Development of hard X-ray imaging detectors for the High Energy Focusing Telescope

Thesis by

Chi Ming Hubert Chen

In Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy



California Institute of Technology Pasadena, California

2008

(Defended Thursday, 10th January, 2008)

© 2008 Chi Ming Hubert Chen All Rights Reserved

Acknowledgements

This thesis describes an experiment involving a large group of researchers, of which the PhD candidate is but one member, but the one bearing the task of documenting the experiment. Credits are due to everyone involved.

My advisor, Prof Fiona Harrison, is the principal investigator of HEFT, and the leader of the detector development team. The project would not have been devised nor carried on to the flight campaign without her planning and execution. I am also in debt to her for the many opportunities she has given me to develop my own scientific career, opportunities of which perhaps I have not fully taken advantage.

Members of the HEFT detector development team at Caltech also include Drs Stephen Schindler, Walter Rick Cook, Aleksey Bolotnikov, Steven Boggs and Wayne Baumgartner, engineers Jill Burnham, Branislav Keçman and Vincent Nguyen, members of the technical staff Ara Badalian, Matvey Farber and John Klemic, and graduate student Lewis (Luke) Kotredes.

Dr Schindler served as the manager of the detector team until his retirement in November, 2002; he was also responsible for the design of various mechanical aspects of HEFT. Dr Cook was the chief designer of the HEFT read-out ASIC, while Jill implemented Rick's circuit schematics into VLSI layouts. The good spectral performance of the HEFT detectors was due to a large part to their effort. Dr Cook also designed and programmed the MISC processors for detector operation. Dr Bolotnikov carried out the bulk of the testing of detectors and hybrids until his departure in June, 2002. Even to this day, much of our knowledge of the material properties of cadmium zinc telluride came from his experiments. I am also in debt to him for the frequent mentoring and personally care he has given me during the first three years of my graduate studies. The detector model I describe in Chapter 6 is an extension (from 2D to 3D) of the graduate work of Dr Boggs; he also analysed the expected performance of the graded-Z shield in its initial design. Dr Baumgartner took over many critical tasks in June, 2004, shortly before the first flight campaign, at a time when we were seriously short-handed. He carried out the testing and calibration of the shield assemblies, he completed the fabrication and testing of the focal-plane pressure vessel, he, with Matvey Farber, fabricated the focal-plane insulation, he set up the cooling apparatus for in-field calibration and arranged for calibration sources to be available during the campaigns, he devised the target list for in-flight observation, and he was involved in the post-flight data analysis, in particular, the analysis of the background level and various attenuation factors. The modelled background spectrum in Fig. 7.12 was his work. I also thank him for his frequent guidance and friendship during and after the tough years of the flight campaigns.

The indium-bump bonding experiment I describe in Section 3.5.1.1 was the painstaking work of Branislav and Vincent. In addition, Jill and Vincent designed and implemented the electrical system—circuit boards, cabling, sensors and power supply—for the HEFT focal-plane, while Branislav handled mechanical engineering tasks. Jill also served as an effective manager during much of the flight campaigns. I thank her in addition for lending a sympathetic ear during those hard times, giving me spiritual support whenever I needed it.

Ara made mechanical drawings of focal-plane components until 2001. John took over many mechanical tasks in 2002, and relieved me of many hours of painstaking precision work under a microscope and with AutoCAD. He provided Fig. 7.1 and modified Ara's drawing for Fig. 2.6 for this thesis. Matvey joined our group in June, 2004, at the same critical time as Dr Baumgartner. It is by no means an exaggeration to say that the HEFT flight in 2005 would not have been possible without Matvey, who executed so many crucial tasks with great enthusiasm, from electrical work such as cabling and soldering, to mechanical work on the entire payload, as well as logistics and transportation. All photographs of instruments in Chapter 2 were Matvey's. His high spirit affected us all.

Luke wrote the routines to extract housekeeping data from the science data stream between the two flight campaigns. He was also involved in post-flight data analysis; he provided the data for the in-flight background spectrum in Fig. 7.12. We, the HEFT detector team, are also in debt to Prof Brent Fultz of the Department of Material Science at Caltech and his graduate students, Jiao Lin and Ryan Monson. Jiao and I were equal partners when we tested detectors and hybrids throughout 2003, producing the work I describe in Chapters 4 and 5. Ryan continued their project with cadmium telluride hybrid detectors as the HEFT team proceeded to prepare for the flight campaigns. Drs Peter Mao and Sarah Yost also helped with the grudging task of detector leakage measurements in 2002 and 2003, respectively. Dr Andrew Davis wrote an excellent real-time data browser that made up the ground station equipment suite; I thank him for bailing me out at a time of crisis. SURF student Alden Waters contributed to leakage measurement, hybrid testing and flight preparation in 2004, and to detector modelling (with the objective of depth sensing) in 2005. Carl Ge of Troy High School in Fullerton was a great help during the three months before we shipped out for our 2004 campaign.

HEFT is a multi-institute collaboration; our colleagues at Columbia University, the Lawrence Livermore National Laboratory (LLNL) and the Danish Space Research Institute (DSRI) have all played important roles. Chonko (2006) and Kruse Madsen (2007) acknowledge contributions to the HEFT mirrors and attitude control system. To that, I add my thanks to Todd Decker of Livermore for being an effective and understanding campaign manager. Fig. 7.6 was the work of Dr Michael Pivovaroff of Livermore; Dr Jason Koglin of Columbia contributed to post-flight data analysis and provided Figs. 7.15 and 7.16, the right halves of Figs. 7.14 and 7.17, and the data for Fig. 2.3.

Chapter 8 of this thesis is a study of relic radio sources, of which I have no prior knowledge. I thank Dr Daniel Harris of the Harvard-Smithsonian Center for Astrophysics, who patiently guided me through the 19-month long project before our paper was finally accepted for publication.

I thank Prof Thomas Tombrello, who served on my candidacy committee, Prof Bradley Filippone, who served on my thesis committee, and Profs Marc Kamionkowski and Jonas Zmuidzinas, who served on both.

During my nine long years at Caltech, I am fortunate to be a member of the Space Radiation Laboratory. Ithank everyone at SRL, for the sense of community amongst us. I have had the pleasant company of Dr Peter Mao, Dr John Cortese, Dr Megan Eckart, Bradley Cenko, Dr Ryan Ogliore, Luke Kotredes and Dr Kristin Kruse Madsen as officemates. Dr Mao and Dr Eckart, in particular, gave me much spiritual support. The HEFT project would not have been possible without the dedicated administrative service of Cherylinn Rangel, Marjorie Miller, Stacia Rutherford, Caprece Anderson, Pamela Nickson and Debby Miles. Minerva Calderon and Kimberly Rubal have done a great job maintaining a stable computer network at SRL, enabling us to concentrate on our research. Elias Villegas kept our offices clean everyday; he never failed to cheer me up when I much needed it.

My work on HEFT involved four months of field work in Fort Sumner, New Mexico, nine weeks in Denmark, and about two weeks in Livermore. I thank Dr James Chonko, Dr Kristin Kruse Madsen, Jill, John, Luke, Matvey and Wayne, who formed the core campaign crew, for their comradery through hard times. Thanks to Dr Finn Christensen, Dr Carsten Jensen, everybody at DSRI, and our colleagues at LLNL for their kind hospitality during my visits.

My family is my anchor. I express the greatest gratitude towards my parents, my grandparents and my sisters. Hong Kong was not yet a prosperous society in my parents' times. Neither of them had the opportunity to higher education. Yet, through hard work and ingenuity, my parents managed to provide me and my sisters a comfortable life, excellent education, and freedom to pursue our loftiest dreams without having to worry about money. Meanwhile, my paternal grandparents took care of our daily needs, making sure that we had food to eat, clothes to wear, and kept us safe, while my parents toiled day and night. Various people dear to my heart, including Aunt Dida, Aunt Tina and the late Mrs Lin, also played significant roles in bringing me up. Without all that they have given unto me, I would never have made it this far.

We had been living with my paternal grandparents until we emigrated to Vancouver when I was sixteen. Since then, I had not been a good grandchild, seldom returning to visit my grandparents, thinking too importantly of myself, especially after I entered university. All this became irreparable in 2004, when my grandfather went into a coma in March, and later passed away in June. With the first HEFT campaign just a few months away, and components for four out of the six flight detectors still no where to be found (see Section 3.8), it was not possible for me to return to Hong Kong to spend time with my ailing grandfather without jeopardising the flight campaign. Yet, that did not stop me from falling into clinical depression, an abyss of despair that mere observers will never understand. I am grateful to Dr Karen Payne, who helped me to get through my depression and the subsequent bereavement during those months of hopelessness. I nearly quitted my studies. I do not know how I could have gone on without her help.

Living alone, thousands of miles from my family, I am grateful for the company of various 'dinner groups' over the years—the graduate residents of Avery House in 1999–2000, international students entering in the same year, the Hongkonger community at Caltech, and the many old friends who visited me in Southern California and in Fort Sumner—too many of you to name individually. It may seem trivial to some, but dinner groups mean a lot to me. (Think Virginia Wolfe's *To the Lighthouse.*) Amongst the many friends who walked with me over these nine years, a few were my surrogate family. I give special thanks to Dr Peter Mao and Dr Deirdre Scripture-Adams, Tommy Tam, Dr Ching Law, and David Lee, without whose spiritual support, it would have been much harder.

Abstract

The High Energy Focusing Telescope (HEFT) is a balloon-borne instrument carrying one of the first focusing telescopes for the hard X-ray band (20–70 keV). It makes use of tungsten-silicon multilayer coatings to extend the reflectivity of nested grazing-incidence mirrors beyond 10 keV. It also carries novel semiconductor pixel detectors on its focal planes to match the capabilities of the multilayer mirrors. After a decade of research and development, we achieve with HEFT an angular resolution of 1.5 arcminutes in half-power diameter, and an energy resolution of 1.0 keV full width at half maximum at 60 keV. We launched HEFT for a 25-hour balloon flight in May, 2005; the instrument performed within specification, and observed Cyg X-1, the Crab Nebula, and other celestial hard X-ray sources.

In this thesis, I lay out the scientific motivations for HEFT, and give an overview of the experiment. I report on our detector development effort in depth, and document the balloon flight of 2005. I also describe a study of two relic radio sources, 0917+75 and 1401-33, with data from the soft X-ray XMM-Newton observatory. With a hard X-ray focusing telescope like HEFT, one can improve the sensitivity and extend the scope of such studies to other classes of objects.

Contents

A	cknowledgements 111			
A	bstra	.ct		vii
1	\mathbf{Intr}	oducti	on	1
	1.1	The h	ard X-ray band: a new window for astronomy	1
		1.1.1	The X-ray background	1
		1.1.2	Core-collapse supernova remnants	2
		1.1.3	Inverse Compton scattering of cosmic background photons	4
		1.1.4	Cyclotron Resonance Scattering Features	5
		1.1.5	Non-thermal emission in the Crab Nebula	6
	1.2	Telesc	ope technology in hard X-ray astronomy	7
		1.2.1	The present: non-focusing technologies	9
		1.2.2	The future: focusing technologies	11
	1.3	The re	est of it all	15
2	\mathbf{The}	High	Energy Focusing Telescope (HEFT)	17
	2.1	A ball	oon-borne experiment	17
	2.2	Perfor	mance objectives	19
	2.3	Overv	iew of the payload	20
	2.4	The X	-ray mirrors	20
		2.4.1	Geometry	20
		2.4.2	Multilayer coating	23
		2.4.3	Performance	24
	2.5	The X	-ray focal plane	25
		2.5.1	The focal-plane modules	27
		2.5.2	The hard X-ray detector	28
		2.5.3	Digital circuitry	28
		2.5.4	Background shield assembly	29

	2.5.5	Environmental control	0
2.6	The p	ointing platform	2
	2.6.1	The telescope truss	2
	2.6.2	Attitude control	2
2.7	Summ	ary	6
3 Det	tector o	development for HEFT 3	7
3.1	Perfor	mance requirements	7
	3.1.1	Spatial requirements	7
	3.1.2	Spectral requirements	7
	3.1.3	Efficiency requirements	8
3.2	Choice	e of technology: semiconductor detectors	8
3.3	Noise	components in semiconductor detectors	9
	3.3.1	Counting statistics and shot noise 3	9
	3.3.2	Johnson (thermal) noise	9
	3.3.3	Leakage current	0
	3.3.4	Charge trapping	0
3.4	Choice	e of material $\ldots \ldots 4$	1
3.5	Achie	ving low noise	5
	3.5.1	Bonding techniques	6
		3.5.1.1 Indium bump bonding	6
		3.5.1.2 Flip-chip bonding with conductive epoxy and gold stud bumps 4	7
	3.5.2	Anode geometry	8
3.6	The fi	nal detector geometry	0
3.7	Proces	ssing signals from the HEFT detectors	1
	3.7.1	Basic concepts	1
	3.7.2	Implementation of hardware	1
	3.7.3	Systematics and noise components and their corrections	4
3.8	Fabric	ation of the focal planes	5
	3.8.1	ASIC production and screening	5
	3.8.2	CdZnTe detector production	7
	3.8.3	Detector leakage measurement	8
	3.8.4	Interconnections	1
3.9	Chara	cterization of the HEFT detectors	2
	3.9.1	Final detector ensemble	3
	3.9.2	Spectral resolution	3

		3.9.3 Charge induction amplitude	69
		3.9.4 X-ray event statistics	71
		3.9.5 Flatfielding	76
		3.9.6 Summary	78
4	Pro	ototype detector characterization: a best-case scenario	79
	4.1	Introduction	79
	4.2	Detector Characterization	80
		4.2.1 Electronic Noise	80
		4.2.2 Leakage Current	81
		4.2.3 Spectral Resolution	83
		4.2.4 Count Rate Uniformity	84
	4.3	Summary	87
5	Cor	mparison of two detector architectures	88
	5.1	Introduction	88
	5.2	Detector characterization	89
		5.2.1 Electronic noise	90
		5.2.1.1 Temperature dependence	90
		5.2.1.2 Leakage current contribution	91
		5.2.2 Spectral resolution for X-ray events	91
		5.2.2.1 Single-pixel events	91
		5.2.2.2 Charge-sharing events	94
		5.2.3 Optimization of operating biases for the gridded detector $\ldots \ldots \ldots \ldots$	96
	5.3	Summary	96
6	Nui	merical modelling of charge sharing in CdZnTe pixel detectors	98
	6.1	Introduction	98
	6.2	Detector geometry	99
	6.3	Phenomenology of charge-sharing	.00
		6.3.1 Experimental setup	.00
		6.3.2 Observations from our experiment	00
	6.4	The numerical model	.03
	6.5	Results and discussion	.05
	6.6	Conclusion	.09

7	The	first	HEFT flight campaigns	110
	7.1	Integr	ation \ldots	110
	7.2	Calibr	ation	112
		7.2.1	Detector calibration	112
		7.2.2	Optics-detector alignment	117
			7.2.2.1 Hybrid-hybrid alignment	117
			7.2.2.2 Focal plane-truss alignment	121
			7.2.2.3 Focal plane-mirror and X-ray-optical alignment	122
	7.3	The o	bservation run	126
	7.4	Result	${ m ss} { m from flight}$	131
		7.4.1	General quality of data	131
		7.4.2	In-flight background measurement	135
		7.4.3	Cyg X-1	135
		7.4.4	The Crab nebula	139
8	A s	cientif	ic case: study of the inverse Compton scattering of cosmic backgrour	ıd
	pho	tons b	y relativistic electrons	141
	8.1	Introd	\mathbf{v}	142
	8.2	The o	bservations	143
	8.3	Data	m analysis	143
		8.3.1	Event filtering	143
		8.3.2	Background determination and flux upper limit analysis	147
			8.3.2.1 0917+75	147
			8.3.2.2 1401-33	151
			8.3.2.3 Flux and field limits	152
	8.4	Discus	ssion	153
		8.4.1	0917+75	153
		8.4.2	1401-33	155
		8.4.3	On the flux and field limits	155
	8.5	Summ	ary	158
9	Cur	rent s	tatus and future outlook	160
Bi	bliog	graphy		164
А	CdZ	ZnTe d	letector specification	172

В	Ana	lysis o	of the HEFT amplifier and read-out circuitry	173
	B .1	Stage-	by-stage circuit analysis	173
		B.1.1	The preamplifier	173
		B.1.2	The postamplifier	175
		B.1.3	The shaping amplifier	177
		B.1.4	The discriminator	179
		B.1.5	The sampling capacitor bank	179
		B.1.6	The read-out amplifier	180
		B.1.7	Analogue to digital conversion	181
	B .2	Conve	rsion between quantities at different stages	181
		B.2.1	Signal amplitude \ldots	181
		B.2.2	Leakage current	181
		B.2.3	Saturation of the preamplifier	182
\mathbf{C}	Pro	cessing	g signals from the HEFT detectors: the details	183
	C.1	Impler	nentation of hardware	183
	C.2	System	natics and noise components and their corrections	184
		C.2.1	Truly random (thermal) noise	185
		C.2.2	Time of rise	185
		C.2.3	Postamplifier gain	187
		C.2.4	Temperature variation of gain	187
		C.2.5	'Detector gain'	187
		C.2.6	Zero energy offset	188
		C.2.7	Transfer function: a unified view	189
		C.2.8	MISC wrap around	189
		C.2.9	Capacitor offset ('baseline')	191
		C.2.10	Software discrimination	192
		C.2.11	Common-mode noise	192
		C.2.12	Depth sensing	193
	C.3	The re	esulting data reduction sequence	193
D	HE	FT scie	ence data formats	196
	D.1	Telem	etry format	196
	D.2	'Crush	\mathbf{r} ed' format	201
		D.2.1	Introduction	201
		D.2.2	Crushed event-data format	201
		D.2.3	Crushed pixel-data format	202

	D.2.4	Additional remarks	202
D.3	The H	EFT FITS file format	203
	D.3.1	Introduction	203
	D.3.2	Header keywords	203
	D.3.3	Column specification for event lists in BINTABLE HDUs	207
	D.3.4	References	209

List of Figures

1.1	Decay sequence from 44 Ti to 44 Ca	3
1.2	Cross-section of the Wolter-I mirror geometry. Grazing angles are exaggerated for	
	illustration.	12
2.1	Attenuation of hard X-rays at various energies by the atmosphere. The top graph shows	
	the column density of air at various altitudes, according to the 1976 US Standard	
	Atmosphere. The bottom graph shows the fraction of celestial X-rays remaining at	
	various altitudes. The five curves correspond, from right to left, to X-rays at 10, 25,	
	$40,55$ and $70{\rm keV}.$ Scientific balloons reach typical altitudes between 30 and $40{\rm km}.$	18
2.2	An annotated photograph of the HEFT instrument.	21
2.3	Total effective area of the three mirror modules of HEFT. Values shown here repre-	
	sent the efficiency of the mirror geometry, two-bounce reflectivity of the multilayer	
	coatings, and the various obscurations associated with the mirror geometry and me-	
	chanical structures.	24
2.4	Photographs of the telscope truss, viewed from the focal-plane side. The picture on the	
	left shows the single-piece aluminium flange with integrated strongbacks that attachs	
	to the truss and forms the focal plane	25
2.5	Photographs of the HEFT focal plane and its components. The left-hand and right-	
	hand pictures show the sides facing away from and towards the mirrors, respectively.	
	The right-hand picture also shows one of the two kevlar domes that bolt to the flange	
	to form the focal-plane pressure vessel.	26
2.6	Diagramme showing the cross-section of the background-shield module and the focal-	
	plane module within.	26
2.7	- Photographs of the HEFT focal-plane module and a bottom view of the shield module	
	well, where the focal-plane module resides.	27
2.8	Diagramme and flow chart of the detector system.	27
		· ·
3.1	A model of thermal noise and stray capacitance at the input of the preamplifier. \ldots	40

3.2Cumulative fraction of charge induced in a semiconductor detector across opposite electrode surfaces, with and without pixels on the anode plane. Distance from the cathode is normalized to the thickness of the detector between the cathode and anode planes. The presence of pixels forces charge induction to take place mostly when tranport electrons get near the pixels, thus reducing the hole signal relative to the electron signal. 423.3 Fraction of hard X-ray photons absorbed by a typical 2 mm-thick detector volume made of various semiconductor materials. 43 Count rate histogram from the flood illumination of a circular area ten pixels in diam-3.4eter on our cadmium telluride detector, with Am-241. 443.5Photograph of indium bumps deposited on the pixel contacts of a CdZnTe detector with the gridded anode geometry described in Section 3.5.2. The indium bumps are $50\,\mu\mathrm{m}$ in diameter, while the pixel pitch is $498\,\mu\mathrm{m}$. 473.6Left: stencil-printed conductive epoxy bumps on the $498 \,\mu$ m-pitch pixel contacts of a CdZnTe detector, prior to flip-chip bonding. The bumps are intentionally offset from the centre of the square contacts to accommodate the row of narrow pixels adjacent to the mating edge (see Section 3.6). Right: cross-section of a gold stud-conductive epoxy bond, made by slicing a mechanical sample hybrid vertically apart and imaging with a scanning electron microscope (SEM). 49Anode plane patterns. The diagramme on the left is a mechanical drawing, to-scale, of 3.7the final anode plane, with a raster array of 44×24 pixels per detector. The top circle shows an experimental pattern with steering electrodes described in Section 3.5.2; the bottom circle shows the pattern on the final HEFT detectors. Note that the guard ring and the first row of pixels from the mating edge in both designs are contracted for two detectors to be placed side by side with minimal dead area in between. 503.8Schematic diagramme of the HEFT ASIC read-out chip. Subcircuit A is the shaping and peak detection stages in a conventional amplifier chain; although this subcircuit exists in the HEFT ASIC, only a simple version is implemented for triggering. Instead, accurate pulse height information is captured in the bank of 16 switch capacitors in Subcircuit B, from which the pulse height is determined. Each pixel in the ASIC contains a copy of the circuit shown in this figure; all pixels share the same serial 52

3.9	Waveform sampled by sampling capacitors in the HEFT read-out circuitry. We recorded	
	this waveform in an experiment where we placed an americium-241 radioactive source	
	in front of a HEFT detector. The detected monoenergetic gamma rays at $59.54\mathrm{keV},$	
	roughly $30{\rm keV}$ (missing the energy of the K-shell escape photons), and below $20{\rm keV}$	
	(the neptunium lines) each produces a discrete waveform whose amplitude records the	
	energy of the detected photon. Multiple events, each triggering at a random time offset	
	with respect to the system clock, provide samples that fill the gaps between discrete	
	clock ticks. The dip immediately after the rising edge is caused by the undesirable	
	discharge of a subcircuit component	53
3.10	Flowchart of steps in the production of a HEFT focal plane.	56
3.11	Wiring diagramme of the apparatus for leakage measuremnts on the HEFT CdZnTe	
	detectors	58
3.12	$Photograph\ through\ a\ microscope\ of\ micromanipulator\ probe\ tips\ connecting\ the\ anode$	
	contacts of a CdZnTe detector to a picoammeter and ground connections	59
3.13	The decay of leakage current at a detector pixel over two hours.	60
3.14	Map of leakage current values across a HEFT CdZnTe detector. Squares with different	
	levels of shading represent the size of the leakage current measured at that particular	
	location. Where individual pixel measurement is not made, the background shading	
	$\operatorname{corresponds}$ to the average leakage current per pixel, inferred from the aggregate leakage	
	current measured from the cathode.	60
3.15	Spectra of Am-241 at -15 (left column) and -5 °C (right column), from each detector	
	with events from all pixels grouped together. Each row of figures show data from the	
	same flight detector.	65
3.16	Spectra of Am-241 at $+5$ °C (left column) and of Co-57 at -15 °C (right column), from	
	each detector with events from all pixels grouped together. Each row of figures show	
	data from the same flight detector	66
3.17	Spectra of Am-241 from individual pixels of the flight hybrid on Module A Side A.	
	The spectra are arranged in the same relative positions as the corresponding pixels	
	appear on the hybrid. The red and blue curves represent single-pixel and charge-	
	sharing events, respectively. The abscissas are in pulse height units, without correction	
	of varying charge induction amplitudes	68

- 3.18Pulse height of the 60 keV line of Am-241, measured at each pixel on the flight detectorhybrids. The left and middle columns show the histogram and map of the pulse heights, repsectively. The right column shows the amplifier gain of the read-out ASICs, which has been compensated in the pulse height maps and histograms, except for half of Module C Side A. To show the dead areas in various hybrids with higher contrast, all maps in this figure show higher values as darker pixels, and lower values as brighter pixels. Each row of figures contains data from the same flight detector. 703.19Histograms and maps of event statistics for the calibration run at $-15\,^{\circ}\mathrm{C}$ measuring an uncollimated source of Am-241. The three columns, from left to right, show the histogram of event counts at each pixel, the map of single-pixel event counts, and the map of two-pixel event counts. Each row of graphs and images are from the same detector hybrid. The histograms show one- and two-pixel events by the red and green curves, respectively; the blue curve represents the sum of both types of events at each pixel. Both maps on each row are drawn with the same greyscale, with black pixels representing zero count, and white pixels representing values at or beyond the midpoint value on the abscissa of the histogram to the left. 733.20Histograms and maps of event statistics for the calibration run at -5 °C measuring an uncollimated source of Am-241. See the caption of Figure 3.19 for an explanation of the various graphs and images. 74Histograms and maps of event statistics for the calibration run at +5 °C measuring an 3.21uncollimated source of Am-241. See the caption of Figure 3.19 for an explanation of 75Comparison of event statistics in the three calibration runs at -15, -5 and +5 °C, 3.22measuring an uncollimated source of Am-241.... 774.1(a) Map and (b) histogram of the electronic noise distribution at room temperature
- (23°C), with the voltage bias off. The mean and standard deviation of the distribution are 540 eV and 56.4 eV, respectively (108 ± 11.2 electrons). The pixels at the bottom-left corner are unconnected, and thus have lower noise amplitudes, between 0.3 and 0.4 keV FWHM (i.e., 60-80 electrons). All anode contacts on Row 1 are narrower than the ones in the remaining rows (see Section 3.6); note the lower noise level across Row 1. 80
 4.2 Leakage current (a) map and (b) histogram at room temperature (23°C) and detector bias HV = -250 V. There is substantial variation in the leakage current at room temperature, with strong spatial correlation. The circles indicate the two regions at which Am-241 spectra are obtained at 0°C. The spectra are displayed in Figs. 4.4 and 4.5.

4.3	Current-voltage relations for the two regions circled in Fig. 4.2, at two different temper-	
	atures. There is substantial difference in the two regions at room temperature, but the	
	leakage current becomes negligible at 0 °C. From these measurements, the apparent	
	bulk resistivity at room temperature is $3.4 \times 10^{10} \Omega$ -cm in the high leakage region and	
	$6.3 \times 10^{10} \Omega$ -cm in the low leakage region. At 0 °C, it is above $5 \times 10^{11} \Omega$ -cm in both	
	regions.	82
4.4	Am-241 spectrum at 0° C, obtained from the high-leakage region (at room temperature)	
	circled in Fig. 4.2. The cathode-to-anode potential is $HV = -400 V$. Data collection	
	took place for 69 minutes.	83
4.5	Am-241 spectrum at 0°C, obtained from the low-leakage region (at room temperature)	
	circled in Fig. 4.2. The cathode-to-anode potential is $HV = -400 V$. Data collection	
	took place for 58 minutes	83
4.6	Count rate (a) map and (b) histogram. The data was taken at room temperature and	
	detector bias $HV = -80$ V. This bias is high enough for the detector to be operational,	
	but just low enough for none of the pixels to saturate at room temperature. Only	
	events corresponding to the 60 keV line of Am-241, and events triggering one or two	
	pixels, are included in this figure. The mean and standard deviation of the distribution	
	are 61.8 and 19.5 counts/pixel, respectively.	85
4.7	(a) Map and (b) histogram of the count rate (displayed in Fig. 4.6) convolved with	
	a boxcar of three pixels in width in each direction. Pixels on Row 1 are narrower	
	(see Section 3.6), and the count rates are thus lower accordingly. The lower rates on	
	Row 2 and at the remaining three edges of the detector are due to the convolution	
	of a finite data set. Pixels at the lower left corner are unconnected. These edge	
	pixels and unconnected pixels are not included in the histogram; specifically, only	
	pixels within the interior rectangle marked by Rows $3-23$ and Columns $2-43$, and	
	outside the unconnected region marked by Rows $20-23$ and Columns $2-6$, are included	
	in the histogram. The included pixels are outlined in the pixel map in (a) above.	
	Mean = 60.98 counts; $1\sigma = 8.48$ counts.	86
4.8	Leakage current-count rate correlation for the HEFT detector. The count rates are	
	from Fig. 4.6 (room temperature, $HV = -80 V$), and the leakage currents are from	
	Fig. 4.2 (room temperature, $HV = -250 V$). Unconnected pixels and pixels registering	
	fewer than 10 event counts are not shown in this figure. The cause of the low count rates	
	at these pixels are likely unrelated to the leakage current, and distorts the correlation.	
	The remainder of the data points, shown in dots, are fitted with a straight line relation,	
	$counts = 0.54 \times (current/pA) + 21$, which is also shown in the figure	87

5.1	Distribution of electronic noise in a 11×10 pixel area of the gridded detector, at	
	four different bias and temperature combinations. The inset map shows the spatial	
	distribution of the same data, with brightness proportional to the noise magnitude. $% \left({{{\bf{n}}_{{\rm{m}}}}} \right)$.	90
5.2	Distribution of line widths from the 59.54 keV line of $^{241}\mathrm{Am},$ at $0^{\circ}\mathrm{C},$ as measured at	
	11×10 pixels of each of the two detector designs. The histogram on the left shows	
	a slightly better performance of the detector without a grid (solid) than the gridded	
	detector. The intensity maps on the right show the spatial distribution of the line	
	widths. Both maps are drawn with the same greyscale, with brightness proportional	
	to the line width. Note that pixels at the corners are shielded from the source by the	
	circular collimator opening. These pixels are not included in the histogram	92
5.3	Spectra of $^{241}\mathrm{Am},$ from all single-pixel events summed over an 11×10 pixel area of	
	the HEFT detectors. The Gaussian line widths as fitted from the spectra are listed.	
	The non-gridded detector in (b) measures a more symmetric $59.54 \mathrm{keV}$ line than the	
	gridded detector in (a). \ldots	93
5.4	Spectra of $^{241}\mathrm{Am},$ from all 2-pixel charge-sharing events summed over an 11×10 pixel	
	area of the HEFT detectors. For comparison, the spectra from single-pixel events are	
	scaled to the same count rate and displayed here in dotted lines. \ldots \ldots \ldots \ldots	94
5.5	Spectra of $^{241}\mathrm{Am},$ from all events summed over an 11×10 pixel area of the HEFT	
	detectors	95
5.6	Intensity maps of the FWHM (left) and skewness (right) of the $59.54\mathrm{keV}$ line of $^{241}\mathrm{Am}$	
	at various electrode bias pairs. Darkness is proportional to the FWHM magnitude	
	on the left and to the skewness magnitude on the right, so that the most desirable	
	configurations are the brightest ones in each case. There is a partial trade-off between	
	FWHM and skewness.	96
6.1	The anode-plane configuration of the CdZnTe detector. Drawn to scale. The pixels	
	studied in Section 6.3 are labelled '1' and '2'	99
6.2	Spectra of $^{241}\mathrm{Am}$ from single-pixel events, when the source is collimated at the centre	
	of a pixel (solid lines) and in between two (dotted lines). The $59.5~\mathrm{keV}$ line is both	
	widened and shifted to lower channels when the source is moved from the centre to the	
	edge of each pixel.	101
6.3	Plot of pulse-height pairs (E_1, E_2) measured at adjacent pixels in charge-sharing events	
	(the dash lines delimit the region included when calculating the energy and width at	
	the centre of the 'track'). \ldots	101
6.4	Cross-section of the $CdZnTe$ detector showing the electric field and potential near the	
	anode plane	105

6.5	Regions corresponding to single-pixel and coincidence events at the corners of four ad-	
	jacent pixels. The number of pixels triggered is labelled in each region. The rows and	
	columns of crosses indicate the starting positions of charge trajectories; 40×40 tra-	
	jectories cover this (667 $\mu {\rm m})^2$ plane. Also shown in shade are the positions of gaps	
	between the anodes and the steering electrode grid.	106
6.6	Comparison of charge loss scenarios as predicted by our model and as observed in	
	experiment. The model assumes $(\mu \tau)_{surface} = (\mu \tau)_{bulk}$. The experimental data are	
	the same as the ones shown in Fig. 6.3 ; the line marked with crosses along its length is	
	the modelling result	107
6.7	Variation in the extent of charge loss computed with $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}}$ varied between 0 and	
	1. Curves shown are for $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}} = 0, 0.05, 0.10, 0.16, 0.40$ and 1.0.	107
6.8	Modelling results with the value of $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}}$ adjusted to fit the data.	108
7.1	Mechanical drawing of the masks used in the determination of the relative positions of	
	the two hybrids on a HEFT focal plane	119
7.2	X-ray images of a radioactive point source through a slit mask, recorded by four of the	
	six flight detector-hybrids, for the determination of the relative positions between the	
	two hybrids on a HEFT focal plane.	120
7.3	$Photograph \ of \ a \ fo cal-plane \ module \ within \ a \ background \ shield \ assembly, \ with \ the \ cover$	
	of the shield removed, revealing the translucent plastic scintillator within. Photographs	
	like this one are used to measure the relative angle between each focal-plane module	
	and the axes of the telescope truss.	122
7.4	Angular offsets of the focal plane modules with respect to the axes of the telescope truss.	123
7.5	Experimental set-up in the alignment procedure.	124
7.6	Actual (left) and simulated (right) images of the X-ray beam used for aligning the	
	X-ray mirrors and detectors.	125
7.7	Flight path during the first flight of HEFT.	130
7.8	Time profiles of the payload altitude, pointing elevation, event rate and temperature	
	in flight. The observation times of the various targets and the time of passage through	
	the Pfotzer Maximum are indicated.	132
7.9	Histograms of two periodic time counters—the Sequence IDs (left) and the Time Since	
	Reset (right)—of events recorded in flight (solid line) and during ground calibration	
	(dashed line). Note the increased fraction of abnormal events in flight in both histograms.	133

- 7.10 Correlation between the two synchronous periodic time counters, Sequence ID (Figure 7.4.1, left panel) and Time Since Reset (Figure 7.9, right panel), in flight and during ground calibration. Each of these figures is a '3D histogram', where the coordinates represent the values of the two time counters (modulo their respective periods), and the darkness of each 'pixel' represents the rate of events recorded at those time values. Note that the usual one-to-one correlation between the two periodic time counters is broken in flight by spurious events caused by an unknown systematic problem. . . . 133
 7.11 Histograms of the number of pixels triggered per event in flight (solid lines) and dur-

- 7.13 Livetime of the X-ray detectors, shield event rates and detector-shield coincidence rates for the three telescope modules. The three curves in each panel show the minimum, average and maximum values, respectively, in 5-minute time bins.
- 7.15 Measurement of the point-spread function (PSF) of the HEFT telescopes through the observation of Cyg X-1 and ground calibration. Data from two of the three modules (A and B) are combined in these figures. The top and bottom panels show fits to the PSF; the middle panel shows the residuals of the fit. Flight data is in blue; the dotted line at the top is the background. The red curves show simulated PSF based on ground X-ray and mechanical measurements. They include an 8" aspect error and the off-axis mirror response. The green curves show the on-axis response with no aspect error. These calculations account for all events out to a 6'0 radius, which includes 99% of the simulated photons.
- 7.16 Spectrum of Cyg X-1, measured by Modules A and B HEFT on 2005- 5-19 for 52 minutes. 139 $\,$
- 7.17 Left: sky image of the Crab Nebula, measured by HEFT on 2005- 5-19 for 15 minutes.Right: projected position of the target on the focal plane of telescope Module A. . . 140

xxii

8.1

by solid lines in (a); the background region is also shown in (b) by dotted lines. Xray point sources detected within the *XMM-Newton* field of view are marked by solid lines in (b), with the size of each region proportional to the X-ray flux of the point source. The boundary of the source region is identical to the mean+5 σ contour in the NVSS image, excluding all overlapping X-ray point source regions. Image (a) is made by (1) scaling the filtered MOS (0.2–3.7 keV) and PN (0.2–2.3 keV) images, according to each instrument's response to the same source model and flux, (2) correcting the scaled images for exposure and combining them with the SAS routine *emosaic*, and (3) smoothing the combined image with a Gaussian function of width $\sigma = 4$ arcsec, about the size of the instrument's point spread function.

8.2 XMM-Newton (a) and NVSS (b) images of 1401-33. See Fig. 8.1 for a description of the various features. In addition to the source and background regions at the centre and top of the images, respectively, we also mark the extent of the radio emission at 330 MHz, according to fig. 1 in Subrahmanyan et al. (2003). Note that the source and background regions here are equidistant from the bright galaxy NGC 5419 at the centre of the Abell S753 cluster, about 10 arcmin to the west (right) of the source region. Image (a) combines the filtered MOS images of 0.2-6.5 keV and the filtered PN image of 0.2-4.1 keV.

144

8.3 Proton flux and ratio of oxygen charge states O^{7+}/O^{6+} in the solar wind, measured by *ACE* during the time of our observations. The two horizontal lines in each panel indicate the mean and 90-percentile levels of each quantity over the first 100 days of 2002. The two vertical lines in each panel delimit the period of the *XMM-Newton* observation. The 64 s-averaged proton flux data are from the SWEPAM instrument on *ACE*, while the hourly charge state ratios are from the SWICS/SWIMS instrument. Note that the break in the O^{7+}/O^{6+} curve at the time of our observation of 0917+75 was due to missing data. Yet, from the values before and after the break, one can infer that the O^{7+}/O^{6+} ratio during the observation is probably higher than the 90percentile. There is a ~1 hour travel time by the solar wind from the L1 point, the location of *ACE*, to the Earth. The time shown here is the time measured by *ACE*. 147

- 8.4 Spectra of 0917+75 and of the blank sky data. The solid line shows the spectrum of 0917+75 from MOS2, convolved with a 100 eV-wide boxcar. The dashed and dashdotted lines show the blank sky spectra (also convolved with a 100 eV-wide boxcar) from Nevalainen et al. (2005) and from Carter and Read (2007), respectively. Both blank sky spectra contain events extracted from the same detector coordinates of MOS2 as the source region in this study, and scaled to the spectrum of 0917+75 by equating the GTI-filtered livetimes. Data from Carter and Read (2007) is also refiltered with (FLAG == 0) for consistent comparison. Note the excess of counts at discrete 'spectral lines' at 0.5-1.0 keV when compared to the data from Nevalainen et al. (2005), and the overall deficit of counts when compared to Carter and Read (2007) (see text for an explanation). The count rates displayed here are livetime corrected. 1488.5Light curves of high-energy events (left) and of events at the periphery of the field of view (right) during the observation of 0917+75. The three rows of graphs show, from top to bottom, events from MOS1, MOS2 and PN, respectively. The abscissae show the time since the start of the observation, in seconds. The ordinates show the count rate, in counts s^{-1} . The three solid lines in each graph on the left indicate the mean and $(\text{mean}\pm 20\%)$ high-energy count rates. The long-dashed lines show the same quantities with the median. The two short-dashed lines delimit the 'normal' range of count rates reported by Nevalainen et al. (2005), which is exceeded during most of this observation. On the right, the solid (long-dashed) lines delimit the acceptable range of count rates computed with the mean (median), according to Nevalainen et al. (2005). The short-dashed lines delimit the 'normal' range observed by Nevalainen et al. (2005); only the top line is seen in the top two graphs. 150**B.1** B.2The HEFT preamplifier circuit.... 175The HEFT postamplifier circuit. B.3176B.4A capacitor connected in series at the input of an amplifier serves as a high-pass filter. 176 B.5178**B.6** Voltage 'memory' connected across the postamplifier. 179B.7The HEFT shaping amplifier.... 179The HEFT discriminator. B.8179B.9The HEFT sampling capacitor bank.... 180B.10The HEFT read-out amplifier. 180

List of Tables

1.1	List of X-ray binary systems displaying CRSFs and their properties	6
1.2	Hard X-ray astronomical instruments currently in operation.	10
1.3	Hard X-ray focusing telescopes currently under development.	15
2.1	Performance targets of HEFT.	19
2.2	Housekeeping sensors within the HEFT focal-plane pressure vessel.	31
2.3	Positional and motion sensors on HEFT.	34
3.1	Properties of various common semiconductor detector materials.	38
3.2	Production yield of the HEFT read-out ASICs on each wafer	57
3.3	List of HEFT detector hybrids fabricated up to the 2005 flight campaign	64
4.1	Summary of Some Quantitative Properties of the HEFT detector	87
5.1	Comparison of the characteristics of Hybrid 5 S/N 3 and S/N 4	89
6.1	Estimates of $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}}$ from numerical modelling.	108
7.1	Timeline of events during flight, 2005- 5-19.	127
7.2	Celestial targets observed during the HEFT campaign of 2005	129
8.1	Source and XMM observation parameters.	143
8.2	Numerical results from the X-ray observations.	145
8.3	Physical parameters inferred from observations.	154
8.4	List of relics with field limits from published IC/CMB measurements.	156

Chapter 1 Introduction

This thesis describes the development and first deployment of an astronomical instrument for the hard X-ray band (10–100 keV), called the High Energy Focusing Telescope (HEFT). The hard X-ray band is the lowest energy band where the sky is not significantly dominated by thermal emission from stars and diffuse astrophysical plasma. Thus, it is also an ideal energy band for the study of *non-thermal processes*, many of which are of high astrophysical significance. Examples of non-thermal phenomena detectable in hard X-rays include the so-called diffuse X-ray background, the up-scattering of lower-energy photons by relativistic electrons in various energetic environments, and various nuclear and relativistic processes observable in stellar remnants, including young supernova remnants, and pulsar and neutron star environments. We do not have a good understanding of many of these astrophysical phenomena; it is the goal to further study these non-thermal processes that motivates us to develop the High Energy Focusing Telescope.

1.1 The hard X-ray band: a new window for astronomy

In the following paragraphs, I describe a number of hard X-ray phenomena in more detail. The intent to understand these phenomena provides the motivation of the work described in subsequent chapters.

1.1.1 The X-ray background

On 1962- 6-19, the first X-ray astronomy experiment onboard an Aerobee rocket discovered that the X-ray sky is filled with an ambient diffuse emission that is spatially fairly homogeneous, at 1.7 photons/s/cm²/sr within 1.5–6.2 keV (2–8 Å) (Giacconi et al., 1962). This diffuse emission is known as the X-ray background (XRB). Surveys in recent years in soft X-rays (0.2–10 keV), with much improved angular resolution enabled by the focusing telescopes *Chandra* and *XMM-Newton*, have revealed that most of the X-ray background at those energies are actually previously unresolved point sources, from about 90% of the XRB at 0.5–2.0 keV, to 80–90% at 2–6 keV and 50–70% at 6–

10 keV (Brandt and Hasinger, 2005). Over 70% of these point sources are found to be active galactic nuclei (AGNs)—the nuclei of 'active galaxies', which emit a large amount of radiation across a broad range of frequencies, in comparison to normal galaxies, such as our own Milky Way. However, at higher energies, the resolved fraction of the X-ray background remains low: observations with IBIS on *INTEGRAL*, the newest hard X-ray astronomical instrument currently in orbit, resolve only 3% of the XRB into point sources at 20–50 keV (Krivonos et al., 2005). It is believed that highly obscured types of AGNs, e.g., Type-II Seyfert galaxies, contribute to the XRB at high energies. Yet, due to the lack of instrument in the hard X-ray band with high angular resolution, observational evidence remains elusive. This is significant, because the spectral density of energy flux (i.e., flux per unit frequency, νF_{ν}) for the XRB is concentrated roughly in the range 20–40 keV (Fabian and Barcons, 1992), and at these energies, the XRB is as yet only 3% resolved. In other words, we have yet to account for the complete origin of the XRB at the frequencies where its effect is the most prominent. If the conjecture that highly obscured AGNs make up the majority of the XRB at these energies, then by resolving this population of hard X-ray sources, we would advance our understanding of these interesting sources tremendously. Of course, such a feat requires that hard X-ray telescopes with high resolving power be available. Unfortunately, the current generation of astronomical instruments in hard X-ray does not yet allow us to achieve such a feat.

1.1.2 Core-collapse supernova remnants

Supernovae can be classified into two main categories—Type Ia supernovae are those formed from the collapse of carbon-nitrogen-oxygen white dwarfs, when they accrete matter beyond the Chandrasekhar mass limit of $1.4M_{\odot}$; core-collapse supernovae are those formed from stellar explosions, when stars use up their fuel for fusion and reach the end of their lives. In the scenario of a corecollapse supernova, the stellar explosion ejects the outer layers of the star into the surrounding medium, forming an expanding shell of ejecta. In contrast, the inner layers shrink inwards due to gravity and form a compact object—either a neutron star or a black hole. Whether a neutron star or a black hole is formed depends on the amount of matter that falls inwards to form the compact object; this in turn depends on the radius at which the infalling matter and ejecting matter separate. We call this critical radius the 'mass cut'.

As much as astronomers have studied core-collapse supernovae and their remnants (SNRs), we still do not know much about the mass cut, and consequently, the initial mass of neutron stars and black holes. According to present theories of core-collapse supernovae, there is an important diagnostic tool that can tell us the location of this mass cut: at the mass cut, an ample supply of alpha particles at extremely high temperatures (above 3.5×10^9 K) fuses into carbon-12 and heavier nuclei to form a sequence of isotopes (up to nickel-56) with equal and even numbers of protons and neutrons, called the alpha nuclei. A prominent member of the alpha nuclei is the



Figure 1.1: Decay sequence from 44 Ti to 44 Ca.

radioactive isotope titanium-44, with a half-life of $t_{1/2} = (58.9 \pm 0.3)$ years (Ahmad et al., 2006). It decays, by way of scandium-44 ($t_{1/2} = 3.97$ hours), to the stable isotope calcium-44. Figure 1.1 illustrates this decay sequence. Because of the relatively long lifetime of Ti-44, we should be able to detect the radioactive emission of Ti-44 in young core-collapse supernova remnants, such as the three-century old Cassiopeia A (Cas A) and the twenty-year old SNR 1987A. In fact, previous observations of Cas A by Comptel onboard the *Compton Gamma-Ray Observatory* (*CGRO*) has detected a 1.156 MeV nuclear spectral line, produced when newly created Ca-44 de-excites to its nuclear ground state (see Figure 1.1). Further observations of Cas A by the Phoswich Detection System (PDS) onboard *BeppoSAX* also detected excess line emission above a continuum in the hard X-ray band, corresponding to two spectral lines of equal strength at 67.9 and 78.4 keV of Sc-44, as it de-excites to its ground state (Vink et al., 2001). Subsequent observations with IBIS/ISGRI onboard *INTEGRAL* confirmed the *BeppoSAX* detection, and also spectrally resolved the two lines (Renaud et al., 2006). If we are able to spatially resolve these nuclear line emissions associated with the decay of Ti-44 in nearby, young, core-collapse SNRs, we would be able to tell the location of the mass-cut, and thus, shed light on the initial mass of neutron stars and black holes.

Unfortunately, none of the gamma-ray and hard X-ray instruments in operation today has sufficient angular resolution to spatially resolve Cas A and other nearby young supernova remnants: Cas A and the Crab Nebula are each about 6' across, which is much smaller than the angular resolutions of Comptel on CGRO (1°7–4°4 FWHM), PDS on BeppoSAX (1°3 FWHM), and IBIS/ISGRI on INTEGRAL (12'). While the angular resolution of MeV-range astronomical instruments are still nowhere near the size of nearby young SNRs, the hard X-ray band shows better promise, as shown by the improvement from PDS on *BeppoSAX* (in operation from 1996 to 2002) to IBIS/ISGRI on INTEGRAL (in orbit since 2002). Thus, the goal of spatially resolving Ti-44 in nearby SNRs by detecting either or both of the emission lines at 67.9 and 78.4 keV still awaits a new hard X-ray telescope with arcminute or better angular resolution.

1.1.3 Inverse Compton scattering of cosmic background photons

Many large celestial systems are populated with relativistic particles (mostly electrons, positrons and protons). These relativistic particles interact with photons from the cosmic microwave background (CMB) via inverse Compton scattering. As a result, the relativistic particles transfer energy to the CMB photons, upscattering the latter from the microwave band to the X-ray band (IC/CMB). Some of these large celestial systems, such as many clusters of galaxies and jets of active galaxies, are also filled with a net magnetic field. In these cases, the relativistic particles also release energy via interaction with the magnetic field and the emission of synchrotron radiation at radio frequencies. Oftentimes, however, critical numbers such as the number density of relativistic particles, $n_{\rm e}$, or the strength of the magnetic field, B, are unknown. Yet, these values are important quantities that tell us the energy budget of the host system. Because these systems are amongst the largest-scale systems in the universe, the understanding of these systems are important in the study of large-scale structures.

Through radio observations, one can measure the strength of synchrotron emission from these systems. However, because synchrotron emission depends on both n_e and B, one cannot untangle these quantities without making some conventional assumption, such as minimum energy or the equipartition of energy density between particles and the magnetic field. In contrast, IC/CMB emission depends only on n_e . Thus, we can deduce n_e via observations of IC/CMB in the X-ray band, without making the previously mentioned assumptions. By combining IC/CMB observations with radio observations, one can also deduce B unambiguously, as well as quantities such as the energy density of the host system (Harris and Grindlay, 1979).

Recent observations of active galaxies, relic radio lobes and clusters of galaxies have shown that minimum energy and equipartition are not always valid assumptions about these celestial systems. This shows the importance of X-ray observations of IC/CMB emission. Unfortunately, current observations of galaxy clusters and AGNs in the soft X-ray band are often dominated by their much stronger thermal emission, making the measurement of IC/CMB emission either impossible, or at best imprecise. Also, with the typically μ G-level field strength in these environments, relativistic particles that emit IC/CMB radiation at soft X-rays emit synchrotron radiation at radio frequencies below about 10 MHz—a band inaccessible from Earth due to reflection by the ionosphere. As a concrete example, Chapter 8 in this thesis describes a study of IC/CMB emission from two relic radio sources, 0917+75 and 1401-33, using the *XMM-Newton* observatory in the soft X-ray band. Thermal emission unrelated to these targets and the mismatch of observation frequencies in Xray and in radio both cause problems that result in large uncertainties in the measurement of the magnetic field strength.

To mitigate the problem of contamination by thermal emission in the soft X-ray band, and to measure the IC/CMB emission of the same particles whose synchrotron emission we observe at hundreds of MHz, one must study these systems in the hard X-ray band. Such effort is again limited by current technologies in hard X-ray astronomy. With the sizes of many clusters of galaxies on the order of a few arcminutes (comparable to the size of a young SNR mentioned above), the current generation of hard X-ray instruments are unable to spatially resolve these systems, making IC/CMB observations as yet impossible.

1.1.4 Cyclotron Resonance Scattering Features

In an X-ray binary star system, a compact object and a normal star companion are gravitationally bound to each other. Particles escape from the companion star either in a stellar wind or through Roche lobe overflow. These particles accrete onto the compact object; in the process, the charged particles accelerate and emit X-rays. In X-ray binary systems where the compact object is a highly magnetized neutron star, the strong magnetic field around the neutron star confines the charged particles to travel in helical trajectories along magnetic field lines, due to the Lorentz force. When one considers the motion of a charged particle in the direction perpendicular to the field lines, one finds that its energy is quantized in Landau levels. The energies E_n of these levels are uniformly spaced, as in a quantum harmonic oscillator:

$$E_n = \frac{\hbar\omega}{1+z} \left(n + \frac{1}{2}\right),$$

where \hbar is the reduced Planck constant, ω the (angular) cyclotron frequency, and z the gravitational redshift for the neutron star. The magnetic field strength B and the charge-to-mass ratio q/m of the particle in motion determine the cyclotron frequency, $\omega = \frac{Bq}{m}$ (in SI units), and thus the separation between successive Landau levels. Much like bound electrons within an atom, charged particles in a strong magnetic field transition between the Landau levels via photon-induced excitations and deexcitations; in the process, they preferrentially absorb and scatter photons of energies equal to the energy differences between Landau levels. Because Landau levels are uniformly spaced, the energy differences between levels are always multiples of $\hbar Bq/m/(1+z)$. For electrons $(q = e, m = m_e)$ travelling in the magnetosphere of a highly magnetized neutron star with a typical field strength of 10^{12} gauss, the energy differences are multiples of $\hbar Bq/m/(1+z) = \hbar Be/m_e/(1+z) = 12 \text{ keV}/(1+z)$,

Object	Class	Line energy of fundamental (keV)
Her X-1	LMXB	40.4
4U 1626-67	LMXB	39.3
Cen X-3	HMXB	30.4
Vela X-1	HMXB	24.4
4U 1907+09	HMXB	18.3
$4U \ 1538{-}52$	HMXB	20.66
GX 301-2	HMXB	42.4
XTE J1946+274	Transient	34.9
$4U \ 0115 + 63$	Be transient	11.6
4U 0352+309 (X Per)	Be persistent	28.6

Table 1.1: List of X-ray binary systems displaying CRSFs and their properties.

well within the hard X-ray band. The absorption and scattering of photons at these discrete energies form observable absorption lines in the hard X-ray spectrum of certain X-ray binaries. We call these lines cyclotron resonance scattering features (CRSFs).

Because the energies of CRSFs are a pure function of the magnetic field strength (for electrons), CRSFs provide the only known direct diagnostic of the magnetic field at the surface of neutron stars and pulsars, where the field is strongest, due to the convergence of field lines at the magnetic poles. As the field strength decreases with distance from the neutron star surface, so does the energies of absorption and scattering of photons by electrons. Thus, while the centroid energy of a CRSF tells us the magnetic field strength at the surface of the compact object, the energy profile of the absorption line provides further information on the structure of the magnetic field. To date, CRSFs have been discovered in 13 accreting X-ray binary systems; Table 1.1 lists these systems. In addition, a putative 73 keV emission line may exist in the spectrum of the Crab pulsar (even though it is not in a binary system), possibly also due to cyclotron emission (Ling et al., 1979, amongst others). With the current generation of X-ray astrophysical instruments, the spectral resolution is only sufficient in detecting a deficiency of counts in a spectrum of an X-ray binary, and thus the precence of a CRSF. One obtains the centroid frequency of the absorption line by fitting the line to a theoretical model line profile, with various assumptions on the structure of the field. For further progress in our understanding of the strong magnetic fields around neutron stars and pulsars, there is a need for a new hard X-ray instrument with much higher spectral resolution than the present ones. With increased spectral resolution, we would be able to resolve the line profiles of CRSFs and study the structure of these magnetic fields.

1.1.5 Non-thermal emission in the Crab Nebula

The Crab Nebula, with a strong X-ray pulsar within a young supernova remnant, is a valuable natural laboratory for the study of pulsars and their non-thermal emission. According to mag-

netohydrodynamic (MHD) models of the Crab Nebula (Pelling et al., 1987), the pulsar dissipates the energy of its rotation via the acceleration and ejection of particles. These accelerated particles become relativistic; they travel away from the pulsar in bulk motion and form a stellar (or pulsar) wind. Being a supersonic flow of plasma, the pulsar wind produces an MHD stand-off shock as it comes into contact with the ambient medium of supernova ejecta. The supersonic flow decelerates through the shock, transferring its energy to particles in the ambient medium. As a result, the postshock medium is compressed, increasing the densities of both charged particles and magnetic field lines in the postshock plasma. With the strengthened magnetic field, the high-density and high-energy charged particles emit synchrotron radiation, which we observe from the Crab Nebula from the optical band through the X-ray band.

In the Crab Nebula, synchrotron emission shows structures dependent on both the position and energy of emission. The pulsar wind is not isotropic; rather, it flows preferrentially in the equatorial plane, perpendicular to the spin axis of the pulsar. As a result, synchrotron emission from the postshock medium is also the most prominent in the equatorial plane; it appears as a torus around the central pulsar in X-ray observations. The inner perimeter of this torus marks the position of the stand-off shock, about 10" across in soft X-ray observations. In contrast, the outer extent of the toroidal emission varies with the band of observation—as the postshock distance increases, emission at high energies decreases faster than at low energies, due to energy loss by the emitting particles through synchrotron radiation. Thus, information on the extent of this emission as a function of energy helps us to characterize the pulsar wind, and to constrain current MHD models of the Crab pulsar.

In soft X-rays, the *Chandra X-ray Observatory* has performed direct imaging of the Crab Nebula with subarcsecond resolution, showing structures of the toroidal emission zone in spectacular detail. In contrast, this central part of the Crab Nebula has only been imaged indirectly in the hard X-ray band, up to 64 keV and with an angular resolution of 15" (Pelling et al., 1987). These hard X-ray images are reconstructions by deconvolving one-dimensional scans of the Crab Nebula with a non-imaging instrument. They show the toroidal emission to span about 1.5 in half-power diameter at hard X-ray energies (22–64 keV). For continuing progress in this work, first done in the 1980s, both direct imaging in the hard X-ray band and imaging at energies beyond 64 keV are much desirable.

1.2 Telescope technology in hard X-ray astronomy

As evident from the various examples detailed in the previous section, many non-thermal astrophysical processes create phenomena in the hard X-ray band with interesting spatial features at the arcminute level, and sometimes with discrete spectral features as well. The set of resolvable targets in hard X-rays is restricted by the limiting angular and spectral resolutions of the current generation of hard X-ray astrophysical instruments. Progress in this field relies on improvements in instrument technologies on several fronts—higher angular resolution, higher spectral resolution, as well as higher sensitivity (that is, higher signal-to-noise ratio and lower minimum detectable flux).

Sensitivity considerations To date, hard X-ray astrophysical instruments come in two main categories—collimators and coded aperture systems (both of which I shall further describe in later sections). Although these technologies enable viable instruments with large collecting areas and reasonable angular resolution, the detectors employed are inevitably large when compared to the collecting area. This is not simply a matter of inconvenience—in the X-ray and gamma ray bands, spurious particles in the environment of an astronomical instrument interact with matter within photon detectors to create noise events; thus, a large detector volume leads to a high noise level, reducing the signal-to-noise ratio (S/N) and thus the sensitivity of the instrument to weak signals from faint sources. Because of their large detector volumes, collimators and coded aperture systems have intrinsically high noise levels, leading to only moderate S/N values and limited capabilities to detect faint sources. For further progress, we need to look into ways of improving the sensitivity.

Without a fundamental change of technology, one can only increase the sensitivity of an observation in two ways—by extending the integration time, or by building a larger instrument with more collecting area. Both approaches increase exposure to the signal source. While the amount of signal (from a time-invariant source) increases linearly with exposure, random noise characterized by Poisson statistics scales only in proportion to the one-half power of exposure. The net result is a one-half power increase of S/N: faint targets require four times as long an observation period per halving of the source flux. To do better than this, we must turn to alternative instrument technologies, such as hard X-ray focusing devices.

Focusing is the concentration of signal from a large collecting area into a relatively small collecting device by way of mirrors or lenses. The principal merit of focusing is to increase the signal-tonoise ratio (S/N) in an observation by decoupling the signal-collecting area of an instrument—now determined by the size of the mirrors or lenses—from the dimensions of a major noise source the photon detector. Thus, we increase the amount of signal (light originating from the target pointing direction) collected by the photon detector without necessarily increasing the amount of local background noise. The result is a substantial increase in the signal-to-noise ratio. In addition, focusing enables the use of photon detectors considerably more compact than those employed in collimator and coded aperture systems. Without the need to cover a large collecting area with detectors, we are able to pursue alternative detector technologies that are not always scalable in size, but with lower noise and better spectral performance. These low-noise detectors lead to further improvement in sensitivity.

In addition to sensitivity improvement, focusing systems also make direct imaging possible, which

is much desirable when observing spatially extended sources. This enables further investigations of the various celestial targets with fine spatial and spectral structures, which I described in Section 1.1.

1.2.1 The present: non-focusing technologies

In this section, I survey the hard X-ray astrophysics instruments currently in operation to provide the context in which we developed our focusing telescopes. Currently, several astrophysical instruments with sensitivity in hard X-rays are in orbit. Table 1.2 lists the names of the missions and their properties. As stated above, all of these instruments are collimator systems and coded aperture systems.

Collimator A collimator system consists of hard X-ray counting detector(s) placed behind a mechanical structure (the collimator) that limits the field of view of the detectors to a selected part of the sky. Most of the astronomical instruments for detecting hard X-rays in the past have been collimator systems, and so are the two instruments onboard the Rossi X-ray Timing Explorer (RXTE), the Proportional Counter Array (PCA) and the High Energy X-ray Timing Experiment (HEXTE). In a collimator system, the detectors of choice are usually gas proportional counters (e.g., high-pressure xenon) for low energies, and solid scintillation counters (e.g., NaI and CsI) for high energies. These detectors are not position-sensitive—they do not provide information about the location where a detected photon enters the detector, and thus the direction of the X-ray. Therefore, these systems can be regarded as simple photon counters. To make an image of the sky using a collimator system, one must scan the target field of view with the collimator. This way, the field of view of the collimator determines the angular resolution of the resulting image, and the observation time period is split into chunks dedicated to different 'pixels' of the image. Thus, imaging with a collimator system is known as 'temporal multiplexing'. For the purpose of comparison, we note that the field of view of both PCA and HEXTE are 1°(full width at half maximum, FWHM); this value determines the angular resolution of X-ray images from RXTE. The spectral resolution of the detectors are 15–18% from 6 to 60 keV.

Coded Aperture In contrast, coded aperture systems make use of 'spatial multiplexing'. A coded aperture system consists of position-sensitive detectors placed behind a 'coded mask'—a special collimator with a complex pattern of holes and plates that allows X-rays incident from different directions to cast different patterns of shadows on the detectors. The detectors do not record a sky image directly; rather, they record the sky image *convolved* with the coded mask pattern. By deconvolving the detected pattern with the mask pattern after an observation (essentially two-dimensional inverse Fourier transforms), one obtains the sky image. Compared to collimator systems, coded aperture systems have the advantage that exposure of multiple 'pixels' of the sky is done simultaneously.

		$0 \mathrm{keV}$	$6 \mathrm{keV}$	0 keV	$0 \mathrm{keV}$	$0 \mathrm{keV}$		$6 \mathrm{keV}$	$8 \mathrm{keV}$	thout	
ergy olution		% at 6	% at	!% at 133(% at 10	% at 1		% at	% at	% throug	
V) res		15	18	0.2	00 8	13		25	22	9	
Energy range (ke		15 - 250	2-60	20 - 800	20 - 100	3-35		2 - 10	2 - 25	15 - 150	
Angular resolution		1° FWHM	1° FWHM	150' FWHM	12' FWHM	3/		3' imes 15'	40'	17'	eration.
Field of view		1° FWHM	1° FWHM	16° FWFC	8° FWFC	4°8 FWFC		$6^{\circ} \times 90^{\circ} \text{ FWHM}$	$90^{\circ} \times 90^{\circ} \text{ FWZR}$	120°	ats currently in o _f
Continuum sensitivity (nominal)		1 Crab	$0.1 \mathrm{mCrab}$	$8.8 imes 10^{-4} \ \mathrm{ph/s/cm^{-2}/MeV}$	$2.3 imes 10^{-6} \mathrm{ \ ph/s/cm^{-2}/keV}$	1.2×10^{-4} ph/s/cm ⁻² /keV	•	30 mCrab	$8 imes 10^{-9} m erg/cm^2/s$	$1 \mathrm{mCrab}$	-ray astronomical instrume
Detector technology		NaI/CsI	Xe	Ge	CdTe+CsI	Xe		Xe	Xe	CdZnTe	1.2: Hard X
Imaging technology		Collimator	Collimator	Coded aperture	Coded aperture	Coded aperture		Coded aperture	Coded aperture	Coded aperture	Table
Instrument	bservations:	HEXTE	PCA	SPI	IBIS	JEM-X	initors:	ASM	WXM	BAT	
Mission	For pointed o	RXTE	RXTE	INTEGRAL	INTEGRAL	INTEGRAL	Wide-field mo	RXTE	HETE-2	Swift	

	g
•	8
	77
	2
	άŪ
	ŏ.
	$\overline{}$
	\sim
•	Ξ.
	ĸ
	<u>-</u>
	نب
	ΈŪ
	Я.
	Ы
	2
	\circ
	ro.
	23
	Ξ.
	Π.
	z.
	Ľ
	Ξ.
	5
	÷
	Ω,
	F.
•	-
	-
	_ ت
	đ,
	Сa
•	nca
•	mica
•	omica
•	nomica
•	onomica
•	ronomica
	tronomica
	stronomica
•	astronomica
•	/ astronomica
•	astronomica
•	tay astronomica
•	-ray astronomica
	^- ray astronomica
	A-ray astronomica
	d A-ray astronomica
	rd A-ray astronomica
	ard A-ray astronomica
	tard A-ray astronomica
	Hard A-ray astronomica
	: Hard A-ray astronomica
	2: Hard A-ray astronomica.
	.Z: Hard A-ray astronomica
	1.2: Hard A-ray astronomica
	e 1.2: Hard A-ray astronomica
	le 1.2: Hard A-ray astronomica
	ble 1.2: Hard A-ray astronomica
	able 1.2: Hard A-ray astronomica
	Lable 1.2: Hard A-ray astronomica

Because monetary cost is always commensurate with observing time, spatial multiplexing is desirable, especially when observing a large field of view, such as in surveys. However, whereas the noise background at each pixel of an image made with a collimator system consists of stray light from the limited field of view due to the collimator, and of noise in the single detector, the noise background at each pixel of a coded-aperture image comes from the entire field of view of the coded mask, as well as a much larger detector volume. Therefore, although coded aperture systems are efficient in imaging large areas of sky, their background levels are high, which is undesirable for the imaging of faint sources or in small fields of view.

Despite the drawbacks mentioned above, most astronomical imaging systems in hard X-rays today are coded aperture systems. They include wide field instruments such as the All-Sky Monitor (ASM) onboard *RXTE*, the Burst Alert Telescope (BAT) onboard the *Swift* satellite, and the Wide Field X-ray Monitor (WXM) onboard the *High Energy Transient Explorer* (*HETE-2*). These instruments make use of the wide fields characteristic of coded aperture to perform efficient surveys of the sky or to monitor transient events, such as gamma-ray bursts, from a large fraction of the sky simultaneously. A mission concept, the Energetic X-ray Imaging Survey Telescope (EXIST), is also under study as a possible implementation of NASA's Black Hole Finder Probe mission in the future, to search for black holes through a survey of the hard X-ray sky using coded aperture.

Another class of coded aperture systems in use today are three instruments onboard the International Gamma-Ray Astrophysics Laboratory (INTEGRAL) mission. The two imaging instruments of INTEGRAL—the Imager on Board the INTEGRAL Satellite (IBIS) and the Joint European X-ray Monitor (JEM-X)—give it overlapping imaging capability at 15 keV–10 MeV and at 3–35 keV, respectively. The third instrument, the Spectrometer on INTEGRAL (SPI), mainly provides spectral capabilities with only limited imaging performance when comparaed to IBIS and JEM-X. IBIS has a field of view of $8^{\circ} \times 8^{\circ}$, with an angular resolution of 12'; JEM-X has a smaller field of view of 4.8 sq. deg, but a better angular resolution of 3'. The spectral resolutions of IBIS and JEM-X in the hard X-ray band are both around 10%, while SPI is optimized for higher energies.

1.2.2 The future: focusing technologies

Focusing of light is usually achieved using either of two kinds of optical devices—mirrors and lenses. Yet, the high penetrating power of X-rays, which makes X-ray astronomy so interesting, also poses problems to optics builders in the field. Because X-rays, and especially hard X-rays, penetrate through most media, it is difficult to change their paths with ordinary optical devices. To effectively focus X-rays, astronomers have made use of a special property of X-rays when designing X-ray optics: at X-ray frequencies, the indices of refraction for all materials are slightly less than unity, the index of refraction of vacuum. This means that when an X-ray is incident on a vacuum-material boundary from the vacuum side, it travels from a medium of *high* index of refraction to a medium of *low*



Figure 1.2: Cross-section of the Wolter-I mirror geometry. Grazing angles are exaggerated for illustration.

index of refraction. In this scenario, near total external reflection is possible at angles of incidence θ beyond the critical angle of incidence, θ_c , or equivalently, at grazing angles $\phi = 90^\circ - \theta$ smaller than the critical grazing angle $\phi_c = 90^\circ - \theta_c$. Note, however, that because the indices of refracton of materials at X-ray frequencies are only very slightly less than unity [on the order of $(1 - 10^{-5})$ at 10 keV], the critical grazing angles are also very small (just a few arcminutes, or less, at 10 keV). In other words, X-rays reflect only at very small grazing angles. Nevertheless, by positioning mirrors almost parallel to incident X-rays (i.e., parallel to the optical axis), one can design effective optics to focus X-rays.

The Wolter-I mirror geometry For a comprehensive survey of X-ray optics designs, I refer the reader to Joensen (1995). Here, I describe only the design that is the most relevant to X-ray astronomy to date—the Wolter-I mirror geometry (Wolter, 1952a,b; Van Speybroeck and Chase, 1972). A Wolter-I mirror configuration is cylindrically symmetric about its optical axis. It consists of the frustum of a paraboloid, followed by the frustum of a hyperboloid (see Figure 1.2). When an X-ray enters a set of Wolter-I mirrors, parallel or almost parallel to the optical axis, it first reflects off the paraboloid (the primary reflection), then off the hyperboloid (the secondary reflection), and travels further onto the focal plane. While a Wolter-I optics system focuses on-axis rays the same way as a simple paraboloid mirror, the presence of its secondary hyperboloid mirror reduces astigmatism in the focused image when rays enter off-axis. In addition, the presence of the secondary mirror significantly reduces the focal length of the system—an important issue in practice. The Wolter-I geometry also allows us to fabricate multiple mirror shells of different radii but the same focal length, and nest them together in concentrically about a common optical axis, so that rays from all mirror
shells focus onto the same focal plane to produce a single image. This nesting effectively increases the collecting area of the overall optics system, and is an important idea, given the small collecting area of a single mirror shell, due to the almost parallel orientation of the mirrors to the optical axis.

In the soft X-ray band, the first focusing telescope came into existence in 1978 as the Einstein Observatory. In subsequent years, many other focusing soft X-ray telescopes have been launched, including the European Space Agency's X-ray Observatory (EXOSAT, 1983–1986), the Broad Band X-ray Telescope (BBXRT, 1990), the Röntgen Satellite (ROSAT, 1990–1999), Japan's Advanced Satellite for Cosmology and Astrophysics (ASCA, 1993–2001), BeppoSAX (1996–2002), the Chandra X-ray Observatory (1999–present), and the X-ray Multi-Mirror Mission (XMM-Newton, 1999–present). All these soft X-ray telescopes employ Wolter-I optics (or conical approximations to it) coated with gold, iridium or similar reflective materials. Unfortunately, the reflectivity of these mirrors reduces significantly beyond about 10 keV, as the critical grazing angle, ϕ_c , is inversely proportional to the energy, $E = hc/\lambda$, of the reflecting X-ray: $\phi_c \propto \lambda \propto 1/E$ (Zombeck, 1990, Chapter 14, pp. 322–323). To extend focusing optics technologies to hard X-rays (10–100 keV), scientists have taken two approaches:

- 1. To further reduce the grazing angles of the mirrors, or
- 2. To coat the mirrors with multilayer coatings to increase the mirror reflectivity by having Xrays reflect off the multiple interfaces between layers in such a manner that the reflected rays are amplified as a result of constructive interference.

The experiment described in this thesis, the High Energy Focusing Telescope (HEFT), takes the second approach.

Focal-plane detectors While the Wolter-I optics system is an important innovation that makes X-ray focusing possible, it makes up only half of a focusing X-ray telescope. Equally important are position-sensitive detectors on the focal plane which capture the focused X-rays. To take full advantage of the concentration of light by focusing mirrors, focal-plane detectors must be efficient in capturing all the photons that are incident on the focal plane. The high angular resolution enabled by focusing mirrors is only achieved with matching small pixel sizes on the focal-plane detector. Focusing requires fine pixels and a compact focal plane; this means that detectors for focusing systems are completely different from the large-area detector systems required in the past for collimator and coded-aperture systems. This enables instrument builders to employ new technlogies—namely, compact semiconductor industry in the past two decades, state-of-the-art very large scale integration (VLSI) technologies can now be used to build low-noise electronics within application-specific integrated circuits (ASICs) on the focal plane, e.g., for detector read-out. This allows us to measure

photon energies to much higher accuracy than ever before. Thus, the union of focusing technologies and the semiconductor revolution has produced position-sensitive focal-plane detectors with high spectral resolution, for the first time in the short history of hard X-ray astronomy.

In the soft X-ray band, charge-coupled devices (CCDs) on silicon are the detector technology of choice. The most advanced soft X-ray observatories at present, *Chandra* and *XMM-Newton*, both have CCDs as their primary focal-plane detectors—namely, the *AXAF* Charged-coupled Imaging Spectrometer (ACIS) on *Chandra*, and the European Photon Imaging Camera (EPIC) on *XMM-Newton*. Yet, due to the high penetrating power of hard X-rays, low-density silicon does not make a good material for hard X-ray detectors. When a beam of 60 keV X-rays is incident on a 2 mm-thick silicon detector, only 13.88% of the X-rays is absorbed. The second most well-researched semiconductor detector material, germanium, also poses problems. Although its density is higher than that of silicon, germanium has a relatively small band gap—the difference in energy between the conduction band and valence band of electrons are in the conduction band, making germanium unsuitable as a detector at room temperature. In order to operate a germanium detector, one has to cool the detector down to low temperature with liquid nitrogen (77 K), to place most electrons in the valence band. Thus, experimentation of new detector materials is required to fully utilize the advantages of hard X-ray focusing.

An ideal material for hard X-ray detectors should have high electron density, to make it efficient in absorbing incident photons; also, its band gap should be large, so that it does not require substantial cooling to operate. It should also be stable in air (in other words, non-reactive with water vapour or oxygen), non-toxic, abundant, easy to handle and easy to pattern, and non-brittle. In recent years, either new or revived interest has arisen on a number of compound semiconductor materials that satisfy a number of these good characteristics. They include cadmium zinc telluride (CdZnTe, CZT, or $Cd_{1-x}Zn_xTe$, with x = 0.9 being the most common), cadmium telluride (CdTe) and mercury (II) iodide (HgI₂, also known as mercuric iodide, in contrast with mercury (I) or mercurous iodide, Hg₂I₂). The bulk of this thesis is a treatise on the development of cadmium zinc telluride as the focal-plane detector on the High Energy Focusing Telescope.

Focusing telescopes of hard X-rays In the past decade, a number of experiments have been underway to develop and build the first generation of focusing telescopes for the hard X-ray band. These experiments include:

- the High Energy Focusing Telescope (HEFT), our experiment,
- the International Focusing Optics Collaboration for μ Crab Sensitivity (InFOC μ S), lead by NASA's Goddard Space Flight Center (GSFC) and Nagoya University in Japan, and

Mission	\mathbf{HEFT}	${ m InFOC}\mu{ m S}$	HERO	
Mirror technology	W/Si multilayer	Pt/C multilayer	Small grazing angle	
	$\operatorname{coatings}$	$\operatorname{coatings}$	mirrors	
Detector technology	CdZnTe detector	CdZnTe detector	High-pressure	
			gas scintillator	
Performance to date:				
Continuum sensitivity	$2 \mathrm{mCrab}$ for $20 \mathrm{ks}$	5×10^{-6} for $28.8 \mathrm{ks}$	2×10^{-5} for 10.8 ks	
$(\mathrm{ph/s/cm^2/keV})$				
Collecting area (cm^2)	$108 \text{ at } 30 \mathrm{keV}$	$42 \text{ at } 30 \mathrm{keV}$	$80 \text{ for } 20-50 \mathrm{keV}$	
	$30 \mathrm{~at}~ 69 \mathrm{~keV}$		20 at 70 keV	
Field of view	12.'6	9.'6	9' at 40 keV	
			5' at 60 keV	
Angular resolution	1.'5	$2'_{2}$	0.25 (ie, $15''$)	
Energy range (keV)	20-69	20 - 40	20-70	
Energy resolution	$0.7\mathrm{keV}$ at $14\mathrm{keV}$	$2.3\mathrm{keV}$ at $22\mathrm{keV}$	$5\%~{ m for}~20{ m -}70{ m keV}$	
(FWHM)	$1.0{\rm keV}$ at $60{\rm keV}$	$4.8{\rm keV}$ at $60{\rm keV}$		

Table 1.3: Hard X-ray focusing telescopes currently under development.

• the High Energy Replicated Optics (HERO) experiment, lead by NASA's Marshall Space Flight Center (MSFC).

The capabilities of the three experiments are listed in Table 1.3. While HEFT and InFOC μ S employ multilayer mirrors and CdZnTe detectors, HERO develops mirrors with small grazing angles and matches them with high-pressure gas (96% xenon, 4% helium) scintillators. Judging simply from the names of these experiments, one can easily sense that hard X-ray mirror technology is the main driving force behind these research projects, while detector development—the topic of this thesis—is often relegated to a back seat. Notwithstanding such negative opinion, we have devoted a considerable amount of time and effort in detector development for HEFT, and photon detectors remain an integral component of a hard X-ray telescope.

While soft X-ray astronomy is arguably in its prime at present, with the two great focusing telescopes *Chandra* and *XMM-Newton* in full operation, hard X-ray astronomy with focusing technologies is still in its early years, with much space for further advances. This thesis documents just a small part of its development in the past eight years.

1.3 The rest of it all

This thesis describes the development of the HEFT hard X-ray pixel detectors, their use, and a science project that can potentially be extended by HEFT technologies. It begins in Chapter 2 with an overview of the HEFT experiment: the objectives, specifications and implementation approaches. In Chapter 3, I describe the development of focal-plane detectors for HEFT. Chapter 4 summarizes recent operations of HEFT in two balloon flight campaigns, and the results obtained. Chapter 5 describes our study of IC/CMB emission in two relic radio sources—a work extendable to other celestial objects through observations in hard X-rays. Finally, I conclude in Chapter 6 by summarizing the current status of HEFT, of experimental hard X-ray astrophysics in general, and giving an outlook into the near future.

Chapter 2

The High Energy Focusing Telescope (HEFT)

2.1 A balloon-borne experiment

The High Energy Focusing Telescope (HEFT) is an instrument carrying several focusing telescopes of hard X-rays on a balloon-borne platform. Celestial X-rays are scattered and absorbed by molecules in the atmosphere, leading to attenuation with decreasing altitude. Figure 2.1 shows the fraction of celestial hard X-rays reaching various altitudes of the atmosphere. Because celestial X-rays are completely absorbed well before reaching the ground, hard X-ray telescopes can only operate above the bulk of the atmosphere, and thus the need for the balloon platform.

Scientific ballooning is done on a scale quite different from everyday life experiences. For HEFT, the 4175-lb payload (excluding ballast) is lifted 40 km up into the sky, above about 99.7% of the atmosphere by mass, by an enormous, 39.57 million cubic-foot zero-pressure balloon made of 20 μ m-thick polyethylene sheets, specially designed for scientific ballooning. At the final altitude, helium gas within expands the balloon to its maximum capacity and to the size of a football field, subtending a circle 60° in diameter as seen from the payload. (In that part of the sky, X-rays penetrate through the two polyethylene layers on their way into the telescopes.) Excess helium leaks out of the balloon until pressure is balanced inside and outside the balloon.

Scientific ballooning is a relatively cheap way to lift a payload above the atmosphere. Twice a year, during a period called turnaround, when wind in the upper atmosphere subsides as it reverses direction, a balloon flight lasts 10–30 hours, enabling the observation of several targets per flight. This compares favourably to rockets, whose flights last for only tens to hundreds of minutes. (Although rockets go higher in altitude, we do not need that extra, but short, boost.) Compared to satellites, which would be the other alternative, balloon platforms are retrievable and much more economical, costing hundreds of thousands of dollars, instead of hundreds of millions. In the United States, scientific ballooning is well-established. The Columbia Scientific Ballooning Facility (CSBF,



Figure 2.1: Attenuation of hard X-rays at various energies by the atmosphere. The top graph shows the column density of air at various altitudes, according to the 1976 US Standard Atmosphere. The bottom graph shows the fraction of celestial X-rays remaining at various altitudes. The five curves correspond, from right to left, to X-rays at 10, 25, 40, 55 and 70 keV. Scientific balloons reach typical altitudes between 30 and 40 km.

Property	Perform	Science motivation	
	HEFT (target)	INTEGRAL/IBIS	
Angular resolution	1′ HPD	12' FWHM	Extended objects
			(SNRs, clusters of galaxies),
			resolving the XRB
Field of view	10'	8°	Extended objects
Energy resolution	$1 \mathrm{keV}$ at $68 \mathrm{keV}$	$8 \mathrm{keV}$ at $100 \mathrm{keV}$	$CRSFs, {}^{44}Ti$ lines in $SNRs$
Energy range	$20 - 100 \mathrm{keV}$	$20\mathrm{keV}{-}10\mathrm{MeV}$	$30 \mathrm{keV}$ peak of the XRB,
			$68 \text{ and } 78 \text{ keV} {}^{44}\text{Ti lines},$
			the Crab Nebula
Continuum sensitivity	2×10^{-7}	2.3×10^{-6}	Resolving the XRB
$(\mathrm{ph/s/cm^2/keV}, 3\sigma)$	(14 modules, 40 keV)	$(100{ m keV})$	
Time resolution	$1~\mu{ m s}$	$61~\mu{ m s}$	X-ray pulsars
Pointing stability	20"	_	

Table 2.1: Performance targets of HEFT.

formerly the National Scientific Ballooning Facility) organises multiple balloon campaigns each year and provides support to many balloon projects. All these factors make scientific ballooning a viable testbed for many new and emerging space technologies, including all three experiments of hard X-ray focusing telecopes (HEFT, InFOC μ S and HERO), prior to potential long-term implementation on satellites.

2.2 Performance objectives

As a successor of collimator and coded-aperture systems for hard X-ray astrophysics, HEFT is designed to meet the following performance targets: HEFT must image the hard X-ray sky with an angular resolution of 1' or better in half-power diameter. At the same time, its field of view must be at least 10' wide, yielding images of at least 10×10 pixels. In terms of spectral performance, we set the target of having a 1 keV or finer resolution (in full width at half-maximum, FWHM) in the hard X-ray band of 10–100 keV. At the same time, HEFT must be sensitive to a substantial fraction, if not all, of this energy band with reasonable efficiency. These targets are listed in Table 2.2, alongside the corresponding performance of the coded-aperture instrument IBIS on *INTEGRAL*, which is the latest space-based instrument for pointed observations in the same energy band as HEFT (see also Table 1.3 for the realised performance to date).

The performance targets enumerated above impose requirements on both the optics and detector technologies. On the optics side, achieving high angular resolution requires high-precision shaping and positioning of the mirrors; a large field of view requires proper design of the mirror configuration. To cover a broad energy band with high efficiency, we must make multilayer coatings with a proper choice of materials and coating profiles. On the detector side, the 1' requirement on angular resolution means that detector pixels must match the half-power diameter of the optics and the focal length of the telescope in size. The 10' field of view imposes a minimum spatial extent on the focal plane to be covered by active detector material. The efficiency and sensitivity requirements affect our choice of detector materials, and the requirements on energy range and spectral resolution must be achieved by proper design of the detector and read-out electronics.

2.3 Overview of the payload

Figure 2.2 is an annotated photograph of the final HEFT instrument, taken during its first launch on Wednesday, 18th May, 2005. It was hanging from a crane on the launch vehicle on the righthand side of the picture, and connected to the inflating polyethylene balloon in the background on the left, through a parachute. The X-ray imaging part of the HEFT instrument is a set of three co-aligned telescopes—mirrors and focal-plane detectors. These telescopes are placed on a roughly cylindrical telescope truss, with the mirrors at one end, and the focal-plane detectors at the opposite end, separated by a focal length of $6.0 \,\mathrm{m}$. The truss is connected to a gondola platform through an azimuth-elevation mount that enables the telescopes to be pointed in any direction in the sky. A set of motors and flywheels controls the motion of the truss with respect to the gondola, and a variety of sensors detects the direction and motion of the telescope truss as well as the location and altitude of the payload. A flight control computer (FCC) onboard provides attitude control by processing the sensor data and signalling the motors and flywheels in a negative feedback loop at a frequency of 10 Hz. The FCC also records and relays science data from the focal plane and attitude data from the motion sensors down to the ground station in real time by radio transmission. Electrical power for the payload comes from lithium batteries. Lithium batteries are preferrable due to their high capacity-to-weight ratio, and their well-regulated voltage and current, which do not decrease gradually throughout the lifetime of the batteries. Fifty battery cells onboard the payload provide at least 50 hours of electrical power to the entire instrument.

The following sections describe in further detail the three main components of the HEFT payload: the X-ray mirrors, the focal-plane system, and the pointing platform.

2.4 The X-ray mirrors

2.4.1 Geometry

Similar to the focusing soft X-ray telescopes in operation today, each HEFT optics module consists of concentrically nested mirror shells providing two-bounce grazing-angle reflectance. Yet, unlike most of the soft X-ray telescopes, the HEFT mirrors are shaped as conical frusta instead of parabolic and hyperbolic ones as in the Wolter-I mirror geometry (see Section 1.2.2). This two-cone configuration



Figure 2.2: An annotated photograph of the HEFT instrument.

21

is an approximation to the true Wolter-I geometry, as parabolae and hyperbolae are much more difficult and expensive to shape compared to conical mirrors. The trade-off is increased blurring in the focused image, but the size of this aberration is insignificant when compared to the point spread function of the HEFT mirrors, given the 1' requirement on angular resolution.

We position the primary conical mirror segments at half opening angles $\alpha = R/4f$, where f =6.000 m is the focal length of the telescope, measured from the focal plane to the plane of intersection between the primary and secondary mirrors, by convention (Van Speybroeck and Chase, 1972), and R is the distance between the mirror segment and the optical axis (which coincides with the symmetry axis of the cone) on the same intersection plane. Similarly, we position the secondary conical mirror segments at half opening angles 3α . This arrangement provides both primary and secondary reflections at grazing angle α for on-axis rays (i.e., those incident parallel to the optical axis). Due to the nature of grazing-incidence optics, a single set of primary and secondary mirrors (called a shell) does not provide much collecting area at all. In order to have a large enough collecting area, we assemble multiple concentric mirror shells at different radii with the same focal length together into mirror modules. We set the separation between consecutive mirror shells to be small enough to prevent excessive stray light from passing through (but note that it is impossible to block stray light entirely), yet large enough to provide us with the 10' field of view imposed by the performance requirements. Given these constraints, we place in each mirror module up to 70 mirror shells, each $0.3 \,\mathrm{mm}$ thick and $2 \times 20.0 \,\mathrm{cm}$ long (from the top of the primary mirror segment to the bottom of the secondary mirror segment). The distances between the shells and the optical axis, R, range from 4.0 cm to 12.0 cm. Both inside the 4 cm inner radius and in between the 21st and 22nd shells, mechanical mandrels hold the mirror module together, with the intermediate mandrel taking up space for two shells. In between mirror shells, meticulously machined graphite spacer 'spokes' maintain each shell at the correct radius without machining error accumulating from one shell to the next. At a pace of 4 shells/week, it takes several months to completely assemble a HEFT mirror module. For the first flight campaign, we have built three mirror modules—two full modules with 70 shells each, and a third module with the inner dozen shells omitted and the intermediate mandrel replaced by two extra shells (i.e., 60 shells total). With mechanical obscurations from the intermediate mandrel and graphite spacers taken into account, each mirror module provides roughly $100\,\mathrm{cm}^2$ of on-axis geometric area.

The development of mirror modules for HEFT has been a decade-long process. From substrate selection to the forming of the mirrors (slumping) and to precise positioning of the mirror shells, each procedure has involved much research and experimenting documented extensively in our reports over the years. Chonko (2006, Chapters 2–3) summarizes these efforts and includes references to our earlier reports.

2.4.2 Multilayer coating

The key characteristic that sets the HEFT mirrors apart from mirrors of existing soft X-ray focusing telescopes is the multilayer coatings on the HEFT mirrors, which extend the reflectivity range to the hard X-ray band. We coat the HEFT mirrors with alternating layers of tungsten and silicon (W/Si), with about 100 to 150 bilayers per coating (two consecutive layers of W and Si form a bilayer). Each layer interface is an abrupt change in electron density, where a fraction of incident X-rays are reflected. With proper adjustment of the thickness of the bilayers, we produce constructive interference of the reflected rays from different bilayers, as in a Bragg crystal. The result is an enhancement in the total reflectivity. To produce reasonable reflectivity for a broad energy band, the thickness of the bilayers is graded (i.e., varies from one bilayer to the next) according to an inverse power-law profile (Joensen, 1995; Mao, 2002).

W/Si multilayer provides reasonable reflectivity up to 69 keV, compared to the upper limits of around 10 keV for single-layer coatings used on soft X-ray mirrors, most notably gold and iridium. Also, it is well-tested and is mechanically stable. The main drawback for W/Si is that its reflectivity decreases to an insignificant level at 69 keV, where the K-edge of tungsten is situated. Fortunately, with W/Si coatings, we might still have a small chance in catching the 67.9 keV line of Ti-44 in young core-collapse supernova remnants, assuming that the motion of the expanding shell towards us does not blueshift the line beyond 69 keV. While other material combinations, such as platinum-carbon (Pt/C) and nickel alloys are also possible, we conclude after our investigations that W/Si is the best option. Although nickel coatings do not suffer from a sudden drop in reflectivity in the hard X-ray band of 10–100 keV (the nickel K-edge being at 8.33 keV), it takes several hundred more bilayers of a nickel-based coating to achieve the same two-bounce reflectivity as W/Si (Christensen et al., 2000; Chen et al., 2003), and producing stable coatings that thick is a big practical challenge. Besides, current sputtering technologies produce W/Si multilayer coatings with smoother interfaces than nickel-based multilayer coatings; this reduces scattering and thus further improves reflectivity. As for Pt/C, the high cost of platinum is the main problem. With our limited budget, we are unable to allocate enough funds to produce Pt/C coatings. Also, platinum deposits much slower than tungsten through magnetron sputtering, and thus poses practical challenges as well. Nevertheless, it should be mentioned that with Pt/C, we would be able to extend the reflective energy band up another 10 keV to 79 keV, the energy of the platinum K-edge. This would guarantee us the spectral coverage of at least one of the Ti-44 emission lines in supernova remnants.

Recently, we have also studied slight variations of the material pairs. We find that by using the carbides of tungsten and silicon (WC/SiC and W/SiC), we can increase the smoothness of the layer interfaces. This translates into higher reflectivity from the overall coatings (Windt et al., 2003). We arrived at this finding after we coated the mirror segments for the first three HEFT modules with W/Si, but we plan to switch from W/Si to W/SiC for future coatings.



Figure 2.3: Total effective area of the three mirror modules of HEFT. Values shown here represent the efficiency of the mirror geometry, two-bounce reflectivity of the multilayer coatings, and the various obscurations associated with the mirror geometry and mechanical structures.

For further information on our multilayer studies, I refer the readers to our reports elsewhere: Mao (2002, Chapters 3–4) describes the design of multilayer coatings for HEFT. Jensen et al. (2002, 2003) describe our magnetron sputtering facility and the coating process. Madsen et al. (2004) details the characterization of the HEFT multilayer coatings. Windt et al. (2003) and Christensen et al. (2006) report on the comparison of W/Si with alternative material pairs. Christensen et al. (2000), Chen et al. (2003) and Kruse Madsen (2007, Chapter 4) present case studies on multilayer designs for other telescope geometries.

2.4.3 Performance

Figure 2.3 shows the combined effective area of the three mirror modules of HEFT at various incident angles. These numbers show the efficiency of the mirror modules, which accounts for the mirror geometry and its associated obscurations, as well as the two-bounce reflectivity of the multilayer coatings. At soft X-ray energies, total external reflection gives the mirror modules effective area comparable to their geometric area ($\approx 300 \,\mathrm{cm}^2$). In the bulk of the hard X-ray band (20–70 keV), reflection off the multilayer coatings maintains the effective area above 10% of the geometric area. The HEFT energy band is bound at the high-energy end by the tungsten K-edge at 69 keV. Note the sizable loss of area through the two reflections, despite the multilayer mirror coating. This highlights the importance that the focal-plane detectors must be highly efficient in capturing the focused X-rays to prevent further loss of effective area and sensitivity.

The angular resolution of the HEFT mirror modules is 1.5 in half-power diameter (HPD).



Figure 2.4: Photographs of the telescope truss, viewed from the focal-plane side. The picture on the left shows the single-piece aluminium flange with integrated strongbacks that attachs to the truss and forms the focal plane.

2.5 The X-ray focal plane

Figures 2.4 and 2.5 are photographs showing the focal plane without the pressure vessel and insulation material that hide it from view in the final configuration of the payload in Figure 2.2. The X-ray detectors are housed within three shield modules, positioned exactly behind the three X-ray mirrors at the opposite end of the telescope truss. Figure 2.6 is a cross-sectional diagramme of the shield module. Each shield module is shaped like a well, with dimensions designed to match the placement of the optics. The shield modules block X-rays outside the telescopes' field of view from entering the detectors, and detect charged particles passing through the shield, which are a major source of background photons. Through these measures, the shields reduce the background event rate and increase the signal-to-noise ratio. The focal-plane flange also hosts a microprocessor system (the 'central MISC') that communicates with additional microprocessors within the shield modules (the 'module MISCs'). The central MISC serves as a branching node along the data path between the module MISCs and the FCC. It sends instructions for booting and detector operation to the module MISCs, and relays X-ray event data from the module MISCs to the FCC. In addition, it is responsible for 'housekeeping tasks' of monitoring and regulating the focal-plane environment. It operates two fans circulating air within the pressure vessel that encapsulates the focal plane, and reads sensors that monitor the pressure within the vessel, the temperature at various locations on the focal-plane flange, both within and outside the shield modules, and voltages at various points on the electrical power network. The focal-plane flange itself is fabricated with integrated strongbacks



Focal-plane/shield modules, bottom view

Kevlar pressure vessel

Figure 2.5: Photographs of the HEFT focal plane and its components. The left-hand and righthand pictures show the sides facing away from and towards the mirrors, respectively. The right-hand picture also shows one of the two kevlar domes that bolt to the flange to form the focal-plane pressure vessel.



Figure 2.6: Diagramme showing the cross-section of the background-shield module and the focalplane module within.



Figure 2.7: Photographs of the HEFT focal-plane module and a bottom view of the shield module well, where the focal-plane module resides.



Figure 2.8: Diagramme and flow chart of the detector system.

carved out of the same piece of aluminium, for rigidity. Designed to pack 14 telescope modules (the number of telescopes in the original design) in the least area, it is 1.3 m wide, and weighs 38.4 kg.

2.5.1 The focal-plane modules

Tugged within each shield module is the focal-plane module, which consists of the hard X-ray detector, the module MISC microprocessor that controls the detector, and support electronics. Figure 2.7 shows a focal-plane module alongside a bottom view of the well in the shield module where the focalplane module resides; Figure 2.8 is a schematic diagramme of the same system.

27

2.5.2 The hard X-ray detector

Central to the focal plane is the hard X-ray detector, a semiconductor detector made up of a cadmium zinc telluride (CdZnTe) detector crystal, and an amplifier and read-out circuit implemented as an application-specific integrated circuit (ASIC). The two components are flip-chip bonded together to form a 'hybrid'—a semiconductor device made up of two different materials (CdZnTe and silicon in this case). Each hybrid has an active area 12.9 mm by 23.6 mm on the focal plane, where we implement a raster array of 24 × 44 pixels of pitch 0.498 mm. With a 6-m focal length, each pixel subtends 17.''12 from the X-ray mirrors, and oversamples the point spread function of the mirrors roughly by a factor of five [HPD / (pixel size)]. Although the size of the hybrid detector is smaller than the focal-plane area that needs to be covered, the size of the CdZnTe crystal was the maximum available from the vendor, eV Products, given the specification of single-crystal, highly uniform material at the time of our first design. In additional, the size of the ASIC is also the maximum permitted by the VLSI fabrication process we use. Therefore, we have no choice but to tile each focal plane with two CdZnTe sensors and supporting electronics. This results in dead area where the two sensors butt against each other, which we seek to minimize by additional measures.

The HEFT hybrid detector is sensitive to photons within the energy range 8–200 keV (1600–40000 electrons), with 15 spectral channels corresponding to 1 keV. The spectral resolution is 0.75 keV FWHM at 14 keV and 1 keV FWHM at 60 keV. The maximum event rate before saturation is about 50 counts/s/module. The detailed implementation of the HEFT hybrid detector is further discussed in Chapter 3.

2.5.3 Digital circuitry

While the CdZnTe-ASIC hybrid detector detects incident X-rays and measures their position and energy, it requires external digital logic for its continual operation, to handle tasks such as clocking, state transitions, and signal read-out. The 'module MISC' in each focal-plane module performs these operations. It is a P24 microprocessor—a 24-bit Minimal Instruction Set Computer (MISC) designed to interpret instructions written in the Forth programming language. Like the central MISC on the focal-plane flange, we implement the module MISC on an Actel A54SX72A field programmable gate array (FPGA). The MISC runs on a 7.3728 MHz clock cycle, driven by a 14.7456 MHz oscillator chip. The number of clock periods in each sampling period is programmable, and we set it to eight. This setting gives the HEFT detectors a 1 μ s timing resolution, fine enough to sample timingsensitive celestial targets, including the many millisecond-period accreting X-ray pulsars. Three 128 kB SRAMs provide the MISC with 128k of 24-bit memory, where we store the kernel, the Forth instructions, as well as events read-out from the ASIC. The output of the ASIC read-out line is digitized by an 80 mW, 12-bit ADC12062 analogue-to-digital converter. At the end of every 1 s frame period, the module MISC pipes the digitized data out to the central MISC on an EIA-422 serial line through a level shifter, at a maximum rate of 240 kbps. The entire focal-plane module consumes about 700 mW of power.

2.5.4 Background shield assembly

On a balloon-borne platform at the upper atmosphere, much of the noise background originates locally from the interaction of cosmic rays (both primary and secondary) with the atmosphere and with passive materials on the payload itself. Because the upper atmosphere 'glows' in hard X-rays and gamma rays from these interactions, we place each focal-plane module inside the well-shaped background-shield module to limit the exposure of the detectors to the noisy ambient environment.

Each shield module is made up of an open-ended cylindrical well of passive blocking materials, surrounded by an active plastic scintillator. The passive shield consists of three layers of materials with graded atomic numbers. From the outside to the inside, they are a 0.040-inch layer of lead (Pb, Z = 82), a 1.274-inch layer of tin (Sn, Z = 50), and a 0.141-inch layer of copper (Cu, Z = 29). When a background hard X-ray photon enters the passive shield, it is blocked by the lead layer. The lead layer releases the energy absorbed in the form of characteristic lead X-rays, at 88 keV for K X-rays, 16 keV for L1 X-rays, and at lower energies. These characteristic lead X-rays are then stopped by the tin layer, which in turn releases the absorbed energy via its characteristic X-rays, at 29 keV (K), 4.5 keV (L1) and lower energies. The same process repeats at the copper layer, and produces copper K X-rays at 9.0 keV. Because this energy is below the threshold of the CdZnTe-ASIC hybrid detectors, the copper K X-rays will not affect the operation of the detector system, and thus can be safely ignored.

Because this passive lead-tin-copper shield does not have perfect efficiency, and is itself a source of background events, we surround the passive shield with an active scintillator to improve the overall efficiency in detecting noise events. The choice of a plastic scintillator (as opposed to NaI or CsI) is an economical measure, as high spectral resolution is unnecessary in this application. When a charged particle travels through the plastic scintillator, it deposits energy along its trajectory, creating scintillation light. The scintillation light reflects off the reflective inner surfaces of the scintillator housing towards the base of the shield assembly, where a photo-multiplier tube (PMT) amplifies the light and converts it into an electrical signal. This process happens in just tens of nanoseconds, with most of the delay occuring in the PMT and the associated electronics. The electrical signal thus produced serves as a veto pulse, which is read by the microprocessor to filter out 1 μ s clock periods coincident with the presence of charged particles through the shield. During operation in flight, each shield module typically produces about 600 veto pulses per second, with 6 pulses coinciding with an event at the X-ray detector on average.

2.5.5 Environmental control

As the air density in the upper atmosphere during flight is only a tiny fraction (0.3% at 40 km) of that at ground level, high-voltage electronic components on the focal plane, notably the semiconductor detector (HV_{Cathode} ≈ -500 V) and the photomultiplier tube (HV_{PMT} ≈ -1200 V), may arc (short circuit across air gaps with a gradient of electric potential) if left in such an environment, due to the reduced emissivity, and thus capacitance, of the thin air. Also, temperature in the upper atmosphere swings dramatically in a diurnal cycle, due to varied exposure to the Sun. This affects the performance of the focal-plane electronics as well, due to changing resistance with temperature. To prevent any damage to the focal-plane electronics, and to provide thermal stability, we must keep the focal-plane assembly under atmospheric pressure, shielded and insulated during flight. We do this by encapsulating the entire focal-plane assembly in an insulated pressure vessel. (Another solution is to pot all electronic components—cover all exposed conductors with a high emissivity material, usually some sort of epoxy, thus eliminating any pure air gap. Yet, due to the delicacy of the semiconductor-ASIC hybrid and lack of testing, we decided against this idea.)

The pressure vessel consists of two hemispheric domes made of kevlar, each bolted independently to the focal-plane flange, as shown in Figure 2.5. These domes are fabricated to match the dimensions of the aluminium mid-flange, with matching bolt holes on both sides and grooves around the midflange where silicone 'O-rings' seal the dome-flange interface. We coat the interior of the kevlar domes with epoxy to completely seal the kevlar fabric.

To reduce temperature variation between daytime and night-time in the electronics, we cover the entire pressure vessel with custom-made insulation. This insulation consists of 1 inch of fiberglass insulation of density 0.6 pounds per cubic foot, sandwiched between two layers of aluminized biaxially-oriented polyethylene terephthalate (boPET) polyester film, commonly known by the trade name Mylar. To protect the insulation, and for ease of handling, we wrap the insulation in brown craft paper on both sides. We also paint the exterior of the paper cover with high refractive-index titanium white paint (TiO₂, absorptivity-to-emissivity ratio $\alpha/\epsilon = 0.29$) to keep off energy from the Sun. At the three areas where the optical axes of the X-ray telescopes intersect the pressure vessel, we cut out the outer paper cover to prevent metals in the paint from absorbing the focused X-rays.

To monitor the 'housekeeping' condition within the pressure vessel, we place multiple sensors of temperature, pressure and voltage across the focal-plane flange. Table 2.2 lists the housekeeping sensors onboard HEFT, their locations on the focal-plane flange, and their ranges and resolutions.

Table 2.2: Housekeeping sensors within the HEFT focal-plane pressure vessel.

Sensor type	Location	Normal range ^{a}	Resolution
Pressure	Mid-flange	10—16 psi	$0.008\mathrm{psi}$
		-	-
Temperature	Focal-plane module A	-25 — $+35$ $^{\circ}$ C	$0.11^{\circ}\mathrm{C}$
Temperature	Focal-plane module B	-25 — $+35$ $^{\circ}$ C	$0.11^{\circ}\mathrm{C}$
Temperature	Focal-plane module C	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Shield module A: HV power supply	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Shield module B: HV power supply	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Shield module C: HV power supply	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Shield module A: housing	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Shield module B: housing	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Shield module C: housing	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	DC-to-DC power supply	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Analogue I/O circuit board	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Fan 1 and heater	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Fan 2	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Mid-flange	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Temperature	Valve inlet baffle	-25 — $+35$ ° C	$0.11^{\circ}\mathrm{C}$
Voltage	$+5\mathrm{V}$ power supply	$+4.75 - +5.27 \mathrm{V}$	$1.8\mathrm{mV}$
Voltage	$-5\mathrm{V}$ power supply	$-5.684.75 \mathrm{V}$	$5.3\mathrm{mV}$
Voltage	+12 V power supply	$+11.55 - +12.20{\rm V}$	$4.6\mathrm{mV}$
Voltage	$-12\mathrm{V}$ power supply	-12.43— -11.77 V	$3.1 \mathrm{mV}$
Voltage	$+2.5\mathrm{V}$ power supply	$+2.16 - +2.72 \mathrm{V}$	$1.1 \mathrm{mV}$
Voltage	+28 V power supply	$+26.50 - +29.65 \mathrm{V}$	$10.0\mathrm{mV}$
Voltage	Digital-to-analogue converter reference	$+4.00 - +5.00 \mathrm{V}$	$1.1~{ m mV}$
Voltage	$+4.5\mathrm{V}$ reference	$+4.00 - +5.00 \mathrm{V}$	$1.1\mathrm{mV}$

^aAlarms turn on at the GSE computer when any quantity is outside its normal range. Implemented alarms are both visual (a red background) and audio (a bell).

2.6 The pointing platform

2.6.1 The telescope truss

The telescope truss, shown in Figure 2.4, provides mechanical structure for the three co-aligned X-ray telescopes. The truss itself is a hollow frame that is roughly cylindrical in shape, and made of metal tubes along its length. At the detector end, it connects to the matching 1.3 m wide circular focal-plane flange with integrated strongbacks carved out of a single piece of aluminium, for rigidity; at the mirror end, an adjustable metal plate holds the mirrors in place and allows for fine alignment between the X-ray mirrors and detectors through the insertion of metal shims between the mirror modules and the holding plate. This hollow truss provides only structural support to the telescopes; functions such as shielding (for the rejection of background events) and insulation (for thermal stability) depend on components in the mirror and focal-plane assemblies. The cylindrical truss is 6 m in length, and 1.3 m in diameter. Six metres is the focal length of the telescopes—the distance between the mirrors and focal planes, measured from the top surface of the semiconductor detectors to the intersection plane of the primary and secondary mirrors by convention. The $1.3 \,\mathrm{m}$ disks at the two ends of the truss are large enough to accommodate 14 telescope modules—the number of telescopes in the original design, allowing HEFT to eventually reach a 3σ continuum sensitivity of $2 \times 10^{-7} \,\mathrm{ph/s/cm^2/keV}$ in a typical 20 ks observation period. Yet, due to various limitations in reality, we have fabricated only three telescopes for the first flight campaign of HEFT (see Section 3.8). We place these telescopes towards one side of the truss, reserving the remaining half of the truss for the two star trackers (optical CCD cameras and baffles) of the attitude control system. The final configuration of the payload in Figure 2.2 shows these placements.

We place the telescope truss on an azimuth-elevation mount. On the balloon-borne gondola platform that carries the HEFT instrument, the natural motions are aligned with its yaw-pitch-roll axes. Because this platform does not have any reference direction unchanged with time (analogous to the celestial poles for ground-based telescopes), the usual equatorial mount for most ground-based telescopes does not provide the same advantage in this case; in contrast, the azimuth-elevation mount is the simplest configuration to implement and to operate. The tracking of field stars (in the visible band) and active attitude control keeps the telescope truss pointed at the intended direction.

2.6.2 Attitude control

Motors and motion sensors form a negative feedback loop that controls the pointing of the X-ray telescopes. With three rotational degrees of freedom, one motor and two flywheels onboard the gondola steer the azimuth-elevation mounted telescope truss in orthogonal directions: a motor at the pivot of the truss controls the elevation of the telescopes through motions about the gondola pitch axis; a flywheel at the top of the gondola steers the entire gondola about its yaw. Figure 2.2

shows the location of these components. A third flywheel, whose axis is horizontal in the reference frame of the gondola and orthogonal to the pitch and yaw axes, dampens any roll of the gondola. The yaw and roll flywheels couple together to control the azimuth pointing of the telescope truss. A final motor links the payload to the cable connecting to the parachute and balloon above; it rotates about the yaw to release shear stress about the cable whenever necessary, as the balloon rotates naturally by an arbitrary amount during flight.

As a balloon platform moves and swings, sensors are needed to monitor its movement, and to provide feedback to the motors and flywheels that control the pointing direction of the telescopes. Table 2.3 lists the positional and motion sensors available on HEFT.

Position sensing: GPS To locate the geographic position of the payload, we rely on satellites of the Global Positioning System (GPS). GPS provides the longitude, latitude and altitude of the payload about once every five minutes, as well as the yaw, pitch and roll of the gondola, and the current coordinated universal time (UTC) about once every second. We need the geographic coordinates of the payload to convert equatorial coordinates (RA, Dec) of celestial targets into horizontal coordinates (Az, El) for the pointing directions. The altitude gives us a measure of the residual airmass above the payload, and thus the amount of attenuation. Accurate and precise timing is useful when converting sky coordinates, and for observing timing-sensitive targets, notably X-ray pulsars, to extract absolute phase information through barycentric correction of photon arrival times and comparison with standard timing data at other (mainly radio) wavelengths. Attitude information from the GPS is not fine enough for use in active real-time control; nevertheless, they provide a long-term, wide-field and absolute reference for calibrating other, higher-resolution sensors that either have a narrow field of view and thus are not always available, or detect motion and thus have references drifting with time.

Absolute but slow direction sensing: the star trackers To determine the pointing direction of the telescope truss, we have two star trackers installed on the truss to image the sky in the optical band and identify the coordinates of stars within view. These star trackers consist of high-quality optical CCD cameras behind long baffles, and function essentially as low-magnification optical telescopes, with fields of view $2-3^{\circ}$ wide, angular resolutions of about 7".5, and a spectral response similar to the optical R band. The long baffles block out stray light from outside the field of view of the cameras; we paint their interiors with a non-reflective black paint to further reduce background noise. In this setting, the cameras can operate both in daytime and night-time, with limiting magnitudes of about +8 in the day and +10 at night for a 200 ms exposure. We install the cameras and baffles on the telescope truss, with one of them on-axis (i.e., co-aligned with the X-ray telescopes) to image the X-ray target field of view whenever possible. We place the other star

m Sampling period	$\approx 1 \text{s}$ when available	ms $\approx 5 \text{ min when available}$	ms $\approx 5 \text{ min when available}$	$ms \approx 1 s$ when available	$ms \approx 1 s$ when available	xel $\approx 3 s$, variable	xel $\approx 3 s$, variable	t 0.1s	t 0.1s	t 0.1s	it 0.1s	it 0.1s	
Resolutic	$1\mu{ m s}$	50m n	70 m	9′ п	15' n	7"79/pi	7″39/pi	20" /bit	20" /bit	9",9 /bid	1''1/s/b	1''1/s/b	
Reference frame	UTC	Earth	Earth	Earth	Earth	Celestial sphere	Celestial sphere	Gondola	(relative)	(relative)	(relative)	(relative)	
Moving frame	1	Gondola	Gondola	Gondola	Gondola	Telescope truss	Telescope truss	Telescope truss	Gondola	Telescope truss	Telescope truss	Gondola	
Sensitive direction(s)		Longitude, latitude	Altitude	Y_{aW}	Pitch, roll	Elevation, azimuth	Elevation, azimuth (partial)	Elevation	Pitch, yaw, roll	Pitch	Pitch, coupled yaw-roll	Yaw, roll	
Type	Time	Position		Orientation		Orientation	Orientation	Orientation	Motion	Motion	Motion	Motion	
Sensor	GPS					On-axis star tracker	Off-axis star tracker	Pitch motor encoders	Magnetometers	Accelerometer	Gyroscope 1	Gyroscope 2	

tracker 29.4° off-axis, so that when we point the X-ray telescopes above 60° elevation, where the polyethylene balloon blocks the view of the on-axis star tracker, we can still track stars in a field offset 30° from the target.

To operate the star trackers, we equip them each with a computer in the compact PC104 format. The computer automates the operation of the camera, reads out images, identifies stars in the images, looks up a star catalogue, and run algorithms to match the stars in the images to stars in the catalogue. It records and sends the matching solutions to the FCC, where data from all attitude sensors are gathered and processed. In the current implementation, the star trackers return a solution once every three seconds or so. These solutions are too infrequent for full attitude control, but the star trackers provide high-resolution solutions that serve as absolute references for other faster but relative sensors.

Relative but fast direction sensing: motion sensors To determine attitude information in between star tracker solutions, we rely on a variety of motion sensors onboard. They include two magnetometers, two gyroscopes (each sensing motion about two orthogonal axes), a rotary encoder for the elevation motor, and an accelerometer. The magnetometers are essentially digital compasses that tell direction by sensing the magnetic field of the Earth. As shown in Figure 2.2, we place them on opposite sides of the gondola, as far away as possible from the metal frame. Nevertheless, metals on the payload still affect the magnetometers, causing noticeable change in the readings, for instance, when the telescope truss slews. Therefore, we only use the magnetometers as motion sensors by measuring only changes with time. The gyroscopes detect angular accelerations of the instrument, and we convert the accelerations into relative changes in direction. We place one of the gyroscopes on the gondola, with its orthogonal axes aligned with the yaw and roll of the gondola, and the other gyroscope on the telescope truss, with its axes aligned with the pitch axis and the optical axis of the truss (which is coupled to the yaw and roll axes). This set of four sensors thus provides redundant information on the gondola and truss motion in all three angular degrees of freedom. The gyroscopes and magnetometers function well under difference circumstances, and they both have drifting biases that reset occasionally. Together, the two sets of sensors complement each other to provide prompt attitude solutions at 10 Hz for active control in between star tracker solutions. Also, unlike the GPS and star trackers, which can go off-line (due to lack of signal for the former and excessive motion, glare or other imaging problems for the latter), the magnetometer and gyroscope readings are always available. In addition, the rotary encoder at the elevation motor provides an absolute and high-resolution measure of the angle between the truss and the gondola. We also install an accelerometer on the telescope truss to detect its relative motion, although its data were unused during the first flight of HEFT.

For further information on the implementation details of the attitude control system, I refer the readers to Gunderson et al. (2004), Kruse Madsen (2007, Chapters 2–3) and Chonko (2006, Chapter 4). Chonko (2006, Section 4.3) describes the star trackers specifically.

2.7 Summary

The design and implementation of the HEFT instrument meet its scientific objectives. The W/Simultilayer coated X-ray mirrors provide effective area up to 69 keV, covering most of the hard X-ray band. The conical-approximation mirror configuration provides 1.5 angular resolution and 10' field of view, sufficient for imaging and resolving extended celestial targets. The CdZnTe-ASIC hybrid detectors and the rest of the focal-plane system efficiently detect the focused X-rays, suppress the noise background, preserve the angular resolution and field of view enabled by the mirrors, and provide the 1 keV spectral and 1μ s timing resolutions needed for relevant science targets. The balloon-borne platform enables access to celestial X-rays from the upper atmosphere, and the pointing platform provides stability and control of the X-ray telescopes.

Chapter 3

Detector development for HEFT

3.1 Performance requirements

As with the HEFT payload as a whole, the design and implementation of focal-plane detectors for HEFT are driven by the science objectives listed in Table 2.2. The resulting technical requirements can be classified into three main categories—spatial, spectral and efficiency requirements.

3.1.1 Spatial requirements

A number of observation targets for HEFT are extended objects, including young supernova remnants (SNRs) and clusters of galaxies. To achieve the science objectives, HEFT must spatially resolve these objects. HEFT mirror technology achieves an angular resolution of about 1.5 in halfpower diameter (HPD). With a focal length of 6 m, this translates to a beam spot 1.7 mm wide on the focal plane, and the HEFT detectors must oversample this area by at least factors of a few. In other words, the pixel pitch must be on the order of 0.5 mm.

The mirrors also limit the field of view. When X-rays enter the mirrors off-axis, graze angles between the X-rays and the mirrors increase, resulting in decreased reflectivity and also decreased effective area. The area eventually also vanishes as the off-axis angle approaches the half-opening angle of the mirror shell. The geometry of the HEFT mirrors and the reflectivity of its multilayer coatings provide a field of view about 6' in full width at half maximum (FWHM) at 20 keV, and 4' at 70 keV. This means that the HEFT detector(s) must tile a circular area 10 mm in diameter on the focal plane, 6 m from the mirrors.

3.1.2 Spectral requirements

Although the residual atmosphere above the payload at balloon altitudes (about 40 000 m) cuts off celestial X-rays below about 20 keV (see Figure 2.1), extending the HEFT energy range down to around 10 keV will be beneficial if and when HEFT is one day implemented on a satellite, as current

Material	Si	${ m Ge}$	CdTe	$Cd_xZn_{1-x}Te$	HgI_2	
Atomic number(s), Z	14	32	$48,\!52$	$48,\!30,\!52$	$53,\!80$	
Mass density, ρ	2.33	5.32	6.06	5.86	6.30	$ m g/cm^2$
Bandgap, $E_{\rm gap}$	1.12	0.74	1.47	1.572	2.13	eV
Energy loss per e ⁻ -h ⁺ pair	3.61	2.98	4.43	4.64	4.22	eV
Fano factor, F	0.084 - 0.16	0.058 - 0.129	0.11	0.14	0.46	
	(77 K)	(77 K)		$(-40 ^{\circ} \mathrm{C})$		
Mobility-lifetime products:						
for electrons, $(\mu \tau)_{\rm e}$	> 1	> 1	3.3×10^{-3}	$(3-5) \times 10^{-3}$	10^{-4}	cm^2/V
for holes, $(\mu \tau)_{\rm h}$	≈ 1	> 1	2×10^{-4}	5×10^{-5}	4×10^{-5}	cm^2/V
Electric resistivity	$< 10^{4}$	50	10^{9}	3×10^{10}	10^{13}	Ω -cm
Electric permittivity, ϵ	11.7	16	11	10.9	8.8	

Table 3.1: Properties of various common semiconductor detector materials

soft X-ray telescopes such as *Chandra* and *XMM-Newton* cease to be sensitive to X-rays above this energy. On the high energy end, the objective to image the 68 and 78 keV lines of Ti-44 from young SNRs requires that HEFT be sensitive to X-rays at at least 69 keV, and ideally beyond 79 keV. Again, while challenges in mirror coatings determine the limit on the instrument as a whole, the detectors must provide sensitivity beyond 79 keV to match the mirror performance.

Measurement of discrete line features in target objects requires fine spectral resolution. In particular, to measure the blueshift and redshift of the Ti-44 lines in young SNRs and to resolve the spectral profiles of cyclotron resonance scattering features from accreting X-ray pulsars require a spectral resolution of about 1 keV in FWHM across the sensitive energy range of 10–80 keV.

3.1.3 Efficiency requirements

Because most scientifically interesting celestial sources of hard X-rays are faint, it is important for the HEFT detectors to have high collecting efficiency, especially at high energies, to make observations within a time-limited balloon flight feasible.

3.2 Choice of technology: semiconductor detectors

Given the requirements of fine spatial and spectral resolutions, semiconductor detectors are the most promising technology to meet these requirements. Compared to semiconductor detectors, scintillators are bulky and low in spatial resolution; their spectral resolutions are also at least factors of a few short of the HEFT requirement. While high-density gas chambers can meet the spatial requirement of HEFT, their efficiency is much lower than semiconductor detectors at the highest energies of the hard X-ray band. Table 3.1 lists some physical and electrical properties of the handful of semiconductor materials commonly used as hard X-ray and gamma-ray detectors.

For the general working principles of a semiconductor detector, including the photo-electric effect, charge transport and signal induction, I refer the readers to references such as Knoll (2000,

Section 2.III, Chapters 11–13, and Appendix D) and Melissinos (1966, Chapter 5). In the following, I discuss only topics specific to the development of the HEFT detectors.

3.3 Noise components in semiconductor detectors

While the pixel pitch of a (multi-pixel) semiconductor detector essentially determines the spatial resolution, the spectral resolution of a semiconductor detector depends on quite a few factors that affect either the precision or accuracy in the determination of the energy of incident photons. Good spectral resolution thus requires proper design of the detectors to minimize or eliminate the effects of these limiting factors. I briefly describe these factors in the following.

3.3.1 Counting statistics and shot noise

Measurement of photon energies by semiconductor detectors is made through the conversion of a photon energy into a quantity of charges. Discrete variation due to counting statistics, also known as shot noise, is intrinsic in energy measurements by semiconductor detectors. The Fano factor, defined as the ratio of the variance to the mean in an energy measurement,

$$F = \frac{\mathrm{VAR}(N)}{\mathrm{Mean}(N)},$$

characterizes the counting statistics for each semiconductor detector material. For pure Poisson statistics, F = 1. Yet, in semiconductor detectors, electron-hole pairs produced in the cascade of collisions following the photo-electric effect take on only a restricted range of energies. Also, the energies of all electron-hole pairs plus loss mechanisms, including the energy imparted to the crystal lattice as phonons, must add up to the incident photon energy. Thus, the energies of charges from each photon interaction are correlated, and the number of induced charges is not statistically independent. As a result, F < 1 for typical semiconductor detectors. Table 3.1 shows the Fano factors of materials common in semiconductor detectors.

3.3.2 Johnson (thermal) noise

Thermal noise is introduced into the signal through all stages of the amplifier and read-out circuitry. As noise and signal are indistinguishable from each other, both components are amplified by the same amount along the amplifier chain. Thus, thermal noise in early stages has the greatest effect on the final waveform, and it is important to keep the contribution of this component low at the input of the amplifier chain. One can model thermal noise at the input as a random voltage source in series at the input of the preamplifier (See Figure 3.1). An input-to-ground capacitor represents the capacitance of all conductors seen by the input electrode, which is the anode contact of each



Figure 3.1: A model of thermal noise and stray capacitance at the input of the preamplifier.

pixel in the case of the HEFT read-out circuit. As noise is introduced at the input, it is amplified, independent of the signal strength, in proportion to the size of the input capacitance. The reduction of thermal noise is thus achieved by reducing the input capacitance.

3.3.3 Leakage current

Because the resistivity of a semiconductor material is finite, a continuous leakage current flows between the grounded anodes and other biased electrodes in contact with the leaky semiconductor volume. Because the amplitude of noise due to leakage current is independent of the size of the event signal, reduction of the leakage current increases the overall signal-to-noise ratio.

Under typical operating conditions of HEFT, the size of the leakage current is comparable with the signal amplitude. When a hard X-ray photon is absorbed within a semiconductor material, the energy of the photon, $E_{\rm ph}$, and the bandgap of the material determine the size of the signal, $S = E_{\rm ph}/\epsilon_{\rm e-h}$, where $\epsilon_{\rm e-h}$ is the energy loss per electron-hole pair. For the absorption of a 10 keV photon in a high-bandgap material ($\epsilon_{\rm e-h} \approx 4.5 \,\mathrm{eV}$), the signal consists of about 2000 eletron-hole pairs. In comparison, within an integration period of 1 μ s, a leakage current of 100 pA/pixel generates 624 electrons at each detector anode, a sizable fraction of the 10 keV photon signal. Thus, decreasing leakage effectively reduces shot noise.

3.3.4 Charge trapping

In addition to the normal process of recombination of transport charges with residual free charges in the semiconductor volume, transport charges are also trapped along their trajectories towards collecting electrodes in a spatially inhomogeneous manner due to characteristics of the semiconductor material. In particular, the mobility of holes is an order of magnitude lower than that of electrons, especially in compound semiconductors (i.e., those made up of more than one element). Table 3.1 lists the mobility-lifetime products $\mu\tau$ for holes and electrons for common semiconductor detector materials. Within the finite time of measurement, if the transport holes have travelled little from the site of photon absorption, this appears as a deficit of induced charges at the collecting electrodes, and a surplus at the neighbouring electrodes, distorting the measured pulse height value. In common detector geometries where collecting electrodes are positioned perpendicular to the optical axis, and thus the direction of photon incidence, the main effect of hole trapping is a dependence of the measured pulse height on the depth of photon interaction. In addition, material defects in the crystal lattice also produce site-specific charge traps and inhomogeneities in the electric field that either prevent transport charges (of both polarities) from reaching the collecting electrodes altogether or significantly prolong their travel duration. The main effect of material defects is a spatially inhomogeneous response of the detector across the focal plane. The adverse effect of charge trapping can be alleviated in a number of ways, as described below.

The small pixel effect To reduce the effect of hole trapping, one can modify the pattern of charge induction within a detector by proper design of the electrodes on the detector surfaces, thus changing the boundary conditions to which the electric field must satisfy. For single-sided pixel detectors in particular, where an array of segmented electrode contacts covers one surface of the semiconductor while a monolithic electrode contact covers the opposite surface, one can force charge induction to take place near the pixels by reducing the pixel pitch relative to the perpendicular distance between the two electrode surfaces, as shown in Figure 3.2. This phenomenon is called the small pixel effect. In a detector with pixel anodes and a monolithic cathode on opposite sides, the bulk of an induced signal is created as transport electrons drifts near the anodes. This reduces the amplitude of the hole signal relative to the electron signal, thus also the deficit of induced charges due to hole trapping.

Independent of the small-pixel effect, one can also read out and compare signals induced at the collecting electrode as well as its neighbours. The relative amplitudes of these signals give us information on the depth of photon interaction for each event. To reduce material defects, one must prescreen detector materials for uniformity and electrical properties prior to fabrication. After fabrication, extensive calibration of the gain and efficiency of every pixel is still needed to compensate trapping by defects not rejected during the initial material selection.

3.4 Choice of material

At present, there are only a handful of semiconductor materials with sufficiently mature technologies to be used as X-ray detectors. These include silicon, germanium, cadmium telluride, cadmium zinc telluride ($Cd_xZn_{1-x}Te$, $0.1 \le x \le 0.3$), and mercury (II) iodide (HgI₂). Table 3.1 lists some of



Figure 3.2: Cumulative fraction of charge induced in a semiconductor detector across opposite electrode surfaces, with and without pixels on the anode plane. Distance from the cathode is normalized to the thickness of the detector between the cathode and anode planes. The presence of pixels forces charge induction to take place mostly when tranport electrons get near the pixels, thus reducing the hole signal relative to the electron signal.

their relevant physical and electrical properties. We choose to implement the HEFT detectors with cadmium zinc telluride; I describe here the various factors influencing our decision.

Silicon With great advances in integrated circuits (ICs) in the past two decades, silicon is the most well developed of these materials. Its high purity and uniformity are desirable as a detector material; in fact, detectors for the soft X-ray band are implemented as silicon charge-coupled devices (CCDs). Yet, in the hard X-ray band, the photon absorption efficiency of silicon decreases significantly with increasing energy, as the photo-electric cross-section scales approximately with the fifth power of the atomic number: $\sigma_{\rm ph} \propto Z^5$ (Melissinos, 1966). Figure 3.3 shows the fraction of photons at various hard X-ray energies absorbed by typical 2 mm-thick detector volumes made of silicon and other semiconductor materials. At 60 keV, silicon (Z = 14) absorbs only 12% of incident X-rays. This shows that materials with high atomic numbers are essential for hard X-ray detectors.

Germanium The photon absorption efficiency of germanium (Z = 32) is high enough to collect the majority of incident hard X-rays within a reasonable thickness, and it is a common material in monolithic hard X-ray and gamma ray detectors. Yet, due to its small bandgap, an unacceptably high number of charges are in the conduction band at room temperature, and germanium detectors must be cooled, typically to the liquid nitrogen temperature of 77K to be operational. While this is relatively easily achieved in ground applications, the cryogen required to operate a germanium



Figure 3.3: Fraction of hard X-ray photons absorbed by a typical 2 mm-thick detector volume made of various semiconductor materials.

detector causes many difficulties on an air- or space-borne platform, especially if the HEFT detectors are placed on an extended satellite mission in the future. These challenges include the additional mass of the cryogen (and thus substantially increased cost), complications of a dedicated cooling system, and the hard limit on mission lifetime imposed by the amount of cryogen onboard the spacecraft. In addition to the challenges of cooling, contact electrode technology for germanium is also an issue for our application: Metal contacts deposited on germanium tend to diffuse beyond their boundaries with time. Consequently, adjacent pixel contacts can potentially be shorted together. Germanium strip detector has found its use in certain position-sensitive applications, often with orthogonal electrode strips on two opposite faces of the detector crystal to provide positional information in both directions. Yet, the pitch between strips is typically larger (typically 1 mm or wider) than the dimensions required for HEFT. Besides, the large area of strip contacts increases the input capacitance to amplifier electronics. Making pixel detectors out of germanium also involves additional problems in connecting the pixel contacts to a read-out circuit (typically in a silicon IC). The large difference in the thermal expansion rates of germanium and silicon means that a germanium detector cannot be bonded directly to a silicon read-out chip with rigid bonds, and cooled from room temperature to 77 K, without stressing and breaking the bonds. Alternatively, if one fans out the pixels to the edge of the detector, the lengthy metal traces produce extra input capacitance to the amplifier circuit that degrades the spectral quality of the detector (see Section 3.3.2), nulling the advantage of germanium as a detector material in the first place. In addition, with the large number of pixels on the HEFT focal plane, a fan out is also infeasible in practice.



Figure 3.4: Count rate histogram from the flood illumination of a circular area ten pixels in diameter on our cadmium telluride detector, with Am-241.

Compound semiconductors In general, compound semiconductors are less desirable as a detector material, because their purities are lower than those of silicon and germanium. A choice of element at each lattice site also carries the problem of imperfections in the crystal structure. Both impurities and lattice defects produce inhomogeneities in the electric field, and are potential charge traps. Also, hole mobility is substantially lower in compound semiconductors than in instrinsic semiconductors, and special design and analysis are needed to address this complication. Yet, with both silicon and germanium deemed unsuitable, we must consider compound semiconductors of high atomic numbers. In fact, the serious choices are between cadmium telluride and cadmium zinc telluride, as mercury (II) iodide is reactive with moisture in the air, requires hermetic sealing, is hard to handle, and provides no significant advantage over the tellurides.

Cadmium telluride and cadmium zinc telluride (hereby CdZnTe) have similar material properties. Because zinc only makes up about 10% of CdZnTe by number, and even less by mass, the photon absorption efficiency of CdZnTe is mainly determined by the cadmium (Z = 48) and tellurium (Z = 52) atoms. In certain situations, the two materials can be used in place of each other without major modifications to the rest of the detector system. In fact, we have bonded both cadmium telluride and CdZnTe detectors to the HEFT read-out chip to investigate and compare the two materials.¹ The cadmium telluride crystals we obtained from its main vendor, AcroRAD, are highly uniform. Flood illumination of a circular area ten pixels in diameter and with 0.498 mm pixel pitch, using an Am-241 radioactive source, shows that few corrections are needed of the relative photon detection efficiency from pixel to pixel (see Figure 3.4). This is a major advantage over CdZnTe, whose detection efficiency varies significantly over the pixel length scale of 0.5 mm. Yet, cadmium telluride has a smaller bandgap than CdZnTe; given the same type of contact electrocdes

¹Our study of cadmium telluride detectors were done in collaboration with the Institute of Space and Astronautical Science (ISAS) of Japan; Oonuki et al. (2004) gives further details.

and identical experimental conditions (electrical bias and temperature), the leakage current is much higher in cadmium telluride than in CdZnTe. While HEFT detectors made with CdZnTe are operable (though with poor spectral performance) at room temperature with tens of volts of bias across a 2 mm thick detector volume, leakage current through the cadmium telluride detector with identical settings but only 4 V across 0.5 mm of cadmium telluride from cathode to anode is high enough to saturate the HEFT preamplifiers (which, admittedly, are designed to accommodate the leakage current level of CdZnTe and not of cadmium telluride). For the cadmium telluride detector to be operational, we must cool it down to 0°C, and below -20° C to get good spectral performance. Even so, the electrical bias required to deplete the detector volume is also higher for cadmium telluride. With the HEFT settings, 0.5–1.0 mm of cadmium telluride requires about 1000 V at low temperatures for good energy resolution; in contrast, we only need 400–500 V across 2 mm of CdZnTe around 0°C for the HEFT detectors.

It should be noted that the contacts between the semiconductor and metal electrodes are Ohmic in both cases here. Alternative contact technology exists for cadmium telluride to place it in a p-i-n diode configuration, with a blocking Schottky contact made with indium as a monolithic anode. However, in this case, pixel contacts must be placed on the cathode side of the crystal, and signals read out from the cathode pixels. The low mobility of holes is then a more significant issue. Besides, our HEFT read-out chip is not designed with the capability of cathode read-out. Thus, until such issues in contact technology for cadmium telluride are resolved, CdZnTe remains the most suitable material for the HEFT detectors, despite its shortcomings (most notably the relatively high density of defects and inhomogeneity).

In addition to its high non-uniformity, CdZnTe also possesses other properties requiring special attention from detector developers. It is very brittle, and special care is needed when handling the material. To achieve 1 keV spectral resolution across the hard X-ray band, moderate cooling to around 0°C is needed. To address the issue of low hole mobility, the vendor of our CdZnTe crystals, eV Products, has tuned their 'spectroscopic grade' material to have essentially no hole mobility at all. We make use of this property to deduce the depth of photon interaction, and consequently improve the signal-to-noise ratio for the HEFT detectors.

3.5 Achieving low noise

In a semiconductor detector system with the read-out circuit implemented as an application specific intergrated circuit (ASIC), the largest component of input capacitance is the parasitic capacitance seen by long leads and traces connecting the semiconductor crystal to the silicon ASIC. One can significantly reduce the amplification of thermal noise at the electronics front end by bonding the ASIC directly to the semiconductor, eliminating high-capacitance traces and cables. In this con-

figuration, called a hybrid, one positions the read-out ASIC directly underneath the semiconductor crystal, facing the pixel contacts (the anodes in the case of HEFT) and opposite to the face where X-ray enters the crystal. Circuitry for each pixel is packed into an area that matches the corresponding pixel on the semiconductor. The pixel electrode on the semiconductor is connected to the input of the preamplifier, typically an opening on the ASIC tens of micrometres wide, using a malleable and conductive material through a general technique called flip-chip bonding or hybridization. For a hybrid detector, the main parasitic capacitance seen by each pixel electrode is those of the ASIC back-plane underneath and of the electrode(s) of opposite polarity on the opposite face of the semiconductor volume. Typically, the ASIC back-plane contributes much more, as its separation from the collecting electrodes (i.e., the bond height) is on the order of micrometres, hundreds to a thousand times smaller than the cathode-anode separation (i.e., the thickness of the semiconductor), while the dielectric constant of typical semiconductor materials are only about ten times that of air and vacuum, insufficient to compensate the difference. Thus, to minimize the input capacitance of the HEFT read-out circuit, we develop techniques to maximize the bond height and minimize the area of pixel electrodes.

3.5.1 Bonding techniques

In the course of development of the HEFT detector hybrids, we have experimented with two bonding techniques:

- 1. Indium bump bonding, and
- 2. Flip-chip bonding with conductive epoxy and gold studs,

with quite different results. We eventually settled on the second technique, but I describe both of them here for the record.

3.5.1.1 Indium bump bonding

Indium bump-bonding was the first of the two methods we tried. In this process, we cover the two sides of the hybrid—the CdZnTe detector and the read-out ASIC—each with a 'shadow mask' containing holes at the intended positions of the physical connections. We then place the two components in a thermal chamber, with the masked surfaces facing an indium target. We heat the entire apparatus to evaporate the indium and to deposit it onto the hybrid components at the intended positions for the connections, through the shadow masks. This produces indium bumps $50 \,\mu\text{m}$ in diameter and $10-12 \,\mu\text{m}$ in height on each of the hybrid components. Figure 3.5 shows a photograph of the indium bumps deposited on pixel contacts of a CdZnTe detector. After the indium bumps are deposited, we properly align the detector and the ASIC, and press them together to form



Figure 3.5: Photograph of indium bumps deposited on the pixel contacts of a CdZnTe detector with the gridded anode geometry described in Section 3.5.2. The indium bumps are 50 μ m in diameter, while the pixel pitch is 498 μ m.

the hybrid. The indium bumps on the two sides fuse together upon pressure to form conductive connections from the anode plane to the input pads on the ASIC that measure $8-10\,\mu\text{m}$ in height. Finally, we apply drops of a silicone adhesive (room temperature vulcanization RTV-162 by General Electric) at the four corners of the hybrid for added mechanical strength.

At the end of 3.5 years of bump-bonding experiments, mostly with 8-by-8 pixel or smaller prototype hybrids, we produced three full-size hybrids via indium bump-bonding. Unfortunately, our track record is not satisfactory at all. One hybrid came apart when we applied excessive heating to hasten the RTV to dry. Although we completed the production steps for two other hybrids successfully, subsequent electrical tests revealed that the indium bonds were not connecting the two sides of the hybrid properly. Only a handful of pixels in the first hybrid were connected; the result for the other hybrid was better—a third of all pixels were connected—but still far behind the expected result of close to 100% connectedness. Although indium bump-bonding is a proven practice that the consumer electronics industry employs routinely to produce, for instance, plasma television sets, they do so on a mass-production scale, with millions of dollars of investment into research and development. With our limited financial budget and manpower, we simply do not have the expertise or time to master the same skill. [Entry alone into the Microdevices Laboratory (MDL) at the Jet Propulsion Laboratory, where we held our indium bump-bonding trials, costs USD 100 000 per annum.] Therefore, we had no choice but to abandon our indium bump-bonding trials and to look for alternative hybridization methods.

3.5.1.2 Flip-chip bonding with conductive epoxy and gold stud bumps

A novel idea that influenced our second hybridization method originated from Professor Tadayuki Takahashi's group at the Institute of Space and Astronautical Science (ISAS) in Japan (Takahashi et al., 2001). In addition to indium bumps connecting the two sides of a hybrid, they first construct gold 'studs' measuring tens of micrometres tall out of the input pads on an ASIC. Then, they deposit indium bumps onto the top of these studs, and onto a detector correspondingly, and bump-bond the detector to the ASIC through the indium-stud connection. The gold studs are essentially the clipped ends of wirebonds; they are sturdy enough to serve as pillars to uphold the detector, which sits on top of the ASIC in the final configuration of the hybrid. The addition of the studs effectively increases the separation between the anode plane and the ASIC back-plane, thus reducing the capacitance of the latter, as seen by each anode contact. However, the technical challenge of successful indium bump-bonding still eludes us, and we must seek an alternative method.

The second hybridization method we have studied makes the connections using conductive epoxy. By that time, in 2002, CdZnTe detector companies had experience using conductive epoxy to glue CdZnTe detectors directly onto printed circuit boards. Working with Polymer Assembly Technology (PAT, formerly the Polymer Flip-Chip Corp, PFCC), an outside vendor specialized in applying conductive epoxy to CdZnTe detectors, we developed a modified hybridization process. In this process, we first populate the ASIC side with gold studs, just like the process developed by ISAS. Then, instead of indium bump-bonding, we screen the detector side with a viscous conductive epoxy, and flip the detector onto the stud-populated ASIC, align the two components together, so that the epoxy makes a physical connection at the tip of the studs. We then leave the hybrid to cure, with the detector on top, so that the conductive epoxy wicks down the surface of the stud further to strengthen the physical connection. Finally, we apply RTV-162 silicone adhesive to the four corners of the hybrid for additional mechanical strength.

This flip-chip bonding method with gold studs and conductive epoxy proves to be successful. For the first hybrid produced this way, 996 pixels out of a total of $44 \times 24 = 1056$ were connected. We have since employed this method to hybridize all our detectors, including all detectors fabricated for the first HEFT campaign. They number at 17 as of the end of 2006.

Figure 3.6 shows a photograph of the stencil-printed conductive epoxy bumps on a CdZnTe detector prior to flip-chip bonding with a stud-populated ASIC, and a photograph of the cross-section of the resulting gold stud-conductive epoxy bond. From here, we measure the height of the overall connection to be $45-50\,\mu\text{m}$, a five-fold improvement from indium bump-bonding. In other words, the input capacitance contributed by the ASIC back-plane is reduced by the same factor.

3.5.2 Anode geometry

In addition to increasing the bond height, we have also experimented with decreasing the area of the pixel anodes. In the early development of HEFT, we designed and tested pixel detectors with anode contacts of various sizes, ranging from 20% to 90% of the pixel pitch (500 μ m). Although detectors with small anodes have low input capacitance, we found that when an X-ray event occurs at a position below which there is no metal contact on the anode plane, the induced charge is smaller than expected, indicating charge loss at the gaps between anodes and incomplete charge


Figure 3.6: Left: stencil-printed conductive epoxy bumps on the $498 \,\mu$ m-pitch pixel contacts of a CdZnTe detector, prior to flip-chip bonding. The bumps are intentionally offset from the centre of the square contacts to accommodate the row of narrow pixels adjacent to the mating edge (see Section 3.6). Right: cross-section of a gold stud-conductive epoxy bond, made by slicing a mechanical sample hybrid vertically apart and imaging with a scanning electron microscope (SEM).

collection (Bolotnikov et al., 1999b). In a further attempt to keep the anode areas small, we added 'steering electrodes' to the anode plane. Figure 3.7 shows a diagramme of the anode plane with and without these steering electrodes. We placed these steering electrodes in between all anode contacts, on the previously bare CdZnTe surface. They form a rectangular grid delineating the boundaries of the pixels. We set them at a slightly negative potential with respect to the anodes, to direct the field lines at the gaps towards the anodes. This measure reduces charge loss at the bare CdZnTe surfaces. However, we eventually have to abandon this idea, due to practical limitations with the detector materials. The critical issue is excessive leakage current brought about by the potential difference between the steering electrode and the anodes required to mitigate the charge loss at the bare CdZnTe gaps. Although the surface resistivity of CdZnTe crystals we obtained in early years are low enough for steering electrodes to be a practical idea, we found that material properties have changed over the years, and that the crystals we obtained since circa 2001 no longer have the low-enough surface resistivity for us to operate the steering electrodes without saturating the readout electronics with excessive input current. To add to this practical barrier, our CdZnTe vendor, eV Products, reported that the narrow width of our steering eletrodes (14 μ m wide) was causing frequent breakages during the lithography that patterns our anode plane, significantly reducing the yield, and thus increasing both cost and lead-time. As a result, we have no choice but to abandon the steering electrode idea. While we do not adopt this design in the final HEFT detectors, we have done a considerable amount of work on this gridded design, which produced results that significantly influenced our final design decisions. I refer the readers to two papers, Chen et al. (2002) and Chen et al. (2004b) for the details of this work.

To minimize charge loss at the CdZnTe gaps without the steering electrodes, we minimize the



Figure 3.7: Anode plane patterns. The diagramme on the left is a mechanical drawing, to-scale, of the final anode plane, with a raster array of 44×24 pixels per detector. The top circle shows an experimental pattern with steering electrodes described in Section 3.5.2; the bottom circle shows the pattern on the final HEFT detectors. Note that the guard ring and the first row of pixels from the mating edge in both designs are contracted for two detectors to be placed side by side with minimal dead area in between.

area of the gaps to 30 μ m. (This number is determined by the minimum feature size that eV Product felt comfortable with their lithography process.) In other words, we have to increase the size of the anode contacts. This has the adverse side effect of increasing the input capacitance from the ASIC back-plane, the primary component of the source of noise for the HEFT detectors. It was exactly for the purpose of countering this capacitance increase that we developed the flip-chip bonding technique with gold studs and conductive epoxy mentioned in Section 3.5.1.2 to increase the separation between the anode plane and the ASIC back-plane. When both changes were in place, we increased the anode contact size by $(468/384)^2 = 149\%$, and the anode-ASIC separation by a factor of 5. The net change in input capacitance from the ASIC back-plane is then a 70% decrease, which was still a substantial improvement over the previous prototype.

3.6 The final detector geometry

After the various experiments with detector design described above, we settled on a final detector geometry. This is the geometry of the HEFT detectors that I discuss in the remainder of this chapter.

We populate each focal plane of HEFT with two identical detector hybrids, tiled side by side with little spacing in between. The HEFT detector hybrid consists of a cadmium zinc telluride detector crystal and a matching custom amplifier and read-out ASIC. We fuse the two components together by flip-chip bonding, using conductive epoxy on the CdZnTe side and gold stud bumps on

50

the ASIC side. We mount both hybrids on a single printed circuit board (PCB) with Staystik 472 thermoplastic, and connect signal terminals on the ASIC to traces on the PCB through 1 mil wide wirebonds (made of gold or aluminium). Figure 2.8 shows a schematic of this big picture.

Each CdZnTe crystal is 2 mm thick between the opposite cathode and anode faces. Both faces measure 12.9 mm by 23.6 mm (plus tolerance). The cathode, through which focused X-rays enter the detector, is a monolithic platinum contact. The anode plane facing the ASIC is segmented into a raster array of 44 by 24 square pixels with 498 μ m pitch and centred, 468 μ m wide, square anode contacts, also made of platinum. A contiguous platinum guard ring surrounds the entire pixel array; it is 99 μ m wide at the 'critical edge' facing the adjacent detector on the same focal plane, and at least 1 mm wide on the other three edges. To retain the uniform pixel spacing across the two detectors on each focal plane, the first row of pixels from the critical edge is reduced in size (in one direction only) so that the contacts measure only 300 μ m. The bottom half of Figure 3.7 shows this anode plane geometry.

The HEFT read-out ASIC contains 48 by 24 pixels of $498 \,\mu$ m pitch matching those on the CdZnTe crystal. Due to the smaller width of the CdZnTe crystals (a limit imposed by the vendor), pixels in Columns 0–3 of the ASIC are unconnected to any semiconductor material, while pixels in Columns 4–47 each connects to a pixel on the detector. (For consistency, we label pixel columns on the CdZnTe from 4 to 47, but in the reverse direction to match the column numbering scheme on the ASIC.) Besides the input to the preamplifier in each pixel, all input and output pads (for digital control, analogue voltage bias, and signal read-out) reside at the edge adjacent to pixel column 47.

3.7 Processing signals from the HEFT detectors

3.7.1 Basic concepts

The design of the ASIC circuitry is the key to achieving low power, low noise and good spectral resolution. In order to achieve low power, we have chosen a design that is different from the conventional amplifier chain. In our design, the signal shaping and peak detection stages of the conventional chain are replaced by a bank of 16 switch capacitors arranged to continuously capture successive samples of the preamplifier output. The result is a large reduction in power dissipation—from $250 \,\mu\text{W}$ to $50 \,\mu\text{W}$ per pixel—while allowing off-chip digital signal processing to extract near optimal energy resolution.

3.7.2 Implementation of hardware

The implementation of the sample and store mechanism with a bank of 16 capacitors is illustrated in Figure 3.8. The preamplifier consists of two stages—a charge-sensitive amplifier with a 40 pF



Figure 3.8: Schematic diagramme of the HEFT ASIC read-out chip. Subcircuit A is the shaping and peak detection stages in a conventional amplifier chain; although this subcircuit exists in the HEFT ASIC, only a simple version is implemented for triggering. Instead, accurate pulse height information is captured in the bank of 16 switch capacitors in Subcircuit B, from which the pulse height is determined. Each pixel in the ASIC contains a copy of the circuit shown in this figure; all pixels share the same serial read-out line for reading out the bank of switch capacitors.

capacitor (i.e., a gain of $4 \mu V$ per electron), followed by a second amplification stage with a voltage gain of 6. The preamplifier output is converted to a current and is integrated by the switch capacitors, cyclically one by one, with a 1μ s integration time. This process gives us a record of the current level during the previous $15-16 \mu$ s at any given time. The preamplifier output is also fed into a simple shaping amplifier with 300 ns shaping time, and then to a simple peak detect circuit that generates event triggers. When a trigger is detected, sampling of the preamplifier output continues for 8 more samples, after which the circuit freezes while the voltage levels at the switch capacitors in the triggered pixel are read out of the ASIC chip. The off-chip ADC then digitizes these 16 preamplifier output samples, and send them to the microprocessor. The pulse height of the event is recovered from these 16 samples and related state information at a later time in software. Figure 3.9 shows the continuous waveform sampled by the sampling capacitors in the HEFT read-out circuitry.

In addition to reading out the 16 samples from each triggered pixel, we also read out samples from a collection of other pixels for additional information to help with the pulse height recovery process. These pixels include all the ones neighbouring any triggering pixel (i.e., those sharing an edge or a corner with a triggering pixel). Samples at these pixels contain any systematic noise that is common in the vicinity of the triggering pixel. These pixels may also have collected a small fraction of the charge from the X-ray event, if the event has occurred near the edge of the triggering pixel, and this charge may be too small to have triggered the neighbouring pixel. With these additional samples, we implement a second discriminator in software with a much lower threshold, and also remove systematic noise from the triggered pixels.

With this scheme, an event triggering a single pixel involves reading out 16 samples from $3 \times 3 =$ 9 pixels ($9 \times 16 = 144$ 12-bit numbers). An event where two adjacent pixels trigger (which we term



Figure 3.9: Waveform sampled by sampling capacitors in the HEFT read-out circuitry. We recorded this waveform in an experiment where we placed an americium-241 radioactive source in front of a HEFT detector. The detected monoenergetic gamma rays at 59.54 keV, roughly 30 keV (missing the energy of the K-shell escape photons), and below 20 keV (the neptunium lines) each produces a discrete waveform whose amplitude records the energy of the detected photon. Multiple events, each triggering at a random time offset with respect to the system clock, provide samples that fill the gaps between discrete clock ticks. The dip immediately after the rising edge is caused by the undesirable discharge of a subcircuit component.

a charge-sharing event) requires reading out 16 samples from $4 \times 3 = 12$ pixels ($12 \times 16 = 192$ 12-bit numbers). With additional information (such as pixel coordinates, time information, etc.), each event produces about 0.5 kilobyte of information. The read out process takes about 20 ms. As implemented for HEFT, with one ADC for two hybrid sensors, the focal plane can read out rates of up to 50 counts/s before saturating.

Appendix B gives further details on the amplifier and read-out ASIC.

3.7.3 Systematics and noise components and their corrections

With much signal processing delegated off-chip, software processing becomes an important procedure that determines the quality of pulse height recovery. The software processing sequence starts with a preliminary screening of the event record, which screens out abnormal events. Most of these abnormal events have triggers at a large number of pixels, or at two or more pixels that are spatially disjoint (i.e., they do not share a common edge). Noise is the most common cause for these events. Within a single 1 μ s integration period, the rate of having two photon events occuring is also low. Also, according to our previous detector modelling results (Chen et al., 2002), we expect no more than 3% of all events to have charge shared amongst three or more pixels in the absence of material defects. When material defects are present, however, the defects trap transport electrons and prevent them from reaching the collecting anode. This leads to residual induced charges at the neighbouring pixels, and consequently, a large number of (contiguous) triggering pixels. Thus, we make an exception in this case and allow for a higher number of pixels triggered per event in regions known to contain defects. Other events that are screened out by this process includes the occasional (a few times in a million) occurrence of bad encodings, and events occuring too soon after a circuit reset, when the quiescent signal level has not stabilized yet.

The next stage of processing involves the removal of systematic noise common to all pixels. We obtain the noise level as the average of samples taken at pixel(s) not sharing an edge with any triggered pixels. Most commonly, these are pixels at the corners of the read-out area. For each of the 16 integration periods in an event record, the average noise level is subtracted from the samples measured at the triggering and neighbouring pixels. We then calculate a pulse height from each pixel that has a noise-corrected sample sequence, as the difference between the average of the last six samples and the average of the first six samples in the sequence. The pulse heights are then passed through a software discriminator, in search of charges deposited in any pixel below the threshold of the hardware discriminator. Events with false triggers are also eliminated at this stage. Because the common noise has been removed, we are now able to detect (and reject) signals below 1 keV at any of the pixels read out. While the hardware threshold requires that more than 8 keV of charges is deposited at a single pixel for an event to be detected, our software discriminator guarantees that residual charges at other pixels are also recovered down to less than 1 keV.

For the single-pixel and charge-sharing events that remain, we recalculate their pulse heights with a more sophisticated formula that takes the values of all 16 noise-corrected samples into account. Finally, the pulse heights are shifted and scaled to compensate for differences in amplifier gains and capacitor offsets across pixels. The pulse height pairs in charge-sharing events are summed, and all events are binned to produce spectra and other related information, which we present in Section 3.9.

Appendix C gives further details on the data processing chain.

3.8 Fabrication of the focal planes

After almost a decade of studies of detector-related issues and the development of the HEFT hybrid design, the HEFT focal-plane team gradually transition into a production stage starting in 2002. Figure 3.10 shows in a flowchart the production steps that lead to a complete focal plane. The production of focal-planes for flight began with the procurement of the various hardware components. Amongst these, we expected the spectroscopy-grade CdZnTe crystals and the read-out ASICs to have the longest lead-times. At the end, each of these components posed their own challenges.

3.8.1 ASIC production and screening

We produced the read-out ASICs at the AMI Semiconductor foundry through the C5N process (CMOS $0.5 \,\mu\text{m}$ N-well). With each of our ASICs measuring $26.20690 \,\text{mm}$ by $13.24390 \,\text{mm}$ (plus tolerance), we fit 63 chips on each 8-inch wafer. In total, we purchased seven wafers containing 441 ASICs; AMI produced them in two wafer runs. Although we were not aware of it at the beginning, the HEFT ASICs are extremely large when compared to most ICs in production, which typically measure only a few millimetres in length. The unfortunate result is that our production yield is significantly lower than that in typical VLSI processes, as the yield is affected mainly by material defects, the number of which is directly proportional to the area of each IC. When the first batch of 126 diced ASICs, from two wafers, arrived in our laboratory in August, 2002, we found their yield to be only 12%. Because it is practically infeasible to mount every single ASIC to a printed circuit board (PCB) and wirebond all 72 external connections for testing, we tested the ASICs with a multi-stage screening process, with most of them involving probing only a small number of the end terminals manually using micromanipulator probes. The first test was the measurement of the resistance across the power (d5V) and ground terminals of the digital circuitry. About 50% of the batch failed due to a short circuit between these two terminals. To further assess the functionality of the remaining ASICs, at first, we had to mount each of them to a focal plane PCB with removable thermal plastic tapes, and wirebonded the connections between the ASIC and the PCB. For every working ASIC, we had to carefully remove the wirebonds by hand, one by one under a microscope, and remove the working ASIC from the PCB. As our testing procedure matured, we purchased an





		V										
Wafer ID	Numb	er of dies	Number of dies with									
	total	passing DC test ^{a}	functional	0	1	2	3	bad rows				
1	63	$\geq 17 (??\%)$	23~(18%)	19	2	2	0					
2	63	$\geq 15 \ (??\%)$										
А	63	?? (??%)	$3\ (\ 5\%)$	1	2	0	0					
В	63	?? (??%)	8~(13%)	4	1	1	2					
\mathbf{C}	63	19~(30%)	6(9%)	2	3	1	0					
D	63	26~(41%)	7 (11%)	1	3	2	1					
\mathbf{E}	63	22~(35%)	8~(13%)	1	4	2	1					
All wafers	441	?? (??%)	55~(12%)	28	15	8	4					

Table 3.2: Production yield of the HEFT read-out ASICs on each wafer.

^aStatistics are incomplete due to poor record keeping at the early stage of development.

expensive probe station and a probe card that allowed us to apply probing needles to the ASICs without the need of first mounting and wirebonding to a PCB. This accelerated our ASIC testing work tremendously. After we tested all 441 ASICs from the seven wafers, we found 55 of them to be functioning, including 27 with one or more row(s) or column(s) of dead or malfunctioning pixels. Table 3.2 summarizes the yield statistics for each wafer.

3.8.2 CdZnTe detector production

The production of CdZnTe detectors for flight turned out to be equally time-consuming, although the challenges were quite different. In terms of time requirement, we found that the actual lead time for the high-quality crystals was substantially longer than the 6-8 weeks that the vendor, eV Products, had promised in their quotations. We placed an order of eight detectors for flight on 2003- 8- 6, with the material specification enumerated in Appendix A of this report. eV Products shipped the first batch of four detectors to our laboratory on 2003-10-27. These four detectors all passed our time-consuming leakage current test, but one of them marginally failed our visual inspection, as the thin guard ring at the mating edge was broken at one place, although the entire metal ring was still contiguous. Unaware of the even longer lead-time that would await us, we were perhaps unnecessarily stringent on the acceptance criteria. We rejected the detector that failed the visual inspection, only to wait for another four months for the second batch of five detectors (the replacement part included) to be shipped to us on 2004-2-27. To add to the complication, the second batch of detectors fared worse than the first batch in the leakage current test. Only one amongst the five have low-enough leakage currents across the entire chip, while two had high leakage at regions small in extent, but close to where the centre of the focal plane would be. Although eV Products sent us three more replacement parts on 2004-5-11, these parts did not meet our required level of leakage current either. At the same time of our wait for detectors to arrive, the HEFT optics team had been successful in producing three mirror modules in-house. Ultimately, to gather enough detectors for



Figure 3.11: Wiring diagramme of the apparatus for leakage measuremnts on the HEFT CdZnTe detectors.

three matching focal planes, we had no choice but to make use of a mechanical sample of unknown spectral performance but acceptable leakage current, and an old, rejected detector with the gridded anode pattern (see Section 3.5.2), but with the steering electrodes turned off (i.e., set at the same potential as the anodes) to avoid excessive surface leakage current, the problem that invalidated this detector at the beginning. By the time we had to make these desperate decisions, it was June, 2004, two months prior to the start of the impending flight campaign, while the fabrication lead-time towards the completion of a functioning focal plane was expected to be almost two months. At the end, we only completed the third and final focal plane one week before our much delayed departure date, leaving little time for laboratory testing and characterization before it was time to leave for Fort Sumner to start the field campaign in the last week of August, 2004.

3.8.3 Detector leakage measurement

Due to the stringent limit on leakage current that the HEFT ASIC can tolerate, we perform extensive measurements of leakage current for each CdZnTe detector that we receive. Figure 3.11 shows a diagramme of our leakage measuring apparatus. First, we apply the expected voltage bias during operation (-300 to -500 V) across the cathode and the anodes, and measure the pA-scale current from a number of pixels individually using a micromanipulator probe connecting to a picoamplifier. In this process, all four pixels adjacent to the one being measured are grounded, also with micromanipulators. Figure 3.12 shows a photograph, taken through a microscope, of the micromanipulator probe tips pressing on the anode contacts of a CdZnTe detector. Next, we set all anodes at ground (by pressing them all against a common, grounded conductive rubber pad), apply the same voltage bias, and measure the nA-scale current out of the cathode. This represents the overall bulk current between the cathode and anode planes. Because the current typically decays with time after we apply the voltage bias across the anodes and cathode, each measurement involves at least a few minutes of data accumulation, as we are interested in measuring the stable current during long



Figure 3.12: Photograph through a microscope of micromanipulator probe tips connecting the anode contacts of a CdZnTe detector to a picoammeter and ground connections.

operations. Figure 3.13 shows the decay of leakage current at a pixel over the course of two hours. Because the current decay lasts for at least two hours, and because we only have a primitive test set-up with manually controlled micromanipulator probes, it is impossible for us to measure the current at every single pixel for a long time. As a practical compromise, we measure the currents as follows:

- 1. for six minutes at pixels in a 6-by-4 array, roughly uniformly distributed across each detector,
- 2. for 15 minutes at pixels in a 5-by-3 array, also evenly distributed,
- 3. for two hours at two pixels on opposite sides of the 6-by-4 array,
- 4. for two hours at the cathode, and
- 5. for six minutes at alternating pixels in areas showing abnormal trends (such as abrupt changes in pixel current), if necessary.

With this measurement scheme, we believe that we are able to characterize the detector leakage over the extent of each detector to sufficient spatial resolution, and at the same time, obtain representative values of the stable current during long operations. Figure 3.14 shows a map of the leakage current values thus obtained across one of the HEFT detectors. Accounting also for the time requiring to adjust the set-up from one pixel to the next, the overall measurements take 14 hours per detector. Based on the result of these measurements, we reject detectors with excessive leakage, as stated above, and rank the remaining detectors to match low-leakage detectors with ASICs with few defective pixels.



Figure 3.13: The decay of leakage current at a detector pixel over two hours.



21-11-2003: CZT-5 621407_K52_03 (no grid, 30µm-gap), HV=-300V---Leakage current at the 6th minute

Figure 3.14: Map of leakage current values across a HEFT CdZnTe detector. Squares with different levels of shading represent the size of the leakage current measured at that particular location. Where individual pixel measurement is not made, the background shading corresponds to the average leakage current per pixel, inferred from the aggregate leakage current measured from the cathode.

3.8.4 Interconnections

In our effort to adopt hybrid semiconductor detectors on the HEFT focal plane, we have to experiment with new and at times challenging methods to connect various pairs of components together.

Flip-chip bonding of the CdZnTe-ASIC hybrids Challenges in our flip-chip bonding process include the accurate alignment of the CdZnTe detector with the ASIC, and making gold stud bumps of uniform height. Stud bumps much shorter than their neighbours cannot make connection with the conductive epoxy on the detector side, thus leaving the pixel disconnected. On the other hand, stud bumps much taller than their neighbours push up the CdZnTe detector within their local region, leaving the neighbouring pixels disconnected. Because the conductive epoxy bumps on the detector side are about 20 μ m tall, we specify a height of (45 ± 3) μ m when we have the ASICs populated with gold stud bumps (by an outside vendor). Stud bumps as much as 7 μ m taller or shorter than the mean are enough to cause disconnections.

Mounting hybrid detectors to printed circuit boards In the HEFT focal-plane design, each focal-plane module consists of a single print circuit board (or a 'motherboard') with support circuitry for two hybrid detectors, which are both glued down to the motherboard. Because the yields of both the ASICs (Section 3.8.1) and the CdZnTe detectors (Section 3.8.2) are low, and their fabrication involves long lead times and characterization labour-intensive, a working detector hybrid does not come easily. If one of the two hybrid detectors on a single motherboard fails for any reason, we have to remove the failed detector without damaging its functional neighbour. This requirement calls for a reversible technique for mounting the hybrids to the motherboard. Besides, in the early days of our ASIC development effort, we relied on mounting and wirebonding the ASICs onto motherboards to test their functionality. Working dies also required a reversible mounting technique so that they could be removed safely for subsequent processing. (We have since acquired a wafer probing station and eliminated the need of repeatedly mounting and wirebonding the ASIC dies.) The usual mounting practice using epoxy does not work in these cases, as epoxy is very difficult to dissolve after cure. Thus, we experiment with alternative adhesive substances, and we decided on thermoplastic interposer films.

The typical application of these thermoplastic films is to place them in between the components to be bonded, and to heat the entire assembly to the melting temperature of the film, at which point, the film liquefies to a glue and bonds the components together. Because our intent is to keep the hybrids removable, we heat the films only to temperatures slightly below their melting point, so that the solid films remain intact. For the particular film we use, Staystik 472 from Cookson Electronics, we heated the motherboard assembly to $90-100^{\circ}$ C, while the specified 'attach temperature range' is $125-200^{\circ}$ C. To remove a failed hybrid detector, we reheat the assembly to the same temperature, at

which point we gently remove the hybrid. For testing ASIC dies, which must be removed regardless of the testing result, we cut the thermoplastic films to lengths slightly longer than the ASIC dies, leaving a tab protruding from underneath one side of the mounted die. To remove the die, we simply pull the tab while grounding the die, without reheating.

Wiring connection between circuit boards As shown in Figure 2.8, the HEFT focal-plane module consists of a stack of two printed circuit boards (PCBs), hosting the hybrid detectors and the module MISC, respectively. A large number of digital signals travel between the two PCBs, for the module MISC to control both hybrid detectors and for data to be read from the detectors to the MISC (through the external analogue-to-digital converter). Because the entire focal-plane module must reside in the small space at the bottom of the shield module well, we use a Nanonics miniature connector with 65 30-gauge wires to make these connection. Excluding the hybrid detectors and the module MISC, this item is the most expensive component in the focal-plane assembly (costing over USD 500) and requires a few months of lead time. The installation of this connector also requires meticulous soldering by hand under a microscope.

Mechanical connection between the focal-plane module and shield module Placement of the focal-plane module within the shield module must be secure and unchanging with time, so that one can measure the position of the detector pixel array with respect to the rest of the instrument once and for all. Because changing temperature has a big effect on the leakage current and spectral performance of the CdZnTe detectors, it is also desirable to thermally couple the small focal-plane modules to the much more massive shield modules for thermal stability. For these purposes, we place the two PCBs that make up a focal-plane module on a custom designed copper stand, as shown in Figure 2.7. The bottom of the stand is surrounded by 'claws' that fit snugly into the interior wall of the shield module. The copper connection also provides some thermal coupling between the PCBs and the shield, although we have not assessed its effectiveness for this purpose. In the actual experiment, when the instrument is in a cold environment $(-28-0 \,^{\circ}\text{C})$ while the focal-plane modules dissipate heat, the temperature at the PCB hosting the hybrid detectors is about 10 $^{\circ}$ C higher than that at the exterior of the shield housing.

3.9 Characterization of the HEFT detectors

In this section, I present the results of our characterization of the HEFT detector hybrids.

3.9.1 Final detector ensemble

Since the first successful, full-scaled detector hybrid we fabricated in 2002, and up till the first two flight campaigns of HEFT in 2004 and 2005, we have fabricated a total of 17 hybrids. Table 3.3 lists their part numbers and some representative properties of each hybrid. We have characterized two of the prototype hybrids, one with the final anode geometry (S/N 4) and one with a gridded anode geometry (S/N 3), in earlier studies, and I refer the readers to Chen et al. (2004b) and Chen et al. (2004a) for further details. Here, I discuss only the six hybrids (S/N 7, 9-12 and 16) that we fabricated to populate the three focal planes of HEFT for its first flight campaign.

We made the four detector hybrids in Modules A (S/N 10 & 11) and B (S/N 7 & 9) from spectroscopy-grade material with the final electrode geometry on the anode plane (see Section 3.6 and Appendix A). The fifth hybrid (S/N 12), on Module C Side A, consists of a CdZnTe detector that eV Products sent us as a mechanical sample. It has the same no-grid anode plane pattern, but its spectral quality is unknown to us. The last hybrid (S/N 16), on Module C Side B, consists of a spectroscopy-grade detector with a gridded anode-plane pattern as described in Section 3.5.2. However, we apply no voltage difference between the grid and the anode. This results in the hybrid having a large effective inter-anode gap of $50 + 14 + 50 = 114 \,\mu\text{m}$.

We took the data from which we produced these results during a 110-hour long continuous calibration run during the Spring 2005 campaign, in a high bay in Fort Sumner, New Mexico, with environmental conditions simulating those in flight (see Section 7.2.1 for more details of the calibration run). These results represent the *average* performance for a population of CdZnTe-ASIC detector hybrids fabricated under realistic conditions.

3.9.2 Spectral resolution

As part of the calibration procedures, we illuminate each of the flight detectors with uncollimated Am-241 and Co-57 radioactive sources. We collect Am-241 X-rays at all three temperatures for calibration $(-15, -5 \text{ and } +5 \degree \text{C})$, and for Co-57, only at $-15 \degree \text{C}$. For each of the four species and temperature combinations, we create a single spectrum from all events recorded at each of the six flight detector hybrids. In the single case where enough events are available (Am-241 at $-5 \degree \text{C}$), we also create spectra for individual pixels. To distinguish the effect of charge-sharing from those of other causes, we show the single- and two-pixel events separately in all spectra.

Comparison amongst different hybrids Figures 3.15 and 3.16 show the single spectrum from each detector for each species and temperature combination, after all calibration steps are applied. The quality of these spectra represents the overall spectral performance of each detector. We observe the following:

Application	Prototype development, HEFT (gridded) Destroved in handling (gridded)	Prototype development, PCB4A, HEFT (gridded)	Prototype development, PCB5A, HEFT	Prototype development, ISAS, Japan	Mössbauer spectroscopy, Fultz's Group, Caltech	Flight detector, PCB13A, HEFT Module B	Destroyed in handling	Flight detector, PCB13B, HEFT Module B	Flight detector, PCB11A, HEFT Module A	Flight detector, PCB11B, HEFT Module A	Flight detector, PCB10A, HEFT Module C (mechanical sample)	Mechanical sample, shake test at JPL for NuSTAR	Mössbauer spectroscopy, Fultz's Group, Caltech		Flight detector, PCB10B, HEFT Module C (gridded)	Mechanical sample, shake test at JPL for NuSTAR
Date of completion	2002 - 3 - 22 2002 - 8 - 16	2002- 9-29	2003 - 4 - 25?	2003 - 7 - 31?	2003 - 7 - 31?	2004- 2-18	2004- 2-18	2004- 2-18	2004 - 6 - 29	2004 - 6 - 29	2004 - 6 - 29	2004 - 6 - 29	2004- 6-29	2004-7-12	2004-7-12	2004-7-12
Bond type	Indium Indium	Indium	Stud+epoxy	Stud+indium	Stud+indium	Stud+epoxy	Stud+epoxy	Stud+epoxy	Stud+epoxy	Stud+epoxy	Stud+epoxy	Stud+epoxy	Stud+epoxy	Stud+epoxy	Stud+epoxy	Stud+epoxy
Detector Serial No.	611293-K33-02 612518-K37-01	611293-K33-05	617119-M03-2-2			621142- $K40$ - 02	621538-K 40 -01	621407-K 52 -03	617742-K34-05	623939-H61-06	621072- $K50$ - 02	621072- $K50$ - 03	0312 - 0104 - 1	611293-K 33 -0 3	611293-K 33 -04	623939-H61-10
Detector Material	CdZnTe CdZnTe	CdZnTe	CdZnTe	CdTe	CdTe	CdZnTe	CdZnTe	CdZnTe	CdZnTe	CdZnTe	CdZnTe	CdZnTe	CdTe	CdZnTe	CdZnTe	CdZnTe
ASIC Serial No.	D10-SN1 C12-SN7	M4	E7	01		B01	B22	06	04	07	A01	B08	B11	M3	B15	H5
Hybrid Serial No.	Hybrid-5 S/N 1 Hybrid-5 S/N 2	Hybrid-5 S/N $\frac{1}{3}$	Hybrid-5 S/N 4	Hybrid-5 S/N 5	Hybrid-5 S/N 6	Hybrid-5 S/N 7	Hybrid-5 S/N 8	Hybrid-5 S/N 9	Hybrid-5 S/N 10	Hybrid-5 S/N 11	Hybrid-5 S/N 12	Hybrid-5 S/N 13	Hybrid-5 S/N 14	Hybrid-5 S/N 15	Hybrid-5 S/N 16	Hybrid-5 S/N 17

Table 3.3: List of HEFT detector hybrids fabricated up to the 2005 flight campaign.

I



Figure 3.15: Spectra of Am-241 at -15 (left column) and -5 °C (right column), from each detector with events from all pixels grouped together. Each row of figures show data from the same flight detector.



Figure 3.16: Spectra of Am-241 at +5 °C (left column) and of Co-57 at -15 °C (right column), from each detector with events from all pixels grouped together. Each row of figures show data from the same flight detector.

- Amongst the six detectors onboard, only one (Module A Side B) performs as well as the 0.9 keV FWHM prototype (Hybrid 5 S/N 4). Without bias from the grid to the anodes, the large effective gap between anode contacts for the gridded hybrid, Module C Side B, substantially degrades the spectral resolution of the detector as a whole. In contrast, the 'mechanical sample' in Module C Side B actually performs better than expected.
- 2. With the gridded detector (Module C Side B), trying to obtain any spectral information from charge-sharing events is a hopeless pursuit. Note that charge-sharing events account for about half of all events, so that we are losing 1/12 = 8.3% overall efficiency for the entire experiment. Unfortunately, we had no choice but to accept this loss (See Section 3.8.2 for the explanation).
- 3. Charge-sharing was more severe for Module B and C at +5 °C. This was in part the result of us setting the cathode voltage bias at -100 V only for Module B (to reduce bad pixels bursting), as opposed to a nominal -350 V at -5 °C. This emphasizes the important of applying a high-enough cathode voltage to obtain as many single-pixel events as possible, and thus as good spectral resolution as possible. This in turn translates to the importance of low-bulk-leakage material.
- 4. Linear temperature correction of the amplifier gain coefficients, as applied to these data, preserves the general shape of the spectra at different temperatures, but the linewidths are still broadened when one measures the linewidths quantitatively.

Spectral performance at individual pixels Figure 3.17 shows spectra of Am-241, each obtained from an individual pixel, from half of the hybrid 'Module A Side A' (Columns 24–47). We show these spectra here in the same relative positions as the corresponding pixels appear on the detector. The red and blue curves represent single-pixel and charge-sharing events, respectively. We have applied all standard calibration steps to these data, except for the correction of charge induction efficiency differences across the detector (to exemplify a phenonmenon described below). Because we have corrected for variations in the amplifier gain but not for variations in detector properties, the position of each event on the abscissa indicates the number of charges induced by the X-rays; thus, the relative position of the prominent 60 keV line at each pixel shows the differences in charge induction efficiency across this detector hybrid.

From these spectra, we observe that for some pixels, the 60 keV line appears only in chargesharing events. Oftentimes, the 'gain' also appears to be reduced at these pixels. For some other pixels, no 60 keV line is seen at all. One possible explanation is that electron transport is slow or halts before the charge reaches the anode at the pixels that appear to have low gain, due to low field or strong trapping. This gives the charge clouds more time to diffuse, which explains the increased charge-sharing. Similarly, increased charge loss along the path of drift to the anode explains the





Figure 3.17: Spectra of Am-241 from individual pixels of the flight hybrid on Module A Side A. The spectra are arranged in the same relative positions as the corresponding pixels appear on the hybrid. The red and blue curves represent single-pixel and charge-sharing events, respectively. The abscissas are in pulse height units, without correction of varying charge induction amplitudes.

reduction in 'gain'. As a consequence, gain correction/verification with the 60 keV line is not possible at these pixels using the calibration data we took in the field. Had the detectors survived the landing, we would have placed them in a cold environment for more calibration runs (but as I shall explained in Section 7.3, the landing was not as clean as we had hoped).

3.9.3 Charge induction amplitude

Variations in the spectra from different pixels on a single detector, as shown above in Figure 3.17, indicate that the number of charges induced by mono-energetic X-rays at each anode across a detector hybrid is not uniform. The number of induced charges is not only a function of the incident X-ray energy, but also of detector properties—specifically, how efficient a detector is in converting the energy of an incident X-ray into an electrical signal at the anode, through charge cloud production, charge transport across the depth of the detector, and induction at the anode. Variations in the pixel spectra indicate that some of these detector properties are localized, showing characteristic length scales smaller than the size of the HEFT detectors (2.4 cm by 1.2 cm), and thus substantial variation across each detector. To assess the amount of variation across the flight detectors, we fit the high-energy half of the spectrum from each pixel with a Gaussian profile, to measure the pulse height value (i.e., the position on the abscissa) of the 60 keV line of Am-241 as recorded by each pixel. These pulse heights are an indicator of the number of charges induced at each pixel by 60 keV X-rays. Figure 3.18 shows a histogram and a map of the pulse heights across each flight detector. Also, they show two distinct types of variations:

- In Module A Side A, change in the pulse heights is abrupt. It appears as a distinct, depressed pattern at the centre of the detector, shaped like a dumb-bell.
- In Module A Side B and Module B Side A, the change is gradual. It appears as two elevated plateau at the centre of each of these hybrids.

While we can expect a certain natural variation in material properties across a detector, as the gradual changes appear to be, the abrupt change in Module A Side A indicates significant inhomegeneity in the detector that is not as tolerable. We believe that there are a large number of localized charge traps at pixels where the pulse height is significantly reduced (as much as halved at some pixels), and that these traps are situated far from the cathode. As the mean free path of 60 keV X-rays in CdZnTe is only $256 \,\mu$ m, and the photons enter the 2 mm-thick detector from the cathode side, the majority of X-rays create charge clouds after travelling only one-eighth the depth of the detector. As the electron clouds drift towards the anode plane, they are trapped, possibly completely, at the lower depth of the abnormal charge traps. If the majority of charges remains at that depth for the duration of signal sampling, the net effect would be an effective reduction of the number of charges



Figure 3.18: Pulse height of the 60 keV line of Am-241, measured at each pixel on the flight detectorhybrids. The left and middle columns show the histogram and map of the pulse heights, repsectively. The right column shows the amplifier gain of the read-out ASICs, which has been compensated in the pulse height maps and histograms, except for half of Module C Side A. To show the dead areas in various hybrids with higher contrast, all maps in this figure show higher values as darker pixels, and lower values as brighter pixels. Each row of figures contains data from the same flight detector.

induced at the anodes; also, the amount of reduction would be roughly proportional to the energy of the incident photons. Note that the electron clouds produce the bulk of the induced signal, due to the close proximity of the photon interaction sites to the cathode, and much stronger electron transport properties compared to holes (manifested in $(\mu\tau)_e$ values two orders of magnitude higher for electrons than $(\mu\tau)_h$ values for holes, as shown in Table 3.1).

A few artefacts in the maps of Figure 3.18 require some explanation. Note that white areas in the maps of the remaining hybrids represent pixels with no data. Many of these pixels, especially those on Module C Side B, the gridded detector with a large effective inter-anode gap, show no clearly identifiable spectral line in their spectra of Am-241. Although the pixels still detect X-ray events, the events appear as a continuum in the spectra, indicating significant charge loss when the steering electrode grid is not in operation. In contrast, most of the pixels with no data on Module B Side B are pixels with excessive leakage, either in the form of a persistently high current or as frequent bursts of charges. In addition, some other pixels do not function altogether.

The uniform increase in pulse height values orthogonal to the rows and columns of Module C Side A is another artefact. It was the result of incorrect determination of the amplifier gain from pulser data for the read-out ASIC in this hybrid. [For unknown reason(s), the pulser data show significant reduction of the amplifier gain for half of the read-out ASIC, which appears to be a systematic error in the only set of calibration data available to us given our extremely rushed flight campaign schedule.] Because the amplifier gain and the amplitude of charge induction have similar multiplicative effects on the pulse heights of X-rays, the error in gain measurement also appears in the X-ray pulse height data, as a change of equal magnitude in the reverse direction. As our measurement of the charge induction efficiency is very sensitive to its strong coupling with the amplifier gain, we also show corresponding maps of the amplifier gain in Figure 3.18 for verification and comparison.

3.9.4 X-ray event statistics

As I discussed in Chen et al. (2004a), the prototype CdZnTe detector-hybrid exhibits substantial variation in the count rate at each pixel across the detector, upon illumination by an uncollimated radioactive source. Unfortunately, but as expected, the flight detectors show similarly large count rate variations.

Using the same data from which we construct the spectra of Am-241 in Figures 3.15 and 3.16, we calculate the total event count at each pixel during each of the three calibration runs involving Am-241 (undertaken at -15, -5 and +5 °C). We then tally the event counts for all pixels on the same detector hybrid into a set of maps and histograms for each hybrid and each temperature setting. Again, we handle single- and two-pixel events separately, and ignore events with more than two pixels triggered. (Note, however, that in the analysis of science targets observed during the first

flight of HEFT, we include events triggering up to five pixels.) Figures 3.19, 3.20 and 3.21 show the resulting graphs and images of event statistics for the three Am-241 measurement sessions conducted at -15, -5 and +5 °C, respectively. We carried out most of the calibration at -5 °C, the expected average temperature range in flight; consequently, the Am-241 data set at -5 °C has the highest statistics. Comparing the three sets of data, we are able to distinguish persistent features in the figures from one-time peculiarities due to an imperfect experimental setup. The following features are common in all three sets of data, and thus are persistent features of the detector hybrids:

- Module B Side B contains a large area of low event counts (with a sizable number of pixels within turned off due to high leakage currents). The high- and low-count regions are delineated by a straight diagonal line that is likely associated with some kind of boundary of the crystal structure. In the corresponding histograms, pixels from the high- and low-rate regions appear as a bimodal distribution (which is the most prominent for the blue curve, tracing the total event count at each pixel).
- A diagonal track of low-count pixels lies across Module C Side B, the hybrid with the gridded detector. This persistent feature is also likely associated with some kind of boundary of the CdZnTe crystal.
- The count rate at the corner of Module B Side A is lower than for the remainder of the hybrid, at all three temperatures, indicating a true reduction of photon detection efficiency.
- In Module A Side A, the collection of pixels with low event counts form a dumb-bell like pattern that is observable in the maps of both single-pixel and two-pixel events. When events triggering up to five pixels are included, however, this pattern disappears. This is strong evidence that transport electrons at these pixels are trapped within the bulk of the CdZnTe crystal well before they reach the collecting anode contact. In this case, charges induced by the trapped electrons are distributed over multiple pixels at the vicinity of the trapping site, and not at the collecting anode alone. Thus, events with multiple triggering pixels at these sites are legitimate X-ray events, unlike those in the other detectors, which are predominantly noise events. In our science analysis, we include events with up to five triggering pixels.

On the other hand, the following feature does not appear in all three sets of data, and is thus due to imperfect experimental setup:

• The reduction of count rate at the corner of Module A Side B at -15 °C is likely due to simple shadowing of the radioactive source by the background shield.

Note that the circular feature at the centre of Module C Side A marks a region of excessive leakage current, and pixels within are all disabled from event triggering. Thus, this feature indicates missing



Figure 3.19: Histograms and maps of event statistics for the calibration run at -15 °C measuring an uncollimated source of Am-241. The three columns, from left to right, show the histogram of event counts at each pixel, the map of single-pixel event counts, and the map of two-pixel event counts. Each row of graphs and images are from the same detector hybrid. The histograms show one- and two-pixel events by the red and green curves, respectively; the blue curve represents the sum of both types of events at each pixel. Both maps on each row are drawn with the same greyscale, with black pixels representing zero count, and white pixels representing values at or beyond the midpoint value on the abscissa of the histogram to the left.



Figure 3.20: Histograms and maps of event statistics for the calibration run at -5 °C measuring an uncollimated source of Am-241. See the caption of Figure 3.19 for an explanation of the various graphs and images.



Figure 3.21: Histograms and maps of event statistics for the calibration run at +5 °C measuring an uncollimated source of Am-241. See the caption of Figure 3.19 for an explanation of the various graphs and images.

data as opposed to low count rates, although its sharp boundary is remarkable in itself—immediately outside the disabled region, pixels show no higher leakage current than for the rest of the detector. Thus, the perimeter of this feature is also likely associated with crystal boundaries. On the other hand, the lack of counts in Column 38 of Module B Side B and an apparent excess of 2-pixel events on Row 6 of Module C Side A are caused by defective circuitry within the read-out ASICs. (Due to severe constraints in time and resources, we had to tolerate up to one dead row or column in the ASICs that we selected to build the hybrids; see Section 3.8.1.)

In general, the distributions of count rates in the various histograms are all broader than expected from pure Poisson statistics. However, apart from the abrupt features described above, the variation in count rates does not show any distinct spatial dependence. Rather, high- and low-rate pixels are both fairly evenly distributed across each detector, resulting in better uniformity when considering length scales larger than the HEFT pixels. Although undesirable, these observations are consistent with those made from the prototype detector Hybrid 5 S/N 4 in Chen et al. (2004a). All these observations further emphasize our need to develop an effective method to screen detectors before they are committed to ASICs and for flight, based not only on the resistivity, but also on their quality as X-ray detectors.

3.9.5 Flatfielding

To test whether we can adjust for the variance in photon statistics observed above, we took the three sets of data, scaled the median pixel event count on each detector in each set to 1, and compared the numbers in each pair of sets. The differences in (normalised) event statistics is displayed in Figure 3.22, as histograms and greyscale maps. The average difference is about 20–25%, not as small as one would have hoped. Note that at +5 °C, there was only -100 V applied to Module B, which explains the very much larger differences. These results show that the relative efficiency of pixels within a single detector is not consistent and stable enough for us to flatfield the detectors using the flood-illumination data we have analysed above.

In order to correct for efficiency differences at different pixels properly, one must scan each detector using a collimated X-ray source of known intensity and with a beam size smaller than the pixel pitch. One must collect sufficient events (a few hundred events per pixel or more) to achieve good statistics at individual pixels. To obtain the absolute quantum efficiency, one must also calibrate the collimated X-ray source used in these measurements using a reference detector of known efficiency (such as a sodium iodide scintillator or a good silicon or germanium detector). We have not made these measurements with the HEFT detectors, and this impairs our ability to correct for efficiency variations from pixel to pixel.



Relative event count comparison for ²⁴¹Am runs, 2005/ 3/19-- 3/22: 2nd column / blue: (-15°C counts/median) - (-5°C counts/median) 3rd column / red: (+5°C counts/median) - (-5°C counts/median) -- column / green: (+5°C counts/median) - (-15°C counts/median)

Figure 3.22: Comparison of event statistics in the three calibration runs at -15, -5 and $+5^{\circ}$ C, measuring an uncollimated source of Am-241.

3.9.6 Summary

From the maps and spectra in this section, we conclude that the quality of the HEFT flight detectors are by no means close to optimal. However, we are limited by serious constraints of time and manpower in a realistic production environment. Various delays prevented us from obtaining enough high-quality detector materials and setting a high rejection threshold during the screening of detector uniformity. Thus, we had no choice but to proceed with the experiment with whatever devices we have available to us. Nevertheless, judging from the satisfactory performance of Module A Side B (Hybrid-5 S/N 11) and the prototype hybrid S/N 4 (Chen et al., 2004a), we can say that given good detector materials, our design does yield good detectors. The remaining challenge is in material selection, and this has yet to be overcome.

Chapter 4

Prototype detector characterization: a best-case scenario

4.1 Introduction

Section 3.9 of this thesis reports on the characterization of the six CdZnTe-ASIC hybrid detectors onboard the HEFT instrument. As explained in Section 3.8, we built these detectors under extremely hasty circumstances, and at times, we sacrificed quality in order to meet compulsory deadlines. As a result, not every detector on HEFT achieve the 1 keV spectral resolution of which we know our detector design is capable. Evidence of this capability comes from our study of an earlier prototype detector, labelled Hybrid 5 S/N 4 in Table 3.3, which we fabricated without the time constraint we faced when producing the detectors for flight. Hybrid 5 S/N 4 has the same detector geometry as five of the flight detectors, except Module C Side B (see Section 3.9). It is mounted on Side A of a prototype motherboard on its own, with Side B left empty. The module MISC controlling this detector implements a slower and more conservative read-out algorithm, with a saturation event rate of 20 events/s. Besides these differences, this prototype system is identical in design to the focal-plane modules of HEFT.

In this chapter, I present the characterization study of this prototype detector. The remainder of this chapter previously appeared as part of the journal article 'Characterization of a large-format, fine-pitch CdZnTe pixel detector for the HEFT balloon-borne experiment' by C. M. Hubert Chen, Walter R. Cook, Fiona A. Harrison and Jiao Y. Y. Lin in the *IEEE Transactions on Nuclear Science*, volume 51, number 5, part 1, pages 2472–2477, in October, 2004.



Figure 4.1: (a) Map and (b) histogram of the electronic noise distribution at room temperature $(23 \,^{\circ}\text{C})$, with the voltage bias off. The mean and standard deviation of the distribution are $540 \,\text{eV}$ and $56.4 \,\text{eV}$, respectively ($108 \pm 11.2 \,\text{electrons}$). The pixels at the bottom-left corner are unconnected, and thus have lower noise amplitudes, between 0.3 and 0.4 keV FWHM (i.e., $60-80 \,\text{electrons}$). All anode contacts on Row 1 are narrower than the ones in the remaining rows (see Section 3.6); note the lower noise level across Row 1.

4.2 Detector Characterization

In this section, we present the results we have obtained to date from the characterization of our first detector of this design.

4.2.1 Electronic Noise

We measure the eletronic noise by sending pulser signals to the input of the preamplifier at each pixel (see Fig. 3.8), which simulate a 75 keV spectral line. We generate these pulser signals with a precision DC voltage source that is accurate to at least 3 p.p.m., so that the intrinsic variation of the pulser signal is negligible. We then assess the electronic noise as the linewidth of the spectral line produced by the pulser signals, as measured by the ASIC circuitry.

Fig. 4.1 is a map of the electronic noise at each pixel at room temperature with the bias off. The dark spots represent quieter pixels, whereas the bright spots represent noisier pixels. The same data is tallied here into a histogram, whose horizontal axis is aligned with the colour scale for the map.

80

For our detector, the electronic noise amplitude is dominated by the input capacitance between the anode contacts on the CdZnTe and the ASIC backplane. So, unconnected pixels such as the ones at the bottom left corner of the map have lower noise amplitudes, as indicated in their positions in the histogram, between 300 and 400 eV (60–80 electrons). Also, the contacts are narrower in the row of pixels next to the mating edge of the detector. Here, the electronic noise amplitude is also smaller than the other pixels on the detector. On the other hand, pixels with relatively high noise amplitudes are found interspersed over the entire detector rather than concentrated in the same place. This suggests random defects in the VLSI fabrication process as a potential cause of the increased noise magnitude. However, the majority of the pixels display good uniformity in noise characteristics across the entire pixel array. On average, the electronic noise in our detector is $(540 \pm 56.4) \, eV$, or 108 ± 11.2 electrons, at room temperature. At lower temperatures, the noise amplitude decreases. On the other hand, it increases when a bias is applied across the detector. The combined effect of temperature and bias at $0 \, {}^{\circ}$ C and $HV = -400 \, V$ is a net increase of the noise amplitude by about $100 \, eV$.

4.2.2 Leakage Current

Another property of interest is leakage current. We obtain the leakage current at a pixel by measuring the amount of charge accumulated in the sampling capacitors (see Fig. 3.8) within a fixed duration when the bias across the detector is off, and when it is on. We then derive the leakage current going into the preamplifier input from the difference in charge accumulated with and without the bias, from the time period, and from the preamplifier gain.

Fig. 4.2 shows the map and histogram of leakage current at each pixel, at room temperature; again, the brighter pixels indicate higher leakage currents, while the darker pixels indicates lower leakage currents. The map and histogram both show that there is substantial variation in the leakage current in this $12 \times 22 \text{ mm}^2$ area of CdZnTe. One also sees in the map a strong spatial correlation in the size of the leakage current. However, if one cools the detector down to 0 °C, then the leakage current becomes negligible. We plot simple i-v curves at the two temperatures in Fig. 4.3.

For practical reasons (to collect sufficient statistics in a reasonable time), we further study the performance of a subset of pixels with an X-ray source at 0 °C. Although the leakage current is negligible at this temperature, we have nevertheless chosen two regions of pixels with very different leakage characteristics at room temperature for further study, as this choice can potentially reveal further differences in material properties in the two regions. These two regions are indicated by the two circles in Fig. 4.2. The circle on the left marks a high-leakage region, while the one on the right marks a low-leakage region.



Figure 4.2: Leakage current (a) map and (b) histogram at room temperature $(23 \,^{\circ}\text{C})$ and detector bias $HV = -250 \,\text{V}$. There is substantial variation in the leakage current at room temperature, with strong spatial correlation. The circles indicate the two regions at which Am-241 spectra are obtained at 0 $^{\circ}\text{C}$. The spectra are displayed in Figs. 4.4 and 4.5.



Figure 4.3: Current-voltage relations for the two regions circled in Fig. 4.2, at two different temperatures. There is substantial difference in the two regions at room temperature, but the leakage current becomes negligible at 0 °C. From these measurements, the apparent bulk resistivity at room temperature is $3.4 \times 10^{10} \,\Omega$ -cm in the high leakage region and $6.3 \times 10^{10} \,\Omega$ -cm in the low leakage region. At 0 °C, it is above $5 \times 10^{11} \,\Omega$ -cm in both regions.



Figure 4.4: Am-241 spectrum at 0 °C, obtained from the high-leakage region (at room temperature) circled in Fig. 4.2. The cathode-to-anode potential is HV = -400 V. Data collection took place for 69 minutes.



Figure 4.5: Am-241 spectrum at 0 °C, obtained from the low-leakage region (at room temperature) circled in Fig. 4.2. The cathode-to-anode potential is HV = -400 V. Data collection took place for 58 minutes.

4.2.3 Spectral Resolution

For each of the two regions selected, we illuminate it with an Am-241 source that is collimated to within a circle of about ten pixels in diameter. We then measure the spectrum at each region at 0° C, which are shown in Fig. 4.4 for the region with high leakage currents at room temperature, and in Fig. 4.5 for the region with low leakage at room temperature. Both of these spectra include all events where one or two adjacent pixels are triggered (we refer to them as one- and two-pixel events, respectively, in short). A tally of the experimental data indicates that one- and two-pixel events account for more than 90% of all the events when the detector is illuminated with a source flooding the region under study, with the remainder being events with three or four adjacent pixels triggered. These statistics are consistent with detector modelling results we found previously (Chen et al., 2002).

The processing of events with multiple pixels triggered requires additional explanation. For twopixel events, one has to reconstruct the photon energy by adding the pulse heights from the two triggered pixels together, with the different gains and offsets of the two pixels properly corrected. These corrections are done in software, after the pulse height from each triggered pixel is calculated individually from the samples read out and digitized from the ASIC circuitry (see Section 3.7 and Appendix C). For events with more than two pixels triggered, we currently do not process them, due to their relatively small contribution, and they do not enter the spectra shown in Figs. 4.4 and 4.5.

At low energies, electronic noise is the dominating component of the X-ray line widths. At high energies, the line width is further affected by variation in the depth of photon interaction, and the energy resolution is close to about 1 keV (200 electrons) FWHM at 60 keV at 0 °C. As was indicated in the previous section, leakage current turns out not to be an issue at 0 °C.

4.2.4 Count Rate Uniformity

Yet another important property for a pixel detector is the uniformity of count rate. We measure the count rate by illuminating the entire detector with an uncollimated point source of Am-241, and measuring the number of counts at each pixel. This measurement was done at room temperature, and the detector bias was set to HV = -80 V, a value that is high enough for the detector to be operational, but just low enough that no pixel is saturated by the leakage current.

Fig. 4.6 shows the map and histogram of the count rate for one- and two-pixel events (i.e., over 90% of all events) that correspond to the 60 keV line; as this histogram indicates, we have observed high variation in the count rate that is clearly more than merely Poisson statistics, with the variation being close to the mean. On the positive side, these variations are stable—the count rate at a given pixel remains the same over time—thus, the variation is correctable when the detector is put into use. However, we have tested an identical ASIC that is bonded to a cadmium telluride sensor instead of CdZnTe, and we have not seen such high count rate variation with the cadmium telluride detector. So, this variation in the count rate seems to be a CdZnTe property. Also, when an equal number of test pulses are sent to the input of the electronics at each pixel, no such variation is seen in the number of test pulses detected; this excludes the behaviour of the ASIC circuitry as a source of the variation.

Per our detector specification to eV Products (reprinted in Appendix A), this detector should be a single crystal, and should contain no precipitate greater than $100 \,\mu$ m in size. eV Products supplied us with a photograph of the detector taken with infrarad microscopy before the contacts were deposited, and the photograph shows no significant feature that can explain this count rate variation. We suspect that this variation may be due to a non-uniform electric field within the CdZnTe, which causes each pixel to have a different 'effective collecting area'. If this is indeed the case, then a gain in collecting area by a pixel will be compensated by a corresponding loss of area by


Figure 4.6: Count rate (a) map and (b) histogram. The data was taken at room temperature and detector bias HV = -80 V. This bias is high enough for the detector to be operational, but just low enough for none of the pixels to saturate at room temperature. Only events corresponding to the 60 keV line of Am-241, and events triggering one or two pixels, are included in this figure. The mean and standard deviation of the distribution are 61.8 and 19.5 counts/pixel, respectively.



Figure 4.7: (a) Map and (b) histogram of the count rate (displayed in Fig. 4.6) convolved with a boxcar of three pixels in width in each direction. Pixels on Row 1 are narrower (see Section 3.6), and the count rates are thus lower accordingly. The lower rates on Row 2 and at the remaining three edges of the detector are due to the convolution of a finite data set. Pixels at the lower left corner are unconnected. These edge pixels and unconnected pixels are not included in the histogram; specifically, only pixels within the interior rectangle marked by Rows 3–23 and Columns 2–43, and outside the unconnected region marked by Rows 20–23 and Columns 2–6, are included in the histogram. The included pixels are outlined in the pixel map in (a) above. Mean = 60.98 counts; $1\sigma = 8.48$ counts.

neighbouring pixels. So, if one averages the count rate at each pixel with the rates at its neighbours, the resulting distribution should be fairly uniform. This distribution is shown in Fig. 4.7, where we convolve the rates in Fig. 4.6 with a boxcar of three pixels in width in each direction. The distribution after this convolution agrees with Poisson statistics (mean = 60.98 counts; $1\sigma = 8.48$ counts). Based on the same logic, one would expect to see a correlation between the count rate and the leakage current at each pixel. The correlation is plotted in Fig. 4.8, and we do see some correlation.

We note that this investigation of the count rate is still under active investigation as of the writing of this paper; a scan of the detector with a fine X-ray beam and the calibration of the beam with a detector of known efficiency, both to be done in the near future, will provide us with more concrete information on the count rate. Until we report further on the results of these investigations, this report on count rate variations can only be considered preliminary, and any conclusion made in

86



Figure 4.8: Leakage current-count rate correlation for the HEFT detector. The count rates are from Fig. 4.6 (room temperature, HV = -80 V), and the leakage currents are from Fig. 4.2 (room temperature, HV = -250 V). Unconnected pixels and pixels registering fewer than 10 event counts are not shown in this figure. The cause of the low count rates at these pixels are likely unrelated to the leakage current, and distorts the correlation. The remainder of the data points, shown in dots, are fitted with a straight line relation, counts = $0.54 \times (current/pA) + 21$, which is also shown in the figure.

Table 4.1: Summary of Some Quantitative Properties of the HEFT detector

Property	Range	
Leakage current	$10-140+\mathrm{pA/pixel}$ at $23^{\circ}\mathrm{C}$	
	$0-10 \text{ pA/pixel at } 0^{\circ}\text{C}$	
Eletronic noise	$(540 \pm 56.4) \text{ eV FWHM at } 23^{\circ}\text{C}$, bias off	
Spectral resolution	$\approx 1.0 \mathrm{keV}$ FWHM at $60 \mathrm{keV}$ at $0^{\circ}\mathrm{C}$	
Count rate	variation \approx mean	

this section speculative.

4.3 Summary

In summary, for the High Energy Focusing Telescope, we have developed CdZnTe pixel detectors with large format, fine pitch and low power consumption. The leakage current varies substantially at room temperature, but is negligible at 0 °C. The electronic noise is around 500 eV at 23 °C, whereas the spectral resolution is about 1 keV FWHM at 60 keV at 0 °C. Table 4.1 summarizes the values of various quantities for the detector. We are still actively investigating the high variation in the count rate, and we shall report new results on this topic in a future paper.

Chapter 5

Comparison of two detector architectures

5.1 Introduction

In Section 3.5 of this thesis, I describe two hybrid bonding techniques and two anode geometries that we have studied through the development of the HEFT detectors. Our assessment of these competing detector configurations is based on a comparison of two prototype CdZnTe-ASIC detectors—Hybrid 5 S/N 3 and Hybrid 5 S/N 4 in Table 3.3. Hybrid 5 S/N 3 has a grid of steering electrodes on the anode plane of the CdZnTe detector, positioned at the pixel boundaries and in between small anode contacts. Indium-bump bonds connect about one-third of the pixel anodes on the CdZnTe crystal to the preamplifier input pads on the ASIC. In contrast, Hybrid 5 S/N 4 has no steering electrode grid on its anode plane; instead, large anode contacts cover 88% of the pixel pitch (limited by our detector vendor's capability in contact deposition). Conductive epoxy and gold stud bumps connect 94% of the anode contacts to preamplifier input pads. Figure 3.7 shows these two anode geometries, with and without grid, side by side. Figure 3.5 shows a photograph of the indium bumps in Hybrid 5 S/N 3, and Figure 3.6 shows photographs of the epoxy bumps and an epoxy-stud bond in S/N 4. Finally, Table 5.1 lists the characteristics of the two hybrids side by side.

These two detector configurations implement different approaches to tackle the same problem to reduce charge loss in 'multiple-pixel events' or 'charge-sharing events'. These are events occuring near pixel boundaries, where collection of transport electrons are shared by pixels on both sides of the boundaries, inducing signal in each. Yet, electrons can be trapped at the bare CdZnTe surface in between pixel contacts on the anode plane, due to the low electric field strength in this region. These trapped electrons appear as a decrease in the measured event energy, thus degrade the spectral resolution of the detector. To reduce charge loss in Hybrid 5 S/N 3, we apply a voltage difference between the grid of steering electrodes and the anode contacts. This creates an electric field across the gaps of bare CdZnTe surfaces in between the electrodes, and steers the otherwise

1	0	1 1
Hybrid 5 Serial No.	S/N 3	S/N 4
Materials	CdZnTe-Si ASIC	CdZnTe-Si ASIC
Strategy to address	Electric field across gaps	Minimize gap area
charge loss at gaps		
0 01		
Anode dimensions (μm) :		
Pixel pitch	498	498
Anode width	384	468
Anode area ratio	1	1.5
Steering electrode width	14	
Bare CdZnTe gap width	50	30
Inter-anode gap width	$14 + 2 \times 50 = 114$	30
01		
Hybrid bonds	Indium	Gold stud + silver epoxy
bond height (μm)	8-10	45-50
Input capacitance ratio	1	0.3
1 1		

Table 5.1: Comparison of the characteristics of Hybrid 5 S/N 3 and S/N 4.

trapped transport electrons towards the anodes. In contrast, for S/N 4, we minimize charge loss by minimizing the gaps in between anode contacts. Yet, the 1.5 times larger anode contacts in S/N 4 increase the input capacitance of the amplifier circuit, amplify thermal noise at the input (see Section 3.3.2), and degrade the spectral resolution for all events (single-pixel and charge-sharing). Thus, we develop the epoxy-stud bonds for S/N 4, which are five times taller than the indium bump bonds in S/N 3, to compensate the effect of increased anode area on the input capacitance. As mentioned in Section 3.5.2, the net result is a 70% decrease in the input capacitance from S/N 3 to S/N 4, a substantial improvement.

To compare the performance of these two different detector architectures, here I present a characterization study of the two hybrids, S/N 3 (hereby the 'gridded detector') and S/N 4 (hereby the 'detector with no grid'). The remainder of this chapter previously appeared as part of the article 'Characterization of the HEFT CdZnTe pixel detectors' by C. M. Hubert Chen, Walter R. Cook, Fiona A. Harrison, Jiao Y. Y. Lin, Peter H. Mao and Stephen M. Schindler in the Proceedings of SPIE, volume 5198: *Hard X-Ray and Gamma-Ray Detector Physics V*, pages 9–18, in January, 2004.

5.2 Detector characterization

In this section, we present the detector performance under a range of temperatures and bias voltages for both architechtures: the grided detector with indium bumps, and the detector with no grid, bonded with epoxy studs. For practical reasons (to collect sufficient statistics in a reasonable time) we present detailed results from a contiguous area of 10×11 pixels.



Figure 5.1: Distribution of electronic noise in a 11×10 pixel area of the gridded detector, at four different bias and temperature combinations. The inset map shows the spatial distribution of the same data, with brightness proportional to the noise magnitude.

5.2.1 Electronic noise

The intrinsic energy resolution of the detector hybrid at low X-ray energies is predominantly determined by the electronic noise in the ASIC circuitry. We measure the electronic noise as the full width at half maximum (FWHM) of a Gaussian spectral line produced by electronic pulses with energies equivalent to 75 keV photons. The distribution of electronic noise for the gridded detector is shown in Fig. 5.1.

5.2.1.1 Temperature dependence

The dotted and dashed lines in Fig. 5.1 show the distributions of the electronic noise for 110 pixels, at room temperature (≈ 22 °C) and at 0 °C, respectively, when all electrode biases are set to zero. Zero bias ensures that we are measuring the noise component from the electronics itself, rather than shot noise caused by leakage current through the CdZnTe crystal that is channelled into the preamplifier inputs. At room temperature, our ASIC design has an energy resolution of (624 ± 29) eV FWHM at 75 keV, due to thermal noise in the circuitry; at 0 °C, jitter in the circuitry decreases, and the resolution is improved slightly to (542 ± 34) eV FWHM at 75 keV. If pulses at these pixels are all summed together, the resulting 75 keV line has a FWHM of 623 eV at room temperature and 540. eV at 0 °C.

5.2.1.2 Leakage current contribution

For the grided detector, leakage current is introduced by surface leakage between the grid and contact, as well as by bulk leakage. The magnitude of the contribution depends on the surface and bulk resistivities, which vary from detector to detector, and the operating bias voltage. For the grided hybrid evaluated here, when the biases are set to nominal values of -300 V at the cathode and -4 V at the steering electrode grid, both relative to the anodes, noise is introduced by surface leakage current, and the resolution degrades to (791 ± 99) eV at room temperature and (757 ± 36) eV at 0° C. The resolutions in the summed spectra are 779 eV and 756 eV FWHM, respectively. The solid and dash-dotted lines in Fig. 5.1 show the FWHM distribution at the same 110 pixels at these biases.

For the detector without a grid, the leakage current results from the bulk leakage component only.

5.2.2 Spectral resolution for X-ray events

To characterize the response of the detectors to X-ray events, and thus the performance of the CdZnTe sensors, we tested the detectors with an Am-241 source collimated into a circular beam with a 10- to 11-pixel (5 mm) diameter.

5.2.2.1 Single-pixel events

Figure 5.2 shows the distribution of the 59.54 keV line widths for the 110 pixels under the collimator at each detector. The measurements were made at 0 °C, the targeted operating temperature for HEFT. The plot shows only events triggering one pixel. For the gridded detector with 10 μ m indium bumps, the energy resolution ranges from 0.8 to 1.4 keV FWHM at 59.54 keV, with the majority of the pixels having 1.0 keV FWHM. For the non-gridded detector with 40 μ m high epoxy and stud bumps, the average energy resolution is improved by an average of ~ 0.1 keV.

Figure 5.3 shows the summed spectra from all 110 pixels. Also indicated in the figure are the measured Gaussian line widths of the various spectral lines of Am-241. At low energies, the line widths are comparable to the electronic noise (the pulser line width), indicating the absence of systematic effects (e.g., incomplete charge collection) in the CdZnTe. For the gridded detector, the 59.54 keV line shows some residual low-energy tailing, with a FWHM of 931 eV. If one ignores the low-energy tail and fits the line only down to the lower half-maximum point, then the FWHM becomes 863 eV. In contrast, the non-gridded detector produces a more symmetric 59.54 keV line with a FWHM of 825 eV (756 eV if the low-energy tail is ignored).



Figure 5.2: Distribution of line widths from the 59.54 keV line of 241 Am, at 0 °C, as measured at 11×10 pixels of each of the two detector designs. The histogram on the left shows a slightly better performance of the detector without a grid (solid) than the gridded detector. The intensity maps on the right show the spatial distribution of the line widths. Both maps are drawn with the same greyscale, with brightness proportional to the line width. Note that pixels at the corners are shielded from the source by the circular collimator opening. These pixels are not included in the histogram.



Figure 5.3: Spectra of ²⁴¹Am, from all single-pixel events summed over an 11×10 pixel area of the HEFT detectors. The Gaussian line widths as fitted from the spectra are listed. The non-gridded detector in (b) measures a more symmetric 59.54 keV line than the gridded detector in (a).



Figure 5.4: Spectra of 241 Am, from all 2-pixel charge-sharing events summed over an 11×10 pixel area of the HEFT detectors. For comparison, the spectra from single-pixel events are scaled to the same count rate and displayed here in dotted lines.

5.2.2.2 Charge-sharing events

The recovery of charge-sharing events is important for detectors with pixels of this small size, since these events account for as much as 50% of the total (Chen et al., 2002). Figure 5.4 shows the spectra from the two detectors obtained from charge-sharing events only. The 59.54 keV line measures a FWHM of 1.83 keV from the gridded detector, and 1.18 keV from the non-gridded detector. It has come as a surprise to us that the non-gridded detector, with its narrower interpixel gaps, does better than the gridded detector, with its steering electrodes. When single-pixel and charge-sharing events are both summed together, the spectra thus produced are shown in Figure 5.5. The FWHM measurements are 1.23 keV and 973 eV for the gridded and non-gridded detectors, respectively.



Figure 5.5: Spectra of $^{241}\mathrm{Am},$ from all events summed over an 11×10 pixel area of the HEFT detectors.



Figure 5.6: Intensity maps of the FWHM (left) and skewness (right) of the 59.54 keV line of ²⁴¹Am at various electrode bias pairs. Darkness is proportional to the FWHM magnitude on the left and to the skewness magnitude on the right, so that the most desirable configurations are the brightest ones in each case. There is a partial trade-off between FWHM and skewness.

5.2.3 Optimization of operating biases for the gridded detector

To minimize charge trapping in the CdZnTe crystal, electrode biases have to be tuned appropriately. We measured the energy resolution of the 59.54 keV line at various bias combinations at the cathode and steering electrode of the gridded detector. For each configuration we measured the width (FWHM) and skewness of the line. Figure 5.6 shows their trend as a function of the bias voltages. We see that while greater grid-to-anode biases are always desirable up to -4 V (producing a surface field strength of $4 \text{ V}/50 \,\mu\text{m} = 80\,000 \text{ V/m}$), cathode biases that are too large will reduce the relative grid bias ($E_{\text{grid}-\text{anode}}/E_{\text{cathode}-\text{anode}}$), and thus reduce the steering effect of the grid and increase tailing of the spectral line. On the other hand, cathode biases that are too low decrease the electron mean free path through the crystal, thus increasing charge trapping and subsequently the line width. Therefore, we need to find a balance between the two effects; according to our measurements, a -450 V cathode bias is appropriate (producing a bulk field strength of $450 \text{ V}/2 \text{ mm} = 225\,000 \text{ V/m}$).

5.3 Summary

The HEFT detectors are optimized to achieve good energy resolution with low power. By careful design of the ASIC circuitry and off-chip digital data processing, we have achieved energy resolutions below 1 keV FWHM at hard X-ray energies, for a large-format CdZnTe-ASIC hybrid detector at 0 °C.

In addition, the results above show that the two detector architectures—one with steering electrodes and indium bump bonds, and the other with narrow inter-anode gaps and epoxy-stud bondsproduce very similar spectral performances, both meeting the HEFT requirements.

Chapter 6

Numerical modelling of charge sharing in CdZnTe pixel detectors

In this chapter, I present my earlier work on the modelling of charge-sharing events (see Sections 3.5.2 and 5.1) in 'gridded detectors' similar to, but smaller in size than, the HEFT prototype detector Hybrid 5 S/N 3 in Table 3.3. The remainder of this chapter previously appeared as the article 'Numerical modeling of charge sharing in CdZnTe pixel detectors' by C. M. Hubert Chen, Steven E. Boggs, Aleksey E. Bolotnikov, Walter R. Cook, Fiona A. Harrison and Stephen M. Schindler in the *IEEE Transactions on Nuclear Science*, volume 49, number 1, part 2, pages 270–276, in February, 2002.

6.1 Introduction

Our group at the California Institute of Technology has been developing CdZnTe pixel detectors for use in astrophysical applications, and specifically, the balloon-borned High Energy Focusing Telescope (HEFT) (Harrison et al., 2000). These detectors operate at hard X-ray energies, from 5 keV to 150 keV. Our goal is to achieve the highest spectral resolution possible by having a customdesigned, low-noise VLSI readout circuit, by optimizing the pixel contact geometry, and by careful processing of the pulse height data.

In an earlier paper (Cook et al., 1999) we reported studies of the detector performance for single-pixel events—events in which incident photons are absorbed by photoelectric effect within a single pixel, with all the induced charges collected by the pixel. This paper describes subsequent effort to study charge-sharing events—events in which photons are absorbed in between pixels, inducing charges on two or more pixels. The phenomenon of charge-sharing has been studied by other groups, for both CdZnTe (Kalemci and Matteson, 2000) and silicon detectors (Arodzero et al., 1996). Previous studies of charge-sharing in the HEFT detectors were described in Bolotnikov et al. (2000a). Here, we first summarize the experiments, and then describe a numerical model we have

developed in an attempt to understand and interpret the experimental results.

6.2 Detector geometry

The prototype HEFT CdZnTe detector, shown in Fig. 6.1, contains an array of 8×8 pixels. The pixel pitch has dimensions 680 μ m × 650 μ m (centre-to-centre spacing), and the CdZnTe crystal is 2 mm thick between the cathode and anode planes. The cathode plane is covered entirely with a single platinum contact, set at the negative bias voltage. The geometric configuration of the anode plane is shown in Fig. 6.1. At the centre of each pixel is a rectangular metal contact for the anode, held at ground potential. At the perimeter, the pixel contact is surrounded by a 'steering electrode', held at a negative potential relative to the anode. The steering electrodes of all pixels are connected into a grid, as shown in Fig. 6.1. Between the anode contacts and the grid of steering electrodes are bare surfaces of CdZnTe, which we refer to as 'the gap' (between the pixel contact and steering electrode). To investigate the effect of different gap geometries, we have divided the anode plane into four quadrants of 4×4 pixels, each having a different gap size. The widths of the gaps in the four quadrants are 100 μ m, 150 μ m, 200 μ m and 250 μ m. In all four quadrants, the common grid of steering electrodes is 50 μ m wide. The CdZnTe crystal is indium-bump bonded at the anodes to a low-noise VLSI circuit for read-out. On the opposite side, the cathode plane is irradiated with X-rays.



Figure 6.1: The anode-plane configuration of the CdZnTe detector. Drawn to scale. The pixels studied in Section 6.3 are labelled '1' and '2'.

The VLSI chip provides, for each pixel, a complete precision analog signal processing chain, including a charge sensitive preamplifier, shaping amplifier and a peak detect and hold circuit. For a 40 μ s peak time setting and simulated detector leakage near 0.1 pA, the measured noise referred to the input is equivalent to approximately a 0.25 keV FWHM energy loss (or about 50 electrons) in the CdZnTe crystal. More details on the circuitry and performance of the chip can be found in

6.3 Phenomenology of charge-sharing

We refer the readers to an earlier paper (Cook et al., 1999) for the analysis of single-pixel events in this CdZnTe pixel detector. We now summarize the observations from studies of charge-sharing events.

6.3.1 Experimental setup

To investigate charge-sharing events, we placed a collimated source of ²⁴¹Am on the cathode side of the detector, opposite to two adjacent pixels (denoted 1 and 2 in Fig. 6.1) on the anode side, and measured the pair of pulse heights, (E_1, E_2) , generated at the pixels at each event. We set the bias voltage at the cathode to HV = -250 V relative to ground at the anode, and the grid of steering electrodes to various values between $V_{steer} = -16$ V and -10 V. The ambient temperature was about 13°C. Both pixel 1 and 2 are in the quadrant with gap size 150 μ m. First, we positioned the collimated source at the centre of pixel 1, measured the pulse height generated at each event, and obtained an energy spectrum, S_1 , of single-pixel events at pixel 1. Next, we did the same steps at pixel 2, obtaining spectrum S_2 of single-pixel events at pixel 2. Then, we placed the collimated source midway between the centres of the two pixels, and measured pairs of pulse heights, (E_1, E_2) , in this configuration, in which the charge from most events is shared between pixels. Finally, we drew a scatter plot of E_2 against E_1 for all charge-sharing events in our data. We performed this sequence of steps repeatedly, varying the potential at the steering electrodes from -10 V to -16 V, in increments of -2 V, producing a series of E_2 against E_1 scatter plots.

The issue of energy calibration requires some special explanation. We found that in single-pixel events, the 59.5-keV spectral line of ²⁴¹Am appeared as a much wider line and at lower channels when the source is collimated between the pixels (i.e., when all (E_1, E_2) are measured) than when it is collimated at the centre of either (when S_1 and S_2 are measured). Fig. 6.2 shows the relevant spectra. This observation prompted us to measure from spectrum S_1 the conversion factor between pulse height and energy for pixel 1, and apply it to the E_1 values. To calibrate the E_2 values, we went through the same procedure using the ²⁴¹Am line in spectrum S_2 as the reference. This change in pulse height is further discussed in Section 6.3.2 below.

6.3.2 Observations from our experiment

Fig. 6.3 shows a typical plot of the energy-calibrated pulse heights, E_2 against E_1 , for charge-sharing events. The 'track' diagonally across the plot represents pairs of pulse heights from the 59.5-keV



Figure 6.2: Spectra of ²⁴¹Am from single-pixel events, when the source is collimated at the centre of a pixel (solid lines) and in between two (dotted lines). The 59.5 keV line is both widened and shifted to lower channels when the source is moved from the centre to the edge of each pixel.



Figure 6.3: Plot of pulse-height pairs (E_1, E_2) measured at adjacent pixels in charge-sharing events (the dash lines delimit the region included when calculating the energy and width at the centre of the 'track').

photons in charge-sharing events. The two square symbols, joined together by a solid line, indicate the reference points at $E_1, E_2 = 59.5$ keV, as obtained from spectra S_1 and S_2 . The triangular symbols, also joined by another solid line, indicate the pulse height positions for single-pixel events measured when the collimated source is in between the pixels. We have identified three main features in these graphs:

- 1. 'Parallel shift' of the 'track': For both pixels, the pulse height measured for single-pixel 59.5keV photons is lower when the collimated source is at the edge of the pixel than when it is at the centre. Thus, there is an apparent shift to lower energies in single-pixel events occurring near the edge of the pixel, compared to those near the centre (this is also the reason for our method of energy calibration when the source is in between the pixels, as described at the end of Section 6.3.1 above).
- 2. 'Curvature' of the track: In charge-sharing events, if all charges are collected by either of the pixels, then the sum of the pulse heights, $E_1 + E_2$, should be equal to that in single-pixel events. In other words, the track of data points should coincide with the straight line joining the triangular symbols in Fig. 6.3. This is clearly not the case for our data, as observed in the figure. The curvature of the track implies that more energy is unaccounted for—i.e., more charges are lost—when E_1 and E_2 are comparable than when they are not. Note that this effect pertains only to charge-sharing events, and is distinct from the 'parallel shift' we have just discussed, which affects both single-pixel and charge-sharing events.
- 3. Escape photon 'kinks': Rather than appearing as a narrow line, the track appears as a broad distribution of points, extending to two 'kinks' at $(E_1, E_2) \approx (25 \text{ keV}, 35 \text{ keV})$ and (35 keV, 25 keV) on the high energy side. We note that one of these values, 25 keV, is close to the energies of the K-shell electron in cadmium (23.2 and 26.1 keV) and tellurium (27.5 and 31.0 keV), and that the energies in each pair sum to 59.5 keV, the ²⁴¹Am photon energy. We attribute this phenomenon to the escape of the K-shell photon, produced as the atom deexcites after the initial photoelectric interaction. The mean free path of these K-shell photons ranges from 60 to 135 μ m, which is a reasonable fraction of the pixel size of our detector. Thus, it is possible for the K-shell photons to not be deposited in the same pixel, resulting in the observed distribution of points along the track in Fig. 6.3.

It is clear from observations 1 and 2 that in charge-sharing events, the measured pulse height does not completely account for the energy of the incident photon. In other words, charge loss for events near the gaps in the CdZnTe crystal degrades the detector spectral resolution. While we have plotted only the results for pixels 1 and 2 here, the same results are repeatable at other pairs of pixels in all four quadrants. In order to understand the detailed charge transport for charge-sharing events, which may shed light on ideas to improve the spectral resolution, we have developed a numerical model of the detector. We now describe the latest version of the model and the results obtained thus far.

6.4 The numerical model

The numerical model emulates the transport of electrons and holes in the detector. To simplify the computation in the first stage of analysis, we assume that the effects of drift and diffusion can be considered separately and combined at the end. We calculate the electric and weighting fields within the CdZnTe detector by solving the three-dimensional Poisson's equation numerically on a rectangular grid of Cartesian coordinates. We set the cathode and anode planes parallel to the xy-plane, and represent each pixel by 40×40 grid points. In the perpendicular (z-) direction 80 grid points cover the 2 mm depth of the CdZnTe crystal, and an additional 10 cover the 8- μ m layer of air between the anode plane and the VLSI back plane. For simplicity and speed, we model the pixels as a 3×3 array of 667 $\mu m \times 667 \mu m$ square pixels, instead of the 8×8 rectangular ones of similar dimensions in reality. As for boundary conditions, we assume the steering electrodes surrounding the 3×3 array of pixels extend to infinity. The models are modified versions of programmes originally used for modelling charge transport in germanium detectors (Boggs, 1998; Philhour et al., 1998). The algorithm we use to solve Poisson's equation is the Gauss-Seidel method with 'simultaneous over-relaxation', as described by Press et al. (1986, \S 17.5). Note that we convert the equations given by Press et al. into their equivalents in three dimensions. All derivatives are approximated by first-order finite differences. The boundary condition at the gap is set by the conservation of surface and bulk leakage currents. Having obtained the electric and weighting potentials from Poisson's equation, we calculate their gradient fields by using fourth-order finite differences to approximate the derivatives (third-order at the boundaries between CdZnTe and the metal electrodes, using only potential values on the CdZnTe side). This gives us accurate determination of the weighting field, $\vec{W}(\vec{x})$, as well as the electric fields, $\vec{E}(\vec{x})$, for any combination of biases we apply on the cathode and the steering electrodes.

Besides the fields, the drift of charges in the detector also depends on their mean free paths, $\lambda_{n,p}(\vec{x})$, in CdZnTe. For electrons (n), $\lambda_n(\vec{x}) = v_n(\vec{x})\tau_n$, where τ_n is the mean trapping time and $v_n(\vec{x})$ is the local drift velocity. From earlier experiments (Bolotnikov et al., 1999a), we have found that the drift velocity is not saturated up to field strengths of at least 3.0 kV/cm, so that $v_n(\vec{x}) = \mu_n |\vec{E}(\vec{x})|$, where μ_n is the electron mobility in CdZnTe. The same formulae apply for holes (p), with the corresponding hole parameters. We have determined from the same experiments (Bolotnikov et al., 1999a) the value of $\mu\tau$ for our detectors to be 1.5×10^{-3} cm²/V for electrons, and 1.0×10^{-5} cm²/V for holes. In this model, we allow for different values of $\mu\tau$ at the detector surface, to account for the possibility of increased trapping. The presence of crystal surface defects that inhibit charge transport is well-known in other semiconductor material such as silicon, so we expect to see similar effects in CdZnTe as well. As one measures the product $\mu\tau$ together experimentally, $(\mu\tau)_{surface}$ serves as a convenient indicator of the charge transport quality on the crystal surface. In this model, we use $(\mu\tau)_{surface}$ instead of $(\mu\tau)_{bulk}$ only on the anode plane; i.e., this surface has no thickness in the z-direction.

With the fields and $\mu\tau$ determined, we next find the charge (electrons and holes) trajectories. We start tracing charges at a depth of 256 μ m (the mean free path of 59.5 keV photons in CdZnTe) from the cathode plane. We cover this xy-plane uniformly with 40×40 starting positions for each pixel, displaced from the electric field grid nodes (described earlier in this section) by half the grid spacing in both the x- and y-directions. We then trace the charges in constant spatial steps (as opposed to constant time steps) of length $|d\vec{x}| = 0.1 \ \mu m$ in the direction of the electric field lines. We find the field in between grid points by trilinear interpolation. At each step, we compute the contribution of the traced charge to the charge induced at the anodes as $dQ(\vec{x}_i) = q_i \vec{W}(\vec{x}_i) \cdot d\vec{x}_i$, where q_i is the amount of charge traced at step *i*, between \vec{x}_i and $\vec{x}_i + d\vec{x}_i$. The sum of these contributions along a trajectory thus gives the charge induced at each pixel due to the charge traced along that trajectory. We also account for charge trapping by modifying q_i after each step, so that for electrons, $q_{i+1} = q_i \exp\left[-\frac{|d\vec{x}_i|}{\lambda_n(\vec{x}_i)}\right]$, where $\lambda_n(\vec{x}_i) = v_n(\vec{x}_i)\tau_n = \mu_n\tau_n|\vec{E}(\vec{x}_i)|$, as described in the previous paragraph. If the electric field guides an electron to land on the bare CdZnTe surface on the anode plane, the $\mu\tau$ value is modified from $(\mu\tau)_{bulk}$ to $(\mu\tau)_{surface}$, and tracing continues along the CdZnTe surface until an anode is reached. By tracing the holes and electrons on all trajectories from the photon interaction depth (at 256 μ m from the cathode plane) to the cathode and anodes respectively, we know the total charge induced, $Q_j(\vec{x}_0) = \sum_i dQ_j(\vec{x}_i) = \sum_i q_i \vec{W}_j(\vec{x}_i) \cdot d\vec{x}_i$, at each pixel j, as a function of the starting position, \vec{x}_0 , of the drift (with $q_i = q_{i-1} \exp\left[-\frac{|d\vec{x}_{i-1}|}{\mu\tau|\vec{E}(\vec{x}_{i-1})|}\right]$, as shown above). To convert $Q_j(\vec{x}_0)$ into energy values $E_j(\vec{x}_0)$, we pick the trajectory that goes through the centre of the pixel, and let the charge induced due to this trajectory, Q^* , correspond to the incident photon energy, namely 59.5 keV. Thus, $E_j(\vec{x}_0) = (59.5 \text{ keV}) \frac{Q_j(\vec{x}_0)}{Q^*}$. This conversion method is consistent with the energy calibration in our experiment, as described at the end of Section 6.3.1.

So far we have only considered drift. We incorporate the effect of diffusion by the following method: For each starting position \vec{x}_0 , we weigh the pulse height function, $E_j(\vec{x})$, by a twodimensional Gaussian distribution of starting positions (in the xy-plane) centred at \vec{x}_0 . Then, we sum all the weighted pulse heights to give the pulse height measured at pixel j, due to a photon interaction at \vec{x}_0 (in practice, we implement this as the convolution of $E_j(\vec{x})$ with the Gaussian distribution). The Gaussian distribution represents the charge cloud produced by diffusion; we set its width to be $\sigma_x = \sigma_y = 26 \ \mu m$, an estimate from experiments (Bolotnikov et al., 1999a) and from the Einstein relation. Finally, we plot pairs of the convolved pulse heights measured at adjacent pixels due to the same photon interaction against each other. We then compare the resulting plot with pulse-height plots obtained experimentally, such as the one shown in Fig. 6.3.



6.5 Results and discussion

Figure 6.4: Cross-section of the CdZnTe detector showing the electric field and potential near the anode plane.

We show the result of the electric field calculation (for $V_{steer} = -10$ V) in Fig. 6.4. Note that far from the anode plane the field lines are essentially straight lines. Fig. 6.4 also shows the structure of the electric field near the low-field region adjacent to a steering electrode. We find the lowest field strength to be 102 V/cm, while the field in the bulk of the detector is on the order of 1 kV/cm. Also according to the model, a 10 V difference across 'the gap' suffices only to prevent electrons from being collected at the steering electrodes, but fails to prevent them from drifting onto the bare CdZnTe surface. In fact, the model indicates that one cannot steer all field lines away from the bare CdZnTe surface even with higher potential differences across the gap. On the other hand, excessively high voltages increase the surface leakage current, which in turn degrades the performance of the readout electronics.

Based on analyzing the pulse heights due to a given charge trajectory, we classify the pixel plane into regions of photon interaction: single-pixel events, or 2-, 3- or 4-pixel coincidence. This allows us to estimate the proportion of each type of event, assuming constant efficiency across the detector, when an uncollimated source illuminates the entire detector. These regions are shown in Fig. 6.5. For $V_{steer} = -10$ V, and with a $2\sigma_{baseline}$ event threshold, 64% of all events should be single-pixel events, 33% are 2-pixel coincidence, while 3- and 4-pixel coincidences are 1.5% each. This shows that charge-sharing events are a significant fraction (about one-third) of all detections. In order to recover useful spectral information from these events, a clear understanding of charge sharing is essential.



Figure 6.5: Regions corresponding to single-pixel and coincidence events at the corners of four adjacent pixels. The number of pixels triggered is labelled in each region. The rows and columns of crosses indicate the starting positions of charge trajectories; 40×40 trajectories cover this $(667 \ \mu m)^2$ plane. Also shown in shade are the positions of gaps between the anodes and the steering electrode grid.

Fig. 6.6 shows a major result of our calculation—a plot of the pulse heights at adjacent pixels in charge-sharing events, with the assumption that $(\mu\tau)_{surface} = (\mu\tau)_{bulk}$. On top of a scatter plot of the experimental data, we superpose the curve obtained from the numerical model, marked by crosses along its length. Each cross marks a data point obtained from a drift trajectory. Recall that we have placed 40 trajectories uniformly along each side of a pixel (see Fig. 6.5), so the distance between adjacent trajectories is 667 μ m/40. Having eight crosses on the curve indicates that both pixels are triggered only when the initial photon interaction occurs within seven trajectory spacings of the boundary between pixels, i.e., 667 μ m/40 × 7 = 117 μ m, or when $\pm 2.24\sigma_x$ of the charge cloud intersects the pixel boundary. The calculated curve lies on the higher energy side of the 'track', and it is apparent that our model with no free parameters predicts lower charge loss than our measurements indicate. To account for this discrepancy, we vary the value of $\mu\tau$ on the surface of the CdZnTe crystal, as mentioned in Section 6.4, and find the value of the ratio $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}}$ that makes the model and the data agree.

Fig. 6.7 shows the variation in the calculated curve as $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}}$ varies between 0 and 1. As Fig. 6.7 indicates, the amount of charge loss is very sensitive to the value of $(\mu\tau)_{surface}$. This is to be expected, as the weighting field is strongest near the anode plane, so that the induced charge is also the most dependent on the drift of the traced charges in this region. Unfortunately, there is yet no method to accurately determine this quantity experimentally. On the other hand, by finding the values of $(\mu\tau)_{surface}$ that make our modelling results agree with the experimental data, we are able to produce independent estimates of $(\mu\tau)_{surface}$, given in Table 6.1.



Figure 6.6: Comparison of charge loss scenarios as predicted by our model and as observed in experiment. The model assumes $(\mu\tau)_{surface} = (\mu\tau)_{bulk}$. The experimental data are the same as the ones shown in Fig. 6.3; the line marked with crosses along its length is the modelling result.



Figure 6.7: Variation in the extent of charge loss computed with $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}}$ varied between 0 and 1. Curves shown are for $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}} = 0, 0.05, 0.10, 0.16, 0.40$ and 1.0.

Table 6.1: Estimates of $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}}$ from numerical modelling. $\frac{V_{steer}}{(\mu\tau)_{surface}}$ -10 V 9-10%

 $\begin{array}{rrrr} -10 \ \mathrm{V} & 9{-}10\% \\ -12 \ \mathrm{V} & 14{-}17\% \\ -14 \ \mathrm{V} & 17{-}30\% \end{array}$

We must emphasize that when obtaining these estimates, we fit the calculated curves to the data by eye only. To exemplify this, we show a series of calculated curves that 'fit' the data in Fig. 6.8.



Figure 6.8: Modelling results with the value of $\frac{(\mu\tau)_{surface}}{(\mu\tau)_{bulk}}$ adjusted to fit the data.

A clear problem exists with the values given in Table 6.1—if the additional charge loss is indeed solely due to reduced $\mu\tau$ on the CdZnTe crystal surface, then $(\mu\tau)_{surface}$ should be constant, independent of the electric field, and thus the bias applied at the steering electrodes. We believe this problem can be caused by some change loss mechanism yet unaccounted by this model. The $(\mu\tau)_{surface}$ estimates here may well be 'effective values' only, shielding other charge loss effects. Potential causes include the possibility that on the crystal surface, the drift velocity saturates, and is no longer proportional to the very strong electric field at the gap (contrary to our assumptions). On the other hand, the trapping time may also depend on the field in some unknown manner. While we may be able to dismiss the problem by introducing various surface defects in the model, we are reluctant to incorporate any feature that is not undisputably observable in experiments. However, we do think that our simplified treatment of diffusion may not be sufficient to account for its entire effect. We note in particular that as the charges pass through the low-field region below the steering electrode shown in Fig. 6.4, diffusion may cause charge to enter and be trapped there, causing additional charge loss. We have roughly estimated the drift velocity there to be still about ten times greater than the diffusion speed. Yet, more detailed study of the effects of diffusion will be needed, and this will be the next step in extending our numerical model.

6.6 Conclusion

In this study, we have demonstrated that a simplified numerical model of the charge transport processes can reproduce the general trend of the charge-sharing in CdZnTe pixel detectors. Our calculation has also yielded estimates of the surface $\mu\tau$ of our CdZnTe detector, as shown in Table 6.1. However, more effort is needed in order to account for the exact amount of charge loss observed experimentally. The next step in our modelling effort will be to study diffusion in a more detailed way, so that we can reach a complete understanding of drift and diffusion in CdZnTe pixel sensors.

Chapter 7

The first HEFT flight campaigns

The HEFT project began in 1996, under the NASA Office of Space Science's 'Supporting Research and Technology' programme. The initial flight was to happen circa 2000, but challenges in technology development, some of which I mentioned in Chapter 3, caused delays and pushed back the first flight campaign eventually to the autumn of 2004. Even then, we continued to encounter problems during that campaign, and missed the time window when wind direction in the upper atmosphere turned around, which enabled flights longer than a few hours. We returned for our second campaign in the spring of 2005, and finally launched HEFT on 2005- 5-18, from Fort Sumner, New Mexico. This chapter documents these two flight campaigns, including the 25-hour balloon flight.

7.1 Integration

The many components of the HEFT payload were fabricated at various institutions within our collaboration—the focal-plane modules were built at Caltech in Southern California; the mirrors were formed at Columbia University in New York, coated at the Danish Space Research Institute, and assembled into modules in Colorado; the gondola was built at the Lawrence Livermore National Laboratory in Northern California. Eventually, we must bring all these components together for integration and overall testing and calibration at one location. Originally, integration was to take place in Livermore, where the bulky gondola was located, starting at the beginning of April, 2004. However, as we had not completed assembling the three focal plane modules until mid-August, only days before our much delayed departure to the launch site in Fort Sumner, New Mexico, we had no choice but to abandon the original plan, and instead, integrate the payload at the launch site.

The integration process began as soon as the gondola hardware arrived in Fort Summer in late August, 2004. We first reassembled the two-storey tall gondola, which had been taken apart for ease of transportation. We decided on the final positions to put the various attitude sensors, including the gyroscopes, magnetometers, the accelerometer and the star trackers, and installed them onto the gondola frame and the truss. We also installed onto the gondola additional components, including the pressure vessel hosting the flight control computer and power supplies, the 50 lithium batteries, various antennae and transmitters for communication with the ground station, and a latch to hold the truss in a stowed position during ascent and descent.

For the focal plane, the mechanical design of our focal-plane flange governed that we had to close the pressure vessel by installing the kevlar domes onto the flange, sealing all focal-plane components within, prior to installing the entire 220-lb focal-plane pressure vessel onto the telescope truss. Thus, integration of the focal-plane assembly began with the installation of the three focal-plane and shield modules, as well as other electronic components and sensors, onto the focal-plane flange. Calibration and testing of the housekeeping sensors followed. At the same time, we began building the pressure vessel by coating the interior of three prefabricated kevlar domes (two on each side of the midflange and one for spare) with epoxy to seal all leaks across the kevlar fabric. We left the epoxy coating to dry, after which we bolted one dome to each side of the mid-flange with 45 bolts equally spaced all around the 1.3 m-wide flange, with a torque of 180 in-lbs. A lubricated silicone O-ring in between each dome and the mid-flange sealed the dome-flange interface. A meticulous leak check then followed: we injected freon into the pressure vessel at about 2 psi overpressure, monitored the pressure-to-temperature ratio over about a day, and scanned every patch of the kevlar dome surfaces with a hand-held freon detector to detect possible leaks. With the pressure vessel in place and tested, we refilled it with a gas mixture of nitrogen and carbon dioxide in a 3-to-1 ratio (the carbon dioxide was there to prevent arcing of the electronics). Finally, the entire launch team, with the help of a crane, installed the entire 220-lb focal-plane assembly onto the telescope truss, with 30 equally spaced bolts, nuts and insulating spacers all around.

At the opposite end of the telescope truss, we built an adjustable frame to hold the three mirror modules in positions matching those of the focal plane modules. We installed this frame onto the much wider telescope truss, originally designed to hold 14 modules. We then installed the mirror modules onto the frame, with metal shims in between. We roughly aligned the boresights of the mirrors to the axis of the truss, using a laser pointer and a low-power telescope a few metres away to shine and look through a hole that goes through the boresight of each mirror module to measure its projected position on the focal plane flange. Subsequent fine adjustment is discussed later in Section 7.2.2.3.

As the various components of the payload came together, we discovered unforeseen mechanical incompatibilities between the components, and made adjustments as needed. Examples of these adjustments included repositioning the star trackers to avoid conflict with the X-ray telescopes and to properly balance the telescope truss, readjustment of the distance between the mirrors and the detectors to maintain the correct focal length, and sawing out parts of the mechanical structure that protrude and prevent free movement of the telescope truss. To keep the payload below the weight limit imposed by the capacity of the carrying balloon, we also stripped the gondola of unnecessary

panels and containers to the bare minimum. With all core components in place, we made and applied thermal insulation, consisted mainly of fiberglass, aluminized boPET films and paper around all electronic components (mostly within pressure vessels), and around the mirror modules. We also touched up the titanium white paint on most of the payload. These measures were aimed at preventing failure of the electronics and deforming of the mirrors and other machanical parts, due to excessive temperature variation in flight.

7.2 Calibration

Before we can put the HEFT instrument into use, we have to perform certain calibrations and obtain a set of reference data. It is only when we compare data obtained in flight with this reference data set do we know the in-flight measurements in absolute scale and proper physical units. To do so, we need to obtain conversion coefficients (and formulae) for converting measured quantities from hardware-specific scales and units to absolute scales and physical units. These conversions include:

- **Energy** We convert measured pulse heights, in arbitrary pulse-height channels or analogue-todigital units (ADUs), to absolute energy values, in keV.
- **Flux** We convert measured counts, in events per second, to photon flux or energy flux, in photons or keV per unit time (e.g., seconds) per unit area (e.g., cm²).
- **Time** We convert timing data from relative clock periods or 'frame counts', in seconds since system start-up, to absolute time, in UTC or similar universally recognised timing standards.
- **Direction** We convert pixel coordinates, where events are registered, to sky positions, in horizontal (El,Az) or equatorial (RA, Dec) coordinates.

Calibration of energy and time involves only the focal-plane system, by and large, as the mirrors and pointing affect these quantities little. Reconstruction of directional and positional information involves the focal plane (for pixel information), the mirror-detector alignment, and the attitude system (for pointing information). Calculation of flux also involves all three subsystems of HEFT the mirrors, focal plane (both for efficiency information), and attitude control (for determining the density of the intervening medium, including the column densities of the Earth's atmosphere and of our galaxy in the pointed direction). Here, I present two calibration procedures relevant to the focal plane: calibration of the detector hybrids themselves, and alignment of the mirrors with the detectors.

7.2.1 Detector calibration

Scope Calibration of the detectors involves mainly the measurement of the following properties:

- Gain and offset of the read-out electronics, which provide the conversion factor between photon energy and pulse height channel number.
- Linearity of the gain. Because CMOS electronics is linear only in small change approximations, whereas the voltage swings in the HEFT amplifier ASICs are large (0-4.5 V), in order to cover an over 100 keV wide energy band, we need to measure the linearity of the gain in addition to its absolute value, at least to second order.
- Temperature dependence of the gain. The balloon-borne platform, on which HEFT operates, is subject to heating and cooling due to the diurnal temperature cycle and exposure to the Sun. Thus, the ambient temperature of the electronics must vary inevitably, regardless of temperature controlling measures. Because electrical properties such as resistance vary substantially with temperature, so does the gain, and we have to meaure this variation for the expected range of temperatures in flight. We make measurement at three temperatures (-15, -5 and +5 °C), and correct temperature variations to first order only.
- Quantum efficiency of the detector, which is an input into flux calculations. We rely on in-flight observations of bright targets with known and non-variable flux only to obtain the efficiency of the overall system.
- Area differences of the pixels (i.e., flatfielding). As was pointed out in Sections 3.9.4 and 3.9.5, the collecting areas of the detector pixels vary across each detector. Thus, uniform illumination of a HEFT focal plane will not produce a uniform map of measured event counts. We must measure this variation in order to differentiate variations in sky brightness from variations due to detector inhomogeneity.

Each pixel of the read-out ASIC is almost a complete and independent circuit of its own: each pixel has its own gain, offset, efficiency and other properties. Thus, we have to repeat measurements of the quantities above at every single pixel (i.e., $6 \times 48 \times 24 = 6912$ pixels in total). Note that although some of the properties, such as the gain of the single read-out amplifier on each ASIC, are common to all pixels of the same hybrid detector, and some other properties, such as the efficiency of the CdZnTe crystal, are very similar, there are other components, such as the capacitor in the pulsing circuitry, that vary significantly across the chip. Therefore, although one can systematize and automate the calibration and data reduction software to a large degree, a pixel-by-pixel calibration is ultimately unavoidable.

Methodology To measure most of the quantities enumerated above, we perform three experiments:

- 1. First, we operate the detector hybrids without any source of event trigger—no source of X-ray photons or artificially generated test pulse. Then, we force a trigger at each pixel to generate a read-out of the ambient levels of the sampling capacitors. This data provides us with the offset values of the pixel electronics.
- 2. Next, we inject electric pulses of five different but well-calibrated amplitudes at the input of the preamplifier at each pixel. These electric pulses simulate charges induced by monoenergetic X-rays, and cause the pixels to trigger, producing controlled events recording the signature of X-ray events at five predetermined energies.
- 3. Finally, we flood irradiate each focal-plane module with a 'point' source of radioactivity to measure the *relative* collecting areas of the pixels on each focal plane. (Calibration of the *absolute* efficiency will be addressed later in this section.) For this procedure, we use americium-241 and cobalt-57. Americium-241 is suitable for our calibration purpose, due to its well-positioned spectral lines at 59.54 and 13.90 keV, which cover both the high and low ends of the HEFT energy band, 10–70 keV (determined at the low end by the hardware triggering threshold of the focal-plane electronics, and at the high end by the reflectivity drop of the tungsten multilayer mirror coating). Cobalt-57 is our choice for an in-flight calibration source, as its spectral lines at 14, 122 and 136 keV surround the HEFT energy band, are all within the sensitive range for the electronics, and there is no line within the $20-70 \,\mathrm{keV}$ band of observation to contaminate events from the sky. One disadvantage with Co-57 is its short half-life of 271 days. This adds to administrative complications during the busiest of times, as we cannot purchase sources well ahead of a field campaign. Originally, we planned to install the calibration sources at the top of the shield modules, about 30 cm from the CdZnTe detectors. Yet, reassessment of the source strength revealed that Compton scattering of the 122 keV photons (and at 136 keV, to a much smaller degree) within the CdZnTe detectors produces a continuum background up to $\frac{2E^2}{2E+m_ec^2} = 39.4 \,\mathrm{keV}$ (the Compton edge for 122 keV photons, where the scattering angle is 180°), and that this continuum would substantially increase the in-flight background event rate. Thus, we had to place the $1.63 \,\mu$ Ci in-flight source at the location of the mirrors 6 m away, much farther from the focal planes than was originally planned, to lower the flux at the detectors. We estimate that this continuum background contributes roughly 2×10^{-5} counts/s/cm² to the total background of 1.5×10^{-4} counts/s/cm²/keV, determined by measurements on the HEFT graded-Z shield and plastic scintillator.

In practice, we perform the first two steps above at the same time, alternating between forced triggers and pulsed triggers. We do this not only for the sake of higher efficiency, but also to ensure that these two closely related sets of data are taken at identical environmental conditions (temperature, voltage settings, and etc.). Throughout these three steps, we maintain the temperature of the entire apparatus at (-5 ± 1) °C, the expected average operating temperature of HEFT at balloon-borne altitudes. To measure the effect of temperature variations to first order, we repeat all three steps twice, at (-15 ± 1) and $(+5 \pm 1)$ °C.

The actual process Unfortunately, due to the many delays and an extremely tight production schedule for the focal planes, as mentioned previously in Section 3.8, we were only able to complete fabricating the three focal planes in late April, late July and mid-August of 2004 for the field campaign of Fall, 2004. In addition, we had decided on the final data format only one month prior to the campaign, and implemented the supporting hardware and software on both the transmitting and receiving ends only days before we left. Although we had taken a number of partial data sets with the focal planes that were ready early, no complete systematic calibration was possible until all hardware and software components of the focal plane were ready. With the last focal plane completed only one week prior to our move to the field, we had no choice but to perform the definitive calibration of the focal planes during the field campaign, outside our well-equipped laboratory at Caltech.

The most debilitating consequence of not performing the detector calibration within the confines of our laboratory was that we had no means to directly measure the absolute quantum efficiency from pixel to pixel. In the laboratory, we could have used a highly collimated ($200 \,\mu$ m wide) radioactive source to illuminate each detector pixel; we could have compared the count rate response of each pixel with the response from a carefully calibrated photon counter, such as a sodium iodide scintillation detector. We could have placed the collimated X-ray source on an automated translation stage, and scanned the X-ray beam across all pixels on each detector. However, in a high bay in the field, we could not afford such luxury. So, to deduce the flux of celestial targets, we had no choice but to rely on in-flight measurements of a well-studied target of known brightness. For this purpose, we chose to observe the Crab nebula, which has a well-known and stable power-law spectrum:

$$\frac{\mathrm{d}N}{\mathrm{d}E} = 10E^{-2.05} \exp(-\sigma n_{\rm H}) \,\mathrm{photons/s/cm^2/keV}, \text{for } 0.1\text{--}100 \,\mathrm{keV}$$
(7.1)

(Zombeck, 1990, pp. 192). Note that this target must not exhibit a time-varying flux. Thus, the well-known high-mass X-ray binary and black-hole candidate Cyg X-1 is unsuitable as a calibration source, despite being the brightest hard X-ray source in the sky. The final flux calibraton will depend on both the relative pixel area measurements performed in the field, and comparison with an observation of the Crab nebula.

Another big challenge associated with running the calibration in the field is the issue of temperature control. Because we had not put in any active cooling measure in the design of the focal-plane pressure vessel, we had to install ad-hoc hosing to pipe liquid nitrogen (LN2) through the pressure

vessel in order to lower the temperature of the focal plane to (-5 ± 10) °C for the calibration process. Note that cooling with LN2 was the only practical measure, due to the extremely large thermal mass of the mostly metal structure of the focal plane (the mid-flange alone weighed about 90 lbs, and the shield modules each about 15 lbs). Also, the environment must be dry to prevent water vapour or frost from building up on and around the delicate detector hybrids; the pressure vessel naturally provided such protection, with the expense that we had to cool down its entire volume. Our first in-field cooling system, implemented during the campaign of Fall 2004, linked the pressure vessel to an LN2 tank via a bronze pipe a few metres in length (and coiled into a spring for easier handling). The idea was to let the LN2 evaporate as it flowed through the long pipe, producing cold and dry nitrogen gas at a high rate. The cold and dry gas passed through the pressure vessel and exited via an open valve on the mid-flange at the opposite end from the long pipe; in the process, the gas carried heat out of the focal-plane system. Unfortunately, this simple system was inadequate for our task. Due to a combination of hasty planning, lack of practice and established procedures, poor visual temperature display and errors incurred in an urgent emergency situation, we failed to notice LN2 overflowing the input pipe, spilling into the pressure vessel itself on 2004-9-12, the second day into an over-the-clock, continuous calibration opration. This was not a situation that we foresaw in advance. As a result of this debacle, the LN2 created cracks in the interior coating of one of the kevlar domes that made up the pressure vessel, and damaged it beyond repair. On the other hand, we were extremely fortunate in that this was the only major damage incurred during this hapless accident. All three focal-plane and shield modules were unharmed, and we had one spare kevlar dome ready for replacement.

A positive consequence of this accident was that we stepped back and slowed down afterwards, making sure that each of our subsequent steps were thoughtful, conservative and low-risk. In the weeks that followed, we concentrated on testing, debugging and improving the software components of the focal-plane and GSE systems extensively. With the expense of time, we discovered and fixed a number of bugs, and made the overall system much more robust, preventing similar accidents from happening in the future. Because of the amount of time involved in getting the system ready, we missed the turnaround of Fall 2004. With the maximum flight time reduced to five hours postturnaround, we decided against a short engineering flight, and opted to retreat for more preparations for the next field campaign of Spring 2005 instead. We stored the payload in the field in the meantime, for reasons of cost and efficiency.

In the months that followed, we concentrated our effort on devising a viable and robust cooling system for in-field calibration. By March 2005, the beginning of the Spring campaign, it was ready and put to the test. Instead of an open system where LN2 evaporated and cooled the focal plane directly, the new system contained a closed LN2 system to prevent accidental contact of LN2 with the focal-plane components. Much like the liquid cooling system in modern-day automobiles, the new system passed LN2 through the pressure vessel in a daisy chain of hoses; the LN2 absorbed heat indirectly via conduction of the metal hose. A 'radiator' (used here as a heat absorber) further helped in the heat transfer. With this new system, it took about one hour for the pressure vessel to cool from +25 °C to -5 °C; in the process, we monitored the temperature continuously and adjusted the rate of LN2 flow accordingly to prevent overshooting the target temperature. After the temperature stabilized, a change in the LN2 flow rate manifested in change of the system temperature only after a 1-hour delay, due to the large thermal mass of the focal-plane assembly. Thus, careful control of the temperature by fine tuning the LN2 flow rate was crucial. Nevertheless, this system successfully and safely provided us with stable temperature control over the continuous week-long calibration process.

Due in part to the difficulty of cycling the focal-plane temperature between room temperature and -5 °C, we devised a plan to take all the data required for calibration in one single continuous run. Manned with a six-person detector team, we divided each day into three shifts of at least two persons each to attend to the data collecting and cooling systems (no one was allowed to man a shift by themselves to prevent operation errors). After 110 hours of continuous data collection, we finally gathered all the data we needed for a complete and systematic calibration of the detector hybrids.

7.2.2 Optics-detector alignment

Accurate determination of the coordinates of incoming radiation requires the proper alignment of the various parts of the HEFT telescope system. The overall alignment consists of the following components:

- determination of the relative displacement between the two detector-hybrids on each focal plane;
- determination of the angle of rotation of each focal plane relative to the axes of the telescope truss (i.e., the pitch axis and the yaw-roll coupled boresight);
- alignment between the optical axes of the mirror modules and the boresight of the telescope truss; and
- alignment between the optical axes of the mirror modules and the optical axis of the on-axis star tracker.

Each of these items, except the last two, is a separate procedure in itself.

7.2.2.1 Hybrid-hybrid alignment

The positioning of detector-hybrids onto a HEFT motherboard is an imprecise procedure. The pixel positions on each focal plane is defined by the position of the semiconductor detector relative to

the motherboard. As these two components are not directly in contact with each other, it is really the positioning of the detector relative to the ASIC during the flip-chip bonding process and the positioning of the ASIC relative to the motherboard when mounting a hybrid to the motherboard that determine the final quality of pixel-focal plane alignment. However, the flip-chip bonding is a process that we outsource to an outside vendor experienced in the technique; as a consequence, we have little control over the precision of relative alignment between the two parts, apart from the physical constraint that the hybrid bonds must lie within the extent of the anode contacts, each a square 468 μ m wide. Unfortunately, once a detector is bonded to an ASIC, the anode plane pattern, which defines the pixels, is no longer visible from the outside. Because the width of the detector area that lies between the pixel array and the physical edges of the detector varies naturally, there is an intrinsic uncertainty in the pixel-focal plane position that we cannot compensate by proper positioning of the hybrid onto its motherboard. In addition, when designing and fabricating the motherboard circuitry, we have not put in alignment marks to make the placement of hybrids onto the motherboards precise and repeatable. Consequently, we can only do our best to position the two hybrids on each focal plane as close to each other as possible, in order to minimize the dead area between them. We must then determine the pixel positions on each focal plane in experiments afterwards.

To determine the relative positions of the two pixel arrays on each focal plane, we illuminate each completed focal plane with an X-ray source through a mechanical mask in between the source and the detectors. The idea is to measure the positions of incident X-rays on each detector, and thus the projected positions of the mask openings onto each detector, then deduce the relative positions of the two hybrids from their respective positions relative to the common mask. We use two masks one with parallel slits placed perpendicular to the inter-hybrid gap, and one with a raster array of small holes also placed orthogonal to the inter-hybrid gap (and thus also to the pixel arrays). Figure 7.1 shows a mechanical drawing of these masks. Note that we do not have precise control over the positioning of the masks in the experimental apparatus; consequently, we cannot guarantee that features on the masks are perfectly orthogonal to the pixel arrays. However, because errors in the angle of rotation of the detectors are relatively insignificant in the final determination of sky coordinates (see Section 7.2.2.2 below), simple alignment of the masks to the hybrids by eye suffices. With each of these masks over the motherboards, we record the positions of incident X-ray events through the mask openings. Figure 7.2 shows the images recorded by four of the six flight detectors from an experiment using the slit mask. From the pixel positions where the event count transits between high and low, we determine the offset between the two detector-hybrids.

We note that our method, as described above, is not the only way to determine the relative positions between the two hybrids. By scanning each focal plane with a highly collimated X-ray source mounted on a translation stage, we would have been able to determine the position as well



Figure 7.1: Mechanical drawing of the masks used in the determination of the relative positions of the two hybrids on a HEFT focal plane.



Figure 7.2: X-ray images of a radioactive point source through a slit mask, recorded by four of the six flight detector-hybrids, for the determination of the relative positions between the two hybrids on a HEFT focal plane.
as the efficiency of each pixel in one single experiment. However, this experiment would have taken much longer time; given the one week or less between the completion of the HEFT focal planes and our departure date to the launch site, we simply did not have time to perform such an experiment. Our simple test using the masks is a desperate attempt to quickly solve a problem given our time and manpower constraint.

7.2.2.2 Focal plane-truss alignment

The second component of the overall alignment scheme is the determination of the angle between each focal-plane module and the axes of the truss. While the translational positions of the focal planes within the truss are fixed simply by the slots on the flange which we designed to hold the focal-plane and shield modules, we have not provided similar mechanical constraint on the rotation of the motherboards within each shield module in the design of these components. Because the rotation of a focal plane translates to a constant offset in roll and yaw, we have to measure the angles between the focal planes and the truss in order to eliminate this potential systematic error.

While this is a necessary procedure, we note that it is not one requiring the utmost accuracy. The angular offset of a focal plane is only observable when its magnitude is large enough to shift the majority of events at one pixel to the adjacent pixel. For this to be the case at pixels at the periphery of the 48-by-44-pixel focal plane, the angular offset has to be at least 1/24 = 42 rad $= 2^{\circ}4$. Given that the observation targets will be much closer to the centre of the focal plane most of the time, the tolerance level for this error is in fact even higher (i.e., less stringent). Thus, we decide to measure the angular offsets using a rather primitive method.

To measure the angular position of each focal plane relative to the truss, we remove the metal cover at the top of each shield module, with the focal-plane module installed within, to reveal the translucent plastic scintillator. We then take a photograph of the focal plane through the well of the shield module (see Section 2.5.4), where X-rays enter the focal plane. Figure 7.3 shows one of these photographs. These photographs capture the orientation of each focal plane, as indicated by the inter-hybrid gap, and the orientation of the shield module, as indicated by the opening in the scintillator where cables pass through to carry signals between the focal-plane module and components on the focal-plane flange outside the shield. We measure the angle between the two features from the photographs by hand using a ruler and a protractor. This exercise gives us a rough estimate of the angular offset from each focal-plane module to its shield assembly. As there are tabs at fixed positions around each shield assembly, which we designed to match sockets also at fixed positions on the focal-plane flange, we obtain the relative orientation between each shield and the truss simply from specifications on their mechanical drawings. Combining the offsets from the focal plane shown in Figure 7.4.



Reflective scintillator housing Plastic scintillator Lead-tin-copper background shield

CdZnTe-ASIC hybrid detectors (within well)

Tunnel for electrical cables

Orientation of the Inter-hybrid gap

Figure 7.3: Photograph of a focal-plane module within a background shield assembly, with the cover of the shield removed, revealing the translucent plastic scintillator within. Photographs like this one are used to measure the relative angle between each focal-plane module and the axes of the telescope truss.

7.2.2.3 Focal plane-mirror and X-ray-optical alignment

For focused X-rays from a mirror module to land on the focal plane behind the mirrors, we must position the mirror modules on the truss so that the optical axes of the mirrors are parallel to the boresight of the truss. This is an important and challenging task, due to the small focal planes and long focal length involved—an error as small as $22 \times 0.498 \text{ mm}/6 \text{ m} = 7'$ is enough to point the optical axis completely off the detector-hybrids. Besides, the throughput of the mirror modules decreases significantly as the incident ray deviates from the optical axis (as shown in Figure 2.3). Therefore, the quality of alignment between the optics and the focal planes directly affects the throughput of the telescopes. In addition, for us to tell the direction at which the X-ray telescopes are pointing in flight, we need to align the optical axis of the on-axis star camera to the optical axes of the X-ray telescopes, so that the X-ray and optical systems image the same field of view simultaneously. We align these three components—the X-ray optics, the X-ray focal planes (which are locked to the truss), and the star trackers—in a single experiment. Chonko (2006, Section 3.3.3) and Kruse Madsen (2007, Section 1.7) describe this alignment process in detail; for completeness of this thesis, I summarize the process in brief below.

We carried out this alignment process in January, 2005, during the winter in between the two flight campaigns, while the payload remained in Fort Sumner. We built a metal plate (hereby called the alignment plate) with circular openings that mirror the positions of the X-ray optics and on-axis star tracker on a cross-section of the truss. We positioned the alignment plate 72 m away from the payload, facing the optics, such that the optical axes of the X-ray and optical systems on the payload passed through the respective openings on the alignment plate. Figure 7.5 shows a diagramme of the



Figure 7.4: Angular offsets of the focal plane modules with respect to the axes of the telescope truss.



Figure 7.5: Experimental set-up in the alignment procedure.

experimental set-up. We achieved fine alignment by imaging the alignment plate with the on-axis star camera, and adjusting the alignment plate according to the pixel position where the opening on the alignment plate appeared. With the alignment plate in proper position, we shone an X-ray beam through the openings of the alignment plate towards the X-ray telescopes, using an X-ray generator that the High-Energy Replicated Optics (HERO) group at the Marshall Space Flight Center kindly lent us. The X-ray mirrors focused the X-ray beam, which formed an image on each of the X-ray focal planes. We then varied the orientation of each X-ray mirror module by inserting metal shims between the mirror module and its holding plate. In each orientation, we illuminated the X-ray imaging system with an X-ray beam and obtained an image. We relied on the fact that incident X-rays from a specific direction and at a specific off-axis angle will produce a beam spot on the focal plane with a specific pattern (characterized by the amount of elongation and its pixel position). By comparing the appearances of the images with simulated images we had previously prepared, we determined the optimal orientation of each X-ray mirror module with respect to the telescope truss (and thus the focal plane). Figure 7.6 shows the actual and simulated images of the X-ray beam obtained from all three modules, side by side. In this process, the optical axes of the on-axis star tracker and the X-ray telescopes are also automatically aligned. Because the mechanical structure of the truss flexes depending on the weight on its two ends, and expands and shrinks by different amounts on different sides depending on the position of the Sun and the amount of exposure, we allow for fine adjustment in flight through ground control, by storing the offset between the X-ray and optical systems as a user-specified parameter in the radio uplink data stream. To adjust the alignment in flight, ground control simply specifies a different value for this parameter through the communication software running on ground station equipments.

This alignment process, as described above, is a somewhat qualitative one. We were comparing the appearances of images, instead of measuring any numeric quantity. A more quantitative align-



Figure 7.6: Actual (left) and simulated (right) images of the X-ray beam used for aligning the X-ray mirrors and detectors.

ment method would have been to measure the count rates registered on the focal planes as we vary the orientations of the optics. This method relies on the principle that the mirror response is highest on-axis. Thus, the optics orientation that yields the highest count rate on the focal plane must have the optical axis aligned with the line connecting the centre of each optics module and the centre of each focal plane. Yet, it also relies on knowing the efficiency of the detectors and its variation from pixel to pixel; when we carried out the alignment process, we had not yet successfully calibrated the detectors to obtain this information.

7.3 The observation run

After almost a decade of hardware development, HEFT was ready for launch in mid-May, 2005. The final payload weighed 4175 lbs, excluding ballast and crush pads. At 18:55 UTC on 2005- 5-18, a 39.57 million cubic-foot zero-pressure balloon carried the payload to the sky from the NSBF launch site in Fort Sumner, New Mexico. Table 7.1 lists the significant events during flight. Ascent through the Pfotzer Maximum of elevated cosmic-ray background at 20 km and to the final altitude of 39.0 km took 3.5 hours. During the ascent, we slowly increased the cathode bias on each detector, as the ambient temperature dropped with increasing altitude. Note that this was necessary to avoid saturating the detector-hybrids with a steady leakage current and with burst events at temperatures too high for our detectors. At the same time, we carried out in-flight calibrations of the altitude control system. At 21:25 UTC, the on-axis star tracker successfully locked on a single star, during daytime.

After reaching flight altitude at 22:30 UTC, we pointed the telescopes in the directions of the Crab Nebula and other known point-like sources of hard X-rays in the sky (Her X-1, Mkn 421 and Cyg X-1), in an attempt obtain X-ray detections of these bright hard X-ray sources, and thus to confirm the proper alignment of the X-ray telescopes with the star trackers and other components of the attitude control system. Due to insufficient preparation and a tight schedule, we did not have very good tools to reduce and analyse data in real-time during the flight. Although we made a hasty attempt on real-time aspect reconstruction, it proved not to be successful. This left us with real-time images only in detector coordinates, in which point sources were smeared out by any undesirable motion of the payload in the upper atmosphere. Instead of relying solely on image data, we measured the event rates when the telescopes were on- and off-target. Note that this was the opposite approach to what we had done for ground alignment using the X-ray generator (see Section 7.2.2.3). Using sky coordinates obtained from the star trackers, we divided the sky in the vicinity of our target into a raster grid centred at the nominal target coordinates. We pointed the telescopes at grid points 5-10' apart, for 2-10 minutes at a time, and measured the event rate from each patch of the sky. The idea was to find the sky coordinates that gave the highest count

Time		Event	Payload	Target
UTC	Local		altitude (km)	elevation
18:55	12:55	Launch, from Fort Sumner, New Mexico.	0.0	
21:25	15:25	On-axis star tracker tracked in single-star mode.	31.8	
22:30	16:30	At float; started searching for the Crab Nebula.	39.0	63° .7
23:47	17:47	White Sands Missile Range started jamming signals.	38.9	—
01:00	19:00	Resumed search for the Crab Nebula.	38.7	$33^\circ5$
02:30	20:30	Stopped search for the Crab Nebula;	38.2	$15^{\circ}_{\cdot}4$
		started search for Her X-1.		$22^{\circ}_{\cdot}2$
03:30	21:30	Started search for Mkn 421.	36.6	$74^{\circ}_{\cdot}8$
		(No one took record during the night.)		
11:10	05:10	First light: Cyg X-1.	37.5	87.9
11:18	05:18	Observation of GRS $1915 + 105$ started.	37.6	$65^\circ_{\cdot}2$
11:40	05:40	Observation of GRS $1915 + 105$ ended.	37.8	$62^{\circ}_{\cdot}7$
12:08	06:08	Observation of Cyg X-1 started.	39.0	$78^{\circ}_{\cdot}9$
13:00	07:00	Observation of Cyg X-1 ended.	39.3	$68^{\circ}_{\cdot}6$
13:04	07:04	Observation of $3C$ 454.3 started.	39.3	$67^{\circ}_{\cdot}1$
13:54	07:54	Observation of $3C 454.3$ ended.	39.4	$72^{\circ}_{-}5$
14:11	08:11	Observation of $3C$ 454.3 started.	39.4	73° 1
15:00	09:00	Observation of $3C$ 454.3 ended.	39.5	$70^{\circ}_{\cdot}4$
15:06	09:06	Observation of $3C$ 454.3 started.	39.5	69 ho7
15:42	09:42	Observation of $3C$ 454.3 ended.	39.5	$64^{\circ}_{\cdot}5$
15:57	09:57	Observation of $3C$ 454.3 started.	39.6	$61^\circ_{\cdot}9$
16:07	10:07	Observation of $3C$ 454.3 ended.	39.6	$60^\circ.1$
16:53	10:53	Observation of the Crab Nebula started.	39.7	$34^{\circ}_{\cdot}8$
17:10	11:10	Observation of the Crab Nebula ended.	39.7	$38^\circ_{\cdot}2$
17:39	11:39	Observation of the Crab Nebula started.	39.7	44°1
17:44	11:44	Observation of the Crab Nebula ended.	39.7	$45^{\circ}_{.}1$
18:10	12:10	Observation of the Crab Nebula started.	39.5	$50^\circ_{\cdot}3$
18:36	12:36	Observation of the Crab Nebula ended.	39.1	$55^{\circ}6$
18:39	12:39	Observation of X Per started.	39.1	$80^{\circ}_{\cdot}0$
19:00	13:00	Observation of X Per ended;	38.5	$83^{\circ}_{\cdot}9$
		telescope latched, electronics switched off.		
19:39	13:39	Termination of flight.		
20:20	14:20	payload landed 35 nmi ENE of Holbrook, Arizona.	0.0	

Table 7.1: Timeline of events during flight, 2005- 5-19.

rate, which would signify the presence of a celestial hard X-ray source. If the payload remained steady, we might also see images of a point source simultaneously on all three focal planes. Once we located the hard X-ray source, we would know the relative offset between the optical star trackers and the X-ray telescopes. We could then input this information into the attitude control system and automate subsequent observations. We expected this offset to vary only slightly during the flight, thus requiring only occasional readjustments.

It took us a long time to confirm our first X-ray detection, in part because this was the maiden flight of HEFT and we had no previous experience on aligning a hard X-ray telescope, and in part because we did not have very good tools at our disposal. To add to our problems, the nearby White Sands Missile Range jammed all radio signals in its vicinity, including ours, beginning at 23:47 UTC, five hours into the flight and 1.5 hours after HEFT was at float, while we were searching for the hard X-ray emission from the Crab Nebula. As a result, GPS signals were unavailable to the GPS sensors onboard, and attitude control was unable to locate the position of the payload; also, the real-time radio telemetry we received from the payload was completely scrambled. This incident cost us valuable flight time, during which we could neither command nor listen to the payload. With the help of NSBF personnel in communicating with White Sands, the situation was resolved in about one hour, and we resumed our search for the Crab in hard X-rays at 01:00 UTC. By 02:30 UTC, the Crab Nebula had set below 20° elevation, while our other calibration source of choice, Cyg X-1, was not yet observable. We attempted to search for the X-ray binary system Her X-1, but soon realised that its X-ray source was in a regular period of eclipse by its optical companion star. We continued our alignment effort with Mkn 421, a blazar, but we could not detect its X-ray emission either. Finally, at 11:10 UTC¹ on 2005-5-19, 16 hours after launch, we successfully imaged Cyg X-1. Using this information, we calculated the offset between the optical axes of the X-ray telescopes and of the on- and off-axis star trackers. Then, we continued the observations of our planned targets, now with enough aspect information to pinpoint our target fields exactly.

During the rest of the flight, we spent less than an hour (< 3.6 ks) each on three galactic compact objects: the Crab, GRS 1915+105 and X Per. Instead of these obvious targets, we took a chance and devoted the majority of our time (about three hours, 11 ks) to the blazar 3C 454.3 in a simultaneous target-of-opportunity (TOO) observation with *Swift*, *Chandra* and *INTEGRAL*. There was news that the blazar was in outburst at most wavelength/energy bands, including the report of an X-ray outburst on 2005- 5-11 (Remillard, 2005), just days prior to our launch.

Table 7.2 lists all the targets we observed during the maiden flight of HEFT. We chose our targets based on the scientific value of hard X-ray observations of each; we scheduled the observations based on their elevations in the sky. Here is a brief description of each of the targets:

¹This was the time recorded in the campaign notebook dedicated to focal plane matters, but it is uncertain whether this was the time our sole record keeper returned from his scheduled break, or if it was truly the time of first light.

Target	Species	Right Ascension (J2000)	$\begin{array}{c} { m Declination} \ ({ m J2000}) \end{array}$	${f Duration}\ ({ m minutes})$
Mkn 421	Active galactic nucleus (BL Lac)	$11^{ m h}04^{ m m}27^{ m s}31$	$+38^{\circ}12'31''_{\cdot}8$	
Her X-1	Eclipsing binary pulsar	$16^{ m h}57^{ m m}49^{ m s}83$	$+35^{\circ}20'32''_{\cdot}6$	
Cyg X-1	High mass X-ray binary	$19^{ m h}58^{ m m}21^{ m s}_{\cdot}68$	$+35^{\circ}12'05''_{\cdot}8$	52
GRS $1915 + 105$	High mass X-ray binary & microquasar	$19^{ m h}15^{ m m}11^{ m s}_{\cdot}60$	$+10^{\circ}56'44''_{\cdot}0$	22
3C 454.3	Active galactic nucleus (blazar)	$22^{ m h}53^{ m m}57^{ m s}75$	$+16^{\circ}08'53''_{\cdot}6$	
Crab	Pulsar & young supernova remnant	$05^{ m h}34^{ m m}31^{ m s}97$	$+22^{\circ}00'52''_{\cdot}1$	15
X Per	High mass X-ray binary	$03^{ m h}55^{ m m}23 .08$	$+31^{\circ}02'45''_{\cdot}0$	21

Table 7.2: Celestial targets observed during the HEFT campaign of 2005.

- The Crab Nebula, or the Messier object M 1, is a young supernova remnant, left behind from a supernova in 1054. It contains a 30 Hz pulsar at its centre; both the pulsar and the nebula emit broadband electromagnetic radiation from radio waves to gamma rays. The Crab Nebula is very well studied, with a well-defined power-law spectrum (see Equation 7.1) characteristic of synchrotron emission. Because of its brightness and its well-defined spectrum, it is the best candidate target for our in-flight calibration of flux measurements. Apart from being a good calibration source, we are also interested in investigating several scientifically interesting aspects of the Crab, such as its spatial/angular profile at hard X-ray energies, which provides information on pulsar-wind driven toroidal shocks, as explained in Section 1.1.5.
- Her X-1 is an eclipsing X-ray binary system of high interest to astronomers. It contains an X-ray pulsar with a pulsating period of 1.24s, while the X-ray intensity is modulated in a 35-day quasi-period. We observe Her X-1 in order to study its cyclotron resonance scattering feature (CRSF) around 41 keV. With the high spectral resolution of the HEFT detectors, compared with previous hard X-ray missions, we intend to resolve the profile of the CRSF absorption line in order to study the magnetic field structure close to the surface of the pulsar (see Section 1.1.4.
- Cyg X-1 is the brightest celestial source of hard X-rays. As such, it is an easy target for verifying that the payload functions properly. Cyg X-1 is a high-mass X-ray binary (HMXB) system with a compact object accreting from a normal star. The compact object is a stellar-mass black hole candidate. Due to the nature of accretion-powered emission, the X-ray intensity of Cyg X-1 fluctuates substantially and on all time-scales. Nevertheless, because it is a bright point source, it is a useful tool for us to measure the imaging performance (i.e., angular resolution) of the HEFT telescopes.
- GRS 1915+105 is another HMXB and a black hole candidate. Its emission is transient, but at times stronger than that from Cyg X-1. It is also known as a microquasar—although it is a stellar-mass system galactic in origin, it forms jets of radio emission on opposite sides of its X-ray emission, similar to the appearance of quasars.



Figure 7.7: Flight path during the first flight of HEFT.

- X Per, also known as 4U 0352+309, is another HMXB. Like Her X-1, its hard X-ray spectrum shows a CRSF absorption line, at 29 keV, well within the energy band observable by HEFT.
- Mkn 421 is a BL Lac-type active galactic nucleus (AGN) whose emission spans an impressive 14 magnitudes, from radio waves to gamma-rays.
- 3C 454.3 is a blazar (a subtype of AGNs) that was flaring in the months preceding the launch of HEFT. Specifically, the All Sky Monitor on *RXTE* detected the outburst in soft X-rays just days prior to our launch (Remillard, 2005). We managed to schedule a simultaneous target-of-opportunity observation of 3C 454.3 with the *Swift, Chandra* and *INTEGRAL* missions.

Throughout the flight, the payload drifted slowly westward. Around 24 hours after launch, the payload had entered Arizona and was gradually getting out of the range of our radio command signals. Although we had set up a down-range station in Holbrook, Arizona, it was not utilised at the end. At 19:00 UTC on 2005- 5-19, we stowed the telescope truss into the latched position; at 19:39 UTC, NSBF cut down the HEFT payload. It later landed at 20:20 UTC 35 nautical miles east northeast of Holbrook. Overall, HEFT stayed afloat for 24h45m. Figure 7.7 shows the path of the flight.

Unfortunately, the landing was far from perfect. For reasons unbeknownst to us, the latch that held the telescope truss in the stowed position disengaged at some point during descent. During landing, the truss was in a vertical position, with the focal-plane assembly at the bottom. As a result, the momentum of the descent and the weight of the payload bent the truss out of shape, crushed the kevlar dome, and dislodged the shield modules out of position. The strong impact tore the hybrid detectors from the motherboards and broke them into pieces. This unfortunate outcome prevented us from performing any post-flight recalibration of the hybrid detectors to compensate for our hasty pre-flight effort in the field.

7.4 Results from flight

At the time of this writing, we have not yet completed the analysis of all flight data. Thus, the following represents only a status report of our major findings to date.

7.4.1 General quality of data

Figure 7.8 shows the time profiles of the payload altitude, pointing elevation, event rate and focalplane temperature in flight. The big drop in altitude at around 02:30 UTC was due to nightfall, and the big rise circa 12:00 UTC was due to daybreak. The release of ballast cirCa 07:15 UTC also caused an abrupt rise in altitude. The event-rate profile includes only events with energy within 20-70 keV, the bandpass bound by atmospheric absorption on the low-energy side, and by the abrupt drop in reflectivity of the W/Si multilayer coatings on the high-energy side. In addition, we have filtered out invalid events and most of the noise events due to bursts of leakage current at individual pixels. From the event rate profile, we see clear elevation of counts during the time intervals when we observed Cyg X-1 and the Crab Nebula. Note that the relatively high event rate during the first two hours of flight was due to the payload crossing the Pfotzer Maximum at about 20 km. From the temperature profile, we note that the mostly metal payload formed a large thermal mass, preventing the focal plane from quickly cooling down to the expected operating temperature range of $[-15, +5]^{\circ}$ C. Once the temperature reached this range, it gradually fell at night and rose in the day, due to varying exposure to the Sun.

Compared to data taken in ground operation, the flight data exhibit the amplification of a number of systematics. Figure 7.4.1 shows histograms of events in two periodic time counters in flight and during ground calibration. As the arrival of X-rays is a truly random event, we expect a uniform event rate across all time bins in each of these histograms. From the ground calibration data (the dashed lines), we see an elevated rate at certain time bins, indicating the presence of systematics that we have yet to understand. In the flight data, we see a two- to three-time increase in the occurrence rate of these spurious events. Given the low (source and background) event rate in flight, the spurious events with sequence ID 8 account for as much as half of all events recorded in flight. In addition, the usual correlation between these two synchronous periodic time counters are sometimes lost, as shown in Figure 7.10. Because real events with sequence ID 8 are indistinguishable from the spurious events that dramatically increase the event rate, we have no choice but to filter out all



Figure 7.8: Time profiles of the payload altitude, pointing elevation, event rate and temperature in flight. The observation times of the various targets and the time of passage through the Pfotzer Maximum are indicated.



Figure 7.9: Histograms of two periodic time counters—the Sequence IDs (left) and the Time Since Reset (right)—of events recorded in flight (solid line) and during ground calibration (dashed line). Note the increased fraction of abnormal events in flight in both histograms.



Figure 7.10: Correlation between the two synchronous periodic time counters, Sequence ID (Figure 7.4.1, left panel) and Time Since Reset (Figure 7.9, right panel), in flight and during ground calibration. Each of these figures is a '3D histogram', where the coordinates represent the values of the two time counters (modulo their respective periods), and the darkness of each 'pixel' represents the rate of events recorded at those time values. Note that the usual one-to-one correlation between the two periodic time counters is broken in flight by spurious events caused by an unknown systematic problem.



Figure 7.11: Histograms of the number of pixels triggered per event in flight (solid lines) and during ground calibration (dashed lines). The event triggers are those computed by the software discriminator during data reduction (see Section 3.7.3 and Appendix C). We present the event counts as fractions of the total event counts in the observation run and the calibration run, respectively, as the totals differ significantly. There is a small but definite surplus of events triggering a large number of pixels in flight when compared to ground calibration.

events with sequence ID 8, including the real events. As a result, we lose 1/16 of source and true background event counts due to this systematic problem.

Figure 7.11 shows another difference between the flight and ground calibration data sets. In flight, we find an increase in events triggering large numbers of pixels. While the size of this increase is small, the presence of this difference is definite and consistent across all six flight detectors. We do not yet know the true cause of this difference, and further analysis is needed. However, the small fraction of events involved means that this difference has little effect on any conclusion we make from either set of data.

The mean livetimes of the X-ray detectors are 85.38%, 84.02% and 85.91% for Modules A, B and C, respectively.

134



Figure 7.12: Background spectrum at balloon altitude, as measured in flight (red) and simulated (green). The emission line at 122 keV is from the 57 Co calibration source onboard the payload. The emission features within 70–90 keV are due to lead fluorescence from the background shields.

7.4.2 In-flight background measurement

Figure 7.12 shows our measured background spectrum at balloon altitude. This spectrum contains events recorded from most of the flight, excluding the initial ascent to the final altitude and time periods when the telescopes were on a known target. The measured background agrees well with our simulation results.

The plastic scintillator shield also functioned in flight per specification. Figure 7.13 shows the detector livetime and event rates from the background shields of all three telescope modules as functions of time. The shields produced about 1400 veto pulses/module/s during passage of the Pfotzer Maximum, and 600 veto pulses/module/s at flight altitude. The mean rates of coincidence between (rejected) X-ray events and shield veto pulses are 5.98, 5.53 and 6.02 events/module/s in Modules A, B and C, respectively, compared to typical source event rates of less than one event/module/s.

7.4.3 Cyg X-1

We observed Cyg X-1 for about 52 minutes with full aspect information. During this period, Cyg X-1 was at elevations above 60° , where the carrying balloon hides the X-ray field of view from the on-axis star tracker. Thus, attitude control for this observation relied on the off-axis star tracker, which tracked stars in a field about -30° in elevation relative to Cxg X-1. Note that the azimuth coordinate of the off-axis star tracker field translated only partially to the azimuth coordinate of the







Figure 7.14: Left: image of Cyg X-1, measured by Modules A and B of HEFT on 2005- 5-19 for 42 minutes. Right: Projected position of Cyg X-1 on the focal plane of telescope Module B over the same observation period.

X-ray field, with a $\cos 30^{\circ}$ scaling, and real-time attitude control with the off-axis star tracker is less accurate compared to on-axis star tracking. While the second axis of Gyroscope 1 measured purely the motion of the telescope truss in the azimuth of the X-ray field of view (see Section 2.6.2), this information was unfortunately not recorded onboard nor transmitted down to the ground station. As a result, we had to infer from indirect measurements only to deduce the azimuth coordinates during the Cyg X-1 observation (and all other observations made above 60° elevation). This greatly complicated the aspect reconstruction and increased the uncertainty in the azimuth coordinate. Kruse Madsen (2007, Section 3.3) describes the aspect reconstruction in detail.

Figure 7.14 shows a sky image of Cyg X-1 on the left, constructed with events from two of the three telescope modules (A and B) in a 42-minute observation period, and the projected position of Cyg X-1 on the focal plane of Module B on the right. Because Cyg X-1 is a point source, we measure the point-spread function (PSF) of the telescopes from this observation. We find the half-power diameter (HPD) of the PSF to be 96". Figure 7.15 shows the details of this measurement. Based on ground calibration and measurement of the X-ray mirror modules using 8 keV X-rays and LVDT mechanical scanning, we find the on-axis mirror response to have an HPD of 81". We believe the difference between the in-flight measurement and the on-axis response is due to the target (Cyg X-1) being off-axis, as the target moves across the focal plane substantially during the 52 minutes of observation, and the off-axis mirror response is worse than the on-axis response.

Figure 7.16 shows our measured spectrum of Cyg X-1.



Figure 7.15: Measurement of the point-spread function (PSF) of the HEFT telescopes through the observation of Cyg X-1 and ground calibration. Data from two of the three modules (A and B) are combined in these figures. The top and bottom panels show fits to the PSF; the middle panel shows the residuals of the fit. Flight data is in blue; the dotted line at the top is the background. The red curves show simulated PSF based on ground X-ray and mechanical measurements. They include an 8" aspect error and the off-axis mirror response. The green curves show the on-axis response with no aspect error. These calculations account for all events out to a 6'0 radius, which includes 99% of the simulated photons.



Figure 7.16: Spectrum of Cyg X-1, measured by Modules A and B HEFT on 2005- 5-19 for 52 minutes.

7.4.4 The Crab nebula

In contrast with Cyg X-1, we observed the Crab Nebula at elevations below 60°, without the carrying balloon obscuring the view of the on-axis star tracker. Thus, the on-axis star tracker tracked the same patch of sky as the X-ray field of view during this observation, providing pointing information referenced entirely to celestial coordinates, without dependence on the moving gondola frame of reference. This greatly reduces the complexity of aspect reconstruction.

Figure 7.17 shows a sky image of the Crab Nebula on the left, measured by HEFT over a 15minute integration period, and the projected position of the target on the focal plane of Module A. We are still analysing this data and trying to improve the quality of the Crab image by eliminating systematics due to inhomogeneous detector response across the focal plane, and by experimenting with new aspect reconstruction schemes. Thus, while there are tantalizing signs suggesting that we might be detecting the extended structure of the nebula at hard X-ray energies, we cannot yet confirm or refute this hypothesis at the writing of this thesis.

In addition to the work reported here, we have yet to analyse a large portion of the flight data. This includes the many hours of data taken prior to our confirmation of detection of Cyg X-1. Because the offset between the optical cameras and X-ray telescopes was unavailable for this time period,



Figure 7.17: Left: sky image of the Crab Nebula, measured by HEFT on 2005- 5-19 for 15 minutes. Right: projected position of the target on the focal plane of telescope Module A.

aspect reconstruction is not a straight forward task, and in fact, may or may not be possible. For instance, we do not know how much time we had on each target, and how much time we had in offset pointings. In addition, we are attempting to salvage events from many detector pixels that require special attention. Most of these pixels have either unusually low gain or excessive charge-sharing. Because of such peculiarities, they require time-consuming manual inspection. As we continue to improve our methods in both aspect reconstruction and detector efficiency correction, we cannot rule out the possibility of extracting new results from the flight data. However, as of now, there was no clear detection for the remaining targets: Her X-1, X Per, GRS 1915+105, Mkn 421 and 3C 454.3.

Chapter 8

A scientific case: study of the inverse Compton scattering of cosmic background photons by relativistic electrons

Although the HEFT payload performed within specifications during its 2005 flight, the excessive time we took to calibrate the pointing in flight left us with insufficient observation time to produce any new science result thus far. To exemplify the type of astronomy that can be enriched by HEFT and focusing telescopes of hard X-rays in general, I present here a study of nonthermal X-ray emission from two relic radio sources, which I mentioned in Section 1.1.3, in the context of celestial targets for HEFT. This study was done with data from the *XMM-Newton* observatory, which operates in the soft X-ray band (0.2–10.0 keV). As I describe below, thermal contamination and current limits in the capabilities of low-frequency radio telescopes prevent us from making definite conclusions at present on the presence and nature of X-ray emission from relic radio sources. Observations in hard X-rays with HEFT or other focusing telescopes will eliminate these problems, and greatly enlarge the sample of observable targets to encompass radio haloes in clusters of galaxies and other relics, most of which are covered in thermal emission in soft X-rays.

The remainder of this chapter appears as the journal article 'Constraints on the average magnetic field strength of relic radio sources 0917+75 and 1401-33 from XMM-Newton observations' by C. M. Hubert Chen, D. E. Harris, Fiona A. Harrison and Peter H. Mao in the Monthly Notices of the Royal Astronomical Society.¹

¹Currently in print.

8.1 Introduction

In a minority of clusters of galaxies, there are isolated diffuse sources of radio emission with steep spectra and without definite associations with an optical galaxy counterpart. These sources are classified as haloes or relics, depending on their location and size: haloes reside at the centre of a cluster of galaxies, and are usually circular and not known to be highly polarized. In contrast, relics are found in cluster outskirts, and are often elongated and linearly polarized. Given the general phenomenological designation, the term 'relic' actually refers to a few different and probably distinct types of sources, including old radio lobes from dead radio galaxies and regions where particles are reaccelerated in cluster mergers (Kempner et al., 2004; Giovannini and Feretti, 2004). Statistics on relics are growing; to date, over 30 relics have been discovered.

The various types of relics share common observational characteristics, including a power-law radio spectrum, indicative of synchrotron emission by relativistic particles in a magnetic field. The radio spectrum alone constrains only the product of the magnetic field and particle density. The standard approach to estimating the magnetic field strength and particle density individually is to assume minimum energy or equipartition of energy densities between the two components. A direct observational approach is to measure the nonthermal X-ray spectrum due to inverse Compton scattering of photons from the cosmic microwave background (IC/CMB) by relativistic particles in the relic (Harris and Grindlay, 1979). This provides a measurement of the particle density, and thus decouples the contributions of magnetic field and particles to the synchrotron flux.

While IC/CMB measurements have been made on other, stronger radio sources of synchrotron emission (e.g., Croston et al., 2005), IC/CMB measurements on relics remain scarce. The difficulty stems from the weak and diffuse nature of relic emission, as well as contamination by thermal emission, often from a nearby cluster. To date, there are only three claims of IC/CMB detections in relic radio sources (Bagchi et al., 1998; Fusco-Femiano et al., 2003; Kempner et al., 2003), but the statistics are poor, and later studies refute two of these detections (Durret et al., 2005; Henriksen et al., 2004). For a number of other relics, nondetections have yielded lower limits on the magnetic field strength. Further X-ray measurements with sensitivity improvement over past studies are interesting, as they strengthen the constraints on the magnetic field and particle content in relics.

In this paper we present XMM-Newton observations of two relics, 0917+75 and 1401-33. Our X-ray observations are the most sensitive to date for these sources. Combined with previous radio observations, we derive constraints on their magnetic field and particle content. §8.2 describes the X-ray and radio data, §8.3 details our analysis, and §8.3.2.3 presents the numerical results. We discuss each source individually in §8.4, and give a summary of our findings in §8.5. We adopt the currently accepted cosmology with $H_0 = 70 \text{ km s}^{-1} \text{Mpc}^{-1}$, $q_0 = -0.55$ and $\Omega_{\Lambda} = 0.7$. At z = 0.125 and 0.0136, respectively, the distances to 0917+75 and 1401-33 are 502 Mpc and 57.8 Mpc, and the

Table 8.1: Source and X	MM observation	parameters.	
Target	0917 + 75	1401 - 33	
Right Ascension $(J2000)$	$09^{h}22^{m}11.40$	$14^{h}04^{m}16^{s}.70$	
Declination $(J2000)$	$74^{\circ}59'31''00$	$-34^{\circ}02'22''_{\cdot}00$	
Angular Size	$4' \times 8'$	$9' \times 20'$	
Galactic H I Column Density, n_H	$2.09 imes 10^{20}$	$5.44 imes 10^{20}$	cm^{-2}
Cosmological Redshift, z	0.125	0.0136	
XMM Observation Date	$2002 {\rm \ Mar\ } 20$	$2002 {\rm Feb} 14$	
	07:00:21.0 UT	$07:59:00.0 \ \mathrm{UT}$	
XMM Observation Length	26744	20433	sec

linear scales are 134 kpc/arcmin and 16.7 kpc/arcmin.

8.2 The observations

We observed each of the targets, 0917+75 and 1401-33, in a ~ 20 ks pointing with XMM-Newton. We obtained radio images of the targets at 1.4 GHz from the NRAO VLA Sky Survey (NVSS) catalogue (Condon et al., 1998), via the *SkyView* facility (McGlynn et al., 1998).² We also used published data from Harris et al. (1993) and Goss et al. (1987). Table 8.1 lists some properties of the targets and the parameters of the observations.

8.3 Data analysis

We first reprocessed the X-ray data sets using the routines *emproc* and *epproc* in the XMM-Newton Science Analysis Software (SAS), version 6.5.0. Then, we filtered the processed data based on temporal, spectral, and spatial criteria, as described below.

8.3.1 Event filtering

For temporal filtering, we generated good-time intervals (GTIs) based on the double-filtering technique described in Appendix A of Nevalainen et al. (2005): we inspected light curves of high-energy events (above 9.5 keV for MOS and 10.0 keV for PN) of 1-ks resolution, and screened out time intervals where the high-energy count rate at any of the three detectors exceeded 120 per cent of the mean level. We also inspected light curves of events within 1–5 keV at the periphery of the field of view (12 arcmin < r < 15 arcmin, r being the distance from the nominal pointing direction), and screened out time intervals where any of the count rates, summed over the periphery annulus, exceeded mean+0.028 counts s⁻¹ for MOS and mean+0.056 counts s⁻¹ for PN. The final GTI for each observation is the intersection of the GTIs from the high- and low-energy filters. It is about

² http://skyview.gsfc.nasa.gov/



Figure 8.1: XMM-Newton (a) and NVSS (b) images of 0917+75. The two images have the same spatial extent. The source (centre) and background (annulus) regions are indicated by solid lines in (a); the background region is also shown in (b) by dotted lines. X-ray point sources detected within the XMM-Newton field of view are marked by solid lines in (b), with the size of each region proportional to the X-ray flux of the point source. The boundary of the source region is identical to the mean+5 σ contour in the NVSS image, excluding all overlapping X-ray point source regions. Image (a) is made by (1) scaling the filtered MOS (0.2–3.7 keV) and PN (0.2–2.3 keV) images, according to each instrument's response to the same source model and flux, (2) correcting the scaled images for exposure and combining them with the SAS routine emosaic, and (3) smoothing the combined image with a Gaussian function of width $\sigma = 4$ arcsec, about the size of the instrument's point spread function.

89 per cent the length of the entire observation for 0917+75, and about 80 per cent for 1401-33. Table 8.2 lists the temporal filtering parameters in detail.

We restricted the region for photon extraction to coincide with the radio source extent, and we filtered out regions containing point source detections. From the NVSS radio image of each target, we made a mask by selecting only pixels in the vicinity of the target whose flux densities are at least 5σ above the background level in each image. To mask out X-ray point sources in each field, we took the three source lists (one from each detector on *XMM-Newton*) provided with the Pipeline Products, and ran them through the SAS routine *region*. This produced three FITS region files for each field, specifying regions that excluded all the detected point sources, with the size of each excluded region determined by the level of flux for that particular source. We combined the radio mask and X-ray point source regions to form the spatial filtering expression for each of the targets. These masks are shown in Figs. 8.1 and 8.2.

We determined the optimal energy band for signal extraction by comparing (1) the expected source flux density at equipartition and (2) the average background for *XMM-Newton* at various

	Table 8.2: N ₁	umerical results	from the X-ray	· observations.			
Target		0917 + 75			1401 - 33		
Detector	MOS1	MOS2	PN	MOS1	MOS2	PN	
Observation Mode	Full Window	Full Window	Extended Full Window	Full Window	Full Window	Extended Full Window	
Filter	Medium	Medium	Medium	Medium	Medium	Medium	
Hard-band Count Rate	0.188 - 0.378	0.207 - 0.383	0.770 - 1.65	0.152 - 0.208	0.151 - 0.237	0.560 - 0.805	$counts s^{-1}$
Mean	0.247	0.259	0.962	0.184	0.188	0.671	$counts s^{-1}$
Soft-band Annulus Count Rate	0.119 - 0.202	0.118 - 0.190	0.210 - 0.340	0.166 - 0.266	0.177 - 0.251	0.368 - 0.516	$counts s^{-1}$
Mean	0.148	0.150	0.262	0.214	0.216	0.446	$counts s^{-1}$
Observation Length	27825	27824	23792	19327	$19\ 326$	15291	sec
Flare-filtered Length	$24\ 715$	24720	20693	15483	15484	12082	sec
Good-time Fraction	%68	89%	87%	80%	80%	29%	
Optimal Energy Band	0.2 - 3.7	0.2 - 3.7	0.2 - 2.3	0.2 - 6.5	0.2 - 6.5	0.2 - 4.1	keV
Source Region							
Event Count in This Study, Optimal Band	444	405	1441	488	529	1244	events
in Blank-sky Data, Optimal Band	÷	÷	:	10122	9840	24308	events
in This Study, Hard Band	:	÷	÷	40	42	120	events
in Blank-sky Data, Hard Band	:	÷	÷	1403	1432	6698	events
Exposure Map, Median	$24\ 054.6$	23764.0	17680.7	14218.5	14594.4	10030.4	sec
Mean	23847.3	23061.4	16870.2	14064.4	$14\ 295.8$	9850.77	sec
Number of 4.35" Pixels	1856	1856	1785	1406	1406	1384	pixels
Event $Rate^{a}$	0.0186	0.0176	0.0854	0.0142	0.0168	0.0821	$counts s^{-1}$
Background Region							
Event Count in This Study, Optimal Band	352	383	1196	739	868	1840	events
in Blank-sky Data, Optimal Band	•	:	:	16422	17797	41911	events
in This Study, Hard Band	:	:	:	63	80	230	events
in Blank-sky Data, Hard Band	•	÷	:	3159	3751	11558	events
Exposure Map, Median	19570.7	19612.4	$14\ 241.8$	6043.56	6656.07	4656.47	sec
Mean	19268.2	19505.8	13634.9	5976.04	6701.31	4615.42	sec
Number of 4.35" Pixels	1881	1881	1763	3128	3208	3096	\mathbf{pixels}
Event Rate ^{<i>a</i>} , Scaled to Source Region	0.0180	0.0194	0.0888	0.0310	0.0319	0.0974	
	± 0.0016	± 0.0015	± 0.0016	± 0.0043	± 0.0035	± 0.0129	$counts s^{-1}$
Source Count Rate, 3 <i>a</i> Upper Limit (MOS Co-added)	$\leq 4.70 \times 10^{-3}$ < 3.86	$\leq 4.59 \times 10^{-3} \times 10^{-3}$	$\leq 4.78 \times 10^{-2}$	$\leq 1.28 \times 10^{-2}$ < 1.11	$\leq 1.05 \times 10^{-2} \times 10^{-2}$	$\leq 3.86 \times 10^{-2}$	counts s ⁻¹ counts s ⁻¹
	I			1			

^{*a*}Event rates are blank-sky substracted for 1401–33 but not for 0917+75. See text for details.

145



Figure 8.2: XMM-Newton (a) and NVSS (b) images of 1401-33. See Fig. 8.1 for a description of the various features. In addition to the source and background regions at the centre and top of the images, respectively, we also mark the extent of the radio emission at 330 MHz, according to fig. 1 in Subrahmanyan et al. (2003). Note that the source and background regions here are equidistant from the bright galaxy NGC 5419 at the centre of the Abell S753 cluster, about 10 arcmin to the west (right) of the source region. Image (a) combines the filtered MOS images of 0.2-6.5 keV and the filtered PN image of 0.2-4.1 keV.

energies. To estimate the source flux density at equipartition, we followed the analytical formulation in Harris and Grindlay (1979) and deduced a power-law model for the IC/CMB emission from each source, with model parameters based on data from past radio observations (Harris et al., 1993). Tables 8.1 and 8.3 lists these parameters. To estimate the background, we took blank sky event files obtained from the web site of the XMM-Newton Science Operation Centre³ as representative background measurements, and computed a background spectrum from the files. Carter and Read (2007) described these blank sky data sets in detail. With these source and background spectra, we selected an energy band $[E_{min}, E_{max}]$ and calculated the signal-to-noise ratio (S/N) associated with the band. We then varied E_{min} and E_{max} at 0.1-keV increments until we obtained the energy band with the maximum S/N. Using this method, we arrived at the optimal MOS bands 0.2–3.7 keV for 0917+75 and 0.2–6.5 keV for 1401–33. Table 8.2 shows the selected energy bands for PN as well.

In addition to the GTIs and energy bands obtained as described above, we further applied the following event filtering expressions: (FLAG == 0 && PATTERN <= 12) for MOS and (FLAG == 0 && PATTERN <= 4) for PN. These expressions retained only events triggering 1 or 2 pixels, and only those flagged as valid by the standard processing routines. From the filtered event list, we generated images and exposure maps. We then used the masks and regions described above to select the region

 $^{^{3}\,\}rm http://xmm.vilspa.esa.es/external/xmm_sw_cal/background/blank_sky.shtml$



Figure 8.3: Proton flux and ratio of oxygen charge states O^{7+}/O^{6+} in the solar wind, measured by *ACE* during the time of our observations. The two horizontal lines in each panel indicate the mean and 90-percentile levels of each quantity over the first 100 days of 2002. The two vertical lines in each panel delimit the period of the *XMM-Newton* observation. The 64 s-averaged proton flux data are from the SWEPAM instrument on *ACE*, while the hourly charge state ratios are from the SWICS/SWIMS instrument. Note that the break in the O^{7+}/O^{6+} curve at the time of our observation of 0917+75 was due to missing data. Yet, from the values before and after the break, one can infer that the O^{7+}/O^{6+} ratio during the observation is probably higher than the 90-percentile. There is a ~1 hour travel time by the solar wind from the L1 point, the location of *ACE*, to the Earth. The time shown here is the time measured by *ACE*.

of sky corresponding to each target. The total counts divided by the mean of the exposure map in the selected region gave the event rate.

8.3.2 Background determination and flux upper limit analysis

The two observations in this paper pose very different challenges for background determination. We discuss our approach for each observation separately in the following.

8.3.2.1 0917+75

For 0917+75, the time of our observation unfortunately coincided with a coronal mass ejection (CME) from the Sun. Snowden et al. (2004) reported an enhancement in the XMM-Newton background, especially at 0.5–1.0 keV, concurrent with an enhancement in the solar wind measured by the Advanced Composition Explorer (ACE) and other monitoring spacecraft. When we obtained data from ACE^4 and inspected the light curves at the time of our observations (Fig. 8.3), we found that our observation of 0917+75 coincided with an episode of significant enhancement of the O⁷⁺/O⁶⁺ ratio in the solar wind, characteristic of a CME (R. C. Ogliore 2006, private communication).

To assess the impact of this event on the background level in our observation, we compared the event rate in our data set with the rates in two standard blank-sky data sets independent of our

⁴ http://www.srl.caltech.edu/ace/ASC/



Figure 8.4: Spectra of 0917+75 and of the blank sky data. The solid line shows the spectrum of 0917+75 from MOS2, convolved with a 100 eV-wide boxcar. The dashed and dash-dotted lines show the blank sky spectra (also convolved with a 100 eV-wide boxcar) from Nevalainen et al. (2005) and from Carter and Read (2007), respectively. Both blank sky spectra contain events extracted from the same detector coordinates of MOS2 as the source region in this study, and scaled to the spectrum of 0917+75 by equating the GTI-filtered livetimes. Data from Carter and Read (2007) is also refiltered with (FLAG == 0) for consistent comparison. Note the excess of counts at discrete 'spectral lines' at 0.5-1.0 keV when compared to the data from Nevalainen et al. (2005), and the overall deficit of counts when compared to Carter and Read (2007) (see text for an explanation). The count rates displayed here are livetime corrected.

observation: one set was from Nevalainen et al. (2005) and the other from Carter and Read (2007). We chose these two sets because we employed the flare-filtering recipe of Nevalainen et al. (2005) in this study, while Carter and Read (2007) provides a much larger data set, with better event statistics. We refiltered the Carter and Read data set with the (FLAG == 0) expression so that comparison of the three data sets is consistent. However, we have not refiltered the Carter and Read data set for flares using the recipe in Nevalainen et al. (2005), as its full-view hard-band event rates (0.26)and $2.7 \text{ counts s}^{-1}$ for MOS and PN, respectively) are above the range deemed consistent with the sample of Nevalainen et al. (2005), making the hard-band filtering method inapplicable. To minimize the variability of any other systematic parameters (such as vignetting), we applied the masks and regions that we obtained for the source regions of 0917+75 in this blank-sky measurement, thus selecting the same detector pixels. We computed count rates at various energies from the filtered blank sky event files, and compared them with our source-region count rates for 0917+75, both without exposure map correction. Fig. 8.4 shows spectra of 0917+75 and of the blank sky data thus obtained. It shows that the continuum level in our observation of 0917+75 is consistent with that from Nevalainen et al. (2005), but lower than the level from Carter and Read (2007). We believe the difference between the two blank sky data sets is due to different levels of flare filtering. As we employed the more stringent double-GTI filtering recipe of Nevalainen et al. (2005) in this study, our spectrum of 0917+75 should be directly comparable with a spectrum of the same detector region made from their data set. We note that the CME during our observation explains the excess line emission (relative to the continuum) at multiple energies below 1.0 keV in the spectrum of 0917+75, which is not seen in either blank sky spectrum.

To further investigate this background discrepancy, we looked at light curves generated from our observation of 0917+75. Fig. 8.5 shows the light curves of all events above about 10 keV from the entire field of view, which we generated to determine the high-energy GTIs (see §8.3.1). This figure shows that the count rates over the entire field of view above about 10 keV were above the 'acceptable' range of high-energy particle background rate reported by Nevalainen et al. (2005) for the majority of our observation period, and for all three detectors. This is evidence that our entire observation was plagued with low-amplitude long-duration flares, and that residual flares still lurk in the background after our double-GTI filtering. Thus, these light curves are consistent with the ACE data, suggesting an elevated background.

Because our observation was contaminated by flares, count rates from blank-sky data sets do not correctly reflect the nature of the background in our data. Besides, according to Fig. 8.4, the source rate in our data does not exceed the blank-sky count rates. Both of these reasons make it impossible for us to use blank-sky data for quantitative background subtraction. Thus, we measured the background level and its variation (for setting a flux limit in the case of non-detection) solely from multiple local background regions. We obtained the local background regions using a similar approach as we did for the source regions, applying temporal and spectral filters described in §8.3.1. For the spatial filters, we inspected the radio image of 0917+75 from NVSS, as well as X-ray images from the Pipeline Products, to find regions of sky within the XMM-Newton field of view but securely outside the target, as seen in radio frequencies, and where there are few detectable X-ray point sources. We selected an annulus immediately outside the target, with area similar to the source region. This annulus is inscribed within the central CCDs of MOS1 and MOS2, which simplifies the comparison between the source and background regions, as potentially different background levels at different CCDs (Pradas and Kerp, 2005) become irrelevant for the two MOS modules. With the background region thus selected, we excluded X-ray point sources from the region in the same way as was done for the source region. Fig. 8.1 shows the resulting background annulus. We then applied this local background region to the filtered images and exposure maps described in $\S8.3.1$, and obtained the total background count, mean exposure, and background count rate scaled to the total exposure in the source region. To assess the variation of the background level in our observation, we obtained additional background regions in the vicinity of the annulus and scaled the event rates within them in a similar fashion, and calculated the standard deviation σ of scaled event rates in this sample.

To assess the amount of degradation in our data due to residual flares, we calculated the count rates of individual pointings that make up the blank sky data set of Nevalainen et al. (2005). We



Figure 8.5: Light curves of high-energy events (left) and of events at the periphery of the field of view (right) during the observation of 0917+75. The three rows of graphs show, from top to bottom, events from MOS1, MOS2 and PN, respectively. The abscissae show the time since the start of the observation, in seconds. The ordinates show the count rate, in counts s^{-1} . The three solid lines in each graph on the left indicate the mean and (mean±20%) high-energy count rates. The long-dashed lines show the same quantities with the median. The two short-dashed lines delimit the 'normal' range of count rates reported by Nevalainen et al. (2005), which is exceeded during most of this observation. On the right, the solid (long-dashed) lines delimit the acceptable range of count rates computed with the median), according to Nevalainen et al. (2005). The short-dashed lines delimit the 'normal' range observed by Nevalainen et al. (2005); only the top line is seen in the top two graphs.

then compared the background variation (over space) in our data with the variation of the blanksky count rates (over time). We found the variation in our background to be greater than that in the blank sky data by factors of 1.26, 1.27 and 2.05 for the MOS1, MOS2 and and PN detectors, respectively. The flux limit we set in this study is affected accordingly.

$8.3.2.2 \quad 1401 - 33$

Compared to 0917+75, the nature of the background in our observation of 1401-33 is quite different. The observation was made at a time of average solar wind activity, as shown by ACE data during this period (see Fig. 8.3); the overall count rates above about 10 keV are also well within the accepted quiescent levels. However, an X-ray halo of thermal emission centred at the neighbouring bright galaxy NGC 5419 about 10 arcmin away west of northwest (Subrahmanyan et al., 2003) dominates the local background, producing excess flux across the entire field of view when compared to blank sky data. This thermal emission is so strong that spectral lines caused by the solar wind, which are normally present below 1.0 keV, are not observable. To assess the amount of thermal emission from the halo in our source region for 1401-33, we made the assumption that the thermal emission is spherically symmetric about the centre of the cluster (i.e., NGC 5419), and chose a local background region equidistant from the centre of the thermal emission as the source region. Our assumption is based on a previous X-ray observation of the cluster with the ROSAT PSPC (in which IC/CMB emission from the relic was not detected); fig. 8 in Subrahmanyan et al. (2003) is an X-ray image from the observation, showing rotationally symmetric contours about NGC 5419. The regular symmetric shape of the cluster justifies our choice of the background region. In addition, Subrahmanyan et al. (2003) also revealed that a large portion of the XMM-Newton field of view in our observation overlaps with the extent of the radio source at 330 MHz. Thus, we chose a circular region just to the north of the target, instead of the usual annulus surrounding the source region. Fig. 8.2 shows the resulting background region, with the extent of the 330-MHz emission for reference. To calculate the local background event rate and its variation, we applied the background region to the filtered image and exposure map, and obtained additional background regions, following the same recipe as outlined above for 0917+75.

Because our background region for 1401-33 is not close to the centre of the telescopes' field of view, vignetting is a potential issue. Pradas and Kerp (2005) reported an overcorrection of vignetting towards the rim of each detector by the routine *eexpmap* that generates exposure maps in SAS versions 5.3.3 and 5.4.1. To avoid overlapping with the extent of the radio emission of 1401-33 at 330 MHz, we have no choice but to place our background region at the periphery of the *XMM-Newton* field of view. Thus, the overcorrection of vignetting affects our analysis, potentially overestimating the count rate in the background region. On the other hand, the particle background that dominates at high energies is not vignetted. To correct for vignetting of the local background properly, we followed a two-step background subtraction method described in Appendix A of Arnaud et al. (2002): We calculated event rates in the blank-sky data from Nevalainen et al. (2005) at pixels in the source and background regions and within the optimal energies in this study. We scaled these source- and background-region event rates according to the ratio of hard-band rates in our data and in the blank-sky data. We then subtracted the scaled rates from the event rates in our data to remove the particle background from both the source and background regions. Having eliminated the non-vignetted component, we then scaled the event rate in the background region to the source region according to the ratio of summed exposure in each region, for vignetting correction. Finally, we arrived at the source event rate by substracting the scaled background-region rate from the source-region rate. Table 8.2 shows the relevant numbers in this analysis.

8.3.2.3 Flux and field limits

Table 8.2 summarizes our quantitative results. In this paper, we discuss only results from the MOS detectors, as the MOS and PN results yield the same conclusions, and the MOS results provide more stringent limits. For 0917+75, the event rates within the source region in the optimal band of 0.2–3.7 keV are 0.0186 and 0.0176 counts s⁻¹ for MOS1 and MOS2, respectively, while the corresponding area-scaled event rates in the background regions are (0.0180 ± 0.0016) and (0.0194 ± 0.0015) counts s⁻¹, respectively. For 1401–33, the blank-sky substracted event rates in the on-axis source region $(0.0142 \text{ and } 0.0168 \text{ counts s}^{-1}$, respectively) are only a fraction of the blank-sky subtracted and area-scaled rates in the off-axis background region [(0.0310 ± 0.0043) and (0.0319 ± 0.0035) counts s⁻¹, respectively]. These numbers, together with the appearance of the X-ray images in Figs. 8.1 and 8.2, indicate that there is not a detection of the IC/CMB emission from either source.

We converted the measured event rate to the unabsorbed energy flux with PIMMS (Mukai, 1993), assuming a power-law spectrum and accounting for redshift and galactic absorption. We set the X-ray spectral index to be the same as the radio spectral index (Harris and Grindlay, 1979), assuming any X-ray emission to be IC/CMB in origin. We took the spectral indices ($\alpha = 1.0$ for 0917+75 and 1.4 for 1401-33) and redshifts (z = 0.125 and 0.0136, respectively) from previous radio measurements, and the galactic H_I column densities (2.09×10^{20} and $5.44 \times 10^{20} cm^{-2}$) from the FTOOLS programme nH.

To assess the maximum source flux, we adopted three times the r.m.s. variation in the local background as an upper limit. When converted to energy flux at 0.3-10.0 keV, the $+3\sigma$ upper limits are $5.22 \times 10^{-14} \text{ erg s}^{-1} \text{ cm}^{-2}$ for 0917+75 and $1.47 \times 10^{-13} \text{ erg s}^{-1} \text{ cm}^{-2}$ for 1401-33. The measured flux in the source region of 0917+75 is consistent with this 3σ limit, while the source flux for 1401-33 is about 4σ below the measured mean background. Using Equation 11 in Harris and Grindlay (1979), and radio measurements reported in Harris et al. (1993) and Goss et al. (1987),

we obtained 3σ lower limits on the magnetic field strength at each source; they are 0.81 μ G for 0917+75 and 2.2 μ G for 1401-33. For comparison, the field strength obtained with the 'classical' equipartition formula (with a low-frequency cutoff of the synchrotron emission at $\nu = 10$ MHz) are 0.63 μ G and 1.3 μ G for 0917+75 and 1401-33, respectively. Table 8.3 provides further details on these calculations. Alternatively, using a revised equipartition formula from Brunetti et al. (1997, Equation A4), with the low-energy electron population cutoff at $\gamma_{\rm min} = 1069\sqrt{E_{\rm x,min}/\text{keV}} = 500$ (Harris and Grindlay, 1979) to match the energy range of XMM-Newton, and with $D(\delta) \sim 1$, the field strengths are 0.79 μ G and 1.6 μ G, respectively. We note, however, that these values are very sensitive to the choice of $\gamma_{\rm min}$, whose true value is unknown without better knowledge of the low-energy synchrotron spectrum.

8.4 Discussion

In this section, we first discuss each target separately, and then comment on the flux and field limits.

$8.4.1 \quad 0917 + 75$

Dewdney et al. (1991) originally thought that 0917+75 was the radio halo of an uncatalogued cluster of galaxies, but Harris et al. (1993) recategorized it as a relic radio galaxy on the basis of its high polarization. It is apparently associated with two galaxies of redshift z = 0.125 within the Rood 27 supercluster, whose members include the Abell clusters A762, A786, and A787. Harris et al. (1993) reported a spectral flux density of $S_r(\nu_r) = (1.2 \pm 0.2)$ Jy at $\nu_r = 151$ MHz, and an energy spectral index of $\alpha = 1.0$ at low frequencies, with a spectral break close to 150 MHz. The spectrum flattens to $\alpha = 0.6$ between 325 and 1500 MHz, and steepens again above that. This complex spectral shape could be the superposition of synchrotron emission from two populations of relativistic electrons; it would be the lower-energy population (478 < γ < 3380, 1 MHz < ν_r < 48 MHz) that is responsible for IC/CMB emission in the energy range of XMM-Newton (0.2–10.0 keV). Unfortunately, the spatial distribution of the lower-energy population is as yet undetermined (below 300 MHz), due to limited angular resolution in low-frequency measurements in the past. In this study, we have made the assumption that the IC/CMB emitting electrons have the same spatial extent as those in the higherenergy population, which was given by the NVSS radio maps at 1.4 GHz.

Harris et al. (1995) observed 0917+75 in a 15.8-ks *ROSAT* PSPC observation (with the same assumption on the spatial extent of the IC/CMB electrons). They found the source isolated in X-ray, and not contaminated by thermal emission. They placed a 2σ upper limit of 2.3×10^{-14} erg s⁻¹ cm⁻² (27 counts / 15827s) on the 0.5-2.0 keV emission from the region. With the better sensitivity of *XMM-Newton*, our result improves on this limit: when converted to 0.5-2.0 keV and 2σ , our upper limit becomes 1.38×10^{-14} erg s⁻¹ cm⁻². The 2σ minimum magnetic field strength also increases

Table 8.3	3: Physical param	eters inferred fro	m observations.		
Target	-2160	$+75^{a}$	140	$1-33^{b}$	
Measurement	MOS	PN	MOS	PN	
Equipartition: $B_{eq} \mapsto flux$					
Energy spectral index, $\alpha (S_{\nu} \propto \nu^{-\alpha})$	1	0.		1.4	
$c_{12}(\alpha)^c$	6	$.3 \times 10^7$		1.6×10^8	
Radio luminosity, $\log_{10}(\tilde{L}_r)^d$	41.	4	4	1.0	$\rm erg~s^{-1}$
Emitting volume, $\log_{10}(V)^d$	72	4.	2	1.3	cm^3
Magnetic field strength, B_{eq}^{e}	0	.63		1.3	μG
IC/CMB flux $(0.3-10 \text{ keV})^f$	œ	$.7 \times 10^{-14}$		5.5×10^{-13}	$\mathrm{erg}~\mathrm{s}^{-1}\mathrm{cm}^{-2}$
This study: flux $\mapsto B_{\text{measured}}$					
Energy spectral index, $\alpha (S_{\nu} \propto \nu^{-\alpha})$	1	0.		1.4	
$C(\alpha) = \sqrt{c_1} / \left[4\pi 2^{\alpha} c_5 (2\alpha + 1) \right]^c$	1	$.32 \times 10^{31}$		1.30×10^{31}	
$G(\alpha)^f$	0	.500		0.653	
Redshift, z	0	.125		0.0136	
Synchrotron flux density, $S_r(\nu_r)$	1	.2		0.5	$\mathbf{J}\mathbf{y}$
Frequency of radio measurement ν_r	0	.151		1.49	m GHz
$3-\sigma$ IC/CMB flux $(0.3-10 {\rm keV})$	$\leq 5.22 \times 10^{-14}$	$\leq 1.91 \times 10^{-13}$	$\leq 1.47 \times 10^{-13}$	$\leq 1.49 \times 10^{-13}$	${\rm erg~s^{-1}cm^{-2}}$
$S_x \nu_x^{\alpha} = \operatorname{flux} / \int \nu^{-\alpha} d\nu$	1.5×10^{-14}	5.4×10^{-14}	4.3×10^{-7}	4.4×10^{-7}	
3σ magnetic field strength, $\vec{B}_{\text{measured}}^{f}$	≥ 0.81	≥ 0.42	≥ 2.2	≥ 2.2	μG
^a Harris et al. 1993 ^b Goss et al. 1987 ^{c1} , $c_5(2\alpha + 1)$ and $c_{12}(\alpha)$ are defined by Pa	acholczyk (1970, App	endix 2, pp. 231–25	3).		

-4 Ą . --; Table 2. Dh ^dValues were from table 4 of Harris et al. 1993, scaled to the currently accepted cosmology. They assumed $H_0 = 50 \text{ km s}^{-1} \text{Mpc}^{-1}$, $q_0 = +0.5$, $\Omega_{\Lambda} = 0$, a lower cutoff frequency of 10 MHz, and simple spherical and cylindrical geometries for the volumes of emission. The original values were $\log_{10}(L_r) = 41.6$, $\log_{10}(V) = 72.8$ and $B_{\text{eq}} = 0.5 \,\mu\text{G}$ for 19017+75, and $\log_{10}(L_r) = 41.3$, $\log_{10}(V) = 71.7$ and $B_{\text{eq}} = 0.8 \,\mu\text{G}$ for 1401-33, although a recalculation gave $B_{\text{eq}} = 1.1 \,\mu\text{G}$ for 1401-33 instead.

Values were calculated using the classical equipartition formula: $B_{\text{eq}} = \left[\frac{8\pi(1+K)c_{12}L_{r}}{\phi V}\right]^{2/7}$. Values were calculated using Equation 11 in Harris and Grindlay 1979: $B^{\alpha+1} = \frac{(5.05 \times 10^4)^{\alpha}C(\alpha)G(\alpha)(1+z)^{\alpha+3}S_{r}\nu_{\alpha}^{\alpha}}{10^{47}S_{s}\nu_{\alpha}^{\alpha}}$, with $G(\alpha)$ from table 1 $10^{47} S_x \nu_x^{\alpha}$ therein.

154

from $0.76 \,\mu\text{G}$ to $0.99 \,\mu\text{G}$, as inferred from this and the radio measurements, but the difference is small compared to uncertainties in the equipartition value of $0.63 \,\mu\text{G}$.

$8.4.2 \quad 1401 - 33$

1401–33 extends over a 20 arcmin×9 arcmin region in the poor cluster Abell S753 around NGC 5419. The morphology is strongly suggestive of a relic radio galaxy. However, the parent galaxy has not been identified. Goss et al. (1987) reported a spectral flux density of $S_r(\nu_r) = 0.5$ Jy at $\nu_r = 1.4$ GHz, with a spectral index of $\alpha = 1.4$. Alternatively, Subrahmanyan et al. (2003) reported $\alpha = 1.4$ between 330 MHz and 1.4 GHz, possibly increasing to $\alpha = 1.9$ between 1.4 and 2.4 GHz, but only for the brightest part of the relic (the 'NE rim', which we consider in this study). When a large area of faint emission to the southwest was included, they reported $\alpha = 2.0$ between 330 MHz and 1.4 GHz, and 2.4 GHz. To reconcile with an 85-MHz measurement of 57 Jy from 1960, they suggested that there is a spectral break between 100 and 300 MHz. Assuming a break at 165 MHz, $\alpha = 0.7$ below the break frequency.

The only *ROSAT* observations of the region around 1401-33 were offset by 45 arcmin and are both less than 5 ks long, placing no interesting limit on the magnetic field strength. In our observation, we found an elevated flux level in the source region when compared to blank sky data, but no excess when compared to a local background region equidistant from the bright galaxy NGC 5419. Subrahmanyan et al. (2003) attributed this excess flux to thermal emission from a halo centred at NGC 5419. When we inspect our data with XSPEC, we are able to confirm that this excess emission fits much better to a thermal spectrum than to a power-law spectrum, although the statistics are too low to produce a high-quality fit. The result from this study places a 3σ upper limit on any excess X-ray flux specific to 1401-33 at 1.47×10^{-13} erg s⁻¹ cm⁻² within 0.3–10.0 keV, and a minimum magnetic field of $2.2 \,\mu$ G. These are the first limits on IC/CMB emission reported for 1401-33.

8.4.3 On the flux and field limits

While the number of known relics has grown beyond 30 (see, e.g., Giovannini et al., 1999; Kempner and Sarazin, 2001, for representative samples), studies of their IC/CMB emission remain scarce, and not a single convincing detection has been made so far. Table 8.4 lists the relics whose IC/CMB emission has been studied, together with the reported field limits. For consistent comparison, we have scaled the equipartition fields from previous studies to the currently acceptable cosmology, but we make no attempt to correct for differences in other parameters (e.g., the frequency band integrated), as the required information is not always available.

There are three claims of possible detections to date, in the Abell clusters A 85, A 754 and A 2034.

		$\operatorname{Ref.}^a$		2	4	9	2	6	11	12	12	
its.	IC/CMB	$\operatorname{Instrument}$		ROSAT	BeppoSAX	RXTE	Chandra	XMM-Newton	XMM-Newton	XMM-Newton	XMM-Newton	
measurement		$B_{ m measured}$	(ηG)	0.95^c	0.1^c	>0.84	$0.3-0.9^{d}$	>1.5	>1.05	>0.81	>2.2	F
Table 8.4: List of relics with field limits from published IC/CM	Classical equipartition	$B_{\rm eq}$ (μG)	cosmology updated	1.1(6)	0.3(3)	1.(8)	:	14.0	0.6(6)	0.63	1.2(5)	
		$B_{ m eq}~(\mu{ m G})$	in references ^{b}	1.3	0.3	2	÷	14.4^e	0.6	0.5	1.1^f	
		H_0		100	50	100	:	75	50	50	50	- (0) -
		ĸ		0.0555	0.0542	0.022	0.113	0.0562	0.023	0.125	0.0136	•••••••••••••••••••••••••••••••••••••••
		$\operatorname{Ref.}^a$		1	က	5 L	:	×	10	10	10	.7
	Host	cluster ⁻		A 85	A 754	A 1367	A 2034	A 133	\mathbf{Coma}	Rood 27	A S753	1 1 1 1 1
	Relic			0038 - 096	'East relic'	1140 + 203	'A 2034 '	'A 133_7a'	1253 + 275	0917 + 75	1401 - 33	9 CL

Table 8.4: List of relics with field limits from published IC/CMI	3 measurement	
Table 8.4: List of relics with field limits from published IC/C	CIME	
Table 8.4: List of relics with field limits from published I	C/C	
Table 8.4: List of relics with field limits from publis	hed I	
Table 8.4: List of relics with field limits from]	publis	
Table 8.4: List of relics with field limits	from]	
Table 8.4: List of relics with field l	imits	
Table 8.4: List of relics with field	Ξ	
Table 8.4: List of relics with	field	
Table 8.4: List of relics	with	
Table 8.4: List of	relics	
Table 8.4: List	of	
Table 8.4:	List	
Table	8.4:	
	Table	

^aReferences. - (1) Feretti and Giovannini 1996; (2) Bagchi et al. 1998; (3) Bacchi et al. 2003; (4) Fusco-Femiano et al. 2003; (5) Gavazzi and Trinchieri 1983; (6) Henriksen and Mushotzky 2001; (7) Kempner et al. 2003; (8) Slee et al. 2001; (9) Fujita et al. 2004; (10) Harris et al. 1993; (11) Feretti and Neumann 2006; (12) this paper. ^bAll cited B_{eq} values assume $q_0 = +0.5$, $\Omega_{\Lambda} = 0$, and H_0 as shown above. ^cTC/CMB detection has since been refuted by later studies. ^dThe authors commented that the errors are 'admittedly quite large.' ^cMinimum energy instead of equipartition is assumed. ^fRecalculation of B_{eq} using luminosity and volume from Ref.(12) gives this value instead $0.8 \,\mu$ G from Ref.(12).
Bagchi et al. (1998) detect IC/CMB emission from 0038-096 near Abell 85, using the PSPC on *ROSAT*. However, the statistics are poor, and the observation suffers from thermal contamination. More recently, Durret et al. (2005) determine from a more sensitive observation with *XMM-Newton* that the X-ray and radio emissions do not coincide spatially, and that the X-ray spectrum fits equally well to a thermal model with and without an added power-law component. Fusco-Femiano et al. (2003) find nonthermal emission above 45 keV in the aggregate spectrum of A 754 from the PDS on *BeppoSAX*; they attribute it to the central halo and 'east relic' in the cluster, but do not rule out contamination from the radio galaxy 26W 020 in the PDS field of view. Subsequently, Henriksen et al. (2004) are able to account for this hard X-ray component with a power-law model of 26W 020. Finally, Kempner et al. (2003) find weak evidence for IC/CMB emission from the relic in A 2034 with *Chandra*, but comment that the errors on the measurement are quite large. With the uncertanties in these studies, the nature of IC/CMB emission from relics remains elusive.

While none of the IC/CMB field strengths and upper limits in Table 8.4 is more than a factor of a few from the values derived from the 'classical' equipartition formula, it is interesting, nevertheless, to note that the three most recent measurements – on 0917+75, 1253+275 and 1401-33 – all give lower limits for the field strength which are slightly larger than the equipartition values. It is as yet unclear whether these deviations are significant, as is evident in the comparison of our field limits with the equipartition values obtained from the revised fomula of Brunetti et al. (1997). Both methods of determining the field strength involve parameters with large uncertainties, as we explain below.

The equipartition magnetic field, B_{eq} , is given by numerous authors in the past, for instance, as in Harris et al. (1995):

$$B_{eq} = \left[\frac{8\pi (1+K)c_{12}L_r}{\phi V}\right]^{2/7},$$

where L_r is the radio luminosity, V the emitting volume, K the ratio of proton energy to electron energy, ϕ the filling factor for the emitting plasma, and c_{12} a weak function of the spectral index and frequency band. We note that uncertainty in the cosmological distance scale contributes to uncertainty in L/V. In addition, we have assumed the most conservative values for K (= 0) and ϕ (= 1), which results in high X-ray flux and low B estimates at equipartition. In order to reconcile with the limits imposed by our results, $(1 + K)/\phi$ has to be 2.4 times larger for 0917+75 and 6.8 times larger for 1401-33.

In calculating the IC/CMB flux, we have assumed that we can infer the nature of the IC/CMBemitting electrons from past radio observations, while this may not indeed be the case. For instance, although the spectral indices for IC/CMB and synchrotron emissions should be the same, electrons responsible for IC/CMB emission in the energy range of XMM-Newton (0.3–10.0 keV) have 586 < $\gamma < 3381$, and in a 1- μ G magnetic field, they emit synchrotron radiation at frequencies between 1 and 50 MHz, where radio observations are either poor in spatial resolution, or entirely unavailable from Earth, due to reflection off the ionosphere. Thus, we can only rely on extrapolations of the spectral indices at higher frequencies to infer the IC/CMB index. Yet, as is evident from the discussion above, radio spectra of relic galaxies often exhibit complex shapes with multiple breaks. For both 0917+75 and 1401-33, there are indications that a break may exist in their radio spectra between 100 and 300 MHz, and it is not guaranteed that these are the only ones. If the spectrum of a synchrotron source flattens at low frequencies, that would imply lower IC/CMB flux as well, leading to closer agreement of IC/CMB measurements and equipartiton.

Various authors have proposed alternatives to the classical equipartition formula: Brunetti et al. (1997) integrate over a fixed particle energy range instead of a fixed synchrotron frequency range; Beck and Krause (2005) replace the energy density ratio of ions to electrons with their number density ratio, and give estimates of the number density ratio for various types of objects; Pfrommer and Enßlin (2004) consider minimum energy and equipartition criteria in the particular scenario where the synchrotron-emitting electrons are produced in inelastic collisions between cosmic-ray protons and ambient thermal gas in the cluster. Yet, uncertainties in the source geometry and electron population can only be resolved with new measurements.

One way to better determine the various uncertain paremeters is to make new measurements at as yet unexplored wavelengths. A new generation of low-frequency radio telescope arrays could provide radio measurements of relic galaxies at low frequencies (10-250 MHz) in the future, thus better constraining both the spectral index and spatial extent of IC/CMB emitting particles. Alternatively, a new hard X-ray / soft gamma-ray telescope would enable us to observe the IC/CMB emission by particles whose synchrotron emission is currently observable (particles emitting synchrotron radiation at 330 MHz will emit IC/CMB radiation at 70 keV), thus eliminating the problems of uncertain geometry and an uncertain spectral index. Hard X-ray observations also have the virtue that thermal emission no longer dominates in this energy range, so that sources thermally contaminated in soft X-rays can be observed, substantially increasing the sample size. Better determination of the filling factor ϕ could come from radio measurements of relics with higher angular resolution.

8.5 Summary

With new X-ray measurements of the two relic sources 0917+75 and 1401-33 using the XMM-Newton observatory, we detected no IC/CMB flux from either target. We set the 3σ upper limits on the IC/CMB flux within 0.3-10 keV at $5.22 \times 10^{-14} \text{ erg s}^{-1} \text{ cm}^{-2}$ for 0917+75 and $1.47 \times 10^{-13} \text{ erg s}^{-1} \text{ cm}^{-2}$ for 1401-33. The corresponding 3σ lower limits on their magnetic field strengths are $0.81 \,\mu\text{G}$ and $2.2 \,\mu\text{G}$, respectively, both slightly larger than the classical estimates of the equipartition field. Our study adds to the small sample of relics with limits on their IC/CMB emission. X-ray studies on more relics are needed for a broader picture of this diverse class of radio sources. Further constraints on the particle density and magnetic field strength can be made by IC/CMB measurements at higher X-ray energies (above 10 keV), and by synchrotron measurements at lower radio frequencies (below 300 MHz).

Chapter 9

Current status and future outlook

In this thesis, I have reported on the development of focal plane detectors for the High Energy Focusing Telescope, which took place at our laboratory at Caltech in the past decade or so. I have made the scientific case in Chapters 1 and 5 for the need of a focusing telescope in the hard X-ray band. Chapter 2 summarized the HEFT project in general, while Chapter 3 described the principles and our implementation of the HEFT focal plane detectors in depth. The first flight of HEFT, described in Chapter 4, confirmed that this focal plane system performed as specified and as expected. It is my hope that through this report, the reader can be convinced that detector development for HEFT is an interesting, meaningful and important project.

Although this thesis is a report on HEFT, any treatise on focusing technologies for hard Xray astrophysics cannot be complete without a report on the other two major projects for the development of focusing hard X-ray telescopes: InFOC μ S and HERO, which we introduced in Chapter 1. A summary of the development of focal-plane detectors for each project and the overall project status thus follows.

InFOC μ **S** The InFOC μ S team has also developed their focal plane using CdZnTe pixel detectors. Each focal plane hosts a single 2 mm thick CdZnTe detector with a 12-by-12 array of pixels of 2.1 mm pitch. As opposed to a CdZnTe-ASIC hybrid design, technical difficulties forced them to fall back to a more traditional approach of bonding the CdZnTe and ASIC each to a printed circuit board (PCB). Conductive epoxy bonds the detector anodes to its PCB, and a flex cable carries signals between the detector PCB and the ASIC PCB. As of 2001, the InFOC μ S detectors operated with a spectral resolution of 2.3 and 4.8 keV FWHM at 22 and 60 keV, respectively (Baumgartner et al., 2003a).

As of the writing of this report in 2007, $InFOC\mu S$ has had four flight campaigns. It was the first project to launch a partial payload. On 2000- 8-29, $InFOC\mu S$ launched its focal plane (without mirrors) from Palestine, Texas to measure the in-flight background at 36.3 km altitude for seven hours. Its first launch with both detector and mirrors was on 2001- 7- 5, from Palestine, Texas. With a single telescope module of 42 cm^2 collecting area, it observed Cyg X-1 for two hours. However, its attitude control system failed to function properly to lock the target continuously within the field of view of the telescope. This problem substantially affected the quality of the Cyg X-1 image and spectra from this flight. InFOC μ S subsequently made launches in August and September of 2004, from Fort Sumner, New Mexico. During its fourth launch on 2004- 9-16, it observed the recurrent transient Be/X-ray binary pulsar 4U 0115+634 in outburst for 12.9ks, obtaining the first focused image of this object in the hard X-ray band.

HERO Compared to HEFT and InFOC μ S, HERO has taken different approaches to tackle the problem of hard X-ray focusing. As mentioned in Section 1.2, the HERO team develops mirrors with high graze angles and no multilayer coating, and matches them with high-pressure gas (96% Xe, 4% He) scintillators instead of semiconductor pixel detectors such as CdZnTe. Despite this initial configuration, they have also started developing CdZnTe pixel detectors for their focal planes circa 2002. Because the angular resolution of the HERO mirrors (15") is much higher than that of HEFT and InFOC μ S, HERO also requires matching focal-plane detectors with finer pixel pitch and more pixels per unit area on the focal plane. Their ultimate goal is to have a 64-by-64 pixel array of 200 μ mpitch behind each of their mirror modules. As of 2004, they had fabricated and were studying a number of CdZnTe-ASIC hybrids with 16-by-16 pixel arrays of 300 μ mpitch on 1 mm and 2 mm thick CdZnTe detectors. Gaskin et al. (2003b) reported the spectral resolution to be 2% (1.2 keV) at 60 keV at a typical pixel, with individual pixels having even better resolution. Because of the fine pitch of these detectors, the phenomenon of charge sharing, as described in Section 3.9 and in Chen et al. (2002), is a much bigger concern for the HERO project.

HERO was first launched on 2000- 9-19 and again on 2001- 5-23 from Fort Sumner, New Mexico. During the 17-hour long 'demonstration' flight in 2001 (Ramsey et al., 2002), it carried three engineering modules of high-graze-angle mirrors and gas scintillators, providing a modest effective area of 4 cm² (at 40 keV) and 45" HPD angular resolution. It observed the Crab Nebula, Cyg X-1 and GRS 1915+105, and made focused images of these objects—in fact, the first focused images of any object at all—in the hard X-ray sky. In 2004, HERO resumed its flight campaigns with eight full modules of mirrors and gas scintillators, having over 80 cm² of collecting area within 20-50 keV, and better than 15" HPD angular resolution. It was launched on 2005- 5- 9 and again on 2006- 9-25 and 2007- 5-27. The first two flights were marred by an unknown electronics problem and a balloon failure, respectively; science results from the latest flight are yet to be published.

When compared to $InFOC\mu S$ and HERO, it took HEFT the longest to reach a flight-ready status. On the other hand, HEFT has had the most success in the development of focal plane detectors. We have fabricated the highest number of detector hybrids, each hybrid hosts the highest number of pixels, and the spectral performance of some of our detector hybrids is also the best in the field of focal-plane detectors for hard X-ray astronomy to date. The HEFT detector hybrids have even found their use in other fields of science. For instance, the materials science group at Caltech is adopting our detector design (modified with more economical CdTe detectors) in a new generation of Mössbauer powder diffractometers (Lin and Fultz, 2003; Fultz and Lin, 2003), with the expectation of improving the signal-to-noise ratio from about 1:1 in previous experiments to about 10:1 or better; this will be extremely beneficial to the development of the novel technique of Mössbauer diffractometry, for the determination of material structures. Clearly, what remains to be achieved is to put the HEFT detectors into their intended use, to study interesting phenomena in hard X-ray astronomy, by making new observations of celestial targets of scientific value with HEFT.

At the time of writing of this thesis, there is active planning for a second flight of HEFT. The Columbia Scientific Ballooning Faculity (CSBF, formerly NSBF) is planning to run a southernhemisphere campaign in Alice Spring, Australia within the next few years. We are in the process of fabricating more HEFT focal planes, and redesigning and rebuilding the deformed HEFT gondola, in order to participate in this upcoming campaign. HEFT is also collaborating with the Institute of Space and Astronautical Science (ISAS) of Japan, who will contribute an instrument for the focal plane of one of the three HEFT telescopes.

While the characteristics and quality of the HEFT detectors described in this thesis suffice to meet the need of the current generation of hard X-ray focusing instruments, there is still a need for continuing development of these focal-plane detectors. As hard X-ray mirror technology matures, the angular resolution of mirrors gets finer and finer. As the HERO mirror modules testify, the halfpower diameters (HPDs) of new mirrors are reaching the subarcminute level. Besides, the typical point spread function (PSF) of hard X-ray mirrors contains a narrow core flanked by broad wings. What matters the most is often not the HPD, which measures the extent of the wings, but the narrower width of the core instead. These angular scales, when projected onto the focal plane, will soon be comparable to the pixel sizes of the current generation of focal-plane detectors. Thus, the need for further development of detectors with finer spatial resolution is imminent, and the HERO detector team is forging their way in this direction. In addition, there is always a quest for ever finer spectral resolution, as even the best hard X-ray detectors today are still faring much worse than well-established semiconductor technologies such as silicon and germanium at lower energies. Also, as Chapter 3 reveals, several other technical difficulties are currently impeding progress in the field—detector quality can improve tremendously by reducing material inhomogeneity; one can relax constraining operating temperature ranges by reducing leakage current in detector materials, or with a new design of read-out electronics that can tolerate higher leakage (e.g., through AC coupling) without sacrificing good spectral resolution; larger spectroscopy-grade semiconductor crystals will also eliminate the need to have multiple hybrids on each focal plane, thus removing dead areas at the physical gap between hybrids. In a mechanical engineering standpoint, focal plane assembly methods also requires improvement. Take HEFT as an example, as explained in Section 3.8.4, each focal plane consists of a single motherboard with support circuitry for two hybrids, which are both glued down to the motherboard. Yet, a working detector hybrid does not come easily, and given the design of the HEFT motherboard, failure of a single hybrid could potentially yield the other one unusable, as the removal of a bad hybrid from a motherboard is risky and untidy. To mitigate this issue, our group is currently developing smaller, pluggable motherboards that each hosts a single detector hybrid. These focal-plane 'halves' connect to remaining focal-plane electronics via a welltested removable interface. With such improved engineering design, we expect to lower our hardware failure rate, and thus to reduce wasted resources, including manpower, time and monetary expenses.

Looking further into the next decade, prospects for further development of hard X-ray focusing astronomy is promising. In 2003, we proposed to NASA a Small Explorer mission called the *Nuclear Spectroscopic Telescope Array (NuSTAR)*. *NuSTAR* is a modified version of HEFT on a satellite platform, and the HEFT collaboration has played a central role in the development of *NuSTAR*. So far, *NuSTAR* has gone through a tumultuous journey. On 2003-11- 4, NASA selected *NuSTAR* for a five-month 'Phase A' implementation feasibility study, and on 2005- 1-26, for an 'Extended Phase A' study. Due to NASA's budgetary constraint in subsequent fiscal years, it cancelled NuSTAR in February 2006 solely for financial reasons. Yet, on 2007- 9-21, NASA restarted *NuSTAR*, with an expected launch date in 2011. If all goes well from now onwards, *NuSTAR* will provide access to the hard X-ray sky for years of new and exciting high-energy astrophysics.

In fact, NuSTAR is not the only hard X-ray focusing mission in development at this time. The Japanese space agency (ISAS/JAXA) is also developing a similar mission called the *New X-ray Telescope* (*NeXT*), proposed to be launched in 2010, while the French and Italian space agencies (CNES and ASI) are studying yet another mission called *Simbol-X*, with a target launch period in early 2013. With potentially three hard X-ray focusing missions in orbit within the next five years, this is an exciting time for hard X-ray astrophysics.

Bibliography

Astrophysics literature

- Ahmad, I., Greene, J. P., Moore, E. F., Ghelberg, S., Ofan, A., Paul, M., and Kutschera, W. (2006). Improved measurement of the ⁴⁴Ti half-life from a 14-year long study. *Physical Review C (Nuclear Physics)*, 74(6):065803.
- Arnaud, M., Majerowicz, S., Lumb, D., Neumann, D. M., Aghanim, N., Blanchard, A., Boer, M., Burke, D. J., Collins, C. A., Giard, M., Nevalainen, J., Nichol, R. C., Romer, A. K., and Sadat, R. (2002). XMM-Newton observation of the distant (z=0.6) galaxy cluster RX J1120.1+4318. A&A, 390(1):27-38.
- Bacchi, M., Feretti, L., Giovannini, G., and Govoni, F. (2003). Deep images of cluster radio halos. A&A, 400(2):465–476.
- Bagchi, J., Pislar, V., and Lima Neto, G. B. (1998). The diffuse, relic radio source in Abell 85: estimation of cluster-scale magnetic field from inverse Compton x-rays. *MNRAS*, 296(1):L23–L28.
- Beck, R. and Krause, M. (2005). Revised equipartition and minimum energy formula for magnetic field strength estimates from radio synchrotron observations. Astronomische Nachrichten, 326(6):414-427.
- Brandt, W. N. and Hasinger, G. (2005). Deep extragalactic x-ray surveys. Ann. Rev. Astron. Astrophys., 43:827–859.
- Brunetti, G., Setti, G., and Comastri, A. (1997). Inverse Compton x-rays from strong FRII radiogalaxies. A&A, 325(3):898–910.
- Carter, J. A. and Read, A. M. (2007). The XMM-Newton EPIC background and the production of background blank sky event files. A&A, 464(3):1155–1166.
- Condon, J. J., Cotton, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., and Broderick,
 J. J. (1998). The NRAO VLA sky survey. AJ, 115(5):1693-1716.

- Croston, J. H., Hardcastle, M. J., Harris, D. E., Belsole, E., Birkinshaw, M., and Worrall, D. M. (2005). An x-ray study of magnetic field strengths and particle content in the lobes of FR II radio sources. ApJ, 626(2):733-747.
- Dewdney, P. E., Costain, C. H., McHardy, I., Willis, A. G., Harris, D. E., and Stern, C. P. (1991). An x-ray and radio study of steep-spectrum radio sources. ii. four fields from a 22 MHz polar-cap survey. ApJS, 76(4):1055–1066.
- Durret, F., Lima Neto, G. B., and Forman, W. (2005). An XMM-Newton view of the cluster of galaxies Abell 85. A&A, 432(3):809-821.
- Fabian, A. C. and Barcons, X. (1992). The origin of the x-ray background. Ann. Rev. Astron. Astrophys., 30:429–456.
- Feretti, L. and Giovannini, G. (1996). Diffuse cluster radio sources. In Ekers, R. D., Fanti, C., and Padrielli, L., editors, *Extragalactic Radio Sources*, volume 175 of *The International Astronomical* Union Symposium Proceedings, pages 333–338, Boston. Kluwer Academic Publishers.
- Feretti, L. and Neumann, D. M. (2006). XMM-Newton observations of the Coma cluster relic 1253+275. A&A, 450(3):L21-L24.
- Fujita, Y., Sarazin, C. L., Reiprich, T. H., Andernach, H., Ehle, M., Murgia, M., Rudnick, L., and Slee, O. B. (2004). XMM-Newton observations of A133: A weak shock passing through the cool core. ApJ, 616(1):157–168.
- Fusco-Femiano, R., Orlandini, M., De Grandi, S., Molendi, S., Feretti, L., Giovannini, G., Bacchi, M., and Govoni, F. (2003). Hard x-ray and radio observations of Abell 754. A&A, 398(2):441-446.
- Gavazzi, G. and Trinchieri, G. (1983). Radio and x-ray observations of the radio halo source in A1367. ApJ, 270(2):410-416.
- Giacconi, R., Gursky, H., Paolini, F. R., and Rossi, B. B. (1962). Evidence for x rays from sources outside the solar system. *Phys. Rev. Lett.*, 9(11):439-443.
- Giovannini, G. and Feretti, L. (2004). Radio relics in clusters of galaxies. Journal of the Korean Astronomical Society, 37(5):323–328.
- Giovannini, G., Tordi, M., and Feretti, L. (1999). Radio halo and relic candidates from the NRAO VLA Sky Survey. New Astronomy, 4(2):141–155.
- Goss, W. M., McAdam, W. B., Wellington, K. J., and Ekers, R. D. (1987). The very low-brightness relic radio galaxy 1401-33. MNRAS, 226(4):979–988.

- Harris, D. E. and Grindlay, J. E. (1979). The prospects for x-ray detection of inverse-Compton emission from radio source electrons and photons of the microwave background. MNRAS, 188:25– 37.
- Harris, D. E., Stern, C. P., Willis, A. G., and Dewdney, P. E. (1993). The megaparsec radio relic in supercluster, Rood #27. AJ, 105(3):769–777.
- Harris, D. E., Willis, A. G., Dewdney, P. E., and Batty, J. (1995). Constraints on the average magnetic field strength in a relic radio galaxy derived from limits on inverse Compton x-rays. MNRAS, 273(3):785-789.
- Henriksen, M. and Mushotzky, R. (2001). X-ray measurements of nonthermal emission from the Abell 1367 galaxy cluster using the Rossi X-ray Timing Explorer. ApJ, 553(1):84–89.
- Henriksen, M. J., Hudson, D. S., and Tittley, E. (2004). A soft x-ray excess in the A754 cluster. ApJ, 610(2):762-771.
- Kempner, J. C., Blanton, E. L., Clarke, T. E., Enßlin, T. A., Johnston-Hollitt, M., and Rudnick, L. (2004). Conference note: A taxonomy of extended radio sources in clusters of galaxies. In Reiprich, T. H., Kempner, J. C., and Soker, N., editors, *The Riddle of Cooling Flows in Galaxies and Clusters of Galaxies*, page E25.
- Kempner, J. C. and Sarazin, C. L. (2001). Radio halo and relic candidates from the Westerbork Northern Sky Survey. ApJ, 548(2):639–651.
- Kempner, J. C., Sarazin, C. L., and Markevitch, M. (2003). Chandra observation of the merging cluster A2034. ApJ, 593(1):291–300.
- Krivonos, R., Vikhlinin, A., Churazov, E., Lutovinov, A., Molkov, S., and Sunyaev, R. (2005). Extragalactic source counts in the 20–50 kev energy band from the deep observation of the Coma region by INTEGRAL IBIS. ApJ, 625(1):89–94.
- Ling, J. C., Mahoney, W. A., Willett, J. B., and Jacobson, A. S. (1979). A possible line feature at 73 kev from the Crab Nebula. *ApJ*, 231:896–905.
- McGlynn, T., Scollick, K., and White, N. (1998). SkyView: The multi-wavelength sky on the internet. In McLean, B. J., Golombek, D. A., Hayes, J. J. E., and Payne, H. E., editors, New Horizons from Multi-Wavelength Sky Surveys, volume 179 of The International Astronomical Union Symposium Proceedings, pages 465-466, Boston. Kluwer Academic Publishers.
- Mukai, K. (1993). PIMMS and Viewing: proposal preparation tools. Legacy, 3:21-31.

- Nevalainen, J., Markevitch, M., and Lumb, D. (2005). XMM-Newton EPIC background modeling for extended sources. ApJ, 629(1):172–191.
- Pacholczyk, A. G. (1970). Radio astrophysics. Nonthermal processes in galactic and extragalactic sources. A Series of Books in Astronomy and Astrophysics. W. H. Freeman, San Francisco.
- Pelling, R. M., Paciesas, W. S., Peterson, L. E., Makishima, K., Oda, M., Ogawara, Y., and Miyamoto, S. (1987). A scanning modulation collimator observation of the high-energy x-ray source in the Crab Nebula. ApJ, 319:416–425.
- Pfrommer, C. and Enßlin, T. A. (2004). Estimating galaxy cluster magnetic fields by the classical and hadronic minimum energy criterion. *MNRAS*, 352(1):76–90.
- Pradas, J. and Kerp, J. (2005). XMM-Newton data processing for faint diffuse emission. proton flares, exposure maps and report on EPIC MOS1 bright CCDs contamination. A&A, 443(2):721-733.
- Remillard, R. (2005). X-ray outburst of the blazar, 3C 454.3. The Astronomer's Telegram, 484.
- Renaud, M., Vink, J., Decourchelle, A., Lebrun, F., den Hartog, P. R., Terrier, R., Couvreur, C., Knödlseder, J., Martin, P., Prantzos, N., Bykov, A. M., and Bloemen, H. (2006). The signature of ⁴⁴Ti in Cassiopeia A revealed by IBIS/ISGRI on INTEGRAL. ApJ, 647(1):L41–L44.
- Slee, O. B., Roy, A. L., Murgia, M., Andernach, H., and Ehle, M. (2001). Four extreme relic radio sources in clusters of galaxies. AJ, 122(3):1172–1193.
- Snowden, S. L., Collier, M. R., and Kuntz, K. D. (2004). XMM-Newton observation of solar wind charge exchange emission. ApJ, 610(2):1182–1190.
- Subrahmanyan, R., Beasley, A. J., Goss, W. M., Golap, K., and Hunstead, R. W. (2003). PKS B1400-33: An unusual radio relic in a poor cluster. AJ, 125(3):1095–1106.
- Vink, J., Laming, J. M., Kaastra, J. S., Bleeker, J. A. M., Bloemen, H., and Oberlack, U. (2001).
 Detection of the 67.9 and 78.4 kev lines associated with the radioactive decay of ⁴⁴Ti in Cassiopeia
 A. ApJ, 560(1):L79–L82.
- Zombeck, M. V. (1990). Handbook of Space Astronomy and Astrophysics. Cambridge University Press, Cambridge, 2nd edition.

Engineering literature

- Arodzero, A., Brau, J. E., Frey, R. E., Gao, D., Kollipara, R. T., Langston, M., Mason, D., Sinev, N., Strom, D., Yang, X., Brooks, M., Lee, D., and Mills, G. (1996). Beam test of prototype 18 cm silicon-strip detectors with high speed electronics. *IEEE Trans. Nucl. Sci.*, 43(3):1180–1187.
- Baumgartner, W. H., Tueller, J., Krimm, H., Barthelmy, S. D., Berendse, F., Ryan, L., Birsa, F. B., Okajima, T., Kunieda, H., Ogasaka, Y., Tawara, Y., and Tamura, K. (2003a). InFOCµS hard x-ray telescope: pixellated CZT detector/shield performance and flight results. In Truemper, J. E. and Tananbaum, H. D., editors, X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy, volume 4851 of Proc. SPIE, pages 945–956.
- Boggs, S. E. (1998). Gamma-ray spectroscopy: galactic center observations and germanium detector development. PhD thesis, University of California, Berkeley.
- Bolotnikov, A. E., Boggs, S. E., Chen, C. M. H., Cook, W. R., A, F., Harrison, and Schindler, S. M. (2000a). Optimal contact geometry for CdZnTe pixel detectors. In James, R. B. and Schirato, R. C., editors, *Hard X-Ray, Gamma-Ray, and Neutron Detector Physics II*, volume 4141 of *Proc. SPIE*, pages 243–252.
- Bolotnikov, A. E., Boggs, S. E., Cook, W. R., Harrison, F. A., and Schindler, S. M. (1999a). Use of a pulsed laser to study properties of CdZnTe pixel detectors. In Doty, F. P., editor, *Penetrating Radiation Systems and Applications*, volume 3769 of *Proc. SPIE*, pages 52–58.
- Bolotnikov, A. E., Cook, W. R., Harrison, F. A., Wong, A.-S., Schindler, S. M., and Eichelberger, A. C. (1999b). Charge loss between contacts of CdZnTe pixel detectors. *Nucl. Instr. Meth.*, A432(2-3):326-331.
- Chen, C. M. H., Boggs, S. E., Bolotnikov, A. E., Cook, W. R., Harrison, F. A., and Schindler, S. M. (2002). Numerical modeling of charge sharing in CdZnTe pixel detectors. *IEEE Trans. Nucl. Sci.*, 49(1):270–276.
- Chen, C. M. H., Christensen, F. E., Harrison, F. A., Mao, P. H., and Windt, D. L. (2003). Design of a soft gamma-ray focusing telescope for the study of nuclear lines. In Truemper, J. E. and Tananbaum, H. D., editors, X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy, volume 4851 of Proc. SPIE, pages 1356–1365.
- Chen, C. M. H., Cook, W. R., Harrison, F. A., and Lin, J. Y. Y. (2004a). Characterization of a large-format, fine-pitch CdZnTe pixel detector for the HEFT balloon-borne experiment. *IEEE Trans. Nucl. Sci.*, 51(5):2472–2477.

- Chen, C. M. H., Cook, W. R., Harrison, F. A., Lin, J. Y. Y., Mao, P. H., and Schindler, S. M. (2004b). Characterization of the HEFT CdZnTe pixel detectors. In Franks, L. A., Burger, A., James, R. B., and Hink, P. L., editors, *Hard X-Ray and Gamma-Ray Detector Physics V*, volume 5198 of *Proc. SPIE*, pages 9–18.
- Chonko, J. C. (2006). Development and first flight of the High Energy Focusing Telescope and analysis of spectral features in the isolated neutron star 1E1207.4-5209. PhD thesis, Columbia University, New York.
- Christensen, F. E., Chakan, J. M., Harrison, F. A., Boggs, S. E., Mao, P. H., Prince, T. A., Craig, W. W., Hailey, C. J., and Windt, D. L. (2000). Grazing incidence optics designs for future gamma-ray missions. In Truemper, J. E. and Aschenbach, B., editors, X-Ray Optics, Instruments, and Missions III, volume 4012 of Proc. SPIE, pages 278–283.
- Christensen, F. E., Jensen, C. P., Madsen, K. K., Pivovaroff, M. J., Chen, H., Dariel, A., and Høghøj, P. (2006). Novel multilayer designs for future hard x-ray missions. In Turner, M. J. L. and Hasinger, G., editors, *Space Telescopes and Instrumentation II: Ultraviolet to Gamma Ray*, volume 6266 of *Proc. SPIE*, page 626611.
- Cook, W. R., Boggs, S. E., Bolotnikov, A. E., Burnham, J. A., Harrison, F. A., Kecman, B., Matthews, B., Schindler, S. M., and Fitzsimmons, M. J. (1999). First test results from a highresolution CdZnTe pixel detector with VLSI readout. In Doty, F. P., editor, *Penetrating Radiation* Systems and Applications, volume 3769 of Proc. SPIE, pages 92–96.
- Cook, W. R., Burnham, J. A., and Harrison, F. A. (1998). Low-noise custom VLSI for CdZnTe pixel detectors. In Siegmund, O. H. W. and Gummin, M. A., editors, EUV, X-Ray, and Gamma-Ray Instrumentation for Astronomy IX, volume 3445 of Proc. SPIE, pages 347–354.
- Fultz, B. and Lin, J. Y. Y. (2003). Mössbauer diffractometry. In Mashlan, M., Miglierini, M., and Schaaf, P., editors, Material Research in Atomic Scale by Mössbauer Spectroscopy, volume 94 of NATO Science Series: II: Mathematics, Physics and Chemistry, pages 285-295.
- Gaskin, J. A., Richardson, G. A., Mitchell, S., Sharma, D. P., Ramsey, B. D., and Seller, P. (2003b). Assessment of cadmium-zinc-telluride detectors for hard-x-ray astronomy. In 2003 IEEE Nuclear Science Symposium Conference Record, volume 5, pages 3404–3407.
- Gunderson, K. S., Chen, C. M. H., Christensen, F. E., Craig, W. W., Decker, T. A., Hailey, C. J., Harrison, F. A., McLean, R., Wurtz, R. E., and Ziock, K. (2004). Ground performance of the High-Energy Focusing Telescope (HEFT) attitude control system. In Flanagan, K. A. and Siegmund, O. H. W., editors, X-Ray and Gamma-Ray Instrumentation for Astronomy XIII, volume 5165 of Proc. SPIE, pages 158–168.

- Harrison, F. A., Boggs, S. E., Bolotnikov, A. E., Christensen, F. E., Cook, W. R., Craig, W. W., Hailey, C. J., Jimenez-Garate, M. A., Mao, P. H., Schindler, S. M., and Windt, D. L. (2000).
 Development of the High-Energy Focusing Telescope (HEFT) balloon experiment. In Truemper, J. E. and Aschenbach, B., editors, X-Ray Optics, Instruments, and Missions III, volume 4012 of Proc. SPIE, pages 693–699.
- Jensen, C. P., Christensen, F. E., Chen, H. C., Smitt, E. B. W., and Ziegler, E. (2002). Multilayer coating facility for the HEFT hard x-ray telescope. In Gorenstein, P. and Hoover, R. B., editors, X-Ray Optics for Astronomy: Telescopes, Multilayers, Spectrometers, and Missions, volume 4496 of Proc. SPIE, pages 104–108.
- Jensen, C. P., Madsen, K. K., Chen, H. C., Christensen, F. E., and Ziegler, E. (2003). Coating of the HEFT telescope mirrors: method and results. In Truemper, J. E. and Tananbaum, H. D., editors, X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy, volume 4851 of Proc. SPIE, pages 724-733.
- Joensen, K. D. (1995). Design, fabrication and characterization of multilayers for broad-band, hard X-ray astrophysics instrumentation. PhD thesis, H. C. Ørsted Instituttet, Copenhagen, Denmark.
- Kalemci, E. and Matteson, J. L. (2000). Charge splitting among anodes of a CdZnTe strip detector. In James, R. B. and Schirato, R. C., editors, *Hard X-Ray, Gamma-Ray, and Neutron Detector Physics II*, volume 4141 of *Proc. SPIE*, pages 235–242.
- Knoll, G. F. (2000). *Radiation Detection and Measurement*. John Wiley & Sons, New York, 3rd edition.
- Kruse Madsen, K. (2007). X-ray Telescopes: Development, Flight and Analysis. PhD thesis, Københavns Universitet, Copenhagen, Denmark.
- Lin, J. Y. Y. and Fultz, B. (2003). Site-specific long-range order in ⁵⁷Fe₃Al measured by Mössbauer diffractometry. *Phil. Mag.*, 83:2621–2640.
- Madsen, K. K., Christensen, F. E., Jensen, C. P., Ziegler, E., Craig, W. W., Gunderson, K. S., Koglin, J. E., and Pedersen, K. (2004). X-ray study of W/Si multilayers for the HEFT hard x-ray telescope. In Citterio, O. and O'Dell, S. L., editors, Optics for EUV, X-Ray, and Gamma-Ray Astronomy, volume 5168 of Proc. SPIE, pages 41–52.
- Mao, P. H.-J. (2002). Hard x-ray observations of the extragalactic sky: the High Energy Focusing Telescope and the Serendipitous Extragalactic X-ray Source Identification survey. PhD thesis, California Institute of Technology, Pasadena, California.

Melissinos, A. C. (1966). Experiments in Modern Physics. Academic Press, New York.

- Oonuki, K., Inoue, H., Nakazawa, K., Mitani, T., Tanaka, T., Takahashi, T., Chen, C. M. H., Cook, W. R., and Harrison, F. A. (2004). Development of uniform CdTe pixel detectors based on Caltech ASIC. In Holland, A. D., editor, *High-Energy Detectors in Astronomy*, volume 5501 of *Proc. SPIE*, pages 218–228.
- Philhour, B., Boggs, S. E., Primbsch, J. H., Slassi-Sennou, S., Lin, R. P., Feffer, P. T., McBride, S., Goulding, F. S., Madden, N. W., and Pehl, R. H. (1998). Simulations of pulse shape discrimination (PSD) techniques for background reduction in germanium detectors. *Nucl. Instr. Meth.*, A403(1):136-150.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1986). Numerical Recipes: The Art of Scientific Computing. Cambridge University Press, Cambridge.
- Ramsey, B. D., Alexander, C. D., Apple, J. A., Benson, C. M., Dietz, K. L., Elsner, R. F., Engelhaupt, D. E., Ghosh, K. K., Kolodziejczak, J. J., O'Dell, S. L., Speegle, C. O., Swartz, D. A., and Weisskopf, M. C. (2002). First images from HERO, a hard x-ray focusing telescope. *ApJ*, 568(1):432–435.
- Takahashi, T., Watanabe, S., Kouda, M., Sato, G., Okada, Y., Kubo, S., Kuroda, Y., Onishi, M., and Ohno, R. (2001). High-resolution CdTe detector and applications to imaging devices. *IEEE Trans. Nucl. Sci.*, 48(3):287–291.
- Van Speybroeck, L. P. and Chase, R. C. (1972). Design parameters of paraboloid-hyperboloid telescopes for x-ray astronomy. Applied Optics, 11(2):440–445.
- Windt, D. L., Donguy, S., Hailey, C. J., Koglin, J., Honkimaki, V., Ziegler, E., Christensen, F. E., Chen, H., Harrison, F. A., and Craig, W. W. (2003). W/SiC x-ray multilayers optimized for use above 100 KeV. Applied Optics, 42(13):2415-2421.
- Wolter, H. (1952a). Spiegelsysteme streifenden einfalls als abbildende optiken f
 ür r
 öntgenstrahlen (mirror systems with glancing incidence on image producing optics for x-rays). Annalen der Physik, 445(1):94–114.
- Wolter, H. (1952b). Verallgemeinerte schwarzschildsche spiegelsysteme streifender reflexion als optiken f
 ür r
 öntgenstrahlen (generalized schwarzschild mirror systems with glancing incidence as optics for x-rays). Annalen der Physik, 445(4):286–295.

Appendix A CdZnTe detector specification

The following is the specification of the CdZnTe detectors we purchased from its manufacturer and vendor, eV Products. A mechanical drawing, whose details are shown in Figure 3.7, accompanies this specification.

CdZnTe radiation detector, $13 \text{ mm} \times 23.7 \text{ mm} \times 2 \text{ mm} (+/-0.1 \text{ mm})$ with clear aperture area of 12.3 mm \times 22.6 mm located along one edge of part as displayed in customer supplied drawing fig [sic]. Part to be metalized as per customer specifications. Slice to be scanned for grain structure, polished for viewing, photographed for inclusions. Pictures of grains and inclusions to be sent with part.

- 1. Single crystal; twins allowed.
- 2. Scratch/dig per mil spec MIL-C-48497A < B scratch and < A dig.
- 3. Planarity of surfaces to be $< 5 \,\mu$ m.
- 4. Parallelism to be $< 50 \,\mu$ m.
- 5. Surface roughness to be < 10 nm, but will verify with measurements on only 10% of batch.
- 6. Maximum precipitate size to be $< 100 \,\mu$ m.
- 7. $\mu \tau_{\rm e}$ product to be > 2 × 10⁻³ cm²/V as measured during ingot characterization with voltage from 50 V–200 V.
- 8. Bulk resistivity > $1 \times 10^{10} \Omega$ -cm at 25 °C as measured during ingot characterization.
- 9. Pattern to be $< 100 \,\mu\text{m}$ from mating edge of part. Pattern to be $30 \,\mu\text{m}$ gap.

Appendix B

Analysis of the HEFT amplifier and read-out circuitry

B.1 Stage-by-stage circuit analysis

Figure B.1 shows the top-level diagramme of the HEFT amplifier and read-out ASIC. The circuitry in each pixel contains a generic sequence of a preamplifier, a postamplifier, a shaping amplifier and a discriminator. In addition, the postamplifier outputs a current whose waveform is continuously captured by a bank of sampling capacitors. The pixel circuitry is duplicated (48×24) -times for each pixel on the ASIC. Upon the detection of an event, the sampling capacitors at selected pixels (based on the event trigger) are piped in sequence to the single read-out amplifier on-chip. The read-out amplifier sends the samples out of the ASIC as a sequence of analogue voltage levels, which are digitized by a separate 12-bit analogue-to-digital converter IC.

B.1.1 The preamplifier

Figure B.2 shows the schematic diagramme of the preamplifier. The preamplifier, with the capacitor C_{preamp} placed in a negative feedback loop, is essentially an integrator of the input signal:

$$v_{\text{preamp}}(t) = -\frac{1}{C_{\text{preamp}}} \int_{t_{\text{reset}}}^{t} \left[i_{\text{signal}}(t') + I_{\text{leakage}} \right] \, \mathrm{d}t' \tag{B.1}$$

The output voltage, $v_{\text{preamp}}(t)$, is a ramp signal with slope determined by the amplitude of the leakage current.

The reset period, τ_{reset} , is set by the module MISC controller, and it must be an integer multiple of the 136 ns clock cycle (7.3728 MHz). We set:

$$\tau_{\text{reset}} = 8000 \text{ clock periods} = 1.09 \text{ ms}.$$







Figure B.2: The HEFT preamplifier circuit.

B.1.2 The postamplifier

G

Figure B.3 shows the schematic diagramme of the postamplifier. This circuit is essentially a voltage differentiator followed by a voltage regulated current source. The capacitor C_{postamp} serves as a high-pass filter (see Figure B.4), and screens out the constant ramping component in v_{preamp} resulting from the leakage current. Current coming out of C_{postamp} must pass through the resistor R_{postamp} , which serves as a converter from current to voltage:

$$v_{\text{postamp}} - v_{\text{DD}} = -iR_{\text{postamp}}$$
$$= -R_{\text{postamp}}C_{\text{postamp}}\frac{\mathrm{d}V_{\text{preamp}}}{\mathrm{d}t}.$$
(B.2)

Combining the integrator at the preamplifier stage (Equation B.1) and this differentiator at the postamplifier stage (Equation B.2), we have the gain of the two-stage transresistance amplifier:

$$v_{\text{postamp}} - v_{\text{DD}} = -iR_{\text{postamp}}$$

$$= -R_{\text{postamp}}C_{\text{postamp}}\frac{\mathrm{d}}{\mathrm{d}t} \left\{ -\frac{1}{C_{\text{preamp}}} \int_{t_{\text{reset}}}^{t} \left[i_{\text{signal}}(t') + I_{\text{leakage}} \right] \mathrm{d}t' \right\}$$

$$= R_{\text{postamp}} \left(\frac{C_{\text{postamp}}}{C_{\text{preamp}}} \right) i_{\text{signal}}$$
ain, $G = \frac{v_{\text{postamp}} - v_{\text{DD}}}{i_{\text{signal}}} = R_{\text{postamp}} \left(\frac{C_{\text{postamp}}}{C_{\text{preamp}}} \right).$
(B.3)

This gain gives the relation between the charge induced at the CdZnTe anode and the change in v_{postamp} as a result.

The time evolution of v_{postamp} following an X-ray event and the subsequent charge induction can be described as follows. The spike in input current causes a step in the ramping v_{preamp} . This abrupt step causes current to pass through C_{preamp} , also abruptly. Within the short time instant, the resistor R_{postamp} acts like an open circuit, while the capacitor C_{decay} in parallel acts like a short

175



Figure B.3: The HEFT postamplifier circuit.



Figure B.4: A capacitor connected in series at the input of an amplifier serves as a high-pass filter.

circuit, and allows the current through. This current causes a potential difference across C_{decay} , and C_{decay} has to discharge through R_{postamp} , with the discharge time constant given by:

$$\tau_{\text{postamp}} = R_{\text{postamp}}C_{\text{decay}}$$

= $(60 \text{ M}\Omega)(0.5 \text{ pF}) = 30 \,\mu\text{s}$

On the other hand, the input potential to the amplifier also changes, and its output potential is raised. This lifts the gate of the NFET 'instantaneously', and the gate is lowered slowly according to the decay time constant τ_{postamp} as C_{decay} discharges.

The current through the NFET, while its gate is open, is mainly determined by R_{DD} :

$$i_{\text{postamp}} = \frac{V_{\text{DD}} - v_{\text{postamp}}}{R_{\text{DD}}}.$$
(B.4)

The 'resistor' R_{postamp} in Figure B.3 requires further explanation. It is in fact a more complex circuit of transistors with a Thévenin equivalent resistance $R_{\text{Th}} = R_{\text{postamp}}$. Figure B.5 shows the schematic diagramme of this transistor circuit. The current divider thereiN provides DC feedback to the postamplifier and acts like a resistor. For noise-related reasons, we implement the postamplifier not as an operational amplifier, but as a 'single input amplifier'. The *v*-*i* switch is normally on, allowing 5 μ A of current through. The current divider divides it by 125 times, giving an output of 40 nA, which is balanced by -40 nA from the current source when the circuit is quiescent. Due to this quiescent current, the input to the shaping amplifier is normally at:

$$v_{\text{postamp}} = V_{\text{DD}} - R_{\text{DD}} i_{\text{postamp}}$$

= 5 V - (500 k Ω)(5 μ A) = 4.75 V.

However, the ramping of the postamplifier input produces a DC output, making the shaping amplifier input slightly below 4.75 V instead. This voltage is 'remembered' by the subcircuit in Figure B.6, connected across the postamplifier between the A and B terminals in Figure B.5. The left amplifier Figure B.6 has small transconductance

 $Deltai_{out}/\Delta v_{in}$. The capacitor at the input to the right amplifier stores the value of the quiescent voltage.

B.1.3 The shaping amplifier

Figure B.7 shows a block diagramme of the shaping amplifier. The shaping amplifier differentiates the step-decay profile of v_{postamp} , and adds the resulting derivative profile to a DC level at $V_{\text{reference}}$ to keep the voltage range within the bound [GND, V_{DD}].



Figure B.5: Details of the transistor circuit abstracted as R_{postamp} .



Figure B.6: Voltage 'memory' connected across the postamplifier.



Figure B.7: The HEFT shaping amplifier.

B.1.4 The discriminator

Figure B.8 shows a block diagramme of the discriminator. The discriminator compares v_{shaped} and $V_{\text{threshold}}$, and triggers if $v_{\text{shaped}} < V_{\text{threshold}}$. Because the quiescent value of v_{shaped} is $V_{\text{reference}}$, the real threshold is given by the difference of two input voltages into the ASIC:

True threshold =
$$V_{\text{reference}} - V_{\text{threshold}}$$

 $\approx 2.5 \text{ V} - 1.0 \text{ V} = 1.5 \text{ V}$

This corresponds to 8–10 keV in photon energy.

B.1.5 The sampling capacitor bank

Figure B.9 shows the schematic diagramme of the bank of sampling capacitors. Because $\frac{\tau_{\text{integration}}}{\tau_{\text{postamp}}} = \frac{1\,\mu s}{30\,\mu s} \ll 1$, we can approximate the input current i_{postamp} to be constant over any integration period, and that the formula for i_{postamp} in Equation B.4 holds for the roughly eight samples measured after



Figure B.8: The HEFT discriminator.



Figure B.9: The HEFT sampling capacitor bank.



Figure B.10: The HEFT read-out amplifier.

an event trigger. Then, the charge stored in a sampling capacitor is:

 $q_{\text{sampled}} = i_{\text{postamp}} \tau_{\text{integration}}$

Upon read-out, the sampling capacitors each gives an output voltage:

$$egin{array}{rll} v_{ ext{sampled}}&=&rac{q_{ ext{sampled}}}{C_{ ext{sample}}}\ &=&rac{i_{ ext{postamp}} au_{ ext{integration}}}{C_{ ext{sample}}}. \end{array}$$

B.1.6 The read-out amplifier

Figure B.10 shows the schematic diagramme of the read-out amplifier. Upon read-out, the sampling capacitors are discharged to a read-out capacitor. Equating the amount of charge transferred, we have the gain of the read-out amplifier stage:

$$q_{\text{sampled}} = C_{\text{sample}} v_{\text{sampled}} = C_{\text{readout}} v_{\text{readout}}$$

$$181$$

$$G_{\text{readout}} = \frac{v_{\text{readout}}}{v_{\text{sampled}}} = \frac{C_{\text{sample}}}{C_{\text{readout}}}$$

$$= \frac{1 \text{ pF}}{0.5 \text{ pF}} = 2$$

B.1.7 Analogue to digital conversion

The ADC chip we select for HEFT is a 5 V to 12-bit converter, which encodes output values from 0 to 4095 $(2^{12} - 1)$. This determines the output resolution to be 5 V/4096 ≈ 1.22 mV per step.

B.2 Conversion between quantities at different stages

Putting everything together, we can convert amongst signals at various stages of the HEFT amplifier and read-out circuit. The following are a few examples.

B.2.1 Signal amplitude

A 60 keV photon from an Am-241 radioactive source deposits $60 \text{ keV} / 4.64 \text{ eV} = 12\,931$ electron-hole pairs in a CdZnTe detector (see Table 3.1). The corresponding change in the preamplifier output is

$$\begin{aligned} \Delta v_{\text{preamp}} &= \frac{12\,931e}{C_{\text{preamp}}} \\ &= \frac{12\,931 \times (1.602 \times 10^{-19}\,\text{C})}{40 \times 10^{-15}\,\text{F}} = 5.2 \times 10^{-2}\,\text{V} = 52\,\text{mV}. \end{aligned}$$

Similarly, a 122 keV photon from a Co-57 radioactive source alters v_{preamp} by slightly more than twice this amount, $\Delta v_{\text{preamp}} = 105 \text{ mV}$, which is about 5% of the 2 V dynamic range.

B.2.2 Leakage current

Due to the presence of leakage current, the quiescent level of v_{postamp} is slightly below V_{DD} . This causes a nonzero current

$$i_{\text{postamp}} = \frac{V_{\text{DD}} - V_{\text{postamp}}}{R_{\text{DD}}}$$

where

$$V_{\rm DD} - v_{\rm post\,amp} = -R_{\rm p\,ost\,amp} \left(\frac{C_{\rm postamp}}{C_{\rm preamp}}\right) i_{\rm leakage},$$

according to Equations B.3 and B.4. From the raw samples, we see that this quiescent level is about 500 ADC channels when the integration period is set to $\tau_{\text{integration}} = 2 \,\mu\text{s}$. Working backwards, we have:

$$v_{\text{readout}} = 5 \text{ V} \left(\frac{500}{4096}\right) \approx \frac{5 \text{ V}}{8}$$

$$\begin{split} v_{\text{sampled}} &= \frac{v_{\text{readout}}}{G_{\text{readout}}} = \frac{5 \text{ V}/8}{2} = \frac{5 \text{ V}}{16} \\ V_{\text{DD}} - v_{\text{postamp}} &= i_{\text{postamp}} R_{\text{DD}} = \frac{C_{\text{sample}} v_{\text{sampled}} R_{\text{DD}}}{\tau_{\text{integration}}} \\ &= \frac{1 \text{ pF}\left(\frac{5 \text{ V}}{16}\right) 500 \text{ k}\Omega}{2 \, \mu \text{s}} = \frac{5 \text{ V}}{64} \\ -I_{\text{leakage}} &= (V_{\text{DD}} - v_{\text{postamp}}) \frac{C_{\text{pramp}}}{C_{\text{postamp}}} \frac{1}{R_{\text{postamp}}} \\ &= \frac{5 \text{ V} 40 \text{ fF}}{64} \frac{1}{3 \text{ pF}} \frac{1}{60 \text{ M}\Omega} \approx 17 \text{ nA}, \end{split}$$

which is on the same order as expected from our pre-bonding leakage current measurement (see Section 3.8.3).

B.2.3 Saturation of the preamplifier

When will leakage current saturate the preamplifier? In other words, how much leakage current can the preamplifier tolerate? In the absence of an event $(i_{signal} = 0)$, the maximum tolerable leakage current saturates the preamplifier at the same time as one reset period:

$$v_{ ext{preamp}}(t=t_{ ext{reset}}+ au_{ ext{reset}})=V_{ ext{preamp,max}}$$
 .

Substituting these quantities into Equation B.1, and assuming a constant leakage current $i_{\text{leakage}}(t) = I_{\text{leakage}}$, we have,

$$2.0 V = -\frac{1}{40 \text{ fF}} I_{\text{leakage}} \tau_{\text{reset}}$$
$$-I_{\text{leakage}} = 80 \text{ pA}.$$

If $|I_{\text{leakage}}|$ exceeds this value, then the preamplifier saturates before it resets. However, as v_{preamp} approaches $V_{\text{preamp,max}}$, the dynamic range decreases linearly with time, which limits the detector's ability to measure photons with high energies. Thus, in the presence of a signal, say, of 122 keV, the maximum tolerable leakage current is further decreased by 5%: $|I_{\text{leakage}}| < 80 \text{ pA} \times (1-5\%) = 76 \text{ pA}$.

Appendix C

Processing signals from the HEFT detectors: the details

This appendix explains the implementation details of the HEFT data reduction software routines [written in the Interactive Data Language (IDL), an array-based interactive computing environment widely used within the astronomy community]. It originally appeared as an internal memorandum I wrote and distributed within the HEFT collaboration in 1/2006.

In order to understand the HEFT data reduction package, it is important to understand the physical process through which a photon turns into a pulse height number, as well as all the systematics and noise components along this physical process. Only with this knowledge can one appreciate the rationale behind all the painstaking data reduction steps.

C.1 Implementation of hardware

Appendix B describes the HEFT amplifier and read-out circuit in terms of its circuit parameters and time constants. Here, I briefly reiterate that description, but with the emphasis on the evolution of a signal through the circuitry.

The ASIC circuitry consists of several subcircuits connected mostly in series (with a fork that will be further explained). The first stage, the preamplifier, is essentially an integrator circuit. When events are absent, the input to the preamplifier is a constant DC current; the output is a voltage ramp (with negative slope), going from 4.5 V to 0 V. Photon events appears in the input as impulse functions superposed above the DC level, and appears in the output as a step function on top of the ramp. The entire circuit is reset once every 1.4 ms to prevent the ramp from saturating the signal line. The reset takes $300 \,\mu$ s in MISC version 8j, and is substantially shorted in 8o, the flight version of 2004 and 2005.

The second stage, the postamplifier, performs two functions: 1. It differentiates the input voltage

ramp signal into a step decay output, with a $30 \,\mu s$ decay lifetime. 2. It converts the signal from a voltage in the input to a current in the output.

After the second stage, the signal line forks into several different branches, two of which will be explained here. One branch goes into a simple comparator circuit (peak detect), which compares the postamplifier output level with a reference level (discTh); its output is a digital signal, the trigger, which is on iff the postamplifier output is above dTh.

Another branch of the postamplifier output is channelled into a bank of 16 sampling capacitors through a multiplexor. The MUX connects the postamplifier output to one sampling capacitor at a time; when connected, the postamplifier output current charges the sampling capacitor for a sampling period of 1 μ s, thus recording the integrated current within the sampling period into the capacitor. Then, the next capacitor in line is reset (discharged), and the connection switched to it to record the integrated current in the next sampling period. The MUX cycles the connection around the 16 sampling capacitors, thus creating a record of the current in the previous 16 μ s at any time.

Signals from an event is registered in the following manner—when the trigger subcircuit senses an event, it sets its output to the on state. This digital line is read by the microprocessor (the 'module MISC'). When the module MISC senses a positive trigger at any pixel within the hybrid detectors, it continues the cycling of the sampling capacitors for eight more sampling periods. After that, it freezes all operation in the ASIC. At this point, the sampling capacitors in the triggered pixel(s) contain a record of its postamplifier output 8 periods before the event, and 8 periods after the event. Then, while sampling operations are still frozen, the MISC inspects the trigger signals at all pixels to determine the location of the event. It then makes a second pass through the pixels, this time reading out the 16 samples from all triggered pixels, as well as from all neighbours of triggered pixels. A neighbour is defined as a pixel sharing either an edge or a corner with the triggered pixel.

C.2 Systematics and noise components and their corrections

In an ideal world without noise, without counting statistics, where amplifier parameters at all pixels are identical, signal thresholds can be set at epsilon above zero, sampling period is infinitesimally fine, and the gondola pointing is perfect, the construction of an event list will be very simple one converts the coordinates of the triggered pixels to celestial coordinates of an event using a time-constant formula, and one obtains the energy of the event by finding the amplitude of the stepdecay output of the postamplifier (e.g., by calculating the difference in the sample level immediately before and immediately after the rise, and converting this level difference to an energy using another time-constant formula). However, in reality, this basic scheme has to be extensively augmented.

C.2.1 Truly random (thermal) noise

Due to the presence of electronic noise in the samples recorded by the capacitors, there is an uncertainty in the amplitude of the step-decay function, calculated as the difference between any sample before and any sample after the step. The standard solution to increase the signal-to-noise ratio in such a scenario is by averaging multiple samples before, and multiple samples after, the step, so that the noise level is reduced by \sqrt{N} , where N is the number of samples averaged, while the signal level remains unchanged. Note that even the samples recorded by the capacitors themselves are, in fact, averages—charge integrated over the sampling period; viewed in this manner, our averaging scheme is simply a natural extension of the functonality of the hardware. In performing this averaging, however, one has to be cautious not to include samples too close to the rising edge of the step, as these samples do not represent the true level before or after the rise. Experience has shown that averaging the first six and last six samples of the 16-sample sequence for the levels pre- and post-rise, respectively, works. This is our baseline scheme, as implemented in the current analysis routines.

Recently (mid 2005), the designer of the HEFT ASIC, Walter R. Cook (WRC hereby), has done additional experiments to understand the exact waveform of this step-decay. He found that,

- 1. The true waveform is actually not a simple step-decay, but rather, two step decays with different times-of-rise and decay time scales superposed together. The first step decays quickly relative to the sampling period, and the second step follows one sample later, and decays slowly $(RC = 30 \,\mu\text{s})$. All this complexity is masked in our baseline scheme, as we are ignoring the middle four samples.
- 2. For signals produced by monoenergetic photons, the amplitudes of the steps are very consistent (i.e., the variance is low), while the decay rate varies more. In other words, the later samples after the step are noisier than the samples immediately after, and there is potential spectral improvement by throwing away the later post-step samples. WRC has been averaging only three samples after the step, but this has not been implemented in this analysis package yet. (This change, by itself, involves only changing the value of a symbolic constant in the analysis code; however, because it occurs very early in the data processing sequence, additional changes will propagate a very long way down, and we do not yet have the man-hours to experiment with this and verify every step of the propagation yet.)

C.2.2 Time of rise

The random nature of the photon arrival time, coupled with the finite width of the sampling period, together pose another problem. Random photon arrival time propagates to become random time of rise, through the workings of the electronic circuit as described above. Although there is the peakdetect circuitry to identify the time of rise, it can only do so to the accuracy of one sampling period $(1 \ \mu s)$. This translates to a $1 \ \mu s$ uncertainty in the position of the step on the time axis. Nothing would be wrong if the signal at hand were a simple step, but because the signal is a step-decay, the levels of the post-rise samples are dependent on the time elapsed from the time of rise to the time a sample is recorded. For monoenergetic photon signals, if the time of interaction is early within a sampling period, the post-rise sample levels will be lower, as more time has elapsed from the time of rise to the time the post-rise samples are taken. Conversely, if a photon interacts late within a sampling period, the said elapse time is short, and the post-rise sample levels appear higher. This phenomenon results in a variation in the computed pulse height, with strong correlation with the time of rise. Proper correction is needed.

Fortunately, the time of rise can be inferred from the level(s) of the middle sample(s) in a 16sample sequence. Because the samples are integrals of the postamplifier current over time, an early rise results in the middle sample being higher, while a later rise does the reverse. Sometimes, when the rise is very close to a boundary between two sampling periods, the levels of both samples are affected. Due to this latter observation, as well as for noise reduction and robustness, we have devised a measure of the time of rise as an average, just as we did for the pulse height. Here, we use the change in the sample levels from Sample[i] to Sample[i + 1] as a weight, to weigh the time index midway between the two samples, (i + 0.5). Then, we average these weighted time indices at the vicinity of the rise, which gives our measure of the time of rise:

$$t_{\text{rise}} = \sum_{i=5}^{9} (i+0.5) (\text{Sample}[i+1] - \text{Sample}[i]).$$

This works satisfactorily, even when some of the level changes are negative (due to noise or any pathologic signal).

With a time-of-rise measure and a raw pulse height for every signal, an almost linear relation can be plotted for monoenergetic events at any given pixel. In our basic calibration scheme, a linear relation is fit to calibration data, and is subsequently used to correct for unknown pulse heights. The fit employs the IDL routine ladfit.pro, which does a robust absolute error minimization.

The amplitude of this correction depends on the amplitude of the rise, and thus the energy of the unknown photon signal. We approximate the latter by the amplitude of the raw pulse height, which is close enough for our purpose, and scale the multiplicative calibration factors linearly. There are at least two potential problems with this scheme:

- 1. Because the gain of the amplifiers are not exactly linear (see below), this scaling should actually not be linear either. We have not assessed how important this is, but this is a second-order correction, a correction on a correction.
- 2. Zero pulse height does not actually correspond to zero energy (see below). Thus, for low-energy

signals (i.e., low-amplitude steps), the uncertainty is larger. This has been seen as a potential problem below about 20 keV, but nothing additional to the above has been done to address it yet.

C.2.3 Postamplifier gain

A third issue has to do with the postamplifier gain at each pixel, and their differences. The gain of the postamplifier is determined by the resistance in between its drain and the +5 V supply, V_{DD} . Due to finite accuracy in VLSI fabrication, component size varies, and so does resistance. Thus, monoenergetic photons captured at different pixels produce step-decay signals of slightly different amplitudes. To correct for such differences, electronic pulses simulating monoenergetic photons are sent into each pixel, their signals recorded and processed. A spectrum is made of all signals from each pixel, which shows a single spectral line at the intended energy, convolved with the amplifier gain. The amplifier gain is thus measured for each pixel, and applied as a correction to photon data. However, in fact, this has been a simplified picture so far. Kristin Kruse Madsen has studied the nonlinearity of the amplifier gain, and has found its magnitude to be non-negligible for the purpose of HEFT/NuSTAR. Thus, instead of electronic pulses of a single energy, we sent pulses of five different (but known) energies to each pixel, producing five spectral lines for each pixel. Their centroid channels in the spectra are recorded, and the channel-energy relation is fit to a quadratic for each pixel.

C.2.4 Temperature variation of gain

Yet another complication about the amplifier gain is its variation with temperature. We thus repeated the same calibration procedures at three different temperatures $(-15, -5 \text{ and } +5^{\circ}\text{C})$ bracketing the expected operating temperature of the instrument (although taking data at only three energies but for -5°C). A linear relation is fit to the temperature-pulse height relation.

C.2.5 'Detector gain'

As if the list is not long enough, here are some more issues with the pulse height values as measured above—

1. The size of electronic pulses simulating photons during calibration depends on the size of a resistor in the pulse-generating circuit (the pulser) that has a copy at each pixel. This resistor itself is also proned to variation due to the VLSI fabrication process (although not as much as the resistor controlling the postamplifier), causing the electronic pulses to vary in energy from pixel to pixel.

2. Inhomogeneity of detector material, nonuniformity and defects, all affect the charge-collecting efficiency of the detector material. (CdZnTe is much poorer than Si or Ge in these aspects.) Thus, even if the read-out circuitry is perfect, there will still be variations in the measured pulse height for monoenergetic photons interacting at different pixels.

Ultimately, calibration with a radioactive source with known-energy lines is the only way to ensure that the pulse height values are correct. In our case, this is done using Am-241 sources, with its 59.54 and 13.90 keV lines, both of which are far enough from crowded fields of spectral lines to be identified unambigiously. A flood source is positioned above each focal plane, data collected, and spectra generated for each pixel. The two line positions are measured, and a final linear correction is made to the pulse height.

C.2.6 Zero energy offset

In reality, we are unable to collect enough counts at 14 keV to (1) see the 14-keV line and (2) to clearly separate it from the bunch of lines at 17 keV or so. As there are over 2000 pixels on each focal plane, the maximum count rate is about 20 counts/s (40 with some adjustment of the MISC serial output line), which means 100 s / (1 count/pixel), 50 000 s / (500 counts/pixel), just to get a good measurement of the 60-keV line, and that is $50\,000\,\mathrm{s} = 14$ hours of data collection already. Add in variations due to Poisson statistics, hassle with changing CO_2 and LN2 tanks every few hours and controlling the temperature, chronic fatigue due to the need to take all the pulser data described above, etc. and etc., and taking enough photon data to see the 14-keV line at every pixel becomes impractical—at least, we did not do it. The drawback, then, is that we can only convert the pulse height scale to energy scale by assuming that Channel 0 correspond to zero energy, while in fact, this is absolutely not true. For some complicated reason, when one extrapolates the pulser data down to 0V, the extrapolated pulse height channel does not correspond to zero energy; it seems that when the pulser circuit is turned on, something causes the voltage to be shifted. So, a 'zero-energy correction' is needed. As I shall describe later, we force the ASIC circuitry to read out repeatedly when there is no event at the pixels read out. Using these 'events', we obtain pulse heights at each pixel when there is truly no charge-induced signal in it. These events then serve the purpose of the 14-keV line, and gives us the true zero-energy offset.

As of this writing, the study of the zero-energy correction is still in progress. While scripts have been written to experiment with and to assess this correction, it has not been intergrated permanently into the analysis code yet. The challenge at hand is that a simple energy shift does not seem to solve all the problem. By inspecting multiple lines in the Am-241 spectrum and by comparing the performance of the correction applied to the six HEFT flight detectors, we found that in addition to a shift in the zero-energy position, there also seems to be a change in the slope of the channel to energy scale. Also, the sign and magnitude of the shift are different for detectors in different modules, but are very similar for the two detectors in the same module. This may indicate that the hardware origin of this shift may lie outside the ASICs but in the pulser voltage supply lines (or something like that), although the true identity of the cause remains beyond our reach.

C.2.7 Transfer function: a unified view

At this point, it should be pointed out that the process described above—of transforming a sequence of 16 samples representing a waveform to a single pulse height value—this process can be viewed as a signal passing through a system with a transfer function. In this view, the computation of the pulse height is really just the convolution of the input signal with the transfer function; we assign weights to each of the 16 samples—some negative, some positive, some fractional and some zeroes add the weighted samples together and arrive at a single number as the output. Thus, there is the possibility of a calibration scheme that is simpler that what we have now as a baseline. For instance, one can isolate signals from the Am-241 60-keV photons for each pixel, average them sample by sample to remove the noise and obtain a 'template' waveform for each pixel. This template (with some additional manipulations) will then give a 16-number transfer function for each pixel, without having to deal with all the complexities described above. What we have done in the baseline scheme is to guess at the transfer function using our knowledge of the underlying electronics, noise and photon interaction. It will be an extremely interesting idea to see if the same, or even better, performance can be obtained if the entire system is treated as a black box. (Internet search companies have found that some of their search algorithms that handle data in a similar 'semantics-blind' fashion out-perform those that rely on understanding the semantics, of human languages, in their case.)

The following corrections are more specific to the idiosyncracies of our system, and has less to do with the general idea of sampling a step-decay waveform and a transfer function.

C.2.8 MISC wrap around

The samples from the sampling capacitors are read out of the ASIC circuit as voltage readings, via the read-out amplifier. These analogue voltages pass through an off-chip analogue-to-digital converter, which converts each voltage into a number between 0 and 4095 (12 bits), represented by a sequence of 12 digital voltage levels (0 V or 4.5 V) on 12 output lines. These numbers are not yet the samples though. Interlaced between the 16 sample voltage levels, the read-out amplifier also sends out the voltage at a reference level. The voltage of the reference level is digitized in the same way, and the resulting number subtracted from the number obtained from the sampling capacitor voltage. (In other words, to produce a 16-sample sequence, 32 readings are made from the read-out amplifier and digitized.) Complication arises when the sample level is lower than the

reference level—the sample value after the subtraction is negative. Although the module MISC performs this substraction properly, it truncates the result and retains only the lower 12 bits at some point downstream, before packing the numbers into the data stream. Thus, (readers familiar with 2s complement arithmetics will immediately realise that) negative sample values are converted into very large values by the truncation, something close to the maximum 12-bit value, 4095. It should be noted that for the majority of pixels and under normal circumstances, negative values do not arise. However, for the few pixels with very low sampling capacitor bias points ('baselines') and for pathological signals, these negative values have to be properly treated. This must be done as the first step of the calibration process—before any calibration step is taken, but after the raw data stream is 'crushed' into transmission-error free, event-based format. Because the negative values have to do with low capacitor bias points, and because the latter vary with temperature, the size of this correction depends on temperature.

At first glance, one would suspect that these negative-turned-large values would 'alias' with truly large values, resulting in ambiguity and thus an attempt of correction might be futile. Fortunately, upon inspecting the calibration data, we found that:

- 1. While aliasing does seem to occur in a sample value histogram, when one looks at a 2D histogram in (sample value, temperature) space, the aliasing goes away (i.e., varying temperature causes the maximum and minimum sample values to shift, but at any fixed temperature, the maximum value is always less than the minimum value (mod 4096).
- 2. There are features at both the maximum and minimum points that identify these boundaries the number of counts is higher, appearing as a bump in the histogram. This is probably due to true saturation.
- 3. The exact locations of the boundaries vary from detector to detector, as expected, and presumably due to variations in the VLSI fabrication process.

Currently, this is the way we correct for this problem—we find the midpoint between the maximum and minimum (mod 4096) values manually, by eye, by inspecting a graph of sample values against temperature. (The 'midpoint' is represented by a slanted line going through the empty area between the two more densely populated regions at the maxima and minima.) This midpoint is then shifted to zero and all numbers are modulo 4096. Due to the last point above, this operation has to be performed individually for each detector hybrid. For HEFT and NuSTAR, this will not be a problem, as there are only six flight detectors (plus spare) in each case. In the event that this technology one day is mass-produced, then some kind of automation will be required.

C.2.9 Capacitor offset ('baseline')

We found that even when the capacitor samples contain no signal (e.g., in the first six samples in each 16-sample sequence, in the samples from neighbouring pixels and reference pixels far away from the event and those from pixel circuits unconnected to any detector material, and from forced readouts when there is no event trigger), there are still systematic variations in the sample levels. These systematic variations are correlated strongly with the capacitors themselves—the same capacitor that gives lower levels will always give lower levels, regardless of where it is in the 16-sample sequence (i.e., the order in which the capacitors are read). Also, there is a clear 4-period (i.e., $4 \mu s$) repeating pattern in this variation, and not so clear 2-period and 8-period ones. Note that in the physical layout of the VLSI circuitry, these 16 capacitors are arranged in a 4×4 array. However, they are cycled not always from top to bottom and left to right, but in a daisy chain, according to Jill A. Burnham, the engineer who implemented the VLSI layout of the HEFT ASIC from WRC's schematic diagrammes. (So, the position of each capacitor, relative to other capacitors, cannot explain this systematic.) When the hardware discriminator threshold is lowered, one sees increasing number of noise events. In a histogram of these noise events, tallied by their starting capacitor numbers, a preference is shown to certain starting capacitors, and a similar pattern of a 4-period cycle is also observed. In any case, WRC believes that this noise is caused by some digital, switching signal line, but we still do not know how to remove it. Currently, we can only remove it after the fact by making a correction.

To correct for this capacitor offset (or 'baseline') variation, we repeatedly force read-outs from each pixel when there is no event. (WRC calls these microP events, as the 'trigger' comes not from the trigger signal lines, but from the microP—i.e., the MISC—itself. I have retained the usage of this term in the analysis code, to avoid multiple names and confusions.) These events are then gathered together, the mean of the 16 samples in each sequence removed (more on this later), the (relative) sample values assigned to each capacitor, and an average baseline for each capacitor computed. Early on, the mean of the distribution was computed, with iterations to remove the outlyers to avoid skewing. Currently, the median is used, which is more robust to the effect of outlyers. The removal of the mean in each 16-sample sequence is required because there is a component of the noise that drifts over long timescales, and without removing the mean, the size of the noise distribution is substantially increased and so is the uncertainty in the average computed. These relative baseline levels serve as calibration data; they are substracted from the sample values obtained in true events from the same capacitors in the same pixels. Note that this step must be done before the waveform to pulse height conversion; otherwise, waveforms from a given pixel have a strong dependence on the starting capacitors of the 16-sample sequences, too strong to be ignored.

192

C.2.10 Software discrimination

Currently, the hardware threshold is set at approximately 5 keV. (We are unable to lower it further due to the systematics explained above, but as of 1/2006, WRC is experimenting on this front.) Consider an event where a 60-keV photon is incident close to the boundary between two adjacent pixels. If the amount of charge deposited in the two pixels are such that there is a 56:4 split, then this event will be detected as a single-pixel trigger at the pixel with the 56-keV share, while the true charge-sharing nature of the event is masked. If we ignore this problem, then spectral lines will be smeared on their low-energy sides down to (E-5) keV. To correct for this problem, all neighbouring pixels are read out, and each inspected for residual signals after the fact, in software. We have been calling this procedure software discrimination. For events in which residual signals are found at neighbouring pixels, the charge from multiple pixels is added together to form the pulse height of the event as a whole. Currently, this step is performed somewhat crudely, with an absolute Channel +50 threshold set for all events from all pixels. Note that the optimal threshold has to be determined by the trade-off between throwing away small true signals and accepting large noise, and Channel +50 is by no means optimal, just something that works. The author believes that improvements can be made on this front, but this is just not a high-priority item to be studied by anyone yet.

C.2.11 Common-mode noise

There is an additional component of noise that is either aperiodic or at a frequency much higher than is observable by our sampling, but one that goes up and down across the entire detector hybrid simultaneously. We have been using the term common-mode noise to describe this phenomenon. (Earlier, the term crosstalk was also used, but it is not clear whether this phenomenon is actually the same as crosstalk in the conventional sense of the word in electronics.) To remove this noise, samples at neighbouring pixels containing no signal are averaged and substracted from samples from pixels containing signals. Specifically, this operation is performed for each of the 16 samples in each event, subtracting the mean of the *i*th samples at neighbouring pixels from the *i*th sample at triggering pixel(s). The issues with software discrimination complicates this matter, however—in order to perform the software discrimination well, this common-mode noise correction must be performed first. However, without software discrimination, we do not really know which pixels contain signals, and which not. This is 'Catch-22' circular logic. Our current approach to this problem is to assume at first that all 'corner pixels' contain no signal, use them to correct for common-mode noise, perform the software discrimination, and redo the common-mode noise correction base on the result of the software discrimination. 'Corner pixels' in this case are defined as pixels positioned at corners of the rectangular read-out areas. Because we read out all pixels sharing an edge or a corner with any
C.2.12 Depth sensing

Yet another complication to the common-mode noise correction is the discovery of hole signals detected at the neighbouring pixels. To explain this point, we need to augment the description of the photon A common property of solid-state photon detector materials (Si, Ge, CdZnTe and etc.) is that hole mobility and lifetime are much lower than electron mobility and lifetime. In fact, eV Products intentially fabricate their CdZnTe (for our orders, at least) so that hole mobility is significantly reduced. This means that when a photon interacts within the detector and creates an electron-hole pair, the hole stays there for some time while the initial electron creates an electron cloud and the whole cloud drifts towards the anode plane. As the electron cloud drifts towards the anode plane, it at first induces an electron signal at all anodes—more the closer a pixel is from the cloud—however, as the cloud approaches the anode(s) that it is eventually collected, the electron signal is reversed at non-collecting anodes and eventually cancels out the original induction by the electrons, while the signal at the collecting anode(s) continue to increase until the induced charge equals the charge in the charge cloud. As a result, only the collecting pixels 'see' the electron signal only. All this happens in about 100 ns, much shorter than our sampling period, and the dynamics is thus invisible. However, in contrast, the hole mobility and lifetime are so low that the hole at the photon interaction point stays there for the duration of our sampling, and the hole signal is still seen by the neighbouring pixels while the samples are recorded. By measuring this hole signal (as apparent in the pulse heights from neighbouring pixels only), one can perform depth sensing—the magnitude of the hole signal depends on the distance of the hole from the anodes, and thus the photon interaction point. This is advantageous in cutting back background photons, as background photons come in from all directions, while celestial photons through the optics enter the detectors only from the top through the cathode, and are thus preferrentially located near the cathode. Applying the same reasoning in the orthogonal direction, one can argue that the 'XY' position of the photon interaction point on the focal plane may also be detected more finely than to the pixel scale by comparing the hole signals induced at neighbouring pixels in various directions. Obviously, if one is able to deduce the focal plane position of the photon further, then this will affect the common-mode noise correction as well, because not all neighbour pixels are equal anymore.

C.3 The resulting data reduction sequence

The following is the final list of all the corrections, in the proper order as they have to be applied, with this order governed by everything described above. In parentheses are the IDL routines in the analysis package that deals with each step. Note that this is not an exhaustive list of routines, but rather, a list of high-level routines arranged by function; low-level and auxiliary routines are not included.

I.	Cru	to remove data error in t	to remove data error in telemetry, bad events, etc., and						
		to format events in fixed	length records. (crush/crush.pro)						
II.	Pul	se height computation	(common/calibrateEvents.pro,						
			<pre>common/calibrateFiles.pro)</pre>						
	1.	Fix wrap around at Channel 4096	(common/fixWrapAround2.pro)						
	2.	Capacitor offset (baseline) correction							
		(obtain calibration data:	baselines/baseline_crushanal.pro,						
		verify calibration data:	baselines/baseline_crushconfirm.pro,						
		apply calibration data:	<pre>baseline_subtract2() in common/tools3.pro)</pre>						
	3.	Common-mode noise correction using co	rner pixels as reference						
			$(\texttt{common/subtract_common2.pro})$						
	4.	Compute raw pulse height	(ph_reconstruct() in common/tools3.pro)						
	5.	Software discrimination	(done within calibrateEvents.pro)						
	6.	Common-mode noise correction using so	oftware-discriminated pixels						
			$(\texttt{common/subtract_common2.pro})$						
	7.	Compute pulse height post-software disc	rimination						
			(ph_reconstruct() in common/tools3.pro)						
	8.	Time of rise correction							
		(obtain calibration data:	pulser/risetime_anal.pro,						
		verify calibration data:	<pre>pulser/risetime_confirm.pro,</pre>						
		apply calibration data:	risetime_calc() in common/tools3.pro,						
			risetime_correct() in common/tools3.pro)						
	9.	Amplifier gain correction							
		(obtain calibration data:	<pre>pulser/gaincoeffs_anal.pro,</pre>						
			auxiliary: pulser/fit_pulsers.pro,						
			auxiliary: pulser/gather_coeffs.pro,						
		verify calibration data:	<pre>pulser/gaincoeffs_confirm.pro,</pre>						
		apply calibration data:	<pre>pulser/correct_gain.pro)</pre>						
	10.	Temperature correction of amplifier gain							
		(obtain calibration data:	<pre>pulser/tempcoeffs_anal.pro,</pre>						
			auxiliary: pulser/fit_pulsers.pro,						
		verify calibration data:	<pre>pulser/tempcoeffs_confirm.pro,</pre>						
		apply calibration data:	<pre>risetime_correct() in common/tools3.pro)</pre>						

11. 'Detector gain' (collecting efficiency) correction

			(done within calibrateEvents.pro)
	12.	Zero-energy point correction	(not yet implemented as a permanent routine)
	13.	Reconstruction of charge-split events	(spectrum/addPulseHeights.pro)
	14.	Event position computation	(aspect/pixel2celestial.pro,
			auxiliary: aspect/aspectLookup.pro)
III.	Εv	ent filtering and data product generation	

energy filter, time filter, depth sensing, triggered-pixel count, starting capacitor (systematics), etc. (background/goodEventFlag.pro)

195

Appendix D HEFT science data formats

This appendix documents the format of the HEFT science data at various stages of data reduction and analysis. The raw data exits the central MISC as an asynchronous, variable-length bit stream. In flight, the flight control computer receives this bit stream and sends it to the ground, unaltered, via radio telemetry. Computers of the ground station equipments save a copy of the raw data before further processing. The standard HEFT data reduction scheme then cleans the raw data of transmission errors and bad data, and produces a fixed-length version of the data we call the 'crushed' format. Further reduction proceeds afterwards, with the final data encoded as standard FITS event lists.

Section D.1 below describes the telemetry format; it originally appeared as an internal memorandum I sent to the HEFT detector team on 2004- 3- 3. Section D.2 describes the 'crushed' format; it was another internal memorandum written and archived on 2006-11-22. The FITS format for HEFT has yet to be finalized; I present an interim design document in Section D.3; it was written and archived on 2006- 5-19.

D.1 Telemetry format

The HEFT science data stream is segmented into 'frames'. Each frame contains a frame header, followed by module data blocks, one block for each existing detector module. Each module data block repeats the same hierachy as in a frame—it contains a module header, followed by event records, one for each event.

The HEFT MISCs, which produce the science data stream, are 24-bit systems. Thus, the science data stream is also formatted into 24-bit words, big-endian (see Note 1 below).

Frame header

Word order in block	Most least sig.bit of word	Size in bits	Content / Description
1	230	24	$r_{r_{2}}$
2	230	24	ditto
3	230	24	Frame count. Resets to zero when MISC reboots. Increments +1 unit per frame.
			GPS '1 pulse per second' (1pps) signal receipt times Sample 1:
4	230	24	Frame count (see Word 3).
5	2312	12	 1.25 msec count, 0799, resets to 0 when frame count increments; +1 when bits 110 pass through 0.
	110	12	 1.8432MHz count 1.25 msec/2¹² steps, decrements from 2303 to 0, extra 1/2 cycle at 0, 1/2 cycle shorter at 2303, resets to 2303 when 1.25 msec count increments.
67			Sample 2: Repeat of format in Words 45.
89			Sample 3: Repeat of format in Words 45.
1011			V Sample 4: Repeat of format in Words 45.
			GPS/QNX time codes.
1213 14	230	24	Receipt time: Repeat of format in Words 45. Upper 24 bits of a 32-bit float in IEEE 754 spec, representing the second-of-week for GPS times, and the second-of-hour for QNX times.
15	2316 150	8 16	<pre>Lower 8 bits of the above. If value equals the ASCII codes of 'QX' (0x51,0x58), then this 'GPS time code' should be interpreted as a QNX time. Utherwise this is the week number since 1980-1-6</pre>
16	230	24	<pre> Upper 24 bits of a 32-bit float in IEEE 754 spec, representing the UTC offset for GPS times. OR</pre>
	2316	8	Hour (023) of QNX time.
	15 8	8	Day of month (024) of QNX time.
	7 0	8	Month (112) of QNX time.
17	2316	8	<pre> Lower 8 bits of the UTC offset above for GPS times, or the year (mod 100) of QNX time.</pre>
	150	16	V Unused; all zeros.
			Housekeeping data: 12-bit ADC readings, 0(2^12)-1.
18	2312 11 0	12 12	Dome pressure Printed circuit board (PCB) A temperature
19	2312	12	Photomultiplier tube (PMT) A temperature
	11 0	12	PCB C temperature
20	2312	12	PMT C temperature
	11 0	12	PCB B temperature

Word order in block	Most least sig.bit of word	Size in bits	Content / Description
21	2312	12	PMT B temperature
	11 0	12	Analogue board temperature
22	2312	12	Unused
	11 0	12	Fan/heater temperature
23	2312	12	DC->DC transformer temperature
	11 0	12	Second-fan temperature
24	2312	12	Shield A exterior temperature
	11 0	12	Shield B exterior temperature
25	2312	12	Shield C exterior temperature
	11 0	12	Mid-flange temperature, equidistant from shields
26	2312	12	Inlet baffle temperature
	11 0	12	+5V power supplymeasured voltage
27	2312	12	-5V power supplymeasured voltage
	11 0	12	+12V power supplymeasured voltage
28	2312	12	-12V power supplymeasured voltage
	11 0	12	+2.5V power supplymeasured voltage
29	2312	12	-28V power supplymeasured voltage
	11 0	12	DAC referencemeasured voltage
30	2312	12	+4.5V referencemeasured voltage
	11 0	12	Unused.
31			Unused.
32			Unused.
33			V Unused.
34+			[3 module data blocks follow.]

Module data block

Word order	Most least	Size in	Content / Description				
in	sig.bit	bits					
block	of word						
1	230	24	Delimiter, 0xF0 0xA0 0xA0 for Module A,				
			$0xF0\ 0xA1\ 0xA1$ for Module B, $0xF0\ 0xA2\ 0xA2$ for Module C				
2N			Residual event data belonging to the previous frame.				
			This entry can be absent.				
(N+1)(N+8)			Module header (see separate listing).				
(N+9) and on			[O or more event records follow.]				

Module header

Word	Most	Size	Content / I	Descri	ptior	ı
order	least	in				
in	sig.bit	bits				
block	of word					
						-
1	230	24	Delimiter,	0xF0	0xF1	0xF2

198

Word order	Most least	Size in	Content / Description
ln	sig.bit	bits	
plock	of word		
2	230	24	Frame count by module MISC
			(different from frame count in frame header)
3	230	24	Live time, 0921600, 1/921600 sec per step
4	2316	8	detector event count
	150	16	Shield event count
5	17	1	Shield reject flag; 1 when shield veto is on.
	16	1	Test pulser flag; 1 when pulser is on.
	15 8	8	Accepted event count (see note 2).
	7 0	8	Rejected event count (see note 2).
6	2316	8	Photon event count
	15 8	8	Shield-detector coincidence event count
	7 0	8	Microprocessor-generated event count
7	2316	8	Pulser event count
	15 8	8	Side A event count
	7 0	8	Side B event count
8	2316	8	Both side event count
	15 8	8	Neither side event count
9+			<pre>[0 or more event records follow.]</pre>

Event record

Word order in block	Most least sig.bit of word	Size in bits	Content / Description
 14			Event herder (and listing heley)
14			Pivel data:
			Pixel 1
5	2318	6	Row ID, 023
	1712	6	Column ID, 047
	11 0	12	Sample O
6	2312	12	Sample 1
	11 0	12	Sample 2
13	2312	12	Sample 15
			Pixel 2
	11 6	6	Row ID
	5 0	6	Column ID
14	2312	12	Sample 0
	11 0	12	Sample 1
			V [Pixel data continues in similar fashion.]
N	11 0	12	If an odd number of pixels are read out,
			the pixel data ends in the middle of a word.
			In this case, the last 12 bits of the last word
			in the event record is unused (all zeros).

Event header

Word order in block	Most least sig.bit of word	Size in bits	Content / Description
1	23 0	24	Delimiter, OxFO OxAA OxAA
			Event time (see GPS receipt times for format)
2	2312	12	1.25 msec count
	11 0	12	V 1.8432MHz count
3	2316	8	Number of pixels read out in this record
	15 0	16	Time since reset, unit=1/(7.3728 MHz)
4	2323	1	1 for 'Side A' hybrid, O for 'Side B' hybrid
			Event tags: 5 bits each, contains timing information,
			nonzero == assertive
	2218	5	test pulse(s)
	1713	5	microprocessor-triggered event(s)
	12 8	5	V shield event(s)
	7 4	4	Second lower 4 bits???
	3 0	4	Starting capacitor ID
5+			[Pixel data follows.]

Notes:

- The HEFT MISC system is only big-endian starting with the first central-MISC/module-MISC system, circa July, 2004. Prior to that, the system was actually small-endian, up to version MISC8j.
- 2. The positions of the accepted and rejected event counts were swapped on 2004- 8-15. This document shows the encoding after the change.
- 3. The following formulae convert ADC readings to pressure, temperature and voltages:

Dome pressure	=	$[(ADC\text{-}reading/4096) \times 4.46 V - 1 V] \times 7.5 psi/V$
Temperatures	\approx	$[1 - (ADC\text{-reading}/4096)] \times 450 \mathrm{K/V}$
+2.5 V supply voltage	=	$(ADC\text{-reading}/4096) \times 4.46 \mathrm{V}$
$+5\mathrm{V}$ supply voltage	=	$(\text{ADC-reading}/4096) \times 4.46\text{V}/0.594$
$+12\mathrm{V}$ supply voltage	=	$(\text{ADC-reading}/4096) \times 4.46\text{V}/0.239$
$+28\mathrm{V}$ supply voltage	=	$(\text{ADC-reading}/4096) \times 4.46\text{V}/0.109$
$-5\mathrm{V}$ supply voltage	=	$[(ADC-reading/4096) \times 4.46 V - (+5V voltage) \times 0.794]/0.206$
$-12\mathrm{V}$ supply voltage	=	$[(ADC-reading/4096) \times 4.46 V - (+12V voltage) \times 0.647]/0.353$

DAC and +4.5 V references will always be at ADC channel 4095, due to the way that the electronics was implemented.

D.2 'Crushed' format

D.2.1 Introduction

This memorandum documents the format of data files produced by the 'crush' programme for HEFT. The crush programme converts raw data streams output by the 'Central MISC' (and often sent through telemetry and RS422 links), which are in 1-second frames, into 'crushed' data files, which are organized in fixed-length event records and pixel records. Apart from discarding data with encoding errors, no information is lost in the crush process. Subsequent, lossy, data reduction and analysis should be performed on the crushed data, which should be error-free and well-formatted.

Each crushed data set consists of two files: An event-data file and a corresponding pixel-data file. The event-data file contains value fields specific to each event, while the pixel-data file contains fields specific to each pixel read out. Multiple pixels are read out in each event. Therefore, one event record corresponds to multiple pixel records. All pixel records from the same event are listed consecutively, and all these pixel record groups are listed in the pixel-data file in the same order as the corresponding event records are listed in the event-data file.

As of this writing (11/2006), crushed data files are produced on a Linux PC. Thus, the binary crushed files are in little-endian format.

D.2.2 Crushed event-data format

Each 'event record' is 17 bytes long: -The first byte in each event record contains an integer (nPix) recording the number of pixels read out in this event. nPix ranges from 0 to 31. -The second byte contains two items---The upper 4 bits (0xf0) contain the 'detector ID'. Range: 0--(number of detectors - 1). The lower 4 bits (0x0f) contain the 'sequence number', a.k.a. the 'starting capacitor number'. Range: 0--15. -The next two bytes contain a 16-bit unsigned integer encoding the 'time since reset', in unit of ~0.125 microsec. Caution: the unit of time since reset has changed in the past (circa 2004). -The next two bytes contain a 16-bit integer encoding the following tags: Bits 10--14 (0x7c00): Shield tag Bits 5-- 9 (0x03e0): Pulser tag Bits 0-- 4 (0x001f): Baseline tag The remainder of the event record contains time tags, and is identical to the time tags in a pixel record (see below).

The next 4 bytes encode a 32-bit unsigned long word encoding the last modification time of the raw data file, in unix time.
The next 4 bytes encode a 32-bit unsigned long word encoding the frame count.
The next 3 bytes encode the coarse subframecount in the upper 12 bits, and the fine subframecount in the lower 12 bits.

D.2.3 Crushed pixel-data format

Each 'pixel record' contains 38 bytes: -Byte 0 contains the horizontal (x) coordinate of the pixel. -Byte 1 contains the vertical (y) coordinate of the pixel in the lower 6 bits (0x3f), and the 'position' tag in the upper two bits (0xc0). position == 3: Pixel is an interior pixel in the read-out area. position == 2: Pixel is an edge pixel in the read-out area. position == 1: Pixel is a corner pixel in the read-out area. position == 0: No geometry information. -Byte 2 encodes the 'order' of the pixel in the event. Nth pixel read out for that event: order=(n-1). -Bytes 3--26: The sample (ie, capacitor) values, in sequence. 16 12-bit unsigned integers |--> 24 bytes. The remainder of the pixel record contains time tags, and is identical to the time tags in an event record (see above). -The next 4 bytes encode a 32-bit unsigned long word encoding the last modification time of the raw data file, in unix time. -The next 4 bytes encode a 32-bit unsigned long word encoding the frame count. -The next 3 bytes encode the coarse subframecount in the upper 12 bits, and the fine subframecount in the lower 12 bits.

D.2.4 Additional remarks

The number of pixel records in the pixel-data file should match the sum of all nPix fields in the corresponding event-data file. If one reads both the event-data file and the pixel-data file sequentially (looking up nPix in the event record to determine how many pixel record to read for each event), the event records and pixel records should naturally be in the right order. In the event that this is not true, then one can match the pixel records to each event record by comparing the time tags, which are identical in both kinds of records for the same event. The order tag in the pixel record also specifies the pixel record groups, as the pixel records for each event are listed in consecutive

'order'.

D.3 The HEFT FITS file format

HDUVERS = '1.0.0'

D.3.1 Introduction

Scientific data from the first HEFT flight is being written to FITS files. This memorandum records the header keywords and binary table columns that are present in the HEFT FITS files, as well as their expected values and the rationale for these items, when appropriate. As far as this author can tell, the HEFT FITS format conforms to both the FITS and OGIP standards. However, these standards alone are insufficient to completely define our FITS file format, due to the presence of many instrument-specific information. This document is intended to be a complement to those standards. The hope is that it will be useful to future implementers of HEFT and related missions.

When deciding which keywords in the FITS standard to include for HEFT, I have studied the ASC FITS convention for *Chandra*, and its FITS event files; I have attempted to include in the HEFT convention as much of the standard FITS information in ASC-FITS as possible, while leaving out 'ASC keywords' that are designed only to be interpreted by the DM data model in CIAO. Rather, considerations are made for the HEFT FITS files to be easily interpreted by the IDL astronomy library routines and by SAOimage DS9; this requires nothing more than the FITS/OGIP/HFWG standards. Despite the philosophical difference, the ASC-FITS document is very well written, and should be a model for future revisions of this document.

D.3.2 Header keywords

Mandatory keywords The presence of these keywords is governed by the FITS standard (NOST 100-2.0). If the value of a keyword is always the same for HEFT, then this value is also listed below.

Primary unit		Binary table	e extension	n
Keyword	Constant value	Keyword	Const	ant value
SIMPLE	Т	XTENSION	'BINT	ABLE'
BITPIX		BITPIX		
NAXIS	n	NAXIS	2	
NAXISn				
PCOUNT	0	PCOUNT	0	(Pointers not used.)
GCOUNT	1	GCOUNT	1	(Ditto.)
		TFIELDS	n	

Keywords that should be present in all HDUs The following keywords are either explicitly defined in the FITS standard, or are commonly used in the astronomy community. They were obtained from the 'FITS Keyword Dictionaries' page at the FITS Support Office web site (http://fits.gsfc.nasa.gov/fits_dictionary.html) as of 11/2005.

TELESCOP= 'HEFT'

Because there is only one instrument on HEFT, and we have not given our detectors acronyms, other keywords such as MISSION and INSTRUME are unnecessary. Reference: NOST 100-2.0, Section 5.4.2.2

DATE = 'yyyy-MM-ddThh:mm:SS.sss'

Value should be automatically generated by software. Reference: NOST 100-2.0, Section 5.4.2.1

DATE-OBS= '2005-05-18T19:MM:SS.sss'

Reference: NOST 100-2.0, Section 5.4.2.2

DATE-END= '2005-05-19T19:MM:SS.sss'

Reference: OGIP/93-003, Section 4.1

HDUCLASS= 'OGIP' or 'HEFT'?

According to (my interpretation of) HFWG Recommendation R8, this identifier specifies the data format for the *entire* FITS file. In this sense, the value 'HEFT' is appropriate. However, according to ASC-FITS-2.1.0, Section 2.3, pp. 17, 'HDUCLAS [sic] gives the origin of the hierachy' (of HDUCLASn), and its value is either "OGIP" or "ASC".' If interpreted this way, then we should always use 'OGIP' only, as we have not invented a new HDUCLASn hierachy. We opt for adhering to the HFWG Recommendation.

HDUDOC ='groa.srl.caltech.edu:/u/heft/analysis/cvsroot/analysis/doc/HEFT-FITS'

Reference: HFWG Recommendation R8

If HDUCLASS= 'HEFT', then this document should be specified.

If HDUCLASS= 'OGIP', then event lists can use: OGIP/93-003; images can use: NOST 100-2.0.

HDUVERS = (Current value is at the top of this article.)

Reference: HFWG Recommendation R8. Obsolete according to ASC-FITS-2.1.0, Section 2.3, pp. 18, but no mention elsewhere.

HDUCLASn

Reference: HFWG Recommendation R8. Entries relevant for us are all listed below:

```
HDUCLAS1= 'EVENTS'

HDUCLAS2= 'ALL'

HDUCLAS2= 'ACCEPTED'

HDUCLAS2= 'REJECTED'

HDUCLAS1= 'IMAGE'

HDUCLAS2= 'TOTAL'

HDUCLAS2= 'BKG'

HDUCLAS2= 'NET'

EQUINOX = 2000.0

RADECSYS= 'FK5'
```

HFWG Recommendation R3 requires both EQUINOX and RADECSYS.

CHECKSUM= (Value produced by fits_add_checksum in the IDL astronomy library) DATASUM = (Value produced by fits_add_checksum in the IDL astronomy library)

Reference: Seaman et al. (2002), "FITS Checksum Proposal".

Forbidden from use:

HDUNAME

Chandra includes this keyword in addition to EXTNAME, although their values seem always to be identical. According to the ASC FITS File Designer's Guide, the CIAO data model DM reads and writes HDUNAME. On the other hand, the FITS Keyword Dictionaries page:

(http://heasarc.gsfc.nasa.gov/docs/fcg/common_dict.html) points to a HEASARC page that states the following:

It is recommended that the HDUNAME and EXTNAME keywords should not both be given in the same HDU header, but if they are, then the HDUNAME keyword will have precedence.

As EXTNAME is in the FITS standard, while HDUNAME is not, we decide not to specify it for HEFT.

Keywords for event list (binary table) extensions

Mandatory:

```
EXTNAME = 'EVENTS' Reference: OGIP/93-003, Section 4, pp 5
MJDREF = Reference: OGIP/93-003, Section 4.2, pp 7--8
```

CLOCKCOR= 'YES'	Reference:	OGIP/93-003,	Section	4.2, pp	78
TIMESYS = 'JD'	Reference:	OGIP/93-003,	Section	4.2, pp	7
TIMEUNIT= 'd'	Reference:	OGIP/93-003,	Section	4.2, pp	7
TIMEZERO=	Reference:	OGIP/93-003,	Section	4.2, pp	78
TSTART =	Reference:	OGIP/93-003,	Section	4.2, pp	7
TSTOP =	Reference:	OGIP/93-003,	Section	4.2, pp	7
TIMEREF = 'LOCAL'	Reference:	OGIP/93-003,	Section	4.4.1,]	рр 9
Mandatory for each column (values given later in this article):					
TTYPEn, TFORMn, TUNITn	Reference:	NOST 100-2.0,	, Section	ns 8.1.2	and 8.3.1.
Optional for each column:					
TZEROn, TSCALn	Reference:	NOST 100-2.0,	, Sectior	n 8.1.2	
TDMINn, TDMAXn, TLMINn,	TLMAXn				
	Reference:	HFWG Recommer	ndation H	36	

For columns with sky coordinates only:

TCTYPn

We are using the gnomonic (TAN) projection, though not about zenith, but about the point of observation instead. There is no such thing defined in the FITS celestial coordinates definition paper though. Chandra FITS files claim that they use TAN as well. Reference: "Representations of Celestial Coordinates in FITS"

TCRVLn, TCRPXn, TCDLTn, TCUNIn

Optional keywords The following keywords, while in the FITS/OGIP standards, have different values almost every time a new FITS file is written. As a result, their inclusion into FITS files cannot, and should not, be done by automated routines. There are two options:

1. Do it in 'macros' batch files.

Pros: a record of the file-creating session.

Cons: No guarantee that the file-creating operator will remember to put in all the information.

2. Interactive routines.

Pros: Provides a mindless way for an operator to include all items.

Cons: No automation.

Perhaps a compromise would be to write an interactive routine, whose invocation is optional (thus solves the automation problem), and have batch files calling the FITS writing routines, documenting the session.

OBJECT = [not included] (default)

Should be included only for subset of data corresponding to a specific target object only.

TITLE = [not included] (Paper/proposal title)
OBSERVER= [Fiona A. Harrison]? (PI)

Reference: ASC-FITS-2.1.0, Section 2.2.1, pp 15

ORIGIN = ['The HEFT Collaboration---Caltech/Columbia/DSRI/LLNL']?(institution)
AUTHOR = [\$whoami]
CREATOR = ['IDL Version 5.5 (linux x86)']
HISTORY

Should record calibration files and routines used.

COMMENT

*FILEn?

D.3.3 Column specification for event lists in BINTABLE HDUs

XDR	quantity	TTYPE	TFORM	TZERO	TSCAL	TUNIT	TCRPX/TCRVL/TCDLT/TCUNI
RA,	a[*,0]	x			200.0	pixel	0/0/0.005/deg
Dec,	a[*,1]	У			200.0	pixel	0/0/0.005/deg

Our sky coordinates are handled in degrees, but the FITS and OGIP standards want them to be in pixels. What is even worse is that the author of ds9 will not allow zooming to subpixel level, which means that we cannot simply set x = RA with 1 pixel = 1 deg (or else nothing useful is rendered in ds9). Given the above, one should set the values stored within the FITS files to have a size such that one FITS pixel indeed corresponds roughly to one HEFT pixel (17".12). Because this is less than 5% away from 18'' = 0.005 = (1/200) deg, it makes sense to set the FITS pixel size to 18'', a round number and simple fraction of a degree. Also, this way, we can preserve the native values of the coordinates as they come out of the analysis routines, by utilising the TSCAL keyword. References: Legacy, No 6, pp 36 (no designation)

Note that while some OGIP document explicitly lists the column names 'RA' and 'Dec', they are not automatically recognised by widely-used software such as ds9. Although one can still get coordinates in columns named 'RA' and 'Dec' displayed by explicitly specifying [bin ra,dec], one can do that for any column with any name; in other words, the column names 'RA' and 'Dec' are nothing special and not recognised in reality. Reference: OGIP/94-003

XDR quantity	TTYPE	TFORM	TZERO	TSCAL	TUNIT	TCRPX/TCRVL/TCDLT/TCUNI
detectorID	detecto	r			(unitle	ess)

Reference: Legacy for *BBXRT*, OGIP/94-003 for all others. Alternatives are CCDID (from *ASCA*) and CCD_ID (from *Chandra*, but the HEFT detectors are not CCDs), and DET-NAM (but this quantity is supposed to be a string, not a number).

XDR quantity detector X detector Y	TTYPE TFORM detX (rawX?) detY (rawY?)	TZERO	TSCAL	TUNIT pixel pixel	TCRPX/TCRVL/TCDLT/TCUNI		
{spectra}.sph OR	PI			chan			
{spectra}.sph	energy		1/15.0	keV			
rawPulseHeight	PHA			chan			
neighbourPulseH	eight						
	neighbour_PHA			chan			
	time						
frame count				S			
filemodtime				S			
integer frame co	ount			S			
coarse sub frame count							
fine sub fram	e count						
t-t_reset	time_since_rese	t					

Note that the IDL routine mwrfits will convert all minus signs and spaces to underscores when writing a header string array to file. In other words, while the FITS standard and OGIP recommendations do not forbid the use of minus signs and spaces, they are effectively forbidden by the lack of support in commonly used software. HFWG Recommendation R15 explicitly forbids the use of the minus sign.

XDR quantity TTYPE TFORM TZERO TSCAL TUNIT TCRPX/TCRVL/TCDLT/TCUNI start cap startCap (unitless)

The name of this quantity, while not standard, is well-understood within the HEFT detector team, and there is no similar precedent.

XDR quantity TTYPE TFORM TZERO TSCAL TUNIT TCRPX/TCRVL/TCDLT/TCUNI time of rise t_rise

There is a conscious effort here against the use of the term 'risetime', an unfortunate misnomer circulated in the HEFT detector team; risetime should mean the length of time it takes the signal to go from a low state to a high state, not the time instance at which the rise occurs.

XDR quantity npixadded npixread	TTYPE grade readout	TFORM _grade?	TZERO	TSCAL	TUNIT count count	TCRPX/TCRVL/TCDLT/TCUNI
microP tag microF shield tag shield pulser tag pulser		tag tag tag			(unitles (unitles (unitles	55) 35) 35)

Although more meaningful names can be attributed to these quantities (eg, charged particles instead of shield), there is no well-established standard in place and the names x in x_{tag} have been routinely used by the HEFT detector team, and that is one reason to keep them in the FITS files.

XDR quantity	TTYPE	TFORM	TZERO	TSCAL	TUNIT	TCRPX/TCRVL/TCDLT/TCUNI
temperature	tempera	ture	273.15		K	
event baseline					chan	
burst tag					(unitles	ss)

Unit reference: OGIP/93-001

D.3.4 References

As spelt out by the *Chandra* personnel who performed this same task of defining FITS format, 'there is no good document describing [the HEASARC FITS Working Group's] recommended conventions'. Besides, I have found this to be true for the 'OGIP standard' as well. (Ironically, many FITS files claim to be OGIP-compliant, but I cannot find a definitive OGIP document anywhere!) In light of this, I am listing here all the sources that I have found on the web that contain either authoritative or unique FITS/OGIP/HFWG-related information.

The FITS standard, NOST 100-2.0

```
ftp://legacy.gsfc.nasa.gov/fits_info/fits_office/aa_fits_standard.pdf
```

Also appears in other versions and formats, linked from this page:

http://fits.gsfc.nasa.gov/fits_documentation.html

FITS WCS standard

List of documents:

http://fits.gsfc.nasa.gov/fits_wcs.html

Specifically, the relevant documents for us are:

Representations of World Coordinates in FITS,

Greisen, E. W., and Calabretta, M. R., Astronomy & Astrophysics, 395, 1061-1075, 2002. Representations of celestial coordinates in FITS,

Calabretta, M. R., and Greisen, E. W., Astronomy & Astrophysics, 395, 1077-1122, 2002.

Common keywords

List of documents:

http://fits.gsfc.nasa.gov/fits_dictionary.html

in FITS standard:

http://heasarc.gsfc.nasa.gov/docs/fcg/standard_dict.html

not in FITS standard:

http://heasarc.gsfc.nasa.gov/docs/fcg/common_dict.html

OGIP/HFWG recommendations

List of documents:

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/ofwg_recomm.html

This page links to a whole myriad of files—original Legacy articles, revised Legacy articles, completely Legacy-unrelated articles, etc, in an unsystematic (in fact, chaotic) manner—some with an OGIP/xx-yyy designation, some as HFWG Recommendation Rzz, some with nothing at all. The following attempts to clarify things a bit. (The idea is that each page is linked once and only once below, and all cross-references spelt out.)

Unique (and relevant) documents listed here are:

OGIP/94-003

Guidelines for Defining FITS Formats for Event Lists

Ian M George, Arnold Rots and Koji Mukai

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/

events/ogip_94_003/ogip_94_003.html

Related: Legacy 6, 36.

OGIP/92-007

The OGIP Spectral File Format

Arnaud, George and Tennant.

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ summary/ogip_92_007_summary.html

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ summary/ogip_92_007a_summary.html

Original: Legacy 2, 65.

OGIP/93-003

The Proposed Timing FITS File Format for High Energy Astrophysics Data Angelini Pence and Tennant

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/

summary/ogip_93_003_summary.html

Original: Legacy 3, 32.

OGIP/93-001

Specification of Physical Units within OGIP FITS Files

George and Angelini

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/

summary/ogip_93_001_summary.html

Original: Legacy 4, 57.

OGIP/93-013

Standard Strings for Mission, Instrument, Filter, Detector, & Grating Names for OGIP FITS Files

George and Angelini

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/

summary/ogip_93_013_summary.html

R1 Underscore vs hyphen

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r1.html Original: Legacy 4, 72. OFWG Report 93_001

- R2 Superseded by R8.
- R3 RA and Dec

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r3.html Original: Legacy 4, 72. OFWG Report 93_001

R4 CATIDn

Original: Legacy 4, 72. OFWG Report 93_001

- R5 Same as OGIP/93-001Original: Legacy 4, 72. OFWG Report 93_001
- R6 TLMINn, TLMAXn, TDMINn and TDMAXn

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r6.html Original: Legacy 4, 72. OFWG Report 93_001

R7 CREATOR

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r7.html Original: Legacy 5. OFWG Report 93_002

$\mathbf{R8}$ HDUCLASn

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r8.html http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/hduclas.html Original: Legacy 5. OFWG Report 93_002

 ${\rm R9} \quad {\rm Quality} \ {\rm flags}$

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r9.html

Original: Legacy 5. OFWG Report 93_002

R10 Channel and energy boundaries

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r10.html Original: Legacy 5. OFWG Report 93_002

R11 Exposure times

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r11.html Original: Legacy 5. OFWG Report 93_002

R12 Same as OGIP/93-013

Original: Legacy 5. OFWG Report 93_002

R13 LONGSTRN

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r13.html Original: Legacy 5. OFWG Report 93_002

R14 TSORTKEY

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r14.html

R15 FITS table column naming conventions

http://heasarc.gsfc.nasa.gov/docs/heasarc/ofwg/docs/ofwg_recomm/r15.html

Legacy Journal (where artcles were originally published)

Main page:

http://heasarc.gsfc.nasa.gov/docs/journal/journals.html

List of all online Legacy articles on FITS:

Issue 1 (5/1992): nil

```
Issue 2, pp 65 (11/1992):
```

OGIP/92-007

The OGIP Spectral File Format

Keith A. Arnaud, Ian M. George and Allyn F. Tennant.

http://heasarc.gsfc.nasa.gov/docs/journal/ogip2.html

Issue 3, pp 32 (5/1993):

OGIP/93-003

The Proposed Timing FITS File Format for High Energy Astrophysics Data

L. Angelini, W. Pence and A. F. Tennant

http://heasarc.gsfc.nasa.gov/docs/journal/timing3.html

Issue 4, pp 57 (2/1994):

OGIP/93-001

Specification of Physical Units

Ian M. George and Lorella Angelini

http://heasarc.gsfc.nasa.gov/docs/journal/units4.html

Issue 4, pp 72 (2/1994):

OFWG Report 93_001 (R1, R3–R6)

The Role & Initial Recommendations of the OGIP FITS Working Group

Ian M. George

http://heasarc.gsfc.nasa.gov/docs/journal/ofwg4.html

Issue 5, pp ?? (11/1994):

OFWG Report 93_002 (R7-R13)

OGIP FITS Working Group Report

Ian M. George

http://heasarc.gsfc.nasa.gov/docs/journal/ogip5.html

Issue 6, pp 36 (9/1995):

(No designation)

The Recommended Columns and Keywords for a FITS Event List

W. Pence and I. George

http://heasarc.gsfc.nasa.gov/docs/journal/fits6.html

Issue 7 (6/1998): nil

FITS information scattered elsewhere

FITS Checksum Proposal

Seaman, Pence, and Rots (2002)

http://heasarc.gsfc.nasa.gov/docs/heasarc/fits/checksum.html

Chandra conventions

Main page of links:

http://cxc.harvard.edu/cda/doc.html

ASC FITS File Designers' Guide, ASC-FITS-2.1.0:

http://cxc.harvard.edu/contrib/arots/fits/ascfits.ps

Table of keyword values / column names:

http://cxc.harvard.edu/contrib/arots/fits/content.txt

Who cares? It's only a thesis.

—Fiona A. Harrison (1993-12-17)