

12-2012

Design, Calibration, and Evaluation of Depth-of-Interaction-Capable PET Detector Modules

Chad J. Bircher

Follow this and additional works at: https://digitalcommons.library.tmc.edu/utgsbs_dissertations



Part of the [Medicine and Health Sciences Commons](#)

Recommended Citation

Bircher, Chad J., "Design, Calibration, and Evaluation of Depth-of-Interaction-Capable PET Detector Modules" (2012). *The University of Texas MD Anderson Cancer Center UTHealth Graduate School of Biomedical Sciences Dissertations and Theses (Open Access)*. 300.
https://digitalcommons.library.tmc.edu/utgsbs_dissertations/300

This Dissertation (PhD) is brought to you for free and open access by the The University of Texas MD Anderson Cancer Center UTHealth Graduate School of Biomedical Sciences at DigitalCommons@TMC. It has been accepted for inclusion in The University of Texas MD Anderson Cancer Center UTHealth Graduate School of Biomedical Sciences Dissertations and Theses (Open Access) by an authorized administrator of DigitalCommons@TMC. For more information, please contact digitalcommons@library.tmc.edu.

DESIGN, CALIBRATION, AND EVALUATION
OF DEPTH-OF-INTERACTION-CAPABLE PET
DETECTOR MODULES

by

Chad Jacob Bircher, B.S., M.S.

APPROVED:

Supervisory Professor, Yiping Shao, Ph.D

Richard Wendt III, Ph.D

Tinsu Pan, Ph.D

Osama Mawlawi, Ph.D

John Cook, Ph.D

APPROVED:

Dean, The University of Texas

Graduate School of Biomedical Sciences

DESIGN, CALIBRATION, AND EVALUATION
OF DEPTH-OF-INTERACTION-CAPABLE PET
DETECTOR MODULES

A

DISSERTATION

Presented to the Faculty of
The University of Texas
Health Science Center at Houston
and
The University of Texas
M.D. Anderson Cancer Center
Graduate School of Biomedical Sciences
in Partial Fulfillment
of the Requirements
for the Degree of

DOCTOR OF PHILOSOPHY

by

Chad Jacob Bircher, B.S., M.S.
Houston, Texas

December, 2012

Dedication: This work is dedicated to my beautiful and supportive wife. Thank you for your relentless dedication during the years I have spent in pursuit of this degree. Thank you for the time and effort you have spent with on our boys. I could not have had a better support than what you have given me year in and year out.

Acknowledgments: This work was supported by my advisor Dr. Yiping Shao, who helped guide my direction, and teach me the intricacies of this work, Dr. Xishan Sun, whose electronics expertise was critical for the data acquisition, Dr. Allen Lan, who helped design the ASICs that were so critical in this work, Kai Lou who developed the reconstruction techniques used in this work, and Agile Technologies, who donated many of the crystals and arrays used in this study, and whose expertise in array construction was critical to the development of materials used in these studies. The research was partially supported by the award R21EB007581 from the National Institute of Biomedical Imaging and Bioengineering, and an institutional research grant from the University of Texas MD Anderson Cancer Center.

Abstract: High-resolution, small-bore PET systems suffer from a tradeoff between system sensitivity, and image quality degradation. In these systems long crystals allow mispositioning of the line of response due to parallax error and this mispositioning causes resolution blurring, but long crystals are necessary for high system sensitivity. One means to allow long crystals without introducing parallax errors is to determine the depth of interaction (DOI) of the gamma ray interaction within the detector module. While DOI has been investigated previously, newly available solid state photomultipliers (SSPMs) well-suited to PET applications and allow new modules for investigation. Depth of interaction in full modules is a relatively new field, and so even if high performance DOI capable modules were available, the appropriate means to characterize and calibrate the modules are not. This work presents an investigation of DOI capable arrays and techniques for characterizing and calibrating those modules. The methods introduced here accurately and reliably characterize energy, timing, and event interaction positioning. Additionally presented is a characterization of the spatial resolution of DOI capable modules and a measurement of DOI effects for different angles between detector modules. These arrays have been built into a prototype PET system that delivers better than 2.0 mm resolution with a single-sided-stopping-power in excess of 95% for 511 keV γ 's. The noise properties of SSPMs scale with the active area of the detector face, and so the best signal-to-noise ratio is possible with parallel readout of each SSPM photodetector pixel rather than multiplexing signals together. This work additionally investigates several algorithms for improving timing performance using timing information from multiple SSPM pixels when light is distributed among several photodetectors.

Table of Contents

Table of Figures	xii
Table of Tables	xvii
Abbreviations	xviii
Chapter 1 Introduction and Background	1
1.1 Positron Emission Tomography Basics	1
1.2 Scintillators Used in PET	2
1.2.1 BGO	2
1.2.2 L(Y)SO	3
1.2.3 Other scintillators.....	4
1.3 Depth of Interaction	4
1.3.1 Parallax Errors	4
1.3.2 DOI Detector Design	7
1.3.3 DOI Gradient Generation.....	10
1.4 Solid State Photomultipliers.....	13
1.4.1 Basics	13
1.4.2 Manufacturer comparisons	14
Chapter 2 Single Pixel Surface Studies.....	21
2.1 Introduction	21
2.2 Experimental Setup	21
2.3 Results	27
2.3.1 1.5x1.5x20mm 5.0 μ Finish	29

2.3.2	1.5x1.5x20mm 30 μ Finish	31
2.3.3	1.5x1.5x20mm Saw Cut	33
2.3.4	1.5x1.5x30mm Saw Cut	35
2.3.5	Summary of Results	37
2.4	Discussion.....	39
Chapter 3 Depth Of Interaction Response Function Calculation		43
3.1	Introduction	43
3.2	Method	45
3.2.1	Experimental Setup	45
3.2.2	Electronic Collimation DOI Calculation	46
3.2.3	Flood Source DOI Calculation	48
3.2.4	Data Collection and Analysis	51
3.3	Results	53
3.3.1	1.5x1.5x20mm ³ LSO scintillator with a 5.0 μ surface finish.....	53
3.3.2	1.5x1.5x20mm ³ LSO scintillator with a 30.0 μ surface finish.....	55
3.3.3	1.5x1.5x20mm ³ LSO scintillator with a saw cut surface finish.....	57
3.3.4	1.5x1.5x30mm ³ LSO scintillator with a saw cut surface finish.....	59
3.3.5	Individual Crystal Summary Results	61
3.3.6	Array Validation	63
3.4	Conclusion	63
Chapter 4 Depth of Interaction Array Design		65

4.1	Scintillator and Reflector.....	65
4.1.1	Initial Prototype	67
4.1.2	Second Prototype.....	71
4.1.3	Third Prototype.....	77
4.2	Light Guide.....	81
4.2.1	Light Guide for Prototype Array	83
4.2.2	Light Guide for High Resolution Array	84
4.3	Conclusions.....	86
Chapter 5 Electronics Setup.....		89
5.1	Parallel Readout.....	89
5.2	Front End Amplifiers	91
5.3	Custom ASIC	93
5.4	FPGA Processing.....	95
5.5	Conclusion	102
Chapter 6 Module Calibration and Performance.....		104
6.1	Introduction	104
6.2	Electronics Calibration.....	104
6.2.1	Electronic Time Walk Calibration.....	106
6.2.2	Electronic Charge Conversion Linearity.....	109
6.3	Full Detector Array Calibration.....	114
6.3.1	Crystal Segmentation	115

6.3.2	Crystal Energy Linearity	119
6.3.3	Time Walk Correction Calibration	123
6.3.4	DOI Response Function	131
6.4	Summary and Discussion	132
Chapter 7 Non-Timing Module Performance		136
7.1	Module Level Energy Resolution	136
7.2	Depth Effects on Crystal Flood Map	139
7.3	Crystal Identification	142
7.4	DOI Resolution and Function Comparison.....	145
7.5	Light Level Uniformity	150
7.6	Energy Resolution Uniformity	151
7.6.1	Depth Effects on Energy Resolution	153
7.7	Light Loss Linear Deviation	156
7.8	Background Counts Uniformity	157
7.9	Summary.....	159
Chapter 8 Timing Algorithms		161
8.1	Introduction	161
8.2	Coincidence Timing Method	163
8.2.1	Selected SSPM Pixel Timing Algorithms	163
8.2.2	Multi-Pixel Average Timing algorithms.....	165
8.2.3	Experimental Setup	168

8.3	Coincidence Timing Results	168
8.3.1	Crystal to Crystal Coincidence Timing	168
8.3.2	Module-to-Module Coincidence Timing	172
8.4	Event Level Timing Error Analysis Method	174
8.5	Event Level Timing Error Analysis Results	177
8.5.1	Crystal to Crystal Timing Error Results	177
8.5.2	Module to Module Timing Error Results	178
8.5.3	Full Module Expected Resolving Time	179
8.6	Summary and Discussion	183
Chapter 9 System Performance		186
9.1	Introduction	186
9.2	Procedure	186
9.2.1	Spatial Resolution	186
9.2.2	Point Source Measurements	189
9.2.3	Phantom	189
9.3	Results	190
9.3.1	Spatial Resolution	190
9.3.2	Point reconstruction & resolution	193
9.3.3	Phantom	197
9.4	Conclusions	197
Chapter 10 Summary and Discussion		199

10.1	Single Pixels	199
10.2	Detector Array Results	200
10.3	Calibration	202
10.3.1	Current Calibration	202
10.4	Future Work	206
10.4.1	Array Design	206
10.4.2	Computational Investigations	207
10.5	Alternate SSPMs	209
Chapter 11 Conclusions		211
11.1	Single Crystal Studies	211
11.2	Module Studies	211

Table of Figures

Figure 1 Parallax Error Effects	6
Figure 2 Concept of DOI designs	10
Figure 3 DOI Gradient Generation.....	12
Figure 4 Devices made by SensL.....	17
Figure 5 Hamamatsu SSPM Array	18
Figure 6 Electronic Collimation Test Setup.....	23
Figure 7 Geometric Beam Spread of Collimated Beam	25
Figure 8 Collimated Beam Profile	26
Figure 9 Energy Spectrum at center of 1.5x1.5x20mm crystal with 30 μ surface finish.....	28
Figure 10 R space distributions for 5x1.5x20mm crystal with 30 μ surface finish.....	28
Figure 11 Light Detection properties for 1.5x1.5x20 crystal with 5 μ finish.....	30
Figure 12 Resolution properties for 1.5x1.5x20 crystal with 5 μ finish.....	30
Figure 13 Light Detection properties for 1.5x1.5x20 crystal with 30 μ finish.....	32
Figure 14 Resolution properties for 1.5x1.5x20 crystal with 30 μ finish.....	32
Figure 15 Light Detection properties for 1.5x1.5x20 crystal with saw cut finish.....	34
Figure 16 Resolution properties for 1.5x1.5x20 crystal with saw cut finish.....	34
Figure 17 Light Detection properties for 1.5x1.5x30 crystal with saw cut finish.....	36
Figure 18 Resolution properties for 1.5x1.5x30 crystal with saw cut finish.....	36
Figure 19 Plot of light loss compared to the crosssectional area to length ratio.....	38
Figure 20 Count profile for 1.5x1.5x20mm crystal	47
Figure 21 R-space histograms for 1.5x1.5x20mm crystal with 5.0 μ finish.....	54
Figure 22 DOI response function comparison for 1.5x1.5x20mm ³ crystal with 5.0 μ finish....	55
Figure 23 Histogram of R-space distribution for 1.5x1.5x20mm ³ crystal with 5.0 μ finish.....	56
Figure 24 DOI response function comparison for 1.5x1.5x20mm ³ crystal with 5.0 μ finish....	57
Figure 25 R-space histograms for 1.5x1.5x20mm ³ crystal with 5.0 μ finish	58

Figure 26 DOI response function comparison for 1.5x1.5x20mm ³ crystal with 5.0μ finish....	59
Figure 27 R-space histograms for 1.5x1.5x30mm ³ crystal with 5.0μ finish	60
Figure 28 DOI response function comparison for 1.5x1.5x30mm ³ crystal with 5.0μ finish....	61
Figure 29 Flood Map of 9x9 LSO array with ESR glued to a rough surface	68
Figure 30 Global R-space histogram from 9x9 array with ESR glued to a rough surface	69
Figure 31 Conceptual drawing of photon scattering.....	70
Figure 32 Images of first generation prototype DOI array	71
Figure 33 Flood Map from 8 individual 1x8 arrays held together and separated by ESR	72
Figure 34 Global R-space histogram immediately after setup.....	73
Figure 35 Global R-space histogram two weeks after setup	74
Figure 36 Images of migration of optical grease into the array	75
Figure 37 Images of 8x8 array.....	76
Figure 38 Flood Map of 8x8 array Optical guide glued to each end	76
Figure 39 Global R-space spectrum from 8x8 array.	77
Figure 40 Flood Map 12x12 array with light guide	79
Figure 41 Global R-space histogram from 12x12 array with light guide	80
Figure 42 Single Pixel R-space histogram within 12x12 array with light guide	80
Figure 43 High Resolution array and light guide	81
Figure 44 Prototype Module with dedicated light guide.....	84
Figure 45 High Resolution Module with Dedicated light guide	86
Figure 46 Multiplex Schematic used for initial testing	89
Figure 47 Left-Signal Degradation due to RC coupling interference	90
Figure 48 Front end polarity inverting Amplifier Schematic	92
Figure 49 Signal before and after op-amp	92
Figure 50 ASIC Process concept.....	94
Figure 51 Picture of electronics	95

Figure 52 Electronics Standard Deviation of T1 as a function of TQ.....	97
Figure 53 Electronics Energy Linearity	98
Figure 54 Electronics Standard Deviation of TQ as a percentage of TQ.....	99
Figure 55 Timing Variations of various comparisons	101
Figure 56 Electronics TWC curve for 32 channels measured simultaneously	107
Figure 57 Average of 10 low, medium, and high energy pulses.....	108
Figure 58 Electronics Linearity overlain with counts spectrum.....	111
Figure 59 Flood Map without electronics correction.....	113
Figure 60 Flood Map for same data set with the electronics calibration applied.....	113
Figure 61 Flood Map with ^{137}Cs source	117
Figure 62 Flood Map with ^{22}Na source	117
Figure 63 Flood Map with Background events.....	118
Figure 64 Four single crystal results for energy spectra fit.....	122
Figure 65 Time Walk Correction, 8 iterations.....	129
Figure 66 Calculated TWC against measured TWC	130
Figure 67 Time Time Walk Error, 8 iterations	130
Figure 68 Four single crystal results for DOI response function.....	132
Figure 69 Energy Spectrum for Background Data Before Alignment	136
Figure 70 Energy Spectrum for Background Data After Alignment	137
Figure 71 Energy Spectrum for ^{22}Na Data After Alignment	138
Figure 72 Energy Spectrum for ^{137}Cs Data After Alignment.....	138
Figure 73 Flood Map at 5mm Depth of Interaction.....	140
Figure 74 Flood Map at 15mm Depth of Interaction.....	140
Figure 75 Flood Map at 25mm Depth of Interaction.....	141
Figure 76 Difference between Flood Maps at 5 and 25mm DOI	141
Figure 77 Crossectional Profile along x-axis. Peak-to-valley is ~8:1	143

Figure 78 Crossectional Profile along y-axis. Peak-to-valley is ~10:1	144
Figure 79 Crossectional Profile at 45° angle to x-axis.....	144
Figure 80 Derivative of count rate with source position. FWHM is 1.8mm	146
Figure 81 Comparison of DOI response function for intrinsic and external measurement..	147
Figure 82 DOI Resolution for three representative crystals within the array	148
Figure 83 Uniformity of Light Detection Efficiency for all crystals in a single Module	151
Figure 84 Energy Resolution Uniformity for all crystals in a single module	152
Figure 85 Comparison of Light Level to Energy Resolution for all 512 crystals.....	153
Figure 86 Energy Resolution for three representative crystals within the array	155
Figure 87 Background Counts Uniformity	158
Figure 88 Histogram of number of Pixels Triggered for good events.	161
Figure 89 Single Crystal Selected SSPM Pixel Timing Algorithms CRT Spectra	170
Figure 90 Single Crystal Multi-pixel average Timing Algorithms CRT	170
Figure 91 Full Module Selected SSPM Pixel Timing Algorithms CRT Spectra.....	172
Figure 92 Full Module Multi-pixel average Timing Algorithms CRT.....	173
Figure 93 Predicted and Measured Timing Resolution spectra.....	177
Figure 94 Predicted Timing Resolution Spectrum compared to measured CRT	179
Figure 95 Predicted Timing Resolution for three representative crystals within the array ..	182
Figure 96 Measurement for Spatial resolution	187
Figure 97 Conceptual drawing of Spatial Resolution measurement.....	188
Figure 98 Spatial Resolution for detectors at 180°.....	191
Figure 99 Spatial Resolution for detectors at 135°.....	191
Figure 100 Spatial Resolution for detectors at 90°.....	192
Figure 101 Spatial Resolution for rotated detectors.....	192
Figure 102 Axial resolution for various radial point source positions.....	193
Figure 103 Tangential resolution for various radial point source positions	194

Figure 104 Radial resolution for various radial point source positions	194
Figure 105 Axial resolution for various axial point source positions	195
Figure 106 Tangential resolution for various axial point source positions	196
Figure 107 Radial resolution for various axial point source positions.....	196
Figure 108 Derenzo Phantom image reconstruction with and without DOI	197

Table of Tables

Table 1 Comparison of Common Scintillators	4
Table 2 Crystal Properties as the Geometry and Surface Finish are Varied	39
Table 3 Differences between various measurements of DOI response function	62
Table 4 Determination of timing correlation between channels within the same ASIC, PCB board, and FPGA.	102
Table 5 DOI resolution of all crystals within an 8x8 module, averaged over 11 positions	149
Table 6 Maximum-Minimum DOI resolution at four depths within each crystal in an 8x8 array	149
Table 7 Summary of Depth effects on various crystals within a detector module.....	150
Table 8 Energy Resolution for each crystal within an 8x8 array, averaged among various different depths.	154
Table 9 Depth Effect on groups of crystals throughout module	155
Table 10 Maximum Light Loss Linear Deviation in keV for each of 64 crystals in a single detector module	156
Table 11 Comparison of Measured and Calculated Coincidence Resolving Time for several crystal pairs with DOI restricted data sets.....	178
Table 12 Expected Timing Resolution (ns) of all crystals in an 8x8 array	181
Table 13 Difference between maximum and minimum ERT (ns) among four mm segments of each crystal.	181
Table 14 Expected Resolving Time for groups of crystals at different depths within the detector module	182

Abbreviations

- ADC - Analog to Digital Converter
- APD - Avalanche Photodiode
- ASIC - Application Specific Integrated Circuit
- BGO - Bismuth Germinate
- CDF - Cumulative Density Function
- CFD - Constant Fraction Discriminator
- CRT - Coincidence Resolving Time
- DOI - Depth of Interaction
- ERT - Expected Resolving Time
- ESR - Enhanced Specular Reflector
- FBP - Filtered Back Projection
- FDG - Fludeoxyglucose
- FPGA - Field Programmable Gate Array
- FWHM - Full Width Half Maximum
- iCDF - Inverse Cumulative Density Function
- LOR - Line of Response
- LYSO - Cerium doped Lutetium Yttrium Orthosilicate
- Op-amp - Operational Amplifier
- PDF - Probability Density Function
- PET - Positron Emission Tomography
- PMT - Photomultiplier Tube
- PSAPD - Position Sensitive Avalanche Photodiode
- SSPM - Solid State Photomultiplier
- TDC - Time to Digital Converter

- TIMPIC - Timing Pick-off ASIC designed by our group for analyzing SSPM signals
- TAC - Time to Amplitude Converter
- TOF - Time of Flight

Chapter 1 Introduction and Background

1.1 Positron Emission Tomography Basics

Positron emission tomography (PET) is a nuclear medical imaging technique that takes advantage of the properties of certain radionuclides in order to image the distribution of radio-labeled biomarkers within a subject. Some low-Z, proton-rich isotopes decay in a process where a proton converts into a neutron and emits a positron. The positron then travels through the medium losing kinetic energy through Coulomb interactions, and once it has lost sufficient energy it pairs with an electron to form positronium. This exotic atom then decays into two annihilation photons each with 511 keV of energy, and an opening angle of 180° in the rest frame of the positronium atom(1). A PET scanner is a system built to detect both of the decay γ 's from the annihilation. While several devices have been studied as detectors for the annihilation photons (2, 3), the most widespread device used in PET systems is a scintillator, which is a material that transfers the kinetic energy of a single high energy particle to a large number of low energy scintillation photons, typically in the visible portion of the electromagnetic spectrum. The scintillator is coupled to a photodetector, which interacts with the scintillation photons and generates an electronic signal proportional to the number of detected photons. This electronic signal is then converted into a digital signal that may be stored and manipulated on computers. The data stored are the specific lines of response (LORs), or line segment in Time of Flight (TOF) capable systems, which is then used in backprojection or iterative reconstruction to reconstruct an image of the biodistribution of the radiotracer (4).

Several nuclides that decay via positron emission (Oxygen-15, Fluorine-18, Carbon-11, Gallium-68, Rubidium-82, and Nitrogen-13) are well suited for use as radiotracers in biomarkers (1, 5). A few examples of chemicals used in PET studies include FDG, which serves as a biomarker for measuring metabolism (5), ^{11}C -PIB, which serves as a biomarker for amyloid β for Alzheimer's studies (6), and ^{82}Rb , which serves as a biomarker for measuring myocardial perfusion (7). A recent study showed FDG based PET imaging is able to raise sensitivity and specificity of the detection of distant metastases in Breast Cancer from 85.9% and 67.3% to 97.4% and 91.2% respectively(8). Modern reconstruction techniques have improved the image reconstruction of PET scans (9), and TOF capabilities have further improved new clinical scanners (10, 11). However, the physical performance of many small bore systems, such as dedicated breast, brain, or animal scanners, is limited more by blurring due to parallax errors than by spatial blurring that may be corrected by incorporating time of flight (TOF) information.

1.2 Scintillators Used in PET

While a few studies have investigated the possibility of using detectors other than scintillation based detectors (3, 12-14), the vast majority of PET detectors use scintillation conversion. A large number of scintillators are available for various purposes, however modern PET scanners almost exclusively use BGO, GSO, or L(Y)SO (5).

1.2.1 BGO

BGO (Bismuth germanate) is a high-stopping-power scintillator used in many PET systems. This scintillator has a density of 7.13 g/cm^3 , two decay modes, a 60ns fast

decay mode and a 300ns slow decay mode, with 90% of the optical decay occurring in the slow mode (15), and an energy resolution of ~20% at 511 keV. The widespread adoption of BGO is due to its relatively low cost to manufacture and its excellent stopping power. However, low light output, relatively slow decay time, and associated poorer timing and energy resolution compared to other PET scintillators has led to a shift away from the use of BGO in high performance scanners in recent years.

1.2.2 L(Y)SO

Lutetium oxyorthosilicate (LSO) may incorporate yttrium to replace some of the lutetium atoms, and the resulting crystal is then abbreviated as LYSO. This scintillator is a doped scintillator, and is usually doped with cerium, although calcium has been investigated more recently (16). While lutetium based scintillators have a higher density and high stopping power than many scintillators, the stopping power is not as high as BGO. However, the light output, timing resolution, and energy resolution are much improved over BGO (40ns decay time for LSO, 53ns decay time for LYSO). This family of scintillators does not match the energy or timing resolution of some other possible scintillators such as Lanthium Bromide (LaBr), but the peak emission is well suited to available photodetectors, and unlike LaBr, LYSO is non-hygroscopic which simplifies the manufacture and processing of LYSO compared to other scintillators, and the higher effective Z in LSO reduces the amount of internal scatter compared to LaBr. Additionally, some recent measurements using LSO and SSPMs for detection have found coincidence resolving times in the range of 200-350ps (17, 18), which are well within the range that has previously justified the use of the more costly LaBr.

1.2.3 Other scintillators

Nal has been used in PET scanners historically, (1, 19, 20) , and several other scintillators have been investigated at one time or another for potential use in PET systems, including BaF₂, CeF₃, LaBr₃, and Csl (21-25). These scintillators have various tradeoffs among timing response, light output, cost, detection efficiency, and ease of production. A current research interest is LaBr₃, which has high light output, excellent timing and energy resolution, but is very expensive and highly hydroscopic. Calcium-doped LSO is another relatively novel scintillator that is similar to LSO, albeit with different timing properties. Due to its novelty Calcium doped LSO has not been studied for full detector modules at this time.

Table 1 Comparison of Common Scintillators

Scintillator	Light Yield (relative to Nal)	Energy Resolution (at 662keV)	Decay Time (ns)	Hydroscopic
LSO	70%	9.8%	40	No
BGO	15%	10.5%	60/300	No
Nal(Tl)	100%	5.6%	230	Yes
Csl(Tl)	130%	4.3%	1000	Slightly
LaBr ₃	160%	2.9%	30	Yes
GSO	31%	7.5%	60/600	No

Data compiled from(15, 26, 27).

1.3 Depth of Interaction

1.3.1 Parallax Errors

The reconstruction technique, either filtered back projection (FBP) or iterative reconstruction, relies on an accurate measurement of the line of response (LOR). The LOR is determined by measuring the interaction location of each γ , and calculating the line that connects the two interaction locations. Most systems incorporate pixilated

detector modules, which gives a discrete number of elements over the detector face. Although it is possible to build modules from monolithic LSO crystals, monolithic modules have severe limitations. Specifically, monolithic modules are very difficult to calibrate (28), have non-uniform resolution, and are generally limited to depths of 10mm or less(29). Due to these limitations, we have chosen to focus on pixilated detectors. Errors in placement of the LOR may arise from the size of the pixel, miscalculation of the pixel in which the interaction initially occurred from either statistical variation or a miscalculation of the interaction location due to Compton scatter within the detector, or parallax error due to an event occurring away from the central axis and at a depth within the detector module. The effect of parallax error on the LOR determination is conceptualized in Figure 1.




Figure 1 Parallax Error Effects. The red region contains all regions which may be traced to the same LOR in each system. **Top left - long crystals produce large parallax errors.** **Top right - reducing the crystal length reduces parallax errors, but at the cost of sensitivity.** **Bottom - Obtaining information about the depth of interaction reduces parallax error without loss of sensitivity**

Currently small bore systems reduce parallax error by reducing the crystal length. This has a limit in its application, due to the sensitivity loss in detector modules with short crystals. This effect is amplified in systems due to the requirement of detecting both γ rays from a positron-electron annihilation. An alternative method would be to continue to use long crystals, and design a large bore, but use only the central region of the bore for the study. Since the DOI effects increase farther from the central axis of a cylindrical scanner, this would also reduce parallax errors. However, the sensitivity for a system with an equivalent axial extent would decrease as the ring would cover a smaller solid angle, and so to achieve an equivalent system sensitivity the axial

coverage would need to be extended as well. The cost of a PET system scales with the amount of scintillator material used and the associated photodetectors and electronics necessary for the system, and so it is not practical to design a system with such an inefficient use of scintillator material. An ideal system would consist of long crystals to improve sensitivity and accurate DOI information to reduce parallax error, and the combination used to reduce the space between the crystal and the subject surface in order to reduce the overall cost of the system.

1.3.2 DOI Detector Design

Several methods have been studied, or proposed, for generating DOI information, including the following (illustrated in Figure 2):

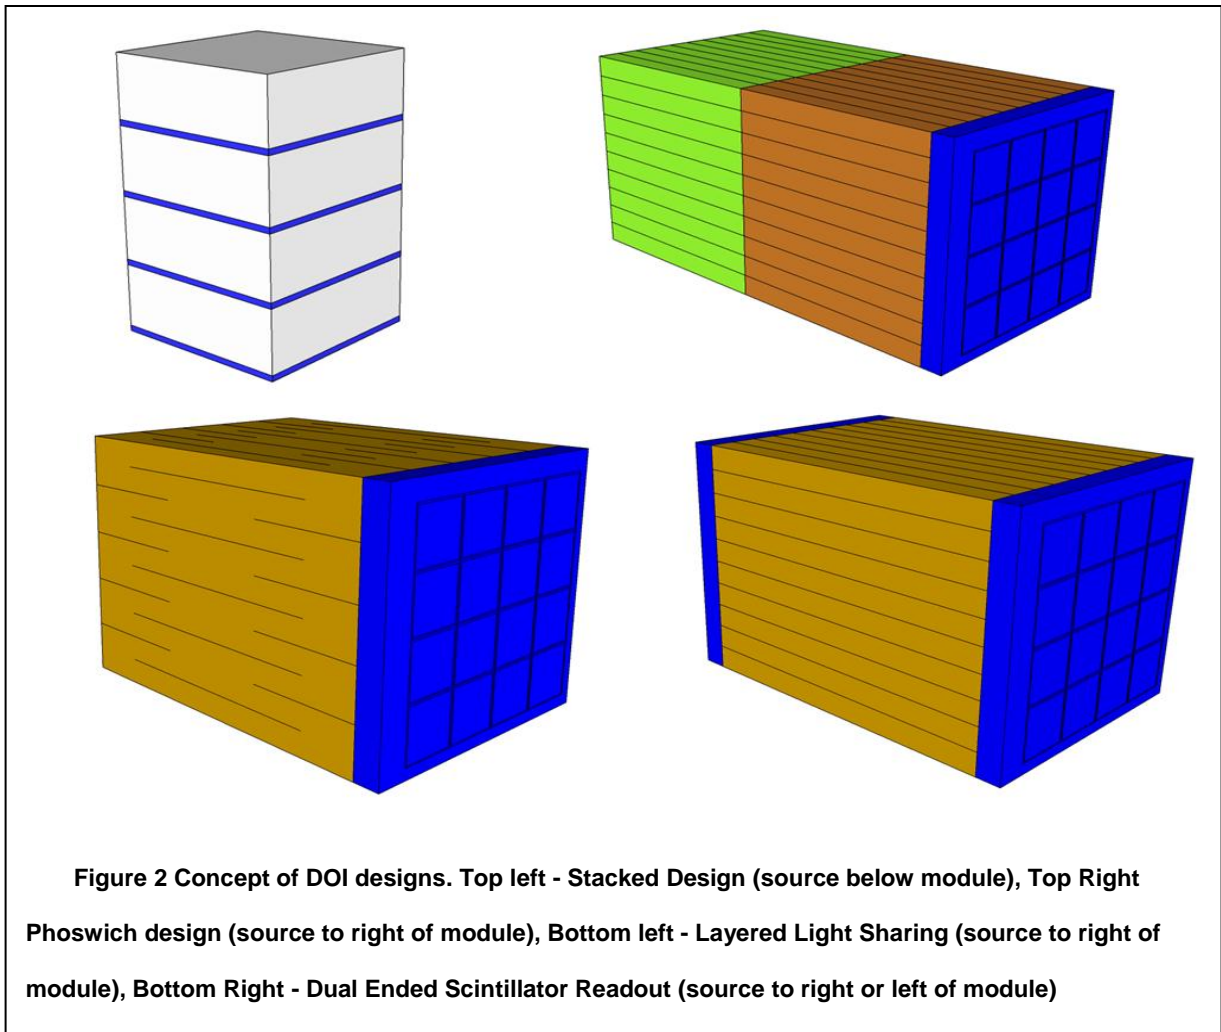
- a) Stacked detectors - In this design a scintillator is coupled directly to a thin photodetector, followed by another scintillator and another photodetector. The active photodetector layer determines both the originating layer, and the in plane position. This method requires the processing of a very large number of electronic channels and appropriate cabling for photodetectors within a stack, which will limit the geometric efficiency of any design based on this model(30-32). However, the scintillator material is frequently monolithic, reducing the processing cost for each layer.
- b) Phoswich detectors - In this design two scintillators with different decay times are stacked on top of one another and optically coupled to a photodetector. The signal decay properties determine which crystal layer the γ interacted in, while a separate calculation determines the x-y position. This method is limited in its resolution to the number of layers that may be

resolved using this method. While there has been work with modules using up to four layers, in practice most systems only involve two layers, and so a limited DOI resolution(33).

- c) Layered Light Sharing - In this design the light spread pattern depends on both the x-y interaction position and the original interaction layer. In these arrays, interactions from the same x and y positions, but different layers will map to a different location on a 2D flood map. This technique involves very complicated light sharing designs, and is limited by blurring between various layers in the 2D flood map (34). Additionally, the crystals must be separated sufficiently in the flood map to prevent overlap between crystals mapped to different layers.
- d) Dual Ended Scintillator Readout - In this design, photodetectors are placed on either end of a scintillator block, and the light from any interaction is shared between the two modules in a manner dependent on the depth. The depth is encoded in the ratio of light collected by the two ends. This approach involves more electronic channels than a single ended readout, but far fewer than a stacked detector design. Additionally this design has no detectors within the block, and so may be tightly packed, depending on the cabling method employed. The limiting factor in this method is the ability to determine the depth of the event from the ratio of the signal amplitudes (35, 36).

The design employed in this work is a Dual Ended Scintillator Readout (DESR) approach. This approach gives a continuous DOI profile, in contrast to the other most common approaches. Compared to a Stacked Detector design, DESR involves far

fewer electronics channels, and allows close packing geometry. Compared to layered detectors, both Phoswich and Layered Light Sharing, DESR avoids the need to optically align several different layers with high precision. The electronic requirements are simpler for DESR compared to a phoswich design because the DESR does not require a measurement of the decay time. DESR allows a detector with higher spatial resolution than a Layered Light Sharing design, because the DESR allows a design in which the pixel size is limited by the ability of the photodetector to resolve two adjacent crystals. In contrast, the Layered Light detector design requires the separation between two adjacent crystals to be large enough to map additional crystals from alternate layers. The reason a DESR approach has not been widely implemented is that DESR requires photodetectors to be placed on the subject side of the detector module. Traditional photomultiplier tubes (PMTs) cannot be placed on the subject side due to their very large depth profile. Some previous attempts had been made to use PIN photodiodes in DESR designs, but the performance of the photodiodes was not sufficient for PET systems. A new generation of photodetectors called Solid State Photomultipliers (SSPMs), or alternatively Position Sensitive Avalanche Photodiodes (PSAPDs) now allows for practical DESR modules (37-40).



1.3.3 DOI Gradient Generation

When light is generated in a scintillator, it is distributed isotropically, and if allowed to propagate undisturbed an equal amount of light will exit each end of the scintillator regardless of the interaction location. However, DESR requires a photon transmission gradient across the depth of the scintillator such that a different amount of light exits each end of the scintillator based on the depth of interaction. This DOI-dependent gradient may be generated either by absorbing light along the optical path, or diffusing

light at reflective boundaries within the DESR module or by a combination of the two methods.

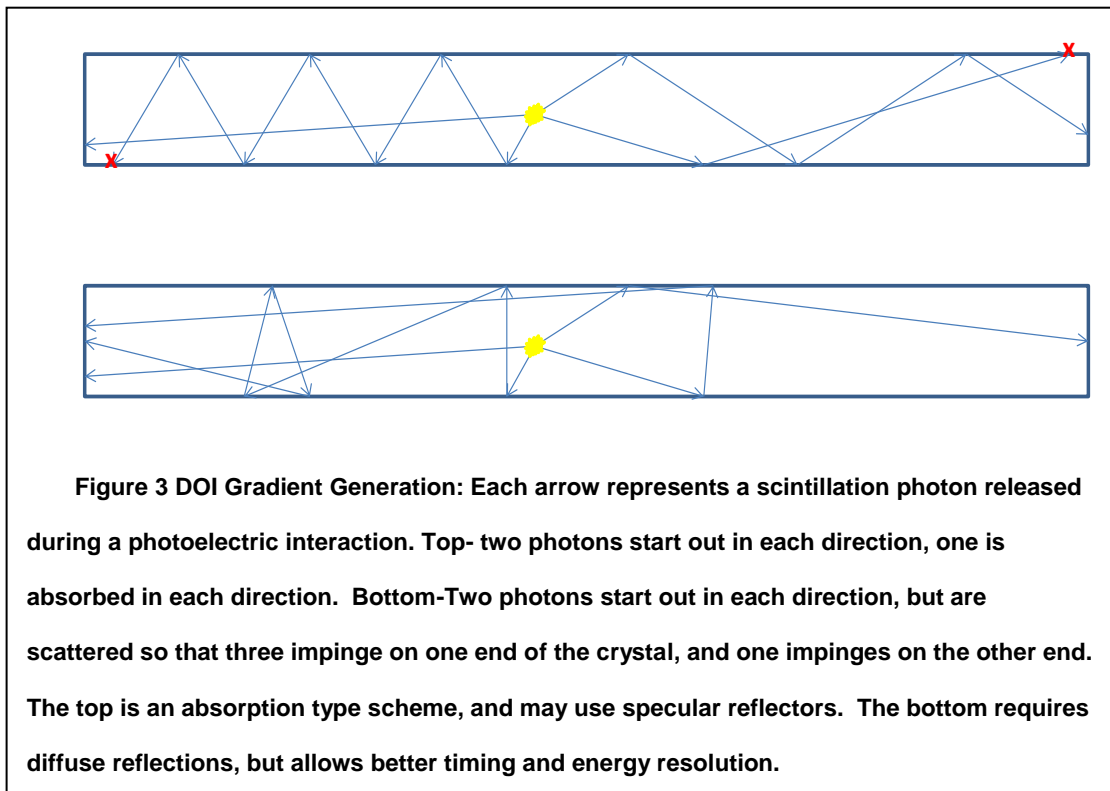
Light Absorption

A DOI gradient may be generated by absorbing some of the light propagated in each direction. If a certain fraction of light is absorbed over a given optical pathlength, the amount of light detected by each detector will depend on the average optical pathlength from the photodetector to the interaction location within the scintillator. In theory this may be accomplished either by introducing a reflector that absorbs some small fraction of light at each reflection (40, 41) or by using a scintillator that is not transparent over its own emission spectrum. However, scintillators appropriate for PET applications are transparent in their own optical emission spectrum, so internal absorption has not been investigated. The drawback of a light absorption approach is that the overall characteristics of a PET detector module tend to be highly dependent on the number of photons collected. That is, systems with a higher photon detection efficiency tend to have better resolution of both timing and energy. If a system absorbs light by design, that system may be expected to have poorer resolution of both timing and energy than a system that preserves those photons.

Optical Diffusion

The other method used to generate a DOI gradient is to cause optical diffusion within the scintillator module. Diffusion will cause a random walk type photon transmission, where the probability of exiting from one particular end depends not only on the initial direction of the photon, but also on the distance between the scintillation photon's origin, and each exit face. Ideally the diffusion would be within the crystal

itself, but in practice the diffusion typically takes place at crystal boundaries. The diffusion may be caused either by a diffuse reflector, such as Teflon, Torray, or Lumirror (42), or by a rough surface boundary and some fraction of total internal reflection at the boundary(35). In these cases a proper selection of reflector and surface treatment is critical, since these properties determine the strength of the DOI generating gradient, and subsequently the accuracy limits on DOI determination. This accuracy limit will then carry through to any other measurements that are dependent on accurate DOI determination.



Tapered Arrays

One additional method is to use a tapered crystal. Crystals of this geometry have a large end, and a small end, with a reflector along the sides. If the reflector is specular,

each reflection redirects the photon towards the large end by the slope of the taper. This leads to a crystal in which photons with identical velocity vectors, but originating at different locations will exit different ends of the scintillator. If the reflector is diffuse, the photon direction is randomized, but has a higher probability of being directed at the large exit face than towards the smaller exit face(38). This method could be combined with either diffuse or slightly absorptive reflectors in order to improve the overall DOI resolution. However, manufacture of these crystals is very labor intensive, and involves a low yield of completed arrays compared to initial material, and so would be prohibitively expensive for commercial systems.

1.4 Solid State Photomultipliers

Solid State Photomultipliers (SSPMs) alternatively referred to as Silicon Photomultipliers (SiPMs), Multi-Pixel Photon Counters (MPPC), or Geiger Mode Avalanche Photodiodes (GAPD), are relatively new photodetectors ideally suited to nuclear imaging techniques.

1.4.1 Basics

A single microcell in a SSPM is composed of an Avalanche Photodiode (APD) APD operated in Geiger mode. When any visible light interacts with the microcell, the microcell saturates, generating an electronic signal independent of either the energy or number of photons impinging on the surface (43). A SSPM pixel is formed from a group of SSPM microcells read out in parallel, forming a single analog signal proportional to the number of microcells that fire for any particular photon burst. The microcells vary in size, but are generally square with a side length of 25-100 μ , depending on the specific application. Smaller area microcells require a higher number of incident photons before

measurable saturation effects, in which two photons impinge on the same microcell, but give an identical signal amplitude to a single photon detection. However, each microcell requires an inactive insulating region around its perimeter, and so small microcell based SSPMs require a large proportion of the pixel area to be dedicated to insulation, thereby giving a lower overall geometric efficiency. While the number of microcells that fires is digital in nature, the SSPM in practice has a noise level that makes impractical individual photon counting unless the device is cooled significantly or if the device area is sufficiently small, generally on the order of 1mm^2 . One exception to this general case is the Digital SSPM (DiSSPM) designed by Phillips (Aachen, Germany) which records the microcell ID number of each microcell that fires (44).

A single SSPM pixel, with the exception of the DiSSPM and a position sensitive device under investigation by RMD (Watertown, MA), is not position sensitive. So, any position information must be generated by building an array of SSPM pixels, and then looking at comparisons among the signal amplitudes of pixels within the array. The techniques needed for these measurements are similar to light sharing techniques developed in traditional PET block detectors using Anger Logic (45).

1.4.2 Manufacturer comparisons

SSPMs are in the early phase of scaling for commercial applications. Many of the early issues have been resolved, and several commercial devices are now available. Due to the novelty of these devices, several generations have become available over the last four years, and not all of the same tests have been performed on every device. As such, it is difficult to perform a full comparison between devices. Additionally, the cost of devices should be considered in any comparison, but it is difficult to compare

costs since the costs have dropped by an order of magnitude in the last four years while the quality of devices produced in the same time period has significantly improved. Finally, many additional products have been produced by research groups and by producers with low volume. Although a comparison of all available devices would be interesting, this section will be limited to a brief discussion of some of the products made by each of the largest commercial manufacturers, rather than a detailed comparison of all devices.

SensL

SensL (Cork, Ireland) is a company that specializes in SSPM production. Their earliest array, SSPMArray was an array of 16 individual SSPM pixels, each with a physical size of $3 \times 3 \text{ mm}^2$ and with an active area of $2.85 \times 2.85 \text{ mm}^2$. This device was mounted in a large electronic housing, and while not scalable certainly gave the impression that future designs would be scalable. The next version (Array2) was based on the same technology, but the pixels were attached to a glass slide with flip chip technology, and the devices were tilable on three sides. The dead space between modules on the three tilable sides was equal to the dead space between pixels within the array, $\sim 0.5 \text{ mm}$. The next version was again based on the same basic technology, but was housed on a ceramic package with 0.5 mm dead space on three sides, and 1.5 mm dead space on the fourth side. The geometry of the new housing was a tradeoff that allows fully scalable 2D panels, but complicates the design of any high performance detectors due to the increased dead space around the edges of the detector. All these devices have similar intrinsic performance with rise times of 40-50ns, uniformity between pixels of $\sim 10\%$, dead space between active areas of pixels of 0.5 mm , and dark count rates of $\sim 8 \text{ MHz}$ per pixel.

The newest array produced by SensL is named ArraySL and has a geometry identical to Array4, but different intrinsic performance. The rise and peaking times of the new device are $\sim 10\text{ns}$ and 2ns respectively. The uniformity between pixels has been improved to better than 5%, and the dead space between pixels is $\sim 0.2\text{mm}$, although the dark count rate remains $\sim 8\text{MHz}$. SensL also groups these new arrays into larger panels with dedicated electronic systems. The vendor electronics determine the timing, channel number, and signal level for single pixel events within the module. However, the vendor electronics are not suitable for systems in which multiple neighboring pixels fire nearly simultaneously, such as in a system with light sharing or a system with dual ended scintillator readout. In those cases, custom electronics must be designed.

In the period from 2008 the price has dropped from \$2000 for the original arrays to $\sim \$100$ for the current 16 channel arrays.




Figure 4 Devices made by SensL. Top left and middle - original SSPM Array. Top Right - Array2. Bottom left - Array4 and Array SL. Bottom Right - larger panel made from 9 Array SL devices

Hamamatsu

Hamamatsu introduced a family of SSPM arrays for purchase in 2010, referred to as MPPC (multi-pixel photon counter). Their devices have an active area of 3x3 mm, with an interpixel gap of 1 mm in one direction, and 1.3 mm in the other, with a dark count rate of 2-3MHz, and a rise time of 15-20ns, caused by the intrinsic capacitance within each microcell and the overall quenching resistance for the full pixel(46). The overall performance of this family of SSPM is similar to that of SensL's ArraySL.




Figure 5 Hamamatsu SSPM Arrays Top Left - Monolithic Array Top Right - Discrete Element array Bottom - Large Area Discrete Element Array

Hamamatsu has recently developed a new monolithic array with an interpixel gap of $\sim 0.2\text{mm}$ which would give an identical form factor to those manufactured by SensL. As an improvement, the packaging of Hamamatsu's device does not include the same amount of dead space around the edges generated in the ceramic housing of

SensL's devices, but does include a large dead zone on one end for the cableing. The uniformity between pixels is within $\pm 5\%$ with a proper bias setting.

Additionally, Hamamatsu has developed some large area arrays built by assembling arrays of monolithic SSPMs, or from an assembly of a large group of individual discrete element pixels, similar to the devices manufactured by SensL. The properties of these arrays are identical to the smaller arrays formed of the same discrete elements.

Other Manufacturers

Other producers of commercially available devices include Philips (Aachen Germany), RMD (Watertown, MA), and AdvanSid (Povo, Italy). Philips device is unique in that it is able to determine which microcells fire in any given burst, allowing a flood

map of individual scintillation photons for any scintillation event. Also, their device includes a TDC which time encodes the first microcell fired in a burst, generating timing accuracy of ~ 100 ps (47, 48). However, Philips' device is orders of magnitude more expensive than other devices and is available only in limited quantities. While there may be large interest in this sort of information, the cost makes the device impractical for current system design. RMD's current work shows promise in delivering high quality, low cost devices. The current devices generate 1ns timing resolution for small devices and 9ns for larger devices. The spatial resolution of their detectors is below $100\mu\text{m}$, and their detectors are based on CMOS technology which has the potential to produce ultra-low cost components when manufactured in large quantities (49, 50). However, the position sensitive devices are not currently available in sufficient quantities to allow system developments, and are still undergoing development. AdvanSid's device performance is comparable to ArraySL, albeit with superior timing resolution (51, 52). However, the AdvanSid device is both more expensive, and also has a large geometric dead space due to individual ceramic packaging of pixels, rather than packaging of full arrays.

Choice For Investigation

This work investigates modules using SensL's family of devices as the photodetector. These devices were chosen due to availability and affordability, and have remained the detector of choice throughout the investigation based on availability, affordability, and compatibility with the developed components. Not all devices operate at the same bias voltage, or have the same physical parameters, and any new photodetector would need to be compatible with the existing system. The work

presented includes data collected on SSPM Array2, SSPM Array4, and Array-SL, as each new detector became available.

Chapter 2 Single Pixel Surface Studies

2.1 Introduction

One essential area of investigation before designing an expensive new detector module is the effect of the surface on the critical detector properties of the scintillator. For PET modules, the essential crystal level properties are: energy resolution, timing resolution, and DOI resolution. In order to guide the production of modules, we first studied the crystal response across the entire crystal length with varying crystal surface treatments. Single crystal measurements cannot test light sharing, Compton scatter, or light capture of escaping photons by neighboring crystals. However, single crystal measurements can give guidance to array designs by characterizing the expected properties of single crystal effects, such as light absorption (both bulk and surface effects), light transport within a crystal, and expected properties of scintillator/photodetector combinations on energy, timing, and DOI resolution.

2.2 Experimental Setup

Crystal surfaces may be finished either mechanically or through acid etching which produce similar overall surface properties (53). In this work the polishing was performed mechanically using lapping films purchased from PSI Dragon (Houston, TX) with nominal grades of 30 μ , 12 μ , 9 μ , and 5 μ . Additional samples were polished by the manufacturer to better than 0.5 μ using a wet lapping technique, where the abrasive particles are in a suspension rather than bonded to a film. A final group of crystals was left with a saw cut surface.

A mechanical approach was used in this study instead of a chemical bath technique because chemical etching produces a surface that depends both on the repeatable procedure (i.e. acid concentration, bath temperature, duration of etching, rinse procedure) and also the less repeatable initial surface finish, which will be affected by the cutting method and the blade condition, if cut with a saw or wire. Thus, different vendors, and the same vendor at different times may be expected to produce crystals with slightly different surfaces. In order to verify the starting point a careful Atomic Force Microscopy, or Scanning Electron Microscopy measurements of the crystal surface would be needed for every batch. However, mechanical finishing does not have this limitation, and the surface depends only on the lapping film properties, not the initial manufacturer produced finish.

After surface preparation, the individual crystals were then wrapped with Teflon tape (PTFE), a high efficiency Lambertian reflector (42), which is commonly used in PET applications (35, 53-55). Teflon tape is an ideal reflector for this experiment for several reasons. First, Teflon tape may be wrapped around the exterior of the crystal without being cut, which allows coverage of the corners. Any reflector which would need to be cut to shape would likely have small gaps at the corners. Second, Teflon tape need not be glued to the crystal, allowing an investigation of the surface effects without having to consider the interaction of optical coupling components. Third, a Lambertian reflector will allow an investigation of whether the crystal polishing has an effect. If the surface does not play a role, all surface finishes will have identical DOI performance, which would be determined by the diffusion properties of the reflector. This would indicate that DOI generation is possible due to surface scatter, but that a Lambertian reflector is critical in such a situation. If a specular reflector were used, and the surface finish does

not bear on the DOI performance, all surfaces would lack a DOI gradient. This would give a negative result, but would not clarify whether the lack of DOI information was due to surface finish or some other factor, such as crystal preparation. In this case a DOI gradient may be possible with a diffuse reflector, but that could not be determined from the test.

Each crystal was optically coupled to a single SSPM pixel in a SensL Array2 photodetector

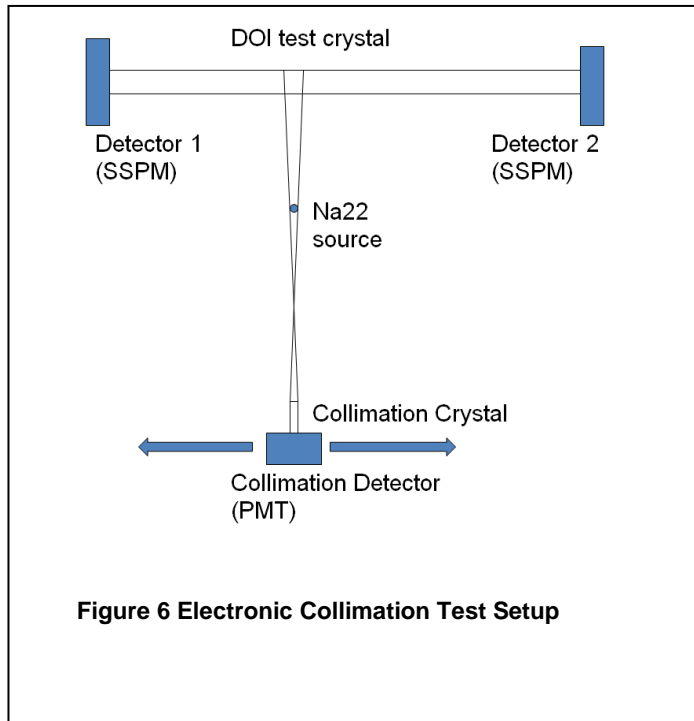


Figure 6 Electronic Collimation Test Setup


using BC-630 optical grease (Saint-Gobain, Hiram, OH). A single pixel readout was used to reduce electronic noise and to optimize timing, energy, and DOI measurements. A Na-22 point source with a 1mm cross-section was electronically collimated(35, 55, 56) to irradiate a narrow DOI region through coincidence with an additional LYSO crystal of 0.5x1.0x5.0mm³ optically coupled to a single channel PMT (Hamamatsu R-7400U), as shown in Figure 6. Each of the two SSPM signals, one from either side of the LYSO test crystal, was split and one output was summed and fed into a CFD (Canberra 454 quad unit) with a trigger threshold set just above the electronic noise for timing measurements. The signal from the collimating PMT was also split, with one output passed into another channel of the same CFD. The trigger signals from the CFD were then passed to a TAC (Canberra 2145) to convert the timing difference between the

sum SSPM and PMT signals. The other outputs of the split SSPM and PMT signals, were passed to a shaping amplifier (CAEN N568LC). The shaped signals and the analog signal from the TAC were then digitized with an ADC (Datel PCI-416) so that each SSPM signal could be analyzed separately. Labview (National Instruments, Austin, TX) was used then captured the 4 digital signals, which included a single TAC signal for coincidence timing and three energy signals from the shaping amplifier, one for each photodetector.

This experiment was designed to test the DOI resolution, energy resolution, and timing resolution for all crystal surface finishes at each depth within the crystals. The DOI ratio is defined as $R = S_1 / (S_1 + S_2)$ (35, 55, 57), where S_1 and S_2 are the signal amplitudes from the first and second SiPM detectors respectively. At any given depth the DOI ratio was calculated for the photopeak events above a 350 keV threshold. All events for this collimation position were then binned into a histogram for further analysis. Both the mean and FWHM were determined from the DOI histogram and used to determine the DOI resolution, which is defined as the physical distance needed between two collimation positions such that the FWHM of the DOI histograms do not overlap.


In order to accurately determine the true DOI resolution the geometric beam spread must be accounted for, so a measurement of the beam spread was performed as well. The beam spread is typically estimated from the setup geometry and estimated source distribution (35, 40, 55). However, we implemented a method of measuring the geometric beam spread directly for an electronically collimated beam similar to that developed by Burr(58). The geometric spread was measured by advancing the

collimated beam in steps of 0.1mm along the end of the test crystal, and the number of counts was measured at each position for a fixed time of 15 minutes per point, as shown in Figure 7.



The resulting count measurement is a convolution of the crystal detection efficiency (a step function for a rectangular crystal shape) with the geometric beam spread. A deconvolution of a step function with any other function is a simple derivative, so the difference between the number of counts between adjacent points was then plotted and taken to be the beam profile. A Gaussian function is then fit to the data, and the geometric beam spread is measured as 1.2mm. When calculating the final DOI resolution this beam width was quadratically subtracted in order to determine the true intrinsic DOI resolution. This method accounted for any errors associated with source

misalignment, where a slight error could lead to a large change in the beam spread or non-uniformities of the radioactivity distribution within the nominal source boundaries. The basic method of determining the number of counts at the edge of the detector could work for more complex crystal geometries (59), but a simple derivative would not appropriately deconvolve the beam geometry and the detection efficiency of a non-rectangular crystal.



In order to estimate the detection and gain difference for the two readout ends of the setup, the signal level for 511 keV events was measured for events originating at either end of the crystal. A straight line was fit to the 511 keV photopeak signal level response of the two end points on either side of the crystal. This line is referred to as the gain balance fit line, and distance between this fit line and the 511 keV signal level was

calculated for each point, and is referred to as Light Loss Linear Deviation. This measure is particularly useful for crystals in which there is a straight line difference between one end of the crystal and the other, but no other variation. In those cases a large absolute difference may be found in the 511keV peak value, but this may correspond to very little actual light loss, and be caused by a simple gain imbalance. An alternative measure of the light loss can be made by comparing the maximum signal level to the minimum signal level, however the usefulness of that method is limited for cases where there is both a detector imbalance between the two edges, and also a light loss gradient across the crystal, such as in the 1.5x1.5x20 mm³ crystal with a saw cut finish. Once the Light Loss Linear Deviation was determined, a correction factor was determined to remove the deviation, and the data were reprocessed before determining final characteristics.

2.3 Results

The crystals investigated included 1.5x1.5x20mm³ crystals with surfaces ranging from saw cut to highly polished, and 2x2x20mm³ crystals ranging from saw cut to highly polished. Unpolished crystals included 1.5x1.5x20mm³, 2x2x20mm³, 1.5x1.5x30mm³, 2x2x30mm³, 1.5x1.5x40mm³, 2x2x40mm³, and 4x4x40mm³. The 4x4x40mm³ crystal was not analyzed because the crystal area was much larger than the pixel active area, and the light collection efficiency was so low that the data were inappropriate for analysis. The 16mm³ crystal area coupled to a 8.1mm² SSPM active area allows a collection of less than 50% of the scintillation light. A sample energy spectrum from the 1.5x1.5x20mm³ 30μ surface finish crystal is shown in Figure 9, and the R space histogram distribution for the same crystal is shown in Figure 10. A detailed description of several crystals is given, followed by a summary from the full test set of crystals.




Figure 9 Energy Spectrum at center of 1.5x1.5x20mm crystal with 30 μ surface finish. Energy Resolution is 14.7%




Figure 10 R space distributions for 5x1.5x20mm crystal with 30 μ surface finish. Distributions are collected with a source step size of 1mm.

2.3.1 1.5x1.5x20mm 5.0 μ Finish

This crystal showed an excellent uniformity in the 511keV photopeak position, the counts detected at each location, the timing resolution, and the energy resolution. The energy resolution varied between 11.8% and 17.3% (15.5% if the edge is neglected). The timing resolution varied between 2.2 ns, and 2.6 ns (2.4 ns if the edge is neglected). The R-space coverage included only the range from 0.3 to 0.7, and the corresponding DOI resolution was poor at 3.5-4.5 mm (>20% of the overall crystal length). This indicates the combination of a polished surface and a Lambertian reflector is not ideal for DOI detectors, and that total internal reflection plays a significant role in this crystal. The DOI resolution would be expected to degrade further with the use of a specular reflector. Although the DOI resolution for this crystal is poor, it displays significantly better DOI resolution than the manufacturer polished crystals of the same size, which give a DOI resolution of 5.8mm. This indicates even a fairly smooth surface can produce some DOI information.




Figure 11 Light Detection properties for 1.5x1.5x20 crystal with 5 μ finish. Top-Normalized total counts, Bottom-Signal level of 511 keV photopeak




Figure 12 Resolution properties for 1.5x1.5x20 crystal with 5 μ finish. Top-Timing Resolution, Middle-Energy Resolution, Bottom-DOI resolution

2.3.2 1.5x1.5x20mm 30 μ Finish

In this crystal the R-values coverage included the range from 0.12 to 0.84, twice the space covered with the 5 μ polish, and the DOI response function is relatively steep, as shown in Chapter 3. There was no measurable light loss in this crystal, with the 511 keV signal level flat across the full crystal length, as shown in Figure 13. The combination of a steeper DOI response gradient and similar R-space resolution gave rise to a raw DOI resolution of better than 2mm, which improved further to 1.57mm after the beam geometry subtraction, show in Figure 14. However, this improvement in DOI resolution did not come at the expense of either light uniformity over the crystal, energy resolution (11.9% to 15.0% neglecting the edge), or timing resolution (2.0ns to 2.4ns), each of which is comparable to the 5.0 μ finish crystals. This indicates that DOI resolution can be improved without a necessary loss in other crystal performance metrics.




Figure 13 Light Detection properties for 1.5x1.5x20 crystal with 30 μ finish. Top-Normalized total counts, Bottom-Signal level of 511 keV photopeak




Figure 14 Resolution properties for 1.5x1.5x20 crystal with 30 μ finish. Top-Timing Resolution, Middle-Energy Resolution, Bottom-DOI resolution

2.3.3 1.5x1.5x20mm Saw Cut

This crystal displayed some interesting light loss phenomena. As the collimated source was advanced across the length of the crystal, the signal level corresponding to the 511keV photopeak shifted significantly. This shift included both a dip at around 8mm, and also a difference in the light collection for 511 keV events at the two ends of the detector module, shown in Figure 15. This combination of gain imbalance and light collection non-uniformities along the crystal length is an example of the cases in which the Light Loss Linear Deviation calculation becomes necessary. The R-space coverage was fairly extensive as shown in Chapter 3. Also noticeable in this crystal was a slight timing resolution degradation compared to both the 5 μ and 30 μ surface finishes. Notice that the depths where the light loss was most severe correspond to the worst timing resolution points (4-14mm). This indicates that the timing resolution of a crystal does depend on the light detected at the photodetector, as previously demonstrated by Moses and Ullisch (60).




Figure 15 Light Detection properties for 1.5x1.5x20 crystal with saw cut finish. Top-Normalized total counts, Bottom-Signal level of 511 keV photopeak




Figure 16 Resolution properties for 1.5x1.5x20 crystal with saw cut finish. Top-Timing Resolution, Middle-Energy Resolution, Bottom-DOI resolution

2.3.4 1.5x1.5x30mm Saw Cut

This crystal showed severe light loss effects with one third less light collected for 511 keV interactions occurring at the center of the crystal compared to similar events occurring at the edge of the crystal. This crystal showed significant energy resolution and timing resolution degradation as well. The best timing resolution was at the extreme edges of this crystal, and was as good as 2.2ns, similar to the edge of the 20mm crystals, but in the center the timing resolution was as high as 5.2ns. The best energy resolution in this crystal was 13.8%, but the response was as poor as 22% in the central region of the crystal. In this crystal it was easier to compare the light loss to the energy and timing degradation, and there is a direct correspondence between the points with the greatest light loss, and the points with the worst energy and timing resolution. The R-space coverage was excellent ranging from 0.05 to 0.95 shown in Chapter 3.




Figure 17 Light Detection properties for 1.5x1.5x30 crystal with saw cut finish. Top-Normalized total counts, Bottom-Signal level of 511 keV photopeak




Figure 18 Resolution properties for 1.5x1.5x30 crystal with saw cut finish. Top-Timing Resolution, Middle-Energy Resolution, Bottom-DOI resolution

2.3.5 Summary of Results

Both the energy and timing resolution were generally depth independent, unless the crystal demonstrated significant light loss. Additionally, neither the energy nor timing showed a strong dependence on the surface finish of the crystals, again unless the crystal generated significant light loss. This readout method was able to achieve better than 15% energy resolution (defined as the full width at half maximum of the photopeak divided by the photopeak signal level) and better than 2.3 ns timing resolution, both independent of surface roughness to the accuracy of this measurement. However, the DOI resolution heavily depended on the surface treatment of the crystal. The DOI resolution varied from better than 2 mm for the roughest surface to nearly 10mm for the polished crystals.

A comparison of the light loss characteristics of the saw cut crystals to the crossection/length ratio shows that the light loss was more severe when the crossection/length ratio decreased, as shown in Figure 19. That is, longer, thinner crystals displayed increasing light loss compared to the gain balance fit line. This indicates that the light loss is a function of the average pathlength of the optical photons. However, the complete remediation of the effect in the hand lapped crystals indicates that changing the surface properties can reduce, and nearly eliminate the effect, and so the light loss must be a surface phenomenon, rather than a bulk phenomenon. It should also be noted that among the hand polished crystals, increasing the surface roughness did not affect the light loss, indicating that the light loss characteristics are not a phenomenon of the mechanical properties of the crystal surface. This is shown by the tight grouping of all hand prepared crystals.




Figure 19 Plot of light loss compared to the crossectional area to length ratio for all crystals.

The lab treated crystals include surfaces from 5μ to 30μ . Manufacturer Polished crystals have a nominal roughness of 0.5μ . The saw cut were thoroughly cleaned with alcohol and wipes, but not treated other than leaning.

Table 2 Crystal Properties as the Geometry and Surface Finish are Varied

Crystal Properties		DOI Resolution (mm)			Energy Resolution (%)			Timing Resolution (ns)			Light Loss (%)	
Geometry (mm)	Surface Finish	Avg	Max	Min	Avg			Avg			Total	Linear Dev
1.5x1.5x20	rough	1.47	2.01	1.21	16	21	8	2.5	2.8	2.1	34.4	25.5
1.5x1.5x20	30 μ	1.57	2.15	1.36	14	16	12	2.2	2.4	2.0	3.2	2.5
1.5x1.5x20	12 μ	2.57	2.97	2.19	15	22	14	2.4	2.6	2.2	6.3	1.7
1.5x1.5x20	9 μ	2.20	2.75	1.87	14	15	13	2.3	2.4	2.1	7.5	0.5
1.5x1.5x20	5 μ	3.29	4.38	2.11	14	16	12	2.4	2.6	2.2	7.8	2.5
1.5x1.5x20	0.5 μ	5.81	7.79	2.41	15	18	14	2.4	2.7	2.2	14.0	11.2
2.0x2.0x20	Rough	1.59	2.07	1.25	14	18	13	2.3	2.5	2.2	6.5	2.3
2.0x2.0x20	30 μ	1.83	2.64	1.54	15	22	12	2.4	2.7	2.1	27.9	2.0
2.0x2.0x20	12 μ	2.78	4.02	2.36	14	15	13	2.2	2.4	2.1	6.6	0.5
2.0x2.0x20	9 μ	3.73	4.10	2.93	14	18	12	2.2	2.5	2.0	9.0	0.7
2.0x2.0x20	5 μ	5.85	7.30	3.85	15	16	14	2.4	2.7	2.2	6.5	4.2
2.0x2.0x20	0.5 μ	7.61	9.93	3.81	15	19	13	2.4	2.7	2.2	14.2	7.3
2.0x2.0x30	rough	2.78	3.85	2.34	17	19	14	2.6	2.8	2.2	26.7	19.3
2.0x2.0x40	rough	3.10	5.10	2.19	17	20	14	2.6	3.3	2.0	34.9	31.0
1.5x1.5x30	rough	2.49	3.84	1.86	20	39	11	3.2	4.7	2.2	42.9	39.6
1.5x1.5x40	rough	4.25	6.94	2.68	23	28	16	--	--	--	62.4	56.9

2.4 Discussion

While photopeak shifting for events at different depths will not generate a large difference in the calculated energy resolution at any one collimated position, it should be noted that an uncollimated, single-ended readout would have a significant degradation in the energy resolution. This would be the case since the overall measured energy spectrum would be an aggregate of the energy spectra at each depth in the crystal, and so the sum photopeak would be broadened by the amount of shifting along the crystal length. In a single ended readout this effect will be limited to differing amounts of light loss for events originating at different depths, but in a DESR technique an additional component from signal mismatch between the two ends may be added.

The difference at the ends could be caused by a mismatch of the crystal to the active area of the test pixel, a mismatch between the gain of the two photodetectors, or associated electronics, a gap in the reflector at the crystal detector interface, grease absorbing into the Teflon tape and reducing the reflectivity, light absorption along the length of the crystal, or other factors. These variations are setup dependent and difficult to reproduce. However, what is unlikely is that the amount of scintillation light changes significantly over the length of the crystal. Any complete characterization of the crystal should account for the signal shifting associated with the interaction location if possible.

The 30 micron surface finish performed nearly identically to the saw cut finish in terms of DOI resolution, although the 1.5mm crosssection crystals showed a minor loss of timing resolution for the saw cut crystals. The 30 μ finish has the best overall performance, when considering both timing and DOI resolution. This indicates that a lossless Lambertian reflector coupled to LSO with an air gap combined with a crystal with a saw cut surface is an appropriate setup for measuring DOI. However, this does not indicate that a Lambertian reflector is necessary for DOI generation. In the polished crystals the Lambertian reflector was unable to generate a DOI gradient, indicating that total internal reflection plays a significant role in the polished crystals. If total internal reflection at the rough surface can act to diffuse the propagated light, then a Lambertian reflector is unnecessary, and a specular reflector may be used. Due to the difficulty of seamlessly bonding ESR to single pixels, and the significant amount of light that would exit along the axial length of the crystal due to the 2% transmission rate of ESR, individual crystal measurements with ESR-wrapped crystals would not give a good indication of performance within an array, and therefore no single crystal measurements using ESR were made; ESR will be tested in an array structure only.

In previous studies Slates, Huber, and Strul all found that a short acid etching increased light output and/or energy resolution in individual crystals (61-63), with little improvement for longer etching times. For single-ended readout, an improvement in energy resolution and an improvement in light output uniformity are indistinguishable (in these cases a difference in light output uniformity will cause a blurring in the overall Energy Spectrum, which can be thought of as a summation of energy spectra at each individual depth), indicating that the improvement in energy resolution may be a result of improving light collection uniformity in Slate's work. It is the opinion of the author that the timing and energy resolution degradation typically attributed to surface roughness (61, 64-66) is not in fact a property of surface roughness, but of surface contamination of the scintillating crystal. It is possible that the contamination is physical with a residue lying on top of the crystal, however subsequent cleaning of the saw cut crystals did not mitigate the light absorption effects. Another possibility is that the contamination is chemical in nature, in which an interaction between the LSO crystal and the chemicals used to lubricate the machines during manufacturing has caused the material at the surface to be less than 100% transparent. If there is a surface contamination, a short acid bath could remove the residue or the defective crystal material without affecting the mechanical properties of the surface determining DOI effects, and similarly a rough grinding of the surface (30 μ in this work) would cause the same effect. The mechanically polished surface has far fewer crevices than the saw cut surfaces, allowing much less area for surface defects to remain and reducing the effect of any defects. If the defect is chemical in nature, it should be noted that the same material used to lubricate the saws during cutting is also used during polishing, and the same defects should be apparent in both cases, however the decrease in surface area would

reduce the effects of surface defects, since the surface area to volume ratio would be lower with a polished surface. This would explain not only why Huber, Strul, and Slates each observed that the short acid bath makes a large difference in light collection, while subsequent time has very little effect, but also why both Strul, Slates, and Pilcher each noticed an improvement in the samples etched for a short time compared to the mechanically polished samples(61, 63, 66), even though these cases would not significantly change the surface roughness, and did not noticeably change DOI resolution. This would also explain the difference typically attributed to rough surfaces, since the rough surfaces in surface treatment investigations are typically manufacturer treated, and not customer treated (53, 65, 66). One suggestion for further investigation is to determine a method to remove any surface defects from the crystals with a low cost, low labor method. Possible methods would include very short acid baths, acid spraying during manufacture, or investigation of different wetting agents during processing.

Chapter 3 Depth Of Interaction Response Function Calculation

3.1 Introduction

The electronic calibration method of the Depth of Interaction response function described in 2.2 is impractical for the characterization of a full system. The experimental setup requires access to the axial direction of the crystal, is time consuming, and is capable of characterizing only a single crystal using a pencil beam geometry, or a select few crystals using a fan beam geometry, at any one time. A full system with thousands of individual crystals will require the ability to determine the DOI response function for all crystals in a relatively short period of time and ideally will not require accurate mechanical positioning of equipment. One advantage of a layered detector design is the lack of a need for this sort of calibration. In those designs depth is determined directly with a knowledge of the physical boundaries between layers. However, a DESR design does not have easily definable physical locations in the measured data, except a measure of the end points of the crystal. Thus a DESR module must incorporate a calibration technique in which the DOI response function is determined for each point. Most investigations of DOI response functions in modules with continuously variable DOI response functions have centered on numerical methods including neural networks(28, 67-71), Monte Carlo simulations(72), and maximum likelihood calculations(73-75). These methods appear to be error prone, complicated, highly dependent on the learning algorithms applied, difficult to calibrate, and degrade significantly in modules longer than 10mm, precisely the modules in which DOI is most critical. Specifically, maximum likelihood and neural network approaches require highly collimated setups in which the entrance angle of the collimated beam is known, and Monte Carlo simulations are time consuming and do a poor job of simulating variances

in electronics over time and optical transport properties. These limitations are especially true when the exact nature of reflections is unknown, such as in a system with a saw cut surface that will vary between detector modules.

A method originally investigated by Shao (76) involved using a uniform external flood source to generate scintillation events over the length of the crystal. Another method investigated by Yang used the end points of an internally measured R-space histogram to determine the end points of a DOI response function, and fit a straight line to those curves (39). Yang's method has the advantage of calculating the DOI response function based on only internal scintillation events, which means the experimental setup is incredibly simple, and can be performed with the individual detector modules in any system configuration, thus removing the need for any external source, either uniform or collimated. However, his method assumes that the DOI response function is linear in nature, an assumption that is not the case for all crystal configurations, as demonstrated later in this chapter. Shao's method as validated in the original work has the drawback of requiring an experimental geometry in which the entire length of the scintillator is exposed to a uniform distribution of external radiation. This would be difficult to achieve in a PET system, since the axial length of the crystal is generally inaccessible, and perimeter pixels would shield the interior pixels from the external flood source. However, as suggested in the original article, the method should hold for any uniform flood source, whether an external flood source or from intrinsic decay events, and a further validation was required to demonstrate the accuracy of the method when using intrinsic radiation.

The method investigated in this builds on Shao's original work by validating the same procedure using intrinsic radiation events. Specifically, this work collects a measurement of internal scintillation events, and calculates a DOI response function by assuming that an inverse cumulative distribution function, a native function to Matlab, was equivalent to the DOI response of the crystals (57). Effectively this assumes scintillation event distribution from intrinsic decays is equivalent to the distribution from an external flood source. This method allows a PET system to calibrate the DOI response function for all crystals simultaneously using only intrinsic decay events, but makes no assumptions about the shape of the DOI response function. The only assumption of this method is that the distribution of radioactive materials is uniform throughout the volume of the scintillator, and that events originating at different locations have equal probabilities of being detected. This method preserves the simplicity of Yang's experimental setup, without making the assumption that the DOI response function is linear over the length of the crystal.

3.2 Method

3.2.1 Experimental Setup


The experimental setup is identical to that described in 2.2, with the exception that these experiments included two additional data sets. The first data set was taken with a 5mm disk source placed sufficiently far enough away (14cm) to mimic a uniform flood source. The second data set was a collection of internal scintillation events in order to compare uniform flood results to internal scintillation distribution measurements. Additionally we tested the DOI capable detector module described in 4.1.2, but with a step size of 2mm, and a collimation crystal size of $1.0 \times 1.0 \times 10 \text{ mm}^3$.

3.2.2 *Electronic Collimation DOI Calculation*

At each point the electronically collimated data may be considered a sample of the overall DOI response function. The width of the R-space value is the uncertainty of the measurement for a given collimated beam position. This uncertainty will include both the geometric beam uncertainty (described and measured in 2.2), and the intrinsic uncertainty of the DOI response function. The intrinsic uncertainty is a combination of statistical effects in the production and propagation of the conversion photons, Compton scatter effects, travel of electrons after energy transfer, photon detection efficiency of the SSPM, and signal-to-noise ratio in the SSPM due to dark counts, electronic noise, and noise within the front end electronics and digitization process. As the point source is advanced over the length of the test crystal, the error in DOI resolution is taken as the physical space traversed such that two measurements of R are separated at the FWHM of each distribution. The distinction that the two measurements are each separated at their own FWHM is necessary because the shape of the R-space distributions is different at various points in the crystal, and it is possible for two distributions to overlap below the half maximum of one distribution, and above the half maximum of the other distribution, see Figure 10 in which the edge points are noticeably non-symmetric. The data points are oversampled, and there is always a distribution close enough that the half maxima overlap, and so the physical point at which the two do not overlap is interpolated for all points. Once this physical space is determined, the intrinsic resolution is calculated by subtracting the geometric distribution width.

One critical item in any electronically collimated setup is the determination of the location of the edges of the crystal. Because the crystal changed between each measurement, and the source was removed and replaced, it cannot be assumed that

the edges remained in the exact same location from one measurement to the next. In order to determine the edge locations, the source is advanced an extra 5mm past the nominal edge of the crystal in each direction. A group of points corresponding to the physical length of the crystal were taken such that the count rate is a plateau. In practice, 20 points were taken for a 20mm crystal, corresponding approximately to physical locations from 0.5mm to 19.5mm. This method is sketched in Figure 20.



It should be noted that when the collimated beam profile begins to fall outside the physical area of the crystal, the DOI response function data from the electronically collimated beam are no longer an accurate measurement of the DOI response function.

The electronically collimated measurement is effectively a convolution of the crystal DOI response function and the distribution of the detected annihilation photons. In the central region of the test crystal the effective beam profile is constant as the collimated beam traverses the crystal. In this region it can be assumed that the weighted average of the distribution is centered at the center of the electronically collimated beam, and the center of the collimated beam position corresponds with the average location at which the DOI response function is being sampled. However, at the edges of the test crystal the effective beam profile truncates at the edge of the crystal, and the weighted average will no longer be centered at the center of the collimated beam, but shifted towards the center of the test crystal. In this experiment the beam spread was 1.2mm wide, and so this effect is only evident for the last point on either side of the crystal, but this indicates that differences between the intrinsic DOI response function, and the externally collimated beam DOI response function at the edge of the crystal may be caused by inaccuracies of the collimated beam method.

3.2.3 Flood Source DOI Calculation

After the collection of the electronically collimated data, a measurement of the intrinsic scintillation events was taken as well. LSO, and LYSO have a significant amount of Lutetium, which exists in a natural abundance of 2.59% as Lu196. This isotope has a half life of 3.7×10^{10} years, and has a decay scheme including gammas of 88, 202, and 307 keV (77), and also betas with a maximum probable energy of 595 keV (78), and an end point of 1193 keV(79). The DOI response function is calculated from the intrinsic flood data using an inverse cumulative distribution function calculation, whose end points are taken to be the edge of the crystal. This method assumes a uniform linear distribution of the radioactive sources, an assumption which may

generally be held for rectangular crystals, but not for more complicated designs (59, 80). In those designs the non-uniform crosssectional area of the crystal along the axial length of the crystal will cause a non-uniform distribution of the radiation events, while the volumetric distribution remains uniform. While a linear model does not apply to these more complicated crystal geometries, the method could be adapted and applied by renormalizing the calculated inverse cumulative distribution function weights or bin widths in order to incorporate the geometry into the model.

One potential failing point of the internal decay method arises when light loss is significant in the crystal. In those cases, low energy events from the areas with high light loss may not trigger the electronics, while events with the same energy at low light loss will generate sufficient signal to cause a trigger. In this case the distribution of detected events will not accurately match the distribution of generated events. For external flood sources the effect is simple to mitigate by applying a cutoff around the measured photopeak. Unless the light loss is so severe that counts are missed from the 511 keV photopeak, the software threshold will be appropriate to normalize the crystal. The background data do not have a photopeak at an appropriate energy, and so cannot be used to self-determine an appropriate energy threshold. Therefore, in order to mitigate this effect, an energy cut is applied to the data at 68% of the separately measured 511 keV photopeak signal. This cutoff corresponds to 350 keV if there is no low level offset, but may vary somewhat for locations with a significant offset. Ideally the linearity of the crystal should be mapped out with several data points for the energy. However, in this investigation the simple 68% energy cut was sufficient to mitigate any non-uniformity effects in all crystals that are appropriate for further investigation.

One noteworthy thought to consider about the experiment is that the source used for determining the energy cutoff need not be uniformly distributed across the crystal length. Rather, all that is needed is sufficient statistics to determine where the energy cutoff ought to lie for each depth. That is, in a practical scanner the source may be a line source at the center of rotation of the scanner, which would give more counts at the detector entrance than at the exit. Since all that is needed is a reference point for the energy cut, the energy histograms from different depths may have vastly different statistics as long as the energy histograms all have sufficient statistics to determine the appropriate cutoff. Thus although an external source may be necessary for some crystals in which light loss is noticeable, the source distribution over the crystal need not be uniform.

Some trends are expected from the detected R-value distribution histograms. Other things being equal, measurements covering a greater proportion of R-space correspond to a better DOI resolution, and so an ideal crystal will have an R-space distribution that is wide. However, if the R-space distribution extends to the full boundaries where $R=0$ and $R=1$, this will correspond to a signal of 0 on one or the other end of the detector. In these cases, a degradation in the DOI resolution may be expected due to insufficient statistics associated with the measurement. A tradeoff between these two is necessary, with a coverage from 0.2-0.8 generally offering the widest coverage without saturation effects.

Next, an ideal R-space distribution will have sharp edges at either side. The R-space distribution will depend on a number of factors, including the DOI response function, the DOI resolution, and the beam geometry. However, the shape of the edge

of the profile will depend on the R-space resolution of the detector system, with sharp edges indicating narrow R-space resolution, and gently sloped edges indicating a broad R-space resolution. The ratio of the slope at the edge of the R-space distribution compared to the slope of the DOI response function may provide some guidance on estimating the DOI resolution of crystals without an external source. However, this method was not rigorously investigated since it would only be able to provide an estimate at a single point (the edge) that may be limited by competing edge effects.

3.2.4 Data Collection and Analysis

Due to the electronics setup, it is possible for the signal level from one SSPM to be calculated as a negative, or zero; all such events were rejected. Once the internal DOI response function was calculated, the nominal location of the collimated source position in the electronically collimated data was allowed to shift by $\pm 0.5\text{mm}$ to the point where the DOI response function curves from the collimated and intrinsic curves aligned most closely. The electronically collimated data were taken with 1mm steps, and so it is assumed that the edges of the crystal are known only to $\pm 0.5\text{mm}$ accuracy. This shift was introduced to remove any global bias that may be introduced by a slight error in the determination of the true edge of the test crystal. For all points in R-space where the DOI response function may be interpolated for the electronically collimated data and a value exists for the internal flood data, the distance between the response function in physical space was calculated. A similar analysis was performed for the intrinsic data with the depth of interaction dependent energy cutoff applied for determining the need of an energy filter.

At a later date the crystals were cleaned, rewrapped, and coupled to the SSPM pixels again. At this point an intrinsic measurement was made once again, and a measurement with a 5mm disk source placed ~14cm from the test crystal. This setup approximated a uniform flood source distribution and was used to verify the intrinsic measurement matched results produced using a uniform external flood source as originally demonstrated by Shao. The external flood data were processed identically to the intrinsic data, and once again distributions were made both with and without energy cuts. This gave five separate determinations of the DOI response function, 1) Electronically Collimated, 2) External Flood Source, 3) External Flood Source with a depth dependent energy cut, 4) Intrinsic Flood Source, and 5) Intrinsic Flood Source with a depth dependent energy cut. Since the electronically collimated measurements were taken on a different date than the external flood source measurement, the electronics gain, optical coupling, and reflector wrapping characteristics may be different. These variations in the experimental setup may be expected to cause slight differences in the DOI response function during each measurement. Therefore, while we can compare internal flood to external step and shoot, and internal flood to external flood, we cannot compare external flood to external step and shoot. However, a validation of the method for an external flood source was verified in the original work(76). A simulated flood file was generated by combining all data from the individual step and shoot files and treated as an external flood source. However, any errors in the setup would carry through into both the electronically collimated treatment of the data and in the simulated flood source treatment. This limits the usefulness of any comparison between these two data treatment methods.

3.3 Results

3.3.1 $1.5 \times 1.5 \times 20 \text{ mm}^3$ LSO scintillator with a 5.0μ surface finish

This crystal showed a very poor R-space distribution. The space covers only 0.4 to 0.65. Notice also that the distribution looks relatively Gaussian in nature without the energy cuts, and only moderately sharpens when the energy cuts are applied. The R-space distributions were incredibly close for the background compared to the external flood source when no energy cut is applied, but a significant difference between the two methods became visible when the energy cut was applied. It is unclear what caused the side lobes on the external flood source measurement, but the effect is present in all crystal measurements.

The R-space distribution corresponded to a relatively flat DOI response function. However, the DOI response function measured from the externally collimated beam and the DOI response function from the intrinsic decay events are similar in shape. The difference in the calculated DOI position between the two methods for any given value of R was less than 1mm. This difference is dominated by differences within the last 2mm of each edge. The DOI resolution is 3.3 mm, which may not be sufficient for some ultrahigh resolution DOI applications. However, the error in R space is 0.052, indicating that the mean R-value measurement has little error and is primarily limited by the range of R space coverage rather than the accuracy of the R determination.






Figure 21 R-space histograms for 1.5x1.5x20mm crystal with 5.0 μ finish. Top-internal and external flood data without energy cuts. Bottom - same data with energy cuts




3.3.2 1.5x1.5x20mm³ LSO scintillator with a 30.0μ surface finish

In this crystal there was a much broader R-space distribution compared to the 5.0μ surface finish, covering 0.18 to 0.82. This distribution also had areas near the edge of the distribution that were higher than the central region, while the central region is relatively flat. This indicates a good overall DOI response function, which is indeed the case, with a measured DOI resolution of 1.57 mm. The overall difference between the internally calculated DOI response function and the externally collimated and external flood source measurements is once again below 1mm. This difference is smaller than either the DOI resolution of the crystal, or the beam spread, and indicates any difference in the determination of the DOI between the two methods should be a

negligible component of the overall system error. Here the average error in the R value was 0.066, quite similar to the value measured for the 5.0 μ crystal of the same geometry. This indicates that a change in the surface finish has only a minimal effect on the resolution in R-space, but has a very large effect on the width of the distribution in R-space. This result supports the previous assertion that an improvement in R-space coverage will generally lead to an improvement in DOI resolution.







3.3.3 1.5x1.5x20mm³ LSO scintillator with a saw cut surface finish

As discussed in 2.3.3, this crystal demonstrated some light loss. However, the R-space distribution was similar to that of the 30μ finish crystal with the same geometry, covering 0.1 to 0.9, and with similar edge peaks. This crystal showed an excellent DOI resolution at 1.47 mm, and the difference between the calculation based on internal scintillation events and external sources was again below 1mm. Once again, this difference was much smaller than the measured DOI resolution. As shown in Chapter 2, crystals with more severe light loss become inappropriate for PET applications as

performance of both energy and timing resolution degrade. The close match between the DOI response function when a moderate amount of light loss is present indicates that the proposed method is accurate for crystals that retain sufficient resolution for PET applications and is not limited to ideal crystals with no light loss.







3.3.4 1.5x1.5x30mm³ LSO scintillator with a saw cut surface finish

In this crystal there was relatively severe light loss, limiting the crystal's timing and energy resolution, as discussed in 2.3.4. The amount of light lost recorded in this crystal was likely near the limit of usability for PET applications. However, although the light loss limits the timing and energy resolution, the DOI is once again excellent at 2.47 mm, or 8.2% of the crystal length. Once again the R-space distribution was broad, in this case covering 0.05 to 0.92. The peaks are high at the edges, and the central region demonstrated some bowing. This bowing indicates a non-linear response of the crystal in the central region, which is apparent in the non-linear DOI response function shown in Figure 28. This indicates that Yang's method would be inappropriate for these

crystals, although the crystals may be useful for some PET systems. Note that despite the light loss, the various calculations are once again close together, giving a mean difference below 1mm between the internal and external determinations.





3.3.5 Individual Crystal Summary Results

A summary of the DOI response function calculations for the remaining crystals is given in Table 3. The difference between various methods used to determine the DOI response function was smaller than the DOI resolution in all crystals, including the 1.5x1.5x40mm³ crystal, which had light loss so severe that it is inappropriate for any PET system (see Figure 19). The difference between the methods is generally less than 1mm, with a maximum difference between the methods of 1.55mm, in a crystal with a DOI resolution of 5.9mm. This discrepancy is likely due to random coincidences in the collimated data, which would tend to weigh externally collimated data towards the center of the crystal, a trend visible in the DOI response function histograms for all

individual crystals. In general, applying an energy filter did not improve the agreement between the internal calculation and the electronically collimated measurement. Additionally, the differences between the intrinsic and external flood sources is exceedingly small, with a largest difference of 0.74mm, in a crystal with a DOI resolution of 4mm.

Although the differences in the various methods may be compared, it is impossible to determine which is more accurate. However, since the differences between the methods is exceedingly small compared to the intrinsic DOI resolution, the effect on the overall system resolution composed of any such crystals is expected to be a minor portion of the final error.

Crystal	EC-I(0)	EC-I(350)	F(350)-F(0)	F(350)-I(0)	F(350)-I(350)
1.5x1.5x20 0.5 μ	1.1	0.39	0.52	0.41	0.30
1.5x1.5x20 5.0 μ	0.57	0.89	0.16	0.42	0.12
1.5x1.5x20 9.0 μ	0.77	0.97	0.34	0.27	0.41
1.5x1.5x20 12 μ	0.56	0.63	0.25	0.40	0.27
1.5x1.5x20 30 μ	0.09	0.55	0.34	0.27	0.26
1.5x1.5x20 saw	0.07	0.59	0.36	0.39	0.17
2.0x2.0x20 0.5 μ	0.75	0.46	0.59	0.41	0.34
2.0x2.0x20 5.0 μ	1.55	1.14	0.58	0.30	0.39
2.0x2.0x20 9.0 μ	0.32	0.57	0.44	0.33	0.32
2.0x2.0x20 12 μ	0.78	1.04	0.40	0.29	0.31
2.0x2.0x20 30 μ	0.15	0.45	0.23	0.40	0.15
2.0x2.0x20 saw	0.33	0.51	0.28	0.36	0.18
1.5x1.5x30 saw	0.17	0.44	0.31	0.40	0.13
1.5x1.5x40 saw	0.35	0.45	0.99	0.55	0.74
2.0x2.0x30 saw	0.25	0.58	0.39	0.41	0.17
2.0x2.0x40 saw	0.81	1.24	0.47	0.43	0.40

Table 3 Differences between various measurements of DOI response function. The values reported are the means of the absolute value of the difference between measurement. EC is electronically collimated, and has a nominal 350 keV cutoff, F(350) is the external flood source with a 350 keV cutoff, F(0) is external flood source without any energy cutoff, I(350) is the internal distribution with a 350 keV cutoff, and I(0) is the internal distribution without any energy cutoff.

3.3.6 Array Validation

The method described above would be interesting, but would have limited utility if it were accurate only in single crystals, so a further validation was investigated in a DOI capable LYSO array. The method is as accurate within array settings as within single crystals. The details of the array investigation are discussed in detail in 6.3.4. At the system level 202 and 307 keV γ 's from other detectors must be considered, which will give a non-uniform distribution within the detector module, and so at the system level the energy cuts are necessary.

3.4 Conclusion

The entire concept of using intrinsic scintillation events was previously proposed by Shao (76). This work has demonstrated the accuracy of Shao's method for single crystal experiments. The method works well in all crystals tested regardless of either crystal geometry or surface finish. The difference between the DOI response function calculated from either externally collimated or intrinsically generated scintillation events was much less than the intrinsic error of the DOI determination. This held both in crystals that were appropriate for PET systems, systems with inadequate DOI resolution, and in crystals with such severe light loss that they are inappropriate for PET systems. Additionally, the method was validated for crystals within a DOI capable prototype array module. This method extends the mathematical model developed by Shao (76) to data collected using internal scintillation events. This method is fast, easy, accurate for non-linear DOI response function, requires no complicated setup, and uses widely available and well known mathematical tools. Energy cutoff information did not improve the results in single crystal measurements, but may prove critical in full

systems which may have non-uniformly distributed events originating from other crystals within the system.

Chapter 4 Depth of Interaction Array Design

4.1 Scintillator and Reflector

The individual data detailed in Chapter 2 and Chapter 3 indicate that a saw cut crystal of 2x2x30, or 1.5x1.5x20 would be appropriate for array investigations, while the amount of light lost in surface interactions for the 1.5x1.5x30mm crystals may limit their usefulness in an array setting. The scintillator used for DOI arrays is LYSO. While Teflon is ideal for single crystal studies, Enhanced Specular Reflector (ESR, manufactured by 3M) is a better reflector in array settings. The reflectivity of Teflon depends on the original manufacturer and the thickness of the material, with a thickness of 230 μ required to achieve 99% reflectivity. ESR achieves 98% reflectivity for the LYSO emission spectrum at a thickness of 65 μ (42), and unlike Teflon may be glued directly without impairing the reflectivity of the material(81). The Giant Birefringent Optics (82) phenomenon underlying the reflectivity of ESR gives rise to several beneficial qualities of the reflector, including that the light that is not reflected is transmitted unlike metallic reflectors such as Aluminum or Silver, or white reflectors such as Torray or Lumirror, which absorb a significant amount of light, and also that the reflectivity is independent of the angle of incidence in contrast to dielectric mirrors. Given the suitability of ESR for PET detector modules, ESR was adopted for use in PET detector modules since within a year of its original release (83).

The lossless property of the reflector is desirable in PET arrays, since any light escaping one crystal will be captured in an adjacent crystal and continue until it exits the module either at an edge of the array or at a photodetector. The preservation of photons is beneficial, due to the general degradation of timing resolution and energy

resolution of nuclear detectors as the number of detected photons is reduced. However, the transmissivity of ESR makes any attempt to model a full array by investigating single crystals impossible, since in single crystals any light escaping from the single crystal is lost altogether, while those photons will be preserved in a full array. In arrays with long thin pixels the number of reflections may be very high, and in these cases it is especially beneficial to have a lossless reflector. Additionally, ESR is relatively stiff when compared to other PET-appropriate reflector materials, which simplifies the manufacturing process for large arrays. However, the stiffness of ESR makes the material very difficult to bond to single crystals, since it cannot be easily wrapped around the sharp corners of single crystals. Therefore, although Teflon was used in the single pixel experiments, ESR is used in the array assemblies.

Other reflectors that have been used previously in PET modules include Toray Lumirror sheets(84, 85) and MgO, and TiO powders(4, 86, 87). The Toray Lumirror sheets are white sheets ranging from 50-240 μ . They have a reflectivity that depends on the thickness. In order to generate a 99% reflectivity a 240 μ sheet must be used(42), which takes a substantial portion of the volume of a high resolution array. A similar problem is encountered with powders, where channels are etched into the crystal and powder is packed into the notches(4, 88). Some very interesting new work involves generating microcracks within the LSO crystal array using subsurface laser etching(89-91) or laser etching air gaps between crystals (92). In this technique a laser is focused to a point within the material. The high energy density of the laser disrupts the local properties of the material. In some materials this disruption is optically diffuse and so can be used as a reflector. In contrast laser cutting introduces an air gap between crystals of $\sim 70\mu$ thickness, and total internal reflection guides the light along the path of

the crystal. This surface laser etching has several advantages, but generates smooth surfaces which do not scatter photons, and so this method is not appropriate for DOI capable detector modules. The sub-surface laser etching technique produces excellent results, and warrants further investigation, specifically into how the material within the cracked region performs as a scintillator, since the method necessarily disrupts the crystalline structure. Due to the specialized equipment required to perform sub-surface laser etching, the need for DOI, and the small pixel size in our arrays, ESR was chosen as the reflector in this work. Lumirror, while capable of producing better DOI resolution in an array, also degrade both spatial and energy resolution(38), and was not chosen.

4.1.1 Initial Prototype

The initial design called for a 9x9 array of LYSO crystals covering a footprint equal to the SensL Array4, 15.3x15.8mm. The detector is 20mm long and rough cut on all four sides. The flood map clearly shows 7x7 of the 9x9 crystals as distinguishable, as shown in Figure 29. Additional crystals may be present at the edges of this image, but the outside double row may be an image artifact instead. The perimeter crystals remained indistinct regardless of attempts to shift or rotate the module, light guides, and photodetectors.





Figure 29 Flood Map of 9x9 LSO array with ESR glued to a rough surface

Additionally, Figure 30 shows a lack of sufficient DOI information. The R-space distribution is narrow, covering only the space between 0.4 and 0.6; what is more the R-space distribution has gentle edges compared to the width, almost Gaussian in its distribution rather than an ideal step function. The overall response of the array is similar to that of the $1.5 \times 1.5 \times 20 \text{ mm}^3$ individual crystal with the 0.5μ surface finish, indicating an expected DOI resolution of $\sim 10 \text{ mm}$.



In the case of the single crystals, the reflector was Lambertian, but such a high number of photons were trapped in total internal reflection that the effective surface was mirror-like, as evident from the lack of DOI information in the polished crystals. It was hypothesized that the most likely cause of the complete loss of DOI information in this array was the glue used in the assembly process. In this case the glue (OP-20, index of refraction 1.456) between the ESR and the LSO (index of refraction 1.82 (93)) could act as an anti-reflective coating, increasing the critical angle for total internal reflection from 33° to 53° (94), and allowing the photons to reflect from the specular ESR without interacting with the diffuse crystal surface through total internal reflection. As shown below, this was indeed the case.


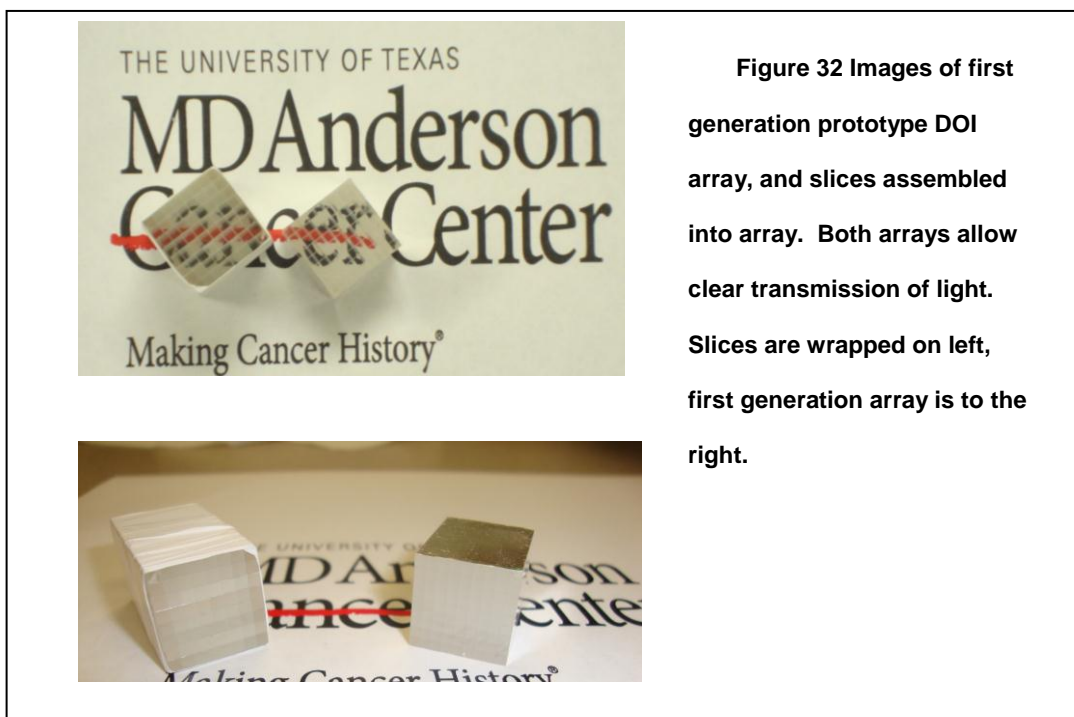


Figure 31 Conceptual drawing of several photons on a smooth surface (top) and a rough surface (bottom). Notice that the rough surface tends to generate more interactions with an incident angle near normal than in the case of a smooth surface.

This difference is especially critical when reflecting from a rough surface, since the geometry of the surface causes more shallow reflection angles than would generally be expected in a polished crystal, as demonstrated in Figure 31. Although the crystal separation and energy resolution was acceptable, the module is unacceptable without DOI information. As is clear in Figure 31, if the critical angle increases by 20° , the number of photons above the critical angle will decrease by a much greater amount than would be expected for a flat surface.

An additional area of interest is the complete lack of DOI information in this array, while an array designed by Yang using ESR glued to a rough surface generated a DOI resolution of $\sim 2\text{mm}$ (38). However, the distribution in R-space for this module is similar to the module evaluated by Yang et al, and so some of the differences in the final performance may have come from our group abandoning this module when it became

clear that the DOI resolution would not be sufficient for our purposes, rather than completing a full study of DOI resolution in a module that was clearly suboptimal.



4.1.2 Second Prototype

Individual Slices

The proposed solution to the loss of DOI information in the original prototype array was to introduce an air gap between the crystals in one direction. If the loss of DOI information was due to optical coupling by of the glue, then an air gap should allow total internal reflection at the saw cut surface boundaries and restore some of the DOI information. In order to test this we designed a 1-D array of 8 individual LYSO crystals glued together with ESR between crystals. The transition was made to 8x8 from 9x9 due to the inability to view all 9x9 crystals in the original prototype array. Eight individual 1D arrays were then stacked together with ESR placed between the layers

with no additional optical coupling, and the stack was then optically coupled to a SensL Array4 on each end using BC-630 optical grease. The DOI information was restored once again, shown in Figure 34, although only 6x6 of the 8x8 crystals were visible, indicating that the double peak at the edge of the 9x9 array was more likely an artifact than additional crystals. Various light-sharing schemes were attempted to allow a visualization of all crystals, but no simple scheme allowed the visibility of all the crystals. The final approach required a dedicated light guide as described in 4.2.






Figure 33 Flood Map from 8 individual 1x8 arrays held together and separated by ESR



Over the course of an experiment the optical grease (BC-630 index of refraction 1.465) tended to migrate between the layers of LSO and to fill in the small surface structures of the 1D arrays. As this happened, the DOI resolution of the arrays degraded, showing that indeed optical coupling between the LSO and ESR was degrading performance of the DOI capabilities of the detector. This was a particularly useful comparison, since the only difference over time was how much optical grease had wicked into the array, and the index of refraction of the grease was almost identical to that of the adhesive used in array assembly. Notice the contraction in coverage of R space between Figure 34 and Figure 35. These data sets were taken with the exact same setup, but with a 2 week delay between the initial setup and the acquisition of the

second data set. During this time grease wicked into the array as shown in Figure 36. The grease wicked to cover less than 1/3 of the face of each of the linear arrays, with a corresponding drop of 1/3 in the corresponding coverage of R space. This effect is similar to that seen by Huber and Moses in single crystal studies (95). In a glued setting the glue will cover the entire extent of the array, and the R-space coverage would degrade further.






Figure 36 Images of individual 1D arrays assembled into a block, and the migration of optical grease into the array

Assembled Array

The results from the individual slice investigation indicated that an array formed from 8 slices would be appropriate for our detector design. Once a working prototype array was designed, two sets of 8 detectors were ordered. Each detector consisted of an LYSO array glued on each end to a dedicated light guide, which is detailed in 4.2.1. This design allowed all 8x8 crystals to be distinguished clearly and good energy resolution for each individual crystal. The complete treatment of timing is complicated and will be presented in detail in Chapter 6 and Chapter 8. The DOI information is restored in the array, with the R-space histogram covering values from 0.3 to 0.7. This coverage is not ideal but it is sufficient to warrant a full investigation of this module. The DOI performance is detailed in detail in 7.4. So, while the crystal separation, timing, and energy were all excellent for the detector, the moderate DOI resolution justified an additional redesign of the array.




Figure 37 Images of 8x8 array assembled by Agile Technologies





Figure 38 Flood Map of 8x8 array assembled by Agile Technologies.
Optical guide glued to each end. Note all 64 crystals are clearly visible.



4.1.3 Third Prototype

As the number of diffuse reflections increases the DOI resolution should increase as well, since each reflection adds an extra step in the random walk to a crystal exit. The most recent design attempted to improve the DOI resolution by increasing the mean number of scattering reflections for scintillation photons. The final design investigated in this study used a 12x12 array of 30mm long LYSO crystals with air gaps on all four sides. The pitch is reduced to 1.25mm in each direction both to increase the in plane sampling and to reduce the crosssectional area to length ratio in order to increase the number of surface reflections. Additionally, an air gap is introduced on all four sides because if the number of diffusing sides is increased, the number of diffuse reflections will increase. This is expected to be crucial since many of the unscattered photons in

the second prototype likely traveled a path in which they reflected only from the glued surfaces. Including an air boundary in the second direction will eliminate any path that would allow this effect. Both of these factors are expected to increase DOI resolution due to the increase in the number of diffuse reflections. However, if the crystals have air gaps on all four sides, the overall module will be structurally weak. In order to improve the mechanical strength of the module a glass clamp was wrapped around the exterior of the array. While this formed a practical single module for investigation, it would introduce a large inactive area between modules in a full system, making such a system impractical. Thus, this module is for investigative purposes only, and a module designed for a full system would need a different means to build sufficient mechanical stability. Also, the cross-sectional area to length ratio is small, and for a saw cut crystal the expected light loss is expected to be fairly high, as shown in Figure 19

The crystals were then attached to a newly designed light guide (detailed in 4.2) and coupled to an Array-SL on each side. Notice in Figure 41 the excellent coverage in R-space. While the sides of the histogram have gentle slopes, this may either be due to poor resolution, or a difference in the R-space coverage for each of the 144 crystals. A histogram for a single sample crystal is shown in Figure 42, which indicates the R-space coverage and edge slope are excellent for individual crystals within the full array. All 144 crystals are visible, although the crystals in the top and bottom rows have low counts. This is largely due to light loss on those crystals. While the air coupling is now symmetric on all four sides of the module, the reflector remains continuous in one direction, and broken in the other direction. This allows extra light sharing in one direction and an overall asymmetric flood map. This asymmetry is apparent in the disappearance of crystals in one direction only. When the energy spectra of the edge

crystals are separated, the energy resolution is indeterminate due to a very high effective lower level discriminator, with events of >300 keV being rejected. The poor performance of these crystals may be mitigated with the redesign of the light guide, or by assembling the modules into an array and allowing light to pass from one detector module into an adjacent module for detection, rather than complete loss as is the case in the current module.






Figure 40 Flood Map 12x12 array with light guide, assembled by Agile Technologies.

10x12 crystals are clearly visible, with poorer separation for the top and bottom rows.

Nuclear spectra for the top and bottom rows shows significant light loss.



**Figure 41 Global R-space histogram from 12x12 array with light guide,
assembled by Agile Technologies**



**Figure 42 Single Pixel R-space histogram within 12x12 array with light
guide. assembled by Agile Technologies**




Figure 43 High Resolution array and light guide. Top Left - Array before Assembly, Top Right - Light Guide before final grinding and Assembly, Bottom Left - Assembled Array, Bottom Right - Final Light Guide

4.2 *Light Guide*

The large inactive areas around the edges of the SensL SSPM array caused a severe degradation in the properties of the perimeter crystals in the arrays without light guides. These crystals were physically coupled to inactive area, which limited the detection efficiency for the perimeter crystals. The problem could be mitigated slightly by shifting the detector, at which point an additional row and column would become visible (a change from 6x6 to 7x7 for the 8x8 slices). It might be supposed that a shift for one detector in one direction, and a shift for the other detector in the other direction would completely correct this problem, but that cannot work in a dual ended detector.

When the crystal is shifted, the non-visible section of the array is moved from an area with ~10% geometric coverage to an area with 0% coverage. In those cases the signal from one end of the detector is zero and the crystals are only read from one end. While the pixels may become visible within a flood map, the zero light detection on one end eliminates the ability to calculate the DOI. Additionally, the edge remains on a low efficiency detector and the other end is in a 0% efficiency region, and so the energy and timing resolution are severely degraded. For detectors in which the scintillator area is larger than the detector area, a means to transport light from the edge of the scintillation array to the active area of the photodetector is required.

Some previous works have investigated a simple light spreading technique where light is spread through a thin glass plate or other monolithic device(96-98). The SNR in an SSPM is intrinsically lower than that of a PMT, and so the relative threshold in terms of keV must be set higher. In our work, we found that in order to transport enough optical photons to the active area of the detector in these modules so that the trigger may be set above the thermal electron signal, the light spreader needed to be 5mm thick. When the light spreader passed ~3mm in thickness, the central pixels blurred together in the flood map, making a simple light spreader impossible to use for good performance in both the central and perimeter regions. Several attempts were made to build a light guide with a central light spreading region and a tapered edge which directed light away from the inactive area of the SSPM. After several such light guide designs failed, it was decided that the best approach for this investigation would be a tapered, segmented light guide, similar to those developed in other works(98-100) in order to direct the light sharing in a specific, controlled manner.

4.2.1 Light Guide for Prototype Array

The light guide designed for this work is based on simple optical sharing and is formed from trapezoidal elements that both share light between modules and direct light away from the inactive and toward the active areas of the SSPM array. The trapezoidal pixels are 1-1 LSO-Light guide coupled for the perimeter crystals and 3-2 LSO-Light Guide coupled elsewhere. Overall, this design has 36 trapezoidal light guide elements for the 64 LSO crystals. Each trapezoidal element has unique angles for all four sides, and the design of the overall light guide takes into account the asymmetric geometry of the SensL arrays.

The model used when designing the light guide assumed perfect reflectors at all points and no redirection of the light at reflective boundaries. While this is not the case, it was assumed that the difference between a light guide modeling these properties, and a light guide with the most rudimentary assumptions would be negligible. Optical modeling of complicated systems tends to have large uncertainties, but the excellent separation of crystals in the flood maps demonstrates that the assumptions were sufficient for this work. The basic diagram of the detector module outlining the light direction in the light guide is shown in Figure 44.




Figure 44 Prototype Module with dedicated light guide. Edge and central crystals are 1-1-1 LSO-Light Guide-SSPM coupled, while the remaining crystals have directed light sharing within the light guide.

4.2.2 Light Guide for High Resolution Array

When the third prototype module was designed, that design included a new light guide. This light guide was based on the same optical assumptions as the original light guide. If the optical separation were perfect, a light guide based on this design would merge the fourth and fifth rows, the sixth and seventh rows, the fourth and fifth columns, and the sixth and seventh columns. This is because both the third and fourth crystals from one end are one-to-one directed onto the same SSPM pixel. An alternative design would remove the central reflector in the light guide and allow merging between the sixth and seventh rows and columns. ESR has ~2% transmittance for each reflection, and so there should be some ability to distinguish the potentially merged rows and columns, as demonstrated by more than 8 SSPM pixels firing for most scintillation events in the 8x8 array (see Figure 88). This is the case even though no crystal has light directed onto more than 4 SSPM pixels on each side, and the optical guidance would indicate a maximum of 8 pixels triggering for any event. Optical simulations are

prone to error, and it was unclear before array testing whether or not the optical separation between adjacent crystals would be sufficient for either design. A tradeoff was made, and the central reflector in the light guide was perforated to allow ~30% transmission, thereby directing light sharing across the two central SSPM pixels. The resulting flood map allows clear separation among all crystals, with an extra wide gap between the sixth and seventh rows and even more so for the sixth and seventh columns. Clearly the separation between the central rows and columns is exaggerated, and a future design could keep the light guide the same, but include an additional column and row in the scintillation array to increase the number of crystals to 13x13. A significant concern with this change, however, would be that the crystals along the top and bottom rows already lose a considerable amount of light, and reducing the cross-sectional area to length ratio may cause these pixels to become unreadable. Any further reduction in crystal cross-section would need to be accompanied by either a reduction in the SSPM and electronics noise or a reduction in the crystal length.




Figure 45 High Resolution Module with Dedicated light guide. This design incorporates both directed light sharing in the light guide, which controls sharing, and also uses the inherent optical crosstalk through the ESR in the LSO array.

After further consideration, it was noticed that if the edge crystal were not 1-1 coupled to a single element in the light guide, the light guide could be redesigned in a way that would eliminate any expected merging. However, that redesign would cause less light from the perimeter crystal to be transported to interior SiPM pixels and might cause separation at the edge of the module to degrade. Due to the long lead times and costs involved in redesigning light guides, a redesign was not investigated. However, incorporation of subsurface laser etching in the manufacturing of light guides may allow for a rapid comparison of several different possible light guides in the future.

4.3 Conclusions

Dual-Ended readout based modules with saw-cut LSO scintillation crystals, ESR reflectors, and SensL ArraySL photodetectors performed well for DOI measurements.

However, the ESR cannot be glued on all four sides of rectangular LSO crystals without an unacceptable loss of DOI resolution. This loss is due to glue serving as an optical coupling material, as shown by the effect of optical grease wicking into the 1D Array slices. Adding an air gap, or other optical diffusion method is essential to preserve DOI information. An alternative method of generating a DOI gradient is to employ a tapered array or to use a diffuse reflector. In either case, the gradient is produced without concern for surface roughness. However, the simplicity and adaptability of the current module design would recommend this design over a tapered array.

Both decreasing the cross-section-to-length ratio and increasing the number of sides with the combination of a rough surface and an air gap improves the overall DOI response of crystals within an array. However, this should be tempered by the investigation which showed that untreated saw-cut crystals have significant light loss for small cross-sectional area to length ratios, and so sufficiently long, thin crystals may show the same effect unless properly treated. Also, careful consideration must be given to light escape at the edge of the detectors from a slightly transmissive reflector such as ESR. The light loss, along with the external glass clamp, precludes the high resolution 12x12 array from being implemented in its current design in a full detector system, even though the energy resolution is comparable to the 8x8 array for many of the interior crystals, and the spatial and DOI resolution are both far superior.

In order to approach more nearly ideal modules, some investigation is needed in how to manufacture saw-cut crystals without photon absorption at crystal boundaries and how best to treat the boundaries of the completed module in order to reduce light escape. Surface improvement is possible as shown both previously in this work (Figure

19), and in separate acid etching experiments (61-63, 65, 66). Subsurface laser etching deserves significant future investigation, especially in light guide design. Improved manufacturing techniques would allow the development of high performance, low cost detector modules that could be adapted for more applications. The techniques developed for the light guide should be suitable for light guides using other reflectors, including subsurface laser etching. Surface level laser engraving may also allow new light guide designs which improve light collection efficiency by directing the light away from inactive areas within the photodetector array.


Chapter 5 Electronics Setup

Custom front end electronics circuit designs and circuit boards were developed by other lab members, and a custom Application Specific Integrated Circuit (ASIC) was developed in collaboration with Tsinghua University, China. The ASIC miniaturized a large portion of the electronics used to process the detector signals into a $2 \times 2 \text{ mm}^2$ silicon chip to reduce size, power consumption, and cost of an overall system. The basics of the system will be described here, while more detailed information is given in other articles (101-103).


5.1 Parallel Readout

Traditional PMT readout methods in PET systems use resistor multiplexing circuits to reduce the number of electronics channels for processing(87, 104-107). Having fewer electronics channels simplifies computation, lowers the overall cost, and may increase the maximum count rate of a system. In this system parallel readout is used to

reduce errors, especially timing errors. Dark counts due to thermal electronics limit the SNR of the SSPM pixels, and the dark count rate scales directly with detector area. This indicates that the SNR would be limited in a multiplexed network due to the addition of dark counts on SSPM pixels collecting no scintillation photons. In order to measure the



degradation, we analyzed the coincidence timing resolution between two LSO crystals irradiated by a ^{22}Na source. One LSO crystal was coupled to a fast PMT, and the other crystal was coupled to a single pixel within an SSPM array. Additional SSPM channels were tied to the signal channel, but with LSO coupled only to the initial SSPM pixel. This test was performed to determine the effect of adding noise from channels without changing the original signal. The additional channels were directly tied together using a solder bond. The coincidence timing was then measured as a function of the number of included SSPM pixels and found to degrade as the number of pixels with contributing noise increases, as shown in Figure 47.



**Figure 47 Left-Signal Degradation due to RC coupling interference,
Green is without multiplexing, Blue is with the multiplex network.
Right - Timing resolution as a function of connected SSPM Pixels**

Additionally, we compared the basic signal shape for a single channel SSPM pulse with one passed through a resistor network designed to minimize pincushion artifacts, Figure 46. Pincushioning is reduced in this network by ensuring through simulation that even steps are produced in the calculated x and y positions with movements of 1 SSPM pixel. Notice that the shape of the pulse in Figure 47 is much smoother for the signal measured without the resistive network. Each SSPM has an internal resistance and capacitance, and when these RC characteristics are coupled to a resistive network, oscillatory effects emerge as this network combines four coupled RC circuits for each SSPM pixel. Additionally, when the pixels are all tied into the network, they tend to interact with each other, further complicating the network. The effects increase as the number of SSPM pixels connected to the network increases, eventually degrading the SSPM output signal beyond what is useful for PET. For these reasons we chose to develop a parallel readout technique whereby each SSPM pixel signal is analyzed simultaneously.

5.2 Front End Amplifiers

The custom ASIC was designed to directly process the signal generated by SiPM Array2 from SensL. When redesigning the photodetector for the housing used in Array4, SensL changed the polarity of the readout signals from negative to positive. To accommodate the newer sensor array, polarity flipping op-amps became necessary in our system and were introduced with the configuration shown in Figure 48. The signal shape was similar before and after the op-amp, as shown in Figure 49, for an average of 10 pulses from before and after the op-amps on different channels after gain and polarity changes. The reduced noise in the pulse after the amplifier is caused by the

650MHz bandwidth of the op-amp which has the effect of smoothing variations of under




Figure 48 Front end polarity inverting Amplifier Schematic

2ns; in this system such fast signals are primarily electronic noise.





Figure 49 Signal before and after op-amp. Results are average of 10 measurements. Data are taken from two different SiPM pixels triggering on the same scintillation event, and results are scaled to the same amplitude. Y-axis is arbitrary scale.

One major drawback to the inclusion of the op-amp board to the system is that the system now has a shaping time limited by the 650 Mhz bandwidth of the op-amp. Components were chosen that optimized the performance of the Array4 device. However, Array SL has a much faster rise time, and the timing resolution should be proportional to the slope to noise ratio of the leading edge. The intrinsic slope of the ArraySL device is higher, but the slope of the signal processed by the ASIC is limited by the bandwidth of the op-amp, and so the timing resolution may now be limited due to the polarity inversion, rather than the intrinsic properties of the detector. This effect may become even more crucial as SensL is developing new SSPM devices capable of better than 300ps timing resolution. A new ASIC has been designed that allows a negative input and removes the need for the polarity inverting amplifiers, but the new ASIC is still under investigation.

5.3 Custom ASIC

Most ASICs currently under investigation for PET system acquisition split the signal into two paths, one for fast timing, and the other with a slow shaping time and a peak sensing Analog to Digital Converter (ADC) (108-114). We developed an ASIC designed as described in (103) that converts the initial analog signals from the SSPM into two TTL level signals, one corresponding to the time above a set threshold (TT) and the other with a width proportional to the total charge in the SSPM pulse (TQ). The leading edge of the time over threshold signal is referred to as T1 and is a leading edge trigger signal. The falling edge is T1' and is not used at this time. The second TTL signal is produced by charging a capacitor with the SSPM signal as long as the time over threshold signal is active, and then discharging the capacitor through a constant current source. The second pulse is active for the entire time of the discharge. The leading

edge is T2, the falling edge is T3, and the total time to discharge is $TQ = T3 - T2$, TQ is a measure of the charge produced in the original signal. A basic diagram of the circuit function is shown in Figure 50. Notice that TT is active while the signal is above a threshold, and TQ is active while the capacitor discharges. While T1' is not used at this time, if there is a desire to adapt this chip for a phoswitch type detector, the comparison of TT to TQ should allow a differentiation of the crystal layer based on the decay time.



In comparison to standard techniques, this method has a lower power consumption, lower noise, better linearity, and a higher potential count rate. In this chip the electronic dead time is the sum of the hold time, which is a constant, and the charge time and the discharge time, which both depend on the signal amplitude. In a light-sharing environment most channels have a signal that is much lower than the maximum signal

range, and so a dead time correlated to the signal amplitude is beneficial. A standard ADC has a fixed dead time based on the shaping time, which may be on the order of microseconds, and is generally much longer than the discharge time in this setup. The current setup typically resets in less than $2\mu\text{s}$ from the leading edge to the full conversion of TQ and has a maximum electronics dead time of $5\mu\text{s}$. Additionally, the energy measure is proportional to the total charge rather than the peak amplitude. This is particularly important in SSPMs, since the pulse shape changes with amplitude. This is discussed in more detail in 6.2.1

Full module electronics have been assembled from the polarity flipping amplifier circuitry, the custom ASIC, and a commercially available Field Programmable Gate Array (FPGA) board to digitize the TTL time stamps. A picture of this system is shown in Figure 51.





Figure 51 Picture of electronics


5.4 *FPGA Processing*

The digital signals are then passed to an FPGA for processing. The FPGA used in this work is the Altera Cyclone IV, with a clock rate of 40MHz. The FPGA is


programmed to include 8 phase-shifted clocks, which gives a nominal bin width of 0.3125 ns for the T1 measurement. The intrinsic resolution of the electronics portion of the system was measured by first sending a model pulse (described in 6.2) from a pulse generator, Agilent 22330A 20 MHz Arbitrary Function Generator, into a 32-channel custom built Fan-In-Fan-Out board, and then into the electronics. A trigger was also passed directly to the FPGA board for reference. The time differences between each channel compared to the reference trigger were measured and compared the others for correlation. The intrinsic electronics error ranged from 22 bins at the lowest amplitude to 1.8 bins at the highest amplitude, and 2-3 bins on average, or approximately 800ps. The large error for the low signals was due to the broad flat nature of the electronic model pulse at its peak, rather than an inherent limitation in the ASIC for low amplitude pulses. When the system triggers in the flat region of the peak, any small variation in the electronic signal will correspond to a relatively large time difference in the trigger. In the sharp leading edge of the signal, a small fluctuation of the electronics will lead to a small time difference.



The energy linearity was measured using the same data set. Using the same FIFO pulse, the error of the electronics in determining energy may be measured. Signal amplitudes varying over two orders of magnitude were passed to the electronics system, with 5000 individual pulses at each amplitude. The system response was linear over values of TQ ranging from the lower threshold value to 2000 bins, and this range covered the useful data from nuclear signals. The electronics energy response varied by as much as 35% between channels, although this included differences among the front end amplifiers, the ASIC channels, and the channels in the FIFO board. From these data it is impossible to determine whether the differences are due to variations within the ASIC or among the amplifiers. Future versions of the electronics will eliminate this uncertainty.



The variance between data for one channel at one energy may be taken as an upper estimate of the electronics error in determining energy. This is an upper estimate, because the error will include both the error within the custom electronics and the error within the pulser. However, while the two sources cannot be clearly distinguished, the overall error is typically below 1%, except at the extreme lower limits of the measurement. The error of the physical measurement may be estimated as the energy resolution of the 511 keV photopeak signal, which is never below 10%. In this case, the 1% error from the electronics can be considered small and neglected, regardless of whether the majority of that 1% is from the pulser or the system electronics.




An additional area of interest was the correlation between channels. In this system several channels may fire simultaneously, and so it is critical to know if the 2-3 bin error between channels on a board is correlated. If the jitter is correlated, then the error of an average of several channels is equal to the full 800ps, while if the jitter is uncorrelated, the total error drops as the standard error of the mean. We measured this by comparing the T1 measured from various channels to each other and found no correlation between channels. Each module consists of one FPGA board, 2 ASIC boards, and 2 ASIC chips on each board. Comparisons were made between channels on the same ASIC, between channels on the same ASIC board, but processed on different ASICs, and channels processed on different ASIC boards. The time stamp for each channel was recorded as the time between the reference trigger and T1.

The mean standard deviation of the difference in the T1 time measurement between channels on the same ASIC was 2.0 ± 0.3 bins, between channels on the same board, but different ASICs was 2.1 ± 0.3 bins, and for channels on the same FPGA but different ASIC boards was 2.3 ± 0.3 bins. These small differences indicate that there is no correlation between time stamps on the same board or between those processed with the same ASIC compared to those events on different boards or ASICs. However, this measure does not account for correlation that may be present within the FPGA. All triggers are generated in the same manner within the FPGA, and so if there is any correlation within the FPGA that may be hidden in the measure. In order to determine whether the trigger has a larger error than is measured between the reference trigger and the individual events, a comparison was made between sequential events. The differences were compared both for the reference trigger and for individual channels. Since the reference trigger has a rise time of 5ns, but the pulse has a rise time of 40ns, we expect the variance for the reference trigger between subsequent events to be much smaller than for the electronics channels. We found the standard deviation of the difference between the reference trigger between two events to be 1.2 bins, with the average standard deviation of the difference for each channel between events measured as 2.3 ± 0.2 bins.

This demonstrates that there is no measurable correlation between the various measurements of the timing stamp. Any correlation within an ASIC board, or within an ASIC would cause a decrease in the deviation between channels within the board or ASIC compared to cross board or cross ASIC measurements. The comparison between events separated in time demonstrated that the expected error in the reference trigger was on the order of 1.2 bins, which is smaller than the observed error within the

measurement, demonstrating there is little or no correlation between the reference trigger time stamp and the model pulse triggers. Since this is the case, we can assume the errors of each channel is independent of the errors of other channels. In our system, 8-15 channels commonly fire for each nuclear event, and so the standard error of the mean from a pure electronics measurement may be below 100ps if all channels are weighed equally. However, this ignores noise in the SSPM which is larger than the electronics noise as demonstrated in Chapter 6. This indicates that the electronics used in this work are sufficient for analyzing the performance of our system and may be applicable to TOF systems as well with minimal modifications.



Comparison Type	Standard Deviation of Timing Difference
Channel to Trigger	0.83 bins
2 Channels, same ASIC	1.08 bins
2 Channels, same ASIC board, different ASIC	1.15 bins
2 Channels, same FPGA, different ASIC boards	1.31 bins

Table 4 Determination of timing correlation between channels within the same ASIC, PCB board, and FPGA.

5.5 Conclusion

The custom electronics and ASIC setup worked very well in this system. The polarity flipping op-amp altered the shape minimally, as demonstrated with direct measurements of the waveform. The electronics energy resolution, at better than 1%, far surpasses what is necessary in this system. The timing resolution is sufficient at 800ps for each channel, with a potential for improvement with a redesign or removal of the front end amplifier or redesign of the FPGA programming. The electronics jitter is uncorrelated between channels, indicating that the expected final error of an average of several channels should be less than the individual error of any single channel, and may be sufficient for time-of-flight measurements. This could lead to a timing resolution better than the bin width of the FPGA timing bins, although such a result has not been realized with the current system. The bandwidth of the polarity-inverting op-amp may currently be the limitation for timing with the new family of SensL devices, but that will remain as an unknown until the new ASIC chips have been fully tested and implemented in the system. The power consumption is low enough to allow a system where no additional cooling is needed for the system. Overall, the electronics are modular, scalable, and practical for a full system. Improvements are needed to remove the polarity inverting op-amp, and to improve the timing resolution for new, high

accuracy SSPMs that are capable of better than 300ps timing resolution. The ASICs have separately been validated to 50 ps accuracy, and so the largest improvement is needed for the FPGA programming architecture, or the integration of a high-accuracy TDC into the ASIC.

Chapter 6 Module Calibration and Performance

6.1 Introduction

A fundamental issue for any PET system is the ability to calibrate individual elements of the system. Since DOI capable systems are new, there has been little published about calibrating full detector modules, and what has been published involves linear estimates of the DOI response function(39, 84, 115, 116). Ideally the system should be calibrated to yield the best possible energy, timing, positioning, and DOI data possible for an event originating anywhere within any module using a non-linear method(57, 76, 117, 118). For systems with a large number of crystals, it is critical to automate the calibration method in order to allow the system to recalibrate at regular intervals without a large overhead time commitment to the calibration. The method developed here describes the procedure developed to calibrate the prototype modules involving only a limited amount of user interaction. The described method calibrates electronics level signal uniformity, electronics pixel level timing corrections, crystal identification, crystal level energy linearity and offset, and the crystal-level DOI response function.

6.2 Electronics Calibration

Before the crystals can be calibrated, the electronics must first be calibrated. Any gain difference between electronic channels will generate a degradation in the overall energy resolution of the module, due to the position dependence of the light distribution pattern onto each of the SSPM pixels, which are then coupled to electronics channels with differing electronic gains. This can be pictured by considering the overall energy resolution of the 2x2x20mm crystal with the 30 μ finish. In that crystal, although the

intrinsic energy resolution was 15%, the gain imbalance between the electronics on either side would cause a 27% blurring and give a final resolution of 31% over the full crystal. The solution to this problem can lie either in developing an energy linearity map for each depth of interaction or through numerically balancing the electronic gain response with either a linear or quadratic function(117, 119). The modules developed in this work show negligible light loss over the length of the crystal, and so either method would work equally well, and non-linear corrections are not needed. However, the method developed here has the goal of usability in a full system in which there is limited access to the crystals and sources need to be placed within the bore. In these geometries the number of total impinging gamma rays needed to generate sufficient statistics for 511 keV photopeak analysis at the exit window of the 30mm crystal would be prohibitively high. Therefore, an electronic gain balance is preferred. Once the electronics are calibrated, any light loss effects within the crystal may be corrected separately if necessary.


In order to calibrate the electronics a test pulse is passed into a 32 channel custom fan-in-fan-out board, and then to each of the 32 electronics channels for a module. This pulse is based on the average of 100 separate signal pulses measured directly from the post-SSPM polarity inverting pre-amplifier. The signal amplitude of the test pulse was varied over 69 steps ranging over two orders of magnitude and covering the electronics from below the signal threshold to well into the saturation region. A total of 5000 pulses at each amplitude were recorded. A line was fit to the linear region of a plot of TQ against nominal pulser voltage, and the slope of the line was taken as the electronic gain of the specified electronics channel. This gain measure included effects of the fan-in-fan-out board, the polarity inverting preamplifier, the ASIC chip, and the FPGA

conversion. The inclusion of the fan-in-fan-out board does add some error to this measure, since the fan-in-fan-out board is not in the final system. However, the measured variance between channels on the fan-in-fan-out board of 3% is smaller than the reported variance in the SSPM pixel gain of 5%, which cannot be included in this measure. Future versions of the electronics will eliminate the need for a model pulse or a fan-in-fan-out board.

6.2.1 *Electronic Time Walk Calibration*


The ASIC developed for this work uses a leading edge trigger, which is known to have a time walk error that varies with the signal amplitude (103, 120, 121). A Time Walk Correction (TWC) curve is a correction of the recorded trigger time T_1 to a corrected detection time, T_1^* , which removes the effects of time walk. This curve can be stored either as a parametric curve or a lookup table. One method of determining the TWC curve is to measure the time between a known trigger and an appropriately shaped pulse with a constant time between the trigger and the pulse and then varying the amplitude of the pulse. As the pulse amplitude is varied, the TWC is sampled across the full range of the system. In our measurement, a reference trigger was generated with each pulse, and the time between the reference trigger and the time at which the electronics triggered for each channel was recorded. The map over the full time range for each channel then should be the time walk correction, shown in Figure 56. There is a clear discontinuity in the TWC curve between pulser voltages of 520, and 540mV. This is due to a change in the operating range of the pulser. As the pulser changes from one mode to another it may incur a noticeable change in the time between the reference trigger and the pulse. The magnitude of this change is not constant, and the effect is not always present. This change is not due to the PET

system electronics setup but to the pulser itself. In our system this discontinuity occurs outside the operating range of the system and does not cause an error in the calibration. However, it demonstrates a potential problem with this method.



An electronics measurement requires that the pulse be appropriate at all amplitudes and for all channels. If the shape of the pulse varies from one SSPM to another or for pulses of different amplitudes measured on the same SSPM, then a modeled pulse method is not appropriate. This was indeed clearly the case in SSPMs used in this work, as shown by the change in the signal shape as the number of fired microcells changes(122). This was due to the overlap of the individual microcell tails. In a pulse with more photons more microcell tails overlap and the rising edge broadens (123). The effect is illustrated in Figure 57, which demonstrates the effects of this pulse pileup

by comparing the pulse shape for signals on the same SSPM, but having different amplitudes. Each pulse in this figure was an average of 10 separate pulses with nearly identical amplitude, as measured by the total charge TQ. Three different ranges were compared, and the result indicates that the pulse shape clearly depended on the amplitude, with larger amplitude signals having a faster edge and a lower peak to total charge ratio. The rise times (10-90% of amplitude) are 28.4ns and 33.6ns for the maximum and minimum amplitude groups respectively. The large amplitude data set carries ~3 times as much charge as the low amplitude data set.



This indicates two problems with using an electronics pulse to calibrate the timing. First, signals of different magnitudes will have a different characteristic leading edge shape, leading to systematic errors in measuring the system response time for signals


of different charge (y-axis of Figure 56). Second, the pulse shape is accurate at no more than a single energy, while at other energies the amplitude of the electronic signal is not equivalent to a scintillation-based signal of the same charge. This would lead to errors in the energy measurement of the TWC curve (x-axis of Figure 56). These problems would be compounded if using a Constant Fraction Discriminator (CFD), which assumes that the pulse shape is independent of the pulse amplitude, to measure the timing as is common in PET systems. In SSPMs, the pulse shape does indeed vary with the pulse magnitude, due to both pulse pileup effects and leading edge rise time variations.

Additionally, any difference in the internal capacitance or quenching resistance between pixels in an SSPM array will cause a different characteristic microcell reset time and also change the shape of the SSPM pulse between pixels. This indicates that the electronically modeled pulse can be an accurate representation of no more than one energy point on each SSPM, and may in fact not accurately model any point for some of the SSPM pixels. Therefore, while the direct electronics measurement shown in Figure 56 is an appropriate approximation to the TWC curve, it has limitations. For these reasons, a more advanced method was developed and is described in 6.3.3.

6.2.2 *Electronic Charge Conversion Linearity*

While the TWC curve generated by the electronics measurement is not ideal, the electronics level charge conversion linearity measurement is an excellent measure of the electronics response to determining the charge of SSPM signals, since the electronics charge conversion technique does not make any assumptions about the signal shape, but depends solely on the total collected charge. This is in contrast to

more commonly used peak-sensing ADCs which operate on the assumption that the peak amplitude is proportional to the total signal. A peak-sensing ADC would not be able to accurately account for the pulse pileup effects of the shape, and rather than measuring the total charge generated by the SSPM, it would give the maximum amplitude value of the pulse shape, which is not linearly related to the total charge generated in a pulse, as demonstrated in Figure 57. However, in contrast to the SSPM pixels the electronics pulser does not have a pulse pileup effect, and so the total charge generated by the pulser is proportional to the amplitude of the pulser. In this case the amplitude of the pulser may be taken as a measure of the injected charge. As the nominal amplitude of the pulse changes, the mean of TQ is recorded. The response is clearly linear over a large range from the minimum trigger threshold to TQ~2000, after which saturation effects begin to become noticeable, as shown in Figure 58. On the same figure a histogram of the signal level on a single channel for 511keV events (energy window 400-600 keV) from the entire module is shown. This is not a complete energy spectrum, since it includes the response of only a single SSPM pixel, while the light is distributed over 8-16 pixels. This histogram includes higher energies when the event occurs close to the particular pixel and lower energy for events further away. For all such events, 99.99% have signal levels below TQ=2000, and so the electronics are linear over the full range of useful data.



The current limitation of this method of energy linearity measurement is the fan-in-fan-out board, which has a slight (3%) variation between output channels for any input signal. The ASIC response is not sensitive to the shape of the pulse and gives identical ratios of corrections between electronics channels when identical amounts of charge are injected regardless of the shape of the pulse. Since that is the case, future electronics boards will allow the FPGA to deliver a digital step pulse of varying time widths to the ASIC input. The response of the ASIC to the digital signal will then be able to determine the electronics gain calibration. This setup will eliminate the need for a pulser entirely for the energy linearity measurement. An additional benefit of the proposed method is that the new method will calibrate the electronics gain while the SSPM is connected to the ASIC, allowing a measurement of the noise and signal to noise ratio. The contribution of the noise will be determined by measuring the variance

in TQ both with the SiPM turned off and with different applied voltages and different trigger thresholds. A comparison of the SNR for various setup parameters will allow a determination of an optimized bias voltage and trigger threshold for each individual SSPM pixel by maximizing the SNR. This method has not been used previously, because when the ASIC board was designed, it was assumed that a pulser signal would need to be passed to the ASIC channels in order to perform the TWC measurement. Investigation of the proposed method will require the design and fabrication of a new electronics board.

The importance of the electronics charge conversion linearity calibration may be demonstrated by the improvement in the calculated flood map with and without the electronics calibration applied. Figure 59 shows a flood map calculated without the electronics calibration, that is using the raw data from the TQ measurement. Figure 60 is a plot of the same data set with the electronics properly calibrated. There is a noticeable improvement in the flood map once an electronics calibration is applied, allowing a much more accurate determination of the crystal boundaries. Notice especially the improvement in the separation of columns 1 and 2. The data from the corrected map are what indicated a smaller pixel pitch would be possible, resulting in the design of our third prototype array. Without the electronics gain improvement, it would be difficult to add more pixels without merging some of the pixels, or causing larger artifacts at the corners. A similar improvement was realized in the energy spectrum for each module.




Figure 59 Flood Map without electronics correction




Figure 60 Flood Map for same data set with the electronics calibration applied

6.3 *Full Detector Array Calibration*

Once the electronics were calibrated, the full module could be calibrated. After the crystals were bonded to the SSPM array, the individual SSPM pixels became inaccessible, and so could not be calibrated. This is unfortunate, since the manufacturer's specified variation between pixels of 5% is large enough to cause some degradation both in the energy and in the calculated position of individual pulses. A DOI based calibration may be able to reduce any errors caused by differences in the gain between SSPM pixels, but for reasons described in 6.2, the first approach was to assume a uniform SSPM response, which proved to be sufficient for this work. It may be possible to develop an algorithm that minimizes the overall energy resolution of the module by varying the nominal gain of each SiPM in software, but algorithms with only one variable for the gain at each end of the detector module (118, 119) are inappropriate because a complete description would need to contain a characterization of each SSPM pixel. A DOI-based detector may balance the SSPM array gain through centering of the R-space histogram. The inclusion of 64 crystals and 32 detectors may allow such balances to treat each SSPM pixel separately and thereby determine the relative gain variance between crystals. Quadratic corrections are only needed if the light detection efficiency varies with the depth of interaction(96). Ideally, the intrinsic gain can be measured directly by looking at the characteristics of single microcell events (124-126). In the current module, the dark count rate is high enough that these events overlap and are indistinguishable, but cooling significantly reduces the dark count rate and may allow a measurement of single microcell events (127). The limitation of the single photon measurements is the need to cool the SSPM. However, if active cooling were added to this system, and the SSPMs were cooled to a level where

the difference between events which vary by only one microcell were measurable, a conversion of TQ to the number of detected photons could become possible. Additionally, SensL is in the process of producing a much more advanced SSPM module with a faster microcell discharge time that may allow a separation of individual dark counts, and so be capable of individual microcell counting. In the future, such methods and algorithms may be investigated, but are not a part of this work.

6.3.1 Crystal Segmentation

Flood maps are initially calculated using a center of mass type algorithm, which is similar to but not identical to standard Anger logic.

$$X_n = \frac{\sum_i x_i * E_i}{\sum_i E_i} \quad \text{Equation 6-1}$$

$$Y_n = \frac{\sum_i y_i * E_i}{\sum_i E_i} \quad \text{Equation 6-2}$$

Here X and Y are the final calculated positions, x_i and y_i are the known positions of the individual SSPM pixels, and E_i is the calibrated energy deposited on each pixel. The crystals are clearly separable, and the current calibration involves the user applying a global energy cut to the flood map and manually locating the apparent centroid for each crystal in a flood image. Once the crystal centroids have been located, a perpendicular bisector between each crystal centroid and its nearest neighbors is calculated. The area within the four perpendicular bisectors is taken to be the portion of the flood map corresponding to an individual crystal. Crystals at the perimeter are assumed to extend to the boundary of the field. If the pixels are on a rectangular grid, the perpendicular bisector of the next-nearest neighbors will intersect with the perpendicular bisector of

the two nearest neighbors. However, in a physical detector the points are not on a perfect grid, and so the next nearest neighbor needs to be considered as well. In an improved calculation the next nearest neighbor would be included if its perpendicular bisector enters the box formed by the perpendicular bisectors of the nearest neighbors, and neglected otherwise. Currently only the nearest neighbors are included because the corrections are in low count rate regions of the flood map that correspond to Compton scatter. A correction in this area of the map would likely not improve the overall results. These secondary corrections may be more important in arrays with poorer crystal separation.

The initial user selection is made on a flood map generated from data collected using a ^{137}Cs flood source. One validation of the map is to verify that the segmentation produced with this data set works well on data sets generated with other sources. The separation generated from the ^{137}Cs source is shown in Figure 61 followed by the ^{22}Na in Figure 62, and background shown in Figure 63. This is the most user intensive portion of the calibration procedure, and while it produces excellent results, it remains the area with the most room for improvement through automation.




Figure 61 Flood Map with ^{137}Cs source. Energy linearity correction applied to individual crystals. Crystal boundaries overlaid




Figure 62 Flood Map with ^{22}Na source. Energy linearity correction applied to individual crystals. Crystal boundaries overlaid




Figure 63 Flood Map with Background events. Energy linearity correction applied to individual crystals. Crystal boundaries overlaid

The background calibration has the poorest peak-to-valley separation between crystals of the three measurements. This is likely due to mispositioning of events arising from the simultaneous emission of a low energy γ with the intrinsic β decay. The intrinsic γ will travel some distance within the module before either escaping or interacting with the detector. If the γ interacts within the module, the measured position will be an energy weighted mean of the beta and gamma interactions, and the determined position will vary slightly from the true event location. However, in spite of this unavoidable blurring, the separation is still good, with all points clearly separable and the calculated crystal boundaries from the ^{137}Cs flood map accurately separating the crystals.

6.3.2 Crystal Energy Linearity

Once the crystal boundaries have been determined, data from three different data sets were used to determine the crystal level energy linearity. The three data sets included data from an intrinsic background measurement, ^{137}Cs , and ^{22}Na flood sources, segmented to correspond to the boundaries of a single crystal. The initial electronics calibrated event level energy E_n is calculated as below.


$$E_n = \sum_i E_i \quad \text{Equation 6-3}$$

In this equation E_i is the electronics calibrated signal level from each fired SSPM pixel. These data were binned into a histogram, and background subtraction is performed using a background histogram that has been normalized based on collection time, but without live time corrections. This background correction is applied to both the ^{137}Cs and ^{22}Na data sets. An initial Gaussian fit is calculated for the 662 keV peak in the ^{137}Cs flood data and the 511 keV photopeak in the ^{22}Na data set. Additionally a half-hyperbola is fit to the beta tail of the background data with a 1198 keV vertex and coupled to zeros after 1198 keV. While there is not a physical model to justify the use of a hyperbola for the end point of the β tail, it does an excellent job of approximating the data within the error bounds using a limited number of variables. A line is fit to these three points, and if sufficient statistics are available, a Gaussian is fit to the ^{22}Na data set in the region of the 1275 keV photopeak. The 1275 keV gammas typically undergo several Compton interactions and so are rarely mapped to the edge crystals of the module. The ultra-low count rate of 1275 keV gammas at the corner of the detector array makes fits to a 1275 keV photopeak for the corner crystals impossible. In more central crystals, if sufficient counts are available to fit the 1257 keV photopeak, the

linearity is once again calculated by a linear fit of these 4 points. However, this initial fit has certain limitations. First, it is most accurate over the range above 511 keV, precisely the region which should never apply in true PET data. Second, there is no error weighting in the linear fit, so the low statistics 1275 keV data point is treated a weight equal to that of the high statistics 511 and 662 keV data points. The only way to appropriately weigh those three points would be to weight based on the area of the Gaussian curves, a metric which would necessarily exclude the non-Gaussian 1198 keV β tail. It would be possible to remove the 1198 point altogether, but there are many instances where the 1198 keV is a useful data point, but the 1275 keV photopeak does not have sufficient statistics for a meaningful fit, and so preserving the 1198 keV data point is preferred, if possible.

In order to improve the fit, the energy spectra are passed to a fitting function designed to error-weight the four data points listed above and also the 202 and 307 keV photopeaks in the background spectrum, which are clearly visible in the background energy spectra of many crystals. The fitting routine fixes the centroids of the five photopeaks at their standard energies in keV, and the vertex of the hyperbola marking the endpoint of the beta tail at 1198 keV. The other parameters, such as Gaussian amplitude and width, are allowed to vary, but are initialized based on the initial fit, and constrained to physically meaningful boundaries (for example, the Gaussian curves can never have a negative amplitude). The linearity response that maps the energy spectra onto a physical energy scale is also initialized based on the three or four data point initial linear fit, but it is allowed to vary as well. A residual between each of the physical spectra and the calculated fit applicable to that spectrum is calculated. The fit then minimizes this residual by varying the linearity scale, the width and amplitude of the

individual Gaussians, and the slope and curvature of the hyperbola. In minimizing the residual, this method inherently gives less weight to the beta tail, which has a large range over which there are near fits with the hyperbola, and weights the photopeaks according to their count statistics. This matches what should be the case as the β tail should be under-weighted, since the endpoint is a low count region, and has more uncertainty. Additionally, this method extends the range of the fit to as low as 202 keV, which covers the meaningful range of PET data. In some crystals the energy threshold extends into the 202 keV region, and the counts in that region are 0. This method fits a Gaussian of amplitude 0 in that region, which gives this region an effective weight of 0, thus appropriately treating regions where the 202 keV data is unavailable. The 202 and 307 keV portions are fit together, while every other energy zone has its own separate histogram. This is due to the intrinsic overlap of these two peaks in the background data. Results for several representative crystals are shown in Figure 64.



While this method is admittedly complicated, it does improve the system performance. Every attempt to calculate the energy linearity using only 662 and 511 keV photopeaks ran into consistency limitations due to the initial choice of the range of each spectrum to fit. A different choice of the energy ranges caused a large variance in the calculated position of the 662 keV photopeak. The 662 keV data point could be used to narrow the range for the 511 keV, but that selection had mistakes in crystals with a high energy threshold. The 1198 keV data point was available to improve the fit, but has some intrinsic uncertainty that calls for a lower weighing of this data point. Additionally, an appropriate fit could only be calculated for the 1275 keV data point in about half the crystals. These limitations caused more simple automated methods to

fail in the fitting of some crystals or to severely overestimate the energy resolution and underestimate the performance of many of the crystals. The current method was developed to overcome these limitations, and the overall efficacy of this method is demonstrated by the excellent correlation between single crystal energy resolution and the full module energy resolution, as described in 7.1.

6.3.3 Time Walk Correction Calibration

The electronics data gives an initial approximation of the TWC curve, but they have fundamental inaccuracies that have been described previously. An ideal time walk calibration is a correction such that the corrected signal time is independent of the number of photons collected. This means that such a calibration, when applied to two photodetectors observing the same gamma-ray interaction but each collecting a different amount of light and so each having a different raw time measurement, T_1 , will generate correct times that vary only by the noise of the two detectors and have no systematic offset due to the SSPM pixel, or signal level. The method described below begins with this fundamental assumption and calculates a TWC curve based on the differences between measured time stamps within individual scintillation events. The differences are weighted according to the relative goodness of the data, either by assuming the goodness is proportional to the total charge collected or by estimating the error of each channel for each signal level. Weighing the differences according to the total charge collected is based on the assumption that the error in timing will be proportional to the slope of the signal at the leading edge pickoff and to the intrinsic noise. In the pickoff region the noise is dominated by the dark counts of the SSPM and will be roughly independent of the signal amplitude, while the slope of the pulse will roughly scale as the total charge of the pulse. However, it has been found that some

SSPM pixels have more intrinsic noise than others. This means that although two pixels may have an identical amount of collected charge, one of the pixels could perform much more poorly due to intrinsic noise differences. This difference cannot be accounted for in an energy-weighted approach, but can be accounted for by using an energy dependent error estimate of the error for each SSPM pixel.

The calibration begins with a TWC curve with a correction of 0 for all channels at all energies. An energy-weighted average time is calculated for each event, and the difference for each channel from the energy weighted average time is recorded according to the channel number and the channel signal level. All differences for a given channel at a given energy are then averaged, and the median of these values is taken as the new TWC for that channel at that energy. The median is chosen in order to reduce the effect of event overlap. The time window that allows signals on various channels to be grouped into a single event must be at least as wide as the 40-60ns rise time of the signal. In this broad time, multiple events may occur, and these events would skew a calculation of the mean for the TWC. A median reduces this skew and provides a more stable estimate of the TWC response of the system. The region corresponding to the central 68.2% of the data is taken to correspond to $\pm\sigma$, which is then used to calculate a variance estimate for the channel at that energy. This central region calculation is used rather than a standard deviation in order to reduce the effect of overlapping events skewing the estimated error. The process is then repeated, but using the variance rather than the energy for weighting. The process is detailed below.

$$T_{m,n}^* = T1_{m,n} - TWC(Ch_{m,n}, E_{m,n}) \quad \text{Equation 6-4}$$

A full data set will consist of ~2M scintillation events, each with 3-32 individual pixels firing. The index m refers to the scintillation event number, while the index n refers to the individual signal within the group of signals corresponding to a single scintillation event. The TWC is determined for each signal within each event according to both the channel number of the signal and the total charge collected. During the first iteration, the TWC value is zero for all signals in all events.

$$T_{ref_m} = \frac{\sum_n T_{m,n}^* * W_{m,n}}{\sum_n W_{m,n}} \quad \text{Equation 6-5}$$

A weighted average of the TWC-corrected time is then calculated for each scintillation event. The initial weighting W is defined as the collected charge, TQ , while the weights in subsequent iterations are set equal to the error estimate associated with the channel and energy of the signal. Alternatively the weighting could be left as the energy in all cases. The technique is identical, and the error can still be estimated for all channels at all energies. However, it will be the case that some SSPM pixels will have a higher dark count rate or slower peaking time, indicating that the error of that channel is higher for a given collected charge than a neighboring pixel. Even if the non-linear shape of the leading edge of the signal is accounted for, an energy weighted average will not attempt to correct for SSPM level effects, while an error weighted average does attempt to incorporate all sources of error. The results of each of the variations as measured by coincidence resolving time are detailed in Chapter 8. During the first iteration there is no estimate of the energy and channel dependent error, and so either a raw average or energy weighted average must be used.

$$\Delta T_{m,n} = T_{m,n}^* - T_{ref_m} \quad \text{Equation 6-6}$$

The next step calculates the difference between the TWC corrected signal time and the weighted average for all signals within each scintillation event.

$$TWC'(i, j) = TWC(i, j) + P_{50}(\{\Delta T_{m,n} | Ch_{m,n} = i, E_{m,n} = j\}) \quad \text{Equation 6-7}$$

The next step is to determine the median of the distribution of all events for each channel at each energy and to use this to calculate the correction for the TWC for the given channel and energy. The median is the fiftieth percentile of the distribution and is referred to as P_{50} . Note the mixed indices in this equation. The index i refers to the channel number, while the index j refers to the total charge collected. This equation determines the new TWC value for a given channel and energy as the TWC value for that given channel and energy minus the median of the difference between the TWC-corrected time stamp for all signals of the given channel and energy within all scintillation events containing the appropriate signal level and channel.

$$Err(i, j) = (P_{84} - P_{16})(\{\Delta T_{m,n} | Ch_{m,n} = i, E_{m,n} = j\})/2 \quad \text{Equation 6-8}$$

The next step is to estimate the error for the given channel and energy. If the data are normally distributed a standard deviation would be an appropriate estimate. However, the difference between the T^* and T_{ref} is not always normally distributed or even symmetric. Additionally, the data set as collected contains outliers. While the outliers can be rejected through standard outlier rejection methods, those methods can prove to be time consuming when calculated for data sets of the size used in this analysis. Specifically, there are 512 electronic channels, each with 4096 energy levels analyzed in our current small system. If an outlier rejection method were scaled to a full system the time commitment to calibration would be significant. Therefore an estimate has been made that the error is proportional to half of the width of the distribution covering 16-84% central data range. If the data are normally distributed this is exactly

equivalent to two standard deviations. This method also does a decent job of rejecting outliers, since those are relatively few in number and occur well outside the 16-84% probability range. In Equation 6-7 the median value was used to determine the new value for the TWC for a given channel and energy. The median was selected rather than the mean in order to reduce the effect of outliers on the final result and also due to computation time requirements. However, since a central region is calculated in order to estimate the energy dependent error for each channel an alternative method of determining the updated TWC value would be to take the mean of the data within the probability region. If the distribution is symmetric then the two methods should produce identical results to within the error of the determination. However, the two methods may have different levels of stability for the event distribution and if the distribution is not symmetric then the values will not necessarily be identical as the mean minimizes the norm of the second order deviation while the median minimized the norm of the first order deviation. The differences in the results of these two methods has not been investigated.

During the iterative process, the higher energy events show less error and so have a higher weighting. In any given event group there is typically only a single channel with an event energy greater than $TQ=1000$. The method tended to push the calculated energy dependent error estimate in this energy range to zero, which then produced event level timing error estimates of near zero. This was an artificial result, not reflective of the true physical performance of the system. The electronics only calibration data measured the electronics only error at each energy level, and so the overall variance for each energy was determined to be the larger of the electronics-only error or the distribution variance. This restriction prevented any points from being

calculated with an expected variance of zero. Ideally an accurate measure of the electronics only error could be determined by passing a step pulse to the system directly from the FPGA and measuring the jitter in the ASIC processed pulse with the SSPM active. This would accurately include the SSPM dark count noise, which is currently ignored.


$$TWC(i,j) = TWC'(i,j) \quad \text{Equation 6-9}$$

Finally, the TWC curve is updated with the newly calculated TWC values, and the process is run again. Equation 6-4 through Equation 6-9 are repeated until the calculation converges. Currently a robust convergence test is not in place, although a test can be developed if necessary.

While this result works well when sufficient statistics are available, in the range above TQ~1200 the statistics are poor, and the results tend to be noisy. In order to suppress this noise, a cutoff is applied when fewer than 20 events occur within the energy band used for the calculation. A parametric fit is applied to the higher energy data where sufficient statistics are not available to perform the calculations as described.

The initial calculation of the energy-weighted time will be inaccurate insomuch as the true TWC is not zero, and so a single iteration does not calculate the true TWC, but merely a closer approximation. Once the TWC has been calculated, the values are used to determine T* for each time stamp, and the process repeats. Over the range of experimental data the result converges to within 1 bin of the final calculation within two iterations, and within 0.1 bins (31ps) within 5 iterations. The results from the first eight

iterations of the calculation for a single module are shown in Figure 65, and are compared to the electronic calibration, which has been given an arbitrary offset to vertically align the two measurements in Figure 66. The two methods show clear similarities, providing a check on the developed scintillation event based calculation, but show differences as well. These differences are expected, since the pulse used in the electronically measured TWC curve does not appropriately vary with signal amplitude, or between pixels.






Figure 66 Calculated TWC against measured TWC. Above TQ=2000 the counts statistics are low and not very meaningful


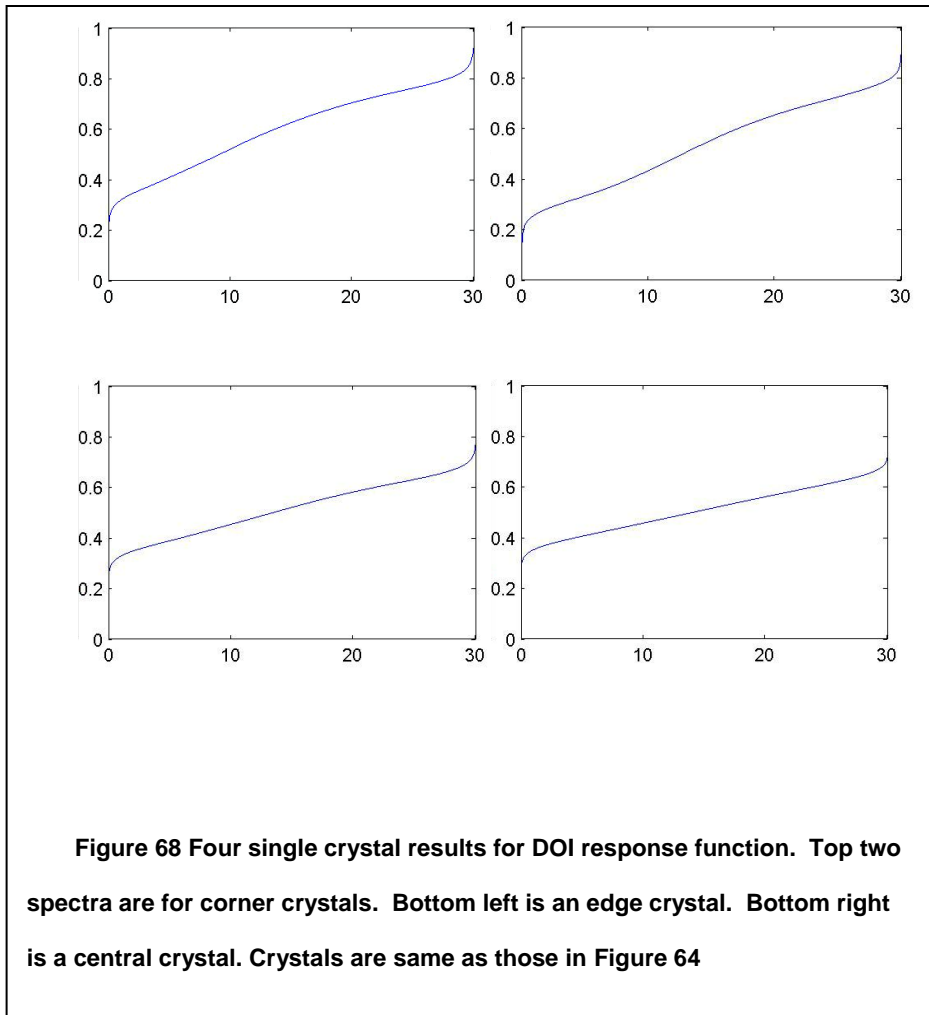


Figure 67 Time Time Walk Error, 8 iterations. Left-Full curve, Right-Exploded region of same plot.

6.3.4 DOI Response Function

After the crystals have been segmented, and the energy linearity at the crystal level calibrated, the DOI response function may be measured. In order to perform this calculation, the background measurement is cut on energies from 400 to 1100 keV, and the method described in 3.2.3 is applied to the subset. While the energy cut is not necessary in the single crystal studies, it is critical in the array due to the inclusion of nearby modules. The 88, 202, and 307 keV gamma events that are emitted from nearby modules are not uniform in their detected distribution. Although the distribution of detected gamma-rays originating from contiguous modules is uniform along the crystal length, the distribution of detected gamma-rays originating from modules on the opposite side of the bore compared to the test module will be weighted toward the bore side of the test module. This difference will lead to an error in the overall determination of the DOI resolution insomuch as the counts at the entrance window increase due to emitted radiation from other modules. In our prototype system of two banks with 4 modules each, this is not significant. However, in a full system with hundreds of modules, this may generate a much larger effect.



6.4 Summary and Discussion

The methods described in this chapter are able to calibrate both the electronic and crystal properties of the module. These are fast, reliable methods to calculate the channel level electronic gain, the crystal interaction location, the electronic Time Walk Correction curve, the scintillation event based Time Walk Correction curve, the depth of interaction response function, and the light level to signal level conversion. The methods are robust and require operator intervention only to determine the location of

the crystal centers and to verify or correct a very small number of energy spectrum fits. These methods are able to perform all additional calculations with no intervention.

The automated energy calibration works for all 256 crystals with no failures, with an implied success rate of greater than 98%. The energy calibration treats all collected data in a statistically appropriate manner and includes a fit from 202 through 1275 keV. In the future, non-linear energy response models may be investigated, but currently they do not appear necessary. The DOI response function calibration treats each crystal separately, and does not force the assumption of a linear DOI response function. The timing calibration is based on true scintillation events and does not make the incorrect assumption that the shape of the electronic pulse is independent of pulse pileup in the microcell responses or that the shape is identical for all crystals. The electronic and scintillation response shapes are similar enough to provide validation to the scintillation event based method but are different enough to warrant the use of a scintillation event based method if it is more accurate. This method also allows a calculation of the error of the TWC at each energy, which may add further usefulness. For more on both of these topics see the in-depth discussion of timing results in Chapter 8.

The areas that may still be improved include automation of the crystal location algorithm, the electronics energy calibration, and the error calculation of the TWC curve. The crystal segmentation portion of the calibration requires a large amount of user interaction, and inconsistencies may carry through to other parts of the calibration. Algorithms that find 2D peaks should be developed, but modified to also include prior information such as the uniform distribution of crystals in physical space. The energy


linearity should be updated to use an input pulse from the FPGA in order to remove the FIFO board and allow a full electronics calibration in less than a minute. This would be in contrast to the current technique which involves a large overhead time commitment when attaching the FIFO module. The energy calibration could be performed before each data acquisition, allowing a thorough diagnostic of the system over time. Additionally an effort needs to be made to account for the gain normalization of the SSPM pixels. This may be done through single microcell event counting using a cooled system, but such approaches are not currently feasible with the described system. New photodetectors with faster microcell response time and lower dark count rates may improve the ability to do single microcell processing and allow this variance to be calibrated. Alternatively new algorithms may be developed that calculate nominal gain differences for each SSPM pixel by minimizing the energy resolution of each of the 64 crystals, and centering the R-space histograms of each of the 64 crystals. The error bounds of the TWC determination are somewhat arbitrary, and require a test pulse in order to determine the electronic only error of the system. An energy weighting removes this problem, since the weighting is not updated each iteration and so cannot drift toward zero. However, an energy weighting does not account for variations between pixel response. An improved method may be to use an FPGA based digital pulse for the energy response function and include a comparison with the SSPM arrays set to various bias voltages. The dataset with the arrays off will give an indication of the electronics only contribution, while the data set with arrays at each potential bias will give an indication of the SSPM contribution. This estimate could then be used as the lower bound of the TWC error, rather than the electronics only measurement. This may also allow an improvement of the determination of the operating conditions of the


system by setting the bias of each SSPM pixel to the point which optimizes the SNR for the individual pixel.

Chapter 7 Non-Timing Module Performance

7.1 Module Level Energy Resolution

Once the energy linearity measurement has been performed for each individual crystal, the total energy is scaled by the linearity factor, and results from all crystals are combined for the full module. The results from the background measurement before and after the crystal level energy scale alignments were applied are shown in Figure 69 and Figure 70 respectively. The 202 and 307 keV peaks are clearly separable after the alignment.





Once the energy scale has been calibrated using the six energy points, the energy resolution of the module was calculated according to Equation 7-1. The energy resolution of this module was 15% at 662 keV, and 19% at 511keV.

$$\text{Energy Resolution} = \frac{\text{Full Width at Half Maximum}}{\text{Centroid Energy}} \quad \text{Equation 7-1}$$

This calculation was made from raw data, and the 511keV energy resolution improved to 15% for data collected in coincidence with a Na²² source which eliminated almost all background events. In contrast, without the crystal level alignment applied, the 511 keV photopeak was not distinguishable on all modules, and in some cases the energy resolution could not be determined.




Figure 71 Energy Spectrum for ^{22}Na Data After Alignment. 511 and 1275 keV photopeaks are clear, energy resolution at 511 keV is 19%.




Figure 72 Energy Spectrum for ^{137}Cs Data After Alignment. Energy resolution at 662 keV is 15%.

7.2 *Depth Effects on Crystal Flood Map*

In order to test depth effects on the flood map a collimated source was used to generate interactions at different depths within the crystal. The x and y positions for events at each depth were calculated and plotted. The results at 5mm are displayed in Figure 73, the results at 15mm are displayed in Figure 74, and the results at 25mm are displayed in Figure 75. The three flood maps appear almost identical, however a significant shift occurred in the calculated position of events corresponding to individual crystals from one end of the module to the other. This shift can be seen by plotting the difference between Figure 73 and Figure 75. The result is shown in Figure 76. However, in the current modules the shift is smaller than the space separating two crystals and so does not affect the ability to separate crystals in a global flood map. The shift was not uniform in either magnitude or direction among the various crystals. The likely cause of the shift is a difference in the alignment of the light guide and SSPM array for the two ends. A small difference in the alignment of the light guide and photodetector could cause a significant difference in the light sharing pattern for the two ends.




Figure 73 Flood Map at 5mm Depth of Interaction




Figure 74 Flood Map at 15mm Depth of Interaction




Figure 75 Flood Map at 25mm Depth of Interaction





Figure 76 Difference between Flood Maps at 5 and 25mm DOI

7.3 *Crystal Identification*

Once all crystals have been identified, and the energy for each crystal has been calibrated, the ability to resolve individual crystal elements must be determined. The metric used to describe the ability to resolve crystals is the peak-to-valley ratio, or the maximum to minimum ratio of a cross-section through the counts histogram taken to run through the crystals being measured. This cross-section should be only as wide as the visible pixel center in order to minimize the effect of the background counts and should include an energy cut in keV in order to isolate photopeak events. A cut across the 4th row of the flood map in Figure 62, with a global energy cut of 350-600 keV is shown in Figure 77. The peak-to-valley ratio ranges from 7:1 to 9:1 over this cut. A similar analysis for the fifth column is shown in Figure 78 with a peak-to-valley ratio of 10:1 to 20:1. However, the flood map shows streaks in the x and y directions, which are visible in the cross-section profile as a low level plateau between the peaks, in contrast to a more u-shaped valley which would correspond to the overlap of the crystals. This plateau is most likely caused by Compton scattering between pixels. In a Compton event the reduced energy gamma will travel a short distance within the module and deposit the remainder of its energy in a nearby crystal. The detected signal will be the sum of the two or three different interactions, each producing light, and the final position calculation will effectively be an energy weighted average of the various interactions. In order to reduce the effect of Compton scatter, I also look at the peak-to-valley separating the next nearest neighbors, that is at a 45° angle to the x and y axes. This result is shown in Figure 79 and again demonstrates a peak-to-valley ratio of greater than 10:1, with a clear plateau between peaks. Only seven crystals are visible in this image, due to a shifting of a corner crystal that caused it to lie off of the rotated axis.

The limiting factor in crystal identification is not the physics of optical transport or separation, but Compton scatter within the array. This indicates that if desired, a higher resolution array design could be designed, especially if the crystal segmentation map is altered to incorporate DOI information.






Figure 78 Crossectional Profile along y-axis. Peak-to-valley is ~10:1





Figure 79 Crossectional Profile at 45° angle to x-axis.


Peak-to-valley is ~15:1

7.4 DOI Resolution and Function Comparison

The prototype array was scanned along its axial length in a setup similar to that described in 2.2. However, in this case the collimation crystal was 1x2x10mm, and a 1mm Na²² point source was placed in the center of a collimation crystal and test module separated by 140 mm. The FPGA communication card has a lower limit of the amount of data it can pass, so it was impossible to measure the beam spread as described in 2.2 for individual crystals, since the count rate is unreliable once it passes into a range where the data transfer drops to a low rate. However, the raw event rate could be determined from the FPGA board before transfer. This allowed a determination of the average beam width across the entire detector module. Ideally information would be available for the geometric beam spread at each crystal. In this setup the average beam width was measured as 1.8mm FWHM as show in Figure 80. For this narrow beam and large separation the expected difference in the geometric beam spread between the entrance and exit is less than 0.3 mm. In this case the difference in the geometric beam spread for different crystals can be neglected. Data were taken at 11 points across the crystal depth in 2.5mm steps. Crystals were analyzed at each depth for the DOI resolution, Energy resolution, and 511 keV photopeak position on the energy calibrated scale.




The DOI response function for the central crystals was approximately linear, and so for these crystals Yang's method (39) could have been applied. However, the edge and corner crystals are non-linear, and so Yang's method would not be appropriate for those crystals. Since a more complicated method must be applied to some crystals, it is preferred to use the more accurate method for all crystals. A comparison of the DOI response function measured from an externally collimated beam and as calculated from the externally collimated point source is shown in Figure 81.



The DOI resolution may be calculated in a form similar to that described in 3.2.2 for each crystal. However, in the array data the non-normal distribution of the data is more significant and a Gaussian fit is not appropriate for the electronically collimated data. The DOI error at each position is calculated as 2.35 times half of the range covering the central 68.2% of the data. If the data are normally distributed this calculation is exactly equivalent to the FWHM of a Gaussian fit. In this investigation the DOI resolution was found to depend on the crystal location, with crystals at the perimeter generating a better DOI resolution than crystals near the center of the array. This is likely due to a combined effect of both light diffusion at the rough boundaries with air gaps and also light loss via transmission through the ESR reflector. Additionally, the DOI resolution varied across the length of any individual crystal, with a better performance in the center

than close to the edge. During the determination of the position, we have *a priori* knowledge of the physical boundaries of the crystal module and that the scintillation event must have taken place within these physical boundaries. Thus, while the flat nature of the DOI response function would indicate a very poor DOI resolution, the event is known to have occurred within the 0 to 30mm range, although the DOI error would allow ranges from -7 to 38 mm. At the extreme edge, the DOI resolution is approximately half of what would be expected from the spread of the function in R-space.



As shown in Table 5, the DOI resolution varied within the range of 6-8 mm throughout the array. The mean resolution of central crystals was 6.7mm, while the

mean resolution of the edge crystals was 5.9 mm and the mean resolution of the corner crystals was 6.5 mm. When calculating the averages, the crystals at the top left of the tables (DOI resolution of 10.3 and 9.9) were omitted. The collimated beam did not fully intersect with these crystals, and so the DOI resolution determination is inaccurate.

Table 5 DOI resolution of all crystals within an 8x8 module, averaged over 11 positions

10.3	9.9	9.1	8.3	7.6	7.4	7.8	6.9
6.3	6.3	6.6	6.3	5.9	6.1	5.6	5.3
6.7	6.9	7.2	7.4	7.1	7.2	6.6	6.0
6.2	6.4	6.3	6.0	5.8	6.1	5.6	5.1
6.7	7.2	7.7	7.4	7.0	7.2	6.7	6.3
6.1	6.7	7.2	6.5	6.1	6.3	5.6	5.0
6.7	7.4	8.1	7.7	7.2	7.3	6.7	6.3
5.8	6.4	6.4	6.3	6.0	5.8	5.6	5.1

The difference between the maximum and minimum resolution varied between 2.2 mm and 5.6 mm. The interior crystals had an average range of 3.5 mm, the edge crystals had an average range of 3.8mm, and the corner crystals had an average range of 4.2mm, as shown in Table 6.

Table 6 Maximum-Minimum DOI resolution among depths within each crystal in an 8x8 array

8.6	7.1	5.7	5.5	4.9	3.0	5.1	3.6
2.8	2.0	3.3	2.6	2.4	2.2	1.3	2.4
3.4	2.7	4.3	3.7	3.9	3.7	2.7	3.0
3.6	3.1	4.4	3.7	3.9	3.2	2.9	3.8
4.3	3.8	4.7	4.6	4.4	4.3	3.1	3.9
4.2	2.9	5.0	3.7	3.8	3.3	2.5	3.2
4.4	3.0	4.4	3.1	3.0	2.7	2.6	3.7
4.0	4.1	4.2	3.9	3.4	3.4	2.7	3.8


The variation among crystals at each particular depth is summarized in Table 7.

Table 7 Summary of Depth effects on DOI resolution for various crystals within a detector module

Depth	Central Crystals			Non-Corner Edge Crystals			Corner Crystals		
	Avg	Min	Max	Avg	Min	Max	Avg	Min	Max
2.5	5.5	3.6	7.9	5.4	4.0	10.3	4.7	4.1	5.4
5	6.8	5.3	8.7	6.8	5.6	8.4	6.8	5.9	7.7
7.5	7.5	6.0	9.6	7.6	6.1	8.7	7.4	6.7	8.1
10	7.4	5.6	9.0	7.1	5.5	8.6	7.0	6.1	7.6
12.5	7.2	5.4	8.7	6.7	5.2	8.4	5.9	5.0	6.5
15	7.1	5.3	8.7	6.8	5.1	10.8	5.5	4.6	6.0
17.5	7.1	5.0	8.7	6.6	4.8	11.0	5.3	4.5	5.7
20	7.1	5.5	8.6	6.7	5.0	11.6	5.5	5.1	6.5
22.5	7.2	5.8	8.2	6.6	5.2	8.7	6.0	5.1	7.3
25	6.6	5.8	7.5	6.8	5.0	10.4	6.2	5.5	7.6
27.5	4.3	2.9	5.6	4.3	2.7	8.0	4.9	2.9	8.2

7.5 *Light Level Uniformity*

The signal level of the 511 keV photopeak was determined using the linearity fit for each individual crystal. This uniformity is plotted for each crystal in the module in order to determine the overall module uniformity. Notice that the light output is lower at the edges than in the center. The most likely cause of this effect is light escaping through the ESR at the edge of the module.




0.67	1.14	1.06	1.18	1.24	1.23	1.19	0.64
0.79	1.03	1.24	1.31	1.38	1.36	1.15	0.76
0.97	1.21	1.32	1.44	1.52	1.50	1.28	0.88
1.17	1.34	1.50	1.50	1.54	1.49	1.34	1.09
1.22	1.39	1.47	1.53	1.52	1.53	1.34	1.12
1.22	1.32	1.46	1.49	1.50	1.44	1.29	1.06
1.13	1.21	1.38	1.40	1.40	1.41	1.22	0.95
0.89	1.06	1.25	1.26	1.29	1.20	1.25	0.69

Figure 83 Uniformity of Light Detection Efficiency for all crystals in a single Module

7.6 Energy Resolution Uniformity


The energy resolution of the 511 keV photopeak was determined for each crystal after all corrections had been applied. This is an interesting measurement because it can determine the effect of light loss on the energy resolution of the individual crystals. There is no correction for depth effects, and so based on the results discussed in 3.3.5, I expect that areas with greater light loss will suffer greater energy resolution degradation. The energy resolution uniformity plot is shown in Figure 84, and a plot

showing the correlation between the signal level of the 511 keV photopeak, and the energy resolution is shown in Figure 85. In that figure it is clear that crystals with poorer keV to electronic signal level conversion ratios have poorer energy resolution.



31%	25%	21%	18%	19%	16%	25%	30%
17%	19%	15%	16%	16%	16%	19%	20%
19%	18%	17%	15%	17%	18%	21%	22%
14%	16%	17%	16%	16%	18%	20%	22%
16%	19%	17%	15%	16%	16%	18%	21%
18%	18%	17%	17%	17%	17%	18%	22%
20%	17%	17%	14%	16%	14%	15%	19%
26%	18%	18%	16%	17%	18%	19%	21%

Figure 84 Energy Resolution Uniformity for all crystals in a single module



7.6.1 Depth Effects on Energy Resolution

The results presented thus far assume there is no light loss within the crystal. If there is a difference in the light collection efficiency for events originating at different depths in the array, then the full crystal energy resolution will be worse than the energy resolution for a single depth. In order to determine the depth corrected energy resolution the crystal the energy resolution within a 4mm window around 27 different locations were averaged. This correction should account for depth variations within each crystal. Due to the reduced counts in the DOI restricted set, background subtractions could not be made. In Table 8 the energy resolution is the calculated energy resolution of a Gaussian fit applied to the 511 keV photopeak without a

background correction. Additionally, this method treats different depths with equal weight regardless of the number of counts observed at each depth. The method used in 7.6 included a background subtraction and assumed each event carries equal weight regardless of depth, thereby weighing the energy spectrum towards the entrance face of the detector module. These differences should give an expectation of slight differences between Table 8 and Figure 84. However, the overall resolution numbers are similar in magnitude and trends.

Table 8 Energy Resolution for each crystal within an 8x8 array, averaged among various different depths.

16%	15%	23%	19%	18%	18%	17%	24%
22%	15%	16%	14%	14%	15%	15%	18%
20%	15%	15%	15%	15%	15%	16%	19%
17%	15%	14%	15%	14%	15%	15%	17%
17%	14%	15%	15%	15%	15%	16%	17%
16%	14%	15%	14%	14%	15%	15%	17%
17%	14%	14%	15%	14%	15%	16%	19%
22%	16%	17%	17%	16%	16%	14%	20%

Figure 86 shows the energy resolution for different depths for three representative crystals within the array. Notice the general trend of the corner crystal having poorer energy resolution than the edge crystal at almost all depths. This poorer baseline energy resolution combined with a greater light loss linear deviation yields even poorer energy resolution for the corner crystals when the full length is considered without depth dependent corrections.




Table 9 Depth Effect on energy resolution for different groups of crystals throughout module

Depth	Central Crystals			Non-Corner Edge Crystals			Corner Crystals		
	Avg	Min	Max	Avg	Min	Max	Avg	Min	Max
2.5	15%	13%	17%	17%	13%	22%	18%	14%	21%
5	14%	13%	17%	17%	13%	22%	21%	17%	29%
7.5	14%	13%	16%	17%	14%	22%	21%	17%	28%
10	14%	12%	16%	17%	14%	21%	21%	17%	23%
12.5	15%	13%	17%	18%	15%	27%	23%	17%	31%
15	15%	13%	17%	18%	14%	32%	20%	17%	22%
17.5	15%	13%	16%	18%	13%	26%	23%	17%	30%
20	15%	13%	18%	18%	14%	23%	21%	16%	23%
22.5	15%	13%	20%	18%	14%	22%	20%	15%	23%
25	14%	13%	16%	17%	14%	22%	21%	15%	24%
27.5	17%	15%	20%	17%	12%	23%	18%	14%	22%

7.7 *Light Loss Linear Deviation*

The linear light loss of each crystal in the detector was determined as described in 2.2 using a 4mm sliding window, and omitting the data within 1mm of the crystal edges. The differences were converted into keV using the calibrated energy response of the full crystal. The deviations for the central crystals ranged from 7 keV to 38 keV. The edge and corner crystals displayed much higher deviation ranging from 11keV to 70 keV for one of the corners. This was largely due to an array setting in which light escaping a central crystal is captured by adjacent crystals and directed to the photodetectors. However, light escaping an edge or corner crystal is lost completely. Photons originating at different depths in an individual crystal will undergo a different number of reflections and this difference in number of reflections changes the likelihood of escaping a particular crystal. The results from Chapter 2 indicate that the timing and energy resolution for these crystals should be expected to perform poorly. That is indeed the case for both energy, as shown in Table 8, and for timing, as will be demonstrated in Chapter 8.

Table 10 Maximum Light Loss Linear Deviation in keV for each of 64 crystals in a single detector module

35	24	9	15	14	15	26	70
42	38	30	33	33	29	34	58
11	15	13	9	7	9	7	31
17	19	13	7	7	8	7	19
18	21	17	10	11	10	8	22
16	20	17	11	11	12	12	23
15	16	13	10	9	9	10	25
31	23	12	12	14	17	20	50

7.8 *Background Counts Uniformity*

The corrected background flood map is segmented into crystals, and a cut of 400-1100 keV is applied. The number of counts in each pixel is taken as an initial determination of the detector uniformity. This uniformity will differ from detector uniformity as measured by external sources. This difference is due to the different distribution of scintillation events within the detector using intrinsic decay events, which will be uniformly distributed, and the distribution of scintillation events originating from an external source, which will be more concentrated on the side of the detector facing the source. This measurement, therefore, is not useful for correcting data acquired from external sources, such as in phantom or patient studies, but may be useful in determining whether the system is performing as expected. Notice the more frequent occurrences of scintillation events in the center of the detector module compared to the edges of the module.




Figure 87 Background Counts Uniformity

The trigger level for the SSPM pixels is independent of the originating location of the event. In these modules the edges tend to display a greater amount of light loss than the center of the array. The intrinsic counts extend to as low as 88 keV, while the effective cutoff at the edges is as high as 300 keV in some cases. Therefore, we expect the edges to display a lower count rate than the center, although this might be somewhat mitigated with a appropriate energy selection criteria. This is different from the energy blurring effects discussed in Chapter 3, which focused on count rate differences between different points within the same crystal. Additionally, events with a relatively low energy β , and in which the simultaneous γ travels a relatively long distance will tend to be weighted toward the center of the array. This miscalculation of the location of events will tend to suppress the count rate at the edge of the detector

module. Future work may investigate how to improve the application of these data to system monitoring.

7.9 Summary


The fundamental results of the modules investigated in this chapter showed excellent promise for high performance applications in which DOI information is necessary. The module and electronics performed well, with all crystals clearly distinguishable, and a peak-to-valley ratio of 10:1 or better in most cases. The flood map data are clear enough that no correction is necessary based on the depth of interaction or recorded event energy. The energy calibration performed well and gives a full detector energy resolution of better than 19% at 511 keV, with individual crystal energy calibration values ranging from 10 to 30%. The DOI resolution is acceptable at around 6mm, but it is not as good as expected based on single pixel measurements. It is suspected that the glue bonds on two sides are inhibiting the DOI resolution and further examination of the 12x12 array composed of 1.25x1.25x30mm individual crystals may show that having air gaps on all four sides allows DOI resolution within an array to be closer to that measured in single crystal experiments. Initial characterization of the 12x12 array have demonstrated DOI resolution of ~2-3mm in the interior region of the module. However, the DOI calibration method works well when compared to externally collimated measurements, in spite of the physical properties that would be expected to hamper these measurements in an array configuration. The two methods matched within 2 mm for all points in all crystals examined and generally were within 0.75 mm. The light level uniformity of the module was acceptable, although there is room for improvement at the edges and corners of the module. Some of these effects

may be mitigated by building larger panels and allowing light to be shared across different modules within a single panel.

Chapter 8 Timing Algorithms

8.1 Introduction

Due to the parallel readout technique used, each event has multiple time stamps. A typical histogram of the number of channels fired for each event in a dataset consisting of 2M events is shown in Figure 88. This shows we can expect to have 8-16 time stamps for each event in a coincidence pair. One critical question is how best to determine the final time stamp for an event with multiple time stamps.



Traditional PMT-based systems use either winner takes all or first over the threshold algorithms, while PSPMT's use a combined signal from the last dynode, which is in effect a non-normalized energy-weighted algorithm of the entire detector face. While some work has been presented at conferences examining a weighted average of 2 or 3

channels, no publications are available on the results. The vast majority of timing studies with new SSPM detectors utilize 1-1 coupling between the crystal and the photodetector (46, 128-133), and when light sharing is incorporated the method of timing is generally poor (134-136), or unreported, although one paper demonstrates 2.6 ns timing resolution in a charge and light shared environment (137), and some work demonstrating better than 1ns timing resolution has been performed(87, 138, 139). However, the specific methods used in the high performance studies are poorly suited to the scales necessary in a full PET detector system, and not all of the publications include pulse shape effects, which should improve timing resolution(133). Although the methods described for the sub nanosecond timing resolution are impractical for a full system due to the inclusion of digital oscilloscopes for pulse shape analysis, they do demonstrate the fastest time currently achieved in a light sharing environment with SSPM detectors and should be the end goal of light sharing detector systems. However, these high resolution timing measurements have been made with Hamamatsu's MPPC array which is known to have a superior timing resolution compared to SensL's current devices due to lower intrinsic noise properties. Thus, while the ultimate goal of a SSPM based module may be below 1ns, the goal for these modules is a much more modest 2.2ns coincidence resolving time, which is better than other results published using these photodetectors (137, 140) and comparable to the single crystal studies performed using CFD analysis shown in 2.3.5.

Single channel coupling is impractical for a high resolution system, since the crystal size would be limited to 3mm, the current standard SSPM pixel size. Our system shares light over many photodetectors, and so a careful investigation of how best to use the large amount of data is appropriate. Additionally, pulse shape corrections appear to

be crucial to improving the response of SSPM based systems. As discussed in Chapter 6, the event analysis and TWC calculation both include pulse shape effects in the calculation process. Additionally, the method used in this work does not make the assumption that the leading edge has a linear characteristic, which is an assumption of most pulse shape analysis methods. Previously published pulse shape methods measure several data points and fit a straight line to the data and extrapolate to the x-intercept (128, 139).

8.2 Coincidence Timing Method

8.2.1 Selected SSPM Pixel Timing Algorithms

The timing response may be calculated using single pixel approaches. In these approaches, a single pixel is chosen as the most accurate representation of the timing. These methods differ in the selection of the most accurate pixel and in whether further corrections are applied to the data collected from the selected pixel. Any of these algorithms could be implemented in the initial electronics, simplifying the system architecture and data structure. The methods without a TWC correction can be implemented at the level of the ASIC, while the methods applying a TWC correction would need to be implemented at the level of the FPGA, which would need to have the TWC stored on board.

First Uncorrected (First-UC)

A baseline for the investigation is the uncorrected time stamp from the first triggered channel. In this case the raw time difference between the first event on one module and the first event on the other module is determined. This may be considered as a

module level leading edge discriminator with only active channels considered. No time walk correction is applied to these data.

$$T_{first UC} = \min(T_i) \quad \text{Equation 8-1}$$

First Corrected (First-TWC)

The next algorithm is to take the first event in each module and apply the TWC to the data. The pixel selected in this algorithm is identical to the pixel selected in the Raw Electronic method. If multiple pixels fire with the same time stamp the pixel with the maximum energy is selected for analysis.

$$T_{first TWC} = TWC(\min(T_i)) \quad \text{Equation 8-2}$$

Corrected First (TWC-First)

The next algorithm is to apply the TWC to all values and then determine the earliest time stamp. In this case, the selection is made after the TWC is applied, whereas in the First Corrected method the pixel is chosen, and then the TWC is applied. The order of selection allows this method to consider low energy pixels which fire later in sequence, but have an earlier corrected T^* .

$$T_{TWC first} = \min(TWC(T_i)) \quad \text{Equation 8-3}$$

Maximum Energy Uncorrected (E Max)

The next algorithm is similar to a winner-takes-all approach. In this algorithm the raw time stamp from the channel with the greatest signal amplitude is taken as the best time estimate. If two channels collect the same charge, the channel that fires first is selected for analysis.

$$T_{E\ max} = T_i(i = \max(E_i)) \quad \text{Equation 8-4}$$

Maximum Energy Corrected (E Max-TWC)

The final single pixel algorithm again uses the Maximum Energy pixel, but applies the TWC correction to the selected data.

$$T_{E\ max} = TWC(T_i(i = \max(E_i))) \quad \text{Equation 8-5}$$

8.2.2 Multi-Pixel Average Timing algorithms

This section describes the methods used to combine the data from multiple channels. These methods require the timing information of all channels to be fully analyzed. These methods should improve upon the timing results of the single pixel methods. Since these methods cannot be implemented at the ASIC level, only methods incorporating a TWC correction are investigated.

Raw Average

The first average algorithm is a simple average of all the time stamps within an event after the Time Walk Correction has been applied. This is a baseline for the averages and is expected to perform poorly. This method assumes all TWC corrected time stamps are equally accurate measurements of the interaction time, even though the error analysis in 6.3.3 demonstrated that higher energy signals tend to have lower intrinsic error.

$$T_{avg} = \frac{1}{n} \sum_{i=1}^n T_i^* \quad \text{Equation 8-6}$$

Energy-Weighted Average

The next algorithm calculates a weighted average time stamp in which the weights are proportional to the total charge collected by the ASIC.

$$T_{E \text{ weight}} = \frac{\sum_{i=1}^n T_i * E_i}{\sum_{i=1}^n E_i} \quad \text{Equation 8-7}$$

Error-Weighted Average

The final algorithm takes advantage of the error estimate made in 6.3.3. This algorithm is nearly identical to the energy-weighted algorithm with the exception that the average is weighed by the error, rather than by the collected energy. The weighting is equal to the inverse of the square of the error estimate made in 6.3.3. If the error estimates are accurate, this method should produce the best overall timing, since error-weighted averages minimize the overall error of the final determination. However, since the error estimate cannot be independently verified it is not certain that this method will generate the best overall results. If this method does produce superior results to the other methods, it will provide one point of validation for the error estimate.

$$T_{error} = \frac{\sum_{i=1}^n T_i * w_i}{\sum_{i=1}^n w_i} \quad \text{Equation 8-8}$$

Channel Number Selection

The methods produced by averaging several signals may choose to include fewer than all signals in the evaluation. If the light is nearly evenly distributed across a few pixels, the best time may very well be a simple or weighted average. However, when one channel contains a miniscule portion of the total energy, or it lies well away from the time determined by the other channels, it may be best to exclude that channel from the evaluation. In order to measure this effect for a group with n signals, n separate averages were performed. The largest energy events tend to have less error in the time stamp, and cross the leading threshold trigger earlier than lower energy events. In the simple average algorithm, these general aspects of the data were used to guide the

method of choosing which events to group for calculations. The signals within an event are sorted according to the uncorrected time stamp, and averages are made of the first two channels to fire, the first three channels to fire, and so on until all channels are included. In the energy-weighted data, a basic assumption was made that the error is proportional to the signal level, and so in this algorithm events were sorted according to the total collected energy. The groups then include the two channels that collected the greatest charge, the three channels that collected the greatest charge, and so on until all channels are included. In the error weighted algorithm the error associated with the particular channel and energy are used for sorting and grouping. In this case, the two channels with the smallest error are averaged, the three channels with the smallest error, and so on until all channels are included.

While n averages were performed, the full range of the comparison will vary from 3 to 32 channels, where 3 is the minimum number of channels that can properly place an event, and 32 is the maximum number of channels available. However, as demonstrated in Figure 88, very few events record more than 16 channels, and it is likely that a 15 channel event may be in coincidence with an 8 channel event. For comparison purposes timing determinations were made for the full range of 3-32 channels fired, in which all values for channels greater than the total number of channels fired are equal to the calculation for the number of channels fired. For example, if 13 channels fire, the calculated time for averages of 13-32 channels are all equal to the time calculation for 13 channels.

Additionally, as discussed in 6.3.3, the TWC curve may be calculated using either energy-weighted or error weighted-values during the calculation or with a test pulse

modeled on an average of several nuclear test pulses. In order to compare the effects of different TWC calculation methods, we compare the timing for all methods using an energy weighted TWC algorithm, an error-weighted TWC algorithm, and a TWC curve measured using an electronic test pulse.

8.2.3 *Experimental Setup*

Two detector modules were placed ~20cm apart with a 1mm Na²² point source placed equidistant between the two modules. Eight million coincident pairs of events were recorded for evaluation. For single crystal analysis only events occurring within a particular crystal near the center of each module were selected for analysis. For full module analysis all data were selected for analysis.

8.3 *Coincidence Timing Results*



8.3.1 *Crystal to Crystal Coincidence Timing*

A single crystal was selected for investigation within one of the detector modules along with the crystal from the opposing module with which it was most frequently in coincidence. In order to reduce edge effects, central crystals were chosen. The coincidence time between the two crystals within the module were determined using each of the described methods. A histogram of the coincidence resolving time was determined for the First-UC, First-TWC, TWC-First, E Max, and E Max-TWC methods, and shown in Figure 89. A Gaussian function was fit to each resultant histogram in order to determine the CRT for each method. The uncorrected first time stamp and uncorrected maximum energy time stamp histograms are nearly identical as are the corrected first time stamp and the corrected maximum energy. This is the case due to

the very high likelihood that the first event is the maximum energy event, which is true for ~65% of events.

Taking the first after the TWC is applied generates poor performance due to the noise of low energy events. Low energy signals have larger error than high energy signals, and the timing method based on the largest energy signal never selects the low energy, high error measurement. Low energy events trigger the system later, and the First-UC and First-TWC methods do not select the later triggered events. However, once the TWC correction has been applied, the low energy events will be centered around the same overall calculated event time but have a larger error. The TWC-First method will select the earliest event after the TWC correction regardless of its error, and in a distribution the high error signals are more likely to occur farther from the mean time. This means the TWC-First will frequently choose a low energy, high error signal as the reference time and so generates a poorer overall CRT.

When the TWC was applied to the first or maximum energy signals, the CRT improved from 4.8ns to 3.2ns. Although the TWC correction is large for most of the signals, the earliest, highest energy signal within any event group is typically above $TQ=1000$. In this region, the TWC curve is relatively flat, and the correction factors are small compared to other ranges of the TWC.



The averages were calculated for groups of 2-32 channels, and a separate Gaussian was calculated for each pairing, giving 31 CRT measurements for each of the Simple Average, Energy Weighted Average, and Error Weighted Average methods. In all cases, separate calculations were made for each of the three methods of determining the TWC curve: electronically measured, error-weighted, and energy-weighted. The simple average is not shown due to its exceedingly poor performance ($>18\text{ns}$). The single channel methods are inferior to the multi-channel averages, and the multi-channel weighted averages improved with each additional channel included. There was a slight improvement of 10 channels compared to 32 channels, but the improvement was minor, and differs in both magnitude and location when different crystals were selected for analysis. The simple average degraded in performance with each additional channel. When comparing the weighted averages, the electronically measured TWC underperformed the scintillation measured TWC curves. The error-weighted average method was better than the energy-weighted average, and the error-weighted TWC performed better than the energy-weighted TWC, although the magnitude of the difference was small ($\sim 200\text{ps}$ for the weighting techniques and $\sim 50\text{ps}$ for the TWC techniques). The optimal performance ($\sim 2.4\text{ns}$) was found using an error-weighted TWC, and an error-weighted time calculation. This matched the expectation that a properly weighted average should generate a better reference time compared to other estimates of weighting. Additionally, the superior performance of the error-weighted algorithms to the energy-weighted algorithms gives some confidence in the method used to estimate the channel and energy dependent timing error. The largest source of error in the method lies not with the choice of TWC algorithm, but with the wide error bars for the distribution demonstrated in Figure 67.

8.3.2 Module-to-Module Coincidence Timing

The results from both selected single SSPM pixel algorithms and multipixel averages for the full module data set are shown below. The analysis is identical to that of the crystal to crystal coincidence timing analysis, but without the single crystal selection criteria.





Figure 91 Full Module Selected SSPM Pixel Timing Algorithms CRT Spectra



At the module-to-module level, the multi-channel weighted averages produced a narrower CRT than the single pixel algorithms. In the weighted averages, the CRT narrowed further with the inclusion of each additional channel, with only a very slight dip at around 12 channels, while the simple average again degraded with each additional channel included. The error-weighted average generated superior CRT compared to the energy-weighted average, and the error-weighted TWC calculation was superior to either the energy-weighted or electronically measured TWC curves (~200ps difference with weighting, ~50ps difference with TWC method). Once again, the narrowest overall CRT was obtained using an error-weighted average and an error-weighted TWC calculation (~2.9 ns). This gives further confidence in the method used to generate the error estimate during the scintillation event based TWC calculation.

The overall CRT was slightly worse at the module level (2.9ns) compared to the crystal level (2.4ns). This is mostly due to the higher accuracy of the single crystal data, which is restricted to events occurring near the center of each module. That region has the greatest light collection efficiency, and so is expected to have the best timing resolution. The module level data includes events along the perimeters and corners which have poorer light collection efficiency, and so are expected to have poorer timing resolution.

8.4 *Event Level Timing Error Analysis Method*

Each event within a module has several time stamps, each of which includes its own error estimate. These error estimates enable a calculation of the error of the error weighted mean used for coincidence timing measurements. Additionally, the variance of the distribution of signals for the SSPM pixels triggered within a single scintillation event may also be calculated. The error of the weighted mean allows for an event level determination of the timing error, rather than either a system level or line-of-response level error analysis. The variance of the time stamps may allow an improved event selection by rejecting events that have an appropriate energy and timing, but for which one of the timing stamps lies too far away from the weighted average time. These calculations are only possible for techniques with multiple time stamps, and so they are only used in the average methods.

The techniques used for determining the error of the mean and the variance within a group of data, each with its own error estimate, are standard statistical techniques(141).

$$V_1 = \sum_{i=1}^N w_i \quad \text{Equation 8-9}$$

$$V_2 = \sum_{i=1}^N w_i^2 \quad \text{Equation 8-10}$$

$$s^2 = \frac{V_1}{V_1^2 - V_2} \sum_{i=1}^n w_i (T_i^* - T_{ref})^2 \quad \text{Equation 8-11}$$

$$\sigma_{T_{ref}}^2 = \frac{1}{V_2} \quad \text{Equation 8-12}$$

$$ERT = \sqrt{\frac{\sum \sigma_{T_{ref}}^2}{m}} \quad \text{Equation 8-13}$$

$$CRT = \sqrt{ERT_1^2 + ERT_2^2} \quad \text{Equation 8-14}$$

In these calculations σ_i^2 is the estimated variance of each of the individual corrected times T^* , w_i is the inverse of σ_i , s is the variance of the distribution of corrected times T^* for a given event, $\sigma_{T_{ref}}^2$ is the variance of the weighted mean, ERT is the expected resolving time for a given group of events, and CRT is the predicted resolving time between two distinct groups of events (typically corresponding to different crystals or arrays). When the weights are equal and the size of the distribution is large, Equation 8-12 simplifies to the standard error of the mean. ERT is the single detector expected resolving time for the distribution of events included in the dataset. If the event level coincident timing errors are normally distributed, the CRT may be calculated by a quadratic addition of the individual event level coincident timing errors for all events with the dataset. However, in our case the data are not normally distributed, and outliers are included from random events. It is possible for a second gamma to interact in one of the modules during the coincidence timing window, causing a detection of three events. These events are not normally distributed and should be rejected. Currently these


events are rejected by calculating the CRT using the central 95% region of the dataset. A further improvement could be possible by using the event level variance calculation and rejecting individual signals more than three standard deviations outside the weighted average time. The event level time and event level error could then be recalculated and the data restricted to signals generated by the true 511 keV coincident gamma rays.

The calculated event level timing error will allow the generation of reconstruction algorithms that may reconstruct each LOR only over the region appropriate given its own timing estimate rather than applying a global timing resolution to all LORs. In non TOF systems this information will have no effect since the line will be distributed over the entire reconstructed image. However, in TOF systems this will allow low error event pairs to be distributed over a narrow range and high error pairs to be distributed over a broader range in the reconstruction. It is unclear what effect this approach will have on the final reconstructed image, but it may improve the image by giving the high timing resolution data a greater weight in the high frequency domain of the reconstructed image and effectively blurring the low timing resolution data. While some current systems may be able to estimate an expected timing resolution for a given LOR, not all events along that LOR will have the same intrinsic timing resolution. Specifically, two events originating at the same point and traveling along the same LOR may have different expected timing resolutions if the interaction point within the detector module changes or the number of Compton interactions is different for the two event pairs. The described method can distinguish between two events along the same LOR with different intrinsic timing errors.

Once again, there is no standard with which to compare the event level error estimates. However, the event-level timing errors may be used to calculate an expected CRT for a given data set. If the predicted and measured coincidence resolving times are not equal for a given data set, then the prediction fails. While this comparison is not a robust proof of the method, it may give a high level of confidence in the error estimate depending on the level of agreement between the predicted and measured timing resolutions.

8.5 Event Level Timing Error Analysis Results

8.5.1 Crystal to Crystal Timing Error Results



The calculated CRT for the single crystal dataset was 2.36ns, the CRT for the measured spectrum was 2.41ns, and the CRT for the predicted spectrum was 2.37ns.

These measures of the CRT were very close, and indicate that this method of determining the crystal level CRT based on error estimates has validity.

The method was further validated at the single crystal level by restricting the region in DOI space to the first 5mm, the second 5mm, the third 5mm, and the final 15mm of a pair of crystals in the centers, a pair of crystals at the edges, and a pair of crystals at the corners of two modules. The edge comparison from 10-15mm depths includes only 62 counts. If that low count group is omitted the measured and calculated resolving times had less than a 14% difference in all cases.


Table 11 Comparison of Measured and Calculated Coincidence Resolving Time for several crystal pairs with DOI restricted data sets.

Crystal Pair	DOI Range	Measured CRT (ns)	Calculated CRT (ns)	Difference (ns)	Difference (%)	Counts
Center	0-5 mm	2.4	2.4	-0.02	0.7%	2247
	5-10 mm	2.2	2.3	0.11	-4.6%	584
	10-15 mm	2.3	2.4	0.10	-4.2%	267
	15-30 mm	2.5	2.5	-0.03	1.1%	885
Edge	0-5 mm	2.7	2.6	-0.12	4.7%	458
	5-10 mm	2.4	2.6	0.21	-8.3%	154
	10-15 mm	2.3	2.7	0.42	-17.0%	62
	15-30 mm	2.8	2.7	-0.15	5.4%	155
Corner	0-5 mm	3.0	3.2	0.22	-7.1%	981
	5-10 mm	3.0	3.3	0.38	-12.0%	419
	10-15 mm	3.2	3.7	0.47	-13.5%	305
	15-30 mm	2.9	3.1	0.19	-6.1%	422

8.5.2 Module to Module Timing Error Results

The calculated CRT for the full module dataset was 2.52 ns, the CRT for the measured spectrum was 2.99ns, and the CRT for the predicted spectrum was 2.67ns.

The residual difference between the measured CRT and calculated CRT is fairly small at 320ps. This method has not fully investigated the effects of timing offset due to crystal identification, which may account for a portion of the residual difference between the predicted and observed module level CRT. Overall it is unclear why the module level data has a greater difference between the measured and predicted CRT spectra than the single crystal data.



8.5.3 Full Module Expected Resolving Time

Once the Expected Resolving Time (ERT) has been verified the method may be used to determine the ERT at each depth within each crystal in a full detector module. This may then be used to examine different crystals to determine areas of the module

which have the poorest performance and should be treated most carefully in updated module designs.

Table 12 shows the average ERT across the full crystal length for each of the 64 crystals in a detector module. The crystal level ERT varied from 1.77 to 2.70 ns among different crystals in the detector module. The average ERT for the interior crystals was 1.9ns while the average ERT was 2.1ns for the perimeter crystals and 2.4 ns for the corner crystals.

The data selected for analysis may then be further restricted based on the calculated DOI of each event and a new ERT calculated for each depth. In this way it is possible to investigate DOI effects in the timing resolution. The difference between the minimum and maximum ERT is show in Table 13. The difference between the ERT varied from less than 200ps for most of the interior crystals to greater than 1ns for one of the corner crystals. This is consistent with the large linear light loss measured in this crystal, as demonstrated in Table 10, and the correlation between timing degradation and linear light loss found in single crystal measurements detailed in Chapter 2. This would indicate that any future module designs should make an effort to improve the light collection efficiency of the edge and corner crystals in order to improve the overall timing resolution.

Table 12 Expected Resolving Time (ns) of all crystals in an 8x8 array. Edge and corner crystals demonstrate a degradation of timing resolution compared to the center of the array


2.40	2.63	2.07	1.96	1.90	1.95	2.48	2.70
2.12	2.21	1.90	1.76	1.77	1.88	2.20	2.28
1.98	2.10	1.87	1.84	1.83	1.87	2.05	2.06
1.89	1.97	1.86	1.83	1.82	1.86	1.96	1.88
1.93	2.04	1.87	1.85	1.87	1.92	2.00	1.93
1.97	2.13	1.91	1.85	1.90	1.96	2.07	1.91
1.97	2.22	1.95	1.80	1.90	1.98	2.11	1.95
2.15	2.26	1.91	1.90	1.94	1.91	2.07	2.18

Table 13 Difference between maximum and minimum ERT (ns) among four mm segments of each crystal. Notice the greatest difference is again at the edges and corners with remarkable uniformity near the center of the module

0.79	0.74	0.39	0.34	0.26	0.22	0.65	1.11
0.60	0.36	0.23	0.13	0.15	0.22	0.32	0.69
0.25	0.15	0.13	0.11	0.09	0.09	0.19	0.23
0.28	0.15	0.20	0.24	0.17	0.11	0.19	0.15
0.20	0.05	0.04	0.03	0.05	0.05	0.12	0.10
0.28	0.07	0.10	0.04	0.04	0.03	0.14	0.12
0.42	0.08	0.09	0.16	0.10	0.11	0.09	0.44
0.54	0.39	0.26	0.26	0.25	0.25	0.45	0.57

Table 14 Expected Resolving Time in ns for groups of crystals at different depths within the detector module

Depth	Central Crystals			Non-Corner Edge Crystals			Corner Crystals		
	Avg	Min	Max	Avg	Min	Max	Avg	Min	Max
2.5	1.89	1.60	2.30	1.87	1.73	2.37	2.01	1.88	2.28
5	1.88	1.61	2.26	1.89	1.73	2.41	2.05	1.89	2.32
7.5	1.90	1.63	2.27	1.93	1.72	2.50	2.12	1.94	2.40
10	1.92	1.64	2.28	1.98	1.72	2.57	2.23	2.01	2.50
12.5	1.94	1.65	2.28	2.03	1.72	2.64	2.35	2.12	2.58
15	1.95	1.67	2.27	2.06	1.72	2.66	2.42	2.20	2.64
17.5	1.95	1.67	2.24	2.06	1.72	2.65	2.40	2.18	2.67
20	1.94	1.65	2.21	2.02	1.71	2.55	2.29	2.08	2.59
22.5	1.91	1.58	2.17	1.97	1.71	2.42	2.14	1.96	2.44
25	1.89	1.57	2.13	1.92	1.71	2.28	2.04	1.87	2.31
27.5	1.92	1.61	2.17	1.88	1.73	2.14	1.98	1.79	2.32



8.6 *Summary and Discussion*

The timing properties of the arrays are adequate for PET systems. The timing resolution of the module at 2.99ns showed some degradation from single crystal results of 2.46ns. The method used to determine the TWC curve made a statistically significant but small difference in the coincidence resolving time of the system. For a non-TOF system, the differences likely would not generate differences in reconstructed images, but the differences may be much more meaningful in a TOF system. However, due to the similar time necessary to determine the TWC for each of the three methods, there is no intrinsic tradeoff involved with using the error weighted algorithm, and that method should be adopted.

The timing resolution is currently limited by the characteristics of the SSPM array, and as higher performance arrays are available the overall performance of the system is expected to improve dramatically. Most importantly for this study, the timing resolution found for single crystal readout shown both in this chapter and also in Chapter 2 demonstrates that the system is capable of matching the timing resolution of single crystal studies. As the system studied in this work shifts from the current generation Array4 and ArraySL devices to higher performance devices from Hamamatsu, SensL, or AdvanSID, these timing techniques will need to be applied in order to optimize the performance of the system.

The method used to predict the timing resolution is novel and appears to be accurate. Consequently, its potential impact for TOF reconstruction should be explored in system simulation studies. New reconstruction techniques are needed that incorporate error analysis in list mode reconstruction in order to broaden or narrow the

projected region of the LOR for each coincident pair. Currently there is no way to prospectively determine the effect this information can contribute to a reconstructed image.

Investigation is necessary to determine how to match the timing performance of the full module with that of the individual crystals within the array. Some of the difference in performance of the crystal-level CRT compared to the module-level CRT is due to the selection of central crystals for the crystal studies. The central crystals have the best light detection properties, and so some improvement is expected based on increased light detection efficiency. Additional work is still needed to improve the light detection efficiency at the borders of the array. However, even including this effect the performance of the module-level CRT is only predicted to degrade by ~200ps rather than the 550ps observed. While the difference is not critical at the level of 2.5-3.0ns observed with the current detector modules, it would be critical at the level of 400-900 ps more typical of time-of-flight systems.

Finally, the event-level variance data are not currently being used. A broad detector-level timing window is necessary due to the relatively long peaking time of the SSPM. In the current configuration, detected events sometimes include two separate scintillation events with interaction times separated by 75ns. An event-level variance calculation could allow the identification and removal of these event groups. In the current system the rate of occurrence is low due to the size and low level of activity that may be introduced into the field of view. However, a system dedicated to breast imaging will not only include more activity within the field of view, but also include a large random rate from events originating outside of the field of view. The ability to

remove event groups that include overlapping random interactions would be beneficial to such a system.

Chapter 9 System Performance

9.1 Introduction


Any potential PET improvements are ultimately evaluated by system performance and image evaluation. If an alteration in module design or system design improves a metric, such as timing resolution, energy resolution, or crystal separation, without affecting the reconstructed image, the alteration may be interesting, but it is not helpful. The metric that we are introducing is DOI information, which is introduced to improve the resolution uniformity and allow systems in which the useful field of view extends close to the detector face. The most direct measurements to test the usefulness of this additional parameter are the spatial resolution determined from reconstructed images of a points source placed at various points within the FOV and a comparison of images with and without DOI information.

9.2 Procedure

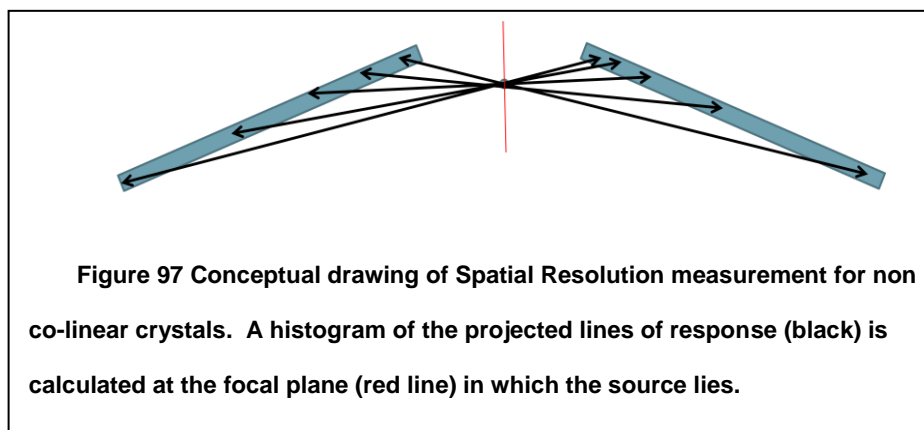
9.2.1 Spatial Resolution

The means of determining the spatial resolution of standard PET detectors is well understood(142), but there is no standard experimental setup for determining the spatial resolution of a DOI-capable PET detector module. One method introduced by Yang was to reconstruct an image with two detectors rotated with respect to one another (39). While his method demonstrates the importance of DOI information, it does not produce a metric that may be used to compare two detectors. The method used in this work was first proposed by Dokhale in 2006 (143). In this work we measure a reference spatial resolution by aligning two detector modules at 180° and measuring the number

of coincidence counts as a Na^{22} point source is advanced between the two detectors, in the same manner as generally used for spatial resolution measurements. We then rotate each of the detectors by 22.5° around the center of the two modules, and advance the source along the midline of an isosceles trapezoid with a detector module along each side. The same experiment is repeated with each of the detectors rotated by an additional 22.5° . This generates coincidence data between two modules with opening angles of 0° , 45° , and 90° . A diagram of the experimental setup is shown in Figure 96. The data are collected across the full range of the detector face, although in a cylindrical design this may include certain angles outside the central bore. These angles are included because alternate system geometries, such as a larger bore, or box shaped systems, will allow lines of response in this region.



The modules are fully calibrated as described in Chapter 6. In the data set collected at 180° , the data were processed in the standard method, in which the number of counts in a selected pair of opposing coincident crystals was determined as the source was advanced. In the remaining data sets the spatial resolution is defined as the physical distance that must be traversed by a point source such that the FWHM of the distributions of the profile of counts in the focal plane of the two acquisitions do not overlap, shown in Figure 97. This definition is similar to the definition used for DOI resolution in Chapter 2. This method is only used for the non-collinear opening angles, since the DOI information contributes no information for the 180° angle, and all lines of response would be tracked to the same location during the backprojection. In these experiments a single crystal is selected in each module for analysis, thereby limiting the data to events that interact only in the selected crystals. This is not a necessary requirement, since the projection can be performed for any pair of events. While future studies may analyze full module response, this restriction of the data set is done so that the method will more closely match what is currently used in non-DOI PET systems, which determine the spatial resolution based on pairs of coincidence crystals rather than full module calculations.



The calculation is verified by plotting histograms of the calculated line densities when the source is at several locations. The distance necessary to separate two source positions should be equivalent to the average of the FWHM of the data at the two positions. This is shown in Figure 100 for detectors rotated with a 90° opening angle, and Figure 99 for detectors with a 135° opening angle.

9.2.2 Point Source Measurements

In a full bore system parallax errors degrade the image away from the central axis of the scanner. One method used to determine the improvement in image quality is to compare the resolution uniformity across the full field of view of a PET system when DOI information is included, to that of a reconstruction of the same data without including DOI information. A prototype system consisting of two panels separated by 106.6mm each with four detector modules assembled in a 1x4 array was built. A Na²² source was advanced over 35 axial locations and 60 radial locations covering 17mm axially, and 29.5mm radially. The measured area was smaller than the useful field-of-view due to saturation effects when the source was placed closer than 23mm from the face of the detector modules. Data were collected at 18 rotation angles with 20° between rotation angles. Data were collected for 8s at each angle for a total number of counts of ~100k at each angle. The reconstruction was done filtered back projection with a ramp filter applied, both with and without DOI information.

9.2.3 Phantom

An Ultra-micro phantom (Data Spectrum Corporation) was filled with 40 µCi of F¹⁸ and imaged for two minutes at each of eighteen rotation angles separate by 20°. The system was set in a configuration with two banks of four detectors each, with each

detector bank covering 68x16mm. The banks were separated by 47mm. The reconstruction method was OSEM with 40 subsets, 1 iteration, and a 1.0mm resolution model. Images were reconstructed both with and without DOI information in order to compare the results.

9.3 Results

9.3.1 Spatial Resolution

As shown in Figure 98, the spatial resolution at 180° was 1.6mm, which compares favorably with the crystal pitch of 1.9mm. The measured spatial resolution at 135° was ~2mm, as shown in Figure 99. The measured spatial resolution at 90° was ~3mm, as shown in Figure 100. The three histograms are from data sets with the source advanced by 2.0mm. In Figure 101 the calculated FWHM of the projected source profile is plotted for all relevant source positions. Notice that at the ends of the data collection range the calculated spatial resolution improves due to edge effects.




Figure 98 Spatial Resolution for detectors at 180°. FWHM=1.6mm
Collected by advancing 1mm Na²² point source by 0.1mm increments.




Figure 99 Spatial Resolution for detectors at 135°. FWHM~2mm. Data
sets are collected with source advanced by 2mm increments




Figure 100 Spatial Resolution for detectors at 90°. FWHM~3mm. Data sets are collected with source advanced by 2mm increments.





Figure 101 Spatial Resolution for rotated detectors with 1mm Na²² source advanced by 0.1 mm increments

9.3.2 Point reconstruction & resolution

When DOI information is included the resolution of a point source varies from 2-3mm radially as the source is moved from the center of the field of view to 29.5mm off axis. The tangential resolution improves marginally, but was approximately 2.0-2.1mm for all locations. The axial resolution remained 1.6-1.8mm throughout all locations as well. Without DOI information the radial resolution varied between 2.5 to 3.5mm, the tangential resolution varied between 2.8 to 2.1mm, and the axial resolution remained between 1.6 and 2.0mm.






Figure 103 Tangential resolution for various radial point source positions





Figure 104 Radial resolution for various radial point source positions

The measurements as the point source was advanced axially were much more constant. When DOI was included the two measures of resolution across the axis varied from 1.8 to 2.0 and 2.0 to 2.25mm. The axial resolution showed some regular structure with a spatial frequency equal to the crystal size. The axial resolution varied between less than 1mm to 2.0mm. When DOI information was not included all three measures degraded slightly.






Figure 106 Tangential resolution for various axial point source positions




Figure 107 Radial resolution for various axial point source positions

9.3.3 Phantom

A total of 9.8 M counts were used in the reconstruction of this image. The 1.7mm rods are clearly separable in the DOI reconstructed image, and most of the 1.35mm rods are separable as well. The non-DOI image shows severe artifacts, and no rods are visible in the reconstructed image. Each of the six sections of the phantom is distinguishable, but each section appears nearly identical to the other sections composed of different sizes of diameter rods.




Figure 108 Derenzo Phantom image reconstruction with and without DOI information included. Rod sizes are 0.75mm, 1.0mm, 1.35mm, 1.70mm, 2.0mm, and 2.4mm.

9.4 Conclusions

Depth of interaction information limits the degradation of the spatial resolution of the detector modules to ~3mm at a 90° opening angle, whereas without DOI information the spatial resolution would increase to the entire 21mm crosssectional profile. The spatial

resolution at 135° is better than at 90° , and it may be that the spatial resolution at small angles near 180° may improve beyond the nominal 1.6mm spatial resolution at 180° . This may also be true for non-collinear parallel crystals, and so further studies of the full module spatial resolution are warranted at more angles. If the spatial resolution were calculated for the full module rather than individual crystals, the system could use the defined method at 180° due to the large number of crystals that will be non-collinear. However, a full module spatial resolution calculation will depend heavily on the module separation, and careful work will be needed to determine the best method of characterizing the full system spatial resolution.

The point source resolution remained relatively uniform across the entire field of view when DOI was included, demonstrating an improvement in axial, tangential, and radial resolution compared to non-DOI data. These effects are expected to be more significant as the source is moved closer to the detector face. The phantom clearly shows 1.7mm rods, with some separation of the 1.35mm visible as well. This compares favorably to the 1.6mm spatial resolution of the modules at 180° , and suggests that the spatial resolution at very large opening angles (145° - 180°) may be better than initial measurements suggest. We expect that the overall results may improve further with an updated calibration of the system.

The system designed based on these modules is well suited for small animal imaging due to the high spatial resolution and high sensitivity. The modules should also be suitable to dedicated systems for breast and neurological imaging with limited modifications.

Chapter 10 Summary and Discussion

10.1 Single Pixels

Single pixel studies are useful for initial investigation and developing an intuition of how devices behave, but data gathered from single crystal studies cannot fully predict the behavior of full arrays. The complete loss of DOI information in Array 1 came as a surprise, and it would have been impossible to predict from single pixel experiments. Even if ESR had been used in single pixel experiments, the light escape would have generated a DOI gradient similar to that produced in crystal and reflector combinations that generate light loss. The results of the single pixel experiments likely would have suggested that DOI information is generated regardless of the surface finish. Additionally, single pixel studies allow a simple coupling between the crystals and the photodetectors, reducing any complications of combining time stamps between multiple pixels.

However, it was useful to show in single pixel experiments that the DOI response function calculated based on intrinsic events corresponded to within 1 mm to a DOI response function measured using an electronically collimated source regardless of the performance of the crystal in terms of energy and timing resolution. Also, the similar energy and timing performance of all crystals regardless of surface roughness is useful in deciding array properties. The development of the Linear Light Loss Deviation parameter may prove useful in determining which crystals within an array have enough light loss to warrant depth dependent corrections, and what those corrections need to be. Finally, the results from single crystal arrays indicate that improvements in array performance may be expected if superior methods of array preparation are developed

that mitigate light loss along the axial length of long narrow crystals. If surface absorptions are avoided, the timing and energy resolution of high-resolution, high-sensitivity PET scintillation arrays may improve as well.

10.2 Detector Array Results

The improvement in DOI resolution for edge crystals in an array was unexpected, although it is easily explained. At the edge, light that is transmitted through the ESR reflector is lost, and that light loss produces a second means of generating a DOI response function gradient. In the center there is no measurable light loss, and so only diffusion at the surface contributes to the DOI response function gradient. Edge effects generally degrade detector performance, so the improvement in the DOI was not anticipated.

Careful consideration should always be given to proper light guide design if there is a desire to match the physical scintillator area to the physical detector area. The concepts of light guide design are relatively simple and, while requiring great care, do not require advanced knowledge of optical transport properties or simulation of surface phenomena. However, the basic assumptions used in this work to guide light guide design may lead to light guide design that include a trapezoid with an entrance window far larger in area than the exit window. These high aspect ratio trapezoids should be avoided, since they will have a tendency to redirect light back into the array and away from the photodetector. The development of procedures to produce light guides through subsurface laser etching would allow a much more systematic study of several potential light guides rather than having to decide prospectively which light guide design would perform the best. This approach is not currently feasible due to the labor and

cost required to fabricate each light guide. The current light guide for the 12x12 array is likely suboptimal and alternative geometries should be investigated.

When developing DOI-capable detectors, air gaps combined with a non-absorptive reflector is the best method of generating DOI gradients, since this method is capable of producing excellent crystal separation, energy resolution, and DOI resolution. Absorptive reflectors such as Lumirror should be avoided, since the diffusion properties of the reflector are not needed as long as an air gap is included in the scintillation array.

The DOI resolution of the detector improves with the number of air gaps present, although there are tradeoffs between the number of air gaps and mechanical stability. The degradation in the crystal identification ability of the 12x12 array along two of the edges is unacceptable and makes these modules inappropriate for PET systems. One method to improve the performance of this module would be to build the modules into arrays where light sharing across modules is facilitated. The current performance degradation is due to light loss through the edge of the module. If the escaping light were to be captured by an adjacent module, the performance along the edge of each module would be enhanced. This improvement may allow a more uniform response of the module performance both in crystal separation and in crystal-level energy resolution. Currently light sharing between modules in the 12x12 array is impossible due to the external glass clamp, so any system designed based on the 12x12 arrays should incorporate a different method to achieve mechanical stability. Additionally the current technique of encasing the scintillation array in a glass clamp limits the overall system sensitivity, since a significant portion of the detector volume is composed of non-scintillating glass. If a glass clamp were to be used in full systems, those systems

should be composed of larger panels, each with a scintillation array equal in size to the full detector panel, rather than the smaller arrays of the current modular design. Alternatively, subsurface laser etching applied to the scintillation array has promise to improve the DOI resolution, provide necessary mechanical stability, simplify production, and substantially lower the cost of the system. Subsurface laser etching may allow a large panel to be composed of replaceable detector modules, while still allowing light sharing across modules.

10.3 Calibration

10.3.1 Current Calibration

The calibration method proposed is both computationally complicated and computationally time consuming. However, the mechanical setup is simple, does not require the accurate placement of any sources, and is repeatable. The developed calibration method could be run overnight in a fully developed system. The use of one or two external sources along with the internal radiation of the crystal material is sufficient to provide all the information needed for accurate timing, energy, and DOI calibration of each crystal. The DOI calibration without an external source is a necessary requirement of DOI-capable systems, and the ability to fully calibrate the DOI response functions with less than 20 minutes of data allows for a simple calibration that could be part of a daily system startup. The ability to calibrate the timing effects without an electronic pulse is a major advancement for accurate timing calibrations in systems that may have slight differences among thousands of elements.

It should be noted that this entire calibration is limited to the module level and does not attempt to calibrate system level effects, such as calculating a system matrix or

system uniformity. The calibration, insofar as it is implemented, is useful and adaptable to any system configuration incorporating DOI-capable detector modules, but a full system calibration will need to include system level effects.

Electronics Energy Linearity Calibration

Currently the electronics linearity must be measured separately from the crystal-level linearity. The current method uses a model pulse varied over two orders of magnitude, and measures the system's response to that pulse. This work has suggested that the electronics energy response be calibrated using a pulse sent from the detector level FPGA, removing the need for an externally modeled pulse. However, a pulse from the FPGA cannot account for the gain uniformities of the individual SSPM elements. In the future, an iterative approach may be developed that performs some combination of optimizing the energy resolution of all 64 crystals and balancing the R-space distribution of all 64 crystals by varying the nominal gain of each of the 32 SSPM pixel in a detector module, after all other electronic corrections have been determined using a direct measurement from the FPGA. Alternatively, if the detector performance improves to the point where individual thermal events may be measured, the SSPM pixel gain could be measured directly by determining the difference in collected charge between different numbers of thermal electrons. This measure would be capable of including the entire electronic chain simultaneously and thereby generating a single gain correction that includes the SSPM, electronics, ASIC, and FPGA. If the SSPM pixel level gain were known, the system could be improved by correcting this parameter, which would affect energy resolution, crystal identification, and DOI resolution, but should have no effect on the timing performance of the system. Additionally, if a pulse from the FPGA were passed to the electronics in order to

calibrate the electronics gain, the effect of the SSPM bias on noise characteristics could be measured as well. With this additional information it would be possible to tune the system to the optimal performance by adjusting the bias of each SSPM pixel individually in order to maximize its signal-to-noise ratio.

Crystal Identification Calibration

The current crystal identification method requires the user to locate the center of each crystal in a flood source image. This is time consuming in a system with 512 crystals and would be unacceptable in a full system. Watershed techniques have been used previously to determine crystal boundaries but frequently require user correction. It may be possible to adapt other methods from astronomy used to identify stars or from topographical mapping algorithms in order to better automate crystal selection. This has not been done in this work, but it should be performed on any system designed to be recalibrated frequently. It is the opinion of the author that topographical techniques such as watershed algorithms may be adapted with *a priori* knowledge of the detector module. Specifically, a final algorithm should include the knowledge of the number of crystals, the knowledge that the number of counts within each crystal should be similar, and a penalty function that encourages the algorithm to distribute the crystal boundaries uniformly throughout the entire space of the collected data.

LSO Crystal Energy Linearity Calibration

The crystal-level energy linearity calibration allows a crystal-dependent energy calibration over the range of 202-1275 keV. The described method includes 202 and 307 keV gamma rays and 1198 keV beta interactions from the decay of Lutetium, extending the calibration range to the lowest energies used in PET reconstruction. This

method minimizes the residual between the measured data and a fit at each of the energies. This corrects for the statistical uncertainties of each photopeak and for the different statistical relevance of the 1198 keV beta tail. The hyperbolic fit of the 1198 keV beta tail endpoint allows a good fit of the data with a relatively large movement of the calculated endpoint. This is appropriate due to the low count statistics at the endpoint of the beta tail. The developed method works well across all 512 crystals with zero failures when the modules are performing as expected, but failures become common as other performance properties degrade. If the 511 keV gamma ray source is changed from Na^{22} to Ge^{68} , it may be possible to omit the Cs^{137} data acquisition. However, in the opinion of the author it is beneficial to include the 1275 keV gamma ray from Na^{22} .

Depth of Interaction Response Function

The DOI response function calculation has been validated on both individual crystals and full detector modules. The method performs best when the crystal-level energy linearity has been determined and an energy cutoff applied. The method matched within 1mm of an externally measured DOI response function. The function can be calculated using about 20 minutes' acquisition of intrinsic decay events.

Time Walk Correction and Timing Error Calibration

The investigation of timing algorithms demonstrated that an appropriate TWC determination and timing algorithm has the potential to significantly improve timing. The SSPM arrays used in this work require a parallel readout in order to reduce errors due to thermal dark count noise. The algorithms investigate how best to determine timing in a system in which multiple timing measurements are made for each scintillation event.

While multi-anode PMT based systems do not have the requirement for parallel readout techniques, the techniques developed should be applicable to those systems as well. It is clear that the information gathered is useful, and all channels should be considered in any measurement of timing.

The timing response calibration for a TOF-capable system is critical. If new SSPM arrays are capable of TOF-level timing resolution, an error estimate of the timing may be beneficial. Most current TOF systems have one overall time resolution for the system that is applied to every line of response. However, with the multiple stamps and error weighting, this system is capable of providing not only an overall time stamp for measuring coincidence time, but an estimate of the error for that overall time stamp for each scintillation event in a coincidence pair. The current method determines the TWC response of each electronic channel including the SSPM pixel. Additionally, the described method calculates a channel and energy dependent timing error that allows a determination of the event level timing error. There are no standards by which to validate the developed method at the event level, and so expected values for large data sets were calculated and compared to measured values. The developed TWC and timing error response were accurate to within our ability to validate each. These newly developed techniques have the potential to improve TOF systems, but new reconstruction algorithms would be required.

10.4 Future Work

10.4.1 Array Design

The subsurface laser etching technique (89, 144) is an excellent technique and should be investigated further. This method could also be adapted to developing

inexpensive light guides, allowing a rapid comparison of a variety of designs in order to correct for effects such as reflections due to the tapered design of the light guide pixels. Additionally, a single light guide could have multiple layers of directed light sharing. Subsurface laser etching could also allow the light guide to be incorporated into the scintillation block and thus reduce the distance between the scintillator face and the edge of the field of view. Surface level laser engraving could further improve light guide design by etching channels in the light guide at the locations of the inactive areas of the SSPM array. This improvement would direct light away from the inactive areas of the SSPM module and improve overall light collection efficiency.

10.4.2 *Computational Investigations*

Spread Function

Several attempts have been made to develop detector modules that calculate the depth of interaction by using a single ended readout design and comparing the light distribution across the photodetector(71, 73, 145, 146). The same data are collected in our system and are available for analysis of light spread. Furthermore, the light spread at the two exit faces may be compared for a more complete measurement. If the DOI resolution based on light spread is comparable to the DOI resolution based on the signal ratios, the overall DOI resolution should improve further by including both the ratio and the light spread information in the DOI determination. If two independent measurements of the DOI are possible for each event, the DOI resolution of the module could be measured across the full module by comparing depths calculated by the two methods. This would allow the system calibration to include a DOI and DOI error for each event pair. This information could be useful if an event-level error analysis were

included in image reconstruction. With this information, the reconstruction could be updated from lines of response to tubes of response, with the tube radius varying based on the calculated error of the interaction location.

Compton Scatter Calculation

Interacting gamma rays may scatter within the module multiple times before depositing all their energy. The center of mass calculation of the event interaction will cause a miscalculation of the line of response, since the center of mass calculation will average light generated at multiple points within the detector module. One possible method to determine which events are Compton events is to determine whether they are mapped to a high or low count density region of the flood map. While this approach would work in the 8x8 module, the crystal separation in the 12x12 module is not sufficient to allow this distinction. An event that Compton scatters within a detector module can be expected to give a slightly different distribution of detectable light compared to a pure photoelectric interaction with the same amount of generated light and the same calculated position. The distribution of light may allow a determination of the likelihood of Compton scattering for a given event group, and if the light distribution has sufficiently accurate data it may be possible to determine where the interaction originally occurred. Even without knowledge of the interaction location, this information may allow image reconstruction methods that incorporate an error analysis including Compton likelihood.

Event Level Error Analysis Reconstruction

The error of the timing data may allow reconstruction algorithm that varies the region over which events are distributed during reconstruction at the event level.

Additionally, event-level DOI resolution data, along with Compton scatter probabilities, may allow further event-level error analysis. Events that are likely to have undergone Compton scatter, or occur in poor DOI resolution regions of the detector have intrinsically greater uncertainty than other events which, are likely to be photoelectric interactions or to occur in superior DOI resolution regions of the scintillation array. If this error analysis were included, the reconstruction could distribute the event probability not only along the line segment (timing uncertainty), but could also distribute the probability around the line segment (location uncertainty). The possibility of combining the Compton scatter, DOI resolution, and event level timing resolution for an event level error analysis is exciting and deserves further investigation in order to determine what improvements are possible in image reconstruction. The practical significance of such algorithms has not yet been demonstrated, but in the opinion of the author these data should allow reconstruction algorithms that produce far more uniform images with far less noise. At worst, an event level error treatment of the data will make no difference in the reconstructed image.

10.5 Alternate SSPMs

The module-level studies were primarily carried out using SensL's Array4. This detector is already obsolete and has been replaced by the ArraySL. SensL is currently bringing to market a more advanced SSPM with ~300ps timing resolution. The application of the developed methods to detector modules with improved photodetectors is crucial. These devices must be studied in electronic arrangements best suited to optimize their performance. As discussed in this work, our electronics setup was not optimized for the current detectors due to the manufacturing changes implemented at SensL between Array2 and Array4. The electronic timing accuracy of

the system including the electronic pulser, pre-amps, ASIC, and FPGA based TDCs was near 800ps, limiting the expected improvement from transitioning to newer, higher performance devices. The effect on TOF capable detector modules would require a new electronic setup in order to meet the timing requirements of advanced new photodetectors. It would also be critical to determine whether the timing calibration method proposed would perform similarly in higher performance environments.

Chapter 11 Conclusions

11.1 Single Crystal Studies

This work has investigated the effects of surface finish on the properties of single crystal that are relevant to PET applications. This work has found negligible effects of surface roughness on light loss, energy resolution, and timing resolution. The surface finish - machine polished, saw cut, or polished with lapping film - has a large effect on light loss, the effects of which carry through into energy resolution and timing resolution. The surface roughness does effect DOI resolution with rougher surfaces corresponding to better DOI resolution.

A method was developed for determining the DOI response function using internal scintillation events. The DOI response function using internal scintillation events, an external uniform flood source, or an externally collimated coincident source match within 1mm across the axial length of the test crystals regardless of crystal geometry, surface finish, or surface roughness. The degree of correlation between the methods is well within the DOI resolution at all cases, indicating that the method is appropriate for further investigation in modules.

11.2 Module Studies

Several detector modules were developed in this work. The front end electronics for these modules uses parallel readout with separate timing and energy information available for each SSPM pixel. The DOI characteristics of the arrays were dependent on the coupling method introduced between the scintillation crystal and the reflector material. Simple optical transport theories were adequate for designing light guides,

with no consideration of the effects of the tapered light guide pixels or transmittance through the ESR. The prototype module has energy resolution varying from 10-30% for individual crystals, with the vast majority of crystals having an energy resolution of 15% or better. The poorer energy resolution crystals correspond to edge crystals, which have a poorer light collection efficiency due to light loss at the boundaries of the arrays. The crystals are all well separated with a peak-to-valley ratio of $\sim 7:1$. The DOI resolution is crystal dependent, and ranges from 3-6mm for various crystals. Edge crystals have superior DOI resolution due to the combination of optical scatter and optical loss in these crystals.

A novel method of determining the time walk correction (TWC) has been developed using only scintillation events. The results of this method are similar to an electronically measured TWC. The scintillation based TWC generates a superior timing resolution compared to the electronically measured TWC. The method for determining the TWC using scintillation events also allows an estimate of the error in timing for each channel at each energy. These error estimates allow a calculation of the timing uncertainty in the system on an event level basis. The timing uncertainty was verified by comparing the predicted overall coincidence resolving time (CRT) to the directly measured CRT. The match between the expected and measured CRT was within 200ps.

The spatial resolution was measured at various opening angles for the detector module and found to vary based on the detector angle. The best spatial resolution was measured for collinear arrays at 1.6mm. At an opening angle of 90° the spatial resolution degraded to 3.5mm. The spatial resolution was measured for a pair of

detector crystals within the array, although an estimate from the full array may be developed for future comparisons to allow an identical method to be used at all angles.

A prototype system was constructed from the detector modules and used for point source and phantom studies. The point source was found to have a uniform reconstructed resolution throughout the FOV. The 1.35mm rods are visually separable in the reconstructed phantom image.

Bibliography

1. Bushberg, J., J. Seibert, E. Leidholdt, and J. Boone. 2002. *The Essential Physics of Medical Imaging*. Lippincott Williams & Wilkins, Philadelphia.
2. Couceiro, M., A. Blanco, N. C. Ferreira, R. Ferreira Marques, P. Fonte, and L. Lopes. 2007. RPC–PET: Status and Perspectives. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 580:915-918.
3. Park, S.-J., L. W. Rogers, S. Huh, H. Kagan, K. Honscheid, D. Burdette, E. Chesi, C. Lacasta, G. Llosa, M. Mikuz, A. Studen, P. Weilhammer, and N. H. Clinthorne. 2007. A Prototype of Very High-Resolution Small Animal PET Scanner Using Silicon Pad Detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 570:543-555.
4. Cherry, S., J. Sorenson, and M. Phelps. 2003. *Physics in Nuclear Medicine*. Saunders, Philadelphia.
5. Zimmerman, R., and M. King. 2006. *Nuclear Medicine*. In *Advances in Medical Physics*. A. Wolbarst, R. Zamenhof, and W. Hendee, editors. Medical Physics Publishing, Madison.
6. Rabinovici, G. D., A. J. Furst, J. P. O'Neil, C. A. Racine, E. C. Mormino, S. L. Baker, S. Chetty, P. Patel, T. A. Pagliaro, W. E. Klunk, C. A. Mathis, H. J. Rosen, B. L. Miller, and W. J. Jagust. 2007. ¹¹C-PIB PET Imaging in Alzheimer Disease and Frontotemporal Lobar Degeneration. *Neurology* 68:1205-1212.

7. Lortie, M., R. Beanlands, K. Yoshinaga, R. Klein, J. DaSilva, and R. deKemp. 2007. Quantification of Myocardial Blood Flow with ^{82}Rb Dynamic PET Imaging. *European Journal of Nuclear Medicine and Molecular Imaging* 34:1765-1774.
8. Niikura, N., C. M. Costelloe, J. E. Madewell, N. Hayashi, T.-K. Yu, J. Liu, S. L. Palla, Y. Tokuda, R. L. Theriault, G. N. Hortobagyi, and N. T. Ueno. 2011. FDG-PET/CT Compared with Conventional Imaging in the Detection of Distant Metastases of Primary Breast Cancer. *The Oncologist* 16:1111-1119.
9. Visser, E. P., J. A. Disselhorst, M. G. J. T. B. van Lier, P. Laverman, G. M. de Jong, W. J. G. Oyen, and O. C. Boerman. 2011. Characterization and Optimization of Image Quality as a Function of Reconstruction Algorithms and Parameter Settings in a Siemens Inveon Small-Animal PET Scanner Using the NEMA NU 4-2008 Standards. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 629:357-367.
10. Kolthammer, J. A., T. Jing, A. E. Perkins, and R. F. Muzic. 2010. Time-of-Flight Precision and PET Image Accuracy. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 3657-3660.
11. Surti, S., A. Kuhn, M. E. Werner, A. E. Perkins, J. Kolthammer, and J. S. Karp. 2007. Performance of Philips Gemini TF PET/CT Scanner with Special Consideration for Its Time-of-Flight Imaging Capabilities. *Journal of Nuclear Medicine* 48:471-480.
12. Clemencio, F. M. C., C. F. M. Loureiro, and C. M. B. A. Correia. 2007. The Design of a Digital Coincidence-Detection Trigger System for a RPC-Based PET. In *IEEE Nuclear Science Symposium Conference Record*. 827-830.

13. Clemencio, F. M. C., C. F. M. Loureiro, and J. Landeck. 2011. Online Trigger Processing for a Small-Animal RPC-PET Camera. *IEEE Transactions on Nuclear Science* 58:1766-1770.
14. Daube-Witherspoon, M. E., S. Surti, A. Perkins, C. C. M. Kyba, R. Wiener, M. E. Werner, R. Kulp, and J. S. Karp. 2010. The Imaging Performance of a LaBr 3 - Based PET Scanner. *Physics in Medicine and Biology* 55:45-64.
15. Pepin, C. M., P. Berard, A. L. Perrot, C. Pepin, D. Houde, R. Lecomte, C. L. Melcher, and H. Dautet. 2004. Properties of LYSO and Recent LSO Scintillators for Phoswich PET Detectors. *IEEE Transactions on Nuclear Science* 51:789-795.
16. Pepin, C. M., M. Bergeron, C. Thibaudeau, C. Bureau-Oxton, S. Shimizu, R. Fontaine, and R. Lecomte. 2010. Digital Identification of Fast Scintillators in Phoswich APD-Based Detectors. *IEEE Transactions on Nuclear Science* 57:1435-1440.
17. Szczesniak, T., M. Moszynski, M. Grodzicka, D. Wolski, L. Swiderski, M. Szawlowski, and M. Kapusta. 2010. Time Resolution of Scintillation Detectors Based on SiPM in Comparison to Photomultipliers. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 1728-1735.
18. Powolny, F., E. Auffray, S. E. Brunner, E. Garutti, M. Goettlich, H. Hillemanns, P. Jarron, P. Lecoq, T. Meyer, H. C. Schultz-Coulon, W. Shen, and M. C. S. Williams. 2011. Time-Based Readout of a Silicon Photomultiplier (SiPM) for Time of Flight Positron Emission Tomography (TOF-PET). *IEEE Transactions on Nuclear Science* 58:597-604.

19. Adam, L.-E., J. S. Karp, M. E. Daube-Witherspoon, and R. J. Smith. 2001. Performance of a Whole-Body PET Scanner Using Curve-Plate NaI(Tl) Detectors. *Journal of Nuclear Medicine* 42:1821-1830.
20. Freifelder, R., C. Cardi, I. Grigoras, J. R. Saffer, and J. S. Karp. 2001. First Results of a Dedicated Breast PET Imager, BPET, Using NaI(Tl) Curve Plate Detectors. In *IEEE Nuclear Science Symposium Conference Record*. 1241-1245 vol.1243.
21. Anderson, D. F. 1989. Properties of the High-Density Scintillator Cerium Fluoride. *IEEE Transactions on Nuclear Science* 36:137-140.
22. Schotanus, P., C. W. E. Van Eijk, and R. W. Hollander. 1988. A BaF₂-MWPC Gamma Camera for Positron Emission Tomography. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 269:377-384.
23. Kuhn, A., S. Surti, J. S. Karp, P. S. Raby, K. S. Shah, A. E. Perkins, and G. Muehllehner. 2004. Design of a Lanthanum Bromide Detector for Time-of-Flight PET. *IEEE Transactions on Nuclear Science* 51:2550-2557.
24. Surti, S., J. S. Karp, G. Muehllehner, and P. S. Raby. 2002. Investigation of Lanthanum Scintillators for 3D PET. In *IEEE Nuclear Science Symposium Conference Record*. 1177-1181.
25. Worstell, W., O. Johnson, H. Kudrolli, and V. Zavarzin. 1998. First Results with High-Resolution PET Detector Modules Using Wavelength-Shifting Fibers. *IEEE Transactions on Nuclear Science* 45:2993-2999.
26. Knoll, G. 2000. *Radiation Detection and Measurement*. John Wiley and Sons, Hoboken, NJ.

27. van Loef, E. V. D., P. Dorenbos, C. W. E. van Eijk, K. W. Krämer, and H. U. Güdel. 2002. Scintillation Properties of LaBr₃:Ce³⁺ Crystals: Fast, Efficient and High-Energy-Resolution Scintillators. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 486:254-258.
28. Bruyndonckx, P., S. Leonard, S. Tavernier, C. Lemaitre, O. Devroede, W. Yibao, and M. Krieguer. 2004. Neural Network-Based Position Estimators for PET Detectors Using Monolithic LSO blocks. *IEEE Transactions on Nuclear Science* 51:2520-2525.
29. Lewellen, T. K. 2010. The Challenge of Detector Designs for PET. *American Journal of Roentgenology* 195:301-309.
30. Levin, C. S., A. M. K. Foudray, P. D. Olcott, and F. Habte. 2003. Investigation of Position Sensitive Avalanche Photodiodes for a New High Resolution PET Detector Design. In *IEEE Nuclear Science Symposium Conference Record*. 2262-2266.
31. Moehrs, S., A. D. Guerra, D. J. Herbert, and M. A. Mandelkern. 2006. A Detector Head Design for Small-Animal PET With Silicon Photomultipliers (SiPM). *Physics in Medicine and Biology* 51:1113.
32. Moses, W. W., M. Janecek, M. A. Spurrier, P. Szupryczynski, W. S. Choong, C. L. Melcher, and M. Andreaco. 2010. Optimization of a LSO-Based Detector Module for Time-of-Flight PET. *IEEE Transactions on Nuclear Science* 57:1570-1576.

33. Seidel, J., J. J. Vaquero, S. Siegel, W. R. Gandler, and M. V. Green. 1999. Depth Identification Accuracy of a Three Layer Phoswich PET Detector Module. *IEEE Transactions on Nuclear Science* 46:485-490.
34. Tsuda, T., H. Murayama, K. Kitamura, T. Yamaya, E. Yoshida, T. Omura, H. Kawai, N. Inadama, and N. Orita. 2004. A Four-Layer Depth of Interaction Detector Block for Small Animal PET. *IEEE Transactions on Nuclear Science* 51:2537-2542.
35. Yang, Y., P. A. Dokhale, R. W. Silverman, K. S. Shah, M. A. McClish, R. Farrell, G. Entine, and S. R. Cherry. 2006. Depth of Interaction Resolution Measurements for a High Resolution PET Detector Using Position Sensitive Avalanche Photodiodes. *Physics in Medicine and Biology* 51:2131.
36. Dokhale, P. A., R. W. Silverman, K. S. Shah, R. Grazioso, R. Farrell, J. Glodo, M. A. McClish, G. Entine, V.-H. Tran, and S. R. Cherry. 2004. Performance Measurements of a Depth-Encoding PET Detector Module Based on Position-Sensitive Avalanche Photodiode Read-Out. *Physics in Medicine and Biology* 49:4293.
37. Godinez, F., A. J. Chaudhari, Y. Yang, R. Farrell, and R. D. Badawi. 2012. Characterization of a High-Resolution Hybrid DOI Detector for a Dedicated Breast PET/CT Scanner. *Physics in Medicine and Biology* 57:3435.
38. Yang, Y., S. S. James, Y. Wu, H. Du, J. Qi, R. Farrell, P. A. Dokhale, K. S. Shah, K. Vaigneur, and S. R. Cherry. 2011. Tapered LSO Arrays for Small Animal PET. *Physics in Medicine and Biology* 56:139.

39. Yang, Y., J. Qi, Y. Wu, S. S. James, R. Farrell, P. A. Dokhale, K. S. Shah, and S. R. Cherry. 2009. Depth of Interaction Calibration for PET detectors with Dual-Ended Readout by PSAPDs. *Physics in Medicine and Biology* 54:433.
40. Moses, W. W., S. E. Derenzo, C. L. Melcher, and R. A. Manente. 1995. A Room Temperature LSO/PIN Photodiode PET Detector Module that Measures Depth of Interaction. *IEEE Transactions on Nuclear Science* 42:1085-1089.
41. Bartzakos, P., and C. Thompson 1991. A PET Detector with Depth-of-Interaction Determination. *Physics in Medicine and Biology* 36:735-748.
42. Janecek, M., and W. W. Moses. 2008. Optical Reflectance Measurements for Commonly Used Reflectors. *IEEE Transactions on Nuclear Science* 55:2432-2437.
43. Buzhan, P., B. Dolgoshein, L. Filatov, A. Ilyin, V. Kantzerov, V. Kaplin, A. Karakash, F. Kayumov, S. Klemin, E. Popova, and S. Smirnov. 2003. Silicon Photomultiplier and Its Possible Applications. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 504:48-52.
44. Degenhardt, C., G. Prescher, T. Frach, A. Thon, R. de Gruyter, A. Schmitz, and R. Ballizany. 2009. The Digital Silicon Photomultiplier - A Novel Sensor for the Detection of Scintillation Light. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 2383-2386.
45. Anger, H. O. 1958. Scintillation Camera. *Review of Scientific Instruments* 29:27-33.

46. Chang Lyong, K., W. Gin-Chung, and S. Dolinsky. 2009. Multi-Pixel Photon Counters for TOF PET Detector and Its Challenges. *IEEE Transactions on Nuclear Science* 56:2580-2585.
47. Degenhardt, C., G. Prescher, T. Frach, A. Thon, R. de Gruyter, A. Schmitz, and R. Ballizany. 2009. The digital Silicon Photomultiplier — A novel sensor for the detection of scintillation light. In *Nuclear Science Symposium Conference Record (NSS/MIC)*, 2009 IEEE. 2383-2386.
48. Frach, T., G. Prescher, C. Degenhardt, and B. Zwaans. 2010. The Digital Silicon Photomultiplier - System Architecture and Performance Evaluation. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 1722-1727.
49. McClish, M., P. Dokhale, J. Christian, E. Johnson, C. Stapels, R. Robertson, and K. S. Shah. 2010. Characterization of CMOS Position Sensitive Solid-State Photomultipliers. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 624:492-497.
50. McClish, M., P. Dokhale, J. Christian, C. Stapels, E. Johnson, F. Augustine, and K. S. Shah. 2011. Performance Measurements from LYSO Scintillators Coupled to a CMOS Position Sensitive SSPM Detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 652:264-267.
51. Puill, V., C. Bazin, D. Breton, L. Burmistrov, V. Chaumat, N. Dinu, J. Maalmi, J. F. Vagnucci, and A. Stocchi. 2012. Single photoelectron timing resolution of SiPM as a function of the bias voltage, the wavelength and the temperature.

Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment In Press.

52. Mazzuca, E., M. Benetti, S. Mariazzi, R. S. Brusa, G.-F. D. Betta, and C. Piemonte. 2012. Compact Gamma Detectors Based on FBK SiPMs for a Ps Time of Flight Apparatus. *Journal of Instrumentation* 7:P05006.
53. Heinrichs, U., A. Blume, N. Bußmann, R. Engels, G. Kemmerling, S. Weber, and K. Ziemons. 2002. Statistical Studies on the Light Output and Energy Resolution of Small LSO Single Crystals with Different Surface Treatments Combined with Various Reflector Materials. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 486:60-66.
54. Shao, Y., R. W. Silverman, R. Farrell, L. Cirignano, R. Grazioso, K. S. Shah, G. Vissel, M. Clajus, T. O. Tumer, and S. R. Cherry. 2000. Design Studies of a High Resolution PET Detector Using APD Arrays. *IEEE Transactions on Nuclear Science* 47:1051-1057.
55. Shao, Y., H. Li, and K. Gao. 2007. Initial Experimental Studies of Using Solid-State Photomultiplier for PET Applications. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 580:944-950.
56. Huber, J. S., W. W. Moses, S. E. Derenzo, M. H. Ho, M. S. Andreaco, M. J. Paulus, and R. Nutt. 1997. Characterization of a 64 Channel PET Detector Using Photodiodes for Crystal Identification. *IEEE Transactions on Nuclear Science* 44:1197-1201.

57. Bircher, C., and Y. Shao. 2012. Use of Internal Scintillator Radioactivity to Calibrate DOI Function of a PET Detector with a Dual-Ended-Scintillator Readout. *Medical Physics* 39:777-787.
58. Burr, K. C., A. Ivan, D. E. Castleberry, J. W. LeBlanc, K. S. Shah, and R. Farrell. 2004. Evaluation of a Prototype Small-Animal PET Detector with Depth-of-Interaction Encoding. *IEEE Transactions on Nuclear Science* 51:1791-1798.
59. James, S. S., Y. Yang, S. L. Bowen, J. Qi, and S. R. Cherry. 2010. Simulation Study of Spatial Resolution and Sensitivity for the Tapered Depth of Interaction PET Detectors for Small Animal Imaging. *Physics in Medicine and Biology* 55:N63-74.
60. Moses, W. W., and M. Ullisch. 2006. Factors Influencing Timing Resolution in a Commercial LSO PET Camera. *IEEE Transactions on Nuclear Science* 53:78-85.
61. Slates, R., A. Chatziioannou, B. Fehlberg, L. Taekyeung, and S. Cherry. 1999. Chemical Polishing of LSO Crystals to Increase Light Output. In *IEEE Nuclear Science Symposium Conference Record*. 939-942 vol.932.
62. Huber, J. S., W. W. Moses, M. S. Andreaco, M. Loope, C. L. Melcher, and R. Nutt. 1999. Geometry and Surface Treatment Dependence of the Light Collection from LSO Crystals. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 437:374-380.
63. Strul, D., J. Sutcliffe-Goulden, P. Halstead, and P. K. Marsden. 2001. Optimisation of Fibre-Optic Readout of LSO Scintillation Crystals with Acid Etching. In *IEEE Nuclear Science Symposium Conference Record*. 1955-1959.

64. Levin, C. S. 2002. Design of a High-Resolution and High-Sensitivity Scintillation Crystal Array for PET with Nearly Complete Light Collection. *IEEE Transactions on Nuclear Science* 49:2236-2243.
65. Shao, Y., K. Meadors, R. W. Silverman, R. Farrell, L. Cirignano, R. Grazioso, K. S. Shah, and S. R. Cherry. 2002. Dual APD Array Readout of LSO crystals: Optimization of Crystal Surface Treatment. *IEEE Transactions on Nuclear Science* 49:649-654.
66. Pichler, B. J., G. Boning, M. Rafecas, M. Schlosshauer, E. Lorenz, and S. I. Ziegler. 1999. LGSO Scintillation Crystals Coupled to New Large Area APDs Compared to LSO and BGO. *IEEE Transactions on Nuclear Science* 46:289-291.
67. Maas, M. C., D. R. Schaart, D. J. van der Laan, P. Bruyndonckx, C. Lemaître, F. J. Beekman, and C. W. E. van Eijk. 2009. Monolithic Scintillator PET Detectors with Intrinsic Depth-of-Interaction Correction. *Physics in Medicine and Biology* 54:1893.
68. Clement, D., R. Frei, J. F. Loude, and C. Morel. 1998. Development of a 3D Position Sensitive Scintillation Detector Using Neural Networks. In *IEEE Nuclear Science Symposium Conference Record*. 1448-1452 vol.1443.
69. Tavernier, S., P. Bruyndonckx, S. Leonard, and O. Devroede. 2005. A High-Resolution PET Detector Based on Continuous Scintillators. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 537:321-325.
70. Bruyndonckx, P., C. Lemaître, D. Schaart, M. Maas, D. J. van der Laan, M. Krieguer, O. Devroede, and S. Tavernier. 2007. Towards a Continuous Crystal

- APD-based PET Detector Design. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 571:182-186.
71. Bruyndonckx, P., S. Leonard, L. Jiangui, S. Tavernier, P. Szupryczynski, and A. Fedorov. 2003. Study of Spatial Resolution and Depth of Interaction of APD-Based PET Detector Modules Using Light Sharing Schemes. IEEE Transactions on Nuclear Science 50:1415-1419.
 72. DeVol, T. A., W. W. Moses, and S. E. Derenzo. 1993. Monte Carlo Optimization of Depth-of-Interaction Resolution in PET Crystals. IEEE Transactions on Nuclear Science 40:170-174.
 73. Ling, T., T. K. Lewellen, and R. S. Miyaoka. 2007. Depth of Interaction Decoding of a Continuous Crystal Detector Module. Physics in Medicine and Biology 52:2213.
 74. Gagnon, D., N. Pouliot, L. Laperriere, M. Therrien, and P. Olivier. 1993. Maximum Likelihood Positioning in the Scintillation Camera Using Depth of Interaction. IEEE Transactions on Medical Imaging 12:101-107.
 75. Tomitani, T., Y. Futami, Y. Izeki, S. Kouda, T. Nishio, T. Murakami, A. Kitagawa, M. Kanazawa, E. Urakabe, M. Shinbo, and T. Kanai. 1999. Depth Encoding of Point-of-Interaction in Thick Scintillation Cameras. In IEEE Nuclear Science Symposium Conference Record. 1182-1186.
 76. Shao, Y., R. Yao, and T. Ma. 2008. A Novel Method to Calibrate DOI Function of a PET Detector with a Dual-Ended-Scintillator Readout. Medical Physics 35:5829-5840.

77. Yao, R., T. Ma, and Y. Shao. 2008. Lutetium oxyorthosilicate (LSO) Intrinsic Activity Correction and Minimal Detectable Target Activity Study for SPECT Imaging with a LSO-Based Animal PET Scanner. *Physics in Medicine and Biology* 53:4399.
78. Moszynski, M., M. Balcerzyk, M. Kapusta, D. Wolski, and C. L. Melcher. 2000. Large Size LSO:Ce and YSO:Ce Scintillators for 50 MeV Range Gamma-Ray Detector. *IEEE Transactions on Nuclear Science* 47:1324-1328.
79. Dhanasopon, A. P., C. S. Levin, A. M. K. Foudray, P. D. Olcott, and F. Habte. 2005. Scintillation Crystal Design Features for a Miniature Gamma Ray Camera. *IEEE Transactions on Nuclear Science* 52:1439-1446.
80. Ramirez, R. A., L. Shitao, A. Shaohui, Z. Yuxuan, L. Hongdi, H. Baghaei, W. Chao, and W. Wai-Hoi. 2010. New Ultra High Resolution LYSO Pentagon Detector Blocks for Lower-Cost Murine PET-CT (MuPET/CT). In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 3366-3371.
81. Huber, J. S., W. W. Moses, M. S. Andreaco, and O. Petterson. 2000. A LSO Scintillator Array for a PET Detector Module with Depth of Interaction Measurement. In *IEEE Nuclear Science Symposium Conference Record*. 14/46-14/50.
82. Weber, M. F., C. A. Stover, L. R. Gilbert, T. J. Nevitt, and A. J. Ouder Kirk. 2000. Giant Birefringent Optics in Multilayer Polymer Mirrors. *Science* 287:2451-2456.
83. Pepin, C. M., P. Berard, and R. Lecomte. 2001. Assessment of Reflective Separator Films for Small Crystal Arrays. In *IEEE Nuclear Science Symposium Conference Record*. 879-883 vol.872.

84. Delfino, E. P., S. Majewski, R. R. Raylman, and A. Stolin. 2010. Towards 1mm PET Resolution Using DOI Modules Based on Dual-Sided SiPM Readout. In IEEE Nuclear Science Symposium Conference Record (NSS/MIC). 3442-3449.
85. Du, H., Y. Yang, and S. R. Cherry. 2008. Comparison of Four Depth-Encoding PET Detector Modules with Wavelength Shifting (WLS) and Optical Fiber Read-Out. *Physics in Medicine and Biology* 53:1829.
86. Mackewn, J. E., S. F. Keevil, W. A. Hallett, P. Halsted, R. A. Page, M. Kelly, S. C. R. Williams, and P. K. Marsden. 2006. An MR Compatible LSO-PET Scanner for Molecular Imaging Studies. In IEEE Nuclear Science Symposium Conference Record. 2986-2989.
87. Song, T. Y., H. Wu, S. Komarov, S. B. Siegel, and Y.-C. Tai. 2010. A Sub-Millimeter Resolution PET Detector Module Using a Multi-Pixel Photon Counter Array. *Physics in Medicine and Biology* 55:2573.
88. Casey, M. E., and R. Nutt. 1986. A Multicrystal Two Dimensional BGO Detector System for Positron Emission Tomography. *IEEE Transactions on Nuclear Science* 33:460-463.
89. Moriya, T., K. Fukumitsu, T. Sakai, S. Ohsuka, T. Okamoto, H. Takahashi, M. Watanabe, and T. Yamashita. 2010. Development of PET Detectors Using Monolithic Scintillation Crystals Processed with Sub-Surface Laser Engraving Technique. *IEEE Transactions on Nuclear Science* 57:2455-2459.
90. Shimizu, K., H. Uchida, K. Sakai, M. Hirayanagi, S. Nakamura, and T. Omura. 2011. Development of an MPPC Detector for PET. In IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC). 2991-2994.

91. Yamaya, T., T. Mitsuhashi, T. Matsumoto, N. Inadama, F. Nishikido, E. Yoshida, H. Murayama, H. Kawai, M. Suga, and M. Watanabe. 2011. A SiPM-Based Isotropic-3D PET Detector X'tal Cube with a Three-Dimensional Array of 1 mm³ Crystals. *Physics in Medicine and Biology* 56:6793.
92. Nagarkar, V. V., S. V. Tipnis, K. Shah, I. Shestakova, and S. R. Cherry. 2003. A High Efficiency Pixelated Detector for Small Animal PET. In *IEEE Nuclear Science Symposium Conference Record*. 2288-2290 Vol.2284.
93. Moses, W. W., and S. E. Derenzo. 1999. Prospects for Time-of-Flight PET Using LSO Scintillator. *IEEE Transactions on Nuclear Science* 46:474-478.
94. Jenkins, F., and H. White. 1976. *Fundamentals of Optics*. McGraw-Hill, New York.
95. Huber, J. S., W. W. Moses, M. S. Andreaco, and O. Petterson. 2000. A LSO scintillator array for a PET detector module with depth of interaction measurement. In *Nuclear Science Symposium Conference Record*, 2000 IEEE. 14/46-14/50 vol.12.
96. Popov, V., S. Majewski, J. Proffitt, J. McKisson, B. Kross, and A. Weisenberger. 2007. Study of Photonis XP1470 Nine-Channel Photomultiplier Tube for Applications in High Resolution Single Gamma and PET Imagers. In *IEEE Nuclear Science Symposium Conference Record*. 2870-2874.
97. Majewski, S., R. Raylman, B. Kross, V. Popov, A. G. Weisenberger, B. Welch, and R. Wojcik. 2003. Application of New Burle Multipad Flat Panel PMTs in Compact Gamma and Positron Imagers-Initial Evaluations. In *IEEE Nuclear Science Symposium Conference Record*. 2376-2378 Vol.2374.

98. Doshi, N. K., Y. Shao, R. W. Silverman, and S. R. Cherry. 2000. Design and Evaluation of an LSO PET Detector for Breast Cancer Imaging. *Medical Physics* 27:1535-1543.
99. Hayashi, T. 1989. New Photomultiplier Tubes for Medical Imaging. *IEEE Transactions on Nuclear Science* 36:1078-1083.
100. Chang Lyong, K., A. Ivan, and A. Ganin. 2010. A Compact SPECT Detector Based on a Quad PMT. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 2557-2559.
101. Sun, X., A. K. Lan, C. Bircher, Z. Deng, Y. Liu, and Y. Shao. 2011. Energy and Timing Measurement with Time-Based Detector Readout for PET Applications: Principle and Validation with Discrete Circuit Components. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 641:128-135.
102. Shao, Y., X. Sun, K. A. Lan, C. Bircher, Z. Deng, and L. Yinong. 2010. Energy and Timing Measurement of a PET Detector with Time-Based Readout Electronics. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 2504-2509.
103. Zhi, D., A. K. Lan, X. Sun, C. Bircher, L. Yinong, and Y. Shao. 2011. Development of an Eight-Channel Time-Based Readout ASIC for PET Applications. *IEEE Transactions on Nuclear Science* 58:3212-3218.
104. Janecek, M., J. Walder, P. J. McVittie, B. Zheng, H. von der Lippe, M. McClish, P. Dokhale, C. J. Stapels, J. F. Christian, K. S. Shah, and W. W. Moses. 2012. A High-Speed Multi-Channel Readout for SSPM Arrays. *IEEE Transactions on Nuclear Science* 59:13-18.

105. Stickel, J. R., J. Qi, and S. R. Cherry. 2007. Fabrication and Characterization of a 0.5-mm Lutetium Oxyorthosilicate Detector Array for High-Resolution PET Applications. *Journal of Nuclear Medicine* 48:115-121.
106. Kolb, A., E. Lorenz, M. S. Judenhofer, D. Renker, K. Lankes, and B. J. Pichler. 2010. Evaluation of Geiger-Mode APDs for PET Block Detector Designs. *Physics in Medicine and Biology* 55:1815.
107. Dolinsky, S., and S. Zelakiewicz. 2010. Multiplexing Requirements for Solid State Photomultipliers in Time-of-Flight PET. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 1763-1766.
108. Callier, S., F. Dulucq, R. Fabbri, C. de La Taille, B. Lutz, G. Martin-Chassard, L. Raux, and W. Shen. 2009. Silicon Photomultiplier Integrated Readout Chip (SPIROC) for the ILC: Measurements and Possible Further Development. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 42-46.
109. Corsi, F., M. Foresta, C. Marzocca, G. Matarrese, and A. Del Guerra. 2009. BASIC: An 8-Channel Front-End ASIC for Silicon Photomultiplier Detectors. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 1082-1087.
110. Dorn, M., T. Harion, W. Shen, G. Sidlauskas, and H. C. Schultz-Coulon. 2012. KLauS – A Charge Readout and Fast Discrimination Chip for Silicon Photomultipliers. *Journal of Instrumentation* 7:C01008.
111. Bagliesi, M. G., C. Avanzini, G. Bigongiari, R. Cecchi, M. Y. Kim, P. Maestro, P. S. Marrocchesi, and F. Morsani. 2011. A Custom Front-End ASIC for the Readout and Timing of 64 SiPM Photosensors. *Nuclear Physics B - Proceedings Supplements* 215:344-348.

112. Wei, S., and H. C. Schultz-Coulon. 2009. STIC - A Current Mode Constant Fraction Discriminator for Positron Emission Tomography Using SiPM (MPPC). In IEEE Nuclear Science Symposium Conference Record (NSS/MIC). 364-367.
113. Meier, D., S. Mikkelsen, J. Talebi, S. Azman, G. Mahlum, and B. E. Patt. 2010. An ASIC for SiPM/MPPC Readout. In IEEE Nuclear Science Symposium Conference Record (NSS/MIC). 1653-1657.
114. Pavlov, N., G. Mashlum, and D. Meier. 2005. Gamma Spectroscopy Using a Silicon Photomultiplier and a Scintillator. In IEEE Nuclear Science Symposium Conference Record. 173-180.
115. Wang, G. C., J. S. Huber, W. W. Moses, W. S. Choong, and J. S. Maltz. 2004. Calibration of a PEM Detector with Depth of Interaction Measurement. IEEE Transactions on Nuclear Science 51:775-781.
116. Cosentino, L., P. Finocchiaro, A. Pappalardo, and F. Garibaldi. 2012. High-Resolution Time-of-Flight PET with Depth-of-Interaction becomes Feasible: a Proof of Principle.
117. Taghibakhsh, F., C. S. Levin, and J. A. Rowlands. 2010. Analytic Pulse Height Correction in Dual-Ended Readout PET Detectors. In IEEE Nuclear Science Symposium Conference Record (NSS/MIC). 2151-2154.
118. Ur-Rehman, F., B. McIntosh, and A. L. Goertzen. 2010. Calibration of Dual-Ended Readout of Axially Oriented 100 mm Long LYSO Crystals for Use in a Compact PET System. In IEEE Nuclear Science Symposium Conference Record (NSS/MIC). 3470-3474.
119. Agodi, C., A. Bassi, R. Bassini, G. Bellia, M. Benovic, C. Boiano, S. Brambilla, R. Coniglione, L. Cosentino, P. Finocchiaro, S. Hlavac, I. Iori, W. Koenig, A.

- Kugler, C. Maiolino, T. Marek, P. Piattelli, R. Pleskač, P. Sapienza, S. Spataro, M. Suk, A. Taranenko, P. Tlustý, I. Turzo, D. Vassiliev, V. Wagner, and D. Zovinec. 2002. The HADES Time-of-Flight Wall. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 492:14-25.
120. Spanoudaki, V. C., D. P. McElroy, and S. I. Ziegler. 2006. An Analog Signal Processing ASIC for a Small Animal LSO–APD PET Tomograph. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 564:451-462.
 121. Mann, A. B. 2010. Sampling ADC Based Data Acquisition for Positron Emission Tomography. In *Physik-Department. Fakultät für Elektrotechnik und Informationstechnik Garching*. 182.
 122. Jarron, P., E. Auffray, S. E. Brunner, M. Despeisse, E. Garutti, M. Goettlich, H. Hillemanns, P. Lecoq, T. Meyer, F. Powolny, W. Shen, H. C. Schultz-Coulon, and M. C. S. Williams. 2009. Time Based Readout of a Silicon Photomultiplier (SiPM) for Time of Flight Positron Emission Tomography (TOF-PET). In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 1212-1219.
 123. Spanoudaki, V. C., and C. S. Levin. 2010. Photo-Detectors for Time of Flight Positron Emission Tomography (ToF-PET). *Sensors* 10:10484-10505.
 124. Akiba, M., K. Inagaki, and K. Tsujino. 2012. Photon Number Resolving SiPM Detector with 1 GHz Count Rate. *Optical Express* 20:2779-2788.
 125. Bonanno, G., P. Finocchiaro, A. Pappalardo, S. Billotta, L. Cosentino, M. Belluso, S. Di Mauro, and G. Occhipinti. 2009. Precision Measurements of Photon Detection Efficiency for SiPM Detectors. *Nuclear Instruments and*

Methods in Physics Research Section A: Accelerators, Spectrometers,
Detectors and Associated Equipment 610:93-97.

126. Stewart, A. G., V. Saveliev, S. J. Bellis, D. J. Herbert, P. J. Hughes, and J. C. Jackson. 2008. Performance of 1-mm² Silicon Photomultiplier. *IEEE Journal of Quantum Electronics* 44:157-164.
127. Akiba, M., K. Tsujino, K. Sato, and M. Sasaki. 2009. Multipixel Silicon Avalanche Photodiode with Ultralow Dark Count Rate at Liquid Nitrogen Temperature. *Optical Express* 17:16885-16897.
128. Vinke, R., H. Löhner, D. R. Schaart, H. T. van Dam, S. Seifert, F. J. Beekman, and P. Dendooven. 2009. Optimizing the Timing Resolution of SiPM Sensors for Use in TOF-PET Detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 610:188-191.
129. Llosa, G., N. Belcari, M. G. Bisogni, G. Collazuol, A. Del Guerra, S. Marcatili, S. Moehrs, and C. Piemonte. 2007. Silicon Photomultipliers and SiPM Matrices as Photodetectors in Nuclear Medicine. In *IEEE Nuclear Science Symposium Conference Record*. 3220-3223.
130. Llosa, G., N. Belcari, M. G. Bisogni, G. Collazuol, S. Marcatili, P. Barrillon, C. de la Taille, S. Bondil-Blin, N. Dinu, M. Melchiorri, A. Tarolli, C. Piemonte, and A. Del Guerra. 2009. Energy, Timing and Position Resolution Studies With 16-Pixel Silicon Photomultiplier Matrices for Small Animal PET. *IEEE Transactions on Nuclear Science* 56:2586-2593.

131. Solf, T., V. Schulz, B. Weissler, A. Thon, P. Fischer, M. Ritzert, V. Mlotok, C. Piemonte, and N. Zorzi. 2009. Solid-State Detector Stack for ToF-PET/MR. In IEEE Nuclear Science Symposium Conference Record (NSS/MIC). 2798-2799.
132. Gola, A., C. Piemonte, and A. Tarolli. 2012. The DLED Algorithm for Timing Measurements on Large Area SiPMs Coupled to Scintillators. IEEE Transactions on Nuclear Science 59:358-365.
133. Spanoudaki, V. C., A. B. Mann, A. N. Otte, I. Konorov, I. Torres-Espallardo, S. Paul, and S. I. Ziegler. 2007. Use of Single Photon Counting Detector Arrays in Combined PET/MR: Characterization of LYSO-SiPM Detector Modules and Comparison with a LSO-APD Detector. Journal of Instrumentation 2:P12002.
134. Schulz, V., B. Weissler, P. Gebhardt, T. Solf, C. W. Lerche, P. Fischer, M. Ritzert, V. Mlotok, C. Piemonte, B. Goldschmidt, S. Vandenberghe, A. Salomon, T. Schaeffter, and P. K. Marsden. 2011. SiPM Based Preclinical PET/MR Insert for a Human 3T MR: First Imaging Experiments. In IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC). 4467-4469.
135. Llosa, G., J. Barrio, J. Cabello, C. Lacasta, M. Rafecas, P. Barrillon, S. Bondil-Blin, C. de La Taille, C. Piemonte, G. Ambrosi, P. Azzarello, M. G. Bisogni, and A. Del Guerra. 2010. First PET Imaging Results with Continuous LYSO Crystals and Monolithic, 64-Pixel SiPM Matrices. In IEEE Nuclear Science Symposium Conference Record (NSS/MIC). 3376-3379.
136. Dokhale, P., R. Robertson, C. Stapels, J. Cristian, M. Kaul, S. Surti, J. Karp, P. Vaska, and K. Shah. 2011. Continuous LYSO-SSPM Array Based PET Detectors for Clinical and Small Volume Imaging studies. In IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC). 2350-2354.

137. Wang, Y., Z. Zhang, D. Li, B. Wang, L. Shuai, B. Feng, P. Chai, S. Liu, H. Tang, T. Li, Y. Liao, X. Huang, Y. Chen, Y. Liu, Y. Zhang, and L. Wei. 2012. Design and Performance Evaluation of a Compact, Large-Area PET Detector Module Based on Silicon Photomultipliers. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 670:49-54.
138. Schaart, D. R., H. T. v. Dam, S. Seifert, R. Vinke, P. Dendooven, H. Löhner, and F. J. Beekman. 2009. A Novel, SiPM-Array-Based, Monolithic Scintillator Detector for PET. *Physics in Medicine and Biology* 54:3501.
139. Vinke, R., S. Seifert, D. R. Schaart, F. P. Schreuder, M. R. de Boer, H. T. van Dam, F. J. Beekman, H. Lohner, and P. Dendooven. 2009. Optimization of Digital Time Pickoff Methods for LaBr^3 -SiPM TOF-PET Detectors. In *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*. 2962-2968.
140. Sanjani, S. S., F. Taghibakhsh, and C. S. Levin. 2011. Energy and Time Characterization of Silicon Photomultiplier Detector Blocks. In *IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. 3045-3047.
141. Zar, J. H. 1984. *Biostatistical Analysis*. Prentice-Hall INC, Englewood Cliffs, N.J.
142. Levin, C. 2004. Basic Physics of Radioisotope Imaging. In *Emission Tomography*. M. N. Wernick, and J. N. Aarsvold, editors. Elsevier Inc., San Diego.
143. Dokhale, P. A., R. W. Silverman, K. S. Shah, R. Farrell, M. A. McClish, G. Entine, and S. R. Cherry. 2006. Intrinsic Spatial Resolution and Parallax Correction Using Depth-Encoding PET Detector Modules Based on Position-Sensitive APD Readout. *IEEE Transactions on Nuclear Science* 53:2666-2670.

144. Yoshida, E., H. Tashima, N. Inadama, F. Nishikido, T. Moriya, T. Omura, M. Watanabe, H. Murayama, and T. Yamaya. 2012. Intrinsic Spatial Resolution Evaluation of the X'tal cube PET Detector Based on a 3D Crystal Block Segmented by Laser Processing. *Radiological Physics and Technology*:1-7.
145. Ito, M., J. S. Lee, M.-J. Park, K.-S. Sim, and S. J. Hong. 2010. Design and Simulation of a Novel Method for Determining Depth-of-Interaction in a PET Scintillation Crystal Array Using a Single-Ended Readout by a Multi-Anode PMT. *Physics in Medicine and Biology* 55:3827-3841.
146. Yongfeng, Y., W. Yibao, and S. R. Cherry. 2009. Investigation of Depth of Interaction Encoding for a Pixelated LSO Array with a Single Multi-Channel PMT. *IEEE Transactions on Nuclear Science* 56:2594-2599.

Vita

Chad Jacob Bircher was born April 16, 1981, in Denver CO to James and Nanette Hagle. After graduating from Cooper High School in Abilene, TX in 2000, he entered Abilene Christian University on an academic scholarship. He earned the Bachelor of Science degree in Physics in 2003, and enrolled in the Graduate Physics program at the University of Tennessee in Knoxville in 2004. He earned the Master of Science degree in Physics in 2007 and after graduation worked in industry until the fall of 2008. In August 2008, he entered the University of Texas Health Science Center at Houston Graduate School of Biomedical Science in the Medical Physics Program. He married Kristine Nicholls in 2004 and is the father of Hunter and Brentyn.