

Algorithms and Electronics for Processing Data from Pixelated Semiconductor Gamma-Ray Detectors

by

Matthew Petryk

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Nuclear Engineering)
Nuclear Engineering and Radiological Sciences
in the University of Michigan
2023

Doctoral Committee:

Professor Zhong He, Chair
Professor Michael Flynn
Adjunct Research Scientist Gianluigi DeGeronimo
Professor Sara Pozzi
Associate Research Scientist Yuefeng Zhu

Matthew Petryk
mpetryk@umich.edu
ORCID iD: 0000-0001-6856-8389

© Matthew Petryk 2023

Dedication

To my family.

*There is our wildest mount- a headless horse.
But though it runs unbridled off its course,
And all our blandishments would seem defied,
We have ideas yet that we haven't tried.*

- Robert Frost

Acknowledgments

I would like to thank my mentor Dr. Zhong He for his advice, guidance, and the many conversations we've had over the years, academic and otherwise. Thank you to the other members of my thesis committee. I also owe a debt to members of the Orion Group, past and present. James Berry provided invaluable feedback and support, both mechanical and theoretical. Thank you especially to Erik Hall, Sara Abraham, Alex Rice, Peter Hotvedt and Max Robbins for reviewing the chapters of this thesis and providing invaluable feedback. Jack Kuchta provided much support, hosting weekend get-togethers and organizing other outings, as well as making sure the lab was never quiet and always exciting. Drs. Daniel Shy, Niral Shah, and Charles Leak were always willing to explain concepts (especially Daniel who helped make clear many complicated imaging concepts). I will remember our travels in Australia, England and Scotland along with Dr. William Steinberger very fondly. The past members of the Orion Group, including Drs. Bennet Williams, Jiawei Xia, Sean O'Neal, and David Goodman, laid the foundations on which I stand. Dr. Zhuo Chen's talent in mastering new concepts and digital ASIC design was critical for the DSP project to proceed as far as it has. Chris Meert provided much advice and occasional experimental support, particularly anytime those sneaky neutrons were involved. Dr. Stefano Marin introduced me to many fascinating concepts, and was a worthy opponent in our weekend tennis matches. Thank you as well to Kimmy Beers and Roxanne & Hayden Walker for being my friends through the last several years of the Ph.D. Most of all, thank you to Colleen for your help, support, and friendship, especially over the last several months of writing. Your kindness and baking prowess both are without limit.

I owe a very large debt of gratitude to Dr. Aleksey Bolotnikov at Brookhaven National Laboratory. You introduced me to CZT over eleven years ago, but also to the work of scientists and engineers. Your example led me to follow the path that has taken me to where I am now. Thank you, for your guidance then and your support ever since.

Thank you to my friends from Binghamton - Alex, Mike, Olivia, Connie, Laura, and Connor. Between our yearly Christmas parties, our ten year reunion, and many pandemic Zoom sessions, you were a bright constancy in my life. I look forward to many more reunions to come!

Lastly, but most of all, thank you to my family. Your love and support over the years have helped me through many challenges. Margie & Suzie are the best sisters anyone could ask for.

Greg and Magda for Kona!




Table of Contents

Dedication	ii
Acknowledgments	iii
List of Figures	viii
Abstract	xv
Chapter 1. Introduction	1
1.1 Radiation and Matter Interactions	1
1.2 Semiconductor Radiation Detectors	3
1.3 CdZnTe 3-D Position Sensitive Pixelated Detectors	3
1.3.1 Device Overview	3
1.3.2 Shockley-Ramo Theorem and Single-Polarity Sensing	4
1.3.3 Electronics	6
1.4 Basic Digital Signal Processing and Event Reconstruction	8
1.4.1 Waveform Processing and Coarse Position Reconstruction	8
1.4.2 Subpixel Position Reconstruction	10
1.5 Imaging Gamma-Rays with 3D CZT Detectors	11
1.5.1 Coded Aperture Imaging	12
1.5.2 Compton Imaging	15
1.6 Overview of Thesis and Contributions to Science	17
Chapter 2. Digital Signal Processing ASIC	19
2.1 ASIC development and signal processing	19
2.1.1 Digital Data Processing	19
2.1.2 Integrated Circuits	21
2.1.3 Digital Logic Design Flow	21
2.2 DSP ASIC Overview	23
2.2.1 Current hardware description	23
2.2.2 DSP ASIC Goals	23
2.2.3 DSP ASIC Top-level Chip Architecture	24

2.3	DSP Architecture	26
2.3.1	DSP Architecture Overview	26
2.3.2	DSP Stage 1	29
2.3.3	DSP Stage 2	33
2.3.4	DSP Stage 3	34
2.3.5	Control and Pipelining	36
2.4	Synthesis, Layout, and Verification	38
2.4.1	Synthesis and Layout	38
2.4.2	Verification Simulations	39
2.5	Test Systems	42
2.5.1	Simplified Test System	42
2.5.2	Integrated Test System	43
2.6	Experimental Evaluation of v1.0 and v1.1	45
2.6.1	v1.0 Debugging	45
2.6.2	v1.1 Debugging	46
2.7	Modifications for Future Version	48
Chapter 3. Advanced Subpixel Position Calibrations		50
3.1	Charge Transport and Position Reconstruction in CZT	50
3.1.1	Subpixel Position - Accuracy and Precision	50
3.1.2	Factors Affecting Subpixel Resolution	51
3.1.3	Mechanisms of Charge Loss and Charge Steering	53
3.2	Subpixel Precision Improvements with Machine Learning	57
3.2.1	Machine Learning Concepts	57
3.2.2	Radial Position Reconstruction	59
3.2.3	XY Position Reconstruction	61
3.3	Charge Steering Calibration Techniques	65
3.3.1	Collimator Scan Corrections	65
3.3.2	Count Density Mapping	68
3.4	Future Work for Improving Subpixel Position Reconstruction	70
Chapter 4. Imaging of Pair Production Events		72
4.1	Pair Production Physics	72
4.1.1	Pair and Triplet Production and Annihilation	72
4.1.2	Cross Sections for Pair Production	73
4.2	Pair Production Telescopes	74
4.2.1	Pair Production Telescope Design	74
4.3	Previous Work on High Energy Imaging with Pixelated CdZnTe	76
4.3.1	Detection of High Energy Photons with 3-D Pixelated CdZnTe	76
4.3.2	Compton Imaging of High Energy Photons >3 MeV	77

4.4	Pair Production Imaging Algorithm	78
4.4.1	Relevant Signal Induction Physics	78
4.4.2	Algorithm Description	80
4.4.3	Pair Production Event Simulations	82
4.4.4	Pair Production Imaging Simulation Results	88
4.4.5	Experimental Verification	88
4.5	Future Work: Filtered Backprojection, MLEM, and Particle Tracking	96
Chapter 5. Perovskite Radiation Detectors		100
5.1	Alternative Materials to CdZnTe	100
5.1.1	The Search for the Ideal Detector	100
5.1.2	TlBr, HgI ₂ and Other Materials	101
5.2	Introduction to Perovskites	102
5.2.1	Perovskite Structure	102
5.2.2	Applications	104
5.3	Perovskite Detectors for Gamma Rays	105
5.3.1	Early planar samples	105
5.3.2	Medium thickness Pixelated Detectors	108
5.3.3	Large Pixelated Detectors	111
5.3.4	Hybrid Perovskite Radiation Detectors	116
5.4	Conditioning in Perovskite Radiation Detectors	120
5.4.1	Conditioning of Furnace Grown Samples	120
5.4.2	Conditioning of Solution Grown Samples	126
5.5	Electron-Hole Conversion Effects and PCA	128
5.5.1	Anomalous Gain and Electron-Hole Conversion	128
5.5.2	Furnace Grown Samples - Corrections with PCA	135
5.5.3	Solution Grown Samples - Corrections with PCA	137
5.6	Future of Perovskite Radiation Detectors	138
Chapter 6. Concluding Remarks		140
6.1	Review of Work	140
6.2	Possible Subsequent Work and Visions for the Future	142
Bibliography		144

List of Figures

1.1	There are many types of particles that fall under the umbrella of radiation.	2
1.2	Metalization pattern of CZT detectors used in this work.	4
1.3	Simulated weighting potential for the cathode (red) and a central anode pixel (blue) along the Z axis perpendicular to the anode plane.	6
1.4	Example anode pixel and cathode waveforms from a radiation interaction output from the VAD_UMv2.2 digital ASIC.	7
1.5	Standard attach and direct attach detectors and ASICs.	8
1.6	Collecting pixel, cathode, and noncollecting neighbor waveforms from a charge cloud which passed close to the top-left edge of the collecting pixel.	11
1.7	Operating principle of a simple pinhole camera.	12
1.8	MURA mask made of tungsten, for a single CZT module.	13
1.9	Figure courtesy of Alexander Rice.	14
1.10	Compton imaging of the 372 keV line from a PuBe source.	16
2.1	Examples of computers that fall into analog/digital and application specific/general purpose categories.	20
2.2	Design flow for digital logic. Starting from HDL descriptions, through to RTL synthesis and layout, and system simulation and verification.	22
2.3	Current hardware layout of the Orion 3D CZT detector array systems.	23
2.4	Diagram of a proposed detector system based on new DSP ASIC.	24
2.5	Top-level architecture of the first version of the DSP ASIC.	25
2.6	Top level DSP core block diagram.	28
2.7	Timing diagram of the operation of the pipelined DSP.	29
2.8	Detailed block diagram of DSP stage 1.	30
2.9	Progression of the rolling average value for a single channel after many input waveforms. From the reset value of zero, it gradually converges to the true value. Fixing the reset value close to the expected baseline reduces the time required for convergence.	31
2.10	Operation of the front-end decay deconvolution module on an experimental H3DD-UM waveform output from the rolling average baseline subtraction module.	32

2.11	Detailed block diagram of DSP stage 2.	33
2.12	Implementation in hardware of the FIR filters. Each filter is 256-tap. A single operation involves 256 multiplies and 255 additions. The blue blocks represent individual 13-bit registers that connect up into a single large shift register bank. Individual waveform samples x_i are clocked in to the shift registers and multiplied by the corresponding filter taps. . .	34
2.13	An experimental waveform from the H3DD-UM ASIC after processing in stage 1 and stage 2 of the DSP.	35
2.14	Detailed block diagram of DSP stage 3.	35
2.15	Operation of the timing pickoff module.	37
2.16	Physical layout of the assembled v1.0 chip. The chip area is 4.2 x 4.2 mm, including the seal ring.	39
2.17	Floorplanning of the DSP core to help guide amebaplace and limit DRC violations from via/metal spacings.	40
2.18	Layout of the DSP after APR. Six large FIR units are clearly visible. The next largest block of area is taken up by pipeline registers, followed by stage 1 and 3 logic.	41
2.19	Example partial output of the integrated testbench of the DSP ASIC system processing data during readout of the H3DD-UM front end. Post-APR delays are included. In this simulation, a single pixel event in channel 65 of the front end is simulated. The very bottom line is the data valid signal to the DSP from the H3DD_Core; here, it is high three times, twice for cathode signals in channels 0 and 64 and once for the triggered pixel in channel 65 immediately after.	42
2.20	Reset bridge structure used to create an asynchronous-assert/synchronous-release reset.	42
2.21	Simplified DSP ASIC test system.	43
2.22	System diagram of the integrated test system and peripherals.	44
2.23	Integrated test system hardware.	45
2.24	Fabricated ASIC incorporated in a 7x7 mm QFN package.	46
2.25	Focused ion beam fixing of the DSP ASIC v1.1.	47
2.26	Output signals from the on-chip H3DD_Core module at -30°C in response to an ASIC_gReset command sent over SPI. On the bottom is the expected output from behavioral simulation. The measured matches the simulated in broad morphology, but several timings are modified and glitches are present in some of the signals.	48
3.1	Difference between accuracy and precision.	51
3.2	Distribution of opposing neighbor ratios in X and Y, showing the “X” distribution which is due to the nonlinearity of equation 1.5.	52

3.3	Subpixel distribution of photopeak events in CZT following irradiation through a $100\mu\text{m}$ wide slit collimator. Ideally, a straight line should be observed. In practice, nonuniform electric fields cause distortions in the reconstruction.	54
3.4	Fine detail of the pixel pattern used with the CZT detectors in this work. The pitch between pixels is 1.72 mm , while the pixel size is $1.66 \times 1.66\text{ mm}^2$. The resulting gap is $60\ \mu\text{m}$ wide.	55
3.5	Pixel pattern with inter-pixel grid, used by older generation CZT detectors.	56
3.6	Diagram of a neural network architecture with three hidden layers. . .	58
3.7	Algorithm for training a machine learning model to predict the subpixel radial position of interaction purely from the collecting waveform. . . .	60
3.8	Radial positions predicted from only collecting waveforms vs. those calculated from opposing neighbor ratios.	61
3.9	XY distributions (calculated from neighbor ratios) of events with particular values of radial position (predicted by GPR model).	62
3.10	Radial distribution of events with predicted values of R between 0.4 and 0.5. Converting to mm, the FWHM of the distribution is about $200\ \mu\text{m}$	63
3.11	Performance of the trained GPR model for predicting radial subpixel positions.	64
3.12	Performance of classic techniques benchmarked in a crystal from Orion-gamma using a $100\ \mu\text{ m}$ slit collimator.	65
3.13	Performance of the trained neural networks for predicting the subpixel XY position.	66
3.14	Demonstration of accuracy error in cathode/anode ratio depth reconstruction for a corner pixel.	67
3.15	Accuracy improvement after applying collimator-based calibrations. . .	68
3.16	Count density-based voxel reshaping algorithm applied to a set of calibration data. (Left) initial count density overpressures; (right) after voxel boundary shifting to accomplish a uniform distribution.	69
3.17	Accuracy error after applying count density-based iterative correction. .	70
3.18	Absolute error between calculated and measured Z coordinates after applying the count density-based calibration to edge and interior pixels. There is generally improvement in the regions where the cathode/anode ratio performs the worst, while little to no additional errors are introduced in the good regions of the crystal.	71

4.1	Approximate cross section for daughter particle emission following pair production into a particular angle away from the incident photon direction. The photon here had 100 MeV of energy. At lower energies, the cross section is more complicated and the preference for forward emission weaker, but still present.	74
4.2	Approximate cross section for emission of a positron with a particular fraction of the total available kinetic energy, for a 10 MeV photon. At lower energies the cross section becomes more uniform.	75
4.3	SAS-2 pair production imaging detector assembly.	76
4.4	Diagram of EGRET instrument. [48]	77
4.5	Example electron charge cloud after thermalization from a Compton scatter depositing 3.64 MeV of total energy. Each circle marks the site of ionization of valence electrons to create electron-hole pairs, with the size of the circle indicating how much energy the primary particle lost to electron-hole pair generation. Simulated with Geant4.	79
4.6	Mean charge cloud centroid location as a function of total energy deposited, from Geant4 simulations. The initial electron was emitted at the origin in the +X direction. As the initial electron energy increases, the charge centroid is biased further in the +X direction. The variance in Y at higher energies comes from poorer statistics at higher energies, as well as an intrinsically greater overall variance in the final position.	80
4.7	Simulated charge cloud detail for a pair production event in which the positron had 17% of the total available kinetic energy. The pink arrow shows the initial photon direction, and the green arrow shows the lever arm constructed by the algorithm described above.	83
4.8	Simulated charge cloud detail for a pair production event in which the positron had 72% of the total available kinetic energy.	84
4.9	Average charge centroid displacement as a function of initial particle energy.	85
4.10	Simulated charge cloud detail for a pair production event in which the positron had 43% of the total available kinetic energy.	86
4.11	Charge cloud displacement from 511 axis vs. positron energy fraction for simulated pair production events from a 6 MeV photon beam.	86
4.12	Charge cloud size (defined as the maximum separation between any two energy depositions in the same cloud) vs. positron energy fraction for simulated pair production events from a 6 MeV photon beam.	87
4.13	Performance of the imaging algorithm for simulated pair production events.	89
4.14	Orion-Gamma detector used for pair production imaging experiments.	90
4.15	Spectroscopic performance of detector 6RID-36 in Orion-Gamma.	91
4.16	Experimental setups for run 1 (left) and run 2 (right).	91

4.17	PuBe spectra measured with Orion-Gamma.	93
4.18	2-interaction spectrum for experiment 1 and Compton images corresponding to several energy lines: 372, 413, 843 and 1014 keV.	94
4.19	2-interaction spectrum for experiment 2 and Compton image of the 413 keV peak.	95
4.20	Distribution of calculated lever arm lengths for experiment 1.	96
4.21	Example events observed in Orion-Gamma under irradiation by a PuBe source.	98
4.22	Pair production images reconstructed using experimental data from PuBe irradiations.	99
5.1	Orthorhombic phase of the perovskite CaTiO_3 , the original naturally occurring mineral known as perovskite. Dark blue are the Ca ions, light blue is Ti, and red is O. From materialsproject.org	103
5.2	History of perovskite testing at the University of Michigan.	105
5.3	Thin planar CsPbBr_3 sample grown at Northwestern University and tested at the University of Michigan.	106
5.4	^{57}Co spectra from a thin planar CsPbBr_3 detector.	107
5.5	^{137}Cs spectra from a thin planar CsPbBr_3 detector.	109
5.6	Details of YH-UM3 and initial ^{137}Cs spectrum.	110
5.7	Time behavior of YH-UM3.	111
5.8	^{137}Cs spectra from YH-UM3 after conditioning.	112
5.9	Refabricated detector YH-UM6R after it was shaved down to dimensions of 8.5 x 8.5 x 12 mm and new contacts deposited. The detector is encapsulated in parylene to prevent moisture intrusion.	113
5.10	Depth-corrected spectrum from all 9 pixels of YH-UM6R. All resolutions are below 2% FWHM, except for channel 9 which is degraded due to a DAQ system problem (not the crystal itself). The best resolution is from channel 5 (the central one) at 1.0%. Overall resolution is 1.4% from all pixels combined.	114
5.11	Per-pixel and overall corrected ^{57}Co and ^{137}Cs spectra from the doubly refabricated sample YH-UM-3-VK.	116
5.12	Detector IP21023-2-1(6), 12.1 x 11.8 x 10.4 mm.	117
5.13	Per-pixel and overall corrected ^{137}Cs spectra from the large crystal IP21023-2-1(6).	118
5.14	Example early solution-grown hemispherical detector and example waveforms.	119
5.15	$\text{FA}_{0.9}\text{Cs}_{0.1}\text{PbBr}_3$ solution grown pixelated detector from UNC. These samples had Bi pixels and Au planar contacts, and were biased negatively on the planar contact (electron collection mode).	119

5.16	I-V scan of the solution grown perovskite samples of Figure 5.15. (Left) entire range, (right) reverse bias region. Bias was applied to the planar Au contact. The sweep tracked from 0V to +50V, back down to 0V and -50V, then returned to 0V (blue to purple to pink in the figure).	120
5.17	Raw ^{137}Cs spectra from the 3x3 solution grown sample UNC-1. The Y scale is logarithmic to highlight the small photopeak and escape peak features in some pixels.	121
5.18	Drift time corrected ^{137}Cs spectra from pixels 8 and 9 of the solution grown sample UNC-1. Resolutions are 4.8% and 5.3% at 662 keV.	122
5.19	Bias and leakage current histories of YH-UM-3-VK and IP21023-2-1(6). These two furnace grown CsPbBr_3 samples from Northwestern have very different material and fabrication histories, but show similar leakage current vs. time characteristics.	123
5.20	Conditioning of IP21023-2-1(6) at 500V.	123
5.21	Pixel 3 signal amplitude vs. drift time over several time slices during the initial 500V bias period for IP21023-2-1(6).	124
5.22	Waveforms from IP21023-2-1(6) immediately as 500V bias is applied (blue) and after bias has been held for several hours (red). The stabilization of the baseline slope is evident.	124
5.23	IP21023-2-1(6) pixel 3 signal amplitude vs. measurement time once the bias is raised to 1000V. There is an initial period of reconfiguring, then the response stabilizes again.	125
5.24	Pixel 3 signal amplitude vs. drift time at different time slices after the bias is raised to 1000V. The trapping is significant for a while, then the crystal recovers.	126
5.25	Conditioning of the solution grown perovskite UNC-1.	127
5.26	Bias history of sample UNC-1. The conditioning phases can be seen as a jump in leakage current immediately after the bias is raised, which exponentially decays as the crystal relaxes. After the crystal is conditioned at -600V, the leakage current begins climbing at about 3.5-4 nA/day; this is separate from the initial conditioning phase and represents polarization of the crystal.	128
5.27	IV sweep of sample UNC-1 after several days under bias. The blocking behavior of Figure 5.16 has vanished. Inset: the pixels of UNC-1 upon arrival at UM (left) and after several days at bias (right).	129
5.28	Example depth-gain plots for CZT and furnace grown CsPbBr_3	129
5.29	Amplitude vs. drift time distributions at different biases from detector YH-UM6R.	130

5.30	Average photopeak waveforms from the bulk and anode side of YH-UM6R under 750V bias. An additional step is clearly visible after the main cloud is collected. The amplitude of the step is delineated by the red dashed lines, while the step delay is indicated in blue.	131
5.31	Anode-side e^- - h^+ conversion gain as a function of bias.	132
5.32	Echo signal delay versus photohole drift time at 750V bias in YH-UM6R. Linear fit line is also shown.	134
5.33	Photopeak event echo signal amplitude vs. electron drift time calculated from echo signal delay. Exponential fit is also shown.	134
5.34	^{137}Cs spectra from YH-UM6R under 750V bias after depth correction by drift time (top) and PCA corrections (middle). The best pixel (bottom) is also shown. Resolutions achieved are 2.4% after depth correction, 1.4% after PCA corrections, and 1.2% for the best pixel.	136
5.35	PCA-corrected spectra from pixels 8 and 9 of UNC-1. Resolutions are 2.8% and 4.8% at 662 keV, an improvement over the standard drift time corrections of Figure 5.18.	137
5.36	Relationship between drift time and principal component 2 for pixels 8 and 9 of UNC-1.	138

Abstract

Pixelated semiconductor gamma-ray detectors are a versatile tool for the characterization of radiation in a wide range of environments. The latest low noise, waveform sampling electronics coupled with innovative reconstruction algorithms allow for year-on-year improvements in energy and position resolution. Meanwhile, CdZnTe's unparalleled combination of high energy resolution and 3D position sensitivity enables many advanced algorithms for imaging and other applications. The development of radiation detectors based on new materials further ensures a bright horizon for this technology.

This dissertation presents new work covering all three aspects of pixelated semiconductor radiation detection in the Orion Group of the University of Michigan. New electronics development is presented which aims to condense all the front end control logic and some of the digital processing currently done by field programmable gate arrays and computers down onto a single silicon chip. This can enable significant further reductions in radiation detector size, weight and power while maintaining the high performance and imaging capabilities characteristic of CdZnTe.

Subpixel position reconstruction accuracy and precision was improved by a series of new algorithms. Machine learning models trained to predict subpixel positions from signal waveforms demonstrated improvement over the resolution that could be achieved with classical techniques. The radial position was predicted to less than 200 μm precision using only the collecting pixel waveform. Additionally, an improvement of 50 μm position resolution is demonstrated for the XY subpixel coordinates. Meanwhile, a range of algorithms are considered for correcting position distortions due to space charge, nonideal weighting potentials and other effects. An iterative technique is presented that successfully corrects some of the distortions in the cathode/anode ratio depth reconstruction.

A technique is presented to image pair production events in CdZnTe. This is typically accomplished with electron tracking. Here, a method to reconstruct images without

electron tracking is presented, which exploits asymmetries in pair production and charge induction physics. The imaging capability is demonstrated both in simulation and in experiment with the 4.44 MeV photons from a PuBe source.

Extensive work is done to characterize and improve the performance of perovskite-based radiation detectors. Thanks to waveform digitization and principal component analysis techniques, world-record performance of less than 1% full width at half maximum for the ^{137}Cs 662 keV line is achieved with pixelated CsPbBr₃-based detectors. Hybrid organic-inorganic solution grown perovskites are also demonstrated to work, with a resolution of less than 3% achieved for the same ^{137}Cs line.

Chapter 1

Introduction

Life is a state of constant radiation and absorption; to exist is to radiate; to exist is to be the recipient of radiations.

- William George Jordan

1.1 Radiation and Matter Interactions

The term radiation encompasses many forms of energy, both naturally occurring and human-made. It is composed of various particles, such as photons, electrons, protons, neutrons, muons, and nuclear fragments. These particles may be organized according to their identity, energy, polarization, mass, or electric charge.

The wide range of particles that fall under the umbrella of radiation follow a diverse range of physical processes. These depend on the nature of the radiation, the energy of the particle, and the type of matter it interacts with. Photons are oscillations of the electromagnetic field characterized by a wavelength and a polarization, and travel at the speed of light. Their energies range from radio waves, through the visible spectrum of light, up to gamma rays that can have many TeV of energy. They are generally characterized as ionizing or nonionizing based on their ability to strip electrons from atoms. This depends on the atom they interact with, but the boundary is generally set in the UV to x-ray range. This thesis will deal with ionizing photon radiation only, although the detectors described are sensitive to other forms of radiation as well. Ionizing photons will typically either be absorbed by an atom in photoabsorption, undergo Compton or Rayleigh scattering, or convert into a matter-antimatter pair. They may be produced




Figure 1.1: There are many types of particles that fall under the umbrella of radiation.

from a wide range of sources, from relaxation of atomic or nuclear energy levels to annihilation of antimatter or the acceleration of other charged particles.

Electrons (betas) may constitute ionizing radiation when freed from an atom and given sufficient kinetic energy. They typically scatter and thermalize in matter, and may produce photons through Bremsstrahlung radiation when their path is deflected. Their antimatter twins, positrons, undergo the same interactions, and may additionally annihilate with electrons to release 511 keV photons.

Neutrons come in two main varieties, thermal and fast. Fast neutrons have energies of keV scale or greater, while thermal neutrons have energies around the thermal energy at room temperature, $kT = 125$ meV. They may scatter off matter, or induce nuclear reactions. These reactions can produce further photons, neutrons, alphas, betas, and various fission fragments.

Other charged particles including alphas, protons, muons, and fission fragments interact similarly to betas, with the main differences coming from their masses. Heavier particles are stopped sooner in matter and deposit more of their energy in a narrow range, a phenomenon called Bragg peaking [1].

1.2 Semiconductor Radiation Detectors

Of the many ways radiation can interact with matter, the three most common modalities utilized for radiation detection are absorption, scatter, and pair production. Following each of these, some energy is deposited inside the radiation detector. This typically comes in the form of heat, light, or mobile electric charge. The detectors in this work are semiconductor detectors, which primarily operate by sensing the electric charge produced.

A semiconductor crystal is made of many atoms arranged in a regular lattice. The outermost valence shells of these atoms overlap, leading to a splitting of energy levels into a broad energy band called the valence band. The interaction of radiation with the crystal bulk can raise some electrons from the valence band to a higher conduction band. The amount of electrons freed is proportional to the energy deposited by the radiation. Electrons in the conduction band are delocalized and free to drift inside the crystal. They leave behind in the valence band an empty region of positive space charge called a hole, which acts as a quasiparticle with similar mass to an electron and opposite charge. Holes are also free to drift while inside the valence band.

If left alone, random thermal disturbances will eventually push a conduction band electron to collide with a valence band hole. The electron will fill the hole and restore local charge neutrality, becoming immobile once more. If an external electric field is applied to the crystal, however, the electrons and holes may drift in opposite directions and induce measurable signals on the attached electrical contacts. This is the principle of operation for semiconductor radiation detectors. Mobile electron-hole pairs generated by radiation quanta are drifted towards sensing contacts, and characteristics of the radiation are reconstructed from the observed electronic signals [2].

1.3 CdZnTe 3-D Position Sensitive Pixelated Detectors

1.3.1 Device Overview

Cadmium zinc telluride (CdZnTe, or CZT) is a ternary semiconductor with a bandgap of 1.6 eV. It has been used successfully for radiation detection by patterning the crystal

surfaces with metal contacts. Many contact geometries are used, but in this work the primary one is the pixelated arrangement shown in Fig. 1.2, with 121 anode pixels and a single monolithic cathode. When a CZT crystal is exposed to ionizing radiation, electron-hole pairs are generated inside and are swept out to the contacts by an externally applied electric field. As they drift, they induce signals on charge-sensitive preamplifiers connected to the various electrodes. These signals are stored and analyzed to reconstruct the details of the radiation interactions inside the crystal.

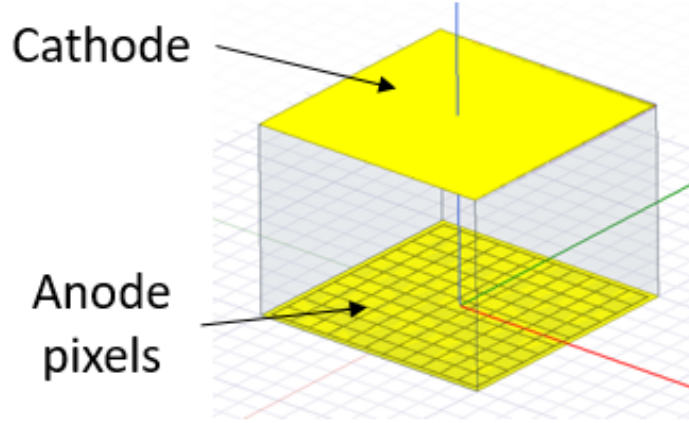


Figure 1.2: Metalization pattern of CZT detectors used in this work.

As the electrons and holes drift, they may encounter crystal defects that trap the mobile carriers and degrade the signal quality. Since hole trapping can be severe, detector geometries such as that of Fig. 1.2 are used that ensure the detection signal is primarily due to the movement of electrons. The electrons travel from the interaction location down to several of an 11x11 grid of anode pixels. As they drift, they induce signals on all electrodes of the device, though typically only a few anode pixels and the planar cathode have any significant contribution [3].

1.3.2 Shockley-Ramo Theorem and Single-Polarity Sensing

The signal induced on an electrode by a drifting charge may be calculated by integrating the electric field over a surface enclosing the contact:

$$Q = \oint_S \epsilon \mathbf{E} \cdot d\mathbf{S} \quad (1.1)$$

In practice, because the electric field \mathbf{E} is constantly changing as the charge drifts, the calculation is very difficult. The Shockley-Ramo Theorem (SR theorem) was developed to ease this computational problem. According to the SR theorem, the differential charge induced on a particular contact of interest is:

$$\frac{dq}{dt} = -q \frac{d\phi_0}{dt} \quad (1.2)$$

which may then be integrated with respect to either time or charge position to find the total signal induced. ϕ_0 is called the weighting potential, and is the electric potential that would exist inside the detector under the following conditions [4]:

- The electrode of interest is at 1 volt.
- All other contacts connected to power supplies are grounded.
- All charge is removed.

The weighting potential is determined purely by the geometry of the contacts and the uniformity of the detector's dielectric constant. It may be calculated once for every electrode of interest and then used to find the signal induced by a wide variety of possible radiation interactions.

Note that the weighting potential is often different from the actual electric potential. The real electric field must still be calculated to find the drift path of the charge, which determines the weighting potential profile the charge sweeps through to induce a signal. However, often only the final signal amplitude is of primary interest. The integral of eq. 1.2 has the same final value regardless of the drift path, and the final signal amplitude may be calculated as:

$$Q = -q\Delta\phi_0 \quad (1.3)$$

where $\Delta\phi_0 = \phi_0(\vec{x}_f) - \phi_0(\vec{x}_i)$ is the difference in weighting potential between the charge's starting point and the collecting electrode.

For the detector geometry shown in Fig. 1.2, the planar cathode on the $Z=15\text{mm}$ plane is typically biased at -3000V . Since the pixels and guard ring at $Z=0$ are held close to ground, the resulting electric potential is nearly linear. Example weighting potentials calculated using ANSYS Maxwell are shown in Fig. 1.3. The anode weighting potential is nearly flat in the bulk for $Z>3\text{ mm}$. Since the operating electric potential is linear,




Figure 1.3: Simulated weighting potential for the cathode (red) and a central anode pixel (blue) along the Z axis perpendicular to the anode plane.

electrons liberated by the interaction of radiation will drift downwards to the anode pixels. The term $\Delta\phi_0$ in eq. 1.3 will be very close to 1 for the electron contribution, resulting in a signal amplitude almost independent of depth. The holes will also induce a signal on the pixels as they drift towards the cathode. However, $\Delta\phi_0$ for the hole drift path will be close to zero, and the holes will make almost no contribution to the anode signal. Therefore, this arrangement of contacts gives rise to a signal determined almost entirely by the drift of electrons. This type of configuration is called single-polarity charge sensing. It is advantageous in CZT since the hole transport properties are often far worse than those of electrons [5].

1.3.3 Electronics

In the past, the signals induced by radiation were read out by analog electronics consisting of charge amplifiers, subsequent gain stages, and finally binning by a multi-channel analyzer. With the development of pixelated detectors with many tens of channels, discrete electronics became impractical. The preamplifiers, shapers, peak hold circuitry, and timing circuits were integrated onto application-specific integrated circuits (ASICs). These ASICs had >124 channels and could provide the signal amplitude and timing

from each of them in response to a trigger from a radiation interaction or an external command. Examples of such ASICs include the VAS-TAT [3] and BNL-H3D systems [6].

However, these two data points per channel failed to capture the full information present in the electron signal. To this end, waveform digitizing ASICs (called “digital ASICs”) were created. These store the preamplifier output directly onto banks of switched capacitor arrays. During readout, they present the stored samples sequentially as analog voltage levels, which are then digitized by an external analog-to-digital converter (ADC) and saved for further processing. In this way, the full signal waveform can be obtained. The first such ASIC used with the detectors in this work was the VAD_UM system [7]. The most recent iteration, VAD_UMv2.2 used in this work, presents 160 samples per waveform at sampling rates up to 80 MHz, and allows for dynamic ranges from 700 keV to 9 MeV per channel. A next-generation digital ASIC, the H3DD-UM ASIC, is under development [8]. An example of the waveforms recorded from a radiation interaction by the VAD_UMv2.2 ASIC is shown in Fig. 1.4.




Figure 1.4: Example anode pixel and cathode waveforms from a radiation interaction output from the VAD_UMv2.2 digital ASIC.

Initially, the CZT crystals were housed on separate carrier boards from the front end

ASICs, as in Fig. 1.5a. They were connected by a row of headers to form a full detector module. However, any additional capacitance on the preamplifier inputs increases the electronic noise. To minimize the trace length and lower the capacitance, the CZT crystal and front end ASIC were bonded to the same substrate board. An example of this direct attach module is shown in Fig. 1.5b.




Figure 1.5

In addition to the front end ASIC, a number of other electronics are required. These consist of ADCs to digitize the ASIC outputs, field-programmable gate arrays (FPGAs) to configure and control the front end ASICs, and a computer to process the saved waveforms, apply calibration parameters, and construct the final spectra and images.

1.4 Basic Digital Signal Processing and Event Reconstruction

1.4.1 Waveform Processing and Coarse Position Reconstruction

With the full waveform data obtainable from the latest digital ASICs, the process of event reconstruction from recorded signals can be further optimized. Digital filters can implement designs not possible with simple analog electronics and can be more easily tuned to achieve the lowest noise. Y. Zhu studied the optimization of digital filters

for pixelated CZT detectors, and arrived at the trapezoidal filter as a good balance between computational complexity and noise performance. Zhu also developed more advanced algorithms centered around the system response function concept, which can achieve better energy resolution than simple digital filters [9]. B. Williams developed an alternative reconstruction technique based on principal component analysis, which can achieve similar performance to system response function fitting at a lower computational cost [10].

A key aspect of 3D-position sensitive CZT detectors is their ability to determine the 3D position of each radiation interaction. Given that the electric potential is linear inside the crystal, a coarse XY position can be determined by examining which anode pixel collected any significant charge. The Z coordinate, along the axis between anode and cathode, can be determined in two ways. Both rely on the dissimilarity between the anode and cathode weighting potentials.

As discussed in sec. 1.3.2, the profile of the anode weighting potential leads to a signal that is almost independent of depth. The same is not true for the cathode. From Fig. 1.3, the cathode weighting potential is nearly linear through most of the bulk. The signal amplitude induced by electrons on the cathode is therefore a linear function of both the depth of interaction and the energy deposited. For the case of a single interaction event, taking the ratio of cathode to anode signal amplitude cancels out the energy dependence and yields a quantity proportional only to the depth:

$$Z = D \frac{S_C}{S_A} = D \frac{zE}{E} = Dz \quad (1.4)$$

where D is the full thickness of the crystal and z varies from 0 to 1. This only works for single interaction events.

For multiple interactions in the same crystal, the anode signal will typically be due to just one of the two interactions, while the cathode will have a contribution from every charge drifting in the crystal. The charge dependencies no longer cancel out, and eq. 1.4 cannot be used directly. Some information about relative depths may still be derived, but the ambiguity cannot be resolved entirely.

In this case, signal timing may be used. From eq. 1.2, the instantaneous charge induced on a contact depends on the time derivative of the contact's weighting potential a drifting charge experiences. Assuming a uniform operating field, the charge drifts at

constant velocity and $d\phi_0/dt$ is proportional to $d\phi_0/dx$, where x is the position of the charge. From Fig. 1.3, $d\phi_0/dx$ for the anode is close to zero in the bulk (Z between 3 and 15 mm), and rises swiftly closer to the pixel. Therefore, no significant signal is induced on the pixel until the drifting charge enters its vicinity. However, for the cathode $d\phi_0/dx$ is approximately constant everywhere in the bulk. The cathode signal begins to develop immediately as the charge begins drifting.

The time difference between the start of the cathode signal and the saturation of the anode signal gives the time for the electron cloud to drift from the interaction location to the collecting pixel. Assuming a uniform electric field, this is proportional to the depth of interaction. In many cases the field is not uniform, but even in that situation the depth calculated from eq. 1.4 is relatively unaffected for single interaction events. These can be used to calibrate the relationship between drift time and depth, which can then be used as a lookup table for multi-interaction events.

1.4.2 Subpixel Position Reconstruction

The coarse XY resolution is limited by the pixel pitch. In the CZT detectors used in this work, that is 1.72 mm. However, it is possible to achieve sub-pixel XY position resolution by considering the signals from neighboring pixels around those which collect the main charge [11]. For these pixels, the weighting potential is nonzero in the bulk but zero on the collecting anode face. Eq. 1.3 yields a negative amplitude which trends to zero for cathode-side events. In practice, there is a transient signal induced on the neighboring pixels which rises to a maximum when the drifting charge passes close to the pixel boundary. The closer the charge passes to the pixel edge, the greater the amplitude of this transient signal. Example neighbor and collecting pixel waveforms are shown in Fig. 1.6. For this event, the charge cloud passed close to the top-left of the central pixel. As a result, the transient signals of the top and left neighbors have greater amplitudes, and those of the bottom right neighbors have lower amplitudes.

The center of gravity of the transient amplitudes can yield the subpixel position of interaction:

$$S_x = \frac{S_R - S_L}{S_R + S_L} \quad (1.5)$$

where S_R and S_L are the sum of the transient signal amplitudes of the right three and left three neighbors, respectively. A similar expression exists for the Y axis. In

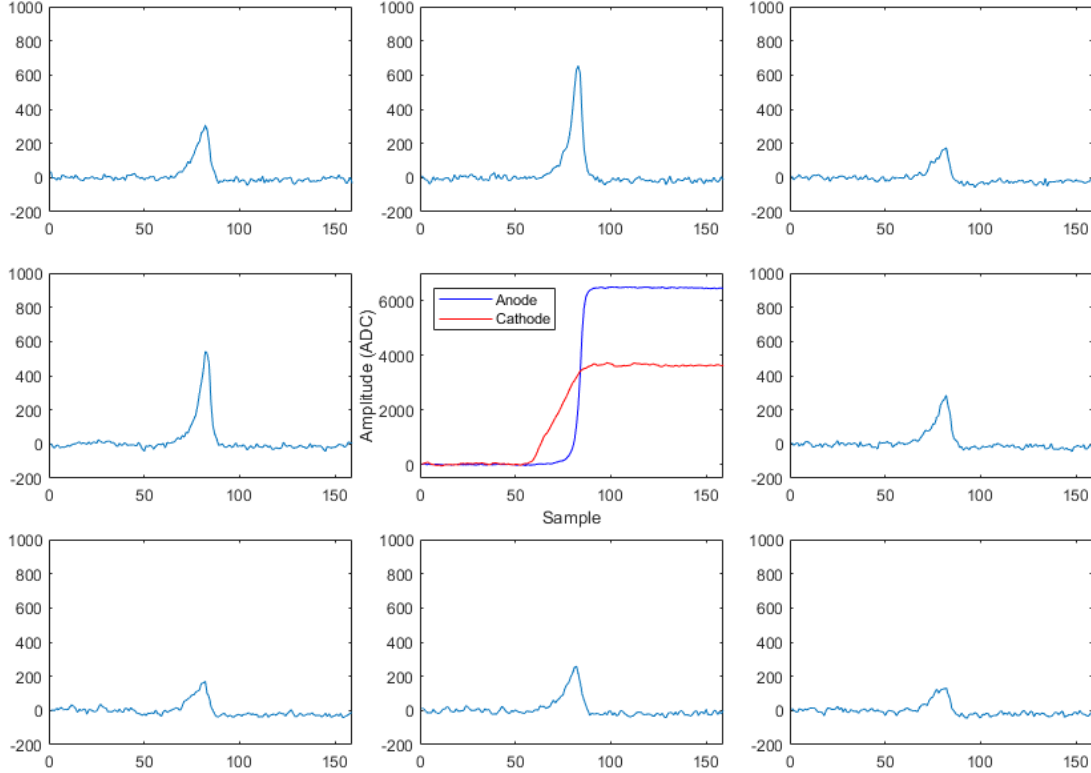


Figure 1.6: Collecting pixel, cathode, and noncollecting neighbor waveforms from a charge cloud which passed close to the top-left edge of the collecting pixel.

practice, the relationship between the calculated S_x and the true subpixel position is nonlinear. Additionally, subpixel reconstruction for edge and corner pixels, with fewer than 8 neighbors, is more complicated. Lastly, subpixel reconstruction for events with a charge cloud spread across multiple adjacent collecting pixels presents a further difficulty. Zhu developed a method to correct for these problems and reconstruct such events by comparison with simulated weighting potentials [12].

1.5 Imaging Gamma-Rays with 3D CZT Detectors

3D position sensitive gamma-ray detectors have the capability to reconstruct the 3D coordinates of each radiation interaction inside a detector volume. In addition to enabling position-dependent calibrations to improve energy resolution, this capability enables the

imaging of gamma rays by various means.

1.5.1 Coded Aperture Imaging

At low gamma-ray energies, photoabsorption is the dominant mode of interaction in CZT. If gamma-rays could be focused like optical photons, imaging could be done the same way using the same method as optical cameras: by directing rays from a specific focal distance onto an imaging plane and observing the cumulative distribution of interactions. Focusing of gamma-rays is possible with Wolter mirrors [13] or Laue lenses [14], but is not typically practical for terrestrial-based or wide field-of-view observations. Instead, aperture imaging is utilized.

The simplest aperture is the pinhole camera. The principle is illustrated in Fig. 1.7. Light scattered or emitted by an object is streamed towards a detector. On the way, an opaque mask blocks all radiation except that which passes through a small pinhole opening. The resulting distribution of counts on the imaging plane of the detector is an inverted representation of the object.




Figure 1.7: Operating principle of a simple pinhole camera.

The open fraction is defined as the ratio of the open areas in the mask to the total

mask area. Since a pinhole blocks all light except in a very small region, its open fraction is very small, and few photons pass through to strike the detector, lowering the overall efficiency. In the case of optical cameras, photon fluxes are high enough that poor efficiency is not a significant drawback. For ionizing radiation imagers, fluxes are typically far lower and interaction probabilities may be significantly less than unity. The severe attenuation caused by a pinhole mask is a more significant drawback in this case. The coded aperture mask was developed to remedy these shortcomings.

A coded aperture mask operates on a similar principle to a pinhole mask, but with a larger open fraction [15]. There are many possible mask designs, but a popular choice is based on modified uniformly redundant arrays (MURAs). One such MURA-patterned mask is shown in Fig. 1.8. The MURA has the property of rotational anti-symmetry: under a 90° rotation, the empty spaces are replaced with filled spaces and vice versa. This property can be used for mask-antimask imaging to suppress background artifacts. Additionally, it has a 50% open fraction, maximizing photon flux received by the detector while maintaining an acceptable signal-to-background ratio.




Figure 1.8: MURA mask made of tungsten, for a single CZT module.

A radiation source will emit photons towards the MURA, some of which will pass through the openings while others will be absorbed. A shadow of the MURA pattern will be cast on the crystal (Fig. 1.9a). The position of the shadow will shift based on the angle of the incident radiation. By cross-correlating the observed mask shadow with the proper decoding pattern (eq. 1.6), the incident source position can be reconstructed, as in Fig. 1.9b.

$$F = S \star G \quad (1.6)$$

Where F is the source image, S is the distribution of recorded counts, G is the decoding pattern (chosen so it convolves with the mask pattern to a delta function), and \star is the cross-correlation operator.

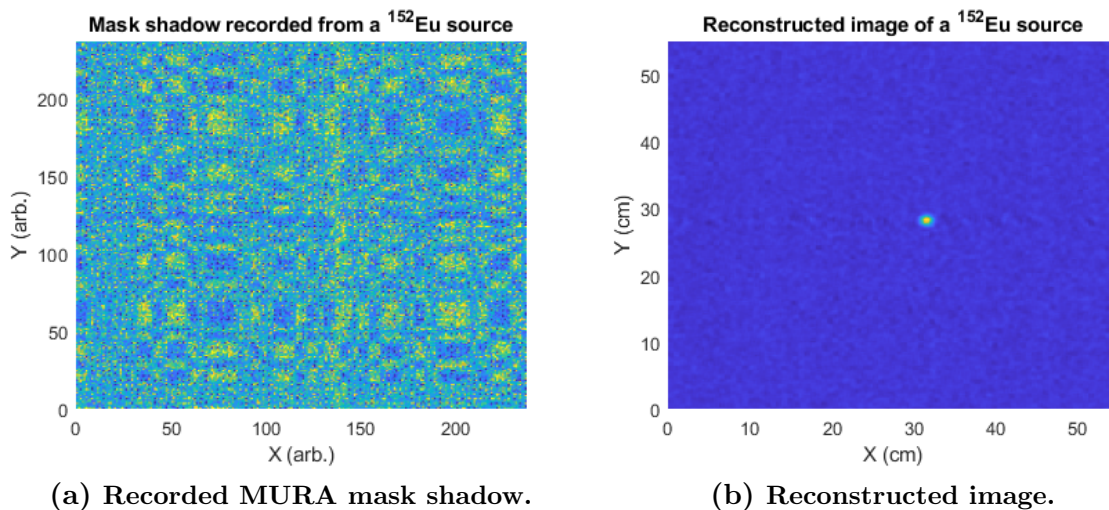


Figure 1.9: Figure courtesy of Alexander Rice.

Coded aperture systems typically have a limited field of view less than 4π , caused by the requirement for a source to be behind the mask relative to the detector. Sources outside the fully coded field of view will cause artifacts in the reconstructed image. Also, the resolution of the MURA image is limited by the spatial resolution of the detector. A way to overcome this limitation is to translate the MURA mask in time, resulting in temporal as well as spatial coding of the observed radiation.

1.5.2 Compton Imaging

At medium energies above several hundred keV, photons will more often penetrate through the closed elements of a MURA mask or scatter rather than photoabsorb in the CZT crystal. As a result, coded aperture imaging becomes increasingly impractical. A different imaging modality becomes possible, based on the kinematics of Compton scattering.


When a gamma ray Compton scatters in CZT, it does so according to the Klein-Nishina cross section. This leads to the Compton scattering equation, which gives the angle at which the photon scattered [16]:

$$\theta = \arccos\left(1 - \frac{m_e c^2 E_1}{E_0 E'}\right) \quad (1.7)$$


where θ is the photon scattering angle, E_0 is the initial photon energy, E_1 is the energy deposited in the Compton scatter, E' is the energy of the photon after scattering, and m_e is the electron mass. 3D position sensitive radiation detectors are capable of measuring the location of the Compton scatter. If the scattered photon then photoabsorbs, a lever arm axis can be drawn from the absorption to the scatter location. The gamma ray can be said to have been incident from somewhere in 4π along a cone centered on this axis and with an opening angle of θ . Knowledge of the recoil electron direction can further restrict the range of possible incident angles to a segment of the cone. If a second photon is detected from the same source, its Compton cone will overlap with that of the first. With sufficient statistics, the intersection of a large number of cones will reveal the distribution of radiation around the detector. Examples of backprojected cones and the final image are given in Figs. 1.10a, 1.10b, and 1.10c.

The 3D CZT detectors described in this work can function as both scattering and photoabsorption media. As such, they are intrinsically capable of Compton imaging radiation sources in a full 4π field of view. Image artifacts may stem from incorrectly identifying the sequence of Compton scatter and photoabsorption events. Multiple scatter events may be used to improve statistics, but the sequencing problem becomes correspondingly harder.


Photons at higher energies above 1.022 MeV may convert into electron-positron pairs instead of scattering. These events cannot be used for Compton imaging and lead to



(a) Backprojection of a single Compton cone.



(b) 500 backprojected Compton cones. The source hotspot can begin to be localized by the overlap of several cones.



(c) Compton image reconstructed by the overlap of many cones in 4π .

Figure 1.10: Compton imaging of the 372 keV line from a PuBe source.

further artifacts. Treatment of these events is addressed in Chapter 4.

1.6 Overview of Thesis and Contributions to Science

In the Orion Group at the University of Michigan, there are several informal subgroups or specialties. They consist of:

- The Detector Group, with two particular focuses:
 - Development of hardware for 3D CZT detectors
 - Development of event reconstruction algorithms to determine information about the radiation interactions in CZT from the signals recorded by the electronics
- The Imaging Group, which focuses on image construction from the event data provided by the detector, and applications for imaging spectrometers in various fields
- The Alternative Materials Group, which investigates other semiconductors which show promise and may eventually come to replace CZT in certain applications

Echoing this structure, this work is divided into four chapters, covering all the subgroups' different interests. Each focuses on a particular research area:

- Chapter 2 describes work to create a digital signal processing (DSP) ASIC. This ASIC aims to integrate on a single piece of silicon the digitization, control logic, and processing required to interface with front end detector modules and derive energy and position information from the raw signals. Particular focus is given to the DSP module and its structure and function. Simulations are done to verify its behavior, and it is incorporated in a first version of the chip synthesized in the TSMC 65nm process node.
- Chapter 3 describes work on event reconstruction with a focus on subpixel position calibrations. Machine learning is used to derive a model to reconstruct radial subpixel position from collecting waveforms, as well as improve XY subpixel position resolution. The chapter also addresses charge steering problem, in which nonuniform electric fields in the bulk lead to distorted estimates of the interaction location. Several techniques to build a calibration correcting for these effects are presented and explored.

- Chapter 4 deals with the imaging of high energy photons that pair produce inside the CZT crystal. Other detectors that image such events rely on electron tracking capabilities. These capabilities have been demonstrated in CZT, but only on a rudimentary level insufficient for image reconstruction. Exploiting the physics of pair production and the charge induction process, a technique to image pair production events without requiring electron tracking capability is presented. The algorithm is verified both in simulation and in experiment.
- Chapter 5 presents results and analysis of detectors based on perovskite materials, primarily CsPbBr₃ but including some hybrid perovskites as well. These semiconductor crystals offer the promise of significantly lower cost compared to CZT, among other benefits. However, they struggle with crystal stability issues. In this chapter, world-record energy resolution is demonstrated from this family of detectors for the 662 keV line from Cs-137. Conditioning effects are explored in furnace and solution grown crystals. A novel effect is discovered whereby electrons reaching the anode release a cloud of holes, providing a second chance to measure the signal amplitude and enabling characterization of both carrier drift properties from a single measurement.

The thesis concludes with a retrospective of work accomplished, a summary of possible future paths to take from the work presented, and thoughts about the future of 3D position sensitive radiation detectors.

Chapter 2

Digital Signal Processing ASIC

The obsidian flake and the silicon chip are struck by the light of the same campfire that has passed from hand to hand since the human mind began.

- George Dyson

2.1 ASIC development and signal processing

2.1.1 Digital Data Processing

Computers, as machines designed to aid in computations, have been a part of human life since antiquity. They may be classified in several ways (Figure 2.1) [17][18][19][20][21]. Application specific computers are designed to solve particular problems. They may have some flexibility in configuration by changing some problem parameters, but are typically limited in scope. By focusing on addressing a particular problem space, they derive other benefits such as faster calculation or lower power consumption. General purpose computers can be configured to solve a much wider array of problems, but often at the cost of slower computation or greater power requirements. Analog computers utilize continuous variables and accomplish tasks by modeling specific systems. Digital computers use discrete units to represent concepts and perform calculations to desired precision.

Well into the 20th Century, the slow speed, bulk, and power requirements of digital computers limited them to simple calculating aids or research curiosities. More complex functions, such as early synthetic aperture radar or modeling of ocean tides [22] and

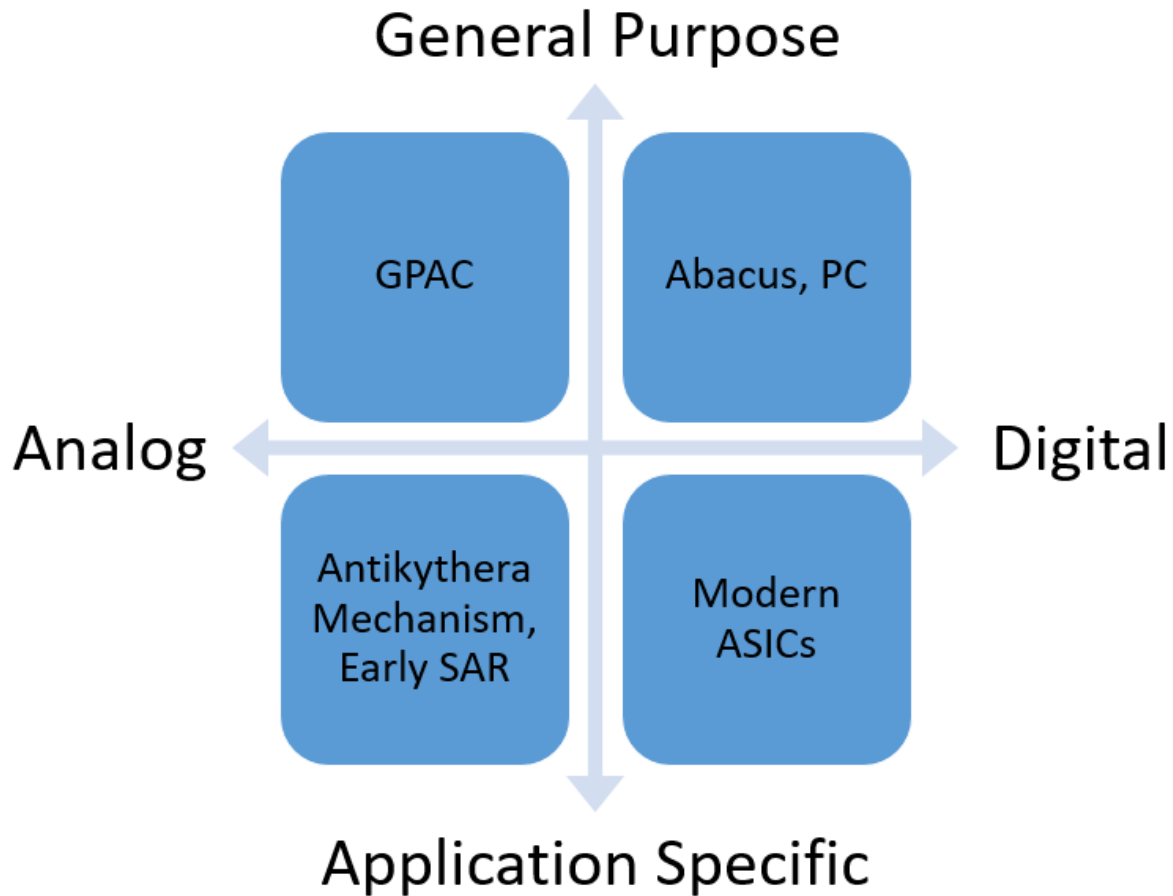


Figure 2.1: Examples of computers that fall into analog/digital and application specific/general purpose categories.

national economies [23], were carried out by application specific analog computers. The advent of all-electronic integrated circuits (ICs) in the mid-20th Century began to change this. What previously required the switching and management of hundreds of punch cards could be done instead with the transfer of electrons; what required a room full of power-hungry vacuum tubes could be shrunk down into a desktop assembly of silicon or germanium chips operating at a fraction of the voltage. As the technology further developed, the increasing speed and ease of interfacing with digital circuits drove them to dominate more and more of the data processing application space.

2.1.2 Integrated Circuits

The first all-electronic computers relied on vacuum tubes for logical operations. Though faster than punch cards, these machines were still bulky and consumed significant amounts of power. Efforts in the 1940's onwards to consolidate operations in a single piece of silicon culminated in the invention by Jack Kilby of the integrated circuit in 1958 [24], for which he was awarded the 2000 Nobel Prize in Physics.

Modern computers composed of integrated circuits use the transistor as the building block of logical operations. This is a three- or four-terminal device which controls the flow of electrical current through two terminals, the source and drain, by the application of a voltage on a different terminal, the gate. By switching between open and closed states, the transistor can represent the fundamental digital units of 1 and 0. By arranging together many of these transistors in particular ways, specific or general logic operations can be carried out.

2.1.3 Digital Logic Design Flow

Modern digital logic ICs contain hundreds of thousands to even billions of transistors. These are far too many to arrange by hand, so a complex suit of design software has been developed over the decades to enable the design of these large systems. Figure 2.2 shows the general flow of modern digital logic design [25]. The functional design is described in a hardware description language (HDL) such as VHDL or SystemVerilog. Rather than explicitly describing transistors, HDLs capture the structure and function of digital modules. Synthesis tools analyze HDL descriptions and use them to generate register transfer level (RTL) netlists: schematics which specify the arrangement of transistors, wires, and other components necessary to implement the described behavior. Automatic place-and-route (APR) tools combine synthesized netlists with predefined technology libraries that provide building blocks of transistors implemented in a particular technology node. They arrange these blocks, called standard cells, on a lattice of power rails and perform routing and insertion of specialized structures such as clock trees. At both the synthesis and APR stage, timing analysis is done to ensure all timing requirements are met.

This work used VHDL and SystemVerilog for top-level system design and behavioral simulations of the various digital components of the DSP ASIC. Design Compiler from

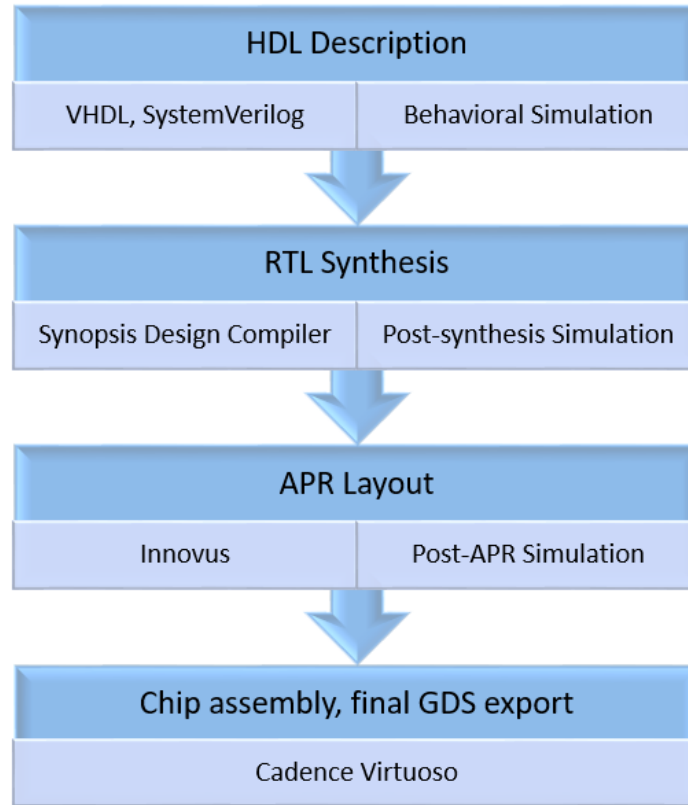


Figure 2.2: Design flow for digital logic. Starting from HDL descriptions, through to RTL synthesis and layout, and system simulation and verification.

Synopsis [26] was used for digital logic synthesis, in conjunction with the ARM library of standard logic cells implemented in the 65nm technology node from Taiwan Semiconductor Manufacturing Company (TSMC). Innovus was used to perform APR. Final violation fixing, integration with mixed signal circuitry, and incorporation within an input/output ring structure and seal ring was done with Cadence Virtuoso. Design rule checking and layout vs. schematic verification was done with Calibre DRC and LVS.

2.2 DSP ASIC Overview

2.2.1 Current hardware description

Section 1.3.3 briefly described the hardware used by current state of the art CZT detectors in the Orion Group. A more complete description follows here.

The broad system diagram is given in Figure 2.3. There are several components. The detector modules are composed of crystals of CZT with 121 anode pixels, a guard ring, and a cathode. These are each directly attached to a VAD_UMv2.2 waveform digitizing ASIC through an interposer board. High voltage is provided through an onboard high voltage generator utilizing a step-up transformer and Cockroft-Walton ladder. The output from the front end ASICs is digitized by 1-3 ADCs. Interfacing with the front end ASICs, packaging of the waveform data, and control of peripheral systems is handled by a FPGA. The significant heat generated by the ASICs and FPGA is dissipated by a thermal management system consisting of a cold finger, Peltier cooler, heatsink, and fans. The data packaged by the FPGA is sent by Ethernet to an external computer for processing. The computer performs some waveform preprocessing, digital filtering, extraction of amplitude and timing parameters, and application of the complicated calibration procedure. The final result is the energy deposition and 3D coordinates of each interaction in the crystal array. The computer may then use this information for downstream applications such as imaging.

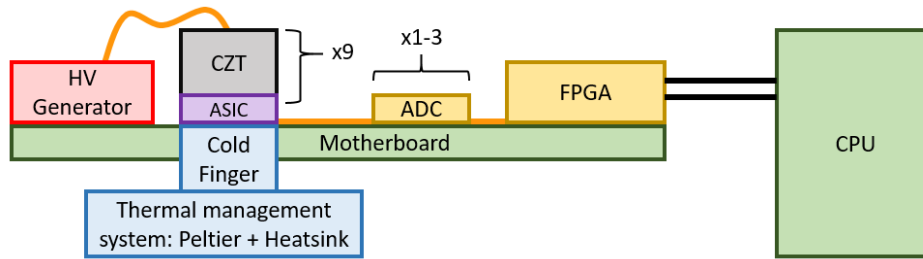


Figure 2.3: Current hardware layout of the Orion 3D CZT detector array systems.

2.2.2 DSP ASIC Goals

The system of Figure 2.3 is bulky and draws considerable power, on the order of several watts depending on CPU and thermal load. Progress has been made in reducing both

the size and power requirements by developing algorithms to correct the effects of temperature drift. These allow for Peltier-free operation, which accounts for a significant fraction of the bulk and power draw. However, the performance is still degraded by the excess heat, and the system still needs the external CPU for digital filtering and other waveform operations.

The DSP ASIC project began with the goal of incorporating all the major functionality of the ADCs and FPGA, and some of the functionality of the CPU, onto a single silicon chip. This would eliminate the ~ 5 watts drawn by the FPGA. A CPU would still be required for final analysis, but migrating the preprocessing, filtering, and pickoff onto a dedicated chip would enable the use of a far smaller, less power-hungry CPU and free cycles for the implementation of more advanced algorithms such as imaging. A diagram of such a system is given in Figure 2.4.




Figure 2.4: Diagram of a proposed detector system based on new DSP ASIC.

2.2.3 DSP ASIC Top-level Chip Architecture

The first version of the chip was designed according to the architecture in Figure 2.5. There are several functional blocks. For the first iteration, the chip is only designed to interface with one front-end ASIC. For simplicity of control and communication, the chip is designed to interface only with the next-generation H3DD-UM waveform digitizing front end ASICs, rather than the VAD_UMv2.2 chips currently in use in the 3x3 array Orion systems. The H3DD-UM chips have a much simpler I/O interface compared to the Norwegian ASICs, reducing the number of I/O pads required on the DSP chip.

The DSP chip consists of an ADC, a DSP core, a H3DD-UM control module, and supporting logic. The supporting logic consists of a main controller (MC), a clock control module (CCM), an event parser and an ADC parser, and a transceiver for

handling communications with an external CPU over SPI. Not shown is an additional monitor module, which multiplexes out several critical internal signals for debugging and verification purposes. The flow of control instructions and signals is shown in blue, and the flow of waveform event data is shown in green.

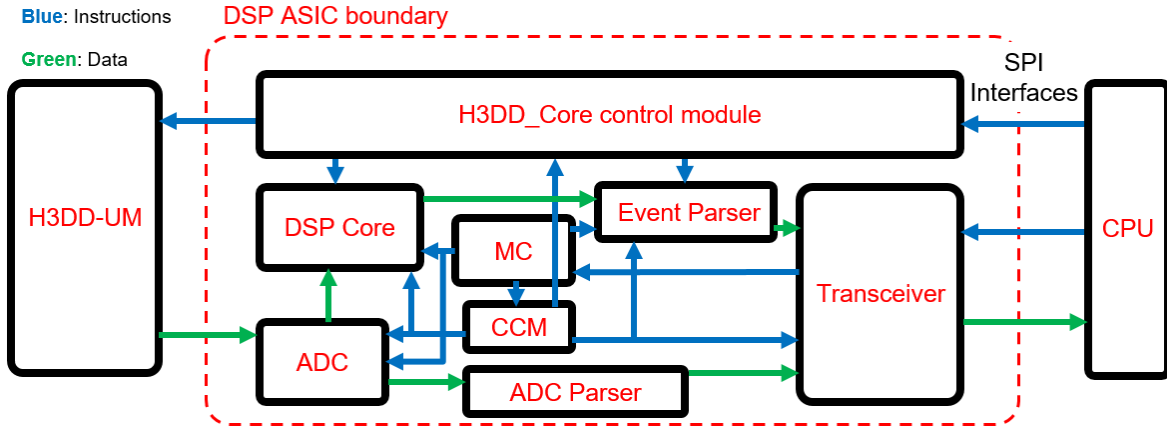


Figure 2.5: Top-level architecture of the first version of the DSP ASIC.

The main controller is responsible for coordinating the overall operation of the digital core, with the exception of the H3DD_Core which operates independently. It fetches communicated instructions from the SPI transceiver, parses them, and reconfigures the chip internal state accordingly. The transceiver contains a first-in-first-out (FIFO) buffer that enables the clock domain crossing between the core master clock and the external SPI clock. The chip is designed to operate as a SPI slave relative to an external master. Also included in the supporting circuitry is a clock control module that provides the many different clocks required in the chip. The H3DD_Core itself requires a couple of clocks of different frequencies, and generates several more. This stems from the requirements of the front end H3DD-UM ASIC, which may use different clock speeds for configuration, waveform sampling, and readout. The DSP requires a different clock from the master clock, since it is designed to synchronize with the front end readout. The ADC also requires an independent clock since its phase and duty cycle may need to be tuned for optimal performance. The CCM is a programmable block that generates all these clocks from the single chip master clock. There are also two parsers, related to SPI transactions. The event parser is used during normal operation, and packages data output by the DSP core into 32-bit wide packets for SPI transfer. The ADC is a debug

tool which captures data from the ADC independent of the DSP core, packages it into 32 bit packets, and sends it out over SPI for verification.

The H3DD_Core is an independent unit based heavily on VHDL currently implemented on FPGAs in the Orion systems. Its main task is to act as an intermediary between an external CPU and the front end ASICs and motherboard peripherals like the high voltage generator and Peltier. For the DSP chip, only the front-end related functionality is included, with high voltage and other peripheral control delegated off-chip. It is responsible for configuration of the front end ASICs, performing the mapping between ASIC channel number and pixel XY coordinates, and running the data acquisition and readout processes. Since this module was already implemented on FPGAs as an independent unit, that independence was preserved when migrating to silicon. The on-chip H3DD_Core has its own dedicated SPI input channel and front end IO pads. The only connections to the rest of the chip are clock inputs from the CCM and control inputs to the DSP core and event parser.

The ADC is a two-stage pipelined SAR ADC with 13 bits of output. It achieves greater than 11.5 effective number of bits, with $\sim 600 \mu\text{W}$ of power consumption and a $350 \times 200 \mu\text{m}$ area. It is designed to operate at 12.5 MHz sampling rate, matching a faster commonly used readout speed of the H3DD-UM front end.

The DSP core is the primary subject of this chapter, and is described in detail in the following section.

2.3 DSP Architecture

2.3.1 DSP Architecture Overview

The DSP core must achieve the following objectives:

- Correct the input waveforms' baseline offset
- Deconvolve the preamplifier decay of the input waveforms on a per-channel basis
- Apply several different digital filters based on whether the waveform is from a collecting pixel, neighbor pixel, or cathode channel, as well as whether it will be used to pick off event amplitude or timing

- Pick off the amplitude of waveforms after a slow digital filter
- Pick off the timing of waveforms after a fast digital filter, including interpolation for sub-sample precision
- Discriminate between collecting pixel, neighbor pixel, and cathode waveforms
- Perform all the above functions with sufficient bandwidth to not present a bottleneck to the readout process in the worst case of maximum H3DD-UM readout speed and waveform content
- Output channel, amplitude, timing, and signal identity data to the supporting logic and signal when processing begins and finishes
- Maintain programming flexibility to reconfigure the baseline lengths, decay and filter coefficients, amplitude and timing thresholds, and other runtime variables to support a wide range of detector types and operating conditions

The requirement to not increase the dead time of the overall system can be met in a simple way by running the DSP at the same clock speed as the readout and ensuring data passes through it at the same rate. The waveform samples are output serially, at the rate of one sample per clock cycle, so there is only one datastream input to the DSP. If the required processing can be accomplished at the same rate, there will not be any increase in the dead time.

Some tasks such as the application of different digital filters, or amplitude and timing pickoff, can be accomplished in parallel. However, some other tasks rely on previous ones to be complete. The decay deconvolution requires baseline subtraction to be done to avoid integrating a slope into the waveform, for example. The digital filters assume the signal to be zero prior to the start of filtering, which also requires baseline subtraction. Baseline subtraction itself presents a problem: a minimum of several samples are needed to determine the baseline in order to subtract it, but the subtraction needs to be applied starting from the very first sample.

To overcome these difficulties, a pipelined structure is designed as in Figure 2.6. The main functions are divided up into three stages. A brief summary of each stage is presented here. Greater detail is presented in the next three sections.

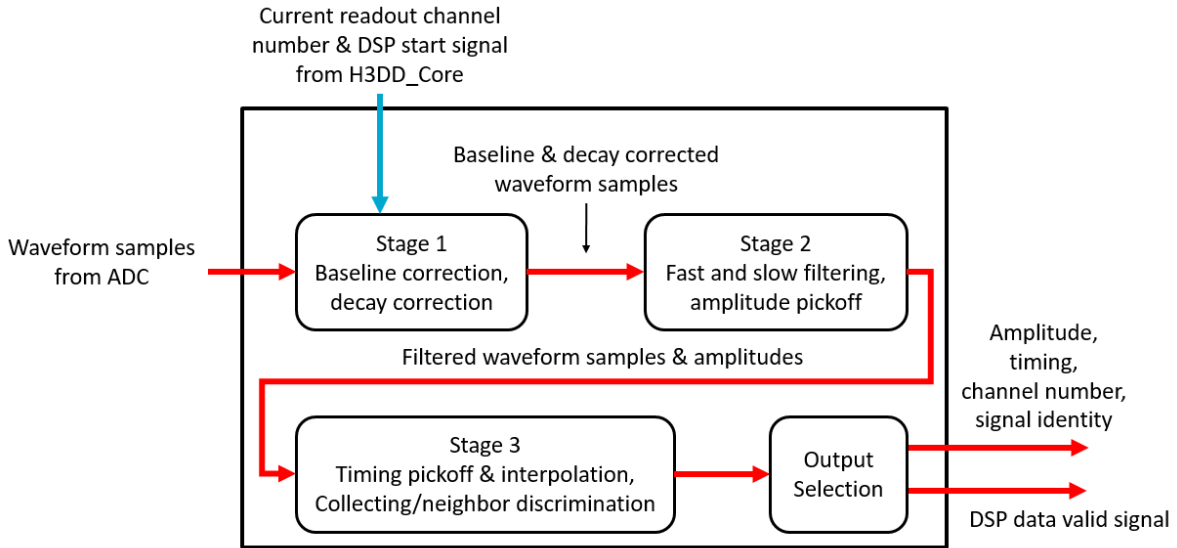


Figure 2.6: Top level DSP core block diagram.

Stage 1 is responsible for calculating the baseline of incoming waveforms. It does this using a rolling average scheme, where the current baseline is averaged together with the previous 127 events from the same channel to obtain a lower noise estimate of the true channel baseline. At the same time, the previously calculated baseline average is used to correct the offset in the current input samples. A decay correction is done here as well. Stage 2 contains the bulk of the digital filters. These implement fast and slow filters for timing and amplitude pickoff, respectively. Amplitude pickoff is done in parallel as the filters operate by a digital peak-hold mechanism. Fast filtered waveform samples and amplitudes are passed on to the third stage. Stage 3 performs a coarse timing pickoff, then a fine 8x linear interpolation. It also performs a discrimination between collecting and noncollecting signals, used to inform the final signal selection. Different processing is required based on whether the signal being analyzed is from a collecting, neighbor, or cathode electrode. The decision between collecting/neighbor can only be made after some processing is already accomplished, so in the interest of maintaining high bandwidth, all processing is done for all waveforms, and the appropriate output is selected only at the very end by a multiplexer.

The waveforms from the H3DD-UM front end ASIC may be of variable length up to 256 samples. They are output at the rate of 1 sample per clock cycle, so a waveform of

length N will take N cycles to fully read out. Each stage is designed to also finish its task in N clock cycles, as in the timing diagram of Figure 2.7. Stage 1 begins operating immediately as the samples arrive, while later stages stay idle. Once stage 1 finishes processing a single waveform, its output is clocked in parallel into a pipeline buffer that acts as an input for stage 2. Stage 2 then begins its operation, while stage 1 resets and begins processing the next waveform. In this way, several waveforms may be in the pipeline simultaneously, increasing bandwidth and preventing a rise in downtime.

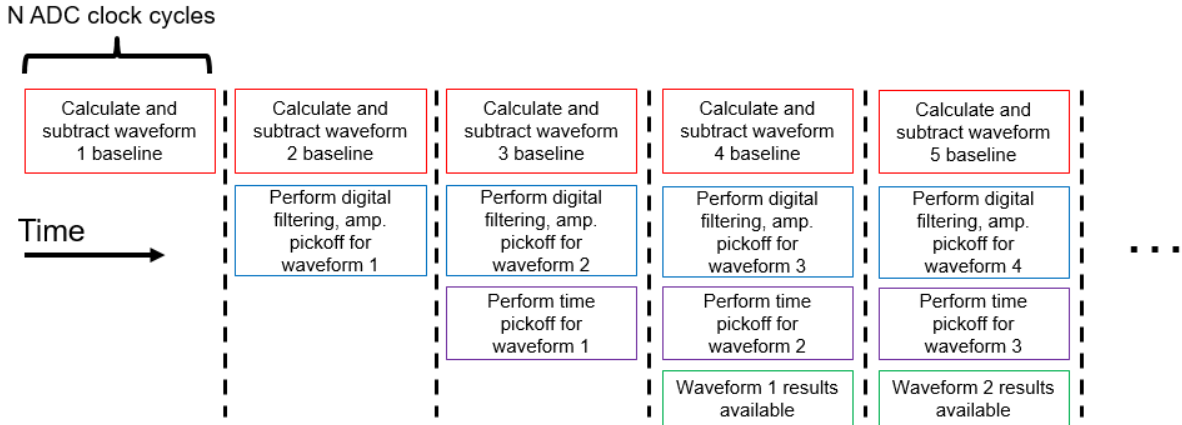


Figure 2.7: Timing diagram of the operation of the pipelined DSP.

2.3.2 DSP Stage 1

The task of stage 1 is to calculate and subtract the baseline from the input waveform, and correct the exponential decay of the waveform from the front-end preamplifiers of the H3DD-UM ASIC. The structure of stage 1 is shown in Figure 2.8. There are three main tasks: baseline calculation, baseline subtraction, and decay deconvolution.

The baseline is a constant offset which varies per channel and must be measured and subtracted. This offset comes mostly from the electrical characteristics of a particular front-end channel, in the absence of low frequency noise. To obtain the best results with the current version of the H3DD-UM front end ASIC, it's necessary to get a high accuracy estimate of the true channel baseline by averaging the baseline over many events. It may shift with environmental factors like temperature over the course of the measurement, so a simple pre-calculated channel baseline will not suffice. Instead, a rolling average baseline of the past 128 events per channel is used.

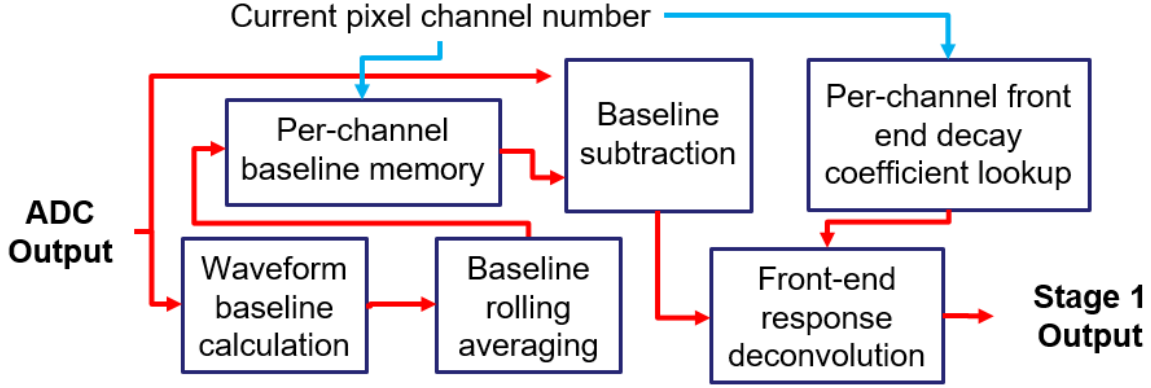


Figure 2.8: Detailed block diagram of DSP stage 1.

The baseline must be subtracted from samples immediately as they arrive, but several samples are first required to calculate it. Stage 1 incorporates a memory block which stores previously calculated per-channel baseline averages. A lookup from this memory is used to correct samples as they arrive, which are also used to update the rolling average.

A sample skip feature is supported, where a programmable number of samples at the beginning of a waveform will be ignored before baseline calculation begins. The event baseline itself is calculated as the average of a programmable number of samples, either 8, 16, 32, or 64:

$$B_i = \sum_{n=skip}^{skip+N_B} \frac{w[n]}{N_B} \quad (2.1)$$

where B_i is the baseline value of the i^{th} event, $skip$ is the start point for the average calculation, w_n is the n^{th} waveform sample, and N_B is the baseline length setting. The rolling average of the past 128 events could be calculated similarly, but to avoid the layout cost of memory to store 128 baseline values for each channel of each ASIC, only the cumulative rolling average is stored. To avoid loss of precision, the bitlength of the stored value is extended:

$$StoredValue_i = 127 \cdot StoredValue_{i-1} + 8 \cdot B_i \quad (2.2)$$

$$AvgBaseline_i = \frac{StoredValue_i}{8} + StoredValue_i(MSB - 8) \quad (2.3)$$

where the extra addition term in 2.3 is a rounding step to avoid lock-in during rolling averaging. The operation of the rolling average calculation module is shown in Figure 2.9




Figure 2.9: Progression of the rolling average value for a single channel after many input waveforms. From the reset value of zero, it gradually converges to the true value. Fixing the reset value close to the expected baseline reduces the time required for convergence.

The front end preamplifiers of the H3DD-UM ASIC have an exponential decay, with a transfer function of $H(z) = 1/(1 - z^{-1}e^{-\frac{1}{\tau_s}})$, where τ_s is the decay parameter. To correct for this decay, the waveform must be convolved with the inverse transfer function times that of an ideal preamplifier with infinite decay. In the z-domain:

$$F(z) = \frac{1 - z^{-1}e^{-\frac{1}{\tau_s}}}{1 - z^{-1} \cdot S(z)} \quad (2.4)$$

where $S(z)$ is the input signal and $F(z)$ is the corrected signal. This is implemented as an infinite impulse response filter (IIR, or recursive filter) in discrete logic:

$$y_i = \frac{1024 \cdot x_i - (1024 \cdot a)x_{i-1} + 1024 \cdot y_{i-1}}{1024} \quad (2.5)$$

where x_i is the input signal at sample i , y_i is the corrected signal, and a is the discretized decay parameter. IIR filters are the most efficient way to implement some complicated transfer functions, but are susceptible to stability issues due to feedback if there is insufficient precision in the calculations [27]. The factors of 1024 in equation 2.5 avoid loss of precision that would otherwise accumulate as the filter operated. The decay coefficients are also per-channel variables, and are fetched from an on-chip programmable memory.

Both the rolling average baseline correction and the decay deconvolution involve division, among the most costly of basic arithmetic operations to implement in digital logic. To simplify the design, all divisions in the DSP chip are by powers of 2. This enables the division to be accomplished by simple bit shifting. However, in recursive processes like some filtering and the rolling average calculation, errors that come from truncation during bit shifting can cause lock-in of the final results to incorrect values. For this reason, the bit shifting is modified to round 0.5 values and higher up to 1, by adding back the most significant truncated bit.

The operation of the decay deconvolution module is shown with a baseline-corrected H3DD-UM waveform from the baseline subtraction module in Figure 2.10.

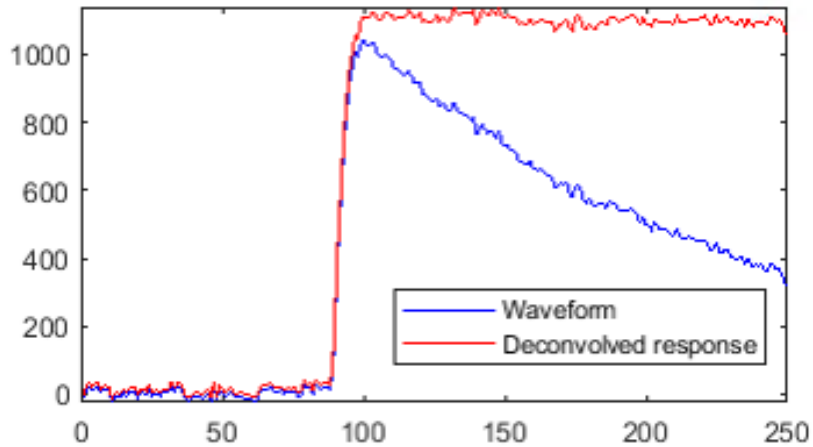


Figure 2.10: Operation of the front-end decay deconvolution module on an experimental H3DD-UM waveform output from the rolling average baseline subtraction module.

2.3.3 DSP Stage 2

The baseline and decay corrected waveforms are output from stage 1 through pipelining registers to stage 2. The task of this stage is to perform the fast and slow digital filtering and pick off the amplitude of the waveform after each different type of filter. The final amplitudes of the fast and slow filtered waveforms, and the actual samples of the fast filtered waveforms, are output to stage 3 for timing pickoff. Figure 2.11 gives an overview of stage 2 internals.




Figure 2.11: Detailed block diagram of DSP stage 2.

The filters are implemented as 256-tap finite impulse response (FIR) filters. These consume more power and area resources than IIR filters. They have the benefit of being intrinsically stable by not incorporating any feedback path, as well as being easily reconfigurable to implement a wide range of transfer functions by reprogramming the tap coefficients [28]. A FIR implements the function

$$F = C * S \quad (2.6)$$

where C is called the filter kernel, and each element of C is called a tap. In hardware, they are implemented as shown in Figure 2.12. Stage 2 comprises of six FIR filters. They can each be independently programmed to any 256-tap transfer function, but typical usage involves one set of trapezoidal and $CR-(RC)^4$ filters for amplitude and timing pickoff, respectively, of the collecting anode and cathode waveforms. One filter is used to pick off the signal amplitude of neighbor waveforms, which is bipolar rather than unipolar. A final filter implements a simple subtraction algorithm that picks off the amplitude by subtraction of the tail from the baseline region. This is used in the

following stage for collecting/neighbor signal discrimination.




Figure 2.12: Implementation in hardware of the FIR filters. Each filter is 256-tap. A single operation involves 256 multiplies and 255 additions. The blue blocks represent individual 13-bit registers that connect up into a single large shift register bank. Individual waveform samples x_i are clocked in to the shift registers and multiplied by the corresponding filter taps.

The amplitude pickoff is folded in with the filter operation by a peak-hold unit included with the filter. It continually monitors the filter output and updates a register if the current value exceeds the running maximum. When the filter operation is fully complete, the stored value is the amplitude of the waveform. Figure 2.13 shows an experimentally obtained collecting pixel waveform from the H3DD-UM front end ASIC as it passes through various parts of the DSP pipeline. The blue and orange waveforms are from stage 1, and the yellow and purple are the outputs of the stage 2 slow and fast filters, respectively.

2.3.4 DSP Stage 3

The amplitude and samples of the fast filtered waveforms are output to stage 3, as well as neighbor signal and simple subtraction amplitude information. The simple subtraction amplitude is compared against a programmable threshold to make the distinction




Figure 2.13: An experimental waveform from the H3DD-UM ASIC after processing in stage 1 and stage 2 of the DSP.

between collecting/neighbor pixel waveforms. The fast filtered waveform samples are sent to three sets of coarse/fine timing pickoff modules. Figure 2.14 gives an overview of stage 3 timing pickoff modules.

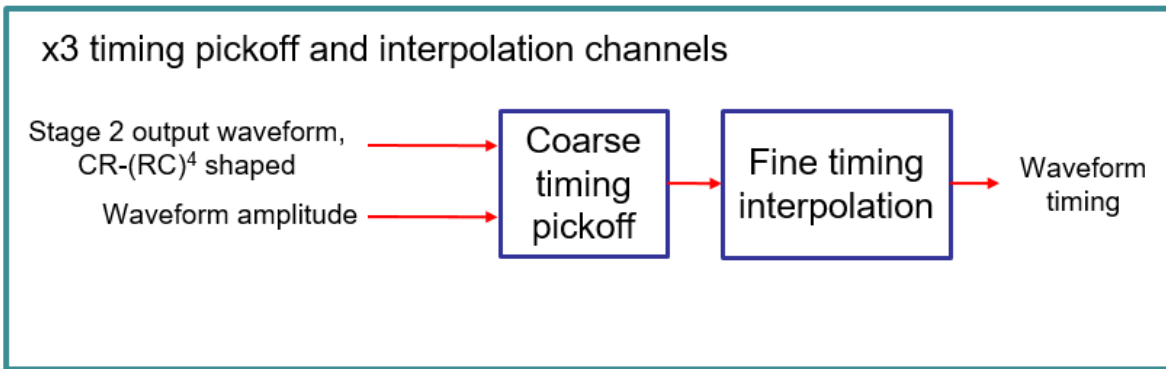


Figure 2.14: Detailed block diagram of DSP stage 3.

The task of timing pickoff is to identify the precise time when the filtered waveform

crosses an amplitude threshold. To avoid problems of time-amplitude walk, a constant fraction scheme is used where the amplitude threshold is defined as 50% of the maximum waveform amplitude provided by stage 2. The threshold is determined by simple bit shifting. The coarse pickoff module starts from the beginning of the fast filtered samples and searches forward for the first sample that rises above this threshold. To avoid potential noise or crosstalk pickoffs for low amplitude events, a programmable counter start value is provided that ignores all samples before a given point.

The ultimate purpose of determining the waveform timing is to find the event depth coordinate. The maximum depth of the waveform is 256 samples, but with typical sampling rate values, rising edges tend to be only up to around 40 samples long. The quantization error associated with this limited bit width can significantly degrade the vertical position resolution. To avoid this, a linear interpolation must be done to achieve a finer timing resolution with 16 bits of precision.


Once the coarse timing pickoff is found, the amplitudes of the filtered waveform before and at the pickoff point are multiplied by 8 through bit shifting, and saved. The threshold is also multiplied by 8. A fine counter starts at 1, and with each update increments the pre-pickoff amplitude by the slope of the waveform:

$$y_i = y_{i-1} + (y_1 - y_0) \tag{2.7}$$

where y_1 and y_0 are the filtered amplitudes at and before the coarse pickoff, prior to multiplication by 8. The starting value $y_{i=0} = 8 \cdot y_0$. When y_i exceeds the modified threshold, it is added to 8 x the coarse pickoff. The result is an 8x linear interpolation of the timing pickoff without complicated division or multiplication, and 16 bit precision. This interpolation scheme is verified with experimental waveforms in Figures 2.15a and 2.15b.

2.3.5 Control and Pipelining

Several control signals are required by the DSP. The H3DD_Core unit provides a data valid signal that indicates the start and stop of waveform readout. The first transition begins the operation of stage 1, as well as starts a counter that tracks the number of valid waveform samples. The second transition ends the counter, stores the results, and transfers them to the following stages. Banks of shift registers after the first and



(a) Timing pickoff interpolation with real experimental waveform from the VAD_UMv2.2 ASIC. Operation with deconvolved, filtered H3DD-UM waveforms is similar. (b) Detail of the timing pickoff point, showing the fine sub-sample pickoff from linear interpolation.

Figure 2.15

second pipeline stages catch the filtered waveform samples into a staging area as they are output. Once the data valid transitions, indicating a new waveform is being read out, the samples are transferred in parallel to another two banks which function as shift register inputs for stages 2 and 3.

After the end of stage 3 operation, a final set of N clock cycles is taken to collate the data and make a decision between collecting/neighbor/cathode processing outputs. The anode/cathode check is made by comparing the channel number provided from H3DD_Core. Cathodes have channel numbers of 0 and 64; all others are anodes. To distinguish between collecting/neighbor anodes, the simple subtraction check from stage 3 is used. If the data is either cathode or collecting anode information, the appropriate amplitude and timing values as well as channel number are passed on to the main controller. If the data fails the simple subtraction threshold test, the channel number and neighbor amplitude is output along with a maximum value for timing, used as a marker indicating the channel's status as a noncollecting neighbor. Finally, a data output strobe is sent to the main controller indicating valid data present on the DSP lines should be packaged by the event parser.

2.4 Synthesis, Layout, and Verification

2.4.1 Synthesis and Layout

The synthesis and layout toolchain is described in the Ph.D. thesis of Z. Chen [29]. In short, Synopsis Design Compiler is used to translate VHDL and SystemVerilog descriptions to netlists. Innovus is used to translate these netlists to a physical layout, making use of the ARM standard cells designed for the TSMC 65nm GP mixed signal process. Clock tree and reset net insertion is handled by the tools to meet timing and floorplan constraints. Once the digital core is assembled, it is mated with an analog section including the ADC and decoupling capacitors. The IO ring is joined at the same junction, with shared ground between analog and digital regions but separate power domains. The final layout of the chip is shown in Figure 2.16. The chip is surrounded by a seal ring structure. This is a purely mechanical component with no electrical function; it protects the integrity of the inner chip area as it is cut from a wafer using a saw.

The DSP, in particular the filters, represent a high density of interconnects and many high-fanout nets for reset lines and clock tree distribution. During APR, it was difficult to achieve DRC violation-free results. In the end several strategies were used. Hierarchical place and route (PnR) during the APR phase allowed for easier timing solving within FIR modules first, given timing constraints on filter inputs and outputs, which could then be used as-is in wider hierarchical synthesis and APR. Additionally, floorplanning tools could be used to help guide the initial solutions of the PnR algorithms to support lower densities. Figure 2.17 shows the impact of floorplanning, which together with hierarchical synthesis and APR is able to minimize the number of final DRC violations. Some few (<10) violations existed in the final design relating to via and metal spacing; these were fixed by hand in Cadence Virtuoso. Figure 2.18 shows the final DSP layout after APR and DRC fixing. The final size of the DSP core is 2.9 x 2.9 mm. Tool estimates for power consumption of the DSP core are on the order of 200 mW at 25 MHz processing clock rate. The DSP is expected to be the dominant source of power draw in the chip, which has a target of less than 1 watt overall. A 200 mW estimate is well within the allotted power budget.




Figure 2.16: Physical layout of the assembled v1.0 chip. The chip area is 4.2 x 4.2 mm, including the seal ring.

2.4.2 Verification Simulations

It is important to run simulations of the chip operation to catch errors prior to tapeout, since it is difficult to probe internal signals in silicon. To this end, a full testbench was built in SystemVerilog, including the digital core, a digital model of the ADC, a mockup of the H3DD-UM front end ASIC including all its sampling and readout logic, units which mimicked CPU function by loading commands from file and sending them over SPI interfaces to the DSP chip, and units which loaded experimentally measured waveforms from the H3DD-UM ASIC for digitization by the ADC. Figure 2.19 shows an example output from a post-APR run of the integrated testbench with back-annotated delays.

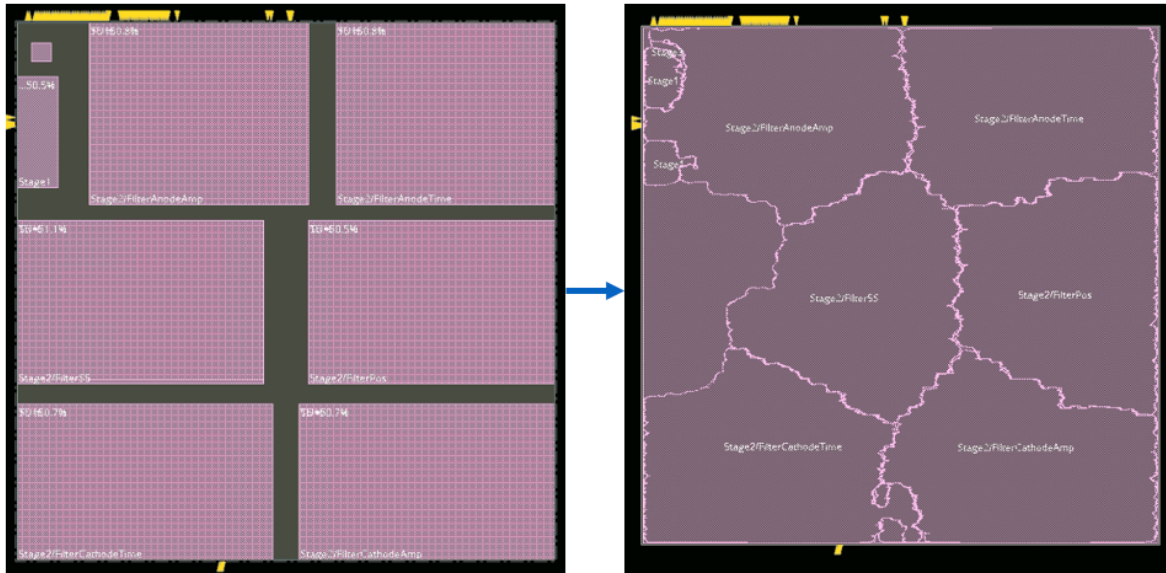


Figure 2.17: Floorplanning of the DSP core to help guide amebaplace and limit DRC violations from via/metal spacings.

There are three levels of simulation: behavioral, post-synthesis and post-APR. Behavioral simulations assume infinite signal speeds, and are purely tests of the HDL logic as written. These catch design flaws related to the logic structure. Synthesis tools like Design Compiler have some information about IO-level delays on standard cells that will be used, and can estimate some signal propagation delays. Post-synthesis simulations include these back-annotated delays and can catch some logic errors relating to physical behavior of the standard cells, such as setup and hold violations. Post-APR simulations are the highest fidelity, modeling the extracted timing delays from parasitic resistances and capacitances including cross-coupling across wire traces. They fold in models of transistor performance across several temperatures and process corners, allowing for best- and worst-case simulations. Different timing models can be used for different rule checks, enabling the testing of worst-case scenarios for different logic rules (such as setup/hold violations).

Several common groups of violations arose in post-synthesis and post-APR simulations. Many were related to setup/hold violations (or reset recovery/removal violations). These arise when dealing with sequential elements like registers. The data input needs to be held constant for some amount of time both before (setup time) and after (hold time)

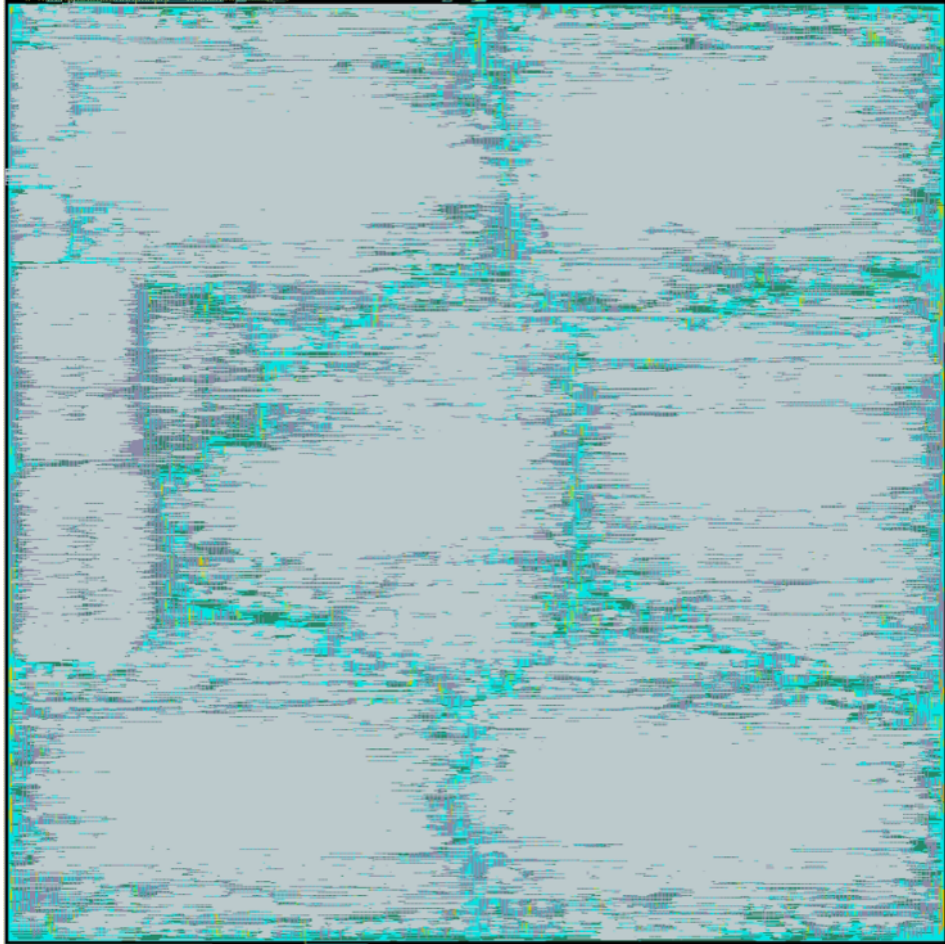


Figure 2.18: Layout of the DSP after APR. Six large FIR units are clearly visible. The next largest block of area is taken up by pipeline registers, followed by stage 1 and 3 logic.

the active clock edge. Violating these constraints leads to unpredictable behavior and possibly metastable states. Most of these were fixed by staggering signal assignments across opposite clock edges, giving the maximum possible combination of setup/hold slack.

A similar requirement exists for resets. If a reset is asserted too close to an active clock edge, it may or may not function. This typically is avoided by the use of asynchronous resets that take priority over any clocking path. Asynchronous resets are used everywhere in the DSP chip. A problem comes when the reset is de-asserted too close to an active clock edge. This gave rise to many recovery time violations. These were resolved

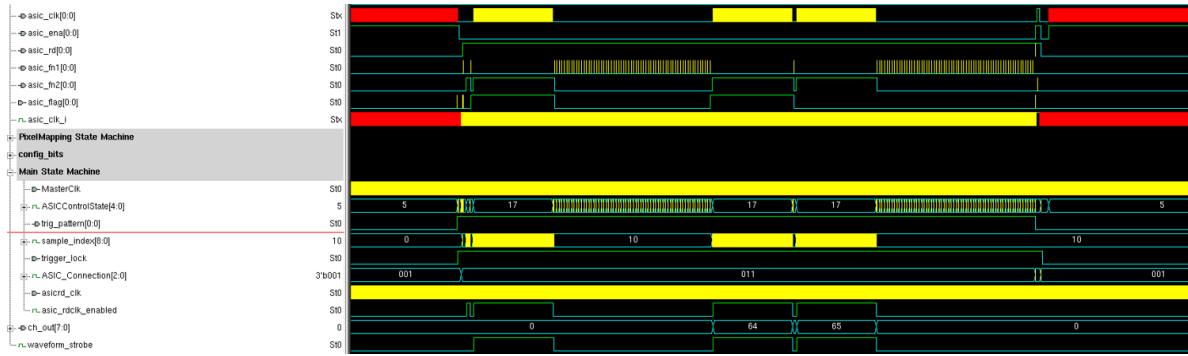


Figure 2.19: Example partial output of the integrated testbench of the DSP ASIC system processing data during readout of the H3DD-UM front end. Post-APR delays are included. In this simulation, a single pixel event in channel 65 of the front end is simulated. The very bottom line is the data valid signal to the DSP from the H3DD_Core; here, it is high three times, twice for cathode signals in channels 0 and 64 and once for the triggered pixel in channel 65 immediately after.

by implementing reset bridge structures (Figure 2.20), creating a hybrid asynchronous-assert/synchronous-release reset structure [30]. By synchronizing the reset release to the inactive clock edge, the maximum slack is given and errors are avoided.




Figure 2.20: Reset bridge structure used to create an asynchronous-assert/synchronous-release reset.

2.5 Test Systems

2.5.1 Simplified Test System

The DSP ASIC requires several input signals, and there are many variables to tune. Independent access to each of the three power rails (analog, digital core, and digital

IO pad) is important to verify the power consumption of different parts of the chip, as well as gauge the level of activity in the digital circuits. Debug of the ADC requires the input of well defined signals such as ramps and sine waves that are not generated by the front-end H3DD-UM ASIC. Interfaces between monitor lines and a digital signal analyzer are needed to properly debug the chip, and the SPI lines need to be monitored for any bugs that arise during communication.

To address these needs, a simplified test system was designed, shown in Figure 2.21. This contains a socket to mount the DSP ASIC, termination networks for better signal quality, and various headers and sockets to attach test equipment.




Figure 2.21: Simplified DSP ASIC test system.

2.5.2 Integrated Test System

The simplified test system enabled fine monitoring of the signals entering and exiting the ASIC, as well as precise monitoring of the power consumption. However, the long traces and aerial leads to different components degraded the signal integrity, limiting the maximum clock rate that could be reliably achieved to about 6.25 MHz. Also, it would not be possible to test with an actual H3DD-UM front end ASIC.

An integrated test system was designed based on a prior H3DD-UM readout board. Significant modifications were made. A diagram of the integrated test system is shown in Figure 2.22. It is split into two boards: the motherboard and the front end board. The motherboard hosts a Trenez 0720 FPGA with a Zynq-7 series processor, several power regulators, and connectors for an external CPU. Only a single 12V power input is needed. Ribbon cables carry critical signals over to the front end board, which can be isolated from the noisier motherboard. On the front end board, level translators shift logic levels to ranges suitable for the DSP ASIC, the Trenez FPGA, and the H3DD-UM front end ASIC. Filter networks and low noise regulators further clean the power rails. A high-speed 14 bit DAC is included, which can be programmed via a module developed and implemented in the FPGA. It also includes the single-to-differential converter circuitry required to interface between the single-ended H3DD-UM output and the differential ADC input on the DSP chip. An image of the integrated test system is shown in Figure 2.23. Thanks to the shorter traces and low noise regulators, clock rates of up to 100 MHz can be achieved.

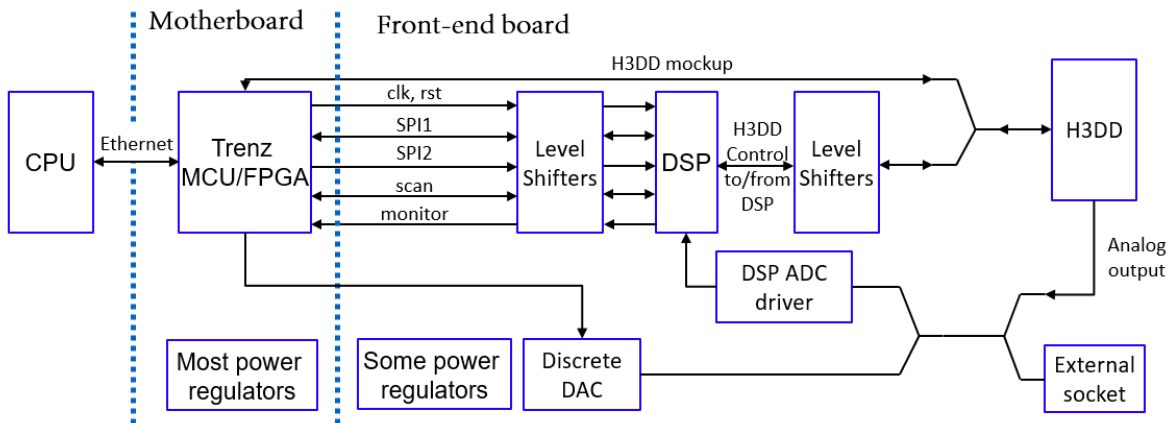


Figure 2.22: System diagram of the integrated test system and peripherals.




Figure 2.23: Integrated test system hardware.

2.6 Experimental Evaluation of v1.0 and v1.1

2.6.1 v1.0 Debugging

The complete chip designs were submitted to TSMC and fabricated in the 65nm GP process. 100 chips were packaged and returned. A fabricated DSP ASIC is shown in Figure 2.24. Taped-lid QFN packages were chosen to avoid interfering with the analog input wirebonds to the ADC. However, they are weaker than molded encapsulations, which resulted in many chip lids cracking when inserted in the test sockets. A specialized pusher plate was designed that transmitted force along the package boundary, minimizing the number of cracked ASICs.

However, even with sufficient pressure, the ASIC did not respond to commands and drew practically no power. Measurements of the gate voltages on the analog input pins showed the process fabrication was executed correctly. The error was traced to a mistake in the power distribution of the I/O ring on the digital side. A revision was made to mask layers M8-AP, and a resubmission was made.


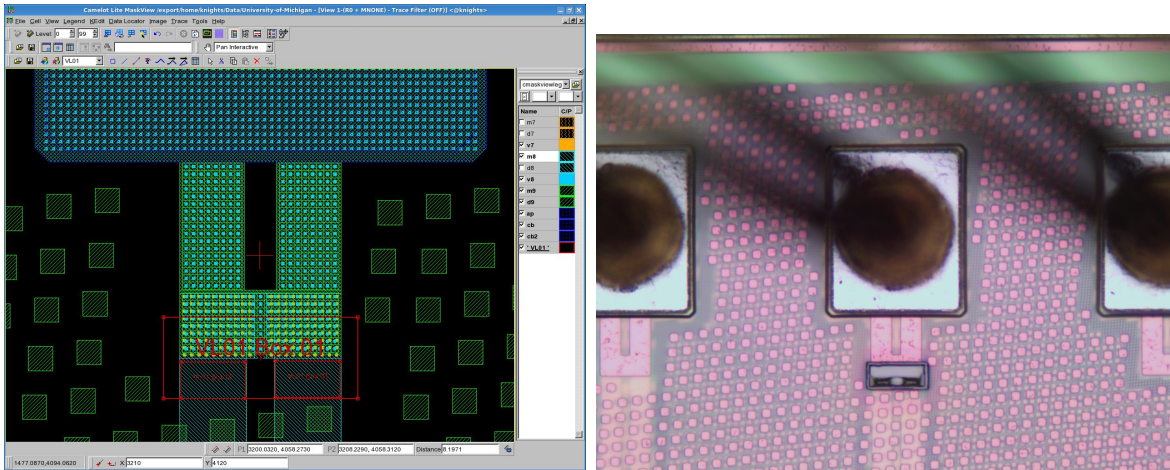


Figure 2.24: Fabricated ASIC incorporated in a 7x7 mm QFN package.

2.6.2 v1.1 Debugging

After the power rail fix, the modified v1.1 ASICs were fabricated and returned. Power consumption was still negligible and the ASIC did not respond. The cause was traced to a mistake on the revised submission tapeout sheet; the revision was implemented for M8 and up but skipped over via 8, with the result that no power was still being delivered to the digital core. Rather than resubmit for another fabrication run, focused ion beam (FIB) technology was used to add in the missing vias to several V_{DD} pads around the chip. Figure 2.25a shows the plan submitted by the FIB company to achieve the work, and Figure 2.25b shows a micrograph of the region after the FIB fix is implemented.

With the FIB fix in place, the chip finally responded. Details of the testing of the top level supporting circuitry and the ADC is in the thesis of Z. Chen. Debugging of the DSP could not be done without a functioning H3DD_Core module, since the DSP relies



(a) Plan for the focused ion beam fix for the missing power via under a V_{DD} pad submitted by the FIB company ABER. (b) Micrograph of the region with the FIB fix implemented.

Figure 2.25

on control signals from the H3DD_Core to control the pipeline progression.

Unfortunately, testing revealed many problems in the H3DD_Core logic that were not present in post-APR simulations. These stemmed from marginal timing, which was sufficient in simulation but slightly worse when implemented in silicon. Though the H3DD_Core could respond as expected to SPI commands initially, after several minutes of operation its responses became erratic and eventually ceased altogether. It was found that cooling the system could stabilize the response and allow for extended testing. The integrated test system was placed in a thermal chamber, and several test sequences were run. The most reliable operation was achieved at temperatures of -20° to -30° C. Figure 2.26 shows the results of one such test. The output signals from the H3DD_Core were monitored, and a command to execute the ASIC_gReset sequence was sent over SPI. The H3DD_Core responded immediately, and the signals matched in broad strokes those expected from simulations. However, some timings were inconsistent with the simulation and there were some replicable glitches on a few of the signals. The conclusion is that narrow timing margins cause counters that control the H3DD_Core state machines to miss certain bits, temporarily resetting them. At higher temperatures, the counters become more unreliable, and remain stuck in endless loops once a command is received.



Figure 2.26: Output signals from the on-chip H3DD_Core module at -30°C in response to an ASIC_gReset command sent over SPI. On the bottom is the expected output from behavioral simulation. The measured matches the simulated in broad morphology, but several timings are modified and glitches are present in some of the signals.

2.7 Modifications for Future Version

The timing margin of the H3DD_Core could be improved with hierarchical synthesis, which should be enough to stabilize the counters. Many design-for-test changes were made to the DSP and H3DD_Core, informed by the debugging experience. Critical internal signals are routed out to the debug monitor, which has been extended from 3 bit select/4 bit output to 4 bit select/16 bit output. The DSP architecture has also

been modified to catch the control inputs from H3DD_Core and optionally bypass them with external inputs. This will allow for the debugging of the DSP core even in the case when the H3DD_Core does not function.

Though the current chip design could be used for 4 H3DD-UM front end ASICs by multiplexing their communications, it is an awkward solution. For future iterations, the DSP core will be scaled to process the input from 4 ADCs simultaneously. This would improve the readout of a 4-module array of H3DD-UM front end ASICs, but will require either significant changes to the pipeline architecture or a great reduction in size. The reduction of size can be accomplished by switching to IIR instead of FIR filters. This fixes the choice of filter to trapezoidal and CR-(RC)⁴, but in practice this is the most widely used configuration, and some programmability will still be possible. A test synthesis of an IIR trapezoidal filter reduced the footprint by a factor of ten compared to the FIR variant.

Chapter 3

Advanced Subpixel Position Calibrations

One accurate measurement is worth a thousand expert opinions.

- Grace Hopper

3.1 Charge Transport and Position Reconstruction in CZT

3.1.1 Subpixel Position - Accuracy and Precision

A particular strength of 3D CZT technology is the capability to reconstruct the position of gamma-ray interactions. This enables many imaging algorithms [31][32] as well as high fidelity energy calibrations [33]. Coarse position resolution in XY can be achieved by observing the pattern of triggered pixels. Smaller pixel pitches can be used for better position resolution. However, the pixel pitch may be optimized for electron cloud collection at a particular photon energy rather than for position reconstruction. This optimization can conflict with that for fine pixel-level position resolution. 50 μm -pitch pixels may achieve very good position resolution, but will suffer from extensive crosstalk and charge sharing that would degrade the energy resolution at 662 keV.

Sub-pixel position reconstruction allows for the the same level of position resolution using larger pixels that support better energy reconstructions. The basic techniques described in section 1.4.2 have been implemented with both analog (BNL-H3D, VAS-TAT)

and digital ASICs (VAD-UM, H3DD-UM). However, digital ASICs have considerable advantages. Access to the full waveform allows for lower trigger thresholds and better charge sharing corrections, both of which benefit subpixel reconstruction which relies on low amplitude transients.

Degradation mechanisms can affect two different aspects of the reconstruction: accuracy and precision. Accuracy refers to the error in the mean of a group of measurements of the same quantity from the true value. It is impacted by systematic errors. Precision, meanwhile, is characterized by the standard deviation of the group of measurements; it is affected by stochastic errors.



Figure 3.1: Difference between accuracy and precision.

3.1.2 Factors Affecting Subpixel Resolution

Subpixel position resolution is affected by both accuracy and precision errors. Precision problems stem largely from electronic noise. The transient signals induced by a drifting charge are much smaller in amplitude than the collecting waveform itself. The electronic noise is constant, leading to a worse signal-to-noise ratio. This imprints itself as a source of stochastic error in the reconstructed positions. These problems are worse at lower energies, with the result that the subpixel position resolution is worse at 32 keV than it is at 662 keV.

Accuracy issues arise from multiple causes, which cause distortions in the measured positions on large (multi-pixel) and small (sub-pixel) scales. On the small scale, charge sharing may pollute the neighbor signals and skew the reconstructed position. Charge sharing is more likely to occur for events near the pixel edge, causing systematic distortions on a subpixel level. Charge sharing is also more likely at higher energies. This

causes the subpixel resolution to be worse at several MeV than it is at 662 keV. Even in the absence of charge sharing, however, the relationship between transient signal amplitude and subpixel position is nonlinear. The classic center of gravity algorithm, equation 1.5, fails to account for this. Distortions will arise even at 662 keV from this nonlinearity, resulting in an “X” pattern of subpixel positions as shown in Figure 3.2. Zhu developed a more accurate method that tackles the nonlinearity issue by comparing measured signal ratios against expectations from weighting potential simulations [12]. However, this technique may still suffer from model mismatch.




Figure 3.2: Distribution of opposing neighbor ratios in X and Y, showing the “X” distribution which is due to the nonlinearity of equation 1.5.

On the large scale, nonuniform electric fields may cause systematic distortions in reconstructed positions. Owing to the same weighting potential shielding effects that cause collecting anode signals to be mostly independent of depth, the neighbor signals are also mostly insensitive to the electron drift path until the charge cloud nears the pixel.

The reconstructed positions are therefore not the XY coordinates of the interaction location, but the coordinates at which the charge cloud entered the near-pixel region. These must be backprojected along the charge drift path to the right interaction depth, which requires knowledge of the electric field distribution inside the crystal. Typically, it is assumed that the electric field is uniform, with field lines running vertically above pixels straight through to the cathode. In this case, the coordinates reconstructed from neighbor signals are the same as the coordinates of the interaction. In practice, this is rarely the case. Nonuniform fields in X and Y can push the charge around as it drifts, and in extreme cases can cause it to collect on a different pixel than the one it originated over. In this case, it is called pixel jumping [34]. If the charge is steered to a dead region, such as the crystal edges, counts may be lost entirely. Even when the effect is not severe enough to push charges to a different pixel or the crystal edge, it will still degrade algorithms that rely on accurate subpixel positions. Both pixel jumping and sub-pixel distortions caused by nonuniform fields are referred to as charge steering in this thesis.

Figure 3.3 demonstrates the effects of charge steering on a subpixel level. A 100 μm -wide slit opening in a block of tungsten was used to collimate 662 keV gamma rays from a Cs source into a fan beam. This was shone down onto the cathode surfaces of an array of CZT detectors, and the subpixel positions were reconstructed according to the simulation-comparison technique developed by Zhu. A straight line should be observed, with potentially some slanting or blurring from a misalignment between detectors and collimator. However, the observed distribution deviates nonmonotonically from a straight line, bending back and forth from left to right. This would lead to distortions in the recorded positions that would imprint as artifacts on Compton or coded aperture images.

3.1.3 Mechanisms of Charge Loss and Charge Steering

When events are recorded on the pixels of CZT detectors, it is assumed they followed a linear trajectory straight down from the initial radiation interaction position. The depth can be calculated either by cathode/anode ratio or by drift time. The 3D location is identified by projecting the calculated 2D location from the pixel plane back up by the measured depth value into the bulk. In practice, there are numerous scenarios where




Figure 3.3: Subpixel distribution of photopeak events in CZT following irradiation through a $100\mu\text{m}$ wide slit collimator. Ideally, a straight line should be observed. In practice, nonuniform electric fields cause distortions in the reconstruction.

charge may be diverted or lost altogether before it reaches the anodes.

To prevent interference among electronic channels connected to the pixels, and for ease of manufacturing, there is always a gap between pixels. In the CZT detectors used in this thesis, that gap is $60\mu\text{m}$ (Figure 3.4) [35]. The electric potential above this gap is lower than that of the pixel surface. If any part of the electron cloud lies between

pixels, the low potential in the interpixel gap will cause it to be attracted to one or the other pixel. In between the pixels the electric field is also very low, so any charge which falls close to the middle of the gap may pile up there and be very slow to collect. If the collection time becomes much slower than the shaping time of the filters used, it can be lost altogether. The presence of trapped or slow charge in the interpixel gap will also modify the electric field there, strengthening the steering effect unpredictably. The magnitude of this effect can vary with temperature, flux, and other factors, making calibration fairly challenging.




Figure 3.4: Fine detail of the pixel pattern used with the CZT detectors in this work. The pitch between pixels is 1.72 mm, while the pixel size is 1.66 x 1.66 mm². The resulting gap is 60 μm wide.

Earlier CZT detectors attempted to explicitly control the degree of steering between pixels by including an inter-pixel grid, which was biased in the range of -20 to -80 V [36]. This arrangement is shown in Figure 3.5. However, challenges in consistent manufacturing of the grid and extra noise from grid-pixel leakage current resulted in its abandonment in favor of simple pixel arrays.

Similar to the inter-pixel grid, modern CZT detectors surround the entire array of pixels with a guard ring (also shown in Figure 3.4). The guard ring makes the weighting potential of the peripheral pixels more uniform, as well as absorbing any surface leakage

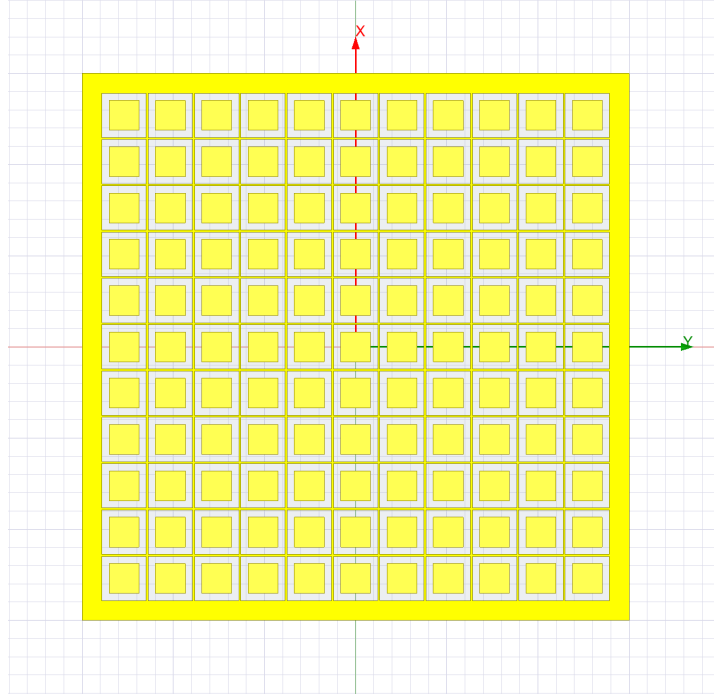


Figure 3.5: Pixel pattern with inter-pixel grid, used by older generation CZT detectors.

current that would otherwise degrade the energy resolution. The guard ring is AC-coupled to its readout channels, so any leakage current absorbed by it is directed through a grounding resistor. This causes a bias to be present on the guard ring which can steer drifting charges near the edges of the crystal to the peripheral pixels. The strength of the steering effect depends on the amount of surface leakage, which varies with temperature. If the guard ring is directly grounded, it may even steer charge away from the pixels.

Throughout the crystal there can be defects which trap electrons or holes for some amount of time. If the detrapping time is long compared to the collection time, the charge will be lost entirely. Along the crystal surfaces, for example, there may be regions of high trapping which result in a dead layer several dozen microns thick. Under high radiation flux there can be a buildup of trapped charge in the crystal which significantly alters the electric field. Even at moderate fluxes, there can still be fixed space charge inside the crystal that alters electron drift trajectories. These can severely distort the reconstructed positions of interactions [37].

3.2 Subpixel Precision Improvements with Machine Learning

3.2.1 Machine Learning Concepts

Machine learning refers to the capability of some computer algorithms to incorporate experience from a set of training data to improve their performance at a particular task [38]. Machine learning algorithms broadly fall into two categories: supervised and unsupervised. In supervised learning, the training set includes labeled examples, and the model is optimized to predict the labels from the example inputs. In unsupervised learning, the algorithm attempts to learn patterns intrinsic to the training data without the guidance of explicit labels.

There exist large families or classes of machine learning algorithms, which operate on similar principles but differ in the details of their implementations. For example, a family of neural network algorithms may have different numbers of hidden layers. These large-scale parameters define the structure and operation of the machine learning model, and are independent of the training set used. They are called hyperparameters. Often, the hyperparameters of a model may be optimized to achieve the best performance for a given task. There are many families of supervised and unsupervised machine learning algorithms, such as support vector machines (SVMs) and Gaussian process regressions (GPRs).

One of the most widely used algorithms is the neural network, such as the multi-layer perceptron [39]. A diagram of a neural network architecture is shown in Figure 3.6. The network consists of an input layer, several hidden layers with varying connectivity, and an output layer. Each layer represents a function or group of functions. Borrowing terminology from neuroscience, each circle is called a “neuron”. Each neuron represents a linear combination of its inputs, with an optional bias term. The linear combination is followed by an activation function which introduces some nonlinearity. The activation function is a hyperparameter, and can vary from model to model. The output of each layer is passed to another layer further down the line. By compositing many layers using simple activation functions, the network can learn to represent and calculate complex relationships among the input data.

The principal component analysis (PCA) algorithm of B. Williams [10] is a form of




Figure 3.6: Diagram of a neural network architecture with three hidden layers.

supervised learning. Pixel signals from photoabsorptions of a known gamma-ray energy, the 662 keV line from ^{137}Cs , are grouped into a data matrix and eigenvectors are extracted using singular value decomposition. The waveforms are then projected along these eigenvectors to yield principal component values, that are regressed against the amplitude reconstructed by traditional means. Since the ground truth energy is known, the regressions can be made to match up with the calibration value. A normalization and time alignment preprocessing step allows signals from other energies to be projected against the same eigenvectors and the same calibration parameters applied. Williams demonstrated a 10% improvement in the energy resolution for CZT detectors, and similar or greater gains have been demonstrated with TlBr and perovskite detectors.

Subpixel position reconstruction presents a different sort of challenge. Unless a highly collimated beam is used, the ground truth of interaction locations is often not known. Even with a collimator, the travel length of the photoelectron represents some extra position blurring. It is not clear if a supervised approach may yield good results. On the other hand, it is difficult to guarantee that the variations learned by unsupervised methods will correlate clearly with the subpixel position.

3.2.2 Radial Position Reconstruction

Current subpixel position reconstruction algorithms use the neighbor pixel signals to derive the subpixel position. Any information embedded in the collecting waveform is discarded. In an attempt to improve the resolution by extracting more information, an algorithm was developed to find the subpixel position from the collecting waveform. Owing to the 90° rotational symmetry of the pixel layout, the subpixel position may not be uniquely determined, but the radial coordinate can be obtained.

A procedure for calibrating a supervised algorithm to predict the radial position from the collecting waveform was developed, diagrammed in Figure 3.7. A calibration measurement is done which flood irradiates a detector with a standard source, such as ^{137}Cs . Classic techniques described in Section 1.4.2 are used to calculate the XY subpixel position. These coordinates are used to find an estimate of the radial position, which serves as a training label for the set of associated collecting waveforms. This training set of collecting waveforms and approximate radial labels is used to train a regression model. The model learns to predict the radial position from the collecting waveform. Because the labels are only estimates rather than the ground truth, any systematic errors in the classic reconstruction will embed itself in the model as well. However, the model may learn a more precise representation of the subpixel position, and improve on the stochastic error in the training set.

The procedure was applied to a calibration set of ^{137}Cs data. To achieve the necessary statistics, data from all 9x9 inner pixels of a single detector were combined. However, since the collecting waveform can also vary significantly with depth, individual models were trained on data from different coarse depth slices throughout the crystal. Matlab's Regression Learner Toolbox was used to perform the training. Aside from neural networks, several different model families were explored, including SVMs, GPRs, and decision trees. Ultimately, GPRs provided the best performance for predicting radial position.

The resulting models were applied to a separate validation dataset. Opposing neighbor ratios were used to estimate the subpixel position and compared against the predicted radial values. Figure 3.8 shows the radial positions predicted by the trained GPR models versus those calculated by opposing neighbor ratios. The red trace shows the line of 1:1 correspondence. As a further demonstration, events with different ranges of predicted

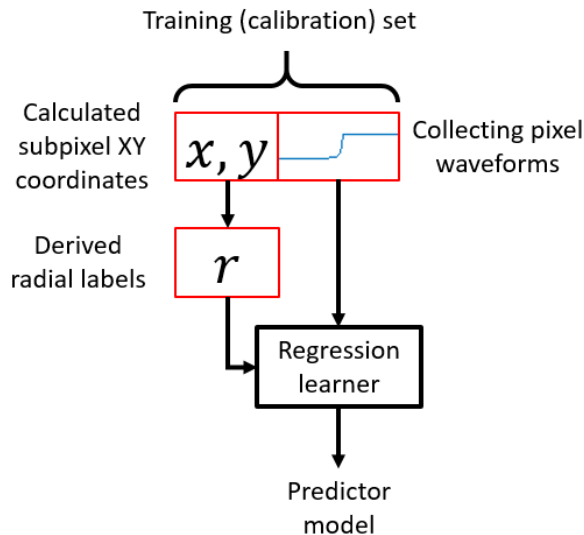


Figure 3.7: Algorithm for training a machine learning model to predict the subpixel radial position of interaction purely from the collecting waveform.

R values were selected from the bulk dataset, and their XY distribution calculated from neighbor ratios is plotted in Figure 3.9. As the selected value of R increases, the distributions form rings of increasing radius.

The radial position resolution can be estimated by choosing events with a particular range of R values predicted by the GPR model, calculating the distribution of corresponding R values from opposing neighbor ratios, and fitting it to a Gaussian. This analysis is shown in Figure 3.10. Converting to mm, the FWHM of the distribution is about $200 \mu\text{m}$. Deconvolving the $100 \mu\text{m}$ selection window, the radial position resolution is found to be about $170 \mu\text{m}$. This estimate is not the true radial resolution, since it is not made using a collimator. Rather, it is the resolution for the radial position of the charge cloud centroid.

The technique is demonstrated for single interaction events. Calibration of multi-interaction events is more difficult, since there is cross-contamination of information between different interactions, especially if events are side-neighboring. About half of events are not side neighboring at 662 keV. To these, the models trained on single-interaction events can be applied directly. Figures 3.11a and 3.11b show similar plots for non-neighboring 2-pixel events, where both interactions are cathode side. The model trained on single pixel interactions was used. Since there was minimal cross-




Figure 3.8: Radial positions predicted from only collecting waveforms vs. those calculated from opposing neighbor ratios.

contamination, the predictions still were relatively accurate.

3.2.3 XY Position Reconstruction

The success of the algorithm in predicting radial position led to considerations whether it could provide useful predictions of X and Y positions as well. To be of benefit, it would need to predict subpixel position to better accuracy than what is represented in the training set. This may be possible, since the regression learning process trains the model to predict the underlying function mapping waveforms to position.

The algorithm for predicting XY position is similar to that for radial position. Instead of only the collecting waveform, the four nearest neighbors are included in the data matrix as well. PCA is used to preprocess the waveforms, along with normalization and time alignment steps to help ensure the trained model can generalize to other energies as well as possible.

Rather than train on opposing neighbor ratios, Zhu’s subpixel algorithm that compares against simulated weighting potentials was used to generate the training labels. This helps the process in two ways. First, it removes some systematic errors from the




Figure 3.9: XY distributions (calculated from neighbor ratios) of events with particular values of radial position (predicted by GPR model).

training labels that would otherwise be imprinted on the trained model. Second, the regressions will be against the true positions, rather than charge centroid estimates. The outputs can be used directly, and may be more precise if the model is able to learn any information about the electron cloud orientation.

A 6-hour calibration with a ^{137}Cs source in close proximity generated sufficient statistics to attempt training per-pixel models. All depths above 2mm were merged into one to improve the per-pixel statistics, since away from the anode there is less variation with depth of the neighbor signals. Per-depth calibrations were also studied, in which data


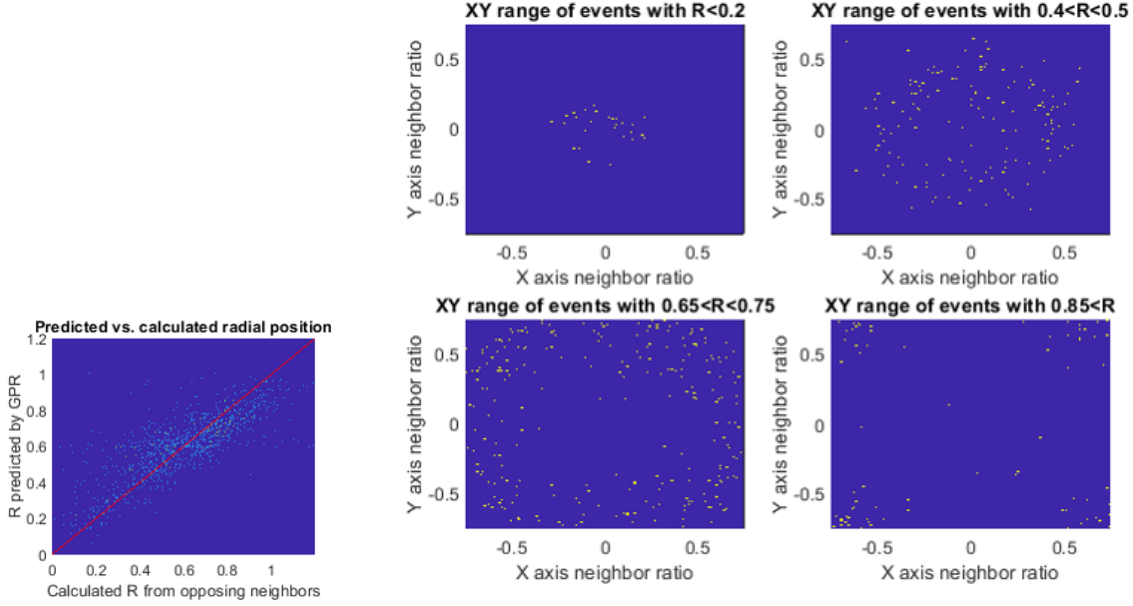


Figure 3.10: Radial distribution of events with predicted values of R between 0.4 and 0.5. Converting to mm, the FWHM of the distribution is about $200 \mu\text{m}$.

from all 9x9 interior pixels were merged into a single training set.

To validate the algorithm, a $100 \mu\text{m}$ wide tungsten slit collimator was used to irradiate a slice across the detectors. Figure 3.12a shows the distribution of events from this collimator, while Figure 3.12b shows the X cross section for the second pixel from the bottom. This pixel is in a region with less steering than the rest of the crystal, making it more suitable for evaluating the limiting position resolution. The response of the collimator is simulated at different depths using raytracing, and deconvolved from the measured distribution. After this deconvolution, the position resolution found using classical techniques is $445 \mu\text{m}$. This is significantly worse than the best subpixel resolutions previously reported at this energy for these types of detectors [9]. The difference is likely from crystal variations - even though the region of severe steering is further to the top of the crystal, there can still be significant diffusion of the electron cloud in this detector.

As with the radial prediction, several different models were tried. Traditionally, the training process aims to minimize some cost function with respect to the training labels, such as the mean squared difference or the cross-entropy. However, in this case the goal is not to predict the training set as accurately as possible but to predict an underlying




(a) Predicted vs. calculated radial positions for 2-pixel non-side-neighboring cathode side events. (b) Calculated XY distributions of 2-pixel events with particular values of radial position predicted by the single-pixel-trained GPR model.

Figure 3.11

distribution that more closely approximates the true distribution than do the labels themselves. The models with the greatest success in matching the training labels may only perform as well as, and not better than, the traditional methods used to generate the training labels in the first place.

In this case, it was found that a narrow neural network, with few neurons in the hidden layers, performed the best. Figure 3.13a shows the XY distribution of collimated events predicted by a narrow neural net trained on aggregate data from all inner 9x9 pixels. This provides plenty of training statistics, but degrades some of the quality of the training data by including data from poorly performing pixels with significant steering problems. Nevertheless, after deconvolving the collimator response function, the subpixel X resolution is found to be $422 \mu\text{m}$, an improvement of $\sim 25 \mu\text{m}$ over the traditional techniques (Figure 3.13b). When a network is trained on data purely from this pixel, the improvement is even greater, down to $395 \mu\text{m}$ or $50 \mu\text{m}$ better (Figure 3.13c). It is also able to do this with only four neighbor signals per event, while the



(a) Distribution of subpixel XY positions for irradiation with a $100\ \mu\text{m}$ slit collimator, calculated using classical techniques. (b) X cross-section of events in pixel (7,3) (second lit up pixel from the bottom in Figure 3.12a) and Gaussian fit. The measured position resolution in X is $445\ \mu\text{m}$ after deconvolving the response of the collimator.



Figure 3.12

classical techniques require eight to achieve the precision they do. Being able to use only four neighbors and achieve comparable or better position resolution can decrease the system dead time. This is an important consideration for applications such as medical imaging, which also stand to benefit from the increased position resolution.

3.3 Charge Steering Calibration Techniques


3.3.1 Collimator Scan Corrections

Problems of accuracy are dominated by space charge and nonlinear electric fields. These fields deflect the charge as it drifts down to the pixels, and cause distortions in the final reconstructed XY positions. The Z positions tend to be less affected. The cathode/anode ratio which is used to determine the depth of single interaction events depends primarily on the weighting potentials, which are calculated with all space charge removed. Drift times can be severely affected by nonuniform fields, but a calibration can correlate drift time with cathode/anode ratio for single interaction events. This correlation can then be used to correct the estimated depth of multi-pixel events. The depth accuracy



(a) Distribution of subpixel XY positions for irradiation with a $100\ \mu\text{m}$ slit collimator, predicted using a narrow neural net trained on aggregate data from all inner 9×9 pixels.

(b) X cross-section of events in pixel (7,3) predicted by a narrow neural net trained on data from all inner 9×9 pixels. The measured position resolution in X is $422\ \mu\text{m}$ after deconvolving the response of the collimator, a $\sim 25\ \mu\text{m}$ improvement over the traditional techniques.




(c) X cross-section of events in pixel (7,3) predicted by a narrow neural net trained on data from this particular pixel. The measured position resolution in X is $395\ \mu\text{m}$ after deconvolving the response of the collimator, a $\sim 50\ \mu\text{m}$ improvement over the traditional techniques.

Figure 3.13

is therefore as good as what can be achieved with the cathode/anode ratio, which is distorted only in the case of nonideal weighting potentials or significant material nonuniformities.

To investigate the uniformity of the cathode/anode ratio, a 100 μm slit collimator was scanned in 200 μm steps across the Z axis of the crystal. Figure 3.14a shows the calculated position from the cathode/anode ratio vs. the true collimator position. There is significant deviation from the true position, especially in the middle of the crystal. Most of the deviation is likely due to the nonideal weighting potential of the cathode, which is distorted by edge effects and the boundary of the crystal. These may be difficult to completely correct by simulation, since the weighting potential distortion depends on the details of the surface treatment and encapsulation, which can vary from crystal to crystal.

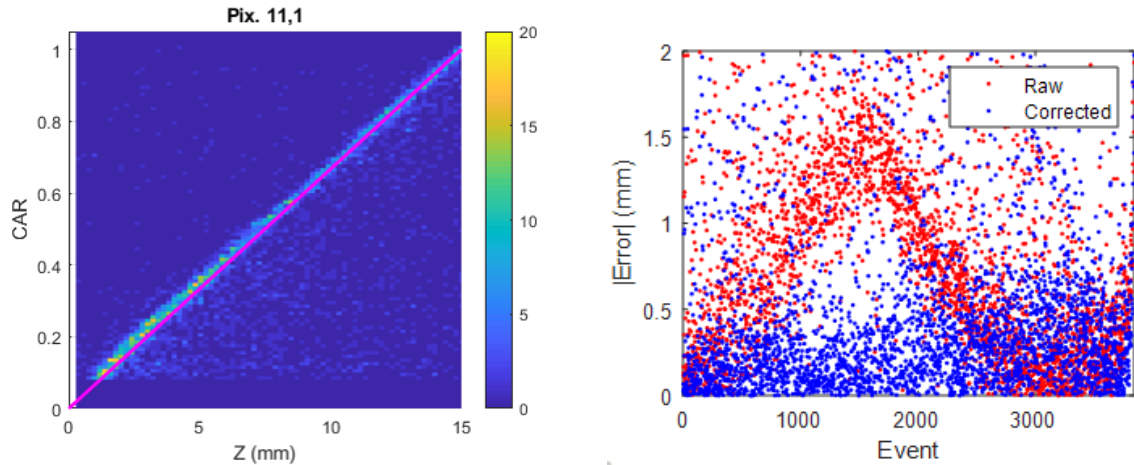


(a) Calculated position from cathode/anode ratio vs. true collimator position for a corner pixel, (11,1). There is significant deviation in the middle, mostly due to the nonideal weighting potential. (b) Absolute error between calculated and measured Z coordinates above the corner pixel (11,1).

Figure 3.14

With data from the collimator scan, a very accurate calibration can be done to correct the distortions. Figure 3.15a shows the calculated vs. measured positions after applying the collimator-based calibration, while Figure 3.15b shows the residual error. The large deviation in the middle is corrected.

However, collimator corrections for position errors are costly in terms of calibration



(a) Calculated position from cathode/anode ratio vs. true collimator position after correction with collimator information. The deviations in the middle are corrected. (b) Absolute error between calculated and measured Z coordinates after applying the collimator-based calibration.

Figure 3.15

time. Also, they are not easily performed in the field, since they require additional hardware. The position distortions may change with temperature, flux, relative humidity, and other parameters, so a simpler, collimator-less way to perform these calibrations would be valuable.

3.3.2 Count Density Mapping

One possibility stems from the observation that the underlying distribution of gamma-ray interactions inside the crystal is not affected by nonuniform internal fields or other defects, but instead follows simple laws of physics. The observed distribution of counts depends on both the physics of photon interactions as well as the nonuniformities in the reconstruction. By comparing the observed with the expected distributions, it should be possible to cancel out the common factor of photon interaction physics and gain insight about the internal nonuniformities.

Specifically, the count densities should follow exponential attenuation distributions. At the ^{137}Cs energy of 662 keV, attenuation should be relatively minor through the 15 mm thickness of the CZT detector. Dividing the observed count density in a given

region by the expected count density gives a value that should be unity if the positions are reconstructed accurately. The distribution of nonuniformities in these dimensionless quantities should correspond to the distribution of nonuniformities in the position reconstruction.

Given these nonuniformities, a calibration can be built up by iteratively perturbing the reconstructed positions with the aim of making the observed-to-expected density ratios uniform. Some inspiration is taken from the behavior of air pressure differentials in a sealed room. Areas of overpressure will expand outwards, and regions of underpressure will contract inwards, until a steady state is reached. In the same way, boundaries of detector voxels can be pushed by overdensities of counts on either side, while making sure to enforce boundary conditions on the edges of the crystal. This does not rely on collimator data, being accomplished with a flood irradiation with ^{137}Cs . Figure 3.16 shows this algorithm applied to calibration data above pixel (11,1) from Figure 3.14a.





Figure 3.16: Count density-based voxel reshaping algorithm applied to a set of calibration data. (Left) initial count density overpressures; (right) after voxel boundary shifting to accomplish a uniform distribution.

Figure 3.17a shows the calculated vs. true positions after applying the voxel reshaping algorithm in the Z dimension. Figure 3.17b shows the absolute error. The improvement is not as good as with the collimator-based calibration. However, there is still a $500\ \mu\text{m}$ improvement in the central, worst part of the pixel. The performance on the cathode side is about the same as with the collimator.

The edge pixels represent the worst case scenarios for the cathode/anode ratio reconstruction, since the weighting potentials are the least ideal. There may still be some deviations in the edges and the center of the crystal. The count density based algorithm



(a) Calculated position from cathode/anode ratio vs. true collimator position after correction with the count density-based voxel reshaping algorithm. (b) Absolute error between calculated and measured Z coordinates after applying the count density-based calibration. The performance in the central region is not as good as with the collimator-based calibration, but still represents an improvement of 500 μm .

Figure 3.17

still allows for some improvement in those regions, while leaving good pixels undisturbed. Figure 3.18 shows the absolute error for three other pixels, one on the edge, one in the center, and one halfway in between. In general, there is some improvement in regions where the cathode/anode ratio performs poorly. Meanwhile, additional error in regions where it performs well is minimized.

3.4 Future Work for Improving Subpixel Position Reconstruction

The narrow neural net of Section 3.2.3 improves upon classical techniques for predicting the subpixel XY position of interactions. However, simple neural nets like the one used may fail to include the temporal information present in the waveforms. That is, individual waveform samples are not random, but highly correlated with those immediately preceding and succeeding them. More advanced models like convolutional neural networks, such as ResNet [40], may perform better at capturing this information, since




Figure 3.18: Absolute error between calculated and measured Z coordinates after applying the count density-based calibration to edge and interior pixels. There is generally improvement in the regions where the cathode/anode ratio performs the worst, while little to no additional errors are introduced in the good regions of the crystal.

convolution is inherently a correlated operation.

The count density-based voxel redistribution algorithm of Section 3.3.2 showed some promise, but could not match the calibrations achieved with collimators. One possible cause is event loss either away to the side of the pixel in X and Y or to traps which remove it from the spectrum altogether. The impact of these types of losses is minimized by the iterative scheme which tends to smooth the errors out over the entire column. Nevertheless, greater success may be realized by a full 3D implementation of the iteration algorithm. This will require solving the problems of charge sharing subpixel reconstruction, which will otherwise skew the densities between pixels. Alternatively, the gaps between pixels may be ignored altogether during the iteration phase.

Other techniques for finding the true subpixel voxel coordinates which do not rely on count densities may provide better, or at least complementary, information. Two possibilities are attenuation mapping and Compton imaging. If the detector is illuminated by sources with different energies, photopeak counts at different energies will drop off at different rates inside the crystal. In any given voxel, the ratio of counts among the different photopeaks can act as a signature of the voxel location independent from charge loss mechanisms. Alternatively, the Compton image is sensitive to the quality of position reconstruction [41]. The FWHM of the Compton image of a calibration source can be a source of additional feedback to iteration algorithms which attempt to arrive at the true voxel locations.

Chapter 4

Imaging of Pair Production Events

The important thing about electrons and protons is not what they are but how they behave, how they move.

- Paul A.M. Dirac

4.1 Pair Production Physics

4.1.1 Pair and Triplet Production and Annihilation

When a gamma ray of sufficiently high energy interacts with matter, it may convert into a pair of matter-antimatter particles. To satisfy conservation of momentum, the process must take place near a tertiary particle which can absorb excess photon momentum. This is frequently an atomic nucleus or, possibly, an electron. Additionally, the incident photon must have an energy above a certain threshold which depends on the type of particles created during conversion and the mass of the tertiary particle. Since electrons are the lightest charged particles, they have the lowest pair creation threshold. When the photon conversion occurs in the vicinity of a nucleus, the threshold is negligibly higher than twice the electron rest energy, $2m_e c^2$. This case is often referred to as pair production, and results in an electron, positron, and recoil nucleus. When conversion happens in the vicinity of an electron, the threshold is raised to $4m_e c^2$ to satisfy conservation of momentum and energy. This situation is called triplet production, and results in two electrons and a positron being emitted [42].

Owing to the higher threshold, the cross section for triplet production in the field

of an electron compared to the cross section for pair production in the nuclear field is about two orders of magnitude smaller in CdZnTe (CZT) for photon energies ranging up to 9 MeV. Through the rest of this chapter, pair production will refer exclusively to electron-positron creation in the field of a nucleus.

The positron will rapidly annihilate in dense matter, usually releasing two photons with about 511 keV of energy. It is sometimes possible for only one or more than two photons to be emitted, but those cases are rare and are neglected here. However, the daughter particles are often created with some significant kinetic energy. One common way for positrons to annihilate in matter is via the formation of an intermediate state called positronium, which sees the positron in a temporary bound state with an atomic electron. The capture of the positron and formation of positronium is inhibited by the positron's initial kinetic energy. Only once the positron thermalizes in the material is there a significant chance for positronium formation. There can therefore be a time delay between positron formation and annihilation, which depends on the energy of the positron [43].

4.1.2 Cross Sections for Pair Production

The cross section for pair production is a complicated function of the photon energy and polarization, target Z number, and energy and emission angle of the daughter particles relative to the parent photon. Because of the large mass difference between the recoil nucleus and the daughter particles, most of the excess energy from the photon after conversion goes into the kinetic energy of the daughters. Since most photons studied here have energies above 3 MeV, the daughter particles are often relativistic. This tendency results in a strong preference for forward emission of the daughters relative to the incident photon direction. For high photon energies much greater than $m_e c^2$, the cross section for emission at angle θ_0 is proportional to

$$\sigma \propto \frac{\theta_0}{\left(\frac{m_e c^2}{E_\gamma}\right)^2 + \theta_0^2}, \quad (4.1)$$

and is shown in Figure 4.1 for $E_\gamma = 100 \text{ MeV}$. For lower energies, the preference for forward emission is weaker, but still present [44].

After the creation of the electron and positron, the remaining photon energy is not




Figure 4.1: Approximate cross section for daughter particle emission following pair production into a particular angle away from the incident photon direction. The photon here had 100 MeV of energy. At lower energies, the cross section is more complicated and the preference for forward emission weaker, but still present.

necessarily split evenly between the electron and the positron. The cross section is approximately symmetric as a function of electron or positron kinetic energy, with only small asymmetries at low energies. For photon energies significantly greater than $m_e c^2$ and less than $137m_e c^2/Z^{1/3}$, the cross section is given by

$$\sigma \propto \frac{(E_e E_p)^2 + \frac{2}{3} E_e E_p}{(h\nu)^3} \cdot \left(\log\left(\frac{2E_e E_p}{h\nu m_e c^2}\right) - \frac{1}{2} \right) \quad (4.2)$$

where E_e and E_p are the electron and positron kinetic energies and $h\nu$ is the photon energy [45]. This relation is plotted for a 10 MeV photon in Figure 4.2.

4.2 Pair Production Telescopes

4.2.1 Pair Production Telescope Design

The task of determining incident photon direction is typically accomplished by detectors capable of reconstructing the tracks of the daughters. In the earliest orbital observatories these were spark chambers, which have since been superseded by silicon strip detectors.






Figure 4.2: Approximate cross section for emission of a positron with a particular fraction of the total available kinetic energy, for a 10 MeV photon. At lower energies the cross section becomes more uniform.

Since these have low density, and so a low cross section for photon interactions, they are sandwiched with thin foils of a high-Z material like tungsten to achieve the conversion. When the daughters stream out from the conversion point, their positions are tracked in 3D by successive layers of 2D-position-sensitive trackers. At the end of the track, they are stopped by thicker calorimeters (typically scintillators) so the total photon energy can be measured. By solving the conservation of energy and momentum from the measured tracks, the incident photon direction can be derived. Since the recoil nucleus is typically not detected, it presents a limit on the angular resolution achievable with these detectors.

The first gamma-ray observatory operating on these principles launched into space was the Small Astronomy Satellite-2 (SAS-2) in 1972, which achieved an angular resolution of $1.5 - 5^\circ$ depending on incidence angle. The satellite's spark chamber assembly is shown in Figures 4.3a and 4.3b. It was succeeded by the EGRET instrument on the Compton Gamma Ray Observatory, shown in Figure 4.4 [46], and then by FERMI-LAT which made the transition to silicon strip detectors from spark chambers [47].



(a) Diagram of SAS-2 spark chamber assembly. (b) SAS-2 spark chamber assembly before launch.
 Credit: NASA

Figure 4.3

4.3 Previous Work on High Energy Imaging with Pixelated CdZnTe

4.3.1 Detection of High Energy Photons with 3-D Pixelated CdZnTe

Previous generations of analog front-end readout ASICs coupled with the 3D CZT detectors used in this study only had dynamic ranges up to 3 MeV. This limited the detection of higher energy gamma rays to inter-module events only. Starting with the VAD_UM waveform digitizing ASIC, dynamic ranges up to 9 MeV became available. J. Xia performed spectroscopy of various Cl lines up to 7.4 MeV following activation by thermal neutrons from a PuBe source [49]. Various effects cause a degradation of energy resolution at high energies compared to 662 keV. Large charge cloud sizes cause more complex charge sharing and weighting potential crosstalk effects which are difficult to calibrate. Channel-by-channel nonlinearity further degrades the resolution at high energies.




Figure 4.4: Diagram of EGRET instrument. [48]

4.3.2 Compton Imaging of High Energy Photons >3 MeV

The VAD_UM ASIC also enabled imaging of high energy photons. No pair production imaging capability has previously been developed, so Compton imaging was the only practical modality. D. Shy studied in depth the artifacts arising in Compton images of the 4.44 MeV line from PuBe [50]. Prominent checkered-shaped artifacts were observed from charge sharing and weighting potential crosstalk effects. These were more pronounced at higher energies and bias interaction depth reconstruction to the anode side, skewing the Compton lever arms. Additionally, traditional sequencing algorithms such as mean squared difference (MSD) resulted in many more mis-sequenced events, partly due to pair production events which deposit charge on only two pixels instead of three. The sequencing algorithms also tended to erroneously reconstruct pair production events as Compton events with cone opening angles of 10° or 83.3° depending on

the event geometry. By developing FIL-MSD (“first-is-largest”) in place of MSD and using opening angle filters to detect and reject pair production events, Shy was able to obtain hotspots in the Compton image for 4.44 MeV photons. However, the improvements came by rejecting pair production events, which account for a significant fraction of interactions at energies above 4 MeV.

4.4 Pair Production Imaging Algorithm

4.4.1 Relevant Signal Induction Physics

Proper imaging of pair production events in CZT requires an understanding of how the energy deposition following photon conversion ultimately results in the detected signal. The discussion of Section 1.1.2 treated the charge liberated by a radiation interaction as a point. In reality, most of the photon energy is transferred to one or two initial particles. These then rapidly thermalize, mostly by further interactions with valence band electrons and lattice phonons. As they do, they liberate more electrons from the valence band and leave behind energetic holes, each of which may excite yet more electron-hole pairs. The final result is a cloud of mobile electrons and holes, which is separated by the externally applied electric field [51].

Fortunately, the Shockley-Ramo theorem still allows for easy calculation of the total induced charge by considering the signal as the superposition of the signals from each of the individual mobile electrons.¹ The calculated 3D interaction location is then actually the position of the charge-weighted centroid of the electron cloud. This presents a difficulty for imaging applications, in which the original interaction location is often desired but only the charge-weighted cloud centroid may be calculated.

The initial charged particles generated by photons are usually single electrons (in the case of photoabsorption and Compton scattering) or an electron-positron pair. At low energies, both of these have similar effective masses to the valence electrons they liberate during thermalization. As a result, large momentum transfers are possible. The path taken by the primary particle can be erratic, and the resulting electron cloud shape is chaotic. At higher energies, the primary particles have a higher effective mass due to relativistic effects, and smaller momentum transfers are possible until they have a

¹This is the principle behind the electron cloud reconstruction algorithms of Zhu and Xia.

chance to significantly thermalize. On average, this results in longer, more elongated electron clouds for high energy events, which may still be quite random in shape.

Because the magnitude of momentum transfers is less at higher energies, some of the information about the initial energetic particle direction is contained in the final electron cloud charge centroid, as well as the final resting position of the primary particle. An example charge cloud from Geant4 simulations illustrating this behavior is shown in Figure 4.5, while the distribution of the charge cloud centroid displacement is shown in Figure 4.6. As the energy deposited by the photon in Figure 4.6 increases, the charge cloud centroid is pushed further out in the initial direction of the primary energetic electron.




Figure 4.5: Example electron charge cloud after thermalization from a Compton scatter depositing 3.64 MeV of total energy. Each circle marks the site of ionization of valence electrons to create electron-hole pairs, with the size of the circle indicating how much energy the primary particle lost to electron-hole pair generation. Simulated with Geant4.




Figure 4.6: Mean charge cloud centroid location as a function of total energy deposited, from Geant4 simulations. The initial electron was emitted at the origin in the +X direction. As the initial electron energy increases, the charge centroid is biased further in the +X direction. The variance in Y at higher energies comes from poorer statistics at higher energies, as well as an intrinsically greater overall variance in the final position.

4.4.2 Algorithm Description

From equation 4.1, the daughter particles' momenta tend to be aligned with the incident photon momentum. Without electron tracking capability, if the direction in which the daughter particles were emitted could be estimated, the incident photon direction could also be inferred. From Figure 4.6, there is some information about the daughter particle direction embedded in the charge cloud centroid offset. The charge cloud centroid can be measured by the traditional methods described in chapter 1. The offset may be calculated by:

$$\vec{x}_o = \vec{x}_c - \vec{x}_{pp} \quad (4.3)$$

where \vec{x}_o is the charge centroid offset, x_c is the charge centroid and x_{pp} is the pair production site. However, there is no straightforward way to determine the initial pair production site.

From Figure 4.2, it can be seen that for some pair production events, the positron will receive very little energy compared to the electron. In this case, it will not travel far from the original pair production site before it is annihilated. The annihilation position

can therefore serve as a proxy for the pair production site. The two 511 keV photons may also be detected in CZT. While the annihilation site can then be inferred to lie on the axis connecting the two 511 keV depositions, where along the axis it lies is still unclear.

The algorithm is as follows:

1. Identify 3-interaction events that correspond to pair production and annihilation photon interactions.
 - One interaction is required to have the photopeak energy - 1022 keV
 - The other two interactions are required to be either 511 keV depositions or below the Compton edge energy corresponding to 511 keV, and on either side of the larger interaction in an approximate line.
2. Draw a line connecting the two smaller energy depositions (the “511 axis”).
3. Draw a line from the main interaction charge centroid to the nearest point on the 511 axis; this will be perpendicular to that axis. This line will serve as a lever arm for the backprojection of a cone of possible incident photon directions.
4. Backproject a cone with a particular fixed opening angle, only if the lever arm meets a length criterion.
5. The superposition of many cones will overlap with the source location.

There are two main complications. First, it is not guaranteed that the positron energy fraction will be small. A way to determine the positron energy fraction is required for the algorithm to work. Second, there is no guarantee that the closest point on the 511 axis to the main cloud centroid will be near the actual positron annihilation site. Since there is no correlation between the angle the annihilation photons are emitted at and the initial incident gamma ray direction, it could even be that the 511 axis passes straight through the main cloud centroid. In this case the lever arm will be undefined. In general, the lever arm may point in almost any direction based on the orientation of the 511 axis. It will be shown that both of these difficulties may be resolved by making an event cut based on the lever arm length. To illustrate how and why, three cases are investigated in detail: when the positron energy fraction is low, when it is high, and when the energy sharing is equal between daughters.

4.4.3 Pair Production Event Simulations

Modifications were made to the Orion Group’s in-house Geant4-based simulation code, T2. These enabled the full reconstruction and analysis of a radiation interaction. For every event, data is saved at every location of a physical process hit inside a crystal. The data includes the hit coordinates, energy deposited, interaction type (photoabsorption, pair conversion, Bremsstrahlung emission, ionization, etc.), the particle responsible for the interaction and its momentum before and after the interaction, and pointers to interactions immediately prior and subsequent in the history of the particle. This capability was used to investigate the geometries of pair production events in CZT.

Figure 4.7 shows the reconstructed event history of a pair production event, resulting in a main charge cloud and two 511 keV photoabsorptions from the annihilation photons. The main charge cloud has two primary components, ionization due to the daughter electron (blue) and due to the daughter positron (orange). In this particular interaction, the positron only had 17% of the total electron-positron kinetic energy. The orange track is correspondingly small, remains near the initial pair conversion site, and has little contribution to the overall charge cloud centroid of the main interaction (black + symbol).

The lever arm of Figure 4.7 (in green) is constructed from the estimated annihilation position on the 511 axis to the main charge cloud centroid. With this convention, an opening angle slightly less than 180° will result in a cone overlapping the true source location. The precise angle to use is determined partly by the distribution of equation 4.1 and partly by the kinematics of the electron cloud formation.

In the simulated event, the closest point on the 511 axis to the main charge centroid happens to be close to the true annihilation location. In general, since the 511 axis must intersect the annihilation site, the length of the lever arm is given by

$$L = D \sin \theta \tag{4.4}$$

where D is the displacement of the charge centroid from the annihilation site, and θ is the angle between the 511 axis and the vector connecting the annihilation site and the main charge centroid. The maximum length occurs at $\theta = 90^\circ$. In this orientation, $L = D$ and the closest point on the 511 axis to the charge centroid is exactly the annihilation site. So, by choosing events with long lever arms, the algorithm selects for




Figure 4.7: Simulated charge cloud detail for a pair production event in which the positron had 17% of the total available kinetic energy. The pink arrow shows the initial photon direction, and the green arrow shows the lever arm constructed by the algorithm described above.

those where the annihilation site is accurately determined.

If the positron energy fraction is high, the situation is similar, with some key differences. This scenario is illustrated in Figure 4.8. In this case, the electron had very little energy, so the electron contribution to the charge cloud (blue) is small while the positron's (orange) is large. The same definition of lever arm, from the estimated annihilation site to the main charge centroid, results in a vector pointing opposite the photon momentum. The ideal opening angle to use for backprojection in this case is a little more than 0° , opposite the case when the positron energy was low. Again, selecting events with the longest lever arms screens for those where the 511 axis is nearest the charge centroid close to the annihilation site.

In both of these cases, most of the available kinetic energy went to one or the other daughter particle. For purposes of charge cloud formation, the positron and electron behave identically. There is some slight difference from the Coulombic influence of the lattice nuclei on the positron, but this is screened by the shell electrons and is negligible here. The overall size of the charge cloud, or the displacement of the charge centroid, is




Figure 4.8: Simulated charge cloud detail for a pair production event in which the positron had 72% of the total available kinetic energy.

therefore about even in both cases. Figure 4.9 presents the same data as Figure 4.6 but explicitly as a function of initial charged particle energy. If a 4.44 MeV photon produces an electron/positron pair with uneven sharing of kinetic energy, the mean displacement of the charge centroid from the pair production site will be roughly 1 mm regardless of whether the electron or the positron received the most energy.

The situation is slightly different if the energy fraction is close to 50%, as in Figure 4.10. In this case, for a 4.44 MeV incident photon, the mean displacement of the electron and positron sections of the charge cloud are both around 0.5 mm. There is still a correlation between the photon direction and the emitted daughters, even when their energies are lower. As a result, they contribute about equally to the total measured charge centroid location, and the overall displacement will be less than when the kinetic energy went mostly to one daughter. This is the reason charge clouds from pair production events are on average more compact than those from other types of interactions, an effect first observed by S. Anderson [52].

In the geometry of Figure 4.10, the lever arm is reconstructed pointing 90° away from the incident photon direction. This can be considered an intermediate case between Figures 4.7 and 4.8. Choosing the events with the longest lever arms still selects for




Figure 4.9: Average charge centroid displacement as a function of initial particle energy.

those where the annihilation site is reconstructed most accurately, but in this case the overall cloud is smaller and the centroid displacements are less. This provides a clue on how the energy fraction may be determined: by choosing the events with the longest lever arms overall, the algorithm prefers ones where the energy fraction was uneven. The optimum length threshold increases with the photon energy. Figure 4.11 shows the lever arm length as a function of positron energy fraction with a fixed photon energy of 6 MeV. A threshold of about 1.3 mm will cut out most events with even energy fraction. Higher thresholds will be more selective for uneven energy fractions, at the cost of reduced statistics.

In fact, there is an asymmetry in Figure 4.11. Even discounting events with balanced energy fractions, those with low positron energies tend to have longer lever arms than those with low electron energies. For example, when the positron receives 10% of the kinetic energy, the mean lever arm length is 1.3 mm, while when the electron receives 10% of the kinetic energy the length is only 0.9 mm. This asymmetry comes from the extra variance in the annihilation position when the positron has significant energy. The positron annihilates at the end of a random walk, which has a chance of bending back in towards the cloud center. Put another way: the lever arm length depends on the charge centroid location and the annihilation site. When the energy fraction is uneven, the centroid location is affected by the variance of either the electron or positron track (whichever is longer). This is symmetric with energy fraction, since




Figure 4.10: Simulated charge cloud detail for a pair production event in which the positron had 43% of the total available kinetic energy.




Figure 4.11: Charge cloud displacement from 511 axis vs. positron energy fraction for simulated pair production events from a 6 MeV photon beam.

positrons and electrons are identical when it comes to forming the overall charge cloud. The annihilation site is always affected only by the variance in the positron track. This is much lower when the positron has little kinetic energy. Overall, there is less variance in the lever arm length at low positron energy fractions. This results in less blurring

down to shorter lever arm lengths.

The asymmetry is fortunate. From Figures 4.7 and 4.8, the optimum cone opening angle to use is different when the positron energy is low or high. If the energy fraction is misassigned, and the wrong angle is used, it will result in image artifacts 140° to 180° away from the source location. This is analogous to missequencing artifacts in Compton imaging. Setting a lever arm length threshold solves these problems by selecting for events with low positron energy fractions, so the right opening angles can be used most of the time. Note from Figure 4.11 that the cut will not be perfect, and there will still be some missequencing artifacts expected.

The discovery of S. Anderson that pair production events tend to have smaller charge cloud sizes could also be used to identify uneven charge sharing fractions. This is simulated for 6 MeV photons in Figure 4.12. As expected, events with balanced charge sharing between electron and positron have the smallest cloud sizes. However, there is not much asymmetry between low and high β^+ kinetic energy as in Figure 4.11. Therefore, the electron cloud size is not a useful metric to determine the exact energy sharing fraction.




Figure 4.12: Charge cloud size (defined as the maximum separation between any two energy depositions in the same cloud) vs. positron energy fraction for simulated pair production events from a 6 MeV photon beam.


4.4.4 Pair Production Imaging Simulation Results

The algorithm of section 4.4.2 was applied to 4.4 MeV and 6.1 MeV photon pair production events simulated with Geant4. These correspond to prominent lines resulting from neutron interactions with Be and Cl respectively. The results are shown in Figures 4.13a and 4.13b. A red cross marks the true source location. In both cases, an opening angle of 140° was used, determined from simulation. Additional blurring of about 30° for the 4.4 MeV data and 40° for the 6.1 MeV data was applied to each cone to help form an image with limited statistics. Only a single $2\times 2\times 1.5$ cm³ CZT module was simulated, with the source illuminating the detector from the cathode side.


The algorithm is able to properly reconstruct a hotspot in both cases. For the 4.4 MeV data, a lever arm threshold of 1.3 mm was used. This removed most of the balanced energy fraction events, but left a considerable number of high positron energy events. These are responsible for the “missequencing” artifact at $\theta = 270^\circ$, about 180° away from the true hotspot. For the 6.1 MeV photons, a much more aggressive threshold was used, around 2 mm. This removed more of the high positron energy events, resulting in less of an artifact at the $\theta = 270^\circ$ location. However, it caused the overall statistics to be lower as well, and greater cone blurring was required to form a comparable image. Also, the opening angle to use in theory depends on the energy fraction. The higher length threshold selected events with lower energy fractions, but the same opening angle was used, resulting in a slightly misshapen hotspot compared to Figure 4.13a.

4.4.5 Experimental Verification

To verify the imaging algorithm experimentally, several experiments were carried out with a PuBe source. This source produces 4.44 MeV photons from the (α, n) reaction with Be and the relaxation of the resulting $^{12}\text{C}^*$. A new 3×3 system, Orion- γ , was assembled similar to the 3×3 systems (Orion- α , Orion- β) previously described. This system has an aluminum housing, several ADCs for digitizing ASIC and environmental sensor output data, and an aluminum detector chamber containing a 3×3 array of detectors. Thermal control is provided by a Peltier cooler to dissipate the heat produced by the ASICs and keep the detectors at room temperature. A bias of -3000V is produced by an onboard high voltage generator and distributed to the detector cathodes. An FPGA provides an interface between the front end detectors and an external data acquisition



(a) Pair production image reconstructed from simulated 4.4 MeV photons. The red cross marks the true source position. Note the missequencing artifact at $\theta = 270^\circ$.




(b) Pair production image reconstructed from simulated 6.1 MeV photons. A significantly higher lever arm length threshold was used, resulting in lower statistics but a much suppressed missequencing artifact at $\theta = 270^\circ$ compared to Figure 4.13a.


Figure 4.13

computer. Figure 4.14a shows a CAD model of the detector system with some top components hidden, and 4.14b shows the assembled system.

The system used 6th generation Redlen CZT detectors directly attached to VAD_UMv2.2 waveform digitizing ASICs. The ASICs were operated at 40 MHz sampling rate, and



(a) 3D CAD model of the 3x3 Orion detector system, courtesy of Y. Zhu.




(b) Orion- γ used in the pair production imaging experiments.

Figure 4.14

up to 8 neighbor pixels were read out for each triggered pixel. The subpixel position algorithms described in section 1.4.2 were used, including the reconstruction techniques developed by Zhu for edge, corner, and 2-pixel charge sharing events. Proper subpixel position reconstruction is crucial for the reconstruction algorithm.

When choosing detector modules for Orion- γ , the best detectors were intentionally avoided to preserve them for any future experiments that require the absolute best performance. Nevertheless, careful cleaning of the motherboard and high voltage components, assembly of a new high voltage distribution board, creation of new high voltage connectors and replacement of several motherboard components resulted in state-of-the-art energy resolution matching the best previously achieved with other detectors, 0.34% FWHM at 662 keV for individual modules (Figures 4.15a and 4.15b). When all 9 crystals are used simultaneously, the overall resolution from all modules combined degrades to 0.39% at 662 keV for single pixel events and 0.55% for all events due to increased readout noise from loading on the shared output trace. Further details of the precise detector characterization are given in chapter 3.

Two experiments with the PuBe source were run. A 7 MeV dynamic range was used for both. The first measurement had the cathode plane facing directly towards the PuBe, which was about 42 cm away. The second measurement had Orion- γ rotated by about 30° and placed about 36 cm away. With this arrangement, the hotspot should shift by the amount rotated between the two datasets, showing true imaging capability.



0.35	0.35	0.32	0.33	0.32	0.34	0.35	0.35	0.33	0.32	0.34
0.35	0.33	0.31	0.34	0.34	0.33	0.32	0.34	0.35	0.33	0.34
0.36	0.33	0.35	0.35	0.35	0.34	0.34	0.35	0.36	0.33	0.33
0.35	0.34	0.34	0.33	0.33	0.34	0.32	0.34	0.34	0.36	0.34
0.34	0.35	0.34	0.34	0.34	0.34	0.34	0.34	0.31	0.34	0.33
0.34	0.33	0.34	0.34	0.36	0.34	0.34	0.35	0.34	0.34	0.34
0.35	0.32	0.33	0.35	0.33	0.35	0.35	0.33	0.34	0.34	0.36
0.35	0.35	0.33	0.35	0.34	0.34	0.34	0.33	0.31	0.33	0.33
0.34	0.35	0.34	0.35	0.34	0.35	0.33	0.34	0.32	0.33	0.32
0.36	0.36	0.33	0.35	0.34	0.36	0.34	0.32	0.32	0.33	0.36
0.36	0.36	0.35	0.34	0.36	0.34	0.35	0.33	0.34	0.33	0.31

(a) Single pixel spectrum from detector 6RID-36 in Orion- γ tested individually. FWHM is 0.34% at 662 keV. (b) FWHM of the spectrum from each pixel of 6RID-36 in Orion- γ , in % at 662 keV.

Figure 4.15

The experimental setups of both runs are shown in Figure 4.16. Both runs lasted for approximately 50 hours.




Figure 4.16: Experimental setups for run 1 (left) and run 2 (right).


The spectra for the two experiments are shown in Figures 4.17a and 4.17b. An

experimental error led to a much higher threshold, ~ 150 keV, being used for the second experiment. This caused problems in the analysis, described later. The 4.44 MeV lines are present in both sets of spectra. 3-interaction events were used to reconstruct the pair production image. Side neighbor charge sharing events were combined into single interactions and their subpixel position reconstructed using the opposing neighbors. To verify the imaging capability, the pair production images were compared to Compton images of the 372 and 413 keV lines from the same source.


Many low energy lines are present in the spectrum below 1.1 MeV. These can be Compton imaged to find the source location and compare with the pair production images. Figure 4.18 shows this region of the 2-interaction spectrum and Compton images reconstructed by simple backprojection from several of the peaks: 372, 413, 843, and 1014 keV. The 372 and 413 keV peaks are emitted by the PuBe source directly, and the hotspot represents the source location. The 843 and 1014 keV images (as well as images of many other peaks in the 500 - 1100 keV range) have dark spots on the cathode- and anode-facing sides, and hot counts all around. From Figure 4.14a, the detectors are enclosed in an aluminum chamber which is significantly thicker around the sides than it is on the cathode- and anode- planes. These peaks in the spectrum originate from neutron activation in the aluminum of the detector housing. For example, a (n,p) reaction with ^{27}Al yields ^{27}Mg , which undergoes beta decay back to ^{27}Al and releases either an 843 or 1014 keV γ -ray.

Figure 4.19 shows the 2-interaction spectrum for the second experiment and the Compton image of the 413 keV peak. The high threshold used for the second run cut into the useful statistics and prevented the formation of any usable image from the 372 keV line. The 413 keV line could still be imaged, but the Compton edge test could not be used, since at least one interaction would always be below the threshold. The Compton edge test typically guarantees the correct sequencing in Compton imaging. Since it could not be used here, the FIL paradigm was followed instead, in which the larger energy deposition was sequenced first. This is typically used for higher energies, but gives acceptable results at lower energies. A hotspot is still reconstructed, albeit with some missequencing artifacts 180° away. The hotspot is shifted about 30° upwards in ϕ compared to Figure 4.18.

Because of extra detector uncertainties and blurring effects, the calculated lever arm lengths are biased to longer lengths than seen in simulation. The distribution is given



(a) PuBe source spectra from experiment 1 with Orion- γ .



(b) PuBe source spectra from experiment 2 with Orion- γ .

Figure 4.17

in Figure 4.20. Example event geometries are given in Figures 4.21a and 4.21b. The thresholds used for lever arm cuts are correspondingly increased. Since the threshold depends on energy, the same threshold of 1.5 mm was used. An upper threshold of 3 mm was also used; events with arms longer than this were rejected. The great majority of these events either have improbably long positron travel paths, candidate annihilation sites outside the crystals, or both, as in Figure 4.21b. These are likely to result from

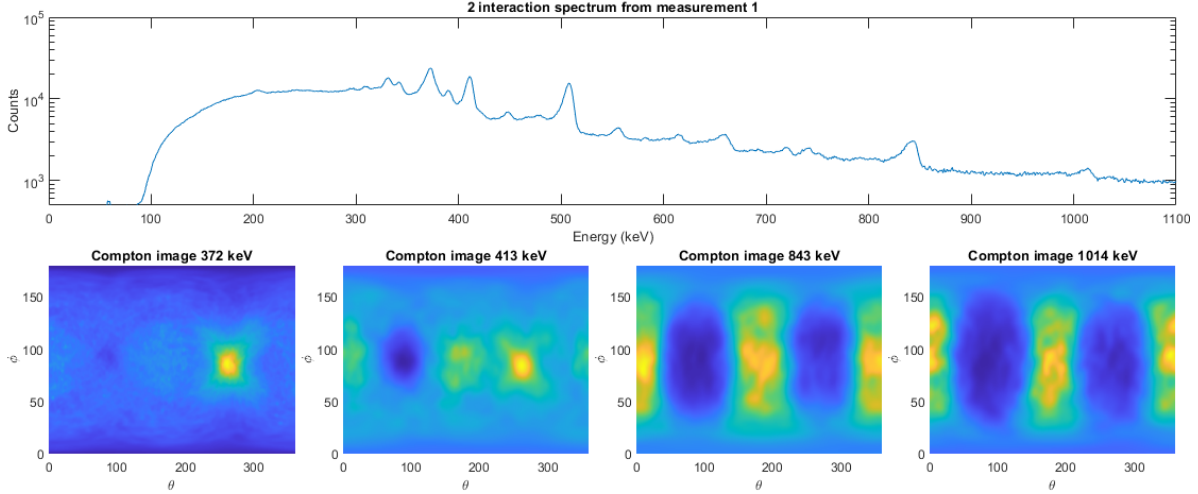


Figure 4.18: 2-interaction spectrum for experiment 1 and Compton images corresponding to several energy lines: 372, 413, 843 and 1014 keV.

multiple scatters rather than pair production.

After all event cuts, the final pair production images from the two experiments are shown in Figures 4.22a and 4.22b. In both images, the red cross marks the location of the hotspot in the corresponding Compton images at lower energies. As the source is shifted 30° off the cathode plane, the hotspot in the pair production image follows it. 20° of blurring was used for each cone to compensate for the low statistics.

There are two significant sources of bias in the system that account for some aspects of the images. First, there is the “missequencing” artifact that is expected to result from the misidentification of the positron kinetic energy fraction. As in simulation, this is expected to produce artifacts surrounding a point 180° away from the true source location. These are present in both images; in Figure 4.22a they cluster around (90,90) while in 4.22b they are seen as a ring of cones around (60,90).

The second effect stems from the detector system geometry. The 3x3 array of CZT crystals has a very nonuniform aspect ratio. The lever arm cut tends to select events where the 511 axis is oriented perpendicular to the incident photon direction. If the photon is incident from the cathode or anode side, there is on average ~ 40 mm of CZT for each 511 keV annihilation photon to interact with. If the photon is incident from the side, however, a perpendicularly emitted 511 keV photon has a much greater chance of escaping entirely. The overall detector system therefore has higher sensitivity in the

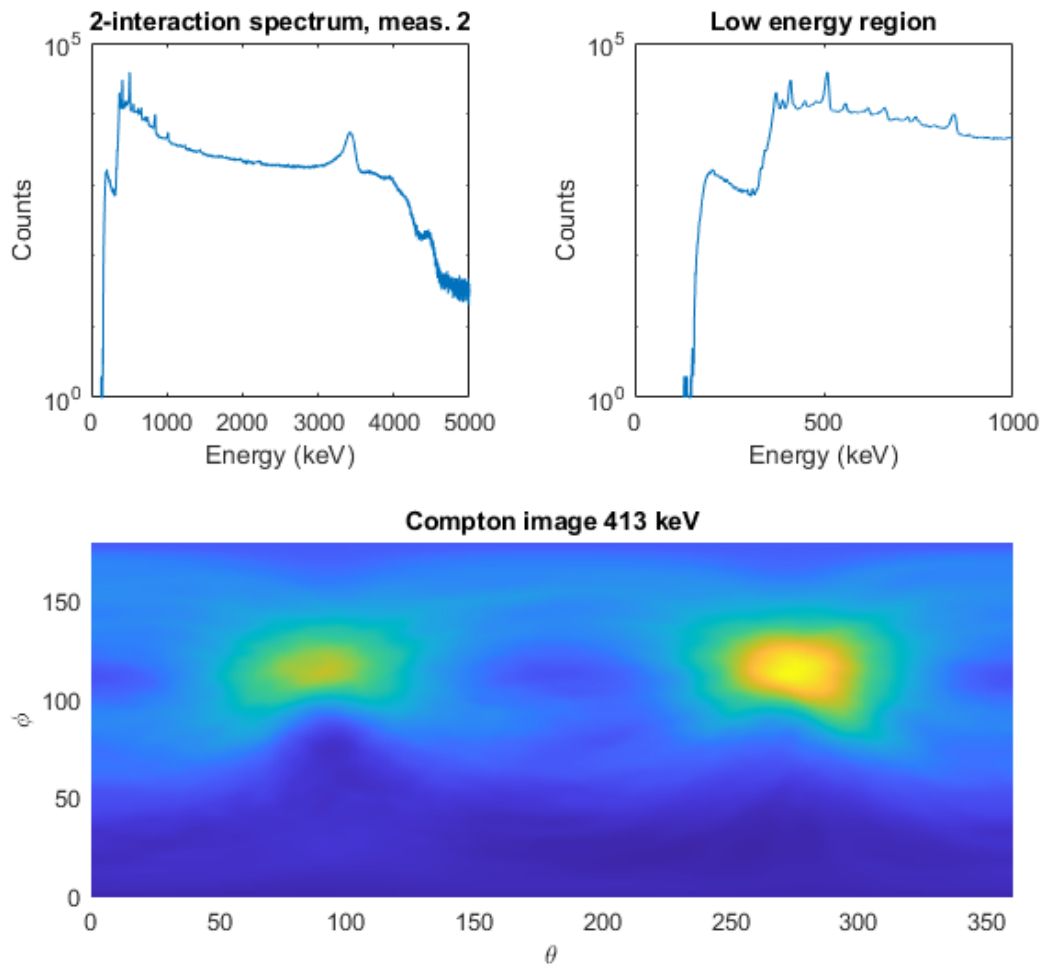


Figure 4.19: 2-interaction spectrum for experiment 2 and Compton image of the 413 keV peak.

cathode and anode directions than along the sides. Uncorrected, this tends to bias the images towards the cathode and anode sides. This effect is likely responsible for the oblong hotspot of Figure 4.22b that stretches down towards the cathode side.




Figure 4.20: Distribution of calculated lever arm lengths for experiment 1.


4.5 Future Work: Filtered Backprojection, MLEM, and Particle Tracking

The biases described in the previous paragraph may be addressed in several ways. Either by simulation or calculation, a point spread function can be derived and used to filter out the detector response. This should correct the biases in the reconstructed figures and lead to an improvement in the image resolution and accuracy, similar to that achieved for Compton imaging. Analytical derivation of the point spread function is difficult due to the highly nonlinear nature of the algorithm, in particular the effects of detector position reconstruction blurring and the lever arm length cut. A simulated point spread function would be more straightforward to use. Additionally, with Compton imaging, simulated point spread functions typically outperform analytical ones. With Compton imaging there is also the possibility of using experimentally measured point spread functions, by characterizing the detector response to known sources at fixed calibration angles. For pair production imaging, this is likely not feasible due to the low interaction cross sections at high energies.


Another option is an iterative reconstruction algorithm based on maximum likelihood expectation maximization (MLEM). MLEM methods incorporate system models to de-

termine the most likely source distributions given an observed array of interactions. These outperform filtered backprojection for point source reconstruction in Compton images. However, they rely on a good system model to accurately model the probabilities of detecting particular sequences of events. This system model may be challenging to construct for pair production interactions.

Lastly, charge cloud shape reconstruction has been demonstrated in 3D pixelated CZT detectors by Zhu and Xia. With sufficient reconstruction fidelity, this may be used to find the original emission direction of the electrons or positrons, and perform pair production imaging in a similar way to the true electron-tracking detectors. Even if the initial direction cannot be accurately determined, coarse information about the electron cloud shape may help inform the event selection process. This may complement or improve on the statistics from the simple lever arm length test described in this chapter. However, the cloud reconstruction algorithms have only been tested with low energies or with muon events, and only for single-pixel events. They will need to be validated at higher energies, and may need to be extended to charge sharing interactions.




(a) Geometry of an observed candidate pair production event with a lever arm length of 1.55 mm.




(b) Geometry of an observed candidate pair production event with a lever arm length of 34 mm. Aside from an improbably long positron travel path, the candidate annihilation location is outside any crystal, in air. This event is rejected on these two grounds, and is likely instead a multiple scatter.

Figure 4.21



(a) Pair production image at 4.44 MeV reconstructed from the experiment 1 data. The red cross marks the hotspot location in the corresponding 372 and 413 keV Compton images.



(b) Pair production image at 4.44 MeV reconstructed from the experiment 2 data. The red cross marks the hotspot location in the corresponding 413 keV Compton image. The maximum of the pair production hotspot is shifted about 30° off the cathode plane, matching the Compton image. Low statistics and detector sensitivity cause an oblong hotspot stretching towards the cathode side.

Figure 4.22

Chapter 5

Perovskite Radiation Detectors

One shouldn't work on semiconductors, that is a filthy mess; who knows whether any semiconductors exist.

- Wolfgang Pauli

5.1 Alternative Materials to CdZnTe

5.1.1 The Search for the Ideal Detector

Radiation detection is a heavily application-driven field. Detectors are used for scientific research, nuclear nonproliferation [53], power plant monitoring [54], and various other ends. The requirements on detectors vary greatly between different application spaces, and sometimes different optimizations conflict with each other. A common tradeoff is between efficiency and resolution. High absolute efficiency is often achieved by large volume detectors such as scintillators. This is required for calorimeters on spacecraft, which must fully stop high energy particles at a relatively low flux [55]. However, scintillators do not have the best energy resolution. For applications where resolution is important (such as grading of special nuclear material) high purity germanium (HPGe) is commonly used [56].

CZT is in the semiconductor family of detectors, which tend to target high energy resolution applications. The best resolution demonstrated from large pixelated CZT is about 2 keV for 662 keV photons, which is about 2-4 times worse than HPGe. The advantage of CZT is that it achieves this resolution at room temperature, while HPGe

requires cooling. The bulk of nitrogen dewars limits the portability of HPGe detectors that rely on cryogenic cooling, while the mechanical coolers of handheld units add complexity and require extra maintenance and startup time. The desire for a handheld, room temperature, high resolution imaging spectrometer drove the early development of CZT technology to the level it is now.

However, CZT is not without any drawbacks. The crystal growth process is difficult, and the yield of detector-grade crystals is low compared to silicon or HPGe. The largest volume achieved for single crystal pixelated CZT detectors is 40x40x15 mm [57]. Electron mobilities are around 100 times less than for HPGe, which also limits the size of practical detectors. Larger detector volumes are achieved by tiling arrays of crystals. Coupled with the difficulty of production, this leads to expensive instruments. A range of other semiconductors are being investigated which offer the same promise of higher energy resolution than scintillators and room temperature operation, while having a cheaper or simpler manufacturing process. Some have characteristics that may make them more suitable than CZT in some applications, and others may open new application spaces entirely.

5.1.2 TlBr, HgI₂ and Other Materials

TlBr is a binary ionic semiconductor with a higher bandgap and density than CZT, about 2.6 eV and 7.56 g/cm³ respectively [58]. It presents a higher effective Z, leading to a greater stopping power. It has a lower melting point than CZT, and detector quality crystals may be made from 3N or 5N purity material. It is self-compensating, leading to a higher resistivity without doping [59], avoiding the segregation problems that make CZT growth difficult. These factors could give a lower cost of manufacture, an advantage over CZT. Energy resolutions of around 1% at 662 keV have been demonstrated with some TlBr detectors for single pixel events [60]. This does not match the best resolution achieved with similar CZT detectors, but is sufficient for many practical applications.

However, TlBr suffers from persistent problems with ionic mobility. TlBr reacts with some common contact metals, decreasing the stability and detector lifetime under bias. For example, Au from electrodes may migrate into the bulk and form AuBr, degrading the detector response. The uniformity of detector response is also a problem. Some attempts have been made to mitigate these problems by using higher purity material,

different surface treatments, different contact metals, and bias switching schemes. Bias switching limits the practicality of these detectors. Tl metal for contacts can help with contact degradation and polarization issues, but at the cost of contact toxicity [61].

TlBr is transparent to visible yellow light, and has a high index of refraction of around 2.6, which enables its use as a dual charge induction-Cherenkov detector [62]. Classic pixel and planar contacts on two faces read out the charge induced by interacting photons, while silicon photomultipliers on the crystal side walls can capture a signal from the Cherenkov light produced by the interaction of fast charged particles, such as the initial kinetic electrons produced by Compton scatters or photoabsorptions. The Cherenkov signal is very fast, giving these detectors superior timing resolution compared to CZT, while achieving better energy resolution than what is provided by fast scintillators. This may allow for unique applications in PET. This is an example of the unique properties of an alternative semiconductor material enabling techniques not feasible with CZT.

Other materials have been investigated as alternatives to CZT. Many come from one of two categories. Some are CZT-like covalent compounds, which aim for qualities similar to CZT while making the synthesis easier. These include CdMnTe (CMT) and CdZnTeSe (CZTS). Detector-grade material of this type has been synthesized, with spectra reported from CMT [63] and CZTS [64] showing similar performance to early CZT. Other materials are ionic compounds like TlBr. The most successful from this category was historically HgI₂. Large pixelated HgI₂ detectors were fabricated, up to 1 cm thick. Energy resolutions as good as 0.9% at 662 keV were obtained from individual pixels [65]. HgI₂ suffered from problems of hygroscopicity, sublimation, and polarization, leading to challenges with device stability.

5.2 Introduction to Perovskites

5.2.1 Perovskite Structure

Perovskites comprise a family of crystal structures with the chemical formula ABX₃, where A and B are positively charged cations and X is a negatively charged anion. The first perovskite discovered was the naturally occurring mineral named perovskite, CaTiO₃ [66]. Others have since been discovered and synthesized, which swapped out the cations or anion, leading to a diverse range of physical and electrical properties.

Perovskites may have cubic, tetragonal, or orthorhombic phases. Figure 5.1 shows the structure of the orthorhombic phase of CaTiO_3 .




Figure 5.1: Orthorhombic phase of the perovskite CaTiO_3 , the original naturally occurring mineral known as perovskite. Dark blue are the Ca ions, light blue is Ti, and red is O. From materialsproject.org.

5.2.2 Applications

Perovskite materials have found a wide range of applications. One of the most prominent is in photovoltaics. The invention of perovskite solar cells in 2009 [67] led to a large research focus on the material family. As a result, the efficiency of perovskite solar cells increased from 3.8% in 2009 to 25.8% in 2021 [68].

The function of solar cells and x-ray and gamma radiation detectors are very similar. Both respond to photons with electrical signals. Many materials used for solar cells, such as silicon and germanium, also succeed in radiation detection applications. Following this trend, several perovskites are under investigation as semiconductor radiation detectors. The most successful in this category is the all-inorganic CsPbBr₃ [69], with hybrid detectors CH₃NH₃PbBr₃ and CH₃NH₃PbI₃ also showing promising results. Detection of 59.5 keV gammas from ²⁴¹Am has been demonstrated, as well as 662 keV gammas from ¹³⁷Cs. The reported energy resolutions are on the order of 4% at 662 keV [70].

Radiation detectors based on perovskites have several possible advantages over CZT. The synthesis may be easier, giving higher yield at lower cost. Some perovskite crystals are grown by vertical Bridgeman processes, similar to CZT, but at lower temperature, with lower cost precursors [71]. Others are grown by precipitation from solution, offering the potential of fast crystal growth at very low cost [72]. Additionally, many perovskites evaluated for radiation detection have decently high mobility-lifetime products for both electrons or holes, though one will typically dominate over the other. The chance to obtain significant signals from both polarity carriers may open up new algorithms for improving the detector response, though the classic techniques described in chapter 1 will need to be modified.

Both the furnace grown and solution grown perovskites have been studied at the University of Michigan. The solution grown samples were primarily produced by the University of North Carolina at Chapel Hill, while the furnace grown samples were produced by Northwestern University. Figure 5.2 shows the timeline of perovskite testing at the University of Michigan. Rapid progress was made on the furnace grown samples, with crystal volumes and energy resolution both improving greatly in the period from 2018-2022. Solution grown samples were difficult to obtain spectra from, but good ¹³⁷Cs and ⁵⁷Co spectra were finally obtained in 2022.

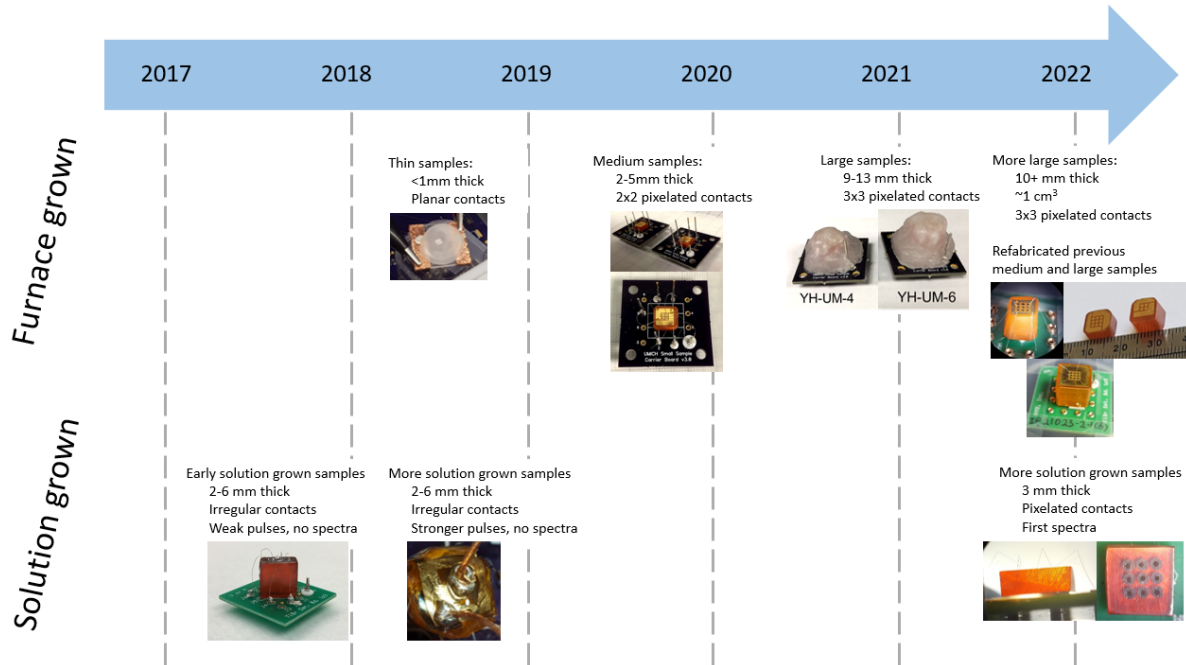


Figure 5.2: History of perovskite testing at the University of Michigan.

The remainder of this chapter is divided into six sections. The following four cover the experimental results from the three generations of furnace grown perovskite detectors (small, medium, and large) and the hybrid solution grown samples. The fifth section discusses conditioning effects, and the final section covers principal component analysis (PCA) techniques applied to perovskite detectors.

5.3 Perovskite Detectors for Gamma Rays

5.3.1 Early planar samples

Beginning in 2018, thin samples of CsPbBr_3 have been tested at the University of Michigan. A representative example of one such detector is shown in Figure 5.3. This sample was $0.9\mu\text{m}$ thick, and had an irregular shape about $3\times 3\text{mm}$ in size. The bulk resistivity of this CsPbBr_3 is fairly low, about $10^8\Omega\text{-cm}$ compared to about 10^{10} for detector-grade CZT. This would normally lead to a very high leakage current when bias is applied, making the output signal unusable. The leakage current is limited by an asymmetric contact structure, with one 3×3 face covered in Au and the other in either

Ga or a eutectic Ga-In alloy. These form a Schottky barrier at the Ga interface. The device is reverse biased so the Au contact is positive relative to the other. In the case of this detector, biases from 50 - 500V were placed on the Au contact. The actual detector in Figure 5.3 is the small white rectangle in the center of the grey circle. The top face is covered in Ga, leading to the pale color; CsPbBr₃ itself is either orange or red depending on purity and dopants. It is somewhat hygroscopic, so the sample was covered in wax to protect against moisture intrusion. Small wires were routed out and attached to contact pads made of copper tape at the ends of the plastic substrate, to which alligator clips connected the crystal to a biasing and readout system.





Figure 5.3: Thin planar CsPbBr₃ sample grown at Northwestern University and tested at the University of Michigan.


Signals were read out from eV-5093 discrete preamplifiers capacitively coupled to the electrodes. These were digitized by GaGe card high speed digitizers at up to 125 MHz. The same preamplifiers and digitizers were used for all perovskites in this study. The detector was exposed to ²⁴¹Am, ¹³⁷Cs, and ⁵⁷Co from anode and cathode sides.

The depth estimation techniques described in section 1.4.1 cannot be used with this simple planar geometry. Both anode and cathode waveforms are dependent on depth of interaction, so their ratio does not give a simple measure of depth. Likewise, since in these perovskites there may be a significant contribution to the signal from drifting electrons as well as holes, the waveform rise time does not have a monotonic mapping to the depth of interaction. Nevertheless, there is some mapping from rise time to depth which can be used to correct the signal amplitude and improve the resolution, as well as obtain some information about the material properties. Figure 5.4a shows the 2D

histogram of anode waveform rise time vs. amplitude for a ^{57}Co measurement. The bright trace of the 122 keV photopeak is visible, as well as the 134 keV peak and various escape peaks from Pb. Similar to the depth-gain curves for pixelated detectors, there is a region of short rise time which correlates with lower signal amplitude and worse resolution. In the case of pixelated detectors, this is usually due to weighting potential deficit and the small pixel effect. In the case of planar contact geometries, this is more due to different carrier properties.



(a) 2D histogram of planar waveform rise time vs. signal amplitude following exposure to a ^{57}Co source. (b) Raw ^{57}Co spectrum from 0.9mm thick planar CsPbBr_3 sample.



(c) ^{57}Co spectrum of Figure 5.4a after cutting out events with rise times less than 110 samples. The 122 and 134 keV peaks become visible and separable.

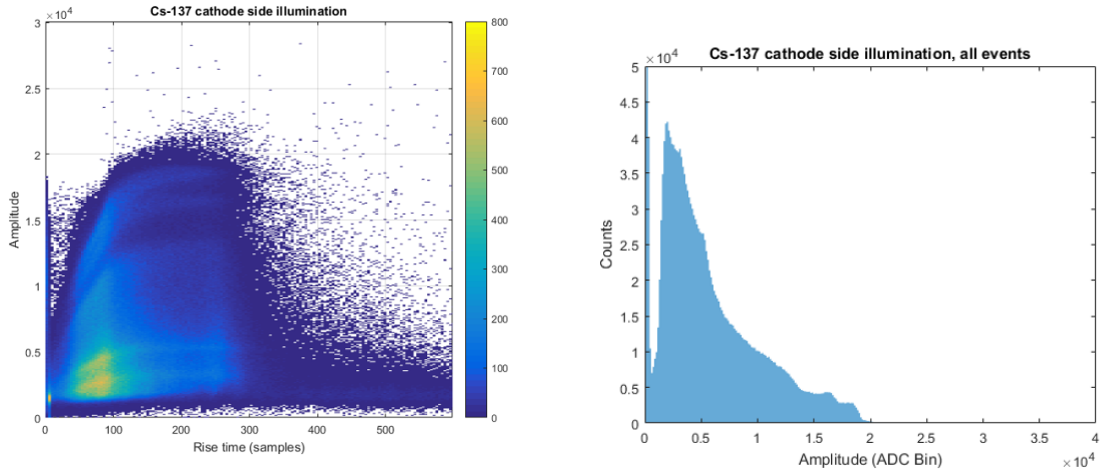
Figure 5.4

The thin planar samples sufficed to show the promise of the material. Good preliminary spectra were obtained from ^{57}Co , showing the ability to separate the 122 keV and 134 keV peaks. However, the small thickness hindered its ability for higher energy spectroscopy. Geant4 simulations showed that the average size of the electron-hole clouds created by 662 keV energy deposition in CsPbBr_3 was about $400\mu\text{m}$, which is significant compared to the sample thickness of $900\mu\text{m}$. In other words, if the interaction happens within a few hundred microns of the crystal surface, the electron may stand a good chance of escaping the crystal. This will lead to incomplete energy deposition and a degraded spectral response. Figures 5.5a - 5.5c show the Cs spectra obtained with this device. There is indeed a region of poor charge collection efficiency, corresponding to events with shorter rise times.


5.3.2 Medium thickness Pixelated Detectors

Following the results from the thin samples, several thicker detectors were fabricated with 2x2 pixel patterns. The pixels had a pitch of 1.2 mm and were made of Au. Since the Ga containing anodes have very low melting points, they are impractical for large-scale production of devices. Instead, Bi contacts were used. Ga was initially selected for its work function forming a good Schottky barrier with CsPbBr_3 and for being relatively chemically inert when in contact with the perovskite. Bi is similar, though slightly less ideal, in these respects. Sample thicknesses ranged from 2.75 mm (sample YH-UM1) to 4.64 mm (YH-UM3). The Bi contact was biased positively, so the devices operated in hole collection mode. The maximum bias each device could withstand was correlated with its thickness. YH-UM1 could be biased up to 250 V, while YH-UM3 was biased to 500 V. Past this point, anode waveforms became unstable, precluding depth estimation, though pixel waveforms remained stable. YH-UM1 achieved about 2.2% resolution at 662 keV after depth corrections for three out of four pixels, with the fourth suffering from connection issues. YH-UM3 also suffered from connection problems on one pixel, but overall performed better, reaching 1.57% at 662 keV after depth correction from three out of four pixels. YH-UM3 and a 3.57 mm thick crystal (YH-UM2) are shown in Figure 5.6a, as well as the initial corrected spectrum in Figure 5.6b.

Both YH-UM1 and YH-UM3 showed similar behaviors regarding leakage current and stability. Immediately upon applying bias, leakage current was high and waveforms were



(a) 2D histogram of planar waveform rise time vs. signal amplitude following exposure to a ^{137}Cs source. (b) Raw ^{137}Cs spectrum from 0.9mm thick planar CsPbBr_3 sample.





(c) ^{137}Cs spectrum of Figure 5.5a after cutting out events with rise times less than 165 samples. The 662 keV photopeak and Pb escape peak become visible, though the resolution is about 7%.

Figure 5.5

somewhat unstable. The most affected the waveform was from the large, planar anode contact. With the bias held high, the leakage current decreases exponentially over time and eventually the waveforms stabilize. This is an initial short-term conditioning period different from the long-term conditioning seen with TlBr .

After a couple of days at bias, the anode waveforms from YH-UM3 became unstable,




(a) 4.64 and 3.57 mm thick 2x2 pixelated CsPbBr₃ detectors before encapsulation with wax. The pixels and guard ring are Au, while the back planar anode is Bi. In operation, the anode was biased up to +500V.

(b) Initial corrected ¹³⁷Cs spectrum from the 4.64 mm thick sample (YH-UM3). The energy resolution is 1.57%.

Figure 5.6

rendering depth correction impossible and limiting the useful characterization that could be done. The crystal was removed from bias, allowed to rest for a few days, and then tested again in a thermal chamber at a slightly elevated temperature of 30°C. Figure 5.7a shows the leakage current vs. time as the detector is being ramped from room temperature up to 30°C. Once it reaches a stable temperature, the leakage current stops increasing, but still shows some instabilities during a brief conditioning period. Figure 5.7b shows the energy resolution during the same period of time. When the temperature and leakage current are changing, the energy resolution is poor. Previous experiments with CZT have shown that the photopeak position depends on temperature, and that careful temperature calibrations can improve the performance even when the temperature is left floating. These corrections were not applied here. Once the temperature stabilizes, the resolution improves during the conditioning period until it reaches about 1.4% overall and 1.2% from the best pixel.

Figure 5.8a shows the overall energy spectrum from this period after conditioning at 30°C. Figures 5.8b and 5.8c following it show the spectrum from the best pixel, which achieved 1.2% FWHM at 662 keV. The $k\alpha$ and $k\beta$ escape peaks from Pb can be individually resolved, showing that the energy resolution is indeed significantly less



- (a) Leakage current vs. time for YH-UM3 as it was heated from room temperature up to 30°C . The leakage current increases steadily with the temperature until thermal equilibrium is reached. Afterwards, there is a brief period of conditioning characterized by unstable leakage current, until it settles down at hour 4.
- (b) Energy resolution (as % FWHM of the 662 keV photopeak) for each of the three working pixels and overall as a function of time. The resolution is worse when the temperature and leakage current is changing, then begins to improve during the conditioning period at 30°C , and stabilizes at a low level. The best resolution achieved is 1.2% on pixel 7.


Figure 5.7

than the ~ 10 keV separation of these peaks.


5.3.3 Large Pixelated Detectors

The encouraging results with medium-thickness detectors led to the growth and characterization of a series of thicker crystals with pixelated contacts. The first group of these was delivered and tested in early 2021. The pixels were arranged in a 3×3 pattern with a pitch of 1.2 mm, surrounded by a guard ring. The far planar contact was again Bi. Crystal cross-sections ranged from 8×8 to 14×14 mm, with three thicknesses of 9.8 mm (YH-UM4), 13.1 mm (YH-UM5) and 13.2 mm (YH-UM6). They were encapsulated with parylene to guard against moisture intrusion.

Of the three, only YH-UM4 was stable under applied bias. Spectra were recorded from all nine pixels, but two suffered from fabrication or connection problems. The best two pixels reached 1.8% and 1.9% FWHM after depth correction, and the overall resolution



- (a) Overall Cs spectrum from YH-UM3 after the brief conditioning phase, while at 30°C. The energy resolution achieved is 1.4% at 662 keV.
- (b) Depth-corrected spectrum from the best pixel at 30°C. The resolution is 1.2% at 662 keV.



- (c) High energy side of the spectrum of Figure 5.8b. The $k\alpha$ and $k\beta$ escape peaks of Pb can be clearly separated, confirming that the resolution is better than 10 keV.

Figure 5.8

was 3%. Since it was suspected that the instability of the other crystals was due to problems during contact fabrication, YH-UM6 was sent back to the crystal growers at Northwestern University. There, the contacts were shaved off (reducing the thickness from 13.2 mm to ~ 12 mm) and new contacts with the same geometry were deposited under carefully controlled conditions. This refabricated crystal is shown in Figure 5.9.




Figure 5.9: Refabricated detector YH-UM6R after it was shaved down to dimensions of 8.5 x 8.5 x 12 mm and new contacts deposited. The detector is encapsulated in parylene to prevent moisture intrusion.

The newly refabricated crystal, named YH-UM6R, was returned to UM in October 2021 and tested again. This time, it was stable up to 750V bias. However, good results were achieved at only 200V bias. The pixel resolutions ranged from 1.0% at 662 keV to 1.9%, except for pixel 9 which had a degraded response due to an error with the experimental apparatus. The overall resolution after depth correction was 1.4% at 662 keV from a 12 mm thick crystal at 200V bias. Spectra from all nine pixels after depth correction are shown in Figure 5.10

The modified two-bias can be used to estimate the hole $\mu\tau$ product. Given measurements of photopeak positions at two different biases, the carrier $\mu\tau$ can be calculated as:

$$\mu\tau_h = k \cdot \frac{D^2}{\ln \frac{N_1}{N_2}} \left(\frac{1}{V_2} - \frac{1}{V_1} \right) \quad (5.1)$$

where D is the thickness of the crystal, N_1 and N_2 are the photopeak positions at the

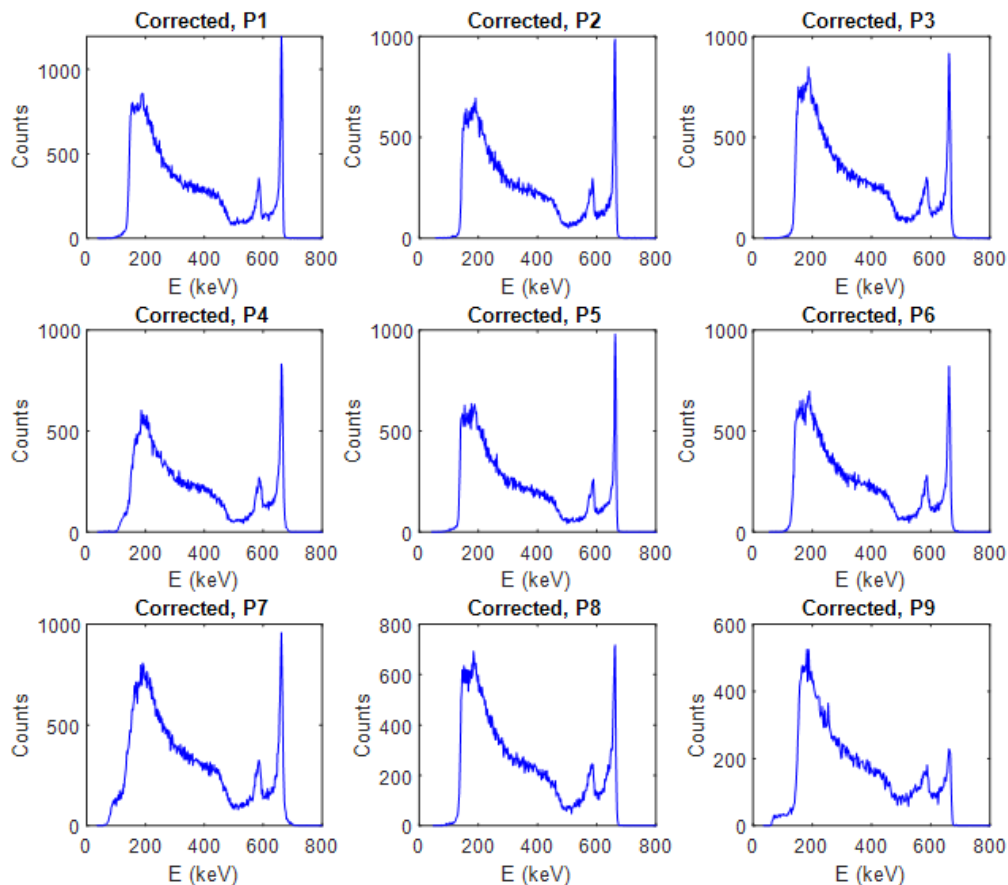


Figure 5.10: Depth-corrected spectrum from all 9 pixels of YH-UM6R. All resolutions are below 2% FWHM, except for channel 9 which is degraded due to a DAQ system problem (not the crystal itself). The best resolution is from channel 5 (the central one) at 1.0%. Overall resolution is 1.4% from all pixels combined.

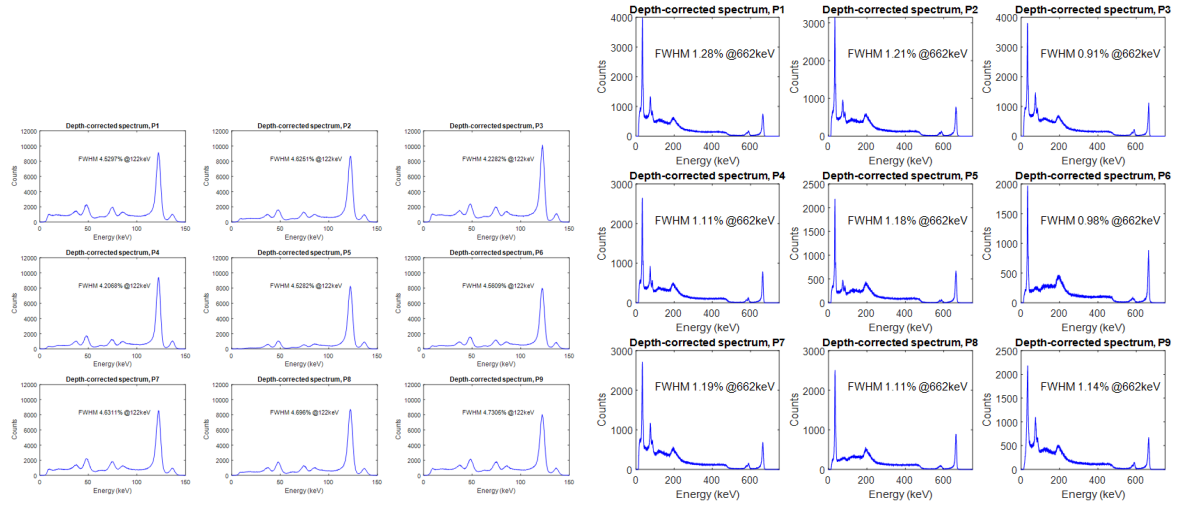
two different biases for planar-side events, V_1 and V_2 are the two different biases used, and k is a correction factor that depends on the pixel pitch to thickness ratio [73]. For this detector, $k = 0.95$. The hole $\mu\tau$ calculated by this method was $5.85 \times 10^{-2} \frac{\text{cm}^2}{\text{V}}$, as good as high quality CZT crystals. Meanwhile, the hole mobility can be found from time-of-flight measurements for planar side events. This yielded a value of $\mu_h = 65 \frac{\text{cm}^2}{\text{V}\cdot\text{s}}$, which together with the $\mu\tau$ measurement gives a hole lifetime of $\tau_h = 900 \mu\text{s}$, much greater than for CZT or TlBr.

Following the success of refabrication of the contacts in greatly improving the performance of YH-UM6R, sample YH-UM3 of Figures 5.6a was also sent back for refabrication of the contacts. Unlike YH-UM6 which did not work at all until after refabrication, YH-UM3 gave good results of about 1.4% FWHM at 662 keV before slowly degrading over time. The heat treatment recovered the performance temporarily, but eventually the crystal entered a breakdown state whenever bias was applied. At Northwestern University, the contacts were shaved off and the thickness reduced to about 4.5 mm. This refabricated sample became YH-UM-3-VK. The same 3x3 pixel pattern as YH-UM6 was used. Co-57 and Cs-137 spectra from the refabricated sample are shown in Figures 5.11a and 5.11b. The overall spectra from all nine pixels combined are in Figures 5.11c and 5.11d. These represent the world record for the best resolutions obtained from perovskite gamma-ray detectors at these energies.

Future efforts will focus on obtaining larger volume crystals and investigating the eventual failure mechanism of these detectors. A 1-inch diameter ingot was cut down into a 12.1 x 11.8 x 10.4 mm detector which was also fabricated with the same 3x3 pixel pattern. This detector, IP21023-2-1(6), is shown in Figure 5.12.

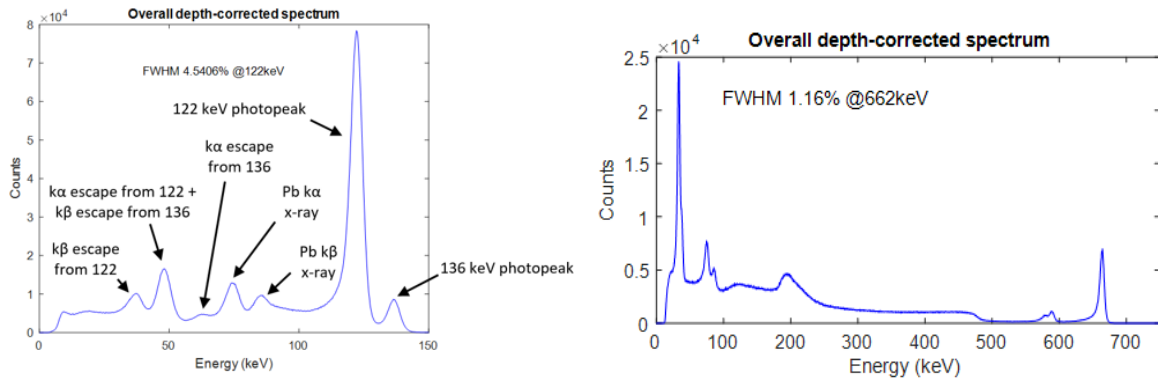
Similar to all the previous pixelated detectors, this sample displayed a high leakage current and unstable or noisy waveforms when the bias was immediately applied. Unlike other detectors, the initial conditioning phase of this sample was far longer. This may partially have been due to the increased volume, and partially due to slightly different chemical makeup. Further discussion of this, and other conditioning effects, is contained in section 5.4.

After the thick detector stabilized, there was significant variance in the performance among all pixels. Depth-corrected spectra are shown in Figures 5.13a and 5.13b. Energy resolutions at 662 keV range from 1.5% for two pixels up to 2.26%, with the overall resolution being 1.9%. The wide guard ring meant that many radiation interactions did not deposit all their energy solely above the pixels. In practice, such a large guard ring is not needed, and the pixelation pattern can extend to cover the majority of the detector face. To obtain the spectral shape without these events (approximating the response from a detector fully covered by pixels) the anode-cathode ratio was used to remove events with incomplete charge collection. This suppressed the characteristic Pb peaks at 70-80 keV, but had minimal impact on the full energy peak.



(a) Depth-corrected ^{57}Co spectra from all 3x3 pixels of YH-UM-3-VK. The resolutions for the 122 keV peak range from 5.1 to 5.7 keV FWHM.

(b) Depth-corrected ^{137}Cs 3x3 spectra from YH-UM-3-VK. The best pixels achieved 0.9% FWHM, while all pixels are at or below 1.2%. Pixels 6 and 8 were corrected using the anode/cathode ratio rather than drift time because of the triggering setup on the DAQ system.



(c) Overall depth-corrected ^{57}Co spectrum from YH-UM-3-VK. The FWHM of the 122 keV peak is 5.5 keV.

(d) Overall depth-corrected ^{137}Cs spectrum from YH-UM-3-VK. The FWHM of the 662 keV peak is 7.67 keV, or 1.16%.

Figure 5.11

5.3.4 Hybrid Perovskite Radiation Detectors

Parallel to the all-inorganic furnace grown samples from Northwestern University, the University of Michigan has also been testing solution grown samples from UNC. These samples mix inorganic anions (PbBr_3) with organic and inorganic cations, Cs mixed with




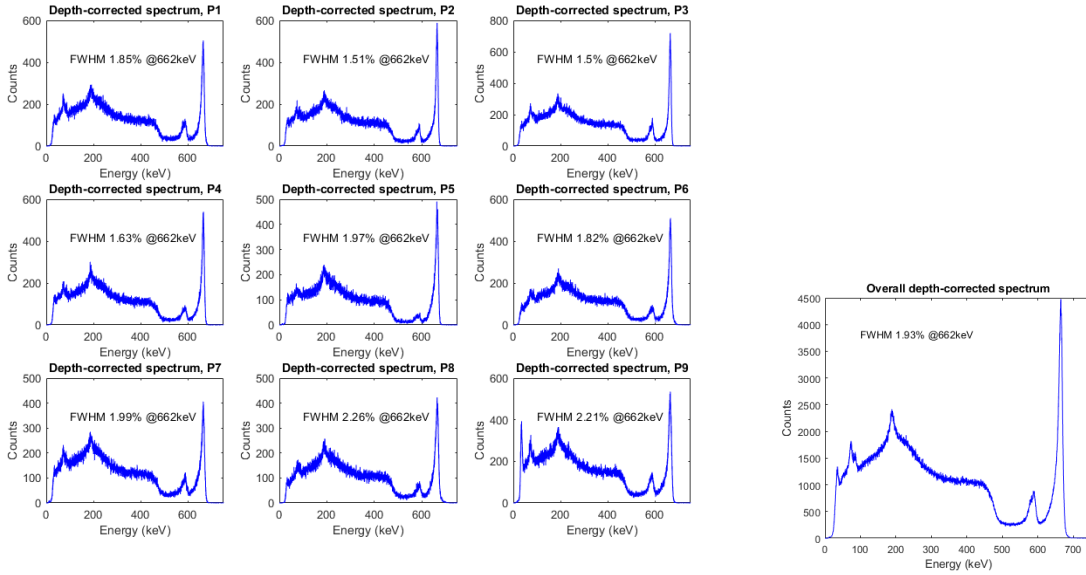
Figure 5.12: Detector IP21023-2-1(6), 12.1 x 11.8 x 10.4 mm.

or replaced by either formamidinium (FA) or methylammonium (MA), to form hybrid perovskites. The growth process is simpler and lower temperature than furnace methods, and promises even lower costs for fabrication. The inclusion of organic elements opens up the possibility of high neutron sensitivity. However, solution grown materials have not yet been able to achieve the same stability and resolution as furnace grown ones.

Solution grown samples have been tested at Michigan since 2017. Initial devices showed some response to radiation, as in Figure 5.14a-5.14b. However, no recognizable features could be identified in the spectra. Devices were generally unstable, with initially transparent bulks clouding over within days and electrodes visibly degrading. Mild electrical shorts between contacts also complicated the response functions.

In 2022, a new generation of solution grown hybrid perovskites based on $\text{FA}_{0.9}\text{Cs}_{0.1}\text{PbBr}_3$ were delivered from UNC. These had much cleaner pixels made of Bi, with Au planar contacts (Figure 5.15). At UNC, they demonstrated good spectra from ^{57}Co , and measured resistivities on the order of $10^{10}\Omega - \text{cm}$, mobility-lifetime products of $2.9 \times 10^{-3} \frac{\text{cm}^2}{\text{V}}$ for holes and $1.8 \times 10^{-3} \frac{\text{cm}^2}{\text{V}}$ for electrons, and mobilities of $\mu_h = 262 \frac{\text{cm}^2}{\text{Vs}}$ and $\mu_e = 189 \frac{\text{cm}^2}{\text{Vs}}$. These are sufficient for radiation detector applications.

The bulk resistivity of these perovskites tends to be relatively low. Similarly to the furnace grown devices, Schottky contacts are used to reduce the leakage current to





(a) Per-pixel depth corrected ^{137}Cs spectra from detector IP21023-2-1(6) after the response had largely stabilized. Resolutions vary from 1.5% FWHM at 662 keV for pixels 2 and 3, to 2.26% for pixel 8. Pixels 6 and 8 were corrected using the anode-cathode ratio because of the triggering setup on the DAQ system; other pixels used the drift time.

(b) Overall ^{137}Cs spectrum after depth-corrections. Overall energy resolution at 662 keV is 1.9%.

Figure 5.13

manageable levels. An I-V scan of these solution grown samples shows the asymmetric nature of the contacts (Figure 5.16). Bias was applied to the planar Au contact, and the leakage current was far lower in the reverse bias regime.

Though hole charge transport parameters in these samples are slightly better than those for electrons, the energy band structure results in a Schottky barrier when the sample is biased negatively on the planar Au contact. This causes the device to operate in electron collection mode. Since the transport properties are similar, however, then if there is good electron collection there will also be good hole collection on the opposite contact. The pixel pitch is about 1mm, while the thickness is 3mm, so these devices are closer to an array of individual planar detectors rather than a true pixelated detector making use of the small-pixel effect. Raw ^{137}Cs spectra from some pixels showed the suggestion of a photopeak, though with the presence of significant trapping (Figure 5.17). Only pixels 8 and 9 provided clean enough signatures of photopeak vs. drift time



(a) Early solution grown sample from UNC with hemispherical contact geometry. Au pixel and guard ring are shown here. (b) Recorded pixel, planar, and guard ring waveforms from a radiation interaction in this early sample. Occasional responses to radiation were observed, but no spectrum with recognizable features could be generated.

Figure 5.14




Figure 5.15: $\text{FA}_{0.9}\text{Cs}_{0.1}\text{PbBr}_3$ solution grown pixelated detector from UNC. These samples had Bi pixels and Au planar contacts, and were biased negatively on the planar contact (electron collection mode).

to perform drift time corrections. Corrected spectra are shown in Figure 5.18.




Figure 5.16: I-V scan of the solution grown perovskite samples of Figure 5.15. (Left) entire range, (right) reverse bias region. Bias was applied to the planar Au contact. The sweep tracked from 0V to +50V, back down to 0V and -50V, then returned to 0V (blue to purple to pink in the figure).

5.4 Conditioning in Perovskite Radiation Detectors

5.4.1 Conditioning of Furnace Grown Samples

Conditioning refers to a change in the response of a detector as a function of time. It may depend on crystal impurities, temperature, bias applied, contact structure, and other factors. Since a stable crystal response is required to perform and apply a calibration, this instability represents a challenge for the successful operation of radiation detectors. CZT detectors are relatively stable, with crystals operating successfully for many years without major changes in performance. HgI_2 and TlBr present a different story. Upon first application of bias, the properties of these detectors often shift, with changes in leakage current and trapping being most noticeable. Various explanations for these phenomena have been proposed, from redistribution of mobile ions in the bulk to chemical changes in the contact interfaces. Conditioning behavior in TlBr has been observed to vary widely among different samples, with some requiring almost no conditioning and others needing long periods under bias before good performance can be obtained.

Perovskites, as ionic materials, may experience similar effects. Understanding what happens during conditioning, and why the detector response degrades, is critical to correcting these effects and achieving the best stability and resolution.




Figure 5.17: Raw ^{137}Cs spectra from the 3x3 solution grown sample UNC-1. The Y scale is logarithmic to highlight the small photopeak and escape peak features in some pixels.

A coarse indicator of what is going on inside the crystal is the leakage current. Among all the pixelated furnace grown samples, the same qualitative behavior is observed when bias is first applied. The leakage current initially increases very quickly and erratically, and the planar contact signal is very unstable. In some samples, the pixel signals are also unstable. After some time, the leakage current stabilizes and begins decreasing exponentially, and the waveforms stabilize. This has been observed in all pixelated CsPbBr_3 samples from Northwestern thus far. The variations among samples are in the degree of initial waveform instability, the time constant for leakage current relaxation, and how much the planar contact stabilizes as the leakage current decreases. This behavior is therefore likely due to fundamental material aspects of the perovskite itself rather than impurities introduced in particular batches.




Figure 5.18: Drift time corrected ^{137}Cs spectra from pixels 8 and 9 of the solution grown sample UNC-1. Resolutions are 4.8% and 5.3% at 662 keV.

Figure 5.19 shows this effect. The bias histories shown are for YH-UM-3VK and IP21023-2-1(6). These two crystals have very different individual histories. YH-UM-3VK was originally grown by Yihui He in 2019, and was refabricated several times by different people. IP21023-2-1(6) was grown by Vladislav Klepov in 2022, has different concentrations of impurities, and has a much larger volume and spectroscopic performance. Despite these differences, the IV histories have similar features. Immediately upon applying bias, the leakage current increases rapidly with time and shows high variations. After a few days, the leakage current quickly spikes and then settles into a gradual, exponential decay. During this period of decaying leakage current the waveform stability increases and spectroscopic performance is maximized. Afterwards, each time the bias is raised the leakage current jumps and then exponentially settles.

More information about the conditioning phase can be gleaned by examining the detector responses in detail. Figure 5.20a shows the amplitude of signals on pixel 3 from IP21023-2-1(6) when the bias on the crystal was raised from 0V to 500V. In the first couple days, a clear, photopeak is visible at about 5000 ADC, though it is very broad. A broad photopeak would usually be due to trapping. However, Figure 5.21 shows the photopeak amplitude doesn't decrease significantly with depth during this part of the measurement, meaning trapping cannot be the cause of the poor resolution. Additionally, from Figure 5.21, the signal velocity changes: initially anode-side events have short drift times, then quickly slow down over the course of the first two days. The drift time is a measure of the internal electric field. The changes in it suggest the





Figure 5.19: Bias and leakage current histories of YH-UM-3-VK and IP21023-2-1(6). These two furnace grown CsPbBr_3 samples from Northwestern have very different material and fabrication histories, but show similar leakage current vs. time characteristics.

internal field is reconfiguring, possible due to some redistribution of mobile charges or a buildup of generated space charge.



(a) Pixel 3 signal amplitude vs. measurement time for crystal IP21023-2-1(6) when biased from 0 to 500V. (b) Pixel 7 signal drift time vs. measurement time for crystal IP21023-2-1(6) when biased from 0 to 500V.

Figure 5.20

Examination of the signal waveforms reveals that extra noise is the cause of the poor resolution. Figure 5.22 shows pixel waveforms immediately when bias is applied, and when it has been held on for several hours. The pixel signals are very noisy immediately when bias is applied, while the leakage current and drift times are increasing. As the bias is held on, the leakage current, drift time, and waveforms all stabilize at the same time, and the resolution after depth correction improves. Paradoxically, this is accompanied

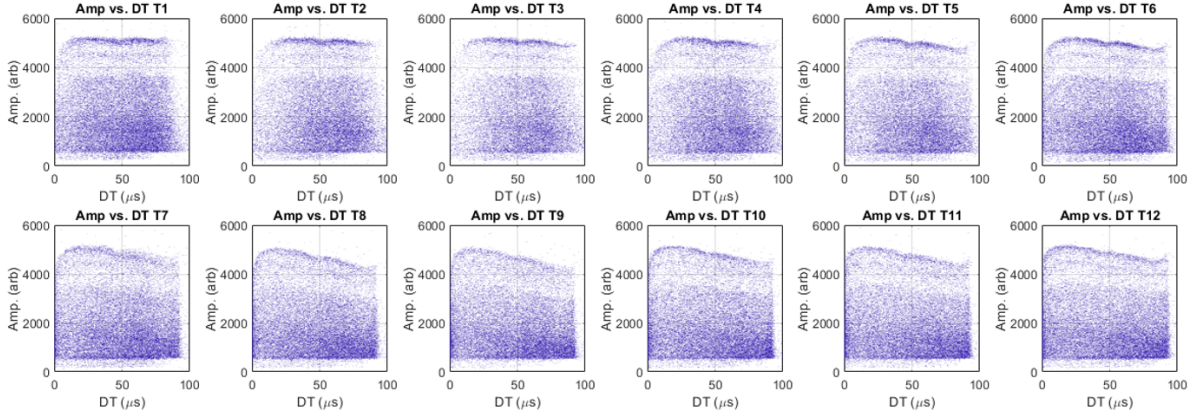


Figure 5.21: Pixel 3 signal amplitude vs. drift time over several time slices during the initial 500V bias period for IP21023-2-1(6).

by an increase in trapping, as in the latter parts of Figure 5.21.




Figure 5.22: Waveforms from IP21023-2-1(6) immediately as 500V bias is applied (blue) and after bias has been held for several hours (red). The stabilization of the baseline slope is evident.

When the bias is raised to 1000V, some similar effects occur. Figure 5.23 shows pixel 3 signal amplitude vs. measurement time. Similarly to the latter part of Figure 5.20a, there is a period of increased trapping, after which the detector recovers and the resolution improves. Figure 5.24 shows the curve of amplitude vs. drift time during

this time. Initially, there is some high trapping similar to the latter part of Figure 5.21. At a certain moment, however, the crystal recovers. This recovery is not marked by a significant change in leakage current, which steadily decreases through the entire time at 1000V. A study of drift times shows that maximum drift times (a measure of electric field uniformity) vary among pixels initially, but eventually converge to a common value.




Figure 5.23: IP21023-2-1(6) pixel 3 signal amplitude vs. measurement time once the bias is raised to 1000V. There is an initial period of reconfiguring, then the response stabilizes again.

Taken together, a hypothesis may be drawn about the nature of the conditioning in IP21023-2-1(6) and similar samples. The electric field redistribution immediately when bias is applied is likely due to some mobile charges inside the bulk. These drift and eventually accumulate against the anode and cathodes, weakening the internal field and increasing the drift times. They also contribute to the leakage (and may potentially modify the Schottky barrier), causing the leakage current to increase. As the charges pile up, they weaken the internal field and slow down any remaining charges, causing the leakage current to decrease exponentially. Eventually, they reach a steady state. The empty sites they left behind may act as trap centers, leading to the increased trapping over time observed in Figure 5.21. When the bias is raised again, the stronger field is able to sweep more charges out of the bulk. Eventually, the trap sites fill and the trapping decreases. On the pixel side, the space charge initially accumulates unevenly,

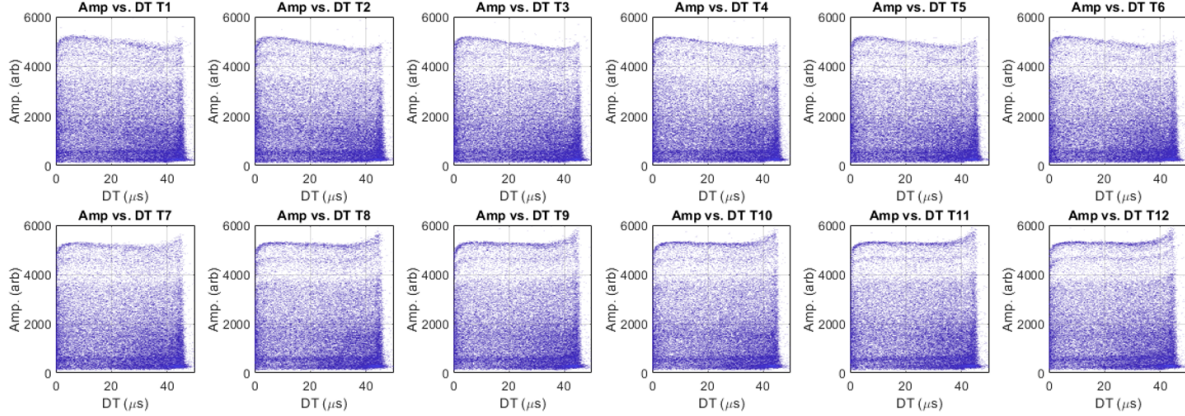



Figure 5.24: Pixel 3 signal amplitude vs. drift time at different time slices after the bias is raised to 1000V. The trapping is significant for a while, then the crystal recovers.

possibly because it was initially distributed nonuniformly in the bulk. Over time, the charges diffuse laterally, causing the drift times to become more uniform among the pixels.

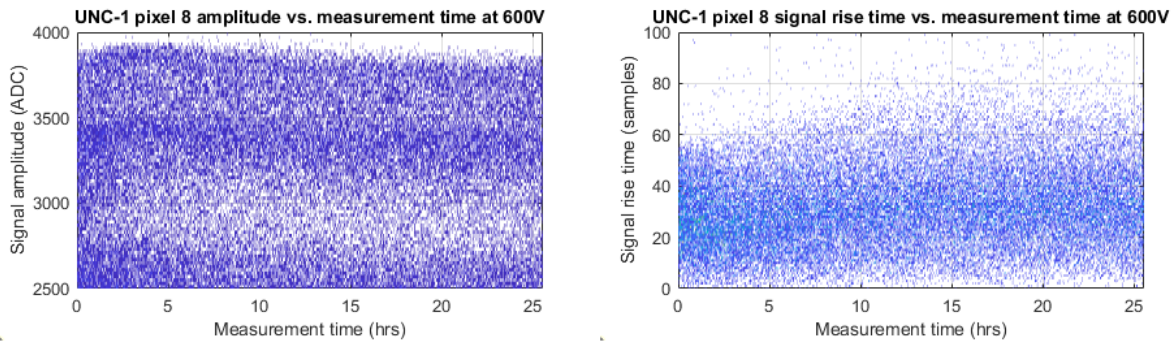
5.4.2 Conditioning of Solution Grown Samples

Similar conditioning effects are seen in solution grown samples. In the case of UNC-1 of Figure 5.15, the device changed rapidly when bias was first applied. With -300V bias on the planar contact, the signal amplitude under ^{137}Cs irradiation first increased rapidly, stabilized, then started to decrease slightly. At the same time, the rise time of pixel signal waveforms increased almost linearly. These trends are shown in Figure 5.25a. Once the response was seen to stabilize somewhat, the bias was increased to -600V. Immediately upon changing the bias, there was a small shift in the amplitude and rise time, but it settled out relatively soon compared to the initial conditioning at -300V (Figure 5.25b). This is similar to the behavior of the furnace grown samples. Once they are conditioned, the response is largely stable except for some slow polarization over the course of many hours. This is mirrored in the leakage current over time (Figure 5.26).

The increase in leakage current suggests a breakdown of the blocking contacts that had limited it. Indeed, the pixel contacts became discolored after several days at bias (Figure 5.27). A new IV sweep showed that the blocking behavior has vanished, and the leakage current in the forward bias regime is much lower than in Figure 5.16. If



(a) UNC-1 pixel 8 signal amplitude vs. measurement time (right) and signal rise time vs. measurement time (left) at -300V bias, the first time bias was applied to the sample. The steady changes are the effects of conditioning.



(b) UNC-1 pixel 8 signal characteristics over time at -600V bias. The sample has conditioned and mostly stabilized.

Figure 5.25

the leakage current is low under forward bias, the detector could be operated in hole collection mode. When the crystal was forward biased, radiation pulses were indeed observed. However, the response was only stable for a few minutes under forward bias before the leakage current increased rapidly back to pre-conditioning levels. A possible explanation is an accumulation of positive space charge under reverse bias due to mobile ions shifting the energy levels and modifying the Schottky barrier slowly over time. Under forward bias, the ions may have been driven back into the bulk and the Schottky barrier restored, though the pixel degradation remained. Such bias switching has been known to temporarily restore the performance of polarized TlBr detectors. It may be that the mobile ions react with the contacts, degrading them irreparably. In the future, some interfacial layer may possibly be used to protect the contacts. Alternatively, if the ions are drifting impurities then higher quality precursors may slow the degradation




Figure 5.26: Bias history of sample UNC-1. The conditioning phases can be seen as a jump in leakage current immediately after the bias is raised, which exponentially decays as the crystal relaxes. After the crystal is conditioned at -600V, the leakage current begins climbing at about 3.5-4 nA/day; this is separate from the initial conditioning phase and represents polarization of the crystal.

process.

5.5 Electron-Hole Conversion Effects and PCA

5.5.1 Anomalous Gain and Electron-Hole Conversion

Plots of pixel signal amplitude versus depth can be used to derive depth correction parameters, but also may reveal a wealth of information about the detector properties. A typical amplitude vs. depth plot for a CZT crystal is shown in Figure 5.28a. At shallow depths close to the pixel, the signal decreases due to the weighting potential deficit. Further into the bulk, more of the signal is induced, and the amplitude reaches a peak somewhere around a quarter of the way to the cathode. Further into the bulk, electron trapping may cause the signal to lose amplitude slightly with increasing depth. This also causes the photopeak trace to widen slightly as the resolution degrades. Fitting this





Figure 5.27: IV sweep of sample UNC-1 after several days under bias. The blocking behavior of Figure 5.16 has vanished. Inset: the pixels of UNC-1 upon arrival at UM (left) and after several days at bias (right).

signal loss to the expected exponential can yield the electron mobility-lifetime product with just a single measurement, rather than the two-bias method or multi-bias Hecht fitting otherwise used. However, when trapping is low this slope is close to zero, giving large uncertainties in the $\mu\tau$ product.



(a) Typical plot of pixel signal amplitude vs. depth for a CZT detector exposed to ^{137}Cs . (b) Plot of pixel signal amplitude vs. drift time for YH-UM6R, the furnace grown CsPbBr_3 detector from Figure 5.9, exposed to ^{137}Cs .

Figure 5.28

Figure 5.28b shows the same type of plot for crystal YH-UM6R, the furnace grown sample shown in Figure 5.9. This detector was biased to 300V and exposed to ^{137}Cs . Similar to the CZT detector, there is a deficit near to the pixel, though it may not be as severe. The signal rises to a maximum and then falls, indicating hole trapping. However, an unusual behavior is observed near to the far planar contact side: the signal amplitude increases again, and the photopeak trace becomes much broader. This has not been observed in CZT.

To investigate this, a series of measurements were carried out with YH-UM6R. Figure 5.29 shows the amplitude vs. drift time plots at different biases. At each bias, there is some extra signal gain for near-anode events. As the bias increases, the region where this gain is present extends deeper and deeper into the bulk. At high enough biases, the near-pixel signal deficit from the weighting potential disappears entirely.

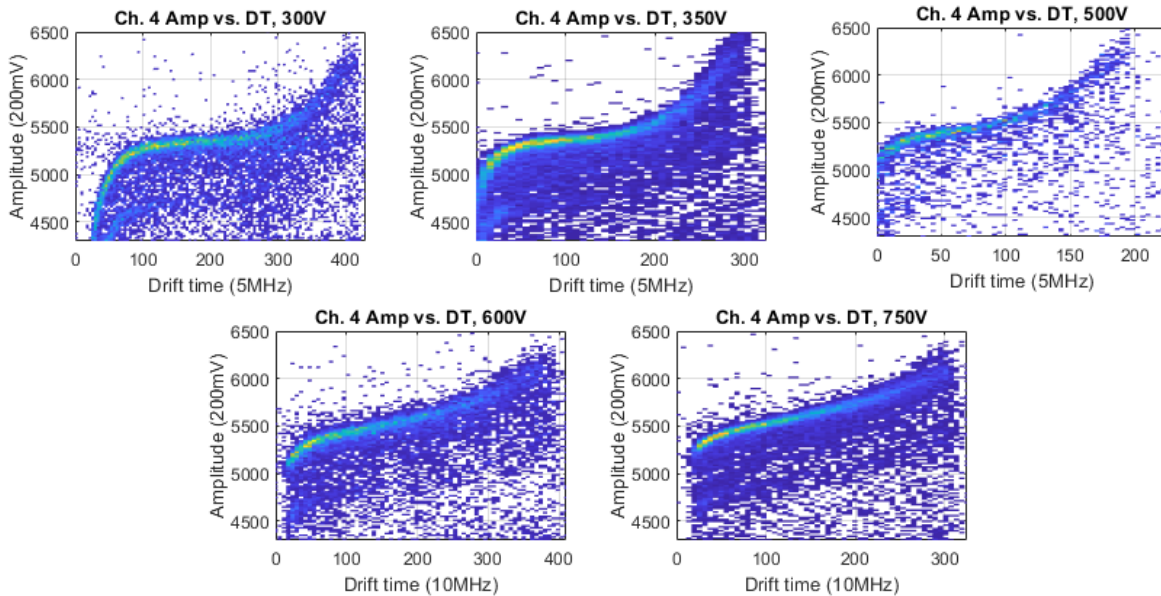


Figure 5.29: Amplitude vs. drift time distributions at different biases from detector YH-UM6R.

Individual and average photopeak waveforms from anode-side, cathode-side and bulk events were examined. Figure 5.30 shows average photopeak waveforms from two regions, one closer to the anode side and one deeper in the bulk. In both, an additional step is seen after the main charge collection. This is the source of the additional signal gain. Two properties are evident: closer to the anode side, the amplitude of the step

increases and the delay of the extra signal grows shorter.




Figure 5.30: Average photopeak waveforms from the bulk and anode side of YH-UM6R under 750V bias. An additional step is clearly visible after the main cloud is collected. The amplitude of the step is delineated by the red dashed lines, while the step delay is indicated in blue.

A discrete step in the waveform can only be produced by a discrete cloud of charge arriving at the pixel. Because of the small pixel effect, the detector is sensitive primarily to the drift of holes, so the extra step must be due to an extra cloud of holes arriving after the main cloud is collected. It is too systematic to be random pile-up. It cannot be some discrete cloud of holes from the main charge cloud which were delayed by some bulk traps, since the prompt signal has the same amplitude regardless of the magnitude of the extra step cloud.

This delayed signal can be interpreted as resulting from the conversion of drifting electrons to holes at the far anode surface. Anode side events would present the shortest distance for photoelectrons to drift before being converted, resulting in the shortest signal delay. They would also have the least amount of electron trapping, resulting in the largest delayed signal amplitude. If the electron detrapping time is sufficiently long, the cloud of electrons that arrives at the anode will be relatively concentrated, resulting in an equally concentrated cloud of holes that is released to drift back to the pixel cathodes. Electron drift can also explain the recovery of the near-pixel region at

high bias, as the electrons drift away from the pixel and recover the signal originally lost to weighting potential deficits.

There are several possibilities for the physical mechanism behind this conversion process. A very high field region close to the anode might result in multiplication and the creation of additional electron-hole pairs. Auger recombination has been known to result in signal gain in HgI_2 [74] and TlBr [75] detectors by the release of electrons from trap levels. Lastly, several Schottky barrier related effects may result in signal gain [76]. To distinguish between these effects, the trend of the signal gain with bias voltage was studied. Multiplication and Schottky barrier effects would both be expected to increase gain with bias. Auger recombination relies on the trapping of minority carriers in deep traps that serve as recombination centers, and so should result in lower gain with increased bias.

The gain is calculated as the ratio of initial to delayed signal amplitude using near-anode-side events. Using near-anode events ensures the impact of electron trapping in the bulk is minimal, and the resulting hole clouds undergo similar levels of trapping which cancel out in division. Figure 5.31 shows the result. The gain decreases with increasing bias in an exponential fashion, consistent with Auger recombination in a narrow layer very near the anode contact.




Figure 5.31: Anode-side e^-h^+ conversion gain as a function of bias.

Assuming the delayed signal is an “echo” from electrons releasing holes from Auger

centers very near the anode, it can be exploited to determine the electron drift properties. The time delay of the echo signal should be the summation of the electron drift time from the location of gamma-ray interaction (a distance d from the cathode surface) to the anode surface and the hole drift time from the anode surface, where the electron-to-hole conversion happens, to the original gamma-ray interaction position (d), thus we have

$$\Delta T = \frac{D-d}{\mu_e E} + \frac{D-d}{\mu_h E}, \quad (5.2)$$

This results in a linear relationship between the echo signal delay ΔT versus photohole drift time $T_h = \frac{d}{\mu_h E}$, with a slope of $-(1 + \frac{\mu_h}{\mu_e})$ and intercepts depending on the electron and hole mobilities:

$$\begin{aligned} \Delta T &= \frac{D}{\mu_h E} + \frac{D}{\mu_e E} - \frac{d}{\mu_e E} - \frac{\mu_h}{\mu_e} \frac{d}{\mu_e E} \\ &= \frac{D}{\mu_h E} + \frac{D}{\mu_e E} - (1 + \frac{\mu_h}{\mu_e}) T_h, \end{aligned} \quad (5.3)$$

where ΔT is the echo signal delay, T_h is the initial photohole drift time, μ_e and μ_h are the electron and hole mobilities, d is the location of the initial gamma-ray interaction, and D is the thickness of the detector (here, 12 mm).

Figure 5.32 shows the plot of echo signal delay versus photohole drift time. The linear fit is a good match to the data. The calculated electron mobility is $61.6 \text{ cm}^2/\text{V} \cdot \text{s}$, similar to the hole mobility. This electron mobility agrees with that measured in similar samples of CsPbBr₃ by time of flight measurements, $63 \pm 5 \text{ cm}^2/\text{V} \cdot \text{s}$, published previously [77].

With electron drift times from the echo signal delay, electron lifetime can be calculated from the echo signal amplitude vs. electron drift time as

$$N(t) = N_0 e^{-\frac{t}{\tau_e}}, \quad (5.4)$$

where $N(t)$ is the echo signal amplitude at a given electron drift time, N_0 is the maximal electron signal amplitude for anode side events, t is the electron drift time, and τ_e is the electron lifetime. Figure 5.33 shows the fit of the data to the exponential of equation 5.4. The resulting electron lifetime is calculated to be $23.4 \text{ } \mu\text{s}$, for an electron $\mu\tau$ of $1.4 \times 10^{-3} \text{ cm}^2/\text{V}$, approximately 40 times lower than the hole lifetime value in this sample.




Figure 5.32: Echo signal delay versus photohole drift time at 750V bias in YH-UM6R. Linear fit line is also shown.




Figure 5.33: Photopeak event echo signal amplitude vs. electron drift time calculated from echo signal delay. Exponential fit is also shown.

5.5.2 Furnace Grown Samples - Corrections with PCA

The extra signal amplitude from electrons converting to holes in principle provides a second opportunity to measure the deposited energy. This opens the possibility for achieving energy resolution better than the limit imposed by the Fano factor. However, since electron trapping is more severe than hole trapping, the echo signal is more stochastic than the primary hole signal. This causes the resolution after depth correction to be degraded by the conversion effect. A more sophisticated method of extracting the information inherent in the waveform shape is required to make the best use of the extra electron signal.

One possibility is system response function fitting. Since the echo signal is a systematic effect that varies as a function of position, it is possible to collect average system response waveforms from different voxels in the detector and construct a matched filter to best estimate the signal amplitude. This can be computationally costly, however.

An alternative method is principal component analysis (PCA). Briefly, PCA decomposes a set of waveforms into orthogonal eigenvectors. Only a few eigenvectors are needed to account for the majority of the variance in the dataset. Projecting the waveforms along these eigenvectors yields a set of scalars, called principal components, which encode some information about the waveform shape. Regressing a known calibration energy, like the 662 keV line from ^{137}Cs , against the principal components sequentially can result in a calibration that takes into account latent variables not addressed by standard depth corrections.

Principal component analysis was applied to the data collected from YH-UM6R at 750V bias. The drift times were divided into six coarse depth bins, and singular value decomposition yielded eigenvectors for each depth bin. Figure 5.34 shows the overall spectra after simple depth correction and after PCA corrections. After depth correction, the resolution is relatively poor, 2.4% FWHM at 662 keV. This is largely due to the extra broadening of the electron signal. PCA can more appropriately weight the information in this region of the waveform, as well as perform intrinsic corrections for radial subpixel distributions and other latent variables. This improves the overall resolution to 1.4%, with the best pixel reaching down to 1.2%.




Figure 5.34: ^{137}Cs spectra from YH-UM6R under 750V bias after depth correction by drift time (top) and PCA corrections (middle). The best pixel (bottom) is also shown. Resolutions achieved are 2.4% after depth correction, 1.4% after PCA corrections, and 1.2% for the best pixel.

5.5.3 Solution Grown Samples - Corrections with PCA

The responses from the solution grown samples of Figure 5.15 may also be improved with PCA. Only the pixel waveform is used here. PCA can be performed directly on the pixel waveforms, but different preprocessing operations may condition the problem better for corrections. The goal of preprocessing is to maximize the dynamic range of the resulting principal component values and achieve the cleanest separation between different spectral features. A search through different hyperparameters identified the 5th root as a good choice of preprocessing function. Like with standard depth correction, the resulting relationship is nonlinear, so the sequential linear fitting employed by B. Williams is not used. Rather, the principal component space is divided into a number of bins, and the spectra from each bin are aligned to a photopeak energy of 662 keV. Only the second principal component is used to avoid overfitting. The result is shown in Figure 5.35.




Figure 5.35: PCA-corrected spectra from pixels 8 and 9 of UNC-1. Resolutions are 2.8% and 4.8% at 662 keV, an improvement over the standard drift time corrections of Figure 5.18.

Some meaning can be assigned to the principal components by examination of the eigenvectors, individual waveforms, or bulk statistics. In this case, the second principal component contains significant information about drift time (Figure 5.36). Correction against the second principal component is therefore similar to correction against depth of interaction, but the principal component does a better job of separating spectra from different regions of the crystal than the drift time itself.




Figure 5.36: Relationship between drift time and principal component 2 for pixels 8 and 9 of UNC-1.

5.6 Future of Perovskite Radiation Detectors

Future research into perovskite based radiation detectors can take many paths. A primary task is to improve the long-term stability of furnace grown and solution grown detectors. This will involve precise identification of the polarization mechanism. If it is indeed due to mobile impurities, higher quality precursors may extend the lifetime of these crystals. More complex contact engineering may also help increase the longevity of the devices.

Larger pixelation patterns and digital ASIC readout may be employed relatively soon. IP21023-2-1(6) already has sufficient area and volume to be patterned with a grid of 9x9 or even 11x11 pixels. It was cut from a larger, 2.54 cm diameter ingot, with the main limitation being the small sample readout board it needed to be mounted onto. The next generation H3DD-UMv4 ASIC recently developed will be an excellent tool, since it is known to perform better with slower materials than the alternative Norwegian VAD-UMv2.2 ASIC. Larger pixelation patterns will provide considerably more information about material uniformity and conditioning behavior.

The electron-hole conversion effect has been seen in every single furnace grown CsPbBr_3

delivered by Northwestern, including those produced by different crystal growers and with different quality precursors. The gain is still relatively small, on the order of 10-20% for reasonable biases. More work can be done to improve algorithms for best extracting the extra information in the electron signal. On the contact engineering side, if the signal is indeed due to Auger emission from a thin near-anode layer, it may also be possible to boost the gain by modifying the anode interface.

Solution grown perovskites provide particular promise due to their unique method of growth. Scaling up the volume of solution grown perovskites to match furnace grown samples is critical. Additionally, since they are precipitated from solution, it may be possible to embed metal contacts within the crystal as it forms. This may open up the possibility for entirely new weighting potential profiles and reconstruction techniques.

Chapter 6

Concluding Remarks

6.1 Review of Work

Radiation detector technology is a field that exists at the intersection of many different disciplines in engineering and physics. This thesis worked aimed to make contributions to many different aspects of pixelated semiconductor detector technology.

The digital logic associated with front-end control and waveform processing was incorporated into silicon in the form of a digital signal processing ASIC. This first version of the ASIC incorporated all the functions previously performed by an array of ADCs and an FPGA, as well as some of the waveform preprocessing and digital filtering done on a CPU. The chip was designed in the TSMC 65nm silicon process, and testbenches were built to verify its proper operation. A dedicated test system for the chip was developed, capable of supporting high speed 100 MHz signals and clock lines. The first version of the chip was received and debugged using this integrated test system. A power distribution error was identified and resolved with focused ion beam circuit editing services. The FIB-fixed chip responded to communications, and showed responses that matched the expected outputs, with some glitches. Revisions were made to the design to reduce these glitches and second version of the chip was prepared for tapeout.

Charge steering effects were studied in CZT detectors with the aid of fine-width collimator scans. Several algorithms for correcting charge steering effects were explored. An iterative algorithm based on comparison of experimentally observed count density variations with theoretical expectations was presented. A one-dimensional version of this algorithm was applied to the Z-dimension for single pixel events. This is typically

calculated using the cathode-anode ratio, which is assumed to be linear with depth. Collimator scans revealed significant nonlinearities for some pixels. These were effectively corrected by the iterative technique, while pixels with linear relationships were not disturbed. Absolute Z position error improved by $500 \mu\text{m}$.

Machine learning models were developed to predict the subpixel XY position from the observed waveforms. An algorithm for performing supervised learning to predict subpixel positions was developed and applied, first to radial coordinate and then to the overall XY position. The subpixel radial position of interaction was predicted only using the anode waveform, to a precision of about $200 \mu\text{m}$. The true XY subpixel position was predicted by a neural network trained by the algorithm and its performance was validated by slit collimator measurements. An improvement of $50 \mu\text{m}$ in the position resolution over classical techniques was demonstrated.

A new technique was presented to image pair production events in pixelated CZT. This has typically been done by resolving the tracks of the daughter particles. However, this capability is not yet fully demonstrated in 3D position sensitive CZT detectors. A simple algorithm was developed that is able to reconstruct an image of a radioactive source using only pair production events, which does not rely on resolving the particle tracks explicitly. Instead, correlations among particle momenta and charge cloud centroid positions were exploited, along with an asymmetry in the event geometry relating to positron thermalization and annihilation. The algorithm is able to draw some information about the split of kinetic energy between electron and positron per-event, and uses this to reconstruct the radiation image. The technique is verified in simulation and in experiment using a PuBe source emitting 4.44 MeV photons.

Pixelated detectors based on perovskite materials are fabricated and tested, achieving world record performances for furnace- and solution-grown crystals. Large volume detectors based on Bridgeman-grown CsPbBr_3 are developed, and their conditioning behaviors are studied. The best energy resolution at 662 keV yet recorded from perovskite detectors, 0.9% FWHM, was obtained, owing to digital processing techniques. Spectroscopy of ^{137}Cs using solution-grown crystals is also achieved, with a best resolution of 2.8% FWHM at 662 keV. A new phenomenon is discovered and characterized, whereby signals are internally amplified in a consistent manner across many different CsPbBr_3 detectors. This is attributed to drifting photoelectrons being converted to holes upon reaching the anode interface by an Auger process. Principal component analysis is

used to efficiently extract the information in this extra detection signal, improving the energy resolution beyond that achieved with simple 3D position correction techniques.

6.2 Possible Subsequent Work and Visions for the Future

There are many opportunities for improvement of the presented techniques, and for follow-on development of new technologies. The first version of the DSP ASIC presented in chapter 2 aimed to incorporate digital filtering and other simple waveform processing operations into a single piece of silicon. These operations result in amplitude and timing information extracted from each waveform. In practice, these values are then converted into 3D coordinates and energies of interactions via complex calibration procedures. These calibration procedures themselves may be implemented in silicon in a future iteration of the ASIC. It may even be possible to incorporate more complicated algorithms in silicon, such as Compton and coded aperture imaging or radioisotope ID.

Charge steering continues to present a challenge for position reconstruction. The one dimensional iterative algorithm presented was able to demonstrate improvements in depth resolution, but a full 3D version will be needed to perform simultaneous XY calibration. Proper subpixel reconstruction of charge sharing events will be needed for the successful application of this type of algorithm. Alternatively, some other algorithms not relying on count density distributions may be able to correct for steering effects. Two possibilities, relying on the energy-dependent attenuation of different lines or Compton imaging, were suggested and will be the focus of follow-on studies. An automated collimator scan may ultimately prove to be the most rigorous solution.

There remains much room for improvement in event reconstruction using machine learning techniques. The algorithm presented in this work for performing supervised training of a model can be applied to any supervised learning algorithm. The neural networks and Gaussian process regressions presented here are simple supervised learning algorithms, but other machine learning models may make better use of the structure of the data. For example, convolutional neural networks like ResNet might perform better since they may natively exploit the correlation in time between different parts of the waveform. Unsupervised methods are also powerful tools, but challenges in the

interpretation of their outputs will need to be solved to tie them to physical quantities like subpixel position.

The pair production imaging presented would benefit greatly from even a little information about the electron cloud shape. Techniques to reconstruct the electron cloud shape have been presented in the past, but have not been successfully applied to multi-pixel high energy events. Such events are the majority of pair production interactions. The improvement in subpixel XY position resolution may help improve the performance of these techniques purely from a signal-to-noise standpoint. Additionally, more refined algorithms for electron cloud shape estimation may give useful results at energies of several MeV. Even without electron cloud shape estimation, the presented algorithm could benefit from filtered backprojection or maximum likelihood estimation maximization techniques.

Perovskite detectors represent a very promising family of materials, but more work needs to be done on device stabilization. Challenges remain with contact stability and optimal conditioning and operation of the samples. Fabrication of CsPbBr₃-based detectors needs to expand to larger pixel arrays of 6x6, 9x9 or 11x11 pixels. Procedures for bonding the crystals to carrier boards for interfacing with front-end ASICs will need to be developed. The H3DD-UM waveform digitizing ASIC is an optimal choice, but modifications may need to be made to it to achieve the best resolution with materials like perovskites which have long carrier lifetimes but low mobilities. The electron-hole conversion effect may be enhanced by careful design of the thin layer on the anode side. The internal gain could result in improved energy and position resolution, giving a second opportunity to measure the energy and depth of each interaction. Separately from furnace grown samples, solution grown samples also show promise for their low manufacturing cost. Their energy resolution needs to catch up to that of furnace grown samples. However, their unique method of preparation opens the possibility of more exotic detector structures, such as contacts embedded within the crystal bulk.

Bibliography

- [1] G. F. Knoll. *Radiation Detection and Measurement*. en. John Wiley & Sons, Aug. 2010. ISBN: 978-0-470-13148-0.
- [2] G. Lutz et al. *Semiconductor radiation detectors*. Springer, 2007.
- [3] F. Zhang and Z. He. “3D position-sensitive CdZnTe gamma-ray spectrometers: improved performance with new ASICs”. en. In: ed. by A. Burger, R. B. James, and L. A. Franks. Denver, CO, Oct. 2004, p. 135. DOI: [10.1117/12.563903](https://doi.org/10.1117/12.563903).
- [4] Z. He. “Review of the Shockley–Ramo theorem and its application in semiconductor gamma-ray detectors”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 463.1 (May 2001), pp. 250–267. ISSN: 0168-9002. DOI: [10.1016/S0168-9002\(01\)00223-6](https://doi.org/10.1016/S0168-9002(01)00223-6).
- [5] B. Thomas, M. C. Veale, M. D. Wilson, P. Seller, A. Schneider, and K. Iniewski. “Characterisation of Redlen high-flux CdZnTe”. en. In: *Journal of Instrumentation* 12.12 (Dec. 2017). ISSN: 1748-0221. DOI: [10.1088/1748-0221/12/12/C12045](https://doi.org/10.1088/1748-0221/12/12/C12045).
- [6] F. Zhang, C. Herman, Z. He, G. De Geronimo, E. Vernon, and J. Fried. “Characterization of the H3D ASIC Readout System and 6.0 cm³ 3-D Position Sensitive CdZnTe Detectors”. en. In: *IEEE Transactions on Nuclear Science* 59.1 (Feb. 2012), pp. 236–242. ISSN: 0018-9499, 1558-1578. DOI: [10.1109/TNS.2011.2175948](https://doi.org/10.1109/TNS.2011.2175948).
- [7] Y. Zhu and Z. He. “Performance of a 2-keV digitizer ASIC for 3-D position-sensitive pixellated semiconductor detectors”. In: *2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*. ISSN: 1082-3654. Oct. 2012, pp. 4109–4112. DOI: [10.1109/NSSMIC.2012.6551939](https://doi.org/10.1109/NSSMIC.2012.6551939).
- [8] Z. He. “Pixelated semiconductor detectors for gamma-ray spectroscopy”. In: *Hard X-Ray, Gamma-Ray, and Neutron Detector Physics XXIII*. Vol. 11838. SPIE, Aug. 2021, p. 1183802. DOI: [10.1117/12.2596224](https://doi.org/10.1117/12.2596224).
- [9] Y. Zhu. “Digital Signal Processing Methods for Pixelated 3-D Position Sensitive Room-Temperature Semiconductor Detectors”. PhD thesis. University of Michigan, 2012.

- [10] B. Williams, Y. F. Zhu, and Z. He. “Applications of principal component analysis for energy reconstruction in position-sensitive semiconductor detectors”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. Symposium on Radiation Measurements and Applications XVII 954 (Feb. 2020), p. 162047. ISSN: 0168-9002. DOI: [10.1016/j.nima.2019.04.002](https://doi.org/10.1016/j.nima.2019.04.002).
- [11] Y. Zhu, S. E. Anderson, and Z. He. “Sub-Pixel Position Sensing for Pixelated, 3-D Position Sensitive, Wide Band-Gap, Semiconductor, Gamma-Ray Detectors”. In: *IEEE Transactions on Nuclear Science* 58.3 (June 2011). Conference Name: IEEE Transactions on Nuclear Science, pp. 1400–1409. ISSN: 1558-1578. DOI: [10.1109/TNS.2011.2132738](https://doi.org/10.1109/TNS.2011.2132738).
- [12] Y. Zhu and Z. He. “Subpixel Sensing for Charge-Sharing Events in Pixelated CdZnTe Detectors”. In: *IEEE Transactions on Nuclear Science* 69.6 (June 2022). Conference Name: IEEE Transactions on Nuclear Science, pp. 1385–1388. ISSN: 1558-1578. DOI: [10.1109/TNS.2022.3176794](https://doi.org/10.1109/TNS.2022.3176794).
- [13] H. Wolter. *Mirror systems with grazing incidence as image-forming optics for X-rays*. Tech. rep. Jan. 1975, pp. 94–114.
- [14] E. Virgilli, H. Halloin, and G. Skinner. *Laue and Fresnel lenses*. en. arXiv:2208.12362 [astro-ph, physics:physics]. Aug. 2022.
- [15] S. R. Gottesman and E. E. Fenimore. “New family of binary arrays for coded aperture imaging”. EN. In: *Applied Optics* 28.20 (Oct. 1989). Publisher: Optica Publishing Group, pp. 4344–4352. ISSN: 2155-3165. DOI: [10.1364/AO.28.004344](https://doi.org/10.1364/AO.28.004344).
- [16] A. H. Compton. “A Quantum Theory of the Scattering of X-rays by Light Elements”. In: *Physical Review* 21.5 (May 1923). Publisher: American Physical Society, pp. 483–502. DOI: [10.1103/PhysRev.21.483](https://doi.org/10.1103/PhysRev.21.483).
- [17] C. E. Shannon. “Mathematical Theory of the Differential Analyzer”. en. In: *Journal of Mathematics and Physics* 20.1-4 (1941), pp. 337–354. ISSN: 1467-9590. DOI: [10.1002/sapm1941201337](https://doi.org/10.1002/sapm1941201337).
- [18] A. C. Kay. “Microelectronics and the Personal Computer”. In: *Scientific American* 237.3 (1977), pp. 230–245. ISSN: 0036-8733.
- [19] J. H. Seiradakis and M. G. Edmunds. “Our current knowledge of the Antikythera Mechanism”. en. In: *Nature Astronomy* 2.1 (Jan. 2018), pp. 35–42. ISSN: 2397-3366. DOI: [10.1038/s41550-017-0347-2](https://doi.org/10.1038/s41550-017-0347-2).
- [20] W. M. Brown. “Synthetic Aperture Radar”. In: *IEEE Transactions on Aerospace and Electronic Systems* AES-3.2 (Mar. 1967), pp. 217–229. ISSN: 1557-9603. DOI: [10.1109/TAES.1967.5408745](https://doi.org/10.1109/TAES.1967.5408745).
- [21] N. Einspruch. *Application Specific Integrated Circuit (ASIC) Technology*. en. Academic Press, Dec. 2012. ISBN: 978-0-323-15323-2.

- [22] E. Fischer. “The Coast and Geodetic Survey Tide Predicting Machine No. 2”. In: *Popular Astronomy vol. 20 pp.269-285* (1912).
- [23] C. Bissell. “Historical perspectives - The Moniac A Hydromechanical Analog Computer of the 1950s”. In: *IEEE Control Systems Magazine* 27.1 (Feb. 2007), pp. 69–74. ISSN: 1941-000X. DOI: [10.1109/MCS.2007.284511](https://doi.org/10.1109/MCS.2007.284511).
- [24] J. Kilby. “Invention of the integrated circuit”. In: *IEEE Transactions on Electron Devices* 23.7 (July 1976), pp. 648–654. ISSN: 1557-9646. DOI: [10.1109/T-ED.1976.18467](https://doi.org/10.1109/T-ED.1976.18467).
- [25] N. Dutt and D. Gajski. “Design synthesis and silicon compilation”. In: *IEEE Design & Test of Computers* 7.6 (Dec. 1990), pp. 8–23. ISSN: 1558-1918. DOI: [10.1109/54.64954](https://doi.org/10.1109/54.64954).
- [26] P. Kurup and T. Abbasi. *Logic Synthesis Using Synopsys®*. en. Springer Science & Business Media, Dec. 2012. ISBN: 978-1-4613-1455-4.
- [27] R. Vuerinckx, Y. Rolain, J. Schoukens, and R. Pintelon. “Design of stable IIR filters in the complex domain by automatic delay selection”. In: *IEEE Transactions on Signal Processing* 44.9 (Sept. 1996), pp. 2339–2344. ISSN: 1941-0476. DOI: [10.1109/78.536690](https://doi.org/10.1109/78.536690).
- [28] R. A. Roberts and C. T. Mullis. *Digital signal processing*. USA: Addison-Wesley Longman Publishing Co., Inc., 1987. ISBN: 978-0-201-16350-6.
- [29] Z. Chen. “Quantitative Measurement and Development of Back-end Processing System-on-Chip for the Pixelated CdZnTe Detector”. PhD thesis. University of Michigan, 2022.
- [30] C. Cummings. *Asynchronous and Synchronous Reset Design Techniques - Part Deux*. Tech. rep. Sunburst Design, 2003.
- [31] D. Xu, Z. He, C. E. Lehner, and F. Zhang. “4-pi Compton imaging with single 3D position-sensitive CdZnTe detector”. In: *Hard X-Ray and Gamma-Ray Detector Physics VI*. Vol. 5540. SPIE, Oct. 2004, pp. 144–155. DOI: [10.1117/12.563905](https://doi.org/10.1117/12.563905).
- [32] D. Shy, M. Petryk, Z. Chen, S. Thompson, and Z. He. “Time encoded imaging of extended gamma-ray sources in 3D using the OrionUM pixelated CdZnTe with subpixel estimation and depth of interaction corrections”. en. In: *Journal of Instrumentation* 16.03 (Mar. 2021). Publisher: IOP Publishing, P03013. ISSN: 1748-0221. DOI: [10.1088/1748-0221/16/03/P03013](https://doi.org/10.1088/1748-0221/16/03/P03013).
- [33] F. Zhang, Z. He, D. Xu, G. Knoll, D. Wehe, and J. Berry. “Improved resolution for 3-D position sensitive CdZnTe spectrometers”. In: *IEEE Transactions on Nuclear Science* 51.5 (Oct. 2004), pp. 2427–2431. ISSN: 1558-1578. DOI: [10.1109/TNS.2004.835635](https://doi.org/10.1109/TNS.2004.835635).
- [34] S. Joshi. “Coded Aperture Imaging Applied to Pixelated CdZnTe Detectors”. PhD thesis. University of Michigan, 2014.

- [35] Z. Chen, Y. Zhu, and Z. He. “Intrinsic photopeak efficiency measurement and simulation for pixelated CdZnTe detector”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 980 (Nov. 2020), p. 164501. ISSN: 0168-9002. DOI: [10.1016/j.nima.2020.164501](https://doi.org/10.1016/j.nima.2020.164501).
- [36] Y. Zhu, F. Zhang, Y. A. Boucher, J. Mann, and Z. He. “Performance comparison of steering-grid and simple-pixel CdZnTe detectors”. In: *2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC)*. Oct. 2013, pp. 1–5. DOI: [10.1109/NSSMIC.2013.6829842](https://doi.org/10.1109/NSSMIC.2013.6829842).
- [37] A. E. Bolotnikov, G. S. Camarda, Y. Cui, A. Hossain, G. Yang, H. W. Yao, and R. B. James. “Internal Electric-Field-Lines Distribution in CdZnTe Detectors Measured Using X-Ray Mapping”. In: *IEEE Transactions on Nuclear Science* 56.3 (June 2009), pp. 791–794. ISSN: 1558-1578. DOI: [10.1109/TNS.2008.2007904](https://doi.org/10.1109/TNS.2008.2007904).
- [38] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. en. MIT Press, Nov. 2016. ISBN: 978-0-262-33737-3.
- [39] K. Hornik, M. Stinchcombe, and H. White. “Multilayer feedforward networks are universal approximators”. en. In: *Neural Networks* 2.5 (Jan. 1989), pp. 359–366. ISSN: 0893-6080. DOI: [10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).
- [40] K. He, X. Zhang, S. Ren, and J. Sun. “Deep Residual Learning for Image Recognition”. In: 2016, pp. 770–778. (Visited on 12/02/2022).
- [41] J. C. Kim, W. R. Kaye, W. Wang, F. Zhang, and Z. He. “Impact of drift time variation on the Compton image from large-volume CdZnTe crystals”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 683 (Aug. 2012), pp. 53–62. ISSN: 0168-9002. DOI: [10.1016/j.nima.2012.04.057](https://doi.org/10.1016/j.nima.2012.04.057).
- [42] R. D. Reymond. “Pair and triplet production revisited for the radiologist”. In: *American Journal of Roentgenology* 114.3 (Mar. 1972), pp. 639–644. ISSN: 0361-803X. DOI: [10.2214/ajr.114.3.639](https://doi.org/10.2214/ajr.114.3.639).
- [43] R. W. Siegel. “Positron Annihilation Spectroscopy”. en. In: (), p. 34.
- [44] C. Köhn and U. Ebert. “Angular distribution of Bremsstrahlung photons and of positrons for calculations of terrestrial gamma-ray flashes and positron beams”. en. In: *Atmospheric Research* 135-136 (Jan. 2014), pp. 432–465. ISSN: 0169-8095. DOI: [10.1016/j.atmosres.2013.03.012](https://doi.org/10.1016/j.atmosres.2013.03.012).
- [45] H. Bethe, W. Heitler, and P. A. M. Dirac. “On the stopping of fast particles and on the creation of positive electrons”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 146.856 (Aug. 1934), pp. 83–112. DOI: [10.1098/rspa.1934.0140](https://doi.org/10.1098/rspa.1934.0140).

- [46] G. Kanbach, D. L. Bertsch, A. Favale, C. E. Fichtel, R. C. Hartman, R. Hofstadter, E. B. Hughes, S. D. Hunter, B. W. Hughlock, D. A. Kniffen, Y. C. Lin, H. A. Mayer-Hasselwander, P. L. Nolan, K. Pinkau, H. Rothermel, E. Schneid, M. Sommer, and D. J. Thompson. “The project EGRET (energetic gamma-ray experiment telescope) on NASA’s Gamma-Ray Observatory GRO”. en. In: *Space Science Reviews* 49.1 (Jan. 1989), pp. 69–84. ISSN: 1572-9672. DOI: [10.1007/BF00173744](https://doi.org/10.1007/BF00173744).
- [47] R. Garner. *Fermi Spacecraft and Instruments*. und. Text. Oct. 2017.
- [48] *Appendix G to the NASA Research Announcement for the CGRO Guest Investigator Program*.
- [49] J. Xia. “Interaction Reconstruction in Digital 3-D CdZnTe Under Various Circumstances”. PhD thesis. University of Michigan, 2019.
- [50] D. Shy, J. Xia, and Z. He. “Artifacts in High-Energy Compton Imaging With 3-D Position-Sensitive CdZnTe”. en. In: *IEEE Transactions on Nuclear Science* 67.8 (Aug. 2020), pp. 1920–1928. ISSN: 0018-9499, 1558-1578. DOI: [10.1109/TNS.2020.3005834](https://doi.org/10.1109/TNS.2020.3005834).
- [51] W. van Roosbroeck. “Theory of the Yield and Fano Factor of Electron-Hole Pairs Generated in Semiconductors by High-Energy Particles”. In: *Physical Review* 139.5A (Aug. 1965), A1702–A1716. DOI: [10.1103/PhysRev.139.A1702](https://doi.org/10.1103/PhysRev.139.A1702).
- [52] S. Anderson. “Event Classification for 3-D Position Sensitive Semiconductor Detectors”. PhD thesis. University of Michigan, 2011.
- [53] H. Al Hamrashdi, S. D. Monk, and D. Cheneler. “Passive Gamma-Ray and Neutron Imaging Systems for National Security and Nuclear Non-Proliferation in Controlled and Uncontrolled Detection Areas: Review of Past and Current Status”. en. In: *Sensors* 19.11 (Jan. 2019), p. 2638. ISSN: 1424-8220. DOI: [10.3390/s19112638](https://doi.org/10.3390/s19112638). (Visited on 12/01/2022).
- [54] M. D. Alam, S. S. Nasim, and S. Hasan. “Recent progress in CdZnTe based room temperature detectors for nuclear radiation monitoring”. en. In: *Progress in Nuclear Energy* 140 (Oct. 2021), p. 103918. ISSN: 0149-1970. DOI: [10.1016/j.pnucene.2021.103918](https://doi.org/10.1016/j.pnucene.2021.103918).
- [55] Y. Wei, Z. Zhang, Y. Zhang, C. Wang, S. Wen, J. Dong, Z. Li, X. Wang, Z. Xu, G. Huang, and S. Liu. “Performance of the BGO Detector Element of the DAMPE Calorimeter”. In: *IEEE Transactions on Nuclear Science* 63.2 (Apr. 2016), pp. 548–551. ISSN: 1558-1578. DOI: [10.1109/TNS.2016.2541690](https://doi.org/10.1109/TNS.2016.2541690).
- [56] R. Carchon, M. Moeslinger, L. Bourva, C. Bass, and M. Zendel. “Gamma radiation detectors for safeguards applications”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. Proceedings of the 11th Symposium on Radiation Measurements and

- Applications 579.1 (Aug. 2007), pp. 380–383. ISSN: 0168-9002. DOI: [10.1016/j.nima.2007.04.086](https://doi.org/10.1016/j.nima.2007.04.086).
- [57] Y. Zhu and Z. He. “Performance of Larger-Volume $40 \times 40 \times 10$ - and $40 \times 40 \times 15$ -mm³ CdZnTe Detectors”. In: *IEEE Transactions on Nuclear Science* 68.2 (Feb. 2021), pp. 250–255. ISSN: 1558-1578. DOI: [10.1109/TNS.2021.3052133](https://doi.org/10.1109/TNS.2021.3052133).
- [58] K. Hitomi, T. Shoji, and K. Ishii. “Advances in TlBr detector development”. en. In: *Journal of Crystal Growth. Compound Semiconductors and Scintillators for Radiation Detection Applications: A Special Tribute to the Research of Michael Schieber* 379 (Sept. 2013), pp. 93–98. ISSN: 0022-0248. DOI: [10.1016/j.jcrysgro.2013.03.002](https://doi.org/10.1016/j.jcrysgro.2013.03.002).
- [59] V. Lordi. “Point defects in Cd(Zn)Te and TlBr: Theory”. en. In: *Journal of Crystal Growth. Compound Semiconductors and Scintillators for Radiation Detection Applications: A Special Tribute to the Research of Michael Schieber* 379 (Sept. 2013), pp. 84–92. ISSN: 0022-0248. DOI: [10.1016/j.jcrysgro.2013.03.003](https://doi.org/10.1016/j.jcrysgro.2013.03.003).
- [60] C. Leak. “Techniques for Pixelated Ambipolar-Sensitive Semiconductor Gamma-Ray Spectrometers”. PhD thesis. University of Michigan, 2021.
- [61] K. Hitomi, Y. Kikuchi, T. Shoji, and K. Ishii. “Polarization Phenomena in TlBr Detectors”. In: *IEEE Transactions on Nuclear Science* 56.4 (Aug. 2009), pp. 1859–1862. ISSN: 1558-1578. DOI: [10.1109/TNS.2009.2013349](https://doi.org/10.1109/TNS.2009.2013349).
- [62] G. Ariño-Estrada and J. R. Ellin. “High accuracy gamma-ray detection using Cherenkov charge induction sensors (Conference Presentation)”. In: *Hard X-Ray, Gamma-Ray, and Neutron Detector Physics XXIV*. Vol. PC12241. SPIE, Oct. 2022, PC122410C. DOI: [10.1117/12.2635910](https://doi.org/10.1117/12.2635910).
- [63] U. e. a. Roy. “Growth and characterization of detector-grade CdMnTe by the vertical Bridgman technique”. In: *AIP Advances* (2018). DOI: [10.1063/1.5040362](https://doi.org/10.1063/1.5040362).
- [64] U. N. Roy, G. S. Camarda, Y. Cui, and R. B. James. “Advances in CdZnTeSe for Radiation Detector Applications”. en. In: *Radiation* 1.2 (June 2021), pp. 123–130. ISSN: 2673-592X. DOI: [10.3390/radiation1020011](https://doi.org/10.3390/radiation1020011).
- [65] L. Meng, Z. He, B. Alexander, and J. Sandoval. “Spectroscopic performance of thick HgI/sub 2/ detectors”. In: *IEEE Transactions on Nuclear Science* 53.3 (June 2006), pp. 1706–1712. ISSN: 1558-1578. DOI: [10.1109/TNS.2006.873714](https://doi.org/10.1109/TNS.2006.873714).
- [66] J. Paul Attfield, P. Lightfoot, and R. E. Morris. “Perovskites”. en. In: *Dalton Transactions* 44.23 (2015), pp. 10541–10542. DOI: [10.1039/C5DT90083B](https://doi.org/10.1039/C5DT90083B).
- [67] A. Kojima, K. Teshima, Y. Shirai, and T. Miyasaka. “Organometal Halide Perovskites as Visible-Light Sensitizers for Photovoltaic Cells”. In: *Journal of the American Chemical Society* 131.17 (May 2009), pp. 6050–6051. ISSN: 0002-7863. DOI: [10.1021/ja809598r](https://doi.org/10.1021/ja809598r).

- [68] H. Min, D. Y. Lee, J. Kim, G. Kim, K. S. Lee, J. Kim, M. J. Paik, Y. K. Kim, K. S. Kim, M. G. Kim, T. J. Shin, and S. Il Seok. “Perovskite solar cells with atomically coherent interlayers on SnO₂ electrodes”. en. In: *Nature* 598.7881 (Oct. 2021), pp. 444–450. ISSN: 1476-4687. DOI: [10.1038/s41586-021-03964-8](https://doi.org/10.1038/s41586-021-03964-8).
- [69] Y. He, M. Petryk, Z. Liu, D. G. Chica, I. Hadar, C. Leak, W. Ke, I. Spanopoulos, W. Lin, D. Y. Chung, B. W. Wessels, Z. He, and M. G. Kanatzidis. “CsPbBr₃ perovskite detectors with 1.4% energy resolution for high-energy gamma-rays”. en. In: *Nature Photonics* 15.1 (Jan. 2021), pp. 36–42. ISSN: 1749-4893. DOI: [10.1038/s41566-020-00727-1](https://doi.org/10.1038/s41566-020-00727-1).
- [70] H. Wei and J. Huang. “Halide lead perovskites for ionizing radiation detection”. en. In: *Nature Communications* 10.1 (Mar. 2019), p. 1066. ISSN: 2041-1723. DOI: [10.1038/s41467-019-08981-w](https://doi.org/10.1038/s41467-019-08981-w).
- [71] D. Y. Chung, M. G. Kanatzidis, F. Meng, and C. D. Malliakas. “Synthesis, purification, and crystal growth of perovskite semiconductor CsPbBr₃ as a new candidate for gamma-ray detector (Conference Presentation)”. In: *Hard X-Ray, Gamma-Ray, and Neutron Detector Physics XVIII*. Vol. 9968. SPIE, Nov. 2016, pp. 188–188. DOI: [10.1117/12.2238221](https://doi.org/10.1117/12.2238221).
- [72] D. N. Dirin, I. Cherniukh, S. Yakunin, Y. Shynkarenko, and M. V. Kovalenko. “Solution-Grown CsPbBr₃ Perovskite Single Crystals for Photon Detection”. In: *Chemistry of Materials* 28.23 (Dec. 2016). Publisher: American Chemical Society, pp. 8470–8474. ISSN: 0897-4756. DOI: [10.1021/acs.chemmater.6b04298](https://doi.org/10.1021/acs.chemmater.6b04298).
- [73] W. Koehler, M. Streicher, S. O’Neal, and Z. He. “A Correction Factor to the Two-Bias Method for Determining Mobility-Lifetime Products in Pixelated Detectors”. In: *IEEE Transactions on Nuclear Science* 63.3 (June 2016), pp. 1832–1838. ISSN: 1558-1578. DOI: [10.1109/TNS.2016.2548420](https://doi.org/10.1109/TNS.2016.2548420).
- [74] V. Gerrish. “Polarization and gain in mercuric iodide gamma-ray spectrometers”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 322.3 (Nov. 1992), pp. 402–413. ISSN: 0168-9002. DOI: [10.1016/0168-9002\(92\)91205-N](https://doi.org/10.1016/0168-9002(92)91205-N).
- [75] S. O’Neal, Z. He, and C. Leak. “Analysis of High-Energy Tailing in TlBr Detectors”. In: *IEEE Transactions on Nuclear Science* 65.3 (Mar. 2018), pp. 955–960. ISSN: 1558-1578. DOI: [10.1109/TNS.2018.2804165](https://doi.org/10.1109/TNS.2018.2804165).
- [76] S. F. S. S. F. Soares. “Photoconductive Gain in a Schottky Barrier Photodiode”. en. In: *Japanese Journal of Applied Physics* 31.2R (Feb. 1992), p. 210. ISSN: 1347-4065. DOI: [10.1143/JJAP.31.210](https://doi.org/10.1143/JJAP.31.210).

- [77] Y. He, Z. Liu, K. M. McCall, W. Lin, D. Y. Chung, B. W. Wessels, and M. G. Kanatzidis. “Perovskite CsPbBr₃ single crystal detector for alpha-particle spectroscopy”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 922 (Apr. 2019), pp. 217–221. ISSN: 0168-9002. DOI: [10.1016/j.nima.2019.01.008](https://doi.org/10.1016/j.nima.2019.01.008).