

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Measuring the polarization of compact objects with the Compton Spectrometer and Imager

Permalink

<https://escholarship.org/uc/item/0zn566rj>

Author

Sleator, Clio

Publication Date

2019

Peer reviewed|Thesis/dissertation

Measuring the polarization of compact objects with the Compton
Spectrometer and Imager

by

Clio Sleator

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 Boggs, Co-Chair

Professor Stuart Bale, Co-Chair

Professor Eliot Quataert

Professor Bernard Sadoulet

Summer 2019

Abstract

Measuring the polarization of compact objects with the Compton Spectrometer and Imager

by

Clio Sleator

Doctor of Philosophy

University of California, Berkeley

Professor Steven Boggs, Co-Chair

Professor Stuart Bale, Co-Chair

γ -rays are an indispensable probe of neutron star and black hole systems, some of the most extreme environments in our universe. The penetrating power of γ -rays allows us to probe deep within these often obscured systems. In particular, polarization measurements of neutron stars and black holes can give essential clues about γ -ray emission mechanisms, source geometry, and magnetic field structure where spectral, temporal, and imaging analysis fall short.

The Compton Spectrometer and Imager ([COSI](#)) is one of the few γ -ray telescopes designed as a polarimeter. [COSI](#) is a wide-field balloon-borne soft γ -ray (0.2-5 MeV) telescope, and one of its main goals is to measure the polarization of γ -rays emitted by compact objects. As a Compton telescope, [COSI](#) is inherently sensitive to polarization: polarized photons preferentially Compton scatter orthogonally to their polarization direction. In May of 2016, [COSI](#) was launched from Wanaka, New Zealand, on NASA's new super pressure balloon and flew for 46 days before the flight was terminated in Peru. The Crab nebula, Cygnus X-1, and Centaurus A are among the compact objects detected during the 2016 flight.

A key step to performing imaging, spectral, and polarization analysis of the sources detected during the 2016 flight is to accurately simulate the detector response. To do so, I developed a detailed detector effects engine which applies the intrinsic detector performance to Monte Carlo simulations. With

accurate simulations of the instrument in place, I developed a spectral analysis pipeline for sources detected by COSI. As needed for spectral analysis and polarimetry, I developed a background subtraction technique for broadband, persistent sources that utilizes the COMPTEL data space. I verified these new analysis methods using the COSI observation of GRB 160530A: after subtracting the background using the COMPTEL data space method and fitting the spectrum, the spectral parameters are consistent with those measured by another instrument.

I used the COMPTEL data space background subtraction algorithm to fit the Crab spectrum, but concluded that the algorithm is limited in estimating the background accurately enough in cases where the source is background-dominated. I used simulations to assess the prospects for polarimetry of COSI's observation of the Crab and inferred that we require a better mechanism of albedo radiation background rejection to perform polarimetry of this particular observation. I note that similar observations from the COSI instrument on a satellite platform would lead to more success in measuring the spectra and polarization properties of compact objects, and that the analysis methods developed in this work could be easily applied.

To my father, Tycho Sleator, for beginning my physics education before I
can even remember, and for always keeping it fun.

ACKNOWLEDGMENTS

Thank you to my advisor, Professor Steve Boggs, for your support, your encouragement, and your constant positivity – I needed it! Thanks for your excellent insights into my work with the detectors and shields and for your helpful comments on this thesis. Throughout my time working with you, I always felt like you had my back, and I hope to emulate your thoughtfulness toward any students that I work with in the future.

Thank you to Dr. Andreas Zoglauer for the [MEGALib](#) tech support and for helping me through each painstaking step of the analysis. More importantly, thank you for teaching me how to view my work in a positive light. I don't think I could have gotten through this without you, and the confidence you have helped me instill in myself is a much more important lesson than the details of detector effects and the [COMPTEL](#) Data Space. I'll miss having you right across the hall.

Thank you to Dr. John Tomsick for teaching me about compact objects and how to analyze [NuSTAR](#) data, and for helping me get my papers out the door in a timely manner. Thanks for planning group lunches, holiday parties, and best of all, the group ski trips (and thanks for some great runs down Burnout!).

Thank you to Steve McBride for coming out of retirement to help make [COSI](#) happen, for supporting me during my time in grad school, and for telling me about the tradition of the [GSE](#) help widget. Your hilarious stories of ballooning from before I was born kept campaigns extremely entertaining.

Thank you to Brent Mochizuki for your hard work on [COSI](#) and for your constant support and friendship. Working with you has been an absolute joy. I appreciate the time you spent teaching me about [FPGAs](#), telemetry streams, and other parts of your job just because I was curious. Also, thank you for always being there to listen when the going got rough; it helped a lot. I hope we work together again in the future (although let's try to avoid thermal vacuum tests going forward).

Thank you to Dr. Alex Lowell and Dr. Carolyn Kierans, my fellow grad students for most of my time on [COSI](#). Alex, thank you for keeping me in the loop about the latest cool thing you've realized that [scipy](#) can do, for never tiring of my many questions about [COSI](#), and for laughing with me about vans and really powerful fans. Carolyn, after sitting next to you every day for years

and sharing rooms on many travels, we could only either become great friends or unable to stand each other, and I'm so glad it was the former. Thank you for working out many COSI hardware and analysis problems with me and for introducing me to the exciting worlds of yarn and loose leaf tea. I'm excited to collaborate in the future and to write you many GSEs in the years to come.

Thank you to Dr. Thomas Siebert, Jacqueline Beechert, and Hadar Lazar for helping me proofread and improve this thesis. Thanks to Dr. Jeremy Hare for the excellent advice of starting to write this thesis early.

Thank you to everyone else at SSL who helped me out over the years. In particular, thanks to Steve Marker for allowing our dirty balloon hardware in your vacuum chamber, for helping me fix the pump cart again and again, and for trading mountain bike stories. Thanks to Lissvett Fields for your friendly presence in the lab and for giving me so many free lunches. Thanks to Anne Takizawa and Joelle Miles in the physics department for helping me navigate the logistics of grad school.

Thank you to my friends in Berkeley, New York, and wherever else you may be. Thanks to my brother Isaac for thinking that COSI is cool even though you're basically a rock star. Thank you to my parents for supporting me in being whatever type of scientist I want to be and for not freaking out too much when I had to miss Christmas to go to Antarctica. And lastly, thank you to Ian Hill for all of your support and for helping me generate enough contentment to stave off the dread.

CONTENTS

I	BACKGROUND MATERIAL	1
1	POLARIZATION OF COMPACT OBJECTS	2
1.1	Polarization in astrophysics	2
1.2	Pulsars and pulsar wind nebulae	3
1.2.1	Pulsed emission	4
1.2.2	Pulsar wind nebulae	6
1.2.3	γ -ray polarization of the Crab pulsar and nebula . .	7
1.2.4	Future polarization measurements of pulsars	9
1.3	Galactic black holes	10
1.3.1	Identifying black hole binaries	10
1.3.2	Accretion onto black holes	12
1.3.3	Spectral states	13
1.3.4	The origin of the γ -rays	13
1.3.5	γ -ray polarization of Cyg X-1	15
1.3.6	Future polarization measurements of Galactic black holes	17
1.3.7	Active galactic nuclei	17
1.4	Summary	17
1.5	Thesis outline	18
II	THE COMPTON SPECTROMETER AND IMAGER: INSTRUMENT OVERVIEW, CALIBRATIONS AND SIMULATIONS, AND THE 2016 FLIGHT	21
2	THE COMPTON SPECTROMETER AND IMAGER	22
2.1	Compton telescopes	22
2.1.1	Types of Compton telescopes	23
2.1.2	Angular resolution measure	25
2.1.3	Compton imaging	27
2.2	The COSI Instrument	27
2.2.1	Germanium detectors	28
2.2.2	Cryostat and cryocooler	29
2.2.3	Anti-coincidence shielding	31
2.2.4	Signal readout	33

	2.2.5	Gondola systems	35
2.3		Instrument calibrations	40
	2.3.1	Energy calibration	40
	2.3.2	Crosstalk and charge loss	40
	2.3.3	Strip pairing	41
	2.3.4	Depth calibration	45
2.4		Data Analysis	46
	2.4.1	Data analysis pipeline	46
	2.4.2	Event selections	47
3		COSI'S DETECTOR EFFECTS ENGINE	49
	3.1	Overview of the simulation pipeline	49
		3.1.1 Mass model	49
		3.1.2 Monte Carlo Simulations	52
		3.1.3 Detector effects engine	53
	3.2	The COSI DEE	53
		3.2.1 Shield veto	54
		3.2.2 Physical position to strip and detector ID	55
		3.2.3 Charge sharing	55
		3.2.4 Depth to timing	59
		3.2.5 Charge loss	62
		3.2.6 Crosstalk	65
		3.2.7 Energy to pulse height	66
		3.2.8 Thresholds and dead strips	68
		3.2.9 Guard ring veto	70
		3.2.10 Trigger conditions	70
		3.2.11 Dead time	70
	3.3	Simulation benchmarking	72
		3.3.1 Trigger rate comparison	72
		3.3.2 Spectrum comparison	74
		3.3.3 Angular resolution measure comparison	81
	3.4	Conclusions	86
4		THE 2016 COSI BALLOON CAMPAIGN	88
	4.1	Historical context	88
	4.2	2016 flight overview	90
	4.3	Flight operations	92
		4.3.1 GSE	93
		4.3.2 Alert software	95
		4.3.3 COSI plotter	96

4.4	Thermal Environment	96
4.5	High voltage failures	97
4.6	Astrophysical sources detected	100
III	ASTROPHYSICS ANALYSIS	102
5	SPECTRAL ANALYSIS WITH COSI	103
5.1	Spectral fitting and the response matrix	103
5.1.1	COSI data with XSPEC	104
5.1.2	Pipeline verification with simulations	106
5.2	Modeling the atmospheric absorption	107
5.3	Spectral analysis of GRB 160530A	108
5.3.1	The COSI observation of GRB 160530A	109
5.3.2	Spectral analysis	110
5.4	Challenges of persistent sources	113
6	THE COMPTEL DATA SPACE	119
6.1	Introducing the COMPTEL data space	119
6.1.1	On-axis sources in the CDS	121
6.1.2	Galactic coordinates and off-axis sources in the CDS	123
6.2	Background subtraction with the CDS	125
6.2.1	Background in the CDS	128
6.2.2	Scaling the background in the CDS	129
6.2.3	CDS-ARM tailing correction	131
6.3	Spectral background subtraction validation	133
6.3.1	Background only	134
6.3.2	Point source with background	134
6.3.3	Dependence of the scale factor	142
6.4	Background subtraction applied to flight data	143
6.4.1	Measured background compared to the Ling Model	143
6.4.2	Simulated and measured CDS-ARM distributions	144
6.4.3	Moving forward with the simulated CDS-ARM distribution	150
6.4.4	CDS background subtraction of GRB 160530A	151
6.5	Conclusions	153
7	THE COSI OBSERVATION OF THE CRAB NEBULA	154
7.1	Observation overview	154
7.2	Spectral analysis	158
7.3	Polarization analysis	162
7.3.1	Compton polarimetry with COSI	163

7.3.2	Polarimetry of simulated Crab observation	166
7.4	Conclusions	171
8	THE REFLECTION SPECTRUM OF THE LOW-MASS X-RAY BI- NARY 4U 1728–34	172
8.1	Background	173
8.2	Observations and data reduction	173
8.3	Analysis and results	174
8.4	Discussion	184
9	SUMMARY AND OUTLOOK	187
	BIBLIOGRAPHY	189

LIST OF FIGURES

Figure 1.1	A schematic of the pulsar electrodynamics model from Goldreich & Julian (1969) . The magnetic field is aligned with the pulsar spin axis. The pulsar is in the bottom left-hand corner. The closed magnetic field lines corotate with the pulsar within the light cylinder. The magnetic field lines that extend past the light cylinder are open, and particles are accelerated along them into the nebula. <i>Figure from Goldreich & Julian (1969).</i>	4
Figure 1.2	A schematic of a pulsar whose rotation axis (marked by the spin frequency Ω) is not aligned with its magnetic field \mathbf{B} . Three models for emission sites of the high energy radiation are marked. <i>Figure from Harding (2004).</i>	5
Figure 1.3	The Crab nebula pictured (a) in optical light by the Hubble Space Telescope (HST) and (b) in soft X-rays by Chandra X-ray Observatory (Chandra). The nebular structure clearly varies based on wavelength.	6
Figure 1.4	The γ -ray polarization angle superimposed onto the Chandra image (blue) and the HST image (red) of the Crab nebula, drawn such that the vector passes through the pulsar position. The polarization angle is very well aligned with the inner jet structure present in the Chandra image. <i>Figure from Dean et al. (2008).</i>	8
Figure 1.5	A schematic of 21 stellar mass black hole binary candidates within our galaxy drawn to scale relative to the distance between the Sun and Mercury. <i>Figure from Narayan & McClintock (2015).</i>	11
Figure 1.6	A schematic of the spectral states of black hole binaries. Since the figure was made, the high state and low state have been renamed the soft state and hard state, respectively. The dots represent the accreting material and the horizontal bars represent the accretion disk. The amount of mass transfer is indicated by the \dot{m} scale. <i>Figure from Esin et al. (1997)</i>	14

Figure 1.7	Spectral fits of Cyg X-1 data using a model in which (a) the corona and (b) the jet is the emitter of the observed high energy radiation. The two models fit the data equally well, indicating that spectral analysis is not able to distinguish between the two models. <i>Figure from Markoff et al. (2005).</i>	15
Figure 1.8	The spectrum of Cyg X-1 as measured by Imager on Board the INTEGRAL Satellite (IBIS) along with the two components required to fit this spectrum. The modulation curves corresponding to the polarization measurements are shown for two energy ranges (more modulation means a higher polarization fraction). The energy dependence of the polarization signatures indicates two distinct emission mechanisms that produce γ -rays. <i>Figure summarizing Figures from Laurent et al. (2011), from ESA press release.</i>	16
Figure 2.1	The cross sections of various interactions in Germanium. Compton scattering dominates between 0.2 and 10 MeV. The cross section data is from NIST XCOM.	23
Figure 2.2	A schematic of COMPTTEL. A photon is shown scattering in the top plane and traveling to the bottom plane, where it is photoabsorbed. Figure from Schönfelder et al. (1993).	24
Figure 2.3	A schematic of a compact Compton telescope. A photon originates from the yellow star and interacts three times in the detector. First the photon Compton scatters at \vec{r}_1 and deposits energy E_1 . Next it Compton scatters at \vec{r}_2 and deposits energy E_2 . Finally the photon is photoabsorbed at \vec{r}_3 and deposits its remaining energy E_3 into the detector. From the energy deposited at each interaction location, we can determine the initial Compton scatter angle ϕ and restrict the origin of the photon to somewhere on the red circle.	26
Figure 2.4	(a) A schematic demonstrating how the angular resolution measure (ARM) is calculated. The source location is the blue dot and the Compton circles for three events are represented in red. (b) An example ARM distribution measured by Compton Spectrometer and Imager (COSI). The data is from a ^{137}Cs calibration source that emits γ -rays at 661.66 keV.	27

Figure 2.5	A backprojection of about 10 events. The Compton circles are visible in red. The bright hot spot at the center of the image is at the source location, where most of the Compton circles overlap.	28
Figure 2.6	(a) A single COSI germanium detector (GeD) in front of a mirror, showing the orthogonal cross strips: the 37 strips on the front are vertical and the 37 strips on the back (in the mirror) are horizontal. Each GeD is $8 \times 8 \times 1.5$ cm ³ . (b) The COSI GeD detector array, showing the $2 \times 2 \times 3$ configuration. The GeDs are mounted to a copper cold finger, which is the U-shaped bar dividing the array vertically.	29
Figure 2.7	A photograph of the COSI cryostat. The footprint of the aluminum shell of the cryostat is about 19×18 inches. The cooling fluid jacket around the cryocooler body enables fluid to be pumped around the cryocooler as part of the liquid cooling system. The preamplifiers, not attached to the cryostat in this photo, connect to the flex circuits.	30
Figure 2.8	A photograph of the cryocooler before integration into the cryostat and without the cooling fluid jacket.	30
Figure 2.9	(a) A schematic and (b) a photograph of the liquid cooling system mounted on the gondola along with the cryostat and CsI shields. The pumps circulate cooling fluid from the reservoir around the cryocooler and through the copper radiator plate, where the heat is radiated away. The four side CsI shield pieces are visible, and two more sit underneath the cryostat.	32
Figure 2.10	Count rate versus threshold DAC for ²⁴¹ Am, ⁵⁷ Co, and background as measured by shield piece 1. After subtracting the background from the source distributions, we determine the energy to threshold DAC relationship by finding the mean of the resulting complementary error functions. The mean corresponds to the energy emitted by the calibration source in question.	34

Figure 2.11	The relationships between energy and threshold DAC for each of the six shield pieces. The starred points represent the two calibration source energies at 59.5 keV and 122.0 keV. The energy- DAC relationship is shown for the full threshold DAC range. Shield piece 3 is a clear outlier due to its exceptionally poor energy resolution.	34
Figure 2.12	(a) A full preamplifier box (left) next to a partially built box (middle) and a single preamplifier board (right). Each box contains four boards. (b) The cryostat with the preamplifiers mounted and connected to the flex circuits. Each detector connects to two preamplifiers, one for each side. The coaxial ribbon signal cables are wrapped in white teflon and are plugged into the preamplifier output.	36
Figure 2.13	A single card cage with one of the analog boards pulled out, showing the 10 pulse shaping circuits. Above the eight analog boards is the DSP board, which handles the trigger logic and relays the data to the flight computer. The next board is the low voltage power supply, which powers both the rest of the card cage and the preamplifiers. The top board is the high voltage power supply which provides 1000 – 1500 V to the detectors.	36
Figure 2.14	A photograph of the fully integrated COSI gondola before the 2016 flight from Wanaka, New Zealand. The side panels are removed to provide a view of the electronics bay and Support Instrument Package (SIP) bay. . .	37
Figure 2.15	(a) Spectrum of multiple lines from a variety of calibration sources that is used to determine the ADC -energy relation and (b) the ADC -energy relation for strip 7 on the AC side of detector 7 (the specific strip was chosen arbitrarily).	41
Figure 2.16	(a) If one interaction occurs in a single detector, determining the interaction location (the green circle) is straightforward. (b) If two interactions occur in a single detector, there are two solutions: the interaction locations could be at the red or the green circles. By comparing the energies deposited on each strip, it is evident that the green circles are the correct solution.	42

Figure 2.17	Strip pairing is made more complicated when (a) charge sharing between two adjacent strips occurs or (b) there are multiple interactions on one strip. The green circles indicate the interaction positions, which the algorithm recovers by allowing two strips on one side to be paired with one strip on the other.	43
Figure 2.18	Plot of the DC side energy versus the AC side energy for each event of a ^{22}Na calibration source. The hot spots indicate the line energies at 511 keV and 1.274 MeV. The horizontal and vertical lines emerging from the hot spots are due to charge loss. The offsets near the 1.274 MeV line are the result of crosstalk.	44
Figure 2.19	Schematic of the analysis pipeline used for <i>COSI</i>	46
Figure 3.1	The mass model of the full cryostat, including the 12 <i>GeDs</i> , cryocooler, preamps, etc., as well as the shields and photomultiplier tubes (<i>PMTs</i>).	50
Figure 3.2	The mass model of the full gondola frame and antenna booms, including the source holder and source salt. The source holder is floating above the gondola and is positioned (a) directly above the cryostat and (b) 50° off axis.	51
Figure 3.3	Photograph of the 80-20 structure used for calibrations. The structure allows for three sources to be placed simultaneously; two of the source holder pieces are shown here.	52
Figure 3.4	Schematic showing the data acquisition time for the <i>GeD</i> hits.	54
Figure 3.5	The fraction of total events that have n strip hits per event, as a function of the number of strip hits per event, for measurements (blue) and simulations (red). The included events are from the 662 keV line of a ^{137}Cs source; their total energy is between 640 and 680 keV to include events affected by charge loss and crosstalk. No charge sharing is explicitly simulated in (a), whereas in (b), charge sharing due to diffusion was added to the simulations. Charge sharing due to diffusion brings the simulated distribution slightly closer to the measured distribution, but does not have a large effect.	57

Figure 3.6	Measured (blue) and simulated (red) distribution of the number of strip hits per event. The included events are from the 662 keV line of a ^{137}Cs source. Charge sharing was empirically applied to the simulations as described in Equation 3.2, and there is good agreement between measurements and simulations.	59
Figure 3.7	Measured (blue) and simulated (red) distribution of the number of strip hits per event for a variety of calibration runs. The energy cut for each line is given in parentheses. The good agreement between measurements and simulations shown here indicates that the empirical charge sharing simulation method described in this section works well for all energies.	60
Figure 3.8	Measured (blue) and simulated (red) depth distributions for the top layer of detectors (a) only including hits with one strip hit per side and (b) including all hits. The distributions for the middle and bottom detector layers have the same shape.	61
Figure 3.9	Sum-difference histogram of two site events from the AC side of detector 0 within a depth of 0-0.5 cm for (a) the 662 keV line of a ^{137}Cs source and (b) the 122 keV line of a ^{57}Co source. In (a) charge loss is evident by the hot spot at large differences: if no charge loss were present, the sum of all the counts in this hot spot would be at ~ 662 keV. In (b), charge loss is evident in the parabolic dip below 122 keV.	63
Figure 3.10	The measured (blue) and simulated (red) spectra of the 511 keV line from a ^{22}Na source (a) before and (b) after charge loss is applied to the simulation. The spectra only contain hits with two adjacent strip hits from the AC side of the detector. Adding charge loss to the simulation significantly improves the agreement between measurements and simulations.	64

Figure 3.11	Measured (blue) and simulated (red) spectra (a) before and (b) after crosstalk is added to the simulation. The spectra, made before the crosstalk correction, only contain hits with two adjacent strip hits on the DC side of the detectors to clearly view the effect of crosstalk. When crosstalk is added to the simulation, the same line shift is present as in the measurement.	66
Figure 3.12	The spectra of events with energy only and events with both energy and timing for two arbitrary strips. Events below the fast threshold do not have a timing measurement. The fast threshold is noisy, so there is no clear cut-off between energy-only events and energy-and-timing events. The energy threshold is not noisy, and results in the sharp cutoff to the spectrum located at around analog-to-digital converter (ADC) value 90 for both strips.	69
Figure 3.13	The energy-and-timing spectrum divided by the total spectrum for two strips. These distributions represent the fraction of events with timing as a function of ADC value. The distributions are fit with an error function, which is used to simulate the fast threshold.	69
Figure 3.14	A comparison of the measured (blue) and simulated (red) energy-only spectrum (a,b) and energy-and-timing spectrum (c,d) for two strips. There is good agreement between data and simulations.	71
Figure 3.15	A comparison of the measured (blue) and simulated (red) trigger rate for a variety of calibration sources and off-axis angles. The trigger rates differ between the top, middle, and bottom layers of detectors, as detectors closer to the top collect more photons. The detector placement is as follows: top: 0, 5, 6, 11; middle: 1, 4, 7, 10; bottom: 2, 3, 8, 9.	73
Figure 3.15	<i>Continued</i> – A comparison of the measured (blue) and simulated (red) trigger rate for a variety of calibration sources and off-axis angles. The trigger rates differ between the top, middle, and bottom layers of detectors, as detectors closer to the top collect more photons. The detector placement is as follows: top: 0, 5, 6, 11; middle: 1, 4, 7, 10; bottom: 2, 3, 8, 9.	74

Figure 3.16	A comparison of the measured and simulated spectrum of (a) a ^{133}Ba source and (b) a ^{137}Cs source. The line shapes match very well, but the number of counts in the lines differs.	75
Figure 3.17	The ratio of measured to simulated counts in the lines as a function of energy and off-axis angle.	76
Figure 3.18	A comparison of the measured and simulated spectra of (a) a ^{133}Ba calibration source and (b) a ^{137}Cs calibration source, after some of the fully absorbed events have been removed. The line shapes and heights match very well. . .	76
Figure 3.19	A comparison of the ratio of measured to simulated counts in the lines as a function of energy and off-axis angle, after some of the fully absorbed events are removed from the simulations.	77
Figure 3.20	A comparison of the measured and simulated full width at half maximum (FWHM) of each calibration source line as a function of energy and off-axis angle. The relatively large difference between the measured and simulated FWHM at 511 keV is due to the assumption in the simulations that the electron-positron pair that produces the 511 keV γ -rays is at rest, when in reality the electron and positron have some momentum; this non-zero momentum leads to the broadening of the line.	77
Figure 3.21	The distribution of $z = (N_M - N_S)/\sigma_M$ when (a) the simulations are processed with the standard Medium Energy Gamma-ray Astronomy library (MEGALib) detector effects engine (DEE) and (b) the simulations are processed with the COSI -specific DEE . The match between measurements and simulations is much better when the simulations are processed with the COSI -specific DEE : the distribution is about four times narrower and centered near zero.	81
Figure 3.22	A comparison of the measured and simulated ARM distributions of the 662 keV line of a ^{137}Cs source.	82
Figure 3.23	A comparison of the measured and simulated ARM FWHM as a function of energy and off-axis angle.	83

Figure 3.24	A comparison of the measured and simulated distribution of distances between the first and second interaction for (a) the 511 keV line of a ^{22}Na source and (b) the 1.836 MeV line of a ^{88}Y source. The discrepancies below 2.5 cm contribute to the differences between the measured and simulated the ARM FWHM	85
Figure 3.25	A comparison of the measured and simulated distribution of initial Compton scatter angles for (a) the 356 keV line of a ^{133}Ba source and (b) the 1.3 MeV line of a ^{60}Co source. These distributions agree better at high energies than at low energies. The measured and simulated ARM FWHM are in better agreement at low energies when only events with Compton scatter angles $> 40^\circ$ are used.	86
Figure 4.1	COSI on the flight line hanging from the launch vehicle and connected to the parachute and super pressure balloon (SPB). Only the top of the SPB is filled on the ground, but as the balloon reaches the float altitude of ~ 33.5 km, the helium expands and fills the entire balloon, resulting in a distinctive pumpkin-shaped balloon.	90
Figure 4.2	COSI 's flight path during the 2016 flight. The green path shows the first circumnavigation, and the red path shows the second partial circumnavigation. The flight began in Wanaka, New Zealand and ended in southern Peru 46 days later. The balloon drifted further north than expected, allowing for some exposure of the northern skies.	91
Figure 4.3	Altitude profile of COSI during the 2016 flight. The launch is shown on May 17. The drops in altitude early in the flight can be attributed to cold storms. The drops in altitude toward the middle and end of the flight are due to an issue with the balloon.	92
Figure 4.4	A photograph of the COSI payload in at its landing site in Peru, taken by the recovery team. Though the gondola, antenna booms, and solar panels were completely destroyed, the rest of the instrument was unharmed. The GeDs , cesium iodide (CsI) shields, card cages, and other subsystems have been tested in the lab since the flight and are fully operational.	93

Figure 4.5	A sampling of the ground support equipment (GSE) widgets.	94
Figure 4.6	An example COSI monitor text message, alerting us that the cryocooler collar temperature (temperature # 25) is outside of the acceptable range of 25° to 45°.	96
Figure 4.7	The temperatures of two subsystems within the electronics bay, card cage # 7 and the flight computer CPU board, and the exterior of one of the side panels during one 24-hour period. The large difference in day- and night-time ambient temperatures is apparent by the large swing in the side panel temperature, yet the temperatures of the subsystems within the electronics bay remain relatively constant. A heater was attached to the flight computer CPU board and was automatically toggled on and off between 12:00 and 16:00 Coordinated Universal Time (UTC).	97
Figure 4.8	The temperatures of the cryocooler and the sides of the radiator plate where the fluid is input and output during one 24-hour period. The difference in day- and night-time temperatures is apparent in the radiator plate temperatures. The oscillations in the cryocooler and radiator input temperature are due to the pump turning on and off, which occurred more frequently during the day when it was warmer.	98
Figure 4.9	The preamp currents and high voltage supply reading during the first failure of Detector 8 in flight. Until ~18:55, the preamp currents behaved normally. The failure occurred between ~18:55 and 19:20, when the currents oscillated between nominal and high values (when the preamp currents are higher than normal, the detector is inoperable). We ramped the high voltage down at ~19:20, at which point the positive preamp current went high and the negative preamp current went low; this is normal behavior during high voltage ramp down.	99
Figure 4.10	The exposure map from the entire flight, assuming an effective area of 20 cm ² . The labeled sources are known γ -ray sources in COSI's field of view (FOV) but were not necessarily detected by COSI.	100

Figure 5.1	The (a) redistribution matrix and (b) effective area of a simulated on-axis point source emitting a power law spectrum with a photon index of 2 detected by COSI. The event selections used for this response matrix are shown in Table 5.1. Note that the measured energy in the redistribution matrix is cut off after 3 MeV; this is due to the energy cut in the event selections.	105
Figure 5.2	Spectral fit of a simulated Crab-like source without background. A comparison of the simulated and fit parameters is shown in (b). The fit results are consistent with the simulated parameters.	107
Figure 5.3	A schematic of our column density calculation. We assume a spherical atmosphere with an altitude of 100 km. We then integrate the NRLMSISE-00 model (Picone et al., 2002) along the dotted line between COSI and the source to compute the column density.	108
Figure 5.4	The effects of (a) the zenith angle of the source relative to COSI and (b) COSI's altitude on the column density. In (a), we assume an altitude of 33 km. In (b), we assume two source positions relative to COSI, on-axis (blue) and a 40° zenith angle (green).	109
Figure 5.5	The COSI image of gamma-ray burst (GRB) 160530A using 10 iterations of the list mode maximum-likelihood expectation-maximization (LM-MLEM) algorithm. The color scale is in arbitrary units.	110
Figure 5.6	The measured light curve of GRB 160530A, where (a) shows the times used for source and background events and (b) shows a zoomed-in version of the GRB window from (a), revealing the structure in the light curve.	111
Figure 5.7	The (a) redistribution matrix and (b) effective area computed for spectral analysis of GRB 160530A, using the event selections detailed in Table 5.2. These plots are similar to those shown in Figure 5.1 but with different event selections and a source zenith angle of 43.5° instead of 0°.	113

Figure 5.8	The COSI spectrum of GRB 160530A fit with the atmospheric absorption model multiplied by the Band function. The spectrum was binned such that each bin had the same signal-to-noise ratio $S/N = 3$	114
Figure 5.9	A schematic of the image-space background subtraction method. The source region is represented in red and the background regions are represented in blue. Many background regions are chosen for more statistics. The background regions are at the same zenith angle as the source region and are spaced 30° apart in azimuth, though there is a 60° gap between the source region and the nearest background regions, resulting in a total of nine background regions.	116
Figure 5.10	The results of the image-space background subtraction method applied to a simulation of a point source <i>without</i> background. The events in the background spectrum are source events that are consistent with background regions. The background spectrum is averaged over the nine background regions. The subtracted source spectrum contains significantly fewer counts than the unsubtracted source spectrum, indicating that we are unable to correctly distinguish between source and background events in image space.	117
Figure 6.1	(a) The first two interactions of a Compton event, showing the angles that define the COMPTEL Data Space (CDS). The photon originates from the source at position (χ_0, ψ_0) and scatters at \vec{r}_1 and \vec{r}_2 . The scatter direction \vec{d}_γ is the vector between \vec{r}_1 and \vec{r}_2 . The angles χ and ψ represent the polar and azimuth angles of \vec{d}_γ , respectively. The initial Compton scatter angle ϕ is defined as the angle between the initial trajectory of the photon and \vec{d}_γ . (b) A schematic of a point source in the CDS . The apex of the cone is at $(\chi = \chi_0, \psi = \psi_0, \phi = 0)$, where (χ_0, ψ_0) is the position in detector coordinates from which the photon originated. <i>Figures from Kierans (2018).</i>	120

Figure 6.2	An on-axis source in the CDS. (a) An event that scatters at a small Compton scatter angle ϕ_1 with the resulting χ_1 . (b) An event that scatters at a large Compton scatter angle $\phi_2 > \phi_1$ with the resulting χ_2 . (c) The $\chi - \phi$ plane in the 3D CDS. (d) The source region defined by the ideal line $\phi = \chi$ in the 2D CDS, which results from integrating the 3D CDS along the ψ direction. The source region has some width related to the instrument's angular resolution. The locations of the two events shown in (a) and (b) in the 2D CDS are marked. <i>Figures from Kierans (2018).</i>	122
Figure 6.3	A simulation of an on-axis point source in the 2D CDS. The line at $\phi = \chi$ is clearly visible and has a thickness related to COSI's angular resolution. The events exterior to this line are likely incompletely absorbed or improperly reconstructed events.	123
Figure 6.4	A schematic of the source and background regions in the 2D CDS. The source region includes events within Δ of the ideal line $\phi = \chi$. The background regions are on either side of the source region and include events between Δ and 2Δ from the ideal line.	126
Figure 6.5	A distribution of distances of events from the ideal line $\phi = \chi$ in the 2D CDS. This distribution, referred to as the CDS-ARM distribution, is analogous to the ARM distribution in image space and describes COSI's angular resolution. This particular CDS-ARM distribution was made from a simulation of a point source. The source and background regions SR and BR are marked for a Δ of 5.1° , which is the FWHM of this particular distribution.	127

Figure 6.6	The background from a single day of flight (June 13, 2016) in the 2D CDS. The CDS has been rotated such that the Crab position is at the origin. The locations of the source and background regions, defined here with $\Delta = 10^\circ$, are indicated. The blue points represent all of the Compton events from that day. The green points are all the Compton events with the Earth horizon cut (EHC) applied to the data. The red points represent the remaining Compton events after applying a 40° pointing cut around the Crab position.	128
Figure 6.7	The results of the background subtraction performed on a simulation of background only for a variety of “source” positions. The “source” positions were chosen by selecting positions that had good exposure during the 2016 flight from the exposure map, other than position (a) where the “source” position is the Crab position. The subtracted spectrum for each of the positions is consistent with no source. The N_S value written in the upper right hand corner of each plot is the total number of subtracted counts (i.e. an integral over the background-subtracted spectrum).	135
Figure 6.8	The values of α and the 0.1 - 3 MeV flux as a function of the number of iterations performed. The dashed red lines indicate the simulated values: $\alpha = -0.93$ and 0.1-3 MeV flux = $2.2 \gamma/\text{cm}^2/\text{s}$. The parameters converge after six iterations.	137
Figure 6.9	The final results of the background subtraction and spectral fit of the simulation of GRB 160530A. The spectrum is binned such that each bin has a signal-to-noise ratio $S/N = 3$. A comparison of the simulated and best fit parameters is shown in (b); the fit results are consistent with the simulated values. The uncertainties on the fit are 99% confidence.	139

Figure 6.10	The results of the background subtraction and spectral fit of the simulation of the “strong” Crab-like source. The spectrum is binned such that the signal-to-noise ratio in each bin is $S/N = 3$. A comparison of the simulated and best fit parameters is shown in (b); the fit results are consistent with the simulated values. The uncertainties on the fit are 99% confidence.	141
Figure 6.11	The results of the background subtraction and spectral fit of the simulation of the “weak” Crab-like source. The spectrum is binned such that the signal-to-noise ratio in each bin is $S/N = 2$. A comparison of the simulated and best fit parameters is shown in (b); the fit results are consistent with the simulated values. The uncertainties on the fit are 99% confidence.	141
Figure 6.12	The results of the background subtraction of the simulation of background only for two different “source” positions when the dependence of f on Θ, Φ, E , and ϕ is ignored. The resulting spectra are not consistent with a flat spectrum at zero. The N_S value in the upper right hand corner of the plots is the total number of subtracted counts.	142
Figure 6.13	A comparison of (a) the first power law index Γ_1 , (b) the second power law index Γ_2 , and (c) the break energy E_{break} for each day of flight and the Ling Model simulations. The red line represents the best fit value for the Ling Model simulation, and the red shaded region represents the error on the fit. The best fit parameters of the Ling Model simulations fall within the spread of the flight data, indicating that the Ling Model is a good approximation of the atmospheric background.	145
Figure 6.14	The (a) measured and (b) simulated CDS-ARM distributions of GRB 160530A. The measured distribution is significantly wider.	146

Figure 6.15	A schematic of a cross section of the simulation setup for atmospheric absorption, not drawn to scale. The height of the atmosphere is 200 km, and each 1 km layer (not drawn) has a composition determined from the NRLMSISE-00 atmosphere model. The blue spherical detectors span the entire plane at COSI's altitude during the observation. The red source is positioned such that a direct beam from the source intersects the center sphere at (Θ, Φ) , or the GRB position in detector coordinates.	147
Figure 6.16	The simulated CDS-ARM distribution with atmospheric scattering included in the simulations. Incorporating the atmospheric scattering barely changes the central part of the simulated CDS-ARM distribution, indicating that the impact of the atmospheric scattering is only marginal. . .	148
Figure 6.17	A 2D histogram in which each cell contains the number of counts in the CDS-ARM distribution peak (from -2° to 2°) corresponding to the source position given by the offset in question. An offset of (0,0), corresponds to the hotspot of the image of GRB 160530A, or $l = 243.4^\circ$, $b = 0.4^\circ$ (theta refers to offset in Galactic latitude b and phi refers to offset in Galactic longitude l). If the hotspot in the image of the GRB data were shifted due to the image deconvolution, we would expect to see a result similar to the simulation result with the bright spot shifted from (0,0).	149
Figure 6.18	(a) The simulated CDS-ARM distribution fit with the ARM function. (b) The measured CDS-ARM distribution fit with the ARM function convolved with a Gaussian. The data points are drawn in blue and the fit is drawn in red. . . .	150
Figure 6.19	The COSI spectrum of GRB 160530A, with the background subtraction performed using the CDS method. The spectrum was fit with the atmospheric absorption model multiplied by the Band function, and was binned such that each bin had the same signal-to-noise ratio $S/N = 3$. See Table 6.5 for the best fit parameters.	152

Figure 7.1	The zenith angle of the Crab in detector coordinates throughout the 2016 flight, where a zenith angle of $\Theta = 0^\circ$ corresponds to an on-axis source. The dashed red line marks a zenith angle of 40° , used as a cutoff point to determine periods in which the Crab was in COSI's FOV. There are 12 days in which the Crab was briefly above $\Theta = 40^\circ$	155
Figure 7.2	The COSI image of the Crab using data from June 12 and 13. To make this image, we used the event selections detailed in Table 7.1 and performed 35 iterations of LM-MLEM algorithm. The color scale is in arbitrary units.	155
Figure 7.3	The ARM distribution of a background observation from the 2016 flight fit with a single Gaussian. The ARM distribution was calculated with respect to $(l = 210^\circ, b = -30^\circ)$, a background location in COSI's FOV, and we used a 40° pointing cut around $(l = 210^\circ, b = -30^\circ)$ to generate the event list. The Gaussian fit is a good approximation of the data, with a $\chi^2/\text{dof} = 2.5$	157
Figure 7.4	The ARM distribution from the Crab observation fit (a) with a single Gaussian and (b) with the sum of two Gaussians. The sum of two Gaussians clearly fits the data better, with a $\chi^2/\text{dof} = 1.4$ in comparison to the fit of a single Gaussian which results in a $\chi^2/\text{dof} = 15.3$	157
Figure 7.5	A schematic of a Compton event with the largest possible Compton scatter angle if the EHC and a 40° pointing cut are applied. The source at the yellow star is at a zenith angle of 40° and emits a photon that interacts in the detector at \vec{r}_1 and then at \vec{r}_2 . The red Compton circle extends to just above the horizon; if it extended any lower, the event would be rejected by the EHC. The resulting maximum Compton scatter angle, 65° , is half of the angle between the horizon and the initial photon direction, which is 130°	159
Figure 7.6	The ideal modulation of the azimuthal scatter angle as a function of Compton scatter angle for a variety of energies. The modulation is highest at low energies with Compton scatter angles $\phi \sim 90^\circ$. <i>Figure from Lowell (2017).</i>	164

Figure 7.7	A spectrum of flight data background compared with a spectrum of the simulated background, scaled such that both spectra contain the same number of events. The 511 keV line is removed since it is a large source of atmospheric background, and we remove it for all analysis done in this section. There is good agreement between the two spectra.	166
Figure 7.8	μ_{100} as a function of polarization angle for two sets of event selections, with the best fit μ_{100} shown as the blue line. (a) Compton scatter angles from 0° - 65° , distance between first two interactions ≥ 1 cm, distance between any two interactions ≥ 0 cm, and 2-7 interactions; the best fit μ_{100} is 0.195 ± 0.006 . (b) Compton scatter angles from 0° - 30° , distance between first two interactions ≥ 0.5 cm, distance between any two interactions ≥ 0.3 cm, and 2-7 interactions; the best fit μ_{100} is 0.084 ± 0.009	169
Figure 8.1	A schematic of photons emitted by the neutron star (black) reflecting off of the accretion disk (red). The photons emitted directly by the neutron star form the primary component of the spectrum and are drawn in green, while the photons that form the reflection component of the spectrum are drawn in blue.	173
Figure 8.2	MAXI and BAT light curves with the time of the Nuclear Spectroscopic Telescope Array (NuSTAR) observations marked by the dashed lines. The arrow marks the time of the Swift observation.	175
Figure 8.3	(a) The NuSTAR Focal Plane Module A (FPMA) (black) and Focal Plane Module B (FPMB) (red) light curve, including the four type I X-ray bursts. The time of the Swift observation is marked in blue. (b) A close-up of the first type I X-ray burst. The green lines indicate the data removed for spectral analysis.	176
Figure 8.4	The hardness ratio throughout the NuSTAR observation, with NuSTAR FPMA in red and NuSTAR FPMB in black. The time of the Swift observation is marked in blue. The hardness ratio is defined as the count rate from 12-25 keV / the count rate from 4.5-12 keV.	176

Figure 8.5	The 4.5-78.4 keV <i>NuSTAR</i> FPMA (black) and <i>NuSTAR</i> FPMB (red) spectra and the 0.7-10 keV <i>Swift</i> (blue) spectrum of 4U 1728-34, fit with the continuum model: disk blackbody + blackbody + power-law. The residuals between 6-8 keV indicate the presence of iron emission.	177
Figure 8.6	The 4.5-78.4 keV <i>NuSTAR</i> FPMA (black) and <i>NuSTAR</i> FPMB (red) spectra and the 0.7-10 keV <i>Swift</i> (blue) spectrum of 4U 1728-34, fit with the <i>relionx</i> relativistically blurred reflection model (model 1). (a) Shows the residuals of the reflection model, specifically the lack thereof between 6-8 keV. (b) Shows the νF_ν plot with individual model components: blackbody (dashed and dotted), cutoff power-law (dashed), and <i>relionx</i> (dotted).	179
Figure 8.7	The 4.5-78.4 keV <i>NuSTAR</i> FPMA (black) and <i>NuSTAR</i> FPMB (red) spectra and the 0.7-10 keV <i>Swift</i> (blue) spectrum of 4U 1728-34, fit with the <i>relionx</i> relativistically blurred reflection model where the column density is fixed to $2.9 \times 10^{22} \text{ cm}^{-2}$ (model 2). (a) Shows the residuals of the reflection model. (b) Shows the νF_ν plot with individual model components: blackbody (dashed and dotted), cutoff power-law (dashed), and <i>relionx</i> (dotted).	181
Figure 8.8	The 4.5-78.4 keV <i>NuSTAR</i> FPMA (black) and <i>NuSTAR</i> FPMB (red) spectra and the 0.7-10 keV <i>Swift</i> (blue) spectrum of 4U 1728-34, fit with <i>relxill</i> (model 5). (a) Shows the residuals of the reflection model. (b) Shows the νF_ν plot with individual model components: blackbody (dashed and dotted), <i>comptt</i> (dashed), and <i>relxill</i> (dotted).	182
Figure 8.9	The change in χ^2 versus the change in inner radius, showing that the model is consistent with the disk extending to the ISCO. The dashed blue line marks 90% confidence and the dashed red line marks 99.7% confidence (3σ). (a) Shows the results of the <i>relionx</i> model (model 1). (b) Shows the results of the <i>relxill</i> model (model 5).	184

LIST OF TABLES

Table 1.1	A summary of all measurements to date of the polarization fraction (PF) and polarization angle (PA) of the Crab nebula. Unless otherwise indicated, the polarization measurements come from the off-pulsed emission only. There is likely some dependence of the polarization angle on energy, indicating that the emission at different energies originates from different areas of the nebula.	9
Table 3.1	COSI calibration sources.	51
Table 3.2	Interpolated values of N used in the empirical method of simulating charge sharing. The charge sharing distribution is defined by a Gaussian with $\sigma = N\sigma_{\text{diffusion}}$. These values of N were chosen such that the measurements and simulations have the same fraction of events with adjacent strips. For each energy shown in this table, N is averaged over all 12 GeDs. The standard deviation of N over all 12 GeDs is computed to give a sense of the variation of N between detectors.	58
Table 3.3	The fraction of flagged events in the strip pairing and depth calibration for three different calibration source line energies. The simulation without charge sharing, with charge sharing, and with charge loss and crosstalk is compared to the measurement. As these effects are applied to the simulation, the fraction of flagged simulated events increases and is closer to the fraction of flagged measured events, as expected.	67
Table 3.4	A comparison of the line widths and centroids of a Gaussian fit of each calibration source line for measurements and simulations.	78

Table 3.5	A comparison of how well the measured and simulated spectra match for a variety of calibration sources and off-axis angles. We compare the measurements to simulations processed with the standard MEGALib DEE and to simulations processed with the COSI-specific DEE . We show the mean and FWHM of the Gaussian fit to the distribution of the difference in number of sigmas $z = (N_M - N_S)/\sigma_M$, where z is computed for each spectral bin. . . .	80
Table 3.6	The systematic error that must be applied to the measured flux, due to the remaining discrepancies between measurements and simulations.	82
Table 3.7	A comparison of how well the measured and simulated ARM histograms match for a variety of calibration sources and off-axis angles. We compare the measurements to simulations processed with the standard MEGALib DEE and to simulations processed with the COSI-specific DEE . We show the mean and FWHM of the Gaussian fit to the distribution of the difference in number of sigmas $z = (N_M - N_S)/\sigma_M$, where z is computed for each bin in the ARM histogram.	84
Table 5.1	Event selections for the redistribution matrix shown in Figure 5.1a and the effective area shown in Figure 5.1b . We chose these relatively open event selections to calculate the best case effective area. The distance cuts could be further opened, which would increase the effective area but decrease the overall quality of accepted events. .	106
Table 5.2	Event selections used for the spectral fitting of GRB 160530A . See Section 2.4.2 for a description of the event selections.	112
Table 5.3	The COSI and Konus best fit parameters of GRB 160530A . The Konus parameters are from Svinkin et al. (2016) . The errors on the COSI parameters are 99% confidence. The parameters marked with an asterisk were fixed during the fit. The COSI parameters are in agreement with the Konus parameters, confirming COSI's spectral abilities.	114
Table 6.1	The event selections used for the background-subtracted spectra shown in Figure 6.7 . See Section 2.4.2 for a description of the event selections.	136

Table 6.2	The event selections for the spectral fits of the simulation of GRB 160530A shown in Figure 6.9. See Section 2.4.2 for a description of the event selections.	136
Table 6.3	The iterations of the algorithm to correct for the source counts present in BR due to the tails of the CDS-ARM distribution of a point source, applied to the simulation of GRB 160530A. Only the values for the first power law index α and the flux are allowed to vary, while β and E_c are fixed to the simulated values of -3.5 and 638 keV, respectively. For each iteration, we show the best fit α and $0.1 - 3$ MeV flux values, as well as their relative deviation from the simulated parameters $(P_{\text{fit}} - P_{\text{sim}})/P_{\text{sim}} \times 100$, where P is the parameter in question. First we fit the subtracted spectrum without applying any correction. To perform the first iteration of the algorithm, we subtracted a spectrum with S_{BR} counts and $\alpha = -0.7657$ from the background spectrum and fit the new background-subtracted spectrum. The parameters converged after six iterations. The final converged parameters, $\alpha = -1.1112^{+0.3326}_{-0.3533}$ and $0.1 - 3$ MeV flux = $3.0353^{+0.4647}_{-0.8503} \text{ } \gamma/\text{cm}^2/\text{s}$, are consistent with the simulated parameters. To prove that the parameters have converged, we continued the analysis out to 20 iterations and show the fit parameters from the 10 th and 20 th iterations. The uncertainties on the fit are to 99% confidence.	138
Table 6.4	The event selections used in the background subtraction and spectral fit of the Crab-like source simulation. See Section 2.4.2 for a description of the event selections. . . .	140
Table 6.5	The COSI best fit parameters of the spectrum of GRB 160530A, using the CDS background subtraction method and the time-domain background subtraction method from Chapter 5. The Konus best fit parameters of GRB 160530A from Svinkin et al. (2016) are shown for comparison. The uncertainties on the COSI parameters are 99% confidence. The COSI parameters from the CDS background subtraction method are consistent with those from the other fits, confirming that this background subtraction method works for COSI flight data.	152

Table 7.1	The event selections used to make the image of the Crab shown in Figure 7.2. See Section 2.4.2 for a description of the event selections.	156
Table 7.2	The event selections used to generate the ARM distributions in Figures 7.3 and 7.4. See Section 2.4.2 for a description of the event selections.	158
Table 7.3	The number of counts after the initial background subtraction of the Crab data for each set of event selections tested. These results indicate that we are not estimating the flight data background with sufficient accuracy.	161
Table 7.4	The best fit constant μ_{100} value for each set of event selections, in addition to the reduced χ^2 of the fit over polarization angle. The Compton scatter angle range used has the largest effect on the resulting μ_{100}	168
Table 7.5	The number of source and background counts S and B from the simulated Crab observation, which is 24.61 ks long, and the minimum detectable polarization (MDP) for each set of event selections. MDP values above 100% indicate that we are unable to measure the polarization properties of the Crab from this observation. Analyzing the source observed at a smaller zenith angle over a longer period of time, combined with an alternate albedo radiation background rejection algorithm accommodating of Compton scatter angles of $\sim 90^\circ$, would significantly improve the MDP for this observation.	170
Table 8.1	Spectral parameters with varied emissivity using the reflionx model	178
Table 8.2	Spectral parameters varying the column density and the energy band using reflionx	180
Table 8.3	Spectral Parameters using relxill	183

ABBREVIATIONS

COSI	Compton Spectrometer and Imager
NCT	Nuclear Compton Telescope
OSO-8	8th Orbiting Solar Observatory
CGRO	Compton Gamma Ray Observatory
COMPTEL	Imaging Compton Telescope
INTEGRAL	International Gamma Ray Astrophysics Laboratory
SPI	Spectrometer on INTEGRAL
IBIS	Imager on Board the INTEGRAL Satellite
BeppoSAX	Beppo Satellite per Astronomia a raggi X
CZTI	Cadmium Zinc Telluride Imager (on <i>AstroSat</i>)
SGD	Soft Gamma-ray Detector (on <i>Hitomi</i>)
PoGO+	Polarized Gamma-ray Observer
Chandra	Chandra X-ray Observatory
RXTE	Rossi X-ray Timing Explorer
XMM-Newton	European Space Agency's X-ray Multi-Mirror Mission
XRT	X-ray Telescope (on <i>Swift</i>)
BAT	Burst Alert Telescope (on <i>Swift</i>)
MAXI	Monitor of All-sky X-ray Image
NuSTAR	Nuclear Spectroscopic Telescope Array
FPMA	NuSTAR Focal Plane Module A

FPMB	NuSTAR Focal Plane Module B
HST	Hubble Space Telescope
OPTIMA	Optical Pulsar Timing Analyser
AMEGO	All-sky Medium Energy Gamma-ray Observatory
ESA	European Space Agency
NASA	National Aeronautics and Space Administration
SSL	Space Sciences Laboratory
MEGAlib	Medium Energy Gamma-ray Astronomy library
DEE	detector effects engine
LM-MLEM	list mode maximum-likelihood expectation-maximization
ARM	angular resolution measure
CTD	collection time difference
EHC	Earth horizon cut
GSE	ground support equipment
CDS	COMPTEL Data Space
ASAD	azimuthal scatter angle distribution
MDP	minimum detectable polarization
AGN	active galactic nucleus
PWN	pulsar wind nebula
HMXB	high-mass X-ray binary
LMXB	low-mass X-ray binary
GRB	gamma-ray burst
REP	relativistic electron precipitation
FOV	field of view

PMT	photomultiplier tube
ADC	analog-to-digital converter
DAC	digital-to-analog converter
GeD	germanium detector
PDU	power distribution unit
FPGA	field-programmable gate array
IDPU	instrument and detector processing unit
PSU	power supply unit
CsI	cesium iodide
SPB	super pressure balloon
ZPB	zero pressure balloon
ULDB	ultra-long duration balloon
CSBF	Columbia Scientific Ballooning Facility
SIP	Support Instrument Package
LOS	line of sight
OP	openport
dGPS	differential Global Positioning System
UTC	Coordinated Universal Time
FWHM	full width at half maximum

PART I

BACKGROUND MATERIAL

POLARIZATION OF COMPACT OBJECTS

Neutron stars and black holes are formed when a massive star explodes as a supernova. These compact objects, so-called because of their extreme densities and relatively small sizes, are some of the most fascinating bodies in our universe. Studying compact objects gives us a window into how matter behaves in ultra-dense environments with strong gravitational and magnetic fields, which we are unable to create in the laboratory.

Astrophysical γ -rays are an excellent probe of non-thermal emission from neutron stars and black holes. In particular, γ -ray polarimetry offers a unique look into these extreme and exotic objects and is a powerful technique for determining source geometries and emission mechanisms, information not always attainable with current spectral, timing, or imaging analysis. In this chapter, I will discuss the potential of studying neutron stars and black holes with γ -ray polarimetry, as well as the current state of polarimetric studies.

1.1 POLARIZATION IN ASTROPHYSICS

I will first briefly introduce the role of polarization studies within astrophysics. When performing polarimetry, the goal is to measure two parameters: the *polarization fraction*, or degree of linear polarization, and the *polarization angle*, or the direction of the polarization vector on the sky. Measuring the polarization fraction can determine the emission mechanism of the γ -rays, since certain radiative processes emit highly polarized light whereas others do not. Measuring the polarization angle can determine certain aspects of the source and magnetic field geometry. Note that the polarization angle is generally defined relative to celestial North in the Easterly direction (e.g. [Chauvin et al. \(2018\)](#)).

[Lei et al. \(1997\)](#) provides a comprehensive overview of emission mechanisms that emit polarized γ -rays, briefly summarized here. Generally, polarized γ -rays are emitted by non-thermal processes. Synchrotron radiation

can lead to highly polarized light with a maximum polarization fraction of 65-85%, depending on the spectral parameters of the source. This maximum assumes a homogeneous magnetic field; certain structure in the magnetic field will reduce the degree of linear polarization. Compton scattering and inverse Compton scattering, meanwhile, can both polarize and depolarize a photon beam. The degree of linear polarization after Compton scattering depends on the polarization fraction of the beam before it scatters, the spatial distribution of the photons, and the initial energy and Compton scatter angle of the photons. Generally, synchrotron radiation from astrophysical objects is expected to be more highly polarized than that from Compton or inverse Compton scattering (Lei et al., 1997).

At γ -ray energies, we measure polarization through Compton scattering. Compton scattering is the dominant interaction process for photons at these energies, and photons are more likely to Compton scatter perpendicular to their polarization direction. If a beam is partially polarized, a modulation in the angle between the direction of the scattered photon and the initial electric field vector of the photon is apparent in the detector. By measuring the amount and phase of this modulation, we can determine the polarization fraction and angle, respectively. See Lei et al. (1997), Lowell (2017), and Chapter 7 of this work for more details about Compton polarimetry.

1.2 PULSARS AND PULSAR WIND NEBULAE

In 1968, Hewish et al. (1968) detected objects emitting rapid periodic pulsations at radio wavelengths. Gold (1968) proposed that these so-called pulsars could be rotating neutron stars, incredibly dense objects held together by neutron degeneracy pressure and previously proposed by Baade & Zwicky (1934). More detections of pulsars, an increased theoretical understanding of neutron stars, and a lack of alternative models have led the community to accept this explanation (Shapiro & Teukolsky, 1983).

Pulsars are found in both supernova remnants and within binary systems accreting mass from a companion star. Isolated pulsars within supernova remnants are also referred to as rotation-powered pulsars, so-called because the radiated energy comes from the rotation of the pulsar. The pulse period of rotation-powered pulsars gradually slows down over time as rotational kinetic energy is lost. Pulsars within binary systems are also referred to as accretion-powered pulsars, since the radiated energy comes from the accretion of the companion star. The pulse period of accretion-powered pulsars

can increase or decrease over time, likely depending on the details of the angular momentum transfer between the neutron star and the accretion disk (Psaltis, 2006). In this section, I will focus on rotation-powered pulsars and discuss their observed properties and surrounding nebular emission.

1.2.1 PULSED EMISSION

In an early model of pulsar electrodynamics, Goldreich & Julian (1969) noted some key points: pulsars cannot exist in a vacuum since a large surface charge will build up on the surface of the star, and as a consequence, the dipole magnetic field will corotate with the star. Since particles cannot travel faster than the speed of light, however, at large distances from the neutron star the dipole magnetic field can no longer corotate. The exact distance and mechanism of the cease in corotation is not currently well understood (Kaspi et al., 2006), but corotation must cease by the light cylinder radius $r_{\text{LC}} = c/\Omega$ where Ω is the spin frequency of the pulsar. Figure 1.1 shows a schematic of the Goldreich & Julian (1969) model, in the case that the pulsar’s magnetic field is lined up with the rotation axis. At $r < r_{\text{LC}}$, the magnetosphere corotates with the pulsar and thus the magnetic field lines are closed; at $r > r_{\text{LC}}$, the magnetosphere no longer corotates and the magnetic field lines are open.

The magnetic field is not necessarily aligned with the rotation axis; in fact, pulsed emission occurs when the magnetic field is misaligned from the rotation axis since the emitted radiation originates at a location within the corotating magnetosphere (Kaspi et al., 2006). Figure 1.2 shows a schematic of a pulsar in which the magnetic field and rotation axis are misaligned, along with three proposed particle acceleration sites; this particle acceleration leads to the observed X-ray and γ -ray emission. In the *polar cap* model, the γ -rays are produced in the region directly above the polar cap. This model has most likely been ruled out, however, because it predicts a narrow pulse profile disfavored by many observed wide γ -ray pulse profiles with multiple components (Ravi et al., 2010). In the *outer gap* model, the γ -rays are produced in between the light cylinder and the null charge surface. In the *slot gap* model, the γ -rays are produced along the entire length of the last open magnetic field line. See Dyks et al. (2004), Harding (2004) and references therein for a more detailed description of these models.

Polarimetry of the pulsed emission in the radio and optical wavebands has significantly furthered our understanding of pulsars (Kaspi et al., 2006). Polarization measurements of the pulsed γ -rays could help determine the emission




Figure 1.1: A schematic of the pulsar electrodynamics model from [Goldreich & Julian \(1969\)](#). The magnetic field is aligned with the pulsar spin axis. The pulsar is in the bottom left-hand corner. The closed magnetic field lines corotate with the pulsar within the light cylinder. The magnetic field lines that extend past the light cylinder are open, and particles are accelerated along them into the nebula. *Figure from Goldreich & Julian (1969).*

mechanism and geometry within the magnetosphere, invaluable information for attaining a correct understanding of the pulsed high-energy radiation.

1.2.2 PULSAR WIND NEBULAE

The pulsed emission described above accounts for less than 10% of the rotational energy of a pulsar; the majority of the energy is emitted as a magnetized wind that powers the pulsar wind nebula (PWN) ([Kaspi et al., 2006](#)). The charged particles that form this magnetized wind flow outward and are accelerated to very high energies. The resulting nebula emits synchrotron radiation from radio to γ -rays ([Gaensler & Slane, 2006](#)).

The structure of PWNe varies significantly across different wavelengths, as shown in Figure 1.3. With [Chandra](#)'s imaging capabilities, significant variations in the structure across PWNe have been detected in the soft X-ray band ([Kargaltsev & Pavlov, 2008](#)). The angular resolution of current γ -ray telescopes is much too poor to determine from where in the nebula the γ -rays originate. γ -ray polarimetry, however, provides a lens through which we can study the γ -ray nebular emission.




Figure 1.2: A schematic of a pulsar whose rotation axis (marked by the spin frequency Ω) is not aligned with its magnetic field \mathbf{B} . Three models for emission sites of the high energy radiation are marked. *Figure from Harding (2004).*




Figure 1.3: The Crab nebula pictured (a) in optical light by the Hubble Space Telescope (HST) and (b) in soft X-rays by Chandra. The nebular structure clearly varies based on wavelength.

1.2.3 γ -RAY POLARIZATION OF THE CRAB PULSAR AND NEBULA

The Crab pulsar and nebula is one of the brightest sources in the high energy sky. The Crab is within the remnant of a supernova that occurred in 1054 and was bright enough to be seen clearly with the naked eye. Pulsations from the Crab pulsar were detected in 1969 with a period of 33 ms (Comella et al., 1969). Because of their brightness, the Crab pulsar and nebula comprise an extremely well-studied source that is an excellent candidate for γ -ray polarimetric measurements.

Novick et al. (1972) performed the first polarization measurement of high energy emission from the Crab using a lithium polarimeter (sensitive to X-rays between 5 and 20 keV) and a Bragg crystal polarimeter (sensitive to X-rays at 2.6 keV) launched into space on a sounding rocket. The resulting detection of polarized emission confirmed that the X-rays are emitted via synchrotron radiation (Novick et al., 1972). Weisskopf et al. (1976) confirmed this result using Bragg crystal polarimeters aboard the 8th Orbiting Solar Observatory (OSO-8), measuring a polarization fraction of 16% at 2.6 keV and 18% at 5.2 keV. Shortly thereafter, Weisskopf et al. (1978) measured the X-ray polarization of the nebula – without the contribution of pulsed emission – using the OSO-8 data: the result was a polarization fraction of $\sim 19\%$ at both energies.

It was not until decades later that polarization of the γ -ray radiation from the Crab was measured. Though not designed as polarimeters, both the Spectrometer on INTEGRAL (SPI) and the IBIS can perform polarimetry of bright sources. Dean et al. (2008) measured the polarization of the unpulsed (i.e. nebular) emission from the Crab using SPI and reported a polarization fraction of $46 \pm 10\%$ and a polarization angle of $123^\circ \pm 11^\circ$. Interestingly, the polarization angle lines up almost perfectly with the rotation axis at $124.0^\circ \pm 0.1^\circ$ (Ng & Romani, 2004) (see Figure 1.4); this is also the case for the optical polarization angle (123° ; Kanbach et al. (2005)), but not for X-ray polarization angle (156.4° at 2.6 keV and 152.6° at 5.2 keV; Weisskopf et al. (1978)). Forot et al. (2008) performed phase-resolved polarimetry of the Crab pulsar using IBIS. Their reported polarization angle of the nebular emission, $120.6^\circ \pm 8.5^\circ$, is consistent with that of Dean et al. (2008), yet their reported polarization fraction of greater than 72% is only marginally consistent with that from Dean et al. (2008). Forot et al. (2008) also detected a high degree of polarization integrating over the entire pulse profile: $42^{+30}_{-13}\%$ and a polarization angle of 70° .




Figure 1.4: The γ -ray polarization angle superimposed onto the [Chandra](#) image (blue) and the [HST](#) image (red) of the Crab nebula, drawn such that the vector passes through the pulsar position. The polarization angle is very well aligned with the inner jet structure present in the [Chandra](#) image. *Figure from [Dean et al. \(2008\)](#).*

By comparing the γ -ray polarization angle to optical polarization angles throughout the nebula, [Forot et al. \(2008\)](#) were able to determine that the unpulsed γ -rays originate from the center of the nebula, close to the pulsar (the angular resolution of γ -ray instruments makes this currently impossible to determine with imaging analysis). The γ -ray polarization angle is closest to the optical polarization angle at the very center of the nebula, e.g. 118.9° within $1.15''$ of the pulsar ([Slowikowska et al., 2008](#)).

In the last few years, many γ -ray instruments that were designed explicitly with polarimetry in mind were launched and have since successfully measured the polarization of the Crab. The polarization fraction and angle of all reported X-ray and γ -ray off-pulse measurements are summarized in [Table 1.1](#) for ease of comparison. The differences in polarization angle could be due to energy-dependent polarization properties, likely indicating that photons of different energies are emitted in different locations of the nebula where the magnetic field is more or less tangled.

The Cadmium Zinc Telluride Imager ([CZTI](#)) on *AstroSat* team performed a detailed phase-resolved polarimetry study of the Crab emission. They discovered that the polarization angle swings rapidly during the peaks in the pulse profile ([Vadawale et al., 2018](#)). Additionally, they detected significant variation in polarization angle within the off-pulsed emission, which indicates

Instrument	PF (%)	PA (°)	Energy	Reference
OSO-8	19.19±0.97	156.36±1.44	2.6 keV	Weisskopf et al. (1978)
	19.50±2.77	152.59±4.04	5.2 keV	
INTEGRAL/SPI	46±10	123±11	0.1-1 MeV	Dean et al. (2008)
INTEGRAL/IBIS	>72	120.6±8.5	200-800 keV	Forot et al. (2008)
PoGO+	17.4 ^{+8.6} _{-9.3}	137±15	20-160 keV	Chauvin et al. (2017)
AstroSat/CZTI	39±10	140.9±3.7	100-380 keV	Vadawale et al. (2018)
Hitomi/SGD ^a	22.1±10.6	110.7 ^{+13.2} _{-13.0}	60-160 keV	Aharonian et al. (2018)

Table 1.1: A summary of all measurements to date of the polarization fraction (PF) and polarization angle (PA) of the Crab nebula. Unless otherwise indicated, the polarization measurements come from the off-pulsed emission only. There is likely some dependence of the polarization angle on energy, indicating that the emission at different energies originates from different areas of the nebula.

^a These results are phase-integrated, meaning that they include the nebular and pulsed emission. Phase-resolved polarimetry was not reported.

that some of the off-pulsed emission may be due to the pulsar itself rather than the nebula. The polarization fraction also varies with phase. Chauvin et al. (2018) performed a similar analysis with the Polarized Gamma-ray Observer (PoGO+) dataset and also found that the polarization fraction and angle vary with phase. These results call into question some of proposed theoretical models of pulsar emission (Vadawale et al., 2018).

1.2.4 FUTURE POLARIZATION MEASUREMENTS OF PULSARS

Future measurements of polarized emission from the Crab pulsar and nebula are potentially interesting for a few reasons. Since γ -ray polarimetry is a relatively young field, how the γ -ray polarization properties of the Crab vary with time is currently unknown. Future measurements would clarify the consistency, or lack thereof, of the polarized γ -ray emission. Additionally, more measurements by instruments sensitive to different energy ranges could illuminate the dependence of polarization fraction and angle on energy, and help determine the magnetic field properties in various locations of the nebula.

Since the Crab is the brightest isolated pulsar system at these energies, it is not surprising that it is the first pulsar whose γ -ray polarization properties were measured. Performing polarimetry of other pulsars and PWNe and

comparing these results to those of the Crab could determine the universality of the Crab results, important for furthering our understanding of pulsar emission and nebular structure.

1.3 GALACTIC BLACK HOLES

A black hole is a region of spacetime that cannot communicate with the outside world, as light within the event horizon of a black hole cannot escape the exceedingly strong gravitational field. Black holes are thought to be created when a massive star explodes as a supernova and the resulting compact object is above the neutron star mass limit (the neutron star mass limit is the mass above which the neutron degeneracy pressure no longer balances the gravitational self-attraction of the object). A black hole can be completely characterized by three parameters: its mass M , angular momentum J , and charge Q (Shapiro & Teukolsky, 1983). Often the angular momentum per unit mass, referred to as the spin $a = J/M$, is used when discussing black holes.

Isolated stellar mass black holes are difficult to detect because their emitted radiation is faint (Sartore & Treves, 2012). Thus, the majority of known black holes within our galaxy are part of binary systems, in which the black hole orbits a companion star and accretes matter from the companion. This accretion process emits radiation that can be bright enough to detect. Currently, there are tens of black hole binary candidates within our galaxy. Figure 1.5 shows a schematic of 21 black hole binary candidates identified as of 2014.

1.3.1 IDENTIFYING BLACK HOLE BINARIES

A black hole binary is only one type of X-ray binary, or binary systems with one compact object orbiting a main sequence star. The majority of known X-ray binaries contain a neutron star as the compact object. To observationally determine the nature of the compact object, we can compute the mass function:

$$f(M_1, M_2, i) = \frac{(M_1 \sin i)^3}{(M_1 + M_2)^2} = \frac{P_{\text{orb}} K_2^3}{2\pi G} \quad (1.1)$$

where M_1 is the mass of the compact object, M_2 is the mass of the companion star, i is the inclination of the orbital plane, P_{orb} is the orbital period of the system, and K_2 is the projection of the orbital velocity of the companion

1.3 GALACTIC BLACK HOLES

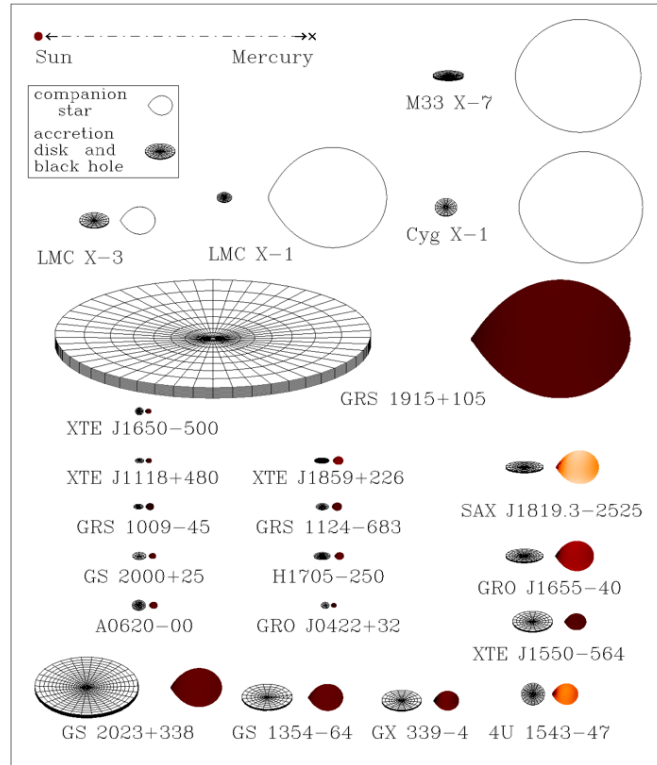


Figure 1.5: A schematic of 21 stellar mass black hole binary candidates within our galaxy drawn to scale relative to the distance between the Sun and Mercury. *Figure from Narayan & McClintock (2015).*

star along the line of sight (see e.g. [McClintock & Remillard 2006](#); [Shapiro & Teukolsky 1983](#)). Since P_{orb} and K_2 are observable quantities, the value for the mass function can be computed. A minimum value for M_2 can be estimated if the companion star type is identified. From the mass function and a minimum value of M_2 , we can determine a minimum value for M_1 , the mass of the compact object. If the minimum possible M_1 is larger than the mass limit of a neutron star ($\sim 3M_{\odot}$, e.g. [Rhoades & Ruffini 1974](#); [Kalogera & Baym 1996](#)), the compact object must be a black hole. For ten of the systems in Figure 1.5, the mass function has been measured and yields a minimum M_1 larger than $3M_{\odot}$, indicating that these systems definitively hold a black hole ([McClintock & Remillard, 2006](#)).

In cases where the companion stars are so faint that the mass function cannot be measured, black hole binary candidates can be identified based on other observational properties. The spectrum must include both an ultrasoft X-ray spectrum between 1–10 keV and a non-thermal component that extends beyond 20 keV. Black hole binaries also exhibit a variety of spectral states (see Section 1.3.3 for more details) and rapid temporal variability ([McClintock & Remillard, 2006](#)). Note that these candidates are likely black holes, but their nature cannot be determined definitively without measuring the mass function. Nevertheless, if an X-ray binary system exhibits periodic pulsations or thermonuclear X-ray bursts, then the compact object can be definitively ruled a neutron star.

1.3.2 ACCRETION ONTO BLACK HOLES

A black hole can accrete matter from the companion either through the inner Lagrange point of the orbit (referred to as Roche lobe overflow) or by capturing mass ejected from the companion star as a stellar wind. Generally, binary systems that accrete via Roche Lobe overflow have a low mass companion star; thus, these systems are referred to as low-mass X-ray binaries (LMXBs) in contrast to high-mass X-ray binary (HMXB) which hold a relatively high mass companion star and accrete via a stellar wind ([Psaltis, 2006](#)).

The Roche lobe is defined as the first common equipotential ([Shapiro & Teukolsky, 1983](#)). In the case of Roche lobe overflow, the companion star grows to fill the Roche lobe and mass transfers from the gravitational potential well of the companion to the gravitational potential well of the compact object through the inner Lagrange point. The transferred material is not accreted directly onto the black hole due to nonzero angular momentum, but

rather forms an accretion disk around the black hole. Many models of accretion flow have been proposed, and are becoming more sophisticated with time; these include thin disks (Shakura & Sunyaev 1973; Novikov & Thorne 1973; Pringle & Rees 1972), slim disks (Abramowicz et al. 1988; Beloborodov 1998), and advection-dominated accretion flow (Narayan & Yi 1994, 1995; Narayan et al. 1996; Quataert & Narayan 1999). See e.g. Abramowicz & Fragile (2013) for a more in-depth review.

In the case of stellar wind accretion, the companion star ejects a stellar wind in all directions, only $\leq 0.1\%$ of which is accreted onto the black hole (Shapiro & Teukolsky, 1983). The accreted material may or may not have enough angular momentum to form an accretion disk; if no accretion disk is formed, the accretion is close to spherical (e.g. Bondi & Hoyle 1944). Observed rapid variability in the light curves of HMXBs have shown the need for more complex models of stellar wind accretion that include highly structured winds and accretion instabilities (Negueruela, 2010).

1.3.3 SPECTRAL STATES

Black hole binaries exhibit a variety of spectral states in which the spectral shape and relative intensities differ. Here I will give a brief overview of the observed parameters of each spectral state and the physical processes thought to be occurring in each spectral state (see e.g. McClintock & Remillard (2006) for more details). Figure 1.6 shows a schematic of the accretion disk configuration and amount of mass transfer onto the black hole.

Hard state—the spectrum of a black hole binary in the hard state is well fit with a power law with a high energy cutoff at around 100 keV in addition to a faint multi-temperature blackbody model. To account for the non-thermal emission, there is likely a corona or a jet present. Spectral features from photons reflecting off of the accretion disk, such as the iron emission line at $\sim 6-7$ keV, are often observed. The multi-temperature blackbody model is thought to be emission from the accretion disk, and its relative faintness is indicative that the disk is truncated from the black hole.

Soft state—the spectrum of a black hole binary in the soft state is dominated by thermal emission and is well fit with a multi-temperature blackbody model and a faint steep power law. The most likely physical picture is that the emission is primarily from the accretion disk, the inner radius of which is close to the black hole.

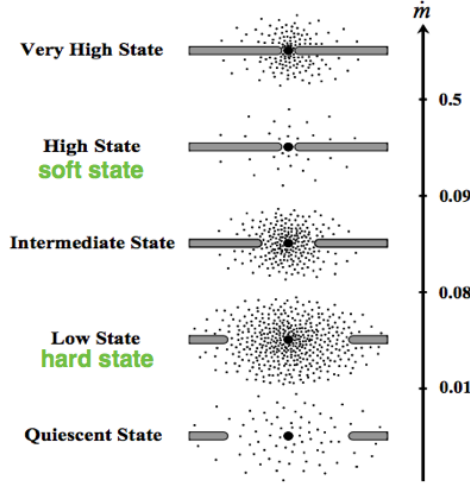


Figure 1.6: A schematic of the spectral states of black hole binaries. Since the figure was made, the high state and low state have been renamed the soft state and hard state, respectively. The dots represent the accreting material and the horizontal bars represent the accretion disk. The amount of mass transfer is indicated by the \dot{m} scale. *Figure from Esin et al. (1997)*

Quiescent state—black hole binaries in the quiescent state are very faint with spectra well fit by a hard power law (i.e. the emission is non-thermal).

Very high state—black hole binaries in the very high state, also referred to as the steep power law state, are highly luminous and exhibit a spectrum well fit with a steep power law and no high energy cutoff. The physical picture remains unclear.

Intermediate state—the intermediate state describes the state of black hole binaries during transitions between the other states.

1.3.4 THE ORIGIN OF THE γ -RAYS

The origin of the non-thermal high energy emission observed in the hard state of black hole binaries is currently unclear. A source of optically thin hot plasma is required to produce non-thermal emission, which has been proposed to originate from a corona (Dove et al. (1997) and references therein). The general idea is that emission from the accretion disk is Comptonized by the corona, resulting in the observed non-thermal high energy emission. This Comptonized radiation can in turn interact with the accretion disk, resulting in the observed reflection features (Dove et al., 1997).

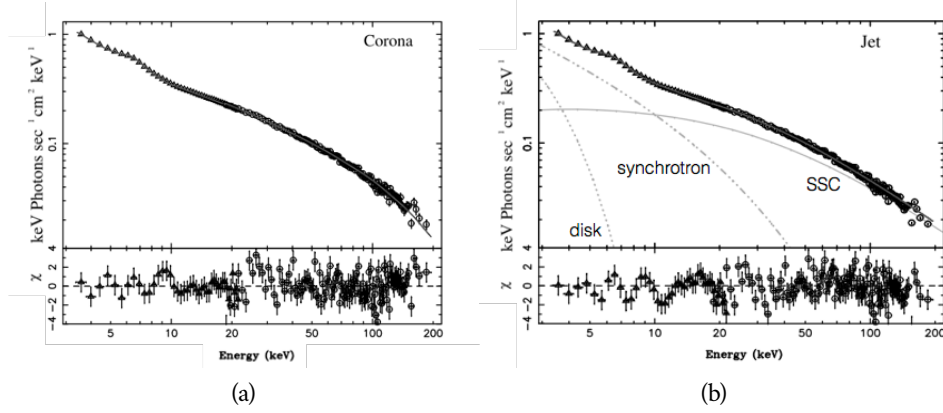


Figure 1.7: Spectral fits of Cyg X-1 data using a model in which (a) the corona and (b) the jet is the emitter of the observed high energy radiation. The two models fit the data equally well, indicating that spectral analysis is not able to distinguish between the two models. *Figure from Markoff et al. (2005).*

Markoff et al. (2005) proposed that jets could replace the role of the corona as the emitter of the high energy radiation. Jets, or bipolar plasma outflows, have been observed in radio observations of black hole binaries. Comparing radio and X-ray observations of black hole binaries has shown a correlation between the radio and X-ray luminosity in the hard state, $L_{\text{radio}} \sim L_X^{0.7}$; this relationship remains consistent over long periods of time, even if the system leaves and re-enters the hard state (Gallo et al., 2003; Corbel et al., 2013). This correlation indicates that the X-ray emission is related to the jets, and thus the jets must be considered in models of black hole accretion. To explore whether jets could replace the corona as the emitter of high energy radiation, Markoff et al. (2005) modeled both scenarios and fit observed spectra of black hole binaries with both models. Both models fit the data equally well (see Figure 1.7 for an example), leaving the source of the high energy emission unclear.

Measuring the polarization of the high energy emission from black hole binaries has the potential to resolve this mystery. In the case that the high energy emission originates from a corona, photons from the accretion disk are inverse Compton scattered by electrons in the corona (Dove et al., 1997; Markoff et al., 2005). This process results in an unpolarized or moderately polarized beam depending on the alignment of the incident vectors of the photons (Lei et al., 1997). In the alternate case that the high energy emission originates from the jet, the jet radiates via synchrotron and synchrotron self Compton emission (see Markoff et al. (2005) for details of the jet model

used in Figure 1.7). Synchrotron emission is expected to be highly polarized (Lei et al., 1997). Thus, though the two models have very similar spectral signatures, their polarization signatures differ significantly. By measuring the polarization of high energy emission from black hole binaries, we can determine the origin of the high energy emission and further our understanding of accretion physics and how matter behaves in ultra-strong gravitational fields.

1.3.5 γ -RAY POLARIZATION OF CYG X-1

The only γ -ray polarization measurements of a black hole binary to date were performed by the SPI and IBIS telescopes. Both SPI and IBIS observed Cygnus X-1, a bright, well-known HMXB with an O-type companion star, when it was primarily in the hard state. The Cyg X-1 spectrum measured by IBIS and SPI has two clear separate components, shown (from IBIS) in Figure 1.8, which led the teams to separately measure the polarization of the low energy and high energy emission. The IBIS team measured weakly polarized emission between 250 and 400 keV with a polarization fraction upper limit of 20%, and strongly polarized emission – $67 \pm 30\%$ – between 400 and 2000 keV (Laurent et al., 2011). Figure 1.8 shows the azimuthal angle modulation curves for each region of the spectrum along with the spectrum fit with two distinct components. The SPI team performed polarimetry in three energy bands, and measured the following polarization fractions: upper limit of 20% between 130 and 230 keV, $41 \pm 9\%$ between 230 and 370 keV, and greater than 75% between 370 and 850 keV (Jourdain et al., 2012).

These results indicate that there are two emission mechanisms that produce γ -rays. The higher energy strongly polarized component to the emission is likely due to synchrotron radiation from the jet, whereas the weakly polarized component could be consistent with the Comptonization of accretion disk photons by hot plasma in a corona (Laurent et al., 2011; Jourdain et al., 2012).

Rodriguez et al. (2015) measured the polarization of the soft state emission of Cyg X-1 between 400 and 2000 keV with IBIS. The results are somewhat inconclusive: the modulation curves are consistent with an unpolarized source, but the upper limit to the polarization fraction is $\sim 70\%$ (Rodriguez et al., 2015). If the strongly polarized emission present in the hard state is due to the jet, we perhaps should not expect to see polarized emission in the soft state, when the jet is thought to be quenched (Stirling et al. 2001 and references therein).

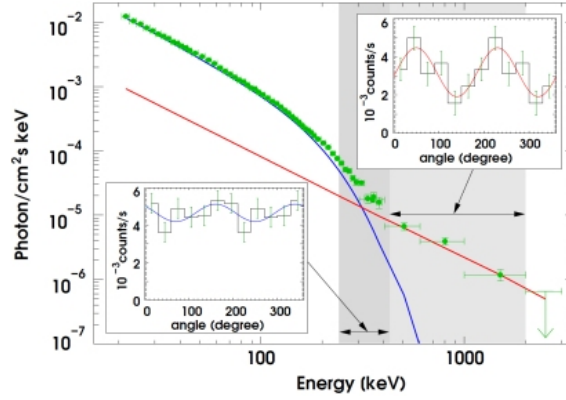


Figure 1.8: The spectrum of Cyg X-1 as measured by *IBIS* along with the two components required to fit this spectrum. The modulation curves corresponding to the polarization measurements are shown for two energy ranges (more modulation means a higher polarization fraction). The energy dependence of the polarization signatures indicates two distinct emission mechanisms that produce γ -rays. *Figure summarizing Figures from Laurent et al. (2011), from ESA press release ^a.*

^a <http://sci.esa.int/integral/48589-cygnus-x-1-s-energy-spectrum-and-polarized-signal>

1.3.6 FUTURE POLARIZATION MEASUREMENTS OF GALACTIC BLACK HOLES

Future polarization measurements of Cyg X-1 could help determine the consistency of the polarization with time. If the polarization properties of Cyg X-1 are changing on short timescales, *SPI* and *IBIS* are insensitive to it; the datasets used in the analysis described in the previous section span about six years (Jourdain et al., 2012).

Measuring the polarization of other black hole binaries will help us determine whether the Cyg X-1 results are widespread across this source class or unique to Cyg X-1. If the polarization properties of other black hole binaries differ from those of Cyg X-1, the features and magnitude of such differences is of interest. Furthermore, measuring the polarization of black hole binaries in other spectral states has the potential to determine their respective emission mechanisms and thus the underlying physical processes.

1.3.7 ACTIVE GALACTIC NUCLEI

While I have focused this section on Galactic black holes, it is important to note that performing γ -ray polarimetry on active galactic nuclei (AGNs) is also of interest. An AGN is a supermassive black hole ($M \sim 10^8 M_\odot$) at the center of a galaxy with a high luminosity that results from the black hole accreting matter (Frank et al., 2002). The origin of the γ -ray emission of AGN remains unclear, and could be inverse Compton scattering in a corona (Beckmann et al., 2011) or synchrotron self-Compton in the observed jets (Abdo et al., 2010). Similarly to Galactic black holes, measuring the polarization can constrain the emission mechanism and thus further our understanding of AGN. To date, there have been no γ -ray polarization measurements of AGN.

1.4 SUMMARY

Studying the polarization of the γ -ray emission from compact objects is a compelling way to deepen our understanding of emission mechanisms, magnetic field structure, and source geometries. Since γ -ray polarimetry is a relatively new analysis technique, so far it has only been performed on the brightest sources: the Crab and Cyg X-1. These results have hinted at the power of γ -ray polarimetry in both confirming and questioning current theoretical models of pulsars and black hole binaries. Developing new telescope technologies with improved polarization sensitivity is an important step towards measuring the polarization of the γ -ray emission from more and fainter compact objects, enabling an enhanced understanding of these fascinating phenomena.

1.5 THESIS OUTLINE

In this chapter, I have given an overview of compact objects and the current state of studying these objects with γ -ray polarimetry. In this section I will briefly outline the remainder of this thesis.

In Chapter 2, I introduce and describe the Compton Spectrometer and Imager (COSI), a balloon-borne γ -ray telescope and polarimeter designed and built at the Space Sciences Laboratory (SSL) at UC Berkeley. I will discuss the operating principle and the components of the instrument. I will also describe the instrument calibrations and the data analysis pipeline.

A large aspect of building any new instrument is ensuring that we understand the data that we collect. This understanding involves careful calibrations and accurate simulations with which we compute the instrument response. Additionally, new analysis methods and pipelines must be developed and implemented. The majority of this thesis describes the copious work done to develop the accurate simulations and new analysis methods required to eventually perform γ -ray polarimetry of compact objects detected by COSI. This work is also the foundation for performing other types of analysis (e.g. spectral) on broadband, persistent sources.

In Chapter 3, I will describe the simulation pipeline developed for COSI, which we use to calculate the instrument response and to better understand the instrument performance. I will focus on the detector effects engine (DEE) which applies the intrinsic detector performance to Monte Carlo simulations. Careful modeling of the detector effects and readout electronics has led to vast improvements in the agreement between simulations and calibration data. I will describe the comparisons we made between calibration data and simulations to assess the accuracy of the simulations.

In Chapter 4, I will provide a brief overview of COSI's 2016 balloon campaign and flight. I will focus on the operations software and strategy, the thermal environment, and the failure of three detectors because of high voltage issues in addition to the work done after the flight to diagnose these issues. I will also discuss the sources detected during the flight and briefly review the status of the analysis progress on each source.

Part III describes the astrophysics analysis work done. In Chapter 5, I describe the spectral analysis pipeline developed for COSI. By comparing the best fit spectral parameters of a source measured by COSI to those of the same source measured by another instrument, we can confirm that we have a good understanding of our instrument. I will describe the components of the spectral analysis pipeline, including the response matrix and the atmospheric absorption model, and will present the spectral fit of GRB 160530A, detected by COSI during the 2016 flight. I will also discuss the challenges of subtracting the background from persistent sources with Compton telescopes such as COSI; because of these challenges, traditional background subtraction approaches using source and background regions defined by locations on the sky are difficult to implement properly.

I have developed an alternate background subtraction method based on a three-dimensional data space pioneered by the COMPTEL collaboration, thus referred to as the COMPTEL Data Space (CDS). The CDS is defined by the pho-

ton scatter direction in detector coordinates and the initial Compton scatter angle. In Chapter 6, I will describe the CDS and the background subtraction algorithm, and discuss the tests I performed to validate this algorithm on both simulations and flight data. Note that Kierans (2018) developed a similar background subtraction algorithm based on the CDS for line sources detected by COSI, whereas the algorithm presented in Chapter 6 in this work is designed for broadband sources.

In Chapter 7, I will discuss COSI’s observation of the Crab nebula. As one of the brightest sources in the soft γ -ray sky, the Crab makes an excellent test case for new analysis methods, such as the CDS background subtraction algorithm and γ -ray polarimetry. Even though the source is bright, however, COSI’s observation of the Crab is still background-dominated because of the profuse atmospheric background. I will discuss our detection of the Crab and our results of the CDS background subtraction algorithm and spectral fitting. Unfortunately, we find that our treatment of the background is limited when it comes to background-dominated sources, and are thus unable to constrain the parameters of the spectral fit. I will then describe how we perform Compton polarimetry with COSI, and use simulations to determine whether we would be able to measure the polarization of the Crab with this observation if we had an acceptable background subtraction.

I change course in Chapter 8 to describe a study of the reflection features in the spectrum of the neutron star LMXB 4U 1728–34 detected by the X-ray telescope NuSTAR. Reflection features arise when high energy radiation emitted by the accreting neutron star interacts with the accretion disk by way of scattering or fluorescence, and measuring these features is a current topic of interest within X-ray astrophysics. NuSTAR’s energy range of 3–80 keV makes it an ideal instrument for measuring key reflection features including the iron emission line around 6–7 keV and the Compton hump at higher energies. By measuring the shape of this iron line, we can determine interesting parameters about the accretion disk, such as the disk’s inner radius. The inner radius of the disk is an upper limit of the radius of the neutron star, a parameter of interest in determining the equation of state for neutron stars. In the NuSTAR observation of 4U 1728–34 discussed in Chapter 8, we find an upper limit for the neutron star radius of 23 km.

In Chapter 9, I conclude with a summary of the work presented here and discuss next steps into the field of γ -ray polarimetry.

PART II

THE COMPTON SPECTROMETER AND IMAGER:
INSTRUMENT OVERVIEW, CALIBRATIONS AND
SIMULATIONS, AND THE 2016 FLIGHT

2

THE COMPTON SPECTROMETER AND IMAGER

The Compton Spectrometer and Imager ([COSI](#)) is a balloon-borne γ -ray Compton telescope designed to study astrophysical sources. As a Compton telescope, [COSI](#) is capable of imaging, spectroscopy, and polarimetry in the scientifically rich MeV range. The primary science targets include compact γ -ray sources such as pulsars and black holes, gamma-ray bursts ([GRBs](#)), Galactic nucleosynthesis, and positron annihilation. In this chapter, I will discuss the operating principle of Compton telescopes and describe the [COSI](#) instrument.

2.1 COMPTON TELESCOPES

The cross section for Compton scattering dominates over that of photoabsorption and pair production in the 0.2 - 10 MeV range for most detector materials (see [Figure 2.1](#)); thus it is natural to utilize our understanding of Compton scattering to detect MeV photons, especially as we are unable to focus γ -rays at these energies with current technologies. When a photon undergoes a single Compton scatter, the scattered energy E' is related to the initial energy E_0 and the Compton scatter angle ϕ as follows ([Compton, 1923](#)):

$$E' = \frac{E_0}{1 + \frac{E_0}{m_e c^2} (1 - \cos \phi)} \quad (2.1)$$

Consider a photon that Compton scatters in a detector one or more times and is then photoabsorbed. By measuring the energy of the recoil electron $E_e = E' - E_0$ for each Compton scatter, we can determine the initial energy of the photon before it interacted with the detector. We can also determine the Compton scatter angle ϕ , which constrains the origin of the photon to a circle on the sky, called the Compton circle. When a source emits multiple photons, we can determine the source location at the point where the Compton circles overlap.




Figure 2.1: The cross sections of various interactions in Germanium. Compton scattering dominates between 0.2 and 10 MeV. The cross section data is from NIST XCOM.

Compton telescopes are inherently sensitive to polarization. The Klein-Nishina differential cross section (Klein & Nishina, 1929) describes the relationship between Compton scatter angle ϕ , azimuthal scatter angle η between to the photon's electric field vector and scatter direction, and initial and final energies E_0 and E' , respectively:

$$\frac{d\sigma}{d\Omega} = \frac{r_e^2}{2} \left(\frac{E'}{E_0} \right)^2 \left(\frac{E'}{E_0} + \frac{E_0}{E'} - 2 \sin^2 \phi \cos^2 \eta \right) \quad (2.2)$$

Due to this dependence of the differential cross section on η , a polarized beam of photons incident on the detector will be modulated in η . By measuring this modulation, Compton telescopes can determine the polarization properties of a source.

2.1.1 TYPES OF COMPTON TELESCOPES

Compton telescopes can be divided into two types: scatterer-absorber Compton telescopes and compact Compton telescopes. Scatterer-absorber Compton telescopes consist of two separated detector planes. Photons Compton




FIG. 2.—Schematic view of COMPTEL.

Figure 2.2: A schematic of [COMPTEL](#). A photon is shown scattering in the top plane and traveling to the bottom plane, where it is photoabsorbed. Figure from [Schönfelder et al. \(1993\)](#).

scatter in the top plane and are photoabsorbed in the bottom plane. The detector planes are placed far apart so that the time difference between scattering and absorption can be measured. This ensures that background photons moving from the bottom to the top can be rejected during analysis.

[COMPTEL](#) is the most well known example of a scatterer-absorber Compton telescope; a schematic of the instrument is shown in Figure 2.2. The scatterer plane was made of low- Z scintillators and was 1.5 meters above the absorber plane, made of high- Z scintillators. Despite the scientific success of [COMPTEL](#), the scatterer-absorber design led to very a low efficiency of 1%, as only events that scattered in the top plane and were photoabsorbed in the bottom plane were used for analysis ([Schönfelder et al., 1993](#)). This design does not allow for photons that Compton scatter multiple times or events that have large Compton scatter angles, and is thus quite restrictive.

Compact Compton telescopes achieve higher efficiencies than scatterer-absorber Compton telescopes by forgoing the concept of a separate scatterer and absorber. Instead, they consist of one active detector volume in which

a photon can Compton scatter multiple times in any direction before being photoabsorbed. Thus, compact Compton telescopes can detect events within the full range of Compton scatter angles. Figure 2.3 shows an example event in a compact Compton telescope. The photon Compton scatters twice, once at \vec{r}_1 depositing energy E_1 and once at \vec{r}_2 depositing energy E_2 , before being photoabsorbed at \vec{r}_3 and depositing energy E_3 (note that the initial energy of the photon $E_0 = E_1 + E_2 + E_3$). These interaction positions and energy deposits, coupled with the Compton equation, specify the initial Compton scatter angle ϕ and constrain the origin of the photon to the surface of a cone, which can then be projected onto the red circle on the sky.

The disadvantage of compact Compton telescopes is that the interactions occur closer together in space, and consequently, given current achievable timing resolution, it is impossible to determine the order of the interactions by measuring the interaction times. There are a variety of techniques to deduce the interaction order, hereafter referred to as the *event reconstruction*; these include Compton kinematic reconstruction (Boggs & Jean, 2000) and a Bayesian method (Zoglauer, 2005), as well as a machine learning method based on neural networks (Zoglauer & Boggs, 2007).

2.1.2 ANGULAR RESOLUTION MEASURE

The angular resolution measure (ARM) describes the angular resolution of a Compton telescope. The ARM of each event is defined as the smallest angular distance from the Compton circle to the known source location (see Figure 2.4). The full width at half maximum (FWHM) of the distribution of all ARM values from an event list defines the achievable angular resolution after the event reconstruction. The ARM values of each event can be positive or negative, depending on whether the source location is outside or inside the Compton circle, respectively.

The ARM is highly dependent on the energy and position resolution of the detector: a detector with improved resolution will result in a smaller ARM FWHM. This dependence on detector resolution makes the ARM an excellent way to benchmark Compton telescopes. Nonetheless, there is a fundamental limit to the ARM FWHM due to the fact that the Compton equation (Equation 2.1) assumes that the γ -rays scatter on electrons that are at rest. In reality, the electrons are bound to a nucleus and have some momentum. As a consequence, the relationship between the scattered γ -ray energy and the electron energy is not exactly what is assumed by the Compton equation, which intro-




Figure 2.3: A schematic of a compact Compton telescope. A photon originates from the yellow star and interacts three times in the detector. First the photon Compton scatters at \vec{r}_1 and deposits energy E_1 . Next it Compton scatters at \vec{r}_2 and deposits energy E_2 . Finally the photon is photoabsorbed at \vec{r}_3 and deposits its remaining energy E_3 into the detector. From the energy deposited at each interaction location, we can determine the initial Compton scatter angle ϕ and restrict the origin of the photon to somewhere on the red circle.

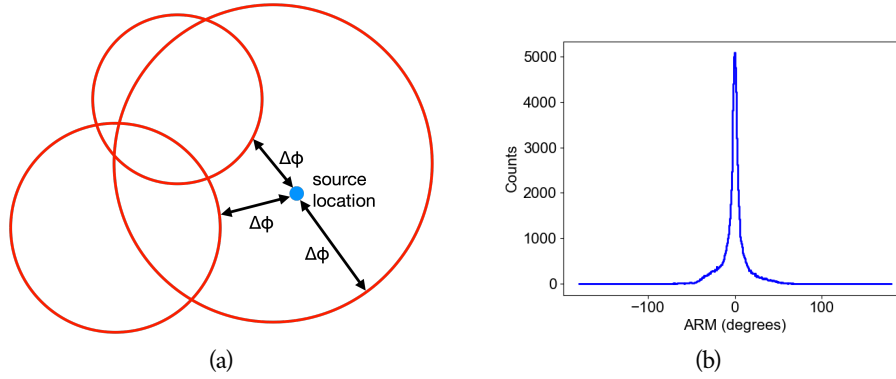


Figure 2.4: (a) A schematic demonstrating how the ARM is calculated. The source location is the blue dot and the Compton circles for three events are represented in red. (b) An example ARM distribution measured by COSI. The data is from a ^{137}Cs calibration source that emits γ -rays at 661.66 keV.

duces broadening into the measured γ -ray spectrum. This effect, referred to as Doppler broadening, fundamentally limits the ARM FWHM.

2.1.3 COMPTON IMAGING

Without measuring the energy of the recoil electron, it is not possible to determine the photon's origin along the Compton circle. Nevertheless, the location of a source emitting multiple photons can be determined via Compton imaging techniques. For two dimensional imaging applications, the Compton circles are backprojected on the imaging plane, as shown in Figure 2.5. The Compton circles overlap at the source location.

Iterative deconvolution techniques can be employed to determine the most likely source distribution given the backprojection and the instrument response. These techniques include list mode maximum-likelihood expectation-maximization (LM-MLEM) (Wilderman et al., 1998) and maximum entropy (Gull & Skilling, 1984). The images from COSI data in this work were made with the LM-MLEM algorithm, as described in Zoglauer et al. (2011).

2.2 THE COSI INSTRUMENT

The Compton Spectrometer and Imager (COSI) is a compact Compton telescope designed to study astrophysical sources from a balloon platform. COSI





Figure 2.5: A backprojection of about 10 events. The Compton circles are visible in red. The bright hot spot at the center of the image is at the source location, where most of the Compton circles overlap.

has been in development for the past 10+ years: the precursor instrument Nuclear Compton Telescope (NCT) first flew as a prototype in 2005. In this section I will describe the COSI instrument, including the detectors, the data acquisition system, and the other gondola systems.

2.2.1 GERMANIUM DETECTORS

COSI utilizes an array of high purity, cross strip germanium detectors (GeDs) (Amman & Luke, 2000). These semiconductor detectors are operated as fully depleted, reverse-biased diodes. The detectors are fabricated at the Lawrence Berkeley National Laboratory using an amorphous semiconductor contact technology (Amman et al., 2007).

COSI's detector array consists of twelve such GeDs stacked in a $2 \times 2 \times 3$ configuration. Each GeD has a volume of $8 \times 8 \times 1.5 \text{ cm}^3$, leading to a total active detector volume of $1,152 \text{ cm}^3$. The electrodes on the anode and cathode sides of each GeD are segmented into 37 strips with a 2 mm strip pitch. A 2 mm guard ring surrounds the strips on each side to prevent surface leakage current from flowing between the anode and cathode. Each GeD operates at a voltage of 1000 V, 1200 V, or 1500 V. The readout electronics on the high



(a) A single COSI GeD with mirror



(b) The COSI GeD detector array

Figure 2.6: (a) A single COSI GeD in front of a mirror, showing the orthogonal cross strips: the 37 strips on the front are vertical and the 37 strips on the back (in the mirror) are horizontal. Each GeD is $8 \times 8 \times 1.5 \text{ cm}^3$. (b) The COSI GeD detector array, showing the $2 \times 2 \times 3$ configuration. The GeDs are mounted to a copper cold finger, which is the U-shaped bar dividing the array vertically.

voltage side, or cathode, are AC-coupled to the detector, and so we refer to this side as the AC side. Similarly, the readout electronics on the anode are DC-coupled, and so we refer to the anode as the DC side.

The COSI GeDs have good energy and position resolution, of 0.2 - 1% and 2 mm³, respectively. The excellent energy resolution provided by germanium is the reason to use this detector material. The $x - y$ position resolution is determined by the number of strips and the strip pitch, as the $x - y$ position is the intersection point between the AC and DC strips (see Section 2.3.3). The z position resolution is determined by the timing resolution, as we deduce the z position from the difference in charge carrier arrival time (see Section 2.3.4).

2.2.2 CRYOSTAT AND CRYOCOOLER

Germanium detectors must be kept at cryogenic temperatures to operate: at room temperature, the band gap of $\sim 0.6 \text{ eV}$ is too small to prevent electrons from reaching the conduction band via thermal excitations, and thus it is difficult to distinguish between electrical noise and a true signal. The COSI GeD array is integrated into an aluminum cryostat (Figure 2.7) and cooled to cryo-




Figure 2.7: A photograph of the COSI cryostat. The footprint of the aluminum shell of the cryostat is about 19×18 inches. The cooling fluid jacket around the cryocooler body enables fluid to be pumped around the cryocooler as part of the liquid cooling system. The preamplifiers, not attached to the cryostat in this photo, connect to the flex circuits.

genic temperatures with a Sunpower CryoTel CT mechanical cryocooler (Figure 2.8). The cryostat is evacuated to $\sim 10^{-6}$ Torr as any residual gas will add heat to the GeDs. We use a mechanical cryocooler rather than liquid nitrogen cooling so that the balloon flight duration is not limited by the amount of consumable liquid nitrogen on board, thus allowing for ultra-long duration balloon (ULDB) flights.

The detectors are attached to a U-shaped copper cold finger (see Figure 2.6b) that is thermally connected to the cryocooler cold tip. There is an observed differential between the cold tip and cold finger temperatures; setting the cold tip to 77 K typically leads to a cold finger (and thus detector) temperature of 84 K. To maintain these temperatures, the cryocooler dissipates ~ 100 W, which causes the cryocooler itself to heat up. As the cryocooler heats up, it dissipates more power, which further heats the cryocooler, causing a positive feedback loop. Once the cryocooler reaches $\sim 70^\circ\text{C}$, parts inside its





Figure 2.8: A photograph of the cryocooler before integration into the cryostat and without the cooling fluid jacket.


body begin to melt. The problem worsens at balloon altitudes where there is no convective cooling due to the lack of atmosphere. To prevent this problem, we developed a liquid cooling system to cool the cryocooler body, shown in Figure 2.9. 3M Novec 7200 cooling fluid is pumped around the cryocooler and then through a copper radiator plate. The radiator plate is painted white to increase its radiative cooling capabilities and rests on top of the gondola frame, pointing straight up. To connect the cryocooler, radiator plate, and fluid reservoir, we use nylon hoses that attach via safety wire to barbed fittings on the aforementioned parts. To pump the fluid, we connect two Fluid-o-Tech FG209 magnetic drive pumps, but only run one pump at a time: if the primary pump breaks, the secondary pump can take over. The liquid cooling system successfully maintains a safe cryocooler temperature both in the lab and in flight (see Section 4.4).

2.2.3 ANTI-COINCIDENCE SHIELDING

Six cesium iodide (CsI) scintillators that constitute both passive and active anti-coincidence shielding surround the cryostat on its bottom and sides (Figure 2.9b). Each shield piece is $40 \times 20 \times 4 \text{ cm}^3$ and is instrumented with two photomultiplier tubes (PMTs) to detect the scintillation light. The thickness of the shield pieces helps prevent atmospheric γ -rays from reaching the GeDs. Additionally, any events detected in coincidence between the shields and the GeDs are rejected. This anti-coincidence system not only reduces the atmospheric γ -ray background, but also reduces the number of incompletely absorbed events detected in the GeDs. Incompletely absorbed events are photons that scatter out the GeDs and thus do not deposit all of their energy in the germanium; they are impossible to reconstruct and are effectively background. The CsI shields constrain COSI's field of view to 25% of the sky.



(a) A schematic of the liquid cooling system.



(b) A photograph of the cryostat, CsI shields, and liquid cooling system.

Figure 2.9: (a) A schematic and (b) a photograph of the liquid cooling system mounted on the gondola along with the cryostat and CsI shields. The pumps circulate cooling fluid from the reservoir around the cryocooler and through the copper radiator plate, where the heat is radiated away. The four side CsI shield pieces are visible, and two more sit underneath the cryostat.

The two **PMT** signals from a single shield piece are summed together and shaped using a bipolar shaping circuit. Then the signals from all the shield pieces are logically OR'ed together to generate a single output signal. If the output signal is high 0.7–1.1 s after the **GeD** signal, the event is vetoed. The threshold of each shield piece is set to 80 keV.

To set the threshold to a particular energy, we must calibrate each shield piece to determine the relationship between energy and threshold **DAC** value. To do so, we turn on a single shield piece and illuminate it with a calibration source that emits monoenergetic γ -rays. We set the threshold to each possible **DAC** value (0 through 255) and measure the count rate. As the threshold approaches the energy emitted by the source, the measured count rate decreases. At a threshold equal to the energy emitted by the source, the count rate is half as high as it was at low thresholds. The data of count rate versus threshold **DAC** can be fit with a complementary error function where the mean is the energy emitted by the source and the sigma is the energy resolution of the shield piece. Using two calibration sources, ^{241}Am and ^{57}Co , that emit γ -rays at 59.5 keV and 122.0 keV respectively, we determine the threshold **DAC** value at two energies. Figure 2.10 plots the count rate versus threshold **DAC** for the two calibration sources and the background. With these two data points, the linear relationship between energy and threshold **DAC** is established (see Figure 2.11), and we can determine the **DAC** value for any energy.

Though their primary purpose is shielding, the **CsI** scintillators can also be used for **GRB** detection: a short spike in the shield count rate is indicative of a potential **GRB**. During the 2016 flight, the operations team responded to triggers in the shields, which led to the team's prompt detection of GRB 160530A (see Section 4.6).

2.2.4 SIGNAL READOUT

Each strip on each side of each detector ($37 \times 2 \times 12 = 888$ strips total) has an individual readout channel where the signal is amplified and shaped. The strip electrode signals leave the cryostat via Kapton-Manganin flex circuits that are coupled to charge-sensitive preamplifiers. The preamplifiers are mounted to the sides of the cryostat as shown in Figure 2.12 and convert the charge induced on each strip into a voltage signal. The output of the preamplifiers are sent via coaxial ribbon cables to the “card cages”, which execute the trigger logic and shape the pulse.

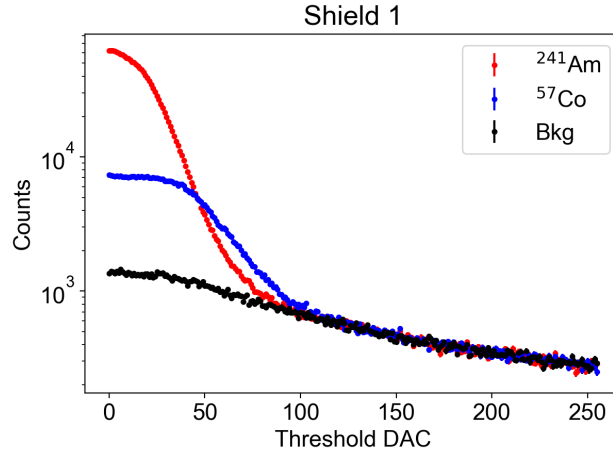


Figure 2.10: Count rate versus threshold DAC for ^{241}Am , ^{57}Co , and background as measured by shield piece 1. After subtracting the background from the source distributions, we determine the energy to threshold DAC relationship by finding the mean of the resulting complementary error functions. The mean corresponds to the energy emitted by the calibration source in question.

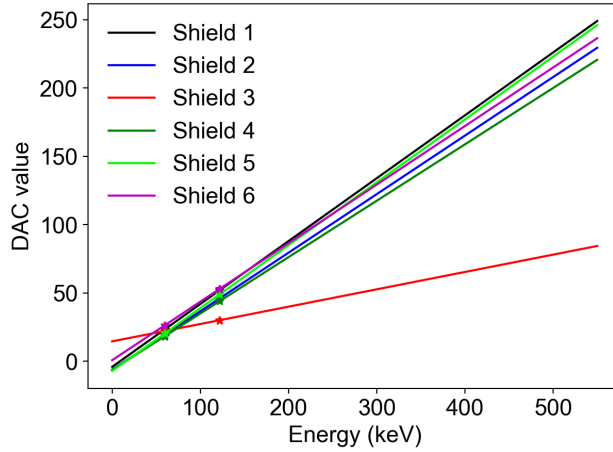


Figure 2.11: The relationships between energy and threshold DAC for each of the six shield pieces. The starred points represent the two calibration source energies at 59.5 keV and 122.0 keV. The energy-DAC relationship is shown for the full threshold DAC range. Shield piece 3 is a clear outlier due to its exceptionally poor energy resolution.

A single card cage holds eight “analog boards”, each of which holds 10 pulse shaping amplifier circuits; one card cage can fully outfit one detector. These circuits consist of a fast bipolar shaper that measures the relative collection time of the charge carriers with 5 ns resolution and a slow unipolar shaper that precisely measures the pulse height of the signal with minimal noise. Each card cage also contains one “DSP board”, which communicates with the flight computer and handles the trigger logic. For the event to proceed from the card cage to the flight computer, at least one strip on each side of the detector must trigger the fast shapers within a 360 ns window. After the 360 ns window, the card cage waits $2.4 \mu\text{s}$ for veto signals from the anti-coincidence shields or the detector guard ring. If no veto signals are present during the veto windows, the pulse heights are accurately measured with the slow shapers and the event information is sent from the card cage to the flight computer.

A single card cage is shown in Figure 2.13. In addition to the data acquisition, the card cages also provide power to the preamplifiers and provide the high voltage power to the GeDs. The 12 card cage internal clocks are synchronized so that photons that interact in multiple detectors are correctly grouped as a single event and then properly reconstructed in the analysis pipeline. To synchronize the card cages, the flight computer sends a “sync” signal indicating to what the internal card cage clock counters should be set. The flight computer also sends each card cage a 10 MHz clock pulse; upon receiving the rising edge of the clock pulse, the card cages increment their internal clock counters by one.

2.2.5 GONDOLA SYSTEMS

The COSI gondola is a three-tiered aluminum structure that holds the instrument together. The cryostat and CsI shields sit at the very top of the gondola and have an unobstructed view of the sky. The middle gondola tier is the thermally insulated electronics bay, which holds the card cages, flight computer, and other electronics boxes described in this subsection. To fully insulate the electronics bay, the floor is separated from the gondola frame with insulating washers, and panels made of foam and aluminum surround the sides and the top. The bottom gondola tier holds the Support Instrument Package (SIP) provided by Columbia Scientific Ballooning Facility (CSBF). Figure 2.14 shows the fully integrated COSI instrument ready for flight.




Figure 2.12: (a) A full preamplifier box (left) next to a partially built box (middle) and a single preamplifier board (right). Each box contains four boards. (b) The cryostat with the preamplifiers mounted and connected to the flex circuits. Each detector connects to two preamplifiers, one for each side. The coaxial ribbon signal cables are wrapped in white teflon and are plugged into the preamplifier output.



Figure 2.13: A single card cage with one of the analog boards pulled out, showing the 10 pulse shaping circuits. Above the eight analog boards is the DSP board, which handles the trigger logic and relays the data to the flight computer. The next board is the low voltage power supply, which powers both the rest of the card cage and the preamplifiers. The top board is the high voltage power supply which provides 1000 – 1500 V to the detectors.




Figure 2.14: A photograph of the fully integrated COSI gondola before the 2016 flight from Wanaka, New Zealand. The side panels are removed to provide a view of the electronics bay and SIP bay.

2.2.51 FLIGHT COMPUTER

The low-power, dual-core COSI flight (or instrument and detector processing unit (IDPU) in Figure 2.14) is responsible for receiving and storing the raw data from the card cages and other subsystems, and for controlling most of the gondola subsystems. The flight computer uses a standard Linux operating system. The flight code is written in C and runs autonomously. Data forwarded from the gondola subsystems are stored on three redundant solid state hard drives. These data include housekeeping data from each subsystem as well as event data from the card cages.

As the interface between the other gondola subsystems and the ground software, the flight computer sends data through the telemetry streams and receives commands. From the ground, we are able to turn on and off each subsystem, change the configuration of the card cages (i.e. change the thresholds set for each strip), and toggle heaters and the liquid cooling system on and off for thermal control.

The flight computer also does some real time, on-board processing of the event data. The telemetry streams do not have large enough capacity to send down all of the event data generated, so the flight computer's on-board processing selects the most scientifically useful data for transmission. Particularly, the algorithm identifies Compton events, which are events where the photon has interacted in the detectors at least twice. Compton events are more useful than single site events because we need at least two interactions to get any sky position information from the event. Removing the single site events from the telemetry data stream reduces the data rate sufficiently so that everything else can be sent down in real time. The single site events are saved on the flight computer hard drives, so we retrieve that information if the instrument is recovered.

2.2.52 TELEMETRY

There are multiple telemetry streams available to the gondola. The fastest is line of sight (LOS), a direct radio link between COSI and the CSBF ground station at the launch site. Though LOS provides downlink speeds of up to 1 Mbps, we can only use it at the beginning of the flight while the gondola is within the line of sight; at some point mountains or the curvature of the Earth block the LOS stream. Despite the short use time of LOS, it is prudent to have this reliable and fast connection at the beginning of the flight when we

are the most uncertain about the instrument health and performance at float altitudes.

After COSI physically leaves the line of sight, the primary telemetry stream is the openport (OP) provided by Iridium, a satellite communications company. There are two OP transceivers on the gondola for redundancy, each of which can send and receive data at a maximum rate of 128 kbps and an average rate of 85 kbps. The OP transceivers are connected to the flight computer via ethernet. The flight computer sends two copies the data directly to the COSI control unit linux machine in Berkeley. Each data copy sent over the internet link is provided by a single OP transceiver.

The gondola also holds two low speed Iridium modems as backup telemetry systems. These modems transmit 255 bytes of data every 15 minutes, an insufficient rate for transmitting science data. Instead, we transmit relevant housekeeping data to help us diagnose potential problems with the OP connection. The low speed Iridium modems can also be used in a dial-up mode, with a data transmission rate of 2 kbps. This slightly improved speed can relay the housekeeping information for faster diagnosis of a failed OP connection and continuous monitoring of instrument health.

2.2.5.3 POWER SYSTEM

COSI is powered by 15 SunCat photovoltaic arrays that both directly power the instrument and charge 24 Odyssey PC1100 lead acid batteries. At night, the instrument is powered by the batteries. There is enough battery capacity to survive nights that last ~15 hours with COSI's average power consumption of ~450 W. The power system is controlled by a Charge Controller from MPPT by Morningstar Corporation (the power supply unit (PSU) in Figure 2.14). The power distribution unit (PDU) distributes power to the gondola subsystems. The individual PDU channels can be switched on and off by the flight computer.

2.2.5.4 ASPECT SYSTEM

Real-time knowledge of COSI's position and orientation is useful both for flight operations and positional information of science data in Galactic coordinates. To acquire this information, we use a Magellan ADU5 differential Global Positioning System (dGPS). The dGPS provides the heading, pitch, and roll of the gondola in addition to latitude, longitude, and altitude. As backup

aspect systems, we use a Trimble BX982 GPS and an Applied Physics Systems Model 544 magnetometer.

2.3 INSTRUMENT CALIBRATIONS

We perform calibrations to convert from the measured parameters of ADC value, timing, detector ID, and strip ID to the physical parameters of deposited energy and three-dimensional interaction position. To determine the conversion between measured and physical parameters, we take extensive data in the lab using calibration sources that emit γ -rays at known line energies, and place these sources at known positions relative to the detectors. Note that COSI does not fly with a calibration source, so all calibration data is taken on the ground before the launch. These measurements also allow us to characterize our instrument and to benchmark our simulations as needed for determining the instrument response, described in Chapter 3. In this section, I will briefly discuss each step in the calibration pipeline.

2.3.1 ENERGY CALIBRATION

The goal of the energy calibration is to determine the relationship between ADC value, or pulse height, and energy. Each strip is calibrated individually due to variations in the gain between strips. First we measure the ADC value at about 15 different line energies using the calibration sources. Then we fit the ADC-energy relation for each strip with an empirical model – a third or fourth order polynomial – that accounts for any non-linear deviations. Figure 2.15 shows the measured spectrum of multiple lines and the resulting ADC-energy relation. Once we have determined the ADC-energy relation, events passing through the calibration pipeline can be easily converted from ADC to energy.

The preamplifiers are temperature sensitive and can shift the line up to 0.5 keV per degree Celcius at 662 keV. We determined experimentally that the peak shift changes linearly with temperature. With this information, we were able to correct for the temperature dependence of the preamplifiers.

2.3.2 CROSSTALK AND CHARGE LOSS

Crosstalk is an effect where one electronics channel influences others. In the case of the COSI detectors, if two nearby strips trigger, the energy recorded is




Figure 2.15: (a) Spectrum of multiple lines from a variety of calibration sources that is used to determine the ADC-energy relation and (b) the ADC-energy relation for strip 7 on the AC side of detector 7 (the specific strip was chosen arbitrarily).

amplified for both strips and the total energy is higher than it would be if all the energy were deposited on a single strip. Previous work by Bandstra (2010) indicates that the effect is strongest for adjacent strips, evident for strips with one strip in between, and negligible with any more distance between strips. We correct for this effect following the method described in Bandstra (2010), which is a straightforward procedure because the offset due to crosstalk is linear with energy.

Charge loss occurs when not all of the charge is collected on the electrodes, and thus the energy measured is lower than its true value. Due to a small lateral electric field in the detectors, any charge carriers in between two strips get stuck and do not make it to the strip electrodes quickly enough to be collected by the readout electronics. It is possible to correct for charge loss, as done in Bandstra (2010). However, since charge loss is only observed on the AC side of the COSI detectors, we can assign each hit the DC side energy and forgo the charge loss correction. We note that we only observe charge loss on the AC side of the detectors most likely because the two sides of the GeDs were processed differently during fabrication.

2.3.3 STRIP PAIRING

Strip pairing determines the x - y interaction position from the AC and DC strip IDs. If there is only one interaction in a single detector, this process is straightforward: the interaction occurred where the two strips intersect (see Figure 2.16a). If there are multiple interactions in a single detector, there are

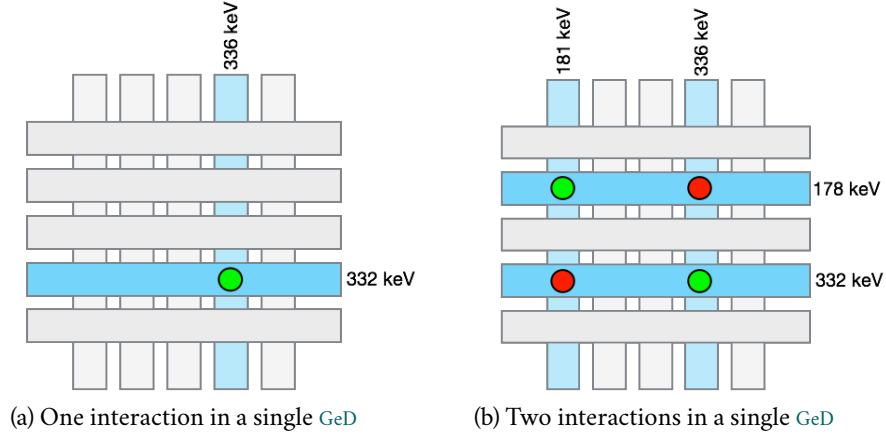


Figure 2.16: (a) If one interaction occurs in a single detector, determining the interaction location (the green circle) is straightforward. (b) If two interactions occur in a single detector, there are two solutions: the interaction locations could be at the red or the green circles. By comparing the energies deposited on each strip, it is evident that the green circles are the correct solution.

multiple solutions to where the interactions took place. Figure 2.16b shows a schematic of two interactions in the detector, where the interactions either occurred at the red circles or at the green circles. As the number of interactions increases, so does the number of possible solutions.

To determine the correct interaction locations, or correct pairing of the strips, we use a greedy algorithm that compares the energies recorded by the AC and DC strips. Greedy algorithms make the locally optimal choice without considering the global consequence of that choice. The advantage of a greedy algorithm is that it approximates an optimal solution with fewer computation steps. In the case of strip pairing, the algorithm begins with a pool of AC and DC strips that triggered in a single detector. It pairs the AC and DC strips that are closest in energy, removes those strips from the pool, and then repeats this process with the remaining strips in the pool. To determine which two strips are the closest in energy, the algorithm uses χ^2 -statistics and calculates the quality $Q_{i,j}$ for each possible pair:

$$Q_{i,j} = \frac{(E_i^{\text{AC}} - E_j^{\text{DC}})^2}{(\sigma_i^{\text{AC}})^2 + (\sigma_j^{\text{DC}})^2} \quad (2.3)$$

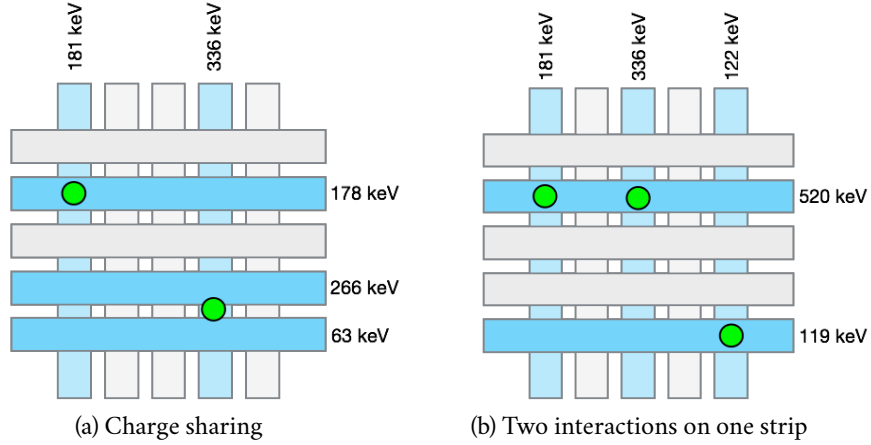


Figure 2.17: Strip pairing is made more complicated when (a) charge sharing between two adjacent strips occurs or (b) there are multiple interactions on one strip. The green circles indicate the interaction positions, which the algorithm recovers by allowing two strips on one side to be paired with one strip on the other.

where i and j denote the specific strips on the AC and DC sides, E is the energy deposited on the strip, and σ is the energy resolution of the strip. A low $Q_{i,j}$ means that the energies are in agreement within measurement errors. Thus the AC and DC strips that are closest in energy are those with the lowest $Q_{i,j}$.

The strip pairing process becomes more complicated due to finite energy resolution, charge sharing between adjacent strips, multiple interactions occurring on a single strip, charge loss and crosstalk effects between adjacent strips, sub-threshold energy deposits, and dead strips (this last phenomenon has a small effect; less than 1% of COSI's strips are dead, due to either a problem with the detector, the readout electronics, or the signal cable between the two). Effects that modify the energy of the strips, such as crosstalk and charge loss, are not currently taken into account in the strip pairing algorithm.

To take into account charge sharing (Figure 2.17a) and multiple interactions occurring on a single strip (Figure 2.17b), the algorithm allows n strips on one side of the detector to be paired with m strips on the other, for $n, m = 1, 2, 3$. However, we observed instances when the algorithm paired more than one strip on one side with more than one strip on the other when the event could be paired more simply with multiple pairs of one strip per side. To avoid such cases, we have modified the algorithm to first pair the

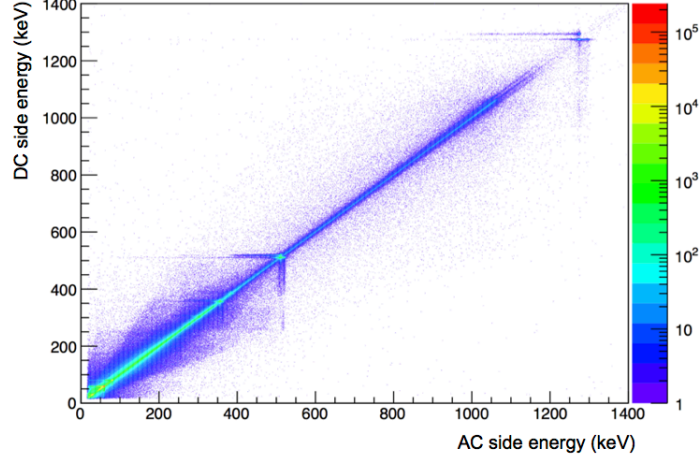


Figure 2.18: Plot of the DC side energy versus the AC side energy for each event of a ^{22}Na calibration source. The hot spots indicate the line energies at 511 keV and 1.274 MeV. The horizontal and vertical lines emerging from the hot spots are due to charge loss. The offsets near the 1.274 MeV line are the result of crosstalk.

strips in the most simple manner possible, only allowing one strip per side in a single pairing. Then we calculate the quality of all of the pairs:

$$Q_A = \frac{1}{N} \sum_{i,j=1}^N Q_{i,j} = \frac{1}{N} \sum_{i,j=1}^N \frac{(E_i^{\text{AC}} - E_j^{\text{DC}})^2}{(\sigma_i^{\text{AC}})^2 + (\sigma_j^{\text{DC}})^2} . \quad (2.4)$$

If Q_A is less than the maximum acceptable value $Q_{\text{max}} = 25$, then the event is paired well, and moves on to the next step in the pipeline. If $Q_A > Q_{\text{max}}$, then we allow for charge sharing between two adjacent strips, repeat the strip pairing process, and reevaluate Q_A . If Q_A is still greater than Q_{max} , we allow for two interactions occurring on a single strip, pair the strips again, and reevaluate Q_A . If Q_A is still too large, we allow for charge sharing across three strips and repeat the process. At this point, if $Q_A > Q_{\text{max}}$, the event is too complicated to pair, most likely due to multiple sub-threshold energy deposits or an interaction that occurred on a dead strip. The event is assigned the pairing that has the lowest Q_A and is flagged so that it can be rejected during later analysis. The fraction of events too complicated to pair is energy dependent and ranges from $\sim 17\%$ at 511 keV to $\sim 28\%$ at 1.274 MeV.

Figure 2.18 plots the AC side energy against the DC side energy for each event for a ^{22}Na calibration source. Ideally, the points should be on the line $y = x$. The offsets due to crosstalk and the tailing due to charge loss are evident around the two line energies, 511 keV and 1.274 MeV. At low energies, the line is broader due to sub-threshold energy deposits.

The algorithm described here is not necessarily the best approach to the problem. Though the greedy algorithm is advantageous because it saves computation time, repeating the algorithm multiple times to allow for charge sharing and multiple interactions on one strip diminishes that advantage. An alternative approach is to calculate Q_A for every possible combination of paired strips and choose the combination with the lowest Q_A . This approach has not yet been implemented for COSI, so the degree to which it improves the performance and the resulting cost to the computation time are currently unknown. Additionally, work is ongoing to develop a machine learning approach to the strip pairing problem using artificial neural networks. A neural network has the potential to perform better than algorithmic approaches because the network can take information into account that is not programmed into an algorithm. For example, it can encode modifications to the energy of adjacent strips due to crosstalk and charge loss. Initial results using a deep neural network for events with up to two strips on each side of the detector currently outperform the above algorithmic methods for simulated data.

2.3.4 DEPTH CALIBRATION

The depth calibration determines the z position, or depth in the detector, as a function of the collection time difference (CTD), which is the difference in collection times of the electrons at the anode and the holes at the cathode. For example, if an interaction occurs close to the anode, the electrons will arrive at the anode before the holes arrive at the cathode. Due to variation in the timing electronics and drift velocities from channel to channel, the depth calibration must find a relationship between depth and CTD for each of the $37^2 \times 12 = 16428$ pixels individually. The methods for determining this relationship are discussed in [Lowell et al. \(2016\)](#).

Applying the depth calibration is fairly straightforward when there is only one triggered strip hit per side of the detector: the CTD is simply the difference in the timing of the two strips. However, as discussed in the previous section, charge sharing can cause n strips on one side of the detector to be paired with m strips on the other, where $n, m = 1, 2, 3$. In situations where $n > 1$ or

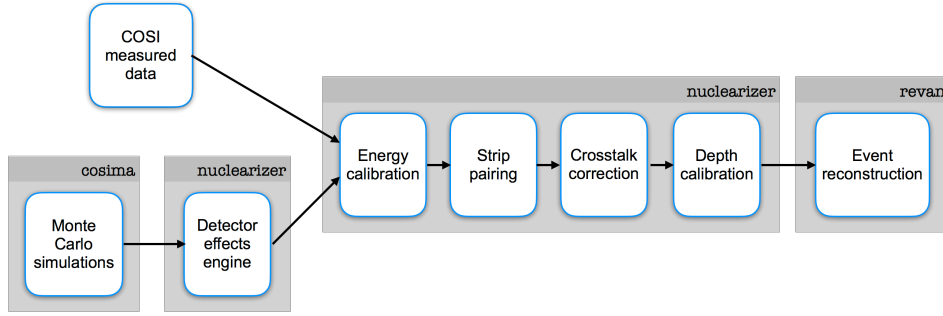


Figure 2.19: Schematic of the analysis pipeline used for COSI.

$m > 1$, it is not obvious how to determine the CTD. Due to the aforementioned variation in the timing electronics, the timing measured by each strip is not an absolute time. As a consequence, the CTD - depth relationship only holds for individual pixels and there is no clear way to combine the timing on two adjacent strips. Instead, in the case of two adjacent strips, the algorithm uses the timing from the strip that has the higher energy. In the significantly more rare case of three or more adjacent strips, the algorithm does not compute a depth but rather flags the event so that it can be rejected during later analysis.

2.4 DATA ANALYSIS

The COSI analysis tools are built on Medium Energy Gamma-ray Astronomy library (MEGALib) (Zoglauer et al., 2006), software written in C++ and ROOT (Brun & Rademakers, 1997) that is designed for the analysis of Compton telescope data. The main MEGALib programs include geomega, the geometry tool used for building the mass model of the instrument, cosima, the Monte Carlo simulation tool based on Geant4 (Agostinelli et al., 2003), revan, the event reconstruction tool, and mimrec, the tool for imaging and other high level analysis.

2.4.1 DATA ANALYSIS PIPELINE

Figure 2.19 describes the data analysis pipeline used for COSI. The analysis pipeline for measured data and simulated data differs: the simulated data first goes through the detector effects engine (DEE), which reverses the event calibration and adds detector effects so that the simulations resemble real data. The simulation pipeline and the DEE are discussed in Chapter 3. After the DEE,

the pipeline for measurements and simulations is the same, and is described here.

For the event calibration, we use `nuclearizer`, a program built on `MEGALib` that performs the `COSI`-specific calibrations. `nuclearizer` takes the raw data or simulated data as input and first applies the energy calibration, which converts from analog-to-digital converter (`ADC`) to energy. The next step is strip pairing, followed by the crosstalk correction. Lastly the depth calibration is applied, which converts from `CTD` to `z` position in the detector.

After the event calibration, we perform the event reconstruction using `revan`. With `revan`, we can choose to cluster nearby hits. The next step is to determine whether an event is single-site or Compton (multi-hit). If the event is a Compton event, `revan` then determines the correct order of interactions and thus the initial Compton scatter angle. There are multiple algorithms to determine the event sequence available in `revan`, including Compton kinematic reconstruction (Boggs & Jean, 2000) and Bayesian reconstruction (Zoglauer, 2005).

After the events are reconstructed, we can proceed with high level analysis. We perform Compton imaging with `mimrec`. Other analysis, such as the spectral analysis pipeline discussed in Chapter 5, is done with software built on but not yet integrated into `MEGALib`.

2.4.2 EVENT SELECTIONS

When analyzing data in Compton telescopes, choosing particular event selections can reduce background, increase detection sensitivity, and optimize the imaging, spectral, and polarization response. Each event can be parameterized by the following:

- **Total photon energy:** A single photon makes multiple energy deposits as it interacts in the detector; the sum of these deposits is the total photon energy. Selecting on the energy is primarily useful for excluding background γ -ray lines (such as the strong atmospheric background line at 511 keV) or for selecting a specific γ -ray line to analyze.
- **Compton scatter angle:** Selecting on the Compton scatter angle of the first interaction (ϕ in Figure 2.3) can increase the overall event quality. For example, backscatter events with Compton scatter angles greater than 90° are more difficult to reconstruct correctly.

- **Distance between first two interactions:** This distance, measured in cm, determines the accuracy of the initial Compton scatter angle and thus the direction of the γ -ray. For smaller distances, determining the Compton scatter angle is subject to more positional uncertainty. Events with larger distances are more accurately reconstructed.
- **Distance between any two interactions:** Interactions that occur very close together – particularly on the same strip – are more difficult to reconstruct correctly. Limiting this distance to larger than one strip width can filter out incorrectly reconstructed events.
- **Number of interaction sites:** Events with more interaction sites are more likely to be properly reconstructed. Two-site events are the most common, so eliminating these will reduce the number of incorrectly reconstructed events at cost to the detection efficiency.
- **Angular resolution measure (ARM):** When the source position is known, any events whose Compton circles are consistent with the source position (within the selection radius) are kept for analysis. In other words, if the ARM of an event is less than the selection radius, the event is kept.

In general, when choosing event selections for a particular analysis, we must balance between stringent event selections that only let through the best events and letting through enough events to build good statistics. A good approach is to choose the event selections that optimize the signal significance relative to the background:

$$S = \frac{r_S}{\sqrt{r_S + r_B}} \quad (2.5)$$

where S is the significance, r_S is the source count rate, and r_B is the background count rate.

In addition to choosing better events with the selections above, we can also perform the following selections to generally reduce background:

- **Earth horizon cut (EHC):** The EHC rejects any event whose Compton circle overlaps the horizon. This is currently the best way to reject albedo radiation background. Work is currently ongoing to replace the EHC with a machine learning approach to background rejection.

- **Pointing cut:** The pointing cut is a selection in the time domain: it selects events that occur during time periods in which a given position is within a given zenith angle relative to the instrument. This cut is useful when analyzing astrophysics data with [COSI](#), as it allows us to select events that are detected when a particular source is within [COSI](#)'s field of view ([FOV](#)).

3

COSI'S DETECTOR EFFECTS ENGINE

Well-benchmarked simulations are essential for high level COSI data analysis. Primarily, we use simulations to compute the instrument response as required for spectral, polarization, and imaging analysis. We also use simulations to benchmark and improve the data analysis pipeline and to better understand the instrument calibrations and in-flight performance. As required for accurate simulations of the COSI instrument, we have developed a detailed detector effects engine (DEE) which applies the intrinsic detector performance to Monte Carlo simulations. With the addition of the DEE, the simulations closely resemble the measurements and the standard analysis pipeline used for measurements can also be applied to the simulations. In this chapter¹, I will describe the simulation pipeline, the DEE, and the benchmarking tests performed with calibrations.

3.1 OVERVIEW OF THE SIMULATION PIPELINE

The simulation pipeline consists of three main steps: the mass model, Monte Carlo simulations, and the DEE. These components will be outlined here.

3.1.1 MASS MODEL

The mass model dictates that the correct amount of material is at the correct position and in the correct shape, and thus determines where in the detectors the simulated interactions occur. Because γ -rays interact with passive material as well as with active detectors, it is important to model all objects near the detectors, included but not limited to the cryostat shell, cryocooler, and preamplifiers. The mass model is implemented in *goomega*, a MEGALib tool to

¹ This chapter is largely based on a publication accepted pending revisions titled “Benchmarking simulations of the Compton Spectrometer and Imager with calibrations” (Sleator et al., 2019). ©Nuclear Instruments and Methods 2019. Reprinted with permission.

3.1 OVERVIEW OF THE SIMULATION PIPELINE




Figure 3.1: The mass model of the full cryostat, including the 12 GeDs, cryocooler, preamps, etc., as well as the shields and PMTs.


manage geometries. Figure 3.1 shows the mass model of the full cryostat and shields. To most accurately match the simulations to calibration data taken in the lab, the mass model used for the following analysis also includes the gondola frame, antenna booms and antennas, card cages, and the reservoir and radiator plate from the liquid cooling system. This mass model containing the entire gondola structure, shown in Figure 3.2, is not always used for COSI analysis, as the simulations take longer to run with this larger mass model.

Accurately modeling the calibration sources in the mass model is essential for comparing simulations to measured data, as γ -rays leaving the source could interact with nearby passive material. Initial scatters off of this material could, for example, change the observed flux of X-ray lines, such as the 32.3 keV line from ^{137}Cs . The calibration sources are Type D disk γ -ray sources from Eckert & Ziegler. Table 3.1 lists the COSI calibration sources and their γ -ray line energies. We included the evaporated metallic salt material and the plastic casing of the calibration sources in the mass model. The sources are held in place by a metal structure made out of 80/20 T-slot material and attached to the structure with a piece of plastic, shown in Figure 3.3. The plastic source holder piece is the closest part of the structure to the source, so γ -rays leaving the source could scatter off of it or be absorbed in it; thus we include it in the mass model. Figure 3.2 shows the source holder in two different positions for two different calibration runs.

3.1 OVERVIEW OF THE SIMULATION PIPELINE

Source	Half life	Activity (μCi)	Energy (keV) (branching ratio)
^{241}Am	432 yr	99.4	59.54 (0.36)
^{57}Co	271 d	2.18	122.05 (0.86)
			136.47 (0.12)
			81.00 (0.33)
^{133}Ba	10.7 yr	8.64	276.40 (0.07)
			302.85 (0.18)
			356.01 (0.62)
			383.85 (0.09)
^{22}Na	2.6 yr	6.34	511.00 (0.90)
			1274.54 (1.0)
^{137}Cs	30.2 yr	91.0	661.66 (0.85)
^{88}Y	107 d	18.7	898.04 (0.94)
			1836.05 (0.99)
^{60}Co	5.3 yr	58.3	1173.23 (0.99)
			1332.49 (0.99)

Table 3.1: **COSI** calibration sources.



(a) Source holder directly above the cryostat

(b) Source holder 50° off axis

Figure 3.2: The mass model of the full gondola frame and antenna booms, including the source holder and source salt. The source holder is floating above the gondola and is positioned (a) directly above the cryostat and (b) 50° off axis.

3.1 OVERVIEW OF THE SIMULATION PIPELINE




Figure 3.3: Photograph of the 80-20 structure used for calibrations. The structure allows for three sources to be placed simultaneously; two of the source holder pieces are shown here.

3.1.2 MONTE CARLO SIMULATIONS

To make Monte Carlo simulations of calibration and astrophysical sources, we use *cosima*, the γ -ray simulation tool in *MEGALib* based on Geant4 (Agostinelli et al., 2003). *cosima* performs Monte Carlo simulations of various source spectra and geometries and can simulate space, balloon, and lab environments. As input, *cosima* requires the mass model, the source position, which can be an astrophysical or local position (i.e. far field or near field), and the source emission properties, including the energy spectrum, flux, and polarization. *cosima* outputs an event list describing interactions in the detectors as defined by the mass model.

When simulating the calibration sources used to benchmark the simulations, we calculated the source activity on the day that the calibration data was taken. For sources with short half-lives, changes in the source activity over the two month period in which the calibration data was taken could potentially have an effect. We also considered the relationships of the emitted γ -ray lines from a single source. For example, in the case of ^{22}Na , although the branching ratio of the 1275 keV line from Na-22 is 100%, 90% of these

decays are effectively coincident with a β^+ decay, which results in two diametrically emitted 511 keV γ -rays. This process, and similar processes for other calibration sources, were included in the Monte Carlo simulations generated by `cosima`.

3.1.3 DETECTOR EFFECTS ENGINE

`MEGALib` includes two ways of processing the `cosima` output: a perfect reconstruction and the standard `MEGALib DEE`. The perfect reconstruction takes the event list and applies the event reconstruction without adding any noise, e.g. energy resolution. The standard `MEGALib DEE` applies the average energy resolution, depth resolution, and thresholds per detector to the simulations, but does not apply any other specific detector and readout electronics effects.

We have developed a `COSI-specific DEE` that includes many of these other effects, including charge sharing, charge loss, crosstalk, and dead time. The `COSI-specific DEE` also uses the measured energy resolution and threshold values of each individual strip rather than an average value over the entire detector. Lastly, the `COSI-specific DEE` reverses the event calibration by converting the physical parameters of energy and position into the measured parameters of pulse height, timing, strip ID, and detector ID. After this conversion, the simulation format mimics the data format and the simulations are run through the `COSI` event calibration pipeline. Thus, any imperfections present in the calibration pipeline will affect the simulations as well as the measurements. With the `COSI-specific DEE`, described in detail in Section 3.2, the simulations much more accurately reproduce the measurements.

3.2 THE COSI DEE

This section describes each step of the `COSI-specific DEE`, hereafter referred to simply as the `DEE`, in the order in which it occurs in the code. In this section, I will adopt the following terminology: an *event* is a single photon that interacts in the detectors. Each interaction is referred to as a *hit*, and an event can have one hit (if the γ -ray is immediately photoabsorbed) or multiple hits (if the γ -ray Compton scatters in the detector before being photoabsorbed). Each hit contains multiple *strip hits*, which refer to the individual strips that trigger during an interaction. A hit must contain at least one strip hit per detector




Figure 3.4: Schematic showing the data acquisition time for the GeD hits.

side, but can also contain multiple strips on each detector side in the case of charge sharing between multiple strips.

3.2.1 SHIELD VETO

If there is a coincident interaction in the shields and in the GeDs, the event is vetoed, as it is not possible to properly reconstruct events that do not deposit all of their energy in the GeDs. If a simulated event interacts in both the GeDs and the shields, that information is readily available and it is straightforward to veto that event. This method, however, does not take into account chance coincidences between the shields and the GeDs. To do so, we need to take into account the nonzero data acquisition time for both the shields and the GeDs.

Figure 3.4 shows the card cage data acquisition timeline. The γ -ray hits the GeDs at $t = 0$. After 700 ns, the shield veto window starts, which lasts for 0.40 μ s. If the shield signal is active at any point during the shield veto window, the event is vetoed. The shield data acquisition timeline is not documented, so the electronics delay time between a γ -ray hitting the shields and the card cage receiving the shield active signal is unknown. We estimated this delay time as 900 ns, as that corresponds with the middle of the shield veto window. The duration of the shield active signal is also unknown. We estimated the shield active signal duration as the average shield dead time per event, 1.7 μ s. This parameter is energy dependent, and could be better estimated by measuring the output of the shield electronics box. Unfortunately, such measurements were not taken before the 2016 COSI balloon campaign.

With our understanding of the data acquisition timeline, we can allow for chance coincidences in the DEE. For every simulated shield hit, the time of that shield hit is recorded as the last shield time. For every GeD hit, we de-

termine the start time of the shield veto window as [GeD](#) hit time + 700 ns. Then we check whether the shield signal is active within the shield veto window, where the shield signal is active between $t_{\text{shield active start}} = \text{last shield time} + \text{shield delay}$ and $t_{\text{shield active stop}} = t_{\text{shield active start}} + \text{shield active signal duration}$. If the signal is active during the shield veto window, the event is discarded.

To assess the impact of chance coincidences, we compared the simulated event rate with and without allowing for chance coincidences. For a simulated ^{137}Cs source, one of the strongest calibration sources with a high count rate, we determined that 13% of the simulated events are rejected due to chance coincidences with the shields. For weaker calibration sources with count rates closer to what is measured during flight, only $\sim 3\%$ of the simulated events are rejected due to chance coincidences with the shields.

3.2.2 PHYSICAL POSITION TO STRIP AND DETECTOR ID

The event list that `cosima` outputs provides a physical (x, y, z) position for each hit. The physical position is converted into the corresponding detector ID and the closest AC and DC strip ID. Each of the two resulting strip hits is assigned the energy deposited in this interaction. We also calculate the depth in the detector for this interaction: based on the z position of interaction within the instrument, we determine the depth within the activated detector based on its relative position in the mass model.

3.2.3 CHARGE SHARING

Charge sharing describes the phenomenon of a single interaction triggering multiple adjacent strips. This effect is attributed to several physical processes. Thermal diffusion and charge carrier repulsion can cause the charge cloud to spread laterally ([Knoll, 2000](#)). [Amman & Luke \(2000\)](#) measured charge sharing in cross strip GeDs and attributed the effect to interactions that physically occur between strips. Work by [Looker et al. \(2015\)](#) indicates that different detector fabrication techniques can increase or mitigate charge sharing.

Including charge sharing in the [DEE](#) is important for a number of reasons. It can be more difficult to properly reconstruct events that contain charge sharing hits, where charge is collected over two or more adjacent strips. Charge loss and crosstalk are effects that occur in the [COSI GeDs](#) and distort the measured energy of adjacent strips. If the measured energy of the hit on one side

of the detector is significantly higher or lower than the measured energy of the hit on the other side, the strip pairing calibration is not able to pair the strips properly and flags the event; these flagged events are discarded later in the analysis pipeline. Additionally, low energy hits may deposit an energy above the strip threshold if all the charge is collected on one strip, but below the threshold if the energy is split into two strips. Thus, charge sharing can cause sub-threshold hits which are not measured by the [GeDs](#), mimicking incompletely absorbed events that are difficult to reconstruct.

Simulating the majority the effects that cause charge sharing is difficult without a detailed charge transport simulation that takes into account variations in the electric field, charge carrier repulsion, and any effects caused by the fabrication method. It may be possible to characterize the charge sharing effect empirically by illuminating each [GeD](#) with a collimated beam, as done in [Amman & Luke \(2000\)](#), but this experiment was not performed before the 2016 flight. Simulating charge sharing due to thermal diffusion, however, is relatively straightforward. Thermal diffusion effectively introduces some spread in the arrival position of the charge carriers. The spread can be characterized as a Gaussian with a width of

$$\sigma_{\text{diffusion}} = \sqrt{\frac{2kTz}{eE}} \quad (3.1)$$

where k is Boltzmann's constant, T is the detector temperature, z is the distance that the charge carrier travels along the electric field direction (the \hat{z} direction), e is the electron charge, and E is the electric field applied to the detector ([Knoll, 2000](#)). To simulate this effect, we divide the deposited energy into charge carriers. For each charge carrier, we randomly draw an x - y drift position from a 2-dimensional Gaussian distribution with width $\sigma_{\text{diffusion}}$ and centered at the x - y interaction position. We then determine the AC strip and DC strip that correspond to the x - y drift position, and add the energy of the charge carrier to those strips.

The effects of charge sharing due to thermal diffusion, however, are small in the [COSI GeDs](#). For a detector temperature $T = 84$ K, an electric field $E = 1000$ V/cm, and an interaction occurring in the middle of the detector at $z = 0.75$ cm, the resulting Gaussian has a width of $\sigma = 0.033$ mm. With a strip pitch of 2 mm, charge sharing due to diffusion is unlikely to have a significant effect. Figure 3.5 compares the simulated and measured distribution of the number of strip hits per event from the 662 keV line of a ^{137}Cs source. Charge sharing affects this distribution, as it will cause more events to have a larger

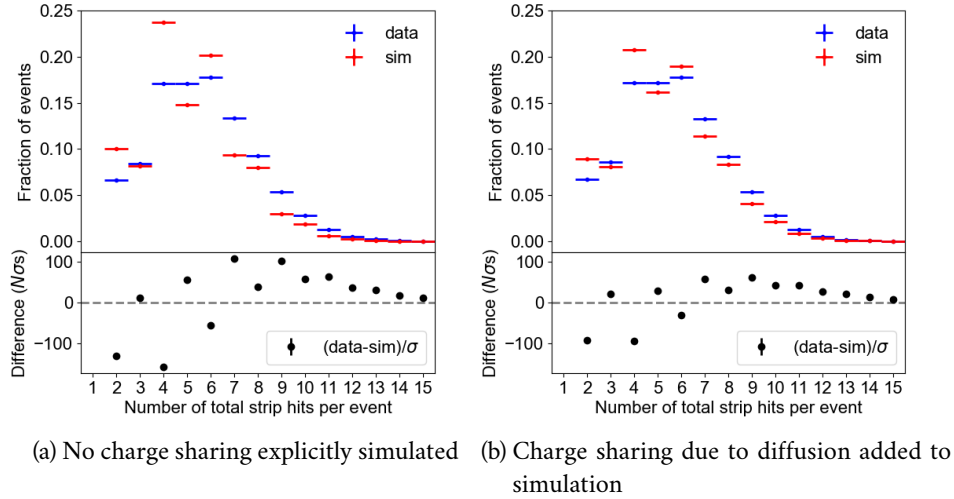


Figure 3.5: The fraction of total events that have n strip hits per event, as a function of the number of strip hits per event, for measurements (blue) and simulations (red). The included events are from the 662 keV line of a ^{137}Cs source; their total energy is between 640 and 680 keV to include events affected by charge loss and crosstalk. No charge sharing is explicitly simulated in (a), whereas in (b), charge sharing due to diffusion was added to the simulations. Charge sharing due to diffusion brings the simulated distribution slightly closer to the measured distribution, but does not have a large effect.

number of strip hits per event. In Figure 3.5a, no charge sharing is explicitly added, and there is a large discrepancy between the measured and simulated distribution. In Figure 3.5b, charge sharing due to thermal diffusion has been added. As expected, the effect is small; however, the simulated distributions are slightly closer to the measured distribution.

Nevertheless, this method can be used to empirically simulate all of the charge sharing in the detectors by increasing σ by a factor N :

$$\sigma = N\sigma_{\text{diffusion}} = N\sqrt{\frac{2kTz}{eE}}. \quad (3.2)$$

The scaling factor N depends on the detector and deposited energy. For each detector, N was empirically selected at each calibration source line energy so that the number of adjacent strip hits in the simulation is equal to that in the data. To find N at any energy, we linearly interpolate between the values at

Energy (keV)	Side	Mean N	Standard deviation of N
122	AC	3.28	1.00
	DC	1.70	0.33
662	AC	3.62	0.63
	DC	1.82	0.39
1274	AC	4.02	0.67
	DC	1.95	1.04

Table 3.2: Interpolated values of N used in the empirical method of simulating charge sharing. The charge sharing distribution is defined by a Gaussian with $\sigma = N\sigma_{\text{diffusion}}$. These values of N were chosen such that the measurements and simulations have the same fraction of events with adjacent strips. For each energy shown in this table, N is averaged over all 12 GeDs. The standard deviation of N over all 12 GeDs is computed to give a sense of the variation of N between detectors.

the calibration source energies. Table 3.2 shows the interpolated N values at a few calibration source energies.

Figure 3.6 shows the measured and simulated distributions of the number of strip hits per event for a ^{137}Cs source, using this empirical method of simulating charge sharing. These distributions are now much closer than in Figure 3.5, in which only charge sharing due to diffusion was present in the simulations. Figure 3.7 compares the measured and simulated distributions of the number of strip hits in each event for a number of different calibration line energies. The agreement between the measured and simulated distributions indicates that this method of simulating charge sharing is a good approximation of charge sharing in the COSI GeDs.

Table 3.3 shows the number of events that are flagged for later removal during the strip pairing and depth calibration components of the event calibration in `nuclearizer`. The table compares the percentage of flagged events in the simulations with and without charge sharing to the percentage of flagged events in the measured data for three calibration source energies. As expected, including the charge sharing in the simulations increases the percentage of flagged events. More strip hits per event inherently complicates the strip pairing process, which can lead to more events flagged by strip pairing. Additionally, charge sharing implicitly leads to more charge loss due to creating more sub-threshold hits, which can also contribute to more events flagged by strip pairing. The depth calibration flags events when the depth of the interaction




Figure 3.6: Measured (blue) and simulated (red) distribution of the number of strip hits per event. The included events are from the 662 keV line of a ^{137}Cs source. Charge sharing was empirically applied to the simulations as described in Equation 3.2, and there is good agreement between measurements and simulations.

cannot be well determined; this occurs when there are too many strip hits and when there are multiple interactions on the same strip. Adding charge sharing to the simulations will increase the number of strip hits per interaction and could potentially cause more instances in which multiple interactions occur on the same strip.

3.2.4 DEPTH TO TIMING

We invert the depth calibration (see Section 2.3) to convert depth into collection time difference (CTD), or the difference between the charge collection times on the AC and DC strips. Though only the CTD is used in the event calibration, each strip hit must be assigned a timing value such that the output format of the DEE accurately mimics the real data format. Because the timing shaper of each strip has a unique offset that is not individually calibrated, the timing value of each strip hit is not meaningful². The relative offsets between two strips that make up a single pixel is automatically taken into account in the CTD during the depth calibration. Thus, we assign each strip an arbitrary

² Note that the event time is determined independently of the strip hit time: each card cage receives a 10 MHz clock signal from the flight computer, which is used to determine the event time.

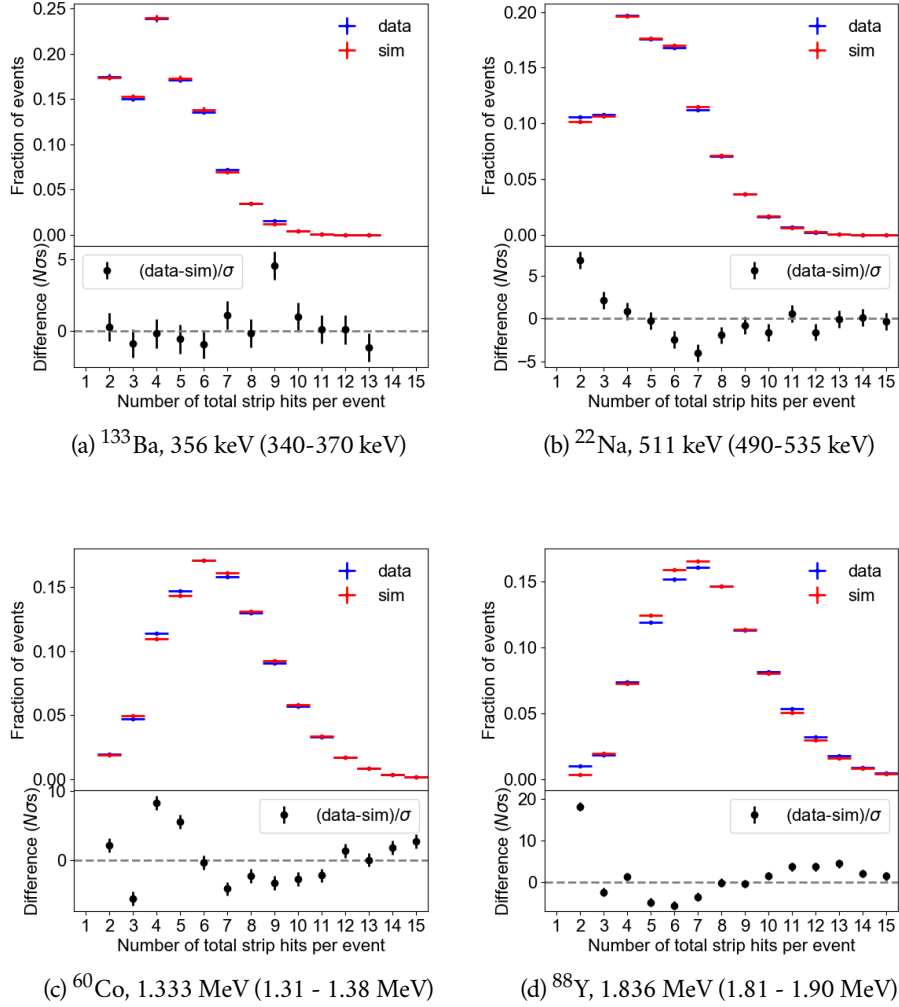


Figure 3.7: Measured (blue) and simulated (red) distribution of the number of strip hits per event for a variety of calibration runs. The energy cut for each line is given in parentheses. The good agreement between measurements and simulations shown here indicates that the empirical charge sharing simulation method described in this section works well for all energies.

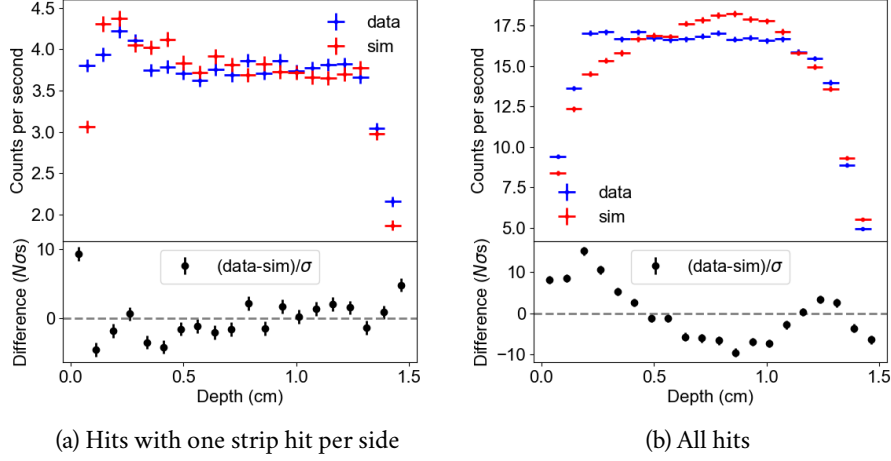


Figure 3.8: Measured (blue) and simulated (red) depth distributions for the top layer of detectors (a) only including hits with one strip hit per side and (b) including all hits. The distributions for the middle and bottom detector layers have the same shape.

timing while ensuring the correct CTD. We then apply Gaussian noise to the timing by randomly drawing a number from a Gaussian distribution centered at zero with a width of 12.5 ns and add that randomly drawn number to the strip hit timing value. The 12.5 ns Gaussian width is determined by the depth calibration (Bandstra, 2010).

Figure 3.8 compares the measured and simulated depth distributions after the depth calibration step. The distributions match well for hits with only one strip hit per side. When all hits are included, however, the differences between the measured and simulated distributions increase. Improvements to the depth calibration could potentially alleviate this issue. As discussed in Section 2.3, the CTD-depth relationship is determined with hits that contain only one strip hit per side, meaning that we do not have a precise depth calibration for charge sharing hits. It is likely that applying the CTD-depth relationship to charge sharing hits skews the measured depth distribution, especially as initial investigations have shown that the timing measurement on a single strip changes when neighboring strips also collect charge.

3.2.5 CHARGE LOSS

Charge loss occurs when not all of the charge deposited in the detector crystal is read out by the electrodes. This effect occurs for a number of reasons. The charge carriers from interactions that occur in between two strips may not arrive at the electrodes within the charge collection time, due to the lack of an electric field in the plane of the strips (Amman & Luke, 2000). Phenomena that cause the charge cloud to spread, such as thermal diffusion and the repulsion of like charge carriers, could result in some charge carriers spreading into the region between two strips and not being collected. Charge loss can also occur due to crystal impurities which can trap charge carriers as they traverse the detector volume (Knoll, 2000; Amman & Luke, 2000).

Charge loss is observed in charge sharing hits, or hits with two adjacent strip hits, on the AC side of the COSI GeDs. Charge loss is not observed on the DC side most likely because the two sides of the GeDs were processed differently during fabrication. Since the effect is only present on the AC side, we do not need to correct for it during the event calibration; we instead choose the hit energy from the DC side during the strip pairing calibration. Nevertheless, it is important to simulate charge loss because its presence causes more events to be flagged during the strip pairing process. These flagged events are discarded after the event calibration stage of the pipeline.

To simulate charge loss, we reverse the correction process described in Bandstra (2010) and summarized here. We select hits that contain two adjacent strip hits and relate the sum of the energies $S = E_1 + E_2$ of the two adjacent strips to the difference of the energies $D = |E_1 - E_2|$. For each hit, (S, D) is plotted, as in Figure 3.9. At high energies (662 keV, Figure 3.9a), the two hot spot clusters are caused by different physical processes: the cluster at differences of ~ 250 keV is made up of backscatter events, where the two adjacent strip hits represent two separate interactions (or hits) in the detector with a Compton scatter angle $\phi \sim 180^\circ$ between the two interactions, while the cluster at differences between 500 and 662 keV is made up of charge sharing hits. Charge loss is evident by the downward slope in the cluster at large differences, where the sum is less than the line energy of the source (in this case, 662 keV). If no charge loss were present in the detector, the sum would be ideally equal to the line energy. At low energies, the photons are much less likely to backscatter at $\phi \sim 180^\circ$ due to the kinematics of Compton scattering, and so all low energy hits with two adjacent strip hits are most likely

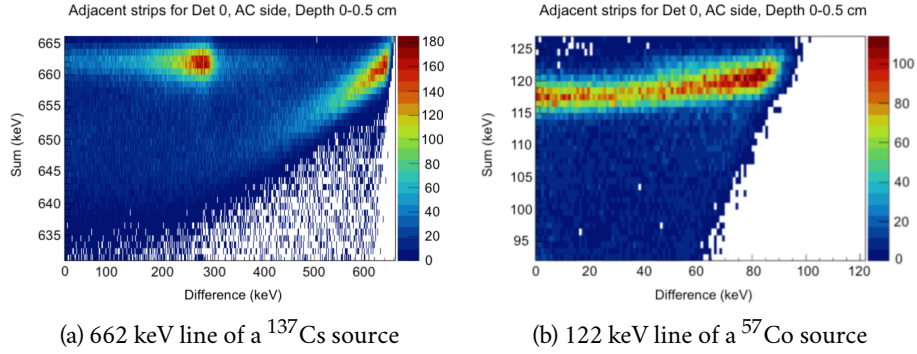


Figure 3.9: Sum-difference histogram of two site events from the AC side of detector 0 within a depth of 0-0.5 cm for (a) the 662 keV line of a ^{137}Cs source and (b) the 122 keV line of a ^{57}Co source. In (a) charge loss is evident by the hot spot at large differences: if no charge loss were present, the sum of all the counts in this hot spot would be at ~ 662 keV. In (b), charge loss is evident in the parabolic dip below 122 keV.

charge sharing hits. Charge loss for low energy γ -ray sources is evident by the parabolic dip in the sum below the line energy, as in Figure 3.9b.

To simulate the charge loss effect, we consider a phenomenological model for the sum S as a function of difference D :

$$S(D) = \begin{cases} E_0 - \frac{B}{2E_0}(E_0^2 - D^2), & E_0 < 300\text{keV} \\ E_0 - B(E_0 - D), & E_0 \geq 300\text{keV} \end{cases} \quad (3.3)$$

where E_0 is the true energy of the hit. The parameter B is the slope of the cluster of charge sharing hits where it meets the line $S(E_0) = E_0$. The model takes into account the linear shape of the outer hot spot at high energies and the curved shape of the hot spot at low energies. B depends slightly on the energy and depth in the detector. We find B as a function of energy for three depth bins in the detector: 0-0.5 cm, 0.5-1 cm, and 1-1.5 cm (the depth is 0 cm at the AC side of the detector and 1.5 cm at the DC side). For each depth bin, we determine B at four different energies (122 keV, 356 keV, 662 keV, and 1333 keV) by fitting the sum-difference plots of four calibration source lines with this model. We then interpolate B linearly as a function of energy to find B at any E_0 . Once the interpolation function is defined, modeling the charge loss effect is straightforward: as each simulated charge sharing hit initially

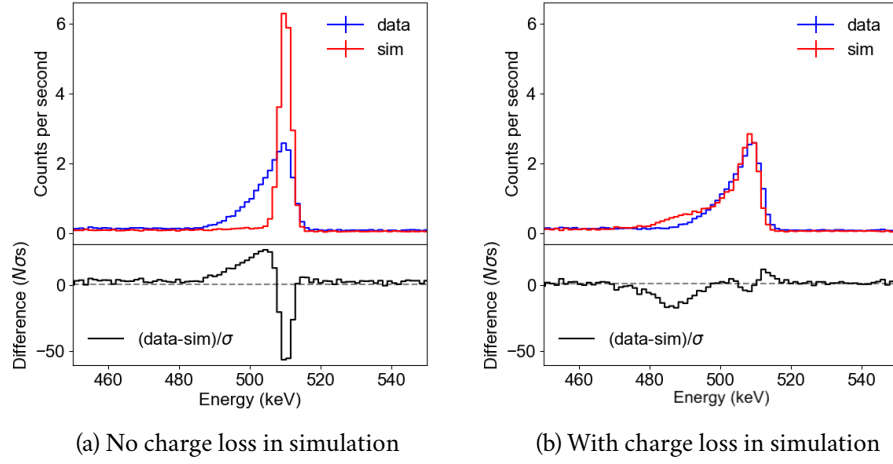


Figure 3.10: The measured (blue) and simulated (red) spectra of the 511 keV line from a ^{22}Na source (a) before and (b) after charge loss is applied to the simulation. The spectra only contain hits with two adjacent strip hits from the AC side of the detector. Adding charge loss to the simulation significantly improves the agreement between measurements and simulations.

has a sum equal to E_0 , we apply Equation 3.3 to estimate the S due to charge loss. To assign the hit a reduced energy of S , the energy of each contributing strip hit must be reduced. We subtract $(E_0 - S)/2$ from the energy of each strip hit to preserve D between the strip hit energies.

Figure 3.10 compares the measured and simulated spectra of the 511 keV line from a ^{22}Na source before and after charge loss is applied to the simulation. These spectra only contain hits with two adjacent strip hits from the AC side of the detector to best see the charge loss effect. After charge loss is included in the simulations, the measured and simulated spectra are in better agreement. The excess tailing on the low-energy side in the simulated spectrum is likely due to the presence of simulated hits with a large D . As the strip hit energies for charge sharing hits were determined empirically as described in Section 3.2.3, a more physical charge sharing simulation could potentially alleviate this excess tailing.

3.2.6 CROSSTALK

Crosstalk is the influence of one electronics channel on another. In the COSI GeDs, crosstalk affects hits that contain two adjacent strip hits by amplifying the energy of each strip hit. In the event calibration, we correct for crosstalk after the strip pairing calibration, so crosstalk still affects the strip pairing process: because the energies of the strip hits affected by crosstalk are higher than they should be, the strip pairing algorithm may not pair the strips correctly or may flag these events for later removal. Since these effects to strip pairing are present in the data, it is important to add crosstalk to the simulations.

To simulate crosstalk, we reverse the correction method described in Section 2.3 and detailed in Bandstra (2010) and Kierans (2018). The crosstalk effect is linear, meaning that the energy increase due to crosstalk is proportional to the energy of the original strip plus an offset due to sub-threshold adjacent strips. In other words, if $E_{1,T}$ and $E_{2,T}$ are the true energies of the strip hits, the measured energies $E_{1,M}$ and $E_{2,M}$ are

$$\begin{aligned} E_{1,M} &= E_{1,T} + \beta E_{2,T} - \frac{\alpha}{2} \\ E_{2,M} &= E_{2,T} + \beta E_{1,T} - \frac{\alpha}{2} \end{aligned} \tag{3.4}$$

where α and β are the sub-threshold offset and correction factor, respectively. These factors are determined for each side of each detector by fitting this model to a spectrum of a known γ -ray line including only hits with two adjacent strip hits on one side of the detector and two non-adjacent strip hits on the other side of the detector. This configuration of strip hits most likely corresponds to two separate interactions in the detector, rather than a charge sharing hit that is adversely affected by charge loss (note that charge sharing hits are affected by crosstalk, but because they are also affected by charge loss, it is best to leave them out when determining α and β).

Once the offset and correction factor are determined, crosstalk can be added to the simulations using Equation 3.4, which estimates the new energy of each strip hit due to crosstalk $E_{n,M}$ as a function of the original strip hit energy $E_{n,T}$. Figure 3.11 compares the measured and simulated spectra of the 511 keV line from ^{22}Na before and after crosstalk is applied to the simulation. The spectra only include hits with two adjacent strip hits from the DC side of the detector to best see the crosstalk effect without interference from the charge loss effect. These spectra are made after the energy calibration but before the strip pairing and crosstalk correction, so the crosstalk effect is still

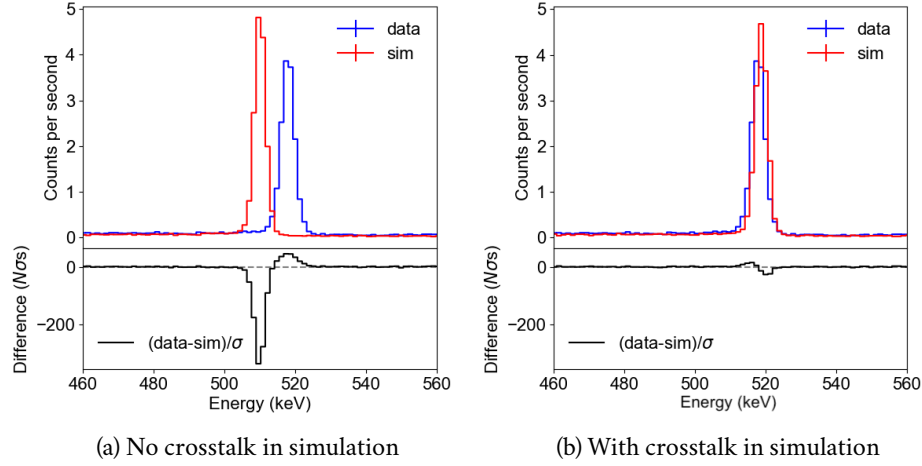


Figure 3.11: Measured (blue) and simulated (red) spectra (a) before and (b) after crosstalk is added to the simulation. The spectra, made before the crosstalk correction, only contain hits with two adjacent strip hits on the DC side of the detectors to clearly view the effect of crosstalk. When crosstalk is added to the simulation, the same line shift is present as in the measurement.

present, causing the line to shift from 511 keV to ~ 519 keV. The same line shift is present in the simulations after incorporating the crosstalk effect.

Table 3.3 shows the result of adding crosstalk and charge loss to the simulations. These effects change the energy of charge sharing hits, which can cause strip pairing to pair the strips incorrectly or flag the events for later removal. The table compares the percentage of flagged events in the simulations before and after crosstalk and charge loss are added to the simulations to the percentage of flagged events in the measured data for three calibration source energies. The flagged events cannot be properly reconstructed and are removed later in the analysis pipeline.

3.2.7 ENERGY TO PULSE HEIGHT

We invert the energy calibration (see Section 2.3) to convert energy into pulse height, or analog-to-digital converter (ADC) value. Each strip has an individual calibration in which the energy is related to the ADC value with a third or fourth order polynomial. We then apply Gaussian noise to the ADC value.

3.2 THE COSI DEE

-	Total events	Flagged events	Flagged strip pairing	Flagged depth calibration
Simulation without charge sharing	1329819	201578 (15.2%)	198569 (14.9%)	3252 (0.2%)
Simulation with charge sharing	1310708	256473 (19.6%)	244687 (18.7%)	1574 (0.1%)
Simulation with crosstalk and charge loss	1274954	262438 (20.6%)	238674 (18.7%)	25199 (2.0%)
Measurement	1094735	255519 (23.3%)	219581 (20.1%)	40359 (3.7%)

(a) ^{137}Cs , 662 keV line, including events with energies between 630 and 670 keV

-	Total events	Flagged events	Flagged strip pairing	Flagged depth calibration
Simulation without charge sharing	269092	35146 (13.3%)	34636 (12.9%)	551 (0.2%)
Simulation with charge sharing	265261	44459 (16.8%)	42411 (16.0%)	2174 (0.8%)
Simulation with crosstalk and charge loss	262633	47287 (18.0%)	43106 (16.4%)	4341 (1.7%)
Measurement	245598	49277 (20.1%)	42346 (17.2%)	7596 (3.1%)

(b) ^{22}Na , 511 keV line, including events with energies between 480 and 520 keV

-	Total events	Flagged events	Flagged strip pairing	Flagged depth calibration
Simulation without charge sharing	64447	12547 (19.5%)	12307 (19.1%)	273 (0.4%)
Simulation with charge sharing	63282	15683 (24.8%)	14915 (23.6%)	833 (1.3%)
Simulation with crosstalk and charge loss	59105	15313 (25.9%)	13821 (23.4%)	1828 (3.1%)
Measurement	57974	19203 (33.1%)	16411 (28.3%)	3482 (6.0%)

(c) ^{22}Na , 1.274 MeV line, including events with energies between 1.240 and 1.284 MeV

Table 3.3: The fraction of flagged events in the strip pairing and depth calibration for three different calibration source line energies. The simulation without charge sharing, with charge sharing, and with charge loss and crosstalk is compared to the measurement. As these effects are applied to the simulation, the fraction of flagged simulated events increases and is closer to the fraction of flagged measured events, as expected.

For each individual strip, the width of the Gaussian as a function of energy is determined during the energy calibration of that strip.

3.2.8 THRESHOLDS AND DEAD STRIPS

Each strip has two pulse shaping circuits. The “fast” shaper measures the time of the interaction to an accuracy of 5 ns resolution and has a threshold of about 40 – 50 keV. The “slow” shaper precisely measures the pulse height of the signal (which corresponds to the deposited energy) with minimal noise, and has a threshold of about 20 keV. Strip hits with energies below the slow threshold are not recorded by the card cage, and strip hits with energies below the fast threshold do not have timing (and therefore depth) information.

The threshold values for each channel differ slightly due to noise and gain variations and must be calibrated separately. Each threshold is set with a digital-to-analog converter (DAC) value between 0 and 255, but as the relationship between DAC and energy is unknown, we must determine the slow and fast threshold values in keV for each strip in order to simulate the thresholds. The slow threshold for each strip is determined by the sharp cutoff in the low-energy spectrum, as shown in Figure 3.12. To apply the slow threshold to the simulations, any strip hits with energy below the slow threshold of that strip are removed.

The energy determination of the fast shaping channel is less precise, so determining the threshold value is more difficult. To do so, we consider two separate spectra for each strip: one containing energy-only events, and one containing events with energy and timing, as shown in Figure 3.12. The spectrum of energy-only events arises because events below the fast threshold lack a timing measurement. The noisiness of the fast threshold is evident in Figure 3.12 as there is no one point where the energy-only spectrum ends and the energy-and-timing spectrum begins. Consider dividing the energy-and-timing spectrum by the total spectrum (energy-and-timing added to energy-only), shown in Figure 3.13. For each ADC value, the corresponding y axis value represents the fraction of events that have timing information. These curves can be well fit with an error function. This error function is then used to determine whether or not a simulated strip hit should have timing: if a random number drawn between 0 and 1 is greater than the value of the error function at the ADC value of the strip hit, the strip hit’s timing is removed.

Figure 3.14 shows a comparison between measured data and simulations for the energy-only spectrum and the energy-and-timing spectrum at low en-




Figure 3.12: The spectra of events with energy only and events with both energy and timing for two arbitrary strips. Events below the fast threshold do not have a timing measurement. The fast threshold is noisy, so there is no clear cutoff between energy-only events and energy-and-timing events. The energy threshold is not noisy, and results in the sharp cutoff to the spectrum located at around ADC value 90 for both strips.




Figure 3.13: The energy-and-timing spectrum divided by the total spectrum for two strips. These distributions represent the fraction of events with timing as a function of ADC value. The distributions are fit with an error function, which is used to simulate the fast threshold.

ergies. There is fairly good agreement between data and simulations, confirming that our implementation of the thresholds in the DEE is accurate.

A very small fraction of the strips (7 out of 888) are dead. These dead strips could be due to a problem with the detector, the readout electronics, or the signal cable between the two. The DEE discards any simulated strip hits that occur on a dead strip.

3.2.9 GUARD RING VETO

Each side of each detector has a guard ring that surrounds the strips to prevent surface leakage current from flowing between the anode and cathode strips. The guard ring is also used to veto events in which interactions occur too close to the edge of the detector, where non-uniformities in the electric field can degrade the detector response. If a strip hit occurs on the guard ring in a detector above the guard ring threshold, all other coincident strip hits that occurred in that detector are discarded. To determine the threshold value in keV, we performed an energy calibration of each guard ring channel; the threshold values range from 14 keV to 47 keV, with an average value of 36.6 keV.

3.2.10 TRIGGER CONDITIONS

The card cages only record events that have at least one strip hit above the timing threshold on both sides of the detector. After all effects (except for dead time) have been applied, the DEE ensures that in each triggered detector there is at least one strip hit per side above the timing threshold. If the trigger conditions are not met for a single detector, all strip hits from that detector are removed. Strip hits from other detectors that do meet the trigger conditions are not removed and continue to the next stage.

3.2.11 DEAD TIME

The dead time is the time after an event is recorded during which the system cannot record another event. Each card cage is dead for $\sim 10 \mu\text{s}$ after a single event occurs in the detector: this is the amount of time it takes for the analog boards to trigger and the coincidence logic to be performed. In the DEE code,

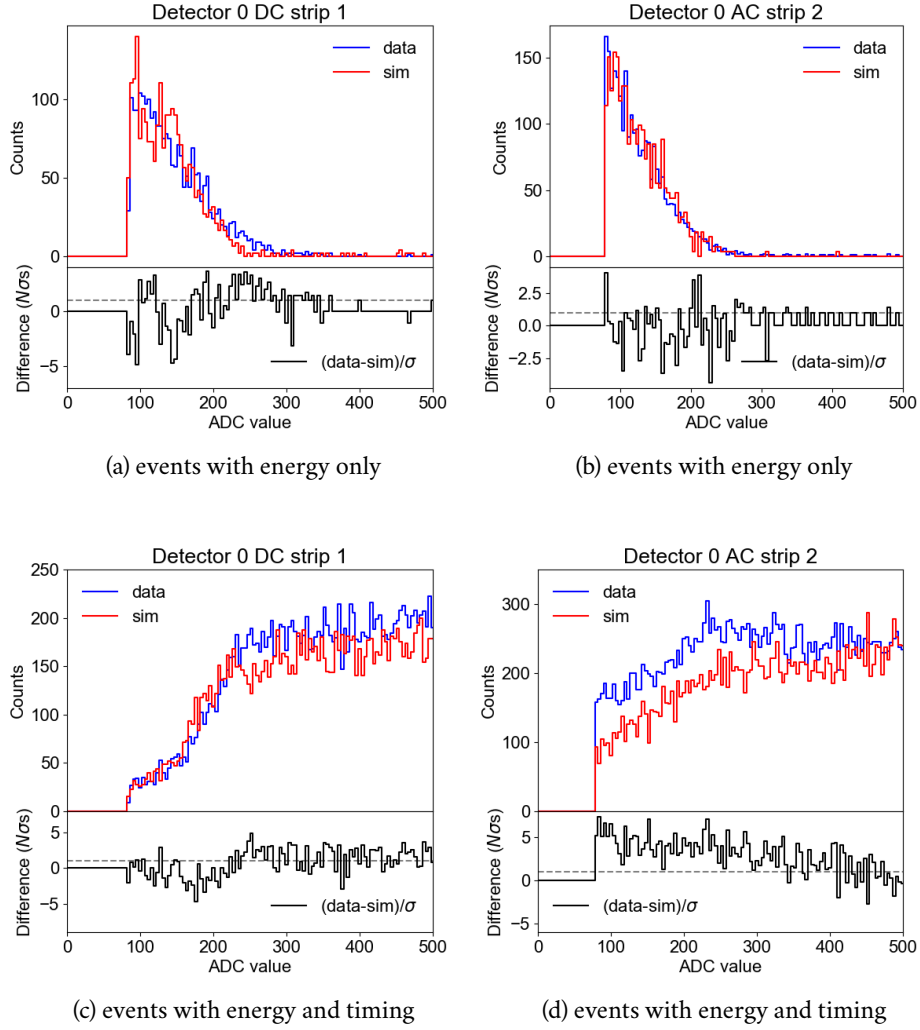


Figure 3.14: A comparison of the measured (blue) and simulated (red) energy-only spectrum (a,b) and energy-and-timing spectrum (c,d) for two strips. There is good agreement between data and simulations.

if a hit occurs in a certain detector less than $10\ \mu\text{s}$ after the previous hit, the hit is discarded.

Additional dead time is introduced by the software on the DSP board in the card cage that takes the information about the triggered event from the [FPGA](#) and parses it into a dataframe packet to send to the flight computer. The card cage takes $625\ \mu\text{s}$ to process one event, which is determined by inverting the fastest measureable data rate per card cage, measured as 1.6 kHz. Each card cage can process up to 16 events simultaneously. The cause of this dead time is modeled in the [DEE](#) as follows: each time a new event occurs in a particular detector and passes the trigger conditions, a timer begins in one available buffer slot. To update the timer, for every new event, the timer is incremented by the time difference between the current and previous event. Once the timer reaches $625\ \mu\text{s}$, the buffer slot is available again. If an event occurs when there are no available buffer slots, that event is discarded.

3.3 SIMULATION BENCHMARKING

To benchmark the simulations once all detector effects have been added, we compare calibration data to simulations at various stages of the analysis pipeline. Rather than looking at individual effects as we did throughout [Section 3.2](#), we compare the overall event rate and distributions such as spectra that include all events.

3.3.1 TRIGGER RATE COMPARISON

The trigger rate is the number of events per second that come out of each card cage. In the simulations, we can approximate the trigger rate as the number of events per second per detector after the [DEE](#) but before the event calibration. By comparing these rates, we can confirm that the shield and guard ring vetoes, trigger conditions, and dead time are all well approximated in the [DEE](#). [Figure 3.15](#) shows a comparison of the measured and simulated trigger rates for each detector as a function of calibration source and off-axis angle. The difference in trigger rate across the detectors is due to the detector placement in the cryostat: the top detectors collect more photons than the middle detectors, which collect slightly more photons than the bottom detectors. The residuals $(r_M - r_S)/\sigma_M$ in [Figure 3.15](#) are the difference between the measured and simulated trigger rate (r_M and r_S , respectively), expressed

3.3 SIMULATION BENCHMARKING

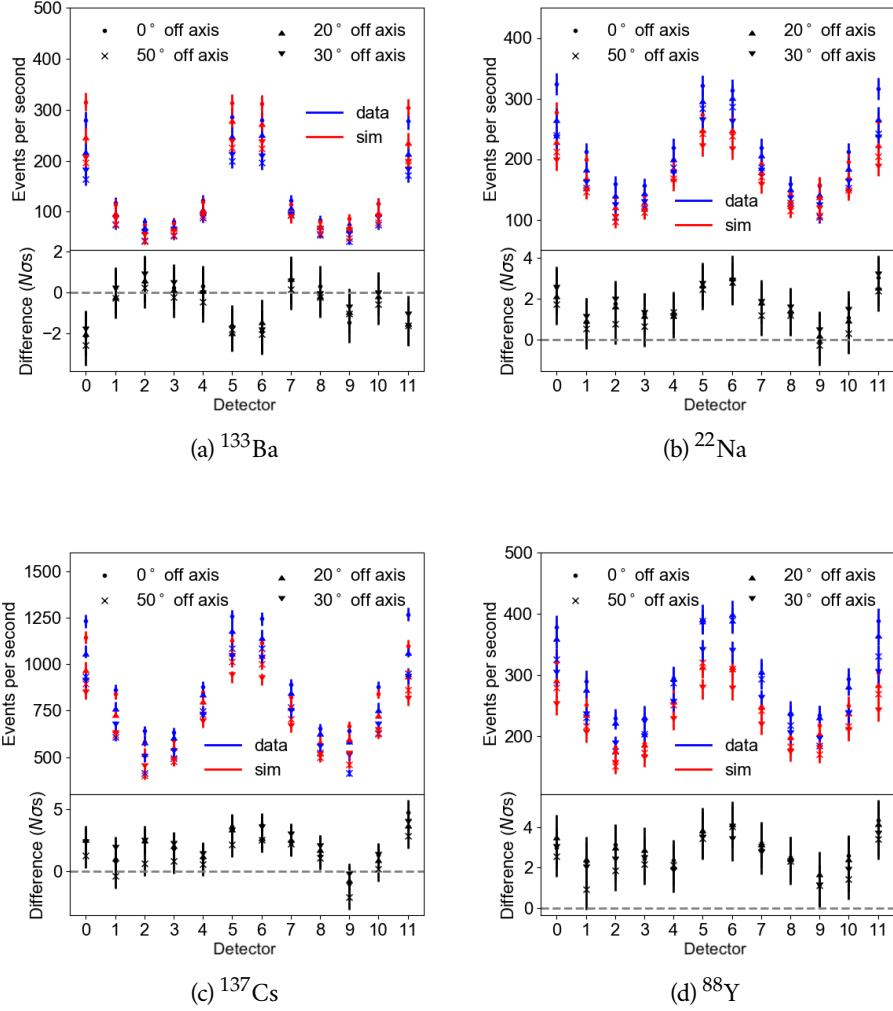


Figure 3.15: A comparison of the measured (blue) and simulated (red) trigger rate for a variety of calibration sources and off-axis angles. The trigger rates differ between the top, middle, and bottom layers of detectors, as detectors closer to the top collect more photons. The detector placement is as follows: top: 0, 5, 6, 11; middle: 1, 4, 7, 10; bottom: 2, 3, 8, 9.

3.3 SIMULATION BENCHMARKING




Figure 3.15: *Continued* – A comparison of the measured (blue) and simulated (red) trigger rate for a variety of calibration sources and off-axis angles. The trigger rates differ between the top, middle, and bottom layers of detectors, as detectors closer to the top collect more photons. The detector placement is as follows: top: 0, 5, 6, 11; middle: 1, 4, 7, 10; bottom: 2, 3, 8, 9.

in the number of standard deviations σ_M , or the statistical error on r_M . The relatively small residuals indicate good agreement between the measured and simulated trigger rates.

3.3.2 SPECTRUM COMPARISON

Figure 3.16 shows a comparison of the measured and simulated spectra of a ^{133}Ba and a ^{137}Cs calibration source. The spectra were made after the event calibration and event reconstruction steps of the analysis pipeline, and therefore include the removal of flagged events from the event calibration. The measured and simulated spectral continuum and line shapes match very well, illustrating that the DEE is accurately simulating the detector performance. However, there is a discrepancy in the peak height, but not in the continuum, indicating that too many simulated fully absorbed events pass through the pipeline. This discrepancy exists for all calibration sources, as shown in Figure 3.17, which plots the ratio of counts in the measured line to counts in the simulated line as a function of energy and off-axis angle for many calibration sources. The off-axis angle has a weak effect on the differences between mea-




Figure 3.16: A comparison of the measured and simulated spectrum of (a) a ^{133}Ba source and (b) a ^{137}Cs source. The line shapes match very well, but the number of counts in the lines differs.

measurements and simulations. To correct for this discrepancy, a certain fraction of fully absorbed events – one minus the ratio shown in Figure 3.17 – are removed in the DEE. The ratio is linearly interpolated across the energy range. For each fully absorbed event in the DEE, a random number between zero and one is drawn; if that number is greater than the ratio of measured to simulated line counts at the energy of the event, the event is discarded.

Figure 3.18 shows a comparison of the measured and simulated spectra of two calibration sources after removing a fraction of the fully absorbed events, as described above, and the peak heights match very well. Figure 3.19 shows an updated plot of the ratio of counts in the measured line to counts in the simulated line as a function of energy and off-axis angle; the discrepancy of peak heights is now minimized.

To compare the shape of the measured and simulated spectra, we fit each calibration source line with a Gaussian function. The results are shown in Table 3.4, and the full width at half maximum (FWHM) of each line as a function of energy and off-axis angle is plotted in Figure 3.20. The relatively large difference between the measured and simulated FWHM at 511 keV is due to the fact that the simulations assume that the electron-positron pair that annihilates to produce the 511 keV γ -rays is at rest, when in reality the electron

3.3 SIMULATION BENCHMARKING




Figure 3.17: The ratio of measured to simulated counts in the lines as a function of energy and off-axis angle.




Figure 3.18: A comparison of the measured and simulated spectra of (a) a ^{133}Ba calibration source and (b) a ^{137}Cs calibration source, after some of the fully absorbed events have been removed. The line shapes and heights match very well.




Figure 3.19: A comparison of the ratio of measured to simulated counts in the lines as a function of energy and off-axis angle, after some of the fully absorbed events are removed from the simulations.

and positron have some momentum. This nonzero momentum results in the broadening of the line.

To quantitatively assess the improvements we have made with the COSI-specific DEE, we compare the measurements to the simulations as follows: for each bin in the spectrum, we compute the difference in number of sigmas $z = (N_M - N_S)/\sigma_M$, where N_M and N_S are number of measured and simulated counts in the bin and σ_M is the error on N_M (with photon counting statistics, $\sigma_M = \sqrt{N_M}$). We then make a histogram of the z values for each bin in the spectrum and fit the resulting distribution with a Gaussian function, as in Figure 3.21. If the measurements and simulations are in good agreement with each other, the mean of the Gaussian fit will be close to zero and the FWHM of the Gaussian fit will be small. Note that we include the entire spectral continuum in this analysis as well as the lines; it is important for both to match well.

We performed this analysis for two cases: first, the simulations were processed with the standard MEGALib DEE, and second, the simulations were processed with the COSI-specific DEE. Figure 3.21 compares the distributions of z values for these two cases, using a calibration run of a ^{137}Cs source directly on-axis. The distribution is four times narrower and the mean of the distribution is five standard deviations closer to zero when we use the COSI-specific DEE compared to when we use the standard MEGALib DEE. Table 3.5 compares the Gaussian fit parameters for a number of different spectra as a function of calibration source and off-axis angle. On average, the FWHM is 1.9 times

3.3 SIMULATION BENCHMARKING

Energy Source	Off-axis angle (°)	Measurement		Simulation	
		Centroid (keV)	FWHM (keV)	Centroid (keV)	FWHM (keV)
356.01 ^{133}Ba	0	355.36 ± 0.02	4.43 ± 0.05	355.35 ± 0.02	4.39 ± 0.05
	20	355.37 ± 0.02	4.28 ± 0.05	355.36 ± 0.02	4.29 ± 0.05
	30	355.34 ± 0.03	4.34 ± 0.06	355.35 ± 0.02	4.33 ± 0.05
	50	355.37 ± 0.02	4.17 ± 0.06	355.38 ± 0.02	4.16 ± 0.05
511.0 ^{22}Na	0	510.30 ± 0.02	5.56 ± 0.04	510.38 ± 0.02	4.61 ± 0.04
	20	510.29 ± 0.02	5.82 ± 0.05	510.38 ± 0.02	4.74 ± 0.04
	30	510.21 ± 0.02	5.44 ± 0.05	510.40 ± 0.02	4.50 ± 0.05
	50	510.15 ± 0.03	5.62 ± 0.06	510.38 ± 0.02	4.59 ± 0.05
661.66 ^{137}Cs	0	660.87 ± 0.01	5.10 ± 0.02	661.02 ± 0.01	4.95 ± 0.02
	20	660.66 ± 0.01	5.17 ± 0.02	661.02 ± 0.01	4.89 ± 0.02
	30	660.69 ± 0.01	5.21 ± 0.02	661.04 ± 0.01	4.88 ± 0.02
	50	660.78 ± 0.01	5.40 ± 0.03	661.03 ± 0.01	5.04 ± 0.02
898.04 ^{88}Y	0	897.17 ± 0.02	5.63 ± 0.07	897.37 ± 0.03	5.53 ± 0.07
	20	897.07 ± 0.03	5.55 ± 0.07	897.44 ± 0.03	5.54 ± 0.07
	30	897.08 ± 0.03	5.50 ± 0.07	897.34 ± 0.03	5.88 ± 0.08
	50	896.97 ± 0.03	5.67 ± 0.08	897.44 ± 0.03	5.58 ± 0.08
1173.23 ^{60}Co	0	1172.15 ± 0.02	7.36 ± 0.05	1172.57 ± 0.01	6.85 ± 0.05
	20	1172.08 ± 0.02	6.89 ± 0.05	1172.56 ± 0.02	6.99 ± 0.05
	30	1172.16 ± 0.02	6.86 ± 0.05	1172.58 ± 0.02	6.88 ± 0.06
	50	1172.24 ± 0.02	6.74 ± 0.05	1172.63 ± 0.02	6.94 ± 0.06
1274.54 ^{22}Na	0	1273.62 ± 0.04	6.42 ± 0.10	1273.83 ± 0.04	6.13 ± 0.11
	20	1273.26 ± 0.05	6.92 ± 0.13	1274.00 ± 0.04	6.28 ± 0.12
	30	1273.31 ± 0.05	6.73 ± 0.13	1273.90 ± 0.05	6.16 ± 0.13
	50	1273.22 ± 0.05	6.44 ± 0.11	1273.98 ± 0.04	6.06 ± 0.13
1332.49 ^{60}Co	0	1331.43 ± 0.02	6.95 ± 0.05	1331.89 ± 0.02	6.45 ± 0.04
	20	1331.37 ± 0.02	6.50 ± 0.04	1331.85 ± 0.02	6.37 ± 0.05
	30	1331.43 ± 0.02	6.46 ± 0.05	1331.88 ± 0.02	6.39 ± 0.05
	50	1331.50 ± 0.02	6.27 ± 0.05	1331.95 ± 0.02	6.34 ± 0.05
1836.05 ^{88}Y	0	1834.78 ± 0.05	8.08 ± 0.12	1835.61 ± 0.05	7.83 ± 0.15
	20	1834.42 ± 0.05	8.10 ± 0.13	1835.49 ± 0.04	7.23 ± 0.13
	30	1834.40 ± 0.05	7.98 ± 0.14	1835.57 ± 0.05	7.15 ± 0.14
	50	1834.48 ± 0.06	8.14 ± 0.15	1835.64 ± 0.05	7.73 ± 0.15

Table 3.4: A comparison of the line widths and centroids of a Gaussian fit of each calibration source line for measurements and simulations.




Figure 3.20: A comparison of the measured and simulated [FWHM](#) of each calibration source line as a function of energy and off-axis angle. The relatively large difference between the measured and simulated [FWHM](#) at 511 keV is due to the assumption in the simulations that the electron-positron pair that produces the 511 keV γ -rays is at rest, when in reality the electron and positron have some momentum; this non-zero momentum leads to the broadening of the line.

smaller and the mean is 1.5 standard deviations closer to zero when we use the [COSI-specific DEE](#), indicating that we have achieved better agreement between measurements and simulations.

When analyzing [COSI](#) data of astrophysical sources, we compute the instrument response from simulations. Thus it is important to determine the systematic error that comes from the remaining discrepancies between measurements and simulations that we must apply to quantities such as the measured flux. To do so, for each calibration source line energy, we calculate the integral of the measured and simulated spectra between $E_{\text{centroid}} - 2\sigma$ and $E_{\text{centroid}} + 2\sigma$, where E_{centroid} is the center and σ is the width of the Gaussian fit to the line (see Table 3.4 for E_{centroid} and σ values). We then determine the systematic error required such that the number of measured counts N_M in the line is consistent with the number of simulated counts N_S in the line. The results are shown in Table 3.6, where the error bars on the systematic errors are the statistical errors of the number of measured counts, or $\sqrt{N_M}$. We performed this analysis for four off-axis angles in [COSI](#)'s field of view (FOV) (0° , 20° , 30° , and 50°), and the results shown in Table 3.6 are the results averaged over the four off-axis angles.

3.3 SIMULATION BENCHMARKING

Source <i>Energies (keV)</i>	Off-axis angle(°)	Standard MEGAlib DEE		COSI-specific DEE	
		Mean	FWHM	Mean	FWHM
¹³³ Ba <i>276.04, 302.85</i> <i>356.01, 383.83</i>	0	-3.7	18.9	-1.1	7.17
	20	-3.1	17.3	-0.38	6.82
	30	-1.8	16.5	-0.58	5.69
	50	-3.1	11.4	-0.43	5.82
²² Na <i>511.00</i> <i>1274.54</i>	0	1.1	3.49	0.29	4.42
	20	0.53	4.32	0.23	4.20
	30	0.8	3.88	0.26	4.02
	50	0.52	4.33	0.33	3.83
¹³⁷ Cs <i>661.66</i>	0	-5.2	23.2	-1.8	5.97
	20	-5.3	20.3	-1.8	5.20
	30	-4.7	16.7	-1.5	5.43
	50	-5.9	15.1	-2.0	5.15
⁸⁸ Y <i>898.04</i> <i>1836.05</i>	0	0.93	-3.57	0.47	3.59
	20	0.86	-3.76	0.34	3.56
	30	0.71	-3.87	0.23	3.57
	50	0.43	3.98	0.09	3.42
⁶⁰ Co <i>1173.23</i> <i>1332.49</i>	0	-1.7	8.37	-0.4	5.38
	20	-1.0	7.24	-0.36	5.20
	30	-0.74	6.61	-0.29	4.80
	50	-2.2	6.96	-0.72	4.99

Table 3.5: A comparison of how well the measured and simulated spectra match for a variety of calibration sources and off-axis angles. We compare the measurements to simulations processed with the standard [MEGAlib DEE](#) and to simulations processed with the [COSI-specific DEE](#). We show the mean and [FWHM](#) of the Gaussian fit to the distribution of the difference in number of sigmas $z = (N_M - N_S)/\sigma_M$, where z is computed for each spectral bin.

3.3 SIMULATION BENCHMARKING




Figure 3.21: The distribution of $z = (N_M - N_S)/\sigma_M$ when (a) the simulations are processed with the standard `MEGALIB DEE` and (b) the simulations are processed with the `COSI-specific DEE`. The match between measurements and simulations is much better when the simulations are processed with the `COSI-specific DEE`: the distribution is about four times narrower and centered near zero.

Energy (keV)	Source	Systematic error (%)
356.01	^{133}Ba	$+20.2 \pm 0.3$
511.00	^{22}Na	$+1.6 \pm 0.2$
661.66	^{137}Cs	$+4.4 \pm 0.1$
898.04	^{88}Y	-5.3 ± 0.3
1173.23	^{60}Co	$+6.3 \pm 0.1$
1274.54	^{22}Na	-4.8 ± 0.4
1332.49	^{60}Co	$+7.2 \pm 0.2$
1836.05	^{88}Y	-0.3 ± 0.4

Table 3.6: The systematic error that must be applied to the measured flux, due to the remaining discrepancies between measurements and simulations.




Figure 3.22: A comparison of the measured and simulated [ARM](#) distributions of the 662 keV line of a ^{137}Cs source.

3.3.3 ANGULAR RESOLUTION MEASURE COMPARISON

The angular resolution measure ([ARM](#)) defined in Section 2.1.2 is used to characterize the spatial resolving capabilities of a Compton telescope and is a useful benchmarking tool, as it depends on the energy and position resolution and the results of the event reconstruction. The [FWHM](#) of the [ARM](#) defines the angular resolution of the instrument. Figure 3.22 shows a comparison of the measured and simulated [ARM](#) distributions for the 662 keV line of a ^{137}Cs source.

Table 3.7 compares the Gaussian fit parameters of the distributions of z values computed by comparing the measured and simulated [ARM](#) distribution. Again, we compare the simulations processed with the standard [MEGALib DEE](#) to the simulations processed with the [COSI-specific DEE](#). Since the shape of the [ARM](#) distribution is dependent on energy, we make a separate [ARM](#) distribution for each calibration source line energy. On average, the [FWHM](#) of the Gaussian fits to the z distribution is 1.8 times smaller and the mean is 1.6 standard deviations closer to zero when the simulations are processed with the [COSI-specific DEE](#), indicating that we have achieved better agreement between measurements and simulations in the [ARM](#) distributions as well as in the spectra.

Figure 3.23 compares the measured and simulated [ARM FWHM](#) as a function of energy and off-axis angle. The measured and simulated [ARM FWHM](#) values are fairly close with an average residual of 3σ , but better agreement

3.3 SIMULATION BENCHMARKING

Energy (keV)	Off-axis	Standard MEGAlib DEE		COSI-specific DEE	
Source	angle (°)	Mean	FWHM	Mean	FWHM
Source		Mean	FWHM	Mean	FWHM
356.01 ^{133}Ba	0	1.9	16.1	0.26	4.46
	20	4.1	14.4	0.17	2.93
	30	0.20	5.09	0.57	4.33
	50	42.4	35.3	-0.09	7.09
511.0 ^{22}Na	0	-0.1	6.62	0.69	2.94
	20	-1.3	4.52	0.65	2.77
	30	-0.24	4.66	0.57	3.07
	50	-1.1	5.47	0.26	3.20
661.66 ^{137}Cs	0	-2.7	13.1	0.99	4.15
	20	-0.95	9.70	0.73	3.34
	30	-1.0	6.16	0.67	3.56
	50	-1.1	6.14	0.03	3.59
898.04 ^{88}Y	0	0.80	6.72	0.94	3.21
	20	0.56	3.12	0.79	2.94
	30	0.59	3.53	0.57	3.01
	50	0.54	3.13	0.82	2.59
1173.23 ^{60}Co	0	-1.2	9.22	-1.2	5.60
	20	-0.27	6.79	-0.53	3.98
	30	-0.35	-4.53	-0.84	3.52
	50	-0.41	3.54	-0.83	3.63
1274.54 ^{22}Na	0	0.81	4.04	0.69	2.62
	20	0.66	3.23	0.65	2.51
	30	0.80	4.46	0.90	3.90
	50	0.51	3.16	0.51	2.72
1332.49 ^{60}Co	0	-0.75	6.56	0.25	4.30
	20	-0.43	3.45	0.41	4.00
	30	-0.58	4.65	0.23	3.30
	50	-0.49	3.60	0.17	3.61
1836.05 ^{88}Y	0	0.31	4.90	0.70	3.08
	20	0.60	3.77	0.61	3.05
	30	0.53	3.11	0.65	2.39
	50	0.34	2.60	0.28	3.06

Table 3.7: A comparison of how well the measured and simulated [ARM](#) histograms match for a variety of calibration sources and off-axis angles. We compare the measurements to simulations processed with the standard [MEGAlib DEE](#) and to simulations processed with the [COSI-specific DEE](#). We show the mean and [FWHM](#) of the Gaussian fit to the distribution of the difference in number of sigmas $z = (N_M - N_S)/\sigma_M$, where z is computed for each bin in the [ARM](#) histogram.




Figure 3.23: A comparison of the measured and simulated [ARM FWHM](#) as a function of energy and off-axis angle.

is desirable. One of the major contributions to the differences between the measured and simulated [ARM FWHM](#) is the discrepancy in the distribution of distances between the first two interactions. Figure 3.24 shows the measured and simulated distribution of distances between the first two interactions at 511 keV and 1836 keV. At distances smaller than ~ 2.5 cm, the shape of the measured and simulated distributions differ significantly. At low energies, there are more measured events at small distances, and at high energies, the simulated distribution appears offset compared to the measured distribution. Selecting events with a distance between the first two interactions greater than 2.5 cm leads to better agreement between the measured and simulated [ARM FWHM](#) values, with an average residual of 1.9σ .

It is unclear why this discrepancy occurs in the distance distributions. We note that distances smaller than 2.5 cm are likely due to multiple hits in the same detector, whereas larger distances are likely due to multiple hits in different detectors. One possible explanation for the observed discrepancy is that the strip pairing algorithm groups what should be two simulated hits as one. This could cause a dearth of events with a small interaction distance between the first two hits. This issue could potentially be improved by more physical models of charge sharing and charge loss, which requires a detailed charge transport simulation. Additionally, the differences in the measured and simulated depth distributions of charge sharing hits (Figure 3.8) should add discrepancies to the overall distance distributions. Improvements to the

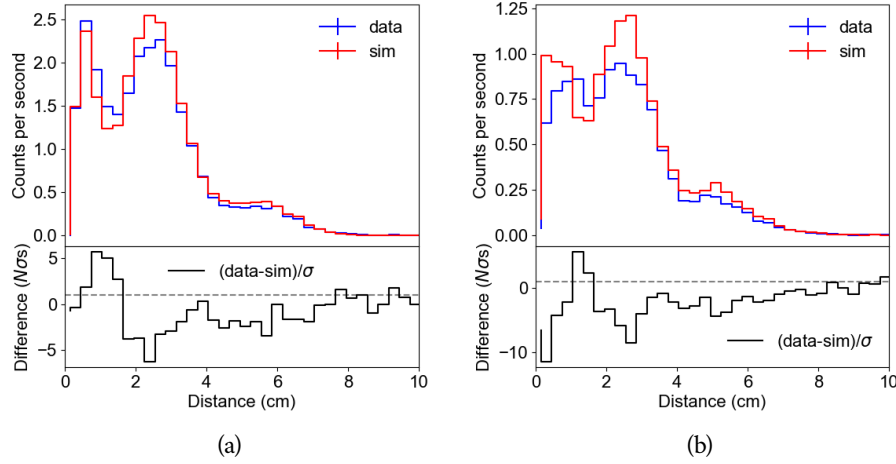


Figure 3.24: A comparison of the measured and simulated distribution of distances between the first and second interaction for (a) the 511 keV line of a ^{22}Na source and (b) the 1.836 MeV line of a ^{88}Y source. The discrepancies below 2.5 cm contribute to the differences between the measured and simulated the [ARM FWHM](#).

depth calibration algorithm could potentially resolve some of the discrepancies in the distance distributions.

Another contribution to the discrepancies in the measured and simulated [ARM FWHM](#) is the discrepancy in the initial Compton scatter angle distribution at low energies. Figure 3.25 compares the measured and simulated initial Compton scatter angle (ϕ in Figure 2.3) distribution at 356 keV and 1333 keV. Selecting events with Compton scatter angles above 40° and distances greater than 2.5 cm results in excellent agreement between the measured and simulated [ARM FWHM](#) values at 356 keV, with a residual of 0.95σ (the residual at 356 keV is 4.6σ with open selections and 2.6σ when selecting events with distances greater than 2.5 cm). The observed difference in Compton scatter angle distribution at low energies could potentially be related to the distance and depth discrepancies. Low energy events are more likely to only interact twice in the detector, and these two-site events are difficult to properly reconstruct. Any systematic error on the interaction distance is likely to translate to an incorrect Compton scatter angle when reconstructing two-site events, which could explain the observed discrepancies in Compton scatter angle dis-

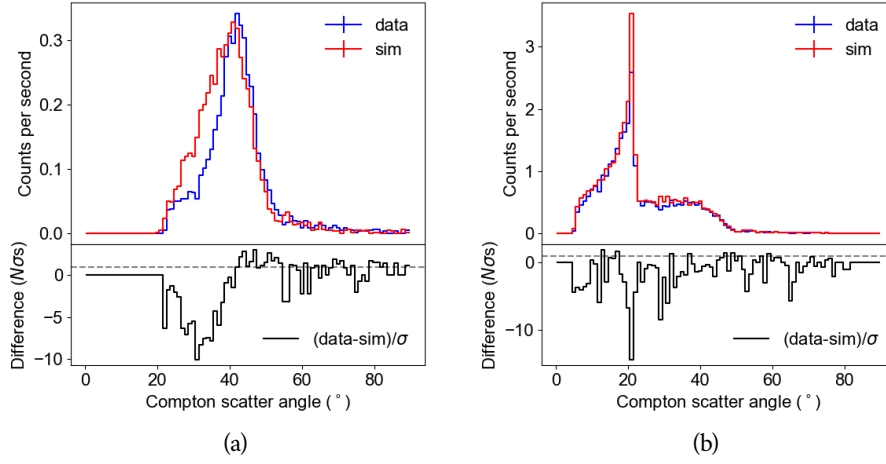


Figure 3.25: A comparison of the measured and simulated distribution of initial Compton scatter angles for (a) the 356 keV line of a ^{133}Ba source and (b) the 1.3 MeV line of a ^{60}Co source. These distributions agree better at high energies than at low energies. The measured and simulated [ARM FWHM](#) are in better agreement at low energies when only events with Compton scatter angles $> 40^\circ$ are used.

tributions. However, it is unclear why smaller Compton scatter angles are affected more than larger Compton scatter angles.

3.4 CONCLUSIONS

The [COSI](#)-specific [DEE](#) described in this chapter is an important addition to the [COSI](#) analysis pipeline. The instrument response for imaging, spectral, and polarization analysis is computed from simulations, and so it is imperative that the simulations closely resemble the measured [COSI](#) data. Careful modeling of the effects in the [COSI](#) detectors and readout electronics brings us significantly closer to this goal. We have compared the simulations to calibrations at various stages throughout the analysis pipeline to ensure that they match well. Our comparisons of the spectra and [ARM](#) after the event reconstruction indicate very good agreement between measurements and simulations in comparison to the standard [MEGALib DEE](#). We have determined the systematic error that must be applied to measured fluxes due to remaining discrepancies between measurements and simulations. The systematic error

for almost every line is less than 10%, and the average of the magnitude of the systematic error is 6.3%. Better agreement could potentially be achieved with more physical models of the detector effects, which requires a detailed charge transport simulation. Work on a detailed charge transport simulation is currently ongoing by other members of the [COSI](#) team.

4

THE 2016 COSI BALLOON CAMPAIGN

COSI was launched from Wanaka, New Zealand in May 2016 as the first science payload to fly on NASA’s super pressure balloon (SPB). The 46-day flight yielded a large data set of several astrophysical γ -ray sources, including the Crab Nebula and GRB 160530A which are discussed in Part III. In this chapter, I will provide a summary of the 2016 flight and describe the flight operations and in-flight instrument performance.

4.1 HISTORICAL CONTEXT

The overall success of the COSI instrument during the 2016 flight can in part be attributed to its precursor instrument, the Nuclear Compton Telescope (NCT). NCT had two successful flights from Fort Sumner, New Mexico, in 2005 and 2009. In 2005, the team flew a prototype version of the instrument with two germanium detectors (GeDs) and measured the atmospheric and instrumental γ -ray background, which matched closely with predictions (Bowen et al., 2006; Bowen, 2009). In 2009, the flight of the full NCT instrument with all 10 GeDs resulted in the first detection and image of an astrophysical source, the Crab Nebula, with a compact Compton telescope (Bandstra et al., 2011).

In 2010, NCT was scheduled to fly from Alice Springs, Australia, but a launch mishap completely destroyed the gondola; however, the detectors and readout electronics remained relatively unharmed. The team used the partial destruction of NCT as an opportunity to redesign the instrument for ultra-long duration balloon (ULDB) flights on NASA’s 18 million ft³ super pressure balloon (SPB). The instrument was renamed the Compton Spectrometer and Imager (COSI).

The main difference between the SPB and the conventional zero pressure balloon (ZPB) is that the SPB is completely closed whereas the ZPB has an opening on the bottom; the completely enclosed volume is what allows the SPB to

fly for longer than the ZPB. ZPB flights are limited in duration because the helium inside the balloon slowly leaks out of the opening. The helium inside the balloon is heated from the sun during the day, causing the volume of the balloon to expand and the balloon to rise, as the buoyancy force is proportional to the volume of the balloon. At night, the helium cools down, which causes the reverse process. The day-night altitude cycles result in a net loss of helium, and thus an overall altitude decrease throughout the flight. The SPB is a closed system that is slightly pressurized compared to the ambient atmosphere. The balloon material is rigid enough to maintain the same volume despite changes in helium temperature, and so the altitude remains constant. As no helium can escape, the SPB can fly until the material breaks down. To date, the SPB flight record is 54 days, but NASA is continuing to push this new balloon technology as long constant-altitude flights are desirable to science teams.

ULDB flights on the SPB require a relatively lighter payload, no consumables, and the ability to telemeter down all data in case the payload lands in the ocean (ZPBs generally only fly over land). These requirements guided most of the changes from NCT to COSI. In particular, a new, lighter gondola was designed for COSI, and the COSI cryostat is cooled by a mechanical cryocooler rather than a liquid nitrogen dewar so that the flight duration is not limited by the amount of liquid nitrogen on board. The new COSI flight software includes on-board data processing to select the more useful events to telemeter down in real time. The instrument was also simplified to run as easily and autonomously as possible: NCT's pointing system was replaced with a zenith-centered field of view (FOV) on COSI, and sources are observed as they pass through the FOV. Additionally, the efficiency of the instrument was improved by adding two more GeDs.

In 2014, COSI had its first flight from McMurdo Station, Antarctica. COSI was the first science payload to fly on the SPB. COSI launched on December 28, 2014 and the flight was expected to last several weeks. Unfortunately, a leak in the balloon caused NASA to terminate the flight after ~40 hours. The payload was recovered within a few weeks and was not damaged, allowing for the team to quickly get ready for a flight in Spring 2016.

Although the 2014 flight was disappointingly short and no astrophysical objects were detected during the flight, the team realized that some aspects of the project needed to be improved. First, the cryocooler did not effectively dissipate heat at float altitude, which led the cryocooler to overheat and made it difficult to cool the GeDs. Because of this problem, we developed the liquid

cooling system described in Section 2.2.2 for the 2016 flight. Second, it was apparent that our ground software was not sufficient for a long flight: the monitoring software crashed constantly and there was no alert system, so someone had to be watching the monitoring software at all times. I wrote entirely new ground software that ran consistently and allowed for more autonomous operations, described in Section 4.3.

4.2 2016 FLIGHT OVERVIEW

The team arrived in Wanaka, New Zealand for the 2016 balloon campaign in mid-February, and the instrument was flight ready by April 1. During those 6 weeks, we integrated and tested the instrument and then performed the calibrations described in Chapter 3. Meanwhile, NASA's Columbia Scientific Ballooning Facility (CSBF) team integrated their parts of the instrument, including the batteries, solar panels, and telemetry. In late March, we tested for compatibility with CSBF, in which we made sure that their systems did not affect our systems and vice versa. After the compatibility test, COSI was declared flight ready.

Low surface winds are required to launch the SPB. For six weeks, all launch attempts were aborted as the weather did not cooperate. Finally, COSI was launched on May 16, 2016 at 23:35 Coordinated Universal Time (UTC). Figure 4.1 shows COSI on the flight line with the SPB almost fully inflated on the launch day. The launch was successful and all systems performed flawlessly during launch and ascent. It took about two hours for the balloon to reach its float altitude of ~ 33.5 km.

The flight lasted a total of 46 days, in which the balloon circumnavigated the Earth one and one-half times before landing in Peru on July 2 at 19:54 UTC. The flight trajectory is shown in Figure 4.2. The wind patterns were highly variable throughout the flight. Initially, the balloon drifted west over Australia rather than the expected direction of east. After completing the first circumnavigation in about two weeks, the balloon drifted further north than was expected, reaching a northernmost latitude of 5.65° S. For the remainder of the flight, the balloon drifted back and forth over the Pacific Ocean with no clear timeline for reaching land. During this time period, the balloon began exhibiting altitude drops at night, as shown in Figure 4.3. The altitude drops effectively reduced the total observation time: most astrophysical γ -rays are attenuated when the payload is at low altitudes. The cause of these altitude drops is not definitively known, but the symptoms are consistent with a small



Figure 4.1: **COSI** on the flight line hanging from the launch vehicle and connected to the parachute and **SPB**. Only the top of the **SPB** is filled on the ground, but as the balloon reaches the float altitude of ~ 33.5 km, the helium expands and fills the entire balloon, resulting in a distinctive pumpkin-shaped balloon.

leak in the **SPB**, causing it to behave like a **ZPB**. When the balloon finally flew over land, the flight was immediately terminated.

The balloon and payload were located within a few days after the flight ended, and a small recovery team went to retrieve the instrument. A photograph of the payload after landing is shown in Figure 4.4. Though the gondola, antenna booms, and solar panels were damaged, the rest of the instrument, including the cryostat, **GeDs**, and the flight computer hard drives, was unharmed. Overall, the flight was a success: 46 days is long for a scientific balloon flight (the world record is currently 55 days), and the **COSI** 2016 flight marked the longest scientific balloon flight at-mid latitudes. Nine out of 12 **GeDs** performed flawlessly during the flight. Because we were able to recover the relatively unharmed instrument, we have the capability of flying it again and increasing our dataset, which we plan to do during Spring 2020.

4.3 FLIGHT OPERATIONS

Flight operations must be relatively autonomous, as it is impossible for the small operations team of four people to be actively monitoring the instrument health at all times during a **ULDB** flight. The operations ground software primarily consists of a real time monitoring program and an alert system, described in more detail in this section.

During the first day of flight, the instrument health was most uncertain, but the payload was also within the line of sight (**LOS**). The **LOS** telemetry stream is fast and reliable as long as the payload is within range. The team




Figure 4.2: *COSI*'s flight path during the 2016 flight. The green path shows the first circumnavigation, and the red path shows the second partial circumnavigation. The flight began in Wanaka, New Zealand and ended in southern Peru 46 days later. The balloon drifted further north than expected, allowing for some exposure of the northern skies.




Figure 4.3: Altitude profile of *COSI* during the 2016 flight. The launch is shown on May 17. The drops in altitude early in the flight can be attributed to cold storms. The drops in altitude toward the middle and end of the flight are due to an issue with the balloon.



Figure 4.4: A photograph of the [COSI](#) payload in at its landing site in Peru, taken by the recovery team. Though the gondola, antenna booms, and solar panels were completely destroyed, the rest of the instrument was unharmed. The [GeDs](#), [CsI](#) shields, card cages, and other subsystems have been tested in the lab since the flight and are fully operational.

actively monitored the instrument for the first 24 hours, taking advantage of the [LOS](#) telemetry stream. This also gave us an opportunity to check our alert system. We then switched to checking on the payload every hour or so for another 24 hour period. By this point, everything seemed stable and it was clear that the alert system worked well. We then began the shift schedule that lasted until the end of the flight: each team member was on shift for a 24-hour period during which she or he was responsible for responding to any alerts that arose. For the most part, the instrument was able to run autonomously, and it was not uncommon for days to go by without a single instrument alert.

4.3.1 GSE

The primary ground software, referred to as the ground support equipment ([GSE](#)), provides comprehensive monitoring of instrument health, commanding capabilities, and minimal low level analysis. The [GSE](#) was designed to run on Linux and Mac OS, but so far runs successfully on every operating system we have tried, including Windows and Raspian. The [GSE](#) can connect to any of the data streams ([LOS](#), [OP](#), etc.), and can switch between data streams while the program is running. The bulk of the [GSE](#) is written in Python and uses the Qt framework as well as `pyqtgraph` to display the data. For additional processing speed, all of the parsing is written in C and connects to the Python

frontend using the Python module `ctypes`. The code is modular and it is relatively straightforward for users to add additional displays.

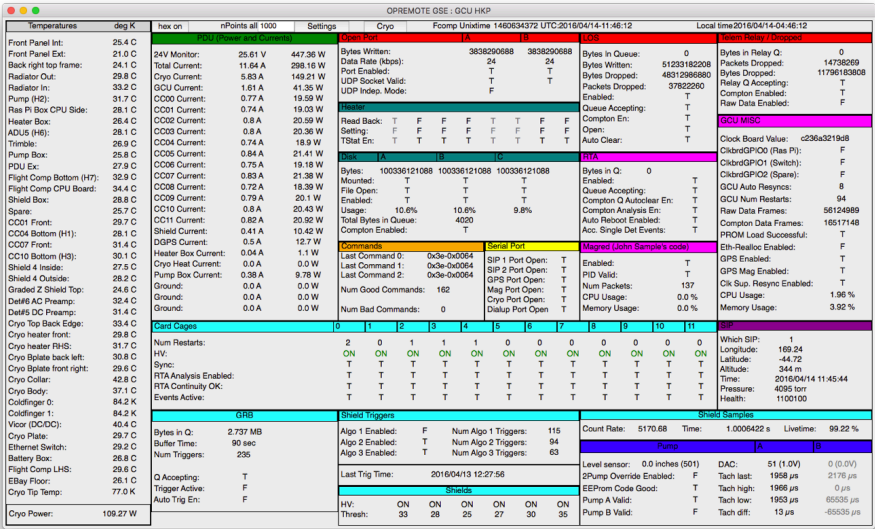
The [GSE](#) contains a number of displays, or widgets, that give insight into instrument health and basic performance. For example, the flight computer housekeeping (Figure 4.5a) shows the temperatures of the subsystems, the current draw of each subsystem, telemetry information, flight computer hard drive status, and other general housekeeping information. The user can click each value displayed in the [GSE](#) to see a strip chart showing the recent history (Figure 4.5b), and the length of the displayed history can be changed by the user. The spectral display (Figure 4.5d) shows the spectrum of each strip and contains a basic line fitting algorithm. With this widget, we can get a baseline for the instrument performance in real time. The [GSE](#) also provides the interface for sending commands to the instrument (Figure 4.5c). Commands can be selected by a drop-down menu. When a specific command is selected, instructions pertaining to that command appear on the screen.

It is important that the [GSE](#) software can remain running for long periods of time. The [LOS](#) telemetry stream in particular sometimes sends packets where one or more bits are flipped. These garbled packets have the potential to crash the [GSE](#) software, particularly in the parsing functions. To test the [GSE](#), we used a data playback program that reads old data from a file and sends it to the [GSE](#) as if it were a live telemetry stream. We added a function to flip one random bit in every single packet and made sure that the [GSE](#) could recognize a bad packet and continue running. Due to this work, the [GSE](#) is able to run for days or weeks at a time without any issues. We also used the data playback program to determine the maximum data rate that the [GSE](#) could handle, which is 10 Mbps, more than enough for [COSI](#) calibrations and flight data. Due to the modular structure of the [GSE](#) code, we were able to use the parsing backend for other ground software, described in the next few sections.


4.3.2 ALERT SOFTWARE

The alert software, called [COSI](#) monitor, texts the operations team if any values go out of specification, particularly if any subsystem gets too hot or too cold (see Figure 4.6). With an ultra-long duration balloon flight and a small operations team of four people, an alert system like [COSI](#) monitor is crucial for autonomous operations. [COSI](#) monitor also texts the operations team when the instrument receives a gamma-ray burst ([GRB](#)) trigger from the shields so that we can respond immediately if [COSI](#) detected a real [GRB](#).

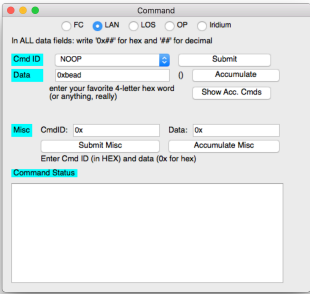
4.3 FLIGHT OPERATIONS



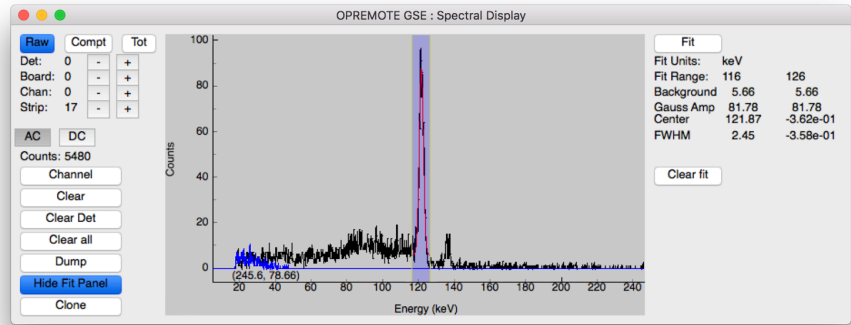
(a) Flight computer housekeeping widget: clicking on each value brings up a history



(b) History of the cryocooler collar temperature

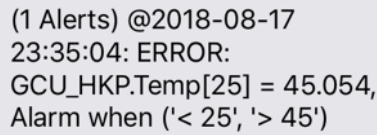


(c) Commanding widget



(d) Spectral display with basic line fitting, showing the spectrum of a ⁵⁷Co source

Figure 4.5: A sampling of the GSE widgets.



```
(1 Alerts) @2018-08-17
23:35:04: ERROR:
GCU_HKP.Temp[25] = 45.054,
Alarm when ('< 25', '> 45')
```

Figure 4.6: An example COSI monitor text message, alerting us that the cryocooler collar temperature (temperature # 25) is outside of the acceptable range of 25° to 45°.

4.3.3 COSI PLOTTER

The last piece of the ground software, called COSI plotter, allows the operations team to quickly look at any housekeeping value during any period of time during the flight. The GSE does retain histories for each housekeeping value, but only for a few hours. With COSI plotter, the user can select a time period and plot any value found in the GSE, and thus have access to GSE-like strip charts but from farther back in time than the last few hours. The back-end of COSI plotter connects to the data stream in real time and uses the GSE parser to parse the housekeeping packets into ASCII files. The ASCII files are then moved to a website, from where the COSI plotter frontend downloads the files from the requested time period onto the user's personal computer and plots the data.

4.4 THERMAL ENVIRONMENT

Payloads on mid-latitude ULDB trajectories must be able to withstand a large range of temperatures, as the ambient temperature at these latitudes can range from -100°C to 80°C. COSI was equipped to handle these large temperature variations: aluminum side panels covered with foam insulated the electronics bay temperature. Figure 4.7 shows the temperature of two subsystems within the electronics bay and the temperature of the outside of one of the side panels, demonstrating how well the electronics bay was insulated. Additionally, kapton heaters were attached to many subsystems. A heater box toggled the heaters on and off according to the temperature of the subsystem in question, measured by LM335 temperature sensors. During the flight, all subsystems stayed within their required temperature range.

The cryocooler thermal performance was also of some concern; however, the liquid cooling system (described in Section 2.2.2) worked extremely well




Figure 4.7: The temperatures of two subsystems within the electronics bay, card cage # 7 and the flight computer CPU board, and the exterior of one of the side panels during one 24-hour period. The large difference in day- and night-time ambient temperatures is apparent by the large swing in the side panel temperature, yet the temperatures of the subsystems within the electronics bay remain relatively constant. A heater was attached to the flight computer CPU board and was automatically toggled on and off between 12:00 and 16:00 UTC.

during the flight. In fact, the liquid cooling system was so effective that the flight computer had to turn the pump on and off throughout the flight to prevent the cryocooler from getting too cold according to the performance specification. Figure 4.8 shows the temperature of the cryocooler and radiator plate during one day of flight; the oscillating cryocooler temperature between 30°C and 40°C is a result of the pump turning on and off. The cryocooler temperature stayed well within operable range during the entire flight.

4.5 HIGH VOLTAGE FAILURES

During the flight, three of the 12 detectors failed due to problems with the high voltage. Detectors 8 and 5 failed within the first 20 and 42 hours of flight, respectively. They exhibited similar symptoms: the preamp currents were higher than normal and then began oscillating quickly between the nominal and high values (see Figure 4.9). The preamp currents normally only deviate from nominal values when the high voltage is being ramped up and down; however, during the flight, the high voltage supply monitor indicated that the high voltage was still on. The combination of these symptoms indicated that




Figure 4.8: The temperatures of the cryocooler and the sides of the radiator plate where the fluid is input and output during one 24-hour period. The difference in day- and night-time temperatures is apparent in the radiator plate temperatures. The oscillations in the cryocooler and radiator input temperature are due to the pump turning on and off, which occurred more frequently during the day when it was warmer.

high voltage breakdown was occurring somewhere between the high voltage supply and the detector. With the high preamp currents, the detectors are inoperable, so they were turned off shortly after these symptoms manifested. We tried turning the detectors back on a few more times throughout the flight, but the symptoms persisted.

Once the flight was over and the instrument was back in Berkeley, we did significant work to determine the underlying cause of the problem. First, the card cages (that contain the high voltage supply), the high voltage cables, and the cryostat itself were all tested on the lab bench, but the problem could not be reproduced. Next, each component in the high voltage path was separately placed in a vacuum chamber. The card cages and high voltage cables performed without issue in vacuum, but the problem was reproduced when the cryostat was placed in the vacuum chamber and the card cages with the high voltage supplies were outside the vacuum chamber. We saw the preamp currents oscillating between nominal and high values for both detectors, as we had seen during the flight. Changing the high voltage cable going from the supply to the cryostat did not resolve the problem. Thus, by process of elimination, the failure point must have originated from the cryostat.

The high voltage filters (shown in Figure 2.7) were deemed the most likely candidates for the failure point, since the cryostat itself is evacuated to a high vacuum of $\sim 10^{-6}$ Torr and high voltage breakdown is much less likely to oc-




Figure 4.9: The preamp currents and high voltage supply reading during the first failure of Detector 8 in flight. Until $\sim 18:55$, the preamp currents behaved normally. The failure occurred between $\sim 18:55$ and $19:20$, when the currents oscillated between nominal and high values (when the preamp currents are higher than normal, the detector is inoperable). We ramped the high voltage down at $\sim 19:20$, at which point the positive preamp current went high and the negative preamp current went low; this is normal behavior during high voltage ramp down.

cur at high vacuum. A leak in a high voltage filter, however, would not affect the cryostat pressure (and thus functionality) but would evacuate the filter to the ambient pressure of the vacuum chamber or of the atmosphere at balloon altitudes; high voltage breakdown is much more likely to occur at these pressures (~ 5 Torr). We replaced the high voltage filters on detectors 5 and 8 with spare filters that had been thoroughly tested in the vacuum chamber and did not show signs of breakdown. We then tested the original filters from Detectors 5 and 8 in the vacuum chamber and witnessed high voltage breakdown. From these tests, we have concluded that the problem from the flight was caused by faulty high voltage filters, which we have replaced.

Detector 0 also failed in flight, but with different symptoms, indicating a different cause. On June 6, the high voltage supply reading suddenly dropped to 0 V and the preamp currents responded as if we had ramped down the high voltage. These symptoms pointed to a problem with the high voltage supply. This problem was also reproduced in the lab after the flight by placing the card cage with the high voltage supply in a vacuum chamber. Further investigation of the high voltage supply led to the discovery of a design flaw: a critical resistor had exceeded its current rating, which prevented the supply

4.6 ASTROPHYSICAL SOURCES DETECTED




Figure 4.10: The exposure map from the entire flight, assuming an effective area of 20 cm^2 . The labeled sources are known γ -ray sources in COSI’s FOV but were not necessarily detected by COSI.

from working. The faulty resistor was replaced with a new resistor that has a higher current rating, and the supply now reliably works in vacuum.

Unfortunately, two of the three failed detectors (detectors 0 and 5) are in the top stack of the GeD array, in which the majority of events are detected. The loss of the failed detectors, especially the two top detectors, significantly lowered COSI’s efficiency: by simulating a calibration source emitting γ -rays at 511 keV, we determined that only 60% of events detected with all 12 detectors were detected without detectors 0, 5, and 8.

4.6 ASTROPHYSICAL SOURCES DETECTED

Despite the loss of three detectors, the flight was overall a success. The dataset has so far yielded detections of five astrophysical γ -ray sources, briefly described in this section. Figure 4.10 shows the exposure map from the entire flight. COSI had excellent exposure of the Galactic center region and thus the 511 keV positron annihilation signature. The analysis of the positron annihilation signal as detected by COSI is described in detail in Kierans (2018).

COSI also detected four compact γ -ray sources, the subject of this work. GRB 160530A is a long, bright GRB that we discovered with COSI (Tomsick & the COSI team, 2016); for details on the observation and analysis of GRB 160530A, see Lowell et al. (2017a), Lowell (2017), and Section 5.3 of this work. COSI had excellent exposure of and detected Cen A, a bright, well-known active galactic nucleus (AGN). Additionally, COSI detected the Crab pulsar wind nebula (PWN) and the Galactic black hole binary Cyg X-1, despite limited exposure of these two sources. The methods described in Part III are ideally designed for the analysis of all of these sources, and in Chapter 7, we apply these methods to COSI’s observation of the Crab.

PART III

ASTROPHYSICS ANALYSIS

5

SPECTRAL ANALYSIS WITH COSI

Fitting the observed spectra of astrophysical sources can tell us about the γ -ray emission mechanisms and thus help us further understand the physical processes occurring in and around the source in question. Additionally, as many other γ -ray instruments are spectrometers, comparing the COSI spectrum of a source to the spectrum of that same source measured by another instrument can help us confirm COSI's spectral capabilities.

We have developed a pipeline to perform spectral fitting of COSI data using XSPEC (Arnaud, 1996), a well known spectral fitting program in high energy astrophysics. As required for any spectral fitting, we have developed a spectral response matrix using our well-benchmarked simulations described in Chapter 3. As COSI is a balloon-borne instrument, there is still a small amount of atmosphere between COSI and the astrophysical sources in question; we have created a model of the atmospheric absorption to use in the spectral fits. In this chapter, I will describe the COSI spectral fitting pipeline, including the response matrix, atmospheric absorption model, and the challenges of background subtraction.

5.1 SPECTRAL FITTING AND THE RESPONSE MATRIX

A real spectrometer does not directly measure the spectrum emitted by the source; due to detector effects, the measured spectrum is somewhat distorted. For example, when measuring the spectrum of a monoenergetic source, the resulting line is broadened due to the finite energy resolution of the detector. Incompletely absorbed events result in a continuum component to the spectrum at energies below the photopeak. In general, the efficiency of a detector is energy dependent, so the measured fluxes at different energies will not represent the emitted flux as a function of energy.

The spectral response matrix describes the measured energy as a function of the incident energy of the photons when they reach the detector. The spec-

tral response matrix $R(I, E)$ relates the measured spectrum to the emitted spectrum as follows:

$$C(I) = \int f(E)R(I, E)dE \quad (5.1)$$

where $C(I)$ is the counts in measured energy bin I and $f(E)$ describes the emitted spectrum of the source as a function of energy E . For an observed counts distribution $C(I)$, it is difficult to invert Equation 5.1 to find a unique $f(E)$. Instead, we can choose a model that depends on a number of parameters $f(E, p_1, p_2, p_3, \dots)$. Using Equation 5.1, we can determine what counts distribution the instrument would measure, $C_{p_i}(I)$, assuming that the source emitted a spectrum in the shape of $f(E, p_1, p_2, p_3, \dots)$ for a given set of parameter values p_i . By varying p_i and comparing the resulting $C_{p_i}(I)$ to the observed $C(I)$, we find the parameter values and model that can best describe the emitted spectrum of the source.

The most common method to determine the best fit $C_{p_i}(I)$, which we use for all spectral fits in this work, is the χ^2 statistic:

$$\chi^2 = \sum \frac{(C(I) - C_{p_i}(I))^2}{\sigma(I)^2} \quad (5.2)$$

where $\sigma(I)$ is the error on the count rate in measured energy bin I . Generally, the model is a good fit to the data when the reduced χ^2 is approximately equal to one (Arnaud et al., 2018). The reduced χ^2 is defined as χ^2/dof , where “dof” stands for the number of degrees of freedom which is equal to the number of measured energy bins minus the number of model parameters.

There are multiple software packages to perform spectral fitting as described here. For spectral fitting with COSI, we use XSPEC, a commonly used program in high energy astrophysics.

5.1.1 COSI DATA WITH XSPEC

We generate the spectral response matrix for COSI using simulations, where we know the incident and measured energy of each photon. XSPEC requires that the spectral response matrix $R(I, E)$ be split into two components: the redistribution matrix $R_d(I, E)$ and the effective area $A_{\text{eff}}(E)$. The redistribution matrix describes how photons emitted at energy E are measured in the detector at energy I . The effective area, defined as the area of an ideal detector that detects the same number of photons as the real detector, describes




Figure 5.1: The (a) redistribution matrix and (b) effective area of a simulated on-axis point source emitting a power law spectrum with a photon index of 2 detected by [COSI](#). The event selections used for this response matrix are shown in Table 5.1. Note that the measured energy in the redistribution matrix is cut off after 3 MeV; this is due to the energy cut in the event selections.

the efficiency of the telescope. The spectral response matrix is the product of these two components:

$$R(I, E) = R_d(I, E) \times A_{\text{eff}}(E) . \quad (5.3)$$

We make the response matrices in [MEGALib](#) using the `responsecreator` tool, which gives $R(I, E)$ as a function of position relative to the instrument. We then split the resulting matrix into $R_d(I, E)$ and $A_{\text{eff}}(E)$ and convert the components into the FITS format ([FITS Working Group, 2016](#)) as required by [XSPEC](#).

Figure 5.1a shows an example redistribution matrix for [COSI](#). The matrix entries below the line $E = I$ are due to incompletely absorbed events, which cause the spectral continuum for line calibration sources. The redistribution matrix is a function of source zenith angle relative to [COSI](#) and of the event selections used in the analysis. For every change in event selections, a new redistribution matrix must be computed. The redistribution matrix in Figure 5.1a is for the case of an on-axis point source and 12 working detectors, and the event selections used are shown in Table 5.1.

The effective area is defined by

$$A_{\text{eff}}(E) = A_{\text{start}} \frac{N_{\text{meas}}(E)}{N_{\text{start}}(E)} \quad (5.4)$$

5.1 SPECTRAL FITTING AND THE RESPONSE MATRIX

Parameter	Allowed range
Photon energy	100 - 3000 keV
Compton scatter angle	0 - 180°
Number of interactions	2-7
Distance between first two interactions	≥ 0.5 cm
Distance between any two interactions	≥ 0.3 cm
Earth horizon cut	None
ARM	$\leq 11.9^\circ$

Table 5.1: Event selections for the redistribution matrix shown in Figure 5.1a and the effective area shown in Figure 5.1b. We chose these relatively open event selections to calculate the best case effective area. The distance cuts could be further opened, which would increase the effective area but decrease the overall quality of accepted events.

where A_{start} is the area over which the photons are simulated, or the area surrounding the mass model, $N_{\text{start}}(E)$ is the number of photons emitted by the source at energy E , and $N_{\text{meas}}(E)$ is the number of photons measured by the instrument at energy E . The effective area curve is a projection of the redistribution matrix normalized by the number of started photons and multiplied by A_{start} , and is also dependent on the zenith angle of the source relative to COSI and the event selections. An example effective area curve for COSI, also for an on-axis point source, 12 working detectors, and computed with the event selections shown in Table 5.1, is shown in Figure 5.1b.

When performing spectral analysis of a source that moves relative to COSI, we must take into account that both components of the spectral response matrix are dependent on the zenith angle of the source. The most correct way to address the dependence on zenith angle is to make a separate spectrum and response matrix for each zenith angle. An alternative solution is to compute an average of the response matrices at each zenith angle weighted by the time the source spent at each angle, and use the average matrix to fit the combined spectrum. The success of the weighted average method depends on the zenith angle range of the observation in question.

5.1.2 PIPELINE VERIFICATION WITH SIMULATIONS

To ensure that the spectral response matrix is calculated correctly and to confirm the conversion to FITS format, we simulated a Crab-like point source

5.2 MODELING THE ATMOSPHERIC ABSORPTION

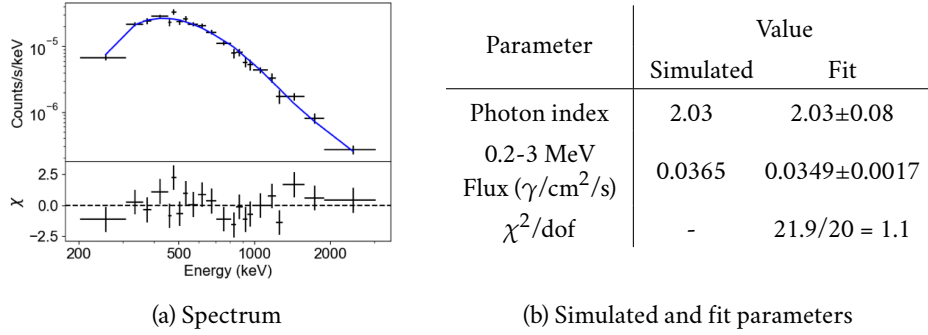


Figure 5.2: Spectral fit of a simulated Crab-like source without background. A comparison of the simulated and fit parameters is shown in (b). The fit results are consistent with the simulated parameters.

without background. The simulation has a power law spectrum with a photon index of 2.03 and a 0.2-3 MeV flux of $0.0365 \gamma/\text{cm}^2/\text{s}$. We fit the resulting spectrum in XSPEC, shown in Figure 5.2. As expected, we were able to recover the simulated spectral parameters; the results are detailed in the table in Figure 5.2. The simulation mimics the path of the Crab through COSI's field of view (FOV) during the 2016 flight and in doing so undergoes a zenith angle range of 28° to 40° . We used a weighted average of the response matrices calculated for each zenith angle. The successful fit indicates that the weighted average approach is a reasonable approximation for this zenith angle range.

5.2 MODELING THE ATMOSPHERIC ABSORPTION

When performing spectral fitting of sources observed by COSI, we must take into account that γ -rays are attenuated in the atmosphere. To do so, we have developed a model of this atmospheric attenuation to include in the spectral fits. The model takes the form

$$f(E) = e^{-n_A \sigma(E)} \quad (5.5)$$

where n_A is the column density in g/cm^2 and $\sigma(E)$ are the attenuation coefficients, which are dependent on the γ -ray energy. We multiply this model by the model with which we wish to fit the spectrum of the source; for example, if we are fitting the source with a power law $f(E) = KE^{-\alpha}$, the absorbed power law is then $f(E) = KE^{-\alpha} e^{-n_A \sigma(E)}$.




Figure 5.3: A schematic of our column density calculation. We assume a spherical atmosphere with an altitude of 100 km. We then integrate the NRLMSISE-00 model (Picone et al., 2002) along the dotted line between COSI and the source to compute the column density.

To calculate the column density, we integrate the NRLMSISE-00 atmosphere model (Picone et al., 2002) along the line of sight towards the source through a spherical atmosphere out to an altitude of 100 km, as shown in the schematic in Figure 5.3. We also use the NRLMSISE-00 model to determine the fraction of nitrogen, oxygen, and argon in the atmosphere, and use the NIST Photon Cross Sections Database (Berger et al., 2010) to determine the attenuation coefficients of the mixture.

The column density depends on COSI’s altitude and the zenith angle of the source (see Figure 5.4), as both of these values change the amount of atmosphere that the source photons travel through to reach the detector. When performing spectral analysis of persistent sources, it may be necessary to split the atmospheric attenuation model and observed spectrum by altitude and zenith angle.

5.3 SPECTRAL ANALYSIS OF GRB 160530A




Figure 5.4: The effects of (a) the zenith angle of the source relative to *COSI* and (b) *COSI*'s altitude on the column density. In (a), we assume an altitude of 33 km. In (b), we assume two source positions relative to *COSI*, on-axis (blue) and a 40° zenith angle (green).

5.3 SPECTRAL ANALYSIS OF GRB 160530A

COSI detected the bright, long duration gamma-ray burst GRB 160530A (Tom-sick & the *COSI* team, 2016) 14 days into the 2016 flight. Spectral analysis of transient sources such as GRB 160530A is a good test of the spectral analysis pipeline applied to flight data for a number of reasons. First, GRB 160530A was short enough to stay in roughly the same position in detector coordinates (i.e. relative to the instrument), preventing the added complication of needing multiple response matrices and atmospheric attenuation models. Secondly, subtracting the background is straightforward, as we can collect a background spectrum from time periods before and after the GRB. Finally, the spectrum of this GRB was measured by the Konus instrument on the Wind spacecraft, so we can compare our spectral fit results to those of Konus.

5.3.1 THE *COSI* OBSERVATION OF GRB 160530A

COSI observed GRB 160530A on May 30, 2016 at 07:03:46 Coordinated Universal Time (UTC). *COSI*'s geographic coordinates at the time of the observation were 56.79° S, 82.31° E, and *COSI* was floating at an altitude of 32.6 km. The close proximity of the South Magnetic Pole resulted in a relatively high background. At the time of the GRB, a relativistic electron precipitation (REP) event (Millan & Thorne, 2007) was occurring, which introduced some vari-




Figure 5.5: The COSI image of GRB 160530A using 10 iterations of the LM-MLEM algorithm. The color scale is in arbitrary units.

ability into the background count rate. Two of the detectors had failed by this point in the flight (detectors 5 and 8; see Section 4.5 for more details), resulting in 10 working detectors. Figure 5.5 shows the COSI image of GRB 160530A. The peak of the image is at $l = 243.4^\circ$, $b = 0.4^\circ$, which is currently the best known localization for this GRB. In the local instrument coordinate system, the GRB position corresponded to a zenith angle of 43.5° and an azimuth of -66.1° . See Lowell et al. (2017a) for more details about COSI’s observation of GRB 160530A and the polarimetric analysis of this GRB.

5.3.2 SPECTRAL ANALYSIS


The Konus team fit the spectrum of GRB 160530A with a Band function (Band et al., 1993) using 38.9 seconds of data between 20 keV and 5 MeV (Svinkin et al., 2016). To compare the COSI spectrum with the Konus spectrum, we selected events from the same 38.9 seconds for our analysis, shown in Figure 5.6. We chose background events from a 573 second period that spans before and after the GRB while leaving a 100 second gap on both sides of the GRB (see Figure 5.6).

The best event selections for this analysis are those that result in the highest signal significance:


$$S = \frac{r_S}{\sqrt{r_S + r_B}} \quad (5.6)$$

where r_S is the background-subtracted source count rate and r_B is the background count rate multiplied by a scaling factor to account for the difference

5.3 SPECTRAL ANALYSIS OF GRB 160530A



(a)



(b)

Figure 5.6: The measured light curve of GRB 160530A, where (a) shows the times used for source and background events and (b) shows a zoomed-in version of the GRB window from (a), revealing the structure in the light curve.

Parameter	Allowed range
Photon energy	100 - 3000 keV
Compton scatter angle	17.7-176.8°
Number of interactions	2-7
Distance between first two interactions	≥ 0.6 cm
Distance between any two interactions	≥ 0.14 cm
Earth horizon cut	None
ARM	$\leq 11.9^\circ$

Table 5.2: Event selections used for the spectral fitting of GRB 160530A. See Section 2.4.2 for a description of the event selections.

in exposure times. We used a Differential Evolution optimization algorithm (Storn & Price, 1997) to find the Compton scatter angle range, minimum interaction distances, and number of interactions that maximize S ; the resulting event selections are detailed in Table 5.2. The lower bound of the first interaction distance was set to 0.5 cm in the optimization algorithm, as events with a larger lever arm are more accurately reconstructed. We originally included the angular resolution measure (ARM) value in the list of parameters to optimize, but the ARM value returned by the optimizer was consistently pegged at the upper bound, even if the upper bound value was much larger than the ARM expected from a point source. Instead, we chose the ARM value of 11.9° used in the polarization analysis described in Lowell et al. (2016). We did not use the Earth horizon cut (EHC) since the GRB is not background-dominated.

Figure 5.7 shows the redistribution matrix and effective area curve computed for the spectral fits of GRB 160530A, using the event selections detailed in Table 5.2. We note that the effective area is lower than the effective area curve showed in Figure 5.1b. As the event selections used for GRB 160530A are only slightly more constrained than the event selections used for the effective area in Figure 5.1b, the majority of the observed decrease in effective area is due to the large zenith angle (43.5°) of GRB 160530A and the loss of two detectors.

We binned the background-subtracted spectrum of GRB 160530A such that each bin had the same signal-to-noise ratio $S/N = 3$. We fit the spectrum in XSPEC with the atmospheric absorption model multiplied by a Band function




Figure 5.7: The (a) redistribuition matrix and (b) effective area computed for spectral analysis of GRB 160530A, using the event selections detailed in Table 5.2. These plots are similar to those shown in Figure 5.1 but with different event selections and a source zenith angle of 43.5° instead of 0° .

(Band et al., 1993). We calculated the column density as described in Section 5.2; the result is $n_A = 9.2 \text{ g/cm}^2$. The Band function is defined as follows:

$$f(E) = \begin{cases} K(E/100)^\alpha e^{-E/E_c} & E < E_c(\alpha - \beta) \\ K[(\alpha - \beta)E_c/100]^{\alpha-\beta} (E/100)^\beta e^{-(\alpha-\beta)} & E > E_c(\alpha - \beta) \end{cases} \quad (5.7)$$

where α and β are the first and second power law indices, E_c is the characteristic energy, and K is the normalization constant. We first fit the spectrum with all parameters in the Band function free, but we were unable to constrain the parameters. We then froze the characteristic energy and second power law index to the best fit values as determined by the Konus team, and left the first power law index free. We were able to constrain the first power law index and the value is consistent to the Konus value. The model is a good fit to the COSI data, evidenced by the reduced χ^2 which is quite close to one. The fitted COSI spectrum is shown in Figure 5.8, and the fit results are compared to the Konus best fit parameters in Table 5.3.

Konus reported the GRB fluence, or integral of the flux over the duration of the burst, in the energy range of 0.2-10 MeV. To compare with their results, we extrapolated the energies in XSPEC to compute the 0.2-10 MeV unabsorbed flux, and multiplied it by the GRB exposure time. As shown in Table 5.3, the COSI fluence is in agreement with the Konus fluence.

This result is an important confirmation of COSI's spectral abilities. The fact that the COSI best fit parameters are consistent with the Konus best fit parameters confirms that the spectral analysis pipeline described in this chapter

5.3 SPECTRAL ANALYSIS OF GRB 160530A




Figure 5.8: The COSI spectrum of GRB 160530A fit with the atmospheric absorption model multiplied by the Band function. The spectrum was binned such that each bin had the same signal-to-noise ratio $S/N = 3$.

Parameter	COSI value	Konus value
Atmospheric N_A (g/cm ²)	9.2*	–
α	$-1.14^{+0.30}_{-0.32}$	-0.93 ± 0.03
β	-3.5*	< -3.5
E_c (keV)	638*	638^{+36}_{-33}
0.2-10 MeV Fluence (erg/cm ²)	$1.49^{+0.23}_{-0.21} \times 10^{-4}$	$1.3 \pm 0.04 \times 10^{-4}$
χ^2/dof	$27.54/30 = 0.92$	$73/73 = 1$

Table 5.3: The COSI and Konus best fit parameters of GRB 160530A. The Konus parameters are from Svinkin et al. (2016). The errors on the COSI parameters are 99% confidence. The parameters marked with an asterisk were fixed during the fit. The COSI parameters are in agreement with the Konus parameters, confirming COSI’s spectral abilities.

and the instrument itself work as expected. As the response matrix is calculated from simulations, this result also attests that due to the detector effects engine (DEE) described in Chapter 3, the simulations accurately mimic the measurements.

5.4 CHALLENGES OF PERSISTENT SOURCES

Performing spectral analysis of persistent sources with COSI is significantly more challenging for two reasons: first, the sources move within COSI’s FOV and COSI itself moves around the Earth, complicating the response matrix and atmospheric absorption model, and second, the background subtraction becomes nontrivial. Possible solutions to the movement of the sources and the instrument include splitting up the spectrum based on the source zenith angle and COSI’s geographic position or using a weighted average of the response matrix and atmospheric model in the spectral fits; these potential methods are discussed in Sections 5.1 and 5.2.

With a persistent source, choosing background photons from before or after the observation, as we did with GRB 160530A, is not possible. A potential method in a similar vein is to choose background events during time periods in which COSI was not observing the source in question. Due to the large variation of measured background rates as a function of COSI’s geographical position and altitude, however, this background subtraction strategy is difficult to implement. For every combination of altitude and geomagnetic cutoff rigidity, a measure of the magnetic shielding of the Earth from cosmic rays that is closely related to geographic latitude, we must sample the background at a time when the source is not in the FOV but COSI is either physically located at the same geographic position and altitude or located at a different location with comparable background rates. Because COSI is a balloon-borne instrument, we have no control over its geographic position, and there is no guarantee that we can sample the background as required for this method for any particular observation of a source, especially with the altitude drops that occurred during the 2016 flight. Thus, it is best to select background events from time periods in which the source was in COSI’s FOV.

One such method is to subtract the background in image space. To do so, we choose a source region around the source and a background region elsewhere in the FOV and generate spectra from the photons consistent with these regions. We then scale the background region spectrum by the difference in area on the sky between the source region and the background region, and fi-




Figure 5.9: A schematic of the image-space background subtraction method. The source region is represented in red and the background regions are represented in blue. Many background regions are chosen for more statistics. The background regions are at the same zenith angle as the source region and are spaced 30° apart in azimuth, though there is a 60° gap between the source region and the nearest background regions, resulting in a total of nine background regions.

nally subtract the scaled background region spectrum from the source region spectrum. Because the background depends strongly on zenith angle, as does *COSI*'s effective area, it is important that the background region have the same zenith angle as the source region. To fulfill this requirement, we choose the source and background regions in detector coordinates rather than Galactic coordinates. In particular, we choose nine background regions at the same zenith angle as the source region, spaced 30° apart in azimuth with a 60° gap between the source region and the nearest background regions, as shown in Figure 5.9. The size of all regions is the same and is determined by the *ARM* full width at half maximum (*FWHM*) of the source in question.

Each event has an associated rotation matrix to convert between Galactic and detector coordinates at that instant in time. For each event, we convert the known source position in Galactic coordinates (determined from the image, or from other instruments if it is a previously observed source) into detector coordinates. Once we know the source position in detector coordinates, we define the source and background regions at that instant. We determine if the event in question is consistent with any of these regions by calculating the *ARM* of the event with respect to the region center and comparing it to the region size; if the *ARM* is smaller than the region size then the event is




Figure 5.10: The results of the image-space background subtraction method applied to a simulation of a point source *without* background. The events in the background spectrum are source events that are consistent with background regions. The background spectrum is averaged over the nine background regions. The subtracted source spectrum contains significantly fewer counts than the unsubtracted source spectrum, indicating that we are unable to correctly distinguish between source and background events in image space.

consistent with that region. Because each event is localized to a circle on the sky rather than a point, it is possible for events to be consistent with multiple regions.

There are two major drawbacks to this background subtraction method. The first is that for small zenith angles, i.e. sources that are almost directly on-axis, the regions begin to overlap. The second is that many events are consistent with multiple regions, which results in an oversubtraction of the background. To determine the effects of events that are consistent with multiple regions, we simulated a point source with a power law spectrum with index 2 without background and performed the image-space background subtraction. We set the radius of the regions to 7.74° , the [ARM FWHM](#) of the point source simulation. In this analysis, all “background” events, or events consistent with background regions, are actually from the source, as no background was included in the simulation. The results of this analysis are shown in Figure 5.10, where it is evident that the “background” is quite prominent and almost half of the source counts are missing from the background-subtracted spectrum. These results indicate that we cannot effectively separate source events from background events in image space.

We have developed a third background subtraction method for persistent sources using the [COMPTEL Data Space \(CDS\)](#), a data space in which each

event is represented as a point rather than a circle. Similarly to image space, a point source in the CDS makes a clearly defined region. Since events in the CDS are represented as points, they cannot be consistent with both the source and background regions, and thus it is theoretically straightforward to separate source and background events. Similarly to the image-space background subtraction, background events are chosen concurrently with source events, eliminating the difficulties introduced by the highly variable background. The CDS background subtraction method will be discussed in detail in Chapter 6.

Though I have framed the challenges of background subtraction of persistent sources in terms of spectral analysis, note that these same challenges exist for other types of astrophysics analysis, such as polarimetry. To proceed with both spectral and polarization analysis of persistent sources detected by COSI, a good background subtraction method is essential.

6

THE COMPTTEL DATA SPACE

The **COMPTTEL** collaboration pioneered a three-dimensional data space specific to Compton telescopes as a useful analysis tool (Knödlseider et al., 1996), hereafter referred to as the **COMPTTEL** Data Space (**CDS**). Since each event can be uniquely mapped to a point in the **CDS** rather than a circle in the more intuitive image space, the **CDS** is a promising tool for background subtraction, fitting and filtering. Kierans (2018) developed a background subtraction method using the **CDS** for sources detected by **COSI** that emit spectral lines, such as the positron annihilation signature from the Galactic Center; however, this method cannot be used on broadband sources since it assumes that there are no source counts at energies higher than the line energy. I have developed a more general background subtraction method using the **CDS** designed specifically for broadband sources. In this chapter, I will describe the **CDS**, the background subtraction algorithm for broadband sources, and the verification of this background subtraction method with simulations. Note that while I discuss the **CDS** background subtraction method in terms of spectral analysis, the algorithm can also be applied to polarimetry.

6.1 INTRODUCING THE COMPTTEL DATA SPACE

The **CDS** is a three-dimensional data space defined by the photon scatter direction in detector coordinates $\vec{d}_\gamma = (\chi, \psi)$ and the initial Compton scatter angle ϕ . Figure 6.1 shows the first two interactions of a Compton event and the three angles, ϕ , χ , and ψ , that define the **CDS**. Each event is represented in the **CDS** as a single point at (χ, ψ, ϕ) .

Events from a point source fill the **CDS** as a three-dimensional hollow cone with an apex at the source position (χ_0, ψ_0) in the $\chi - \psi$ plane (see Figure 6.1 for a schematic). As the Compton scatter angle ϕ approaches zero, \vec{d}_γ approaches (χ_0, ψ_0) ; in other words, the photon is barely scattered from its initial trajectory. Thus at $\phi \rightarrow 0$, all events will have a $\vec{d}_\gamma \rightarrow (\chi_0, \psi_0)$, forming

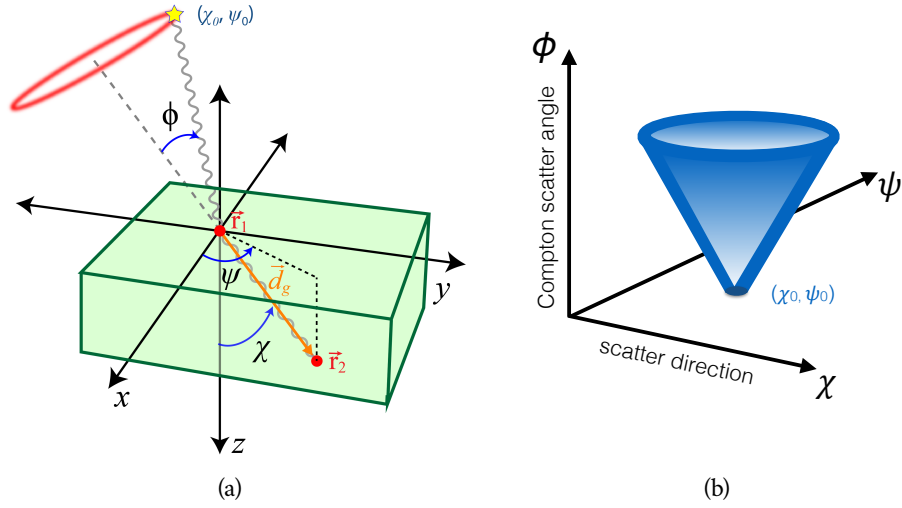


Figure 6.1: (a) The first two interactions of a Compton event, showing the angles that define the CDS. The photon originates from the source at position (χ_0, ψ_0) and scatters at \vec{r}_1 and \vec{r}_2 . The scatter direction \vec{d}_γ is the vector between \vec{r}_1 and \vec{r}_2 . The angles χ and ψ represent the polar and azimuth angles of \vec{d}_γ , respectively. The initial Compton scatter angle ϕ is defined as the angle between the initial trajectory of the photon and \vec{d}_γ .

(b) A schematic of a point source in the CDS. The apex of the cone is at $(\chi = \chi_0, \psi = \psi_0, \phi = 0)$, where (χ_0, ψ_0) is the position in detector coordinates from which the photon originated.

Figures from Kierans (2018).

the apex of the cone at the source position. As ϕ increases, the photons are scattered further from the direction to the source, resulting in the conical shape. The opening angle of the cone is 90° since ϕ and the polar scatter direction χ increase at the same rate.

6.1.1 ON-AXIS SOURCES IN THE CDS

In the case of a point source directly above the instrument at $(\chi_0, \psi_0) = (0, 0)$, the apex of the cone will be at the origin of the CDS, and instead of forming a cone, the events will form a two-dimensional plane at 45° as shown in Figure 6.2. The events form a plane because the polar angle of the scatter direction χ by definition ranges from 0° to 90° and cannot be negative. Therefore with changes in azimuth, the point source shape does not circle around the origin like a cone, but rather forms the aforementioned plane. By integrating along the ψ direction, we can project the CDS into the $\chi - \phi$ plane. We note that in the case of a source directly on-axis, ψ is the azimuthal scatter angle η from Equation 2.2, and thus encodes the polarization information of the source. The CDS could also be used as a tool for polarization analysis by searching for modulation in ψ . Here, however, we are using the CDS as a tool for background subtraction: for spectral analysis and for performing polarimetry of persistent broadband sources using the standard method as described in Section 7.3.1, we use the CDS to provide source and background event lists. In this case it is unnecessary to preserve the polarization information in the CDS. Therefore projecting the CDS into the $\chi - \phi$ plane – hereafter referred to as the 2D CDS – by integrating along the ψ direction is an acceptable geometric simplification.

Here I will explain how the 2D CDS is populated in the case of an on-axis source at $(\chi_0, \psi_0) = (0, 0)$. In the following section, I will generalize this example to arbitrary positions. Figure 6.2a shows an event that Compton scatters with angle ϕ_1 to scatter direction $\vec{d}_\gamma = (\chi_1, \psi_1)$. Figure 6.2d shows the 2D CDS (the $\chi - \phi$ plane) and indicates the location of (χ_1, ϕ_1) in this projection. Because ϕ_1 is close to zero, this event in the CDS will be near the origin of the dataspace, which is at $(\chi, \phi) = (0, 0)$ in the 2D CDS. Figure 6.2b shows an event that Compton scatters with an angle $\phi_2 > \phi_1$ to scatter direction $\vec{d}_\gamma = (\chi_2, \psi_2)$. Again, one can see the location of this event in the 2D CDS in Figure 6.2d. Because χ is the polar angle of \vec{d}_γ , χ is always equal to ϕ ; thus all events from the on-axis source fill the plane $\phi = \chi$ (Figure 6.2c), which becomes the straight line defined by $\phi = \chi$ in the 2D CDS (Figure 6.2d).

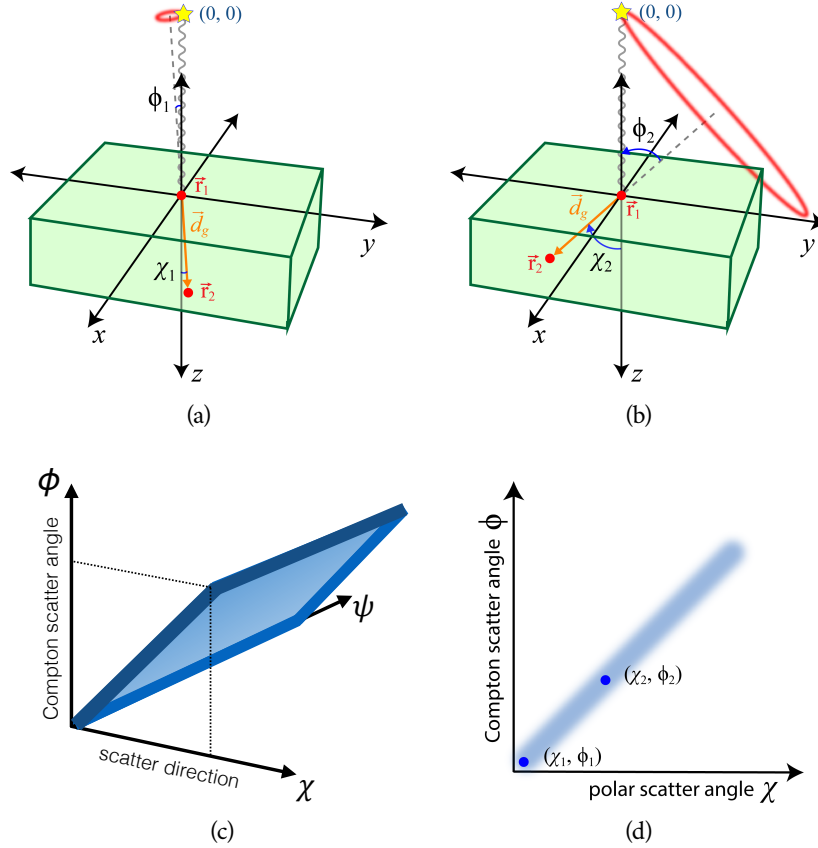


Figure 6.2: An on-axis source in the CDS. (a) An event that scatters at a small Compton scatter angle ϕ_1 with the resulting χ_1 . (b) An event that scatters at a large Compton scatter angle $\phi_2 > \phi_1$ with the resulting χ_2 . (c) The $\chi - \phi$ plane in the 3D CDS. (d) The source region defined by the ideal line $\phi = \chi$ in the 2D CDS, which results from integrating the 3D CDS along the ψ direction. The source region has some width related to the instrument's angular resolution. The locations of the two events shown in (a) and (b) in the 2D CDS are marked. *Figures from Kierans (2018).*




Figure 6.3: A simulation of an on-axis point source in the 2D CDS. The line at $\phi = \chi$ is clearly visible and has a thickness related to COSI’s angular resolution. The events exterior to this line are likely incompletely absorbed or improperly reconstructed events.

Figure 6.3 shows a simulation of an on-axis point source in the 2D CDS. The line $\phi = \chi$ is clearly visible and has a thickness related to COSI’s angular resolution. There are many events filling the 2D CDS that are not on the line $\phi = \chi$, despite the fact that this simulation is of a point source without background. These events are attributed to incompletely absorbed events or events that were not properly reconstructed. Using the event selections described in Section 2.4.2 to select better quality events would reduce the relative number of events far from the line $\phi = \chi$. The thickness of the point source line at $\phi = \chi$ and the effects of the incompletely absorbed and improperly reconstructed events will be discussed in Section 6.2.

It is relatively straightforward to use the 2D CDS to extract source and background events. In the next section, I will describe how we can rotate the CDS such that point sources at any position form the plane $\phi = \chi$.

6.1.2 GALACTIC COORDINATES AND OFF-AXIS SOURCES IN THE CDS

When analyzing astrophysical sources, the source position in detector coordinates changes as the source moves relative to COSI. By rotating \vec{d}_γ into Galactic coordinates for each event, we effectively rotate the entire CDS into Galactic coordinates. In this scenario, a point source still forms a cone with the apex at its source position, but its source position is in Galactic coordinates at (χ_G, ψ_G) .

To rotate \vec{d}_γ into Galactic coordinates for a single event, we utilize our knowledge of COSI's position when that event was detected. Each event has an associated rotation matrix \mathbf{R}_G that allows us to rotate between detector and Galactic coordinates, as follows:

$$\vec{d}_{\gamma,G} = \mathbf{R}_G \times \vec{d}_\gamma \quad (6.1)$$

where $\vec{d}_{\gamma,G}$ is the scatter direction of the photon in Galactic coordinates.

It is convenient to use the 2D CDS; to do so, however, the source position must be at $(0,0)$. It is possible to rotate the CDS such that any source position (χ_G, ψ_G) is rotated into Galactic coordinates $(l, b) = (0,0)$. By performing this rotation, all events from the source will be mapped to the line $\phi = \chi$ in the 2D CDS, no matter the source position (χ_G, ψ_G) .

To perform this rotation, we can use predefined formulae to rotate one vector into another on a sphere, but we must use spherical coordinates rather than Galactic coordinates. These two coordinate systems differ as follows: in Galactic coordinates, the polar angle b is equal to zero when the polar angle Θ in spherical coordinates is equal to $\pi/2$. In other words, $(0,0)$ in Galactic coordinates is at the Galactic center, or where the prime meridian crosses the equator, whereas $(0,0)$ in spherical coordinates is at the top of the sphere. Thus, to rotate an arbitrary source position (χ_G, ψ_G) into the Galactic coordinates position $(l = 0, b = 0)$, we must rotate the source position into the spherical coordinates position $(\Theta = \frac{\pi}{2}, \Phi = 0)$. Similarly, we transform the source position (χ_G, ψ_G) into spherical coordinates: $(\Theta, \Phi) = (\frac{\pi}{2} - \chi_G, \psi_G)$.

To rotate one vector \vec{a} into another \vec{b} in spherical coordinates, we find the angle between the two vectors and the vector about which we are rotating. In this case, \vec{a} represents the arbitrary source position (χ_G, ψ_G) , or $\vec{a} = \hat{r} + \Theta\hat{\theta} + \Phi\hat{\phi} = \hat{r} + (\frac{\pi}{2} - \chi_G)\hat{\theta} + \psi_G\hat{\phi}$ on the unit sphere. \vec{b} represents the Galactic center, which is $\vec{b} = \hat{r} + \frac{\pi}{2}\hat{\theta} + 0\hat{\phi}$ on the unit sphere. We transform \vec{a} and \vec{b} into Cartesian coordinates:

$$\begin{aligned} \vec{a} &= \sin \Theta \cos \Phi \hat{x} + \sin \Theta \sin \Phi \hat{y} + \cos \Theta \hat{z} \\ \vec{b} &= \hat{x} \end{aligned} \quad (6.2)$$

To rotate \vec{a} into \vec{b} , we find the angle α about which we are rotating using the dot product:

$$\alpha = \arccos \left(\frac{\vec{a} \cdot \vec{b}}{\|\vec{a}\| \|\vec{b}\|} \right) = \arccos(\sin \Theta \cos \Phi) \quad (6.3)$$

and the vector \vec{k} about which we are rotating using the cross product:

$$\vec{k} = \frac{\vec{a} \times \vec{b}}{\|\vec{a} \times \vec{b}\|} = \frac{\cos \Theta \hat{y} - \sin \Theta \sin \Phi \hat{z}}{\sqrt{\cos^2 \Theta + \sin^2 \Theta \sin^2 \Phi}} . \quad (6.4)$$

Once we have α and \vec{k} , we can use the Rodrigues Rotation formula (Rodrigues, 1840) that defines a rotation matrix \mathbf{R}_{To0} we use to rotate \vec{a} into \vec{b} :

$$\mathbf{R}_{\text{To0}} = \mathbf{I} + (\sin \alpha) \mathbf{K} + (1 - \cos \alpha) \mathbf{K}^2 \quad (6.5)$$

where

$$\mathbf{K} = \begin{bmatrix} 0 & -k_z & k_y \\ k_z & 0 & -k_x \\ -k_y & k_x & 0 \end{bmatrix} = \frac{1}{\|\vec{k}\|} \begin{bmatrix} 0 & \sin \Theta \sin \Phi & \cos \Theta \\ -\sin \Theta \sin \Phi & 0 & 0 \\ -\cos \Theta & 0 & 0 \end{bmatrix} \quad (6.6)$$

and

$$\|\vec{k}\| = \sqrt{\cos^2 \Theta + \sin^2 \Theta \sin^2 \Phi} . \quad (6.7)$$

For each event in the observation, we rotate the scatter direction with the rotation matrix \mathbf{R}_{To0} :

$$\vec{d}_{\gamma, G_0} = \mathbf{R}_{\text{To0}} \times \vec{d}_{\gamma, G} = \mathbf{R}_{\text{To0}} \times (\mathbf{R}_G \times \vec{d}_{\gamma}) . \quad (6.8)$$

Here, \vec{d}_{γ, G_0} represents the scatter direction of the photon in rotated Galactic coordinates such that the apex of the CDS point source cone is at $(l, b) = (0, 0)$.

Using Equation 6.8, we can rotate the CDS such that a point source makes a plane $\phi = \chi$ for any source position in Galactic coordinates. Then we integrate along the ψ direction to form the 2D CDS.

6.2 BACKGROUND SUBTRACTION WITH THE CDS

With our understanding of the theoretical shape of a point source in the 2D CDS, the line $\phi = \chi$, we can define source and background regions relative to this ideal line. In reality, the line made from a point source will have some thickness Δ related to COSI's angular resolution. From this thickness, we can define the source and background regions. We define the source region **SR**




Figure 6.4: A schematic of the source and background regions in the 2D CDS. The source region includes events within Δ of the ideal line $\phi = \chi$. The background regions are on either side of the source region and include events between Δ and 2Δ from the ideal line.

as the region containing events whose distances from the ideal line are between 0 and Δ . We define the background region **BR** in the 2D CDS as the region containing events whose distances from the ideal line are between Δ and 2Δ . The background region is adjacent to the source region. See Figure 6.4 for a schematic of these regions. Note that there are two components to the background region, one on either side of the source region, referred to as the inner background region **BR_{in}** and the outer background region **BR_{out}**. Events are in **BR_{in}** when $\phi > \chi$; this region corresponds to a cone inside the source cone in the 3D CDS. Likewise, events are in **BR_{out}**, corresponding to a cone outside the source cone in the 3D CDS, when $\chi > \phi$.

For each event, we determine the shortest distance D_i from that event's location in the 2D CDS to the ideal line. This shortest distance is along a line perpendicular to the ideal line, or:

$$D_i = (\phi_i - \chi_i)/\sqrt{2} \quad (6.9)$$

where ϕ_i and χ_i represent the Compton scatter angle and polar angle of the scatter direction in Galactic coordinates of event i . If $|D_i| \leq \Delta$, then the event is in **SR**, and if $\Delta < |D_i| \leq 2\Delta$, then the event is in **BR**.

To determine Δ , or the width of **SR**, we simulate a point source without background. For each event in the simulation, we calculate D_i and make a distribution of all D_i values. This distribution is analogous to the angular resolu-




Figure 6.5: A distribution of distances of events from the ideal line $\phi = \chi$ in the 2D CDS. This distribution, referred to as the CDS-ARM distribution, is analogous to the ARM distribution in image space and describes COSI’s angular resolution. This particular CDS-ARM distribution was made from a simulation of a point source. The source and background regions **SR** and **BR** are marked for a Δ of 5.1° , which is the FWHM of this particular distribution.

tion measure (ARM) distribution in image space, as the full width at half maximum (FWHM) of the distribution describes COSI’s angular resolution; thus, we will refer to the distribution as the CDS-ARM distribution, where Δ is the FWHM of the CDS-ARM distribution. Figure 6.5 shows an example CDS-ARM distribution calculated for a simulation of a point source with a power law spectrum. Note the resemblance to a typical ARM distribution, e.g. Figure 2.4. The source and background regions are marked for a Δ of 5.1° , which is the FWHM of this particular CDS-ARM distribution.

To perform the background subtraction, we generate spectra from **SR** and **BR**. Let $N_{\text{SR}}(E)$ denote the spectrum from the source region and $N_{\text{BR}}(E)$ denote the spectrum from the background region. If the CDS were evenly populated and if the regions were the same size, we could find the background-subtracted source spectrum $N_S(E)$ as follows:

$$N_S(E) = N_{\text{SR}}(E) - N_{\text{BR}}(E) . \quad (6.10)$$

However, the background within the CDS is not evenly distributed throughout the data space, as we will show in the next section. In particular, **SR** and **BR** are not populated evenly by background data, which requires adjustments to the basic algorithm. Additionally, as shown in Figure 6.5, there are

6.2 BACKGROUND SUBTRACTION WITH THE CDS




Figure 6.6: The background from a single day of flight (June 13, 2016) in the 2D CDS. The CDS has been rotated such that the Crab position is at the origin. The locations of the source and background regions, defined here with $\Delta = 10^\circ$, are indicated. The blue points represent all of the Compton events from that day. The green points are all the Compton events with the EHC applied to the data. The red points represent the remaining Compton events after applying a 40° pointing cut around the Crab position.

some source counts present in the background regions due to the tailing of the CDS-ARM distribution, requiring further adjustments to the algorithm.

6.2.1 BACKGROUND IN THE CDS

Figure 6.6 shows how the background fills the 2D CDS on June 13, 2016, a single day during the flight during which COSI had particularly good exposure of the Crab. The rotations from Equation 6.8 have been applied to the dataset such that the Crab position ($l = 184.56$, $b = -5.78$) has been rotated to the origin of the CDS. As is typical for a balloon-borne γ -ray telescope, the observation is background dominated; the line of source photons at $\phi = \chi$ is not at all visible.

Figure 6.6 shows the background with three different sets of events. The blue points represent all of the Compton events from June 13, 2016. The green points represent the remaining Compton events after the Earth horizon cut (EHC) is applied. The red points represent the Compton events after both the EHC and a 40° pointing cut around the Crab position is applied (see

Section 2.4.2 for a description of these event selections). The **EHC** and pointing cut significantly reduce the number of background events.

A consequence of these event cuts, however, is that the **CDS** is not uniformly populated and so we cannot use Equation 6.10 to subtract the background. Instead, we must scale the background, which I will describe in Section 6.2.2. Additionally, the source and background regions are not exactly the same size because of differences at low Compton scatter angles (see Figure 6.4). The scaling method naturally takes into account the differences in region size.

6.2.2 SCALING THE BACKGROUND IN THE CDS

We scale the background with a scale factor f such that $f \times N_{\mathbf{BR}}$ represents the number of background counts in **SR**. Then we correctly recover the number of source counts:

$$N_S = N_{\mathbf{SR}} - f \times N_{\mathbf{BR}} . \quad (6.11)$$

The scale factor f depends on the position of the source relative to **COSI** and the Compton scatter angle and energy of the event. Thus, we must perform the scaled background subtraction as a function of these parameters:

$$N_S = \sum_{\Theta, \Phi, \phi, E} \left(N_{\mathbf{SR}}(\Theta, \Phi, \phi, E) - f(\Theta, \Phi, \phi, E) \times N_{\mathbf{BR}}(\Theta, \Phi, \phi, E) \right) \quad (6.12)$$

where Θ and Φ are the zenith and azimuth angles, respectively, of the source position relative to **COSI** (i.e. the source position in detector coordinates), ϕ is the Compton scatter angle, and E is the energy.

To determine the scale factor $f(\Theta, \Phi, \phi, E)$, we use a large background simulation that covers **COSI**'s entire field of view (**FOV**) to make a five-dimensional scaling matrix: the dimensions are Θ , Φ , ϕ , E , and source region **SR** or background region **BR**. For each event in the simulation, we iterate over all Θ and Φ values. We designate each (Θ, Φ) as the source position in detector coordinates, and use that source position to define the source and background regions **SR** and **BR**. We then determine if the event in question is in either **SR** or **BR**. Next, we increment the value of the bin defined by this particular (Θ, Φ) , the Compton scatter angle ϕ of the event, the energy E of the event, and **SR** or **BR**, by one.

The result of this process is a matrix $\mathbf{M}(\Theta, \Phi, \phi, E, R)$, where R can be either **SR** or **BR** to designate the source or background region, respectively. The value of $\mathbf{M}(\Theta, \Phi, \phi, E, R)$ is the number of counts in the source or background region with Compton scatter angle ϕ and energy E when the source position is at (Θ, Φ) in detector coordinates. Thus the scale factor f is defined as:

$$f(\Theta, \Phi, \phi, E) = \frac{\mathbf{M}(\Theta, \Phi, \phi, E, \mathbf{SR})}{\mathbf{M}(\Theta, \Phi, \phi, E, \mathbf{BR})} . \quad (6.13)$$

To apply this scale factor to the background, we extract the counts from the background region and generate a background region spectrum for each Θ , Φ , and ϕ bin (for data in Galactic coordinates, e.g. flight data, we use the event's associated Galactic pointing rotation matrix to determine the source position in detector coordinates at the time of the event). Handling the energy axis is slightly more complicated. When performing spectral fits in XSPEC, the initial binning of the spectra must correspond to the binning of the spectral response matrix, which is 500 logarithmic bins between 10 keV and 20 MeV. The spectral response matrix binning is not necessarily the same as the binning of the scaling matrix \mathbf{M} . Thus, we perform the algorithm as follows: for each background spectrum corresponding to a particular Θ , Φ , and ϕ , we bin it in energy according to the spectral response matrix. We iterate over each energy bin in the spectrum and determine whether or not that energy bin is fully contained in one scaling matrix energy bin or split across two scaling matrix energy bins. If the energy bin is fully contained in one scaling matrix energy bin $E_{\mathbf{M}}$, we scale that bin value with $f(\Theta, \Phi, \phi, E_{\mathbf{M}})$. If the energy bin is split across two scaling matrix energy bins $E_{\mathbf{M}_1}$ and $E_{\mathbf{M}_2}$, we scale that bin value with an average scaling factor $f_{\text{avg}} = 1/2 \times (f(\Theta, \Phi, \phi, E_{\mathbf{M}_1}) + f(\Theta, \Phi, \phi, E_{\mathbf{M}_2}))$.

When scaling the background, we also propagate the errors from the scaling matrix. If there are n counts in a given bin in the background spectrum, the scaled number of counts becomes

$$n_{\text{scaled}} = f(\Theta, \Phi, \phi, E) \times n = \frac{\mathbf{M}(\Theta, \Phi, \phi, E, \mathbf{SR})}{\mathbf{M}(\Theta, \Phi, \phi, E, \mathbf{BR})} \times n . \quad (6.14)$$

The error on each bin of the scaling matrix is Poissonian, as is the error on n ; i.e. $\sigma_{\mathbf{M}(\Theta, \Phi, \phi, E, R)} = \sqrt{\mathbf{M}(\Theta, \Phi, \phi, E, R)}$ where R represents **SR** or **BR**, and $\sigma_n = \sqrt{n}$. The error on n_{scaled} is then

$$\sigma_{n_{\text{scaled}}} = n \times f \sqrt{\left(\frac{\sigma_f}{f}\right)^2 + \frac{1}{n}} \quad (6.15)$$

where

$$\sigma_f = f \times \sqrt{\frac{1}{\mathbf{M}(\Theta, \Phi, \phi, E, \mathbf{SR})} + \frac{1}{\mathbf{M}(\Theta, \Phi, \phi, E, \mathbf{BR})}}. \quad (6.16)$$

6.2.3 CDS-ARM TAILING CORRECTION

One additional complication with the CDS background subtraction method is that not all source events are within the source region, or within $\pm\Delta$ of the ideal line $\phi = \chi$; see e.g. Figure 6.5. When we perform the background subtraction, these source counts are included in the background spectrum, resulting in an oversubtraction of background. We will refer to this effect as the tailing of the CDS-ARM distribution.

A naïve solution to this problem is to choose background regions farther away from the source region; i.e. to define \mathbf{BR} as the region including events that satisfy $n\Delta < D_i < (n+1)\Delta$ where $n > 1$. Because background unevenly populates the CDS, however, a large n results in very few background events in \mathbf{BR} ; see the red events in Figure 6.6. The tails of the point source CDS-ARM distribution can extend quite far, so even with a reasonable $n > 1$, we will still oversubtract the background. Since using $n > 1$ to define \mathbf{BR} will not solve this problem and will provide fewer statistics at small Compton scatter angles, we proceed with $n = 1$.

Instead of avoiding the tails of the CDS-ARM distribution, we correct for this oversubtraction by subtracting the oversubtracted source counts from the background spectrum. To do so, however, we require prior knowledge of the CDS-ARM distribution so that we know how many counts to remove. Ideally, we could retrieve the CDS-ARM distribution from simulations of any point sources detected during flight. In practice, this is not quite so straightforward. The difficulties will be discussed in Section 6.4.2. In what follows, we assume the ideal case where the shape of the CDS-ARM distribution obtained from simulations is sufficient.

Let S be the total number of source counts detected by COSI. We will refer to the number of source counts in the source region as $S_{\mathbf{SR}}$ and the number of source counts in the background region as $S_{\mathbf{BR}}$. The goal is to determine $S_{\mathbf{BR}}$, the number of oversubtracted counts that we will subtract from the background spectrum. Note that we do not know S until we have correctly subtracted the background.

The first step of the correction is to use the simulated source CDS-ARM distribution to determine the fraction of source counts in the source region,

$r'_{\text{SR}} = S'_{\text{SR}}/S'$, and the fraction of source counts in the background region, $r'_{\text{BR}} = S'_{\text{BR}}/S'$. The primed variables indicate that these values come from the simulated **CDS-ARM** distribution rather than from the **CDS-ARM** distribution of the detected point source from which we are attempting to subtract the background. Since the *shapes* of these distributions are the same, r'_{BR} should equal r_{BR} even if S' is not equal to S .

Now consider the dataset with both a point source and background. Let B_{SR} and B_{BR} refer to the number of background counts in the source and background region, respectively. Then the total number of counts in the source region is $N_{\text{SR}} = S_{\text{SR}} + B_{\text{SR}}$ and the total number of counts in the background region is $N_{\text{BR}} = S_{\text{BR}} + B_{\text{BR}}$. When we perform the background subtraction, we calculate $N_{\text{sub}} = N_{\text{SR}} - f \times N_{\text{BR}}$ where f is the scale factor. We use N_{sub} here rather than N_{S} since the quantity $N_{\text{SR}} - f \times N_{\text{BR}}$ is not actually the number of source counts because of the oversubtraction that we are attempting to correct. Note that $B_{\text{SR}} - f \times B_{\text{BR}} = 0$; this is the background scaling method described in Section 6.2.2. Then we have:

$$\begin{aligned}
 N_{\text{sub}} &= N_{\text{SR}} - f \times N_{\text{BR}} \\
 &= S_{\text{SR}} - f \times S_{\text{BR}} + B_{\text{SR}} - f \times B_{\text{BR}} \\
 &= S_{\text{SR}} - f \times S_{\text{BR}} + 0 \\
 &= r_{\text{SR}} \times S - f \times r_{\text{BR}} \times S \\
 N_{\text{sub}} &= (r_{\text{SR}} - f \times r_{\text{BR}}) \times S .
 \end{aligned} \tag{6.17}$$

$N_{\text{sub}} = (r_{\text{SR}} - f \times r_{\text{BR}}) \times S$ is a quantity that we measure from performing the background subtraction using Equation 6.12. We can substitute r_{SR} and r_{BR} with the values determined from the simulated **CDS-ARM** distribution, r'_{SR} and r'_{BR} . Note that f is determined from the scaling matrix \mathbf{M} . Now the total number of source counts, S , is the only unknown in the equation. We can then solve for S :

$$S = \frac{N_{\text{sub}}}{r'_{\text{SR}} - f \times r'_{\text{BR}}} . \tag{6.18}$$

Recall that the goal is to determine S_{BR} , the number of source counts in the background region, or the number of counts being incorrectly subtracted from the source. S_{BR} relates to S as follows: $S_{\text{BR}} = r_{\text{BR}} \times S$. We can again replace r_{BR} with the value from the simulated **CDS-ARM** distribution, r'_{BR} . Then we have determined S_{BR} as a function of known variables:

$$S_{\text{BR}} = r'_{\text{BR}} \times S = r'_{\text{BR}} \times \frac{N_{\text{sub}}}{r'_{\text{SR}} - f \times r'_{\text{BR}}} . \tag{6.19}$$

The next step is to subtract S_{BR} counts from the background spectrum. To do so, we must determine the spectral shape of the source to know how many of the S_{BR} counts to subtract at each energy; however, as the end goal is to fit the background-subtracted spectrum, presumably the spectral shape is unknown. Consequently, we use an iterative process to together fit the spectrum and correct for the oversubtraction of background, described in the following steps:

1. Subtract the background using Equation 6.12, resulting in N_{sub} counts. This will lead to an oversubtracted background as described above. Fit the background-subtracted spectrum in XSPEC and get the parameters describing the spectral shape. These fits include the instrument response and thus result in parameters describing the spectrum emitted by the source (see Chapter 5 for more details).
2. Subtract S_{BR} counts distributed according to the best fit spectrum to the background spectrum, and subtract the background spectrum from the total spectrum (i.e. the spectrum from the source region, containing source and background events).
3. Re-fit the background-subtracted source spectrum now that the S_{BR} counts have been removed from the background spectrum, and get updated best fit parameters.
4. Repeat steps 2 and 3 until the best fit parameters converge; i.e. until the change in parameter values between iterations is less than the threshold accuracy of 10^{-4} .

We will show in Section 6.3 that an average scale factor

$$f_{\text{avg}} = \frac{1}{n} \sum_{\Theta} \sum_{\Phi} \sum_E \sum_{\phi} f(\Theta, \Phi, E, \phi) \quad (6.20)$$

is sufficient to reproduce input parameters for simulations when correcting for the CDS-ARM tails. The dependency of the f used for the CDS-ARM tail correction on the source position in detector coordinates (Θ, Φ) , the energy E , and the Compton scatter angle ϕ has a marginal influence on the inferred spectral parameters. Additionally, using an average value f_{avg} significantly simplifies the CDS-ARM tail correction algorithm.

6.3 SPECTRAL BACKGROUND SUBTRACTION VALIDATION

We will use simulations to verify the background subtraction method described in this chapter. To simulate the atmospheric background encountered during COSI’s flight, we used the atmospheric background model from Ling (1975), an empirical model derived from the measured flux of a balloon-borne germanium detector. The Ling Model gives the γ -ray flux as a function of atmospheric depth, zenith angle, and energy. We determine the atmospheric depth input to the Ling Model using the method described in Section 5.2. With the Ling Model parameters as input to `cosima`, we can simulate the atmospheric background in COSI’s entire FOV for the duration of the 2016 flight.

6.3.1 BACKGROUND ONLY

As an initial verification, we first performed the background subtraction on simulations of background only and confirmed that the subtracted spectrum is consistent with no source. We applied the background subtraction algorithm to five simulated “source” positions, including the Crab position $(l, b) = (184.56, -5.78)$ and four random points in the sky with good exposure according to COSI’s 2016 flight exposure map. Note that the background simulations used mimic COSI’s movement relative to the Earth and thus simulate the same view of the sky that COSI had during the 2016 flight. Each “source” position in Galactic coordinates went through a variety of positions in detector coordinates.

For each position, we show a subtracted spectrum and the total number of subtracted counts in Figure 6.7. The event selections used for this analysis are shown in Table 6.1. Note that the same event selections were used to make the scaling matrix \mathbf{M} as well as the background-subtracted spectra. We limited the Compton scatter angle range to angles below 65° as that is a natural consequence of the EHC (see Figure 6.6). For each of the five test positions, the number of subtracted counts is consistent with no source, as expected.

6.3.2 POINT SOURCE WITH BACKGROUND

The natural next step to verify the CDS background subtraction technique is to simulate a point source with background and ensure that we recover the simulated spectral parameters. We first performed this analysis with a simulation of GRB 160530A (see Section 5.3 for a description of COSI’s observation

6.3 SPECTRAL BACKGROUND SUBTRACTION VALIDATION

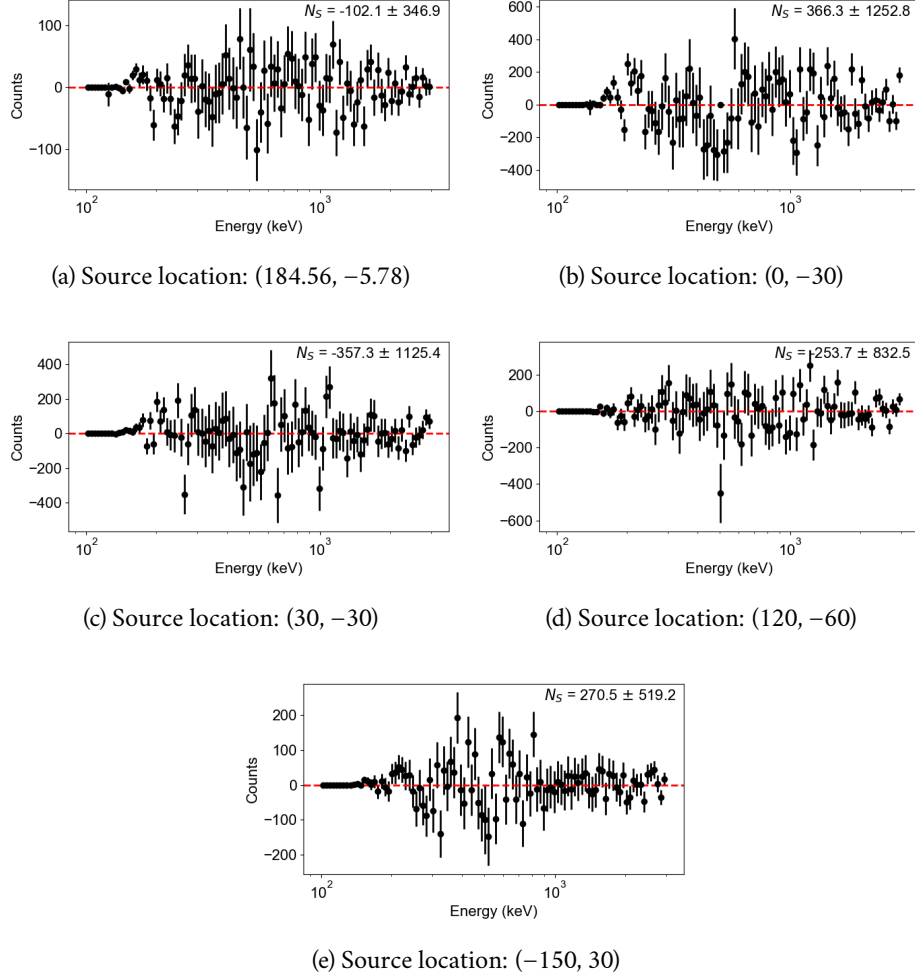


Figure 6.7: The results of the background subtraction performed on a simulation of background only for a variety of “source” positions. The “source” positions were chosen by selecting positions that had good exposure during the 2016 flight from the exposure map, other than position (a) where the “source” position is the Crab position. The subtracted spectrum for each of the positions is consistent with no source. The N_S value written in the upper right hand corner of each plot is the total number of subtracted counts (i.e. an integral over the background-subtracted spectrum).

6.3 SPECTRAL BACKGROUND SUBTRACTION VALIDATION

Parameter	Allowed range
Photon energy	100 - 3000 keV
Compton scatter angle	0 - 65°
Number of interactions	2 - 6
Distance between first two interactions	≥ 0.5 cm
Distance between any two interactions	≥ 0 cm
Earth horizon cut	Reject events whose full cone circle intersects Earth
Pointing cut	$\leq 40^\circ$ of source position
CDS-ARM	$\leq 7.7^\circ$

Table 6.1: The event selections used for the background-subtracted spectra shown in Figure 6.7. See Section 2.4.2 for a description of the event selections.

of this GRB). We simulated GRB 160530A with a Band function (Band et al., 1993) using the best fit parameters from the Konus spectral fit of this GRB (Svinkin et al., 2016). We did not simulate the structure in the GRB light curve (see Figure 5.6), but rather simulated a flat light curve for the 38.9 second duration of the GRB. We ensured that the number of simulated counts matched the number of counts measured by COSI. We also simulated the atmospheric background concurrent with the GRB using the Ling Model.

The FWHM of the CDS-ARM distribution is $\Delta = 9.4^\circ$ for this particular point source simulation; recall that Δ determines the size of **SR** and **BR**. The event selections used for this analysis are detailed in Table 6.2. These event selections are the same selections determined in Section 5.3.2 using an optimization algorithm to optimize the signal significance, but with a CDS-ARM cut determined by the value of Δ for this particular simulation. Note that these event selections were used for both constructing the scaling matrix **M** as well as the spectrum.

We performed the background subtraction following the steps described in Section 6.2.3 to account for the tails of the CDS-ARM distribution of the point source that are present in **BR**. Note that when fitting the spectrum with a Band function, we froze the characteristic energy E_c and second power law index β to the known values, as we did in Section 5.3.2. Thus only the first power law index α and the flux were free parameters. The α and flux after each iteration are shown in Table 6.3. First we fit the spectrum with no CDS-ARM tail correction; the resulting flux is significantly lower than the simulated flux due to the oversubtraction of the background. After the first

6.3 SPECTRAL BACKGROUND SUBTRACTION VALIDATION

Parameter	Allowed range
Photon energy	100 - 3000 keV
Compton scatter angle	17.7 - 176.8°
Number of interactions	2 - 7
Distance between first two interactions	≥ 0.6 cm
Distance between any two interactions	≥ 0.14 cm
Earth horizon cut	None
Pointing cut	N/A
CDS-ARM	$\leq 9.4^\circ$

Table 6.2: The event selections for the spectral fits of the simulation of GRB 160530A shown in Figure 6.9. See Section 2.4.2 for a description of the event selections.

iteration, the flux is much closer to the simulated value. We continue through the procedure from Section 6.2.3 until the sixth iteration, at which point the parameters have converged: both parameters are equal to the values from the previous iteration within the required tolerance of 10^{-4} . To prove that the parameters have truly converged, we continued the process for a total of 20 iterations. Figure 6.8 shows the values for the α and the flux evolve over the course of the 20 iterations; it is clear that the parameters have converged.

Figure 6.9 shows the fitted COSI spectrum of the simulation of GRB 160530A using the CDS background subtraction method and the best fit parameters. The best fit parameters are consistent with the simulated parameters, proving that the CDS background subtraction method performed successfully for this simulation.

The simulation of GRB 160530A provides a good test case: because of the relatively short duration of the burst, the burst barely moves within COSI’s FOV. For a persistent source, the change of COSI’s FOV must be taken into account. To test this algorithm for persistent sources, we simulated a Crab-like source with a power law spectrum mimicking the Crab’s movements relative to COSI during the 2016 flight. We simulated the point source with two different fluxes, $0.0365 \text{ } \gamma/\text{cm}^2/\text{s}$ and $0.0183 \text{ } \gamma/\text{cm}^2/\text{s}$, to test the algorithm on a “strong” and a “weak” source. In both cases, the spectral shape of the simulation is a power law with a photon index of 2. In comparison, the parameters of the Crab spectrum as measured by Spectrometer on INTEGRAL (SPI) are a photon index of 2.169 ± 0.008 and a power law normalization of $14.44 \pm 0.44 \text{ } \gamma/\text{cm}^2/\text{s}/\text{keV}$ (Sizun et al., 2004), which gives a 0.2-3 MeV flux of 0.0330

6.3 SPECTRAL BACKGROUND SUBTRACTION VALIDATION

Iteration	α	Deviation in α (%)	0.1 - 3 MeV Flux ($\gamma/\text{cm}^2/\text{s}$)	Deviation in flux (%)
No correction	$-0.7657^{+0.6591}_{-0.7390}$	-17.67	$1.2549^{+0.4091}_{-1.2549}$	-42.96
1 st iteration	$-1.0522^{+0.3305}_{-0.3485}$	13.14	$2.9216^{+4604}_{-0.7686}$	32.8
2 nd iteration	$-1.0973^{+0.3327}_{-0.3535}$	17.99	$3.0130^{+0.4740}_{-0.8450}$	36.95
3 rd iteration	$-1.1082^{+0.3321}_{-0.3529}$	19.16	$3.0287^{+0.4743}_{-0.7917}$	37.67
4 th iteration	$-1.1115^{+0.3533}_{-0.3326}$	19.52	$3.0353^{+0.4757}_{-0.8063}$	37.97
5 th iteration	$-1.1112^{+0.3326}_{-0.3533}$	19.48	$3.0353^{+0.4647}_{-0.8503}$	37.97
6 th iteration	$-1.1112^{+0.3326}_{-0.3533}$	19.48	$3.0353^{+0.4647}_{-0.8503}$	37.97
10 th iteration	$-1.1112^{+0.3326}_{-0.3533}$	19.48	$3.0353^{+0.4647}_{-0.8503}$	37.97
20 th iteration	$-1.1112^{+0.3326}_{-0.3533}$	19.48	$3.0353^{+0.4647}_{-0.8503}$	37.97
Simulated	-0.93		2.2000	

Table 6.3: The iterations of the algorithm to correct for the source counts present in **BR** due to the tails of the **CDS-ARM** distribution of a point source, applied to the simulation of **GRB** 160530A. Only the values for the first power law index α and the flux are allowed to vary, while β and E_c are fixed to the simulated values of -3.5 and 638 keV, respectively. For each iteration, we show the best fit α and 0.1 - 3 MeV flux values, as well as their relative deviation from the simulated parameters $(P_{\text{fit}} - P_{\text{sim}})/P_{\text{sim}} \times 100$, where P is the parameter in question. First we fit the subtracted spectrum without applying any correction. To perform the first iteration of the algorithm, we subtracted a spectrum with S_{BR} counts and $\alpha = -0.7657$ from the background spectrum and fit the new background-subtracted spectrum. The parameters converged after six iterations. The final converged parameters, $\alpha = -1.1112^{+0.3326}_{-0.3533}$ and 0.1 - 3 MeV flux = $3.0353^{+0.4647}_{-0.8503} \gamma/\text{cm}^2/\text{s}$, are consistent with the simulated parameters. To prove that the parameters have converged, we continued the analysis out to 20 iterations and show the fit parameters from the 10th and 20th iterations. The uncertainties on the fit are to 99% confidence.

6.3 SPECTRAL BACKGROUND SUBTRACTION VALIDATION




Figure 6.8: The values of α and the 0.1 - 3 MeV flux as a function of the number of iterations performed. The dashed red lines indicate the simulated values: $\alpha = -0.93$ and 0.1-3 MeV flux = $2.2 \gamma/\text{cm}^2/\text{s}$. The parameters converge after six iterations.

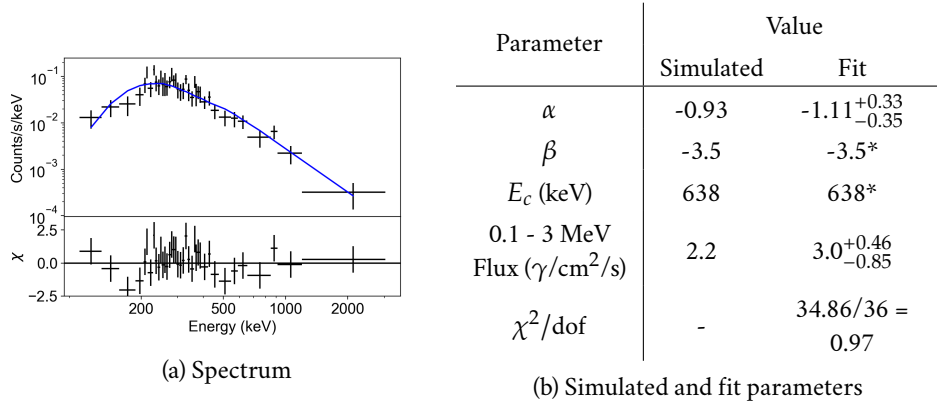


Figure 6.9: The final results of the background subtraction and spectral fit of the simulation of GRB 160530A. The spectrum is binned such that each bin has a signal-to-noise ratio $S/N = 3$. A comparison of the simulated and best fit parameters is shown in (b); the fit results are consistent with the simulated values. The uncertainties on the fit are 99% confidence.

$\gamma/\text{cm}^2/\text{s}$. We included the full flight Ling Model background simulation to these point source simulations. These Crab-like simulations are background-dominated: the counts from the strong source make up 6.3% of the total counts and the counts from the weak source make up 2.9% of the total counts. In contrast, the source counts from the simulation of GRB 160530A make up 44% of the total counts.

The FWHM of the CDS-ARM distribution is $\Delta = 7.7^\circ$ for these point source simulations. The event selections used for this analysis are detailed in Table 6.4; these selections were used to construct both the scaling matrix \mathbf{M} and the spectra. Note that these selections differ from those used to fit the simulated spectrum of GRB 160530A, particularly in the Compton scatter angle range. Since the Crab-like simulations are background-dominated, using the EHC can significantly reduce the background, whereas in the case of the GRB, the source is so bright that the EHC is unnecessary. As shown in Figure 6.6, using Compton scatter angles between 0° and 65° is a natural choice when the EHC has been applied; the events within the range of larger Compton scatter angles in Figure 6.6 are most likely backscatter events or events that come from the bottom of the detector since they are not rejected by the EHC, and thus either contribute to the background or are likely not completely absorbed in the detector.

Again we performed the background subtraction following the steps outlined in Section 6.2.3 to account for the source counts present in BR. Seven iterations were required for the strong simulation, and six were required for the weak simulation. Figure 6.10 shows the fitted background-subtracted spectrum and best fit parameters for the strong source simulation, and Figure 6.11 shows the fitted background-subtracted spectrum and best fit parameters for the weak source simulation. In both cases, the fit parameters are consistent with the measured parameters, indicating that the CDS background subtraction method is successful when applied to simulations of point sources that move within COSI's FOV and are highly background-dominated.

6.3.3 DEPENDENCE OF THE SCALE FACTOR

In Section 6.2.2, we stated that the scale factor f is dependent on the source position in detector coordinates (Θ, Φ) , the measured energy E , and the initial Compton scatter angle ϕ : $f = f(\Theta, \Phi, E, \phi)$. Here we will use the above simulations to explore these dependencies and determine an acceptable bin size for each variable.

6.3 SPECTRAL BACKGROUND SUBTRACTION VALIDATION

Parameter	Allowed range
Photon energy	200 - 3000 keV
Compton scatter angle	0 - 65°
Number of interactions	2 - 6
Distance between first two interactions	≥ 0.5 cm
Distance between any two interactions	≥ 0 cm
Earth horizon cut	Reject events whose full cone circle intersects Earth
Pointing cut	$\leq 40^\circ$ of Crab position
CDS-ARM	$\leq 7.7^\circ$

Table 6.4: The event selections used in the background subtraction and spectral fit of the Crab-like source simulation. See Section 2.4.2 for a description of the event selections.




Figure 6.10: The results of the background subtraction and spectral fit of the simulation of the “strong” Crab-like source. The spectrum is binned such that the signal-to-noise ratio in each bin is $S/N = 3$. A comparison of the simulated and best fit parameters is shown in (b); the fit results are consistent with the simulated values. The uncertainties on the fit are 99% confidence.

6.3 SPECTRAL BACKGROUND SUBTRACTION VALIDATION




Figure 6.11: The results of the background subtraction and spectral fit of the simulation of the “weak” Crab-like source. The spectrum is binned such that the signal-to-noise ratio in each bin is $S/N = 2$. A comparison of the simulated and best fit parameters is shown in (b); the fit results are consistent with the simulated values. The uncertainties on the fit are 99% confidence.

Figure 6.12 shows two background-subtracted spectra from two different “source” positions of a simulation of background only. When performing this background subtraction, we ignored all dependence on Θ , Φ , E , and ϕ from f ; in other words, we used one very large bin for each variable on which f depends. The spectra are clearly not flat, confirming that f does in fact depend on these variables.

To determine the required binning of f (and of \mathbf{M}) to achieve the desired results, we tested a variety of bin sizes on the background-only simulations and the strong Crab-like point source simulation described in the previous section. We decreased the bin size for each variable and applied the CDS background subtraction method until the background-only simulations resulted in a flat spectrum and we recovered the simulated parameters of the strong Crab-like point source. Further decreasing the bin size did not have an effect on the results. For the previous analysis and any future analysis using this method in this work, we use the following bin sizes: 2° zenith angle bins, 10° azimuth angle bins, 10 logarithmically spaced energy bins between 100 keV and 3 MeV, and 5° Compton scatter angle bins.

We note that the simulations used to make the scaling matrix \mathbf{M} , from which we calculate f , must be simulations of the same spectral shape as the background data that we are trying to subtract. We initially made \mathbf{M} with a flat spectrum where the flux is constant as a function of energy and tested the

6.4 BACKGROUND SUBTRACTION APPLIED TO FLIGHT DATA




Figure 6.12: The results of the background subtraction of the simulation of background only for two different “source” positions when the dependence of f on Θ , Φ , E , and ϕ is ignored. The resulting spectra are not consistent with a flat spectrum at zero. The N_S value in the upper right hand corner of the plots is the total number of subtracted counts.

background subtraction method with the background simulations made from the Ling Model. We did not recover the subtracted spectra that we expected. For all background subtractions using \mathbf{M} previously shown in this chapter, \mathbf{M} was made with a Ling Model simulation. This allowed us to recover the expected background-subtracted spectral shape.

6.4 BACKGROUND SUBTRACTION APPLIED TO FLIGHT DATA

Two potential difficulties arise when using the CDS background subtraction algorithm with COSI flight data. First, we must ensure that the shape of the spectrum from the Ling Model simulations is close enough to the shape of the measured background spectrum such that we can use the scaling matrix \mathbf{M} computed from simulations to scale the measured background. Second, discrepancies between the shape of the measured and simulated CDS-ARM distribution may have a non-negligible impact on r_{SR} and r_{BR} , required to perform the CDS-ARM tail correction. In this section, we will explore these two potential difficulties and verify the background subtraction algorithm with the COSI observation of GRB 160530A.

6.4.1 MEASURED BACKGROUND COMPARED TO THE LING MODEL

We compared the shape of the measured background spectrum to the spectrum from the Ling Model simulations by fitting the spectra with a phenomenological broken power law model and comparing the resulting parameters. The broken power law model is defined as:

$$f(E) = \begin{cases} KE^{-\Gamma_1} & E \leq E_{\text{break}} \\ KE_{\text{break}}^{\Gamma_2 - \Gamma_1} E^{\Gamma_2} & E > E_{\text{break}} \end{cases} \quad (6.21)$$

where K is a normalization constant. This model was fit to the raw [COSI](#) data for each day of the flight using least-squares fitting. We excluded the energy range surrounding the 511 keV atmospheric background line, 490 - 530 keV, in these fits.

Figure 6.13 shows the three parameters of the broken power law, the first and second power law indices Γ_1 and Γ_2 and the break energy E_{break} , for the background during each day of the [COSI](#) 2016 flight and for a Ling Model simulation. The parameters from the simulation are well within the spread of the parameters from the different days of the flight, indicating that the Ling Model is a good approximation of the observed atmospheric background.

6.4.2 SIMULATED AND MEASURED CDS-ARM DISTRIBUTIONS

We verified whether the simulated [CDS-ARM](#) distribution closely matches the measured [CDS-ARM](#) distribution by comparing the simulated and measured distributions of [GRB 160530A](#). Ideally, the simulations match the measurements closely enough that the simulated [CDS-ARM](#) distribution can be used to estimate r_{SR} and r_{BR} . Note that for most of the sources detected during [COSI](#)'s 2016 flight, we cannot determine the shape of the [CDS-ARM](#) distribution without a background subtraction because the observations are heavily background-dominated. Since [GRB 160530A](#) is a transient source, however, we are able to use the time-domain background subtraction discussed in Section 5.3.2 to acquire a measured, background-subtracted [CDS-ARM](#) distribution.

To perform the background subtraction in the time domain, we chose background events from a 573-second period that spans before and after the [GRB](#) while leaving a 100 second gap on both sides of the [GRB](#); the [GRB](#) itself lasts 38.9 seconds. See Figure 5.6 for the [GRB](#) lightcurve with the times used for

6.4 BACKGROUND SUBTRACTION APPLIED TO FLIGHT DATA

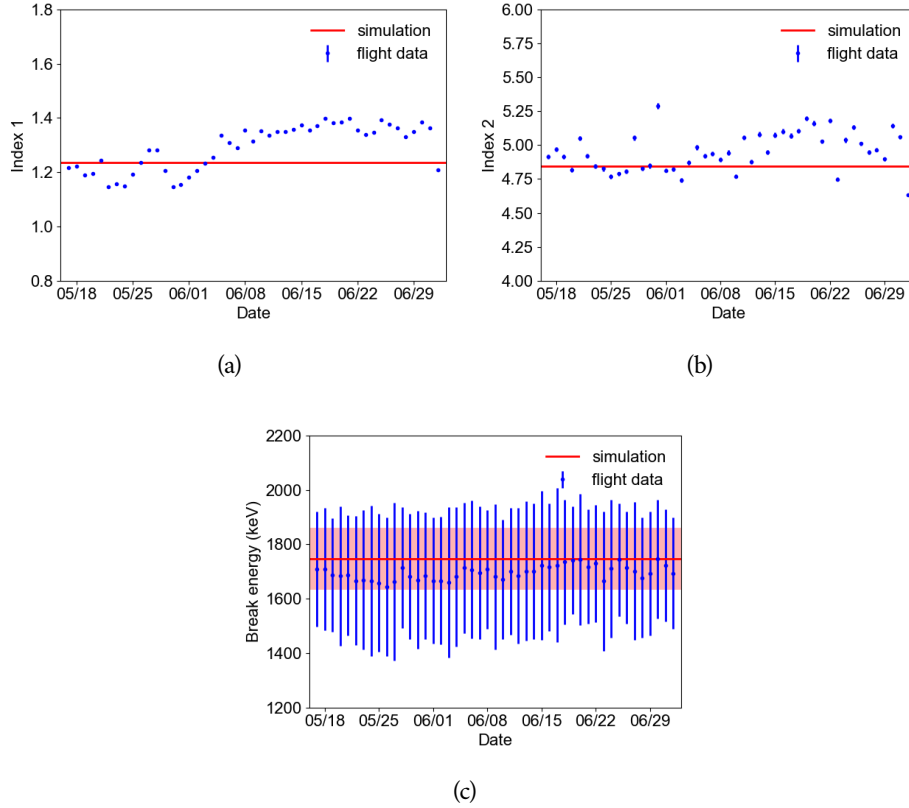


Figure 6.13: A comparison of (a) the first power law index Γ_1 , (b) the second power law index Γ_2 , and (c) the break energy E_{break} for each day of flight and the Ling Model simulations. The red line represents the best fit value for the Ling Model simulation, and the red shaded region represents the error on the fit. The best fit parameters of the Ling Model simulations fall within the spread of the flight data, indicating that the Ling Model is a good approximation of the atmospheric background.




Figure 6.14: The (a) measured and (b) simulated [CDS-ARM](#) distributions of [GRB 160530A](#). The measured distribution is significantly wider.

source and background events marked. To get the background-subtracted [CDS-ARM](#) distribution of [GRB 160530A](#), we make a [CDS-ARM](#) distribution from the background events (the events before and after the [GRB](#)), scale that by the ratio in observation times (38.9 s/573 s), and subtract the scaled background distribution from the [CDS-ARM](#) distribution of events that occurred during the [GRB](#) time. As we can see in [Figure 5.6](#), the count rate stays constant over the background period.

The simulation of [GRB 160530A](#) lasts for 38.9 s and has an input flux chosen such that the number of simulated counts is equal to the number of background-subtracted measured counts. The spectral shape of the simulation is a Band function with parameters from [Svinkin et al. \(2016\)](#). We added background to the simulation using the Ling Model, and performed the same time-domain background subtraction on the simulated [CDS-ARM](#) distribution.

[Figure 6.14](#) shows a comparison of the measured and simulated [CDS-ARM](#) distribution. From a visual comparison, it is clear that the measured distribution is wider. The [FWHM](#) of the measured distribution is $\Delta = 10.6^\circ$, whereas the [FWHM](#) of the simulated distribution is $\Delta = 6.4^\circ$. While for strong transient sources we can directly determine the measured [CDS-ARM](#) distribution, persistent broadband sources involve a more careful treatment. We investigate possible causes for the discrepancy in the simulated and measured [CDS-ARM](#) distributions in the following sections. We find, however, that our treatment of the background is limited in this regard.




Figure 6.15: A schematic of a cross section of the simulation setup for atmospheric absorption, not drawn to scale. The height of the atmosphere is 200 km, and each 1 km layer (not drawn) has a composition determined from the NRLMSISE-00 atmosphere model. The blue spherical detectors span the entire plane at COSI’s altitude during the observation. The red source is positioned such that a direct beam from the source intersects the center sphere at (Θ, Φ) , or the GRB position in detector coordinates.

6.4.2.1 ATMOSPHERIC SCATTERING

One potential explanation that is not included in the previously described simulation setup is that γ -rays from the source scatter in the atmosphere before being collected by COSI. This scattering could cause γ -rays from the GRB to appear as if they are coming from elsewhere in the sky, which would broaden the measured CDS-ARM distribution. To determine if this explanation can account for the observed width of the measured CDS-ARM distribution, we included atmospheric scattering into our simulation of GRB 160530A.

We approximate the atmosphere as a rectangular prism with a height of 200 km, spaced by 1 km thick layers. Each layer is composed of nitrogen, oxygen, and argon. The relative quantities of these elements in addition to the total densities of the layers are determined from the NRLMSISE-00 atmosphere model (Picone et al., 2002). An array of spherical detectors fills the plane at COSI’s altitude; these detectors are used to determine from which directions scattered source photons approach the instrument. A source is placed at the top of the atmosphere and oriented such that a beam traveling directly from the source to the center sphere intersects that sphere at the GRB’s position relative to COSI (e.g. at detector coordinates $\Theta = 43.5^\circ$, $\Phi = -66.1^\circ$). See Figure 6.15 for a schematic of this setup. To avoid potential edge effects, the rectangular atmosphere is significantly wider in both the x and y directions than the lateral distance of the source from the center.

6.4 BACKGROUND SUBTRACTION APPLIED TO FLIGHT DATA




Figure 6.16: The simulated **CDS-ARM** distribution with atmospheric scattering included in the simulations. Incorporating the atmospheric scattering barely changes the central part of the simulated **CDS-ARM** distribution, indicating that the impact of the atmospheric scattering is only marginal.

The source emits photons in the shape of the **GRB** spectrum, with the same model we used in Section 6.4.2. Now, the emitted photons can interact with the layers of atmosphere above and below the spheres. Each time a photon enters any of the spheres, its position and direction relative to the center of the sphere as well as its energy are recorded. If a photon enters multiple spheres, the information from only one of the entry points, chosen at random, is recorded. This information is used to generate an event list to use as input for a second simulation. In the second simulation, each photon from the event list enters the sphere surrounding the regular **COSI** mass model at the recorded position and direction with the recorded energy. Thus, the simulated photons do not necessarily travel on a direct path from the source to **COSI** and can be scattered by the atmosphere, allowing some fraction of the photons to appear as though they were emitted from a different location.

With the simulation of **GRB** 160530A that includes atmospheric scattering, we made a **CDS-ARM** distribution, shown in Figure 6.16. The distribution is still narrow with a large central peak, contrary to the measured **CDS-ARM** distribution shown in Figure 6.14. Including the atmospheric scattering in the simulations has only a marginal impact on the simulated **CDS-ARM** distribution.

6.4.2.2 SHIFT IN THE GRB POSITION

Since **GRB** 160530A was far off-axis (43.5°), it is possible that the image deconvolution algorithm shifted the hotspot in the image a few degrees from the

true position. This shift could potentially explain the broad CDS-ARM distribution: if we are using the incorrect GRB position, the GRB events will not be centered around the ideal line in the CDS, and so a histogram of the distances of each event to the ideal line (the CDS-ARM distribution) would be broader than expected. With an incorrect position for the GRB, the best events that should lie very close to the ideal line would instead be somewhat scattered throughout the data space, causing a broadened distribution.

The best known localization of this GRB is from the COSI image itself (Lowell et al., 2017a), so we performed a systematic search of best fit positions. For this purpose, we compare the background-subtracted CDS-ARM distribution calculated with respect to a variety of potential source positions arranged in a grid around the center of the image hotspot. The center of the image hotspot for GRB 160530A is at $l = 243.4^\circ$, $b = 0.4^\circ$. We used a grid of potential source positions ranging from the center of the image hotspot plus or minus 20° in both directions with 2° steps, or $l = [-223.4^\circ, -225.4^\circ, \dots, 261.4^\circ, 263.4^\circ]$ and $b = [-19.6^\circ, -17.6^\circ, \dots, 18.4^\circ, 20.4^\circ]$. For each potential source position, we determined the number of counts in the very center of the CDS-ARM distribution by integrating the distribution between -2° and 2° . With this integral over the center of the CDS-ARM distribution, we can establish the peakiness of the distribution. If the broadening of the distribution is due to using the wrong source position, the distribution should narrow and contain more counts in the very center when we calculate it from the true position.

Figure 6.17 shows the results of this analysis for both the GRB data and the simulation. The results are displayed in a 2D histogram, in which each cell contains the number of counts in the CDS-ARM distribution peak corresponding to that offset position. The histograms were then smoothed. In the case of the simulation, the 2D histogram clearly resembles an image of a point source. In the case of the data, the brightest point in the histogram is close to the center, but its relative brightness compared to the surrounding positions is significantly less. Note the difference in the z axis scales of both images, and recall that the measured and simulated observations have the same total number of counts: this is another way of showing that the measured CDS-ARM distribution is relatively broad. We find that the highest point on the grid is consistent with the previously ascertained position of the GRB. Otherwise, we would expect to see a result similar to the simulation result, with the bright spot at a position other than (0,0).

6.4 BACKGROUND SUBTRACTION APPLIED TO FLIGHT DATA




Figure 6.17: A 2D histogram in which each cell contains the number of counts in the **CDS-ARM** distribution peak (from -2° to 2°) corresponding to the source position given by the offset in question. An offset of (0,0), corresponds to the hotspot of the image of **GRB 160530A**, or $l = 243.4^\circ$, $b = 0.4^\circ$ (theta refers to offset in Galactic latitude b and phi refers to offset in Galactic longitude l). If the hotspot in the image of the **GRB** data were shifted due to the image deconvolution, we would expect to see a result similar to the simulation result with the bright spot shifted from (0,0).

6.4.3 MOVING FORWARD WITH THE SIMULATED CDS-ARM DISTRIBUTION

The physical mechanisms discussed above only have marginal influence on the width of the simulated **CDS-ARM** distribution. To precisely determine r_{SR} and r_{BR} using simulations, we empirically include additional broadening with a Gaussian convolution of the simulated **CDS-ARM** distribution. With this empirical approach, it is possible to determine r_{SR} and r_{BR} for other, particularly non-transient, sources.

To determine the width of the Gaussian with which we convolve the simulated **CDS-ARM** distribution, we use the data from **GRB 160530A**. First, we fit the **CDS-ARM** distribution from the simulation of **GRB 160530A** with the **ARM** fit function: a double Lorentzian and a Gaussian. We convolve the fitted **ARM** function with a Gaussian to broaden the distribution, and fix the parameters of the **ARM** function. We then fit the convolved function to the measured **CDS-ARM** distribution of **GRB 160530A** to determine the width of the Gaussian, which is $\sigma = 3.9^\circ$. To approximate the measured **CDS-ARM** distribution of another source, we fit the simulated **CDS-ARM** distribution of that source and

6.4 BACKGROUND SUBTRACTION APPLIED TO FLIGHT DATA




Figure 6.18: (a) The simulated CDS-ARM distribution fit with the ARM function. (b) The measured CDS-ARM distribution fit with the ARM function convolved with a Gaussian. The data points are drawn in blue and the fit is drawn in red.

then convolve the fitted ARM function with a Gaussian with $\sigma = 3.9^\circ$. Then we use this convolved function to estimate r_{SR} and r_{BR} . Figure 6.18 shows the fitted simulated CDS-ARM distribution and the ARM function convolved with a Gaussian fit to the data of GRB 160530A.

Note that the difference in widths between the simulated and measured CDS-ARM distributions may vary based on certain parameters of the source, such as the spectral shape or the location in COSI's FOV. Since we only detected one source during the 2016 flight with an attainable measured CDS-ARM distribution, determining any such dependencies is difficult. Thus we will continue under the assumption that the width of the convolution Gaussian determined from GRB 160530A can be used for all point sources detected by COSI.

6.4.4 CDS BACKGROUND SUBTRACTION OF GRB 160530A

Here we will use the CDS background subtraction algorithm to subtract the background from COSI's observation of GRB 160530A, and fit the resulting background-subtracted spectrum. We constructed the scaling matrix \mathbf{M} using a Ling Model simulation. We determined $\Delta = 21.4^\circ$ as well as r_{SR} and r_{BR} from the best fit function of the simulated CDS-ARM distribution convolved with a Gaussian, as described in the previous section. We used the optimized




Figure 6.19: The [COSI](#) spectrum of [GRB 160530A](#), with the background subtraction performed using the [CDS](#) method. The spectrum was fit with the atmospheric absorption model multiplied by the Band function, and was binned such that each bin had the same signal-to-noise ratio $S/N = 3$. See Table 6.5 for the best fit parameters.

event selections for this analysis that we determined in Section 5.3.2, detailed in Table 5.2, except with a [CDS-ARM](#) cut of $\leq 21.4^\circ$.

As we did in Section 5.3.2, we fit the background-subtracted spectrum with a Band function ([Band et al., 1993](#)) and fixed the second power law index and the break energy to the best fit values according to Konus. The fit spectrum is shown in Figure 6.19. Table 6.5 compares the best fit parameters found with the [CDS](#) background subtraction method to those found with the time-domain background subtraction method, and also to the best fit parameters reported by the Konus team ([Svinkin et al., 2016](#)). The best fit [COSI](#) parameters found with the [CDS](#) background subtraction method are consistent with the other two spectral fits, indicating that the [CDS](#) background subtraction method works for [COSI](#) flight data as long as the shape of the [CDS-ARM](#) distribution can be estimated.

6.5 CONCLUSIONS

We have developed the [CDS](#) background subtraction method described in this chapter particularly for broadband, persistent sources. We have verified this method using simulations as well as the [COSI](#) data of [GRB 160530A](#). [Kierans \(2018\)](#) developed a background subtraction method using the [CDS](#) for line sources, in which the high energy range of the spectrum is used to scale the background in place of a scaling matrix. Though the [Kierans \(2018\)](#) method

Parameter	COSI value <i>CDS</i>	COSI value <i>time-domain</i>	Konus value
Atmospheric N_A (g/cm ²)	9.2*	9.2*	—
α	$-0.90^{+0.44}_{-0.46}$	$-1.14^{+0.30}_{-0.32}$	-0.93 ± 0.03
β	-3.5*	-3.5*	<-3.5
E_c (keV)	638*	638*	638^{+36}_{-33}
0.2-10 MeV Fluence (erg/cm ²)	$1.58 \pm 0.31 \times 10^{-4}$	$1.49^{+0.23}_{-0.21} \times 10^{-4}$	$1.3 \pm 0.04 \times 10^{-4}$
χ^2/dof	17.67/22 = 0.93	27.54/30 = 0.92	73/73 = 1

Table 6.5: The COSI best fit parameters of the spectrum of GRB 160530A, using the CDS background subtraction method and the time-domain background subtraction method from Chapter 5. The Konus best fit parameters of GRB 160530A from Svinkin et al. (2016) are shown for comparison. The uncertainties on the COSI parameters are 99% confidence. The COSI parameters from the CDS background subtraction method are consistent with those from the other fits, confirming that this background subtraction method works for COSI flight data.

requires less computation time since there is no scaling matrix to compute, the method detailed in this chapter is more general and can be applied to both line and continuum sources.

THE COSI OBSERVATION OF THE CRAB NEBULA

The Crab nebula is a pulsar wind nebula (PWN) surrounding the Crab pulsar within the remnant of a supernova that occurred in 1054. Because the Crab pulsar and nebula are extremely bright in X-rays and γ -rays, the system has been studied by numerous telescopes over the years. Most projected flight paths from Wanaka, New Zealand, did not put COSI far enough north during the 2016 flight to observe the Crab, a northern hemisphere source. However, the balloon drifted farther north than expected (see Figure 4.2), allowing for some exposure of the northern skies, and we were able to detect the Crab in the COSI data.

7.1 OBSERVATION OVERVIEW

Figure 7.1 shows the zenith angle of the Crab in detector coordinates for the entire COSI 2016 flight, in which $\Theta = 0^\circ$ indicates that the source is directly on-axis. On June 12 and June 13, the Crab reached the highest zenith angles of the 2016 flight. Figure 7.2 shows the COSI image of the Crab using the data from those two days. The event selections used to create this image are detailed in Table 7.1, and were found by trial and error (it is difficult to numerically describe what makes the best, clearest image, preventing us from using an optimization algorithm to determine the best event selections). Including the other days in which the Crab was briefly at zenith angles $\Theta < 40^\circ$ does not enhance the detection of the Crab in the image, but rather degrades the image clarity. Thus we will use the data from June 12 and 13 for the analysis going forward. Applying a 40° pointing cut around the Crab position to the data from June 12 and 13 results in 24.61 ks of exposure time.

During the observation of the Crab, COSI went through a relatively small altitude range from 33.4 - 33.6 km. These relatively high altitudes (see Figure 4.3) allow for the best atmospheric transmission probability we can expect during a balloon flight, which is $\sim 49\%$ at 511 keV. By this time in the flight,




Figure 7.1: The zenith angle of the Crab in detector coordinates throughout the 2016 flight, where a zenith angle of $\Theta = 0^\circ$ corresponds to an on-axis source. The dashed red line marks a zenith angle of 40° , used as a cutoff point to determine periods in which the Crab was in COSI's FOV. There are 12 days in which the Crab was briefly above $\Theta = 40^\circ$.




Figure 7.2: The COSI image of the Crab using data from June 12 and 13. To make this image, we used the event selections detailed in Table 7.1 and performed 35 iterations of LM-MLEM algorithm. The color scale is in arbitrary units.

Parameter	Allowed range
Photon energy	0 - 480 keV, 530 - 1500 keV
Compton scatter angle	0 - 30°
Number of interactions	3-7
Distance between first two interactions	≥ 0.5 cm
Distance between any two interactions	≥ 0 cm
Earth horizon cut	Reject events whose full cone circle intersects Earth
Pointing cut	$\leq 70^\circ$ of source position

Table 7.1: The event selections used to make the image of the Crab shown in Figure 7.2. See Section 2.4.2 for a description of the event selections.

however, three detectors had failed, limiting COSI’s detection efficiency (see Section 4.5 for more details). Throughout the observation, the Crab was fairly stationary relative to the instrument; the source position in detector coordinates ranged from a zenith of 28° to 40° and from an azimuth of 3° to 12°.

To ensure that the hotspot in the Crab image is a detection and not just an artifact of the image deconvolution algorithm, we used the angular resolution measure (ARM) distribution of the observation (see Section 2.1.2 for an explanation of the ARM distribution). Since the observation is severely background-dominated as is typical for a balloon-borne γ -ray telescope, the Crab is not immediately visible above the background in the ARM distribution. Instead, we fit the ARM distribution with a single Gaussian and the sum of two Gaussians and compare the results. A Gaussian can approximately describe the shape of an ARM distribution consisting solely of background events. For example, Figure 7.3 shows the ARM distribution calculated with respect to $(l = 210^\circ, b = -30^\circ)$, a background location in COSI’s FOV, fit with a Gaussian. The reduced χ^2 of the fit is $\chi^2/\text{dof} = 143.5/57 = 2.5$, indicating that a Gaussian is a reasonable approximation to the shape of the background ARM distribution. If a source were present, this would not be the case; rather, the sum of two Gaussians would be a better fit to the ARM distribution than a single Gaussian, since the second Gaussian would fit the excess that arises from the source events.

We first fit the ARM distribution from the Crab observation with a single Gaussian, testing the hypothesis that no source is present. The fitted ARM distribution is shown in Figure 7.4a and the resulting reduced χ^2 is $\chi^2/\text{dof} = 628.8/41 = 15.3$, implying that a single Gaussian is not a good representation




Figure 7.3: The **ARM** distribution of a background observation from the 2016 flight fit with a single Gaussian. The **ARM** distribution was calculated with respect to $(l = 210^\circ, b = -30^\circ)$, a background location in **COSI's FOV**, and we used a 40° pointing cut around $(l = 210^\circ, b = -30^\circ)$ to generate the event list. The Gaussian fit is a good approximation of the data, with a $\chi^2/\text{dof} = 2.5$.




Figure 7.4: The **ARM** distribution from the Crab observation fit (a) with a single Gaussian and (b) with the sum of two Gaussians. The sum of two Gaussians clearly fits the data better, with a $\chi^2/\text{dof} = 1.4$ in comparison to the fit of a single Gaussian which results in a $\chi^2/\text{dof} = 15.3$.

Parameter	Allowed range
Photon energy	250 - 480 keV, 530 - 1500 keV
Compton scatter angle	0 - 30°
Number of interactions	2-7
Distance between first two interactions	≥ 0.5 cm
Distance between any two interactions	≥ 0 cm
Earth horizon cut	Reject events whose full cone circle intersects Earth
Pointing cut	$\leq 40^\circ$ of source position

Table 7.2: The event selections used to generate the ARM distributions in Figures 7.3 and 7.4. See Section 2.4.2 for a description of the event selections.

of the data. We then fit the ARM distribution with the sum of two Gaussians to test the hypothesis that there is a source, shown in Figure 7.4b. With a reduced χ^2 of $\chi^2/\text{dof} = 52.1/38 = 1.4$, the sum of two Gaussians is clearly a better fit to the data than a single Gaussian. We conclude that we have detected the Crab in the COSI dataset.

The event selections used in this analysis are documented in Table 7.2. The ARM distributions of the observations do not span from -180° to 180° because of the Earth horizon cut (EHC) and pointing cut; these event cuts enhance any already existing exposure effects. We fit the ARM distributions between -60° and 60° because we are most interested in how the distribution behaves around 0° where an excess of source counts should appear. Additionally, the fits to the data were improved when this range was used.

7.2 SPECTRAL ANALYSIS

The first step in performing spectral analysis of the Crab observation is to determine the best event selections to use for the analysis. Ideally, we would optimize the signal significance $S = r_S / \sqrt{r_S + r_B}$ where r_S and r_B are the source and background count rates, respectively. Using an optimization algorithm, however, becomes impractically time consuming when we use the COMPTEL Data Space (CDS) background subtraction method for broadband, persistent sources described in Chapter 6, because we must compute a new scaling matrix for each set of event selections. Instead, we search through a grid of event selections to determine the best ones for the analysis.




Figure 7.5: A schematic of a Compton event with the largest possible Compton scatter angle if the [EHC](#) and a 40° pointing cut are applied. The source at the yellow star is at a zenith angle of 40° and emits a photon that interacts in the detector at \vec{r}_1 and then at \vec{r}_2 . The red Compton circle extends to just above the horizon; if it extended any lower, the event would be rejected by the [EHC](#). The resulting maximum Compton scatter angle, 65° , is half of the angle between the horizon and the initial photon direction, which is 130° .

The [EHC](#) already restricts the Compton scatter angle range to angles between 0° and 65° . Figure 7.5 shows a Compton event emitted by a source that is 40° off axis, which is the largest zenith angle allowed by the 40° pointing cut. The particular Compton event shown scatters with the largest possible Compton scatter angle, as any larger Compton scatters will lead to a section of the red event circle dipping below the horizon and rejection by the [EHC](#). The angle from the initial source direction to the horizon is $40^\circ + 90^\circ = 130^\circ$, half of which (65°) is the Compton scatter angle for this particular event. Thus, 65° is the largest Compton scatter angle that we can measure with the pointing cut and [EHC](#).

With this restriction on the Compton scatter angles, we choose three potential Compton scatter angle ranges for the grid of event selections: the full range of 0° - 65° and the low and high ranges of 0° - 30° and 30° - 65° . We choose two values for the lower limit of the distance between the first two interactions, 0.5 cm and 1 cm, and two values for the lower limit of the distance between any two interactions, 0 cm and 0.3 cm. Events with distances that are very close together are more difficult to reconstruct, but a large distance cut will exclude more events. We also try both including and excluding two-site events, or events that only interact twice in the detector; these events are more difficult to reconstruct but also comprise about half of the total events (we note that excluding two-site events from the Crab image significantly improved the image quality). In all cases, we used the energy range of 0-480 keV and 530-1500 keV to cut out the 511 keV atmospheric background line. Note that this is the energy range that yields the highest quality Crab image.

For each set of event selections, we generated the scaling matrix \mathbf{M} as well as a source and background event list. To determine the region size Δ , we used the procedure described in Section 6.4.3, in which we convolve the simulated [CDS-ARM](#) distribution with a Gaussian. The [CDS-ARM](#) distribution from a simulation of the Crab has a width of 6.6° . When that distribution is convolved with a Gaussian, the resulting width is 17.6° , so we proceed with $\Delta = 17.6^\circ$. We performed the initial background subtraction according to Equation 6.12, and the results are shown in Table 7.3. Recall that this initial subtraction is an oversubtraction of the background because of the tailing present in the [CDS-ARM](#) distribution of a point source. We correct for this oversubtraction and fit the spectrum together in an iterative process (see Section 6.2.3 for the detailed procedure). However, the initial subtraction resulted in a negative number of counts for all but one set of event selections.

With negative counts, we cannot correct for the CDS-ARM tails as outlined in Section 6.2.3 nor can we fit the spectrum.

With a Compton scatter angle range of 30° - 65° , the distance between the first two interactions ≥ 1 cm, the distance between any two interactions ≥ 0 cm, and 2-7 interactions, we get 43 ± 534 counts after the initial subtraction. With these event selections, we attempted to fit the spectrum with a power law model multiplied by the atmospheric absorption model discussed in Section 5.2, and correct for the CDS-ARM tails according to the procedure in Section 6.2.3. Unfortunately, with so few counts we were unable to constrain the fit parameters.

The results of this analysis indicate that we are not correctly estimating the background of the COSI observation of the Crab. While the CDS background subtraction method in described in Chapter 6 performed well on simulations of background-dominated sources and on the COSI data of gamma-ray burst (GRB) 160530A, a very strong source, the analysis described in this section marks our first attempt at using the CDS background subtraction algorithm on a severely background-dominated source detected by COSI. Evidently there is something about the flight background that we are not correctly estimating with the methods employed here. More accurate simulations of the flight background could potentially improve the results. Additionally, using the entire data space rather than a background region to estimate the background beneath the source is worth investigating in the future.

7.3 POLARIZATION ANALYSIS

Without a good background subtraction of the Crab observation, we cannot measure the polarization properties of the Crab as detected by COSI. We can, however, use simulations to determine whether we would have been able to measure polarization from the Crab using the data from the 2016 flight, given a better estimate of the background data during the observation. As shown in Chapter 6, the CDS background subtraction algorithm performs well for simulations of background-dominated observations, so we can employ it for this analysis. In this section, I will first describe how we perform polarimetry with COSI¹, and then discuss the results of doing so on the simulation of the Crab observation.

¹ See [Lei et al. \(1997\)](#) and [Lowell \(2017\)](#) for more detailed treatments of Compton polarimetry.

7.3 POLARIZATION ANALYSIS

Compton scatter angles	Distance between first two interactions	Distance between any two interactions	Number of interactions	Counts after subtraction
0°-65°	≥ 0.5 cm	≥ 0 cm	2-7	-433 ± 1708
0°-65°	≥ 0.5 cm	≥ 0 cm	3-7	-194 ± 1108
0°-65°	≥ 0.5 cm	≥ 0.3 cm	2-7	-405 ± 1268
0°-65°	≥ 0.5 cm	≥ 0.3 cm	3-7	-163 ± 656
0°-65°	≥ 1 cm	≥ 0 cm	2-7	-240 ± 1286
0°-65°	≥ 1 cm	≥ 0 cm	3-7	-121 ± 881
0°-65°	≥ 1 cm	≥ 0.3 cm	2-7	-273 ± 929
0°-65°	≥ 1 cm	≥ 0.3 cm	3-7	-143 ± 523
0°-30°	≥ 0.5 cm	≥ 0 cm	2-7	-178 ± 806
0°-30°	≥ 0.5 cm	≥ 0 cm	3-7	-169 ± 545
0°-30°	≥ 0.5 cm	≥ 0.3 cm	2-7	-109 ± 561
0°-30°	≥ 0.5 cm	≥ 0.3 cm	3-7	-96 ± 301
0°-30°	≥ 1 cm	≥ 0 cm	2-7	-161 ± 618
0°-30°	≥ 1 cm	≥ 0 cm	3-7	-164 ± 430
0°-30°	≥ 1 cm	≥ 0.3 cm	2-7	-115 ± 421
0°-30°	≥ 1 cm	≥ 0.3 cm	3-7	-110 ± 235
30°-65°	≥ 0.5 cm	≥ 0 cm	2-7	-253 ± 1102
30°-65°	≥ 0.5 cm	≥ 0 cm	3-7	-21 ± 665
30°-65°	≥ 0.5 cm	≥ 0.3 cm	2-7	-310 ± 846
30°-65°	≥ 0.5 cm	≥ 0.3 cm	3-7	-69 ± 407
30°-65°	≥ 1 cm	≥ 0 cm	2-7	-82 ± 819
30°-65°	≥ 1 cm	≥ 0 cm	3-7	43 ± 534
30°-65°	≥ 1 cm	≥ 0.3 cm	2-7	-162 ± 616
30°-65°	≥ 1 cm	≥ 0.3 cm	3-7	-33 ± 330

Table 7.3: The number of counts after the initial background subtraction of the Crab data for each set of event selections tested. These results indicate that we are not estimating the flight data background with sufficient accuracy.

7.3.1 COMPTON POLARIMETRY WITH COSI

The Klein-Nishina equation (Klein & Nishina, 1929) gives the differential cross section of photons Compton scattered from a free electron:

$$\frac{d\sigma}{d\Omega} = \frac{r_e^2}{2} \left(\frac{E'}{E_0} \right)^2 \left(\frac{E'}{E_0} + \frac{E_0}{E'} - 2 \sin^2 \phi \cos^2 \eta \right) \quad (7.1)$$

where r_e is the classical electron radius, E_0 and E' are the initial and final energies of the photon, ϕ is the Compton scatter angle, and η is the azimuthal scatter angle, or the angle between the initial electric field vector and the direction of the scattered photon. In other words, $\eta = 0$ when the photon scatters along the direction of the initial electric field vector.

It follows from the Klein-Nishina equation that a polarized beam will be modulated in η . This modulation can be described with a simple cosine:

$$f(\eta) = A + B \cos(2(\eta - \eta_0)) . \quad (7.2)$$

To determine the polarization properties of the incident beam, we measure the azimuthal scatter angle distribution (ASAD), or distribution of η values, and fit the ASAD with $f(\eta)$. η_0 is the polarization angle of the beam, and from A and B we can determine the polarization fraction. We define the measured modulation of the ASAD as

$$\mu = \frac{B}{A} \quad (7.3)$$

which can be converted into a polarization fraction Π :

$$\Pi = \frac{\mu}{\mu_{100}} \quad (7.4)$$

where μ_{100} is the observed modulation for a 100% polarized beam, i.e. in the case that $\Pi = 1$. For an ideal detector and a given energy E and Compton scatter angle ϕ , μ_{100} is given by

$$\mu_{100}^{\text{ideal}}(E, \phi) = \frac{d\sigma(\eta = 90) - d\sigma(\eta = 0)}{d\sigma(\eta = 90) + d\sigma(\eta = 0)} = \frac{\sin^2 \phi}{\frac{E'}{E} + \frac{E}{E'} - \sin^2 \phi} . \quad (7.5)$$

Figure 7.6 shows the the ideal modulation as a function of energy and Compton scatter angle. It is clear that the modulation is larger at low energies and for $\phi \sim 90^\circ$.




Figure 7.6: The ideal modulation of the azimuthal scatter angle as a function of Compton scatter angle for a variety of energies. The modulation is highest at low energies with Compton scatter angles $\phi \sim 90^\circ$. *Figure from Lowell (2017).*

With a real detector, the μ_{100} values will be lower than the ideal modulation primarily because of geometric effects. For example, if more detector material is present at certain azimuthal scatter angles, the efficiency at those azimuthal scatter angles will be higher and more counts will be collected, skewing the measured [ASAD](#) from the ideal. The placement and orientation of the strips is another geometry-related factor that can skew the [ASAD](#): if a photon Compton scatters twice on the same strip, that event will be discarded, since we are unable to assign a depth to such events. The removal of such events from the pipeline could artificially reduce the number of counts in certain directions, affecting the measured [ASAD](#). Thus it is imperative to correct for geometric effects when performing polarimetry.

To determine the polarization fraction and angle of a source detected by [COSI](#), we use the following procedure:

1. Choose the event selections that minimize the minimum detectable polarization ([MDP](#)), defined as

$$\text{MDP} = \frac{4.29}{\mu_{100} r_s} \sqrt{\frac{r_s + r_b}{T_{\text{obs}}}} \quad (7.6)$$

where r_S and r_B represent the source and background count rates, respectively, and T_{obs} is the observation time. The factor of 4.29 corresponds to 99% confidence (Weisskopf et al., 2010).

2. Make a background-subtracted **ASAD** by generating a source and background **ASAD** and subtracting the background **ASAD** from the source **ASAD**. We correct for geometric and other systematic effects by dividing the background-subtracted **ASAD** by a correction **ASAD**. The correction **ASAD** is the **ASAD** generated from the simulation of an unpolarized source, rescaled by its mean value.
3. Fit the corrected background-subtracted **ASAD** with Equation 7.2. Then we have determined the polarization angle η_0 and the polarization fraction $\Pi = \frac{\mu}{\mu_{100}} = \frac{(B/A)}{\mu_{100}}$.

In the case of a broadband persistent source, we use the **CDS** background subtraction method detailed in Chapter 6 to generate the background-subtracted **ASAD** (step 2 of the above procedure). The oversubtraction of background due to the **CDS-ARM** tails complicates the process of fitting the **ASAD**, since we need to correct for this oversubtraction (see Section 6.2.3 for a thorough description of this phenomenon). We adopt an iterative algorithm similar to the method described for spectral analysis in Section 6.2.3 and modify the above procedure for the polarization analysis as follows (step 1 remains the same, so we do not repeat it here):

2. Subtract the background **ASAD** from the source **ASAD** using Equation 6.12 (this will lead to an oversubtracted background). Perform the geometry correction and fit the corrected background-subtracted **ASAD** with Equation 7.2 to determine best fit values for η_0 , A , and B .
3. Subtract S_{BR} counts, or the number of source counts in the background region, from the background **ASAD** (recall that S_{BR} is estimated from the shape of the point source **CDS-ARM** distribution). The S_{BR} counts are distributed according to Equation 7.2 with the best fit parameters plugged in and multiplied by the correction **ASAD** so that the counts resemble the uncorrected data. After the S_{BR} counts are subtracted from the background **ASAD**, subtract the background **ASAD** from the source **ASAD** and apply the geometry correction.

4. Re-fit the background-subtracted, geometry corrected **ASAD** with the S_{BR} counts removed from the background **ASAD** to get updated best fit parameters.
5. Repeat steps 4 and 5 until the best fit parameters η_0 , A , and B converge.

73.2 POLARIMETRY OF SIMULATED CRAB OBSERVATION

We simulated a point source with a power law index of 2.169 and an 0.2-3 MeV flux of $0.0330 \text{ } \gamma/\text{cm}^2/\text{s}$, mimicking the parameters of the Crab spectrum as observed by the Spectrometer on **INTEGRAL** (SPI) (Sizun et al., 2004). To simplify the analysis, we simulated the source in detector coordinates at position ($\Theta = 34^\circ$, $\Phi = 7^\circ$), which is the average position of the Crab in the **COSI FOV** during the 2016 flight observation. For the background component, we used the Ling Model simulations introduced in Section 6.3. To ensure an accurate source to background ratio, we scaled the background simulations such that they contained the same number of counts as measured background. Figure 7.7 shows a comparison of the measured and simulated background spectrum, with the atmospheric background line at 511 keV removed (we exclude this line for all analysis performed in this chapter). The simulations are in good agreement with the measured background. After scaling the background, the source counts make up about 5% of the total counts in the simulated observation.

The first step of the polarization analysis is to find the event selections that minimize the **MDP**, defined in Equation 7.6. Minimizing the **MDP** is a trade-off between the maximizing the μ_{100} and the source count rate r_s . The best approach to this step is to use an optimization algorithm to minimize the **MDP**. As discussed in Section 7.2, using an optimization algorithm becomes too time consuming with the **CDS** background subtraction routine. To find the best set of event selections, we calculated the **MDP** for each combination of event selections in the same grid that we used in Section 7.2. Note that some of these sets of event selections are unlikely to result in a low **MDP**: for example, the modulation is higher for Compton scatter angles closer to 90° (see Figure 7.6), so using the Compton scatter angles between 0° and 30° is unlikely to result in a low **MDP**. However, that scatter angle range led to a clearer image, so we included it in this analysis. We used an **ARM** cut of 5° for this analysis, which is the full width at half maximum (**FWHM**) of the **ARM** distribution of the simulated point source.




Figure 7.7: A spectrum of flight data background compared with a spectrum of the simulated background, scaled such that both spectra contain the same number of events. The 511 keV line is removed since it is a large source of atmospheric background, and we remove it for all analysis done in this section. There is good agreement between the two spectra.

To calculate the [MDP](#), first we must determine μ_{100} with a simulation of a 100% polarized source. We simulated multiple 100% polarized point sources at $(\Theta = 34^\circ, \Phi = 7^\circ)$ with various polarization angles, ranging from 0° to 170° with 10 degree steps, to assess whether the μ_{100} varies based on polarization angle. For each set of event selections, we determined the μ_{100} as a function of polarization angle and fit the μ_{100} values with a constant. Figure [7.8](#) shows the μ_{100} value as a function of polarization angle for two sets of event selections, and Table [7.4](#) displays the best fit constant value in addition to the reduced χ^2 for each set of event selections. The reduced χ^2 values are close to 1, indicating that there is no significant variation in the modulation as a function of polarization angle. Thus we use the best fit μ_{100} value for each set of event selections when calculating the [MDP](#). Note that the Compton scatter angle range has the largest effect on the μ_{100} values; as expected (see e.g. Figure [7.6](#)), the range of smaller Compton scatter angles from 0° - 30° results in the least modulation whereas the range of larger Compton scatter angles from 30° - 65° results in the most modulation.

With the best fit μ_{100} values, we proceed with calculating the [MDP](#). For each set of event selections, we determined r_S and r_B from the simulations. Table [7.5](#) shows the source and background count rates for each set of event

Compton scatter angles	Distance between first two interactions	Distance between any two interactions	Number of interactions	best fit μ_{100}	χ^2/dof
0°-65°	≥ 0.5 cm	≥ 0 cm	2-7	0.188 ± 0.005	1.19
0°-65°	≥ 0.5 cm	≥ 0 cm	3-7	0.180 ± 0.007	1.35
0°-65°	≥ 0.5 cm	≥ 0.3 cm	2-7	0.188 ± 0.005	1.30
0°-65°	≥ 0.5 cm	≥ 0.3 cm	3-7	0.180 ± 0.007	1.48
0°-65°	≥ 1 cm	≥ 0 cm	2-7	0.195 ± 0.006	0.92
0°-65°	≥ 1 cm	≥ 0 cm	3-7	0.183 ± 0.007	0.88
0°-65°	≥ 1 cm	≥ 0.3 cm	2-7	0.195 ± 0.006	0.98
0°-65°	≥ 1 cm	≥ 0.3 cm	3-7	0.183 ± 0.007	0.89
0°-30°	≥ 0.5 cm	≥ 0 cm	2-7	0.085 ± 0.009	0.55
0°-30°	≥ 0.5 cm	≥ 0 cm	3-7	0.095 ± 0.011	0.89
0°-30°	≥ 0.5 cm	≥ 0.3 cm	2-7	0.084 ± 0.009	0.55
0°-30°	≥ 0.5 cm	≥ 0.3 cm	3-7	0.093 ± 0.011	0.92
0°-30°	≥ 1 cm	≥ 0 cm	2-7	0.091 ± 0.010	0.53
0°-30°	≥ 1 cm	≥ 0 cm	3-7	0.097 ± 0.012	0.75
0°-30°	≥ 1 cm	≥ 0.3 cm	2-7	0.090 ± 0.010	0.51
0°-30°	≥ 1 cm	≥ 0.3 cm	3-7	0.095 ± 0.012	0.75
30°-65°	≥ 0.5 cm	≥ 0 cm	2-7	0.236 ± 0.007	0.79
30°-65°	≥ 0.5 cm	≥ 0 cm	3-7	0.233 ± 0.012	0.85
30°-65°	≥ 0.5 cm	≥ 0.3 cm	2-7	0.236 ± 0.007	0.89
30°-65°	≥ 0.5 cm	≥ 0.3 cm	3-7	0.234 ± 0.012	1.04
30°-65°	≥ 1 cm	≥ 0 cm	2-7	0.240 ± 0.008	0.94
30°-65°	≥ 1 cm	≥ 0 cm	3-7	0.236 ± 0.013	0.75
30°-65°	≥ 1 cm	≥ 0.3 cm	2-7	0.241 ± 0.008	1.00
30°-65°	≥ 1 cm	≥ 0.3 cm	3-7	0.236 ± 0.013	0.83

Table 7.4: The best fit constant μ_{100} value for each set of event selections, in addition to the reduced χ^2 of the fit over polarization angle. The Compton scatter angle range used has the largest effect on the resulting μ_{100} .

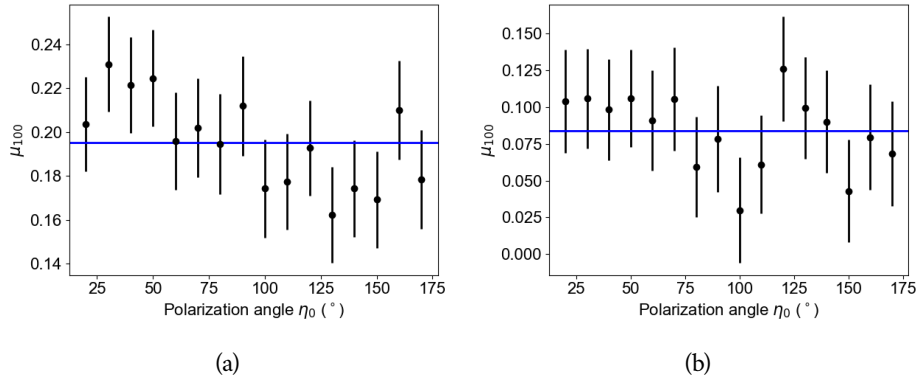


Figure 7.8: μ_{100} as a function of polarization angle for two sets of event selections, with the best fit μ_{100} shown as the blue line.

(a) Compton scatter angles from 0° - 65° , distance between first two interactions ≥ 1 cm, distance between any two interactions ≥ 0 cm, and 2-7 interactions; the best fit μ_{100} is 0.195 ± 0.006 .

(b) Compton scatter angles from 0° - 30° , distance between first two interactions ≥ 0.5 cm, distance between any two interactions ≥ 0.3 cm, and 2-7 interactions; the best fit μ_{100} is 0.084 ± 0.009 .

selections in addition to the **MDP**. All of the **MDP** values are larger than 100%, indicating that we are not sensitive to the polarization of the Crab observation with these event selections. Consequently, there is no need to fit the modulation curve of the simulated Crab emission as we are unable to measure the polarization signature. A longer observation time would lower the **MDP**. Additionally, if the Crab were further on-axis, we would have a larger effective area and thus the detection of more source counts, increasing r_s and lowering the **MDP**. The **EHC** has a large effect on the **MDP** values: the **EHC** cuts out Compton scatter angles above 65° , but as shown in Figure 7.6, the modulation is highest at Compton scatter angles of $\sim 90^\circ$. The effect of the Compton scatter angle range on the modulation and thus the **MDP** is clear by comparing the **MDP** values as a function of Compton scatter angle in Table 7.5. The **MDP** values are significantly larger when we used a Compton scatter angle range of 0° - 30° than when we used ranges that include higher Compton scatter angles.

The **EHC** is currently the best algorithm we have to reject albedo radiation background present at balloon altitudes. The **COSI** team is currently working on a machine learning approach to reject the albedo radiation background to

Compton scatter angles	Distance between first two interactions	Distance between any two interactions	Number of interactions	S	B	MDP
0°-65°	≥ 0.5 cm	≥ 0 cm	2-7	1511	26363	251 ± 9%
0°-65°	≥ 0.5 cm	≥ 0 cm	3-7	728	11695	364 ± 19%
0°-65°	≥ 0.5 cm	≥ 0.3 cm	2-7	1489	26111	254 ± 9%
0°-65°	≥ 0.5 cm	≥ 0.3 cm	3-7	706	11442	373 ± 20%
0°-65°	≥ 1 cm	≥ 0 cm	2-7	1153	19892	277 ± 11%
0°-65°	≥ 1 cm	≥ 0 cm	3-7	593	9224	391 ± 22%
0°-65°	≥ 1 cm	≥ 0.3 cm	2-7	1134	19687	280 ± 11%
0°-65°	≥ 1 cm	≥ 0.3 cm	3-7	574	9019	401 ± 23%
0°-30°	≥ 0.5 cm	≥ 0 cm	2-7	594	13974	1022 ± 112%
0°-30°	≥ 0.5 cm	≥ 0 cm	3-7	313	6254	1170 ± 152%
0°-30°	≥ 0.5 cm	≥ 0.3 cm	2-7	583	13845	1050 ± 117%
0°-30°	≥ 0.5 cm	≥ 0.3 cm	3-7	302	6125	1221 ± 163%
0°-30°	≥ 1 cm	≥ 0 cm	2-7	421	10227	1154 ± 134%
0°-30°	≥ 1 cm	≥ 0 cm	3-7	237	4685	1311 ± 183%
0°-30°	≥ 1 cm	≥ 0.3 cm	2-7	412	10128	1189 ± 141%
0°-30°	≥ 1 cm	≥ 0.3 cm	3-7	228	4586	1371 ± 197%
30°-65°	≥ 0.5 cm	≥ 0 cm	2-7	917	12389	229 ± 10%
30°-65°	≥ 0.5 cm	≥ 0 cm	3-7	415	5440	340 ± 24%
30°-65°	≥ 0.5 cm	≥ 0.3 cm	2-7	906	12266	230 ± 10%
30°-65°	≥ 0.5 cm	≥ 0.3 cm	3-7	404	5317	343 ± 24%
30°-65°	≥ 1 cm	≥ 0 cm	2-7	732	9665	249 ± 12%
30°-65°	≥ 1 cm	≥ 0 cm	3-7	356	4540	363 ± 27%
30°-65°	≥ 1 cm	≥ 0.3 cm	2-7	722	9559	250 ± 12%
30°-65°	≥ 1 cm	≥ 0.3 cm	3-7	356	4540	358 ± 27%

Table 7.5: The number of source and background counts S and B from the simulated Crab observation, which is 24.61 ks long, and the MDP for each set of event selections. MDP values above 100% indicate that we are unable to measure the polarization properties of the Crab from this observation. Analyzing the source observed at a smaller zenith angle over a longer period of time, combined with an alternate albedo radiation background rejection algorithm accommodating of Compton scatter angles of $\sim 90^\circ$, would significantly improve the MDP for this observation.

replace the EHC; if this approach is successful, we could potentially measure the polarization of the Crab with this particular dataset. With thicker shielding or COSI in a space (rather than balloon) environment, the EHC would no longer be necessary, and COSI would perform better as a polarimeter. For extremely bright sources like GRBs, the EHC is not necessary, and with the ability to use Compton scatter angles of $\sim 90^\circ$, we were able to perform polarimetry of GRB 160530A with the current iteration of COSI (Lowell et al., 2017a).

Lowell et al. (2017b) implemented a maximum likelihood approach to Compton polarimetry for sources detected by COSI, based on the method described in Krawczynski (2011) for an ideal detector. The maximum likelihood approach decreases the MDP by $\sim 20\%$, allowing for more sensitive polarization measurements than the standard method described in this chapter. However, the increase in polarization sensitivity due to the maximum likelihood method is insufficient to perform polarimetry of COSI’s observation of the Crab.

7.4 CONCLUSIONS

During the 2016 flight, we detected the Crab nebula in the COSI data in both the image and the ARM distribution. The CDS background subtraction algorithm detailed in Chapter 6 resulted in too few counts to successfully fit the spectrum and constrain the parameters. These results indicate that our treatment of the flight background is limited when studying severely background-dominated sources. A more accurate model of the flight background is needed to fully analyze broadband, persistent, background-dominated sources detected by the current iteration of COSI.

We proceeded to use simulations to determine the prospects of performing polarimetry of COSI’s observation of the Crab, if we had an acceptable background subtraction. With MDP values above 100%, due (at least in part) to the EHC, we conclude that we would not have been able to perform polarimetry of this observation. An alternative albedo radiation rejection algorithm that accommodates Compton scatter angles of $\sim 90^\circ$ would significantly improve the prospects of polarimetry with COSI. Additionally, deploying COSI in a space environment with thicker shielding would minimize background and remove the need for such an algorithm, leading to a vastly improved polarization sensitivity.

8

THE REFLECTION SPECTRUM OF THE LOW-MASS X-RAY BINARY 4U 1728–34

Though the majority of this thesis is focused on the goal of studying compact objects with γ -ray polarimetry, particularly with [COSI](#), we can also deepen our understanding of these objects with other analysis techniques. In this chapter¹, I will discuss measuring the reflection features in the spectrum of the neutron star low-mass X-ray binary ([LMXB](#)) 4U 1728–34. By studying the reflection features, we can measure the inner radius of the accretion disk, and if the disk is not truncated, we can use this measurement to constrain the radius of the neutron star itself. Measurements of neutron star radii are critical to understanding the neutron star equation of state ([Lattimer & Prakash, 2007](#); [Cackett et al., 2010](#)). The equation of state of the ultra-dense matter found in neutron stars is currently unknown, and as we cannot create matter at such densities in the laboratory, observing neutron stars is the only method available to us to study how matter behaves at ultra-high densities.

A reflection spectrum is thought to arise when hard X-ray emission reflects onto an area of the accretion disk close to the neutron star ([Fabian et al., 1989](#)) (see Figure 8.1 for a schematic). The primary signature of this phenomenon is an iron emission line with energies from 6.4 to 7.1 keV, caused by the fluorescent $K\alpha$ transition of iron. These lines have been detected in both neutron star and black hole X-ray systems. Relativistic and Doppler effects distort the profile of the line, broadening it significantly and skewing it to low energies ([Miller, 2007](#); [Reynolds & Nowak, 2003](#)). From this unique shape, one can measure interesting properties of the accretion disk, including its inner radius.

¹ This chapter is largely based on a previous publication titled “A NuSTAR observation of the reflection spectrum of the low-mass X-ray binary 4U 1728–34” ([Sleator et al., 2016](#)). ©Astrophysical Journal 2016. Reprinted with permission.




Figure 8.1: A schematic of photons emitted by the neutron star (black) reflecting off of the accretion disk (red). The photons emitted directly by the neutron star form the primary component of the spectrum and are drawn in green, while the photons that form the reflection component of the spectrum are drawn in blue.

8.1 BACKGROUND

4U 1728–34 is a neutron star low-mass X-ray binary (LMXB) of the atoll class (Lewin et al., 1976; Hasinger & van der Klis, 1989), with an estimated distance of 4.1–5.1 kpc (Di Salvo et al., 2000; Galloway et al., 2003). It exhibits Type-1 X-ray bursts caused by thermonuclear burning on the stellar surface (Galloway et al., 2010). From burst oscillations, Strohmayer et al. (1996) measured a spin frequency of 363 ± 6 Hz.

A broad emission line at 6.7 keV has been detected in several spectral analyses of 4U 1728–34 performed with satellites such as *XMM-Newton* (Ng et al., 2010; Egron et al., 2011), *INTEGRAL* (Tarana et al., 2011), *RXTE* (Piraino et al., 2000; Seifina & Titarchuk, 2011) and *BeppoSAX* (Di Salvo et al., 2000; Piraino et al., 2000; Seifina & Titarchuk, 2011). The continuum spectrum is generally composed of a soft component described as thermal emission from the stellar surface or accretion disk, and a hard component described as Comptonization.

The Nuclear Spectroscopic Telescope Array (*NuSTAR*) (Harrison et al., 2013) has been successful in revealing iron $K\alpha$ lines and reflection spectra in neutron stars (e.g. Miller et al. 2013, Degenaar et al. 2015, King et al. 2016). In this work we analyze a coordinated *NuSTAR* and *Swift* observation of 4U 1728–34 in an effort to further constrain its reflection spectrum and thus the neutron star radius.

8.2 OBSERVATIONS AND DATA REDUCTION

NuSTAR observed 4U 1728–34 on 2013 October 1 for 33.5 ks (Figure 8.2; ObsID 80001012002). The data were processed with the *NuSTAR* Data Analysis Software (*NuSTARDAS*) version 1.4.1 and CALDB 20150702, resulting in 27 ks of net exposure time. We extracted spectra from the Focal Plane Module A (*FPMA*) and Focal Plane Module B (*FPMB*) using a circular region with a $180''$ radius centered at the source position. We used a $671'' \times 114''$ rectangular background region centered $388''$ away from the source position. At 5 keV, the ratio of source count rate to background count rate is 200, and at 50 keV, this ratio is 3, indicating that the spectrum is not very sensitive to the background estimate. Due to a known shift in gain offset that has not been properly accounted for in the current response files used for the reduction in this paper, we restrict our analysis to the 4.5–78.4 keV band².

An observation (ObsID 00080185001) of 4U 1728–34 was made with the *Swift* (Gehrels et al., 2004) X-ray Telescope (*XRT*) (Burrows et al., 2005) near the beginning of the *NuSTAR* observation (see Figures 8.2 and 8.4). The *XRT* was operated in Windowed Timing mode to avoid photon pile-up. Although the *XRT* observation lasted for two *Swift* orbits, the source was near the edge of the active area of the detector for the first orbit, and we produced an energy spectrum using only the second orbit, giving an exposure time of 1068 s. As this source has a relatively high column density, we used only grade 0 photons as this is the recommended procedure for high column density sources³. The 0.7–10 keV spectrum was extracted from a circular region with a 20 pixel ($47''$) radius, and background was taken from a region far from the source. We obtained a source count rate in the 0.7–10 keV band of 17.7 counts per second. For spectral fitting, we calculated a response matrix appropriate for grade 0 photons (file `swxwt0s6_20130101v015.rmf`) and used `xrtmkarf` with an exposure map to produce the ancillary response file.

All spectra were analyzed using XSPEC version 12.8 (Arnaud, 1996). All fits were made assuming Verner et al. (1996) cross sections and Wilms et al. (2000) abundances. The spectra were binned such that the signal-to-noise ratio in each bin is 15σ . To better constrain the low energy spectrum, particularly the column density, we fit the *Swift* spectrum together with the *NuSTAR* spectra. Due to flux variations between the instruments, we added a multiplicative constant in each fit. We fixed the constant for the *NuSTAR* *FPMA*

² Confirmed by *NuSTAR* instrument team (Kristin Madsen, private communication)

³ see http://www.swift.ac.uk/analysis/xrt/digest_cal.php#abs




Figure 8.2: *MAXI* and *BAT* light curves with the time of the *NuSTAR* observations marked by the dashed lines. The arrow marks the time of the *Swift* observation.

spectrum to unity and allowed the constants for the *NuSTAR* FPMB and *Swift* spectra to vary.

8.3 ANALYSIS AND RESULTS

4U 1728–34 is known to exhibit Type-1 X-ray bursts. Using the light curves made by the `nuproducts` FT00L, we detected and removed four bursts in the *NuSTAR* data, each lasting about 20 s (Figure 8.3). No bursts were detected during the *Swift* observation. Mondal et al. (2017) did a full analysis of the bursts. To check the stability of the energy spectrum during the observation, we looked at the hardness ratio, defined here as the count rate from 12–25 keV divided by the count rate from 4.5–12 keV (Figure 8.4). The hardness ratio only changes slightly, softening gradually with time, indicating a fairly stable spectrum.

We first fit the combined *Swift* and *NuSTAR* continuum spectra similarly to Ng et al. (2010), with a model consisting of a neutral absorption component `tbabs`, a single temperature blackbody `bbodyrad`, a disk blackbody `diskbb`, and a powerlaw component `cutoffpl`. This model takes into account ther-




Figure 8.3: (a) The NuSTAR FPMA (black) and FPMB (red) light curve, including the four type I X-ray bursts. The time of the *Swift* observation is marked in blue. (b) A close-up of the first type I X-ray burst. The green lines indicate the data removed for spectral analysis.




Figure 8.4: The hardness ratio throughout the NuSTAR observation, with NuSTAR FPMA in red and NuSTAR FPMB in black. The time of the *Swift* observation is marked in blue. The hardness ratio is defined as the count rate from 12-25 keV / the count rate from 4.5-12 keV.




Figure 8.5: The 4.5-78.4 keV *NuSTAR* FPMA (black) and *NuSTAR* FPMB (red) spectra and the 0.7-10 keV *Swift* (blue) spectrum of 4U 1728-34, fit with the continuum model: disk blackbody + blackbody + power-law. The residuals between 6-8 keV indicate the presence of iron emission.

mal emission from the boundary layer between the neutron star surface and disk, thermal emission from the disk, and non-thermal emission from Comptonization. We found a blackbody temperature of 1.41 ± 0.01 keV, a disk blackbody temperature of 0.32 ± 0.02 keV, a photon index of 1.29 ± 0.05 , and a cutoff energy of 18.3 ± 0.7 keV (all errors are 90% confidence). This model fits most of the energy band fairly well ($\chi^2/\text{dof} = 1700/1257 = 1.35$), but there are significant residuals between 6-8 keV (Figure 8.5).

We added a Gaussian line to model the excess at ~ 6 keV. Leaving the line parameters free gives unphysical results (the line energy is at $\simeq 5.7$ keV with a width $\sigma \simeq 1.6$ keV). Similarly to D’Ái et al. (2006) and Eggen et al. (2011), we limited the line width to $\sigma = 0.5$ keV, resulting in a line energy of 6.56 ± 0.05 keV and an equivalent width of 74 eV. The addition of a Gaussian significantly improved the fit, with a $\chi^2/\text{dof} = 1529/1253 = 1.22$. Moreover, the `diskbb` component became statistically insignificant, its exclusion from the fit resulting in $\Delta\chi^2 = 5$ and only slight changes to the other parameters. We continued our analysis without the `diskbb` component.

Having confirmed the presence of the iron line with the Gaussian model, we replaced the Gaussian with a physical reflection model, `reflionx` (Ross & Fabian, 2005). The `reflionx` model used here is a version of Ross & Fabian (2005) that assumes reflection of a power law with a high-energy exponential cutoff. To take relativistic blurring into account, we convolved `reflionx` with `kerrconv` (Brenneman & Reynolds, 2006).

The `reflionx` parameters include the photon index and cutoff energy (tied to those of the `cutoffpl` component), the ionization ($\xi = L/nr^2$, or

the ratio of the flux to the column density, where L is the luminosity, r is the distance, and n is the column density), the iron abundance (A_{Fe}), and the flux. The `kerrconv` parameters include the compact object spin parameter, disk inclination, disk inner and outer radius, and inner and outer emissivity indices. The spin parameter $a \equiv cJ/GM^2 = 0.15$ (where J is angular momentum) can be calculated from previous measurements of the neutron star spin frequency, 363 Hz (Strohmayer et al., 1996), assuming a typical neutron star mass of $M = 1.4M_{\odot}$. We fixed the disk outer radius $R_{\text{out}} = 400 R_{\text{ISCO}}$ (where R_{ISCO} is the radius of the innermost stable circular orbit) because the emissivity profile drops steeply with radius, causing the fits to be insensitive to the exact value of this parameter.

As it was difficult to constrain the emissivity of the reflecting disk when it was left free, we considered fits with $q = 1, 2, 3, 4$ and 5. We also considered modeling the emissivity as a broken power law with the outer index fixed to 3, the inner index left free, and the break radius fixed to $25R_g$, but we were unable to constrain the inner index. As shown in Table 8.1, the changes in emissivity negligibly affect the χ^2 . Furthermore, the values for the inner radius are very similar across the models (other than $q = 1$, where the inner radius has very large error bars). As the parameters, particularly the inner radius, seem to be insensitive to the emissivity, we continued our analysis with the emissivity fixed at $q = 3$, consistent with a Newtonian geometry far from the neutron star.

Adding the reflection component (Figure 8.6) improves the fit over the Gaussian line model ($\chi^2/\text{dof} = 1430/1254 = 1.14$; $\Delta\chi^2 = 99$). Table 8.2 lists the best fit parameters for the relativistically blurred reflection model (model 1). We find an inclination of $\sim 37^\circ$, consistent with the lack of dips in the light curve which implies a low viewing angle. From the normalization of `bbbodyrad`, we find a blackbody source radius of 1.4 km, consistent with thermal emission from the boundary layer. We find a higher column density, $N_{\text{H}} \sim 4.5 \times 10^{22} \text{ cm}^{-2}$, than what has been previously measured for this source ($N_{\text{H}} \sim 2.6\text{--}2.9 \times 10^{22} \text{ cm}^{-2}$, e.g. Di Salvo et al. 2000; Piraino et al. 2000; Narita et al. 2001; Egron et al. 2011). To be consistent with past measurements, we fixed the column density to $2.9 \times 10^{22} \text{ cm}^{-2}$ (Figure 8.7; model 2 in Table 8.2). This significantly worsens the fit, resulting in a $\chi^2/\text{dof} = 1712/1255 = 1.36$ ($\Delta\chi^2 = 282$), yet does not cause a large change to the main parameter of interest, R_{in} .

Due to the proximity of 4U 1728–34 to the Galactic plane ($b = -0.15^\circ$), it is reasonable to consider the column density of molecular hydrogen. Galactic

8.3 ANALYSIS AND RESULTS

Model	Parameter ^a	Units	$q = 1$	$q = 2$	$q = 3$	$q = 4$	$q = 5$	q broken
constant	FPMA	–	1 ^b	1 ^b	1 ^b	1 ^b	1 ^b	1 ^b
	FPMB	–	1.05 ± 0.002	1.05 ± 0.002	1.05 ± 0.002	1.05 ± 0.002	1.05 ± 0.002	1.05 ± 0.002
	XRT	–	1.12 ± 0.01	1.12 ± 0.01	1.15 ± 0.01	1.13 ± 0.01	1.13 ± 0.02	1.12 ± 0.02
tbabs	N_{H}^c	10 ²² cm ⁻²	4.6 ± 0.1	4.5 ^{+0.1} _{-0.2}	4.5 ± 0.1	4.4 ± 0.1	4.4 ± 0.1	4.4 ± 0.1
bbodyrad	kT	keV	1.54 ± 0.02	1.69 ^{+0.03} _{-0.09}	1.53 ^{+0.02} _{-0.05}	1.52 ± 0.04	1.53 ± 0.04	1.49 ^{+0.06} _{-0.01}
	norm	$R_{\text{km}}^2/D_{10\text{kpc}}^2$	6.26 ^{+1.4} _{-0.5}	4.2 ^{+0.7} _{-0.4}	7.9 ^{+2.6} _{-0.7}	8.5 ± 2	7.7 ^{+3.3} _{-1.7}	9.4 ^{+1.7} _{-2.1}
cutoffpl	Γ	–	1.52 ^{+0.04} _{-0.03}	1.51 ^{+0.02} _{-0.01}	1.54 ^{+0.04} _{-0.05}	1.52 ^{+0.05} _{-0.02}	1.51 ± 0.05	1.52 ± 0.04
	HighECut	keV	26 ⁺³ ₋₁	26 ⁺² ₋₄	25 ± 2	25 ⁺² ₋₃	26 ⁺³ ₋₂	26 ⁺³ ₋₁
	norm	$\gamma \text{ keV}^{-1} \text{ cm}^{-2} \text{ s}^{-1} \text{ at } 1 \text{ keV}$	0.19 ^{+0.02} _{-0.04}	0.19 ^{+0.13} _{-0.03}	0.29 ^{+0.03} _{-0.07}	0.26 ^{+0.04} _{-0.07}	0.23 ^{+0.05} _{-0.07}	0.20 ± 0.06
kerrconv	q_{in}	–	1 ^b	2 ^b	3 ^b	4 ^b	5 ^b	> 5.4
	q_{out}	–	1 ^b	2 ^b	3 ^b	4 ^b	5 ^b	3 ^b
	a	–	0.15 ^b	0.15 ^b	0.15 ^b	0.15 ^b	0.15 ^b	0.15 ^b
	Incl.	deg.	35 ⁺¹³ ₋₅	19 ⁺⁸ ₋₄	37 ⁺¹ ₋₂	37 ⁺² ₋₃	39 ⁺² ₋₄	40 ⁺¹ ₋₃
	R_{in}	ISCO	3.2 ^{+3.45} _{-2.2}	1.0 ^{+0.56} ₋₀	1.0 ^{+0.77} ₋₀	1.1 ^{+0.50} _{-0.1}	1.1 ^{+0.48} _{-0.1}	1.2 ^{+0.4} _{-0.2}
	R_{out}	ISCO	400 ^b	400 ^b	400 ^b	400 ^b	400 ^b	400 ^b
reflionx	ξ	erg cm s ⁻¹	3142 ⁺⁹⁸⁷ ₋₃₀₂	4564 ⁺²⁰⁰⁰ ₋₁₅₅₆	796 ⁺⁷⁹⁸ ₋₂₂₉	984 ⁺⁷²² ₋₄₇₂	1208 ⁺⁵⁵⁹ ₋₅₃₂	1135 ⁺³⁵⁰ ₋₂₆₅
	A_{Fe}	–	0.22 ± 0.1	0.46 ^{+0.01} _{-0.08}	0.19 ± 0.1	0.17 ^{+0.06} _{-0.04}	0.19 ^{+0.07} _{-0.08}	0.14 ± 0.01
	norm	10 ⁻⁶	1.3 ^{+0.3} _{-0.2}	0.9 ^{+0.3} _{-0.2}	3.8 ^{+2.6} _{-1.4}	3.6 ^{+2.3} _{-0.6}	3.2 ^{+1.8} _{-0.9}	4.1 ^{+1.3} _{-1.2}
χ^2/dof	–	–	1431/1254	1427/1254	1430/1254	1427/1254	1425/1254	1422/1253

- ^a The errors on the parameters are 90% confidence.
^b These parameters were fixed.
^c The column density is calculated assuming [Wilms et al. \(2000\)](#) abundances and [Verner et al. \(1996\)](#) cross sections.

Table 8.1: Spectral parameters with varied emissivity using the `reflionx` model

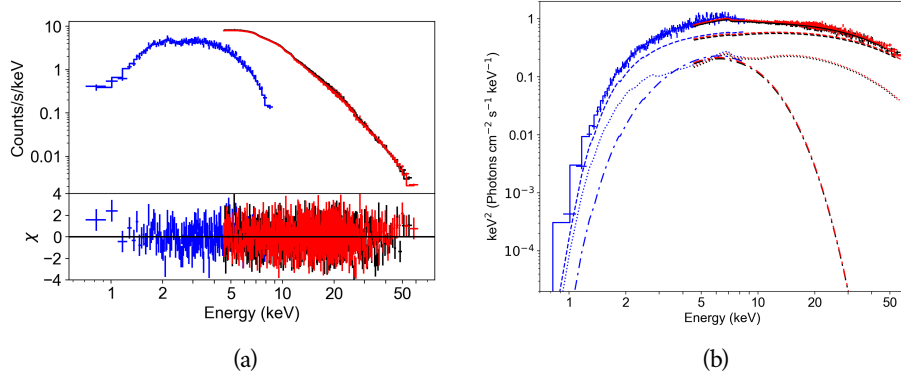


Figure 8.6: The 4.5–78.4 keV *NuSTAR* FPMA (black) and *NuSTAR* FPMB (red) spectra and the 0.7–10 keV *Swift* (blue) spectrum of 4U 1728–34, fit with the `reflionx` relativistically blurred reflection model (model 1). (a) Shows the residuals of the reflection model, specifically the lack thereof between 6–8 keV. (b) Shows the νF_ν plot with individual model components: blackbody (dashed and dotted), cutoff power-law (dashed), and `reflionx` (dotted).

8.3 ANALYSIS AND RESULTS

Model	Parameter ^d	Units	Model 1	Model 2	Model 3	Model 4
			Full band	N_{H} fixed	line only	<i>NuSTAR</i> only
constant	FPMA	–	1 ^b	1 ^b	1 ^b	1 ^b
	FPMB	–	1.05 ± 0.002	1.05 ± 0.002	1.05 ± 0.002	1.05 ± 0.002
	XRT	–	1.12 ± 0.02	1.15 ± 0.01	1.15 ± 0.02	–
tbabs	N_{H}^c	10 ²² cm ^{−2}	4.5 ± 0.1	2.9 ^b	4.0 ± 0.1	2.9 ^b
bbodyrad	kT	keV	1.53 ^{+0.02} _{−0.05}	1.42 ± 0.01	1.47 ± 0.01	1.45 ^{+0.06} _{−0.03}
	norm	$R_{\text{km}}^2/D_{10\text{kpc}}^2$	7.9 ^{+2.6} _{−0.7}	20.6 ^{+0.8} _{−0.4}	10.7 ^{+0.5} _{−0.6}	10.8 ± 1.8
cutoffpl	Γ	–	1.54 ^{+0.04} _{−0.05}	1.0 ^{+0.02} _{−1.0}	1.34 ^{+0.04} _{−0.02}	1.69 ± 0.08
	HighECut	keV	25 ± 2	18.0 ^{+0.6} _{−0.1}	19.1 ^{+0.7} _{−0.6}	36 ⁺⁶ _{−5}
	norm	γ keV ^{−1} cm ^{−2} s ^{−1} at 1keV	0.29 ^{+0.03} _{−0.07}	0.065 ^{+0.008} _{−0.005}	0.29 ± 0.02	0.15 ^{+0.08} _{−0.05}
kerrconv	q	–	3 ^b	3 ^b	–	3 ^b
	a	–	0.15 ^b	0.15 ^b	–	0.15 ^b
	Incl.	deg.	37 ⁺¹ _{−2}	29 ⁺³ _{−2}	–	33 ⁺³ _{−2}
	R_{in}	ISCO	1.0 ^{+0.77} _{−0}	1.0 ^{+0.44} _{−0}	–	1.7 ^{+1.0} _{−0.7}
	R_{out}	ISCO	400 ^b	400 ^b	–	400 ^b
reflionx	ξ	erg cm s ^{−1}	796 ⁺⁷⁹⁸ _{−229}	995 ⁺³⁰ _{−62}	–	909 ⁺³³⁷ _{−189}
	A_{Fe}	–	0.19 ± 0.1	0.09 ± 0.01	–	0.07 ^{+0.07} _{−0.01}
	norm	10 ^{−6}	3.8 ^{+2.6} _{−1.4}	3.3 ^{+0.4} _{−0.2}	–	8.8 ^{+4.2} _{−5.6}
relline	lineE	keV	–	–	7.1 ± 0.1	–
	q	–	–	–	3 ^b	–
	a	–	–	–	0.15 ^b	–
	Incl.	deg	–	–	18 ⁺⁵ _{−3}	–
	R_{in}	ISCO	–	–	1.1 ^{+0.2} _{−0}	–
	R_{out}	R_{g}	–	–	400 ^b	–
	norm	10 ^{−3}	–	–	1.8 ± 0.2	–
$\Omega/2\pi$	–	–	0.43	0.76	–	–
χ^2/dof	–	–	1430/1254	1712/1255	1532/1255	1256/1096

^a The errors on the parameters are 90% confidence.

^b These parameters were fixed.

^c The column density is calculated assuming [Wilms et al. \(2000\)](#) abundances and [Verner et al. \(1996\)](#) cross sections.

Table 8.2: Spectral parameters varying the column density and the energy band using **reflionx**




Figure 8.7: The 4.5–78.4 keV *NuSTAR* FPMA (black) and *NuSTAR* FPMB (red) spectra and the 0.7–10 keV *Swift* (blue) spectrum of 4U 1728–34, fit with the *reflionx* relativistically blurred reflection model where the column density is fixed to $2.9 \times 10^{22} \text{ cm}^{-2}$ (model 2). (a) Shows the residuals of the reflection model. (b) Shows the νF_ν plot with individual model components: blackbody (dashed and dotted), cutoff power-law (dashed), and *reflionx* (dotted).

surveys indicate that at the position of 4U 1728–34, $N_{\text{H}_2} \sim 1.8 \times 10^{22} \text{ cm}^{-2}$ (Dame et al., 2001) and $N_{\text{H}_I} \sim 1.24 \times 10^{22} \text{ cm}^{-2}$ (Kalberla et al., 2005). Thus, the expected total column density is $N_{\text{H}_{\text{total}}} = N_{\text{H}_I} + 2N_{\text{H}_2} \sim 4.84 \times 10^{22} \text{ cm}^{-2}$, close to our measured value. We conclude that the model with the column density as a free parameter, which is a better fit to the data, is more correct. From this model, we find an upper limit for the disk inner radius $R_{\text{in}} \leq 1.77 R_{\text{ISCO}}$, with the best value at $R_{\text{in}} = 1.00^{+0.77}_{-0} R_{\text{ISCO}}$.

We consider if the reflection parameters are mostly constrained by the iron line as opposed to the reflected continuum by fitting the data with the *relline* model, a relativistic line profile excluding broadband features such as the Compton hump. The best fit parameters are shown in model 3 of Table 8.2. The value for the inner radius, $R_{\text{in}} = 1.1^{+0.2}_{-0} R_{\text{ISCO}}$, is consistent with our above upper limit of $R_{\text{in}} \leq 1.77 R_{\text{ISCO}}$ and is even better constrained. However, the *relline* model, with $\chi^2/\text{dof} = 1532/1255 = 1.22$, is not as good a fit as the broadband reflection model described above ($\chi^2/\text{dof} = 1.14$), suggesting that the broadband reflection spectrum does make some contribution, at least statistically.

To verify that the *Swift* spectrum helps constrain the spectral shape, we fit only the *NuSTAR* data using the *reflionx* model (model 4 in Table 8.2). We

fixed the column density to $2.9 \times 10^{22} \text{ cm}^{-2}$ as the *NuSTAR* data is unable to measure this parameter. Fitting only the *NuSTAR* spectrum with this model resulted in a $\chi^2/\text{dof} = 1256/1096 = 1.15$ and is statistically better than including the *Swift* data (model 3 in Table 8.2; $\chi^2/\text{dof} = 1712/1255 = 1.36$). However, without the low-energy coverage offered by *Swift*, the column density cannot be well measured and thus the inaccuracy of fixing $N_{\text{H}} = 2.9 \times 10^{22} \text{ cm}^{-2}$ is not revealed in the fit statistics. Without the *Swift* spectrum, we find $R_{\text{in}} \leq 2.7 R_{\text{ISCO}}$ as an upper limit for the inner radius (contrasting $R_{\text{in}} \leq 1.4 R_{\text{ISCO}}$ found in model 3), indicating that the *Swift* spectrum is useful in evaluating the main parameter of interest as well as the column density.

Mondal et al. (2017) have also analyzed this coordinated *NuSTAR* and *Swift* observation, fitting the data instead with *relxill*, another relativistically blurred reflection model (*García et al., 2014*). To compare to their results, we also fit the spectrum with *relxill*, and in addition replaced *cutoffpl* for the more physical *comptt*. The *relxill* parameters include the ionization, iron abundance, compact object spin parameter, disk inclination, disk inner and outer radius, and inner and outer emissivity indices. We fixed $R_{\text{out}} = 400 R_{\text{g}}$ and $a = 0.15$. Similarly to *Mondal et al. (2017)*, we fixed $q_{\text{out}} = 3$ and left q_{in} free, but we were only able to find a lower limit on q_{in} . The *relxill* model (Figure 8.8) is statistically similar to model 1 with a $\chi^2/\text{dof} = 1412/1250 = 1.13$, and the best fit parameters shown in Table 8.3 are comparable to those of model 1. In particular, the column density $N_{\text{H}} = 3.9 \times 10^{22} \text{ cm}^{-2}$ is considerably higher than past measurements and the inclination $i = 37^\circ$ is the same as model 1. We find an upper limit of $R_{\text{in}} \leq 2 R_{\text{ISCO}}$ for the disk inner radius, with the best value at $R_{\text{in}} = 1.6 \pm 0.4$. This upper limit is close to our previous upper limit of $1.77 R_{\text{ISCO}}$, but as it is slightly higher, we continue our analysis considering $R_{\text{in}} \leq 2 R_{\text{ISCO}}$ as the upper limit for the disk inner radius.

With spin parameter value of $a = 0.15$, we calculate $1 R_{\text{ISCO}} = 5.5 R_{\text{g}}$ (*Bardeen et al., 1972*). Assuming a typical neutron star mass of $1.4 M_{\odot}$, we find an upper limit for the neutron star radius $R_{\text{NS}} \leq 2 \times 5.5 R_{\text{g}} = 23 \text{ km}$ (Figure 8.9).

8.4 DISCUSSION

The unabsorbed flux extrapolated in the 0.1-100 keV energy range is $F_{0.1-100} \sim 6 \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}$. Assuming a distance to the source of 5.1 kpc (*Di Salvo et al., 2000*; *Galloway et al., 2003*), we calculate a luminosity $L_{0.1-100} = 1.9 \times 10^{37} \text{ erg s}^{-1}$, corresponding to 8% of the Eddington luminos-

Model	Parameter ^a	Units	Value
constant	FPMA	–	1 ^b
	FPMB	–	1.05 ± 0.002
	XRT	–	1.15 ± 0.02
tbabs	N_{H}^c	10 ²² cm ^{−2}	3.9 ^{+0.4} _{−0.3}
bbody	kT	keV	0.49 ^{+0.06} _{−0.05}
	norm	$L_{10^{39}\text{ergs/s}}/D_{10\text{kpc}}^2$	0.007 ^{+0.002} _{−0.001}
comptt	kT_{seed}	keV	1.13 ^{+0.05} _{−0.03}
	kT_e	keV	15 ⁺⁵ _{−3}
	optical depth	–	1.6 ^{+0.3} _{−0.6}
	norm	–	0.022 ^{+0.002} _{−0.007}
relxill	q_{in}	–	> 4
	q_{out}	–	3 ^b
	a	–	0.15 ^b
	Incl.	deg.	37 ⁺⁶ _{−2}
	R_{in}	ISCO	1.6 ± 0.4
	R_{out}	R_{g}	400 ^b
	log ξ	–	3.0 ^{+0.1} _{−2.1}
	A_{Fe}	–	0.8 ± 0.1
	norm	10 ^{−3}	1.2 ⁺¹ _{−0.2}
χ^2/dof	–	–	1412/1250

^a The errors on the parameters are 90% confidence

^b These parameters were fixed

^c The column density is calculated assuming [Wilms et al. \(2000\)](#) abundances and [Verner et al. \(1996\)](#) cross sections.

Table 8.3: Spectral Parameters using `relxill`




Figure 8.8: The 4.5–78.4 keV *NuSTAR* FPMA (black) and *NuSTAR* FPMB (red) spectra and the 0.7–10 keV *Swift* (blue) spectrum of 4U 1728–34, fit with *relxill* (model 5). (a) Shows the residuals of the reflection model. (b) Shows the νF_ν plot with individual model components: blackbody (dashed and dotted), *comptt* (dashed), and *relxill* (dotted).




Figure 8.9: The change in χ^2 versus the change in inner radius, showing that the model is consistent with the disk extending to the ISCO. The dashed blue line marks 90% confidence and the dashed red line marks 99.7% confidence (3σ). (a) Shows the results of the *reflionx* model (model 1). (b) Shows the results of the *relxill* model (model 5).

ity, $L_{\text{EDD}} = 2.5 \times 10^{38} \text{ erg s}^{-1}$ (van Paradijs & McClintock, 1994). We note that this value is around an average luminosity for 4U 1728–34 compared to previous observations. Di Salvo et al. (2000) measured $L_{0.1-100} = 2 \times 10^{37} \text{ erg s}^{-1}$. Others have used different energy ranges; for the sake of comparison, we recalculate the unabsorbed flux and luminosity in different ranges. Eggen et al. (2011) measured $L_{0.1-150} = 5 \times 10^{36} \text{ erg s}^{-1}$, in comparison to our value of $L_{0.1-150} = 1.9 \times 10^{37} \text{ erg s}^{-1}$. Piraino et al. (2000) measured $L_{0.2-50} = 3.7 \times 10^{37} \text{ erg s}^{-1}$, while we measure of $L_{0.2-50} \sim 1.7 \times 10^{37} \text{ erg s}^{-1}$. Ng et al. (2010) measured $L_{2-10} = 7.8 \times 10^{35} \text{ erg s}^{-1}$, compared to our measurement of $L_{2-10} \sim 6.6 \times 10^{36} \text{ erg s}^{-1}$.

With our upper limit on the accretion disk inner radius $R_{\text{in}} \leq 23 \text{ km}$, we infer an upper limit on the neutron star magnetic field strength by equating the magnetic field pressure to the ram pressure of the accretion disk (Illarionov & Sunyaev, 1975). We use the following relation:

$$R_{\text{in}} = 4 \times 10^8 B_{11}^{4/7} \dot{m}_{15}^{-2/7} M^{-1/7} \text{ cm} \quad (8.1)$$

where B_{11} is the magnetic field in units of 10^{11} G , \dot{m}_{15} is the mass accretion rate in units of 10^{15} g s^{-1} , and M is the neutron star mass in solar masses. We calculate the mass accretion rate from the luminosity given above with the relation $L = \eta \dot{m} c^2$, where $\eta = GM/Rc^2$. Assuming $M = 1.4M_{\odot}$, we find $B \leq 2 \times 10^8 \text{ G}$.

We compare our results with those obtained for other neutron star LMXBs also observed with NuSTAR. Miller et al. (2013) found that Serpens X-1, a persistent source, has a disk extending almost to the ISCO when observed at a luminosity of $L \sim 0.44L_{\text{EDD}}$. From their measured values of L and R_{in} , we estimate $B \leq 2.2 \times 10^8 \text{ G}$. Degenaar et al. (2015) found that 4U 1608-52, a transient source with a spin frequency of 620 Hz, also has a disk extending close to the ISCO when observed at a luminosity of $L \sim 0.02L_{\text{EDD}}$; we estimate $B \leq 1.2 \times 10^8 \text{ G}$. Aql X-1, also a transient observed by King et al. (2016) at a luminosity of $L \sim 0.08L_{\text{EDD}}$, has a truncated disk ($R_{\text{in}} \sim 15R_g$), a spin frequency of 550 Hz, and a magnetic field $B \leq 5 \pm 2 \times 10^8 \text{ G}$.

In using Equation 1, we assume that the magnetic pressure truncates the accretion disk (Illarionov & Sunyaev, 1975). If this is the case, we expect to see pulsations, yet of the above sources only Aql X-1 has been observed to exhibit pulsations during outburst (Casella et al., 2008). Papitto et al. (2013) note that exhibiting pulsations is rare among LMXBs and consider various explanations, including the possibility that the spin axis is aligned with the magnetic axis (Lamb et al., 2009), that pulsations do exist but the pulse amplitude is below

the detectable threshold, or that magneto-hydrodynamical instabilities cause material to accrete onto the neutron star at random places instead of being channeled by the magnetic field lines (Romanova et al., 2008). The lack of pulsations in 4U 1728–34, 4U 1608-52 and Serpens X-1 may indicate that our assumption of the magnetic pressure truncating the disk is not correct or at least not the complete physical picture. In that case, the true magnetic field values of these sources are most likely somewhat lower than the upper limits quoted above.

According to White & Zhang (1997), a higher spin frequency should imply a lower magnetic field: with a low magnetic field, the disk can extend deeper into the potential, spinning up the neutron star. While it is consistent that the disk extends closer to the ISCO in the above sources with lower magnetic fields, the correlation between spin frequency and magnetic field is not always followed. While 4U 1608-52 has the highest spin frequency and lowest magnetic field, Aql X-1 has both a higher spin frequency and a higher magnetic field than 4U 1728–34. (The spin frequency of Serpens X-1 is unknown, so we leave it out of this comparison.)

It is likely that this discrepancy is due to the expected variability in LMXB spin periods and magnetic fields. Additionally, as we are considering upper limits on the magnetic fields, we recognize that better estimates of the magnetic field values could negate this discrepancy. It is possible, however, that the discrepancy is real, in which case we consider the effects of magnetic field screening (Cumming et al., 2001) as an explanation. Magnetic field screening is a process by which the accreting matter becomes magnetized slowly compared to the accretion rate, causing the field implied by equating the pressures to be orders of magnitude smaller than the true field of the neutron star. Because Aql X-1 is a transient, we consider the possibility that the magnetic field emerges during quiescence and is not screened immediately after the outburst begins. The time it takes for the magnetic field to become significantly screened depends on accretion rate; for an accretion rate of $0.08L_{\text{EDD}}$, the field will be screened by one order of magnitude five days or less after the outburst began (Cumming, 2008). The observation of Aql X-1 occurred less than five days after the outburst began; hence it is possible that Aql X-1 is not fully screened (though a better estimate of the screening timescale could further support, or refute, this hypothesis). This possibility provides an explanation for the low spin frequency and magnetic field of 4U 1728–34 as compared to Aql X-1.

We consider possible mechanisms for accretion onto the neutron star. Because the magnetic field is relatively low, it is possible that the material is getting to the neutron star via a magnetic gate (Lamb et al., 1977). Due to chaotic accretion on the stellar surface, type II X-ray bursts are expected in the magnetic gate model, yet are not exhibited in 4U 1728–34. Instead, the material could be channeled along the magnetic field lines to the poles (Lamb et al., 2009). This scenario would cause a hot spot on the magnetic pole. However, it is possible that the magnetic axis is aligned with the rotation axis as 4U 1728–34 does not emit regular pulsations.

SUMMARY AND OUTLOOK

Observations of neutron stars and black holes allow us to study how matter behaves in extremely high gravitational and magnetic fields that we are currently unable to create in the laboratory. These relatively small but considerably dense compact objects often emit strongly in the γ -rays, and this high energy radiation is an excellent probe of the non-thermal processes occurring within these systems. While spectral, temporal, and imaging analysis can help us discern what we currently do not know about these systems (see e.g. Chapter 8), γ -ray polarimetry offers a new look into compact objects. With γ -ray polarimetry, we can determine source geometries, emission mechanisms, and magnetic field structure, information that is not always attainable with other types of analysis.

Historically, γ -ray instruments were not designed as polarimeters, leading to a dearth of γ -ray polarization measurements of astrophysical sources. In recent years, new instruments have been proposed, built, and launched designed to perform these intriguing observations. The Compton Spectrometer and Imager ([COSI](#)) is one such instrument. In 2016, [COSI](#) was launched from Wanaka, New Zealand, on [NASA's](#) new [SPB](#) and flew for 46 days. The Crab nebula, the Galactic black hole Cygnus X-1, and the active galactic nucleus ([AGN](#)) Centaurus A are among the sources detected in the [COSI](#) data.

Analyzing data collected by a new instrument is challenging: new algorithms and methods must be developed and written, and we must ensure that we understand the data that we collect. I developed a detailed detector effects engine ([DEE](#)) to apply the measured detector performance to Monte Carlo simulations, ensuring that the simulations we use to calculate the instrument response are accurate. I also developed a spectral analysis pipeline for [COSI](#) data and successfully fit the spectrum of [GRB 160530A](#), recovering best fit parameters that are consistent with those from other instruments. To perform spectral and polarization analysis of persistent sources, we must extract the signal from the copious atmospheric background present in balloon environ-

ments. To this end, I developed a background subtraction algorithm based on methods from the COMPTEL collaboration.

Using simulations, I established that the minimum detectable polarization (MDP) from the Crab observation is too high to perform polarimetry on this dataset. The large MDP values stem from the Earth horizon cut (EHC), an algorithm that reduces the albedo radiation background at the expense of rejecting the events that give us the clearest polarization signal. The abundant atmospheric background also leads to severely background-dominated observations of even the brightest astrophysical sources (such as the Crab), which further complicates extracting the source events from the background. With COSI deployed on a satellite rather than a balloon, we would be above the large atmospheric γ -ray background and would thus have much more success with spectral and polarization analysis. Thicker shielding around the bottom and sides of the detector would remove the need for the EHC. Alternatively, a new approach to the albedo radiation background rejection to replace the EHC could allow us to make these measurements from a balloon platform, potentially with the COSI data from the 2016 flight.

A version of COSI deployed on a satellite, referred to as COSI-SMEX at present, is currently being proposed. If COSI-SMEX were to fly, the prospects for polarimetry would be significantly improved for the reasons discussed above. Additionally, the analysis techniques developed in this thesis would be easily applied to the satellite-based version of the instrument. More generally, this work can be applied to any future Compton telescopes, such as the All-sky Medium Energy Gamma-ray Observatory (AMEGO), another proposed satellite-based γ -ray instrument that utilizes Compton scattering and is thus inherently a polarimeter. Hopefully a space-based γ -ray polarimeter will emerge from the proposed mission concepts to brighten the future of γ -ray polarimetry against the background.

BIBLIOGRAPHY

- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2010, *The Astrophysical Journal*, 720, 912, doi: [10.1088/0004-637x/720/1/912](https://doi.org/10.1088/0004-637x/720/1/912)
- Abramowicz, M. A., Czerny, B., Lasota, J.-P., & Szuszkiewicz, E. 1988, *The Astrophysical Journal*, 332, 646
- Abramowicz, M. A., & Fragile, P. C. 2013, *Living Reviews Relativity*, 16
- Agostinelli, S., Allison, J., Amako, K., et al. 2003, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506, 250, doi: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8)
- Aharonian, F., Akamatsu, H., Akimoto, F., et al. 2018, *Publications of the Astronomical Society of Japan*, 70, doi: [10.1093/pasj/psy118](https://doi.org/10.1093/pasj/psy118)
- Amman, M., & Luke, P. N. 2000, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 452, 155
- Amman, M., Luke, P. N., & Boggs, S. E. 2007, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 579, 886
- Arnaud, K., Gordon, C., & Dorman, B. 2018, *XSPEC: An X-ray Spectral Fitting Package*, HEASARC Astrophysics Science Division, NASA / GSFC
- Arnaud, K. A. 1996, *Astronomical Data Analysis Software and Systems V*, ASP Conference Series, 101
- Baade, W., & Zwicky, F. 1934, *Physical Review*, 45, 138
- Band, D., Matteson, J., Ford, L., et al. 1993, *Astrophysical Journal*, 413, 281
- Bandstra, M. S. 2010, PhD thesis, University of California, Berkeley
- Bandstra, M. S., Bellm, E. C., Boggs, S. E., et al. 2011, *The Astrophysical Journal*, 738, 8

- Bardeen, J. M., Press, W. H., & Teukolsky, S. A. 1972, *The Astrophysical Journal*, 178, 347
- Beckmann, V., Jean, P., Lubinski, P., Soldi, S., & Terrier, R. 2011, *Astronomy and Astrophysics*, 531, A70
- Beloborodov, A. M. 1998, *Monthly Notices of the Royal Astronomical Society*, 297, 739
- Berger, M. J., Hubbell, J. H., Seltzer, S. M., et al. 2010, XCOM: Photon Cross Section Database (version 1.5). <http://physics.nist.gov/xcom>
- Boggs, S. E., & Jean, P. 2000, *Astron. Astrophys. Suppl. Ser.*, 145, 311, doi: [10.1051/aas:2000107](https://doi.org/10.1051/aas:2000107)
- Bondi, H., & Hoyle, F. 1944, *Monthly Notices of the Royal Astronomical Society*, 104, 273
- Bowen, J. D. 2009, PhD thesis, University of California, Berkeley
- Bowen, J. D., Zoglauer, A., Bandstra, M. E., et al. 2006, in *Bulletin of the American Astronomical Society*, Vol. 38, High Energy Astrophysics Division Meeting #9, 385
- Brenneman, L. W., & Reynolds, C. S. 2006, *The Astrophysical Journal*, 652, 1028
- Brun, R., & Rademakers, F. 1997, in *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Vol. 389, 81–86
- Burrows, D. N., Hill, J. E., Nousek, J. A., et al. 2005, *Space Science Reviews*, 120, 165, doi: [10.1007/s11214-005-5097-2](https://doi.org/10.1007/s11214-005-5097-2)
- Cackett, E. M., Miller, J. M., Ballantyne, D. R., et al. 2010, *The Astrophysical Journal*, 720, 205, doi: [10.1088/0004-637X/720/1/205](https://doi.org/10.1088/0004-637X/720/1/205)
- Casella, P., Altamirano, D., Patruno, A., Wijnands, R., & van der Klis, M. 2008, *The Astrophysical Journal*, 674, L41
- Chauvin, M., Floén, H.-G., Friis, M., et al. 2017, *Scientific Reports*, 7
- . 2018, *Monthly Notices of the Royal Astronomical Society: Letters*, 477, L45

- Comella, J. M., Craft, H. D., Lovelace, R. V. E., Sutton, J. M., & Tyler, G. L. 1969, *Nature*, 221, 453
- Compton, A. H. 1923, *Physical Review*, 21, 483
- Corbel, S., Coriat, M., Brocksopp, C., et al. 2013, *Monthly Notices of the Royal Astronomical Society*, 428, 2500
- Cumming, A. 2008, *AIP Conference Proceedings*, 1068, doi: [10.1063/1.3031186](https://doi.org/10.1063/1.3031186)
- Cumming, A., Zweibel, E., & Bildsten, L. 2001, *The Astrophysical Journal*, 557, 958
- D'Aí, A., Di Salvo, T., Iaria, R., et al. 2006, *Astronomy and Astrophysics*, 448, 817, doi: [10.1051/0004-6361:20053228](https://doi.org/10.1051/0004-6361:20053228)
- Dame, T. M., Hartmann, D., & Thaddeus, P. 2001, *Astronomy and Astrophysics*, 547, 792
- Dean, A. J., Clark, D. J., Stephen, J. B., et al. 2008, *Science*, 321, 1183
- Degenaar, N., Miller, J. M., Chakrabarty, D., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 451, L85
- Di Salvo, T., Iaria, R., Burderi, L., & Robba, A. 2000, *The Astrophysical Journal*, 542, 1034, doi: [10.1086/317029](https://doi.org/10.1086/317029)
- Dove, J. B., Wilms, J., & Begelman, M. C. 1997, *The Astrophysical Journal*, 487
- Dyks, J., Harding, A. K., & Rudak, B. 2004, *The Astrophysical Journal*, 606, 1125
- Eggen, E., Di Salvo, T., Burderi, L., et al. 2011, *Astronomy and Astrophysics*, 530, 7, doi: [10.1051/0004-6361/201016093](https://doi.org/10.1051/0004-6361/201016093)
- Esin, A. A., McClintock, J. E., & Narayan, R. 1997, *The Astrophysical Journal*, 489, 865
- Fabian, A. C., Rees, M. J., Stella, L., & White, N. E. 1989, *Monthly Notices of the Royal Astronomical Society*, 238, 729
- FITS Working Group. 2016, *Definition of the Flexible Image Transport System (FITS)*, 4th edn.

- Forot, M., Laurent, P., Grenier, I. A., Gouiffés, C., & Lebrun, F. 2008, *The Astrophysical Journal*, 688, L29
- Frank, J., King, A., & Raine, D. 2002, *Accretion Power in Astrophysics* (Cambridge University Press)
- Gaensler, B. M., & Slane, P. O. 2006, *Annual Review of Astronomy and Astrophysics*, 44
- Gallo, E., Fender, R. P., & Pooley, G. G. 2003, *Monthly Notices of the Royal Astronomical Society*, 344, 60
- Galloway, D. K., Psaltis, D., Chakrabarty, D., & Munro, M. P. 2003, *The Astrophysical Journal*, 590, 999, doi: [10.1086/375049](https://doi.org/10.1086/375049)
- Galloway, D. K., Yao, Y., Marshall, H., Misanovic, Z., & Weinberg, N. 2010, *The Astrophysical Journal*, 724, 417, doi: [10.1088/0004](https://doi.org/10.1088/0004)
- García, J., Dauser, T., Lohfink, A., et al. 2014, *The Astrophysical Journal*, 782, 76
- Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, *The Astrophysical Journal*, 611, 1005, doi: [10.1086/422091](https://doi.org/10.1086/422091)
- Gold, T. 1968, *Nature*, 218, 731
- Goldreich, P., & Julian, W. H. 1969, *The Astrophysical Journal*, 157
- Gull, S. F., & Skilling, J. 1984, *IEEE Proceedings F (Communications, Radar and Signal Processing)*, 131, 646
- Harding, A. K. 2004, in *22nd Texas Symposium on Relativistic Astrophysics at Stanford University*
- Harrison, F. A., Craig, W. W., Christensen, F. E., et al. 2013, *The Astrophysical Journal*, 770, 103, doi: [10.1088/0004-637X/770/2/103](https://doi.org/10.1088/0004-637X/770/2/103)
- Hasinger, G., & van der Klis, M. 1989, *Astronomy and Astrophysics*, 225, 79
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709
- Illarionov, A. F., & Sunyaev, R. A. 1975, *Astronomy and Astrophysics*, 39, 185

- Jourdain, E., Roques, J. P., Chauvin, M., & Clark, D. J. 2012, *The Astrophysical Journal*, 761, 27
- Kalberla, P. M. W., Burton, W. B., Hartmann, D., et al. 2005, *Astronomy and Astrophysics*, 440, 775
- Kalogera, V., & Baym, G. 1996, *The Astrophysical Journal*, 470, L61
- Kanbach, G., Slowikowska, A., Kellner, S., & Steinle, H. 2005, *AIP Conference Proceedings*, 801, 306
- Kargaltsev, O., & Pavlov, G. G. 2008, *AIP Conference Proceedings*, 983, 171
- Kaspi, V. M., Roberts, M. S. E., & Harding, A. K. 2006, *Compact Stellar X-ray Sources*, ed. W. H. G. Lewin & M. van der Klis (Cambridge University Press)
- Kierans, C. 2018, PhD thesis, University of California, Berkeley
- King, A. L., Tomsick, J. A., Miller, J. M., et al. 2016, *The Astrophysical Journal*, 819, L29
- Klein, O., & Nishina, Y. Z. 1929, *Zeitschrift für Physik*, 52, 853
- Knödseder, J., von Ballmoos, P., Diehl, R., et al. 1996, in *Proceedings of the SPIE*, Vol. 2806, *Gamma-Ray and Cosmic-Ray Detectors, Techniques, and Missions*
- Knoll, G. F. 2000, *Radiation Detection and Measurement* (John Wiley and Sons, Inc.)
- Krawczynski, H. 2011, *Astroparticle Physics*, 34, 784 , doi: <https://doi.org/10.1016/j.astropartphys.2011.02.001>
- Lamb, F. K., Boutloukos, S., Van Wassenhove, S., et al. 2009, *The Astrophysical Journal*, 706, 417
- Lamb, F. K., Fabian, A. C., Pringle, J. E., & Lamb, D. Q. 1977, *The Astrophysical Journal*, 217, 197
- Lattimer, J. M., & Prakash, M. 2007, *Physics Reports*, 442, 109
- Laurent, P., Rodriguez, J., Wilms, J., et al. 2011, *Science*, 332, 438

- Lei, F., Dean, J., & Hills, G. L. 1997, *Space Sciences Reviews*, 82, 309
- Lewin, W. H. G., Clark, G., & Doty, J. 1976, *IAUC*, 2922
- Ling, J. C. 1975, *Journal of Geophysical Research*, 80, 3241
- Looker, Q., Amman, M., & Vetter, K. 2015, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 781, 20, doi: <https://doi.org/10.1016/j.nima.2015.01.069>
- Lowell, A., Boggs, S., Chiu, J. L., et al. 2016, in *Proc.SPIE*
- Lowell, A. W. 2017, PhD thesis, University of California, Berkeley
- Lowell, A. W., Boggs, S. E., Chiu, C. L., et al. 2017a, *The Astrophysical Journal*, 848, 119
- . 2017b, *The Astrophysical Journal*, 848, 120
- Markoff, S., Nowak, M. A., & Wilms, J. 2005, *The Astrophysical Journal*, 635, 1203
- McClintock, J. E., & Remillard, R. A. 2006, *Compact Stellar X-ray Sources*, ed. W. H. G. Lewin & M. van der Klis (Cambridge University Press), 157–214
- Millan, R. M., & Thorne, R. M. 2007, *Journal of Atmospheric and Solar-Terrestrial Physics*, 69, 362
- Miller, J. M. 2007, *Annual Review of Astronomy and Astrophysics*, 45, 441
- Miller, J. M., Parker, M. L., Fuerst, F., et al. 2013, *The Astrophysical Journal Letters*, 779, L2, doi: [10.1088/2041-8205/779/1/L2](https://doi.org/10.1088/2041-8205/779/1/L2)
- Mondal, A. S., Pahari, M., Dewangan, G. C., Misra, R., & Raychaudhuri, B. 2017, *Monthly Notices of the Royal Astronomical Society*, 466, 4991
- Narayan, R., & McClintock, J. E. 2015, *Relativistic Astrophysics*, ed. A. Ashtekar, B. K. Berger, J. Isenberg, & M. MacCallum (Cambridge University Press), 97–161
- Narayan, R., McClintock, J. E., & Yi, I. 1996, *The Astrophysical Journal*, 457, 821

- Narayan, R., & Yi, I. 1994, *The Astrophysical Journal Letters*, 428, L13
- . 1995, *The Astrophysical Journal*, 444, 231
- Narita, T., Grindlay, J. E., & Barret, D. 2001, *The Astrophysical Journal*, 547, 420
- Negueruela, I. 2010, in *High Energy Phenomena in Massive Stars ASP Conference Series*, ed. J. Marti, P. L. Luque-Escamilla, & J. A. Combi, Vol. 422
- Ng, C., Diaz Trigo, M., Cadolle Bel, M., & Migliari, S. 2010, *Astronomy and Astrophysics*, 522, 25, doi: [10.1051/0004-6361/200913575](https://doi.org/10.1051/0004-6361/200913575)
- Ng, C.-Y., & Romani, R. W. 2004, *The Astrophysical Journal*, 601, 479
- Novick, R., Weisskopf, M. C., Berthelsdorf, R., Linke, R., & Wolff, R. S. 1972, *The Astrophysical Journal*, 174, L1
- Novikov, I. D., & Thorne, K. S. 1973, *Black Holes*, ed. C. DeWitt & B. DeWitt (Gordon and Breach)
- Papitto, A., D'Ai, A., Di Salvo, T., et al. 2013, *Monthly Notices of the Royal Astronomical Society*, 459, 3411
- Picone, J. M., Hedin, A. E., Drob, D. P., & Aikin, A. C. 2002, *Journal of Geophysical Research*, 107, 1468
- Piraino, S., Santangelo, A., & Kaaret, P. 2000, *Astronomy and Astrophysics*, 360, L35
- Pringle, J. E., & Rees, M. J. 1972, *Astronomy and Astrophysics*, 21, 1
- Psaltis, D. 2006, *Compact Stellar X-ray Sources*, ed. W. H. G. Lewin & M. van der Klis (Cambridge University Press)
- Quataert, E., & Narayan, R. 1999, *The Astrophysical Journal*, 520, 298
- Ravi, V., Manchester, R. N., & Hobbs, G. 2010, *The Astrophysical Journal Letters*, 716, L85
- Reynolds, C. S., & Nowak, M. A. 2003, *Physics Reports*, 377, 389
- Rhoades, C. E., & Ruffini, R. 1974, *Physical Review Letters*, 32, 324
- Rodrigues, O. 1840, *J Mathematiques Pures Appliquees*, 5, 380

- Rodriguez, J., Grinberg, V., Laurent, P., et al. 2015, *The Astrophysical Journal*, 807, 17
- Romanova, M. M., Kulkarni, A. K., & Lovelace, R. V. E. 2008, *The Astrophysical Journal*, 673, L171
- Ross, R. R., & Fabian, A. C. 2005, *Monthly Notices of the Royal Astronomical Society*, 358, 211, doi: [10.1111/j.1365-2966.2005.08797.x](https://doi.org/10.1111/j.1365-2966.2005.08797.x)
- Sartore, N., & Treves, A. T. 2012, *Astronomy and Astrophysics*, 539, A52
- Schönfelder, V., Aarts, H., Bennett, K., et al. 1993, *The Astrophysical Journal Supplement Series*, 86, 657
- Seifina, E., & Titarchuk, L. 2011, *The Astrophysical Journal*, 738, 128, doi: [10.1088/0004-637X/738/2/128](https://doi.org/10.1088/0004-637X/738/2/128)
- Shakura, N. I., & Sunyaev, R. A. 1973, *Astronomy and Astrophysics*, 24, 337
- Shapiro, S. L., & Teukolsky, S. A. 1983, *Black Holes, White Dwarfs, and Neutron Stars: the Physics of Compact Objects* (John Wiley and Sons, Inc.)
- Sizun, P., Shrader, C. R., Attié, D., et al. 2004, *Proceedings of the 5th INTEGRAL Workshop on the INTEGRAL Universe* (ESA SP-552)
- Sleator, C. C., Tomsick, J. A., King, A. L., et al. 2016, *The Astrophysical Journal*, 827, 134
- Sleator, C. C., Zoglauer, A., Lowell, A. W., et al. 2019, *Nuclear Instruments and Methods Section A*, accepted pending revisions
- Slowikowska, A., Kanbach, G., Kramer, M., & Stefanescu, A. 2008, *AIP Conference Proceedings*, 984, 51
- Stirling, A. M., Spencer, R. E., de la Force, C. J., et al. 2001, *Monthly Notices of the Royal Astronomical Society*, 327, 1273
- Storn, R., & Price, K. 1997, *Journal of Global Optimization*, 11, 341
- Strohmayer, T. E., Zhang, W., Swank, J. H., et al. 1996, *The Astrophysical Journal*, 469, L9, doi: [10.1086/310261](https://doi.org/10.1086/310261)
- Svinkin, D., Golenetskii, S., Aptekar, R., et al. 2016, *GRB Coordinates Network, Circular Service*, 19477

- Tarana, A., Belloni, T., Bazzano, A., Méndez, M., & Ubertini, P. 2011, *Monthly Notices of the Royal Astronomical Society*, 416, 873
- Tomsick, J. A., & the COSI team. 2016, *GRB Coordinates Network, Circular Service*, 19473
- Vadawale, S. V., Chattopadhyay, T., Mithun, N. P. S., et al. 2018, *Nature Astronomy*, 2, 50
- van Paradijs, J., & McClintock, J. E. 1994, *Astronomy and Astrophysics*, 290, 133
- Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, *The Astrophysical Journal*, 465, 487, doi: [10.1086/177435](https://doi.org/10.1086/177435)
- Weisskopf, M. C., Cohen, G. G., Kestenbaum, H. L., et al. 1976, *The Astrophysical Journal*, 208
- Weisskopf, M. C., Elsner, R. F., & O'Dell, S. L. 2010, *Proc. SPIE 7732, Space Telescopes and Instrumentation 2010: Ultraviolet to Gamma Ray*, 77320E
- Weisskopf, M. C., Silver, E. H., Kestenbaum, H. L., Long, K. S., & Novick, R. 1978, *The Astrophysical Journal*, 220, L117
- White, N. E., & Zhang, W. 1997, *The Astrophysical Journal*, 490, L87
- Wilderman, S., Clinthorne, N. H., Fessler, J. A., & Rogers, W. L. 1998, in *IEEE Nuclear Science Symposium and Medical Imaging Conference*, Vol. 3, 1716–1720
- Wilms, J., Allen, A., & McCray, R. 2000, *The Astrophysical Journal*, 542, 914, doi: [10.1086/317016](https://doi.org/10.1086/317016)
- Zoglauer, A. 2005, PhD thesis, Technische Universität München
- Zoglauer, A., Andritschke, R., & Schopper, F. 2006, *New Astronomy Reviews*, 50, 629
- Zoglauer, A., & Boggs, S. E. 2007, in *IEEE Nuclear Science Symposium Conference Record*, Vol. 6, 4436–4441
- Zoglauer, A., Boggs, S. E., Galloway, M., et al. 2011, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 652, 568