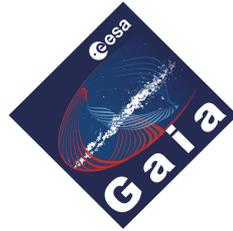


**CHARACTERISATION AND MITIGATION OF RADIATION DAMAGE
ON THE GAIA ASTROMETRIC FIELD**



This dissertation is submitted for the degree of
Doctor of Philosophy

by

SCOTT WILLIAM BROWN

Institute of Astronomy
&
Magdalene College

University of Cambridge
December 3, 2010



SUMMARY

Characterisation and mitigation of radiation damage on the Gaia Astrometric Field

Scott William Brown

In November 2012, the European Space Agency (ESA) is planning to launch Gaia, a mission designed to measure with microarcsecond accuracy the astrometric properties of over a billion stars. Microarcsecond astrometry requires extremely accurate positional measurements of individual stellar transits on the focal plane, which can be disrupted by radiation-induced Charge Transfer Inefficiency (CTI). Gaia will suffer radiation damage, impacting on the science performance, which has led to a series of Radiation Campaigns (RCs) being carried out by industry to investigate these issues. The goal of this thesis is to rigorously assess these campaigns and facilitate how to deal with CTI in the data processing.

We begin in Chapter 1 by giving an overview of astrometry and photometry, introducing the concept of stellar parallax, and establishing why observing from space is paramount for performing global, absolute astrometry. As demonstrated by Hipparcos, the concept is sound. After reviewing the Gaia payload and discussing how astrometric and photometric parameters are determined in practice, we introduce the issue of radiation-induced CTI and how it may be dealt with.

The on board mitigating strategies are investigated in detail in Chapter 2. Here we analyse the effects of radiation damage as a function of magnitude with and without a diffuse optical background, charge injection and the use of gates, and also discover a number of calibration issues. Some of these issues are expected to be removed during flight testing, others will have to be dealt with as part of the data processing, e.g. CCD stitches and the charge injection tail.

In Chapter 3 we turn to look at the physical properties of a Gaia CCD. Using data from RC2 we probe the density of traps (i.e. damaged sites) in each pixel and, for the first time, measure the Full Well Capacity of the Supplementary Buried Channel, a part of every Gaia pixel that constrains the passage of faint signals away from the bulk of traps throughout the rest of the pixel.

The Data Processing and Analysis Consortium (DPAC) is currently adopting a 'forward modelling' approach to calibrate radiation damage in the data processing. This incorporates a Charge Distortion Model (CDM), which is investigated in Chapter 4. We find that although the CDM performs well there are a number of degeneracies in the model parameters, which may be probed further by better experimental data and a more realistic model. Another way of assessing the performance of a CDM is explored in Chapter 5. Using a Monte Carlo approach we test how well the CDM can extract accurate image parameters. It is found that the CDM must be highly robust to achieve a moderate degree of accuracy and that the fitting is limited by assigning finite window sizes to the image shapes. Finally, in Chapter 6 we summarise our findings on the campaign analyses, the on-board mitigating strategies and on how well we are currently able to handle radiation damage in the data processing.

Cambridge, December 3, 2010

DECLARATION

I hereby declare that my thesis entitled *Characterisation and mitigation of radiation damage on the Gaia Astrometric Field* is not substantially the same as any that I have submitted for a degree or diploma or other qualification at any other University. I further state that no part of my thesis has already been or is being concurrently submitted for any such degree, diploma or other qualification. This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text.

I note that the Gaia logo on the title page is credited to ESA. I note that Chapter 1 is intended as a review, and contains a number of figures and tables that have been extracted from or consist of other works, all of which are clearly cited in the appropriate caption. Those parts of this thesis which have been published within the Gaia project are as follows:

- Much of the work contained in Chapter 2, section 2.2 formed the basis of the technical note published as **Brown, S.** and van Leeuwen, F. (2008) ‘Analysis of the CCN10 data set: CTI effects on a Gaia astrometric CCD.’, *Gaia Livelink* and **Brown, S.** and van Leeuwen, F. (2009) ‘Further analysis of Radiation Campaign 1: an extension to SWB-001.’, *Gaia Livelink*.
- Much of the work contained in Chapter 2, section 2.3 formed the basis of the technical note **Brown, S.** and van Leeuwen, F. (2009) ‘An independent analysis of Astrium’s Radiation Campaign 2 astrometric tests.’, *Gaia Livelink*.
- The work contained in Chapter 3 formed the basis of the technical note published as Prod’homme, T. and **Brown, S.** (2009) ‘Determination of Gaia Astrometric CCD Physical Parameters using Radiation Campaign 2 Charge Injection Profiles.’, *Gaia Livelink*, which was completed in collaboration with T. Prod’homme. I provided the information for section 3.1. Section 3.2 (except part 3.2.2) was put together during a collaborative visit to Leiden Observatory. Afterwards, some of the interpretation and tables were extended by T. Prod’homme. In section 3.3, the plots were produced during the same visit. I performed the curve fittings and associated table of values. The interpretation of the results were done together.
- The work contained in Chapter 4 formed the basis of the technical note published as Prod’homme, T., Weiler, M. and **Brown, S.** (2009) ‘CDM02 Validation: A comparison between CDM02 outcomes and RC2 data.’, *Gaia Livelink*, and was completed in collaboration with the named authors. My main contribution was in providing the reduced data as described in section 4.2 and helping to set up the simulation as in section 4.3. I include some of their results and interpretation, which I have added to, for completeness. I also add to their conclusions focusing on the quality of the test data and possible improvements that could be made.

This thesis contains fewer than 60,000 words.

Scott W. Brown

Cambridge, December 3, 2010

ACKNOWLEDGEMENTS

Opportunities to work at the forefront of scientific research on a European space mission do not come along everyday. So for this I have many people to thank. Of course it all began when my supervisor Floor van Leeuwen took me on as a PhD Student and ensured that I became involved with the Gaia project as soon as possible. For this I am extremely grateful. The guidance you've given me and the enthusiasm you've shown in our work has been invaluable. Thank you Floor.

Throughout the Gaia community there are many people with whom I've had the opportunity to work with directly, communicate with or present results to at various meetings. In particular, my thanks go to Thibaut Prod'homme for inviting me to work in Leiden for a short time as well as for collaborating on several key Gaia-related issues. This includes Michael Weiler, and the rest of the CTI dream team, with whom I also had a chance to work with.

I am grateful to Alex Short, Lennart Lindegren, Anthony Brown, Patricio Ortiz, Nicholas Cross and George Seabroke for their questions and communications on the analysis of the radiation campaigns as well as all of the other members of the Radiation Task Force for their feedback. To those at ESA and at EADS Astrium, I would like to thank you all for your help; in particular, Jos de Bruijne, Laurent Georges and Jean-Francois Pasquier.

At the Institute of Astronomy I am very grateful to Francesca de Angeli for helping me out with Java in the early days and Nic Walton for providing feedback on my progress throughout the PhD as well as Sian Owen and Joy McSharry for organising deadlines and training courses. The funding for this research was kindly provided by the UK Science and Technology Facilities Council (STFC), which has allowed me to participate in many Gaia-related schools, workshops and meetings.

Closer to home, I cannot begin to thank my Mum and Dad enough for their love and support throughout my time not only as a student but also down the many different paths in life I have decided to venture. I'm sure Dominique would agree with me when I say how much you mean to us both and how you've enriched our lives in so many ways - *alla famiglia!* Speaking of sisters, I cannot help thinking that in some way mine may have lent a helping hand when I decided to 'reach for the stars'. Writing from halfway around the world you certainly planted a seed in my mind that helped lead me here. For this I cannot thank you enough. To my friends both in Cambridge and those scattered around the world it goes without saying how good it feels to share this with you all.

Elena, you've helped me countless times throughout my PhD, so in a way this has also been a journey for you. Thank you so much for the laughs, the distractions, for your patience, warmth and love. Now, above all, I look forward to our adventure.

Scott W. Brown

Cambridge, December 3, 2010

For my family.

CONTENTS

Summary	iii
Declaration	v
Acknowledgements	vii
Contents	xi
List of Tables	xv
List of Figures	xvi
Acronyms	xxiv
1 Introduction	1
1.1 Astrometry	2
1.1.1 Measuring the distances to the stars: stellar parallax.	3
1.2 Photometry	6
1.3 Ground vs space-based astrometry	7
1.4 The legacy of Hipparcos	8
1.5 Introducing Gaia: 'The Billion Star Surveyor'	9
1.5.1 The payload	10
1.5.2 The focal plane	11
1.5.3 The scanning law	13
1.6 Determining astrometric parameters	14
1.7 Determining photometric parameters	15
1.8 Processing the data: from the raw data to a Gaia Catalogue	15
1.9 Radiation-induced Charge Transfer Inefficiency: a major challenge	17
1.9.1 The solar cycle	19
1.9.2 Dealing with radiation damage: learning from other missions	20
1.9.3 Mitigating radiation damage on board Gaia	22
1.9.4 Mitigating radiation damage in the data processing	23
1.10 A template for this thesis	26
2 The Radiation Campaigns	29
2.1 Methodology	30
2.1.1 Charge-Coupled Device (CCD) readout	30
2.1.2 Extracting the background	32
2.1.3 Extracting the raw image	33

2.1.4	Combining 1D stellar images: sub-pixel sampling	34
2.1.5	Defining a robust Global Reference Curve	35
2.1.6	Fitting the GRC to the stellar field: extracting positions and fluxes	36
2.1.7	Calibrating the geometry of the mask and extracting centroid biases	38
2.1.8	Calibrating the flux response across the mask and extracting charge losses	40
2.1.9	Comparing local non-irradiated and irradiated image profiles	41
2.2	Radiation Campaign 1	42
2.2.1	Hardware and test configuration	42
2.2.2	Background analysis	45
2.2.3	The reference data	54
2.2.4	The irradiated data	59
2.2.5	Relative centroid bias and charge loss in the leading column	62
2.2.6	Conclusions	70
2.3	Radiation Campaign 2	73
2.3.1	Charge Injection	73
2.3.2	Hardware and test configuration	74
2.3.3	Background analysis	78
2.3.4	The reference data	90
2.3.5	The irradiated data	94
2.3.6	Relative centroid bias and charge loss in the leading column	110
2.3.7	Conclusions	116
2.4	Radiation Campaign 3	121
2.4.1	Hardware and test configuration	121
2.4.2	Background analysis	123
2.4.3	The reference data	131
2.4.4	The irradiated data	133
2.4.5	Relative centroid bias and charge loss in the leading column	140
2.4.6	Conclusions	148
2.5	Next steps	151
2.6	Impact on astrometric performance	152
3	Probing an Astrometric Field CCD at the Microscopic Level	155
3.1	The data: radiation campaign 2 charge injection profiles	156
3.2	Trap densities	157
3.2.1	First Pixel Response	157
3.2.2	Reference levels	159
3.2.3	Results	161
3.2.4	Total charge loss	164
3.3	Measuring the Supplementary Buried Channel Full Well Capacity	165
3.4	Conclusions	168

4	Charge Distortion Model Validation	171
4.1	Charge Distortion Models	171
4.2	The test data	172
4.2.1	Input to the simulation	172
4.2.2	Comparative output for the simulation	173
4.3	The simulation procedure	174
4.3.1	The Downhill-Simplex method	175
4.3.2	How many trap species are required?	177
4.4	Initial parameters	177
4.5	Results	178
4.5.1	Test 1: fitting to a single damaged image profile	178
4.5.2	Test 2: fitting to a set of damaged image profiles with different CI delays	180
4.5.3	Test 3: fitting to a set of damaged image profiles with different magnitudes	183
4.5.4	Test 4: fitting to every available damaged image profile	185
4.6	Conclusions	186
5	Image Parameter Extraction	189
5.1	Monte Carlo Profile Generator	190
5.1.1	The seed profile	190
5.1.2	Random locations	191
5.1.3	Random fluxes	191
5.2	Extracting image parameters	191
5.2.1	A first approximation	191
5.2.2	A final determination	193
5.3	The observed image shape	194
5.4	The predicted image shape	196
5.5	Varying the window size	198
5.6	Conclusions	200
6	Conclusions and Final Remarks	203
6.1	On the hardware and test setup	203
6.2	On the process of calibrating the campaign data	205
6.3	On the effect of the diffuse optical background (DOB)	206
6.4	On the effects of charge injection (CI)	206
6.5	On the microscopic properties of an irradiated Astrometric Field CCD	207
6.6	On Charge Distortion Model (CDM) validation	208
6.7	On extracting image parameters in the data pipeline	208
6.8	On current radiation campaigns	209
6.9	Outside Gaia: a study of HST images	209
6.10	Final remarks	210

A Analytical tools	211
A.1 Combining 1D stellar images	211
A.1.1 Cross-correlation	211
A.1.2 Quadratic Interpolation	211
A.2 Model fitting: extracting image parameters	212
A.2.1 First-order Taylor expansions and Householder least squares	212
A.2.2 Maximum Likelihood 1	213
A.2.3 Maximum Likelihood 2	215
B Further results from the Radiation Campaigns	217
B.1 Radiation-induced CTI image parameters from RC1	217
B.2 Results from the analysis of RC2 at $G = 15$	218
B.3 Radiation-induced CTI image parameters from RC2	221
B.4 Radiation-induced CTI image parameters from RC3	222
C Trap Physics and Charge Distortion Models	225
C.1 Shockley Read Hall Theory	225
C.2 Capture and emission probabilities	226
C.3 CDM01	226
C.3.1 Propagation of the state variable	228
C.3.2 Initialization of the state variable	229
C.4 CDM02	229
C.4.1 Electron capture	230
C.4.2 Electron release	231
C.4.3 Setting the initial trap occupancy state	231
C.5 CDM03	232
Bibliography	233

LIST OF TABLES

1.1	A selection of stars and their astrometric parameters from the Hipparcos catalogue	5
1.2	Gaia requirements on end-of-mission parallax standard errors	9
1.3	DPAC Coordination Units.	16
2.1	RC1: signal intensities in terms of the total flux and G-band magnitude.	43
2.2	RC1: sampling scheme and associated CCD readout area.	45
2.3	RC1: statistics on the noise in the background.	48
2.4	RC1: global statistics on the background in the reference data.	50
2.5	RC1: correlation matrix for background variations in the reference data.	51
2.6	RC1: detecting AC image profiles in neighbouring samples.	52
2.7	RC2: sampling scheme and associated CCD readout area.	76
2.8	RC2: information on the first CI test.	77
2.9	RC2: information on the second CI test.	78
2.10	RC2: information on the third CI test.	78
2.11	RC2: statistics on the noise in the background.	83
2.12	RC2: global statistics on the background in the reference data.	84
2.13	RC2: AC variability in the CI level.	96
2.14	RC3: information on the MF star tests.	122
2.15	RC3: information on the bright star tests.	122
2.16	RC3: statistics on the background.	126
3.1	RC2 CI levels available for probing the trap density.	156
3.2	RC2 CCD temperatures available to probe the trap density.	156
3.3	Trap intervals as a function of temperature using the FPR.	163
3.4	Trap intervals as a function of temperature using the total charge loss.	164
3.5	Values for curve minima and maxima for estimating the SBC FWC.	167
4.1	Simulation parameters for CDM validation	177
4.2	Simulation initialisation parameters for CDM validation	178
4.3	CDM02 best fit parameters at $G = 15$ from the first CDM test.	178
4.4	CDM02 best fit parameters at different magnitudes from the first CDM test.	179
4.5	CDM02 best fit parameters for a specific magnitude with different CI delays	181
4.6	CDM02 best fit parameters for a CI delay ≤ 150 and different sets of magnitudes.	184

4.7	CDM02 best fit parameters for a CI delay > 150 and different sets of magnitudes.	184
4.8	CDM02 best fit parameters to fit all of the available test data.	185
B.1	RC1: CTI image parameters.	217
B.2	RC2: CTI image parameters measured in the first CI test.	221
B.3	RC2: CTI image parameters measured in the second CI test.	221
B.4	RC2: CTI image parameters measured in the third CI test.	221
B.5	RC3: CTI effects measured in the MF star tests.	222
B.6	RC3: CTI effects measured at $G = 20$	222
B.7	RC3: CTI effects measured at $G = 20$ and the lowest possible DOB.	222
B.8	RC3: CTI effects measured in the bright star tests.	223

LIST OF FIGURES

1.1	Stellar parallax	4
1.2	Stellar motion	4
1.3	Improvements in astrometric accuracy.	5
1.4	The principle of absolute parallax measurement.	7
1.5	Schematic of the Hipparcos telescope.	9
1.6	Gaia's reach.	10
1.7	The Gaia payload.	11
1.8	The Gaia focal plane.	12
1.9	The Gaia scanning law.	13
1.10	CCD proton damage.	18
1.11	An example of Hubble CTI charge release trails.	18
1.12	Solar Cycle 24.	19
1.13	Calibrating radiation damage in the data pre-processing.	23
1.14	Modelling radiation-induced image distortion I	25
1.15	Modelling radiation-induced image distortion II	25
1.16	The current radiation calibration strategy.	26
2.1	An example of a raw data file read out from the CCD.	30
2.2	An example of the sampling scheme used in the CCD readout.	31
2.3	How the background is extracted using a flat cut.	32
2.4	How a raw image is extracted from the CCD readout.	33
2.5	How to construct a GRC using the reference data.	35

2.6	How to extract relative centroid positions and fluxes from across the stellar field.	37
2.7	How to extract relative centroid biases from across the stellar field.	40
2.8	How to extract relative charge losses from across the stellar field.	40
2.9	RC1: details of the test setup.	44
2.10	RC1: running means of samples 1 and 7 in the reference data.	46
2.11	RC1: running mean of sample 6 at 60000 and a DOB level of 0, 5, and 10.	47
2.12	RC1: background noise distributions in the reference data.	47
2.13	RC1: Mean background flux as a function of AC position in the reference data.	49
2.14	RC1: the overlap of AC image profiles with neighbouring transits.	51
2.15	RC1: the overlap of AC image profiles in the raw data.	52
2.16	RC1: constructing the GRC.	54
2.17	RC1: unexpected features in the GRC.	55
2.18	RC1: geometric calibration.	56
2.19	RC1: flux response calibration I.	56
2.20	RC1: flux response calibration II.	57
2.21	RC1: total flux calibration.	57
2.22	RC1: regular vs custom-spaced knots.	58
2.23	RC1: investigating the asymmetry in the GRC.	59
2.24	RC1: finding a new GRC for irradiated data with a DOB.	60
2.25	RC1: illustrating the preferential damage seen in scan 1.	61
2.26	RC1: integrated flux as a function of AC position in the irradiated data.	63
2.27	RC1: relative centroid bias as a function of signal intensity and DOB level.	64
2.28	RC1: relative charge loss as a function of signal intensity and DOB level.	65
2.29	RC1: comparing Astrium's results.	66
2.30	RC1: average relative centroid bias.	67
2.31	RC1: average relative charge loss.	68
2.32	RC1: image distortion at 60000 and 650 with and without a DOB.	69
2.33	RC1: image distortion as a function of radiation level.	69
2.34	RC1: image distortion as a function of signal intensity and DOB level.	70
2.35	RC2: charge injection structure	74
2.36	RC2: details of the test setup.	75
2.37	RC2: CI tests performed (level, duration, delay).	77
2.38	RC2: running mean of the reference data (first CI test, $G = 13.3$).	79
2.39	RC2: running mean of the reference data (first part of the third CI test).	80
2.40	RC2: background behaviour in the reference data from the first CI test at $G = 13.3$	81
2.41	RC2: background behaviour in the reference data from the first part of the third CI test.	82
2.42	RC2: background noise distributions in the reference data in the first CI test.	82
2.43	RC2: dependencies in the background noise of the reference data.	84
2.44	RC2: running mean of the irradiated data (third CI test, $G = 13.3$, CI delay = 150).	85
2.45	RC2: running mean of the irradiated data (third CI test, $G = 13.3$, CI delay = 30000).	86

2.46	RC2: additional features seen in the third CI test.	86
2.47	RC2: time variations between scans at $G = 16.75$ with a CI delay of 1000.	87
2.48	RC2: background behaviour in the irradiated data at $G = 16.75$ and a CI delay of 300.	87
2.49	RC2: the effect of image spacing on the estimation of the background level.	88
2.50	RC2: constructing the GRC.	90
2.51	RC2: FWHM of the GRC for each mask column.	91
2.52	RC2: calibrating the geometry of the mask I.	92
2.53	RC2: calibrating the geometry of the mask II.	92
2.54	RC2: geometrical offsets found across the mask.	93
2.55	RC2: calibrating the flux response I.	93
2.56	RC2: calibrating the flux response II.	95
2.57	RC2: relative centroid positions in column 1 and column 50 for a CI of 7.51V.	96
2.58	RC2: investigating which irradiated image centroids appear undamaged.	97
2.59	RC2: calibrated relative centroid positions in column 1 and column 50.	97
2.60	RC2: onset of centroid bias in column 1 at various CI levels.	98
2.61	RC2: relative centroid bias across the stellar field at various CI levels.	99
2.62	RC2: relative centroid bias seen edge-on when the CI is 8.06V.	100
2.63	RC2: relative centroid bias seen edge-on without CI.	100
2.64	RC2: CCD stitching pattern.	101
2.65	RC2: fitting images in channels 1 and 22 in column 1 with a CI of 6.51V.	101
2.66	RC2: onset of charge loss in column 1 at various CI levels.	102
2.67	RC2: relative charge loss across the stellar field at various CI levels.	103
2.68	RC2: relative charge loss seen edge-on when the CI is 8.06V.	104
2.69	RC2: relative charge loss seen edge-on without CI.	104
2.70	RC2: integrated flux and radiation damage as a function of CI level ($G = 13.3$).	105
2.71	RC2: additional step required for calibrating the charge loss.	107
2.72	RC2: integrated flux and radiation damage as a function of CI duration ($G = 13.3$).	108
2.73	RC2: integrated flux and radiation damage as a function of CI delay ($G = 13.3$).	109
2.74	RC2: relative charge loss as a function of CI delay for $G > 16$	111
2.75	RC2: relative centroid bias and charge loss as a function of CI level ($G \lesssim 15$).	112
2.76	RC2: relative centroid bias and charge loss as a function of CI duration ($G \lesssim 15$).	113
2.77	RC2: relative centroid bias and charge loss as a function of CI delay ($G \lesssim 15$).	113
2.78	RC2: relative centroid bias and charge loss as a function of CI delay ($G \gtrsim 16.75$).	114
2.79	RC2: image distortion in the first CI test.	115
2.80	RC2: image distortion in the second CI test.	116
2.81	RC2: image distortion in the third CI test.	117
2.82	RC3: details on the test setup.	122
2.83	RC3: schematic of an AF CCD showing the AL positions of some of the gates.	123
2.84	RC3: running mean of the reference data (first MF test, $DOB = 2$).	124
2.85	RC3: running mean of the irradiated data ($G = 13.6$, $DOB = 2$, CI delay = 1000).	124

2.86	RC3: background behaviour in the reference data from the first MF test with a DOB of 2.	125
2.87	RC3: background noise distributions in the reference data.	125
2.88	RC3: mean background level in the irradiated data as a function of AC position and time.	126
2.89	RC3: global statistics on the reference data in the first MF test with a DOB of 2.	127
2.90	RC3: CI tails observed at CI delays of 10 and 30 in the non-irradiated part of the CCD.	127
2.91	RC3: discrepancies in the CI profiles I.	128
2.92	RC3: discrepancies in the CI profiles II.	128
2.93	RC3: how the CI tail is extracted from the stellar field.	130
2.94	RC3: constructing the GRC.	131
2.95	RC3: geometrical offsets found across the mask.	132
2.96	RC3: flux response calibration.	132
2.97	RC3: sample clustering.	134
2.98	RC3: original geometric calibration.	134
2.99	RC3: refined geometric calibration.	135
2.100	RC3: relative centroid bias as a function of CI delay ($G = 13.6$, $DOB = 7$).	136
2.101	RC3: relative charge loss as a function of CI delay ($G = 13.6$, $DOB = 7$).	136
2.102	RC3: relative centroid bias as a function of DOB ($G = 20$, CI delay = 1000).	137
2.103	RC3: relative charge loss as a function of DOB ($G = 20$, CI delay = 1000).	137
2.104	RC3: relative centroid bias as a function of CI delay ($G = 11.1$, $DOB = 4$, gate 9).	138
2.105	RC3: relative charge loss as a function of CI delay ($G = 11.1$, $DOB = 4$, gate 9).	139
2.106	RC3: relative centroid bias as a function of DOB ($G = 12.6$, CI delay = 1000, gate 11).	139
2.107	RC3: relative charge loss as a function of DOB ($G = 12.6$, CI delay = 1000, gate 11).	140
2.108	RC3: radiation damage as a function of AC position and CI delay ($G = 13.6$, $DOB = 2$).	141
2.109	RC3: radiation damage as a function of AC position and CI delay ($G = 20$, $DOB = 2$).	141
2.110	RC3: radiation damage as a function of AC position and CI delay ($G = 11.1$, no DOB).	142
2.111	RC3: radiation damage as a function of AC position and CI delay ($G = 12.6$, no DOB).	142
2.112	RC3: radiation damage as a function of CI delay and DOB ($G = 13.6$).	143
2.113	RC3: radiation damage as a function of CI delay and DOB ($G = 20$).	143
2.114	RC3: radiation damage as a function of CI delay and DOB ($G = 11.1$, gate 9).	144
2.115	RC3: radiation damage as a function of CI delay and DOB ($G = 12.6$, gate 11).	144
2.116	RC3: radiation damage as a function of magnitude and CI delay ($DOB = 2$).	145
2.117	RC3: radiation damage as a function of magnitude and CI delay ($DOB = 7$).	145
2.118	RC3: image distortion at $G = 13.6$	146
2.119	RC3: image distortion at $G = 20$	146
2.120	RC3: image distortion as a function of CI delay with the lowest possible DOB.	146
2.121	RC3: sky-like test background analysis I.	148
2.122	RC3: sky-like test background analysis II.	148
2.123	RC3: calibrating the stellar field in the sky-like test.	149
3.1	Two examples of unsaturated CI profiles from RC2.	157

3.2	Calculated CI reference levels at 163K (T1).	160
3.3	Calculated CI reference levels at 198K (T5).	160
3.4	FPR and trap density as a function of reference level and CCD temperature.	162
3.5	Relative FPR as a function of reference level and CCD temperature.	163
3.6	Total fractional charge loss as a function of reference level and temperature.	164
3.7	Gaia CCD pixel potential.	165
3.8	Fractional charge loss at T1 (163K) with a cubic spline fit.	166
3.9	Fractional charge loss at all five temperatures with cubic spline fits.	166
4.1	An example of the input to the CDM validation tests.	173
4.2	CDM02 best fits to individual damaged image shapes at different magnitudes without CI	179
4.3	CDM02 best simultaneous fits to damaged image shapes at $G = 15$ with different CI delays	181
4.4	CDM02 best simultaneous fits to damaged image shapes without CI at different mag- nitudes	183
5.1	The seed profile used to generate 10000 image profiles.	190
5.2	Testing the performance of the initial centroid estimates	193
5.3	Observed image shape parameter statistics	195
5.4	Observed formal errors in position and magnitude as a function of magnitude	195
5.5	The predicted image shape vs the observed image shape	196
5.6	Predicted image shape parameter statistics	197
5.7	Predicted formal errors in position and magnitude as a function of magnitude	197
5.8	Image parameter statistics as a function of window size	199
5.9	Observed formal errors in position and magnitude as a function of magnitude with an 18-pixel window	200
A.1	Combining 1D stellar images: cross-correlation and quadratic interpolation.	212
B.1	RC2: integrated flux and radiation damage as a function of CI level ($G = 15$).	218
B.2	RC2: integrated flux and radiation damage as a function of CI duration ($G = 15$).	219
B.3	RC2: integrated flux and radiation damage as a function of CI delay ($G = 15$).	220

ACRONYMS

- 1D** one-dimensional
- 2D** two-dimensional
- AC** across-scan
- ACS** Advanced Camera for Surveys
- AF** Astrometric Field
- AGIS** Astrometric Global Iterative Solution
- AL** along-scan
- BAM** Basic-Angle Monitor
- BAR** Barycentre
- BP** Blue Photometer
- BC** Buried Channel
- CCD** Charge-Coupled Device
- CCN10** Contract Change Notice no. 10
- CDM** Charge Distortion Model
- CEMGA** CTI Effects Models for Gaia
- CI** charge injection
- CTI** Charge Transfer Inefficiency
- CU** Coordination Unit
- DOB** diffuse optical background
- DPAC** Data Processing and Analysis Consortium
- DS** Downhill-Simplex
- DU** Development Unit
- EADS** European Aeronautic Defence and Space company
- ELSA** European Leadership in Space Astrometry
- EPIC** European Photon Imaging Cameras

- ESA** European Space Agency
- ESAC** European Space Astronomy Centre
- ESTEC** European Space Research and Technology Centre
- FOV** Field of View
- FPA** Focal Plane Assembly
- FPR** First Pixel Response
- FWC** Full Well Capacity
- FWHM** Full Width at Half Maximum
- G** G band magnitude
- GASS** Gaia System Simulator
- GCR** Galactic Cosmic Rays
- GIBIS** Gaia Instrument and Basic Image Simulator
- GOF** goodness-of-fit
- GOG** Gaia Object Generator
- GPD** Gaia Parameter Database
- GRC** Global Reference Curve
- HHLS** Householder least squares
- HST** Hubble Space Telescope
- IAU** International Astronomical Union
- ICRF** International Celestial Reference Frame
- ICRS** International Celestial Reference System
- IDA** IDT/IDU Algorithms
- IDT** Initial Data Treatment
- IDU** Intermediate Data Update
- IA** Image Area
- LED** Light Emitting Diode

-
- LSB** Least Significant Bit
- LSF** Line Spread Function
- MCPG** Monte Carlo Profile Generator
- MF** medium/faint
- ML1** Maximum Likelihood 1
- ML2** Maximum Likelihood 2
- MXC** Mirror Cross-Correlation
- OITF** Offset Instability Task Force
- PDR** Preliminary Design Review
- PEM** Proximity Electronics Module
- PSF** Point Spread Function
- RC** Radiation Campaign
- RCWG** Radiation Calibration Working Group
- RON** Readout Noise
- RP** Red Photometer
- RTF** Radiation Task Force
- RVS** Radial Velocity Spectrometer
- SAA** Solar Aspect Angle
- SBC** Supplementary Buried Channel
- SED** Spectral Energy Distribution
- SRH** Shockley-Read-Hall
- SM** Sky Mapper
- SNR** signal-to-noise ratio
- SR** Serial Register
- STFC** Science and Technology Facilities Council
- STIS** Space Telescope Imaging Spectrograph

SVN subversion

TBW Tukey's Biweight

TDI Time-Delay Integration

VLBI Very Long Baseline Interferometry

WFC3 Wide Field Camera 3

WFPC2 Wide Field Planetary Camera 2

WFS Wave-Front Sensor

XMM X-ray Multi-Mirror

1

Introduction

THROUGHOUT history the interplay between science and technology has led to some of our greatest breakthroughs and deepest insights, no more so than in the field of astronomy. To understand the universe and the objects that exist naturally within it, astronomers invariably need to use the best equipment available. This often means pushing current technology to the limit; the further we wish to see, the more refined our instruments must be.

In more recent times, with a host of new observatories that are either online or close to completion, astronomy has begun to make considerable progress. Ground-based as well as space-based programmes have led to an explosion in astronomical data, significantly increasing the chances of new discoveries being made in the near future. For many, a great deal of this data is available now, yet for others there is the challenge of predicting what the data will look like before it arrives. Space programmes, in particular, can take more than 10 years to carry through from concept to launch and as the mission begins it is usually advantageous to process the data autonomously. Covering all eventualities is a complex task, which is why so many astronomers are dedicated to these preparatory stages.

In November 2012 the European Space Agency (ESA) is planning to launch Gaia, a mission which is expected to measure the geometric and kinematic properties of over a billion stars with unprecedented accuracy. Given the extreme demands on the science requirements and on the hardware that must meet those demands, Gaia is a good example of how modern research is being driven by advances in both science and technology. Predicting how Gaia will perform is paramount, which is why the focus of this thesis is concerned with the effects of radiation damage on the mission.

Gaia's legacy will be to build a stereoscopic map of the Milky Way galaxy so that by knowing *where* everything is and *how* it moves, we can begin to ask *why*.

1.1 Astrometry

Astrometry is a branch of science that focuses on measuring precisely and accurately the positions of celestial bodies as a function of time. It has a rich heritage (cf. Høg (2009)) and in more recent times, with access to space and advances in instrumentation, the subject has been given fresh impetus and seen a resurgence in interest from all fields of astronomy.

Astrometry provides the only direct means of measuring the distances to the stars, so precise and accurate astrometric measurements, coupled with information on a star's intrinsic properties, are invaluable for interpreting the dynamics of stars and testing theories on stellar evolution. Take this one step further by conducting an astrometric survey on a Galactic scale and there the Milky Way becomes a study ground. Galactic evolution, the star formation history of the Milky Way, even fundamental physics; these are all applications that astrometry can open doors to.

In total, there are six parameters that define the geometric and kinematic properties of a celestial body. Five of those can typically be found using astrometric measurements:

- α : the right ascension or longitude in the chosen coordinate system.
- δ : the declination or latitude in the chosen coordinate system.
- μ_α : the proper motion or the slow, angular change of a body's position in the direction of α . Units are typically given in milliarcseconds (mas) per year (yr^{-1}).
- μ_δ : the proper motion or the slow, angular change of a body's position in the direction of δ . Units are also typically mas yr^{-1} .
- π : the parallax or the apparent periodic angular displacement of a body's position on the sky caused by the Earth's orbit around the Sun. This will be explained in more detail in the next section. Typical units are mas.

The sixth, v_r , is the radial velocity or the velocity along the line-of-sight, typically measured in km s^{-1} , which is usually determined by the wavelength shift observed in a body's spectrum (a positive shift indicating recession) rather than using astrometric measurements. This does not mean the latter cannot be used; the radial velocity may be inferred using astrometric measurements and indeed there is more than one way of doing this as described in Dravins et al. (1999). The most direct method, which involves measuring the change in parallax, would require very accurate measurements over very long time periods; unfeasible in the short term. However, as astrometric missions become more sophisticated this may change. Gaia will measure the radial velocities of a small fraction of the stars it observes using a spectrometer to determine the Doppler shift of spectral lines as mentioned above.

To know where a celestial body is on the sky in relation to every other object in the universe, the position of the body must be referred to a particular coordinate system. There are many different types (e.g. equatorial, ecliptic, barycentric, galactic) all of which are valid. However, in order to relate all astrometric programmes a standard has to be agreed upon at the international level. The International Celestial Reference System (ICRS) is an idealised barycentric coordinate system, which in practice is

realised by the International Celestial Reference Frame (ICRF) (Ma et al. (1998)) and is composed of at least 212 quasars¹. As quasars exhibit very high redshifts, they are either assumed to be fixed in position or, at the very least, their proper motions and parallaxes are expected to be negligible compared to the errors in measuring their positions. Quasars are also radio emitters, so to realise the first ICRF these objects were observed at radio frequencies using Very Long Baseline Interferometry (VLBI). In 1997 the ICRF was realised at optical wavelengths by Hipparcos, an ESA mission launched to take astrometric measurements of more than 10^5 stars. Hipparcos will be discussed in more detail in section 1.4².

1.1.1 Measuring the distances to the stars: stellar parallax.

Aside from the positions and proper motions of stars, a fifth parameter that can be deduced using astrometric measurements is parallax. By observing the yearly variations in a star's position as the Earth orbits the Sun, it is possible to calculate a star's parallax and from this determine its distance from the Sun. Figure 1.1 illustrates this concept. Stars positioned between the Sun and more distant background stars appear to trace out a repeating path on the sky, which is linked to their distance. The shape of the projected motion will vary, depending on the angle between the direction to the star and the ecliptic.

The parallax is defined as being one-half of the maximum change in angular position a star moves through on the sky as observed from diametrically opposite ends of the Earth's orbit. The stars are so far away that the small angle approximation ($\tan \pi \approx \pi$) may be used:

$$\pi \equiv \frac{r}{d}, \quad (1.1)$$

where r is the radius of the Earth's orbit (1 Astronomical Unit (AU) $\approx 1.5 \times 10^8$ km), and d is the distance between the Sun and the star. d is usually quoted in parsecs. A parsec or pc is an abbreviation for the unit of parallax-arcsecond and is defined according to the definition of stellar parallax. If a star's parallax subtends an angle of 1 arcsecond (arcsec or ") on the sky then the distance to that star is defined as 1 pc, which is equivalent to 206265 AU.

Figure 1.1 over-simplifies the path a star may actually take across the sky. To give a better picture both proper motion and, when present, orbital motion should be included. The proper motion of a star is also seen as a slow, angular change in the star's position. However, the proper motion is not periodic. Over time the star will move across the sky relative to the Sun and superimposed on this will be the apparent motion due to parallax. If the star is in a binary or multiple system, the star may exhibit orbital motion, which again will superimpose itself on the proper motion and parallax. Figure 1.2 attempts to explain these concepts and give a sense of how the stars are in no way static, quantified as they are by the astrometric technique.

Figure 1.2 represents an ideal case where the positions of the stars have been measured with extremely high precision and the errors are far smaller than the astrometric parameters we are trying to measure. In

¹In August 2009 ICRF2 was submitted for adoption by the International Astronomical Union (IAU) at the XXVII General Assembly. ICRF2 contains precise measurements of the positions of 3414 compact radio astronomical sources.

²Before the ICRF was established, a catalogue of fundamental stars was used to provide a standard celestial reference system, the latest being the FK6 system (Wielen et al. (2002)). FK6 combines the measurements of Hipparcos with ground-based data summarised in the previous catalogue, FK5.

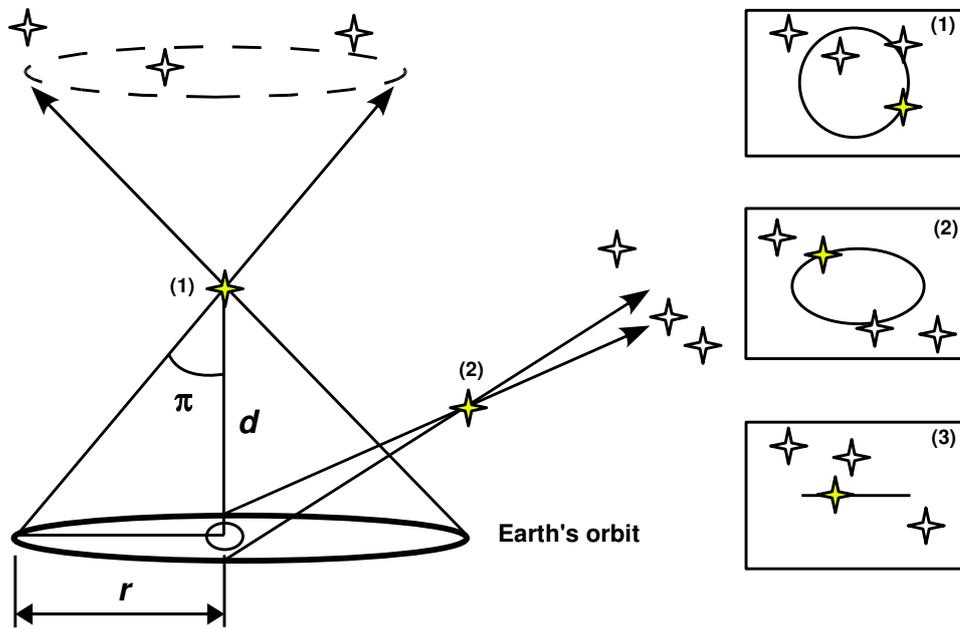


Figure 1.1: When observing a patch of sky from the Earth at different times of the year some stars appear to move along a repeating path relative to more distant stars. This apparent motion is stellar parallax. (1) Here the star is situated towards the pole of the ecliptic and appears to trace out a circle as the Earth orbits the Sun. (2) Here the star lies about 45° above the ecliptic and traces out an ellipse as the Earth orbits the Sun. (3) Here the star lies in the plane of the ecliptic and traces out a line as the Earth orbits the Sun (not shown explicitly on the left). The parameters r (1 Astronomical Unit $\approx 1.5 \times 10^8$ km), π (the parallax), and d (the distance between the Sun and the star) are related by eqn. 1.1.

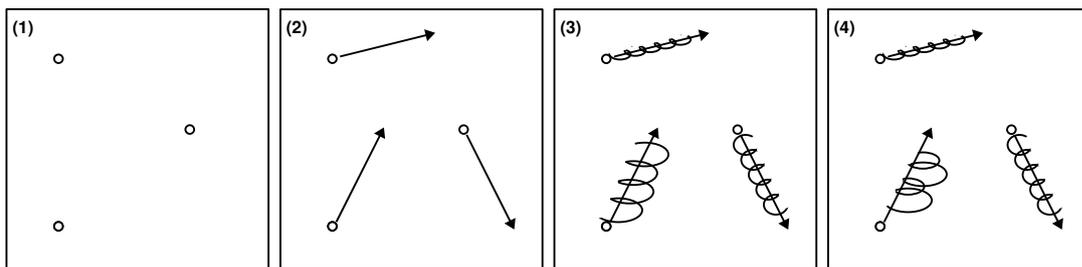


Figure 1.2: Stellar motion. (1) Here a snapshot of some stars are shown. Only their positions are discernible. (2) Here another snapshot of the same stars are shown but at a much later time (i.e. years). The movement is attributed to proper motion. (3) Here the apparent annual displacement, attributed to stellar parallax, is superimposed on the proper motion. (4) Lastly, the star in the bottom left is further perturbed by orbital motion (i.e. it is in an orbit with another star(s)). The star at the top may also be disturbed but the orbital motion is not discernible. This figure was based on Figure 1.3 in Perryman (2008) and Figure 1 in Reffert (2010).

Star	Position (J1991.25)		Proper motion [mas/yr]		π [mas]	d [pc]
	α [h m s]	δ [° ' "]	μ_α	μ_δ		
Proxima Centauri	14 29 47.747	-62 40 52.867	-3775.64 ± 1.52	768.16 ± 1.82	772.33 ± 2.42	1.3
Barnard's Star	17 57 48.966	+04 40 05.837	-797.84 ± 1.61	10326.93 ± 1.29	549.01 ± 1.58	1.8
Sirius A/B	06 45 09.250	-16 42 47.315	-546.01 ± 1.33	-1223.08 ± 1.24	379.21 ± 1.58	2.6
Ross 154	18 49 48.956	-23 50 08.752	637.55 ± 2.22	-192.47 ± 1.45	336.48 ± 1.82	3.0
61 Cygni	21 06 52.192	+38 44 03.890	4107.4 ± 0.43	3143.72 ± 0.59	285.42 ± 0.72	3.5
Groombridge 1830	11 52 55.816	+37 43 58.110	4003.69 ± 0.66	-5813 ± 0.44	109.21 ± 0.78	9.2
Arcturus	14 15 40.347	+19 11 14.172	-1093.45 ± 0.6	-1999.4 ± 0.52	88.85 ± 0.74	11.3
51 Pegasi	22 57 27.851	+20 46 07.263	208.06 ± 0.64	60.96 ± 0.57	65.1 ± 0.76	15.4
Gamma Draconis	17 56 36.378	+51 29 20.224	-8.52 ± 0.48	-23.05 ± 0.51	22.1 ± 0.46	45.2
Beta Hydri	11 52 54.561	-33 54 29.269	-58.02 ± 0.81	2.4 ± 0.67	8.93 ± 0.88	112
Beta Doradus	05 33 37.516	-62 29 23.481	1.06 ± 0.56	12.56 ± 0.66	3.14 ± 0.59	318
X Sgr	17 47 33.627	-27 49 50.745	-3.62 ± 1.01	-10.67 ± 0.71	3.03 ± 0.94	330

Table 1.1: A selection of objects and their astrometric parameters from the Hipparcos catalogue. Distances have been calculated from the Hipparcos parallaxes. As Hipparcos flew between 1989 and 1993, the epoch of observation is given as J1991.25.

reality, this is very hard to achieve and a number of corrections and transformations must be performed. Observations made from the ground have to compensate for the Earth's daily rotation, variations in the orientation of the rotation axis as well as the intrinsic properties of the instruments. The same compensations must be made from space: the rotation of the satellite, instrumental properties as well as a host of other factors that will add noise to the diagrams shown in Figure 1.2. For Gaia, radiation damage to the focal plane is one such factor that can significantly increase the uncertainty in the position (and flux) of an image, and must be corrected for before deriving astrometric parameters.

To give some examples of typical values for astrometric parameters consider Table 1.1. Here a number of objects taken from the Hipparcos catalogue (ESA (1997)) are listed and sorted in order of increasing distance from the Sun. These distances have been calculated from the Hipparcos parallaxes.

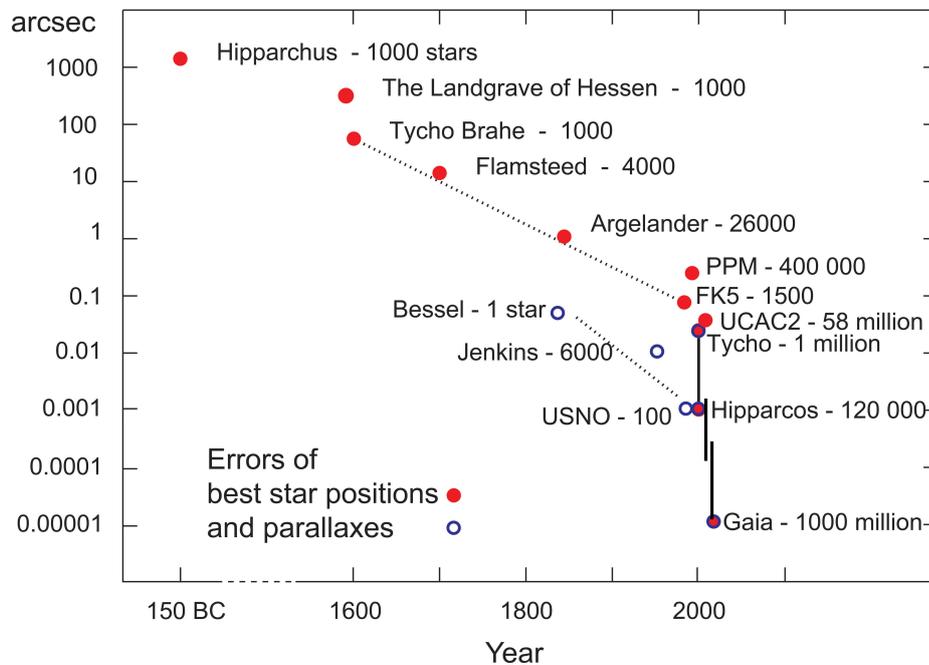


Figure 1.3: Improvements in astrometric accuracy. Credit: Erik Høg (ESA (2010b)).

As one can see, with Hipparcos it is possible to quote positions, proper motions and parallaxes and their errors at the mas level. To probe further and carry out a survey on a Galactic scale, measurements at the microsecond level would be required.

To conclude this section, Figure 1.3 puts the history of astrometry in context. From the time of Hipparchus (c. 190 BC - c. 120 BC), the Greek astronomer, to Hipparcos, positional and parallax accuracies have improved by 6 orders of magnitude. Gaia is expected to reach 1 order of magnitude better even than Hipparcos.

1.2 Photometry

Before moving on to discuss the merits of performing astrometry from space, another cornerstone of astronomy will now be introduced, namely photometry. While astrometry focuses on the geometric and kinematic properties of a star, photometry focuses on measuring the magnitude or brightness of a star, which in practice means measuring the amount of radiant flux received. As detectors tend to have different sensitivities to different wavebands it is common to use filters to allow only certain wavelengths through. For example, the Johnson-Cousins $UBV(RI)_C$ photometric system measures a star's apparent magnitude through five filters (values taken from Bessell (1990)):

- U : this measures a star's apparent ultraviolet magnitude. The filter is centred at 370nm and has an effective bandwidth of about 100nm.
- B : this measures a star's apparent blue magnitude. The filter is centred at 420nm and has an effective bandwidth of about 200nm.
- V : this measures a star's apparent visual magnitude. The filter is centred at 530nm and has an effective bandwidth of about 200nm.
- R : this measures a star's apparent red magnitude. The filter is centred at 600nm and has an effective bandwidth of about 250nm.
- I : this measures a star's apparent infrared magnitude. The filter is centred at 805nm and has an effective bandwidth of about 200nm.

Another example is the *ugriz* Sloan photometric system. Gaia will be recording 'white' light over a broad band in the instrument that focuses on astrometric measurements. For this a G band magnitude (G) is defined. The G band is similar to the V band as defined above. However, due to the inherent properties of the detector, there are some slight differences.

As with astrometry, any system taking photometric measurements requires a reference. For this purpose, the overall standard is Vega (α Lyrae), which is fixed to have zero magnitude in the photometric system.

Absolute magnitudes may be found by determining the distance to a star, which as discussed earlier can be provided by measuring the parallax of the star. Ideally, the distance would also be an absolute

measurement and for this an absolute parallax would be required. The difference between relative and absolute parallax is described in the next section, as performing absolute astrometry is one of the key reasons for moving an observatory from the ground into space.

Combining the measurements from astrometry, photometry and spectrometry can help build up a detailed picture of a star's geometric, kinematic and intrinsic properties. Such parameters include the star's effective temperature, surface gravity, composition and line-of-sight extinction. From these values it is possible to determine the masses and ages of the stars, which ultimately can lead to a description of the composition, formation and evolution of the Galaxy - Gaia's primary objective.

1.3 Ground vs space-based astrometry

In the 21st century, with the option of conducting astrometric studies from space as well as from the ground many would naturally ask which is better. In other words what are the advantages of wide-angle (space-based) astrometry over narrow-angle (ground-based) astrometry? Narrow-angle astrometry typically involves taking several images of the same patch of the sky at different times of the year, combining these images by cross-matching the positions of the background star(s), and measuring the proper motion and parallax of any stars that are relatively close by. If the parallax is large compared to the parallax of the background star then this relative measurement will be close to the absolute value. However, without being able to decouple the parallax of the background star from the image, this technique will always give a relative or differential measurement. If the star of interest, say the target star, and background star could be observed in two very different parts of the sky at the same time then the parallax of the background object could, in principle, be decoupled from the parallax of the target star. This principle is illustrated in Figure 1.4.

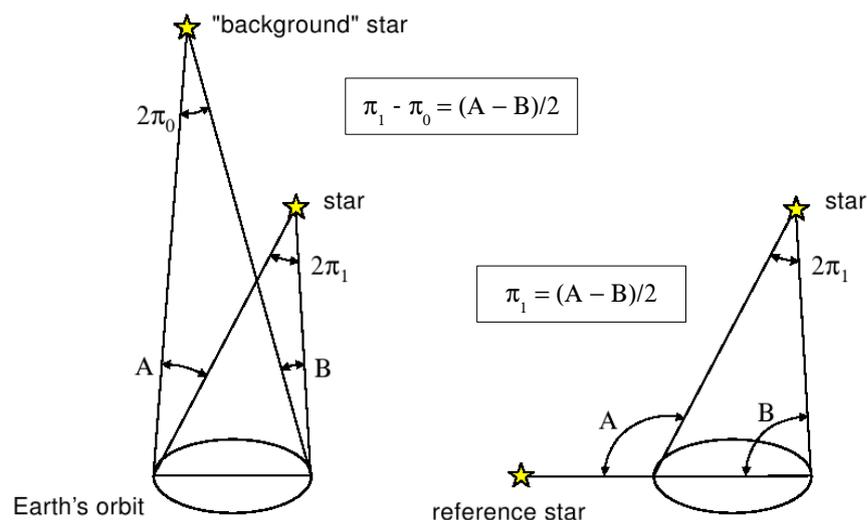


Figure 1.4: The principle of absolute parallax measurement. On the left, measuring small angles A and B provides the relative parallax $\pi_1 - \pi_0 = (A - B)/2$. On the right, measuring large angles A and B provides the absolute parallax $\pi_1 = (A - B)/2$ of the star. Credit: Lindegren (2005).

As the position of the reference star is measured not to change when observed at two diametrically opposite ends of the Earth's orbit, the parallax of the target star no longer depends on the parallax of the reference star. Consequently, with wide-angle astrometry it becomes possible to measure absolute parallaxes. Is it possible to perform wide-angle astrometry from the ground? Not without overcoming the effects of refraction imposed by the Earth's atmosphere, which can dominate the error budget in astrometric ground-based programmes. Due to turbulence in the Earth's atmosphere, light from an astronomical source can be refracted in almost random directions, blurring its observed image. The quality of this image at a given observing location and time is referred to as (astronomical) seeing. From space, astronomical seeing no longer limits the quality of an image. Furthermore, from space the whole sky is visible; thermal and gravitational stability can be far higher, and the level of background light far lower. Both fields-of-view can be accommodated onto a single satellite, which can scan the sky to mimic the Earth's rotation. In other words, space-based observatories are ideal for performing global, absolute astrometry. Hipparcos was the first to realise this concept and Gaia will soon follow in its footsteps.

Precision differential astrometry from the ground (see e.g. Cameron et al. (2009) and Bellini et al. (2009)) as well as from space with the Hubble Space Telescope (HST) (see e.g. Piatek et al. (2007)) is still a valuable tool. Ultimately, ground observations provide the initial input catalogues for space programmes, allowing a very good first estimate to be made of the positions (and fluxes) of the calibration stars in the subsequent data reduction. Often, the results of space programmes feed into detailed follow-up analyses from the ground. For example, Gaia will generate science alerts when "interesting" objects are identified in the data stream, such as supernovae, gamma-ray burst optical counterparts, microlensing events, that could then be followed up from the ground.

1.4 The legacy of Hipparcos

The ability to perform global astrometry from space was successfully demonstrated by Hipparcos, an ESA mission which flew between August 1989 and May 1993. The Hipparcos Catalogue contains ~ 118000 entries and is complete to $V = 7.3$. Formal errors are within 1 to 3 mas (ESA (1997)). Accompanying the Hipparcos Catalogue was the Tycho-1 Catalogue, which contained the positions, proper motions and photometry on $\sim 10^6$ stars. Since its publication Tycho-1 has been updated to the Tycho-2 catalogue (Høg et al. (2000)), incorporating improved data analysis techniques and other complementary datasets. In addition, a new reduction of the raw Hipparcos data has been made (van Leeuwen (2007)), identifying a number of anomalous features in the original analysis. With this new reduction the formal errors have been reduced to < 1 mas for stars of $V < 8$.

In many ways Hipparcos laid the groundwork for Gaia. It was able to make large-angle measurements at high precision by combining the light from two fields-of-view pointing at different parts of the sky using a beam combiner, and through a succession of reflections direct the light onto a single focal plane. It also adopted a scanning law which traded-off the need to maximise the parallax in the scanning direction with the need to maintain thermal stability as well as scan the sky as many times as possible. This blueprint will be discussed again shortly. Figure 1.5 shows a schematic of the Hipparcos telescope. For a thorough description of the Hipparcos payload, the mission and the data see ESA (1997).

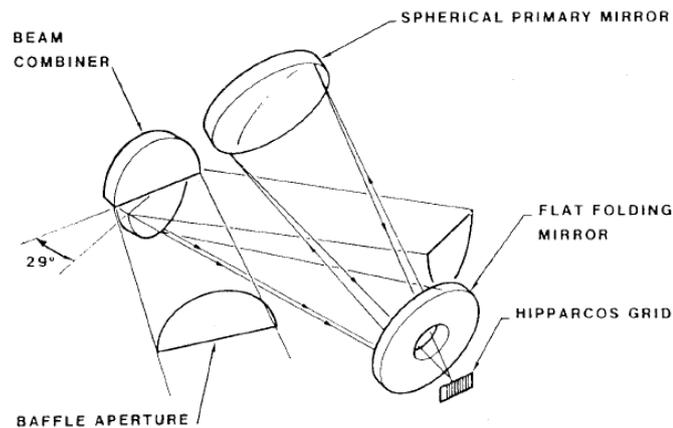


Figure 1.5: Schematic of the Hipparcos telescope. Light entered from two baffle directions and was brought together at the beam combiner. The combined light was reflected from the flat folding and spherical primary mirrors onto the focal surface where the modulating grids were located. Credit: ESA (1997), Volume 2, Ch. 2.

1.5 Introducing Gaia: 'The Billion Star Surveyor'

Hipparcos was designed so that it could perform astrometry at the milliarcsecond level. This led to a host of astronomical applications (Perryman (2008)). Gaia is being designed to take the next step, expecting to achieve microarcsecond (μas) astrometry down to $V = 20$ -22. This is equivalent to a limiting magnitude of $G = 20$ in the Gaia 'G' band, which measures unfiltered (white) light from ~ 350 to 1000nm . Table 1.3 gives the current requirements on the end-of-mission parallax standard errors, averaged over the sky, for unreddened B1V, G2V, and M6V stars in the V band.

	B1V	G2V	M6V
$V < 10$ mag	$< 7 \mu\text{as}$	$< 7 \mu\text{as}$	$< 7 \mu\text{as}$
$V = 15$ mag	$< 25 \mu\text{as}$	$< 24 \mu\text{as}$	$< 12 \mu\text{as}$
$V = 20$ mag	$< 300 \mu\text{as}$	$< 300 \mu\text{as}$	$< 100 \mu\text{as}$

Table 1.2: Gaia requirements on end-of-mission parallax standard errors. Credit: ESA (2010a)

By observing every object down to $G = 20$ it is expected that Gaia will identify at least 1 billion stars as well as an unprecedented number of other celestial bodies. For example, between 10^6 and 10^7 galaxies, ~ 500000 quasars, 10^5 to 10^6 solar system bodies and tens of thousands of exoplanets. Clearly such numbers are bound to have an impact in many areas of astronomy. Figure 1.6 illustrates this by pointing out the depth of the Gaia census.

Since its inception Gaia has gone through various stages of development. After a two-year concept and technology study (Perryman et al. (2001)) detailing its science return and feasibility, Gaia was approved by ESA in 2000 as a cornerstone mission. Following an intense definition/study phase, which included in 2002 a cost-reduction exercise leading to a change of launcher (Ariane 5 to Soyuz) and a redesign of the satellite, the contract for building Gaia was awarded to the European Aeronautic Defence

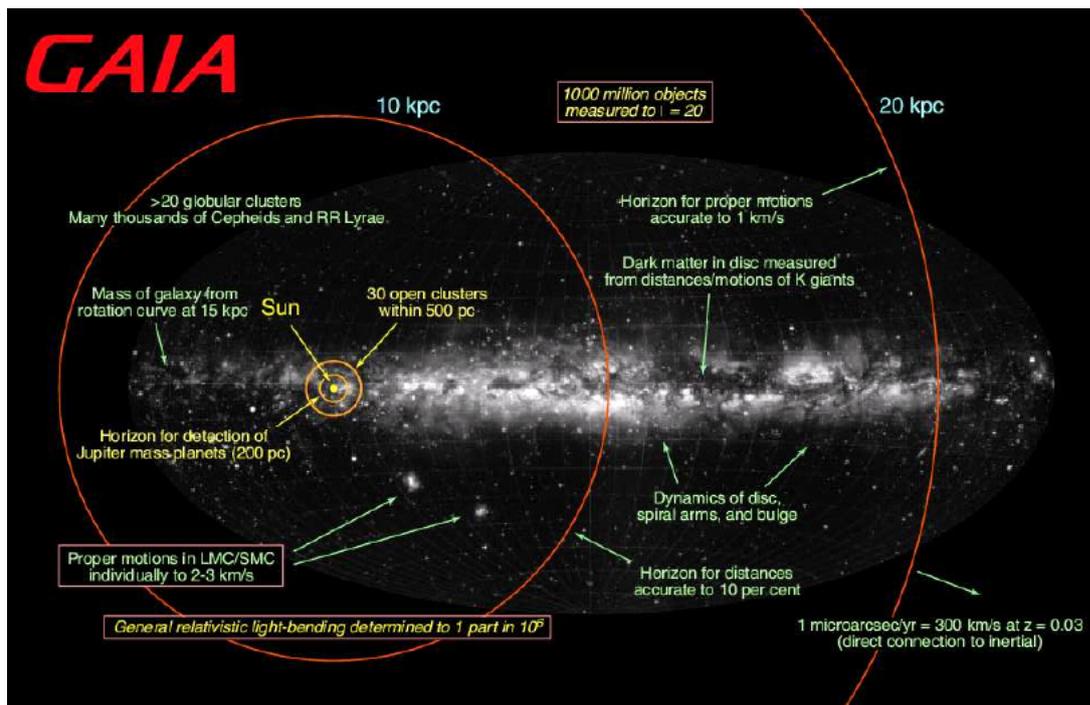


Figure 1.6: A Galactic census. Gaia will be able to observe objects throughout the Milky Way Galaxy. The background to this schematic is the Lund sky map. Credit: Lund Observatory (ESA (2010b)).

and Space company (EADS) Astrium in 2006. In June 2007 the Preliminary Design Review (PDR) was completed and now the project is in Phase C/D; many parts of the spacecraft design have been frozen and several parts of the payload have already been built. Gaia is due to launch in November 2012 and operate for 5-6 years. During this time there will be a continuous effort on ground to process the data, culminating in a final data release around 2020-21.

1.5.1 The payload

Two schematics of the Gaia payload are shown in Figure 1.7. It can be seen that the recipe for performing wide-angle astrometry as demonstrated by Hipparcos has also been adopted here. Two viewports, both measuring $1.45 \times 0.5 \text{ m}^2$, are backed by two large primary mirrors (M1/M'1). These are the first in a sequence of mirrors that direct light down into the telescope. When the light from both Field of Views (FOVs) are projected onto mirrors M4/M'4 the beams are combined so that the two fields can be superimposed onto a single focal plane. The focal length is 35m, which has been folded into an optical bench 3m wide.

The focal plane contains three instruments, which in sequence measure the astrometric, photometric and spectroscopic properties of the incoming light. For the astrometry and broadband G photometry no filter is required; the light is projected immediately onto the focal plane. For the filtered photometry, two low-resolution fused-silica prisms are used to disperse the light into a spectrum in the scanning direction of the satellite. Gaia's scanning law is described in section 1.5.3. One of the prisms, the Blue Photometer (BP), has an effective bandwidth range of 330-680nm; the other, the Red Photometer (RP), operates over

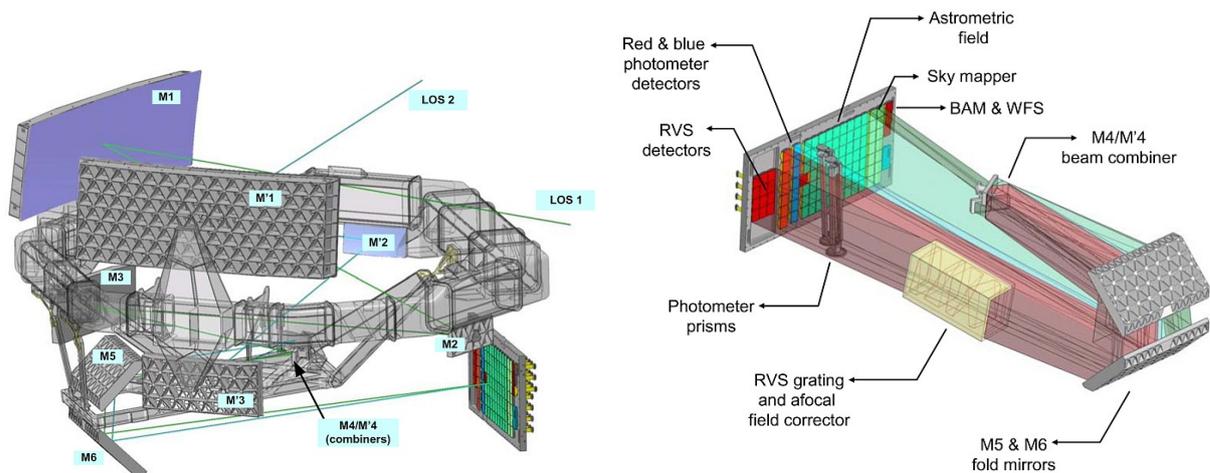


Figure 1.7: The Gaia payload. (left) The light from two different parts of the sky is directed down into the telescope via a sequence of mirrors and projected onto a single focal plane. (right) Once the beams have been combined they are projected onto a focal plane that contains three different instruments. Credit: ESA (2010b).

640-1000nm. By extracting an object's Spectral Energy Distribution (SED), the photometers can correct the astrometric measurements for chromaticity effects and derive astrophysical parameters.

Next to the photometers is the Radial Velocity Spectrometer (RVS). This instrument is a near infrared (847-874nm), medium-resolution ($R \sim 11500$), integral-field spectrograph that disperses light onto the focal plane again in the scanning direction. The actual dispersion is carried out by an optical module which contains a grating plate, a filter plate, and four fused-silica lenses.

1.5.2 The focal plane

The Focal Plane Assembly (FPA) contains a large array of CCDs, which rely on the photoelectric effect to quantify the intensity of the incoming beam of light. The amount of charge generated is proportional to the brightness of the star. Gaia is a scanning satellite, so the CCDs are set to operate in Time-Delay Integration (TDI) mode. This technique matches the spin rate to the pixel clock rate so as an object passes over a CCD the charge that is generated moves with the stellar image until it reaches the end of the CCD and is read out. Figure 1.8 shows a schematic of the focal plane and how it is divided up into several areas. As one can see, the first two complete strips of CCDs constitute the Sky Mappers (SMs). These are used to detect objects as they appear on the focal plane, define a transit time and ordinate position, and communicate this to the CCDs that follow. They also estimate the intensity of the object so that an appropriate window may be assigned to it. SM1 flags the objects that originate from the first FOV, SM2 the second FOV. As there is a limit to the amount of data that can be telemetered back to the ground, Gaia has a windowing scheme that assigns a window around an image. Any data outside this window is flushed out of the data stream. The size of the window depends on the intensity of the image, with fainter stars most of the time receiving smaller windows.

Once the object has passed over the SMs it is picked up by the Astrometric Field (AF). This field constitutes the bulk of the FPA: 62 CCDs out of a total of 106. It is responsible for providing the

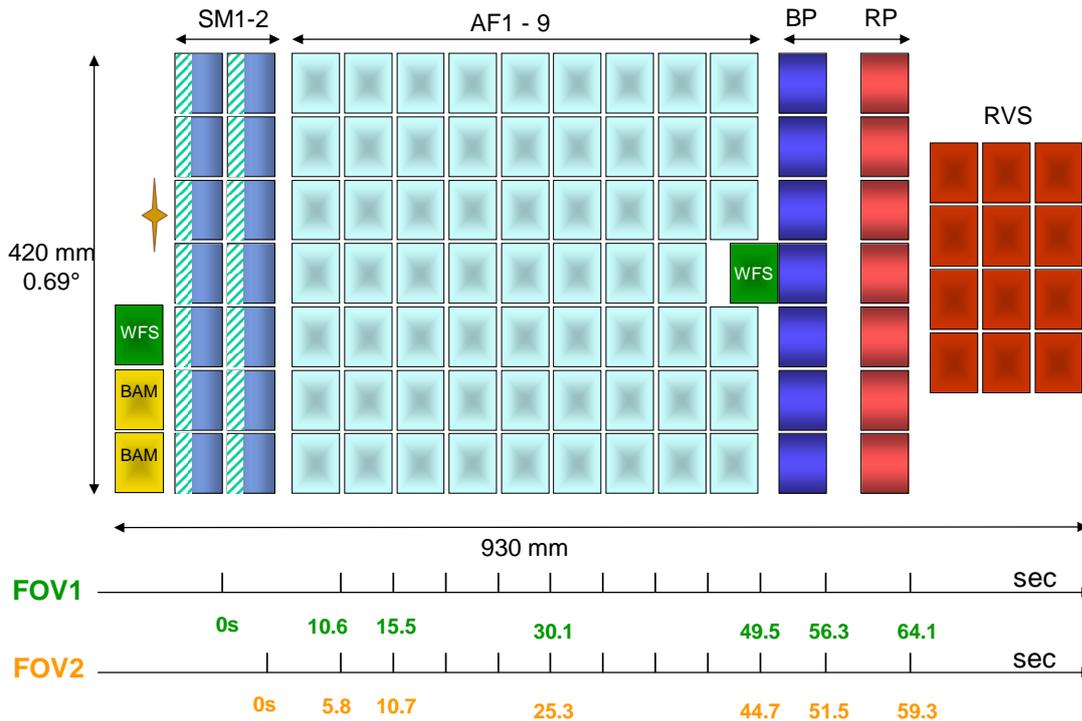


Figure 1.8: The Gaia focal plane. As a star (indicated) traverses the focal plane from left to right it crosses a number of CCDs. The Sky Mappers (SM1 for the first field-of-view (FOV), SM2 for the second FOV) detect the star, define a time and AC position and communicate this to the CCDs that follow. The Astrometric Field (AF) provides the astrometric and broadband photometric data. The Blue and Red Photometers (BP/RP) extract the stars Spectral Energy Distribution (SED) and lastly, the Radial Velocity Spectrometer (RVS) is used to determine the star's radial velocity. A few more CCDs are situated on the focal plane: the Wave-Front Sensor (WFS) monitors optical alignment along both FOVs and the Basic Angle Monitor (BAM) monitors the variation in the basic angle. Credit: Mignard & Drimmel (FM-030).

astrometric and broadband photometric data. AF1 confirms the presence of a stellar image that was identified in the SM. As the satellite spins the image moves over the photometers (BP/RP) and finally the RVS. There are four other CCDs on the focal plane that are used for monitoring purposes. The Wave-Front Sensor (WFS), of which there are two, monitors the alignment of optical elements along both lines-of-sight. The Basic-Angle Monitor (BAM), again of which there are two, monitors the stability of the basic angle. This is the angle separating the two FOVs and is one of the geometric parameters that must be known very precisely in order to achieve the accuracy levels Gaia is aiming for. For this reason, the angle must remain extremely stable and will be monitored continuously. It is expected not to deviate by more than $4\mu\text{as}$ throughout the mission³.

To resolve as many stars as possible each CCD contains 8847000 pixels: 4500 in the along-scan (AL) direction and 1966 in the across-scan (AC) direction. Each pixel measures $10\mu\text{m}$ AL \times $30\mu\text{m}$ AC, so with a plate scale of $170\mu\text{m}/\text{arcsec}$, this corresponds to a $59\text{mas} \times 177\text{mas}$ patch on the sky. To achieve microarcsecond accuracy levels, for example, for a G2V star at $V = 10$, the equivalent error on the focal plane at the end-of-mission would have to be 1.2×10^{-4} pixels. If, on average, such an object is

³This value has been taken from the Gaia Parameter Database (GPD) from the entry labelled 'Satellite: BasicAngle_Stability_RequirementSystematic'.

observed ~ 75 times in 9 CCDs, then the actual centroiding error per transit would have to be 30×10^{-4} pixels (Credit: Alex Short, c.f. Brown (2009)). Gaia places strict constraints on centroiding accuracies, which in turn places strict constraints on the methods used to acquire these accuracies. As will be shown, radiation damage, which can degrade this requirement significantly if left unchecked, has been one of the major challenges facing Gaia.

1.5.3 The scanning law

For Gaia to maximise end-of-mission astrometric accuracies, observe the entire sky and continually calibrate the basic angle, a scanning law has been devised that determines how the spin axis evolves with time. There are two conditions that the scanning law must meet: 1) maximise the Solar Aspect Angle (SAA) (i.e. the angle between the spin axis and the Sun); and, 2) maximise the uniformity of the sky coverage over the course of the mission. As Gaia spins each FOV traces out a great circle on the sky (see Figure 1.9). The along-scan component of the parallax, parallel to the scan direction along the great circle, is proportional to the sine of the SAA, so maximising the SAA maximises the parallax measurement. For Gaia there is a limit to the SAA, which is fixed to 45° , as stray light must not enter the telescope and the solar panels need to be oriented to receive sufficient energy from the Sun. As the SAA is fixed, full-sky visibility is achieved by letting the spin axis precess around the direction of the Sun. For Gaia, the spin axis returns to its original position every ~ 63 days to a) ensure consecutive great-circle scans overlap enough to bring stability to the observations and b) allow the same patch of sky to be observed at least twice every 6 months.

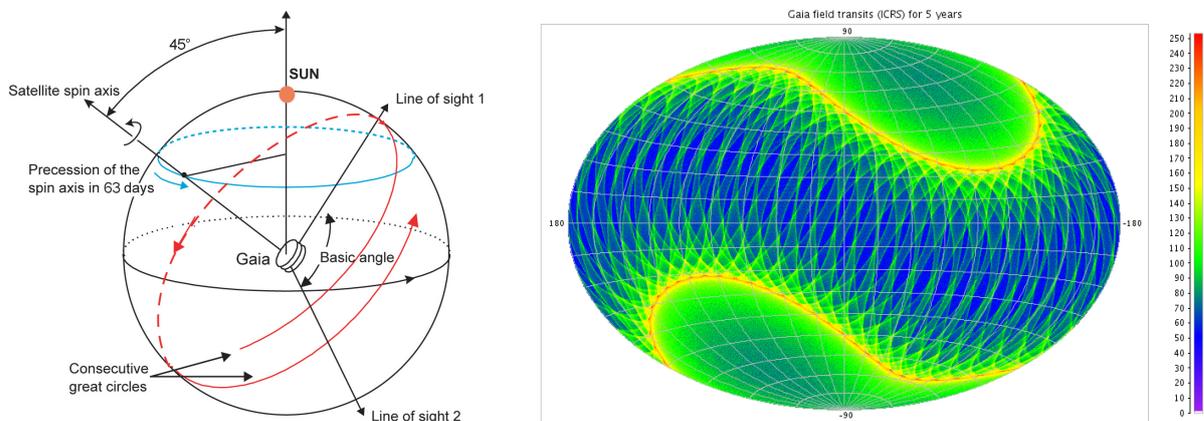


Figure 1.9: (left) The Gaia scanning law. The spin axis evolves such that consecutive great-circle scans overlap enough to bring stability to the observations and allow the same patch of sky to be observed as often as possible (Credit:ESA (2010b)). (right) End-of-mission sky coverage (given in galactic coordinates). The colours indicate the number of times each part of the sky is observed. On average, each object will be observed ~ 75 times. (Credit:Berry Holl (Lund University)).

To calibrate the astrometric measurements the basic angle, which is fixed to 106.5° , must be highly stable and variations must be accurately measured continuously. For this purpose the spin rate has been fixed to 60 arcsec/s to allow great-circle scans to overlap; as the satellite spins what one FOV sees will be seen soon after by the second FOV. With this spin rate Gaia will take 6 hours to complete one revolution.

Figure 1.9 (left) illustrates how the scanning law has been carefully designed to achieve the necessary conditions mentioned above. Figure 1.9 (right) shows the result of a simulation which traces the path of every great circle throughout the course of the mission and tracks the number of times each part of the sky is observed. On average, each object will be observed ~ 75 times.

1.6 Determining astrometric parameters

Transforming observations on the focal plane to a set of astrometric parameters in the ICRF is a complex procedure, requiring the same data to be used at various levels to converge on a final solution. The observed location (centroid) of an image on the focal plane is modelled as follows:

$$\begin{pmatrix} \text{Observed} \\ \text{Centroid} \end{pmatrix} = \begin{pmatrix} \text{Global ref.} \\ \text{frame} \end{pmatrix} + \begin{pmatrix} \text{Source} \\ \text{position} \end{pmatrix} + \begin{pmatrix} \text{Instrument} \\ \text{Attitude} \end{pmatrix} + \begin{pmatrix} \text{CCD/pixel} \\ \text{offset} \end{pmatrix} + \text{noise} \quad (1.2)$$

The global reference frame takes account of large scale perturbations to the centroid, e.g. light-bending, and the CCD/pixel offsets take account of the geometric calibration (i.e. the instrument properties), chromaticity and charge transfer inefficiency effects. Source position describes the centroid in terms of the astrometric parameters and the instrument attitude determines where the satellite was pointing at the time of observation.

More specifically, on the focal plane the image centroid is given in terms of a) an AL and AC pixel coordinate assigned to a particular CCD and b) a transit time defined by the on-board clock. Through modelling the instrument the coordinate on the focal plane can be mapped over to a coordinate on the FOV. This coordinate is given in terms of field angles assigned to a particular FOV. By taking into account all of the large and small scale effects that can arise at this stage (e.g. Point Spread Function (PSF) dependencies on CCD, FOV, AC position and colour, and how the actual image shape may be distorted as a result of radiation damage, AC motion smear and flux loss outside the window), every centroid regardless of where it is in the focal plane can be placed into a single instrument reference system defined by the field angles. Mapping the field angles onto the sky relative to the satellite requires knowing where the satellite was pointing at the time of observation. This is given by modelling the satellite attitude. Finally, the image position may be transformed over to a barycentric coordinate system after correcting for the satellite's orbit, stellar aberration and the bending of light around large objects (Mignard (GAIA-FM-023)).

Gaia is a self-calibrating satellite, so it is important to realise that the observations which are used to derive astrometric and photometric parameters are also used to calibrate the same observations. In practical terms it is the Astrometric Global Iterative Solution (AGIS) that carries out these calibrations (see e.g. Lindegren (GAIA-LL-034) and Lindegren (GAIA-LL-059)). For example, when solving for the instrument parameters it assumes the other parameters (attitude, astrometric, etc) are known. The instrument parameters are then used to solve for the attitude parameters and so on. This iterative procedure continues throughout the course of the mission until a global solution is found. Only after a list of

relative, internally consistent, astrometric parameters for every object observed has been generated is the transfer over to the ICRF made via the observations of astrometric standards.

1.7 Determining photometric parameters

There are two instruments on board Gaia that deal with photometry: the AF and BP/RP. Photometry on the AF is concerned with white light or broadband photometry. Photometry on the BP/RP CCDs are concerned with the SED towards the blue and red ends of the visual spectrum. Here, only photometric parameters derived from broadband photometry will be discussed.

For every image read out of the AF, a flux or magnitude must be extracted. This is performed by fitting either a PSF or a Line Spread Function (LSF) that has been corrected for charge transfer inefficiency (CTI) to the observed image shape/profile. Whether a PSF or an LSF is chosen depends on whether the charge from several pixels was binned in the AC direction on being read out of the CCD. In other words the LSF is a PSF that has been binned in the AC direction.

The PSF or LSF is an idealised image of a point source that has been spread out by the telescope optics before intercepting the focal plane. Its shape is expected to vary as a function of position in the focal plane, and will depend on the colour of the star. Predicting what the PSF or LSF will look like is one of the steps that is required to model the distortion of the image after it has been damaged by radiation. The remaining part is carried out by the CTI correction; it corrects the predicted PSF/LSF to produce a prediction of the damaged image shape which can then be compared to the actual damaged image shape at the time of observation. How the CTI correction is carried out will be described in section 1.9.4. Once a suitable fit has been found the image parameters can be extracted.

The centroid of an image profile and its flux are determined simultaneously in the data processing, so the accuracy of the formal errors on the astrometric and photometric parameters will depend on finding a "good" representation of the image shape at the time of observation. To build a library of predicted PSF/LSFs, calibration standards (i.e. well-behaved stars) are picked from the Gaia data stream and used to define the PSF/LSF for that calibration period. How long this PSF/LSF is valid for will depend on how quickly the conditions on the focal plane change. For any particular period of time, once a calibrated PSF/LSF is found it can be used to fit the observations with the current CTI image correction parameters. Following further observations of calibration standards, the PSF/LSF can be recalibrated and, as with the astrometric approach, the older observations can also be refitted. Convergence is achieved once the residuals remaining between the fits have become sufficiently randomised.

1.8 Processing the data: from the raw data to a Gaia Catalogue

Determining astrometric and spectrophotometric parameters is part of the data processing, which begins once Gaia has started observing. Here a general overview of this process is given.

Gaia will send on average ~27 Gigabytes of compressed data to ground every day (EADS Astrium (GAIA.ASF.TCN.SAT.00133)). This will, for most of the time, be received by an ESA ground station

at Cebreros in Spain. Handling this volume of data and processing it in a reasonable time is a huge software engineering challenge that must be carried out autonomously. To distribute the effort, Table 1.3 shows how a number of Coordination Units (CUs) have been set up to organise the data processing. Each CU has a distinct responsibility and must communicate through clearly defined interfaces. For example, CU5 must provide an estimate of a star's colour before CU3 can correct each observation for chromaticity effects. Governing this entire process is the Data Processing and Analysis Consortium (DPAC). DPAC currently consists of more than 350 scientists, not including those working within industry and is spread across many European institutions. From the ground station the data is channelled through the European Space Astronomy Centre (ESAC) to the various CUs in order to perform their part of the analysis.

CU	Name	Responsibility
0	Management Unit	Overall coordination and running of the DPAC.
1	System Architecture	Help define the overall system architecture. Provide a central database.
2	Data Simulations	Provide input data for simulating reduction software: GIBIS, GASS, GOG.
3	Core Processing	Handle core data reduction and astrometric data, e.g. IDT, AGIS, IDU.
4	Object Processing	Process all ill-behaved objects from CU3, CU5, or CU6.
5	Photometric Processing	Process all photometric data from SM, AF, BP and RP CCDs.
6	Spectroscopic Processing	Process all spectroscopic data from the RVS.
7	Variability Processing	Characterise the photometric and spectral variability.
8	Astrophysical Parameters	Classify objects and estimate their astrophysical parameters.
9	Catalogue Access	Make the final scientific catalogues available to the scientific community.

Table 1.3: DPAC Coordination Units.

In Table 1.3 it can be seen that CU3 is in charge of handling the major data reduction packages such as Initial Data Treatment (IDT), AGIS and Intermediate Data Update (IDU). IDT covers one of the first parts of the data processing, after the telemetry has been ingested, and turns the raw data into 'elementary data objects'. Some of the outputs at this stage include a detection time, a pair of field angles and a background estimation. AGIS carries out many of the processes discussed in section 1.6; in order to map the location of an image on the CCD at the time of observation onto the sky, AGIS iterates through a set of parameters (astrometric, satellite attitude, instrument, global, plus noise) until a global solution has been found. For a scaled-down working version of AGIS (called AGISLab) see Holl et al. (BH-002). After every six months of data collecting, all of the data that has been received up to that point is analysed again as part of IDU. This takes into account the improved calibration parameters that have been found as the number of observations increases.

CU2 provides the input for simulating the data processing pipeline. It handles three packages: the Gaia Instrument and Basic Image Simulator (GIBIS), the Gaia System Simulator (GASS) and the Gaia Object Generator (GOG). GIBIS generates data at the pixel-level. GASS simulates the telemetry stream and GOG produces data in various intermediate forms throughout the data processing. For more information on DPAC and the CUs refer to the GaiaWiki site⁴.

⁴See <http://www.rssd.esa.int/wikiSI/index.php>. Note that a username and password are required to access this site.

1.9 Radiation-induced Charge Transfer Inefficiency: a major challenge

The transfer of charge through a CCD is not without losses. Charge Transfer Inefficiency (CTI) gives a measure of the percentage of charge that is lost between exposure and readout. Modern values can be as low as $\sim 0.0001\%$. However, radiation damage can increase this value significantly.

Radiation damage to CCDs can appear in a variety of forms (see e.g. Howell (2006)). In some cases high-energy photons can produce fast electrons, which can then go on to cause simple, localised defects. In others, high-energy particles can knock atoms out of the CCD lattice and leave behind vacancy-interstitial pairs⁵. These vacancies can move around in the lattice until they encounter doping elements such as phosphorus atoms⁶, whereupon stable states can form. These states give rise to electron traps, which have energy levels between the conduction and valence bands of the silicon. When electrons in the conduction band fall into these traps CTI increases. For Gaia, at the L2 point, high-energy protons (several hundred MeV) emanating from the Sun can cause this type of damage.

More specifically, Gaia will operate in a Lissajous halo orbit around the L2 point, 1.5 million km from Earth outside the geomagnetic field. At this location the energetic particle environment is expected to consist of Galactic Cosmic Rays (GCR), solar protons and solar heavy ions. GCR originate from outside the Solar System and though their rates are quite low (4 to 8 particles $\text{cm}^{-2}\text{s}^{-1}$ depending on the phase of the solar cycle), it is almost impossible to shield against them as they have energies typically in excess of hundreds of MeV. Particles from the Sun on the other hand can be shielded against as long as the Sun remains quiet. When the Sun is fully active coronal mass ejections and solar flares can produce thousands of highly energetic protons $\text{cm}^{-2}\text{s}^{-1}$. If Gaia were to encounter such a maelstrom, the long term damage would be significant.

Figure 1.10 illustrates the processes that can occur when a high-energy proton passes through a CCD. There are three distinct types of damage (see e.g. Hopkinson et al. (1996)):

1. Total Ionizing Dose Effects: this involves mainly the buildup of trapped charge in the oxide (i.e. the insulator or gate dielectric) and the generation of traps at the silicon dioxide/silicon interface.
2. Transient Effects: these are non-permanent effects (the excess charge is removed during readout) and are due to the ionization-induced generation of charge within the silicon epitaxial layer. As the proton passes through the silicon, electron-hole (e-h) pairs are left behind in the protons wake.
3. Displacement Damage: even though most of the energy from the irradiation goes into producing e-h pairs (ionization), a small fraction goes into displacing atoms from their lattice sites and creating vacancy-interstitial pairs. This fraction is called the nonionizing energy loss (NIEL). More than 90% recombine, but those that don't form stable defects that have energy levels between the valence and conduction bands. As an image passes over these defects electrons can fall into the lower energy levels, become trapped, and get left behind as the image continues to be read across

⁵The vacancy is where the original atom in the lattice used to be. The interstitial is a "hole" in the lattice where the atom now sits.

⁶When a vacancy (V) forms a stable state with a phosphorus (P) atom a P-V centre is formed. This also known as an E-centre.

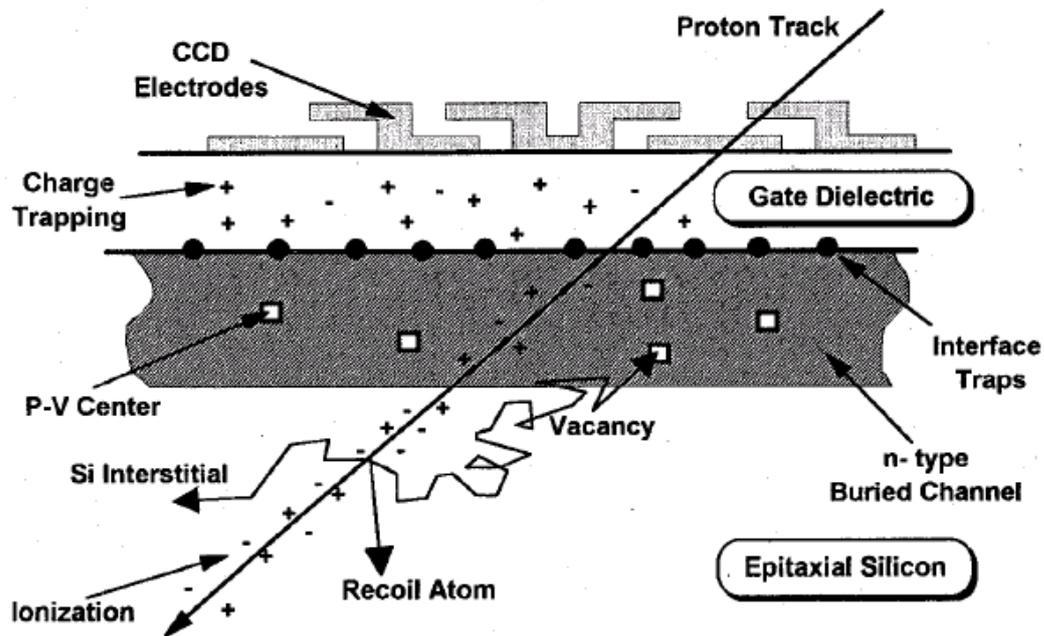


Figure 1.10: CCD proton damage. When a highly energetic proton passes through a CCD a number of events can occur: charge trapping in the gate, interface trap generation (ionization damage), ionization in the silicon (transient effects), and displacement damage in the silicon. Displacement damage can produce vacancy/interstitial pairs, which when combined with doping elements, e.g. phosphorus, create traps (the P-V centre is one such trap). The traps in turn increase CTI and lower the performance of the affected pixel(s). Figure credit: Hopkinson et al. (1996)

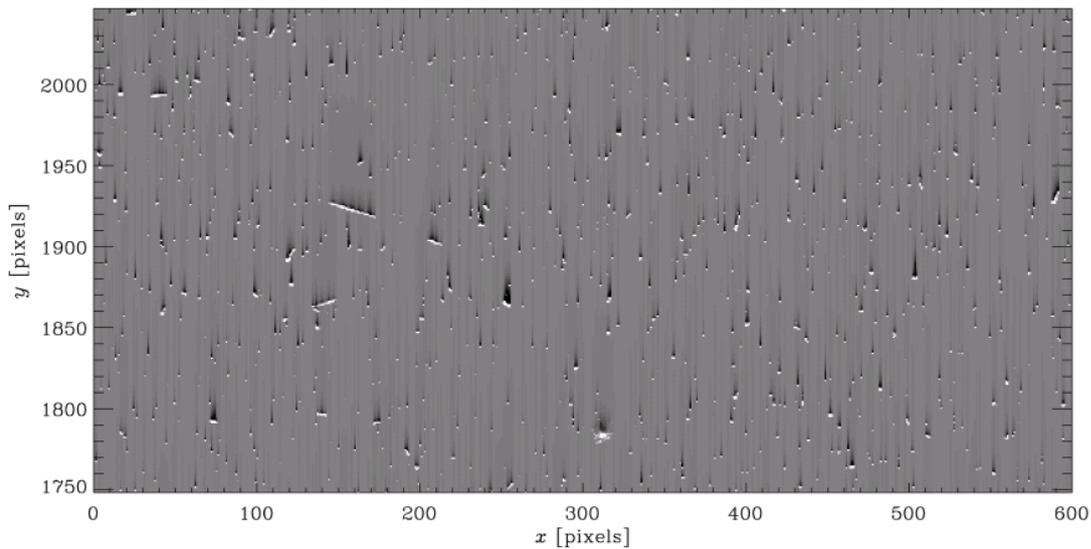


Figure 1.11: An example of the effects of CTI on a Hubble image. This is a difference image that has been formed by subtracting a CTI corrected image from some raw Hubble data. The charge release trails can clearly be seen behind many of the stars in the image as it is read out. Figure credit: Massey et al. (2010)

the CCD. As the electrons are released at a later time, they can manifest as a charge trail behind the image. An example of this is shown in Figure 1.11.

The buried channel is part of the silicon substrate which channels electrons through the CCD to the readout register. CTI increases when traps within the buried channel capture electrons from passing images and release them back into the conduction band at a later time. As an increase in the number of protons passing through the CCD also increases the density of traps, CTI depends upon the how active the Sun is and this is linked to the solar cycle. Gaia will take about one month to reach the L2 point, so there is a chance the CCDs will be damaged before operations begin.

1.9.1 The solar cycle

In order to predict the level of radiation damage Gaia will receive by the end-of-life, a radiation analysis was performed by Astrium, where they used a proton fluence model (JPL-1991 model) to take into account the future behaviour of the solar cycle (EADS Astrium (GAIA.ASF.TCN.SAT.00035)). The current launch date for Gaia is November 2012. However, when the document was initially written the launch date was December 2011. With this earlier date Astrium predicted that Gaia would fly for 3.5 years during solar maximum (predicted to occur between mid-2008 and mid-2015) and the remaining years during solar minimum. Rounding the length of the solar maximum up to 4 years (to allow for

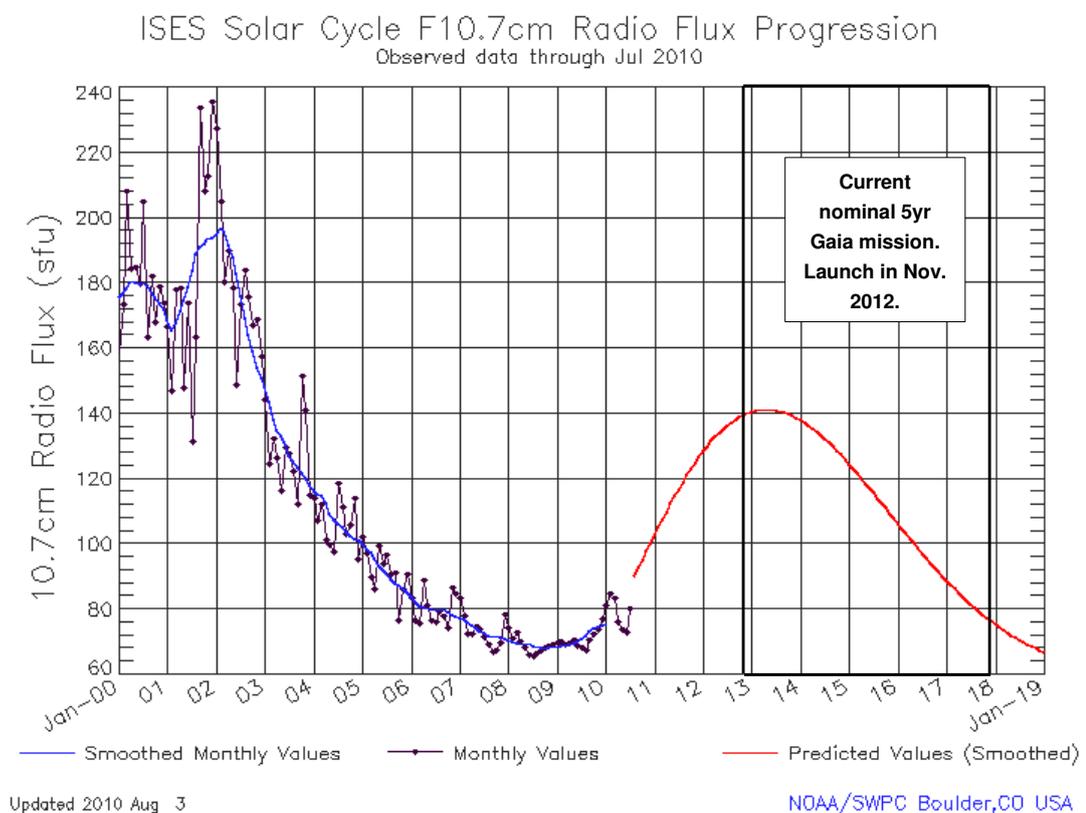


Figure 1.12: Solar Cycle 24. The Gaia mission period is overlaid. Credit: National Oceanic and Atmospheric Administration (NOAA)/Space Weather Prediction Center.

uncertainty in the prediction), the expected end-of-life mean level of radiation damage on the AF was estimated to be $4.8 [4.1 ; 5.3] \times 10^9$ 10MeV proton/cm². The minimum and maximum predicted levels are shown in the square brackets.

Since this time the observed solar cycle has altered expectations. Figure 1.12 shows the current data that has been collected by the National Oceanic and Atmospheric Administration (NOAA) on the solar cycle as well as a projection of future activity. The Gaia mission, now with a nominal launch date of November 2012, has been overlaid. As one can see, since the maximum in the solar cycle is expected to be around May 2013, the amount of time Gaia will spend in a period of solar maximum is less than before, as is the peak intensity of the maximum. As a result, the total mean end-of-life level of radiation damage on the AF has been lowered to $\sim 3.11 [2.55 ; 3.37] \times 10^9$ 10MeV proton/cm² (EADS Astrium (GAIA.ASF.TCN.SAT.00035), current version). The original level was adopted for the radiation campaigns, which were carried out by Astrium to investigate CTI effects on Gaia CCDs. This is important to be aware of as the analyses performed in chapter 2 on these campaigns will most likely be indicative of a radiation environment slightly worse than what Gaia will actually experience.

1.9.2 Dealing with radiation damage: learning from other missions

Although Gaia is unique in many ways, it will not be the first mission to have to deal with the issue of radiation damage on CCDs. CCDs have become a common detector type for space observatories, and as a result a number of laboratory-based tests have been made characterising the behaviour of irradiated CCDs (see e.g. Waczynski et al. (2001), Abbey et al. (2001) and Ambrosi & Denby (2005)). HST, X-ray Multi-Mirror (XMM)-Newton as well as Kepler use CCDs, which all have to deal with radiation damage as part of their operation.

HST resides in low-Earth orbit at an altitude of 580km. Here, the satellite is exposed to electrons and protons trapped in the Van-Allen radiation belts, which intensifies around the South Atlantic Anomaly. At a low inclination of 28° the geomagnetic field tends to shield HST from solar protons. Clampin et al. (2002) discuss the in-flight performance of the Advanced Camera for Surveys (ACS) CCDs focusing on the evolution of hot pixels and their partial mitigation via annealing. Sirianni & Mutchler (2006) describe the radiation damage imposed on the ACS, Wide Field Planetary Camera 2 (WFPC2) and Space Telescope Imaging Spectrograph (STIS) CCDs, in terms of the dark current, hot pixels and CTI. Displacement damage, described earlier, is the root cause of these problems, which is tackled in various ways. By annealing (i.e. heating the CCD up to $\sim 20^\circ\text{C}$) every month a significant number of hot pixels (i.e. individual pixels with very high dark current) can be removed. However, there is an underlying increase in their numbers over time. CTI is mainly removed from Hubble data by modelling radiation damage (see e.g. Massey et al. (2010)). Here a corrective term is applied that returns individual electrons to their original pixels. First Pixel Response (FPR) (Janesick (2001)) is one way to measure CTI that has been used to monitor HST (see e.g. Chiaberge et al. (2006)). This technique is also used in chapter 3 to probe a Gaia prototype AF CCD.

CCDs have also been used in several X-ray space missions including XMM-Newton. Abbey et al. (2003) discusses the mitigation of radiation damage from XMM-Newton European Photon Imaging

Cameras (EPIC) CCDs. As with HST, a corrective piece of software is implemented on-ground. This can remove most of the damage, although the statistical noise from charge trapping persists. Charge injection, a hardware technique that injects charge across the CCD, has been implemented onboard XMM-Newton (see e.g. Smith et al. (2004)) as well as HSTs Wide Field Camera 3 (WFC3) and NASAs Kepler mission. It effectively fills the traps on a CCD before an image is taken. Annealing CCDs to 100-150°C has been shown to help remove some of the traps on a CCD (Hopkinson & Mohammadzadeh (2003)), and XMM-Newton CCDs are qualified up to 150°C (Abbey et al. (2003)). Although annealing to or above 150°C is needed to remove the dominant P-V defect (E-centre), which may not be practical for removing CTI. Lastly, the XMM-Newton CCDs after originally being cooled to 173K, were further cooled to 153K in 2002. This reduced CTI by a factor of 2-3 depending on the CCD considered. The serial CTI was, however, unchanged (Kirsch et al. (2005)).

Every space mission using CCDs may use one or several of the techniques mentioned above: optimising the CCD temperature; annealing; charge injection; CTI mitigation model, to mitigate radiation damage. Are any of these techniques viable for Gaia? Gaia is an astrometric mission, and obtaining highly accurate positional measurements is paramount. If microsecond astrometry is to be realised, it is imperative that high-precision centroiding can be performed per CCD transit (e.g. $\sim 10^{-3}$ pixel at $G \approx 10$ and 10^{-2} pixel for $G \approx 15$). This places stringent demands on CTI correction and essentially drives the choice of mitigation strategy. In the next section a host of possible onboard solutions for Gaia are discussed. Here, the techniques which cannot be used by Gaia are mentioned.

Annealing CCDs to very high temperatures ($>150^\circ\text{C}$) is not practical for Gaia so applying corrective software, developed to deal with CTI in HST and XMM-Newton could be a starting point (see e.g. Bristow (2003) and Massey et al. (2010)). Typically, these models use the data to find which trap types are dominant and model their behaviour at the pixel level using capture and emission times as parameters. However, there are some noticeable differences between the data that is present in an HST image and what will be sent to ground from Gaia. For example, Massey et al. (2010) measure CTI by identifying warm pixels within an HST image. A brief overview of this procedure is as follows: firstly, warm pixels, which are comparable to 'hot' pixels but less saturated, are found by locating the sites of 2D local maxima in the image. Then, the charge trail behind the warm pixel is fitted with a model describing an exponential decay process that is parameterised according to the emission times of each trap species and their relative densities. The number of exponential terms decides the number of trap species that are relevant. Incidentally, whereas Massey et al. (2010) consider the capture time to be instantaneous, this will not be the case for Gaia as the capture time is comparable to the dwell time (i.e. the time a charge packet spends accumulating charge in one phase of a CCD pixel). By integrating under these charge trails the number of exposed traps may be measured, from which the mean density of traps may be found. After this the CCD readout is simulated and the corrected image is obtained by comparing the predicted 'true' image, which has had charge trails added to it, to the actual observed image, subtracting the difference and repeating the process until convergence is achieved. If Gaia were able to read out and transmit to ground the data from every pixel from each of its CCDs as HST does then a great deal of information would be available on warm pixels. However, this is not the case. Gaia's focal plane contains 106 CCDs, each CCD comprises 8847000 pixels. As it takes $\sim 4.4\text{s}$ to read out a full frame and,

uncompressed, each pixel takes up ~ 2 bytes of data, ~ 36 Terabytes of data would have to be transmitted per day. In practice the compression factor is just over 2 so this value could be halved. However, such a value is still unfeasible and the actual allowed maximum daily data rate is on average 2 to 3 Mbits s^{-1} , which amounts to ~ 27 Gigabytes of data per day (compressed), a factor of ~ 600 less than would otherwise be needed for full-frame readout. Because of this constraint, Gaia must adopt a 'window' scheme that only transmits data centred on an image that has been identified as having a particular brightness in the Sky Mappers. All remaining data, including those containing information on warm pixels, are discarded on readout. Additionally, to increase signal-to-noise ratio most of these windowed images are binned on readout so only 1D images are transmitted to ground. Crucially though, as the image formation takes place in TDI mode there is no one-to-one correspondence between the image and its location on the CCD or in the focal plane array. Consequently, the method described in Massey et al. (2010) cannot be applied to Gaia, and instead a unique means to mitigate CTI must be found. To this end, a 'forward-modelling' approach has been adopted by the DPAC whereby the observed image location and shape must be predicted in both 1D and 2D form beyond the confines of the window and the local transit history modelled. This process is properly introduced in section 1.9.4.

1.9.3 Mitigating radiation damage on board Gaia

Understanding the effects of radiation damage on Gaia's FPA and learning how to mitigate it has been a critical part of the design phase. Finding the right hardware solution has required a significant number of industrial tests. These tests, referred to hereafter as the radiation campaigns (RCs), incorporate many of the mitigation strategies discussed below.

Shielding the FPA: stopping the damage before it occurs. Shielding the focal plane from the radiation environment is required in any space mission. It can be effective in reducing the level of ionizing radiation. However, it is difficult to stop non-ionizing radiation (i.e. those that cause displacement damage), if at all, by increasing the thickness of Al shielding. Furthermore, there are strict mass limits to any additional shielding. As mentioned before, from an initial analysis, Astrium calculated the expected end-of-life level of radiation damage on the AF CCDs (EADS Astrium (GAIA.ASF.TCN.SAT.00035)) to be $\sim 4.8 \times 10^9$ 10MeV protons/cm². Since then the launch date has been moved to November 2012 and the solar maximum is expected to be late and less intense, so the level of damage is likely to be lower.

Optimising the CCD operating temperature: suppressing the effects of the damage. As an object is tracked across a CCD in TDI mode it passes from the Image Area (IA), where charge is generated, to a Serial Register (SR), where the charge is read out of the CCD. Tests conducted by Sira on Gaia prototype AF CCDs (Hopkinson (254.DO.28)) found that in order to balance both the smear and charge loss in the IA and the SR, the operating temperature of the CCD should be around 163K. Below this temperature CTI effects became worse in the SR, whereas above 163K the opposite was seen in the IA. Such a definitive result, and coupled with the fact that transfer rates through the CCD were altered following these tests, led Astrium to conduct their own investigations on the CCD temperature. These tests were

incorporated into the radiation campaigns.

Injecting charge across each CCD: keeping traps full. Charge injection, already used on XMM-Newton, is a hardware technique that uses a controlled burst of charge to actively mitigate CTI effects on a CCD. The principle is as follows: a line of charge is injected at the beginning of a CCD and then left to fill as many traps as possible as the charge is read out over the rest of the CCD. As discussed above, this is an option for Gaia and has been tested extensively as part of the radiation campaigns.

Relying on the sky diffuse optical background: keeping traps full. CTI may be reduced by increasing the background level (see e.g. Hopkinson (2000)) and HST, being a direct imaging telescope, can take advantage of this by taking long exposures of the same patch of sky for hundreds of seconds. Gaia, being a scanning telescope, cannot do this so the average sky background is on average very low. de Bruijne (JDB-031) shows this value to be about 2 [~ 0.9 ; ~ 6.2] e^-/pixel in the AF, assuming the sky has a uniform brightness of $V = 22.5$ mag arcsec $^{-2}$ (ESA Gaia Project (GAIA-EST-RD-00553)). The minimum and maximum are shown in the square brackets. Whether this very low level of diffuse optical background has any mitigating effects onboard Gaia was also addressed in the radiation campaigns.

Inserting a supplementary buried channel into the CCD structure: avoiding traps. CTI increases as the brightness of an image decreases (Janesick (2001)). One way that has been developed to offset this trend is to insert a supplementary buried channel or "notch" inside the buried channel. This method has for instance been incorporated into the new device for the WFC3 camera onboard HST (installed May 2009). The Supplementary Buried Channel (SBC) is much narrower than the buried channel and channels faint signals away from many of the traps. Gaia CCDs are also expected to incorporate this structure. How well the SBC is at mitigating radiation damage is investigated in chapter 3.

1.9.4 Mitigating radiation damage in the data processing

DPAC currently plans to remove radiation damage in the pre-processing part of the data pipeline. This is where image parameters (centroid and flux) will be extracted, before the astrometric or photometric processing stages take over. Rather than treat the data directly, the distorted image will be modelled.

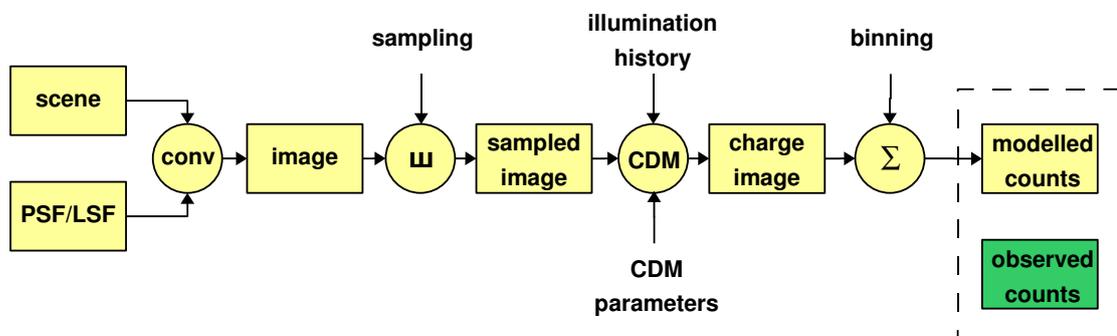


Figure 1.13: Calibrating radiation damage in the data pre-processing. Credit: Lindegren (LL-075).

How we might create a Charge Distortion Model (CDM) that is sufficiently robust to cover all of the possible observing conditions is currently an area of intense research. Overall, the 'forward-modelling' approach being adopted by DPAC is illustrated in Figure 1.13.

The idea here is to iterate over the set of various model parameters (scene, PSF/LSF, CDM) until the best fit between the modelled counts and observed counts is found. In each iteration a better estimate of the undamaged centroid and flux may be obtained. The different stages shown above in the forward-modelling approach are as follows:

- *scene* - an idealised representation of the object, e.g. a scaled delta function on a constant background for a single star. It describes the positions of the source on the sky and the background. Once the PSF/LSF and CDM have been properly calibrated, these are the positions that will be used to determine the sources' astrometric parameters.
- *PSF/LSF* - a point-spread function (i.e. a profile of an astronomical point source imaged onto a 2D array) or line-spread function (i.e. binned PSF). This is extracted from a library of PSF/LSFs that are modelled according to the FOV, the CCD, the AC position, star colour etc.
- *image* - a continuous function expressed in expected number of counts per pixel or sample, obtained by linear superposition of the corresponding (scaled and shifted) PSF or LSF (see below for more details).
- *sampled image* - a set of discrete values representing the intensity of the image in each pixel.
- *charge image* - a sampled image modified by CTI effects. The modification is described by a charge distortion model (CDM) that takes into account the illumination history of the pixels (e.g. time since charge injection, background level) and feeds these values into the CDM as parameters. Together with a set of trap parameters a prediction is made of how the sampled image will appear after being distorted by radiation damage.
- *modelled counts* - these are obtained by binning the charge image and adding known noise. Most images onboard Gaia will be binned on readout.
- *observed counts* - these are the raw data from Gaia with bias (i.e. an electronic offset), gain (i.e. conversion factor between ADUs and electrons per sample), and non-linearity all taken into account.

Figure 1.14 illustrates the process of producing an image when a double star in the scene is convolved with an LSF for that particular position in the focal plane, colour of star(s) etc. The image provides a continuous function of the expected number of counts per sample. The image is then sampled and distorted by the CDM. This is shown for a single star in Figure 1.15.

Note that once the *observed counts* and *modelled counts* are found to agree the centroid of the undamaged image may be found as the distorted image centroid is defined in terms of the undistorted LSF.

The CDM itself is, at this time, in the process of being developed. Two versions exist: one empirical, the other based on physical theory. Both are discussed in chapter 4, in particular the latter, which has

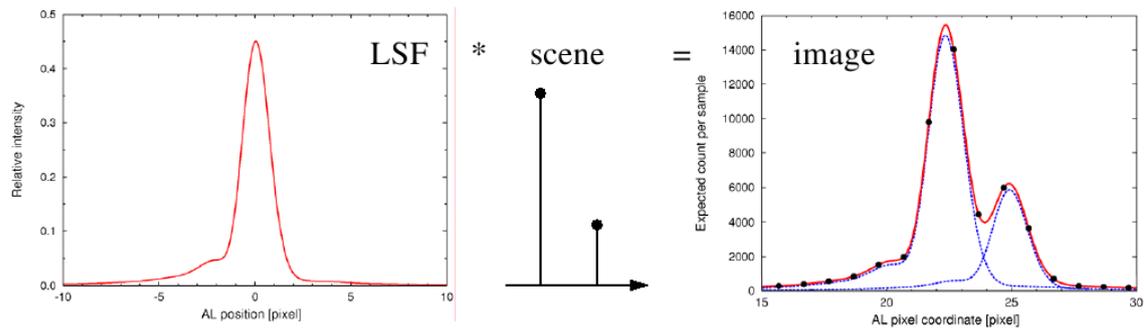


Figure 1.14: Here the *LSF*, modelled for a particular part of the focal plane (i.e. FOV, CCD, gate etc), is convolved with the *scene*, which in this case comprises two stars, to produce an undamaged *image*. Credit: Lennart Lindgren, c.f. ESA Gaia Project (GAIA-EST-MN-05213).

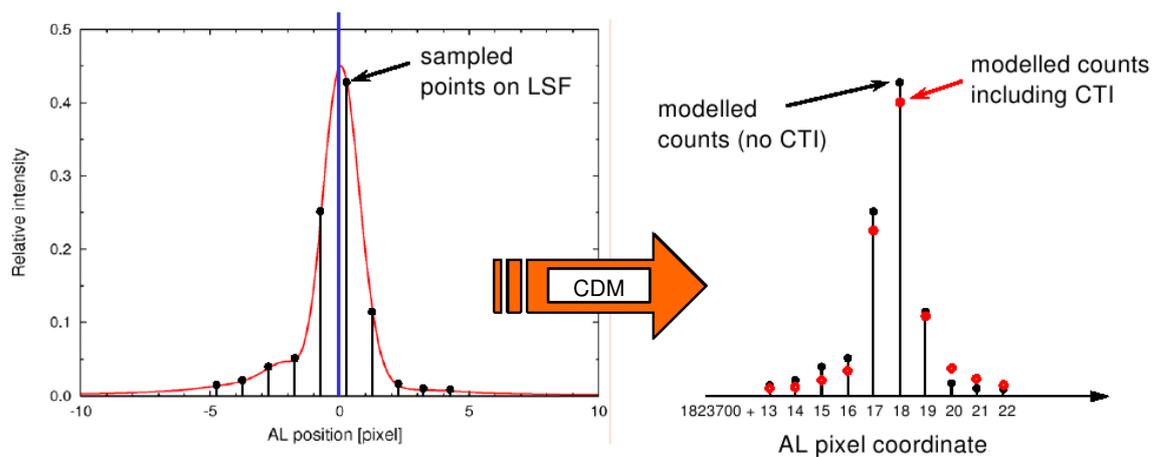


Figure 1.15: Here a single *image* is sampled to produce the *sampled image*. The samples are then distorted to produce the *charge image*. Binning produces the *modelled counts*, which can then be compared to the *observed counts*. Credit: Lennart Lindgren, c.f. ESA Gaia Project (GAIA-EST-MN-05213).

been tested against radiation campaign data. The challenge is to find a unique set of trap parameters that can describe a whole host of test conditions, e.g. different magnitudes, different times since charge injection, which can then be used to implement the radiation calibration shown above.

This is how radiation damage may be implemented as part of the data processing. To see how radiation damage is currently being handled as part of DPACs overall radiation calibration strategy consider Figure 1.16. Here the connection between Astrium's hardware tests and the final CTI processing algorithms is shown. The radiation test plans are formulated and agreed upon by the Radiation Calibration Working Group (RCWG), comprising members of ESA, Astrium and DPAC. The data from these tests are analysed by both industry and DPAC. The results are used to verify microscopic and macroscopic models which have been developed to take account of radiation damage in the data processing and provide input for the data simulations.

Every six months the Radiation Task Force (RTF) meets to discuss the current radiation calibration status within DPAC. Members of ESA, DPAC and the European Leadership in Space Astrometry (ELSA)

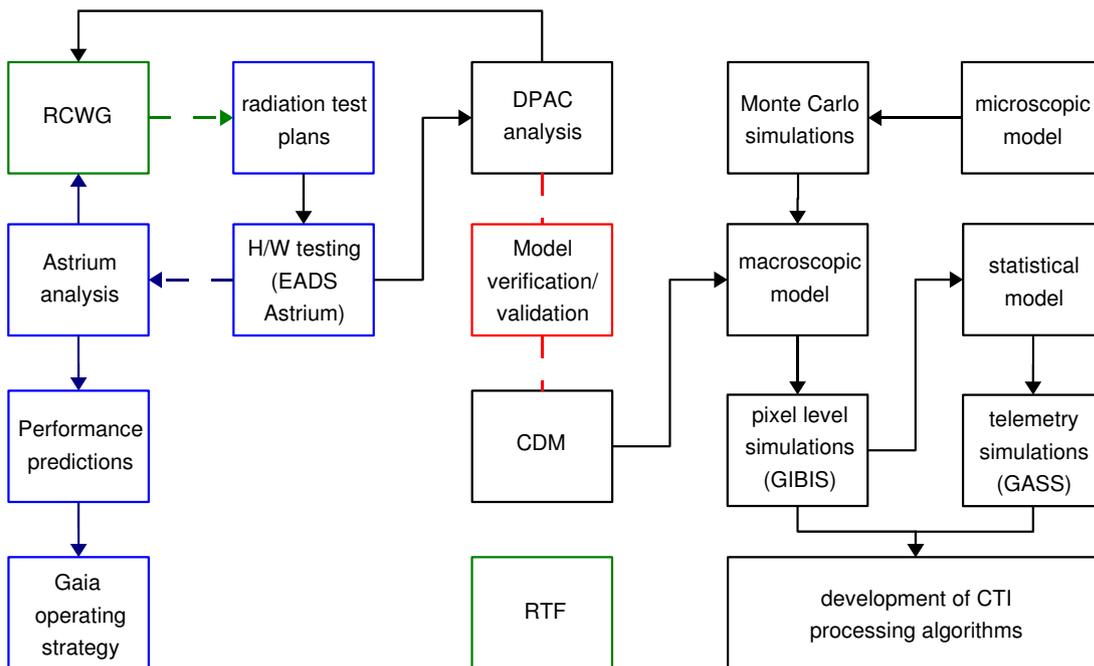


Figure 1.16: The current radiation calibration strategy. The black boxes show DPAC's input. The green boxes highlight the input from the RCWG and the RTF. The blue boxes show Astrium's input and the red box shows where the data analysis and model development cross over. (This figure is based on information found in Short (2009) and Brown (2009)).

network attend. Like the RCWG, the RTF decides where the effort should be directed in the near future in order to incorporate radiation handling into the data processing.

1.10 A template for this thesis

Radiation damage is a significant threat to the Gaia mission, so there are currently five lines of attack proposed by DPAC for calibrating its effects. These are described in van Leeuwen & Lindegren (FVL-006) and are restated here, as they essentially serve as a template for the work that will be presented in the coming chapters:

1. *Minimizing radiation damage by further hardware optimisation.* Chapter 2 looks closely at the first three radiation campaigns and discovers a number of hardware anomalies that need optimising.
2. *Hardware (CCD) testing by industry.* These are the campaigns performed by Astrium.
3. *Analysis of the hardware tests by industry and by the data processing groups.* Chapter 2 analyses the radiation campaigns in detail and compares, where applicable, the results to industrial analyses.
4. *Implementation in the simulation software of the radiation damage effects and production of simulated data using this implementation.* Working with colleagues who are involved with producing simulated data, Chapters 3 and 4 discuss how the analysis of RC2 has been used to probe certain

physical parameters of a Gaia AF CCD including the supplementary buried channel, and test the predictions of a CDM respectively.

5. *Software implementation for the astrometric, photometric and RVS data reductions.* Chapter 5 considers several algorithms that may be chosen to carry out the initial image parameter extraction procedure. With these algorithms it has been possible to test the robustness of one of these CDMs.

2

The Radiation Campaigns

DURING the Gaia definition phase, prior to 2005, it became apparent how critical the issue of radiation damage could be for Gaia. Early testing suggested that if left untreated, radiation-induced Charge Transfer Inefficiency (CTI) could well impede Gaia's potential to reach microarcsecond astrometry. A sizeable part of the catalogue would end up being downgraded or even dismissed completely if a workable solution could not be found. Since then a series of radiation campaigns (RCs) have been conducted to understand and characterise the problem, testing the different conditions Gaia may well experience in operation, optimising predefined parameters and observing the effects CTI will have on Gaia data.

The campaigns have posed a number of significant challenges, from installing and optimising the test equipment to analysing and interpreting the data. To reproduce the Gaia observations a test-bench was built in-house by Astrium to project a diffraction-limited star field onto a prototype Astrometric Field (AF) CCD. The CCD had to be operated within a specific temperature range and remain cold while a large number of scans were performed in order to over-sample the stellar image and generate statistics. Mitigating strategies had to be calibrated and characterised. Once the data had been analysed the results were fed back to Astrium as soon as possible so that improvements to the test configuration could be incorporated into the next campaign.

While all three instruments (AF, RP/BP and RVS) were included in the campaign test plans, this work focuses entirely on the effects of CTI on the AF.

Software tools

Since Java is the programming language that has been chosen to engineer the Gaia data processing software, Java has also been used to code the analyses of the RCs. Where needed, GaiaTools¹ and JFreeChart² have also been used. All of these projects are stored under subversion (SVN)³.

2.1 Methodology

In each RC the test setup may change, e.g. a larger optical mask or a different injection strategy, but the underlying techniques required to reduce, process and analyse the data in many ways remain the same. Refinements, of course, must be made. However, it is advantageous that the core methodology is always followed. For this reason, in this section I will describe the generic aspects of each campaign that need only be discussed once.

2.1.1 CCD readout

In its simplest form the readout from a CCD pairs the number of photoelectrons contained in a single pixel to the time the readout was made. For Gaia the signal will usually be binned on readout over a 12-pixel window in the AC direction and the time will be given by an internal clock. Any charge packet that is formed by binning photoelectrons across several pixels is called a sample. Therefore one sample, in this case, contains the charge collected in 1 AL x 12 AC pixels. In the RCs the resulting sample value, which is converted into a Least Significant Bit (LSB) in the Proximity Electronics Module (PEM), was stored in a tab-delimited text file. An example of this output is shown in Figure 2.1.

Fichier de configuration: D:/projet_gaia/CCN 10/conf de seqts/a cheval 4E9.xml													
Date de fin d'acquisition: 04/11/06 02:46 PM													
Commentaire: default													
SYNCHRO1	SYNCHRO2	TIMETAG1	TIMETAG2	PEMID	Sample 1	Sample 2	Sample 3	Sample 4	Sample 5	Sample 6	Sample 7	...	Sample 22
61610	65450	23	10963	12288	1275	1276	1275	1276	1281	1282	1283	...	1276
61610	65450	23	10964	12288	1277	1274	1275	1277	1279	1281	1282	...	1275
61610	65450	23	10965	12288	1275	1275	1276	1277	1279	1283	1282	...	1274
61610	65450	23	10966	12288	1273	1271	1276	1278	1280	1279	1282	...	1272
61610	65450	23	10967	12288	1276	1274	1272	1278	1283	1283	1281	...	1273
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
61610	65450	23	17762	12288	1276	1276	1280	1278	1281	1283	1283	...	1274

Figure 2.1: An example of a raw data file. The basic file layout comprises a header followed by a series of readouts. The first two columns labelled SYNCHRO1 and SYNCHRO2 are tags which Astrium used to acquire the data. The columns labelled TIMETAG1 and TIMETAG2 represent the most and least significant (16-) bits of a counter matched to the sample TDI period. The column labelled PEMID refers to the PEM index which again Astrium used to acquire the data. Lastly, columns labelled SAMPLE1 through to SAMPLE22 represent the flux in LSBs.

¹Developed by CU1, GaiaTools contains a number of useful routines that are available to the Gaia community.

²See <http://www.jfree.org/jfreechart/>

³<http://gaia.esac.esa.int/dpacsvn/>. For the radiation analysis code see <http://gaia.esac.esa.int/dpacsvn/DPAC.branches/CU5/software/personal/sbrown/>. Note that these projects are username and password protected so membership of DPAC is required.

For our analysis only those columns labelled TIMETAG2 and SAMPLE_x, where *x* denotes the sample index, in Figure 2.1 were important. The TIMETAG2 column refers to the time given by an internal clock/counter and is set by the TDI period. This is the time it takes a packet of charge to cross one pixel in the AL direction, $982.8\mu\text{s}$, and is fixed according to Gaia's scanning rate. The sample value refers to the flux given in LSBs. This can easily be converted back to the photoelectron count by multiplying the flux by the electronic gain. The actual conversion is carried out in the PEM housed behind the CCD. In RC1, a total of 22 samples were read out from the CCD after each TDI period, whereas in RC2 and 3, 32 samples were read out. Whenever a sample is mentioned followed by a number this will refer to the sample index. This indicates the AC position at which the sample was read out. Many samples do not contain any information on the stellar field, so in this analysis sample indices that *do* contain the stellar field are often referred to as channels. Figure 2.2 attempts to illustrate the process of binning the signal on readout and the use of this terminology.

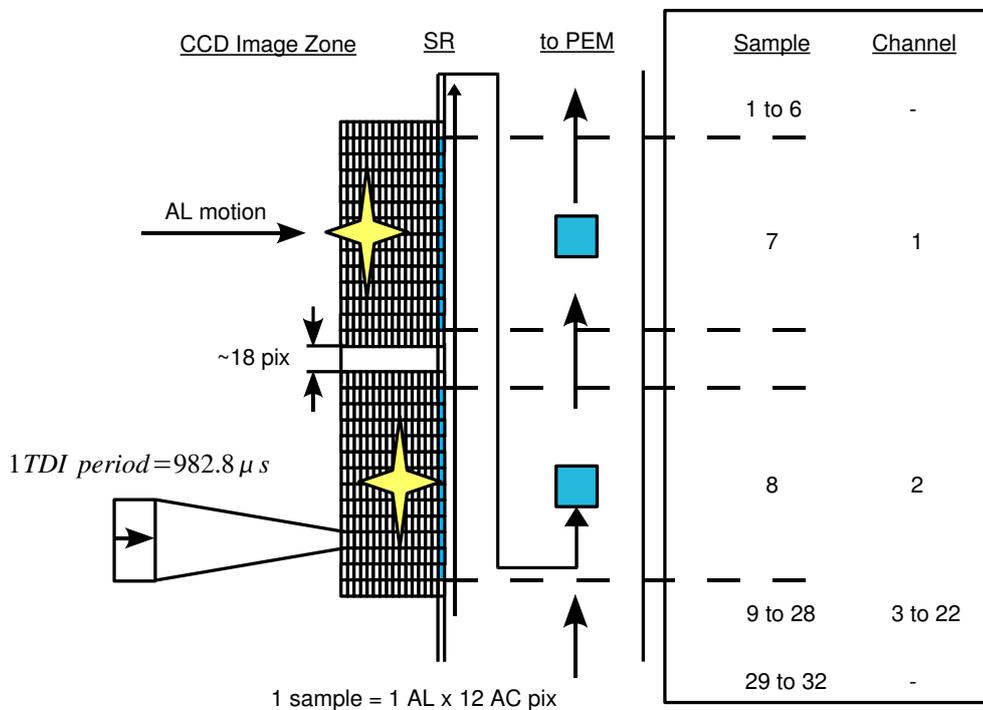


Figure 2.2: Schematic of the sampling scheme used in the CCD readout. Here the procedure is shown for RC2. In total there are 32 samples read out from the CCD after each TDI period. Only samples 7 to 28 contain information on the stellar field. To make this explicit these samples are often referred to as channels numbering 1 to 22. The pixels are binned on-chip at the CCD output node level, i.e. charge is left unbinned as it is transferred through the serial register (SR) and then binned before entering the PEM.

Since binary files are more efficient to process than text files both the TIMETAG2 and the SAMPLE_x columns were extracted and converted into a binary format.

The counter operates on a 16-bit loop such that the maximum range attainable is 65535. Since each test involved the accumulation of several scans, it was often the case that this value was exceeded, and the counter can be seen to reset during a scan. Even though this occurred in all three campaigns the reset

only became an issue in RC2 for tests where the charge injection (CI) delay was very long (≥ 30000 TDI periods). Here the TDI period measured after the clock had to be incremented and appended to the previous readout to avoid scans overlapping.

2.1.2 Extracting the background

One highly important, preliminary part of the analysis on the effects of CTI on the AF is an accurate determination of the background underneath the stellar field. The PEM applies an electronic offset to each sample on readout. In the campaign data, sample 1 contains only this offset. In this particular sample, also called a *prescan* sample, the readout is left unbinned. In RC2 and 3 sample 2 also records only the electronic offset, however, in this case the readout has been binned by 12 AC pixels. Since the electronic offset is expected to be a fixed value in every sample, predefined by the test setup, the prescan sample provides a simple way of directly measuring the offset and associated Readout Noise (RON).

During the analysis of RC1 it became clear that the prescan sample could not be used to determine the electronic offset. The reasons for this are given in the full analysis of RC1 (see section 2.2). Suffice it to say for now that to proceed the RON as measured in sample 1 was assumed to remain constant while the electronic offset and diffuse optical background (DOB) were measured in each individual channel.

Since each scan took several thousands of TDI periods to complete and the stellar field always appeared towards the end of each scan the background was estimated over a region typically between 500 and 1500 TDI periods after the scan had begun. Unless otherwise stated the method chosen to extract the background under the stellar field in each campaign is shown in Figure 2.3. Here the region over which the mean background B_{est} is estimated is indicated and the red line shows the result of applying a flat cut. It will be shown that this method is not optimal. However, given the dimensions of the optical mask in RC2 and 3, in some cases this was the only method that could be used.

In reality it will not be possible to measure the background directly far before the image has passed since the stellar field will be far more complex. In order to get a hold on the background, virtual objects will be used to extract windows over parts of the field that contain no stars. The size of the windows will determine how much contamination will arise from close-by images.

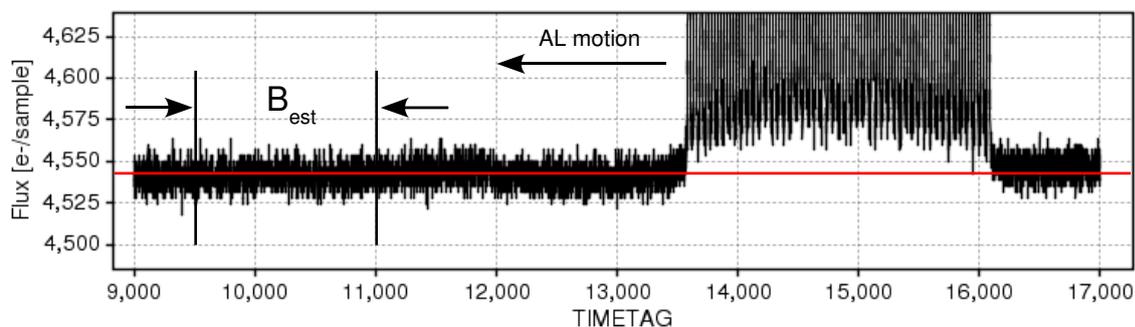


Figure 2.3: Extracting the background using a flat cut. This example is taken from RC2. Readout is to the left as indicated and the stellar field in this channel can be seen on the right. Since the electronic offset must be estimated within each channel, the region marked B_{est} is chosen to extract a representative mean value. The red line extrapolates the flat cut to cover the entire scan.

2.1.3 Extracting the raw image

Once the background has been determined and subtracted, the image itself must be extracted from the rest of the readout. This introduces, in part, the Gaia windowing scheme. As mentioned in section 1.5.2, Gaia assigns a window around an image the size of which depends on the object's magnitude⁴. Any signal falling outside this window and not falling inside a neighbouring window is discarded on readout. For the campaign analyses, the AL window length was set to 15 pixels to constrain curve fitting and modelled over 12 pixels to agree with the Gaia window scheme in the majority of AF CCD columns for bright and medium stars ($G = 13-16$). This window size was also chosen by ASTRUM in their analyses. For fainter stars ($G = 16-20$) the window length in some CCDs is reduced to 6 pixels. However, this reduction was not adopted for this work (or indeed by ASTRUM) since the focus in this chapter is to investigate the effects of CTI on the AF across all magnitudes. Chapter 5 addresses in more detail how a window scheme impacts on the accuracy of extracting image parameters.

In order to extract an image from the readout, the peak flux in that image had to be identified. The peak flux defines the centre of the window and the 7 datapoints either side fill the remainder. As no assumptions were made concerning the consistency in scan timing, the peak flux was identified by defining an isophotal layer, i.e. a flux threshold, above which only the tip of the stellar image would appear. This was a quick and simple method to extract the brightest signals. For the fainter signals greater care had to be taken to define the threshold levels. Here the noise in the background could easily be mistaken for an image and checks became more frequent to ensure the procedure was working properly. Using

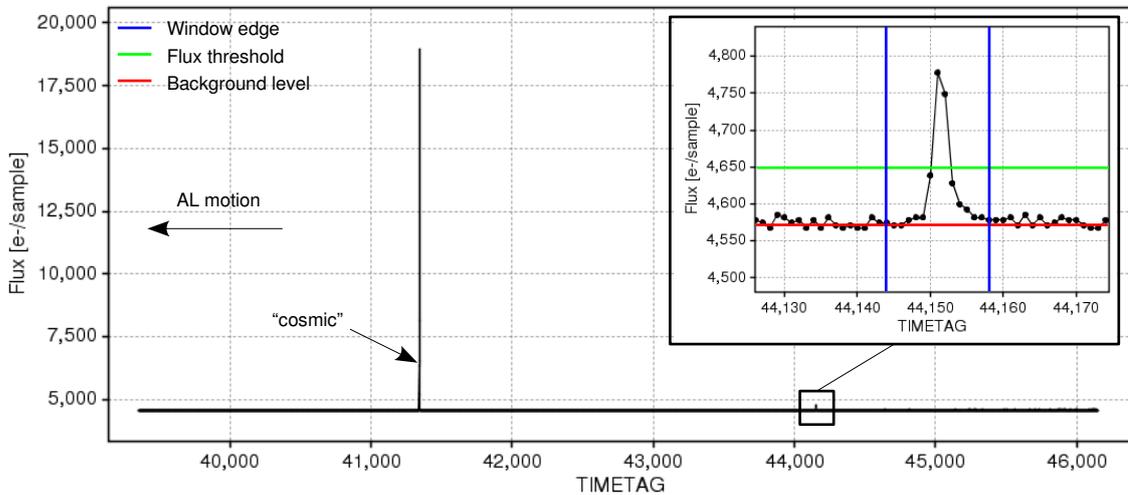


Figure 2.4: Extracting a raw image from the readout. This example is taken from RC1 when $G \approx 20$. In the closeup the 15-pixel window is shown (blue) together with the flux threshold level (green) and the mean background estimate (red). Before the image is readout a "cosmic" is seen. Lying above the flux threshold this extraneous signal must be bypassed in order to extract the real image.

⁴In the AF, which consists of 9 columns of CCDs, the AL window length depends upon the column number and the source magnitude. In AF 2, 5 and 8 every image, including the faint ones, are extracted using 12-pixel windows. In all the other AF CCD columns the faint stars are extracted using 6-pixel windows, the bright stars 12-pixel windows. See e.g. Hoeg & de Bruijne (EH-176).

a threshold to identify the peak flux in the image also increases the chances of picking up anomalous features in the data that appear above this level. For example, light leaking onto the CCD in short bursts occurred often in the first two campaigns. Additional checks had to be included in the procedure in order to ensure these "cosmics" were not being picked up. Figure 2.4 illustrates the technique used to extract the raw image. This particular readout also contains a cosmic.

2.1.4 Combining 1D stellar images: sub-pixel sampling

On the CCD image resolution is limited by pixel size and even though each pixel in the AF measures only $10\mu\text{m AL} \times 30\mu\text{m AC}$ this still does not adequately sample the image shape. In order to go deeper and resolve image shapes at the sub-pixel level, multiple scans are performed, and every comparable image is combined in such a way that the underlying image shape can be over-sampled. For multiple scans the process of combining one-dimensional (1D) stellar images is as follows:

1. Firstly, one scan is chosen to act as the reference. This is an arbitrary choice but may lead to processing problems if the images in this scan are corrupt or an unexpected feature is encountered. This is especially the case with the irradiated data. To avoid any problems the data is reprocessed several times with the choice of reference scan being different in each case⁵. Scan 1 could not be a reference scan since it was damaged far more than any other scan (see "Scan 1: a self-injection" in section 2.2.4). Only in the final analysis stages are the average positions and fluxes calculated over all reference scans.
2. Fix every comparable image to the position of the peak flux. This is not strictly necessary but is convenient when step 3 is carried out. It is also convenient at this stage to fix the zeropoint to the position of peak flux. The reason for this will become apparent when the global reference curve is constructed in the next section.
3. Cross-correlate every comparable image with the reference image, keeping the position of the reference image fixed. Where the cross-correlation reaches a maximum, add the associated value to the position of the comparable image.
4. The cross-correlation performed in step 3 can only provide positional adjustments at the pixel level. To determine sub-pixel positions at the scale of ~ 100 th of a pixel, the region around the peak of the cross-correlation curve can be interpolated using a quadratic function. Use the new peak value to redefine the position of the comparable image.
5. Repeat the above procedure for all remaining images.

Further details on combining 1D stellar images are given in A.1.

⁵In RC1 the number of scans chosen to act as references equalled the available number of scans. This, however, soon became prohibitively slow to process so in RC2 and RC3 the number was reduced to only 3 scans, usually scans 2, 3 and 4.

2.1.5 Defining a robust Global Reference Curve

In order to extract relative positions and fluxes from 1D stellar images a reliable model of the undamaged LSF must be defined. Since this model acts as the reference for every image read out from the CCD it is called the *global* reference curve, and is constructed by, firstly, combining every image in the reference dataset. Though this approach has had to be adapted in each campaign analysis, the objective remains the same: to ensure the undamaged, global reference LSF is very well sampled. At this point a suitable model must be chosen to represent the LSF. Without assuming any particular profile a cubic spline function was chosen to fit the accumulated data. The implementation of this particular spline function has already been tested and used extensively by van Leeuwen in his new reduction of the raw Hipparcos data (van Leeuwen (2007), Appendix B). Customizing the position of the spline knots is a particular feature that has been used to great effect in this analysis to check the reliability of the Global Reference Curve (GRC) fitting.

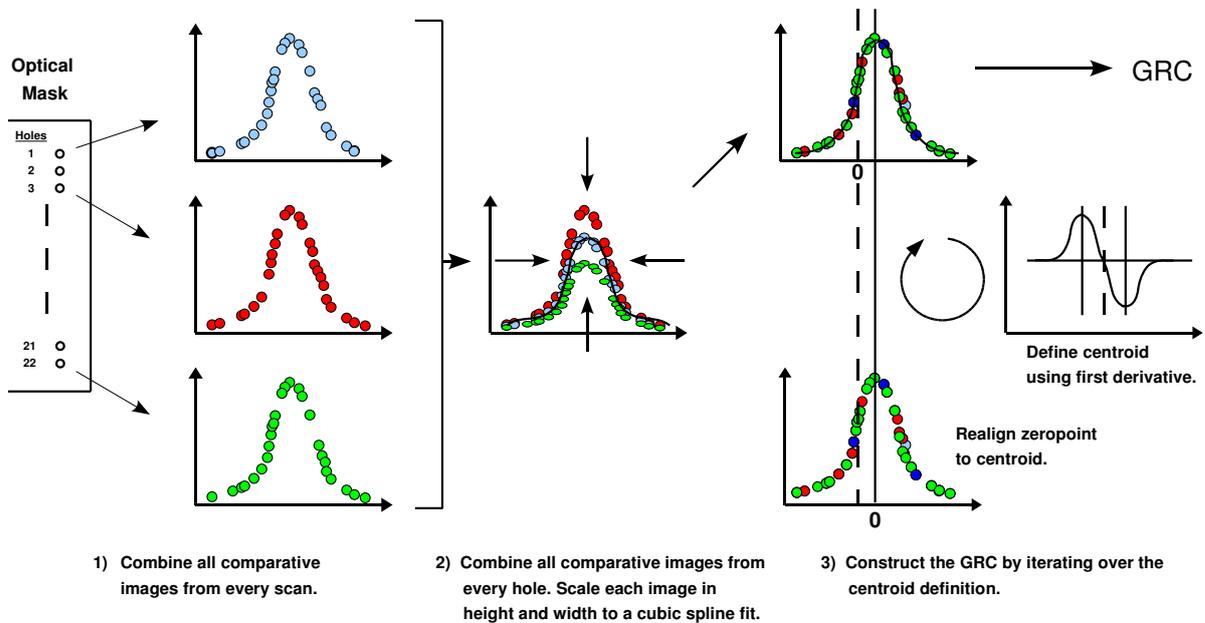


Figure 2.5: How to construct a GRC using the reference data. In this diagram the optical mask is the regular grid used in RC2 and 3. All the available data is combined to form a single image, one which is heavily over-sampled and able to constrain a cubic spline fit. The GRC is the final fit to the accumulated data that terminates the iteration over the centroid definition.

As the optical mask used to generate the stellar field was not illuminated uniformly in each campaign, overlaying a combined set of images from across the mask left variations in both the height and width between images. Assuming the optical setup does not change between scans and the non-uniform illumination is due to the finite distance between the mask and light source, the variations in height and width between images were normalised to the Full Width at Half Maximum (FWHM) and height of a cubic spline fit to the accumulated data. Curve fitting was performed using Householder least squares (HHLS). Further details can be found in A.2. In effect, this assumes the GRC is an average response over the

mask, which can be scaled in both height and width in order to reproduce the images in every hole.

In many respects the definition of the centroid is completely arbitrary. As long as the centroid definition is consistent, image parameters may be related. How robust the centroid definition is depends on how 'well-behaved' the images are in the dataset. In Chapter 5 the performance of three centroid definitions is assessed in terms of their ability to reduce the ambiguity in the centroid estimate. This is not so much of a problem for the GRC since it is heavily over-sampled. So, for this analysis the GRC centroid is defined as follows:

- The GRC centroid is aligned to the midpoint between the positions of the maximum and minimum of the first derivatives of a curve fit to the data.

This definition is sensitive to the gradients of both the leading and trailing edges of the GRC. If the shape of the GRC is asymmetric this definition will take account of it, shifting the centroid away from the symmetric zeropoint. The symmetry of the GRC is expected to be extremely high, since it is constructed from data taken from the non-irradiated part of the CCD. The effects of radiation damage should only be seen after the undamaged and damaged images have been properly calibrated.

Defining the GRC is an iterative process. Initially, the zeropoint is fixed to the peak flux of the reference scan and then a centroid definition is chosen to coincide with the zeropoint. A spline fit is used to realise this definition by repositioning the data until the updates to the iteration become less than a given limit. In this case, a value of 10^{-3} was used. The GRC is the final fit to the data, normalised to the peak flux with a centroid definition coinciding with the zeropoint. Figure 2.5 summarises these steps.

2.1.6 Fitting the GRC to the stellar field: extracting positions and fluxes

Once the GRC has been defined it can be used to extract relative centroid positions and fluxes across the stellar field. To maintain a handle over the relative positions this procedure requires a reference frame. Defining a reference frame is equivalent to choosing a particular channel and scan that can act as a reference both in the AC and AL direction respectively. Which channel and scan to take as a reference is an arbitrary choice, as long as the choice is consistent throughout the analysis. As channel 1 is the first to contain an image in the readout, this channel was chosen to act as the reference in the AC direction and scan 2 to act as the initial reference in the AL direction. As described above, further scans were added to act as references when problems were encountered in the data processing. Averages over these reference scans were taken during the final stages of the analysis.

Fitting the GRC to a set of combined data to extract a position relative to the GRC centroid is again an iterative procedure. The zeropoint of each image is initially defined to be aligned to the position of the peak flux of the image extracted from the reference scan. Feeding first-order Taylor expansions into a HHLS solution outputs a succession of positional corrections which are: a) used to shift the image position closer to the GRC so that the iteration can continue; and, b) summed, until the iteration terminates. The sum total of all the positional corrections is a record of the relative separation between the initial and final zeropoints for that image. By repeating this procedure for every image down each

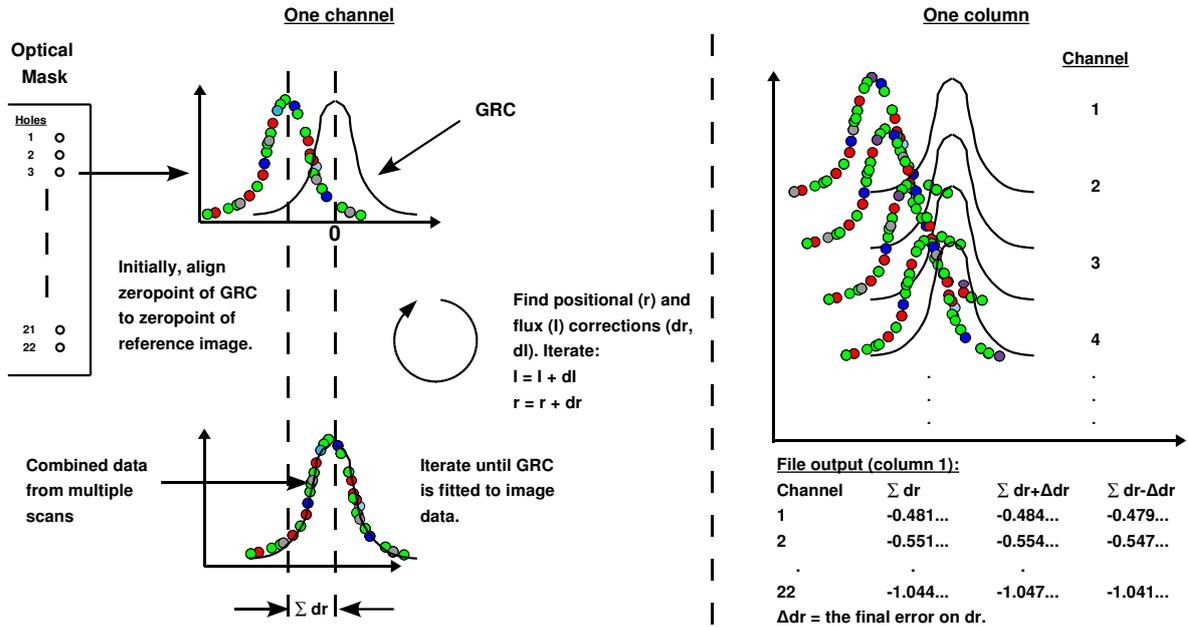


Figure 2.6: Extracting relative centroid positions and fluxes from across the stellar field. For each channel in a single (mask) column the GRC is fitted to the image data. This is an iterative process and in each cycle a positional and flux correction is applied to the GRC. In practice, it is the image data that is repositioned. This is repeated for every image in the column. Part of a generated file is shown on the bottom right that contains a list of the sum total of the positional corrections that are applied to each image together with formal errors. The difference between these measurements in each channel defines the relative centroid positions of each image within that column.

column of the mask a list of relative separations is generated and it is the differences between these summed shifts that give us a measure of the relative position of each centroid across the mask.

There is one additional check that must be made before these positions can be cross-referenced. The initial zeropoint is aligned to the position of the peak flux of the image extracted from the reference scan. If this position is more than one AL pixel away from the equivalent position of the image found in the reference scan in the reference channel, as may be the case in channel 22, an additional n number of pixels must be added to the sum total of the positional corrections for that image, where n is the difference between the absolute locations in the CCD readout. For example, if the absolute position of the peak flux of the image extracted from the reference scan (channel 1) was 4500 and the absolute position of the peak flux of the image extracted from channel 22 was 4501, $n = 1$, i.e. 1 would have to be added to the sum total of positional corrections applied to the image in channel 22. Figure 2.6 illustrates the process of fitting the GRC to each image and extracting relative positions and fluxes.

The technique of feeding first-order Taylor expansions into a HHLS solution allows both positional corrections and flux corrections to be estimated simultaneously (see A.2). To begin the iteration an initial estimate of the peak flux is required. The datapoint that has the highest flux in the combined data from multiple scans is chosen for this purpose. Once the iteration has terminated, a list of relative peak fluxes of the fitted GRC is generated.

To carry out an extended analysis there are many other items of information that are saved while the data is being processed:

- *The FWHM of the GRC.* As will be shown, the mask in every campaign is not illuminated uniformly. Therefore, variations in the peak flux instigate variations in the FWHM of each image. Since the FWHM of the GRC does not change if it is scaled in height: a) the reliability of the GRC must be evaluated for each fitting; and, b) the uniqueness of the GRC over the entire stellar field must be checked.
- *The spline knots and parameters of the GRC.* Rather than construct the GRC from the raw data every time some of the data is processed, it is quicker to generate the GRC using the knots and parameters that have been extracted from the first analysis.
- *The peak flux, FWHM, spline parameters and knot positions of a local fit to the image.* The reliability of the GRC to recover the correct image shape can be checked by fitting a local LSF to an image's accumulated data. The local LSF is a cubic spline that can have the knot positions customised to maximise the chances of recovering the actual image shape. Recovering the local fit is also important for post-campaign analysis on CDM modelling (see Chapter 3) and image parameter extraction performance (see Chapter 4).
- *The absolute AL position of each reference image.* It is very useful to plot the positions, fluxes, and other variables as a function of the absolute position of each image in the stellar field. These values can simply be added to the sum total of the positional corrections for each image in order to recover the actual positions of each image. By doing this it is also possible to compare the actual separation between images in the same channel with the nominal AL pitch.

Once the reference data has been processed the above procedure is repeated for the irradiated data. Now the data is in an intermediate form that can be used to calibrate the test setup and extract CTI image parameters.

2.1.7 Calibrating the geometry of the mask and extracting centroid biases

The relative centroid positions determined in the last section contain several biases: those due to intrinsic misalignments between the holes on the mask, those due to the rotation of the mask with respect to the CCD, and those due to CTI effects. By dealing with the first two, the bias due to CTI effects should naturally fall out of the data analysis. Using the reference data, the calibration proceeds as follows:

1. For every scan that acts as a reference the fitting procedure described above is repeated. On completion the average relative centroid positions and errors are calculated.
2. A linear fit to all the average relative positions in column 1 is made to mark the expected positions. This same fit is refitted to only those channels that remain non-irradiated in the irradiated data. The angle between the refit and the zeropoint gives the rotation of the mask with respect to the CCD

over all channels in the reference data. The zeropoint, in this case, is defined to be the expected position in channel 1.

3. The positional offsets between this fit and the actual observed positions are attributed to misalignments between the holes on the mask and are saved for later use to correct the irradiated relative centroid positions.

In order to calibrate the rotation of the mask in the irradiated data, the following additional steps must be taken when analysing the reference data:

1. The residuals between the (comparable) non-irradiated positions and the linear fit are extracted.
2. In RC2 and 3, the average residual per non-irradiated channel is found.

Once the average relative centroid positions in the irradiated tests have been determined, the geometric calibration proceeds as follows:

1. The linear fit found in the reference data, parameterised by the gradient m , is refitted to the non-irradiated positions in the irradiated data.
2. The residuals between the non-irradiated positions and the linear fit are extracted.
3. In RC2 and 3, the average residual per non-irradiated channel is found.
4. These residuals are compared to those found in the reference data. If the mask has a different rotation over all channels relative to the CCD in the irradiated test the residuals will differ.
5. The rotation found in the reference data is used to calibrate the rotation in the irradiated data. If the rotation were the same, the (average) non-irradiated residuals in both the reference and irradiated data would agree. By plotting the difference between the residuals as a function of (non-irradiated) channel number, the rotation in the irradiated data, parameterised by the gradient m' , can be found by determining the angle offsetting the residual differences:

$$m' = m + \frac{offset}{|channel - pivot|}, \quad (2.1)$$

where *offset* measures the offset in the residual differences in a chosen non-irradiated channel. *channel* is the number of the chosen non-irradiated channel and *pivot* measures the point at which the residual differences are the same, i.e. the pivot point.

6. The positional offsets found in the reference data earlier are removed from the irradiated positions.

What is assumed to remain at this point in the irradiated channels are centroid biases, caused predominantly by radiation-induced CTI effects. There may be second-order effects, e.g. a variable hole size, that have not been removed as part of the calibration. However, these effects will only impart additional noise on the positions. The remaining offsets are attributed, unless otherwise stated, to radiation damage. Figure 2.7 illustrates the steps described above.

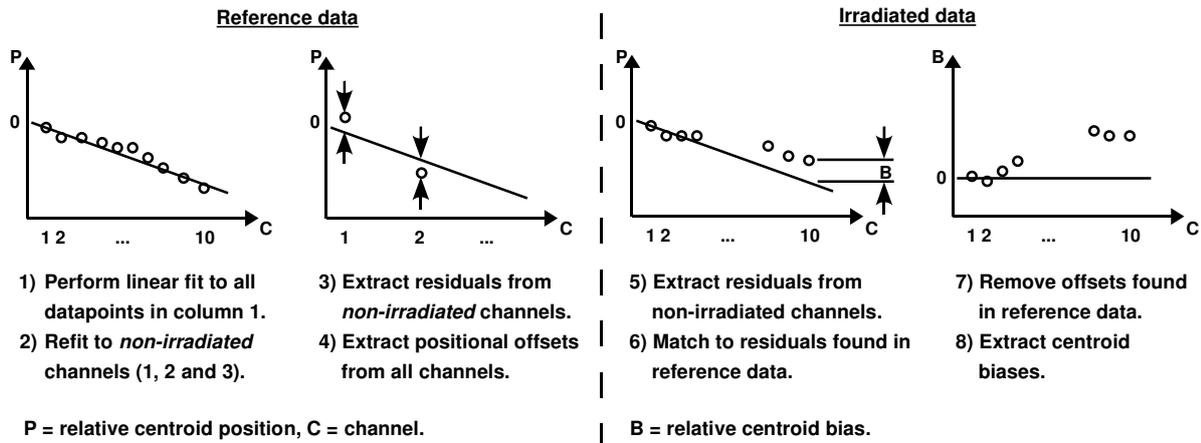


Figure 2.7: Extracting relative centroid biases from across the stellar field. In this diagram the mask is taken from RC1. In the reference data all of the channels are non-irradiated, however, the *non-irradiated* channels highlighted here are those that remain non-irradiated in the irradiated data.

2.1.8 Calibrating the flux response across the mask and extracting charge losses

The peak flux and integrated flux of both the fitted GRC and the local LSF are measures of the flux response across the mask. Since the integrated flux was found to be insensitive to the shape of the image (to be explained shortly in the first campaign analysis) this value is used to calibrate the relative charge loss across the stellar field. When comparing the integrated fluxes in the irradiated data to those found in the reference data, the fluxes measured in the non-irradiated channels should agree. This is how the calibration of the flux response is performed: matching the fluxes in the non-irradiated channels takes

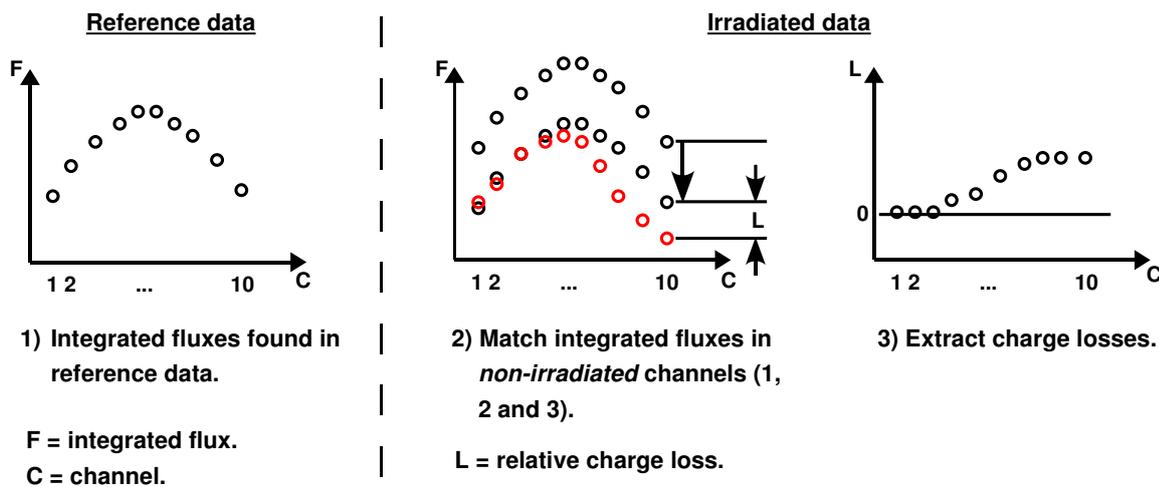


Figure 2.8: Extracting relative charge losses from across the stellar field. In this diagram the mask is taken from RC1. In the reference data all of the channels are non-irradiated, however, the *non-irradiated* channels highlighted here are those that remain non-irradiated in the irradiated data.

account of differential focusing effects. Any depreciation in the flux measured in the irradiated channels relative to the calibrated reference flux is now attributable to radiation-induced CTI effects.

It is worth mentioning here that in all three campaigns there were many tests where the reference data were collected with a magnitude that differed to the magnitude the irradiated tests were performed at. In these cases a further correction had to be made to the calibration. In RC1 the assumption was made that both the reference and irradiated flux levels in the first non-irradiated channel were the same. Then by multiplying the reference flux in every channel by the ratio between these two fluxes in the first channel an approximation for the correct reference illumination level could be found. The approaches taken in RC2 and RC3 are slightly more complex. For this reason, they will be described in later sections when these particular campaigns are discussed in more detail.

Figure 2.8 illustrates how relative charge losses are extracted from the irradiated data when the magnitudes are the same and only a translation is required.

2.1.9 Comparing local non-irradiated and irradiated image profiles

The centroid bias and charge loss that have been extracted from the calibrations as well as the local LSF profiles that have been fitted to corresponding pairs of images in both the unirradiated and irradiated data can all be used to see how the image shape is distorted by radiation damage. This is important when considering how the leading and trailing edges of an image profile are affected by the action of CTI as a function of magnitude, DOB level, CI level, duration and delay, and increasing radiation level. Examples of these comparisons will be shown in each campaign analysis.

2.2 Radiation Campaign 1

Prior to RC1 a significant number of tests, all well documented, had already been carried out on the radiation effects expected for Gaia. In these early tests Astrium, together with Sira, found that their results for both centroid bias and charge loss agreed, but only as far as $G = 15$. Fainter images, it seemed, responded somewhat differently to expected trends. Since this result tallied with suggestions that an uncontrolled DOB may be the cause, Contract Change Notice no. 10 (CCN10), now referred to as RC1, was proposed to specifically address magnitudes fainter than 15 and a possible dependency on the level of DOB. Both EADS Astrium (GAIAFPA.RPT.00700.T.ASTR) and Short (AS-011) carried out early analyses on this dataset. A more detailed history of the motivation behind these tests can be found in these documents.

During the course of this research RC1 was analysed twice. The first (Brown & van Leeuwen (SWB-001)) concentrated on developing: a) an understanding of the CCD readout and calibration procedures; and, b) the tools to process and analyse the data. Much of the groundwork was done during this early period. The follow-up analysis (Brown & van Leeuwen (SWB-003)) appeared after work on RC2 had begun, since it became clear during this time how the calibration of the mask rotation could be improved. As a result, it also became possible to quantify CTI effects over the full range of magnitudes investigated. In the discussions that follow both analyses will be presented concurrently. It will be made clear which analysis the results have come from, although for now the figures and tables are taken from the original analysis.

2.2.1 Hardware and test configuration

EADS Astrium (GAIAFPA.RPT.00445.T.ASTR) give detailed descriptions of the hardware and test setup. Further information can be found in EADS Astrium (GAIAFPA.RPT.00700.T.ASTR) when the same test bench was used during the study phase. Here a general overview is given.

RC1 ran between April and May 2006. In conjunction with previous tests, a back-illuminated irradiated AF CCD prototype was used, which was coupled to a PEM in order to operate the CCD. The PEM has several components including an analogue-to-digital converter. The conversion factor or gain in RC1 was set to $3.56 \text{ e}^-/\text{LSB}$. Placed between the light source (a Light Emitting Diode (LED) operating at 626nm) and the CCD was a rectangular aperture to mimic one of Gaia's viewing windows, an optical mask and a lens. Mounted on a translation table these three components, together with the light source, were able to move both horizontally and vertically relative to the CCD, ensuring that the TDI operating mode could be replicated and optical resolution was diffraction-limited. Earlier studies had shown that if a CCD operating in TDI mode was cooled to 163K the dependency between trapping and transfer rate in both the image area (IA) and serial register (SR) could be balanced, i.e. CTI effects could be minimised in both zones. For this reason the same temperature was chosen for RC1.

In order to extract positional bias and charge loss information, one third of the CCD was irradiated with a fluence of $4 \times 10^9 \text{ 10MeV proton/cm}^2$, and another third with $1 \times 10^{10} \text{ 10MeV proton/cm}^2$. According to predicted irradiation levels (see e.g. EADS Astrium (GAIA.ASF.TCN.PLM.00071): Figure

2/1), the former corresponds closest to the level Gaia is likely to experience in-orbit at the end of life. Therefore, only this part of the CCD was used in the testing. The remaining third was left unirradiated in order to provide reference data.

Figure 2.9 shows details of the AF CCD and optical mask in RC1. The optical mask measures less than 1cm^2 and during irradiated test runs was carefully positioned to pass over the boundary between the non-irradiated and irradiated parts of the CCD. Between these two regions the radiation level varied gradually such that stars E, F and G passed entirely within this transition region. Star D was also reported to have been partially damaged. The design of the mask incorporated a leading column of 10 holes labelled A, B1, C, D1, E, F1, G, H, I and J, followed by a smaller number of intermittent holes. Although there are data on all stars beyond the leading column, including stars D4-6 and J2 that were not identified in EADS Astrium (GAIAFPA.RPT.00700.T.ASTR), these have not been considered in this analysis. A more thorough examination of stars that follow the leading column was made in RC2.

RC1 was concerned with the behaviour of signals fainter than $G = 15$. As such, six different intensity levels were considered: 60000, 7000, 2000, 650, 400 and 200 e^-/window . The reference test was carried out with a signal close to 60000 e^-/window . These levels are equivalent to the flux integrated over each binned PSF. The link between the total number of electrons under the PSF and the corresponding G-band magnitude (G) can be found using the count-rate-magnitude relation defined in the Gaia Parameter Database (GPD). At the time of writing this relationship was defined as follows:

- The magnitude corresponding to 1 photoelectron s^{-1} , including telescope transmission (mirror reflectivity, mirror rugosity, and mirror contamination) and CCD quantum efficiency = 25.6571 mag at $1 \text{e}^- \text{s}^{-1}$.

Eqn. 2.2 shows how an approximation⁶ for G may be found using the preceding relation:

$$G = m_{Ref} - 2.5 \log_{10}\left(\frac{F[\text{e}^- \text{s}^{-1}]}{1[\text{e}^- \text{s}^{-1}]}\right), \quad (2.2)$$

where m_{Ref} is the magnitude corresponding to 1 photoelectron s^{-1} given above and F is the total flux arriving in units of $\text{e}^- \text{s}^{-1}$. This can be determined by dividing the total flux under the PSF by the integration time, t_{int} . This is the time it takes a charge packet to cross a CCD: $t_{int} \approx 4.42\text{s}$. Table 2.1 gives approximate G values for the signal intensity levels used in RC1.

Total flux [e^-/window]	G
60000	15.33
7000	17.66
2000	19.02
650	20.24
400	20.77
200	21.52

Table 2.1: Signal intensities in terms of the total flux and G in RC1.

⁶In reality, there are many factors eqn. 2.2 does not take account of including the star colour, the FOV, the precise location in the focal plane and the telescope wave-front error.

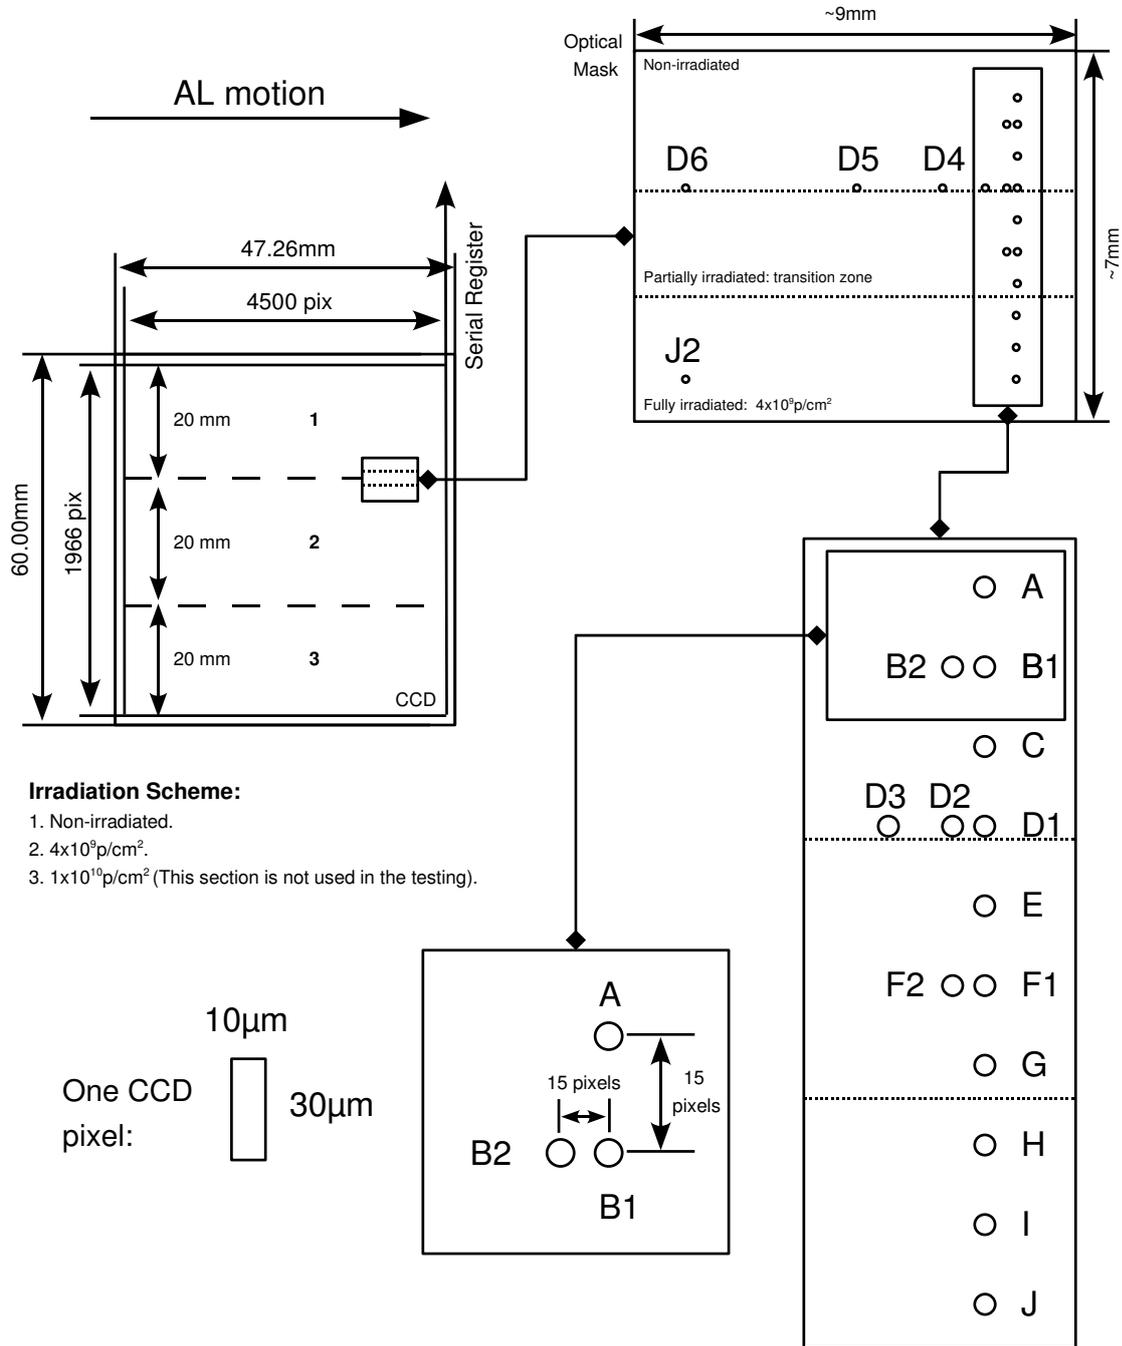


Figure 2.9: Details of the test setup in RC1. In the top left the configuration of the optical mask with respect to the CCD during the irradiated test runs is shown. Stars D2 and D3 are separated by ~30 pixels, stars D3 and D4 by ~60 pixels, stars D4 and D5 by ~120 pixels and lastly, stars D5 and D6 by ~240 pixels. Star J is separated from J2 by ~465 pixels.

Star	Channel	Sample index		AC position [pix]		CCD area	
		Reference	Irradiated	Reference	Irradiated	Reference	Irradiated
Prescan	-	1	1	1	1	-	-
A	1	7	6	420	550	NI	NI
B*	2	8	7	435	565	NI	NI
C	3	9	8	450	580	NI	NI
D*	4	10	9	465	595	NI	NI/PI
E	5	11	10	480	610	NI	PI
F*	6	12	11	495	625	NI	PI
G	7	13	12	510	640	NI	PI
H	8	14	13	525	655	NI	I
I	9	15	14	540	670	NI	I
J*	10	16	15	555	685	NI	I

*Multiple stars are read out from these AC positions.

Table 2.2: Sampling scheme and associated CCD readout area in RC1. The AC position corresponds to the first of 12 pixels binned AC. NI = non-irradiated, PI = partially irradiated, I = irradiated.

In addition to varying the intensity level, three levels of DOB: 0, 5 and 10 $e^-/\text{pixel}/t_{int}$, were also tested. Given the time it takes to cross one CCD, these DOB levels equate to 0, ~ 1.1 and ~ 2.3 $e^-/\text{pixel}/s$. No DOB was applied in the reference case. For brevity, the units for signal intensity and DOB from this point on shall be omitted.

For every signal/DOB test case 20 scans were performed. The only exception being 60000/0 where 10 scans were performed.

Table 2.2 shows which part of the sampling scheme was analysed in RC1. In total, 22 samples were read out after each TDI period, but only 10 samples contained data pertaining to the stellar field. In the reference data the 10 leading stars were found in samples 7 to 16 inclusive, whereas in the irradiated data the leading stars were found in samples 6 to 15. To isolate the stellar field the sample indices containing the stellar field were labelled channels 1 to 10. Recall that a sample corresponds to a single, binned charge packet read out after each TDI period, whereas a sample *index* indicates the AC position in the read out. The two are essentially synonymous in the SR. See Figure 2.2 for an illustrative example of a sampling scheme used in RC2.

2.2.2 Background analysis

To accentuate features in the background, Figure 2.10 shows a running mean of every scan read out from sample 1 (top) and sample 7 (bottom) in the reference data. The running mean ϵ is defined as a set of arithmetic means calculated over the length of each scan:

$$\epsilon = \left[\left(\frac{\sum_{i=1}^{i=n} x_i}{n} \right)_{j=1}, \left(\frac{\sum_{i=n+1}^{i=2n} x_i}{n} \right)_{j=2}, \left(\frac{\sum_{i=2n+1}^{i=3n} x_i}{n} \right)_{j=3}, \dots, \left(\frac{\sum_{i=(N-1)n+1}^{i=Nn} x_i}{n} \right)_{j=N} \right], \quad (2.3)$$

where n is the number of samples per interval j , in this case 20, and $j = 1, \dots, N$ where N is the total number of allowed intervals over the scan. Recalling that sample 1 contains only a record of the

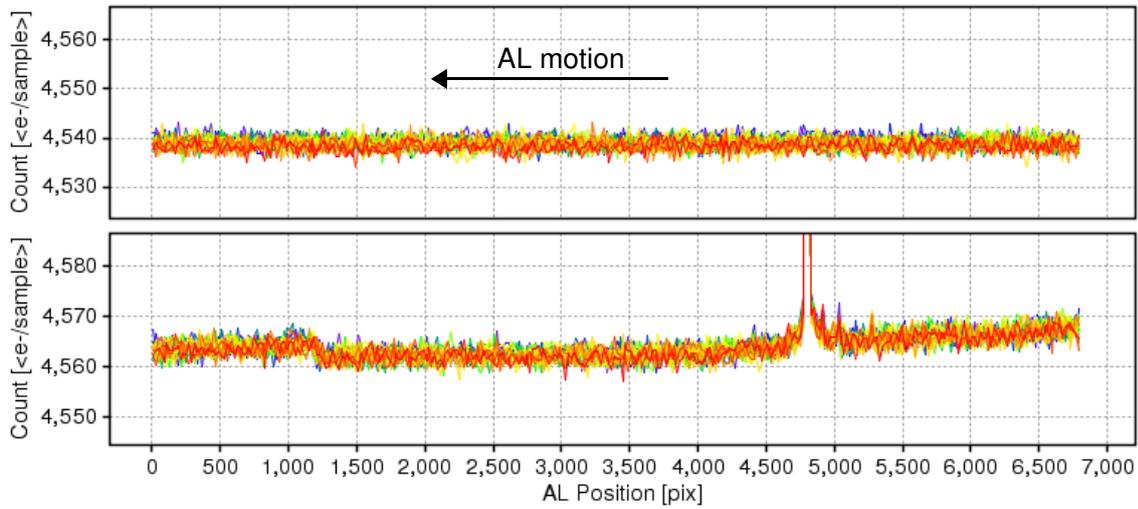


Figure 2.10: Running means of all scans from samples 1 (top) and 7 (bottom) in the reference data.

electronic offset and RON⁷, the running mean in sample 1 appears stable over the entire length of the scan. The mean electronic offset measured in this sample over all scans is $4538.86 \pm 0.13 \text{ e}^-/\text{sample}$. The scatter is given by the standard error on the mean. Sample 7 is the first sample in the reference data to contain a stellar image (star A) that can be seen to appear soon after 4500 pixels have been read out. Compared to the prescan, the background is seen to vary. Though these variations are slight, there are two noticeable features: 1) a discontinuity around 1200 pixels and 2) a small positive gradient appearing roughly halfway down the readout. The overall background level is also different between the two samples with sample 7 showing a background at least 20e^- higher than that seen in sample 1. It was originally expected that the mean count in the prescan sample should be the same in every sample. However, Figure 2.10 clearly shows that this is not the case, and can be verified for all other samples. An estimate of the background must, therefore, be made for each sample. In section 2.1.2 the region over which the mean background is extracted was stated as being between 500 and 1500 pixels after the readout began. In RC1, both the discontinuity and gradient observed in Figure 2.10 must be avoided, so instead, the mean background is estimated between 1500 and 3000 pixels after readout. This makes the mean electronic offset for sample 7 in the reference data to be $4562.29 \pm 0.15 \text{ e}^-/\text{sample}$.

Figure 2.11 shows a running mean of every scan from sample 6 in the irradiated data with a signal intensity of 60000 and DOB of 0 (top), 5 (middle) and 10 (bottom) respectively. Without a DOB the running mean is very similar to that observed in sample 7 in the reference data; the mean electronic offset is $4562.21 \pm 0.14 \text{ e}^-/\text{sample}$. By adding a DOB both the overall background level and the noise are seen to increase. The discontinuity at the beginning of each scan and the gradient underneath each image continue to appear in the irradiated data.

Figure 2.12 shows background noise distributions in both the reference and a subset of the irradiated data. The plot on the left shows the noise level in sample 1 and the overall noise level in samples 7 to 16

⁷The prescan is read out from a part of the CCD that is physically covered by a mask. In flight, the prescans will be used to monitor the electronic offset and RON.

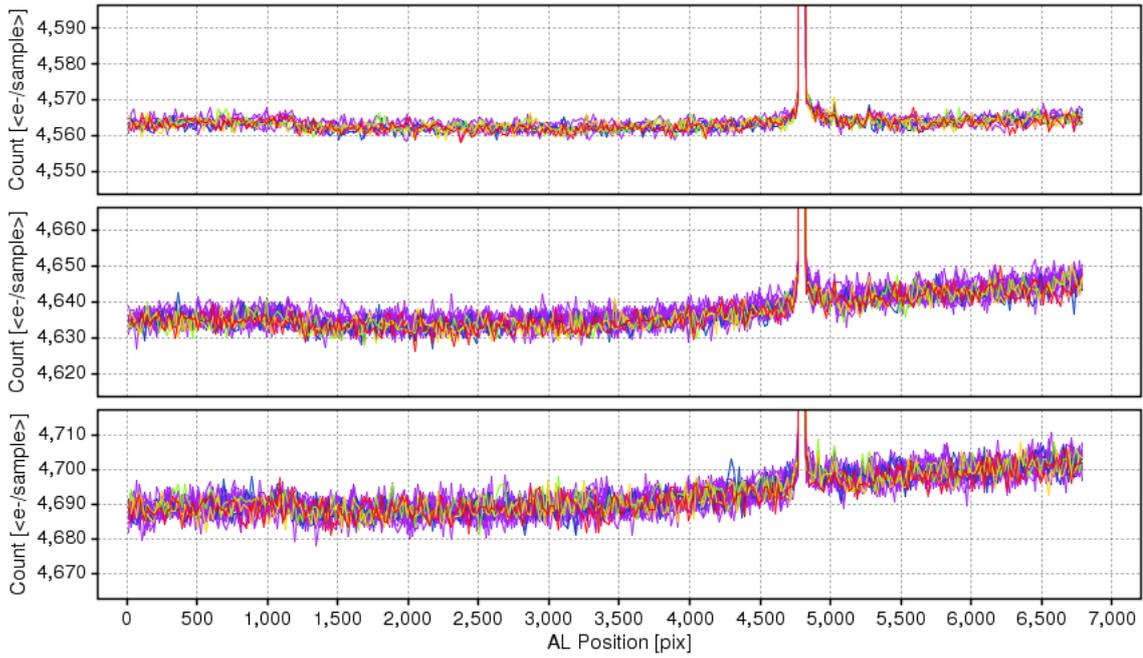


Figure 2.11: Running mean of all scans from sample 6 in the irradiated data for 60000 with a DOB of 0 (top), 5 (middle) and 10 (bottom).

in the reference data. The average noise level in sample 1 ($5.16 \pm 0.10 e^-/\text{sample}$) is equivalent to the RON. In contrast the average noise in samples 7 to 16 is $5.76 \pm 0.12 e^-/\text{sample}$. This indicates a DOB level of $0.55 \pm 0.10 e^-/\text{pixel}$. The relation between the total noise level and its components, including the conversion between samples and pixels, is given by the following equation:

$$DOB = \frac{1}{12} (\sigma_T^2 - RON^2), \quad (2.4)$$

where σ_T is the mean total background noise level over all channels. This assumes, firstly, that the noise on the DOB is purely photon noise and, secondly, that an estimate of the noise in each pixel can be found by dividing the DOB found in each sample by the total number of pixels per sample, i.e. 12.

The error on the DOB has been calculated using standard error analysis relations associated with

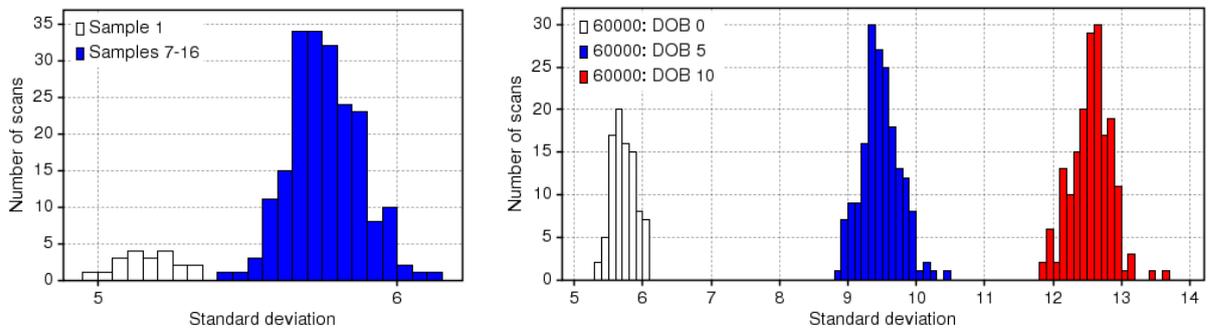


Figure 2.12: Background noise distributions. (left) Noise in sample 1 and samples 7 to 16 from the reference data. (right) Noise in samples 6 to 15 at 60000 and all three levels of DOB.

multiplying and subtracting values with standard deviations. The errors on σ_T and the RON have been determined from the noise distributions as shown, for example, in Figure 2.12.

The second plot in Figure 2.12 shows comparable results for the background noise level in tests with a signal of 60000 and all three levels of DOB. Without a DOB the average noise level is measured to be 5.70 ± 0.16 e⁻/sample, equating to a DOB level of 0.53 ± 0.12 e⁻/pixel. This is essentially the same as the background level seen in the reference data. With a nominal DOB of 5 and 10 the observed background levels are 9.49 ± 0.28 e⁻/sample and 12.60 ± 0.29 e⁻/sample respectively, indicating an observed DOB level of 5.36 ± 0.32 e⁻/pixel and 11.08 ± 0.44 e⁻/pixel respectively. The fact that all the observed DOB levels are slightly higher than expected suggest it may be difficult to completely remove all low-level background light in the test facility.

Table 2.3 gives a summary of the mean background noise levels observed in RC1. Throughout the campaign the RON remains consistent around 5.1 e⁻/sample and the observed DOB levels are very close to their nominal values.

Figure 2.13 shows the mean background flux as a function of AC position. The scatter is given by the standard error on the mean. The plot on the left is for the reference data, which clearly shows a systematic reduction in the background between samples 15 and 16 of the order of $\sim 40e^-$. The plot on the right shows comparable results for a signal of 60000 with all three levels of DOB. Aside from the expected increase in the overall background with increasing DOB, the systematic reduction in the last sample is again clearly present. This effect cannot be occurring in the image zone since the absolute AC position, as indicated in Figure 2.13 (note that the AC positions are different between the left and right plots), has shifted over to cover a different part of the CCD. Astrium suggest that the root of this effect can be attributed to an instability between the CCD and PEM. Now known as the 'PEM anomaly', the

Signal [e ⁻ /window]	Target DOB [e ⁻ /pixel]	RON [e ⁻ /sample]	σ_T [e ⁻ /sample]	Observed DOB [e ⁻ /pixel]
Reference	0	5.16 ± 0.10	5.76 ± 0.12	0.55 ± 0.10
60000	0	5.11 ± 0.10	5.70 ± 0.16	0.53 ± 0.12
	5	5.07 ± 0.14	9.49 ± 0.28	5.36 ± 0.32
	10	5.08 ± 0.14	12.60 ± 0.29	11.08 ± 0.44
7000	0	5.09 ± 0.07	5.34 ± 0.09	0.22 ± 0.07
	5	5.08 ± 0.08	9.44 ± 0.23	5.28 ± 0.26
	10	5.09 ± 0.07	12.60 ± 0.25	11.07 ± 0.37
2000	0	5.05 ± 0.08	5.34 ± 0.09	0.25 ± 0.07
	5	5.08 ± 0.08	9.44 ± 0.24	5.28 ± 0.27
	10	5.11 ± 0.09	12.48 ± 0.27	10.80 ± 0.40
650	0	5.10 ± 0.09	5.32 ± 0.09	0.19 ± 0.08
	5	5.07 ± 0.08	9.27 ± 0.18	5.02 ± 0.20
	10	5.09 ± 0.06	12.48 ± 0.30	10.82 ± 0.44
400	0	5.09 ± 0.08	5.33 ± 0.09	0.21 ± 0.07
	5	5.09 ± 0.08	9.28 ± 0.20	5.02 ± 0.22
	10	5.10 ± 0.08	12.58 ± 0.29	11.02 ± 0.43
200	0	5.13 ± 0.09	5.54 ± 0.12	0.36 ± 0.10
	5	5.07 ± 0.08	9.29 ± 0.19	5.05 ± 0.21
	10	5.09 ± 0.07	12.63 ± 0.19	11.13 ± 0.29

Table 2.3: Statistics on the noise in the background in RC1.

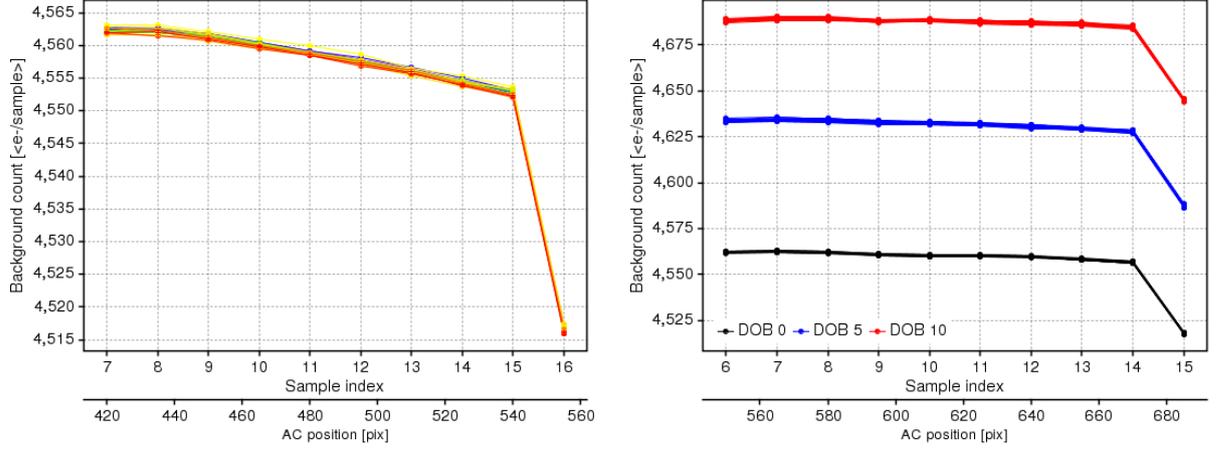


Figure 2.13: Mean background flux as a function of AC position in (left) the reference data and (right) for 60000 with a DOB of 0, 5 and 10.

two causes for this instability (simultaneous transitions in the IA and SR as well as SR flushing) have been investigated by EADS Astrium (GAIA.ASF.TCN.PLM.00368). The impact on astrometry has also been looked at, indicating it is the fainter stars ($G > 15$) that fair worse, even for moderately small offsets (see Figure 4 in de Bruijne (JDB-048)).

In Figure 2.13 the mean background in every scan is shown. To quantify how the mean background in every channel behaves it is firstly assumed that the mean background in each scan \bar{b}_s ($s = 1, \dots, N$) follows a Gaussian distribution, where s is the scan number and N is the total number of scans per star. The mean background in each channel, \bar{b}_c , can then be estimated as follows:

$$\bar{b}_c = \frac{\sum_s^N \left(\frac{\bar{b}_s}{\sigma_s^2} \right)}{\sum_s^N \left(\frac{1}{\sigma_s^2} \right)}, \quad (2.5)$$

where σ_s is the standard error on the mean background per scan. Since N is the same in every channel the mean standard error per channel is estimated by averaging over the standard error in each scan:

$$\bar{\sigma}_c = \frac{\sum_s^N \sigma_s}{N} \quad (2.6)$$

An estimate of the standard deviation, σ , for the mean background per channel is given by:

$$\sigma = \sqrt{\frac{\chi^2}{n}} = \sqrt{\frac{\chi^2}{N-1}} = \sqrt{\frac{1}{N-1} \sum_s^N \left(\frac{\bar{b}_s - \bar{b}_c}{\bar{\sigma}_c} \right)^2}, \quad (2.7)$$

where χ^2 is the weighted, squared sum of the mean background residuals and $n = N - 1$ is the number of degrees of freedom.

Table 2.4 summarises these results for the reference data. The mean background level decreases gradually from channel 1 to 9 dropping sharply in channel 10, while the standard error and standard deviation remain consistent across all channels. The last column is a statistic g that gives a standard Gaussian ap-

Channel	Star	$\bar{b}_c \pm \bar{\sigma}_c$ [<e ⁻ /sample>]	χ^2	σ	g
1	A	4562.29 ± 0.15	132.67	2.64	8.53
2	B	4562.14 ± 0.15	159.32	2.90	9.65
3	C	4561.35 ± 0.15	117.98	2.49	7.86
4	D	4560.07 ± 0.15	123.73	2.55	8.13
5	E	4558.84 ± 0.15	105.51	2.36	7.24
6	F	4557.49 ± 0.15	142.15	2.74	8.95
7	G	4555.98 ± 0.15	93.94	2.22	6.61
8	H	4554.34 ± 0.15	137.35	2.69	8.74
9	I	4552.70 ± 0.15	166.68	2.96	9.93
10	J	4516.55 ± 0.15	138.15	2.70	8.78

Table 2.4: Global statistics on the background in the reference data

proximation for the χ^2 distribution. It is sometimes referred to as the cube-root approximation (Wilson & Hilferty (1931)) and is useful even for small n .

$$g = \frac{1}{\sqrt{2/9n}} \left[\left(\frac{\chi^2}{n} \right)^{\frac{1}{3}} - \left(1 - \frac{2}{9n} \right) \right]. \quad (2.8)$$

For N independent measurements of X_i , which are taken from a normal distribution, the cumulative distribution function of χ^2 ($= \sum X_i^2$), $\Phi(\chi^2)$ approaches that of $\Phi(g)$, where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp(-\frac{t^2}{2}) dt$, t being a (dummy) integration variable.

There are too few measurements to say with any certainty that the background, in the majority of cases, does not follow a Gaussian distribution. However, the strong bias towards positive values does suggest that, even if a greater number of channels were considered, the mean background level would still not be distributed randomly. Evidence to date suggests that such systematics are likely to be linked to the PEM anomaly. This is investigated further as global variations in the background are considered.

To see how the background measured in one channel may depend upon the background measured in another channel a correlation matrix was generated. The correlation coefficient, r_{ij} , has been determined for all possible combinations of the background measured in each channel containing a leading star:

$$r_{ij} = \frac{\sum_s^N (x_{i,s} - \bar{x}_i)(x_{j,s} - \bar{x}_j)}{\sqrt{\sum_s^N (x_{i,s} - \bar{x}_i)^2 \sum_s^N (x_{j,s} - \bar{x}_j)^2}}, \quad (2.9)$$

where $x_{i,s}$ and $x_{j,s}$ are the mean background values for stars i ($i = A, \dots, J1$) and j ($j = A, \dots, J1$) summed over s scans, $s = 1, \dots, N$. \bar{x}_i and \bar{x}_j represent the weighted mean of the background per channel for stars i and j . Table 2.5 shows the correlation matrix for the reference data.

If the background level and noise were consistent with the assumed statistics no correlations would be expected between channels. However, this is clearly not the case here. There are significant positive correlations between the background values found across all 10 leading stars, indicating that there are global variations in the background common to all channels. The fact that the variations appearing in the background are not reflected by the flux levels recorded over the stellar field also suggest that

Channel	Star	A	B	C	D	E	F	G	H	I	J
1	A	1.00	0.77	0.83	0.83	0.83	0.88	0.88	0.85	0.82	0.78
2	B	-	1.00	0.82	0.83	0.78	0.88	0.81	0.75	0.85	0.84
3	C	-	-	1.00	0.91	0.84	0.84	0.81	0.88	0.89	0.91
4	D	-	-	-	1.00	0.85	0.83	0.80	0.89	0.89	0.83
5	E	-	-	-	-	1.00	0.92	0.80	0.89	0.86	0.86
6	F	-	-	-	-	-	1.00	0.86	0.90	0.87	0.87
7	G	-	-	-	-	-	-	1.00	0.82	0.79	0.82
8	H	-	-	-	-	-	-	-	1.00	0.87	0.84
9	I	-	-	-	-	-	-	-	-	1.00	0.89
10	J	-	-	-	-	-	-	-	-	-	1.00

Table 2.5: Correlation matrix for variations in the mean background in the reference data.

these variations are occurring outside the IA. It is known that the PEM adds a systematic variation to the background between channels, which suggests that the PEM anomaly is also the cause of these correlations in the background variations.

Figure 2.14 shows a running mean (see eqn. 2.3) of the background flux from sample 7 in the reference data (left) and sample 6 in the irradiated data (right) with a signal of 60000 and a DOB of 10 after the stellar image has passed. Data from all available scans are shown. The horizontal blue line indicates the estimate of the mean background level. Though the slow rise in the background level has already been seen in Figure 2.10 and Figure 2.11, what is more interesting here are the small jumps in flux located around positions 120, 245 and 470. They appear in every scan. Closer inspection of the raw data has revealed that the locations of these extra peaks coincide with the positions of holes 4, 5 and 6 attributed to star D. Since Figure 2.14 is focused on star A this result implies that by binning the signal, the AC wings of star D are recorded in the transit of star A. In the AC direction these two stars are separated by 45 pixels.

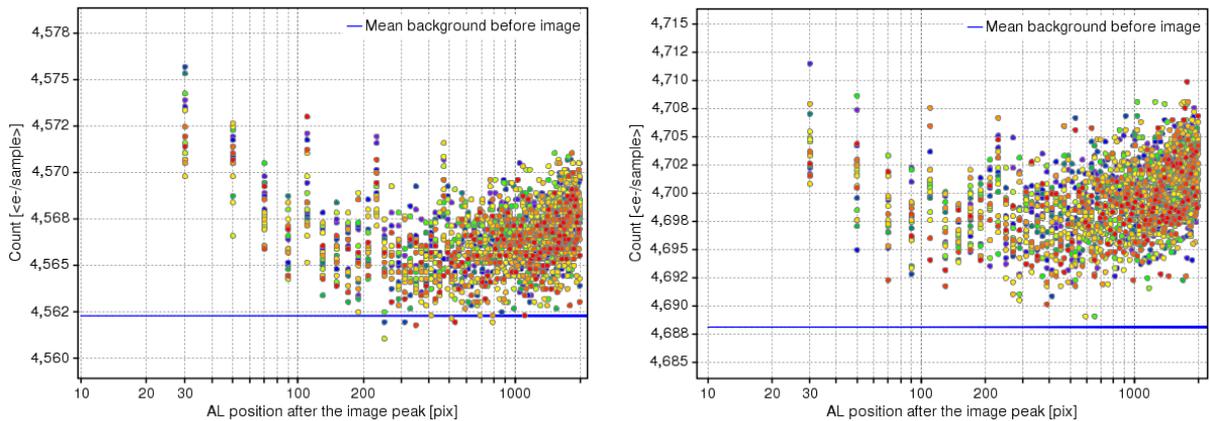


Figure 2.14: A running mean of background fluxes taken from the (left) reference data and (right) irradiated data for 60000 with a DOB of 10 for star A, following the transit of the stellar image.

To illustrate this effect, Figure 2.15 shows a section of the first scan read out from channel 3 in the reference data. Star C, the dominant signal on the left of the plot, is the only signal that is expected. However, 5 additional images are picked up. These images are aligned with the AL positions of stars D2

(plus B2, F2 etc) to D6, and clearly impart a signal 150 to 200 e^- /sample above the background. Why the additional images vary in height remains unclear. It could be that the source of illumination is directly above D5 or that the hole size varies across the mask.

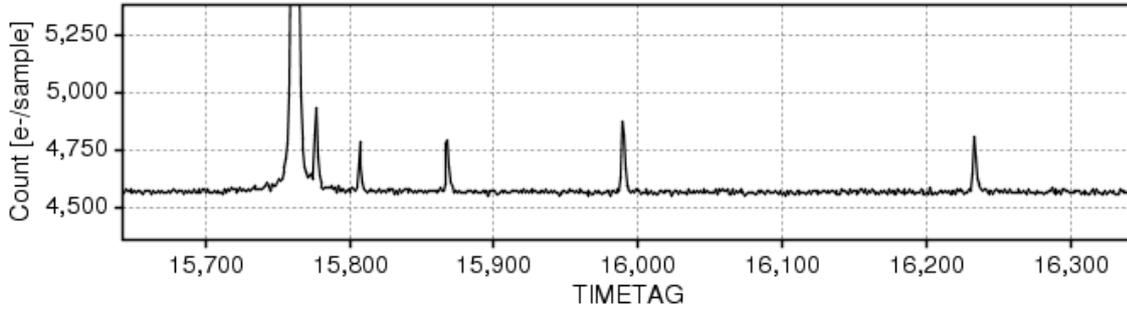


Figure 2.15: Raw data from scan 1 with star C.

Channel	Star	No. expected images	No. additional images	Approximate flux [%]	Implied dominant source
1	A	1	1	1.00	B2
2	B	2	4	0.40	D3, D4, D5, D6
3	C	1	5	1.30	D2, D3, D4, D5, D6
4	D	6	-	-	-
5	E	1	5	1.20	D2, D3, D4, D5, D6
6	F	2	4	0.15	D3, D4, D5, D6
7	G	1	1	0.80	F2
8	H	1	1	0.20	F2
9	I	1	1	0.30	J2
10	J	2	-	-	-

Table 2.6: Consequences of pinhole placement: detecting AC image profiles in neighbouring channels.

Table 2.6 summarises the effects described above. The number of additional images are shown in each channel and their intensities are indicated as a percentage of the flux of the expected image. However small the intensity of the AC profile, stars D2 to D6, which tend to appear always brighter than the surrounding stars, clearly overlap with the transits in other channels.

What is further implied by these overlaps, but which cannot be shown explicitly in this analysis, is the influence of the AC profile of neighbouring stars in the leading column. If star D2 imparts an additional flux in the transit of C of $\sim 1.3\%$ then a significant amount of overlap between stars in the leading column would not be expected; as the stars become fainter the overlap should decrease further. Calculating the correct flux for a single image when partially blended with another is a situation that will have to be dealt with as part of the Gaia data processing. Here though, as the images in the leading column cannot be deblended and the extent of overlap is very small, even for the brightest images, it is assumed that the blending has negligible effect upon the CTI measurements.

To summarise:

- The background signal in the CCD readout comprises a large electronic offset with (Gaussian) RON, and a much smaller (Poisson) DOB.
- The signal in the prescan (sample 1), which contains only the electronic offset, is stable and contains no anomalies. The RON in this sample is consistent throughout the campaign and measures $\sim 5.1 e^-/\text{sample}$.
- In the AL direction the signal in the IA contains a systematic discontinuity after ~ 1200 pixels have been read out and begins to increase linearly halfway through each scan.
- In the AC direction the signal is not uniform and contains a large systematic discontinuity ($\sim 40 e^-/\text{sample}$) between the last two channels. This is attributed to the PEM anomaly.
- Further statistical analysis supports the assertion that an instability between the CCD and PEM is the cause of unexpected systematics, correlations and global variations between the background signals measured between channels. If the PEM anomaly were not to exist it is thought that the background level and noise would be consistent with the expected statistics.
- All of the observed DOB levels are slightly higher than expected. This suggests a limitation; the test facility is unable to completely remove all low-level background light.
- The spacing of pinholes on the mask causes the AC wings of each image to encroach upon the transit of neighbouring images. This may be a problem for very bright images spaced close together as in RC2 and RC3. However, this will be difficult to decouple as the stars are placed on a regular grid and all images are binned on readout. In RC1 this will not be a significant problem, since image intensities are faint, and the flux level is never much greater than 1% of the expected image.

Two alternative methods for subtracting the background

The majority of results presented so far have come from the first analysis of RC1. As shown, it was known that the stellar field always appeared later in the readout. Therefore, it was natural to estimate the background over an initial part of the scan. Subtracting a flat cut ensured the technique was consistent. In the reanalysis of RC1, two other methods for estimating the background were explored. The first looked at fitting a straight line to the background, seen rising in the second half of the readout. The only problem with this approach was the uncertainty surrounding the position at which the increase in the background could be said to begin. To avoid this ambiguity a third method was tried. This is described in Lindegren (LL-069), and uses the fact that at some point along the scan, the wings of an image must merge with the background signal. At this point the background could be estimated locally. Since RC1 concentrates primarily on fainter images, the sample values on the edge of a 15-pixel window lie close enough to the background to warrant use of this method. I mention this here as some of the results presented in the next section have used this method to extract the background.

2.2.3 The reference data

The Global Reference Curve (GRC)

As described in section 2.1.5, the first step in extracting relative positions and charges is to construct a GRC. Figure 2.16 shows four of the steps that have been taken to define the GRC. In the top left, every raw image for every star across the mask has been extracted and stacked. The peak flux is defined to be the initial zeropoint. In the top right, every image has been repositioned following: a) cross-correlation with a reference image; and, b) quadratic interpolation of the cross-correlation curve. The non-uniformity of the mask illumination is clear in this plot. To constrain a curve fit, the images are scaled in both height and width to a temporary reference curve. Once this was done the zeropoint was aligned to the centroid

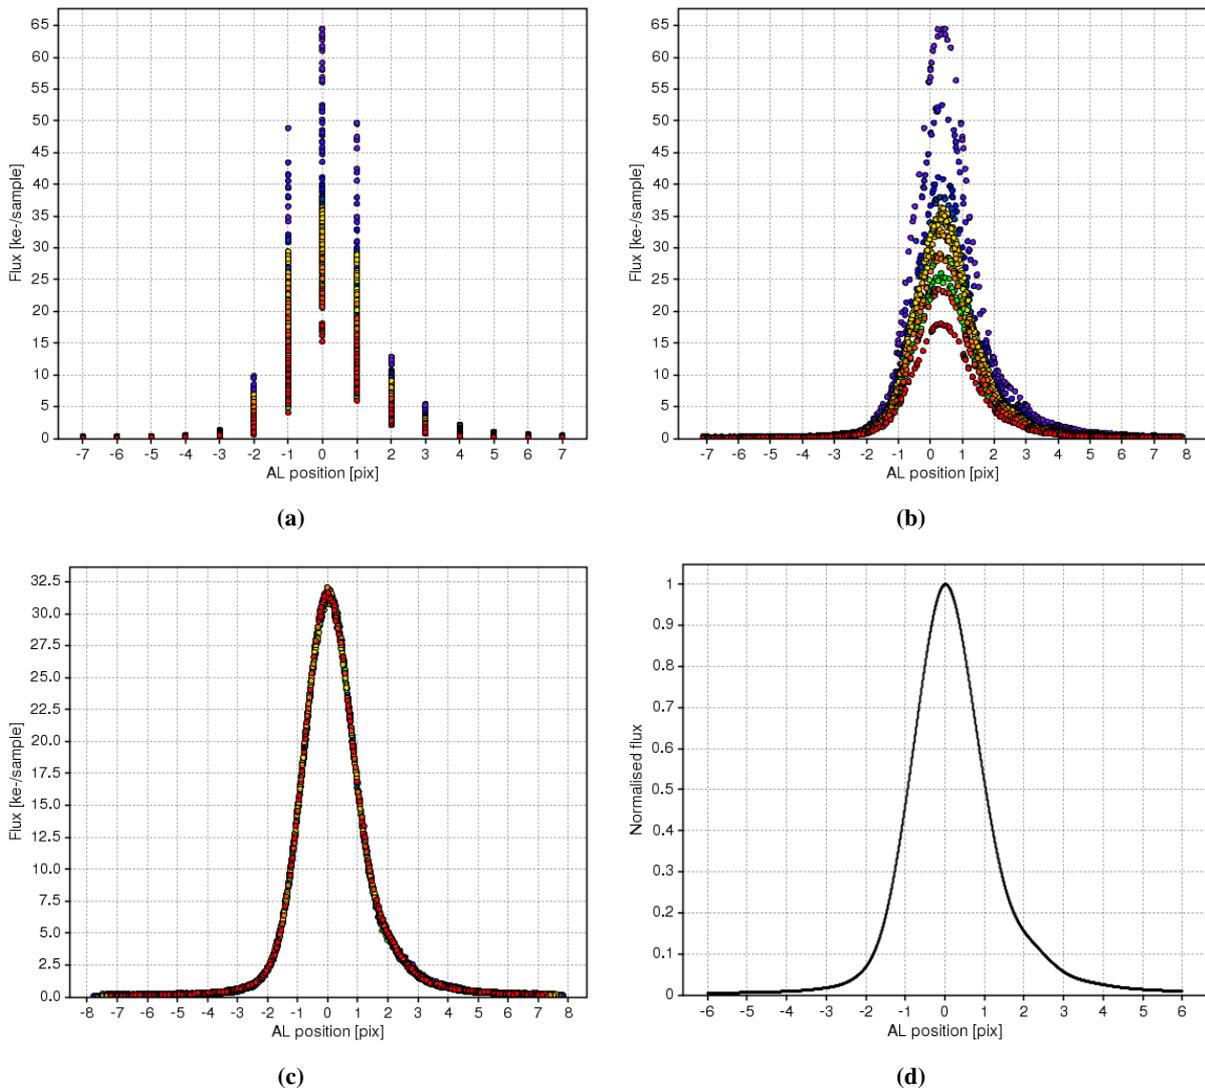


Figure 2.16: Constructing the GRC. (a) Stacking the raw images (column 1). (b) After cross-correlation and quadratic interpolation. (c) After image scaling with iterations to remove positional errors and defining the centroid. (d) The normalised GRC (cubic spline with 15 knots spaced regularly).

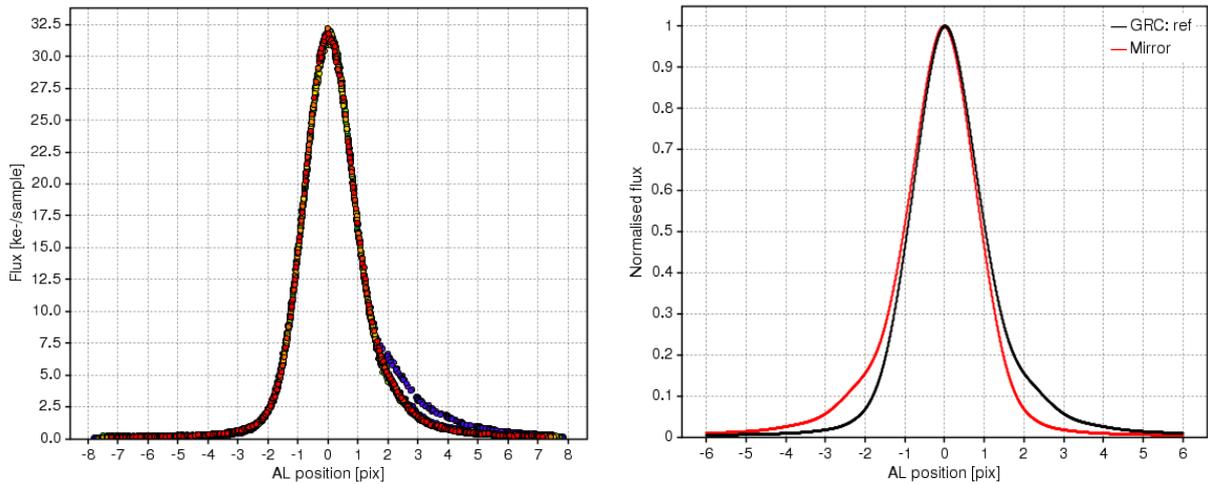


Figure 2.17: Unexpected features in the reference data. (left) With stars D6 and J2 included. (right) The asymmetry of the GRC.

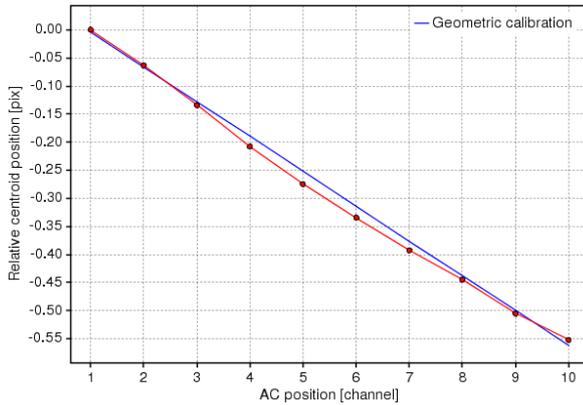
definition. The plot on the bottom left shows the outcome at this stage.

During the reanalysis an additional step was added before the centroid was defined. In parallel with analysing RC2, it was noticed that the final positions of the datapoints were not producing the correct shapes when a local LSF fit was made. To reconcile this, the image data was fitted to a new temporary curve several times over until the datapoint positions could no longer be adjusted to fit the curve any better. This iteration was the final step in determining the sub-pixel positions of the image data. The normalised GRC is shown in the bottom right-hand corner of Figure 2.16.

Figure 2.16 excludes the accumulated data from stars D6 and J2. If these two stars had been retained, the result following step 3 above would have looked like the plot on the left in Figure 2.17. The strong feature attached to the right-hand side of the scaled data is caused by these two stars. As their shapes are damaged, they were not included in the analysis. The plot on the right of Figure 2.17 shows the degree of asymmetry in the GRC, which is clearly prevalent over a large portion of the profile. This asymmetry was originally thought to arise from some form of CTI. If not radiation damage, then design traps, introduced when the CCD was manufactured. Feedback from Astrium explained that the asymmetry was almost certainly caused by a misalignment in the optics. This was indeed the case for stars D6 and J2, where the LSF shapes were known to be degraded, and the optics were not optimised for them.

Geometric and flux response calibration

Once the GRC has been constructed the relative centroid positions can be determined in the reference data. The relative centroid position as a function of channel number is shown in Figure 2.18, together with the geometric calibration. Using the zeropoint as a reference for the pivot point, which is fixed to the observed position of star A, the geometric calibration shows that the mask has been rotated about an axis perpendicular to the CCD and mask by $\sim 0.004^\circ$, such that the stars in channel 10 are read out ~ 0.55 pixels before they would otherwise have been read out. The geometrical offsets between the calibration



Star	Relative centroid position [pix]		AL offset [pix]
	Observed	Expected	
A	0.000	-0.004	0.004
B1	-0.063	-0.066	0.003
C	-0.134	-0.128	-0.006
D1	-0.207	-0.190	-0.017
E	-0.274	-0.252	-0.022
F1	-0.335	-0.314	-0.021
G	-0.392	-0.376	-0.016
H	-0.445	-0.438	-0.007
I	-0.505	-0.500	-0.005
J1	-0.552	-0.562	0.010

Figure 2.18: (left) Relative centroid positions for the 10 leading stars with geometric calibration. (right) Geometrical offsets remaining after the calibration.

and the relative centroid positions are shown on the right-hand side of Figure 2.18. They are attributed to physical misalignments on the mask.

The flux response across the leading column of the mask is calibrated at the same time as the centroid positions. On the left of Figure 2.19 the GRC (red) has been fitted to the accumulated data for star A. To see how well the GRC describes the accumulated data, a local LSF has also been fitted (blue). The two are almost indistinguishable. The residual flux between these two fits is also shown and can be seen to deviate only slightly around the peak region and where the gradient is changing quickly on the slopes. On the right of Figure 2.19 the peak fluxes measured for every raw image in the leading column are plotted against the peak of the fitted GRC to that image. The spread in raw peak values illustrates the variability in the scan timing and/or the rotation of the mask between scans.

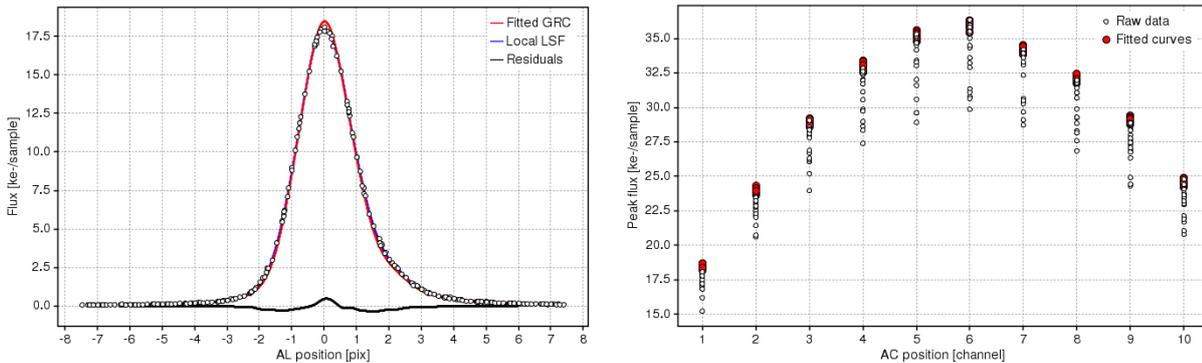


Figure 2.19: Flux response calibration I. (left) Curve fitting to star A in the reference data. (right) Raw peak fluxes are shown against peak fluxes of the fitted GRC to each image.

Figure 2.20 contains two plots. On the left, the result of extracting the residuals between the fitted GRC and local LSF after matching peak fluxes is shown. The variations that remain are primarily due to differences in the width of the two fits. It suggests that the local LSF is wider than the fitted GRC in star A but narrower than the GRC in star J. This is confirmed in the second plot (taken from the reanalysis) which shows that the local LSF FWHM for each of the leading stars are also following this trend. As expected, the FWHM of the GRC does not change when the flux is scaled.

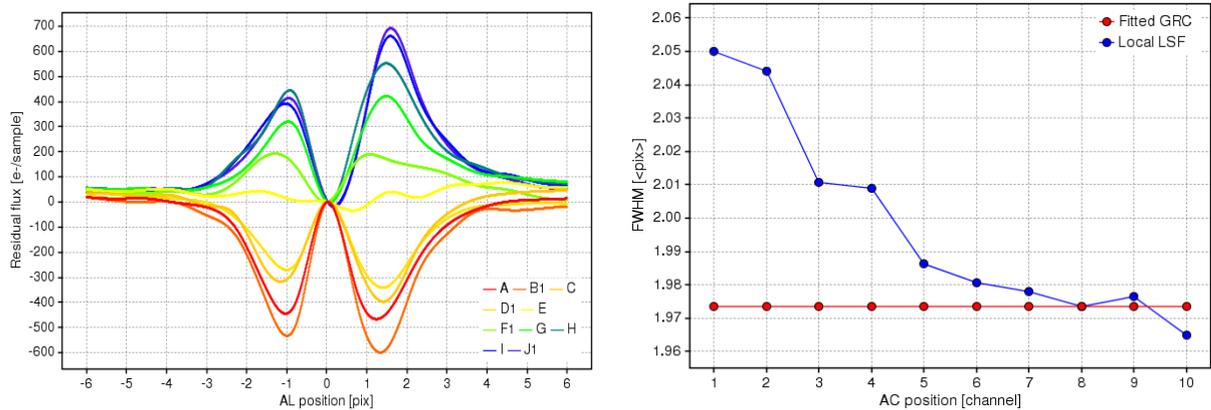


Figure 2.20: Flux response calibration II. (left) Residuals between the fitted GRC and local LSF after scaling both to normalise the illumination conditions across the mask. (right) Variations in the residuals seen on the left are reflected by the variations between fitted GRC and local LSF FWHMs. Note that the first plot is from the original analysis and the second is from the reanalysis.

The variation in the FWHM of the local LSF cannot be recovered by the GRC, which suggests the fitted GRC may not be able to accurately measure the charge under the image. This is not the case however, for if the peak flux of the fitted GRC and the local LSF is plotted as a function of the channel number, as shown on the left of Figure 2.21, the variation between the fitted GRC and local LSF reflects the differences between the FWHM values seen in Figure 2.20.

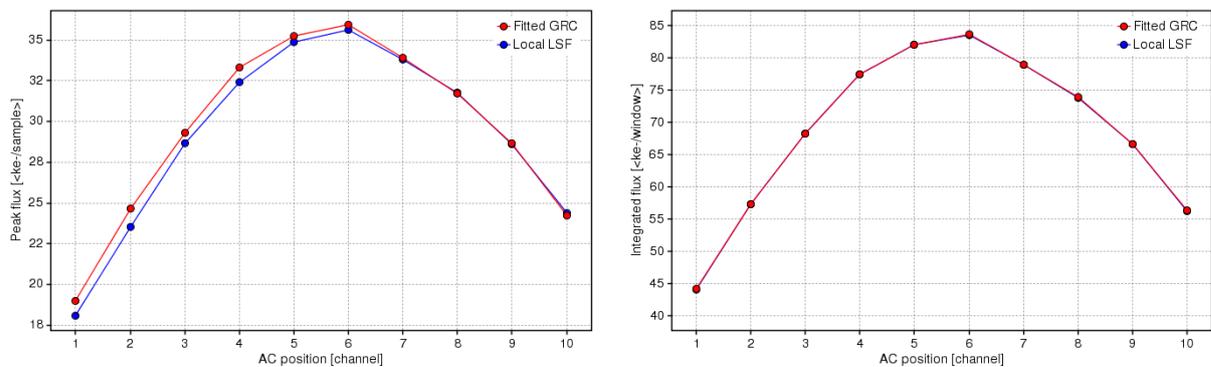


Figure 2.21: (left) The variation in the FWHM is mirrored by the variation in the peak flux between the local LSF and fitted GRC. (right) The fitted GRC replicating the integrated flux under the local LSF. Both plots are taken from the reanalysis.

When the local LSF FWHM is larger than the fitted GRC the local LSF peak flux is smaller than the fitted GRC. This indicates that the fitted GRC is able to accurately measure the *total charge* in the image. The second plot in Figure 2.21 confirms this by showing the variation between the integrated charge under both the fitted GRC and local LSF for all stars in the leading column. Within 1% the two are effectively the same. The parabolic shape of the curves illustrate the non-uniform illumination of the optical mask.

Customising the local LSF

Up to now the local LSF has been portrayed as a good representation of the actual LSF. This statement is now justified by showing how the spline functions used throughout this analysis may be customised to improve the reliability of the local LSF.

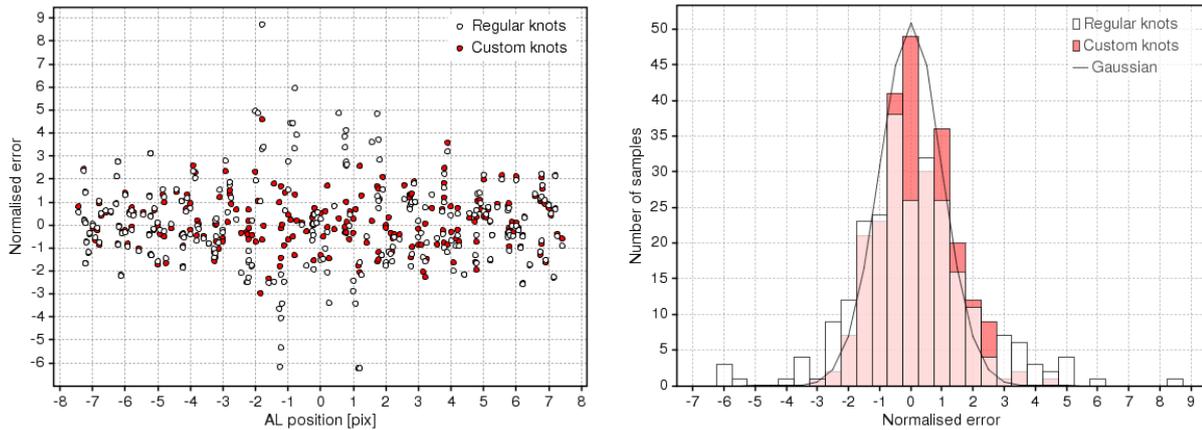


Figure 2.22: Regular vs custom-spaced knots. (left) The normalised error for each datapoint sampling the local LSF. (right) A histogram of the normalised errors compared to a Gaussian distribution containing the same number of samples.

The spline function is initialised by a series of knots, the number and positions of which can be arbitrarily chosen to minimise residual errors. On the left of Figure 2.22 the normalised residuals between the local LSF and accumulated data for star A are shown when the knots, numbering 16, have: a) been regularly spaced; and, b) placed to increase the density in the image core over the length of the window. The effect of customising the knot locations is to generate a noise distribution that can effectively be described as Poissonian and produces a more accurate representation of the LSF. The second plot in Figure 2.22 shows this noise distribution for both knot schemes. The noise in the fit is normalised by the square root of the predicted flux, i.e. the Poisson error in the limit of large counts. The normalised errors approach a Gaussian distribution with zero mean and unit variance, which is also plotted here. The Gaussian has been scaled to contain the same number of samples⁸ as shown on the left in Figure 2.22. As the customised knot sequence concentrates on increasing the density of knots in the image core, the number of outliers outside the distribution is seen to reduce.

Investigating the asymmetry in the GRC

Before Astrium provided feedback on our first analysis of RC1 there was concern that the asymmetry in the GRC could be caused by CTI effects. This would mean the non-irradiated part of the CCD was in fact either irradiated or damaged by some other means so that the reference data could not be used to calibrate the image parameters attributed to CTI. To confirm this, a definitive signature in the data was searched for that could be attributed to CTI effects.

⁸Here, reference is being made to the number of electrons *per sample*.

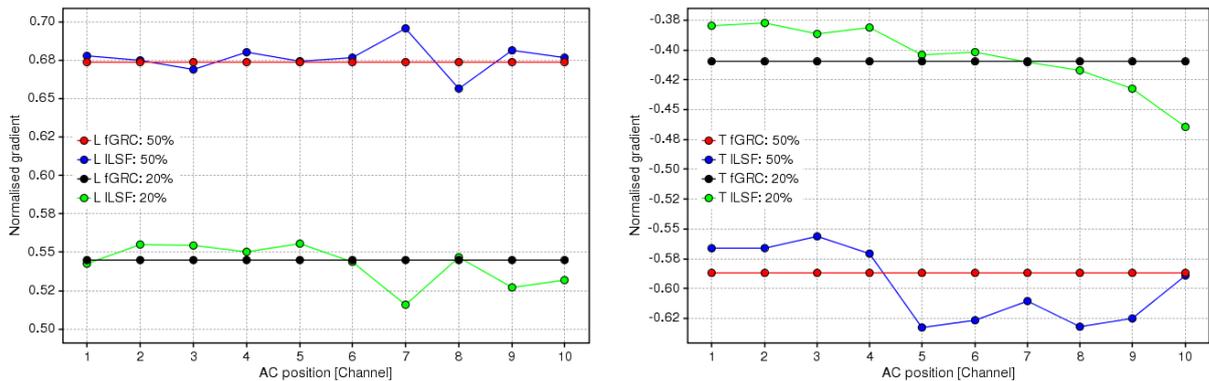


Figure 2.23: If the asymmetry of the GRC were due to CTI effects, it was thought that the (left) leading and (right) trailing edges of the local LSF compared to the fitted GRC would have reflected this. However, any expected signatures were not seen in the data.

Assuming star J is likely to be more damaged than star A, as it is closer to the irradiated zone, the gradients on the leading and trailing edges of the local LSF were expected to be steeper and shallower, respectively, in star J compared to star A. This idea follows on from other related studies suggesting radiation-induced CTI effects tend to take away charge from the leading edge and deposit this charge into the trailing edge. The leading slope should, therefore, become steeper and the trailing slope flatter, as the level of damage increases. The plot on the left of Figure 2.23 shows the gradients of the leading edge at the 50% and 20% level of peak intensity for all stars in the leading column compared with the equivalent gradient for the fitted GRC. The local LSF has been normalised to both the height and width of the fitted GRC before extracting the gradient. The leading edge gradient is expected to increase relative to the GRC. However, this is not seen. The plot on the right in Figure 2.23 shows equivalent results for the trailing edge, which is expected to decrease relative to the GRC. However, this too is not seen. In fact, the opposite appears to be occurring as star J is approached. These results indicated that radiation-induced CTI effects were not a likely cause for the asymmetry seen in the GRC. As previously mentioned, the most likely cause was a misalignment in the optics.

2.2.4 The irradiated data

Investigating the non-irradiated part of the irradiated data

In order to calibrate the effects of CTI in the irradiated data, a section of the mask must pass over the non-irradiated part of the CCD. The data pertaining to these non-irradiated stars should be comparable to the data found in the reference data. To investigate whether this is the case, the irradiated data taken with a signal intensity of 60000 have been analysed.

Figure 2.24(a) shows the accumulated data for stars A to C, i.e. A, B1, B2 and C with no DOB fitted with the GRC and a local LSF. The residuals between the two fits are minimal. Figure 2.24(b) shows the same residuals with all three levels of DOB. The local LSF in each case has been scaled in height to match the peak of the fitted GRC to remove the non-uniformity in illumination. While the residual flux without a DOB is comparatively small and symmetrical about the centroid, the cases with a DOB are

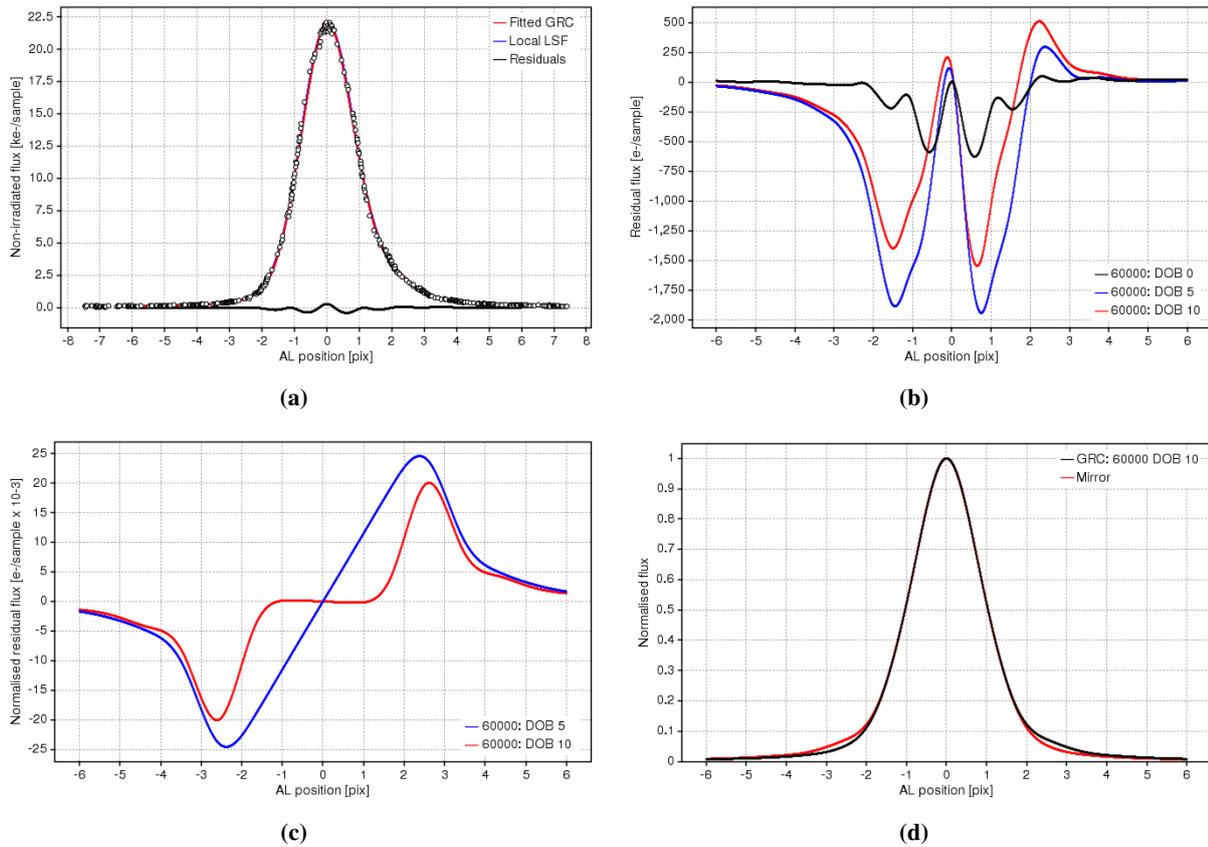


Figure 2.24: Finding a new GRC for irradiated data with a DOB (NI = non-irradiated). (a) Original GRC and local LSF fitted to NI data at 60000 with no DOB. (b) Residuals between the GRC and a scaled local LSF fitted to NI data at 60000 and all three levels of DOB. (c) Normalised residuals between a local LSF fitted to the NI data and its mirror image at 60000 and DOBs of 5 and 10. (d) Normalised fit to NI data at 60000 and a DOB of 10 with its mirror image.

far from producing an accurate fit to the data. On closer inspection it appears that the misalignment in the optics had been removed after tests without a DOB had been conducted but before tests began with a DOB. In these latter tests, the PSF is more symmetric and, therefore, the original GRC was no longer valid (note that the original GRC was still valid for tests without a DOB). To rectify this issue a new GRC had to be found that could be used to calibrate all the irradiated tests that included a DOB. Naturally, the brightest image, having the greatest signal-to-noise ratio (SNR), is better at constraining a curve fit. So a GRC was constructed out of the accumulated non-irradiated data in the irradiated tests at a signal intensity of 60000 and a DOB of 5 and 10, and the residuals between this curve and its mirror image were compared to decide which curve to use. Figure 2.24(c) shows the residuals between the GRC and its mirror image when a DOB of 5 (blue) and 10 (red) are used. The residuals are smaller when a DOB of 10 is used, therefore this GRC must have the highest degree of symmetry. The last plot, Figure 2.24(d), shows the normalised GRC and its mirror image for a signal intensity of 60000 and a DOB of 10. This is the GRC that was used to calibrate all the irradiated data that included a DOB. The original GRC, constructed from the reference data, was used to calibrate all the irradiated data that did not include a DOB.

Scan 1: a self-injection

Once the GRC had been defined for both tests, with and without a DOB, a complete analysis of the irradiated dataset could be carried out. As described in section 2.1.6, in order to extract relative centroid positions and fluxes one scan had to be chosen as the reference. It soon became apparent that if scan 1 was chosen to be this reference, or, indeed, if scan 1 was simply included, the resulting image parameters followed a different trend. Figure 2.25(a) shows the relative centroid positions for every star in the leading column when scan 1 is chosen to be the reference. The signal intensity is 60000 and there is no DOB. For comparison, two other cases are shown where the reference is scan 2 and 10, and scan 1 is not included when accumulating the data. In both of these cases the centroid positions are considerably less damaged than those measured when the reference scan is 1. The centroid positions also appear not to change significantly between scan 2 and scan 10. Figure 2.25(b) shows the same results when a DOB is applied. Incorporating scan 1 generates the same additional damage. Figure 2.25(c) illustrates how the peak flux varies across the leading column of stars when only a single scan is used to fit the GRC. Here, as with the centroid positions, scan 1 is preferentially damaged compared to scans 2 to 10. Figure 2.25(d) shows comparable cases when a DOB is applied.

In the absence of any preceding scans, shielding the leading images from the full range of vacant traps

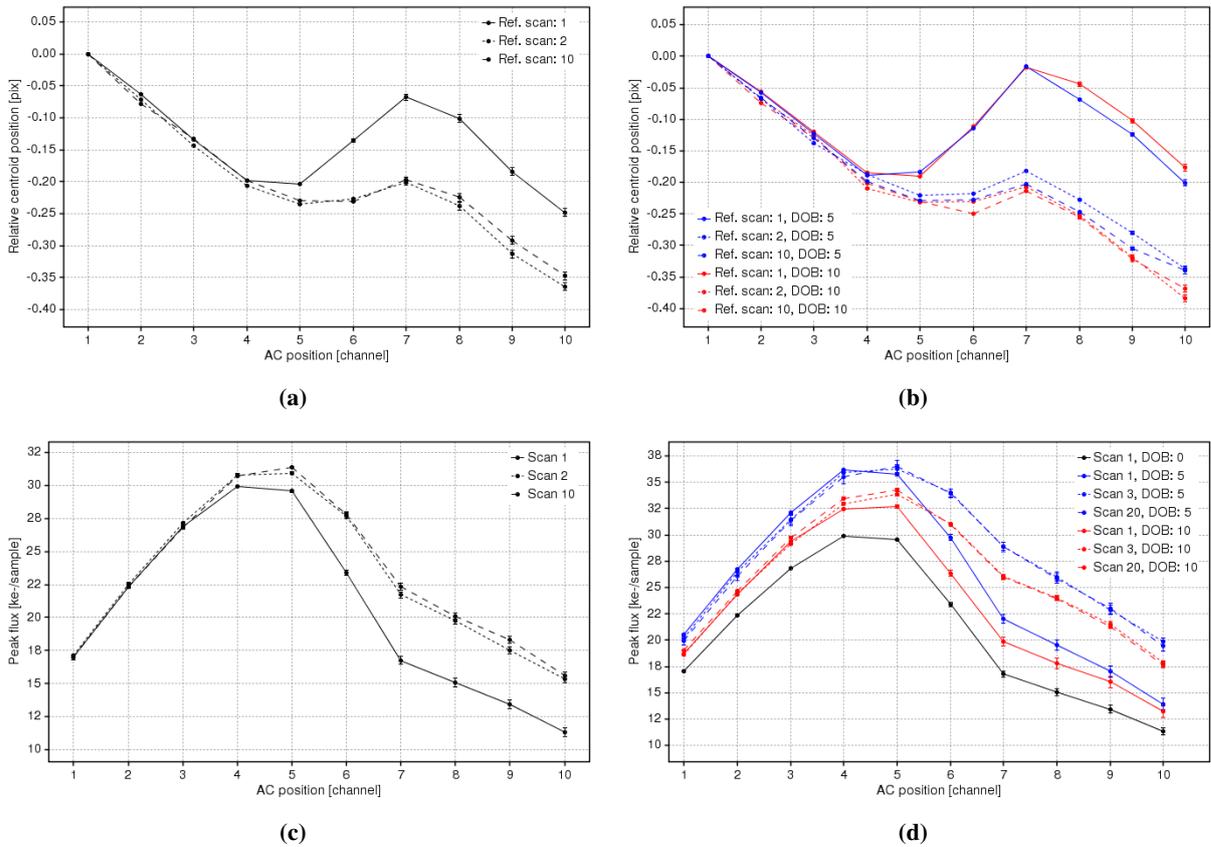


Figure 2.25: Illustrating how scan 1 is preferentially damaged compared to all other scans performed thereafter. These plots are for 60000. (a) Relative centroid positions without a DOB. (b) The same as in (a) with a DOB. (c) The peak flux without a DOB. (d) The same as in (c) with a DOB.

in the first scan is not possible. Traps that have release time-constants longer than the time between scans remain filled after scan 1, such that scan 2 and all scans that follow never experience the trap conditions initially in play on the CCD. Scan 1 therefore acts as a self-injection and limits the damage seen in subsequent scans to traps that have a chance of releasing charge before the next scan begins.

2.2.5 Relative centroid bias and charge loss in the leading column

All of the results presented hereafter come from the reanalysis of RC1. By subtracting the background locally and recalibrating the rotation of the mask with respect to the CCD it became possible to analyse the full range of magnitudes available and make comparisons with Astrium's results.

Integrated fluxes

Plots in Figure 2.26, labelled (a) to (f), show how the integrated fluxes of all the leading stars, at signal intensities of 60000, 7000, 2000, 650, 400 and 200 respectively, vary as a function of channel number and DOB level.

Firstly, the tests with a DOB of 5 register a higher integrated flux than with a DOB of 10. The only exception to this could be the tests performed at a signal intensity of 200. From these plots it would make more sense if the results for a DOB of 5 and 10 were reversed. However, this cannot be the case since the intensity of the DOB in each test has already been verified. Additionally, the fact that in some cases the integrated flux is measured to be the highest when no DOB is applied, tends to suggest it is the repeated calibration, and hence variability, of the light source that is the cause of these discrepancies.

Relative centroid bias

Plots in Figure 2.27 show how the centroid bias in each of the leading stars varies as a function of channel number, signal intensity, and DOB level relative to an undamaged reference. The centroid bias is given in units of AL pixels as well as angular size (mas) on the sky. The conversion between pixel and angular size is taken from the GPD: pixel angular area of a CCD in the AL direction = 58.933 mas.

Without applying a DOB, images with $G = 15.3$ (60000) passing through the fully irradiated area of the CCD suffer a bias on the centroid of ~ 0.22 pixels or equivalently ~ 13 mas on the sky. This is more than one order of magnitude above the accuracy levels required by Gaia on the basis of one transit to reach microarcsecond astrometry at the end of the mission⁹. Even when a DOB of 5 and 10 are applied at this magnitude, the centroid bias is only partially corrected.

This does not seem to be the case for $G = 17.7$ (7000). The action of applying a DOB does appear to rectify the bias, by a factor of just under 2 with a DOB of 5 and by a factor of 3 with a DOB of 10. Clearly, there is a case here for applying a strong DOB¹⁰. However, for even fainter images of $G = 19$ (2000), applying a DOB seems to have no effect on removing CTI. For the faintest images ($G = 20.24$

⁹Assuming that each object is observed 75 times in 9 CCDs (= 675 CCD transits), then the residual centroiding error per CCD transit at $G \approx 15$ is 0.011 pixels (Brown (2009))

¹⁰This is on the assumption that Gaia would have an artificial means of creating a DOB.

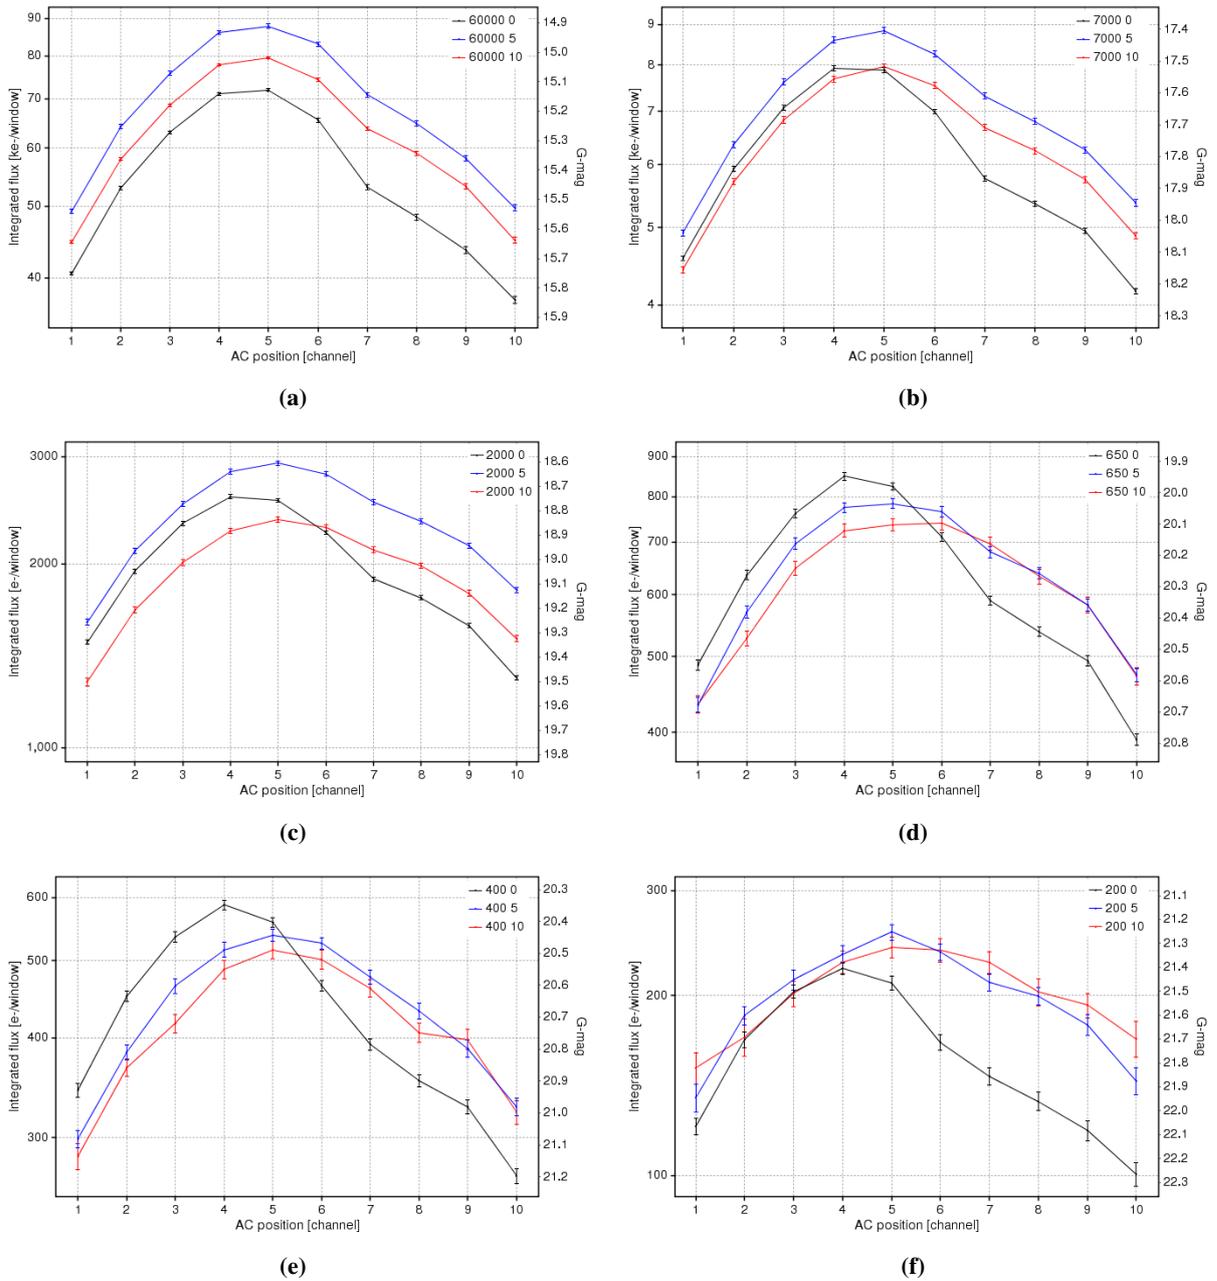


Figure 2.26: Integrated flux as a function of AC position at signal intensities of (a) 60000 (b) 7000 (c) 2000 (d) 650 (e) 400 (f) 200, and at all three levels of DOB.

(650), $G = 20.77$ (400) and $G = 21.52$ (200)) the picture is still not completely clear. In the last two cases, applying a DOB of 10 appears to shift the bias in the direction the image is moving. If correct, this is an effect caused by adding too much charge to the leading edge of the image. Overall, there is a general trend that as the DOB is increased the radiation damage measured as a bias on the centroid, in particular for the very faint images, is observed to decrease.

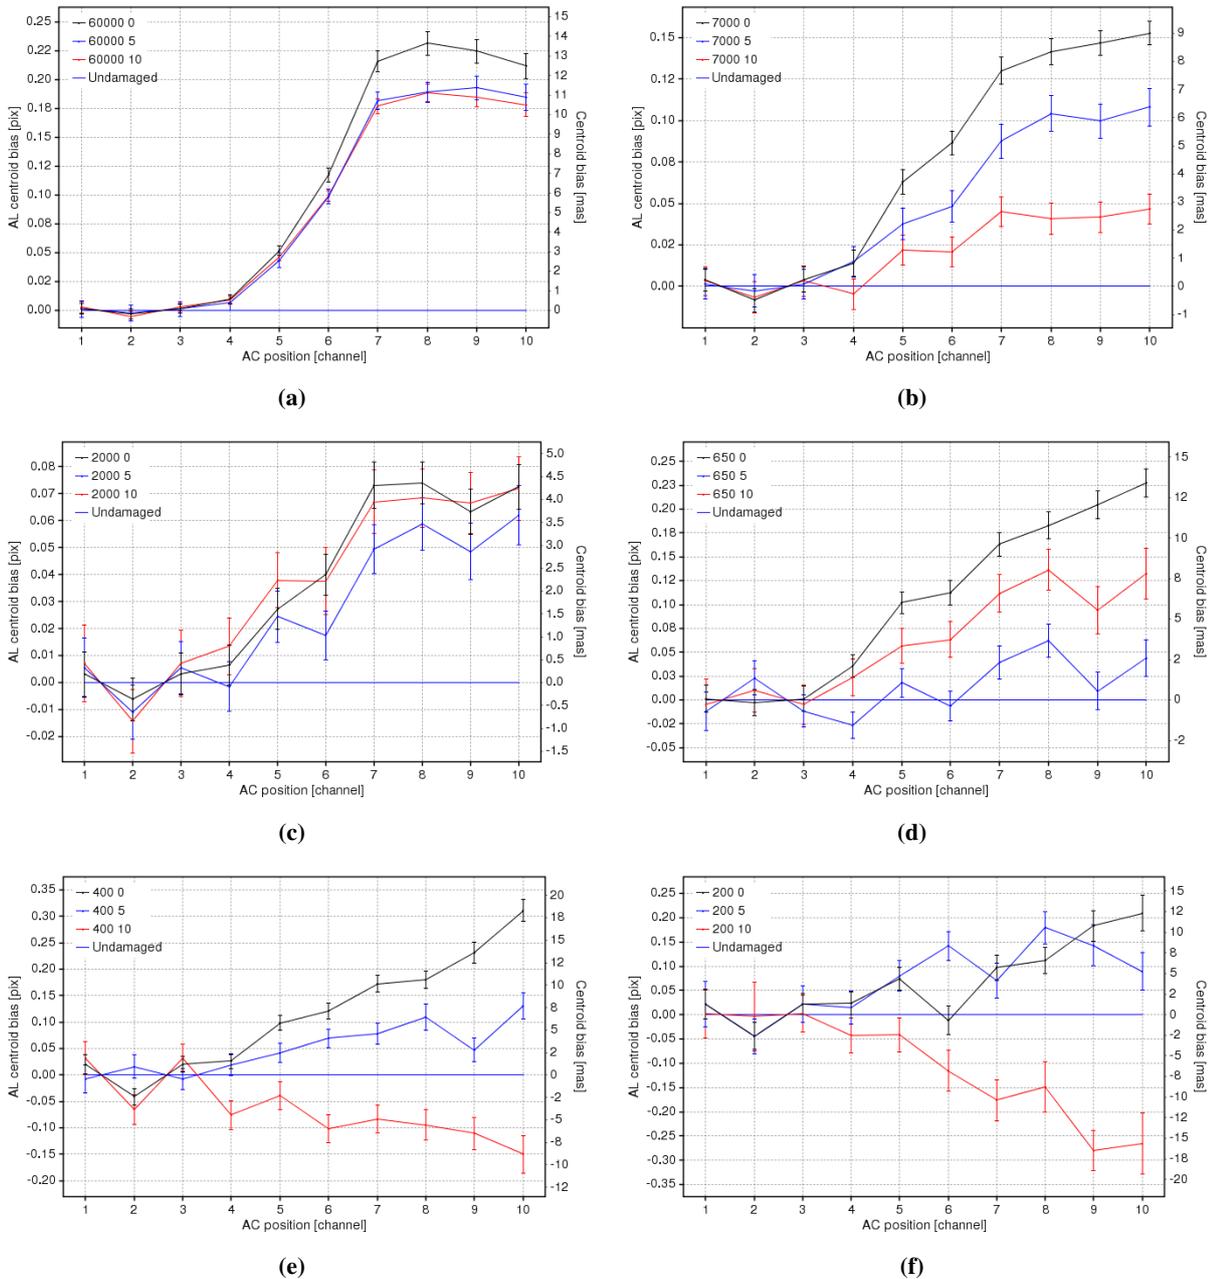


Figure 2.27: Centroid bias for each leading star on the mask relative to an undamaged centroid at signal intensities of (a) 60000 (b) 7000 (c) 2000 (d) 650 (e) 400 (f) 200 and at all three levels of DOB.

Relative charge loss

The relative centroid bias discussed above is intimately tied in with the charge loss suffered by a stellar image. Figure 2.28 shows how, for each signal intensity, the relative charge loss measured in each leading star varies as a function of DOB level.

As with the centroid bias, only a small amount of relative charge, after a DOB has been applied, is recovered for stellar images at $G = 15.3$ (60000). Increasing the level of the DOB from 5 to 10 has negligible effect. This is not the case for the fainter images. Adding a DOB has a cumulative effect on the

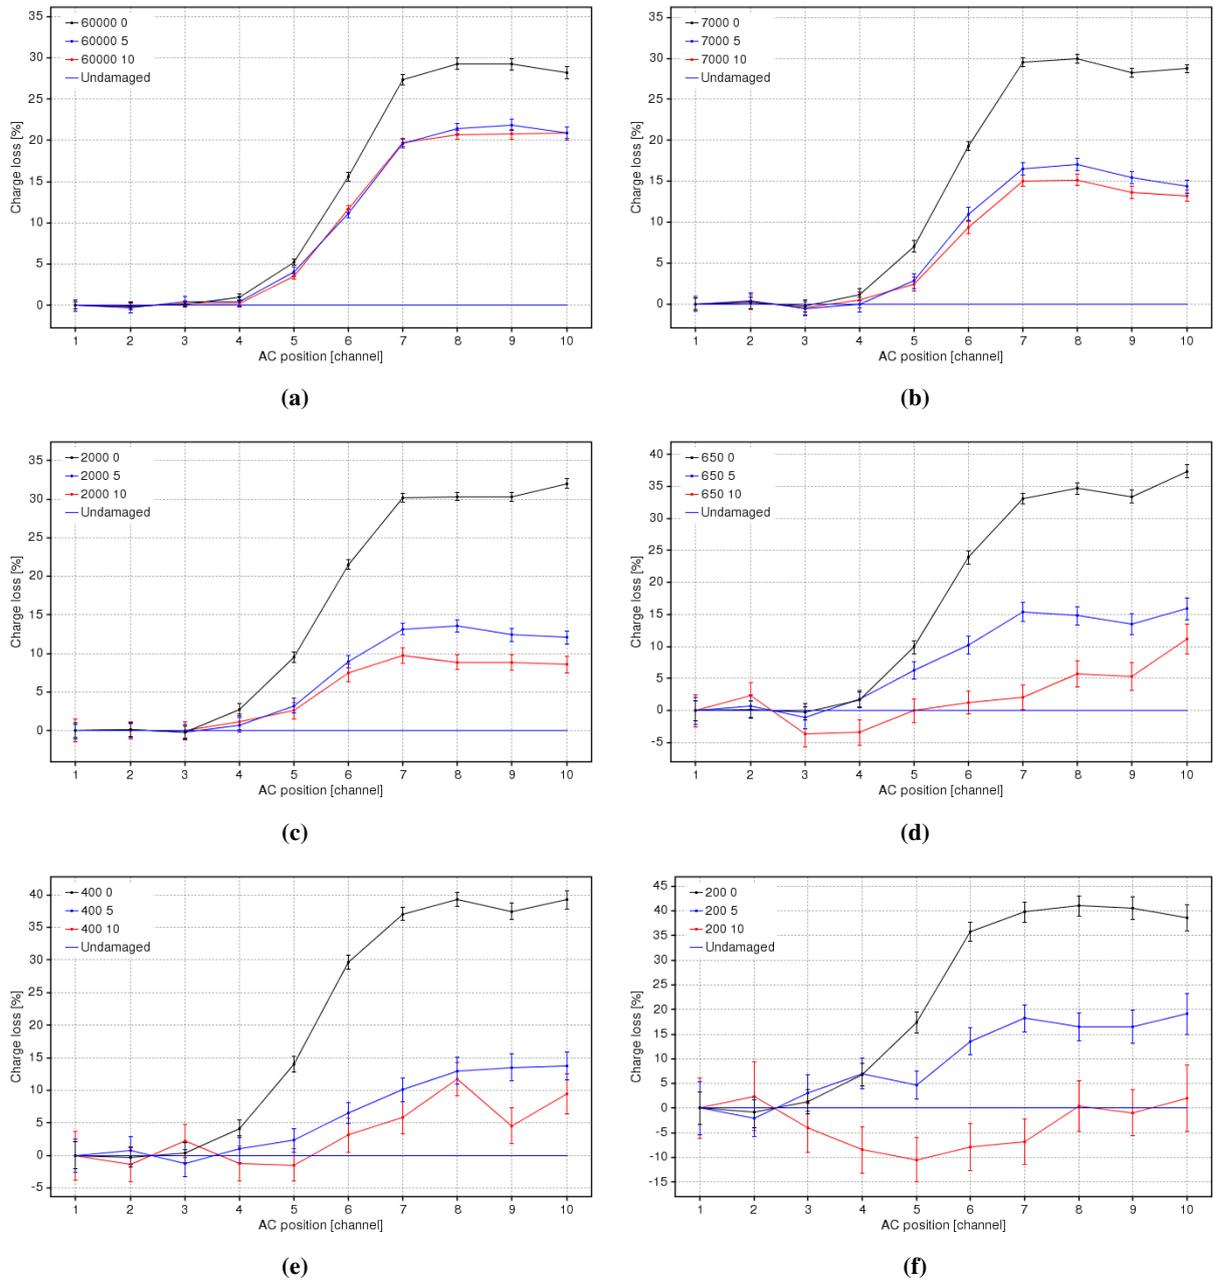


Figure 2.28: Relative charge loss suffered by each leading star on the mask at signal intensities of (a) 60000 (b) 7000 (c) 2000 (d) 650 (e) 400 (f) 200 and at all three levels of DOB.

charge loss as G increases ($G > 15$). Radiation damage is reduced not only as the image grows fainter, but also as the DOB grows stronger.

Radiation-induced CTI effects as a function of magnitude and DOB level

The relative centroid biases shown in Figure 2.29(a), (c) and (e) represent the average biases of the last three stars (H, I, J) passing over the irradiated part of the CCD with a DOB of 0, 5 and 10 respectively.

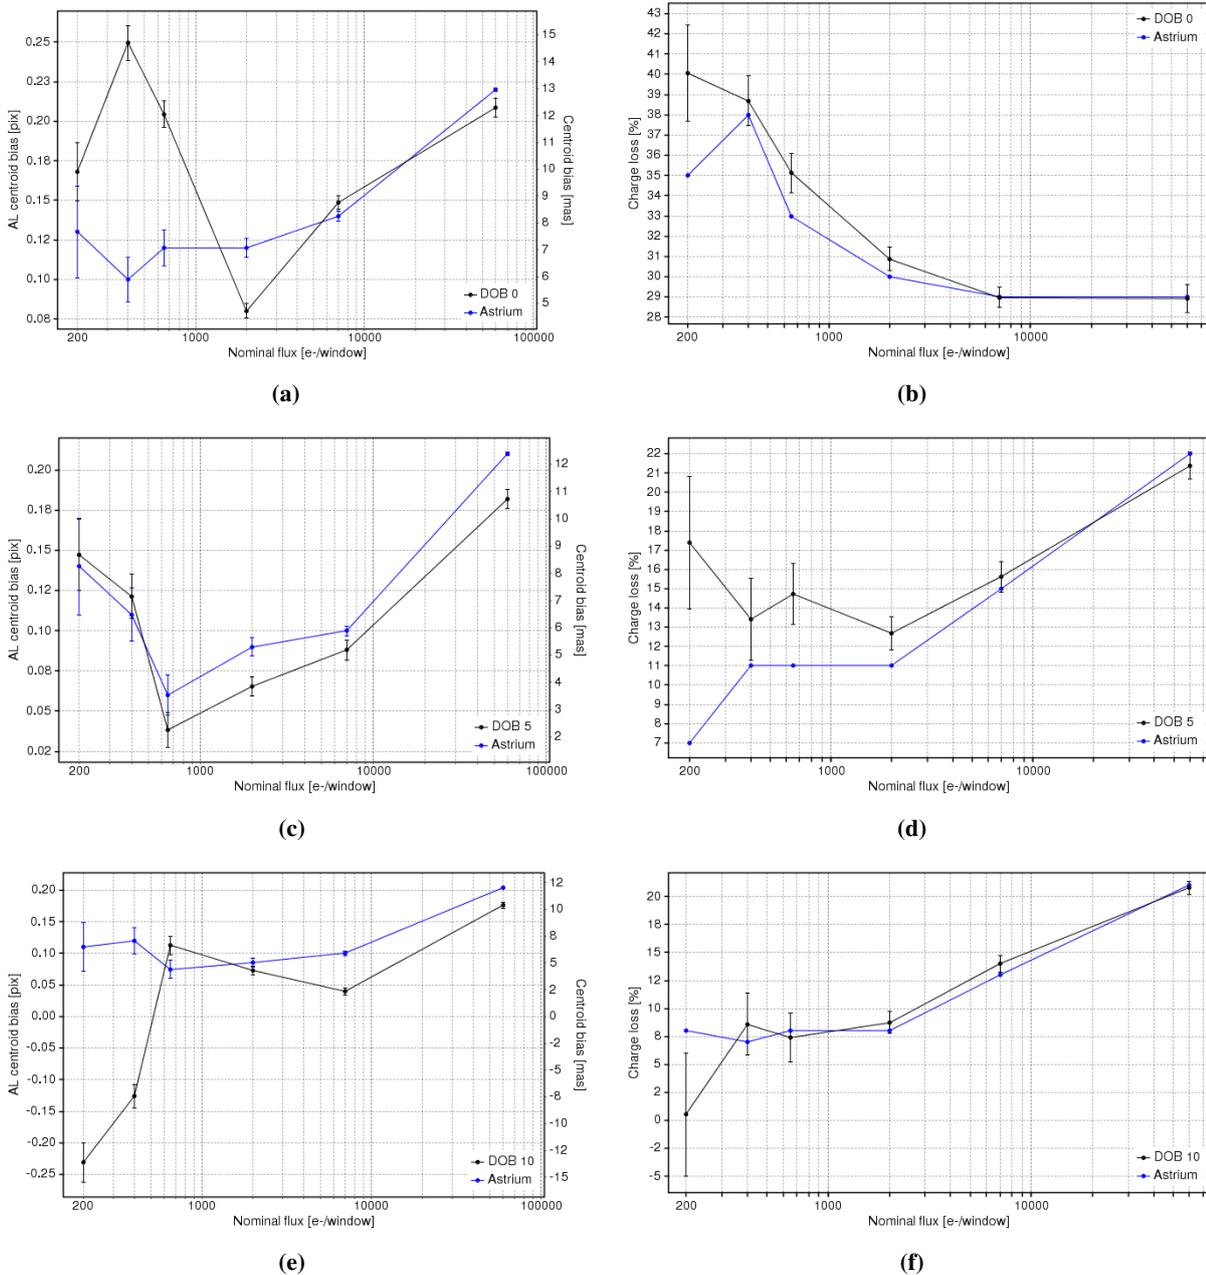


Figure 2.29: Comparing Astrium's results: average relative centroid bias and charge loss as a function of signal intensity level and DOB level.

Figure 2.29(b), (d) and (f) show how the relative charge loss varies as a function of image intensity and DOB. The image intensity is related to the magnitude according to Table 2.1. The average centroid biases of the last three stars have been taken in order to compare our results to those found by Astrium, which are also shown in each plot.

Without a DOB, there are some clear differences between the results for the fainter images. While Astrium's results show a steady reduction in the bias with signal level, our results suggest there is a more complicated behaviour. At some low signal level the bias stops decreasing and increases sharply before

dropping again at the faintest level. For the charge loss Astrium measured a decrease at the faintest end. However, here it was found that the charge loss continues to increase. This may not be significant, given the extent of the error at this magnitude, though no errors are given on the charge losses recorded by Astrium to compare.

With a DOB of 5, both the centroid bias and charge loss are reduced as the signal intensity decreases. This is until the signal intensity drops below $\sim 1000 e^-$ ($G \approx 19.8$), where the radiation damage begins to increase. Only the charge loss appears to disagree with Astrium's results for the faintest signals. Finally, with a DOB of 10, the charge loss appears to gradually improve as the signal intensity decreases. Though this not as clear cut in the centroid bias.

In general, the two datasets tend to diverge as the fainter signals are probed. Astrium used a Gaussian function to fit the images, which may partially explain this. The SNR at the fainter end is naturally worse and the final best fit to the image as well as the errors on final CTI parameters will depend on how the data was accumulated. How noisy the images are will be touched upon in the next section.

Figure 2.30 and Figure 2.31 overlay the results from Figure 2.29 and compare directly how the different levels of DOB affect the centroid bias and charge loss measurements respectively as a function of image intensity. For signal intensities ≥ 2000 ($G \lesssim 19$) the trends observed in Figure 2.30 suggest that the brighter the image, the greater the centroid bias. Below this intensity it is not clear how the centroid bias is being set. Though, in general, applying a DOB does reduce the effects of CTI on the faintest images.

In contrast, Figure 2.31 clearly shows reasonably predictable trends for charge loss as a function of signal intensity and DOB level. Applying a DOB does have an effect on retaining the charge in the

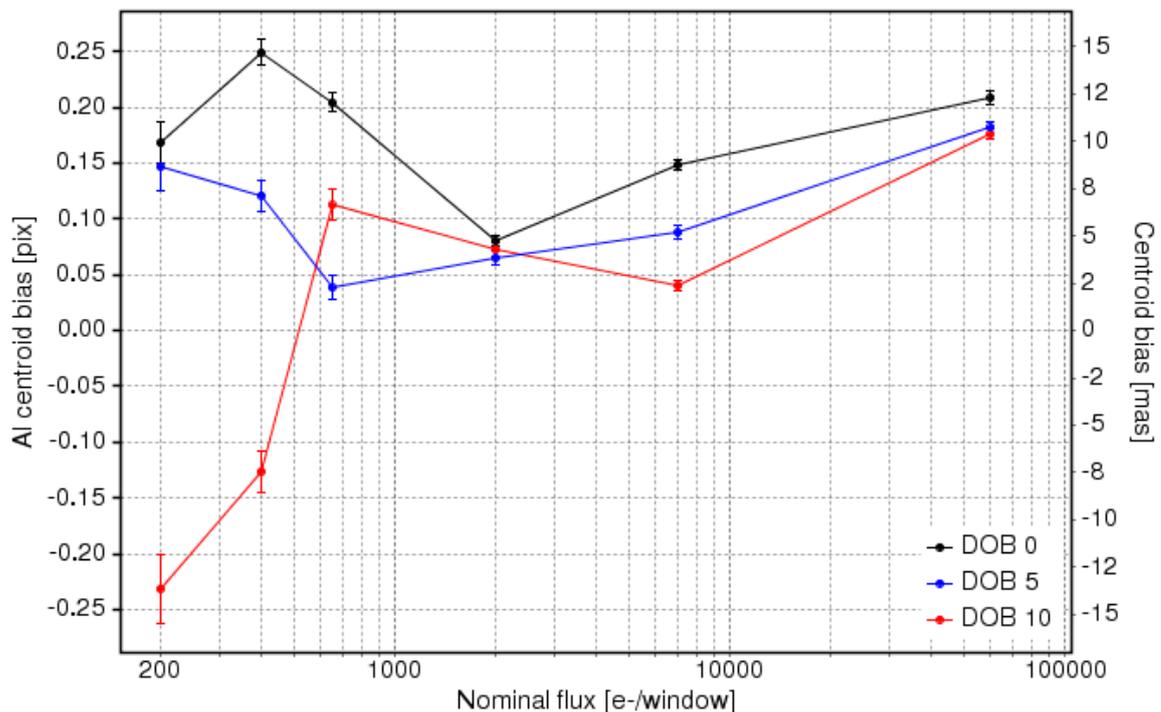


Figure 2.30: Average relative centroid bias as a function of signal intensity and DOB level

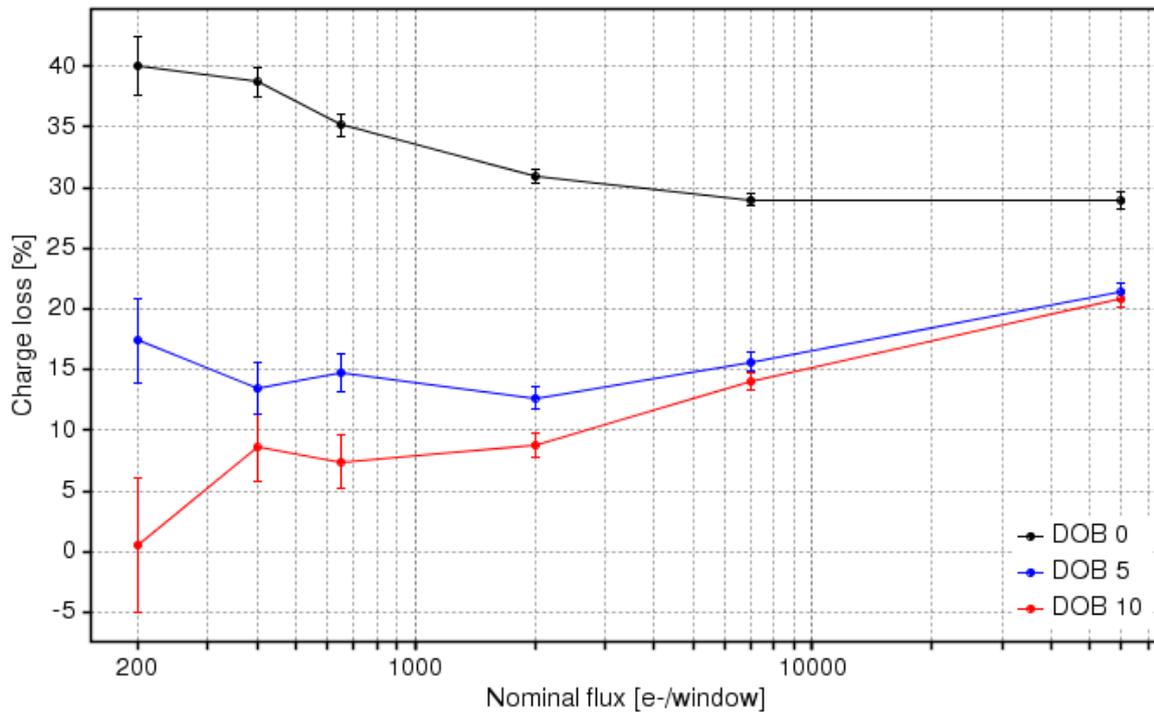


Figure 2.31: Average relative charge loss as a function of signal intensity and DOB level

image, even more so for the fainter images. Doubling the strength of the DOB from 5 to 10 does reduce the amount of radiation damage observed even further, but not as much as when the DOB was increased from 0 to 5. This suggests there is a limit to how far a DOB can go in removing radiation damage.

Appendix B.1 collates the average relative centroid bias and charge loss measurements as a function of signal and DOB level.

Image profiles: undamaged vs damaged

The calibrations performed earlier, leading to values for the centroid bias and charge loss for each star, can be used to determine how the shape of the image profile is affected by radiation-induced CTI and the application of a DOB. In Figure 2.32(a), a normalised fit to both the undamaged and damaged images of star J at a signal intensity of 60000 ($G = 15.3$) and no DOB is shown, together with the image data and charge lost between the two fits. Star J is the last image to be read out of the CCD and passes entirely over the irradiated zone. Since the image is moving to the left in this plot, it is the leading edge that is damaged the most compared to the trailing edge, which appears also to be slightly damaged. Figure 2.32(b) shows the effect of applying the two levels of DOB. Whereas the trailing edge up to a certain point is fully recovered, the leading edge remains depleted close to the level seen without a DOB.

Figure 2.32(c) and (d) show comparative results with a signal intensity of 650 ($G = 20.2$). Without a DOB, the damage is similar to that seen with a bright image, i.e. the leading edge is eaten away first while the trailing edge suffers less damage. With a DOB, around half of the charge lost is recovered - the stronger DOB having even more of an effect. The trailing edge is again recovered, while the leading

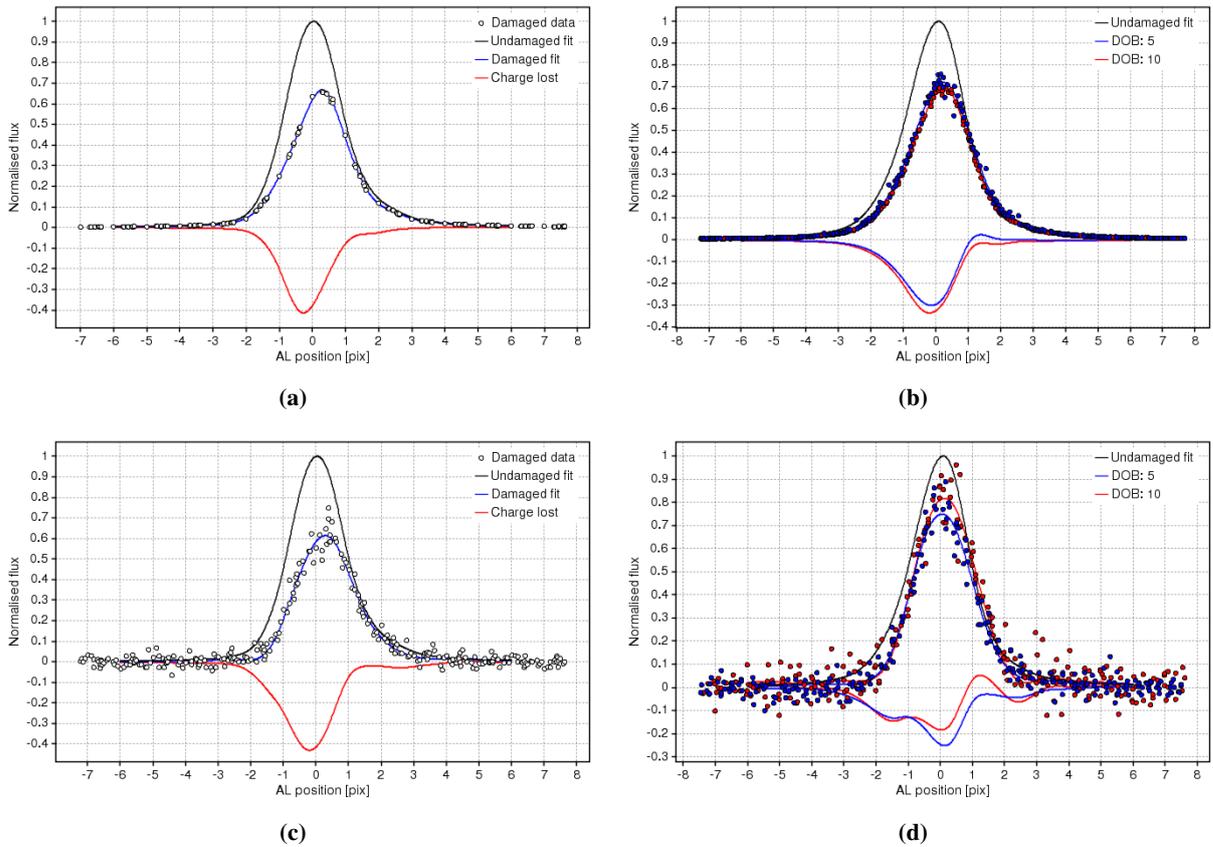


Figure 2.32: Normalised undamaged vs damaged image profiles for star J. (a) At 60000 ($G \approx 15$) without DOB. (b) At 60000 with a DOB of 5 and 10 applied. (c) At 650 ($G \approx 20$) without DOB. (d) At 650 with a DOB of 5 and 10 applied.

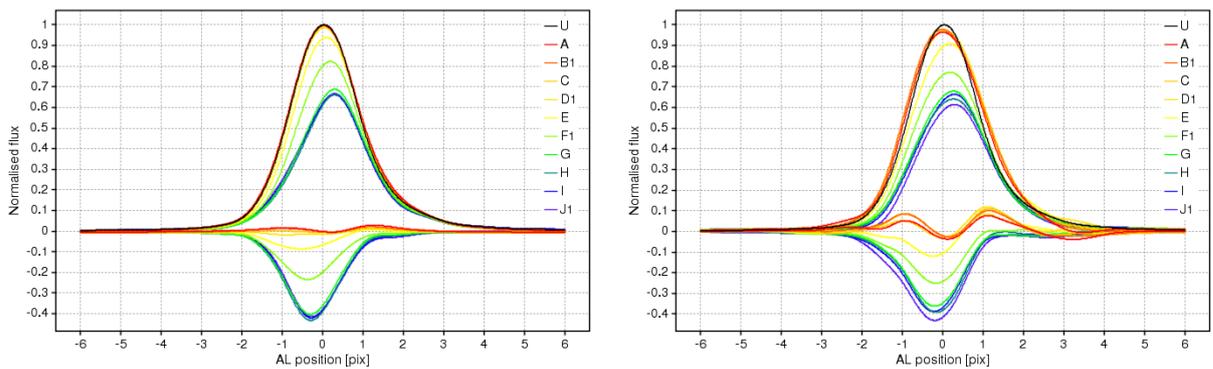


Figure 2.33: CTI effects as a function of radiation level. (left) Normalised images of all leading stars compared to an undamaged reference at 60000 ($G \approx 15$) without DOB. (right) Comparable results to those seen on the left at 650 ($G \approx 20$). U = undamaged reference curve.

edge continues to take the brunt of the damage.

The stars in the leading column pass over three different areas of the CCD: a non-irradiated zone, a transition zone and a fully irradiated zone. In the transition zone the level of radiation is expected to increase gradually. In Figure 2.33 the damaged profile fitted to each star in the leading column is

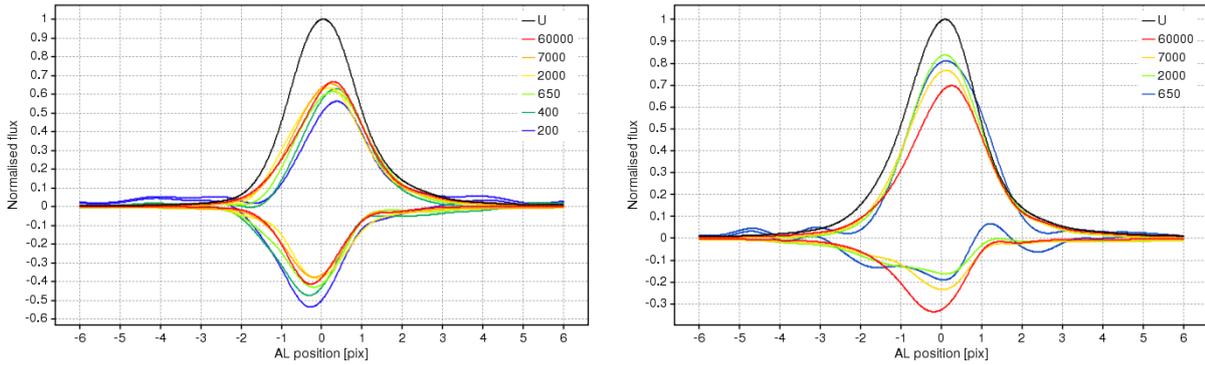


Figure 2.34: CTI effects as a function of signal intensity and DOB level for star J. (left) Normalised images for all intensity levels without a DOB. (right) Normalised images for all intensity levels with a DOB of 10. U = undamaged reference curve.

plotted against the undamaged profile from star J. For clarity, the undamaged images for all the stars are not shown. It is clear how an increase in radiation level progressively damages an image. Charge is removed, primarily, from the peak and leading edge. The width of the damaged image in the non-irradiated channels is wider than the undamaged reference because of the variation in the illumination conditions on the mask.

Lastly, Figure 2.34 attempts to show how CTI effects vary as a function of signal intensity and DOB level. As discussed above, the application of a DOB (of 10), as shown on the right, does little to recover the undamaged profile of bright images, but appears to increase the chances of shielding CTI effects for fainter images. The damaged profiles for signal intensities of 400 and 200 are not shown with a DOB of 10 as a reliable fitting could not be found.

2.2.6 Conclusions

RC1 was carried out to investigate and characterise the effects radiation-induced CTI has on a small number of stellar images, varying in signal intensity and subject to three different levels of DOB. Before image parameters were extracted the background had to be subtracted. As this was found to vary both in the AL and AC directions a local background extraction method was used. The assumption was made that the background could be sampled at the edges of a 15-pixel window. This may not however be the case for images brighter than $G \approx 15$.

After calibrating for the geometry of the test setup and the non-uniform illumination over the CCD, values for the centroid bias and charge loss as a function of magnitude and DOB level could be extracted. The results suggest that:

1. In general, the centroid bias decreases as the signal intensity decreases. Though, at certain faint signal intensities (around 400: $G \sim 20.8$), for a DOB < 10, the centroid bias may actually increase.
2. In general, applying a DOB decreases the centroid bias. However, as the DOB is increased there is no clear way of predicting how this trend will vary as a function of signal intensity.

3. With a DOB of 0, the charge loss gradually increases as the signal intensity decreases.
 - At 60000 ($G = 15.33$) the relative charge loss is $28.91 \pm 0.70\%$.
 - At 200 ($G = 21.52$) the relative charge loss is $40.05 \pm 2.37\%$.
4. With a DOB of 5, the charge loss steadily decreases before slightly increasing as the signal intensity decreases.
 - At 60000 ($G = 15.33$) the relative charge loss is $21.39 \pm 0.69\%$.
 - At 200 ($G = 21.52$) the relative charge loss is $17.39 \pm 3.45\%$.
5. With a DOB of 10, the charge loss gradually decreases as the signal intensity decreases.
 - At 60000 ($G = 15.33$) the relative charge loss is $20.78 \pm 0.60\%$.
 - At 200 ($G = 21.52$) the relative charge loss is $0.52 \pm 5.51\%$.
6. As the DOB level increases so does the noise in the background. This will degrade the SNR and increase the error estimates in the fitting procedures.
7. In terms of relative improvement, increasing the DOB from 0 to 5 has a more pronounced effect than increasing the DOB from 5 to 10.
8. Though applying a DOB mitigates radiation damage at all signal intensities, it is the fainter images that benefit the most.
9. At a CCD operating temperature of 163K, CTI manifests itself by eating away the leading edge and peak region of a stellar image. If the damage is left to continue, or the signal flux is low enough, the trailing edge also begins to erode.
10. By applying a DOB, the trailing edge is either partially or fully recovered, depending on whether the image is bright or faint respectively.

So what does this suggest about the dependency between a DOB and radiation damage? Increasing the DOB reduces the overall charge loss, but the damage continues to vary as a function of signal level. From this it appears that the DOB cannot remove entirely the local transit history and there is some form of radiation damage which is insensitive to a DOB. In terms of the centroid bias it is even less clear what the effects of a DOB are both as a function of intensity and signal level.

Physically, it has been suggested that the results shown here imply the overall reduction and limited improvement seen when doubling the DOB intensity in the charge loss is down to keeping only the slow trap species full. Faster traps, i.e. those with shorter release time-constants, tend to redistribute charge, which has more of an effect on the centroid bias and are not overly susceptible to a constant low-level background. Taking this argument further, the fact that a very small number of electrons are able to interact at all with (slow) traps spread throughout an entire pixel tends to suggest the electrons are not confined to a specific volume, and the extent of trapping is density- rather than volume-driven. This has implications for predicting the shape of damaged images (see Chapter 4). Given a lack of alternatives, I would say this is the best explanation we have at the moment in terms of trap occupancy. The acid test is

whether the charge distortion model (i.e. the model being developed by DPAC to predict the distortion of an image by radiation damage) can reproduce these results. To date, only the data from RC2 has been used, which did not investigate the DOB.

With regards to the shape of the damaged image, it has been clarified how the shape degrades as a function of radiation level, magnitude and DOB level, and an interesting result was found: while the leading edge is "eaten away", the trailing edge appears to remain fairly undamaged, at least for the brighter images. For the fainter images the damage is fairly uniform across the entire image with both leading and trailing edges losing charge. There is no clear evidence for any significant deposit of charge behind the image. This may be the result of operating the CCD at 163K, which was found by Sira to be an optimal temperature to suppress traps both in the image area and the serial register.

To summarise, at one time it may have been thought that an artificial DOB alone could have been all that was required to mitigate radiation damage. What these results show is that this is not the case; the effects of the local transit history cannot be removed by a DOB. That said, the DOB does have an effect upon the severity of the radiation damage and so must be taken into account in the radiation calibration model.

Recommendations made to industry and follow-up

After analysing RC1, DPAC wished to make recommendations to Astrium in order to improve the output from future campaigns. Firstly, it was clear from analysing scan 1 that it was not possible to separate memory effects from the effects of the DOB; the first scan could have masked some of the effects of the DOB. So the time delay between successive scans was asked to be increased from 30s to at least 90s to try and rectify this. Secondly, it had also been shown that the flux response over the mask was different for faint stars compared to bright stars. This required an extra step to calibrate the charge loss in the faint stars. To avoid misinterpreting the calibration of the faint images in the data, calibration data without radiation damage for faint stars was also requested.

These findings were presented to the Radiation Calibration Working Group for the first time on 20th September 2007 (ESA Gaia Project (GAIA-EST-MN-03815), Meeting #5) and then again in the form of Brown & van Leeuwen (SWB-001) on 23rd April 2008 (ESA Gaia Project (GAIA-EST-MN-05213), Meeting #6). Just before this I gave a presentation to the Radiation Task Force on 14th April 2008 on this analysis. Bearing in mind that RC2 had been already carried out between July and August 2007, there was no way these recommendations could have been fed back to Astrium in time. With RC3 however, which began in April 2008, it was possible to suggest updates to the test plan. In this context it is the 'astrometric regular' tests that are now referred to, which ran between July and August 2008. The first recommendation above was not followed, presumably because of the limited time available to perform such tests and the fact that charge injection was, after RC2, shown to effectively reset the transit history. The second request for reference data at the same magnitude as the irradiated data was provided for half of the dataset, namely for bright stars when gates were activated. Unfortunately, the reference data for the medium and faint star tests were all taken at $G = 15$. Since the bulk of RC3 was performed after the analysis of RC2, further feedback on this recommendation shall be discussed in section 2.3.7.

2.3 Radiation Campaign 2

In RC1 a small number of stellar images ($G > 15$) were illuminated onto an AF CCD prototype to investigate the level of image distortion incurred as a result of radiation damage, and how a DOB may reduce this distortion. As a passive form of CTI mitigation, it was found that a DOB could shield images from the effects of radiation damage, though the benefits were limited. To try and fully mitigate the effects of radiation damage another method was required.

RC2 deals with an active form of CTI mitigation: charge injection (CI). With this technique traps are filled, initially, by injecting a line of charge across the CCD, perpendicular to the scanning direction. This charge, injected at the beginning of the CCD, is then read across the CCD filling in as many traps as possible along the way. Any images following the CI should encounter very few, if any, traps thus removing the effects of CTI. The overall benefits of applying a CI depend upon the magnitude of the stellar image, and the properties of the CI; namely, the level and duration of the injection, and the delay between the injection and the image. In this analysis each of these properties were investigated in turn, characterising radiation damage as a function of magnitude with and without CI, and bringing into focus the effects of the injection tail. This tail is composed of electrons that are released back into the background from traps that have been filled by the charge injection. Alongside subtracting the electronic offset, it will be shown that the CI tail must also be dealt with before extracting image parameters. The technical note on this work is Brown & van Leeuwen (SWB-002).

2.3.1 Charge Injection

Charge injection is one way of mitigating CTI, by keeping traps full, that has already been used on CCDs in space (cf. WFC3 on HST (Giavalisco (2003)) and EPIC MOS CCDs on XMM-Newton (Smith et al. (2004)) as well as the soft X-ray detector on SUZAKU (Prigozhin et al. (2007))). After an injection has been made and read across the CCD, many of the traps that are filled are no longer available to distort the images that follow. This is, of course, until the traps release their charge and become active again. The optimum time between charge injections must therefore be found. Ensuring all traps are full while minimising dead-time (this is also increased by extending the duration of the injection) requires investigation, which forms the basis for the second and third radiation campaigns.

The following description of charge injection refers to Figure 2.35. A charge injection structure for the Gaia CCDs comprises a diode (connection labelled ID) and gate (connection labelled IG) electrode, which are provided at either one end or both ends of the CCD image area. If ID is pulsed from "high" to "low" voltage levels, a row of charge can be injected into the image area.

The quantity of charge N [e^-] depends on the voltage difference between the "high" level of the image area clock $I\emptyset 1$ and that on IG as well as whether $I\emptyset 1$ is "high" when $I\emptyset 2$ is "low" (shown on the left of Figure 2.35 and the first equation below) or both $I\emptyset 1$ and $I\emptyset 2$ are "high" (shown on the right of Figure 2.35 and the second equation below):

$$N = S(V_{I\emptyset 1} - V_{IG}) + N_o \quad \text{or} \quad N = S(V_{I\emptyset 1+I\emptyset 2} - V_{IG}) + N_o \quad (2.10)$$

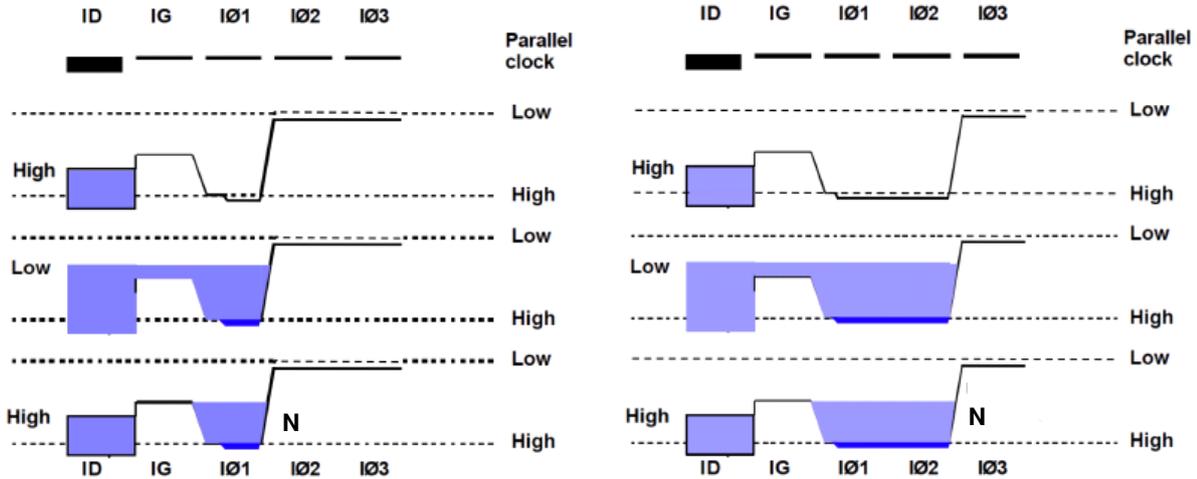


Figure 2.35: Charge injection structure. Charge is injected from the diode electrode (ID) via the gate electrode (IG) into (left) the image area clock IØ1 when IØ1 is "high" and IØ2 is "low" or (right) both IØ1 and IØ2 when they are "high". The Supplementary Buried Channel (see chapter 3) is shown in dark blue. Credit: Seabroke et al. (GMS-002).

where S is the responsivity in electrons per volt and N_o is a fixed offset. Once charge has been injected into IØ1 (or both IØ1 and IØ2), the charge can be clocked over to IØ2 (or IØ3). This is done by lowering the voltage in IØ1 (or both IØ1 and IØ2) and raising the voltage in IØ2 (or IØ3). Gaia pixels in fact have 4 electrodes to clock charge over, so an additional IØ4 could be included in Figure 2.35. In the pixels that follow only the 4-phase clock electrodes are present to clock charge through the image area. In addition to the duration and delay of the charge injection being investigated in RC2 the injection level is also considered. The range of voltages and number of electrons N in this test are shown in Table 2.8.

2.3.2 Hardware and test configuration

The test bench as well as the CCD used in RC2 were identical to those used in RC1. See section 2.2.1 for references to specific documents detailing the test setup. The important changes in RC2 were the use of charge injection and a larger optical mask. Figure 2.36 shows a schematic of this mask and the test configuration adopted during the irradiated tests. The size of the pinholes on the mask have been over exaggerated. Note that the injection was driven from only one end of the CCD (i.e. the lower end), causing the AC profile of the charge injection to decrease over the CCD. In effect, each channel (see Figure 2.36) received a different charge injection.

The effective mask area used to represent the stellar field was a rectangular $22 \text{ AC} \times 50 \text{ AL}$ grid. In total, 1100 stars traversed the CCD in each scan. Each hole on the mask is separated by ~ 20 pixels AC and ~ 50 pixels AL. Figure 3-3 in EADS Astrium (GAIA.ASF.TCN.PLM.00197) test report on RC2 shows a detailed schematic of the illumination pattern. In this document it is also mentioned that in column 16 of the mask the 14th hole is blocked, probably by a dust grain.

Table 2.7 shows the sampling scheme and the position of the mask with respect to the CCD in both the reference and irradiated test cases. In total, 32 samples were read out from the CCD after each TDI

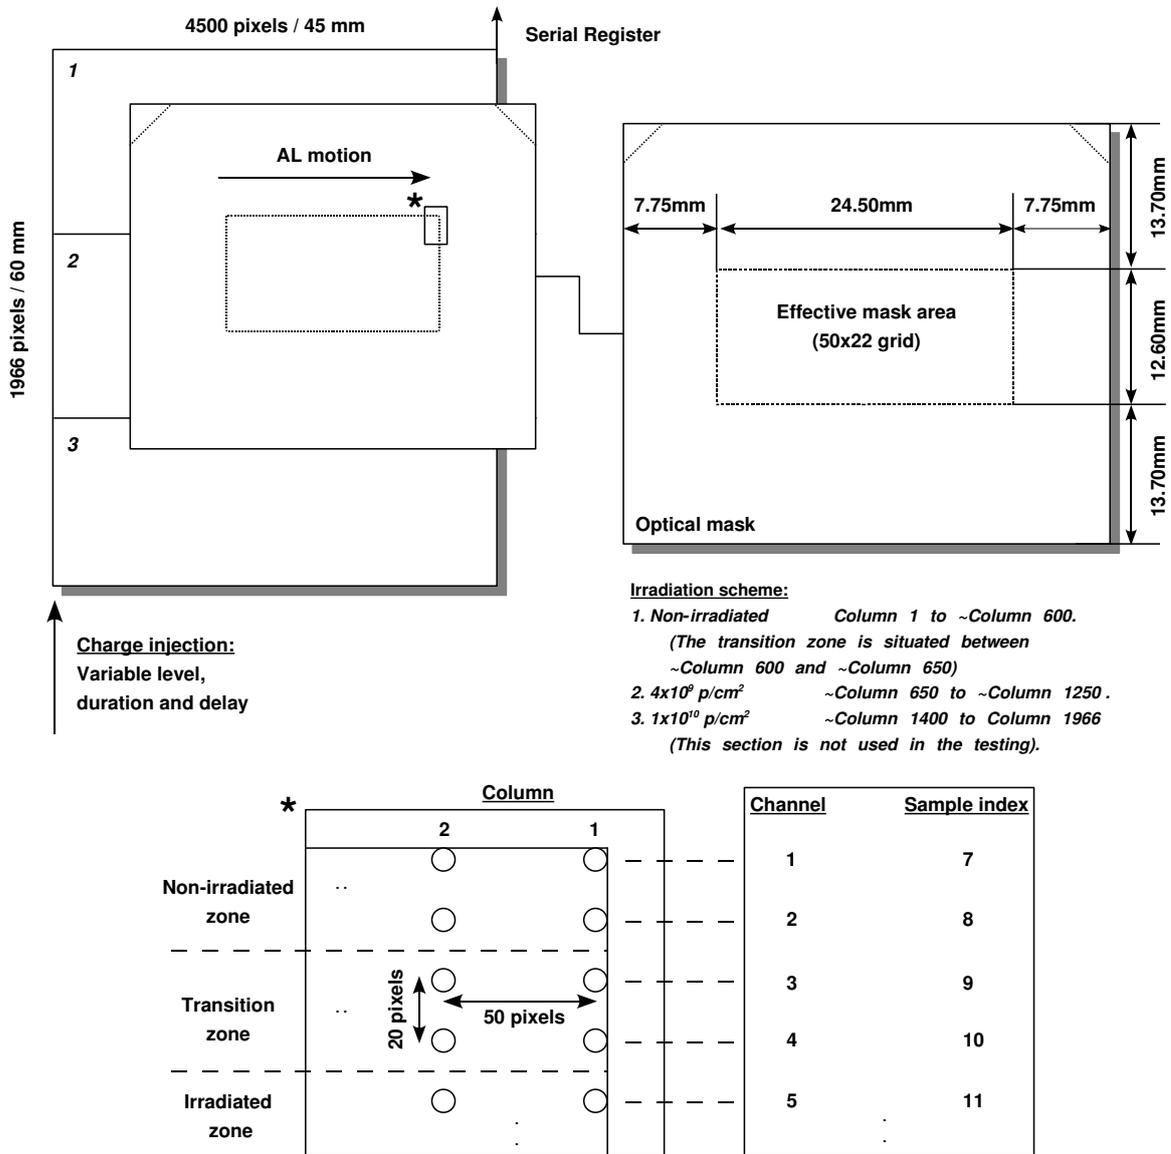


Figure 2.36: Details of the test setup in RC2. This particular configuration shows the mask passing over the CCD during an irradiated test, where only channels 1 and 2 pass over the non-irradiated zone.

period. However, only 22 samples (7 to 28) contain data pertaining to the stellar field. To differentiate these samples from the rest of the readout, they have been labelled channels 1 to 22 respectively. Samples 1 and 2, also known as prescan samples, contain information on the unexposed background (electronic offset + RON), while samples 3 to 6 and samples 29 to 32 may be used to monitor the CI. Samples 7 to 28 (and 2) have been formed by binning 12 pixels AC with the binning approximately centred on the PSF, while all the remaining samples were left unbinned. Associated with the sample index is an AC pixel position closest to the PSF centre. Since the AC pixel position for the last two channels in the reference data do not overlap with the first two channels in the irradiated data, two completely different areas of the CCD were scanned over when the reference and irradiated tests were performed. Note that in both the reference and irradiated data there is a slight drift away from an AC pitch of 20 pixels between

Channel	Sample index	AC position [pix]		CCD area	
		Reference	Irradiated	Reference	Irradiated
Prescan	1	-	-	-	-
1	7	110	564	NI	NI
2	8	130	584	NI	NI
3	9	150	604	NI	PI
4	10	171	625	NI	PI
5	11	191	645	NI	PI
6	12	211	665	NI	I
7	13	231	685	NI	I
8	14	252	706	NI	I
9	15	272	726	NI	I
10	16	292	746	NI	I
11	17	313	767	NI	I
12	18	333	787	NI	I
13	19	353	807	NI	I
14	20	374	828	NI	I
15	21	394	848	NI	I
16	22	414	868	NI	I
17	23	434	888	NI	I
18	24	455	909	NI	I
19	25	475	929	NI	I
20	26	495	949	NI	I
21	27	516	970	NI	I
22	28	536	990	NI	I

Table 2.7: Sampling scheme and associated CCD readout area in RC2. The channel or sample index corresponds to the relative AC position of each charge packet on readout, i.e. channel 1 is readout first, then channel 2, etc. The PSF-centred AC pixel coordinate is determined by the binning procedure. In this table each PSF is binned AC by 12 pixels. NI = non-irradiated, PI = partially irradiated, I = irradiated.

each image centre. Note also that only channels 1 and 2 pass over the non-irradiated part of the CCD in the irradiated tests to calibrate the irradiated data.

The investigations into CI on an AF CCD were divided into three separate tests (referred to as 'astrometric tests' in EADS Astrium (GAIA.ASF.TCN.PLM.00197)). In the first test the duration of the CI was ~ 20 pixels, whilst the delay was set to ~ 25 pixels. Once an optimal level for the CI had been found, the second test adopted this level, and the delay was set to ~ 30 pixels. Likewise, once an optimal duration had been found, the third test adopted this duration as well as the CI level acquired in the first test (see Figure 2.37).

In the first test, the CI intensity was given in terms of an applied voltage as well as the number of electrons injected per pixel. In the analysis that follows the voltage is used. The injection levels shown in Table 2.8 have been calculated by averaging over the injection levels quoted for samples 7 to 28 in the last table in EADS Astrium (GAIA.ASF.TCN.PLM.00212). Note that as the voltage increases, the level (e^-/pixel) decreases. This is due to the way in which an injection is generated. Section 2.3.1 explains in more detail why this is so. Primarily, it is caused by the voltage difference between the diode and gate: the higher the diode voltage, the lower the intensity of the charge injection.

In RC2 the electronic gain was set to $3.56 e^-/\text{LSB}$. Since the maximum count attainable out of the PEM is 65535 LSB, it follows that the maximum number of electrons that can be recorded per

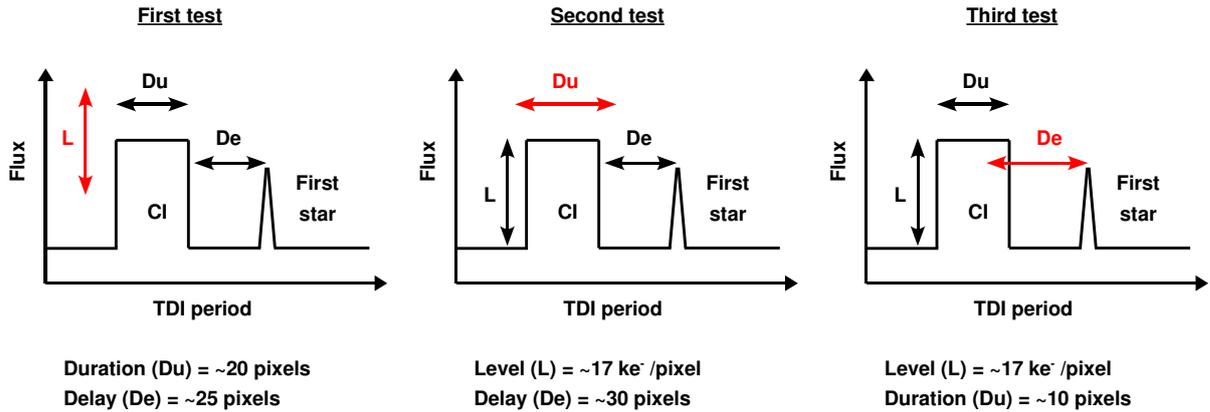


Figure 2.37: The three tests performed in RC2 to investigate the effects of CI on CTI mitigation. These diagrams are based on Figures 2-1, 2-3 and 2-5 in EADS Astrium (GAIA.ASF.TCN.PLM.00212).

sample is ~ 233304 or ~ 19442 e⁻/pixel. Without the electronic offset (nominally set to ~ 1300 LSB), the maximum injection level attainable before saturation sets in becomes ~ 19100 e⁻/pixel. In RC2, when the voltage was ≤ 7.51 V, the injection level was high enough (≥ 42489 e⁻/pixel) to saturate the read out in *every channel* after binning. Below this level, since the injection was made from only one end of the CCD, those channels furthest away from the injecting source were the first to remain unsaturated.

Tables 2.8 and 2.9 summarise the test conditions in the first and second tests respectively. It was Astrium's initial intent to systematically inject ahead of all stars between $G = 11.5$ and 15.0 . This is why only two magnitudes: 13.3 and 15.0, were studied in the first and second tests. In the third test, summarised in Table 2.10, three additional magnitudes (16.75, 18.5 and 20) were included, to consider the effects of CI on the fainter stars. In terms of the number of scans, only 5 were performed in the first two tests to collect irradiated data, whilst in the third test the number of scans varied between 5, 10 and 15, depending on the magnitude. In every test the reference data comprised 20 scans. Since a number of scans were performed on different days, certain injection delays are labelled 1, 2 or 3. This label identifies which set of reference data were used to calibrate the irradiated data.

The reference data was in fact not always the same magnitude as the irradiated data. In the second test the $G = 13.3$ irradiated data had to be calibrated using the $G = 15$ reference data, and in the third test

First test: INJECTION LEVEL.

$G = 13.3$ and 15.0

V ¹	6.51	6.96	7.31	7.51	7.84	8.06	8.26	8.42	8.57	8.65	8.72
<e ⁻ /pix>	118251	79444	54594	42489	26376	16466	9898	6049	3506	2551	1870
T (13.3)	161.2	161.2	161.2	163	163	163	163	163	163	163	163
T (15.0)	161.2	161.2	161.2	161.2	161.2	161.2	161.2	161.2	161.2	-	-

5 scans performed for each injection level.

Tests without injection: T1 (13.3) = 163, T2 (13.3) = 161.2 (performed 1 week later), T (15) = 161.2, 5 scans.

Reference data: T (13.3) = 163, T (15) = 161.2, 20 scans performed.

T (mag) indicates the CCD temperature in K and the magnitude of the illumination source.

¹The test at 7.51V and gate 10 activated was not analysed. Gates are investigated in RC3.

Table 2.8: Information on the first test.

Second test: INJECTION DURATION.

G = 13.3 and 15.0

Pixel duration: 2, 4, 6, 8, 10 and 20.

5 scans performed for each injection duration.

T (13.3) = 161.4 ± 0.2 , T (15) = $161.9 \pm 0.4/-0.6$.

Tests without injection: T (13.3) = 161.4 ± 0.2 , T (15) = $161.9 \pm 0.4/-0.6$, 5 scans performed.

Reference data: T (15) = 164.2 ± 0.7 , 20 scans performed. No reference data was taken at G = 13.3.

T (mag) indicates the CCD temperature in K and the magnitude of the illumination source.

Table 2.9: Information on the second test.

Inj. delay [pix]	G									
	13.30		15.00		16.75		18.50		20.00	
	Temp [K]	S	Temp [K]	S	Temp [K]	S	Temp [K]	S	Temp [K]	S
Ref.	164.38 ± 0.62	20	166.60 ± 0.90	20	163.29 ± 0.64	20	165.63 ± 0.83	20	166.9 ± 0.85	20
Ref.	-	-	170.90 ± 1.08	20	166.28 ± 0.85	20	168.38 ± 0.97	20	Same as 18.5	20
Ref.	-	-	169.44 ± 1.03	20	-	-	-	-	-	-
10	162.17	5	161.80 ± 0.70	10	161.40 ± 0.48	10	162.45 ± 1.68	15	163.05 ± 1.83	15
30	162.17	5	161.80 ± 0.70	10	161.40 ± 0.48	10	162.45 ± 1.68	15	163.05 ± 1.83	15
60	162.17	5	161.80 ± 0.70	10	161.40 ± 0.48	10	162.45 ± 1.68	15	163.05 ± 1.83	15
100	162.18 ± 0.01	5	161.80 ± 0.70	10	161.40 ± 0.48	10	162.45 ± 1.68	15	163.05 ± 1.83	15
150	162.22 ± 0.02	5	161.80 ± 0.70	10	161.40 ± 0.48	10	162.45 ± 1.68	15	163.05 ± 1.83	15
300	162.30 ± 0.04	5	161.80 ± 0.70	10	161.40 ± 0.48	10	162.45 ± 1.68	15	163.05 ± 1.83	15
1k	162.45 ± 0.07	5	161.80 ± 0.70	10	161.40 ± 0.48	10	162.45 ± 1.68	15	163.05 ± 1.83	15
1k ¹	-	-	161.64 ± 0.08	10	161.33 ± 0.09	10	160.93 ± 0.14	15	162.45 ± 1.68	15
5k	162.66 ± 0.10	5	161.60 ± 0.05^1	10	-	-	-	-	-	-
30k	163.03 ± 0.23	5	162.02 ± 0.35^1	10	161.29 ± 0.08^1	10	161.15 ± 0.37^1	15	163.67 ± 1.15^1	15
60k	-	-	161.85 ± 0.60^2	10	161.36 ± 0.11^1	10	-	-	-	-
120k	-	-	164.41 ± 1.93^2	10	162.96 ± 1.46^1	10	-	-	-	-
No inj	162.07 ± 0.10	15	161.80 ± 0.70	5	161.40 ± 0.48	5	162.45 ± 1.68	5	163.05 ± 1.83	5
No inj ¹	-	-	~ 162	15	161.51 ± 0.05	15	161.20 ± 0.12	15	-	-
No inj ²	-	-	161.52 ± 0.12	15	-	-	-	-	-	-

¹Tests which are calibrated using the second reference dataset.

²Tests which are calibrated using the third reference dataset.

Table 2.10: Information on the third test. S indicates the number of scans performed.

all of the reference data were taken at G = 15 regardless of the magnitude of the irradiated data. Since it is known from RC1 that the illumination conditions change across the mask, corrections to the flux calibration would have to be made for these particular test cases. Fortunately, in the first test reference data were provided at G = 13.3 and 15.

To study the effects of altering the CCD operating temperature Astrium repeated the above tests at four different temperatures: 163K, 175K, 185K and 195K. Since they showed conclusively that the nominal temperature of 163K was best suited to minimising CTI effects in the image area (IA), this analysis concentrated on data collected at 163K.

2.3.3 Background analysis

As in RC1, to begin an analysis of the background, underlying features are accentuated by taking a running mean (see eqn. 2.3) of the readout in the first three channels of the reference data taken from the first test. This is shown in Figure 2.38.

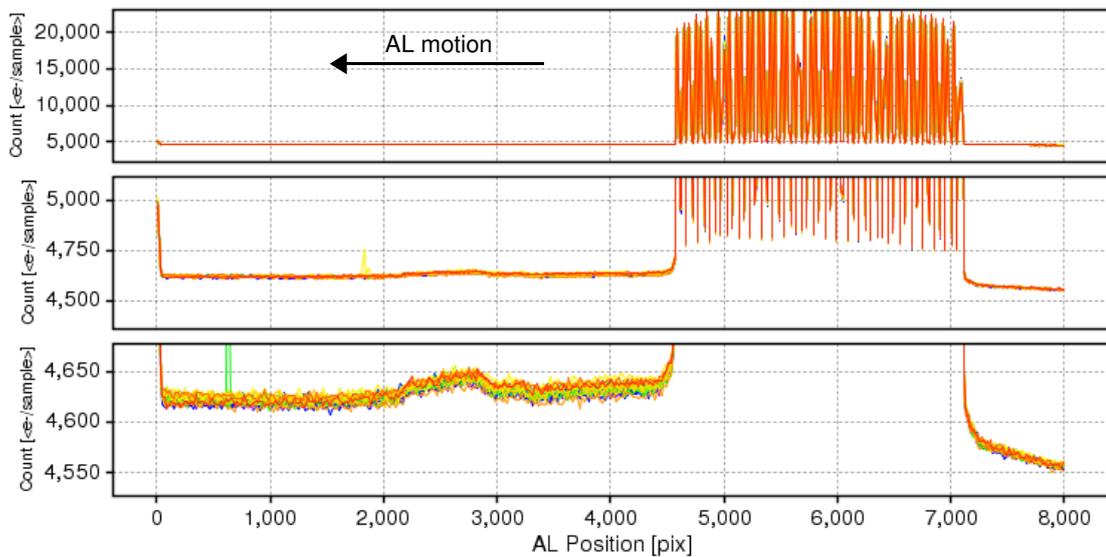


Figure 2.38: Running means of all 20 scans in channels 1 to 3 of the reference data taken from the first test at $G = 13.3$. In each channel (moving down the figure) successive close-ups are made of the data.

The top plot corresponds to channel 1, the second, channel 2, etc. $G = 13.3$ and all 20 scans have been overlaid. In each channel successive close-ups of the data are shown to identify features that become visible at different scales. In all 22 channels many of the features shown here can also be seen. The following observations can be made:

- In channel 1 the electronic offset can be seen as well as the bulk of the stellar field. Over half the scan records the background before any image appears. The level of illumination across the stellar field in the AL direction is non-uniform, peaking towards the centre. Every scan extends for 8000 TDI sample periods.
- Zooming in closer on channel 2 a few more features can be seen. Firstly, every scan begins with a sharp decrease in the count level, falling off rapidly before continuing at a steady rate. After ~ 1900 pixels have been read out, a "cosmic" appears in the background.
- Lastly, zooming in even closer on channel 3, towards the beginning of one scan another cosmic appears followed by a series of "bumps". The largest bump is read out around 2700 pixels after the scan has begun. Afterwards the background continues to increase up to the first stellar image. Once the last image appears the background falls quickly below the level recorded in the first half of the scan.

As described in section 2.1.2 the region chosen to estimate the background lies between 500 and 1500 pixels after readout. This is the region that appears to be the most stable over the length of each scan. Regarding the features described above, Astrium explains that:

- The initial sharp decrease in the count level is attributed to processes occurring off the CCD in the PEM and does not interfere with the transfer of charge across the IA.

- The two bumps appearing in the background before the stellar field has passed are caused by stray light leaking around screws that are holding the mask in place. As shall be shown this leak is also seen on the other side of the mask. As the brightness of the source increases so too does the intensity of the bumps.
- The extraneous signals or cosemics are seen in numerous scans. They are attributed to light randomly striking the CCD from unknown sources.

Figure 2.39 is comparable to Figure 2.38 except in this case the reference data has been taken from the first part of the third test. This is when the nominal G magnitude was 13.3, but the reference data was actually taken at $G = 15$.

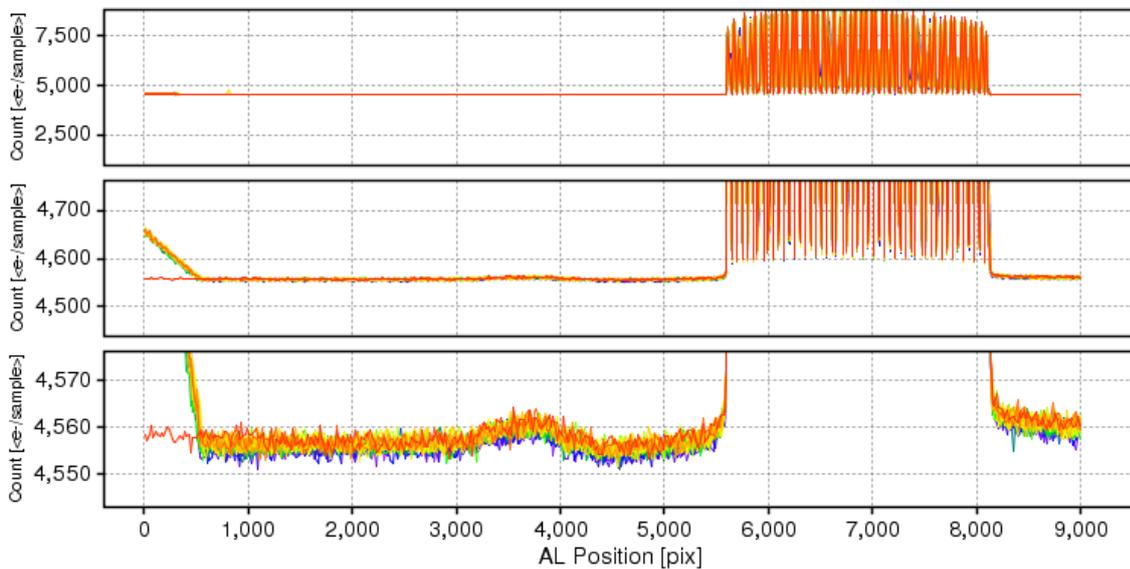


Figure 2.39: Running means of all 20 scans in channels 1 to 3 of the reference data taken from the first part of the third test.

In this dataset some new features appear:

- Compared to channel 1 in Figure 2.38, the intensity of the stellar field seen here is clearly less than before. This is because in the third test all reference data were taken at $G = 15$ regardless of the irradiated test magnitude.
- The length of each scan is 9000 TDI sample periods rather than 8000.
- By zooming in on channels 2 and 3 a new systematic feature at the beginning of each scan, except scan 1 (red line), can be seen. Scan 1 records the electronic offset from the outset. Scans 2 to 20 show a gradual decrease in the count level settling into the background only after ~ 500 pixels have been read out.
- A close-up of channel 3 reveals that the background after the stellar field returns approximately to the pre-field level.

Feedback from Astrium explains that:

- The initial decrease in the count level seen at the beginning of every scan, except the first, is due to the mask travelling back over the CCD after each scan. To stabilise the star intensity from one scan to the next the light source is kept on as the mask position is reset.

One other feature that appears in the data, not shown in the preceding figures, is a systematic drop in the underlying background between samples 21 and 22 (channels 15 and 16). Returning to the reference data in the first test at $G = 13.3$, the plot on the left of Figure 2.40 shows the mean background count as a function of the AC position. This is clearly the same systematic that appeared in the analysis of RC1 (cf. star J (see Figure 2.13 and Corcione et al. (CL-001))), and is again attributed to the PEM anomaly.

This plot also hints at how the background measured in the same channel may vary between scans. To shed light on this behaviour, the plot on the right shows how by sequentialising the AL position of every scan, the background varies within each scan. Since the maximum count attainable out of the PEM is 65535 (after which the clock would reset), it was necessary to sequentialise the AL position. In this plot the background estimate (B_{est}) in scan 1, taken between 500 and 1500 pixels, is placed on the x-axis at 1000 pixels (midway between 500 and 1500). The rest of the background estimates fall on the x-axis according to the sequentialised clock output. As the TDI period is $\sim 982.8\mu s$ the length of time taken up by performing 20 scans showed that the average time between one scan ending and the next scan beginning was $\sim 19.7s$. Since each scan lasted, in this case, for 8000 TDI periods, the average time between scans was $\sim 27.3s$, as reported by Astrium. Whether the background variations are due to the CCD-PEM coupling or a variable light source is unclear. However, it appears that when the first test was performed the test equipment had not yet reached a sufficient state of stability to reproduce the background level recorded in the previous scan.

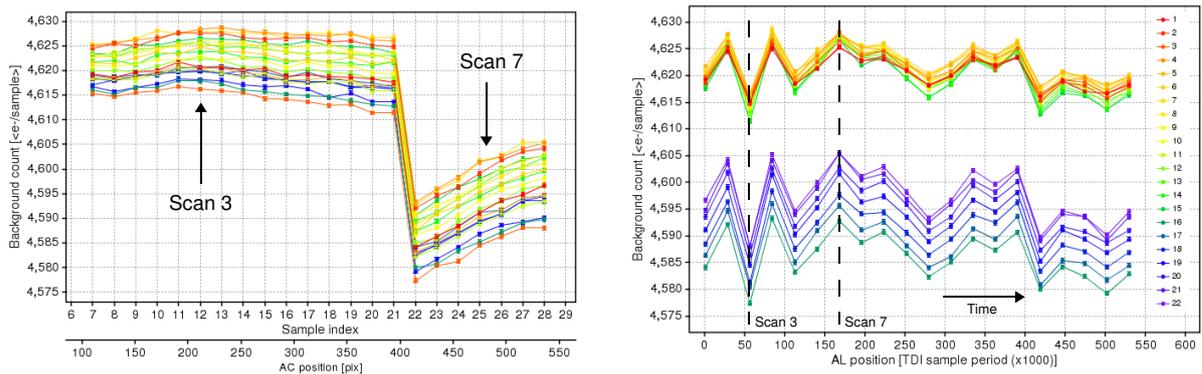


Figure 2.40: Background behaviour in the reference data from the first test at $G = 13.3$. (left) Mean background count as a function of AC position. (right) Mean background count measured over consecutive scans.

Comparable to Figure 2.40 are the two plots shown in Figure 2.41. This time the AC behaviour of the mean background count in the reference data taken from the first part of the third test is shown. As expected, the systematic drop caused by the PEM can be seen between channels 15 and 16. Most notable

is the reduced spread in the mean background count between scans. This suggests that the variability in the background over time is much less compared to the trend seen in the reference data in the first test. The plot on the right in Figure 2.41 verifies this.

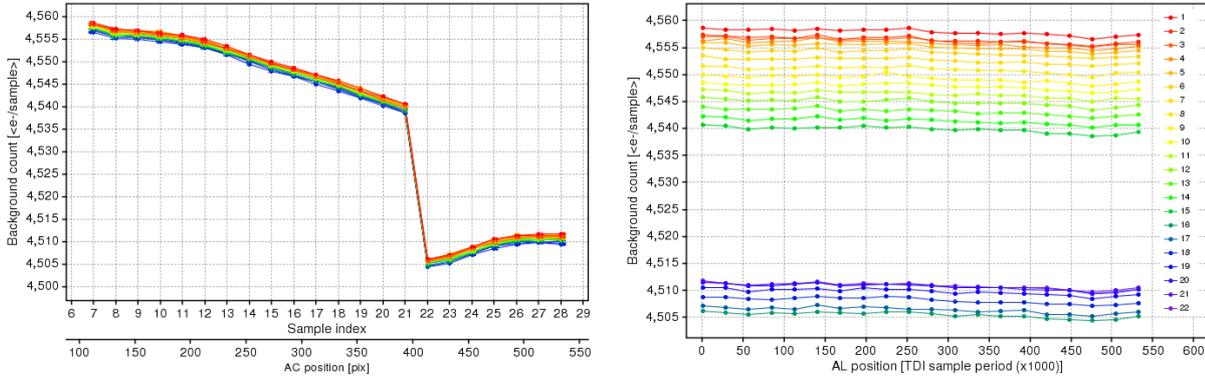


Figure 2.41: Background behaviour in the reference data from the first part of the third test. (left) Mean background count as a function of AC position. (right) Mean background count measured over consecutive scans.

Figure 2.42 shows histograms of the noise distributions observed in two sets of reference data. On the left the noise in sample 1 and the combined noise in samples 7 to 28 are plotted for the reference data in the first test at $G = 13.3$. The mean noise level in sample 1, equivalent to the RON, is $4.87 \pm 0.14 e^-/\text{sample}$. In the samples containing stellar images the mean noise level is $12.67 \pm 0.60 e^-/\text{sample}$. This indicates a DOB level of $11.40 \pm 0.90 e^-/\text{pixel}$ (see eqn 2.4). Since no DOB is expected, there must be an additional source of low-level illumination present that has not been taken into account. On the right of Figure 2.42 comparable results are shown for the reference data in the first part of the third test. In this plot the mean RON is $4.85 \pm 0.10 e^-/\text{sample}$ and the mean noise level in the image readout is $5.72 \pm 0.11 e^-/\text{sample}$. This indicates a DOB level of $0.77 \pm 0.09 e^-/\text{pixel}$, closer to the expected value. Since the third test was conducted some time after the first test, the stability of the test conditions appear to have improved over the course of the campaign.

Table 2.11 gives a summary of the mean background noise levels observed throughout RC2. During the first test, conducted at $G = 13.3$, the DOB level appears to be much higher than expected. This is in fact corrected for by the time the second magnitude in the first test is looked at. Throughout the second

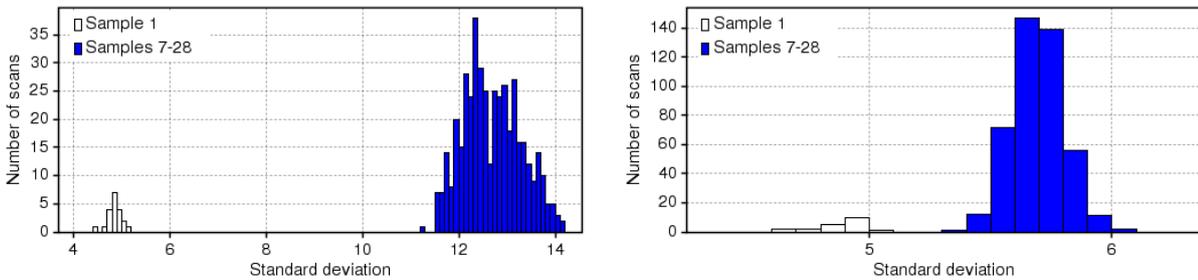


Figure 2.42: Noise in the background in sample 1 and samples 7 to 28. (left) From reference data in the first test at $G = 13.3$. (right) From reference data in the first part of the third test at $G = 15$.

First test				
G	Injection level [V]	Mean prescan RON [e ⁻ /sample]	Mean scan total background noise	Observed DOB [e ⁻ /pixel]
13.3	Reference	4.87 ± 0.14	12.67 ± 0.60	11.40 ± 0.90
	8.06	4.91 ± 0.13	12.95 ± 0.51	11.97 ± 0.78
	No injection ¹	4.83 ± 0.15	13.17 ± 0.52	12.51 ± 0.81
15	Reference	4.87 ± 0.10	6.26 ± 0.14	1.29 ± 0.12
	8.06	4.91 ± 0.06	6.56 ± 0.29	1.58 ± 0.23
	No injection	4.75 ± 0.09	6.48 ± 0.26	1.62 ± 0.20
Second test				
G	Injection duration [pix]	Mean prescan RON [e ⁻ /sample]	Mean scan total background noise	Observed DOB [e ⁻ /pixel]
15	Reference	4.90 ± 0.12	5.74 ± 0.20	0.74 ± 0.15
	10	4.90 ± 0.06	6.02 ± 0.21	1.02 ± 0.15
	No injection	4.80 ± 0.19	5.95 ± 0.23	1.03 ± 0.19
Third test ²				
G	Injection delay [pix]	Mean prescan RON [e ⁻ /sample]	Mean scan total background noise	Observed DOB [e ⁻ /pixel]
13.3	Reference	4.85 ± 0.10	5.72 ± 0.11	0.77 ± 0.09
	1000	4.85 ± 0.09	6.76 ± 0.22	1.85 ± 0.18
	No injection	4.87 ± 0.10	6.85 ± 0.51	1.93 ± 0.42
16.75	Reference	4.88 ± 0.10	5.94 ± 0.15	0.96 ± 0.12
	1000	4.84 ± 0.10	6.19 ± 0.92	1.24 ± 0.67
	No injection ¹	4.89 ± 0.08	6.15 ± 0.18	1.16 ± 0.14
20	Reference	4.87 ± 0.14	6.06 ± 0.15	1.08 ± 0.13
	1000	4.86 ± 0.08	6.11 ± 0.24	1.14 ± 0.18
	No injection	4.87 ± 0.06	5.79 ± 0.16	0.82 ± 0.11

¹Results come from the first of two test runs taken without injection.

²All reference data were taken with a G = 15 source.

Table 2.11: Statistics on the noise in the background in RC2.

and third tests the DOB level hovers around 1.0 e⁻/pixel. In contrast, the RON remains consistent around 4.8 e⁻/sample, which agrees with the value quoted by Astrium.

Since the background observed in the first test performed at G = 13.3 does not reflect the relative stability seen in the background observed thereafter, the statistical analysis presented next is derived from the reference data for the first part of the third test. Performed at G = 15 to calibrate the G = 13.3 irradiated data, the results presented in Table 2.12 corroborate the trends observed in Table 2.4 in RC1. In general, the noise $\bar{\sigma}_c$ (see eqn. 2.6) in the mean background per channel \bar{b}_c (see eqn. 2.5) remains consistent across all 22 channels, while the goodness-of-fit (GOF) parameter g (see eqn. 2.8) hints at a systematic offset existing in the background noise. As in RC1, the background errors are not completely randomised, and evidence to date suggests the PEM anomaly is the probable cause.

The plot on the left of Figure 2.43 shows how the Gaussian variable g in Table 2.12 compares to a standard Gaussian distribution containing the same number of background measurements. It is clear that the expected Poisson errors in the background are offset by a large bias. The plot on the right explicitly shows the underlying correlations (see eqn. 2.9) between the background noise levels in each channel. Compared to the table of correlation coefficients (Table 2.5) shown for the reference data in RC1, the

Channel	Sample index	$\bar{b}_c \pm \bar{\sigma}_c$ [$\langle e^-/\text{sample} \rangle$]	χ^2	σ	g
1	7	4557.82 ± 0.15	259.76	3.80	13.02
2	8	4556.39 ± 0.15	302.20	4.10	14.16
3	9	4556.18 ± 0.15	251.75	3.74	12.79
4	10	4555.66 ± 0.15	344.56	4.38	15.19
5	11	4555.05 ± 0.15	314.80	4.18	14.47
6	12	4553.99 ± 0.15	267.75	3.86	13.24
7	13	4552.52 ± 0.15	223.75	3.53	11.96
8	14	4550.84 ± 0.15	271.05	3.88	13.34
9	15	4549.14 ± 0.15	288.34	4.00	13.80
10	16	4547.72 ± 0.15	276.31	3.92	13.48
11	17	4546.29 ± 0.15	293.05	4.03	13.92
12	18	4544.80 ± 0.15	291.38	4.02	13.88
13	19	4543.08 ± 0.15	284.51	3.98	13.70
14	20	4541.33 ± 0.15	262.22	3.82	13.09
15	21	4539.74 ± 0.15	272.82	3.89	13.38
16	22	4505.42 ± 0.15	226.96	3.55	12.06
17	23	4506.34 ± 0.15	247.73	3.71	12.68
18	24	4508.13 ± 0.15	264.60	3.83	13.16
19	25	4509.70 ± 0.15	272.01	3.89	13.36
20	26	4510.61 ± 0.15	264.80	3.84	13.16
21	27	4510.76 ± 0.15	217.43	3.48	11.76
22	28	4510.73 ± 0.15	287.74	4.00	13.78

Table 2.12: Global statistics on the background in the reference data from the first part of the third test.

correlations in this set of reference data are even higher. Again, this is probably caused by the variable electronic offset added by the PEM.

In general, the irradiated data in the first two tests contained the same features as those already seen in the reference data. To cover all the additional features only the third test needs to be considered. All of the following examples are taken from the first part of this test, when $G = 13.3$.

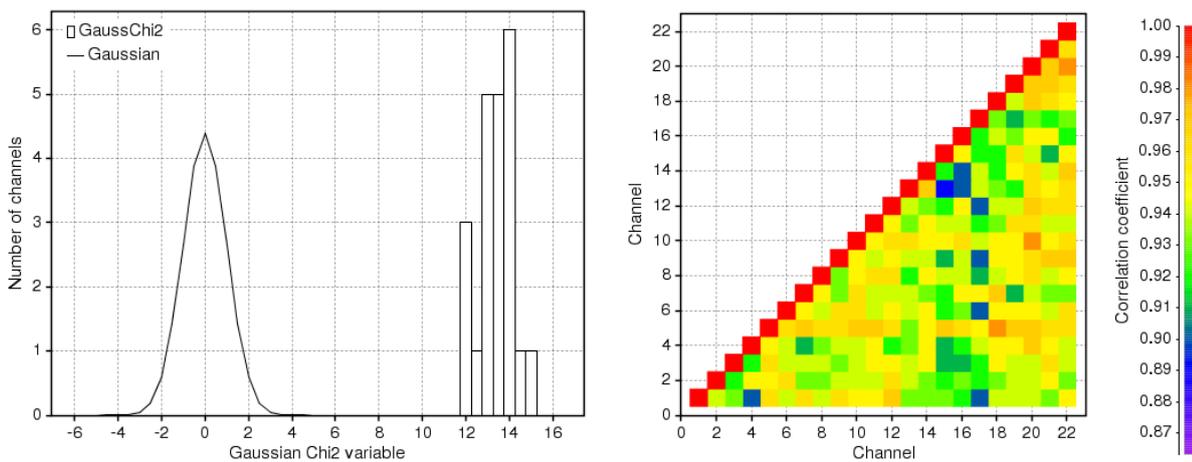


Figure 2.43: Dependencies in the background noise of the reference data found in the first part of the third test. (left) Histogram of the GOF parameter g plotted against a Gaussian distribution of $\sigma = 1$ containing the same number of measurements. (right) Correlation coefficients between variations in the mean background measured in each channel.

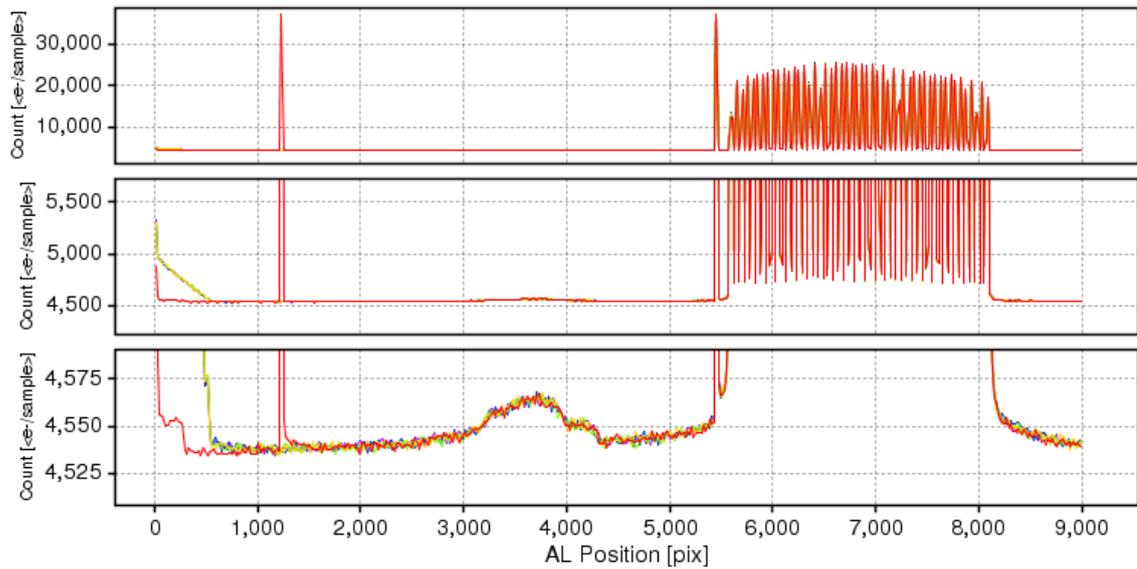


Figure 2.44: Running mean of all 5 scans in channels 1 to 3 of the irradiated data taken from the third test at $G = 13.3$ with a CI delay of 150 pixels.

Figure 2.44 shows a running mean of the raw data from the first three channels taken with a CI delay of 150 pixels. The following observations can be made:

- In this particular test case two injections have been made. The first appears ~ 1200 pixels into scan 1 and the second appears ~ 150 pixels before the stellar field is read out. The additional injection in scan 1 was made in several test cases, e.g. with delays of 30 pixels and 60 pixels, but left out in others. Each scan lasts for 9000 TDI sample periods.
- The charge tail caused by the illuminating source remaining on as the mask returns back over the CCD is again present at the beginning of scans 2 to 20. The flux in scan 1 takes two steps to settle into the background.

Figure 2.45 shows comparable results for a CI delay of 30000 pixels, except here channels 1, 2 and 22 are shown. Notably:

- Each scan lasts for more than 50000 TDI sample periods.
- The light leaking around the mask seen earlier in the reference data now appears with the same form and at an equivalent distance after the stellar field has passed.
- The running mean clearly shows the tail of the CI in channel 22 extending as far as ~ 20000 pixels beyond the actual injection. Note that the injection was made from the bottom of the CCD, so the injection level will be the highest in channel 22¹¹.

¹¹Recall that as the CI is driven from only one end of the CCD, the AC profile of the CI level diminishes across the CCD reducing in intensity between channel 22 and channel 1. In effect, this means each channel receives a different charge injection.

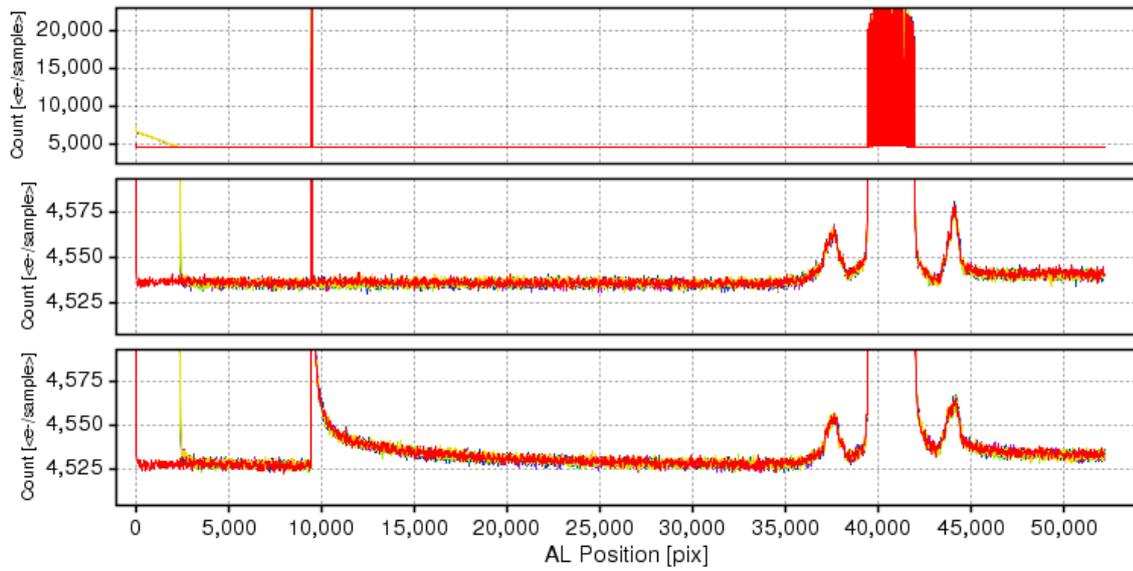


Figure 2.45: Running mean of all 20 scans in channels 1, 2 and 22 of the irradiated data taken from the third test at a $G = 13.3$ with a CI delay of 30000 pixels.

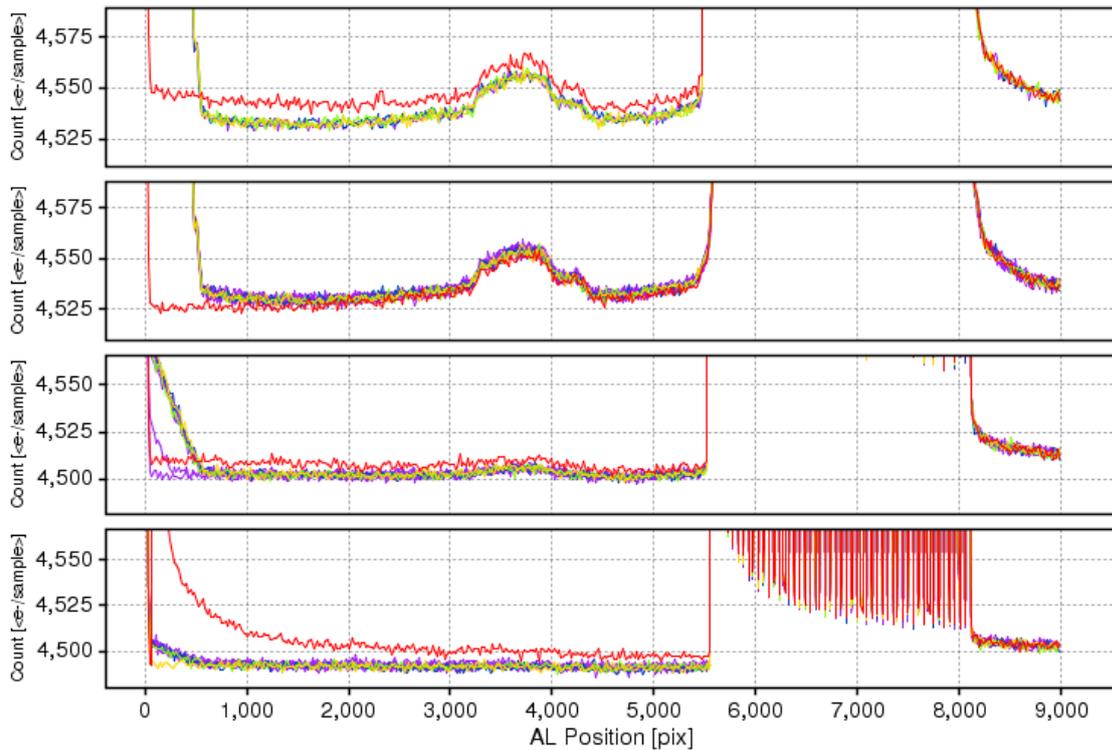


Figure 2.46: Additional features seen in the third test: (top/first) $G = 13.3$, delay 100, 5 scans, (second) $G = 13.3$, no injection, 15 scans, (third) $G = 15$, delay 60, 10 scans, (bottom/fourth) $G = 16.75$, delay 30, 10 scans. All of these plots show the readout in channel 22.

Figure 2.46 shows some additional features in the irradiated data that highlight predominantly the behaviour of the signal in scan 1. It is clear that processes occurring outside the IA are somewhat

unpredictable at the beginning of each test case and it takes a certain period of time for the PEM to stabilise. Since scan 1 (red line) is affected on many occasions by the PEM, it has not been included in any further analysis to extract image parameters.

One other feature to be aware of and illustrated in Figure 2.47 is when the time between scans began to vary. During these instances the CI delay tended to drift considerably from the expected value and scans where this occurred were not included in subsequent analysis.

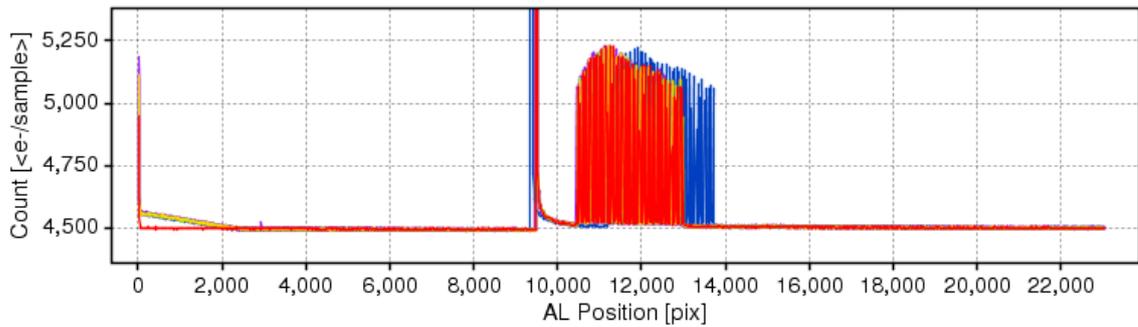


Figure 2.47: Time variations between scans. Here in the third test at $G = 16.75$ with a CI delay of 1000, one scan (out of 10) in channel 22 is read out late relative to the other scans.

In the irradiated data, the anomaly associated with the PEM in the AC direction can be seen in Figure 2.48, where the mean background per scan has been extracted from the third test at $G = 16.75$ and a CI delay of 300. On the left, the systematic drop in the background can be seen to take place between samples 18 and 19 (channels 12 and 13). Sample 14 (channel 8) also shows a distinct jump in the estimated count relative to the other measurements. On the right, the background appears to be relatively stable over consecutive scans.

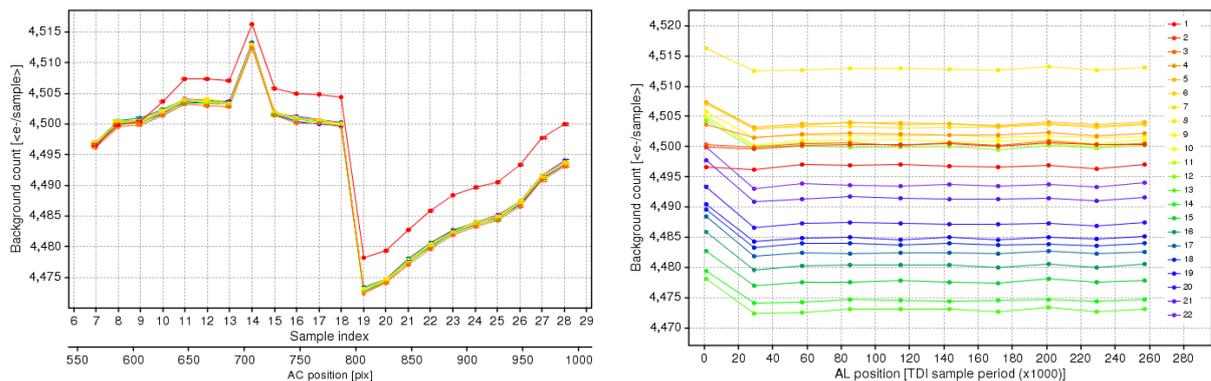


Figure 2.48: Background behaviour in the irradiated data from the third test at $G = 16.75$ and a CI delay of 300. (left) Mean background count as a function of AC position. The large systematic offsets are attributed to the PEM anomaly. Scan 1 (red) can be seen deviating from all the other scans beyond sample 9. This is due to an anomalous feature in the data. (right) Mean background count measured over consecutive scans. The background appears stable from one scan to the next.

Lastly, we show how the spacing between the pinholes on the mask in the AL direction and the intensity of the source affects the resolution of the stellar field. Figure 2.49 shows a portion of the

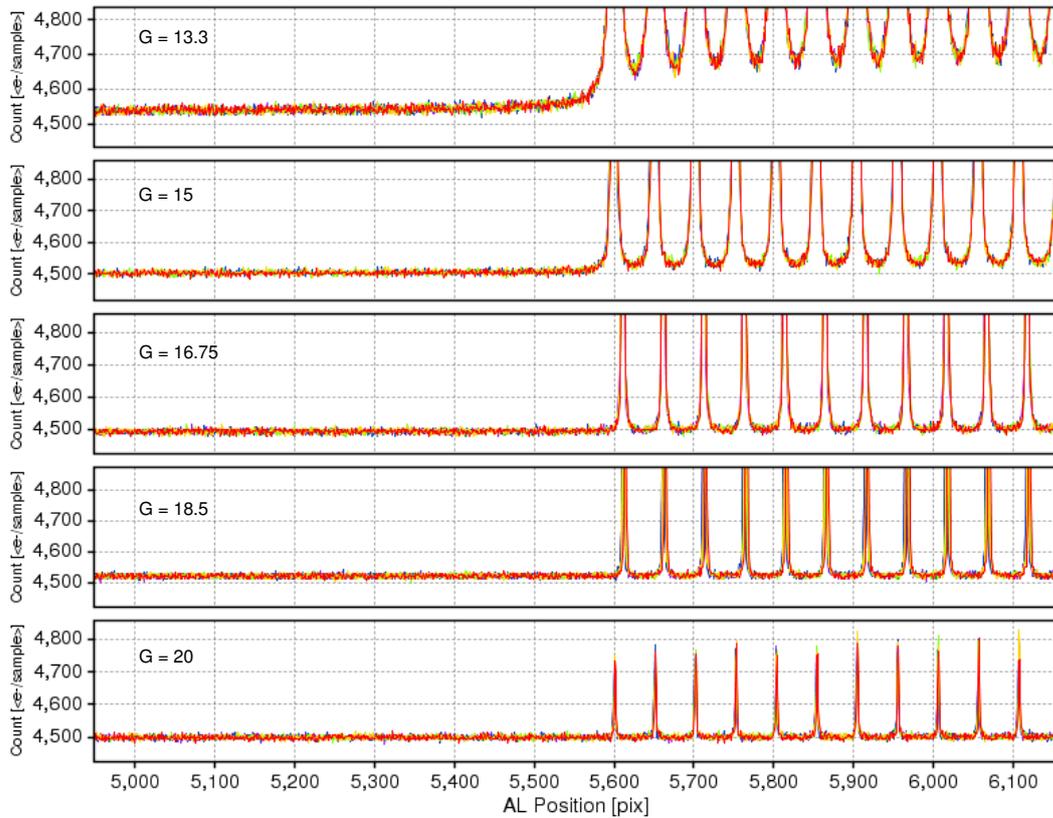


Figure 2.49: How image spacing influences the estimation of the background level.

readout centred at the point where the stellar field is beginning to appear. The data is taken from channel 1 in the third test when no injection was applied. With each star separated by ~ 50 pixels AL there is a high degree of overlap when $G \lesssim 15$, and over a 15-pixel window only the faintest images come close to merging with the local background. The gradual rise in the background under each image reflects the non-uniform level of illumination across the mask.

This result is important because it clarifies which method may be used to estimate the background under the stellar field. In RC1 the flat background cut method was superseded by a local background cut. However, in RC2 we can clearly see a significant degree of overlap between the brightest images. Bearing in mind that each image in the AL direction is separated by 50 pixels (this is the quoted pitch) and every image is being extracted using a 15-pixel wide window, we must estimate the background, at least for the brightest images ($G = 13.3$, 15 and possibly 16.75), using a flat background cut. This includes all of the reference data. For the faintest images we may be able to use the local background cut or introduce a new model that can extract the CI tail, although this model still requires a flat background cut. At the time Brown & van Leeuwen (SWB-002) was being prepared, the reanalysis of RC1 had not yet begun and the flat background cut was the only method available. Once found, the local background cut method was employed to reanalyse RC1 and the CI tail model, developed by Development Unit (DU) 10^{12} , was incorporated into the analysis of RC3. The remainder of this analysis therefore uses the flat

¹²DUs are sub-groups of a CU. Work on characterising the CI tail was carried out at the Institute for Astronomy, Edinburgh.

background cut method to derive CTI image parameters.

At this point it is worth emphasising that the methods discussed above for subtracting the background are not optimal. We must assume that variations in the readout in the AL direction are effectively the same in both the reference and irradiated data, which may not always be the case. The calibration of the charge loss will be wrong if the background in both sets of data are significantly different. Large variations in individual channels are not seen, so the results for the brighter images will be reasonably accurate. The fainter images on the other hand will be far more sensitive to background variations so the final results will have to take account of this. Ideally, the electronic offset would be a unique value over the entire CCD. However, the non-uniformity issue with the current laboratory-model PEM does not permit this.

To summarise:

- There are certain features in the raw data not previously seen in RC1:
 1. A sharp decrease in the flux at the beginning of some scans. This is caused by the PEM and does not interfere with processes occurring in the IA.
 2. A few bumps in the background either side of the stellar field, due to light leaking around the mask.
 3. A linearly decreasing flux at the beginning of some scans, caused by the light being kept on as the test equipment is reset after each scan.
 4. The beginning of scan 1 shows variable behaviour compared to the scans that follow in some test cases. This is probably caused by processes occurring in the PEM.
- Multiple occasions when an extraneous signal (cosmic) is seen.
- The PEM continues to have a major influence on how the background level is set. There are sharp discontinuities ranging between ~ 25 and ~ 40 e^- /sample between two neighbouring channels as well as strong correlations between the variations in the background measured in each channel.
- The observed DOB levels in the first test at $G = 13.3$ are far higher than expected. In the remaining tests the DOB level approaches zero as expected.
- The spacing between pinholes on the mask are close enough to ensure every image overlaps in the AL direction (and probably the AC direction). Nonetheless we are still able to extract images of the individual stars according to the method described in section 2.1.3 as the flux threshold, which is dependent on the background level, remains at a level between the background and the peak fluxes of the images. For $G \leq 15$ the background is never reached between images over a 50-pixel spacing, while over a 15-pixel window only the faintest images may sample the background. In the analysis that follows a flat background cut is made.

2.3.4 The reference data

The Global Reference Curve (GRC)

Comparable with Figure 2.16, Figure 2.50 shows four of the steps taken to define the GRC. Here the first column in the reference data is taken from the first test at $G = 13.3$. The symmetry of the GRC is far better than the original GRC observed in RC1.

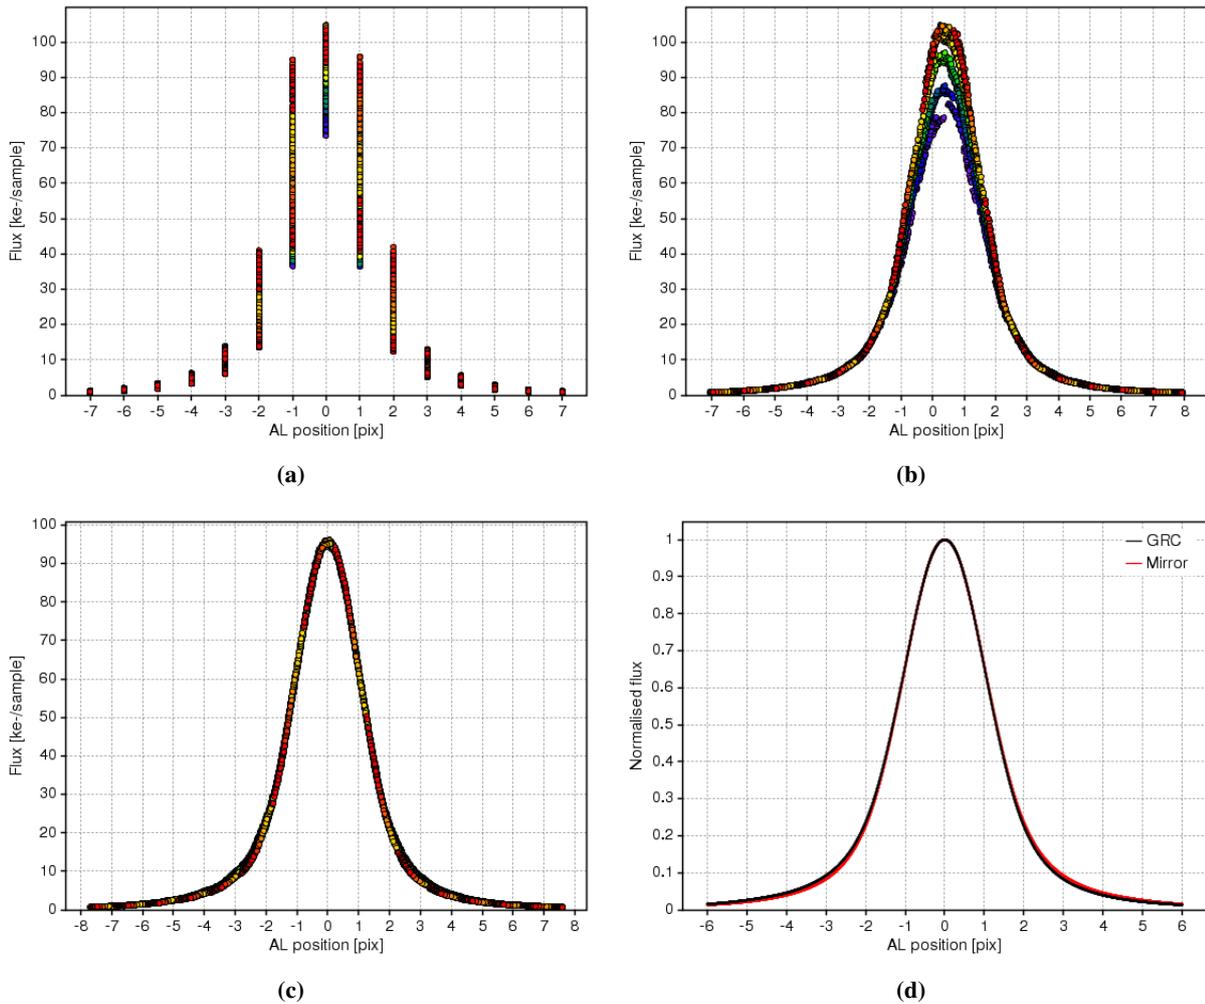


Figure 2.50: Constructing the GRC. (a) Stacking the raw images. (b) Following cross-correlation and quadratic interpolation. (c) Following height and width scaling, plus successive iterations to remove positional errors and define the centroid. (d) The normalised GRC constructed using a cubic spline fit and 16 regularly spaced knots with mirror image.

Geometric and flux response calibration

In RC1 a single GRC was constructed to determine relative centroid positions and fluxes in the leading column of stars. In RC2 we have 50 columns of stars to work with and 22 stars in each column. The

question is: can a single GRC be used to calibrate the geometry of the entire mask? Determining how the width of the average response varies across the mask will decide whether a single GRC can be used. Scaling the GRC does not alter the width, so this parameter would have to be known beforehand. Figure 2.51 shows the FWHM of every GRC constructed for each column along the mask. The results are very similar, since all but one set of reference data was taken at $G = 15$. As the FWHM does vary along the mask we chose to define a GRC for each column. Interestingly, the GRC defined for column 16 (only shown here in three cases) has a FWHM that does not conform to the general quadratic trend down the mask. On closer inspection this feature was caused by a much reduced flux through hole 14, equivalent to the same hole that was reported by Astrium to have been blocked. As this disrupts the GRC the associated image was excluded from further analysis.

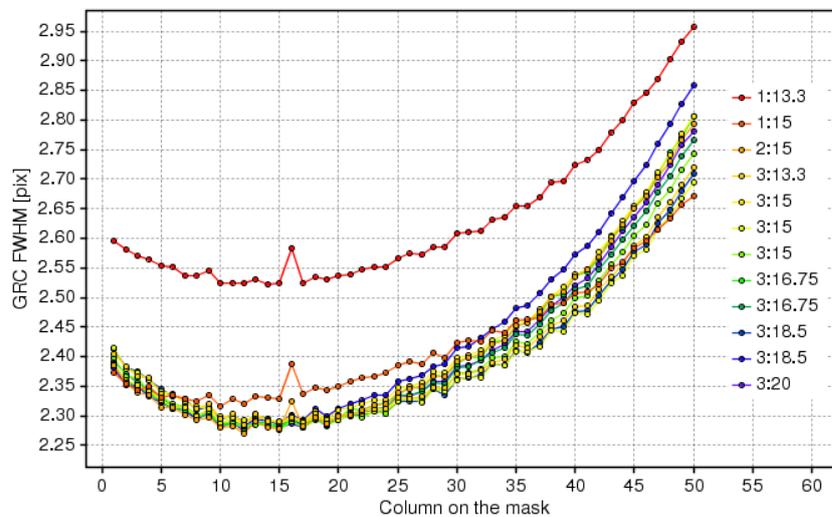


Figure 2.51: The GRC FWHM for each mask column. Legend refers to 'test:magnitude'.

Each GRC is used to extract average relative centroid positions of every image in the reference data. A subset of these positions in column 50 and column 1, taken from the first test at $G = 13.3$, are shown on the left and right respectively in Figure 2.52. To mark the expected positions (see section 2.1.7) a linear fit to all of the datapoints in column 1 should be made. However, the non-linearity of the positions in column 1 do not justify such a fit. A more sensible choice is to use the centroid positions in column 50. In fact, the same ambiguity in column 1 is seen in every set of reference data, so the initial fit in RC2 is always made using column 50. Since the rotation should be the same in all columns, and only the first two channels are used in the irradiated tests to calibrate the data, the initial fit to column 50 is refitted to only those centroids in the first two channels in every column. The refit (blue line) to column 50 and column 1 is shown on the left and right-hand sides of Figure 2.52 respectively. Now it is clearer what is actually occurring in column 1: towards channel 22 a clear geometrical offset emerges. Whether this is caused by a pressure or thermal gradient, or simply a misalignment between the holes on the mask is unclear at this stage. In this example, the rotation of the mask's leading column with respect to the CCD displaces the centroid position of the image in channel 22 by ~ 1.99 pixels, amounting to a rotation of $\sim 0.27^\circ$ about the zeropoint.

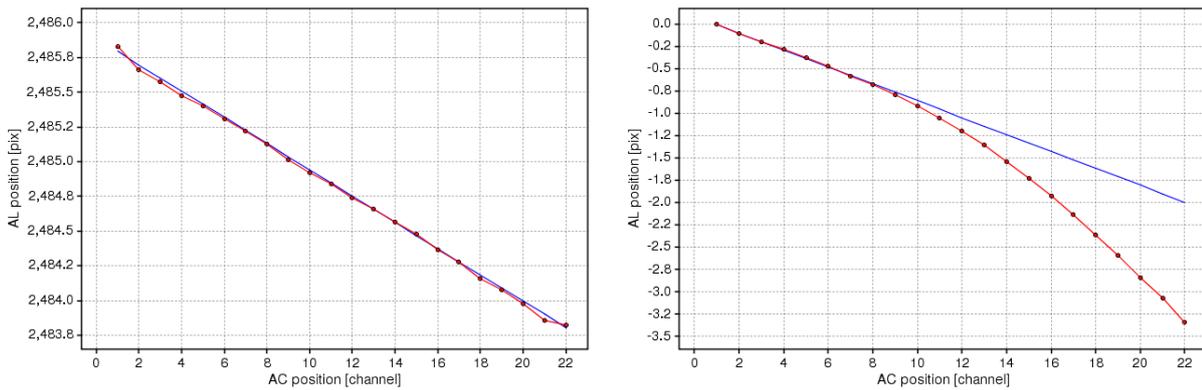


Figure 2.52: Calibrating the geometry of the mask I. This data is taken from the first test at $G = 13.3$. Relative centroid positions extracted from (left) column 50 and (right) column 1 with the final geometric calibration.

Any expansion in the AL direction is automatically taken account of by applying this linear fit to the first two channels in each column. To show this expansion explicitly (see Figure 2.53(a)) a rigid grid is constructed where every linear fit is separated by exactly 50 pixels (the quoted AL pitch). Starting from column 50 and working back up the mask we find that the differences between the centroid positions and the grid accumulate, extending up to ~ 37 pixels beyond where they are expected to be in column 1. The blank square in the middle of the plot corresponds to column 16, channel 14, which was earlier identified as being blocked.

Following on from Figure 2.52, the remaining geometrical offsets between all fifty sets of centroids and their associated fits can be found. Figure 2.53(b) shows how the offsets vary across the mask. Most

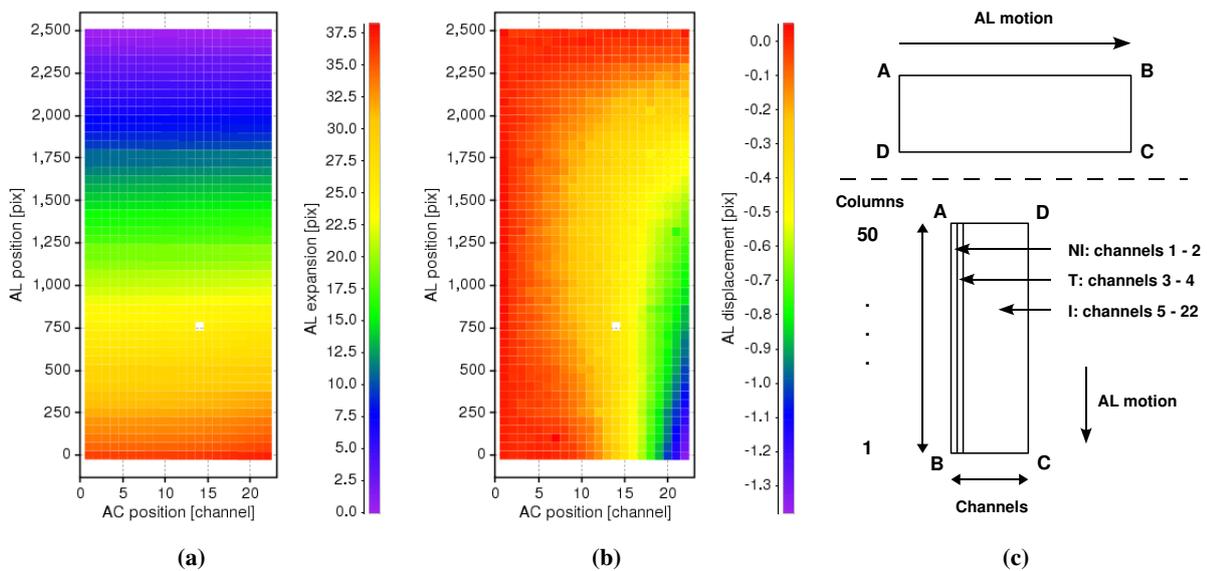


Figure 2.53: Calibrating the geometry of the mask II. (a) An apparent expansion of the mask. (b) Remaining offsets between the centroid positions and the geometric calibration. (c) A diagram explaining how to interpret the preceding plots. Above the dashed line the mask is shown as depicted in Figure 2.36. NI = non-irradiated, T = transition, I = irradiated.

of the holes are well-aligned, apart from those found towards the front of the mask close to channel 22. The same feature is seen in every set of reference data. Figure 2.54 complements Figure 2.53(b) by separating out the geometrical offsets across the mask (column 1 is red at the bottom and column 50 is violet at the top) to give better contrast.

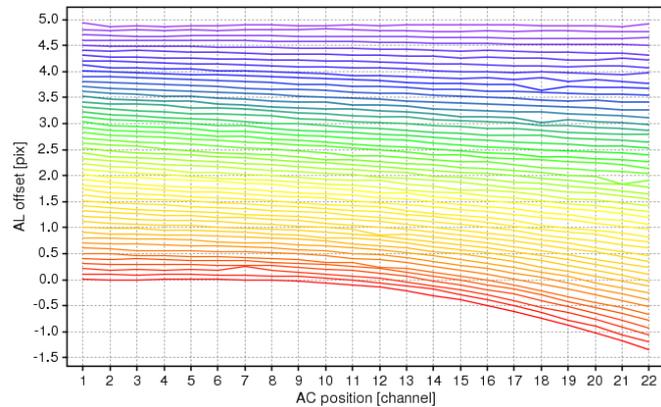


Figure 2.54: Here the geometrical offsets found across the mask in the reference data (first test, $G = 13.3$) have been spaced out to simply highlight how the offsets vary across the mask. The AL offset does not correspond to the AL position on the mask shown in Figure 2.53.

The procedure followed to calibrate the flux response is described in section 2.1.8. Consequently, comparisons were made between the peak fluxes, integrated fluxes and FWHM of each fitted GRC and local LSF across the mask.

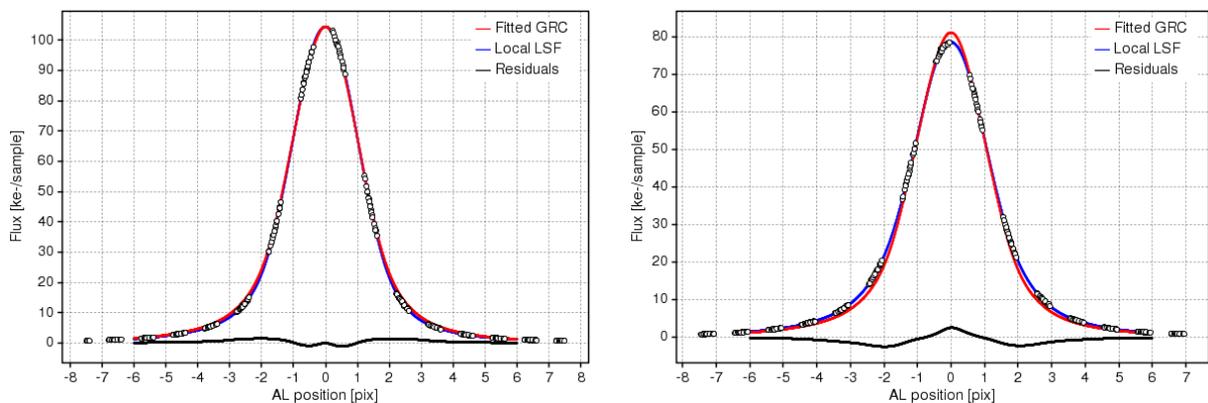


Figure 2.55: Calibrating the flux response across the mask I. Results are taken from the reference data in the first test at $G = 13.3$. (left) Fitting the image in channel 1, column 1. (right) Fitting the image in channel 22, column 1.

Figure 2.55 shows the accumulated reference data (first test, $G = 13.3$) from channel 1 (left) and channel 22 (right) in column 1. The fitted GRC and local LSF are shown as well as the residuals between the two fits. In channel 1 the match between both curves is very close, whereas in channel 22 there are a few places where the match deviates. There are a few reasons why this may be so:

- The datapoints tend to cluster, which suggests that the time between scans was almost identical.

Whilst this ensures the CI timing is the same in each scan, it has the adverse effect of reducing the effective coverage of the sub-pixel samples.

- A reanalysis of RC1 (Brown & van Leeuwen (SWB-003)) has shown how sensitive the flux correction can be to the estimate of the background underneath the stellar image. A better fitting is achieved if the flux level in the wings of the image is closer to zero. Using a flat background cut, the wings in each image in Figure 2.55 remain above zero which may hinder the final flux fitting.
- The width of the actual image in column 1 is varying, which cannot be accommodated by the fitting procedure.

Addressing these issues further, Figure 2.56(a) shows the peak flux of every fitted GRC across the mask. The non-uniformity in the illumination is clearly visible. Figure 2.56(b) shows the same plot, but this time the peak fluxes are taken from every local LSF across the mask. In both of these plots the maximum and minimum peak flux have been fixed to highlight the larger dynamic range observed by the local LSF. Figure 2.56(c) shows the residual peak flux between these two fits. Clearly, the peak of the fitted GRC is lower than the local LSF in the lower channels and higher than the local LSF in the higher channels. This is the case in every column and there is a smooth transition between the two extremes. Comparing this result to comparable plots for the FWHM (Figures 2.56(d), (e) and (f)) we see the reverse trend exists. The peak flux residuals are inversely related to the FWHM values. This variation essentially ensures the integrated count in both the fitted GRC and local LSF in each image are identical, differing by $<0.1\%$, as verified in Figure 2.56(g).

Figure 2.56(g) also contains some anomalous features in columns 9, 12 and 23 that contrast starkly with the rest of the field. To explain these anomalies we must consider how each fitted curve is constructed. In Figure 2.56(g) the knots used to construct each local LSF were placed at regular intervals along a 12-pixel window. In Figure 2.56(h) the positions of the knots have been customised to cluster more around the core of the local LSF. When this is done these anomalous features disappear. Although by increasing the density of knots the local LSF is more sensitive to variations in the accumulated data, which increases the size of the residuals.

We could adopt this customised knot scheme, however, when the irradiated data in the first and second tests are analysed only four scans are available to sample the LSF. This makes the curve fitting numerically unstable and forces the knots to be spaced further apart. For this reason we chose to use a regular knot sequence for fitting the local LSF. It highlights how insufficient sampling of the LSF limits the sensitivity of the fitting procedure.

2.3.5 The irradiated data

In RC2 charge was injected from only one side of the CCD. If the injection was $>7.51\text{V}$, the readout in channel one was no longer saturated. Table 2.13 gives an example of how the flux levels in the last pixel of the injection vary between channels 1 and 22. Note that the strength of the injection in channel 1 is negligible when the CI level is nominally set to $\geq 8.65\text{V}$. A more thorough analysis of the variation

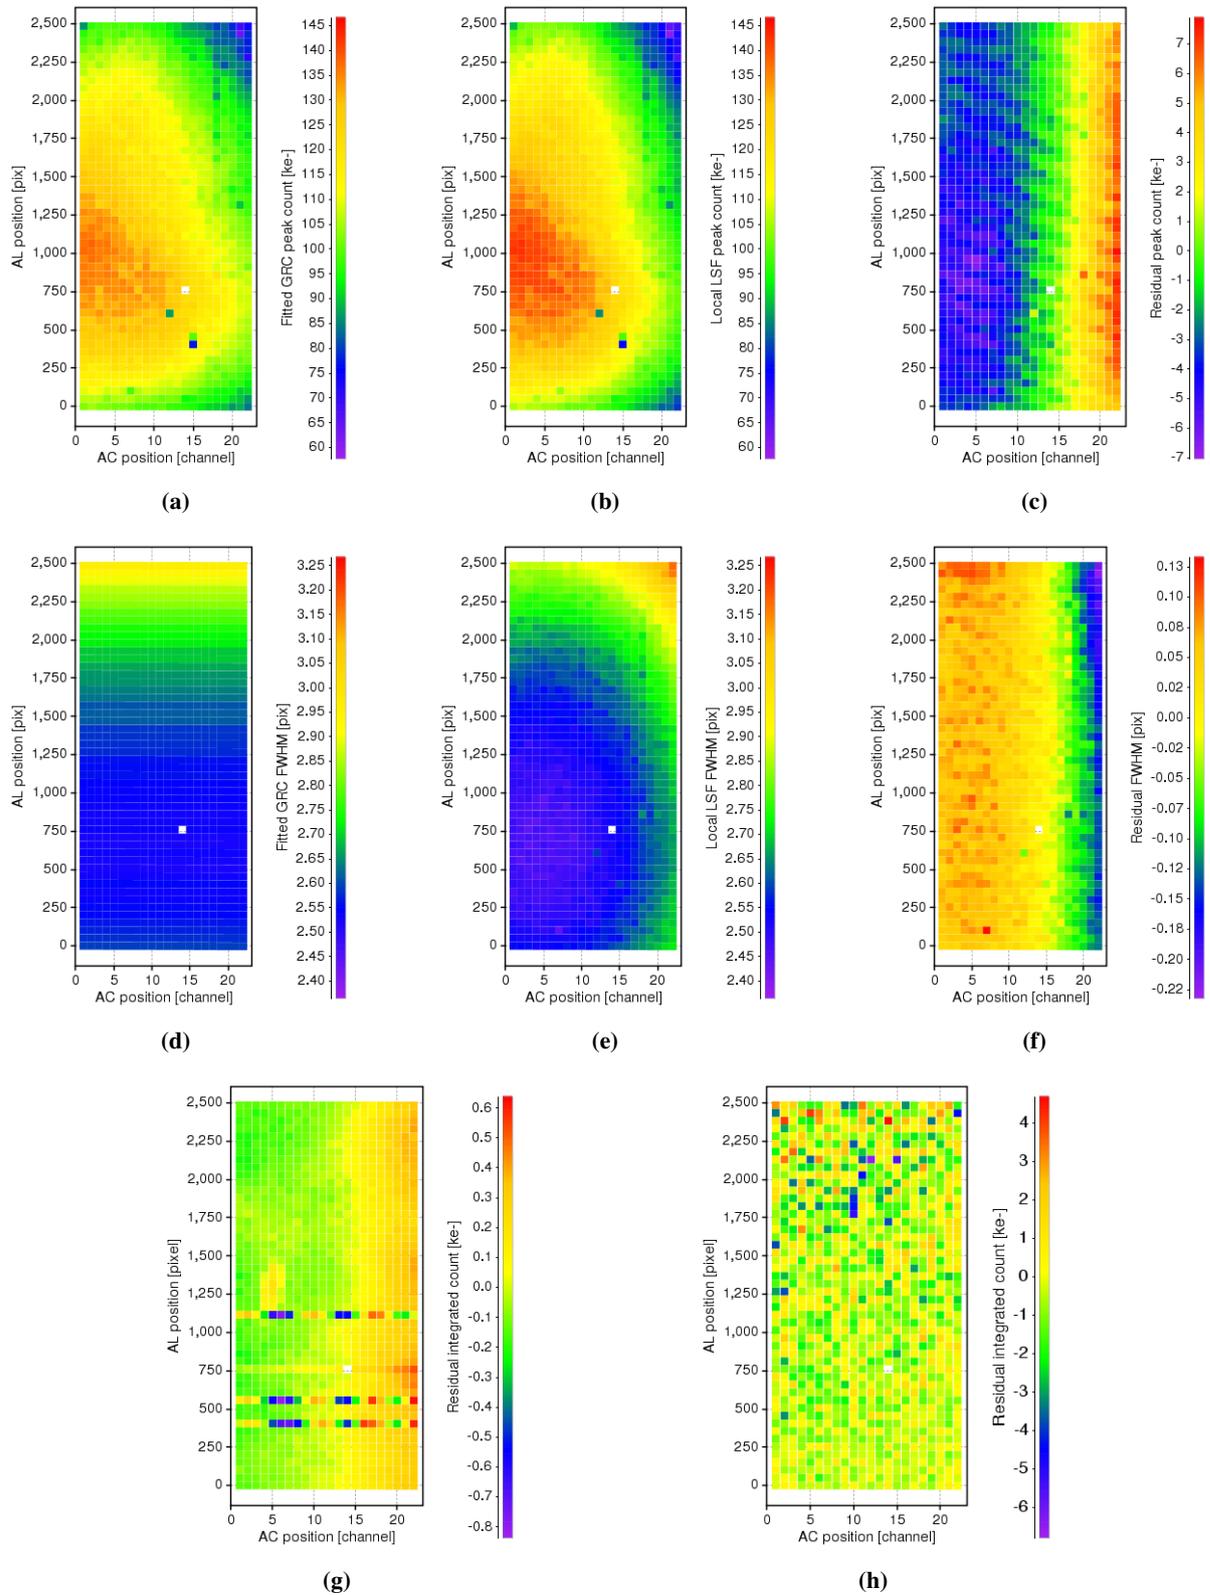


Figure 2.56: Calibrating the flux response across the mask II. These results are taken from the reference data in the first test at $G = 13.3$. Note that the orientation of these plots is the same as that prescribed in Figure 2.53. (a) Fitted GRC peak flux, (b) local LSF peak flux, (c) residual peak flux between (a) and (b), (d) fitted GRC FWHM, (e) local LSF FWHM, (f) residual FWHM between (c) and (d), (g) residual integrated flux using regularly spaced knots, and (h) residual integrated flux using customised knots.

in the CI level AC can be found in EADS Astrium (GAIA.ASF.TCN.PLM.00212) (see section "Extra information added (09 October 2008)").

Inj. level [V]		6.51	6.96	7.31	7.51	7.84	8.06	8.26	8.42	8.57	8.65	8.72
Flux	C1	65535	65535	65535	65535	42493	18934	6794	2669	1440	-	-
[LSB]	C22	65535	65535	65535	65535	65535	65535	65535	58227	41309	33197	27042

Table 2.13: A typical example of the variability in the CI level in each channel observed throughout RC2. C1 and C22 are for channel 1 and channel 22 respectively. 65535 = saturation.

Investigating the CI level

In the first test the brightest stars at $G = 13.3$ and 15 were used to find a CI level that could optimise the mitigation of CTI effects. The results presented in this section will focus on the analysis performed at $G = 13.3$. In total 11 CI levels were tested (see Table 2.8). An additional test at 7.51V with gate 10 activated was also performed. However, this test case was not analysed. Gates will be studied in more detail in RC3.

Recalibrating the rotation of the mask Following the methodology described in section 2.1.7 to calibrate the geometry of the mask in the irradiated data, Figure 2.57 shows the relative positions of the image centroids in column 1 (left) and column 50 (right) overlaid with the geometric calibration.

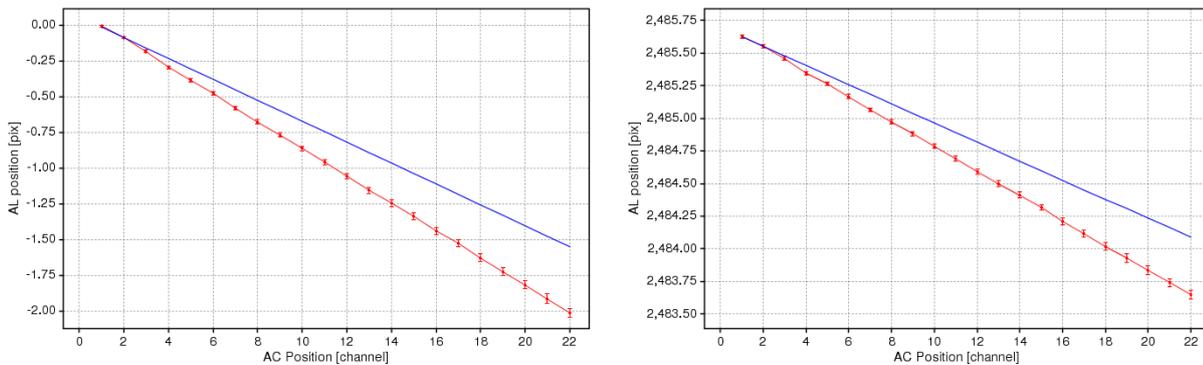


Figure 2.57: Relative centroid positions in (left) column 1 and (right) column 50, corrected for geometrical offsets for a CI of 7.51V.

By following standard procedure (i.e. using channels 1 and 2 in the irradiated data to carry out the calibration), all the other centroids show negative biases. With a CI of 7.51V, where the signal saturates in the PEM, you would expect the centroids to show no damage. The centroid positions should align themselves with the geometric calibration (blue line), which is governed simply by the rotation of the mask with respect to the CCD. As will be shown, these negative biases do not have any physical meaning. The calibration has been disrupted by the way the CCD is manufactured.

Since the rotation of the mask could not be determined using the information in the non-irradiated channels alone, we had to use a part of the irradiated data which we could assume was undamaged.

Naturally the best place to make this assumption was in column 50 and any other columns prior to this that may have been shielded by the preceding stars. This raises the question: how many columns can we include without jeopardising the accuracy of the bias measurements?

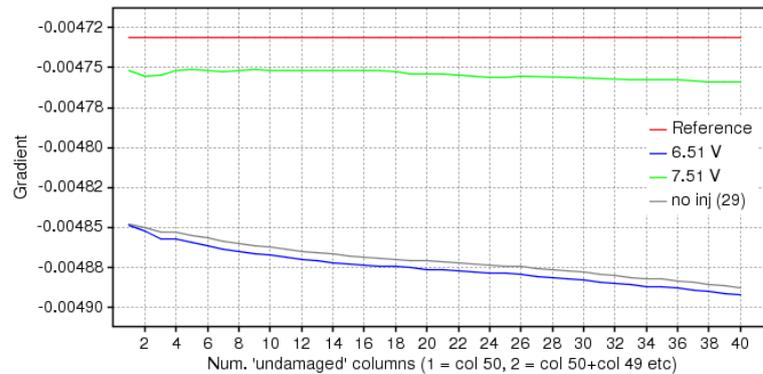


Figure 2.58: The gradient of the line of best fit is used to investigate which irradiated image centroids appear undamaged. In the first test at $G = 13.3$ two test cases were conducted without injection. The test case shown here refers to the second, performed on 29/06/07.

Figure 2.58 shows the gradient of the line of best fit for the geometric calibration as a function of the number of columns counting back from 50. The red line shows the gradient in the reference data. With a CI of 7.51V the gradient appears to settle around one value when the number of columns included lies between 5 and 20. However, for a CI at 6.51V and for a test case without CI the gradient is constantly changing and no particular value stands out. From this plot we chose to include 20 columns (columns 30 to 50) to determine the rotation of the mask. This balanced the need to include a sufficient number of "undamaged" images whilst maintaining a suitable distance from column 1.

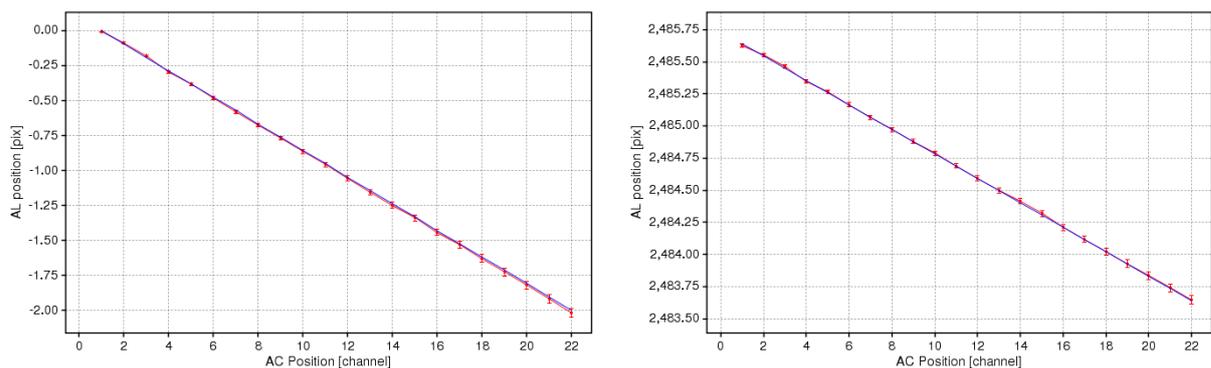


Figure 2.59: Calibrated relative centroid positions in columns 1 (left) and 50 (right). The CI is 7.51V.

Figure 2.59 shows the result of including the extra information from the irradiated centroids to recalibrate the rotation of the mask as well as the offsets found in the reference data. Recall that these offsets include the curvature seen in column 1 in Figure 2.52 (right plot) and Figure 2.54. As expected, the centroids across all 22 channels are now very close to the undamaged positions defined by the calibration.

Relative centroid bias across the stellar field Figure 2.60 shows how the centroid positions in column 1 beginning with channel 3 are successively damaged as the injection level decreases (cf. Figure 2.7).

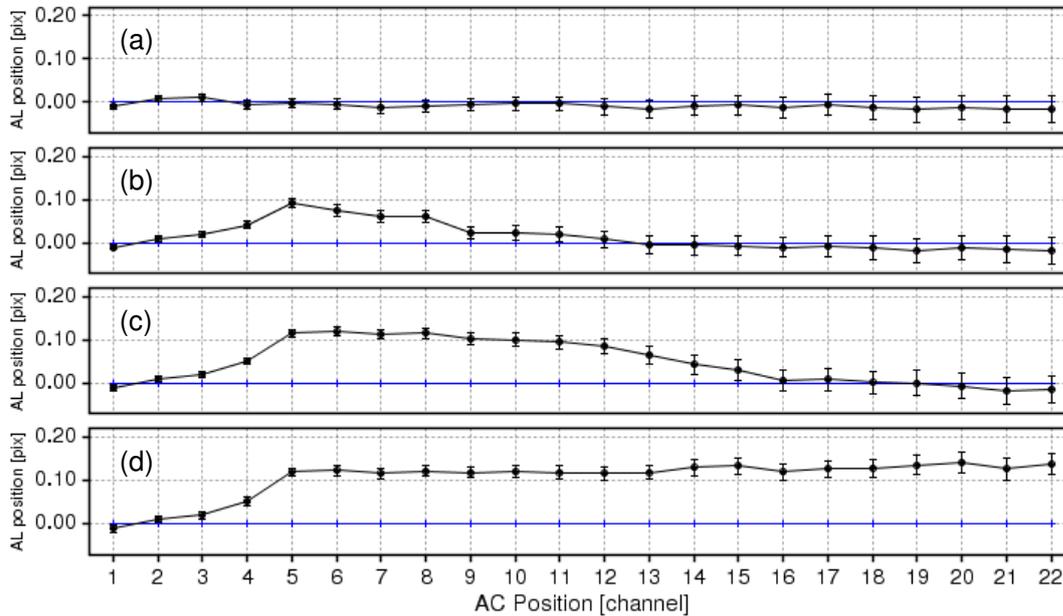


Figure 2.60: The onset of centroid bias in column 1 at (a) 8.06V, (b) 8.42V, (c) 8.65V and (d) without CI (22/06/07).

Now that the centroid positions have been calibrated the remaining offsets indicate the extent of the bias on the centroid. Figure 2.61 contains a selection of plots showing the relative centroid bias observed across the stellar field for various CI levels. Any voltage level higher than 8.06V begins to show damage in the first column, so Figure 2.62 shows how the biases appear at this level. Figure 2.63 shows the worst-case scenario when no injection is applied. Here the bias can reach as high as 0.12 pixels (~ 7 mas). From these three figures several observations can be made:

- Firstly, a clear systematic bias can be seen in channels 2, 3, 14 and 15, appearing in every column. The reason behind this shall be explained shortly.
- Secondly, with intense injections the biases measured closest to the injecting source are positive. This implies that either: (a) there are still underlying offsets in the data that have not been accounted for by the calibration; or, (b) the CI is lifting the leading edge of the image shifting the centroid forward. Since the flat background cut does not exclude the CI tail, it is more likely that (b) is occurring in this case.
- As the CI level decreases the centroid biases in column 1 gradually increase, appearing first in channel 3 and eventually extending all the way down to channel 22.
- CTI effects are most profound in column 1.
- Even with the lowest injection level, CTI effects appear to be reduced considerably in the channels closest to the injecting source.

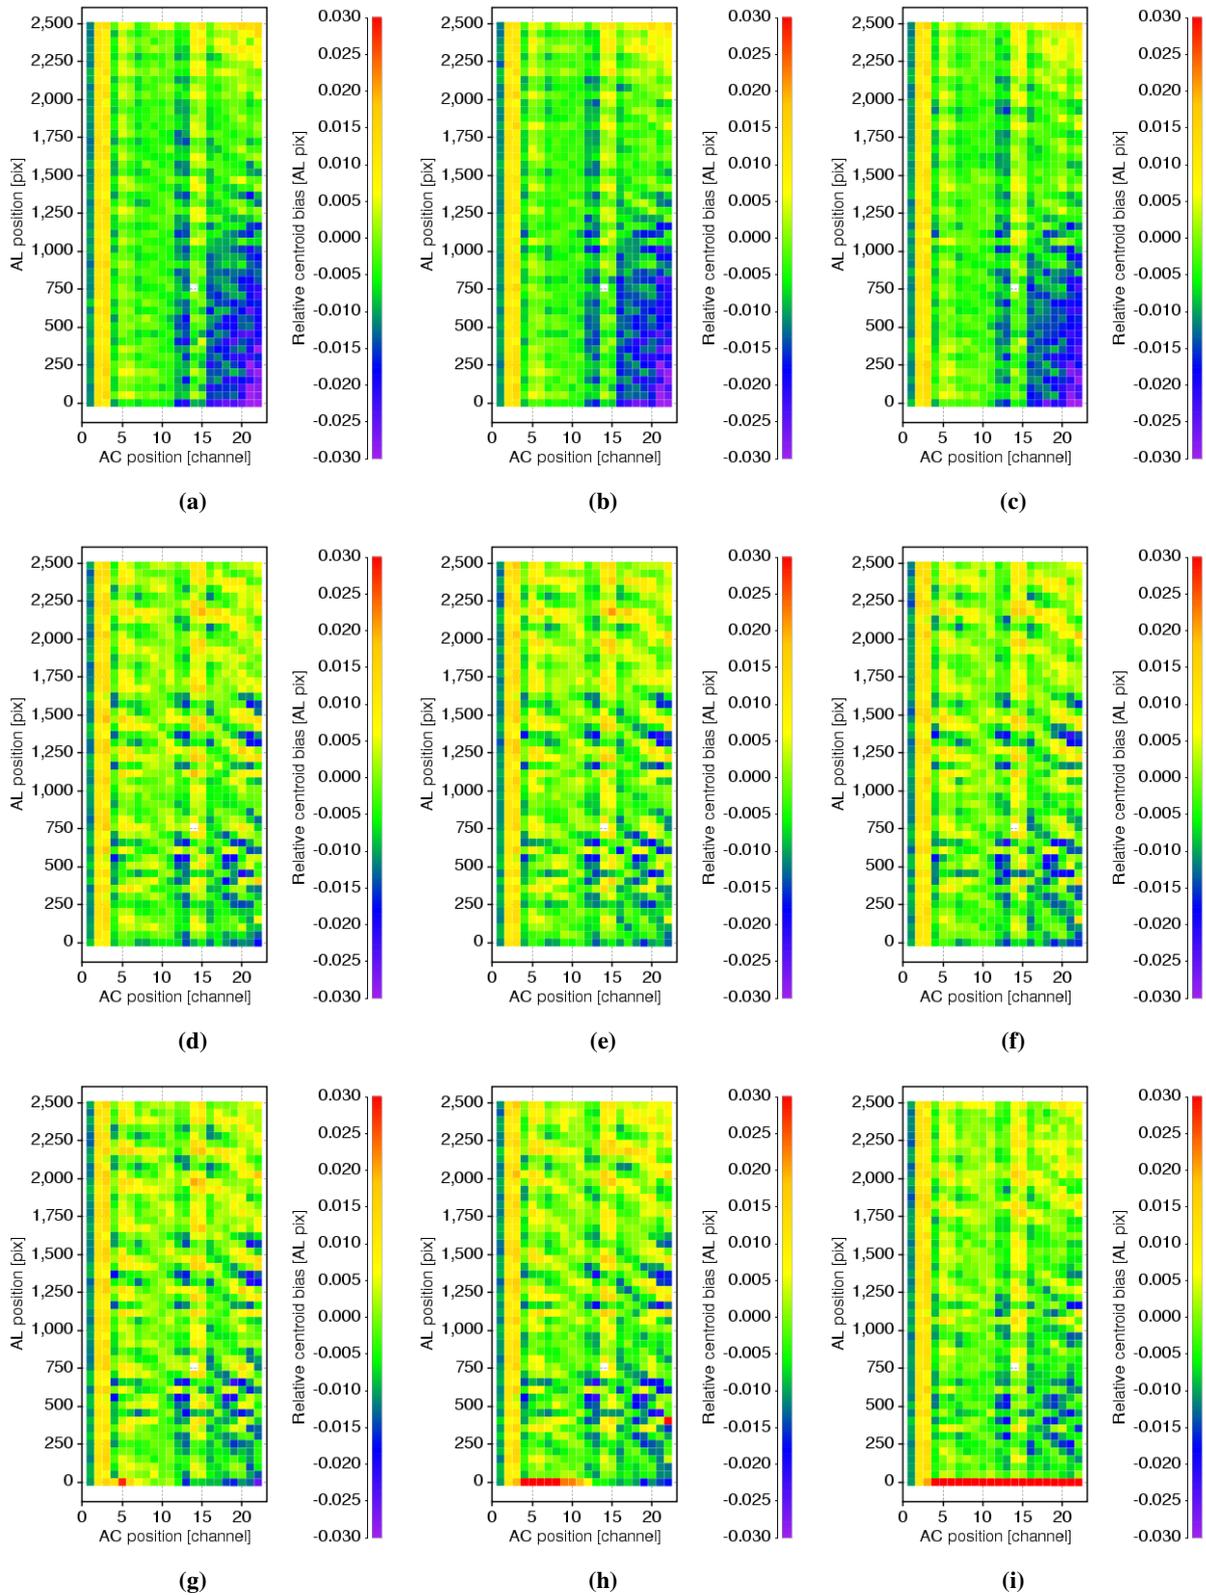


Figure 2.61: Relative centroid biases across the stellar field caused by radiation-induced CTI with a CI at (a) 6.51V, (b) 6.96V, (c) 7.31V, (d) 7.51V, (e) 7.84V, (f) 8.06V, (g) 8.26V, (h) 8.42V and (i) without CI (the first of two performed on 22/06/07). The scale on the bias has been fixed to highlight changes between the different injection voltages.

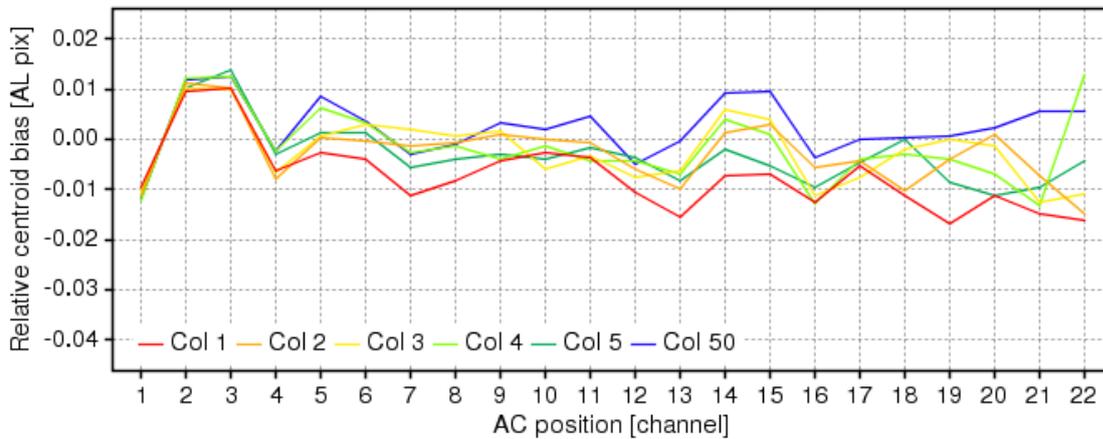


Figure 2.62: Relative centroid biases seen edge-on when the CI is 8.06V.

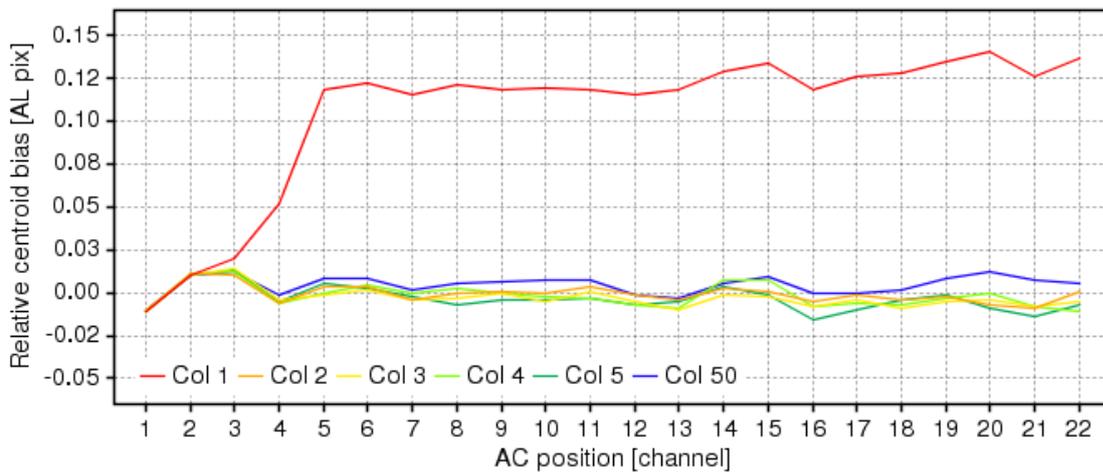


Figure 2.63: Relative centroid biases seen edge-on without CI (22/06/07).

- Regardless of the CI level, after column 1 has passed the centroid bias decreases significantly.

CCD stitching Once the issue concerning systematic centroid biases in channels 2, 3, 14 and 15 was raised, an immediate explanation was found¹³. These biases are caused by CCD stitches. Figure 2.64 illustrates the Gaia CCD stitching pattern. CCDs are not manufactured as one single unit but 'stitched' together by treating smaller, adjacent sections of the silicon wafer. The borders between these blocks are called stitches and, in the case of Gaia, are placed ~ 250 pixels apart in the AC direction, the same distance separating channels 2/3 and 14/15 on the CCD. In RC2 one such stitch happened to coincide with channel 2, systematically displacing the centroids in that channel, and preventing the geometric calibration being performed with just the non-irradiated channels.

This analysis suggests the AL centroid bias caused by a CCD stitch is ~ 0.01 - 0.02 pixels ($\sim 0.15\mu\text{m}$ or $\sim 0.89\text{mas}$ on the sky). However, this may vary depending on the properties of the stitch, the type of

¹³Feedback was provided by J. de Bruijne at the European Space Research and Technology Centre (ESTEC) and G. Seabroke at the Mullard Space Science Laboratory.

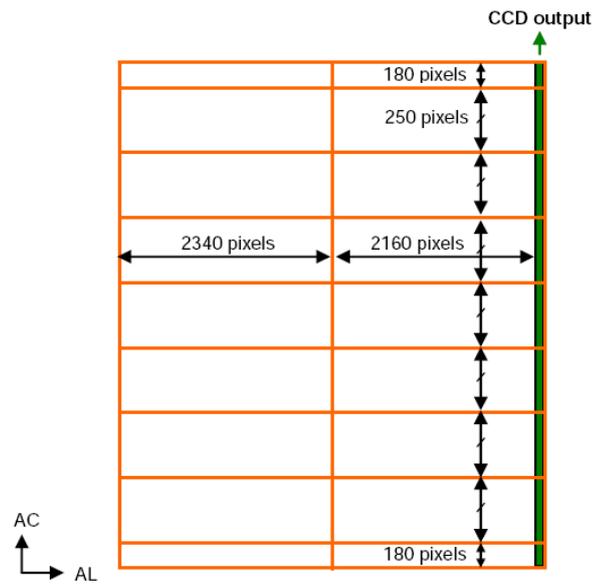


Figure 2.64: CCD stitching pattern. Credit: EADS Astrium (GAIA.ASF.TCN.PLM.00313).

CCD, etc.

For a more detailed discussion on CCD stitches and the effect they may well have on the data processing see de Bruijne (JDB-043). With regards to RC testing, non-irradiated channels falling on CCD stitches is a situation that must be avoided. These effects are further quantified by EADS Astrium (GAIA.ASF.TCN.PLM.00313) where the test bench has been initially characterised as part of RC3.

Relative charge loss across the stellar field Figure 2.65 shows two examples of image profiles read out at the same time from different AC positions in column 1 with a 6.51V injection. Overlaid are the fitted GRC and a local LSF. As with the reference data, the GRC does not always fit precisely to the local LSF, although further analysis has shown that the total charge under each curve is effectively the same.

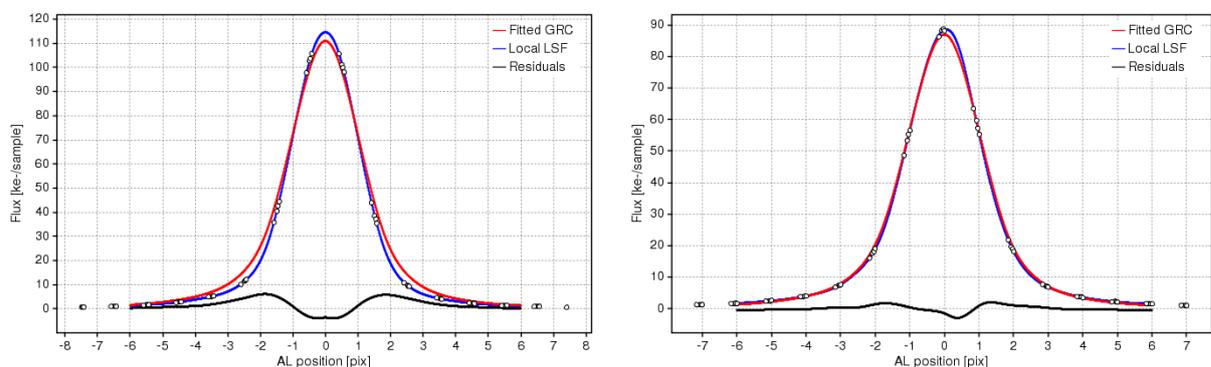


Figure 2.65: Fitting images in channels 1 (left) and 22 (right) in column 1 with a CI of 6.51V.

With a limited number of datapoints that tend to cluster around the same sub-pixel positions, a regular knot scheme again caused resonances in the residual integrated flux. By shifting the positions of the knots

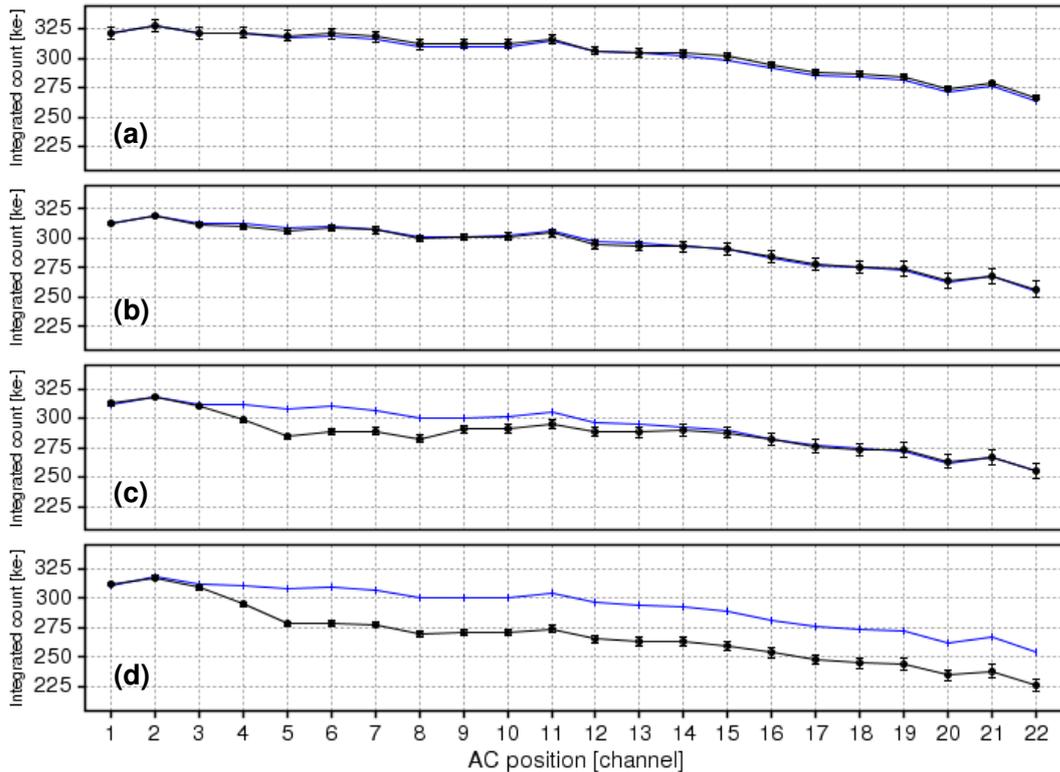


Figure 2.66: The onset of charge loss in column 1 with CI at (a) 8.06V, (b) 8.42V, (c) 8.65V and (d) without CI (22/06/07).

the resonances also shifted. In one case the intensity of the resonance was reduced significantly. Clearly, it is important to avoid correlation effects in the fitting procedure. One way of achieving this in future campaigns is to increase the number of scans in the irradiated tests.

Figure 2.66 shows how the integrated flux as measured in the reference data (black line) is used to calibrate the charge loss in the irradiated data (blue line) for a selection of CI levels (cf. Figure 2.8).

Figure 2.67 shows the relative charge loss observed across the stellar field for various CI levels. The range is fixed between -1.0% and 1.5% to highlight any changes over the CCD. Figure 2.68 and 2.69 show how the charge loss appears when the CI is 8.06V and no injection is applied respectively. The following observations are made:

- From high to intermediate CI levels (6.51V to 8.26V) the relative charge loss in column 1 is negative, indicating more charge is present than expected. This is probably due to the CI tail not being removed.
- Channels 12 and 13, located next to a stitch, always show more charge loss relative to other channels.
- Channels 19 and 20 always show less charge loss relative to other channels.
- Without removing the CI tail, an injection close to 8.06V appears to balance the charge excess and charge loss seen in the first column.

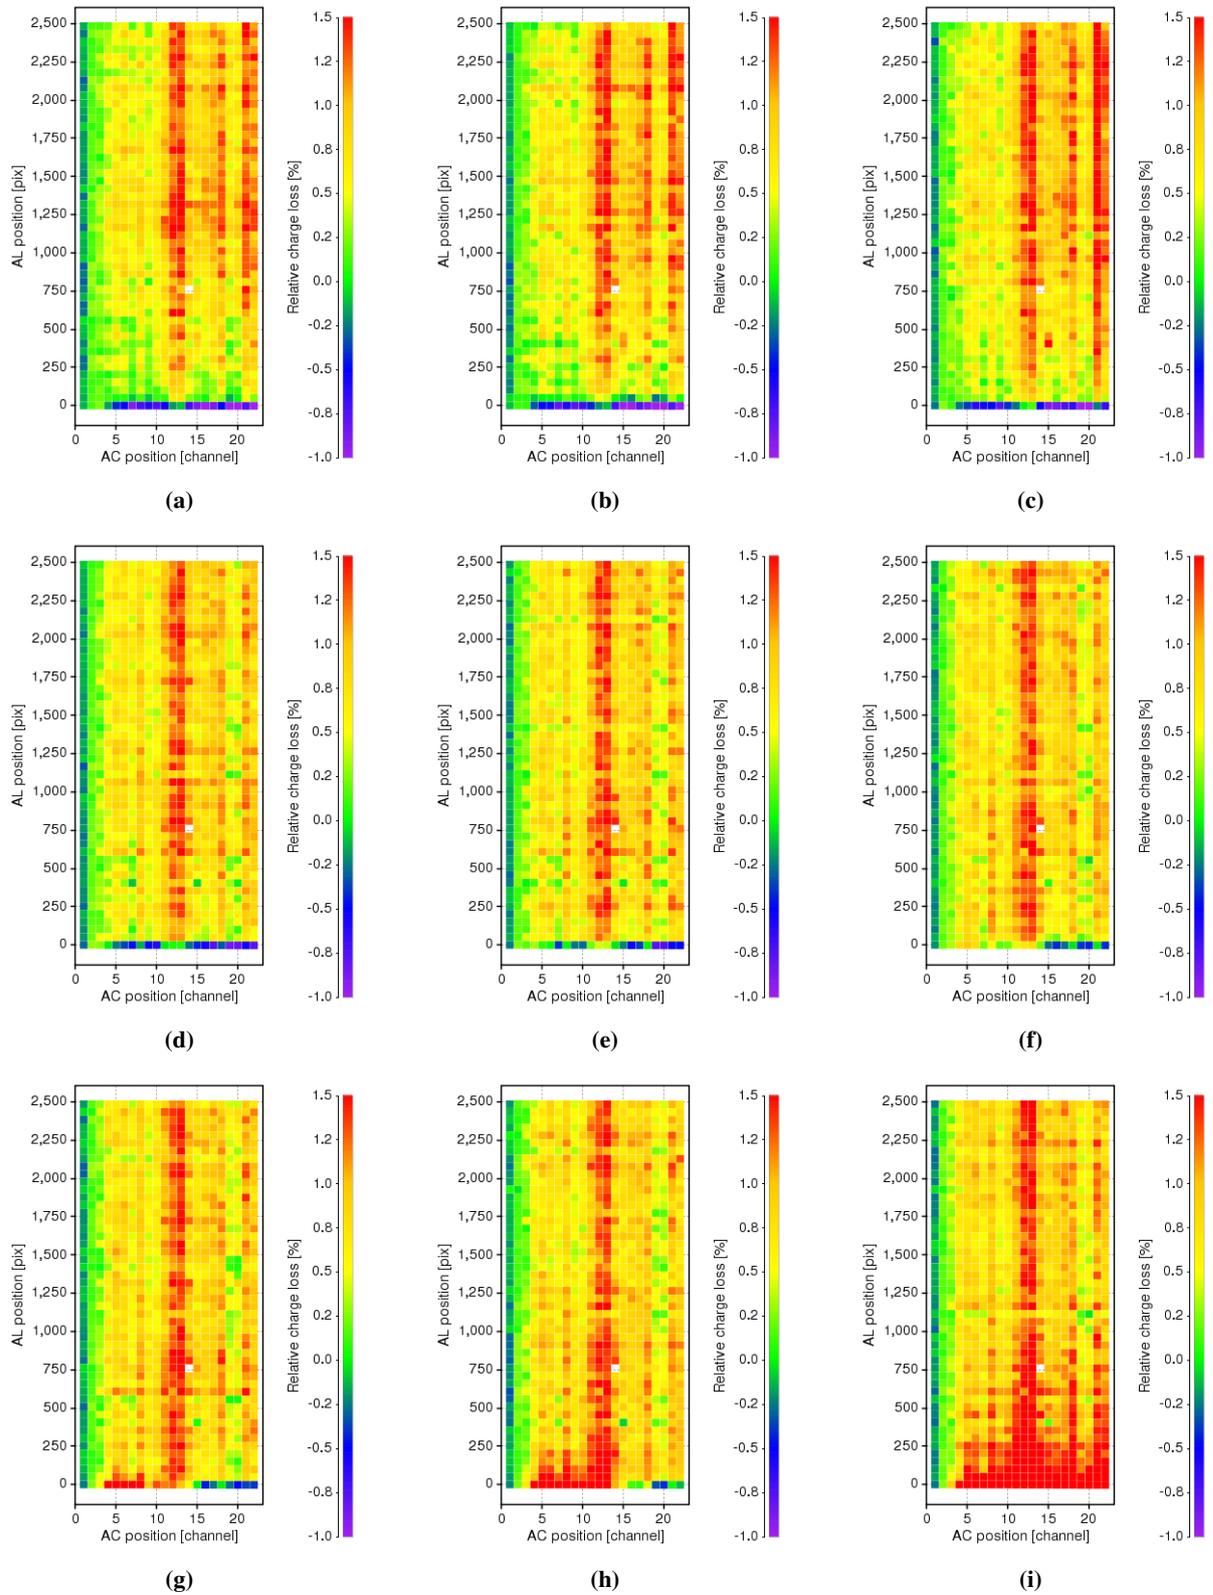


Figure 2.67: Relative charge losses across the stellar field caused by radiation damage at CI levels of (a) 6.51V, (b) 6.96V, (c) 7.31V, (d) 7.51V, (e) 7.84V, (f) 8.06V, (g) 8.26V, (h) 8.42V and (i) without CI (22/06/07).

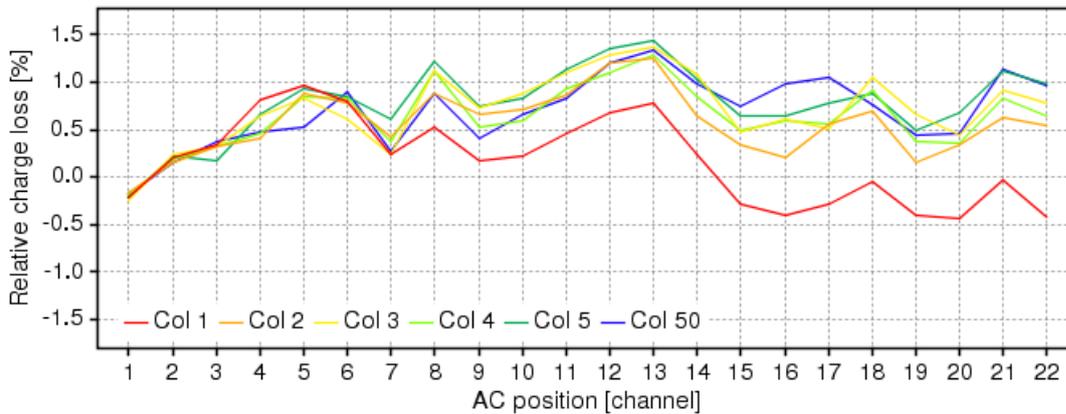


Figure 2.68: Relative charge losses seen edge-on when the CI is 8.06V.

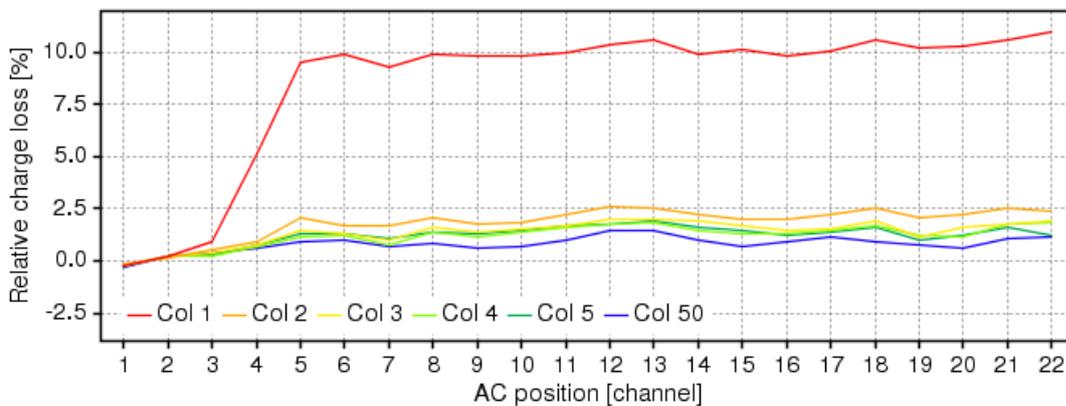


Figure 2.69: Relative charge losses seen edge-on without CI (22/06/07).

- Without injection the charge loss extends beyond column 1, remaining relatively high close to the stitch.

To summarise, the variation in integrated flux, centroid bias and charge loss over all 22 channels in column 1 for all injection levels are shown in the three respective plots in Figure 2.70. Not all of the tests were performed on the same day, so the variation in the fluxes seen in the non-irradiated channels (1 and 2) reflects the need to recalibrate the light source. This is clearer at $G = 15$ in Figure B.3 in Appendix B.2.

Memory effects Due to the inconsistencies found in the background preceding the stellar field, scan 1 was excluded from the analysis. However, this does not mean the memory of scan 1 on the CCD can be ignored. Memory effects were strong in RC1, so to see if this is the case here a test conducted at $G = 13.3$ without CI was reanalysed with scan 1 included. In this particular case it was possible to extract the background. As expected, it was found that the relative charge loss increased from $\sim 10\%$ to $\sim 16\%$.

Clearly, a fraction of the radiation damage is hidden by scan 1. However, with a CI that saturates the signal preceding the stellar field (a test case was reanalysed at 7.51V with scan 1 included) the charge loss

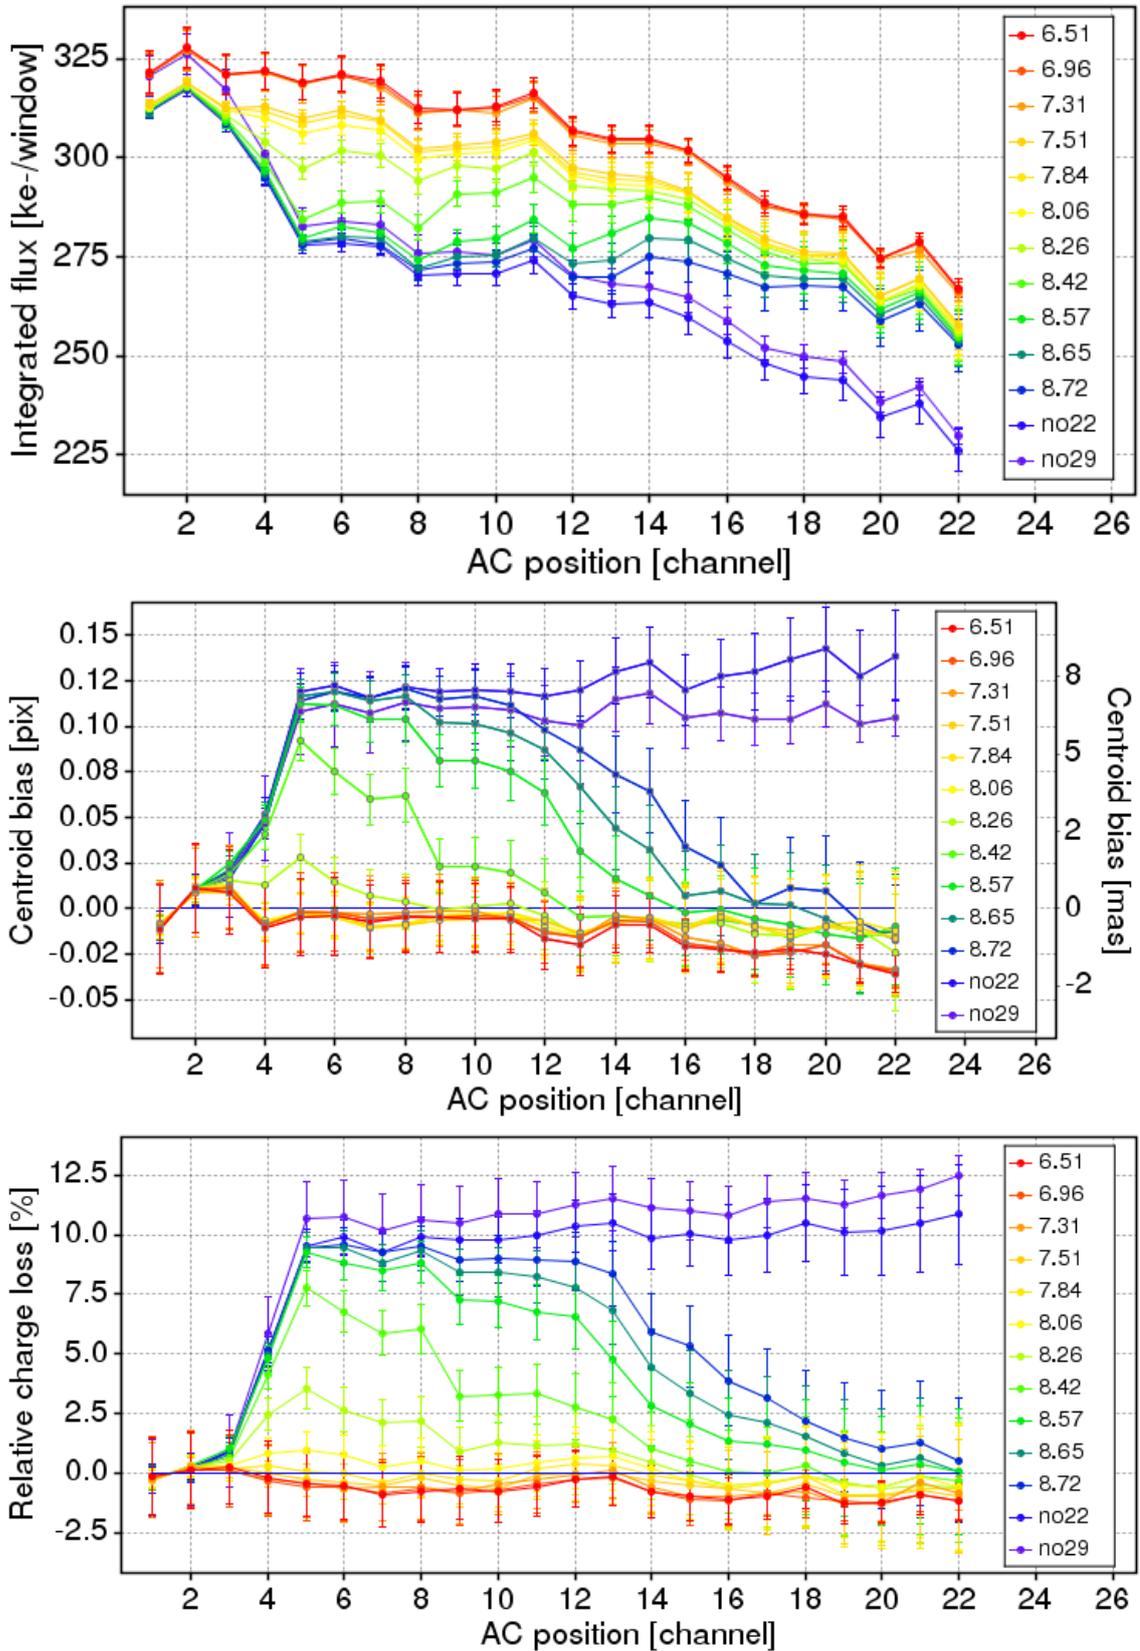


Figure 2.70: (top) Integrated flux ($G = 13.5$ to 13.9), (middle) relative centroid bias, and (bottom) relative charge loss as a function of AC position and CI level in column 1 at a nominal $G = 13.3$. Two test cases were performed without injection: one on the 22/06/07 (referred to as 'no22' in the figures) and the other on 29/06/07 (referred to as 'no29' in the figures).

is the same across the CCD with or without including scan 1. This is to be expected since the memory of traps filled or left unfilled by the preceding scan is removed.

To investigate the radiation conditions scan 1 experiences, the time between scans would have to be increased to allow the CCD to "reset". Astrium have said that it would not be possible to conduct such tests, given the number of other tests to perform and the extended time required. Even so, it raises the question: how much time on average does a CCD column spend observing the lowest sky background level between images? A model of the sky and Gaia's scanning law could be used to estimate this value. However, the work carried out in Holl (BH-001) gives some idea. Statistically, on average each FOV will see a density corresponding to a 50% sky fraction. At magnitudes 17, 18.5 and 20 this corresponds to a cumulative density of 6.35×10^{-3} , 3.51×10^{-3} , and 1.75×10^{-3} star deg^{-2} respectively. Given a spread factor (the spread factor accounts for diffraction and AC motion) of 2.5 (Holl (BH-001) suggest a spread factor between 2 and 2.5 for mags 17 to 20), the average time between stars of magnitude 17, 18.5 and 20 on a CCD column is longer than $\sim 146\text{s}$, $\sim 73\text{s}$, and 40s respectively for two FOVs combined. In the Galactic plane these average times are brought down to 3 to 5s at magnitude 20. However, on the whole it is more likely that on average the time between stars will be longer than 30s. Consequently, the results presented in this analysis, for all scans > 1 , are not entirely indicative of what Gaia may observe most of the time. Therefore, the effects of self-injection may not be as beneficial for Gaia as observed here.

Investigating the CI duration

As with the first test, the two brightest magnitudes ($G = 13.3$ and 15) were chosen to find the optimal CI duration in the second test. Here the injection level was fixed at $\sim 17 \text{ ke}^-/\text{pixel}$ ($\sim 8.06\text{V}$) and the injection delay was set to ~ 25 pixels.

No reference data were collected at $G = 13.3$ so a correction had to be applied to the flux calibration in the $G = 15$ reference data. This correction proceeded as follows. Firstly, it was assumed that if the tests to collect the reference, R , and irradiated data, I , were carried out at the same magnitude the flux in channel 1 (I_1 and R_1) should agree. The scaled reference flux S_i in each channel i ($i = 1, \dots, 10$) should therefore in the first instance be:

$$S_i = \left(\frac{I_1}{R_1} \right) R_i. \quad (2.11)$$

To approximate the final reference flux values in the remaining channels, F_i , the flux ratio, A , between the irradiated data and the scaled reference data in channel 2 relative to the flux in channel 1 is used:

$$F_i = S_1 + (S_i - S_1) \times A \quad \text{where} \quad A = \frac{I_2 - I_1}{S_2 - S_1}. \quad (2.12)$$

Figure 2.71 shows the outcome of this correction for a typical test case before (left) and after (right) the flux calibration in column 1. With only two channels to constrain the flux profile the final result appears ambiguous. However, luckily, the first two channels in column 1 tend to contain integrated fluxes that are sufficiently consistent in their relative intensity to allow the rest of the fluxes in the irradiated channels to be calibrated in this way. Towards the centre of the mask this was not always possible, as the illumination

conditions tended to vary more so than the variability seen in the reference data.

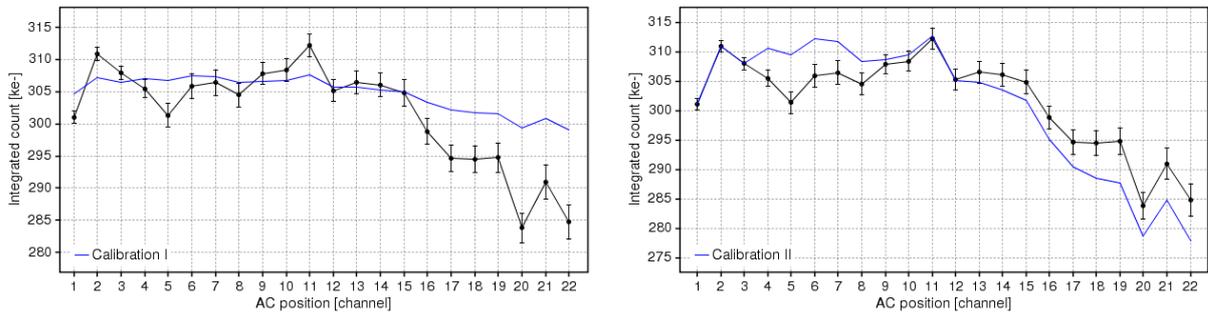


Figure 2.71: Calibrating the charge loss in column 1 for test cases that do not have reference data of the same magnitude. (left) Reference fluxes at $G = 15$ (blue) are overlaid directly with the fluxes at $G = 13.3$ (black) from the irradiated data. As the fluxes in channels 1 and 2 should agree the fluxes are set to match in channel "1.5". (right) References fluxes at $G = 15$ are scaled according to eqn. 2.12 to estimate the undamaged fluxes at $G = 13.3$.

An alternative to using the $G = 15$ reference data is to use the $G = 13.3$ reference data from the first test. However, it is not clear if the conditions on the CCD are similar, since the latter was conducted over 5 weeks before.

The results shown earlier for the first test are also valid here when the CI duration is varied. So to summarise, Figure 2.72 shows the integrated flux (top), relative centroid bias (middle) and relative charge loss (bottom) in column 1 for every CI duration. Comparative plots for $G = 15$ are shown in Appendix B.2.

Clearly, a CI duration of only 2 pixels is enough to remove a considerable part of the radiation damage experienced at both these magnitudes. When the CI is not high enough in the channels furthest from the injecting source an injection duration of 4 pixels is preferred. However, any duration beyond this should suffice. As mentioned before, the CI tail has not been excluded from the image data, so the positive biases and charge losses measured at both magnitudes towards channel 22 may well be over estimated because of this.

Investigating the CI delay

In the third test the injection level was held at $\sim 17 \text{ ke}^-/\text{pixel}$ ($\sim 8.06\text{V}$) and the injection duration was set to 10 pixels. In total, five magnitudes were investigated (13.3, 15, 16.75, 18.5 and 20) covering the fainter end of Gaia's visual range. All the sets of reference data were taken with $G = 15$, so an additional correction had to be made to the flux calibration, as described above.

In this test the smallest injection delay was 10 pixels, which is very close to the 15-pixel window surrounding the peak of the image. Without removing the CI tail this made the curve fitting procedure unstable. As a proper analysis could not be performed, tests with an injection delay of 10 pixels were not considered.

In addition to those features discussed in section 2.3.3, a few more anomalous features appeared in the data, which for completeness are mentioned here. Note that these runs were excluded from the analysis:

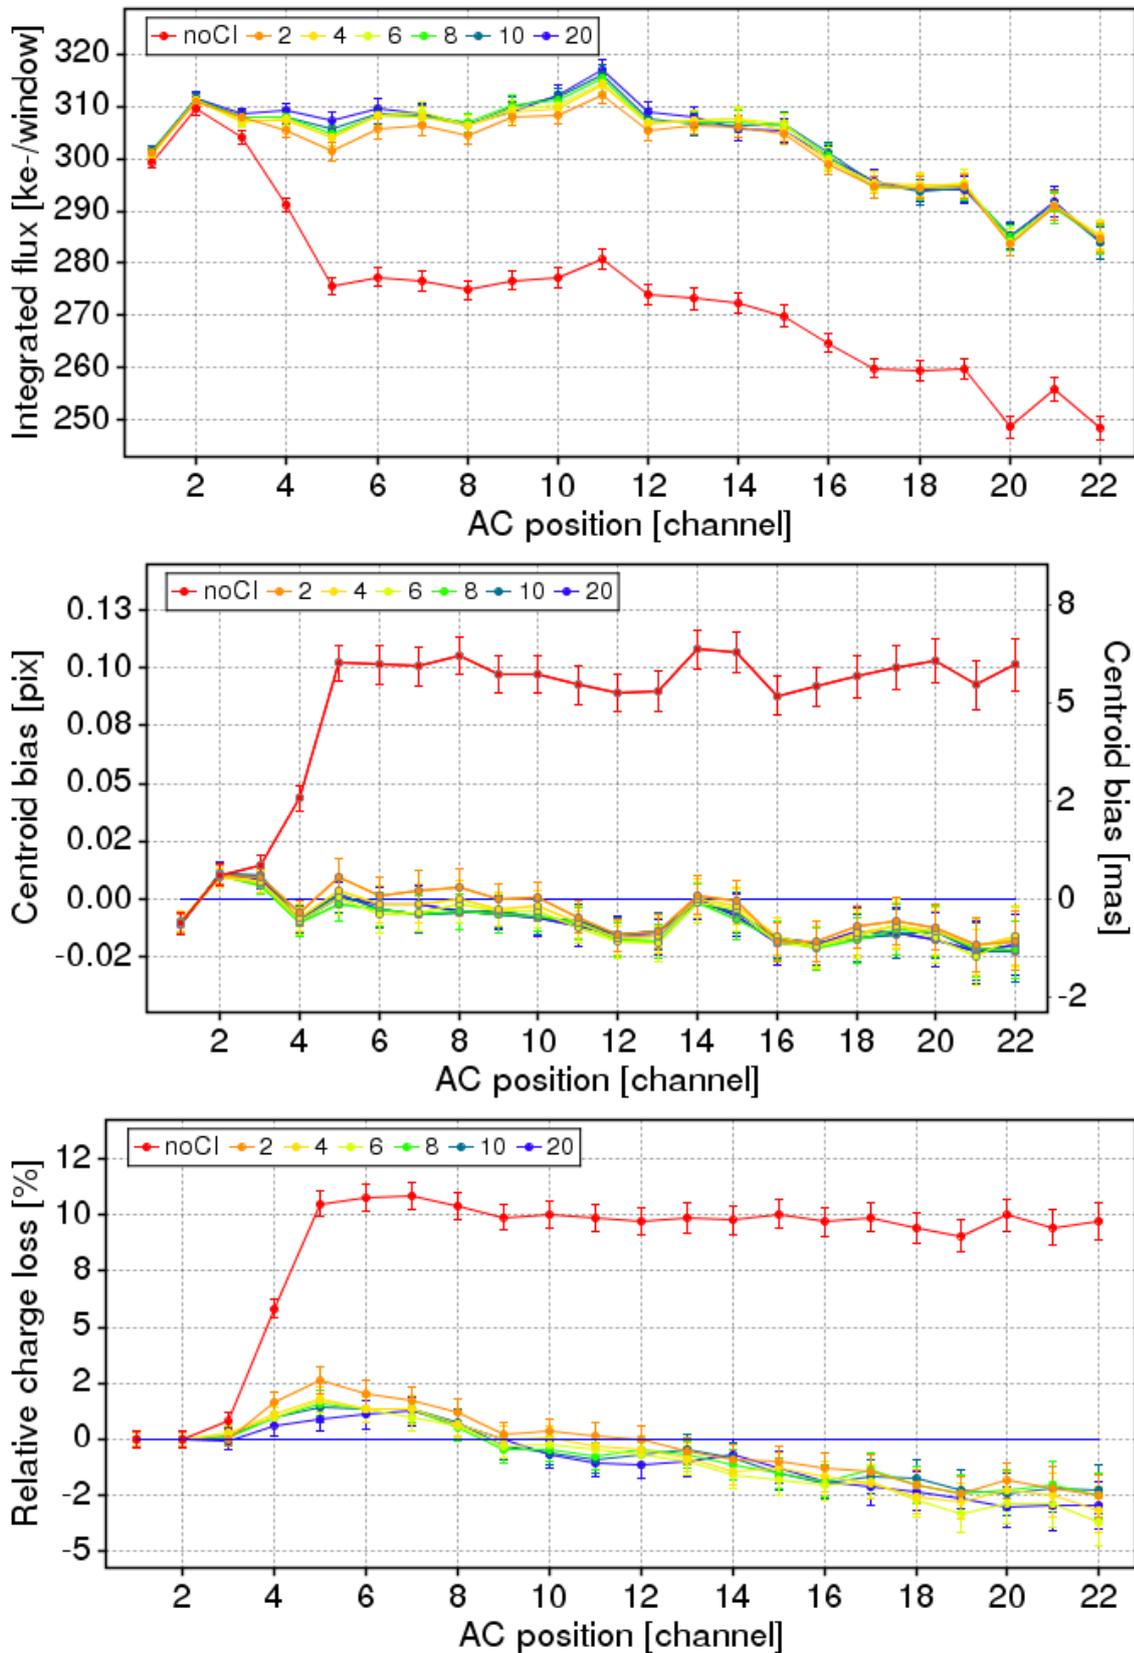


Figure 2.72: (top) Integrated flux ($G = 13.5$ to 13.8), (middle) relative centroid bias, and (bottom) relative charge loss as a function of AC position and CI duration in column 1 at a nominal $G = 13.3$.

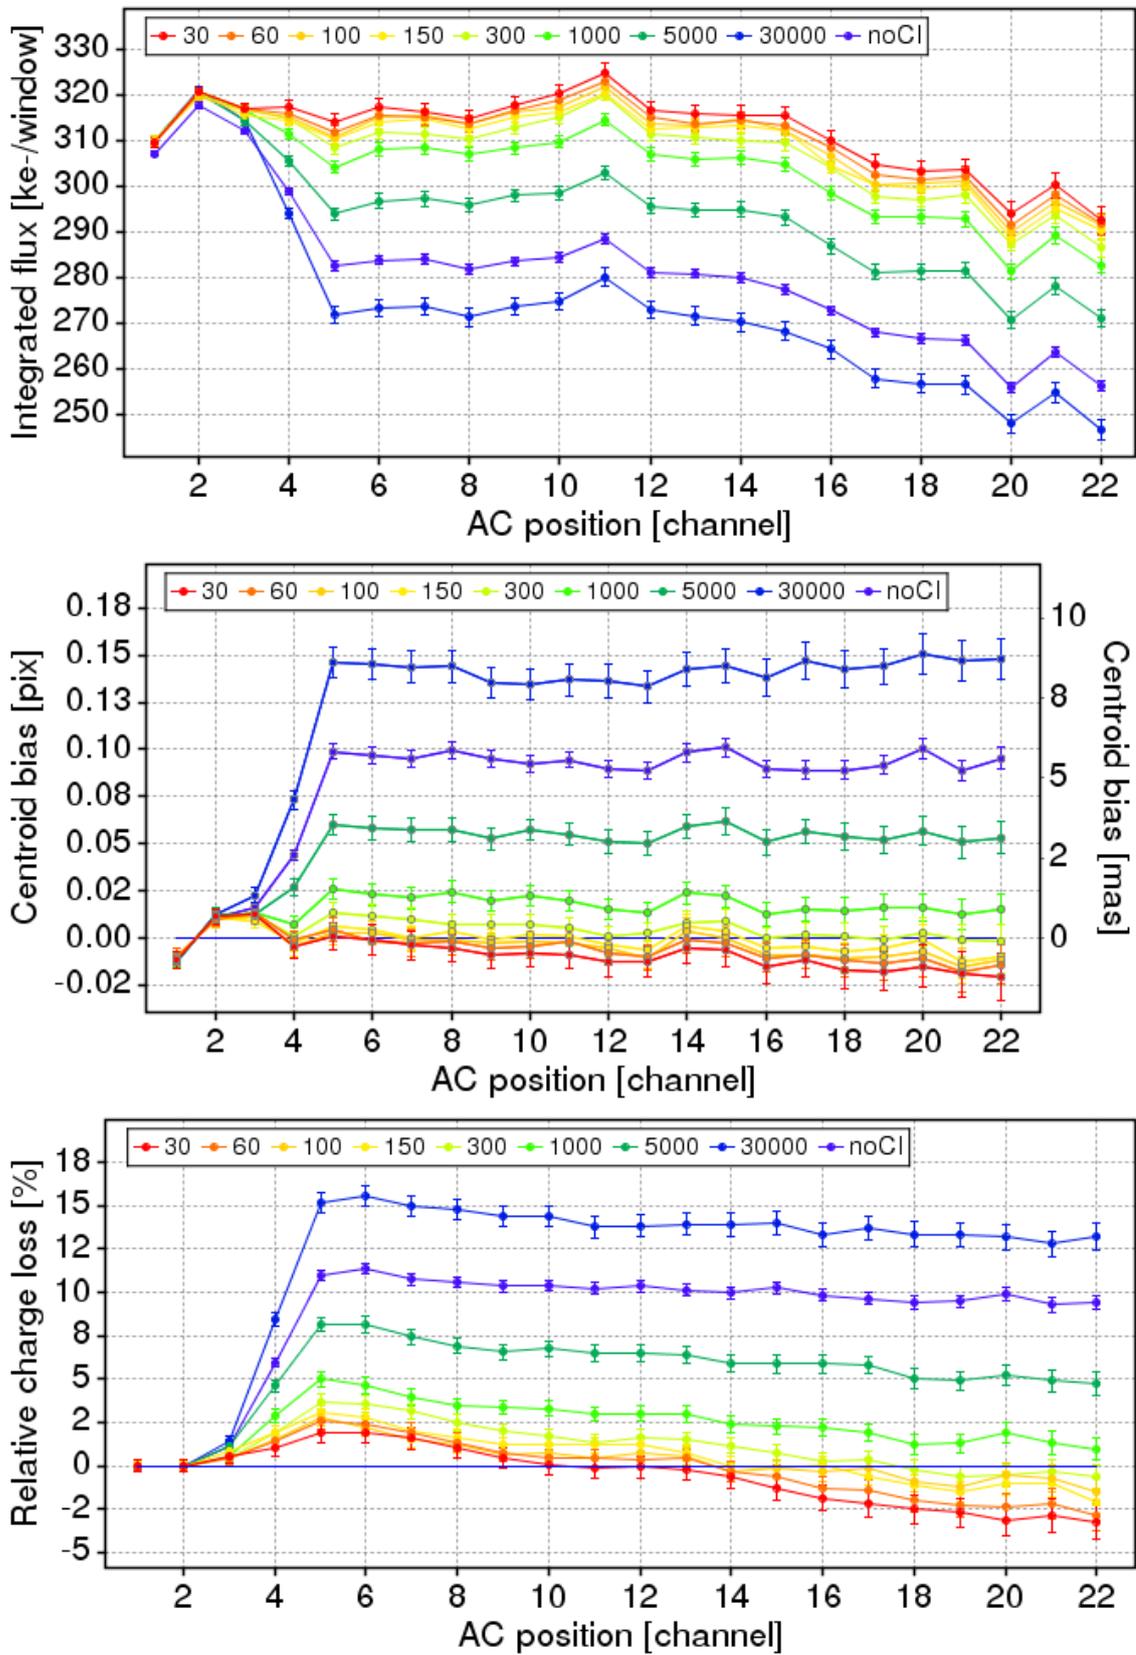


Figure 2.73: (top) Integrated flux ($G = 13.5$ to 13.9), (middle) relative centroid bias, and (bottom) relative charge loss as a function of AC position and CI delay in column 1 at a nominal $G = 13.3$.

- In scan 4 at $G = 15$, without injection (this is the second test case without injection in Table 2.10) a series of fainter images appear adjacent to the main stellar field. Presumably this was caused by interference from a secondary light source.
- In the same test case only 48 columns were read out of the CCD.
- At $G = 15$ with an injection delay of 100 pixels the CI was made after the third image had passed. Likewise with a delay of 150 pixels, the CI was made after the second image had passed.
- In some scans, e.g. $G = 15$, delay 30000, scan 4, the delay was far greater than expected. In this case, 34490 pixels before the stellar field appeared.
- In other scans, e.g. $G = 15$, delay 150, scan 7, the delay was far less. In this case, only 34 pixels before the stellar field appeared.
- In some scans, e.g. $G = 16.75$, delay 150, scan 4, the data does not contain any record of the test.

Figure 2.73 shows how the integrated flux (top), relative centroid bias (middle) and relative charge loss (bottom) vary as a function of CI delay at $G = 13.3$. Comparative plots for $G = 15$ are shown in Appendix B.2. For channels closest to the injecting source (i.e. towards channel 22) an injection delay of 100-150 pixels appears to remove the bulk of CTI from all channels. Beyond this delay radiation damage begins to appear in channels 3/4. Given the additional step required to calibrate the charge loss, it is difficult to make a clear recommendation on which CI is optimal from the bottom plot. If the centroid is recovered after ~ 150 pixels then presumably this should be reflected in the relative charge loss.

For the fainter images ($G > 16$) the noise on the centroid bias hides any clear underlying trend, so Figure 2.74 only shows the charge loss as a function of CI delay. The CI tail has a significant impact on the shape of the image and the charge underneath it. By not removing the CI tail this additional background dominates the signal. When $G = 16.75$ and 18.5 the charge losses are distinctly positive when the CI delay is less than 300 pixels. When $G = 20$ the same feature can be seen when the CI delay is less than 1000 pixels. Clearly there is a need to remove the CI tail to fully recover CTI image parameters, especially when $G > 15$.

2.3.6 Relative centroid bias and charge loss in the leading column

In this section we compare our results to those obtained by Astrium. Since Astrium took the average bias over channels 6 to 17, we also do the same. Note though that this includes the systematic biases observed in channels 14/15 attributed to a CCD stitch. For completeness we show the results for all magnitudes.

Figure 2.75 shows the centroid bias (left) and charge loss (right) as a function of CI level for $G \lesssim 15$. Two test cases were performed without CI at $G = 13.3$, so an average of the two CTI measurements is shown here. The CI level without CI has been given an arbitrary value of 1000. Although Astrium used a local background cut, both results tend to agree within the final errors. At $G = 13.3$, a CI level between $10\text{ke}^-/\text{pixel}$ ($\sim 8.26\text{V}$) and $20\text{ke}^-/\text{pixel}$ ($\sim 8.06\text{V}$) would appear to be sufficient to remove the

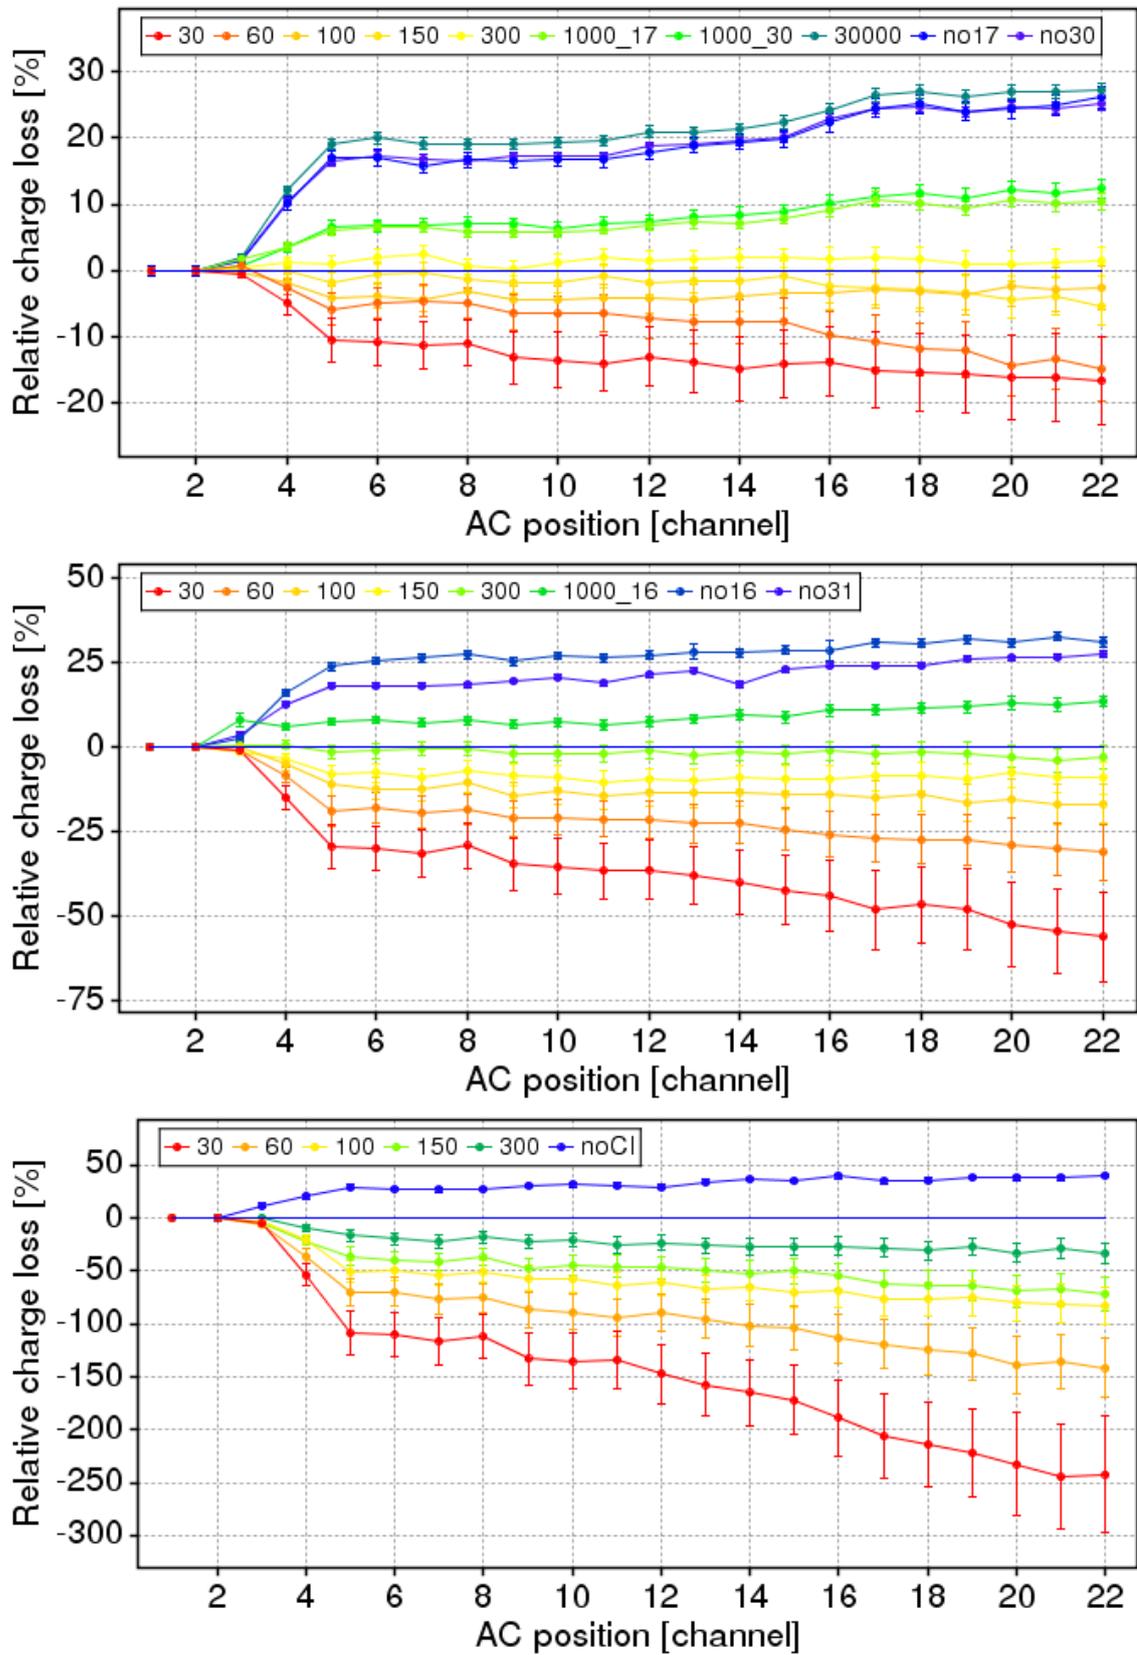


Figure 2.74: Relative charge loss at (top) $G = 16.75$, (middle) $G = 18.5$, and (bottom) $G = 20$ as a function of AC position and CI delay in column 1. Tests without injection: 'no17' performed on 17/08/2007, 'no30' performed on 30/08/2007, 'no16' performed on 18/08/2007, 'no31' performed on 31/08/2007.

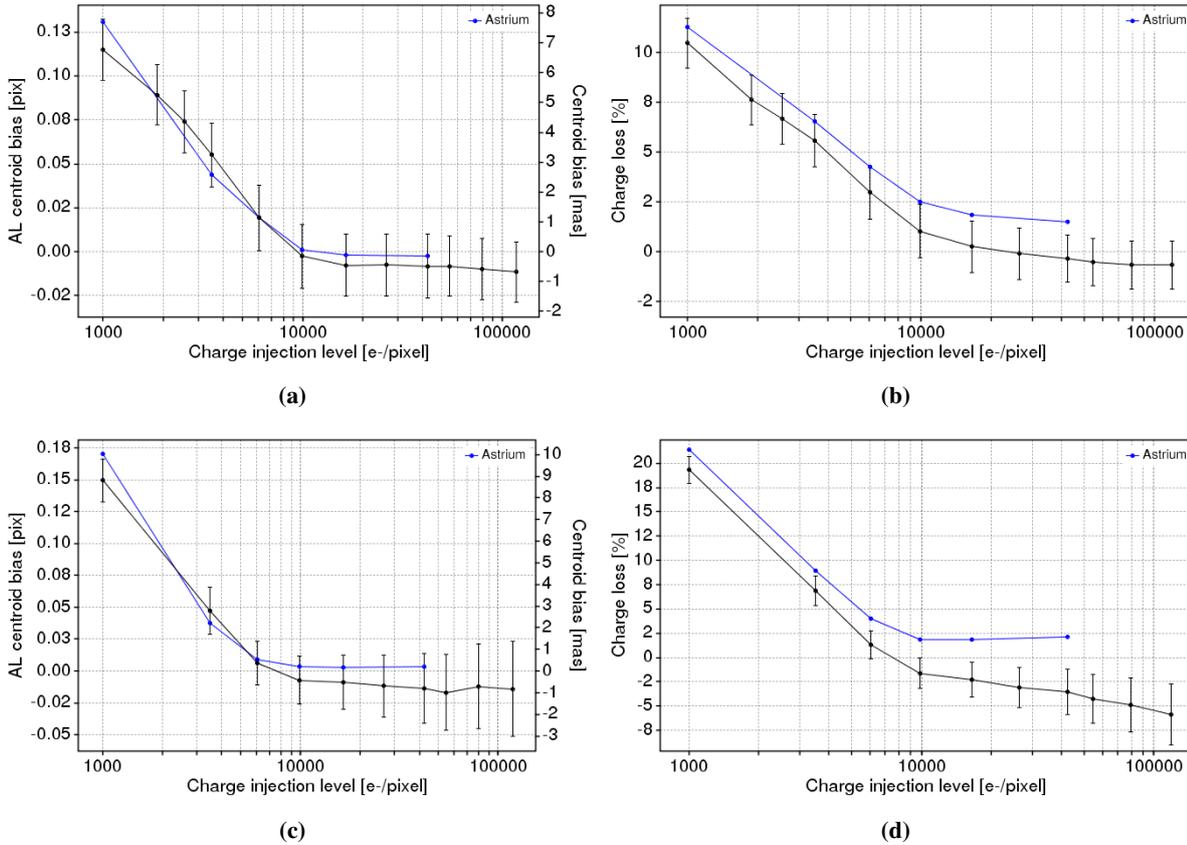


Figure 2.75: Relative centroid bias and charge loss as a function of CI level for (a/b) $G = 13.3$ and (c/d) $G = 15$.

bulk of radiation damage. At $G = 15$ a CI level between $6ke^-/\text{pixel}$ ($\sim 8.42V$) and $10ke^-/\text{pixel}$ ($\sim 8.26V$) appears sufficient to remove CTI effects.

In the second test the centroid bias and charge loss at $G = 13.3$ and 15 as a function of CI duration are shown in Figure 2.76. In both cases the CI duration needs to be only 2 pixels, preferably 4, to remove the bulk of radiation damage.

Figures 2.77 and 2.78 show the effects of radiation damage with various CI delays when $G \lesssim 15$ and $G \gtrsim 15$ respectively. The average result is shown when the same test case has been repeated. For test cases without CI the delay has been given an arbitrarily large value of 1×10^6 pixels. It may seem surprising that the biases and charge losses incurred without CI are less than those seen with the longest delays. However, it should be remembered that the time between scans is typically 27s, which is equivalent to an approximate (self-injection) delay of 27500 pixels. At $G = 13.3$ and 15 it appears any CI delay ≥ 100 pixels would begin to incur increasing amounts of radiation damage. For fainter images it is clear the two methods of analysis are diverging, becoming more pronounced as the magnitude increases. As mentioned earlier, Astrium used the local background method to subtract the background from the stellar field. This method assumes the wings of the image sample the background, so it can partially take account of the CI tail when the image is faint enough. This is probably why Astrium's results do not diverge as the CI delay is reduced. If we assume that any charge gain is associated with the CI tail, Figure 2.78 suggests a

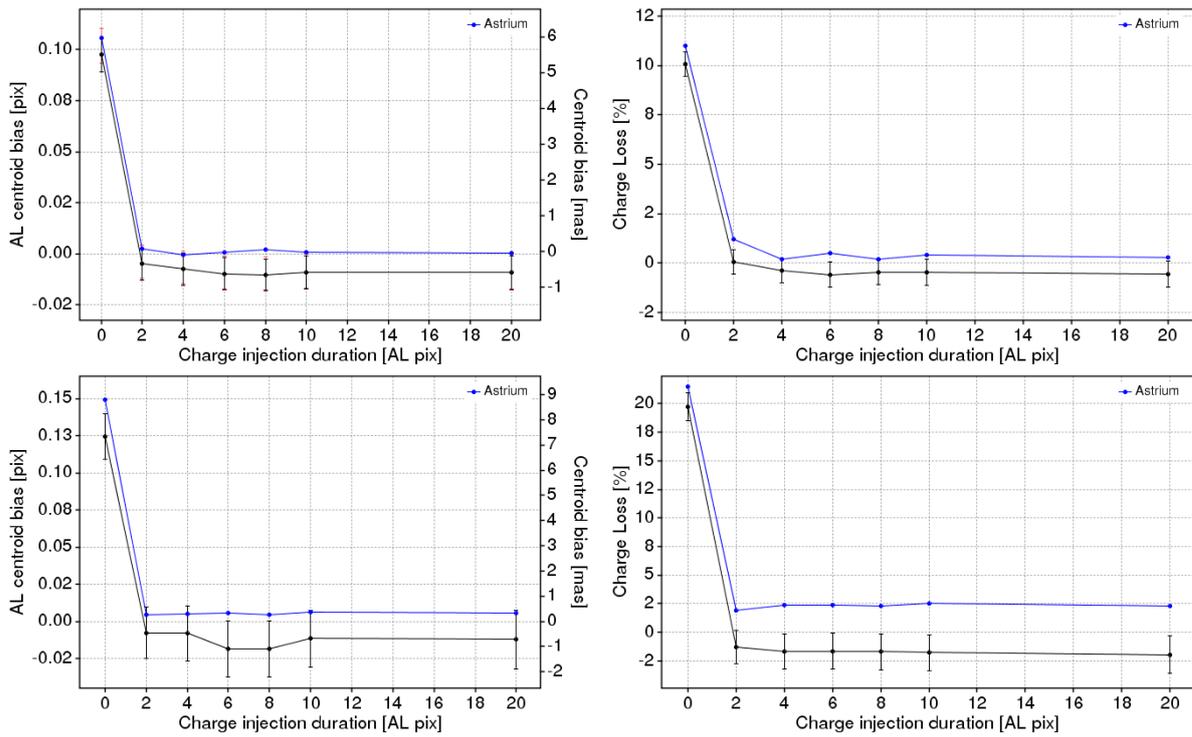


Figure 2.76: Relative centroid bias and charge loss as a function of CI duration for (a/b) $G = 13.3$ and (c/d) $G = 15$.

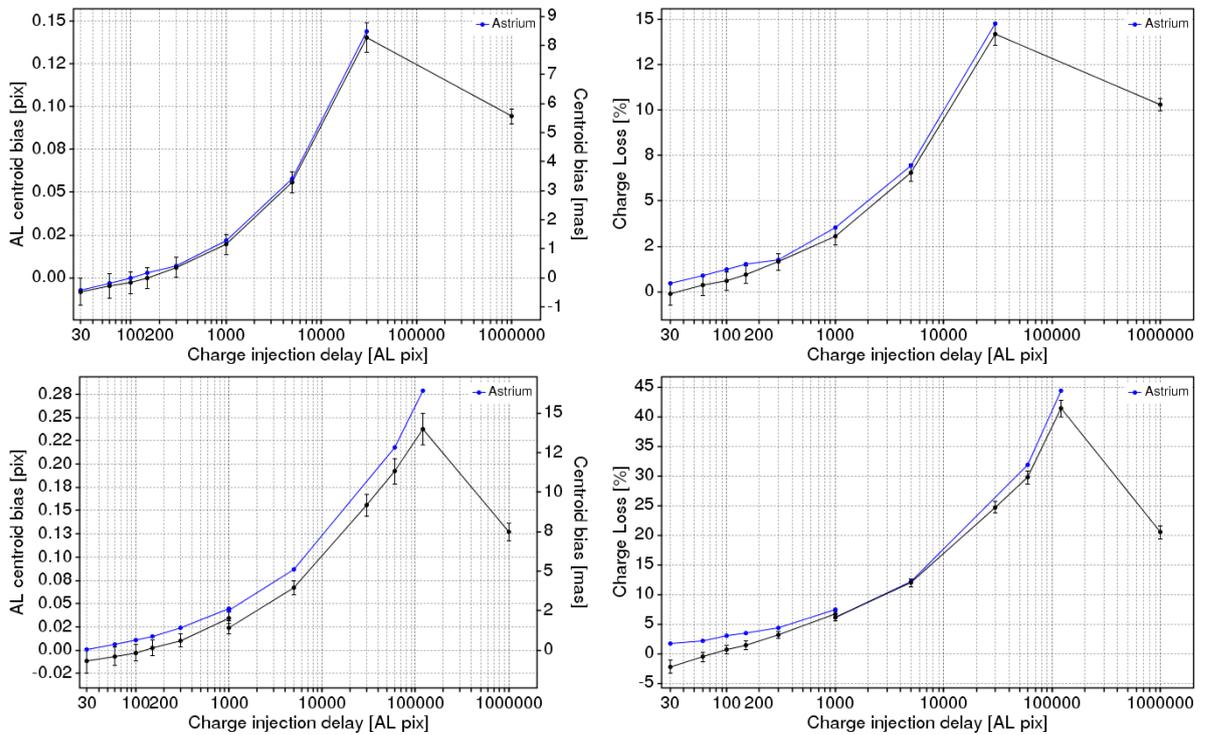


Figure 2.77: Relative centroid bias (left) and charge loss (right) as a function of CI delay for (top) $G = 13.3$ and (bottom) $G = 15$.

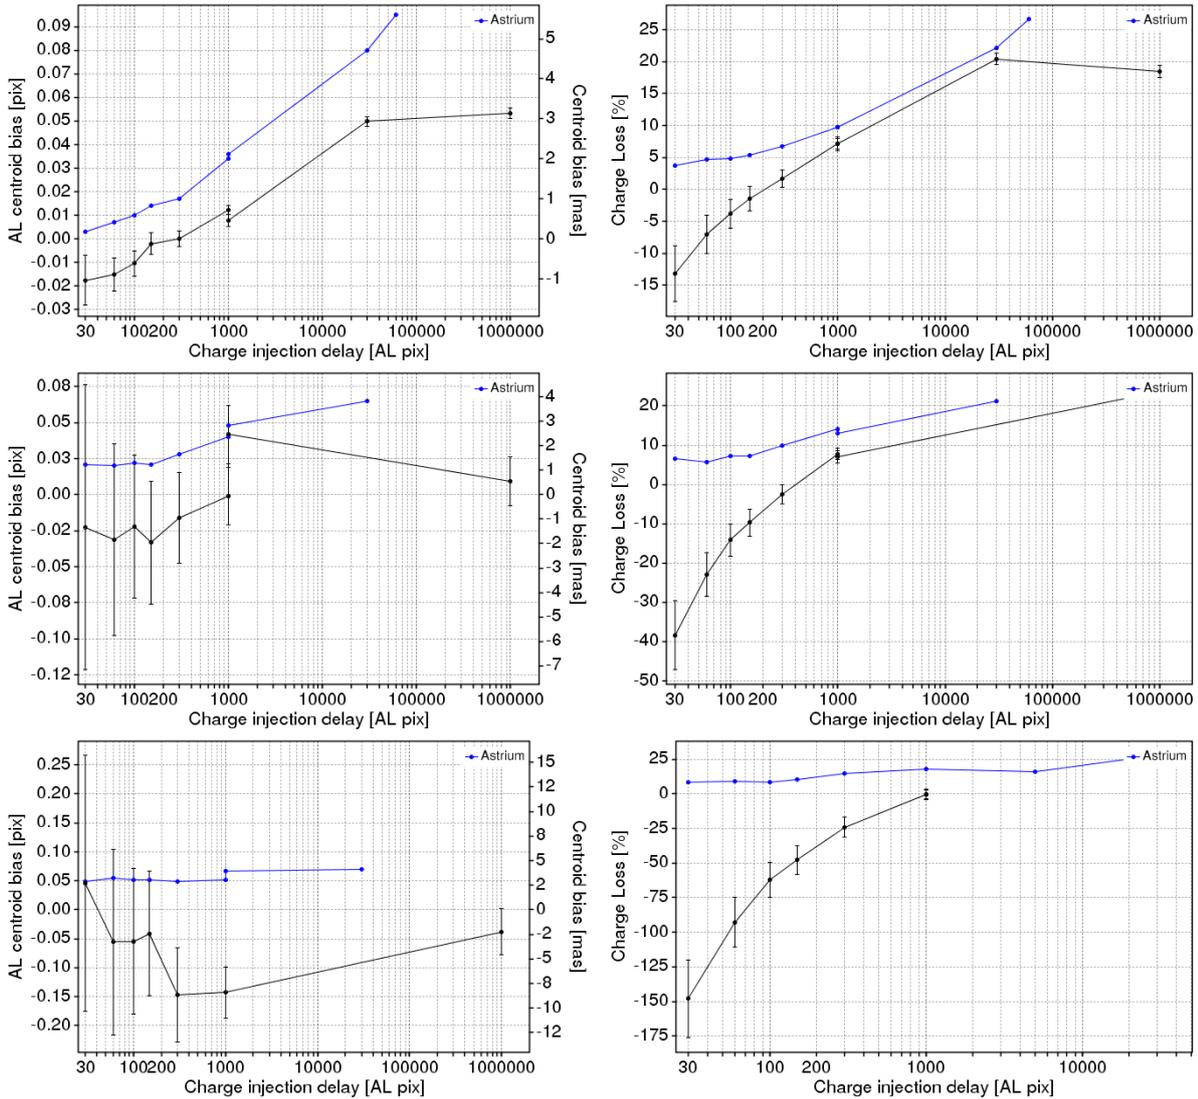


Figure 2.78: Relative centroid bias (left) and charge loss (right) as a function of CI level for (top) $G = 16.75$ and (middle) $G = 18.5$ and (bottom) $G = 20$.

CI delay of <1000 pixels should suffice to remove the bulk of radiation damage from the faintest images.

Tables B.2, B.3 and B.4 in Appendix B.3 tabulate the results shown here.

Image profiles: undamaged vs damaged

In this section the shape of the LSF, having passed over the undamaged part of the CCD, is compared directly to the same image shape after passing over the damaged part of the CCD. The centroid and flux calibrations performed earlier are used to position and scale each image. Since the CI is made from only one end of the CCD the images shown here are taken from channel 22 in column 1, the image closest to the CI source.

First test In Figure 2.79 the undamaged and damaged local LSFs found in column 1, channel 22 at $G = 13.3$ (left) and $G = 15$ (right) as a function of CI level are shown together with the damaged data and residuals between the two fits. Recall that the AL motion is to the left.

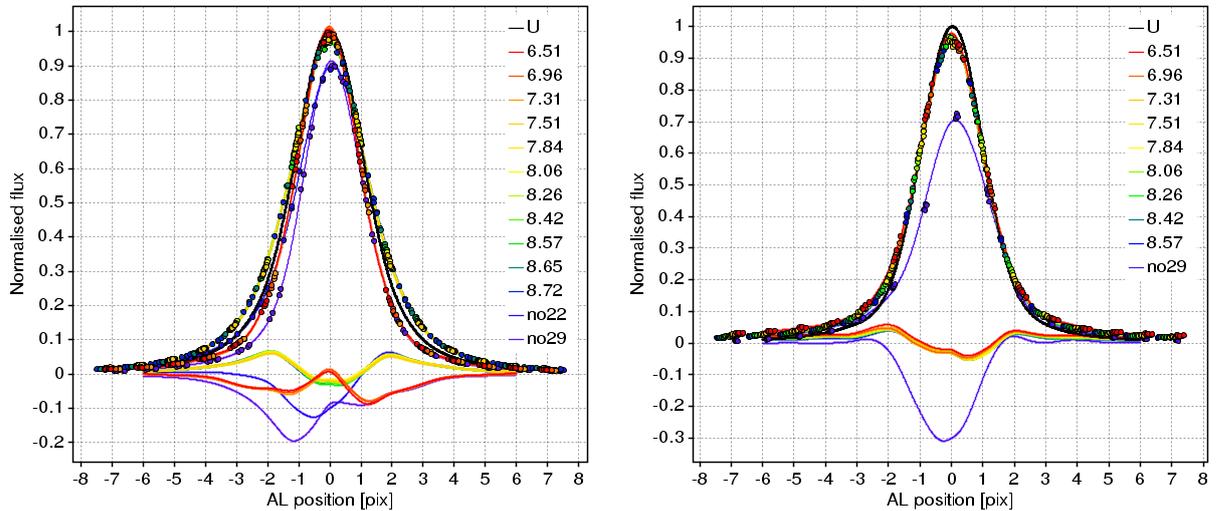


Figure 2.79: Image distortion at (left) $G = 13.3$ and (right) $G = 15$ as a function of CI level. Two test cases were performed without injection: one on the 22/06/07 (referred to as 'no22' in the left plot) and the other on the 29/06/07 (referred to as 'no29' in both plots). U = undamaged curve.

At $G = 13.3$:

- Firstly, without CI the leading edge and peak of the image are eroded as expected, while the trailing edge is left intact.
- Secondly, with very high CI levels the damaged image appears to be narrower than the undamaged image, while the opposite is true for very low CI levels. This is reflected in the residuals. This cannot be caused by the CI tail as the trailing edge is also narrower/wider than the respective undamaged image. We know the equipment was not calibrated as well as it could have been in the early tests in RC2, so these differences are probably due to variations in the intensity of the light source.
- Even though the charge injection tail was not removed from this analysis, and indeed must still be taken into account, the additional charge underneath the image is not evident from this figure at this magnitude.

At $G = 15$:

- Without CI, even though the datapoints are few, the leading edge and peak of the image are again preferentially eroded, while the trailing edge is left intact. As expected, compared to $G = 13.3$, the damage is worse.
- Secondly, with CI the background in the damaged image is slightly higher than expected. This is probably caused by the CI tail.

Second test Figure 2.80 shows comparable results for $G = 13.3$ (left) and $G = 15$ (right) as a function of CI duration. Here the damaged image is almost completely recovered after a 2 pixel CI duration. Increasing this to 10 pixels recovers the peak of the image at $G = 15$.

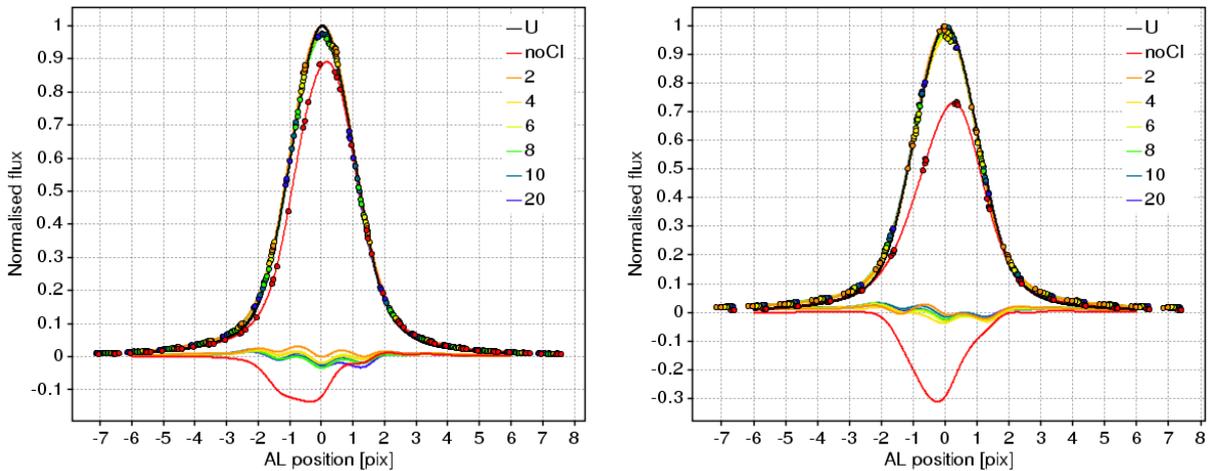


Figure 2.80: Image distortion at $G = 13.3$ (left) and $G = 15$ (right) as a function of CI duration. U = undamaged curve.

Third test Figure 2.81 shows how the CI delay can impact on the shape of an image at various magnitudes. At $G = 13.3$ the leading edge is slightly eroded and requires a CI around 100 pixels before the image to recover the lost charge. At $G = 15$ the bulk of the leading edge is also recovered with a delay of 100 pixels, although the peak remains slightly damaged regardless of the CI delay.

For $G \gtrsim 16.75$, firstly, the trailing edge begins to erode as well as the leading edge. Secondly, the background added by the CI tail is clearly having an impact, and thirdly, it is also clear how in some test cases the samples cluster around the same sub-pixel position, which restricts a reliable curve fit.

2.3.7 Conclusions

In RC2 charge injection has been investigated as a means to mitigate radiation damage. The CI level, duration and delay has been varied to quantify CTI image parameters and find a suitable set of CI parameters that can mitigate radiation damage over a range of signal intensities. To begin with, an analysis of the raw data has identified a number of features that were due to processes occurring both on and off the CCD:

- Light leaking around the mask via screws that were holding the mask in place accounts for the systematic bumps seen either side of the stellar field; their intensity correlates with the variation in test magnitude.
- The light tail seen at the beginning of each scan after scan 1 has been made is caused by the light source being kept on as it is passed back over the CCD.

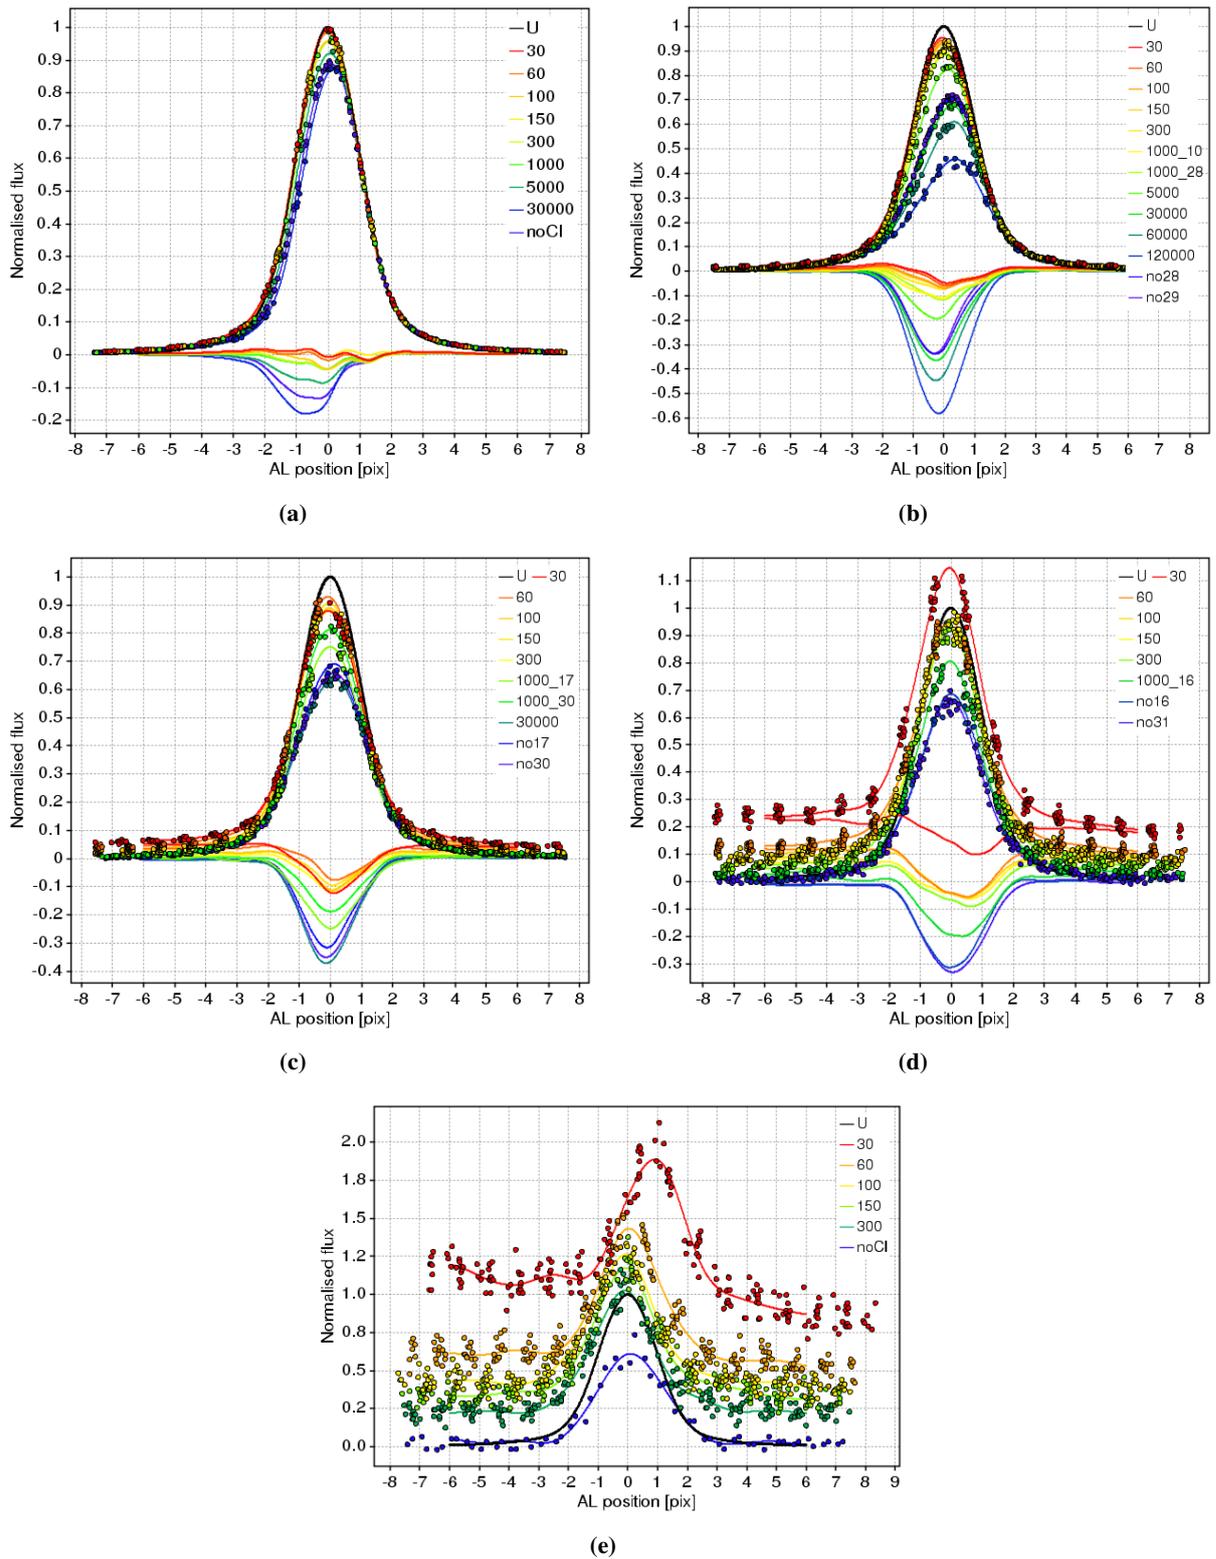


Figure 2.81: Image distortion at (a) $G = 13.3$, (b) $G = 15$, (c) $G = 16.75$, (d) $G = 18.5$, and (e) $G = 20$ as a function of CI delay. At $G = 20$ the residuals are not shown as clearly they are governed by the CI tail and not radiation damage. Multiple test cases were performed with a CI delay of 1000 and without CI. The label in the legend in various plots indicate the date the test was performed, e.g. noX and 1000_X was performed on X/08/2007.

- There are many instances where cosmics are seen that intermittently interfere with the CCD read-out.
- The majority of other anomalous features (e.g. the behaviour of scan 1, the AC variation in the electronic offset) can be explained by the CCD-PEM couple instability.

All of these systematics had to be avoided in order to estimate the background and extract the stellar field. The PEM anomaly has now been flagged alongside radiation damage as a critical risk in the Gaia mission. A dedicated Offset Instability Task Force (OITF), analogous to the RTF, has been set up to address this very issue.

The background has been estimated by averaging over a 1000-pixel scan interval prior to the CI being made and extending this estimate as a flat cut over the entire scan. With hindsight, this does not remove the CI tail, which has a significant impact on the fainter images. A local background cut may be used to analyse the fainter images, though it is not known if this method could properly remove the CI tail. A local background cut cannot be used to analyse the brighter images ($G = 13.3$ and 15) as the background is never reached locally. Choosing a proper means to treat the background is a vital step in the data reduction process and one that must be taken with great care. As we have shown here, the charge injection tail clearly imparts a strong signature on the charge loss and complicates the centroid bias measurement. In RC3 the CI tail will be extracted explicitly using a model developed by DU10.

In the very first test in RC2, i.e. the investigation into the CI level performed at $G = 13.3$, the DOB level was far higher than expected (~ 12 e^- /pixel). This was rectified in the tests that followed with the DOB level hovering around ~ 1 e^- /pixel. In this first test there may also have been an instability in the light source, which may explain some of the variations in the widths between the undamaged and damaged image profiles.

Scan 1 was not included due to: a) anomalous features in the background preceding the stellar field; and, b) the different radiation conditions experienced on the CCD. This left only 4 scans to analyse CTI effects at $G = 13.3$ and 15 . After combining images from each scan it was found that: a) the lack of data limited the coverage of the LSF; and, b) the data tended to cluster around the same sub-pixel positions. Comparisons between the fitted GRC and a local LSF showed that, interestingly, the fitted GRC compensated for differences between the two by scaling the curve height; if the local LSF was wider, the fitted GRC was taller and vice versa. In doing so, the integrated fluxes under each curve were very close to each other ($< 1\%$) and could be used to calibrate the charge loss.

The original calibration of the geometry of the mask produced centroid positions that could not be explained by any physical process. Assuming a certain set of centroid positions recorded further up the mask were undamaged in the irradiated data, a recalibration was possible, which led to the discovery of CCD stitches. CCD stitches have a significant impact on the centroid bias and can systematically offset the centroid by ~ 0.01 - 0.02 pixels. Unfortunately, one of these stitches coincided with the location of channel 2 on the mask so that the intended calibration could not be performed. In RC3 the number of non-irradiated channels in the irradiated tests are increased and their locations chosen to avoid CCD stitches.

Results on the relative charge loss have shown that there may be a preferential loss of charge running along a stitch. If the stitch introduces a potential pocket into which electrons may fall and become trapped, then the stitch itself could be another source of charge loss.

When the CI is too weak, the bulk of the damage is observed in the leading column of stars. This sets up a shielding effect which hides a significant number of traps from all of the stars that follow. As the CI was driven from only one end of the CCD, this induced a degree of variability in the level of damage experienced AC as well as AL. To rectify this inhomogeneity, in RC3 the CI is driven from both ends of the CCD.

Finally, regarding the optimum parameters a CI should have to mitigate radiation damage on board Gaia, and taking into account the CI tail not being removed, the following can be said:

- *Injection level.* An injection close to 8.06V or ~ 16500 e⁻/pixel appears to mitigate the bulk of radiation damage seen in the first test at G = 13.3 and 15. This agrees with Astrium's choice to use an injection level of ~ 17000 e⁻/pixel in the second and third tests.
- *Injection duration.* Although an injection duration of only 2 pixels appears to mitigate the bulk of radiation damage observed at G = 13.3 and 15, an injection duration of 4 pixels is preferred. This compensates for the injection becoming weaker further away from the injecting source and keeps dead-time (i.e. the time over which the charge injection is taking place) as small as possible. Driving the injection from both ends of the CCD should keep the injection duration close to this value.
- *Injection delay.* Considering how the relative charge loss varies as a function of CI delay, it appears an injection delay of <1000 pixels would be needed to remove radiation damage from the faintest images (G = 20) and an injection delay of <100 pixels would be required to remove the bulk of radiation damage from the brightest images (G = 13.3). For all magnitudes in between, it is expected that some intermediate delay would be required. Whether a charge injection delay of 1000 should actually be adopted for the mission at this stage remains unclear as the sky density in this analysis is not indicative of what Gaia will see on average. This will have to be addressed further in campaigns 3 and 4.

Recommendations made to industry and follow-up

After analysing RC2 the following recommendations were made to industry to improve the output from future campaigns:

1. Perform a greater number of scans, preferably 20, in each of the irradiated test cases to increase the coverage of the LSF. This request becomes obsolete, however, if the next request is not followed.
2. Dither the time between scans to avoid clustering samples around the same sub-pixel positions. This counters the need for periodic injection but is important for pre-launch analysis.

3. Collect reference data at the same magnitude as that used in the irradiated tests. This is to avoid the additional step that must be included to calibrate the flux response in the irradiated tests.
4. Non-irradiated images must not traverse the locations of CCD stitches. This can be compensated for by increasing the number of non-irradiated channels in the irradiated tests.
5. Since scan 1 does not experience any memory effects it would be useful to probe these conditions further by increasing the length of time between scans. However, it is understood that time commitments to carry out such tests have become prohibitive and will not be performed.

These findings were presented to the Radiation Calibration Working Group on 16th September 2008 (ESA Gaia Project (GAIA-EST-MN-06347), Meeting #7) and presentations on the analysis were given to the Radiation Task Force on Monday 6th October 2008 and Monday 6th April 2009. As mentioned in section 2.2.6, the astrometric regular tests for RC3 had already been performed by September 2008, which could not have been changed, although some feedback from this analysis had already been fed back to Astrium before this date (this may have resulted in the bright star tests having reference data at the same magnitude as the irradiated data). The bulk of RC3 had yet to be performed in September 2008 so many of the recommendations made here could be applied. Continuing with the analysis of the AF, a number of sky-like tests were performed in RC3, which have begun to be analysed by DPAC. Feedback on the recommendations made above state the following:

1. 60 or 65 scans have been acquired for all 3 illumination levels, i.e. at nominal magnitudes for the reference stars of $G = 15, 17$ and 19 .
2. The time between scans has been dithered such that each scan, after co-adding images, samples the LSF after ~ 0.2 pixels. The setup appears also to be quite stable.
3. Reference data has been collected at the same magnitude as the irradiated data for all tests.
4. To compensate for CCD stitching there are 5 rows of stars on either side of the mask to act as reference stars.
5. The time between scans is similar to previous campaigns. However, in most cases there is a CI in most of the tests.

Clearly, the independent analysis performed here has contributed to Astrium's own analysis and fed back into ways in which the campaigns could be improved.

Another improvement that did not feature in previous campaigns are the inclusion of CI calibration runs in the sky-like tests. All of the irradiated data in this set has a corresponding CI calibration run, composed of the CI without the stellar field, which makes it possible to effectively subtract the injection tail, CI forward-spilling, the PEM non-uniformity offset, hot pixel offsets and sky background all in one go. Furthermore, there are as many CI calibration runs as corresponding irradiated runs, although the unirradiated sky-like data (see the next section) lacked these runs.

2.4 Radiation Campaign 3

RC3 began by investigating the test bench itself (see EADS Astrium (GAIA.ASF.PLN.PLM.00052) for the RC3 test plan). Though this dataset is not analysed here, Astrium report on the bench characterisation tests (EADS Astrium (GAIA.ASF.TCN.PLM.00311)). Amongst other things, they investigated the effects of CCD stitches and the impact of the background removal algorithm. They state that "the worst-case stitching error in the AL direction is of 0.05 AL pixel, which can yield an error in pitch between the two adjacent lines of 0.025 AL pixels". Though this is in fact slightly worse than our findings, clearly stitches must be dealt with.

The next round of tests in RC3 involved using the same mask as in RC2, but with a new CCD irradiation scheme. These are the *Astrometric Regular Tests* performed by Astrium. This new scheme allowed an improved calibration strategy to be tried out as well as a reinvestigation of the effects of a DOB and CI. In the following sections an analysis of a subset of this dataset is presented, concentrating on the brightest and faintest stars that were tested. Astrium's analysis of these tests can be found in EADS Astrium (GAIA.ASF.TCN.PLM.00447). See also EADS Astrium (GAIA.ASF.TCN.PLM.00349) for a description of the datapack.

In addition to the astrometric regular tests, a new sky-like optical mask has been used to gather data on a stellar field that is more representative of the sky that Gaia may well observe. This mask incorporates two regular grids placed either side of the sky-like field to calibrate the mask-CCD system and extract CTI image parameters in the same vein as that performed in RC2. Indeed, RC2 was conducted, in part, to prepare for the sky-like field tests. Some aspects of these tests will be discussed here, but for a more detailed study currently in progress see Lindegren (LL-082). Astrium's analysis can be found in EADS Astrium (GAIA.ASF.TCN.PLM.00497). The datapack description is in EADS Astrium (GAIA.ASF.TCN.PLM.00415).

2.4.1 Hardware and test configuration

An overall description of the hardware and test configuration used in the RC3 astrometric regular tests can be found in EADS Astrium (GAIA.ASF.TCN.PLM.00447). In particular Figure 3-2 in this document shows the AF CCD irradiation scheme and mask positions in both the reference and irradiated tests. The gain in RC3 was set to $3.226 \text{ e}^-/\text{LSB}$.

The most important changes in these tests compared to RC2 are: a) having 3 non-irradiated channels either side of the stellar field to calibrate the irradiated data; and, b) having the charge injection (CI) driven from both ends of the CCD, stabilising the CI level across every channel. These concepts are illustrated in Figure 2.82.

The DOB, CI level, CI delay, and the use of gates were all investigated at various magnitudes between 11.1 and 20 in the astrometric regular tests. However, there was not enough time to analyse all of the available data. Instead, particular subsets were chosen that could cover as much of the testing as possible. For example, the CI level tests performed initially are not covered, since the CI level is always set to $20 \text{ ke}^-/\text{pixel}$ in every test thereafter, and for completeness the CI duration is always set to 20 pixels. Other

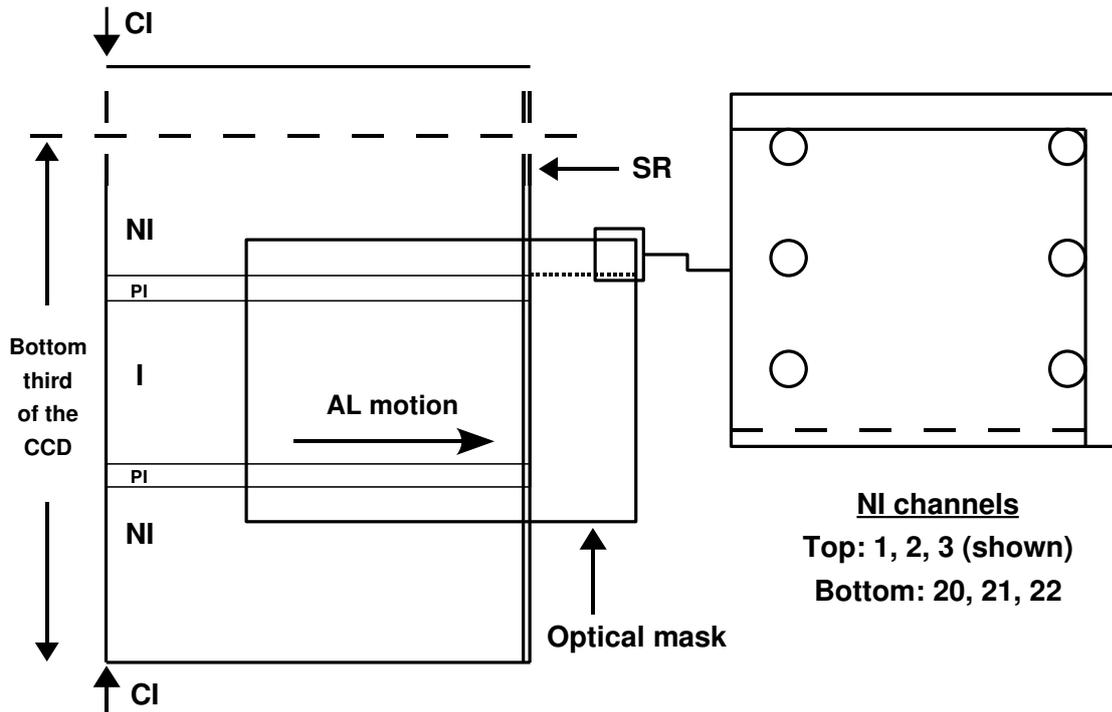


Figure 2.82: Details on the test setup in RC3. This particular configuration shows the mask passing over the CCD during the irradiated tests. NI = non-irradiated, PI = partially irradiated, I = irradiated.

DOB: 2, 4, 7 e ⁻ /pixel					
Short CI delays			Long CI delays		
G	No. scans	Injection Delay [pix]	G	No. scans	Injection Delay [pix]
15	10	Reference	15	10	Reference
13.6	5	No inj., 100, 500, 1000	15	10	No inj., 1000, 5k, 10k
15	10	Reference	15	10	Reference
20	25	No inj., 100, 500, 1000	20	25	No inj., 1000, 5k, 10k

Table 2.14: Information on the medium/faint star tests.

DOB: 0, 2, 4 e ⁻ /pixel					
Gate 11		Gate 10		Gate 9	
G	Inj. delay [pix]	G	Inj. delay [pix]	G	Inj. delay [pix]
12.6	Reference	11.9	Reference	11.1	Reference
12.6	100, 500, 1k, 5k, 10k	11.9	100, 500, 1k, 5k, 10k	11.1	100, 500, 1k, 5k, 10k

Table 2.15: Information on the bright star tests. 10 scans were performed in each test case.

tests are split into two groups: those focusing on medium/faint (MF) stars ($G > 13$) and those focusing on the bright stars ($G < 13$). In both groups, both the DOB and CI delay have been varied. Tables 2.14 and 2.15 summarise the MF star tests and the bright star tests respectively. Note that in the MF tests the reference data is always taken at $G = 15$, regardless of the irradiated test magnitude. Fortunately, this is not the case with the bright star tests.

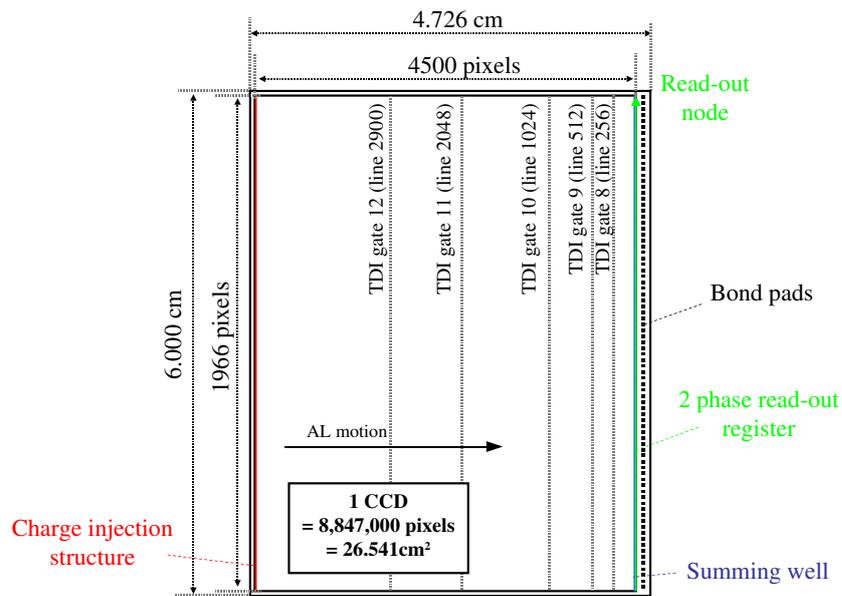


Figure 2.83: A schematic of an AF CCD showing the AL positions of some of the gates. Adapted from the Picture of the Week: Gaia Focal Plane (See ESA (2010b)).

For tests with bright stars, certain gates have been activated. Gates come into play when a signal is too bright to cross the CCD without saturating the readout. In every pixel there is an embedded structure called the antiblooming drain that stops charge from saturating the pixel. In 12 locations along an AF CCD the image electrodes which clock charge across the CCD are separated from the rest of the image area (IA). These are connected to drains or gates that can be clocked to remain at a relatively low voltage to stop charge passing through to the next pixel. Once a gate is activated, any excess charge that has been generated can be held in the gate for a given time so that the image may carry on integrating across the CCD as if it were passing over it for the first time. After a given time the excess charge is clocked across the CCD along with the rest of the field. As the excess charge is held back it spreads out across a CCD pixel column in the gate structure, behaving therefore somewhat analogously to a charge injection. Activating a gate reduces the effective length of a stellar transit to the distance between the gate and the serial register (SR). Figure 2.83 illustrates the positions of the some of the gates on a Gaia AF CCD. There are in fact 12 gates, the other 7 are situated at pixel lines 2, 4, 8, 16, 32, 64 and 128 respectively.

2.4.2 Background analysis

As in RC1 and RC2 there are some features in the raw data that are not expected. Figure 2.84 shows a running mean (see eqn. 2.3) of the reference data in the first of the MF test cases with a DOB of 2 in channel 1. The two signals seen either side of the stellar field are thought to be due to light leaking around the mask. Their intensity correlates with the magnitude of the source to support this assertion. These leaks were also seen in RC2, but more spread out over the scan, appearing as bumps in the background. If both originate from the same leak, the peak positions of the resulting pollution have shifted from ~ 1800 pixels to ~ 100 pixels away from either side of the stellar field. The background itself appears stable from

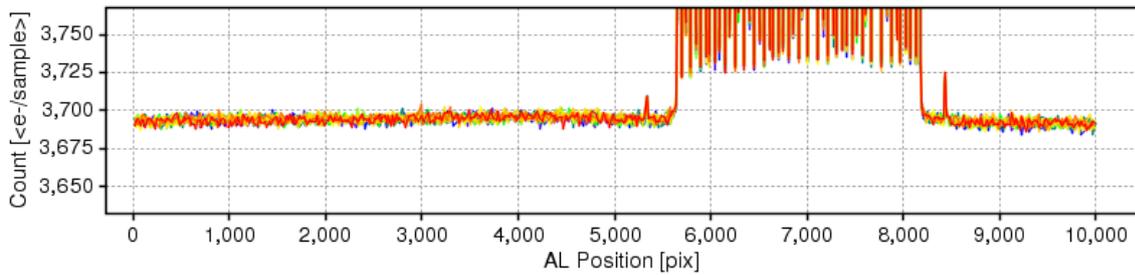


Figure 2.84: Running mean of all 10 scans in channel 1 from the reference data in the first ($G = 13.6$) test with a DOB of 2. Note that the reference data is actually at $G = 15$.

one scan to the next but does contain a slight variation, rising and falling prior to the stellar field, peaking around 4500 pixels.

Figure 2.85 shows comparable results for an irradiated test case taken at $G = 13.6$ with a DOB of 2 and a CI delay of 1000. The sharp decrease seen at the beginning of each scan was also seen in RC2 and is attributed to the PEM. Astrium state that it should not interfere with the stellar field as it is an effect added outside the IA. After this the background signal appears to gradually rise and fall in every scan prior to the stellar field, akin to what is seen in Figure 2.84. However, here the amplitude of the rise is greater. Since this data is taken at $G = 13.6$ rather than $G = 15$, it could be that this "bump" in the background is the same as the bump seen in RC2. In which case, the gradual variation in the background could also be due to light leaking around the mask. At first glance it appears the background is ramping up towards the CI, however, there is one scan that does not contain a CI. It is simply a coincidence that a CI is made when the rise in the background peaks.

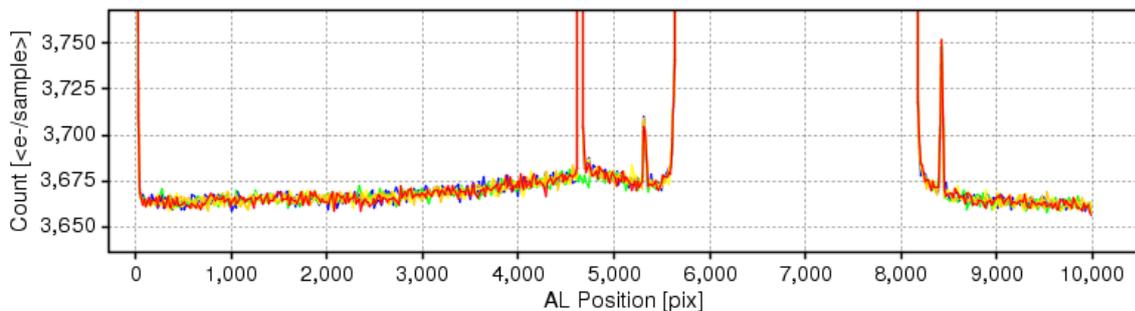


Figure 2.85: Running mean of all 5 scans in channel 1 from the irradiated data at $G = 13.6$ with a DOB of 2 and a CI delay of 1000.

Figure 2.86 shows on the left how the mean background level behaves in each of the channels in the first set of reference data. On the right the background level in consecutive scans is shown. The systematic offsets seen in this figure have already been observed in RC1 and RC2 and are again attributed to the PEM anomaly. Between scans the background level remains fairly stable.

On the left of Figure 2.87 the noise distribution from the first set of reference data in sample 1 (prescan) and samples 7 to 28 is shown. The mean noise level in sample 1, equivalent to the RON, is $5.63 \pm 0.16 \text{ e}^-/\text{sample}$. In the image channels the mean noise is $9.76 \pm 0.29 \text{ e}^-/\text{sample}$, which indicates a

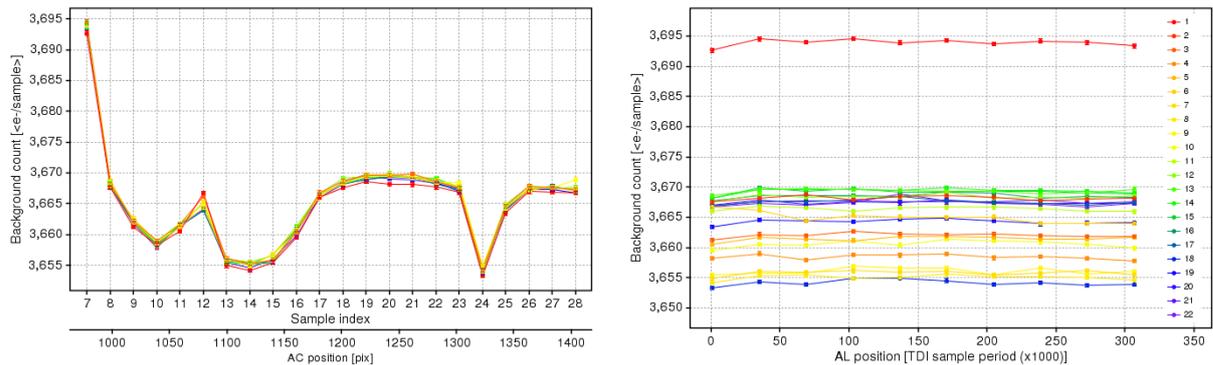


Figure 2.86: Background behaviour in the reference data from the first MF test with a DOB of 2. (left) Mean background count as a function of AC position. (right) Mean background count measured over consecutive scans.

DOB of $5.30 \pm 0.35 e^-/\text{pixel}$ (see eqn. 2.4). The DOB level is expected to be $2 e^-/\text{pixel}$. On the right of Figure 2.87 comparable results for the noise levels in the IA are shown when the DOB level is nominally 2, 4 and 7 in the reference data. For a quoted DOB of 4, the observed DOB level is $10.46 \pm 0.55 e^-/\text{pixel}$ and for a quoted DOB of 7, the observed DOB level is $24.42 \pm 1.00 e^-/\text{pixel}$.

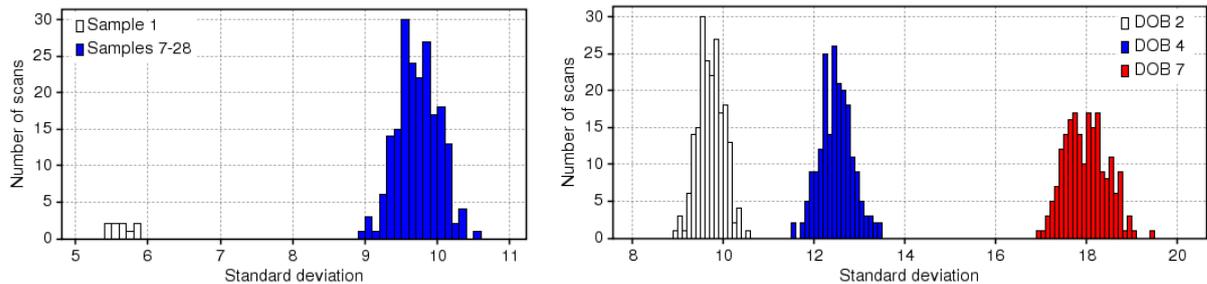


Figure 2.87: Noise in the background observed in the reference data: (left) With a DOB of 2, in sample 1 and samples 7 to 28, (right) with a DOB of 2, 4, and 7 in samples 7 to 28.

Table 2.16 shows the DOB level observed in other test cases. It is clear that the DOB level is, in general, much higher than the expected level. Presumably, this is caused by a miscalibration of the DOB level.

In some cases in the irradiated data, there are in fact two or three distinct noise distributions. For these test cases the mean of the densest group of values is given in Table 2.16. The outliers appear to originate from two distinct background signals seen in samples 21 and 27. The plot on the left of Figure 2.88 illustrates the anomalous behaviour in these samples, using the irradiated data taken at $G = 15$ with a DOB of 2 and a CI delay of 1000. The mean background in samples 21 and 27 jump ~ 100 to $\sim 150 e^-/\text{sample}$ above the surrounding background level. Recent evidence from an analysis of the sky-like data suggest that jumps of this magnitude could be caused by hot pixels rather than the PEM anomaly. In either case, a flat background cut removes such variation in the background level. The plot on the right of Figure 2.88 shows that the background is consistent in each scan.

Examining the global statistics on the noise distribution in both the reference and irradiated data gives further evidence that the noise in the background cannot be explained entirely by a standard Gaussian

Medium and Faint star tests					
G-mag	Injection delay [pix]	Target DOB [e ⁻ /pixel]	Mean prescan RON [e ⁻ /sample]	Mean scan total background noise	Observed DOB [e ⁻ /pixel]
13.6	Reference ¹	2	5.63 ± 0.16	9.76 ± 0.29	5.30 ± 0.35
		4	5.63 ± 0.20	12.54 ± 0.36	10.46 ± 0.55
		7	5.60 ± 0.13	18.01 ± 0.47	24.42 ± 1.00
	1000	2	5.67 ± 0.15	10.36 ± 1.56	6.27 ± 1.91
		4	5.69 ± 0.13	13.03 ± 1.49	11.45 ± 2.29
		7	5.71 ± 0.12	18.07 ± 1.12	24.49 ± 2.39
20	Reference ¹	2	5.69 ± 0.14	9.76 ± 0.30	5.24 ± 0.36
		4	5.61 ± 0.13	12.59 ± 0.38	10.59 ± 0.57
		7	5.68 ± 0.22	17.55 ± 0.59	22.98 ± 1.23
	1000	2	5.78 ± 0.12	10.11 ± 1.72	5.73 ± 2.05
		4	5.67 ± 0.18	12.67 ± 1.22	10.70 ± 1.83
		7	5.69 ± 0.15	16.68 ± 0.98	20.49 ± 1.93

Bright star tests					
G-mag	Injection delay [pix]	Target DOB [e ⁻ /pixel]	Mean prescan RON [e ⁻ /sample]	Mean scan total background noise	Observed DOB [e ⁻ /pixel]
12.6	Reference	0	5.63 ± 0.22	6.72 ± 0.31	1.12 ± 0.29
		2	5.61 ± 0.21	8.38 ± 1.23	3.23 ± 1.22
		4	5.68 ± 0.19	10.33 ± 0.40	6.20 ± 0.50
	1000	0	5.71 ± 0.15	7.17 ± 1.84	1.57 ± 1.56
		2	5.74 ± 0.12	8.95 ± 3.07	3.93 ± 3.24
		4	5.74 ± 0.18	10.44 ± 1.41	6.34 ± 1.74
11.1	Reference	0	5.68 ± 0.11	6.69 ± 0.59	1.04 ± 0.47
		2	5.74 ± 0.14	7.13 ± 0.61	1.49 ± 0.52
		4	5.65 ± 0.21	7.59 ± 0.49	2.14 ± 0.46
	1000	0	5.51 ± 0.15	6.58 ± 1.36	1.08 ± 1.06
		2	5.53 ± 0.19	7.01 ± 1.34	1.55 ± 1.11
		4	5.59 ± 0.21	7.51 ± 1.29	2.10 ± 1.15

¹Reference data were taken at G = 15.

Table 2.16: Statistics on the background noise for a selection of test cases.

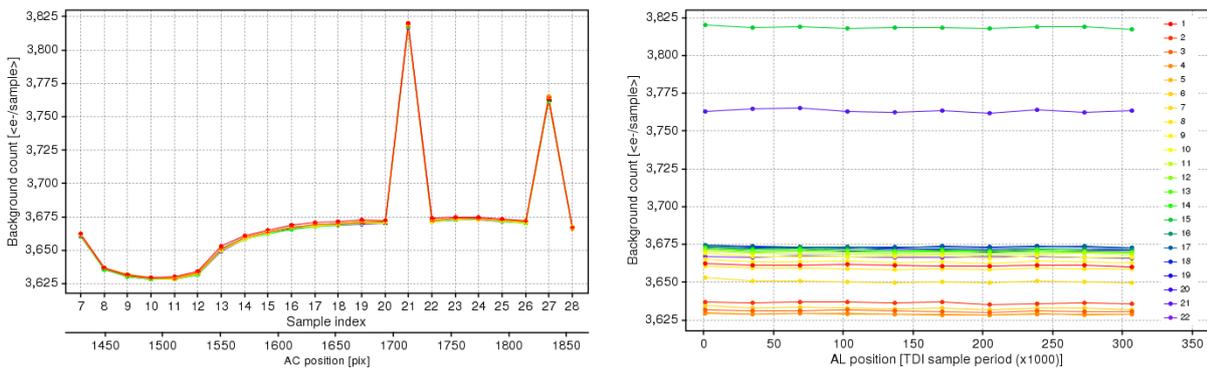


Figure 2.88: Mean background level in the irradiated data at G = 15 with a DOB of 2 and CI delay of 1000 as a function of (left) AC position and (right) time.

distribution. There is an underlying systematic bias in the noise that in previous campaigns has been attributed to the PEM anomaly. This is illustrated in Figure 2.89 which shows, on the left, the set of

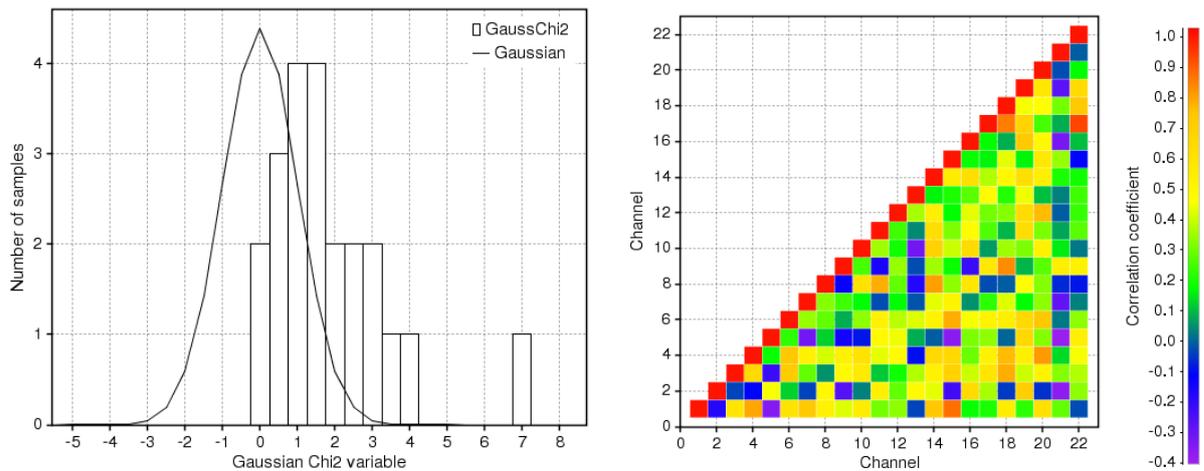


Figure 2.89: Global statistics on the reference data in the first MF test with a DOB of 2. (left) Background noise distribution represented as a Gaussian variable g compared to a Gaussian curve of $\sigma = 1$ containing the same number of samples. (right) Cross-correlation matrix for variations in the background.

Gaussian variables g (see eqn: 2.8) for the reference data from the $G = 13.6$ test case with a DOB of 2. The plot on the right of Figure 2.89 shows the correlation matrix for variations in the background (see eqn: 2.9). Significant correlations continue to exist, though for the first time this does not apply to every channel. A different AF CCD has been used in RC3 compared to that used in RC1 and RC2, which may partly explain the reduced dependency between the CCD and PEM seen here.

Test cases with CI delays of 10 and 30 were not analysed because of an additional feature found in the raw data associated with the CI. This feature is illustrated in Figure 2.90. Here, the raw data from two scans taken at $G = 20$ and read out of channel 1 at these two delays have been overlaid. Only the base of the CI profile and the first two images are shown. Clearly, the CI has a tail, which after a delay of 10 distorts the shape of the first image and, after a delay of 30, also appears to disturb the first image. Since channel 1 passes over the *non-irradiated* part of the CCD, traps releasing charge back into the

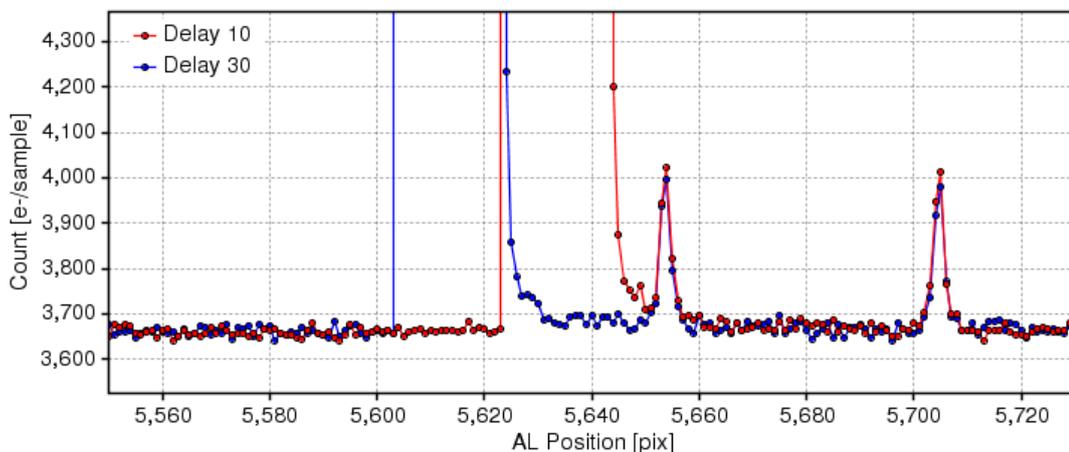


Figure 2.90: CI tails observed at delays of 10 and 30 in channel 1 impinging upon the first two images ($G = 20$) in the stellar field.

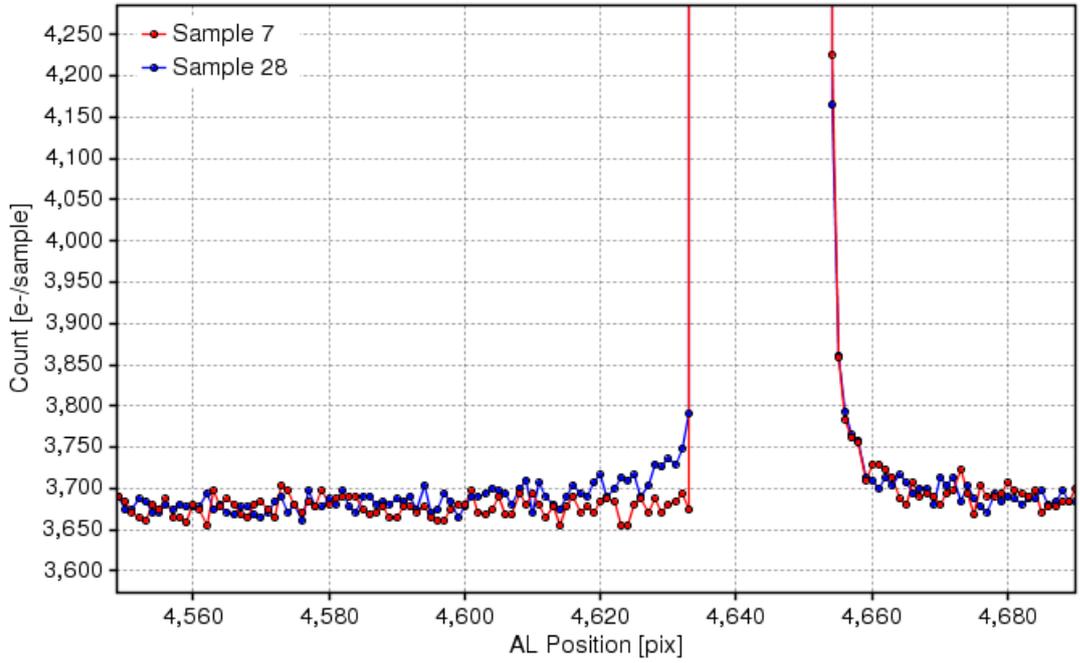


Figure 2.91: The background signal in sample 28 (channel 22) appears to "ramp-up" to the CI.

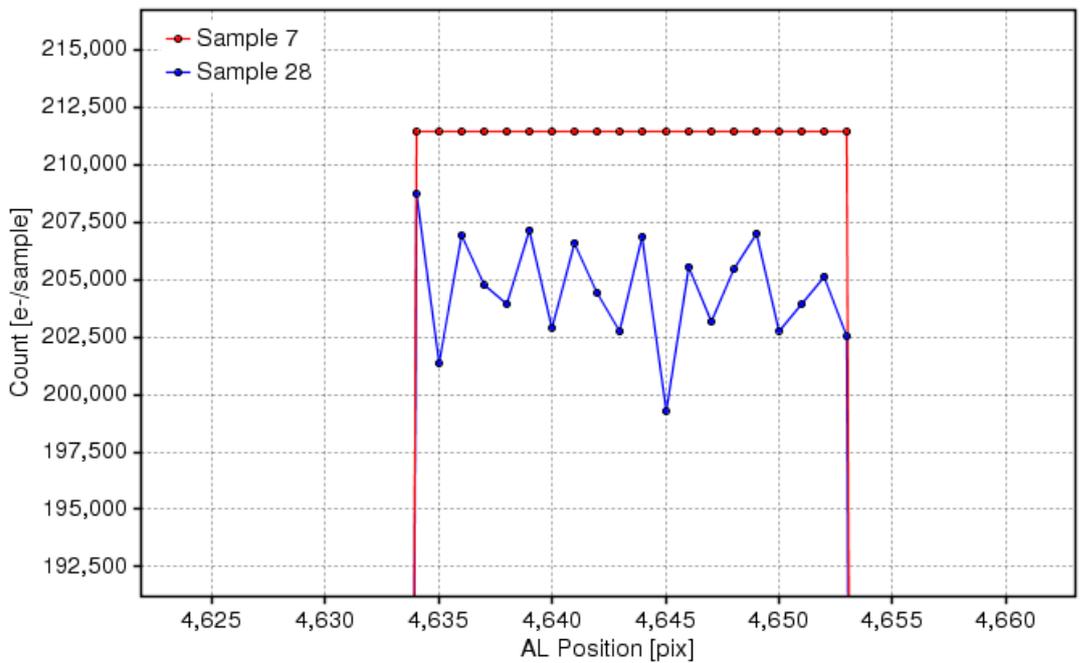


Figure 2.92: The CI profile appears damaged in sample 28 (channel 22) relative to the same signal seen in sample 7 (channel 1).

background can not originate from radiation damage. They must, instead, be attributed to some form of design trap. Currently, no model has been put forward to handle design traps. Therefore, it is not possible to remove the CI tail in this part of the CCD and properly calibrate the irradiated data with these short delays. For longer delays, the CI tail no longer impinges on the first image, which is why the analysis begins at a delay of 100.

Other features that appear to distinguish the behaviour of the CI profile in channels 1 (sample 7) and 22 (sample 28) are shown in Figure 2.91 and Figure 2.92. In Figure 2.91 the base of the CI profile in channel 1 has been overlaid with the CI profile in channel 22, both of which are part of the same scan and read out of the CCD at the same time. It appears the background signal prior to the CI in channel 22 "ramps-up" to the CI, whereas in channel 1 the CI is instantaneous. The cause of this "ramp-up" is currently unknown. It has been labelled the "divination anomaly" by Astrium. They also mention that when the flight model PEMs are used this anomaly is no longer there.

In Figure 2.92, the top of the CI profile in each channel has been overlaid. Full saturation is expected in both channels. However, the signal read out of channel 22 is eroded. Since channel 22 passes through more of the SR, it is now thought that this could be the cause of the damage seen here. Further tests are being conducted to characterise these anomalies.

To summarise:

- There are some features in the raw data that have already been observed in RC2: a) a sharp decrease at the beginning of certain scans due to processes occurring outside the IA; and, b) light leaking onto the CCD via screws holding the mask in place.
- Extraneous signals or "cosmics" continue to appear but are not as prevalent as in RC2.
- The background level in each channel continues to vary depending on the electronic offset added by the PEM. Variations in the background are correlated to this offset. Dramatic jumps ($>100 e^-/\text{sample}$) in the background level seen in the irradiated data may be caused by hot pixels.
- The CI profile contains a tail in the *non-irradiated* part of the CCD that interferes with stellar images if the CI delay is 10 or 30 pixels. As this charge is thought to be caused by design traps, these tests have not been analysed.
- On two counts the CI profile in channel 22 is noticeably different to that seen in channel 1: a) a small part of the signal in front of the main CI is "ramping-up" to the CI; and, b) the actual CI profile is eroded in channel 22. In the former case the cause is unknown, and in the latter case the damage is thought to arise from an extended time spent passing through the SR.

Extracting the raw image

In earlier analyses the raw image was extracted by setting a flux threshold below the peak of the first image. In RC3 this practice was bypassed due to: a) time constraints; and, b) the repeatability of the data, i.e. the position of the peak flux in the first image could, on most occasions, always be found close

to a predefined pixel position after readout. Though checks were still made to ensure this was the case, the overall time taken to analyse the data was reduced.

Extracting the CI tail

Details on the function used to model the CI tail can be found in Hambly & Cross (NCH-003). The probability that a trap will release charge after a certain amount of time is usually modelled by an exponential function. A sum of exponentials accounts for more than one trap species. Hambly & Cross used a weighted non-linear least squares fit via a Levenburg-Marquardt algorithm to model exponential charge release curves using test data provided by Astrium to characterise the CI release profile. The fit must solve for the amplitude and time constant for i number of exponentials that describe as closely as possible the CI tail. In this case $i = 3$. Figure 2.93 shows the parameters that are used to calculate the flux released by the CI.

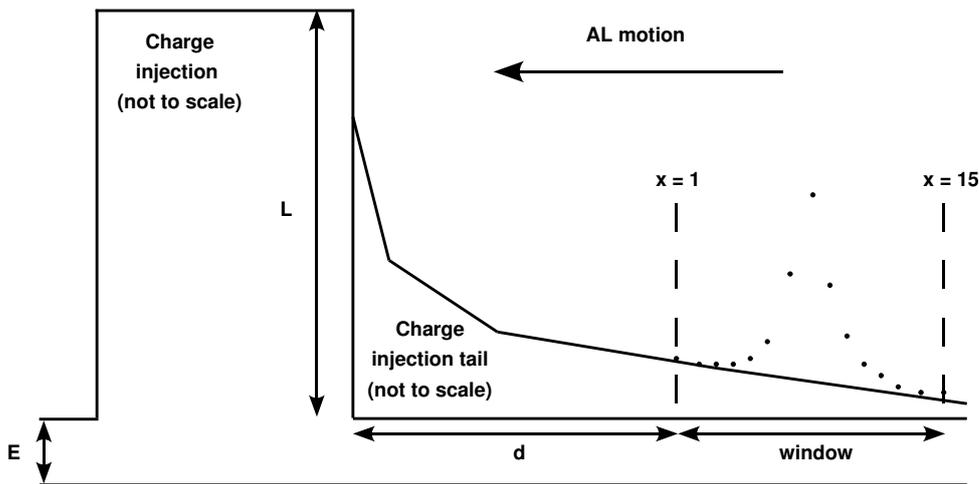


Figure 2.93: Illustrating how the CI tail is extracted from the stellar field. E = electronic offset added by the PEM, L = the CI level (e^-/pixel), d = the CI delay up to the first pixel in the window, x = the pixel position in the window.

Firstly, to extract the CI tail under the stellar field the electronic offset E has to be subtracted. Then, for each pixel position x , where $x = 1, \dots, 15$, in the window, the flux in the CI tail f_x can be calculated as follows:

$$f_x = L^{0.6742} \times \sum_{i=1}^{i=3} a_i e^{-(x+d)/t_i}, \quad (2.13)$$

where L is the CI level ($20 \text{ ke}^-/\text{pixel}$) and d is the delay between the CI ending and the window beginning. The scaling parameters a_i are weighted by the sum of another set of parameters A_i found via the test fitting:

$$a_i = \frac{A_i}{\sum_{j=1}^{j=3} A_j}. \quad (2.14)$$

Following discussion with Nicholas Cross¹⁴, the following values for A_i and t_i were adopted:

- $A_1 = 763.9807$, $A_2 = 22.9981$, $A_3 = 6.900426$.
- $t_1 = 0.6685$, $t_2 = 26.54$, $t_3 = 490.4$.

If y_x is the uncorrected flux at position x , the corrected background counts b_x after the CI tail has been removed are:

$$b_x = y_x - (N \cdot f_x). \quad (2.15)$$

In the data y_x is given in units of e^-/sample . Therefore, f_x has to be multiplied by an extra factor N where N is the number of pixels per sample, which in this case is 12. This assumes the charge being released in each pixel is the average over a 12-pixel AC window.

2.4.3 The reference data

The Global Reference Curve (GRC)

For every irradiated test conducted at a different magnitude, a new set of reference data was collected. There is no need to show the GRC in every case, since the GRC was highly consistent in every test. Figure 2.94 shows the GRC with its mirror image for the first column, taken from two different sets of reference data at two different magnitudes. The GRC in both cases is very symmetric and remained so throughout the tests that were considered in this analysis.

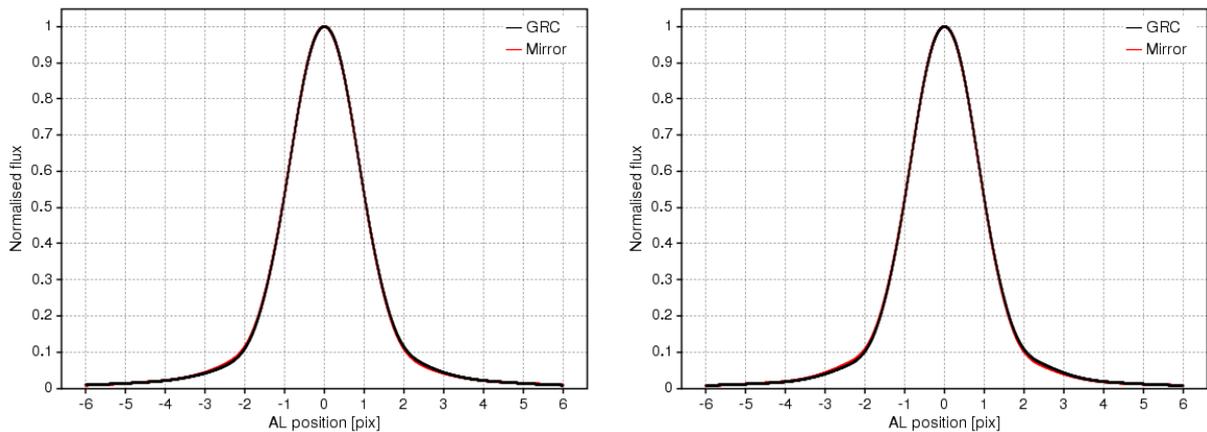


Figure 2.94: (left) The normalised GRC with its mirror image taken at $G = 15$ from the MF star tests. (right) The normalised GRC with its mirror image taken at $G = 11.1$ from the bright star tests.

¹⁴Institute for Astronomy, Edinburgh

Geometric and flux response calibration

Since the mask in the astrometric regular tests is the same as that used in RC2, the same procedure was followed to extract relative centroid positions and fluxes across the stellar field. This incorporated the extended set of non-irradiated channels either side of the mask in the irradiated dataset (i.e. rather than refitting the linear fit to column 50 to only the first two channels, this linear fit was refitted to all six equivalent non-irradiated channels). Figure 2.95(a) shows how the relative centroid positions drift away from a rigid AL pitch of 50 pixels between each pinhole. In total, 35 pixels are accumulated over the field, roughly the same (~ 37 pixels) as observed in RC2. The geometric offsets are shown in Figure 2.95(b). Here the pattern of offsets are quite different to those seen in RC2 (Figure 2.95(c)). Crucially, the range

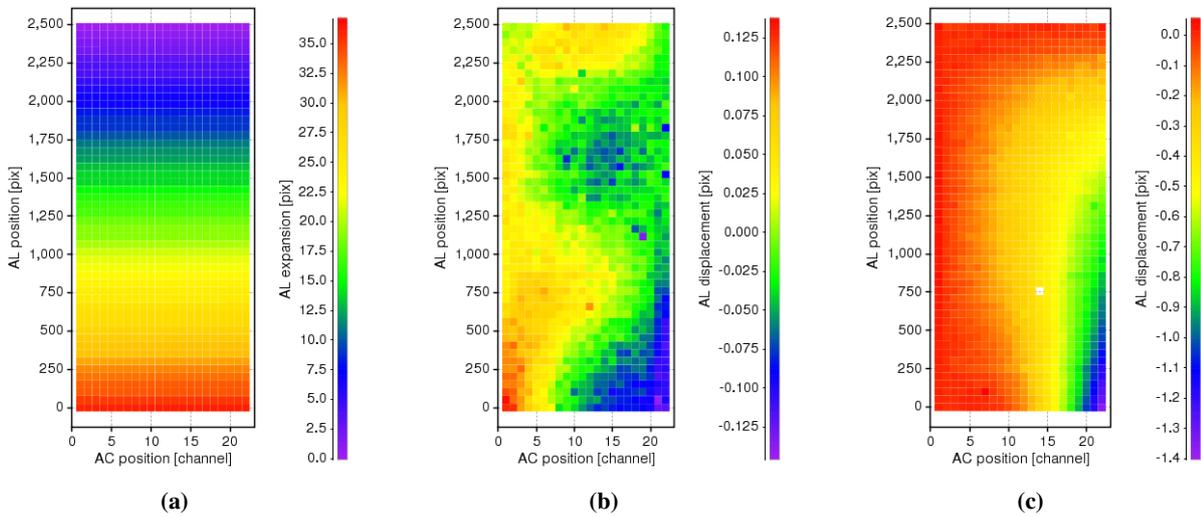


Figure 2.95: Geometric calibration. (a) Deviations from a rigid 50 pixel AL pitch. (b) Offsets remaining between the calibration and the centroid positions. (c) Comparable result to (b) taken from RC2.

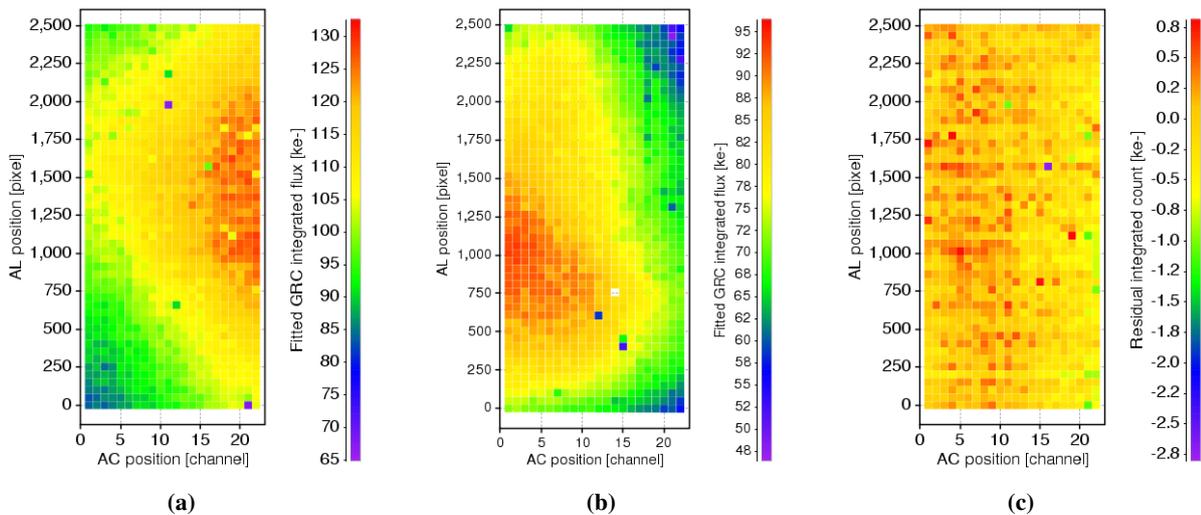


Figure 2.96: Flux response calibration. (a) Fitted GRC integrated flux in the reference data at $G = 15$. (b) Comparable to (a) from RC2. (c) Residual integrated flux between the fitted GRC and local LSF.

is far smaller. This tends to support the idea that in RC2 it was probably a pressure gradient across the mask, stressing the bottom right-hand corner, that was causing the large geometrical offset.

In RC2 it was observed that the peak flux of the fitted GRC was scaled in such a way that it compensated for variations in the FWHM of the local LSF down and along each column. This inverse correlation meant the integrated flux could be found using the fitted GRC. This is also observed in RC3, though there are a few differences between the nominal and the observed magnitude in each campaign. Figure 2.96(a) shows the integrated flux over the fitted GRC in the reference data at $G = 15$ (the first MF star test). The maximum flux of $\sim 130000 e^-$ is closer to $G = 14.5$. For comparison, Figure 2.96(b) shows an example of a comparable set of reference data from RC2. The maximum flux here of $\sim 95000 e^-$ is closer to $G = 14.8$. It is also clear from these two plots how the position of the light source relative to the mask has shifted between the two campaigns. Regardless of these differences, Figure 2.96(c) shows how the residuals in the integrated flux between the fitted GRC and local LSF across the stellar field in RC3 differ by $<1\%$.

2.4.4 The irradiated data

Sample Clustering

One of the issues that was raised after analysing RC2 was sample clustering. Each scan tended to sample the LSF at the same point leaving the image under-sampled, and negating the use of having multiple scans. This issue reappeared in the astrometric regular tests, though only in certain test cases. Figure 2.97 shows four examples of the same image taken from different test cases in RC3. It is clear where the samples are clustering.

In general, the reference data satisfied the recommendations made in RC2. However, samples tended to cluster when a CI had been made. The need to inject charge at regular intervals is a requirement in the Gaia mission, though for the campaign analyses a certain degree of dither is required between scans.

Recalibrating the rotation of the mask

In RC3 the original methodology is used (i.e. using only equivalent non-irradiated images in both the reference and irradiated data) to calibrate the geometry of the mask.

Figure 2.98 shows the result of applying this calibration to one of the first irradiated tests performed. Here $G = 13.6$ and no CI has been made. On the left the relative centroid positions are shown for all images in column 1 with the geometric calibration. On the right the relative centroid biases are shown across the stellar field. Some observations can be made:

- Firstly, the mask is tilted in the opposite direction over the leading column of stars relative to the CCD compared to the same images seen in RC2.
- There are systematic features in every column. The sharp decrease in the bias between channels 10 and 11 suggest a CCD stitch is co-located with images in these channels. There are similar

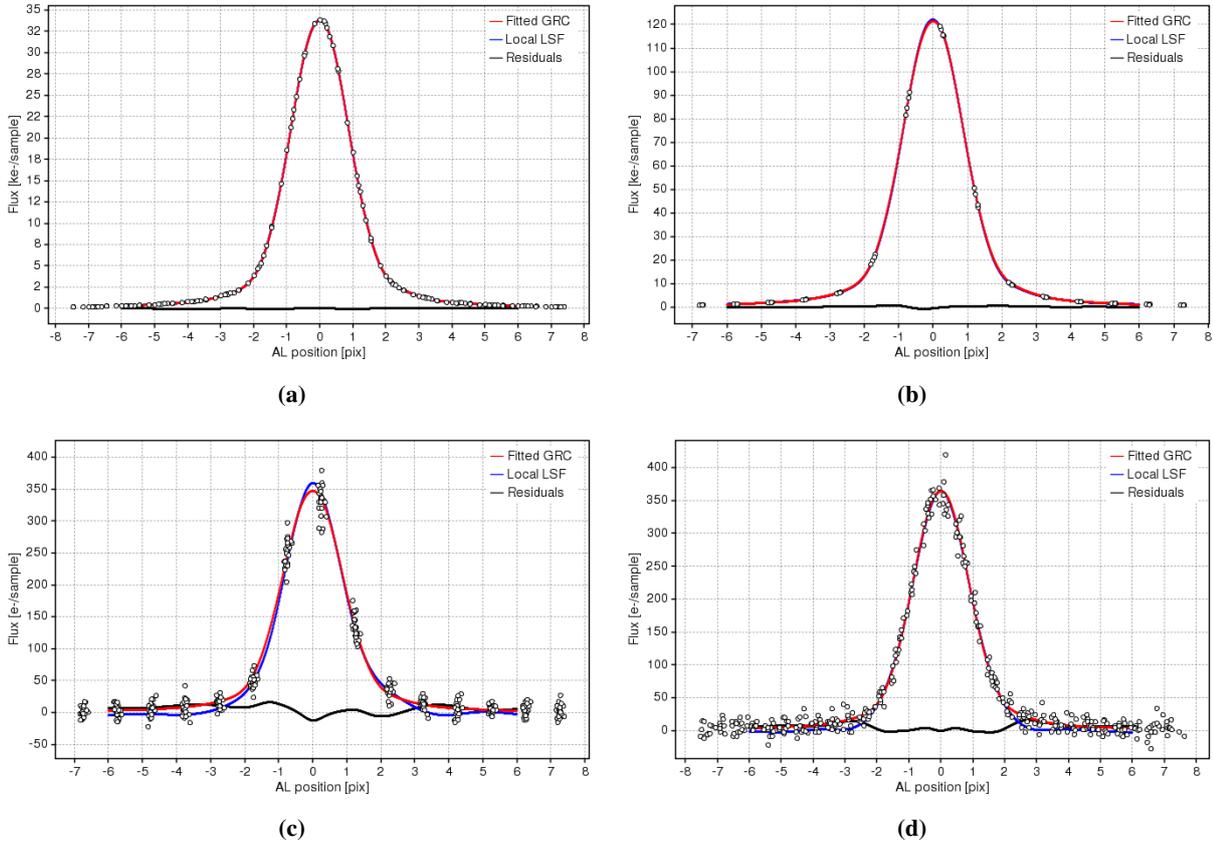


Figure 2.97: Sample clustering is observed in RC3 but only when a CI has been made. Here sample 7 is shown from (a) reference data at $G = 15$ (b) irradiated data with a CI delay of 30, $G = 13.6$, (c) irradiated data with a CI delay of 1000, $G = 20$ and (d) irradiated data without CI, $G = 20$.

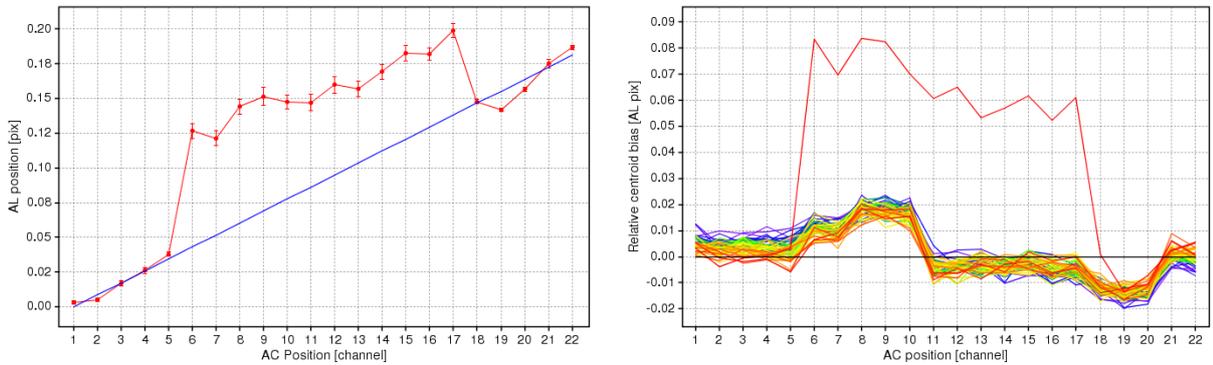


Figure 2.98: Original geometric calibration. (left) Relative centroid positions in column 1 at $G = 13.6$ without CI. (right) Relative centroid biases across the stellar field.

features in channels 6 to 9. The systematic offsets seen in channels 19 to 22 appear to correlate with these features.

- After column 1 (red line) the bias suffered is considerably reduced.

According to documentation on CCD stitches, each stitch is separated AC by ~ 250 pixels or $\sim 12/13$ channels. The offset seen in channel 10 should therefore be repeated in channel 22/23. The repeating pattern across all 22 images does suggest there are systematic biases adjacent to a CCD stitch and also 80 pixels away from the stitch. Furthermore, the sensitivity of the CCD between stitches is not identical. As a result, the geometric calibration is distorted by the systematic features seen in channels 20, 21 and 22. Taking account of this, the rotation of the mask was recalibrated using only the first 3 and the first of the last 3 non-irradiated channels. Figure 2.99 shows the result of doing this.

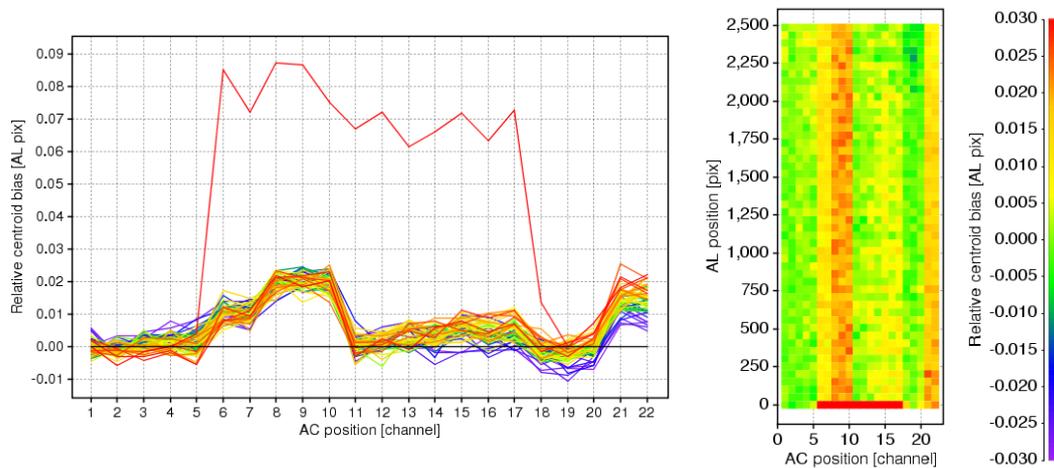


Figure 2.99: Refining the geometric calibration. (left) The geometric calibration has been redone using only channel 1, 2, 3 and 20. (right) Relative centroid bias across the stellar field.

Relative centroid biases and charge losses across the stellar field

MF star tests In the MF star tests the level of DOB and the CI delay were varied for magnitudes between 13.6 and 20. Figure 2.100 shows the relative centroid bias found across the stellar field as a function of CI delay for a $G = 13.6$ source and a DOB level of 7. The CCD stitch and pixel column displacements can clearly be seen in Figure 2.100(a). When the delay is increased to 500 (Figure 2.100(b)) the biases measured in column 1 begin to increase relative to those measured in Figure 2.100(a) as expected. Although surprisingly, the systematic biases seen further up the mask are less than those seen in Figure 2.100(a). This effect is more pronounced when the delay is increased to 1000. There may be slight variations in the calibration between the different test cases which could emphasise features such as these (e.g. the charge injection tail not being removed properly). If this is not the case, then it may be that the effects of the DOB, which were known to be higher than reported, gradually come into play after a suitable amount of time has passed after a charge injection. According to the analysis of campaign 1, the greatest effect of applying a DOB is reducing the charge loss. However, according to

the next figure (Figure 2.101) there appears to be no reduction in charge loss with increasing time since charge injection, so it is not entirely clear if the DOB is the cause of these changes.

Since the reference data were collected at mag 15 in these tests a correction had to be applied as in RC2 before determining the charge loss. For this, an alternative procedure was followed. Since the fluxes in the non-irradiated channels on both sides of the irradiated area should agree, the flux ratios between what was measured in the the non-irradiated channels in the irradiated data and the non-irradiated data were found, and a linear fit was applied to these ratios to find the remaining scale factors. A quadratic fit was also tried, but its use could not be justified as the level of uncertainty in the fit was too high. With 4 datapoints it was assumed the flux could scale linearly in each column. The reference fluxes were multiplied by these scale factors to estimate the undamaged flux levels.

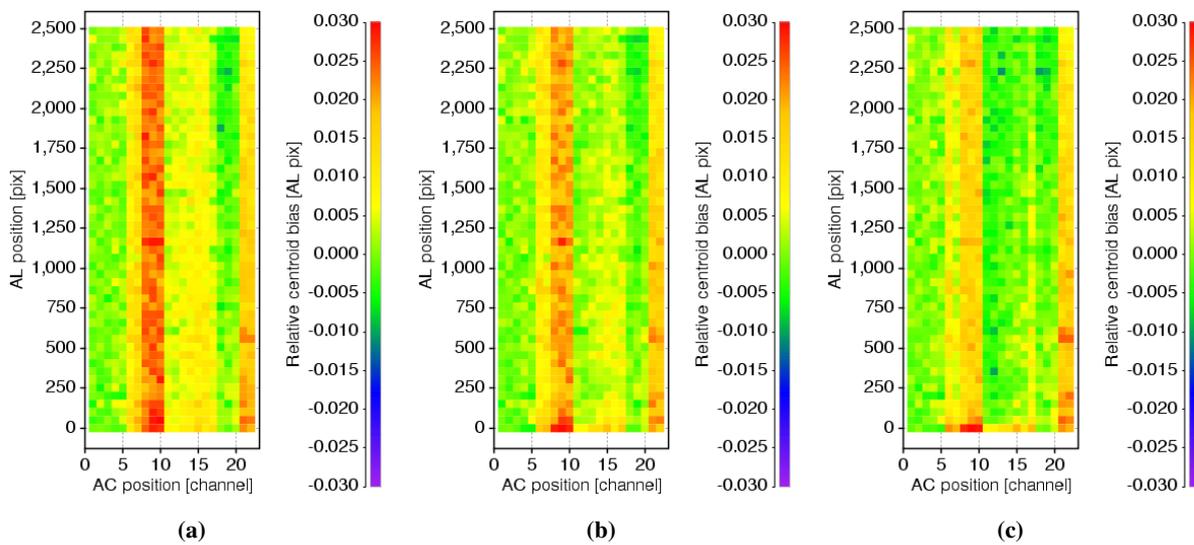


Figure 2.100: Relative centroid bias at $G = 13.6$, DOB of 7, and CI delay of (a) 100, (b) 500, (c) 1000.

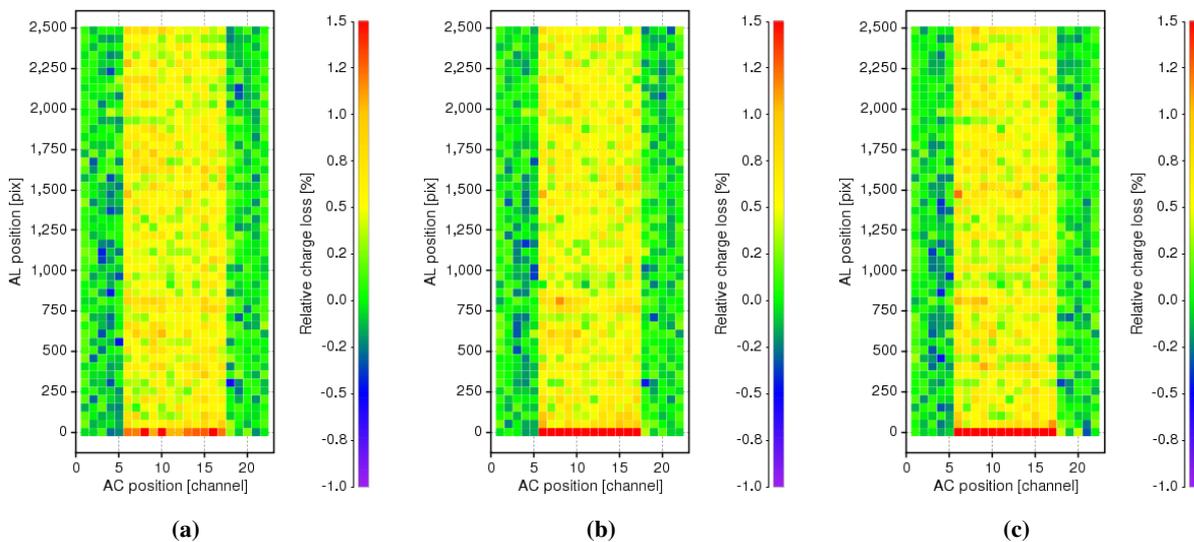


Figure 2.101: Relative charge loss at $G = 13.6$, DOB of 7, and CI delay of (a) 100, (b) 500, (c) 1000.

Figure 2.101 shows comparable results for the relative charge loss across the stellar field. Note that the actual charge loss measured in column 1 is not shown here, since the range of the z-axis is fixed to emphasise the charge lost across the rest of the stellar field. The charge loss seen after column 1 has passed appears not to change.

Now the biases observed for the faintest images at $G = 20$ with a CI delay of 1000 and a variable DOB level are considered. From Figure 2.102 note that, firstly, with this CI delay the damage in column 1 is reduced to levels that are comparable with the biases seen in all the other columns. Secondly, with a DOB of 2 the bias can reach ~ 0.12 pixels in some images, but not in those images that are in column 1. Lastly, increasing the DOB to 4 and 7 reduces the bias even more, not only in column 1 but in all images.

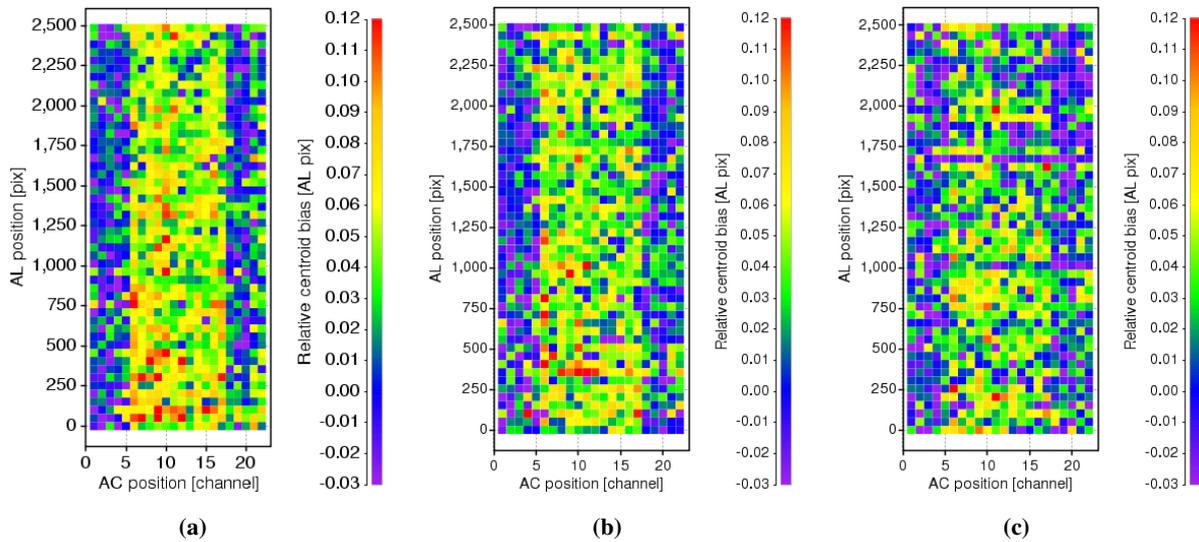


Figure 2.102: Relative centroid bias at $G = 20$, CI delay of 1000, and DOB level of (a) 2, (b) 4, (c) 7.

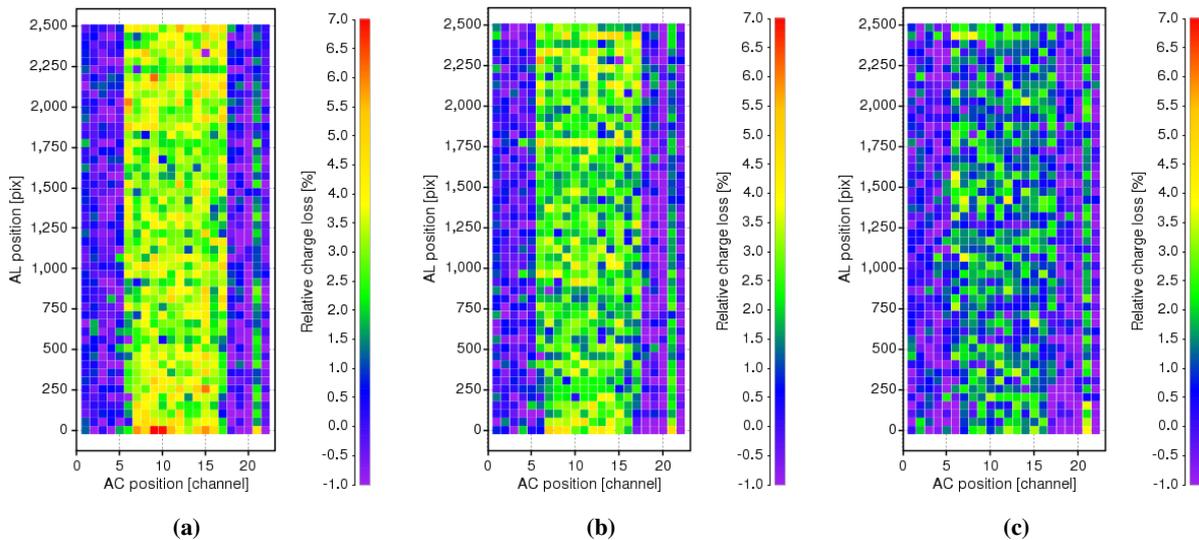


Figure 2.103: Relative charge loss at $G = 20$, CI delay of 1000, and DOB level of (a) 2, (b) 4, (c) 7.

This may be down to the DOB alone but could also be due to slight variations in the calibration between the different test cases.

The charge loss seen at $G = 20$, with a CI delay of 1000 and a variable DOB level is shown in Figure 2.103. With a DOB level of 2 there is some noticeable charge loss in column 1. However, this is reduced when the DOB level is increased to 4, and even more so when the DOB level becomes 7. According to the analysis of RC1 the DOB does mitigate the effects of radiation damage, more so for the faintest images. This is again the case here. Recall though that the levels of the DOB in absolute terms are far higher than the nominal levels of DOB and not representative of the DOB level Gaia may typically observe.

Bright star tests In the bright star tests the level of DOB and the CI delay were varied for magnitudes between 11.1 and 12.6. To avoid saturating the CCD readout, gates were activated.

In Figure 2.104 the centroid bias across the stellar field is shown as a function of CI delay. $G = 11.1$ and the DOB level is 4. Gate 9 has been activated. The largest biases are found next to the stitch between channels 9 and 10. For delays of 100 and 1000 the bias remains negligible, and only becomes slightly more pronounced when the delay is increased to 10000.

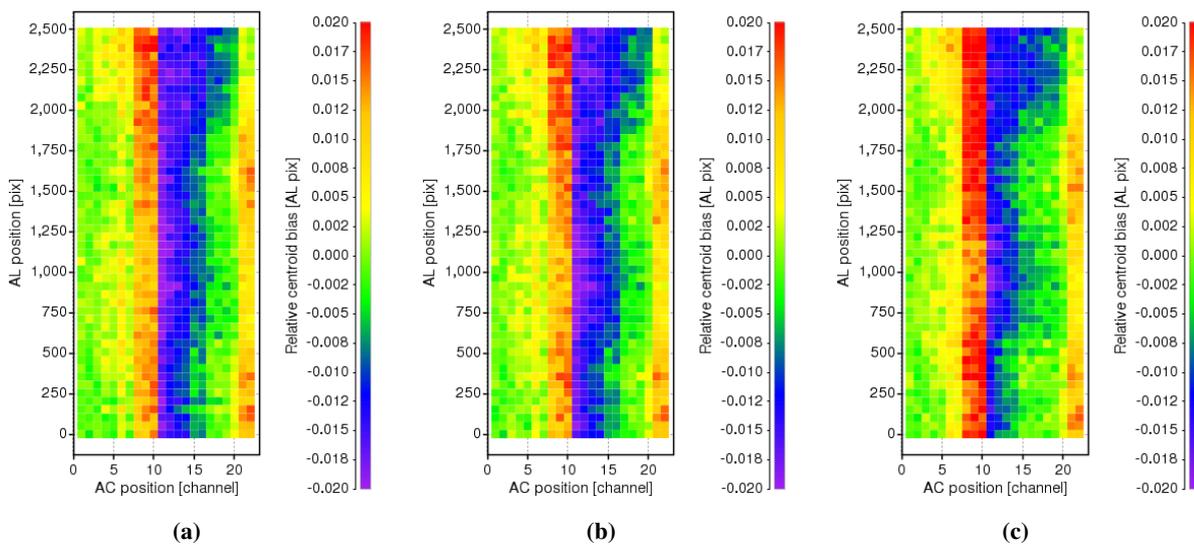


Figure 2.104: Relative centroid bias at $G = 11.1$ with a DOB of 4 and a CI delay of (a) 100, (b) 1000, (c) 10000.

The relative charge loss at $G = 11.1$, under the same test conditions, is shown in Figure 2.105. For such a bright source the charge lost in every image is almost always $<1\%$ for delays <1000 . Even with a delay of 10000 the charge loss is minimal ($\sim 1\%$). Interestingly, in all three test cases there appears to be a particular group of images that are preferentially damaged. It is not entirely clear why this is but appears to have some dependency on the magnitude of the source. Another interesting feature can clearly be seen in Figure 2.105(c). Here the discrete variations between the charge loss seen in adjacent channels is similar to the variations that appeared in the analysis of RC2 (see Figure 2.56) when a resonance arose between the spacing of the spline knots and the positions around which samples tended to cluster. In

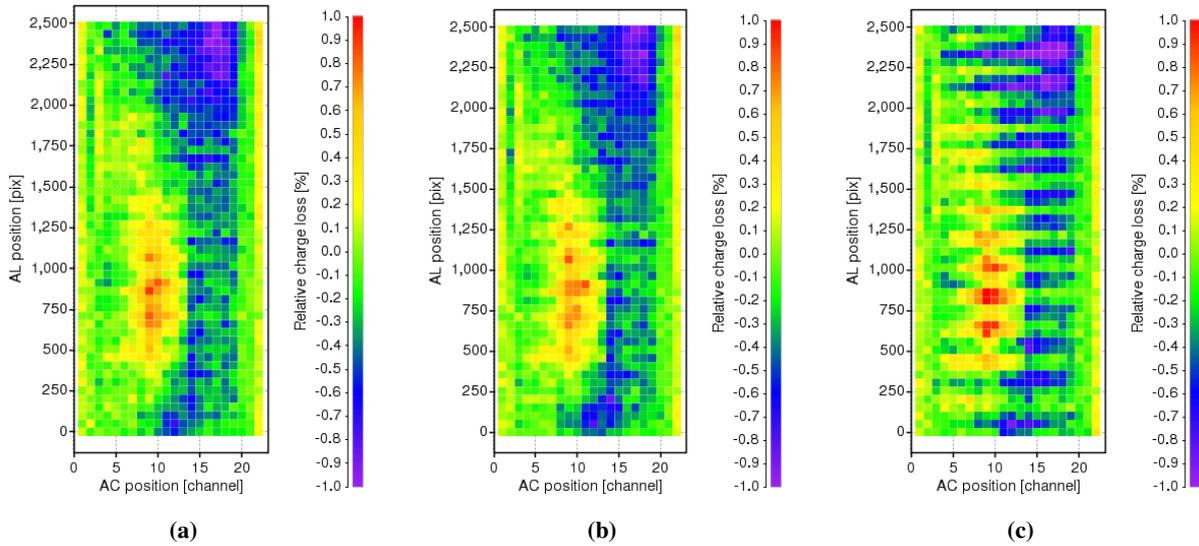


Figure 2.105: Relative charge loss at $G = 11.1$ with a DOB of 4 and a CI delay of (a) 100, (b) 1000, (c) 10000.

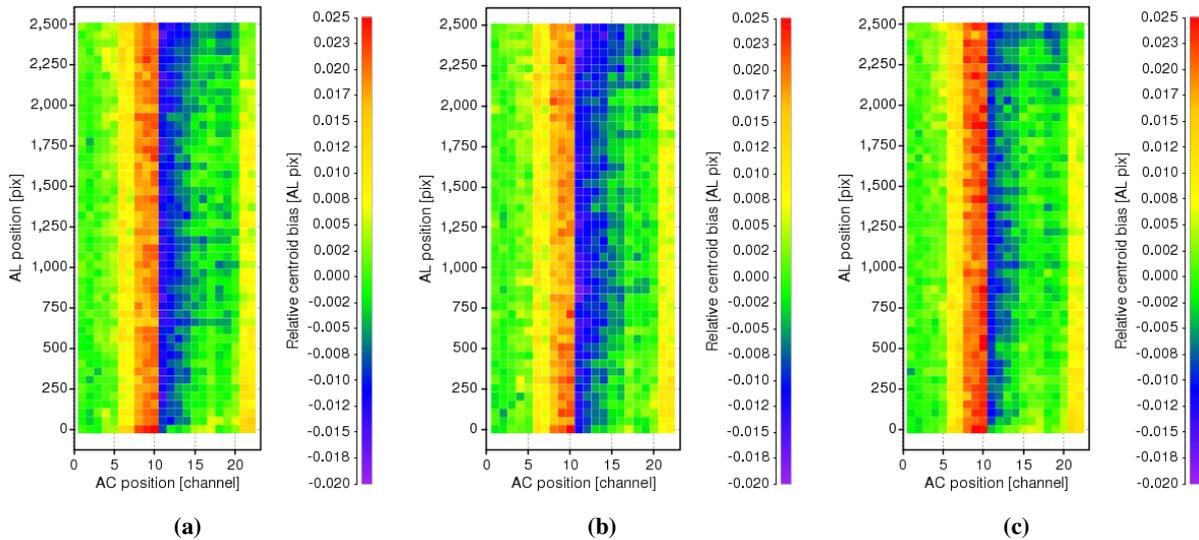


Figure 2.106: Relative centroid bias at $G = 12.6$ with a CI delay of 1000 and a DOB level of (a) 0, (b) 2, (c) 4. Note that the actual DOB levels are ~ 1.6 , ~ 3.9 and ~ 6.3 respectively.

this particular test case 10 scans were performed and the samples do cluster around the same sub-pixel positions after they have been combined, suggesting the same type of resonance is occurring.

For the faintest images in the bright star tests the effect of an increasing DOB level is shown. At $G = 12.6$ with a CI delay of 1000 and gate 11 activated, Figure 2.106 shows the centroid bias across the stellar field as a function of DOB level. Column 1 shows preferential, albeit reduced, bias as the DOB increases. The bias across the stellar field appears to decrease, as expected, with a DOB of 2. However, this effect appears to reverse when the DOB is increased to 4. The reason for this is not entirely clear. There may be some residual effects left over from the calibration, namely a mis-calibration of the charge

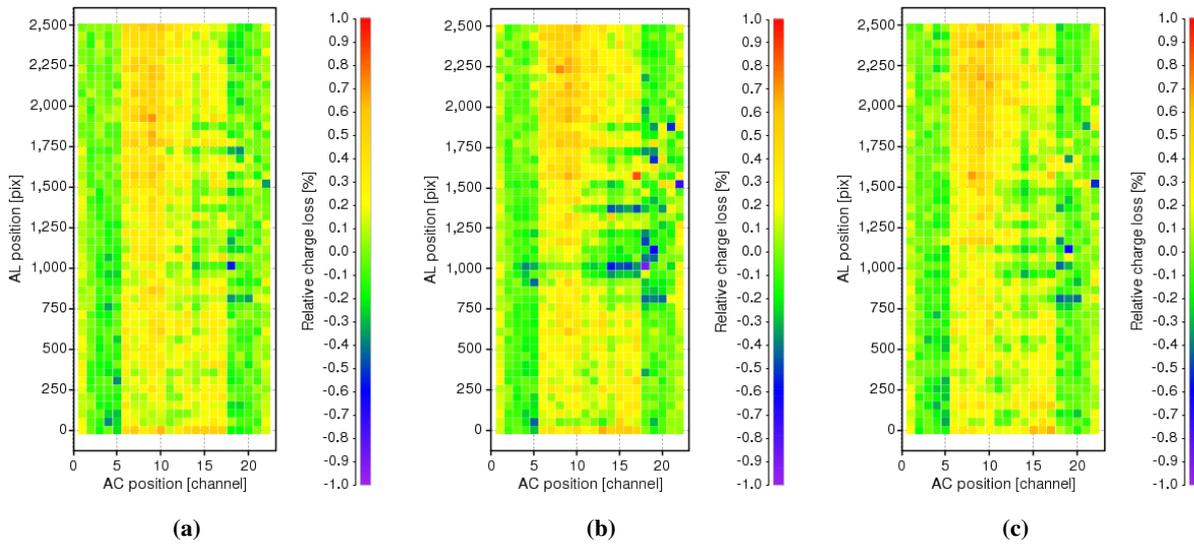


Figure 2.107: Relative charge loss at $G = 12.6$ with a CI delay of 1000 and DOB level of (a) 0, (b) 2, (c) 4.

injection tail or the rotation of the mask (the latter assumes the fourth, isolated non-irradiated channel is a reliable indicator of the mask rotation). If this is not the case, the charge injection itself in this particular test case may have been less than expected. However, checking this is not possible as the readout for the injection is saturated.

Figure 2.107 shows how the charge loss appears across the stellar field at $G = 12.6$ under the same test conditions. Here, again, the charge loss is never greater than 1%. For all levels of DOB, column 1 suffers the most damage, though the loss is partially mitigated when the DOB is increased to 4. In each plot there is a group of images that are relatively less damaged. The location of this group appears to be where the illumination level is highest. In which case, the relative effects of radiation damage appear slightly less.

2.4.5 Relative centroid bias and charge loss in the leading column

In general, it is the leading column of stars that exhibit the most damage, so now the analysis focuses on this part of the stellar field.

Medium and faint star tests

Figure 2.108 shows how the centroid bias (left) and the charge loss (right) in the leading column at $G = 13.6$ vary as a function of CI delay with a DOB of 2. In RC2 it was found that a CCD stitch could introduce a systematic bias of ~ 0.02 pixels at the stitch boundary, and this is also seen here in channels 10/11 when the CI delay is 100. In fact, rather than just the stitch producing a systematic offset, every channel appears to shift the centroid by a measurable amount. As the injection delay is increased, these offsets are still seen. However, the effects of radiation damage are also added. Underlying systematic centroid biases must be caused by the relative displacement of each pixel line on the CCD as opposed to

radiation damage, so a correction is required. As for the flux calibration, the stitching can be accounted for as part of the calibration. Hence, the charge loss in the non-irradiated channels is negligible.

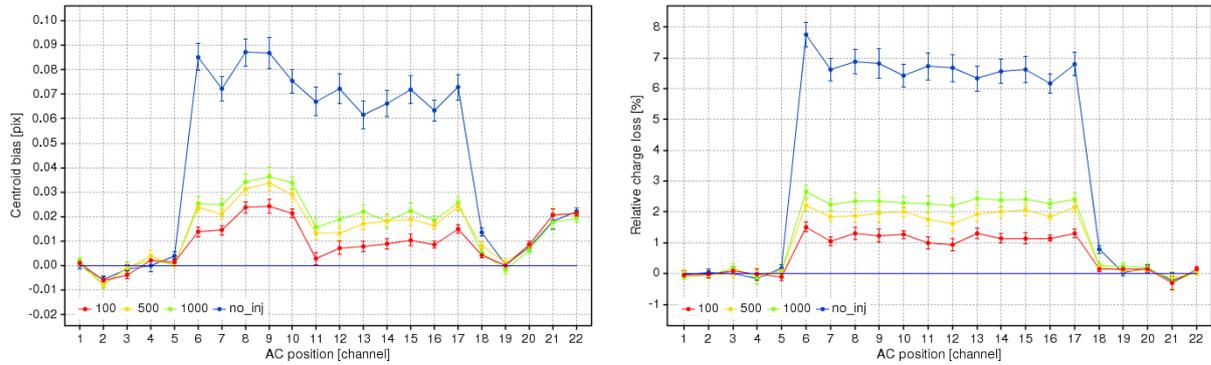


Figure 2.108: Relative centroid bias (left) and charge loss (right) in the leading column as a function of AC position and CI delay at $G = 13.6$ and DOB level of 2.

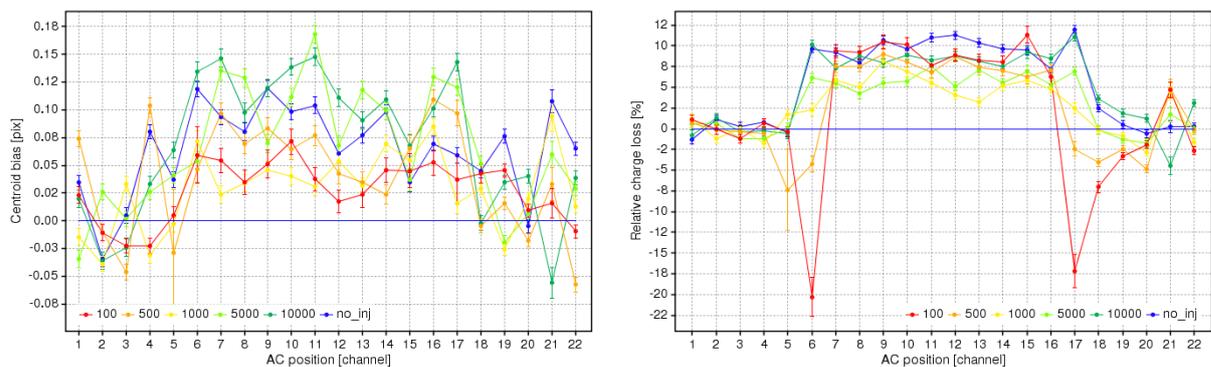


Figure 2.109: Relative centroid bias (left) and charge loss (right) in the leading column as a function of AC position and CI delay at $G = 20$ and DOB level of 2.

Figure 2.109 shows comparable results at $G = 20$. Both the bias and charge loss are worse and a strange behaviour is observed in channels 6 and 17 in the charge loss, especially when the CI delay is 100. This feature arises because the CI tail has not been removed in the channels that pass over the transition zone on the CCD, namely channels 4, 5, 6, 17, 18 and 19. In these channels the radiation conditions are unknown, so it is not possible to apply the charge release model.

Bright star tests

Figure 2.110 shows how the centroid bias (left) and the charge loss (right) in the leading column at $G = 11.1$ vary as a function of CI delay with no DOB. In these test cases gate 9 has been activated. The bias remains the same for every CI delay and tends to focus on the stitching effects. Note that the size of the systematic centroid bias on a CCD stitch is ~ 0.03 to 0.035 pixels whereas at $G = 13.3$ and no gate activated the centroid bias is ~ 0.02 . This suggests that by activating a gate the length of time a signal has to smear out is reduced, which increases the centroid bias seen around a stitch. In contrast, the charge

loss is very similar for every CI delay and, more importantly, remains very close to zero.

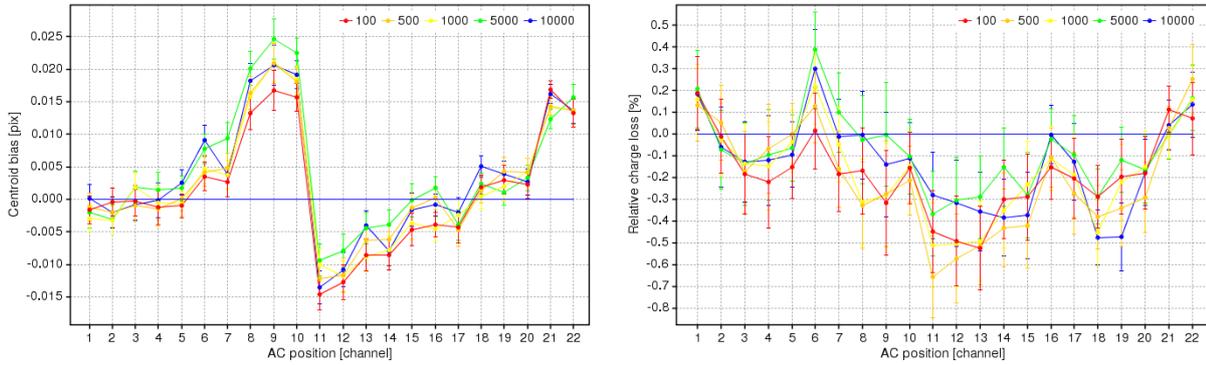


Figure 2.110: Relative centroid bias (left) and charge loss (right) in the leading column as a function of AC position and CI delay at $G = 11.1$ and no DOB.

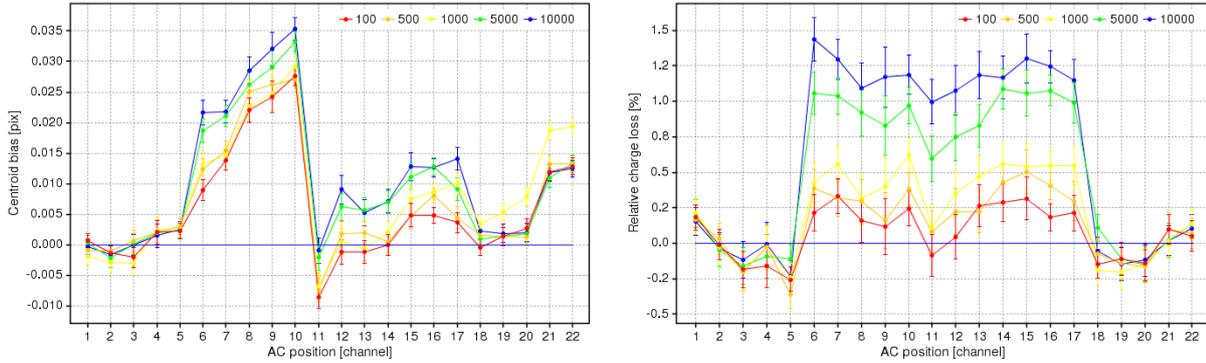


Figure 2.111: Relative centroid bias (left) and charge loss (right) in the leading column as a function of AC position and CI delay at $G = 12.6$ and no DOB.

In Figure 2.111 comparable results are shown for $G = 12.6$. Gate 11 has been activated. Here the damage is worse and increases as the CI delay increases. Again the systematic bias incurred by passing over a stitch is enhanced by activating a gate.

Correcting for systematic biases: pixel displacements and stitches

As mentioned above, a correction must be made to the observed biases to take into account the systematics incurred by pixel displacements and stitches across the CCD. For the MF star tests, it is assumed that the biases observed at the brightest magnitude, 13.6, with the closest CI delay, 100, and the lowest DOB level, nominally 2, are undamaged. The lowest DOB level is chosen since the SNR is higher and the observed biases with the highest DOB level are actually worse than with a lower DOB. This is clearly an approximation, since $\sim 1\%$ of charge is still being lost. Recall that test cases where the CI delay is shorter cannot be used, since a CI tail that interfered with the leading column of stars was seen in the non-irradiated area of the CCD.

In the bright star tests the stellar field does not pass over the entire length of the CCD without a gate

being activated. Staggering the transit length appears to alter the size of the systematic biases observed in each channel. Therefore, individual corrections must be made for each gate. It is assumed again that the images in the leading column taken with a CI delay of 100 and, in this case, the highest DOB, nominally 4, are undamaged. The highest DOB is chosen since the observed biases in these test cases are less than the biases observed with a lower DOB.

Once this correction has been made, the average bias and charge loss are determined for the fully irradiated area (channels 7 to 16).

Average biases and charge losses over the leading column

Figure 2.112 shows the average relative centroid bias and charge loss over the leading column at $G = 13.6$ as a function of CI delay and DOB level. Without CI the delay has been given an arbitrary value of 1×10^6 pixels. Clearly, decreasing the delay decreases CTI. However, increasing the DOB level makes no difference, especially when the delay is < 1000 .

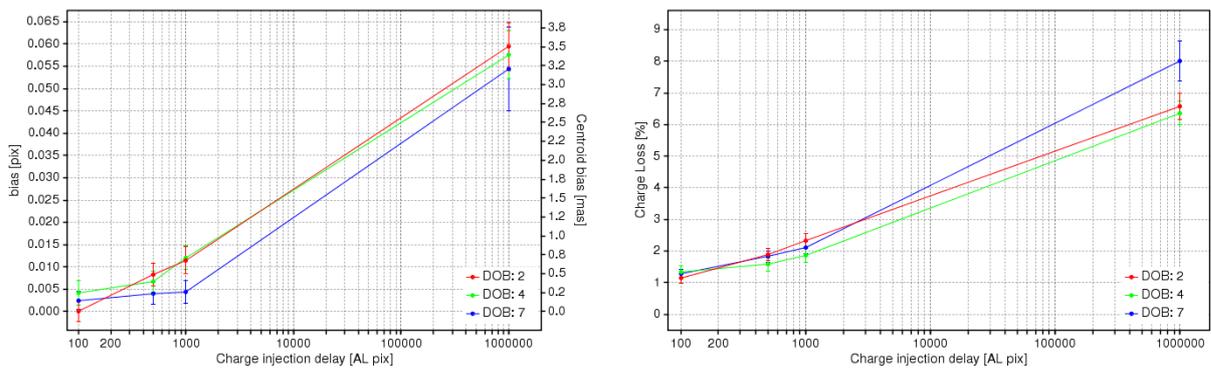


Figure 2.112: Relative centroid bias (left) and charge loss (right) as a function of CI delay and nominal DOB level at $G = 13.6$.

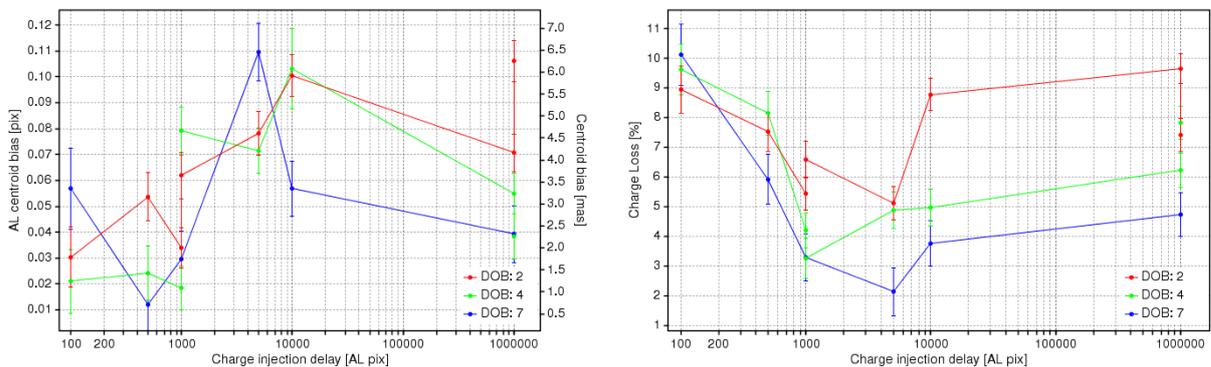


Figure 2.113: Relative centroid bias (left) and charge loss (right) as a function of CI delay and nominal DOB level at $G = 20$.

Figure 2.113 shows comparable results at $G = 20$. Here, longer CI delays were incorporated (5000, 10000) and in two test cases (namely with a delay of 1000 and without CI) the test was performed

multiple times. Two of these cases at nominal DOB levels of 2 and 4 were analysed. With a CI delay of 1000 the biases differ considerably. The charge loss also shows variations, which may increase or decrease depending on the DOB level. Without CI a strange behaviour is seen: as the bias increases the charge loss decreases and vice versa. Recalling that the repeat test was conducted two weeks after the first test suggests that either the test setup has changed or, providing no tests were conducted in the interim, the radiation conditions on the CCD vary appreciably as a function of time.

In terms of charge loss, increasing the DOB level appears to mitigate a noticeable part of the damage. Interestingly though, the charge loss increases as the delay length is decreased. This may be caused by an over-correction of the background under the faintest images, which are highly sensitive to the background component introduced by the CI tail. Determining whether the background has been over-corrected is, however, not obvious from the accumulated data profiles shown in Figure 2.97.

For $G = 11.1$ in the bright star tests, increasing the DOB level has a negligible effect. This is shown in Figure 2.114. It is difficult to differentiate even the benefits of decreasing the CI delay. With a CI and gate 9 activated, the effects of CTI are relatively minor.

Figure 2.115 shows comparable results at $G = 12.6$ and gate 11 activated. The bias and charge loss here remain extremely small. The benefits of decreasing the delay length and increasing the DOB level are clearer in terms of the charge loss.

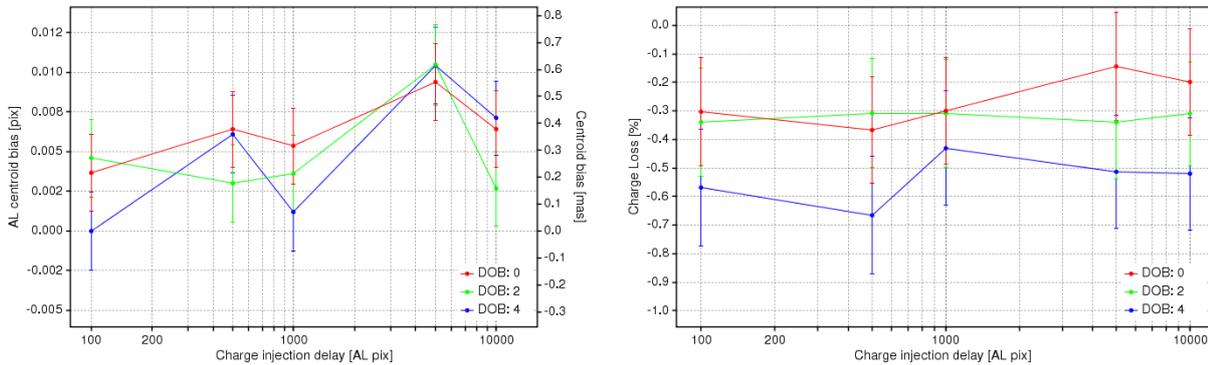


Figure 2.114: Relative centroid bias (left) and charge loss (right) as a function of CI delay and nominal DOB level at $G = 11.1$. Gate 9 is active.

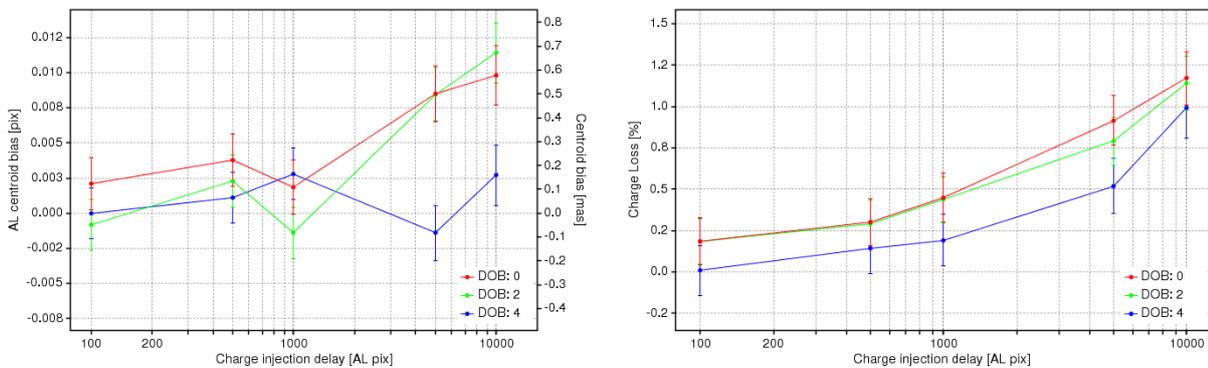


Figure 2.115: Relative centroid bias (left) and charge loss (right) as a function of CI delay and nominal DOB level at $G = 12.6$. Gate 11 is active.

Lastly, it is shown how the relative centroid bias and charge loss vary as a function of magnitude at a DOB level of 2 (see Figure 2.116) and at a DOB level of 7 (see Figure 2.117). Clearly, it is the faintest stars that suffer the most damage, but which also benefit the most from a high DOB level, as was observed in the analysis of RC1.

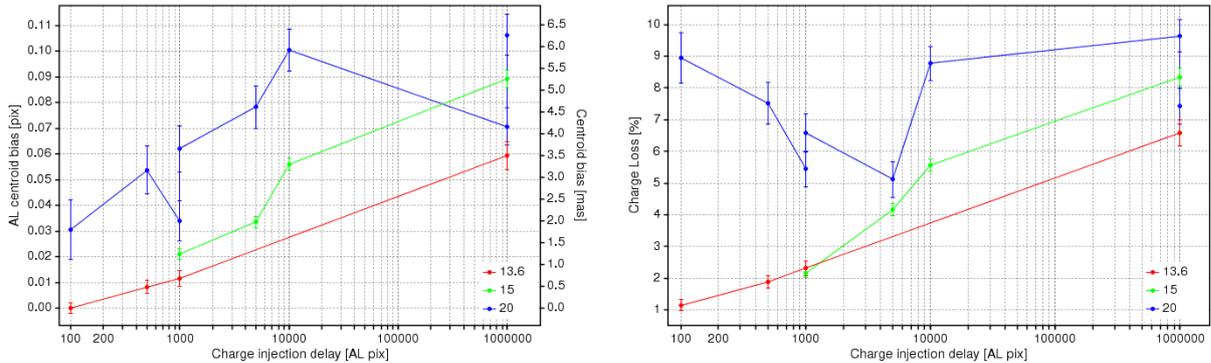


Figure 2.116: Relative centroid bias (left) and charge loss (right) as a function of CI delay and magnitude at a DOB level of 2.

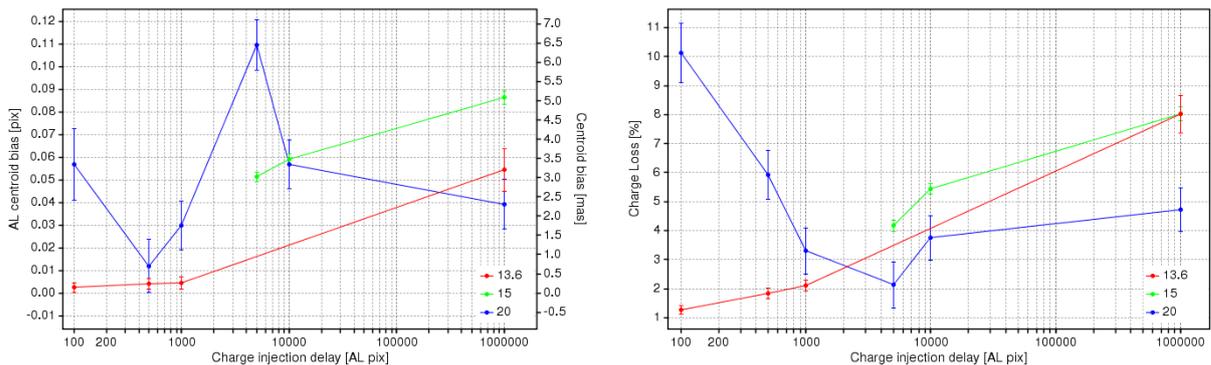


Figure 2.117: Relative centroid bias (left) and charge loss (right) as a function of CI delay and magnitude at a DOB level of 7.

Tables B.5, B.7, B.6 and B.8 tabulate most of the CTI image parameters observed in both the MF star tests and bright star tests. They can be found in Appendix B.4. Note that Astrium’s results are not compared as they used the tests with a CI delay of 10 to calibrate their results.

Image profiles: undamaged vs damaged

Taking corresponding pairs of local images from the reference and irradiated data, it is shown how the shape of the LSF is damaged by radiation-induced CTI effects and how the CI delay and DOB level may be used to mitigate this damage. Since the charge has been injected from both ends of the CCD images taken from channel 14 in column 1 are compared.

On the left of Figure 2.118 the local LSFs acquired at G = 13.6 for each CI delay at a nominal DOB of 2 are shown. As expected, the damage is worse without CI. So on the right of Figure 2.118, it is

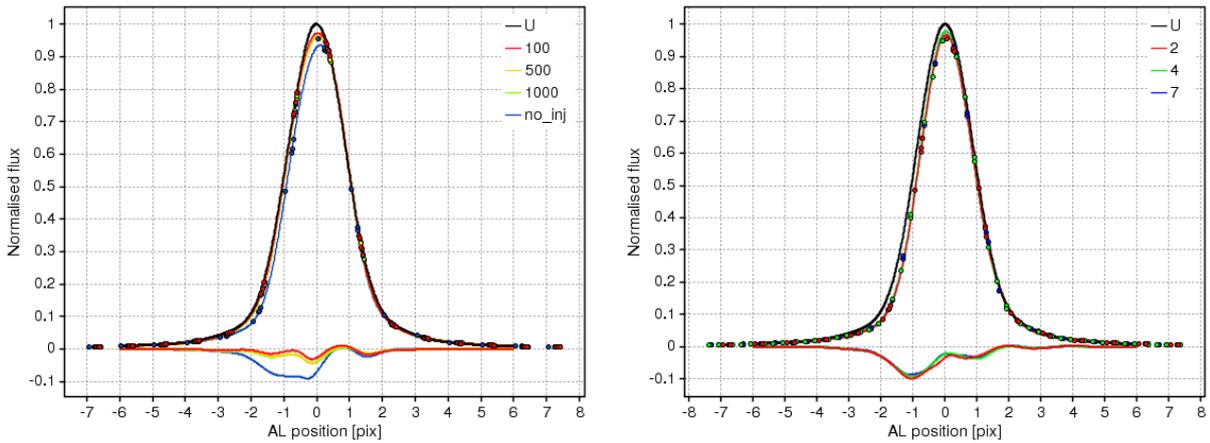


Figure 2.118: Undamaged vs damaged image profiles at $G = 13.6$. (left) Shown as a function of CI delay at a nominal DOB level of 2. (right) Shown as a function of DOB level without CI.

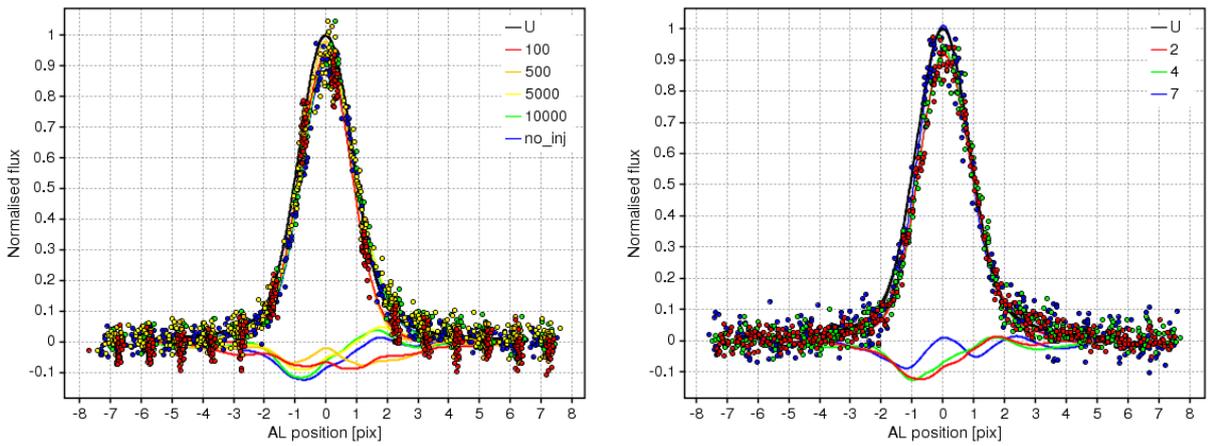


Figure 2.119: Undamaged vs damaged image profiles at $G = 20$. (left) Shown as a function of CI delay at a nominal DOB level of 2. (right) Shown as a function of DOB level without CI.

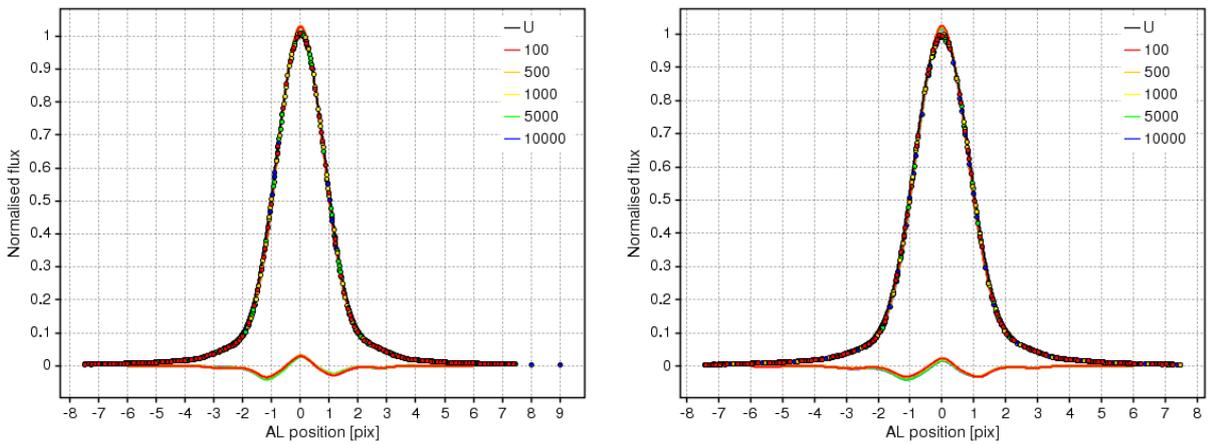


Figure 2.120: Undamaged vs damaged image profiles shown as a function of CI delay with the lowest possible DOB level at (left) $G = 11.1$ and (right) $G = 12.6$.

shown how the DOB level affects the local LSF shape when no CI has been made. Increasing the DOB makes little difference to the damage incurred.

Figure 2.119 shows comparable results at $G = 20$. Here, when a CI is made, the sample clustering tends to inhibit fitting a reliable curve to the data. Without CI, the coverage is more spread out and the action of increasing the DOB level does make more of an impact on recovering the original image.

Figure 2.120 shows results from the bright star tests. On the left, local images at $G = 11.1$ are shown as a function of CI delay and on the right comparable results are shown at $G = 12.6$. Long CI delays as well as activating gates only slightly perturb the passage of the bright images. The effect of increasing the DOB level is not shown since the impact is negligible in both cases.

Sky-like tests

As part of RC3 a series of tests were conducted using a sky-like mask to generate a more realistic stellar field. Figure 3-2 in EADS Astrium (GAIA.ASF.TCN.PLM.00415) shows a schematic of this new mask. Since the calibration of the stellar field is actually carried out by analysing the positions and fluxes of two regularly spaced sets of stars placed either side of the field (2×300), it is possible to perform a brief analysis of this subset.

For the results presented here, the reference data taken to calibrate a field with $650000 \text{ stars/deg}^2$ with the lowest possible DOB¹⁵ have been analysed. The mask was in the "mask up" position¹⁶ and 60 scans were performed.

In the CCD readout, 28 samples are recorded. The first is a prescan, while samples 2 to 6 and samples 23 to 27 contain the regularly spaced stars. Samples 7 to 22 record the sky-like field.

In the prescan the RON is found to be $5.72 \pm 0.17 \text{ e}^-/\text{sample}$ while in the IA the background noise is $\sim 6.00 \pm 0.87 \text{ e}^-/\text{sample}$. AC binning is again equal to 12, so the DOB level is $0.27 \pm 0.63 \text{ e}^-/\text{pixel}$ (see eqn. 2.4) in the IA or the lowest possible, as expected.

Figure 2.121 shows the mean background level as a function of AC position (left) and the background noise distribution in the IA (right). As observed in RC1 and RC2, the systematic variations in the mean background level are attributed to the PEM anomaly. The split in the background noise distribution suggests the noise level in some scans is systematically higher than in the majority of other scans.

Further evidence to suggest that there are distinct correlations between variations in the background is shown in Figure 2.122 on the left and right respectively. The correlations remain high between all channels, except around channels 11 and 12. The mean backgrounds in these two channels are in fact quite distinct, as seen in Figure 2.121. It may be that the jump in the background between channels 11 and 12 is not caused by the PEM anomaly but is the result of a hot pixel.

Finally, the shape of the GRC along each mask column is shown along with geometrical offsets across the mask. On the left of Figure 2.123, the GRC from every column is shown. The transition from the GRC in column 1 (red) to column 60 (violet) follows a rainbow gradient. Immediately it becomes clear

¹⁵Three densities (650000 , 150000 and $20000 \text{ stars/deg}^2$) and four DOB levels (1 , 2 or $3 \text{ e}^-/\text{pixel}$ or the lowest possible) were considered in these tests.

¹⁶Tests were performed in the "mask up" and "mask down" position, i.e. rotated 180° .

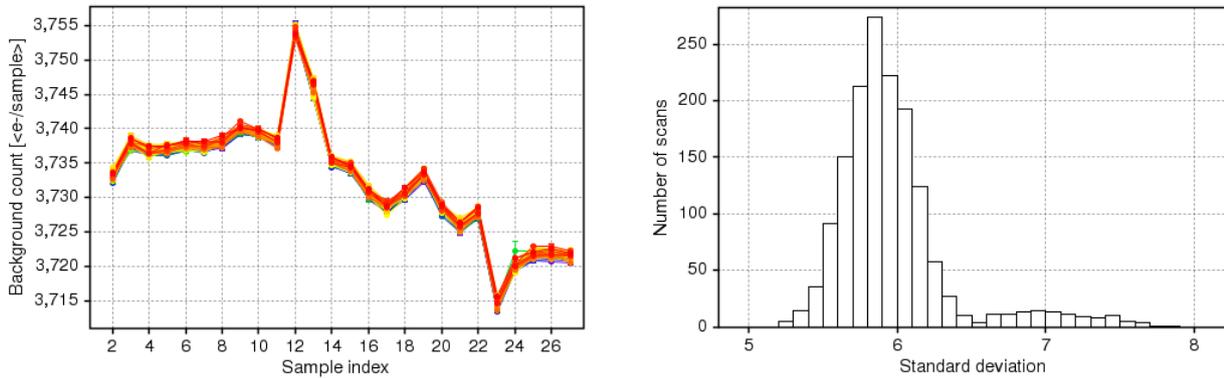


Figure 2.121: (left) Mean background level as a function of AC position. (right) Background noise distribution in the IA.

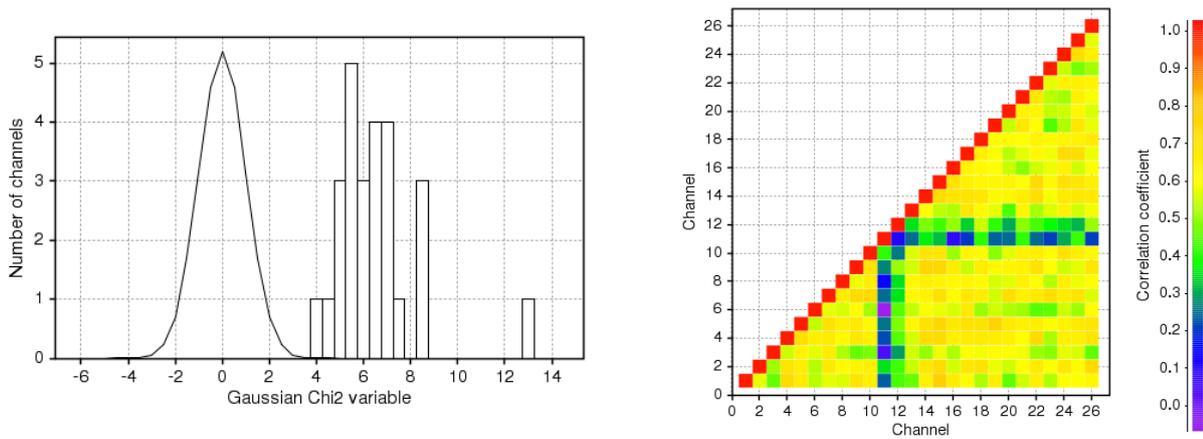


Figure 2.122: (left) Global noise distribution compared to a Gaussian of $\sigma = 1$ containing the same number of measurements. (right) Cross-correlation matrix for variations in the background level between channels.

that there are some anomalies in the base of each GRC. The GRC in column 1 could almost be said to be mirrored in column 60, though the feature on the left is more pronounced than the feature on the right. Whether these features are caused by the illumination conditions or a malformed pinhole is unclear.

On the right of Figure 2.123, the remaining displacements between the centroid positions and the geometric calibration across the mask are shown. The displacements are smaller than those observed in RC2, but larger than those seen during the regular astrometric mask tests in RC3. Since the sky-like mask is larger it will be more difficult to equalise the pressure and angle of illumination over the entire mask.

2.4.6 Conclusions

Astrometric regular tests

Some of the features observed in the raw data in RC2 have again been observed in RC3, e.g. light leaking around the screws holding the mask in place and a sharp decrease in the background signal when

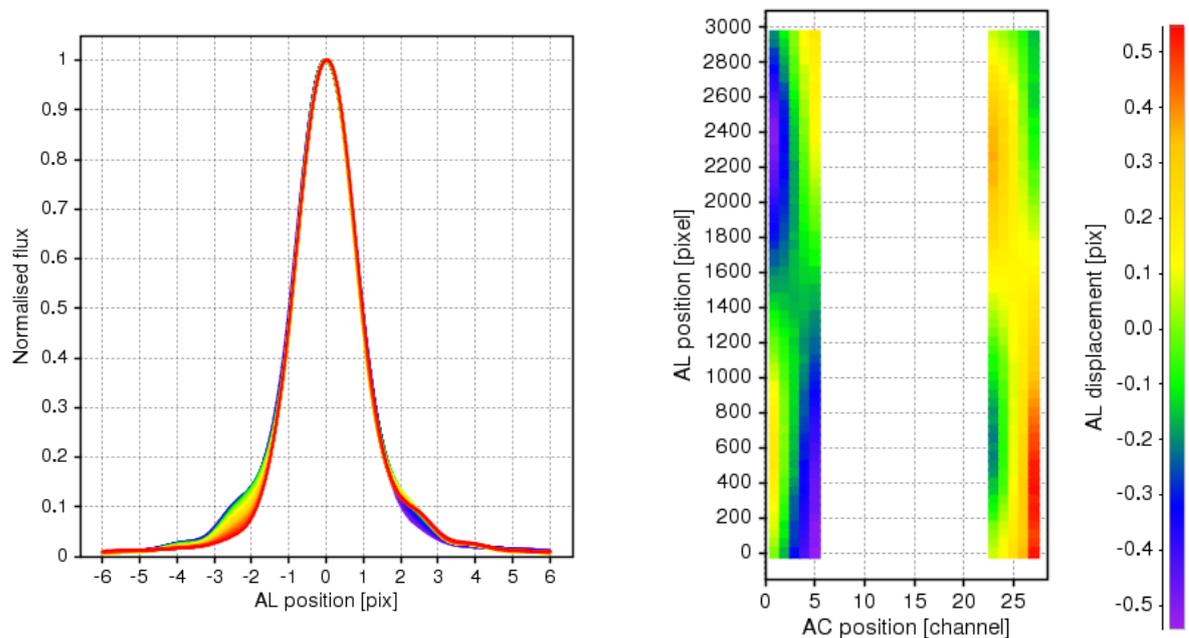


Figure 2.123: (left) GRCs from across the stellar field. (right) Offsets remaining between the calibration and the centroid positions.

the readout begins. The background continues to be affected by the PEM anomaly, which introduces distinct correlations between variations in the background levels. As this non-uniformity has also been observed with a flight PEM (in RC4), the PEM anomaly will have to be dealt with as part of the data processing.

Two new features have been found that are connected to the CI. From the top of the CCD the CI appears instantaneous, distinct from the preceding background. From the bottom of the CCD the CI and the background preceding it are connected by a small, gradual increase in the count level. This has also been flagged by Astrium and been given the name, the "divination anomaly". Astrium state that the effect has not been seen so far in tests with the flight-model PEM, so they are currently assuming it is related to the laboratory electronics. The other feature concerns the CI profile itself observed in the first and last channels read out from the CCD. The latter appears preferentially damaged. The only difference between the two profiles is the number of pixels that are passed over in the SR. Even though a part of the SR was irradiated, further tests have shown that these pixels do not have much of an effect. The damage therefore could be due to native serial CTI.

The DOB was again studied in RC3 at nominal values of 0, 2, 4 and 7 e^- /pixel in the MF star tests and at 0, 2 and 4 in the bright star tests. In the former the observed DOB levels were far higher than expected around 2, 6, 11 and 23 e^- /pixel respectively. In the latter the observed DOB levels were also slightly higher (around 1, 3 and 6 e^- /pixel respectively), and in some cases lower. Though the DOB level was set incorrectly, the relative effect of increasing the DOB level could still be analysed. As observed in RC1, it was the faintest images that benefited the most from an increase in the DOB level. The medium stars showed little improvement. This was also seen in the bright star tests where the DOB made little difference. For all CI delays with a gate activated the charge lost was never more than 1.5%. This figure

dropped to <1% for the brightest image ($G = 11.1$).

Although tests were conducted with CI delays of 10 and 30 pixels, a CI tail was discovered in the non-irradiated part of the CCD, interfering with the stellar field. Currently, it is not possible to model the profile induced by charge release from native (design) traps so these tests were not analysed.

In general, the closer a CI is to an image the less damage incurred, as long as the CI tail can be properly estimated and removed. The fact that the relative charge losses measured at $G = 20$ become progressively worse as the injection delay decreases, suggests a mis-calibration of the charge injection tail. As found in RC2, a CI delay of no more than 1000 TDI periods would appear to mitigate the bulk of images at different magnitudes. Note though that this cannot be taken as a definitive result. The optimal injection delay may have to be fine-tuned once Gaia is launched.

The geometric calibration identified a number of systematic features in the centroid positions that were attributed to CCD stitches and individual pixel-line displacements between channels. One of these features, co-located with the non-irradiated channels 21 and 22, was particularly strong and offset from the biases observed in the first 4 non-irradiated channels. For this reason, the geometry of the mask was recalibrated using only channels 1, 2, 3, and 20. As observed in RC2, the largest systematic bias was ~ 0.02 pixels measured close to a CCD stitch. However, pixel-line displacements, not observed in RC2, can also show biases as high as ~ 0.01 pixels between stitches. The centroid bias around a stitch increases if a gate is activated, as the image has less time to spread out before being read out.

Sky-like tests

A brief analysis of some of the reference data in the sky-like tests has revealed again the non-uniformity in the background and has also identified an anomalous feature in the image profiles in the regular part of the mask. The mask is larger than that used in RC2, therefore, it is expected that the illumination conditions will vary over the mask. This could be the cause of the asymmetry seen in the GRC. The geometric calibration shows a more symmetric distribution of offsets compared to RC2, illustrating better control of the mask fixtures.

Recommendations made to industry and follow-up

Following the analysis of RC3, the same recommendations arose following the analysis of RC2. However, as was described in section 2.3.7, all of these recommendations bar one, which was not necessary with CI being implemented, were carried through into the sky-like tests in RC3.

These findings were initially presented to the Radiation Task Force on Monday 26th October 2009 and summarised in a follow-up presentation on Monday 26th April 2010. Following on with the feedback from the sky-like tests in RC3, an initial analysis of the pollution mask tests in RC4 returned similar comments:

1. 50 scans have been acquired per dataset at $G = 15$ all the way up to 250 scans per dataset at $G = 20$.

2. The time between scans has been dithered such that each scan, after co-adding images, samples the LSF after ~ 0.25 pixels.
3. Reference data has been collected at the same magnitude as the irradiated data for all tests.
4. The non-irradiated channels in both the reference and irradiated tests have been carefully placed such that CCD stitches have been avoided.
5. The time between scans is similar to previous campaigns. However, in all cases there is a CI made before the stellar field.

This concludes the analysis of the radiation campaigns. Although, there are many more tests planned as part of RC4, which should be optimised according to the recommended criteria discussed above. These are discussed in the next section.

2.5 Next steps

As mentioned in the introduction to this campaign, there are other tests that were carried out as part of RC3. For example, in the last part of the astrometric regular tests the readout was left unbinned allowing, for the first time, the PSF to be viewed. These images gave a glimpse of CTI effects arising in the serial register (SR) and have since sparked another set of additional tests (EADS Astrium (GAIA.ASF.TCN.PLM.00439)). Surprisingly, it was found that native CTI in the unirradiated part of the SR introduced significant distortion to the image. The fainter the image, the worse the distortion (EADS Astrium (GAIA.ASF.TCN.PLM.00442)). In contrast, there were no conclusive results regarding the effects of radiation damage in the SR. The action of this native CTI on the location bias and charge loss as a function of G magnitude has been analysed (e.g. see de Bruijne (JDB-061)) and the results suggest only a minor across-scan charge loss is seen in the AF for faint signals around $G = 20$. Whether and how SR CTI degrades after irradiation is to be investigated in RC4 (EADS Astrium (GAIA.ASF.PLN.PLM.00073)).

Throughout the radiation campaigns the analyses has been focussed on the AF. This is not to say that the other instruments have not been tested. Running in parallel with these analyses have been other studies concentrating on the RP/BP as well as the RVS. These instruments will also suffer radiation damage and calibrating a distorted spectrum is also proving to be a major challenge.

The first batch of tests in RC4 have been carried out (see EADS Astrium (GAIA.ASF.PLN.PLM.00073) for the test plan), which focus on astrometric "pollution", i.e. how CTI effects may be polluted/alterd by signals from neighbouring stars (EADS Astrium (GAIA.ASF.TCN.PLM.00516)). Tests with the sky-like mask are also expected to be carried out throughout the remainder of 2010.

Following on from the analyses presented here, a need was identified to create a dedicated position based at ESTEC to carry on with the analysis of the radiation campaigns. Since October 2009, this position has been filled and work has commenced on completing the analysis of the campaigns for the AF. I have been liaising directly with this person, describing our approach, so that the work may continue as before.

2.6 Impact on astrometric performance

Ultimately, the errors incurred by radiation damage will have an impact on the astrometric performance of the mission, and it would be useful to see how CTI expected for the Gaia CCDs might feed through into the degradation of the astrometry. This has been hinted at in some of the plots presented in this chapter. For example, Figure 2.27 shows that at $G = 15$, with a negligible DOB the centroid bias per CCD transit is of the order of 0.21 pixels, which is equivalent to a positional error on the sky of 12.4 mas, a factor of 20 higher than the required centroiding error per transit. However, this is without charge injection. With charge injection after a delay of 1000, according to Figure 2.116, the systematic bias at $G = 15$ is reduced to ~ 0.02 pixels, a factor of only 2 higher than the expected centroiding error per transit. Taking these biases and converting them into final astrometric performance values is a possible next step. However, these numbers would not be fully representative as there is still an important factor that has not been assessed in any systematic way in the radiation campaigns to date. This is how the noise that is introduced by radiation damage affects the image shape and whether it is correlated with magnitude or entirely stochastic.

Taking a step back, according to the campaign analyses, the systematic biases on the centroids of single stars as incurred by radiation damage over the Astrometric Field appear to be removable for the most part via charge injection. In this case, it is expected that AGIS will be able to calibrate for biases on the centroid as long as the distortion of the image can be modelled well enough. Updates to the model will follow in each iteration (cf. the Intermediate Data Update) and will be aided by staggering the CI delays in the same FOV transits. Without charge injection, and again it is now known from the campaign analyses, the biases would be many times higher than required and probably too high to be removed in AGIS. With charge injection, there is justification to treat radiation damage early on in the data processing. This is where the image parameters (centroid and flux) are extracted and fed through into AGIS. AGIS will effectively check whether the forward-modelling approach (cf. Figure 1.13) to extract radiation damage has been successful and, if there are any dependencies left, propose corrections to be implemented in the next iterative loop. With successive iterations, the remaining residuals should be very small at the end of the mission and there should be no contributions from radiation damage. Note that the process of image parameter extraction will be explored more fully in chapter 5.

The centroid bias and subsequent treatment of that bias described above has not, however, considered explicitly the noise that is introduced by radiation damage, and at this stage it is not clear just how much this noise will impact on the astrometric performance. Poisson noise is the statistical limit and as such is the best anyone can expect. Modelling noise is limited effectively by two factors: a) the amount of data (brighter stars are better sampled but there are less of them); and, b) the complexity of the model(s). Gaia has a complex instrument model that must take account of the FOV, star colour, the CCD, the AC position, pixel column, and gates, which all increase the chances of dependencies arising between the parameters. So as the parameter space for the instrument calibration grows in size so too does the uncertainty on the formal errors. With Gaia, there is an additional source of noise arising from radiation damage, which is effectively caused by the stochastic nature of traps as they capture and release charge back into the conduction band of the CCD. Is this noise magnitude dependent? How does

it compare to the other sources of noise? To assess this issue it would be useful to compare a set of mean images taken at different magnitudes under the exact same test conditions with different charge injection intervals. However, due to the variations in the radiation campaigns to date (e.g. non-uniform mask illumination, sample clustering, very high DOB levels) it would be difficult to address these issues. Radiation Campaign 4 may prove otherwise. If not, checking the data (of which there will be far more) during the mission itself may be ultimately the only way to truthfully assess the impact of radiation damage on the final astrometric performance.

3

Probing an Astrometric Field CCD at the Microscopic Level

THE fourth line of attack in DPAC's radiation calibration strategy (van Leeuwen & Lindegren (FVL-006)) is to implement a software solution that may simulate the calibration of radiation damage in the data pre-processing. Simulating the damage requires an understanding of the causes underlying image distortion, which at the pixel level brings into focus the 'trap domain' on the CCD at the time of observation. Each trap species, of which there could be several, are typically described in terms of several trap parameters including the capture cross-section, the release time-constant and the trap density. One way of fixing these parameters is to compare simulated results with real observations, which is where the analysis of the radiation campaigns (RCs) comes in.

In this chapter a technique is adopted that measures the charge loss in charge injection (CI) profiles taken from RC2, both in the first pixel and over the entire profile, to probe the density of traps on the CCD. The First Pixel Response (FPR) (Janesick (2001)) is in fact a technique that can be used during the Gaia mission to gauge the level of radiation damage on the CCD and, as such, provide a further means to constrain CDM model parameters when they are updated after the Initial Data Treatment (IDT). Knowing how to perform such an experiment, the assumptions involved, and how to interpret the results, all of which will be presented in this chapter, is therefore very important for knowing what to expect later. For the original technical note on this work see Prod'homme & Brown (TPH-001).

3.1 The data: radiation campaign 2 charge injection profiles

To observe how the profile of a CI at different intensity levels is affected by CTI, the set of data corresponding to the first test in RC2 was used. This test is where the CI level was varied while the duration and delay were fixed to 20 pixels and ~ 25 pixels respectively. Table 3.1 shows the 11 CI levels that were tested in RC2. As charge was injected from only one side of the CCD, the intensity of the CI tended to vary across-scan (AC). Saturation in the readout could occur after the count level had been binned and converted into a Least Significant Bit (LSB) with the electronic offset added. If the LSB was higher than 65535, then the readout took this limiting value. In RC2 the gain was $3.56 \text{ e}^-/\text{LSB}$ so this value is equivalent to $233304.6 \text{ sample}^{-1}$ or $19442.05 \text{ pixel}^{-1}$. In fact for any $\text{CI} \leq 7.51\text{V}$ the CI profile was saturated in every channel, so these levels could not be used to determine the charge loss.

CI Level (V)	CI Level ($\langle \text{e}^-/\text{pixel} \rangle$)	Saturated?
6.51	114147	Yes
6.96	76846	Yes
7.31	52921	Yes
7.51	41477	Yes
7.84	26691	Partially
8.06	17834	Partially
8.26	11725	Partially
8.42	7973	No
8.57	5316	No
8.65	4220	No
8.72	3399	No

Table 3.1: A list of CI levels tested in RC2 and whether the read out was saturated.

For every CI level the test was repeated with a different CCD temperature. Table 3.2 shows the temperatures that were experimented with in RC2. In total, five scans were performed at each magnitude, CI level and temperature.

G	CI Level (V)	Temperature (K)				
		T1	T2	T3	T4	T5
13.3	6.51 to 7.31	161.2	171.1	188.6	198.0	-
	7.51 to 8.72	163.0	170.0	180.0	187.0	198.0
15.0	6.51 to 8.57	161.2	171.1	188.6	198.0	-

Table 3.2: CI levels and associated CCD temperatures tested in RC2. Originally, the two sets of data taken at T2 and T3 for CI levels between 7.51V and 8.72V at $G = 13.3$ were reported to have been at the same temperature: 180K. However, distinct differences between the two sets of data were observed, which indicated T3 was actually higher than T2. Since the original publication of this work Astrium have clarified that in fact T2 was 170K and T3 was 180K.

Figure 3.1 shows two examples of the way in which CI profiles at different levels can appear damaged after they have passed over a CCD at different temperatures in RC2. Clearly, the first pixel in both cases is far more damaged compared to the pixels that have been read out at the end of the CI. In general, if two

CI profiles are taken from the same channel and with the same CI level, the higher the temperature the more damaged the CI profile appears. As can be seen on the right of Figure 3.1 with a high temperature the leading edge of the CI struggles to transit the CCD without appearing almost completely depleted.

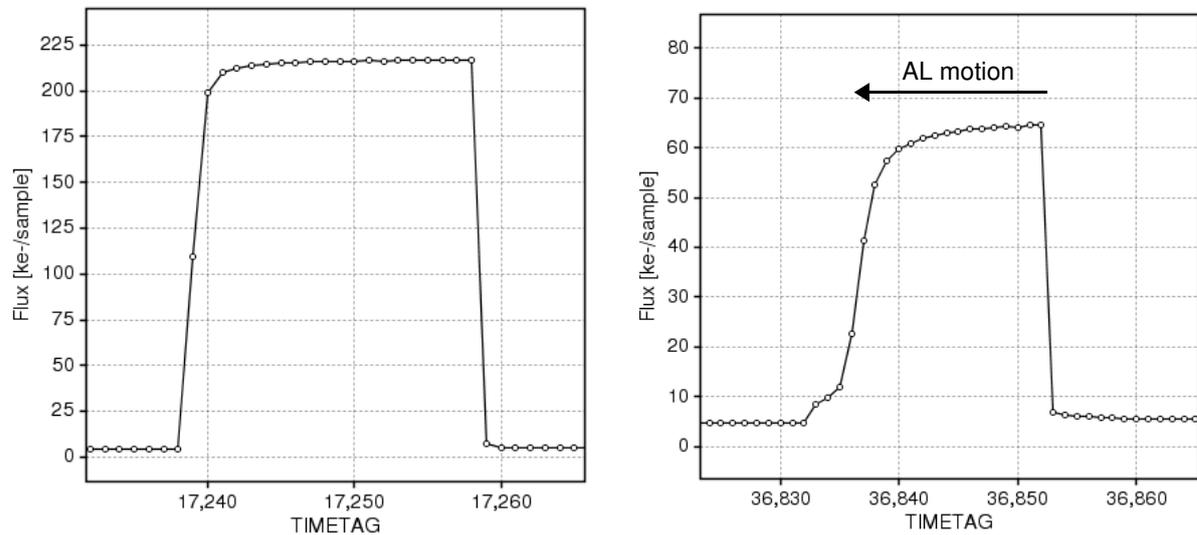


Figure 3.1: Two examples of damaged CI profiles from RC2. In both cases scan 2 through channel 18 is shown. On the left, the CI level is set to 8.26V and the CCD temperature is 163K and on the right, the CI level is set to 8.65V and the CCD temperature is 180K.

In the results that follow the CI profiles have been taken from the $G = 13.3$ test in RC2. Plots extracted from CI profiles accompanying $G = 15$ images show similar results. These can be found in Appendices 1 to 3 in Prod'homme & Brown (TPH-001).

3.2 Trap densities

Currently it is a requirement to fix certain properties of traps present in a Gaia CCD in order to model radiation damage in the data. The density of traps is one such property. By determining the charge loss over two distinct parts of a CI profile it has been possible to measure intervals for the number of traps per pixel as a function of both temperature and signal level. Why intervals and not a unique value will be shown shortly.

3.2.1 First Pixel Response

The First Pixel Response (FPR) of the CI profile is a value that can be used to probe the trap density on a Gaia AF CCD. In any particular pixel column, the first pixel of the CI profile is the first signal that will interact with the population of empty traps. As a consequence, it is the charge loss in this pixel which is more likely to indicate the number of traps in that pixel column. Given certain assumptions shortly to be explained, the observed charge loss can be equated to the number of traps per pixel.

In Chapter 2 it was shown how charge was released back into the background following a CI. A

charge release function was used in the analysis of RC3 to model the CI tail. This model was based on the assumption that a CI tail corresponds to a linear combination of exponentials, each representative of a different trap species. Unfortunately, the process of fitting the tails from experimental data is somewhat degenerate, since one can arbitrarily add species to obtain a better fit, or obtain the same quality of fit with different parameters. These uncertainties have been pointed out in several reports (cf. Hopkinson (254.DO.28), and Hambly & Cross (NCH-003)). In contrast, FPR does not consider the charge released behind the CI but the charge that is captured at the beginning of the CI. In this way, the same parameters may be investigated without the need to find a reliable fit.

FPR is a technique that must follow a precise protocol (cf. Janesick (2001)). However, the RCs were not designed to achieve the conditions required. For example, the protocol specifies that a CCD must be exposed to a flat field followed by several full frame readouts before a CI is made. This is so that the CCD can reach a state of unbiased equilibrium before extracting the FPR. Here, it is assumed that this condition has been met, since only those scans that were made after scan 1 are considered, i.e. it is assumed that the CCD has reached a state of equilibrium after scan 1. In RC2 scan 1 is known to have experienced a different set of conditions compared to all the scans that follow, so the results shown in the following sections concentrate on the average values over scans 2 to 5. Another key aspect of the FPR protocol is the decoupling of the FPR in the image area (IA) and the serial register (SR), so as not to confuse parallel CTI with serial CTI measurements. Unfortunately, this implicit assumption cannot be completely justified here given what is now known about the effects of CTI in the SR (see section 2.5). Aside from the IA and SR both being damaged in RC2, native serial CTI is also present and must be taken into account when interpreting the final results. Native serial CTI is expected to become more pronounced as the injection level decreases; the effects of radiation-induced SR CTI effects are yet to be studied as part of RC4.

When a CI is made, it is clocked across the CCD, filling traps along the way. The leading edge of the CI will be the first to encounter the bulk of traps. Providing that the electrons belonging to the leading pixel fill most of the empty traps, the pixels that follow will see less empty traps. The first pixel is therefore the best probe of the trap density. Assuming one electron fills one trap and traps extracting charge from the first pixel do so without releasing that charge back into the first pixel¹, the charge lost in the first pixel can be equated directly with the density of traps in that CCD pixel column². The number of traps per pixel is then simply the charge lost divided by the number of pixels in that column. So, for one pixel column:

$$N_t = \frac{\Delta C}{N_{pAL}}, \quad (3.1)$$

where N_t is the mean number of detected traps per pixel, ΔC is the charge loss in one pixel (preferably

¹This is an approximation because the release probability $(1 - e^{-t/\tau_r})$ for traps with a release time-constant greater than 10 TDI sample periods ($\sim 982.8\mu s$) is 8% when $t = 1$ TDI period.

²Even though FPR is concerned with the charge loss in a pixel, in reality, a Gaia pixel has a 4 phase structure. So the time each charge packet remains in the same location on the CCD, close to any particular trap, is of the order of a 1/4 of the TDI period. This is known as the dwell time. At the electrode level, this information is useful. However, the locations of traps within each pixel are of no concern here, just their density. So the results from this study would probably be the same as those that came from an electrode-level study.

the leading pixel), and N_{pAL} is the number of pixels in the AL direction in the IA³.

$$\Delta C = E_{reference} - E_{measured}, \quad (3.2)$$

where $E_{reference}$ is the expected electron count for an undamaged profile, i.e. the number of electrons injected, and $E_{measured}$ is the measured electron count. If ΔC is measured in e⁻/sample it must also be divided by a binning factor, i.e. the number of pixels that constitute one sample. For example, in RC2 each sample was comprised of 12 pixels binned AC.

For a trap with a release time-constant τ_r close to the TDI sample period, the probability an electron will be emitted within the same pixel ($t \approx \tau_r$) is close to 0.6. With such traps, a disproportionate number of them may not be picked up by the FPR, but will have a cumulative effect on the rest of the CI profile.

As can be seen in eqn. 3.2, calculating the charge loss requires a reference level, $E_{reference}$, which in RC2 is not known. No reference data exists that measures the level of the CI for each of the levels in every AC position seen here. To compensate for this, it is assumed that the last n pixels of the CI, which lasts for 20 pixels, are undamaged. Setting n equal to 3, 4 or 5 did not change the results, so for now $n = 4$. In the AC direction, since the readout is binned AC by 12 pixels and the problems with serial CTI are known to exist, the profile of the CI is also unknown. The calculated reference levels must therefore be taken as first-order approximations, especially when the charge injection level becomes faint (serial CTI increases as the signal level decreases, though this may be mitigated to some extent in the last samples in the injection by the preceding samples via self-injection).

3.2.2 Reference levels

Figure 3.2 and 3.3 show the calculated reference levels (e⁻/pixel) for each of the CI levels analysed at 163K, the Gaia operating temperature, and 198K, the maximum temperature tested respectively. Errors are representative of the standard deviations on the calculated mean. Recall that scans 2 to 5 have been averaged to determine these levels.

Summarising the results on the observed reference levels at all temperatures (T1 to T5):

- As the charge is injected from only one side of the CCD, the CI level varies AC. Channel 1 is the furthest from the injecting source and registers the lowest value.
- The electronic offset added by the PEM saturates the signal readout for a CI level of 8.26V in channels 20, 21 and 22.
- The average reference level calculated over all channels for each voltage is in good agreement with the expected values given in Table 3.1.
- Similar reference levels are found for the same tests with a $G = 15$ source, indicating a degree of stability in the test setup.

³In the results that follow this value was set to 4494. However, it should have included the shielded pixels defined in the GPD (Satellite:CCD_TDILine_Blocked = [1,2,5,6,9,10]). Note then that the trap density will be slightly over-estimated.

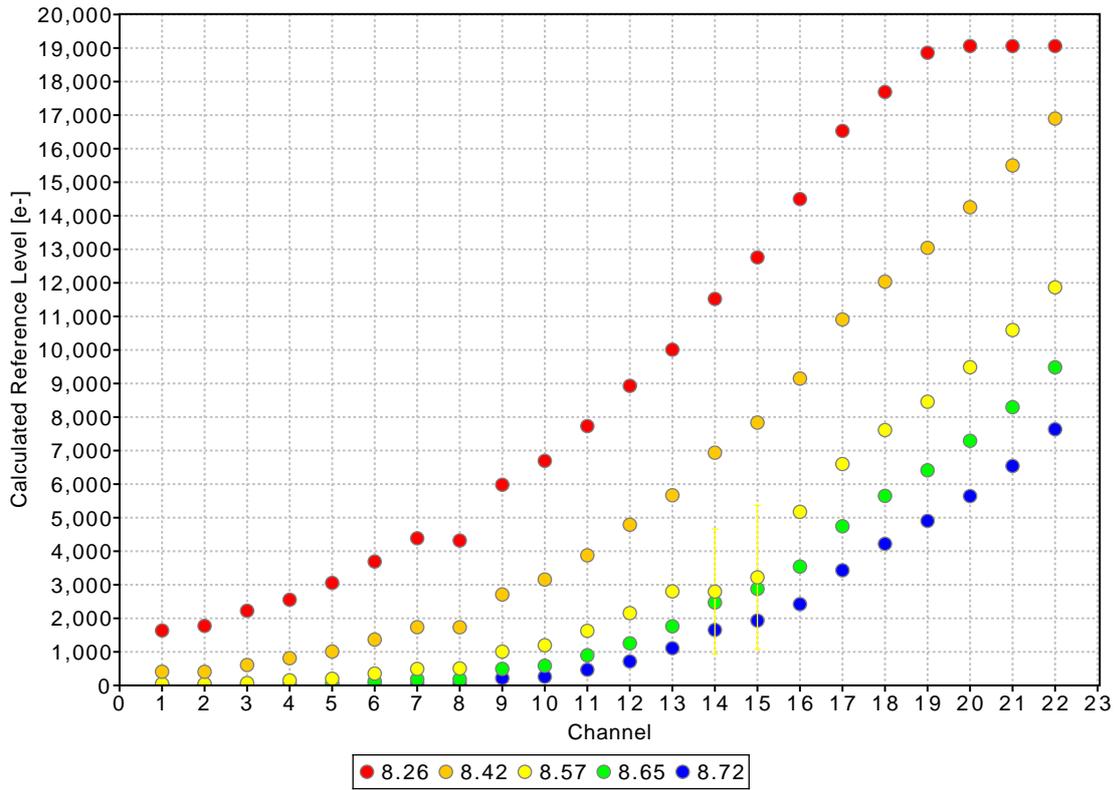


Figure 3.2: Calculated CI reference levels at 163K (T1).

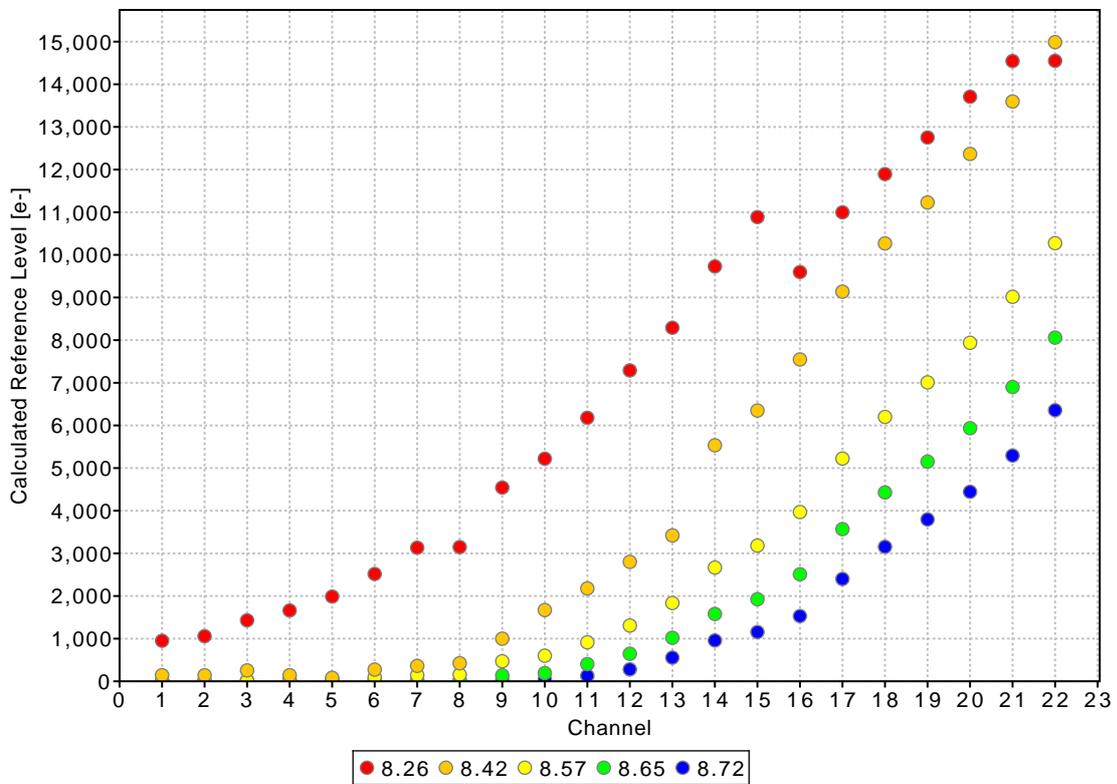


Figure 3.3: Calculated CI reference levels at 198K (T5).

- From T1 to T5 a general decrease in the reference level is seen. This is probably caused by an increase in CTI in the image area as observed by Sira (which is why a temperature of 163K (T1) was chosen for the CCD).
- Though fluctuations are seen in the saturation level at different CI levels at T5, the saturation level for 8.26V ($\approx 14550e^-/\text{pixel}$) is less than the equivalent saturation level at T1 ($\approx 19000e^-/\text{pixel}$). This may be linked to the PEM anomaly.
- In some test cases the reference level can be seen to decrease for several levels of CI in channels 14 and 15. These particular channels lie adjacent to a CCD stitch.
- Why there is a drop in channel 8 is less clear. Figure 2.67 shows that there may be a higher degree of damage in this channel compared to the surrounding channels.
- The reference levels at 8.26V at T5 drop sharply in channel 16 and above. As this is a localised anomaly it is assumed that these levels are faulty. They are not included when interpreting the results.

3.2.3 Results

The FPR and, therefore, the number of traps per pixel is shown in Figure 3.4 as a function of reference level and CCD temperature. Errors are again representative of the standard deviation. Those datapoints with relatively large errors are attributed to faulty scans. The datapoints on the right of the plot that appear to measure different trap densities for a similar reference level, originate from the readout saturating. Interestingly, the number of traps per pixel increases as the signal level increases, rather than the density of traps remaining constant throughout the CCD. At the Gaia operating temperature of 163K (T1) the number of traps varies between 0.0040 ± 0.0003 and 2.0637 ± 0.0003 traps/pixel. The latter value seems reasonable and agrees closely with the different analyses performed so far on the CI tail (e.g. see GAIA.ASF.TCN.PLM.00071). In any case, an interval must be given for the trap density. Results at $G = 15$ show similar values. Appendix 4 in Prod'homme & Brown (TPH-001) shows comparable results for scan 1, where the trap density appears to be slightly higher (e.g. $N_t = [0.0039; 2.3976]$ traps/pixel at 163K). This can be explained by having traps with release time-constants longer than the time between scans (~ 30 s) initially unfilled. After scan 1, these traps remain filled for the duration of the test.

Compared to the other temperatures, the FPR at T5 contains some seven datapoints that behave differently to the others, dropping out just below a calculated reference level of $15000 e^-/\text{pixel}$. These particular values correspond to the same seven values shown in Figure 3.3 for T5 at 8.26V. It is not clear what the cause of this drop is and why the signal saturates at a reduced level. However, since it only occurs in these channels at this temperature, it was assumed these levels were faulty (calculated wrongly) and not representative of the FPR.

The underlying reason(s) as to why the trap density should depend on the signal level are currently not very well understood. Figure 3.4 suggests that the first pixel of the CI does not fill all of the traps throughout the pixel. This is probably because at small signal levels the density maximum of the electron

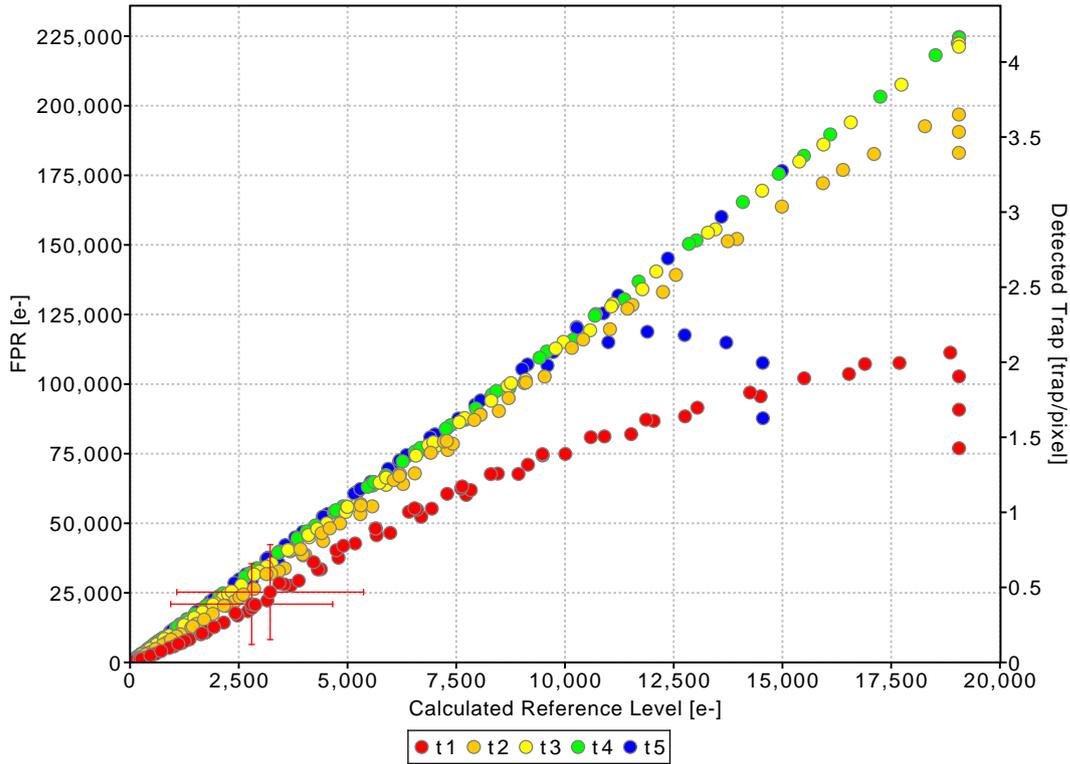


Figure 3.4: FPR and trap density as a function of reference level and CCD temperature.

cloud is too far away from some of these empty traps. If this were the case, the trap density should saturate (i.e. all traps are filled) when the CI level is close to the pixel Full Well Capacity (FWC). It would also imply that the density of the electron cloud decreases very sharply towards the edge of the signal confinement volume. Recently, the best modelling results, especially when interpreting the effects of the DOB, have been obtained with a homogeneous electron density (cf. Short (AS-015)), not a sharply defined boundary. Therefore, it is more likely that both density and confinement volume vary simultaneously.

Figure 3.5 shows the relative FPR ($\text{FPR} / \text{Calculated Reference Level}$) as a function of reference level and CCD temperature⁴. In certain areas the response is distinctly non-linear, suggesting again that a volume-only driven model does not describe the full picture.

One other approach to performing an FPR experiment could be to extend the exposure or dwell time in each pixel. This would test whether the same signal level could fill more traps just by clocking charge more slowly across the CCD.

If traps with release time-constants less than the time it takes a CI to be made (20 TDI periods ($\sim 20\text{ms}$)) are ignored, then the traps that were not filled by the CI in the first pixel could be filled by the CI in the pixels that follow. Charge loss measured across the entire CI profile could, therefore, also be indicative of the trap density on the CCD. This is explored in the next section.

⁴Note that in Figure 3.4 the calculated reference level is measured in e^-/pixel whereas the FPR is measured in e^-/sample (1 sample = 12 pixels binned AC)

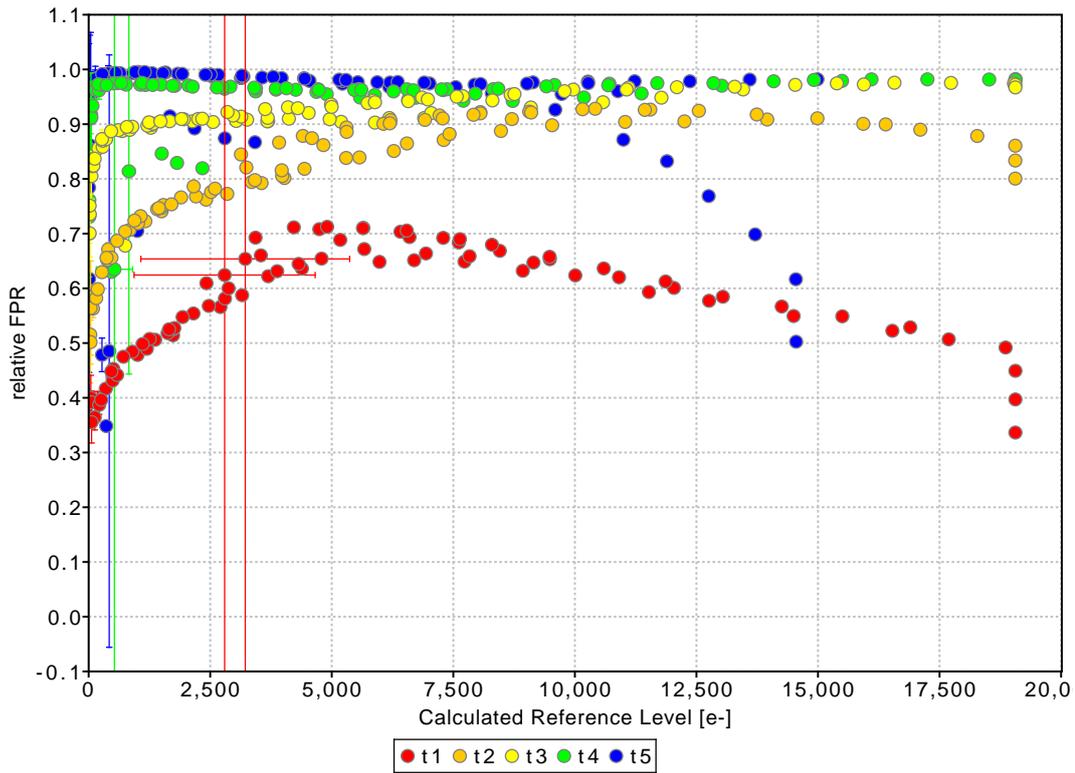


Figure 3.5: Relative FPR as a function of reference level and CCD temperature.

Table 3.3 summarises how the trap density varies as a function of temperature. In general, the number of traps detected increases as the temperature increases, at least as far as 198K. Why this should be can be inferred from highlighting those parameters that are temperature dependent; namely the release time-constant and the capture cross-section. If, at higher temperatures, the release time-constant for certain traps becomes comparable to the time domain the FPR is probing, more electrons may be captured. If the probability of capture increases with temperature, traps located far from the electron cloud density maximum may have more chance of capturing electrons at a higher temperature. Similar results are found in the $G = 15$ tests.

G = 13.3 test.			
Temperature		Number of detected traps	
Reference	Value (K)	Minimum [\langle trap/pixel \rangle]	Maximum [\langle trap/pixel \rangle]
T1	163.0	0.0040 ± 0.0003	2.0637 ± 0.0003
T2	170.0	0.0036 ± 0.0006	3.6500 ± 0.0020
T3	180.0	0.0032 ± 0.0002	4.1260 ± 0.0020
T4	187.0	0.0017 ± 0.0001	4.1667 ± 0.0006
T5	198.0	0.0011 ± 0.0005	3.2750 ± 0.0010

Table 3.3: Average number of detected traps as a function of temperature using the FPR.

3.2.4 Total charge loss

Figure 3.6 shows the *total* charge loss as a function of the reference signal level. As expected, the trap density increases as the reference signal level increases. On the far right some datapoints drop sharply for the same signal level. This is a saturation side effect as observed previously. Compared to Figure 3.4 non-linearity appears to set in at a lower signal level. No saturation is seen in the trap density either, implying the FWC has not been reached.

As with the FPR analysis, Table 3.4 summarises the temperature impact on the trap density determined via the total charge loss. Negative values for the average number of detected traps per pixel are not physical. They are caused by an improper estimate of the reference level, and the invalid assumption that at very low signal levels the number of electrons released is not relevant during one TDI step.

G = 13.3			
Temperature		Number of detected traps	
Reference	Value (K)	Minimum [$\langle \text{trap/pixel} \rangle$]	Maximum [$\langle \text{trap/pixel} \rangle$]
T1	163.0	0.02 ± 0.70	2.80 ± 0.02
T2	180.0	0.02 ± 0.40	5.61 ± 0.01
T3	180.0	0.02 ± 0.80	7.85 ± 0.04
T4	187.0	-0.50 ± 0.90	8.37 ± 0.04
T5	198.0	-13.00 ± 4.00	4.02 ± 0.02

Table 3.4: Average number of detected traps as a function of temperature taken using the total loss.

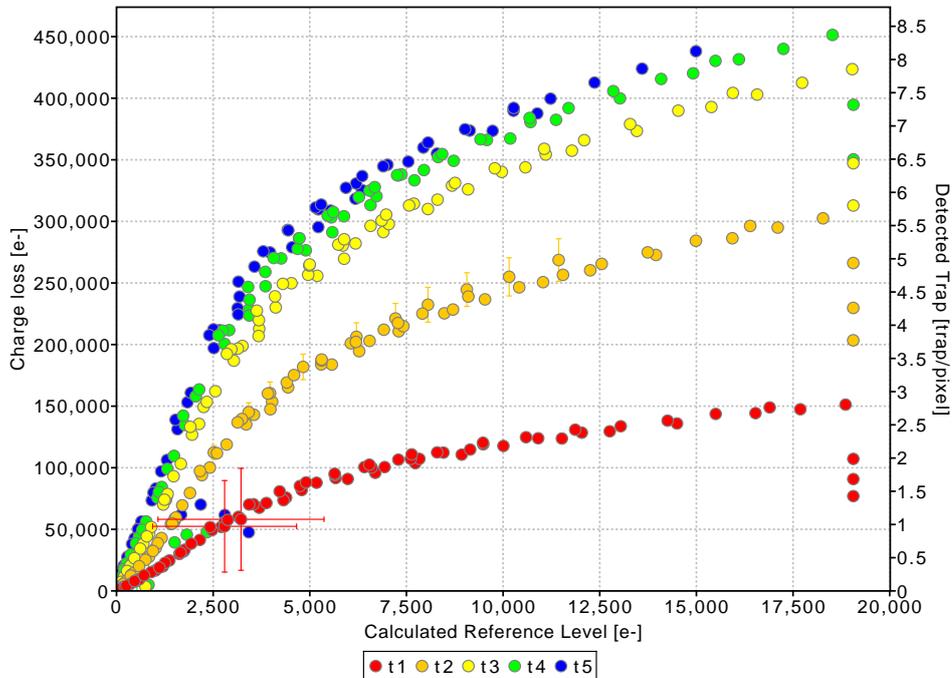


Figure 3.6: Total fractional charge loss as a function of reference level and temperature.

3.3 Measuring the Supplementary Buried Channel Full Well Capacity

Measuring the total charge loss in a CI profile provides not only information on traps but also reveals one of the most important structural features within a Gaia CCD: the Supplementary Buried Channel (SBC). The SBC is included as part of the Gaia CCD design in order to channel small signals through the CCD so that they encounter fewer traps. As shall be shown, without the SBC the level of damage would increase significantly as stars become fainter.

Why is it important to know the Full Well Capacity of the SBC? Primarily, it tells us how extensive the magnitude range of Gaia will be. At this transitional point, the effects of radiation damage are suppressed, so there is an increased chance these faint objects will be identifiable in the final catalogue.

It is usually very hard to measure the FWC of this narrow channel under the main signal channel, the Buried Channel (Buried Channel (BC)), because signal packets containing more electrons than the SBC FWC overflow into the BC. Physically speaking, the SBC potential is slightly lower than the BC potential, which attracts and retains small signals. A schematic of the pixel potentials in a Gaia CCD is illustrated in Figure 3.7.

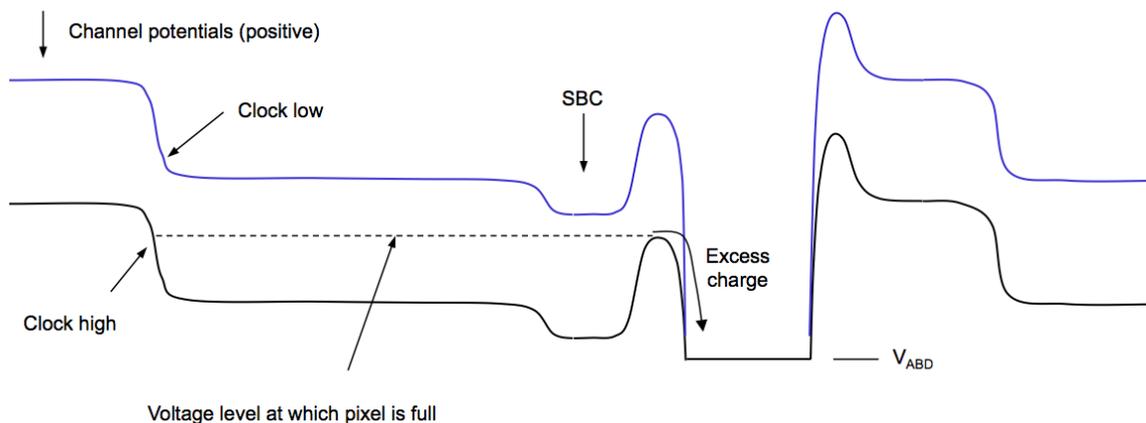


Figure 3.7: Gaia CCD pixel potential. Any excess charge falls into the antiblooming drain (ABD). Credit: Seabroke et al. (2008).

Figure 3.8 shows the fractional charge loss as a function of signal level at 163K. There is indeed a distinct reduction in the charge loss at low signal level, which can provide an indirect measurement of the SBC FWC. Hopkinson (254.DO.28) produced similar measurements (cf. Fig. 8.2-2 to 8.2-4), which led to the same conclusions. Note though that the levels of irradiation they tested ranged between 1.8×10^9 to 2×10^{10} protons per cm^{-2} , not just 4×10^9 protons per cm^{-2} as was used in the radiation campaigns.

What would happen if the SBC were not included? This situation is also highlighted in Figure 3.8 by an extrapolation of the charge loss observed in the BC as the signal level decreases. Even though it is purely an indication, it is clear that without the SBC, the charge loss would continue to increase significantly, and the fainter images beyond $G \approx 19$ would suffer to the extent that they may not be retrievable. The analysis of RC1 has shown that radiation damage can remove tens to hundreds of electrons from faint images ($G \approx 19 - 20$) depending on the level of the DOB. Knowing the SBC FWC therefore allows us to know just how far the Gaia catalogue will go.

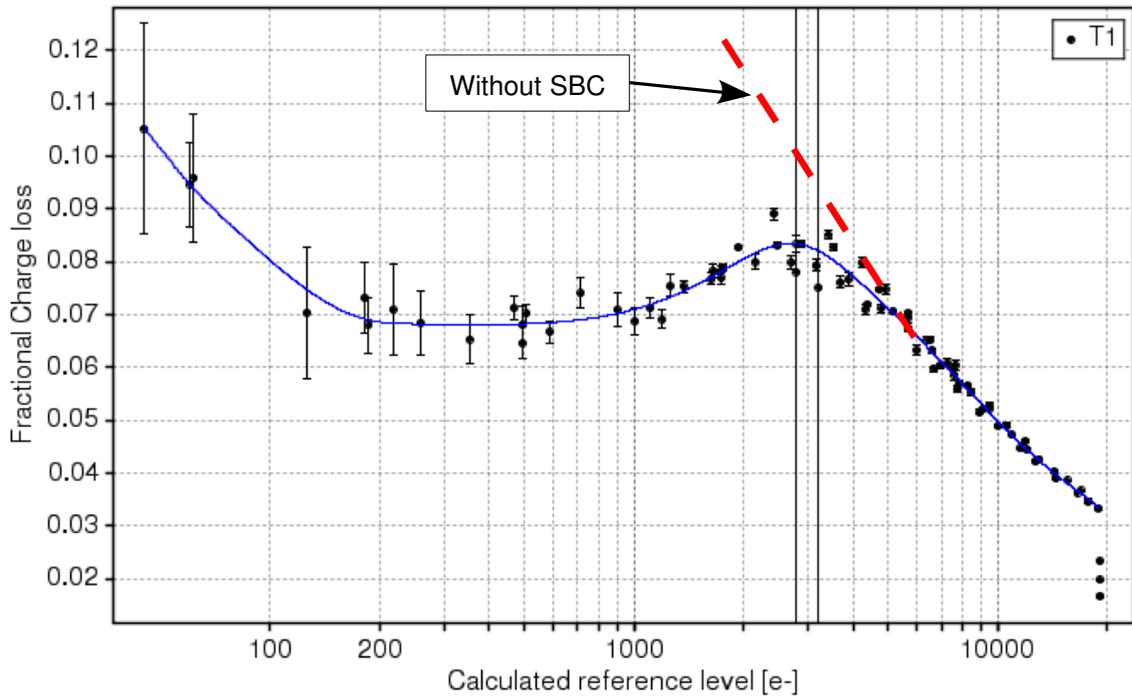


Figure 3.8: Fractional charge loss at T1 (163K) with a cubic spline fit. The red-dashed line simply gives an indication of how the charge loss could increase if the SBC were not present.

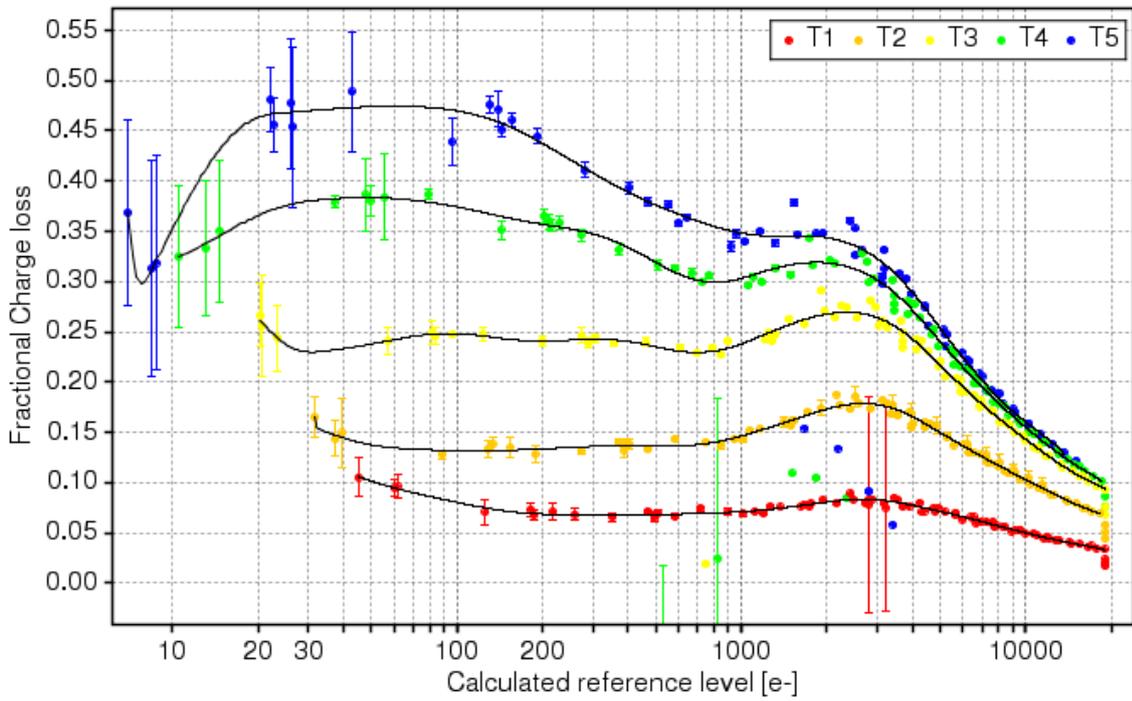


Figure 3.9: Fractional charge loss at all five temperatures with cubic spline fits to each.

The transition between the BC and the SBC can be seen occurring around a signal level of $\sim 2500 e^-$. Estimating the size of the SBC FWC can be done at two locations: a) where the charge loss is minimal, i.e. when the bulk of electrons are resident in the SBC⁵; and, b) when the mitigating effects of the SBC overcome the characteristic charge loss seen in the BC as the signal level decreases. To approximate the latter point it is assumed that the SBC takes over when the amount of charge is almost equally shared between the BC and SBC⁶. To extract these values a cubic spline function was fitted to the datapoints shown in Figure 3.8 using a logarithmically spaced 12-knot sequence and interpolating the minimum and maximum of the curve at these two locations respectively. The fits for all five temperatures are shown in Figure 3.9. Clearly, the temperature is having an effect upon the fractional charge loss⁷.

It appears the location of the minimum in fractional charge loss in Figure 3.9 fluctuates with temperature. Moreover, the location of this minimum appears to become more and more extended as the temperature decreases. This could partly be explained by the non-homogeneity of the SBC in the AL direction. Due to manufacturing constraints, e2v, the company building the Gaia CCDs, increase the width of the SBC in the first half of the CCD from $3\mu\text{m}$ to $4\mu\text{m}$ up to TDI line number 2340⁸. After this, in the second half of the CCD (the closest to the SR), the SBC width supposedly remains at a constant value of $4\mu\text{m}$ (see e2v technologies (GAIA-E2V-TN-020)). This makes it difficult to extract a single value for the SBC FWC using only the location of the minimum in fractional charge loss.

Table 3.5 summarises the maximum and minimum values extracted from these plots. An approximation for the depth of the SBC FWC is made using the half of the maximum value as described above. The GPD states that the FWC of the SBC is $1300 e^-$, which is in good agreement with the results at T1 and T2.

Temperature	Fractional charge loss		Calculated reference level [e-/pixel]		SBC FWC [e-/pixel]
	Minimum	Maximum	Minimum	Maximum	
T1	0.068±0.001	0.083±0.001	374	2680	1340
T2	0.132±0.004	0.179±0.002	113	2675	1338
T3	0.235±0.003	0.269±0.002	697	2349	1175
T4	0.300±0.004	0.319±0.002	812	1847	924
T5	0.345±0.003	0.346±0.003	1352	1639	820

Table 3.5: Values for curve minima and maxima extrapolated from Figure 3.9. Formal errors from the spline fittings are given for the fractional charge losses. For T5 the values for the local minimum situated just to the left of the SBC/BC interface are given in the table. Values for the second local minimum on the far left of the figure are 0.298 ± 0.008 and 8 for the fractional charge loss and calculated reference level respectively. The FWC of the SBC is estimated to be half of the maximum calculated reference level.

⁵This is what is assumed in Hopkinson (254.DO.28). The minimum in the fractional charge loss is "interpreted as" the point at which the SBC is full.

⁶This is an intuitive approach based on the simple "two-box" idea: it is assumed that the charge loss starts to decrease when there is a greater amount of charge in the SBC as opposed to the BC. As this may not be the case, the interpretation remains open for discussion.

⁷Communication with G. Seabroke suggests the following: according to eqn. 2 in Seabroke et al. (GMS-001) the SBC potential is not temperature dependent and remains effectively fixed. This suggests that the result seen in Figure 3.9 may be due to the volume of the charge packet increasing with temperature, which then overflows the SBC at lower signal levels. Further simulations are required to confirm this.

⁸As the CCDs are stitched together there is a need to avoid charge collecting at the single AC CCD stitch (see Figure 2.64).

It is important to recall that due to the lack of reference data it is impossible to determine accurately the reference signal level. As stated in section 3.2.1, the number of e^-/pixel were calculated according to several assumptions. To test the validity of these assumptions, a Monte-Carlo pixel-level model for CTI effects was implemented using the CTI Effects Models for Gaia (CEMGA) platform⁹. A comparison between the fractional charge loss measured in the test data and in the simulated data was made. The results of this simulation are shown in Figure 14 in Prod'homme & Brown (TPH-001). It suggests that the method used to determine the CI reference level may well have introduced a slight systematic bias that diminished as the signal increased¹⁰.

3.4 Conclusions

Using charge injection profiles from the first test in RC2 it has been possible to derive certain physical parameters of a Gaia CCD that are required in order to better understand the trap domain Gaia may experience, and help DPAC find a means to implement a better constrained radiation calibration strategy.

Firstly, using the First Pixel Response technique, intervals have been determined for the trap density as a function of temperature. Intervals had to be given rather than a single unique value, since the number of detected traps depended on the signal level. Coupled with the fact that this dependence was found not to be linear at very low and very high signal levels suggests that (along with other modelling efforts) both density and signal confinement volume evolve with the signal level. The maximum trap density averaged over scans 2 to 5 at 163K was measured to be 2.0637 ± 0.0003 trap/pixel. This is in good agreement with charge injection tail analyses, which instead use curve fitting to estimate trap densities. As expected, the trap density in scan 1 is slightly higher: $N_t \approx 2.40$ trap/pixel, due to a number of traps being unfilled before the test began. These traps become filled after scan 1 and remain filled by the time scan 2 begins.

More generally, the charge loss in the first pixel as well as the charge loss over the whole injection profile was found to increase with temperature over the tested range (between 163K and 198K). Assuming these traps remove and retain charge beyond the first pixel, the FPR preferentially probes trap species that have release time-constants $>10 \times \text{TDI}$ period. As the release time-constant varies with temperature, the drop in the charge loss seen at 198K might be due to a release time-constant shifting below the TDI period for one or several of the trap species involved.

For the first time for a Gaia CCD the Full Well Capacity (FWC) of the Supplementary Buried Channel (SBC) has been measured. This is an important parameter to know as it essentially flags the magnitude at which Gaia will continue to observe less damaged images. Signals identified at or below the FWC will see less traps and the level of radiation damage should be suppressed up to and slightly beyond this magnitude. In other words they should be flagged by the Sky Mappers, as there is a chance they could be recovered. The SBC has been observed by analysing the total charge loss at very low signal levels. The

⁹CEMGA is a Java software package being developed at Leiden University by T. Prod'homme and A. Brown that can perform Monte-Carlo pixel level modelling of the transits of images through damaged CCDs over each of the three instruments (AF, BP/RP and RVS) onboard Gaia. See Prod'homme (2010) for a fuller discussion of this modelling platform.

¹⁰T. Prod'homme also states that at the time this simulation was made the electron density distribution variation with signal level was poorly described. This could explain why there is a strong discrepancy between the data and the model in the SBC signal regime.

FWC has been measured using both the local minimum and maximum of a spline fit, corresponding to a full and half-full SBC respectively. By dividing the maximum reference level by two at T1, where half the charge is assumed to reside in the SBC and the other half in the BC, the SBC FWC measures $\sim 1.1 \times 10^3 \pm 0.2 \times 10^3 e^-/\text{pixel}$. According to eqn. 2.2, this is equivalent to a magnitude of ~ 19.7 . The nominal value for parameter in the GPD is $1300 e^-/\text{pixel}$. This is equivalent to a magnitude of ~ 19.5 .

Lastly, in section 3.2.1 it was mentioned that conventionally FPR is indicative of the damage observed purely in the image area. However, in RC2 the serial register was damaged alongside the image area so can this assumption be justified? In RC3 no discernible effects could be seen over this small, damaged part of the CCD. However, a notable native serial CTI (i.e. CTI in the unirradiated area of the CCD) was observed. This would suggest the serial register cannot be ignored, even though self-injection in the serial register may shield those profiles that were read out later. At this stage, it is not clear how much of an effect the serial register has had on the results. To reconcile this, in future FPR laboratory tests, it may be useful to inject a series of closely-packed very intense charge injections through the CCD before repeating the analysis described here.

4

Charge Distortion Model Validation

PART of the DPAC's current radiation calibration strategy involves implementing a 'forward-modelling' approach to handle radiation damage in the data processing. This was introduced in section 1.9.4 in Chapter 1. A core component in this approach is the Charge Distortion Model (CDM), which predicts how an image will be distorted given the current state of the CCD, time since charge injection (CI), background level and the parameters attributed to the CCD. This includes trap parameters.

4.1 Charge Distortion Models

To date, three CDMs have been proposed:

1. *CDM01*. A phenomenological, macroscopic model that attempts to describe the distortions observed in the test data at an empirical level. Described in Lindegren (LL-075), CDM01 constrains the illumination history by defining a single state variable called the 'equivalent fill level'. The number of electrons captured and released are calculated in each time step along a CCD column to output a set of distorted counts. CDM01 was intended to commence the development of CDM models.
2. *CDM02*. A fast, analytical model based upon accepted physical theory (Shockley-Read-Hall (SRH) theory) that evaluates the number of electrons captured and released over an entire CCD column in a single calculation (see Short (AS-015)). In its original form CDM02 operated in Time Delay Integration (TDI) mode only on 1D data. It could be made to work with two-dimensional (2D) data (as is the case in this analysis), operate in imaging mode, and handle CTI in the serial register

(SR) but these modifications were left to the user¹.

3. *CDM03*. With *CDM03* these additions have been included automatically. It can now handle both TDI and imaging mode to process transits through the image area (IA) and SR, as well as clock a charge injection (CI) through the IA in imaging mode. It no longer assumes that all traps are filled by the CI.

For further details on the CDMs as well as SRH theory refer to Appendix C.

Before any CDM is implemented in the data processing it must be rigorously tested and validated. One of the first parts of this validation is to compare the predictions of the CDM against the observations seen in the radiation campaigns (RCs). In order to carry out such a comparison a procedure has been developed using the CEMGA platform (Prod'homme (2010)) to simulate the transit of an image over an AF CCD, and implement a CDM to output a distorted image. *CDM02* is the first version of this model to be tested (see Prod'homme et al. (TPH-003)).

4.2 The test data

When work on CDM validation began, *CDM02* could only operate in TDI mode. This meant it was not possible to conduct a First Pixel Response or a charge trail analysis of the simulated CI profile. It was also assumed that a CI filled every trap on the CCD regardless of the level or duration of the CI. So the first two tests in RC2, which investigated these two factors, could not be used. This left the third RC2 test where the CI delay was varied. In this dataset several magnitudes ranging between $G = 13.3$ and 20 as well as several CI delays (30 to 120000 TDI periods) were investigated. Test cases without injection could still rely on self-injection, limiting the time since CI to ~ 30000 TDI periods.

The test data in RC2 were collected for a range of different temperatures. However, the data used here corresponds to the set which were taken at 163K - the Gaia operating temperature.

4.2.1 Input to the simulation

Undamaged images, used as input to these simulations, were taken from a single pinhole located on the mask in RC2 for two reasons: a) the undamaged images in the leading column varied in both height and width due to non-uniform illumination, and b) the CI level varied across-scan as the injection was made from only one end of the CCD (see Figure 2.36). In reality, CI is to be driven from both ends of the CCD, so the image closest to the CI source was chosen. In the context of the RC2 analysis this was channel 22 in column 1.

The undamaged image in RC2 was modelled by a cubic spline function. To use this spline function in the simulation it had to be generated without refitting the undamaged data in RC2. A spline function is parameterised by a set of knots (or nodes) and a set of coefficients that describe the set of polynomials connecting the knots. Loading both of these into an algorithm² that can sample the spline function at any

¹See Radiation Task Force: Meeting 6: Developments on *CDM3* in <http://www.rssd.esa.int/wikiSI/index.php>

²Refer to the Gaia tool box available via DPAC for many of the algorithms used to generate and fit spline functions.

point it was possible to generate an array of pixelated values that corresponded to the undamaged image as observed in RC2.

As only binned data was available in the RCs, the spline function is 1D. However, since Gaia will observe in 2D, the simulation also requires a 2D image as input. To extrapolate the AC profile it was assumed that this profile was the same as that observed in the AL direction:

$$PSF(x, y) = LSF(x) \times LSF(y), \quad (4.1)$$

where LSF (Line Spread Function) is the undamaged image in RC2 and PSF (Point Spread Function) is the input to the simulation. An example of one of these PSFs at $G = 15$ is shown in Figure 4.1. The flux under the PSF has been scaled such that binning in the AC direction recovers the original LSF.

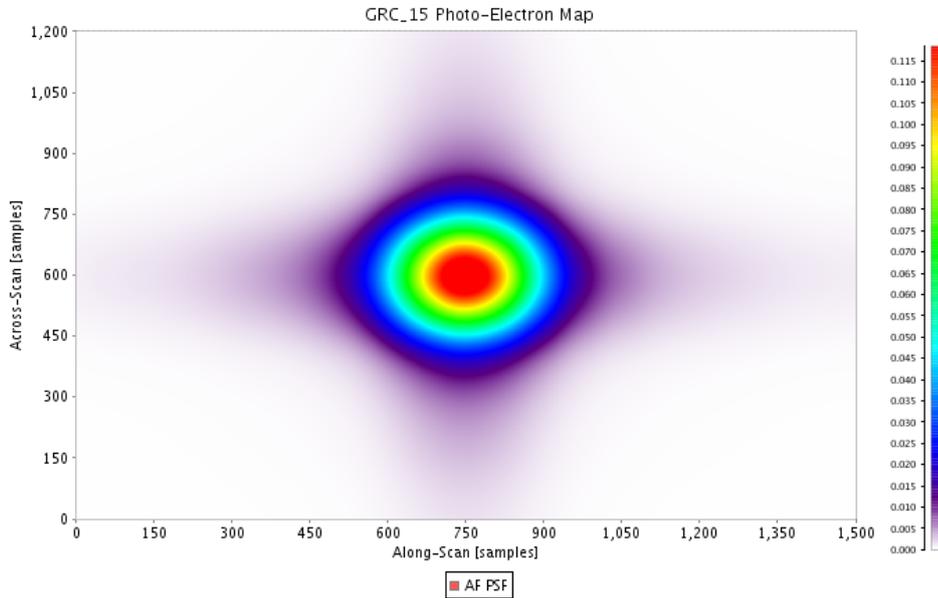


Figure 4.1: An example of a 100x over-sampled normalised PSF (12AC x 15AL pixels) extrapolated from an undamaged image in RC2 at an equivalent magnitude of $G = 15$.

The input must be scaled to the correct magnitude before commencing the simulation. This was achieved by normalising the PSF to the total flux and using eqn. 2.2 to find a close approximation for the actual flux.

4.2.2 Comparative output for the simulation

Rather than compare cubic spline functions to the output of CDM02, the accumulated damaged data itself has been compared. This avoids the ambiguity that may arise from an ill-defined fit to some of the RC2 data, caused by a lack of scans and/or samples clustering around particular sub-pixel positions. These effects were described in Chapter 2.

The accumulated data that have been extracted over a 15-pixel AL window are provided in the form as

shown in Figure 2.81. Except here the datapoints are scaled back to their absolute values for comparative purposes. In the campaign analyses the window was defined over ± 7.5 pixels. Here, the window is redefined to lie between 0 and 15 pixels, for both the input and output. Applying the CDM requires no repositioning so the input and output can be directly compared, i.e. the bias is implicitly accounted for in the CDM02 prediction of the image profile. The only repositioning that is required is when the profile is being sampled at the sub-pixel level, and this process is described in the next section.

4.3 The simulation procedure

Regardless of the different test conditions (e.g. variable magnitude), this is the procedure that is followed in every simulation in order to test the performance of CDM02:

1. To begin with, a 2D undamaged image (modelled as a spline function) is generated from the data archive. The corresponding damaged image profile from RC2 is also loaded. Typically, the damaged image profile will be over-sampled, say n times, so the output of the simulation will also have to be over-sampled n times to make the same comparison.
2. The simulation requires a discrete (i.e. pixelated) input, so the 2D spline function is sampled once per pixel. The positions of where the spline function is sampled depends on the sub-pixel positions found in the over-sampled damaged image profile. For example, if in one scan the RC2 damaged profile was found to be sampled at 0.3, 1.3, 2.3 pixels etc, the spline function would also be sampled at these positions. Then the 2D pixelated image would be fed through the simulation (here each datapoint in the image would have to be positioned at 0, 1, 2 pixels, etc) and the pixel positions of the output would be shifted back to the original sampling positions (i.e. 0.3, 1.3, 2.3 pixels, etc).
3. As the damaged RC2 image profile has been over-sampled n times, the simulation is repeated n times in order to generate n CDM02 damaged profiles. On readout these profiles are binned to form 1D CDM02 damaged profiles and then combined to form a single 1D CDM02 damaged profile over-sampled n times. This is the CDM02 prediction that will be compared to the observed RC2 damaged image profile.
4. Note that at this stage each simulation (out of a total of n) is performed with a single predefined set of model parameters.
5. The over-sampled CDM02 damaged image profile and RC2 damaged image profile are compared using a customised goodness-of-fit (GOF) parameter, shortly to be explained.
6. The GOF is minimised using a Downhill-Simplex (DS) method, also shortly to be explained.
7. When an abort condition is met the minimisation procedure terminates.

Quantifying the quality of the fit between the two damaged image shapes is carried out using a GOF parameter g , based on the standard χ^2 statistic:

$$g = \frac{1}{S} \times \frac{1}{F} \times \sum_{i=0}^{S-1} \frac{(\lambda(x_i) - N(x_i))^2}{\sigma^2}, \quad (4.2)$$

where λ is the CDM02 damaged profile, N the RC2 damaged profile, x_i ($i = 0, \dots, S - 1$) the sub-pixel position of the i^{th} datapoint, S the total number of datapoints, σ the noise, and F the total integrated flux. The noise, σ , is considered to be the quadratic sum of the photon noise and the readout noise r . The photon noise is assumed to follow a Poisson distribution, which for large counts is limited by a Gaussian distribution with $\sigma_g = \sqrt{N}$:

$$\sigma^2 = N(x_i) + r^2. \quad (4.3)$$

r was fixed to the mean value, $4.8 \text{ e}^-/\text{sample}$, found in the RC2 analysis. The scaling factor S is required when different RC2 damaged profiles with different numbers of datapoints per pixel are fitted simultaneously. In the same way, F is needed when including RC2 damaged profiles of different magnitudes. If these scaling factors are not present, the fit can be biased towards the brightest magnitudes and/or the most over-sampled profiles, which in each case return the highest values of g .

4.3.1 The Downhill-Simplex method

A Downhill-Simplex (DS) algorithm was used (cf. Nelder & Mead (1965)) to minimise the parameter given by eqn. 4.2 with respect to the free parameters of the model. If CDM02 uses N free parameters to fit the test data (e.g. trap densities, capture cross-sections, and release time-constants for a number of different trap species), the DS method initially creates $N + 1$ different sets of these parameters. The $N + 1$ sets of parameters represent points in the N -dimensional parameter space that can be illustrated as the nodes of a simplex. One set contains the (estimated) initial values, the other N sets contain variations on each parameter in turn, as follows:

$$\begin{aligned} (i = 0 : & (\rho_1, \quad \sigma_1, \quad \tau_1), (\rho_2, \quad \sigma_2, \quad \tau_2), \dots, (\rho_{N_t}, \quad \sigma_{N_t}, \quad \tau_{N_t}), \dots) \\ (i = 1 : & (\rho_1 \times A, \quad \sigma_1, \quad \tau_1), (\rho_2, \quad \sigma_2, \quad \tau_2), \dots, (\rho_{N_t}, \quad \sigma_{N_t}, \quad \tau_{N_t}), \dots) \\ (i = 2 : & (\rho_1, \quad \sigma_1 \times A, \quad \tau_1), (\rho_2, \quad \sigma_2, \quad \tau_2), \dots, (\rho_{N_t}, \quad \sigma_{N_t}, \quad \tau_{N_t}), \dots) \\ & \vdots \\ (i = N : & (\rho_1, \quad \sigma_1, \quad \tau_1), (\rho_2, \quad \sigma_2, \quad \tau_2), \dots, (\rho_{N_t}, \quad \sigma_{N_t}, \quad \tau_{N_t} \times A), \dots) \end{aligned} \quad (4.4)$$

Each line represents a set of parameters assigned to each node in the initial simplex, where only one parameter has been changed. i represents the i^{th} set of parameters and P_j is the fitting parameter, e.g. the trap density ρ , the cross-section σ or the release time-constant τ , where j represents the j^{th} trap species ($j = 1, \dots, N_t$ where N_t is the total number of trap species). A is the factor by which the parameter has been varied. For these simulations A was set to 1.1 to keep the initial size of the simplex quite small. Note that the β parameter and the diffuse optical background (DOB) level are also fitted in the DS method.

The DS method operates as follows:

1. First, $N + 1$ sets of parameters are set up by the DS method to create the initial simplex.
2. The GOF for each of the nodes is determined.
3. Then, the worst, second worst, and the best node are selected (i.e the node with the worst, etc GOF parameters).
4. Mirror the worst node on the "face" of the simplex.
5. If the mirror node is now the best node, try to go further in this direction by expanding the simplex in that direction by a factor (2 was used).
6. If the mirror node is better than the original worst point, but worse than the second-worst point, look for a point between the original and the mirror point; make a 1D contraction of the simplex by a given factor (0.5 was used).
7. If a 1D contraction still doesn't help to find a real improvement, do a multidimensional contraction towards the best point you have so far (a factor of 0.5 was used for this).

As the simulations continue the simplex changes its form to move "downhill" and contract towards the minimum node, i.e. the node that describes a set of parameters that minimise g . If the relative difference of the GOF parameters of the remaining nodes of the simplex is less than a given limit, the algorithm terminates. In this case, a value of 10^{-8} was used.

There is a risk that the algorithm will get stuck in a local minimum, rather than converging on the global minimum. To reduce this risk several simulations were run with different initial sets of parameters and it was checked whether the algorithm in all cases converged on the same solution.

Furthermore, the option was included to constrain all free parameters to remain within certain specified intervals. If the algorithm suggested a parameter should leave its allowed interval, the parameter was set to the corresponding limit of that interval. There are two reasons why this is useful. Firstly, it excludes the possibility that a parameter may adopt an unphysical value (e.g. a negative trap density). Secondly, it stops certain parameters (e.g. release time-constants) for different trap species from overlapping.

All of the parameters, both fixed and fitted, are shown in Table 4.1. As can be seen, all of the trap parameters were fitted as well as the β parameter. The level of the DOB was also left unconstrained for fitting. In the first instance this was to see whether the parameters that were found for the DOB were in any way physical, and also to see if there were any correlations between parameters. In retrospect, this approach, whereby many unknowns are being fitted, requires as many parameters as possible to be fixed. Since the DOB was measured during the analysis of the second radiation campaign, this parameter could have been constrained, and in future simulation runs this could be done. The time since charge injection t_{CI} was fixed according to the RC2 test case that was being chosen for comparison.

Parameter	Description	Default value	Allowed interval
CCD parameters			
T	temperature	163K	-
V	signal confinement volume	$30\mu\text{m} \times 5\mu\text{m} \times 1\mu\text{m}$	-
FWC	CCD full well capacity	$190000e^-$	-
P	TDI period	0.9828ms	-
β	electron cloud growth parameter	-	0 - 1
Trap parameters			
J	number of trap species	4	-
Nt_j	density of trap species j	-	$10^{-10} - 10^{30}\text{m}^{-3}$
σ_j	capture cross-section	-	$10^{-35} - 10^{01}\text{m}^2$
τ_j	release time-constant	-	$10^{-12} - 10^6\text{s}$
Other parameters			
t_{CI}	time since last charge injection	0 - 10^6s depending on test case	-
DOB	diffuse optical background	-	$0 - 10^6 e^- \cdot \text{pixel}^{-1} \cdot \text{s}^{-1}$

Table 4.1: Simulation parameters

4.3.2 How many trap species are required?

One of the trap parameters that is fixed beforehand in the simulations is the number of trap species J . As shown in Table 4.1, the default value is 4. Why has 4 been chosen and no other value, at least for this initial round of tests? The choice extends from previous work on the subject, primarily performed on the analysis of charge injection trails. See, for example, Hambly & Cross (NCH-003) and Astrium's own analyses. In EADS Astrium (GAIA.ASF.TCN.PLM.00071) the statistical arguments for understanding trap behaviour as a function of time are laid down. The average number of electrons released by a family of traps into the trail of a charge injection are modelled as an exponential decay process or the sum of decaying exponentials, each term representative of a different trap species. A non-linear fitting procedure is performed to estimate the trap amplitudes (i.e. the number of traps) and release time-constants. To evaluate the performance of the fitting procedure Astrium adopted a χ^2 statistic, weighted such that $\langle \chi^2 \rangle = 1$. By successively increasing the number of trap species from 1, Astrium found that at least 5 trap species were required to return a χ^2 value sufficiently close to unity. Later, following an analysis of the radiation campaign 2 data, in EADS Astrium (GAIA.ASF.TCN.PLM.00197), it was stated that only 4 trap species were required to reach a sufficiently good fit. This is why in the first instance 4 trap species were chosen to perform the fitting procedure.

4.4 Initial parameters

As mentioned above, to initiate the simulation a set of parameters had to be chosen that would later be fitted by the DS method. To avoid the risk of a particular set of parameters settling on a local minimum rather than the global minimum, the initial parameters were varied. Five different sets of initial parameters were chosen and these are shown in Table 4.2 labelled A to E.

Parameter	A	B	C	D	E
β	10^{-4}	10^{-3}	10^{-2}	10^{-4}	5×10^{-1}
dob [$e^-/s/pixel$]	10^{-3}	10^{-2}	10^{-1}	10^{-3}	10^1
ρ_1 [per CCD line]	10^4	10^5	10^6	10^3	10^3
σ_1 [m^2]	10^{-22}	10^{-21}	10^{-20}	10^{-19}	10^{-19}
τ_1 [s]	10^2	10^3	10^4	10^{-1}	10^{-1}
ρ_2 [per CCD line]	10^4	10^5	10^6	10^4	10^4
σ_2 [m^2]	10^{-22}	10^{-21}	10^{-20}	10^{-20}	10^{-20}
τ_2 [s]	10^0	10^1	10^2	10^{-1}	10^{-1}
ρ_3 [per CCD line]	10^4	10^5	10^6	10^5	10^5
σ_3 [m^2]	10^{-22}	10^{-21}	10^{-20}	10^{-21}	10^{-21}
τ_3 [s]	10^{-2}	10^{-1}	10^0	10^{-1}	10^{-1}
ρ_4 [per CCD line]	10^4	10^{-5}	10^6	10^6	10^6
σ_4 [m^2]	10^{-22}	10^{-21}	10^{-20}	10^{-22}	10^{-22}
τ_4 [s]	10^{-3}	10^{-2}	10^{-1}	10^{-1}	10^{-1}

Table 4.2: Sets of free parameters defined to initiate the simulation.

4.5 Results

4.5.1 Test 1: fitting to a single damaged image profile

Taking a set of damaged images from RC2 at different magnitudes and at a specific time since CI, and fitting them individually, is the first and simplest test that was performed. Five magnitudes were available, nominally $G = 13.3, 15, 16.75, 18.5$ and 20 , and the CI delay was set to the average time between scans in RC2 ($\sim 27.3s$ or 27778 TDI periods). Tests without CI in RC2 tended not to have the problem of samples clustering on the same sub-pixel positions. Therefore, the datapoints were more evenly spread over the damaged image profile and could be used to better constrain the fitting procedure in this analysis.

Magnitudes	15.29	15.29	15.29	15.29	15.29
starting configuration	A	B	C	D	E
goodness-of-fit					
g	1.76×10^{-4}	1.96×10^{-4}	1.42×10^{-4}	2.51×10^{-4}	8.41×10^{-5}
final parameters					
β	5.26×10^{-9}	2.62×10^{-3}	0.0182	1.27×10^{-3}	0.904
dob [$e^-/s/pixel$]	8.47×10^{-4}	4.75×10^{-3}	7.575	2.96×10^{-3}	660.0
ρ_1 [per pixel]	3.91×10^{-6}	2.23×10^{-4}	510.6	2.70×10^{-6}	5.0×10^{-27}
σ_1 [m^2]	1.0×10^{-35}	3.28×10^{-31}	2.83×10^{-20}	2.56×10^{-19}	1.87×10^{-19}
τ_1 [s]	103.0	2.608	509.4	0.123	3.42×10^{-3}
ρ_2	1.61	100.6	1.97	5.00×10^{-27}	58.44
σ_2	4.13×10^{-22}	1.57×10^{-21}	5.32×10^{-22}	9.45×10^{-20}	3.68×10^{-24}
τ_2	0.328	74.26	13.63	2.632	3.054
ρ_3	1.56×10^{-7}	0.942	5.0×10^{-27}	5.0×10^{-27}	12.03
σ_3	6.37×10^{-23}	1.91×10^{-22}	1.55×10^{-23}	1.0×10^{-35}	3.21×10^{-20}
τ_3	3.42×10^{-4}	1.0×10^{-12}	4.435	0.491	1.25
ρ_4	0.60	5.0×10^{-27}	1.74×10^5	1.582	3.05×10^3
σ_4	7.62×10^{-23}	5.31×10^{-24}	1.86×10^{-23}	4.73×10^{-22}	1.0×10^{-35}
τ_4	4.90×10^{-8}	0.122	4.27×10^{-4}	0.249	3.663

Table 4.3: Best fit parameters obtained with different starting configurations at $G = 15.29$.

Magnitudes	13.63	15.29	16.96	18.65	20.25
starting configuration	B	E	E	E	E
goodness-of-fit					
g	2.98×10^{-5}	8.41×10^{-5}	9.21×10^{-5}	4.72×10^{-4}	2.00×10^{-3}
final parameters					
β	6.14×10^{-5}	0.904	0.0332	0.368	0.253
dob [e^- /s/pixel]	0.0	660.0	0.308	2.75×10^{-6}	1.85×10^{-3}
ρ_1 [per pixel]	93.50	5.0×10^{-27}	2.06×10^{-3}	0.0212	4.13×10^{-4}
σ_1 [m^2]	2.11×10^{-30}	1.87×10^{-19}	5.51×10^{-19}	1.72×10^{-20}	8.40×10^{-19}
τ_1 [s]	27.40	3.42×10^{-3}	0.0134	0.677	2.44×10^{-6}
ρ_2	5.0×10^{-27}	58.44	0.967	1.06×10^{-8}	0.0222
σ_2	1.0×10^{-35}	3.68×10^{-24}	7.81×10^{-22}	5.06×10^{-20}	1.17×10^{-19}
τ_2	483.5	3.054	2.569	0.546	3.614
ρ_3	1.733	12.03	5.0×10^{-27}	120.4	15.13
σ_3	1.44×10^{-22}	3.22×10^{-20}	7.13×10^{-20}	1.36×10^{-22}	7.41×10^{-22}
τ_3	0.526	1.249	0.256	0.223	1.362
ρ_4	32.07	3.05×10^3	0.420	315.1	709.7
σ_4	6.49×10^{-25}	1.0×10^{-35}	4.79×10^{-22}	1.84×10^{-22}	1.40×10^{-21}
τ_4	1.0×10^{-12}	3.663	7.364	1.718	2.093

Table 4.4: Best fit parameters for individual damaged profiles of different magnitudes.

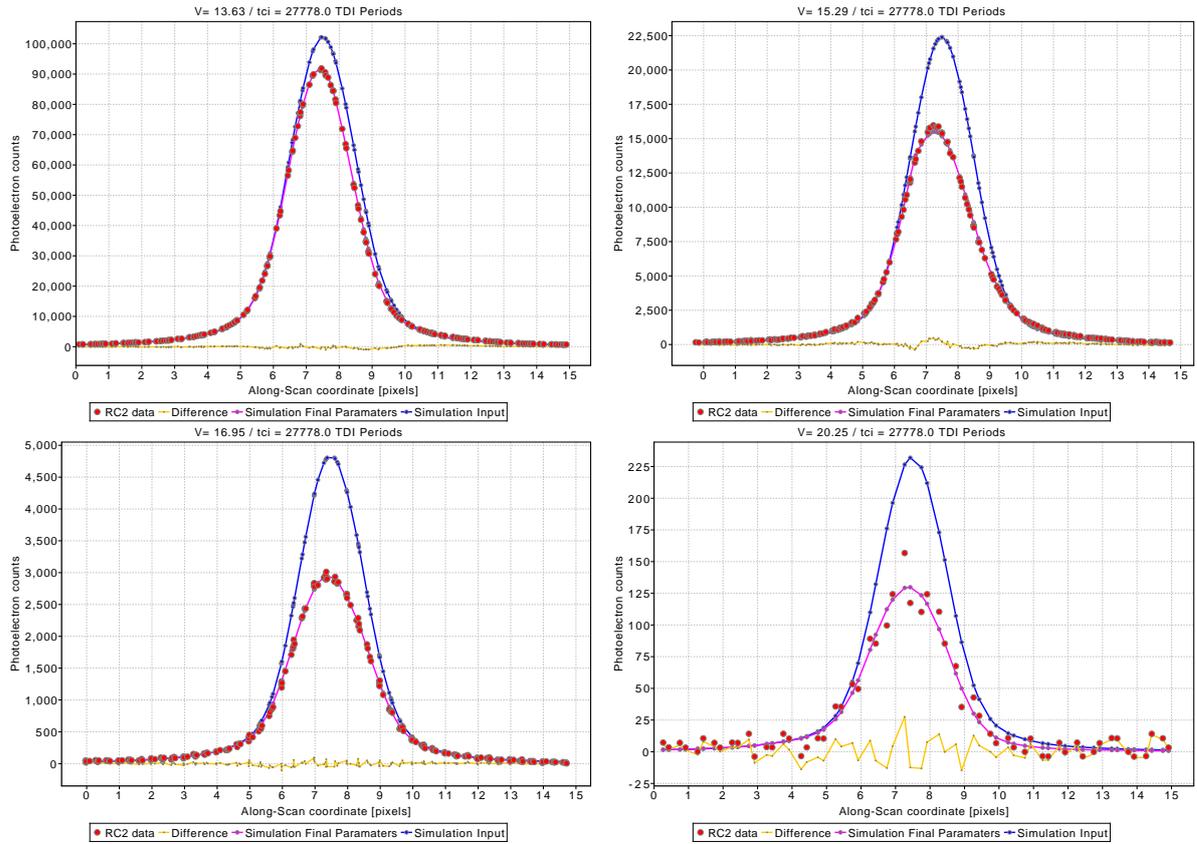


Figure 4.2: Best fits to individual damaged profiles at $G = 13$ to 20 without CI. Red dots show the RC2 damaged profile, blue dots the undamaged simulation input, pink dots the best CDM02 fit, and yellow dots the difference between the CDM02 predicted image shape and the RC2 damaged profile.

For a single image shape spread over a 15-pixel window, one can expect radiation damage to arise from a host of trap species. Some may take charge away from the image and release it again after the image has passed. Others may take charge away from the image and release it back into the same location or further down the image. Which trap species are dominant under different levels of DOB will be investigated here, though there is a certain degree of degeneracy in the fitting parameters. For example, a trap species with a large density can also have a small capture cross-section, which makes it less likely to interact with an electron and distort the image shape. While a trap species with a small density and a large capture cross-section may be filled by a DOB at all times, thereby reducing the effect of this trap species on the image.

Table 4.3 shows the results for a $G = 15$ star for each of the different sets of initial parameters shown in Table 4.2. Notably, a different set of final parameters are found for each different set of initial parameters, though the GOF is roughly the same in each case. It can be observed that a low DOB level with small trap densities (A, B and D) gives good fits as well as a high DOB level with high trap densities (C and E). Furthermore, trap species with a high density (e.g. 510.6 or 58.44) and a relatively moderate capture cross-section (e.g. 2.8×10^{-20} or 3.68×10^{-24}) have release time-constants that are large (e.g. 509.4s or 3.054s) compared to the length of the telemetry window (15 TDI periods (~ 0.015 s)). This suggests that the damage observed on individual image shapes could be modelled by assuming only the traps with release time-constants longer than the length of the image are present³. In other words, as the charge is lost completely from the image, a single dominant trap species could be used to model the distortion of a single image that has a slow trap release time-constant.

Table 4.4 summarises the best results at each magnitude. Note that at $G = 15$ the high trap density compensates for the unrealistically high DOB level. Note also that the β parameter varies in each case. Figure 4.2 compares the final predicted image shapes against the actual damaged image shapes from RC2. The binned input is also shown. At face value CDM02 appears to be performing well, though the final set of parameters in each case are different.

4.5.2 Test 2: fitting to a set of damaged image profiles with different CI delays

How well CDM02 can fit a large set of damaged images simultaneously is addressed in this second test. Here, for each available magnitude and every CI delay, all of the damaged images were fitted simultaneously using a single set of model parameters. The delay could vary between 30 and 120000 TDI periods for each magnitude, depending on the number of test cases performed in RC2.

As it is assumed that the CI fills all traps in the CCD, only those traps with release time-constants less than or close to the time since CI will be empty as a star transits the CCD. As the CI delay increases, so too will the number of vacant traps that can extract charge from the image. It is expected that certain trap species will become more dominant as the CI delay increases.

Table 4.5 shows the results for the best fits at each magnitude as well as the initial set of parameters that were required to achieve these fits. Some observations:

³Follow-up tests with 1, 2 and 3 trap species have shown that the results with 1 trap species are comparable to those found with 4, with 2 trap species the results were better, and with 3 no noticeable improvement was seen.

Magnitude	13	15	17	18	20
CI delays	30 60 100 150 300 1000 5000 30000	30 60 100 150 300 1000 5000 30000 30000 60000 120000	30 60 100 1000 30000	100 150 300	30 60 100 150 300
goodness-of-fit					
$\langle g \rangle$	1.52×10^{-4}	2.23×10^{-4}	2.86×10^{-4}	6.25×10^{-4}	0.0041
starting conf.	C	A	A	B	D
final parameters					
β	0.036	7.47×10^{-7}	7.39×10^{-5}	1.68×10^{-5}	1.77×10^{-5}
dob [$e^-/s/pixel$]	0.0	0.0012	0.25	0.082	0.0026
ρ [per pixel]	4.38	3.31	0.11	5.86	0.0029
σ [m^2]	7.42×10^{-22}	8.30×10^{-22}	2.02×10^{-20}	6.74×10^{-19}	5.07×10^{-20}
τ [s]	35.74	43.35	1.48×10^4	6.67	8.04×10^{-4}
ρ	212.32	504.16	1.11	4.78	0.79
σ	1.0×10^{-35}	1.31×10^{-22}	7.85×10^{-22}	2.58×10^{-18}	1.37×10^{-20}
τ	32.06	797.10	3.23	78.89	0.38
ρ	0.086	0.99	0.41	5.0×10^{-27}	5.0×10^{-27}
σ	5.89×10^{-30}	1.37×10^{-22}	4.91×10^{-22}	1.11×10^{-21}	4.28×10^{-21}
τ	0.024	0.30	0.069	43.64	0.42
ρ	5.29×10^{-5}	4.43	5.0×10^{-27}	0.24	0.22
σ	1.15×10^{-23}	1.95×10^{-27}	1.24×10^{-20}	1.28×10^{-21}	3.54×10^{-23}
τ	0.099	0.0052	0.54	0.076	0.027

Table 4.5: Best fit parameters for a specific magnitude with different CI delays

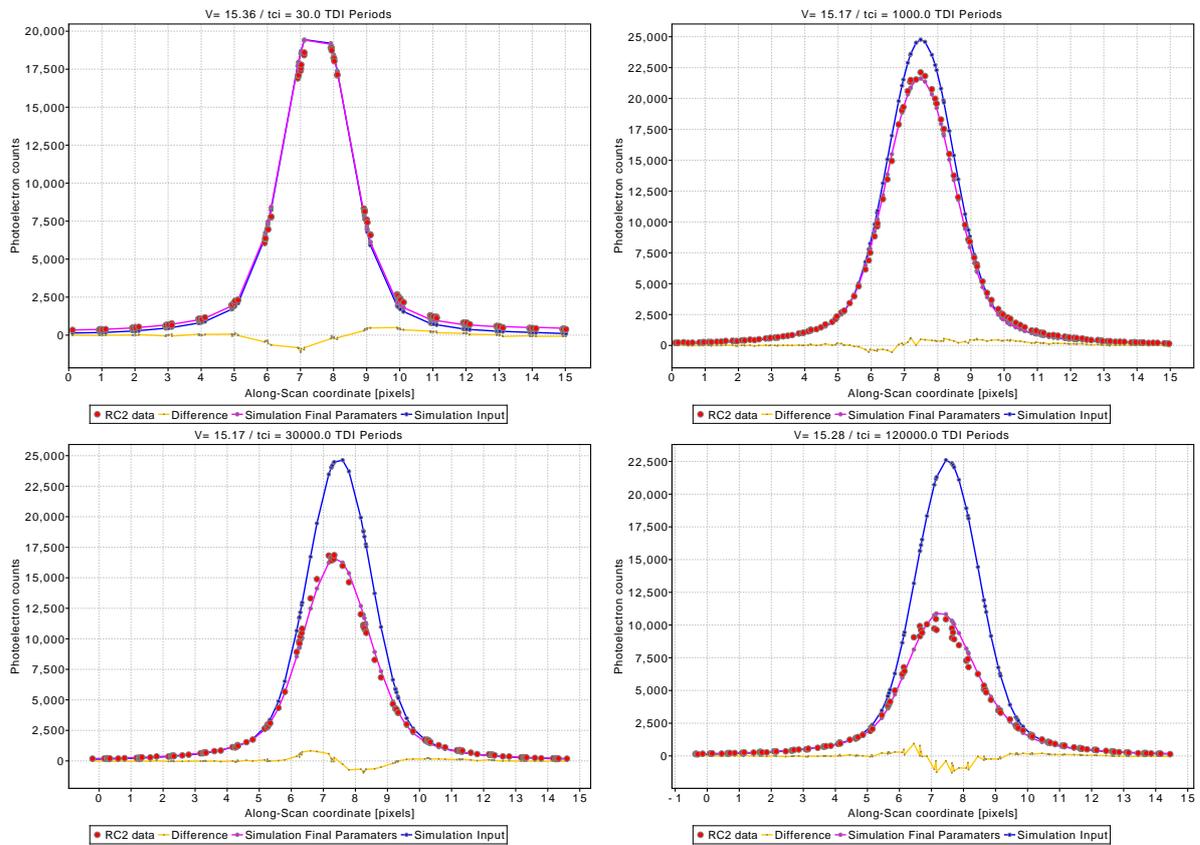


Figure 4.3: Best simultaneous fits to damaged image shapes at $G = 15$ with different CI delays.

- The quality of the final best fit, when fitting to a larger set of damaged images simultaneously, is not in any noticeable way worse compared to when fitting single images. This is surprising given the number of profiles that must be fitted, the problem with samples clustering, and the variable number of scans performed at each magnitude. Yet it may confirm the limitations inherent in the test data itself. These issues just mentioned may limit the accuracy of the fit, which sits above the level required to probe the accuracy of the CDM itself.
- When fitting simultaneously a large number of damaged image profiles over a variable time since CI, CDM02 requires 2 to 4 dominant trap species, whereas as was shown previously (for high density traps with moderate cross-sections) only one dominant trap species (with a long release time-constant compared to observational timescales (>15 TDI periods)) could, in principle, be required to fit a single damaged image profile. If the trap species with extremely low densities or capture cross-sections are discarded, at $G = 13$ only two trap species are necessary to obtain a good fit to the RC2 test data, at $G = 15$ four are required (note that this is the largest set of data spanning the largest range of CI delays), and at the fainter end of the magnitude scale, 3 trap species are required.
- When the CI delays are the same at different magnitudes one would expect the trap parameters to be the same, since the same release time-constants are being probed. However, the results show that different trap parameters are found in each case. This may be because both the DOB level and the β parameter were also fitted as part of the algorithm and these values change depending on the magnitude. That said, in all cases in Table 4.5, both the DOB level and the β parameter remain low. Clearly, some ambiguity remains in the test.
- If the trap densities for the dominant trap species (i.e. excluding those with very low capture cross-section or $\tau \ll 1$ TDI period) are summed, the total trap density at each magnitude (13: 4.38, 17: 1.63, 18: 10.88, 20: 1.01), excluding $G = 15$, are in agreement with the trap density intervals derived using an FPR analysis of the CI profiles (see Chapter 3 for the results at $G = 13.3$). The high trap density at $G = 15$ may be explained by the lack of constraint on trap species with long release time-constants (see section 4.5.4).

Figure 4.3 shows some of the final fits at various CI delays between the CDM02 damaged image profiles and the RC2 damaged image profiles at $G = 15$.

The points made above are general remarks to try and interpret the results. However, it is clear that there are some difficulties in this area, especially when posing specific questions. For example, how are the traps that are not probed by these experiments accounted for in the interpretation of the analysis? To answer this question the actual number of traps per pixel per trap species would have to be known. As these values are not known it is difficult to know how close these results are to the actual conditions on the CCD. In which case further constraints are required. Also, it should not be assumed that all the traps are filled by the charge injection.

4.5.3 Test 3: fitting to a set of damaged image profiles with different magnitudes

How well CDM02 can model a large set of damaged image shapes at different magnitudes was addressed in this test. The same time since CI was used for each magnitude and all damaged image shapes were fitted simultaneously. Figure 4.4 shows some of the best fitted images at different magnitudes when no injection was made (starting configuration D).

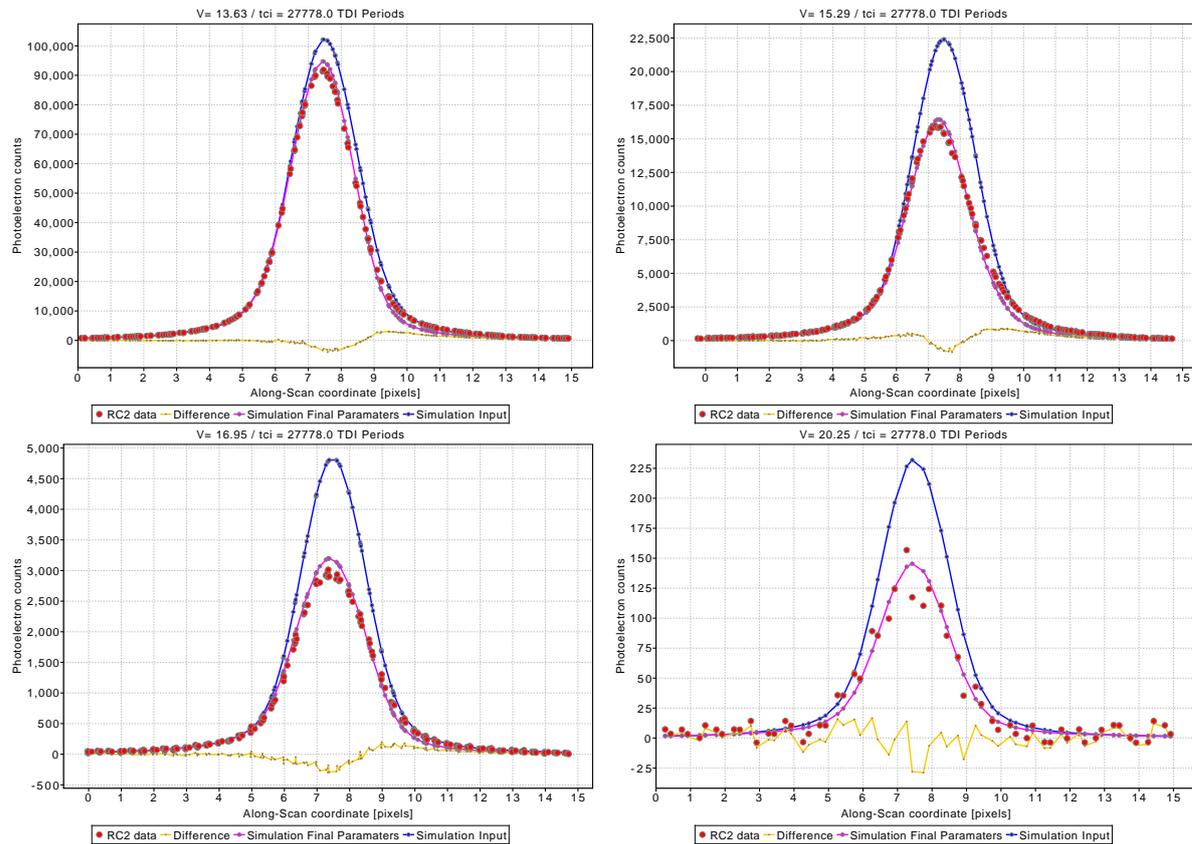


Figure 4.4: Best simultaneous fits to damaged image shapes without CI at different magnitudes.

Tables 4.6 and 4.7 summarise the best results for each time since CI. Some observations:

- For CI delays of 60, 100, 1000, 5000, and 30000 TDI periods, the GOF parameter is of the same order as that found in Test 1 when individual damaged images were fitted. However, for the other delays (30, 150, 300 and No Injection) the quality of the final fit is relatively worse (> 0.001). These differences may be due to temperature variations between tests in RC2, rather than the performance of CDM02 itself. For example, when the magnitude was changed, the CCD temperature would in some cases vary by more than 1K. However, when the CI delay was changed the CCD temperature was far more stable (see Table 2.10 in Chapter 2). Since these simulations assume the CCD temperature is kept fixed at 163K, the impact of any temperature fluctuations in the data has not been taken into account. The release time-constants are temperature dependent. Therefore, a single trap species may be misinterpreted as two if the temperature is changed. This may also explain why local minima were found to be common and different sets of initial parameters led to

CI delay	30	60	100	150
Magnitudes	13 15 17 20	13 15 17 20	13 15 17 18	13 15 18 20
goodness-of-fit				
$\langle g \rangle$	0.0028	8.56×10^{-4}	4.58×10^{-4}	0.0011
starting configuration	B	A	B	E
final parameters				
β	0.0	0.0018	0.0	0.023
dob [e^- /s/pixel]	0.042	1.26×10^{-4}	0.22	46.50
ρ [per pixel]	0.0099	29.74	21.14	0.23
σ [m^2]	1.67×10^{-21}	5.21×10^{-22}	1.55×10^{-20}	4.09×10^{-20}
τ [s]	1.0×10^{-12}	931.48	4637.78	0.55
ρ	2.00	0.20	1.45	0.20
σ	3.44×10^{-27}	8.43×10^{-23}	2.81×10^{-20}	4.45×10^{-20}
τ	11.24	0.075	0.81	0.60
ρ	0.40	0.23	5.0×10^{-27}	0.016
σ	5.54×10^{-22}	7.87×10^{-23}	6.31×10^{-21}	7.49×10^{-22}
τ	0.054	0.087	1.31	0.60
ρ	1.83	0.0057	130.61	3.09×10^{-16}
σ	4.82×10^{-29}	2.23×10^{-24}	1.0×10^{-35}	4.69×10^{-22}
τ	2.75×10^{-4}	3.32×10^{-8}	0.0092	0.0055

Table 4.6: CDM02 best fit parameters for a CI delay ≤ 150 and different sets of magnitudes.

CI delay	300	1000	5000	30000	No CI
Magnitudes	13 15 18 20	13 15 17	13 15	13 15 17	13 15 17 18 20
goodness-of-fit					
$\langle g \rangle$	0.0018	4.69×10^{-4}	5.35×10^{-4}	1.95×10^{-4}	0.0011
starting configuration	D	E	C	E	D
final parameters					
β	1.23×10^{-6}	0.0	2.34×10^{-6}	0.0090	0.010
dob [e^- /s/pixel]	0.0015	17.70	0.0040	19.50	0.24
ρ [per pixel]	8.75	8.35	1.54	1.41	0.20
σ [m^2]	2.06×10^{-18}	3.47×10^{-21}	5.53×10^{-22}	2.016×10^{-22}	3.02×10^{-18}
τ [s]	3.07	1.38	8.59	0.36	11.06
ρ	0.021	1.70	191.16	6.19×10^{-6}	1.14
σ	5.24×10^{-21}	3.84×10^{-21}	8.62×10^{-33}	2.41×10^{-17}	6.17×10^{-21}
τ	0.0081	3.07	188.65	7.88	0.37
ρ	0.017	0.15	3.07	1.09	0.033
σ	2.20×10^{-20}	1.59×10^{-23}	1.60×10^{-33}	2.01×10^{-21}	6.29×10^{-22}
τ	2.61	0.0058	5.81	7.56	1.69×10^{-11}
ρ	1.92	9.64	99.09	81.50	0.049
σ	7.94×10^{-20}	3.57×10^{-28}	9.14×10^{-33}	2.26×10^{-25}	8.91×10^{-23}
τ	0.10	0.11	0.31	1.0×10^{-12}	0.0012

Table 4.7: CDM02 best fit parameters for a CI delay > 150 and different sets of magnitudes.

different sets of fitted parameters, even though the quality of the fits were the same in most cases. To avoid these degenerate problems, it may be better to fix the temperature in each simulation, according to the prescribed temperature given for each individual test case.

- On average, 3 trap species were required in Test 2 to explain the damage observed at different CI delays. In Test 3, this number rises to 4 even when the number of damaged profiles available for a specific CI delay are few. This may also be due to the degeneracy in the data discussed above. So far, no more than 4 species have been used to fit the data.
- The β parameter remains consistently very low for every CI delay. However, the DOB level can be seen to vary significantly. For CI delays of 150, 1000 and 30000, the DOB level is $> 17 e^-/\text{pixel}/\text{s}$.
- The sum of the trap densities for the dominant trap species at each CI delay are in half of the cases generally higher than those found in Test 2. Furthermore, as the CI delay increases, the numbers are expected to increase. However, this is not seen. There is a considerable degree of variation as the CI delay increases.

In summary, it is difficult to identify a single set of trap species that can explain the results for all magnitudes and for all CI delays. Even if it appears at first that a consistent set of trap species with release time-constants of the order of a second, tens of seconds, ten times the TDI period have been found, their capture cross-sections and densities in each case are different.

4.5.4 Test 4: fitting to every available damaged image profile

In this final test every damaged image profile for every magnitude and at every CI delay is fitted simultaneously. In effect, this involves fitting 35 over-sampled damaged image profiles at the same time using a single set of fitted parameters. Table 4.8 shows the results that gave the best fits. The starting configuration is C.

β	DOB	Trap species	ρ	σ	τ
0.090	0.42	1	1381.99	6.86×10^{-20}	17558.31
		2	2.37	3.93×10^{-22}	2.34
		3	0.036	9.42×10^{-20}	0.27
		4	0.27	1.36×10^{-22}	0.042

Table 4.8: The set of parameters corresponding to the best fit to all of the data. Units: dob [$e^-/\text{pixel}/\text{s}$], ρ [per pixel], σ [m^2], τ [s].

By fitting to all of the data:

- $\langle g \rangle \approx 0.0014$ is higher than the corresponding value found in Test 2, but comparable to some values obtained in Test 3.
- In general, the fainter the image the poorer the fit.
- In general, four trap species are required.

- The DOB level and β parameter are both low.
- The total density of traps is very high due to the presence of one trap species that has a density of 1382 traps per pixel and an extremely long release time-constant. This may be linked to the assumption that the charge injection fills all of the traps, which has been removed in CDM03.

4.6 Conclusions

To begin the process of rigorously testing the performance of each CDM, the passage of single stellar image profiles through a damaged CCD have been modelled both individually and simultaneously over a range of magnitudes and times since charge injection in order to compare the output to a set of real damaged image profiles. To date, only RC2 has been fully reduced. Summarising:

1. The GOF parameters found throughout the simulations suggest that CDM02 can fit the RC2 data well. Whether this performance is sufficient must be assessed in terms of the ability of CDM02 to extract accurate image parameters. This will be evaluated in the next chapter.
2. In the majority of cases, the quality of the final fit found when fitting individual damaged image profiles to the RC2 data was comparable to that found when fitting many damaged image profiles simultaneously. Since there are no significant differences between the two, it may be that the GOF is limited by the quality of the RC2 data itself and not the performance of CDM02.
3. The exception is when many damaged image profiles are fitted simultaneously at different magnitudes at the same CI delay. Here, the quality of the final fit was found to be relatively worse compared to that found when fitting individual damaged image profiles. A varying CCD temperature may cause this, since a traps release time-constant is temperature dependent. To avoid such inconsistencies in RC2 the CCD temperature would have to be fixed for each individual test simulation.
4. No particular set of trap species has been identified that can describe consistently the results found in each test performed. Changing the initial set of parameters resulted in a different set of final fitted parameters.

In effect, the current parameter space is too large, which cannot be constrained enough to return a unique set of trap species. If this exercise is to be continued some means must be found to further constrain the fitting parameters. To begin with the value for the diffuse optical background as well as the temperature of the CCD in each individual test case can be provided. In RC2 this information is available, albeit in the form of an estimate. If this looks promising, other minimisation routines other than the Downhill-Simplex method could be implemented that may be more robust at this type of problem (cf. genetic algorithms).

There is reason to believe that the quality of the RC2 data is an issue in this simulation. Sub-pixel sample clustering and temperature fluctuations between tests have been avoided in RC4, which should now supersede this dataset.

At present these simulations are not giving a definitive picture of the number of trap species and their associated properties. However, there is room for improvement. Given the importance being placed on the performance of the CDM, it would be useful to implement these changes and see how CDM03 performs.

5

Image Parameter Extraction

At the time of writing the current approach being taken by DPAC to remove radiation damage from Gaia data is to instigate a Charge Distortion Model (CDM) in the pre-processing part of the data pipeline. This is where image parameters (centroid and flux) are extracted and corrected for CTI effects before the astrometric and photometric processing stages take over. The original counts recorded by Gaia are never treated directly, i.e. they are not corrected by some inversion process, which may complicate calibrating both the signal and noise. Instead, the observed signal is modelled by a CDM, which may recover the 'true' centroid and flux by comparing sampled counts to the actual observed counts and iterating until the two converge. Once an accurate fit has been found both the state of the CCD at the time of observation and the (undamaged) LSF/PSF image parameters are understood to be calibrated.

To date, three CDMs have been proposed to describe the distortion of an image caused by radiation-induced CTI effects. These were described at the beginning of Chapter 4. In this study CDM02 is adopted for predicting the image shape since this is the model that was adopted initially (see Prod'homme et al. (TPH-003)) to simulate radiation damage on a Gaia AF CCD.

Extracting image parameters requires a fast, robust technique that can estimate the centroid and flux simultaneously. The performance of several techniques that may well be used in the data processing are tested by generating statistics on several image parameters, including formal errors on the position and flux as a function of G-magnitude. For complementary studies see, for example, Castaneda & Fabricius (JC-019) and Gardiol et al. (DG-013). The latter describes the structure of the current algorithm being proposed for IDT/IDU Algorithms (IDA).

5.1 Monte Carlo Profile Generator

To perform a statistical study on the performance of image parameter extraction techniques a Monte Carlo Profile Generator (MCPG)¹ was used to generate a large number of single, sampled image profiles varying in position and flux. Each profile is generated from a heavily over-sampled 'seed profile', which has been taken from the analysis of radiation campaign 2 (RC2).

5.1.1 The seed profile

The MCPG is initialised by a seed profile, which is based on real test data taken from the analysis of RC2. Real data is used because this is the actual profile the Charge Distortion Model discussed in the previous chapter is trying to recover. There were many different test conditions that were analysed in RC2. However, for this work a specific test case has been chosen. Compared to other tests, a reasonable number of scans were performed with a nominal $G = 13.3$ source without charge injection and a suitable degree of dither was introduced into the scan timing to sufficiently sample the LSF. A cubic spline function fitted to the test data over a knot sequence that reduced the separation of the knots towards the core of the image provided a reliable representation of the LSF, which could then be sampled to provide a seed profile. Figure 5.1 shows this seed profile occupying a 15-pixel window normalized to contain unit flux. The centroid bias introduced by radiation damage has already been calibrated for as part of the campaign analysis such that the profile is slightly asymmetric relative to the zeropoint, which is aligned to coincide with the undamaged definition of the centroid (see Chapter 2, Section 2.1.5).

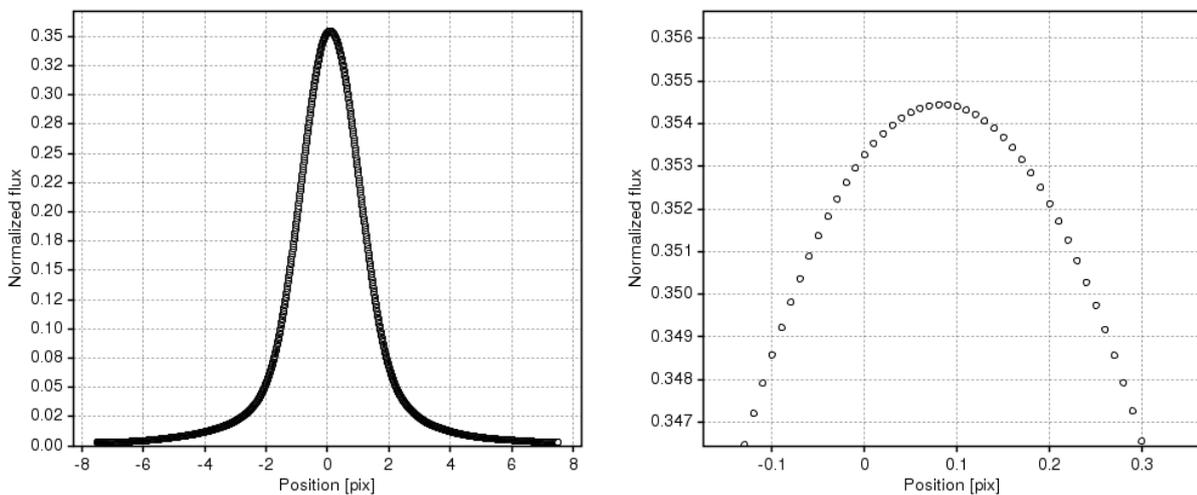


Figure 5.1: (left) A seed profile taken from a test case without CI in RC2 at a nominal $G = 13.3$. (right) A close-up of the peak of the seed profile.

¹Developed by Patricio Ortiz at the University of Leicester.

5.1.2 Random locations

The plot on the right of Figure 5.1 shows how the seed profile has been over-sampled 100x in each pixel. Oversampling the seed profile means fluxes can be extracted at various sub-pixel positions and then reassigned to individual pixel values. The generated profiles must have the same shape as the seed profile to test the fitting procedures, so the randomised positions generated by the MCPG must be rounded up to the nearest 1/100th pixel to allow the seed profile to be sampled. Note that the resolution of the seed profile could easily be extended to say 1/1000th pixel. However, this is a limiting case and it is more useful in the first instance to address a less demanding condition. The seed profile cannot be sampled outside the 15-pixel window so the width of the generated profile is typically 15 pixels minus the maximum predefined allowed shift. For example, for a 12-pixel window the maximum shift that can be applied is ± 1.5 pixels. Note that by generating profiles over 12 pixels, the flux values found at sub-pixel positions within the seed profile are eventually reassigned to a predefined set of positions separated by 1 pixel, e.g. $-5.5, -4.5, -3.5, \dots, 3.5, 4.5, 5.5$ pixels. Every generated profile is thus sampled at the same pixel positions, however, their location is randomised over those pixels.

5.1.3 Random fluxes

Before the MCPG begins the seed profile is normalised to contain unit flux. This is done so that the random fluxes generated by the MCPG can easily be converted into their corresponding magnitudes and the fitting procedure, which is iterative and relies on updates, can generate flux corrections that can easily be fed into the next round of iterations.

In order to generate random fluxes over a predefined magnitude range (e.g. between 13 and 20), the range was divided into bins of width 1 mag and subsequently filled according to a flat distribution by the corresponding fluxes. Choosing a flat distribution was purely a way of ensuring a sufficient spread of profiles over the magnitude range. Later on this distribution could be modified to instigate a predefined luminosity function. However, for the purposes of this work the shape of the flux distribution was not important. Calculating the number of profiles per bin is simply found by dividing the number of profiles to generate by the number of bins. If this value turned out to be a non-integer and rounded down, the total number of profiles to generate may not have been reached by the time each bin is full. In this case, any remaining profiles needing to be generated were placed in the faintest bin. Once the flux levels are set Poisson noise can be added to each profile.

Both the shifts and scaling factors produced by the MCPG were stored for later use to evaluate the performance of the image parameter extraction techniques.

5.2 Extracting image parameters

5.2.1 A first approximation

The fitting procedures require initial estimates of the image parameters, so to provide an initial boost several methods are considered here to estimate the initial centroid. These are Tukey's Biweight (TBW),

the Barycentre (BAR) and Mirror Cross-Correlation (MXC).

In Lindegren (LL-068), several constraints are stated that the line spread function $L(x)$ must satisfy. One of them is that "the origin of x in $L(x)$ corresponds to a suitably defined centre of the LSF". For this the following constraint is adopted:

$$\int_{-\infty}^{\infty} L(x)w(x) dx = 0, \quad (5.1)$$

where $w(x)$ is a weighting function. This defines the centroid as the location of the LSF centre among the samples. The choice to use Tukey's Biweight as a weighting function in this description of the centroid of an AF LSF was made in this technical note, where a trade-off study was performed between a number of other centroid weighting functions. TBW is defined as follows:

$$w(z) = \begin{cases} z(1 - z^2)^2 & \text{if } |z| < 1 \\ 0 & \text{otherwise} \end{cases} \quad (5.2)$$

where $z = x/s$, a normalised argument, and s is a scale parameter determining the width (in x) over which the function operates. This was originally fixed to be 2.7 pixels but later updated to 2.9 pixels for better performance (see Lindegren (LL-069) (a FORTRAN algorithm is also given here for estimating the centroid)). This approach employs Newton-Raphson iteration to update the centroid definition where the initial estimate corresponds to the position of the largest flux in the image.

The BAR method equates the initial estimate of the centroid to a point at which the average flux can be found. Mathematically, each position in the image is multiplied by the flux at that position and divided by the total flux in the image.

Lastly, the MXC method estimates the centroid of an image by finding the midpoint between itself and its mirror image. It begins by cross-correlating the mirror image with itself, then interpolating the cross-correlation curve with a quadratic function. The position of the peak of the interpolation corresponds to twice the centroid estimate, therefore, an initial estimate is found by dividing the measured delay by two.

An initial estimate of the scaling factor is found simply by summing over all discrete fluxes in the window.

If the observed image² is used (i.e. the original cubic spline function that was sampled 100x to produce the seed profile) to refit the generated profiles the problems associated with not knowing the profile exactly (the observed image and seed profile are exactly the same shape) can be ignored; the quality of each initial estimate procedure can be evaluated instead. The plot on the left of Figure 5.2 shows a histogram of the remaining offsets between the initial estimate of the centroid and the original shift that was applied to the seed profile. 10000 profiles have been generated between $G = 11$ and 20. By subtracting these two values an approximation for the offset of the centre of the seed profile from the zeropoint may be recovered, which equates to ~ 0.09 , ~ 0.05 and ~ 0.08 pixels for TBW, BAR and MXC respectively. In the campaign analysis the centroid bias due to radiation damage at $G = 13.3$ without injection was found

²The "observed" image here refers to what we actually "observe" in the real test data, as opposed to the "predicted" image which will be discussed later.

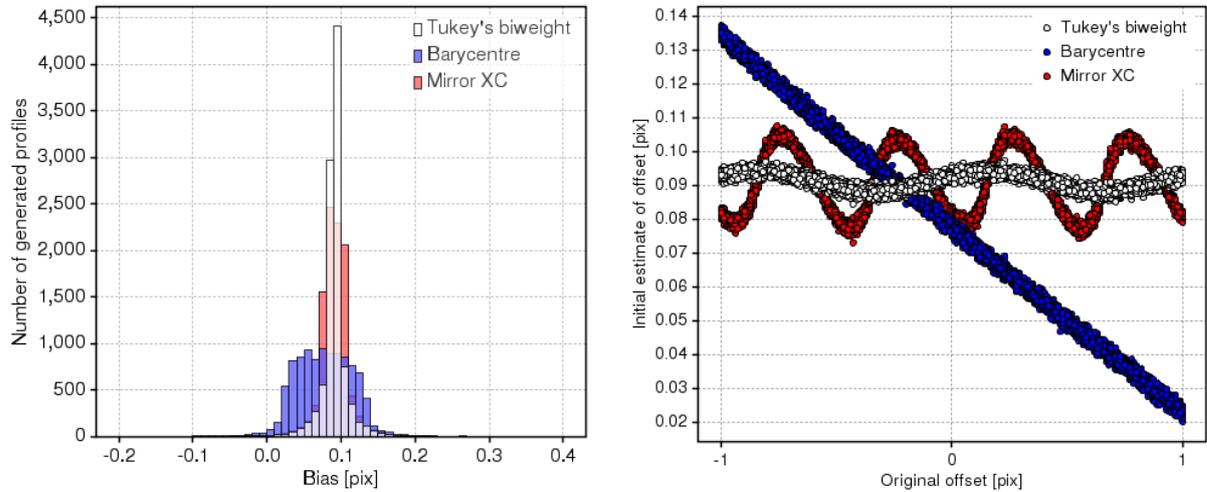


Figure 5.2: (a) Histogram of the offsets remaining between the initial estimate of the centroid and the original (random) shifts that were applied to the seed profile. The profiles are generated between $G = 11$ and 20. The position the histogram converges on tends towards the centroid of the seed profile (b) Comparable biases shown in (a) plotted as a function of the original shift. Here, the magnitude range is limited to lie between 11 and 12 to suppress the noise component.

to be 0.094 pixels. TBW shows the least ambiguity in the centroid measurement, closely followed by MXC. On the right of Figure 5.2 the remaining offsets between the initial estimate of the centroid and the original shift are plotted as a function of the original shift. Here, 10000 profiles have been generated between $G = 11$ and 12 to reduce the noise level³. On a sub-pixel level, there are clear dependencies between where the image is located and the initial estimate of the centroid, which must be taken into account when deciding which method to use. As TBW appears to be the most robust of the three this definition shall be adopted for the remainder of this study.

5.2.2 A final determination

There are again several methods available that have been used to finalise the determination of the image centroid and flux. They are Householder Least Squares (HHLS) with first-order Taylor expansions, a Maximum Likelihood method (ML1) that was developed specifically for centroiding on 1D sample data, not taking into account CTI effects, and a generalised Maximum Likelihood method (ML2) that was developed to handle fitting 1D and 2D data as well as distortions caused by CTI. In terms of feeding the iteration all three methods perform essentially the same procedure: a) evaluate the model value and its first derivative at the observed location relative to the current centroid estimate; b) evaluate the residual between the observed and model value; and, c) determine weights, which are set as the inverse of the observation errors. Having acquired these values, this is where the three methods diverge. Householder least squares with first-order Taylor expansions essentially linearises the problem of finding the corrections to the centroid and flux. The maximum likelihood methods maintain the non-linearity of the problem and attempt to solve for the corrections via the steps outlined in Appendix A.2. Whereas ML1

³The resonances are still present at fainter magnitudes but they are lost in the noise.

is specified in terms of 1D non-overlapping images, ML2 is constructed to be more general and handle situations where CTI may be present.

In accordance with the Gaia windowing scheme for bright images, both the observed (RC2) profile and predicted (CDM02) profile have been fitted over a 12-pixel window. The effect of varying this window size is studied in section 5.5.

The final determination of the image parameters is an iterative procedure so the cutoff has been set to 10^{-3} pixels. For simplicity, the DOB has been set to zero e^-/pixel and the readout noise (RON) to $10^{-3} e^-/\text{sample}$. If the RON is set to zero, it appears Maximum Likelihood 2 (ML2) struggles to calculate the formal error at $G > 20$.

5.3 The observed image shape

To begin with, how well each fitting procedure performs is evaluated by fitting the observed (RC2) image profile to the generated data. Since this profile is also used to generate the sample data, both the seed profile and the observed image profile are the same. The final errors in the fitting should therefore match the original Poisson noise added to the generated profiles. Using TBW to estimate the initial centroid Figure 5.3 shows how each of the fitting procedures perform when extracting image parameters.

At first glance it would appear that the formal errors on both the position (plot a) and flux (plot b) for Maximum Likelihood 1 (ML1) are far better than for either of the other two methods. However, this is not supported by plots (c) and (d), which show the standard deviation of the fit and Gaussian variable distributions (see eqn. 2.8) respectively. Both of these plots indicate ML1 is not converging on the best fit. In particular, the Gaussian variable illustrates how random the residuals are after the final fit has been made. With a bad fit the normalised errors are no longer random and stray outside the Gaussian curve. To explain plots (a) and (b), ML1 must be producing the wrong normalised errors, over-estimating the formal errors. This will be made clearer in the next figure.

Figure 5.3 also shows that:

- Across the whole magnitude range ($G = 11$ to 21) HHLS and ML2 produce very similar results.
- The normalised errors for a small percentage of fits are always larger than expected, creating wings (see plots (a) and (b)) in the error distributions. If the HHLS and ML2 methods are correctly estimating the formal errors, then there must be an additional factor limiting the accuracy of the final solution. It is thought the window size of the fitted image may be the cause of this limitation.
- In plot (d) the Gaussian variable continues to appear beyond the highest positive value shown. This range has been limited so all values beyond 7.5 are placed in the right-most bin.

The same results are produced when either of the three initial centroid estimation methods are used.

Lastly, Figure 5.4 shows how the formal errors on the position (left) and magnitude (right) vary as a function of magnitude. The magnitude range ($G = 11$ to 21) has been split into bins, each containing 100 flux values. The datapoints in Figure 5.4 show the arithmetic mean of the fluxes in each bin. The fluxes have been converted into magnitudes according to eqn 2.2.

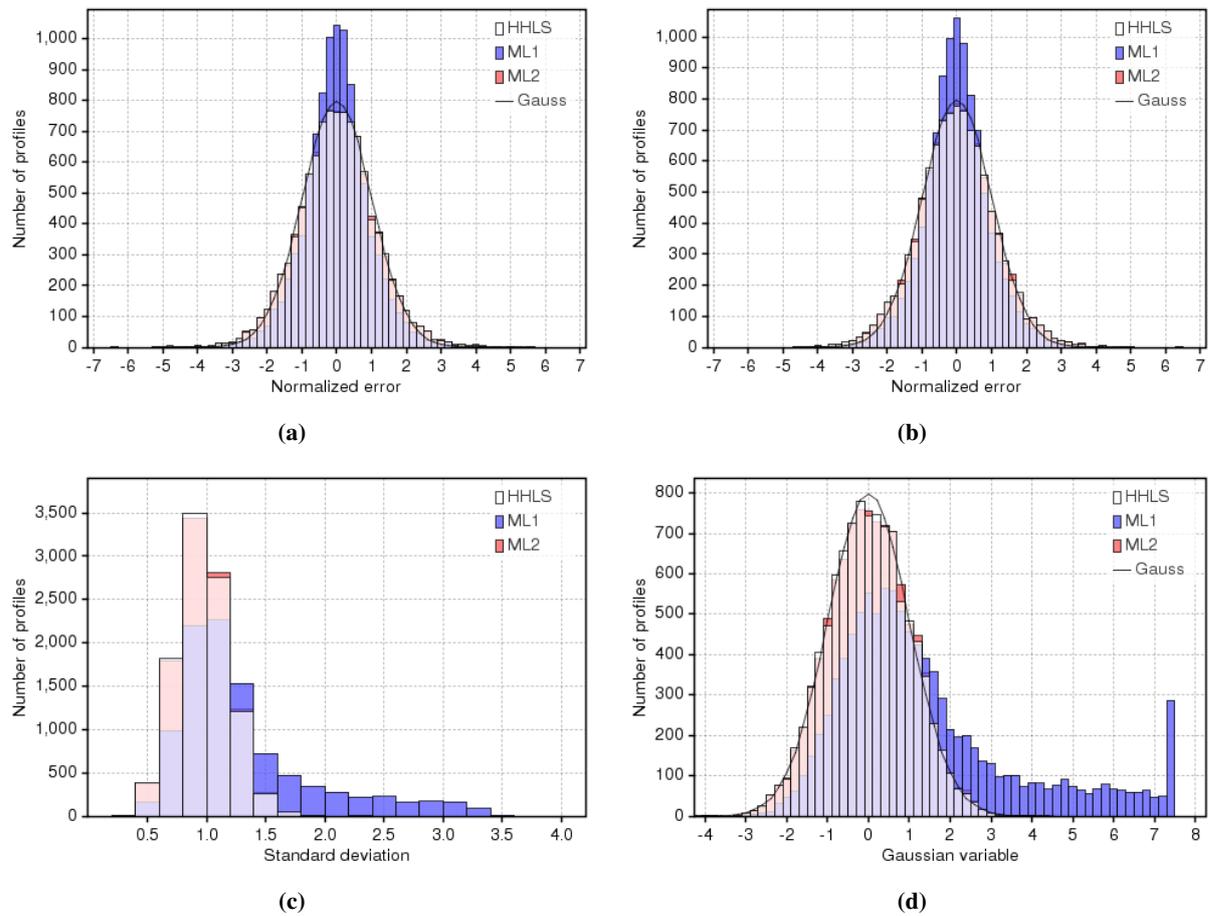


Figure 5.3: Observed image shape parameter statistics. Histograms for (a) the normalised positional errors, (b) the normalised scale factor (total flux) errors, (c) the standard deviation of the fit, and (d) the Gaussian variable for $G = 11$ to 21. The curve is a $\sigma = 1$ Gaussian containing the same number of profiles.

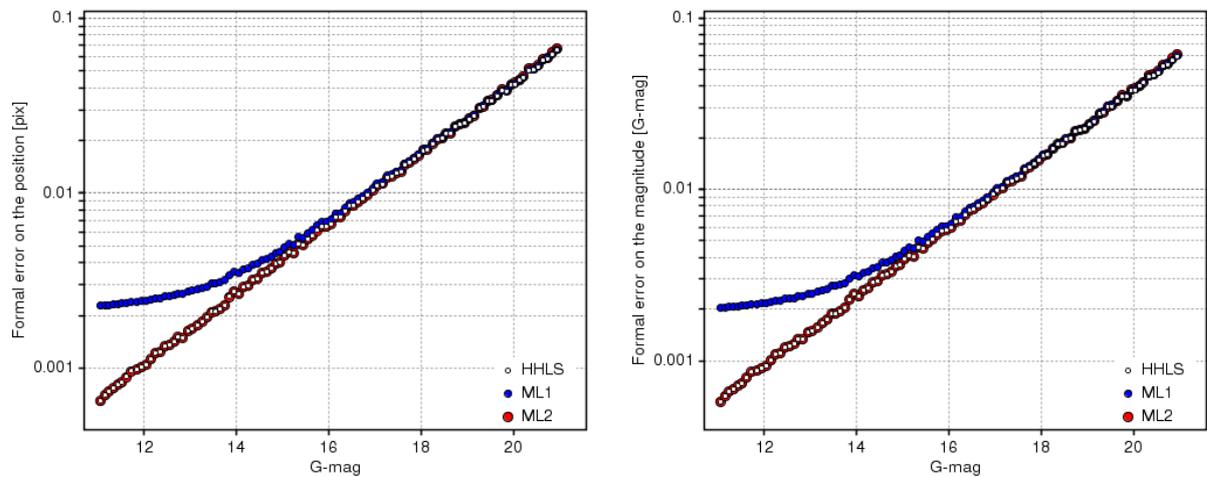


Figure 5.4: Observed formal errors in position (left) and magnitude (right) as a function of magnitude.

Following on from the discussion above, it is in fact the formal errors on the bright images that ML1 over-estimates. This is a problem with the ML1 method and would require improvement if it were to be used. Both HHLS and ML2 recover the Poisson noise added to each image profile in the MCPG. At $G = 12$ the formal error on the position and magnitude is expected to be around 10^{-3} pixel and 1 millimag respectively.

5.4 The predicted image shape

The predicted image shape⁴ has been derived on the CEMGA platform using CDM02 in the same way as was described in Chapter 4 when CDM02 was being evaluated. To produce the predicted image, in the first instance, an undamaged LSF taken from the same test case in the radiation campaign analysis as the observed image (i.e. at $G = 13.3$ and without injection) was used to create an artificial PSF, which could then be fed through CEMGA to simulate this test case. CEMGA models the effects of CTI over the Gaia focal plane so using CDM02, the output image is distorted by radiation damage. If the model and its parameters are correct, this image should match the image observed in the real test data. Thus, in this section it can be shown how well the "predicted" image (i.e. the output image from CEMGA using CDM02 to include CTI effects) matches the actual "observed" image (the same seed profile is used as before) and how the formal errors in position and flux may appear when an insufficient fit is made. Figure 5.5 shows how well the predicted image shape matches the observed image shape, both at face value (left) and in terms of residuals between the two (right). The output from CEMGA is in practice a series of samples and the fitting procedure here requires a continuous function in order to interpolate the flux at any position. For this reason, the predicted image in Figure 5.5 is a cubic spline fit to the sampled output from CEMGA. The number and positions of the knots have been chosen to try and reproduce as closely as possible the original output.

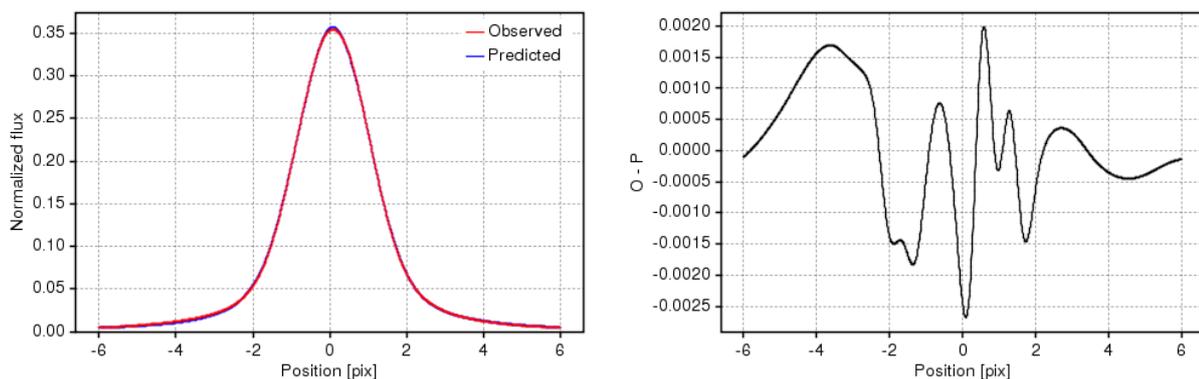


Figure 5.5: The predicted vs observed image shape. (left) Normalised image shapes. (right) Residuals between the two image shapes.

As Figure 5.5 shows there is little difference between the two fits. Now the predicted image profile is fitted to 10000 generated profiles using the same seed profile as before. The image parameter statistics

⁴Provided by Thibaut Prod'homme (Leiden Observatory) and Michael Weiler (Meudon Observatory).

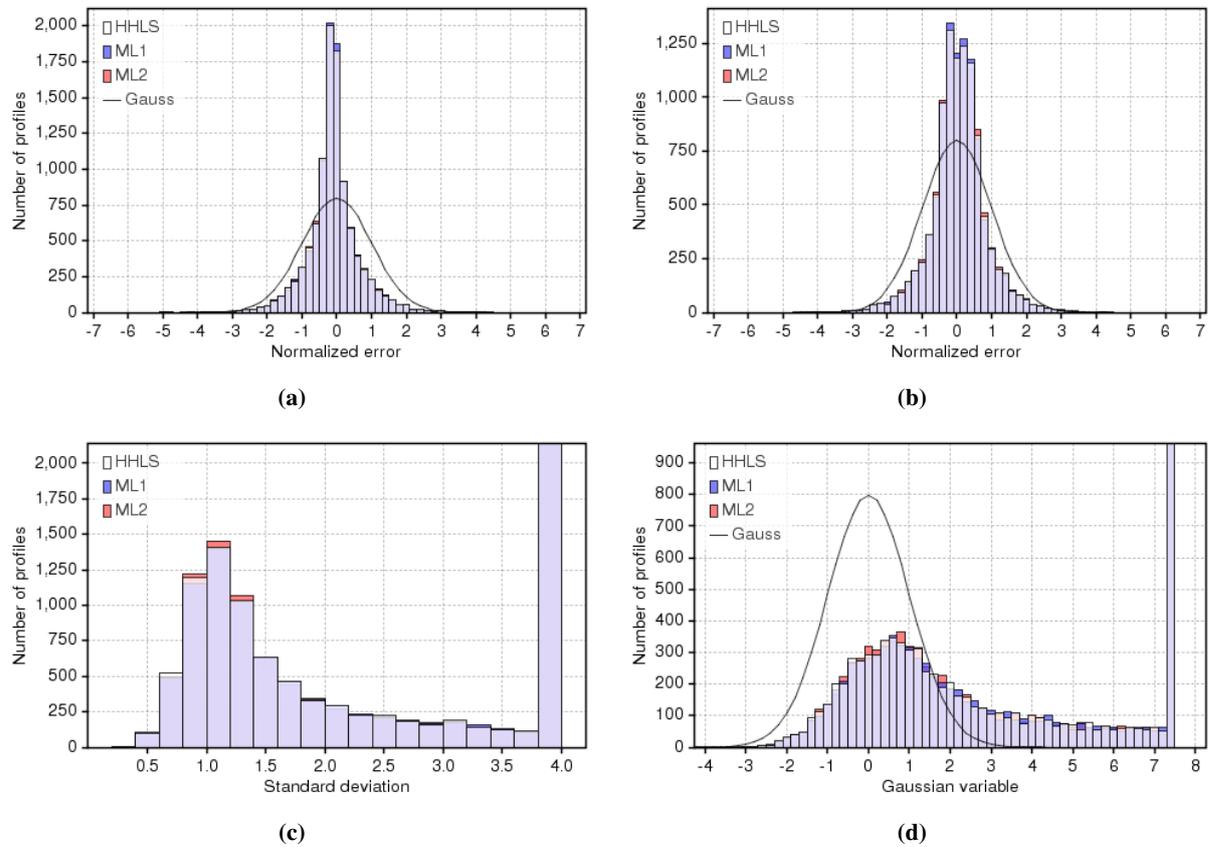


Figure 5.6: Predicted image shape parameter statistics. Histograms for (a) the normalised positional errors, (b) the normalised scale factor (total flux) errors, (c) the standard deviation of the fit, and (d) the Gaussian variable for $G = 11$ to 21. The curve is a $\sigma = 1$ Gaussian containing the same number of profiles. In plots (c) and (d) the right-most bin contains all profiles that returned higher values.

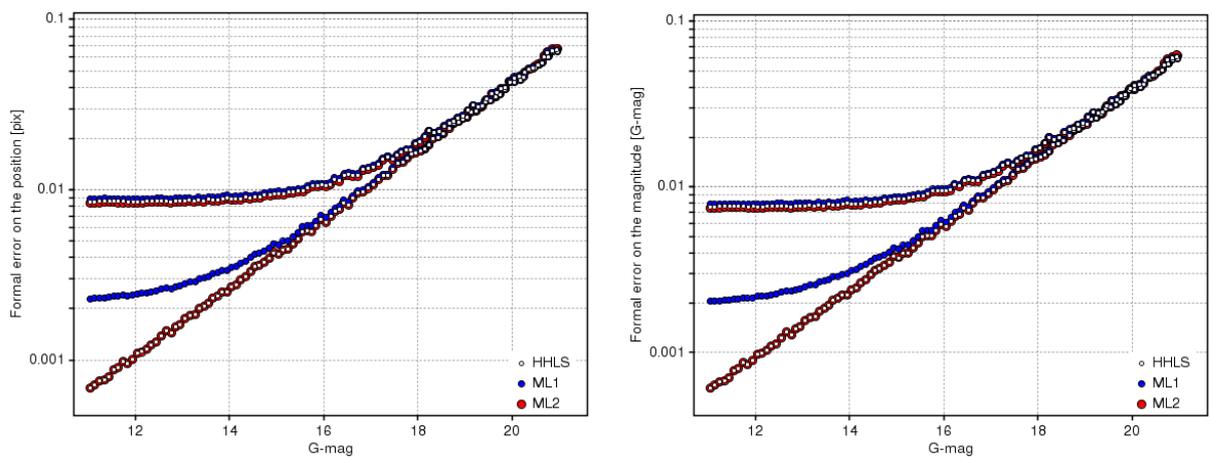


Figure 5.7: Formal errors in position (left) and magnitude (right) as a function of magnitude when both the predicted and observed image shapes (taken from Figure 5.4) are used to fit the generated profiles. The predicted formal errors when using any of the three fitting procedures can be seen to begin to level off when $G \lesssim 18$.

are shown in Figure 5.6 using TBW as the initial estimate of the centroid.

Again, plots (a) and (b) suggest the current CDM02 prediction fits very well to the observed image shape, however, plots (c) and (d) suggest the opposite. If the normalised errors on the image parameters are correct, the formal errors on both the position and flux are being over-estimated. In this case, the cause is probably the systematic differences that still exist between the predicted and observed image shapes. Plot (d) shows that the residuals in many fits are no longer Gaussian, which is due to a systematic offset remaining between the two image shapes. The image parameters may be found but if the two shapes are not identical the formal errors will not validate those image parameters.

The same results are produced if BAR or MXC are used to estimate the initial centroid. The formal errors in both position (left) and magnitude (right) as a function of magnitude are shown in Figure 5.7. Only slight deviations between the predicted and observed image shapes are required to reduce the relative accuracy of the centroid determination from $\sim 1 \times 10^{-3}$ pixel to $\sim 8-9 \times 10^{-3}$ pixels at $G = 12$.

5.5 Varying the window size

Up to now the width of each of the generated profiles has been fixed at 15 pixels and a cubic spline function has been constructed in a window overlaying the central 12 pixels. To investigate the effects of varying this window size it is firstly assumed that the wings of the original seed profile continue to decrease linearly. Two straight lines are fitted to a subset of samples between positions ± 7.5 to ± 7.4 (10 points), and these lines are extended out to cover an extra 3 pixels either side. This extended (original) seed profile is fitted again with a cubic spline and resampled every 1/100 of a pixel to generate a continuous new seed profile of width 21 pixels. The window size of the fitted image may then take on any value close to this width.

With this new seed profile the window size has been varied between 6 and 18 pixels to see the effect this has on extracting accurate image parameter statistics. Using TBW with HHLS, Figure 5.8 shows how the image parameter statistics vary as a function of window size for image shapes between $G = 11$ and 12. From Figure 5.8:

1. The position and flux normalised errors (plots (a) and (b) respectively) are smaller and better match the expected distribution when the window size is larger.
2. A reduction in the number of degrees of freedom, brought about by a smaller window size, increases the spread in the standard deviation (plot(c)).
3. Regardless of the window size, the noise in the final fit is sufficiently randomised when the generated profile and fitted image are the same shape (plot (d)).

Which suggests the following:

1. The window size has a direct impact on the accuracy of the final formal errors.
2. Placing a bright image in a small window that truncates the profile's wings increases the formal errors on the image parameters.

3. The wings in the error distributions seen in Figure 5.3 can be attributed to having a 12-pixel window.
4. Both the core and the wings of an image are required to estimate an accurate final formal error.

Figure 5.9 shows how the formal errors in position (left) and magnitude (right) vary as a function of magnitude when the window size is 18 pixels.

It is known that ML1 is incorrectly estimating the formal errors for the bright stars and this is seen again here, though the extra information in the window provides greater constraints on the fitting. There is also a hint that the formal errors from HHLS and ML2 diverge for very faint images. Further investigation shows this is the case when $G \gtrsim 23$.

An attempt was also made to extend the length of the CDM02 predicted image shape out to 21 pixels.

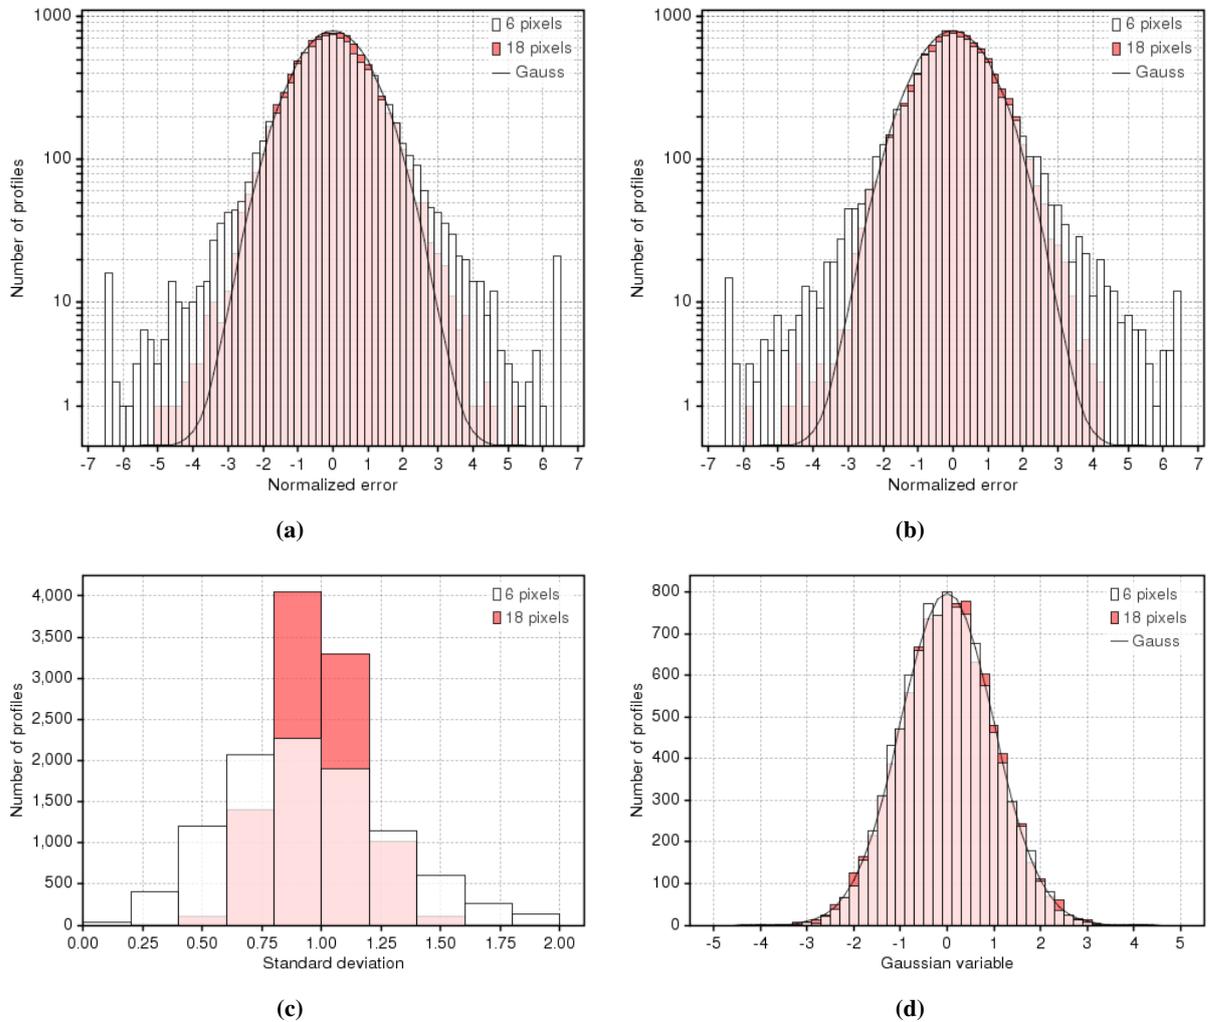


Figure 5.8: Image parameter statistics as a function of window size. (a) Histograms of the normalised positional errors with a $\sigma = 1$ Gaussian containing the same number of images for a window size of 6 and 18 pixels. (b) Comparable results to (a) for the scaling parameter (total flux). (c) Comparable results to (a) for the standard deviation. (d) Comparable results to (a) for the Gaussian variable.

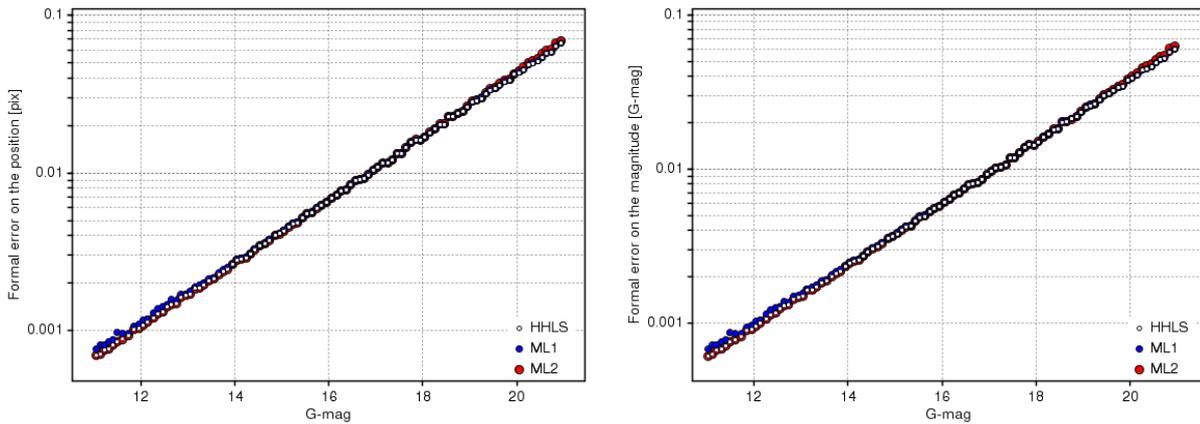


Figure 5.9: Formal errors in position (left) and magnitude (right) as a function of magnitude using the observed image to fit the generated profiles over an 18-pixel window.

However, since a spline function must be fitted to the extended profile in order to fit it to the generated profiles (i.e. some ambiguity remained), no clear differences could be seen when compared to having a 12-pixel window.

5.6 Conclusions

Using a Monte Carlo profile generator together with the output from RC2 and CDM02 it has been possible to investigate the performance of several image parameter extraction techniques, as well as initiate a study on DPAC’s current radiation calibration strategy.

Before extracting final image parameters, three methods (TBW, BAR and MXC) were used to estimate the initial centroid. For the least ambiguity on the estimate as well as minimal dependency on the sub-pixel position, TBW gave the most convincing results, closely followed by MXC. That said, one advantage MXC has over TBW is the absence of any predefined scaling parameter.

Firstly, the observed (RC2) image shape has been fitted to 10000 generated profiles of the same shape as the observed image over a 12-pixel window to evaluate the robustness of final image parameter extraction techniques. The findings from these tests are as follows:

- ML1 is unable to estimate correctly the formal errors for $G \lesssim 18$. Presumably, for fainter images the noise is too great to disentangle the formal errors on the fitting.
- ML2 requires the RON to be greater than zero when estimating the formal errors on very faint images.
- HHLS and ML2 produce similar statistics over the full magnitude range tested ($G = 11$ to 21). The formal errors recover the expected Poisson noise.
- Using HHLS and ML2 the normalised errors for a small percentage of fits are larger than expected. This has been found to be related to the size of the window.

Secondly, the predicted (CDM02) image shape has been fitted to 10000 generated profiles modelled on the observed (RC2) image shape to evaluate the output of CDM02. The findings from these tests are as follows:

- All three methods over-estimate the formal errors on the image parameters. This is attributed to the systematic offsets between the two image shapes.
- These offsets remove the randomness in the final residuals.
- The best formal errors in the centroid and magnitude are found to level off at $\sim 8-9 \times 10^{-3}$ pixels and $\sim 7-8$ millimags respectively, towards the brighter end of the magnitude scale. This can be compared to 10^{-3} pixels and 1 millimag, which are the expected noise on the position and magnitude at $G = 12$ respectively.

Even though it appears the prediction of the CDM02 model is not able as yet to completely recover the Poisson errors on the position and flux, there are some factors that should be taken into account and could be improved upon in future work:

- A cubic spline fit had to be made to the CEMGA output to allow the iterative image parameter extraction process to proceed. It was through trial and error that such a fit was made and slight differences remained between the two shapes. The fit therefore was not entirely representative of the simulated, damaged image.
- The undamaged PSF fed into the CEMGA simulation is an approximation of the true PSF that existed in the second radiation campaign. No information was retained on the AC profile of the PSF so the output from CEMGA using CDM02 is again an approximation to the actual damaged image observed in RC2. To compensate for this, it may be useful to feed in the undamaged 1D LSF directly, assume the binning does not change the shape of the LSF significantly, and compare this output to the real test data.

Extending the seed profile to 21 pixels, and thus extending the (fitted) image profile to cover 18 pixels, identified the effects of varying the window size. Compared to a 6-pixel window there was a notable reduction in the number of normalised error outliers. The window size does have an impact on image parameter extraction and in order to produce an accurate fit the wings of the image must be included or estimated.

Lastly, there are several aspects of this work that could be taken further:

- The test case chosen for this study does not include charge injection. Therefore, the formal errors observed here could be taken as the worst-case scenario. Consequently, it would be expected that including charge injection would improve the accuracy of the final formal errors.

- The shape of the seed profile has been taken from a $G = 13.3$ damaged image, which is known not to have the same shape as a $G = 20$ damaged image. With better quality experimental data, as expected to come from radiation campaign 4, the seed profile could be customised for each magnitude bin.
- In practice the DOB and RON will be non-zero and the window size will vary depending on the magnitude. These are parameters that have not been investigated but are expected to be understood as part of the data calibration.
- Processing time is also a key issue not considered here. It may help decide which procedure, both for the initial centroid estimate and final image parameter extraction, to adopt for the data processing.

6

Conclusions and Final Remarks

IN this brief chapter I will pull together the results and conclusions discussed in previous chapters and reiterate why a detailed study of radiation damage on the Gaia Astrometric Field (AF) was required, what we have learnt from the radiation campaigns (RCs), and how they have helped shape the radiation mitigation strategy both on board Gaia and on-ground as part of the data processing.

During the development phase of Gaia it became apparent that radiation damage threatened the scientific requirements of the mission and had the potential to endanger a sizeable portion of the catalogue at the faint end of the magnitude scale. Without the possibility of performing microarcsecond astrometry Gaia would not be able to meet its primary objective. Faced with such a challenge ESA, DPAC and industry went about planning and performing a series of radiation campaigns. EADS Astrium conducted the tests in-house and replicated as closely as possible the conditions the Gaia focal plane would be under when situated at the Earth-Sun L2 point at the end-of-mission. Only then could it be understood: a) how the radiation domain Gaia was expected to operate in could affect every instrument on board; b) how the damage could be characterised; and, c) how image distortion could be modelled and calibrated in the early stages of the data processing.

6.1 On the hardware and test setup

The radiation campaigns are the primary pre-launch source of information for understanding and developing the means to handle radiation damage. Irradiated CCDs have been analysed in the laboratory for other space missions. However, there are some features that Gaia in particular must deal with. First and foremost by analysing the raw data, many unexpected technical features were found:

Non-uniform background In RC1 the background in the AL direction was seen to increase linearly halfway down each scan and the background in the AC direction varied systematically; in some cases the difference between two adjacent samples in the readout could be as high as $\sim 40e^-$. The fluxes in the stellar field did not reflect the large differences seen in the background so it was believed this effect was being introduced outside the image area (IA). Behind each CCD a Proximity Electronics Module (PEM) converts the readout into a digital unit and also adds an electronic offset. It was thought this electronic offset was causing the variations in the underlying background and as a result could not be predicted using only the prescan sample. Further statistical analysis supported the assertion that an instability between the CCD and PEM was the cause of unexpected systematics, correlations and global variations between the background signals measured between sample readout positions. If the "PEM anomaly", caused by simultaneous transitions in the IA and serial register (SR) as well as SR flushing, were not to exist it is thought that the background level and noise would be consistent with the expected statistics. As the PEM anomaly is also seen with a flight-model PEM, these effects will have to be dealt with in the data processing.

In RC2 the background was found to contain several sources of light pollution: light leaking around the mask onto the CCD before and after the stellar field; many "cosmics" hitting the CCD at random times, and a charge trail appearing at the beginning of every scan. The latter can be explained by having to keep the light source on (and therefore stable) between scans as the test bench was reset to its initial position. In RC3 these features were mostly removed, though not entirely.

CCD stitches In RC2 CCD stitches were discovered to have a significant impact on the final centroid positions in the readout, amounting to ~ 0.01 to 0.02 AL pixels. In RC3 pixel-line displacements between CCD stitches were found to shift centroid positions as much as 0.01 AL pixels. In other words, systematic variations of this order could be expected to arise between any pixel line, which presumably could vary between CCDs. In effect, each pixel line must be individually calibrated. Activating a gate increased the systematic bias observed close to a CCD stitch by reducing the integration time. For this reason, gates must also be calibrated individually.

Charge injection structure In RC3 the charge injection (CI) was driven from both sides of the CCD and here two effects were noticed for the first time: a) the bottom part of the CI profile that spent longer in the SR was damaged more than the top part; and, b) the part of the CI profile that was read out closer to the bottom part of the CCD showed the background level increasing prior to the CI. The first is thought to be due to CTI in the SR. The second is currently unexplained. Labelled the "divination anomaly", this feature has not been seen in tests with the flight PEM. Since a charge injection involves lowering the diode voltage to spill charge over the gate into the image area, it may be that some form of "over-spill" is occurring, i.e. the voltage in the gate is not high enough initially to contain the collected charge, just prior to the injection.

6.2 On the process of calibrating the campaign data

Before extracting centroid and flux values a number of preliminary steps had to be taken to calibrate the campaign data.

Background subtraction Due to the PEM anomaly the background had to be removed on a per channel basis, each channel corresponding to a part of the readout that contained a part of the stellar field. Whether or not a flat or local background cut could be made depended on the spacing of the images. In RC1 a local cut was viable as the stars were faint enough to avoid any noticeable overlap between neighbouring images. However, in RC2 and RC3 a flat cut had to be made, at least for $G \lesssim 15$, as the wings in each image overlapped. For stars fainter than $G = 15$ a local cut could have been made, but this did not take full account of the charge released behind the CI. The data for RC3 became available while the RC2 analysis was being carried out, so the removal of the CI tail was left for the analysis of RC3. To make a flat background cut it was necessary to avoid non-linear regions in the background, e.g. bumps appearing where light had leaked around screws holding the mask in place. So the mean background level was estimated towards the beginning of each scan, several thousands of TDI periods before the first image would appear. In reality, estimating the background in this way will not be possible as only windows of data onboard Gaia are transmitted to ground. Instead, virtual windows and windows around very faint images will provide the local background.

Geometric calibration The geometrical calibration of the test setup required a set of "well-behaved" non-irradiated stars in both the reference and irradiated data. However, in RC2 it was discovered that one of these sets (out of a total of 2) were aligned with a CCD stitch. This had the effect of systematically displacing their centroid positions and returning unphysical centroid biases for the rest of the field. To compensate for this, stars appearing further up the mask that had passed over the damaged part of the CCD were assumed to be undamaged and included in the calibration. Fortunately, this turned out to be a reasonably valid assumption as the stars that had passed over the CCD previously could mitigate a significant amount of the radiation damage. In RC3 CCD stitches were avoided and the calibration proceeded as originally planned. However, pixel-line displacements were discovered to be just as prevalent as CCD stitches, so again two of the non-irradiated sets of stars had to be discarded. Of course with hindsight the mask would be designed to avoid the situation where non-irradiated stars pass over CCD stitches and pixel-line displacements. So in RC4, at least in some tests, the *entire* CCD is being used, with two appropriately sized non-irradiated sections either side of the irradiated region, as in RC3, to deal with the biases incurred by stitches.

Flux calibration As the illumination conditions over the mask were non-uniform a careful analysis of the variations in the fitted Global Reference Curve (GRC) and local LSF image shape (e.g peak flux and FWHM) across the stellar field were required. This found that the total charge under the fitted GRC was accurate to within 1%.

A number of recommendations were made after each campaign to improve the reliability of the

output. All of these, time allowing, have been incorporated into the latter stages of RC3 and now RC4. These include reference data taken at the same magnitude as irradiated data; a sufficient number of scans to sample the LSF; dither between each scan to sufficiently sample the LSF, as well as avoidance of stitches. Providing feedback to industry from an independent analysis has been one of the primary goals of this thesis. Based on the quality of the data that is now being observed in RC4, I believe this requirement has been met.

6.3 On the effect of the diffuse optical background (DOB)

As Gaia is a scanning satellite, the average level of background for any one observation will be very low. It was unclear, therefore, if a DOB would in any way mitigate radiation damage. Tests performed prior to RC1 hinted that the DOB could mitigate some of the damage incurred by radiation, in particular at the fainter end of Gaia's visual range. So in RC1 a range of magnitudes fainter than $G = 15$ were tested alongside three different levels of DOB: 0, 5 and $10 \text{ e}^-/\text{pixel}$. The results on the charge loss, which were clearer than those for the centroid bias, indicated that even with a low-level background some form of mitigation did occur. The fainter the star, the more pronounced the effect. A further improvement was seen when the DOB level was raised from 5 to 10 but less so, indicating a limit to the effect. In conclusion, the DOB cannot be relied on to mitigate radiation damage entirely, but must be included in the radiation calibration. Physically, it may be said that many of the traps filled by a DOB have long release time-constants, which explains why there is a limit to its effects: there are a finite number of these types of traps.

Image distortion: longevity of the trailing edge As part of the RC analysis it was possible to show how radiation damage and possible mitigating strategies could affect the shape of the distorted image. Tests prior to the RCs suggested that charge was taken from the leading edge and added to the trailing edge, creating a charge trail. However, in the RCs something slightly different was found. The leading edge was indeed eaten away but the trailing edge remained intact without a trail appearing behind the image. Even beyond $G \gtrsim 17$ when the trailing edge would begin to erode no charge trail was seen behind the image. This may be linked to operating the CCD at 163K. Different types of traps become dominant at different temperatures. At this temperature CTI effects in both the IA and SR were found by Sira to balance out, which may also have suppressed those traps in the IA with trap release time-constants comparable to the time required to release charge back into the trailing edge. Charge recovered by increasing the DOB would appear predominantly in the trailing edge of the image before the peak and leading edge.

6.4 On the effects of charge injection (CI)

In the RC2 analysis the charge injection tail was not removed as part of the background subtraction and the injection was only made from one end of the CCD. This is not optimal for making conclusive remarks on the charge injection strategy. So bearing this in mind, the following can be said:

- A CI close to 8.06V or $\sim 16500 e^-/\text{pixel}$ appeared to mitigate the bulk of radiation damage seen in the first test at $G = 13.3$ and 15.
- In the second test, with the same magnitudes, the injection duration needed to be only 4 pixels to remove the bulk of the damage. This is currently the nominal injection duration for when Gaia is launched, which balances mitigating radiation damage with minimising dead-time.
- It appears a CI delay of <1000 pixels would be needed to remove radiation damage from the faintest images ($G = 20$) and a CI delay of <100 pixels would be required to remove the bulk of radiation damage from the brightest images ($G = 13.3$). For all magnitudes in between it is expected that some intermediate CI delay would be required. Having an injection delay of <1000 pixels is supported by the RC3 analysis, where charge injection was driven from both ends of the CCD.

While the level and duration are now fairly well understood, the choice for the optimal delay remains unclear. For this reason, it may be the case that the delay will have to be decided upon once Gaia is launched; 1000 pixels being taken as an upper limit. It may turn out that a fixed value is not what is needed and depending on the stellar density, the injection will have to be fine-tuned to particular areas of the sky.

6.5 On the microscopic properties of an irradiated Astrometric Field CCD

By carrying out a First Pixel Response (FPR) analysis of the CI profiles in RC2 it was possible to investigate the trap density as a function of CCD temperature and CI signal level. From this a range of trap densities were observed that grew as the signal level increased. The maximum number of detected traps averaged over four scans at 163K, the Gaia operating temperature, was measured to be ~ 2 traps/pixel. Recall that the CCD is composed of 4500 AL \times 1966 AC pixels and the radiation level was 4×10^9 MeV protons cm^{-2} . Evaluating the total charge loss over the CI profile agreed with the FPR analysis, although the trap density was higher as expected due to a larger signal level. The trap density was also found to increase with CCD temperature. This is thought to be related to the trap parameters that are temperature dependent.

Supplementary Buried Channel Full Well Capacity In addition to determining the trap density the Supplementary Buried Channel (SBC) embedded in each Gaia pixel has been identified. This channel funnels very small signals through the CCD so that the majority of traps spread throughout the rest of the CCD can be avoided. For the first time the Full Well Capacity (FWC) of the SBC has been measured and found to be $\sim 1.1 \times 10^3 \pm 0.2 \times 10^3 e^-/\text{pixel}$. This is in good agreement with the current nominal value of $1300 e^-/\text{pixel}$.

Down to this signal level, which corresponds to $G \approx 19.5$, Gaia can expect to observe signals that are partially mitigated against radiation damage by the action of the SBC. Beyond this magnitude the fainter signals will begin to again suffer radiation damage and it will become more difficult to recover the signal.

This ultimately sets a limit to the Gaia catalogue since CTI effects are a dominant source of noise for Gaia.

6.6 On Charge Distortion Model (CDM) validation

With DPACs approach radiation damage calibration requires a CDM, which must be rigorously tested before it can be implemented. To date, only CDM02 has been tested, having been incorporated into the CEMGA platform to assess its performance against RC2 data.

CDM02 and trap parameters While CDM02 appears to be able to fit the data well, both in terms of individual images and fitting a set of images simultaneously, there are certain ambiguities and degeneracies remaining in the simulations that make it difficult to identify a consistent set of trap species that can describe all of the RC2 damaged image shapes together. That said, these limitations may be improved upon. RC4 should record damaged image profiles that are better sampled than those in RC2, the experimental conditions should be more stable, and CDM02 has since been updated to CDM03, removing the unrealistic assumption that a CI fills every trap in the CCD. This study is the beginning of a series of tests to be performed. With better constraints (e.g. fixing the DOB level) the reliability of the tests and the estimation of the trap parameters should improve.

6.7 On extracting image parameters in the data pipeline

The ability of CDM02 to predict the shape of a damaged image as observed in RC2 was investigated statistically using a Monte Carlo approach.

Initial estimate This analysis has been coupled with a performance check on several initial centroid estimate and final fitting procedures. Of the three initial centroid estimate methods, Tukey's Biweight (TBW) appeared to be the most robust and least ambiguous in the centroid determination.

Final estimate Of the final fitting procedures, both Householder Least Squares (HHLS) with first-order Taylor expansions and a generalised non-linear Maximum Likelihood method (ML2) showed the best results and could not be separated. The final formal errors were very similar.

CDM prediction As for the accuracy of the CDM02 prediction, it became clear how slight the variations need to be between the observed and predicted image shapes to introduce noticeable formal errors in the position and flux. For example, the best formal errors in the centroid at $G = 12$ were found to level off at $\sim 8\text{-}9 \times 10^{-3}$ pixels, while the expected noise on the position at the same magnitude is 10^{-3} pixels. This still gives good confidence in the model, which has yet to be optimised.

Window size The effect of the window size on extracting accurate image parameters has also been investigated. It was found that by increasing the size of the window there was a notable reduction in the number of normalised error outliers; the window size does have an impact on the accuracy of the image parameter extraction. Now that this limitation is known it can be included as one of the sources of error in the data processing.

6.8 On current radiation campaigns

RC4 is currently underway which is addressing a number of new and previously identified issues from earlier investigations. A flight-model PEM should be used in some tests as opposed to a laboratory PEM, which caused the divination anomaly. Additional tests are planned for the AF (e.g. possible impact of AC motion and image shape (to be confirmed) as well as the effects of very bright stars or planetary transits on the data) and the RVS instrument as well as taking an extended look at the radiation damage in the serial register. RC4, which includes RVS and photometer sky-like mask tests, is expected to take place throughout the remainder of 2010.

The quality of these tests should exceed what was observed in campaigns 1 and 2, when aided by the recommendations that have been made throughout this thesis.

6.9 Outside Gaia: a study of HST images

Throughout the RC analyses several mathematical tools have been used to evaluate the astrometry and photometry on the AF. The scope for applying these tools to other areas has been looked at. In particular, a study was initiated along with Maya Belcheva at the University of Athens to extract relative positions and fluxes from several Hubble images of the Small Magellanic Cloud star cluster NGC121 in two different bands: V and I. The approach was very similar to that described in Chapter 2, section 2.1.

To begin with we identified the brightest, single, most undisturbed stars in the field, extracted them over a window of given size and stacked them one on top of the other according to the peak flux. Rather than fit a PSF, each image was treated one-dimensionally, firstly in the x-direction, then in the y-direction. To sample a reference LSF at the pixel level each 1D star profile was cross-correlated with an arbitrary reference star profile. Then the cross-correlation curve was interpolated using a quadratic function to oversample the reference LSF at the sub-pixel level (1/100 pixel). A curve was fitted (in this case a double Gaussian) to form the reference LSF. The centroid of this curve was aligned with the zeropoint and all other images were fitted with this curve to extract relative positions and fluxes. The process was repeated for the y-direction. Finally, the results were to be combined to build up a 2D map of the image.

Since HST data was available in both the V and I bands the above procedure could be repeated for both bands. Then by subtracting the fluxes for each star in both images a colour-magnitude diagram could be constructed, giving some insight into stellar evolution within the cluster.

One other application that could stem from this data is a study of the proper motions of star clusters within the SMC. Images are available at multiple epochs, which may distinguish clusters moving relative

to the cloud. The ACS HRC pixel size is $0.028'' \times 0.025''$. With typical magnitudes in NGC121 of $V > 19$, the positional accuracy would be no better than 1/100 to 1/10 pixel. This means the best accuracy on the sky would be around 0.1 to 3 mas. This is in good agreement with Piatek et al. (2008) who found the mean proper motion for the SMC to be $(\mu_\alpha, \mu_\delta) = (75.4 \pm 6.1, -125.2 \pm 5.8)$ mas per century. In the fields of red giants observed by Propris et al. (2010) in the SMC, the radial velocities can vary between 100-200 kms^{-1} . If a proper motion of 1 mas yr^{-1} at 1000 pc is equivalent to a velocity of 4.74047 kms^{-1} , then at a distance of $\sim 60 \text{ kpc}$ (i.e. the distance to the SMC) the velocity would have to be about 280 kms^{-1} to produce a proper motion of 1 mas yr^{-1} or about one tenth of this to produce a proper motion of 0.1 mas yr^{-1} . Since the radial velocities are above this limit and the HST pixels could handle this resolution, HST could, in principle, provide proper motion measurements of star clusters within the SMC.

6.10 Final remarks

This thesis was motivated by the need to understand and characterise the effects radiation damage would have on the Gaia Astrometric Field at the pixel level. These studies were performed in parallel with industry in order to provide an independent view of CTI effects for Gaia and raise the level of awareness on what to expect in the data once it arrives. Gaia requires highly accurate positional measurements in order to reach microarcsecond astrometry levels by the end of the mission. This essentially requires all known forms of CTI mitigation to be applied including cooling the CCD, using charge injection and developing a CTI mitigation model that is robust and can cope with changing conditions across the focal plane. This model must also include the effects of a diffuse optical background that is now known, even at a very low level, to partially mitigate radiation damage. Optimising the charge injection strategy may have to be done during the mission itself as the time between charge injections has not been definitively found here. However, this may be better in some ways as only in the early stages of the mission does the real picture emerge.

Clearly, CCD radiation damage is an area in the future that will hinder astrometric missions beyond Gaia using the same hardware to reach nanoarcsecond astrometry. Many of the features highlighted in this thesis will almost certainly have to be taken into account, if radiation damage cannot be dealt with by more advanced CCD designs.



Analytical tools

A.1 Combining 1D stellar images

A.1.1 Cross-correlation

The first stage in combining 1D stellar images is to cross correlate an image x_i ($i = 1, \dots, N$ where N is the number of datapoints within the image) with a reference image y_i . While the position of the reference image remains fixed the comparable image passes through a series of displacements. For each displacement τ the correlation coefficient r is calculated:

$$r(\tau) = \frac{\sum_i (x_i - \bar{x})(y_{i+\tau} - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2 \sum_i (y_i - \bar{y})^2}}, \quad (\text{A.1})$$

where \bar{x} and \bar{y} are the means of the corresponding image fluxes. r is normalised to lie in the range ± 1 . When the maximum in r is found, the value for τ can be extracted and used to reposition of the comparable image so that it is combined with the reference image.

A.1.2 Quadratic Interpolation

By fitting a quadratic function to the peak of the cross-correlation curve it is possible to increase the correlation between the two images. The point at which the first derivative of the function vanishes determines the additional shift that must be applied to the comparable image in order to increase the resolution of the sub-pixel positioning to $\sim 100^{\text{th}}$ pixel.

Figure A.1.2 shows an example of how to combine two 1D stellar images using cross-correlation and quadratic interpolation.

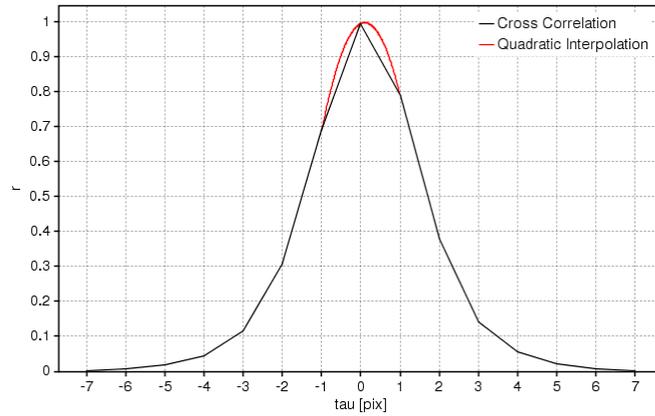


Figure A.1: A more exact measure of the maximum correlation can be found using quadratic interpolation.

A.2 Model fitting: extracting image parameters

A core part of analysing the radiation campaign data as well as testing the CDM prediction was image parameter extraction. In order to extract the centroid and flux of a stellar image simultaneously, which is optimal in the case of Gaia, there are several fitting procedures that have been used. Here these procedures will be discussed in more detail.

A.2.1 First-order Taylor expansions and Householder least squares

If O is the observation (e.g. the campaign test data) and P the prediction (in our case, a cubic spline function), we relate the two in terms of a scaling parameter, the intensity I , and a position x :

$$O = IP(x). \quad (\text{A.2})$$

Ensuring that the prediction is as close as possible to the observation, i.e. a best fit is realised between the data and the model, a correction is applied to I and x , namely dI and dx respectively:

$$O = (I + dI)P(x + dx). \quad (\text{A.3})$$

Estimating these corrections is performed via Householder least squares. However, before we implement this process we approximate P as a first-order Taylor expansion:

$$P(x + dx) \approx P(x) + \frac{\partial P(x)}{\partial x} dx, \quad (\text{A.4})$$

By substituting eqn. A.4 into eqn. A.3, multiplying out terms and rearranging we can identify the coefficients that need to be fed into a Householder least squares algorithm in order to find the associated corrections:

$$O - IP = PdI + I \frac{\partial P}{\partial x} dx. \quad (\text{A.5})$$

In the formalism of Householder least squares: $O - IP$ are the 'observations', P is the first coefficient and $I \frac{\partial P}{\partial x}$ is the second coefficient (see van Leeuwen (2007) Appendix C).

When estimating more than one unknown simultaneously, the resulting vector of the least squares estimates are given in part by evaluating the inverse of a matrix. However, inverting a matrix can lead to problems with accuracy and rounding. Introducing orthogonal matrices into the least squares problem makes the solution less susceptible to numerical and rounding errors. The Householder transformation is one such choice for an orthogonal matrix that aids this process. This transformation can permit the triangularization of matrices, which forms one part of the least squares problem. For further information on the properties of the Householder transformation and its use in least squares problems see Bierman (1997).

Extracting image parameters is an iterative process. The iteration terminates once the update to the positional correction has reached a given limit. Usually this limit was set to 10^{-3} pixels.

A.2.2 Maximum Likelihood 1

This method, described in Lindegren (SAG-LL-032), estimates both the centroid location ξ_0 and intensity N of a single stellar image on a constant background b by maximising the probability of measuring the observed counts n_j ($j = 1, 2, \dots, m$). Here is a description of this method very much in line with what is described in Lindegren (SAG-LL-032).

The predicted count or rather the expected number of electrons in the j th sample is:

$$E(n_j) = r^2 + b + NL(j - \xi_0), \quad (\text{A.6})$$

where r is the RMS readout noise and $L(\xi)$ the line-spread function (normalised to $\int L(\xi) d\xi = 1$). As the observed counts obey the Poisson distribution the probability of observing the mean value $E(n_j)$ is:

$$P(n_j; E(n_j)) = \frac{[r^2 + b + NL(j - \xi_0)]_j^n}{n_j!} \exp\{r^2 + b + NL(j - \xi_0)\} \quad (\text{A.7})$$

Strictly, $L(\xi)$ must be known to infinite precision, which is not a reasonable assumption in the case of Gaia. So to approximate $L(\xi)$, a calibrated line-spread function is used instead: $C(\xi)$. This makes the estimation slightly sub-optimal. The (quasi-) likelihood function to maximise is therefore:

$$l(\xi_0, N) = \prod_{j=1}^m \frac{[r^2 + b + NC(j - \xi_0)]_j^n}{n_j!} \exp\{r^2 + b + NL(j - \xi_0)\}. \quad (\text{A.8})$$

Using $\ln N$ prevents N from becoming negative during the iterations, so Eqn. A.8 is differentiated with respect to ξ_0 and $\ln N$ after taking the logarithm of the likelihood:

$$g(\xi_0, \ln N) \equiv \frac{\partial \ln l}{\partial \xi_0} = N \sum_j^m = 1 \left(1 - \frac{n_j}{s_j}\right) C'_j = 0 \quad (\text{A.9})$$

$$h(\xi_0, \ln N) \equiv \frac{\partial \ln l}{\partial \ln N} = -N \sum_j^m = 1 \left(1 - \frac{n_j}{s_j}\right) C_j = 0 \quad (\text{A.10})$$

where $C_j = C(j - \xi_0)$ and the first derivative $C'_j = (j - \xi_0)$, and s_j represents the fitted counts:

$$s_j = r^2 + b + NC_j \quad (\text{A.11})$$

Eqns. A.9 and A.10 may be solved by the Newton-Raphson method, but for this the second derivatives of the log-likelihood are required. These quantities may be gathered in the Fisher information matrix:

$$\mathbf{F} = \begin{bmatrix} -\frac{\partial^2 \ln l}{\partial \xi_0^2} & -\frac{\partial^2 \ln l}{\partial \xi_0 \partial \ln N} \\ -\frac{\partial^2 \ln l}{\partial \xi_0 \partial \ln N} & -\frac{\partial^2 \ln l}{\partial (\ln N)^2} \end{bmatrix} \quad (\text{A.12})$$

Therefore,

$$F_{11} = N^2 \sum_j^m = 1 \left[\frac{n_j}{s_j^2} C_j'^2 + \frac{1}{N} \left(1 - \frac{n_j}{s_j}\right) C_j'' \right] \quad (\text{A.13})$$

$$F_{12} = -N^2 \sum_j^m = 1 \frac{n_j}{s_j^2} C_j C_j' \quad (\text{A.14})$$

$$F_{22} = N^2 \sum_j^m = 1 \frac{n_j}{s_j^2} C_j^2. \quad (\text{A.15})$$

Improved values to ξ_0 and $\ln N$ are given by $\xi_0 + \Delta \xi_0$ and $\ln N + \Delta \ln N$. The updates are given by:

$$\begin{bmatrix} \Delta \xi_0 \\ \Delta \ln N \end{bmatrix} = \mathbf{F}^{-1} \begin{bmatrix} g(\xi_0, \ln N) \\ h(\xi_0, \ln N) \end{bmatrix} \quad (\text{A.16})$$

The iterations continue until the (absolute) size of the updates fall below a given limit.

Using the Newton-Raphson method requires reasonable initial estimates for the centroid and intensity. This is especially true for ξ_0 , as Lindegren (SAG-LL-032) states "If the initial ξ_0 is wrong by more than $\approx 0.5 \times \text{FWHM}$ of the line-spread function, then the method could easily converge to the completely wrong solution..... However, normally a good initial estimate will be available from even a primitive search of the maximum in n_j ".

For estimating the errors, this model uses the Poisson error property: $V(n_j) = E(n_j) \approx n_j$, and takes the errors of the individual n_j as independent. A very good estimate of the covariance of $[\xi_0, \ln N]$ is given by \mathbf{F}^{-1} such that:

$$\sigma_{\xi_0} = \left(\frac{F_{22}}{F_{11}F_{22} - F_{12}^2} \right)^{1/2} \quad (\text{A.17})$$

$$\sigma_N = N \left(\frac{F_{11}}{F_{11}F_{22} - F_{12}^2} \right)^{1/2}. \quad (\text{A.18})$$

A.2.3 Maximum Likelihood 2

This method, described in Lindegren (LL-078), generalises the ML method described above to handle not only line-spread functions but also point-spread functions and estimate CDM model parameters. Here is a description of this method very much in line with what is described in Lindegren (LL-078).

For a model of the sample data N_k , the model is correct when:

$$E(N_k) = \lambda_k(\boldsymbol{\theta}), \quad (\text{A.19})$$

where λ_k are the predicted sample values and $\boldsymbol{\theta}$ are a set of model parameters. The variance of N_k is given as:

$$V(N_k) = \lambda_k(\boldsymbol{\theta}) + r^2, \quad (\text{A.20})$$

where r is the standard deviation on the readout noise; the readout noise is modelled as a continuous gaussian process with zero mean and a standard deviation r . The probability density function for N should be a convolution between the Poisson and Gaussian distributions, however, this function is not suited to ML estimation. An approximation is made to simplify the estimation process. Trading off a (compound) gaussian model against a (modified) poissonian model shows that the latter gives simpler likelihood equations. The following probability function is thus used in the ML fitting:

$$p(N; \lambda, r) = \text{const} \times \frac{(\lambda + r^2)^{N+r^2}}{\Gamma(N + r^2 + 1)} e^{-\lambda - r^2}, \quad (\text{A.21})$$

where *const* is used to normalise the integral of $p(N; \lambda, r)$, assuming it is independent of λ .

Before initiating the ML process N_k has been corrected for bias (i.e. electronic offset) and gain, but still includes dark current and the background. r is assumed to be known. In the context of Gaia, the predicted sample values $\lambda_k(\boldsymbol{\theta})$ are constructed from the various source, attitude and calibration models, the set of parameters linked to the the type of application being performed. Here, the application is left undefined.

In practice, it is easier to maximise the logarithm of L . So using Eqn. A.21 the expression to maximise is:

$$\ln L(\boldsymbol{\theta}|\{N_k\}) = l(\boldsymbol{\theta}|\{N_k\}) = \text{const} + \sum_k \left[(N_k + r^2) \ln(\lambda_k(\boldsymbol{\theta}) + r^2) - \lambda_k(\boldsymbol{\theta}) \right] \quad (\text{A.22})$$

To find the set of parameters that makes the probability of obtaining N_i as large as possible, the position of the maximum may be found by solving:

$$\frac{\partial l(\boldsymbol{\theta}|\{N_k\})}{\partial \boldsymbol{\theta}} = \sum_k \frac{N_k - \lambda_k(\boldsymbol{\theta})}{\lambda_k(\boldsymbol{\theta}) + r^2} \frac{\partial \lambda_k}{\partial \boldsymbol{\theta}} = \mathbf{0}. \quad (\text{A.23})$$

As Eqn. A.23 is non-linear in the parameters $\boldsymbol{\theta}$ the equations must be solved by iteration. In the m^{th} iteration:

$$\boldsymbol{\delta}^{(m)} \equiv \left(\frac{\partial l(\boldsymbol{\theta} | \{N_k\})}{\partial \boldsymbol{\theta}} \right)_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(m)}} \neq \mathbf{0}. \quad (\text{A.24})$$

Here an update $\Delta\boldsymbol{\theta}^{(m)}$ is required that will reduce $\boldsymbol{\delta}^{(m)}$ to zero. Writing $\boldsymbol{\delta}^{(m+1)}$ as a Taylor expansion and setting the next $\boldsymbol{\delta}$ term in the iteration to zero:

$$\boldsymbol{\delta}^{(m+1)} \simeq \boldsymbol{\delta}^{(m)} + \left(\frac{\partial^2 l(\boldsymbol{\theta} | \{N_k\})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \right)_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(m)}} \Delta\boldsymbol{\theta}^{(m)} \quad (\text{A.25})$$

identifies the update that must be applied after each iteration:

$$\Delta\boldsymbol{\theta}^{(m)} = - \left(\frac{\partial^2 l(\boldsymbol{\theta} | \{N_k\})}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \right)_{\boldsymbol{\theta}=\boldsymbol{\theta}^{(m)}}^{-1} \quad (\text{A.26})$$

The Hessian $\mathbf{H} = \frac{\partial^2 l}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'}$ is a symmetric $n \times n$ matrix and $\Delta\boldsymbol{\theta}$ and $\boldsymbol{\delta}$ are n -dimensional vectors. As Lindgren (LL-078) states: "If the parameter estimation is a well-posed problem, then $-\mathbf{H}$ is positive definite and there is a unique ML solution".

Assuming the iteration is close to the true solution so that $E(N_k) \simeq \lambda_k$ the expectation of the Hessian $E(H_{ij})$ given in terms of its elements ij is:

$$E(H_{ij}) \simeq - \sum_k \frac{1}{\lambda_k + r^2} \frac{\partial \lambda_k}{\partial \theta_i} \frac{\partial \lambda_k}{\partial \theta_j}. \quad (\text{A.27})$$

This does not involve any second derivatives in λ_k , which is often favourable for the stability of the iteration. If \mathbf{A} , a symmetric, positive definite matrix is equated to:

$$A_{ij} = \sum_k \frac{1}{\lambda_k(\boldsymbol{\theta}) + r^2} \frac{\partial \lambda_k}{\partial \theta_i} \frac{\partial \lambda_k}{\partial \theta_j}, \quad (\text{A.28})$$

and

$$\delta_i = \sum_k \frac{1}{\lambda_k(\boldsymbol{\theta}) + r^2} \frac{\partial \lambda_k}{\partial \theta_i}, \quad (\text{A.29})$$

the linear set of equations to be solved in each iteration is:

$$\mathbf{A}^{(m)} \Delta\boldsymbol{\theta}^{(m)} = \boldsymbol{\delta}^{(m)} \quad (\text{A.30})$$

As \mathbf{A} is symmetric and positive definite LU factorization via Cholesky's method may be a way to solve Eq. A.30 for $\Delta\boldsymbol{\theta}^{(m)}$.

To calculate the covariance the Cramér-Rao limit is used as given by the matrix \mathbf{A} after convergence:

$$\text{Cov}(\hat{\boldsymbol{\theta}}) \simeq \mathbf{A}^{-1} \quad (\text{A.31})$$

B

Further results from the Radiation Campaigns

B.1 Radiation-induced CTI image parameters from RC1

Total flux [e-]	G mag	DOB	Bias		Bias error		Charge loss		Charge loss error	
			AL pixel	mas	AL pixel	mas	e ⁻	%	e ⁻	%
60000	15.33	0	0.209	12.23	0.006	0.36	17569	28.91	425	0.70
		5	0.181	10.72	0.006	0.35	15695	21.39	502	0.69
		10	0.176	10.38	0.005	0.30	13741	20.78	389	0.60
7000	17.66	0	0.149	8.77	0.004	0.25	1972	28.98	34	0.50
		5	0.088	5.18	0.006	0.37	1147	15.61	57	0.78
		10	0.040	2.38	0.005	0.31	921	13.99	47	0.71
2000	19.02	0	0.080	4.72	0.005	0.28	689	30.88	13	0.57
		5	0.065	3.85	0.006	0.35	307	12.67	21	0.86
		10	0.073	4.28	0.007	0.39	169	8.75	19	1.00
650	20.24	0	0.204	12.05	0.008	0.49	256	35.13	7	0.97
		5	0.039	2.27	0.011	0.64	97	14.73	10	1.57
		10	0.112	6.63	0.014	0.82	44	7.41	13	2.18
400	20.77	0	0.250	14.71	0.011	0.66	200	38.68	6	1.23
		5	0.121	7.15	0.014	0.82	36	13.41	11	2.11
		10	-0.126	-7.44	0.018	1.08	40	8.58	14	2.77
200	21.52	0	0.168	9.90	0.019	1.09	79	40.05	5	2.37
		5	0.147	8.67	0.022	1.31	36	17.39	7	3.45
		10	-0.231	-13.63	0.031	1.80	1	0.52	10	5.51

Table B.1: Average radiation-induced CTI image parameters as a function of magnitude and DOB level.

B.2 Results from the analysis of RC2 at $G = 15$

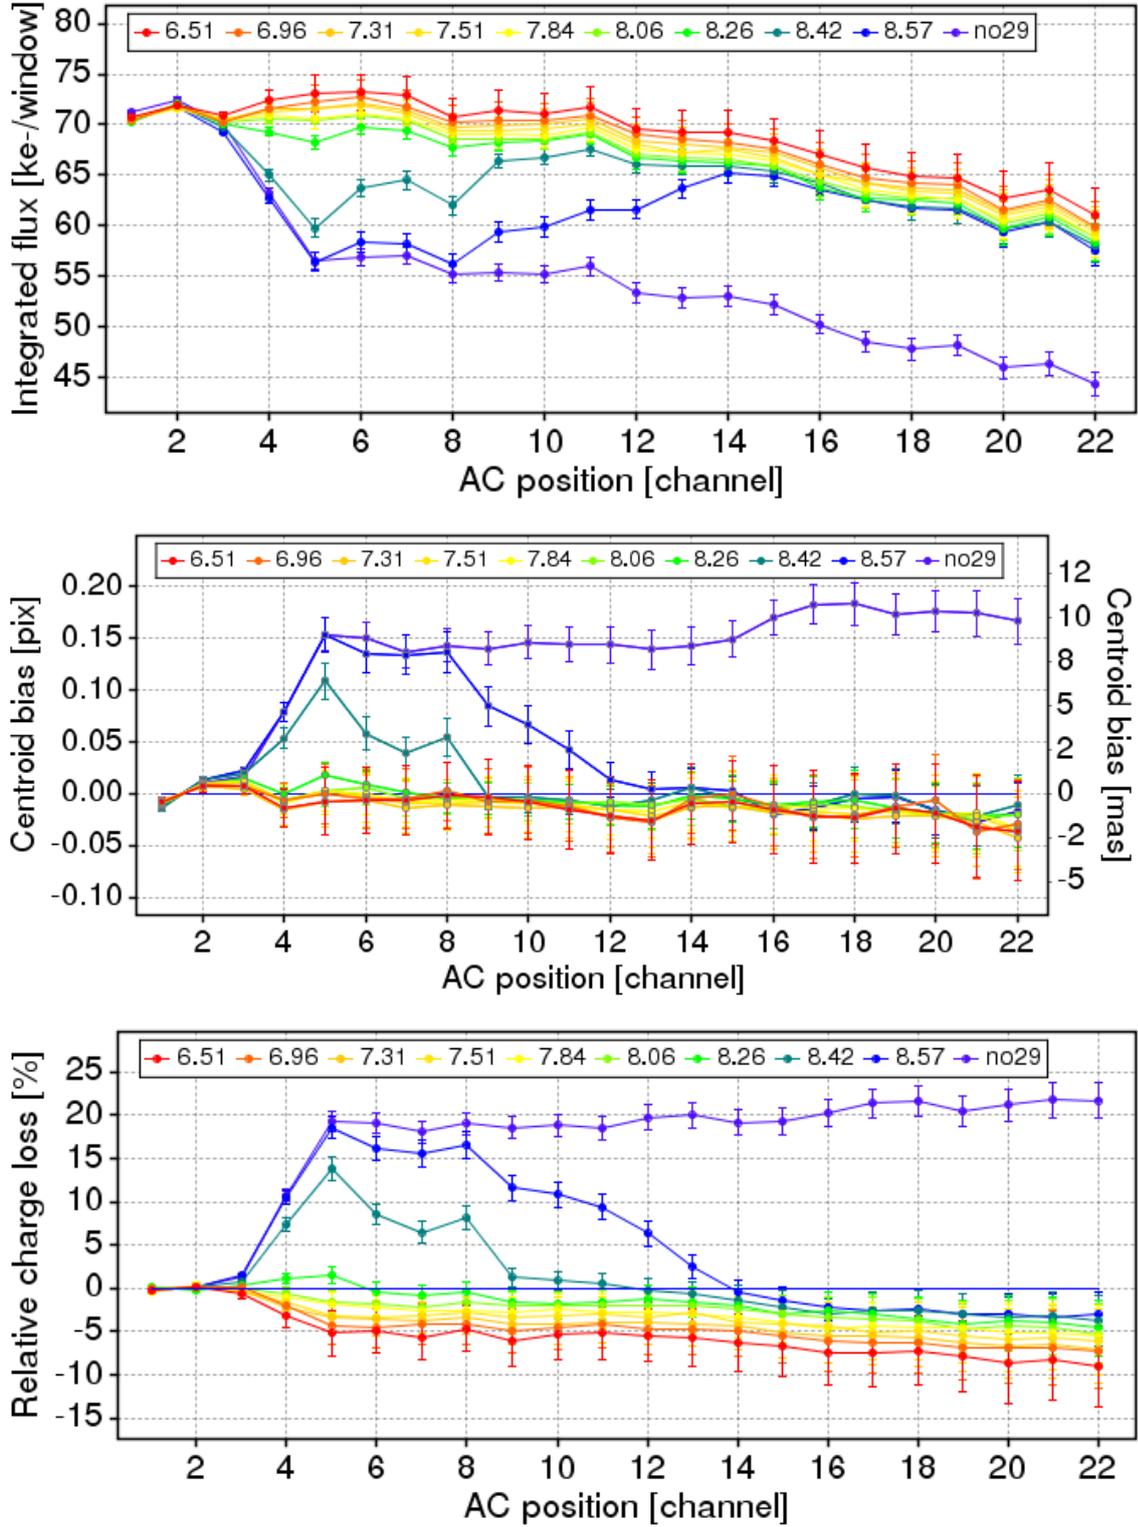


Figure B.1: (top) Integrated flux, (middle) relative centroid bias, and (bottom) relative charge loss as a function of AC position and CI level in column 1 at $G = 15$.

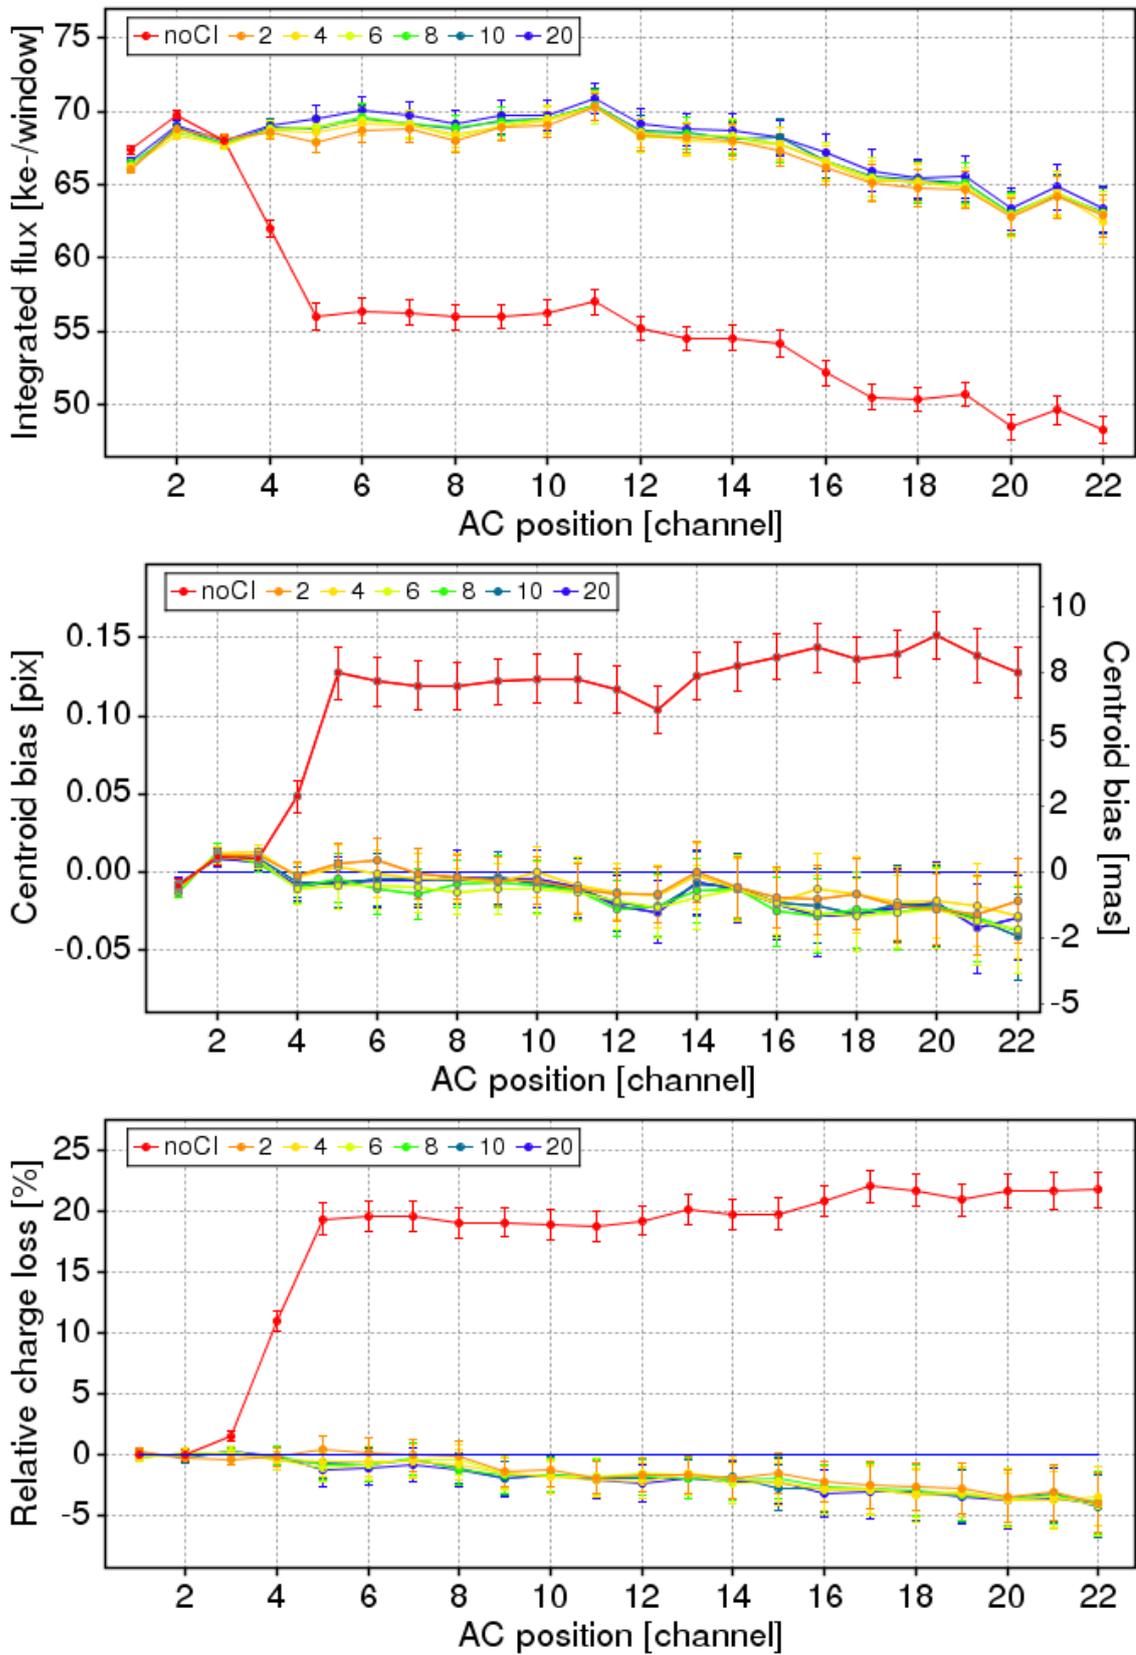


Figure B.2: (top) Integrated flux, (middle) relative centroid bias, and (bottom) relative charge loss as a function of AC position and CI duration in column 1 at $G = 15$.

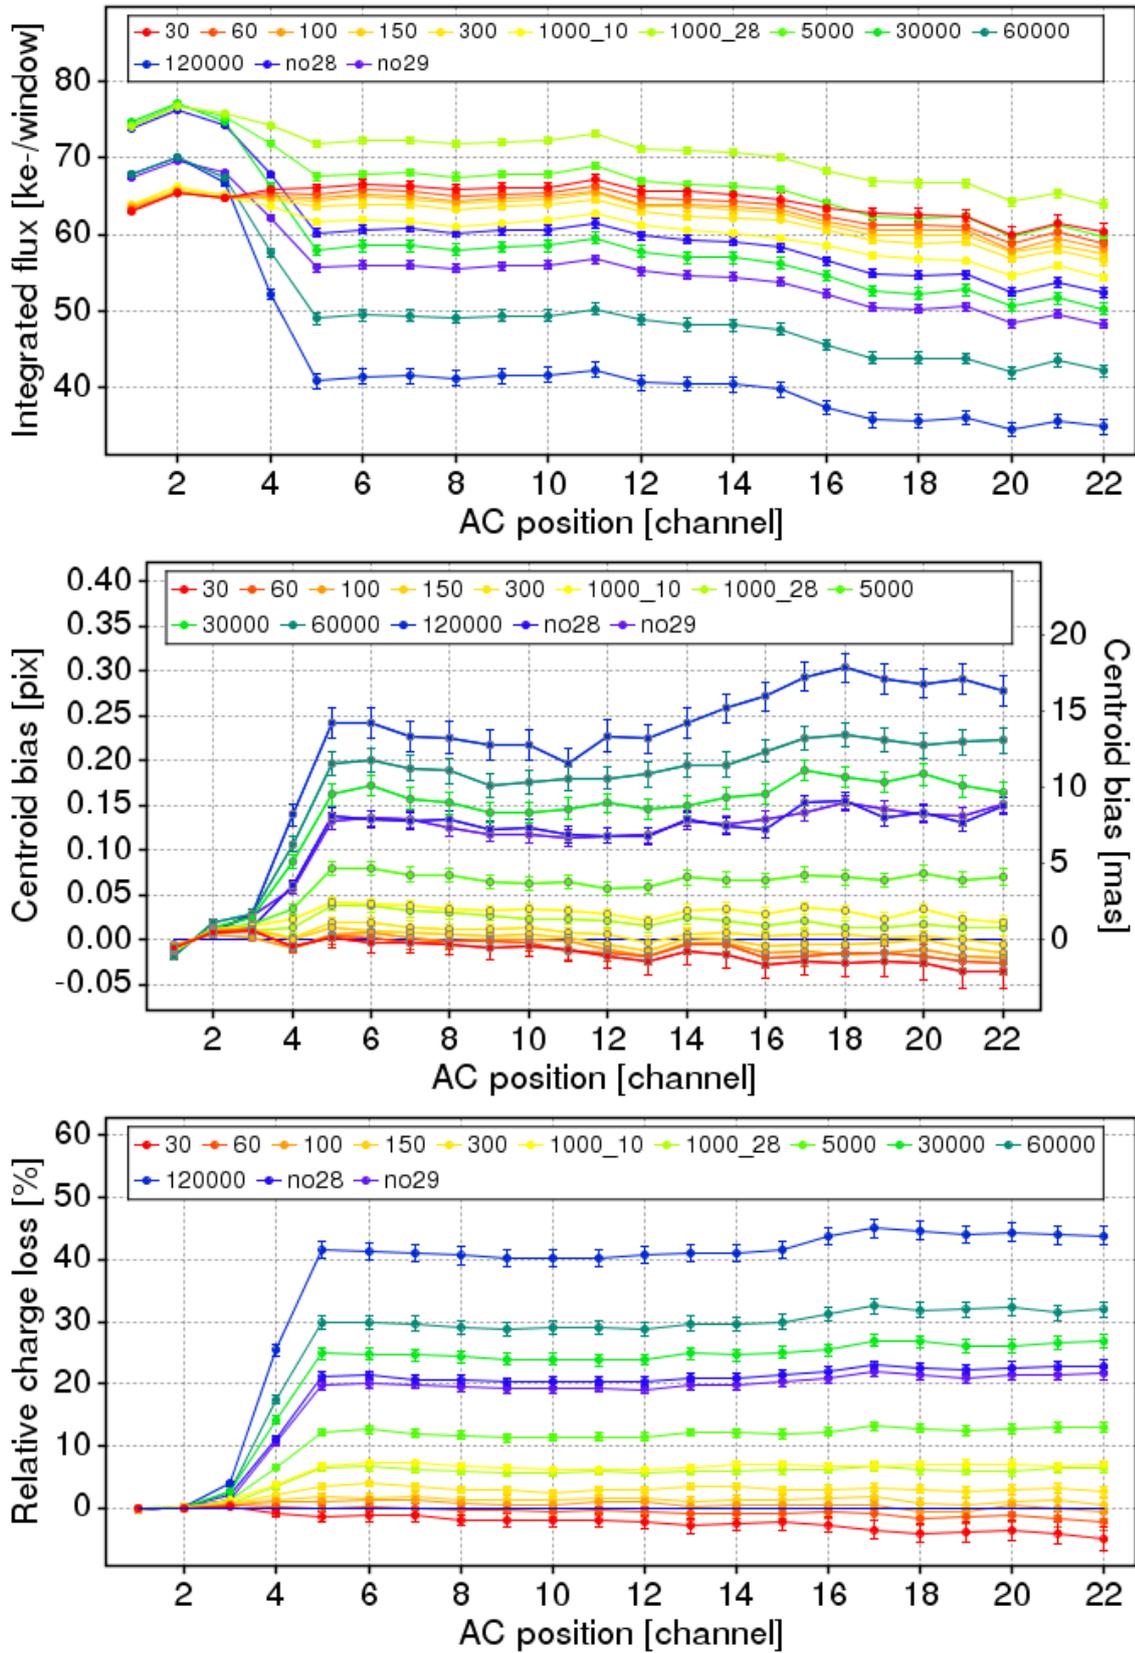


Figure B.3: (top) Integrated flux, (middle) relative centroid bias, and (bottom) relative charge loss as a function of AC position and CI delay in column 1 at $G = 15$.

B.3 Radiation-induced CTI image parameters from RC2

G = 13.3												
[V]	6.51	6.96	7.31	7.51	7.84	8.06	8.26	8.42	8.57	8.65	8.72	0
B [pix]	-0.012	-0.010	-0.008	-0.008	-0.008	-0.008	-0.003	0.019	0.055	0.074	0.089	0.115
bE [pix]	0.017	0.017	0.017	0.018	0.018	0.018	0.018	0.019	0.018	0.018	0.017	0.017
C [%]	-0.66	-0.66	-0.51	-0.34	-0.09	0.26	1.04	2.97	5.56	6.68	7.61	10.46
cE [%]	1.20	1.22	1.20	1.16	1.29	1.29	1.32	1.35	1.32	1.28	1.26	1.24

G = 15												
B [pix]	-0.014	-0.012	-0.017	-0.014	-0.012	-0.009	-0.007	0.006	0.047	-	-	0.150
bE [pix]	0.037	0.033	0.030	0.027	0.024	0.021	0.019	0.017	0.018	-	-	0.017
C [%]	-5.89	-4.90	-4.29	-3.54	-3.09	-2.29	-1.63	1.34	6.89	-	-	19.32
cE [%]	3.13	2.78	2.51	2.30	2.04	1.79	1.57	1.44	1.53	-	-	1.39

B = Bias, bE = Bias error, C = Charge loss, cE = Charge loss error.

Table B.2: CTI image parameters measured in the first test.

		Inj. duration [pix]	0	2	4	6	8	10	20
G = 13.3	Bias [pix]		0.098	-0.005	-0.008	-0.010	-0.010	-0.009	-0.009
	Bias Error [pix]		0.008	0.008	0.008	0.008	0.008	0.008	0.008
	Charge loss [%]		10.07	0.04	-0.37	-0.58	-0.46	-0.47	-0.57
	Charge loss error [%]		0.63	0.62	0.62	0.62	0.63	0.64	0.66
G = 15	Bias [pix]		0.124	-0.008	-0.008	-0.019	-0.018	-0.011	-0.012
	Bias Error [pix]		0.015	0.017	0.018	0.019	0.019	0.019	0.020
	Charge loss [%]		19.73	-1.28	-1.69	-1.64	-1.68	-1.78	-1.95
	Charge loss error [%]		1.25	1.44	1.52	1.56	1.56	1.59	1.63

Table B.3: CTI image parameters measured in the second test.

G = 13.3											
[pix]	30	60	100	150	300	1000	5k	30k	60k	120k	No inj.
B [pix]	-0.008	-0.005	-0.003	0	0.006	0.020	0.056	0.140	-	-	0.094
bE [pix]	0.008	0.007	0.006	0.006	0.006	0.006	0.006	0.009	-	-	0.004
C [%]	-0.12	0.36	0.58	0.94	1.66	3.04	6.54	14.20	-	-	10.31
cE [%]	0.63	0.56	0.49	0.47	0.45	0.45	0.48	0.62	-	-	0.33

G = 15											
B [pix]	-0.012	-0.007	-0.003	0.002	0.010	0.029	0.067	0.156	0.192	0.238	0.127
bE [pix]	0.013	0.010	0.009	0.008	0.007	0.007	0.008	0.012	0.013	0.017	0.009
C [%]	-2.18	-0.50	0.73	1.49	3.16	6.38	11.97	24.72	29.79	41.41	20.46
cE [%]	1.13	0.84	0.74	0.73	0.56	0.54	0.65	0.97	1.12	1.39	0.77

G = 16.75											
B [pix]	-0.018	-0.015	-0.010	-0.002	0	0.010	-	0.050	-	-	0.053
bE [pix]	0.011	0.007	0.005	0.005	0.003	0.002	-	0.002	-	-	0.002
C [%]	-13.18	-7.02	-3.82	-1.43	1.64	7.20	-	20.44	-	-	18.43
cE [%]	4.40	2.98	2.25	1.91	1.37	0.96	-	0.88	-	-	0.94

G = 18.5											
B [pix]	-0.023	-0.032	-0.022	-0.033	-0.016	0.020	-	-	-	-	0.009
bE [pix]	0.099	0.066	0.050	0.043	0.032	0.020	-	-	-	-	0.017
C [%]	-38.41	-23.00	-14.168	-9.67	-2.54	7.38	-	-	-	-	23.69
cE [%]	8.77	5.59	4.17	3.43	2.48	1.49	-	-	-	-	1.12

G = 20											
B [pix]	0.046	-0.056	-0.055	-0.041	-0.147	-0.143	-	-0.940	-	-	-0.038
bE [pix]	0.221	0.160	0.126	0.108	0.081	0.044	-	0.025	-	-	0.039
C [%]	-148.08	-92.86	-61.93	-47.74	-24.16	-0.34	-	16.56	-	-	31.83
cE [%]	28.22	18.03	12.65	10.44	7.23	3.52	-	3.06	-	-	2.17

B = Bias, bE = Bias error, C = Charge loss, cE = Charge loss error.

Table B.4: CTI image parameters measured in the third test.

B.4 Radiation-induced CTI image parameters from RC3

Nom. DOB	Delay [pix]	G = 13.6		G = 15		G = 20	
		B [pix]	C [%]	B [pix]	C [%]	B [pix]	C [%]
2	100	$\sim 0 \pm 0.002$	1.15 ± 0.17	-	-	0.030 ± 0.012	8.95 ± 0.80
	500	0.008 ± 0.002	1.89 ± 0.19	-	-	0.054 ± 0.009	7.52 ± 0.66
	1000	0.012 ± 0.003	2.32 ± 0.23	0.021 ± 0.002	2.17 ± 0.16	see Table B.6	see Table B.6
	5000	-	-	0.034 ± 0.002	4.17 ± 0.19	0.078 ± 0.008	5.12 ± 0.56
	10000	-	-	0.056 ± 0.002	5.57 ± 0.20	0.100 ± 0.008	8.78 ± 0.54
	No inj.	0.059 ± 0.005	6.58 ± 0.41	0.089 ± 0.003	8.34 ± 0.27	see Table B.6	see Table B.6
4	100	0.004 ± 0.003	1.33 ± 0.21	-	-	0.021 ± 0.012	9.63 ± 0.85
	500	0.007 ± 0.002	1.58 ± 0.21	-	-	0.024 ± 0.010	8.16 ± 0.74
	1000	0.012 ± 0.003	1.86 ± 0.22	0.026 ± 0.002	2.31 ± 0.16	see Table B.6	see Table B.6
	5000	-	-	0.039 ± 0.002	4.30 ± 0.19	0.071 ± 0.009	4.88 ± 0.63
	10000	-	-	0.045 ± 0.002	5.38 ± 0.22	0.103 ± 0.016	4.97 ± 0.62
	No inj.	0.058 ± 0.005	6.37 ± 0.38	0.086 ± 0.003	8.34 ± 0.25	see Table B.6	see Table B.6
7	100	0.003 ± 0.002	1.27 ± 0.15	-	-	0.057 ± 0.016	10.13 ± 1.02
	500	0.004 ± 0.002	1.83 ± 0.18	-	-	0.012 ± 0.012	5.92 ± 0.85
	1000	0.004 ± 0.003	2.11 ± 0.19	-	-	0.030 ± 0.011	3.30 ± 0.80
	5000	-	-	0.051 ± 0.002	4.17 ± 0.18	0.110 ± 0.011	2.13 ± 0.80
	10000	-	-	0.059 ± 0.002	5.43 ± 0.18	0.057 ± 0.011	3.75 ± 0.77
	No inj.	0.054 ± 0.009	8.02 ± 0.64	0.086 ± 0.003	8.03 ± 0.23	0.039 ± 0.011	4.73 ± 0.75

Table B.5: CTI effects measured in the MF star tests. B = Bias, C = Charge loss.

There are certain instances during the MF star tests when some test cases were repeated on different days. We have analysed two cases when the charge injection CI delay is 1000 and when no CI has been made. Table B.6 shows the biases and charge losses measured in both cases. These values are shown plotted in Figure 2.113.

G = 20			CI delay [pix]		
Nominal DOB	Observed DOB	Test date		1000	No inj.
2	~ 5.6	02/07/08	Bias [pix]	0.034 ± 0.008	0.071 ± 0.007
		16/07/08		0.062 ± 0.009	0.106 ± 0.008
		02/07/08	Charge loss [%]	5.44 ± 0.56	9.65 ± 0.50
		16/07/08		6.59 ± 0.61	7.42 ± 0.56
4	~ 10.6	03/07/08	Bias [pix]	0.019 ± 0.008	0.055 ± 0.008
		17/07/08		0.079 ± 0.009	0.038 ± 0.009
		03/07/08	Charge loss [%]	4.20 ± 0.59	6.22 ± 0.58
		17/07/08		3.26 ± 0.68	7.83 ± 0.56

Table B.6: Multiple instances of the same test case at G = 20 with a varying DOB level.

	Inj. delay [pix]	100	500	No inj.
G = 20	Bias [pix]	0.058 ± 0.022	0.126 ± 0.012	0.148 ± 0.008
	Charge loss [%]	16.22 ± 1.37	11.65 ± 0.78	-

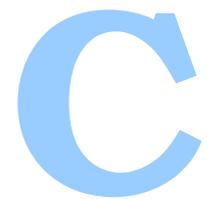
Table B.7: CTI effects measured at G = 20 and the lowest possible DOB.

A brief analysis was made of the G = 20 test cases when the DOB level was the lowest possible. The results are summarised in Table B.7. As expected, they are worse than the results seen with a DOB of 2

Nom. DOB	Delay [pix]	G = 11.1		G = 11.9		G = 12.6	
		B [pix]	C [%]	B [pix]	C [%]	B [pix]	C [%]
0	100	0.004 ± 0.002	-0.30 ± 0.19	-0.001 ± 0.002	-0.10 ± 0.18	0.002 ± 0.002	0.19 ± 0.14
	500	0.006 ± 0.002	-0.37 ± 0.19	0.001 ± 0.002	-0.04 ± 0.18	0.004 ± 0.002	0.30 ± 0.14
	1000	0.005 ± 0.002	-0.30 ± 0.19	0.001 ± 0.002	0.08 ± 0.18	0.002 ± 0.002	0.45 ± 0.15
	5000	0.009 ± 0.002	-0.14 ± 0.19	0.005 ± 0.002	0.23 ± 0.18	0.008 ± 0.002	0.92 ± 0.15
	10000	0.006 ± 0.002	-0.20 ± 0.19	0.007 ± 0.002	0.29 ± 0.18	0.010 ± 0.002	1.17 ± 0.16
2	100	0.005 ± 0.002	-0.34 ± 0.19	-0.001 ± 0.002	0.04 ± 0.20	0.001 ± 0.002	0.19 ± 0.14
	500	0.003 ± 0.002	-0.31 ± 0.19	0.003 ± 0.002	-0.20 ± 0.19	0.002 ± 0.002	0.29 ± 0.14
	1000	0.004 ± 0.002	-0.31 ± 0.19	0.001 ± 0.002	0.08 ± 0.20	-0.001 ± 0.002	0.44 ± 0.14
	5000	0.010 ± 0.003	-0.34 ± 0.20	0.001 ± 0.002	0.34 ± 0.20	0.008 ± 0.002	0.79 ± 0.15
	10000	0.003 ± 0.002	-0.31 ± 0.18	-0.005 ± 0.002	0.32 ± 0.20	0.011 ± 0.002	1.14 ± 0.16
4	100	~0 ± 0.002	-0.57 ± 0.20	~0 ± 0.002	-0.16 ± 0.20	~0 ± 0.002	0.04 ± 0.20
	500	0.006 ± 0.002	-0.66 ± 0.21	~0 ± 0.003	-0.17 ± 0.22	0.001 ± 0.002	-0.20 ± 0.19
	1000	0.001 ± 0.002	-0.43 ± 0.20	0.004 ± 0.002	-0.06 ± 0.20	0.003 ± 0.002	0.08 ± 0.20
	5000	0.010 ± 0.002	-0.51 ± 0.20	-	-	-0.001 ± 0.002	0.34 ± 0.20
	10000	0.007 ± 0.002	-0.52 ± 0.20	-	-	0.003 ± 0.002	0.32 ± 0.20

Table B.8: CTI effects measured in the bright star tests. B = Bias, C = Charge loss.

and appear to follow the same trend: the bias gets progressively worse as the CI delay increases, while the charge loss initially decreases.



Trap Physics and Charge Distortion Models

C.1 Shockley Read Hall Theory

The presence of defects in a semiconductor such as silicon, brought about by impurities, imperfections or radiation damage can produce discrete energy levels in the band gap between the valence and conduction band. The interaction between the defect and the valence or conduction band can be viewed as four basic processes, where an "occupied" defect here is understood to have an electron:

1. An unoccupied defect can capture an electron from the conduction band.
2. An occupied defect can emit its electron into the conduction band.
3. An unoccupied defect can emit a hole into the the valence band.
4. An occupied defect can capture a hole from the valence band.

If the defect is a trap either the first two or the last two steps above are realised in combination. For Gaia, it is clearly the first two that are of concern when radiation damage creates trapping centres in the silicon. Shockley & Read (1952) and Hall (1952) formulated a theory of recombination that could describe the capture and emission rates of these defects.

According to their statistical approach, the capture rate for electrons c may be expressed as:

$$c = \sigma v_{th} n, \tag{C.1}$$

where σ is the trap capture cross-section, v_{th} is the electron thermal velocity and n is the electron density in the conduction band (for Gaia, this would be the electron density in the vicinity of the trap). Identifying

the occupation probability with the Fermi-Dirac distribution the emission rate e may be expressed as:

$$e = \sigma v_{th} N_c \exp \left\{ -\frac{E_c - E_t}{k_B T} \right\}, \quad (\text{C.2})$$

where N_c is the effective density of states in the conduction band, E_c is the energy of the conduction band edge, E_t is the trap energy level, k_B is Boltzmann's constant and T is the temperature in Kelvin.

In practice, when simulating the Gaia data stream, both the trap capture cross-section and the trap release time-constant τ_r are found by fitting to experimental data. τ_r is the inverse of e :

$$\tau_r = 1/e \quad (\text{C.3})$$

C.2 Capture and emission probabilities

CTI is dependent upon the probability an electron is captured P_c or released P_e from a trap, which is usually modelled as a decaying exponential:

$$P_i = 1 - \exp \left(\frac{-t}{\tau_i} \right), \quad (\text{C.4})$$

where $i = c$ or e , t is a characteristic transfer time and τ is the trap capture or emission time. There are different types of defects (e.g. the P-V centre) and each one is described by a set of trap parameters. These parameters are implicit to the trap capture or emission times. Identifying which ones may have a dominant effect on Gaia and determining their τ_i has become a significant part of the effort in characterising the effects of radiation damage on the focal plane.

C.3 CDM01

The first Charge Distortion Model (CDM01) developed to describe the distortion of an image caused by radiation damage was put together by Lennart Lindegren at Lund University. CDM01 was meant to act as a starting point from which a more accurate model could be derived. Here is a description of this model very much in line with what is described in Lindegren (LL-075).

In Figure 1.13 the input to the CDM is the *sampled image* represented by a sequence of values s_i in successive TDI periods (or pixels) i . The output, the *charge image*, is a sequence of distorted counts, d_i , which are generated as follows:

$$d_i = s_i - c_i + r_i, \quad (\text{C.5})$$

where c_i and r_i are the number of electrons captured and released in TDI step i respectively. To specify the illumination history, i.e. how c_i and r_i depend on the current and previous pixel values, it is assumed that c_i and r_i are functions of a small number of state variables $f_i^{(0)}, f_i^{(1)}, f_i^{(2)}, \dots$ that are computed recursively. With $\mathbf{f} \equiv (\mathbf{f}^{(0)}, \mathbf{f}^{(1)}, \mathbf{f}^{(2)}, \dots)$:

$$c_i = c(s_i, \mathbf{f}_i) \quad (\text{C.6})$$

$$r_i = r(s_i, \mathbf{f}_i) \quad (\text{C.7})$$

$$\mathbf{f}_i = \mathbf{f}(s_i, \mathbf{f}_{i-1}) \quad (\text{C.8})$$

A single state variable called the *equivalent fill level*, denoted by f_i , is currently used, which could be understood as the minimum s_i required for part of the charges to be captured. If $s_i < f_i$ hardly any traps would be seen.

To compute $r(s_i, f_i)$ we introduce the charge release model:

Let $N_{filled}(s_i, f_i)$ be the number of filled traps 'seen' by the current charge package. Assuming a single trap species with release time constant τ , the release function becomes:

$$r(s_i, f_i) = (1 - e^{-\Delta t/\tau}) N_{filled}(s_i, f_i) = \eta N_{filled}(s_i, f_i), \quad (\text{C.9})$$

where Δt is the TDI period and $\eta = 1 - \exp(-\Delta t/\tau)$ for brevity. For $\tau \gg \Delta t$, $\eta \approx \Delta t/\tau$.

For lack of more detailed information, it is assumed that N_{filled} is independent of s_i and proportional to f_i ; thus

$$N_{filled}(s_i, f_i) = R f_i, \quad (\text{C.10})$$

where $R \geq 0$ is a model parameter characterising the degree of radiation damage, or equivalently the total number of traps. The release function is then:

$$r(s_i, f_i) = \eta R f_i \quad (\text{C.11})$$

To compute $c(s_i, f_i)$ we introduce the charge capture model:

The capture function $c(s_i, f_i)$ should encapsulate the following behaviour: Let N_{empty} be the number of empty traps seen by the charge package, and N_c the number of charge available to them.

- In the limiting case when $N_{empty} \ll N_c$, the number of captures is given by N_{empty} .
- Conversely, when $N_{empty} \gg N_c$, the number of captures is given by N_c .
- When both N_{empty} and N_c are small, the number of captures is proportional to the product $N_{empty} N_c$.

One function that described this behaviour is:

$$c(s_i, f_i) = \frac{N_{empty} N_c}{N_{empty} + N_c + K} \quad (\text{C.12})$$

where $K \geq 0$ is a parameter representing the inverse capture rate in the limiting case of small N_{empty} and N_c .

Let F be the equivalent fill level corresponding to saturation of the traps (i.e. the value of f_i reached after passing a long series of big charge packages across the CCD). Then RF is the total number of traps and consequently

$$N_{empty} = RF - N_{filled} = R(F - f_i) \quad (C.13)$$

In order to estimate the number of charges available to fill the traps it is necessary to consider the different behaviours of the capture process at different signal sizes. This results in different capture functions c_A , c_B , etc. A weighted mean of these functions, using s_i to determine their weights, then gives the effective capture function c .

For small packages (subscript A), all the charges are available ($N_c = s_i$), but the probability of capture per charge is small because of the low charge density. Using Eqn. C.12 and Eqn. C.13 this gives

$$c_A(s_i, f_i) = \frac{R(F - f_i)s_i}{R(F - f_i) + s_i + K_A}, \quad (C.14)$$

where the parameter K_A may be adjusted to reflect the capture rate in this situation.

For large packages the effective fill level f_i acts as a threshold:

$$N_c = \begin{cases} s_i - f_i & \text{if } s_i > f_i \\ 0 & \text{otherwise} \end{cases} \quad (C.15)$$

so that the capture function becomes

$$c_B(s_i, f_i) = \begin{cases} \frac{R(F-f_i)(s_i-f_i)}{R(F-f_i)+(s_i-f_i)+K_B} & \text{if } s_i > f_i \\ 0 & \text{otherwise} \end{cases} \quad (C.16)$$

A smooth transition between the two cases is allowed by writing:

$$c(s_i, f_i) = q(s_i)c_A(s_i, f_i) + (1 - q(s_i))c_B(s_i, f_i) \quad (C.17)$$

where the transition function $q(s)$ goes from 1 for small s to 0 for large s . One model for $q(s)$ could be an exponential transition function:

$$q(s_i) = \exp(-s_i/S) \quad (C.18)$$

where the parameter S is the size of the charge package where the transition occurs. Although other more suitable models could perhaps be used.

C.3.1 Propagation of the state variable

In TDI step i , the net number of charges captured is given by $s_i - d_i = c_i - r_i$. Conservation of charges then requires that

$$N_{filled}(s_{i+1}, f_{i+1}) = N_{filled}(s_i, f_i) + c_i - r_i \quad (\text{C.19})$$

Using Eqn. C.10 and Eqn. C.11 we can then find f_{i+1} :

$$f_{i+1} = (1 - \eta) f_i + c_i/R \quad (\text{C.20})$$

C.3.2 Initialization of the state variable

In principle, f_i depends on the infinitely long prehistory of values s_j ($j \leq i$). In practice, it can be initialized in two different ways:

1. Immediately after a sufficiently large charge injection, it can be assumed that $f_i = F$.
2. After a sufficiently long, constant low-level illumination $s_j = \beta$ (background), the trap filling reaches a constant level $f_i = f_{SS}$ obtained by solving the steady-state equation

$$c(\beta, f_{SS}) = r(\beta, f_{SS}) \quad (\text{C.21})$$

If $\beta \ll s$ and $f_{SS} \ll F$, an approximate solution is

$$f_{SS} = \frac{\beta}{(R + K_A/F)\eta} \quad (\text{C.22})$$

C.4 CDM02

The second Charge Distortion Model (CDM02) developed to describe the distortion of an image caused by radiation damage was put together by Alex Short at ESTEC. As described in Short (AS-015) "In principle, the CDM may be empirical and need not be based upon detailed physical modelling of electron trapping and release. However, any model based upon physical considerations is certain to be superior in terms of application over a broad parameter space using the fewest possible variables". This is why CDM02 is based upon physical (SRH) theory, and developed in such a way that it may be applied in the Gaia data processing pipeline. Here is a description of this model very much in line with what is described in Short (AS-015).

To avoid time-consuming operations that calculate the trapping and release of charge in each individual pixel for each TDI period, this model evaluates the trapping process over the entire CCD column in a single calculation.

It is useful to keep the model flexible so to include the choice as to whether the trapping process is volume- or density driven (or somewhere in between) the confinement volume V_c is defined in terms of the number of electrons N_e , the Full Well Capacity FWC , the geometrical confinement model V_g (i.e. the volume of silicon when $N_e \approx FWC$), and a fitting parameter β :

$$\frac{V_c}{V_g} = \left(\frac{N_e}{FWC} \right)^\beta. \quad (C.23)$$

So along a column, the fraction of silicon "seen" by a given input sample (S) is given by integrating Eqn. C.23 from $N_e = 0$ to $N_e = S$:

$$F_{silicon} = \frac{1}{S} \int_0^S \left(\frac{N_e}{FWC} \right)^\beta dN_e = \frac{1}{1+\beta} \left(\frac{S}{FWC} \right)^\beta \quad (C.24)$$

S here is understood to be the signal that would be read out of the CCD if no traps were present, i.e. an undamaged signal. If N_T traps are spread uniformly throughout the column, the fraction of traps, F_T , S will interact with is $F_{silicon}$. If N_o traps are already occupied, the fraction of traps that interact with S and are vacant is:

$$F_T = \frac{1}{1+\beta} \left(\frac{S}{FWC} \right)^\beta - \frac{N_o}{N_T} \quad (C.25)$$

In the same way, the fraction of signal electrons F_e that see vacant traps is:

$$F_e = \frac{\frac{1}{1+\beta} \left(\frac{S}{FWC} \right)^\beta - \frac{N_o}{N_T}}{\frac{1}{1+\beta} \left(\frac{S}{FWC} \right)^\beta} \quad (C.26)$$

In the limit as $t \rightarrow \infty$ the number of captures N_c is:

$$\lim_{t \rightarrow \infty} N_c = \frac{N_T F_T S F_e}{N_T F_T + S F_e} = \frac{\gamma S^\beta - N_o}{\gamma S^{\beta-1} + 1} \quad \text{where} \quad \gamma = \frac{N_T}{(1+\beta) FWC^\beta} \quad (C.27)$$

C.4.1 Electron capture

Eqn. C.27 may be used when instantaneous trapping is occurring, but this is not the case with Gaia where electron capture times are comparable to the dwell time. For this reason, this model evaluates the probability of capture according to electron density, the latter of which grows as the signal integrates over the CCD.

Assuming traps in a CCD column interact with electrons for a period equal to half the TDI period (t), the probability a vacant trap will capture an electron is:

$$P_c = 1 - e^{-\frac{t}{2\tau_c}}. \quad (C.28)$$

τ_c is the capture time constant:

$$\tau_c = \frac{1}{\sigma v_{th} n_e} \quad (C.29)$$

σ is the capture cross-section, n_e is the electron density in the vicinity of the trap and v_{th} is the electron thermal velocity:

$$v_{th} = \sqrt{\frac{3kT}{m_e^*}}, \quad (C.30)$$

where m_e^* is the effective electron mass. Assuming V_c contains N_e :

$$\tau_c = \frac{V_c}{\sigma v_{th} N_e} \quad (C.31)$$

Implying P_c , the capture probability per vacant trap, is the following after including Eqn. C.23 and assuming t is constant:

$$P_c = 1 - e^{-\alpha N_e^{1-\beta}} \quad \text{where} \quad \alpha = \frac{t \sigma v_{th} F W C^\beta}{2V_g} \quad (C.32)$$

To find the effective capture probability for the entire column an approximation is made that $\bar{P}_c(S) \approx P_c(S/2)$. This is to avoid the unwieldy term that arises from instead integrating over S for each P_c . This leads to the number of electrons captured by a given trap species from S :

$$N_c = \frac{\gamma S^\beta - N_o}{\gamma S^{\beta-1} + 1} \left(1 - \exp\left\{-\alpha \left(\frac{S}{2}\right)^{1-\beta}\right\} \right) \quad (C.33)$$

To stop N_c from becoming negative when N_o exceeds the number of electrons interacting with traps, the following constraint is used:

$$N_c = \text{MAX} \left[\frac{\gamma S^\beta - N_o}{\gamma S^{\beta-1} + 1} \left(1 - \exp\left\{-\alpha \left(\frac{S}{2}\right)^{1-\beta}\right\} \right), 0 \right] \quad (C.34)$$

C.4.2 Electron release

The probability a trap will release an electron into a sample, assuming the trap has one TDI period (t) to do so, is:

$$P_r = 1 - e^{-\frac{t}{\tau_r}} \quad (C.35)$$

So the number of electrons released by a given trap species into S is:

$$N_r = N_o \left\{ 1 - e^{-\frac{t}{\tau_r}} \right\} \quad (C.36)$$

C.4.3 Setting the initial trap occupancy state

In order to determine the state of the initial N_o one of the factors that must be included is the diffuse optical background. Assuming traps will reach an equilibrium state, i.e. $N_c = N_r$ under a constant background:

$$N_{o,init} = \frac{AB}{A+C}, \quad (C.37)$$

where:

$$A = \frac{1 - \exp\left\{-\alpha\left(\frac{S_{dob}}{2}\right)^{1-\beta}\right\}}{\gamma S_{dob}^{\beta-1} + 1} \quad B = \gamma S_{dob}^{\beta} \quad C = 1 - e^{-\frac{t}{\tau_r}} \quad (\text{C.38})$$

The other deciding factor is the time since charge injection t_{ci} . Assuming the charge injection fills all traps and that they then empty at a rate determined by their exponential trap release time constant until they reach the occupancy level determined by the diffuse optical background:

$$N_{o,init} = N_T e^{-\frac{t_{ci}}{\tau_r}} + \frac{AB}{A+C} \left(1 - e^{-\frac{t_{ci}}{\tau_r}}\right). \quad (\text{C.39})$$

C.5 CDM03

The third Charge Distortion Model (CDM03) is a recent extension to CDM02, incorporating many of the features that were left up to the user to include in CDM02. For the most recent documentation on this model refer to Short et al. (2010).

BIBLIOGRAPHY

All of the documents shown here with an initial reference enclosed in bold square brackets ([...]) have been taken from a Livelink repository hosted by ESA. Livelink allows online access to an archive of scientific documents specifically relevant for Gaia and those working on the mission but not, in general, to those outside the Gaia project.

Abbey, A., Ambrosi, R., Smith, D., et al., 2001, 201 – 208

Abbey, A.F., Bennie, P.J., Turner, M.J.L., Altieri, B., Rives, S., 2003, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 513, 136 ,

URL <http://www.sciencedirect.com/science/article/B6TJM-49FGX8X-X/2/019742ce73b256aac7208d28cfe0bd56>,

Proceedings of the 6th International Conference on Position-Sensitive Detectors

Ambrosi, R., Denby, M., 2005, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 547, 346 ,

URL <http://www.sciencedirect.com/science/article/B6TJM-4G4X6TM-3/2/dfef458dbdee9fefee1b67c59f3fc746>

Bellini, A., Piotto, G., Bedin, L.R., et al., 2009, *A&A*, 493, 959, ADS Link

Bessell, M.S., 1990, *PASP*, 102, 1181, ADS Link

Bierman, G., 1997, *Factorization Methods for Discrete Sequential Estimation*, Academic Press, New York

Bristow, P., 2003, ArXiv Astrophysics e-prints, ADS Link

Brown, A., 2009, *Dealing with Gaia CCD radiation damage effects: overview of activities. ELSA CCD Modelling Workshop, Leiden, 19-21 Jan 2009*, Tech. rep., Leiden Observatory

[SWB-001], Brown, S., van Leeuwen, F., 2008, *Analysis of the CCN10 data set. CTI effects on a Gaia astrometric CCD*,

GAIA-CH-TN-IOA-SWB-001,

URL <http://www.rssd.esa.int/l1ink/liveliink/open/2818513>

[SWB-002], Brown, S., van Leeuwen, F., 2009, *An independent analysis of Astrium's Radiation Campaign 2 astrometric tests*,

GAIA-CH-TN-IOA-SWB-002,

URL <http://www.rssd.esa.int/l1ink/liveliink/open/2890023>

[SWB-003], Brown, S., van Leeuwen, F., 2009, *Further analysis of Radiation Campaign 1. An extension to SWB-001*,

GAIA-CH-TN-IOA-SWB-003,

URL <http://www.rssd.esa.int/l1ink/liveliink/open/2885744>

Cameron, P.B., Britton, M.C., Kulkarni, S.R., 2009, *AJ*, 137, 83, ADS Link

[JC-019], Castaneda, J., Fabricius, C., 2009, *Maximum Likelihood Fitting Algorithms: Test Report*, Tech. rep., UB,

- URL <http://gaia.esac.esa.int/dpacsvn/DPAC/CU3/docs/IDA/nonECSS/JC-019/GAIA-C3-TN-UB-JC-019.pdf>
- Chiaberge, M., Riess, A., Mutchler, M., Sirianni, M., Mack, J., 2006, In: A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel (ed.) *The 2005 HST Calibration Workshop: Hubble After the Transition to Two-Gyro Mode*, 36–+, ADS Link
- Clampin, M., Sirianni, M., Hartig, G., et al., 2002, *Experimental Astronomy*, 14, 107,
URL <http://dx.doi.org/10.1023/B:EXPA.00000004358.00269.51>,
10.1023/B:EXPA.00000004358.00269.51
- [**CL-001**], Corcione, L., Lattanzi, M., Terrazas-Vargas, J., et al., 2009, *Analysis of CCN10 radiation test campaign data; photometric performance and charge loss in Gaia irradiated AF CCD*,
GAIA-C3-TN-INAF-CL-001,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2873422>
- [**JDB-031**], de Bruijne, J., 2007, *Diffuse sky brightness and implications for a diffuse-optical-background source*,
GAIA-CA-TN-ESA-JDB-031,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2777186>
- [**JDB-043**], de Bruijne, J., 2008, *CCD pixel-pitch non-uniformity, stitch boundaries and supplementary buried channel*,
GAIA-CH-TN-ESA-JDB-043,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2837394>
- [**JDB-048**], de Bruijne, J., 2009, *PEM offset instability. Implications for astrometric performance derived from Monte Carlo centroiding experiments*,
GAIA-CA-TN-ESA-JDB-048,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2889588>
- [**JDB-061**], de Bruijne, J., 2010, *Serial-register-CTI-induced astrometric biases*,
GAIA-CA-TN-ESA-JDB-061,
URL <http://www.rssd.esa.int/lmlink/livelink/open/3021213>
- Dravins, D., Lindegren, L., Madsen, S., 1999, *A&A*, 348, 1040, ADS Link
- [**GAIA-E2V-TN-020**], e2v technologies, 2003, *AF CCD Design Report*,
GAIA-E2V-TN-020,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2729630>
- [**GAIAFPA.RPT.00445.T.ASTR**], EADS Astrium, 2005, *Irradiated AFCCD Test Report, Issue 1*, Tech. rep.
- [**GAIAFPA.RPT.00700.T.ASTR**], EADS Astrium, 2006, *CCN 10 Radiation Tests - Test Report*,
GAIAFPA.RPT.00700.T.ASTR,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2722564>
- [**GAIA.ASF.TCN.PLM.00071**], EADS Astrium, 2007, *Status at PDR on Radiation Recovery Activities*,
GAIA.ASF.TCN.PLM.00071,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2770066>
- [**GAIA.ASF.TCN.SAT.00035**], EADS Astrium, 2007, *Radiation Analysis Report*,

- GAIA.ASF.TCN.SAT.00035,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2722340>
- [GAIA.ASF.PLN.PLM.00052], EADS Astrium, 2008, *Radiation campaign #3 tests plan*,
GAIA.ASF.PLN.PLM.00052,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2822333>
- [GAIA.ASF.TCN.PLM.00197], EADS Astrium, 2008, *Tests report of the radiation campaign #2*,
GAIA.ASF.TCN.PLM.00197,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2823366>
- [GAIA.ASF.TCN.PLM.00212], EADS Astrium, 2008, *Radiation campaign #2 data delivery*,
GAIA.ASF.TCN.PLM.00212,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2814737>
- [GAIA.ASF.TCN.PLM.00311], EADS Astrium, 2008, *Radiation campaign #3. Bench characterization data delivery*,
GAIA.ASF.TCN.PLM.00311,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2859672>
- [GAIA.ASF.TCN.PLM.00313], EADS Astrium, 2008, *Radiation campaign #3. Astrium bench characterization tests report*,
GAIA.ASF.TCN.PLM.00313,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2863129>
- [GAIA.ASF.TCN.PLM.00349], EADS Astrium, 2008, *Radiation campaign #3. Astrometric regular data delivery*,
GAIA.ASF.TCN.PLM.00349,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2864842>
- [GAIA.ASF.PLN.PLM.00073], EADS Astrium, 2009, *Radiation campaign #4 tests plan*,
GAIA.ASF.PLN.PLM.00073,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2963939>
- [GAIA.ASF.TCN.PLM.00368], EADS Astrium, 2009, *PEM response stability anomaly test report*,
GAIA.ASF.TCN.PLM.00368,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2900733>
- [GAIA.ASF.TCN.PLM.00415], EADS Astrium, 2009, *Radiation campaign #3. Astrometric sky-like data delivery*,
GAIA.ASF.TCN.PLM.00415,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2899661>
- [GAIA.ASF.TCN.PLM.00439], EADS Astrium, 2009, *Radiation campaign #3. Serial register additional tests data delivery*,
GAIA.ASF.TCN.PLM.00439,
URL <http://www.rssd.esa.int/l1ink/livelihood/open/2907680>
- [GAIA.ASF.TCN.PLM.00442], EADS Astrium, 2009, *Radiation Campaign #3. Astrium report of the serial register tests*,
GAIA.ASF.TCN.PLM.00442,

- URL <http://www.rssd.esa.int/lmlink/livelink/open/2911255>
- [GAIA.ASF.TCN.PLM.00447], EADS Astrium, 2009, *Radiation campaign #3. Astrium astrometric regular test report*,
GAIA.ASF.TCN.PLM.00447,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2961322>
- [GAIA.ASF.TCN.PLM.00497], EADS Astrium, 2010, *Radiation campaign #3. Astrium astrometric skylike tests report*,
GAIA.ASF.TCN.PLM.00497,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2975103>
- [GAIA.ASF.TCN.PLM.00516], EADS Astrium, 2010, *Radiation campaign #4. Astrometric data delivery*,
GAIA.ASF.TCN.PLM.00516,
URL <http://www.rssd.esa.int/lmlink/livelink/open/3011179>
- [GAIA.ASF.TCN.SAT.00133], EADS Astrium, 2010, *Dead Time Budget*,
GAIA.ASF.TCN.SAT.00133,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2722338>
- ESA, 1997, *The Hipparcos and Tycho Catalogues*, ESA,
ESA SP-1200
- ESA, 2010a, *Science performance*, ESA,
For further information go to <http://www.rssd.esa.int/gaia/> and click on 'Science performance' on the left-hand information panel.
- ESA, 2010b, *Gaia Image and Movie Gallery*, ESA,
For further information go to <http://www.rssd.esa.int/gaia/> and click on 'Image and Movie Gallery' on the left-hand information panel.
- [GAIA-EST-MN-03815], ESA Gaia Project, 2007, *Meeting on Gaia Radiation Calibration WG#5*,
GAIA-EST-MN-03815,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2824120>
- [GAIA-EST-MN-05213], ESA Gaia Project, 2008, *Meeting on Gaia Radiation Calibration WG#6*,
GAIA-EST-MN-05213,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2824128>
- [GAIA-EST-MN-06347], ESA Gaia Project, 2008, *Meeting on Gaia Radiation Calibration WG#7*,
GAIA-EST-MN-06347,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2853130>
- [GAIA-EST-RD-00553], ESA Gaia Project, 2010, *Mission Requirements Document (MRD)*,
GAIA-EST-RD-00553,
URL <http://www.rssd.esa.int/lmlink/livelink/open/463164>
- [DG-013], Gardiol, D., Russo, F., Busonero, D., 2009, *Description of the Image Parameter Determination Algorithm for IDA*,
GAIA-C3-TN-INAFA-DG-013,
URL <http://www.rssd.esa.int/lmlink/livelink/open/2908441>

- Giavalisco, M., 2003, *Minimizing CTE losses in the WFC3 CCDs: Post Flash vs. Charge Injection*, Tech. rep.
- Hall, R.N., 1952, *Phys. Rev.*, 87, 387
- [NCH-003], Hambly, N., Cross, N., 2008, *Independent Analysis of Astrium Test Campaign 2 Charge Release Data*,
GAIA-C5-TN-IFA-NCH-003,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2862092>
- [EH-176], Hoeg, E., de Bruijne, J., 2006, *Gaia sampling scheme*,
GAIA-CA-TN-NBI-EH-176,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/1112518>
- Høg, E., 2009, *Experimental Astronomy*, 25, 225,
URL <http://dx.doi.org/10.1007/s10686-009-9156-7>,
10.1007/s10686-009-9156-7
- Høg, E., Fabricius, C., Makarov, V.V., et al., 2000, *A&A*, 355, L27, ADS Link
- [BH-001], Holl, B., 2009, *Mean CCD-transit intervals estimated from cumulative star densities*,
GAIA-C3-TN-LU-BH-001,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2852887>
- [BH-002], Holl, B., Lindegren, L., Hobbs, D., 2009, *AGISLab - A facility for experimental astrometric solutions*,
GAIA-C3-TN-LU-BH-002,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2958464>
- Hopkinson, G., 2000, *Nuclear Science, IEEE Transactions on*, 47, 2460
- [254.DO.28], Hopkinson, G., 2006, *Final report on radiation testing of Gaia AF and MBP CCDs*,
254.DO.28,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2731616>
- Hopkinson, G., Mohammadzadeh, A., 2003, *Nuclear Science, IEEE Transactions on*, 50, 1960
- Hopkinson, G., Dale, C., Marshall, P., 1996, *Nuclear Science, IEEE Transactions on*, 43, 614
- Howell, S., 2006, *Handbook of CCD Astronomy, Second edition*, 5, Cambridge University Press, COHRA
- Janesick, J., 2001, *Scientific Charge-Coupled Devices*, SPIE Press, Bellingham, Washington
- Kirsch, M.G., Abbey, A., Altieri, B., et al., 2005, vol. 5898, 58980S, SPIE,
URL <http://link.aip.org/link/?PSI/5898/58980S/1>
- van Leeuwen, F., 2007, *Hipparcos, the New Reduction of the Raw Data*, Springer, astrophysics and space science library. vol. 350 edn.
- [SAG-LL-032], Lindegren, L., 2000, *Centroiding on GAIA CCD sample data*,
SAG-LL-032,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/356853>
- [GAIA-LL-034], Lindegren, L., 2001, *Proposed prototype processes for the GAIA Global Iterative Solution*,
GAIA-LL-034,

- URL <http://www.rssd.esa.int/lmlink/livelihood/open/357825>
- [GAIA-LL-059], Lindegren, L., 2005, *An alternative scheme for the GIS processing*, GAIA-LL-059,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/435451>
- Lindegren, L., 2005, The astrometric instrument of Gaia: principles, ESA : The Three-Dimensional Universe with Gaia (eds. Turon C, O’Flaherty KS, Perryman MAC), SP-576,
URL http://www.rssd.esa.int/index.php?project=Gaia&page=Gaia_2004_Proceedings
- [LL-068], Lindegren, L., 2006, *Centroid definition for the Astro line spread function*, GAIA-C3-TN-LU-LL-068,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2698669>
- [LL-069], Lindegren, L., 2007, *Simplified centroiding algorithm for radiation test data*, GAIA-C3-TN-LU-LL-069,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2742592>
- [LL-075], Lindegren, L., 2008, *Modelling radiation damage effects for Gaia - A first-order Charge Distortion Model (CDM-01)*, GAIA-C3-TN-LU-LL-075,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2824544>
- [LL-078], Lindegren, L., 2008, *A general Maximum-Likelihood algorithm for model fitting to CCD sample data*, GAIA-C3-TN-LU-LL-078,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2861864>
- [LL-082], Lindegren, L., 2009, *Demonstrating the feasibility of a CTI mitigation strategy for astrometry using astrometric sky-like test data*, GAIA-C3-TN-LU-LL-082,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2907018>
- Ma, C., Arias, E.F., Eubanks, T.M., et al., 1998, *The Astronomical Journal*, 116, 516,
URL <http://stacks.iop.org/1538-3881/116/i=1/a=516>
- Massey, R., Stoughton, C., Leauthaud, A., et al., 2010, *MNRAS*, 401, 371, ADS Link
- [GAIA-FM-023], Mignard, F., 2005, *Ephemeris requirements for Gaia*, GAIA-FM-023,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/490679>
- [FM-030], Mignard, F., Drimmel, R., 2007, *DPAC: Proposal for the Gaia Data Processing*, GAIA-CD-SP-DPAC-FM-030-02,
URL <http://www.rssd.esa.int/lmlink/livelihood/open/2720336>
- Nelder, J., Mead, R., 1965, *Computer Journal*, 7, 308
- Perryman, M., 2008, *Astronomical applications of Astrometry: Ten years of exploitation of the Hipparcos satellite data*, Cambridge University Press
- Perryman, M.A.C., de Boer, K.S., Gilmore, G., et al., 2001, *A&A*, 369, 339, ADS Link
- Piatek, S., Pryor, C., Bristow, P., et al., 2007, *AJ*, 133, 818, ADS Link
- Piatek, S., Pryor, C., Olszewski, E.W., 2008, *AJ*, 135, 1024, ADS Link

- Prigozhin, G., Bautz, M., Kissel, S., LaMarr, B., 2007, vol. 3, 2271–2277
- Prod'homme, T., 2010, In: Proceedings of SPIE: Gaia: at the frontiers of astrometry, 7-11 June 2010, Sévres, France.,
Awaiting publication.
- [TPH-001], Prod'homme, T., Brown, S., 2009, *Determination of Gaia Astrometric CCD Physical Parameters using Radiation Campaign 2 Charge Injection Profiles*,
GAIA-C5-TN-LEI-TPH-001,
URL <http://www.rssd.esa.int/l1ink/liveliink/open/2885098>
- [TPH-003], Prod'homme, T., Weiler, M., Brown, S., 2009, *CDM02 validation. A comparison between CDM02 outcomes and RC2 data*,
GAIA-CH-TN-LEI-TPH-003,
URL <http://www.rssd.esa.int/l1ink/liveliink/open/2936008>
- Propris, R.D., Rich, R.M., Mallery, R.C., Howard, C.D., 2010, ApJ, 714, L249
- Reffert, S., 2010, New Astronomy Reviews, In Press, Corrected Proof, ,
URL <http://www.sciencedirect.com/science/article/B6VNJ-50JHC35-6/2/bf893f9c35c17cd742014c6b9e57c231>
- Seabroke, G., Holland, A., Cropper, M., 2008, In: Proceedings of SPIE: High Energy, Optical, and Infrared Detectors for Astronomy III, 23-27 June 2008, Marseille.,
URL <http://oro.open.ac.uk/view/year/2008.html>
- [GMS-002], Seabroke, G., Burt, D., Robbins, M., Holland, A., 2010, *Modelling Gaia pixels: Silvaco 3D simulations to determine SBC capacities*,
GAIA-C5-TN-OU-GMS-002,
draft
- [GMS-001], Seabroke, G., Burt, D., Robbins, M., et al., 2010, *Modelling Gaia pixels. Silvaco planar simulations of buried channels*,
GAIA-C5-TN-OU-GMS-001,
URL <http://www.rssd.esa.int/l1ink/liveliink/open/2966507>
- Shockley, W., Read, W.T., 1952, Phys. Rev., 87, 835
- [AS-011], Short, A., 2006, *Independent analysis of EADS Astrium CCN10 radiation test data*,
GAIA-CH-TN-ESA-AS-011,
URL <http://www.rssd.esa.int/l1ink/liveliink/open/2705128>
- [AS-015], Short, A., 2009, *Charge Distortion Model 02 (CDM02)*,
GAIA-CH-TN-ESA-AS-015,
URL <http://www.rssd.esa.int/l1ink/liveliink/open/2882372>
- Short, A., 2009, *Summary of Radiation Damage Calibration Effort, Gaia Science Team presentation*,
Tech. rep., ESTEC
- Short, A., Prod'homme, T., Weiler, M., Brown, S., Brown, A., 2010, vol. 7742, 774212, SPIE,
URL <http://link.aip.org/link/?PSI/7742/774212/1>
- Sirianni, M., Mutchler, M., 2006, In: Beletic, J., Beletic, J., Amico, P. (eds.) *Scientific detectors for astronomy 2005*, 171–178, Springer Netherlands

- Smith, D.R., Holland, A.D., Hutchinson, I.B., et al., 2004, vol. 5501, 189–196, SPIE,
URL <http://link.aip.org/link/?PSI/5501/189/1>
- [FVL-006], van Leeuwen, F., Lindegren, L., 2007, *Outline of the Radiation Calibration Strategy for Astrometry, Photometry and Spectroscopy*,
GAIA-C5-TN-IOA-FVL-006,
URL <http://www.rssd.esa.int/l1ink/liveliink/open/2775637>
- Waczynski, A., Polidan, E., Marshall, P., et al., 2001, Nuclear Science, IEEE Transactions on, 48, 1807
- Wielen, R., Schwan, H., Dettbarn, C., et al., 2002, VizieR Online Data Catalog, 1264, 0, ADS Link
- Wilson, E.B., Hilferty, M.M., 1931, Proceedings of the National Academy of Sciences of the United States of America, 17, 684,
URL <http://www.pnas.org/content/17/12/684.short>