincidence rate and the solid blue bars represent the rate after the f_{opt} parameter had been applied to the data, i.e. if the total coincidence rate was 2497 Hz and the f_{opt} parameter was 0.11, then the solid blue bar rate was $0.11 \times 2497 = 275$ Hz. The exact values are tabulated in Table 4.1.

It is clear from Figure 4.9 that it was preferable to have no energy filtering and a wide coincidence window. Both the Prompt rate and the final location rate are an order of magnitude higher for the same radial location uncertainty, despite a halving of the f_{opt} parameter. This is a good example of the power of the PEPT algorithm in distinguishing the difference between signal and noisy LORs - including more and more data by widening the processing parameters meant more signal was included at the cost of proportionally more noise, but the PEPT algorithm could be relied

TABLE 4.1: Table displaying four different combinations of coincidence processing parameters (displayed in the columns) and the resulting Prompt coincidence rate, f_{opt} value and location rate. Also shown are the N-values required to get a location uncertainty between 700 - 800 μm , as well as the final uncertainty.

Window (ns)	600	600	5000	5000
Energy Filter	Yes	No	Yes	No
Prompt Rate (Hz)	34	568	111	2497
f_{opt} (%)	21	16	16	11
N	36	98	74	254
Location Rate (Hz)	0.94	5.80	1.50	9.83
Uncertainty (μm)	711	788	772	711

upon to effectively remove the addition noise. This is particularly advantageous in the case of PolarisJ as data-rates could be improved while retaining the precision gained from the higher pixel resolution.

Figure 4.9 showed that a very large proportion of the True coincidences were removed when an energy filter was applied. For both narrow and wide coincidence windows the Prompt coincidence rate decreased by an order of magnitude when an energy window was applied, while f_{opt} increased only slightly. The two reasons for this were scattering cross-section and experimental setup. For 511 keV photons in CZT, the Compton Scattering cross-section is $7.12 \times 10^{-2} \text{ cm}^2/\text{g}$ and the photoelectric absorption cross-section is $1.52 \times 10^{-2} \text{ cm}^2/\text{g}$ [40]. An annihilation photon was therefore over four times as likely to Compton Scatter than to photo-absorb when it interacted with the crystal. Additionally, the mean-free path for 511 keV photons is 1.9 cm and the crystal was only 1 cm thick, so the incoming photons were most likely to scatter once within the crystal before exiting.

The primary disadvantage to no energy filter is that any genuine coincidence photon that scattered in the experimental system before being detected would be included in the Prompt coincidence data, which lead to inaccurate LORs. However, for this particular experiment there was very little scattering material in the FOV of the detector, and so few photons would have in-scattered. This is why there was a very small reduction in f_{opt} when the energy filter was removed. For systems with more scattering material in the FOV, in-scattered events will become increasingly problematic and energy filtering may become preferable to remove these events.

Another important consideration are tracer isotopes that are not pure 511 keV emitters, such as ^{22}Na which emits a 1.275 MeV gamma-ray that is temporally correlated with the annihilation photons. Without an energy filter, these additional gammas could lead to more Random coincidences and additional noise.

Further research is required to determine to what extent energy filtering is advantageous for high-scatter systems and non-pure tracer isotopes. For example, it could be useful to have a low-energy filter based on the angular scattering distribution or back-scatter peak. Changing the energy window to increase data-rates is often done at PEPT Cape Town when using the ADAC, based on the expected amount of scattering from the system being investigated, with multiple windowing to include the upper end of the Compton continuum. Similar techniques could be

investigated for CZT PEPT. It is also not clear whether there is a critical signal-to-noise ratio where the PEPT algorithm can no longer reliably remove the erroneous lines. However, all experiments thus far suggest that the algorithm is resilient to high amounts of noise. Finally, for the types of systems expected to be studied using CZT (e.g. micro-fluid dynamics) the amount of high scattering material is expected to be very low, so it is preferable not to have an energy-filter.

As expected, the wider timing window lead to a significant increase in the coincidence rate as the additional peaks in Figure 4.5 were included. For both the filtered and unfiltered data, this was an order of magnitude increase in the Prompt rate. This also increased the Random rate and hence decreased the f_{opt} parameter, but the trade-off lead to a two-fold increase in the location rate for comparable location uncertainty. Similarly to the energy filter, increasing data-rates at the cost of additional noise was a worthwhile trade-off.

Parameter Comparisons for Low Activity Source

Naturally, if increasing noise in the service of data-rates was preferable for a high-activity source, this would apply to low-activity tracers with inherently lower noise - particularly in the Random rate. The same analysis as above was thus performed for the low-activity experiment described in Section 3.1.1, with the results shown in Figure 4.10 and Table 4.2. The Figures tell a very similar story to the higher activity source and a wide coincidence window and no energy filter produced an order of magnitude increase in the location rates for similar location uncertainties.




FIGURE 4.10: Bar chart comparing the Prompt coincidence rate (x-axis) and location rate (right-hand y-axis) for four different combinations of coincidence processing parameters (left-hand y-axis) that correspond to a location uncertainty of between 700 - 800 μm . The translucent blue bars represent the total coincidence rate and the solid blue bars represent the rate after the f_{opt} parameter had been applied to the data, i.e. if the total coincidence rate was 121 Hz and the f_{opt} parameter was 0.50, then the solid blue bar rate was $0.50 \times 121 = 60.5$ Hz. The exact values are tabulated in Table 4.2.

TABLE 4.2: Table displaying four different combinations of coincidence processing parameters (displayed in the columns) and the resulting Prompt coincidence rate, f_{opt} value and location rate. Also shown are the N-values required to get a location uncertainty between 700 - 800 μm , as well as the final uncertainty.

Window (ns)	600	600	5000	5000
Energy Filter	Yes	No	Yes	No
Prompt Rate (Hz)	5	69	5	121
f_{opt} (%)	90	50	90	52
N	10	22	10	25
Location Rate (Hz)	0.50	3.14	0.50	4.84
Uncertainty (μm)	730	797	740	753

4.3 Improving Data Rates

4.3.1 The Limits of High Resolution

As discussed in Section 3.2, the radial uncertainty of a PEPT location is $u = \sigma / \sqrt{N_f}$, where $N_f = f_{opt} \cdot N$ is the number of accepted LORs and σ is the distribution width of these accepted LORs around the final PEPT location. As shown in Figure 3.7, the location uncertainty width σ can be reduced dramatically by increasing the pixel resolution within the detector, especially for low N values, and as described in the past two chapters it is possible to achieve sub-millimetre location uncertainties with as few as 10 LORs using PolarisJ. Conversely, a detector system can overcome a relatively low pixel resolution if it is able to compensate with high data rates, and the work done at Birmingham and PEPT Cape Town has been very successful at this approach by building very fast and very high efficiency detectors as well as producing tracer particles with increasing activities to take advantage of the high maximum event rates achievable with these systems. Where this strategy fails is in applications that require very small particles, particularly below 500 μm . At this size, current techniques are simply unable to produce tracer particles with enough activity to produce the statistics needed to reduce the location uncertainty to desired values, and this is where the benefit of high pixel resolution detectors is needed.

It is pertinent at this point to reiterate the fundamental limit of improving the accuracy of PEPT by using higher pixel resolution detector systems, as discussed in Section 3.2. Often in PEPT, tracer particles are treated as "point" particles where all the annihilation photons originate from that point, but of course this is never the truth in practice. All tracer particles used in PEPT experiments have a finite size, and positrons are emitted from decaying nuclei with a certain kinetic energy so only annihilate after they thermalise through interactions with the source material or surrounding matter. Annihilation photons are also necessarily assumed to be perfectly collinear, which again they are not. This means that even with an infinitely small pixel size in a detector, a sample of True coincidences will not converge on a single location but instead will be randomly distributed around the centre-of-activity (assuming positron range and acollinearity are Gaussian distributed) of the radioactive material within the tracer, caused by both the finite size of the distribution of source material and positron range. Therefore the limit to improving location accuracy by reducing pixel size is set to these distributions. See [41] for further discussion.

In 1999, Levin and Hoffman released a paper looking at the positron range in water of popular positron-emitting isotopes used in PET imaging [42]. Most relevant to PEPT Cape Town was ^{15}O , which has a mean β^+ kinetic energy of 735 keV compared to ^{68}Ga with a mean β^+ kinetic energy of 829 keV. Water is also not dissimilar to many materials encountered in a typical PEPT experiment. The FWHM distribution width of the positron range was calculated to be 501 μm , with a FWTM of 4140 μm . This is on the same order-of-magnitude of the location uncertainties observed in the results of the experiments presented thus far without considering photon acollinearity and source material distribution⁴, which suggests that the maximum benefit of reducing pixel size has been achieved with the current pixel size of

⁴A FWHM of 500 μm is an order of magnitude higher than the ^{68}Ga sources being produced for this application, and therefore the positron range dominates the distribution width. However, the ^{22}Na calibration sources used have source material distributions on the order of millimetres, which further widens location uncertainty width.

PolarisJ.

PEPT imaging only attempts to locate a single point in space rather than a distribution. Assuming that the spread of annihilation events is isotropic, the uncertainty on a stationary location of the source can be further reduced by increasing the number of lines used for that location. What this means for PEPT using PolarisJ (PolarisPEPT) is that to further reduce the location uncertainty - not to mention track fast moving (> 1 cm/s) tracers - the True coincidence rate will need to be increased. While this is always a goal in producing more accurate PEPT, this discussion aims to point out that trying to further reduce the location uncertainty for a fixed location rate by increasing the pixel resolution is unlikely to be effective.

It is worth pointing out that different source isotopes would be effective in further reducing location uncertainty without increasing data-rates. A good example is ^{18}F , which has a mean β^+ energy of 250 keV, and therefore a FWHM distribution of $102\ \mu\text{m}$ in water ($1030\ \mu\text{m}$ FWTM) - incidentally this is one of the many reasons why it is so widely used in PET imaging. A sub-100 μm ^{18}F tracer particle with an activity of 1 MBq in the current UCT PolarisJ detector array could theoretically provide a 10 Hz stationary location rate with location uncertainty under $200\ \mu\text{m}$. As part of a broader project to increase the scope of tracer particle production, PEPT Cape Town is developing its ^{18}F isotope production capability via beam-activation at iThemba Labs. However, the limit of pixel resolution will still remain an issue and so it is important to try and improve the maximum achievable data-rates with PolarisJ.

4.3.2 Effects on F-Optimisation

One interesting way that CZT PEPT differs from scintillator PEPT is in the f_{opt} parameter. When the first PEPT experiments with PolarisJ were performed, the stand-out feature was that f_{opt} was between 0.85-0.95 - unprecedented for PEPT. At the time this was particularly promising as much of the narrative of PEPT from its inception was to try and improve f_{opt} , from the 0.15-0.35 f_{opt} values in Birmingham to now 0.6-0.7 for the Cape Town HR++ ring. An f_{opt} of 0.85-0.95 represented another step-change in the progression of PEPT.

The reason that PolarisJ was able to obtain an f_{opt} of 0.95 was due to its high energy resolution and pixel resolution. Consider the results of the non-deadtime-limited experiment presented in Section 3.1.3. The energy resolution allowed for a very tight energy window (15 keV), which meant that after filtering, almost all measured coincidences were True coincidences that had not scattered before detection. The low source activity (≈ 20 kBq) meant that with a coincidence timing window of 600 ns, the expected Random coincidence rate was negligible. The high pixel resolution within the crystals meant that all of these detected LORs converged tightly on the source, as shown in Figure 3.4 and explained in Figure 3.7. Looking at Figure 3.4, almost all of these LORs could be used without appreciably increasing σ , and hence an f_{opt} of 0.95 minimises the location uncertainty.

However, the recorded coincidence rates using the energy- and timing-window were almost usually low. As discussed in Section 4.2.2, this was because a significant portion of the True coincidence photons were scattering within the crystals instead of photo-absorbing, and so their recorded energies fell outside the tight 15 keV energy window about the photopeak. Removing the energy window to include

these scattered events naturally lead to including noise from in-scattered events - although for the systems observed in the experiments in this thesis this was not a significant contribution. Further analysis of the coincidence timing resolution (Figure 4.5) showed that even more True coincidences were being lost due to additional delays in recorded event times caused by the detector electronics. Increasing the coincidence window from 600 ns to 5000 ns to compensate increased the Random coincidence rate. This increase in noise reduced f_{opt} from 0.95 to 0.50, but the location rate for the same location uncertainty increased almost ten-fold. This suggested that a high f_{opt} value is thus not necessarily optimal with PolarisJ. Performing the same analysis on the results from Section 4.2 reveals a very similar story, but because the source activity was 100 times higher and the Random coincidence rate grows like the square of the measured singles event rate, the proportion of Random coincidences to True coincidences was much higher. The system was also experiencing deadtime losses, further contributing to lost True coincidences. The combination of these factors further reduced f_{opt} to 0.10.

A further nuance is the mathematical relationship between σ and N_f in the uncertainty $u = \sigma/\sqrt{N_f}$; u scales proportionally with σ but scales inversely proportionally with $\sqrt{N_f}$. For example, if halving the fraction of used lines (N_f) leads to halving σ , then u will change like $0.5/\sqrt{0.5} \approx 0.71$. This means that it is preferable to remove lines (i.e. decrease f_{opt}) to reduce σ until removing lines no longer decreases σ by more than $1/\sqrt{N_f}$. Exactly what the best trade-off is between removing lines to reduce σ versus keeping lines to increase N_f depends on the detector pixel resolution, positron range and photon acollinearity. Provided that the pixel resolution of the detector is high enough, it is preferable to remove lines until σ has reached the physical limits of positron range and photon acollinearity. As per the discussion in Section 4.3.1 above, PolarisJ has a high enough pixel resolution to reach these physical limits, meaning that in general the PolarisJ CZT system favours a lower f_{opt} .

4.4 Conclusions

The deadtime-limiting experiment described in Section 4.1 revealed the biggest weakness in PolarisJ PEPT performance: detector timing and low efficiency for detecting 511 keV gamma-rays. The low coincidence timing resolution shown in Figure 4.5 meant that to maximise coincidence rates a wide window of 5 μ s was necessary, which then lead to a high Random coincidence rate. The deadtime curves shown in Figures 4.2 and 4.6 showed that the detectors became deadtime-saturated at a source activity orders of magnitude lower than BGO, with corresponding coincidence rates similarly low. This placed an upper limit on the maximum achievable True coincidence rate of 0.6 kHz, with a Prompt coincidence rate of 2.2 kHz and Random coincidence rate of 1.6 kHz for a source activity of 0.8-1.5 MBq or a single module raw rate of 15-20 kHz. Compare this to a single pair of BGO modules ($4 \times 8 \times 8$ crystals) that are similarly deadtime-limited: the singles rate is of the order of 1500-2000 kHz while the Random rate is on the order of 40 kHz using a coincidence window of 10 ns.

For the UCT PolarisJ system, the maximum measurable location rates for an optimal activity tracer is on the order of 10s of Hz for sub-millimetre tracking, which limits trackable tracer speeds to mm/s or cm/s. The HR++ ring has accurately measured speeds of 10s of m/s. A simple (but expensive) way to improve coincidence

rates would be to purchase more detector crystals and module pairs to improve the solid angle. However, similarly to the narrative of scintillator PEPT, one of the goals of further research into semiconductor PEPT will be to improve timing resolution and intrinsic efficiency to increase the useful fraction of detected events, i.e. increase f_{opt} . There are certain hardware improvements that could be made to improve the timing performance of the H3D CZT systems. New and improved ASIC designs could speed up the time taken to process each recorded event. Separate readouts for each crystal within the PolarisJ module would allow them to record events independently. At the time of writing, H3D Inc. is working to include such improvements in the next generation of their detector modules. With the next generation of H3D module electronics and a four-module, sixteen-crystal array, it is likely that the performance presented in this chapter could be improved by at minimum an order-of-magnitude. With further developments in optimising tracer particles, two or three orders-of-magnitude may even be possible. An example of such a system is presented in Chapter 5.

Despite these low absolute location rates compared to traditional scintillators, PolarisJ was still able to perform sub-millimetre tracer locating using relatively few lines-per-position and at a relatively low activity. The advantage of PolarisJ PEPT therefore comes down to context. For an experiment that requires a large FOV, using millimetre-size (or bigger) tracers with high activity that move at fast speeds, scintillator detectors easily outperform semiconductor detectors. But as demand for very small tracer particles moving through small systems pushes the maximum achievable tracer activity down, the scintillators cannot obtain the optimally high data-rates that allows them to perform. CZT was able to locate a particle to $730 \mu\text{m}$ using just 23 lines-per-position. The modularity and relatively small size of the detectors allows for greatly increasing achievable solid angle by placing the detectors right on the boundaries of a small system. This shows promise for tracking small, low-activity tracers in small systems when BGO cannot. The next step was therefore to place a moving tracer in a CZT array and quantify its performance. This is the topic of Chapter 5.

4.4.1 Further Work - Multiple Interaction Events

The results in Section 4.2 above demonstrate that useful data-rates were improved by relaxing the energy and timing window processing parameters to include scattering interactions and account for electronic buffering effects. However, it is still possible to extract more useful information from the current PolarisJ data-stream from the multiple scatter interaction events. For the analysis shown in Chapters 3 and 4, when a multiple scatter event was detected in coincidence, one of the possible LORs was chosen at random and the other discarded (or in the case of two multiple events in coincidence, one of the four possible LORs was chosen at random). Not only did this mean that roughly half of the True coincidences that resulted in a multiple scatter event were being lost, but the noise coincidence rate was also being increased by the same proportion.

Not only are coincidences being lost due to single gamma-rays having multiple interactions but True coincidences are also being lost due to "false" multiple interaction events, as highlighted by work done by Polf *et al.* [43]. This discussion focusses on double scatter events, but the same reasoning applies to triples, quadruples, etc.

Monte Carlo simulations of the PolarisJ detectors in high radiation-flux environments showed that a significant number of double scatters (i.e. events that the PolarisJ camera recorded as a single gamma-ray scattering twice within the detector) were "false" doubles. This is when two gamma-rays interact with the detector close enough in time to be recorded as a single event scattering multiple times. This is opposed to a "true" double scatter, when one gamma-ray genuinely scatters within the detector and deposits energy twice. Therefore even if a True coincidence interacts only once with the detector, it could be obscured by another event arriving and causing a false double. Similarly, about half of these occurrences will result in a True coincidence being lost.

A brute force approach to resolving this would be to include all possible permutations of coincident LORs, and then rely on the PEPT algorithm to remove the erroneous lines. A more elegant approach would be to look at the energies being deposited - if one or both interactions in a double scatter have a deposited energy within the photopeak (or on the Compton Edge), then it was a false double. If the deposited energies sum to land under the photopeak (or Compton Edge) then it was a true double. Even more sophisticated would be to use the Compton Scattering formula and known locations and energies of the scatters to determine whether a recorded double is possible or not, which would indicate false and true scatters. There is much research that could be done in this direction.

Chapter 5

Experiment 3 - Moving Tracer

After fully characterising the PEPT performance on a stationary source, the next step was to test it on a moving source. For this, the opportunity was presented to work on an experimental detector array built by H3D Inc. at their labs in Ann Arbor, MI. Primarily built to do small animal PET imaging, it consisted of a four-module PolarisJ array, each with four crystals, arranged in a square geometry with a rotating disc in the centre driven by a stepper motor. Not only did this system have four times the number of crystals, it was also arranged in an optimal geometry, giving equal coverage in both the x- and y-coordinates. The performance of this system provided an excellent concept of what a micro-PEPT CZT modular array would look like. Since these experiments were performed at the H3D Inc. factory, the future capabilities of these detectors and the implications that would have for CZT PEPT were discussed with the engineers at H3D Inc. Additionally, this small animal PET array was taken to the Knoll Lab run by Zong He, situated in the nearby University of Michigan, to run additional experiments and to further discuss this work with his research group there.

It's important to emphasise that this small animal PET imager was a prototype device assembled by H3D Inc. for proof-of-concept experiments, and therefore did not perform at a standard one might expect from a commercially sold device. This meant that a small fraction of the datasets taken on the device were either partially or fully corrupted. The device was also only available for a two week period, and so because many of these corruptions were discovered after the two weeks there was no opportunity to retake these data. However, these corruptions were never catastrophic and overall the device worked as expected.

There were additional limitations on the precision of both the tracer sources and the circular motion performed by the rotating disc (radius and speed specifically). Two ^{22}Na sources were used for the experiments: a 37 kBq source and a 1.2 MBq source. The ^{22}Na source material within its housing was distributed across a diameter of ≈ 1 mm for the 37 kBq source and ≈ 2 -3 mm for the 1.2 MBq source - ideally they would need to be distributed over less than 100 μm to appear more point-like. The precision of the source stand and stepper motor arrangement was also unknown and there was not enough time to properly calibrate it to determine the precision. Therefore the known radii and speeds of the various circular tracks were based on assumptions that were likely to have uncertainties on the order of millimetres.

Ideally, a similar set of experiments would have been performed at the PEPT Cape Town lab using the UCT PolarisJ array. These experiments would have also used more ideal ^{68}Ga tracer particles and the modified 3D printer to more precisely measure the limits of the tracking ability of CZT PEPT. However, due primarily to

the COVID-19 pandemic we were unable to access the PEPT Cape Town lab to perform these experiments, and prior to the pandemic a series of equipment malfunctions meant that we were unable to perform any useful moving source experiments. This should therefore be the initial focus of any follow-up work and will be important for providing the precise benchmarks of PolarisPEPT.

5.1 Experimental Setup

The H3D small-animal PET system consisted of four Polaris detector modules (see Section 2.2), with each module holding four ($20 \times 20 \times 10$) mm³ CZT crystals, giving a total of sixteen crystals. This gave a FOV of slightly less than ($77 \times 77 \times 42$) mm³, since there was no coverage in the corners. Placed in the centre of this FOV was a round plastic disc of radius 30 mm, driven by a stepper motor controlled by an Arduino Uno, with a step resolution of 700 steps/revolution. The source was held by a small plastic stand in approximately the centre of the FOV in the z-direction. The stand fitted into slots cut out of the disc, allowing for placement of the source from the centre of the disc in steps of 0.5 mm from 0.5 - 10 mm. This allowed the source to be rotated in the FOV of the array in 0.5 - 10 mm circles up to a speed of around 50 mm/s. The uncertainty on the known radii was ± 1 mm and the uncertainty on the known speeds was ± 1 mm/s. A more complete description of the system with figures is presented in Section 2.4.

Two sources were used for tracking measurements. The first was a 37 kBq ²²Na source embedded in a 2.5 cm diameter plastic disc with the source material distributed over a 1 mm sphere, owned by H3D Inc. The second was a 1.2 MBq ²²Na source embedded in a 2.5 cm plastic disc with the source material distributed over a 2-3 mm with unknown geometry, supplied by Zong He's group at the University of Michigan. Since these were the only sources available, it was not possible to generate a short-lived isotope to observe the deadtime behaviour or optimise the source activity and geometry for this particular array.

5.2 Stationary Source Performance

5.2.1 Low Activity ²²Na Source

A five minute stationary measurement of a 37 kBq ²²Na source in the centre of the FOV was taken. The total raw singles rate summed over all four modules was 8.2 kHz. Using a coincidence timing window of 5 μ s and no energy filter, the total Prompt coincidence rate was 1 kHz with a True coincidence rate of 450 Hz. A top-down view of the first 1000 coincidences in xy-plane is shown in Figure 5.1.




FIGURE 5.1: A sample of the first 1000 coincidence LORs from a ^{22}Na source of activity 37 kBq, plotted as straight lines between the detected positions of the annihilation gamma-rays. Plotted to show convergence on the source position.

F-optimisation curves were then run on the coincidence data for N-values of 40, 100, 200, 500 and 1000; the resulting curves are shown in Figure 5.2. The f_{opt} values were between 0.35 - 0.45. For $N = 40$ the radial location uncertainty was $580 \mu\text{m}$, giving a final location rate of 25 Hz. For $N = 1000$ the radial location uncertainty was $92 \mu\text{m}$, giving a final location rate of 1 Hz. This was an approximately five-fold improvement in location rate for similar uncertainties compared to the low-activity source in the UCT system presented in Section 3.1.3, primary due to the four-fold increase in the number of crystals¹. The H3D system also had an additional geometric advantage: because it had two pairs of opposing detectors that were orthogonal to each other, the uncertainty between the detectors was significantly reduced where the angles between LORs were very acute. For a visual depiction of this, compare the convergence of the LORs along the x-axis in Figure 5.1 with those in Figure 3.4.

¹Note: at the time of experiment, when one of the four crystals in a module triggered an event, all four crystals went dead while that event was processed. Future generations of CZT detectors from H3D plan to have separate readouts to reduce deadtime effects.




FIGURE 5.2: The 37 kBq ^{22}Na source position uncertainty plotted as a function of the PEPT F-parameter to determine the optimal F-parameter to minimise the position uncertainty. Several different lines-per-position were chosen, shown in the figure legend. A polynomial has been fit to each curve for clarity.

5.2.2 High Activity ^{22}Na Source

A five minute stationary measurement of a 1.2 MBq ^{22}Na source in the centre of the FOV was taken. The total raw singles rate was 61 kHz summed over all four modules. Using a coincidence timing window of $5\ \mu\text{s}$ and no energy filter, the total Prompt coincidence rate was 4.7 kHz with a True coincidence rate of 750 Hz. F-optimisation curves were run on the coincidence data for N-values of 250, 350, 500, 1000 and 2000, with the results shown in Figure 5.3. The f_{opt} values were between 0.05 - 0.10. For $N = 485$ the radial location uncertainty was $515\ \mu\text{m}$, giving a final location rate of 10 Hz. For $N = 4000$ the radial location uncertainty was $160\ \mu\text{m}$, giving a final location rate of 1 Hz.

Surprisingly, the stationary source performance was worse for the 1.2 MBq source than for the 37 kBq source (which had a location rate of 25 Hz for a radial uncertainty of $\approx 580\ \mu\text{m}$ and 4 Hz for a radial uncertainty of $\approx 150\ \mu\text{m}$), despite an increase in the True coincidence rate from 450 Hz to 750 Hz. Ultimately this was caused by the difference in source distribution between the two sources; the 37 kBq source was distributed over 1 mm, whereas the 1.2 MBq source was distributed over 2-3 mm. As outlined by Leadbeater *et al.* in a SAIP conference in 2017 [41], the uncertainty of a PEPT-measured location can be described by

$$u \propto \frac{w^2 + r^2 + S^2 + \theta^2 + FWHM^2}{fN} \quad (5.1)$$




FIGURE 5.3: The 1.2 MBq ^{22}Na source position uncertainty plotted as a function of the PEPT F-parameter to determine the f_{opt} . Several different lines-per-position were chosen, shown in the figure legend. A polynomial has been fit to each curve for clarity.

where u is the uncertainty, w is the pixel resolution of the detectors, S is the spatial distribution of the source, θ is the annihilation photon acollinearity and fN represents the fraction of uncorrupted coincidences (True coincidences). The $FWHM$ term represents other factors not considered in detail here, such as coincidence window, deadtime, higher multiplicity gamma emissions, etc. Note that the uncertainty u scales proportionally with the source distribution S , but scales inversely proportionally with the square-root of fN . Therefore, even though the True coincidence rate (fN) almost doubled from 450 Hz to 750 Hz, this was insufficient to compensate for the more than doubling of the source distribution (S) from 1 mm to 2-3 mm, and hence the radial uncertainty of the 1.2 MBq source was greater than the 37 kBq source for similar location rates.

This result serves to emphasise an important consideration that must be made when performing PEPT using CZT detectors. The low location uncertainties achievable with PolarisJ are due to their high pixel resolution - the w term in (5.1). However, their low timing resolution means that the measured data-rates are low compared to similar BGO arrays. Due to the asymmetry in the relationship between u and the factors in the numerator of (5.1), and u and the denominator of (5.1), increasing data-rates (fN) at the cost of proportionally increasing source attributes such as r , S or θ is unlikely to be advantageous. In scintillator systems, the w term in (5.1) typically dominates the numerator, and so it is usually more beneficial to maximise fN at the cost of r , S or θ . For semiconductor systems w is no longer the dominant term, and since fN is typically limited, researchers must aim to maximally reduce

the r , S and θ terms through careful choice of tracer sources to achieve the best possible results.

Despite the comparatively worse performance of the higher activity source, the results of the experiments performed are still presented in this Chapter. This is because the range of speeds measured using the 1.2 MBq was greater than that of the 37 kBq source, and revealed certain upper limits of trackable tracer speeds.

5.3 Circular Motion Tracking Performance

5.3.1 Low Activity ^{22}Na Source

The first series of experiments used the 37 kBq ^{22}Na source. Using the central stepper-driven disc, six radii were chosen: 0.5 mm, 1 mm, 3.5 mm, 4 mm, 8 mm and 19.5 mm. For each of these radii the disc was run at the following speeds: 0.25 mm/s, 0.5 mm/s, 1 mm/s, 2.5 mm/s and 5 mm/s. A total of five revolutions were run for each combination. After PEPT processing, a sinusoidal curve was fit to the x- and y-coordinates and a straight line to the z-coordinate using the Levenberg-Marquardt least-squares algorithm. The equation used for the sinusoidal fit was $r_{x/y}(t) = A_{x/y} \sin(\omega_{x/y}t + \phi_{x/y}) + B_{x/y}$ where $A_{x/y}$ was the amplitude of the motion (i.e. the radius of the circle), $\omega_{x/y}$ was the angular frequency and $\phi_{x/y}$ and $B_{x/y}$ were to account for the initial position of the particle. The x- and y-speeds were therefore $v_{x/y} = A_{x/y} \cdot \omega_{x/y}$. The fitted radii and speeds of each run were then compared to the known radius and speed. From the stationary F-optimisation results, the PEPT algorithm parameters were chosen to be $N = 200$ and $f_{opt} = 0.40$, which gave a stationary radial location uncertainty of $215 \mu\text{m}$ and location rate of 5 Hz.

PEPT Results for Tracer Moving at 5 mm/s around a 4 mm Radius

The analysis for a single set of motion parameters are presented here to demonstrate how analysis was performed for all the results in this Chapter. A "track" will refer to a tracer moving in a circle with a particular radius at a particular speed. The expected motion was a 4.0 ± 1.0 mm radius circle, with an expected speed of 5.0 ± 1.0 mm/s. The PEPT-measured positions as a function of time are shown in Figure 5.4, with the fitted equation parameters shown in Table 5.1. Qualitatively, all three coordinates appear as expected, and the fitted equations fit the coordinates closely. Both the x- and z-coordinates are slightly shifted from the centre of the FOV, indicating that the rotating disc apparatus did not accurately centre the tracer. From the average of the fitted amplitudes and angular frequencies of the x- and y-coordinates, the radius of the PEPT track was $r = 4.211 \pm 0.028$ mm and the fitted speed was 4.4817 ± 0.0036 mm/s. Both of these were within experimental uncertainty of the known values.




FIGURE 5.4: The three PEPT-measured coordinates of the tracer particle as a function of time for the 4 mm radius track moving at 5 mm/s. Fitted over the x- and y-coordinates are sinusoidal curves and the z-coordinate is a flat line, with the fitted parameters and uncertainties shown in Table 5.1. Plotted to judge qualitatively the appropriateness of the fit and quality of the PEPT positions.

TABLE 5.1: The fitted parameters for the circular motion of the 4 mm circular track with the tracer moving at 5 mm/s. A sinusoidal curve was fitted to the x- and y-coordinates using the equation $r_{x/y}(t) = A_{x/y} \sin(\omega_{x/y}t + \phi_{x/y}) + B_{x/y}$, and a flat line was fitted to the z-coordinate. The fit was performed using the Levenberg-Marquardt least-squares algorithm, which also produces a correlation matrix from which the uncertainties were calculated.

Parameter	Fitted Value
A_x	$4.206 \pm 0.019 \text{ mm}$
ω_x	$1.0651 \pm 0.0006 \text{ rad/s}$
ϕ_x	$-0.932 \pm 0.011 \text{ rad}$
B_x	$0.432 \pm 0.014 \text{ mm}$
A_y	$4.215 \pm 0.020 \text{ mm}$
ω_y	$1.0637 \pm 0.0006 \text{ rad/s}$
ϕ_y	$0.651 \pm 0.011 \text{ rad}$
B_y	$-0.039 \pm 0.014 \text{ mm}$
z	$2.1560 \pm 0.0097 \text{ mm}$

For another qualitative view of the results, the xy-coordinates were plotted on a single axis as well as the instantaneous xy-velocities, shown in Figure 5.5. The velocities were calculated from the PEPT-positions using the following six-point method:

$$\vec{v}_i = 0.1 \left[\frac{\vec{P}_{i+5} - \vec{P}_i}{t_{i+5} - t_i} \right] + 0.15 \left[\frac{\vec{P}_{i+4} - \vec{P}_{i-1}}{t_{i+4} - t_{i-1}} \right] + 0.25 \left[\frac{\vec{P}_{i+3} - \vec{P}_{i-2}}{t_{i+3} - t_{i-2}} \right] + \quad (5.2)$$

$$0.25 \left[\frac{\vec{P}_{i+2} - \vec{P}_{i-3}}{t_{i+2} - t_{i-3}} \right] + 0.15 \left[\frac{\vec{P}_{i+1} - \vec{P}_{i-4}}{t_{i+1} - t_{i-4}} \right] + 0.1 \left[\frac{\vec{P}_i - \vec{P}_{i-5}}{t_i - t_{i-5}} \right] \quad (5.3)$$




FIGURE 5.5: **Left:** The measured xy-coordinates of the tracer particle over the duration of the measurement. **Right:** The xy-coordinates of the instantaneous velocity of the tracer over the duration of the measurement, calculated from the positions using the 6-point method shown at (5.2).

Qualitatively both plots look as expected. The positions follow a well-formed circular track with a slight width to the track indicating the statistical fluctuations of the PEPT positions and the velocities are tangential to the positions and appear equal in magnitude. These plots were primarily used as a sanity-check that everything was operating as expected.

PEPT Results for All Motion Parameters

The same analysis was then performed on the datasets for all of the different speed and radius parameters. Figure 5.6 shows the PEPT xy-positions of all the tracks using $N = 200$ with the source moving at a speed of 0.5 mm/s. Qualitatively, the tracks appear clearly defined and accurate to the known radius except for the 0.5 mm radius circle, which does not show a clearly defined circular track. However, for the 0.5 mm track the centre of the source was moving along the edge of the circle, and the source distribution was 1 mm in diameter, meaning that the edge of the source was at the centre of the circle for the entirety of the motion. This meant that the annihilation events were distributed from the centre of the circle to 0.5 mm past the edge, and there was an insufficient number of events to compensate for the relatively large

distribution of annihilation sites compared to the range of motion.




FIGURE 5.6: The xy-coordinates of all of the tracks measured using the 37 kBq source moving at 0.5 mm/s, as measured by PEPT using $N = 200$. The known radii are shown in the Figure legend. Plotted to show the quality of the PEPT tracking.

A more quantitative comparison of the PEPT performance is shown in Figures 5.7 and 5.8. Figure 5.7 shows the PEPT-measured radii compared to the known radii for all of the measured source speeds². All of the measured radii successfully agreed with the known radii within experimental uncertainty. Figure 5.8 shows the PEPT-measured speeds compared to the known speeds for all of the measured source radii. With one exception, all of the measured speeds agreed with the known speeds within experimental uncertainty.

Investigating Figure 5.7 in more detail, the uncertainty on the radius increased as the speed increased, which is unsurprising since the particle moved further between sequential positions for faster speeds. The uncertainty of each measured radius was consistent for the same speed (e.g. the radius uncertainty for a speed of 0.5 mm/s was the same for the 0.5 mm circle, the 1.0 mm circle, the 3.5 mm circle, etc.), with the exception of the 19.5 mm circle where the uncertainties on the measured speeds were relatively much greater. This is likely because at this radius the circular track is moving around the edge of the FOV of the detector array, which reduces the usable

²Note: not all speeds are shown for all radii. This is either because of a corrupted dataset or that experiment wasn't performed due to time constraints.

area of the detector element pairs for coincidence counting.

Investigating Figure 5.8 in more detail, first note that the measured source speeds were corrected by the measured radius. This was because the known speeds were calculated based on the known radius of the track, but as shown in Figure 5.7, the known radii were slightly inaccurate. Therefore all speeds were multiplied by the ratio of the known and measured radii to correct for this. The measured speed for the 0.5 mm track diverged significantly at 2.5 mm/s from the known speed, indicating an upper limit of tracer speed at that radius. To find the precise limits of tracer speeds at this activity (37 kBq) further investigation is required, preferably with much smaller tracer sizes.




FIGURE 5.7: The PEPT-measured radius of each track, as determined from the fitted parameters like in Table 5.1, plotted as a function of the known radius of the track for each speed run (shown in the Figure legend). Plotted to determine whether the fitted radii agree with the known radii for all of the speeds measured. All fits are from the 37 kBq tracer particle and were processed using $N = 200$ lines per position.




FIGURE 5.8: The PEPT-measured speed of each track, as determined from the fitted parameters like in Table 5.1, plotted as a function of the known speed of the track for each radius run (shown in the Figure legend). All speeds have been corrected for the measured radius. Plotted to determine whether the fitted speeds agree with the known speeds for all of the radii measured. All fits are from the 37 kBq tracer particle and were processed using $N = 200$ lines per position.

5.3.2 High Activity ^{22}Na Source

The next series of experiments used the 1.2 MBq ^{22}Na source. Using the central stepper-driven disc, three radii were chosen: 1 mm, 2 mm and 8 mm. For each of these radii the disc was run at the following speeds: 0.25 mm/s, 0.5 mm/s, 1 mm/s, 2.5 mm/s, 5 mm/s, 7.5 mm/s, 10 mm/s and 15 mm/s. A total of five revolutions were run for each combination. After PEPT processing, a sinusoidal curve was fit to the x- and y-coordinates and a straight line to the z-coordinate using the Levenberg-Marquardt least-squares algorithm. The equation used for the sinusoidal fit was $r_{x/y}(t) = A_{x/y} \sin(\omega_{x/y}t + \phi_{x/y}) + B_{x/y}$ where $A_{x/y}$ was the amplitude of the motion (i.e. the radius of the circle), $\omega_{x/y}$ was the angular frequency and $\phi_{x/y}$ and $B_{x/y}$ were to account for the initial position of the particle. The x- and y-speeds were therefore $v_{x/y} = A_{x/y} \cdot \omega_{x/y}$. The fitted radii and speeds of each run were then compared to the known radius and speed. From the stationary F-optimisation results, PEPT algorithm parameters were chosen to be $N = 1000$ and $f_{opt} = 0.08$, which gave a stationary radial location uncertainty of $330 \mu\text{m}$ and location rate of 5 Hz.

PEPT Results for All Motion Parameters

The same analysis that was performed on the 37 kBq source was applied to all of the 1.2 MBq datasets. Figure 5.9 shows the PEPT xy-positions of all the tracks using $N = 1000$ and $f_{opt} = 0.08$ with the source moving at a speed of 0.5 mm/s. Qualitatively, the tracks appear clearly defined and accurate to the known radius.




FIGURE 5.9: The xy-coordinates of all of the tracks measured using the 1.2 MBq source moving at 0.5 mm/s, as measured by PEPT using $N = 1000$ and $f_{opt} = 0.08$. The known radii are shown in the Figure legend. Plotted to show the quality of the PEPT tracking.

A more quantitative comparison of the PEPT performance is shown in Figures 5.10 and 5.11. Figure 5.10 shows the PEPT-measured radii compared to the known radii for all of the measured source speeds³. All of the measured radii successfully agreed with the known radii and with each other within experimental uncertainty. Figure 5.11 shows the PEPT-measured speeds compared to the known speeds for all of the measured source radii. For speeds between 0.25 mm/s and 6 mm/s, all of the measured speeds agreed with the known speeds within experimental uncertainty. For 7.5 mm/s, 10 mm/s and 15 mm/s, three of the measured speeds significantly overestimated the known speeds.

From Figure 5.10, note that as the speed increases, there is a negative gradient on the measured radii. This is a sampling artefact: as the speed increased, the tracer moved further around the circle between two points, and so the measured circle radius decreased. A similar effect is seen in Figure 5.7 for the lower activity source, although not as pronounced.

³Note: not all speeds are shown for all radii. This is either because of a corrupted dataset or that experiment wasn't performed due to time constraints.




FIGURE 5.10: The PEPT-measured radius of each track, as determined from the fitted parameters like in Table 5.1, plotted as a function of the known radius of the track for each speed run (shown in the Figure legend). Plotted to determine whether the fitted radii agree with the known radii for all of the speeds measured.




FIGURE 5.11: The PEPT-measured speed of each track, as determined from the fitted parameters like in Table 5.1, plotted as a function of the known speed of the track for each radius run (shown in the Figure legend). All speeds have been corrected for the measured radius. Plotted to determine whether the fitted speeds agree with the known speeds for all of the radii measured.

5.4 Conclusions

The experiments performed at H3D Inc. and presented in this Chapter demonstrated that PEPT-tracking of a tracer source was possible and was accurate for most of the experiments performed. The H3D small animal PET detector array showed improved stationary location performance over the UCT PolarisJ array - unsurprising as it had four times the number of crystals and an additional pair of modules. For a 37 kBq central stationary source, a True coincidence rate of 450 Hz was measured for a Prompt rate of 1000 Hz, giving a f_{opt} value of around 0.40. A location uncertainty of 580 μm was achieved at a location rate of 25 Hz and a 150 μm location uncertainty for a location rate of 4 Hz. The 1.2 MBq source had worse performance despite an increase in True rate from 450 Hz to 750 Hz. This was primarily due to the more than doubling of the source material distribution within the tracer (2-3 mm), and it is expected that a tracer with geometry on the order of 100s of microns will produce significantly improved location performance.

For both the 37 kBq and 1.2 MBq sources, the H3D small animal PET system was able to track a particle accurately (to within the uncertainty of the known motion parameters) moving in a circular motion for circle radii between 0.5 mm and 19.5 mm. For the 37 kBq source, it was able to accurately measure tracer speeds between 0.25 mm/s and 5 mm/s. For the 1.2 MBq source, it was able to accurately measure tracer speeds between 0.25 mm/s and 10 mm/s.

As with all PEPT tracking, the limits of the technique depend strongly on the typical speeds that the tracer undergoes over the course of its motion. For the ^{22}Na tracers used for these experiments, tracking was accurate for tracers moving at speeds of millimetres/second. For speeds of centimetres/second the tracking was inaccurate. However, due to the distributed nature of the sources and large uncertainty of the motion of rotating disc used in these experiments, more accurate tracking is possible. Further investigation with more ideal tracer particles is required to find the true limits of PolarisJ PEPT tracking.

The H3D small animal PET detector array is a good example of an initial detector array to acquire to perform PEPT measurements for experiments beyond proof-of-concept, given the options available at the time of writing. Having two pairs of detectors allows for a symmetric coverage of the tracer particle in all three coordinates, which is a large geometric advantage over a single pair. One change to be made is that the array should be modular to allow for better positioning of the module pairs, allowing for custom array geometries suited to each system being investigated. Future generations of CZT detector modules from H3D will also have better timing resolution and deadtime performance, making them a promising technology for micro-scale PEPT tracking.

Chapter 6

Proton Beam Tracking using PEPT

As discussed in Section 1.6.2, a potential application of positron imaging using a CZT detector array is in dose verification in proton radiotherapy. It is well-established that there are a number of positron-emitting isotopes produced when a proton beam (typical energies of 50-250 MeV) is incident on an organic target such as human tissue and there has been much research into using these positron-emitters for PET imaging for both real-time and post-treatment dose verification [32, 33, 44]. In this Chapter we propose using positron information to improve Compton Camera Prompt Gamma Imaging through improved filtering. As outlined by Draeger *et al.* [31], one such filter uses the proton beam direction to remove from image reconstruction reconstructed cones that pass too far from the beam-line as they are likely noise. This is known as the Distance of Closest Approach (DCA) filter - see Section 1.6.2 and the work done by Draeger *et al.* for more details on how the technique works and how effective it is. One drawback of this filter is that it requires assumed knowledge of the incoming beam-direction. However, PEPT provides a promising approach to be able to measure the path of the incoming beam, which offers a way to inform the DCA filter without any assumed knowledge. Since Compton Camera imaging uses high resolution semiconductor detectors like PolarisJ, this beam-tracking can be performed *in-situ* without the need for additional detector arrays. Additionally, since the CZT detector array would be arranged to detect positrons in coincidence, it may be possible to also perform PET imaging for further dose-verification. Finally, beamline tracking has applications outside of this particular medical situation, and so this technique will be broadly applicable to any situation where knowledge of the precise path of a proton-beam through organic materials is required.

6.1 Theory

When a beam of high-energy protons enters an organic target there are a number of positron-emitting isotopes produced, primarily through the (p, xn) reaction type. The precise isotopes that are produced and the rates of production are dependent on the beam parameters and material composition, and since the half-lives of the various isotopes vary across several orders of magnitude their contribution to the overall positron emission rates will be time-dependent. For short-lived isotopes ($T_{1/2} < 1$ s) the activated nuclei will not move appreciably before decaying and the resulting position of the positron annihilation will be approximately that of the interaction between a proton and activated nucleus. It is reasonable therefore to assume the distribution of positron emissions will correlate with the distribution of proton-nuclei interactions in the proton-beam. One can therefore model an incident proton-beam as a line-source of positron-emitting isotopes with a distribution around the centre of the beam-line that depends on the width of the proton-beam and isotopes that are

activated, that will spread out towards the Bragg-peak. If a lab coordinate-system is chosen such that the z-axis is aligned with the incoming beam-line, a cross-section of the xy-plane will appear as a "point" source (with some width), with the emitted annihilation photons converging on that point. Therefore the beam-line can be modelled as the integral of a sequence of "point" sources along the central beam-line, where each "point" corresponds to an infinitesimal slice of the beam.

To see how PEPT could be used to track the beam, suppose there is a proton beam incident on an organic target. To detect the produced annihilation photons, a PEPT-style detector array is placed around the beam-target system such that its FOV fully encompasses the beam inside the target. A measurement is then taken of the annihilation photons with the lab coordinate system chosen such that the z-axis aligns with the direction of the incoming beam-line. Next, the detectors are "sliced" up along the z-axis, with the width of the slices chosen to be wide enough so that sufficient LORs fall entirely within these slices to be able to produce a PEPT-position with the desired accuracy - naturally this will depend on the incoming photon flux, solid-angle, measurement time, etc. The PEPT algorithm is then used to find the mean location of the beam-line within each slice. A linear fit could then be applied to the mean position from each of the slices, which would then correspond to the beam-line. This is the proposed method of measuring the beam-line through a target using PEPT.

If the event rates are too low to produce accurate positions using just the slices along the z-axis (linear slicing), additional slices can be taken diagonally across the field of view to increase the number of used LORs (diagonal slicing). To do this, for each detector slice along the z-axis of the crystals in one detector module, instead of looking directly opposite that slice in the opposing module for LORs (i.e. parallel to the xy-plane), one can look at an offset slice either higher or lower in the opposing module. A PEPT-location can then be determined from these diagonal slices and added to the number of locations along the beam-line, thereby utilising more of the detected LORs.

6.2 Experimental Setup

There were initially plans to use the UCT PolarisJ detector array and the proton-beam at iThemba Labs to perform the proof-of-concept experiments to demonstrate this technique. Unfortunately these experiments had to be cancelled due to the COVID-19 pandemic. Therefore the results presented here are from Monte Carlo simulations produced using the Geant4 toolkit, and focus on the specific case of short-lived isotopes produced by a 55 MeV proton beam incident on PMMA. Primarily this is because of work done by Dendooven *et al.* [35] where measurements of the rates of activation of positron-emitting isotopes were done. Additionally, a simulation of a 200 MeV proton beam was performed to show that tracking could be performed on a beam with this energy, as it is more typical of the beam energies used in PRT.

6.2.1 55 MeV Proton Beam

For the 55 MeV proton beam, four PolarisJ modules with four detector crystals each were placed in an array almost identical to the H3D small-animal PET array described in Section 2.4, except the separation distance of opposing modules was 90 mm, giving a FOV of slightly less than $(90 \times 90 \times 42)$ mm. The PMMA target was a $(30 \times 85 \times 85)$ mm block positioned in the centre of the FOV with the (85×85) mm face in the xy-plane. The incident proton beam was aligned with the z-axis. The beam width was 3 mm and the proton range in the block was 22 mm. See Figure 6.1. A total of 10^{10} protons were simulated.




FIGURE 6.1: Schematic of the Geant4 Simulation used for the 55 MeV proton beam. The FOV was slightly less than $(90 \times 90 \times 42)$ mm. The PMMA target was a $(30 \times 85 \times 85)$ mm block positioned in the centre of the FOV with the (85×85) mm face in the xy-plane. The incident proton beam was aligned with the z-axis. The beam width was 3 mm and the proton range in the block was 22 mm. A total of 10^{10} protons were simulated.

6.2.2 200 MeV Proton Beam

For the 200 MeV proton beam, four PolarisJ module arrays similar to what was described in Section 6.2.1 above were arranged adjacent to each other, giving a FOV of slightly less than $(110 \times 110 \times 200)$ mm. The target was a $(300 \times 105 \times 105)$ mm block positioned in the centre of the FOV with the (105×105) mm face in the xy-plane. The incident proton beam was aligned with the z-axis. The beam width was 3 mm and the proton range in the block was 112 mm. See Figure 6.2. A total of 10^9 protons were simulated.




FIGURE 6.2: Schematic of the Geant4 Simulation used for the 200 MeV proton beam. The FOV was slightly less than $(110 \times 110 \times 200)$ mm. The target was a $(300 \times 105 \times 105)$ mm block positioned in the centre of the FOV with the (105×105) mm face in the xy -plane. The incident proton beam was aligned with the z -axis. The beam width was 3 mm and the proton range in the block was 112 mm. A total of 10^9 protons were simulated.

6.3 Results

6.3.1 55 MeV Proton Beam

A total of 10^{10} protons were simulated to produce 15 317 coincidences using an energy filter 510-512 keV. The simulation was not optimised for highly accurate positron production rates and so this number of detected coincidences was much lower than expected for a real proton beam with this geometry. However for a proof-of-concept this was sufficient. Dividing the FOV along the z -axis into 10 equal slices produced slices of 4.25 mm in length with an average of 110 ± 130 coincidences per slice. The PEPT positions were obtained using $f_{opt} = 0.85$ and N was the total number of coincidences in each slice. The large uncertainty on the number of coincidences per slice was primarily due to the Bragg peak occurring over a single slice, which lead to a large proportion of the coincidences landing in one slice, and some slices occurring over gaps between the detector crystals leading to only one or two coincidences in those slices. The f_{opt} was chosen to coincide with the f_{opt} from a 37 kBq ^{22}Na source inside the H3D small-animal PET array with a coincidence window of 5 μs and energy window of 100 kBq on the photopeak, as this most closely resembled the proton beam setup from the datasets available.

The measured (x,y) coordinates of the positions are shown in Figure 6.3 along with the outline of the incoming proton beam. All slices with fewer than 10 LORs were discarded and hence there were only six positions plotted. As is clear from the figure, all of the measured positions fell easily within the distribution of the incoming beam, and were randomly distributed around the centre of the beamline as expected. Figure 6.3 also shows the (x,z) and (y,z) coordinates of the positions, along with a least-squares linear fit plotted over the positions. Once again, all of the positions fell within the incoming beam and the positions were distributed roughly evenly around the centre of the beamline. The linear fit fell very close to the centre of the beam-line.

Next, diagonal slices were added and the analysis was rerun to produce the plot in Figure 6.4. This effectively increased the total number of positions by an order of magnitude, but the spread of the positions was also significantly bigger, particularly on the left-hand side of the Figure. However, from the simulation the Bragg peak ended at around $z = -2$ mm, and so the positions to the left of this are likely just random coincidences or annihilation from positrons that were emitted in the $-z$ direction.

Looking at Figure 6.4 the points were all clustered around the centre of each linear z -slice. Each of these clusters were then averaged, with the resulting plot shown in Figure 6.5. The resulting beam-line fit was not noticeably better than the result in Figure 6.3, but showed that it was possible to effectively utilise more of the detected LORs to do beam tracking.

To demonstrate the usefulness of the diagonal slicing, the total coincidences were reduced from 15 000 to 1000. With 1000 coincidences and only using the linear slicing, there were not enough LORs per slice to produce PEPT positions. After adding diagonal slicing the results are shown in Figure 6.6. All of the positions fell within the known beam-line and the linear fits were very close to the centre of the beam-line. Therefore with diagonal slicing, accurate beam tracking could be performed with an order of magnitude fewer lines.

6.3.2 200 MeV Proton Beam

From the 200 MeV proton beam simulation, 10^9 produced a total of 3590 detected coincidences using an energy filter 510 - 512 keV. Dividing the FOV into 10 equal slices produced slices of 22 mm in length with an average of 45 ± 21 coincidences per slice. To contrast, using 5 slices produced 163 ± 50 coincidences per slice and 20 slices produced 12.5 ± 9.2 coincidences per slice. For the 10 slice dataset the PEPT positions were obtained using $f_{opt} = 0.85$ and N was the total number of coincidences in each slice, similarly to the simulation in Section 6.3.1. Diagonal slicing was performed on the FOV using the 10 slice dataset to produce a final average of 83 ± 73 coincidences per slice. Note that the coincidence counting was done on every possible pair of opposite slices, and the final average coincidences per slice was calculated using each individual pair, giving a total of 100 (10×10) total slice pairs.

The measured (x,y) coordinates of the positions are shown in Figure 6.7 along with the outline of the incoming proton beam. The results were very similar to the results in Section 6.3.1, with all the points falling within the known beam-line and the linear fits approximating the centre of the beam-line accurately. However, the spread of the points was slightly larger. This is likely due to fewer LORs collected due to fewer protons and smaller solid angle around the beam-line, as well as the protons from the higher energy beam spreading out further along their trajectory. However, it was still accurate enough to reproduce the known beam-line direction.

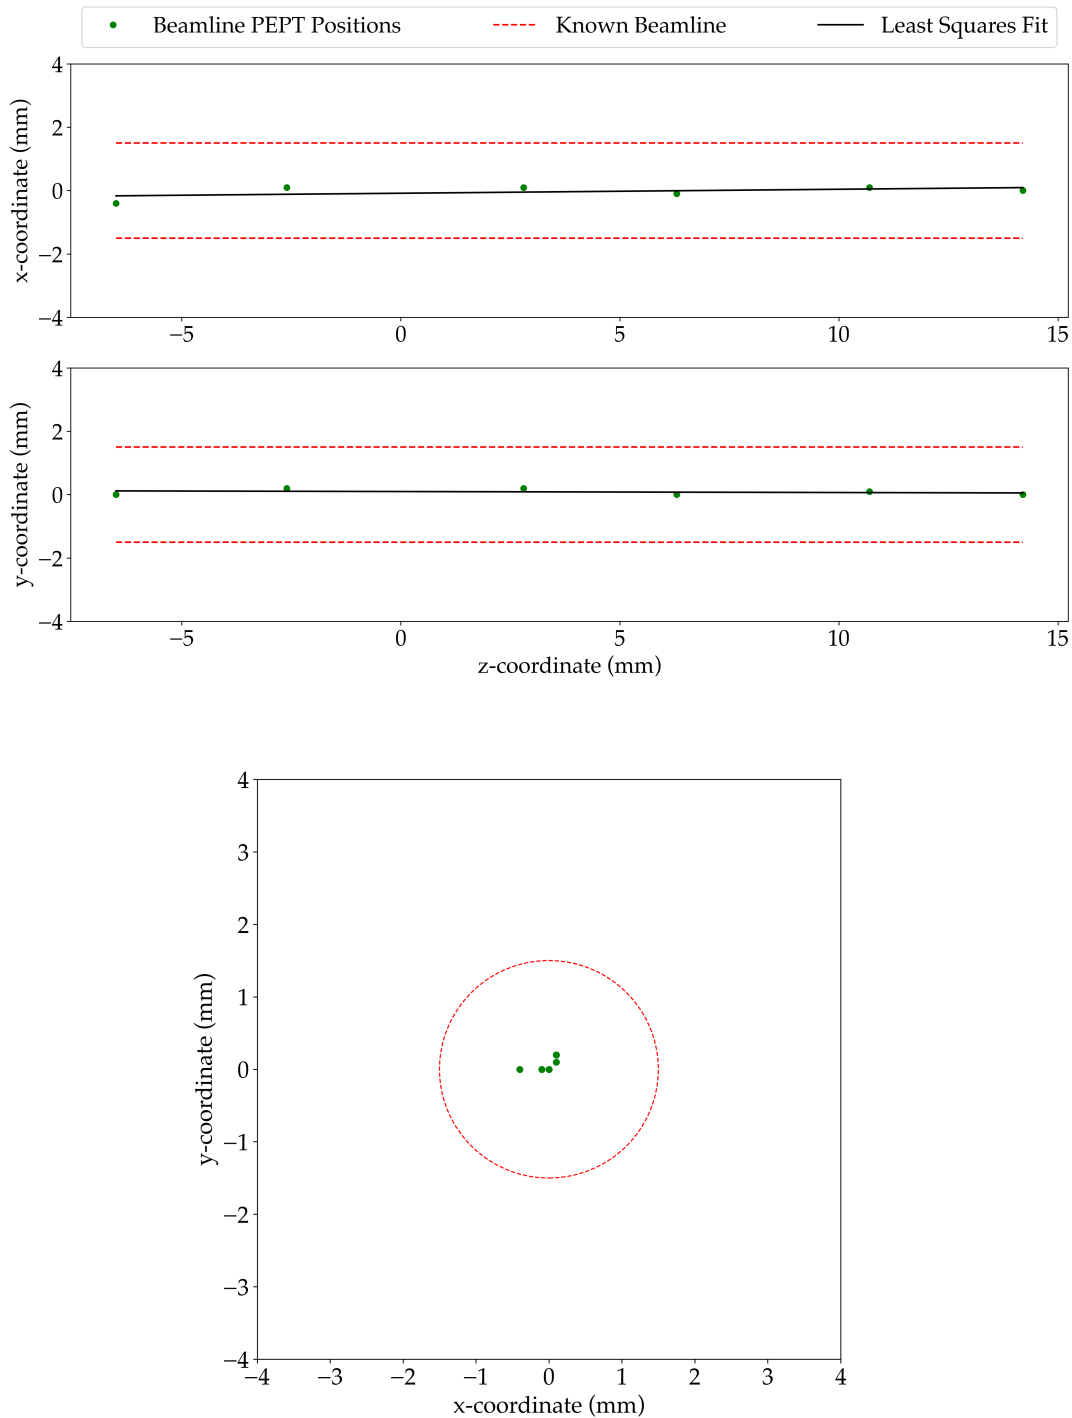


FIGURE 6.3: Results from the 55 MeV proton beam. XZ (top), YZ (middle) and XY (bottom) cross-sections of the PEPT-determined beam-line positions (plotted as green dots) along with the known beam-line (plotted as a dotted red line), determined from 15 317 coincidences using linear slicing. Plotted to verify that that the PEPT-positions fall within the known beam-line and observe how distributed the positions are around the centre. Plotted in black is a linear least-squares fit - this was the final measured beam-line.

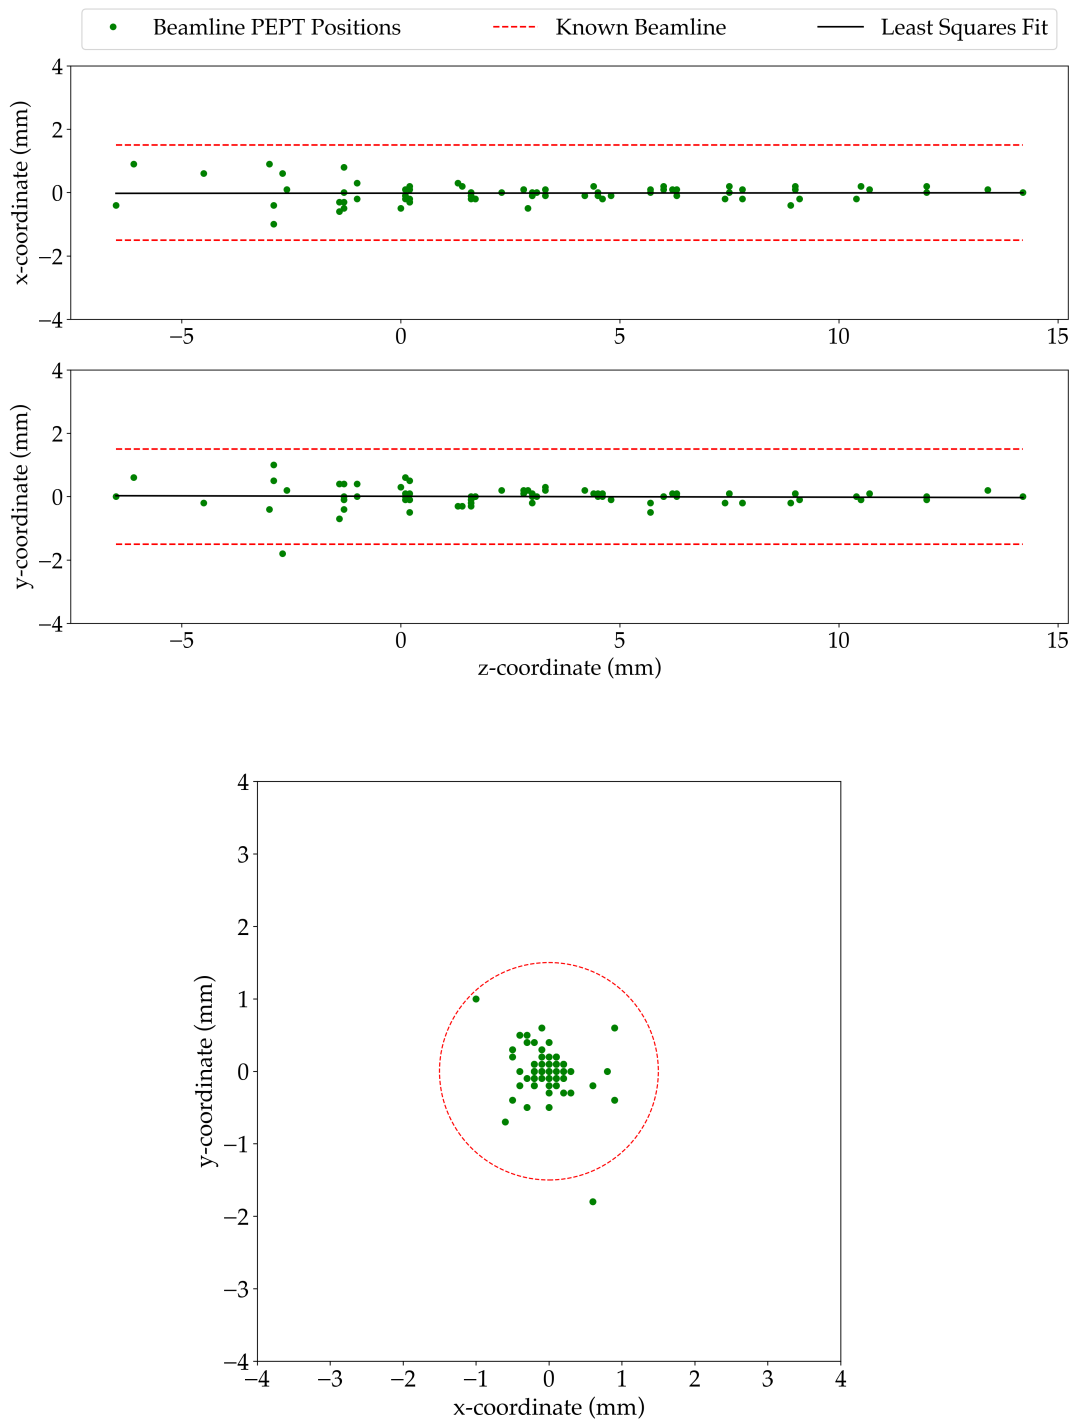


FIGURE 6.4: Results from the 55 MeV proton beam. XZ (top), YZ (middle) and XY (bottom) cross-sections of the PEPT-determined beam-line positions (plotted as green dots) along with the known beam-line (plotted as a dotted red line), determined from 15 317 coincidences using linear and diagonal slicing. Plotted to verify that that the PEPT-positions fall within the known beam-line and observe how distributed the positions are around the centre. Plotted in black is a linear least-squares fit - this was the final measured beam-line.

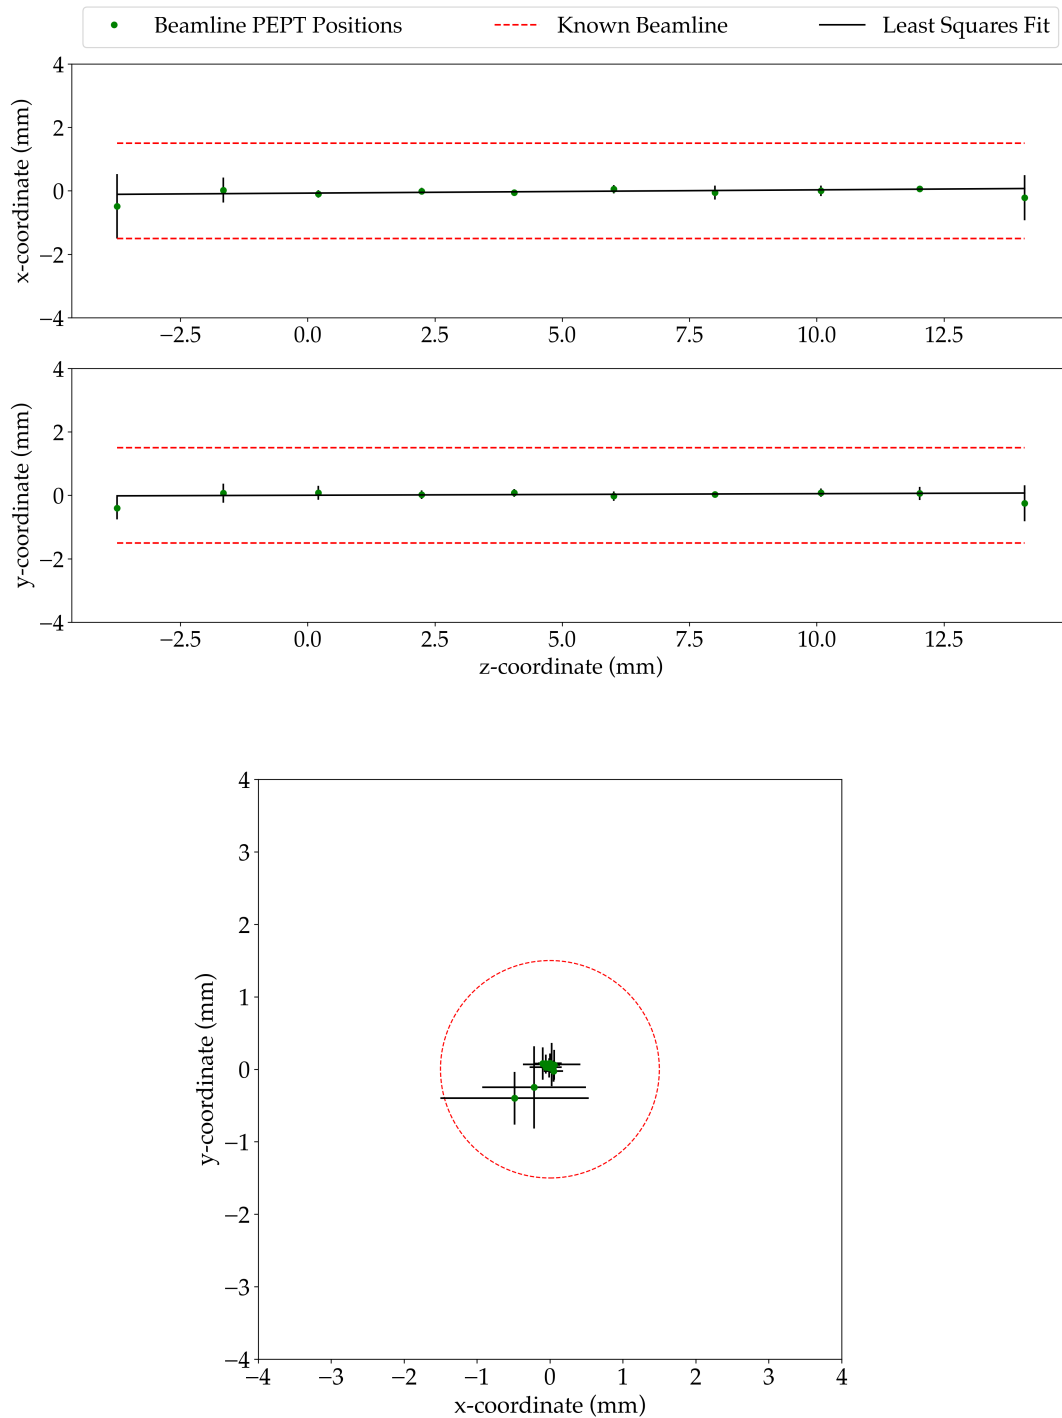


FIGURE 6.5: Results from the 55 MeV proton beam. XZ (top), YZ (middle) and XY (bottom) cross-sections of the PEPT-determined beam-line positions (plotted as green dots) along with the known beam-line (plotted as a dotted red line), determined from 15 317 coincidences using linear and diagonal slicing with positions averaged within the linear slices. The error-bars represent the experimental uncertainty of the averages. Plotted to verify that the PEPT-positions fall within the known beam-line and observe how distributed the positions are around the centre. Plotted in black is a linear least-squares fit - this was the final measured beam-line.

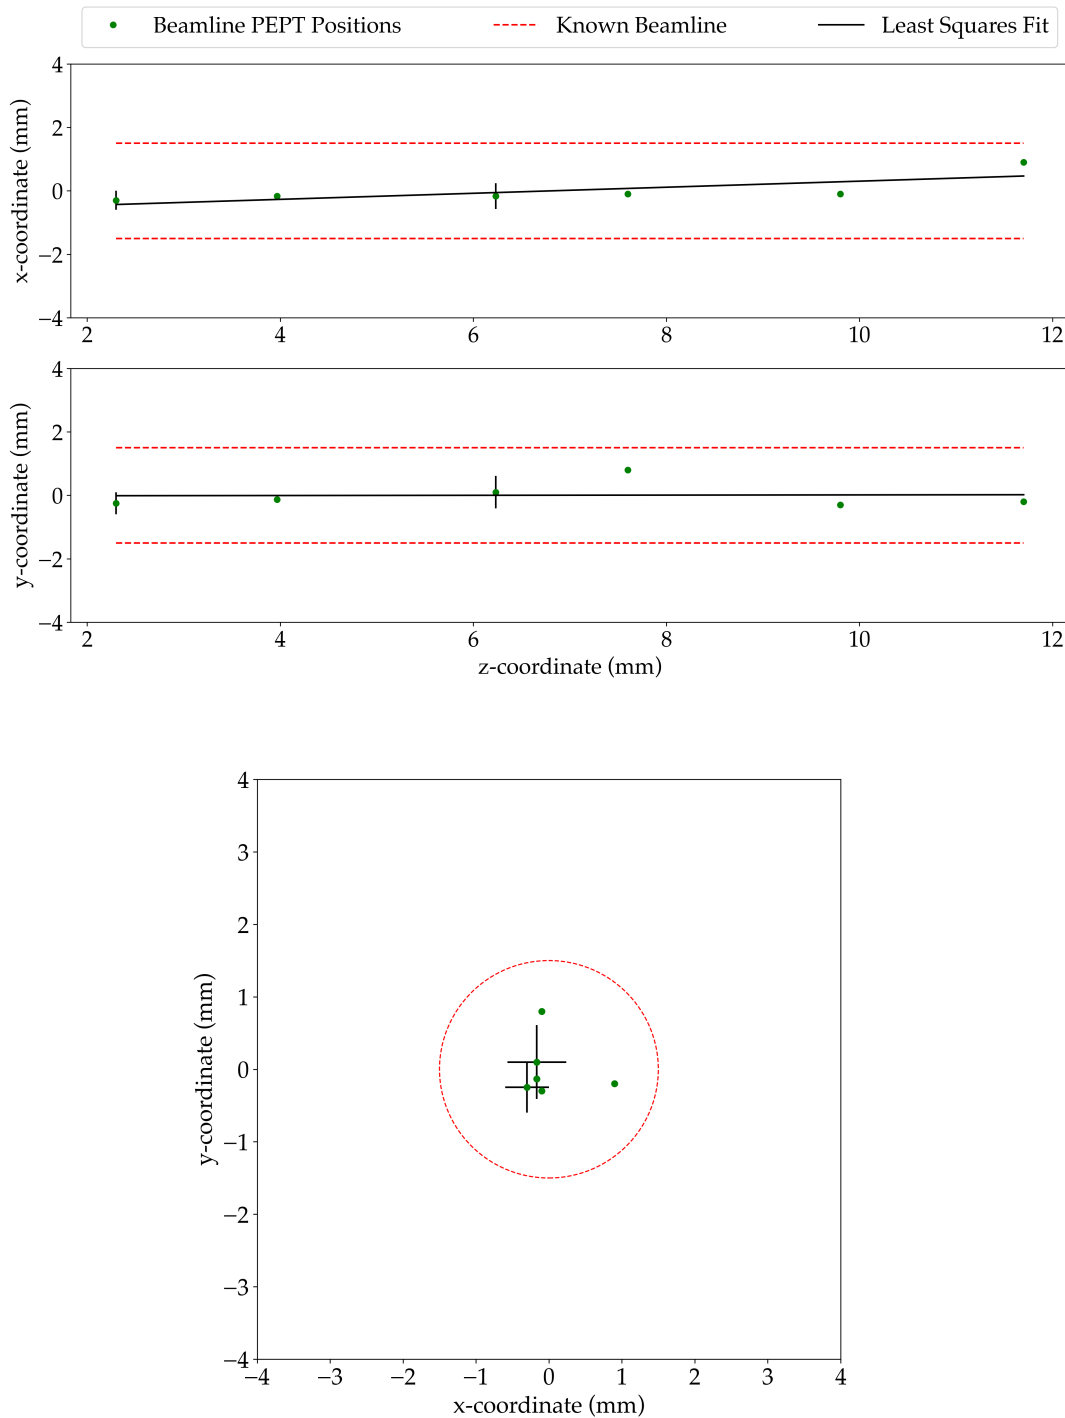


FIGURE 6.6: Results from the 55 MeV proton beam. XZ (top), YZ (middle) and XY (bottom) cross-sections of the PEPT-determined beam-line positions (plotted as green dots) along with the known beam-line (plotted as a dotted red line), determined from 1000 coincidences using linear and diagonal slicing with positions averaged within the linear slices. The error-bars represent the experimental uncertainty of the averages. Plotted to verify that that the PEPT-positions fall within the known beam-line and observe how distributed the positions are around the centre. Plotted in black is a linear least-squares fit - this was the final measured beam-line.

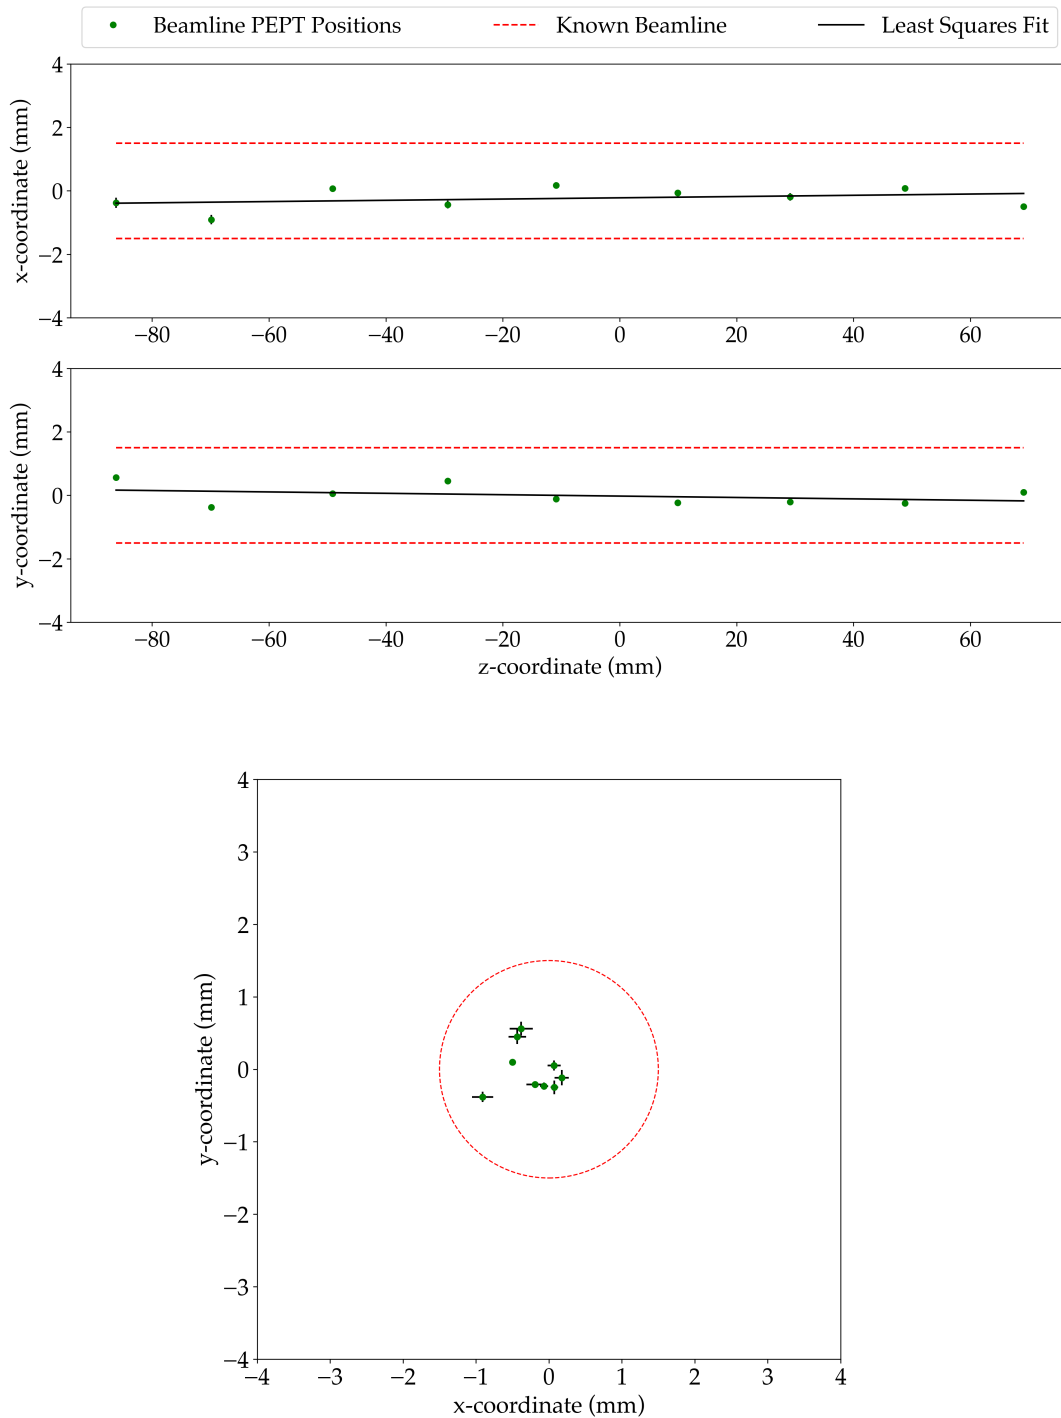


FIGURE 6.7: Results from the 200 MeV proton beam. XZ (top), YZ (middle) and XY (bottom) cross-sections of the PEPT-determined beam-line positions (plotted as green dots) along with the known beam-line (plotted as a dotted red line), determined from 3590 coincidences using linear slicing. The error-bars represent the experimental uncertainty of the averages. Plotted to verify that that the PEPT-positions fall within the known beam-line and observe how distributed the positions are around the centre. Plotted in black is a linear least-squares fit - this was the final measured beam-line.

6.4 Feasibility Estimates

While the results presented were promising, there are several important caveats to these simple Monte Carlo simulations. Firstly, there wasn't any accurate timing simulated, meaning that no positron production rates were obtained from the results. Secondly, as pointed out in a publication by Bongrand *et al.* [44] Monte Carlo simulations produced by the Geant4 toolkit do not accurately simulate β^+ emitting isotope production from a proton beam. In particular, they compared Geant4 simulated rates with measurements done by Dendooven *et al.* [35] and found deviations in the rates between 8% and 174%, as well as significant disagreement in the spatial distribution of the produced positrons, particularly with respect to possible β^+ emitters for "real-time" imaging - i.e. isotopes with very short half-lives like ^{12}N ($T_{1/2} = 11$ ms). What follows here is a discussion based on measured production rates from literature and the PEPT performance results from previous chapters in this thesis to try and provide a rough description of the theoretical real-world beam-line tracking performance using current detector technologies.

This discussion will focus on the work done by Dendooven *et al.*, where researchers performed measurements of positron-emitting isotope production in both water and carbon-rich PMMA with an incident 55 MeV proton-beam. In particular, they focussed on ^{15}O , ^{12}N , ^{30}P , ^{29}P , ^{38m}K , ^{38g}K and ^{11}C . Here we will focus on ^{15}O ($T_{1/2} = 2.04$ min) and ^{12}N ($T_{1/2} = 11$ ms) as they dominate the positron-production rates over their respective decay times and target material.

6.4.1 Real-Time Imaging

"Real-time" imaging refers to being able to perform accurate beamline-tracking with measurements over the time-scale of seconds. This is important if one wants to be able to image the proton beam while the beam is on rather than post-treatment. Dendooven *et al.* measured the production of ^{12}N in PMMA to be 2.4×10^{-4} per 55 MeV proton. A typical proton-beam used in clinical practice is 10^8 protons/second, which translates to 10^4 ^{12}N decays per second when integrated over several seconds. Assuming that all the positrons annihilate within the FOV of the CZT array, this translates to an incoming annihilation photon-pair rate of $\approx 10^4$ pairs/second. From the measurements described in Section 5.2.1, for a 37 kBq ^{22}Na source sitting in the centre of the FOV of the H3D small-animal PET system, the measured Prompt coincidence rate was measured to be 50 Hz using a coincidence window of $5 \mu\text{s}$ and energy window of 461 - 561 keV - an absolute efficiency of $\approx 0.1\%$. Extrapolating, if a PolarisJ array similar to the H3D small-animal PET array surrounds a proton beam incident on PMMA and sufficiently covers the proton range within the PMMA, a Prompt rate of 10 Hz can be expected due to ^{12}N over short time periods (≈ 1 s). This is two orders of magnitude lower than the total number of coincidences used for the beam-tracking results shown in 6.3.2, and so would not be feasible for accurate real-time beam-tracking without a significant increase in the efficiency of the PolarisJ array for a 55 MeV proton beam. For a 200 MeV beam, simulations done by Mashayekhi *et al.* [45] showed that in general the yield of positron-emitting isotopes increased with proton energy, but not by more than an order of magnitude; they also did not run simulations for ^{12}N . Additionally, a 200 MeV proton beam has a much larger range than 55 MeV, and would therefore need many more PolarisJ detectors to effectively cover the area over which the annihilation photons would be emitted.

Therefore it is unlikely that real-time tracking would be possible with a 200 MeV beam either.

6.4.2 Post-Treatment Imaging

"Post-treatment" imaging refers to performing beam-tracking more than 100s after the proton-beam is switched on, which is typically after the proton-beam has been switched off. Over this time period, particularly in water, the dominant positron emitter is ^{15}O . Dedooven *et al.* measured the production of ^{15}O in water to be 5.45×10^{-3} per 55 MeV proton. Integrated over 180 seconds, this produces on the order of 10^{-1} decays per proton. For a typical beam intensity of 10^8 protons/second, this translates to 10^7 total decays, which for the H3D small-animal PET system with a central source translates to a total of 10^4 detected Prompt coincidences. This is 10 times the number of coincidences used for the results in Section 6.3.2, meaning that post-treatment beam-tracking is certainly feasible.

6.5 Conclusions and Further Investigation

The results shown here demonstrate that beam-tracking using PEPT is theoretically possible, particularly over the time-scale of minutes. However, the precision of the tracking will be strongly dependent on the exact specifications of the PolarisJ array used to do the tracking. Naturally, beam-tracking using this technique is independent of the detector system used to implement it, and a more efficient PET system could be used to implement such tracking. However, the appeal of using a CZT array is that these arrays are already being used for Compton Camera imaging and so could also be used to implement this beam-tracking technique without the need for additional detector systems.

Further work to investigate this technique will need to focus on performing experiments to measure the real-life performance possible using the PolarisJ array. For proof-of-concept measurements, the UCT PolarisJ detector array could be used along with the proton beam at iThemba Labs to investigate a proton beam incident on various organic targets over a range of proton energies (50 - 200 MeV). This will be able to address considerations not addressed in the Monte Carlo simulations, such as detector deadtime effects due to incoming gamma-ray and neutron flux, isotope production other than ^{12}N and ^{15}O , absolute efficiency for detecting a line source of positron-emitters (as opposed to a point source), positron range effects and positron isotope production. Additionally, the beam-tracking information can then be used in a DCA filter for a Compton Camera reconstruction of the energy deposition to verify the effectiveness of such a filtering technique.

In addition to a real-world measurement, further work could be done to improve the technique itself. To improve efficiencies, filtering on the Compton edge and looking at multiple scatter events could increase the usable number of coincidence events. The number of slices used for the final tracking is an important parameter that requires careful investigation. For this analysis, all the slices were uniform in size, but non-uniform slices could improve the quality of the tracking. In particular, more and narrower slices could be taken around the Bragg Peak where the proton path is the least easily predictable. Implementing a 2D version of the PEPT

algorithm could potentially speed up the reconstruction and produce more accurate positions with fewer lines. Finally, integrating this technique into the filtering and reconstruction algorithms used in Compton Camera imaging will be important to demonstrate the potential for an all-in-one Compton Camera system for effective dose verification in Proton Radiotherapy.

Chapter 7

Conclusions and Further Investigation

7.1 PolarisPEPT

It is uncontroversial that the Positron Emission Particle Tracking technique has been, and continues to be, a powerful tool in system process analysis, both for industrial process optimisation and pure scientific research across a wide domain. It is also true that this progress was made possible by the many advances in high-efficiency, fast scintillator detectors (NaI and BGO primarily), which remain the gold standard for tracking millimetre-sized particles moving at metre/second speeds with millimetre precision at kHz rates. However, to perform at this optimal level these scintillators require that the tracer particles achieve a certain minimum activity to produce the necessary statistics to reduce the tracer location uncertainty. For millimetre-sized particles, this is easily achieved with current tracer production techniques (for ^{68}Ga and ^{18}F), but as researchers begin to investigate mixed-phase (solid-liquid) systems, the demand for ever-smaller tracer particles grows. Despite advances in tracer production techniques, for particles smaller than $500\ \mu\text{m}$ it is very difficult to produce tracers with activities high enough to achieve optimal performance. This prompted researchers to look into alternative detector technologies for a potential solution to this problem.

One such alternative is the room-temperature semiconductor detector, such as the Cadmium-Zinc-Telluride modules made by H3D Inc. (Ann Arbor, MI, USA). These are modular detectors with energy and pixel resolution orders of magnitude higher than BGO, but with timing resolution, efficiency and physical size orders of magnitude lower. Naturally, several questions arise: how would the PEPT technique perform? Would it be possible to track micron-sized particles to micron-precision? What sort of applications would best suit them? These are the questions that this dissertation set out to answer.

The detectors used in this investigation were the UCT PolarisJ CZT array. They were bought by Steve Peterson for research into Prompt Gamma Imaging used in Proton Radiotherapy, and consisted of two identical modules with two CZT crystals in each, perfect for the initial proof-of-concept experiments. The first experiment showed that these detectors were able to locate a stationary ^{22}Na source in three-dimensional space to 100-micron location uncertainties. The next experiment used a short-lived ^{68}Ga source produced by Mike van Heerdan at the UCT PEPT lab in Cape Town and determined the deadtime-limits of the PolarisJ modules and how this impacted their PEPT performance. The final PEPT experiments were performed

in Ann Arbor, Michigan with collaborators at H3D Inc. and the University of Michigan, Ann Arbor on a four-module, sixteen-crystal array built by H3D Inc. This array was used to track two moving ^{22}Na tracer particles with different activities and determined the limits on tracer speed and position uncertainty using PEPT.

The final Chapter of this thesis investigated a potential application of semiconductor PEPT: could the PolarisJ detectors be used to track a proton beam incident on an organic target, and how accurate would this tracking be for different proton beam energies? The primary motivation for this study was to see if this beam tracking information could be used in Distance of Closest Approach filtering used in Prompt Gamma Imaging to improve image quality. The advantage being that the PolarisJ detectors are being used to perform the imaging, and so this additional information could be obtained without the need for more and/or different detectors. A proof of concept experiment was performed using Geant4 simulations, and the feasibility of beam tracking was discussed using experimental results both from this work and the literature.

7.1.1 Locating a Stationary Point Source

The initial PEPT experiments on a stationary ^{22}Na tracer particle of 17.8 ± 3.6 kBq, described in Chapter 3, demonstrated that the PolarisJ CZT modules are capable of implementing PEPT and successfully located the tracer in three-dimensional space. An energy window of 15 keV and coincidence timing window of 600 ns was chosen for coincidence processing. The F-optimisation produced f_{opt} values of between 0.90 and 0.95 for $N = 30$ to $N = 200$ lines-per-position. The location uncertainty was $546 \mu\text{m}$ for $N = 15$ and $101 \mu\text{m}$ for $N = 200$. These results showed a step-change in performance compared to a similar BGO system, which has typical f_{opt} values of 0.5 to 0.6, and would obtain a location uncertainty between 1-2 mm for N values on the order of 100s.

The improvement in f_{opt} was due to two features of the experiment. Firstly, the high energy resolution of 1.15% in the semiconductor detectors was able to exclude most of the background noise (in-scatter events, Compton-scatter events from the 1274 keV line, cosmic radiation, etc.) at the energy filter. Secondly, because the source activity was low, the detected event rate in the modules was also low, which ensured that the Random coincidence rate was very low. The measured Delayed rate was <1 Hz. The improvement in location uncertainty was mostly due to the 1.8 mm pixel size and sub-millimetre depth-of-interaction information offered by the PolarisJ modules. A similar BGO array has a pixel size of $\approx 4 \times 4$ mm and no depth-of-interaction information. For the CZT modules, the primary source of uncertainty in the location of the tracer was due to source distribution, positron range and acollinearity of the annihilation photons.

While the location accuracy of these results showed an order-of-magnitude improvement over a similar BGO array, the location rate and field-of-view was not comparable. The measured Prompt coincidence rate was 4 Hz, which gave a maximum location rate of the tracer of 0.27 Hz using $N = 15$ (for a stationary location uncertainty of $546 \mu\text{m}$). A similar BGO array would offer a location rate in the 100s of Hz. Similarly the FOV was $(62 \times 42 \times 20) \text{ mm}^3$, which is 2-3 times smaller than a similar BGO array. For this activity source, with this detector array in the described geometry, to track a moving tracer would require a very slowly moving tracer or

very smooth tracer motion, which strongly restricts the range of systems that could be investigated with this array. The next experiment performed therefore focussed on maximising the possible location rate, which required a detailed investigation into the deadtime limits of PolarisJ.

7.1.2 Deadtime Behaviour

For the deadtime experiment, as described in Chapter 4, a $160 \mu\text{m}$ ^{68}Ga source with activity 4.72 ± 0.24 MBq was placed in the centre of the FOV of the two PolarisJ modules and observed over five half-lives until the activity decayed down to 160 kBq. The maximum measured singles event rate in each module was 27 kHz at 3.5 MBq, and the modules primarily exhibited non-paralysable deadtime behaviour, with a fitted deadtime parameter of $24.177 \pm 0.096 \mu\text{s}$. The modules measured a singles event rate of 4 kHz each at the lowest activity of 160 kBq.

The deadtime effects on coincidence rates were then investigated. To successfully model the Prompt, True and Random coincidence rates, no energy window was applied and a timing coincidence window of $10 \mu\text{s}$ was used (see Section 4.1.4 for a detailed explanation). The Prompt coincidence rate reached a maximum of ≈ 7 kHz at 3.5 MBq with the Random rate at a maximum of 6.5 kHz. The True coincidence rate reached a maximum rate of 600 Hz at a source activity of 0.8 - 1.5 MBq, which corresponded to a raw measured singles rate of between 15 and 20 kHz in each module. The measured data fitted closely with the coincidence models as described in Section 1.4.

The PEPT performance on the stationary source was characterised for an activity of 0.8 - 1.5 MBq. The f_{opt} value was between 0.10 and 0.12, primarily due to the wide timing window ($5 \mu\text{s}$) and source activity increasing the Random rate significantly. Sub-millimetre location uncertainty was obtained using $N = 200$, with an uncertainty of $\approx 800 \mu\text{m}$, which had a location rate of 11 Hz. This was a 50-fold improvement over the 18 kBq ^{22}Na source which had an location rate of 0.27 Hz for a location uncertainty of $\approx 550 \mu\text{m}$. Therefore for the four-crystal UCT PolarisJ array, for sub-millimetre PEPT the maximum achievable location rate is ≈ 10 Hz. The source used for this experiment was ^{68}Ga glued to the tip of a hypodermic needle with half the positrons being emitted into air. This result could be improved by producing a positron emitter with lower positron range and annihilation photon acollinearity, such as ^{18}F , or placing the tracer in a dense encasing material.

Further investigation into the coincidence processing parameters showed that for all the experiments performed, it was preferable to use no energy filter and a coincidence timing window of $5 \mu\text{s}$. Compared to a 15 keV energy window and 600 ns timing window, the location rate could be improved by an order-of-magnitude for a similar location uncertainty. The timing window was increased because buffering effects in the modules added a $\approx 1 \mu\text{s}$ time-delay into the modules, significantly frequently to cause around half of the True coincidences to be excluded by the 600 ns timing window. However, the biggest improvement was made by removing the energy filter; because the scattering cross-section for 511 keV gamma-rays in CZT is over four times the photoabsorption cross-section, and because the crystal depth of 10 mm was less than the mean free path of 19 mm, most of the annihilation photons

were scattering out of the detector and thus being excluded by the tight 15 keV energy window.

The results in Chapter 4 showed that PolarisPEPT's biggest weakness was its timing resolution and therefore maximum achievable event rate. The maximum True coincidence rate was 600 Hz (with a Prompt coincidence rate of 2.2 kHz) for a source activity of 0.8-1.5 MBq or a singles rate of 15-20 kHz. Compare this to a single pair of BGO modules ($4 \times 8 \times 8$ crystals) that are similarly deadtime-limited: the singles rate is of the order of 1500-2000 kHz while the Random rate is on the order of 40 kHz using a coincidence window of 10 ns. For the UCT PolarisJ system, the maximum measurable location rates for an optimal activity tracer was tens of Hz for sub-millimetre tracking, which limits trackable tracer speeds to mm/s or perhaps cm/s. A single pair of BGO modules with an optimal activity source could track it at m/s. Therefore if the tracer activity for a particular experiment is not limited, semiconductor PEPT cannot compete with scintillators.

However, despite these relatively low absolute location rates, PolarisJ was still able to perform sub-millimetre tracer locating using relatively few lines-per-position and at relatively low activity. The advantage of PolarisPEPT therefore comes down to context: for an experiment that requires a large FOV, using millimetre-size (or bigger) tracers with a high achievable activity, scintillator detectors easily outperform semiconductor detectors. But as demand for very small tracer particles moving through small systems pushes the tracer activity down, the scintillators cannot obtain the optimally high data-rates that allow them to perform so well. CZT was able to locate a particle to $730 \mu\text{m}$ using just 23 lines-per-position. Additionally, the modularity and relatively small size of the detectors allows experimenters to greatly increasing achievable solid angle by placing the detectors right on the boundaries of a small system. This shows promise for tracking small ($< 500 \mu\text{m}$), low-activity (< 1 MBq) tracers in small systems when BGO cannot.

There are also several ways the PolarisJ PEPT array could be improved to optimised for PEPT. A simple (but expensive) way to improve coincidence rates would be to purchase more detectors to improve the solid angle. However, similarly to the narrative of scintillator PEPT, it will be necessary to improve timing resolution and intrinsic efficiency of PolarisJ to increase the useful fraction of detected events, i.e. increase f_{opt} . New and improved ASIC designs could speed up the time taken to process each recorded event, separate readouts for each crystal within the PolarisJ module would allow them to record events independently, to name a few. At the time of writing, H3D Inc. is working to include such improvements in the next generation of their detector modules. With the next generation of H3D module electronics and a four-module, sixteen-crystal array, it is likely that the performance presented in this chapter could be improved by at minimum an order-of-magnitude. With further developments in optimising tracer particles, two or three orders-of-magnitude may even be possible. An example of what such a system might look like was presented in Chapter 5.

7.1.3 Tracking a Moving Tracer

The moving tracer experiments were performed on the H3D Inc. small animal PET system. This consisted of a four-module, sixteen-crystal array arranged in two vertically opposite pairs that were perpendicular to each other, forming a $(77 \times 77 \times 42)$ mm³ FOV. Two tracers were used: a 37 kBq ²²Na sphere with a 1 mm diameter, and a 1.2 MBq ²²Na source that was roughly circular with a 2-3 mm diameter. Ideally, tracking experiments using the UCT PolarisJ detector array with optimal ⁶⁸Ga sources produced at PEPT Cape Town would have also been performed, but due to issues with faulty equipment and the interruptions caused by the COVID-19 pandemic, these experiments were unfortunately not performed.

The H3D small animal PET detector array showed improved stationary location performance over the UCT PolarisJ array, unsurprising as it had four times the number of crystals and an additional pair of modules. For a 37 kBq central stationary source, a True coincidence rate of 450 Hz was measured for a Prompt rate of 1000 Hz and a f_{opt} value of around 0.40. A location uncertainty of 580 μ m was measured at a location rate of 25 Hz and a 150 μ m location uncertainty was measured at a location rate of 4 Hz. The 1.2 MBq source showed worse performance: despite an increase in True rate from 450 Hz to 750 Hz, the location rate was only 10 Hz for a location uncertainty of 515 μ m, and 1 Hz for a location uncertainty of 160 μ m. This was primarily due to the more than doubling of the source material distribution within the tracer (2-3 mm), and it is expected that a tracer with geometry on the order of 100s of microns would significantly improve performance.

For both the 37 kBq and 1.2 MBq sources, the H3D small animal PET system was able to track a particle moving in a circular motion for circle radii between 0.5 mm and 19.5 mm to within the uncertainty of the known motion parameters. For the 37 kBq source, it was able to accurately measure tracer speeds between 0.25 mm/s and 5 mm/s. For the 1.2 MBq source, it was able to accurately measure tracer speeds between 0.25 mm/s and 10 mm/s.

As with all PEPT tracking, the limits of the technique depend strongly on the typical speeds that the tracer undergoes over the course of its motion. For the ²²Na tracers used for these experiments, tracking was accurate for tracers moving at speeds of millimetres/second. For speeds of centimetres/second the tracking was inaccurate. However, due to the distributed nature of the sources and large uncertainty of the motion of rotating disc used in these experiments, more accurate tracking is possible. Further investigation with more ideal tracer particles is required to find the true limits of PolarisJ PEPT tracking.

The H3D small animal PET detector array is a good example of the ideal starting detector array to acquire to do semiconductor PEPT measurements, given the current options available at the time of writing. Having two pairs of detectors allows for a symmetric coverage of the tracer particle in all three coordinates, which is a large geometric advantage over a single pair. Ideally the array should be modular to allow for better positioning of the module pairs, allowing for custom array geometries suited to each system being investigated. Future generations of CZT detector modules from H3D will also have better timing resolution and deadtime performance, making them a promising technology for micro-scale PEPT tracking.

7.2 PolarisPEPT Further Work

The results presented in the first five Chapters laid the groundwork for semiconductor PEPT. They demonstrated that the technique works with the CZT PolarisJ detectors, and showed promise for tracking micron-sized particles with micron precision moving through small systems. This is currently not possible with the state-of-the-art scintillator PEPT arrays. As these detectors improve their processing electronics and they become cheaper, many potential applications become possible, from investigating mixed-phase floatation-cell-type systems, to micro-scale fluid-dynamics, to combining the outputs of both semiconductor and scintillator systems.

7.2.1 Immediate Follow Up Work

Due to interruptions caused by the COVID-19 pandemic, there are some immediate follow-up experiments that would provide valuable insights into the tracking accuracy possible with PolarisJ. The experiments performed at H3D used tracers that were far from ideal, both in their geometry, isotope and size. A series of experiments with moving sources, using the UCT PolarisJ array in the modified 3D printer, where the tracers are $\approx 100 \mu\text{m}$ ^{68}Ga sources with activities that maximise the True coincidence rate ($\approx 1 \text{ MBq}$) will need to be done to fully describe the limits of tracking precision. A further follow-up could be performed with very small ($\approx 10 \mu\text{m}$) ^{68}Ga sources, to characterise the limits of tracking such particles using PolarisJ.

The current UCT PolarisJ array worked for the proof-of-concept experiments done in this work, but are not ideal for more rigorous PEPT investigations. The author recommends that the PEPT Cape Town laboratory invests in a system similar to the H3D Inc. small animal PET array, with four modules, each with four-crystals to provide a sixteen-crystal system. These modules should be modular to allow for custom arrangement based on the experiments being performed, and if possible should utilise the next generation of H3D Inc. modules with improved ASIC designs and separate crystal readout. Such a system should be capable of tracking 10-100 μm tracers to sub-millimetre precision at speeds of centimetres/second, and would have a FOV on the order of $(100 \times 100 \times 40) \text{ mm}^3$.

7.2.2 Hybrid CZT-BGO Systems

It may be possible to use the advantages of scintillators and semiconductors to complement each other in a combined CZT-BGO array. A simple way to do this would be to place the CZT array around the system-of-interest, and put the whole system in the FOV of a scintillator array like the HR++. The high pixel resolution information provided by the CZT array could then be combined with the high data-rate information from the BGO array. The output of each system could also be used to inform the processing of the other, for example the CZT array could provide an initial position, which is then fed into the PEPT algorithm of the BGO array which uses its high data-rates to further reduce the position uncertainty. There is much promising work that could be done in this direction, and since the UCT PEPT lab already has access to two such arrays, this would not require significant investment to pursue.

7.2.3 Small-Scale Systems

As mentioned, PolarisPEPT only begins to outperform scintillator PEPT for small tracer sizes. An example of where such tracers are necessary are in mixed-phase systems, such as fluidiser beds. In a fluidiser bed, bubbles are used to transport very finely ground mineral particles to the top of a tank filled with a mixture of fluids and these minerals. For certain applications, a PEPT tracer would need to be 10-100 μm in size to behave similarly to these tiny mineral particles. While this is too small to obtain optimally radioactive tracers for current BGO arrays, CZT arrays could potentially be able to track such particles accurately enough to provide insight into the behaviour of such systems.

Another example of a small system where PolarisPEPT could be applied is in lab-on-a-chip systems. These are devices of millimetre to centimetre sizes that implement several laboratory functions in a single integrated system (or chip), utilising microfluidics to automate a wide variety of processes. For example, organ-on-a-chip systems are used to mimic the cellular processes and structure of different human organs for disease research. It is possible that very small PEPT tracers could be placed into these systems and tracked with PolarisPEPT to more deeply investigate the characteristics of such systems.

7.3 Beam Tracking using PolarisPEPT

One application of PolarisPEPT that was performed was using it to perform beam-tracking on a 55 MeV and 200 MeV proton beam incident on a PMMA target. There were plans to use the proton beam available at iThemba Labs along with the UCT PolarisJ array to perform these experiments, but due to interruptions by the COVID-19 pandemic this was not possible. Geant4 Monte Carlo simulations were therefore done to simulate the experiment setup, and results from the measured rates of emitted positrons by a proton beam incident on a PMMA target done by Dendooven *et al.* [35] were used to estimate the feasibility of such an application in a real-world scenario. The reason to develop such a technique was to improve the distance of closest approach filtering technique proposed by Draeger *et al.* [31] to improve imaging quality in prompt gamma imaging used for dose verification in proton radiotherapy.

The results in Chapter 6 demonstrated that beam-tracking using PEPT was theoretically possible over the time-scale of minutes. The measured path of the proton beam moving through the PMMA target agreed with the known path to within the beam width for both the 55 MeV and 200 MeV beams, and the measured positron emission rates presented by Dendooven *et al.* were high enough to obtain these results in a real-world system. However, when looking at the measured positron emission rates for "real-time" beam imaging (performing accurate beamline-tracking with measurements over the time-scale of seconds), it would not be feasible without a significant increase in the efficiency of the PolarisJ array for either the 55 MeV or 200 MeV proton beam.

The precision of beam-tracking was strongly dependent on the exact specifications of the PolarisJ array used to do the tracking. Naturally, beam-tracking using this technique is independent of the detector system used to implement it, and a more efficient PET system could be used to implement such tracking. However, the

appeal of using a CZT array is that these arrays are already being used for Compton Camera imaging and so could also be used to implement this beam-tracking technique without the need for additional detector systems.

7.4 Further Work

Naturally, follow-up experiments to measure the real-world performance of proton beam tracking using the PolarisJ array are crucial. For proof-of-concept measurements, the UCT PolarisJ detector array could be used along with the proton beam at iThemba Labs to investigate a proton beam incident on various organic targets over a range of proton energies (50 - 200 MeV). This will be able to address considerations not addressed in the Monte Carlo simulations, such as detector deadtime effects due to incoming gamma-ray and neutron flux, isotope production other than ^{12}N and ^{15}O , absolute efficiency for detecting a line source of positron-emitters (as opposed to a point source), positron range effects and positron isotope production.

In addition to a real-world measurement, further work could be done to improve the technique itself. To improve efficiencies, filtering on the Compton edge and looking at multiple scatter events could increase the usable number of coincidence events. The number of slices used for the final tracking is an important parameter that requires careful investigation. For this analysis, all the slices were uniform in size, but non-uniform slices could improve the quality of the tracking. In particular, more and narrower slices could be taken around the Bragg Peak where the proton path is the least easily predictable. Implementing a 2D version of the PEPT algorithm could potentially speed up the reconstruction and produce more accurate positions with fewer lines.

Once beam-tracking has been adequately implemented and characterised, the beam-tracking information can then be used in a DCA filter for a Compton Camera reconstruction of the energy deposition of the proton beam to verify the effectiveness of using the DCA filtering technique informed by the tracking information. Integrating this technique into PGI will be important to demonstrate the potential for an all-in-one Compton Camera system for effective dose verification in Proton Radiotherapy.

Appendix A

PolarisJ Module Specifications Sheet




J6400 High-Energy High-Flux Spectrometer

Features

- ✓ Records exact interaction positions and energies from high-energy high-flux gamma rays
- ✓ Better than 1.1% FWHM energy resolution at 662 keV
- ✓ No cryogenic cooling required
- ✓ Compact
- ✓ Operates in high dose rates
- ✓ Precise event time stamps for event correlation
- ✓ Fully user-defined coincidence module
- ✓ 16 independent readout CPUs to maximize event throughput
- ✓ Highly efficient with over 3.9 lbs (1.7 kg) of CdZnTe
- ✓ Two planes of position-sensitive detectors for maximum Compton-imaging efficiency
- ✓ Energy range to over 6 MeV
- ✓ Real-time spectrum view
- ✓ Multi-threaded file transfer over Ethernet
- ✓ Start up in only 2 minutes
- ✓ Software updates included
- ✓ Annual recalibration and software updates included

HY-6



The H3D® J6400 is designed to Compton image high-energy, high-flux gamma rays produced during proton-therapy treatment. Designed for high performance, the compact J6400 has

- High energy resolution
- High timing resolution
- High efficiency
- High count-rate throughput

In a user-friendly package, the J6400 requires only a single power and Ethernet cable and includes PC interface software to download data.

The J6400 can also be used in applications such as

- Security monitoring
- Laboratory measurements
- Fundamental science measurements




Photo and data from Dr. Jeremy Polf

About H3D, Inc.

H3D® is commercializing CZT-based 3D radiation-imaging technologies for nuclear power plant, defense & homeland security, and medical applications. A 2011 spinout from the University of Michigan, we have performed sponsored research for the Defense Threat Reduction Agency, Department of Energy NA-22, and National Institutes of Health. We currently ship H100 to nuclear power plants and research labs around the world, and we have several additional product variants in the development pipeline or undergoing customer feedback trials. Our team has over 100 years of combined experience in Compton Imaging, CZT readout, and system integration. We are privately held, market-driven, and committed to providing our customers with the highest performance and most user-friendly instruments possible.

J6400 Specifications

Dimensions:	15.7 in x 5.3 in x 18.9 in (40 cm x 13 cm x 48 cm)
Weight:	36 lbs (16.3 kg)
Power Supply:	100-240 V, 50-60 Hz
Power Consumption:	100 W at 23° C (73° F)
Startup & Operating Temp.:	10° C to 35° C (50° F to 95° F)
Storage Temperature:	-20° C to 60° C (-4° F to 140° F)
System Cooling:	Peltier cooler and external fan
Energy Resolution:	≤1.1% FWHM at 662 keV
Sensitivity:	Detects ¹³⁷ Cs producing ~3 µR/hr (10 µCi at 1 m) in <3 s
Energy Range:	50 keV to 3 MeV per detector pixel
Crystal Volume:	310 cm ³ CZT (CdZnTe) (smaller and larger volumes available)
Detector Pixels:	7744
Spatial Resolution:	< 1 mm (at 240 kcps) < 2 mm (at 480 kcps)
Timing Resolution:	<100 ns
Coincidence:	FPGA-level master-slave architecture
Coincidence settings:	User defined
Readout Architecture:	16 independent chains, each with CPU, FPGA, 4 ASICs/Crystals
Startup Time:	2 min
Communication Mode:	Ethernet RJ45 port; TCP/IP
Interface Software:	PC-based graphical user interface
Software Capabilities:	Control; real-time spectral view; list-mode data download
Data Storage:	16 removable USB drives (16 GB each) included with system
Warranty:	2 years (includes annual recalibration and software updates)
Includes:	Power cables; mounting brackets



H3D®, Inc. • 812 Avis Drive • Ann Arbor, MI 48108 • USA
Tel +1 734-661-6416 • sales@h3dgamma.com • www.h3dgamma.com

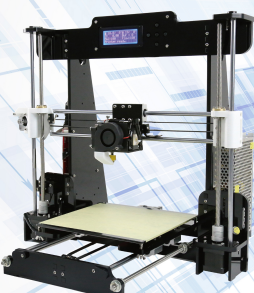
© 2017-2018 H3D, Inc. All Rights Reserved. Polaris-J™ and related systems patent protected by: U.S. Pat No. 7,411,197 & U.S. Pat No. 7,692,155 under license from the University of Michigan.

Specifications, descriptions and images contained in this document were in effect at time of publication. H3D, Inc. reserves the right to change specifications or discontinue products without notice or obligation. All names, logos, and products herein are trademarks of their respective companies.

HY-6

Appendix B

Anet A8 Specifications Sheet



Anet
3D PRINTER

More stable, more finer
Efficient and easy to use Structural stability

A8

the age of 3D virtual to reality is coming , are you ready?
DIRECT FORMING, ALLOWING YOU TO EASILY TEST FOR THE PRECISION OF ART-CRYSTALLIZATION
THE FANTASY JOURNEY OF 3D PRINTED BEGINS FROM HERE

Address: Huizhou, Longhua New District
Shenzhen City, Guangdong Province
No. 341 Avenue, Xiangli Industrial Park,
No. 4 Factory Building


WeChat Scan
Follow us...

National Unified Hotline:
0755 2373 6621



Product Features

- Simple to assemble, which is suitable for DIY customers
- Support 0.1mm printing layer thickness
- Easy to do offline printing from 2004 LCD display with five buttons



High precision
Unique high-precision nozzle design and control structure

Product parameters

Color : Black/Transparent	Package weight : 8.5kg
Structure : Acrylic	Packing Size : 510*310*208mm
Layer Thickness : 0.1-0.4mm	Printing size : 220*220*240mm
Printing speed : 40-120mm/s	Machine size : 500*400*450mm
Printing resolution : ±0.1-0.2mm	Operating System : Windows,MAC
Screen : LCD 2004	XY Axis Position Precision : 0.012mm
Data Input Format : STL,OBJ,JPG	Z Axis Position Precision : 0.004mm
Power Supply : 220V/110V 250W	Printable Filaments : ABS,PLA,HIPS
	Filament Diameter : 1.75mm
	Working Condition : Temp : 10-40°C , Humidity:20-50%

Appendix C

NRW-100 Ion-Exchange Resin Specifications Sheet



Product Data Sheet

PuroLite® NRW100
Nuclear Grade Cation

Nuclear Grade Ion Exchange Resin

PuroLite NRW100 is a nuclear grade strong acid cation exchange resin especially designed for the removal of cations including radioactive isotopes from aqueous solutions. The product contains extremely low levels of other metal cations and is carefully converted to the hydrogen form. It may be mixed in the desired ratio with the strong base anion resin PuroLite NRW400, which is supplied highly converted to the hydroxide form. Such mixtures are capable of removing both cationic and anionic contaminants. All PuroLite nuclear resins are supplied to exacting standards of high purity as given in the specifications below.

Alternatively PuroLite NRW100Li and NRW100Li⁷ are available for removal of radioactive contaminants from primary circuits of nuclear power plants conditioned with natural lithium or lithium⁷.

Typical Physical and Chemical Characteristics

BASIC FEATURES:

Application - Removal of Radioactive Isotopes from Aqueous Solution

Polymer Structure - Gel polystyrene crosslinked with divinylbenzene

Appearance - Spherical beads

Functional Group - Sulfonic Acid

Ionic Form as Shipped - H⁺

PRODUCT INFORMATION:

Total Capacity, (min.)	1.8 eq/l (H ⁺ Form)
Moisture Retention	51 - 55 % (H ⁺ Form)
Particle Size Range	425 - 1200 µm
<425 µm (max.)	2 %
Uniformity Coefficient (max.)	1.7
Ionic Form (min.)	99.9 % (H ⁺ Form)
Impurities Iron (max.)	50 ppm
Impurities Sodium (max.)	40 ppm
Impurities Heavy Metals (max.)	40 ppm
Specific Gravity	1.20 (approx.)
Shipping Weight (approx.)	760 - 800 g/l (47.5 - 50 lbs/ft ³)
Temperature Limit	120 °C (250 °F)

USA
Telephone: +1 610 668 9090
Fax: +1 610 668 8139
Email: info@puro-liteusa.com

Europe
Telephone: +44 1443 229334
Fax: +44 1443 227073
Email: sales@puro-lite.com

Asia Pacific
Telephone: +86 571 876 31382
Fax: +86 571 876 31385
Email: pultalam@puro-litechina.com

Bibliography

- [1] Nicholas Hyslop. "Study of Filtration Techniques for Improving Spatial Resolution and Accuracy of Image Reconstructions used in Prompt Gamma Ray Imaging". B.Sc. Honours in Physics. University of Cape Town, 2018.
- [2] Glenn F. Knoll. *Radiation Detection and Measurement*. Wiley, 2010. Chap. 17, pp. 680–700.
- [3] D.J Parker et al. "Positron emission particle tracking - a technique for studying flow within engineering equipment". In: *Nuclear Instruments and Methods in Physics Research A* 326 (1993). URL: [https://doi.org/10.1016/0168-9002\(93\)90864-E](https://doi.org/10.1016/0168-9002(93)90864-E).
- [4] C.R. Bemrose et al. "Application of positron emission tomography to particle flow measurement in chemical engineering processes". In: *Nuclear Instruments and Methods in Physics Research A* 273 (1988). URL: [https://doi.org/10.1016/0168-9002\(88\)90111-8](https://doi.org/10.1016/0168-9002(88)90111-8).
- [5] D.J. Parker et al. "Positron emission particle tracking using the new Birmingham positron camera". In: *Nuclear Instruments and Methods in Physics Research A* 477 (2002). URL: [https://doi.org/10.1016/S0168-9002\(01\)01919-2](https://doi.org/10.1016/S0168-9002(01)01919-2).
- [6] Alireza Sadrmomtaz, D. Parker, and Larry Byars. "Modification of a medical PET scanner for PEPT studies". In: *Nuclear Instruments and Methods in Physics Research Section A-accelerators Spectrometers Detectors and Associated Equipment - NUCL INSTRUM METH PHYS RES A* 573 (Apr. 2007), pp. 91–94. DOI: [10.1016/j.nima.2006.11.013](https://doi.org/10.1016/j.nima.2006.11.013).
- [7] D.J. Parker et al. "Positron imaging techniques for process engineering: recent developments at Birmingham". In: *Measurement Science and Technology* 19 (2008). URL: <https://doi.org/10.1088/0957-0233/19/9/094004>.
- [8] 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". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 562.1 (2006), pp. 345–350. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2006.03.015>. URL: <http://www.sciencedirect.com/science/article/pii/S0168900206005341>.
- [9] T.W. Martin, J.P.K. Seville, and D.J. Parker. "A general method for quantifying dispersion in multiscale systems using trajectory analysis". In: *Chemical Engineering Science* 67 (2007). URL: <https://doi.org/10.1016/j.ces.2007.02.050>.
- [10] D.J. Parker et al. "Positron emission particle tracking using a modular positron camera". In: *Nuclear Instruments and Methods in Physics Research A* 604 (2009). URL: <https://doi.org/10.1016/j.nima.2009.01.085>.

- [11] T.W. Leadbeater, D.J. Parker, and J. Gargiuli. "Positron emission particle tracking using a modular positron camera". In: *Nuclear Instruments and Methods in Physics Research A* 604 (2011). URL: <https://doi.org/10.1016/j.nima.2009.01.085>.
- [12] T.W. Leadbeater, D.J. Parker, and J. Gargiuli. "Positron imaging systems for studying particulate, granular and multiphase flows". In: *Particuology* 10 (2012). URL: <https://doi.org/10.1016/j.partic.2011.09.006>.
- [13] T J Spinks et al. "Physical characteristics of the ECAT EXACT3D positron tomograph". In: *Physics in Medicine and Biology* 45.9 (2000), pp. 2601–2618. DOI: 10.1088/0031-9155/45/9/313. URL: <https://doi.org/10.1088/0031-9155/45/9/313>.
- [14] A Buffler et al. "PEPT Cape Town: A new positron emission particle tracking facility at iThemba LABS". In: *International Topical Meeting on Nuclear Research Applications and Utilization of Accelerators* (2009).
- [15] K.E. Cole et al. "Positron emission particle tracking measurements with 50 micron tracers". In: *Chemical Engineering Science* 75 (2012). DOI: <https://doi.org/10.1016/j.ces.2012.02.053>. URL: <http://www.sciencedirect.com/science/article/pii/S0009250912001492>.
- [16] T.W. Leadbeater and D.J. Parker. "A high speed PC-based data acquisition and control system for positron imaging". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 604.1 (2009). PSD8, pp. 355–358. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2009.01.184>. URL: <http://www.sciencedirect.com/science/article/pii/S0168900209001922>.
- [17] A. Buffler et al. "Positron emission particle tracking: A powerful technique for flow studies". In: *International Journal of Modern Physics: Conference Series* 48 (Jan. 2018), p. 1860113. DOI: 10.1142/S2010194518601138.
- [18] David Schlyer. "PET tracers and radiochemistry". In: *Annals of the Academy of Medicine, Singapore* 33 (Apr. 2004), pp. 146–54.
- [19] K. Cole et al. "A surface coating method to modify tracers for positron emission particle tracking (PEPT) measurements of froth flotation". In: *Powder Technology* 263 (2014), pp. 26–30. ISSN: 0032-5910. DOI: <https://doi.org/10.1016/j.powtec.2014.04.083>. URL: <http://www.sciencedirect.com/science/article/pii/S0032591014004136>.
- [20] Michael van Heerden. "Optimization of the radio-labelling of ion exchange resin tracers for positron emission particle tracking". Master of Chemical Engineering. University of Cape Town, 2015.
- [21] Glenn F. Knoll. *Radiation Detection and Measurement*. Wiley, 2010. Chap. 4, pp. 121–128.
- [22] Amy Brock and Sui Huang. "Precision Oncology: Between Vaguely Right and Precisely Wrong". In: *American Association for Cancer Research* (2017).
- [23] Nobel Media. *The Nobel Prize in Physics 1901*. 2018. URL: <https://www.nobelprize.org/prizes/physics/1901/summary/>.
- [24] D.J. Shah, R.K. Sachs, and D.J. Wilson. "Radiation-induced cancer: a modern view". In: *The British Journal of Radiology* (2012).
- [25] R.R. Wilson. "Radiological use of fast protons". In: *Radiology* (1946).

- [26] Jeyasingam Jeyasugiththan and Stephen W Peterson. "Evaluation of proton inelastic reaction models in Geant4 for prompt gamma production during proton radiotherapy". In: *Physics in Medicine & Biology* 60.19 (2015), p. 7617. URL: <http://stacks.iop.org/0031-9155/60/i=19/a=7617>.
- [27] J.C. Polf et al. "Detecting prompt gamma emission during proton therapy: the effects of detector size and distance from the patient." In: *Physics in Medicine & Biology* (2014).
- [28] F Fracchiolla et al. "Characterization and validation of a Monte Carlo code for independent dose calculation in proton therapy treatments with pencil beam scanning". In: *Physics in Medicine & Biology* 60.21 (2015), p. 8601. URL: <http://stacks.iop.org/0031-9155/60/i=21/a=8601>.
- [29] Drosoula Giantsoudi et al. "Validation of a GPU-based Monte Carlo code (gPMC) for proton radiation therapy: clinical cases study". In: *Physics in Medicine & Biology* 60.6 (2015), p. 2257. URL: <http://stacks.iop.org/0031-9155/60/i=6/a=2257>.
- [30] M.B. Chadwick et al. "Nuclear interaction cross sections for proton radiotherapy". In: *Nuclear Physics A* 654.1, Supplement 1 (1999), pp. 1051c–1057c. ISSN: 0375-9474. DOI: [https://doi.org/10.1016/S0375-9474\(00\)88598-X](https://doi.org/10.1016/S0375-9474(00)88598-X). URL: <http://www.sciencedirect.com/science/article/pii/S037594740088598X>.
- [31] Emily Draeger et al. "Feasibility Studies of a New Event Selection Method to Improve Spatial Resolution of Compton Imaging for Medical Applications". In: *IEEE TRANSACTIONS ON RADIATION AND PLASMA MEDICAL SCIENCES* (2017). DOI: <https://doi.org/10.1109/TRPMS.2017.2703095>.
- [32] K. Parodi, W. Enghardt, and T. Haberer. "In-beam pet measurements of B+ radioactivity induced by proton beams". In: *Physics in Medicine & Biology* (2002).
- [33] C. H. Min et al. "Clinical application of in-room positron emission tomography for in vivo treatment monitoring in proton radiation therapy". In: *International Journal of Radiation Oncology Biology Physics* (2013).
- [34] Xuping Zhu and Georges El Fakhri. "Proton Therapy Verification with PET Imaging". In: *Theranostics* (2013). DOI: <https://doi.org/10.7150/thno.5162>.
- [35] P Dendooven et al. "Short-lived positron emitters in beam-on PET imaging during proton therapy". In: *Physics in Medicine and Biology* 60.23 (2015), pp. 8923–8947. DOI: [10.1088/0031-9155/60/23/8923](https://doi.org/10.1088/0031-9155/60/23/8923). URL: <https://doi.org/10.1088/0031-9155/60/23/8923>.
- [36] Z. He et al. "3-D position sensitive CdZnTe gamma-ray spectrometers". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 422.1 (1999), pp. 173–178. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(98\)00950-4](https://doi.org/10.1016/S0168-9002(98)00950-4). URL: <http://www.sciencedirect.com/science/article/pii/S0168900298009504>.
- [37] Y. Zhu, S. E. Anderson, and Z. He. "Sub-Pixel Position Sensing for Pixelated, 3-D Position Sensitive, Wide Band-Gap, Semiconductor, Gamma-Ray Detectors". In: *IEEE Transactions on Nuclear Science* 58.3 (2011), pp. 1400–1409.
- [38] Dan Xu. "GAMMA-RAY IMAGING AND POLARIZATION MEASUREMENT USING 3-D POSITION-SENSITIVE CdZnTe DETECTORS". Doctor of Philosophy. University of Michigan, 2006.

- [39] *Automation systems and integration - Numerical control of machines - Program format and definitions of address words - Part 1: Data format for positioning, line motion and contouring control systems*. Standard. Mar. 2009.
- [40] M.J. Berger et al. *XCOM: Photon Cross Sections Database. NIST Standard Reference Database 8 (XGAM)*. NIST, PML, Radiation Physics Division. 2010. URL: <https://www.nist.gov/pml/xcom-photon-cross-sections-database>.
- [41] Thomas Leadbeater. "Uncertainty analysis for Positron Emission Particle Tracking (PEPT) measurements". In: South African Institute of Physics Conference.
- [42] Craig S Levin and Edward J Hoffman. "Calculation of positron range and its effect on the fundamental limit of positron emission tomography system spatial resolution". In: *Physics in Medicine and Biology* 44.3 (1999), pp. 781–799. DOI: 10.1088/0031-9155/44/3/019. URL: <https://doi.org/10.1088%2F0031-9155%2F44%2F3%2F019>.
- [43] Jerimy Polf et al. "The effects of Compton camera data acquisition and readout timing on PG imaging for proton range verification". In: *IEEE Transactions on Radiation and Plasma Medical Sciences* PP (Feb. 2021), pp. 1–1. DOI: 10.1109/TRPMS.2021.3057341.
- [44] A. Bongrand et al. "Use of short-lived positron emitters for in-beam and real-time range monitoring in proton therapy". In: *Physica Medica: European Journal of Medical Physics* 69 (2020), pp. 248–255. DOI: 10.1016/j.ejmp.2019.12.015. URL: <https://doi.org/10.1016/j.ejmp.2019.12.015>.
- [45] Mohsen Mashayekhi, Ali Asghar Mowlavi, and Sayyed Bijan Jia. "Simulation of positron emitters for monitoring of dose distribution in proton therapy". In: *Reports of Practical Oncology and Radiotherapy* 22.1 (2017), pp. 52–57. ISSN: 1507-1367. DOI: <https://doi.org/10.1016/j.rpor.2016.10.004>. URL: <http://www.sciencedirect.com/science/article/pii/S1507136716301444>.