

BLAST: A BALLOON-BORNE, LARGE-APERTURE, SUBMILLIMETRE
TELESCOPE

by

Donald Victor Wiebe

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Physics
University of Toronto

Copyright © 2008 by Donald Victor Wiebe

Abstract

BLAST: A Balloon-borne, Large-Aperture, Submillimetre Telescope

Donald Victor Wiebe

Doctor of Philosophy

Graduate Department of Physics

University of Toronto

2008

BLAST is a balloon-borne large-aperture, submillimetre telescope, which makes large area (1–200 square degree) surveys of Galactic and extragalactic targets. Since BLAST observes in the stratosphere, it is able to make broad-band observations between 200 μm and 550 μm which are difficult or impossible to perform from the ground. BLAST has been designed to probe star formation both in the local Galaxy and in the high redshift ($z = 1\text{--}4$) universe.

Because BLAST is flown on an unmanned stratospheric balloon platform, it has been designed to be able to operate autonomously, without needing operator intervention to perform its scientific goals. This thesis includes an overview of the design of the BLAST platform, with emphasis on the command and control systems used to operate the telescope.

BLAST has been flown on two long-duration balloon flights. The first of these, from Esrange, Sweden in June of 2005, acquired ~ 70 hours of primarily Galactic data. During the second flight, from Willy Field, Antarctica in December of 2006, BLAST acquired ~ 225 hours of both Galactic and extragalactic data. Operational performance of the platform during these two flights is reviewed, with the goal of providing insight on how future flights can be improved.

Reduction of the data acquired by these large-format bolometer arrays is a challenging procedure, and techniques developed for BLAST data reduction are reviewed. The

ultimate goal of this reduction is the generation of high quality astronomical maps which can be used for subsequent portions of data analysis.

This thesis treats, in detail, the iterative, maximum likelihood map maker developed for BLAST. Results of simulations performed on the map maker to characterise its ability to reconstruct astronomical signals are presented. Finally, astronomical maps produced by this map maker using real data acquired by BLAST are presented, with a discussion on non-physical map pathologies resulting from the data reduction pipeline and map making procedures.

Πάνυ καλῶς οἱ γνησίως φιλοσοφήσαντες, ὦ Σύρε, δοκοῦσί μοι κεχωρικέναι τὸ θεωρητικὸν τῶς φιλοσοφίας ἀπὸ τοῦ πρακτικοῦ. καὶ γὰρ εἰ συμβέβηκε καὶ τῷ πρακτικῷ πρότερον αὐτοῦ τούτου θεωρητικῷ τυγχάνειν, οὐδὲν ἦττον ἂν τις εὔροι μεγάλην οὖσαν ἐν αὐτοῖς διαφορὰν, οὐ μόνον διὰ τὸ τῶν μὲν ἠθικῶν ἀρετῶν ἐνίας ὑπάρξει δύνασθαι πολλοῖς καὶ χωρὶς μαθήσεως, τῆς δὲ τῶν ὅλων θεωρίας ἀδύνατον εἶναι τυχεῖν ἄνευ διδασκαλίας,
—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Α':α' [1]

Procœmium

*(In lieu of acknowledgements, I offer the following sentences;
I hope that those mentioned herein will take this to heart.)*

BLAST was not a simple experiment to design nor implement, and all who had an hand in its success should feel rightly proud of its accomplishments.

First I need to mention Matthew Truch, Marie Rex, and Gaelen Marsden, my compatriots in the “fantastic four,” without whom BLAST would not have been what it became. I hope BLAST’s success will continue to reflect well upon them, as it should. Not only have they shown themselves to be researchers of the first rate, but they have also created an environment within the BLAST collaboration within which it has been a delight to grow and develop as a scientist myself. I can only hope that I have managed to repay, at least in part, the support they have given to me. That I am the last of us to sluff this naissance, should not be overlooked as a mark to their quality. I could not have done this with out you guys, nor, in retrospect, would I have wanted to. You are always welcome in my house. I pray we will have the chance for many reunions.

Others have had no less great an impact on my development. First among these is my thesis supervisor, Barth Netterfield. Barth, you are one of the most analytically astute persons I know. I first met Barth when he lectured a course during my first year of graduate studies. I soon found him to be among the most approachable and supportive people of my early graduate career. Having worked with him for the past few years, I have had the opportunity to learn from his pragmatism and keen intuition. That I have come as far as I have should be a testament to his ability. Barth, thank you for you guidance and support. May peace be with you.

Mark Devlin has had an incalculable contribution to the success of BLAST. I hope that some of his savvy has rubbed off on me. Edward Chapin is no less responsible for what I have become than any of those mentioned before. I have spent almost as much time in the field with him as any other, and more time sharing a room with him than with any other perhaps Matt. Similarly, Chris Semisch I have, perhaps, spent more time sharing living quarters with than anyone else on BLAST. The rest of you who joined me during the BLAST field campaigns I remember, too.

Enzo Pascale has been more of a mentor than a fellow graduate student to me. His contributions to BLAST and to my education cannot be measured. I also need to mention Carrie MacTavish and Perseus Watt, who suffered with me through the beer factory—by far the coldest part of my BLAST experience.

This thesis wouldn’t have been possible without the support of my family, especially my parents for their love and guidance, Anna and Beth, who followed me out east, and

who have been a source of most needed companionship, and, finally, my cousin Cynthia, who showed me when I first arrived how to live as a Torontonion.

This document was written in \TeX and \LaTeX using VIM—Vi IMproved¹ and typeset using teTeX ² and additional \LaTeX packages from CTAN—the Comprehensive \TeX Archive Network.³ DVI output from teTeX was further processed by dvips ⁴ and GNU GhostScript⁵ to produce PostScript and PDF versions of this document.

Photographs in this document were prepared in GIMP, the GNU Image Manipulation Program,⁶ and then converted to encapsulated PostScript by ImageMagick.⁷ Most plots were made using kst ⁸; maps produced by the map makers were prepared for publication using Fv ⁹; circuit schematics were produced using gEDA , the GPL'd Electronic Design Automation tools¹⁰; other schematics were produced in Dia ¹¹ or Inkscape .¹² Document layout was performed in Perl .¹³ The complete build process was managed by GNU Make.¹⁴

Mylar is a registered trademark of E. I. du Pont de Nemours and Company.¹⁵ Linux is a registered trademark of Linus Torvalds. Altera and Nios are registered trademarks of Altera Corporation.¹⁶ SHARC is a registered trademark of Analog Devices, Incorporated.¹⁷

Corrections are gratefully accepted by the author at the electronic mail address below. When sending corrections, please indicate the edition you wish to ammend. (This is the third edition.) Please also indicate the particular layout for the given edition. (This is the duplex layout.) Electronic copies of this document may be obtained free of charge over the World Wide Web by visiting <http://ketiltrout.net/thesis/>.

D. V. Wiebe
Toronto, Ontario
25 June, 2008
dvw@ketiltrout.net

¹<http://www.vim.org/>

²<http://www.tug.org/tetex/>

³<http://www.ctan.org/>

⁴<http://www.radical-eye.com/dvips.html>

⁵<http://www.gnu.org/software/ghostscript/>

⁶<http://www.gimp.org/>

⁷<http://www.imagemagick.org/>

⁸<http://kst.kde.org/>

⁹<http://heasarc.nasa.gov/fvtools/fv/>

¹⁰<http://www.geda.seul.org/>

¹¹<http://live.gnome.org/Dia>

¹²<http://www.inkscape.org/>

¹³<http://www.perl.org/>

¹⁴<http://www.gnu.org/software/make/>

¹⁵1007 Market Street, Wilmington, Delaware, USA, 19898

¹⁶101 Innovation Drive, San Jose, California, USA, 95134

¹⁷1 Technology Way, Norwood, Massachusetts, USA, 02062

Short Contents

Abstract	iii
Prooemium	vii
Short Contents	ix
Contents	xi
List of Tables	xvii
List of Figures	xxi
List of Acronyms	xxiii
1 Introduction	1
2 The BLAST Instrument	13
3 Flight Performance	99
4 Data Reduction	129
5 Map Making	151
6 BLAST Maps	179
7 Summary	219
A Watchdog Circuit Schematics	221
B Glossary	227
C Excerpts from Ptolemy's <i>Almagest</i>	235
Bibliography	239

Contents

Abstract	iii
Prooemium	vii
Short Contents	ix
Contents	xi
List of Tables	xvii
List of Figures	xxi
List of Acronyms	xxiii
1 Introduction	1
1.1 Observing in the Submillimetre	1
1.1.1 <i>K</i> Correction	2
1.1.2 Previous Observations	5
1.1.3 Science Goals	10
1.2 Thesis Outline	11
1.3 My Thesis Work	11
2 The BLAST Instrument	13
2.1 Overview	14
2.2 Structural Elements	17
2.2.1 The Inner Frame	17
2.2.2 The Outer Frame	18
2.2.3 The Sun Shields	18
2.3 The Master Control Program	19
2.3.1 Architecture	20
2.3.2 The Main Thread	20
2.3.2.1 Initialisation	21
2.3.2.2 BLASTbus Reception	22
2.3.2.3 Frame Control	22
2.3.2.4 Control Update	23
2.3.2.5 The Command Scheduler	23

2.3.3	The In-flight Pointing Solution	26
2.3.3.1	Pointing Solution Calculation	27
2.3.3.2	Pointing Solution Summation	28
2.3.3.3	Gyro Offset Calculation	28
2.3.4	On-board Data Storage	29
2.4	The Flight Computers	32
2.4.1	Operating System	32
2.4.2	The Pressure Vessel	33
2.4.3	The Watchdog Circuit	35
2.5	Receiver	36
2.5.1	Detectors and Feeds	37
2.5.1.1	The Bolometer Device	37
2.5.2	Read-Out	39
2.6	The Cryostat	40
2.7	The Data Acquisition System	43
2.8	DSP Firmware	45
2.8.1	Lowpass Filter	46
2.8.2	The Phase-locked Loop and Digital Lock-in	48
2.9	The BLASTbus	48
2.9.1	BLASTbus Architecture	49
2.9.1.1	Frames	50
2.9.2	The BLASTbus Controller	51
2.9.2.1	BBC Architecture	52
2.9.2.2	BBC Firmware	52
2.10	Optics	54
2.10.1	The Cassegrain Telescope	54
2.10.2	Thermal Defocusing of the Telescope	56
2.10.3	The Active Focus Mechanism	60
2.10.3.1	Instrumentation	60
2.10.3.2	Operation	62
2.10.4	The Re-imaging Optics	63
2.10.5	Filters and Bandpasses	64
2.11	Pointing Sensors	65
2.11.1	The Star Cameras	66
2.11.1.1	Star Camera Operation	66
2.11.1.2	Communication	67
2.11.1.3	The Trigger	69
2.11.1.4	The Star Camera Cycle	69
2.11.2	The Gyroscopes	71
2.11.2.1	Gyroscope Operation	72
2.11.2.2	Thermal Drift	73
2.11.2.3	Readout	73
2.11.3	Coarse Sensors	74
2.11.3.1	The Differential GPS	74
2.11.3.2	The Sun Sensor	75

2.11.3.3	The Magnetometer	76
2.11.3.4	The Elevation Encoder	77
2.11.3.5	The Elevation Inclinometer	77
2.12	The Attitude Control System	77
2.12.1	Motion Control	78
2.12.1.1	Pointing Modes	78
2.12.1.2	Elevation Control	80
2.12.1.3	Azimuth Control	80
2.12.1.4	Roll Control	81
2.12.2	The Power System	82
2.12.2.1	Solar Arrays	83
2.12.2.2	Batteries	83
2.12.2.3	Power Distribution	84
2.12.3	Auxiliary Systems	85
2.12.3.1	The Balance System	85
2.12.3.2	The Inner Frame Lock	85
2.12.3.3	The Active Cooling System	87
2.13	Telemetry	87
2.13.1	Line-of-Sight Data Transmission	88
2.13.1.1	Biphase Reception	90
2.13.1.2	Interloquendi	91
2.13.1.3	Defile	92
2.13.2	The High-Rate Downlink	92
2.13.2.1	TDRSS Compression	93
2.13.2.2	High Rate Reception	94
2.13.3	The Low-Rate Downlink	94
2.13.4	The Command Uplink	95
2.13.4.1	The Science Stack	95
2.13.4.2	Software Commands	96
2.13.4.3	Groundstation Commanding Software	96
2.13.5	Flight Monitoring Software	97
3	Flight Performance	99
3.1	Scan Strategies	99
3.1.1	Scan Limitations	99
3.1.2	Scanning Modes	101
3.1.3	Cross-linking	103
3.2	Flight Operations	104
3.2.1	Line-of-Sight Diagnostics	104
3.2.2	Calibrators	105
3.3	<i>BLAST05</i>	106
3.3.1	Ascent	106
3.3.2	The Cassegrain Telescope	107
3.3.3	Balloon Environment	110
3.3.4	The Pointing System	112

3.3.4.1	Elevation Control Loss	113
3.4	<i>BLAST06</i>	115
3.4.1	Pre-flight Focus Tests	115
3.4.2	In-flight Focus Tests	116
3.4.3	Balloon Environment	116
3.4.4	The Inner Frame Lock	118
3.4.5	The Pointing System	119
3.4.5.1	The Pointing Failure of T+225 ^h	120
3.4.6	Telemetry	126
4	Data Reduction	129
4.1	Initial Data Preparation	131
4.1.1	Merging and Verification	131
4.1.2	Initial Inspection	133
4.2	Detector Cleaning	133
4.2.1	The Detector Transfer Function	134
4.2.2	Deconvolution and Gap Filling	134
4.3	Post-Flight Pointing Solution	137
4.3.1	Star Camera Analysis	138
4.3.2	Pointing Solution Calculation	138
4.4	Detector Response Profiling	139
4.5	Map Making	141
4.6	Noise Power Spectra	141
4.6.1	Total Noise Estimation	142
4.6.2	Correlated Noise Estimation	143
4.7	Flatfielding and Absolute Calibration	143
4.7.1	Flatfielding	144
4.7.2	Absolute Calibration	145
4.8	Detector Alignment	147
4.9	Pointing Offsets	148
5	Map Making	151
5.1	The Map-Making Equation	151
5.1.1	Naïve Map Making	153
5.1.2	Maximum Likelihood Map Making	153
5.2	Iterative Map Making	154
5.3	The BLAST Map Makers	156
5.3.1	The Pointing Solution	156
5.3.2	Calibration	157
5.3.3	Optibin	158
5.3.4	SANEPIC	158
5.4	Almagest	159
5.4.1	The Algorithm	160
5.4.2	Map Initialisation	160
5.4.3	Data Partitioning	161

5.4.4	Data Pre-Filtering	162
5.4.5	Linear Predictive Gap Filling	163
5.4.6	The Weight Spectrum	165
5.4.7	Correlated Noise	166
5.4.8	Performance	168
5.5	Almagest Simulations	169
5.5.1	Input Data	169
5.5.2	Results	170
5.5.3	Implications	176
6	BLAST Maps	179
6.1	<i>BLAST05</i>	179
6.1.1	Cygnus X	179
6.1.2	Vulpecula	182
6.1.3	Aquila	186
6.1.4	Sagitta	186
6.1.5	CRL 2688	189
6.1.6	Arp 220	189
6.1.7	IC 5146	192
6.1.8	K3-50	192
6.1.9	NGC 4565	196
6.1.10	Cassiopeia A	196
6.2	<i>BLAST06</i>	198
6.2.1	CDFS Deep	198
6.2.2	CDFS Wide	198
6.2.3	Vela	204
6.2.4	Puppis A	207
6.2.5	η Carinae	208
6.2.6	VY Canis Majoris	208
6.2.7	IRAS 08470–4243	212
6.2.8	Nearby Resolved Galaxies	212
6.2.8.1	NGC 1097	212
6.2.8.2	NGC 1291	212
6.2.8.3	NGC 1365	215
6.2.8.4	NGC 1512	215
6.2.8.5	NGC 1566	215
6.2.8.6	NGC 1808	215
7	Summary	219
A	Watchdog Circuit Schematics	221
B	Glossary	227
C	Excerpts from Ptolemy’s <i>Almagest</i>	235

List of Tables

1.1	Submillimetric and FIR instruments	10
2.1	Pointing sensor characteristics	26
2.2	MCC processor speed	32
2.3	Physical properties of M1 and M2	59
2.4	Physical properties of M3, M4, and M5	64
2.5	Star camera specifications	66
2.6	BLAST Pointing Modes	79
2.7	BLAST Power Budget	82
3.1	Pointing solution recovery after gyroscope switch	121
4.1	Absolute flux calibration	145
4.2	Bolometer offset transformation parameters	149
4.3	Star camera—boresight rotation	150
6.1	Observations during <i>BLAST05</i>	181
6.2	Observations during <i>BLAST06</i>	201

List of Figures

1.1	SED of IRAS 22134+5834	3
1.2	SED of BLAST J194136+232325	4
1.3	Negative K-correction	6
1.4	Atmospheric transmission from Mauna Kea	7
1.5	Atmospheric transmission at aircraft and balloon altitudes	8
2.1	The BLAST platform	15
2.2	The warm electronics	16
2.3	The Master Control Program	21
2.4	Gyro offsets calculated in flight	30
2.5	In-flight and post-flight gyro offsets	31
2.6	The pressure vessel	34
2.7	The watchdog circuit	35
2.8	Bolometer and array	38
2.9	Bolometer read-out electronics	40
2.10	The cryostat	41
2.11	ADC card	44
2.12	Digital filter response	47
2.13	The BLASTbus word	50
2.14	The BLASTbus controller card	53
2.15	Optical layout	55
2.16	The two Cassegrain telescopes	56
2.17	Alignment tolerances for optical elements	57
2.18	Estimate of thermal defocusing for <i>BLAST05</i>	58
2.19	Active focusing mechanism	61
2.20	LVDT noise performance	62
2.21	The optics box	64
2.22	Bandpasses	65
2.23	The star cameras	67
2.24	Fibre-optic gyroscope	72
2.25	The sun sensors	75
2.26	BLAST telemetry	89
2.27	The decom card	90
3.1	LDB flight paths	100
3.2	Scan patterns	102

3.3	Cross-linking	103
3.4	A pass over Saturn	108
3.5	<i>BLAST05</i> beam maps	108
3.6	The <i>BLAST05</i> launch event	109
3.7	Altitude variations during <i>BLAST05</i>	110
3.8	In-flight temperature profiles for <i>BLAST05</i>	111
3.9	Altitude variations during <i>BLAST06</i>	117
3.10	In-flight temperature profiles for <i>BLAST06</i>	118
3.11	Pointing solution recovery after gyroscope switch	122
3.12	Catastrophic single bit flip	123
3.13	Loss of azimuthal control after elevation gyro error	125
3.14	An excerpt of the TDRSS high-rate telemetry	127
4.1	The data reduction pipeline	130
4.2	Count of discrepancies between MCC0 and MCC1	132
4.3	Measured detector transfer function	135
4.4	Bolometer cleaning for <i>BLAST05</i>	136
4.5	Power spectrum of the deconvolved data	137
4.6	Detector responsivity correction	140
4.7	Noise decomposition	144
4.8	Flatfielding from <i>BLAST05</i>	146
4.9	Measured bolometer offsets	148
4.10	Field offsets for <i>BLAST05</i>	150
5.1	Typical almagest weight spectrum	167
5.2	Signal-only simulation maps	172
5.3	Noise-only simulation maps	173
5.4	Signal-only radial power spectra	174
5.5	Noise-only radial power spectra	174
5.6	Radial map maker transfer function	175
5.7	Two-dimensional transfer function	177
6.1	Coverage map for <i>BLAST05</i>	180
6.2	Galactic plane survey in Cygnus X at 250 μm	183
6.3	Galactic plane survey in Cygnus X at 350 μm and 500 μm	184
6.4	Difference map for Cygnus X	185
6.5	Galactic plane survey in Vulpecula	187
6.6	Galactic plane survey in Aquila	188
6.7	Galactic plane survey in Sagitta	190
6.8	Proto-planetary nebula CRL 2688	191
6.9	Ultra-luminous infrared galaxy Arp 220	193
6.10	Low-mass star forming region near IC 5146	194
6.11	Compact HII region K3-50	195
6.12	Edge-on spiral galaxy NGC 4565	197
6.13	Supernova remnant Cassiopeia A	199

6.14	Coverage map for <i>BLAST06</i>	200
6.15	Extragalactic BLAST field CDFS Deep	202
6.16	Extragalactic BLAST field CDFS Wide	203
6.17	Galactic plane survey in Vela at 250 μm	205
6.18	Galactic plane survey in Vela at 350 μm and 500 μm	206
6.19	Vela Deep embedded in Vela Wide	207
6.20	Puppis A supernova remnant	209
6.21	η Carinae Nebula	210
6.22	VY Canis Majoris	211
6.23	IRAS 08470–4243	213
6.24	NGC 1097 and NGC 1291	214
6.25	NGC 1365 and NGC 1512	216
6.26	NGC 1566 and NGC 1808	217
A.1	Watchdog Schematic Part 1 of 6	221
A.2	Watchdog Schematic Part 2 of 6	222
A.3	Watchdog Schematic Part 3 of 6	222
A.4	Watchdog Schematic Part 4 of 6	223
A.5	Watchdog Schematic Part 5 of 6	224
A.6	Watchdog Schematic Part 6 of 6	225

List of Acronyms

<i>Notation</i>	<i>Description</i>	<i>Page Ref.</i>
0CTE	Zero Co-efficient of Thermal Expansion	54, 56
AC	Alternating Current	39, 42, 47
ACS	Attitude Control System	12, 22, 43, 49, 51, 52, 67, 69, 70, 73, 77, 78, 80, 82–84, 86, 87, 110, 117, 120, 124, 221, 227
ADC	Analogue-to-Digital Converter. <i>See also</i> $\Sigma\Delta$.	23, 42, 43, 45–53, 73, 77, 78, 84, 110, 132, 227
AFM	Active Focus Mechanism	12, 60, 61, 86
ASCII	American Standard Code for Information Interchange	60
AzTEC	Astronomical Thermal Emission Camera	9
BBC	BLASTbus Controller	21–23, 31, 32, 36, 49, 51–53, 88, 90, 228
BDA	Bolometer Detector Array	37, 42
BLAST	Balloon-borne Large Aperture Submillimetre Telescope	<i>passim</i>
<i>BLAST03</i>	BLAST test flight, New Mexico, USA, August 2003	11, 12, 16, 17, 32, 51, 54, 56, 65, 71, 75, 86, 228
<i>BLAST05</i>	BLAST LDB flight, Sweden—Canada, June 2005	1, 2, 12, 16–18, 24, 25, 32, 54, 56, 58, 62, 65, 66, 68, 71, 75, 76, 82, 86, 87, 99, 104, 107, 109, 110, 112, 113, 117, 119, 129, 131–137, 139–145, 147–149, 152, 168, 169, 179, 186, 189, 198, 219, 229

continued...

<i>Notation</i>	<i>Description</i>	<i>Page Ref.</i>
<i>BLAST06</i>	BLAST LDB flight, Antarctica, December 2006	12, 15–19, 24–26, 29, 32, 49, 53, 54, 58–61, 63, 65, 66, 68, 73, 75, 76, 82, 84, 86, 87, 99, 103–105, 107, 112, 114, 116, 117, 119, 120, 126, 127, 129, 131, 133, 136, 137, 139, 145, 147–149, 152, 179, 198, 208, 212, 219, 220, 228
BOOMERanG	Balloon Observations of Millimetric Extragalactic Radiation and Geophysics	83
CCD	Charge-Coupled Device	65, 66, 75, 82, 87
CDFS	Chandra Deep Field South	198, 204
CITA	Canadian Institute for Theoretical Astrophysics	159, 168
CMB	Cosmic Microwave Background	152, 161
COBE	Cosmic Background Explorer	9
COI	Composite Optics Incorporated	54, 59, 63, 228, 229
CPU	Central Processing Unit	32, 33
CRL	Cambridge Research Labs	144, 145, 179, 189
CSBF	Columbia Scientific Balloon Facility	15, 16, 18, 88, 94–96
CSO	Caltech Submillimeter Observatory	9
CSS	CCD Sun Sensor	75, 76, 87
DAC	Digital-to-Analogue Converter	45
DAS	Data Acquisition System	11, 17, 23, 39, 42, 43, 47, 49, 51, 52, 77, 82–84, 86, 106, 110, 117, 132
DC	Direct current	221
DC–DC	DC to DC power converter	33, 35, 45, 84, 221
dGPS	Differential Global Positioning System	15, 74, 112, 119
DIRBE	Diffuse Infrared Background Experiment	9
DM	Data Memory, one bank of the DSP’s internal memory	47
DSP	Digital Signal Processor	43, 45–47, 49, 51, 78, 80, 81, 119, 120, 122, 227
ECC	Error-Correction Circuitry	122
ELAIS	European Large Area ISO Survey	142, 179
EMF	Electromotive Force	81

continued...

<i>Notation</i>	<i>Description</i>	<i>Page Ref.</i>
EPROM	Erasable, Programmable, Read-Only Memory	45
ESO	European Southern Observatory	212, 215
ext3	Third Extended Filesystem	31
FFT	Fast Fourier Transform	162, 168
FFTW	Fastest Fourier Transform in the West, an FFT library	162
FIB	Far Infrared Background	9
FIC	Field Interface Computer	88, 90, 91
FIFO	First In, First Out	53
FIR	Finite Input Response	9, 28
FOG	Fibre-Optic Gyroscope	71, 72
FPGA	Flux-Programmable Gate Array	43, 52
FTP	File Transfer Protocol	91
FTS	Fourier Transform Spectrometer	63, 65, 230
FWHM	Full Width at Half Maximum	100
GOODS	Great Observatories Origins Deep Survey	198
GPS	Global Positioning System. <i>See also</i> dGPS	15, 22, 26, 74, 82, 112, 119
GRSMC	Galactic Ring Survey molecular cloud	186
GSE	Groundstation Support Equipment	16, 88, 94, 96
GUI	Graphical User Interface	96
HF	High-frequency	16, 87, 96
IC	Index Catalogue [2, 3]	179, 192
IIR	Infinite Input Response	80
IR	Infrared	40
IRAM	Institut de radioastronomie millimétrique	9
IRAS	Infrared Astronomical Satellite	1, 2, 5, 9, 179, 196, 198, 204, 207, 208, 212
ISA	Industry Standard Architecture	51, 52
ISC	Integrating Star Camera	65, 66, 120, 139, 148, 149, 156
IVC	Intermediate Velocity Cloud	169–171, 176, 179
JCMT	James Clerk Maxwell Telescope	5, 9
JFET	Junction Field Effect Transistor	39, 40, 43
LDB	Long Duration Balloon	11, 24, 99, 106, 114
LEDA	Lyon-Meudon Extragalactic Database [4]	215
LHe	Liquid Helium	14, 39, 40, 42
LLNL	Lawrence Livermore National Laboratory	56
LMT	Large Millimeter Telescope	9
LN	Liquid Nitrogen	14, 39, 40
LOS	Line of Sight	16, 49, 52, 82, 87, 88, 95, 104, 116

continued...

<i>Notation</i>	<i>Description</i>	<i>Page Ref.</i>
LP	Linear Prediction	135, 136, 163–165, 168
LSD	Local Sidereal Day	24
LST	Local Sidereal Time	23, 24, 116, 232
LVDT	Linear Variable Differential Transformer	60, 61, 230
M1	Mirror #1 (Primary)	54, 56, 58–60, 63
M2	Mirror #2 (Secondary)	54, 56, 58–60, 62
M3	Mirror #3	54, 56, 63
M4	Mirror #4 (Lyot-stop)	36, 54, 56, 63, 228
M5	Mirror #5	54, 56, 63
MAMBO	Max-Planck Millimetre Bolometer Array	9
MCC	Master Control Computer	14, 31, 36, 66, 67, 71, 74, 77, 131, 221
MCC0	Master Control Computer #0, commonly called frodo	32, 33, 126, 127, 131, 132
MCC1	Master Control Computer #1, commonly called sam	32, 33, 126, 127, 131, 132
mcp	Master Control Program	11, 12, 14, 19–22, 25, 31, 35, 42, 43, 49, 53, 68–71, 74, 76, 78, 80, 84–88, 91, 94–96, 102, 112–115, 118–120, 124, 126, 136–138, 231
MIPS	Multiband Imaging Photometer for Spitzer	2, 9
MSX	Midcourse Space Experiment	2, 9
NASA	National Aeronautic and Space Administration (United States)	15, 54, 92, 232
NEP	Noise Equivalent Power	37
NFS	Network Filesystem	91
NGC	New General Catalogue [5]	179, 182, 196, 198, 208, 212, 215
NiMH	Nickel-Metal Hydride	16, 83
NOAA	National Oceanic and Atmospheric Administration (United States)	76
NSBF	National Scientific Balloon Facility (United States; <i>defunct</i> , now CSBF)	15
NTD	Neutron Transmutation Doped (a type of bolometer)	37, 39
NTP	Network Time Protocol	74
NTSC	National Television Standards Committee	88, 231
OCC	Operations Control Center	88, 94, 95

continued...

<i>Notation</i>	<i>Description</i>	<i>Page Ref.</i>
OSC	Other Star Camera (in contrast to the ISC, which was built first)	65, 66, 120
PCB	Printed Circuit Board	107
PCI	Peripheral Component Interconnect	51–53, 90, 228
PCU	Power Conditioning Unit (also known as a Charge Controller)	82
PET	Polyethylene Terephthalate, a polyester	227
PI	Proportional-Integral Control Algorithm	78, 80
PID	Proportional-Integral-Derivative Control Algorithm	73
PKS	Parkes Catalog [6]	198
PLL	Phase-locked Loop	46, 47
PM	Program Memory, one bank of the DSP’s internal memory	47
PSF	Point-spread function	107, 152, 182, 186
$\psi\mu\text{m}$	“pseudomicron” = 1.0583 μm , used by the AFM	60–62, 116
PSS	Photodiode Sun Sensor	75, 76, 87
PV	Pressure Vessel	33, 110, 221
PVBOB	Pressure Vessel Break-Out Box	33
PWM	Pulse Width Modulated	45, 78, 80, 81
RAM	Random Access Memory	20, 31, 43, 52, 122, 168
RMS	Root Mean-Square	54, 56, 61, 140
ROx	Ruthenium Oxide Thick Film Resistor	42
RS-485	Recommended Standard #485	60, 232
SANEPIC	Signal and Noise Estimation Procedure Including Correlations, one of the BLAST map makers	140–143, 153, 156, 158, 166, 167, 169–171, 176, 178, 182
SCC	Star Camera Computer	65, 66, 68, 70, 71, 74, 137
SCUBA	Submillimetre Common-User Bolometer Array	1, 5, 9
$\Sigma\Delta$	Sigma-Delta, a type of ADC	43
SDL	Sidereal Date Line	24
SED	Spectral Energy Density	1, 2, 145, 179, 192
SEP	South Ecliptic Pole	198
SHA	Secure Hash Algorithm	132
SHARC	Submillimeter High Angular Resolution Camera	9
SHARC	Super Harvard Architecture Single-Chip Computer	43
SIC	Science Interface Computer	88, 94, 96

continued...

<i>Notation</i>	<i>Description</i>	<i>Page Ref.</i>
SIP	Support Instrumentation Package	16, 18, 88, 93–95
SOFIA	Stratospheric Observatory for Infrared Astronomy	5
SPIRE	Spectral and Photometric Imaging Receiver	9, 11, 36, 139
SRC	United Kingdom Science Research Council	212, 215
TCP/IP	Transport Control Protocol / Internet Protocol	88, 91
TDRSS	Tracking and Data Relay Satellite System	16, 29, 87, 88, 92–94, 96, 110, 113, 116, 120, 126, 127, 230, 232
TOD	Time-Ordered Data	92, 93, 100, 134, 140, 142, 151, 170, 182
TTL	Transistor–Transistor (<i>i.e.</i> 5 V) Logic	90
UKT14	United Kingdom Telescope—Cryostat #14	9
ULIRG	Ultra-luminous Infrared Galaxy	9, 145, 179, 189
USA	United States of America	5, 16, 92, 94, 95
UTC	Universal Time, Co-ordinated	24
UV	Ultraviolet	45
VCS	Vapour-Cooled Shield	40
VHDL	VHSIC Hardware Description Language	52
VMR	Vela Molecular Ridge	198
WMM	World Magnetic Model	76
WSGT	White Sands Ground Terminal	92, 94

Chapter 1

Introduction

1.1 Observing in the Submillimetre

The submillimetric regime ($\lambda = 50\text{--}1000 \mu\text{m}$) is ideal for measuring star formation. At these wavelengths we can probe both the earliest stages of star formation within our own galaxy, but also the earliest times of star formation in the early Universe ($z = 1\text{--}6$).

In our Galaxy, star formation is seen to occur in dusty nebulosities. The dust which surrounds these proto-stellar objects is opaque at optical wavelengths, shrouding them from optical detection. Radiation absorbed by the dust heats the surrounding medium to 10–40 K. Greybody emission from dust at these temperatures peaks in the submillimetre, where the surrounding medium is optically thin. Figure 1.1 shows the spectral energy density (SED) of a typical proto-stellar object, the medium to high mass protostar IRAS 22134+5834 [7]. This object is roughly 900 pc distant [8] and its SED, which peaks around 90 μm , is fit by a greybody spectrum:

$$L_\nu \propto \nu^\beta B_\nu(T) \tag{1.1}$$

where $B_\nu(T)$ is the Planck function for a temperature T , and β is the emissivity index of the greybody. The greybody parameters found from this fit are: $T = 21.9 \pm 5.6$ K,

and $\beta = 1.7 \pm 0.2$ [9].

The coldest source detected in the *BLAST05* Galactic plane survey in Vulpecula (Section 6.1.2) is shown in Figure 1.2. This source has a temperature of only 12.3 ± 1.6 K [11]. At this temperature, the greybody SED turns over in the BLAST bands, making it undetectable by previous far-infrared telescopes such as the Infrared Astronomical Satellite [7] (IRAS), the Multiband Imaging Photometer for Spitzer [12] (MIPS), or the Midcourse Space Experiment [13] (MSX).

Measurement of the star formation rate in the high-redshift Universe is also ideally suited to observations in the submillimetre. Greybody emission from dust-enshrouded starburst galaxies, which has a rest-frame peak around $100 \mu\text{m}$ (*cf.* Figure 1.1), will be shifted into the submillimetric range $200\text{--}600 \mu\text{m}$ at redshifts of $z = 1\text{--}5$.

1.1.1 *K* Correction

Observations of extragalactic sources in the BLAST bands, centred at $250 \mu\text{m}$, $350 \mu\text{m}$, and $500 \mu\text{m}$, are aided by a phenomenon often referred to as *negative K correction*. Briefly, *K* correction is the phenomenon by which the luminosity of a source observed, in a particular band, at various different redshifts will vary by an amount which is not a simple function of the source’s distance from the observer.

The term *K* correction was coined by Edwin Hubble, who first pointed out the effect in 1936 [14]. Hubble used the term to refer to the change in luminosity due to the fact that a different part of the source spectrum is being sampled. Since Hubble was working in optical magnitudes, the additive quantity *K*, also measured in magnitudes, is positive when the source is dimmer because of this effects, and negative when this correction results in a brightening of the source. The notion of *K* correction has since been expanded from Hubble’s work to account for arbitrarily different observed and emitted-frame band passes [15, 16].

In submillimetre astronomy, the *K* correction is rarely calculated explicitly; in cases




Figure 1.1: The spectral energy density of the medium to high mass proto-stellar object IRAS 22134+5834. BLAST points measured during *BLAST05* are diamonds, other points included in the fit are triangles, and are from SCUBA (450 μm), and IRAS (100 μm) [10, 7]. The black curve, with associated error-curves, is the best fit greybody SED fit to the indicated points. Asterisks denote IRAS observations (60 μm , 25 μm) which sample warmer components of the SED, and are not fit to. From [9].




Figure 1.2: The spectral energy density of the high-mass proto-stellar object BLAST J194136+232325. This is the coldest object ($T = 12.3$ K) detected by BLAST in the Galactic plane survey in Vulpecula during *BLAST05* (Section 6.1.2). BLAST points are diamonds. The solid black line is the best fit SED for a greybody with fixed emissivity index, $\beta = 1.5$. Since the SED turns over in the BLAST bands, there are no detections by IRAS (100 μm , 60 μm , 25 μm , 12 μm), MIPS (70 μm), or MSX (21.3 μm , 12.13 μm , 8.28 μm). Upper bounds shown are shown for these instruments. From [11].

when the redshift of a submillimetric source is known, the rest-frame luminosity of a source is typically calculated from multi-band data, where the actual values of the K corrections for each band have already been accounted for.

Nevertheless, submillimetre astronomers still mention negative K correction, simply due to the fact that, when observing cold sources on the Rayleigh-Jeans side of the blackbody curve, the effect that K correction measures allows observations at higher redshifts than would be possible at shorter wavelengths. The measured flux density of the source presented in Figure 1.1 is plotted in Figure 1.3 at various wavelengths and redshifts. As can be seen, the negative K correction effect, which occurs on the Rayleigh-Jeans side of the greybody curve, is very beneficial to the detection of high-redshift sources. Conversely, once one is observing on the Wiens side of the greybody curve (for example, at $100\ \mu\text{m}$ in this figure), the K correction is positive, and luminosity drops off very sharply with increasing redshift.

1.1.2 Previous Observations

Due to absorption of submillimetre light by water vapour in the atmosphere, submillimetre observations are difficult to do from the ground. Figure 1.4 shows the atmospheric transmission at submillimetric wavelengths for the James Clerk Maxwell Telescope (JCMT), where the Submillimetre Common-User Bolometer Array (SCUBA) [17] is installed. The BLAST experiment overcomes this by flying on a high-altitude balloon at ~ 38 km altitude, placing it above the bulk of the atmosphere. Figure 1.5 shows typical submillimetric atmospheric transmission at these altitudes. As a result, BLAST is in a unique position to make large submillimetric maps in a relatively short amount of time.

Due to atmospheric absorption, submillimetric observation is still a relatively young science. The first large-scale submillimetric survey was conducted using the Infrared Astronomical Satellite (IRAS) [7], launched in 1983. It observed 96% of the sky in four bands: $12\ \mu\text{m}$, $25\ \mu\text{m}$, $60\ \mu\text{m}$, and $100\ \mu\text{m}$. IRAS was phenomenally successful. It




Figure 1.3: The greybody of Figure 1.1 plotted at different redshifts. The left panel shows the greybody function (Equation 1.1) plotted as a function of wavelength for several different redshift values. The right panel shows the same function, but plotted as a function of redshift at several different wavelengths. The flux density scale is arbitrary. Notably, for observations around millimetre wavelengths, the observed flux of the source remains roughly constant as a function of redshift over the range $z = 2-8$.




Figure 1.4: Atmospheric transmission at submillimetric wavelengths calculated for the top of Mauna Kea, Hawai'i, USA (4.2 km altitude) where the JCMT is located. In descending order, the three curves show transmission profiles for precipitable water vapour of 0.5, 1, and 5 mm. At wavelengths shorter than 500 μm , atmospheric opacity is high. From [18].




Figure 1.5: Atmospheric transmission at submillimetric wavelengths calculated for aircraft altitudes used by the SOFIA instrument, as well as at 35 and 40 km, which bracket the altitude range of the BLAST telescope. At BLAST altitudes, the atmosphere is largely transparent. Plot generated with AT software by Erich Grossman. The BLAST bands are overlain in colour.

detected $>15\,000$ galaxies and hundreds of thousands of point sources, providing the definitive far-infrared source catalogue still used today.

IRAS was followed by the Diffuse Infrared Background Experiment (DIRBE), which flew on the Cosmic Background Explorer (COBE) satellite in 1989. DIRBE observed at $240\ \mu\text{m}$, $140\ \mu\text{m}$, $100\ \mu\text{m}$, and $60\ \mu\text{m}$, and detected the far-infrared background (FIB), a diffuse background due to the large number of unresolved ultra-luminous infrared galaxies (ULIRGs) in the early Universe. The power in the FIB is greater than the total optical power in the Universe, and is an order of magnitude larger than the integrated power obtained from IRAS point sources.

Follow-up experiments to both IRAS and DIRBE are attempting to constrain the composition of the FIB by directly observing high-redshift ($z = 1\text{--}6$) ULIRGs. This work was begun in 1992 using the UKT14 instrument [18], a single bolometer observing through atmospheric windows between $350\ \mu\text{m}$, and $2\ \text{mm}$, on the $15\ \text{m}$ James Clerk Maxwell Telescope (JCMT) on Mauna Kea, Hawai'i. In 1997, UKT14 was replaced by the Submillimeter Common User Bolometer Array, (SCUBA) [17], observing at $450\ \mu\text{m}$ and $850\ \mu\text{m}$, but containing 128 detectors split between the two bands, each an order of magnitude more sensitive than UKT14. A planned replacement for SCUBA, SCUBA-2, will continue this trend of larger numbers of increasingly more sensitive detectors at the JCMT. Similar observations have been carried out by the SHARC [19], Bolocam [20], and SHARC-II [22] instruments at the Caltech Submillimeter Observatory (CSO), also located on Mauna Kea, as well as the Max-Planck Millimetre Bolometer Array (MAMBO) [23] on the 30-metre Institut de radioastronomie millimétrique (IRAM) telescope in Spain.

BLAST's unprecedented mapping speed, and negligible atmospheric contamination, permit it to significantly exceed the performance of the current ground-based submillimetric instruments presented above. Furthermore, BLAST's bands ($200\ \mu\text{m}\text{--}550\ \mu\text{m}$) are well suited to augmentation by existing far-infrared (FIR) data from experiments such as IRAS, MIPS, and MSX. A survey of past, current, and future submillimetric and

<i>Instrument</i>	<i>Platform</i>	<i>Bands (μm)</i>	<i>Beam size ($''$)</i>	<i># of detectors</i>	<i>Years active</i>
IRAS	Satellite	12, 25, 60, 100	30–120	62	1983
DIRBE	<i>COBE</i>	1.25, 2.2, 3.5, 4.9, 12, 25, 60, 100, 140, 240	2500	10	1989–1990
UKT14	JCMT	350, 450, 750, 800, 850, 1100, 1300, 2000	19–27	1	1992–1997
SHARC	CSO	350, 450	9–12	20	1995–2003
MSX	Satellite	8.28, 12.13, 14.65, 21.3	21–23	3456	1996–1997
SCUBA	JCMT	450, 850	7.5–14	128	1997–2008
MAMBO-1	IRAM	1200	11	37	1998—
Bolocam	CSO	1100, 2100	30–60	115	2001—
MAMBO-2	IRAM	1200	11	117	2002—
MIPS	<i>Spitzer</i>	24, 70, 160	31–325	17 448	2003—
SHARC-II	CSO	350, 450	8–9	384	2003—
BLAST	Balloon	250, 350, 500	40–60	288	2003–2006
AzTEC	JCMT, LMT	1100, 2100	30–60	144	2005—
SCUBA-2	JCMT	450, 850	7.5–14	6400	2008—
SPIRE	<i>Herschel</i>	250, 350, 500	35–71	288	2008—

Table 1.1: A summary of several past, current, and future submillimetric and far-infrared instruments. BLAST is well suited to make use of data from experiments observing in a large number of different bands.

far-infrared instruments is presented as Table 1.1.

1.1.3 Science Goals

The scientific goals of BLAST are presented below:

- resolve and identify $\mathcal{O}(10^4)$ high-redshift submillimetre galaxies, roughly 100 times the number of known sources, and provide a statistically significant sample of such objects
- determine photometric redshifts of $\mathcal{O}(10^3)$ observed sources to determine their bolometric rest-frame far-infrared luminosities, and, hence, star formation rates
- measure the large-scale clustering of high-redshift galaxies
- identify pre-stellar cores in local star-forming regions in our own Galaxy

- search for high-mass proto-stellar objects near ultra-compact HII regions
- map Galactic cirrus, and other submillimetre bright Galactic structures

The SPIRE instrument on the Herschel satellite, scheduled to launch in late 2008, has similar scientific goals to BLAST, and will continue the work started by BLAST by mapping large portions of the submillimetric sky in the same bands.

1.2 Thesis Outline

This thesis is further organised as follows. Chapter 2 describes the BLAST instrument in detail, with particular focus on command and control systems. Chapter 3 touches on the in-flight performance of the instrument during the two long-duration balloon (LDB) flights. Chapter 4 outlines the data reduction pipeline implemented to convert the data obtained during the two LDB flights into science-quality maps. Chapter 5 considers map making from the BLAST data in detail. Finally, Chapter 6 presents several actual maps made from BLAST data acquired during the two flights.

1.3 My Thesis Work

For the construction, intergration, and operation of BLAST, I held primary responsibility for the flight software and flight electronics. This gave me the opportunity to interact with all systems on BLAST, to a greater or lesser extent. In addition to construction of the telescope, I have also worked on parts of the data reduction pipeline, most importantly by creating the maximum-likelihood map maker *almagest*.

In particular, before the *BLAST03* test flight, I built the Data Acquisition System (DAS, Section 2.7), and wrote the bulk of the flight software (*mcp*, Section 2.3), telemetry, ground station software (Section 2.13). At this time I also took care of integrating the coarse pointing sensors (Section 2.11.3) with the flight software, as well as developing the

star camera communications protocol (Section 2.11.1) with Ed Chapin, who wrote the star camera software.

Between *BLAST03* and *BLAST05*, I rebuilt the Attitude Control System (ACS, Section 2.12) to provide capabilities for extra systems (primarily the second star camera, and the solar array power system) added after *BLAST03*. I also worked to refine the star camera integration, and performed general testing on the pointing system.

After *BLAST05*, I started working on the data analysis pipeline. Having written the data archival code for `mcp`, my first task was to do the initial merge and verification of the data (Section 4.1). I also wrote the bolometer deconvolver (Section 4.2.2). My primary work during this time was the development of the almagest map maker (Section 5.4). In preparation for *BLAST06*, I also worked on flight code for the active focus mechanism (AFM, Section 2.10.3) and the new secondary lock (Section 2.12.3.2). Following *BLAST06*, data analysis continued, and my primary focus was completion and characterisation of the map maker.

Ἐξ ὧν διανοηθέντες, ὅτι τὰ μὲν ἄλλα δύο γένη τοῦ θεωρητικοῦ μᾶλλον ἢ τις εἰκασίαν ἢ κατάληψιν ἐπιστημονικὴν εἴποι, τὸ μὲν θεολογικὸν διὰ τὸ παντελῶς ἀφανὲς αὐτοῦ καὶ ἀνεπίληπτον, τὸ δὲ φυσικὸν διὰ τὸ τῆς ὕλης ἄστατον καὶ ἄδηλον, ὡς διὰ τοῦτο μηδέποτε ἂν ἐλπῖσαι περὶ αὐτῶν ὁμονοῆσαι τοὺς φιλοσοφοῦντας,
—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Α':α' [1]

Chapter 2

The BLAST Instrument

BLAST is an autonomous platform designed to observe the sky at submillimetric wavelengths. It makes simultaneous measurements in three broad bands: 215–300 μm (*the 250 μm band*), 300–400 μm (*the 350 μm band*), and 400–550 μm (*the 500 μm band*). See Figure 2.22. To mitigate absorption of light at these wavelengths by water vapour in the atmosphere, the BLAST platform is flown from a stratospheric balloon at a nominal altitude of ~ 38 km (“at float”), placing the telescope above 99.5% of the Earth’s atmosphere. While in flight, BLAST may be commanded from the ground by operators monitoring the telescope. However, due to the potential for loss of communications with the ground at this altitude, BLAST is designed to be able to operate without user intervention for the duration of the flight, following a pre-set observing schedule.

BLAST is a mapping experiment. The telescope and pointing systems are optimised for mapping areas of the sky from $\frac{1}{4}$ to hundreds of square degrees. BLAST’s primary science goal is imaging of submillimetre galaxies at high redshift, which requires long integrations in regions of low Galactic cirrus. However, the BLAST platform is flexible enough to permit pointed observations of any submillimetre-bright source, including local regions of star formation, as well as diffuse structures in the Galactic plane.

BLAST is designed to provide accurate pointing for the telescope to within half an

array size ($<3'$) in flight, and better than $\frac{1}{6}$ beam ($<5''$) in post-flight reconstruction. Primary pointing is provided by a pair of redundant daytime star cameras, which give absolute attitude measurements with an accuracy of $\sim 5''$ at a frequency of ~ 1 Hz, and two redundant sets of three orthogonal gyroscopes, which give relative attitude information between star camera solutions. Augmenting these is a suite of coarse absolute pointing sensors used primarily to bootstrap the pointing system when the star cameras have lost lock.

A full description of the instrument is provided in [24]. A summary, and various details are given below.

2.1 Overview

The BLAST platform can be decomposed into two functional components: the telescope, consisting of the optical elements and receiver; and the gondola, providing the pointing control and other support systems for the telescope. Physically, BLAST also consists of two pieces: an inner frame, housing the telescope and the primary pointing sensors: the star cameras and gyroscopes; and an outer frame, containing the pointing motors, secondary pointing sensors, and support electronics. The inner frame is mounted on the outer frame through an elevation bearing which permits free rotation of the inner frame in pitch relative to the outer frame. The fully assembled BLAST platform is presented in Figure 2.1

The BLAST telescope has three bolometric focal plane arrays, each tuned to a different colour and cooled to 300 mK by a LHe/LN cryostat. These are fed by a set of cold (2 K) re-imaging optics, and a ~ 2 m Cassegrain telescope. The bolometers are read out by a custom-built data acquisition system.

A redundant set of flight computers, known as the master control computers (MCCs), provides primary control of BLAST. The MCCs run a single, monolithic program, `mcp`




Figure 2.1: A simplified drawing of the BLAST platform with several major subsystems labelled. A metre-tall penguin is included for scale. From [24].

(the *Master Control Program*), which handles all necessary functions, including data acquisition and archival, motion control, telemetry, and housekeeping. Due to bandwidth restrictions in the telemetry downlink, all data acquired from the telescope are stored on board. Post-flight recovery of the data storage vessel is required for mission success. An overview of the major electronic subsystems for BLAST is presented in Figure 2.2.

Motion control for the gondola is provided by a elevation drive on the elevation axis, which allows rotation of the inner frame in pitch relative to the outer frame; a pivot motor and azimuthal reaction flywheel together providing yaw motion; and a roll damping flywheel, which minimises pendulation in roll. These motors are controlled based on input from the star cameras and gyroscopes, the primary pointing sensors, as well as the secondary pointing sensors: a differential GPS (dGPS), a magnetometer, a sun sensor, an elevation encoder, and various inclinometers.

The outer frame is hung on steel suspension cables from the pivot which, in turn, is attached to the flight train, consisting of a steel-cable ladder and parachute. The flight train is itself connected to a $\sim 1\,000\,000\text{ m}^3$ zero-pressure, helium-filled balloon which is provided by the Columbia Scientific Balloon Facility (CSBF), a NASA contractor, and

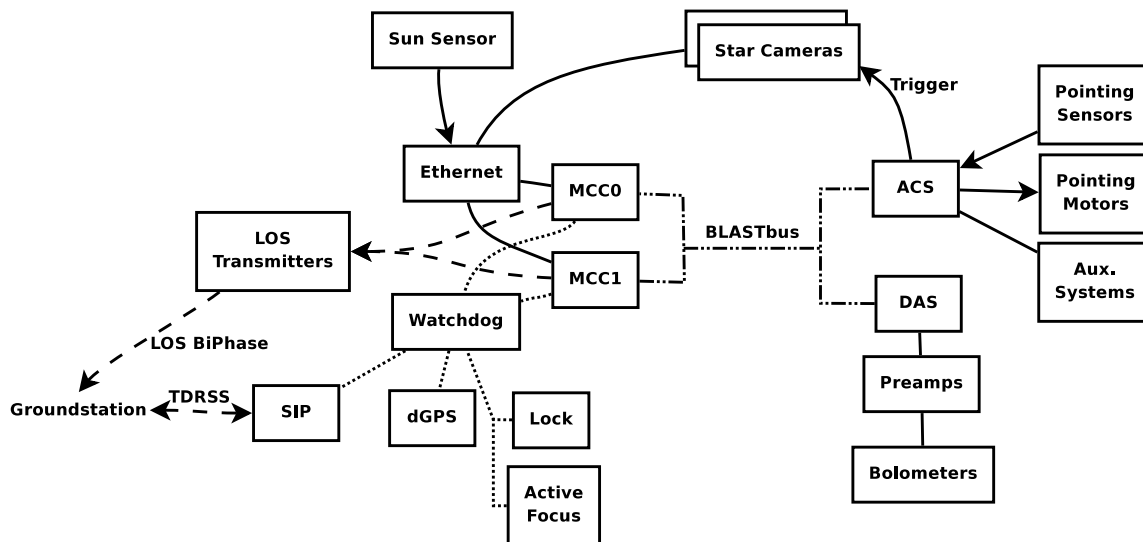


Figure 2.2: A schematic drawing of major components of the electronics for *BLAST06*. Dotted lines represent serial communication; dashed lines are uplink and downlink telemetry; other communication types are labelled.

provides the necessary lift and buoyancy for the stratospheric flight.

Communication with the platform in flight happens via the TDRSS network or, failing that, through the Iridium network. Interface with these satellite systems is provided by CSBF, via the Support Instrumentation Package (SIP) on the gondola, and groundstation equipment (GSE) in the field and at CSBF headquarters in Palestine, Texas, USA. Additionally, while the gondola is within line-of-sight (LOS) of the field camp, higher data rates can be obtained by using L- or S-band transmitters (the *LOS transmitters*) on the gondola, and high-frequency (HF) radio transmitters on the ground. Excluding the LOS transmitters, CSBF provides a route-independent interface to all telemetry channels.

The BLAST platform was designed to be flown in the polar regions, when the sun is always above the horizon. To permit this, the majority of the instrument is surrounded by an aluminised Mylar[®] and polystyrene sun shield with aluminum ribbing. Power for the platform is provided by solar arrays, routed through a pair of nickel/metal-hydride (NiMH) batteries. BLAST is designed to be at float for up to 250 hours, with flight time limited by cryogen boil-off.

Certain aspects of BLAST have been changed for each of its three flights. The rest

of this chapter will describe primarily the configuration of the BLAST platform during its final flight in December of 2006 (*BLAST06*). Important changes from *BLAST05* will be noted where appropriate. The incomplete platform flown in the *BLAST03* engineering test flight will only be mentioned insofar as it aids in describing the subsequent configurations.

2.2 Structural Elements

The BLAST frame provides the structure necessary to point the telescope. It consists of two primary structural components: an inner frame containing the telescope, star cameras, and gyroscopes; and an outer frame housing the remainder of the electronics, and the pointing motors. Attached to the outer frame are sun shields, which protect the telescope and the electronics from the elements. All frame members are composed primarily of aluminum. The dimensions of the frame are constrained by the launch vehicle and the buildings used to integrate BLAST.

2.2.1 The Inner Frame

The inner frame provides the structure to mount the telescope. It is roughly hemispherical in shape. It is attached to the outer frame at two points, defining the elevation axis around which it is free to rotate with respect to the outer frame.

The primary mirror is mounted on the flat front face of the inner frame. A triangular mounting bracket for the primary mirror allows it to be mounted at three points, minimizing deformation of the mirror's surface. The cryostat is mounted within the inner frame, directly behind the mirror. The detector readout electronics, the Data Acquisition System (DAS, Section 2.7) and preamp crates are mounted behind the cryostat. At the top of the inner frame are mounted the two bore-sight star cameras.

In order to balance the inner frame, its attachment points can be moved forwards and

backwards. For *BLAST03* and *BLAST06*, when a heavy aluminum primary was flown (Section 2.10.1), the attachment points were moved as far forward as possible compared to their position during the *BLAST05* flight.

For *BLAST05* a baffle protecting the Cassegrain telescope was flown, cantilevered off the front of the telescope (see Figure 2.1). This baffle was designed to provide a ground shield for the telescope. It was not flown in *BLAST06* due to the need to save weight.

2.2.2 The Outer Frame

The outer frame forms the base of the telescope. Rectangular in footprint, it is suspended from the pivot by four steel suspension cables, one at each corner. The azimuthal reaction wheel is mounted at the centre of the outer frame. Below it, a transverse keel provides strength and forms part of a cage for the CSBF's Support Instrumentation Package (SIP). For flight, a ballast hopper is attached to the bottom of the keel.

The outer frame has a set of legs which permit it to stand freely in the lab. These are flown as well, since they provide convenient attachment points for the CSBF solar arrays, and provide shock absorption on landing.

The top of the outer frame contains two pyramidal structures upon which the inner frame is mounted. On one pyramid is mounted the elevation motor, on the other, the free elevation bearing. This bearing has a hollow shaft, allowing power and signal cables to be passed through from the outer to inner frames with a minimum of torque applied to the inner frame.

2.2.3 The Sun Shields

A set of sun shields is used to protect the BLAST platform from the sun. These sun shields are composed of a tubular aluminum skeleton, and aluminised Mylar[®] covered polystyrene panels. The telescope is designed to observe primarily away from the sun, so the sun shields encompass the entirety of the back and sides of the BLAST platform.

Holes in the sun shield in the front provide apertures through which the electronics can radiate heat.

Additional sun shields provide a ground shielding forward of the outer frame (the *scoop*). The sun shields define the cavity through which the telescope observes the sky. At the top of the sun shields, a shelf provides a place to set the pivot down while it is not hanging. This is useful when transferring the gondola to and from the launch vehicle. This shelf also traps the pivot, preventing it from striking the telescope upon ‘chute shock after termination.

The sun shields determine the telescope’s limits of observation. The chin limits the lower elevation to $\sim 25^\circ$, while the upper elevation is limited to $\sim 55^\circ$ by the pivot shelf. Stipulating that sunlight may not reach the warm optics also limits the maximum azimuth to which the telescope can slew relative to the sun. For *BLAST06*, extensions were added to the top of sun shields to increase the telescope’s azimuthal range and prevent sunlight from directly illuminating the Cassegrain assembly.

2.3 The Master Control Program

Primary control of the gondola resides in the Master Control Program (*mcp*), which runs on a pair of redundant flight computers (Section 2.4). *mcp* is designed to run in the absence of operator intervention, and be robust enough to correct for, or else mitigate, common failure modes of the platform’s subsystems. The job of the Master Control Program is manifold:

- acquisition, collation, and storage of data acquired from the detectors, the pointing, and auxiliary sensors;
- calculation of the in-flight pointing solution;
- motion control and execution of the pre-scheduled flight-plan;

- communication with the ground, both data compression and downlink, and command uplink and processing;
- ^3He fridge cycling;
- focus correction;
- rebalancing of the inner frame;
- regulation of the cooling and heating systems;

2.3.1 Architecture

The Master Control Program is a multithreaded C program. A schematic layout of `mcp` is presented as Figure 2.3. The main thread controls communication with the BLASTbus (Section 2.9). The remaining threads asynchronously handle communication with systems not connected to the BLASTbus, via serial and Ethernet connections.

The main thread collects data from the BLASTbus and stores it locally in RAM for use as inputs by the control subsystems. Data not collected by the BLASTbus is written to the BLASTbus after acquisition. These “writebacks” collect all data on the BLASTbus, permitting the data stream written to disk to be composed solely of the received BLASTbus frames.

In addition to the local copy of the data, `mcp` also keeps a local state record in memory. The Master Control Program writes the state to disk whenever it updates it. The state contains sufficient information so that `mcp` can easily resume operation in cases where it is unexpectedly rebooted by the watchdog circuit (Section 2.4.3).

2.3.2 The Main Thread

The main thread in `mcp` initialises the system, and then handles transmission and reception of the BLASTbus. As a result, the main thread receives the 100 Hz frame sync from

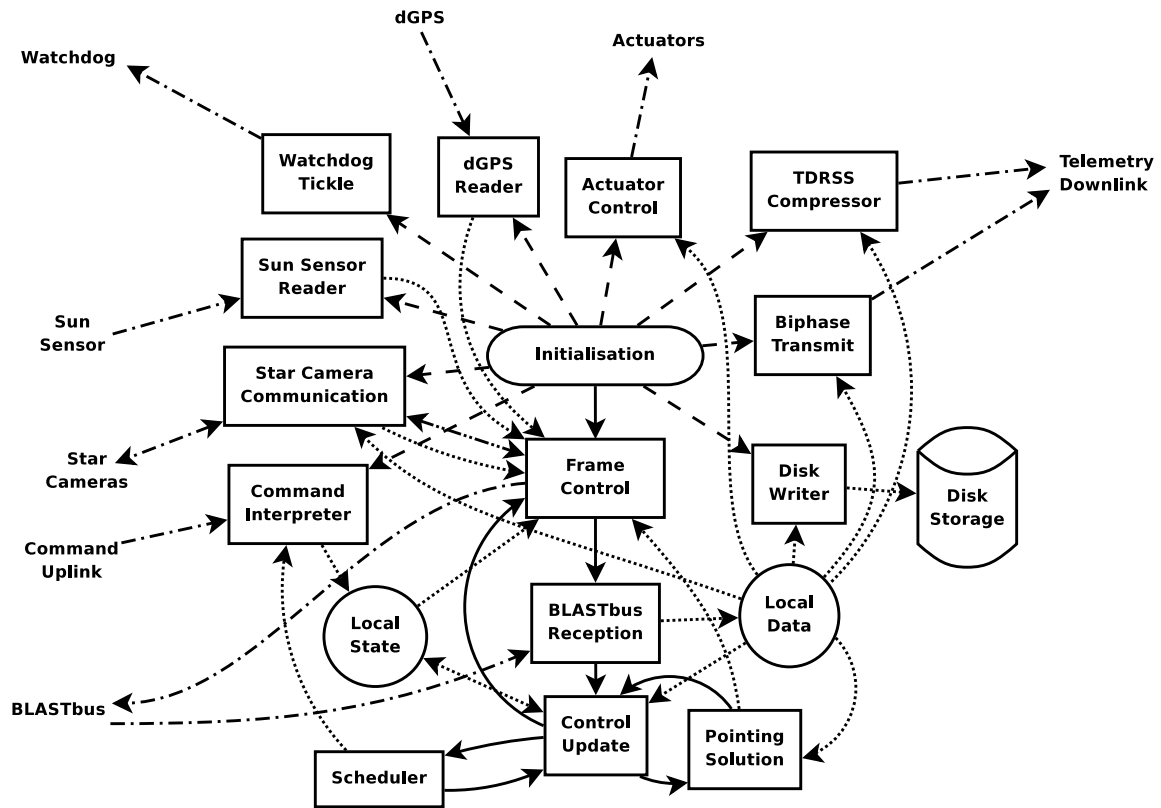


Figure 2.3: A schematic representation of major subsystems of the Master Control Program (mcp). The program entry-point is the initialisation system (oval). Subprocesses are rectangles; in-memory data are circles. Dashed lines denote thread instantiation, all of which is done by the initialisation system. Solid lines indicate program flow in the main thread. Dotted lines indicate internal data flow while dot-dashed lines are external data connections. The double-dot-dashed line between frame control and the star camera thread indicates the special semaphores needed for star camera synchronisation (see Section 2.11.1.2).

the BLASTbus, and all control data sent over the BLASTbus, including the in-flight pointing solution, are generated by the main thread.

2.3.2.1 Initialisation

The initialisation procedure in mcp sets up the system for normal operation. Early on in this procedure, the initialisation instantiates the watchdog thread in order to prevent power cycling by the watchdog circuit (Section 2.4.3). This watchdog thread does nothing other than send the “tickle” to the watchdog. If the main thread detects loss of data from the BLASTbus, it will stop the watchdog thread, ultimately resulting in a reboot

of the computer. Initialisation includes the write of the initial BLASTbus frame to the BLASTbus Controller (BBC). Once the system has been set up, the initialisation procedure spawns all other subthreads.

2.3.2.2 BLASTbus Reception

In order not to fall behind data acquisition, the main thread must keep up with data coming in from the two BLASTbuses. The main thread repeatedly polls the BBC for new data. These data are received one 32-bit sample at a time, which includes a 16-bit address and 16-bits of data, as described in Figure 2.13.

All samples which do not have the write bit set are immediately discarded. These are poll requests generated by the BBC which contain no data, and amount to approximately 48% of all traffic on the BLASTbus. The address portion of the remaining words is used as the input to a look-up table which returns the position of the datum in the local data frame. This look-up is mildly inefficient in memory usage; of the 2^{13} entries in the table, only 1077 positions (13%) are actually filled, but results in very fast address look-up.

Whenever a frame sync word arrives from the ACS side bus, the main thread runs all its control processing routines. As a result, these routines are executed at the frame frequency of 100.16 Hz, providing an accurate time reference for the control subsystems to use. A few controls which only update the 5 Hz slow data are run only once every twenty frames. Execution of these routines is triggered on the reset of the multiplex index to zero. Because the BBC keeps a static copy of the entire frame, only those channels which need to be updated are modified by these control routines.

2.3.2.3 Frame Control

In addition to BLASTbus reception, the main thread also handles writes to the BLASTbus. These writes consist of two classes of data:

- control data generated by the control update subsystem (see Section 2.3.2.4),

- data not directly collected by the BLASTbus, such as data from the star cameras, sun sensor, and differential GPS, as well as data computed by `mcp` itself, such as the in-flight pointing solution. These data are written to the BLASTbus after the fact for archival purposes (the so-called “writebacks”).

Writes to the BLASTbus Controller by frame control happens asynchronously to the transmission of data by the BBC on the BLASTbus itself, so writes which take up more than one BLASTbus channel may be written out on different frames. For the bulk of the control data, this is not problematic for BLAST in flight, since they consist of independent, single word channels. In instances where this does matter, care is taken to ensure proper behaviour in the event of control data being so split across frames. Typically this is done by passing, along with the multi-word control data, a semaphore which indicates when an update of the multi-word data has completed.

2.3.2.4 Control Update

In addition to actually transferring the control data to the BLASTbus, the main thread is also responsible for performing the necessary calculations required to update these data as needed. These data command the systems controlled through the ADC cards, including the pointing motors, heaters, pumps, and DAS tuning. An important part of the control update is the calculation of the in-flight pointing solution. This is necessary to provide information for the motion control system to command the motors. Calculation of the in-flight pointing solution is discussed in Section 2.3.3. Data generated by the control update is recorded in the local state as well as written to the BLASTbus by frame control.

2.3.2.5 The Command Scheduler

The final task of the main thread is the automatic execution of pre-scheduled observations and other tasks. Before launch, the flight computers are loaded with several *schedule*

files which contain a list of commands indexed by local sidereal time (LST) for the entire length of the flight. Because the geographical position of the gondola at any given time is unknown, LST is used instead of solar time to simplify the creation of these schedule files.

Because the flight lasts more than one day, the time in the schedule file also includes a local sidereal day (LSD). To define the LSD, the schedule file contains a date indicating the start of LSD 0. This is the time (in UTC) when LST is zero, at the start of LSD 0, for an observer at 0° longitude. This time defines the *schedule-file epoch*, relative to which all times in the schedule file are measured.

Use of a LSD also requires the definition of a sidereal date line (SDL) where the discontinuity in LST occurs. For both LDB flights, the SDL was along 180° longitude. For *BLAST05*, this was outside the flight path of the platform. For *BLAST06*, the SDL was approximately 13° east of the launch site, a condition designed to ensure that crossing the SDL would occur as late in the flight as possible, the prevailing flight path being westward. When the payload crosses the SDL, the LSD moves back a day, causing one day of the scheduled flight-plan to be repeated. The flight-plan is designed with this in mind.

When the payload is in schedule mode, the command scheduler is queried at 100 Hz. If the command associated with the current LST is different than the last command sent, the new command is submitted to the command interpreter (see Section 2.13.4.2).

Because the flight path of the payload is not known in advance, the flight-plans calculated before flight must be flexible enough to accommodate many different paths. An estimate of the rate of change of longitude is used when creating these flight-plans. Differences in the longitude rate from the estimate will result in more or less time being spent on a given field. As a result, the scheduled scans have a reasonable safety margin in scan time.

Changes in longitude are more problematic, since they affect the visibility of targets.

To help minimise this, three different sets of flight-plans are created, each optimised for a different latitude band. As the platform drifts in latitude, the command scheduler will move between these flight-plans as appropriate.

For each latitude band, two flight-plans are created, one optimistic, one pessimistic, for a total of six different flight-plans. The pessimistic flight plan is designed to accommodate poor sensitivity from the detectors and poor performance from the power system. Accordingly, the pessimistic plan spends observes fewer targets, allowing longer integration time on each, as well as scanning less aggressively from the anti-solar azimuth.

The pessimistic flight-plan is selected at launch. During the early part of the flight, while within line-of-sight, the sensitivity of the detectors is measured to evaluate whether to switch to the optimistic flight-plan (see Section 3.2.1). For *BLAST05*, problems with the optics resulted in the sensitivity being below even the level of the pessimistic plan (see Section 3.3.2). As a result, the command scheduler was not used at all. Instead, targets were selected manually, and commands were sent regularly from the ground. For *BLAST06*, the sensitivity of the detectors was found to be roughly nominal, and the optimistic flight-plan was followed by the scheduler.

The command scheduler is also able to automatically perform the minimum required to prepare the platform for normal operation upon reaching float altitude: unlocking the inner frame, turning on the motors, and opening the pumped pot valve. This ability is provided only in case the command uplink fails. In practice, these tasks are performed from the ground via the uplink.

As a fail-safe measure, the commanding system has a built-in time-out. If a command is not received by the commanding system with the specified time-out period, `mcp` will engage the command scheduler. This protects against unexpected losses in telemetry uplink. The time-out period may be changed from the ground while in flight.

Sensor	Co-ord.	Rate	σ_{sys}^2	w_{samp}^{-1}
El. Encoder	El.	100.2 Hz	400	3600
El. Clinometer	El.	100.2 Hz	3600	3600
Star Camera {	El.	~ 0.5 Hz	0.04	0.04
	Az.	~ 0.5 Hz	0.04	0.09
Magnetometer	Az.	100.2 Hz	8100	3600
Differential GPS	Az.	~ 5 Hz	225	400
Sun Sensor	Az.	~ 10 Hz	3600	900
Null	Az.	—	3.6×10^7	36

Table 2.1: Characteristics of the absolute pointing sensors used in the in-flight pointing solution. Column three gives the rate at which new pointing data arrives from the sensor, assuming nominal operation. Column four lists the systematic variance and column five the inverse of the sample weight for each sensor for *BLAST06*. Units for both columns are $(\prime/s)^2$. These values are empirically determined from pointing tests done before flight.

The Null sensor listed is a special azimuthal solution which is not associated with any physical sensor and, hence, is never updated. As a result, the Null azimuth solution consists simply of the integrated gyros, starting from an arbitrary point. This is useful in the lab where none of the other azimuth sensors work. Its exceptionally large systematic variance ensures that this solution is ignored whenever any other azimuth solution is available. A similar Null elevation solution is not required, since the elevation encoder provides elevation data, even in the absence of external references.

2.3.3 The In-flight Pointing Solution

Because BLAST is a pointed platform, `mcp` must generate an in-flight pointing solution to better than the array size to ensure targets are actually observed. The BLAST in-flight pointing solution is designed to be accurate to roughly $1'$. The primary job of the pointing solver is integration of the gyros. Assuming the orientation of the telescope is known at a given time, the pointing solution is evolved by integrating the gyroscopes, yielding the orientation of the telescope at any subsequent time. The pointing solver is placed in the main thread in order to synchronously integrate the gyros.

The gyros have random walk noise $\sim 4\text{--}5''/\sqrt{s}$. As a result, absolute measurements of the orientation of the telescope are needed every 2 s or better to prevent the integrated gyro error from dominating the pointing solution. Absolute pointing measurements acquired by the suite of BLAST pointing sensors are fed to the pointing solver. Each sensor is characterised by a systematic variance and a sample weight as outlined in Table 2.1.

2.3.3.1 Pointing Solution Calculation

A separate pointing solution is calculated for each absolute pointing sensor. This solution consists of the pointing data received from that sensor, which is then evolved in time by integrating the gyroscopes. Azimuth and elevation solutions are computed independently.

A sensor's pointing solution as follows: if α_{i-1} is the pointing solution from the last time step, the new pointing solution, α_i , is, in the case where there is no new data from the sensor:

$$\alpha_i = \alpha_{i-1} + \dot{\alpha}_i \Delta t \quad (2.1)$$

where $\dot{\alpha}_i$ is the angular velocity read from the gyroscopes (see Equation 2.6) and Δt is the time step between samples (≈ 10 ms). If there is a new datum from the sensor, ξ_{i-n} , which describes the orientation n samples ago, the new pointing solution is, instead:

$$\alpha_i = \frac{\sigma_{\text{sys}}^{-2} \alpha_{i-n} + w_{\text{samp}} \xi_{i-n}}{\sigma_{\text{sys}}^{-2} + w_{\text{samp}}} + \Delta t \sum_{j=0}^{n-1} \dot{\alpha}_{i-j} \quad (2.2)$$

where σ_{sys}^2 and w_{samp} are taken from Table 2.1. For most pointing sensors, $n = 1$, since they are sampled at 100 Hz. The exception is the star cameras, where the pointing data is delayed due to the computation time required by the star camera. The accumulated error in the solution is similarly calculated:

$$\sigma_{\alpha_i}^2 = \begin{cases} \sigma_{\alpha_{i-1}}^2 + \sigma_{\text{gyro}}^2 & \text{with no new data} \\ \sigma_{\text{sys}}^2 + (\sigma_{\text{sys}}^{-2} + w_{\text{samp}})^{-1} & \text{with new data} \end{cases} \quad (2.3)$$

where σ_{gyro}^2 is a measure of the variance of the gyro signal. For BLAST, the gyro variance is dominated by uncertainty in the calculation of the gyro offset (see Section 2.3.3.3), and is set to 0.0072 ($''/s$)²

2.3.3.2 Pointing Solution Summation

After calculation of the individual pointing solutions, the solutions are summed to make the nominal pointing solution. For each angle (azimuth and elevation), the solution, S_i , is computed as:

$$S_i = \frac{\sum (\sigma_{\alpha_i}^2 + \sigma_{\text{sys}}^2) (\alpha_i + \delta)}{\sum (\sigma_{\alpha_i}^2 + \sigma_{\text{sys}}^2)} \quad (2.4)$$

where the sums are over unvetoes sensors, and δ is a operator-settable trim, which may be different for each sensor. In flight, trims for each sensor can be uploaded from the ground. This allows for correction to the sensor alignment done on the ground before flight. Sensors can also be vetoed, which excludes them from the sum in Equation 2.4 in case they aren't working as desired. An early procedure in flight is to set the trims and vetoes to tune the pointing solution.

2.3.3.3 Gyro Offset Calculation

The gyro signal includes a slowly varying bias which must be removed before the gyros can be integrated by the pointing solution. As a result, the pointing solution subsystem must also estimate this gyro offset. Whenever a new reading, ξ_i , arrives from one of the sensors, a new gyro offset estimate, o_i , is generated as:

$$o_i = \frac{\xi_i - \xi_k}{\Delta t(i - k)} - g_{i-k} \quad (2.5)$$

where ξ_k is the previous sensor reading, which occurred on sample k and $g_{i-k} \equiv \sum_{j=0}^k \dot{\alpha}_{i-j}$ is the raw integrated gyro between samples k and i . This estimate is fed into a finite input response (FIR) filter with weighting that approximates a Gaussian in eight steps. The length of this filter is 30 minutes, and a new mean gyro offset estimate, $\langle o_i \rangle$, is emitted from the filter once every filter length.

A separate gyro offset estimate is calculated for each coarse sensor and these estimates are summed analogously to Equation 2.4. As a result, only gyro offsets from unvetoes

sensors are used by the pointing solution, yet vetoed sensors still calculate an estimate, so that if they are later unvetoed, their contribution to the gyro offset will be reasonable at the outset.

The gyro signal fed to Equation 2.1 is now calculated as:

$$\dot{\alpha}_i = \dot{\alpha}_i^{\text{raw}} - \langle o_i \rangle - \dot{\alpha}_{\oplus} \quad (2.6)$$

where $\dot{\alpha}_i^{\text{raw}}$ is the raw gyro signal, and $\dot{\alpha}_{\oplus}$ is the component of the Earth's rotation in the direction that the particular gyroscope measures. A plot of the calculated gyro offset estimates are presented in Figure 2.4.

A comparison between the gyro offsets calculated in-flight during *BLAST06*, and the gyro offsets recovered in post-flight pointing solution calculation (see Section 4.3.2) is presented as Figure 2.5. Scatter in the in-flight gyro offsets is roughly equal to the noise in the post-flight reconstruction of the gyro offsets, indicating that the in-flight gyro offset calculation is sufficient for in-flight pointing.

2.3.4 On-board Data Storage

Because the TDRSS downlink bandwidth (Section 2.13.2) is insufficient to send the entire data stream, all data is stored on local hard-drives by the flight computers. The hard-drives, along with the other components of the flight computers, are housed in a data storage vessel pressurised to one atmosphere (see Section 2.4.2) to provide a suitable environment for proper operation of the drives. Recovery of the pressure vessel after the flight is required for mission success.

A dedicated subsystem, the disk writer, handles dumping of data to disk. This subsystem runs in its own thread so that latencies associated with disk I/O are minimally disruptive to the frame control in the main thread. Data acquired by the BLASTbus reader are written to a ring buffer by the main thread, which is read asynchronously by




Figure 2.4: Gyro offsets calculated from the in-flight pointing solution for *BLAST06*. The high-frequency noise in the calculated offset is around $2''/s$. At $+82^h$ Gyro2 failed. As a result, the calculated offset for Gyro2 rose to to $+4.93''/s$ in an attempt to compensate. At $+83^h$, the in-flight pointing solution was switched to use the second set of gyroscopes. This resulted in recovery of Gyro2, and a step in the gyro offset of all three gyro signals due to different underlying hardware. (See, further, Figure 3.11.) The impulse response time for the gyro offset filter is estimated from these excursions to be approximately 50 minutes, during which time the in-flight pointing solution is poor. Since gyro offsets are recalculated post-flight for both sets of gyroscopes, this does not seriously degrade the post-flight reconstruction of the pointing solution.




Figure 2.5: Gyro offsets calculated in-flight (dotted curve) compared with gyro offsets calculated during post-flight pointing reconstruction (solid curve; see Section 4.3.2) for a 24-hour period of *BLAST06*. The two curves are offset due to being calculated in different reference frames. The 30-minute delay in the in-flight calculated offsets is clearly apparent.

the disk writer. This ring buffer is 3000 frames (~ 30 seconds) long, much larger than any latencies inherent in the system.

Data is written to disk as raw 100 Hz frames to regular files in on a third extended filesystem (ext3). The ext3 filesystem was used to provide robustness under computer crashes while writing data to disk. To allow fast bootstrapping of the MCCs, no filesystem integrity checks are performed at start-up. Lab tests indicated that, while corruption could occur under these conditions, the corrupted ext3 filesystem could still be used for writing new data and reading old data.¹

To keep file sizes manageable, a new file is started every 90 000 frames (~ 15 minutes or ~ 100 Mbytes of data). These file sets, known as *framefiles*, can be subsequently converted to a format readable by analysis software by `defile` (see Section 2.13.1.3).

¹That is, the only problem associated with filesystem corruption occurs when modifying existing files, for instance, deleting them. Since the only operation performed during flight is writing new data, and the only necessary operation to recover the data is reading old data, this sort of corruption of the ext3 filesystem is not mission critical.

Flight	CPU Speed	
	MCC0	MCC1
<i>BLAST03</i>	200	350
<i>BLAST05</i>	366	650
<i>BLAST06</i>	366	850

Table 2.2: Processor speed, in MHz, of the two flight computers (MCC0 and MCC1) for the three BLAST flights. The table shows a general trend towards increasing speed for the flight computers. This was deliberate: we were attempting to measure the rate of CPU upsets at float as a function of CPU speed. To the accuracy we can measure this, it appears flat: approximately one upset every four days.

2.4 The Flight Computers

The Master Control Program (`mcp`) runs on a redundant pair of flight computers (the Master Control Computers, MCC). The two computers are functionally equivalent, and consist of a passive backplane, a single board computer containing a Intel Celeron processor and 512 Mbyte of RAM, a BLASTbus Controller (BBC, Section 2.9.2), and multiport a serial card. The only substantive difference between the two computers is the speed of their processors (see Table 2.2).

2.4.1 Operating System

The computers run a generic Linux[®] operating system. A custom kernel module for the BBC Card was written to allow communication with it. The Linux kernel used is an unmodified 2.6 series kernel. Specifically, it has not been patched to allow real-time execution. Rather, the kernel is optimised for low-latency operation.

The *BLAST03* test flight occurred before the release of the first 2.6 series kernels. As a result, this flight was flown using the 2.4 series kernel. This series had no low-latency configuration, and latencies as high as 250 ms were seen, requiring a much larger BLASTbus latency than used on subsequent flights. This was the primary reason behind the switch to the 2.6 series of kernels. The first 2.6 series kernel used was 2.6.0_test11. A few subsequent upgrades were needed to get full hardware support from these early

2.6 kernels, and both *BLAST05* and *BLAST06* were flown with the 2.6.8.1 kernel of 14 August, 2004.

2.4.2 The Pressure Vessel

The flight computers are housed in a cylindrical pressure vessel (PV) pressurised to one atmosphere, which is required for proper operation of the hard-drives (see Figure 2.6). The flight computers are mounted on a shelf inside the PV which is itself attached to the PV front plate. The entire unit: computers, shelf, and front plate, can slide freely in and out of the pressure vessel, and the shelf can be completely removed to permit working on the flight computers on the bench (as in Figure 2.6.1). Power and signals for the flight computers are all routed through a pair of 50-DD subminiature hermetic feedthroughs on the PV front plate. A break-out box (the PVBOB) is mounted on the outside of the front plate and distributes the signals from the hermetic feedthroughs to numerous external connections for ease of cabling.

Running the flight computers inside a pressure vessel also assists with cooling. The computers used are built of stock industrial components. Specifically, the Celeron CPUs use an air-cooled heatsink and fan to dissipate heat. The air inside the pressure vessel transports this heat to the exterior of the vessel, where it is radiated away. The hard-drives also have cooling fans.

In addition to the two flight computers, the pressure vessel also contains a master Ethernet hub, and the flight computer watchdog circuit (see Section 2.4.3). The PV is supplied with +28 V battery power. Dedicated DC–DC converters inside the PV produce the +5 V, +12 V, –12 V, and +3.3 V required to run the flight computers, as well as the +5 V used for the BLASTbus isolation.



Figure 2.6.1: The flight computers.




Figure 2.6.2: The pressure vessel mounted on the gondola.

Figure 2.6: The BLAST pressure vessel (PV) and flight computers. Figure 2.6.1 shows the flight computers mounted on the PV shelf. MCC0 (*frodo*) is at left, MCC1 (*sam*) is at right. Above MCC0 is the Ethernet hub. Above MCC1 is the watchdog card. The hard-drives and DC-DC converters are mounted underneath the shelf. The PVBOB can also be seen at right, attached to the PV front plate. Figure 2.6.2 shows the PV mounted on the gondola; the front plate is on the right. The PV rests on L-shaped rails and is tied down by hose clamps and aircraft webbing for flight.




Figure 2.7: A schematic representation of the BLAST watchdog circuit. See text for a description of the circuit. In addition to the watchdog circuit itself, the watchdog board also includes the three extra components shown at bottom: a 28V→5V DC–DC converter, which generates the power to drive the 5-V isolated BLASTbus, and ambient pressure and temperature sensors to monitor the pressure vessel. Full component schematics are given in Appendix A.

2.4.3 The Watchdog Circuit

The BLAST watchdog circuit is designed to monitor the flight computers and power cycle them if an error is detected. The watchdog circuit is built of low density components to reduce potential upsets from cosmic rays, and is housed inside the pressure vessel (see Figure 2.6.1).

A block diagram of the watchdog circuit is presented as Figure 2.7. Timing for the circuit is provided by a 555 timer operating in astable mode. This clock signal has a period of 1.162 seconds. The clock drives a pair of countdown timers, one for each computer. When a countdown timer reaches zero, the computer it is monitoring is power cycled by the watchdog by momentarily asserting the reset line for the computer’s DC–DC converter. After power cycling the computer, the countdown timer is inhibited for 4.96 minutes to allow time for the computer to boot.

The computers prevent this reboot by feeding the watchdog a “tickle” which resets the countdown. The “tickle” consists of three lines coming from the parallel port alternating between logic 010 and logic 101. This “tickle” is generated by a dedicated subprocess in `mcp` (Section 2.3.2.1). Interruption of this signal indicates an upset in `mcp`, and is remedied by a power cycle of the affected computer.

The watchdog circuit also selects which of the two computers has control of the platform. This control select is done by enabling the BBC transmitter on one of the computers and simultaneously routing that computer’s serial transmit lines to the rest of the gondola. The computer which is not in charge is still able to listen to both the BBC and the serial lines, but data it tries to send are not propagated past the watchdog.

The algorithm used to select which computer is in charge is simple: when the computer that is currently in charge is rebooted, control passes to the other computer. No change in control happens when the computer not in charge is rebooted. Which computer is in charge at power-up is arbitrary: the watchdog will select whichever computer starts running first. A boot inhibit is enabled at power-up to ensure the countdown timers are properly cleared, and the MCC have time to boot, before the MCC monitors start operating. Full watchdog component schematics are presented in Appendix A.

2.5 Receiver

BLAST has three bolometric detector arrays, one for each band, similar to the ones developed for the SPIRE instrument [25]. These bolometers are fed by smooth-walled conical feeds which provide a tapered illumination of M4 and, hence, the primary. The detectors are read out by a custom-built digital lock-in amplifier.

2.5.1 Detectors and Feeds

The three BLAST arrays are composed of 149, 88, and 43 micromesh (“spider-web”) bolometers at 250 μm , 350 μm and 500 μm respectively (Figure 2.8). The bolometers are coupled to the optics by smooth-walled conical feedhorns spaced at intervals of

$$d = 2F\lambda \quad (2.7)$$

where F is the focal ratio of the complete telescope optics. This optimises the aperture efficiency of the detectors [26, 27]. A 2λ -long waveguide at the base of each horn provides the band-defining high-pass (Section 2.10.5).

2.5.1.1 The Bolometer Device

Each bolometer (Figure 2.8.1) consists of a silicon nitride (Si_3N_4) mesh absorber, which has good absorption efficiency over a wide frequency range, and low heat capacity, while having a relatively small cosmic ray cross-section, an important consideration at balloon altitudes. A neutron-transmission-doped (NTD) germanium thermistor is glued to the absorber and measures the temperature of the mesh. Measurement of the temperature of the mesh gives a direct measurement of the power deposited on the absorber.

The resistance profile of the NTD thermistor is given by:

$$R_b(T) = R_0 e^{\sqrt{\Delta/T}} \quad (2.8)$$

where $R_0 = 55 \Omega$ and $\Delta = 50 \text{ K}$ characterise these bolometers [28]. The thermal response of the absorber is described by the thermal time constant, τ_T , while the electrical response of the thermistor itself is given by a second, electrical time constant, $\tau_e < \tau_T$. The thermal time constant of the BLAST detectors have been measured using cosmic ray hits as described in [29, 30], and found to be $\tau_T \approx 2 \text{ ms}$ (see, further, Section 4.2.1). The



Figure 2.8.1: 500 μm Spider-web BolometerFigure 2.8.2: 500 μm ArrayFigure 2.8.3: Complete 500 μm Assembly

Figure 2.8: The BLAST detectors. Figure 2.8.1 shows an individual 500 μm spider-web bolometer. The device consists of the Si_3N_4 micromesh absorber, and the rectangular NTD thermistor pad at bottom left. Figure 2.8.2 show a portion of the complete 500 μm array. Figure 2.8.3 show the fully assembled 500 μm bolometer detector array (BDA), including the feedhorn block.

bolometers' noise equivalent power (NEP) is:

$$\text{NEP} = \sqrt{4k_B T_0^2 G} \quad (2.9)$$

where G is the thermal conductance of the absorber, and T_0 is the bath temperature. For a given background load, Q , the maximum sensitivity occurs when [31]: $G = \frac{Q}{T_0}$. The BLAST arrays have been optimised for loads, Q , of 55, 40 and 30 pW at 250 μm , 350 μm and 500 μm [28].

2.5.2 Read-Out

A schematic representation of the readout electronics is presented in Figure 2.9. To reduce the effects of $1/f$ noise in the detectors, the bolometers are biased with an AC voltage:

$$V(t) = V_0 \sin(2\pi\nu t) \quad (2.10)$$

where ν , the bias frequency, is chosen to be much higher than the frequency of the signal on the sky. For both BLAST flights, $\nu = 200.32$ Hz. The NTD thermistor is biased in series with a pair of load resistors, with a resistance, $R_L \gg R_b$, designed to decouple the current through the bolometer from the voltage drop across the bolometer. Each bolometer signal is amplified by a differential JFET follower with noise 5–7 nV Hz^{1/2}. The JFETs are located in the JFET cavity (see Section 2.6), to place them as close as possible to the bolometers themselves, in order to capacitively dissociate the transmission lines from the bolometer signal at as early a stage as possible.

From the JFETs, the signals pass out of the cryostat and to the Receiver rack, where they are further amplified and filtered using Analog Devices AD624 instrumentation amplifiers. Here the signals are filtered with a 100-Hz-wide bandpass centred around 208 Hz. The signals are then sent to the Data Acquisition System (DAS; Section 2.7) for digitisation, lock-in, and further filtering (Section 2.8).




Figure 2.9: A schematic representation of the bolometer read-out electronics.

2.6 The Cryostat

The detectors must be cooled to cryogenic temperatures for proper operation. A long hold time cryostat (Figure 2.10) cools the detectors to ~ 300 mK and the re-imaging optics to ~ 1.5 K. This is accomplished in four stages using liquid cryogenes: a 1 atmosphere liquid nitrogen (LN) stage, a 1 atmosphere liquid helium (LHe) stage, a pumped ^4He pot, and a closed-cycle ^3He refrigerator. The cryogen boil-off rate is the limiting factor that sets the maximum flight time of the payload. The BLAST cryostat has been designed to have a hold time of > 11 days. Full details of cryostat construction and operation are presented in [32].

The cryostat holds ~ 43 ℓ of liquid nitrogen (LN) and ~ 32 ℓ of liquid helium-4 (LHe) in two cylindrical tanks in the upper half of the cryostat. These provide cryogenic stages at 77 K and 4.2 K respectively. Boil-off from the LHe tank passes through a heat exchanger attached to an intermediate stage between the LN stage and the LHe stage. This vapour-




Figure 2.10: A schematic cross-section of the BLAST cryostat. The cryostat is divided into two parts by the cold plate. The liquid cryogen tanks and the JFET cavity fill the top half of the cryostat. The optics box fills the bottom half of the cryostat. The vapour-cooled shield separates the LN and LHe stages, reducing the load on the helium. Light from the Cassegrain telescope enters the optics box through the window on the left. The ^3He fridge, which hangs from the cold plate behind the optics box in this view, is not shown. It can be seen in Figure 2.21, where the cryostat has been turned upside-down and the bottom half, containing the optics box, exposed.

cooled shield (VCS) sits around 35 K, and reduces the load on the LHe stage. These three shields fully surround the bottom half of the cryostat containing the optics box. The optical window in each stage contains an IR blocking filter to reduce thermal loading.

A pumped ^4He pot cools the optics box and cold mirrors to 1.5 K. The pot, a small helium bath, is open at float to the atmosphere via a pump line, keeping pressure inside the pot near vacuum (~ 3 millibar). On the ground, the pump line is attached to a vacuum pump. Just before launch, the vacuum pump is removed and the valve on the pump line is closed for ascent. Once reaching float altitudes, the valve is opened again. The pumped pot is connected to the main LHe tank by a capillary. The pressure difference between the pumped pot and the main tank (which is kept near one atmosphere) forces helium into the pot. The length and diameter of the capillary regulate the flow of helium.

From 1.5 K, the BDAs are cooled to 300 mK by the ^3He refrigerator. The fridge contains a charcoal getter which absorbs gaseous ^3He when the charcoal is cold. Once all the ^3He has been absorbed, the charcoal heater is turned on. When the charcoal is heated above ~ 10 K, it starts outgassing. The outgassed ^3He passes through a heat exchanger connected to the pumped pot and condenses, collecting in the ^3He cold stage. This process is repeated once the ^3He re-evaporates. A complete cycle takes approximately two days, and, accordingly, the fridge must be cycled several times during flight to keep the bolometers cool.

In flight, the ^3He fridge cycle is controlled automatically by `mcp` which powers on and off the charcoal heater as appropriate. During a fridge cycle, the temperature of the bolometers rises above 300 mK, and it takes about 2.5 hours before they are back down to their nominal temperature, during which time scientific data cannot be collected.

Cryogenic thermometers provide housekeeping data on the health of the cryostat. Two types of sensors are used. For temperatures above LHe, silicon diodes are used. Below LHe temperatures, ruthenium oxide resistors (ROxes) are used. These housekeeping data are read out by a dedicated ADC card in the DAS crate. The ROxes are AC-biased in

a similar manner to the detectors themselves, and require the same digital lock-in to be read out.

The JFET modules are located in the JFET cavity in the centre of the cryostat. This places them maximally close to the bolometers themselves, reducing capacitive pickup in the bolometer read-out. A heater in the JFET cavity keeps the JFETs above their 100 K minimum operating temperature. The heater is controlled by a simple thermostat routine in `mcp` which set the heater level, L , per:

$$L = \begin{cases} 0 & T_{JFET} > 135 \text{ K} \\ 9 - \frac{T_{JFET}}{15 \text{ K}} & 120 \text{ K} < T_{JFET} < 135 \text{ K} \\ 1 & T_{JFET} < 120 \text{ K} \end{cases} \quad (2.11)$$

where 1 is full scale and T_{JFET} is the measured temperature of the JFET cavity.

2.7 The Data Acquisition System

Most of BLAST's signal processing is done digitally, using custom-built, multipurpose analogue-to-digital (ADC) cards. These cards provide an integrated solution for digitisation, and filtering of analog signals, as well as digital control. They communicate with the flight computers over the BLASTbus (Section 2.9), ensuring synchronisation of data. Although created to fill BLAST's needs, the cards were designed to be sufficiently generic, so that they have been used on subsequent projects, thus amortising their development cost.

The BLAST Data Acquisition System (DAS) is composed of fourteen of these multipurpose read-out cards known as *ADC cards* (Figure 2.11). Two additional ADC cards in the DAS rack provide capabilities for cryostat and bias generator card logic, and general housekeeping. Furthermore, the Attitude Control System (ACS, Section 2.12) is built around four more ADC cards.




Figure 2.11.1: ADC Card

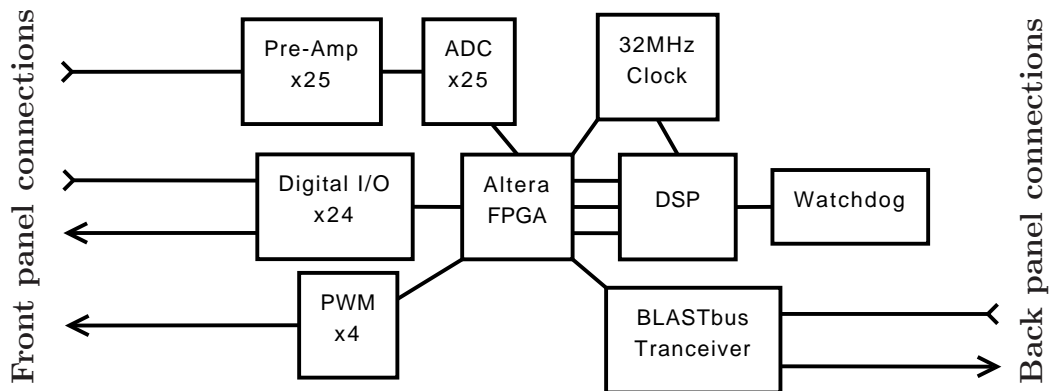


Figure 2.11.2: ADC Card Schematic

Figure 2.11: The BLAST multipurpose ADC card. Figure 2.11.1 shows a typical BLAST ADC card. Important subsystems are highlighted. A block diagram of the major subsystems of the card and their interconnections is presented in Figure 2.11.2. The Altera[®] FPGA provides inter-system communication for most components on the board. The bulk of signal processing occurs on the DSP.

The ADC cards are named for their twenty-five $\Sigma\Delta$ analogue-to-digital converters (ADCs, Burr-Brown ADS1252). The $\Sigma\Delta$ s are fed by the on-board, low-noise differential instrumentation amplifiers (Burr-Brown INA2128), which are configurable for several different analogue input ranges via jumpers on the boards. The 24-bit $\Sigma\Delta$ s are sampled at $10416\frac{2}{3}$ Hz. The digitised signals are processed on a 32-bit digital signal processor (DSP) running at 32 MHz. This DSP (Analog “SHARC [®]” ADSP21062) has 2 Mbits of on board RAM.

In addition to the twenty-five analogue inputs, the ADC card is also equipped with twenty-four digital I/O lines, ganged into three groups of eight. Each group can be used as either an input or an output group. Four pulse-width modulated (PWM) drivers provide further digital outputs. The ADC card was also designed to provide a digital-to-analogue converter (DAC) channel, however, this feature was not needed by BLAST, and the DAC channel has never been implemented.

A single ADC card requires six different external supply voltages: +10 V, +7.5 V, +5 V, +3.3 V, +2.5 V, and -10 V. All voltages are provided by dedicated DC-DC converters located in the ADC card racks. An on-board watchdog can detect faults in the DSP and power cycle the card in such an event.

2.8 DSP Firmware

The DSPs are programmed in standard C. The programs are burned onto UV-erasable EPROMs, making modification of the DSP firmware relatively easy. The basic DSP program is similar for all cards. A $10416\frac{2}{3}$ Hz interrupt handler routine services the ADC interrupt. This interrupt handles the bulk of the analog signal processing including: phase-lock, lock-in, and lowpass filter. A second interrupt handler services the 100 Hz BLASTbus frame-sync interrupt, and takes care of synchronous writes of data to the BLASTbus.

The remainder of the firmware is special purpose code, which varies from card to card, based on the role the card plays, primarily concerning digital I/O. Typically this special purpose code runs in the main loop of the firmware. It is executed whenever the DSP is not processing an interrupt, which results in a typical main loop frequency for this code of ~ 600 Hz.

2.8.1 Lowpass Filter

All analogue inputs, including the bolometer signals, pass through a multistage, boxcar lowpass filter. The i^{th} stage boxcar of the n stage filter has a length, L_i , given by the geometric sequence:

$$L_i = \left\lfloor \frac{F_{\text{SR}}}{F_{\text{Ny}}} \left(\frac{1}{2} \right)^{i/n} \right\rfloor \quad \text{for } i = 0, \dots, n-1 \quad (2.12)$$

where $F_{\text{SR}} = 10416\frac{2}{3}$ Hz is the ADC sample rate, $F_{\text{Ny}} = 50.08$ Hz the filter Nyquist frequency, and $\lfloor \cdot \rfloor$ indicates the floor function. The choice of Equation 2.12 ensures that the first n nulls are equally spaced in $\log(F)$. The limited time available for computation on the DSP allows us at most $n = 4$. This gives filter lengths of 208, 175, 147, and 123 samples. The filter response is shown in Figure 2.12. The response time for such a filter is given by:

$$t_{\text{res}} = \sum_{i=0}^{n-1} \frac{L_i}{2}. \quad (2.13)$$

For the BLAST case, this gives a response time of 31.344 ms.

Important in the design of the filter is its computational simplicity. Only 3072 clock cycles occur on the DSP between successive ADC interrupts, and, in these limited cycles, the phase-locked loop (PLL), lock-in and filter for all twenty-five input channels must be accomplished, with sufficient time left over for foreground processes. On ADC cards which do not need to lock-in their analogue signals (*i.e.* non-bolometer cards), the removal of the PLL and lock-in would provide sufficient time for an additional, fifth stage on the




Figure 2.12: The transfer function of the BLAST digital filter after successive stages. The 50.08 Hz Nyquist frequency is marked. The full four stage filter (solid curve) has a maximum leakage above the Nyquist frequency of 0.1%, which is aliased back down to lower frequencies.

lowpass. However, this is not implemented, in the interest of keeping the filter response the same for all signals, and minimising code differences.

The DSP's 2 Mbits of memory is divided into two independent 1 Mbit banks, designated PM (Program Memory) and DM (Data Memory). In normal operation, the machine code is loaded into PM, while data is accessed only from DM. In certain cases, operations on both PM data and DM data may occur simultaneously in the same clock cycle. As a result, buffers for successive stages of the filter alternate between PM and DM, optimising computation. The filter is the only part of the DSP firmware where this is done, since the filter buffers plus the firmware image itself largely fill up PM.

2.8.2 The Phase-locked Loop and Digital Lock-in

Before being lowpass filtered, the AC-biased bolometer signals must be locked-in to a reference signal to remove the AC modulation. In addition to bolometer signals, each ADC card in the DAS is sent a reference signal from the bias generator. This Bias Reference is read by the DSP and sent through a phase-locked loop (PLL). This PLL detects the phase of the input bias reference, and then generates a synthetic reference signal, in phase with the bias reference. The synthetic reference is then phase-shifted by a user-settable amount.

Given a bolometer signal (*cf.* Equation 2.10):

$$S = S_0 \sin(\omega t + \delta_S) + N_S(t) \quad (2.14)$$

and the synthetic reference:

$$R = \sin(\omega t + \delta_R) \quad (2.15)$$

where S_0 , is the actual bolometer response, δ_S is the phase shift applied to the bolometer signal by the read-out electronics, δ_R is the user-settable phase shift applied to the reference, and N_S is noise in the bolometer trace, then the locked-in signal is produced by multiplying the bolometer by the reference. Time averaging this extracts the bolometer signal, S_0 .

2.9 The BLASTbus

The ADC cards communicate with the flight computers over a custom RS-485 synchronous serial bus known as the BLASTbus. This bus is composed of three differential pairs (data, clock, and strobe) with a 4 MHz bit clock. The cabling for the BLASTbus also supplies +5 V and matching ground for signal isolation on the bus. The protocol used on the bus consists of 32-bit words, each containing a 16-bit address section and a

16-bit data section.

2.9.1 BLASTbus Architecture

The protocol used by the BLASTbus is a poll-and-response architecture, stipulating a 32-bit request word from the controller on the flight computer, followed by a 32-bit response with 16 bits of data from the ADC card polled. Alternately, a single 32-bit write may be sent by the controller to update data on the addressed ADC card. In this case no response is returned. Given the 4 Mhz bit clock, this protocol results in a bandwidth on the bus of 1 Mbit/s, which is the limit imposed by the BLAST LOS transmitters. However, two separate BLASTbuses are implemented on the platform, potentially allowing an internal bandwidth of 2 Mbits, although this extra potential bandwidth is not used by BLAST.

The format of a 32-bit BLASTbus word is presented in Figure 2.13. Since the BLASTbus address consists of a 6-bit node number and a 6-bit channel number, up to sixty-four ADC cards may be present per bus, each with sixty-four read channels (with the write-bit low) and another sixty-four write channels (with the write-bit high).

The two buses are named after the system that they connect to the flight computers. These are the ACS bus and the DAS bus. Because the bulk of the data acquired by BLAST comes from the bolometers themselves, the DAS bus is by far the busier. In the *BLAST06* configuration, the ACS bus had a throughput of ~ 135 kbits/s, while the DAS bus has a throughput of ~ 756 kbits/s. After accounting for overheads, the combined data stream from the two buses is ~ 889 kbits/s. This combined data stream is what is written to disk and transmitted by the LOS transmitters.

The native data type of the BLASTbus is a 16-bit integer, and most data are sampled as such. The firmware on the DSP rescales input data to ensure they are maximally resolved to 16 bits. Since the ADCs are 24-bits, typically this results in the discarding of the bottom eight bits of the sampled signals. For some of the data channels, including the bolometers and gyroscopes, 16 bits is insufficient resolution. In general, in these

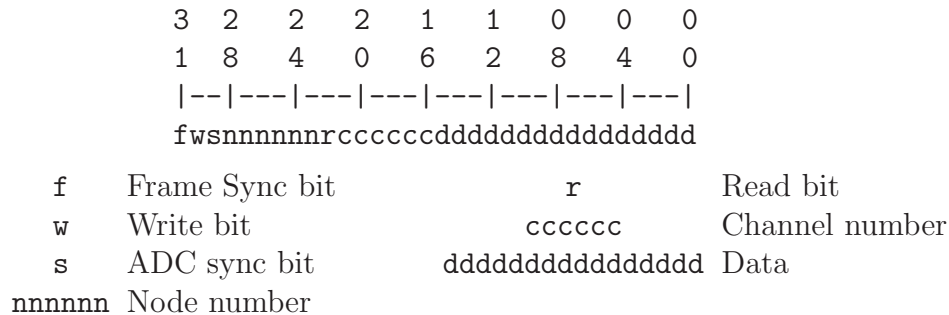


Figure 2.13: A representation of the 32-bit BLASTbus word. The Frame Sync bit is set in the first word of every frame, and generates the 100 Hz interrupt on the DSP. Words can be classified by the state of their Read and Write bits:

<u>r</u>	<u>w</u>	<u>Word Type</u>
0	0	(not used)
0	1	Write by BBC
1	0	Poll by BBC
1	1	Response by ADC card

The ADC sync bit is set when `mcp` detects that an ADC card's DSP has just booted. When encountered by an ADC card, it causes the ADCs to synchronise their outputs.

cases, this is overcome by sampling these channels at 32-bits. This is done by requesting the top 16-bit word and the bottom 16-bit word separately on the BLASTbus and then reconstructing the 32-bit data on the flight computers. Care is needed in this case to ensure that the two 16-bit words obtained are from the same sample of the ADC. This is done by having the DSP write its data only once per 100 Hz frame, accomplished by triggering the writes on the 100 Hz interrupt mentioned above (Section 2.8).

2.9.1.1 Frames

The data stream is broken-up on the flight computers into frames. These frames have a fixed rate of: $4 \text{ MHz}/(384 \cdot 104) = 100.16025641 \text{ Hz}$.² Each signal acquired by the platform is sampled at most once per frame. Specifically, the bolometers, the gyros, and other high-precision data are sampled once per frame.

²The factor of 384 here is set by the ADC hardware (resulting in the $10416\frac{2}{3} \text{ Hz}$ ADC sample rate); the factor of 104 is chosen arbitrarily to give a frame rate near 100 Hz.

Data which are not needed to such high precision, or are slowly varying, are sampled less often, once every twenty frames, resulting in a ~ 5 Hz sample rate. In groups of twenty different channels, these slow data are multiplexed into a single 100 Hz channel, greatly reducing the bandwidth required. As a result there are twenty different 100 Hz frames, all with the same size which form a single 5 Hz superframe. These twenty different frames are distinguished by a multiplex index, which is included in the frames in the data stream.

The original design for the BLASTbus had a frame rate of ~ 60 Hz. As the bolometers were characterised and found to be faster than anticipated, it was decided to increase the frame rate to the 100 Hz used in flight. At 100 Hz, the 288 bolometer channels ($= 12$ DAS cards \times 24 channels/card) from the DAS meant that, at 32-bits, the bolometer data would require a BLASTbus throughput of ~ 923 kbit/s. This was considered too close to the 1 Mbit/s total BLASTbus budget. To overcome this, the bolometer data are rescaled to 24-bits on the DSP, and two bolometer channels are close-packed into three 16-bit BLASTbus channels. In this way, the bolometer portion of the BLASTbus budget drops to ~ 692 kbit/s, without loss in precision. These close-packed bolometer channels are written as-is to disk. They are unpacked later by `defile` (Section 2.13.1.3) in post-processing on the ground.

2.9.2 The BLASTbus Controller

The BLASTbus Controller (BBC) is a PCI-card located on the flight computers. These cards were designed and built after the *BLAST03* test flight. For the *BLAST03* test-flight a different BBC card, which implemented the same BLASTbus protocol, was used, built around the ISA bus. This ISA BBC card suffered several problems, including long latencies in transmission and reception on the bus, as well as being built around an older bus standard. The redesign addressed these problems.

One further problem with the BLASTbus discovered during *BLAST03* was reflections

on the bus. These were tracked down to a flaw in the design of the ADC cards: the BLASTbus line drivers on the cards were placed too far from their connection to the bus. The reflections only became a problem when both the ACS and the DAS were connected in series to the same bus, something not done during development. This problem was worked around for the test-flight by phase-shifting the reflections to a safe part of the bus by increasing the length of one of the BLASTbus cables.

To solve this problem would have required redesigning and rebuilding all ADC cards. This was not feasible. Instead the new PCI BBC card worked around the problem by controlling two electrically independent buses, to which the ACS and DAS were attached separately.

In addition to the two BLASTbuses, the BBC card also generates the biphasic signal used to drive the LOS transmitters (Section 2.13.1), a job that was done in the test-flight by another ISA card which we wanted to replace.

2.9.2.1 BBC Architecture

The BLASTbus Controller is presented in Figure 2.14. It consists of two parts: a PCI card, containing an Altera[®] FPGA, and a daughter card containing the line drivers and other discrete components needed to electrically drive the BLASTbuses. The same PCI card, with no daughter card, is used for the decom card (Section 2.13.1.1).

2.9.2.2 BBC Firmware

The Altera FPGA implements a Nios[®] embedded processor core. This gives the option of writing portions of the BBC firmware in C, rather than writing the complete program in VHDL. The BBC architecture is relatively straight forward: the Nios RAM contains two 5 Hz superframes, one for each BLASTbus. The firmware steps through the superframe, transmitting a channel poll or write as it is encountered in the frame once every 8 ms. At the top of each 100 Hz subframe within the 5 Hz frame is a word with the frame-sync bit




Figure 2.14.1: A BLASTbus Controller Card disassembled.




Figure 2.14.2: An assembled BLASTbus Controller card.

Figure 2.14: The BLASTbus Controller card. Figure 2.14.1 shows the two component cards of the BBC card. On left is the PCI card containing the Altera FPGA and the PCI interface; on the right is the daughter card including the BLASTbus line drivers and external connector. The cards are connected with the 96-pin jumper header on their near edges. Figure 2.14.2 shows the card assembled for insertion in a computer.

set (see Figure 2.13). When the BBC encounters one of these, it waits for a 100.16 Hz interrupt (generated on the BBC card itself) before it continues, ensuring the proper frame rate. Immediately after the end of the superframe is a special end-word which signals the BBC to return to top of the superframe and start again.

All communication on the BLASTbus, both responses from the ADC cards and data generated by the BBC itself (polls and writes) are placed in a FIFO on the Nios. Data from this FIFO are copied off the PCI card by a Linux kernel driver which periodically polls for pending data and presents the data to `mcp`. The full trip latency of the BLASTbus, that is the time from `mcp` transmission to `mcp` reception, is approximately 30 ± 10 ms.

The two superframes in Nios memory may be updated asynchronously by `mcp`, via the kernel driver, in order to send updated data to the ADC cards. Only channels which need to be changed have to be uploaded in this way. In the absence of updated data from `mcp`, the BBC will resend the previous data, mitigating potential problems arising from latencies in the Linux kernel and foreground processes. Kernel latencies in this

configuration rarely exceeded 25 ms. The largest latency observed during *BLAST06* was 147 ms.

2.10 Optics

The optical layout is presented in Figure 2.15. The BLAST optical system is built around a Cassegrain telescope consisting of a ~ 2 m spherical primary mirror (M1) and a correcting secondary (M2). This two element system directs light to the cryogenic re-imaging optics (M3, M4, M5) which includes the Lyot-stop (M4). After passing through the re-imaging optics, the light is split into the three bands by a pair of dichroic beam splitters, which permits simultaneous measurement in all three colours by the bolometric arrays.

2.10.1 The Cassegrain Telescope

For *BLAST05*, a 2 m spherical carbon fibre mirror was used as a primary (M1). This mirror (the *COI mirror*, Figure 2.16.1), was built by Composite Optics Incorporated (COI), and lent to the BLAST Collaboration by NASA. The mirror was gold-plated with a silicon dioxide (SiO_2) coating, providing an RMS $2.4 \mu\text{m}$ surface roughness. A primary backing structure, to allow attachment to the BLAST gondola, as well as a 0.5 m aluminum secondary mirror (M2) and carbon fibre support structure were also built by COI. The secondary mirror corrects for the spherical aberration of the primary.

Due to damage sustained by the COI mirror during the *BLAST05* campaign (see Section 3.3.2), it could not be flown again. Instead a 1.8 m paraboloid aluminum primary mirror was flown in its place for the *BLAST06* campaign. This mirror (the *Bosma mirror*, Figure 2.16.2) was originally fabricated for use in the *BLAST03* campaign, to save the COI mirror from potentially being damaged in the test flight. This mirror was fabricated from a single piece of aluminum by Bosma Machine and Tool Corporation, as was its




Figure 2.15.1: BLAST Optical Layout




Figure 2.15.2: Cold Optics

Figure 2.15: The optical layout of the BLAST telescope. Figure 2.15.1 shows the complete optical system. Figure 2.15.2 shows an expanded view of the cold re-imaging optics. Light collected by the primary (M1) is refocused by M2 onto the Cassegrain focus, behind the primary. The flat M4 is the Lyot-stop, which defines the illumination of the primary for each detector. The dichroic beam splitters, which define the three bands are located between M5 and the focal plane. For *BLAST06*, M2 was able to move in one linear and two rotational dimensions, to permit in-flight focusing.



Figure 2.16.1: The COI mirror assembly



Figure 2.16.2: The Bosma mirror assembly for *BLAST06*, including the recut primary, the 0CTE carbon fibre struts and Active Focus Mechanism.

Figure 2.16: The two Cassegrain telescope assemblies flown on BLAST.

0.4 m aluminum secondary and aluminum secondary support structure.

After *BLAST05*, this mirror was resurfaced by Lawrence Livermore National Laboratory (LLNL) to a surface roughness of $4 \mu\text{m}$ RMS. The secondary, too, was resurfaced by Diamond Turning Incorporated, to near optical quality [33]. This resurfacing was done in order to permit optical alignment tests of the primary telescope. A new support structure, consisting of carbon fibre struts fabricated to have a zero linear co-efficient of thermal expansion (0CTE) along the length of the strut, was also made at this time. The aluminum support structure used in *BLAST03* was replaced by the 0CTE struts to reduce the effects of temperature fluctuations on the focus of the Cassegrain telescope.

2.10.2 Thermal Defocusing of the Telescope

Numerical analyses of the BLAST optical system (Figure 2.17) indicates that, while minor misalignment in some of the optical components is acceptable, the M1–M2 alignment is critical to maximise optical efficiency. Specifically, the M1–M2 distance must be correct to within $\lambda/2$. A redesign of the secondary mirror support structure was the response to these analyses.




Figure 2.17: Results of the numerical analyses of the alignment tolerances of the optical system for all three colours. The quantity plotted is the Strehl Ratio, the ratio of the peak intensity of the beam to the peak intensity of the nominal beam. *Optics Box* refers to movement of the entire cold optics assembly (M3, M4, M5 filters, and detectors). *X* and *Y* form an arbitrary orthogonal basis for the plane perpendicular to the boresight of the Cassegrain telescope.

Even if it were possible to perfectly focus the telescope on the ground before flight,³ both uniform and differential thermal fluctuations of the Cassegrain telescope would have caused relative displacements greater than $\lambda/2$ in flight (see Figure 2.18). The OCTE struts described above eliminate potential displacement due to thermal elongation and contraction of the secondary struts themselves. However, thermal variation of the

³In fact, this is a difficult task, as explained in Section 3.4.1.




Figure 2.18: Estimate of the thermal displacement which would be experienced by the *BLAST06* Cassegrain telescope were the temperatures of the primary and secondary mirrors to follow those measured for the corresponding elements of *BLAST05* (lower panel). The zero point is arbitrary. Thermal gains used are those of Equation 2.17. The trend roughly corresponds to the diurnal trend of T_{M1} (top left panel). Since, for *BLAST05*, the secondary support structure was attached through the hole in M1, no thermometry could be attached to M2. Instead the temperature of the star camera baffle (top right) is used as a proxy for T_{M2} . $T_{SC\ Baffle}$ shows more short timescale variation, due to changes in attitude, than T_{M1} , but these do not have a large effect on the calculated displacement. The excursions at $+15^{\text{h}}$ and $+97^{\text{h}}$ are during scans of Saturn when the telescope was pointed close to the sun, allowing sun on the star camera baffle itself. Their effect can be seen, minutely, in the calculated displacement.

mirrors also changes their radii of curvature, resulting in a focal displacement of the M1–M2 system. Changing the radius of curvature of the primary also has the effect of bowing the secondary support structure, due to the attachment of the support struts at the edges of M1 and M2. This, too, produces a displacement.

These effects can be corrected by displacing the secondary mirror along the boresight (the Z -axis) relative to the primary based on the temperatures of the two mirrors. Assuming the telescope has been found to be in focus with the secondary mirror at some reference position, Z_0 , relative to the primary mirror, when the mirrors have measured

	<i>COI mirror</i>		<i>Bosma mirror</i>	
	M1	M2	M1	M2
Nominal Shape	Sphere	Oblate Ellipsoid	Paraboloid	Hyperboloid
Conic Constant	0.000	8.060	-1.029	-2.853
Diameter	2.000 m	0.478 m	1.816 m	0.399 m
Radius of Curvature	4.200 m	1.585 m	4.186 m	1.154 m

Table 2.3: Physical properties of the Cassegrain telescope elements, M1 and M2, for the two assemblies.

temperatures, $T_{M1,0}$ and $T_{M2,0}$, the in-focus position relative to the primary at any other temperatures, (T_{M1}, T_{M2}) is given by:

$$Z = Z_0 + \alpha_{Al} g_{M1} R_{M1,0} (T_{M1} - T_{M1,0}) + \alpha_{Al} g_{M2} R_{M2,0} (T_{M2} - T_{M2,0}) \quad (2.16)$$

where $R_{M1,0}$ and $R_{M2,0}$ are the fiducial radii of curvature of the mirrors⁴ (see Table 2.3), α_{Al} is the linear co-efficient of thermal expansion of aluminum ($= 2.35 \times 10^{-5} \text{K}^{-1}$ [34]), and g_{M1} and g_{M2} are composite gain co-efficients characterising the magnitudes of the above deformations. These gain co-efficients were found, by further numerical simulation of the optical system, to be:

$$\begin{aligned} \alpha_{Al} g_{M1} R_{M1,0} &= -45.97 \frac{\mu\text{m}}{\text{K}} - 9.9 \frac{\mu\text{m}}{\text{K}} = -55.87 \frac{\mu\text{m}}{\text{K}} \\ \alpha_{Al} g_{M2} R_{M2,0} &= 7.36 \frac{\mu\text{m}}{\text{K}} - 2.2 \frac{\mu\text{m}}{\text{K}} = 5.12 \frac{\mu\text{m}}{\text{K}} \end{aligned} \quad (2.17)$$

where the first term is due to the change in focal distance due to the change in radius of curvature and the second term is due to the displacement of the secondary mirror due to bowing of the secondary support structure. (With this sign convention, larger values are farther from M1.)

As is apparent, the focus of the telescope is much more sensitive to thermal variation

⁴ A careful treatment would have $R_{M1}(T_{M1})$ and $R_{M2}(T_{M2})$ here instead, but the constant approximation used, which discards terms $\mathcal{O}(\alpha_{Al}^2)$, is acceptable over the observed temperature range at our required level of tolerance ($\lambda/2$). Further to this approximation, the $R_{M1,0}$ and $R_{M2,0}$ used are not the radii of curvature at $T_{M1,0}$ and $T_{M2,0}$, but rather the reference values measured at the time of characterisation of the surfaces of the mirrors.

of the primary than of the secondary. Assuming a constant M2 temperature, a thermal variation of the primary of $2\frac{1}{4}$ centigrade degrees will produce a displacement greater than the $\lambda/2$ tolerance at $250\ \mu\text{m}$. To further mitigate defocusing of the telescope, a thermal shield of several layers of aluminised Mylar[®] superinsulation was installed behind the primary mirror for *BLAST06* to better thermally isolate it from the rest of the gondola. The measured temperature fluctuations of the Cassegrain telescope during *BLAST06* is presented in Figure 3.10.

2.10.3 The Active Focus Mechanism

In order to be able to correct for thermal defocusing of the telescope while in flight, an active focus mechanism (AFM, Figure 2.19), capable of moving the secondary mirror, was installed on the telescope for *BLAST06*. With this system, instead of mounting M2 directly on the secondary support structure, it was instead mounted on a leaf-sprung plate attached to the support structure. The leaf springs were tensioned under normal operation to provide resistance for the AFM, and the plate had about five centimetres of total throw.

2.10.3.1 Instrumentation

Three linear actuators provided three degrees of freedom: translational motion in the Z -direction, as well as rotational motion around an arbitrary axis perpendicular to the Z -axis. Each actuator was equipped with an encoder with a resolution of 24 000 counts per inch ($1.0583\ \mu\text{m}/\text{count}$). These *pseudomicrons* ($\psi\mu\text{m}$) were the basic operational unit of the system.

Each linear actuator is driven by a stepper motor controlled by a EZHR23 stepper controller produced by AllMotion Incorporated. Communication with these controllers occurs over a RS-485 serial bus, using a custom ASCII-based communication protocol. Each controller can be assigned a unique address, which allows daisy chaining all




Figure 2.19: The Active Focus Mechanism ready for installation on the *BLAST06* telescope. The mechanism is mounted behind the secondary mirror, which is face down here, at bottom. The three leaf springs are clearly visible at the centre of the mechanism, each with an actuator and stepper motor stack on its radius. The other three long thin devices, positioned 60° from the actuators and uncabled, are the LVDT sensors. The box labelled “SMAK” houses the three stepper controllers. Cabling for the mechanism are run down the outside of carbon fibre struts and around M1. The mechanism installed can be seen in Figure 2.16.2.

three controllers, as well as the similar controller used for the inner frame lock (see Section 2.12.3.2), on a single RS-485 bus.

The actuator encoders effectively provide only a relative displacement measurement, since their count is reset on power cycling the stepper controllers. In order to provide an absolute position measurement, three linear variable differential transformer (LVDT) sensors were also incorporated into the AFM. These had claimed repeatabilities of 0.000025 in ($0.6 \psi\mu\text{m} = 0.635 \mu\text{m}$), but integration into BLAST resulted in higher noise than expected, as high as $\sim 5 \psi\mu\text{m}$ RMS (Figure 2.20). As a result, the LVDTs were not used in the active feedback loop of the control mechanism, but rather as an operator check to allow manual detection and correction of faults in the system.




Figure 2.20: Measured noise of the three LVDTs from an ~ 80 hour portion of the *BLAST06* flight when the secondary mirror was stationary. Although the actuator encoders noticed no movement at the limit of their resolution ($1 \psi\mu\text{m}$) over this period, significant correlated noise ($\sim 5 \psi\mu\text{m}$ RMS) is present in the top two LVDTs. The third LVDT (bottom) had consistently better noise performance ($\sim 1 \psi\mu\text{m}$ RMS).

2.10.3.2 Operation

When active, the system servoed the secondary mirror according to Equation 2.16. A redundant pair of AD592 thermistors attached to the back of the primary mirror, and a similar redundant pair on the secondary provided the necessary temperature readings for the system, while a focus measurement before flight (see Section 3.4.1) provided a first estimate of the reference values. Once at float, further in-flight focus measurements corrected the pre-flight derived numbers (see Equation 2.17). No attempt was made to adjust the tilt of the secondary mirror with the AFM in-flight. With proper tilt measured and set pre-flight, thermal fluctuation in the telescope should not produce a rotational misalignment, assuming the primary and secondary have thermalised.

In order to reduce any spurious signal which a movement of M2 might produce in

the detectors, the system was designed to only move the secondary mirror if the distance between the in-focus position of Equation 2.16 differed from the current position of M2 by more than $100 \psi \mu\text{m}$, a number chosen to be less than the required $\lambda/2$ tolerance for the alignment. Based on data from the *BLAST05* flight (Figure 2.18), it was estimated that this would produce a step at most once an hour, with no perceptible loss in efficiency due to the minute out-of-focus positioning between focus corrections.

2.10.4 The Re-imaging Optics

The re-imaging optics are shown in Figure 2.15.2. These optical elements, along with the detectors, are contained in a light-tight optics box (Figure 2.21). For both Cassegrain assemblies, the Cassegrain focus is formed approximately 0.2 m behind M1, inside the optics box. Beyond the Cassegrain focus, the first of the off-axis re-imaging mirrors, M3, forms an image of the primary on M4. This mirror acts as a Lyot-stop, which defines the illumination of the primary. A second off-axis re-imaging mirror, M5 then refocuses the light on the focal plane, at the detectors.

A hole in the centre of the Lyot-stop mirrors the hole in the centre of the primary. In this hole is a calibrator source (the *cal-lamp*) which is flashed periodically during the flight, producing a recognisable signal in the bolometers. This signal is then used to characterise bolometer responsivity variations over the course of the flight [9] (see Section 4.4).

A different set of re-imaging optics were needed for each Cassegrain assembly (see Table 2.4). The COI mirror re-imaging optics further corrected for the spherical aberration of the primary. The parabolic Bosma assembly, however, required no such correction and the re-imaging optics merely defined the illumination of the primary and formed the second focal plane.




Figure 2.21: The BLAST optics box sitting on the cold plate. Two faces of the box have been removed to allow access. Light from the Cassegrain telescope enters through the window at bottom left. M3 is at bottom right, M5 at top right. The cylindrical M4 is at centre, attached to a copper heat strap. The ring at top centre holds one of the dichroic beam splitting filters (not installed). The second filter ring is also missing, but would be attached just to the left of the one present. From left are the $500\ \mu\text{m}$, $350\ \mu\text{m}$, and $250\ \mu\text{m}$ bolometric arrays. Attached to the cold plate in front of the optics box, bottom centre, is the ^3He fridge.

<i>COI mirror</i>	M3	M4	M5
Nominal Shape	Asphere	Asphere	Asphere
Aperture	$85 \times 75\ \text{mm}$	$\text{Ø}43.2\ \text{mm}$	$65 \times 61\ \text{mm}$
Radius of Curvature	$349.7\ \text{mm}$	$164.0\ \text{mm}$	$391.7\ \text{mm}$
<i>Bosma mirror</i>	M3	M4	M5
Nominal Shape	Sphere	Sphere	Sphere
Conic Constant	0.000	0.000	0.000
Aperture	$95 \times 75\ \text{mm}$	$\text{Ø}38.6\ \text{mm}$	$95 \times 75\ \text{mm}$
Radius of Curvature	$348.6\ \text{mm}$	$174.3\ \text{mm}$	$348.6\ \text{mm}$

Table 2.4: Physical properties of the cold re-imaging optical elements, M3, M4 and M5, for the two assemblies.

2.10.5 Filters and Bandpasses

Two metal-mesh dichroic filters [35] are used split incident radiation into the three BLAST bands. The first of these reflects wavelengths shorter than $300\ \mu\text{m}$, and transmits longer wavelengths. The second dichroic reflects wavelengths shorter than $400\ \mu\text{m}$, and transmits longer wavelengths. For all three bands, the long wavelength cutoff is defined by the 2λ -long waveguide at the end of the bolometer feed horns. The short wavelength




Figure 2.22: Measured spectral response of the three BLAST bands for *BLAST06*. Each band has been normalised to a peak responsivity of one. Absorption lines from residual water in the FTS are present in all bands.

cutoff of the three bands is defined by filters placed directly on top of the feed-horn blocks [28, 33]. The bandpasses for *BLAST06*, measured with a Fourier Transform Spectrometer (FTS), are presented in Figure 2.22.

2.11 Pointing Sensors

BLAST's primary pointing sensors are a pair of daytime CCD star cameras which provide absolute pointing information and two sets of three orthogonal fibre-optic gyroscopes providing rate information. In addition to these, BLAST has a collection of coarse absolute pointing sensors to augment the star cameras during flight. Data from these sensors are combined into the in-flight pointing solution as described in Section 2.3.3.

<i>Camera</i>	<i>ISC</i>	<i>OSC</i>
Model	QImaging PMI 1401	QImaging Retiga EXi
CCD	Kodak KAF-1401e	Sony ICX285AL
Chip Size (pixels)	1312×1024	1360×1036
Pixel Size	6.8 μm ×6.8 μm	6.45 μm ×6.45 μm
Digital Output	14 bit	12 bit
Well Depth	45 000 e ⁻	18 000 e ⁻
Readout Noise	20 e ⁻	8 e ⁻
Plate Scale	7"/px	7"/px
Field of View	2°5×2°	2°5×2°

Table 2.5: Specifications for the two CCD as used in *BLAST05* and *BLAST06*. Both cameras were mated to the same 200 mm $f/2$ Nikkor lens and both cameras were fitted with a lowpass optical filter (Nikon R60) with a cutoff at 600 nm.

2.11.1 The Star Cameras

The two cameras are commonly referred to as the ISC and the OSC. The ISC was built and flown on the *BLAST03* test flight. Its design was later copied to build the OSC and both cameras were flown for *BLAST05* and *BLAST06*. The only substantive difference between the two BLAST star cameras is their CCD camera. Table 2.5 summarises the characteristics of the two cameras. Each star camera is housed in a pressure vessel which also contains a star camera computer (SCC) built from a PC/104 stack. The star cameras are presented in Figure 2.23. Full details of star camera construction are given in [30].

2.11.1.1 Star Camera Operation

The PC/104-based star camera computer (SCC) is able to run the star camera autonomously. The SCCs communicate with the MCCs via Ethernet routed through an Ethernet switch in the MCC pressure vessel. Software running on the SCC acquires images from the camera, determines star positions from the image, and computes the position and rotation of the image on the celestial sphere. The position, in equatorial co-ordinates, and rotation are then transmitted to the flight computer for input into the in-flight pointing solution and data archival. Data, including a subset of captured images, is also logged locally by the SCC for use in post-flight reconstruction of the pointing-




Figure 2.23.1: A Star Camera Disassembled




Figure 2.23.2: The Star Cameras Mounted

Figure 2.23: The BLAST star cameras. Figure 2.23.1 shows one of the cameras disassembled on the bench. The camera and lens assembly is at centre. One of the belt-drive stepper motors which control the focus and aperture of the lens is visible on the right. At right, mounted on the back flange of the star camera pressure vessel is the PC/104 stack star camera computer (SCC). The star camera pressure vessel cylinder is visible on the left behind the monitor. Figure 2.23.2 shows both star cameras mounted on the gondola. The star cameras are mounted by a clamp which encircles the pressure vessel near the front flange. The star cameras have both been fitted with metre-long baffles made from G-10 which block stray atmospheric light.

solution. In addition to the calculation of the pointing solution, the SCC is also able to change the focus and aperture of the camera lens, and operates heaters within the star camera pressure vessel.

The BLAST star cameras can operate in two trigger modes: internal, and external. In internal mode, the star camera autotrigger, independently from the flight computers. In external mode, the camera is triggered by a digital output line from the ACS.

While running on internal trigger, the star cameras periodically check for an external trigger and, if detected, will switch to external triggering. Similarly, while running on external trigger, if the star camera detects a loss of the external trigger, it will switch back to internal mode. The default mode of the star cameras is internal, but normal operation results in them switching to run on external trigger.

2.11.1.2 Communication

Communication between the star cameras and the MCCs is complicated by the need to provide precise timing of camera triggers for accurate post-flight pointing-solution

reconstruction. The star camera trigger is generated by the Attitude Control System (ACS, Section 2.12) which is able to sample the trigger at 100 Hz. All other data to and from the star camera is sent via Ethernet. Ethernet communication with the star camera happens in a dedicated subthread asynchronously from the main thread. Accordingly, a set of semaphores are passed back and forth between this thread and the main thread, which handles the trigger requests to the ACS, in order to provide proper synchronisation and error recovery.

The star camera communication threads, one for each star camera, run in an open loop which, in the absence of signals from the main thread, simply acquires data from the star camera and sends it elsewhere within `mcp` for storage and inclusion in the pointing solution. The main thread contains a trigger routine which is called at 100 Hz during the Control Update task (Section 2.3.2.4). This routine is responsible for the setting timing associated with managing the star cameras.

Three types of packets are passed between `mcp` and the SCCs:

- **The state packet:** sent by `mcp` to the SCCs. It contains the entire state of the running star camera. The star camera's state is stored by `mcp` since it is `mcp` which receives and interprets commands sent from the ground, which includes commands to modify the star camera's state. The state packet also includes the current in-flight pointing solution. For *BLAST05*, this pointing solution was used by the star cameras to limit the size of star catalogue it had to query. Improvements made for *BLAST06* meant this was no longer necessary. The in-flight pointing solution is also used by the star camera to reject obviously bad solutions.
- **The data packet:** sent by the SCCs to `mcp`. It contains the pointing solution determined by the star camera from the most recent exposure. In addition to the pointing solution itself, this packet also contains measured blob positions and fluxes, to permit post-flight recalculation of the pointing solutions.

- **The handshake packet:** also sent by the SCCs to `mcp` to acknowledge the receipt of a state packet. It is identical in form to the data packet, but contains no data. Instead only an acknowledgement flag is set which, correspondingly, is not set in the data packet.

2.11.1.3 The Trigger

The star camera trigger is generated by the ACS to provide high precision timing. To generate a trigger, a trigger request is sent from `mcp` to the ACS. This trigger request consists of a 14-bit exposure length (in 100 Hz frames) and a two-bit request serial number, needed for the ACS to distinguish one request from the next.

The trigger itself is generated in the 100 Hz interrupt handler (Section 2.8) to provide 10 ms accuracy in trigger timing. While the trigger line is high, a flag is also written back to `mcp` over the BLASTbus to precisely record the trigger.

2.11.1.4 The Star Camera Cycle

During normal operation, the star camera cycle runs as follows:

1. The main thread starts the cycle by raising the Write State semaphore. The main thread then enters the ACK Wait state.
2. The Write State semaphore is recognised by the star camera thread, which causes it to send a state packet to the star camera. It then lowers the Write Pointing Semaphore and enters the ACK Wait state itself.
3. The star camera receives the state packet and responds by sending back an handshake packet.
4. The handshake packet is read and recognised by the star camera thread, ending that thread's ACK Wait. It then raises the Write Trigger semaphore.

5. The main thread recognises the Write Trigger semaphore and leaves the ACK Wait state, lowering the semaphore. It then enters the Trigger Wait state. The Trigger Wait state is of fixed length (300 ms) and gives the star camera sufficient time to set up for an exposure.
6. When Trigger Wait times out, the main thread sends the trigger request to the ACS. The main thread enters the Start Wait state.
7. The ACS receives the trigger request and raises the trigger line attached to the star camera for the requested length of time.
8. The star camera takes an exposure and processes it, obtaining a pointing solution. It send the pointing solution back to `mcp` in a data packet.
9. The star camera thread in `mcp` reads and recognises the data packet and raises the Start Cycle semaphore.
10. The main thread recognises the Start Cycle semaphore, lowers it and returns to the top of the loop for a new cycle.

Significant error detection and recovery code exists within this cycle to provide a high degree of robustness to the system. In the main thread, both the ACK Wait and Start Wait will time out if the corresponding semaphore is not set in time. The Start Wait timeout is usually 9 s and the ACK wait timeout is usually 1 s. If either of these states time-out, the trigger routine will assume a packet from the star camera has been lost and proceed as if it had been received (*i.e.* by continuing with the next step). These timeouts are overly generous on purpose: precise tuning of the cycle is performed by the SCC.

The star camera thread needs no time-out of this sort since it reads and recognises every packet received. Any packet received while not in ACK Wait will cause it to raise the Start Cycle semaphore, while any packet received while in ACK Wait will cause it to leave ACK Wait.

Since two packets are received by `mcp` per cycle, it is possible for `mcp` to be half a cycle out of phase. Most commonly this is an issue at start up. This situation is detectable by noticing the Write Trigger going high while in Start Wait (step 10). To correct this situation, the main thread resynchronises by immediately ending the Start Wait and skipping steps 1–4 in the subsequent cycle.

The star camera cycle must also be able to operate in the presence of intermittent or non-existent Ethernet communication. Although no data would be received by the MCCs in this situation, since the SCCs store all data locally, it is important that the star camera keep running on external triggers so that the data can be analysed post-flight. During the cycle, the Ethernet link is flagged either good (whenever a packet is received) or bad (whenever the main thread’s ACK wait times out). If the link is flagged bad, `mcp` modifies the loop by changing the timeouts: the start wait timeout is set to ~ 1.1 s and the ACK Wait timeout is set to 100 ms. Including the trigger delay (300 ms), this results in a fixed time between triggers of approximately 1.5 s, which is a reasonable estimate of the processing time needed by the star cameras.

The converse failure mode: failure of the external trigger in the presence of a good Ethernet link is also handled. In this case, the star camera will revert to internal triggers. As a result, the receipt of data packets by `mcp` will occur asynchronously to the star camera cycle followed by `mcp`. Nevertheless all data will be read and stored appropriately.

2.11.2 The Gyroscopes

The BLAST platform has two sets of fibre-optic gyroscopes (FOGs). Each set contains three FOGs, each measuring the rotation rate of the platform around one orthogonal axis. For *BLAST03*, only one set of gyroscopes was used, three KVH E-core 2000 FOGs. Since *BLAST03* this model has been discontinued, so we were unable to buy spare gyroscopes as potential replacements leading up to the *BLAST05* campaign. As a result, we had to buy a complete set of a new model, KVH DSP 3000, and installed this new set. The old




Figure 2.24: Typical set-up for a fibre-optic gyroscope (FOG). The source laser beam is split in two and fed into a fibre coil in opposite directions. The phase shift induced between the two beams is then measured with an interferometer.

E-core gyroscopes remained on the gondola as a back-up set, although loss of one of the E-core gyroscopes, with no spares, would have rendered this set inoperative. Each set of gyroscopes is mounted in its own box and affixed to the inner frame.

2.11.2.1 Gyroscope Operation

The FOGs operate according to the Sagnac effect. Figure 2.24 shows a typical FOG setup. The FOG detects rotation by measuring the path length difference between light travelling through a fibre-optic coil in the direction of the rotation and counter to the rotation direction. The angular random walk noise for the E-core 2000 gyroscopes is $4.8''/\sqrt{s}$, and for the DSP 3000 is $4.0''/\sqrt{s}$. The angular random walk determines the longest period over which the gyroscopes can be usefully integrated in the absence of an absolute reference. For the BLAST system this is < 2 seconds, and the nominal star camera solution rate is set with this in mind.

Because of Faraday rotation, a phase shift is induced in the FOGs due to the Earth's magnetic field. This results in a spurious signal comparable to the rotation rate of the Earth ($15''/s$) and modulated with azimuth on short timescales. To reduce this spurious signal, the gyroscope boxes are wrapped with a $100 \mu\text{m}$ layer of a high permeability metal

(μ -metal—a nickel-iron alloy). This μ -metal covering reduces the signal induced by the Earth’s magnetic field by 90%.

2.11.2.2 Thermal Drift

Both sets of gyroscopes suffer from offset variations due to temperature gradients. As a result, each gyroscope box contains a heater regulated with a proportional-integral-derivative (PID) control algorithm which can achieve ~ 10 mK stability.

Because the gyroscopes are sensitive to thermal gradients and not temperature directly, the setpoint of the algorithm is not fixed. Rather, it is modulated on long timescales (~ 5 minutes) in one degree increments with the goal of keeping the duty cycle of the heater between 5% and 40%. This is done to reduce the power consumption of the heaters while keeping the temperature of the gyroscopes stable.

The gyroscope boxes are well isolated from the external environment, being mounted behind the primary mirror. As a result, they are very thermally stable. During *BLAST06*, excluding exceptional events, the heater setpoints for the gyroscope boxes stepped only once over the duration of the flight. Gyro offset drift is not well correlated with temperature, and the offsets cannot be removed by fitting gyroscope temperature.

2.11.2.3 Readout

The E-core 2000 gyroscopes provide a single analogue output which is read in by an ADC card in the ACS. The DSP 3000 model, however, has digital output. Integration of the DSP 3000 required special modifications to one of the ACS ADC cards to allow these gyroscopes to be read in over the digital inputs. Regardless of method of input, all gyroscope signals are filtered with the same four stage digital filter (Section 2.8.1). The filter induces a 31 ms delay in the processed signals. For the in-flight pointing solution, a secondary signal is extracted from this filter after passing through only its first stage. This provides the in-flight pointing solution with a gyro signal which has

negligible (~ 2 ms) filter delay.

2.11.3 Coarse Sensors

In addition to the star camera, the BLAST platform contains a number of coarse sensors which provide high-rate, low-accuracy data. A differential GPS, a magnetometer, and a sun sensor provide coarse azimuth data; an elevation encoder and elevation inclinometer provide coarse elevation data. These data are of use primarily for the in-flight pointing solution: the post-flight pointing solution reconstruction doesn't use them.

2.11.3.1 The Differential GPS

BLAST flew a differential Global Positioning System (dGPS). In addition to providing positional and time data, which would be output from a conventional GPS, the dGPS also provides orientation data. BLAST's differential GPS receiver is a Thales Navigation⁵ ADU5. The dGPS receiver is composed of four GPS antennas, mounted on the top of the sun shields to give them an unobstructed view above, and the dGPS receiver unit. The receiver unit takes input from the four antennas and produces a position, orientation, and time solution which is read by `mcp`. BLAST uses the dGPS for coarse azimuth information. The accuracy of the dGPS is approximately $10'$, which, to a large degree, is set by the 2 m separation of the mounted antennas.

The location of the GPS antennas must be surveyed by the dGPS before it can be used. This is done by a proprietary program provided by Thales Navigation. Output from this program is incorporated into `mcp` before flight to properly initialise the dGPS unit.

A secondary output from the dGPS is used to provide accurate time information to Network Time Protocol (NTP) daemons running on the MCCs. The NTP daemon reads the time data from the GPS and updates the computer's local clock accordingly.

⁵now Magellan GPS




Figure 2.25.1: The CCD Sun Sensor




Figure 2.25.2: The Photodiode Sun Sensor

Figure 2.25: The two sun sensors flown on BLAST shown mounted in the same place on the outer frame. The CCD Sun Sensor (CSS), Figure 2.25.1, was flown on *BLAST03* and *BLAST05*. In addition to the two linear slit cameras shown, the CSS also has a separate computer unit, not shown. The Photodiode Sun Sensor (PSS), Figure 2.25.2, was flown on *BLAST06*. The PSS consists of twelve photodiodes mounted in a ring. Significant portions of the PSS are in shadow and only a portion of the sensor was used. A foot long ruler is included in Figure 2.25.2 for scale.

As a result, the MCCs' internal clocks were used as the primary clock for the BLAST platform. The MCCs also broadcast their time over NTP on the local Ethernet, and the SCC clocks are kept synchronous by receiving this signal with their own NTP daemons.

2.11.3.2 The Sun Sensor

The BLAST sun sensor determines coarse azimuth by observing the position of the sun. Two different sun sensors were flown on successive BLAST flights, both of which are described in detail in [32] and presented as Figure 2.25. The sun sensor is mounted on the back of the outer frame.

The first sun sensor, the CCD Sun Sensor (CSS, Figure 2.25.1), was used in *BLAST03* and *BLAST05*. It has a pair of 2048-pixel linear CCD cameras. The two linear CCDs are housed behind thin slit apertures, one oriented horizontally, the other vertically. This provides both azimuth and pitch information. The CSS includes a separate computer unit which is responsible for reading the cameras and computing the azimuth and ele-

vation of the outer frame relative to the position of the Sun. The CSS consumes about 50 W of power and the cameras, in constant sunlight, requires active fluid cooling (see Section 2.12.3.3) when in operation.

The second sun sensor, the Photodiode Sun Sensor (PSS, Figure 2.25.2), has twelve photodiode sensors mounted radially to cover the full 360° of azimuth, although shadowing by structural elements when mounted reduces the usable portion of the sensor to about half of the full circle. The PSS only measures azimuth relative to the sun. It consumes only 5 W of power and requires no fluid cooling, which was the main reason behind choosing to use this sensor for *BLAST06* instead of the older CSS.

From the point of view of `mcp`, the two sensors are similar. They both have a long-term accuracy of $\sim 5^\circ$ and a short-term precision of $\sim 5'$. Like the star cameras, `mcp` communicates with the sun sensors over Ethernet. A trivial protocol is used for both sun sensors: once `mcp` connects to the sun sensor, it simply sends pointing data as fast as possible, which is read by `mcp`. The typical data rate is 5 Hz for the CSS and 10 Hz for the PSS.

2.11.3.3 The Magnetometer

A flux gate magnetometer is also used to provide coarse azimuth pointing by measuring the magnetic field of the Earth. Two magnetometers have been flown. The first was a two-axis model: Schonstedt Instruments SAM-72C. This magnetometer was damaged upon landing after *BLAST05*, and a new magnetometer was fitted for *BLAST06*. The new magnetometer was a three-axis model, Applied Physics Systems Model 534. To reduce of magnetic field induced by the pointing motors, the magnetometer is placed on a long boom cantilevered off the front of the outer frame (see Figure 2.1).

The Earth's magnetic field is modelled with the World Magnetic Model (WMM) provided by the United States' National Oceanic and Atmospheric Administration (NOAA). The model is updated every five years and has an accuracy of 10%. Additionally, in the

polar regions where BLAST flies, field lines of the Earth's magnetic field are highly inclined, reducing the accuracy of the sensor. As a result, although the magnetometer has high precision, the overall accuracy of the sensor is approximately 5° .

2.11.3.4 The Elevation Encoder

The primary coarse elevation sensor is a 16-bit absolute encoder, mounted on the elevation motor, which measures the rotation between the inner frame and the outer frame. Although this provides high-rate data with $20''$ precision, pendulations and uncertainty in the pitch of the outer frame degrade from its usability as an absolute elevation sensor, resulting in an overall accuracy of $\sim 30'$.

2.11.3.5 The Elevation Inclinometer

As a back-up for the elevation encoder, an liquid-bubble inclinometer is also mounted on the inner frame. This sensor is not degraded due to uncertainties in knowing the outer frame's pitch; however, this sensor measures acceleration, so accelerations due to movement of the BLAST platform are indistinguishable from gravitational acceleration. The inclinometer is also subject to thermal drifts. As a result, the sensor is accurate to $\sim 1^\circ$, although its short-term precision is near $30''$.

2.12 The Attitude Control System

The Attitude Control System (ACS) is built of four ADC cards identical to those found in the Data Acquisition System (DAS, Section 2.7). The ACS provides analogue and digital sampling, and control for systems which are neither handled by the DAS nor communicate with the MCCs via serial or Ethernet. Primary systems controlled by the ACS include motion control, the power systems, auxiliary motors, the star camera trigger (see Section 2.11.1.3), and readout of various housekeeping sensors.

2.12.1 Motion Control

The orientation of the BLAST platform is controlled primarily by three direct drive torque motors: an elevation motor, which rotates the inner frame relative to the outer frame; a reaction flywheel motor, which rotates the azimuthal flywheel, and a pivot motor, which rotates the two halves of the pivot. These motors are described in detail in [28]. In addition to these, a fourth motor drives a roll-damping flywheel designed to damp out long-period pendulations of the platform.

All four motors are controlled using pulse-width-modulated (PWM) outputs located on one of the ADC cards in the ACS (see Figure 2.11). The PWM level is set using a proportional-integral (PI) control system implemented in the ACS.

In-flight pointing ability must be better than the array size ($7' \times 14'$) to ensure desired targets are mapped by the experiment. Pointing accuracy $< 10''$ is desired. In order to achieve this, `mcp` handles only the highest levels of motion control. Low-level control is accomplished by the DSP in the ACS.

2.12.1.1 Pointing Modes

The role of `mcp` in motion control is to provide a velocity request to the ACS. This velocity request is composed of two separate parts: an azimuthal velocity and an elevation velocity. The requested velocity is determined by the current pointing mode, of which there are ten. The pointing modes are summarised in Table 2.6. Half the pointing modes are *scanning modes* and are discussed further in Section 3.1.2. The remainder are *auxiliary modes* used for testing and other special circumstances. Each pointing mode specifies independently an elevation mode and an azimuth mode, which are typically either *goto* or *drift*.⁶

In drift modes, the velocity request for the particular direction is set directly by the

⁶A third type of mode, *lock*, exists for elevation only, and is discussed in Section 2.12.3.2.

<i>Mode</i>	<i>El. Mode</i>	<i>Az. Mode</i>	<i>Description</i>
<i>Scanning Modes</i>			
V-Cap	Drift	Drift	Map a circular region by rastering in azimuth while drifting in elevation.
S-Cap	Drift	Drift	Map a circular region by rastering in azimuth while stepping in elevation.
V-Box	Drift	Drift	Map a rectangular region by rastering in azimuth while drifting in elevation.
S-Box	Drift	Drift	Map a rectangular region by rastering in azimuth while stepping in elevation.
Quad	Drift	Drift	Map an arbitrary quadrilateral region by rastering in azimuth while stepping in elevation.
<i>Auxiliary Modes</i>			
Goto	Goto	Goto	Point to a fixed point in azimuth and elevation.
Drift	Drift	Drift	Move at constant velocity in azimuth and elevation.
Track	Goto	Goto	Track a fixed point in right ascension and declination.
Az. Scan	Goto	Drift	Raster in azimuth at a fixed elevation; a building block for the true scanning modes.
Lock	Lock	Drift	Drift at zero velocity in azimuth and move the inner frame to engage the inner-frame lock.
<i>Pseudomodes</i>			
Stop	Drift	Drift	A special case of Drift mode with both velocities set to zero.
Anti-Sun	Goto	Goto	A special case of Goto mode with the target set to the current antisolar azimuth and the current elevation of the telescope.

Table 2.6: A summary of pointing modes for BLAST. The table is divided into two classes: scanning modes, which are pointing modes used for making scientific maps (see Section 3.1.2), and auxiliary modes which are used for other purposes. Also listed are two common “pseudomodes”, which are common special cases of the other pointing modes given. Columns two and three indicate whether the given axis is in “goto” (position) mode or “drift” (velocity) mode. “Lock” is a special type of elevation mode used in conjunction with the inner frame lock (see Section 2.12.3.2).

commanded scan. A common case of this is *stop* mode, which sets the requested velocity in both azimuth and elevation to zero. In *goto* mode, the velocity request is proportional to the error between the current position and the destination position. In azimuth *goto* mode, the direction is also chosen so that the slew of the telescope does not cross the sun. In elevation, the requested velocity is limited to $0.50^\circ/\text{s}$ and the maximum acceleration is limited to $0.15^\circ/\text{s}^2$. In azimuth, the requested velocity is limited to $1.83^\circ/\text{s}$ and the maximum acceleration to $1.83^\circ/\text{s}^2$.

Because the ACS is unaware of the gyro offset (Section 2.3.3.3), the offset is added to the velocity request before transmission. The velocity request is received by a DSP on the ACS where it is compared to the current velocity as measured by the gyroscopes, and processed as described below.

2.12.1.2 Elevation Control

Elevation is controlled by the Elevation motor alone. The elevation velocity request from *mcp* is fed into a PI control algorithm with user-settable gains:

$$I_{el} = P_{el}\Delta v_{el} + I_{el} \int \Delta v_{el} \quad (2.18)$$

where I_{el} is the PWM level sent to the motor, P_{el} and I_{el} are the user-settable gains, and Δv_{el} is the error in elevation velocity. The integral is implemented as an infinite input response (IIR) filter with a characteristic time of five seconds. Since the elevation motor control depends only on the elevation gyroscope, pendulations in pitch of the outer frame are decoupled from the elevation of the inner frame by the control algorithm.

2.12.1.3 Azimuth Control

Azimuth is controlled by both the reaction flywheel and the pivot. In normal operation the flywheel spins slowly. Changes in azimuth are effected by modulation of the flywheel's rotational speed. The pivot regulates the flywheel's speed on long time scales to keep it

near a desired value (9 revolutions per minute) by dumping angular momentum to the flight train.

The reaction flywheel motor control is the same as that of the elevation motor (Equation 2.18), with the addition of a term to compensate for back-EMF in the motor:

$$\Pi_{reac} = P_{reac} \Delta v_{az} + I_{reac} \int \Delta v_{az} + C_1 (v_{reac} + C_2) \quad (2.19)$$

where C_1 and C_2 are empirically determined constants characterising the back-EMF of the motor and v_{reac} is the rotation rate of the reaction flywheel. The pivot control algorithm has only a proportional term:

$$\Pi_{piv} = P_{piv} (v_{reac} - v_{sp}) \quad (2.20)$$

where v_{sp} is the set-point for the reaction wheel rotation rate (9 revolutions per minute).

2.12.1.4 Roll Control

The roll damping flywheel is designed to damp long-period, small-amplitude oscillations in the roll of the outer frame excited by stratospheric winds. Like the pivot, the roll motor has only a proportional gain. However, unlike the pivot, the roll motor gain is automatically scaled in inverse proportion to the amplitude of the roll oscillation.

The goal of the roll damping control algorithm is to reduce the roll velocity to zero. However, since the DSP does not know the gyro offset around the roll axis, it also constrains the requested PWM level to be zero on long timescales:

$$\Pi_{roll} = \frac{P_{roll}}{\langle \alpha_{roll} \rangle} \alpha_{roll} - \langle \Pi_{roll} \rangle \quad (2.21)$$

where α_{roll} is the angular velocity measured by the gyroscopes projected along the roll direction. The characteristic time for the Π_{roll} time average is approximately 2 minutes.

<i>System</i>	<i>Power (W)</i>	<i>System</i>	<i>Power (W)</i>
ACS Side		DAS Side	
Pointing Motors	85	Data Acquisition System	113
Flight Computers	55	Preamps & Cryostat	113
Star Cameras (2)	55		
Attitude Control System	30		
Gyro Heaters (2)	30		
Auxiliary Systems	15		
Differential GPS	5		
Photodiode Sun Sensor	5		
LOS Transmitters	(150)		
ACS Total: 280 Watts		DAS Total: 225 Watts	
Total: 505 Watts			

Table 2.7: The steady-state power budget for *BLAST06*. When slewing, the pointing motors may have short-term excursions up to 560 Watts, nearly doubling BLAST’s power consumption. The LOS Transmitters are only on at the start of the flight and do not contribute to the steady-state budget. The power budget for *BLAST05* was similar, except for the increased power consumption by the CCD Sun Sensor (50 W instead of 5 W), bringing the total up to 550 Watts.

The time averaging of the roll velocity used to scale the proportional gain is asymmetric: it responds quickly (~ 0.5 s) to increases in the roll amplitude, but very slowly (~ 100 s) to decreases in it.

2.12.2 The Power System

During normal, steady-state operation, BLAST consumed 550 Watts of power during *BLAST05* and 505 Watts during *BLAST06*. An additional 150 Watts is drawn by the Line-of-Sight Transmitters when they are on during the first ~ 24 hours of the flight. The power reduction for *BLAST06* is almost exclusively due to replacement of the sun sensor (Section 2.11.3.2). The power breakdown for the BLAST platform is presented as Table 2.7.

The BLAST power system is broken into two independent halves. The larger of these, the *ACS Side*, powers the flight computers, ACS, pointing sensors, motors, and auxiliary systems. The other, the *DAS Side*, powers the DAS, preamp crate, and cryostat. Each

side consists of a solar array, power conditioning unit (PCU, also known as a *charge controller*), battery, and power distribution unit.

2.12.2.1 Solar Arrays

The BLAST solar arrays were designed and constructed by Meer Instruments of Palomar Mountain, California. There are two arrays, each with five panels. Six of the ten panels are connected to the ACS Side; the remaining four are connected to the DAS Side. Each panel consists of 96 solar cells, wired in series. Shunt diodes are installed every eight cells so that a breakdown of one cell in the series will result in the loss of only eight cells and not an entire panel. Sufficient redundancy exists in the panel so that one group of eight cells can be lost with no appreciable drop in current from the panel.

The panels are regulated using charge controllers also built by Meer Instruments and originally flown on the BOOMERanG experiment [29]. The charge controllers isolate the solar panels from each other, and regulate the current supplied by the solar arrays to the system. The current of a given panel is regulated by driving it up towards its open circuit voltage (~ 45 V for the BLAST panels at float). Each charge controller contains a resistor divider network which ensures that only one panel at a time is in this current regulation mode, the other panels being entirely on or entirely off.

The two solar arrays are mounted on either side of the BLAST gondola, attached to structural members of the sun shields and facing rearwards. The backs of the arrays are painted white to permit radiative cooling of the arrays at float.

2.12.2.2 Batteries

The BLAST power system uses Cobasys Series 9500 nickel-metal hydride (NiMH) batteries for situations when the solar panels are unable to provide sufficient power to the platform. Typically this occurs when the solar arrays are pointed significantly away from the sun, or when the gondola begins a slew and the pointing motors require extra power

to accelerate the platform. Each battery consists of ten 1.2 V cells. Two batteries are used in series on each side of the power system to provide 85 A·h at 24 V.

These batteries suffer from a positive feedback cycle during charging: as the batteries reach their lid voltage, their internal resistance increases, causing them to heat up and lowering their lid voltage. Unchecked charging of the batteries leads to thermal runaway. To prevent this, a variable resistor was installed in the charge controllers to allow regulation of the lid voltage through the ACS. The lid voltage is calculated by `mcp` as:

$$V_{lid} = \begin{cases} 31.07 \text{ V} & T_{batt} < -20.41^\circ\text{C} \\ 30.18 \text{ V} - (0.0436 \text{ V}/^\circ\text{C}) T_{batt} - e^{T_{batt}/4^\circ\text{C} - 7.25} & -20.41^\circ\text{C} \leq T_{batt} \leq 30.84^\circ\text{C} \\ 27.25 \text{ V} & T_{batt} > 30.84^\circ\text{C} \end{cases} \quad (2.22)$$

where T_{batt} is the temperature of the battery. Equation 2.22 ensures that, as battery temperature rises, battery charging becomes increasingly less aggressive. During *BLAST06*, the batteries typically ran at 15–20°C while at float.

2.12.2.3 Power Distribution

Power distribution for the two halves of the BLAST power system are accomplished by custom-built modules in the ACS and DAS crates. In addition to power break-out, these units also provide step-down DC–DC conversion between battery voltage and the voltages needed by the ADC cards and other subsystems. Finally, the power modules provide switching capabilities for the various subsystems operated by the science stack commands (Section 2.13.4.1).

2.12.3 Auxiliary Systems

2.12.3.1 The Balance System

The cryostat contains ~ 40 kg of cryogenics which slowly boil off during an 11 day flight. As a result of this boil-off, the inner frame becomes increasingly imbalanced. Further imbalance results from redistribution of cryogenics as the cryostat is tilted at different elevations. To allow finer control in elevation, and to reduce the power drawn by the elevation motor, this imbalance is counteracted by an active fluid balance system on the inner frame.

The balance system consists of two 10 ℓ tanks, placed far away from, and on opposite sides of, the axis of rotation. The tanks are filled with a high density saline solution (Dynalene HC-50) and are connected by copper piping. A reversible fluid pump transfers fluid back and forth between the tanks to actively balance the inner frame.

The inner frame balance is measured by `mcp` by monitoring the steady-state current drawn by the elevation motor. It then controls the balance pump in order to drive this current towards zero using a simple proportional gain on the current. A dead zone around the zero point of ~ 0.5 A ensures the pump is only active when sufficient imbalance has accumulated. Since the current drawn by the elevation motor during telescope slews can be significantly different than at steady-state, the balance system vetoed during these events.

2.12.3.2 The Inner Frame Lock

The inner frame lock inhibits movement of the inner frame relative to the outer frame. This is necessary for launch, when the pointing motors are not powered, to prevent the inner frame from suffering damage due to forces associated with launch and ascent. While the lock is engaged, elevation control is vetoed to prevent the elevation motor from torquing against the lock. Upon reaching float altitude, the inner frame lock is

disengaged to allow pointing of the telescope in elevation.

The inner frame lock consists of a linear actuator mounted to the bottom of one of the pyramids on the outer frame. The actuator moves a hardened steel shaft in and out of a matching hole on a plate attached to the inner frame, preventing the inner frame from moving. For convenience while working in the lab, this plate has several holes allowing the lock to be engaged at various elevations between 5° and 90° . While engaged, the lock permits travel of approximately $\pm 0.25^\circ$ in elevation. The hole at 49° is larger than the rest, and allows travel of approximately $\pm 1^\circ$.

Two inner frame locks have been flown on BLAST. The first, flown during *BLAST03* and *BLAST05*, was controlled via digital I/O lines in the ACS. This lock suffered significant damage upon launch during *BLAST05* (Section 3.3.2) and, consequently, a new inner frame lock was built for *BLAST06*. This new lock motor was controlled by a stepper motor controller on the AFM RS-485 serial bus (Section 2.10.3.1). This lock, too, suffered an in-flight failure, which was not mission critical (see Section 3.4.4).

During a flight, the inner frame lock is disengaged once at the beginning of the flight and then engaged again at the end of the flight, remaining inactive for the bulk of the flight, except during error recovery procedures (see, for instance, Section 3.4.5). When engaging the lock, `mcp` enters a special elevation pointing mode, “lock”, similar to “goto” mode (see Section 2.12.1.1) except that the elevation encoder is used to position the inner frame instead of the fully computed elevation from the in-flight pointing solution.

Although elevation control is by default vetoed while the lock is engaged, provision exists within `mcp` to override this behaviour. This allows operation of the telescope in situations where the readout of the inner frame lock’s state has failed. When launched in the larger, 49° lock position, this would also permit limited movement in elevation, were the lock to fail to disengage upon reaching float at the start of the flight.

2.12.3.3 The Active Cooling System

The Data Acquisition System (DAS) and Preamp crate each dissipate ~ 100 W of heat. Since they are located behind the primary mirror, they are unable to radiate this power to the sky. Instead, a fluid cooling system transfers heat from the electronics crates to the inner frame structure itself.

Aluminum fins, painted white to aid heat transfer, between the cards in the crates absorb heat radiated by the electronics and transfer it to the outer cases of the two crates. Dynalene HC-50 in the cooling system passes through heat exchangers on the outside of the crates and transports this heat to another heat exchanger in contact with the inner frame. A pump controlled by `mcp` runs continuously to move the fluid around the system.

Since the density of electronics in the ACS crate is significantly lower, and because the ACS can radiate to the sky directly, through gaps in the sun shields, no active cooling of the ACS is required. Fans on the flight computers' processors and hard-drives effectively transfer heat from the computers to the outside of the pressure vessel through the air. Similarly to the ACS, the Pressure Vessel is able to radiate directly to the sky and needs no active cooling⁷.

For *BLAST05*, however, the CCD Sun Sensor (CSS) required a similar active cooling system to cool its camera, which produced ~ 50 W of heat and was in direct sunlight for the duration of the flight. The reduced power consumption of the Photodiode Sun Sensor (PSS) removed the need for this this second cooling system during *BLAST06*.

2.13 Telemetry

Although BLAST is designed to operate autonomously, in normal circumstances the platform is in contact with a groundstation where we can monitor its progress and adjust

⁷This is true at float but not on the ground where the ambient temperature is higher. When closed, on the ground, the pressure vessel must be cooled to prevent overheating. As a result, the pressure vessel is usually left open in the lab until the experiment is flight-ready.

its behaviour as needed. BLAST is equipped with three downlink channels and three uplink channels for telemetry. Of these, one downlink channel, the L- or S-band biphasic, and one uplink channel, the High-Frequency (HF) channel, are available only when the gondola is within line-of-sight (LOS) of the launch site. The remaining telemetry channels are nominally available at all times. These are TDRSS and the Iridium network, both of which are bi-directional. A diagram of the available communication channels is provided as Figure 2.26.

Because CSBF handles all telemetry between BLAST and the ground, we provide science interface computers (SICs) to interface with the CSBF groundstation equipment (GSE). Two computers are needed. The primary SIC is installed in the CSBF Operations Control Center (OCC) in Palestine, Texas, and handles TDRSS and Iridium communication. The second SIC, the field interface computer (FIC) is installed at the launch site and handles line-of-sight communication.

2.13.1 Line-of-Sight Data Transmission

BLAST is equipped with three line-of-sight transmitters, supplied by CSBF, and broadcasting in either L- or S-band. Two of these transmitters broadcast NTSC compliant video signals from the star cameras which are used for diagnostic purposes at the start of the flight and are not otherwise available. The third transmitter broadcasts a 1 megabit, biphasic-modulated data channel from the flight computers. The bandwidth of this channel is sufficiently large to accommodate the entire data stream from the platform and is preferred to other downlink channels when available.

Because the bandwidth of the transmitter is greater than the data rate of the experiment, entire 100 Hz frames can be transmitted. Frames received from the BLASTbus are padded to fill out the 1 Mbit data rate, a checksum for the frame is computed, a sync word is inserted at the start of the frame, and then the frame is sent to the BLASTbus Controller (BBC), which biphasic modulates the data stream. The output from the BBC

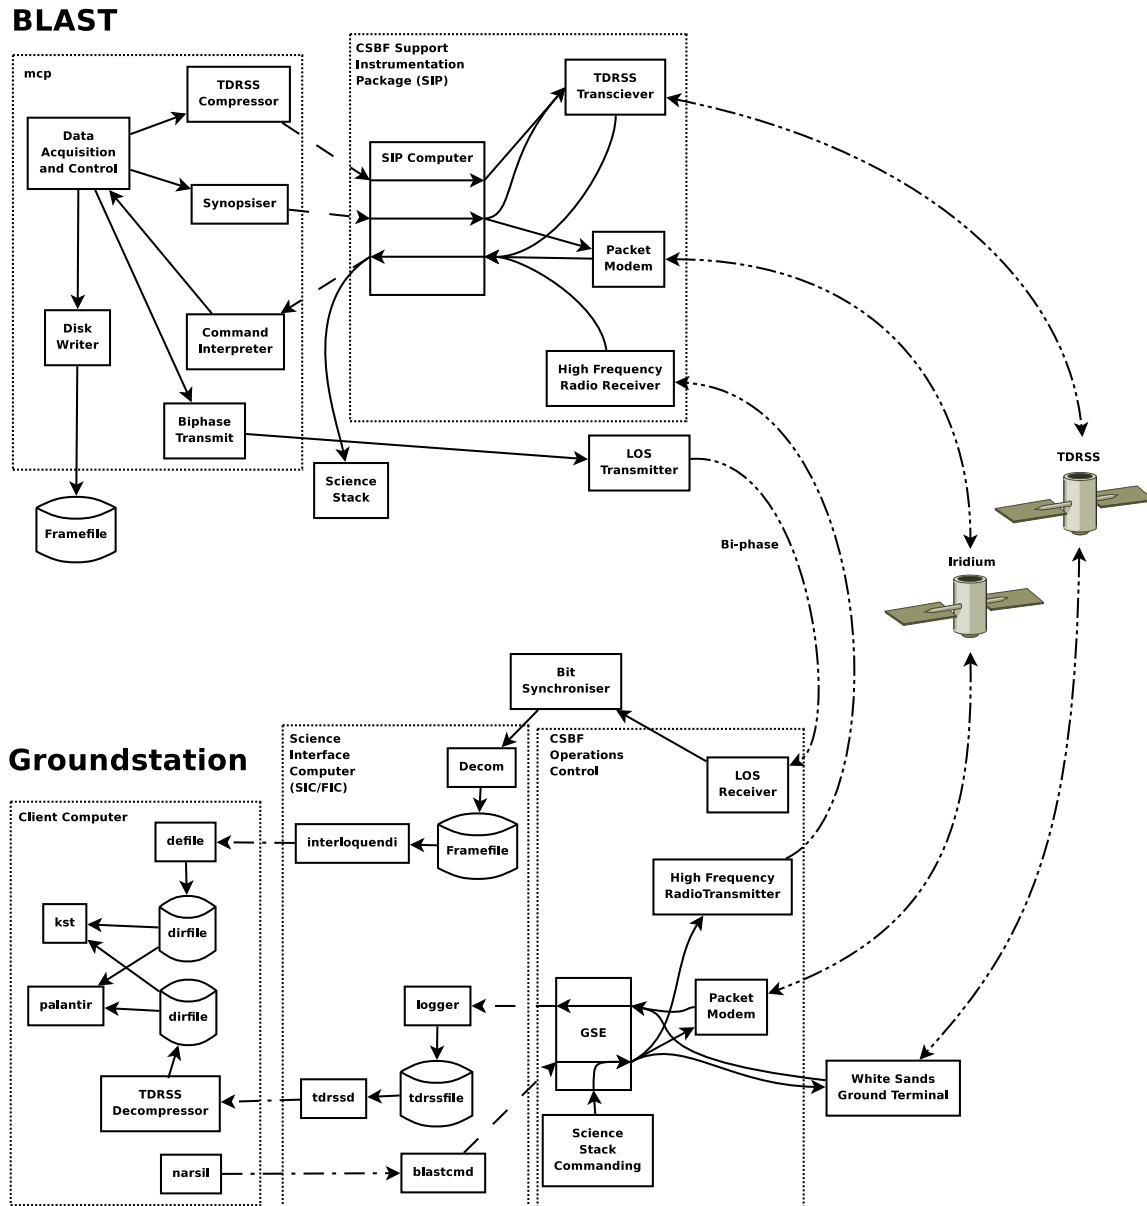


Figure 2.26: A block diagram of the BLAST telemetry paths. Dashed lines show RS-232 serial connections; dot-dashed lines are TCP/IP connections; double-dot-dashed lines indicate radio communication. `mcp` (top left) interfaces with the Support Instrumentation Package (SIP) provided by CSBF for most telemetry channels, except for the LOS Transmitters, which it drives directly. The ground link is handled entirely by CSBF, who relay data and commands between their groundstation equipment (GSE) and our Science Interface Computer (SIC), which decodes the data from `mcp`, stores it on disk and transmits it over TCP/IP to client computers where it can be viewed with quick-look analysis software, such as `palantir` and `kst`. Many clients may be simultaneously connected; only one is shown for clarity. Beyond line-of-sight, the telemetry channels which are not passed through satellites are unavailable. Commanding from client computers is done through `narsil`, which sends command requests to `blastcmd` for encoding and transmission to the GSE. Science stack commands are sent directly from CSBF equipment.




Figure 2.27: An assembled decom card. The card uses the same Altera[®]-based PCI card as the BLASTbus Controller (BBC) card (Figure 2.14). The two co-axial connectors provide input for the data and clock signals produced by the bit synchroniser.

is sent into the LOS Transmitter for broadcast to the ground.

2.13.1.1 Biphase Reception

The biphase stream is received on the ground and sent to a bit synchroniser for demodulation. The output from the bit synchroniser is a TTL bit stream and reference bit-clock. These are input into a decommutation (“decom”) card installed in the FIC. The decom card is built from the same PCI as the BLASTbus Controller card and is shown in Figure 2.27.

The decom card reconstructs data words from the bit stream by recognising the sync words at the starts of the frames and sends a word stream to a daemon running on the FIC for frame reconstruction. The sync word alternates between $1110\ 1011\ 1001\ 0000_b$ ($EB90_h$) and its inverse, $0001\ 0100\ 0110\ 1111_b$ ($146F_h$). This is done to decrease the probability of spurious frame-lock by the decom card.

Because the sync word flips every frame, the decom card may reconstruct the data stream 180° out of phase (*i.e.* with all bits inverted) with equal probability to reconstructing it with 0° phase. Furthermore, data glitches in the telemetry stream may cause this polarity to flip. The decom daemon is designed to detect and correct for polarity

inversions.

The decom daemon creates two copies of each frame it reconstructs, one with normal phase, and one with the phase inverted. It then deletes the sync word from the start of these two frames and replaces it with the original start-of-frame word used by the BLASTbus. The checksum for each frame is then computed. Since the replacement of the sync word ensures that the copies are no longer inverses of each other, only one of the two checksums will match the checksum computed by `mcp` and embedded in the data stream. Only the matching frame is kept and the polarity of the data stream is noted. If both checksums fail, it indicates corruption of the data stream during transmission, which is also noted. In this case, the frame with the same polarity as the last good frame is preferred, since polarity flips are infrequent. Frames are written to disk by the decom daemon with identical format to that used by `mcp` (into a *framefile*, Section 2.3.4).

2.13.1.2 Interloquendi

Data acquired by the decom card on the FIC is distributed to client computers by a network service called Interloquendi. This service implements a protocol similar to the file transfer protocol (FTP). Interloquendi monitors the data being written to disk by the decom daemon. Clients which connect to interloquendi are served a data stream corresponding to the current data output by the decom daemon. Interloquendi gives many computers simultaneous access to the telemetry coming from BLAST, while avoiding deficiencies that arise from exporting the data with the Network Filesystem (NFS) protocol. Although interloquendi implements a TCP/IP protocol, it is typically only used on a local network, due to security concerns. Consequently, the line-of-sight data is only available at the point of reception in the field.

2.13.1.3 Defile

The framefiles used to store data by `mcp` and the `decom` daemon is designed to be a simple, and fast format for storage. Unfortunately it is inefficient for reading. Since reads from spinning hard-drives are typically buffered in kilobyte-sized blocks, entire frames must be read, even when only one channel is of interest⁸. To overcome this deficiency, the `defile` utility converts these framefiles into a second format known as *dirfiles*. Analysis software used for BLAST is designed to read dirfiles.

A dirfile is a filesystem-based database format for time-ordered binary data. A dirfile separates each field in the framefile into a separate time-ordered data TOD stream and writes each of these to its own file, all collected in a central directory. The dirfile also contains a text file describing the format of the database.

Framefiles can be read directly from disk by `defile`, or else it can accept data from the interloquendi network service. Since it is used for in-flight data monitoring, `defile` is designed to be able to read framefiles as they are being written, providing real-time data.

2.13.2 The High-Rate Downlink

The *high-rate downlink* refers to the 6 kbit/s telemetry channel available to BLAST through NASA's Tracking and Data Relay Satellite System (TDRSS). TDRSS consists of nine geosynchronous satellites, three of which are active at any given time, offering near global coverage. The point of contact between TDRSS and the ground is the White Sands Ground Terminal (WSGT) in New Mexico, USA. The primary role of the system is to provide communication capabilities between the ground and low-orbiting satellites, including the space shuttle and International Space Station.

⁸Storing the framefiles on flash drives, where single bytes can be fetched, could overcome this problem, but current flash drive technology has insufficient capacity to fit the BLAST data set.

2.13.2.1 TDRSS Compression

The high-rate downlink is the primary channel used to communicate with BLAST when the platform is not within line-of-sight of the launch site. Because the BLAST data-rate is significantly wider than the 6 kbit/s available, data is aggressively compressed before being transmitted over the high-rate to maximise information density.

Compression occurs in two stages. The first stage, field compression, decimates the data by down-sampling the data stream. Three types of field compressions are implemented:

- **single sample**, which simply records one sample of the field per 15 second TDRSS frame;
- **differential compression**, which differentiates the data, and records the first sample, and the differential stream. Between two and twenty samples are stored per 15 second TDRSS frame; and
- **integral preserving compression**, which differentially compresses the integral of the field, and is used solely on the gyroscope TOD streams, for which the integral is more useful than the TOD stream itself.

In the last two cases, the data are lowpassed to remove frequencies above the resampling Nyquist frequency.

The second stage of compression, sample compression, packs the data streams resulting from field compression. Each field is assigned a certain number of bits in the in the TDRSS frame. A nominal range for the field is mapped into the specified number of bits. Values which lie outside the nominal range are allowed to overflow into additional bits. Were all values to fit within the nominal range, the bits reserved for the field would be underutilised; conversely, were all samples to exceed the nominal range, more bits than necessary would be used. As a result, the compression system automatically adjusts the nominal range to keep the fraction of samples exceeding it between $1\frac{1}{4}\%$ and $11\frac{1}{4}\%$.

The data are packed into a 15 second long TDRSS frame. This frame may be at most 2041 8-bit bytes long. In cases where the frame exceeds this length, whole fields at the end of the frame are discarded to keep the frame length under the maximum. The frame is padded to its maximum length and sent to the SIP for transmission via TDRSS.

2.13.2.2 High Rate Reception

TDRSS data from the gondola is received by the White Sands Ground Terminal (WSGT) in New Mexico and transferred to CSBF's Operations Control Center (OCC) in Palestine, Texas, USA. BLAST's Science Interface Computer (SIC) at the OCC receives the TDRSS packets from the CSBF groundstation equipment (GSE) and stores it on disk. A network daemon on the SIC sends the TDRSS packets over the Internet to client computers. The client computers decompress the data and store them as a dirfile, which can be read by quick-look analysis software.

2.13.3 The Low-Rate Downlink

When the high-rate downlink is unavailable, BLAST relies on the low-rate downlink. The low-rate downlink consists of one 255 8-bit byte packet every 30 seconds, resulting an effective bandwidth of 68 bits/s. Whenever a low-rate packet request arrives from the SIP, `mcp` produces a synopsis of BLAST's current state by compressing approximately two hundred channels from the current frame into the 255 bytes available. Bitfield channels are simply sent down as is. Other data are rescaled to fit a nominal range within the one or two bytes available to the channel. Unlike with the high-rate compression, values of the channel lying outside the nominal range are simply truncated to the extrema of the range.

The low-rate is sent both over TDRSS and, via packet modem, over the Iridium Satellite Telephone network. Regardless of the route of the packets, the low-rate is received at the CSBF Operation's Control Center and interleaved into the TDRSS high-

rate packet stream for transmission to the Science Interface Computer (SIC). Accordingly, it is transferred over the Internet by the same network daemon as the high-rate, and clients decompress and store the low-rate in an analogous way.

Unlike the high-rate, which includes several bolometer traces for monitoring purposes, the low-rate contains only housekeeping state data and is only able to provide state-of-health indication for BLAST's many subsystems.

2.13.4 The Command Uplink

Communication from the ground to the BLAST platform consists solely of commands sent manually by operators on the ground monitoring the health of the experiment. Two types of commands are used by BLAST: science stack commands used to power subsystems on and off, and software commands used to manipulate the operation of `mcp` and the star cameras.

2.13.4.1 The Science Stack

The CSBF provided science stack provides twenty-eight open-collector outputs which provide a 100 ms momentary contact. The BLAST power distribution is divided into fourteen distinct subsections, and these science stack commands are used to switch these subsections on or off via relays in the power modules.

Except for the pointing motors BLAST launches with all subsystems powered. As a result, science stack commands are mostly used at the end of the flight, just before termination to shut the system down. However, they are also used at the start of the flight to turn on the pointing motors, and after loss of line-of-sight to turn off the LOS Transmitters. They are also included to provide emergency power cycling capability for subsystems while in flight in case of an otherwise uncorrectable hardware upset.

Science stack commands are requested through verbal communication with the CSBF operations control center (OCC) in Palestine, Texas, USA.

2.13.4.2 Software Commands

Software commands are used to control all aspects of BLAST's operation. Two hundred and thirty-three commands are recognised by `mcp`. Of these, ninety-six include one or more parameters. Commands arrive in encoded packets from the SIP and range in length from two bytes, for a command with no accompanying parameters, to twenty-six bytes for the longest command defined.

While within line-of-sight, commands are sent through a high-frequency (HF) radio link. Beyond line-of-sight TDRSS is the preferred channel for sending commands, due to its low latency. When a TDRSS uplink is unavailable, commands are sent through Iridium. Long commands may be broken up into smaller packets for transmission. Missing packets result in `mcp` discarding the whole command. Commands are also provided to `mcp`'s command interpreter by the scheduler subsystem (Section 2.3.2.5), by which BLAST's operation is automated.

2.13.4.3 Groundstation Commanding Software

All software commands are routed through CSBF's groundstation equipment (GSE). A commanding daemon, `blastcmd` runs on the Science Interface Computer (SIC), which is responsible for encoding commands for transmission to the GSE. This daemon listens for command requests on a network port, allowing commanding of BLAST from anywhere. Since multiple clients can connect to the daemon at the same time, the daemon also handles command queueing, to ensure simultaneous commands don't collide.

The `blastcmd` program also functions as a command-line client to the daemon, allowing a simple interface to commanding. Since there are a large number of commands, and many commands have multiple parameters, operator errors when using the command-line `blastcmd` client are likely. To simplify the commanding process, a GUI commanding client, `narsil`, is usually used. This program allows command selection from a list, with descriptive parameter names. It also records all commands sent to a central log file.

2.13.5 Flight Monitoring Software

While in flight, the bulk of flight status monitoring is done with `kst`, an open-source program designed to plot and analyse high quantities of real-time, time-ordered data. To augment `kst`, a graphical snapshot program was also written. This program, `palantir`, verbosely displays important quantities from the most recent frame received on the ground and flags values which are found to be outside some defined nominal range. Both of these programs are designed to read the dirfiles created by `defile` (Section 2.13.1.3).

Ἐνεκεν μὲν δὴ τοῦ μηδὲ τοῦτον παραλελειφθαι τὸν τόπον ἰκανῶς ἔχειν καὶ ταῦτα
ἡγούμεθα πρὸς ἔνδειξιν τῶν κατὰ τὴν τοιαύτην θεωρίαν ἐφοδευομένων,
—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Η΄:τ΄ [1]

Chapter 3

Flight Performance

BLAST has been flown in two long-duration balloon (LDB) flights. The first flight, *BLAST05*, was launched from Esrange, Sweden in July, 2005, and acquired ~ 100 hours of data before landing on Victoria Island, Northwest Territories. The second flight, *BLAST06*, was flown in December 2006 from Williams Field, Antarctica, and acquired ~ 250 hours of data. Flight tracks for both flights are presented in Figure 3.1. A few general points will be considered first, before discussing the flights in particular.

3.1 Scan Strategies

BLAST is a mapping experiment, and its scan strategy is designed to optimise post-flight map making. The $1/f$ noise in the bolometers make it desirable to revisit an astronomical object numerous times during the flight. Accordingly, the telescope is designed to raster on the sky, visiting each point on the sky several times with different bolometers.

3.1.1 Scan Limitations

Changes in elevation result in thermal drifts in the bolometers possibly due to sloshing of the cryogenics in the cryostat. Additionally, elevation changes affect the air mass through



Figure 3.1.1: *BLAST05*Figure 3.1.2: *BLAST06*

Figure 3.1: Flight paths for the two LDB flights of BLAST. Figure 3.1.1 shows the *BLAST05* flight path from Sweden to Canada. Figure 3.1.2 shows the *BLAST06* flight path around Antarctica. The scallop features correspond roughly with diurnal cycles. From [37].

which the telescope looks. To reduce these effects, the telescope is scanned primarily in azimuth while rastering maps. While the telescope is scanning in azimuth, slow changes in elevation are made in order to scan the desired region.

The 2 ms time constant of the bolometers (Section 2.5.1.1) places a weak constraint on azimuthal scan velocity. The effect of this time constant is the convolution of the sky signal with the bolometers' exponential response, resulting in an asymmetric smearing of the sky along the scan direction. Slower scan rates decrease this effect. In practice this isn't a serious limitation, since the bolometer time constant is deconvolved from the time-ordered-data (TOD) streams in analysis pipeline (Section 4.2.2).

A tighter constraint results from the finite sampling rate of the detector readout. The 250 μm beam has a full width at half maximum (FWHM) of $\phi = 30''$. Requiring at least n samples of the beam limits the speed of the telescope on the sky to:

$$\omega \leq \omega_{\text{samp}} = \frac{\phi f}{n} \quad (3.1)$$

where $f = 100.16$ Hz is the detector sample rate. For $n = 3$, the absolute minimum to fully sample the beam, this limits the telescope's speed to $0.28/s$.

Since the time constant of the bolometers is significantly smaller than the sampling period, this limitation could be overcome by increasing the sample rate of the detectors. However, the hard limit on BLAST's scan speed is not the telescope, but the pointing system.

In order to constrain the random walk noise of the gyroscopes in post-flight pointing reconstruction, we require the star cameras to be able to detect and centroid at least single stars while in the middle of a scan. The requirement to detect at least one star dictates the minimum possible exposure time for the star cameras. This time is somewhat field dependant since, in areas of bright stars, a shorter exposure time can be used. Typical values for the mid-scan exposure time is $t_{ms} = 30\text{--}60$ ms. Motion of the platform during an exposure results in smearing of stars along the scan direction. This, in turn, limits the accuracy of the star centroiding algorithm on the star camera. Limiting smearing of blobs to a length ζ , limits the speed of the telescope to:

$$\omega \leq \omega_{point} = \frac{\zeta}{t_{ms}}. \quad (3.2)$$

Experience with the star cameras indicates that the centroiding algorithm can reasonably handle $\zeta \approx 18''$ ($\sim 2.5\text{--}7''$ pixels) of smearing, resulting in a maximum scan rate of $\sim 0.1/s$.

3.1.2 Scanning Modes

The five scanning modes have been already tabulated in Table 2.6. Three basic scan shapes on the celestial sphere are implemented: a circle (the *Cap* modes), a rectangle aligned in azimuth and elevation (the *Box* modes), and a generalised quadrilateral defined by its four corner points (the *Quad* mode). These scanning strategies are illustrated in Figure 3.2.




Figure 3.2: Idealised representations of the three BLAST scan strategies plotted on the same region in right ascension and declination. All three scans are done primarily in azimuth, with slowly varying elevation. The plot at left shows a V-Cap or S-Cap scan; the plot at center shows a V-Box or S-Box; the plot at right shows a Quad.

Although all five scanning modes use the low-level drift mode for elevation motion control (see Section 2.12.1.1), they may also be divided according to their high-level elevation control mode. Two options are available. The first is stepping mode, used by S-Cap, S-Box, and Quad, which recalculates the elevation drift rate at each azimuthal turn-around. The second is fixed-velocity mode, used by V-Cap and V-Box, which keeps a fixed elevation drift rate throughout the scan.

As the telescope exits the defined scan area during an azimuthal scan, it enters azimuthal turn-around. During turn-around, the telescope's azimuthal acceleration is fixed at $0^{\circ}1/s^2$. As the scan enters turn-around, a special semaphore is sent to the star camera cycle (Section 2.11.1.4) which inhibits the writing of the trigger until the telescope's velocity drops below $0^{\circ}015/s$.

In normal operation, in-flight star camera solutions are only obtained at azimuthal turn-around. Here the telescope is moving slowly, and integration times on the star cameras are longer, resulting in more stars detected. Between turn-arounds, `mcp` relies mostly on the gyroscopes for pointing. Random-walk noise in the gyroscopes produces pointing jitter. The stepping scanning modes were found to be better able to compensate for




Figure 3.3: The fields of Figure 3.2 cross-linked at 13° , typical of *BLAST06*. Note the fundamental difference between Box and Quad modes. Because only the central portion of a Box mode is cross-linked, the Box mode is used primarily to scan point source targets, especially calibrators. Box mode is preferred to Cap mode for this purpose due to the smaller amount of time spent in turn-around for a given scan area.

random walk in elevation, since the scan trajectory is recomputed at every turn-around. Because of this, stepping scanning modes (*S-Cap*, *S-Box*, *Quad*) are used exclusively during the science flights for data acquisition.

3.1.3 Cross-linking

Because the bolometers have a large $1/f$ noise component, it is important to *cross-link* the maps, in order to characterise this low-frequency noise. Cross-linking involves scanning the target field at different angles. Since scanning occurs primarily in azimuth, cross-linking is accomplished through sky rotation. Figure 3.3 demonstrates this cross-linking.

For a given terrestrial latitude, λ , the maximum possible cross-linking angle, χ_{max} ,

for a field with declination δ is:

$$\chi_{max} = \begin{cases} 180^\circ - 2|\lambda| & \phi_{max} > \{\theta_+, \theta_-\} > \phi_{min} \\ (180^\circ - 2|\lambda|) \sin(\theta_+ - \phi_{min}) & \theta_+ > \phi_{min} > \theta_-, \delta\lambda < 0 \\ (180^\circ - 2|\lambda|) \sin(\phi_{max} - \theta_-) & \theta_+ > \phi_{max} > \theta_-, \delta\lambda > 0 \\ \text{field not observable} & \text{otherwise} \end{cases} \quad (3.3)$$

where $\theta_{\pm} \equiv \pm 90^\circ \mp |\lambda \mp \delta|$ are the upper and lower transit elevations of the source, $\phi_{min} = 25^\circ$ is BLAST's minimum observable elevation, and $\phi_{max} = 55^\circ$ is BLAST's maximum observable elevation. Maximum cross-linking is achieved by observing the targets on opposite sides of the transit points, as far away as possible. For *BLAST05*, the mean terrestrial latitude of the platform was $69^\circ 3$ N, making χ_{max} as high as 40° . For *BLAST06*, the mean terrestrial latitude was $83^\circ 6$ S, reducing χ_{max} to $\sim 15^\circ$.

3.2 Flight Operations

Numerous diagnostic tests, both to optimise in-flight performance, and also for post-flight characterising of the data, are performed while the telescope is in flight. The bulk of flight diagnostics involve monitoring of the housekeeping data transmitted by the platform. This section describes more in-depth tests and checks which must be performed.

3.2.1 Line-of-Sight Diagnostics

The start of the flight, while the platform is within line-of-sight (LOS) of the launch site, is the most important part of the flight, since this is the only time when the entire dataset is available to the ground crew during the flight. During this time, which typically lasts 18–24 hours, a suite of tests are performed to verify and optimise the in-flight performance of the instrument. Additionally, it is desirable to make scientifically useful observations

during this time, to hedge against not being able to recover the data storage vessel.

One early task which must be performed is the determination of the in-flight bias levels and lock-in phase offsets. This is done by varying separately the detector bias (V_0 of Equation 2.10) and the lock-in phase shift (δ_R of Equation 2.15) and maximising the corresponding bolometer responsivity to the calibrator lamp (Section 2.10.4). Due to the volume of data required to be analysed in this process, special routines were written to simplify the process. See [28] for full details.

For *BLAST06*, the focus of the Cassegrain telescope needed to be checked while within line-of-sight as well. Although focus tests had been performed on the ground before launch (Section 3.4.1), in-flight tests were needed to make sure that the focus value determined for flight was correct. These tests were performed by measuring the peak response of the bolometers when scanning a point-source for several different focus positions.

Pointing measurements are also performed early in the flight by mapping known bright point-like sources. These are used to characterise and correct misalignment in the pointing sensors.

Once the instrument has been tuned, a known point source is scanned and a naive map made to determine the telescope's point-source sensitivity. This measurement is used to determine whether to use the optimistic or pessimistic set of schedule files (Section 2.3.2.5).

3.2.2 Calibrators

In addition to the scientific fields, the BLAST flight schedule routinely scans specially selected sources used as calibrators. Two types of calibration measurements are performed: pointing calibration and flux calibration. Calibrators are observed every few hours, alternating between flux calibrators and pointing calibrators.

Pointing calibrators are bright, compact objects with a known position. Shallow maps

are made of these calibrators, which are used in post-flight analysis to help characterise performance of the pointing system (see Section 4.9).

A subset of pointing calibrators are also used as flux calibrators. Sources used as flux calibrators have well characterised spectra and are mapped in more depth than the pointing calibrators, to provide a better measurement of flux. These measurements are used to calculate the conversion factor between detector response in volts and source brightness in Janskys (see Section 4.7.2), as well as to ensure equal response across the array (see Section 4.7.1). Since they are deeper, flux calibrator measurements can also be used for pointing calibration.

3.3 *BLAST05*

BLAST's first long-duration balloon (LDB) flight was launched from Esrange, Sweden (67°53' N, 21°06' E) at 1:10 UTC, 12 June, 2005. After 96 hours at float, it landed on Victoria Island, Northwest Territories (71°07' N, 110°30' W) at 6:15 UTC, 16 June, 2005. The flight suffered from an unexplained problem with the Cassegrain telescope, and corresponding decrease in point-source sensitivity. Upon landing, the telescope fell forward, destroying the Cassegrain telescope and other forward members, but was otherwise largely undamaged.

3.3.1 Ascent

Ascent speed was 2.5 m/s through the tropopause and the platform got quite cold. As a result of this cooling, three of the ADC Cards in the Data Acquisition System (DAS) rack stopped working. This included the housekeeping card controlling the bias generator, resulting in a loss of all bolometer read-out. This failure mode had been seen before, during winter star camera tests in the Toronto high-bay, when the telescope was allowed to cool to -5°C , and should have been investigated then. After getting to float and

activating the pointing motors, the telescope was pointed towards the sun to heat up the electronics. After ~ 30 minutes of doing this, the affected cards started working again. A plot of temperatures is presented as Figure 3.8 indicating the low temperature experienced while going through the tropopause. The DAS case thermometer read around 0°C when the cards stopped working, and approximately $+20^{\circ}\text{C}$ when they started working again.

Following recovery, the temperature-sensitive cards were replaced and investigated to determine the faulty component. Failure occurred at a temperature above the minimum safe operating temperatures of the components of the board, and it's thought that the problem was the printed circuit board (PCB) itself. Although all affected cards were replaced, a single card failed during ascend of *BLAST06* for the same reason, but was not considered a significant problem as it began working again once it heated up.

3.3.2 The Cassegrain Telescope

Significant, but unknown, damage occurred to the Cassegrain telescope either before or during launch of *BLAST05*. Destruction of the Cassegrain telescope upon landing has made precise determination of the nature and cause of the damage impossible. The resultant degradation of the beams resulting from the damage was immediately apparent when making early beam maps during the line-of-sight portion of the flight. Figure 3.4 shows an early pass across Saturn exhibiting the unusual structure.

Post-flight beam maps (Figure 3.5) show a six-fold radial symmetry indicating that the likely damage had to do with the six sections of the carbon-fibre primary mirror coming apart. This may have happened when the gondola came off the pin at launch, when the inner frame experienced a jolt (Figure 3.6).

At launch, the balloon pulled BLAST off the launch pin, giving the inner frame a severe jolt. This broke the inner frame lock (Section 2.12.3.2). The inner frame bounced against a mechanical stop at 56° twice before before coming to rest two degrees lower than the original lock position, indicating deformation of the lock. Despite this, the




Figure 3.4: One of many passes made of Saturn during the early part of the flight showing the complex beam shape in one of the $250\ \mu\text{m}$ bolometers. The time stream is base-line subtracted. The signal shows up as a negative detector response since an increase in temperature of the bolometer results in a decrease in the bolometer's resistance.




Figure 3.5: Point spread functions (PSFs) for the three wavebands created in post-flight analysis from the *BLAST05* data. The small circles in the bottom right corners of the maps represent the nominal beam sizes calculated before flight. The extended beams significantly reduced the point-source sensitivity during the flight. From [9].




Figure 3.6: The event at the launch of *BLAST05* which may have damaged the primary mirror. The top panel shows the elevation encoder reading, which measures relative pitch between inner frame and outer frame. The middle panel shows the pitch gyroscope which measures the angular velocity of the inner frame. The bottom panel shows altitude, indicating the time of launch. The balloon was released at $T-1^{\text{m}}6$, at which time the launch vehicle, on which *BLAST* was hanging began to move. The elevation encoder shows the inner frame rattling in the larger, 49° inner frame lock hole, which provides movement of $\pm 1^\circ$. At $T+0$, as the balloon pulls *BLAST* off the pin, the inner frame experiences a jolt, and bounces against the 56° mechanical stop twice before coming to rest at 46° , two degrees lower than the original lock position.

inner frame remained engaged and was able to disengage upon reaching float. However, we were unable to re-engage the lock for termination, which likely contributed to the subsequent destruction of the Cassegrain telescope upon landing.

Another possibility is that the surface of the mirror became delaminated. While waiting for launch, the telescope was caught in a brief rainstorm. Water which fell on the primary mirror may have seeped under the surface of the mirror. This water would




Figure 3.7: Float altitude variations for *BLAST05*. Diurnal variations are apparent, as are altitude changes due to ballast drops (notably at T+54^h, T+66^h and T+69^h dashed lines), changing albedo between land and water, and climatic effects.

have frozen on ascent in the tropopause, potentially delaminating the surface.

Whatever its cause, the damage resulted in a significantly reduced point-source sensitivity, obligating us to discard both the optimistic as well as the pessimistic pre-loaded flight schedules (Section 2.3.2.5). Instead, new targets were selected and the telescope was operated entirely from the ground via TDRSS uplink.

3.3.3 Balloon Environment

The mean BLAST float altitude for *BLAST05* was 38.2 km. Figure 3.7 shows the altitude variation of the BLAST gondola during *BLAST05*. At this altitude, the telescope is above 99% of the atmosphere, which is 98% transparent across the BLAST bands (Figure 1.5). The near vacuum environment at float means that the BLAST electronics must cool radiatively.

Figure 3.8 shows the temperature profiles of various BLAST systems. The gondola cools after launch during ascent while it goes through the tropopause. Components exposed to direct sunlight, including the sun sensor, the solar arrays and the pivot recover fastest from the cooling of the tropopause, reaching their minimum temperatures approximately 40 minutes after launch. Shaded components reach their minimum temperatures later, one to two hours after launch.




Figure 3.8: Measured temperature profiles during the first sixty-five hours of *BLAST05*. The top panel shows the temperatures of various outer frame electronics: the sun sensor camera, the DAS battery, the exterior of the PV and the ACS case. The middle panel shows the temperatures of inner frame components: the DAS case, the exterior of the cryostat, the fluid cooling system heat exchanger, and the back of the primary mirror. The bottom panel shows the temperatures of several extreme points on the gondola: one of the solar arrays, the tip of one of the star camera baffles, the pivot, and the scoop tip. The gaps in the middle panel, and in the star camera baffle temperature in the bottom panel are due to the temporary loss of the ADC card in the DAS which reads out these thermometers (Section 3.3.1).

At float, the gondola warms up again. Once thermalised, a 10–20°C diurnal variation is apparent in most ambient thermometers. The knee in the inner frame heat exchanger temperature at T+48^h (middle panel, Figure 3.8) occurs when the inner frame cooling pump is turned off. This results in an immediate drop in temperature of the heat exchanger to the level of other inner frame passive components, and a gradual increase in temperature of the DAS. By the time the pump was turned back on at T+63^h the

DAS case temperature had risen to 34°C. The solar array temperature (bottom panel, Figure 3.8) shows a bimodal temperature, which fluctuates between -10°C to $+40^{\circ}\text{C}$, when the panel is in use in conductive mode, and $60\text{--}75^{\circ}\text{C}$ when the panel is disabled in resistive mode.

The scoop tip thermometer, (bottom panel, Figure 3.8) measures the temperature of the most forward member of the BLAST sun shields, the bottom edge of telescope cavity (see Figure 2.1). This thermometer showed the greatest diurnal variation of any ambient thermometer (-50°C to $+50^{\circ}\text{C}$). This indicates that the scoop tip saw direct sunlight in *BLAST05*, when the sun was at its highest elevation.

3.3.4 The Pointing System

Except for some isolated issues, the pointing system behaved well during the flight. The differential GPS (dGPS) unit was unable to provide reliable orientation data to the in-flight pointing solution, and was therefore vetoed. In this vetoed state, it still provided location and time data to *mcp*. It is thought that the cause of this failure was differential flexure in the outer frame, and, consequently, the sun shields on which the GPS antennas are mounted. The locations of the four antennas were characterised (see Section 2.11.3.1) while sitting on its feet on the ground. The sun shields have a different flexure in this orientation than they do while the gondola is hanging from the pivot. This effect could have been removed by characterising the GPS array while the platform was hanging from the launch vehicle, but shadowing of GPS satellites by the launch vehicle itself makes this an unviable solution. The dGPS failed again during *BLAST06*, likely for the same reason. The other, redundant coarse azimuth sensors made this failure non-critical.

The azimuthal gyro offset calculation (Section 2.3.3.3) did not behave as desired. As a result, the gyro offsets were measured by hand and set to constant values for the duration of the flight. Improvements made after the flight to the gyro offset calculation algorithm, to make it conform to the algorithm described previously, meant that this was

not a problem for *BLAST06*.

3.3.4.1 Elevation Control Loss

Elevation control was lost for 3^h9 during the flight due to an unanticipated interaction between the command scheduler (Section 2.3.2.5) and the elevation veto on the inner frame lock (Section 2.12.3.2). Although the inner frame lock disengaged upon reaching float, freeing the inner frame, the limit switch indicating full retraction of the lock pin never closed. This was likely due to damage suffered by the lock upon launch (Section 3.3.2). Because the limit switch did not close, `mcp` kept the elevation motor vetoed, thinking, erroneously, that the inner frame lock was still engaged. To overcome this, a command was sent from the ground to disregard the inner frame lock state, after which the elevation motor could be operated normally.

Secondly, because of the significantly degraded performance of the *BLAST05* telescope (Section 3.3.2), the command scheduler was never engaged upon reaching float. Instead the telescope was operated manually from the ground via the TDRSS uplink.

In order to facilitate this, the commanding system time-out (Section 2.3.2.5) was set to use the maximum possible time-out period, 4 hours. Throughout the flight, operators on the ground needed to make sure that the time between successive uplinked commands was shorter than this time-out period. In instances where no new commands needed to be sent, a no-operation command could be uplinked, which would simply result in the time-out being reset.

Due to a loss in the telemetry uplink, this time-out period was reached at T+27^h67 without an intervening command to reset it. As a result, `mcp` queried the command scheduler for the first time of the flight. This caused the command scheduler to execute the initial float commands: opening the pot valve, unlocking the inner frame and turning on the elevation motor, despite these tasks having already been completed manually.

The design of the command scheduler had taken into account the fact that the initial

float commands would most likely be performed manually. The commands used in the initial float commands are all commands which simply change `mcp`'s operating state. No harm comes from re-executing these commands if `mcp` is already operating in the nominal at-float state.

Unfortunately, due to the inner frame lock override, `mcp` was not in its nominal at-float state: a side effect of executing the command that enables the elevation motor is that it resets the inner frame lock override. Accordingly, when the commanding system timed-out and queried the command scheduler for the first time, the command scheduler submitted the command to enable the elevation motor. This reset the inner frame lock override and had the opposite effect from the one desired: the elevation motor was vetoed by `mcp` which now again considered that the inner frame lock did not register as fully retracted.

This is a simple problem to overcome: simply sending the inner frame lock override command recovers proper operation. However, this was not done immediately upon reacquisition of the telemetry uplink, half an hour after loss of elevation control, due to insufficient feedback from the telescope, and the inexperience of the operators on duty with this particular aspect of the BLAST platform. When the true nature of the problem was discovered, at T+29^h:56, the override command was uplinked, and normal operation resumed.

Due to this failure, improvements were made in the flight code for *BLAST06*. First, a bit indicating whether the lock motor override was enabled or disabled was added to the downlink telemetry from the platform. Secondly, a command was added to the uplink which allows the operator on duty to indicate to the command scheduler that the initial float commands have been performed manually and should be disregarded by the command scheduler.

Because azimuthal control was unaffected by this failure, the telescope was kept away from the sun, and the only serious penalty of this failure was a loss in observing time.

3.4 *BLAST06*

BLAST's second long-duration balloon (LDB) flight was launched from Williams Field, Antarctica ($77^{\circ}52'$ S, $167^{\circ}04'$ E) at 1:54 UTC, 21 December, 2006. After 282 hours at float, the flight was terminated at 1:05 UTC, 2 January, 2007, and the telescope landed on the Antarctic plateau ($83^{\circ}28'$ S, $141^{\circ}36'$ E), 750 km southwest of McMurdo Station. The platform performed well, and successfully completed its mission goals. After landing, the telescope was dragged ~ 200 km after the parachute failed to detach from the payload upon landing. The remains of the gondola came to rest in a crevasse field, but the cryostat, Cassegrain assembly, and data storage vessel were eventually recovered from along the drag path.

3.4.1 Pre-flight Focus Tests

Before launch, a series of focus tests were performed in an attempt to focus the Cassegrain telescope before launch. The far field limit of the telescope is several kilometres but atmospheric absorption limits the range of the telescope on the ground to only a hundred metres or so. In order to deduce the far field focus position, therefore, several sets of precision spacers were manufactured. These were inserted behind the secondary mirror to allow focusing of the telescope on near field targets.

A focusing target was constructed of a liquid nitrogen cold source mounted behind a mechanical shutter, which chopped the source at low frequency. A dry nitrogen purge system was included in the design to prevent ice forming on the cold source window.

The cold source was placed on a translation stage able to move the source in the two dimensions perpendicular to the beam, and the target was placed at 50, 100 and 150 metres from the telescope. Code was added to `mcp` to permit it to both read the shutter signal and the position of the cold source, as well as move the source on the translation stage. Beam maps were created by rastering the source across the beam, first

horizontally, and then vertically.

Beam maps were made at the three distances and the far-field focus position for the secondary was extrapolated.

3.4.2 In-flight Focus Tests

Focus tests were also performed in flight to verify the focus position determined on the ground and the thermal correction co-efficients (Equation 2.17). The achievable precision for these tests is not better than 100 *pseudomicrons* ($\psi\mu\text{m}$. $1 \psi\mu\text{m} = 1.0583 \mu\text{m}$; see Section 2.10.3.1). Most importantly, a focus test was done early in the flight during line-of-sight (LOS) communication (Section 3.2.1). From these LOS tests we determined the telescope was out of focus by 500 $\psi\mu\text{m}$, and the focus position was adjusted accordingly. Further focus tests were performed throughout the flight, using the low-resolution TDRSS data to ensure the telescope remained in focus.

All in-flight test were performed by stepping the focus position back and forth by several hundred microns and measuring the point-source response at each focus position. The optimal focus position could be interpolated by ensuring the focus positions at which the focus tests were performed bracketed the peak response.

With the telescope in focus, the point-source sensitivities were found to be near nominal values ($240 \text{ mJy}/\sqrt{s}$ at $250 \mu\text{m}$). Therefore, the optimistic set of schedule files (Section 2.3.2.5) were used for the flight.

3.4.3 Balloon Environment

The mean *BLAST06* float altitude was 38.5 km. A plot of float altitude variation is presented as Figure 3.9. The 2006–2007 season was anomalistic for the Antarctic stratospheric vortex. The vortex did not set up over the pole, as usual but, rather, over the Transantarctic Mountains, significantly closer to the launch site. This anomalous vortex caused BLAST to stray significantly further south than normal (see Figure 3.1.2). Af-




Figure 3.9: Float altitude variations for *BLAST06*. Overall altitude variations during the flight are less than for *BLAST05* (Figure 3.7), but diurnal variations are equally pronounced. Minima at T+224^h and T+275^h correlate with travel over the Transantarctic Mountains. Ballast drops occurred at T+17^h, T+37^h, T+227^h, and T+275^h (dashed lines).

ter crossing the Transantarctic Mountains, *BLAST* headed almost directly south. This meant that gondola local sidereal time (LST) was changing at a significantly slower pace than anticipated when the schedule file was created. Accordingly, the LST-based observing schedule was spending more time than necessary on a given field. Later, when the gondola reached its minimum terrestrial latitude, the opposite occurred: less time than anticipated was spent on a particular field since, at this low latitude, LST changes more quickly for a given ground speed. The observation schedule had sufficient redundancy built into it that this exceptional LST behaviour did not significantly impact scientific observations.

Figure 3.10 shows the temperature profiles of selected systems during the flight. The overall trend is the same as that of *BLAST05* (Section 3.3.3), but diurnal variations are generally smaller: 5–10°C. Incidents when azimuthal control was lost due to gyroscope and pivot motor upsets (Section 3.4.5) are plainly apparent at T+176^h, T+226^h, and T+252^h. Straying close to the sun in these situations results in a general warming of the entire platform. The exception to this is the temperature of the solar arrays, which decrease in this situation, since they are driven into conductive mode as the angle of incidence of solar radiation becomes less optimal.

The pivot ran significantly hotter during *BLAST06* (bottom panels, Figure 3.10) than




Figure 3.10: Measured temperature profiles during *BLAST06*. The column on the left shows the first thirteen hours of the flight. The column on the right shows the final 228 hours (9.5 days) of the flight. The top panels show the temperatures of various outer frame electronics: the DAS battery, the ACS case, the inner frame lock motor, and the elevation drive housing. The middle panels show the temperatures on the inner frame: the DAS case, a thermometer on the frame itself, the primary and secondary mirrors. The bottom panels show the temperatures of the pivot, and one of the solar arrays.

it had in *BLAST05* (bottom panel, Figure 3.8). This resulted in occasional shutdown of the pivot motor controller, resulting in loss of azimuthal control (see Section 3.4.5). This was due to reflection of sunlight onto the pivot motor controller by the sun shields.

3.4.4 The Inner Frame Lock

The inner frame lock control algorithm suffered from a non-fatal error during flight. Upon reaching float, the lock was disengaged. After disengaging, the lock's stepper

motor controller did not respond to the stop command sent by `mcp`. As a result, it kept the lock motor on, attempting to retract the inner frame lock pin further, driving it up against a hard stop, where it remained spinning. No critical damage to the mechanism resulted, as evidenced by its subsequent proper functioning, but the inner frame lock drew 0.3 A for the first 80 hours of the flight, warming it significantly. At T+83^h the first gyroscope failure occurred (Section 3.4.5). As part of the recovery, the inner frame lock was re-engaged, allowing the inner frame lock to recover. When the lock was again disengaged, it operated properly, shutting down after disengaging. This is the cause of the sudden decrease in inner frame lock motor temperature at T+83^h (top panel, Figure 3.10).

This problem was a result of insufficient information obtainable from the inner frame lock. The position of the inner frame lock is read by `mcp` through a potentiometer incorporated into the lock actuator. The fully disengaged position of the lock is beyond the furthest position measurable by the potentiometer. Because of this, `mcp` is unable to use the potentiometer to determine whether the inner frame lock has retracted sufficiently. Instead, `mcp` relies on a limit switch, which is activated when the inner frame lock is fully disengaged.

There is no robust method for determining whether the inner frame lock has moved passed this limit switch. The stepper controller keeps an integrated count of steps sent to the motor. However, this does not necessarily reflect the actual position of inner frame lock. It also is reset to zero whenever the controller is rebooted, a situation which is difficult for `mcp` to detect. As a result, although `mcp` reads and stores this integrated count, it does not use it in its control algorithm.

3.4.5 The Pointing System

Except for isolated incidents, the *BLAST06* pointing system behaved well. As with *BLAST05*, the differential GPS (dGPS) failed to work for large portions of the flight,

likely due to the same reason (flexure of the antenna mounts, see Section 3.3.4). Both sets of gyroscopes suffered intermittent failures, forcing us to occasionally switch from one set to the other. Figure 3.11 and Table 3.1 show the recovery of the in-flight pointing solution after one such switch. The effect of switching gyroscope sets on the gyro offset calculation (Section 2.3.3.3) is further explored in Figure 2.4. The gyro offset takes approximately 50 minutes to settle on a new value after switching, making in-flight pointing poor during this time. A single bit flip in one of the DSPs, resulting from a cosmic ray hit on the DSP, led to catastrophic failure of the pointing system late in the flight; see Section 3.4.5.1 for details.

The additional sun shields added for *BLAST06* (Section 2.2.3) had the unfortunate side effect of reflecting sunlight onto the pivot motor controller at certain orientations relative the sun. This caused the controller to shut itself down several times during the flight when it exceeded its nominal maximum operating temperature of 65°C. Shutdown of the pivot controller resulted in saturation of the reaction flywheel, and, hence, loss of azimuthal control. Azimuthal control was restored by letting the telescope list in azimuth until sun was not being reflected onto the pivot, and the pivot cooled off. At T+252^h, azimuthal control was lost due to pivot shutdown for approximately an hour, in the midst of a telemetry blackout, during which time the telescope was drifting slowly towards the solar azimuth, resulting in the large temperature increase seen in Figure 3.10.

3.4.5.1 The Pointing Failure of T+225^h

The most catastrophic failure of the pointing system occurred during a TDRSS blackout and resulted in the telescope spinning out of control in azimuth for ~45 minutes. This incident was caused by a cosmic-ray-induced single bit flip in the gyroscope digital filter on one of the ACS DSPs, resulting in a non-physical step in the measured gyroscope velocity. The step occurred in the space of a single 100 Hz sample in the first-stage filtered gyroscopes (see Figure 3.12). This happened during a TDRSS blackout, when the

Flight Time	Event
T+82 ^h 59 ^m	The inner frame is locked at 38°
T+83 ^h 01 ^m	After ~12 minutes without a good solution, the star cameras time out and compute a whole-sky solution. This indicates the in-flight pointing solution is ~5° off in azimuth. The in-flight solution is corrected, and this results in an offset between the magnetometer solution and the star cameras.
T+83 ^h 03 ^m	The in-flight pointing solution is switched from the DSP 3000 gyroscope set over to the E-core 2000 gyroscope set. Immediately all elevation solutions begin drifting due to the different offset between the two elevation gyroscopes (see Figure 2.4). The raw elevation encoder, however, properly reports that the inner frame is not moving. The magnetometer solution also starts drifting, due to the incorrect gyroscope solution. The star camera solutions, which have separately calculated gyro offsets, do not drift.
T+83 ^h 11 ^m	The inner frame lock is disengaged and mcp enters <i>stop</i> mode, during which time it tries to drive the gyro velocity to zero. The incorrect gyro offsets, however, result in movement of the telescope. The inner frame is driven down into the lower hard elevation stop at 20°.
T+83 ^h 16 ^m	The telescope is commanded to point towards anti-sun, a ~25° change in azimuth, at its current supposed elevation (55°).
T+83 ^h 17 ^m	The ISC is vetoed, resulting in a small jump in the pointing solution.
T+83 ^h 18 ^m	The OSC is vetoed. As a result of vetoing both star cameras, the in-flight pointing solution immediately jumps to the elevation encoder and magnetometer solutions, resulting in a slew back to the commanded orientation.
T+83 ^h 21 ^m	The ISC gets a solution and corrects the ISC-only solution. Gyro offset induced drift is still apparent.
T+83 ^h 26 ^m	The OSC gets a solution and corrects the OSC-only solution. Both star camera solutions are now in reasonable agreement with the coarse sensors: 5° off in azimuth (<i>cf.</i> T+83 ^h :02) and quite close in elevation.
T+83 ^h 31 ^m	With the star cameras getting reliable solutions, they are enabled again, resulting in a 5° jump in azimuth. Enabled, they are able to re-lock and get frequent solutions. This overcomes the gyro offset induced drift.
T+83 ^h 32 ^m	The coarse sensors are re-trimmed to the star cameras, fixing the 5° discrepancy in the magnetometer-only solution. The gyro offset filter has had 30 minutes to recalculate the offsets after the gyro switch, and are already pretty good, and the pointing solution is working reasonably well again.

Table 3.1: Events in the recovery of the pointing solution after failure of an azimuthal gyroscope for *BLAST06*. See also Figure 3.11.




Figure 3.11: Recovery of the in-flight pointing solution after switching gyroscope sets due to failure of an azimuthal gyroscope. In both panels, the azimuth and elevation calculated by the in-flight pointing solution is the solid line, the ISC-only solution is dotted, and the OSC-only solution is dashed. In the azimuth plot, the dot-dashed line is the pointing solution for the magnetometer only; in the elevation plot, the dot-dashed line is the elevation encoder. The gyro offsets calculated during this failure has been presented in Figure 2.4. Events are further explained in Table 3.1.

only data arriving from the gondola was the low-rate (Section 2.13.3), which meant the operators on duty were unable to see the spike itself, making the subsequent diagnosis of the problem very difficult. TDRSS was not re-acquired until 22 minutes after the incident, by which time the gondola had started spinning rapidly in azimuth (see below).

The signature of this step (*i.e.* it is apparent in the first-stage filtered gyros; it is persistent; it occurs over a single sample) indicate that the bit flip most likely occurred in the first-stage accumulator of the digital filter for the affected gyroscope. The DSPs'

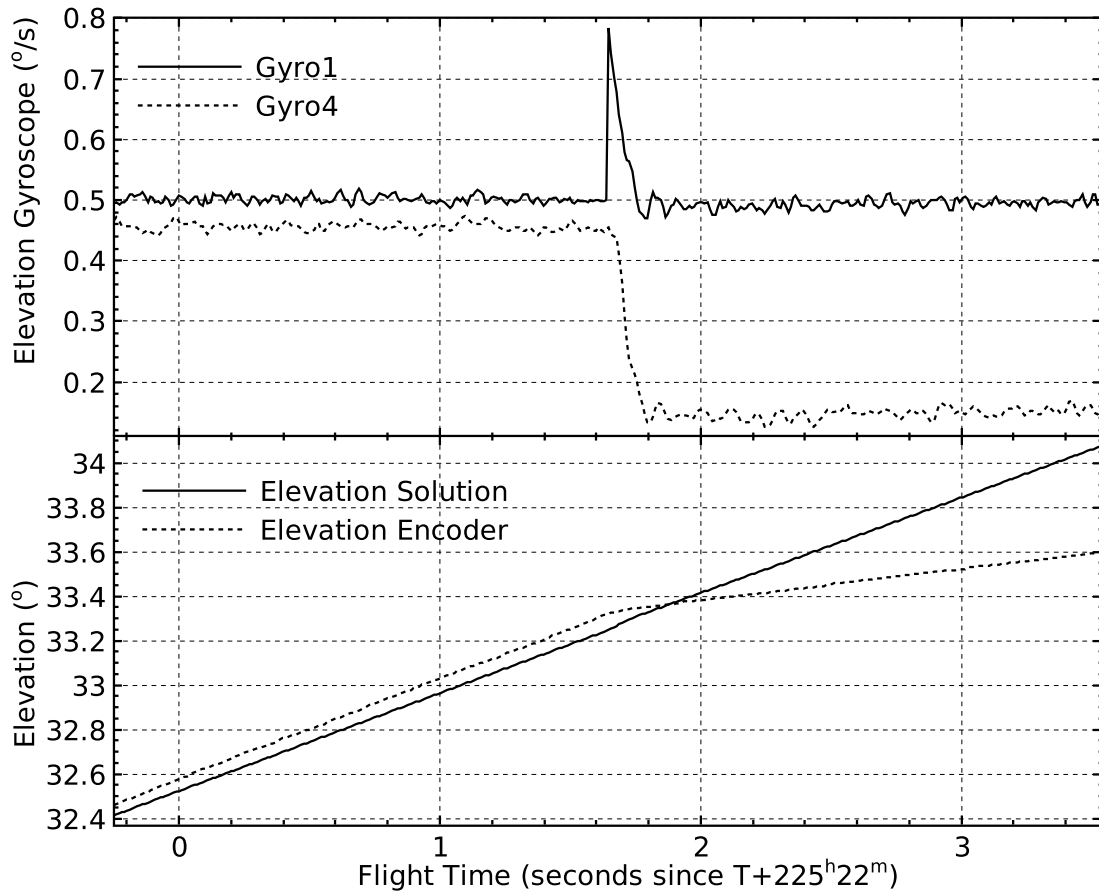


Figure 3.12: The incident that eventually lead to the gondola spinning out of control. The top panel shows the elevation gyroscopes. The solid curve is the E-core 2000 elevation gyroscope (Gyro1), after passing through only the first stage of the digital filter, and which is being used by the in-flight pointing solution. The dotted curve is the DSP-3000 elevation gyroscope (Gyro4), after passing through all four stages of the filter. A step occurs in Gyro1 in a single 100 Hz sample. It is quickly corrected for by the pointing system but, since it is not a real signal, this results in an undetected change in the elevation velocity, as is apparent in the Gyro4 signal. The bottom panel shows the effect on the elevation solution (solid) which immediately starts to diverge from the true elevation as measured by the elevation encoder (dotted).

RAM is not augmented by error-correction circuitry (ECC). Because of this, the error went uncorrected.

The effect of this bit flip was that the elevation gyroscope began to report rotation rates that were $\sim 0.3/s$ larger than the true elevation rate. Loss of elevation control was immediate. First, mcp unwittingly drove the inner frame down against the lower elevation stop. Once the inner frame's motion had been halted by the hard stop, the subsequent effect was to cause an increase of the in-flight elevation solution at the rate of the gyroscope discrepancy.

This runaway of the elevation solution eventually led to loss of azimuthal control. The azimuthal and roll velocities used in control algorithms Equation 2.19 and Equation 2.21 are calculated by rotating the gyroscopes into horizontal co-ordinates:

$$\begin{pmatrix} \omega_{az} \\ \omega_{roll} \end{pmatrix} = \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix} \begin{pmatrix} \omega_2 \\ \omega_3 \end{pmatrix} \quad (3.4)$$

where ϕ is the elevation calculated by the in-flight pointing solution, and $\omega_{2,3}$ are the angular velocities measured by Gyro2 and Gyro3. The ever-increasing elevation solution eventually caused an inversion in the rotation, resulting in a positive feedback in the azimuthal control (Figure 3.13). This feedback eventually saturated the reaction flywheel. The pivot, trying to drive the reaction flywheel rate back down to zero, then began to rotate the gondola in the direction opposite the saturation. Because of the elevation inversion, however, this only caused the reaction wheel to attempt to rotate faster, an unrecoverable situation. The pivot managed to spin the gondola up to $\sim 69^\circ/s$ before itself saturated.

The uncontrolled spinning of the experiment was eventually curtailed by vetoing the entire motion control system, and waiting for the gondola to spin down. This did not, however, correct the gyroscope error, and re-enabling of motion control resulted in a resumption of the erroneous behaviour. The error was eventually fully corrected




Figure 3.13: The positive feedback loop set up in the azimuth control as a result of the loss of the elevation solution. The top panel shows the elevation solution, increasing at $\sim 0.3^\circ/\text{s}$. The middle panel shows the azimuthal motion measured by Gyro2 and Gyro3 induced by the positive feedback. The motion is purely azimuthal, and the relative amplitude between the signal in the gyros indicates the true elevation of the inner frame. The bottom panel shows the rotation rate of the azimuthal flywheel which is driving the azimuthal motion. Shortly after $T+225^{\text{h}}31^{\text{m}}5$, the reaction flywheel saturates at 0.9 revolutions per second, after which the platform begins spinning in azimuth, driven by the pivot.

by an accidental upset of the ACS, caused by a power spike induced by the reaction flywheel motor, which rebooted the affected DSP. Once control of the experiment had been regained, normal scanning could resume.

The bit flip occurred at $T+225^{\text{h}}22^{\text{m}}02^{\text{s}}$, with azimuthal spinning commencing at $T+225^{\text{h}}31^{\text{m}}38^{\text{s}}$. Motion control was vetoed at $T+226^{\text{h}}12^{\text{m}}30^{\text{s}}$, and full control was regained at $T+226^{\text{h}}52^{\text{m}}49^{\text{s}}$, resulting in ~ 45 minutes of uncontrolled spinning and ~ 1.5 hours of loss of motion control. Although normal operation was resumed after this, the Cassegrain

telescope, which had seen direct sunlight during the spin, was very warm, and, as a result, out of focus, resulting in a loss of sensitivity for the remainder of the flight.

It should be noted that, had sufficient error checking been in place for the elevation solution, so that it would have rejected obviously bad solutions (*i.e.* solutions outside the physical limits of the inner frame, 20–60°), azimuthal control would not have been lost, and the gyro offset calculation (Section 2.3.3.3) could have eventually compensated for the bit flip in the elevation gyroscope, with no other adverse effects.

Both this failure and the gyroscope switch recovery of Figure 3.11, indicate the fundamental fragility of this particular star camera—gyroscope-based in-flight pointing solution. When the gyroscopes fail in such a way that they no longer report the true motion of the telescope, the in-flight pointing solution quickly diverges from the true orientation of the telescope. Once in this situation, the star cameras lose lock. They are, then, only able to generate pointing solutions when they time out and compute a whole-sky solution. The rate of these whole-sky solutions (~ 1 per minute at best) is insufficient to recover from the situation. Indeed, in both these failures, control was not regained until the star cameras were vetoed, and `mcp` reverted to a coarse sensor—gyroscope-based pointing solution. Additionally, this bit-flip induced failure could have been resolved by vetoing the star cameras. This would have eliminated the elevation inversion, because elevation would have been based solely on the elevation encoder.

3.4.6 Telemetry

Due to the extremely southerly flight path during *BLAST06* (Figure 3.1.2), the geosynchronous TDRSS satellites were at a very low elevation. The TDRSS high-rate return was at times intermittent, with blackouts lasting as long as sixteen hours. Overall, the TDRSS high-rate downlink (Section 2.13.2) was available for 79.3% of the flight.

An excerpt of the TDRSS high-rate data stream acquired in flight is presented as Figure 3.14. Comparison of the data streams from the two flight computers show the




Figure 3.14: A short section of the TDRSS high-rate telemetry acquired during the *BLAST06* flight showing one of the $250\ \mu\text{m}$ detectors. The half-minute gap in the data, representing the loss of two 15-second long TDRSS packets, was due to control switching from MCC1 (*sam*) on the left to MCC0 (*frodo*) on the right, after an upset of MCC1. MCC0 gained control at $T+107^{\text{h}}09^{\text{m}}875$ (dashed line). Data from MCC0 start before this, since MCC0 was already accumulating data for a TDRSS packet when it acquired control. The plot shows the significant compression artefacts associated with the TDRSS data stream from MCC0, which aren't present in MCC1.

significant compression artefacts associated with data from MCC0 which are not present in the data from MCC1. It is thought that TDRSS data corruption on MCC0 is due to the TDRSS compression thread running out of time on the slower MCC0. These artefacts are most pronounced when the slope of the data is large, indicating the corruption is likely due to mixing of old and new data. The MCC1 data presented contain regular spikes. These artefacts are due to the circulant nature of the low-pass filter.

The artefacts in MCC0's TDRSS high-rate were surprising when first seen in flight, but investigation of TDRSS data acquired before launch show the corruption in that

data as well, when MCC0 was transmitting. The flight team should have been spent practicing with the TDRSS high-rate data before flight, which could have helped discover and correct this problem.

Τὸ δ' ἔνεκεν τῆς ἐν ταῖς παρ' ἕκαστα μεθόδοις κατασκευεῖας ἐκόντα καταφρονῆσαι τινος τῶν τηλικούτων, ἥλικα καὶ παρὰ τὰς ὑποθέσεις καὶ παρὰ τὰς τηρήσεις αὐτὰς ἐνδέχεται παραθεωρεῖσθαι, τοῦ μὲν κατὰ τὸ ἀπλούστερον χρησίμου πλείστην αἴσθησιν ἐμποιεῖ, τοῦ δὲ περὶ τὰ φαινόμενα διαμαρτανομένου ἢ οὐδεμίαν ἢ παντάπασι βραχεῖαν.

—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Γ':ζ' [1]

Chapter 4

Data Reduction

Significant processing of the raw data obtained in flight is required to be able to extract scientific results. The ultimate goal of the data reduction pipeline is the generation of high-quality submillimetric maps. The data reduction pipeline is outlined in Figure 4.1 and consists of ten steps, many of which are run in parallel:

1. Initial data preparation
2. Detector cleaning
3. Pointing solution generation
4. Detector response profiling
5. Preliminary map making
6. Noise estimation
7. Flatfielding and absolute calibration
8. Detector alignment
9. Field alignment
10. Final map making

The data reduction pipeline is summarised below. Significant portions of this pipeline have already been described in detail elsewhere, and references are given when appropriate.

Time-oriented data products created by the reduction pipeline are typically stored as dirfiles. The dirfile format makes adding new fields to an existing database trivial. This, in turn, simplifies the data reduction process, allowing raw and processed data to be worked on in the same database.




Figure 4.1: A block-level diagram of the BLAST data reduction pipeline used for both *BLAST05* and *BLAST06*. Rectangles represent operations; hexagons represent data products. Dashed lines indicate the potential for multiply iterated recursion in the reduction.

4.1 Initial Data Preparation

Two early tasks are performed on the raw data sets recovered from the flight computers with the goal of creating a canonical data set which can be used for the remainder of the data reduction. The first of these tasks is to combine the data sets obtained by the two flight computers into a combined data set. The second task is to annotate the data set to locate all science-grade observations made in flight. The output product of the initial data preparation is a merged, annotated data set.

4.1.1 Merging and Verification

For the sake of in-flight redundancy, both master control computers (MCCs, Section 2.4) record the entire data stream coming from the BLASTbus to disk. This results in two almost-complete data sets, one on each computer. Portions of the flight where only one of the computers has the data are typically due to the reboot of the other computer.

For *BLAST05*, 6.15 minutes of data occur on only one computer. Of this total, 51 seconds were lost on MCC1 due to a reboot, and the remainder, 5.3 minutes, were lost on MCC0, associated with the reboot of that computer upon reaching float. This time is on the order of the time-out period of the watchdog, 4.96 minutes (see Figure A.4), potentially indicating a sequence of back-to-back reboots. An additional single frame was found missing on the not-in-charge computer during the flight, likely due to a missed frame sync word from the BLASTbus.

For *BLAST06*, 20.5 minutes of the data are on only one computer. Fifteen minutes of this is due to a single framefile chunk being unrecoverable from the hard drive of MCC0, due to physical damage suffered by the hard drive during the drag after landing. The remaining 5.5 minutes are associated with five reboot events of one or the other of the two computers, each lasting roughly a minute.

The two datasets need to be merged to create a canonical data set which can be




Figure 4.2: A histogram of discrepant samples in the bolometer channels between MCC0 and MCC1 during the *BLAST05* flight. The solid histogram shows discrepancies for all channels binned by DAS ADC Card node number, which run from five through sixteen. Cards are queried on the BLASTbus for bolometer data in node number order, so the peak around node 10 seems to indicate the cause of the discrepancies are related to BLASTbus timing. The dotted histogram are the discrepant samples for only the bias reference channels. These are read together, before any bolometers are read, and their discrepancy counts show no dependence on node number.

used for data reduction. Before merging the *BLAST05* dataset, the two datasets were byte-compared to check for differences between the two. A number of differences were found, and counts of the differences in the bolometer channels are presented in Figure 4.2. In total, 148 318 discrepant samples were found in the bolometers (0.36% of the total bolometer data). The gyroscopes were also checked for discrepant samples, but none were found. Discrepancies were dealt with by preferring data from the in-charge computer.

Institutions participating in the data reduction were each shipped copies of the raw

data recovered from the flight hardware before the merging process has been finalised. To avoid having to send further datasets through the mail, all participating institutions were asked to perform this merging themselves, from the raw data they had already received. Once the data had been merged into a single data set, it was then verified against a reference set of SHA-1 [38] sums produced from a version of the merged dataset which was byte-compared to the raw data. This ensured that all participating institutions were working from the same dataset.

The same procedure was performed for *BLAST06*, once the data were recovered from the damaged hard drives, with similar results.

4.1.2 Initial Inspection

The second major preliminary task is the annotation of the data set. The goal of this annotation is to identify times during the flight when scientific-quality data was obtained, and to identify the field of each observation. Since all pointing commands are recorded into the data stream, part of this can be automated. However, each observation must be inspected by hand to ensure data quality. The output from this step is a list of observations indexed both by position in the data set and field observed.

For *BLAST05*, 73 hours (72.8%) of the data, in 23 fields (Table 6.1), were found to be of suitable quality for use in the data reduction. For *BLAST06*, 225 hours (78.7%) in 32 fields were selected (Table 6.2).

4.2 Detector Cleaning

Once the merged data set has been created, several subsequent tasks can be accomplished simultaneously. One major task is the cleaning of the detector time streams. This process is described in detail in [39].

Within the portions of science quality data identified during the initial inspection

(Section 4.1.2), bolometer data may be corrupted by three types of unwanted events:

- Cosmic rays which deposit power on the bolometers;
- Pulses from the calibrator lamp (see Section 2.10.4);
- Non-physical glitches arising in the data acquisition pipeline.

A set of detector flags is created which indicate data that exhibit these types of non-astronomical signals. For *BLAST05*, 4.3% of bolometer samples (~ 5 hours per bolometer) were flagged and excluded from the data reduction. Of these flagged samples, 262 samples (~ 1 sample per bolometer) were flagged as non-physical glitches and 0.045% (~ 3 minutes per bolometer) were flagged as pulses from the calibrator lamp.

4.2.1 The Detector Transfer Function

The transfer function of the data acquisition system is dominated by the transfer function of the digital filter (Figure 2.12). To this is added the ~ 2 ms bolometer time constant (Section 2.5.1.1). The full detector transfer function, incorporating both of these factors, is measured from cosmic rays in the flight data. This procedure is explained in detail in [30], and returns, in addition to the transfer function itself, a measurement of the bolometer time constant. After removal of the digital filter, the transfer function extracted from this analysis is best fit assuming two thermal time constants, a fast bolometer time constant of $\tau_1 \approx 2$ ms, and a second, slow, bolometer time constant of $\tau_2 \approx 50$ ms. The measured detector transfer function for *BLAST05* is presented in Figure 4.3.

4.2.2 Deconvolution and Gap Filling

Given the raw bolometer signals, the detector flags, and the transfer function as input, a purpose built program cleans the detector TOD streams in three steps. First, the non-physical glitches are removed, replaced by a straight line fit to the data. The glitches must




Figure 4.3: The measured transfer function for *BLAST05*. The transfer function used to deconvolve the bolometers is purely real, and is dominated by the transfer function of the digital filter (Figure 2.12). The sampling Nyquist frequency (50.08 Hz) is marked with a dashed line.

be removed first, to prevent them from causing ringing during deconvolution. Second, the bolometers are deconvolved by dividing the Fourier transform of the data by the transfer function measured in Section 4.2.1. The final step is to remove the flagged, non-astronomical signals.

Gaps resulting from removal of the corrupt data are filled with simulated noise to reduce the effect of these gaps on subsequent analysis of the data. Two gap-filling methods have been used. For *BLAST05*, gaps were filled using linear predictive (LP) methods, which are explained further in Section 5.4.5. LP-gap filling creates simulated data in the gap with a similar spectrum as data around the gap. This works well when data surrounding the gap has a low signal-to-noise ratio, and performed adequately for *BLAST05*. A




Figure 4.4: A short section of a single $250\ \mu\text{m}$ bolometer channel for *BLAST05*, showing both a cosmic ray hit in the bolometer (left) and a calibration source pulse (right). The dotted curve is the raw data obtained by the data acquisition system. The solid curve is the clean data after deconvolution, removal of these non-astronomical features, and LP gap filling. The two gap-filled regions are demarcated by dashed vertical lines. The high-frequency noise amplification of the deconvolution is apparent outside the gap-filled regions.

sample of deconvolved and gap-filled data is presented in Figure 4.4. The power spectra of data before and after deconvolution and gap-filling is presented in Figure 4.5.

For *BLAST06*, this method was found to fail in situations when gaps were located amongst data with a large signal component. In these situations, LP gap filling fails since the gap is short in comparison to the width of the typical signal. Instead, for *BLAST06*, gaps were filled by a straight-line fit to the surrounding data, to which is added white noise with amplitude characteristic of the particular bolometer being cleaned.




Figure 4.5: The spectral density of a 117-minute-long sample of low signal-to-noise data from a single $250\ \mu\text{m}$ channel acquired during *BLAST05*. The grey, dotted curve is the power spectrum of the data before cleaning, the black, solid curve is the power spectrum after cleaning. The high-frequency amplification of the deconvolution is apparent. The reduction in power on scales between 0.1 and 5 Hz is due to removal of calibrator pulses. The overall result is a whitening of the power spectrum at frequencies above approximately 0.2 Hz.

4.3 Post-Flight Pointing Solution

The generation of a high-quality pointing solution is an important goal of post-flight data reduction. The in-flight pointing solution calculated by `mcp` (Section 2.3.3) is of insufficient precision ($\sim 1'$) to be used in post-flight analysis. The post-flight pointing solution is generated by re-analysing data from the gyroscopes and star cameras. The process is described in detail in [28, 33].

4.3.1 Star Camera Analysis

Data from the star cameras are re-analysed as part of the post-flight pointing solution reduction. This is done firstly to ensure that the pointing solution obtained in flight is correct, and secondly to obtain solutions in one and two star cases for which no solution was calculated in flight. The plate scale and rotation of the star camera is also fit for. For *BLAST05*, star camera exposures stored on the star camera computers (SCCs) were re-analysed in order to verify the blob-centroiding performed in flight. For *BLAST06*, star camera analysis was performed on blob centroids and fluxes stored in the `mcp` data stream. This was necessitated by the loss of the star cameras during the drag after landing.

Computation of the star camera pointing solution proceeds iteratively: the first pass of analysis is only able to reconstruct the pointing solutions for exposures with three or more stars. The first full, post-flight pointing solution computed, as described in Section 4.3.2, is only reliable around these star camera solutions. This full solution can then be used to solve for one and two star exposures near the exposures already solved for. This process is further iterated to successively solve for one and two star solutions increasingly far from the three or more star solutions. When the process finishes, the star camera pointing solution is complete.

4.3.2 Pointing Solution Calculation

The post-flight pointing solution is calculated using a Kalman filter [40] to combine the star camera pointing solution with the integrated gyroscope data. The Kalman filter operates on a prediction-and-correction model with the goal of producing a least-squares solution of the pointing. Given a current orientation, the gyroscopes are integrated to evolve the solution into the future. The periodic solutions from the star camera solution are used to correct the evolved solution, mitigating gyroscope noise. It should be noted

that this Kalman solution proceeds, broadly, in the same manner as the calculation of the in-flight pointing solution does in `mcp` (Section 2.3.3), and, indeed, the in-flight pointing solution is also an implementation of a modified Kalman filter.

In addition to the orientation of the telescope, the post-flight pointing solution also needs to calculate the gyroscope offsets and the rotation of the star cameras in the coordinate system defined by the gyroscopes. The output from the post-flight pointing solution calculation is a three-component telescope orientation, which may be rotated into any necessary fiducial co-ordinate system in which maps need to be made.

The pointing solution may be calculated by either evolving time forwards or backwards. The two solutions calculated this way are independent, and may be combined to further refine the pointing solution. The process by which they are optimally combined is known as *Kalman smoothing* [40].

For *BLAST05*, a pointing solution consisting of only data from the ISC and the DSP-3000 set of gyroscopes was computed. This provided a solution of sufficient precision to make maps with the large beams of *BLAST05*. For *BLAST06*, better precision is needed, and star camera solution combining both star cameras is used instead.

The Kalman filter produces a pointing solution which describes the orientation of the star camera. Pointing calibrators observed throughout the flight can be used to measure the rotation between the star camera orientation and the orientation of the submillimetre beam; see Section 4.9.

4.4 Detector Response Profiling

Variations in detector responsivity over the course of the flight must be corrected in order to meaningfully combine data from multiple passes of the same target. A known calibration source (the *cal-lamp*, Section 2.10.4) is flashed periodically, once every 10 to 15 minutes, over the course of the flight to track these responsivity changes. This

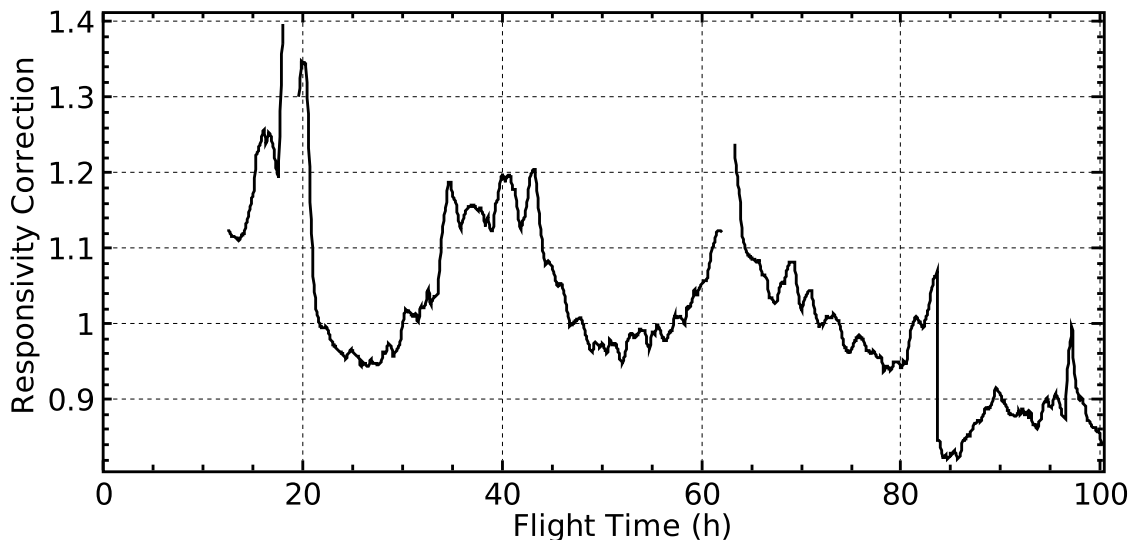


Figure 4.6: The calculated detector responsivity correction for a single $250\ \mu\text{m}$ channel during *BLAST05*. The scale is arbitrary: the choice of scale is reversed by the flatfielding constants determined later in the analysis pipeline (see Section 4.7.1). The calibrator lamp (Section 2.10.4) was not turned on until $T+12^{\text{h}}6$. Before this time, no responsivity calibration is possible. Breaks at $T+19^{\text{h}}$ and $T+63^{\text{h}}$ are during ^3He fridge cycles (Section 2.6), during which time bolometer response drops to zero and the calculated correction becomes ill-defined. The drop in the correction at $T+83^{\text{h}}6$ is due a change in the $250\ \mu\text{m}$ bias voltage, which increased the array’s responsivity.

calibration source is similar to the one incorporated into the SPIRE instrument [41], and has negligible variance over time, meaning variation in detector response to the cal-lamp is due solely to changes in the detector’s responsivity.

A special set of cleaned data is used to analyse the calibration pulses. This data set is cleaned as normal (Section 4.2.2), except that cal-lamp pulses are exempted from removal and gap-filling. The calibration pulses are extracted from these cleaned bolometer TOD streams, and fit to a template pulse. The amplitude fit to this template is directly proportional to the responsivity of the bolometer. The time-oriented detector response profiles generated from this procedure are smoothed and subsequently used to scale the detectors’ response curves during map making. Measured responsivity variations are less than 2% RMS. A sample of responsivity correction calculated for *BLAST05* is presented in Figure 4.6. This process is described in detail in [9, 32].

4.5 Map Making

Map making is discussed in detail in Chapter 5. Map making is important not only for creating the final maps from the data; it is also used extensively during the data reduction process. Three map makers were developed for BLAST: a modified naïve map maker, which simply bins filtered data according to the pointing solution, SANEPIC [39], a matrix solver, and almagest, an iterative map maker; see Section 5.3.

Since the ultimate goal of the data reduction process is well-made maps, it is not unexpected that intermediate maps are used during the data reduction process. Which map maker is used depends on the task at hand. For preliminary maps, the naïve map maker is often the best choice, due to its speed. Tasks such as detector alignment, which is solely a function of the pointing solution, gain no benefit from being run on the optimal map makers, since all map makers use the same pointing solution. Again, for beam profiling, the naïve map maker is often good enough again, since it has good fidelity at small angular scales.

Map-based reduction tasks involving flux or noise extraction, on the other hand cannot use the naïve map maker since, by design, it poorly estimates noise. In these cases, one of the optimal map makers is used instead, typically the particular map maker which is intended to make the final maps.

Final maps of *BLAST05* data were made with SANEPIC; therefore, the final round of reduction was done using SANEPIC maps as well. Intermediate revisions of these data products, however, can still be made with almagest, which is faster and simpler to use, decreasing time between iterations.

4.6 Noise Power Spectra

Proper estimation of the noise in the bolometers is required for the generation of optimal maps. The power spectrum of a representative sample of low signal-to-noise data has

already been presented as Figure 4.5. Noise in the BLAST detectors is observed to be non-stationary and highly correlated at large scales, both in time and between detectors. These factors complicate proper estimation of the noise.

4.6.1 Total Noise Estimation

We model a detector TOD stream, an m -element vector, \mathbf{d}_t , as (*cf.* Equation 5.1):

$$\mathbf{d}_t = P_{tp}\boldsymbol{\Sigma}_p + \mathbf{n}_t \quad (4.1)$$

where $\boldsymbol{\Sigma}_p$, is the true signal map, a n -element vector, P_{tp} is the $m \times n$ pointing matrix, and \mathbf{n}_t is the noise in the data. (In this section we use the linear algebra notation of Chapter 5: \cdot_t indexes bolometer samples, and \cdot_p indexes map pixels.) We estimate the noise with the following procedure:

1. Use \mathbf{d}_t in a low signal-to-noise field as the initial estimate of the noise, $\hat{\mathbf{n}}_t$, for any particular observation of interest. For *BLAST05*, the field used was ELAIS-N1 [42], a blank, extragalactic field (see Table 6.1).
2. Generate an intermediate map, $\hat{\boldsymbol{\Sigma}}_p$, with the optimal map maker (Section 5.1.2) using $\hat{\mathbf{n}}_t$ as the estimate of the noise.
3. Given the intermediate map, $\hat{\boldsymbol{\Sigma}}_p$, which is now the current best estimate of the signal, compute a new best estimate of the noise with:

$$\hat{\mathbf{n}}_t = \mathbf{d}_t - P_{tp}\hat{\boldsymbol{\Sigma}}_p. \quad (4.2)$$

4. Iterate steps two and three until convergence is achieved.

This iterative process can be accomplished by almagest during normal operation. Noise can also be estimated using SANEPIC, by iterating the map maker through several runs.

For SANEPIC operation, convergence is usually achieved with three to six iterations [39].

4.6.2 Correlated Noise Estimation

The noise in the detectors is significantly correlated between bolometers at frequencies smaller than ~ 0.1 Hz. Ignoring this correlated component leads to an underestimation of the noise at on these scales. To account for the correlated noise, the noise is modelled as:

$$\mathbf{n}_{dt} = \mathbf{u}_{dt} + \alpha_d \mathbf{c}_t \quad (4.3)$$

where \mathbf{c}_t is a common mode component of the noise, \mathbf{u}_{dt} is the uncorrelated portion of the noise, and α_d is a scaling constant which is fit for. Here the index \cdot_d runs over detectors.

The common mode is evaluated separately for each of the three BLAST bands. The common mode spectrum for a given band may be estimated by averaging the total noise across all detectors of that band, Fourier transforming the result, subtracting the white noise level from the spectrum and then low-passing the spectrum at 1 Hz. The scaling factor α_d is estimated by requiring the uncorrelated component of the noise to be white through 0.1 Hz. The noise used in this procedure must be derived from properly flatfielded detectors, in order to properly weight the detectors (see Section 4.7.1). Figure 4.7 presents a section of data broken up into correlated and uncorrelated components.

SANEPIC, which is used to generate final common mode spectra, doesn't use the procedure outlined above. Rather, it uses a blind component separation method described in [43]. This method allows simultaneous estimation of all components of the noise decomposition by fitting all the data for a given pass and wavelength band.

4.7 Flatfielding and Absolute Calibration

Once preliminary maps have been made, the final two steps of detector calibration can be performed. Flatfielding involves the rescaling of individual bolometers in order to produce




Figure 4.7: The spectral density of a low signal-to-noise portion of data from a single $250\ \mu\text{m}$ detector acquired during *BLAST05* broken up into correlated and uncorrelated components. The solid curve is the total noise power spectrum of Figure 4.5 after deconvolution. The dashed curve is the correlated noise spectrum, calculated by scaling the common mode noise spectrum, and the dotted curve is the reconstructed uncorrelated component of the noise. Correlated noise dominates on frequencies smaller than 0.1 Hz. The uncorrelated noise is white down to ~ 0.03 Hz, where it is ~ 5 times smaller than the correlated component.

uniform response across the arrays. Absolute calibration involves determining the ratio between detector voltage and submillimetric flux. These operations are described in detail in [32, 9].

4.7.1 Flatfielding

For *BLAST05*, flatfielding was accomplished by making single bolometer maps of the proto-planetary nebula CRL 2688 [44]. This was a point-source calibrator which has been observed early in the flight with a sufficiently dense scan pattern that fully sampled

Colour (μm)	Coefficient (GJy mV^{-1})	Uncertainty (%)
250	7.61	12
350	3.16	10
500	1.56	8

Table 4.1: Absolute flux calibration co-efficients for *BLAST05* obtained from measurements of Arp 220.

maps could be made from single bolometers¹ The bolometers were then re-scaled so that the point source flux extracted from these single bolometer maps were identical. The result of the flatfielding is presented in Figure 4.8.

For *BLAST06*, a similar procedure was done using the hypergiant star VY Canis Majoris. These maps were not as fully sampled as CRL 2688 had been for *BLAST05*, so small groups of detectors were used to make the maps, instead of individual bolometers.

4.7.2 Absolute Calibration

Ideally, BLAST would have observed Uranus or Mars to obtain an absolute calibration. These objects have systematic uncertainties $<5\%$ in the submillimetre [45, 46]. Unfortunately, the ecliptic does not pass through the region of sky observable to BLAST. Instead, for *BLAST05*, the ultra-luminous infrared galaxy (ULIRG) Arp 220 [47] was used as the primary flux calibrator. A template spectral energy density (SED) was fit to measurements of Arp 220 obtained between $60 \mu\text{m}$ and $1100 \mu\text{m}$ [48, 49, 50, 51, 52, 53, 54]. Absolute fluxes determined from this fit are compared to point-source flux measurements extracted from BLAST maps of Arp 220. Absolute flux calibration co-efficients are presented in Table 4.1.

¹Typically, the full array is needed to fully sample the map.




Figure 4.8: A diagram of the flatfielding co-efficients calculated for *BLAST05*. The diagram shows all three arrays, with detectors laid out according to their physical location. The 250 μm array is on top, the 350 μm array in the centre, and the 500 μm array at bottom. The number in the centre of the detector hexagon indicates the applied flatfielding co-efficient. The small number at the top of the hexagon is the detector’s identification number. The central bolometer at 500 μm (#1610) is used as the reference, and is assigned a flatfield co-efficient of 100. Broken or otherwise unused detectors are labelled with a large ‘B’.

4.8 Detector Alignment

Due to the projection of the array on the focal plane, in a given sample, two detectors on the same array sample different parts of the sky. In order to make meaningful multi-detector maps, the rotation between the beam of each detector and the nominal boresight of the telescope must be measured and corrected for.

Before *BLAST05*, the three arrays were surveyed in the lab to determine nominal values for the detector offsets. The results of this survey was a constant fiducial separation between adjacent detectors in the six cardinal directions of the hexagonally packed arrays of $45''.6$, $61''.3$, and $93''.1$ at $250\ \mu\text{m}$, $350\ \mu\text{m}$, and $500\ \mu\text{m}$ respectively.

Due to the large beams of *BLAST05* (Section 3.3.2), the numbers from this pre-flight survey were accurate enough to produce final maps without significant broadening of point sources. For *BLAST06*, however, the significantly smaller beams required these numbers to be refined.

An empirical spacing was generated for *BLAST06* by making maps, one at each colour, of a pointing calibrator (VY Canis Majoris) without taking the bolometer offsets into account. For this procedure, a naive map is sufficient, since this is only dependent on the pointing solution. This procedure results in a map containing an imprint of the bolometer array, from which bolometer offsets may be measured.

In an attempt to match the measured offsets with the fiducial bolometer spacing determined before *BLAST05*, the fiducial spacing was deformed according to the transformation:

$$\begin{pmatrix} \xi'_i \\ \eta'_i \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} Z & 0 \\ S_\xi & Z \end{pmatrix} \begin{pmatrix} \xi_i \\ \eta_i \end{pmatrix} \quad (4.4)$$

where Z , a scale factor, S_ξ , a skew factor, and θ , an array rotation, are fit for. The measured bolometer offsets, and fiducial fit, are presented in Figure 4.9. Best fit values for the transformation parameters are presented in Table 4.2. The fit works reasonably well at $350\ \mu\text{m}$ and $500\ \mu\text{m}$, but the fit for the bottom half of the $250\ \mu\text{m}$ array is




Figure 4.9: Bolometer offsets measured for *BLAST06* (asterisks) compared to the fiducial spacing determined for *BLAST05* transformed according to Equation 4.4 and Table 4.2 (squares). The $250\ \mu\text{m}$ array is presented at top, the $350\ \mu\text{m}$ below it, and the $500\ \mu\text{m}$ array at bottom. Holes in either pattern are due to bolometers which are broken or otherwise unusable. At $250\ \mu\text{m}$ and $500\ \mu\text{m}$, the central bolometer defines the origin. Since there is no central bolometer at $350\ \mu\text{m}$, the origin is defined to be equidistant from the four bolometers surrounding the centre.

systematically poor, indicating that some more complex effect is responsible for the bolometer offsets. For making maps of *BLAST06* data, the empirically measured offsets are used.

4.9 Pointing Offsets

The reconstructed pointing solution (Section 4.3.2), while self-consistent on short timescales, suffers from long timescale drifts due to deflection between the star cameras' orientation and the boresight of the telescope. For pointing purposes, the orientation

Colour (μm)	Fit Parameters			Mean Displacement
	θ	Z	S_{ξ}	
250	2.91°	1.065	-0.093	5.1230''
350	0.73°	1.075	-0.003	3.9645''
500	1.83°	1.050	-0.005	2.5696''

Table 4.2: Best fit values for the transformation parameters of Equation 4.4. The resulting bolometer grid is displayed in Figure 4.9. The mean displacement between the measured bolometer offsets from *BLAST06*, and the fiducial offsets calculated from the transformation is presented in the rightmost column.

ISC is used as the reference. BLAST typically scans a given field several times during a flight. Successive passes over the same field may be separated by long periods of time. As a result, if meaningful multi-pass maps are to be made, corrections to the pointing solution must be calculated to ensure all passes of the field are consistently aligned.

Relative offsets between two passes of the same field are calculated by creating separate maps of the two passes, and then maximising the convolution of the two maps by moving the second relative to the first. In regions of high signal-to-noise, offsets can also be calculated by measuring positions of bright sources across each map. The calculated field offsets for *BLAST05* are presented in Figure 4.10

Although this procedure ensures consistency between passes of any given field, it does not recover the absolute pointing offset of the telescope. Periodic observations of pointing calibrators (Section 3.2.2) are used to calculate the rotation between the star cameras' field of view and the boresight. These offsets are presented in Table 4.3. Further refinement of the absolute offset for particular fields is accomplished by aligning known sources observed in the map. For convenience, these corrections may be folded into the field offsets calculated above.




Figure 4.10: Calculated field offsets for *BLAST05*. Offsets are only calculated for fields observed on multiple passes. Since the first pass of any field is used as the baseline, its offsets are artificially set to zero. Points which were calculated using whole-map cross-correlation typically have precisions of three seconds of arc. Offsets with higher precision are calculated by aligning individual sources in the maps.

Colour (μm)	<i>BLAST05</i>		<i>BLAST06</i>	
	Elevation Offset (')	Cross-Elevation Offset (')	Elevation Offset (')	Cross-Elevation Offset (')
250	-12.996	-17.438	-8.500	-2.833
350	-12.835	-17.097	-7.834	-3.500
500	-13.325	-17.376	-8.333	-3.500

Table 4.3: Measured rotation between the ISC and the submillimetric beam at each colour. Differences between *BLAST05* and *BLAST06* are due to differences in the mounting of the star camera, and changes in the telescope optics (see Section 2.10).

Προσῆκει γὰρ τοῖς τῶ ὄντι φιλαλήθως καὶ ζητητικῶς τῇ τοιαύτῃ θεωρίᾳ προσερχομένοις μὴ πρὸς μόνην τὴν τῶν παλαιῶν ὑποθέσεων διόρθωσιν συγχρησθαι τῇ καινότητι τῶν ἐπὶ τὸ ἀδιστακτότερον εὕρισκομένων ἐφόδων, ἀλλὰ καὶ πρὸς τὴν τῶν ἰδίων, ἂν οὕτως ἔχωσι, μηδὲ αἰσχρὸν ἡγεῖσθαι μεγάλης τινὸς καὶ θείας οὔσης τῆς ἐπαγγελίας, κἂν ὑπ' ἄλλων καὶ μὴ μόνον ὑφ' αὐτῶν τῆς ἐπὶ τὸ ἀκριβέστερον τύχῃ διορθώσεως.

—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Δ':θ' [1]

Chapter 5

Map Making

BLAST’s observing strategy requires the creation of post-flight maps from the time-ordered data obtained in flight. Different data analysis tasks have different requirements on maps produced, and three different map makers are used for BLAST analysis as outlined below.

5.1 The Map-Making Equation

Consider the bolometer time-ordered data streams as a m -element vector, \mathbf{d}_t . Here \cdot_t runs up to m (*i.e.* it indexes detectors and samples of time). We model \mathbf{d}_t as:

$$\mathbf{d}_t = P_{tp}\boldsymbol{\Sigma}_p + \mathbf{n}_t \quad (5.1)$$

where \mathbf{d} is the TOD stream, $\boldsymbol{\Sigma}$ is the observed sky signal we’re interested in recovering, and \mathbf{n} is the noise in the deconvolved bolometer time-ordered data TOD stream (Section 4.2.2). Here $\boldsymbol{\Sigma}$ is a n -element vector of pixels on the sky ordered in any convenient way, and \cdot_p runs up to n (*i.e.* over map pixel indices). The matrix P_{tp} is a known $m \times n$ matrix which converts the observed sky map into a TOD stream. It is simply the pointing solution. P_{tp} is a sparse matrix, having only m non-zero elements, one per row,

representing the pixel observed by any bolometer sample.

It should be pointed out that the pointing matrix, P_{tp} , need not be as simple as the one proposed above. Specifically, the point-spread function PSF could be included in the naïve pointing matrix described above, permitting the map maker to deconvolve the PSF as it generates the map. Leaving this out of the pointing matrix, as is done in all BLAST map-makers, results in the production of maps which are convolved with the beam.

For the symmetric beams of *BLAST06*, the beam can be simply deconvolved from the maps after they are produced. The asymmetric beams of *BLAST05*, however, could have benefited by taking into account the PSF during the map-making stage. This was not done, due to concerns of computational complexity.

The goal of map making is to recover $\boldsymbol{\Sigma}$ given \mathbf{d} ; *i.e.* the map maker computes

$$\hat{\boldsymbol{\Sigma}}_p = \phi(\mathbf{d}_t) \quad (5.2)$$

where $\hat{\boldsymbol{\Sigma}}_p$ is the map produced by the map maker, and $\phi : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a function selected for desirable qualities of the maps it produces. BLAST map making limits itself to functions ϕ of the form:

$$\phi(\mathbf{d}) = \hat{\boldsymbol{\Sigma}} = (P^\top W P)^{-1} P^\top F \mathbf{d} \equiv Q \mathbf{d} \quad (5.3)$$

for some $m \times m$ matrices W and F , which reduces map making to a linear problem. Here \cdot^\top denotes the matrix transpose. This form has been used extensively for CMB data analysis [55, 56, 57, 58, 59], especially in the case of maximum likelihood map making (Section 5.1.2). Particular choices of W and F used for BLAST analysis are discussed below.

Equation 5.3 can be seen as a three step process. Working from the right, first the data is filtered in the time domain by the filter F . Second, the filtered data is binned into a map based on the pointing solution P^\top . Finally, the pixels in the map are rescaled

based on some time-domain weighting function W .

5.1.1 Naïve Map Making

A computationally simple option is to choose $W = I$. Then Equation 5.3 reduces to the naïve map making equation:

$$\hat{\boldsymbol{\Sigma}} = (P^\top P)^{-1} P^\top F \mathbf{d}. \quad (5.4)$$

In this equation, $(P^\top P)$ is simply the hits-per-pixel for the map. If we also choose $F = I$, map making simply becomes the process of calculating the mean bolometer value in each pixel. Due to the large $1/f$ noise component (see Figure 4.5), this is not a very useful choice of F for BLAST. A better choice for making naïve maps is $F = LH$ where H is a high-pass filter which cuts off above the $1/f$ -knee (~ 0.1 Hz) and L is a pixel antialiasing low-pass filter. A typical value for the high-pass cutoff is 0.2 Hz. The low-pass cutoff is calculated from $f_{hp} = v_s/2s_p$ where v_s is the scan velocity on the sky (typically $\sim 0.1^\circ/\text{s}$) and s_p is the diameter of a pixel.

5.1.2 Maximum Likelihood Map Making

The optimal map makers, SANEPIC and almagest, choose $F = W = N^{-1}$, yielding:

$$\hat{\boldsymbol{\Sigma}} = (P^\top N^{-1} P)^{-1} P^\top N^{-1} \mathbf{d} \quad (5.5)$$

where $N = \langle \mathbf{nn}^\top \rangle$ is the time-domain noise correlation matrix.¹ This particular choice has several desirable properties:

1. It maximises

$$\log \mathcal{L}(\mathbf{d} | \boldsymbol{\Sigma}) = -\frac{1}{2} (\mathbf{d} - P\boldsymbol{\Sigma})^\top N^{-1} (\mathbf{d} - P\boldsymbol{\Sigma}) \quad (5.6)$$

which is the logarithmic likelihood of $\boldsymbol{\Sigma}$ given \mathbf{d} , assuming \mathbf{n} is Gaussian and

¹Here we've assumed, without loss of generality, that $\langle \mathbf{n} \rangle = 0$.

stationary. For this reason, Equation 5.5 referred to as *maximum likelihood map making*.

2. Similarly, it minimises

$$\chi^2 = (\mathbf{d} - P\hat{\boldsymbol{\Sigma}})^\top N^{-1} (\mathbf{d} - P\hat{\boldsymbol{\Sigma}}). \quad (5.7)$$

That is, it produces the minimum variance map from the data.

3. The reconstruction error ϵ :

$$\epsilon \equiv \hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma} = (QP - I) \hat{\boldsymbol{\Sigma}} + Q\mathbf{n} \quad (5.8)$$

reduces simply to $\epsilon = Q\mathbf{n}$, which means that $\hat{\boldsymbol{\Sigma}}$ differs from $\boldsymbol{\Sigma}$ by a noise term that is independent of the signal in the map²

4. It minimises $\langle |\epsilon^2| \rangle$, subject to the constraint $W = F$.

Within the choice of Equation 5.5, the quantity $N_{pp}^{-1} \equiv P_{pt}^\top N_{tt}^{-1} P_{tp}$ is sometimes referred to as the *inverse pixel-pixel noise covariance matrix*.

5.2 Iterative Map Making

The matrices to be computed for Equation 5.5 are large: for the 10'' 250 μm BLAST map of Cygnus-X, the number of data samples is $m \approx 6 \times 10^8$ and the number of map pixels $n \approx 3 \times 10^6$. To simplify the computation, an iterative solution to Equation 5.5 can

²This is true whenever $W = F$.

be used:

$$\mathbf{n}^{(j)} = \mathbf{d} - P \hat{\Sigma}^{(j)} \quad (5.9.1)$$

$$N^{(j)-1} = \left\langle \mathbf{n}^{(j)} \mathbf{n}^{(j)\top} \right\rangle^{-1} \quad (5.9.2)$$

$$\hat{\Sigma}^{(j+1)} - \hat{\Sigma}^{(j)} = (P^\top P)^{-1} P^\top N^{(j)-1} \mathbf{n}^{(j)} \quad (5.9.3)$$

where $\cdot^{(j)}$ runs over iteration. This procedure is designed to simultaneously estimate \mathbf{n} and $\hat{\Sigma}$. If the noise power spectrum is known *a priori*, which is typically the case for final BLAST maps (see Section 4.6), N^{-1} can be held fixed, and the procedure will iterate only on $\hat{\Sigma}$.

It is easy to show that Equation 5.9 converges to Equation 5.5 by simple substitution:

$$\begin{aligned} \hat{\Sigma}^{(j+1)} - \hat{\Sigma}^{(j)} &= (P^\top P)^{-1} P^\top N^{(j)-1} \mathbf{n}^{(j)} \\ &= (P^\top P)^{-1} P^\top N^{(j)-1} \left(\mathbf{d} - P \hat{\Sigma}^{(j)} \right) \\ &= (P^\top P)^{-1} P^\top N^{(j)-1} \left[\mathbf{d} - P \left(P^\top N^{(j)-1} P \right)^{-1} P^\top N^{(j)-1} \mathbf{d} \right] \\ &= (P^\top P)^{-1} \left[P^\top N^{(j)-1} \right. \\ &\quad \left. - P^\top N^{(j)-1} P \left(P^\top N^{(j)-1} P \right)^{-1} P^\top N^{(j)-1} \right] \mathbf{d} \\ &= (P^\top P)^{-1} \left(P^\top N^{(j)-1} - P^\top N^{(j)-1} \right) \mathbf{d} \\ &= 0 \end{aligned} \quad (5.10)$$

i. e., if $\hat{\Sigma}^{(j)}$ is the map of Equation 5.5, the map of the subsequent iteration $\hat{\Sigma}^{(j+1)}$ is no different: convergence has been achieved.

Since N^{-1} and $\hat{\Sigma}$ are being estimated simultaneously, care must be taken to ensure that, for a given dataset \mathbf{d} , that convergence actually does occur. The convergence properties of Equation 5.9 may be modified by replacing $(P^\top P)^{-1}$ with $(P^\top C P)^{-1}$, where C is a constant diagonal matrix chosen for the benefit it provides to convergence

[59]. For BLAST, convergence is achieved without this modification, and no such matrix was introduced.

5.3 The BLAST Map Makers

Three map makers have been written for BLAST, each useful for different aspects of analysis:

- *Optibin*, a naïve map maker which computes Equation 5.4.
- *SANEPIC*, a maximum-likelihood map maker which computes Equation 5.5.
- *almagest*, an iterative map maker which computes Equation 5.9.

Several common tasks required by all map makers are discussed below, along with the Optibin and SANEPIC. Almagest is discussed further in Section 5.4.

5.3.1 The Pointing Solution

The pointing solution calculated in post-flight analysis (Section 4.3.2) describes the orientation of the ISC star camera. In order to make submillimetric maps, the map maker must rotate this star camera pointing solution into the reference frame of the submillimetric beam by the rotation, R_{sc} , calculated as described in Section 4.9. In addition to this offset, both the individual bolometer offsets, $R_b(d)$, (Section 4.8) and the field offsets (Section 4.9), $R_f(t)$ must also be applied to the pointing solution. To simplify the procedure, these rotations can be combined into one:

$$R(t, d) = R_f(t)R_b(d)R_{sc} \tag{5.11}$$

which depends on both time, t , and detector number, d .

This rotation can be a computationally expensive portion of the map making procedure. One possible way to get around this would be to store one pointing solution on disk for each bolometer for the entire flight (~ 254 in total). This is expensive both in terms of disk space and disk I/O, and has not been done for the BLAST analysis. However, *almagest* does compute and store $R(t, d)$ for all appropriate values of t and d once, which is relatively inexpensive to read from disk later, and speeds rotation slightly.

5.3.2 Calibration

A second task required for all map makers is the ability to make properly calibrated maps. Many intermediate maps used in the analysis pipeline (Section 4.5) do not require calibration, and this step can be skipped. The calibration consists of three parts: the detector responsivity correction, c_{dr} , (Section 4.4), the flatfielding co-efficients c_{ff} , (Section 4.7.1), and the absolute calibration, c_{abs} , (Section 4.7.2).

Of these, the absolute calibration, which is simply a constant scaling, can be applied to the completed map after production, and is therefore not an integral part of the map making procedure. On the other hand, both the flatfielding coefficients and the responsivity corrections must be applied to the data before inputting them into the map maker, if meaningful maps are to be made.

The dirfile containing the analysis data contains a field which is $\mathbf{d}_{calib} = c_{abs}c_{ff}c_{dr}\mathbf{d}$ which can be used for making calibrated maps. A drawback of using this field is that some of the power in the low-frequency baseline of \mathbf{d} is pushed to higher frequencies when multiplied by the slowly varying c_{dr} . To overcome this problem, the map makers may decide, instead, to apply the calibration co-efficients after first removing a baseline from the data. The dirfile stores \mathbf{d}_{calib} as a *derived channel*, which means that the database contains the time-ordered data \mathbf{d} and c_{dr} , and the constants c_{ff} and c_{dr} , plus rules to compute the product, which is done whenever the field is retrieved. As a result, having the map maker compute this product itself is no more computationally expensive than

reading \mathbf{d}_{calib} from the dirfile.

5.3.3 Optibin

Optibin implements Equation 5.4 with the simple high-pass filter described in Section 5.1.1, typically with a 0.2 Hz cutoff. Performance could be improved with the addition of the pixel-antialiasing low-pass filter as well, which would reduce the noise in the map. This naïve map maker reproduces small scale features quickly and with high-fidelity, at the expense of large-scale structure. As such, it is well suited for point-source analysis, for instance extragalactic maps, while being poor choice for recovering the diffuse structure in galactic maps.

5.3.4 SANEPIC

SANEPIC computes Equation 5.5 directly. This process can be logically divided into two parts:

1. Compute

$$\Xi_p \equiv P_{pt}^\top N_{tt}^{-1} \mathbf{d}_t; \quad (5.12)$$

2. Solve

$$N_{pp} \hat{\boldsymbol{\Sigma}}_p = \Xi_p. \quad (5.13)$$

Computation of Equation 5.12 is straightforward. The calculation is simplified by assuming N is circulant, allowing the multiplication to be efficiently performed in Fourier space. Equation 5.13 is either solved directly, by computing N_{pp}^{-1} , or else by iterating towards a solution using a conjugate gradient method similar to that presented in [60]. SANEPIC is described in detail in [39].

5.4 Almagest

Almagest creates maps by iterating Equation 5.9. It is designed to produce high-fidelity maps quickly. The almagest map making procedure is as follows:

1. Initialise the map, either with a null map, or else with a naïve map generated according to Equation 5.4.
2. Load pointing data and compute the rotation described in Section 5.3.1.
3. Load bolometer data, gap-fill, and prefilter it.
4. Load noise power spectra (if pre-computed).
5. Compute $\mathbf{n}^{(j)} = \mathbf{d} - P\hat{\Sigma}^{(j)}$.
6. Calculate the weight spectrum (essentially $N^{(j)-1}$).
7. Compute a new best-estimate of the signal (Equation 5.16).
8. Bin the new signal estimate into a new map (Equation 5.15.2).
9. Repeat steps 2. through 8. until convergence is achieved.

Because almagest is designed to be used for much of the data analysis pipeline (see Section 4.5), a robust front-end has been implemented, which simplifies operation of the map maker.

Due to the high volume of data required to make a BLAST map, almagest is designed to be run on high-performance parallel computing clusters (*a Beowulf*), such as the McKenzie and Sunnyvale clusters at CITA. Almagest is perfectly parallel in that, provided sufficient compute resources exist, it scales as $\mathcal{O}(1)$ with increasing quantities of input data.

The following sections discuss the design and operation of the almagest map maker.

5.4.1 The Algorithm

In the interests of speed, Equation 5.9 is diagonalised, allowing it to be computed in the Fourier domain. Solving this equation for $\hat{\Sigma}^{(j+1)}$ gives:

$$\hat{\Sigma}^{(j+1)} = (P^\top P)^{-1} P^\top \left[N^{(j)-1} \mathbf{d} + (1 - N^{(j)-1}) P \hat{\Sigma}^{(j)} \right]. \quad (5.14)$$

The quantity in square brackets is the new best estimate of the signal in the map, and the process is broken up into two parts (*viz.* steps 7. and 8. above):

$$\mathbf{e}^{(j)} = N^{(j)-1} \mathbf{d} + (1 - N^{(j)-1}) P \hat{\Sigma}^{(j)} \quad (5.15.1)$$

$$\hat{\Sigma}^{(j+1)} = (P^\top P)^{-1} P^\top \mathbf{e}^{(j)}. \quad (5.15.2)$$

The calculation of the estimate, \mathbf{e} , is now diagonalised, and re-arranged to produce:

$$\epsilon^{(j)} = \delta - (1 - \omega^{(j)}) \nu^{(j)} \quad (5.16)$$

where $\epsilon^{(j)}$ is the Fourier transform of $\mathbf{e}^{(j)}$, δ the Fourier transform of \mathbf{d} , $\nu^{(j)}$ the Fourier transform of $\mathbf{n}^{(j)}$, and $\omega^{(j)} \equiv \nu^{(j)-2}$ is known as the *weight spectrum*, and is unit normalised. Equation 5.16 can be seen as a simple selection function: at frequencies where the noise is small ($\omega \sim 1$), the data are preferred; where the noise is large ($\omega \sim 0$), the previous iteration's map ($= \mathbf{d} - \mathbf{n}$) is preferred.

5.4.2 Map Initialisation

Equation 5.9 makes no indication of the initial map, $\hat{\Sigma}^{(0)}$, which is to be used to start the iterative process. A poorly chosen initial guess may prevent the map from converging at all. However, for any reasonable initial map, the converged map will be independent of the initial guess. Making a good initial guess can significantly speed convergence. (To

take an extreme case, choosing $\hat{\Sigma}^{(0)} = \hat{\Sigma}$ would result in the operation converging in a single iteration.)

In CMB analysis, where this technique was pioneered, typically the initial map is chosen to be $\hat{\Sigma}^{(0)} = \mathbf{0}$. As a result, the noise estimated for the first iteration is simply the data itself. This is a good approximation in the CMB case, where the signal-to-noise ratio is low. For BLAST maps, however, it can be a poor choice, especially in high signal-to-noise galactic maps.

Accordingly, *almagest* has the option of using a naïve map for an initial guess:

$$\hat{\Sigma}^{(0)} = (P^\top P)^{-1} P^\top H(f) \mathbf{d} \quad (5.17)$$

where the high-pass cutoff frequency, f , may be chosen to optimise convergence. Experience indicates that the optimal choice of f depends on the spectral energy density of the map, and that typical values are $f \sim 0.05$ Hz, which is significantly lower than the cutoff value used for a production-quality naïve map (see Section 5.3.3). A side-effect of this feature is that *almagest* can be used to produce naïve maps identical to those made by *Optibin*, but with somewhat more computational overhead than needed for the simpler *Optibin* program.

5.4.3 Data Partitioning

In order to be able to work with the data in a reasonable fashion, the data for a particular map are usually broken up into distinct chunks. Typically a chunk will represent the data from a single detector for a single visit of a field. Long (many hour) visits of a single field may also be artificially broken up to make chunks of more manageable size. The typical upper limit on the size of a single chunk is approximately 1.44×10^6 samples, representing ~ 2 hours of data acquisition with a single bolometer. Longer un-interrupted periods are rare, due to the desire to have regularly scheduled calibrator observations (Section 3.2.2)

throughout the flight.

Because the map making algorithm, Equation 5.14, is a linear operation, maps of each chunk are made independently, and then the maps are summed at the end of each iteration. This model of map making allows perfect scaling on parallel clusters, since each chunk can be independently computed, and summed maps need only to be passed between nodes. The down-side to this method is that correlations between chunks are poorly modelled (see Section 5.4.7).

5.4.4 Data Pre-Filtering

The deconvolved data read in by the map maker undergoes a certain amount of manipulation before being fed into the map maker proper. The input data to the map maker contains a slowly varying bias, which adds broad-spectrum power to the data. To speed the convergence process, a polynomial is least-squares fit to the data and subtracted to remove the bulk of this bias. The order of this polynomial depends somewhat on the pathology of the data itself, but is typically fourth order. Once the polynomial has been subtracted, the baseline is small enough that errors associated with application of the calibration constants (Section 5.3.2) are negligible.

After calibration, the data are padded out to a convenient length. Typically this length is a power of two, which is advantageous for the FFTs needed in the map maker. However, the FFT library used by `almagest` (FFTW [61]) has optimal algorithms even for prime sizes. When not using strict powers of two, `almagest` limits itself to lengths, l , of the form:

$$l = 2^a 3^b 5^c 7^d \quad a, b, c, d \in \mathbb{N}_0 \quad (5.18)$$

where \mathbb{N}_0 is the set of natural numbers including zero. The only potential benefit to using lengths which are not integer powers of two is to save memory during computation: in the worst case, padding to a power of two can double the memory requirements of a

given chunk. Reducing memory consumption may speed the operation of the map maker by permitting more chunks to be computed in-core (see Section 5.4.8).

Most of the data reduction is done in the Fourier domain, which assumes that a given chunk of data are circulant (*i.e.* its last sample and first sample are sequential). The second reason to pad the data is to decouple the effects of Fourier domain filtering across the ends of the data. To accomplish this, the data is always padded by a small amount, even were the data length before padding to be a perfect power of two.

Padding is accomplished using linear-prediction; see Section 5.4.5. Since this padding can result in the first sample and last sample of the chunk having different values, the final step of the data pre-processing is to remove a straight line from the data which ensures that both the first and last sample of the data are zero, which significantly reduces ringing at the ends.

5.4.5 Linear Predictive Gap Filling

Linear prediction (LP) is the process of extrapolating future values of a discrete signal based on the previous samples [62, 63]. Almagest uses linear prediction for both gap filling and padding of the data (Section 5.4.4). Linear prediction is also used by the deconvolver to fill gaps (Section 4.2.2).

Linear prediction considers a segment of data, \mathbf{d} , p samples long. It is assumed these data are of the form (*cf.* Equation 5.1):

$$\mathbf{d}_t = \mathbf{s}_t + \mathbf{n}_t \quad (5.19)$$

where \mathbf{s} is the underlying signal, and \mathbf{n} is a noise component. Again, \cdot_t indexes discrete time samples and runs up to p . Now we wish to estimate the value at any sample, τ , which may be one of the existing $t \in [1, p]$, or it may be a new point outside this range.

The true signal there is s_τ . We calculate an estimate s_τ which is linear in \mathbf{d}_t :

$$s_\tau = \boldsymbol{\lambda}_\tau^\top \mathbf{d} + x_\tau \quad (5.20)$$

where $\boldsymbol{\lambda}_\tau$ are the linear co-efficients which depend on τ , and x_τ is the discrepancy between the estimate and the true signal.

The co-efficients $\boldsymbol{\lambda}_\tau$ which minimise x_τ in the least-squares sense may be calculated by minimising:

$$\langle x_\tau^2 \rangle = \langle (\boldsymbol{\lambda}_\tau^\top \mathbf{d} - s_\tau)^2 \rangle \quad (5.21.1)$$

$$= \boldsymbol{\lambda}_\tau^\top (\langle \mathbf{ss}^\top \rangle + \langle \mathbf{nn}^\top \rangle) \boldsymbol{\lambda}_\tau - 2 \langle s_\tau \mathbf{s}^\top \rangle + \langle s_\tau^2 \rangle \quad (5.21.2)$$

$$= (\boldsymbol{\lambda}_\tau - \hat{\boldsymbol{\lambda}}_\tau)^\top (S + N) (\boldsymbol{\lambda}_\tau - \hat{\boldsymbol{\lambda}}_\tau) - \mathbf{S}_\tau^\top (S + N)^{-1} \mathbf{S}_\tau + \langle s_\tau^2 \rangle \quad (5.21.3)$$

where we've made the substitutions:

$$S \equiv \langle \mathbf{ss}^\top \rangle \quad N \equiv \langle \mathbf{nn}^\top \rangle \quad \mathbf{S}_\tau \equiv \langle s_\tau \mathbf{s} \rangle \quad (5.22.1)$$

$$\hat{\boldsymbol{\lambda}}_\tau \equiv (S + N)^{-1} \mathbf{S}_\tau \quad (5.22.2)$$

and, in Equation 5.21.3, have completed the square. We have also assumed that \mathbf{s} and \mathbf{n} are uncorrelated (*i.e.* $\langle \mathbf{sn}^\top \rangle = 0$). Clearly Equation 5.21.3 is minimised if $\boldsymbol{\lambda}_\tau = \hat{\boldsymbol{\lambda}}_\tau$. Accordingly, the minimum variance estimate for the true signal at τ can now be calculated as:

$$\hat{s}_\tau = \mathbf{S}_\tau^\top (S + N)^{-1} \mathbf{d}. \quad (5.23)$$

The method of linear prediction is, then, as follows:

1. Estimate S and N , the autocorrelations of the underlying signal and noise, using the data \mathbf{d} .
2. Compute $\boldsymbol{\lambda}$, the *linear prediction (LP) co-efficients*, by solving Equation 5.22.2 for

known values of τ (i.e. $\tau \in [1, p]$).

3. Extrapolate the data by computing necessary new points via Equation 5.20.

The algorithm implemented is taken from [65], which is based on method of Burg [66]. For both gap filling and padding, `almagest` uses $p = 200$ data points to estimate $r = 160$ LP co-efficients. This procedure assumes the data are stationary, which is quite reasonable over the span, p , of data which is considered (~ 2 seconds).

The computation time of the co-efficients is $\mathcal{O}(p^2 + r)$, while the calculation of N new points is $\mathcal{O}(Nr)$. As a result, this is often the dominant step in the iteration of the map-maker. If the properly prepared data can be remembered from one iteration to the next by being computed in-core (see Section 5.4.8), this need only be done on the first iteration.

The algorithm also tends to compute co-efficients λ which sum not to 1, but some number slightly smaller. As a result, the extrapolated values tend towards zero as the length of the extrapolated section becomes longer. This is not an issue for filling gaps in the data, which are short, generally < 500 ms. Indeed, these short gaps could have been filled with simple white noise without significant degradation in maps produced. Figure 4.4 shows some gaps of typical length after being filled with LP-generated fake data. For padding the ends of the chunks out to a convenient length, the tendency of the LP-generated data towards zero is apparent. Here too, however, this is not a serious issue, since the LP-generated data is only used to provide a smooth transition between the real and generated data: the nature of the generated data far from the transition point is not important.

5.4.6 The Weight Spectrum

The weight spectrum is central to the iterative process, since it is this vector which largely determines the properties of the converged map. The weight spectrum is, in its essence,

simply the inverse of the noise power spectrum:

$$\omega^{(j)} = \nu^{(j)^{-2}} \quad (5.24)$$

The weight spectrum is unit normalised simply as a convenience, in order not to have to re-normalise the signal estimate (Equation 5.16) during the map-binning step.

At low frequency, where the $1/f$ component in ν dominates, ω tends towards zero. This has the result that Equation 5.16 prefers the current map over the data. In situations where a field is well cross-linked (Section 3.1.3), this is a reasonable assumption, since the low frequency structures in one scan direction will be measured at a higher frequency in the other scan direction. In the absence of cross-linking, however, these modes are cannot be recovered from in the data, and the map will be unconstrained at the largest scales. In order to provide a small constraint to these modes, even in the absence of cross-linking, the calculated weight spectrum is artificially raised from zero to some small number (typically 10^{-5} of full scale) at the lowest frequencies. This very weakly constrains the large scale structure in the estimate to the data, in the absence of any other information. Since this artificial weight floor is very small, more reliable data coming from cross-scans at higher frequency will still be preferred, if such data exist.

The weight spectrum is also rolled off to zero at high frequency to provide a pixel antialiasing low-pass filter. This pixel antialias filter has already been discussed in Section 5.1.1, and has a cutoff frequency which is calculated by $f_{hp} = v_s/2s_p$, where v_s is the scan velocity on the sky ($\sim 0.1^\circ/\text{s}$, Section 3.1.1), and s_p is the diameter of a pixel (typical values are in the range $10''\text{--}60''$). A sample weight spectrum is presented as Figure 5.1.

5.4.7 Correlated Noise

Almagest has difficulty properly correcting for the correlated noise between bolometers. This is not due to the linearisation of the problem (Equation 5.9), but rather a con-




Figure 5.1: A typical weight spectrum, $\omega^{(j)}$, calculated by almagest. Because important features are at low frequency, the weight spectrum is logarithmically binned. It has also been unit normalised. Below 0.02 Hz, the spectrum, which would normally roll off to zero, is elevated to the “weight floor”, 10^{-5} . The feature at 0.26 Hz corresponds to the fundamental scan frequency (the time it takes to raster from one side of the map to the other). At 18 Hz, corresponding to the pixel sampling Nyquist frequency, the spectrum is rolled off to zero artificially.

The weight spectrum plot can be read as a selection function: for a given frequency, the ratio between the space under the curve and the space above the curve indicate the relative weighting between the data, \mathbf{d} , and either the signal recovered from the previous iteration’s map, \mathbf{s} , or, for the case of the pixel-antialias low-pass, formal zero, \emptyset .

sequence of the decision to process data from different bolometer in separate chunks (Section 5.4.3). Correlated noise can be properly handled in an iterative map maker if the map maker reduces the data from all correlated detectors simultaneously, since this amounts merely to making the same assumption about the noise that is made by SANEPIC: that N_{tt} is circulant. Due to the design of almagest, such a procedure is beyond its scope, and has not been attempted.

The effective result of ignoring correlations is that the weight spectrum (Section 5.4.6)

generally overweights the data, \mathbf{d} , at scales where the correlated noise dominates. Several attempts were made to correct for this by reducing the weight spectrum at those frequencies by some function of the correlated noise spectra determined by SANEPIC. Results of these tests have indicated this is not particularly effective at correcting for the correlated noise and, as a result, *almagest* is typically run with no correction in place. Despite this, runs of *almagest* on simulated data indicate that it does still outperform SANEPIC, which fully corrects for correlated noise, in certain regimes. See Section 5.5 for details.

5.4.8 Performance

Almagest run times are dominated by the amount of random access memory (RAM) available to the map maker. To run efficiently, *almagest* requires 33 bytes per bolometer sample, or 11.9×10^6 bytes per hour of data for a single bolometer. Additionally, 32 bytes are required for each map pixel. The largest *BLAST05* map, Cygnus-X, requires approximately 20 Gbytes of main memory to run optimally, which is typically not available on a single machine. To overcome this, *almagest* was designed to run on a parallel processing machine (*beowulf*), where typically more memory is available.

Since it is also useful to be able to run *almagest* on single processors, *almagest* is designed to scale itself to work with whatever amount of memory is available to it. It does this by deciding to schedule one or more chunks of data “out-of-core”. All out-of-core chunks share work buffers, which means memory need only be allocated for the largest of the out-of-core chunks.

Almagest is significantly slower when working out-of-core. In this case, for each chunk, it has to rotate the pointing solution (Section 5.3.1), load the bolometer data from disk, calibrate (Section 5.3.2), and pre-filter it (Section 5.4.4) every iteration. This results in an overall slowdown of a factor of four compared to “in-core” chunk processing, assuming reading from disk is not unusually slow.

Several hundred iterations are typically required to achieve convergence. When running completely in-core on a CITA beowulf, provided sufficient memory exists for further expansion, almagest scales as $\mathcal{O}(1)$ with increases in data size, to first order, and a single iteration can be computed in a few seconds. On a single processor, almagest computation time scales as $\mathcal{O}(n \log n)$ (limited by FFT computation) if sufficient memory exists to compute all data in-core. Once running out-of-core, almagest scales as $\mathcal{O}(n^2)$, limited by the computation of the LP-coefficients needed to pad the data.

5.5 Almagest Simulations

We now present results of testing almagest on the same set of simulated data upon which SANEPIC was tested in [39]. This analysis compares almagest, which doesn't properly correct for correlated noise (see Section 5.4.7), with the full SANEPIC treatment, which does properly account for correlated noise. Conclusions are drawn about the usefulness of the two map makers in their ability to maximally extract science from the BLAST datasets.

5.5.1 Input Data

Twenty simulated maps were created as input to the simulations. The pointing used to generate these maps correspond to the observations of intermediate velocity cloud (IVC) G86.0+38.3 [67] from the *BLAST05* flight. This 2 deg² field was chosen for being well cross-linked, and small enough that running multiple simulations could be accomplished in a reasonable amount of time. The field is well cross-linked and consists of four visits with slightly more than six hours of data in total.

For each of the twenty simulations, a diffuse signal component, uncorrelated noise components for each bolometer, and a single common mode noise component were all generated. The maps are generated at 1'' resolution, which is significantly higher than

the maps created for analysis, to reduce effects of pixel noise. The noise components were generated using power spectra derived according to the procedures outlined in Section 4.6. Typical spectra have been shown in Figure 4.7. The diffuse signal components were generated using a power spectrum proportional to k^{-3} , typical for Galactic cirrus emission [68].

Once the maps have been produced, they are re-scanned according to the pointing solution for the IVC G86 field to produce a time-ordered data set. The time-ordered data (TOD) streams can then be re-input into the map maker to try to reconstruct the simulated maps. Simulations are done with the 500 μm array only. At 500 μm , 41 detectors are used to make BLAST maps. Were these simulations to be performed at 250 μm , where we use 132 detectors, simulations would have taken at least an order of magnitude longer to perform. Because the map-making algorithm is linear, signal-only and noise-only simulations can be considered separately.

5.5.2 Results

The simulated TOD streams are fed into the map maker in exactly the same manner as the real data. Specifically, although the simulated data contain no non-astronomical signals which need to be removed, gaps are still removed from the simulated TOD streams which are identical to those removed from the real data. The simulated TOD streams are also pre-filtered according to Section 5.4.4. The map maker is run with fixed noise, using the total noise power spectra ($\mathbf{n}_d = \mathbf{u}_d + \alpha_d \mathbf{c}$, *cf.* Equation 4.3) used to generate the simulated data.

Figure 5.2 shows the results of one of the twenty signal-only simulations. The input map (Figure 5.2.1) contains a large gradient across the whole map which is roughly aligned with galactic latitude. Both the map reconstructed by *almagest* (Figure 5.2.2) and the map reconstructed by SANEPIC (Figure 5.2.3) are missing this low-frequency mode, although the *almagest* map retains more medium-scales than the SANEPIC map.

This SANEPIC map is produced with the full correlated noise treatment described in [39].

The results of the corresponding noise-only simulation is presented in Figure 5.3. Here `almagest` (Figure 5.3.2) does worse than both a simple naïve highpass (Figure 5.3.1), and the full SANEPIC (Figure 5.3.3). Since `almagest` has no explicit correlated noise correction, it is not surprising that a large-scale residual remains in the noise map.

The mean radial power spectra averaged over all twenty simulations are shown in Figure 5.4, for the signal-only simulations, and in Figure 5.5, for the noise-only simulations. Again, it can be seen that, while `almagest` does better than SANEPIC at recovering large-scale modes in the signal map, `almagest`'s large-scale noise rejection performance is significantly poorer. We then calculate the signal-to-noise ratio for both map makers and also their radial transfer function (the ratio of the reconstructed signal map and the simulated input map). These are presented in Figure 5.6.

The signal-to-noise ratios of the two map makers are very close up to $10'$ scales, rising like $k^{3/2}$ as expected in a white-noise case. Above this scale, `almagest`'s signal-to-noise ratio remains roughly constant, while SANEPIC's continues to rise all the way up to $k = (25')^{-1}$. This behaviour is unsurprising: SANEPIC's ability to properly account for correlated noise allows it to whiten the reconstructed noise out to larger scales than is possible with `almagest`, which does not attempt to compensate for correlated noise.

Considering the transfer functions, the situation is reversed. `Almagest` clearly does better at recovering low modes in the signal-only simulations and, as a result, its transfer function does not drop off nearly as much as that of SANEPIC. `Almagest`'s transfer function reaches a maximum of ~ 0.99 around $k = (23')^{-1}$, and remains relatively flat up to $k = (30')^{-1}$, before dropping off at large-scale. Since these transfer function is calculated from two equally pixelised maps, the pixelisation bias exhibited in the signal-only power spectra of Figure 5.4 should not be present here. Nevertheless, `almagest`'s transfer function drops slowly from its maximum towards smaller scales. SANEPIC's




Figure 5.2.1: Simulated input signal map




Figure 5.2.2: Almagest reconstructed map




Figure 5.2.3: SANEPIC reconstructed map

Figure 5.2: One of the twenty signal-only simulations on the IVC G86 field. Figure 5.2.1 shows the input signal map, rebinned to the pixel size of the reconstructed maps. Figure 5.2.2 is the map reconstructed by almagest. Figure 5.2.3 is the map reconstructed by the full SANEPIC, including correlations. Flux units are arbitrary, but consistent across the three maps. Figure 5.2.1 and Figure 5.2.3 correspond to the left and centre panels, respectively, of Figure 12 in [39].




Figure 5.3.1: Optibin reconstructed map




Figure 5.3.2: Almagest reconstructed map

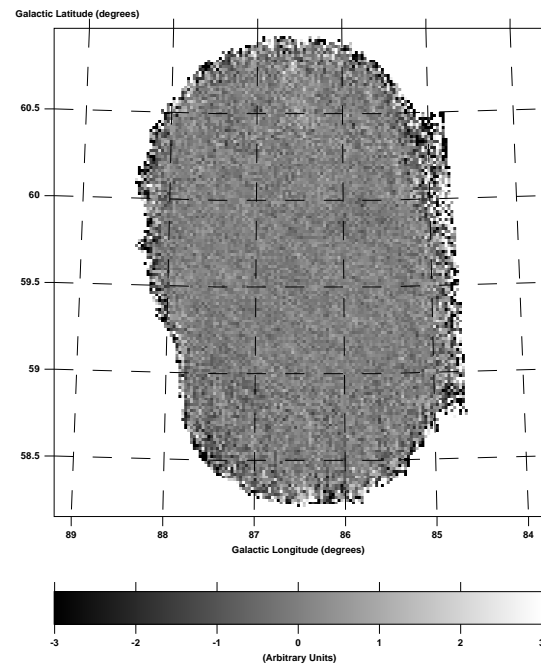


Figure 5.3.3: SANEPIC reconstructed map

Figure 5.3: One of the twenty noise-only simulations on the IVC G86 field. Figure 5.3.1 shows the reconstructed noise map produced by Optibin, after simply highpassing the data at 0.2 Hz. Figure 5.3.2 is the map reconstructed by almagest. Figure 5.3.3 is the map reconstructed by the full SANEPIC, including correlations. Flux units are arbitrary, on the same scale as Figure 5.2. Figure 5.3.3 corresponds to the left panel of Figure 9 in [39].




Figure 5.4: Mean one-dimensional radial power spectra averaged over all twenty signal-only simulations for both almagest (solid curve) and SANEPIC (dashed curve). The mean simulated input power spectrum is also shown (dotted). The spectra have been multiplied by $k^{3/2}$ so that the simulated input has the same power at all scales (*i.e.* no filtering would correspond with a flat line). The fall-off at small angular scales is due to the $1'$ pixelisation of the map.




Figure 5.5: Mean one-dimensional radial power spectra averaged over all twenty noise-only simulations for both almagest (solid curve) and SANEPIC (dashed curve). At small angular scales, both power spectra should be flat, indicating the white-noise level in the maps. However, again, the spectra fall off due to pixelisation.




Figure 5.6: Calculated signal-to-noise ratios (top panel) and transfer functions (bottom panel) calculated from the simulations. Results for both *almagest* (solid curves) and *SANEPIC* (dashed curves) are presented. For both map makers, the signal-to-noise ratios rise like $k^{3/2}$ so long as the reconstructed noise is relatively white. At its greatest, the *SANEPIC* signal-to-noise ratio is three times greater than that of *almagest*. The transfer functions produced with these simulations agree well with the observed performance of the map makers on real data (see Figure 6.4).

transfer function does not show this, and this may be a result of *almagest*'s pixel anti-aliasing filter (Section 5.4.6).

Almagest's large-scale drop-off is due in part to the fourth-order polynomial subtracted from the data during pre-filtering (see Section 5.4.4). Large scales are also suppressed in the map maker due to choosing a naïve map as the initial guess. In the naïve map, these large scale modes are completely missing. The map maker converges logarithmically at large scales from the naïve map to the input data, and the large-scale portion of the transfer function grows as the number of iterations is increased. It is

speculated that the transfer function will eventually rise to a level consistent with only the polynomial subtraction, but this might take thousands of iterations.

SANEPIC's transfer function is relatively flat at small scales up to $k = (15')^{-1}$, above which it drops off. The SANEPIC transfer function has a maximum of ~ 0.91 for the 500 μm band. The SANEPIC bias depends on the particular region being scanned, and is lower in the other bands (3% at 250 μm). It is thought to be due to numerical imprecision in the inversion of N_{pp} [39]. This bias can easily be corrected for through simulations.

To further investigate *almagest*'s transfer function, we plot in Figure 5.7 the full two-dimensional transfer function for the single simulation presented in Figure 5.2. The two-dimensional transfer function for SANEPIC is similar. The prominent line in the two-dimensional transfer function (Figure 5.7.3) corresponds to the primary scan direction of the observation of the field. A much weaker, but similar feature is present along the secondary scan direction. As this plot indicates, the map maker filters large modes most heavily along the scan direction.

A consequence of this is that reconstructed large scale features in maps made by *almagest* and SANEPIC tend to be narrower in the direction of the scan, than perpendicular to it, since the largest modes have been removed along the scan direction. This is most evident in maps which are not cross-linked. A good example of this phenomenon is the *almagest* generated map of Cassiopeia A (Figure 6.13).

The turn-down in the *almagest* transfer function seen in Figure 5.6 is also apparent in the two-dimensional transfer function, where it appears most prominently along the scan direction. This indicates the turn-down may be related to the pixel antialiasing filter which *almagest* applies when generating the maps (Section 5.4.6).

5.5.3 Implications

The above analysis indicates that the two maximum-likelihood map makers are of complementary usefulness for BLAST data analysis. SANEPIC excels in the low signal-to-noise




Figure 5.7.1: Fourier transform of the simulated input
 Figure 5.7.2: Fourier transform of the almagest reconstruction




Figure 5.7.3: Two-dimensional map maker transfer function

Figure 5.7: The two-dimensional almagest transfer function. Figure 5.7.1 and Figure 5.7.2 show the amplitude of the two-dimensional Fourier transforms of Figure 5.2.1 and Figure 5.2.2, respectively, after apodisation to remove edge effects. Figure 5.7.3 presents the resulting two-dimensional transfer function obtained by the division of the upper two plots. Note the parabolic intensity scale on the transfer function plot. Small scale noise in the calculated transfer function is due to misalignment in the pixel gridding of the two input maps.

case. Here it will be able to extract signals embedded in the noise, where *almagest* would be unable to do so. A prime example of this is the analysis done on the (actual) IVC G86 field [67]. This analysis could not have been done with maps generated by *almagest*.

Almagest, on the other hand, does a much better job of retaining large-scale structure in high signal-to-noise maps, typical of the BLAST galactic maps from both flights. In this case, *SANEPIC* will produce maps with lower noise levels, but at the expense of signal attenuation. Accordingly, judicious choice of which BLAST map-maker to use for creation of final maps for a particular field, is required to maximise the science output from the BLAST data.

Καὶ μηδεὶς τὰς τοιαύτας τῶν ὑποθέσεων ἐργώδεις νομισάτω σκοπῶν τὸ τῶν παρ' ἡμῖν ἐπιτεχνημάτων κατασκελές· οὐ γὰρ προσήκει παραβάλλειν τὰ ἀνθρώπινα τοῖς θείοις οὐδὲ τὰς περὶ τῶν τηλικούτων πίστεις ἀπὸ τῶν ἀνομοιοτάτων παραδειγμάτων λαυβάνειν· τί γὰρ ἀνομοιότερον τῶν αἰεὶ καὶ ὡσαύτως ἐχόντων πρὸς τὰ μηδέποτε καὶ τῶν ὑπὸ παντὸς ἀνκωλυθησομένων πρὸς τὰ μηδ' ὑφ' αὐτῶν; ἀλλὰ πειρᾶσθαι μὲν ὡς ἔνι μάλιστα τὰς ἀπλουστέρως τῶν ὑποθέσεων ἐφαρμόζειν ταῖς ἐν τῷ οὐρανῷ κινήσεσιν, εἰ δὲ μὴ τοῦτο προχωροίη, τὰς ἐνδεχομένας.

—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, ΙΓ':β' [1]

Chapter 6

BLAST Maps

We now present a selection of maps of real data created with *almagest* for analysis purposes for both *BLAST05* and *BLAST06*. Also presented are tables of observations for both flights.

6.1 *BLAST05*

During the 100 hour *BLAST05* flight, BLAST made observations of 23 fields, collecting 73 hours of science-quality data. Observations made for *BLAST05* are summarised in Table 6.1. A coverage map of the Galactic surveys is presented as Figure 6.1. Maps for selected fields are presented in the following sections.

6.1.1 Cygnus X

The Galactic plane survey in Cygnus X is presented as Figure 6.2 and Figure 6.3. At ~ 10 square degrees, this is the largest area map made from *BLAST05* data. Three bright point-like objects dominate the map. This is the W 75N/DR21 region [71, 72] resolved as three point-sources, designated W 75N ($\ell = 81^{\circ}87$ $b = +0^{\circ}78$), and the overlapping pair DR21 ($\ell = 81^{\circ}72$ $b = +0^{\circ}57$), and DR21(OH) ($\ell = 81^{\circ}67$ $b = +0^{\circ}54$). These three




Figure 6.1: Coverage map showing selected scientific targets surveyed during *BLAST05*. The background intensity image is $100\ \mu\text{m}$ emission from [69].

<i>Field Name</i>	<i>Centre (J2000.0)</i>		<i>Area (deg²)</i>	<i>Time</i>	<i>Cit.</i>	<i>Observation Type</i>
	<i>R.A.</i>	<i>Dec.</i>				
<i>Cygnus X</i>	20 ^h 36 ^m 1	+41°32'	10.13	12 ^h 08 ^m	*	Galactic plane survey
ELAIS-N1	16 ^h 14 ^m 9	+55°51'	0.63	8 ^h 24 ^m		Extragalactic blank field
ELAIS-N1 HVC A	16 ^h 15 ^m 6	+55°26'	0.79	8 ^h 18 ^m		High velocity cloud
<i>Vulpecula</i>	19 ^h 40 ^m 9	+23°26'	3.23	6 ^h 23 ^m	[11]	Galactic plane survey
G86.0+38.3	14 ^h 40 ^m 7	+49°04'	3.87	6 ^h 12 ^m	[67]	Intermediate velocity cloud
<i>Aquila</i>	19 ^h 15 ^m 7	+11°17'	5.49	6 ^h 04 ^m	*	Galactic plane survey
<i>CRL 2688</i>	21 ^h 02 ^m 6	+36°43'	0.53	5 ^h 06 ^m	[9]	Proto-planetary nebula
Draco IVC	16 ^h 46 ^m 6	+60°08'	4.22	3 ^h 10 ^m		Intermediate velocity cloud
<i>Arp 220</i>	15 ^h 34 ^m 8	+23°30'	0.48	2 ^h 58 ^m	[9]	ULIRG
IRAS 22134+5834	22 ^h 14 ^m 6	+58°47'	1.62	2 ^h 41 ^m	[9]	High-mass proto-star
2 Pallas	12 ^h 22 ^m 8	+20°07'	0.41	2 ^h 08 ^m	[9]	Asteroid
<i>IC 5146</i>	21 ^h 46 ^m 5	+47°34'	0.77	1 ^h 28 ^m	*	Star-forming region
IRAS 21307+5049	21 ^h 32 ^m 9	+51°05'	0.33	1 ^h 17 ^m	[9]	High-mass proto-star
LDN 1014	21 ^h 24 ^m 4	+50°02'	0.27	1 ^h 01 ^m	[9]	Proto-stellar core
IRAS 20126+4101	20 ^h 14 ^m 5	+41°16'	0.44	1 ^h 00 ^m	[9]	High-mass proto-star
IRAS 23011+6126	23 ^h 02 ^m 5	+61°41'	0.36	0 ^h 55 ^m	[9]	Young protostar
<i>Sagitta</i>	19 ^h 29 ^m 1	+17°11'	2.91	0 ^h 50 ^m		Galactic plane survey
<i>NGC 4565</i>	12 ^h 36 ^m 6	+26°00'	0.36	0 ^h 49 ^m	*	Edge-on spiral galaxy
IRAS 21078+5211	21 ^h 09 ^m 1	+52°21'	0.42	0 ^h 41 ^m	[9]	Compact HII region
Mrk 231	12 ^h 56 ^m 3	+56°52'	0.32	0 ^h 26 ^m	[9]	ULIRG
<i>K3-50</i>	20 ^h 02 ^m 0	+33°32'	0.34	0 ^h 24 ^m	[9]	Compact HII region
<i>Cassiopeia A</i>	23 ^h 23 ^m 8	+58°40'	0.60	0 ^h 20 ^m	[70]	Supernova remnant
Saturn	7 ^h 51 ^m 3	+21°09'	0.36	0 ^h 08 ^m	[9]	Planet

Table 6.1: Fields observed during the *BLAST05* flight ordered by quantity of data acquired. Fields which are discussed below have italicised field names. Columns two and three list the equatorial co-ordinates of the centre of the map. Column four lists the area of the final map, which is typically larger than the targeted area. Column five lists the total length of observations used to produce the map of the field. This number does not include poor quality data discarded during the reduction process before the map making stage. Column six gives citations to publications analysing the given BLAST field. An asterisk in this column indicates that such a publication is in preparation. Column seven gives an indication of the type of observation made.

sources are the brightest objects observed during *BLAST05*, with per-pixel luminosities around 4 GJy/sr, 7 GJy/sr, and 6 GJy/sr, respectively at 250 μm , and similarly high values in the other bands. Flux densities and spectral energy density (SED) fits for these three point-like sources are presented in [9].

The high intensities of these sources cause ringing in the time-domain filtering for both BLAST map makers, which induces artefacts. Most noticeable is the large scale depression in flux (a *shadow*) along one scan direction, above and to the right of these sources, but smaller scale striping along the scan directions also occurs. To remove this effect, a fit to the sources could be removed from the TOD streams before map making, and then re-added to the map after map making is complete.

Subtracting the 250 μm map of Cygnus X generated by SANEPIC from the same map generated by *almagest* (Figure 6.2) results in the difference map presented as Figure 6.4. There is good agreement between the behaviour seen here and the transfer functions calculated as a result of simulations run on the map makers (Figure 5.6). The maps agree well at medium scales. At large scales, where the map is very poorly constrained, the maps differ the most. These large scale fluctuations across the maps are unconstrained modes introduced by the map making process. The maps also differ on small scales near bright sources, a result of the extended PSF, and *almagest*'s pixel antialiasing filter.

6.1.2 Vulpecula

The Galactic plane survey in Vulpecula is presented as Figure 6.5. This field includes the open cluster NGC 6823 [5]. These BLAST maps have been extensively analysed in [11, 28, 30]. All three colours have a noticeable artificial (*i.e.* map-maker induced) depression in flux near $\ell = 58^{\circ}5$, $b = +0^{\circ}9$. This is at the end of one of the passes of this field, where the TODs are poorly constrained. The brightest object in the map has a peak per-pixel luminosity of ~ 1 GJy/sr at 250 μm , significantly lower than the brightest sources in the Cygnus X maps (Section 6.1.1). The map is correspondingly much better

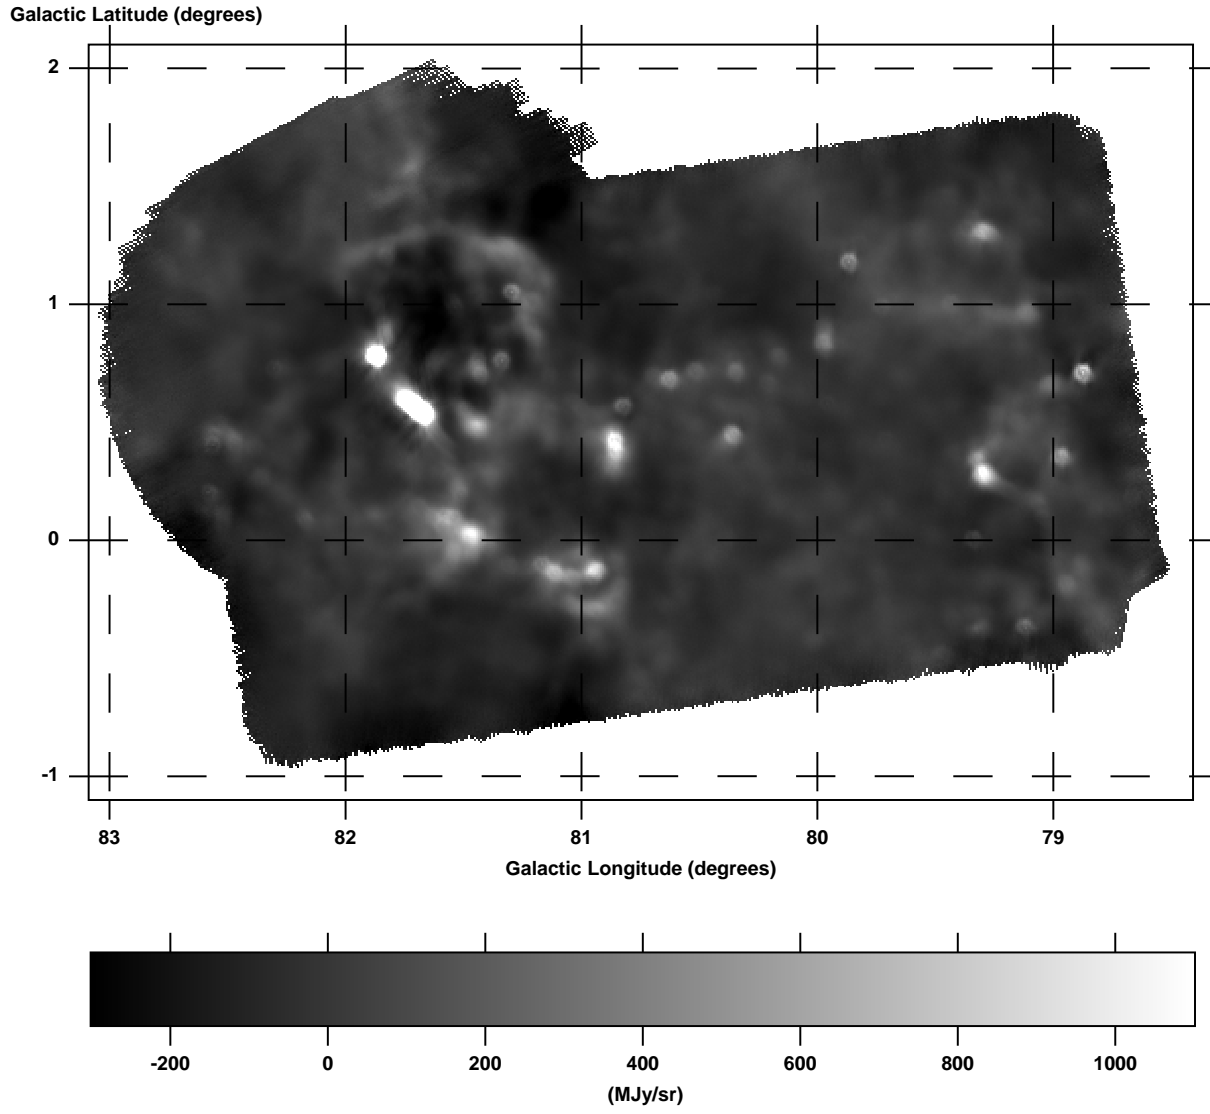


Figure 6.2: The ~ 10 square degree Galactic plane survey in Cygnus X at 250 μm . Pixel size is 15". The region is composed of a well cross-linked rectangular area with a shallower circular area superimposed on it. Outside of the overlap, the circular region is poorly cross-linked. The 350 μm and 500 μm maps of this field are presented in Figure 6.3.



Figure 6.3.1: 350 μm mapFigure 6.3.2: 500 μm map

Figure 6.3: The ~ 10 square degree Galactic plane survey in Cygnus X at 350 μm and 500 μm . Pixel size is $15''$. The 250 μm map of this region has been presented as Figure 6.2.




Figure 6.4: The difference map between the 250 μm Cygnus X map generated by almagest (Figure 6.2) and the same map generated by SANEPIC. Difference maps for the other two bands are similar.

behaved in the neighbourhood of this source, although the depression in flux along the bottom edge of the map is due to shadowing induced by the row of sources above it.

A map of the per-pixel variance for the $250\ \mu\text{m}$ band is presented in Figure 6.5.4. The variance is a strong function of the coverage. The step in variance between the left and right halves of the map is the result of one pass observing only the right-hand side of the field, due to a operator error when setting the parameters for that scan. As a result, the right-hand side of the map has deeper coverage.

Bright point sources also show up strongly in the variance map. This results from observing point sources with the non-symmetric *BLAST05* point-spread function (PSF; cf. Figure 3.5). Observations of the same source with different crossing angles results in rotated copies of the PSF being stacked on top of each other in the map. Since the PSF is not symmetric, the resultant areas of high and low flux do not correlate, increasing the per-pixel variance in the map around these point-sources.

6.1.3 Aquila

The Galactic plane survey in Aquila is presented as Figure 6.6. This field contains the Galactic Ring Survey [73] molecular cloud GRSMC 45.46+0.06 [74, 75], which is the brightest object in the BLAST map. Only two visits were made to this field, both at similar elevations ($\sim 32^\circ$), resulting in a cross-linking angle of only $\sim 18^\circ$. As a result, cross-linking is worse, as is depth. Significant map-maker induced shadowing occurs along the scan direction above GRSMC 45.46+0.06, which has a per-pixel luminosity of $\sim 2\ \text{GJy/sr}$ at $250\ \mu\text{m}$, due to the poorly constrained data.

6.1.4 Sagitta

The Galactic plane survey in Sagitta is presented as Figure 6.7. Designed to be a 6 square degree survey of the plane similar to that of Cygnus X, Aquila, and Vulpecula, the observation was aborted after completing roughly half of the first pass of the field, and








Figure 6.5.1: 250 μm mapFigure 6.5.2: 350 μm mapFigure 6.5.3: 500 μm mapFigure 6.5.4: 250 μm pixel variance map

Figure 6.5: The ~ 3.2 square degree Galactic plane survey in Vulpecula. Pixel size is $15''$. Note the logarithmic intensity scale on the 250 μm per-pixel variance map presented in Figure 6.5.4.

Figure 6.6.1: 250 μm mapFigure 6.6.2: 350 μm mapFigure 6.6.3: 500 μm mapFigure 6.6: The ~ 5.5 square degree Galactic plane survey in Aquila. Pixel size is $15''$

not subsequently re-observed. The resulting map, therefore, is not crosslinked, and poorly covered. At $25''$ resolution, the $500\ \mu\text{m}$ map, observed with the fewest detectors, has holes in the centre of the map, which results in significant map-making artefacts in the cross-scan direction (see Figure 6.7.3). For meaningful science to be extracted from this field, the map would have to be degraded to larger pixel size.

The large amount of structure in this map, despite being only a single pass across the field, indicates that, while the optical damage during *BLAST05* (Section 3.3.2) significantly affected point-source sensitivity, sensitivity of the instrument to diffuse structures was still quite high.

6.1.5 CRL 2688

The BLAST maps of the proto-planetary nebula CRL 2688 [44], also known as “The Egg” are presented as Figure 6.8. This calibrator was observed numerous times throughout the flight, and was used to flatfield the bolometers (Section 4.7.1). The map presented combines nine separate observations of the field at two primary cross-linking angles, corresponding to rising and setting observations. The BLAST measurements of CRL 2688 are discussed in detail in [9].

In Figure 6.8.4 we present a map of the ratio of samples flagged as bad, and, thus, discarded, to the total number of hits in a pixel. The map shows an average flagged sample rate of $\sim 0.5\%$, which rises to $\sim 4\%$ on the actual source. This is an indication that the flagging process (Section 4.2) is erroneously considering some observations of CRL 2688 to be bad data, and removing them. This appears to be a problem for this source in particular, as other fields with point sources do not exhibit this behaviour.

6.1.6 Arp 220

The BLAST observations of the ultra-luminous infrared galaxy (ULIRG) Arp 220 [47] is presented as Figure 6.9. Arp 220 was used as the *BLAST05* primary flux calibrator





Figure 6.7.1: 250 μm mapFigure 6.7.2: 350 μm mapFigure 6.7.3: 500 μm mapFigure 6.7.4: 500 μm coverage map

Figure 6.7: The ~ 3 square degree Galactic plane survey in Sagitta. Pixel size is $25''$. Map coverage is increasingly poor at longer wavelengths. The 500 μm map (Figure 6.7.3) has holes in the map, and the poor coverage induces cross-scan gradients. Note, for instance, the discontinuity in flux near the source at $\ell = 52^\circ 9'$, $b = -0^\circ 7'$. Figure 6.7.4 plots the number of observations per pixel for the 500 μm map and clearly shows the unequal coverage of the map.





Figure 6.8.1: 250 μm mapFigure 6.8.2: 350 μm mapFigure 6.8.3: 500 μm mapFigure 6.8.4: 250 μm flagged samples map

Figure 6.8: The proto-planetary nebula CRL 2688. The map is 0.53 square degrees, and has a pixel size of $15''$. Note the logarithmic intensity scale on the flagged samples per pixel map, in Figure 6.8.4.

(Section 4.7.2) and was observed several times throughout the flight. This map combines four separate visits of the source, all with similar elevations, resulting in only minor cross-linking. The scan pattern is typical of the *BLAST05* isolated point-source maps: a rectangle aligned in azimuth and elevation roughly half a degree wide and a quarter degree high. The BLAST measurements of Arp 220 are discussed in detail in [9].

Due to the small cross-linking angle, there is a cross-scan gradient in the maps for all three colours, which is most pronounced in the 500 μm map. Arp 220 is a relatively dim source for BLAST, and at 500 μm , the intensity of the cross-scan gradient is on the same order as the luminosity of the galaxy itself.

6.1.7 IC 5146

The BLAST maps of the low mass star forming region near the open cluster IC 5146 [3], or Cocoon Nebula, are presented as Figure 6.10. A single pass was made of the field, which is not cross-linked as a result. The 500 μm map shows striping due to the non-uniform coverage of the field. The odd shape of the top and bottom corners of the map is due to the minimum azimuthal scan requirements of the pointing system.

The brightest object is the submillimetric source associated with the FU Orionis-type protostar V1357 Cygni [76]. It has a per-pixel luminosity of ~ 100 MJy/sr at 250 μm , and is shadowed heavily by the map maker to the left, along the scan direction.

6.1.8 K3-50

The BLAST maps of the compact HII region K3-50 [77, 78], also called W 58 [71], are presented as Figure 6.11. BLAST resolves two components, K3-50 A ($\alpha = 20^{\text{h}}01^{\text{m}}46^{\text{s}}$, $\delta = +33^{\circ}32'42''$) and K3-50 C ($\alpha = 20^{\text{h}}01^{\text{m}}54^{\text{s}}$, $\delta = +33^{\circ}34'15''$), with per-pixel luminosities of ~ 2 GJy/sr and ~ 1.5 GJy/sr, respectively. Flux densities and spectral energy density (SED) fits for these two sources are presented in [9].

There is a single observation of this field. Standard map-making procedures produce



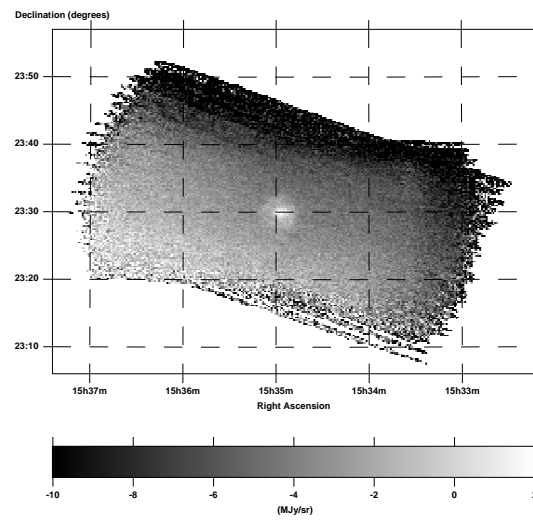
Figure 6.9.1: 250 μm mapFigure 6.9.2: 350 μm mapFigure 6.9.3: 500 μm map

Figure 6.9: The ultra-luminous infrared galaxy Arp 220. The map is 0.48 square degrees, and has a pixel size of $15''$




Figure 6.10.1: 250 μm mapFigure 6.10.2: 350 μm mapFigure 6.10.3: 500 μm map

Figure 6.10: The low-mass star forming region near IC-5146. The map size is 0.77 square degrees and has a pixel size of $15''$




Figure 6.11.1: 250 μm unconstrained map




Figure 6.11.2: 250 μm constrained map




Figure 6.11.3: 350 μm unconstrained map




Figure 6.11.4: 350 μm constrained map




Figure 6.11.5: 500 μm unconstrained map




Figure 6.11.6: 500 μm constrained map

Figure 6.11: The compact HII region K3-50. The map is 0.34 square degrees, and has a pixel size of $15''$. Maps on the right are constrained versions of the corresponding map on the left.

strong shadowing on both sides of these sources along the scan direction (left column, Figure 6.11). This behaviour is also seen in other, similar small maps of bright, isolated sources with no cross-linking. These shadows can be removed by the introduction of an artificial constraint. In the maps on the right of Figure 6.11 a constraint has been applied during the map making stage which zeroes all pixels outside a circle of radius 7.5 centred on the sources. Almagest applies this constraint when the map is rescanned (Equation 5.9.1), not when the map is binned. As a result, the constraint itself is not obvious in the maps produced. Small ripples aligned with the scan direction are visible away from the source in the constrained maps. These are a result of the non-uniform coverage, due to having only one observation of this field.

6.1.9 NGC 4565

The BLAST observation of the edge-on spiral galaxy NGC 4565 [5] is presented as Figure 6.12. These maps have been analysed in [28]. The surface brightness of the galaxy is quite low in the BLAST bands. Only one observation of the field, with no cross-linking is available. To reduce striping in the scan direction, the same 7.5 radius artificial constraint has been applied during the map making procedure that was used for the maps of K3-50 (Section 6.1.8).

6.1.10 Cassiopeia A

The BLAST map of the Cassiopeia A [79] supernova remnant is presented as Figure 6.13. Only a single pass was made of this field which is, therefore, not cross-linked. The 500 μm map shows significant striping along the scan direction as well as a pronounced cross-scan gradient. In the shorter wavelength bands, where the coverage is more uniform, the maps are better behaved. Addition of an artificial constraint similar to that used for the K3-50 maps (Section 6.1.8) could reduce these stripes.

A single point source is detected in the BLAST bands. This is IRAS 23239+5826




Figure 6.12.1: 250 μm mapFigure 6.12.2: 350 μm mapFigure 6.12.3: 500 μm map

Figure 6.12: The edge-on spiral galaxy NGC 4565. The map is 0.36 square degrees, and has a pixel size of $15''$

[7] ($\alpha = 23^{\text{h}}26^{\text{m}}15^{\text{s}}$; $\delta = 58^{\circ}42'49''$). The remainder of the map is low intensity diffuse structure. Analysis of these BLAST maps has been done in [70].

6.2 *BLAST06*

During the 286 hour *BLAST06* flight, BLAST made observations of 37 fields, collecting ~ 225 hours of science-quality data. A preliminary list of observations made for *BLAST06* are summarised in Table 6.2. A coverage map is presented as Figure 6.14. Maps for selected fields are presented in the following sections.

6.2.1 CDFS Deep

Maps of the CDFS Deep extragalactic BLAST field are presented in Figure 6.15. This observation overlaps the Chandra Deep Field South (CDFS) [80], also known as GOODS South [81]. There are relatively few stars available to the star cameras in this region. As a result, the star cameras typically were only able to acquire pointing solutions during turn-around on the left-hand side of this field. This results in several pointing excursions to the right of the map, which should be removed from the final maps.

6.2.2 CDFS Wide

Maps of the CDFS Wide extragalactic BLAST field are presented in Figure 6.16. These maps include all the data for the embedded CDFS Deep field (Section 6.2.1), totalling ~ 80 hours of observations, more than the entire *BLAST05* usable data set. Streaks in the maps near $\alpha = 3^{\text{h}}35^{\text{m}}$ $\delta = -29^{\circ}20'$ and $\alpha = 3^{\text{h}}32^{\text{m}}$ $\delta = -26^{\circ}50'$ are due to bad data which should be removed to make the final maps.




Figure 6.13.1: 250 μm mapFigure 6.13.2: 350 μm mapFigure 6.13.3: 500 μm map

Figure 6.13: The supernova remnant Cassiopeia A. The map is 0.60 square degrees, and has a pixel size of $15''$

<i>Field Name</i>	<i>Centre (J2000.0)</i>		<i>Area</i> (deg ²)	<i>Time</i>	<i>Observation Type</i>
	<i>R.A.</i>	<i>Dec.</i>			
SEP	4 ^h 43 ^m :7	−53°30′	12.00	67 ^h 04 ^m	Extragalactic blank field
<i>CDFS Deep</i>	3 ^h 32 ^m :7	−27°49′	1.69	43 ^h 38 ^m	Extragalactic blank field
<i>CDFS Wide</i>	3 ^h 32 ^m :7	−27°49′	12.98	35 ^h 54 ^m	Extragalactic blank field
<i>Vela Deep</i>	8 ^h 58 ^m :2	−45°26′	55.43	27 ^h 42 ^m	Galactic plane survey
<i>Vela Wide</i>	9 ^h 17 ^m :8	−47°15′	207.89	10 ^h 19 ^m	Galactic plane survey
Bullet	6 ^h 58 ^m :5	−55°57′	0.36	9 ^h 48 ^m	$z = 0.296$ cluster
<i>VY Canis Majoris</i>	7 ^h 23 ^m :0	−25°46′	0.56	4 ^h 15 ^m	Hypergiant star
PKS B0529−549	5 ^h 30 ^m :5	−54°54′	0.54	3 ^h 37 ^m	$z = 2.575$ radio galaxy
<i>IRAS 08470−4243</i>	8 ^h 48 ^m :7	−42°54′	0.36	3 ^h 37 ^m	Compact HII region
Abel 3112	3 ^h 17 ^m :6	−44°19′	1.17	3 ^h 28 ^m	$z = 0.075$ cluster
<i>NGC 1808</i>	5 ^h 07 ^m :7	−37°31′	1.09	2 ^h 59 ^m	Resolved Seyfert galaxy
Puppis A	8 ^h 23 ^m :3	−43°00′	4.24	2 ^h 20 ^m	Supernova remnant
<i>NGC 1291</i>	3 ^h 17 ^m :3	−41°07′	0.45	1 ^h 32 ^m	Resolved lenticular galaxy
<i>η Carinae</i>	10 ^h 44 ^m :6	−59°42′	4.47	1 ^h 30 ^m	Low-mass star forming region
IRAS 08485−4419	8 ^h 50 ^m :4	−44°30′	0.38	0 ^h 54 ^m	Compact HII region
<i>NGC 1097</i>	2 ^h 46 ^m :3	−30°15′	0.31	0 ^h 52 ^m	Resolved Seyfert galaxy
<i>NGC 1566</i>	4 ^h 20 ^m :0	−54°56′	0.76	0 ^h 50 ^m	Resolved spiral galaxy
IRAS 09149−4743	9 ^h 16 ^m :6	−47°56′	0.29	0 ^h 48 ^m	Compact HII region
Gum Glob	8 ^h 06 ^m :9	−35°58′	0.92	0 ^h 45 ^m	Cometary globules
G284.35−0.42	10 ^h 24 ^m :1	−57°53′	0.84	0 ^h 31 ^m	Supernova remnant
<i>NGC 1365</i>	3 ^h 33 ^m :6	−36°08′	0.29	0 ^h 30 ^m	Resolved Seyfert galaxy
<i>NGC 1512</i>	4 ^h 03 ^m :9	−43°21′	0.52	0 ^h 28 ^m	Resolved spiral galaxy
PKS B1138−262	11 ^h 40 ^m :8	−26°29′	0.73	0 ^h 17 ^m	$z = 2.156$ radio galaxy
IRAS 08337−4028	8 ^h 35 ^m :5	−40°39′	0.22	0 ^h 13 ^m	Compact HII region
IRAS 08140−3559	8 ^h 15 ^m :9	−36°09′	0.28	0 ^h 12 ^m	Compact HII region
IRAS 09018−4816	9 ^h 03 ^m :5	−48°28′	0.45	0 ^h 11 ^m	Compact HII region
IRAS 08211−4158	8 ^h 22 ^m :9	−42°08′	0.29	0 ^h 06 ^m	Intermediate mass proto-star
IRAS 08476−4306	8 ^h 49 ^m :5	−43°17′	0.27	0 ^h 05 ^m	Intermediate mass proto-star
IRAS 08307−4303	8 ^h 32 ^m :5	−43°13′	0.27	0 ^h 05 ^m	Unidentified IRAS source
IRAS 09014−4736	9 ^h 03 ^m :1	−47°48′	0.27	0 ^h 05 ^m	Intermediate mass proto-star
IRAS 08247−4223	8 ^h 26 ^m :4	−42°33′	0.27	0 ^h 05 ^m	Compact HII region
IRAS 09002−4732	9 ^h 01 ^m :9	−47°44′	0.29	0 ^h 04 ^m	Compact HII region

Table 6.2: Preliminary list of fields observed during the *BLAST06* flight ordered by quantity of data acquired. Columns are the same as for Table 6.1 except that the column of citations is omitted.




Figure 6.15.1: 250 μm mapFigure 6.15.2: 350 μm mapFigure 6.15.3: 500 μm map

Figure 6.15: The ~ 1.7 square degree extragalactic blank-field overlapping the CDFS. Pixel size is $10''$. These preliminary maps contain data acquired during pointing excursions which result in noise to the right of the circular field itself.




Figure 6.16.1: 250 μm map




Figure 6.16.2: 350 μm map




Figure 6.16.3: 500 μm map

Figure 6.16: The ~13 square degree extragalactic blank-field overlapping the CDFS. Pixel size is 10''

6.2.3 Vela

The $\sim 55^\circ$ Galactic plane survey in Vela is presented as Figure 6.17 and Figure 6.18. The primary Galactic survey in *BLAST06*, Vela Deep contains the Vela Molecular Ridge (VMR) [82]. IRAS point sources embedded in this map which were also mapped separately, and so included in Table 6.2, include:

- IRAS 08337–4028
- IRAS 08470–4243
- IRAS 08476–4306
- IRAS 08485–4419
- IRAS 09002–4732
- IRAS 09014–4736
- IRAS 09018–4816
- IRAS 09149–4743

Data from these targeted observations, totalling 3^h6 for IRAS 08470–4243, 0^h9 for IRAS 08485–5519, and 1^h4 for all others together, are included in the maps presented.

As with CDFS Deep (Section 6.2.1), the Vela Deep field is embedded within the $\sim 208^\circ$ Vela Wide field. The overlap of these two fields is shown in Figure 6.19. Vela Wide was designed primarily as a way to constrain larger scale modes in the Vela Deep field. As such, it was scanned with an azimuthal velocity of 0^o4/s, rather than the typical BLAST mapping speed of 0^o1/s (see Section 3.1.1). This has consequences both for the fidelity of the pointing solution, which is degraded due to the lack of reliable mid-scan star detections by the star cameras, as well as widening of the submillimetric beam itself along the scan direction. Because of these issues, and the field’s large size, it is currently




Figure 6.17: The ~ 55 square degree Galactic plane survey in Vela. Pixel size is $10''$. The $350\ \mu\text{m}$ and $500\ \mu\text{m}$ maps of this field are presented in Figure 6.18.



Figure 6.18.1: 350 μm mapFigure 6.18.2: 500 μm map

Figure 6.18: The ~ 55 square degree Galactic plane survey in Vela. Pixel size is $10''$. The 250 μm map of this field has been presented in Figure 6.17.




Figure 6.19: The $250\ \mu\text{m}$ Vela Deep map (Figure 6.17) and the $250\ \mu\text{m}$ map of Puppis A (Figure 6.20.1), shown with an ellipse overdrawn which roughly corresponds to the area mapped in the $\sim 208^\circ$ Vela Wide survey. The rightmost corner of Vela Deep is outside the larger field, as is all of Puppis A.

unclear whether the available map makers will be able to produce high quality maps of Vela Wide which can be used for analysis.

6.2.4 Puppis A

The ~ 4 square degree BLAST observation of the Puppis A [83] supernova remnant is presented as Figure 6.20. The location of this field in relation to the Vela Deep survey can be seen in Figure 6.19. Two observations of this field were performed, relatively closely in time, so the map is essentially not cross-linked. Nevertheless, almagest does a good job

of recovering the structure. The brightest two sources in the map are IRAS 08211–4158 ($\ell = 259^{\circ}7584$, $b = -2^{\circ}3379$) and IRAS 08247–4223 [7] ($\ell = 260^{\circ}4843$, $b = -2^{\circ}5361$). Both these sources were also targeted individually, and appear in Table 6.2.

Due to the small number of observations made of this field, coverage is quite variable. The $500\ \mu\text{m}$ map contains some mid-map holes and suffers increased striping along the scan direction due to the non-uniform coverage. This problem can be mitigated by generating maps with larger pixel size.

6.2.5 η Carinae

The ~ 4.5 square degree BLAST observation of the η Carinae nebula, also known as the Keyhole Nebula, or NGC 3372 [5], are presented as Figure 6.21. The star η Carinae itself, along with the associated Homunculus Nebula, is the small, isolated point source near the centre of the map. The observation consists of two largely non-overlapping observations. Despite not being cross-linked, the structure is well reconstructed by almagest, although with several dark bands along the scan direction.

6.2.6 VY Canis Majoris

Observations of the hypergiant star VY Canis Majoris [84, 85] are presented as Figure 6.22. VY Canis Majoris was one of the primary calibrators for *BLAST06*, and was observed many times during the flight. The maps presented contain all $4^{\text{h}}25$ of data. Although the cross-linking angle is small, as is typical for *BLAST06*, the map is well constrained with a slight gradient from bottom to top. The source VY Canis Majoris itself is well resolved. The streak heading off the map to the top is due to a pointing glitch.







Figure 6.20.1: 250 μm mapFigure 6.20.2: 350 μm mapFigure 6.20.3: 500 μm map

Figure 6.20: The ~ 4 square degree survey of the Puppis A supernova remnant. Pixel size is $10''$. For the position of this field relative to the Vela Deep field, see Figure 6.19.

Figure 6.21.1: 250 μm mapFigure 6.21.2: 350 μm mapFigure 6.21.3: 500 μm mapFigure 6.21: The ~ 4.5 square degree BLAST survey around η Carinae. Pixel size is $10''$




Figure 6.22.1: 250 μm mapFigure 6.22.2: 350 μm mapFigure 6.22.3: 500 μm map

Figure 6.22: The hypergiant star VY Canis Majoris. Pixel size is $10''$. Note the small side lobes below the source.

6.2.7 IRAS 08470–4243

Beam maps created from a single fifteen-minute-long observation of the compact HII region IRAS 08470–4243 [7] are presented as Figure 6.23. This source served as one of two primary calibrators for *BLAST06*, and was observed many times throughout the flight. The *BLAST06* side lobes are apparent in the maps provided. Extended features to the left and right of the source are map maker induced shadows. The maps are plotted in “telescope co-ordinates”: elevation and cross-elevation relative to some arbitrary point on the sky. Since *BLAST06* was flown near the South Pole, telescope co-ordinate maps are rotated $\sim 180^\circ$ from equatorial co-ordinate maps.

6.2.8 Nearby Resolved Galaxies

Six nearby galaxies were observed by BLAST during *BLAST06*. These are presented as Figure 6.24, Figure 6.25, and Figure 6.26, along with optical counterparts.

6.2.8.1 NGC 1097

The Seyfert galaxy NGC 1097 [5] (Figure 6.24.1) is well detected in all BLAST bands. The centre is quite bright (~ 400 MJy/sr at $250 \mu\text{m}$), which causes some map-maker induced streaking in the $250 \mu\text{m}$ map.

6.2.8.2 NGC 1291

The lenticular galaxy NGC 1291 [5] (Figure 6.24.2) is noted for its inner spiral and outer ring. BLAST faintly detects both the central core of the galaxy and the ring in all bands, although the maps presented suffer from stripes which run right through the galaxy itself and need to be cleaned up.




Figure 6.23.1: 250 μm mapFigure 6.23.2: 350 μm mapFigure 6.23.3: 500 μm map

Figure 6.23: Beam map produced from one fifteen minute long observation of the compact HII region IRAS 08470–4243. Pixel size is $10''$. The map is plotted in “telescope co-ordinates,” so the images appear flipped from the similar beam maps in Figure 6.22. Note the small sidelobes above the source.




Figure 6.24.1: NGC 1097

Figure 6.24.2: NGC 1291

Figure 6.24: NGC 1097 and NGC 1291 as observed by BLAST. BLAST pixel size is $10''$. The BLAST maps are arranged in order of increasing wavelength, with the $250 \mu\text{m}$ map at top. Optical counterparts shown at bottom are, for NGC 1097, a blue (IIIa-J emission) and, for NGC 1291, a red (IIIa-F emission) images taken from the ESO/SRC Southern Sky Survey [86].

6.2.8.3 NGC 1365

The structure of the barred spiral galaxy NGC 1365 (Figure 6.25.1 [5]) is well resolved by BLAST. NGC 1365 is the brightest of the nearby galaxies observed: its central core has a per-pixel luminosity of ~ 1 GJy/sr at $250 \mu\text{m}$.

6.2.8.4 NGC 1512

The dim NGC 1512 [5] (Figure 6.25.2) is weakly detected by BLAST. Cross-scan gradients in the $350 \mu\text{m}$ and $500 \mu\text{m}$ maps degrade these colours. The companion galaxy, NGC 1510, (lower right in the optical image) also appears weakly detected by BLAST at $250 \mu\text{m}$.

6.2.8.5 NGC 1566

NGC 1566 [5] (Figure 6.26.1) is detected in all three bands. All bands show map-making artefacts, both shadowing and cross-scan gradients.

6.2.8.6 NGC 1808

The Seyfert galaxy NGC 1808 [5] (Figure 6.26.2) is one of the brightest nearby galaxies observed BLAST (per-pixel luminosity of ~ 725 MJys/sr at $250 \mu\text{m}$.) Of all the nearby galaxies, it was observed for the longest period of time (~ 3 hours) and the resulting maps are well defined, and relatively free of map-making artefacts. Also detected in this field, but not shown in Figure 6.26.2, are the interacting galaxies LEDA 16801 and 16804 [4] ($\alpha = 5^{\text{h}}08^{\text{m}}5$, $\delta = -37^{\circ}39'5$).




Figure 6.25.1: NGC 1365

Figure 6.25.2: NGC 1512

Figure 6.25: NGC 1365 and NGC 1512 as observed by BLAST. BLAST pixel size is $10''$. The BLAST maps are arranged in order of increasing wavelength, with the $250 \mu\text{m}$ map at top. Optical counterparts shown at bottom are, for NGC 1365 a near-infrared (IV-N emulsion) taken from the ESO/SRC Southern Sky Survey [86] and, for NGC 1512, a red (IIIa-F emulsion through an ON590 filter) image taken from the Second Epoch Survey [87].




Figure 6.26.1: NGC 1566

Figure 6.26.2: NGC 1808

Figure 6.26: NGC 1566 and NGC 1808 as observed by BLAST. BLAST pixel size is $10''$. The BLAST maps are arranged in order of increasing wavelength, with the $250\ \mu\text{m}$ map at top. Optical counterparts shown at bottom are red (IIIa-F emulsion) images taken from, for NGC 1566, the Second Epoch Survey [87] and, for NGC 1808, the ESO/SRC Southern Sky Survey [86].

βλάβει τι τὸ προκείμενον, οὔτε τὰ ἀναποδείκτως ὑποτιθέμενα, ἐὰν ἅπαξ σύμφωνα τοῖς φαινομένοις καταλαμβάνηται, χωρὶς ὁδοῦ τινος καὶ ἐπιστάσεως εὐρῆσθαι δύναται, κἂν δυσέκθετος ἦ ὁ τρόπος αὐτῶν τῆς καταλήψεως, ἐπειδὴ καὶ καθόλου τῶν πρώτων ἀρχῶν ἢ οὐδὲν ἢ δυσερμήνευτον φύσει τὸ αἴτιον,
—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Θ':β' [1]

Chapter 7

Summary

BLAST is a two-metre submillimetric telescope able to make simultaneous observations at 250 μm , 350 μm , and 500 μm . BLAST was built and has been flown on two scientific flights. During the first scientific flight, *BLAST05*, BLAST suffered significant degradation of the optical system. Despite this, it obtained 73 hours of useful scientific data during this flight. The majority of data acquired during *BLAST05* are observations of star forming regions within the Galaxy. Instrumental problems encountered during the first scientific flight were corrected, and BLAST's second scientific flight, *BLAST06*, was a resounding success. During *BLAST06*, BLAST acquired 225 hours of useful scientific data, primarily observations of extragalactic targets.

To analyse the data acquired by BLAST during these two flights, a full analysis pipeline was created to reduce the raw data acquired by the telescope with the ultimate goal of producing scientific maps of the fields observed. The pipeline cleaned and calibrated the data, and generated a pointing solution for the telescope. An iterative, maximum-likelihood map maker was created to optimally generate maps from the data products produced by the reduction pipeline.

Maps for all fields observed by BLAST during the two scientific flights have been produced. Analysis of the data acquired in both flights is underway. A first series

of papers, based on targets observed during *BLAST05*, is in press. Data acquired by BLAST will satisfy the scientific objectives set out for the project.

Although BLAST was largely destroyed after the *BLAST06* landing, the irreplaceable bolometric detector arrays were recovered. Plans are underway to rebuild BLAST, and add to it a polariser. Rebuilt, BLAST should be able to provide polarised submillimetric observations, like no other current experiment can.

Προσαναπληρωθέντων οὖν καὶ τῶν τοιούτων, ὦ Σύρε, καὶ σχεδὸν πάντων κατ' ἑμὸν γε νοῦν ἐφωδευμένων τῶν εἰς τὴν τοιαύτην σύνταξιν ὀφειλόντων θεωρηθῆναι, καθ' ὅσον ὃ τε μέχρι τοῦ δεῦρο χρόνος πρὸς εὕρεσιν ἢ ἐπανόρθωσιν ἀκριβεστέραν συνήργει, καὶ ὁ πρὸς τὸ εὐχρηστον μόνον τῆς θεωρίας, ἀλλ' οὐ πρὸς ἔνδειξιν, ὑπομνηματισμὸς ὑπέβαλλεν, οἰκεῖον ἂν ἡμῖν ἐνταῦθα καὶ σύμμετρον εἰλήφοι τὸ τέλος ἢ παροῦσα πραγματεία.

—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, ΠΓ':α' [1]

Appendix A

Watchdog Circuit Schematics

The watchdog schematic is presented below, broken up into pieces roughly corresponding to the blocks of Figure 2.7. Some passive components have been suppressed from these schematics for brevity.

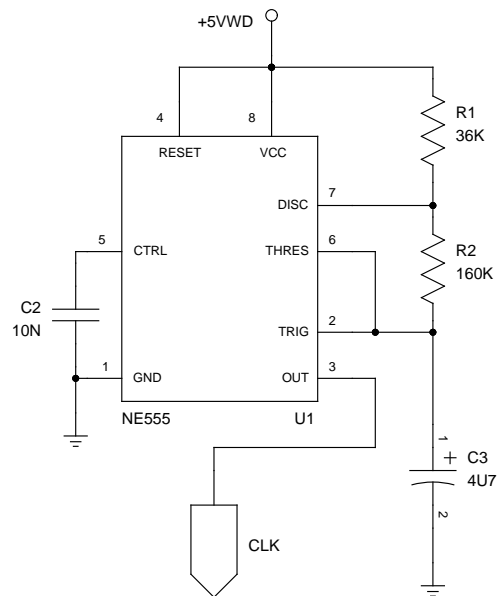


Figure A.1: The watchdog clock. U1 is a 555 operating in astable mode. The clock period is given by: $p = 0.693 \frac{s}{F\Omega} (R1 + 2R2) C3$. With $R1 = 36 \text{ k}\Omega$, $R2 = 160 \text{ k}\Omega$, and $C3 = 4.7 \mu\text{F}$, this gives a clock period of $p = 1.162 \text{ s}$.

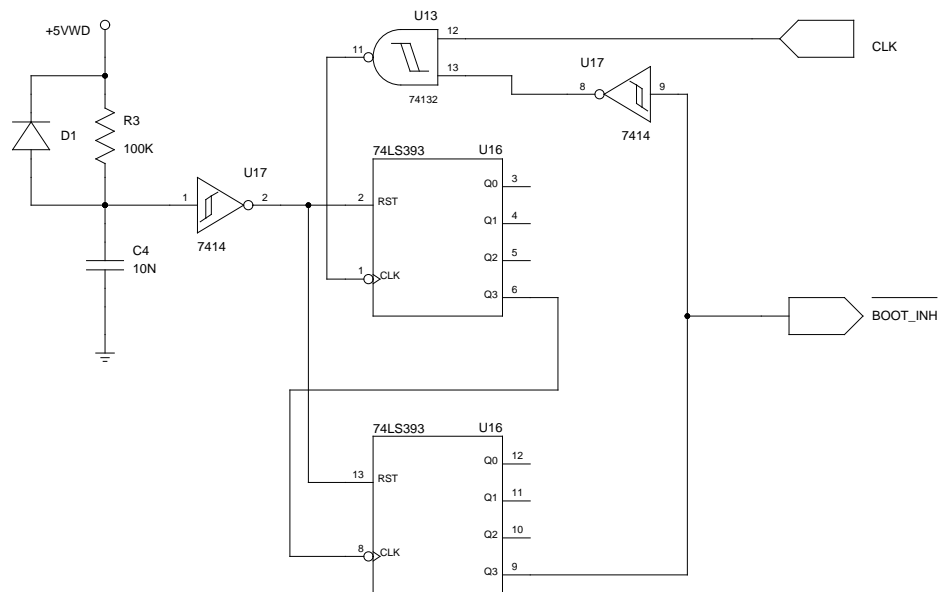


Figure A.2: The watchdog boot inhibit. U16 is an 8-bit counter. The RC-circuit R3–C4 holds the reset lines of U16 high for the first 3 s after the watchdog unit is powered. Once the reset line of U16 goes low, U16 begins counting up. After 256 clock cycles (4.96 minutes), pin 9 of U16 goes high, raising the /BOOT_INH and inhibiting the clock via U13, holding the circuit high for the rest of operation.

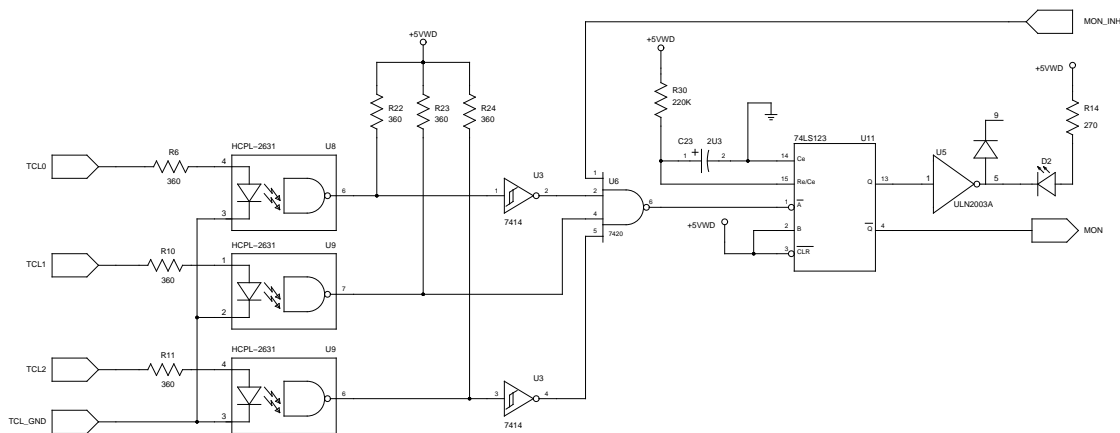


Figure A.3: The watchdog MCC monitor. There are two identical copies of this circuit, one for each computer. U11 is a one-shot retriggerable monostable multivibrator with a 486 ms timeout. The computer tickle is input at left on lines from the parallel port. The four port NAND gate U6 requires the tickle to alternate between low–high–low and any other state to continuously retrigger U11 to prevent it from timing out. If the tickle stops and U11 times out, MON goes high, signalling the countdown timer to begin.

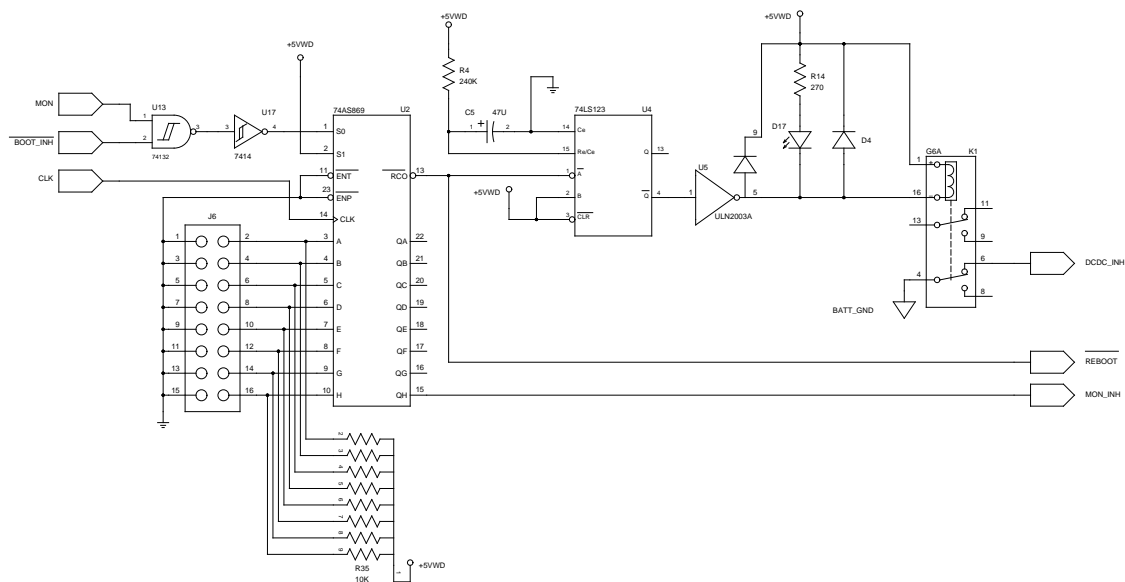


Figure A.4: The watchdog countdown timer. Again, one of two. Once MON goes low, U2 begins counting down. The /BOOT_INH line causes MON to be ignored until the boot inhibit times out. The initial value of the counter is set by J6. For BLAST, all pins of J6 were open, providing the shortest countdown delay, *i.e.* single clock cycle. If the MCC tickles the watchdog while the counter is active, MON will be lowered again, and the counter will be reset. When the counter times out, /RCO goes low. The multivibrator U4 is triggered on the falling edge of /RCO. After being triggered, U4 generates a 3.72 second pulse which temporarily inhibit's the MCC's DC-DC via relay K1. /RCO is also sent to the control select circuit as /REBOOT. The high bit of counter U2 is also fed back to the monitor circuit as MON_INH. This inhibits the monitor for 4.96 minutes, allowing time for the computer to reboot.

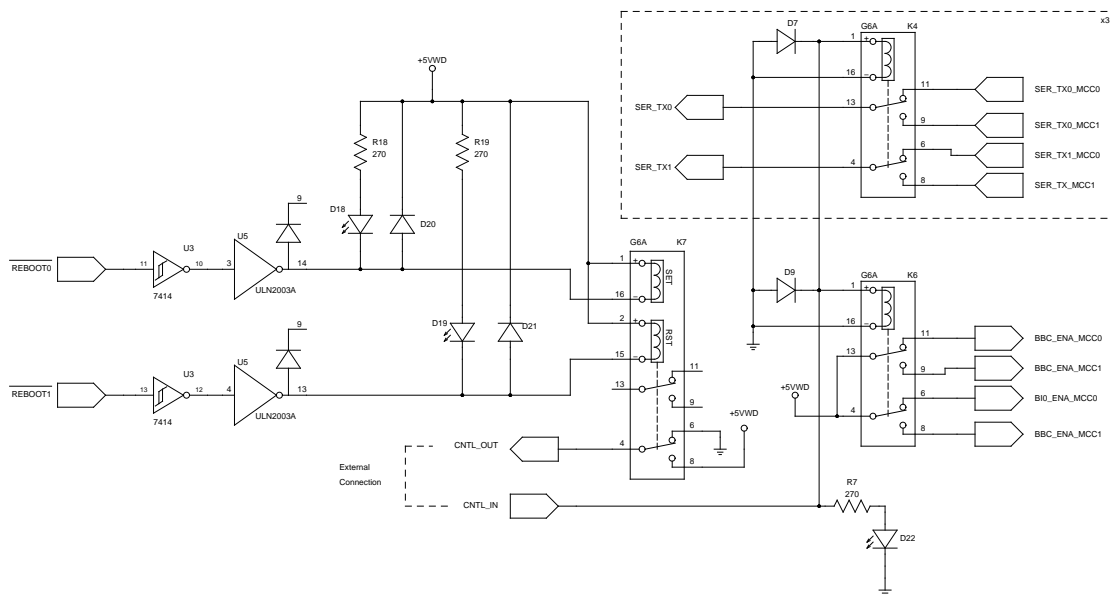


Figure A.5: The watchdog control select. The $/\text{REBOOT}$ lines from the two MCC count down timers are fed into the circuit at left. When one of the computers is rebooted, the corresponding $/\text{REBOOT}$ line goes low, and the relay K7 selects the other coil. K7 connects CNTL_OUT to either +5 V or ground, depending on which computer has control. CNTL_OUT is connected externally to CNTL_IN , permitting the watchdog circuit to be overridden in the lab. The in-charge computer is selected by enabling its BLASTbus Controller via relay K6, and the transmit lines on the six serial ports via three additional relays, only one of which is shown (K4).

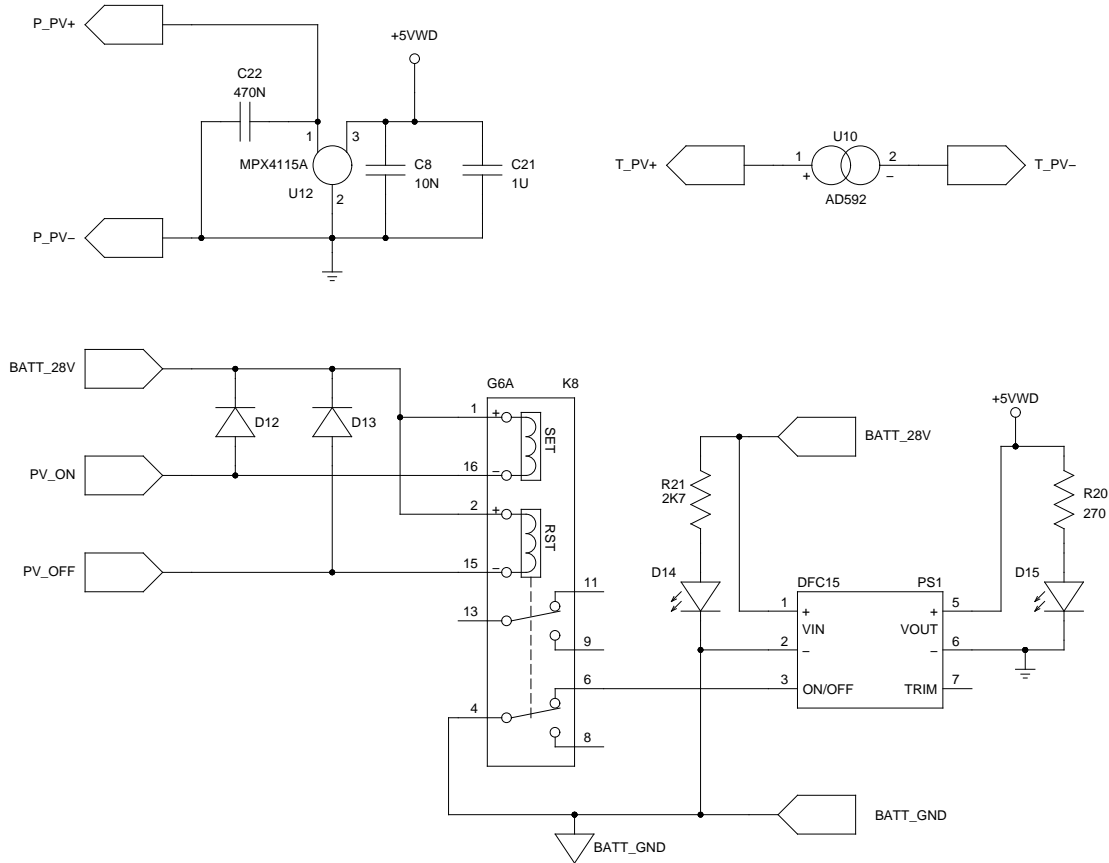


Figure A.6: Other watchdog components. The watchdog circuit contains a 28V→5V DC–DC converter (PS1) which provides power to the watchdog circuit itself, and is also used as the isolation voltage on the BLASTbus. Discrete commands from the science stack (Section 2.13.4.1) switch relay K8 which enables or disables the watchdog DC–DC, and as a result, the other DC–DCs in the pressure vessel. The ambient pressure and temperature sensor, which monitor the PV environment, are also shown. The temperature sensor, U10, an AD592, is completely isolated from the watchdog circuit and included for convenience only. The pressure sensor, U12, is powered by the 5 V generated by the watchdog. Both sensors are read out by the ACS.

Ταῦτα μὲν οὖν αἰεὶ τηροῦντες ἐπὶ τῆς τοιαύτης καταγραφῆς, καθ' ὃν ἂν βουλώμεθα τρόπον,
τοὺς αὐτοὺς λόγους ἐπὶ τῶν ἀριθμῶν εὐρήσομεν φερομένους,
—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Ι':ζ' [1]

Appendix B

Glossary

ADC Card:

A multipurpose data acquisition card designed for BLAST, but now used by other experiments as well. The card provides digitisation of twenty-five analogue inputs, a DSP for data processing, and several digital I/O lines for control. Almost all systems on BLAST are interfaced through the ADC Cards. ADC Cards communicate with the flight computers over the BLASTbus (*q.v.*).

Aluminised Mylar[®]:

Mylar[®] is a bi-axially oriented polyethylene terephthalate (PET) thin polyester film. It is coated with aluminum and used on BLAST to provide insulation in various contexts.

Aperture Efficiency:

The ratio of the amount of power from a point source diffraction pattern that couples to a detector centred on a source, to the total power from the diffraction pattern.

Attitude Control System (ACS):

An electronics crate on the outer frame containing four ADC Cards (*q.v.*), a power distribution module, a break-out panel, and other associated electronics. The ACS is responsible for pointing readout, motion control, and housekeeping.

Biphase:

A digital encoding used to transfer data without the need for a common ground reference. High bits are encoded as rising edges and low bits are encoded as falling edges. Accordingly, the biphase clock is twice as fast as the data clock, and transitions may be necessary between data edges to encode the data.

BLASTbus:

A custom RS-485 synchronous serial bus used by the ADC Cards (*q.v.*) to communicate with the flight computers. Each BLASTbus has a maximum throughput of 1 Mbit/s. The BLAST platform is equipped with two BLASTbuses.

BLASTbus Controller (BBC):

A PCI card installed in the flight computers to manage both the BLASTbus (*q.v.*), and also create the biphase-encoded data stream to be fed to the LOS Transmitters (*q.v.*).

Bosma Mirror:

A 1.8 meter, parabolic primary mirror for BLAST constructed of aluminum. It was built for the *BLAST03* test flight. After destruction of the COI mirror (*q.v.*), it was resurfaced and used again during the *BLAST06* science flight.

Cal-Lamp:

A calibration source located in the hole in the Lyot stop (M4). It is used to measure both the relative response between detectors, and the variance of a

single bolometer over time.

COI Mirror:

A 2 meter, spherical primary mirror used for BLAST constructed of carbon fibre. This mirror was used for the *BLAST05* science flight, and was destroyed upon landing after this flight. COI refers to Composite Optics Incorporated, the firm who produced it. The other BLAST flights used the Bosma mirror (*q.v.*).

Conic Constant:

A constant, K , characterising a class of radially symmetric surfaces defined by:

$$z(r) = \frac{\frac{r^2}{R_c}}{1 + \sqrt{1 - (1 - K)\frac{r^2}{R_c^2}}} \quad (\text{B.1})$$

where R_c is the radius of curvature. Such surfaces may be classified as follows:

$K > 0$	Oblate ellipsoid
$K = 0$	Sphere
$-1 < K < 0$	Prolate ellipsoid
$K = -1$	Paraboloid
$K < -1$	Hyperboloid

Dirfile:

A database format defined originally for BLAST but now used in other contexts as well. A dirfile is organised around fixed-rate records. It consists of a directory containing binary files, one per field in the record. A text file in the directory contains metadata describing the database. Dirfiles are used in BLAST data analysis, created by converting the framefiles (*q.v.*) generated by the flight software with the `defile` program.

Faraday Rotation:

The phenomenon whereby the plane of polarisation of light rotates through interaction with a magnetic field in a dielectric. Also known as the Faraday Effect or the magneto-optical effect. The Earth's magnetic field induces errors in measurements made by fibre-optic gyroscopes due to Faraday rotation. This phenomenon was first observed by Faraday on 13 September, 1845, and was one of the early indications of the relationship between light and electromagnetism [88].

Fourier Transform Spectrometer (FTS):

A Michelson interferometer with a moveable mirror. The interference pattern created by moving the mirror is the Fourier transform of the spectral response of the source.

Framefile:

A simple database format generated by the flight software to store the flight data. A framefile consists of fixed-rate records stored sequentially in 15 minute long files. Framefiles aren't used in analysis; rather, they are converted to dirfiles (*q.v.*) by the `defile` program.

High-Rate Downlink:

The 6 kbit/s telemetry channel transmitted through the TDRSS network (*q.v.*).

Iridium:

A satellite telephone network. BLAST is equipped with a satellite telephone and packet modem which provide bi-directional communication with the ground. Iridium is used as a back-up for the TDRSS link (*q.v.*).

Linear Variable Differential Transformer (LVDT):

A precision device for measuring displacement. It consists of three windings encircling a hollow bore and positioned one next to the other along the bore's length. The centre winding of the three is known as the primary winding and the two on either side are the secondary windings. A transducer rod with high magnetic permeability known as the core is paced inside the bore.

To measure the displacement between two objects, the piece housing the windings is attached to one object and rod attached to the core is fixed to the second. The core is allowed free movement along the axis of the bore. The primary winding is energised with an AC signal. Since the core is free to move, induced current in the two secondary windings will vary according to how much of the rod is encircled by the particular winding. The displacement is found by measuring the difference of the induced current in the two secondary windings.

LOS Transmitters:

A set of three L- or S-band transmitters used to transmit from the gondola to a listening station on the ground at the launch site. One transmitter is used for a 1 Mbit/s biphase encoded data stream. The other two transmitters are used for NTSC compliant signals from the star cameras. Since the LOS transmitters are a significant power user, they are switched off when the gondola is not within line-of-site of the launch site.

Master Control Program (mcp):

A monolithic program which runs on the flight computers and handles all aspects of gondola operation.

μ -Metal:

Any of a number of nickel-iron alloys exhibiting high permeability and used for magnetic shielding.

RS-485:

A two-wire electrical (*i.e.* physical layer) standard for half-duplex serial communication allowing multiple nodes on a single bus. The differential voltage between the two wires conveys the information: one voltage polarity indicating a high bit and the other polarity indicating a low bit.

Sagnac Effect:

The phenomenon whereby light travelling around a rotating circular path will see a different path length whether it is rotating with or against the rotation of the path itself [89]. This phenomenon is used in fibre-optic gyroscopes to measure rotational rates.

Schedule File:

A text file containing a precomputed flight schedule which is uploaded to the flight computers before launch. The file consists of a list of commands to be executed, primarily pointing commands, indexed by gondola local sidereal time (LST).

Strehl Ratio:

A measure of optical efficiency. The Strehl ratio is the ratio of the peak intensity of the beam to the peak intensity of the nominal beam [90].

Tracking and Data Relay Satellite System (TDRSS):

A NASA satellite system consisting of nine geosynchronous satellites, three of which are active at any given time. The primary role of the TDRSS network is

to provide communication between low-orbiting satellites and the ground, primarily the space shuttle and International Space Station. The satellite system was inaugurated in 1983, to replace the collection of ground based listening posts used previously by NASA. For BLAST, TDRSS is used for the high-rate downlink (*q.v.*). The link from the ground to BLAST via TDRSS is known as the *TDRSS forward*. The reverse, from BLAST to the ground, is known as the *TDRSS return*.

Ὅλως δὲ ἡγούμεθα προσήκειν δι' ἀπλουστέρων ὡς ἔνι μάλιστα ὑποθέσεων τὰ φαινόμενα ἀποδεικνύειν, ἐφ' ὅσον ἂν μηδὲν ἀξιόλογον ἐκ τῶν τηρήσεων ἀντιπίπτον τῇ τοιαύτῃ προθέσει φαίνεται.

—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Γ':α' [1]

Appendix C

Excerpts from Ptolemy's *Almagest*

All translations are taken from [91].

Page v:

“The true philosophers, Syrus, were, I think, quite right to distinguish the theoretical part of philosophy from the practical. For even if practical philosophy, before it is practical, turns out to be theoretical, nevertheless one can see that there is a great difference between the two: in the first place, it is possible for many people to possess some of the moral virtues even without being taught, whereas it is impossible to achieve theoretical understanding of the universe without instruction;” (I:i)

Page 12:

“From all this we concluded: that the first two divisions of theoretical philosophy should rather be called guesswork than knowledge, theology because of its completely invisible and ungraspable nature, physics because of the unstable and unclear nature of matter; hence there is no hope that philosophers will ever be agreed about them.”
(I:i)

Page 97:

“We think that the above suffices as an indication of the methods in this type of

theoretical investigation, enough [at least] so it cannot be said that we have neglected this topic.” (VIII:vi)

Page 128:

“While it would be absurd not to recognise any of these effects¹, on the other hand, when one considers the resulting complications in the methods necessary to deal with each of them, deliberate neglect of effects small enough to be overlooked both in theory and observation evokes [in the reader] a strong feeling of the advantage of greater simplicity, and no regret, or little, for the resulting error in representing the phenomena.” (VI:vii)

Page 150:

“For those who approach this science in a true spirit of enquiry and love of truth ought to use any new methods they discover, which give more accurate results, to correct not merely the ancient theories, but their own too, if they need it. They should not think it disgraceful, when the goal they profess to pursue is so great and divine, even if their theories are corrected and made more accurate by others beside themselves.” (IV:ix)

Page 178:

“Now let no one, considering the complicated nature of our devices, judge such hypotheses to be over-elaborated. For it is not appropriate to compare human [constructions] with divine, nor to form one's beliefs about such great things on the basis of very dissimilar analogies. Rather, one should try, as far as possible, to fit the simpler hypotheses to the heavenly motions, but if this does not succeed, [one should apply hypotheses] which do fit.” (XIII:ii)

Page 218:

“...and we know too that assumptions made without proof, provided only that they are found to be in agreement with the phenomena, could not have been found

¹That eclipse maxima don't occur exactly at opposition or conjunction

without some careful methodological procedure, even if it is difficult to explain how one came to conceive them (for, in general, the cause of first principles is, by nature either non-existent or hard to describe);” (IX:ii)

Page 220:

“We have now completed these additional topics, Syrus, and have shown the way to deal with almost all the topics which should, at least to my mind, be subjected to theory for the purposes of this kind of treatise, at any rate as far as the time up to our own days contributed to greater accuracy in our discoveries or in corrections [of earlier discoveries], and as far as was suggested by a memorandum directed only toward scientific usefulness, and not towards ostentation. So at this point our present discussion can be terminated at an appropriate place at the right length.” (XIII:xi)

Page 225:

“If we always observe the above rules when drawing this type of figure, we will find that the same numerical ratios result however we decide to draw it.” (X:vii)

Page 233:

“And in general, we consider it a good principle to explain the phenomena by the simplest hypothesis possible, in so far as there is nothing in the observations to provide a significant objection to such a procedure.” (III:i)

Page 238:

“...but as for assertions of validity ‘for eternity’, or even for a length of time which is many times that over which the observations have been taken, we must consider such as alien to a love of science and truth.” (III:i)

Page 251:

“Rather, we should not judge ‘simplicity’ in heavenly thing from what appears to be simple on earth, especially when the same thing is not equally simple for all even here.” (XIII:ii)

τὰς δὲ περὶ ὅλου τοῦ αἰῶνος ἢ καὶ τοῦ μακρῶ τινι πολλαπλασίου τοῦ κατὰ τὰς τηρήσεις
χρόνου διαβεβαιώσεις ἀλλοτρίας φιλομαθείας τε καὶ φιλαληθείας ἠγάμεθα.

—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, Γ':α' [1]

Bibliography

- [1] C. Ptolemy. “Μαθηματικῆς Συντάξεως.” In J. L. Heiberg, editor, *Claudii Ptolemaei opera quae existant omnia*, volumes 1–2. B. G. Teubner, Leipzig, 1898, 1903.
- [2] J. L. E. Dreyer. “Index Catalogue of Nebulae Found in the Years 1888 to 1894, with Notes and Corrections to the New General Catalogue.” *Memoirs of the Royal Astronomical Society*, **51**:185–228, 1895.
- [3] J. L. E. Dreyer. “Second Index Catalogue of Nebulae Found in the Years 1895 to 1907; with Notes and Corrections to the New General Catalogue and to the Index Catalogue for 1888 to 1894.” *Memoirs of the Royal Astronomical Society*, **59**:105–198, 1908.
- [4] G. Paturel, P. Fouque, L. Bottinelli, L. Gouguenheim. “An extragalactic database. I - The Catalogue of Principal Galaxies.” *Astronomy & Astrophysics Supplements Series*, **80**:299–315, 1989.
- [5] J. L. E. Dreyer. “A New General Catalogue of Nebulae and Clusters of Stars, being the Catalogue of the late Sir John F. W. Herschel, Bart., revised, corrected, and enlarged.” *Memoirs of the Royal Astronomical Society*, **49**:1–237, 1888.
- [6] A. Wright, R. Otrupcek. *Parkes Catalog*. Australia Telescope National Facility, 1990.
- [7] C. A. Beichman, G. Neugebauer, H. J. Habing, P. E. Clegg, T. J. Chester. “Ex-

- planatory supplement." In *Infrared astronomical satellite (IRAS) catalogs and atlases*, volume 1, 1988.
- [8] K. Dobashi, J.-P. Bernard, Y. Yonekura, Y. Fukui. "Molecular clouds in Cugnus. 1: A large-scale (13)CO survey." *The Astrophysical Journal Supplement Series*, **95**:419–456, 1994.
- [9] M. D. P. Truch, P. A. R. Ade, J. J. Bock, E. L. Chapin, M. J. Devlin, S. Dicker, E. Pascale, M. Griffin, J. O. Gundersen, M. Halpern, P. C. Hargrave, D. H. Hughes, J. Klein, C. J. MacTavish, G. Marsden, P. G. Martin, T. G. Martin, P. Mauskopf, C. B. Netterfield, L. Olmi, G. Patanchon, M. Rex, D. Scott, C. Semisch, N. Thomas, C. Tucker, G. S. Tucker, M. P. Viero, D. V. Wiebe. "The Balloon-borne Large Aperture Submillimeter Telescope (BLAST) 2005: Calibration and Targeted Sources." *The Astrophysical Journal*, *in press*, 2007.
- [10] R. Chini, D. Ward-Thompson, J. M. Kirk, M. Nielbock, B. Reipurth, A. Sievers. "Mm/Submm images of Herbig-Haro energy sources and candidate protostars." *Astronomy & Astrophysics*, **369**:155–169, 2001.
- [11] E. L. Chapin, P. A. R. Ade, J. J. Bock, C. Brunt, M. J. Devlin, S. Dicker, M. Griffin, J. O. Gundersen, M. Halpern, P. C. Hargrave, D. H. Hughes, J. Klein, G. Marsden, P. G. Martin, P. Mauskopf, C. B. Netterfield, L. Olmi, E. Pascale, G. Patanchon, M. Rex, D. Scott, C. Semisch, M. D. P. Truch, C. Tucker, G. S. Tucker, M. P. Viero, D. V. Wiebe. "The Balloon-borne Large Aperture Submillimeter Telescope (BLAST) 2005: A 4 deg² Galactic Plane Survey in Vulpecula ($\ell = 59^\circ$)." *The Astrophysical Journal*, *in press*, 2008.
- [12] G. H. Rieke, E. T. Young, C. W. Engelbracht, D. M. Kelly, F. J. Low, E. E. Haller, J. W. Beeman, K. D. Gordon, J. A. Stansberry, K. A. Misselt, J. Cadien, J. E. Morrison, G. Rivlis, W. B. Latter, A. Noriega-Crespo, D. L. Padgett, K. R. Stapelfeldt,

- D. C. Hines, E. Egami, J. Muzerolle, A. Alonso-Herrero, M. Blaylock, H. Dole, J. L. Hinz, E. Le Floch, C. Papovich, P. G. Pérez-González, P. S. Smith, K. Y. L. Su, L. Bennett, D. T. Frayer, D. Henderson, N. Lu, F. Masci, M. Pesenson, L. Rebull, J. Rho, J. Keene, S. Stolovy, S. Wachter, W. Wheaton, M. W. Werner, P. L. Richards. “The Multiband Imaging Photometer for Spitzer (MIPS)” *The Astrophysical Journal Supplement Series*, **154**:25–29, 2004.
- [13] J. D. Mill, R. R. O’Neil, S. Price, G. J. Romick, O. M. Uy, E. M. Gaposchkin, G. C. Light, W. W. J. Moore, T. L. Murdock, A. T. Stair, Jr. “Midcourse space experiment: Introduction to the spacecraft, instruments, and scientific objectives” *Journal of Spacecraft and Rockets*, **31**:900–907, 1994.
- [14] E. Hubble. “Effects of Red Shifts on the Distribution of Nebulae” *The Astrophysical Journal*, **84**:517–554, 1936.
- [15] A. W. Blain, M. S. Longair. “Submillimetre cosmology” *Monthly Notices of the Royal Astronomical Society*, **264**:509–512, 1993.
- [16] D. W. Hogg, I. K. Baldry, M. R. Blanton, D. J. Eisenstein. “The K correction” *arXiv:astro-ph/0210394v1*, 2002.
- [17] W. S. Holland, E. I. Robson, W. K. Gear, C. R. Cunningham, J. F. Lightfoot, T. Jenness, R. J. Ivison, J. A. Stevens, P. A. R. Ade, M. J. Griffin, W. D. Duncan, J. A. Murphy, D. A. Naylor. “SCUBA: a common-user submillimetre camera operating on the James Clerk Maxwell Telescope” *Monthly Notices of the Royal Astronomical Society*, **303**:659–672, 1999.
- [18] W. D. Duncan, E. I. Robson, P. A. R. Ade, M. J. Griffin, G. Sandell. “A millimetre/submillimetre common user photometer for the James Clerk Maxwell Telescope” *Monthly Notices of the Royal Astronomical Society*, **243**:126–132, 1990.

- [19] T. R. Hunter, D. J. Benford, E. Serabyn. “Optical Design of the Submillimeter High Angular Resolution Camera (SHARC)?” *Publications of the Astronomical Society of the Pacific*, **108**:1042–1050, November 1996.
- [20] J. Glenn, J. J. Bock, G. Chattopadhyay, S. F. Edgington, A. Lange, J. Zmuidzinas, P. D. Mauskopf, B. Rownd, L. Yuen, P. A. Ade. “Bolocam: a millimeter-wave bolometric camera.” In Phillips [21].
- [21] T. G. Phillips, editor. *Advanced Technology MMW, Radio and Terahertz Telescopes*, volume 3357 of *Proceedings of the SPIE*, March 1998.
- [22] C. D. Dowell, C. A. Allen, R. S. Babu, M. M. Freund, M. Gardner, J. Groseth, M. D. Jhabvala, A. Kovacs, D. C. Lis, S. H., Jr. Moseley, T. G. Phillips, R. F. Silverberg, G. M. Voellmer, H. Yoshida. “SHARC II: a Caltech submillimeter observatory facility camera with 384 pixels.” In T. G. Phillips, J. Zmuidzinas, editors, *Millimeter and Submillimeter Detectors for Astronomy*, volume 4855 of *Proceedings of the SPIE*, 2003.
- [23] E. Kreysa, H.-P. Gemuend, J. Gromke, C. G. Haslam, L. Reichertz, E. E. Haller, J. W. Beeman, V. Hansen, A. Sievers, Zylka R. “”Bolometer array development at the Max-Planck-Institut für Radioastronomie”” In Phillips [21].
- [24] E. Pascale, P. A. R. Ade, J. J. Bock, E. L. Chapin, J. Chung, M. J. Devlin, S. Dicker, M. Griffin, J. O. Gundersen, M. Halpern, P. C. Hargrave, D. H. Hughes, J. Klein, C. J. MacTavish, G. Marsden, P. G. Martin, T. G. Martin, P. Mauskopf, C. B. Netterfield, L. Olmi, G. Patanchon, M. Rex, D. Scott, C. Semisch, N. Thomas, M. D. P. Truch, C. Tucker, G. S. Tucker, M. P. Viero, D. V. Wiebe. “The Balloon-borne Large Aperture Submillimeter Telescope: BLAST?” *The Astrophysical Journal*, *in press*, 2007.

- [25] J. J. Bock, J. Glenn, S. M. Grannan, K. D. Irwin, A. E. Lange, H. G. Leduc, A. D. Turner. “Silicon nitride micromesh bolometer arrays for SPIRE.” In Phillips [21].
- [26] M. J. Griffin, J. J. Bock, W. K. Gear. “Relative Performances of Filled and Feedhorn-Coupled Focal-Plane Architectures.” *Applied Optics*, **41**(31):6543–6554, 2002.
- [27] W. K. Gear, C. R. Cunningham. “SCUBA: A submillimetre camera.” In B. Kaldeich, editor, *From Ground-Based to Space-Borne Sub-mm Astronomy*, volume 29 of *Liège International Astrophysics Colloquia*, pages 353–358, December 1990.
- [28] G. Marsden. *Observing the Galactic Plane with the Balloon-borne Large-Aperture Submillimeter Telescope*. PhD thesis, University of British Columbia, July 2007.
- [29] B. P. Crill, P. A. R. Ade, D. R. Artusa, R. S. Bhatia, J. J. Bock, A. Boscaleri, P. Cardoni, S. E. Church, K. Coble, P. de Bernardis, G. de Troia, P. Farese, K. M. Ganga, M. Giacometti, C. V. Haynes, E. Hivon, V. V. Hristov, A. Iacoangeli, W. C. Jones, A. E. Lange, L. Martinis, S. Masi, P. V. Mason, P. D. Mauskopf, L. Miglio, T. Montroy, C. B. Netterfield, C. G. Paine, E. Pascale, F. Piacentini, G. Polenta, F. Pongetti, G. Romeo, J. E. Ruhl, F. Scaramuzzi, D. Sforna, A. D. Turner. “BOOMERANG: A Balloon-borne Millimeter-Wave Telescope and Total Power Receiver for Mapping Anisotropy in the Cosmic Microwave Background.” *The Astrophysical Journal Supplement Series*, **148**:527–541, 2003.
- [30] M. Rex. *The Balloon-borne Large Aperture Submillimeter Telescope (BLAST)*. PhD thesis, University of Pennsylvania, 2007.
- [31] J. C. Mather. “Bolometers: ultimate sensitivity, optimization, and amplifier coupling.” *Applied Optics*, **23**(4):584–588, 1984.
- [32] M. D. P. Truch. *The Balloon-borne Large Aperture Submillimeter Telescope*. PhD thesis, Brown University, May 2007.

- [33] E. Pascale. *The Balloon-borne Large Aperture Submillimeter Telescope: BLAST*. PhD thesis, University of Toronto, September 2007.
- [34] E. Oberg, F. D. Jones, H. L. Horton, H. H. Ryffel. *Machinery's Handbook*. Industrial Press, 27th edition, 2004.
- [35] P. A. R. Ade, G. Pisano, C. Tucker, S. Weaver. "A review of metal mesh filters." In Zmuidzinas et al. [36].
- [36] J. Zmuidzinas, W. S. Holland, S. Withington, W. D. Duncan, editors. *Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III*, volume 6275 of *Proceedings of the SPIE*, May 2006.
- [37] Columbia Scientific Balloon Facility. <http://www.csbf.nasa.gov/>, 2005, 2007.
- [38] *Secure Hash Standard*, Federal Information Processing Standard (FIPS) 180-2. United States of America, National Institute of Science and Technology, August 2002.
- [39] G. Patanchon, P. A. R. Ade, J. J. Bock, E. L. Chapin, J. Chung, M. J. Devlin, S. Dicker, M. Griffin, J. O. Gundersen, M. Halpern, P. C. Hargrave, D. H. Hughes, J. Klein, C. J. MacTavish, G. Marsden, P. G. Martin, T. G. Martin, P. Mauskopf, C. B. Netterfield, L. Olmi, E. Pascale, M. Rex, D. Scott, C. Semisch, N. Thomas, M. D. P. Truch, C. Tucker, G. S. Tucker, M. P. Viero, D. V. Wiebe. "SANEPIC: A Map-Making Method for Timestream Data from Large Arrays." *The Astrophysical Journal*, *in press*, 2007.
- [40] E. W. Kamen, J. K. Su. *Introduction to Optimal Estimation*. Springer-Verlag London Limited, London, 1999.
- [41] P. Hargrave, T. Waskett, T. Lim, B. Swinyard. "Performance of flight-model on-board calibration sources on Herschel-SPIRE." In Zmuidzinas et al. [36].

- [42] S. Oliver, M. Rowan-Robinson, D. M. Alexander, O. Almaini, M. Balcells, A. C. Baker, X. Barcons, M. Barden, I. Bellas-Velidis, F. Cabrera-Guerra, R. Carballo, C. J. Cesarsky, P. Ciliegi, D. L. Clements, H. Crockett, L. Danese, A. Dapergolas, B. Drolias, N. Eaton, A. Efstathiou, E. Egami, D. Elbaz, D. Fadda, M. Fox, A. Franceschini, R. Genzel, P. Goldschmidt, M. Graham, J. I. Gonzalez-Serrano, E. A. Gonzalez-Solares, G. L. Granato, C. Gruppioni, U. Herbstmeier, P. Héraudeau, M. Joshi, E. Kontizas, M. Kontizas, J. K. Kotilainen, D. Kunze, F. La Franca, C. Lari, A. Lawrence, D. Lemke, M. J. D. Linden-Vørnle, R. G. Mann, I. Márquez, J. Masegosa, K. Mattila, R. G. McMahon, G. Miley, V. Missoulis, B. Mobasher, T. Morel, H. Nørgaard-Nielsen, A. Omont, P. Papadopoulos, I. Perez-Fournon, J.-L. Puget, D. Rigopoulou, B. Rocca-Volmerange, S. Serjeant, L. Silva, T. Sumner, C. Surace, P. Vaisanen, P. P. van der Werf, A. Verma, L. Vigroux, M. Villar-Martin, C. J. Willott. “The European Large Area ISO Survey - I. Goals, definition and observations.” *Monthly Notices of the Royal Astronomical Society*, **316**:749–767, August 2000.
- [43] J. Delabrouille, J.-F. Cardoso, G. Patanchon. “Multi-Detector Multi-Component spectral matching and applications for CMB data analysis.” *Monthly Notices of the Royal Astronomical Society*, **346**:1089, 2003.
- [44] R. G. Walker, S. D. Price. *AFCRL Infrared Sky Survey. Volume I. Catalog of Observations at 4, 11, and 20 Microns*. AFCRL-TR-75-0373. United States Air Force Cambridge Research Labs, July 1975.
- [45] M. J. Griffin, G. S. Orton. “The near-millimeter brightness temperature spectra of Uranus and Neptune.” *Icarus*, **105**:537–547, 1993.
- [46] E. L. Wright. “Infrared Brightness Temperature of Mars, 1983–2103.” *arXiv:astro-ph/0703640*, 2007.

- [47] H. Arp. “Atlas of Peculiar Galaxies.” *The Astrophysical Journal Supplement Series*, **14**:1, 1966.
- [48] L. Dunne, S. A. Eales. “The SCUBA Local Universe Galaxy Survey—II. 450- μm data: evidence for cold dust in bright IRAS galaxies.” *Monthly Notices of the Royal Astronomical Society*, **327**:697–714, 2001.
- [49] U. Klaas, M. Haas, S. A. H. Müller, R. Chini, B. Schulz, I. Coulson, H. Hippelein, K. Wilke, M. Albrecht, D. Lemke. “Infrared to millimetre photometry of ultraluminous IR galaxies: New evidence favouring a 3-stage dust model.” *Astronomy & Astrophysics*, **379**:823–844, 2001.
- [50] U. Lisenfeld, K. G. Isaak, R. Hills. “Dust and gas in luminous infrared galaxies—results from SCUBA observations.” *Monthly Notices of the Royal Astronomical Society*, **312**:433–441, 2000.
- [51] D. Rigopoulou, A. Lawrence, M. Rowan-Robinson. “Multiwavelength energy distributions of ultraluminous IRAS galaxies—I. Submillimetre and X-ray observations.” *Monthly Notices of the Royal Astronomical Society*, **278**:1049–1068, 1996.
- [52] D. B. Sanders, J. M. Mazzarella, D.-C. Kim, J. A. Surace, B. T. Soifer. “The IRAS Revised Bright Galaxy Sample.” *Astronomical Journal*, **126**:1607–1664, 2003.
- [53] M. Seiffert, C. Borys, D. Scott, M. Halpern. “An upper limit to polarized submillimetre emission in Arp 220.” *Monthly Notices of the Royal Astronomical Society*, **374**:409–414, 2007.
- [54] L. Spinoglio, P. Andreani, M. A. Malkan. “The Far-Infrared Energy Distributions of Seyfert and Starburst Galaxies in the Local Universe: Infrared Space Observatory Photometry of the 12 Micron Active Galaxy Sample.” *The Astrophysical Journal*, **572**:105–123, 2002.

- [55] C. H. Lineweaver, G. F. Smoot, C. L. Bennett, E. L. Wright, L. Tenorio, A. Kogut, P. B. Keegstra, G. Hinshaw, A. J. Banday. “Correlated Noise in the *COBE* DMR Sky Maps.” *The Astrophysical Journal*, **436**:452–455, 1994.
- [56] E. L. Wright, G. Hinshaw, C. L. Bennett. “Producing Mega-pixel CMB Maps from Differential Radiometric Data.” *The Astrophysical Journal*, **458**:L53–L55, 1996.
- [57] M. Tegmark. “How to Make Maps from CMB Data without Losing Information.” *The Astrophysical Journal*, **480**:L87–L90, 1996.
- [58] J. Borrill. “MADCAP—The Microwave Anisotropy Dataset Computational Analysis Package.” *arXiv:astro-ph/9911389v1*, 1999.
- [59] S. Prunet, C. B. Netterfield, E. Hivon, B. P. Crill. “Iterative Map-Making for Scanning Experiments.” In *Energy Densities in the Universe*, Proceedings of the XXXVth Rencontre de Moroid, January 2000.
- [60] M. A. J. Ashdown, C. Baccigalupi, A. Balbi, J. G. Bartlett, J. Borrill, C. Cantalupo, G. de Gasperis, K. M. Górski, V. Heikkilä, E. Hivon, E. Keihänen, H. Kurki-Suonio, C. R. Lawrence, P. Natoli, T. Poutanen, S. Prunet, M. Reinecke, R. Stompor, B. Wandelt. “Making Maps from Planck LFI 30GHz Data.” *Astronomy & Astrophysics*, **471**:361–380, 2007.
- [61] M. Frigo, S. G. Johnson. “The Design and Implementation of FFTW3.” *Proceedings of the IEEE*, **93**(2):216–231, 2005.
- [62] G. B. Rybicki, W. H. Press. “Interpolation, Realization, and Reconstruction of Noisy, Irregularly Sampled Data.” *The Astrophysical Journal*, **398**:169–176, 1992.
- [63] J. Makhoul. “Linear Prediction: A Tutorial Review.” In Childers [64], pages 99–118.
- [64] D. G. Childers, editor. *Modern Spectrum Analysis*. IEEE Press, New York, 1978.

- [65] W. H. Press, S. A. Teukolsky, W. T. Vetterling, B. P. Flannery. *Numerical Recipes in C : The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1988, 1992.
- [66] J. P. Burg. “A New Analysis Technique for Time Series Data.” In Childers [64], pages 42–48.
- [67] M. P. Viero, P. A. R. Ade, K. P. M. Balgrave, J. J. Bock, E. L. Chapin, M. J. Devlin, S. Dicker, M. Griffin, J. O. Gundersen, M. Halpern, P. C. Hargrave, D. H. Hughes, J. Klein, F. J. Lockman, G. Marsden, P. G. Martin, P. Mauskopf, C. B. Netterfield, L. Olmi, E. Pascale, G. Patachon, M. Rex, D. Scott, C. Semisch, M. D. P. Truch, C. Tucker, G. S. Tucker, D. V. Wiebe. “The Balloon-borne Large Aperture Submillimeter Telescope (BLAST) 2005: Dust Properties in the Intermediate Velocity Cloud G86.0+38.3” *In preparation*, 2008.
- [68] M.-A. Miville-Deschênes, G. Lagache, F. Boulanger, J.-L. Puget. “Statistical properties of dust far-infrared emission.” *Astronomy & Astrophysics*, **469**:595–605, 2007.
- [69] D. J. Schlegel, D. P. Finkbeiner, M. Davis. “Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds?” *The Astrophysical Journal*, **500**:525–553, 1998.
- [70] P. C. Hargrave, P. A. R. Ade, J. J. Bock, E. L. Chapin, M. J. Devlin, S. Dicker, M. Griffin, J. O. Gundersen, M. Halpern, D. H. Hughes, J. Klein, G. Marsden, P. G. Martin, P. Mauskopf, C. B. Netterfield, L. Olmi, E. Pascale, G. Patanchon, M. Rex, D. Scott, C. Semisch, B. Sibthorpe, M. D. P. Truch, C. Tucker, G. S. Tucker, M. P. Viero, D. V. Wiebe. “BLAST Observations of the Cassiopeia A Supernova Remnant” *In preparation*, 2008.
- [71] G. Westerhout. “A survey of the continuous radiation from the Galactic System at a

- frequency of 1390 Mc/s." *Bulletin of the Astronomical Institutes of the Netherlands*, **14**(488):215–260, 1958.
- [72] D. Downes, R. Rinehart. "Microwave Observations of the Cygnus X Region." *The Astrophysical Journal*, **144**:937–942, 1966.
- [73] J. M. Jackson, J. M. Rathborne, R. Y. Shah, R. Simon, T. M. Bania, D. P. Clemens, E. T. Chambers, A. M. Johnson, M. Dormody, R. Lavoie, M. H. Heyer. "The Boston University-Five College Radio Astronomy Observatory Galactic Ring Survey." *The Astrophysical Journal Supplement Series*, **163**:145–159, 2006.
- [74] L. Úbeda, A. Pellerin. "Discovery of Two Dust Pillars near the Galactic Plane." *The Astrophysical Journal*, *submitted*, 2007.
- [75] R. Simon, J. M. Jackson, D. P. Clemens, T. M. Bania, M. H. Heyer. "The Structure of Four Molecular Cloud Complexes in the BU-FCRAO Milky Way Galactic Ring Survey." *The Astrophysical Journal*, **92**:747–763, 2001.
- [76] G. Sandell, D. A. Weintraub. "On the Similarity of FU Orionis Stars to Class I Protostars: Evidence from the Submillimeter." *The Astrophysical Journal Supplement Series*, **134**:115–132, 2001.
- [77] L. Kohoutek. "Hamburg Schmidt-Camera Survey of Faint Planetary Nebulae." *Bulletin of the Astronomical Institute of Czechoslovakia*, **16**:221–227, 1965.
- [78] C. G. Wynn-Williams. "Radio Emission from the Vicinity of the Galactic Nebula NGC 6857." *Astrophysical Letters and Communications*, **3**:195, 1969.
- [79] M. Ryle, F. G. Smith. "A New Intense Source of Radio-Frequency Radiation in the Constellation of Cassiopeia." *Nature*, **162**(4116):462–463, 1948.
- [80] P. Rosati, P. Tozzi, R. Giacconi, R. Gilli, G. Hasinger, L. Kewley, V. Mainieri, M. Nonino, C. Norman, G. Szokoly, J. X. Wang, A. Zirm, J. Bergeron, S. Borgani,

- R. Gilmozzi, N. Grogin, A. Koekemoer, E. Schreier, W. Zheng. “The Chandra Deep Field-South: The 1 Million Second Exposure.” *The Astrophysical Journal*, **566**:667–674, 2002.
- [81] M. Dickinson, M. Giavalisco, The GOODS Team. “The Great Observatories Origins Deep Survey.” In R. Bender, A. Renzini, editors, *The Mass of Galaxies at Low and High Redshift*, ESO Astrophysics Symposia, page 324, October 2001.
- [82] R. Liseau, D. Lorenzetti, B. Nisini, L. Spinoglio, A. Moneti. “Star formation in the VELA molecular clouds. I - The IRAS-bright Class I sources.” *Astronomy & Astrophysics*, **265**:577–596, 1992.
- [83] G. J. Stanley, O. B. Slee. “Galactic Radiation at Radio Frequencies. II. The Discrete Sources.” *Australian Journal of Scientific Research—A: Physical Sciences*, **3**:234–250, 1950.
- [84] C. Hoffmeister. “316 neue Veränderliche.” *Astronomische Nachrichten*, **242**(5791):129–142, 1931.
- [85] P. Guthnick, H. Schneller. “Benennung von veränderlichen Sternen.” *Astronomische Nachrichten*, **268**(6419–20):165–190, 1939.
- [86] E. B. Holmberg, A. Lauberts, H.-E. Schuster, R. M. West. “The ESO/Upsalla Survey of the ESO (B) Atlas of the Southern Sky.” *Astronomy & Astrophysics Supplements Series*, **18**:463–489, 1974.
- [87] B. M. Lasker, R. D. Cannon. “A second epoch survey of the southern sky.” *Bulletin d’Information du Centre de Données Stellaires*, **37**:13–14, December 1989.
- [88] M. Faraday. “13 September 1845.” In T. Martin, editor, *Faraday’s Diary*, volumes 1–7. Bell, London, 1932–1936.

- [89] G. Sagnac. “L'éther lumineux démontré par l'effet du vent relatif d'éther dans un interféromètre en rotation uniforme.” *Comptes rendus hebdomadaires des séances de l'Académie des sciences*, **157**:708–710, 1913.
- [90] K. Strehl. “Aplanatische und fehlerhafte Abbildung im Fernrohr.” *Zeitschrift für Instrumentenkunde*, **15**:362–370, October 1895.
- [91] C. Ptolemy. *Ptolemy's Almagest*, translated and annotated by G. J. Toomer. Princeton University Press, 1998.

Μᾶλλον δὲ καὶ αὐτὸ τὸ ἀπλοῦν τῶν οὐρανίων οὐκ ἀπὸ τῶν παρ' ἡμῖν οὕτως ἔχειν
δοκούντων προσήκει κρίνειν, ὅποτε μηδ' ἐφ' ἡμῶν τὸ αὐτὸ πᾶσιν ὁμοίως ἐστὶν ἀπλοῦν·
—Κλαυδίου Πτολεμαίου Μαθηματικῆς Συντάξεως, ΙΓ':β' [1]