

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Polarimetric Studies of the Long Duration Gamma-Ray Burst GRB 160530A with the Compton Spectrometer and Imager

Permalink

<https://escholarship.org/uc/item/3b8664q5>

Author

Lowell, Alexander

Publication Date

2017

Peer reviewed|Thesis/dissertation

**Polarimetric Studies of the Long Duration Gamma-Ray Burst GRB 160530A
with the Compton Spectrometer and Imager**

by

Alexander Lowell

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Steven Edward Boggs, Chair
Professor Daniel Kasen
Professor Kai Vetter

Fall 2017

**Polarimetric Studies of the Long Duration Gamma-Ray Burst GRB 160530A
with the Compton Spectrometer and Imager**

Copyright 2017
by
Alexander Lowell

Abstract

Polarimetric Studies of the Long Duration Gamma-Ray Burst GRB 160530A with the Compton Spectrometer and Imager

by

Alexander Lowell

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Steven Edward Boggs, Chair

Gamma-ray bursts (GRBs) are highly luminous flashes of gamma-rays that last from hundreds of milliseconds up to hundreds of seconds, peak in the soft gamma-ray band (0.1 MeV to 1 MeV), represent ultra-relativistic motion, and originate from cosmological distances. Despite thousands of observations of GRB prompt emission, and hundreds of observations of GRB afterglows, many details pertaining to the nature of gamma-ray bursts remain elusive. A wide variety of models exist that describe the progenitors, central engines, emission mechanism, geometric structure, and magnetization of GRBs. Polarization measurements of the GRB prompt emission are believed to give essential clues where spectral and temporal observations fall short.

The Compton Spectrometer and Imager (COSI) is a balloon borne Compton polarimeter, imager, and spectrometer sensitive in the 0.2 MeV to 5 MeV band. One of COSI's main science goals is measuring the polarization of astrophysical gamma-ray sources, including GRBs. Polarimetric measurements with a Compton polarimeter are possible due to the sensitivity of the Compton scattering cross section upon the orientation of the photon's electric field vector. We verified COSI's polarimetric capabilities in the laboratory using partially polarized and unpolarized beams, and compared the measurements to simulations. Overall, the measurements reflect the simulations closely, and we were able to place upper limits on the systematic error on the polarization level of 3% – 4%.

In May of 2016, COSI took flight from Wanaka, New Zealand aboard a NASA Super Pressure Balloon. COSI continuously telemetered science data for 46 days before landing in Southern Peru. On May 30, 2016, COSI clearly detected the long duration gamma-ray burst GRB 160530A. We performed a polarization analysis of this GRB using both a standard method (SM) and a maximum likelihood method (MLM). The measured polarization level was below the 99% confidence minimum detectable polarization level (MDP) using both analysis methods ($MDP = 72.3\% \pm 0.8\%$ for the SM, and $MDP = 57.5\% \pm 0.8\%$ for the MLM), so we claim a non-detection of polarization for this GRB. Using the MLM, we placed a 90% confidence upper limit of 46% on the polarization level. GRBs with polarization levels

higher than 46% have been reported in the literature, and so this result will be useful in future meta-analyses that consider the statistical distribution of polarization levels across the GRB population.

Contents

Contents	i
List of Figures	v
List of Tables	xvi
1 Gamma Ray Bursts	1
1.1 History of GRB Observations and Results	2
1.1.1 First Observations — Vela Spacecraft	2
1.1.2 Compton Gamma-Ray Observatory	2
1.1.3 First Afterglows and Supernovae Associations in the BeppoSAX and HETE-2 Era	5
1.1.4 Rapid Afterglow Identification with the Swift Gamma-Ray Burst Mission	6
1.1.5 High Energy Observations with the Fermi Space Telescope	8
1.2 Progenitors and Central Engines	9
1.3 Origins of the Prompt Emission	10
1.3.1 Prompt Emission in the “Fireball” Model	10
1.3.2 Prompt Emission in Electromagnetic Models	11
1.4 Polarization of the Prompt Emission	11
1.4.1 Polarization from Synchrotron Emission	13
1.4.2 Polarization from Inverse Compton Scattering and Photospheric Emis- sion	15
1.4.3 Statistical Distribution of Polarization Levels	17
1.5 Summary	18
1.6 Thesis Outline	18
2 COSI Overview	21
2.1 Compton Telescopes	21
2.2 Types of Compton Telescopes	24
2.3 The Compton Spectrometer and Imager	24
2.3.1 Germanium Detectors	25
2.3.2 Vacuum Cryostat and Cryocooler	26

2.3.3	Anti-Coincidence Shield	28
2.3.4	Signal Acquisition and Instrumentation	28
2.4	Gondola Systems	34
2.4.1	Flight Computer (GCU)	34
2.4.2	Power System	40
2.4.3	Aspect and GPS Systems	40
2.5	Ground Software	41
2.6	Calibrations	41
2.6.1	Energy Calibration	41
2.6.2	Depth Calibration	42
2.6.3	Crosstalk Correction	43
2.6.4	Charge Loss Correction	43
2.6.5	Strip Pairing	43
2.7	Data Analysis	44
2.7.1	Simulations and Mass Model	46
2.7.2	Event Selections	47
2.7.3	Compton Imaging	49
2.7.4	Compton Polarimetry	49
3	The 2016 COSI Super Pressure Balloon Campaign	50
3.1	From NCT to COSI	50
3.2	Super Pressure Balloon	51
3.3	2014 Antarctic Flight	52
3.4	2016 Flight from Wanaka, New Zealand	53
3.4.1	Ground Calibrations	53
3.4.2	Flight Summary and Overview	53
3.4.3	Thermal Environment	54
3.4.4	Altitude Variations	55
3.4.5	Background Environment and Relativistic Electron Precipitation	56
3.4.6	High Voltage Anomalies	58
3.4.7	Performance of Telemetry Systems	58
3.4.8	Ground Operations	64
3.5	COSI 2016 Flight Source Catalog	64
3.5.1	The Crab Nebula	66
3.5.2	Cygnus X-1	66
3.5.3	Centaurus A	66
3.5.4	Galactic positron annihilation emission	67
3.5.5	Relativistic electron precipitation	67
3.5.6	GRB 160530A	67
3.6	Summary	68
4	Compton Polarimetry with COSI	69

4.1	Principles of Compton Polarimetry	70
4.1.1	COSI as a Compton Polarimeter	71
4.1.2	Standard Analysis Method	73
4.1.3	Maximum Likelihood Method	74
4.2	Modulation reducing effects	78
4.2.1	Event Selections	79
4.2.2	Event Reconstruction Efficiency	80
4.2.3	Event Resolution with Finite Position Resolution	80
4.2.4	Measurement Error on the Azimuthal Scattering Angle	81
4.3	Minimum Detectable Polarization	82
4.4	Summary and Conclusion	83
5	Polarimetric Validation	84
5.1	Experimental Configuration	84
5.2	Analysis Methods	87
5.3	Polarized Runs	89
5.4	Event Selections	89
5.4.1	Energy Selection	89
5.4.2	Distance Selection	91
5.4.3	ARM Selection	91
5.5	Results with d_1 Distance Selections	91
5.6	Results with d_2 Distance Selections	93
5.6.1	Discussion	102
5.7	Azimuthal Scattering Angle Distributions of Unpolarized, Calibration Sources	102
5.7.1	Event Selections	104
5.7.2	Discussion	104
5.8	Conclusions	104
6	Polarimetric Analysis of GRB 160530A	110
6.1	GRB 160530A	110
6.2	Polarimetric Analysis	114
6.2.1	Standard Method	114
6.2.2	Maximum Likelihood Method	116
6.3	MDP Improvement	121
6.4	Discussion	124
7	Towards a Better Polarization Sensitivity	126
7.1	N-site Maximum Likelihood Method	126
7.1.1	Stokes Vector Framework	127
7.1.2	NMLM Algorithm	129
7.1.3	Simulation Setup	130
7.1.4	Results	130

7.1.5	Discussion	133
7.2	Atmospheric Modulation Effect	134
7.2.1	Simulation Setup	136
7.2.2	Results	137
7.2.3	Discussion	137
	Bibliography	140
A	Positional Calibrations	148
A.1	Depth Calibration: Classic Approach	149
A.2	Depth Calibration: New Approach	151
A.3	Angular Resolution	153
A.4	Sub-Strip Positioning Using Spectator Strips	157
A.5	Summary	158

List of Figures

1.1	Spatial distribution of 1637 BATSE bursts from the fourth BATSE catalog (Figure 5 of Paciesas et al., 1999), showing the isotropic distribution of GRBs in galactic coordinates.	3
1.2	Histograms of Band model fit parameters to a sample of 5500 BATSE GRB spectra (Figures 7, 8, and 9 of Preece et al., 2000). The break energies E_p are clustered around ~ 250 keV, while the average low energy power law index α is ~ -1 , and the high energy power law index β tends to fall between -2 and -2.5.	4
1.3	T_{90} distribution for 2041 BATSE GRBs, along with Gaussian fits to the short and long GRB populations (Figure 1 from Lazzati & Perna, 2009). T_{90} is defined as the time over which the detector collects between 5% and 95% of the total fluence.	5
1.4	A selection of X-ray afterglows for long GRBs (left), and short GRBs (right) as seen by Swift (Figure 6 from Gehrels et al., 2009). A variety of different behaviors can be seen including steep to shallow transitions, jet breaks, and late time flares.	8
1.5	$\Delta\chi^2$ maps for the IKAROS/GAP observations of GRB 110301A (left) and GRB 110721A (right) indicating detections of non-zero polarization levels at significance levels of 3.7σ and 3.3σ respectively (Figure 3 from Yonetoku et al., 2012). These results, in concert with the observed polarization angle changes in GRB 100826A (Yonetoku et al., 2011a) were interpreted as evidence in support of large scale magnetic field structures and synchrotron emission as the dominant GRB emission mechanism.	12
1.6	A cartoon illustrating the synchrotron emission from oriented magnetic fields (SO) model. On the left, the visible region of the jet Γ^{-1} centered along the viewing angle θ_v is contained within the jet opening angle θ_j . The magnetic field lines (blue) are assumed to be helical and result from a Poynting flux from the central engine. Within the visible region, the field lines are rather straight, the corresponding polarization vector directions (red arrows) are well defined, and thus a large polarization level is observed ($\Pi \sim 40\%$). On the right, lower values of Γ result in a larger visible emitting region. Polarization vectors with different orientations are now visible from other regions of the jet, and thus the overall polarization level is significantly reduced. The diagrams are adapted from Toma (2013).	13

- 1.7 A cartoon illustrating the synchrotron emission from random magnetic fields (SR) model. On the left, the visible region Γ^{-1} is contained within the jet opening angle θ_j . Within the visible region, the magnetic field orientations (blue arrows) are random, and the polarization vectors (red arrows) are axisymmetric around the line of sight, leading to no net polarization. On the right, the visible region coincides with the edge of the jet, leading to a loss of symmetry and a net polarization ($\Pi \sim 25\% - 30\%$). However, the situation on the right may be accompanied by a lower observed intensity, as the visible region is not filled with emission. The diagrams are adapted from Toma (2013). 15
- 1.8 A cartoon illustrating the Compton Drag model, which relies on inverse Compton emission to produce a high polarization level. Similar to the SR model, when the visible region Γ^{-1} is contained within the jet opening angle θ_j , the polarization vectors are axisymmetric around the line of sight, and no net polarization results. On the right, the visible region coincides with the jet edge, the symmetry is broken, and a high polarization results (up to $\Pi \sim 80\%$). Once again, the overall intensity of the prompt emission may be lower for the situation on the right, as the visible region contains less of the jet emission. 16
- 1.9 Simulated, cumulative distributions of polarization levels measured by the proposed POETS satellite (60 keV to 500 keV) for the SO, SR, and CD models over a period of time such that 200 bursts are detected. A net polarization level of $\Pi \sim 30\% - 40\%$ is frequently observed for the SO model. The CD model is capable of yielding higher polarization levels $\Pi > 60\%$, albeit less frequently. The SR model appears to be inefficient at producing a net polarization. Cumulative distributions for a real sample of GRBs may be compared to those shown here to help discern between models. 17
- 2.1 A schematic drawing of a canonical Compton event. An incident photon with energy E_0 (purple arrow) Compton scatters at r_1 with Compton scattering angle θ and azimuthal scattering angle η , depositing an energy E_1 . The scattered photon is subsequently photoabsorbed at r_2 , depositing an energy of E_2 . By measuring the energies and locations of the interactions, the origin of the photon can be constrained to a Compton cone (blue). By measuring η for a large sample of events, the polarization level and angle of the incident beam may be inferred. The COSI germanium detector array is shown in the background. 22
- 2.2 Backprojected Compton cones for correctly reconstructed photons in spherical coordinates (image is 4π sr). The arcs are found to intersect at the location of the source, which appears as a yellow hot spot. The smaller “circles” correspond to Compton cones with smaller Compton scattering angles. 23
- 2.3 Left: A single GeD in front of a mirror, showing the orthogonality of the strips between the front and back sides. Right: The twelve GeD array during integration. The copper U-shaped bar dividing the array vertically is the upper segment of the cold finger. 25

2.4	A photograph of the cryostat with various components annotated. The preamplifier boxes, which ordinarily mount where the flex circuits feed through the cryostat wall, are not present in this image.	27
2.5	A photograph showing the COSI cryostat integrated into the top of the gondola, with various component annotated.	27
2.6	A diagram illustrating the signal acquisition process for the COSI GeDs. A photon deposits energy in the germanium, resulting in a number of electron hole pairs. These pairs are swept to the surface electrode by an applied electric field. The preamplifiers convert the collected charge into a voltage pulse, which is then sent to a slow unipolar shaper, and a fast bipolar shaper. The peak value of the unipolar shaper pulses is sampled by an analog to digital converter to facilitate pulse height spectroscopy. The zero crossing time of the bipolar shaper signal is measured by a 200 MHz timer, and the time difference τ between the cathode and anode signal is used to infer the interaction depth.	29
2.7	A schematic of a single preamplifier channel, with several annotations corresponding to the circuit description in Section 2.3.4.	30
2.8	A preamplifier box with the lid removed. The ten preamp channels on the top board are visible. The front side of the preamp box (not visible) connects to a flex circuit, which carries the charge signals from one side of a detector through the wall of the cryostat.	32
2.9	A COSI card cage used to readout a single GeD. The boards from bottom to top are as follows: high voltage power supply board, low voltage power supply board, DSP board, 8x analog boards. One analog board has been extracted so that the ten shaper channels are visible.	32
2.10	A photograph of the entire COSI gondola hanging from the launch vehicle in Wanaka, New Zealand during a test. Various subsystems are annotated for reference.	35
2.11	A flow chart showing the relationship between the flight computer (GCU) and other systems on the COSI gondola. Additional connections not involving the GCU are also shown. The connections are labeled with their function, and the type of connection is in parentheses. A few connections are not labeled, as they are not transparent to the science team.	36
2.12	Spectra of several radioactive sources in channel space for a single strip electrode. The channel number on the x axis is proportional to the voltage pulse height from the slow shaper. Each peak in the spectra is fit with an analytic model of the emission line, thus yielding a curve that maps channel number to photon energy.	42

- 2.13 A diagram illustrating the strip pairing process. In this example, a photon interacts at two locations marked by the green circles. This results in 3 triggered x strips and 2 triggered y strips. The strip pairing must successfully determine that x strip 9 and y strip 17 should be paired, and that x strips 11 and 12 (which shared the total charge) should be paired with y strip 15. Charge sharing, statistical fluctuations on the measured energies, hits below threshold, and strips with multiple hits can complicate the strip pairing process. 44
- 2.14 The data pipelines for measurements (top) and simulations (bottom). The simulated pipeline makes use of the Detector Effects Engine, which adds realism to the Cosima simulation. Also, the simulated pipeline does not include the cross-talk correction, as that is currently not modeled in the simulations. 45
- 2.15 A top view of the COSI mass model, which is used for performing Cosima simulations, as well as global 3D positioning of measured events. 46
- 2.16 A diagram illustrating the definition of the Angular Resolution Measure (ARM). The source of the photons is the small, filled black circle. The open circles represent the arcs traced out by the Compton cones of two separate photons. The smallest angular distance between an event's Compton cone and the origin of the photons corresponds to the ARM for that event. When the photon origin falls outside of the arc, the ARM is positive, and when it falls inside, the ARM is negative. Note that ARM values are always relative to a specified location. . . 47
- 2.17 ARM distribution from a Cs-137 (662 keV) calibration run, along with a fit to the data. The x axis represents the smallest angular distance between a Compton cone and the specified source location. The analytical fit function is a superposition of two Lorentzians and an asymmetric Gaussian. This shape is empirical and accurately fits the ARM peak. The FWHM is determined from the fit rather than the data itself. 48
- 3.1 The COSI flight line on launch day. COSI can be seen hanging from the launch vehicle on the left. The inflated SPB is on the right of the image. In the center of the image is the parachute, as well as the uninflated portion of the balloon. This portion expands significantly as the pressure drops with altitude until the balloon makes a round, pumpkin shape at float. 54
- 3.2 Left: A still from a video shot from the gondola, moments before landing in the Atacama desert. The shadow of the COSI payload and parachute are clearly seen on the desert floor. Right: A photograph taken by the recovery team of COSI at its landing site. The exterior structure sustained damage, but the instrument and electronics were recovered in excellent shape. 55

3.3	The COSI flight trajectory throughout the Southern Hemisphere. COSI launched from Wanaka, New Zealand, flew for 46 days, and then landed in Southern Peru. The color of the flight trajectory alternates for each day of flight. A red circle denotes the position of the South Magnetic Pole. The blue/green circles mark COSI's position at the time that the REP events were detected, as well as GRB 160530A.	57
3.4	From top: The altitude, angular distance to the SMP, count rate (> 80 keV) in the shield system , and count rate (> 50 keV) in detector 6 (top layer) for the first third of the flight. The shield rate has been median filtered with a window size of 5 samples to suppress 1 and 2 sample spikes due to excess rate events. The count rate in detector 6 has been box car averaged with a window size of 5 samples to reduce the Poisson noise in the plot. A clear anti-correlation exists between the distance to the SMP and the count rates. The flaring activity near the SMP is due to REP activity.	59
3.5	The same as Figure 3.4, but for the second third of the flight. A clear anti-correlation can be seen once again between the distance to the SMP and the count rates, as well as the altitude and the count rates.	60
3.6	The same as Figure 3.5, but for the final third of the flight. Significant altitude variations were experienced on a daily basis, with corresponding enhancements in the count rates.	61
3.7	Reconstructed event light curves for all COSI events (black) and Compton events only (blue) for the first 12.5 hours of May 30, 2016. This time period included the second REP event (REP2), which began as fast ~ 100 ms pulses and transitioned into slower and brighter pulses just before 07:00:00. GRB 160530A is the spike just after 07:00:00. The relative amplitude of the REP activity over the pre-REP background is significantly smaller in the Compton event light curve, indicating that the REP photons are too low in energy ($\lesssim 200$ keV) to produce many Compton events.	62
3.8	The cumulative transmitted and received data volume for the first ten hours of May 20, 2016. The red curve representing the received data slowly diverges from the transmitted curve as small numbers of packets are lost in the Iridium network over time. At the end of the ten hour period shown here, 2.4% of the data is lost, which is considered to be acceptable. All lost data was eventually recovered using the packet fetcher just before flight termination.	63

3.9	COSI images in galactic coordinates of four persistent astrophysical sources that were detected during the 2016 flight: The Crab Nebula (a), Cyngus X-1 (b), Centaurus A (c), and 511 keV positron annihilation emission from the galactic center (d). The peaks in the images are slightly offset from the known positions, possibly due to the lack of exposure correction, or uncorrected systematic errors in the aspect reconstruction. All of these are point sources, with the exception of the 511 keV image, which is known to be diffuse from INTEGRAL/SPI observations (Bouchet et al., 2010). Efforts are underway to determine if the central bulge in the COSI image is diffuse, and whether or not the apparent extended structure surrounding the bulge is real.	65
4.1	The modulation of the azimuthal scattering angle probability as a function of the Compton scattering angle for various photon energies. At lower energies and for Compton scattering angles near $\theta \sim 90$ degrees, the modulation is strongest. . .	71
4.2	The ideal probability density function (Equation 4.5) for the azimuthal scattering angle for various polarization levels Π and polarization angles η_0 . In these plots, the modulation is $\mu = 0.5$. Notice that even when $\Pi = 1$, there is a non-zero probability of scattering along the electric field vector.	72
4.3	The full PDF used for the MLM analysis (gray bars), along with the ideal PDF (blue and green lines) for two cases $\Pi = 0$ (left) and $\Pi = 1$ (right). The PDFs are drawn for a photon energy of 337.5 keV and a Compton scattering angle of 92.5 degrees. The slice $g(\eta; E, \theta)$ used for these PDFs is valid over the range $E = 325 - 350$ keV and $\theta = 90 - 95$ degrees. The slice $g(\eta; E, \theta)$ used here was taken from the COSI response $\mathcal{H}[E, \theta, \eta]$ for GRB 160530A, which occurred 43.5° off-axis. Error bars are drawn on the PDFs based on the simulation statistics. . .	76
4.4	The corrected, best fit polarization level $\hat{\Pi}/\Pi_{100}$ (top) and best fit angle $\hat{\eta}_0$ (bottom) as a function of the true, simulated polarization level. The blue line in the top panel is drawn with a slope of 1 for reference. These simulations include a background component based on Ling (1975), with a signal purity of $f = 0.72$. The simulated polarization angle was $\eta_0 = 150^\circ$ and is denoted with the dashed line in the bottom plot. The error bars shown in these plots are 1σ and were obtained using MINOS.	79
4.5	The Modulation reduction factor (MRF) as a function of the 1σ azimuthal scattering angle error $\delta\eta$. The red curve is computed using the discrete Fourier transform of the (assumed) normal distribution on η . The black circles come from toy Monte Carlo simulations. Both methods show that the MRF decreases with increasing $\delta\eta$, resulting in a reduced modulation.	81

5.1	A photograph showing the experimental scattering configuration on March 17, 2016. 662 keV photons originate from the Cs-137 source (lower left corner), Compton scatter in the NaI scintillator (top middle), and proceed to interact in the COSI GeDs. The Compton scattered beam is partially polarized, with the polarization vector perpendicular to the plane of scattering. The lead brick placed between the Cs-137 source and COSI blocks the 662 keV photons from interacting in the CsI shields.	85
5.2	A diagram describing the interconnection between various components in the experimental configuration. The signal from the NaI scintillator is sent to a shaper, then a discriminator, and then finally to a digital input on the sGCU, which latches the 48 bit clock value when a rising edge is detected. The sGCU is connected to the system in parallel with the twelve card cages. This way, the sGCU clock is kept synchronized with the GCU and card cage clocks, so that timestamps from events in the NaI may be aligned with COSI events.	87
5.3	A time difference histogram, where the x axis represents the time between a COSI event, and the time of the next NaI time stamp. A clear coincidence peak is seen between 2300 ns and 2800 ns. Chance coincident events were sampled between 3300 ns and 3800 ns.	88
5.4	Top views (parallel perspective) of the three mass models representing the polarized runs on March 16, 17, and 18, 2016. Each configuration uses a different x and y position for the NaI detector, the Cs-137 source, and the lead brick. Each run also had a different z position.	90
5.5	Measured (blue) and simulated (red) energy spectra for the 316 (top), 317 (middle), and 318 (bottom) runs. The shape of the measured scatter peak is in good agreement with the simulated scatter peak.	92
5.6	Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 316 run using the d_1 distance selections.	94
5.7	Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 317 run using the d_1 distance selections.	95
5.8	Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 318 run using the d_1 distance selections.	96
5.9	Measured (blue) and simulated (red) Compton scattering angle distributions for the 316 (top), 317 (middle), and 318 (bottom) runs. The measured and simulated distributions match well for all three runs.	97
5.10	Measured (blue) and simulated (red) distribution of distances between first two interactions for the 316 (top), 317 (middle), and 318 (bottom) runs. Events with a distance between first two interactions of $\lesssim 1$ cm appear to be overrepresented in the simulations, especially for the 316 and 318 runs.	98

5.11 Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 316 run using the d_2 distance selections. 99

5.12 Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 317 run using the d_2 distance selections. 100

5.13 Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 318 run using the d_2 distance selections. 101

5.14 A summary of the fitted parameters and their differences between measurements and simulation. Each point along the x axis represents one of the three polarized runs with either the d_1 or d_2 selections. In the top three plots, blue circles indicate measurements, and red circles indicate simulations (simulation circles have been offset for comparison). The top two plots show the fitted modulations and polarization angles, respectively. The middle panel shows the χ^2_{red} values for each run. The bottom two panels show the difference (measurement - simulation) in the fitted modulations and polarization angles, with the dashed line at zero drawn as a reference. A χ^2 meta-analysis of the differences suggests that we cannot reject the null hypothesis that there is no difference between measurements and simulations. 103

5.15 Binned ASADs corresponding to the unpolarized runs in Tables 5.7 and 5.8. Measurements are drawn in blue, and simulations are drawn in red. The distributions are clearly non-uniform, which makes them useful for probing the instrument systematics. 105

6.1 Two COSI images of GRB 160530A in Galactic coordinates using zero iterations (top) and ten iterations (bottom) of the LM-MLEM image deconvolution algorithm. Zero iterations corresponds to a simple back-projection of the Compton cones. The event selections used to make this image were different than those used for the polarization analysis, as the goal here is imaging performance. The color scale intensity is in arbitrary units. 111

6.2 (top) COSI and Konus light curves plotted for photon energies between 50 and 200 keV. The Konus light curve has been shifted by +0.651 seconds to account for the COSI-Konus delay time. In this energy band, only 2% of events are Compton events; the rest are single-site events. (bottom) The COSI Compton event light curve, using the event selections shown in Table 6.3. 112

6.3 The modulation factor μ_{100} as a function of the simulated polarization angle, along with the best fit constant value of $\mu_{100} = 0.484 \pm 0.002$, $\chi^2_{red} = 0.42$ (dof = 17). For each simulation, $2.4 - 2.5 \times 10^4$ counts were used. 116

- 6.4 The background subtracted azimuthal scattering angle distribution (ASAD) for GRB160530A (top), an ASAD from an unpolarized simulation of GRB160503A used to correct for systematics (middle), and the corrected ASAD (bottom) showing the best fit modulation curve in blue. 118
- 6.5 The MLM correction factor Π_{100} as a function of the simulated polarization angle, along with the best fit constant value $\Pi_{100} = 0.799 \pm 0.003$, $\chi_{red}^2 = 0.68$ (dof = 17). For each simulation, $3.4 - 3.6 \times 10^4$ counts were used. 119
- 6.6 Bootstrap distribution of the polarization level using the MLM. Each entry in the histogram was computed by bootstrap resampling the 542 qualifying events, drawing a sample of the signal purity f from its underlying distribution, running the minimizer, and dividing the result by a value of Π_{100} drawn from its underlying distribution. The dashed line corresponds to the measured value, and the dot-dashed line is the 90th percentile of the distribution. We interpret this as our 90% confidence upper limit on the polarization level. 121
- 6.7 Confidence contours for GRB 160530A, derived from a map of the log likelihood. The optimal polarization level and angle are denoted by the cross. Confidence contours from MINOS are shown for 1σ to 5σ as well as 90% confidence. The dashed line denotes the 90% confidence upper limit on the polarization level as determined using a Monte Carlo/bootstrap technique. The dotted line represents the MDP, or equivalently, the 99% confidence detection threshold. 122
- 6.8 Distribution of polarization levels from simulated trial observations of GRB 160530A using the MLM (top) and SM (bottom). The true polarization level was set to zero for these simulations. In the MLM case, each trial observation was analyzed in order to determine $\hat{\Pi}$, which was then corrected using the MLM correction factor Π_{100} (Equation 4.14). The resulting polarization level was added to the histogram. For the SM, the fitted modulation $\hat{\mu}$ was divided by μ_{100} and then added to the histogram (Equation 4.8). The MDP, which is the 99th percentile of this distribution, is represented by the dark green vertical lines. The lighter green lines represent the standard deviation of the MDP, which was obtained by bootstrap resampling the underlying data vector. The MLM MDP is 20% lower than the SM MDP. 123
- 7.1 A schematic representation of the NMLM simulations. A cube of germanium with side length 1 m was used as an ideal detector. Monoenergetic photons were started from the center of the cube with initial directions parallel to the $-\hat{z}$ direction. 130
- 7.2 Histograms of the inferred polarization level using the SM, MLM, and NMLM, for a true polarization level of $\Pi = 0$. The dashed lines denote the 99th percentile of the distribution i.e. the MDP. The photon energies used were 200 keV (top) 400 keV (middle) and 800 keV (bottom). Each sample of events contained 500 photons. 131

7.3	Histograms of the inferred polarization level (top) and angle (bottom), where the true values are $\Pi = 0.65$ and $\eta_0 = 0^\circ$. The photon energy used here was 400 keV. Each analyzed sample contained 500 photons. The NMLM gives a more sharply peaked distribution about the true values (denoted by the vertical black lines).	132
7.4	A cartoon illustrating the AME principle. The view into the page is looking down from the top of the atmosphere along the $-\hat{z}$ direction. Photons from a polarized GRB move into the page, with their electric field vectors aligned along the \hat{y} direction. These photons Compton scatter, and the scattered photons are denoted by the purple arrows. The density of the purple arrows is drawn to represent the azimuthal scattering angle probability, which is higher in the directions perpendicular to the electric field vector. The detector thus detects an excess of photons from the $\pm\hat{x}$ directions and a deficit from the $\pm\hat{y}$ directions, resulting in a modulation of the azimuthal arrival angles.	135
7.5	A schematic representation of the AME simulations. The atmosphere was modeled as a cylinder of air 100 km tall, and 100 km in diameter. The atmosphere was divided vertically into 100 layers with densities given by the NRLMSISE00 model. Photons from a standard GRB spectrum were started from the top of the atmosphere with initial directions along the $-\hat{z}$ direction (purple arrows). At 33 km, we placed a “detector layer” (blue disk) made of air at the corresponding density. Photons passing through the detector layer had their directions, energies, and positions recorded.	136
7.6	The azimuthal arrival angle distribution at 33 km along with the best fit modulation curve (red) for a fully polarized GRB incident at the top of the atmosphere. The modulation is $\mu_{100} = 0.173 \pm 0.003$. A higher modulation can be achieved by optimizing the polar arrival angle range. Here, no optimizations were performed.	138
7.7	The same as Figure 7.6, but with a polarization level of $\Pi = 0.5$ (50%), and a shift of 90° in the polarization angle. The modulation in this case is $\mu_{50} = 0.087 \pm 0.003$	138
A.1	$z_n(\tau_{\text{sim}})$ for each detector n , as determined from charge transport simulations. For a given detector with specified thickness, bias voltage, and impurity concentration, charges were initiated at various depths z , and the resulting collection time differences τ_{sim} were recorded.	149
A.2	Best fit CTD templates (red) to the τ_{meas} distributions (blue) for individual pixels from detectors 6 and 10.	152

A.3 A diagram illustrating the new approach to the depth calibration. The top histogram is the depth distribution $g_n(z)$ in detector n , determined from Cosima simulations. The bottom histogram is the measured distribution of collection time differences $f_p(\tau_{meas})$ in pixel p . The area under each histogram has been normalized to 1. Assuming the map between depth and collection time difference is one-to-one, then each area element under $g_n(z)$ corresponds to an area element under $f_p(\tau_{meas})$. The area elements are denoted by the green regions under the histogram, and are not drawn to scale. Using this principle, the relationship between depth and collection time difference $z_p(\tau_{meas})$ can be determined for pixel p by solving Equation A.2. The dashed lines show the choice of boundary conditions; the edges of the detector are mapped to the edges of $f_p(\tau_{meas})$ 154

A.4 $z(\tau_{meas})$ for 24 pixels of detector 6. The red pixels are closer to the center of the detector, while the blue pixels are closer to the edge. There is a clear difference in the shape between edge and center pixels that cannot be resolved by stretching and offsetting. 155

A.5 Comparison between the simulated ARM and measured ARM using both depth calibration approaches as a function of energy. The calibration sources used were Ba-133 (356 keV), Na-22 (511 keV), Cs-137 (662 keV), and Y-88 (898 keV). The error bars are 90% confidence. 156

A.6 Spectator strip amplitude vs x position of deposited charge under the energy collecting strip. The black pulse is the fast shaper signal from the energy collecting strip and the collection of colored pulses are the spectator strip pulses vs x 158

A.7 Spectator strip amplitudes vs x for a spectator strip that is two strips away from the energy collecting strip. 158

List of Tables

5.1	Statistics and SM fit results for the 316 run using the d_1 distance selections. . .	94
5.2	Statistics and SM fit results for the 317 run using the d_1 distance selections. . .	95
5.3	Statistics and SM fit results for the 318 run using the d_1 distance selections. . .	96
5.4	Statistics and SM fit results for the 316 run using the d_2 distance selections. . .	99
5.5	Statistics and SM fit results for the 317 run using the d_2 distance selections. . .	100
5.6	Statistics and SM fit results for the 318 run using the d_2 distance selections. . .	101
5.7	Results for unpolarized runs using d_1 distance selections.	106
5.8	Results for unpolarized runs using d_2 distance selections. Note that the AD p -value for the third entry is greater than one; this a result of the algorithm performing an extrapolation. Nonetheless, this indicates that the p -value is high and far from the threshold of statistical significance.	106
5.9	Upper limits on the systematic error for the polarization level for the three polarized runs and for both distance selections. The methodology for the calculation can be found in Section 5.8.	108
6.1	Optimized time windows in units of seconds. Times are with respect to UTC = 1464591800 seconds (May 30th, 2016 07:03:20 UTC). The total time spanned by the windows is 12.945 seconds.	112
6.2	SM event selections, statistics, and fit results. The number of bins in the azimuthal scattering angle distribution was 30 (12° per bin).	117
6.3	MLM event selections, statistics, and fit results.	120
7.1	MDPs and average number of Compton interactions per sample of 500 photons for the SM, MLM, and NMLM at 200 keV, 400 keV, and 800 keV.	133
A.1	Average and standard deviation of fit parameters λ and Δ along with χ_{red}^2 across all pixels for each detector.	151

Acknowledgments

My time in graduate school has been exceptionally enjoyable and stimulating due in large part to the following people:

- My advisor Professor Steven Boggs for his endless positivity, encouragement, and wisdom, without which I am not sure how I would have managed. His upbeat attitude permeates the entire COSI project and team. I have always been impressed by his ability to be in tune with so many different analysis projects, and give crucial insights at times when they were most needed. Through his exceptional grant writing our project was fully funded for many years, which allowed me and my fellow graduate students to flourish and focus on research without worry. I am also grateful for the postdoctoral opportunity he has offered me and I am thrilled to continue working with him on GAPS and other exciting projects. Infinite thanks for all of the amazing opportunities you have provided me with!
- Thesis committee members Professors Steven Boggs, Dan Kasen and Kai Vetter for their feedback on making this thesis more readable. Thanks to Dan for teaching me about radiative processes in astrophysics, the best astronomy course I took in graduate school. Thanks to Kai for his hospitality during my visits to LBL.
- My co-advisor Dr. John Tomsick for his deep astronomy and statistics knowledge, and always keeping me focused on the science. His guidance on COSI as well as several earlier projects at SSL were instrumental to my development as a scientist. Also, many thanks for organizing group social events and fostering a strong sense of community in our group and at SSL in general.
- Dr. Andreas Zoglauer for building the software backbone that supports virtually all COSI data analysis. His patience as a mentor with the graduate and undergraduate students in our group was truly astounding. Thanks for always answering my MEGAlib tech support questions, and I will sincerely miss having you across the hall from me.
- My fellow graduate students Carolyn Kierans and Clio Sleator, whose simultaneous efforts on COSI have greatly enabled the work in this thesis. I personally feel that we make an excellent team, and I hope that we will find ourselves working together again some day.
- Steve McBride for teaching me most of what I know about practical electronics, for his friendship, and for making our campaigns extremely interesting. And I mean that in the absolute best way possible.
- Brent Mochizuki for his good humor, electronics talent, and his often elegant solutions to tough problems. I look forward to working more with him on GAPS and beyond.

- My colleague Dr. Alan Chiu for his charm and dedication, evidenced by his frequent use of the cot. Thanks also for taking so many useful photographs of our project.
- Jane Hoberman for her meticulous attention to detail, and always insisting in the absolute highest level of quality for the COSI electronics.
- Dr. Mark Amman for fabricating and testing the germanium detectors central to our project. His willingness to work closely with the end users of his technology is admirable.
- Our collaborators in Taiwan Professor Hsiang-Kuang Chang, Chien-Ying Yang, Chao-Hsiung Tseng, Che-Yen Chu, and Chih-Hsun Lin for their extensive work on the COSI electronics, and their help with data analysis and calibrations in the field.
- Our collaborators in France Professors Pierre Jean and Peter von Ballmoos, and engineer Vincent Waegbaert for their design and construction of the COSI anti-coincidence shield, and their help with COSI data analysis. It was a pleasure collaborating with Pierre on various shield related projects. Special thanks to Peter for a highly memorable sled ride in Mürren.
- The technical staff at SSL that played a role in COSI development: Yvette Irwin, Helen Yuan, Bruce Dalen, Chris Cork, Tom Clemons, Paul Turin, Chris Smith, Jerry Olson, Bill Donakowski, Steve Marker, Ron Phillips, Micah Chong, Terry McDonald, and Dean Colomb. Many thanks for your contributions to COSI, it could not have been done without you.
- The staff at Columbia Scientific Balloon Facility for facilitating the COSI launch, flight, landing, and recovery. A special thanks to Bryan Stilwell, Juan Perez, Chris Field, Joseph Jones, Jill Juneau, Dwayne Orr, Hugo Franco, David Sullivan, and the team of riggers.
- Deborah Fairbrother and the Balloon Program Office at NASA Wallops for giving us the opportunity to fly on the Super Pressure Balloon.
- Previous graduate students Eric Bellm, Mark Bandstra, and Jason Bowen. Their theses and papers were incredibly useful during calibration and data analysis efforts. I hope my thesis will be as useful as theirs for future COSI graduate students.
- Will Rachelson for his technical advice over the years, his friendship, and for bailing me out of sticky auto repair situations.
- My undergrad advisor Professor David Smith, who introduced me to gamma-ray detection. His generous reference resulted in my employment at SSL, and subsequent graduate school career. I have many fond memories of working with him and chasing thunderstorms all over Florida.

- My family and friends, for their continual support and love. A very special thanks to my beloved parents, who sparked my interest in science and engineering at a young age.

Chapter 1

Gamma Ray Bursts

Gamma-ray bursts (GRBs) are bright, transient events of cosmological origin. They are aptly named to reflect the fact that the first phase — the prompt emission — is dominated by soft gamma-rays, typically between 0.1 MeV to 1 MeV. Following the prompt emission is an “afterglow” phase which is marked by broadband emission from X-rays down to radio, and can persist for months. Typical GRBs release a total energy in gamma-rays on the order of 10^{51} erg over the course of seconds, and can outshine the entire gamma-ray universe combined as they reach their peak intensities. At the time of this writing, thousands of GRBs have been observed, and hundreds have afterglow measurements. Despite the breadth of observations, many aspects of the GRB phenomenon remain ambiguous. Part of the problem lies in the broad diversity of temporal and spectral characteristics that are seen across the GRB population. This complicates theoretical modeling, as the proposed models must somehow unify the (often conflicting) observational evidence.

Polarization measurements of GRB prompt emission have the potential to discern between competing models. Differences between models include the presence or lack of magnetic fields within the GRB outflow, the coherency scale of magnetic fields, the dominant emission mechanism, and the geometry of the outflow. The polarization of the prompt emission is sensitive to all of these details. To date, the sample size of GRBs with polarization measurements is 23 (McConnell, 2017; Chattopadhyay et al., 2017; Lowell et al., 2017a). Many of these measurements claim a detection of non-zero polarization, but none of these detections exceed a significance level of 4σ . Recent advances in instrumentation and methods will push the sensitivity of upcoming GRB polarization measurements.

In this chapter, we start with a broad overview of GRB observations with major results along the way, extending from their discovery in the late 1960’s up to today. We will then describe the physical models behind the GRB phenomenon. Finally, we will discuss the implications of the theoretical models on the polarization of GRB prompt emission, as well as the current status of observations.

1.1 History of GRB Observations and Results

1.1.1 First Observations — Vela Spacecraft

GRBs were first discovered in the late 1960's (Klebesadel et al., 1973) serendipitously by the Vela satellites, which were deployed by the United States as a monitor for nuclear weapons testing in the wake of the Partial Test Ban Treaty. The initial report from Klebesadel et al. (1973) was with respect to measurements made with four Vela spacecraft; Vela 5A, 5B, 6A, and 6B. The Vela 5 and 6 series were equipped with cesium iodide scintillators, which provided for nearly full-sky sensitivity in the 0.3 MeV to 1 MeV band. The time delay for transient events detected on pairs of Vela spacecraft could be used to constrain the source location along an annulus on the sky. The intersection regions of the annuli were found to exclude both the Sun and Earth as sources of the bursts. Moreover, cosmic rays were ruled out on the basis that the observed time delays between Vela spacecraft never exceeded the maximum light transit time of ~ 0.8 seconds.

An early analysis of Vela burst locations by Strong & Klebesadel (1974) suggested that the majority of bursts were of galactic origin, originating from the local spiral arm of the galaxy, and that a minority could possibly be of extragalactic origin. However, the small sample size of 11 prevented firm conclusions from being made. Through the 1980's, several attempts were made to explain the origin of gamma-ray bursts using the observed spatial distribution. A firm answer to the question of the spatial distribution would not come until the early 1990's.

It was realized early on (Guilbert et al., 1983) that the GRB phenomenon may be highly relativistic; From a simple causality argument, the size scale of the emitting region can be no larger than $c\delta t \sim$, where δt s is the observed variability timescale of the GRB. The opacity to pair production for photons above 1 MeV, assuming cosmological distances to GRBs as well as the observed variability timescale ($\delta t \sim 10$ ms), would be on the order of $\tau \sim 10^{15}$ (Piran, 2005). Such a high opacity would prevent photons above 1 MeV from escaping, which is at odds with the observation of > 1 MeV photons from GRBs. The resolution to this problem is to assume that the source is moving relativistically towards the observer with a bulk Lorentz factor Γ of $\Gamma \sim 10^2 - 10^3$. The opacity is then reduced to a value less than unity if both length contraction and Doppler shifting of the emitted photons are taken into account.

1.1.2 Compton Gamma-Ray Observatory

In 1991, the Compton Gamma-Ray Observatory (CGRO) was carried into orbit by space shuttle Atlantis. CGRO utilized various instruments to make observations of astrophysical sources of gamma-rays across six decades in energy: the Burst and Transient Source Experiment (BATSE, 0.02 MeV to 2 MeV), the Oriented Scintillator Spectrometer Experiment (OSSE, 0.05 MeV to 10 MeV), the Imaging Compton Telescope (COMPTEL, 0.8 MeV to 30 MeV), and the Energetic Gamma-Ray Experiment (EGRET, 20 MeV to 30 GeV).




Figure 1.1: Spatial distribution of 1637 BATSE bursts from the fourth BATSE catalog (Figure 5 of Paciesas et al., 1999), showing the isotropic distribution of GRBs in galactic coordinates.

Of these experiments, BATSE was designed specifically to study the spatial distribution, as well as spectral and temporal characteristics of GRBs and other gamma-ray transients. BATSE’s eight sodium iodide detector modules were capable of performing spectroscopy and timing analysis for transient events. Localization of bursts was achieved by comparing the intensities between detector modules, since the acceptance of each module was strongest for a distinct region of the sky. Using this technique, unambiguous localizations with an accuracy of $< 10^\circ$ could be made for most bursts, and $\sim 1^\circ$ for the brightest bursts (Paciesas et al., 1999). Meegan et al. (1991) first reported that the spatial distribution of 153 BATSE GRBs was isotropic, indicating that the bursts likely are not of galactic origin. The log N - log F distribution of bursts revealed a deficit of weak bursts, which could be explained either by a galactic halo origin, or cosmological effects such as the redshift of distant bursts pushing the burst flux out of the BATSE bandpass (Fenimore et al., 1993). The fourth BATSE GRB catalog (Paciesas et al., 1999) reinforces these earlier findings (Figure 1.1).

Spectral studies of GRB prompt emission were significantly advanced by the capabilities of the BATSE detector modules. Band et al. (1993) found that the majority of BATSE GRBs are well described by a smoothly broken power law model — frequently referred to as a “Band model” — which is parameterized by a peak energy E_p , the power law slope α below E_p , the power law slope β above E_p , and an overall normalization constant. The spectrum is non-thermal and rather sharply peaked in νF_ν space, where the peak energy corresponds to the break energy E_p of the Band model. While there is no physical motivation for this spectral shape, its ability to reproduce the shape of many GRBs is very useful. Further




Figure 1.2: Histograms of Band model fit parameters to a sample of 5500 BATSE GRB spectra (Figures 7, 8, and 9 of Preece et al., 2000). The break energies E_p are clustered around ~ 250 keV, while the average low energy power law index α is ~ -1 , and the high energy power law index β tends to fall between -2 and -2.5.

work on BATSE spectral analysis has assumed variants on the Band model (Goldstein et al., 2013). Figure 1.2 shows the distribution of α , β , and E_p from Band model fits to 5500 BATSE GRBs (Preece et al., 2000). The values of E_p are clustered around ~ 250 keV, while the values of α and β are clustered around -1 and about -2.3 respectively. Efforts to explain the physical origin of the spectral shape have focused largely on the synchrotron shock model (Rees & Mészáros, 1992; Tavani, 1996), where the peak energy E_p is associated with the synchrotron characteristic frequency from radiating, non-thermal electrons in relativistic shocks. While the observed β values are consistent with plausible non-thermal electron distributions and magnetic field strengths, the observed α values are found to violate the so-called “synchrotron line of death” in about 20% of BATSE GRBs (Preece et al., 1998a). The origin of the problem is that the low energy tail of the synchrotron spectrum cannot have a slope harder than -2/3 (in photon number spectra) in the fast cooling regime. However, this limitation can be overcome with tuning certain parameters such as the electron energy distribution and the pitch angle distribution (Lloyd-Ronning & Petrosian, 2002).

Time domain analyses of BATSE bursts revealed that the distribution of GRB time durations is bimodal, with “long” GRBs lasting longer than ~ 2 seconds, and “short” GRBs lasting shorter than ~ 2 seconds (Kouveliotou et al., 1993). The distribution is plotted in Figure 1.3 for 2041 BATSE GRBs. Furthermore, it was discovered that short GRBs tend to have a harder spectrum, whereas the long GRBs tend to have a softer spectrum. This formed the basis of a simple classification scheme — short-hard, long-soft — which was taken as an indication that the population of GRBs is divided amongst two distinct classes of progenitors.

In addition to the BATSE results, EGRET became the first instrument to detect GRBs at high energies, proving that the GRB phenomenon is capable of producing > 100 MeV




Figure 1.3: T_{90} distribution for 2041 BATSE GRBs, along with Gaussian fits to the short and long GRB populations (Figure 1 from Lazzati & Perna, 2009). T_{90} is defined as the time over which the detector collects between 5% and 95% of the total fluence.

gamma-rays (GRB 910503, Schneid et al., 1992), and then > 1 GeV gamma-rays (GRB 930131, Sommer et al., 1994). In the case of GRB 930131, a power law fit described the data well over three decades in energy, which thoroughly ruled out a thermal spectrum with a temperature less than 1 GeV. By the end of the CGRO mission, 5 GRBs with emission above 30 MeV had been detected by EGRET (Dingus et al., 1997).

1.1.3 First Afterglows and Supernovae Associations in the BeppoSAX and HETE-2 Era

While the BATSE spatial distribution and $\log N - \log f$ distribution gave strong evidence that GRBs occur at cosmological distances, it was still not possible to rule out an extended galactic halo origin (Klose, 1995), which could also reproduce the observed results. GRB counterparts at other wavelengths had long been sought in order to confirm possible associations between GRBs and host galaxies, although no convincing evidence had been found. This changed when the X-ray satellite BeppoSAX detected GRB 970228 with both its Gamma-Ray Burst

Monitor (GRBM, 40 keV to 700 keV) and one of its Wide Field Cameras (WFC, 2 keV to 30 keV) (Costa et al., 1997). The WFC detection and subsequent detection with its Medium Energy Concentrator Spectrometer instrument represented the first GRB afterglow detected in any wave band. BeppoSAX continued to detect the afterglow out to seven days post-burst, and found that the intensity in the 2 keV to 10 keV range decayed smoothly with a power law shape. A few months later, GRB 970508 and its X-ray afterglow were detected by BeppoSAX, and rapid follow-up observations in the optical band confirmed a fading optical counterpart (Heise et al., 1997; Galama et al., 1998a). Deeper, spectroscopic observations in the optical band of GRB 970508 revealed absorption features, which were used to infer a redshift of $z = 0.835$, thus confirming that at least some GRBs are extragalactic in origin (Metzger et al., 1997; Bloom et al., 1998).

Another major discovery in this era was the association between some GRBs and Type 1b/c supernovae. The first indication of this came with the long GRB 980425, whose BeppoSAX error region was found to coincide with the type Ic supernova SN1998bw (Galama et al., 1998b). The chance probability of such a coincidence was estimated to be $< 9 \times 10^{-5}$, but was based on a posteriori statistics. Moreover, the optical emission was not accompanied by a fading afterglow component, and so the connection between GRB 980425 and SN1998bw remained controversial. Stronger evidence of a link between GRBs and supernovae came following the launch of the High Energy Transient Explorer 2 (HETE-2) satellite in 2000, whose primary goal was to localize X-ray afterglows to 10 arcseconds and notify ground observers of a burst within minutes. HETE-2 detected the long GRB 030329, which was promptly followed up in the optical band, resulting in a robust detection of a fading optical afterglow (Price et al., 2003). After 7 days, the optical light curve began showing features resembling a supernova light curve, with the optical afterglow superimposed. Subtraction of the afterglow component revealed a spectral evolution very similar to that of SN1998bw, making it clear that a connection does indeed exist between GRBs and supernovae.

1.1.4 Rapid Afterglow Identification with the Swift Gamma-Ray Burst Mission

By the early 2000s, several GRBs had optical counterparts with robust redshift measurements. However, the limiting factor in making these observations was the localization accuracy of the X-ray instruments initially reporting the afterglow position, and also the long delay time for optical follow up. The Swift satellite — a dedicated gamma-ray burst mission — was launched in 2004 with the goal of autonomously detecting GRBs and rapidly providing positions for X-ray afterglows. Swift constantly observes a large portion of the sky with the Burst Alert Telescope (BAT, 15 keV to 150 keV), which is a coded mask imager utilizing a cadmium zinc telluride detector plane. When a transient such as a GRB is detected by the BAT, the spacecraft autonomously re-points itself and begins to tile the BAT error circle region with the X-ray Telescope (XRT, 0.2 keV to 10 keV) and the UV-Optical telescope (UVOT). The XRT is capable of localizing the X-ray afterglow to an accuracy of

5 arcseconds, which enhances the probability of optical detection with UVOT and ground based telescopes. Swift transmits X-ray coordinates to the ground within 1 to 2 minutes, which are then relayed over the Gamma-ray Coordinates Network (GCN). The GCN alert allows optical and radio observers on the ground to quickly point telescopes at the position reported by Swift in an effort to detect the afterglow.

Swift's efficient detection of GRBs with a high likelihood of subsequent afterglow identification marked a new era in GRB studies. Spectral, temporal, and even polarimetric measurements of afterglows localized by Swift yield valuable information that is not available in the prompt emission. Observations of the optical afterglow spectra have resulted in redshift measurements for more than > 360 GRBs ranging from $z = 0.03$ to $z = 9.38$ (Lien et al., 2016), thus eliminating any doubt regarding the cosmological origin of GRBs. Follow ups of the optical afterglows have yielded GRB supernovae connections for 46 GRBs to date Cano et al. (2017), either by direct spectroscopic analysis of the supernova ejecta, or by identification of a late time bump, interpreted as the supernova light curve becoming dominant over the GRB optical afterglow (Woosley & Bloom, 2006). Importantly, observations of well localized short GRBs have yet to reveal supernovae associations (Hjorth et al., 2005; Bloom et al., 2006), reinforcing the notion that the progenitor for short GRBs differs from that of long GRBs. Analyses of host galaxies has revealed another clue regarding short GRB progenitors: long GRBs are associated with star-forming galaxies, while short GRBs are associated with either low star-forming elliptical galaxies, or regions of galaxies where the star formation is low (Schady, 2017). If long GRBs are due to the collapse of massive stars, then it is not surprising that they occur in regions where the star formation rate is high. The link between short GRBs and low star forming rates is consistent with the hypothesis that short GRBs are a byproduct of either a binary neutron star (NS) merger, or a black hole (BH) NS merger, as these types of systems are associated with old stellar populations. Compact object merger scenarios for short GRBs predate the Swift mission (Eichler et al., 1989; Narayan et al., 1992), and remain the most popular explanation. Very recently, a gravitational wave event consistent with a NS-NS merger was detected by both the Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo interferometers (Abbott et al., 2017). A faint short GRB was detected by Fermi/GBM 1.7 seconds after the gravitational wave signal (Goldstein et al., 2017), marking the first association between short GRBs and merger events.

Swift observations of the X-ray afterglows show detailed behavior in the time domain that appears to be relatively common across many GRBs. A selection of X-ray afterglows are shown in Figure 1.4. The canonical X-ray afterglow light curve (Gehrels et al., 2009) starts with an initial steep decay lasting typically $10^1 - 10^2$ seconds, which then transitions in to a shallow decay phase. The interpretation is that the steep decay represents the tail end of the prompt emission, and the transition to the shallow phase happens when the afterglow emission begins to dominate over the fading prompt emission. This is further corroborated by an observed change in spectral index between the steep and shallow phases, which would be expected if the afterglow and prompt emission originate from different regions. After $10^2 - 10^3$ seconds post burst, the light curve steepens again as it enters the classical afterglow

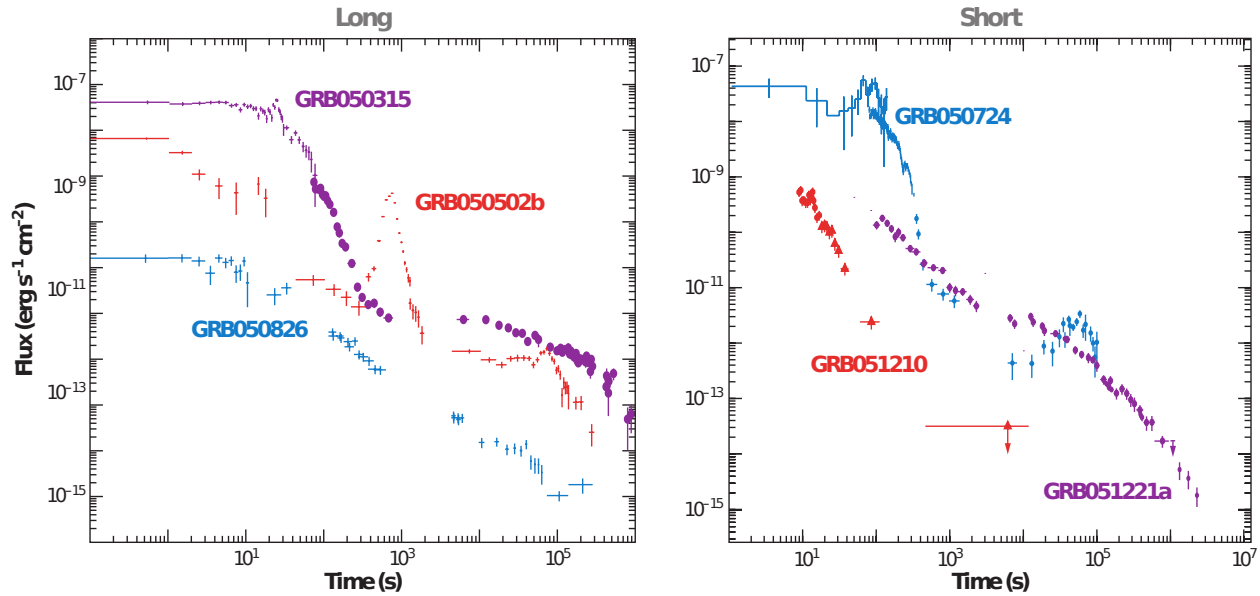


Figure 1.4: A selection of X-ray afterglows for long GRBs (left), and short GRBs (right) as seen by Swift (Figure 6 from Gehrels et al., 2009). A variety of different behaviors can be seen including steep to shallow transitions, jet breaks, and late time flares.

phase. Generally, the afterglow is believed to be powered by synchrotron emission from shock generated magnetic fields as the relativistic outflow plows into the circumburst medium. Occasionally, a “jet-break” is seen at later times ($10^5 - 10^6$ seconds post burst), which has been interpreted as the reduction of relativistic beaming as the outflow decelerates, as well as lateral expansion of the outflow as the jet becomes causally connected (Piran, 2005). Bright X-ray flares are seen during the shallow or classical afterglow phase in roughly half of the observations (Gehrels & Razzaque, 2013), which indicates that the central engine powering the GRB reactivates at later times. Such temporal behavior places strong constraints on the nature of the central engine, as it must have the ability to repeatedly tap a large source of energy.

1.1.5 High Energy Observations with the Fermi Space Telescope

The Fermi gamma-ray space telescope was launched in 2008 and is equipped with two instruments which have been useful for studying GRBs: The Gamma-Ray Burst Monitor (GBM), and the Large Area Telescope (LAT). The GBM is a distributed array of NaI and bismuth germanate (BGO) scintillators sensitive between 8 keV and 40 MeV, and similar in spirit to the BATSE detectors. With its large field of view of 70% of the sky, GBM detects ~ 300 GRBs per year, and is capable of performing spectroscopy of bursts over a much larger energy range than BATSE. Overall, the GBM results reflect the BATSE results closely, with small differences between catalogs attributed to the extended energy of GBM, as well as

instrumental or methodological biases (Goldstein et al., 2012).

The LAT is a pair conversion telescope sensitive to high energy gamma-rays in the range 20 MeV to 300 GeV. The LAT energy range overlaps the upper end of the GBM energy range, providing for unprecedented spectral coverage for GRB observations (8 keV to 300 GeV). The LAT field of view is considerably smaller than that of the GBM, resulting in approximately 10 GRBs detected per year. While previous EGRET results had confirmed that GRB emission extends up to 1 GeV, Fermi observations of GRB 130427A for example prove that the emission can extend up to ~ 95 GeV (Ackermann et al., 2014). Perhaps the most important results from LAT GRBs is that the high energy emission during the prompt phase is delayed and lasts longer than the 100 keV - 1 MeV prompt emission (Gehrels & Razzaque, 2013). GRB 130427A is an example case of both of these characteristics, where > 1 GeV photons were only detected starting 10 seconds after the GBM trigger, and a 32 GeV photon was detected 9 hours post trigger. With regards to the spectral shape of the high energy emission, the Band model fits well for many GRBs over the combined GBM + LAT bandpass, but some GRBs show extended emission all the way up to 10 GeV, which often requires an additional power law component in the fits (Ackermann et al., 2011, 2013).

1.2 Progenitors and Central Engines

The observational evidence for the association of long duration GRBs with the collapse of massive stars is strong, due to the 46 GRBs for which SN have been spectroscopically identified, or for which late time bumps have been observed (Cano et al., 2017). The picture is less clear for short GRBs, although the recently observed association between a NS-NS merger and a short GRB is strong evidence in favor of the merger model (Abbott et al., 2017). Despite this dichotomy, it is believed that the central engine responsible for powering the GRB explosion is the same, regardless of the progenitor.

In the “collapsar” model for long GRBs (Woosley, 1993; MacFadyen & Woosley, 1999), the massive stellar core gravitationally collapses, resulting in a hyper-accreting black hole (BH) or possibly a highly magnetized neutron star (NS) — a “magnetar” (Usov, 1992). In short GRBs, a compact object binary system — NS-NS or NS-BH — merges due to the loss of orbital energy to gravitational waves, forming a hyper-accreting BH. In the case of a BH central engine and assuming high mass accretion rates, an optically thick, compact disk is formed around the BH, which can convert gravitational energy into a relativistic outflow in at least two ways. First, the disk may be in the neutrino dominated accretion flow (NDAF) regime, which results in rapid cooling via the emission of neutrinos and anti-neutrinos (Asano & Fukuyama, 2000; Narayan et al., 2001). Neutrino annihilation creates a hot pair plasma outflow, giving rise to the GRB. A second scheme for energy conversion is the Blandford-Znajek mechanism (Blandford & Znajek, 1977; Lee et al., 2000) in which the spin energy of a rapidly rotating BH is tapped by an external magnetic field and carried away by a Poynting flux. In either case, the central engine variability in the hyper accreting BH scenarios is usually attributed to episodic accretion, possibly due to turbulence in the

accretion disk caused by magnetorotational instabilities (Masada et al., 2007).

In the case that a rapidly rotating magnetar is formed instead of a BH, the energy release process is different. Usov (1992) showed that the total rotational kinetic energy and energy release timescale of such an object would be $\sim 5 \times 10^{52}$ erg and ~ 10 seconds respectively, consistent with long GRB energetics and durations. Initially, neutrino heating of the NS drives a baryon rich wind from the NS surface, which is incapable of attaining relativistic velocities due to the high baryon mass. As the NS cools and the baryonic wind dies down over the course of several seconds, the outflow becomes magnetically dominated, and magnetic instabilities dissipate the energy (Metzger et al., 2011; Ruderman et al., 2000; Thompson et al., 2004).

1.3 Origins of the Prompt Emission

Measured fluences of the prompt emission phase range from $10^{-7} - 10^{-3}$ erg cm^{-2} . For typical cosmological distances to GRBs, these fluences translate to an isotropic energy release of $E_{\text{iso}} \sim 10^{53}$ erg. However, GRB outflows are likely beamed into a solid angle Ω (Piran et al., 1999; Perna et al., 2003), which reduces E_{iso} by a factor of $\Omega/4\pi \sim 0.001$ for jet opening angles of around 0.1 radians (Frail et al., 2001). Even after accounting for beaming, GRBs are the most violent known energy releases in the universe.

1.3.1 Prompt Emission in the “Fireball” Model

In the “fireball” model (Meszaros & Rees, 1993), the central engine taps a reservoir of gravitational energy and generates a hot fireball consisting of photons, electrons, positrons, and a relatively small number of baryons. Under immense thermal pressure, the fireball rapidly expands, accelerating particles up to relativistic speeds until a coasting Lorentz factor of $\Gamma = 10^2 - 10^3$ is achieved. The kinetic energies of the now relativistic particles are then dissipated by some physical mechanism, which results in the GRB prompt emission.

Typically, synchrotron emission is invoked as the dissipation mechanism, where the magnetic fields facilitating this process are either advected by the central engine, or produced in shocks via the Weibel instability (Medvedev & Loeb, 1999; Silva et al., 2003). In order to effectively radiate, the particles must have a high velocity relative to the comoving magnetic field. In the “internal shocks” model (Rees & Meszaros, 1994), variable activity from the central engine launches shells with variable Lorentz factors. As faster shells collide with slower shells, shocks are formed, and electrons and positrons are accelerated via the first order Fermi mechanism within the shock-generated magnetic fields (Spitkovsky, 2008). This results in a non-thermal distribution of electrons and positrons which then emit optically thin synchrotron radiation (Pilla & Loeb, 1998).

While synchrotron cooling of Fermi accelerated particles is a popular choice for the emission mechanism in the fireball model, it suffers from a low radiative efficiency (Kumar, 1999), as well as disagreement with the observed spectrum (Preece et al., 1998b). Several model

variants were suggested to fix the efficiency problem (Beloborodov, 2000; Kobayashi & Sari, 2001), which rely on uncomfortably large differences in Lorentz factors between the colliding shells. To address the spectral discrepancy, (Panaitescu & Mészáros, 2000) suggested that the observed low energy portion of the spectrum could be hardened if synchrotron self Compton plays a significant role (Panaitescu & Mészáros, 2000).

In lieu of synchrotron emission, several other possibilities have been suggested. Thermal radiation from an optically thick fireball photosphere has been suggested as either the primary prompt emission mechanism (Eichler & Levinson, 2000; Beloborodov & Mészáros, 2017), or as an additional component superposed with e.g. synchrotron emission (Meszaros & Rees, 2000; Peer & Ryde, 2016). These models are bolstered by a subset of GRBs whose prompt emission spectra are better described by a quasi-thermal model (Ryde & Pe’er, 2009). Lastly, the Compton drag model invokes bulk Comptonization of ambient soft photons from the relativistic particles as the principle dissipation mechanism (Lazzati et al., 1999a). Compton drag cures the inefficiency problem with synchrotron in internal shocks, but in order to achieve an emergent spectrum resembling a Band function, constraints must be placed on the ambient photon spectral shape (Ghisellini et al., 2000; Lazzati et al., 1999b).

1.3.2 Prompt Emission in Electromagnetic Models

A key assumption in the fireball model is that magnetic fields do not play a dominant role in the GRB jet dynamics. A competing set of models exist, where the internal energy of the outflow is not dominated by the kinetic energy of particles, but rather in the Poynting flux of an advected magnetic field. In these models, electrons are accelerated in the comoving frame either by current-driven instabilities (Lyutikov & Blandford, 2003), or by magnetic reconnection processes (Spruit et al., 2001). In either case, synchrotron emission is readily emitted in the presence of the presumed magnetic fields that permeate the outflow.

An interesting scenario which combines elements of the electromagnetic model with the internal shocks of the fireball model is the Internal Collision induced Magnetic Reconnection and Turbulence (ICMART) model (Zhang & Yan, 2010). Within the ICMART framework, the central engine launches a magnetically dominated outflow, consisting of various shells with different Lorentz factors, similar to the internal shocks case. Each shell starts out with a globally ordered magnetic field, but as shells with higher Lorentz factors collide with other shells, the magnetic fields become more and more tangled. Eventually, the tangling becomes severe and significant magnetic reconnection occurs, resulting in a relativistic population of charged particles which then emit synchrotron radiation.

1.4 Polarization of the Prompt Emission

Polarimetric measurements add two more parameters to the space of observables; the polarization level Π and the polarization angle η_0 . Polarization measurements of the GRB prompt emission are currently not commonplace, but progress has been made in the last 15 years.

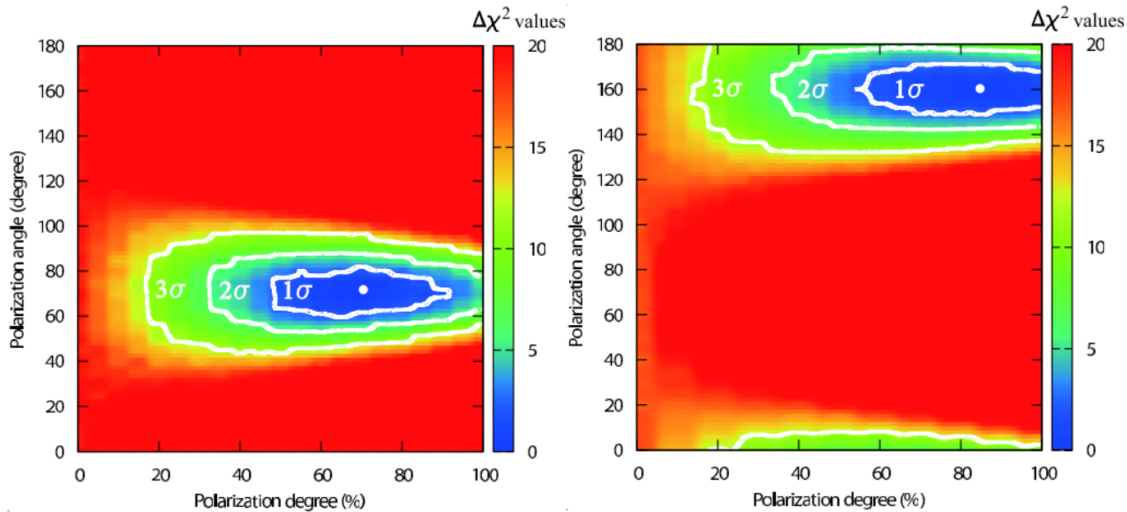


Figure 1.5: $\Delta\chi^2$ maps for the IKAROS/GAP observations of GRB 110301A (left) and GRB 110721A (right) indicating detections of non-zero polarization levels at significance levels of 3.7σ and 3.3σ respectively (Figure 3 from Yonetoku et al., 2012). These results, in concert with the observed polarization angle changes in GRB 100826A (Yonetoku et al., 2011a) were interpreted as evidence in support of large scale magnetic field structures and synchrotron emission as the dominant GRB emission mechanism.

The first reported GRB polarization measurement occurred for GRB 021206, which was made using RHESSI (Coburn & Boggs, 2003). A high polarization level of $\Pi = 80\% \pm 20\%$ was measured for this burst, with a significance level of $> 5.7\sigma$. However, subsequent re-analyses by several groups were unable to confirm these findings, so the initial claims remain controversial (Rutledge & Fox, 2004; Wigger et al., 2004). Nonetheless, this high level of polarization spurred much theoretical interest in explaining the observations, and underscored the potential impact such measurements may have on the field.

Since then, GRB polarization measurements have been made using INTEGRAL/SPI, INTEGRAL/IBIS, IKAROS/GAP, Astrosat/CZTI, and COSI (for a recent review, see McConnell, 2017). All of the aforementioned instruments operated as Compton polarimeters to infer the polarization level and angle of GRBs. Additionally, the BATSE Albedo Polarimetry System (BAPS) attempted to measure GRB polarization by imaging the backscattered flux from the top of Earth’s atmosphere, which should be azimuthally modulated if the GRB were polarized (McConnell et al., 1996; Willis et al., 2005). To date, there are no robust, high significance measurements of GRB polarization (McConnell, 2017). If the GRB 021206 result is excluded (due to its controversial nature), the most significant results to date came from the Gamma-ray Burst Polarimeter (GAP), a dedicated GRB polarimeter on the JAXA IKAROS solar sail (Yonetoku et al., 2011b). The GAP reported $\Pi = 70\% \pm 22\%$ for GRB 110301A and $\Pi = 84\%_{-28\%}^{+16\%}$ for GRB 110721A, at significance levels of 3.7σ and 3.3σ respectively. A $\Delta\chi^2$ map for these two GRBs is shown in Figure 1.5. Furthermore, a change in the

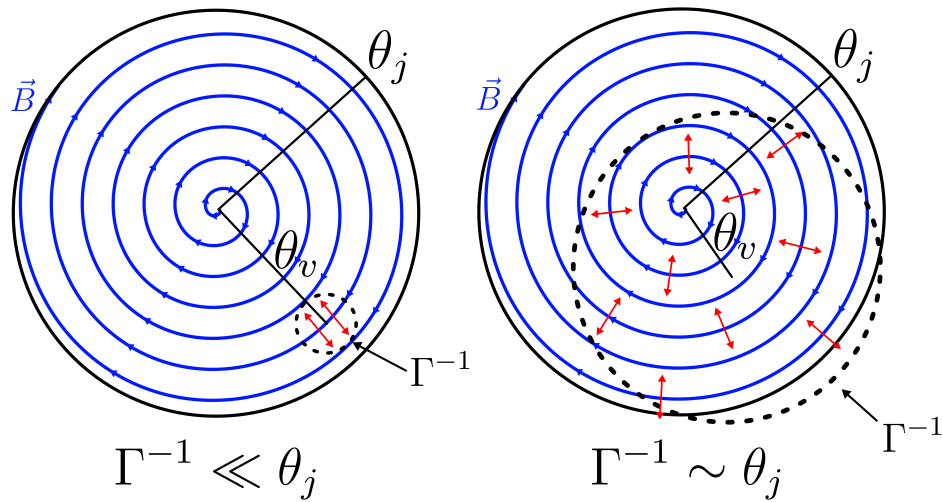


Figure 1.6: A cartoon illustrating the synchrotron emission from oriented magnetic fields (SO) model. On the left, the visible region of the jet Γ^{-1} centered along the viewing angle θ_v is contained within the jet opening angle θ_j . The magnetic field lines (blue) are assumed to be helical and result from a Poynting flux from the central engine. Within the visible region, the field lines are rather straight, the corresponding polarization vector directions (red arrows) are well defined, and thus a large polarization level is observed ($\Pi \sim 40\%$). On the right, lower values of Γ result in a larger visible emitting region. Polarization vectors with different orientations are now visible from other regions of the jet, and thus the overall polarization level is significantly reduced. The diagrams are adapted from Toma (2013).

polarization angle of GRB 100826A between two intervals was identified by the GAP with a significance of 3.5σ (Yonetoku et al., 2011a). Most recently, the Astrosat team has reported GRB polarization measurements for 11 GRBs, four of which are consistent with non-zero polarization levels at significances ranging from 3σ to less than 4σ (Chattopadhyay et al., 2017). Finally, the POLAR instrument aboard the Chinese space station Tiangong 2 is another dedicated GRB polarimeter which was only recently deployed (Sun et al., 2016; Orsi et al., 2014; Produit et al., 2005). POLAR has successfully detected several GRBs in its few months of service (Kole & Marcinkowski, 2016; Kole et al., 2016), although no polarization analyses have yet been published.

1.4.1 Polarization from Synchrotron Emission

Central to many models for the prompt emission is the existence of a magnetic field, which can influence the dynamical evolution of the GRB jet, provide an environment for particles to produce synchrotron radiation, or both. Naturally, the question of polarization arises, since synchrotron emission from a coherent magnetic field is capable of producing a high level of polarization. A related question is that of the structure and dynamical behavior of

the magnetic fields where the synchrotron emission is generated. If a non-thermal population of electrons with kinetic energies distributed as E^{-p} exists in the presence of a large-scale, coherent magnetic field, the emitted synchrotron radiation will be polarized at a level of $\Pi = (p + 1)/(p + 7/3)$, assuming an isotropic pitch angle distribution (Rybicki & Lightman, 2008; Lei et al., 1997). However, inhomogeneities in the magnetic field, and distortions due to relativistic aberration (Lazzati, 2006), can degrade the observed polarization level.

Following Toma et al. (2009), the synchrotron in ordered magnetic fields (SO) model can produce a high level of polarization for a wide range of viewing geometries. A diagram illustrating the SO model is shown in Figure 1.6. In the SO model, it is assumed that a large scale toroidal magnetic field permeates the GRB jet cone, whose opening angle is θ_j . As a result of the high Lorentz factor of the outflow, an observer whose viewing angle θ_v falls within the jet cone only sees emission from within an angle Γ^{-1} due to relativistic beaming. For large values of Γ , the visible region is small (left side of Figure 1.6) and the magnetic field lines are relatively straight. In this situation, a high polarization level is observed ($\Pi \sim 40\%$), where the polarization vector is perpendicular to the local magnetic field lines. If Γ is small enough such that $\Gamma^{-1} \sim \theta_j$ (right side of Figure 1.6), then the observer will see emission from a larger region, and the polarization vectors will cancel out, leading to lower polarization levels. However, this can be counteracted if both $\Gamma^{-1} \sim \theta_j$ and $\theta_v \sim \theta_j + \Gamma^{-1}$, such that the viewing angle falls $\sim \Gamma^{-1}$ outside of the jet cone. In this case, only a small portion of the emitting region is seen. The polarization will be high ($\Pi \sim 60\%$), but the intensity of the emission will be low.

In the synchrotron in random magnetic fields model (SR), shock-generated magnetic fields (such as those in the internal shocks scenario) give rise to synchrotron emission in the presence of Fermi accelerated electrons. A diagram illustrating the SR model is shown in Figure 1.7. Shock generated fields are expected to be random on the scale of the plasma skin depth, and contained within the shock itself (Medvedev & Loeb, 1999). Therefore, the magnetic field lines are parallel to the shock plane. In the emitting electron's comoving frame, an observer at 90° sees maximally polarized emission with a polarization vector perpendicular to the shock plane. After Lorentz boosting into the lab frame, the emission from the electron is beamed, and the emission at 90° in the comoving frame is now emitted at an angle Γ^{-1} in the observer frame, still with a high level of polarization. Emission from angles smaller than Γ^{-1} will still be polarized, albeit at a lower level. Nonetheless, when $\theta_v < \theta_j$ (left side of Figure 1.7), the emission is axisymmetric around the line of sight, and the polarization cancels. If $\theta_v \sim \theta_j + \Gamma^{-1}$ (right side of Figure 1.7), such that the viewing angle falls $\sim \Gamma^{-1}$ outside of the jet cone, the emission is no longer axisymmetric, and a moderately high polarization level results ($\Pi \sim 25\% - 30\%$). In such a geometry, a high polarization level will be accompanied by a lower flux.

In between the SO and SR models is the ICMART model, which includes an explicit time dependence of the observed GRB polarization properties (Zhang & Yan, 2010). Within the ICMART framework, the GRB jet begins with a structured, toroidal magnetic field as in the SO model. Variable activity from the central engine results in internal collisions which gradually distort the magnetic field lines. Eventually, the field lines become so distorted,




Figure 1.7: A cartoon illustrating the synchrotron emission from random magnetic fields (SR) model. On the left, the visible region Γ^{-1} is contained within the jet opening angle θ_j . Within the visible region, the magnetic field orientations (blue arrows) are random, and the polarization vectors (red arrows) are axisymmetric around the line of sight, leading to no net polarization. On the right, the visible region coincides with the edge of the jet, leading to a loss of symmetry and a net polarization ($\Pi \sim 25\% - 30\%$). However, the situation on the right may be accompanied by a lower observed intensity, as the visible region is not filled with emission. The diagrams are adapted from Toma (2013).

that subsequent collisions trigger a runaway magnetic reconnection event which destroys the large scale magnetic field. This is called an ICMART event, and each pulse of a GRB is associated with one ICMART event. Thus, the degradation of the magnetic field coherency over time would result in a decrease of the polarization level with time from $\Pi \sim 50\% - 60\%$ to $\Pi \sim 0\%$ throughout a single GRB pulse. This prediction could be tested using time-resolved polarimetry for the brightest GRBs.

1.4.2 Polarization from Inverse Compton Scattering and Photospheric Emission

Inverse Compton scattering has been suggested as another emission mechanism that may be responsible for the prompt emission. In the Compton Drag (CD) scenario (Lazzati et al., 2004; Toma et al., 2009), ambient ultraviolet photons are inverse Compton scattered by relativistic electrons in an optically thin fireball. In the comoving frame of an emitting electron, blue shifted photons approach from the outflow direction and Compton scatter on the electron. Compton scattering induces a polarization such that the polarization vector is perpendicular to the plane of scattering, and photons with 90° Compton scattering angle are maximally polarized. When the emission is Lorentz boosted into the observer frame, the




Figure 1.8: A cartoon illustrating the Compton Drag model, which relies on inverse Compton emission to produce a high polarization level. Similar to the SR model, when the visible region Γ^{-1} is contained within the jet opening angle θ_j , the polarization vectors are axisymmetric around the line of sight, and no net polarization results. On the right, the visible region coincides with the jet edge, the symmetry is broken, and a high polarization results (up to $\Pi \sim 80\%$). Once again, the overall intensity of the prompt emission may be lower for the situation on the right, as the visible region contains less of the jet emission.

highly polarized emission at 90° in the comoving frame is now emitted at an angle Γ^{-1} in the observer frame due to relativistic beaming. Figure 1.8 shows a diagram of the CD scenario. When $\theta_v < \theta_j$, the polarization vectors are axisymmetric along the line of sight, and the polarization cancels. Similar to the SR model, when $\theta_v \sim \theta_j + \Gamma^{-1}$, only a small sliver of the emission region is seen, and the symmetry is lost. This can result in polarization levels as high as $\Pi \sim 80\% - 90\%$ for favorable geometric configurations. Once again, this situation will likely lead to lower intensities since the viewing angle does not fall within the GRB jet cone.

Photospheric emission models (Lundman et al., 2014) have gained popularity recently as a way to more naturally explain the observed Band model shape of GRB spectra. In these models, the observed spectrum consists mostly of photospheric emission from an optically thick, dissipative fireball. Below the photosphere, photons are up scattered by hot electrons, but subsequently down scattered due to the high optical depth. At the photosphere however, some up scattered photons escape along the GRB jet axis, contributing to the GRB luminosity. If small scale structures exist, or if the photosphere is viewed near the edge of the jet cone, then the inverse Compton scattered photons can have a high degree of polarization (Beloborodov & Mészáros, 2017). Additionally, these models predict synchrotron cooling within the fireball, assuming an advected magnetic field permeates the outflow (Lundman et al., 2016). The synchrotron emission is expected to dominate the low energy part of the

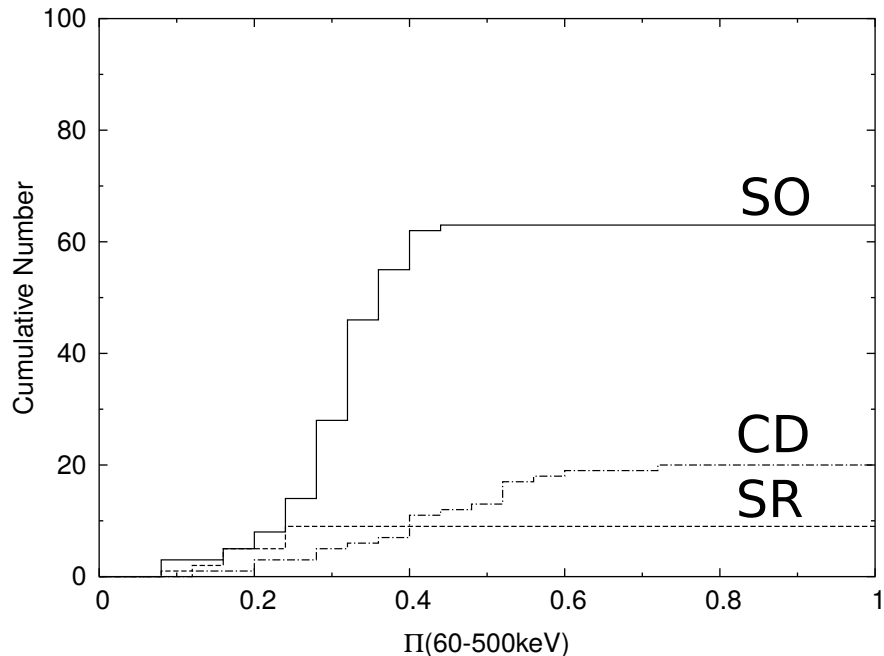


Figure 1.9: Simulated, cumulative distributions of polarization levels measured by the proposed POETS satellite (60 keV to 500 keV) for the SO, SR, and CD models over a period of time such that 200 bursts are detected. A net polarization level of $\Pi \sim 30\% - 40\%$ is frequently observed for the SO model. The CD model is capable of yielding higher polarization levels $\Pi > 60\%$, albeit less frequently. The SR model appears to be inefficient at producing a net polarization. Cumulative distributions for a real sample of GRBs may be compared to those shown here to help discern between models.

observed GRB spectrum, so that a net polarization level may be observed in this model at lower energies, even if the emission is axisymmetric.

1.4.3 Statistical Distribution of Polarization Levels

The observed polarization in the aforementioned scenarios depends on various parameters including spectral characteristics, temporal behavior, the bulk Lorentz factor, the jet opening angle, and the viewing angle with respect to the jet axis. Each GRB can have different values for these parameters, thus making it difficult draw to conclusions from a single measurement. Thus, it may be necessary to consider the statistical distribution of measured polarization levels across a large sample of bursts.

This type of statistical analysis was performed by Toma et al. (2009) as part of a concept study for the POETS satellite. Monte Carlo simulations of GRBs from the SO, SR, and CD models were generated and analyzed with the hypothetical POETS satellite. Figure 1.9 shows the cumulative distributions of polarization measurements for the three models, in the energy range 60 - 500 keV. Clearly, it is much more likely to observe a net polarization

($\Pi \sim 30 - 40\%$) in the SO model than in the SR, and CD model. While the CD model results more rarely in a net polarization, the observed polarization levels can be quite high ($\Pi > 60\%$). The SR model appears to be inefficient in producing a net polarization in the prompt emission.

1.5 Summary

Much progress has been made in understanding the GRB phenomenon since its serendipitous discovery in the 1960's. However, many of the details regarding the nature of the central engine, the emission mechanism, and the degree of magnetization of the outflow remain uncertain. It should be clear that polarization measurements can make a significant impact on our understanding of the GRB phenomenon. Summarizing the current status of the observations, a net polarization level is frequently reported in the literature albeit at low significance. Assuming that a net polarization is frequently detected, this would tend to favor the SO model, as it is efficient at producing a net polarization in many bursts. However, additional measurements with more sensitive polarimeters will help to elucidate the situation. Further increases in detector technologies and polarimetric methods will push the polarization sensitivities of new instruments. Increases in sensitivity are always welcomed, but certainly required for performing time-resolved or energy-resolved polarization analyses. Over time, the significances of GRB prompt emission polarization measurements have been slowly increasing, and it seems we are on the cusp of being able to perform decisive measurements.

1.6 Thesis Outline

In the present chapter, we have given an overview of GRB observations with an emphasis on the polarization of the prompt emission. The following is a brief outline of the remainder of this thesis.

In Chapter 2, we will introduce the Compton Spectrometer and Imager (COSI), a balloon-borne Compton telescope and polarimeter sensitive between 0.2 MeV and 5 MeV. We will cover the operating principles of Compton telescopes and Compton polarimeters, the COSI detector technology, readout electronics, and the overall functioning of the entire instrument including subsystems on the balloon gondola. We finish with a description of the data analysis pipeline for both measured and simulated data.

Chapter 3 will cover the 2016 COSI Super Pressure Balloon campaign from Wanaka, New Zealand. In summary, COSI flew for 46 days and successfully detected five sources of astrophysical gamma-rays: The Crab Nebula, Cygnus X-1, Centaurus A, positron annihilation emission from the galactic center, and the long duration gamma-ray burst GRB 160530A. We will give an overview of the flight, focusing on details such as instrument performance,

background environment, and telemetry performance. We will conclude with a description of each detected source, including the planned analysis for each one.

In Chapter 4, we will outline the fundamentals of Compton polarimetry, focusing on 1) a standard analysis method (SM) which fits a modulation curve to a binned distribution of azimuthal scattering angles and 2) a maximum likelihood method (MLM) which uses more information per photon and can improve the minimum detectable polarization level (MDP) by up to $\sim 21\%$. We will also discuss real world effects of non-ideal polarimeters, and how they impact the polarization performance.

Chapter 5 describes our efforts to validate COSI’s polarimetric performance using both partially polarized beams of gamma-rays and unpolarized beams. For the polarized data sets, we created partially polarized beams by Compton scattering 662 keV photons (Cs-137) from a sodium iodide scintillator triggering in coincidence with COSI. The outgoing, partially polarized beam was then analyzed by COSI. For the unpolarized data sets, we simply illuminated the detector with various sources at different locations within the field of view. While the unpolarized data sets do not result in a modulation in the azimuthal scattering angle distributions, they are still useful probes of the instrument systematics. We compare the measured data sets to simulations of the same experimental configurations. In almost all cases, we find no statistically significant differences between measurements and simulations. One of the polarized data sets does show statistically significant differences, but we argue that this is likely caused by an imperfect transcription of the experimental configuration into the mass model. Using the polarized data sets, we place upper limits on the systematic error on the polarization level measurements of $3\% - 4\%$.

Having verified that COSI performs well as a Compton polarimeter, we turn to our observation of GRB 160530A in Chapter 6, and apply both the SM and MLM polarization analysis techniques. Using Monte Carlo simulations of our observation, we calculate an MDP of $72.3\% \pm 0.8\%$ for the SM and $57.5\% \pm 0.8\%$ for the MLM, confirming that the MLM does indeed improve the MDP by $\sim 20\%$. However, the measured polarization level for each analysis method was below its respective MDP. Therefore, we claim a non-detection of polarization in the prompt emission of GRB 160530A. We used the MLM to place a 90% confidence upper limit on the polarization level of 46%. This measurement alone does not allow us to constrain models for the GRB prompt emission. However, GRBs with polarizations above this level have been reported, so our result will be useful in future meta-analyses that consider a large sample of GRB polarization measurements.

In Chapter 7, we will describe two avenues we have identified for improving the polarization sensitivity of Compton polarimeters. The first is an “N-site” maximum likelihood method, which extends the MLM from Chapter 4 to use all N Compton interactions of an event. To test the MDP improvement of the NMLM over the SM and MLM, we perform simulations with an ideal germanium detector. We find that the NMLM gives an MDP improvement of 32% at 200 keV, 43% at 400 keV, and 47% at 800 keV when compared to the SM. Additionally, we show that the NMLM gives more accurate estimates of the true polarization level and angle than both the SM and MLM. These improvements are attributed to the fact that the NMLM uses significantly more information per photon than the SM

and MLM. The second avenue we have identified for improving the polarization sensitivity amounts to using the tenuous atmosphere at balloon float altitudes as the scattering element of a Compton polarimeter. We carry out simulations using a mass model of the atmosphere and show that photons arriving at an altitude of 33 km display a modulation in their azimuthal arrival angles. We call this the Atmospheric Modulation Effect (AME). After optimization, the modulation of this signal is $\mu_{100} = 0.260 \pm 0.003$, which is high enough to exploit the AME as a tool for studying GRB polarization from within the atmosphere. Detection of the AME with either COSI's imaging capabilities, or possibly a specialized instrument flying on the same gondola as COSI, could be combined with maximum likelihood Compton polarimetry of unscattered photons to improve the overall polarization sensitivity.

Chapter 2

Overview of The Compton Spectrometer and Imager

The Compton Spectrometer and Imager (COSI) is a 0.2 MeV to 5 MeV gamma-ray Compton imager, spectrometer, and polarimeter designed to fly on an 18 million cubic feet Super Pressure Balloon (SPB) (Chiu et al., 2015; Kierans et al., 2017). COSI combines a wide field of view, 3D positioning, and high spectral resolution in order to address a number of science targets including but not limited to: polarization studies of GRBs and compact objects, gamma-ray line spectroscopy of galactic nucleosynthesis sites, and imaging of the 511 keV emission from positron annihilation at the galactic center. In this chapter, we will give a broad overview of the COSI program, with a focus on methods, instrumentation, and data analysis.

2.1 Compton Telescopes

Compton telescopes are a powerful technology for performing imaging, spectroscopic, and polarimetric studies in the 0.2 MeV to 10 MeV band. This is due to the fact that the cross-section for Compton scattering dominates over photoabsorption and pair-production in this energy range for most detector materials of interest. A Compton telescope works by exploiting the phenomenon of Compton scattering, whereby a photon undergoes an inelastic collision with an electron initially at rest. A simple momentum conservation argument results in the classic, kinematic Compton formula:

$$E' = \frac{E_0}{1 + \frac{E_0}{mc^2}(1 - \cos \theta)} \quad (2.1)$$

which relates the incident photon energy E_0 to the scattered photon energy E' and the Compton scattering angle, or polar scattering angle, θ .

A canonical Compton event in a Compton telescope consists of one or more Compton scatters, followed by photoelectric absorption of the scattered photon, all in the active detec-




Figure 2.1: A schematic drawing of a canonical Compton event. An incident photon with energy E_0 (purple arrow) Compton scatters at r_1 with Compton scattering angle θ and azimuthal scattering angle η , depositing an energy E_1 . The scattered photon is subsequently photoabsorbed at r_2 , depositing an energy of E_2 . By measuring the energies and locations of the interactions, the origin of the photon can be constrained to a Compton cone (blue). By measuring η for a large sample of events, the polarization level and angle of the incident beam may be inferred. The COSI germanium detector array is shown in the background.

tor volume. If the detectors of a Compton telescope successfully and accurately measure the energies of all the photon interactions, then Equation 2.1 may be applied to determine the Compton scattering angle θ . If the positions of the photon interactions are also successfully and accurately measured, then a “Compton cone” can be fully specified. The Compton cone surface represents all possible points from which the photon originated. Moreover, the direction of the incident photon must be contained within the Compton cone surface.




Figure 2.2: Backprojected Compton cones for correctly reconstructed photons in spherical coordinates (image is 4π sr). The arcs are found to intersect at the location of the source, which appears as a yellow hot spot. The smaller “circles” correspond to Compton cones with smaller Compton scattering angles.

This technique is illustrated in Figure 2.1. A photon traveling along the z direction with initial energy E_0 Compton scatters with Compton scattering angle θ at \vec{r}_1 , depositing an energy of E_1 . After scattering, the photon’s energy is $E' = E_0 - E_1$. The scattered photon is then photoelectrically absorbed at \vec{r}_2 , depositing an energy of $E_2 = E'$. A Compton cone with half opening angle θ can now be constructed by placing the cone vertex at \vec{r}_1 , and aligning the cone axis with the vector $\vec{r}_1 - \vec{r}_2$.

Without further information (such as the direction of the recoil electron), it is not possible to determine the photon’s origin along the cone surface. However, if multiple photons are detected from the same source, then the Compton cones will be found to overlap at the position of the photon source. For 2D imaging applications, Compton cones are “backprojected” onto the imaging plane, which in the work presented here is typically a spherical coordinate system. An example of this is shown in Figure 2.2; the light blue arcs traced out by the individual Compton cones are found to overlap at the source of the photons, which appears as a bright, yellow hot spot. For 3D imaging applications, a voxelized imaging volume is used instead of an imaging plane, and the Compton cones themselves are projected into the imaging volume.

In addition to the Compton (polar) scattering angle θ , there is also an azimuthal scattering angle η involved in the Compton scattering process (Figure 2.1). The distribution of η from a sample of photons originating from a beam is sensitive to the beam’s polarization level and angle. This forms the basis for using Compton telescopes as polarimeters in the soft gamma-ray band. Compton polarimetry will be discussed in detail in Chapter 4.

2.2 Types of Compton Telescopes

Compton telescopes designs can be broadly divided into two general categories: Scatterer Absorber Compton telescopes (SACTs), and Compact Compton Telescopes (CCTs). SACTs are generally comprised of a scattering detector plane, and an absorbing detector plane. The idea behind the SACT is that an incident photon will Compton scatter once in the scattering plane, and will be photoelectrically absorbed in the absorbing plane. COMPTEL, which flew aboard the Compton Gamma Ray Observatory (CGRO) spacecraft from 1992 to 2000, is the prototypical example of an SACT (Schoenfelder et al., 1993). COMPTEL’s scattering (front) detector plane was comprised of NE 213A liquid scintillator, while the absorbing (rear) plane was comprised of sodium iodide (NaI) scintillator detectors. The separation distance between the two planes was 1.5 m, meaning that only photons with relatively shallow Compton scattering angles ($\theta \lesssim 40^\circ$) could be detected. Photons Compton scattering in the scattering plane with steeper Compton scattering angles would simply miss the absorbing plane, thus reducing the efficiency of detection. One benefit of the COMPTEL design was simplified event sequencing. Due to the fast scintillators used in the detector planes, as well as the large separation distance between the planes, it was possible to use time of flight (TOF) to discriminate between photons originating from the field of view (FOV), and photons originating from the Earth’s albedo from below (a prominent source of background).

CCTs differ from the SACT design in that the detectors are generally not divided into a scattering plane and an absorbing plane. CCTs typically employ a single type of detector arranged into a compact array. As photons propagate through the detector volume, they may Compton scatter multiple times and be subsequently photoabsorbed throughout the entire detector volume. It is not a requirement that particular interactions are segregated into different detector volumes or regions. Typically, the distances between interactions in a CCT are too small to exploit TOF type event sequencing. Therefore, other schemes are used to determine the event sequencing (Boggs & Jean, 2000; Zoglauer, 2005). CCTs tend to have higher efficiency, due to a more uniform distribution of detector mass, which in turn leads to the ability to detect photons which scatter at steeper Compton scattering angles. Another benefit of CCTs is that their ability to measure gamma-ray polarization is better than SACTs. As will be shown in Chapter 4, the ability to detect photons at steeper Compton scattering angles, in particular $\theta \sim 90^\circ$, results in more sensitive polarization measurements.

2.3 The Compton Spectrometer and Imager

Here we introduce the Compton Spectrometer and Imager, a CCT designed to perform astrophysical gamma-ray measurements from a balloon based platform. COSI is the culmination of 20+ years of research and development in germanium detector technology, and Compton telescope data analysis methods.




Figure 2.3: Left: A single GeD in front of a mirror, showing the orthogonality of the strips between the front and back sides. Right: The twelve GeD array during integration. The copper U-shaped bar dividing the array vertically is the upper segment of the cold finger.

2.3.1 Germanium Detectors

The core technology behind COSI are high purity, double-sided strip germanium detectors (GeDs, Figure 2.3a, Amman & Luke 2000). As semiconductor detectors, each GeD is operated as a fully depleted, reverse biased diode. Due to the extremely low impurity concentration of the germanium materials used for these detectors ($n_{\text{imp}} = 1 - 10 \times 10^9 \text{ cm}^{-3}$), large depletion depths can be achieved with bias voltages ranging from $1 - 1.5 \times 10^3 \text{ V}$. This makes it possible to deploy thicker, more efficient detectors.

The COSI GeDs are produced at Lawrence Berkeley Laboratory using a new contact technology based on amorphous germanium (a-Ge) (Amman et al., 2007). Several key benefits arise from using a-Ge contacts in place of the more standard Lithium-diffused n+ and Boron-implanted p+ contacts: 1) The a-Ge layer passivates the surface of the detector, which leads to long-term, stable operation 2) the a-Ge to Ge junction displays good bipolar blocking behavior, thus reducing the overall leakage current 3) the a-Ge material is easily manipulated, allowing for complex electrode patterns with fine inter-electrode spacing (Amman et al., 2007).

The COSI detector array is comprised of twelve such GeDs, which are arranged in a $2 \times 2 \times 3$ array (Figure 2.3b). Each GeD is nominally 8 cm on a side and 1.5 cm thick, with small differences due to crystal processing. The anode (x -side) and cathode (y -side) electrodes on each side of the detector are segmented into 37 strips with a strip pitch of 2 mm

and a gap between strips of 0.25 mm. Each side of the detector has a guard ring surrounding the 37 strips which prevents surface leakage current from flowing between cathode and anode strips. The guard rings are also instrumented and used to veto events that deposit energy in the side-edges of the detector.

Each detector is mounted to a detector board, which facilitates mounting, signal routing, and basic signal conditioning with simple passive electronics. Each strip of the GeD is wire bonded to a copper trace on the detector board. For the strips on the high voltage side of the detector, a decoupling capacitor is placed in series with the strip so as to block the bias voltage from the downstream electronics. Moreover, each high voltage strip is equipped with a bias resistor on the detector board which acts to isolate the high voltage strips from one another. The signals are then routed to a flex circuit, which carries the charge signal through the cryostat wall and into the charge-sensitive preamplifier inputs.

2.3.2 Vacuum Cryostat and Cryocooler

At room temperature, the conductivity of Ge is too high to operate as a detector, due to the small Ge band gap of 0.6 eV. Therefore, Ge must be cooled to cryogenic temperatures for detector applications. The COSI GeD array is integrated into an aluminum cryostat and cooled to ~ 80 K using a Cryotel CT Stirling cryocooler (Figure 2.4). The thermal connection between between the GeDs and the cryocooler is made through a U-shaped, copper “cold finger” (Figure 2.3b), which facilitates both mounting the detectors and conducting heat from the detectors into the tip of the cryocooler. During integration, the detectors are wrapped in a series of thermal blankets that block infrared radiation from the room temperature cryostat walls from being absorbed by the detector assemblies. After integration, the cryostat is connected to a pumping system with brings the pressure inside the cryostat down to a level of approximately 10^{-7} Torr. This is a necessary step, as any residual gas within cryostat will transmit heat to the detectors. The total heat leak into the cryostat is on the order of 5 – 6 W, which is well within the cryocooler’s maximum heat lift specification of 11 W.

Although the heat leak is only 5 – 6 W, the cryocooler is a relatively inefficient device and must dissipate ~ 100 W (external to the cryostat) to manage the cryostat heat leak. This poses a thermal design challenge, as 100 W is a considerable amount of heat to dissipate at balloon float altitudes, where the air pressure is too low for convective cooling to take place. In order to keep the body of the cryocooler from reaching unsafe operating levels (> 80 C°), we developed a liquid cooling system to cool the body of the cryocooler. A cooling fluid (3M Novec 7200) is circulated through a copper sleeve which is thermally connected to the cryocooler body. The heated fluid then travels through a copper heat exchanger, which is thermally connected to a copper radiating plate. The top surface of the plate is painted with high emissivity paint and oriented towards the zenith so that the plate can radiatively cool into space. The fluid loop was driven with a pair of Fluid-o-Tech FG209 magnetic drive pumps. The pumps were wired in series, with only one pump operating normally, and a second pump on standby in case of a failure of the primary pump. Figure 2.5 shows a photograph of the cryostat and cooling system, integrated at the top of the COSI gondola.




Figure 2.4: A photograph of the cryostat with various components annotated. The preamplifier boxes, which ordinarily mount where the flex circuits feed through the cryostat wall, are not present in this image.




Figure 2.5: A photograph showing the COSI cryostat integrated into the top of the gondola, with various component annotated.

2.3.3 Anti-Coincidence Shield

The cryostat is surrounded on the sides and from the bottom by a cesium iodide (CsI, thallium doped) scintillator anti-coincidence shield (ACS) system. Each of six CsI shield modules is comprised of a $40 \times 20 \times 4$ cm block of CsI coupled to two photomultiplier tubes to detect the scintillation light. The two PMT signals from each module are summed and sent to a bipolar shaper circuit with $1 \mu\text{s}$ shaping time. The outputs of the six bipolar shapers are sent to discriminators, whose outputs are OR'ed together to generate a single veto output signal. A tunable threshold per module is set to 80 keV for flight. Particles that deposit more than 80 keV of energy during a time $0.7 - 1.1 \mu\text{s}$ post GeD trigger will result in a veto of any event in the GeDs. The ACS thus actively rejects events that did not deposit their full energy in the GeDs, as well as shields the GeDs from atmospheric background photons from below and from the sides.

2.3.4 Signal Acquisition and Instrumentation

Fig. 2.6 shows a schematic diagram outlining the signal acquisition scheme for the COSI GeDs. The following sections describe the various stages of the signal acquisition process in sequential order.

The Charge Signal

Photons deposit energy in the GeDs primarily through Compton scattering and photoelectric absorption. In both of these physical processes, energy is transferred from the photon to an electron in the Ge lattice, breaking it free of its atomic orbital. The ejected electron will proceed to interact with nearby Ge atoms, primarily by ionizing nearby Ge atoms. Through these ionization losses, a number of electron-hole pairs are generated (approximately 338 pairs per keV of recoil electron energy), which then drift under the influence of the applied electric field toward the nearest strip electrodes. As the charges move through the detector, corresponding image charges are induced on the electrodes, as described by the Shockley-Ramo theorem. It is this induced charge which is actually measured by the preamplifiers. The magnitude of the total induced charge, which is equal to the number of electron hole pairs times the electron charge, is proportional to the energy deposited by the incident photon. This forms the basis of measuring interaction energies in the GeDs.

Charge Sensitive Preamplifiers

The charge sensitive preamplifiers, hereafter referred to simply as the preamps, convert the charge induced on a single strip electrode into a low impedance voltage signal. A schematic of the preamp design is shown in Figure 2.7. The preamp is an inverted cascode type amplifier, which behaves as follows: The circuit node corresponding to the drain of Q1 (an N-channel JFET) and the emitter of Q3 (a PNP BJT) is at a fixed voltage, which is a diode drop above

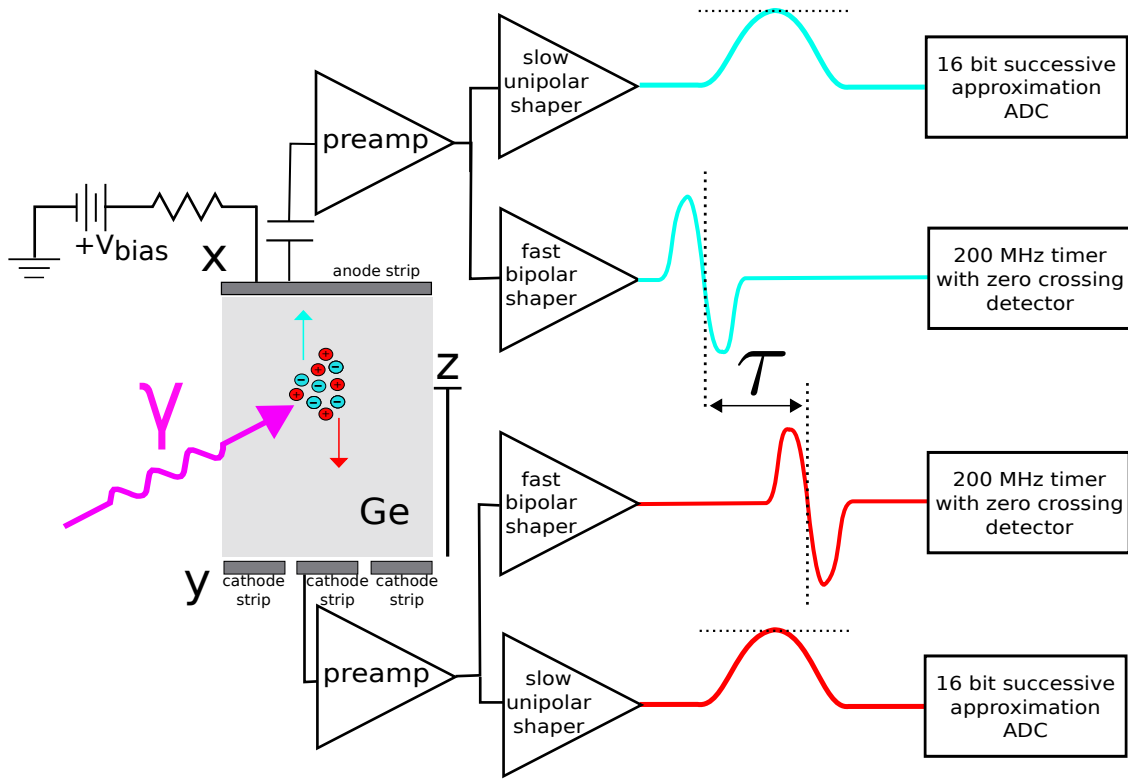


Figure 2.6: A diagram illustrating the signal acquisition process for the COSI GeDs. A photon deposits energy in the germanium, resulting in a number of electron hole pairs. These pairs are swept to the surface electrode by an applied electric field. The preamplifiers convert the collected charge into a voltage pulse, which is then sent to a slow unipolar shaper, and a fast bipolar shaper. The peak value of the unipolar shaper pulses is sampled by an analog to digital converter to facilitate pulse height spectroscopy. The zero crossing time of the bipolar shaper signal is measured by a 200 MHz timer, and the time difference τ between the cathode and anode signal is used to infer the interaction depth.

the base of Q3¹. This voltage does not change appreciably throughout the operation of the amplifier. As a result, the current through R7 is effectively constant. This current, which we will call I_{total} , is divided between the quiescent current through Q1, I_1 , and the quiescent current through Q3, I_3 . Clearly, $I_{\text{total}} = I_1 + I_3$. The magnitude of I_3 is set by the constant current source circuit comprised of Q2, R3, R4, R5, C7, and C36. The overall negative feedback from the output of U1 (a current feedback amplifier) through R1 back to the input will drive the gate voltage of Q1 so that $I_1 = I_{\text{total}} - I_3$. When a positive signal charge $+q$

¹The common base voltage for Q3 is established by the diodes D2 and D3, as well as the resistors R8 and R9. The purpose of the diodes is to compensate the variation in base to emitter voltage of Q3 with changing temperature.

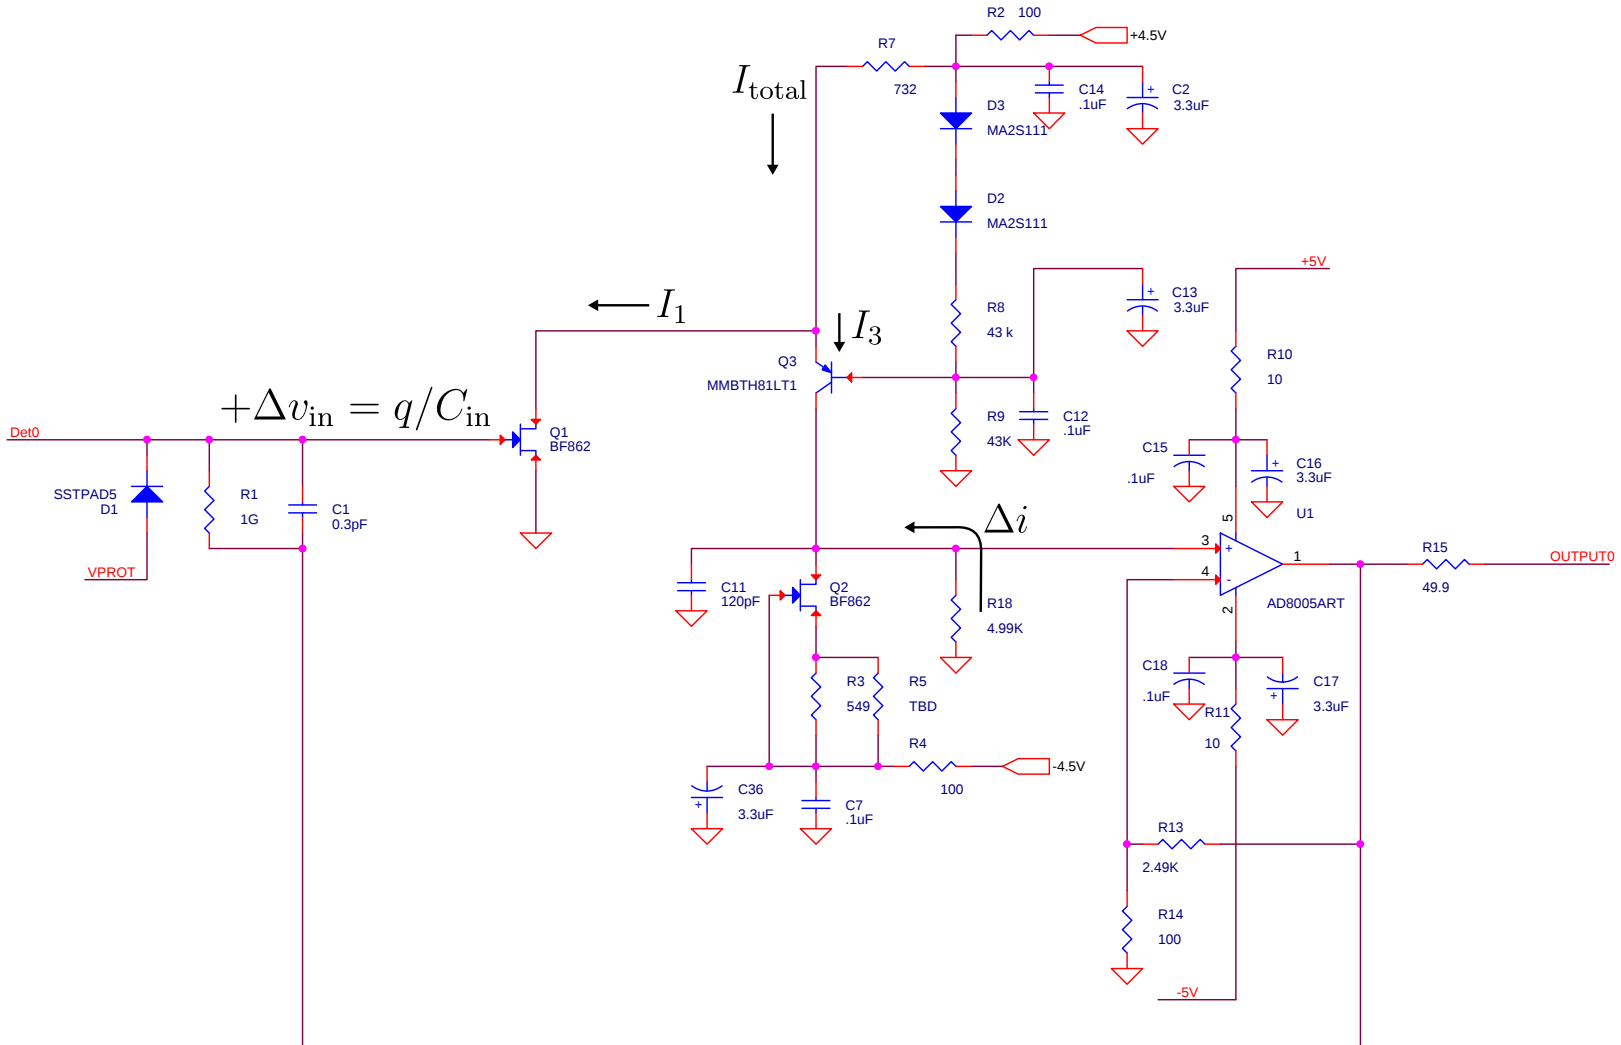


Figure 2.7: A schematic of a single preamplifier channel, with several annotations corresponding to the circuit description in Section 2.3.4.

is deposited at the gate of Q1, the gate voltage will increase by $+\Delta v_{\text{in}} = q/C_{\text{in}}$, where C_{in} is the input capacitance. This increase in the gate voltage results in an increase in the drain to source current of Q1, Δi , so that the total current through Q1 is $I_1 + \Delta i$. Due to the fact that I_{total} is constant, the current flowing through Q3 must be $I_3 - \Delta i$. This can be thought of as Q1 “stealing” a current Δi , which would have otherwise flowed through Q3 and into the constant current source circuit (the drain of Q2). To make up for this current deficit and to maintain a constant current flowing into the drain of Q2, a current with magnitude Δi flows from ground, through the resistor R18, and into the drain of Q2². This produces a negative voltage $-\Delta i R_{18}$ across R18, which is then amplified with a gain of ~ 25 by the non-inverting current feedback amplifier U1. Note that the voltage across R18 is negative, while Δv_{in} at the gate of Q1 is positive. The total open loop gain is $-25g_m R_{18} \Delta v_{\text{in}}$ Volts per Coulomb, where g_m is the transconductance of Q1 ($g_m = 17.3 \text{ mA V}^{-1}$ on average). The negative gain results in a negative feedback loop from the output of U1 back to the input via R1 and C1, which works to counteract any voltage changes at the gate of Q1. As the output of U1 quickly swings in the negative direction, a dV/dt current is seen across the feedback capacitor C1, which results in a current flow through C1. It should be emphasized that the feedback resistor R1 has a very high impedance at high frequencies, so virtually all of the signal current flows through C1, thus charging up C1. The output pulse from U1 is the desired, low impedance voltage signal whose height is proportional to the signal charge. After this whole process has finished, C1 slowly discharges through R1 with a time constant of $R_1 C_1 \sim 0.3 \text{ ms}$. This long decay time is responsible for the tail-pulse shape of the preamplifier output pulses.

The main advantage of using an inverted cascode type design is that the signal voltage does not appear at the drain of Q1. This way, the Miller effect is mitigated, which would otherwise lead to an amplification of the Q1 gate to drain capacitance. In general, it is crucial to minimize the input capacitance at the preamp, as the gain is proportional to the inverse of the input capacitance.

The preamp channels are laid out on printed circuit boards, ten channels per board. Each board also has its own low-noise, linear power supply which supplies the positive and negative rail voltages to the preamp circuits. A “preamp box” contains four preamp boards, which is enough to instrument one side of a detector. An image of a preamp box with the lid removed is shown in Figure 2.8. The preamp boxes are mounted directly to the cryostat where the flex circuit carrying the signal charges feeds through the cryostat wall.

Card cages

The outputs of the preamps are sent down 50Ω coaxial ribbon cables to the “card cages”, which house all of the pulse processing and triggering electronics for a single GeD. Each card cage is made up of four different boards:

²The capacitor C11 is a compensation capacitor, used to tailor the high frequency gain of the amplifier so that instabilities are avoided. At high frequencies, a fraction of Δi may also flow through C11.




Figure 2.8: A preamplifier box with the lid removed. The ten preamp channels on the top board are visible. The front side of the preamp box (not visible) connects to a flex circuit, which carries the charge signals from one side of a detector through the wall of the cryostat.




Figure 2.9: A COSI card cage used to readout a single GeD. The boards from bottom to top are as follows: high voltage power supply board, low voltage power supply board, DSP board, 8x analog boards. One analog board has been extracted so that the ten shaper channels are visible.

1. 8x “Analog boards”. The analog boards contain the analog pulse processing electronics, as well as digital logic for triggering and queuing of the pulse heights and pulse timings.
2. 1x “DSP board”. The DSP board mainly consists of an Altera Stratix FPGA which runs an Altera NIOS II soft core CPU. The FPGA handles the global triggering logic and is responsible for fetching the pulse height and timing values from the analog boards over a parallel data bus. Software running on the CPU relays the raw data as well as housekeeping information and trigger rates to the flight computer over an ethernet connection using the UDP protocol.
3. 1x Low voltage power supply (LVPS) board. All rail voltages used in the card cage circuitry are generated on the LVPS board using DC/DC converters to convert the incoming battery voltage to the desired voltages.
4. 1x High voltage power supply board (HVPS). The HVPS board supplies a single detector with a high voltage ranging from 1000 to 1500 Volts using a Cockroft-Walton type circuit. The high voltage is slowly ramped up and down at a rate of approximately 1.5 V s^{-1} in order to prevent damage to the preamplifiers.

A photograph of a COSI card cage can be found in Figure 2.9.

Absolute time tags are assigned to triggered events via a 48 bit counter driven by a 10 MHz oscillator. The 10 MHz oscillator actually comes from the flight computer, and is shared between all twelve card cages. Synchronization between the card cages is achieved using a “sync” signal that also comes from the flight computer and is shared between all twelve card cages. When a pulse arrives on the sync line, the value in the 48 bit counter in each card cage will jump to a specific value. The register containing this 48 bit specific value can be updated via the ethernet connection prior to the synchronization event. Maintaining the synchronization between the card cages is crucial for identifying coincident events between detectors.

Pulse Shaping

The preamplifier signal is sent to a pair of pulse shaping circuits, which coexist on an analog board within the card cage. The “fast” shaper is a bipolar pulse shaper with a nominal shaping time of 180 ns. The fast shaper signal is sent to a zero crossing detection circuit whose output is used to latch the value of a 200 MHz timer. This way, the zero crossing times of the fast shapers can be measured down to a 5 ns resolution. By computing the time difference τ between zero crossings of a triggered cathode and anode strip, the depth (or z coordinate) of the interaction can be inferred. Beyond measuring the signal timing, the fast shaper output is also used to initiate the trigger logic for the readout system, and latch a timestamp for the event.

The second pulse shaper is a unipolar, “slow” shaper, as its shaping time is nominally set to 6 μs . The slow shaper circuit is designed to precisely measure the pulse height of the

preamp signal, with as little noise as possible. The peak height of the pulse is converted by a successive approximation analog to digital converter using a bit depth of 16 bits, although only the most significant 13 bits are stored. The conversion is triggered by a differentiator circuit, which detects the peak position using the zero crossing of the signal's derivative. The voltage of the peak height, which is proportional to the preamp pulse peak height, is also proportional to the energy deposited in the interaction.

Trigger Logic

A trigger in a detector is initiated when the fast shaper of any strip crosses the fast threshold, which is typically set to around 40 keV - 50 keV, depending on the noise present on a given channel. In order for the readout sequence to continue, another fast shaper trigger on the opposite side of the detector as the first trigger must occur within 360 ns of the first trigger. This coincidence between the two sides of the detector is necessary, as noise will produce many false triggers on only one side of the detector. The triggering logic then spends 2.4 μ s waiting for veto signals from the cesium iodide anti-coincidence shield system or the guard ring. If no veto arrives, the readout sequence finishes by measuring the pulse heights of the slow shaper channels that exceeded their thresholds, which are typically set to around 20 keV. The zero crossing times and pulse heights of all triggered channels are recorded and stored in a hardware queue for future readout by the DSP board.

2.4 Gondola Systems

The COSI gondola is the aluminum structure supporting the entire instrument. As well as providing mechanical support, the large surface area of the gondola also makes it an effective radiator. Thus, various systems on the gondola rely on the gondola structure itself to dissipate heat. The entire COSI gondola hanging from the launch vehicle can be seen in Figure 2.10.

2.4.1 Flight Computer (GCU)

The COSI flight computer, hereafter referred to as the GCU (originally ground control unit), is responsible for performing a variety of critical tasks. Figure 2.11 shows a diagram of the GCU and its relationship to other systems on board the balloon gondola. The main CPU board is an Advantech PCM-9562, which employs an Intel Atom D510 1.66 GHz dual core CPU. A number of PC-104 expansion cards are stacked on the CPU board in order to add needed functionality:

1. 1x CTI Xtreme/104 Opto board, which provides four opto-isolated RS232 serial ports. Note that the main CPU board has 3x RS232 serial ports as well.
2. 1x Diamond Systems DMM-32DX-AT board, which provides 32 analog to digital converter channels, used for reading out 32 LM335 temperature sensors.




Figure 2.10: A photograph of the entire COSI gondola hanging from the launch vehicle in Wanaka, New Zealand during a test. Various subsystems are annotated for reference.

3. 1x custom designed “clock board” which performs several functions. First, the clock board has a 10 MHz crystal oscillator, which is distributed to all of the card cages for time stamping events with a precision of 100 ns. The “sync signal” is also generated by the clock board and sent to the card cages. The clock board maintains its own 48 bit counter, which is synchronized to those in the card cages. Beyond the 10 MHz clock logic, another timing related function of the clock board is to detect edges on the pulse-per-second output (PPS) from the GPS. As mentioned above, when a rising edge arrives on the PPS line, the clock board latches the 48 bit counter value, which can be used later for aligning event times to the absolute time. Finally, the clock board also reads the veto signal from the ACS, and counts the number of rising edges. This way, the GCU can directly measure the count rate of the shield system.

The GCU runs a standard Linux operating system (Slackware 13.1.0). The flight software was written in C with the intent of being fully autonomous, and as reliable as possible. A number of threads run concurrently to perform various tasks such as collecting raw data from the instrument, storing data to disk, sending data down various telemetry streams, real time analysis of COSI events and count rates, and interacting with various subsystems on the gondola.

Data Storage

All data that is generated on the gondola is eventually forwarded to the GCU, where it is stored on three redundant solid state drives (SSDs). The SSDs used in the GCU are off-the-shelf, 1 TB SATA drives from SanDisk and Samsung. The drives are formatted using the Linux ext4 filesystem.

Telemetry

Various telemetry streams are available on the gondola. The highest bandwidth stream is a line of sight (LOS) link, which provides for up to 1 gigabit down link. However, this stream is only available while the gondola is within the line of sight, which typically lasts for only several hours following launch. The LOS connection has a sufficient capacity to transmit all raw COSI data, with the exception of when the gondola passes through the Pfozter maximum. The LOS link is handled by an “LOS box”, which is connected to the GCU via a USB connection. The LOS box places its own headers around data frames and then uses Manchester encoding to transmit the data over a serial line to the LOS L-band transmitter. On the ground side, a similar LOS box decodes the received data, which is then read over a USB connection by a ground computer.

Outside of the LOS, the primary telemetry stream is the Iridium Openport (OP) connection which operates over the Iridium satellite network. The gondola is equipped with two OP L-band transceivers, which are each capable of up to 128 kilobit bidirectional communication. On the gondola side, the OP transceivers are connected to the flight computer over an ethernet connection, and provide full internet access to the GCU. The GCU sends

each packet redundantly down each OP connection using the UDP protocol. These packets are routed onto the open internet and received by a Linux workstation located at the Space Sciences Laboratory. All housekeeping data is sent down the OP links, but the capacity of the link is not large enough to handle the raw card cage event data. Instead, a real time analyzer (RTA) algorithm on the GCU identifies possible Compton events, and these are telemetered over the OP link. Each OP link has sufficient capacity to transmit all Compton events under normal conditions at balloon float altitudes.

In addition to the OP links, the gondola is also equipped with two low speed Iridium modems, which can be used to telemeter 256 byte packets every several minutes. The modems can be “dialed up” to provide up to 2 kilobaud of link capacity. Neither of these links are sufficient to transmit science data. Instead, housekeeping information is sent through these links as a backup in case of a failure of the OP or LOS links.

Real Time Analyzer

The real time analyzer (RTA) is tasked with identifying all possible Compton events so that they may be telemetered through the OP links. This way, the RTA filters out single-site events which cannot be used for Compton imaging or Compton polarimetry. Without the RTA, it would not be possible to analyze data real-time on the ground, as the raw COSI data stream is too large for the OP links.

Two types of Compton events are identified by the RTA: single detector Compton events, and multi-detector Compton events. A single detector Compton event consists of multiple hits in a single detector. In order to qualify as a single detector Compton event, at least two strips on each side of the detector must have triggered both the slow and fast shaper channels, and a separation of at least one strip must exist between any two strips on at least one side of the detector. For example, if strips 4 and 5 triggered on the x side, and strips 16 and 17 triggered on the y side, this would not qualify as a single detector Compton event, since the triggered strips on each side are neighboring, i.e. not separated by at least one strip. On the other hand, if x strips 4 and 5 triggered, and y strips 16 and 18 triggered, this would qualify, since the strips on the y side are separated by a strip. This requirement for strip separation effectively filters out single site events where charge is shared between neighboring strips. Although some of these events may indeed be Compton events, they are generally of lower quality since the hits are not separated by a large distance. A multi-detector Compton event is simply defined as any event that resulted in a trigger of more than one card cage. These events are automatically considered candidate Compton events and telemetered over the OP links.

As a consequence of parsing the raw event data from the card cages in real time and performing the coincidence search, the RTA algorithm is the dominant factor in the overall CPU usage of the flight software. However, the maximum load on the CPU never exceeds 20-30%, even during calibrations when the count rate from the card cages saturates (approximately 1700 triggers per second from each card cage).

Transient Monitor

A primary science goal of the COSI experiment is to observe gamma-ray bursts, and possibly other gamma-ray transients. Transient events are most easily identified by analysis of the light curve. The transient monitor algorithm (TMA) uses the CsI shield light curve as input to identify possible transient events. Each sample of the shield count rate is determined by integrating the number of shield counts over one second time intervals. Before appending the sample to the end of the light curve, the samples are passed through a median filter with a window size of 3 samples in order to suppress single sample spikes which are due to heavy cosmic rays exciting long lived phosphorescence states in the CsI scintillator. After filtering, a background count rate is computed over the previous 100 samples, and each new sample is compared to the background level. If the new sample is 6σ above the background, then this new sample is considered to be a transient. In the event of a trigger, all raw data for the following ninety seconds is logged to a single file (in addition to the file logging that occurs continuously throughout the operation of the flight software) with a standardized name and saved in an easily accessed location for future download using standard programs such as `rsync`, `scp`, etc. Also, a special alert packet is issued with the time of the triggering sample. This alert can be used by operators on the ground to download a data set corresponding to the transient.

In the case that the OP links are fully functioning without any significant delay time, the TMA is not necessary. However, it is not uncommon for the OP links to drop out for a sustained period of time, in which case an autonomous identification of transient events becomes necessary. Once the links come back online, a COSI team member can easily inspect the incoming housekeeping data to determine if a transient event was identified by the TMA during the outage.

Commanding

Various commands can be sent to the GCU to perform a number of functions, for example changing thresholds for the pulse shapers, cycling power to the gondola subsystems, resynchronizing the clock counters, throttling data rates, adjusting the cryocooler power, and more. During ground calibrations, commands are simply sent over an ethernet connection. During flight, commands are primarily issued over the OP links by first establishing a secure SSH tunnel to the GCU. Setting up the command link can take anywhere from less than ten seconds to over a minute, depending on the latency of the OP links, which can be highly variable. In general, the process takes less than a minute, and once the tunnel is established, subsequent commands can be sent relatively quickly. When this approach is not working, commands can be sent over a low speed iridium link, which are forwarded by the CSBF SIP to the GCU over a serial connection. These commands are also variable in their timing, and can take up to a minute or two to arrive at the GCU. Finally, when the payload is still within the line of sight, commands can be sent over an LOS uplink to the CSBF SIP, which then forwards the commands to the GCU over the serial connection. The command links

are shown with blue circles in Figure 2.11.

2.4.2 Power System

The COSI payload is powered using solar panel arrays, or PV (photo-voltaic) arrays. During the day, a rotator keeps the gondola azimuth oriented towards the sun so as to maximize the generated solar power. The current from the solar panels is used to both supply the instrument with power, and to charge an array of lead acid batteries. At night, the instrument operates on battery power. The nominal system voltage is 24 V, but this voltage varies in the range 20 V to 30 V depending on the charge state of the batteries. In order to handle this wide range of voltages, all gondola subsystems are equipped with DC-DC converters which can handle input voltage ranges of 18 V to 36 V. Under normal operating conditions, the instantaneous power consumption of the entire system is between 400 W and 450 W.

Power is distributed to gondola subsystems via the Power Distribution Unit (PDU). Each power channel of the PDU is equipped with a magnetically latched relay that is switchable by the flight computer. Hall effect sensors read out the current draw on each power channel, as well as the total system current draw. This information, along with the system battery voltage is relayed to the flight computer over a serial connection.

2.4.3 Aspect and GPS Systems

As a wide FOV imager, COSI is not required to point at fixed locations in the sky. However, in order to determine the origins of photons with respect to astronomical coordinate systems, the orientation of the gondola must be known at all times. The primary component of the aspect system is an ADU5 differential GPS (dGPS) system, which reports the gondola heading (from true North), pitch, and roll on a one second cadence. The pointing error of the ADU5 system is $< 0.2^\circ$, which is small relative to the COSI's angular resolution (6° FWHM at 662 keV). Two backup aspect systems are available as well: a Trimble BX982, and an Applied Physics Systems model 544 magnetometer and accelerometer. The BX982 is only able to provide the heading from true North ($< 0.1^\circ$ error), while the 544 can provide the magnetic heading ($\pm 1.2^\circ$ error) as well as pitch and roll ($\pm 0.4^\circ$) based on the accelerometer readings. The data from the BX982 and the 544 are also read out at one second intervals.

Both the ADU5 and the BX982 provide standard GPS data, such as latitude, longitude, altitude, and time. Additionally, the CSBF SIP has a GPS system of its own which can be queried by the GCU. All in all, there are three sources of GPS information available to the GCU.

Beyond providing aspect data and GPS data, the ADU5 and the BX982 serve additional purposes. The ADU5 is equipped with a pulse per second (PPS) output, which transmits a TTL level pulse once a second. The rising edge of this pulse is within 200 ns of the UTC second. The COSI system relies on this signal to align the 48 bit timestamps of events in the GeDs to the absolute time. This is accomplished by sending the PPS signal to the clock board in the GCU, which also maintains a synchronized 48 bit counter along with the card

cages. When a rising edge on the PPS line is detected, the value of the 48 bit clock is latched and stored in a queue for future readout by the flight software. During post-processing, these PPS latches are used to align the 48 bit card cage timestamps to the absolute time. The BX982 also serves as an NTP server for the gondola, which the GCU uses to keep its system time synchronized to the GPS time to within several milliseconds. This is useful because each COSI data packet contains the UTC time in the header as measured by the GCU, so having NTP synchronization ensures that the UTC time stamps in the packet headers are synchronized to the GPS time. Without the NTP server, the CPU system time would drift by a few seconds each day.

2.5 Ground Software

The GSE (originally, ground support equipment) is the software that runs on the ground computers, which enables members of the COSI team to monitor the data from the payload as it arrives in real time, as well as send commands to the GCU. The incoming data stream to the GSE can be selected to come in over the LAN (during ground calibrations), over the LOS link, over the OP link, or over the low speed Iridium links, and the commanding up link can also be configured similarly. The GSE is written in Python, and uses the Qt framework along with the pyqtgraph library to perform real time plotting of virtually any parameter over long time records. Additionally, all time series plots can be arbitrarily zoomed and panned, giving an operator the option to examine large scale trends as well as small scale features in the data.

2.6 Calibrations

As mentioned in Section 2.3.4, the slow pulse shapers are equipped with circuitry to measure the pulse heights in volts, and the fast pulse shapers are equipped with zero crossing detectors which measure the time of arrival of the charge signal on the two sides of the detector. Before this data can be used in high level analysis such as Compton imaging or Compton polarimetry, the measured voltages and timings must be converted into interaction energies and detector depths, respectively. The calibrated energies are further processed to correct for cross-talk and possibly charge loss between strip electrodes. Finally, “hits” are assembled by performing the strip pairing, which is the process of finding the triggered strips on each side of the detector that correspond to the same physical interaction.

2.6.1 Energy Calibration

Various factors are at play which cause variations in the gain from channel to channel, including but not limited to: preamplifier gain, detector capacitance, and slow shaper gain. Differences in gain can be on the order of 25-50%. Therefore, each strip electrode must

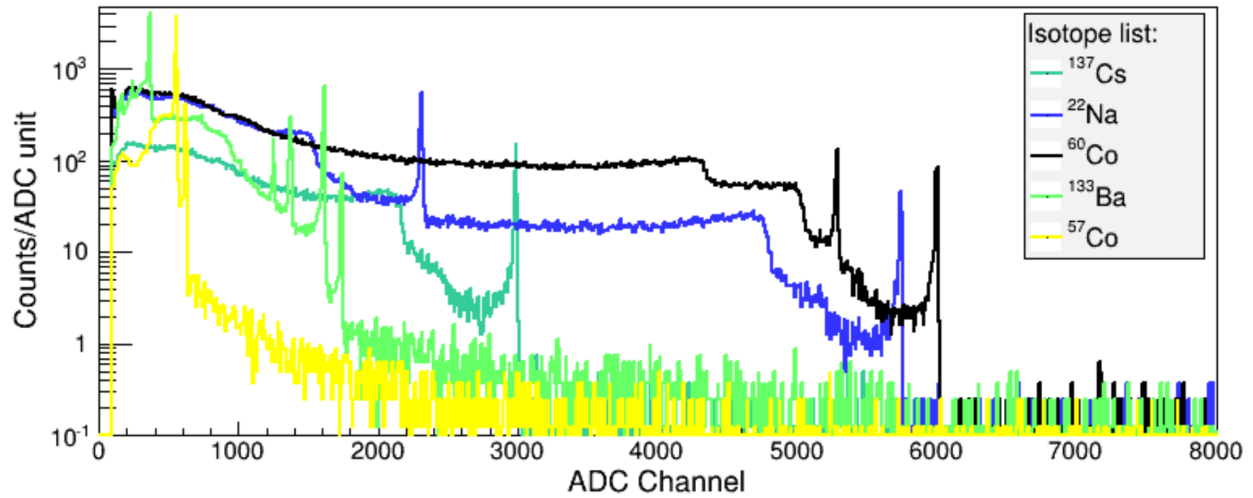


Figure 2.12: Spectra of several radioactive sources in channel space for a single strip electrode. The channel number on the x axis is proportional to the voltage pulse height from the slow shaper. Each peak in the spectra is fit with an analytic model of the emission line, thus yielding a curve that maps channel number to photon energy.

be calibrated individually to determine the relationship between the pulse heights (or ADC channels) from the slow shaper, and the true energy deposited in the interaction.

The calibration is performed by simply illuminating the detectors with a variety of radioactive sources. Each source emits gamma-rays at known energies, so the peaks in the ADC channel spectra for each strip can be associated with the known energies. After doing this for several line energies, a third order polynomial is fit to the set of energy-ADC channel pairs. The best fit polynomial is then used to convert ADC channels to energies. Figure 2.12 shows spectra for several radioactive sources in ADC channel space for a single electrode strip.

2.6.2 Depth Calibration

Determining the relationship between the collection time difference τ and the depth of the interaction z is referred to as the depth calibration. Unlike the energy calibration, where a calibration curve is sought for each detector strip, the depth calibration seeks to find $z_p(\tau)$ for each detector pixel p ³. With 37 strips on each side of the GeDs, the number of pixels is $37^2 = 1369$. Multiplying by twelve detectors, this results in a total of 16428 pixels which must undergo a depth calibration. Two approaches to the depth calibration are discussed in detail in Appendix A.

³A pixel is defined as the intersection of an x strip and a y strip.

2.6.3 Crosstalk Correction

When multiple strips on one side of a detector collect charge, it is possible for the signals to couple to each other, possibly at the detector, in the preamplifier circuitry, in the signal cables, or in the shaper circuitry. The exact source of the effect is unknown. The total charge collected on one side of a detector when multiple strips are triggered will appear to be higher than what it would have been if all charges had been collected to a single strip. Previous work on the cross-talk effect using the COSI detectors has shown that the effect is strongest for adjacent strips, and diminishes as the separation between strips increases (Bandstra, 2010). Fortunately, the enhancement in energy is linear with the measured energy, and so the correction is rather straightforward to apply. Each face of each detector is calibrated with a pair of linear coefficients (slope and intercept) for adjacent strips, and a pair of coefficients for skip-1 neighbor strips.

2.6.4 Charge Loss Correction

Often times, the charge clouds that are generated by interactions in the detectors may be located between strips. In this case, the charge is collected to two adjacent strips. Previous work has demonstrated that when such charge sharing occurs, there can be a deficit in the total collected charge. This is usually attributed to incomplete charge collection, due to a small fraction of the charge being “stuck” in a region of low electric field at the surface of the detector and between the strip electrodes. The stuck charges are eventually collected by the strip electrodes, but over a timescale longer than the shaping time of the slow shaper, so they do not contribute to the measured pulse height.

Recent work has demonstrated that charge sharing can be strongly mitigated by using several specialized processing techniques pertaining to the sputter deposition of the amorphous Ge contacts (Looker et al., 2015). Interestingly, the COSI GeDs do exhibit charge loss on the x side (anode) of the detector, but not on the y side (cathode). Due to this observation, we currently do not perform any charge loss corrections. Instead, we simply rely on the energy measured on the y side of the detectors.

2.6.5 Strip Pairing

Once the energy calibration and cross talk correction have taken place, the triggered strips on the x side of the detector must be paired with the corresponding triggered strips on the y side. The strip pairing process would be simple if each triggered event in a detector represented a single interaction. However, triggered events frequently represent multiple interactions. For example a Compton scatter may be followed by a photoelectric absorption, both occurring in the same detector. In this example case, the two interactions deposit two separate energies, and it is necessary to determine which x and y strip correspond to the Compton scatter, and which x and y strip correspond to the photoelectric absorption. Essentially, the algorithm finds the x and y strips that are closest in energy, and pairs them.

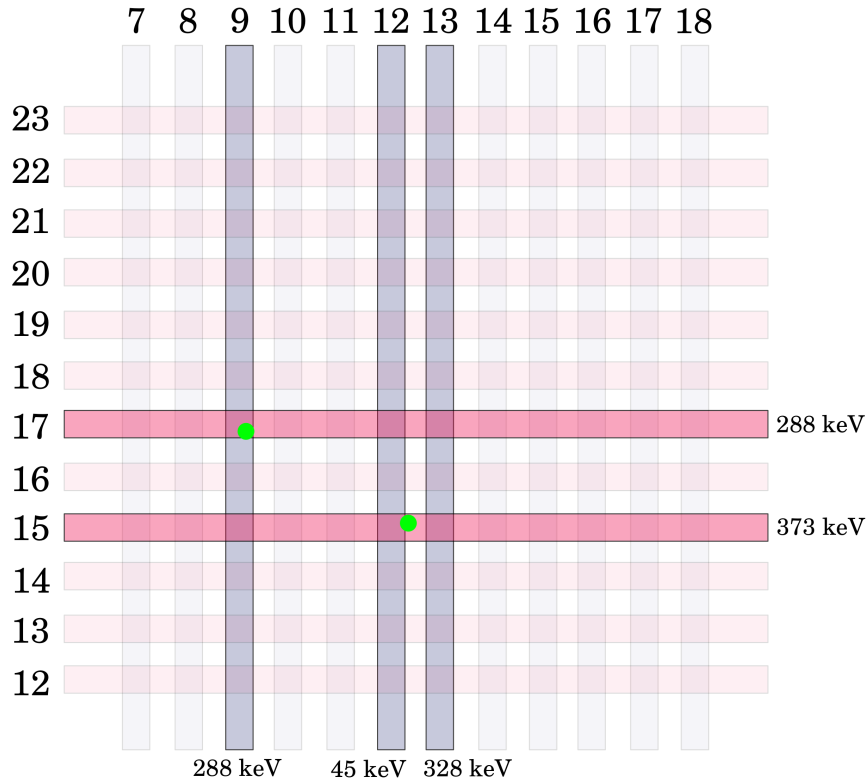



Figure 2.13: A diagram illustrating the strip pairing process. In this example, a photon interacts at two locations marked by the green circles. This results in 3 triggered x strips and 2 triggered y strips. The strip pairing must successfully determine that x strip 9 and y strip 17 should be paired, and that x strips 11 and 12 (which shared the total charge) should be paired with y strip 15. Charge sharing, statistical fluctuations on the measured energies, hits below threshold, and strips with multiple hits can complicate the strip pairing process.


In the presence of charge sharing, statistical fluctuations in number of charge carriers, and sub-threshold energy deposits, this procedure can be difficult to perform. Figure 2.13 gives an example case where two physical interactions result in 3 triggered strips on the x side, and 2 triggered strips on the y side. It is the job of the strip pairing algorithm to associate x strip 9 with y strip 17, and x strips 12 and 13 with y strip 15.

2.7 Data Analysis

COSI software tools are based on MEGALib, the Medium Energy Gamma Ray Astronomy Library (Zoglauer et al., 2006). MEGALib itself is built on top of the Monte Carlo particle simulator GEANT4 (Agostinelli et al., 2003), and the particle physics data analysis library ROOT (Brun & Rademakers, 1997). The first piece of software in the COSI data pipeline is Nuclearizer, which is structured as a set of modules that sequentially manipulate the events.



(a) Measurement pipeline.



(b) Simulation pipeline.

Figure 2.14: The data pipelines for measurements (top) and simulations (bottom). The simulated pipeline makes use of the Detector Effects Engine, which adds realism to the Cosima simulation. Also, the simulated pipeline does not include the cross-talk correction, as that is currently not modeled in the simulations.

The data pipeline for the measured data is shown in Figure 2.14a, where the following Nuclearizer modules are applied sequentially: parser and coincidence search, energy calibration, strip pairing, cross-talk correction, and depth calibration. The simulation pipeline is somewhat different, and is shown in Figure 2.14b. Simulated data is first processed by the Detector Effects Engine (DEE, Sleator et al., 2017) which performs inverse energy and depth calibrations, adds noise, and models the electronics (thresholds, finite bit depth, etc.) in order to add as much realism as possible to the simulated data. After the DEE, the events are passed through the standard calibration modules which are used for measured data, with the exception of the cross-talk correction module, as the cross-talk effect is currently not modeled in the simulations. The final output from Nuclearizer is a list of calibrated hits, or energy deposits, for each event. Each hit is parameterized by a time, energy, a 3D position, and the associated uncertainties for these quantities.

The next piece of software in the pipeline is Revan, which performs the event reconstruction. Event reconstruction involves the following steps: 1) decide whether or not to cluster hits if there are hits that occurred close together 2) decide if an event is a Compton (multi-hit) event or a single-site event 3) for Compton events, determine the proper sequencing. Event sequence determination is a challenging task, and several algorithms have been developed to do this. In the remainder of this thesis, we will use the Bayesian Approach from Zoglauer (2005) for event sequence determination. Once Revan has identified and reconstructed the Compton events, high level data analysis such as Compton imaging and Compton polarimetry may proceed.




Figure 2.15: A top view of the COSI mass model, which is used for performing Cosima simulations, as well as global 3D positioning of measured events.

2.7.1 Simulations and Mass Model

Simulations are run using the MEGAlib program Cosima, which provides a user friendly interface to the GEANT4 simulation framework. A Cosima simulation is parameterized by 1) a particle beam geometry 2) the energy spectrum, flux, and polarization of the particle beam 3) a mass model, and 4) a physics list describing the models used for physical interactions.

We have developed a detailed mass model of the COSI detector system, which is shown in Figure 2.15. The mass model describes the locations of both active detector mass and passive mass in as much detail as possible. Mass model accuracy is essential to the overall success of COSI data analysis. While the mass model is clearly a core component of Cosima simulations, it is also extremely important in the analysis of real data. The reason for this is that the global coordinates of the individual hits are determined by consulting the mass model. Specifically, the absolute x , y , and z coordinates that are inferred from the triggered x strip, the triggered y strip, and the collection time difference, respectively, will be affected by any position offsets or geometric modeling inaccuracies of the detectors.

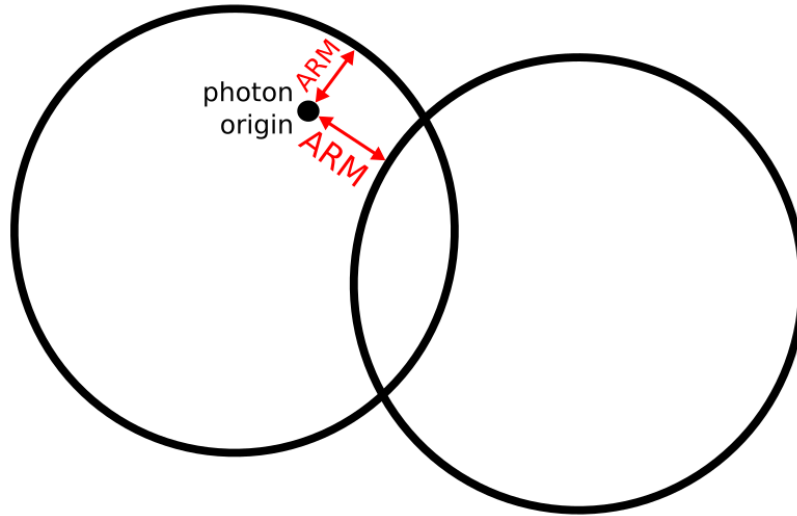


Figure 2.16: A diagram illustrating the definition of the Angular Resolution Measure (ARM). The source of the photons is the small, filled black circle. The open circles represent the arcs traced out by the Compton cones of two separate photons. The smallest angular distance between an event’s Compton cone and the origin of the photons corresponds to the ARM for that event. When the photon origin falls outside of the arc, the ARM is positive, and when it falls inside, the ARM is negative. Note that ARM values are always relative to a specified location.

2.7.2 Event Selections

Compton events are parameterized by an assortment of quantities such as energy, Compton scattering angle, distance between interactions, etc. Data analysis proceeds by imposing selections on these quantities — “event selections” — which act to reduce background, increase signal, and optimize imaging, spectral, or polarization performance. The most relevant event selections for analysis of COSI data are as follows:

1. Total photon energy. All energy deposits are summed to get the total photon energy. It is possible for a photon to escape the detector system without depositing all of its energy, but these events can be identified with reasonable certainty and are rejected during event reconstruction.
2. First Compton scattering angle. This is the Compton scattering angle of the first interaction, and is determined by evaluating the classic Compton scattering formula (Equation 2.1) using the total photon energy and the photon energy after the first Compton scatter.
3. Angular resolution measure (ARM). This quantity is defined as the smallest angular distance between an event’s Compton cone and a position in the sky. A diagram illustrating this definition is shown in Figure 2.16, and an ARM distribution from a

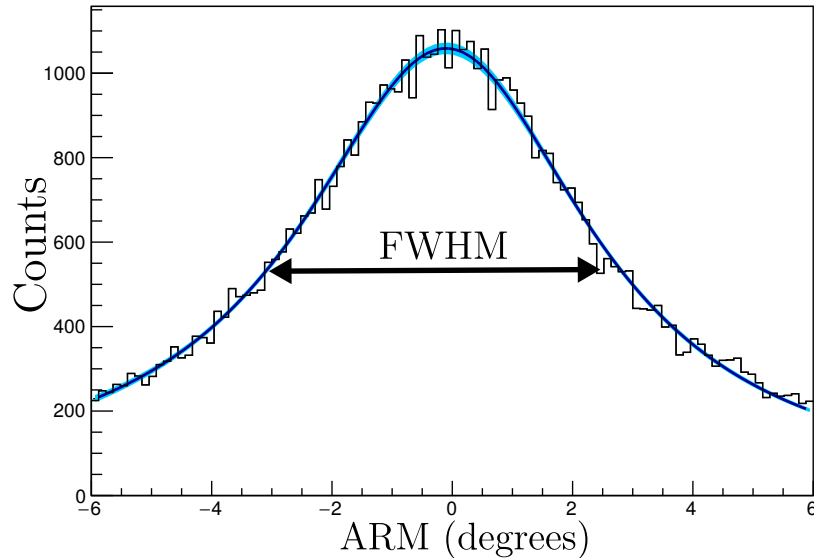


Figure 2.17: ARM distribution from a Cs-137 (662 keV) calibration run, along with a fit to the data. The x axis represents the smallest angular distance between a Compton cone and the specified source location. The analytical fit function is a superposition of two Lorentzians and an asymmetric Gaussian. This shape is empirical and accurately fits the ARM peak. The FWHM is determined from the fit rather than the data itself.

Cs-137 (662 keV) calibration run is shown in 2.17. Placing a selection on the ARM allows one to select events whose Compton cones are consistent with a specified location in the sky. ARM values can be positive or negative, depending on whether the source location falls outside (positive) or inside (negative) the Compton cone. The FWHM of the ARM distribution is interpreted as the angular resolution of a Compton telescope.

4. Distance between first two interactions. This distance determines the accuracy with which the orientation of the Compton cone can be reconstructed. For small distances, the orientation is more sensitive to the position uncertainty, which will have a negative impact on the imaging performance. Small distances also have the effect of increasing the error on the azimuthal scattering angle, which has a negative impact on the polarization performance.
5. Distance between any two interactions. Events with interactions that are too close are harder to reconstruct properly. Therefore, excluding these events will result in a better reconstruction efficiency.
6. Number of interaction sites. This parameter impacts the reconstruction efficiency, since events with more interaction sites can generally be reconstructed more accurately. Events with two interaction sites are the most common, followed by three sites, four sites, etc.

2.7.3 Compton Imaging

Simple Compton images can be made by backprojecting Compton cones into either a 2D imaging plane or a 3D imaging volume. Further deconvolution of the image may be performed using advanced algorithms such as List-Mode Maximum Likelihood Expectation Maximization (LM-MLEM, Wilderman et al., 1998), Maximum Entropy (Gull & Skilling, 1984), or Stochastic Origin Ensembles (Andreyev et al., 2011). The goal of all of these algorithms is to determine the most likely source distribution in the image or volume, given the initial back-projection, as well as the imaging “response,” which for COSI is a cone.

2.7.4 Compton Polarimetry

A useful feature of Compton telescopes, central to this thesis, is that they are inherently sensitive to polarization. This topic will be discussed at greater length in later chapters, but a brief summary follows. Two angles are involved in Compton scattering: the Compton scattering angle θ and the azimuthal scattering angle η (Figure 2.1). Information relating to the polarization state of a gamma-ray beam is encoded in the η distribution. If a source of gamma-rays is polarized, the distribution of η when measured over many events will take the form of an offset sinusoid with a period of π . The amplitude of the sinusoid is proportional to the polarization level Π , and the phase of the sinusoid is the polarization angle η_0 with respect to a reference direction. Polarization measurements provide unique information about source geometries, magnetic fields, and emission mechanisms that often cannot be ascertained from spectral or timing properties.

Chapter 3

The 2016 COSI Super Pressure Balloon Campaign

In 2016, COSI flew from Wanaka, New Zealand on a Super Pressure Balloon, continuously observing the gamma-ray sky for 46 days as the payload traversed the globe. A significant volume of data was generated, which so far has revealed multiple interesting sources of astrophysical gamma-rays. This accomplishment within the COSI collaboration was the product of 1) hard work on behalf of previous and current COSI team members 2) several previous balloon flights, both successful and unsuccessful, and 3) several iterations of the COSI instrument leading up to where we are today. In this chapter, we will start with a brief historical review of COSI and its predecessor, the Nuclear Compton Telescope (NCT). The remainder of this chapter will serve as a summary of the 2016 flight with details such as an instrument performance, background environment, and a quick “catalog” of sources detected by COSI thus far.

3.1 From NCT to COSI

The predecessor to COSI was named the Nuclear Compton Telescope, which took part in two successful balloon flights. In 2005, a two GeD prototype flew from Ft. Sumner, New Mexico on a zero pressure balloon (ZPB) and measured the background environment, which was found to be in close agreement with predictions from background modeling (Bowen et al., 2007). In 2009, a full, 10 detector instrument was flown again from Ft. Sumner New Mexico on another ZPB (Bandstra et al., 2009). NCT successfully achieved its primary science goal by detecting and imaging the Crab Nebula (Bandstra et al., 2011). This represented a milestone for astrophysical Compton telescopes, as NCT became the first compact Compton telescope to image an astrophysical source. The instrument was recovered in good shape from the 2009 flight and quickly turned around for a 2010 flight from Alice Springs, Australia. Unfortunately, a launch mishap resulted in partial destruction of the payload. Luckily, the GeDs and readout electronics were relatively unharmed.

Between 2010 and 2014, NCT was significantly overhauled in several key ways, and a new name was given to the instrument; the Compton Spectrometer and Imager. COSI differs from NCT in a number of key ways:

1. The cryostat was completely redesigned in order to accommodate two more detectors (twelve total, compared to NCT's ten), and a new detector configuration. The COSI GeDs are laid out in a $2 \times 2 \times 3$ array, whereas the NCT GeDs were laid out in a 2×5 array.
2. NCT used liquid nitrogen (LN2) to cool the GeDs to cryogenic temperatures. Gradual evaporation of the LN2 placed a hard constraint on flight durations. As described in the previous chapter, COSI uses a cryocooler to cool the GeDs, which operates continuously with a steady source of electrical power.
3. The BGO anti-coincidence shield system for NCT was replaced with an ACS system consisting of CsI scintillators.
4. NCT had closed-loop azimuth pointing control based on orientation data provided by a 3-axis magnetometer. Additionally, the NCT detector system was tilted so that the FOV was not zenith centered. This was done to optimize exposure to sources that were part of the observing plan, such as the Crab Nebula. In contrast, the COSI FOV is zenith centered, and does not require active pointing control; sources are simply observed as they pass through the FOV. Note that the azimuth of COSI is controlled by a rotator above the gondola in order to keep the solar panels oriented towards the sun.
5. The NCT flight computer (GCU) was retired, and a new GCU system was designed. With virtually all new hardware, the NCT flight software was no longer appropriate. Therefore, the flight software was rewritten, with various features added for enhanced telemetry applications and autonomy.
6. NCT relied solely on LOS telemetry, which was possible due to the short flight durations. COSI's anticipated flight durations of up to 100 days makes LOS telemetry unfeasible. Instead, COSI relies on two Iridium Openport (OP) transceivers to telemeter data from any location on the globe. A significant effort was made to incorporate event identification, filtering, and compression in the flight software in order to transmit all Compton events through the OP links.

3.2 Super Pressure Balloon

Many of COSI's design parameters were driven by the prospect of flying on NASA's 18 million ft³ Super Pressure Balloon (SPB). The primary difference between an SPB and a ZPB is that the SPB is sealed, whereas the ZPB has an opening at the base of the balloon from

which helium gas can escape. When the gas in an SPB is heated, due to solar irradiance as well as albedo irradiance, the pressure increases. In an SPB the balloon is sealed, and the balloon material is rigid enough to mostly counteract the increase in volume that would be associated with a rise in gas pressure. The buoyancy force is proportional to the volume of the balloon, so this volumetric stability leads to altitude stability. In contrast, when the gas is heated in a ZPB, the pressure increases, the volume increases, and the balloon's altitude increases. When the balloon cools back down at night, the reverse process happens. The constant altitude cycling of the balloon results in a gradual loss of helium gas, and thus a gradual decrease in altitude over time. Usually these altitude variations are compensated somewhat by dropping ballast, of which a finite amount is flown. The SPB platform should be able to provide for much longer flight durations due to the altitude stability inherent in the design.

While the altitude stability of the ZPB may not be as good as the SPB, the ZPBs are generally easier to manufacture, as the balloon material does not need to be as rigid as is the case for the SPB. Moreover, NASA has decades of experience successfully flying the ZPB, while the SPB remains somewhat of an experimental platform. Lastly, the current generation of NASA ZPBs are volumetrically larger than the current generation of SPBs (40 million ft³ vs. 18 million ft³), so the maximum attainable float altitude is higher than that for SPBs. Higher altitudes are certainly desirable in high-energy astrophysics, as any residual atmosphere above the telescope will act to scatter and absorb photons. Nonetheless, NASA continues to push the SPB technology, as it is expected that longer flight durations and altitude stability will be desirable to science teams.

3.3 2014 Antarctic Flight

In October 2014, the COSI team traveled to McMurdo Station, Antarctica to fly the COSI payload on the SPB. After several weeks of canceled launch attempts, COSI finally launched on December 28, 2014. This marked the first flight of a science payload on the NASA SPB. The instrument reached float altitudes and proceeded to collect background data. However, after only ~ 36 hours at float, the SPB developed a catastrophic leak, leading the balloon operators to terminate the flight. The payload made a gentle landing on the Antarctic plateau, and was successfully recovered a few weeks later.

While the outcome of the 2014 Antarctic balloon campaign was disappointing, several important lessons were learned. First, the cryocooler was not effectively dissipating heat at float. This led to a situation where the cryocooler body temperature was operating outside of specifications (> 80 °C), which not only risked the health of the cryocooler itself, but also precluded effective cooling of the GeDs. In light of this, we developed the liquid cooling system (Section 2.3.2) to keep the cryocooler within its operating range. Second, it was clear that our ground software was not optimized to support a long duration flight. To address this, we overhauled the ground software completely, adding features such as SMS/email alerts for events of interest (parameters going out of range, transient event detected by the

TMA, etc.), and real time plotting of any parameter over an arbitrarily long time record. Lastly, the COSI team gained invaluable experience from successfully deploying the COSI payload in the field.

3.4 2016 Flight from Wanaka, New Zealand

Following the 2014 Antarctic campaign, arrangements were made to deploy the payload again in 2016 from Wanaka, New Zealand. This would be the first scientific payload to fly from the new NASA balloon base in Wanaka. The COSI team arrived in New Zealand in mid February, 2016 and promptly began to integrate the gondola.

3.4.1 Ground Calibrations

After integrating the instrument and all subsystems into the gondola, we carried out an extensive calibration operation throughout the weeks prior to launch. The detectors were illuminated with the following radioactive sources from different positions within the field of view: americium-241, cobalt-57, cobalt-60, barium-133, sodium-22, cesium-137, and yttrium-88. The resulting datasets were used to perform energy calibrations, depth calibrations (Appendix A), effective area benchmarking, and general benchmarking against simulations. Additionally, a considerable amount of effort was invested into performing a polarization performance validation, which is the subject of Chapter 5.

3.4.2 Flight Summary and Overview

Unexpectedly poor weather conditions through April and May resulted in many aborted launch attempts. However, COSI was finally launched on May 16, 2016 at 23:35 UTC time into clear skies. Figure 3.1 shows COSI hanging from the launch vehicle on the flight line, with the SPB almost completely inflated on launch day. At the time of launch, all systems were fully operational and the health of the system was perfect. After an hour of ascending towards float, the payload traversed through the Pfozter maximum (~ 17 km), where the concentration of cosmic ray secondaries reaches its peak. As a result, our LOS telemetry link saturated, thus complicating data down link for a few minutes. This was expected and has been seen on all other COSI/NCT balloon flights. After just over two hours, the balloon reached its nominal float altitude of ~ 33.5 km. Once at float, the OP links were enabled, and found to be working within their specifications.

The flight path taken by COSI is shown in Figure 3.3, where the color of the trajectory alternates after each day. Initially, and contrary to expectations, the payload headed west towards Australia, flying nearly over Sydney during the 2nd day of flight. During days 3 and 4, the payload's trajectory shifted eastward, where it was carried around the globe by circumpolar winds in ~ 10 days. Upon passing by New Zealand, the payload drifted North towards the equator, reaching a latitude of 5.65° S on the 28th day of flight. The overall




Figure 3.1: The COSI flight line on launch day. COSI can be seen hanging from the launch vehicle on the left. The inflated SPB is on the right of the image. In the center of the image is the parachute, as well as the uninflated portion of the balloon. This portion expands significantly as the pressure drops with altitude until the balloon makes a round, pumpkin shape at float.

wind patterns and resulting direction of the payload during this time were highly variable, with no clear timeline for a future land crossing. Finally, a trend in the wind direction developed and carried the payload over Southern Peru. The flight was terminated, and COSI fell back to Earth on its parachute, making a relatively gentle landing in the Atacama Desert. Figure 3.2a shows a still image from a video that was shot from the gondola as COSI made its landing. The video was downloaded almost immediately over the OP links, which survived the landing. Within three days, a team comprised of COSI and Columbia Scientific Ballooning Facility (CSBF) personnel was able to locate the payload and balloon in a remote area of the desert. Figure 3.2b is a photograph of COSI at its landing site, taken by the recovery crew. Unsurprisingly, the outer structure including the antennae booms, solar panels, and gondola frame sustained damage during the landing. However, the instrument and electronics internal to the gondola were all recovered in excellent condition, including the solid state drives containing the full raw dataset.

3.4.3 Thermal Environment

A major challenge of flying from New Zealand is the highly variable thermal environment, which can result in excessively cold temperatures at various points on the gondola. The coldest temperatures were seen when flying over storm systems with heavy cloud cover, as the cloud tops have a low surface temperature and act as an additional radiative heat sink (in addition to space above). Fortunately, COSI was equipped to handle these situations; the GCU controlled a “heater box”, which made it possible for the flight computer to ther-



Figure 3.2: Left: A still from a video shot from the gondola, moments before landing in the Atacama desert. The shadow of the COSI payload and parachute are clearly seen on the desert floor. Right: A photograph taken by the recovery team of COSI at its landing site. The exterior structure sustained damage, but the instrument and electronics were recovered in excellent shape.

mostatically maintain the temperature of gondola subsystems. The temperatures were read out by LM335 sensors, and kept within a programmable range autonomously by the flight computer. Evidently, the system worked well, with no temperature going out of range, and no temperature related anomalies identified thus far.

3.4.4 Altitude Variations

Approximately 19 days into the flight, the SPB started to exhibit significant day-night variations in altitude. The most extreme drop in altitude occurred on June 21, where the payload dropped to 22 km, a difference of approximately 11.5 km from the nominal float altitude. The altitude profile for each third of the flight is plotted in the top panels of Figures 3.4, 3.5, and 3.6, where the altitude drops are apparent. The reason for this behavior has not been determined conclusively, but one possibility that would result in a similar pattern would be if the balloon had a temporary leak of gas. In this scenario, the nominal float altitude would not be affected, as prior to losing gas the balloon would have been over-pressurized. However, with less gas, the balloon altitude becomes less resilient against drops in pressure, thus leading to altitude dips at night when the temperature decreases. Unfortunately, the altitude dips were significant enough to impact the observation plan. As the payload altitude decreases, the column density due to the atmosphere becomes more severe. Additionally, the background count rate increases at lower altitudes (next section), which impacts the sensitivity of the observations. Although, the altitude variations significantly impacted our

observation time, the payload still rose back to a nominal altitude of 33.5 km each day.

3.4.5 Background Environment and Relativistic Electron Precipitation

The count rates in the CsI shield system as well as detector 6 are shown in the flight summary plots Figures 3.4, 3.5, 3.6 as a function of time. Two trends are readily apparent. The first is that the count rates are anti-correlated with the altitude, which is most evident in Figures 3.5 and 3.6. This is due to the fact that beyond that Pfozter maximum, the flux of cosmic ray secondaries, which dominate the COSI background, decreases with altitude. When the altitude drops, the payload moves close to the Pfozter maximum, thus resulting in an enhancement of the count rates. The second trend is that the count rates are anti-correlated with the distance to the South Magnetic Pole (SMP). The angular distance to the SMP is plotted in the second panel from the top in Figures 3.4, 3.5, 3.6, and mostly tracks the geographic latitude. This is due to the dependence of the cosmic ray primary flux on the cutoff rigidity with respect to Earth’s geomagnetic field. Near the SMP, the geomagnetic cutoff rigidities are reduced, which increases the flux of cosmic ray primaries that reach the upper atmosphere. These cosmic rays interact with air nuclei and produce electromagnetic and hadronic showers in the atmosphere, which constitute the dominant background component for balloon borne gamma-ray detectors.

The count rate histories in Figures 3.4, 3.5, 3.6 also reveal strong variability at three distinct times, all during close approaches to the South Magnetic Pole: May 21 06:37 to 10:15 (REP1), May 30 03:00 to 09:57 (REP2), and May 30 16:18 to 18:23 (REP3). These time durations are approximate and not easy to define, as the main episodes are often pre-empted and followed by episodes of lower level activity. These events are due to relativistic electron precipitation (REP), which is a radiation belt phenomenon in which electrons that are trapped in the outer Van Allen radiation belts are released following geomagnetic disturbances (Parks et al., 1979; Millan & Thorne, 2007). The precipitating electrons travel along the geomagnetic field lines and interact in the upper atmosphere, producing copious amounts of bremsstrahlung X-rays. The time domain behavior of the observed REP activity appears to be divided into two categories: fast (< 100 ms) pulses of x-rays — “microbursts” — only seen during REP2, and slow, minutes long pulses reaching very high count rates which occurred for all three REP events. The points in the flight trajectory where REP activity was observed are denoted by the blue and green circles in Figure 3.3. While REP related activity is certainly an interesting scientific topic, it mainly acts as a source of background for observations of astrophysical sources. Notably, the long duration gamma-ray burst GRB 160530A occurred during the REP2 event. Figure 3.7 shows the total COSI event light curve, as well as the Compton event light curve for the first 12 hours of May 30, 2016 which included both the REP2 event and GRB 160530A. The GRB is the spike just after 07:00:00, most visible in the Compton event light curve. The microburst activity can be seen starting around 03:15:00 and lasting up until about 06:30:00, at which point the

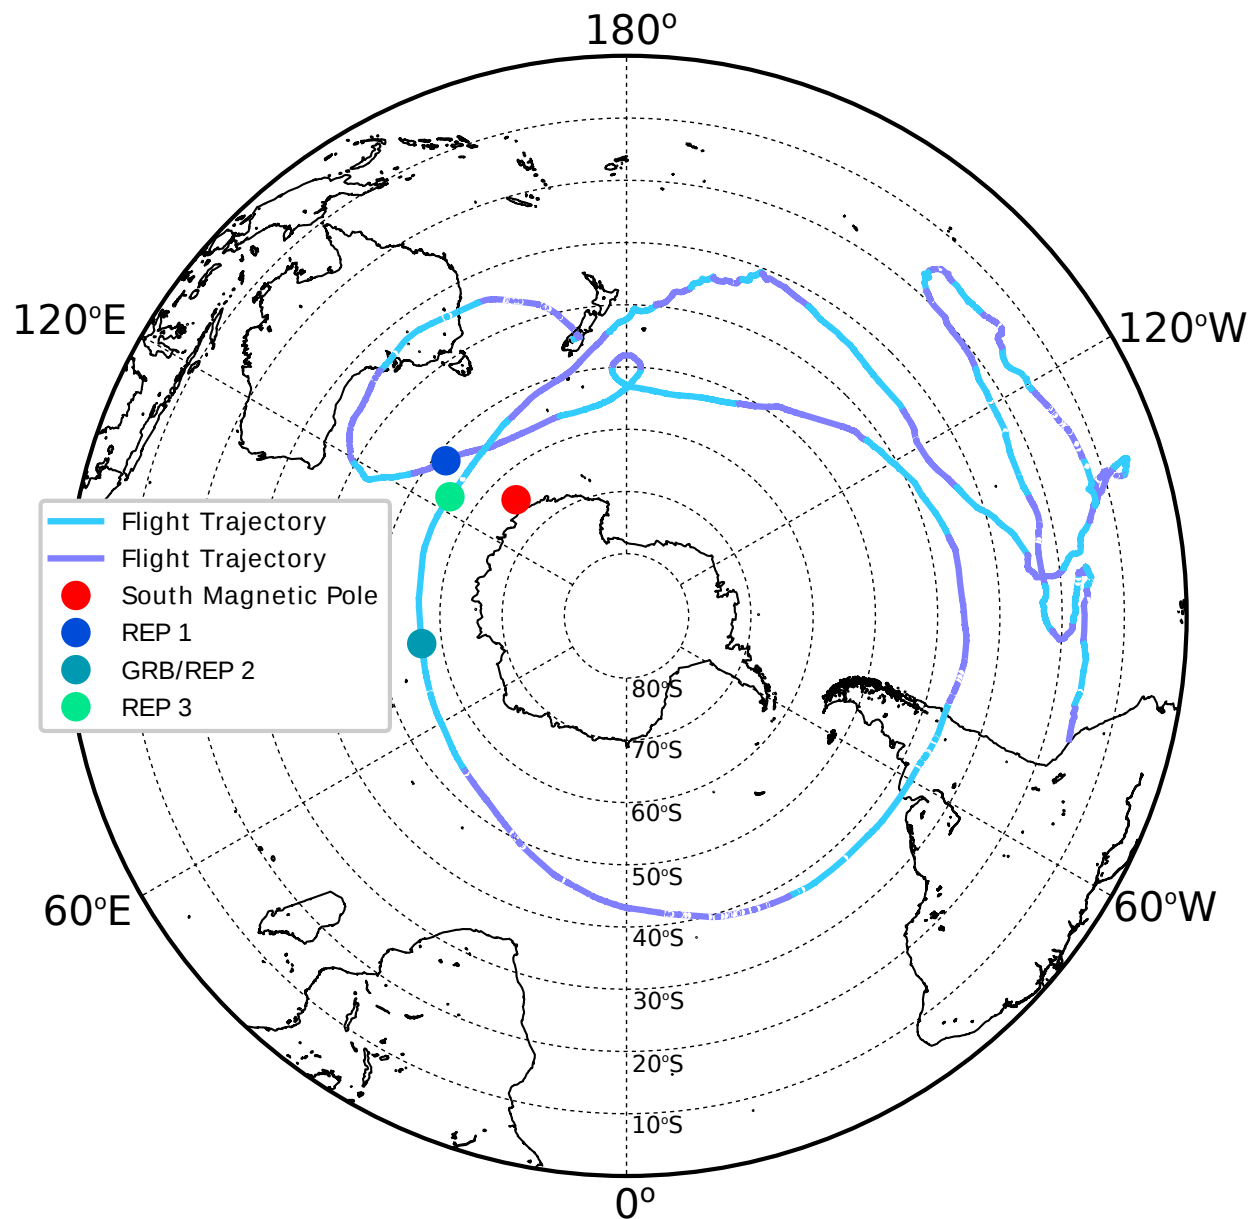


Figure 3.3: The COSI flight trajectory throughout the Southern Hemisphere. COSI launched from Wanaka, New Zealand, flew for 46 days, and then landed in Southern Peru. The color of the flight trajectory alternates for each day of flight. A red circle denotes the position of the South Magnetic Pole. The blue/green circles mark COSI's position at the time that the REP events were detected, as well as GRB 160530A.

behavior transitions into the slow bright pulses.

The shield count rates in Figures 3.4, 3.5, 3.6 were median filtered to suppress 1 and 2 sample “spikes” of extremely high amplitude. The origin of these spikes is likely to be the excitation of numerous, long-lived phosphorescence states in the CsI crystal, probably in response to excitation by heavy cosmic rays. Similar behavior has been seen in space-borne CsI crystals (Hurley, 1978; Segreto et al., 2003). The exact mechanism by which these events produce an artificially high count rate is unclear. One possibility is that the decay of the phosphorescent states results in a bright and steady light output from the CsI crystal over a period of ~ 100 ms. The magnitude of the shot noise then increases, and passes through the shaper. If the noise is not attenuated enough by the shaper, then it may vary in the vicinity of the discriminator threshold and cause an artificially high count rate. We used a median filter in the flight software to suppress 1-sample spikes, but not 2-sample spikes. The 2-sample spikes led to a large false alarm rate in the TMA of roughly 1-2 auto triggers per hour.

3.4.6 High Voltage Anomalies

After approximately 20 hours at float, it became clear that detector 8, a detector in the bottom layer, was malfunctioning. This manifested as sporadic, high current draw from the preamplifier boxes assigned to this detector. It is likely that this problem was due to a malfunction in the high voltage system, which was reflected in the preamp current draw. Specifically, if the high voltage was arcing to the cryostat or some connector along the high voltage path, this would result in an abrupt change in the high voltage, which has the net effect of producing a dV/dt current through the coupling capacitors on the preamp channels connected to the high voltage side of the detector. Detector 5 began to malfunction the following day, with very similar behavior as detector 8. This detector was in the top layer of the detector array. Nineteen days later, another detector in the top layer, detector 0, had a high voltage related malfunction. However, this detector showed distinct behavior from the other two. While it is likely that a malfunction in the high voltage system was at play, the origin of the problem may have been different. All three detectors which displayed high voltage anomalies were shut down upon verifying that these issues were persistent.

During data analysis, these detectors were treated as passive mass in the simulations. Beyond simply reducing the effective area, the loss of detectors is also problematic because disabled detectors will block the incoming gamma rays from interacting in detectors beneath the disabled ones. Consequently, response matrices used for imaging, spectroscopy, and polarimetry must all be regenerated for different time periods throughout the flight corresponding to different detector on/off combinations.

3.4.7 Performance of Telemetry Systems

For the first several hours of flight, while the payload was within the LOS, the LOS telemetry system worked as expected. As the payload approached the horizon, the number of bit errors




Figure 3.4: From top: The altitude, angular distance to the SMP, count rate (> 80 keV) in the shield system, and count rate (> 50 keV) in detector 6 (top layer) for the first third of the flight. The shield rate has been median filtered with a window size of 5 samples to suppress 1 and 2 sample spikes due to excess rate events. The count rate in detector 6 has been box car averaged with a window size of 5 samples to reduce the Poisson noise in the plot. A clear anti-correlation exists between the distance to the SMP and the count rates. The flaring activity near the SMP is due to REP activity.




Figure 3.5: The same as Figure 3.4, but for the second third of the flight. A clear anti-correlation can be seen once again between the distance to the SMP and the count rates, as well as the altitude and the count rates.




Figure 3.6: The same as Figure 3.5, but for the final third of the flight. Significant altitude variations were experienced on a daily basis, with corresponding enhancements in the count rates.

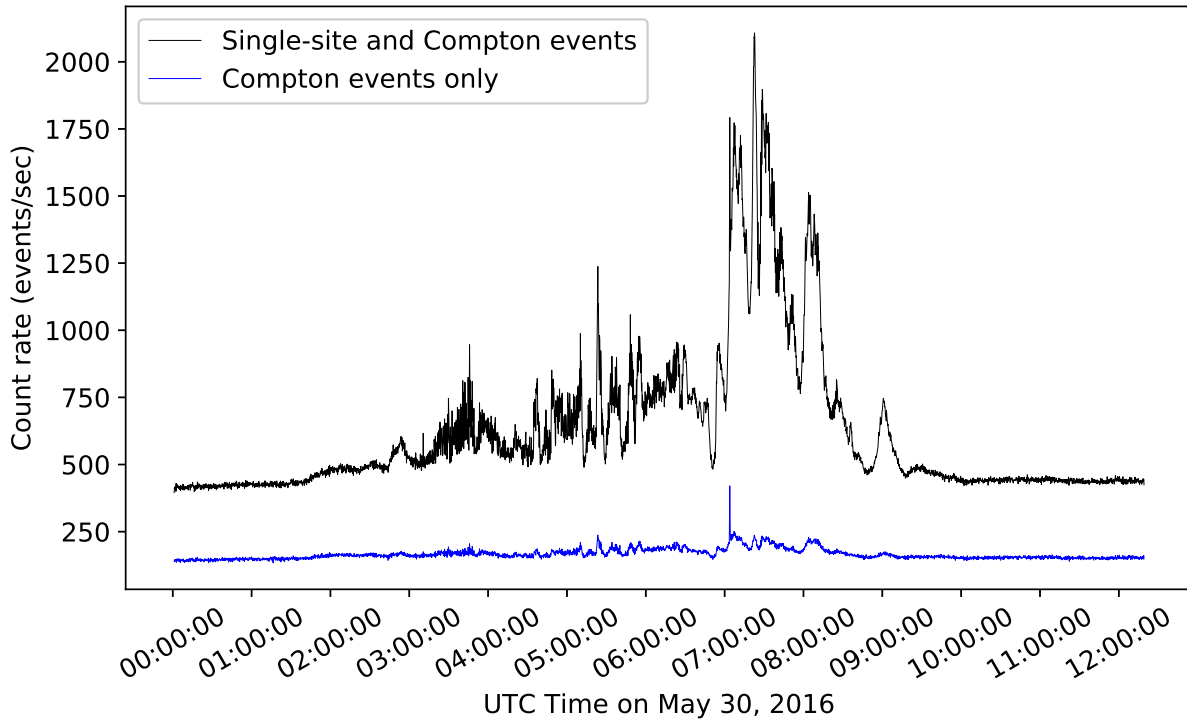


Figure 3.7: Reconstructed event light curves for all COSI events (black) and Compton events only (blue) for the first 12.5 hours of May 30, 2016. This time period included the second REP event (REP2), which began as fast ~ 100 ms pulses and transitioned into slower and brighter pulses just before 07:00:00. GRB 160530A is the spike just after 07:00:00. The relative amplitude of the REP activity over the pre-REP background is significantly smaller in the Compton event light curve, indicating that the REP photons are too low in energy ($\lesssim 200$ keV) to produce many Compton events.

on the LOS stream became unacceptably large, prompting the COSI team to switch over to the OP links for the remainder of the flight.

The two OP links served COSI well overall. For most of the flight, both OP transceivers worked reliably well, with both OP transceivers sending identical copies of the data down via a UDP connection to a server at Space Sciences Laboratory. Occasionally, one or both of the OP links would drop out. These dropouts could last anywhere between a minute, and an hour. When both OP links dropped out, data loss clearly occurred. When one OP link would drop out, the other link would often continue to work well. In this scenario with one OP link active, some data was still lost, as the UDP connection does not guarantee delivery of the data packets. Unfortunately, the OP hardware is essentially a black box, providing little to no diagnostics that would help understand the conditions that lead to drop outs and data loss. Nonetheless, the OP links were essential to our ability of performing real time




Figure 3.8: The cumulative transmitted and received data volume for the first ten hours of May 20, 2016. The red curve representing the received data slowly diverges from the transmitted curve as small numbers of packets are lost in the Iridium network over time. At the end of the ten hour period shown here, 2.4% of the data is lost, which is considered to be acceptable. All lost data was eventually recovered using the packet fetcher just before flight termination.

data analysis on the ground throughout the flight.

Figure 3.8 shows the cumulative transmitted and received data for the first ten hours of May 30, 2016, a time period that included both a REP event and GRB 160530A. The black curve is the cumulative transmitted data in megabytes, determined by scanning the recovered hard drive data for packets that were telemetered over the OP links. The red curve is the cumulative received data over the OP link, determined by scanning the OP telemetry files as received by our ground computer in Berkeley. At the end of the ten hour period, the amount of data received was lower than the amount transmitted by 2.4%, indicating that 2.4% of data was lost. This was typical for most periods of the flight where the OP links were known to be operating. Note that because the UDP protocol was used, entire data packets were either delivered successfully or lost. It is not the case that the 2.4% loss rate was evenly distributed over random bytes; this would be much more of a problem, as segments of complete packets would be missing, thus severely complicating the data parsing process.

A loss rate of 2.4% is a small price to pay for real time access to all Compton events from any location on the globe. However, there were other periods where the OP links simply failed to operate for extended periods time, leading to total data loss. To combat this, we developed a set of software tools before launch to retrieve lost data packets from the GCU. In the header of each COSI packet is a time stamp (with 1 second resolution) as well as a packet counter, specific to the packet type. The time, packet counter, and packet type uniquely specify every COSI packet. Missing packets can be identified by scanning over all received packets and identifying gaps in the packet counter values. To retrieve these lost packets, a special data message containing the missing packet IDs and timestamps are uploaded to the GCU, where a “packet fetcher” script scans through the data stored on disk, collects all of the requested missing packets, and saves them to a special file in a standardized location. The file containing the lost packets can then be downloaded over the OP links using standard file transfer tools such as scp, rsync, etc. Towards the end of flight, we were successful in running the packet fetcher to retrieve all lost packets. This was an important step, as it was not clear whether or not COSI would land in the ocean, in which case recovery of the hard drives would have not been possible.

3.4.8 Ground Operations

At all times throughout the 46 days of flight, at least one COSI team member was “on shift” monitoring the payload health, sending any necessary commands, and responding to alerts from the transient monitor algorithm. The false alarm rate for the transient monitor algorithm was unfortunately quite high due to the excess rate events; approximately two per hour. In response to TMA alerts, the ground computers would run an automated analysis and generate a reconstructed event light curve, as well as LM-MLEM images corresponding to the time of the alert. The COSI operator on shift would inspect these data products to determine if the TMA alert was not a false alarm, and then respond accordingly. In the case of GRB 160530A, this system worked as intended; the TMA issued an alert which was received over the OP links, and the automated analysis of Compton events from the OP links revealed an interesting light curve as well as an image with a clear peak. This prompted the COSI team to issue a Swift target of opportunity observation of the COSI error region for GRB 160530A.

3.5 COSI 2016 Flight Source Catalog

In this section, we will give a brief overview of each gamma-ray source that has been detected by COSI so far using the data from the 2016 flight. A total of four persistent sources, and two transient sources were detected and imaged. One detail pertaining to the detected sources whose positions are known is that the peaks in the COSI images for these sources are slightly offset from their known positions. It is not entirely clear why this happens, but two possible explanations are 1) the images are not exposure corrected, so the peak positions could be




Figure 3.9: COSI images in galactic coordinates of four persistent astrophysical sources that were detected during the 2016 flight: The Crab Nebula (a), Cyngus X-1 (b), Centaurus A (c), and 511 keV positron annihilation emission from the galactic center (d). The peaks in the images are slightly offset from the known positions, possibly due to the lack of exposure correction, or uncorrected systematic errors in the aspect reconstruction. All of these are point sources, with the exception of the 511 keV image, which is known to be diffuse from INTEGRAL/SPI observations (Bouchet et al., 2010). Efforts are underway to determine if the central bulge in the COSI image is diffuse, and whether or not the apparent extended structure surrounding the bulge is real.

biased in the direction of larger exposure and 2) there could be systematic errors in the aspect reconstruction which result in positioning errors.

3.5.1 The Crab Nebula

COSI's Northward excursion later in the flight led to significant exposure to the Northern sky. As a result, COSI detected the Crab Nebula, a nearby supernova remnant powered by a young pulsar. Due to its close proximity (2 kpc) and high intensity, it is an exceptionally well studied source. The spectral shape of the Crab is known to follow a power law with a spectral index of -2.2 above 100 keV, as measured by INTEGRAL/SPI (Jourdain & Roques, 2009). The off-pulse emission from the Crab Nebula above 200 keV is highly polarized at a level of $\Pi > 72\%$ (Forot et al., 2008), with a possible recent change in the polarization angle (Moran et al., 2015).

Figure 3.9a shows the COSI image of the Crab using two days of data where the Crab's position was closest to the center of COSI's field of view. Note that this is the second time our collaboration has imaged the Crab Nebula, the first time being in 2009 with NCT. Currently, efforts are underway to reconstruct the Crab Nebula spectrum (Sleator et al., 2017), perform a polarization analysis, and detect pulsations from the Crab pulsar. These properties are generally well understood for the Crab, and will thus serve as useful benchmarks for COSI.

3.5.2 Cygnus X-1

Another Northern sky source detected by COSI was Cygnus X-1, a galactic high-mass X-ray binary system known for its bright and frequent outburst episodes. Like the Crab, Cyg X-1 is a well studied source. Its spectrum in the soft gamma-ray band consists of a thermal Comptonization component, as well as a hard power law tail extending to over 1 MeV (Jourdain et al., 2012). Energy resolved polarimetric studies of Cyg X-1 with INTEGRAL/IBIS and INTEGRAL/SPI have shown that the power law component is highly polarized, indicating that this component is associated with the relativistic jet outflows that have been seen in the radio band for this source (Laurent et al., 2011; Jourdain et al., 2012).

Figure 3.9b shows the COSI image of Cyg X-1 using the three days where Cyg X-1 entered COSI's field of view. Exposure to Cyg X-1 was not as optimal as for the Crab, so the detection significance is lower. Spectral and polarimetric studies of Cyg X-1 are forthcoming, and a comparison to the INTEGRAL results is anticipated.

3.5.3 Centaurus A

Centaurus A (Cen A) is a well studied active galactic nucleus located 3 - 5 Mpc away. The origin of the high energy emission from Cen A is not clear, but thermal Comptonization and synchrotron self-Compton (SSC) have both been suggested. These two mechanisms produce distinct polarization signatures: $\Pi \lesssim 10\%$ in the thermal Comptonization scenario (Matt,

1993), and Π up to 60% in the SSC case (Krawczynski, 2011a). Therefore, polarization measurements may help to identify the dominant emission mechanism for Cen A.

Figure 3.9c shows the COSI image of Cen A, using all 8 days where Cen A fell into the COSI FOV. Efforts are underway to determine if a polarization analysis will be feasible for Cen A using the COSI data.

3.5.4 Galactic positron annihilation emission

COSI successfully accomplished one of its main science goals by imaging the 511 keV emission from the galactic center. This source of gamma-rays results from the annihilation of a large flux of positrons with galactic electrons. The origin of these positrons, and their transport into the galactic center is a hotly debated topic (Prantzos et al., 2011), and it is believed that imaging of this emission will reveal essential clues. INTEGRAL/SPI observations at 511 keV indicate that the galactic center emission is diffuse (not a point source), and is accompanied by extended emission along the galactic disk, with possibly a galactic halo component (Bouchet et al., 2010).

Figure 3.9d shows the current best COSI image of the galactic center region at 511 keV, where a clear excess can be seen near the center. The offset from the center is probably not real but related to the aforementioned errors in absolute localization. Current work is focused on determining the extent of the emission, and confirming its diffuse nature. Also, spectral analyses with the COSI data reveal a strong detection ($\sim 10\sigma$) of the 511 keV line in the spectrum after the atmospheric 511 keV line has been subtracted, allowing for detailed studies of the line shape. This work is the subject of an upcoming publication (Kierans, 2017, in preparation).

3.5.5 Relativistic electron precipitation

COSI observed 3 episodes of REP activity throughout the flight. A light curve of the second REP event is shown in Figure 3.7. Fast, ~ 100 ms pulses were observed, as well as slow yet extremely bright pulses. The spatial distribution on the sky of REP activity is not known, and forthcoming imaging studies using the COSI data will shed light on this topic.

3.5.6 GRB 160530A

During the second REP event, COSI detected the bright, long duration gamma ray burst GRB 160530A. While the REP itself was quite luminous, the GRB was still clearly detected. This is clear from Figure 3.7, where the spike just after 07:00:00 in the Compton event light curve outshines the concurrent REP activity. A detailed polarization analysis of this GRB was recently published (Lowell et al., 2017a,b), and is the subject of Chapter 6. Work is ongoing to perform a spectral deconvolution of this GRB (Sleator et al., 2017), although preliminary results indicate good agreement between the COSI measured spectrum, and the spectrum measured by KONUS/Wind (Svinkin et al., 2016a)

3.6 Summary

The 2016 COSI SPB flight campaign from Wanaka, New Zealand was successful on many fronts. On the science end, COSI collected data continuously for 46 days and has so far positively identified five astrophysical sources of gamma-rays. Analysis efforts so far have focused on measuring the polarization of GRB 160530A, spectral analysis of GRB 160530A and the Crab Nebula, and imaging and spectroscopy of the 511 keV positron annihilation line from the galactic center. The data collected during the REP events is arguably the richest dataset of its kind. Future collaborations with scholars in the REP field are planned. Lastly, the background environment at float altitudes was sampled by COSI over a wide range of geographical positions, and efforts are underway to deconvolve the background distribution from the measured data. Such background information is crucial for planning science goals for future balloon flights.

With respect to the instrument and payload, the performance and reliability of the entire system exceeded our expectations. Although there are yet to be resolved issues with the high voltage anomalies, no further problems with any other systems on the gondola have been identified. The OP telemetry system worked well, and all data that was lost during normal operations was recovered at the end of the flight. A gentle landing and a successful recovery from the Atacama desert enhances the prospects for more balloon flights in the near future.

Finally, the 2016 campaign marked NASA's most rigorous test of the SPB technology to date. The long flight duration, highly variable thermal environment, and uncertain flight trajectory were challenges that were all faced and overcome by the SPB during the 2016 flight. Future work on SPB technology is centered on maintaining a reliable flight altitude, which has been challenging up to now.

Chapter 4

Compton Polarimetry with COSI

Astrophysical polarization measurements can reveal valuable information relating to the emission mechanism, geometrical configuration, and magnetization of a source, which may not be accessible via spectral, timing, imaging, or photometric analysis. The primary goal of a polarimeter is to measure the polarization level and angle of a beam. Secondary goals include measuring the time, energy, angular, or flux dependence of the polarization level and angle. In the soft gamma-ray band (0.1 MeV to 10 MeV), where the dominant cross-section for matter-light interactions is Compton scattering, Compton telescopes can deliver strong polarimetric performance owing to the fact that the Compton scattering cross-section is sensitive to the orientation of the photon’s electric field vector (Lei et al., 1997).

The standard approach to the polarization data analysis consists of generating a histogram of measured azimuthal scattering angles for qualifying events and fitting a “modulation curve” to the data. While this approach is simple and effective, it disregards information that can be used to further constrain the polarization properties of the beam, such as the Compton (polar) scattering angle, and the initial photon energy. Krawczynski (2011b) has shown that by combining the Compton scattering angle and photon energy measurements with the azimuthal scattering angle measurement in an unbinned, maximum likelihood analysis, the sensitivity of an ideal polarimeter is improved by $\sim 21\%$ over the standard approach.

In this chapter¹, we will start by giving an overview of the physical principles underlying Compton polarimetry, and then review the standard method and the maximum likelihood method for the polarization analysis. Along the way, we will discuss some important details pertaining to polarization analysis. In a few cases, the reader will be asked to refer to the following two chapters, which contain real world examples of polarization analysis.

¹This chapter is largely based on a recently published paper “Maximum Likelihood Compton Polarimetry with the Compton Spectrometer and Imager” (Lowell et al., 2017b). ©Astrophysical Journal 2017, reprinted with permission.

4.1 Principles of Compton Polarimetry

The Klein-Nishina equation gives the differential cross-section for Compton Scattering of photons on free electrons at rest:

$$\frac{d\sigma}{d\Omega} = \frac{r_0^2}{2} \left(\frac{E'}{E} \right)^2 \left(\frac{E}{E'} + \frac{E'}{E} - 2 \sin^2 \theta \cos^2 \eta \right), \quad (4.1)$$

where r_0 is the classical electron radius, E is the initial photon energy, E' is the scattered photon energy, θ is the Compton scattering angle (or polar scattering angle), and η is the azimuthal scattering angle defined such that $\eta = 0$ corresponds to scattering along the direction of the initial photon's electric field vector. After some algebraic manipulation of Equation 4.1, the probability density function (PDF) of scattering with a particular η takes the simple form of an offset cosine:

$$p(\eta; E, \theta) = \frac{1}{2\pi} [1 - \mu(E, \theta) \cos(2\eta)], \quad (4.2)$$

where $\mu(E, \theta)$ is the “modulation”:

$$\mu(E, \theta) = \frac{\sin^2 \theta}{\frac{E'}{E} + \frac{E}{E'} - \sin^2 \theta}. \quad (4.3)$$

E , E' , and θ are all related by the kinematic Compton scattering formula:

$$E' = \frac{E}{1 + \frac{E}{m_e c^2} (1 - \cos \theta)}, \quad (4.4)$$

where $m_e c^2 = 511$ keV. Figure 4.1 shows the dependence of μ on E and θ , where it is clear that μ is larger at lower energies, and for Compton scattering angles near $\sim 90^\circ$.

When a gamma-ray beam is polarized at a level of Π , where $0 \leq \Pi \leq 1$, then a fraction Π of the photons from the beam will have their electric field vectors aligned along a specific direction. The other $1 - \Pi$ fraction of the photons will have their electric field vectors randomly oriented. Thus, for a photon from a beam with polarization level Π and polarization angle η_0 , Equation 4.2 becomes:

$$p(\eta; E, \theta, \Pi, \eta_0) = \frac{1}{2\pi} [1 - \Pi \mu(E, \theta) \cos(2(\eta - \eta_0))]. \quad (4.5)$$

It is clear from Equation 4.5 that photons from a polarized gamma-ray beam will preferentially scatter such that $\eta - \eta_0 = +90^\circ$ or -90° . Equation 4.5 is plotted in Figure 4.2 for various values of Π and η_0 .




Figure 4.1: The modulation of the azimuthal scattering angle probability as a function of the Compton scattering angle for various photon energies. At lower energies and for Compton scattering angles near $\theta \sim 90$ degrees, the modulation is strongest.

4.1.1 COSI as a Compton Polarimeter

The COSI detector system is comprised of twelve cross-strip germanium detectors, which are capable of measuring interaction locations and energies with high accuracy. For polarization analysis, the relevant measurable quantities are the azimuthal scattering angle, the Compton scattering angle, and the total photon energy. COSI measures these quantities in the following ways: The azimuthal scattering angle η is determined geometrically, using the positions of the first and second interaction locations along with the known source position (determined by Compton imaging, or from previous knowledge of the source position). The total photon energy E is determined by simply summing all of the measured interaction energies. Finally, the Compton scattering angle is computed using Equation 4.4, where $E' = E - E_1$, and E_1 is the energy deposited in the first Compton scatter.

COSI departs from the ideal in several ways. Most importantly, the detection efficiency is not uniform for all η , even for specified values of θ and E . This is mostly a geometric effect, as the detector volume is not azimuthally symmetric about the incident photon beam. Thus, in directions where more detector material is present, the efficiency will be higher. Moreover,




Figure 4.2: The ideal probability density function (Equation 4.5) for the azimuthal scattering angle for various polarization levels Π and polarization angles η_0 . In these plots, the modulation is $\mu = 0.5$. Notice that even when $\Pi = 1$, there is a non-zero probability of scattering along the electric field vector.

intervening passive material will have the effect of lowering the efficiency. Another geometry related effect is the electrode strip segmentation and orientation. It is currently not possible to assign a depth to events where the same strip was hit more than once, so these events are excluded from the analysis. Excluding these events reduces the efficiency along directions where a photon would hit the same strip in a detector twice. Besides geometrical effects, there are other detector effects which can impact the efficiency. The most important of these is non-uniformity of energy thresholds across channels. Each strip of each COSI GeD has a threshold which is tuned independently depending on the channel gain and noise. This can impact the efficiency when photons scatter along directions which lead them to interact in channels with higher or lower than average energy thresholds.

Accounting for these effects through simulation is paramount to the success of polarization analysis. Without careful consideration of these instrumental systematics, spurious signals may arise and real signals may be suppressed. In our analysis, we use an accurate mass model of the COSI detector system which faithfully represents the distribution of active detector material as well as any passive mass. Additionally, all simulations are processed by the Detector Effects Engine, which adds noise, applies thresholds, performs inverse calibrations, and discretizes the data on a per channel basis.

4.1.2 Standard Analysis Method

The standard analysis method (SM) aims to fit a generalized form of Equation 4.5 to a measured azimuthal scattering angle distribution (ASAD). The fit function is

$$A - B\cos(2(\eta - \eta_0)), \quad (4.6)$$

where the offset A , the amplitude B , and the polarization angle η_0 are free parameters. The measured modulation $\hat{\mu}$ is

$$\hat{\mu} = \frac{B}{A}, \quad (4.7)$$

which can be converted to a polarization level using

$$\Pi = \frac{\hat{\mu}}{\mu_{100}}, \quad (4.8)$$

where μ_{100} is the “modulation factor” – the observed modulation in the case that $\Pi = 1$. For an ideal polarimeter at a specific photon energy and Compton scattering angle, μ_{100} is given by Equation 4.3. However, a real polarimeter typically must accept a range of energies and Compton scattering angles, in which case μ_{100} represents an averaging over Equation 4.3. Moreover, a real polarimeter may not be able to correctly reconstruct all Compton events, thus reducing the value of μ_{100} . It is usually not possible to determine μ_{100} experimentally for a given source, due to the difficulty of producing polarized beams with the correct angular distribution and spectral shape in the laboratory. Therefore, μ_{100} is estimated using simulations.

Some authors use a slightly different approach when fitting the ASAD: the measured ASAD is first normalized, and then Equation 5 is fit directly with Π and η_0 as the only free parameters. Conceptually, the difference between the two approaches is that the former treats the average counts per bin as a quantity to be fitted, while the latter approach uses the measured average counts per bin as a normalization constant. Both approaches should give practically the same answer when the number of counts in the ASAD is large, and possibly different answers when fewer counts are used. In this work, we treat the average counts per bin as a free parameter by using Equation 4.6 as our fit function.

Before the fits take place, there are three pre-processing steps:

1. Determine the event selections that give the lowest minimum detectable polarization (MDP). The MDP is given by

$$\text{MDP} = \frac{4.29}{\mu_{100}r_s} \sqrt{\frac{r_s + r_b}{t}}, \quad (4.9)$$

where r_s is the average source count rate, r_b is the average background count rate, t is the observation time, and the factor of 4.29 corresponds to a confidence level of 99% (Weisskopf et al., 2010). Note that this equation does not include systematic

errors, and so the true MDP must be obtained using MC simulations which have been carefully benchmarked using calibrations. Nonetheless, Equation 4.9 is a reasonable starting point for event selection optimization.

2. Make a background ASAD and subtract it from the source ASAD. For transient sources such as GRBs, data from before or after the event may be used. For persistent sources, the polarimeter must select data from a region of the data space which has similar characteristics, but does not coincide with the source.
3. Correct the background subtracted ASAD for systematics using an unpolarized ASAD. The unpolarized ASAD is usually obtained by an MC simulation which mimics the source under study with $\Pi = 0$. The correction is performed by first rescaling each bin of the unpolarized ASAD by the mean value, and then dividing each bin of the background subtracted ASAD by each mean-scaled bin of the unpolarized ASAD.

4.1.3 Maximum Likelihood Method

The goal of the maximum likelihood method (MLM) is to find the beam polarization level Π and angle η_0 that maximize the likelihood \mathcal{L}

$$\mathcal{L} = \prod_{i=1}^N p(\eta_i; E_i, \theta_i, \Pi, \eta_0) \quad (4.10)$$

where p is the conditional probability of measuring the azimuthal scattering angle η_i given that we have accurately measured the energy E_i and polar scattering angle θ_i of event i . For event lists longer than several hundred, \mathcal{L} can easily underflow a double precision floating point number. To mitigate this problem, the natural logarithm of the likelihood is used

$$\ln \mathcal{L} = \sum_{i=1}^N \ln p(\eta_i; E_i, \theta_i, \Pi, \eta_0). \quad (4.11)$$

The values of Π and η_0 that maximize $\ln \mathcal{L}$ also maximize \mathcal{L} , since the natural logarithm is a monotonically increasing function. A hat symbol is used to denote the optimal values, i.e. $\hat{\Pi}$ and $\hat{\eta}_0$.

For an ideal polarimeter, $p(\eta; E, \theta, \Pi, \eta_0)$ takes the simple form of Equation 4.5. However, for a real polarimeter, Equation 4.5 no longer holds due to the systematic effects of the detector system (Section 4.1.1). The complexity of the MLM thus lies in determining $p(\eta_i; E_i, \theta_i, \Pi, \eta_0)$ for each event i in such a way so as to include the instrument systematics. Here we outline a simulation based scheme for evaluating $p(\eta_i; E_i, \theta_i, \Pi, \eta_0)$:

1. Carry out a simulation of the instrument mass model subjected to an *unpolarized* gamma-ray beam with the same coordinates and spectrum as the source under study.

2. Define a three-dimensional histogram $\mathcal{H}[E, \theta, \eta]$ indexed by energy E , polar scattering angle θ , and azimuthal scattering angle η . Let the number of E , θ , and η bins be j , k , and l , respectively. This histogram will also be referred to as the “response.”
3. For the i^{th} simulated event, perform the event filtering and reconstruction, determine E_i , θ_i , and η_i , and increment the corresponding cell in $\mathcal{H}[E, \theta, \eta]$ by one.
4. The azimuthal scattering angle probability $p(\eta_i; E_i, \theta_i, \Pi, \eta_0)$ for a real event i can now be computed in the following way: take a one-dimensional slice along the η axis of $\mathcal{H}[E, \theta, \eta]$, and call this slice $g(\eta; E_{j'}, \theta_{k'})$, where j' is the index of the E bin containing E_i and k' is the index of the θ bin containing θ_i . Then the conditional PDF for η is

$$p(\eta; E_i, \theta_i, \Pi, \eta_0) = \frac{1}{A} \left[g(\eta; E_{j'}, \theta_{k'}) \times \frac{1}{2\pi} (1 + \Pi \mu(E_i, \theta_i) \cos(2(\eta - \eta_0))) \right], \quad (4.12)$$

where A is a normalization constant chosen so that the area under the total PDF is equal to unity. Equation 4.12 can then be evaluated at η_i to yield $p(\eta_i; E_i, \theta_i, \Pi, \eta_0)$.

Equation 4.12 is intuitively simple to understand; the slices $g(\eta; E, \theta)$ encode the effects pertaining to the instrument systematics, and the second term - which is just Equation 4.5 - is the ideal PDF. If this analysis were carried out with an ideal polarimeter, the slices $g(\eta; E, \theta)$ would be uniform in η , and Equation 4.12 would collapse to Equation 4.5. In essence, the $g(\eta; E, \theta)$ slices represent the acceptance as a function of η , and parameterized by E and θ .

An alternative scheme for determining the azimuthal scattering angle probability for each event would be to perform a series of MC simulations with various Π and η_0 , and interpolate the responses during the maximization of $\ln \mathcal{L}$ as Π and η_0 are varied. However, such an approach requires significantly more simulation time in order to achieve adequate statistics in every bin of each response. Contrast this with the approach outlined above, where only a single simulation of an unpolarized source is needed, and no interpolation based on the values of Π and η_0 is required. In our procedure, interpolation is avoided because the part of the PDF in Equation 4.12 that depends on Π and η_0 is analytic.

Figure 4.3 shows the total PDF for the azimuthal scattering angle for a photon with $E = 337.5$ keV and $\theta = 92.5$ degrees. The ideal PDF is overplotted for comparison. At this energy and Compton scattering angle, the modulation is relatively high. On the left, where $\Pi = 0$ (unpolarized), the ideal PDF is just a constant, so the full PDF is equivalent to $g(\eta; E, \theta)$. On the right, where $\Pi = 1$ (fully polarized), the ideal PDF is modulated, and so the full PDF is the normalized product of the modulated, ideal PDF (Equation 4.12) with $g(\eta; E, \theta)$. Clearly, the systematic effects of the detector system (Section 4.1.1) distort the PDF from its ideal shape. However, the structure of the ideal PDF still comes through in that where the ideal PDF has peaks, the probability is enhanced, and where the ideal PDF has troughs, the probability is suppressed. Note that the response used in Figure 4.3 is for the COSI observation of GRB 160530A, which occurred 43.5° off-axis.




Figure 4.3: The full PDF used for the MLM analysis (gray bars), along with the ideal PDF (blue and green lines) for two cases $\Pi = 0$ (left) and $\Pi = 1$ (right). The PDFs are drawn for a photon energy of 337.5 keV and a Compton scattering angle of 92.5 degrees. The slice $g(\eta; E, \theta)$ used for these PDFs is valid over the range $E = 325 - 350$ keV and $\theta = 90 - 95$ degrees. The slice $g(\eta; E, \theta)$ used here was taken from the COSI response $\mathcal{H}[E, \theta, \eta]$ for GRB 160530A, which occurred 43.5° off-axis. Error bars are drawn on the PDFs based on the simulation statistics.

In the presence of background, the probability in Equation 4.12 must be modified to include a term that represents the background probability distribution:

$$p_{\text{total}} = f \cdot p(\eta; E, \theta, \Pi, \eta_0) + (1 - f) \cdot p_{\text{bkg}}(\eta; E, \theta), \quad (4.13)$$

where $f = (T - B)/T$ is the signal purity, T is the total number of counts detected, B is the estimated number of background counts in the sample, and p_{bkg} is the probability of measuring the azimuthal scattering angle η , given that we have accurately measured the energy E and Compton scattering angle θ , and that the photon originated from a source of background. A straightforward approach for evaluating p_{bkg} is to generate a background response $\mathcal{H}_{\text{bkg}}[E, \theta, \eta]$ with the same binning as $\mathcal{H}[E, \theta, \eta]$, filled with measured background events or simulated background events. Each η slice of $\mathcal{H}_{\text{bkg}}[E, \theta, \eta]$ is then normalized so that the bin contents along the η axis represent probability densities. Finally, the background probability for event i can be looked up by retrieving the contents of the bin corresponding to η_i , E_i , and θ_i .

Once $\hat{\Pi}$ has been found, $\hat{\Pi}$ must be corrected to account for various imperfections of the detector system such as imperfect reconstruction efficiency and measurement error. Correcting for these effects amounts to determining $\hat{\Pi}$ in the case that $\Pi = 1$, and $B = 0$. The value of $\hat{\Pi}$ returned by the MLM algorithm under these conditions is referred to as the MLM correction factor, denoted as Π_{100} . The corrected polarization is then given by:

$$\Pi = \frac{\hat{\Pi}}{\Pi_{100}}. \quad (4.14)$$

For an ideal polarimeter capable of perfectly reconstructing all events with perfect precision, $\Pi_{100} = 1$. In reality, some events will be improperly reconstructed and yield a random value for η , which effectively reduces the measured polarization level. Additionally, the measurement error on the azimuthal scattering angle will also reduce the measured polarization level.

We used the MINUIT minimizer to determine $\hat{\Pi}$ and $\hat{\eta}_0$, and MINOS to determine the errors (James & Roos, 1975) for these parameters. Confidence contours were drawn in the 2D Π - η_0 space along paths of constant $2\Delta \log \mathcal{L}$, where $2\Delta \log \mathcal{L}$ is twice the difference between the maximum log likelihood and the log likelihood of a trial point. This quantity is asymptotically distributed as χ^2 (Wilks, 1938), so the confidence level corresponding to a particular value of $2\Delta \log \mathcal{L}$ can be calculated using a χ^2 distribution with two degrees of freedom.

However, the signal purity f has an associated uncertainty stemming from the Poisson distributions underlying T and B , which will create additional uncertainty on Π . The MINOS errors do not reflect this source of error, because f is held constant during the minimization. Moreover, it is also possible that Π_{100} will have a non-negligible uncertainty. To determine the total uncertainty on the measured, corrected polarization level Π , we approximated the probability distribution of Π by repeatedly simulating the observation. For each simulated observation, we bootstrap resampled the event list, drew a value of f from its associated

probability distribution, ran the minimizer to determine $\hat{\Pi}$ and $\hat{\eta}_0$, and then divided $\hat{\Pi}$ by a value of Π_{100} drawn from its associated probability distribution. The resulting distribution of Π from the simulated observations could then be analyzed numerically to determine confidence intervals or upper limits.

The MLM has two main advantages over the SM. First, more information is used per event. In the SM, only the azimuthal scattering angles of qualifying events are considered. In the MLM, the photon energy E and Compton scattering angle θ are considered as well. Effectively, each event's contribution to the likelihood statistic is implicitly weighted by Equation 4.3 (Figure 4.1), which is a function of E and θ . Second, for a realistic observation, the MLM can use more counts in the analysis. Consider that the first step in the SM is to optimize the statistical MDP (Equation 4.9). This amounts to choosing event selections on E and θ which yield as high a value of μ_{100} as possible, while still accepting (rejecting) as many source (background) counts as possible². In the MLM however, events with any energy or Compton scattering angle can be used. Therefore, events that were removed during the optimization of the SM analysis can now be included. Although these events generally represent points on the μ profile (Equation 4.3, Figure 4.1) corresponding to lower modulation, they are still meaningful contributors to the likelihood statistic.

We verified the functionality of the MLM by applying the algorithm to simulated data sets, where each simulation had a distinct level of polarization ranging from $\Pi = 0$ to $\Pi = 1$. Figure 4.4 shows the corrected, best fit polarization level $\hat{\Pi}/\Pi_{100}$ (top) and $\hat{\eta}_0$ (bottom) for each simulation run. In the top panel, a blue line with a slope of 1 is drawn for reference. These simulations used the position and spectrum of GRB 160530A, along with a full sky background based on Ling (1975). Events from the background simulation were added to the observation using a signal purity of $f = 0.72$, and a background response was generated from the entire background simulation in order to evaluate p_{bkg} . In the top panel, the relationship between Π and $\hat{\Pi}$ shows the expected linear behavior of Equation 4.14. As for the polarization angle, the true value in the simulations was $\eta_0 = 150$ degrees, and is denoted with a dashed line in the bottom panel of Figure 4.4. As expected, the best fit values for η_0 become more accurate as the simulated polarization level is increased. It is clear from Figure 4.4 that the algorithm yields reasonable estimates for Π and η_0 .

4.2 Modulation reducing effects

It is desirable for a polarimeter to have as high as modulation as possible, as this results in a better sensitivity. In practice, several effects act to reduce the modulation.

² E and θ are the most meaningful selections in this context, but other event parameters can and should be optimized as well.




Figure 4.4: The corrected, best fit polarization level $\hat{\Pi}/\Pi_{100}$ (top) and best fit angle $\hat{\eta}_0$ (bottom) as a function of the true, simulated polarization level. The blue line in the top panel is drawn with a slope of 1 for reference. These simulations include a background component based on Ling (1975), with a signal purity of $f = 0.72$. The simulated polarization angle was $\eta_0 = 150^\circ$ and is denoted with the dashed line in the bottom plot. The error bars shown in these plots are 1σ and were obtained using MINOS.

4.2.1 Event Selections

Equation 4.3 demonstrates that the modulation depends on the photon's energy and the Compton scattering angle. If event selections are chosen such that events with higher energies and Compton scattering angles not close to $\sim 90^\circ$ are used in the analysis, then the modulation will naturally be lower.

4.2.2 Event Reconstruction Efficiency

Events which are reconstructed incorrectly pose somewhat of a problem for Compton polarimeters. The most severe problem comes for 3+ site events, when the interaction sequence is predicted incorrectly. If the true first or second interaction is swapped with a subsequent interaction, then the azimuthal scattering angle will clearly be incorrect. If we assume that the improperly reconstructed event yields a random azimuthal scattering angle, then this event effectively resembles a properly reconstructed event that originated from the unpolarized fraction of the photon beam. Therefore, this effect will clearly reduce the modulation.

For 2-site events, which dominate the population of COSI events, this problem is somewhat less severe. The reason is that the azimuthal scattering angles of the two interaction permutations are related geometrically; they are offset by 180° . Therefore, if the azimuthal scattering angles are computed using $\eta \bmod 180^\circ$, and both event permutations pass the event selections, then reconstruction efficiency should not affect the modulation. In practice, the detector response is not symmetric over the 2π range of azimuthal scattering angles, and a wrong permutation may not pass the event selections. Additionally, even if the wrong permutation does pass the event selections, then the event will be assigned the wrong Compton scattering angle, which can have a negative impact on the analysis methods, especially the MLM.

4.2.3 Event Resolution with Finite Position Resolution

Non-ideal detectors have a finite position resolution, which limits the detector's ability to resolve interactions which occur too close together. The result is that interactions that happen too close together — within the position resolution of the detector — will be indistinguishable from a single interaction. This can have a negative impact on the polarization analysis, which is illustrated in the following example: consider a photon that originates from a polarized beam. When this photon reaches the detector system, it Compton scatters twice within the position resolution of the detector, and then is subsequently photoabsorbed at a location that is farther away relative to the position resolution. Due to the finite resolution of the detector, this event will appear to have been a 2-site event, even though three interactions were involved (two Compton scatters, and a photoabsorption). The azimuthal angle that we would like to measure is the azimuthal angle between the first two Compton scatters, but the azimuthal angle we actually measure will be the azimuthal angle between the position of the two Compton scatters and the position of the photoabsorption. It is known that Compton scattering alters the polarization state of the incoming beam such that an unpolarized beam can become partially polarized after Compton scattering, and a polarized beam becomes less polarized after Compton scattering (Depaola, 2003). Therefore, without information relating to the scattering parameters of the first Compton scatter, there is little we can say about the likely polarization state of the incident photon. Our inability to resolve the two Compton scatters amounts to a loss of polarization related information, which leads to a reduction in the modulation.




Figure 4.5: The Modulation reduction factor (MRF) as a function of the 1σ azimuthal scattering angle error $\delta\eta$. The red curve is computed using the discrete Fourier transform of the (assumed) normal distribution on η . The black circles come from toy Monte Carlo simulations. Both methods show that the MRF decreases with increasing $\delta\eta$, resulting in a reduced modulation.

4.2.4 Measurement Error on the Azimuthal Scattering Angle

Error on the azimuthal scattering angle, $\delta\eta$, has the effect of reducing the modulation. This is best understood in the context of linear systems theory; the ideal PDF in Equation 4.5 is a purely sinusoidal “signal.” If we were to convolve this PDF with some kernel that represents the measurement error, the result would be another sinusoid at the same frequency with possibly different phase and amplitude. If the kernel is assumed to be symmetric - a reasonable choice would be a Gaussian with a standard deviation of $\delta\eta$ - then the convolution will have no effect on the phase³. The reduction in amplitude can be computed by determining the magnitude of the discrete Fourier transform (DFT) of the noise kernel at the frequency of interest.

In general, $\delta\eta$ is determined by two factors: the position uncertainty of the first two interactions, and the distance between the first two interactions. Events whose first two interactions are separated by larger distances have more favorable values of $\delta\eta$, as small variations in the positions have less of an impact on η . Therefore, assuming that the position uncertainty is uniform across the detector system, the modulation will be reduced if events

³Convolution with a symmetric noise kernel is equivalent to filtering a signal with a symmetric impulse response. A filter with a symmetric impulse response is referred to as a zero phase filter, because the output phase is equal to the input phase.

with smaller distances between the first two interactions are used. The degree by which the modulation is reduced is referred to as the modulation reduction factor (MRF). The MRF modifies Equation 4.5 by reducing the amplitude of the cosine term:

$$p(\eta; E, \theta, \Pi, \eta_0) = \frac{1}{2\pi} [1 - \text{MRF}(\delta\eta)\Pi\mu(E, \theta) \cos(2(\eta - \eta_0))], \quad (4.15)$$

where the MRF ranges from zero to one, and we assume a symmetric noise kernel so that the phase of the cosine term is unaffected. Figure 4.5 shows the MRF as a function of the 1σ error on η . The red curve is obtained by evaluating the DFT of the Gaussian noise kernel at a frequency of $1/\pi$ cycles rad^{-1} (the frequency of the ideal PDF in Equation 4.5), and normalizing it by the DC value of the DFT. The normalization step results in a DC gain of 1, which preserves the area under the signal. As a check on this calculation, we performed toy Monte Carlo simulations, the results of which are indicated by the black circles in Figure 4.5. For each value of $\delta\eta$, we drew many samples from the PDF in Equation 4.5 with $\Pi\mu$ set to 1, and added random noise from a normal distribution with a standard deviation of $\delta\eta$. The noised samples were binned into 64 bins and fit with Equation 4.6, which yielded a modulation equal to the MRF. The toy MC simulations agree closely with the DFT method. For small values of $\delta\eta$, the MRF is close to 1, and the modulation is only marginally reduced. As $\delta\eta$ increases, the MRF decreases, thus reducing the observed modulation.

4.3 Minimum Detectable Polarization

The standard measure of a polarimeter's sensitivity is the minimum detectable polarization (MDP), which is the 99th percentile of the distribution of Π when the true value of Π is zero (unpolarized). In other words, it is the polarization level above which there is a 1% chance of measuring a polarization when the source is truly unpolarized. In the SM, the analytic expression for the MDP for statistical errors only is given by Equation 4.9. In the MLM, there is currently no known analytical expression for the MDP. Instead, we calculate the MDP in the MLM by forming many unpolarized trial observations using the same number of source and background counts that were actually observed. Unpolarized source counts are generated using simulations, and background counts may be either simulated or drawn from the observed background. The polarization level of each trial observation is determined using the corresponding values of f and Π_{100} . The 99th percentile of the distribution of Π from the trial observations is then used as the MDP. In order to do a side by side comparison of the MDP between the SM and the MLM, we also use the numerical approach to determine the MDP in the SM, rather than Equation 4.9. Each trial observation is analyzed with the SM, yielding a value of $\hat{\mu}$, which is re-scaled by μ_{100} to yield a polarization level (Equation 4.8). Note that Equation 4.9 gives the MDP when the only source of error is statistical error. Realistically, systematic error is frequently at play in polarization measurements. The numerical approach to determining the MDP has the benefit of including some systematic error, as it is based on realistic simulations of the detector system.

4.4 Summary and Conclusion

In this chapter, we have given an overview of the physical principles underlying Compton polarimetry. We also outlined a standard method and a maximum likelihood method for determining the polarization level and angle of a gamma-ray beam. In the following chapter, the SM will be used extensively to validate COSI's polarimetric performance. In Chapter 6, both the SM and the MLM are used to place upper limits on the polarization level for GRB 160530A. In fact, a literature search indicates that this is the first application of the MLM to a real data set.

The primary benefit of the MLM, which will be shown in more detail in Chapter 6, is that it results in a better MDP, providing for more sensitive polarization measurements. At the outset of this work, our expectation was that the MDP would improve by up to $\sim 21\%$, as Krawczynski (2011b) showed that the MLM improved the MDP by $\sim 21\%$ over the SM with an idealized Compton polarimeter at 100 keV, and under the assumption that the same number of counts was used in both the SM and MLM. However, as outlined in Section 4.1.2, the SM event selections must be optimized to give as low an MDP as possible. After the optimization, some counts will not pass the selections and thus won't be used in the SM analysis. In the MLM, the event selections for the energy and Compton scattering angle are wide open, so the MLM can use more counts and possibly yield an MDP improvement better than $\sim 21\%$. If the source spectrum is not monoenergetic, the relative performance between the SM and MLM will also depend on the energy spectrum of the source, and the effective area of the polarimeter at various energies. This is a consequence of the fact that the modulation of azimuthal scattering angles depends on energy (Equation 4.3, Figure 4.1).

While the MLM does indeed improve the MDP, the algorithm is significantly more time consuming to implement. Specifically, generating the polarization response for the MLM requires significantly more simulation time than generating the unpolarized ASAD for the SM due to the much larger number of bins in the MLM polarization response. However, we believe that the added effort is worthwhile, especially when the number of detected counts is small enough for statistical errors to dominate. In this work, our polarization response had three dimensions: total photon energy (36 bins), Compton scattering angle (36 bins), and azimuthal scattering angle (36 bins). Additional dimensions could be used to describe the response, such as the distance between interactions, or position within the detector system. Including more dimensions would probably improve the performance somewhat, but also substantially increase the number of bins and simulation time needed to fill the bins of the polarization response adequately.

Chapter 5

Validation of the Polarimetric Performance of the Compton Spectrometer and Imager

A shortcoming of astrophysical polarization measurements in the soft gamma-ray band is that many were made with space borne instruments whose polarization performance was not validated in the laboratory. Laboratory measurements using polarized beams can help to identify systematic errors, which can play a significant role in polarization analysis. In this chapter, we will compare data taken in the laboratory with partially polarized beams to simulations which mimic the experimental configuration. In addition to using polarized gamma-ray beams for the validations, we also compare measurements and simulations of unpolarized calibration sources, which can further validate our ability to capture the instrument systematics in the simulations.

5.1 Experimental Configuration

When unpolarized photons Compton scatter from a target, the outgoing beam is partially polarized with a polarization level of Π_U given by (Lei et al., 1997):

$$\Pi_U = \frac{\sin^2 \theta}{\epsilon + \epsilon^{-1} - \sin^2 \theta}, \quad (5.1)$$

where ϵ is the ratio of scattered photon energy to initial photon energy, and θ is the Compton scattering angle¹. The polarization vector of the scattered beam is perpendicular to the scattering plane, defined as the plane that contains both the initial and final direction vector of the photon. Using this principle, we produced partially polarized gamma-ray beams in the laboratory by scattering photons from a Sodium Iodide (NaI) scintillator. A major benefit of

¹Note that Equation 5.1 has the exact same form as Equation 4.3 for the modulation μ . Therefore, Figure 4.1 is also a plot of Π_U vs. Compton scattering angle for various energies.




Figure 5.1: A photograph showing the experimental scattering configuration on March 17, 2016. 662 keV photons originate from the Cs-137 source (lower left corner), Compton scatter in the NaI scintillator (top middle), and proceed to interact in the COSI GeDs. The Compton scattered beam is partially polarized, with the polarization vector perpendicular to the plane of scattering. The lead brick placed between the Cs-137 source and COSI blocks the 662 keV photons from interacting in the CsI shields.

using an instrumented detector (NaI) as the scatterer is that a large amount of background can be rejected if only events coincident between COSI and the NaI detector are considered. Background reduction is essential, as the coincident count rate of the scattering configuration is quite low ($< 1 \text{ count s}^{-1}$).

Figure 5.1 shows a photograph of the geometrical configuration used on March 17, 2016 during the pre-flight calibrations in Wanaka. The NaI detector was suspended above and to the side of the cryostat, while the Cs-137 source (662 keV) was placed at approximately the cryostat's level. A lead brick was placed between the Cs-137 source and the COSI detector system in order to block the direct flux from unnecessarily elevating the shield count rate. Two other permutations of this scattering setup were also used with different xy positions for the Cs-137 source, the lead brick, and the NaI detector. For each permutation, the source and NaI positions were chosen so as to select Compton scattering angles near $\sim 90^\circ$ with respect to the COSI position. This was done to get as high a polarization level in the scattered beam as possible (Equation 5.1). For 662 keV photons Compton scattering with $\theta = 90^\circ$, the outgoing photon energy will be 288 keV, and the polarization level will be

$\Pi = 58\%$. However, due to the size of the NaI scatterer, a range of Compton scattering angles will result. Therefore, we expect events coincident between COSI and the NaI detector to have energies clustered near 288 keV, forming a “scatter peak” in the energy spectrum. The spread in Compton scattering angle will also reduce the polarization level, since Compton scattering with angles higher or lower than $\theta \sim 90^\circ$ results in a lower polarization level.

Coincident events were identified by comparing timestamps between COSI events and events from the NaI detector. In order to time-tag the NaI events with respect to the COSI clock, we modified a spare flight computer and clock board to act as a slave (sGCU) to the master COSI GCU and clock board. The master clock and sync signal were used to synchronize the sGCU clock in the same way as the COSI card cages are synchronized. The procedure for synchronizing the sGCU clock (as well as the COSI card cages) is as follows:

1. GCU loads a 48 bit time stamp into the “sync-to” register of the sGCU via an ethernet connection.
2. GCU pulses the sync line synchronously with the 10 MHz clock signal.
3. Upon arrival of the sync pulse, the sGCU clock jumps to the “sync-to” value and is now synchronized with the GCU clock.

The signal from the NaI PMT was sent to a unipolar shaper with a shaping time of 500 ns. The output of the shaper was then sent to a discriminator whose threshold was set to approximately 100 keV. Finally, the output of the discriminator was sent to a digital input on the sGCU clock board for time stamping. A rising edge on this input latched the 48 bit clock value, which was then stored in a hardware queue for readout by the sGCU CPU. In this way, events in the NaI were time-tagged with respect to the COSI GCU clock at a resolution of 100 ns. A diagram illustrating the interconnection between various systems to perform these validations is shown in Figure 5.2.

This setup was convenient because foreknowledge of the coincidence window was not needed; the coincidence window could be determined during post-processing by looking at a histogram of time differences between COSI events and NaI events. Figure 5.3 shows one such histogram for the polarized run on 18 March, 2016. Each entry in this histogram represents the time between a COSI event, and the next NaI event. A coincidence peak can clearly be seen between 2300 ns and 2800 ns post COSI trigger. Accordingly, only COSI events with a NaI trigger occurring between 2300 ns and 2800 ns later were used for the analysis.

By only considering events coincident between COSI and the NaI detector, the background is significantly reduced. However, two background components still remain within the coincidence peak. The first component is chance coincidences, which occurs when two independent particles deposit energy in both the COSI and NaI detector at the same time. This component can easily be subtracted by selecting events from a time window with the same size as the coincidence window, but not overlapping the coincidence window. We used the time range of 3300 ns to 3800 ns to sample the chance coincidence event population.

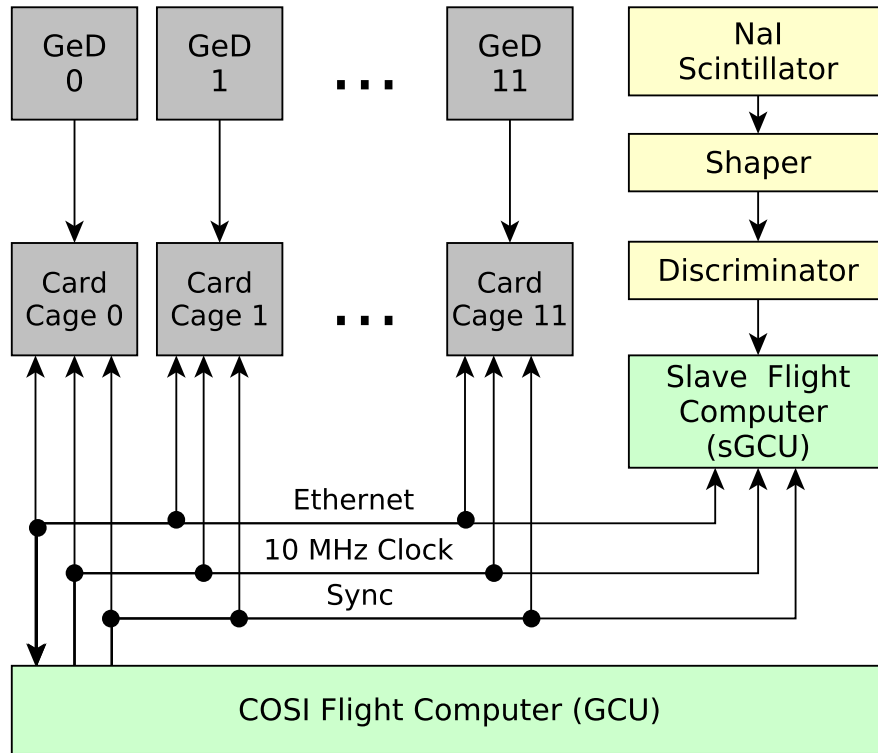


Figure 5.2: A diagram describing the interconnection between various components in the experimental configuration. The signal from the NaI scintillator is sent to a shaper, then a discriminator, and then finally to a digital input on the sGCU, which latches the 48 bit clock value when a rising edge is detected. The sGCU is connected to the system in parallel with the twelve card cages. This way, the sGCU clock is kept synchronized with the GCU and card cage clocks, so that timestamps from events in the NaI may be aligned with COSI events.

The second component is the true coincident background. This component represents events that did not originate from the Cs-137 source, but did deposit energy in both the COSI detectors and the NaI detectors. The most straightforward way to sample the true coincident component would have been to collect data with the Cs-137 source absent, and then select coincident events. Unfortunately, such data sets were not available, so we were unable to subtract this background component from our measurements.

5.2 Analysis Methods

Our goal for these tests was to match the measured modulation and polarization angle to simulations which mimicked the experimental configuration as closely as possible. Thus, we




Figure 5.3: A time difference histogram, where the x axis represents the time between a COSI event, and the time of the next NaI time stamp. A clear coincidence peak is seen between 2300 ns and 2800 ns. Chance coincident events were sampled between 3300 ns and 3800 ns.

used the standard polarization analysis method (SM) to analyze both the measurements and corresponding simulations.

Additionally, we employed two non-parametric statistical tests, the Kolmogorov-Smirnov (KS, Conover, 1980) and Anderson-Darling (AD, Anderson & Darling, 1952; Scholz & Stephens, 1987) tests, to test the hypothesis that the azimuthal scattering angle samples from the measurements and simulations were drawn from the same underlying distribution. In a KS or AD test, the empirical distribution function (EDF) is generated for both the measured and simulated sample of azimuthal scattering angles. Then, the test statistic is computed using the distances between the measured and simulated EDFs. In the KS test, the KS test statistic is equal to the single, largest distance between the two EDFs. The AD test statistic is a more complicated function of various distances between the EDFs. The AD test has been shown to have higher statistical power than the KS test (Hou et al., 2009; Stephens, 1974). Nonetheless, we used both methods to compute p -values for rejecting the null hypothesis that the measured and simulated samples of azimuthal scattering angles were drawn from the same distribution. Here, we set $p < 0.01$ as our threshold for statistical significance.

5.3 Polarized Runs

Three data sets corresponding to three different geometric configurations were collected in Wanaka during pre-flight calibrations. The data was collected on three consecutive days: March 16 (316), March 17 (317), and March 18 (318) of 2016. The difference between the three runs is that the x and y positions of the Cs-137 source, the lead brick, and the NaI detector were varied for each run. In all three cases, the polarization vector was approximately parallel to the xy plane. Figure 5.4 shows a schematic diagram of the three scattering geometries.

The geometries for the polarized runs were modeled closely after the real world configuration. In particular, the positions of the NaI detector, the Cs-137 source, and the lead brick were carefully measured and transcribed into mass models. We then simulated unpolarized, 662 keV photons originating from the location of the Cs-137 source. We reduced the simulation time significantly by only simulating the photons within a cone that fully overlapped with the NaI detector. This way, time was not wasted simulating photons that would have never interacted in the NaI detector. In theory, it is possible that a photon with an initial direction outside of the simulated cone beam could scatter in some passive material, then scatter in the NaI, and subsequently be detected by COSI. However, as will be shown shortly, the final event selections include a selection on the scatter peak energy, which would kinematically prohibit the possibility of accepting a photon that scatters in the aforementioned way.

To correct for geometric and systematic effects, we generated unpolarized azimuthal scattering angle distributions (ASADs) for each run by simply using a physics list in the simulations which does not include the polarization dependencies of the electromagnetic interaction cross sections. In doing this, the outgoing scattered beam from the NaI detector will behave like an unpolarized beam. This approach is the simplest, and appears to give reasonable results. However, it is not ideal, since even photons that originate from an unpolarized beam can become “polarized,” in the statistical sense, after Compton scattering. This can affect the instrument response for 3+ site events, since these events have more than a single Compton scatter. However, most COSI Compton events are 2-site events, so it is unlikely that ignoring polarization effects will have a strong impact on the unpolarized ASAD.

5.4 Event Selections

5.4.1 Energy Selection

In both the real data and the simulations, only events with energy within the “scatter peak” were accepted for the polarization analysis. The energy and shape of the scatter peak depends strongly on the geometrical configuration, as well as the energy of the source. For the runs considered here, the scatter peak was close to ~ 294 keV for all three runs. Figures




Figure 5.4: Top views (parallel perspective) of the three mass models representing the polarized runs on March 16, 17, and 18, 2016. Each configuration uses a different x and y position for the NaI detector, the Cs-137 source, and the lead brick. Each run also had a different z position.

5.5a, 5.5b, and 5.5c show measured and simulated scatter peaks for the 316, 317, and 318 runs respectively. We performed a rough Gaussian fit to the measured scatter peaks to determine their locations and widths. These parameters appear in Tables 5.1, 5.2, and 5.3 under “scatter peak energy”. Events that were detected within $\pm 1\sigma$ of the centroid energy were considered to pass the energy selection. From Figures 5.5a, 5.5b, and 5.5c, it is clear that the energy spectrum matches closely between measurement and simulation for energies that are not too far away from the scatter peak. We note that the shape of the scatter peak for the 316 run appears to have more structure, and is not well modeled by a Gaussian. Nonetheless, we did a rough Gaussian fit to the entire structure to determine a reasonable energy selection.

5.4.2 Distance Selection

Two sets of distance cuts were used. For the first set, we accepted events with a minimum distance between first two interactions of 0.5 cm, and a minimum distance between any two interactions of 0.3 cm. This first set of distance cuts will be referred to as d_1 . The second set of cuts were more restrictive with a minimum distance between first two interactions of 1.0 cm, and a minimum distance between any two interactions of 0.5 cm. This second set of distance cuts will be referred to as d_2 .


5.4.3 ARM Selection

Events with $|\text{ARM}| < 20^\circ$ were accepted for the polarization analysis. The ARM was computed with respect to the peak positions in the images. The images used for determining the peaks were simple back-projections, with no iterations of the LM-MLEM algorithm applied.


5.5 Results with d_1 Distance Selections

The results for the 316, 317, and 318 runs using the d_1 distance cuts are summarized in Figures 5.6, 5.7, 5.8 and Tables 5.1, 5.2, 5.3 respectively. In the figures, the top panel shows the ASAD after subtracting the chance coincident background. The center panel shows the unpolarized ASAD obtained from a simulation using a physics list which ignores polarization. Lastly, the bottom panel shows the corrected ASAD along with the best fit modulation curve.


Overall, the simulations resemble the measurements acceptably well. The fitted modulations and polarization angles from the measurements are within 1 or 2 error bars of the simulated values. The goodness of fit, as indicated by the value of χ_{red}^2 , is good for 318 ($\chi_{\text{red}}^2 = 0.98$) and acceptable for 316 ($\chi_{\text{red}}^2 = 1.71$) and 317 ($\chi_{\text{red}}^2 = 1.82$). In addition to the polarization parameters, we also computed p -values using the KS and AD tests to test the hypothesis that the measured and simulated samples of azimuthal scattering angles were drawn from the same underlying distribution. The p -values were acceptable for 316 and 318,



(a) 316



(b) 317



(c) 318

Figure 5.5: Measured (blue) and simulated (red) energy spectra for the 316 (top), 317 (middle), and 318 (bottom) runs. The shape of the measured scatter peak is in good agreement with the simulated scatter peak.

but marginal for 317. In addition to the low p -values and high χ_{red}^2 for 317, the mismatch in the polarization angle is larger than that found in the 316 and 318 runs. All things considered, we conclude that the simulation corresponding to the 317 run was not modeled with the same degree of geometrical precision as the 316 and 318 runs.

There appears to be a trend whereby the fitted modulations from the simulations are somewhat lower than the measured modulations. We investigated two possible sources for this discrepancy. First, we extracted the Compton scattering angle distributions for the runs; these are shown in Figures 5.9a, 5.9b, and 5.9c. If it were the case that the Compton scattering angle distributions for the measurements were skewed such that the peak was closer to $\sim 90^\circ$ - where the ideal modulation peaks - then this could increase the measured modulation over the simulation. However, as is clear in Figures 5.9a, 5.9b, and 5.9c, no such skewing is observed. Second, we compared the distribution of distances between first and second interactions for the measured and simulated data. The distributions are shown in Figures 5.10a, 5.10b, 5.10c. For smaller distances, the error on the azimuthal scattering angle is larger, which has the effect of reducing the modulation. Indeed, it appears that events with distances $\lesssim 1$ cm are overrepresented in the simulation. This is most evident for the 316 and 318 runs. Interestingly, these two runs have a larger mismatch in the modulation, whereas the 317 run has a tighter match.

5.6 Results with d_2 Distance Selections

In light of this observation, we performed the analysis again with the more restrictive d_2 distance selections. The analysis results are summarized in Figures 5.11, 5.12, 5.13 and Tables 5.4, 5.5, and 5.6 for the 316, 317, and 318 runs, respectively. Now, with the d_2 selections, the quality of the fits is significantly better. The χ_{red}^2 values are lower and closer to 1 for all three polarized runs. Moreover, the measured and simulated modulations appear to be in somewhat better agreement. The largest discrepancy in the modulation is seen for the 318 run, where the modulation from simulations is 11% lower than the measured modulation. As was seen with the d_1 selections, the 317 run has the highest value of χ_{red}^2 (1.36), as well as statistically significant KS and AD p -values.




Figure 5.6: Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 316 run using the d_1 distance selections.

	Measurement	Simulation
Coincident counts	29697	42636
Chance coincident counts	1057	—
Integration time	48199.880	62911.418
Amplitude (counts)	504.8 ± 20.1	672.8 ± 26.1
Offset (counts)	1809.9 ± 14.1	2693.2 ± 18.3
Angle (degrees)	177.92 ± 1.06	178.74 ± 1.04
Modulation	0.279 ± 0.016	0.250 ± 0.014
χ^2_{red} (dof = 13)	1.71	0.76
Scatter peak energy (keV)	294.5 ± 28.6	293.3 ± 28.1
KS p -value	0.18	
AD p -value	0.22	

Table 5.1: Statistics and SM fit results for the 316 run using the d_1 distance selections.




Figure 5.7: Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 317 run using the d_1 distance selections.

	Measurement	Simulation
Coincident counts	12770	38369
Chance coincident counts	382	—
Integration time	62813.404	190929.115
Amplitude (counts)	230.4 ± 10.3	$693.6 \pm 192.$
Offset (counts)	775.5 ± 7.5	2407.3 ± 14.0
Angle (degrees)	3.08 ± 1.33	0.30 ± 0.82
Modulation	0.297 ± 0.020	0.288 ± 0.012
χ^2_{red} (dof = 13)	1.82	1.87
Scatter peak energy (keV)	293.4 ± 20.3	293.4 ± 19.2
KS p -value	0.005	
AD p -value	0.016	

Table 5.2: Statistics and SM fit results for the 317 run using the d_1 distance selections.






Figure 5.8: Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 318 run using the d_1 distance selections.

	Measurement	Simulation
Coincident counts	26265	27374
Chance coincident counts	762	—
Integration time	54345.328	65319.694
Amplitude (counts)	498.1 ± 18.2	468.6 ± 18.8
Offset (counts)	1613.1 ± 13.1	1727.7 ± 13.5
Angle (degrees)	1.11 ± 1.03	3.43 ± 1.12
Modulation	0.309 ± 0.016	0.271 ± 0.016
χ^2_{red} (dof = 13)	0.98	1.19
Scatter peak energy (keV)	295.6 ± 25.9	294.5 ± 25.9
KS p -value	0.14	
AD p -value	0.21	


Table 5.3: Statistics and SM fit results for the 318 run using the d_1 distance selections.



(a) 316




(b) 317




(c) 318


Figure 5.9: Measured (blue) and simulated (red) Compton scattering angle distributions for the 316 (top), 317 (middle), and 318 (bottom) runs. The measured and simulated distributions match well for all three runs.



(a) 316



(b) 317



(c) 318

Figure 5.10: Measured (blue) and simulated (red) distribution of distances between first two interactions for the 316 (top), 317 (middle), and 318 (bottom) runs. Events with a distance between first two interactions of $\lesssim 1$ cm appear to be overrepresented in the simulations, especially for the 316 and 318 runs.




Figure 5.11: Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 316 run using the d_2 distance selections.

	Measurement	Simulation
Coincident counts	20765	28926
Chance coincident counts	738	—
Integration time	48199.880	62911.418
Amplitude (counts)	329.0 ± 16.5	461.3 ± 21.0
Offset (counts)	1252.2 ± 11.7	1809.4 ± 14.9
Angle (degrees)	177.88 ± 1.38	177.76 ± 1.26
Modulation	0.263 ± 0.019	0.255 ± 0.017
χ^2_{red} (dof = 13)	1.10	0.85
Scatter peak energy (keV)	293.5 ± 28.5	293.6 ± 27.9
KS p -value	0.29	
AD p -value	0.30	

Table 5.4: Statistics and SM fit results for the 316 run using the d_2 distance selections.




Figure 5.12: Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 317 run using the d_2 distance selections.

	Measurement	Simulation
Coincident counts	9297	27196
Chance coincident counts	270	—
Integration time	62813.699	190929.115
Amplitude (counts)	158.4 ± 8.8	479.6 ± 16.2
Offset (counts)	560.7 ± 6.4	1692.6 ± 11.7
Angle (degrees)	3.90 ± 1.63	0.58 ± 0.99
Modulation	0.283 ± 0.023	0.283 ± 0.014
χ^2_{red} (dof = 13)	1.36	1.29
Scatter peak energy (keV)	293.1 ± 21.0	293.2 ± 19.4
KS p -value		0.001
AD p -value		0.003

Table 5.5: Statistics and SM fit results for the 317 run using the d_2 distance selections.




Figure 5.13: Chance coincidence subtracted ASAD (top), unpolarized ASAD from simulations (middle), and the corrected ASAD with best fit modulation curve (bottom) for the 318 run using the d_2 distance selections.

	Measurement	Simulation
Coincident counts	18785	18862
Chance coincident counts	523	—
Integration time	54345.328	65319.718
Amplitude (counts)	365.6 ± 15.3	336.3 ± 15.4
Offset (counts)	1144.5 ± 11.1	1181.5 ± 11.1
Angle (degrees)	0.69 ± 1.19	2.69 ± 1.30
Modulation	0.319 ± 0.019	0.284 ± 0.019
χ^2_{red} (dof = 13)	0.93	0.70
Scatter peak energy (keV)	295.4 ± 26.1	294.4 ± 26.1
KS p -value	0.47	
AD p -value	0.65	

Table 5.6: Statistics and SM fit results for the 318 run using the d_2 distance selections.

5.6.1 Discussion

Figure 5.14 shows a summary of the fitted modulations, fitted polarization angles, and χ_{red}^2 values for the measurements and simulations corresponding to the 316, 317, and 318 runs, with both the d_1 distance selections (three leftmost entries) and the d_2 selections (three rightmost entries). The bottom two panels of Figure 5.14 show the differences between the measured and simulated modulations and polarization angles. In each case, a dashed line is drawn at zero. The distance from zero is an indication of how well the measurement and simulation agree. For each of the two distance selection criteria, we performed a meta-analysis using a χ^2 test to test the hypothesis that the differences between measurements and simulations were consistent with zero across the three runs. With respect to the fitted modulations, we calculate $p = 0.18$ (d_1) and $p = 0.62$ (d_2) for $\text{dof} = 3$. With respect to the fitted polarization angles, we calculate $p = 0.12$ (d_1) and $p = 0.23$ (d_2) for $\text{dof} = 3$. None of these p -values are low enough to indicate that the differences in fitted modulations and polarization angles are statistically significant.

While the agreement in fit parameters may be good using either the d_1 or the d_2 distance selections, the goodness of fit is clearly improved when using the d_2 distance selections. This can be seen in the middle panel of Figure 5.14, where the values of χ_{red}^2 are closer to 1 using the d_2 selections. It is not surprising that the d_2 selections show some benefit over the d_1 selections; one area that is currently lacking in the Detector Effects Engine is a detailed handling of interactions that occur very close to one another. Specifically, charge sharing, charge loss, and cross-talk are currently not modeled. Therefore, it is expected that the simulations should differ somewhat from the measurements when it comes to events with small distances between interactions. While the d_1 distance selections give acceptable results, future work should focus on improving the modeling for events with small separation distances. Such improvements will benefit all aspect of COSI data analysis, not just the polarization analysis.

Seeing as the d_2 selections are more restrictive, fewer counts will remain for the analysis when the d_2 selections are used instead of the d_1 selections. For the datasets presented in this chapter, plenty of counts are available to start with, so this is not of much concern. However, for real observations such as our observation of GRB 160530A presented in Chapter 6, this poses a problem, as the number of detected counts is modest.

5.7 Azimuthal Scattering Angle Distributions of Unpolarized, Calibration Sources

As another check on COSI's polarimetric performance, we analyzed several calibration data sets using radioactive sources placed within the field of view. Unlike the data sets analyzed in the previous sections, we did not utilize a scatterer to create partially polarized beams. Instead, we aimed to simply compare the measured and simulated samples of azimuthal scattering angles using the KS and AD tests in order to accept or reject the null hypothesis




Figure 5.14: A summary of the fitted parameters and their differences between measurements and simulation. Each point along the x axis represents one of the three polarized runs with either the d_1 or d_2 selections. In the top three plots, blue circles indicate measurements, and red circles indicate simulations (simulation circles have been offset for comparison). The top two plots show the fitted modulations and polarization angles, respectively. The middle panel shows the χ^2_{red} values for each run. The bottom two panels show the difference (measurement - simulation) in the fitted modulations and polarization angles, with the dashed line at zero drawn as a reference. A χ^2 meta-analysis of the differences suggests that we cannot reject the null hypothesis that there is no difference between measurements and simulations.

that the measured and simulated samples came from the same underlying distribution. While the underlying distribution won't be modulated for unpolarized sources, there is still a fair deal of structure due to various geometric and systematic effects (Section 4.1.1).

5.7.1 Event Selections

The event selections used for these measurements and simulations were as follows:

1. $\pm 1\sigma$ of the photopeak energy. The width of the photopeak was obtained by performing a simple Gaussian fit to the photopeak.
2. ± 1 FWHM of the ARM distribution. The ARM distribution was generated for the photopeak events, and the FWHM of the distribution was determined numerically.
3. The analysis was carried out with both the d_1 and d_2 distance selections.

5.7.2 Discussion

Luckily, collecting data with a calibration source requires very little effort compared to the polarized beam configurations. As a result, many such data sets were available for analysis, whereas only three data sets were available with real, polarized beams. The unpolarized data sets are actually quite useful, as they simply generate an azimuthal scattering angle distribution in the detectors. While this distribution is not modulated, it still contains information relating to the instrument systematics. It should also be clear that the geometrical configuration for the unpolarized runs is far simpler than that used for the polarized runs. Consequently, it is less likely that the mass model will contain geometric errors for the unpolarized runs.

Tables 5.7 and 5.8 show the results for four different gamma-ray lines at several locations within the COSI field of view using both the d_1 and d_2 selections. After running KS and AD tests on the measured and simulated η samples (unbinned) for various unpolarized data sets, we found no cases with $p < 0.01$ so we cannot reject the null hypothesis that the measured and simulated distribution of azimuthal scattering angles are the same. This can be taken as further evidence that COSI is performing well in the polarimetric sense, and that the instrument systematics are accurately captured in the simulation pipeline. While the KS and AD tests were performed on the unbinned samples, we also binned the samples into ASADs as a visual aid. These are shown in Figure 5.15, where good agreement can be seen between measurement and simulation. Additionally, Figure 5.15 illustrates that the unpolarized ASADs are non-uniform and thus useful probes of the instrument systematics.

5.8 Conclusions

Overall, COSI's polarimetric capabilities with respect to the data sets considered here is well reflected in the simulations. We showed that for the polarized runs, the differences between

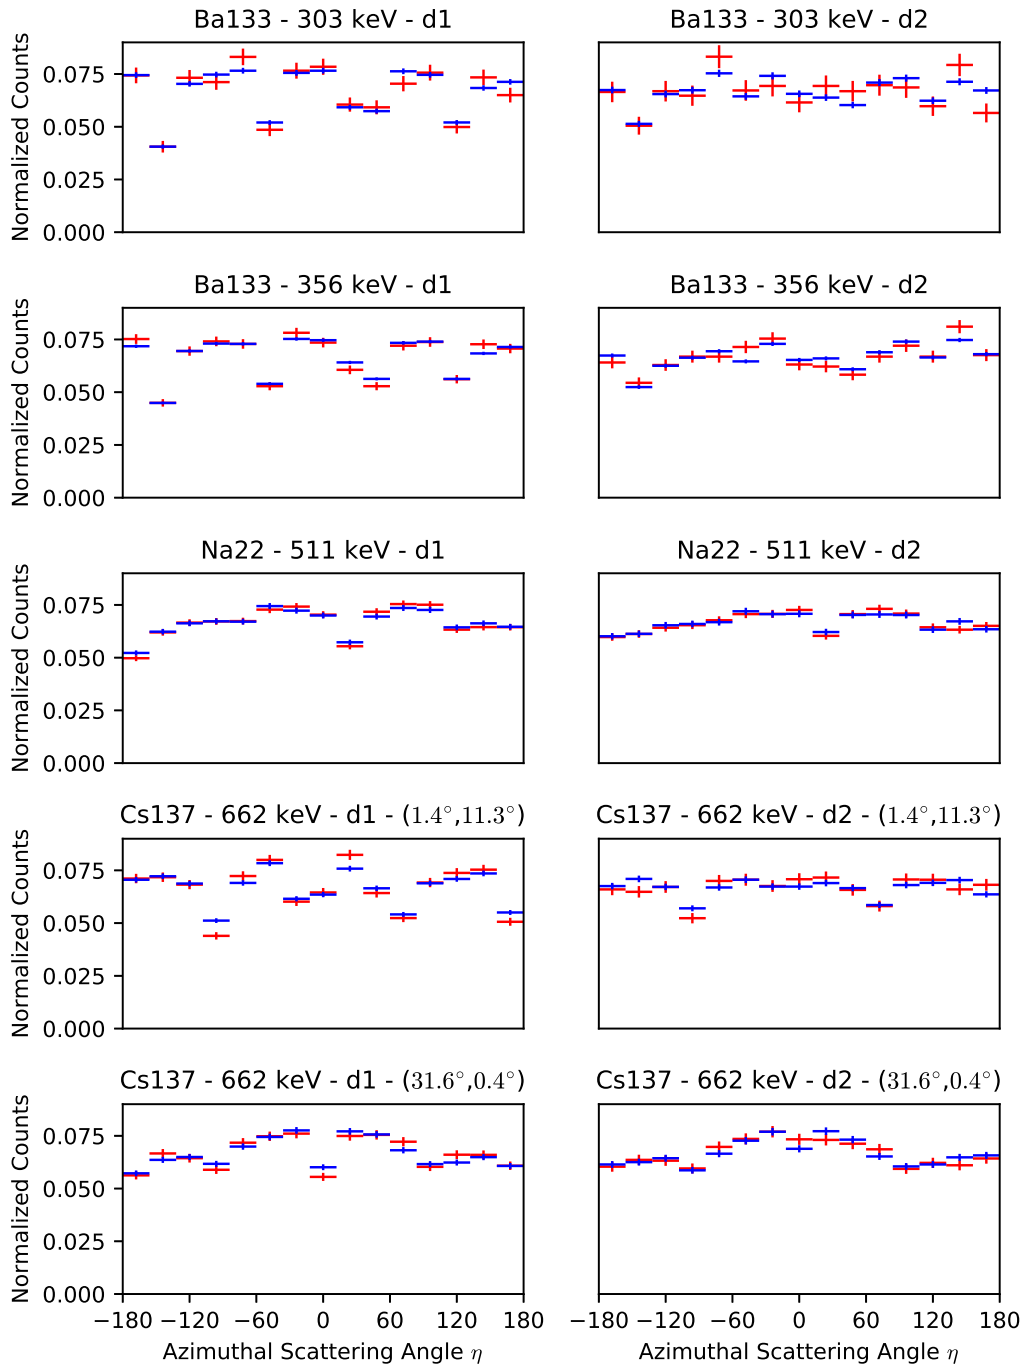


Figure 5.15: Binned ASADs corresponding to the unpolarized runs in Tables 5.7 and 5.8. Measurements are drawn in blue, and simulations are drawn in red. The distributions are clearly non-uniform, which makes them useful for probing the instrument systematics.

Source	Energy (keV)	Location (θ, ϕ)	Photopeak energy (keV)		ARM FWHM		KS p -value	AD p -value
			Measured	Simulated	Measured	Simulated		
Ba-133	303	(7.7°, -127.9°)	302.8 ± 1.6	302.6 ± 1.7	13.0°	10.1°	0.81	0.69
Ba-133	356	(7.7°, -127.9°)	355.9 ± 1.6	355.7 ± 1.6	10.4°	8.0°	0.54	0.25
Na-22	511	(38.8°, -106.8°)	510.3 ± 2.0	510.7 ± 1.7	9.5°	7.9°	0.94	0.63
Cs-137	662	(1.4°, 11.3°)	661.2 ± 2.0	661.7 ± 1.8	7.1°	7.1°	0.36	0.40
Cs-137	662	(31.6°, 0.4°)	661.5 ± 1.8	661.7 ± 1.7	7.2°	7.2°	0.31	0.42

Table 5.7: Results for unpolarized runs using d_1 distance selections.

Source	Energy (keV)	Location (θ, ϕ)	Photopeak energy (keV)		ARM FWHM		KS p -value	AD p -value
			Measured	Simulated	Measured	Simulated		
Ba-133	303	(7.7°, -127.9°)	302.6 ± 1.5	302.5 ± 1.6	12.1°	8.6°	0.75	0.51
Ba-133	303	(7.7°, -127.9°)	355.8 ± 1.6	355.6 ± 1.6	9.2°	7.4°	0.48	0.59
Na-22	511	(38.8°, -106.8°)	510.1 ± 2.0	510.8 ± 1.8	7.9°	7.3°	0.99	1.05
Cs-137	662	(1.4°, 11.3°)	661.1 ± 1.8	661.7 ± 1.8	6.6°	6.3°	0.06	0.12
Cs-137	662	(31.6°, 0.4°)	661.4 ± 1.8	661.6 ± 1.9	6.5°	6.3°	0.45	0.47

Table 5.8: Results for unpolarized runs using d_2 distance selections. Note that the AD p -value for the third entry is greater than one; this a result of the algorithm performing an extrapolation. Nonetheless, this indicates that the p -value is high and far from the threshold of statistical significance.

the measured and simulated modulations and polarization angles were not statistically significant. However, the KS and AD p -values indicate that there are statistically significant differences in the distribution of azimuthal scattering angles for the 317 run, using both the d_1 and d_2 distance selections. It is not entirely clear why this is, but it is certainly possible that geometrical measurements which were transcribed into the mass model for the 317 run may not accurately reflect the true experimental configuration. Previous work on polarization calibrations of NCT used a theodolite for highly accurate measurement of angles between the Cs-137 source, lead bricks, scatterer, and cryostat (Bellm, 2011). Due to a lack of time, we were not able to perform such detailed localizations of all components involved in the polarized runs. Instead, measurements were made using linear measurements, which are probably not as reliable as theodolite localizations. However, it is encouraging that the energy of the scatter peak matches closely between measurements and simulations, as this parameter depends strongly on the scattering geometry.

We also compared samples of azimuthal scattering angles from unpolarized calibration sources in the field of view. The geometric configuration for these runs was far simpler than for the polarized runs, thus reducing the chances of geometric modeling errors in the mass model. For the unpolarized runs we considered, we found no statistically significant differences between measurements and simulations using the KS and AD tests.

If we consider the results from the polarized and unpolarized runs, and assume that the 317 polarized run suffered from geometric modeling errors, then we conclude that COSI’s polarimetric performance is not dominated by systematic errors when the number of detected counts is of the order 10^4 . This is roughly the number of counts within a factor of several in each measured data set presented in this chapter.

Computation of the modulation factor μ_{100} for the polarized runs was not performed here, and thus we do not quote values for the absolute polarization level for the polarized runs. Determining μ_{100} for the polarized runs is actually quite complex; unlike the case for a far field source simulation, where the polarization level and angle of the beam can be specified exactly in the simulation configuration file, the polarization of the beam in the polarized runs arises naturally from scattering in the simulation. Additionally, due to the near field geometry of the simulation, the polarization vector will vary from photon to photon, as its orientation depends on the plane of scattering. Therefore, there is no simple way of setting the polarization state of photons after they have scattered from the NaI in the simulation. One option would be to do a two step simulation. The first step would correspond to our simulations that we have done in this chapter for the polarized runs. The second step would be to loop through the ideal interaction information from the first simulation, identify photons that Compton scatter in the NaI, and set their polarization vector to be perpendicular to the plane of scattering. These “manipulated” photons would be the input to a second simulation. However, our goal for the polarimetric validation was not to determine the absolute polarization level, but rather to verify that there is good agreement between the measured and simulated distributions and fit parameters.

Nonetheless, it is possible to estimate μ_{100} for the polarized runs, which can then be used to place an upper limit on the systematic error for the polarization level $\delta\Pi_{\text{sys}}$. As an

example, we will illustrate this calculation for the 318 polarized run using the d_1 distance selections, and then quote the results for the other runs. The first step is to calculate an upper limit on the polarization level of the scattered beam. We do this by drawing samples from the observed scatter peak shape (normal distribution with mean = 295.6 keV and sigma = 25.9 keV, only accepting samples within 1 sigma of 295.6 keV), calculating the Compton scattering angle using Equation 4.4 and an initial energy of 662 keV, and then computing the outgoing polarization level using Equation 5.1. Averaging over the calculated polarization levels, we compute $\Pi_{\max} = 58\%$. This assumes that the scattering planes of all the photons are aligned, which is not the case due to the near field geometry of the source. Therefore, the true polarization level will be lower than Π_{\max} . This allows us to calculate $\mu_{100,\min} = \mu/\Pi_{\max} = 0.319/0.58 = 0.55$, which is a lower limit for μ_{100} . Finally, if we assume that the magnitude of the systematic error is less than the statistical error $\delta\mu_{\text{sys}} < \delta\mu_{\text{stat}} = 0.019$, then we can place an upper limit on the polarization level systematic error of $\delta\Pi_{\text{sys}} < \delta\mu_{\text{stat}}/\mu_{100,\min} = 0.019/0.550 = 3.5\%$. Table 5.9 tabulates this calculation for the three polarized runs and for both distance selections. The limits on the systematic error for the d_2 selections are slightly less constraining, due to fewer counts (more restrictive selections) and larger statistical errors. The quoted upper limits are conservative, since their calculation assumes that the statistical and systematic errors are of the same magnitude, while we demonstrated in this chapter that our errors are dominated by statistics.

In the future, when COSI or a related instrument is undergoing similar calibration procedures, it would be prudent to take a few additional steps. As was just mentioned, more precise localizations of the components involved in the scattering geometry would be helpful. Additionally, it would be useful to collect data in a geometrical configuration where the NaI detector location stays the same, but the position of the Cs-137 is rotated so that the scattering plane rotates, resulting in different polarization angles. With the data that is currently available, the polarization vector is always parallel to the xy plane. Finally, as discussed in Section 5.1, these data sets include an unsubtracted background component consisting of background photons that deposit energy in both the COSI detectors and the NaI. It is probable that the event selections — $\pm 1\sigma$ of the scatter peak energy, and ARM within $\pm 20^\circ$

Run	μ	$\delta\mu_{\text{stat}}$	Π_{\max}	$\mu_{100,\min}$	$\delta\Pi_{\text{sys}}$
316- d_1	0.279	0.016	58%	0.481	< 3.3%
317- d_1	0.297	0.020	58%	0.512	< 3.9%
318- d_1	0.309	0.016	58%	0.533	< 3.0%
316- d_2	0.263	0.019	58%	0.453	< 4.2%
317- d_2	0.283	0.023	58%	0.488	< 4.7%
318- d_2	0.319	0.019	58%	0.550	< 3.5%

Table 5.9: Upper limits on the systematic error for the polarization level for the three polarized runs and for both distance selections. The methodology for the calculation can be found in Section 5.8.

from the image peak location — will significantly reduce the impact of this component on the analysis. Moreover, the overall agreement between measurement and simulations indicates that this component is not playing a significant role. However, future calibration runs should make an effort to measure the coincident background, which would require collecting data in the absence of a radioactive source. This way, this background component can be properly subtracted.

Chapter 6

Polarization Analysis of GRB 160530A with the Compton Spectrometer and Imager

In Chapter 4 we reviewed the fundamentals of Compton polarimetry, and in Chapter 5 we benchmarked COSI’s polarimetric performance against simulations. In this Chapter¹, we will present our polarimetric analyses of a long duration gamma-ray burst detected by COSI during the 2016 Super Pressure Balloon campaign from Wanaka, New Zealand.

6.1 GRB 160530A

On May 30, 2016, 14 days into the flight, COSI detected the bright, long duration gamma-ray burst GRB 160530A (Tomsick & the COSI team, 2016). At the time, COSI’s geographic coordinates were 56.79° S, 82.31° E, which resulted in a high background level due to the proximity of the South Magnetic Pole. Additionally, a relativistic electron precipitation event (REP) (Parks et al., 1979) coincided with the GRB observation, resulting in low-frequency variations in the background count rate.

Two COSI images of GRB 160530A are shown in Figure 6.1, where the top image corresponds to a simple back-projection of the Compton cones, and the bottom image corresponds to the same image after 10 iterations of the LM-MLEM algorithm. The images were made using 834 events with event selections that were optimized for imaging performance. The peak value in the image occurs at $l = 243.4^\circ$, $b = 0.4^\circ$ (1.5° error circle, 90% confidence), which is the best known localization for this GRB². In the local COSI coordinate frame, the

¹This chapter is largely based on a recently published paper titled “Polarimetric Analysis of the Long Duration Gamma-Ray Burst GRB 160530A with the Compton Spectrometer and Imager” (Lowell et al., 2017a). ©Astrophysical Journal 2017, reprinted with permission.

²This position is different than the position originally reported in Tomsick & the COSI team (2016) of $l = 243.9^\circ$, $b = 2.1^\circ$ due to changes and improvements in the event reconstruction algorithm.




Figure 6.1: Two COSI images of GRB 160530A in Galactic coordinates using zero iterations (top) and ten iterations (bottom) of the LM-MLEM image deconvolution algorithm. Zero iterations corresponds to a simple back-projection of the Compton cones. The event selections used to make this image were different than those used for the polarization analysis, as the goal here is imaging performance. The color scale intensity is in arbitrary units.




Figure 6.2: (top) COSI and Konus light curves plotted for photon energies between 50 and 200 keV. The Konus light curve has been shifted by +0.651 seconds to account for the COSI-Konus delay time. In this energy band, only 2% of events are Compton events; the rest are single-site events. (bottom) The COSI Compton event light curve, using the event selections shown in Table 6.3.

Pulse	Start Time	Stop Time
1	27.845	30.286
2	30.650	32.374
3	32.607	34.653
4	34.738	35.825
5	36.518	37.826
6	39.731	44.070

Table 6.1: Optimized time windows in units of seconds. Times are with respect to UTC = 1464591800 seconds (May 30th, 2016 07:03:20 UTC). The total time spanned by the windows is 12.945 seconds.

GRB was 43.5° off-axis with an azimuth of -66.1° , although a gradual azimuthal motion due to an anomaly in the gondola rotator caused the GRB position to shift in the local frame by approximately 1.7° over the duration of the GRB. The average position in the local frame was used when performing simulations of GRB 160530A.

GRB 160530A was detected by Konus-Wind and INTEGRAL/ACS (Svinkin et al., 2016b), as well as Astrosat/CZTI and the CsI shield system for Astrosat/CZTI (private

communication). Konus and INTEGRAL/ACS are constituents of the Inter-Planetary Network (IPN), a network of gamma-ray detectors scattered around the solar system which use time delays for source triangulation. IPN carried out a partial triangulation of GRB 160530A using the time delay between the Konus and INTEGRAL/ACS light curves. With only two IPN instruments detecting this GRB with a large baseline, IPN was only able to localize the source position to an annulus in the sky. The 1.5° radius COSI error circle (90% confidence) was found to overlap a region of the IPN annulus, as can be seen in Svinkin et al. (2016b). This prompted the COSI team to trigger a Swift target of opportunity observation to tile the overlapping region in search of the GRB afterglow between 19 ks and 158 ks post-burst (Evans, 2016; D’Elia et al., 2016b,a). Unfortunately, Swift/XRT did not detect any fading X-ray sources, and therefore no afterglow was detected. This could be due to several reasons, the most likely being that the reported COSI error circle in Tomsick & the COSI team (2016) was not accurate at the time, and that improvements in the data pipeline since then would give a better position. We note that our updated position of $l = 243.4^\circ$, $b = 0.4^\circ$ still overlaps with the IPN annulus, but that the image in Svinkin et al. (2016b) is no longer up to date.

Figure 6.2 shows the COSI and Konus 50 keV to 200 keV light curves (top), as well as the COSI 100 keV to 1 MeV Compton event light curve (bottom). The light curve is characterized by two phases: a decaying pulse train lasting ~ 16 seconds, and an ensuing period of low-level activity lasting ~ 21 seconds. The pulse train is clearly visible in all three light curves, while the low-level activity is only clearly detected in the Konus data. As a result, we excluded the low-level activity phase from our analysis and only used the six pulses that are clearly seen in the COSI Compton event light curve. We performed a time domain optimization of the Compton event light curve in order to extract signal from the pulses and reject background from between the pulses. The procedure for this was to define six time windows bracketing the main pulses, and then use the Differential Evolution algorithm (Storn & Price, 1997) to find the window start and end times that optimize the significance $S/\sqrt{S+B}$, where B is the expected number of background counts and S is the estimated number of source counts. The quantity B is estimated by calculating the average background count rate from two 200 second time periods bracketing the GRB, and multiplying the background rate by the total time spanned by the windows. Subtracting B from the total number of counts within the time windows yields S . The optimized time windows are shown as the grey regions in the bottom panel of Figure 6.2, and printed in Table 6.1. The total time spanned by the optimized windows is 12.945 seconds.

The time delay between the COSI light curve and the Astrosat/CZTI shield light curve was used to constrain COSI’s absolute timing accuracy. COSI events are time tagged using a 10 MHz oscillator and then converted to an absolute time using the pulse-per-second signal from a GPS receiver. In theory, the absolute timing accuracy of the system is $< 1\mu\text{s}$. However, calibrating the timing accuracy to this precision is challenging, and systematic effects will worsen this figure. The best fit time delay between COSI and the Astrosat/CZTI shield was 1 ± 5 ms (K. Hurley, private communication), with COSI detecting the GRB earlier. Using the position of the COSI and Astrosat instruments along with the best known

COSI position for GRB 160530A, the time delay should be 5 ms, with Astrosat detecting the burst first. Given that the absolute timing accuracy of Astrosat is known to be 2 ± 0.3 ms (D. Bhattacharya, private communication), the difference between the observed and predicted time is then 6 ± 5 ms, where the time uncertainties have been added in quadrature. Therefore, We conclude that that the absolute timing accuracy of COSI is good to 11 ms at a confidence level of 1σ .

The Konus team performed a spectral analysis of GRB 160530A on 38.9 seconds of data between 20 keV and 5 MeV (Svinkin et al., 2016a). Fitting with a Band model yielded $\alpha = -0.93 \pm 0.03$, $E_p = 638_{-33}^{+36}$ keV, and an upper limit on β , $\beta < -3.5$ (errors are 90 % confidence). The 20 keV to 10 MeV fluence was reported to be $1.30 \pm 0.04 \times 10^{-4}$ erg cm⁻². For our MC simulations of GRB 160530A, we used the best fit Konus Band model (Band et al., 1993) with $\beta = -3.5$, along with a multiplicative absorption model to account for extinction in the atmosphere. The effective column density was determined by integrating the NRLMSISE00 (Picone et al., 2002) model along the line of sight towards the GRB through a spherical atmosphere out to an altitude of 100 km. The result was an effective column density of 9.2 g cm⁻². A spectral analysis of GRB 160530A using the COSI data set will be presented elsewhere (Sleator et al., in preparation).

6.2 Polarimetric Analysis

Two different methods were used to determine the polarization level and angle of GRB 160530A: a standard method (SM), where the azimuthal scattering angles are histogrammed and then fit with a “modulation” curve, and a maximum likelihood method (MLM) which does not bin the data and uses more information per photon. A detailed description of both the SM and MLM can be found in Chapter 4.

6.2.1 Standard Method

The first step in the SM is to determine the event selections that optimize the statistical minimum detectable polarization given by Equation 4.9. During the optimization, the observation time t was fixed at 12.945 seconds (from the time domain optimization), while r_s , r_b , and μ_{100} were computed for various energy, Compton scattering angle, and ARM selection ranges. For each ensemble of event selections, r_s and r_b were computed using the real data, while μ_{100} was computed using a polarized simulation of GRB 160530A with an arbitrary polarization angle. We employed the Differential Evolution optimization algorithm to find the energy range, Compton scattering angle range, and ARM cut that optimize Equation 4.9. It is reasonable to expect the range on the Compton scattering angle to close in towards $\sim 90^\circ$, as this is where the ideal modulation peaks (Lei et al., 1997), giving a larger value for μ_{100} . However, a trade-off exists between the magnitude of μ_{100} and the source count rate; the optimized event selections should give as large a value for μ_{100} as possible while still accepting as many source counts as possible. The other event selections – distance between

interactions, and number of interaction sites – were fixed during the optimization. The first distance cut was set to 0.5 cm so as to avoid artifacts that appear in the azimuthal scattering angle distributions when the interactions are too close. These artifacts are a result of the finite position resolution of the detectors. The cut on the minimum distance between any two interactions was set at a more relaxed 0.3 cm, which has the effect of rejecting many 3+ site events where interactions occur on neighboring strips (strip pitch = 0.2 cm). This has no impact on the azimuthal scattering angle measurement, since the azimuthal angle is only measured using the first two interactions which are subject to the more restrictive 0.5 cm cut. Finally, events with two or more interaction sites were used in order to get as many counts as possible, although no events were detected with more than five interaction sites. Table 6.2 shows the event selections used for the SM.

After applying these selections, $T = 445$ total counts remained for the analysis, approximately $B = 123$ of which were background. The number of background counts B was estimated by fitting a third order polynomial to the light curve 200 seconds before and after the GRB, but with the time bins corresponding to the GRB excluded from the fit. Integrating the polynomial over the optimized time windows resulted in $B = 123$. Using Equation 4.9 and assuming no systematic error, the MDP was $58 \pm 2\%$. In order to determine the MDP including systematic errors, we generated $N = 10,000$ trial data sets using $S = T - B = 322$ counts from an unpolarized simulation of GRB 160530A, and $B = 123$ counts from real background data taken from two 200 second intervals bracketing the GRB. For each trial data set, the polarization analysis was performed and the resulting polarization level was stored. Finally, the 99th percentile of polarization levels was determined numerically and found to be $\text{MDP} = 72.3 \pm 0.8\%$. We used this value as our 99% confidence (2.6σ) detection threshold.

We studied the dependence of the modulation factor μ_{100} on the polarization angle by performing MC simulations of a 100% polarized GRB 160530A at various polarization angles. The result of each simulation is shown in Figure 6.3, along with the best fit constant. The χ_{red}^2 for the fit is 0.42 (dof = 17), indicating that the value of μ_{100} is consistent with being uniform over the full range of polarization angles. The best fit constant $\mu_{100} = 0.484 \pm 0.002$ was used in the following analysis.

Once the best event selections and μ_{100} have been identified, the SM analysis may proceed. Three separate azimuthal scattering angle distributions (ASADs) are shown in Figure 6.4, corresponding to a background subtracted ASAD for GRB 160530A (top), a “correction” ASAD generated from an unpolarized simulation of GRB 160530A (middle) which has been rescaled by the mean value, and finally the corrected ASAD (bottom). The final corrected ASAD (bottom) was obtained by dividing the background-subtracted GRB 160530A ASAD (top) by the “correction” ASAD (middle). This was done to correct for systematic effects of the detector system such as non-uniformity of efficiency (due to geometry, channel thresholds, etc.) and measurement uncertainty. Finally, Equation 4.6 was fit to the corrected ASAD. The fit results are shown in Table 6.2. We measured a polarization level of $\Pi = 33_{-31}^{+33}\%$ for this GRB using the SM, which was below the detection limit $\text{MDP} = 72.3 \pm 0.8\%$. In order to determine the 90% confidence upper limit on the polarization level, we used the




Figure 6.3: The modulation factor μ_{100} as a function of the simulated polarization angle, along with the best fit constant value of $\mu_{100} = 0.484 \pm 0.002$, $\chi_{red}^2 = 0.42$ (dof = 17). For each simulation, $2.4 - 2.5 \times 10^4$ counts were used.

MINOS (James & Roos, 1975) algorithm to determine the 90% confidence contour in the 2D parameter space of amplitude vs. offset. Then, the maximum value for the modulation along the contour was found and divided by μ_{100} , which yielded a 90% upper limit on the polarization level of 79%.

6.2.2 Maximum Likelihood Method

The MLM can perform better than the SM because additional information – the photon energy and Compton scattering angle – is used to implicitly weight each event’s contribution to the likelihood statistic (Krawczynski, 2011b). Contrast this with the SM, where the azimuthal scattering angle of each event that passes the event selections is added to a histogram, and all other event information is discarded. Additionally, the MLM benefits from

Event Selections	
Total photon energy	111 – 847 keV
First Compton scattering angle	55.4° – 145.2°
Angular resolution measure	ARM < 12.5°
Smallest distance between first two interactions	0.5 cm
Smallest distance between any two interactions	0.3 cm
Number of interaction sites	2 - 7

Statistics	
Total counts T	445
Background counts B	123
MDP	72.3 ± 0.8%
Modulation factor μ_{100}	0.484 ± 0.002

Fit Results	
Amplitude (counts s ⁻¹)	0.12 ± 0.08
Offset (counts s ⁻¹)	0.76 ± 0.05
Fitted polarization angle $\hat{\eta}_0$	117 ^{+19°} _{-20°}
Fitted modulation $\hat{\mu}$	0.16 ^{+0.16} _{-0.15}
Polarization level $\Pi = \hat{\mu}/\mu_{100}$	33 ^{+33%} _{-31%}
Π 90% upper limit	79%
χ^2_{red} (dof = 27)	0.98

Table 6.2: SM event selections, statistics, and fit results. The number of bins in the azimuthal scattering angle distribution was 30 (12° per bin).

being able to use events with any energy or Compton scattering angle, resulting in more usable counts for the analysis. Consequently, for the MLM event selections, we accepted all photon energies between 100 keV and 1000 keV, as well as all Compton scattering angles 0° – 180°. With these broader selections, we determined that the FWHM of the ARM distribution was 11.9° using simulated data. For the ARM selection range, we used twice the FWHM of the ARM distribution, centered about zero, i.e. |ARM| < 11.9°. The distance cuts used for the MLM were the same as for the SM and for the same reasons as outlined in Section 6.2.1.

The MLM aims to determine the polarization level Π and angle η_0 that maximize the log likelihood (Equation 4.11). In the presence of background, the probability distribution in Equation 4.12 must be mixed with the probability distribution which describes the background, p_{bkg} . The relative contributions of the source and background probability are mixed using the signal purity $f = (T - B)/T$, where T is the total number of counts, and B is the expected number of background counts. The total probability for event i is then




Figure 6.4: The background subtracted azimuthal scattering angle distribution (ASAD) for GRB160530A (top), an ASAD from an unpolarized simulation of GRB160503A used to correct for systematics (middle), and the corrected ASAD (bottom) showing the best fit modulation curve in blue.

given by Equation 4.13. For the event selections in Table 6.3, we compute a signal purity of $f = 0.72 \pm 0.01$. The error bar on the signal purity is the standard deviation of the simulated distribution of f , obtained by randomly sampling the Poisson distributions underlying T and B . We determined p_{bkg} by accumulating background events from a time period which included GRB 160530A. Due to the large number of bins (46656) in the histogram representing p_{bkg} , this time period spanned 48 hours (7 hours pre-GRB, 41 hours post-GRB). Although it would have been preferable to only include background data from a time period closer in time to the GRB itself, this was not possible due to the relatively low absolute count rate. The background environment at balloon float altitudes, which is dominated by cosmic ray secondaries (photons in particular), is known to vary with altitude, or equivalently, atmospheric depth (Ling, 1975; Bowen et al., 2007). Throughout the 48 hour background




Figure 6.5: The MLM correction factor Π_{100} as a function of the simulated polarization angle, along with the best fit constant value $\Pi_{100} = 0.799 \pm 0.003$, $\chi_{red}^2 = 0.68$ (dof = 17). For each simulation, $3.4 - 3.6 \times 10^4$ counts were used.

integration interval, the COSI altitude was quite stable at ~ 33.1 km, with small excursions at a level of $< 2\%$. Thus, we conclude that using background events from this time interval is a reasonable choice.

The simulated data sets used to determine μ_{100} in the SM were also used to determine Π_{100} in the MLM. A plot of Π_{100} vs. polarization angle can be found in Figure 6.5, along with the best fit constant. As in the SM case, the distribution is rather uniform; the best fit constant has a value of $\Pi_{100} = 0.799 \pm 0.003$ with a χ_{red}^2 of 0.65 (dof = 17). Once again, we used the best fit constant for the MLM analysis.

After applying the MLM event selections, $T = 542$ counts remained for the analysis, approximately $B = 152$ of which were background. Once again, B was determined by fitting a third order polynomial to the light curve, and integrating the polynomial over the time windows (Table 6.1). The MDP was computed in a similar fashion to the SM, where

Event Selections	
Total photon energy	100 - 1000 keV
First Compton scattering angle	$0^\circ - 180^\circ$
Angular resolution measure	$ \text{ARM} < 11.9^\circ$
Smallest distance between first two interactions	0.5 cm
Smallest distance between any two interactions	0.3 cm
Number of interaction sites	2 - 7

Statistics	
Total counts T	542
Background counts B	152
Signal purity f	0.72 ± 0.01
MDP	$57.5 \pm 0.8\%$
MLM correction factor Π_{100}	0.799 ± 0.003

Fit Results	
Fitted polarization angle $\hat{\eta}_0$	$141^\circ \pm 47^\circ$
Fitted, uncorrected polarization level $\hat{\Pi}$	$0.13^{+0.22}_{-0.13}$
Corrected polarization level $\Pi = \frac{\hat{\Pi}}{\Pi_{100}}$	$16^{+27}_{-16}\%$
Π 90% upper limit	46%

Table 6.3: MLM event selections, statistics, and fit results.

$N = 10,000$ unpolarized trial observations were generated and analyzed using the MLM. This process yielded $\text{MDP} = 57.5 \pm 0.8\%$.

Table 6.3 summarizes the results of the MLM analysis, while Figure 6.7 shows the confidence contours for GRB 160530A along with the best fit values, detection limit, and 90% confidence upper limit. The corrected polarization level was $\Pi = 16^{+27}_{-16}\%$, which was below the detection limit $\text{MDP} = 57.5 \pm 0.8\%$. We obtained a 90% confidence upper limit on the corrected polarization level Π using a Monte Carlo/bootstrap approach. The procedure was as follows. First we generated $N = 10,000$ bootstrap resamples of the measured data. Second, for each resample, we drew a value of $f = (T - B)/T$ by assuming a Poisson distribution for the total number of counts ($T = 542$) and the expected number of background counts ($B = 152$). Third, we re-ran the minimizer on each resample using the value of f generated in step 2 in order to yield a value of $\hat{\Pi}$. Finally, the $\hat{\Pi}$ values were re-scaled by values of Π_{100} drawn from its associated probability distribution, which was assumed to be Gaussian. The 90th percentile of the resulting distribution on Π was then used as the 90% confidence upper limit. Using this procedure, we arrived at a 90% confidence upper limit of 46% on the polarization level. The benefit of the Monte Carlo/bootstrap approach is that the errors in the signal purity and Π_{100} are correctly propagated into the final result. A plot




Figure 6.6: Bootstrap distribution of the polarization level using the MLM. Each entry in the histogram was computed by bootstrap resampling the 542 qualifying events, drawing a sample of the signal purity f from its underlying distribution, running the minimizer, and dividing the result by a value of Π_{100} drawn from its underlying distribution. The dashed line corresponds to the measured value, and the dot-dashed line is the 90th percentile of the distribution. We interpret this as our 90% confidence upper limit on the polarization level.

of the bootstrap distribution on Π is shown in Figure 6.6, along with vertical lines denoting the measured value as well as the 90% confidence upper limit.

6.3 MDP Improvement

Figure 6.8 shows the distribution of Π using the MLM (top) and SM (bottom) for a simulated polarization level of $\Pi = 0$. The dark green line denotes the 99th percentile of the distribution, and the lighter green lines represent the standard deviation of the 99th percentile, obtained by bootstrap resampling of the underlying data vector. Each trial observation




Figure 6.7: Confidence contours for GRB 160530A, derived from a map of the log likelihood. The optimal polarization level and angle are denoted by the cross. Confidence contours from MINOS are shown for 1σ to 5σ as well as 90% confidence. The dashed line denotes the 90% confidence upper limit on the polarization level as determined using a Monte Carlo/bootstrap technique. The dotted line represents the MDP, or equivalently, the 99% confidence detection threshold.




Figure 6.8: Distribution of polarization levels from simulated trial observations of GRB 160530A using the MLM (top) and SM (bottom). The true polarization level was set to zero for these simulations. In the MLM case, each trial observation was analyzed in order to determine $\hat{\Pi}$, which was then corrected using the MLM correction factor Π_{100} (Equation 4.14). The resulting polarization level was added to the histogram. For the SM, the fitted modulation $\hat{\mu}$ was divided by μ_{100} and then added to the histogram (Equation 4.8). The MDP, which is the 99th percentile of this distribution, is represented by the dark green vertical lines. The lighter green lines represent the standard deviation of the MDP, which was obtained by bootstrap resampling the underlying data vector. The MLM MDP is 20% lower than the SM MDP.

that contributed to the distributions in Figure 6.8 came from realistic simulations of our observation of GRB 160530A. The number of source counts, number of background counts, and event selections used for the trial observations can be found in Table 6.2 (SM) and Table 6.3 (MLM). The MLM MDP is 20% lower than the SM MDP for GRB 160530A. This improvement is consistent with the findings of Krawczynski (2011b), where it was shown that application of the MLM can improve the MDP of an idealized Compton polarimeter by $\sim 21\%$.

6.4 Discussion

Both the SM and MLM give best fit polarization levels below their respective MDPs, so we conclude that COSI did not detect polarized emission from GRB 160530A at a confidence level of 99%. Moreover, from the error bars on the polarization levels, as well as the log likelihood contours, the measured polarization level is consistent with $\Pi = 0\%$ at a confidence level of 1σ . Our most constraining 90% confidence upper limit on the polarization level was 46% using the MLM. This upper limit is not strong enough to favor or disfavor any particular polarization scenario on its own. However, several GRBs have been found to be polarized at a level above our upper limit (McConnell, 2017; Chattopadhyay et al., 2017), so our result is indeed useful in the context of the overall distribution of polarization levels within the GRB population. GRBs with polarization levels exceeding $\Pi = 46\%$ can be generated in the Compton drag model (Lazzati et al., 2004), but are more frequently explained as arising from a scenario where synchrotron emission in an ordered magnetic field dominates. However, the polarization level in such scenarios can be reduced to lower levels due to certain factors, such as an evolution in the polarization angle with time. Such changes in the polarization angle with time were seen in GRB 100826A (Yonetoku et al., 2011a), and it has been theorized that each pulse in a GRB light curve may correspond to a distinct ordering – and thus polarization angle – of a large scale magnetic field within the GRB jet. Moreover, the Internal-Collision-induced MAGnetic Reconnection and Turbulence (ICMART) model for GRBs (Zhang & Yan, 2010) predicts that the polarization angle and level change from pulse to pulse, and also that the polarization level continually decreases throughout the duration of each pulse. These changes in the polarization properties with time are theorized to result from collisions between highly magnetized shells, which are launched intermittently from the GRB central engine. As the shells collide, the magnetic field lines distort and lose uniformity, thus reducing the polarization level of synchrotron emission. GRB 160530A, whose early phase light curve is composed of several clear pulses, would be a good candidate to perform time-resolved polarimetry. Unfortunately the number of counts detected with COSI from this GRB was not large enough to enable such studies.

Several factors reduced COSI’s sensitivity to this GRB. First, attenuation by the atmosphere significantly reduced the source flux; at the average measured photon energy of 360 keV, and with an effective column density of 9.2 g cm^{-2} , approximately 60% of photons were absorbed or scattered by the atmosphere. Moreover, the 100 keV to 1 MeV flux of the Band model with the atmospheric absorption component was reduced by 64% compared to the unabsorbed case. Second, the GRB occurred 43.5° off axis in the COSI local frame, so the effective area was reduced by about 22% due to occultation from the CsI shields. Third, by the time COSI detected this GRB, two germanium detectors were no longer operational due to anomalies in the corresponding high voltage supplies. Accordingly, all MC simulations discussed in this paper treated the two non-functioning detectors as passive mass. A simulation of the fully functioning detector system revealed that about 16% of events were lost due to the malfunctioning detectors. Lastly, as mentioned in Section 6.1, the GRB coincided with a REP event as well as a relatively close approach to the South Magnetic

Pole, both of which had the effect of elevating the background count rate.

Despite these losses, we were able to regain some sensitivity by using the maximum likelihood based approach. Notably, the MDP in the MLM is $\sim 20\%$ lower than the SM, while the 90% confidence upper limit on the polarization level is $\sim 39\%$ lower. This improvement in MDP is consistent with what was reported by Krawczynski (2011b), where the MDP of an idealized Compton polarimeter was shown to improve by $\sim 21\%$ when using the MLM.

Chapter 7

Towards a Better Polarization Sensitivity

Polarization measurements of GRB prompt emission remains a promising avenue for discerning between various models of the GRB phenomenon. With some measurements nearing the 4σ level, we are indeed on the brink of a new era of GRB polarization measurements. However, pushing the sensitivity will require advances in technology and analysis methods. Certain models that predict polarization in the prompt emission also predict that the polarization level may depend on time, energy, or both. To test these models, the improvements in sensitivity must be significant enough such that the MDP in each time or energy bin is sufficiently low.

In this final chapter, we will outline two avenues that we have identified for performing more sensitive GRB polarization measurements. The first is an “N-site” maximum likelihood method, where the azimuthal scattering angles of all Compton interactions — not just the first — of each event are used to infer the polarization properties of the beam. The second is the “Atmospheric Modulation Effect”, where photons that Compton scatter in the atmosphere will display a modulation in their arrival directions when they reach a detector at balloon float altitudes. Both of these techniques are the subjects of upcoming publications (Lowell 2017, in preparation).

7.1 N-site Maximum Likelihood Method

Compton polarimeters are designed to measure the azimuthal scattering angle of the first Compton scatter. Some Compton polarimeters, such as COSI, are also capable of identifying subsequent Compton scatters in the detectors, and measuring their azimuthal scattering angles. In fact, these subsequent Compton scatters contain information regarding the polarization of the incoming beam. The challenge is in determining how the polarization state is transformed after Compton scattering. In this section, we introduce the N-site Maximum Likelihood Method (NMLM), which uses all the Compton scatterings of all detected events

to further constrain the polarization parameters.

7.1.1 Stokes Vector Framework

As we saw in Chapter 5, Compton scattering of an unpolarized beam results in a partially polarized outgoing beam, where the level of polarization is given by Equation 5.1, and the polarization vector is perpendicular to the plane of scattering. This was the basis for forming partially polarized beams for our polarimetric validations. When a polarized beam Compton scatters, the outgoing polarization level and orientation of the polarization vector are also modified. The expression for the polarization level of the outgoing beam is given by (Matt et al., 1996)

$$\Pi_P = 2 \frac{1 - \sin^2 \theta \cos^2 \eta}{\epsilon + \epsilon^{-1} - 2 \sin^2 \theta \cos^2 \eta} \quad (7.1)$$

Unlike Equation 5.1, the outgoing polarization level in this case depends on the azimuthal scattering angle. The outgoing polarization vector is given by (Matt et al., 1996)

$$\vec{P} = \frac{1}{|\vec{P}|} (\vec{P}_0 \times \vec{D}) \times \vec{D} \quad (7.2)$$

where \vec{P}_0 is the incoming polarization vector, and \vec{D} is the direction of the scattered photon.

A more general framework exists to describe the effects of Compton scattering on the polarization state of a beam of photons. The framework makes use of Stokes vectors, defined as

$$\begin{bmatrix} I \\ Q \\ U \\ V \end{bmatrix} \quad (7.3)$$

where I is the beam intensity, Q is the plane polarization component, U is the diagonal plane polarization component, and V is the circular polarization component. Here we consider only linear polarization, so we set $V = 0$. The polarization level of the beam is given by

$$\Pi = \frac{\sqrt{Q^2 + U^2}}{I}, \quad (7.4)$$

and the polarization angle is given by

$$\eta_0 = \frac{1}{2} \arctan \frac{U}{Q}, \quad (7.5)$$

Fano (1949) determined the transformation matrix for Compton scattering, which can be used to transform the Stokes vector of a Compton scattered beam. Here, we follow the

conventions used in McMaster (1961) and Depaola (2003). The matrix is given by:

$$\begin{aligned}
 T(\theta, k, k_0) &= \frac{1}{2} r_0^2 \left(\frac{k}{k_0} \right) \begin{bmatrix} a & b & c & d \\ e & f & g & h \\ i & j & k & l \\ m & n & o & p \end{bmatrix} \\
 a &= 1 + \cos^2 \theta + (k_0 - k)(1 - \cos \theta) \\
 b &= \sin^2 \theta \\
 c &= 0 \\
 d &= -(1 - \cos \theta)(\vec{k}_0 \cos \theta + \vec{k}) \cdot \vec{S} \\
 e &= \sin^2 \theta \\
 f &= 1 + \cos^2 \theta \\
 g &= 0 \\
 h &= (1 - \cos \theta) \vec{n} \times \vec{n}_0 \cdot \vec{k}_0 \times \vec{S} \\
 i &= 0 \\
 j &= 0 \\
 k &= 2 \cos \theta \\
 l &= (1 - \cos \theta) \vec{k} \times \vec{n}_0 \cdot \vec{S} \\
 m &= -(1 - \cos \theta)(\vec{k} \cos \theta + \vec{k}_0) \cdot \vec{S} \\
 n &= (1 - \cos \theta) \vec{n}_0 \times \vec{n} \cdot \vec{k} \times \vec{S} \\
 o &= (1 - \cos \theta) \vec{k}_0 \times \vec{n} \cdot \vec{S} \\
 p &= 2 \cos \theta + (k_0 - k)(1 - \cos \theta) \cos \theta
 \end{aligned} \tag{7.6}$$

where k_0 and k represent the incoming and outgoing photon energies in units of mc^2 , r_0 is the classical radius of the electron, θ is the Compton scattering angle, and \vec{S} is the spin of the electron before scattering. Terms involving \vec{S} only occur in the fourth row and column of T . If we assume an isotropic distribution for \vec{S} , then the fourth row and column can be set to zero (Depaola, 2003). Moreover, if we assume that the initial Stokes vector has no circular polarization ($V = 0$), then the matrix can be reduced to a 3×3 matrix:

$$T(\theta, k, k_0) = \frac{1}{2} r_0^2 \left(\frac{k}{k_0} \right) \begin{bmatrix} 1 + \cos^2 \theta + (k_0 - k)(1 - \cos \theta) & \sin^2 \theta & 0 \\ \sin^2 \theta & 1 + \cos^2 \theta & 0 \\ 0 & 0 & 2 \cos \theta \end{bmatrix} \tag{7.7}$$

An important point about the matrix T is that it relates the incoming and outgoing Stokes vector with respect to the plane of scattering. Therefore, if multiple scatterings are involved, the incoming Stokes vector must first be rotated into the new scattering plane

before T can be applied. The rotation is given by

$$M(\omega) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\omega & \sin 2\omega & 0 \\ 0 & -\sin 2\omega & \cos 2\omega & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (7.8)$$

where ω is the angle between the previous and current scattering plane. Once again, circular polarization is not considered, so the fourth row and column of $M(\omega)$ are ignored.

7.1.2 NMLM Algorithm

The NMLM algorithm is an extension of the MLM algorithm described in Section 4.1.3. One difference is that rather than finding the optimal values of Π and η_0 , we instead search for the optimal values of Q and U . For each Compton interaction, the following parameters are stored: the incoming direction, incoming energy, outgoing direction, and outgoing energy. The analysis for a single event starts with an initial Stokes vector $\vec{S} = (1, Q, U, 0)$ which is with respect to a reference plane¹. For each Compton interaction in the event...

1. Determine the angle ω between the previous scattering plane and the current scattering plane. If this is the first Compton interaction of the event, ω is the angle between the reference plane and the current scattering plane.
2. Rotate the Stokes vector using $M(\omega)$ from Equation 7.8 so that it now represents the polarization state of the incoming beam with respect to the scattering plane of the current Compton interaction.
3. Determine the polarization level and angle for the current Compton scatter using the rotated Stokes vector (Equations 7.4 and 7.5).
4. Compute the azimuthal scattering angle probability (Equation 4.5 for the ideal case) using the calculated polarization level and angle from step 3, as well as the measured Compton scattering angle, azimuthal scattering angle, and initial photon energy for the current Compton interaction. Increment the log likelihood by the log of the computed probability.
5. Transform the rotated Stokes vector using $T(\theta, k, k_0)$ from Equation 7.6 along with the corresponding values of θ , k , and k_0 for the current Compton interaction.
6. Repeat steps 2 through 7 for each Compton interaction.

Performing the above steps for all events, we can compute the total log likelihood for trial values of Q and U . An optimizer such as MINUIT (James & Roos, 1975) can then be used to determine the optimal values of Q and U .

¹For the starting reference plane, we used the plane normal to the \hat{y} direction.





Figure 7.1: A schematic representation of the NMLM simulations. A cube of germanium with side length 1 m was used as an ideal detector. Monoenergetic photons were started from the center of the cube with initial directions parallel to the $-\hat{z}$ direction.

7.1.3 Simulation Setup


To study the performance of the NMLM, we carried out simulations using an ideal germanium detector cube. A diagram illustrating the simulation setup is shown in Figure 7.1. The photons were started from within the cube and along the $-\hat{z}$ direction. The dimensions of the cube were chosen to be very large so that the probability of escape of the photon was negligible. The parameters for each Compton interaction (incoming direction, incoming energy, outgoing direction, outgoing energy) were stored with ideal precision. The total number of simulated photons N was subdivided into groups of size 500, where each of the 500 photons Compton scattered in the germanium at least once. The polarization level and angle were determined for each group of 500 photons using the SM, MLM, and NMLM.

7.1.4 Results


We first set out to determine if the NMLM improves the MDP over the SM and MLM. To do this, we ran simulations of unpolarized monoenergetic sources at 200, 400, and 800 keV. Each sample of 500 photons was analyzed with the three methods, and the polarization level inferred from each method was added to a histogram. Figures 7.2a, 7.2b, and 7.2c show the resulting histograms, where the dashed lines correspond to the MDP — the 99th percentile of the distribution. The MDPs are also printed in Table 7.1. The NMLM outperforms the SM



(a) 200 keV




(b) 400 keV




(c) 800 keV

Figure 7.2: Histograms of the inferred polarization level using the SM, MLM, and NMLM, for a true polarization level of $\Pi = 0$. The dashed lines denote the 99th percentile of the distribution i.e. the MDP. The photon energies used were 200 keV (top) 400 keV (middle) and 800 keV (bottom). Each sample of events contained 500 photons.



(a) Distribution of polarization levels as inferred by the SM, MLM, and NMLM.



(b) Distribution of polarization angles as inferred by the SM, MLM, and NMLM.

Figure 7.3: Histograms of the inferred polarization level (top) and angle (bottom), where the true values are $\Pi = 0.65$ and $\eta_0 = 0^\circ$. The photon energy used here was 400 keV. Each analyzed sample contained 500 photons. The NMLM gives a more sharply peaked distribution about the true values (denoted by the vertical black lines).

Energy (keV)	SM MDP	MLM MDP	NMLM MDP	Compton interactions per sample (avg.)
200	44%	35%	30%	902
400	54%	43%	31%	1402
800	72%	57%	38%	1827

Table 7.1: MDPs and average number of Compton interactions per sample of 500 photons for the SM, MLM, and NMLM at 200 keV, 400 keV, and 800 keV.

and MLM at all three energies. This is because the NMLM uses more information (Compton scatters) per photon. In Table 7.1, the average number of Compton scatters per sample of 500 photons is listed. While the SM and MLM only use 500 Compton scatters per sample of 500 photons, the NMLM uses 902 Compton scatters at 200 keV, 1402 Compton scatters at 400 keV, and 1827 Compton scatters at 800 keV. The relative improvement gained from using the NMLM increases at higher energies, as more Compton scatters per photon are expected for higher energy photons. In the low energy limit, where the average number of Compton scatters per photon approaches one, the NMLM MDP approaches the MLM MDP.

We also sought to determine if the NMLM gives more accurate estimates for the true polarization level and angle. We tested this by running a polarized simulation of 400 keV photons ($\Pi = 0.65, \eta_0 = 0^\circ$), and histogramming the polarization level and angle given by each of the three methods. The polarization level histograms are shown in Figure 7.3a, and the polarization angle histograms are shown in 7.3b. The vertical black lines in these two figures show the true values, and it is clear that the NMLM distributions are more sharply peaked around the true values. Thus, the NMLM also outperforms the SM and MLM in its accuracy of the estimates for the polarization level and angle.

7.1.5 Discussion

It is not surprising that the NMLM outperforms the SM and MLM for a given number of photons, since the NMLM uses significantly more information per photon. Our results indicate that the NMLM improves the SM MDP by 32% at 200 keV, 43% at 400 keV, and 47% at 800 keV. Another interesting detail is that while the MDPs for the SM and MLM get moderately worse with energy, the NMLM only gets slightly worse with increasing energy. The SM and MLM MDPs worsen with energy because the modulation of the azimuthal scattering angle decreases with increasing energy. This is also true in the NMLM, but the increased number of Compton interactions per photon appears to counterbalance the reduction in modulation. The result is that the polarization sensitivity degrades less with energy in the NMLM.

These simulations used an ideal germanium detector with perfect event resolution², per-

²We use the term event resolution to describe a detectors ability to resolve individual interactions. Detectors with coarse segmentation will cluster closely spaced interactions, resulting in a poorer event resolution.

fect measurement accuracy, and perfect efficiency. Such detectors do not exist in the real world, and a considerable amount of work lies ahead to implement the NMLM for a real detector such as COSI. In particular, the application of Equation 4.5 to evaluate the azimuthal scattering angle probability will no longer be valid due to geometric and other systematic effects of the detector system. Unfortunately, the approach outlined in Section 4.1.3 for the evaluation of the probability in the MLM cannot be used for the NMLM. The reason is that the polarization response used in the MLM analysis is generated for a specific incidence angle, where the photons originate from a far field source. In the NMLM, after the photon has Compton scattered once, its direction changes, and so a response with a “burned in” incidence angle is no longer applicable to the subsequent Compton scatters. Thus, future work will focus on methods for evaluating the azimuthal scattering angle probability for each Compton scatter of each event for non-ideal polarimeters.

Finally, we would like to point out one further extension to this technique that may be applicable to low-Z, finely segmented detectors. In photoelectric polarimetry, the azimuthal angle of the ejected photoelectron is related to the linear polarization of the photoabsorbed photon. Thus, if it is possible to track the photoelectron, then the azimuthal angle may be used in conjunction with the azimuthal scattering angles from the previous Compton scatters to further increase the number of measurements used per photon.

7.2 Atmospheric Modulation Effect

Residual atmosphere above balloon borne gamma-ray instruments acts to scatter and absorb photons that may have reached the detector otherwise. Usually, this is considered to be a detrimental effect; we saw in Section 6.4 that 64% of the flux from GRB 160530A was lost to atmospheric attenuation. However, for certain values of atmospheric depth, the atmosphere can be used as the scattering element of a Compton polarimeter. The basic principle is illustrated in Figure 7.4, where the viewer is looking down into the atmosphere ($-\hat{z}$ direction, into the page). Photons from a polarized GRB are incident atop the atmosphere along the $-\hat{z}$ direction, where the polarization vector — or electric field vector — is denoted by the blue arrows, and parallel to the \hat{y} direction. The purple arrows represent the outgoing photon direction in the xy plane after Compton scattering in the atmosphere. The density of the purple arrows, which is higher along the $\pm\hat{x}$ directions, is representative of the statistical distribution of azimuthal scattering angles given by Equation 4.5. It is clear from Figure 7.4 that the detector will collect an excess of photons from the $\pm\hat{x}$ directions, and a deficit from the $\pm\hat{y}$ directions. The distribution of azimuthal arrival angles will have the same offset cosine form familiar to the field of polarimetry. We refer to this as the Atmospheric Modulation Effect (AME).

For very low atmospheric depths, the AME results in a modulation of the backscattered photons from below. In fact, this was the impetus for the BATSE Albedo Polarimetry System which was able to place lower limits on the polarization of two GRBs using this technique (McConnell et al., 1996; Willis et al., 2005). The limiting factor in the BAPS measurements

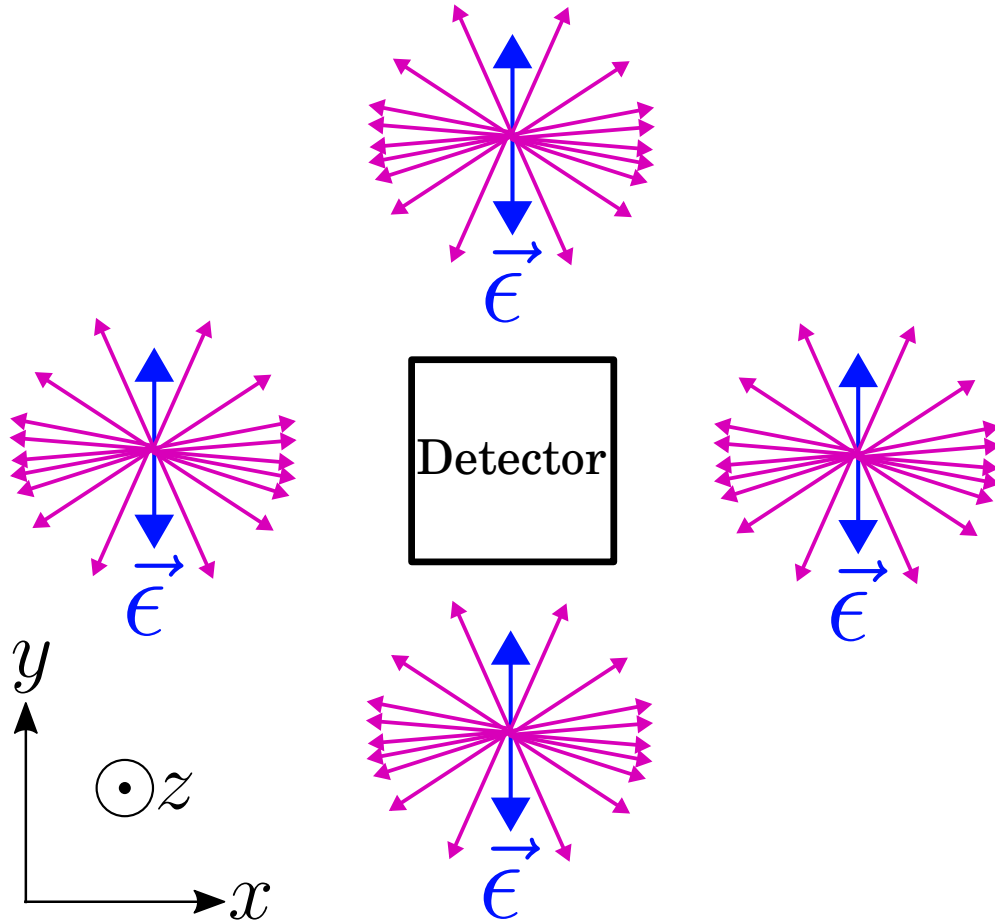


Figure 7.4: A cartoon illustrating the AME principle. The view into the page is looking down from the top of the atmosphere along the $-\hat{z}$ direction. Photons from a polarized GRB move into the page, with their electric field vectors aligned along the \hat{y} direction. These photons Compton scatter, and the scattered photons are denoted by the purple arrows. The density of the purple arrows is drawn to represent the azimuthal scattering angle probability, which is higher in the directions perpendicular to the electric field vector. The detector thus detects an excess of photons from the $\pm\hat{x}$ directions and a deficit from the $\pm\hat{y}$ directions, resulting in a modulation of the azimuthal arrival angles.

was that the BATSE detector modules were not collimated, and so isolating the flux as a function of azimuth was not possible.

At moderate atmospheric depths, a fraction of the photons will reach the detector without interacting in the atmosphere; we will refer to this as the direct component (DC). The scattered component (SC) is the fraction of photons that Compton scatter at least once in the atmosphere before reaching the detector, and is equal to $SC = 1 - DC$, where SC and DC are roughly of the same order. At very high atmospheric depths, $SC \gg DC$, and photons from the SC have Compton scattered more than once. Thus, the modulation due to the




Figure 7.5: A schematic representation of the AME simulations. The atmosphere was modeled as a cylinder of air 100 km tall, and 100 km in diameter. The atmosphere was divided vertically into 100 layers with densities given by the NRLMSISE00 model. Photons from a standard GRB spectrum were started from the top of the atmosphere with initial directions along the $-\hat{z}$ direction (purple arrows). At 33 km, we placed a “detector layer” (blue disk) made of air at the corresponding density. Photons passing through the detector layer had their directions, energies, and positions recorded.

AME becomes vanishingly small. Moreover, much of the photons are severely attenuated at large depths.

7.2.1 Simulation Setup

We investigated the possibility of exploiting the AME as a tool for GRB polarization measurements at a float altitude of 33 km, where the atmospheric depth is 6.3 g cm^{-2} . Specifically, we set out to determine if any modulation was present, and if so, which energies and polar arrival angles gave the best MDP. The mass model for these simulations was simple: 100 cylindrical layers of air representing the atmosphere from 0 to 100 km, where the density of each layer was determined as a function of its altitude using the NRLMSISE00 model. The diameter of the cylindrical atmosphere was also set to 100 km. At 33 km, we placed an infinitesimally thin “detector layer.” The simulation setup is illustrated in Figure 7.5. Note that this not a real detector, it is simply a slab of air with the corresponding density at 33 km. Each photon crossing the detector layer had its position, arrival direction, and energy stored for later analysis. The photon source for the simulation used a typical GRB spectrum from 50 keV to 10 MeV with the following Band parameters: $\alpha = -1, \beta = -2.2, E_p = 250$

keV³. Photons drawn from the Band spectrum were then started at random locations at the top of the atmosphere with initial directions along the $-\hat{z}$ direction.

7.2.2 Results

Our first goal was to determine DC and SC at 33 km altitude. Photons arriving at the detector plane with a polar arrival angle of 180° (corresponding to the $-\hat{z}$ direction) were considered to be constituents of the DC. All other polar arrival angles were considered part of the SC. We find that $SC = 0.57$, and $DC = 1 - SC = 0.43$. Therefore, the SC represents a significant fraction of the total photons that arrive at the detector layer.

Figures 7.6 and 7.7 show two azimuthal arrival angle distributions (AAADs) corresponding to 100% polarization (top), and 50% polarization with a 90° shift in the polarization angle (bottom). We fit Equation 4.6 to the distributions, which resulted in modulations of $\mu_{100} = 0.173 \pm 0.003$, and $\mu_{50} = 0.087 \pm 0.003$ respectively. We also performed a simulation with an unpolarized GRB, which yielded a modulation of $\mu_0 = 0.006 \pm 0.003$. In the context of the AME, we will refer to μ_{100} as the AME modulation factor. Clearly, the modulation tracks the polarization level, and the phase of the AAAD tracks the polarization angle.

The AAADs in Figures 7.6 and 7.7 represent all events crossing the detector layer. We sought to find selection ranges on the energy and polar arrival angle that give the lowest MDP. In the absence of background, the MDP becomes $MDP = 4.29/(\mu_{100}\sqrt{N})$, where N is the number of counts. Here, we are not concerned with the absolute value of the MDP, since we are using an arbitrary flux for the GRB. Rather, the background-less MDP provides a useful optimization target which results in selections that give both a high AME modulation factor and a large number of counts. We used the Differential Evolution (Storn & Price, 1997) algorithm to search the space of energy and polar arrival angle selections. The result is that polar arrival angles between 40° and 149° are preferred, and an energy range of 42 to 2066 keV is preferred. We ran the optimizer several times with different random number seeds and found that the polar arrival angle selection is robust, but that the upper energy changes somewhat between runs. This is due to the fact that there are not many counts above 1 MeV, so the optimizer is not sensitive to the exact value of the upper energy as long as it is high enough. With these selections, we compute an AME modulation factor of $\mu_{100} = 0.260 \pm 0.003$.

7.2.3 Discussion

We have demonstrated through atmospheric simulations that at 33 km altitude, the magnitude of the scattered component is of the same order as the direct component, and that the azimuthal arrival angles are indeed modulated. After optimizing the MDP, the AME modulation factor was $\mu_{100} = 0.260 \pm 0.003$. The optimized polar arrival angle range was

³We used an arbitrary flux as the goal here was to determine if a modulation was present at this altitude, and how strong the modulation would be. For extensive studies including atmospheric background, the flux would be relevant




Figure 7.6: The azimuthal arrival angle distribution at 33 km along with the best fit modulation curve (red) for a fully polarized GRB incident at the top of the atmosphere. The modulation is $\mu_{100} = 0.173 \pm 0.003$. A higher modulation can be achieved by optimizing the polar arrival angle range. Here, no optimizations were performed.




Figure 7.7: The same as Figure 7.6, but with a polarization level of $\Pi = 0.5$ (50%), and a shift of 90° in the polarization angle. The modulation in this case is $\mu_{50} = 0.087 \pm 0.003$.

$40^\circ - 149^\circ$, which is consistent with the idea that the photons arriving from the sides should display a stronger modulation, since those photons Compton scattered in the atmosphere with Compton scattering angles clustered around 90° . Of course, these exact values will change as various parameters of the simulation are changed such as the Band model parameters, the altitude of the detector layer, the incidence angle of the GRB, and the addition of atmospheric background components. Nonetheless, it is clear that the AME is strong enough at 33 km to be used as a tool to study GRB polarization.

COSI is a Compton imager, and so it should be possible to image the SC in addition to performing Compton polarimetry on the DC. This raises the interesting possibility of performing “multi-mode” GRB polarization measurements from within the atmosphere. If the analysis of the SC component were performed in the context of maximum likelihood, then the total likelihood $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{DC}} \cdot \mathcal{L}_{\text{SC}}$ could be maximized to find the best fit polarization level and angle, where \mathcal{L}_{DC} is the likelihood from Compton polarimetry of the DC component and \mathcal{L}_{SC} is the likelihood from analysis of the SC. The combination of the two measurements should result in a better MDP, but this has yet to be tested, and an estimate of the MDP improvement remains to be studied. However there are two limitations to using COSI to study the AME. First, while COSI is capable of imaging, it may prove difficult to separate the SC and the DC in an imaging data space. Further study will shed light on how challenging such a task may be. The second issue is that the CsI shield system blocks photons from the sides, and so the more strongly modulated AME flux cannot reach the GeDs.

A more straightforward way of measuring the modulation of the azimuthal arrival angles would be to construct a system of collimated scintillators to fly as a standalone payload, or as a piggy-back with COSI. A system comprised of 8-12⁴ detector modules is envisioned, where each module consists of a high efficiency NaI scintillator enshrouded in a lead or tungsten collimator such that each detector samples a range of azimuthal and polar arrival angles. The polar arrival angle range of the collimator would be determined by optimizing the simulated MDP, but with all relevant background components included. The azimuthal arrival angle range would correspond to $360^\circ/D$, where D is the number of detectors. This way, the counts in each detector could be interpreted as the counts in each bin of the AAAD. As a standalone payload, the polarization level and angle could be determined from fits to the AAAD. As a piggy-back instrument on the COSI balloon gondola, the likelihood \mathcal{L}_{SC} could be determined from fits to the AAAD, which could then be combined with \mathcal{L}_{DC} from COSI’s Compton polarimetry of the same GRB, thus yielding a lower MDP.

⁴A minimum of 8 detectors — or azimuthal bins — is needed to satisfy the Nyquist sampling criterion.

Bibliography

- Abbott, B. P., Abbott, R., Abbott, T., et al. 2017, *Physical Review Letters*, 119, 161101
- Ackermann, M., Ajello, M., Asano, K., et al. 2011, *The Astrophysical Journal*, 729, 114
- . 2013, *The Astrophysical Journal*, 763, 71
- Ackermann, M., Ajello, M., Asano, K., et al. 2014, *Science*, 343, 42
- Agostinelli, S., Allison, J., Amako, K. a., et al. 2003, *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506, 250
- Amman, M., & Luke, P. 2000, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 452, 155
- Amman, M., Luke, P., & Boggs, S. 2007, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 579, 886
- Anderson, T. W., & Darling, D. A. 1952, *The Annals of Mathematical Statistics*, 23, 93
- Andreyev, A., Sitek, A., & Celler, A. 2011, *Medical physics*, 38, 429
- Asano, K., & Fukuyama, T. 2000, *The Astrophysical Journal*, 531, 949
- Band, D., Matteson, J., Ford, L., et al. 1993, *The Astrophysical Journal*, 413, 281
- Bandstra, M., Bellm, E., Boggs, S., et al. 2011, *The Astrophysical Journal*, 738, 8
- Bandstra, M. E., Bowen, J. D., Zoglauer, A., et al. 2006, in *Nuclear Science Symposium Conference Record, 2006. IEEE*, Vol. 2, IEEE, 770–777
- Bandstra, M. S. 2010, PhD thesis, University of California, Berkeley
- Bandstra, M. S., Bellm, E. C., Chiu, J.-L., et al. 2009, in *Nuclear Science Symposium Conference Record (NSS/MIC), 2009 IEEE*, IEEE, 2131–2139
- Bellm, E. C. 2011, PhD thesis, University of California, Berkeley

- Beloborodov, A., & Mészáros, P. 2017, *Space Science Reviews*, 1
- Beloborodov, A. M. 2000, *The Astrophysical Journal Letters*, 539, L25
- Blandford, R. D., & Znajek, R. L. 1977, *Monthly Notices of the Royal Astronomical Society*, 179, 433
- Bloom, J. S., Djorgovski, S., Kulkarni, S., & Frail, D. 1998, *The Astrophysical Journal Letters*, 507, L25
- Bloom, J. S., Prochaska, J., Pooley, D., et al. 2006, *The Astrophysical Journal*, 638, 354
- Boggs, S., & Jean, P. 2000, *Astronomy and Astrophysics Supplement Series*, 145, 311
- Bouchet, L., Roques, J.-P., & Jourdain, E. 2010, *The Astrophysical Journal*, 720, 1772
- Bowen, J. D., Bandstra, S., Zoglauer, A., et al. 2007, in *Nuclear Science Symposium Conference Record, 2007. NSS'07. IEEE*, Vol. 1, IEEE, 436–444
- Brun, R., & Rademakers, F. 1997, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 389, 81
- Burks, M., Jordan, E., Hull, E., Mihailescu, L., & Vetter, K. 2004, in *Nuclear Science Symposium Conference Record, 2004 IEEE*, Vol. 2, IEEE, 1114–1118
- Cano, Z., Wang, S.-Q., Dai, Z.-G., & Wu, X.-F. 2017, *Advances in Astronomy*, 2017
- Chattopadhyay, T., Vadawale, S. V., Aarthi, E., et al. 2017, arXiv preprint arXiv:1707.06595
- Chiu, J.-L., Boggs, S., Chang, H.-K., et al. 2015, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 784, 359
- Coburn, W., & Boggs, S. E. 2003, *nature*, 423, 415
- Conover, W. 1980, in *Practical nonparametric statistics* (John Wiley & Sons)
- Costa, E., Frontera, F., Heise, J., et al. 1997, *Nature*, 387, 783
- D’Elia, V., Izzo, L., Breeveld, A. A., & Markwardt, C. B. 2016a, *GRB Coordinates Network*, 19501
- D’Elia, V., Izzo, L., Evans, P. A., et al. 2016b, *GRB Coordinates Network*, 19481
- Depaola, G. 2003, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 512, 619
- Dingus, B. L., Catelli, J. R., & Schneid, E. J. 1997, *International Cosmic Ray Conference*, 3, 29

- Eichler, D., & Levinson, A. 2000, *The Astrophysical Journal*, 529, 146
- Eichler, D., Livio, M., Piran, T., & Schramm, D. N. 1989, *Nature*, 340, 126
- Evans, P. A. 2016, *GRB Coordinates Network*, 19472
- Fano, U. 1949, *Journal of the Optical Society of America*, 39, 859
- Fenimore, E., Epstein, R., Ho, C., et al. 1993, *Nature*, 366, 40
- Forot, M., Laurent, P., Grenier, I., Gouiffes, C., & Lebrun, F. 2008, *The Astrophysical Journal Letters*, 688, L29
- Frail, D., Kulkarni, S., Sari, R., et al. 2001, *The Astrophysical Journal Letters*, 562, L55
- Galama, T., Groot, P., Van Paradijs, J., et al. 1998a, *The Astrophysical Journal Letters*, 497, L13
- Galama, T., Vreeswijk, P., Van Paradijs, J., et al. 1998b, *Nature*, 395, 670
- Gehrels, N., Ramirez-Ruiz, E., & Fox, D. B. 2009, *Annual Review of Astronomy and Astrophysics*, 47, 567
- Gehrels, N., & Razzaque, S. 2013, *Frontiers of Physics*, 8, 661
- Ghisellini, G., Lazzati, D., Celotti, A., & Rees, M. J. 2000, *Monthly Notices of the Royal Astronomical Society*, 316, L45
- Goldstein, A., Preece, R. D., Mallozzi, R. S., et al. 2013, *The Astrophysical Journal Supplement Series*, 208, 21
- Goldstein, A., Burgess, J. M., Preece, R. D., et al. 2012, *The Astrophysical Journal Supplement Series*, 199, 19
- Goldstein, A., Veres, P., Burns, E., et al. 2017, *The Astrophysical Journal Letters*, 848, L14
- Guilbert, P., Fabian, A., & Rees, M. 1983, *Monthly Notices of the Royal Astronomical Society*, 205, 593
- Gull, S. F., & Skilling, J. 1984, in *IEE Proceedings F (Communications, Radar and Signal Processing)*, Vol. 131, IET, 646–659
- Heise, J., in 't Zand, J., Costa, E., et al. 1997, *IAU Circular*, 6654
- Hjorth, J., Sollerman, J., Gorosabel, J., et al. 2005, *The Astrophysical Journal Letters*, 630, L117
- Hou, A., Parker, L. C., Harris, W. E., & Wilman, D. J. 2009, *The Astrophysical Journal*, 702, 1199

- Hurley, K. 1978, *Astronomy and Astrophysics*, 69, 313
- James, F., & Roos, M. 1975, *Computer Physics Communications*, 10, 343
- Jourdain, E., & Roques, J. 2009, *The Astrophysical Journal*, 704, 17
- Jourdain, E., Roques, J., Chauvin, M., & Clark, D. 2012, *The Astrophysical Journal*, 761, 27
- Kierans, C. A., Boggs, S. E., Chiu, J.-L., et al. 2017, arXiv preprint arXiv:1701.05558
- Klebesadel, R. W., Strong, I. B., & Olson, R. A. 1973, *The Astrophysical Journal*, 182, L85
- Klose, S. 1995, *Astronomy and Astrophysics*, 303, 75
- Knoll, G. F. 2010, *Radiation Detection and Measurement* (John Wiley & Sons)
- Kobayashi, S., & Sari, R. 2001, *The Astrophysical Journal*, 551, 934
- Kole, M., & Marcinkowski, R. 2016, GRB Coordinates Network, 20301
- Kole, M., Prodit, N., Bernasconi, T., et al. 2016, GRB Coordinates Network, 20243
- Kouveliotou, C., Meegan, C. A., Fishman, G. J., et al. 1993, *The Astrophysical Journal*, 413, L101
- Krawczynski, H. 2011a, *The Astrophysical Journal*, 744, 30
- . 2011b, *Astroparticle Physics*, 34, 784
- Kumar, P. 1999, *The Astrophysical Journal Letters*, 523, L113
- Laurent, P., Rodriguez, J., Wilms, J., et al. 2011, *Science*, 332, 438
- Lazzati, D. 2006, *New Journal of Physics*, 8, 131
- Lazzati, D., Ghisellini, G., Celotti, A., & Rees, M. J. 1999a, *The Astrophysical Journal Letters*, 529, L17
- . 1999b, *The Astrophysical Journal Letters*, 529, L17
- Lazzati, D., & Perna, R. 2009, *Physics of Relativistic Objects in Compact Binaries: From Birth to Coalescence*, 245
- Lazzati, D., Rossi, E., Ghisellini, G., & Rees, M. J. 2004, *Monthly Notices of the Royal Astronomical Society*, 347, L1
- Lee, H. K., Wijers, R. A. M. J., & Brown, G. E. 2000, *Physics Reports*, 325, 83

- Lei, F., Dean, A., & Hills, G. 1997, *Space Science Reviews*, 82, 309
- Lien, A., Sakamoto, T., Barthelmy, S. D., et al. 2016, *The Astrophysical Journal*, 829, 7
- Ling, J. C. 1975, *Journal of Geophysical Research*, 80, 3241
- Lloyd-Ronning, N. M., & Petrosian, V. 2002, *The Astrophysical Journal*, 565, 182
- Looker, Q., Amman, M., & Vetter, K. 2015, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 781, 20
- Lowell, A., Boggs, S., Chiu, J., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9915
- Lowell, A. W., Boggs, S. E., Chiu, C., et al. 2017a, *The Astrophysical Journal*, 848, 119
- . 2017b, *The Astrophysical Journal*, 848, 120
- Lundman, C., Pe'er, A., & Ryde, F. 2014, *Monthly Notices of the Royal Astronomical Society*, 440, 3292
- Lundman, C., Vurm, I., & Beloborodov, A. M. 2016, *ArXiv e-prints*, arXiv:1611.01451
- Lyutikov, M., & Blandford, R. 2003, *arXiv preprint astro-ph/0312347*
- MacFadyen, A., & Woosley, S. 1999, *The Astrophysical Journal*, 524, 262
- Masada, Y., Kawanaka, N., Sano, T., & Shibata, K. 2007, *The Astrophysical Journal*, 663, 437
- Matt, G. 1993, *Monthly Notices of the Royal Astronomical Society*, 260, 663
- Matt, G., Feroci, M., Rapisarda, M., & Costa, E. 1996, *Radiation Physics and Chemistry*, 48, 403
- McConnell, M., Forrest, D., Vestrand, W., & Finger, M. 1996, in *AIP Conference Proceedings*, Vol. 384, AIP, 851–855
- McConnell, M. L. 2017, *New Astronomy Reviews*, 76, 1
- McMaster, W. H. 1961, *Reviews of Modern Physics*, 33, 8
- Medvedev, M. V., & Loeb, A. 1999, *The Astrophysical Journal*, 526, 697
- Meegan, C., Fishman, G., Wilson, R., et al. 1991, in *AIP Conference Proceedings*, Vol. 265, AIP, 61–69
- Meszáros, P., & Rees, M. 1993, *The Astrophysical Journal*, 405, 278

- Meszáros, P., & Rees, M. J. 2000, *The Astrophysical Journal*, 530, 292
- Metzger, B., Giannios, D., Thompson, T., Bucciantini, N., & Quataert, E. 2011, *Monthly Notices of the Royal Astronomical Society*, 413, 2031
- Metzger, M., Djorgovski, S., Kulkarni, S., Steidel, C., et al. 1997, *Nature*, 387, 878
- Millan, R. M., & Thorne, R. M. 2007, *Journal of Atmospheric and Solar-Terrestrial Physics*, 69, 362
- Moran, P., Kyne, G., Gouiffes, C., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 456, 2974
- Narayan, R., Paczynski, B., & Piran, T. 1992, *The Astrophysical Journal*, 395, L83
- Narayan, R., Piran, T., & Kumar, P. 2001, *The Astrophysical Journal*, 557, 949
- Orsi, S., Cadoux, F., Leluc, C., et al. 2014, in *Space Telescopes and Instrumentation 2014: Ultraviolet to Gamma Ray*, Vol. 9144
- Paciesas, W. S., Meegan, C. A., Pendleton, G. N., et al. 1999, *The Astrophysical Journal Supplement Series*, 122, 465
- Panaiteanu, A., & Mészáros, P. 2000, *The Astrophysical Journal Letters*, 544, L17
- Parks, G., Gurgiollo, C., & West, R. 1979, *Geophysical Research Letters*, 6, 393
- Perna, R., Sari, R., & Frail, D. 2003, *The Astrophysical Journal*, 594, 379
- Peer, A., & Ryde, F. 2016, *International Journal of Modern Physics D*, 1730018
- Picone, J., Hedin, A., Drob, D. P., & Aikin, A. 2002, *Journal of Geophysical Research: Space Physics*, 107
- Pilla, R. P., & Loeb, A. 1998, *The Astrophysical Journal Letters*, 494, L167
- Piran, T. 2005, *Reviews of Modern Physics*, 76, 1143
- Piran, T., Halpern, J., et al. 1999, *The Astrophysical Journal Letters*, 519, L17
- Prantzos, N., Boehm, C., Bykov, A., et al. 2011, *Reviews of Modern Physics*, 83, 1001
- Preece, R. D., Briggs, M. S., Mallozzi, R. S., et al. 1998a, *The Astrophysical Journal Letters*, 506, L23
- . 1998b, *The Astrophysical Journal Letters*, 506, L23
- . 2000, *The Astrophysical Journal Supplement Series*, 126, 19

- Price, P., Fox, D., Kulkarni, S., et al. 2003, *Nature*, 423, 844
- Produit, N., Barao, F., Deluit, S., et al. 2005, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 550, 616
- Rees, M., & Mészáros, P. 1992, *Monthly Notices of the Royal Astronomical Society*, 258, 41P
- Rees, M., & Meszaros, P. 1994, *The Astrophysical Journal*, 430, L93
- Ruderman, M., Tao, L., & Kluźniak, W. 2000, *The Astrophysical Journal*, 542, 243
- Rutledge, R. E., & Fox, D. B. 2004, *Monthly Notices of the Royal Astronomical Society*, 350, 1288
- Rybicki, G. B., & Lightman, A. P. 2008, *Radiative Processes in Astrophysics* (John Wiley & Sons)
- Ryde, F., & Pe'er, A. 2009, *The Astrophysical Journal*, 702, 1211
- Schady, P. 2017, *Royal Society Open Science*, 4, 170304
- Schneid, E., Bertsch, D., Fichtel, C., et al. 1992, *Astronomy and astrophysics*, 255, L13
- Schoenfelder, V., Aarts, H., Bennett, K., et al. 1993, *The Astrophysical Journal Supplement Series*, 86, 657
- Scholz, F. W., & Stephens, M. A. 1987, *Journal of the American Statistical Association*, 82, 918
- Segreto, A., Labanti, C., Bazzano, A., et al. 2003, *Astronomy and Astrophysics*, 411, L215
- Silva, L. O., Fonseca, R., Tonge, J., et al. 2003, *The Astrophysical Journal Letters*, 596, L121
- Sleator, C., Boggs, S. E., Chiu, J.-L., et al. 2017, in *AAS/High Energy Astrophysics Division*, Vol. 16, *AAS/High Energy Astrophysics Division*, 304.03
- Sleator, C. C., Boggs, S. E., Chiu, J.-L., et al. 2017, arXiv preprint arXiv:1701.05563
- Sommer, M., Bertsch, D., Dingus, B., et al. 1994, *The Astrophysical Journal*, 422, L63
- Spitkovsky, A. 2008, *The Astrophysical Journal Letters*, 682, L5
- Spruit, H., Daigne, F., & Drenkhahn, G. 2001, *Astronomy & Astrophysics*, 369, 694
- Stephens, M. A. 1974, *Journal of the American Statistical Association*, 69, 730

- Storn, R., & Price, K. 1997, *Journal of Global Optimization*, 11, 341
- Strong, I. B., & Klebesadel, R. W. 1974, *Nature*, 251, 396
- Sun, J., Wu, B., Bao, T., et al. 2016, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9905
- Svinkin, D., Golenetskii, S., Aptekar, R., et al. 2016a, *GRB Coordinates Network*, 19477
- . 2016b, *GRB Coordinates Network*, 19476
- Tavani, M. 1996, *The Astrophysical Journal*, 466, 768
- Thompson, T. A., Chang, P., & Quataert, E. 2004, *The Astrophysical Journal*, 611, 380
- Toma, K. 2013, arXiv preprint arXiv:1308.5733
- Toma, K., Sakamoto, T., Zhang, B., et al. 2009, *The Astrophysical Journal*, 698, 1042
- Tomsick, J. A., & the COSI team. 2016, *GRB Coordinates Network*, 19473
- Usov, V. 1992, *Nature*, 357, 472
- Weisskopf, M. C., Elsner, R. F., & O'Dell, S. L. 2010, in *Space Telescopes and Instrumentation 2010: Ultraviolet to Gamma Ray*, Vol. 7732
- Wigger, C., Hajdas, W., Arzner, K., Güdel, M., & Zehnder, A. 2004, *The Astrophysical Journal*, 613, 1088
- Wilderman, S. J., Clinthorne, N. H., Fessler, J. A., & Rogers, W. L. 1998, in *Nuclear Science Symposium, 1998. Conference Record. 1998 IEEE*, Vol. 3, IEEE, 1716–1720
- Wilks, S. S. 1938, *The Annals of Mathematical Statistics*, 9, 60
- Willis, D., Barlow, E., Bird, A., et al. 2005, *Astronomy and Astrophysics*, 439, 245
- Woosley, S. 1993, *The Astrophysical Journal*, 405, 273
- Woosley, S., & Bloom, J. 2006, *Annual Review of Astronomy and Astrophysics*, 44, 507
- Yonetoku, D., Murakami, T., Gunji, S., et al. 2011a, *The Astrophysical Journal Letters*, 743, L30
- . 2011b, *Publications of the Astronomical Society of Japan*, 63, 625
- . 2012, *The Astrophysical Journal Letters*, 758, L1
- Zhang, B., & Yan, H. 2010, *The Astrophysical Journal*, 726, 90
- Zoglauer, A., Andritschke, R., & Schopper, F. 2006, *New Astronomy Reviews*, 50, 629
- Zoglauer, A. C. 2005, PhD thesis, Technische Universität München

Appendix A

Positional Calibrations for the Compton Spectrometer and Imager¹

COSI’s ability to perform Compton imaging and polarimetry is predicated on its three dimensional positioning capabilities. The segmentation of the electrodes on the two faces of the detectors simplifies localization along the x and y directions, as it is simply a matter of identifying the triggered strips. Localization in the z dimension is more involved. As discussed in Section 2.3.4, each strip electrode is instrumented with a fast bipolar shaper, whose zero crossing time is measured with respect to a 200 MHz timer (5 ns resolution). The time at which the shaper output crosses zero is a measure of the collection time of the induced charges on the electrode. The collection time represents the time taken by the charges to drift from the location of the initial charge cloud to the strip electrodes. By computing the collection time difference τ (Figure 2.6) between triggered x and y strips, the depth, or z coordinate, of interaction may be inferred (Amman & Luke, 2000).

Due to channel variations in time delays and drift velocities, each pixel — defined as the intersection of an x strip and a y strip — must be calibrated so that collection time differences of interactions in said pixel can be converted to depths. The calibration coefficients will in general vary from pixel to pixel. Here, we will review our methods for calibrating each of COSI’s $12 \times 37 \times 37 = 16428$ pixels. We will start with a “classic approach”, which has been previously applied to the COSI GeDs with success (Bandstra et al., 2006; Bandstra, 2010). We will then introduce a newly developed approach which is substantially simpler in its implementation, and leads to similar performance. Finally, we investigate the possibility of using “spectator strips” (Knoll, 2010) to perform sub strip positioning along the x and y directions.

¹This appendix is largely based on a previous publication “Positional Calibrations of the Germanium Double Sided Strip Detectors for the Compton Spectrometer and Imager” by Lowell et al. (2016). ©2016, Society of Photo-Optical Instrumentation Engineers (SPIE). Reprinted with permission.

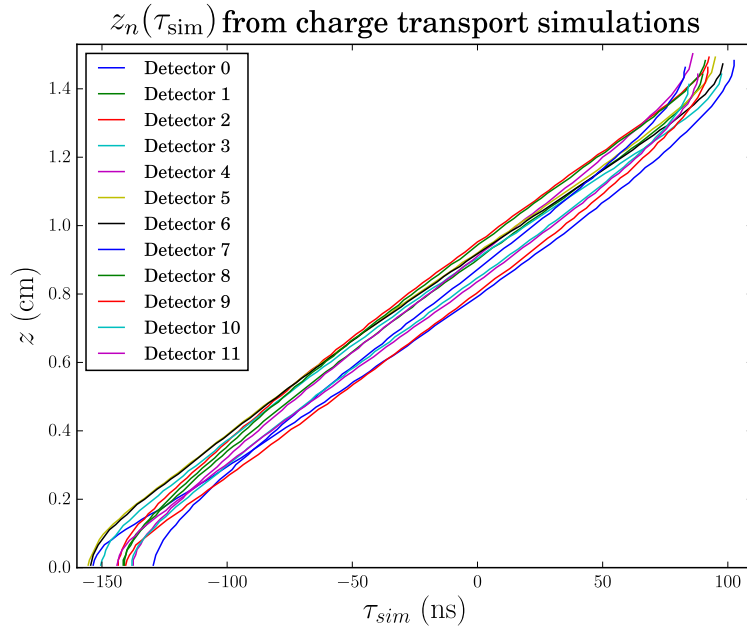


Figure A.1: $z_n(\tau_{\text{sim}})$ for each detector n , as determined from charge transport simulations. For a given detector with specified thickness, bias voltage, and impurity concentration, charges were initiated at various depths z , and the resulting collection time differences τ_{sim} were recorded.

A.1 Depth Calibration: Classic Approach

Previous work on COSI GeD depth calibrations can be found in (Bandstra et al., 2006; Bandstra, 2010), but the procedure will be outlined here for clarity. The steps are as follows:

1. **For each detector, solve Poisson's equation for the electrostatic potential using each detector's impurity concentration, bias voltage, and thickness.** For the COSI GeDs, the ranges for these parameters are impurity concentration $\in [1.1 \times 10^9, 1.0 \times 10^{10}] \text{ cm}^{-3}$, bias voltage $\in [1000, 1500] \text{ V}$, and thickness $\in [1.42, 1.51] \text{ cm}$. A relaxation method is used to solve for the potential in a 3D volume that includes a simplified, 5×5 strip version of the COSI detectors.
2. **Solve for the weighting potential of each strip electrode.** This step is similar to step 1, except different boundary conditions are used: the voltage of a strip electrode is set to 1 V, and all other electrodes are set to 0 V. The weighting potential for each strip allows one to calculate the induced charge on said strip due to drifting charges in the detector via the Shockley-Ramo theorem (Knoll, 2010).
3. **Perform charge transport simulations using the potentials from steps 1 and 2 along with a model of the fast shaper electronics to produce a lookup**

table $z_n(\tau_{sim})$ for each detector n . Charges are deposited at a specific depth in the detector model, and the resulting collection time difference τ_{sim} is recorded. This is done for various depths in the detector, thus yielding a curve $z_n(\tau_{sim})$ for each detector n . Figure A.1 shows $z_n(\tau_{sim})$ for all twelve detectors.

4. **Collect data with a calibration source and mimic the experimental configuration in a Cosima simulation with the COSI mass model.** We used a 100 μCi Cs-137 source (662 keV) positioned on axis at a distance of 64 cm from the cryostat. Only events that triggered one strip on each side and fell in either of the energy ranges 650 keV - 670 keV (photopeak) or 200 keV - 477 keV (continuum) were used for the analysis. Due to the high activity of the Cs-137 source and its close proximity to the cryostat, this calibration run was strongly source dominated. Therefore, the background was ignored in the simulations.
5. **For each detector, extract the simulated interaction depth distribution and transform it into a distribution in τ_{sim} using the inverse of $z_n(\tau_{sim})$ from step 3.** Call these τ_{sim} distributions “CTD templates”.
6. **For each pixel in a detector, fit the detector’s CTD template to the measured τ_{meas} distribution by time-stretching and offsetting the CTD template.** This will yield a “stretching factor” λ and an “offset” Δ which account for variations across the detector in drift velocity and electronic time delays, respectively. Also, a scale factor parameter was left free to vary which simply scaled the CTD template amplitude. The relationship between τ_{sim} and τ_{meas} is thus

$$\tau_{meas} = \lambda\tau_{sim} + \Delta \tag{A.1}$$

Once this depth calibration has been completed, then real, measured interactions in the GeDs can be assigned depths by first transforming τ_{meas} to τ_{sim} via Equation A.1, and then to a depth via $z_n(\tau_{sim})$ from step 3.

Figure A.2 shows examples of CTD template fits to two pixels from detector 6, and two pixels from detector 10. Generally, the fits are acceptable with an average chi-square reduced value of 2.15 for the whole system (16428 pixels). Table A.1 summarizes the fits per detector. One interesting detail that stands out is that the detectors in the top layer (0, 5, 6, and 11) have the highest average χ_{red}^2 and that those in the bottom layer (2, 3, 8, and 9) have the lowest. It is not clear why this is the case, but it could be related to the difference in statistics between the two layers (the top layer has a much higher count rate) so that the top layer uncertainties contain a significant contribution from systematic errors.

The high values of χ_{red}^2 are consistent with the fact that the shape of the CTD template is not a good match to the τ_{meas} distribution in a significant fraction of pixels, even after least squares fitting. Two examples of this are shown in Figure A.2b. For pixel 102203, the amplitude of the peaks at the edges does not match well with the CTD template. For pixel 63314, the region in between the peaks at the edges does not match well. On the other hand,

	Det. 0	Det. 1	Det. 2	Det. 3	Det. 4	Det. 5
λ	1.14 ± 0.01	1.14 ± 0.01	1.25 ± 0.03	0.98 ± 0.02	1.34 ± 0.01	1.07 ± 0.01
Δ	-6.37 ± 12.02	5.97 ± 5.82	0.09 ± 7.74	5.55 ± 4.73	10.98 ± 7.56	9.49 ± 6.15
χ_{red}^2	3.20 ± 0.83	1.64 ± 0.40	1.81 ± 1.48	1.38 ± 0.30	2.09 ± 0.50	2.15 ± 0.49
	Det. 6	Det. 7	Det. 8	Det. 9	Det. 10	Det. 11
λ	1.14 ± 0.02	1.15 ± 0.01	1.15 ± 0.20	1.11 ± 0.02	1.17 ± 0.01	1.20 ± 0.01
Δ	9.87 ± 4.77	8.57 ± 5.81	6.09 ± 19.99	6.24 ± 8.19	10.43 ± 8.79	10.39 ± 5.46
χ_{red}^2	3.61 ± 1.14	2.05 ± 0.51	1.58 ± 0.72	1.55 ± 0.40	1.83 ± 0.45	2.94 ± 0.74

Table A.1: Average and standard deviation of fit parameters λ and Δ along with χ_{red}^2 across all pixels for each detector.

there are some pixels where the template shape is rather accurate, such as those in Fig. A.2a. This seems to imply that a single template shape is not enough to characterize all pixels in a detector. One possibility is that the depth distribution from the Monte Carlo simulation does not accurately represent the true depth distribution in some pixels. This should be expected to some degree since a single depth distribution is used for each detector, when in reality the depth distribution should not be exactly constant over the whole detector due to the near field geometric configuration². If the error is dominated by an incorrect depth distribution, then the overall depth calibration will not be significantly impacted, because the CTD template fit function is not very sensitive to the shape between the edges in the τ_{meas} distribution; besides the amplitude scaling factor, the CTD template can only stretch and offset along the τ axis. Thus, as long as the best fit CTD template edges line up well enough to the data, then the best fit values for λ and Δ will still give a good calibration.


However, if the differences between the τ_{meas} distributions and the CTD templates are dominated by unmodeled inhomogeneities that result in changes to the shape of $z_n(\tau_{sim})$ across a detector, then the best fit values for λ and Δ will not give a good calibration for pixels where the true $z_n(\tau_{sim})$ is significantly different from the $z_n(\tau_{sim})$ from step 3.

A.2 Depth Calibration: New Approach

A new depth calibration technique was developed in order to better account for inhomogeneities in the detector. The steps are as follows:

1. **Collect data with a calibration source and mimic the experimental configuration in a Cosima simulation.** This is identical to step 4 in the classic approach.
2. **For each detector, extract the simulated depth distribution and normalize it.** Call these normalized depth distributions $g_n(z)$ where n denotes the detector.

²Ideally, we would record the depth distribution in each pixel in the simulation, but this would require extensive CPU time to achieve good statistics in all 16428 pixels



(a) The shape of the CTD templates is a good match to the data for these two pixels.

(b) These two pixels are examples of pixels where the CTD template shape does not match up well to the data.

Figure A.2: Best fit CTD templates (red) to the τ_{meas} distributions (blue) for individual pixels from detectors 6 and 10.

3. **For every pixel, extract the τ_{meas} distribution from the real data and normalize it.** Call these normalized τ_{meas} distributions $f_p(\tau_{meas})$ where p identifies the pixel uniquely.
4. **Assume that the mapping between τ_{meas} and depth z is one-to-one; this implies that**

$$g_n(z)dz = f_p(\tau_{meas})d\tau_{meas} \rightarrow \frac{dz}{d\tau_{meas}} = \frac{f_p(\tau_{meas})}{g_n(z)} \quad (\text{A.2})$$

Another way of saying this is that because the mapping is one-to-one, each area element under $g_n(z)$ corresponds to an area element under $f_p(\tau_{meas})$.

5. **Numerically solve the ordinary differential equation in Equation A.2 for $z(\tau_{meas})$ for every pixel, and call the solution $z_p(\tau_{meas})$ where p denotes the pixel.** The boundary conditions were chosen by finding the edges of $f_p(\tau_{meas})$ and associating them with the detector edge z coordinates (See Figure A.3).

Once this depth calibration has been completed, then real, measured interactions are assigned depths by simply using the $z_p(\tau_{meas})$ solution for the corresponding pixel. Figure A.3 is a

graphical representation of this technique.

The algorithm for solving Equation A.2 is straightforward. $g_n(z)$ and $f_p(\tau_{meas})$ can be thought of as binned probability density functions over z and τ_{meas} respectively. The separability of Equation A.2 allows for a simple integration

$$\int_{\tau_{meas,0}}^{\tau_{meas}} f_p(\tau') d\tau' = \int_{z_0}^z g_n(z') dz' \quad (\text{A.3})$$

where $\tau_{meas,0}$ and z_0 are the initial conditions. z_0 is taken to be zero, representing the cathode side edge of the detector, while $\tau_{meas,0}$ is determined by finding the left-most edge of $f_p(\tau_{meas})$. The right hand side is integrated up to the anode edge of the detector to yield a cumulative density function $G_n(z)$. Then, for each 5 ns τ bin in $f_p(\tau_{meas})$, the left hand side integrand is evaluated, and a value of z is found (by inverting and linearly interpolating $G_n(z)$) such that both integrands are equal. This yields a series of (τ_{meas}, z) pairs which constitute the desired $z_p(\tau_{meas})$ profile for pixel p .

The major difference between the new approach and the classic approach is that the new approach does not utilize the electrostatic and charge transport models at all. The only model dependence in the new approach comes in through the Monte Carlo simulations. Therefore, the new approach is particularly reliant upon tight matching between the simulated and experimental configuration, and an accurate mass model.

It should be clear that this new approach is better able to account for inhomogeneities; whereas the classic approach used two parameters per pixel, λ and Δ , and assumed the same $z_n(\tau_{sim})$ for every pixel of a detector, the new approach generates a unique $z_p(\tau_{meas})$ for every pixel. Fig. A.4 shows several $z_p(\tau_{meas})$ profiles for a few pixels in detector 6. These specific pixels were chosen so as to illustrate inhomogeneity that cannot be described by just a stretching factor and an offset. The red cluster of $z_p(\tau_{meas})$ profiles corresponds to pixels closer to the center of the detector, while the blue cluster corresponds to pixels closer to the edge. We note that this trend has not been studied in detail and is currently only an observation. Regardless, this underscores the utility of the new approach, as the ‘‘one size fits all’’ paradigm of the classic approach cannot account for details such as these.

A.3 Angular Resolution

In the domain of Compton telescopes, angular resolution is defined as the FWHM of the ARM distribution, where the ARM for a single event is the smallest angular distance between the Compton cone edge and a specified source location (Figure 2.16). The angular resolution of Compton telescopes such as COSI is determined primarily by three factors:

1. **Position resolution.** The locations of the interactions must be known with high precision in order to accurately determine the orientation of the Compton cones. The 2 mm strip pitch of the COSI GeDs gives a position resolution of 2 mm in the x and y directions. The z position resolution is 0.2 mm RMS for most of the detector




Figure A.3: A diagram illustrating the new approach to the depth calibration. The top histogram is the depth distribution $g_n(z)$ in detector n , determined from Cosima simulations. The bottom histogram is the measured distribution of collection time differences $f_p(\tau_{meas})$ in pixel p . The area under each histogram has been normalized to 1. Assuming the map between depth and collection time difference is one-to-one, then each area element under $g_n(z)$ corresponds to an area element under $f_p(\tau_{meas})$. The area elements are denoted by the green regions under the histogram, and are not drawn to scale. Using this principle, the relationship between depth and collection time difference $z_p(\tau_{meas})$ can be determined for pixel p by solving Equation A.2. The dashed lines show the choice of boundary conditions; the edges of the detector are mapped to the edges of $f_p(\tau_{meas})$.




Figure A.4: $z(\tau_{meas})$ for 24 pixels of detector 6. The red pixels are closer to the center of the detector, while the blue pixels are closer to the edge. There is a clear difference in the shape between edge and center pixels that cannot be resolved by stretching and offsetting.

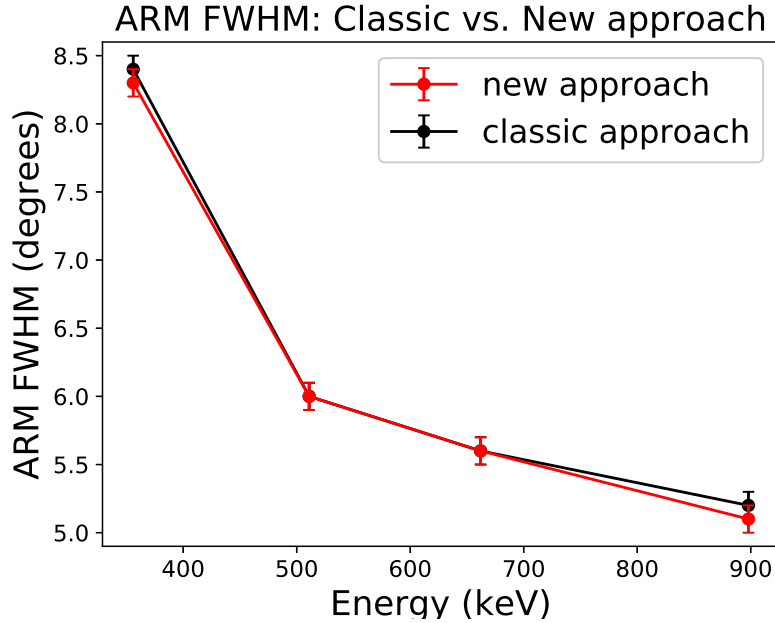


Figure A.5: Comparison between the simulated ARM and measured ARM using both depth calibration approaches as a function of energy. The calibration sources used were Ba-133 (356 keV), Na-22 (511 keV), Cs-137 (662 keV), and Y-88 (898 keV). The error bars are 90% confidence.

and somewhat worse near the edges of the detector, due to the steepening of $z(\tau_{meas})$ near the edges (Bandstra et al., 2006). COSI’s angular resolution is dominated by the position resolution, and in particular the x and y position resolutions.

2. **Energy resolution.** Application of the Compton scattering formula requires knowledge of the incident and scattered gamma-ray energies, which COSI can resolve accurately due to the exceptional spectral resolution of germanium (2.0 - 2.5 keV at 662 keV). COSI’s imaging performance is not significantly impacted by the energy resolution.
3. **Doppler Broadening.** Bound electrons which participate in Compton scattering are not initially at rest, but rather in an orbital state with unknown momentum. This introduces an error in the Compton scattering angle when applying Equation 2.1. Doppler broadening places a theoretical lower limit on the angular resolution, which for germanium is 2.85° at 200 keV, 1.25° at 500 keV, and 0.65° at 1000 keV (Zoglauer, 2005). As can be seen in Fig. A.5, COSI is not close to these limits.

Although the angular resolution for the current COSI GeDs is dominated by the xy position resolution, we set out to determine if the new depth calibration approach had any effect on the ARM FWHM as compared to the classic depth calibration. Figure A.5

shows a comparison of the measured ARM FWHM using both depth calibration approaches. Interestingly, both approaches result in almost the same performance, with the new approach possibly performing slightly better. This could mean that we have reached the best achievable angular resolution with the current COSI GeDs, or it could mean that something else in the analysis pipeline is limiting the angular resolution.

A.4 Sub-Strip Positioning Using Spectator Strips

We investigated the possibility of achieving x and y position resolutions finer than the COSI GeD strip pitch of 2 mm. This was motivated by the fact that the majority of energy collecting strips are accompanied by neighboring strips that do not collect energy, but do in fact trigger the fast shaper electronics and yield a timing value. This happens because charge collected to a strip will induce an instantaneous current on neighboring strips, but not a net charge (Knoll, 2010). The result is a fast pulse from the preamplifier which is attenuated in the slow shaper, but produces a signal large enough in the fast shaper to cross the fast threshold and trigger the zero-crossing detector.

Our charge transport model was used once again to determine if these “spectator” strips might be used to refine our xy positioning. We simulated energy deposits in the detector at various values of x while keeping y and z constant. All of the simulated x positions were below a single energy collecting strip, and the fast shaper signal was recorded for the neighboring spectator strip as a function of x . Figure A.6 shows the results of these simulations. The large black pulse is the fast shaper signal from the energy collecting strip, while the collection of multi-colored signals are fast shaper pulses as the value of x is varied over the energy collecting strip’s width (1.75 mm). Figure A.7 is similar to Figure A.6 except that the signals now correspond to a spectator strip that is two strips away from the energy collecting strip. Also, the pulse from the energy collecting strip is not shown in Figure A.7, although it is of the same amplitude as in Fig. A.6.

The amplitude of the spectator strip adjacent to the energy collecting strip exhibits a variation of a factor of two, with the maximum reaching 0.3 of the energy collecting strip’s pulse amplitude. The spectator strip two strips over from the energy collecting strip is much less sensitive to the x position of the energy deposits, showing almost no variation. Treating the fast shaper threshold as a discriminator threshold, the presence of a spectator strip on say the left side of the energy collecting strip and not on the right side could indicate that the energy was deposited under the left half of the energy collecting strip. Of course, this would only work for a relatively narrow range of energies where the spectator strip amplitude is varying in the vicinity of the fast threshold; about 150 - 300 keV for a fast threshold of 50 keV. Unfortunately, COSI’s current fast shaper channels are not equipped with any sort of pulse height measuring circuitry. If these circuits were present, then an xy position refinement would likely be possible over a larger energy range.

Similar studies using double sided strip germanium detectors have been carried out using digitizers to record the signal on each strip (Burks et al., 2004). The authors were able to




Figure A.6: Spectator strip amplitude vs x position of deposited charge under the energy collecting strip. The black pulse is the fast shaper signal from the energy collecting strip and the collection of colored pulses are the spectator strip pulses vs x .




Figure A.7: Spectator strip amplitudes vs x for a spectator strip that is two strips away from the energy collecting strip.

reduce the xy position resolution from 2 mm (the strip pitch) to 0.5 mm using the digitized spectator strip waveforms to integrate the spectator strip pulse, leading to a factor of two improvement in angular resolution. We note that this approach is probably not feasible for space flight due to the relatively high power consumption of digitizers and the sheer number of channels to read out in an instrument such as COSI. However, as mentioned before, a measurement of the pulse heights in the fast shaper electronics of COSI might be enough to improve the xy position resolution without employing waveform digitization.

A.5 Summary

We performed the depth calibration of the COSI GeDs using a classic, previously used approach, as well as a newly developed approach that is simpler in its implementation. Specifically, the classic approach required detailed charge transport simulations as an intermediate step, whereas the new approach completely avoids this. A benefit of the new approach is that it has the capability to account for inhomogeneities that are not modeled in the classic approach. We computed the ARM FWHM at several energies using both depth calibration approaches and found no significant change.

Finally, our work on xy position refinement is encouraging, but unlikely to yield a significant increase in position resolution outside of a relatively narrow energy band (150 keV - 300 keV) with the current COSI electronics. Future iterations of detectors and electronics may accommodate additional circuitry such as pulse height analyzers on the fast shaper channels, which will improve the prospects for sub-strip positioning.