Universiteit Leiden
The Netherlands

From electrons to stars : modelling and mitigation of radiation damage effects on astronomical CCDs

Prod'homme, T.

Citation

Prod'homme, T. (2011, November 22). *From electrons to stars : modelling and mitigation of radiation damage effects on astronomical CCDs*. Retrieved from https://hdl.handle.net/1887/18136

Version: Corrected Publisher’s Version

License: Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden

Downloaded from: https://hdl.handle.net/1887/18136

Note: To cite this publication please use the final published version (if applicable).

From Electrons to Stars

Thibaut Prod'homme

From Electrons to Stars

Modelling and mitigation of radiation damage effects on astronomical CCDs

Thibaut Prod'homme

INVITATION
UITNODIGING

On Tuesday, November 22 at **15:00h**
I will defend my thesis

From Electrons to Stars
Modelling and mitigation of radiation damage effects on astronomical CCDs

in the **Leiden University**
Academiegebouw
Rapenburg 73

We will celebrate this event from **20:00h**
at **Café de Bonte Koe**
Hooglandsekerk-Choorsteeg 13

You are welcome to attend the defense and the following celebration

Thibaut Prod'homme
thibaut.prodhomme@gmail.com

From Electrons to Stars

Modelling and mitigation of radiation damage effects on astronomical CCDs

From Electrons to Stars

Modelling and mitigation of radiation damage effects on astronomical CCDs

Proefschrift

ter verkrijging van
de graad van Doctor aan de Universiteit Leiden,
op gezag van de Rector Magnificus prof. mr. P.F. van der Heijden,
volgens besluit van het College voor Promoties
te verdedigen op dinsdag 22 november 2011
klokke 15.00 uur

door

Thibaut Prod'homme

geboren te Flers, Normandië, Frankrijk
in 1984

Promotiecommissie

Promotor:	Prof. dr. K.H. Kuijken
Co-promotor:	Dr. A.G.A. Brown
Overige leden:	Prof. dr. M.W. Beijersbergen
	Prof. dr. C.W.M. Fridlund
	Dr. R.J. Massey (Royal Observatory, Edinburgh, UK)
	Prof. dr. H.J.A. Röttgering

ISBN: 978-94-6191-057-8

The paper used for the inside of this thesis is FSC-certified.

‘And the stars look very different today.’

Space Oddity - David Bowie

Cover by Julien Alday
To visualize outer space and immense manifestations such as solar eruptions, French illustrator Julien Alday manipulates photographic elements by shooting oxidation of metallic plates. In recording a small chemical reaction, he ultimately depicts events on a cosmic scale. From top to bottom the graphical elements represent the Gaia spacecraft downlinking data to Earth, the Earth and its magnetosphere, the Sun and the solar wind.
`http://www.julienalday.com`

Table of contents

Page

Chapter 1. Introduction **1**
1.1 Context . . . 2
1.2 Astrometry . . . 4
1.3 Gaia . . . 7
1.4 Charge-Coupled Device . . . 12
1.5 Radiation damage . . . 17
1.6 The Gaia CTI mitigation strategy . . . 20
1.7 This thesis . . . 25
1.8 Lessons learned and outlook . . . 28

Chapter 2. Electrode level Monte Carlo model of radiation damage effects on astronomical CCDs **31**
2.1 Introduction . . . 32
2.2 Model description . . . 33
2.2.1 Simulation process . . . 35
2.2.2 Effective charge capture and release probabilities . . . 37
2.2.3 Charge density distribution modelling . . . 40
2.2.4 CCD illumination history . . . 42
2.3 Radiation tests . . . 45
2.4 Model verification and comparison to experimental data . . . 46
2.4.1 Model comparison to analytical prediction . . . 48
2.4.2 Model comparison to experimental data . . . 49
2.5 Conclusion . . . 59

Chapter 3. The impact of CCD radiation damage on Gaia astrometry part I Image location estimation in the presence of radiation damage **61**
3.1 Introduction . . . 62
3.2 Overall methodology . . . 63
3.3 Generating Gaia-like observations . . . 65
3.3.1 How Gaia observes . . . 66
3.3.2 Construcing a Gaia-like reference image . . . 67
3.3.3 Monte Carlo simulations of observations . . . 68
3.3.4 Simulation of the CTI effects . . . 69
3.4 The Gaia image location estimation procedure . . . 72
3.4.1 Observation model: scene . . . 72
3.4.2 Maximum-likelihood estimation of the image parameters . . . 73
3.4.3 First image parameter estimates and LSF model . . . 74
3.4.4 Iterative image parameter and LSF model improvement . . . 74

3.5 Theoretical and actual limit to the image location accuracy 76
3.5.1 Definition of the astrometric Cramér-Rao bound 76
3.5.2 Location independent error and standard deviation 78
3.5.3 CTI-free location bias results per magnitude 79
3.5.4 CTI-free location accuracy per magnitude 80
3.5.5 Radiation damage intrinsic uncertainty increase 82
3.5.6 Radiation induced image location bias 84
3.5.7 Comparison with experimental data 84
3.5.8 Damaged location estimation standard errors 85
3.6 CTI effects mitigation 86
3.6.1 Potential alternative approaches 87
3.6.2 A complete forward modelling approach 89
3.6.3 Testing the forward modelling approach 91
3.6.4 Current best CDM candidate 91
3.6.5 The forward modelling approach initialization 93
3.6.6 Image location bias and accuracy recovery 95
3.6.7 Image flux bias recovery 98
3.7 Discussion 99
3.8 Conclusions 100
3.9 Future work 102

Chapter 4. The impact of CCD radiation damage on Gaia astrometry part II Effect of image location errors on the astrometric solution 107
4.1 Introduction 108
4.2 Methodology 109
4.2.1 Generating Gaia-like observations 109
4.2.2 Charge transfer inefficiency model 113
4.2.3 Simulating the astrometric solution 124
4.3 Results 133
4.3.1 How a constant bias would affect the astrometric solution 135
4.3.2 Detailed model results 136
4.3.3 Residual based correction 138
4.3.4 Disturbing stars 139
4.4 Discussion 140
4.4.1 Implications for the scientific performance of Gaia 140
4.4.2 CTI mitigation in Gaia 143
4.5 Conclusions 144

Chapter 5. Stress-testing a fast analytical Charge Transfer Inefficiency model 147
5.1 Introduction 148
5.2 Charge Distortion Model 150
5.2.1 Principles 150
5.2.2 Usage 151
5.3 Gaia CCD design and operation 153
5.4 Experimental tests 154

5.4.1 Experimental test setup . . . 154
5.4.2 The test dataset . . . 155
5.4.3 Modelling the CDM input signal . . . 155
5.5 Model-data comparison methodology . . . 156
5.5.1 Comparison procedure . . . 156
5.5.2 Determination of the best-fitting CDM parameters . . . 156
5.5.3 Calibration schemes . . . 157
5.6 Results . . . 158
5.6.1 The best agreement achievable . . . 158
5.6.2 Calibration on a per magnitude basis . . . 161
5.6.3 Calibration on a per CCD basis . . . 162
5.7 Discussion . . . 165
5.7.1 Experimental instabilities . . . 165
5.7.2 The CDM performance in the context of the Gaia approach to CTI mitigation . . . 166
5.7.3 Calibrating CDM . . . 168
5.8 Conclusions and future work . . . 170

Chapter 6. Digging supplementary buried channels: Investigating the notch architecture within the CCD pixels on ESA's Gaia satellite 173
6.1 Introduction . . . 174
6.2 The Gaia CCD pixel architecture . . . 176
6.3 The SBC Full Well Capacity issue . . . 180
6.3.1 Measurements on irradiated CCDs . . . 180
6.3.2 Measurements on non-irradiated CCDs . . . 182
6.4 FPR measurements: Model to Data comparison . . . 185
6.4.1 Principles of the FPR measurement . . . 185
6.4.2 Selection of the experimental data . . . 187
6.4.3 Generation of the synthetic data . . . 188
6.4.4 Comparison results . . . 192
6.4.5 Discussion of FPR measurements . . . 194
6.5 Pocket pumping measurements . . . 195
6.6 Comparison of Radiation Campaigns 3 and 4 test data from the same CCDs 198
6.7 Discussion . . . 200
6.7.1 Non-functional SBC statistics and predictions . . . 200
6.7.2 Impact on the Gaia image location accuracy . . . 205
6.8 Conclusions . . . 207
6.9 Recommendations . . . 209

Bibliography 212
List of acronyms 218
Nederlandse samenvatting 221
Résumé en français 231
Curriculum vitae 241
Acknowledgments 243

Chapter 1

Introduction

Replacing ships by spacecrafts and eyes by their modern digital extension — Charge-Coupled Devices (CCDs) — the scientists that design, build, and use space missions follow in the footsteps of the explorers from the Renaissance who first unravelled the outline of new continents and questioned the world as it was known at the time. Outer space offers astronomers control, stability, and the possibility of observing the entire sky within and without the optical window, free from the disturbing effects of the Earth's atmosphere. It is thus not surprising that many of the most recent advances in our understanding of the universe have been enabled by space-born telescopes and experiments. However space does not only come with advantages; vacuum, high temperature gradients, and most importantly particle radiation render it the most hostile environment ever conquered by mankind. Energetic particles such as protons emitted by the Sun can slowly degrade the performance of instruments on-board satellites as well as instantaneously interrupt their functioning. To carry out accurate measurements from space, shielding and utilizing radiation-hard materials and components often does not suffice so that the effects of radiation damage on the measurements must be taken into account in the processing of the data downlinked to Earth. CCD type detectors are now so effective at detecting and imaging incoming photons that they have become the core of many instruments on board satellites observing the sky as well as the Earth and the many bodies of the solar system, from near-infrared to X-ray wavelengths and in particular in the optical. The work presented in this thesis is part of an on-going effort to understand and mitigate the effects of radiation damage in astronomical CCDs. My research was motivated by and took place in the challenging context of the European Space Agency's (ESA) astrometric mission, Gaia, for which radiation damage has been considered since its conception as one of the most important threats to its scientific performance. In this chapter, I first present the context of this research and then give background information that will support the understanding of the other chapters. Finally I provide an overview of the thesis by summarizing the contents and conclusions of each chapter and close this introductory chapter with an outlook on further development of this work and of the field.

1.1 Context

Gaia aims at creating the largest and most accurate map to date of our galaxy, the Milky Way, by collecting for one billion stars: position, proper motion, and most importantly distance from our host star, the Sun. This map will cover regions we can only dream of physically reaching one day. Gaia is thus not about space tourism but about (scientific) discovery; a stereoscopic and dynamic map of the Milky Way will enable astronomers to explain the formation and evolution of our galaxy. It will also allow us to sharpen our understanding of stars and the solar system as well as provide tests of Einstein's general relativity.

In this quest towards a better understanding of our universe, knowing the distances of celestial objects from the Sun is critical. It is for instance the only means of knowing the intrinsic brightness (absolute magnitude) of an object. Distance can only be measured in a model-independent way by means of the trigonometric parallax: the apparent motion of a foreground object with respect to background objects induced by a change in position of the observer. For a fixed baseline (i.e. distance between two observing locations), the amplitude of the apparent motion of the foreground object is determined by its distance from the observer: the further the object is, the smaller the amplitude. In fact stars are so distant that astronomers must use the one year journey of the Earth around the Sun to be able to detect stellar parallaxes (see Fig. 1.1). It is indeed by performing observations all along the year and thus using the longest baseline available — the diameter of the Earth's orbit (i.e. 2 astronomical units, AU) — that astronomers first measured the long predicted stellar parallax at the beginning of the 19^{th} century. The first parallax measurement was performed by Friedrich Bessel in 1838 for one of the brightest and closest stars in the sky, 61 Cygni. Parallax is usually expressed as an angle: ϖ (see Fig. 1.1). For 61 Cygni, $\varpi = 287$ milli-arcsecond (mas), this is less than $1/10\,000^{\text{th}}$ of a degree. Hence the parallax measurement of distant objects requires: (i) a high stability of the instrument and observing platform over time as measurements of the star's position must be repeated at different epochs and connected together; (ii) accurate image location to determine the exact position of the star on the sky at each epoch; and finally (iii) a large photon collecting efficiency to enable very accurate image location and the detection of faint distant stars. The number of parallax measurements have been constantly increasing and improving in accuracy since the 19^{th} century through the refinement of the techniques used, but most importantly through the technological progress made in light collection and detection. In 1989, ESA's Hipparcos mission brought to its climax high accuracy parallax measurements by the first use of a dedicated satellite in space equipped with photomultiplier type detectors.

Today's CCDs enable the collection of nearly 100% of the incoming photons depending on the wavelength band and generally contain several million pixels. Depending on the focal length of the telescope used and the physical size of the pixel, a CCD's platescale can reach 1 mas or less. Gaia will make use of 106 high-performance custom-made CCDs with the aim of superseding Hipparcos results by several order of magnitude in terms of parallax accuracy (from mas to μas) and number of observed stars (from $100\,000$ to 1 billion). The performance of CCDs is degraded in space by the effects of solar radiation. Especially affected is the capability to transfer charges from one

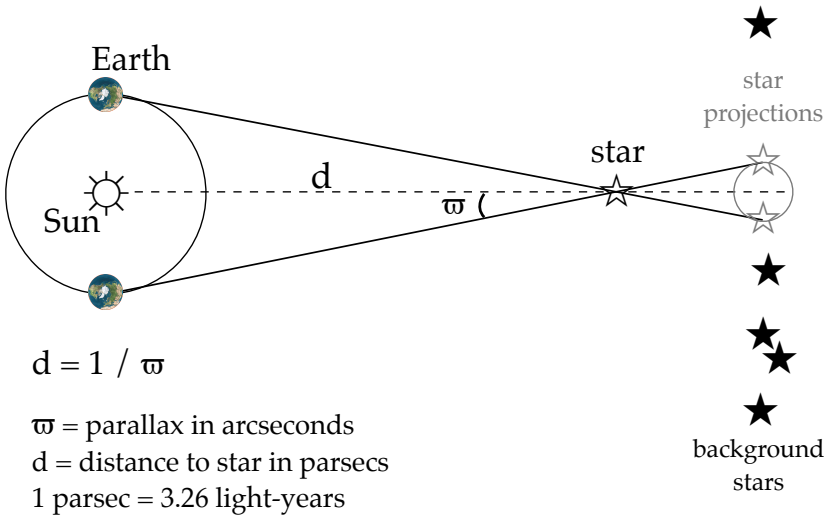


Figure 1.1 — Schematic representation of the stellar parallax and its measurement using the diameter of the Earth's as the baseline. The parallactic motion is depicted by the small grey circle. It is the apparent motion of the star of interest with respect to background (more distant) stars. The parallactic motion is periodic and mirrors the Earth's orbit around the Sun as seen from the star of interest.

pixel to the other in the imaging area, down to the CCD serial register and its output node. In astronomy, the issue of radiation-induced Charge Transfer Inefficiency (CTI) has been identified since the first use of CCDs by the Hubble Space Telescope (HST) in 1990. However CTI and its effects on astronomical measurements have never been studied as thoroughly as in the past decade, due to the more and more demanding use of CCDs for astronomical applications, pushing these devices to their limits. Due to a particular way of operating CCDs, very faint signal levels, and stringent image quality requirements, the preparation of Gaia triggered the need for a more profound understanding of CTI as well as novel solutions to counter and mitigate CTI effects at both the hardware and software level.

Early on in the preparation of the mission, CTI was recognized as potentially a major threat to the Gaia science performance as CTI would not only blind Gaia to the most distant objects but lead to systematic errors in all the Gaia measurements. As a consequence industrial partners have been mandated to carry out several campaigns of experimental tests on irradiated Gaia CCDs to characterize the effects of CTI on the Gaia measurements, to identify and optimize the use of hardware CTI-counter-measures, and to support the modelling efforts and the elaboration of a CTI mitigation strategy by the Gaia Data Processing and Analysis Consortium (DPAC). In this context my research focused primarily on the modelling of the CTI effects supported by the analysis of the experimental test data. I developed the most detailed model to date of CTI in CCDs that enables simulating the operation of irradiated devices (Chapter 2). Using this model I have been able to verify and enhance our current understanding of CTI as well as support the characterization of CCDs and the understanding of experimental results (Chapters 3 and 6). As part of this research I conducted the comprehensive re-assessment of the performance of Gaia taking into account radiation damage (Chapters 3 and 4). Finally I took part in the effort of countering CTI by elaborating, testing

and improving a forward modelling approach to mitigate the CTI effects on the Gaia measurements (Chapters 3 and 5), as well as test and explore the potential of a specific hardware mitigation tool (Chapter 6).
In the following the key elements of this thesis context are presented in more detail: Astrometry and its history and principles (Section 1.2), the Gaia mission (Section 1.3), the CCD (Section 1.4): the device that enables the extreme accuracy of Gaia's astrometric measurements, and radiation damage (Section 1.5) which is a threat to this accuracy.

1.2 Astrometry

Definition
Astrometry is the field of Astronomy that concentrates on the accurate determination and detailed study of the position and motion of stars and other objects on the sky (such as planets, galaxies etc.). The motion of a star as seen from Earth (see Fig. 1.2) is composed of (i) its proper motion, the angular change of the star's position on the sky over time that reflects the motion of the star through space with respect to the sun, (ii) its parallactic motion, the projection on the sky of the orbital path of the Earth around the Sun as seen from the star, and (iii) its orbital motion if the star belongs to a binary or multiple system. The radial velocity (the motion of a star along the line of sight) should also be accounted for in the astrometric motion of the star on the sky, although its determination is in practice not possible through astrometry. The Equatorial coordinate system is usually used as reference for astrometric measurements. In this system, the astrometric parameters of a star are defined and denoted as follows:

1. The position parameters: α and δ, respectively the right ascension and declination.
2. The proper motion parameters: $\mu_{\alpha*}$ and μ_δ respectively the angular changes per year along the right ascension and declination directions. [1]
3. The parallax: ϖ (see below for further details).

History
Although the first stellar parallax (distance) measurement was made during the first half of the 19$^{\text{th}}$ century, the determination of star positions on the celestial sphere is a much older activity. The art of mapping the heavens may be traced all the way back to the Paleolithic period with some of the paintings of the Lascaux caves (15 300 BC), disputably representing several constellations. The Greek Hipparchus (190–120 BC) is considered the father of astrometry as a science. He laid the mathematical groundwork for astrometric measurements and created a catalogue containing positions for about 1000 stars, which was later published during Roman times by Ptolemy in his Almagest. By establishing the observatory and research centre Uraniborg, the danish astronomer Tycho Brahe (1546–1601) triggered the first significant qualitative progress in Astrometry since Roman times. The work of Brahe (a catalogue of 1000 stars with a two orders of magnitude improvement in accuracy with respect to Hipparchus) has

1. The notation $\mu_{\alpha*} = \mu_\alpha \cos\delta$ signifies that the proper motion in right ascension is expressed as a true arc length on the sky (as opposed to $\mu_\alpha = d\alpha/dt$).

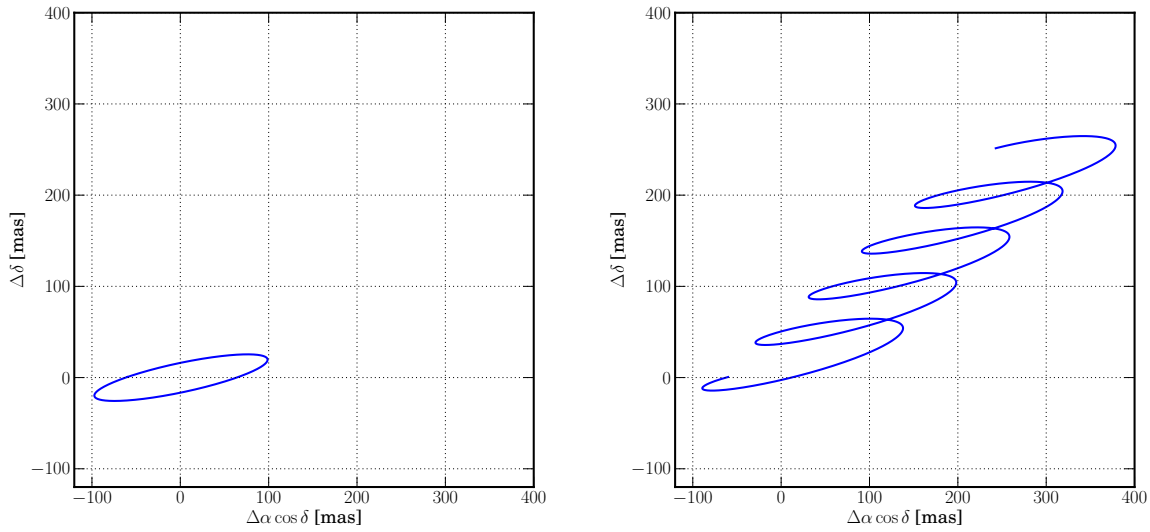


Figure 1.2 — Simulation of the motion of a star on the sky as observed from Earth in the Equatorial coordinate system. The star's astrometric parameters are: Position $(\alpha, \delta) = (60°, 30°)$, Parallax $\varpi = 100$ mas, and Proper motion $(\mu_{\alpha*}, \mu_\delta) = (60, 50)$ mas yr^{-1} **Left:** Parallactic motion only. **Right:** Parallactic and proper motion for a simulated continuous observation of 5 years. Figure courtesy of A.G.A. Brown (Leiden Observatory)

been followed by 400 years of constant progress mainly driven by the use of better equipment. The next quantitative and qualitative jump was enabled by space astrometry first realized by the ESA satellite Hipparcos launched in 1989 and operated for three and a half years. The mission outcomes were the Hipparcos catalogue, containing all astrometric parameters (i.e. including parallax) for about 100 000 stars, and the Tycho catalogue, containing positions and proper motions for 2.5 million stars. See the work of Høg (2008) for further insights into the long history of Astrometry.

Principles of ground-based and space-based parallax measurements
Figure 1.1 depicts the principle of a parallax measurement for which the background stars are assumed to be located at infinity. In this case, the distance d of the star of interest from the Sun, can be computed ignoring the parallax of the reference stars:

$$d = \frac{\text{baseline}}{2} \frac{1}{\tan \varpi} = \frac{1\,\text{AU}}{\varpi}\,, \tag{1.1}$$

where we make use of the small angle approximation. However background stars with respect to which the parallax of the star of interest is measured are in practice never located at infinity. As a consequence a parallax measurement as carried out from the ground is a differential measurement (Fig. 1.3, left). The angular difference ϕ between the star of interest and a background star located in the same field of view is measured at two different epochs (6 months interval in the depicted case) and the parallax of the star of interest can then be computed:

$$\varpi = \varpi_0 + (\phi_2 - \phi_1)/2 \tag{1.2}$$

In this case the measured value of the parallax depends on the parallax of the background star ϖ_0. If ϖ_0 is ignored or unknown, a systematic error is introduced in the measurement of ϖ. One can avoid this systematic error by selecting a reference star located at $\sim 90^\circ$ from the star of interest (Fig. 1.3, right). This technique is referred to as wide-angle measurement and can only be carried out in space thanks to weightlessness and the absence of turbulence from the Earth's atmosphere. Indeed, reaching the highest accuracy on the ground requires that the light path through the atmosphere is the same for the reference star and the star of interest, the same holds for the gravitational pull on the instrument. The practical implementation of the wide-angle technique requires two fields of view separated by a large enough angle (the so-called basic angle). It was used for the first time by Hipparcos and, soon, Gaia and the Japanese mission Nano-JASMINE will also make use of this technique to carry out absolute astrometric measurements. In order to achieve the greatest accuracy possible when using this technique, it is preferable to perform mainly one-dimensional angle measurements along great circles in the sky (Lindegren 2005). This implies that the exact value of $\sim 90^\circ$ for the basic angle should be avoided (see Lindegren & Bastian 2011). Based on design constraints and other considerations, Hipparcos, Gaia, and Nano-JASMINE have the following values for the basic-angle: $\sim 58^\circ$, $106\overset{\circ}{.}5$, and $\sim 99\overset{\circ}{.}5$, respectively. A comprehensive description of the advantages of wide-angle measurements over differential small field measurements is given by Lindegren (2005).

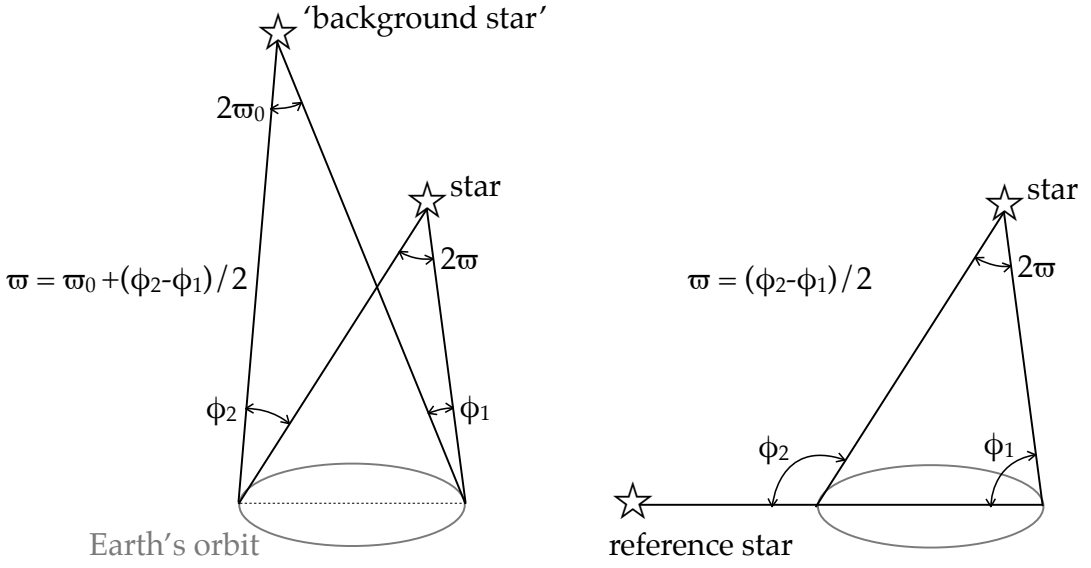


Figure 1.3 — Parallax measurement principle. **Left:** Differential small field measurement, the 'background' star is in the same field of view as the star of interest. **Right:** Wide-angle or absolute measurement, the reference star is separated by a large angle from the star of interest. This technique can in practice only be used in space. Figure adapted from Lindegren (2005).

1.3 Gaia

Gaia is an ESA space mission that aims at carrying out absolute astrometric measurements with unprecedented accuracy for about one percent of the stellar population in the Milky Way (1 billion stars) using the wide-angle technique described above. It is in particular by the use of a very large CCD mosaic camera that Gaia will supersedes the results of Hipparcos in accuracy, statistics, and completeness by several orders of magnitude. The formidable wealth of scientific discoveries that were enabled by the Hipparcos Catalogue can only lead us to expect an even broader and deeper impact of the future Gaia data on almost every field of modern astronomy and even beyond on solar system and fundamental physics. Gaia is scheduled for launch in 2013 by a Soyuz-Fregat rocket from the Guiana Space Centre (French Guiana, South America). Gaia will operate for a nominal lifetime of 5 years (with a potential extension of 1 year) in an orbit around L2, the second Lagrangian point of the Sun-Earth system. The mission is currently in the final phases of its preparation both from the satellite manufacturing and data processing perspectives which constitute equal challenges as shall be explained upon in the following. EADS Astrium is the Gaia prime contractor, the main industrial partner responsible for the coordinating the manufacturing and assembly of the Gaia spacecraft. DPAC is the international consortium of scientists ($\sim$ 400) in charge of elaborating and performing the Gaia data processing. The mission and its

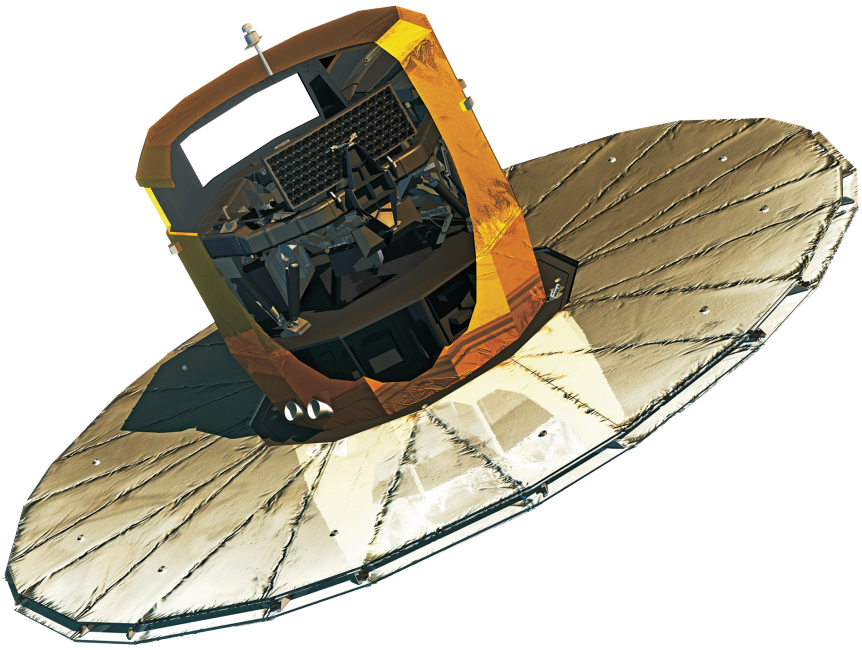


Figure 1.4 — General view of the Gaia spacecraft with the deployed sun shield. The cut-out in the thermal tent allows one to see: the payload and its two primary mirrors mounted on the mechanically and thermally stable torus (top), the service module (bottom). The satellite height is ~ 3 m and, once deployed, the sun shield diameter reaches 10 m. Illustration courtesy of EADS Astrium

scientific case are summarized by Lindegren et al. (2008) and Perryman et al. (2001). In the following I give a brief introduction to the key elements of the mission in order to support the comprehension of the work presented in this thesis.

Spacecraft
The Gaia spacecraft (Fig. 1.4) has been designed to carry out highly accurate absolute astrometric measurements, it thus embeds two telescopes separated by a basic angle of $106\overset{\circ}{.}5$. Incoming light will be imaged by 106 CCDs assembled in a single focal plane of 0.42×0.93 m: the largest ever flown in space. The focal plane and all the optical elements are mounted on a very steady and light-weight structure, a torus of 3 m diameter made out of silicon carbide (see Fig. 1.4 and 1.5). The payload (Fig. 1.5) is supported by a service module and the overall structure protected by a thermal tent. Payload and service module will be maintained thermally stable by a sun shield of 10 m diameter onto which solar panels are placed. The sun shield will only be deployed once in space so that the spacecraft fits in the upper stage of the Soyuz launch rocket (diameter of 4.1 m). The overall mass of the spacecraft is estimated to be $\sim$2 000 kg.

Focal plane
Gaia will not only carry out astrometric measurements but will also provide astronomers with photometric measurements for all observed stars and spectroscopic measurements for stars brighter than magnitude 17. The photometric measurements will be used to infer the astrophysical parameters of stars: luminosity, effective temperature, mass, age, and chemical composition. They are also necessary in the astrometric calibration, e.g. to correct for wavelength dependent effects on the stellar image. The spectroscopic measurements allow the derivation of the radial velocity of stars. The focal plane (Fig. 1.6) thus consists of several instruments: the Astrometric Field (AF), the Blue and Red Photometers (BP and RP, respectively), and the Radial Velocity Spectrograph (RVS). AF consists of 9 columns of 7 CCDs minus one CCD referred to as WFS1 (Wave-Front Sensor 1) dedicated to the initial alignment of the telescopes. For the AF instrument the light is not dispersed. The broad wavelength range set by the telescopes transmission and the detectors quantum efficiency (see Section 1.4) is 300–1000 nm and the associated magnitude scale is denoted G [2]. BP and RP consist of two low dispersion prisms that disperse the light before it reaches 2 columns of 7 photometric CCDs. BP covers the wavelength range 330–680 nm and RP covers 650–1050 nm. The RVS instrument is composed of 3 columns of 4 CCDs and relies on four prisms and a diffraction grating to disperse the incoming light in the narrow wavelength range of 847–874 nm. As explained in the following, Gaia is a spinning satellite, and as a consequence stellar images will not be stationary but will transit over the focal plane. The transit direction is called AL for Along-Scan; the perpendicular direction to the transit is called AC for Across-Scan. Stellar light will first encounter AF and then BP, RP, and RVS. As Gaia comprises two telescopes, light from two Fields Of View (FOV) is projected onto the same focal plane. To detect stars and discriminate between stars observed in one or the other FOV, 2 columns of 7 CCDs precede AF. This additional instrument is called Sky

2. The zero point is fixed by the convention that $G = V$ for an unreddened A0V star (Jordi et al. 2010; Perryman et al. 2001)

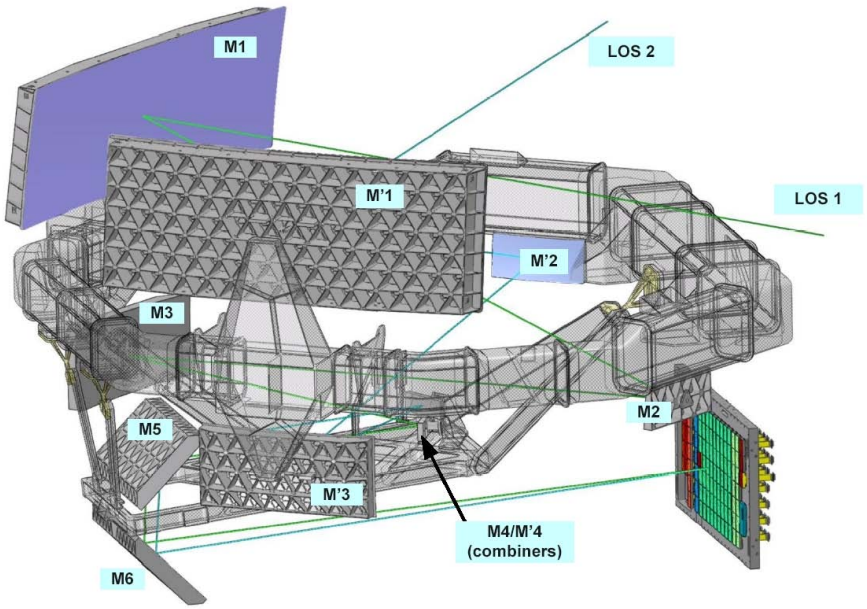


Figure 1.5 — Overview of the Gaia payload. One can recognize the multiple mirrors, including the primary mirrors M1 and M1′ (top left) of the two telescopes, the torus that supports the entire set of optical elements, and the focal plane (bottom right). The optical path for each telescope is also depicted by the two lines of sight LOS 1 and LOS 2. Illustration courtesy of EADS Astrium.

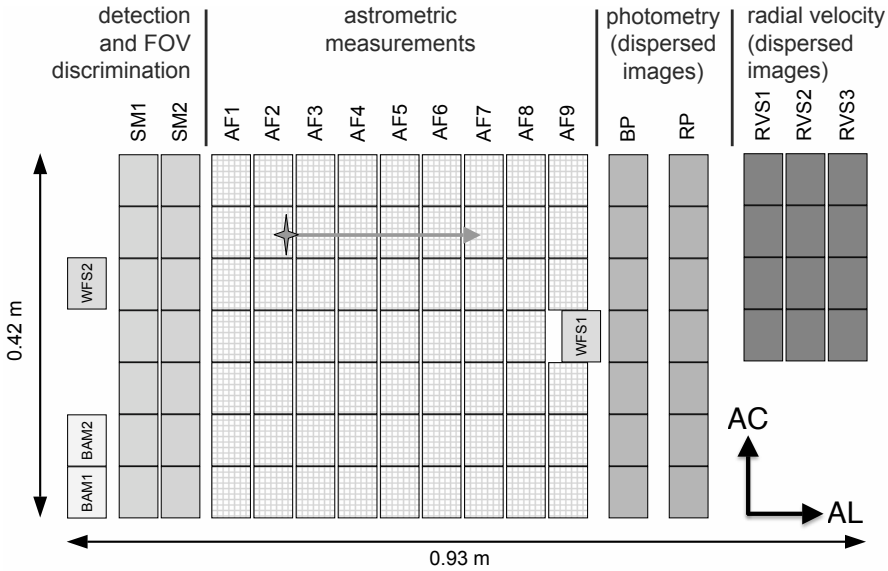


Figure 1.6 — Schematic layout of the CCDs in the focal plane of Gaia. The field of view for one telescope corresponds to 0.74 square degrees. Due to the satellite's spinning motion, a star enters the focal plane from the left and travels across it in the along-scan (AL) direction. The orthogonal direction is called the across-scan (AC) direction. All stars brighter than 20 mag are detected by one of the sky mappers (SM1 or SM2, depending on the field of view) and then tracked over the subsequent CCDs dedicated to astrometry (AF1–9), photometry (BP and RP), and radial-velocity determination (RVS1–3). In addition there are special CCDs for interferometric Basic-Angle Monitoring (BAM), and for the initial mirror alignment using Wavefront Sensors (WFS). Illustration courtesy of A. Short (ESA/ESTEC).

Mapper (SM). Only stars from FOV1 are detected by the first column of SM, SM1. And only stars from FOV2 are detected by SM2. Note that the AF1 CCDs take also part to this on-board detection of stars. There are three additional CCDs preceding SM1. One of them is WFS2 and the other two are dedicated to the constant monitoring of fluctuations in the angle between the two fields of view, they are referred to as Basic-Angle Monitor (BAM) 1 and 2.

Measurement principles
The reconstruction of the astrometric signal (see Fig. 1.2) of a star requires repeated observations of this star in the two fields of view. To collect astrometric measurements for all the stars brighter than $G = 20$, Gaia will constantly scan the sky. Hence, unlike pointing telescopes such as the HST that point at a certain location in the sky and stare for a certain amount of time to integrate light, Gaia will slowly rotate around the axis perpendicular to the two fields of view with a constant angular rate (see Fig. 1.7). As a consequence the integration time is fixed to 4.4 s per CCD independent of the brightness of the observed object. Gaia is thus a spinning satellite that scans the sky along great circles; one great circle will be completed every 6 hours. The spin axis will precess around the spacecraft-sun direction and change the orientation of the consecutive great circles such that Gaia covers the whole sky in about 6 months. Accounting for mission dead-time, on average each position in the sky will transit the combined field of view 72 times. The orbit of L2 around the sun and this combination of spinning motion and precession of the spin axis mainly determine the satellite scanning law (see Chapter 4 for a comprehensive description of the nominal scanning law). When a source (star, planet, galaxy etc.) transits over the focal plane, it is first detected by one of the SM. The detection algorithm is designed to be complete over the range $5.7 \leq G \leq 20$ and for any star in this magnitude range the image is then sampled and read-out by the successive CCDs. Sending to Earth the entire amount of data collected by the 106 CCDs at any moment is impossible due to a limited telemetry rate. For this reason only a truncated sampled image for each detected star is downlink, i.e. only a window of pixels is conserved: e.g., 12 pixels in the AC direction and 6 to 12 AL pixels depending on the observed object brightness for the AF instrument. Moreover, for most sources, the image will be binned in the AC direction, as only the AL information is important due to the nature of the astrometric measurement. As a result the CCD observation of a source by Gaia consists in most cases of a one-dimensional truncated image: a set of photoelectron counts.

Data processing
The Gaia data processing consists of transforming the satellite raw data into scientifically meaningful quantities such as the stellar astrometric and astrophysical parameters (for a comprehensive overview see Mignard & Drimmel 2007). The final outcome will be a catalogue containing all the information collected for each of the one billion observed stars, and similarly for other sources such as exoplanets, solar system bodies, galaxies, quasars etc. Although the final version of this catalogue is expected to be released by 2020, preliminary versions will be available most probably as soon as two years after the mission start. As already mentioned, DPAC is the organization respon-

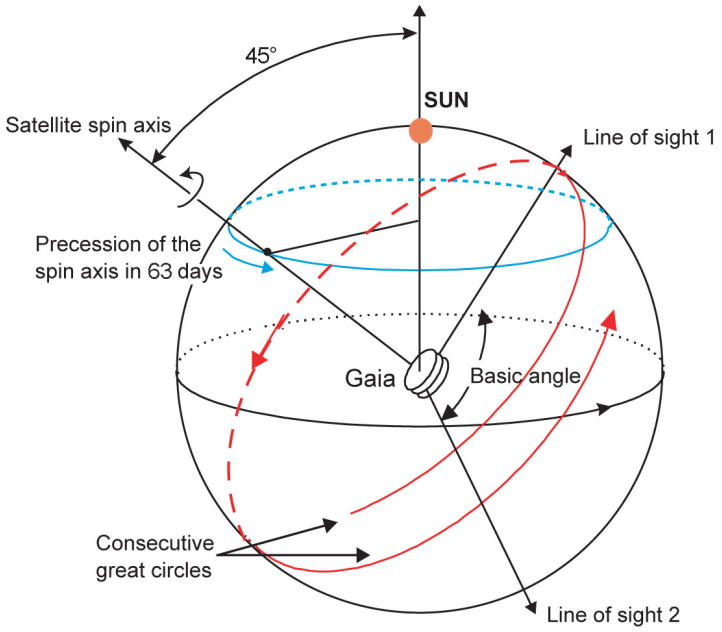


Figure 1.7 — Main principles of the Gaia scanning law. The spacecraft spin axis is perpendicular to the two fields of view and makes an angle of $45°$ with respect to the direction to the sun. The precession of the spin axis around the spacecraft-sun direction is completed in about 63 days. As a result of this scanning law, Gaia scans the sky along great circles, two consecutive great circles overlapping each other. Schematic courtesy of J. de Bruijne (ESA/ESTEC).

sible for this processing. It gathers more than 400 scientists spread over 24 countries. Due to the expected large amount of raw data (~ 100 TB) and the complexity of the task (highly interdependent data mixed in space and time), the Gaia data processing is one of the most challenging components of the Gaia mission. One of the main tasks of the Gaia data processing will be to reconstruct the astrometric signal of each observed star using the sets of photoelectron counts (or image profiles) acquired at different epochs. This task can be crudely summarized in two steps: (i) estimating from each image profile the image location on the CCD as well as the image flux and the background flux (Image Parameter Determination), and then (ii) using the image parameters for all observations to infer the astrometric parameters of all stars simultaneously. This second step will be performed by the Astrometric Global Iterative Solution, AGIS. These two steps are comprehensively explained in the Chapters 3 and 4 of this thesis. Apart from the core initial data treatment, the astrometric solution, and the associated and required photometric and spectroscopic processing, the Gaia data processing also includes the processing of specific astronomical sources (i.e. the identification and classification of binaries, exoplanets, solar system objects etc.), the determination of stellar astrophysical parameters, variability analysis, and ultimately the publication of the Gaia catalogue. Note that the data processing itself and its validation require considerable modelling efforts, a task also undertaken by DPAC. In order to prepare the Gaia data processing, the DPAC members elaborated models that range from the distribution of electrons in a Gaia CCD pixel to the distribution of stars in the Milky Way and its satellite galaxies.

Performance and science outcomes
After having been processed by DPAC, the Gaia measurements are expected to yield an unprecedented astrometric accuracy in the visible. The most recent performance predictions (parallax accuracies), produced just prior to the commencement of the final integration of the spacecraft and payload are as follows: 5–14 μas for $V \lesssim 12$, 9–26 μas for $V = 15$, and 100–330 μas for $V = 20$. For the brightest stars the range in accuracies reflects uncertainties in the calibration precision that can be achieved while for the fainter stars the ranges reflect the different colours of the stars, red stars being brighter in G for a given V magnitude. The astrometric parameters will be collected for one billion stars in the Milky Way, its satellites, and the nearest galaxies. Gaia is also expected to detect 5×10^5 quasars, 5×10^5 asteroids, and 10 000 new exoplanets. The main scientific areas that Gaia will tackle are (as listed by Mignard 2005): the mapping of the Milky Way, stellar physics, galactic kinematics and dynamics, distance scale (geometric to 10 kpc), age of the universe, dark matter, reference frame, planet detection, fundamental physics, solar physics, and solar system science.

As already touched upon in Section 1.1, this amazing performance and expected harvest of discoveries is largely enabled by the extensive use of extremely efficient light detectors called CCDs. In the next two sections, I give a comprehensive introduction to this type of detector and its operating principles and also introduce the problem of radiation damage to CCDs and its effect on space astronomical applications.

1.4 Charge-Coupled Device

A CCD is a light sensor: an electronic device that can detect and digitize light into images. It was invented in 1969 at the Bell Telephone Laboratories by Willard S. Boyle and George E. Smith. Its invention revolutionized astronomy and science, by allowing quantifiable, precise measurements of photons in a variety of wavelengths, as well as the world we live in by enabling digital photography. It is in acknowledgement of this achievement that Boyle and Smith were awarded the 2009 Nobel Prize in Physics.

Top-level description
Most generally, a CCD is a two-dimensional array of discrete elements called pixels. In its simplest form it consists of an imaging area, a serial register, and an output node (see Chapter 2 Fig. 2.1). Both the imaging area and serial register are built up with pixels. Each pixel corresponds to a MOS (Metal-Oxide-Semiconductor) structure, as explained in detail below. The imaging area consists of light sensitive pixels, and is the part of the CCD where the image is formed, i.e. photons are converted into charge packets. The image composed of charge packets located in different pixels is then transferred pixel row by pixel row to the serial register. In the serial register, the charge packets are read out; they are transferred to the output node where their charge is measured. The transfer direction in the imaging area is called parallel and in the serial register serial.

MOS Structure

The MOS structure is the basic building block of a CCD. It is a sandwich of three layers of materials of different nature (see Fig. 1.8): metal (conductor), oxide (insulator), and semiconductor. The metal layer is divided into sub-elements called electrodes (3 or 4 in most pixel architectures) connected to a circuit through which voltages can be applied. The semiconductor layer, referred to as the substrate, is made of silicon. It contains dopant atoms such as phosphorous or boron to create an excess of negative (electrons) or positive (holes) charge carriers. A substrate where the majority carriers are holes is called p-type, and n-type if the majority carriers are electrons. The oxide layer acts as an insulator between the metal and semiconductor layers. When applying a voltage to (biasing) an electrode, majority carriers are either attracted to the oxide-semiconductor interface or driven away from it depending on the sign of the voltage and the substrate type. The biasing of an electrode increases the region deeper in the substrate depleted

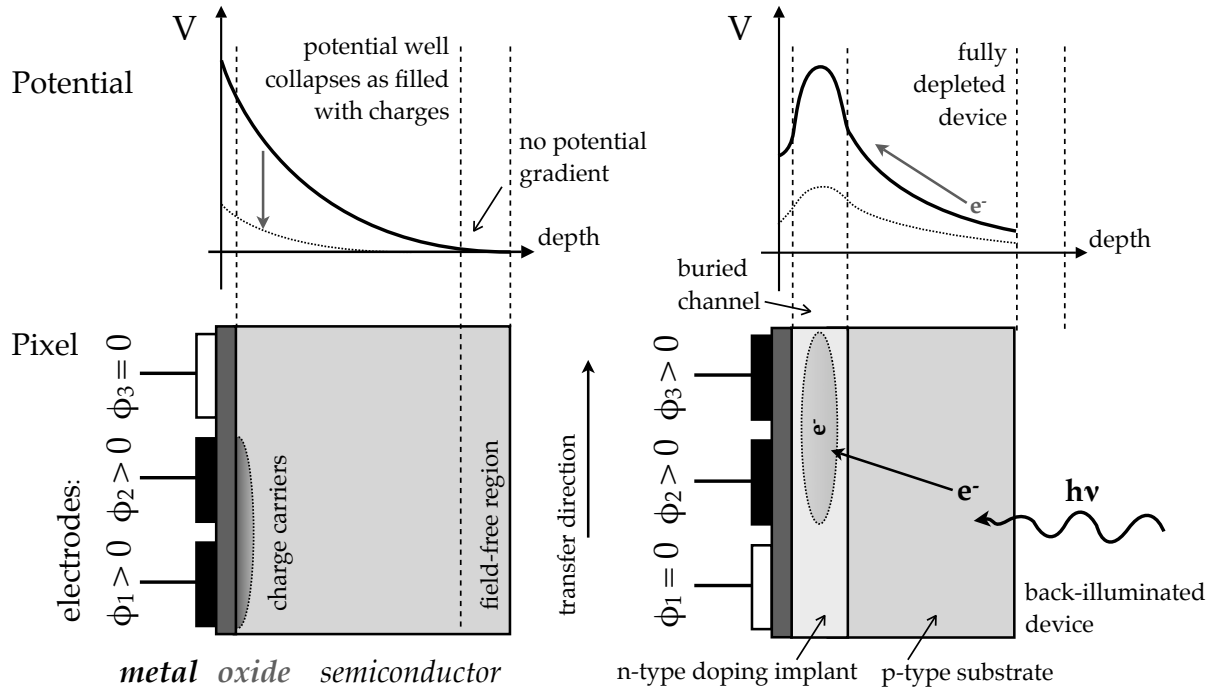


Figure 1.8 — Schematic of two different CCD pixel architectures (bottom panels) with integrating electrodes depicted in black and corresponding electrostatic potentials for these integrating electrodes (top panels). **Left:** Surface device: the CCD pixel is a simple MOS structure with a p-type substrate. A positive voltage is applied to the phase 1 and 2 electrodes (ϕ_1 and ϕ_2) and electrons are attracted to the oxide-semiconductor interface. The substrate is depleted of majority carriers (holes) and an electrostatic potential is formed. In this configuration the charge carriers (electrons) are transferred at the oxide-semiconductor interface where the maximum of the potential is located. **Right:** Buried channel device: n-type doping atoms are implanted beneath the oxide-semiconductor interface in a p-type substrate. The potential maximum is displaced and lies within the substrate. Electrons are transferred in the buried channel. In this configuration, the charge packets are located below electrodes (ϕ_2 and ϕ_3) biased high. The generation of an electron by an incoming photon is also depicted for a back-illuminated device. The photoelectron is driven along the potential following the positive gradient up to the maximum. The potential well gradually collapses as electrons are accumulated. This is depicted with the dotted lines in the top pannels.

from majority carriers which induces an electric field and a corresponding electrostatic potential. The potential can either act as a barrier or as a well for the signal carriers. Doping is used to shape this electrostatic potential in each pixel of a CCD.

Full well capacity
Once generated the charges follow the potential and fill the well. At Full Well Capacity (FWC), the potential well collapses and the pixel saturates. To prevent charge spilling into the neighbouring pixels, an anti-blooming drain can be added to evacuate charges when FWC is reached.

Buried channel
Electrons are transferred along the potential well maximum. Modern CCDs contain a thin layer of dopant atoms located below the oxide-semiconductor interface in the substrate. This layer causes a displacement of the potential well maximum away from the oxide-semiconductor interface into the substrate (Fig. 1.8). This layer is called a buried channel. Signal carriers are confined to this channel to avoid being trapped by interface states. Avoiding surface trapping considerably improves the Charge Transfer Efficiency (CTE). Most CCDs use electrons as signal carriers, they have an n-type buried channel in a p-type substrate (Fig. 1.8 right). These devices are referred as to n-channel CCDs.

CCD Operation Principles
The operation of a CCD can be divided into four fundamental tasks: (i) charge generation (from incoming light), (ii) charge collection, (iii) charge transfer, and (iv) charge measurement. The MOS structure plays an important role in the second and third tasks.

(i) Charge generation: to generate electrons from incoming photons, a CCD makes use of the photoelectric effect in semiconductor materials. Incoming photons with sufficient energy deplete the valence band of the semiconductor from one or more electrons which then end up in the conduction band (electron-hole pair creation). Both holes and electrons can be used as charge carriers, although electrons are chosen in most devices. Electrons generated this way are usually referred to as photoelectrons. The fraction of incident photons that produces a measurable charge is called the Quantum Efficiency (QE). Nowadays CCDs used for astronomy applications achieve nearly 100% QE over a wide range of wavelengths.

(ii) Charge collection: in principle, photoelectrons are free to move in the silicon lattice and to recombine with holes at any time. To avoid recombination and to conserve spatial information, photoelectrons are separated from holes and driven to the nearest pixel (buried channel) by an electric field following the electrostatic potential (see Fig. 1.8). Once an electron is located in the buried channel below an electrode, it is ready to be transferred to the neighbouring pixel as explained in the next item. A CCD for which incoming photons first encounter the metal layer is referred to as front-illuminated, while a CCD for which incoming photons first encounter the substrate is back-illuminated. Back-illuminated devices avoid the absorption of photoelectrons by the front-side structure; they have a higher QE. Depending on the substrate thickness and the value of the applied voltage, the silicon layer can contain a field-free region, i.e. a region for which there is no electrostatic potential gradient. This is undesirable,

as electrons are then free to diffuse and recombine in this region. The thickness of this region can be decreased by a back-thinning process and the voltage level tuned to fully deplete the substrate. To achieve nearly 100% QE, astronomical CCDs are back-illuminated, back-thinned, and also contain an anti-reflection coating.

(iii) Charge transfer: this task consists of the transfer of the generated signal carriers from pixel to pixel to the output node where the fourth and final task will be accomplished. As already mentioned, each pixel contains a set of electrodes; three-phase devices contain 3 electrodes per pixel in the imaging area while four-phase devices have 4 electrodes per pixel. In a three-phase device, every third electrode is connected to the same clock driver, either the phase 1, 2, or 3 clock. By applying voltage to phase 1 and 2, the electrodes no. 1 and 2 of every pixel will be biased-high; a higher electrostatic potential than in the neighbouring (phase 3 for which no voltage has been applied) electrode will immediately form in the buried channel. Photo-electrons are collected under phase 1 and 2 electrodes. And photoelectrons formerly sitting under the phase 3 electrodes are transferred under phase 1 and 2 electrodes of the neighbouring pixel. At the next step, only the phase 2 electrode remains biased-high and all electrons are collected and transferred under it. Then phase 2 and 3 are biased high and so on. See Chapter 2 Fig. 2.1 for a schematic of this process in a four-phase device. Note that the serial register generally contains three phases, two in more exceptional cases. In the imaging area, the columns of pixels are separated by potential barriers to prevent electrons spilling from one column to the next.

(iv) Charge measurement: charge packets in the serial register are transferred to the output node where charges are detected and measured. The output node is composed of a floating diffusion node and an output amplifier. Charges are stored in the floating diffusion node and the output amplifier generates a voltage proportional to the number of charges in a packet. The resulting voltage is then digitized. All the charges are transiting through a single amplifier, this means that it introduces as little noise as possible. The noise level is now reaching less than 1 or 2 electrons rms allowing for almost single photoelectron counting if all the other CCD noise sources are reduced to zero.

Charge Transfer Efficiency

Due to the high number of pixels and corresponding transfer steps, it is critical that all charges are transferred at each step. The CTE is the fraction of transferred charges from one pixel to its neighbour, the opposite is the Charge Transfer Inefficiency: CTI $=$ $1 -$ CTE. CTI is the fraction of charges lost in one transfer. In the first manufactured CCD the CTE was 99%. Although it sounds reasonable, it means that 1% of the signal was lost at each transfer. For a 4000 by 4000 pixel device, it would not be long before all the signal is lost. Modern devices can achieve a CTE of 99.99999%. One of the conditions to achieve such CTE is the use of a buried channel.

The Gaia CCDs

In the following I only summarize the main characteristics and peculiarities of the Gaia CCDs, for a complete description see Short et al. (2005). The Gaia CCDs are back-illuminated, and full frame (i.e. their entire surface contributes to the light detection). They are custom made by e2v technologies and referenced as CCD91-72. The CCD image area is four-phase and contains 4500 $\times$ 1966 pixels (parallel $\times$ serial), each 10 μm

$\times$ 30 μm in size. Among the 4500 pixel rows only 4494 are light-sensitive: six rows are blocked by an aluminium shield. Each pixel of the image area contains an n-type buried channel, a lateral anti-blooming drain, and an extra doping implant on top of the buried channel a so-called supplementary buried channel or notch (a detailed description of the Gaia CCD pixel architecture is provided in Chapter 6). A Gaia CCD pixel reaches FWC at a signal level of more than 190 000 electrons. The serial register is two-phase, it consists of a single row of 1966 pixels. The Gaia CCD contains a single high performance output node; the readout noise is measured to be less than 4 electrons rms. The first CCD row contains a charge injection structure (diode and gate) allowing for the injection of artificial electrons in the CCD. The nominal temperature of operation of 163 K has been selected to minimize the dark current and the radiation damage effects. However due to design considerations, the average focal plane temperature is expected to be slightly warmer (169 K). Gaia is a spinning satellite, this means that the star projections on the focal plane are not stationary. As a consequence, the Gaia CCD will be operated in Time-Delayed Integration (TDI) mode. In this mode the CCD is constantly read out and the satellite scanning rate (and induced light source motion) has been synchronized with the charge transfer period, so that the charge profile continues to build up as the image travels across the CCD avoiding as much as possible image smearing. The charge transfer period is 0.9828 ms and the integration time 4.4 s. A Gaia CCD also comprises 12 so-called TDI gates, that enable the reduction of the integration time and prevent saturation of the images of bright sources. The QE of the Gaia devices has been optimized depending on which instrument the CCD would be assigned to. There are three types of devices: AF (broad-band), BP (blue-enhanced), and RP (red-enhanced). AF CCDs are also used for SM (and WFS) and RP CCDs for the RVS instrument (and BAM). None of the Gaia CCDs are fully depleted devices; the AF and BP CCDs have a thickness of 16 μm and a field-free region of about 4 μm. The RP CCDs are thicker (40 μm) so that their QE is higher towards redder wavelengths, but have a thinner field-free region of 2 μm. All the Gaia CCDs have now been manufactured and Fig. 1.9 shows the fully integrated Gaia CCD mosaic onto its support structure. It will be the largest CCD focal plane array ever flown in space.

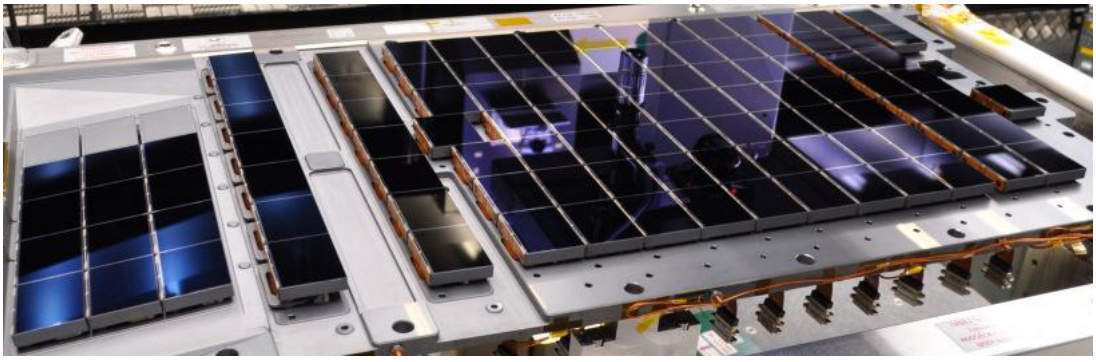


Figure 1.9 — The 106 Gaia CCDs integrated onto the CCD support structure. The columns of CCDs corresponding to the different instruments of the focal plane can be recognized: (from left to right) RVS (first three columns), RP, BP, AF9–1, SM2, SM1, and in the last column BAM1, BAM2 and WFS. A single CCD measures 5.9×4.5 cm, and the entire structure 1×0.5 m. Image credit: EADS Astrium.

1.5 Radiation damage

In the context of this thesis, radiation damage is the degradation of the performance of embedded electronic devices on-board satellites caused by energetic particles in outer space. Energetic particles in space such as the ones constantly released by our star, the Sun, can severely affect the functioning of satellites. This was realized very early on in the history of space conquest; the first loss of a satellite directly attributed to the effect of radiation occurred in 1967, only ten years after the first successful orbit of an artificial satellite, Sputnik. Radiation is thus a common issue in space science, and a radiation budget analysis is one of the mandatory (and routine) steps in the preparation of a satellite. Nowadays every satellite employs countermeasures such as the use of radiation-hard materials and shielding. Regarding the electronics, most components must be certified radiation-hard and include redundant circuits to avoid a single-point failure. This means that for space missions, such as Gaia, radiation is often a challenge because they require novel designs and custom-made devices, in particular to equip their payload. Astronomy missions also often tend to have large openings to collect as much light as possible, which render shielding difficult. Radiation damage was recognized as a one of the major threats to the Gaia science performance early on in the mission preparation due to a new radiation environment at L2, a launch close to the maximum activity of the Sun, large telescope apertures, a severe weight constraint, the extensive use of large CCDs, and very stringent requirements on the image quality. This concern led to a major effort from the Gaia community to mitigate the radiation damage threat to which the work in this thesis represents a significant contribution. In the following, I provide an overview of the radiation environment of Gaia, the type of interactions between energetic particles and matter, as well as the mechanisms that give rise to damage in electronic components. I also give a brief summary of radiation-induced CTI in CCDs and the expected effects on the Gaia measurements.

Space Environment
Although it is nearly impossible to reach space vacuum by artificial means on Earth, space is far from being empty especially in the vicinity of stars such as the Sun. Interplanetary space is composed of neutral particles, plasmas, cosmic rays, micrometeoroids, space debris, and most importantly radiation. By radiation one usually means the entire light spectrum (i.e. radio to gamma-ray photons) and energetic subatomic particles such as electrons, protons, and neutrons. Photons can also be harmful but subatomic particles are the predominant cause of radiation damage, and the solar wind is the main source of these particles. The solar wind (illustrated on the cover of this thesis) carries electrons, protons, and neutrons, with energies ranging from few eV to several MeV, expelled by the Sun during magnetic events at its surface (i.e. the upper layer of its atmosphere), such as coronal mass ejections or flares. In the Earth's vicinity, where the vast majority of the satellites are orbiting, the radiation environment is well understood and predictive models of radiation fluence are accurate. In addition the Earth's magnetic field acts as a shield against radiation and traps particles into belts. However L2 is a rather uncommon location for satellite operation, ESA sent its two first missions to L2, Herschel and Planck, only recently in spring 2009. L2 is located

1.5 million km away from the Earth on the Sun-Earth axis, and thus resides well outside the influence of the Earth's magnetosphere; the outer radiation belt extends only up to 100 000 km away from Earth. At L2 the radiation environment is expected to be dominated by solar wind protons. The reference model used to predict particle fluences at this location is the Interplanetary Proton Fluence Model — JPL (Jet Propulsion Laboratory) 1991 by Feynman et al. (1993). The solar particle events, which originate in magnetic events, are governed by the solar activity cycle usually monitored by Sun spot counting. This cycle has an 11 year period with half a period of intense activity and the other half of quiescent activity. A new solar cycle started in 2010, with a maximum of activity expected to be reached during 2013. Gaia is due for launch in 2013, and as a consequence Gaia is expected to experience most damage during its first year of operation. According to the JPL 1991 model, taking into account the satellite design, and assuming 4 years of operation during the solar maximum (and one year during minimum), the average accumulated radiation dose received by a CCD of the AF instrument is predicted to be $\sim 3\times10^9$ (10 MeV equivalent) protons cm^{-2}. This means that the expected end-of-mission damage will be equivalent to 19 10 MeV-protons per second bombarding every square centimetre of a Gaia CCD during 5 years.

Energy transfer and damage mechanisms
When an energetic particle collides with a satellite part, it transfers a part of or its entire energy to the target material through ionizing and non-ionizing processes. The energy of the incident particle lost by ionization of the target material is quantified by the TID (Total Ionizing Dose), and the energy lost by non-ionizing processes is quantified through the NIEL (Non-Ionizing Energy Loss). TID results in temporary effects such as SEUs (Single Event Upsets) where the state of an electronic device is changed by the incoming particle. TID also results in long-term and cumulative effects such as voltage drift, increasing dark current in light detectors etc. Those effects are important but in the case of Gaia and in particular the Gaia CCDs, the effects induced by non-ionizing processes are the most threatening. NIEL essentially causes long-term cumulative effects by creating defects in the target material through displacement damage, which is the displacement of atoms in the target material. This displacement is due to the recoil of a target atom induced by elastic scattering of the incoming particle; the scattering results form electrostatic or nuclear interaction depending on the energy of the incident particle. The recoiling target atom is usually referred to as the Primary Knock on Atom (PKA). A PKA can cause a cascade of atom displacements creating defect clusters at different locations in the target material. In the case of a CCD the target material is crystalline silicon; displacement damage results in the creation of interstitial atom - vacancy pairs (as illustrated in Fig. 1.10) provided the energy transferred to the displaced atom is higher than a certain value (2.3 eV in Si). Vacancies can diffuse in the silicon lattice and gather with other vacancies or impurities: doping atoms (phosphorus, boron) or polluting atoms present during the CCD manufacturing process (oxygen and carbon etc.). These vacancy-vacancy and vacancy-impurity complexes are referred to as bulk traps, and complexes of different nature are called trap species.

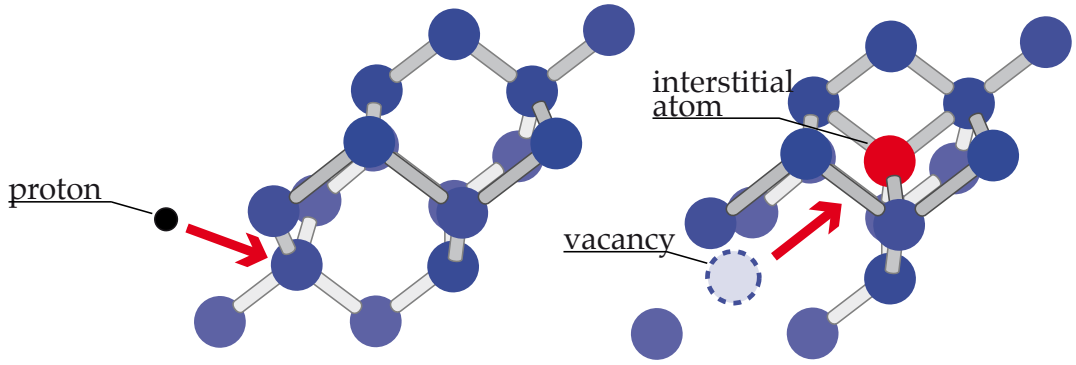


Figure 1.10 — Displacement damage in the silicon lattice. An energetic proton collides with a silicon atom (left). The energy transferred by the proton to the target atom is enough to displace it into an interstitial location (requiring more potential energy than the normal lattice location of the Si atom), the lack of atom at the lattice location is called a vacancy: an interstitial atom-vacancy pair has been created (right). Vacancies can diffuse in the silicon lattice to bind with impurities and create bulk traps. Illustration adapted from G. Lucas and L. Pizzagali (LMP).

CTI induced by bulk traps

Bulk traps decrease the charge transfer efficiency (increase CTI) in CCDs by introducing energy levels in the semiconductor band-gap. These energy levels (simply referred to as traps) can cause the stochastic capture and release of signal carriers. Following the Shockley-Read-Hall formalism (Shockley & Read 1952; Hall 1952) the stochastic capture and release of a charge by a trap can be considered as a decay process with a characteristic capture and release time constant, respectively τ_c and τ_r. In this way, probabilities of capture and release can be computed; they depend on the time constant and the interaction time between charge and trap. According to the Shockley-Read-Hall formalism, statistically 67% of the traps have released their captured charge after an interaction time equal to the release time constant. The capture and release time constants essentially depend on the temperature, the trap capture cross-section, the energy difference between the level and the conduction band if electrons are the signal carriers (valence band in the case of holes), and in the case of the capture, the signal carrier density in the bulk trap vicinity. See Chapter 2 for a detailed derivation of these probabilities. Note that bulk traps are already present in a CCD without radiation due to pollution during the CCD manufacturing. CTI caused by the manufacturing traps is called native CTI. As already mentioned in Section 1.4, CTE in non-irradiated modern CCDs for scientific application reaches 99.99999%. Native CTI is thus lower than $10^{-5}\%$, this means that the density of manufacturing traps is extremely low.

CTI effects

Different trap species have different energy levels and capture cross-sections. As a result they have different capture and release time constants and different effects on the transferred signal (charge packets) in the CCD. These effects depend on the temperature of operation and the pixel to pixel transfer period T, once a charge is captured we can distinguish between three cases: (i) $\tau_r \ll T$: the captured charge is released before the charge packet transfer, there is no net effect on the signal. (ii) $\tau_r \gg T$ the

charge release occurs after the signal of interest has been readout, the signal may be distorted and there is a net charge loss which decreases the signal-to-noise ratio. (iii) $\tau_r \sim T$ the charge release occurs before the signal of interest is completely readout, the signal is strongly distorted and a characteristic charge trail is formed (see Fig. 1.11). As shall be seen all along this thesis, and especially in the Chapter 3 and 4, both the signal distortion and decrease in signal-to-noise ratio are expected to affect considerably the performance of Gaia if not properly taken into account. Fig. 1.11 (right) illustrates such distortion. The transfer rate in the CCD image area and serial register are different, different trap species will thus be of importance in these two CCD regions. CTI occurring in the image area is referred to as parallel CTI, while CTI in the serial register is called serial CTI. The distortion always occurs along the transfer direction. In the case of Gaia, since the signal in the AL direction is the most important for astrometry (cf. Section 1.3) it is the parallel CTI that will mostly degrade the mission performance. The native parallel CTI is very low and radiation damage will drastically increase it. However it is interesting to note that experimental tests carried out on Gaia irradiated and non-irradiated devices have shown that serial CTI is natively high due to the very high rate of serial transfer selected for the Gaia CCD and radiation damage will only marginally increase it. Based on experimental tests, the nominal operating temperature of the Gaia CCD has been selected such that CTI is minimized and especially the strong distortion of the signal in the AL direction avoided for a transfer period of 0.9828 ms (in the CCD imaging area). In the next section, I describe the strategy selected to mitigate CTI effects for the Gaia mission.

1.6 The Gaia CTI mitigation strategy

One can distinguish two ways of mitigating the CTI effects on the Gaia measurements and CCD-based measurements in general: first, avoiding the trapping of the signal carriers by the use of hardware solutions, and second, taking into account the CTI effects in the (on-ground) data processing by the use of software solutions. The Gaia CTI mitigation strategy relies on both approaches. The main components of this strategy are: the periodic injection of charges, a SBC present in each pixel of the Gaia CCD imaging area, and a novel CTI mitigation procedure applied at the image processing level. The elaboration of this strategy results from extensive and detailed test campaigns on irradiated Gaia CCDs as well as an unprecedented modelling effort to simulate the CTI effects at different levels of detail (see Fig. 1.12).

Hardware CTI-countermeasures
Hardware solutions to the CTI problem comprise: (i) shielding, (ii) CCD operation and (iii) CCD design optimization, (iv) injection of sacrificial charges, (v) constant illumination of the CCD by a diffuse optical background, and (vi) annealing. Shielding consists of the placement of extra layers of material around sensitive parts of the satellite payload to stop energetic particles or attenuate their energy and hence avoid displacement damage. The efficiency of this solution on-board satellites is often disputed as strict weight constraints imposed by the launch generally restrict the shield thickness. This limits the attenuation efficiency and enables secondary particles created inside the

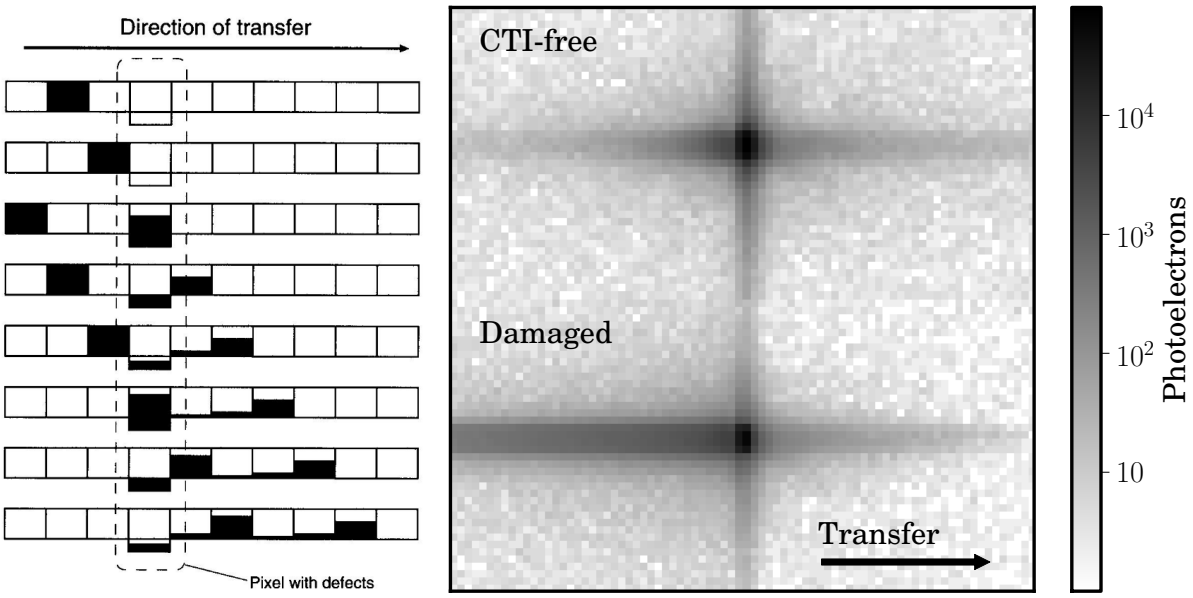


Figure 1.11 — Left: CCD Charge Transfer Inefficiency induced by bulk traps. In the schematic, rows correspond to a pixel row at different step of the charge transfer process. From top to to bottom, a charge packet is transferred from the second to ninth pixel. The fourth pixel contains bulk traps that capture the signal carriers, which are released at each subsequent transfer. This leads to the charge trail characteristic of devices operated with a transfer period similar to the release time constant of the trap species. Schematic courtesy of J. Walder (Lancaster University). **Right:** CTI-induced distortion of the shape of a Gaia-like stellar image. The image distortion was simulated using the physical Monte Carlo model described in Chapter 2. The illustration shows the CTI-free and damaged transits of a $G = 14$ star over a Gaia CCD. The trap density has been artificially increased to clearly bring out the CTI-induced distortion and the resulting photo-centre shift and charge trail.

shielding material to significantly contribute to the radiation damage. Gaia disposes of a relatively low level of shielding due to the two very large apertures of its telescopes and stringent weight constraints; nevertheless shielding has been reinforced behind the CCD mosaic camera.

The CCD operation can be optimized to mitigate the CTI effects; the main parameters tuned to avoid trapping or to minimized the effect of the most abundant trap species are the temperature at which the CCD is operated, the transfer rate, and the clocking and readout scheme.

The CCD design and in particular the pixel architecture can also be optimized, for instance by the use of extra doping implants to shape the electron density distribution in the pixel. This ensures that an electron packet occupies the smallest volume possible at a particular signal level to minimize its potential interaction with bulk traps. The SBC is an example of such feature in the Gaia CCDs (cf. Chapter 6).

The injection of sacrificial charges in the CCD aims at filling the traps to avoid the capture of the signal charges. A charge injection structure, a diode and a gate, is placed before the first pixel column of the Gaia CCD. Charge injections can be performed before any signal of particular interest or periodically to constantly fill the traps with a release time constant greater than the charge injection period. After a charge injection, the trapped charges are released and create a charge injection trail. The duration

and the level of the charge injection (the number of charges injected per pixel) must be carefully optimized to avoid a significant increase in the noise induced by the released charges. Charge injections have been retained as one of the Gaia CTI hardware countermeasures.
By constantly illuminating the CCD with a uniform source of photons, it is possible to generate a low level of background charges that will fill a fraction of the traps. This is because the interaction time between the background generated charges and the traps is infinite. This mean that even the very low density of background charge can decrease the CCD CTI (see Chapter 5). The diffuse optical background option (using artificial light sources) was discarded for Gaia because of significant noise increase and only marginal mitigating effect.
Temperature annealing physically removes the CCD-lattice defects from which the charge traps originate by heating the CCD up to a certain temperature ($T > 20°C$). Because of the extreme requirements on the instrument stability for Gaia and because the efficiency of this method is disputed, annealing was not investigated for Gaia. The study of the impact of radiation damage on the Gaia final astrometric accuracy presented in the Chapter 3 and 4 of this thesis enabled the detailed assessment of the efficiency of the selected hardware CTI-countermeasures.

Software CTI-countermeasures
Software solutions are necessary to address residual CTI as hardware solutions cannot totally prevent trapping. The main aim is hence to avoid the CTI-induced systematic errors and decrease in the precision of the image-based measurements and inferred scientific quantities. Software solutions can be applied at various stages of the data processing chain. Chapter 3 provides a review of the range of solutions available in the literature. Most methods have been developed for photometric and spectroscopic measurements carried out in the optical or at X-ray wavelengths, but very rarely for astrometric measurements. Among the most successful methods, the HST CTI mitigation scheme (see e.g., Bristow 2003; Massey et al. 2010) relies on the correction of the raw pixel data to obtain CTI-free measurements. The rest of the data processing chain is then performed with the altered CCD measurements. Because of the required accuracy in the image location estimation, the lack of full frame data, the operation of the CCD in TDI-mode, and several other reasons (see Chapter 3), HST-like methods cannot be applied in the case of the Gaia measurements. As a result a novel software CTI-mitigation procedure has been specifically designed for the Gaia data processing to be employed at the image processing level. This procedure relies on a forward modelling approach that enables the estimation of the true (CTI-free) image parameters from a damaged observation and avoids any direct correction of the raw data. In this procedure each observation is ultimately compared to a modelled charge profile for which the distortion has been simulated through a fast analytical CTI model. Chapter 3 provides a detailed description of this approach and its multiple advantages as well as a test against synthetic data to demonstrate its potential accuracy. Chapter 5 discusses the applicability of this approach on real data and in provides an assessment of the performance of the fast analytical CTI model required by the software solution described in Chapter 3.

Experimental studies

A series of experimental tests carried out on irradiated Gaia CCDs by the industrial partners in the project served as basis from which to elaborate the Gaia CTI mitigation strategy. The tests aimed at the evaluation of the amplitude and the trends in the CTI effects to be expected in the Gaia operating conditions and the identification, characterization, and optimization of potential hardware CTI countermeasures.

The first series of tests (Hopkinson et al. 2005) was performed by Sira Technology Ltd and focused on the characterization of the Gaia CCDs, the identification of the trap species involved, and the optimization of the CCD operation temperature. Later, Surrey Satellite Technology Limited (SSTL, formerly Sira) investigated the potential difference in the CCD radiation damage characteristics resulting from an irradiation performed at room temperature or at operating temperature (163 K). Hopkinson (2008) concluded that the results obtained for CCDs irradiated at room temperature should be adequate for Gaia performance predictions within the usual experimental uncertainties.

Up to now EADS Astrium has performed four test campaigns on AF and RP CCDs irradiated at room temperature with radiation doses of 2 and 4×10^9 protons cm^{-2} (10 MeV equivalent). A fifth campaign is in preparation. For these tests, Astrium built a special test bench and experimental setup to reproduce the Gaia operating conditions (e.g., temperature, low level of background light) and the motion of the light source to ensure the operation of the tested CCD in TDI mode (see Pasquier 2011). Diffraction limited (AF) and dispersed light (BP, RP, RVS) measurements were reproduced by the use of a variety of custom-made optical masks placed after a light source (LED) and before a rectangular aperture mimicking one of the Gaia openings and a lens. The analysis of these tests (e.g., Georges 2008; Brown 2009b) provided the Gaia community with a detailed characterization of the CTI effects on all the Gaia measurements at different signal levels, a large data set against which models can be verified, and several conclusions of importance regarding the CTI mitigation strategy of Gaia. Some of the conclusion are for instance the confirmation of a nominal operation temperature at 163 K, the use of periodic charge injections (with a 1 s period, 5 pixel duration, and $\sim$17 000 electron level), the rejection of the artificial diffuse optical background, and the increase in AC size of the telemetry window for the dispersed light measurements to mitigate the signal-to-noise ratio decrease induced by native and radiation-induced serial CTI. Also a large variety of anomalies in the functioning of the Gaia CCDs, in particular related to the associated driving and reading electronics (the proximity electronic module) were identified during these studies. Some of these anomalies have been solved by changes in the hardware and the operation of the CCDs and their associated electronics. However some of the anomalies remain and will be addressed in the Gaia data processing.

Modelling efforts

Experimental studies alone do not suffice and modelling efforts are needed to deepen and test our current understanding of the CTI effects on the Gaia measurements, to generate a large variety of synthetic data to support the data processing preparations, and to perform a detailed evaluation of the CTI impact on the Gaia performances.

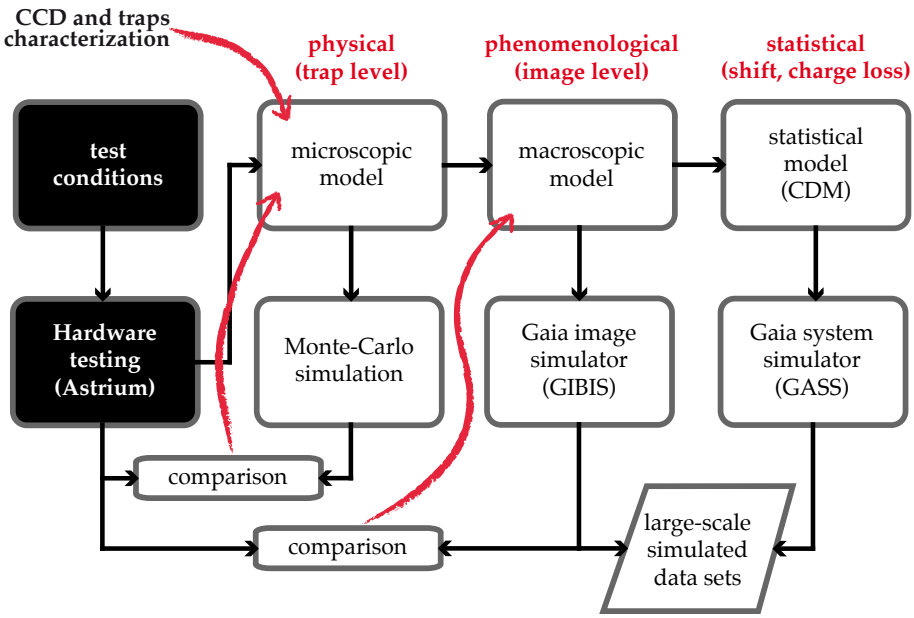


Figure 1.12 — The elaboration and implementation of the Gaia CTI mitigation strategy requires both experimental studies and modelling efforts at different levels of detail. This diagram outlines the purpose and represents the interactions between the experimental tests performed on Gaia irradiated CCDs and the various models of CTI-effects. The work presented in this thesis is relevant to all the different stages of the modelling studies. The models developed as part of this work were used to acquire a better understanding of the experimental data. Diagram courtesy of L. Lindegren (Lund Observatory).

Moreover, as discussed above, the calibration of the CTI effects at the image processing level necessitates a very fast and accurate model of the image distortion. In order to reach these objectives, several levels of CTI effects modelling (see Fig. 1.12) have been developed and in particular in the context of my thesis research:

– *Microscopic* (physical) models that simulate the trapping process as accurately as possible, based on physical principles, and simulate the charge transfer at the pixel or sub-pixel level. These models should be able to reproduce the experimental tests, and support the improvement of the basic understanding of the CTI effects. They can be used for example to generate synthetic data in order to verify the higher levels of modelling and to assess the Gaia image location accuracy (Chapter 3). Such a model is presented in the Chapter 2 of this thesis.
– *Macroscopic* (phenomenological) models capable of reproducing the mean location bias and charge loss of the microscopic models as a function of fewer parameters and in a very efficient way. Such fast and accurate analytical models are required by the Gaia data processing, and the final astrometric accuracy of Gaia depends on their performance. Chapter 5 provides a test of such a model against experimental data.
– *Statistical* models are able to describe the systematic errors and additional noise induced by the CTI effects. This type of model is used to perturb large-scale telemetry simulations used to test the Gaia global data processing and to enable the study of the propagation of the CTI-induced errors in the data processing chain (as performed in Chapter 4). Chapter 4 provides a description of such a model and applies it to the astrometric solution of Gaia.

1.7 This thesis

This thesis presents in Chapter 2 a description and verification of the most detailed model to date of radiation damage on astronomical CCDs. Chapters 3 and 4 present the first detailed assessment of the impact of CCD radiation damage on Gaia astrometry as well as proposed solutions to mitigate the CTI threat. In Chapter 5 the test of the CTI mitigation procedure at the image processing level against experimental test data is described, and in Chapter 6 on one of the Gaia CCD hardware CTI mitigation tools, the supplementary buried channel, is investigated. In the following I give an overview of the contents of each chapter and the associated main conclusions.

Electrode level Monte Carlo model of radiation damage effects on astronomical CCDs
Due to numerous particularities inherent to the mission operation, payload, and measurements, Gaia necessitated a deeper understanding of CCD radiation damage and in particular of the effects of CTI on the acquired images. This understanding was needed in order to elaborate a strategy and new dedicated solutions to counter the effects of CTI on the final astrometric accuracy of Gaia, as well as on the photometric and spectroscopic measurements. In this context both experimental and theoretical work have been and continue to be carried out (see Section 1.6). In **Chapter 2**, I describe a physical Monte Carlo model that simulates CTI effects induced by radiation damage in astronomical CCDs at the pixel electrode level. This model implements a new approach to both the charge density distribution within a pixel and the charge capture and release probabilities, which allows the reproduction of CTI effects on a variety of measurements (astrometric and spectroscopic) for a large signal level range, and in particular for signals of the order of a few electrons. The main lesson learned from this chapter is that to be successful in modelling the CTI effects at very faint signal levels, no detail should be neglected: the simulations should be as realistic as possible, down to the transfer of electrons at the electrode level and the simulation of each individual trap. Although developed to cope with the very demanding case of Gaia, the model was set up to be as general and as flexible as possible. It can be used to simulate any kind of measurements performed in different CCD operating modes, with different clocking schemes, and pixel architectures. It is for instance perfectly suitable to evaluating the impact of radiation damage on the performance budgets of future missions pushing the CCD to its limits, such as ESA's Euclid and Plato missions. Note that the model is readily available on line as part of the CEMGA Java package (Prod'homme 2011) that I developed to support the research described in this thesis. CEMGA stands for CTI Effects Models for Gaia, it is a platform that hosts the different models developed within the Gaia framework. The platform aims to offer a rigorous environment for testing, verifying and comparing the hosted models. It also comes with a set of pre-configured experiments (such as First Pixel Response) and various tools to analyze the simulations (e.g., I/O management methods, image location and minimization algorithms, visualizing routines) and compare them to experimental test data.

The impact of CCD radiation damage on Gaia astrometry

In **Chapters 3 and 4** I present two studies aimed at characterizing and quantifying the impact of CCD radiation damage on the final astrometric accuracy of Gaia. The Gaia astrometric requirements set a very stringent constraint on the accuracy of the estimation of the stellar image location on the CCD for each observation. For instance a final parallax standard error of 30 μas for a G2V type star of magnitude 15 translates into the requirement that the image location error is less than 0.3 mas or 0.005 pixels for each CCD observation. In **Chapter 3**, the image location estimation in the presence of radiation damage is investigated. For this purpose CEMGA and the model presented in Chapter 2 were used to generate a large data set of ($\sim$40 000) synthetic Gaia-like stellar images obtained from a CTI-free and a damaged CCD. We first compute what is the theoretical limit to the image location estimation for CTI-free and damaged Gaia-like images, and find that CTI introduces an intrinsic loss of accuracy independent of any image location estimator. This is due to the decrease in signal-to-noise ratio associated to the CTI-induced charge loss. This intrinsic loss can only be prevented by avoiding the trapping of the signal carriers, hardware mitigation thus plays a very important role. Taking into account all the Gaia hardware CTI countermeasures, we find that this intrinsic loss of accuracy can reach up to 6% for the end of mission radiation dose. Then we apply the Gaia image location estimation to damaged images and evaluate the location bias (systematic error) induced by the image distortion if CTI is not properly taken into account in the image processing. The bias is considerable, up to 0.05 pixels or 3 mas for realistic Gaia operating conditions and the end-of-mission radiation dose. This shows that CTI software mitigation must be applied. However most of the CTI mitigation procedures available in the literature have been developed to correct photometric measurements for the HST or X-ray CCD-based measurements that do not require the same level of accuracy in terms of image location. We thus present a new CTI mitigation procedure that relies on the forward modelling of the image distortion. It offers several advantages among which the capability to handle any kind of complex observed scenes (such as multiple overlapping stellar images). It nevertheless requires the use of a very fast and accurate analytical model of the image distortion, a so-called Charge Distortion Model (CDM). We show that in principle the forward modelling approach enables the complete recovery of the CTI-induced bias, and a recovery of the image location precision down to the theoretical limit in the presence of radiation damage. We test this new approach with the current best CDM candidate for implementation in the Gaia data processing, and show that the bias is reduced by a factor ten in close to ideal conditions regarding the calibration of such a model. Note that **Chapter 3** also offers a review of the most recent methods to correct for CTI in the literature.

Chapter 4 focuses on the effect of the image location errors on the astrometric solution which converts the Gaia image location measurements for all (single) stars into a set of astrometric parameters for each of these stars. For the first time, not only the increased random errors but also the image location bias is rigorously propagated through a realistic astrometric solution for 1 million stars. We found that while the mean of the CTI-induced bias is absorbed in the solution, the variation with magnitude and other factors (e.g., the illumination history) is propagated to the astrometric parameters. This

shows once more that it is absolutely mandatory to calibrate for CTI at some stage to be able to recover the final astrometric accuracy of Gaia. It also means that since the bias is propagated in the astrometric solution, it is possible to use the solution residuals to support the CTI calibration at the image processing level. We thus test the use of these residuals and show that by combining a forward modelling approach at the image processing level with the analysis of astrometric solution residuals, we are able to recover a virtually bias-free estimation of the astrometric parameters from bright to faint magnitudes. This study enabled the identification of all the mechanisms that can contribute the mitigation of the CTI effects on the final astrometric accuracy of Gaia. Hence we give a review of all these mechanisms and their respective efficiency (when possible) including the hardware and software mitigation tools. Taking all the CTI mitigation countermeasures into account, we then demonstrate that the overall astrometric accuracy of Gaia can be preserved to within 10% from the CTI-free case, from bright to faint magnitudes.

Stress-testing a fast analytical Charge Transfer Inefficiency model
As mentioned before the calibration of CTI at the image processing level through a forward modelling approach requires the use of a CDM, a fast analytical charge transfer inefficiency model. The final astrometric accuracy of Gaia is conditioned on the capability of a CDM to reproduce observations affected by CTI. **Chapter 5** presents a study that aims at evaluating the performance of the current best CDM candidate when reproducing experimental test data representative of the future Gaia observations. The main conclusion of this study is that the level of agreement obtained using this model is enough to reduced the CTI-induced image location bias by a factor ten, and thus enables the potential recovery of the required final astrometric accuracy but the calibration of such a model is a complicated enterprise and potentially problematic. We thus study different calibration schemes and identify the most practical to be used in the data processing, we also propose modifications to the current best CDM candidate that would ease its calibration.

Digging supplementary buried channels: Investigating the notch architecture within the CCD pixels on ESA's Gaia satellite
Doping atoms in the CCD silicon substrate are implanted to shape the electron distribution inside the CCD pixel (cf. Section 1.4). An extra implant can be made in the buried channel to concentrate small charge packets into a smaller volume. In this way small charge packets are less likely to interact with bulk traps. This extra implant that in principle runs through each CCD pixel column is called a Supplementary Buried Channel (SBC). The SBC is one of the hardware CTI mitigation tools present in all the Gaia CCDs. Its effect is however limited to the SBC full well capacity of about 3000 electrons, i.e. packets containing much more electrons than this limit will not experience any CTI mitigation due to the SBC. Experiments carried out by Kohley et al. (2009) on a Gaia CCD showed that for the particular device tested the SBC full well capacity in the first half of the CCD appeared to be less than 10 electrons. This led him to conclude on the absence of a SBC in this region of the CCD. **Chapter 6** is a detailed investigation of this issue including its consequences for the Gaia astrometric accuracy

and data processing. Using the model presented in Chapter 2, we simulate First Pixel Response (FPR) experiments for CCDs with a functioning and a non-functioning SBC in their first half. We can then re-analyze FPR experiments carried out on 7 Gaia CCDs manufactured before 2004 and deduce whether or not those CCDs are affected by the SBC issue. We also analyze the data from post-2004 CCDs using the latest experimental tests carried out on irradiated devices. We find that significantly more post-2004 CCDs are affected by the issue. We propose an explanation regarding why in these CCDs the SBC is not properly functioning in the device upper-half. e2v (the Gaia CCD manufacturer) predicts that all CCDs in the same manufacturing batch should have the same SBC characteristics. By comparing the batch numbers of the three affected post-2004 CCDs (three different batches) with those currently assigned to the Gaia satellite, we show that a minimum of 17% of flight CCDs are likely to be affected by the SBC issue. In the absence of further testing, we predict that in the other 29 completely untested batches 69% of the CCDs may be affected (between 11 and 100% with a 99% confidence interval). We also show that the absence of a SBC in the upper half of the CCD does not endanger the image location accuracy as it causes at most a 10% extra loss of accuracy. However the Gaia data processing must account for the SBC issue because the size of the image location bias is doubled at faint magnitudes. A review of the techniques that can be used prior to launch or on-board Gaia to identify which CCD is affected by this issue is also provided.

1.8 Lessons learned and outlook

With the preparation of the Gaia mission and this work in particular, the overall understanding of radiation damage on astronomical CCDs has significantly progressed and the means at our disposal to counter CTI are now well identified and characterized. The Gaia community is now more confident that CTI can be calibrated to the level required to preserve the astrometric accuracy of Gaia. This is only possible by combining hardware and software mitigation solutions. The study of the impact of CCD radiation damage on Gaia astrometry was a considerable effort, it required the use of an accurate model of CTI supported by dedicated experimental tests. This effort also showed that to understand the effect of radiation damage on a measurement carried out by CCDs on-board a space mission, CTI-induced errors at the image level must be propagated through the entire data processing. The particular measurements performed by Gaia required a novel approach to CTI mitigation at the image processing level. We believe that this new approach will also benefit other space astronomy missions, possibly with some adaptations the specific demands of these missions.

Future ESA missions under study such as Euclid and Plato plan to make extensive use of CCDs (e.g., 136 for the Plato focal plane). Euclid will carry out extremely accurate weak gravitational lensing measurements that require a very detailed knowledge of the instrument PSF in order to attain a very high level of accuracy in image shape characterization. CCD radiation damage and the associated CTI have already been identified as one of the main contributors to the overall error budget of the weak lensing measurements. Detailed studies of the impact of CCD radiation damage, such as presented in this thesis for the Gaia astrometric measurements, will have to be per-

formed for Euclid as well. They can greatly benefit from the knowledge acquired in this domain by the Gaia community and in particular by using the various models of CTI presented in this thesis. The CEMGA platform can play an important role in this transfer of knowledge.

CTI modelling has reached an unprecedented level of accuracy in the context of the Gaia mission, yet the models can still be improved. Monte Carlo models are intrinsically demanding in terms of computer resources. This strongly limits their applicability and their use in inferring parameters from experimental results. It is thus important to try to summarize their complexity into faster analytical models. Current fast analytical models such as CDM do not yet fully achieve the required level of performance, this can be achieve by taking into account in more detail the particularities of the CCD pixel architecture, in particular the SBC.

Although it is now demonstrated that the Gaia CTI mitigation strategy enables the recovery of the required astrometric accuracy, its implementation in the Gaia data processing is facing numerous difficulties, making the implementation a challenge in itself. Chapters 3, 4 and 5 address some of the issues currently faced. Nevertheless dedicated work and further experimental tests are still needed. In particular the photometric and spectroscopic processing do not yet completely account for CTI. The principles of the approach developed for the processing of the astrometric measurements will be used but will have to be adapted to the specific demands of the photometric and spectroscopic data processing. The first Gaia data downlink to Earth in 2013 will contain critical information about the actual effect of radiation damage. This information will confirm or alter our expectations, and most likely the CTI countermeasures will need to be tuned during the mission.

CCDs are by nature sensitive to radiation damage because charge must be transferred a long way and multiple times before it reaches the output node where it is converted into bits. Periodical charge injections, that fill the traps with electrons that do not belong to the signal, and a SBC can significantly counter CTI. However these countermeasures are not enough and future even more demanding space missions would benefit from investigating alternative detector technologies. p-channel CCDs seem to be less sensitive to radiation-induced CTI, this is a direction to further explore in the quest of radiation-hard devices. Soon, a Japanese mission, Nano-JASMINE will embed such a device to carry out astrometric measurements. It will be an ideal opportunity to test such a device in action. Also one can think about multiplying the number of output nodes to reduce the amount of required transfers and thus the CTI effects. For instance the CCDs planned to be used on-board the future dark matter mission from ESA, Euclid, will be equipped with 4 output nodes. Nevertheless as long as CCDs will be used in space, radiation damage will be a problem as charge will need to be physically transferred from one pixel to the other. Another type of light detector, a CMOS device, for which an output amplifier is embedded in each pixel, solves the CTI problem as charge transfer is no longer required. Unfortunately multiplying output amplifiers increase the noise. That is mainly why CMOS have not been popular so far regarding Astronomical applications in space that often aim at almost single electron detection. Nevertheless these devices continue to improve and CMOS detectors should ultimately replace CCDs in space for observations carried out in the visible and near

infrared (for more details about the next generation of light detectors see Kohley 2011). So will the successor to Gaia use CMOS-type devices? It might be the case although TDI operation of a CMOS is a very complex task and necessitates novel architectures that are only now starting to be investigated.

Chapter 2

Electrode level Monte Carlo model of radiation damage effects on astronomical CCDs

Current optical space telescopes rely upon silicon Charge Coupled Devices (CCDs) to detect and image the incoming photons. The performance of a CCD detector depends on its ability to transfer electrons through the silicon efficiently, so that the signal from every pixel may be read out through a single amplifier. This process of electron transfer is highly susceptible to the effects of solar proton damage (or non-ionizing radiation damage). This is because charged particles passing through the CCD displace silicon atoms, introducing energy levels into the semi-conductor bandgap which act as localized electron traps. The reduction in Charge Transfer Efficiency (CTE) leads to signal loss and image smearing. The European Space Agency's astrometric Gaia mission will make extensive use of CCDs to create the most complete and accurate stereoscopic map to date of the Milky Way. In the context of the Gaia mission CTE is referred to with the complementary quantity Charge Transfer Inefficiency (CTI $= 1-$CTE). CTI is an extremely important issue that threatens the performance of Gaia: the CCDs are very large so that the electrons need to be transferred a long way; the focal plane is also very large and difficult to shield; the mission will operate at L2 where the direct solar protons are highly energetic (penetrating); and the science requirements on image quality are very stringent. In order to tackle this issue, in depth experimental studies and modelling efforts are being conducted to explore the possible consequences and to mitigate the anticipated effects of radiation damage. We present here a detailed Monte Carlo model which has been developed to simulate the operation of a damaged CCD at the pixel electrode level. This model implements a new approach to both the charge density distribution within a pixel and the charge capture and release probabilities, which allows the reproduction of CTI effects on a variety of measurements for a large signal level range in particular for signals of the order of a few electrons.

T. Prod'homme, A.G.A. Brown, L. Lindegren, A.D.T. Short, S.W. Brown
Monthly Notices of the Royal Astronomical Society, 2011, Vol. 414

2.1 Introduction

We present a detailed physical Monte Carlo model of Charge Transfer Inefficiency (CTI) caused by displacement damage in irradiated CCD detectors. The development of the model took place in the highly challenging context of Gaia, a European Space Agency mission scheduled for launch in 2012. Gaia will operate for 5 years in an orbit around the second Lagrange point (L2) (Perryman et al. 2001; Lindegren et al. 2008) and will measure the parallaxes, proper motions, radial velocities, and astrophysical parameters of over one billion stars. To do so, the satellite will constantly scan the sky, observing the stars with two telescopes focussed on a single focal plane comprising 106 CCDs. The derived astrometric parameters are highly sensitive to the precise image shape, and hence to the effects of CTI.

Gaia will be subjected to the radiation environment at L2 which is entirely dominated by protons emitted during solar flares. The energetic protons collide with and displace atoms in the CCD silicon lattice, leading to the creation of interstitial atom-vacancy pairs. The vacancies thus formed, combine, by diffusion, with other vacancies or impurities (e.g. oxygen, phosphorus, carbon atoms) present in the CCD as doping implants or due to pollution during the fabrication process. The impurity-vacancy complexes introduce energy levels in the semiconductor band gap that stochastically trap and release the transferred signal carriers (electrons from the conduction band in n-type devices). The time-dependent capture and release probabilities vary as a function of several factors; most importantly, the temperature, the local charge density distribution in the vicinity of the trap, and trap parameters such as energy level and capture cross-section.

Based on the standard JPL model (Feynman et al. 1993), the average proton dose received by the CCDs at the end of the 5-year mission lifetime was originally predicted to be 4.14×10^9 protons cm^{-2} (with 90% confidence levels) for a launch in 2011 (Fusero 2007). Current space weather forcasts predict that the next solar maximum may be considerably less severe than average, so that the dose may be rather lower. However the sensitivity of Gaia to radiation damage is such that this in no way reduces the need to calibrate the effects. In addition, the peak of the Solar activity is expected to occur in late 2013 which means that Gaia will receive most of the total proton fluence early in the mission such that all data will be affected.

Based on experimental studies led by the industrial partners in the Gaia project and independent analyses carried out within the Gaia science community, the CTI resulting from radiation-induced traps is expected to affect the mission performance by causing charge loss and image distortion. Mitigating those effects has been recognized as critical to achieving the mission requirements. Several aspects specific to Gaia contribute to the high impact of radiation damage on the mission. The large focal plane is difficult to shield and the CCDs will be exposed to most of the incoming particles. The required image location accuracy is extreme, e.g., the end of mission parallax error is required to be better than 25 micro-arcseconds for a star of magnitude 15. The corresponding requirement on the residual image location error per CCD transit is 0.01 pixels. However the image profile distortion induced by CTI has been measured to cause biases in the image location measurement of up to 0.17 pixels. Gaia will also study very faint objects

down to magnitude 20 and at this signal level only a few electrons comprise the PSF core and the effects of trapping are poorly understood. Gaia will scan the sky by continuously spinning around an axis perpendicular to the plane containing the telescope viewing directions (see Perryman et al. 2001; Lindegren et al. 2008). In order to follow the resulting motion of the stars across the focal plane and integrate the light during the transit, the CCDs will be operated in Time-Delayed Integration mode (TDI mode). In this mode even fairly bright objects will remain faint for a part of their transit. Likewise, the sky background will form a gradient in the CCD parallel direction and, due to the relatively short integration time, Gaia will not benefit from the potential trap filling effects of a bright sky background.
These special operating conditions (high radiation dose, low signal level, low sky background and extremely high accuracy image location) demand a very high level of detail in the simulation of radiation damage effects, and preclude the use of models that assume instantaneous trapping within a certain volume that varies with the signal level (e.g., Massey et al. 2010; Rhodes et al. 2010). The Monte Carlo model presented here was developed to provide the required level of simulation detail. Our model thus simulates charge transfer at the electrode level and simulates the signal carrier trapping thanks to a new approach to the representation of both the charge density distribution and the capture and release probabilities. These simulations are used within the Gaia project to study the effect of CTI on measurements, to generate simulated data with which to verify the future radiation damage mitigation algorithms and to obtain a better understanding of CTI itself. In the following sections we describe the relevant details of our Monte Carlo model and show that it can reproduce the experimental data obtained from irradiated CCDs operated in TDI mode.

2.2 Model description

As described by Janesick (2001), a CCD needs to perform four fundamental tasks to generate an image: charge generation, charge collection, charge transfer, and charge measurement. The primary goal of our model is to simulate the effects of charge traps, induced by displacement damage in the CCD. These traps affect principally the third fundamental task of a CCD, the charge transfer, by stochastically capturing, and releasing charges during their transfer from one set of electrodes to another. The charge collection process can also be affected if a trap present in the CCD field free or depleted region captures a freshly photo-generated charge drifting towards the CCD buried channel. We chose not to take into account this secondary aspect and focus on the signal charge transfer only. Thus our model simulates exclusively the transfer of charges present in the signal confinement region under each electrode in the image section and the serial register (Fig. 2.1a). The signal confinement region is simulated as a box (Fig. 2.1b) of which the dimensions are defined by the manufacturing characteristics of the CCD, i.e. the width of the electrodes biased high in the transfer direction, the width of a pixel perpendicular to the transfer direction (serial direction) and in depth by the depletion of an electrode biased high with no electrons underneath. Fringing fields present at the edges of the signal confinement region reduce the actual volume (Seabroke et al. 2008a).

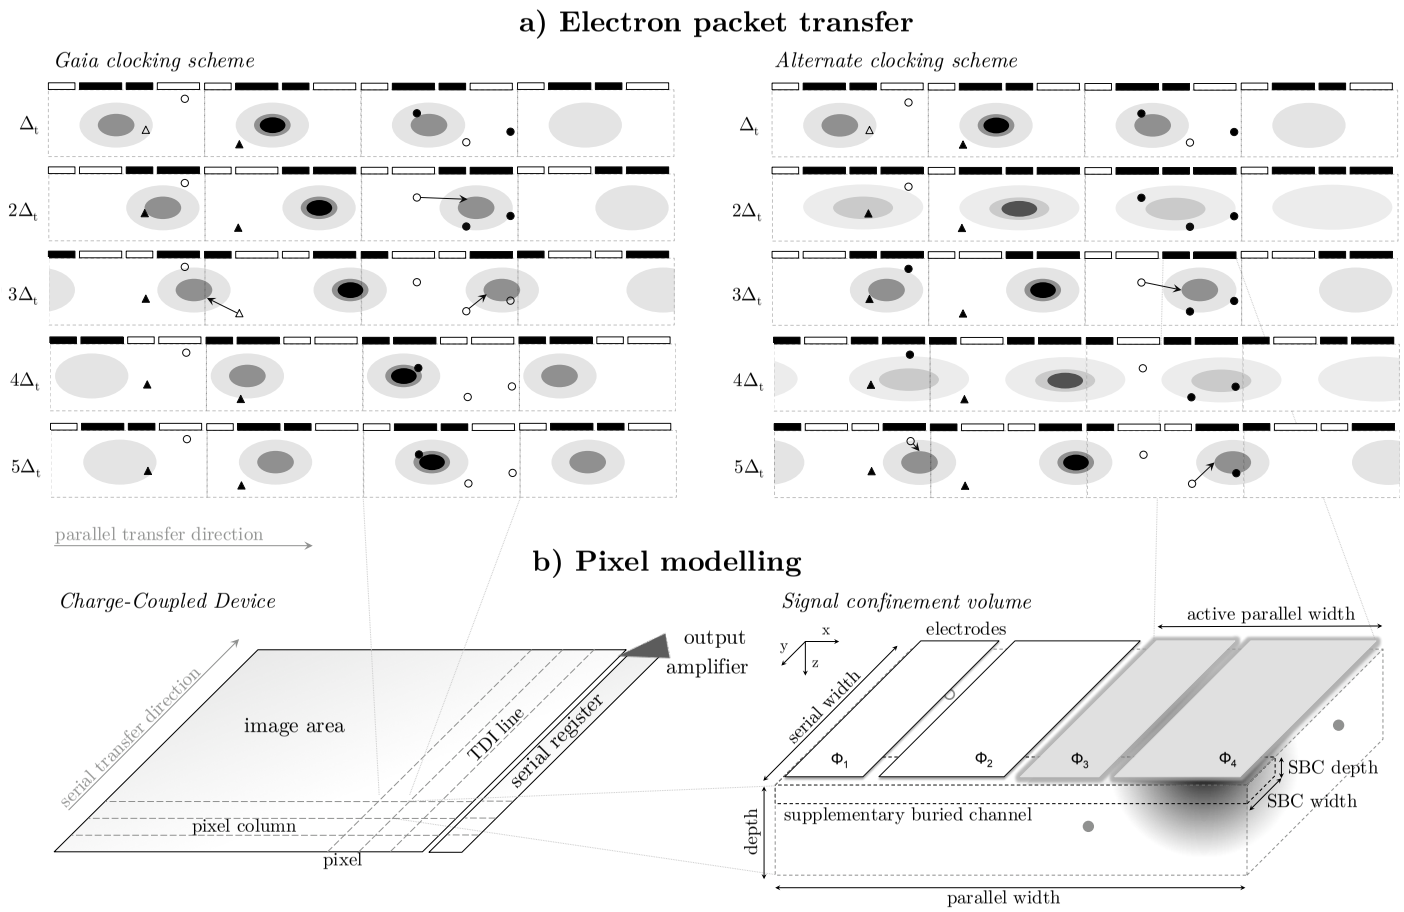


Figure 2.1 — (a) depicts the transfer of several charge packets from pixel to pixel in a four phase device. A high voltage is applied on the black coloured electrodes (later referred to as 'biased high') while no voltage is applied on the empty ones (at rest). The ellipsoidal shapes represent the charge density distribution under the electrodes. The packets interact with the traps of different trap species (circle and triangle) during the dwell time Δ_t. The filled traps (filled black symbols) can release their charge (arrows) at any time. To illustrate the importance of an electrode level simulation, two different clocking schemes are shown; note the charge density distribution successively stretching and contracting in the alternate case compared to the Gaia case. **(b)** illustrates the signal confinement volume, a key point of our modelling approach, and puts it in the context of the CCD as well as the supplementary buried channel. The CCD pixel rows are referred to here as TDI lines.

Parameter	Value
General	
Number of pixels (parallel × serial)	4500×1966
Number of light sensitive pixels	4494×1966
Pixel size (parallel × serial)	10×30 μm^2
Operational temperature	163 ± 3 K
Image section	
Number of phases	4
Transfer period	982.8 μs
Pixel FWC	$190\,000$ e$^-$
SBC FWC*	~ 1300 e$^-$
SBC size* (1$^{\text{st}}$ CCD half)	10×3 μm^2
SBC size* (2$^{\text{nd}}$ CCD half)	10×4 μm^2
Serial register	
Number of phases	2
Transfer period	< 0.5 μs
Pixel FWC	$475\,000$ e$^-$

Table 2.1 — The e2v CCD91-72 parameters. SBC stands for supplementary buried channel and FWC for full well capacity.
$\star$ nominal values; the simulated and measured values differ significantly.

To avoid any arbitrary assumption on the induced volume reduction, the dimensions of the signal confinement volume remain set to the manufacturing characteristics, while, as we shall see later in Sections 2.2.3 and 2.4.2.1, the actual distribution of the charge density within this volume is constrained by experimental measurements. This compensates to some extent our ignoring the fringing fields. Gaia CCDs (Short et al. 2005) are custom made by e2v technologies and referenced as CCD91-72. They are back-illuminated, full frame devices and incorporate a number of specific features such as a charge injection structure and 12 TDI gates to integrate bright stars over a shorter distance and avoid saturation. Each pixel also contains a supplementary buried channel (SBC), and an anti-blooming drain. In Table 2.1 important parameters of the Gaia astrometric CCDs are summarized. Our model is capable of simulating the effect of charge injection and also takes the SBC into account. The TDI gates and anti-blooming drain are not explicitly modelled.

2.2.1 Simulation process

Figure 2.2 presents a simplified version of the whole simulation process. The first step consists of defining an input signal and specifying the CCD characteristics such as the number of pixels in the parallel and serial directions, the number of electrodes per pixel, the clocking scheme, the operating temperature and so on.
Prior to the actual trapping and transfer simulation, empty bulk traps are randomly distributed across the CCD according to their specified concentration. Within each pixel, the traps are assigned a position in space by randomly generating coordinates within the signal confinement volume. If necessary the trap position can be kept fixed

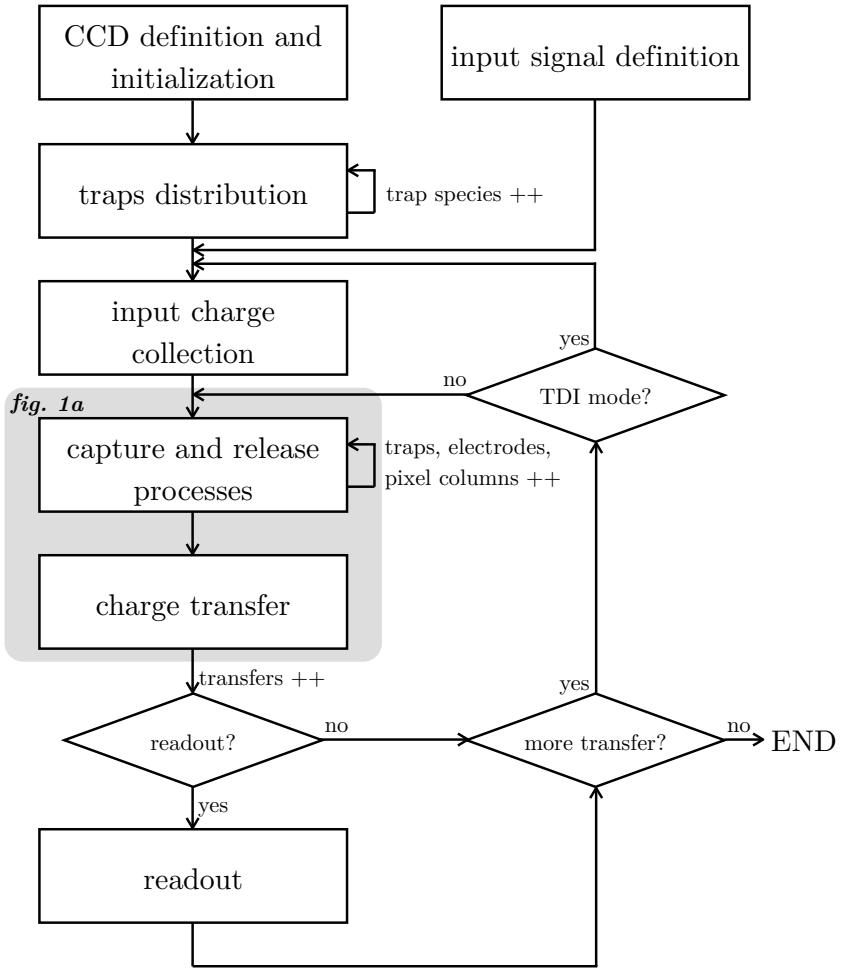


Figure 2.2 — Top level diagram of the simulation process.

to repeat experiments with the exact same simulated CCD. The traps can belong to different species defined by the parameters described in Table 2.2 according to the Shockley-Read-Hall (SRH) formalism.

The charge collection step corresponds to the generation of charges under the CCD electrodes. Photon detection from a light source can be described as a Poisson process, thus we create photo electrons in the CCD using a random generator with a Poisson distribution and a mean equal to the expected number of collected photons within the integration time. As a consequence only integer numbers of electrons are generated so that the actual physical process is reliably simulated, including the variance in the number of electrons generated. This is important in the Gaia context since a considerable number of star images will contain only a few electrons per pixel due to the faint nature of the observed sources and the operating mode of the CCDs. In TDI mode, the light source motion is synchronized with the CCD charge transfer rate so that the charge profile continues to build up as the image travels across the CCD. The transfer period is equal to the integration time at each step, which determines the number of photoelectrons generated. This implies that even for high signal levels the number of

Parameter	Description
E_t (eV)	Energy level in the semiconductor gap
σ_t (m^{-2})	Capture cross section
n_t (traps/pixel or m^{-3})	Concentration

Table 2.2 — The trap species parameters.

charges transferred will be very low, at least during the initial transfers. To simulate the background illumination (sky-light and scattered light) the corresponding background count rate is added to the expected number of collected photoelectrons before invoking the Poisson random number generator. Electronic charge injections (CIs) can also be simulated. In order to save simulation time any previous exposure of the CCD to a continuous level of background illumination can be simulated analytically. This is done by pre-computing (cf. Section 2.2.4) the corresponding trap occupancy level of each of the trap species and filling the corresponding number of traps prior to the transfer of the signal of interest. The capture and release process is simulated as follows. From the trap parameters, the density of charges at the trap position (cf. Section 2.2.3), and the interaction time, it is possible to calculate the probability of capture or release (cf. Section 2.2.2), for a specific trap relative to its state: empty or filled. If a trap is empty, the capture probability p is computed and a random number R generated; if $R < p$, then the capture is triggered and a charge is removed from the charge packet. Correspondingly, for a full trap, the release probability is computed and a random number is generated; if a charge is released, it is added to the closest charge packet (cf. Fig. 2.1a). This procedure is repeated for each trap in the CCD, pixel column by pixel column. The interaction time (cf. Section 2.2.2) is defined by the amount of time a charge packet stays under a given set of electrodes. It defines the temporal resolution of the simulation and depends on the charge transfer period, the number of electrodes and the clocking scheme. CCDs with two, three, or four phases can be simulated, and any kind of clocking scheme applied. This facilitates, for instance, testing the radiation hardness of different CCD configurations, taking into account the specific measurements to be carried out.
During the charge transfer step, the CCD electrodes are biased high or set at rest according to the predefined clocking scheme. The charge packets are then redistributed under the next set of biased high electrodes. Trapping, transfer and charge collection (in TDI mode) are repeated until the last charge packet belonging to the input signal reaches the serial register, at which point the simulation ends.
During read-out the signal charges are collected. It is also possible to simulate the charge and release processes in the serial register by repeating the same procedure as for the image section and making sure the illumination history is respected, i.e. all pixels in a line are processed, ordered by distance from the output amplifier (Fig. 2.1b).

2.2.2 Effective charge capture and release probabilities

In the SRH formalism charge capture and release are described as decay processes. One can derive the charge capture and release probabilities as follows. First let us consider a number of filled traps N_filled. In an infinitesimal time interval dt, the number

of released charges is proportional to N_{filled}. The proportionality constant is the release rate r_{r}. The number of filled traps as a function of time can then be derived:

$$\begin{aligned} \frac{dN_{\text{filled}}}{dt} &= -r_{\text{r}} N_{\text{filled}} \,, \\ N_{\text{filled}}(t) &= N_{\text{full,0}} e^{-r_{\text{r}} t} \,, \end{aligned} \tag{2.1}$$

where $N_{\text{full,0}}$ is the number of filled traps at $t = 0$. The fraction of filled traps remaining after a time t is then statistically equivalent to the probability for any specific trap to remain filled after a time t:

$$\frac{N_{\text{filled}}(t)}{N_{\text{full,0}}} = p_{\text{filled}}(t) = 1 - p_{\text{r}}(t) = e^{-r_{\text{r}} t} \,, \tag{2.2}$$

where p_{r} is the probability that a filled trap releases an electron within a time interval t:

$$p_{\text{r}}(t) = 1 - e^{-r_{\text{r}} t} \,. \tag{2.3}$$

The release rate constant is given by:

$$r_{\text{r}} = \frac{1}{\tau_{\text{r}}} = X \chi \sigma_{\text{t}} v_{\text{th}} n_{\text{c}} e^{-E_{\text{t}}/kT} \,, \tag{2.4}$$

where τ_{r} is the release time constant, σ_{t} the capture cross-section, E_{t} the trap energy level in the semiconductor forbidden gap, X the entropy factor, χ the field enhancement factor, T the CCD operating temperature, k the Boltzmann constant, and v_{th} the electron thermal velocity:

$$v_{\text{th}} = \sqrt{\frac{3kT}{m_{\text{e}}^*}} \,. \tag{2.5}$$

with $m_{\text{e}}^* = 0.5\, m_{\text{e}}$ the effective electron mass. The effective density of states n_{c} in the conduction band is:

$$n_c = 2 \left(\frac{2\pi m_e^* kT}{h^2} \right)^{3/2} \,, \tag{2.6}$$

where h is the Planck constant.
Likewise one can derive the probability that an empty trap captures an electron within a time interval t:

$$p_{\text{c}}(t) = 1 - e^{-r_{\text{c}} t} \tag{2.7}$$

$$r_{\text{c}} = \frac{1}{\tau_{\text{c}}} = \sigma_{\text{t}} v_{\text{th}} n_{\text{e}} \tag{2.8}$$

where r_c is the capture rate, τ_c the capture time constant, and n_e the electron density at the trap location (cf. Section 2.2.3).
The characteristic interaction time between the traps and a charge packet is the dwell time Δ_t. It corresponds to the elapsed time between two charge redistributions, during which n_e remains constant. The dwell time is proportional to the transfer period and depends on the selected clocking scheme. Δ_t is greater in the CCD image area than in the serial register, since a complete TDI line (~ 2000 pixels) must be read out during a pixel to pixel transfer in the image area. For the Gaia CCDs, Δ_t varies from several tens of micro-seconds (serial register) up to a forth of milli-second (image section). The release time constants for certain trap species can be as short as several hundreds of nano-seconds like the A centre (oxygen-vacancy complex). These 'fast' traps can lead to multiple capture and release events during a dwell time and play an important role in the CTI associated with the serial register. As a consequence we shall follow Lindegren (1998) and introduce effective probabilities of charge capture and release that take into account the possibility of multiple capture and release cycles within one time interval. We consider the probability p_{filled} for a trap of unknown state to be filled after a time interval t. N is the total number of traps and N_{empty} the number of empty traps:

$$N = N_{\text{filled}} + N_{\text{empty}} \tag{2.9}$$

We follow the same derivation as in (eq. 2.1) but now accounting for the empty traps

$$\begin{aligned} \frac{dN_{\text{filled}}}{dt} &= r_c N_{\text{empty}} - r_r N_{\text{filled}} \,, \\ \frac{1}{N}\frac{dN_{\text{filled}}}{dt} &= \frac{dp_{\text{filled}}}{dt} \,, \\ &= r_c\, p_{\text{empty}} - r_r\, p_{\text{filled}} \,, \\ &= r_c\,(1 - p_{\text{filled}}) - r_r\, p_{\text{filled}} \,. \end{aligned} \tag{2.10}$$

In these equations p_{empty} is the probability for a trap of unknown state to be empty after the time interval t. One can then derive the following general solution, assuming that r_c and r_r are constants:

$$p_{\text{filled}}(t) = \frac{r_c}{r_r + r_c} + C \exp\left[-\left(r_r + r_c\right) t\right] \,. \tag{2.11}$$

It is now possible to derive the effective capture and release probabilities, p_c and p_r. If the trap is empty at $t = 0$ then $p_{\text{filled}}(0) = 0$ and $C = -r_c/(r_r + r_c)$, thus:

$$p_c \equiv p_{\text{filled}}(t) = \frac{r_c}{r_r + r_c}\left(1 - \exp\left[-\left(r_r + r_c\right) t\right]\right) \,. \tag{2.12}$$

Similarly, if the trap is filled at $t = 0$ then $p_{\text{filled}}(0) = 1$ and $C = +r_c/(r_r + r_c)$. Hence:

$$\begin{aligned} 1 - p_r \equiv p_{\text{filled}}(t) &= \frac{r_c + r_r \exp\left[-\left(r_r + r_c\right) t\right]}{r_r + r_c} \,, \\ p_r &= \frac{r_r}{r_r + r_c}\left(1 - \exp\left[-\left(r_r + r_c\right) t\right]\right) \,. \end{aligned} \tag{2.13}$$

2.2.3 Charge density distribution modelling

The CTI effects model we describe in this chapter is a density driven model to be contrasted with the more commonly used volume driven models. The volume driven models assume instantaneous trapping within a certain volume that varies with the signal level. A density driven model necessitates the computation of the capture and release probabilities for each trap — taking the charge density in the trap vicinity into account — regardless of its location (no trap is a priori ignored). The density driven model thus requires the evaluation of the charge density distribution as a function of the signal level and location within the pixel signal confinement region. The confinement region is defined by the electrodes biased high. As can be seen from the work of Hardy et al. (1998) and more recently Massey et al. (2010) and Rhodes et al. (2010), the volume driven approach is fairly successful in explaining experimental data, in particular HST data. However the Gaia operating conditions differ significantly from the HST ones. The Gaia CCDs will be operated in TDI mode. In this mode the exposure time equals the charge transfer period. Thus the charge-trap interaction time is significantly decreased and instantaneous trapping cannot be assumed anymore. Moreover, as already stated in the introduction, Gaia will deal with very low levels of background and source signal. Trapping in these particular conditions was investigated for the first time in studies related to the Gaia mission. The density driven approach proved to be necessary to explain and reproduce experimental results (Short 2007; Seabroke et al. 2008a) which show that the CTI effects are modified by an extremely small level of background light (of the order of a few photons per second). The volume occupied by these very few electrons should be negligible and thus prevent any trapping from occurring according to volume-driven models. However the Gaia experimental studies showed that very few electrons are capable of filling a significant amount of traps. This can be explained by a density driven approach. Indeed, as the background light constantly illuminates the CCD, the long effective charge-trap interaction time for these very few electrons compensates for the very small charge density in the vicinity of each trap and trapping then becomes likely to occur. Our simulation method is also particularly convenient for accurately simulating CTI effects over a wide range of signal levels and for CCDs in which extra doping implants (such as a supplementary buried channel) modify the pixel potential and induce non-linearities between the CTI effects and the signal intensity (cf. Section 2.4).

The charge density distribution within the signal confinement volume, which varies strongly with the CCD architecture, is thus a key parameter in CTI modelling. Although the distribution cannot be directly measured, in principle it can be accurately determined for a specific CCD architecture and signal level by solving simultaneously the Poisson equation and the charge continuity equation (e.g. Seabroke et al. 2009) in order to find a consistent electrode potential and charge carrier distribution. This requires a detailed knowledge of the CCD implant characteristics (i.e., the nature and concentration of the dopants). This information is often commercially sensitive and in order to keep our model flexible regarding its application to other cases than Gaia, we use an analytical description of the charge density distribution which is roughly consistent with the modelling results by Seabroke et al. (2009). This analytical description

consists of a normalized Gaussian function in the three space directions of which the complexity increases with the number of the CCD potential characteristics included. The distribution parameters are listed in Table 2.3. The density is defined as:

$$n_{\mathrm{e}}(x,y,z) = S \times \rho(x,y,z) = S \times \rho(\mathbf{x})\,, \tag{2.14}$$

where

$$\rho(\mathbf{x}) = \frac{\exp\left[-\frac{1}{2}\left(\mathbf{x}-\mathbf{x}_0\right)^T \mathbf{C}^{-1}\left(\mathbf{x}-\mathbf{x}_0\right)\right]}{\left(\sqrt{2\pi}\right)^3 |\mathbf{C}|^{1/2}}\,, \tag{2.15}$$

with

$$\mathbf{C} = \begin{pmatrix} \sigma_x^2 & 0 & 0 \\ 0 & \sigma_y^2 & 0 \\ 0 & 0 & \sigma_z^2 \end{pmatrix}\,. \tag{2.16}$$

Hence

$$\rho(\mathbf{x}) = \frac{\exp\left[-\frac{1}{2}\left(\left(\frac{x-x_0}{\sigma_x}\right)^2 + \left(\frac{y-y_0}{\sigma_y}\right)^2 + \left(\frac{z-z_0}{\sigma_z}\right)^2\right)\right]}{\left(\sqrt{2\pi}\right)^3 \sigma_x\sigma_y\sigma_z}\,. \tag{2.17}$$

The supplementary buried channel (SBC) corresponds to an additional doping implant, which generates a deeper potential well and narrows the charge distribution so that electrons are transferred through a smaller volume of silicon and encounter fewer traps. In the Gaia CCDs, the SBC potential well collapses at signal levels of a few thousand electrons. As a consequence, the SBC only improves CTI in the small signal regime. To implement the SBC, we simply introduce a second, low signal regime with a different charge density distribution and ensure a smooth transition between them so that (eq. 2.17) becomes:

$$\rho(\mathbf{x}) = \frac{\exp\left[-\frac{1}{2}\left(\left(\frac{x-x_0}{\sigma_x}\right)^2 + \left(\frac{y-y_0^\star}{\sigma_y^\star}\right)^2 + \left(\frac{z-z_0^\star}{\sigma_z^\star}\right)^2\right)\right]}{\left(\sqrt{2\pi}\right)^3 \sigma_x\sigma_y^\star\sigma_z^\star} \tag{2.18}$$

where the parameters indicated with a $\star$ now vary as a function of signal level:

$$P^\star = P\left(1 - e^{-\frac{S}{S_{\mathrm{SBC}}}}\right) + P_{\mathrm{SBC}} e^{-\frac{S}{S_{\mathrm{SBC}}}}\,, \tag{2.19}$$

where P refers to the parameter value in the buried channel and P_{SBC} to the corresponding value in the SBC.
At high signal levels, saturation effects limit the linear growth of the charge density; while more charges can still be added, the maximum density cannot be overcome and the distribution expands in the three spatial directions. Saturation occurs when the charge density becomes larger than the doping concentration. Taking into account

the presence of the anti-blooming drain we consider that locally, the charge density cannot exceed $n_{\text{sat}} = S_{\text{sat}}/V_e$, with $V_e = (2\pi)^{3/2}\sigma_x\sigma_y\sigma_z$. Following the method proposed by Lindegren (1998), the model for the saturation process is:

$$n_{\text{e}}(\mathbf{x}) = \frac{n_{\text{sat}} S' \rho(\mathbf{x})}{n_{\text{sat}} + S'\rho(\mathbf{x})} \tag{2.20}$$

where S' has to be adjusted to give the correct total charge S:

$$\begin{aligned} S &= \iiint n_{\text{e}}(x,y,z)\,dx\,dy\,dz\,, \\ S &= \sqrt{\frac{2}{\pi}} \int_0^\infty \frac{ur^2}{u + \exp\left(r^2/2\right)}\,dr\,, \end{aligned} \tag{2.21}$$

with $u = S'/S_{\text{sat}}$ and $S_{\text{sat}} = n_{\text{sat}} \times (2\pi)^{3/2}\sigma_x\sigma_y\sigma_z$. In the linear growth case:

$$u \ll 1 \quad \text{and} \quad S \simeq S_{\text{sat}} u\,. \tag{2.22}$$

In the saturation case:

$$u \gg 1 \quad \text{and} \quad S \simeq S_{\text{sat}} \left(4/3\sqrt{\pi}\right) (\ln u)^{3/2}\,. \tag{2.23}$$

The adjustment of S' is then analytically constructed to provide the right behaviour in the linear and saturation cases and a reasonable approximation for $u \approx 1$, i.e. the transition region:

$$S' = S\left(1 + (S/S_{\text{sat}})^{0.8}\right)^{-1.25} e^{\left(\frac{3\sqrt{\pi}}{4} S/S_{\text{sat}}\right)^{2/3}}\,. \tag{2.24}$$

Figure 2.3 shows an example of the evolution of the charge density distribution as a function of the signal level as it is simulated by the presented model. It illustrates the particular features that the model aims to reproduce. At low signal level the charge density distribution shows a narrow profile in accordance with the expected SBC effect. While the signal level increases, the density profile widens and shifts to the centre of the pixel. It corresponds to the transition from the SBC regime to the BC regime in the case of a SBC implant not located at the centre of the pixel as depicted in Fig. 2.1b (which corresponds to the Gaia case). And finally one can clearly observe the saturation occurring for the large signal levels.

2.2.4 CCD illumination history

The CCD illumination history determines the CCD state (i.e. the trap occupancy level) prior to a star transit. Every star brighter than magnitude 20 will be observed ~ 80 times by Gaia (on average). During each observation of a given star, the satellite will be scanning the same part of the sky but from a different direction i.e. with a different orientation of the focal plane. Hence the occupancy of the traps in the CCD will be different prior to each transit of a given star. As a consequence the CTI effects (image location estimation bias and charge loss) are likely to be different from one observation of a

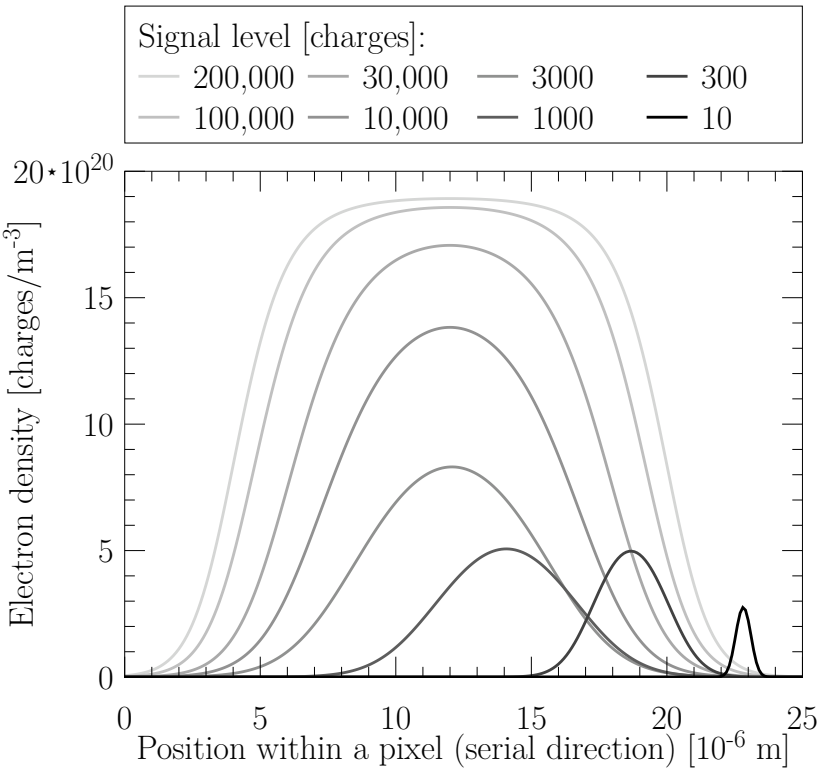


Figure 2.3 — Example of a charge density distribution as it is simulated by the presented analytical model: the charge density at the center of the charge cloud (in the parallel direction and in depth) is plotted along the CCD serial direction for different signal levels. The simulation includes an off-centred SBC and a saturation occurring at high signal level.

particular star to the next (even if the radiation damage to the CCDs were unchanged). Apart from the trapping process and the electron density distribution, the CCD illumination history is therefore another key element of the CTI effects modelling. The illumination history is determined by discrete events such as star transits and charge injections, which are directly reproduced during the simulation. Our simulations also account for a continuous optical background comprising light from unresolved stars and scattered light within the spacecraft. The light from the sky background constantly illuminates the Gaia CCDs at the level of ~ 42 e^- s^{-1} arcsec^{-2} (taking the contributions of both telescopes into account), i.e. each CCD pixel will on average receive $\sim 4 \times 10^{-4}$ e^- per pixel transfer step. In TDI mode the charges generated by the sky background are not only integrated but also continuously transferred from one pixel to another, to form a signal gradient along the parallel direction, the last light-sensitive pixels in the CCD effectively receiving about 2 e^-.

Brown (2009a) showed that such low levels of constant illumination noticeably modify the trap occupancy level. Indeed even if the local electron density distribution remains approximately unchanged, the interaction time of these background electrons with the traps is effectively much greater than the one of a usual transiting source. To simulate the effect on the trap occupancy level, one can run the model for a certain amount of time with the background light as the only input signal. The star signal of interest is

then inserted once an equilibrium trap occupancy level has been attained. However, our model is computationally very intensive and using it to simulate several minutes of CCD operation would take too much time. For this reason prior to each simulation, we determine the state of each trap using a random generator and considering the probability $p_{\text{filled}\,\infty}$, the probability for a particular trap to be filled by an electron from the background light. $p_{\text{filled}\,\infty}$ is computed from eq. 2.12 assuming an infinite time of interaction:

$$p_{\text{filled}\,\infty} = p_{\text{filled}}(t = \infty) = \frac{r_{\text{c}}}{r_{\text{r}} + r_{\text{c}}}. \tag{2.25}$$

In the simulations, the background contribution is also taken into account during the charge collection step by adding the background count rate to the expected number of collected photoelectrons before invoking the Poisson random number generator (cf. Section 2.2.1). Note that the dark current can be taken into account in the same way. However in the Gaia case, the contribution of the dark current to the global background at operational temperature is not significant ($\sim 10^{-7}$ e^-per pixel transfer step), we thus ignore it in the following.

Table 2.3 — Charge density distribution parameters

Parameter	Description
S (e^-)	number of e^-in the signal confinement volume
n_{sat} (m^{-3})	density saturation level
Signal confinement volume	
x_{max} (m)	parallel width
y_{max} (m)	serial width
z_{max} (m)	depth
V (μm^3)	signal confinement volume $x_{\text{max}} \times y_{\text{max}} \times z_{\text{max}}$
Buried channel regime	
σ_x (m)	parallel distribution standard width
σ_y (m)	serial distribution standard width
σ_z (m)	depth distribution standard width
x_0 (m)	distribution centre parallel coordinate
y_0 (m)	distribution centre serial coordinate
z_0 (m)	distribution centre depth coordinate
S_{sat}(e^-)	Signal level at which the buried channel saturates $\equiv$ FWC
Supplementary buried channel regime	
$\sigma_{y,\text{SBC}}$ (m)	parallel distribution standard width
$\sigma_{z,\text{SBC}}$ (m)	depth distribution standard width
$y_{0,\text{SBC}}$ (m)	distribution centre serial coordinate
$z_{0,\text{SBC}}$ (m)	distribution centre depth coordinate
S_{SBC} (e^-)	Signal level at which the SBC saturates $\equiv$ FWC_{SBC}

2.3 Radiation tests

We describe the verification of our model against experimental data. These data were obtained by the industrial partners in the Gaia project.

Several experimental studies were carried out on irradiated CCDs in order to evaluate the impact of CTI on Gaia's scientific requirements, to define the optimal operating temperature, to prepare the CCD calibration activities, and to elaborate a radiation damage mitigation strategy. Sira electro-optics acquired the first sets of CCD radiation test data (Hopkinson et al. 2005), focusing on the determination of trap parameters and CCD characterization. Later, Surrey Satellite Technology Limited (SSTL, formerly Sira) investigated the potential difference between a CCD irradiated at room temperature and a CCD irradiated at 163K and kept at that temperature. SSTL concluded that the results obtained for CCDs irradiated at room temperature should be adequate for Gaia performance predictions within the usual experimental uncertainties (Hopkinson 2008).

Up to now the prime contractor for Gaia, EADS Astrium, has performed three test campaigns on Gaia CCDs irradiated at room temperature with a radiation dose of 4×10^9 protons cm^{-2} (10 MeV equivalent), and a fourth one is on-going. The experimental setup includes a translation stage which enables it to reproduce the star motion and hence to operate the CCD in TDI mode. The CCD was cooled and operated at constant temperature throughout the test campaigns. During the first campaign, Astrium concentrated on determining CTI effects on Gaia's astrometric measurements: the charge loss and the image profile distortion leading to a biased evaluation of the image location. The experiments were carried out as a function of stellar brightness (i.e. signal level) and background light level. The purpose of this first campaign was also to evaluate the viability of an artificial diffuse optical background source as a CTI mitigation device. The second campaign (RC2) allowed an alternative mitigation tool, Charge Injection (CI), to be thoroughly studied. Each of Gaia's CCDs contains a row of diodes and gate electrodes before the first pixel electrodes in the image section. A row of charges can thus be injected and transferred to fill the traps prior to a star transit. The CI level or the quantity of injected electrons is defined by the difference in voltage applied to the pixel and the gate electrodes. To study the influence of the CI parameters on the CTI effects, RC2 data has been acquired for different CI levels, durations (number of CI rows at a time), and delays (elapsed time between the CI and the first star transit), at different temperatures. The third campaign used a realistic sky-like illumination pattern to simulate the star transits instead of the uniform illumination grid from RC2.

Not only astrometric tests were performed during these campaigns but photometric and spectrometric issues were also addressed. The CCDs used by the red photometer and the radial velocity spectrometer (RVS) instruments on board Gaia differ slightly from the astrometric ones. They are based on the same architecture but are thicker devices and use a modified anti-reflection coating to enhance their quantum efficiency in the red wavelength band. A set of preliminary tests was performed during RC2 on such a device irradiated with a lower radiation dose (2×10^9 protons cm^{-2}), followed by a more detailed study during the radiation campaign 3 (RC3). This study included

a mask mimicking realistic G2V stellar spectra and a very tight control over the background light. As a result, the red-enhanced CCD was established to be more sensitive to radiation damage (most likely due to a greater depletion depth) and the expected shift and distortion of the spectral features induced by CTI were characterized as a function of stellar brightness and level of background light.
The raw data acquired during these campaigns were made available to the Gaia Data Processing and Analysis Consortium in order to be re-analyzed and to support the CTI modelling efforts.
To compare our model results with the experimental data, we relied mostly on the RC2 data (Georges 2008; Brown 2009b). We now describe in more detail RC2 experiments. The CCD irradiation scheme of the image section consists of three areas of similar size. The first area is non-irradiated, the second is irradiated at the level of 4×10^9 protons cm^{-2} (10 MeV equivalent) and the third, irradiated at a higher dose, was not used during testing. The serial register is non-irradiated. The reference and CTI affected data are obtained by translating the light source and a mask over respectively the first and second area of the CCD. The mask contains 50×22 (parallel $\times$ serial) holes, and the projected stellar images are separated by 50 pixels in the parallel direction and 20 pixels in the serial direction. Each stellar image is binned along its serial dimension which necessitates some assumptions in the input signal modelling (detailed in Section 2.4). For each test at least 5 consecutive scans with identical configuration are performed. The mean time interval between the end of a scan and the beginning of a new one is 29.7 s. The first scan has a different illumination history from the others and is usually excluded from the analysis. We used two different astrometric tests out of the three that were performed during RC2. In the first test the CI delay and duration are kept fixed while the CI levels vary from ~ 4000 to $\sim 115\,000$ e$^-$. Each injection level was tested at different temperatures (from 163 K to 198 K) and for different illumination levels (corresponding to stellar magnitudes 13.3 and 15). In the second test the delay between the CI and the first star transit is varied from 30 to at most $120\,000$ pixels (i.e. from ~ 29 ms to ~ 118 s) while the CI level is fixed to $20\,000$ e$^-$ and the CI duration to 20 pixels. The test temperature was kept close to 163 K. The whole sequence of tests was repeated for different illumination levels (star magnitudes 13.63, 15.29, 16.96, 18.65, 20.25). In the following section, we explain how we validated our MC model against these experimental tests.

2.4 Model verification and comparison to experimental data

To validate the model we proceed in two steps. First the unit blocks of the model, the capture and release processes, are tested individually by comparing the results of Monte Carlo simulations to statistical analytical predictions derived from the equations in Section 2.2.2. The second step consists of a direct comparison between the model and a subset of the RC2 data. The variety of tests performed under different experimental conditions during RC2 allows us to separately verify the different features of our model. We first address the capability of our electron density distribution model to reproduce the CTI effects over a large range of signal levels and in particular to mimic the effect of the CCD supplementary buried channel. Subsequently we

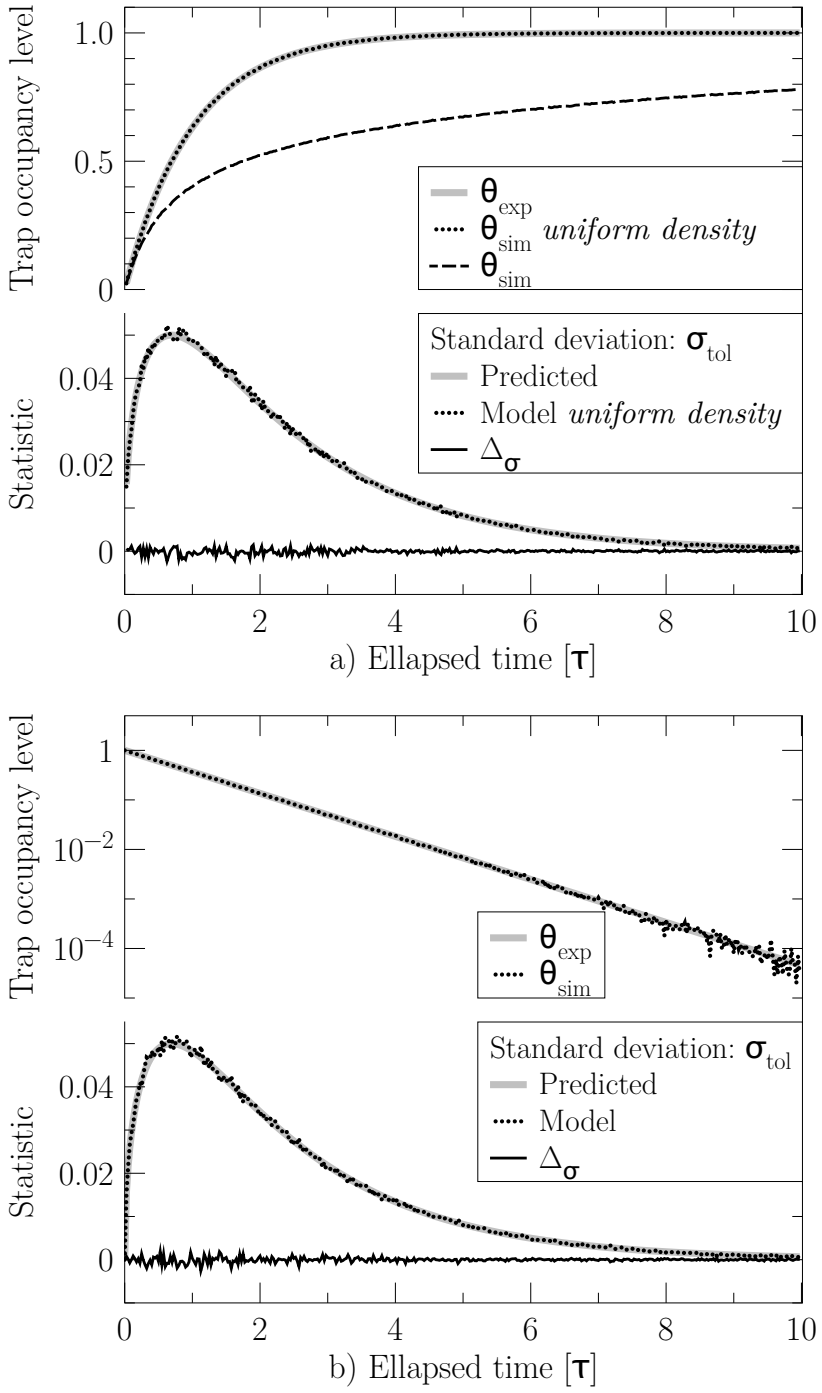


Figure 2.4 — Tests of the capture (a) and release (b) processes. Top panels: comparison between the expected trap occupancy level (grey) and the mean trap occupancy level for 1000 Monte Carlo realizations for a CCD containing 100 traps (black). For the capture test (a), the dotted line depicts the results obtained for a uniform electron density distribution and the dashed line for a 3D Gaussian distribution including a SBC (eq. 6.6). Bottom: the expected (grey line) and actual standard deviation (dotted black line, measured from the Monte Carlo experiments) for the trap occupancy levels, and the difference between them (black line).

investigate how well the model reproduces individual measurements acquired in TDI mode.

2.4.1 Model comparison to analytical prediction

During each transfer step of a simulation, it is possible to monitor the trap occupancy level θ_{sim} (or fraction of filled traps in the CCD):

$$\theta_{\text{sim}} = \frac{N_{\text{filled}}}{N}\,, \tag{2.26}$$

with N_{filled}, the number of filled traps and N, the total number of traps in the CCD. For a simple experimental configuration, one can analytically predict the evolution of the trap occupancy level θ_{exp} as a function of time and compare it with θ_{sim} to verify the reliability of the basic steps in the Monte Carlo simulation. The capture module in the simulation is tested by simulating a CCD under a constant illumination, which contains a unique trap species with a capture time constant τ_{c} and an infinite release time constant so that no electron release can occur. We also make all traps in the CCD empty at the beginning of the simulation, i.e. $\theta_{\text{sim}}(0) = 0$. Under these conditions, the expected trap occupancy level θ_{exp} is equivalent to the capture probability:

$$\theta_{\text{exp}}(t) = p_{\text{filled}} = p_{\text{c}} = 1 - e^{-\frac{t}{\tau_{\text{c}}}}\,. \tag{2.27}$$

To compute τ_{c} analytically (eq. 6.2), we assume a uniform electron density distribution, which we also implemented in our model for the purpose of this test. Similarly, to test the release module in the simulation we simulate a CCD in complete darkness, which contains a unique trap species with a release time constant, τ_{r} and an infinite capture time constant. All the traps are artificially filled at the beginning of the simulation so that $\theta_{\text{sim}}(0) = 1$. Under these conditions, one can write:

$$\theta_{\text{exp}}(t) = p_{\text{filled}} = 1 - p_{\text{r}} = e^{-\frac{t}{\tau_{\text{r}}}} \tag{2.28}$$

The number of traps in a CCD is finite. Thus, θ_{sim} and θ_{exp} are necessarily different due to the discrete nature of the capture and release process (no fractions of electrons can be captured or released). Yet we want to assess how accurately our model reproduces the expected value in these particular conditions. After a time t each trap in a CCD can only be empty or filled, with the probability of occupancy given by p_{filled}. Hence the simple simulations of trapping or release considered here constitute a Bernoulli trial where one counts the number of times n that a trap is filled or emptied after time t. The value of θ_{sim} is then given by n/N and will follow a binomial distribution with the following expectation value and variance:

$$\begin{aligned} E(\theta_{\text{sim}}) &= p_{\text{filled}} = \theta_{\text{exp}}\,, \\ \text{Var}(\theta_{\text{sim}}) &= \frac{1}{N} p_{\text{filled}}(1 - p_{\text{filled}}) = \frac{1}{N}\theta_{\text{exp}}(1 - \theta_{\text{exp}})\,. \end{aligned} \tag{2.29}$$

We repeat the Monte Carlo simulation N_{run} times in order to compute the mean trap occupancy level $\langle\theta_{\text{sim}}\rangle$ and its variance $\text{Var}(\theta_{\text{sim}})$ and to verify whether these values

are the same as the expected ones given in (eq. 2.29). As can be seen from Fig. 2.4, for the capture and release tests the model reproduces the expected results. It is also interesting to note that even in this simplistic experimental configuration there is a clear difference (Fig. 2.4a) between the capture profile obtained for a uniform electron density distribution and a more realistic distribution as described in Section 2.2.3.

2.4.2 Model comparison to experimental data

2.4.2.1 Fractional charge loss

The charge loss induced by CTI is directly connected to the trap capture process, which is particularly sensitive to the electron density distribution. As a consequence, reproducing charge loss measurements over a large signal range implies an accurate modelling of the electron density distribution. That is why we chose to extract fractional charge loss measurements from RC2 data to verify our particular approach to the electron density distribution modelling described in Section 2.2.3.
Thanks to its simple signal shape, it is particularly convenient to study the charge loss occurring in a charge injection (CI) line. An undamaged CI profile gives a constant signal with a mean value corresponding to the CI level (cf. Section 2.3), whereas a damaged CI profile typically shows a strong electron deficit in the leading edge corresponding to electron captures and a slight electron surplus in the end of the profile as well as a trailing edge after the CI profile due to the release of electrons from the traps. To obtain the charge loss one needs to compare the damaged and undamaged profiles. RC2 data does not contain undamaged CI profile since the reference data were acquired without CI, it is therefore impossible to know with any accuracy, the number of electrons injected for a specific injection level. To compute the charge loss we then assume that the average number of electrons over the last 4 CI lines (CI duration $= 20$ pixels) would constitute an acceptable reference. Those last pixels undergo the least charge loss and therefore remain closest to the actual number of electrons injected per line. The simulations, for which the true CI level is known, show that this assumption leads to a slight underestimation of the fractional charge loss at all signal levels. Hence, the last 4 CI lines are also used in the simulations to compute the CI reference level and subsequently the fractional charge loss to avoid any bias in our comparison. The fractional charge loss, formally the charge loss divided by the charge injection level times the CI duration, allows us to study the fraction of signal that is lost due to CTI in a CI profile as a function of the signal level.
In Figure 2.5 the black crosses show the fractional charge loss as a function of the CI level extracted from RC2 data, we obtain results similar to Hopkinson et al. (2005) (for a direct comparison note that the devices were irradiated at different doses in the two studies). Fig. 2.5 reveals the complex structure of the pixel by showing a clear break (close to the nominal SBC full well capacity 1500 e^-) in the increase of the fractional charge loss as the signal level diminishes.
To reproduce these measurements with our model we simulate a simplified version of the experimental setup. For each injection level we simulate the operation of a CCD with a single pixel column of 4494 pixels, only 2 scans are performed (instead of 5), the CCD contains a unique trap species, no DOB is added, and the serial register is

Density parameters	
$\mathrm{FWC}_{\mathrm{SBC}}\star$	$2824.89\ \mathrm{e}^-$
$\mathrm{FWC}_{\mathrm{BC}}$	$190\,000\ \mathrm{e}^-$
$\eta_{\mathrm{BC}}\star$	0.101
$\eta_{\mathrm{SBC}}\star$	0.111
σ_x	$1.11\,\mu\mathrm{m}$
σ_y	$2.42\,\mu\mathrm{m}$
σ_z	$0.076\,\mu\mathrm{m}$
$\sigma_{y,\mathrm{SBC}}$	$0.22\,\mu\mathrm{m}$
$\sigma_{z,\mathrm{SBC}}$	$0.01\,\mu\mathrm{m}$
$y_{0,\mathrm{SBC}}$	$23\,\mu\mathrm{m}$
$z_{0,\mathrm{SBC}}$	$0.05\,\mu\mathrm{m}$
Trap species parameters	
$\rho\star$	4.08 traps/pixel
σ	$5.00\,10^{-20}\ \mathrm{m}^2$
τ	18.06 ms
Goodness-of-fit	
χ^2	338
χ^2_{red}	21

Table 2.4 — Simulation parameters and goodness-of-fit for the model result example. The fitted parameters are indicated with a $\star$.

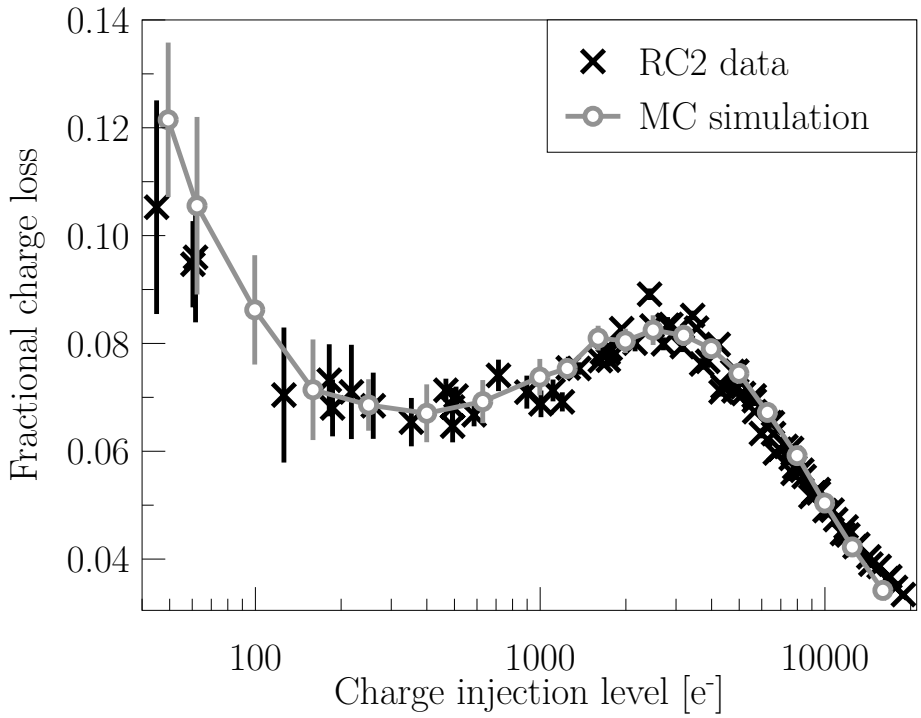


Figure 2.5 — Comparison between the modeled fractional charge loss (grey) and experimental data extracted from RC2 (black) as a function of the signal level.

not simulated. The CCD is operated at the Gaia operational temperature (163 K), it is operated in TDI mode as in the RC2 tests. The selected electron density distribution includes the description of the SBC but no saturation. A charge injection is performed at the beginning of each scan. For each considered CI level, we perform this simulation with the same CCD (i.e. the same CCD characteristics and trap locations), we then measure the fractional charge loss following the same approach as in the test data analysis. The whole procedure is repeated 8 times and the charge loss measurements are then averaged so that the modeled fractional charge loss is independent from the noise due to the stochastic nature of the trapping processes.
The model outcomes are then compared to an analytical fit to the experimental data and χ^2 is used as a comparison criterion.

$$\chi^2 = \sum_{i=0}^{S-1} \frac{\left(\lambda\left(x_i\right) - \Phi\left(x_i\right)\right)^2}{\sigma^2} \tag{2.30}$$

where λ is the simulated fractional charge loss, Φ the fitted fractional charge loss derived from RC2 data, x_i a particular signal level, and σ the noise. For a particular signal level, the experimental data is the mean of 4 measurements (cf. Section 2.3). We assume that the standard deviation from this mean measurement (black bars in Fig. 2.5) encompasses the experimental noise, the readout noise as well as the injection noise and is therefore equivalent to σ. The set of parameters which minimizes our comparison criterion is found by using the downhill simplex minimization method (Nelder & Mead 1965). The free parameters in our simulation are the trap density, and the density distribution parameters: the SBC full well capacity, η_{BC}, and η_{SBC}. We introduce η in order to preserve in the electron density distribution the ratio between the signal confinement volume dimensions $\mathbf{x}_{max}$. In this way the standard width of the distribution $\sigma_{\mathbf{x}}$ in each direction corresponds to a fixed fraction of the predefined dimension of the signal confinement volume in that direction $\mathbf{x}_{max}$:

$$\sigma_{\mathbf{x}} = \eta \times \mathbf{x}_{max} \tag{2.31}$$

To avoid local minima, we first randomly probe the parameter space and establish a χ^2 map. The downhill simplex is then initialized with the set of parameters for which the χ^2 is the smallest. The grey line and circles in Fig. 2.5 (left) give an example of the model results, representative of the quality of fit that can be achieved using this method. The simulation parameters for this particular fit are summarized in Table 2.4. The grey bars indicate 1-sigma deviation from the average over the 8 simulation iterations. Within the error bars, the model reproduces the experimental data over a wide range of signal levels (three orders of magnitude). The CTI mitigation effect of the SBC and the transition between the BC and SBC signal regimes are particularly well handled. At higher signal levels the performance of the model slightly degrades. We notice a deviation that would ultimately lead to underestimate the charge loss at very high signal levels. It is not possible to confirm this tendency as the charge injections with a signal level higher than 17 ke$^-$could not be studied due to saturation of the output amplifier.

A comparison between our analytical description of the electron density distribution and the more detailed model of a Gaia pixel (Seabroke et al. 2009, 2010) shows that in the serial direction the Gaussian approximation may be too crude at high signal level, where the distribution saturates quickly and takes a box-like shape. However, when using Seabrokes model we were not able to reproduce satisfactorily the fractional charge loss measurement. The simulations deviated significantly from the experimental data in particular during the transition between the BC and SBC signal regimes. One possible explanation is that the single e2v Gaia CCD in question has a non-nominal pixel architecture due to manufacturing tolerances (Seabroke's model uses nominal e2v design values).

2.4.2.2 Astrometric images

The verification of the elementary units of the model being satisfactory as well as the validation of our electron density distribution approach, we are now interested in estimating the capabilities of the model to reproduce the CTI induced distortion of stellar images acquired in TDI mode.
Thanks to the second test of RC2 (cf. Section 2.3), the model performance can be investigated for different stellar brightnesses. To do so we compare the model outcomes with the stellar images acquired in the irradiated part of the CCD. This comparison is performed at the sub-pixel level thanks to damaged over-sampled stellar images built using the multiple scans performed for each set of experimental parameters (temperature, star brightness, CI delay). The model input signal must be representative of the CCD illumination conditions. Therefore, we used the accumulated data in the non-irradiated part of the CCD to create a reference undamaged image profile. In order to stay as close as possible to the original experimental set-up, the transit of the investigated star is simulated in two dimensions, i.e. over a virtual CCD with 12 pixel columns of 4494 pixels. This requires a two-dimensional input. However during RC2, the images were binned in the serial direction to recreate in-flight conditions such that the reference curves are LSFs (Line Spread Function, the PSF integrated in the serial direction). Since we have no information on the two-dimensional PSF, to generate a PSF from the original reference curve we assume that the profiles in the parallel and serial directions are the same: $P(x,y) = L(x) \times L(y)$, where L is the undamaged reference curve, x and y the positions in pixel respectively in the parallel and serial directions, and P the model input image. The integrated flux of the reference curve is scaled to produce an input image for each level of illumination (or artificial star magnitude).
To perform a direct comparison between the model and the over-sampled damaged profiles, we first extract the sampling scheme specific to each of the damaged profiles. We then apply it to the PSF generated from the reference curves to create the model input signal. In this way the required number of simulations for a single set of parameters to generate an n times over-sampled simulated damaged profile is n. Once the simulations are completed, the individual predictions are binned in the serial direction and each data point is then placed at the correct sub-pixel position according to the original sampling scheme so as to form an oversampled simulated damaged profile. We present the results of this detailed comparison (Fig. 2.6, Table 2.5) for two different cases: (i) the model is first used to reproduce a single image profile (i.e. for a unique set

Fit to	single damaged profile		different signal levels
Figure	2.6 a)	-	2.6 b)
Magnitude	13.67	17.03	13.67, 15.17, 17.03
CI delay	30s	30s	30s
Goodness-of-fit			
χ^2_{red}	15.4	14.4	70.0, 41.4, 18.2
Trap parameters			
ρ_1 [per pixel]	0.57	4.59	3.97
σ_1 [m^2]	$6.56\,10^{-21}$	$3.99\,10^{-21}$	$2.33\,10^{-21}$
τ_1 [s]	$2.23\,10^{-3}$	0.60	0.56
ρ_2	4.11	4.39	4.10
σ_2	$1.17\,10^{-20}$	$1.89\,10^{-22}$	$2.38\,10^{-22}$
τ_2	97.34	2.92	78.91

Table 2.5 — MC model fitting parameters corresponding to the examples shown in Figs. 2.6 a) and b). Note that χ^2_{red}, the reduced χ^2, indicates here the goodness-of-fit. In the three presented cases, a short and a long release time constant species are needed to explain the experimental data.

of experimental parameters: temperature, CI delay, and magnitude); (ii) and then, with a single set of simulation parameters, the model is used to simultaneously reproduce a set of damaged image profiles with different magnitudes.
When fitting to an individual profile (Fig. 2.6 (a)), χ^2 is used as our goodness-of-fit criterion:

$$\chi^2 = \sum_{i=0}^{S-1} \frac{\left(\lambda\left(x_i\right) - N\left(x_i\right)\right)^2}{\sigma^2} \tag{2.32}$$

where λ is the simulated damaged profile, N the RC2 profile, x_i a particular sub-pixel position, S the total number of data points. σ, the noise, is considered to be equivalent to the quadratic sum of the photon-noise and the readout noise r. The photon-noise is assumed to be Gaussian with a standard deviation of $\sqrt{N}$ (an approximation of the Poisson statistics for large photon counts) and r is assumed to have the constant value of 4.8 electrons:

$$\sigma^2 = N\left(x_i\right) + r^2 \tag{2.33}$$

When fitting to a set of damaged profiles (Fig. 2.6 (b)), χ^2 is altered in order to avoid any fitting bias towards the brightest magnitudes and the most over-sampled profiles. The new comparison criterion g is thus defined as follows:

$$g = \frac{\chi^2}{S \times F} \tag{2.34}$$

where F is the total integrated flux.
In each case the free parameters of the simulation are the trap parameters (ρ the density, σ the capture cross-section, and τ the release time constant). The number of trap

species is fixed to 2. The electron density distribution parameters are fixed to the values summarized in Table 2.4. Compared to the fitting of the fractional charge loss measurement, fitting astrometric measurements is a complex task:

1. The parameter space is larger.
2. The parameter space is more degenerate; a single image profile does not set high constraints on the trap parameters.
3. Both the capture and release processes are involved.
4. The illumination history plays an important role.
5. As we will explain in the following, the number of experimental uncertainties is higher as well as the number of assumptions required by the data processing and the simulation.
6. The fitting also necessitates a high number of two-dimensional simulations of stellar transits, which is computationally very intensive, and sets a high requirement on the minimization procedure efficiency.

The preliminary χ^2 maps also appeared to be hard to characterize as they contained a lot of local minima, none of which presented a satisfactory agreement with the data. In order to better sample the parameter space we decided to perform the first step of our minimization procedure by means of an evolutionary algorithm [1]. We applied two evolutionary mechanisms, mutation and cross-over, with optimal occurrence probabilities on a restricted initial population of parameter sets and let it evolve towards smaller comparison criteria. After a limited amount of generations, we initialized the downhill simplex algorithm with the set of parameters for which our comparison criterion was the smallest. As previously mentioned, simulating two-dimensional stellar transits is a relatively time-consuming process. For long charge injection delays (> 20s) one can neglect the effect of the charge injection release trail, the star then crosses a CCD with almost all the traps empty (from short to fairly long release time-constants). This enables us to drastically reduce the number of simulated TDI steps. Additionally the stellar transits acquired in these conditions offer the advantage of presenting important CTI effects that set a higher constraint on the fitted parameters. We thus selected the subset of RC2 damaged profiles with a CI delay of ~ 30 s to perform our comparison. Finally, although the profiles were ultimately compared at the sub-pixel level, in order to decrease the total number of simulations, only one sample per pixel was simulated during the fitting procedure

Figure 2.6 (a) presents an example of a simulated damaged profile (circles), this profile is representative of the best-fitting achievable in these particular conditions at a magnitude of 13.67 and should be compared to the black line profile (RC2 data). As can be seen from the residuals normalized by the photon-noise, the simulation is in fairly good agreement with the data over the whole profile. However we note a slightly larger disagreement in the image leading edge. This is not surprising as for bright images the wings still contain a fairly large amount of electrons (e.g., 800 e^-in the first simulated sample of the leading edge), which may play an important role in the self-illumination history by filling a number of traps and mitigating the CTI effects in the

1. `http://watchmaker.uncommons.org/`

image core. Hence the fit may be further improved by simulating a profile wider than 15 pixels. The best individual fits obtained at fainter magnitudes (cf. example in Table 2.5) shows a slightly better overall agreement. The relative amplitude of the CTI effects in the core compare to the wings is very sensitive to the input image shape in the CCD serial direction. The assumption made in order to build a two-dimensional input signal may thus also limit the ultimate goodness-of-fit.
Figure 2.6 (b) shows the resulting fits at three different magnitudes (13.67, 15.17, 17.03) obtained for a single set of simulation parameters (cf. Table 2.5). The overall agreement is remarkable. As can be seen in the bottom part of the plot, the amplitude of the CTI effects is qualitatively well reproduced as a function of the signal brightness. And as expected the charge loss peak slowly shifts towards the image core with fainter signal. However one should notice that when a single set of trap parameters is fitted to several profiles at different signal levels the resulting goodness-of-fit for each individual profile decreases, see Table 2.5. This is also illustrated in the bottom part of Fig. 2.6 (b), where the resulting charge loss curves at the magnitude 13.67 for the individual fit case (black dotted line) can be compared to the multiple magnitudes fit case (black dashed line). This difference in goodness-of-fit may be ascribed to variations in the experimental conditions between two tests with different signal level (slightly different temperatures and illumination background). Also the intervals between the tests were several days. In between two tests the CCD was stored at room temperature. Astrium later acknowledged that this may have resulted in a change of the CCD state, as discrepancies in test results were observed for similar experimental conditions. As already discussed in Prod'homme et al. (2010), this may set a limit on the ultimate goodness-of-fit achievable with a single set of trap parameters to several profiles by any model. As a final remark on this part of the comparison, we would like to state that even if the trap species parameters found as a result of the fitting procedure could be associated with known trap species from the literature, the uncertainties in the experiments are too large and the assumptions in the simulation process too numerous to conclude that such species were indeed present in the tested CCD. If one wants to infer trap species parameters by using this model to reproduce test data, the experiment should be carefully designed to serve this purpose. Artificial charge injections would be particularly suited for such an experiment as one can infer trap densities and capture cross-sections from the charge loss that occurs in their profile and release time constants from their trailing edge. We would in particular recommend: (i) to repeat the tests at different temperatures in order to break the potential degeneracies in the time release constant space; (ii) to use different levels of charge injection from a few electrons to the pixel full well capacity, in order to make sure that no trap is missed as well as to set the highest constraint possible on the charge density distribution model over a complete signal range.

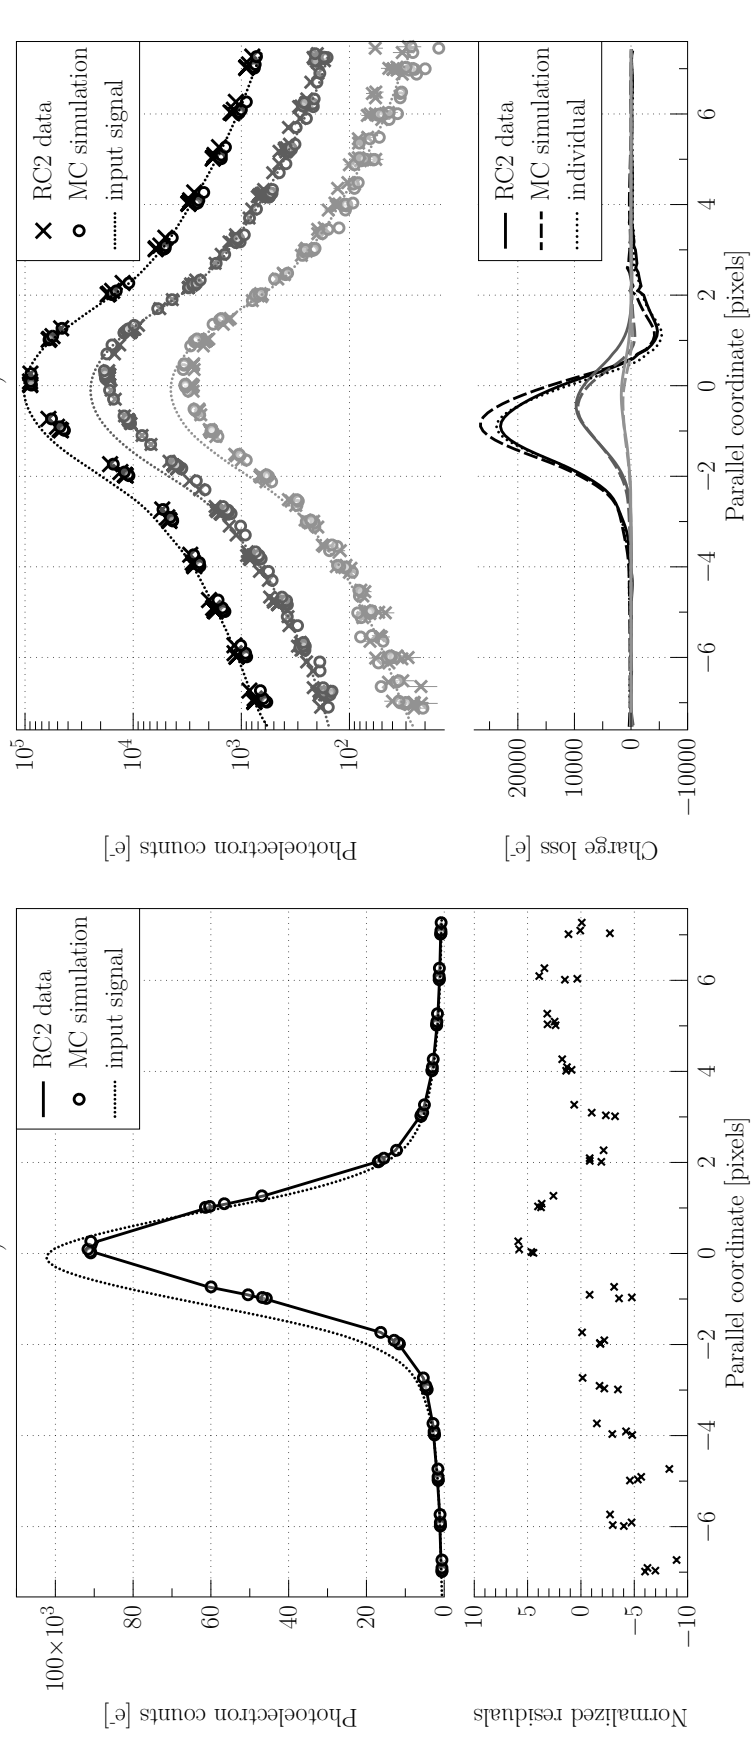


Figure 2.6 — (a) *Top left:* Comparison between the RC2 test data (black line) as reduced by Brown (2009b) and the model simulation (black circles) for a particular signal level and CI delay (Magnitude: 13.63; CI delay: ~ 30 s). The charge transfer direction is from right to left. The amplitude of the radiation damage can be appreciated by comparing the damaged profiles (black line and circles) to the input signal (dotted line) based on the non-irradiated test data. *Bottom left:* difference between the simulation and the experimental data normalized by the photon-noise. The quality of fit is similar in the profile core and wings. Note the clustering of the sampling: each data point within a pixel corresponds to a different scan.
(b) *Top right:* Comparison between the RC2 test data (crosses) and the model simulations (circles) for several signal levels and a particular CI delay (Magnitude: 13.63 (black), 15.17 (dark grey), 17.03 (light grey); CI delay: ~ 30 s). The model simulations are obtained for a single set of parameters (Table 2.5), which has been optimized to find a reasonable agreement with the test data. Note that this time the ordinate scale is logarithmic in order to facilitate the visual evaluation of the quality of fit for each magnitude and in the core as well as in the wings of the image profiles. The charge transfer direction is from right to left. *Bottom right:* Comparison between the resulting experimental (continuous line) and simulated (dashed line) charge loss as a function of magnitude. The dotted line corresponds to the resulting charge loss for the simulation parameter obtained when fitting to magnitude 13.63 only.

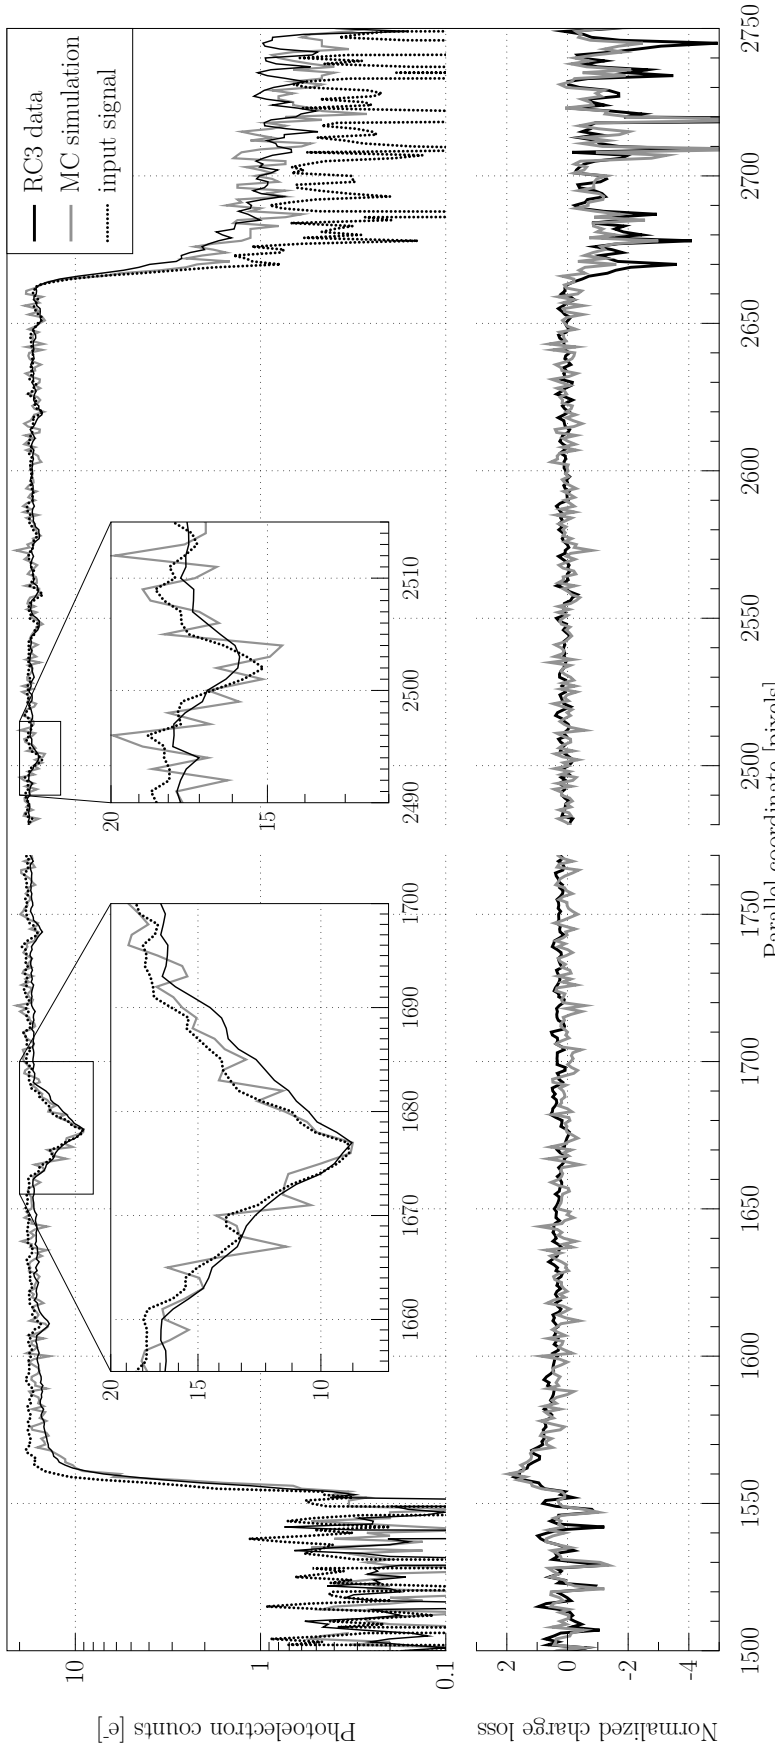


Figure 2.7 — *Top:* Comparison between the RC3 test data (black line) as reduced by Hall (2010) and the model simulation (grey line) for a star of magnitude 12.5 (No CI, and lowest level of background light ($< 1\,e^-$/pixel)). The charge transfer direction is from right to left. The reference averaged spectrum (acquired in the non-irradiated part of the CCD) was used as input signal (dotted). *Bottom:* The resulting charge loss normalized by the photon-noise for the simulation (grey) and the test data (black).

2.4.2.3 Radial velocity spectrometer images

The radial velocity of stars will be measured on-board Gaia thanks to a medium resolution spectrometer (Katz et al. 2004) that will enable the analysis of the Doppler shift of the stellar spectral lines, in particular the Ca II triplet. The spectrometer is composed of a series of prisms and a diffraction grating that will disperse the stellar light before it reaches the red-enhanced Gaia CCDs. The combined effect of TDI and high light dispersion will result in a very faint spectral signal, down to a few electrons per spectral feature. The signal-to-noise ratio for a single measurement is expected to be very low (e.g., 7 for a star of magnitude 16), and, at the faint end, the radial velocity measurement will rely on the co-addition of multiple spectra to recover the spectral features. Until recently CTI had never been studied at such low signal levels, and it was uncertain if one could indeed recover spectral features by co-adding damaged spectra. RC3 has been designed to address this particular issue by the use of a mask manufactured to reproduce in detail the spectrum of a G2V star. To illustrate the capacity of our model to reproduce the CTI effects in TDI mode at extremely low signal levels (< 20 e^-) we use the RC3 test data and compare the model outcomes to the stellar spectra acquired in the irradiated part of the CCD.

The model input signal is built from the accumulated data in the non-irradiated part of the CCD and the transits are simulated over a virtual CCD containing 6 pixel columns of 4494 pixels. During the tests the spectra were binned over 12 pixels in the serial direction, we thus again had to assume the serial signal profile to perform two-dimensional simulations. We assumed the same parallel spectral profile in every pixel column but with different intensities. The fractions of the total intensity for each pixel column were chosen according to the serial profile of the simulated spectrum provided by the Gaia Instrument and Basic Image Simulator, GIBIS (0.04, 0.14, 0.26, 0.29, 0.20, 0.07). This time we do not perform any optimization of the trap parameters but we selected two trap species from Seabroke et al. (2008a) (divacancy and unknown: $\rho_{1,2} =$ 1 trap/pixel, $\sigma_{1,2} = 5\,10^{-20}\,\mathrm{m}^2$, $\tau_1 = 20$ s, and $\tau_2 = 80$ ms). Once again we fixed the electron density distribution parameters to the values summarized in Table 2.4.

In these conditions, one can obtain a spectrum such as the one presented in Fig. 2.7 for a star of magnitude 12.5 (grey line). In a similar fashion to the depicted experimental damage spectrum (black line), 15 transits were averaged to obtain the simulated spectrum. In both the RC3 data and the model outcome, one can notice that even at such low signal level the CTI effects do not completely wash away the spectral features; they can be recovered by co-adding several spectra. However, as can be seen in more detail from the blown-up portions of the presented spectra, CTI lowers the continuum and affects the absorption features by increasing their width and shifting them. These effects lead to a significant decrease of the overall spectral contrast. In the bottom part of the figure, we show the charge loss normalized by the photon-noise. The model reproduces remarkably well the first significant charge loss bump (only few electrons) occurring at the leading edge of the spectrum (left). The electron release that occurs at the trailing edge after the spectrum transit (right) is also well reproduced. The simulated averaged profile is binned in the serial direction over 6 pixels instead of 12 during the tests, as a consequence the simulation is slightly noisier than the RC3 data. It is thus

hard to distinguish, but within the spectrum itself one can notice that the model does widen and shift the spectral features (see the blow-ups). The amplitude of the shift and increase in width is not quite matched though, this could be imputed to a difference in the release time constant for the trap species.

2.5 Conclusion

We have described a physical Monte Carlo model that simulates CTI effects induced by radiation damage in astronomical CCDs at the electrode level. This model has been elaborated in the challenging context of ESA's Gaia mission. The operating conditions, the extreme astrometric requirements and a novel CCD pixel architecture, necessitated the development of a new approach in the representation of the electron density distribution as well as a more detailed description of the trapping probabilities. These new features have been combined in a comprehensive simulation of CCD charge collection and transfer at the electrode level. In order to verify the model predictions, we first validated the unit blocks of the simulation and then proceeded to a detailed comparison between the model and experimental test data. We showed that the model is able to accurately reproduce the CTI effects for a wide range of signal levels down to a few electrons, hence validating our electron density distribution representation. We finally validated the global model operation by assessing the model capability to reproduce the CTI induced distortion on stellar images and spectra acquired in TDI mode at different magnitudes.

The model elaboration contributed greatly to our present understanding of CTI effects. In particular, by showing that the implementation of a density driven approach of the electron packet growth enables the reproduction of experimental data and that to be successful in modelling the CTI effects at very faint signal levels, no detail should be neglected. The simulations should be as realistic as possible, down to the transfer of electrons at the electrode level and the simulation of each individual trap.

Due to the complexity of the CTI effects (including the varying illumination history) and the extreme accuracy required, in particular on the estimation of the stellar image location, one cannot apply a conventional correction of the raw data. Even if this were possible, the process would likely be unstable and lead to poorly understood error propagation. Instead of a direct correction, the Gaia Data Processing and Analysis Consortium (DPAC) is developing a scheme which relies on a forward modelling approach that enables the estimation of the true image parameters from a damaged observation (e.g., Prod'homme 2011). In this scheme each observation is ultimately compared to a modelled charge profile in which the distortion of the CTI-free image (PSF) will be simulated through an analytical CTI model. The Monte Carlo model described in this chapter is used to generate large synthetic data sets of both damaged and damage-free observations. In the next two chapters, these simulations are used to re-assess the final performance of the mission, taking the effects of radiation damage into account, as well as to verify the DPAC CTI mitigation scheme.

Although developed in the Gaia context, we tried to keep the model as general and as flexible as possible. It can be used to simulate any kind of measurements performed with a CCD operated in imaging mode or TDI mode. Different clocking schemes can

be applied and the description of the electron density distribution should be flexible enough to simulate different pixel architectures. The use of this model is particularly relevant in the frame of space experiments that aim at very accurate measurements at low signal levels. ESA missions under study Euclid and Plato, the astrometric measurements performed on board HST or future X-ray missions sent to L2 may benefit from the use of such a model to evaluate the impact of radiation damage on their performance budgets.
A running version of the model as well as a brief documentation and a few examples are readily available at `http://www.strw.leidenuniv.nl/~prodhomme/cemga.php` as part of the CEMGA java package (CTI Effects Models for Gaia) developed at Leiden Observatory. Please contact the authors for more information on how to use the package.

Acknowledgments

The authors would like to acknowledge EADS Astrium for kindly supplying the experimental test data necessary to accomplish the verification of our model. We would like to thank F. van Leeuwen from the Institute of Astronomy in Cambridge, and D.J. Hall from the e2v centre for electronic imaging at the Open University for providing the results of their reduction and analysis of the test data to support the model verification effort. This model and study greatly benefited from the numerous discussions with all the participants of the Gaia Radiation Task Force in particular M. Weiler (Observatoire de Paris-Meudon), B. Holl (Lund Observatory), G.M. Seabroke (MSSL), R. Kholey (ESA-ESAC), and F. Raison (ESA-ESAC). This work has been funded by the European Community's sixth framework programme (FP6) through the Marie Curie research training network ELSA (European Leadership in Space Astrometry) contract No MRTN-CT-2006-033481.

Chapter 3

The impact of CCD radiation damage on Gaia astrometry: I. Image location estimation in the presence of radiation damage

The Gaia mission has been designed to perform absolute astrometric measurements with unprecedented accuracy; the end-of-mission parallax standard error is required to be of the order of 10 micro-arcseconds for the brightest stars ($V \leq 10$) and 30 micro-arcseconds for a G2V type star of magnitude 15. These requirements set a stringent constraint on the accuracy of the estimation of the location of the stellar image on the CCD for each observation: e.g., 0.3 milli-arseconds (mas) or 0.005 pixels for the same $V = 15$ G2V star. However the Gaia CCDs will suffer from charge transfer inefficiency (CTI) caused by radiation damage that will degrade the stellar image quality and may degrade the astrometric performance of Gaia if not properly addressed. For the first time at this level of detail, the potential impact of radiation damage on the performance of Gaia is investigated. In this chapter we focus on the evaluation of the CTI impact on the image location accuracy using a large set of CTI-free and damaged synthetic Gaia observations supported by experimental test results. We show that CTI decreases the stellar image signal-to-noise ratio and irreversibly degrades the image location estimation precision. As a consequence the location estimation standard errors increase by up to 6% in the Gaia operating conditions for a radiation damage level equivalent to the end-of-mission accumulated dose. We confirm that in addition the CTI-induced image distortion introduces a systematic bias in the image location estimation (up to 0.05 pixels or 3 mas in the Gaia operating conditions). Hence a CTI mitigation procedure is critical to achieve the Gaia requirements. We present a novel approach to CTI mitigation that enables, without correction of the raw data, the unbiased estimation of the image location and flux from damaged observations. We show that its current implementation reduces the maximum measured location bias for the faintest magnitude to 0.005 pixels ($\sim 4 \times 10^{-4}$ pixels at magnitude 15) and that the Gaia image location estimation accuracy is preserved. In the next chapter of this thesis we will investigate how the CTI effects and CTI mitigation scheme affect the final astrometric accuracy of Gaia by propagating the residual errors through the astrometric solution.

T. Prod'homme, B. Holl, L. Lindegren, A.G.A. Brown
Accepted for publication in Monthly Notices of the Royal Astronomical Society, 2011

3.1 Introduction

Gaia is a European Space Agency mission that aims to create the most complete and accurate stereoscopic map to date of the Milky Way, containing parallaxes, proper motions, radial velocities, and astrophysical parameters for one billion stars, one percent of the estimated stellar population in our galaxy (Perryman et al. 2001; Lindegren et al. 2008). Due to the satellite's constant spinning motion, the determination of the astrometric parameters ultimately comes down to measuring very precisely the time $t_{\rm obs}$ at which a particular star crosses a fiducial line on the focal plane (Lindegren & Bastian 2011; Bastian & Biermann 2005). The required astrometric precision is extreme, e.g., the end of mission parallax uncertainty for a star of magnitude $V = 15$ is required to be better than 25 micro-arcseconds (μas). In order to determine $t_{\rm obs}$, one needs to measure the image location on the Charge-Coupled Device (CCD) relative to the instrument axes. As a consequence the required astrometric accuracy sets a direct and stringent requirement on the residual image location uncertainty per CCD star transit. In the left part of Table 3.1 we detail the end-of-mission parallax standard error, σ_ϖ, as function of stellar magnitude and type [1] computed using de Bruijne (2009b). These predicted standard errors do include the increased photon noise due to the radiation damage induced charge loss, but not the residual bias-calibration errors considered in the present chapter (except for a general contingency margin of 20%). For this chapter we are interested in the mean image location uncertainty per CCD star transit σ_κ that would be needed to reach a given targeted parallax accuracy. Based on de Bruijne (2005) we estimate the corresponding 'requirement' on the image location uncertainty shown in the right part of Table 3.1 and computed as: $\sigma_\kappa = \sqrt{N_{\rm obs}}\,\sigma_\varpi/(m\,g_\varpi)$, with $N_{\rm obs}$ the average number of astrometric observations per star (662), m the end-of-mission scientific contingency margin which is 1.2, and g_ϖ the geometrical parallax factor which is 2.08 for the Gaia solar aspect angle $\xi = 45^\circ$. This formula has also been used to compute the 'requirement' curve as function of Gaia G-band [2] magnitude shown in several figures throughout this chapter. No spectral type distinction is needed when these uncertainties are expressed in G because they are virtually independent of spectral type. Note that the computed location uncertainties do not contain the 20% contingency margin, making them very stringent.

During the 5 year mission life time, solar wind protons will collide with Gaia's focal plane and create electron traps in the CCDs by displacement damage. These radiation-induced traps drastically increase the CCD charge transfer inefficiency (CTI) and will lead to a significant loss of signal for all Gaia measurements by stochastically capturing and releasing signal electrons. The resulting electron redistribution will also distort each stellar image. The CTI effects are expected to significantly contribute to the error budget of all the Gaia measurements (astrometric, photometric, and spectroscopic), especially if not properly taken into account in the data processing.

We present here the first part of a detailed evaluation of the impact of the radiation

1. Updated estimates of the science performance are given on: `http://www.rssd.esa.int/index.php?project=GAIA&page=Science_Performance`

2. The Gaia G-band magnitude is a broad-band, white-light magnitude in the wavelength range 300 – 1000 nm defined by the telescope transmission and CCD quantum efficiency. $G = V$ for an un-reddened A0V star (Jordi et al. 2010; Perryman et al. 2001).

	Parallax accuracy target standard error σ_ϖ [μas]			Corresponding CCD image location uncertainty σ_κ [mas]		
Type	B1V	G2V	M6V	B1V	G2V	M6V
$V-G$	0.03	0.16	2.18	0.03	0.16	2.18
$V=10$	6.9	7.0	7.3	0.072	0.073	0.076
$V=15$	25	24	10	0.26	0.25	0.11
$V=20$	322	300	102	3.3	3.1	1.1

Table 3.1 — On the left side we tabulate the Gaia requirements of the sky averaged end-of-mission parallax standard error, σ_ϖ (in μas), as function of spectral type and Johnson V-band magnitude (see de Bruijne 2009b). $V-G$ allows for the conversion from V to the Gaia G-band magnitude. On the right we give the corresponding mean CCD image location uncertainty, σ_κ (in mas), that would result in the parallax standard error on the left when observed 662 times (the average number of astrometric observations per star). See Section 3.1 for more details.

damage effects on the final accuracy of the Gaia astrometric measurements. This chapter focuses on the effect of CTI on the image location accuracy. Studying the accuracy of a measurement is a rather complex enterprise. Hence we present in the following section the overall applied methodology and the different steps of this study. In the next chapter of this thesis, we will investigate how the final Gaia astrometric accuracy is affected by the CTI-induced errors at the image processing level as characterized in this study.

3.2 Overall methodology

The use of synthetic data

To evaluate the impact of the CTI effects on the image location accuracy, we apply the Gaia image parameter estimation procedure (Section 3.4) to a large data set of simulated CTI-free and CTI-affected observations (from here on the latter are referred to as 'damaged observations'). The use of synthetic data presents several fundamental advantages compared to the use of experimental data: while in experimental studies, the true image parameters, the instrument model or Point Spread Function (PSF), and the different noise contributions need to be estimated, in the simulation, these are known parameters. Hence the uncertainties related to the estimation of such parameters cannot bias the result of our study. Furthermore only simulation can allow the determination of the absolute image location bias and the associated standard errors as this requires the knowledge of the true image location. Finally, by using synthetic data one can compute the statistical uncertainties on the measured image location bias and standard errors, and this at any precision level just by increasing the number of simulated observations for a particular set of conditions. In Section 3.3, we detail the simulation of Gaia-like observations in the absence and presence of radiation damage for different stellar magnitudes, image widths, background and readout noise levels.

The bias allows for the quantification of the trueness of an estimation, and the standard errors for the quantification of the estimator precision. If an estimator delivers bias-free estimations, then its standard errors can also be regarded as a means to quantify

the accuracy of this estimation. In the following we will thus make the important distinction between precision and accuracy when it is justified.

The Gaia image parameter estimation procedure
In the Gaia data processing, the image location and flux are estimated or 'self-calibrated' through the use of an iterative procedure, that allows for the successive determination and improvement of the PSF and the image parameters without prior knowledge. A detailed description of this procedure is provided in Section 3.4. In Sections 3.5.3 and 3.5.4, the procedure is applied to the data set of simulated CTI-free observations, in order to verify that, in the absence of CTI, the Gaia image parameter estimation procedure performs efficiently, according to expectations.

The Cramér-Rao bound
Assessing the efficiency of the Gaia image parameter estimation procedure necessitates the computation of the theoretical limit to the image location accuracy in the Gaia observing conditions. This theoretical limit corresponds to the ultimate accuracy achievable by any bias-free estimator. It is set by the Cramér-Rao bound, described in Section 3.5.1. We thus compute the Cramér-Rao bound as a function of magnitude (G), image width, background, and readout noise level (Section 3.5.4) and subsequently use it as a reference in a comparison with the standard errors of the estimated image parameters. The Cramér-Rao bound is also required to assess the impact of the CTI effects independently from any estimation procedure.

The radiation damage impact on the image location estimation
In Section 3.5.5, we use the set of damaged observations to demonstrate that the image distortion and the charge loss induced by CTI imply an irrevocable loss of accuracy in the image location determination. This loss of accuracy, which directly affects the performance budget of Gaia, is independent from any estimation method and can only be avoided by physically preventing charge trapping. This is done by optimizing the hardware (e.g., the CCD operating temperature) and using hardware countermeasures such as the periodic injection of artificial charges, or the use of a supplementary buried channel, an extra doping implant in each pixel that confines small charge packets. Taking into account these countermeasures in the simulation of the damaged observations allows us to verify that they indeed substantially contribute to diminish the CTI effects on the theoretical image location accuracy. However, Sections 3.5.6 and 3.5.8 show that it is not possible to rely solely on these countermeasures: indeed we find that if the CTI effects are not properly taken into account in the image parameter determination the image location bias can be as large as 10 mas for a star of magnitude 15 (to be compared to the requirement of $\sim$0.08 mas, see Table 3.1). This is in agreement with the experimental tests performed on Gaia irradiated CCDs (see Section 3.5.7) and confirms that the CTI effects need to be addressed by the Gaia data processing in order to achieve the mission requirements.

Mitigating the CTI effects

The software mitigation of CTI effects is a complicated task. Several schemes have been discussed in the literature; they usually imply the direct correction of the raw data in the context of photometry-based measurements. In Section 3.6.1, we review the different potential schemes, and where they intervene in the data processing chain. Then, in Section 3.6.2, we present and motivate a novel approach to CTI mitigation that does not involve a direct correction, so that the noise properties of the raw observations remain unchanged. This approach, developed by the Gaia Data Processing and Analysis Consortium (DPAC, Mignard et al. 2008), relies on the forward modelling of each observation including the CTI distortion, such that the true image parameters can be directly estimated from the damaged observations. The modelling of the distortion of the stellar image is performed thanks to a fast analytical CTI effects model, a so-called Charge Distortion Model (CDM). The success of this CTI mitigation approach depends on the performance of such a model. In Section 3.6.3, the potential accuracy of this approach is assessed in the case of an ideally calibrated CDM. Finally, using the current best CDM candidate (Short et al. 2010; Prod'homme et al. 2010,and Section 3.6.4), we apply the image parameter estimation procedure and the DPAC CTI mitigation scheme to the set of damaged observations and show that one can recover the CTI-induced image location and flux bias (see Sections 3.6.6, 3.6.7). Only then are we able to answer the question: does the current Gaia image location procedure combined with the presented CTI mitigation scheme allow an unbiased estimation of the image location with a sufficient accuracy in the presence of radiation damage?

Methodology summary

1. Generation of CTI-free and damaged Gaia-like observations.
2. Determination of the theoretical limit to the image location accuracy in the absence of CTI by computing the Cramér-Rao bound.
3. Performance assessment of the Gaia image parameter estimation procedure in the absence of CTI.
4. Evaluation of the intrinsic loss of accuracy in the image location estimation induced by radiation damage by computing the Cramér-Rao bound for a damaged LSF.
5. Characterization of the CTI effects on the image parameter estimation procedure.
6. Performance assessment of the Gaia image parameter estimation procedure in the presence of CTI and including a forward modelling approach to mitigate the CTI effects.

3.3 Generating Gaia-like observations

In the following, we first describe the main principles of the Gaia observations, and explain how we generate Gaia-like reference images. These images are used to simulate thousands of observations for different stellar magnitudes and operating conditions, and thus constitute the basis of our study. Then we detail how we simulate the stellar

transits over a CCD. To achieve a high level of realism, we use a physically motivated Monte Carlo model that simulates the CCD charge collection and transfer, as well as the trapping processes, at the pixel-electrode level (see Chapter 2). Finally we summarize the expected CTI effects on the stellar images in the Gaia operating conditions and explain the choices we made regarding the radiation damage parameters of the simulation (trap species, level of radiation).

3.3.1 How Gaia observes

The Gaia spacecraft will orbit around the second Lagrange point (L2) and constantly spin around its own axis such that its two telescopes scan a great circle on the sky several times a day. The precession of the spin axis changes the orientation of the consecutive great circles, allowing for the coverage of the whole sky in about six months. The measurements are recorded in a single focal plane consisting of 106 CCDs. Due to the satellite spinning motion, the star projections will not remain stationary during an observation but will transit the focal plane in the along-scan (AL) direction. The orthogonal direction is called across-scan (AC). To integrate the stellar images along the star transits, the CCDs are operated in time-delayed integration (TDI) mode. In this mode the CCD is constantly readout and the satellite scanning rate (and induced light source motion) has been synchronized with the charge transfer period, so that the charge profile continues to build up as the image travels across the CCD avoiding as much as possible image smearing. The charge transfer period is 0.9892 ms and the integration time 4.4 s. The observing principle of Gaia relies on differential positional measurements among the stars simultaneously visible in the two superposed fields of view. In particular the differential measurements *between* the two fields of view (covering arcs of about 106.5° on the sky) are essential for the construction of a global reference frame and the determination of absolute parallaxes. For these measurements the AC component of the differential positions is largely degenerate with the instrument pointing, and mainly the AL component matters (Lindegren & Bastian 2011). Gaia is therefore primarily optimized for AL measurements and the image location accuracy in that direction is the most critical one for the performance.

Because of limitations on the amount of data that can be sent to ground only a small window around each source is read out. For sources fainter than magnitude 13 the windowing scheme of the astrometric instrument is simple. In the AL direction, the window size is 12 pixels for stars brighter than magnitude 16, and 6 pixels for fainter stars. In the AC direction the size of the on board readout windows is 12 pixels (the pixel size is 58.9 mas AL and 176.8 mas AC). The observations are then binned across-scan before being sent to the ground. Hence, for a particular star, a Gaia observation results in a one-dimensional along-scan set of electron counts that corresponds to the sampling of a one-dimensional point spread function or Line Spread Function (LSF). Note that Gaia will observe sources as bright as $G = 5.7$, but due to the relatively low number of stars between magnitude 5.7 and 13 ($\sim$1% of the expected 10^9 sources) and because of the use of a complicated gating scheme (to avoid pixel saturation) these magnitudes are ignored in this study.

3.3.2 Construcing a Gaia-like reference image

In this study we only analyze observations of point sources because these will be used in the Astrometric Global Iterative Solution (AGIS, Lindegren et al. 2011) to calibrate the instrument and satellite attitude, which are then used to estimate the astrometric parameters for all sources. The PSF, the actual two-dimensional flux distribution of an unresolved star that illuminates a CCD, depends on the spectral energy distribution of the star, the CCD properties, and the optics and its associated wavefront errors. Because the spectral energy distribution depends on the type of star, and the wavefront errors depend on the position in the focal plane, we cannot construct a single PSF that would be representative of the whole range of possible profile shapes.

To construct a set of PSFs that samples the range of possible profile shapes we make use of a study by Lindegren (2009) in which a set of 20,000 one-dimensional line spread functions (LSFs) were generated that are representative of the Gaia optics and wave front errors, stellar spectral energy distributions, and CCD effects (e.g., smearing due to TDI). A set of basis functions was extracted from this dataset using a Principle Component Analysis (PCA) in order to describe the data with a minimum number of parameters. A set of ten basis functions was found to be sufficient to represent any of the LSFs with an RMS error of 10^{-4}. To be readily usable the basis functions were subsequently fitted by a special quartic spline [3] that is flexible enough to fit the data while being smooth enough to avoid overfitting that could result in sub-pixel position bias when the function is used for location estimation on pixel sampled data.

For our study we decided to use three reference images. First of all we construct a reference profile based on a selection of all profiles that have a FWHM within $\pm$ 1% of the mean FWHM of 1.958 pixels. These profiles are subsequently symmetrized (i.e. $L_{symm} = (L(x) + L(-x))/2$), averaged, and then fitted with the first four even basis functions, resulting in a symmetric mean profile with FWHM 1.957 pixels, from hereon referred to as the 'typical' reference profile. Four components are used to get a profile that is sufficiently close to the target mean FWHM. To represent the extremes we introduce a 'narrow' and a 'wide' reference profile. The narrow and wide profiles are constructed in exactly the same way as the typical reference profile, only differing in the selection of LSF samples, which are: 90 $\pm$ 1% and 110 $\pm$ 1% of the mean FWHM respectively. This results in a FWHM of 1.767 pixels for the narrow, and 2.161 pixels for the wide reference profile. All three reference images are normalized as shown in Fig. 3.1.

Although the two-dimensional PSF in the Gaia focal plane is wider in the across-scan direction than in the along-scan direction, the pixels are shaped such that in pixel units the PSF is nearly identical in both directions. Because we are only interested in how the pixels are illuminated we can therefore construct the two-dimensional reference image

3. The special quartic spline used to represent the LSF can be described as the convolution of an ordinary cubic spline, defined on a regular knot sequence with a knot separation of half a pixel, with a rectangular function of width equal to one pixel. This spline has the property that the sum of points sampled at one pixel separation is independent of the sub-pixel phase of the sampled points. Any 'effective' LSF, being the result of an optical LSF convolved with the pixel response function (Anderson & King 2000), should have this property. Details of the special quartic spline are found in a technical note by Lindegren (2003).

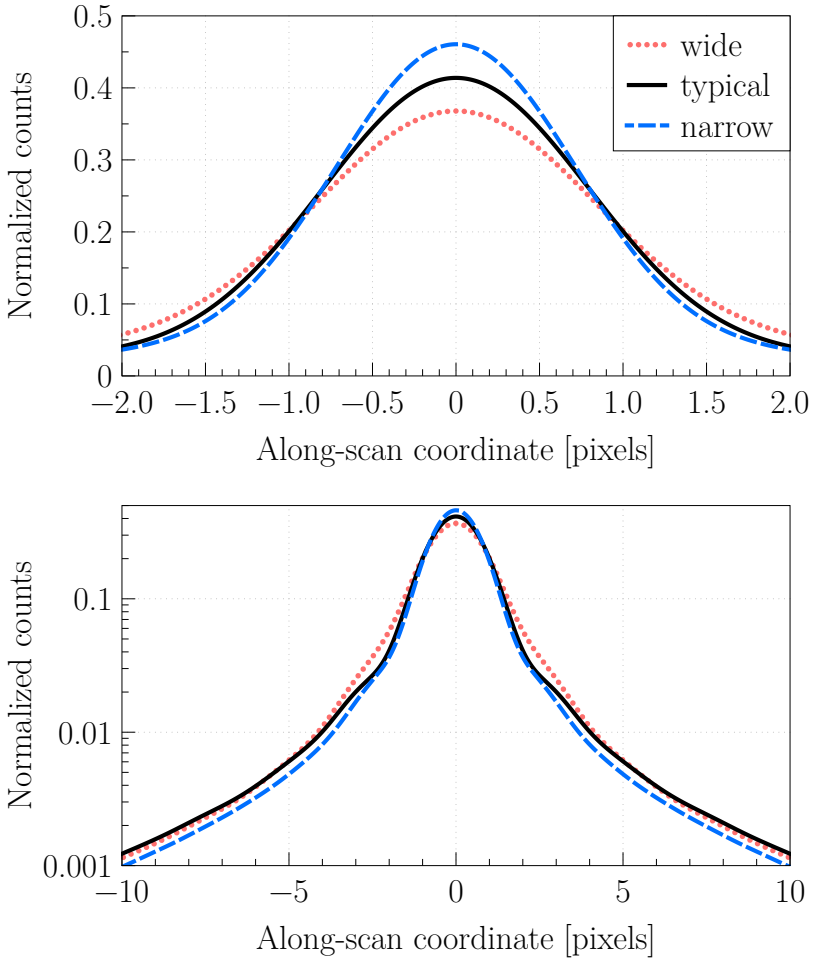


Figure 3.1 — One-dimensional Gaia-like reference images generated from a symmetrized subset of realistic Gaia-like line spread functions (LSFs). These reference images are used to reflect the range of narrow, typical, and wide profiles that results from different stellar types and wave front errors. The bottom figure shows the same profiles as the top figure for a wider AL coordinate range and with a logarithmic ordinate scale.

from simply multiplying the one-dimensional reference image $L(x)$ in two dimensions:

$$P(x, y) = L(x - \kappa) \times L(y - \mu) \tag{3.1}$$

In all our further analyses we assume $\mu = 0$. Because we defined the zero-point of the symmetric profile L to be at the symmetry point, this means that the PSF is always in the centre of the window in the across-scan direction.

3.3.3 Monte Carlo simulations of observations

The approach we have chosen for this study is to simulate a fully synthetic dataset using a detailed physical simulation of the photoelectron collection and transfer in CCDs at the pixel-electrode level (Chapter 2), available through the CEMGA software package [4] (Prod'homme 2011). The model also allows for a detailed treatment of radiation

4. www.strw.leidenuniv.nl/~prodhomme/cemga.php

induced traps that capture and release electrons and thereby distort the charge profile transferred through the CCD (see Section 3.3.4). The observations are simulated in two dimensions: 4494 pixels in along-scan and 12 pixels in the across-scan direction. In the software we illuminate the CCD with a two-dimensional reference image $P(x, y)$ described in Section 3.3.2. The normalized reference image is scaled to produce an illumination that corresponds to a particular stellar magnitude. The photon detection is modelled as a Poisson process: at each transfer step, the photoelectrons are generated using a random generator with a Poisson distribution and a mean equal to the expected number of collected photons (given by the reference image) within the integration time ($\sim$1 ms $\times$ 4494 pixels in the case of Gaia). Note that we control the exact along-scan location of the reference image, therefore allowing us to determine the exact error when a location estimate from the observation has been made. In the simulations we can optionally include a constant background. The electron packet transfer in the readout register is not simulated.
The raw two-dimensional observation counts are cropped in along-scan direction to 80 pixels (centred around the signal) and the resulting 80×12 pixels are stored. All used reference images are zero for $|x| > 20$, therefore any relevant signal is always contained in the cropped raw observation data.
When processing an observation we load the raw two-dimensional pixel counts. Depending on the windowing scheme for this particular magnitude and CCD we crop the data around the signal to the relevant window size and optionally bin the pixel counts in the across-scan direction resulting in a one-dimensional sample of the transit photoelectron counts $\{N_k\}$. Readout noise can be added to the counts using a normally distributed random generator Normal$(0, r^2)$: having zero mean and standard deviation r (the readout noise value).
Our total synthetic observational dataset consists of:

(i) 3 different CCD states: CTI-free, damaged with 1 trap pixel^{-1} and damaged with 4 traps pixel^{-1} (see Section 3.3.4 for details about the damaged cases),

(ii) 3 different reference images: narrow, typical and wide (see Section 3.3.2),

(iii) 2 different levels of sky background: 0 and 0.44698 e$^-$pixel^{-1}s^{-1} (the latter corresponding to the average sky surface brightness),

(iv) 9 different magnitudes: G= 13.3, 14.15, 15.0, 15.875, 16.75, 17.625, 18.5, 19.25, 20.0,

(v) the (two-dimensional) photoelectron counts of 250 CCD transits, each with the reference image incrementially shifted by $1/250^{\text{th}}$ of a pixel in the along-scan direction.

In almost all of the processing we select a unique combination of (i), (ii), (iii) and magnitude (iv), containing all 250 transits. The selection of all across-scan binned transits for a given magnitude is denoted as: $\{\{N_k\}\}_G$.

3.3.4 Simulation of the CTI effects

At L2 the radiation environment is dominated by energetic protons emitted during solar flares. The proton fluence is thus governed by the cyclic activity of the Sun which is usually monitored through sunspot counts. According to the latest predictions (see

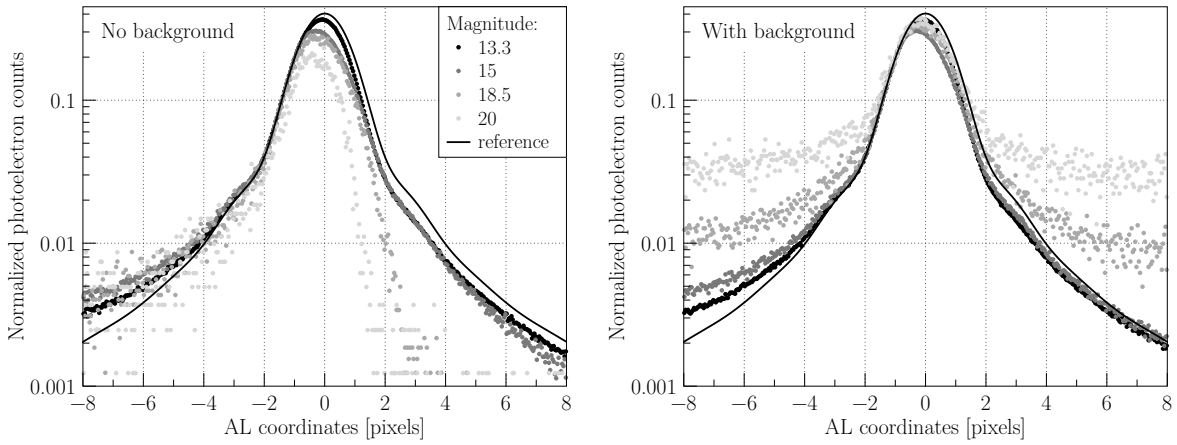


Figure 3.2 — The two plots show the simulated damaged observation resulting from multiple transits over an irradiated CCD containing 4 traps pixel^{-1} at different magnitudes with no background (left) and a background set to the sky brightness (right). In both cases the typical reference image is used and no readout noise has been added. The transfer direction is from left to right. The ordinate scale corresponds to the photoelectron counts normalized by the total expected CTI-free flux at each magnitude. Only 25 observations per magnitude (out of 250) are shown. The reference image (line) enables us to appreciate the variation of amplitude in the CTI induced distortion and charge loss as a function of the signal level. Also the overall profile centroid shift is clearly visible. Note that, as one can observe from the left plot, a radiation dose leading to 4 traps pixel^{-1} translates into a severe distortion of the Gaia measurements, particularly for fainter stars. The expected end-of-mission dose, based on the latest prediction of the next solar cycle, is however three quarters of the one we chose to simulate. Moreover the right plot shows that even a low level of background ($\sim$1 e$^-$pixel^{-1}) strongly mitigates the CTI effects by filling the traps prior to the star transit.

e.g, SIDC-team 2011), the next peak of activity will occur during 2013 coinciding with the launch of Gaia. Using the JPL 1991 model, the reference interplanetary proton fluence model by Feynman et al. (1993), taking into account the satellite design, and assuming 4 years of operation during the solar maximum (and one year during minimum), the average accumulated radiation dose received by a CCD of the astrometric instrument is predicted to be $\sim 3{\times}10^9$ (10 MeV equivalent) protons cm^{-2}. These protons will collide with and displace atoms in the Gaia CCD silicon lattice, and lead to the creation of electron traps. These traps stochastically capture and release the electrons transferred in the CCD. For more information concerning the trapping processes see Chapter 2 and references therein. The traps originate from different chemical complexes generally referred to as a trap species: a summary of the expected trap species in the Gaia CCDs is provided by Seabroke et al. (2008a) and Hopkinson et al. (2005). One usually distinguishes between trap species with short and long release time constants relative to the characteristic trap-electron interaction time ($\sim$1 ms for the Gaia CCDs), as they have different effects on the measurements. The traps with short release time constants capture electrons from the image leading edge and redistribute them within the telemetry window, which induces a distortion of the charge profile. The traps with longer release time constants capture electrons from the stellar profile and release them outside the telemetry window, which implies a charge loss that reduces the signal to noise ratio (see Fig. 3.2). The Gaia CCDs comprise two hardware CTI mitigation tools:

a charge injection (CI) structure and a supplementary buried channel (SBC). The CI structure is located all along the first CCD pixel row; it is composed by a diode capable of generating artificial charges and a gate that controls the number of electrons to be injected in the first pixel row and subsequently transferred across the whole CCD. Charge injections temporarily fill a large fraction of the traps present in the CCD and effectively prevent the trapping of the following generated and transfered photoelectrons. The SBC is a second and narrower doping implant on top of the buried channel. It creates a deeper potential that disappears into the shallower but wider buried channel for charge packets larger than 1500 e^-. By concentrating the electron distribution into a smaller volume it minimizes the electron-trap interactions in the rest of the pixel volume, effectively reducing the fraction of trapped electrons at low signal levels ($\leq$ 1500 e^-or $G > 15$).

In order to obtain representative results from our study, it is critical to achieve a high level of realism in the simulation of the CTI effects on each observations. This is why we make use of the most detailed CTI effects model to date (Chapter 2) verified against experimental tests performed on Gaia irradiated CCDs. At each transfer step, this model simulates the capture and release of electrons by computing for each trap the capture and release probabilities according to the trap characteristics and the local electron density distribution taking into account the Gaia pixel architecture and in particular the presence of the SBC. In the following we detail the considerations that led us to choose to simulate a unique trap species and two different radiation levels.

During the mission CI will be performed at periodic intervals. This means that most of the traps with release time constants greater than the injection period will be permanently filled, as only a very small fraction of them will have the time to release an electron. The current most likely value to be selected for the injection period is 1 s. If one neglects the serial CTI effects (occurring during the charge transfer in the CCD readout register), and take for reference the trap species as presented in Seabroke et al. (2008a), the only trap species that remains significantly active corresponds to the so-called 'unknown' with a release time constant $\tau \sim 90$ ms at the Gaia operational temperature and a capture cross-section $\sigma = 5 \times 10^{-20}$ m^2. We thus decided to generate the damaged observations using a virtual irradiated CCD containing a unique trap species, with these parameters. Note that the release and capture time constants vary exponentially with the temperature. The temperature over the entire Gaia focal plane is expected to deviate at most 5 K from the nominal operating temperature. This means that for different CCDs the effect of a single trap species will be different. However the temperature variation over a single device is expected to be negligible, hence our assumption regarding a single trap species with a unique release time constant still holds.

The traps filled by the CI, release their electrons and induce a characteristic 'release' trail after the CI. This trail changes the uniform nature of the background and must be carefully taken into account in the background estimation procedure. In order to prevent the CI background estimation and removal from affecting the results of our analysis, no CI was performed before the stellar transit during the simulations. Hence, the trap density has to be carefully selected to reproduce the trap density as it is perceived by a star after a CI delay (or time since last CI) comparable to the CI period. In

this way, the simulated amplitude of the CTI effects corresponds to the one observed in the experimental tests with CI performed in similar conditions.
In a series of four different campaigns of experimental tests carried out on irradiated Gaia CCDs, the prime contractor for Gaia, EADS Astrium, investigated the performance of potential hardware mitigation tools and characterized the trend and amplitude of CTI effects on Gaia-like measurements. These campaigns are referred to as radiation campaigns (RC). The RCs were performed in simulated Gaia operating conditions: a CCD operated in TDI mode at a temperature of 163 K with a low level of background light. The devices were irradiated at room temperature with a radiation dose of 4×10^9 protons cm^{-2} (10 MeV equivalent) that corresponds to an upper limit to the predicted Gaia end-of-life accumulated radiation dose. A trap density of 4 traps pixel^{-1} is necessary to reproduce the amplitude of the CTI effects, in particular the fractional charge loss as observed in the second RC (RC2) from first pixel response measurements (Chapter 2 Fig. 5). This test was performed with charge injections occurring every ~ 27 s. The relative image location bias was measured in similar conditions during the same campaign for a star transit occurring 1 and 27 s after the last CI. These results are summarized in Fig. 3.12 along with the absolute location bias computed in this study. For a CI delay of 1 s the location bias is clearly smaller than for a longer CI delays (e.g., CI period of ∼27 s), this is due to the fact that shorter CI periods maintain a larger portion of the traps constantly filled. As a consequence, our simulations were performed for two different active trap densities (or level of radiation damage), 1 and 4 traps pixel^{-1}. By active we mean empty before the transit of the star of interest over the CCD. These densities reproduce the amplitude of the CTI effects as observed for short and long CI periods in the experimental tests.
Figure 3.2 shows the resulting simulated CTI-induced distortion and charge loss by comparing, for different illumination levels, the simulated CTI-free observations and damaged observations (4 traps pixel^{-1}) after a normalization. Note that for a unique damaged CCD containing a single trap species with fixed parameters, the distortion varies significantly from one signal level to the other and not linearly. This is due in particular to the SBC, which mitigates the CTI effects at low signal level.

3.4 The Gaia image location estimation procedure

3.4.1 Observation model: scene

To model the flux distribution that illuminates a CCD we need a model of the instrument response to a point-like source, and a model of the actual distribution of (point) sources on the sky. The former has already been parameterized in Section 3.3.2: it is given by the line spread function L when considering one dimension, or the PSF when considering two dimensions. Because we will mainly deal with one-dimensional data in this study we will hereafter only refer to the one-dimensional LSF. For the purpose of this study a simple observation model is sufficient: $E(N_k)$ is the expected number of photoelectron counts in pixel k, λ_k is the modelled photoelectron count given by a flat background β plus a single point source with flux α at location κ:

$$E(N_k) \equiv \lambda_k = \alpha L\left(k - \kappa\right) + \beta \tag{3.2}$$

Here α, κ, and β are called the scene- or image parameters.

3.4.2 Maximum-likelihood estimation of the image parameters

In the Gaia data processing the image location κ and flux α will be estimated by fitting the modelled photoelectron counts $\{\lambda_k\}$ (Eq. 3.2) to the observed photoelectron counts $\{N_k\}$ using a Maximum-Likelihood (ML) algorithm, and this for each observation.The image background β is not determined using the ML algorithm but beforehand by a more adequate method. This method makes use of empty telemetry windows to estimate separately the different main components of the background: astrophysical background (zodiacal light, faint stars and galaxies i.e. $G > 20$), CI trails, and the CCD electronic offset. In this study we always assume that β is known. Another parameter that is considered to be known beforehand is the CCD readout noise r. Therefore, when estimating any of the image parameters, the true values of β and r will be used. The ML algorithm is comprehensively described in Lindegren (2008), therefore only the main assumptions and equations are detailed in this chapter.
According to the ML principle, the best estimate of the parameter vector θ (here $\theta_1 = \kappa$ and $\theta_2 = \alpha$) maximizes the likelihood function or equivalently the log-likelihood function:

$$\ell\left(\theta|\{N_k\}\right) = \sum_k \ln p\left(N_k|\lambda_k, r\right) \tag{3.3}$$

with p the probability density function of the sample value given the modelled count and readout noise. We hence need to adopt a probability model for the sample values. To do so we assume (i) that the noise is not correlated from a sample to another (already implicit in the sum in Eq. 3.3); (ii) that the variance of the noise is $E[(N_k - \lambda_k)^2] = \lambda_k + r^2$; and (iii) that the sample value including the readout noise can be modelled as Poissonian random variables, $N_k \sim \text{Poisson}(\lambda_k + r^2) - r^2$ (see Lindegren 2008). That the Poisson distribution is discrete, while N_k (obtained by correcting the digitized values for bias and gain) are in general non-integer, is not a problem as long as $N_k + r^2 \geq 0$. The continuous probability density function derived from the Poisson distribution is:

$$p\left(N_k|\lambda_k, r\right) = \text{const} \times \frac{(\lambda_k + r^2)^{N_k + r^2}}{\Gamma(N_k + r^2 + 1)} e^{-\lambda_k - r^2} \tag{3.4}$$

and Eq. 3.3 can then be re-written:

$$\ell\left(\theta|\{N_k\}\right) = \text{const} + \sum_k [(N_k + r^2)\ln(\lambda_k(\theta) + r^2) - \lambda_k(\theta)] \tag{3.5}$$

which is maximized by solving the following system of equations:

$$\frac{\partial \ell\left(\theta|\{N_k\}\right)}{\partial \theta} = \sum_k \frac{N_k - \lambda_k(\theta)}{\lambda_k(\theta) + r^2} \frac{\partial \lambda_k}{\partial \theta} = \mathbf{0} \tag{3.6}$$

These equations are non-linear and must be solved by iteration. Given an initial estimate $\theta^{(0)}$, the linear system to be solved in iteration m is:

$$\mathbf{A}^{(m)} \Delta\theta^{(m)} = \delta^{(m)} \tag{3.7}$$

whereupon

$$\theta^{(m+1)} = \theta^{(m)} + \Delta\theta^{(m)} \tag{3.8}$$

A is a symmetric positive definite matrix computed from the expectation of the Hessian matrix; its elements are:

$$A_{ij} = \sum_k \frac{1}{\lambda_k(\theta) + r^2} \frac{\partial \lambda_k}{\partial \theta_i} \frac{\partial \lambda_k}{\partial \theta_j} \tag{3.9}$$

and

$$\delta_i = \sum_k \frac{1}{\lambda_k(\theta) + r^2} \frac{\partial \lambda_k}{\partial \theta_i} \tag{3.10}$$

The iterations converge quickly if the initial estimate is reasonably close to the ML solution.

3.4.3 First image parameter estimates and LSF model

The ideal image model L (the true underlying flux distribution for each observation) corresponds to the reference image that is used to generate the data. During the mission, L will not be known. Therefore we have to estimate an image model $\tilde{L}$ using the observations themselves. This estimation is an iterative process (Section 3.4.4), and successive iterations are denoted with the superscript (n) for $n = 0, 1, \ldots$ (not to be confused with the iterations in Eqs. 3.7 and 3.8).
Given a set of transits for a certain reference image, background and G, denoted by $\{\{N_k\}\}_G$, how do we make the first estimate of the image parameters and generate the first LSF model, $L^{(0)}$? As mentioned in Section 3.4.2, the background β and readout noise r are assumed to be known already. The most straightforward initial flux estimate $\alpha^{(0)}$ can be made by simply taking the sum of the observed counts after subtracting the background. The initial estimate for the image location $\kappa^{(0)}$ is determined using Tukey's Biweight centroiding algorithm (Press et al. 1992; Lindegren 2006).
To generate the first estimate $\tilde{L}^{(0)}$ we use the initial location estimates $\{\kappa^{(0)}\}_G$ to relatively align the photoelectron counts of all selected profiles and create an oversampled profile. The creation of the oversampled profile is possible because each count results from the sampling of the reference image at a different sub-pixel position (Section 3.3.3). After a background subtraction the oversampled profile is fitted by the special quartic spline to obtain $\tilde{L}^{(0)}$. This profile estimation procedure is illustrated in Fig. 3.3.

3.4.4 Iterative image parameter and LSF model improvement

Once the first image model $\tilde{L}^{(0)}$ is available, an improved estimate of the image parameters of each individual transit can be made using the ML algorithm (described in Section 3.4.2). This is illustrated in Fig. 3.4. Based on these improved image parameters an improved image model can be constructed, leading to the iterative scheme

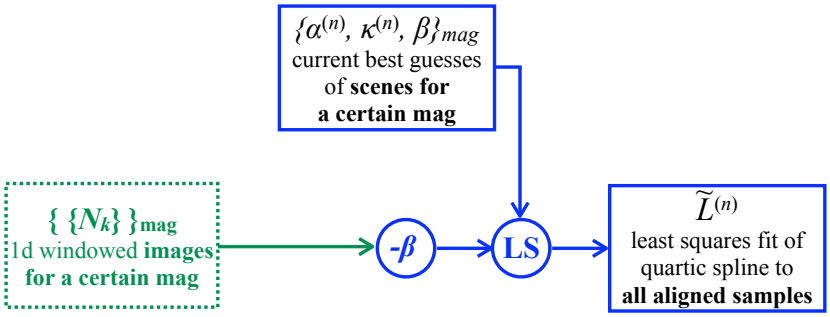


Figure 3.3 — Diagram of the construction of the estimated LSF model $\tilde{L}$. For a particular reference image, background level and G, the photoelectron counts of all selected profiles are aligned relative to each other using the current best estimates of the scene parameters to create an oversampled profile, which is then fitted by the special quartic spline using a Least Squares algorithm, after removal of the background.

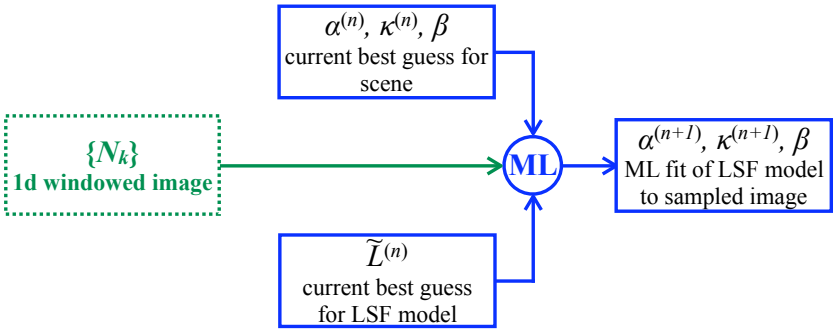


Figure 3.4 — Diagram summarizing the estimation of the scene (or image) parameters from a single observation. Modelled counts, computed using the latest LSF model, are fitted to the observed counts using our Maximum Likelihood (ML) algorithm (see Section 3.4.2). Note that the background is not determined by the ML algorithm but by a dedicated procedure not detailed in this chapter. In our study, the background β is assumed to be known.

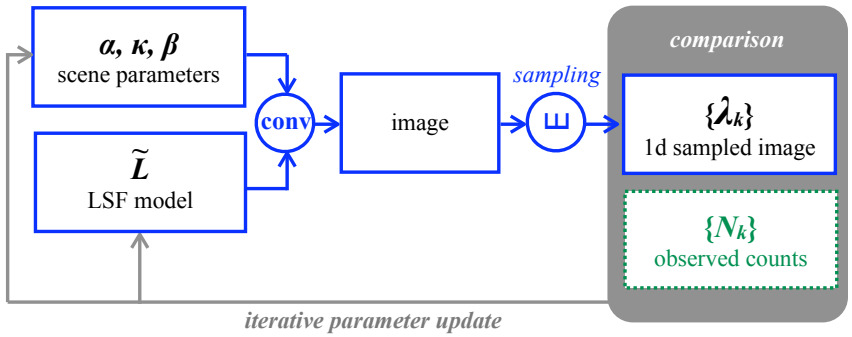


Figure 3.5 — Resulting top level diagram of the Gaia image parameter estimation iterative procedure. The LSF model, $\tilde{L}$, and the scene parameter estimates, κ, α (respectively the image location and flux), are iteratively improved by fitting the modelled counts $\{\lambda_k\}$ to the observed counts $\{N_k\}$. The modelled counts are predicted by our observation model (Eq. 3.2) based on the current best scene parameter estimates. In our study the background β is assumed to be known with a high accuracy (we use the true value).

shown in Fig. 3.5 where the image parameters and image model are improved one after the other. Note that in the whole procedure we have not used any prior knowledge: everything is estimated from the observed photoelectron counts (i.e. 'self-calibrating'). After each iteration the residuals between the modelled and the observed photoelectron counts are monitored through the computation of the χ^2:

$$\chi^2_G = \sum_{t=0}^{T-1} \sum_{k=0}^{K-1} \frac{(\lambda_{tk} - N_{tk})^2}{\sigma^2_{tk}} \text{ with } \sigma^2_{tk} = N_{tk} + r^2 \tag{3.11}$$

where T is the total number of transit profiles for a certain G, and K the number of along-scan pixels in each profile. For transit t and pixel k, λ_{tk} and N_{tk} are the predicted and observed photoelectron counts respectively. The uncertainty σ_{tk} is considered to be equivalent to the quadratic sum of the photon noise and the readout noise r.
The agreement between observed and modelled counts (Fig. 3.6) does not significantly improve after two iterations, however the agreement between the LSF model and the reference image (see Fig. 3.7), and the average image location bias as well as the location estimator standard errors (Sections 3.5.3 and 3.5.4) can still improve after a certain number of iterations that essentially depends on the stellar magnitude. As a consequence, we stop the iterative procedure after a particular number of iterations that is determined for each magnitude beforehand.

3.5 Theoretical and actual limit to the image location accuracy

To be able to evaluate the accuracy of the image location estimation procedure, we first need to determine what the theoretical limit of any image location estimator is. This is done by computing the Cramér-Rao bound (Section 3.5.1), which shows that it depends uniquely on the image shape, flux, background and noise. Subsequently, and first in the absence of CTI, we verify that any potential bias of the Gaia image location estimator does not depend on the image location (Section 3.5.2). And then the estimator bias and standard errors are calculated as a function of G, image reference width, and for different operating conditions (Sections 3.5.3 and 3.5.4). We compare the latter to the theoretical limit and evaluate the efficiency of our estimation procedure in the absence of radiation damage. Then we estimate the irreversible loss of accuracy intrinsic to radiation damage, by computing the Cramér-Rao bound for a 'damaged LSF' generated from the data set of damaged observations (Section 3.5.5). Ultimately, we apply the Gaia image location estimator to the damaged observations without any CTI mitigation. This allows us to characterize the radiation damage induced location bias (Section 3.5.6), and check the consistency of our CTI effects simulation by comparing our results to experimental test results (Section 3.5.7).

3.5.1 Definition of the astrometric Cramér-Rao bound

For a dataset with a known underlying probability density function the Cramér-Rao minimum variance bound theorem gives the minimum reachable variance of a free parameter using any estimation procedure. In the case of estimating the location κ of a one-dimensional image containing N_p detected photons, the Cramér-Rao bound σ^2_κ

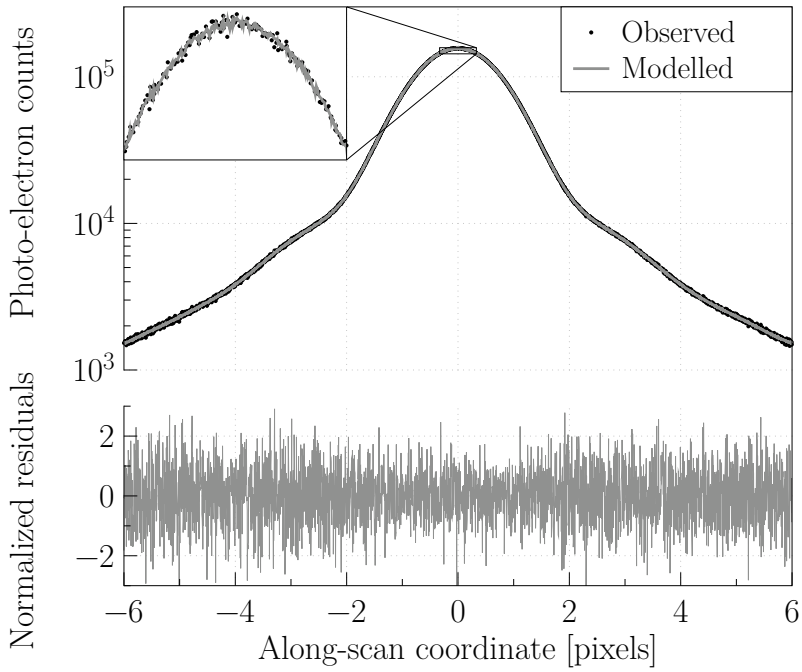


Figure 3.6 — Top: Comparison between the observed CTI-free observations $\{\{N_k\}\}_{13.3}$ (black dots), and the modelled counts $\{\{\lambda_k\}\}_{13.3}$ (grey line) computed following the presented iterative procedure. The transits were generated from the typical reference image, at $G = 13.3$, with readout noise and background. **Bottom:** Residuals normalized by the noise: $(N_{tk} - \lambda_{tk})/\sigma_{tk}$, at the last stage of the image parameter estimation iterative procedure.

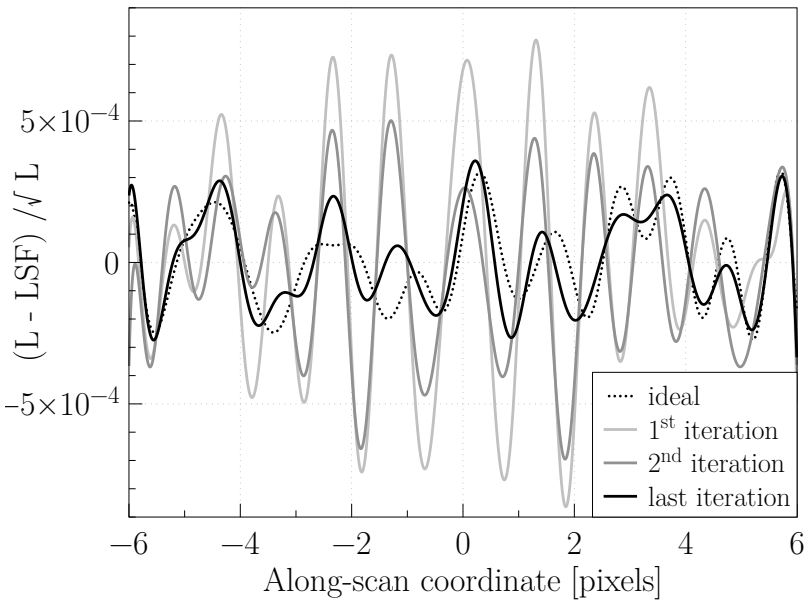


Figure 3.7 — Normalized difference between the LSF model and the true underlying flux distribution at different stages of the image parameter estimation iterative procedure ($G = 13.3$, readout noise and background, and typical reference image). The ideal case corresponds to a LSF model constructed from the observed data and the true image location.

can be expressed as follows (Lindegren 1978):

$$\sigma_\kappa^2 = \frac{1}{N_p}\left(\int \frac{|I'(x)|^2}{I(x)+(\beta+r^2)/N_p}dx\right)^{-1} \tag{3.12}$$

with $I(x)$ a normalized one-dimensional flux distribution of the image along x, β the background and r the CCD readout noise.

3.5.2 Location independent error and standard deviation

The iterative procedure on a transit (described in Section 3.4) provides us with an image location estimation $\kappa^{(n)}$ for iteration n. One can compute the image location error δ_κ by directly comparing κ to the true image location κ_{true}:

$$\delta_\kappa = \kappa - \kappa_{\text{true}} \tag{3.13}$$

Before averaging over all the image location estimates (or the corresponding errors) of a particular magnitude (for a particular reference image, window size, background level and readout noise value), one first needs to check that the error does not significantly fluctuate as a function of the relative location offset from the pixel grid. The latter is simply given by the collection of estimated image locations $\{\kappa\}_G$. Fig. 3.8 shows an example of this variation: each point corresponds to the average location error over 25 adjacent sub-pixel positions, the error bars represent the standard deviation of the points with respect to this mean. At the last stage of the iterative procedure, the set of estimated locations $\{\kappa^{(n)}\}$ shows that there is virtually no significant systematic variation across the relative location offsets. A certain number of iterations (7 for the brightest and 2 for the faintest) is needed to remove the variation introduced during the procedure initialization by the Tukey's Biweight centroiding algorithm.
Having established that there is no significant error as function of the relative location offset from the pixel grid it is allowed to average over all the transits within a particular magnitude (for a particular reference image, window size, background level and readout noise value) to find the bias:

$$\langle\delta_\kappa\rangle = \frac{1}{T}\sum_{t=0}^{T-1}\delta_{\kappa_t} \tag{3.14}$$

with t going through all transits of the transit selection. We will indicate the average over all transits of a particular magnitude as $\langle\delta_\kappa\rangle_G$. For this subset of transits we can now also compute the corresponding standard deviation:

$$\sigma_\kappa = \sqrt{\frac{1}{T-1}\sum_{t=0}^{T-1}(\delta_{\kappa_t} - \langle\delta_\kappa\rangle_G)^2} \tag{3.15}$$

The statistical uncertainty of this standard deviation is:

$$\upsilon_{\sigma_\kappa} = \frac{\sigma_\kappa}{\sqrt{2T}} \tag{3.16}$$

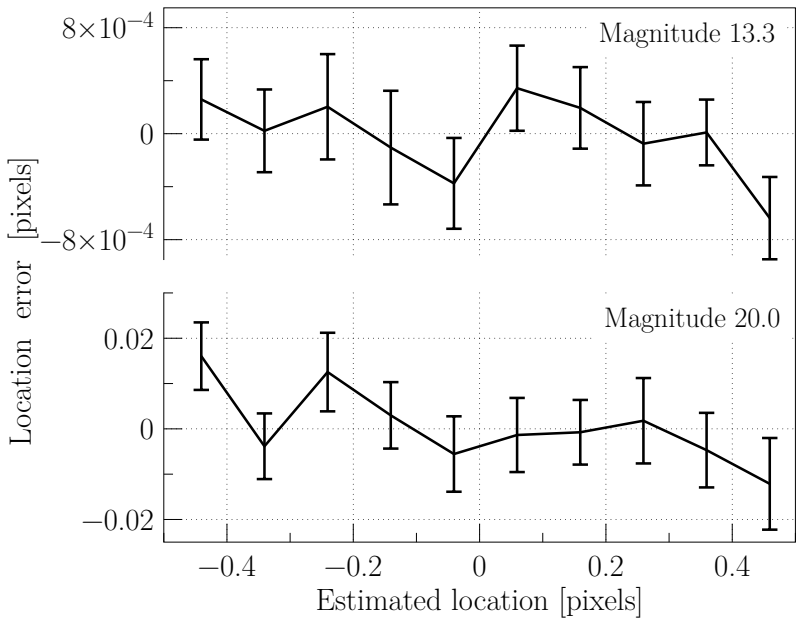


Figure 3.8 — Image location error $\{\delta_\kappa^{(n)}\}_G$, as a function of the location offset from the pixel grid $\{\kappa^{(n)}\}_G$, for the brightest (top) and the faintest magnitude (bottom) at the last stage of the iterative procedure (respectively 7 and 2 iterations).

and the statistical uncertainty of the bias is:

$$\upsilon_{\langle\delta_\kappa\rangle} = \frac{\sigma_\kappa}{\sqrt{T}} \tag{3.17}$$

Summarizing, we can for all transits of a particular magnitude quantify the location bias as: $\langle\delta_\kappa\rangle_G \pm \upsilon_{\langle\delta_\kappa\rangle}$, and the location standard deviation, hereafter called location precision, as: $\sigma_{\kappa,G} \pm \upsilon_{\sigma_\kappa}$. In the absence of any significant bias, the latter can be referred to as the location accuracy.

3.5.3 CTI-free location bias results per magnitude

Figure 3.9 shows the image location bias as a function of G for the three different reference image widths, a sky background level set to the average sky brightness and the Gaia CCD operating conditions regarding the readout noise value (4.35 e^-) and the size of the telemetry windows in the along-scan direction (12 pixels for G <16 then 6 AL pixels G >16). This set of conditions, hereafter referred as to Gaia operating conditions, constitutes the most realistic case of our study and also the most unfavourable case for the image parameter estimation procedure. Yet, one can observe from Fig. 3.9 that the location bias, $\langle\delta_\kappa\rangle_G$, for none of the magnitudes exceeds the level of 5 milli-pixels (~ 0.3 mas). Moreover: within the uncertainty of our measurement, and for the three different image widths, $\langle\delta_\kappa\rangle_G$ does not significantly deviate from zero. Hence we can establish that the Gaia image location estimator is a bias-free estimator in the absence of radiation damage. Increasing the window size in the along-scan direction or reducing the readout noise has no significant effect on $\langle\delta_\kappa\rangle_G$. Only setting the back-

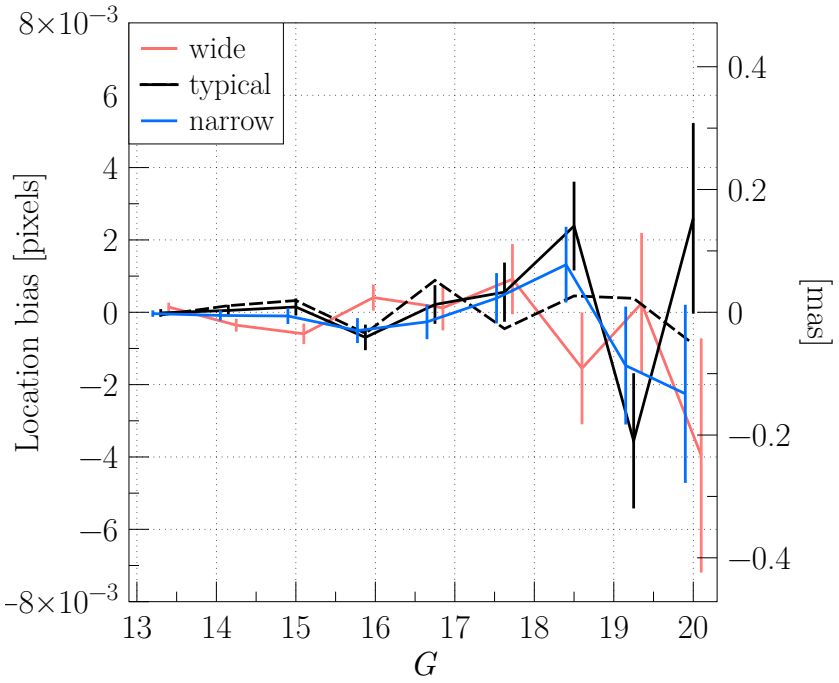


Figure 3.9 — Comparison of the location bias $\langle \delta_\kappa \rangle \pm \upsilon_{\langle \delta_\kappa \rangle}$ for different reference image widths as a function of G measured in the Gaia operating conditions (cf. Section 3.5.3). The dashed line corresponds to a measurement realized in the same conditions for the typical image width but without background. The error bars correspond to $\upsilon_{\langle \delta_\kappa \rangle}$, the statistical uncertainty on the location bias (see Eq. 3.17). Note that for readability a slight offset has been introduced on the G axis for the narrow and wide reference image results. There are two ordinate axes, left is the location bias in units of pixels and the right in units of mas. The same holds for the following figures.

ground level to zero seems to slightly decrease the bias for the faintest magnitudes. This effect is illustrated by the dashed line in Fig. 3.9.

3.5.4 CTI-free location accuracy per magnitude

To evaluate the efficiency of our estimator, we compare the measured standard errors, σ_κ, to the astrometric Cramér-Rao bound, the theoretical limit to the image location accuracy of any bias-free estimator (see Section 3.5.1). The comparison results are summarized in Table 3.2 for different values of G, image widths, window size, background levels, and values of CCD readout noise. In Fig. 3.10, we compare the accuracy of the Gaia image location estimator in the Gaia operating conditions, the Cramér-Rao bound computed for the same reference image width and level of readout noise, and the requirements as presented in Table 3.1. As one can see, the Gaia image location estimator performs remarkably well. The estimator standard errors are always below the requirements and this for any reference image width. Also the standard errors, within the measurement statistical uncertainty υ_{δ_κ}, strictly follow the Cramér-Rao bounds at every signal level. Note how stringent the Gaia requirements are: for the wide reference image, the actual and theoretical limits to the image location accuracy are very close to the required accuracy.

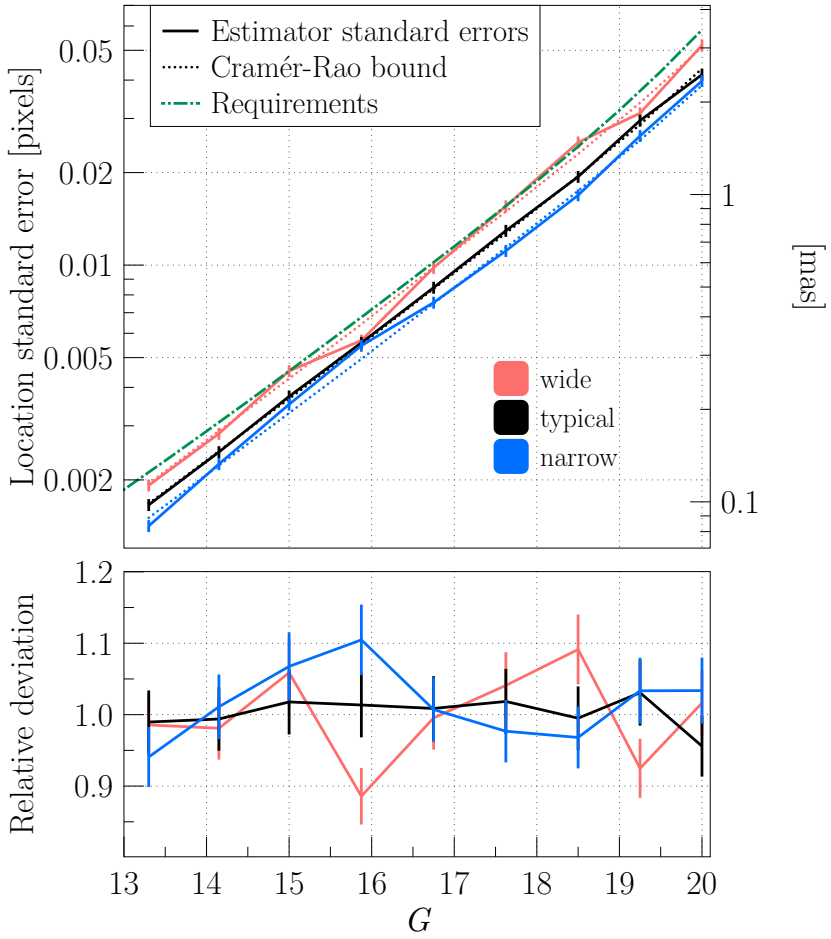


Figure 3.10 — Comparison between the actual and theoretical limit to the image location accuracy, as a function of G. **Top:** The continuous lines correspond to the Gaia image location estimator accuracy standard errors, $\sigma_{\kappa,G} \pm \upsilon_{\sigma_\kappa}$, measured in the Gaia operating conditions for the three different reference images: narrow, typical and wide (same grey scale coding as in Fig. 3.1). The associated error bars corresponds to the statistical uncertainty on the location accuracy (see Eq. 3.16). The dotted line represents the Cramér-Rao bounds computed for the same reference image widths, window size, background level, and readout noise value. Finally, the dashed and dotted line shows the Gaia required image location accuracy. **Bottom:** The ratio between $\sigma_{\kappa,G}$ and the Cramér-Rao bounds and the associated error bars are depicted as a function G. For the three reference image widths, the relative deviation does not exceed 10%. The Gaia image location estimator can thus be considered efficient in the absence of radiation damage.

As mentioned in Section 3.1, the targeted performance predictions of Gaia contain a margin of 20% to take into account unmodelled on-ground calibration errors including for instance residual bias. In this context we consider an estimator efficient if its standard errors are within 10% of the Cramér-Rao bound and thus not consuming more than half the margin. The Gaia estimator rigorously fulfills this criteria. This is illustrated in Fig. 3.10 (bottom) for the three reference image profiles: the ratio between the estimator standard errors and the Cramér-Rao bounds remain below 1.1 (i.e. 10% relative deviation). As expected, in both the theoretical and actual cases, an increase in the image width is directly translated into a loss in location accuracy. This loss varies

linearly with the image FWHM. Table 3.2 shows that increasing the readout noise, the background level, and/or decreasing the window size also increases the Cramér-Rao bound and the Gaia estimator standard errors.
In the absence of radiation damage, we established that in realistic operating conditions and from bright to faint magnitudes, the Gaia image location estimator is bias-free, efficient, and performs within the requirements, with a high accuracy close to the theoretical limit. It is now important to characterize in detail the impact of radiation damage on the image location uncertainty.

3.5.5 Radiation damage intrinsic uncertainty increase

Computing the Cramér-Rao limit (Eq. 3.12) for a flux distribution including the CTI distortion and taking into account the charge loss allows to quantify this intrinsic uncertainty increase induced by the radiation damage. As one can observe from Fig. 3.2 the CTI induced distortion sharpens the image profiles and renders them more asymmetric but the charge loss significantly decreases the signal-to-noise ratio. The latter effect prevails and, at a given G, causes an increase in the image location uncertainty. To generate L_D, the damaged flux distribution, we proceed in a similar fashion to the construction of $\tilde{L}$ (cf. Sections 3.4.3 and 3.4.4). First we place each data point from the damaged observations at the right sub-pixel position to create an oversampled damaged profile. Then the over-sampled profile is fitted by the special quartic spline so that we can use an analytical representation. The resulting minimum variances on the estimate of an image location acquired by a damaged CCD, and thus accounting for the CTI effects, are summarized in Table 3.3 for different image widths, background levels and levels of radiation damage.
In the Gaia operating conditions, the relative intrinsic uncertainty increase (or accuracy loss) can be as large as 23% (see Fig. 3.11) for the highest trap density and 6% for the lowest. Here we recall that this drop in active trap density results from the use of a more frequent CI: from a CI period of 27 s to 1 s. In both cases, this increase is more pronounced for narrower stellar profiles and peaks at a signal level of $G = 15.875$. Then, due to the mitigating effects of the SBC at lower signal levels, one clearly observes a flattening of the uncertainty increase. This illustrates the critical importance of the two hardware mitigation tools (see Section 3.3.4), which are the only mitigation countermeasures capable of reducing the CTI induced intrinsic loss of accuracy, by physically preventing the electron trapping and thus the image distortion and charge loss.
The Cramér-Rao bound computed for the damaged flux distribution now constitutes the maximum achievable accuracy by any unbiased image location estimator in the presence of radiation damage. Although the loss of accuracy can be quite large, Fig. 3.11 shows that the Gaia requirements would still be fulfilled, if an image location estimator that is bias-free and efficient enough can be elaborated (excluding the wide reference image and highest trap density case). In the next section, in order to assess the efficiency of the Gaia image location estimator without CTI effects mitigation (Section 3.4) in the presence of radiation damage, we directly apply it to the data set of damaged observations.

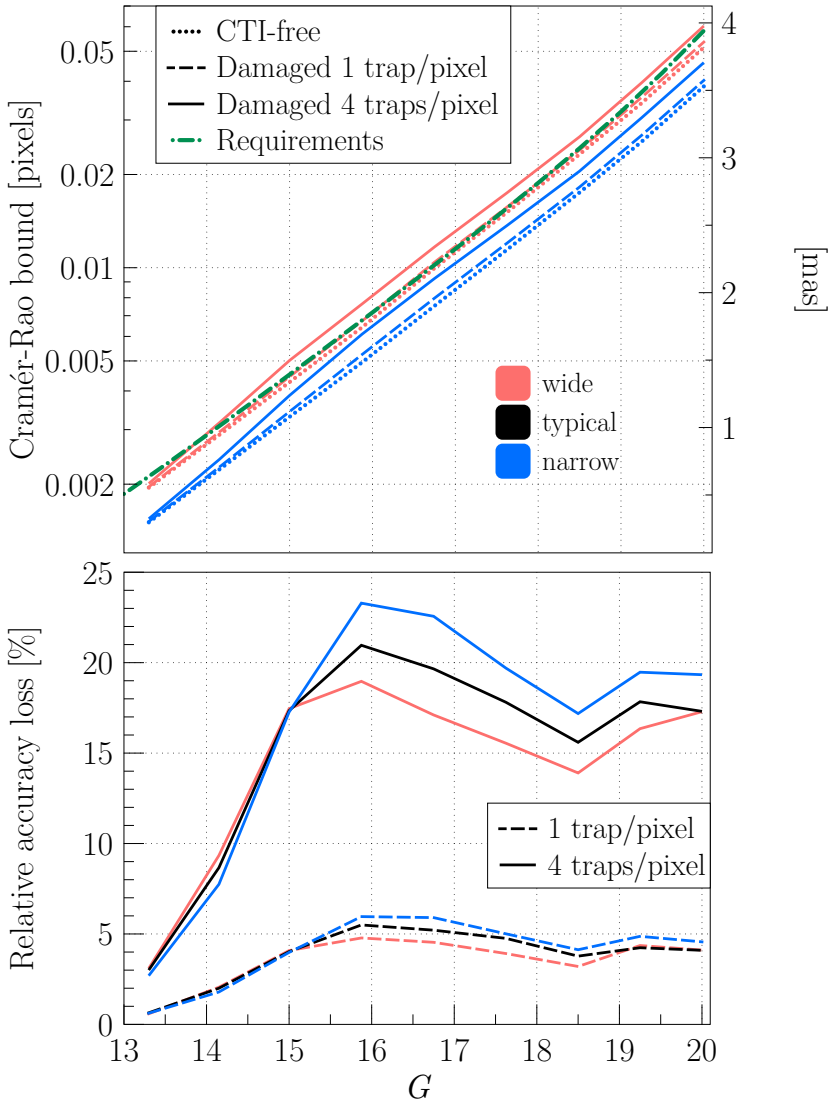


Figure 3.11 — Top: Comparison between the Cramér-Rao bounds computed from the original flux distribution (dotted lines) and the constructed damaged flux distribution for the two trap densities: 1 trap pixel^{-1} (dashed line), 4 traps pixel^{-1} (continuous line). Two reference images are here considered: the narrow and the wide. The window sizes, background level and readout noise value correspond to the Gaia operating conditions. The computed required image location accuracy is also shown (dashed-dotted line). **Bottom:** The relative intrinsic loss of accuracy induced by radiation damage as a function of G for the three reference images: narrow, typical (black), wide, and the two trap densities: 1 trap pixel^{-1} (dashed line), 4 traps pixel^{-1} (continuous line). The relative loss of accuracy corresponds to the relative difference between the Cramér-Rao bound computed from the original flux distribution and the constructed damaged flux distribution. Note the important difference in loss amplitude, for the two different trap densities: a reduction in the active trap density (e.g., by the means of CI) is directly translated into a gain in location accuracy of a similar factor. Similarly, the flattening of the intrinsic loss for $G > 15$ is due to the effect of the SBC.

3.5.6 Radiation induced image location bias

In this section we are interested in exploring the consequences of not accounting for the CTI effects during the image location estimation. We thus apply the Gaia image location estimator as presented in Section 3.4 to the data set of damaged transits. In this case the image distortion shall not be accounted for in the LSF model construction. This is achieved by using $\tilde{L}_U$, the LSF model generated from the CTI-free transits. After applying the procedure, one eventually obtains an estimated location κ_D for each transit, which after subtraction of the true image location κ_{true}, gives us the error δ_κ. After averaging for a particular magnitude and CCD operating conditions, we obtain the image location bias induced by the CTI effects as a function of signal level, $\langle \delta_\kappa \rangle_{G,D}$. The location bias results from the mismatch between the observed profile shape and the modelled LSF used to estimate the location. Only one iteration of the scheme from Fig. 3.5 is performed since the LSF model cannot be improved using the damaged counts without taking into account the CTI effects.
The results for different image widths, window sizes, and background levels are summarized in Table 3.4 and depicted in Fig. 3.12 (left) for a trap density of 4 traps pixel^{-1} and in Fig. 3.12 (right) for 1 trap pixel^{-1}. The bias strongly varies as a function of G. In the Gaia operating conditions including background, the location bias reaches a maximum for $G = 15$. The mitigation effects of the SBC is clearly noticed for $15 < G < 18$ as the bias is either reduced or levels off. For $G > 18$, the background plays an important role in limiting the image distortion and reducing the bias as can be seen by comparing the dashed and solid lines. From these results we can conclude that in the presence of radiation damage, and without any attempt at any stage to correct or mitigate the CTI effects, the estimator is strongly biased. Indeed the image location can be shifted from a tenth of a pixel up to half a pixel for the fainter stars in the no-background case. When the estimator is applied to the damaged observations simulated with a background level set to the average sky brightness, the location bias is not as dramatic at low signal level. Nevertheless the image location bias for any image width and any signal level is constantly higher than $\sim$0.1 pixels in the 4 traps pixel^{-1} case, which is not acceptable. Changing the telemetry window size has no significant effect on the location bias. As can be seen from Fig. 3.12 (right), for a shorter CI delay (or CI period), and thus less active traps (here 1 trap pixel^{-1}), the location bias is significantly lowered with a minimum level of $\sim$0.02 pixels. It is interesting to note that decrease in bias is scaled by the same factor (~ 4) as the decrease in trap density. Regarding the required performance, for the faintest magnitude this level of bias might be acceptable in a limited amount of cases (e.g., the bluest stars). However, in most cases, and especially for the bright stars this level of bias inevitably requires a software-based CTI mitigation scheme.

3.5.7 Comparison with experimental data

In order to check how representative the results obtained from synthetic data are in terms of the overall amplitude of the CTI effects and also fluctuation as a function of signal level, Fig. 3.12 shows results obtained experimentally from RC2 (Georges 2008; Brown 2009b). In the experimental case the location bias does not correspond to an absolute image location bias since the true image location is by definition unknown.

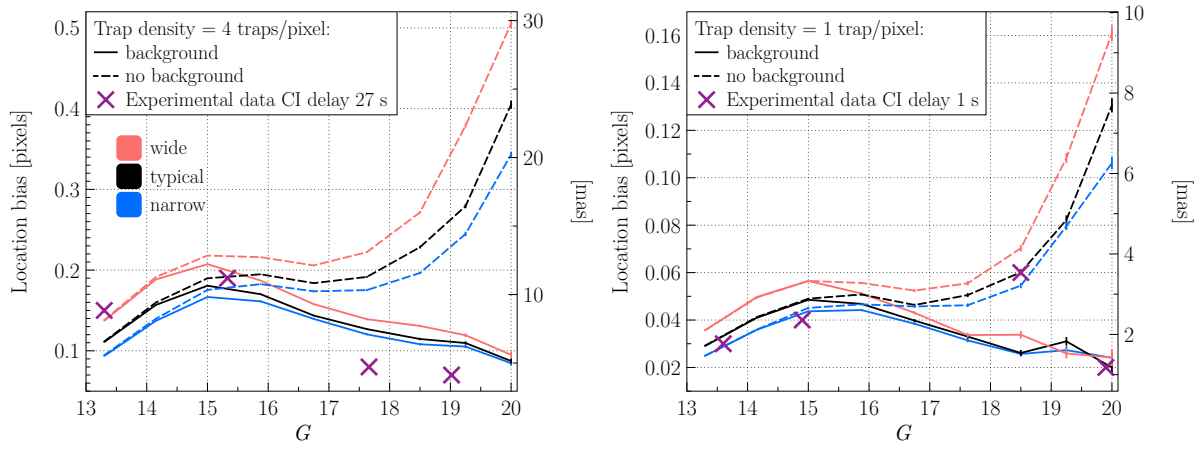


Figure 3.12 — Image location bias, $\langle \delta_\kappa \rangle_{G,D}$, resulting from the radiation damage effects on the stellar image and the use of an image location estimator that does not take into account these effects. $\langle \delta_\kappa \rangle_{G,D}$ corresponds to the mean location error for all transits at a particular G. It is measured in the Gaia operating conditions (window size and readout noise value) for the three reference image widths and with no background (dashed line) and a background level set to the averaged sky brightness (continuous lines). The left figure shows the results obtained for a trap density of 4 traps pixel^{-1}, representative of a CI delay of ∼27 s as demonstrated by the comparison to experimental results (crosses). The right figure shows the results obtained for a trap density of 1 trap pixel^{-1}, representative of a CI delay of 1 s. The very small error bars correspond to the computed statistical uncertainty on the measured bias, υ_δ (see Eq. 3.17). The crosses represent the average relative location bias computed from experimental tests carried out during RC2, for a CI delay of ∼27 s and 1 s and the lowest background level. The large experimental uncertainties are not shown, but as can be noticed the overall trend and amplitude of the measured bias is well reproduced by our model.

The presented bias is thus the relative location bias. It is computed by comparing the stellar transits over the irradiated part of the CCD and the same stellar transits over the non-irradiated part of the same CCD. Taking into account the differences between real and synthetic data, as well as experimental uncertainties, the overall agreement between the results obtained from the RC2 and our simulations is remarkable. The combined mitigating effects of the SBC and the background are also noticeable in the test data at low signal levels. Hence not only the amplitude of the location bias for different CI delays (or densities of active traps) is reproduced by our model but also the overall bias evolution over a wide range of signal levels: 7 magnitudes. The simulations suggests that the illumination setup (and resulting PSF width) as well as slight differences in background light between experiments can have a significant impact on the measured CTI effects. This may explain observed discrepancies between the results from different RCs and within a RC.

3.5.8 Damaged location estimation standard errors

Finally we show the resulting standard errors, σ_{κ_D}, as a function of G in Fig. 3.13: the standard errors are larger than the theoretical minimum variance, especially for intermediate magnitudes. For the most severe radiation level, the standard errors are larger than the requirements, and for the lowest radiation level the requirements are barely met; the mismatch between modelled and observed line spread function implies

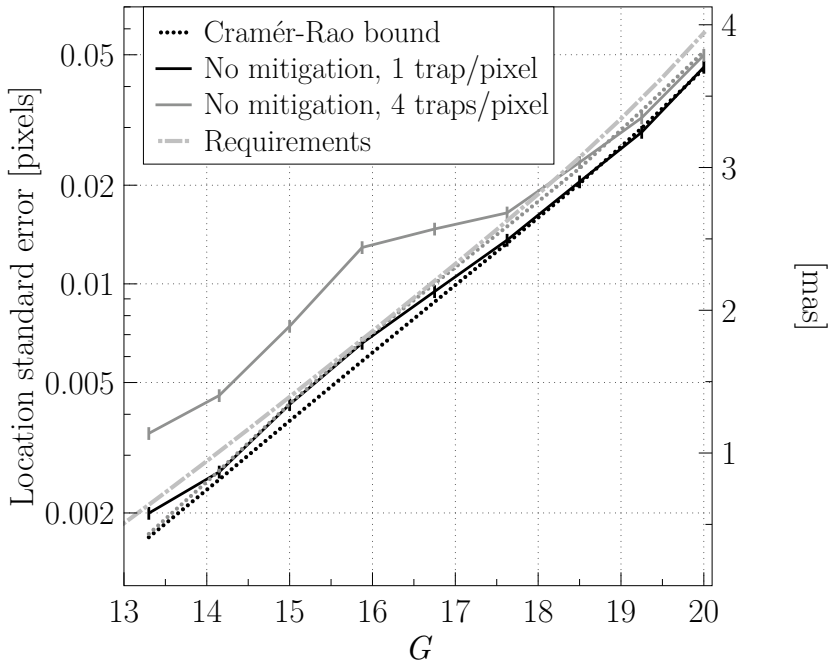


Figure 3.13 — Comparison between the Cramér-Rao bound in the presence of radiation damage (dotted lines) and the standard errors (continuous lines) obtained by applying the Gaia image location estimation procedure to the set of damaged observations without any CTI mitigation (i.e. using the LSF model generated from the CTI-free observations). The Cramér-Rao bounds were computed and the standard errors measured for the typical reference image considering the Gaia operating conditions and two different levels of radiation damage: 1 trap pixel^{-1} (black) and 4 traps pixel^{-1} (grey). In the presence of radiation damage, and without CTI effects mitigation, the Gaia image location estimator cannot be considered efficient anymore as its standard errors deviate significantly from the Cramér-Rao bound (in addition to the estimator being biased c.f. Fig. 3.12).

a broader spread in the locations estimated by the ML algorithm. This effect is less pronounced for the lowest level of radiation as the distortion, and thus the mismatch, is less important. The overall variance remains quite low as compared to the bias. Table 3.5 summarizes these results for the three different image widths.

3.6 CTI effects mitigation

Correcting for CTI is a complicated task and not only because the induced charge loss and distortion are considerable (Fig. 3.2). The trapping probabilities (e.g., Chapter 2) depend on the electron density and thus the CTI effects vary with G. This variation is not linear, in particular due to the presence of the SBC that mitigates the CTI effects only at low signal levels. This can be clearly observed from Fig. 3.12 from both simulations and experimental data. An important consequence is that the stellar core and wings (in the CCD serial direction) will not experience the same distortion. These different contributions to the global stellar image distortion will nevertheless be collapsed into a one-dimensional signal. In addition, the location bias and charge loss will not be repeatable for a particular star or signal level as the CTI effects depend on the state (empty or filled) of the traps prior to the stellar transit. During the mission each star

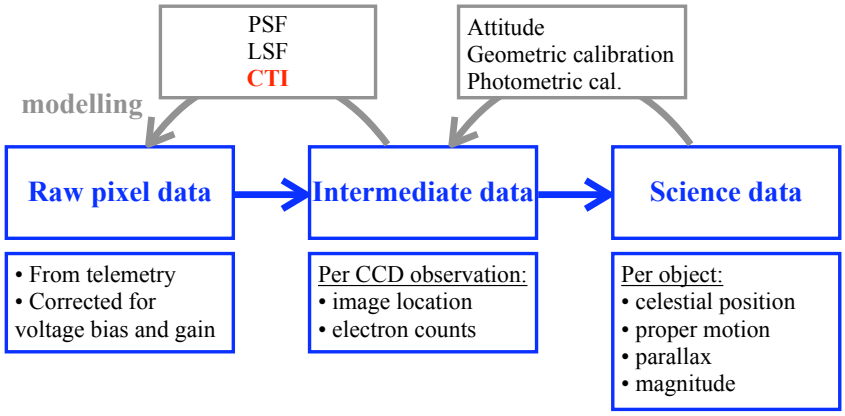


Figure 3.14 — From raw to science data: summary of the Gaia data processing chain in three stages (middle boxes). The top boxes contain the required models/assumptions to go from one stage to the other. The top arrows symbolize the feedback occurring at each stage which enables the iterative improvements of both models and data. The bottom boxes give some details about the data delivered/used at each stage. In principle, each set of data can be empirically corrected for CTI, however we choose not to perform any correction but to model the CTI distortion as part of the image parameter estimation procedure (see Section 3.6.2 and below).

will transit on average ∼72 times over the focal plane of Gaia. For each of these stellar transits the scanning direction of the satellite will differ, and thus also the CCD illumination history that determines the trap state. It is also likely that the trap density will have increased between two consecutive transits. The CI will play an important role here, not only by decreasing the active trap density but also by simplifying the illumination history by reseting it every 1 s. Finally, it is important to note that, as already mentioned, Gaia's launch and first year of operations coincide with the predicted peak of the Sun's activity for the current solar cycle, and that none of the Gaia measurements will be free of radiation damage. This will strongly limit our knowledge of the exact instrument LSF/PSF in space.

3.6.1 Potential alternative approaches

Figure 3.14 summarizes the Gaia data processing chain in three different stages at which a different set of data is available: (i) the raw data, (ii) the intermediate data, (iii) the science data. Each set of data is further explained in the figure. Different ways of handling the CTI effects in this chain are possible, and the literature provides us with a handful of correction procedures for photometric, spectroscopic, and (very rarely) astrometric measurements carried out in the optical or at X-Ray wavelengths.

CTI correction at the level of the raw pixel data:
One can correct the raw pixel data to obtain artificial CTI-free data and perform the rest of the data processing using the corrected raw data. This constitutes one of the most common approaches, and has been successfully used to correct the CTI effects on HST data for instance. Its main advantage is that is minimizes the impact of CTI on the remaining data processing chain, the correction being performed very close to the source of the problem. Either the photoelectron count correction is directly per-

formed by means of a parametric empirical or semi-empirical formula (e.g., Goudfrooij & Kimble 2002; Dolphin 2009) that determines the CTI induced charge loss as a function of signal level, background, radiation dose, and source position on the CCD. Or it is performed by 'comparing' the damaged observation to a simulated observation, for which the damage is simulated by an empirical or physically-motivated analytical forward model of the charge transfer and trapping (e.g., Bristow 2003; Massey et al. 2010; Anderson & Bedin 2010). Bristow (2003) provides a detailed comparison between direct empirical and model-based corrections: while the direct correction can only correct photometric and spectroscopic point source measurements, a model-based correction allows for astrometric correction of arbitrary complex sources (extended, binaries etc.). The latter is more complex, i.e. computationally intensive, but versatile and potentially more accurate. In principle, the model-based correction requires the generation of a synthetic undamaged observation to be subsequently distorted by the CTI model. However, as comprehensively described in Massey et al. (2010), and first proposed by Bristow et al. (2005), one can avoid this step and iteratively remove the CTI induced image distortion by subtracting actual and simulated observations, assuming, in a first step, that the actual damaged observation is the CTI-free input signal. This relies on the assumption that the CTI effects correspond to a slight perturbation around the true image. Although promising, the model-based correction of the raw data at the pixel level has only been tested against the empirical direct correction, and mostly for photometric and spectroscopic data. Massey et al. (2010) go one step further and assess the astrometric correction induced shift as a function of signal level and distance from serial register. Although the correction performs as expected, the accuracy of such a method cannot be guaranteed yet due to the lack of reference or CTI-free data that prevents the measurement of the method absolute bias and standard errors. On top of this uncertainty regarding the final accuracy of this method, two other considerations preclude the direct use of this approach in the Gaia data processing before more investigations. First, the noise properties of a corrected pixel value are no longer simple and may introduce hard-to-track effects in the image location estimation procedure, and subsequently in the astrometric global iterative solution (AGIS) that combines all observations to infer absolute astrometry for each observed object. In particular, the assumptions on which the maximum likelihood estimation of the image parameters is based, namely that the individual samples are statistically independent and described by the Poissonian model (Section 3.4.2), no longer hold for the corrected samples. Secondly, the lack of full frame data and the binning of most telemetry windows implies that we lack the information required to perform a full pixel-based correction.

CTI correction at the level of the intermediate data:
At this level, the correction is performed thanks to a parametric ad-hoc model (e.g., Rhodes et al. 2007; Schrabback et al. 2010). It offers the advantage of being simple and fast to apply, and once formulated the model should be relatively simple to calibrate. However, the elaboration of such model is not trivial. It first requires a careful study of the CTI effects on the parameters extracted from the raw measurements as a function of a finite number of pre-selected variables. Subsequent to this study, the dependency of the CTI induced bias on the pre-selected variables must be mathemati-

cally described for each estimated parameter of interest. It is not guaranteed that such a mathematical formulation is possible and the resulting models have by definition no predictive power. In the case of Gaia, the CTI-induced image location bias and charge loss could be parametrized as function of the signal level, background, radiation dose (or observation time), source position on the CCD, and illumination history (or time since last CI). Fig. 3.12 shows an example of the image location bias dependence on the signal level and background. Comparison between Fig. 3.12 left and right, also provides additional information about the dependence on the time since last CI. Such an approach was studied by EADS Astrium, but does not constitute the current baseline approach of the Gaia Data Processing and Analysis Consortium (DPAC) as it cannot handle complex scenes but only single stars.

CTI correction at the level of the science data:
This last potential approach is the most impractical. It also requires a parametric ad-hoc model, most likely impossible to formulate as the CTI effects are too entangled at the level of the science data. Moreover, the calibration of such approach would require the use of reference data, which in the case of Gaia will be mostly not available.

3.6.2 A complete forward modelling approach

Due to the complexity of the CTI effects and the extreme accuracy required in the image location estimation, as well as for the reasons mentioned above, the DPAC adopted a forward modelling approach. Thus in contrast to the solution applied to HST data, no direct correction of the raw data shall be performed, essentially to preserve the simple noise properties and avoid arbitrary assumptions. Instead, the true image parameters are estimated in an iterative scheme, in which each observation is ultimately compared to a modelled charge profile for which the distortion has been simulated through an analytical CTI model, a so-called charge distortion model (CDM). This approach is illustrated by the schematic depicted in Fig. 3.15, where the modelled counts are now described as follows:

$$\lambda_k = D\left[\alpha L\left(k-\kappa\right)+\beta \mid \mathbf{c}, \mathbf{h}\right] \tag{3.18}$$

with D the CTI distortion applied to the sampled image using CDM, $\mathbf{c}$ a set of CDM parameters (e.g., trap species characteristics, electron density distribution parameter), and $\mathbf{h}$ a set of parameters that describes the illumination history (the most obvious being the time since the last charge injection).
As illustrated in the Fig. 3.15, the scene, the CDM, and the instrument (LSF/PSF) parameters are iteratively adjusted until the modelled counts $D\left[\{\lambda_k\}\right]$ agree with the observation $\{N_k\}$. Fig. 3.16 gives the details of the CDM parameter update. It is important to note that the model LSF cannot be directly generated from the observations anymore as they are now affected by CTI. During the mission, the LSF model will thus be extracted from a LSF library composed partly by modelled LSFs and by a subset of observations: mostly the single bright stars that are the least affected by radiation damage (i.e. early mission data and/or observations close to a charge injection). If the CDM and the instrument model are properly calibrated, the estimated scene parameters subsequently used to determine the stellar parallaxes should be unbiased and free of CTI.

Since no direct correction is performed the noise properties of the observation should remain dominated by the photon and readout noise and thus a complex contamination of the rest of the data processing chain and its products is avoided. This approach can handle arbitrarily complex scenes and offers the advantage being in accordance with the general Gaia data processing principle of self-calibration. A similar approach was successfully used to handle CTI effects on photometric and spectroscopic X-ray measurements performed by Chandra (Townsley et al. 2000, 2002; Grant et al. 2004).

In the following (Section 3.6.3), we demonstrate the ability of the Gaia CTI mitigation approach to reach the best achievable image location estimation accuracy for damaged observations (Section 3.5.5) in the case of an ideal CDM, and ideally calibrated LSF and

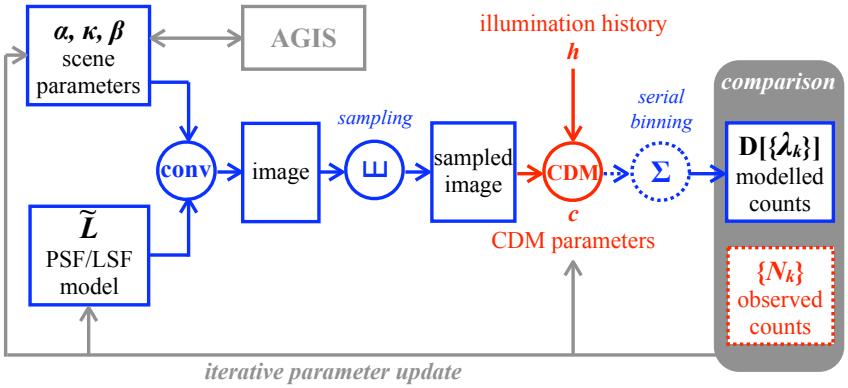


Figure 3.15 — Forward modelling approach to CTI mitigation (Lindegren 2008): a CTI-free sampled image is generated following the method explained in Section 3.4 and is subsequently distorted by a fast analytical CTI effects model, so-called CDM (Charge Distortion Model). The distorted counts are then compared to the observed counts. In an iterative procedure, the scene, the LSF model, and the CDM parameters are successively improved. Note that if CDM takes as input a two-dimensional signal, the CDM output needs to be binned in the CCD serial direction before comparison to the observed counts. AGIS, the Gaia Astrometric Global Iterative Solution, uses the scene parameters to estimate the astrometric parameters. AGIS also provides an updated estimate of the scene parameters that is corrected for 'nuisance' parameters such as the satellite attitude.

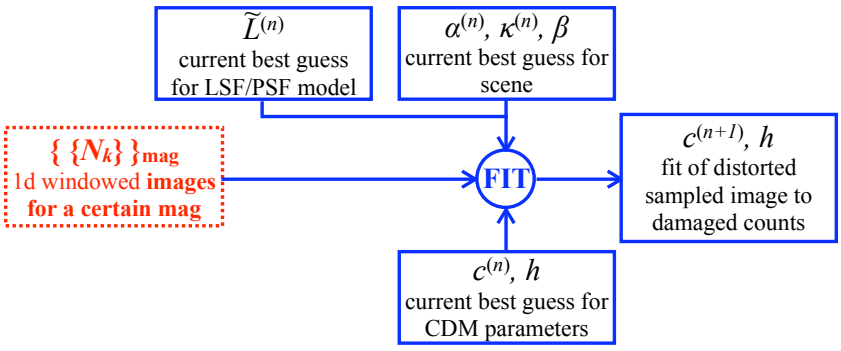


Figure 3.16 — CDM parameter estimation procedure: a set of CDM input signals is generated by sampling the current best LSF or PSF model using the best estimate of the scene parameters at a particular G. CDM simulates a set of damaged modelled observations that is subsequently compared to the set of observed counts corresponding to a particular G. A fitting algorithm provides us with a set of CDM parameters that optimizes the agreement between the CDM predictions and the observed counts.

CDM parameters. Then we assess the actual performance of this approach regarding the recovery the image location estimate bias (Section 3.6.6) and image flux estimate bias (Section 3.6.7), using the current best CDM candidate (Short et al. 2010).

3.6.3 Testing the forward modelling approach

In a first step towards a more complete validation of our approach, we would like to ensure that this approach, if perfectly calibrated, enables an unbiased estimation of the image location with high enough precision. To do so we estimate the (unknown) scene parameters for the set of damaged observations, in the case of an ideal CDM and ideally calibrated LSF and CDM parameters. This ideal case is simulated by using $\tilde{L}_D$, the damaged LSF (cf. Section 3.5.5). This is allowed because in this scheme the true LSF and CDM parameters correspond to a model that is capable of fully explaining the image distortion and the charge loss in the damaged observations.
Figures 3.17 and 3.18 show the location bias and the estimator standard errors obtained in these conditions, for the two different levels of radiation damage, and for the typical reference image. The results obtained for the two other reference images can be found in Tables 3.4 and 3.5. As one can see, in the case of ideal CTI mitigation, the location bias in the presence of radiation damage is now comparable to the one obtained for the CTI-free observations (see Fig. 3.9); the bias does not exceed 5 milli-pixels and does not significantly deviate from zero within the error bars ($\upsilon_{\langle\delta_\kappa\rangle}$, the statistical uncertainty), and this even for the most severe level of damage. Regarding the estimator standard errors, they comply with the Gaia requirements, even in the case of the most severe level of damage for most of the magnitudes. The bottom part of Fig. 3.18 shows that the location estimator including CTI mitigation performs efficiently for the lowest level of damage (i.e. less than 10% relative deviation). However it is interesting to note that even in this favourable case (ideally calibrated LSF model and CDM parameters), the relative deviation of the estimator precision from the best achievable one can reach 20% for the intermediate magnitudes and the strongest level of damage.
From these results we can conclude that a forward modelling approach to CTI mitigation, as presented in the previous section, allows the recovery of the CTI-induced location bias and enables the bias-free estimation of the image location at the required precision, close to the theoretical limit. This level of performance is achieved in the favorable conditions of a very good LSF model and the CDM parameter calibration, but for the strongest expected image distortion in the Gaia operating conditions, i.e. stars located the furthest away from the last CI and a density of traps equivalent to the predicted upper limit to the Gaia end-of-life accumulated radiation dose.

3.6.4 Current best CDM candidate

The elaboration and calibration of a CDM that allows to reach the level of performance presented in Section 3.6.3 is challenging. The presented mitigation scheme requires a CDM that must be both accurate and fast, as the iterative procedure is performed for each observation, and the CDM distortion applied at each iteration. The DPAC strategy regarding the elaboration of such CDM is detailed by van Leeuwen & Lindegren (2007) and van Leeuwen (2007a) and a short summary is given by Prod'homme (2011). In this

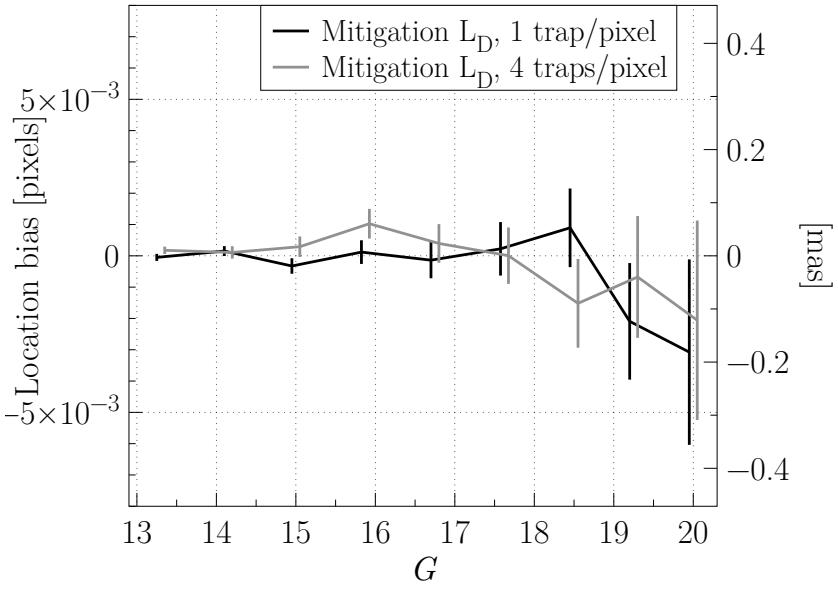


Figure 3.17 — Location bias $\langle \delta_{\kappa,D} \rangle \pm \upsilon_{\langle \delta_\kappa \rangle}$ after a CTI mitigation using the LSF model $\tilde{L}_D$, representative of an ideally calibrated forward modelling approach. These results are obtained in the Gaia operating conditions, considering the typical reference image only and for two different levels of radiation damage: 1 trap pixel^{-1} (black) and 4 traps pixel^{-1} (grey). As one can see, in these conditions, the forward modelling approach allows for a full recovery of the CTI-induced location bias. Note that for readability a slight offset on the G axis as been introduced between the results for the two trap densities.

study, to demonstrate the validity of our CTI mitigation approach including a CDM and thus assess its present actual performance, we use the current best CDM candidate (later referred as to CDM for simplicity) as it is described in Short et al. (2010), and for which a first comparison of its outcomes to experimental test data is presented in Prod'homme et al. (2010).

CDM is based on the common Shockley Read Hall formalism (Shockley & Read 1952; Hall 1952) and describes the capture and release processes in a statistical way. To cope with the computational speed requirement, it suppresses the treatment of the numerous charge transfer steps required to transfer the signal from one CCD end to the other, but computes the signal transit in a single calculation making use of several assumptions (Short et al. 2010). CDM is able to simulate the CTI effects in TDI and imaging mode for any kind of signal (single, double stars, spectrum etc.). The CDM free parameters are: γ which determines how the volume of the electron packet grows as electrons are added, β the background light (respectively denoted β and S_{dob} in Short et al. (2010), but changed herein for disambiguation), and three trap parameters per trap species, ρ, σ and τ, respectively the trap density, the capture cross-section, and the release time constant. It has to be noted that a more recent version of this model has been elaborated. This newer version incorporates a better handling of the charge injection modelling and the possibility of simulating the serial CTI that occurs in the readout register. As charge injections are not explicitly simulated in our synthetic data set and the serial CTI was not simulated, the hereafter demonstrated performances remain representative of the current performances of our mitigation scheme.

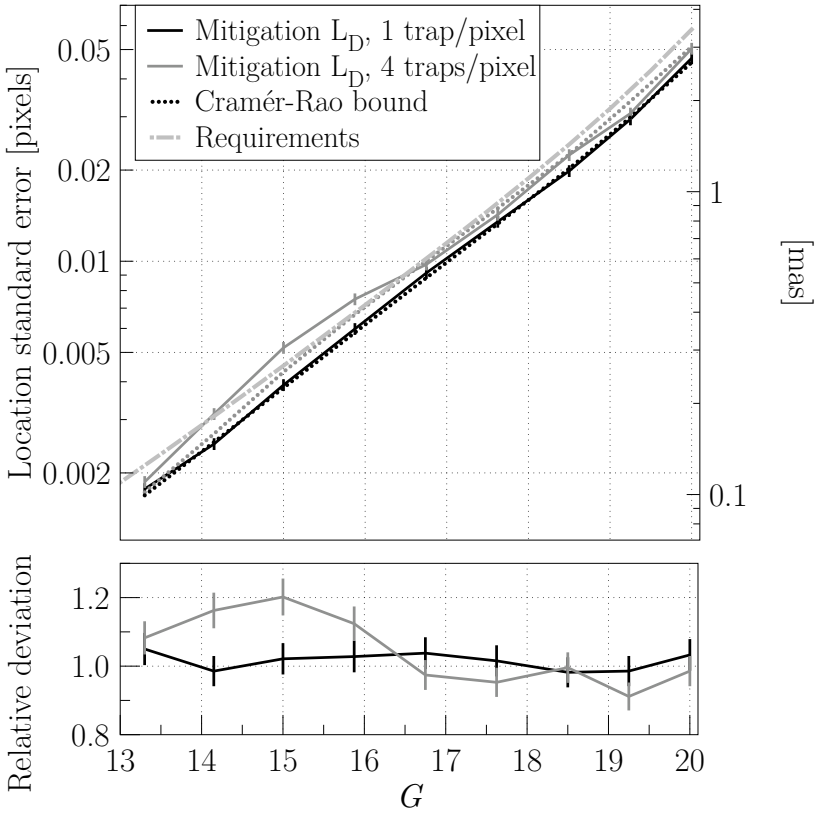


Figure 3.18 — Comparison between the Cramér-Rao bound in the presence of radiation damage (dotted lines) and the measured standard errors (continuous lines) when applying the Gaia image location estimation procedure including CTI mitigation to the set of damaged observations. These results are obtained in the Gaia operating conditions, considering the typical reference image only and for two different levels of radiation damage: 1 trap pixel^{-1} (black) and 4 traps pixel^{-1} (grey). The CTI mitigation is performed following the forward modelling approach, using $\tilde{L}_D$ (i.e. and ideally calibrated CDM and LSF model). In these conditions the estimator standard errors are below the requirements (dashed and dotted line) even for the most severe damage (top). The relative deviation from the best achievable accuracy (bottom) is reasonable (below 10%) in the 1 trap pixel^{-1} case but can reach 20% in the 4 traps pixel^{-1} case for the intermediate magnitudes.

3.6.5 The forward modelling approach initialization

The iterative image parameter estimation procedure including CTI mitigation now involves three different sets of parameters to be successively improved: the scene, the LSF/PSF, and the CDM parameters (see Fig. 3.15). Reaching a stable solution in these conditions is complex; each set of parameters needs to be initialized with values not too far off from the 'true' ones for the iterative procedure to converge.

LSF model

As already mentioned, the LSF model cannot be generated from the damaged observations, as they are not directly representative of the instrument anymore. During the mission the LSF model will partly be generated using the least damaged observations of single bright stars. In the following we thus use $\tilde{L}_U$, the LSF model generated from the CTI-free observations at a particular G. This constitutes a favourable yet realistic

case.

Scene parameters
The initial estimate for the image location $\kappa^{(0)}$ is determined using the Tukey's Biweight centroiding algorithm, the initial flux estimate $\alpha^{(0)}$ corresponds to the sum of the observed counts after background subtraction (and as in the rest of the study the background β is considered to be known). Hence the initial location and flux estimates are biased by the CTI effects. However one should note that due to its construction the LSF model contains some information about the true location of the observations in its zero-point. This is still reasonable as we have so far ignored that during the mission the astrometric solution (AGIS) will provide extra information about the true location of each observation through a feedback mechanism (see Fig. 3.15).

CDM parameters
It is first important to realize that several fundamental differences exist between CDM and the detailed Monte Carlo CTI effects model that we used to simulate the damaged observations: the most important ones being related to the charge transfer simulation, the computation of the capture and release probabilities, and the modelling of the electron density distribution. Hence, in this context, no 'true' CDM parameters exist but only CDM parameters that allow the reproduction of the simulated damaged observations. This actually constitutes a similar situation to the one that will be experienced during the operation of Gaia. Indeed, due to the simplifications intrinsic to the elaboration of a fast analytical model of a complex phenomenon, even with the right parameterization, the agreement between the damaged observations and the CDM predictions will not be perfect.
Furthermore: although the trapping occurs during the transfer of two-dimensional stellar images, only one-dimensional information is accessible from the binned observations. The CDM distortion can be applied to a one- or a two-dimensional CTI-free signal. In the latter case, one needs to reconstruct a PSF and the resulting modelled counts must be binned prior to a comparison with the observed damaged counts. In our study we generally obtained a significantly better agreement between CDM predictions and the damaged observations by applying the CDM distortion to a one-dimensional signal. In the following we thus only present results obtained in this case. In reality this might be different, in particular due to the serial CTI that was not taken into account here. Yet, if a comparable performance level can be achieved, the one-dimensional option would still be preferred during the mission for the 1D binned data as it presents the advantage of saving a significant number of computations.
To obtain an initial set of CDM parameters that describes reasonably well the damaged observations, we use $\tilde{L}_U$ as input signal, and fit the CDM predictions to the damaged observations $\{\{N_k\}\}_G$ for a particular G and set of operating conditions (i.e. windowing scheme, background and readout noise level). The fitting procedure minimizes the χ^2 (Eq. 3.11) between the CDM predictions and the damaged observations. The fitted parameters are γ, ρ, σ, and τ (see Section 3.6.4), β is fixed to the true value. At this stage, the fitting procedure is an evolutionary algorithm [1] that uses two mechanisms,

1. `http://watchmaker.uncommons.org/`

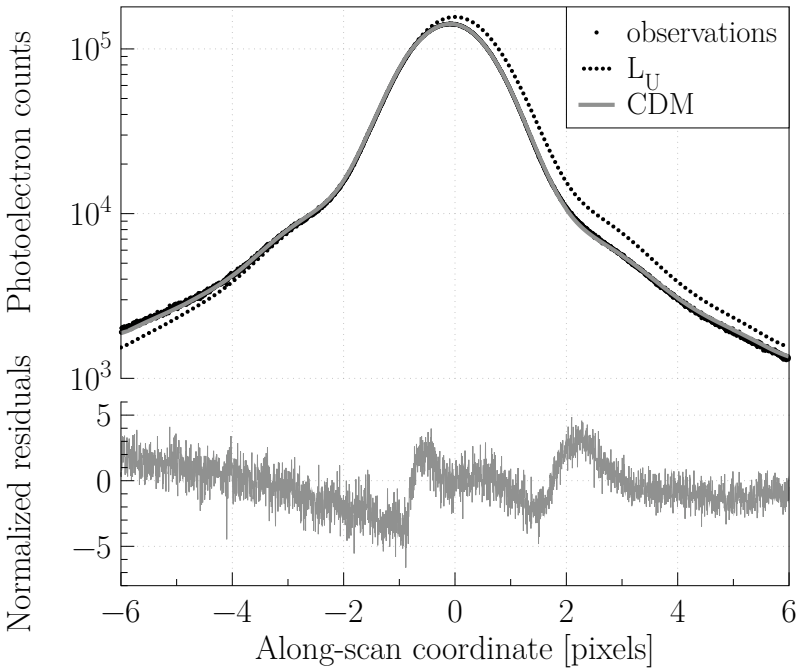


Figure 3.19 — Top: Comparison between the CDM predictions after initialization of the CDM parameters (grey) and the damaged observations $\{\{N_k\}\}_{13.3}$ (black dots) under the Gaia operating conditions and for a radiation level of 4 traps pixel^{-1}. To enable the observation of the CTI induced distortion the input signal $\tilde{L}_U$ is also shown (black dotted line). **Bottom:** Residuals normalized by the photon noise, the reduced χ^2 is $\sim$3.0.

mutation and cross-over. It is applied on an initial population of 100,000 parameter sets and evolves towards smaller χ^2 generation after generation. After 10 generations, we select the set of parameters with the smallest χ^2. This set of parameters can be further improved by using the downhill simplex minimization method (Nelder & Mead 1965). Fig. 3.19 gives an example of the obtained agreement between the CDM outcomes and the damaged 'observations' (generated with the Monte Carlo model described in Section 3.3) at a particular value of G. This example is representative of the best level of agreement achieved after applying the described initialization procedure. The illumination history parameters, **h**, will be fixed to the reconstructed illumination history. Here **h** is only the time since last CI that is set to infinity as no CI has been explicitly simulated. The effect of not calibrating for disturbing stars, i.e. stars located between the last CI and the star of interest, will be studied in the second part of this study. It can however already be mentioned that stars located between a CI and the star of interest are only disturbing if they are located in the same pixel column (or an adjacent one) and that, for a CI period of 1 s, the number of disturbing stars is expected to be very low even for the densest parts of the sky (see Chapter 4).

3.6.6 Image location bias and accuracy recovery

As mentioned in the previous section, after initialization, the CDM parameters can still be further improved by the use of the downhill simplex method. In the following

we thus distinguish between two different cases: (*a*) the CDM parameters have been fully optimized and the scene parameters are then estimated, no more iterations are performed, (*b*) the CDM parameters have not yet been fully optimized and are refined as part of the image location iterative procedure (Fig. 3.15). In the latter case, once each set of parameters is initialized, the iterative procedure is performed as follow: (i) the scene parameters are first estimated using $\tilde{L}_U$ and the CDM parameters $\gamma^{(0)}$, $\rho^{(0)}$, $\sigma^{(0)}$, $\tau^{(0)}$ and β, then (ii) the CDM parameters are updated as presented in Fig. 3.16 using the newly estimated scene parameters, and (iii) a new scene parameter update is performed. In our study, the CDM parameter update (Fig. 3.16) is performed using the downhill simplex minimization method (Nelder & Mead 1965) only; as we shall see it proves to be quite inefficient at this stage of the procedure. A maximum likelihood based procedure would be better suited and is currently being developed to perform this task in the Gaia data processing.

Figure 3.20 shows the remaining image location bias after using the CTI mitigation forward modelling approach including the current best CDM candidate for the two different initial optimizations of the CDM parameters. These results should be compared to Fig. 3.12 that shows the location bias when no CTI mitigation is applied, and

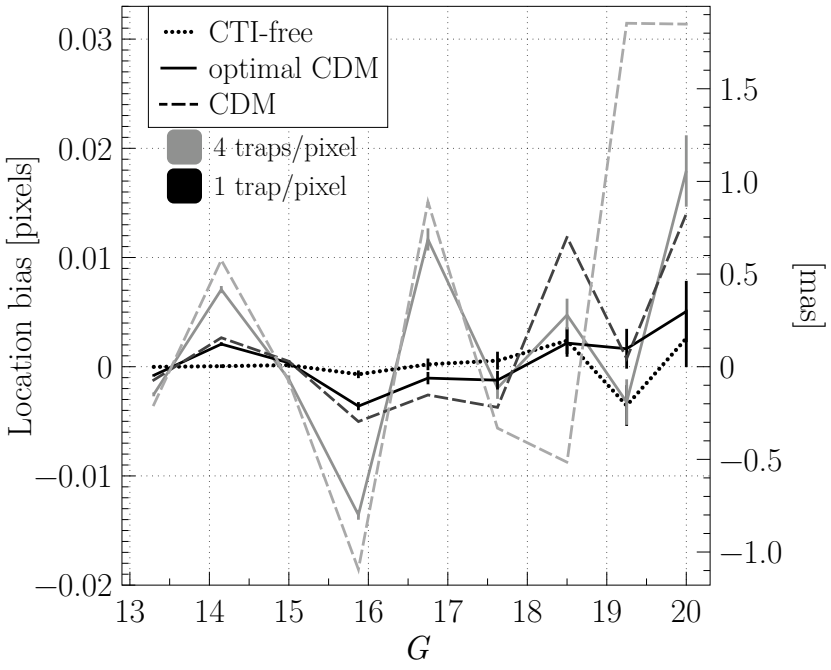


Figure 3.20 — Image location bias recovery using the CTI mitigation forward modelling approach including the current best CDM candidate. These results are obtained in the Gaia operating conditions, considering the typical reference image only, for two different levels of radiation damage (1 trap pixel^{-1} in black, and 4 traps pixel^{-1} in grey), and for two different initial optimizations of the CDM parameters: fully optimized (continuous line) and iteratively improved (dashed line). For comparison the average location bias measured for the CTI-free observations is also shown (dotted line). The CDM parameters are calibrated per magnitude and as a result there is a different set of CDM parameters for each magnitude and the resulting final agreement between modelled and observed counts varies from one signal level to the other. This explains the bias oscillations (in particular for the highest trap density).

to Fig. 3.17 that shows the ideal performance of the presented mitigation scheme. All these results are summarized in Table 3.4 for the three different reference images. The current best CDM candidate does not allow for a total recovery of the location bias at each magnitude, however this bias is considerably reduced for both levels of radiation damage. For instance, in the case of the lowest trap density and an optimal CDM optimization, the bias does not exceed the level of 0.005 pixels, while without any mitigation and in the same conditions the bias reaches 0.05 pixels (see Fig. 3.12). Fig. 3.20 shows that at the faint end the bias is significantly lowered by a better optimization of the CDM parameters, and that the CDM parameter update performed during the iterative procedure in these conditions is too limited to recover the full potential of CDM. This means that the CDM initialization is a crucial step of the CTI mitigation scheme, especially at faint magnitudes, and if the CDM parameters are iteratively refined, the simplex method seems not to be efficient enough. The bias oscillations as a function of G are due to the fact that the CDM parameters have been calibrated per magnitude. As a result a different level of agreement between observed and modelled counts is

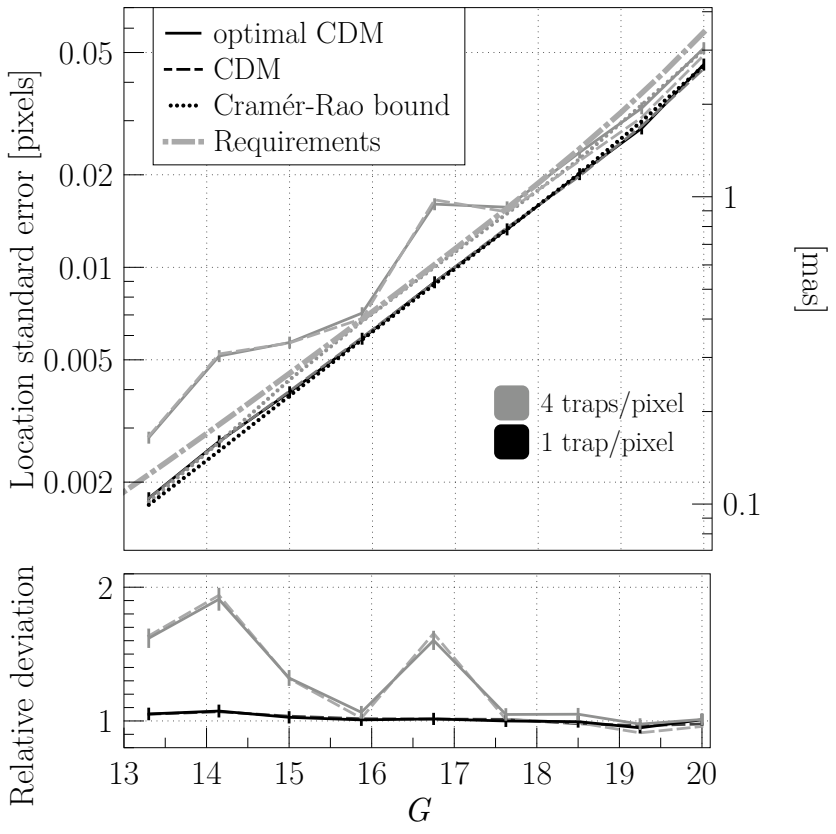


Figure 3.21 — Image location standard errors resulting from the use of the CTI mitigation forward modelling approach including the current best CDM candidate for two different levels of radiation damage (1 trap pixel^{-1} in black, and 4 traps pixel^{-1} in grey), and for two different initial optimization of the CDM parameters (fully optimized as a continuous line, and iteratively improved as a dashed line). The lower panel shows the relative deviation from the best achievable accuracy. The Cramér-Rao bound in the presence of radiation damage and the Gaia requirements are also shown (respectively the dotted lines and the dashed and dotted line).

achieved for each signal level.

Figure 3.21 shows the measured location standard errors for the typical reference image (see Table 3.5 for the other reference images). As expected (see Fig. 3.18) for the most severe level of radiation damage, the Gaia requirements are not met for the bright magnitudes, and the relative deviation from the best achievable accuracy is large: it almost reaches 100% at $G = 14.15$. However when considering the more realistic trap density of 1 trap pixel^{-1}, the standard errors are safely below the requirements, and the relative deviation from the Cramér-Rao bound remains below 10%. In these conditions our estimator thus remains efficient, even if, as already explained, biased. Finally, it has to be noted that the overall precision of the location estimation seems to be quite insensitive to the fine tuning of the CDM initial parameters.

3.6.7 Image flux bias recovery

This study focused on the estimation of the image location parameter as it is the most critical image parameter to be determined for astrometry. However the accurate estimation of the integrated image counts is also important as this forms the basis for the G-band photometry. The photometry, when combined with the parallax measurements, will provide the absolute luminosities of the stars observed by Gaia. In addition to this fundamental parameter the multiple observations of each source constitute an all sky variability survey, providing another treasure trove of astrophysical information. For both applications high photometric accuracy, complementary to the astrometric accuracy is required. We refer to Jordi et al. (2010) for more details on Gaia's photometric capabilities. As can be observed from Fig. 3.2, CTI induces not only an image distortion but also an important charge loss that, if not properly taken into account, biases the image flux estimation. In order to judge the capability of our approach to CTI mitigation to achieve high photometric accuracy we show in Fig. 3.22 the image flux estimation bias (in units of magnitude) achieved with and without CTI mitigation using CDM. These biases are compared to the photometric performance predictions in Jordi et al. (2010) including the intrinsic loss in photometric precision induced by CTI (Jordi et al. 2010), where the numbers in their figure 19 have been translated to the photometric errors expected for a transit across a single CCD. The 'safety margin' in equation (6) in Jordi et al. (2010) was omitted.

Despite the rather strong bias induced by CTI in the image flux estimation (see Fig. 3.22 top), the presented CTI mitigation scheme allows to eliminate most of it (see Fig. 3.22 bottom). It thus allows an unbiased estimation of the image flux within the requirements if the estimation procedure is precise enough. In this context the current CDM performances are remarkable: for the highest trap density and when no mitigation is applied, the flux bias can reach $\sim$0.25 magnitudes at $G = 15.875$, after mitigation we measure a flux estimation bias of 0.0029 magnitudes for the same magnitude, well below the requirement of 0.0067 magnitudes. It is also interesting to note that the image flux estimation is much less sensitive to the calibration of CDM than the image location estimation. Indeed Fig. 3.22 (bottom) shows that a similar level of performance is obtained for a fully optimized or an iteratively improved CDM calibration.

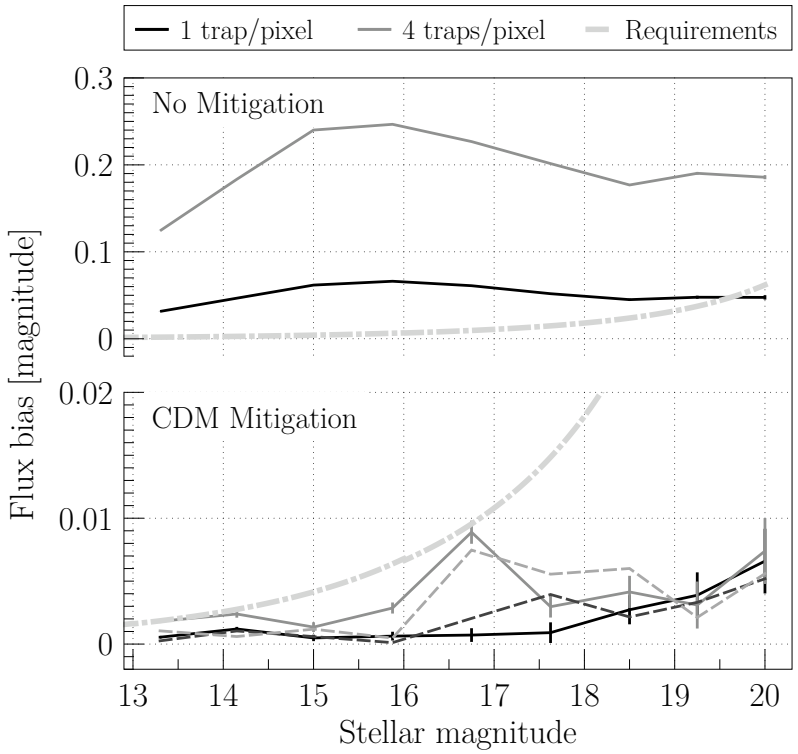


Figure 3.22 — Comparison between the image flux estimation bias (in units of magnitude) before (top) and after CTI mitigation (bottom) in the Gaia operating conditions considering the typical reference image for two different levels of trap density: 1 trap pixel^{-1} (black) and 4 traps pixel^{-1} (grey). These results are compared to the Gaia requirements (Jordi et al. 2010) for a transit across a single CCD (dashed and dotted line). In the case of the CTI mitigation (bottom) we considered two different initial optimizations of the CDM parameters: fully optimized (continuous lines) and iteratively improved (dashed lines).

3.7 Discussion

Throughout this chapter we assumed a fully functional SBC. Nevertheless the manufacturing of a SBC is a complex process and the CCDs of the Wield Field Camera channel of the Advanced Camera for Surveys on board HST have been reported not to contain the SBC present in their design (Anderson & Bedin 2010). Fig. 3.23 shows the importance of the SBC CTI mitigation for achieving the Gaia requirements at low signal levels. If the SBC was not present (dashed line), the intrinsic loss of accuracy at low signal levels would reach 200% instead of only 5% in the presence of a SBC (continuous line and Fig. 3.11, bottom). The extra location bias induced by a missing SBC is also significant although there is no particular reason for which the presented CTI mitigation scheme would not be able to recover it. In the case of Gaia the SBC has been demonstrated to be functional using experimental tests. For instance, Fig. 5 in Chapter 2 shows the effect of the SBC on the CTI-induced fractional charge loss as a function of signal level. However a recent study by Kohley et al. (2009) identified a non-functional SBC in the upper half of a Gaia CCD. Based on more tested devices, Chapter 6 shows that a significant number of the Gaia CCDs could be affected by this issue. Using the same methods as in this chapter, we evaluate the extra loss of accu-

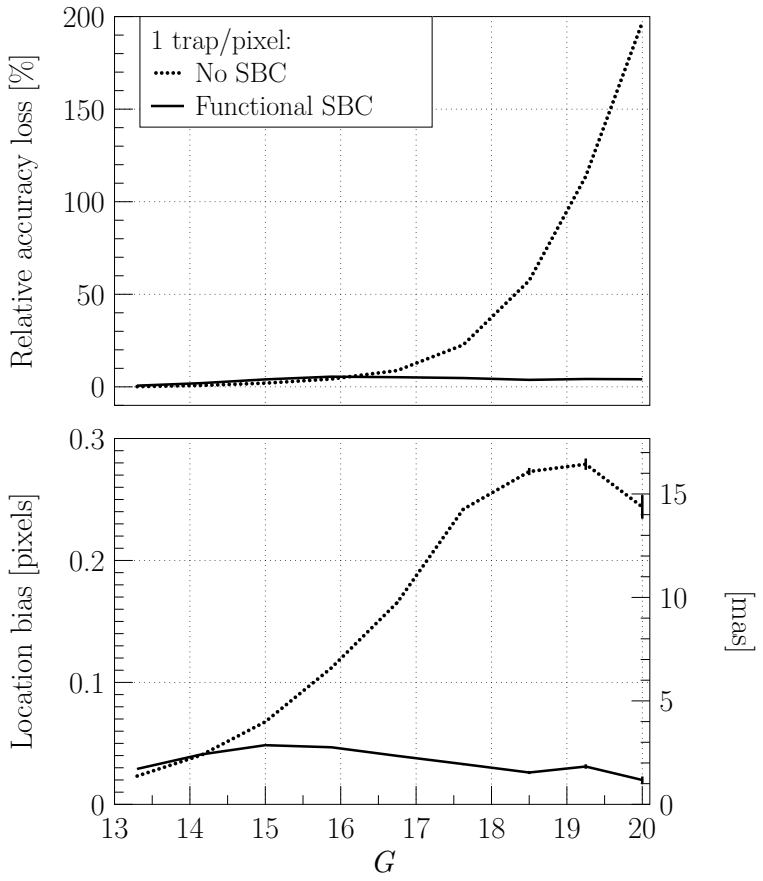


Figure 3.23 — Top: The relative intrinsic loss of accuracy induced by radiation damage as a function of G and as computed in Section 3.5.5. The continuous line is the same as the black dashed line in Fig. 3.11 and corresponds to the case where the SBC is functional. The dashed line is obtained using simulated transits for a CCD containing no SBC. Note the difference in ordinate scale range with Fig. 3.11 (bottom). **Bottom:** Absolute image location bias as a function of G and as computed in Section 3.5.6. The continuous line is the same as the black continuous line in Fig. 3.12 (typical image width, Gaia operating conditions, 1 trap pixel^{-1}) and corresponds to a functional SBC case. The dashed line was obtained for simulating star transits in the same conditions but with a CCD containing no SBC.

racy and location bias induced by a non-functional SBC in the CCD upper half only; in contrast to the no SBC case investigated here, Chapter 6 shows that there is only at most 10% extra loss of accuracy.

3.8 Conclusions

We have presented a detailed characterization and evaluation of the impact of CCD radiation damage on the Gaia image location accuracy. The underlying principle of this study consisted of a systematic comparison between the computed theoretical limit to the image location accuracy, the Cramér-Rao bound, and the actual performance of the Gaia image parameter estimation procedure under realistic Gaia operating conditions. The image location estimation bias and the associated standard errors were measured by applying the Gaia image parameter estimation procedure to a large set of synthetic

data accounting for different stellar image widths, magnitudes, and background levels. We considered two different active trap densities; they are representative of two different levels of hardware mitigation in the presence of a radiation dose equivalent to the upper limit to the expected end-of-mission accumulated dose. The lowest active trap density constitutes the most realistic case because of a lower expected radiation dose and more importantly because of the planned injection of artificial charges in the CCD every second. In this context, a total of 41,472 synthetic two-dimensional Gaia-like observations have been generated using a detailed Monte Carlo model of the CCD charge collection and transfer, and the radiation induced trapping. The dataset is readily available for the Gaia scientists to continue to test and further improve these critical steps in the Gaia data processing that are the image parameter estimation and the mitigation of the CTI effects. The main conclusions we can draw from this study are:

The Gaia image location estimation procedure is bias-free and efficient in the absence of radiation damage. We showed that under realistic operating conditions, and from bright to faint magnitudes, the Gaia location estimator performs within the requirements at an accuracy close to the theoretical limit for CTI-free observations.

The radiation damage effects induce an irreversible loss of accuracy that is independent of any image location estimator. It can only be avoided by the use of hardware CTI countermeasures that physically prevent the trapping. In the theoretical limit (i.e. perfect CTI calibration at the image processing level), the location accuracy loss is still acceptable when compared to the Gaia requirements: it can reach 6% in the lowest trap density case and 24% in the highest case. Due to the presence of a supplementary buried channel in each of the Gaia CCD pixels, the accuracy loss stops increasing for stars fainter than $G \sim 16$.

A CTI mitigation procedure is critical to achieve the Gaia requirements. We showed that if CTI is not taken into account in the image parameter estimation procedure, the resulting image location estimations are significantly biased. In the Gaia operating conditions, the most important bias is obtained for the widest type of stellar images at $G = 15$: 0.05 pixels in the lowest trap density case and 0.2 pixels in the highest case. For comparison, at this magnitude the requirement on the image location accuracy for a G2V type star is 0.0045 pixels, at least an order of magnitude smaller than the measured bias.

The CTI-induced image location bias varies significantly with the stellar image width and the background level. This is particularly relevant for experimental studies in which the image flux distribution and the background level cannot be absolutely known. At faint magnitudes small differences in the experimental setup can lead to significant differences in the measured CTI effects.

In principle, a complete forward modelling approach to CTI mitigation allows for an accurate and bias-free estimation of the true image location from a damaged observation. We demonstrated that the forward modelling of a damaged observation

using ideally calibrated models for both the PSF/LSF and the CTI effects provides a location estimate that on average never exceeds 0.003 pixels and does not deviate from zero within the error bars from bright to faint stars and for the two considered levels of radiation damage. The accuracy reached using this CTI mitigation scheme complies with the Gaia requirements. In the case of the lowest active trap density, this method even allows for the recovery of the theoretical limit to the image location accuracy in the presence of radiation damage.

If calibrated well enough, the current best candidate for the charge distortion model (CDM) associated with the forward modelling approach allows significant image location and flux estimate bias recovery. In these favourable conditions (simple illumination history, 1 trap species, no serial CTI, well calibrated LSF model, and close to optimal CDM parameters) yet for a trap density level representative of the end-of-mission accumulated radiation dose, the Gaia image location accuracy is preserved. In the Gaia operating conditions and after CTI mitigation using the current best CDM at our disposal, the maximum measured location bias is $\sim$0.005 pixels for the lowest radiation level and $\sim$0.017 pixels for the highest.

3.9 Future work

Estimating the location of an image to milli-pixel accuracy is an extremely challenging exercise, in which no detail must be neglected. This is especially true in the presence of radiation damage as shown in this chapter. The work presented here is not the final word on the Gaia image parameter estimation procedure. Indeed the estimation procedure must be tested and improved further using synthetic and experimental data. In addition the elaboration and the calibration of the charge distortion model is a key element in the success of the presented CTI mitigation scheme. In this study, we have established the level of agreement with the damaged observation that any CDM must achieve to recover a bias-free image location estimation. In Chapter 5 we test if the current best CDM candidate is capable of reaching such agreement with experimental data and suggest necessary improvements. Regarding the calibration of the radiation damage parameters, the periodic charge injections will enable us to monitor and characterize the radiation damage during the mission. In addition the charge injections will act to reset the illumination history. We intend to study what are the parameters that one can infer from the study of the CTI effects on the charge injection signal, and how one can use these parameters to initialize and calibrate the charge distortion model.
This study constitutes the first step in evaluating the impact of the CCD radiation damage on the final astrometric accuracy of Gaia. It is indeed not yet clear in detail how a biased and less precise estimation of the image location, as induced by CTI, propagates into the astrometric parameters derived by the Gaia astrometric global iterative solution (AGIS). In the next chapter of this thesis we investigate this particular question by using a small-scale version of AGIS, AGISLab (Holl et al. 2009), that will allow us to perform a careful error propagation analysis for different cases (no CTI mitigation and optimal mitigation versus CTI-free). The study presented here is used to construct for

each case a model that provides location bias and uncertainty as function of magnitude, shape of the stellar profile, illumination history (time since last charge injection), and mission time (or trap density). These models are used to disturb the image locations (observation times) processed by AGIS. The study of the resulting astrometric parameters allow us to characterize and evaluate the impact of CCD radiation damage on Gaia's astrometry. Also the effect of disturbing stars between a charge injection and a target star is assesed.
The next steps outlined are crucial ingredients in the successful radiation damage mitigation strategy for Gaia, enabling the extraction of the best scientific performance from this exciting and much anticipated mission.

Acknowledgments

The work of TP and BH was supported by the European Marie-Curie research training network ELSA (MRTN-CT-2006-033481). LL acknowledges support by the Swedish National Space Board. AB acknowledges support by the Netherlands Research School for Astronomy (NOVA). The authors would like to kindly thank Carme Jordi for providing the Gaia photometric requirement values and Jos de Bruijne for the detailed feedback given on the conversion of the Gaia astrometric requirements. The authors would also like to thank the referee L.R. Bedin for his comments and suggestions that significantly improved this chapter and the corresponding paper to be published in MNRAS.

Reference image type	typical		typical		typical		typical	
Background level (e^-pixel^{-1} s^{-1})	0		0		0.44698		0.44698	
Readout noise (e^-)	0		4.35		4.35		4.35	
Window size (pixels)	40		40		40		18	
Trap density (traps pixel^{-1})	0		0		0		0	
Magnitude (G-band)	Cramér-Rao bound (10^{-3} pixel)	$\sigma_\kappa \pm \upsilon_{\sigma_\kappa}$ (10^{-3} pixel)						
13.3	1.67	1.668 ± 0.075	1.67	1.672 ± 0.075	1.68	1.655 ± 0.074	1.67	1.662 ± 0.074
14.15	2.48	2.49 ± 0.11	2.48	2.50 ± 0.11	2.48	2.46 ± 0.11	2.48	2.45 ± 0.11
15.0	3.66	3.69 ± 0.17	3.67	3.71 ± 0.17	3.67	3.77 ± 0.17	3.67	3.75 ± 0.17
15.875	5.48	5.95 ± 0.27	5.50	6.00 ± 0.27	5.50	5.45 ± 0.24	5.51	5.63 ± 0.25
16.75	8.29	7.65 ± 0.34	8.33	7.65 ± 0.34	8.34	8.46 ± 0.38	8.30	8.34 ± 0.37
17.625	12.40	12.65 ± 0.57	12.54	12.73 ± 0.57	12.57	12.63 ± 0.56	12.61	12.70 ± 0.57
18.5	18.55	17.88 ± 0.80	18.99	18.03 ± 0.81	19.09	18.84 ± 0.84	19.37	19.34 ± 0.86
19.25	26.29	25.87 ± 1.16	27.49	26.56 ± 1.19	27.62	27.83 ± 1.24	28.53	28.86 ± 1.29
20.0	37.06	37.12 ± 1.66	40.26	39.99 ± 1.79	40.99	38.08 ± 1.70	43.40	39.08 ± 1.75
Reference image type	typical		narrow		typical		wide	
Background level (e^-pixel^{-1} s^{-1})	0		0.44698		0.44698		0.44698	
Readout noise (e^-)	0		4.35		4.35		4.35	
Window size (pixels)	telemetry*		telemetry		telemetry		telemetry	
Trap density (traps pixel^{-1})	0		0		0		0	
Magnitude (G-band)	Cramér-Rao bound (10^{-3} pixel)	$\sigma_\kappa \pm \upsilon_{\sigma_\kappa}$ (10^{-3} pixel)						
13.3	1.67	1.680 ± 0.075	1.50	1.416 ± 0.063	1.67	1.658 ± 0.074	1.95	1.920 ± 0.086
14.15	2.48	2.54 ± 0.11	2.23	2.25 ± 0.10	2.48	2.46 ± 0.11	2.88	2.83 ± 0.13
15.0	3.66	3.69 ± 0.17	3.30	3.52 ± 0.16	3.67	3.74 ± 0.17	4.27	4.52 ± 0.20
15.875	5.48	5.98 ± 0.27	4.95	5.47 ± 0.24	5.51	5.59 ± 0.25	6.42	5.69 ± 0.25
16.75	8.29	7.69 ± 0.34	7.50	7.55 ± 0.34	8.38	8.45 ± 0.38	9.85	9.80 ± 0.44
17.625	12.40	12.66 ± 0.57	11.39	11.13 ± 0.50	12.70	12.94 ± 0.58	14.94	15.55 ± 0.70
18.5	18.55	17.99 ± 0.80	17.44	16.88 ± 0.76	19.49	19.39 ± 0.87	22.99	25.09 ± 1.12
19.25	26.29	25.87 ± 1.16	25.48	26.33 ± 1.18	28.65	29.53 ± 1.32	33.73	31.20 ± 1.40
20.0	37.06	37.12 ± 1.66	38.43	39.73 ± 1.78	43.57	41.66 ± 1.86	51.11	51.94 ± 2.32

Table 3.2 — Comparison between the theoretical and the actual limit to the image location accuracy in the absence of radiation damage as a function of G, sky background, readout noise, image width, and size of the telemetry windows in the along-scan direction. The theoretical limit corresponds to the Cramér-Rao bound, and the actual to the Gaia image location estimator standard errors σ_κ with υ_{σ_κ} the statistical uncertainty (Eq. 3.16). While the image width has a significant impact on those limits (e.g., for a 20% increase in FWHM, one can note a $\sim$25% decrease in accuracy for the faintest magnitude), the window size, readout noise, and background level only slightly affects the image location accuracy.
* 'telemetry size' refers to the size of the windows as they will be transmitted to the ground segment during the operational phase of Gaia (Paulet 2009): 12 pixels in the along-scan direction for $G < 16$, and 6 pixels for $G > 16$.

Reference image type	typical			narrow			typical			wide		
Background level (e^-pixel^{-1} s^{-1})	0			0.44698			0.44698			0.44698		
Readout noise (e^-)	4.35			4.35			4.35			4.35		
Window size (pixels)	telemetry			telemetry			telemetry			telemetry		
Trap density (traps pixel^{-1})	0	1	4	0	1	4	0	1	4	0	1	4
Magnitude (G-band)	Cramér-Rao bound (10^{-3} pixel)											
13.3	1.67	1.69	1.72	1.51	1.52	1.55	1.68	1.69	1.73	1.95	1.96	2.01
14.15	2.48	2.53	2.69	2.23	2.27	2.40	2.48	2.53	2.70	2.88	2.94	3.15
15.0	3.67	3.81	4.28	3.30	3.43	3.87	3.67	3.82	4.31	4.27	4.45	5.02
15.875	5.50	5.79	6.61	4.95	5.25	6.10	5.51	5.82	6.67	6.42	6.73	7.64
16.75	8.33	8.72	9.97	7.53	7.98	9.23	8.41	8.85	10.06	9.97	10.43	11.68
17.625	12.54	13.08	14.89	11.44	12.01	13.69	12.74	13.35	15.01	15.11	15.70	17.46
18.5	18.99	19.71	22.16	17.50	18.22	20.51	19.54	20.28	22.59	23.24	23.98	26.47
19.25	27.49	28.54	33.12	25.54	26.78	30.51	28.72	29.93	33.84	34.01	35.50	39.57
20.0	40.26	43.27	53.28	38.51	40.27	45.96	43.64	45.43	51.19	51.43	53.54	60.32

Table 3.3 — Comparison between the Cramér-Rao bounds computed for different image widths, background levels, and radiation damage levels. This comparison allows to characterize the intrinsic loss of precision and ultimately accuracy induced by radiation damage. This loss is relatively more important for the narrowest image, and increase with the trap density as expected.

Reference image type	narrow		typical		wide	
Background level (e^-pixel^{-1} s^{-1})	0.44698		0.44698		0.44698	
Readout noise (e^-)	4.35		4.35		4.35	
Window size (pixels)	telemetry		telemetry		telemetry	
Trap density (traps pixel^{-1})	1	4	1	4	1	4
Mitigation	none		none		none	
Magnitude (G-band)	Location bias $\langle\delta_\kappa\rangle_G \pm \upsilon_{\langle\delta_\kappa\rangle}$ (10^{-3} pixel)					
13.3	24.85 ± 0.11	93.48 ± 0.27	29.06 ± 0.11	110.84 ± 0.22	35.64 ± 0.13	137.18 ± 0.19
14.15	35.79 ± 0.14	137.10 ± 0.22	40.93 ± 0.17	156.83 ± 0.29	49.68 ± 0.20	188.40 ± 0.28
15.0	43.54 ± 0.24	166.65 ± 0.66	48.53 ± 0.23	180.64 ± 0.47	56.48 ± 0.27	207.34 ± 0.54
15.875	44.06 ± 0.44	161.29 ± 1.11	46.81 ± 0.43	170.11 ± 0.82	51.19 ± 0.48	187.06 ± 0.74
16.75	37.70 ± 0.56	139.54 ± 1.12	39.75 ± 0.60	143.66 ± 0.93	43.03 ± 0.64	157.81 ± 0.97
17.625	31.26 ± 0.81	120.41 ± 1.18	33.02 ± 0.87	126.87 ± 1.04	33.66 ± 0.99	139.18 ± 1.20
18.5	26.09 ± 1.07	108.10 ± 1.29	26.11 ± 1.28	114.60 ± 1.48	33.79 ± 1.58	130.81 ± 1.50
19.25	28.67 ± 1.64	105.25 ± 1.92	31.02 ± 1.81	109.66 ± 2.03	25.92 ± 2.14	119.33 ± 2.20
20.0	24.92 ± 2.51	84.61 ± 2.73	19.97 ± 2.86	87.34 ± 3.15	24.16 ± 3.57	95.08 ± 3.71
Mitigation	ideal		ideal		ideal	
Magnitude (G-band)	Location bias $\langle\delta_\kappa\rangle_G \pm \upsilon_{\langle\delta_\kappa\rangle}$ (10^{-3} pixel)					
13.3	0.01 ± 0.10	0.02 ± 0.12	-0.05 ± 0.11	0.18 ± 0.12	0.09 ± 0.13	-0.19 ± 0.13
14.15	-0.16 ± 0.14	0.05 ± 0.16	0.15 ± 0.16	0.11 ± 0.20	-0.17 ± 0.20	-0.15 ± 0.20
15.0	0.27 ± 0.22	0.55 ± 0.36	-0.32 ± 0.25	0.29 ± 0.33	-0.25 ± 0.28	0.11 ± 0.32
15.875	-0.12 ± 0.36	-0.30 ± 0.52	0.12 ± 0.38	1.03 ± 0.47	-0.33 ± 0.45	0.44 ± 0.47
16.75	0.00 ± 0.50	0.51 ± 0.63	-0.14 ± 0.58	0.40 ± 0.62	-0.07 ± 0.65	0.81 ± 0.70
17.625	-1.63 ± 0.77	0.25 ± 0.87	0.22 ± 0.85	0.00 ± 0.90	-1.82 ± 0.97	1.04 ± 1.10
18.5	-0.54 ± 1.07	-1.94 ± 1.24	0.90 ± 1.26	-1.52 ± 1.42	2.04 ± 1.58	-0.38 ± 1.51
19.25	-0.28 ± 1.66	-1.98 ± 1.91	-2.09 ± 1.86	-0.67 ± 1.95	1.30 ± 2.15	-0.75 ± 2.18
20.0	4.84 ± 2.44	0.30 ± 2.91	-3.08 ± 2.96	-2.06 ± 3.18	-2.63 ± 3.53	6.64 ± 3.58
Mitigation	CDM		CDM		CDM	
Magnitude (G-band)	Location bias $\langle\delta_\kappa\rangle_G \pm \upsilon_{\langle\delta_\kappa\rangle}$ (10^{-3} pixel)					
13.3	-1.64 ± 0.11	-6.55 ± 0.17	-0.79 ± 0.11	-2.53 ± 0.18	0.89 ± 0.13	5.41 ± 0.17
14.15	0.13 ± 0.16	3.67 ± 0.33	2.07 ± 0.17	7.06 ± 0.33	3.04 ± 0.20	12.49 ± 0.34
15.0	-0.71 ± 0.27	-0.69 ± 0.48	0.35 ± 0.25	-1.16 ± 0.36	3.42 ± 0.28	3.94 ± 0.36
15.875	-1.18 ± 0.32	-7.15 ± 0.49	-3.61 ± 0.37	-13.56 ± 0.45	-3.60 ± 0.43	-18.86 ± 0.48
16.75	1.21 ± 0.50	22.88 ± 1.14	-1.04 ± 0.57	11.65 ± 1.02	-2.98 ± 0.63	7.38 ± 1.04
17.625	-3.11 ± 0.76	5.06 ± 0.96	-1.23 ± 0.85	-1.95 ± 0.99	-9.33 ± 0.96	-3.82 ± 1.19
18.5	8.54 ± 1.09	-4.25 ± 1.24	2.16 ± 1.25	4.73 ± 1.50	9.13 ± 1.54	-10.50 ± 1.60
19.25	2.12 ± 1.62	35.19 ± 1.95	1.65 ± 1.82	-3.24 ± 2.08	-9.38 ± 2.03	23.92 ± 2.20
20.0	23.62 ± 2.39	25.12 ± 2.89	5.06 ± 2.80	17.92 ± 3.27	20.99 ± 3.50	11.74 ± 3.84

Table 3.4 — Summary of the measured image location biases in the Gaia operating conditions for different stellar image widths, radiation levels, and different levels of mitigation: (i) 'none' corresponds to no mitigation (Section 3.5.6), (ii) 'ideal' to the presented forward modelling approach associated to an ideal CDM and LSF model and calibration (Section 3.6.3), and (iii) 'CDM' to the presented forward modelling approach including the current implementation of CDM (Section 3.6.6). Note that in the latter case, the optimization of the CDM corresponds to the fully optimized case as described in Sections 3.6.5 and 3.6.6. This optimization was performed for the typical reference image only, the same CDM parameters were used for the two other reference images.

Reference image type	narrow		typical		wide	
Background level (e^-pixel^{-1} s^{-1})	0.44698		0.44698		0.44698	
Readout noise (e^-)	4.35		4.35		4.35	
Window size (pixels)	telemetry		telemetry		telemetry	
Trap density (traps pixel^{-1})	1	4	1	4	1	4
Mitigation	none		none		none	
Magnitude (G-band)	Standard errors $\sigma_\kappa \pm \upsilon_{\sigma_\kappa}$ (10^{-3} pixel)					
13.3	1.80 ± 0.08	4.26 ± 0.19	2.00 ± 0.09	3.50 ± 0.16	2.08 ± 0.09	3.08 ± 0.14
14.15	2.31 ± 0.10	3.55 ± 0.16	2.67 ± 0.12	4.56 ± 0.20	3.16 ± 0.14	4.43 ± 0.20
15.0	3.75 ± 0.17	10.50 ± 0.47	4.29 ± 0.19	7.43 ± 0.33	4.40 ± 0.20	8.59 ± 0.38
15.875	7.10 ± 0.32	17.58 ± 0.79	6.59 ± 0.29	12.91 ± 0.58	7.62 ± 0.34	11.69 ± 0.52
16.75	9.02 ± 0.40	17.78 ± 0.80	9.48 ± 0.42	14.71 ± 0.66	10.05 ± 0.45	15.36 ± 0.69
17.625	13.04 ± 0.58	18.63 ± 0.83	13.60 ± 0.61	16.47 ± 0.74	15.79 ± 0.71	18.95 ± 0.85
18.5	16.96 ± 0.76	20.39 ± 0.91	20.49 ± 0.92	23.48 ± 1.05	24.99 ± 1.12	23.72 ± 1.06
19.25	26.44 ± 1.18	30.43 ± 1.36	29.02 ± 1.30	32.12 ± 1.44	33.86 ± 1.51	34.77 ± 1.56
20.0	37.29 ± 1.67	43.24 ± 1.93	45.83 ± 2.05	49.85 ± 2.23	56.45 ± 2.52	58.65 ± 2.62
Mitigation	ideal		ideal		ideal	
Magnitude (G-band)	Standard errors $\sigma_\kappa \pm \upsilon_{\sigma_\kappa}$ (10^{-3} pixel)					
13.3	1.57 ± 0.07	1.91 ± 0.09	1.77 ± 0.08	2.01 ± 0.09	1.87 ± 0.08	2.05 ± 0.09
14.15	2.24 ± 0.10	2.55 ± 0.11	2.49 ± 0.11	3.09 ± 0.14	3.13 ± 0.14	3.12 ± 0.14
15.0	3.45 ± 0.15	5.90 ± 0.26	3.90 ± 0.17	5.01 ± 0.22	5.18 ± 0.23	4.37 ± 0.20
15.875	5.70 ± 0.25	8.20 ± 0.37	5.98 ± 0.27	7.49 ± 0.34	7.49 ± 0.34	7.10 ± 0.32
16.75	7.92 ± 0.35	9.72 ± 0.43	9.15 ± 0.41	11.04 ± 0.49	9.76 ± 0.44	10.28 ± 0.46
17.625	12.19 ± 0.55	13.36 ± 0.60	13.51 ± 0.60	17.43 ± 0.78	14.26 ± 0.64	15.39 ± 0.69
18.5	16.89 ± 0.76	18.50 ± 0.83	19.86 ± 0.89	23.80 ± 1.06	22.42 ± 1.00	24.94 ± 1.12
19.25	26.21 ± 1.17	30.35 ± 1.36	29.43 ± 1.32	34.44 ± 1.54	30.76 ± 1.38	34.00 ± 1.52
20.0	38.64 ± 1.73	45.00 ± 2.01	46.82 ± 2.09	56.58 ± 2.53	50.36 ± 2.25	55.76 ± 2.49
Mitigation	CDM		CDM		CDM	
Magnitude (G-band)	Standard errors $\sigma_\kappa \pm \upsilon_{\sigma_\kappa}$ (10^{-3} pixel)					
13.3	1.67 ± 0.07	2.72 ± 0.12	1.77 ± 0.08	2.79 ± 0.12	2.03 ± 0.09	2.65 ± 0.12
14.15	2.47 ± 0.11	5.21 ± 0.23	2.72 ± 0.12	5.15 ± 0.23	3.18 ± 0.14	5.38 ± 0.24
15.0	4.33 ± 0.19	7.54 ± 0.34	3.92 ± 0.18	5.69 ± 0.25	4.42 ± 0.20	5.71 ± 0.26
15.875	5.02 ± 0.22	7.74 ± 0.35	5.86 ± 0.26	7.10 ± 0.32	6.77 ± 0.30	7.55 ± 0.34
16.75	7.91 ± 0.35	17.99 ± 0.80	8.95 ± 0.40	16.05 ± 0.72	9.96 ± 0.45	16.51 ± 0.74
17.625	12.02 ± 0.54	15.12 ± 0.68	13.30 ± 0.59	15.69 ± 0.70	15.11 ± 0.68	18.74 ± 0.84
18.5	17.23 ± 0.77	19.66 ± 0.88	20.12 ± 0.90	23.65 ± 1.06	24.31 ± 1.09	25.32 ± 1.13
19.25	25.55 ± 1.14	30.77 ± 1.38	28.33 ± 1.27	32.95 ± 1.47	32.16 ± 1.44	34.73 ± 1.55
20.0	37.83 ± 1.69	45.73 ± 2.05	45.74 ± 2.05	51.67 ± 2.31	55.35 ± 2.48	60.64 ± 2.71

Table 3.5 — Summary of the measured image location estimator standard errors in the Gaia operating conditions for different stellar image widths, radiation levels, and different levels of mitigation: (i) 'none', (ii) 'ideal', and (iii) 'CDM' (see Table 3.4).

Chapter 4

The impact of CCD radiation damage on Gaia astrometry: II. Effect of image location errors on the astrometric solution

Gaia, the next astrometric mission of the European Space Agency, will use a camera composed of 106 CCDs to collect multiple observations for one billion stars. The astrometric core solution of Gaia will use the estimated location of the stellar images on the CCDs to derive the astrometric parameters (position, parallax, and proper motion) of the stars. The Gaia CCDs will suffer from Charge Transfer Inefficiency (CTI) mainly caused by radiation damage. CTI is expected to significantly degrade the quality of the collected images which ultimately affects the astrometric accuracy of Gaia. This chapter is the second and last part in a study aiming at characterizing and quantifying the impact of CCD radiation damage on Gaia astrometry. Here we focus on the effect of the image location errors induced by CTI on the astrometric solution. We apply the Gaia Astrometric Global Iterative Solution (AGIS) to simulated Gaia-like observations for 1 million stars including CTI-induced errors as described in Chapter 3. We show that a magnitude-dependent image location bias is propagated in the astrometric solution, biassing the estimation of the astrometric parameters as well as decreasing its precision. We demonstrate how the Gaia scanning law dictates this propagation and the ultimate sky distribution of the CTI induced errors. The possibility of using the residuals of the astrometric solution to improve the calibration of the CTI effects is investigated. We also estimate the astrometric errors caused by (faint) disturbing stars preceding the stellar measurements on the CCDs. Finally we show that, for single stars, the overall astrometric accuracy of Gaia can be preserved to within 10 per cent of the CTI-free case for all magnitudes by appropriate modelling at the image location estimation level and using the solution residuals.

B. Holl, T. Prod'homme, L. Lindegren, and A.G.A. Brown
Submitted to Monthly Notices of the Royal Astronomical Society, 2011

4.1 Introduction

This chapter, in combination with the previous study presented in Chapter 3, provides the first detailed evaluation of the impact of radiation damage effects on the astrometric performance of Gaia, the European Space Agency astrometric mission scheduled for launch in 2013. Gaia will observe one billion stars to produce a catalogue of positions, parallaxes, proper motions, and photometric data, as well as radial velocities and astrophysical parameters for many of the stars (Perryman et al. 2001; Lindegren 2010).

Radiation damage of the CCD detectors in the space environment has been identified as a potential threat to achieving the most demanding scientific goals of Gaia, requiring astrometric accuracies in the 10–20 μas range for the brighter stars (e.g., Lindegren et al. 2008). It is therefore important to study the impact of the radiation damage on all stages of the data acquisition and analysis, from the individual CCD pixels to the final catalogue.

In Chapter 3 we investigated the effects of the radiation damage on the image location process. This is central to the astrometric performance of the mission, as all positional information is ultimately derived from the very precise estimation of the locations of (stellar) images in the CCD pixel stream. For this we used detailed Monte-Carlo simulations of the charge build-up and transfer in a Gaia-like CCD containing localized electron 'traps' caused by the particle radiation. These traps were found to significantly increase the Charge Transfer Inefficiency (CTI) in the CCD, causing the location estimation to become both biased and less precise. Also (part of) the proposed approach to calibrate the effect was studied, i.e., using forward modelling of the CCD signal by means of a so-called Charge Distortion Model (CDM). To mitigate the damage at the hardware level, the CCDs to be used by Gaia are equipped with a Supplementary Buried Channel (SBC), a doping profile that runs along each pixel column to confine the volume of the charge cloud at low signal levels, thereby drastically reducing the number of traps encountered. In Chapter 3 we showed that the CTI effects for the faintest magnitudes are significantly reduced thanks to the presence of the SBC. We also discussed the use of periodic Charge Injections (CI) in Gaia. This regular injection of artificial charges fills a large fraction of the traps, thereby reducing the CTI. It also eases the calibration of the remaining CTI by resetting the illumination history of the pixels at each CI. By synthesizing these various aspects, Chapter 3 resulted in a realistic assessment of the radiation effects at the image location level under a variety of conditions.

In this chapter we use the statistical results from Chapter 3 to model radiation-damaged observations, and study their effect on a simulated Gaia-like astrometric solution for one million stars. Although much smaller than foreseen for the actual astrometric solution (about 100 million stars), this number is large enough to model quite realistically the diffusion of the errors as function of position on the sky as well as in magnitude. Previous attempts to assess the impact of radiation damage on the astrometric performance of Gaia have mainly focussed on the increased photon noise due to the charge loss in the stellar images (which is another manifestation of the CTI). What sets this study apart from previous studies is that the image location bias as well as the in-

creased random errors are rigorously propagated through a realistic astrometric solution, including the spacecraft attitude determination.
The main goals of this chapter are: (i) to characterize and quantify the impact of radiation damage effects on the estimated astrometric parameters; (ii) to compare this with current scientific requirements for Gaia; (iii) to investigate whether the solution residuals can be used to improve the calibration of the CTI effects; and (iv) to estimate the typical astrometric errors due to (faint) disturbing stars preceding stellar measurements.

4.2 Methodology

To assess the impact of the radiation damage on the astrometric parameters, several elements of the mission need to be modelled in some detail: how Gaia operates by scanning the sky, including how the stars are distributed in position and magnitude (Section 4.2.1), how the observations are modified by the CTI, based on the results from Chapter 3 (Section 4.2.2), and how these observations are used to estimate the satellite attitude and the astrometric parameters (Section 4.2.3). Based on this information we then proceed to interpret the simulation results in Section 4.3 and discuss the implications in Section 4.4.

4.2.1 Generating Gaia-like observations

4.2.1.1 How Gaia observes

The main instrument of Gaia is an optical telescope with two fields of view imaged on the same focal plane and CCD mosaic. On the celestial sphere, the two fields are separated by 106.°5 (the basic angle). The spacecraft will orbit around the second Lagrangian point (L2) of the sun–earth system during a nominal science mission duration of five years. It will continuously spin around its own axis with a period of six hours, allowing its two fields of view to scan the sky approximately along great circles. The spin axis of the satellite is constantly pointed 45° away from the sun and has a precession-like motion around the solar direction with a period of 63 days. The combined motion due to the spin, precession, and the annual (apparent) motion of the sun is called the Nominal Scanning Law (NSL), and the spacecraft is commanded to follow this NSL to within 1 arcmin in all three axes. The precession of the spin axis changes the orientation of the consecutive great-circle scans, allowing the whole sky to be covered in about six months. For a nominal mission lifetime of five years the number of field-of-view transits as function of position on the sky is shown in Fig. 4.1. Accounting for mission dead-time, on average each position will transit the combined field of view 72 times.
The measurements are recorded in a single focal plane consisting of 106 CCDs (Fig. 4.2). Due to the satellite spinning motion, the star images will not remain stationary on the CCDs during an observation but will transit the focal plane in the along-scan (AL) direction. The orthogonal direction is called across-scan (AC). The charges accumulated in the pixels during the exposure are transferred in the AL direction over the CCD until they are read out at the edge of each CCD. When a star enters the focal plane it first passes over one of the sky mappers (SM), followed by the nine CCDs in the astrometric

field (AF1–9), and finally the photometric and spectroscopic CCDs. We will only consider the SM and AF CCDs because their observations are used to estimate the attitude (pointing) of the telescope and the astrometric parameters of the stars. Each field of view transit thus results in 10 individual CCD observations.
The SM and AF CCDs observe in a very broad wavelength range 300–1000 nm (defined by the telescope transmission and CCD quantum efficiency), and the corresponding magnitude scale is denoted G. The zero point is fixed by the convention that $G = V$ for an unreddened A0V star (Jordi et al. 2010; Perryman et al. 2001). Gaia will in principle observe all point-like objects with $5.7 \leq G \leq 20$, or about 1 billion stars (Drimmel et al. 2005; Robin et al. 2009).
As discussed in Chapter 3, for the CCD observation of each star, only a rectangular window of a few pixels around the star (typically 6 AL×12 AC pixels) is kept from the read-out stream; for the majority of the stars, these pixels are moreover binned in the AC direction, resulting in a one-dimensional set of electron counts that is sent to the ground. From these counts, an AL location estimation is performed and subsequently converted into a precise 'observation time' t (Section 4.2.2). The complete set of observation times for all the stars, CCDs and field-of-view transits constitutes the main input for the astrometric solution (Section 4.2.3).
Additionally to the along-scan measurements, represented by the observation times t, there is a measurement of the across-scan angle ζ for every SM observation, and for stars brighter than $G = 13$ also for each AF observation. The AC measurements are needed for the full three-axis attitude determination, although the requirements in the AC direction are much relaxed compared to the AL measurements. Typically, the uncertainty of the AC measurement is a factor 5–10 times larger than in the AL direction. Because of this, and the fact that most stars will have just one AC measurement for every ten AL measurements, the across-scan measurements hardly contribute at all to the final astrometric parameter estimates, except indirectly via the AC attitude. Therefore we will concentrate in this chapter on the AL observation times when we speak about observations.

4.2.1.2 Transit characteristics of the scanning law

An important property of the nominal scanning law used by Gaia (described in Section 4.2.1.1) is that the range of angles under which a star is scanned by the fields of view depends on its position on the sky. We already noted that the observations can be regarded as virtually only along-scan measurements, so the relevant direction to consider is the position angle of the direction in which the field is moving at the time of transit. For brevity we refer to the position angle of the scan as the 'scan angle'. The distribution of scan angles affects the determination of the different components of the astrometric signal, and is largely determined by the ecliptic latitude, β_{e}, of the star. For $|\beta_{\mathrm{e}}| > 45^\circ$ the scan angles have a rather uniform distribution over 360°, while for $|\beta_{\mathrm{e}}| < 45^\circ$ they are getting more aligned with the (ecliptic) north–south direction. The most extreme case is on the ecliptic ($\beta = 0$), where the scan angles are all within 45° of the north–south direction. Also for $|\beta_{\mathrm{e}}| < 45^\circ$ the number of transits varies a lot depending on the exact position on the sky, as can be seen from Fig. 4.1. The points with the smallest number of transits can be found in this region as well. Another im-

portant property of the ecliptic region is that the coverage in time is a lot more irregular than further away from the ecliptic. For example, the (almost vertical) arc-like structures around ecliptic plane with the higher number of transits typically receive most of their observations in a period of 32 days of consecutive observations (half the spin axis precession period), when the scan angles are all in a one particular direction along the arc. Although these stars will also receive scans in the other (opposite) directions at other times, this will not erase the single-directional signature of the bulk of the observations. Therefore some of the arcs will have scans predominantly in the northern direction, and others in the southern direction. This loss of symmetry in the scan angles will be important when interpreting the sky distribution of the astrometric errors in Section 4.3.

4.2.1.3 Star distribution model

The astrometric solution will only use a subset of well-behaved (apparently single) 'primary' stars (see Section 4.2.3). A rough estimate of the expected number of primary stars is $\sim 10^8$, i.e., 10% of all stars that Gaia will observe. Although it is difficult to predict the precise number and distribution of these primary stars beforehand, it is expected that the brighter stars ($G \lesssim 15$), despite their relatively small number, will dominate the astrometric solution due to their high statistical weights (small uncertainties).

Given available computer resources our simulations cannot conveniently handle more than about 1 million stars. On the other hand, the astrometric solution requires a certain minimum density of primary stars to work, and one million stars, if uniformly distributed, is just above this minimum density. Since the primary stars should also have a reasonable coverage in magnitude, the possible choices of primary star distributions is rather restricted, as discussed below.

As a starting point for the *primary* star distribution model we introduce the star density function of *all* stars on the sky: $A(G,p)$, in stars mag^{-1} deg^{-2} as function of G and position p. The adopted model has a spatial resolution of about 1 deg^2 and is based on GSC-II counts for the brighter stars (Drimmel et al. 2005; Drimmel et al. 2005) and the Besançon Galaxy model (Robin et al. 2003) for the fainter. The model provides the star density in bins of 0.5 mag in the range $4 \leq G \leq 21$ for 32,768 HTM pixels (O'Mullane et al. 2001) on the sky. Drawing a sample of 10^6 stars from this model (restricted to the magnitude range observed by Gaia) would however give far too few stars at high galactic latitudes for the attitude determination to work properly (this requires several stars in each field of view at any time). Thus we have chosen to adopt for the primary stars a model distribution with uniform spatial density. At each position p, half of the primary stars are drawn from the magnitude distribution of $A(G,p)$, but normalized to the $12.45 \leq G \leq 15$ range, while the other half are uniformly distributed both in position and in the magnitude range $12.45 \leq G \leq 20$. This split of the magnitude distribution combines a realistic bright-star contribution with the possibility to study the impact of CTI effects on the full range of magnitudes.

The bright magnitude limit of $G = 12.45$ is due to the availability of CTI data from Chapter 3, in which brighter magnitudes were avoided due to a complicated (and not yet fixed) bright-star gating scheme, together with the relatively small number of stars

Table 4.1 — Overview of our primary star distribution model ('selected') versus the total number of stars according to the Galaxy model described in the text. Column 3 shows the effect of our two-distribution selection: an increasing number of stars up to $G = 15$ plus a flat distribution over the whole magnitude range. The last two columns show that the density of the selected stars near the galactic poles is only a fraction of the real sky densities even at the bright end.

G [mag]	All sky [stars]		$\lvert b\rvert > 80$ deg [stars deg^{-2}]	
	Total	Selected	Total	Selected
12.45–13.30	6 471 080	142 773	30.9	3.5
13.30–14.15	12 315 019	210 312	48.7	5.1
14.15–15.00	22 667 442	315 436	71.0	7.6
15.00–15.88	42 148 217	57 971	102.6	1.4
15.88–16.75	72 526 612	58 287	140.6	1.4
16.75–17.63	123 786 875	58 097	192.6	1.4
17.63–18.50	202 585 466	58 188	258.1	1.4
18.50–19.25	277 964 109	49 083	296.9	1.2
19.25–20.00	425 876 271	49 853	387.8	1.2
12.45–20.00	1 186 341 090	1 000 000	1529.2	24.2

in the range $5.7 \leq G \leq 12.45$ (∼0.5% of the expected 10^9 sources). Furthermore, as will be detailed in Section 4.2.2, the data from Chapter 3 is available for nine magnitudes between $G = 13.3$ and $G = 20$ with a typical separation of 0.85 mag. Since it is preferable not to interpolate these data we bin the stars in our model into nine magnitude bins, and assign to each star the magnitude of the fainter limit of the bin for the purpose of calculating the CTI effects. The resulting magnitude distribution is given in Table 4.1.
With this model of the primary stars distribution it is not obvious that the real sky number density of bright stars is not exceeded in the galactic pole regions. That this is not the case is shown in the last columns of Table 4.1, giving the mean densities for galactic latitude $|b| > 80$ deg. While the total number of primary stars is less than 0.1% of the total numbers predicted from the Galaxy model, about 1% of the brighter ($G \leq 15$) stars are selected, and more than 10% of the bright stars in the polar regions.

4.2.1.4 Computing synthetic Gaia observations

For this study we simulate a fully synthetic set of observations using a 'scanner' program contained in the AGISLab software (see Section 4.2.3.3). Based on detailed input models (further described in Section 4.2.3) and a specific scientific mission time window, the scanner produces the true (i.e., noise-free) AL observation times and AC angles, t^{true} and ζ^{true}, for a set of stars with given astrometric parameters.
To these noise-free data we can apply any kind of perturbations. The most basic perturbation model is to add Gaussian noise. In the absence of radiation damage the location uncertainty of a star image with a particular magnitude is well described by a normal distribution with zero mean and a standard deviation that depends on G and the type of observation (AL/AC, SM/AF). In this chapter we model the radiation damage-induced bias and increased location uncertainty by a normal distribution with

non-zero mean and a widened standard deviation, as detailed in the next section.

4.2.2 Charge transfer inefficiency model

Radiation damage will drastically increase the Charge Transfer Inefficiency (CTI) in the Gaia CCDs. The distortion and overall charge loss caused by CTI in the one-dimensional stellar images collected by the astrometric instrument on-board Gaia introduce a bias in the image location estimation procedure and decrease its ultimate precision. In order to propagate these errors to the astrometric solution, we need to disturb the times of observation for each of the stars that constitute the input of the solution. Due to the very large number of observations considered (about 800 per primary star, or 8×10^8 observations in total) it would be practically infeasible to simulate each CCD observation through a detailed Monte-Carlo model of the CTI effects such as the one used in Chapter 3. We need a simple and fast model, capable of predicting in a realistic way the CTI induced bias on the observation times and the decrease in precision as functions of star brightness, accumulated radiation dose (i.e., time in the mission), illumination history, and level of mitigation. In the following we detail such a model.

4.2.2.1 Model formulation and principle

The satellite spin rate, expressed in AL pixels per second, is effectively kept at its nominal value to better than 1 part in 10^4 through the active attitude control of the spacecraft. As a consequence the image location biases δ_κ and standard errors σ_κ derived in Chapter 3 can be directly expressed in angular units by using the nominal AL pixel size as projected on the sky (58.93 mas). From here on, they are for simplicity denoted δ, σ (stellar declination will also be denoted by δ, but the distinction will always be clear from the text).
In Chapter 3 we characterized the variation of δ and σ as a function of the stellar magnitude G,

$$\delta\left(G|\rho^\star, w, r, \beta\right) \quad \text{and} \quad \sigma\left(G|\rho^\star, w, r, \beta\right) \tag{4.1}$$

for different fixed values of the image width w, level of background β, readout noise r, and density of active traps $\rho^\star$.
The active trap density $\rho^\star$ corresponds to the number of empty traps per pixel immediately before the stellar observation of interest. It depends on: (i) the total density of traps in the CCD, ρ, set by the accumulated dose of radiation, and (ii) the trap occupancy level θ, or the fraction of filled traps, set by the CCD illumination history:

$$\rho^\star = (1 - \theta)\,\rho \tag{4.2}$$

Both theory and experiments show that the CTI effects increase monotonically with $\rho^\star$. In order to temporarily fill a large fraction of the traps in the CCD image area (and thus mitigate the CTI effects), artificial charges will be periodically injected in the first CCD pixel row and transferred throughout the CCD (see Chapter 3). It is currently planned to perform a Charge Injection (CI) every second during the mission, and we

therefore adopt a nominal value for the CI period, $T_{\rm CI} = 1$ s. This periodic CI will be phased in such a way that each of the nine AF observations in a given star transit will see a different time since CI. This de-phasing of the individual observations is made on purpose to allow the CTI effects to be adequately mapped as function of $t_{\rm CI}$, the time elapsed since the preceding CI. Due to the high frequency of charge injections, they will dominate the CCD illumination history for the vast majority of observations. Once filled by a CI, the traps will release their electrons following an exponential decay process with a time constant that depends on the nature of the trap (i.e., trap species). This means that for a given CI period θ reaches a minimum, and $\rho^\star$ a maximum, for a CI delay equal to the CI period:

$$\rho^\star(t_{\rm CI} = T_{\rm CI}) = \rho^\star_{\rm max} \tag{4.3}$$

Consequently, δ and σ also reach a local maximum for $t_{\rm CI} = T_{\rm CI}$. The global maximum values of CTI-induced bias and standard errors, $\delta_{\rm max}$ and $\sigma_{\rm max}$, will thus be reached for $t_{\rm CI} = {\rm T_{CI}}$ and for the Gaia end-of-life accumulated radiation dose, after 5.5 years of operation (including half a year of pre-science phase), when the total trap density reaches its maximum.

In Chapter 3 (cf. Section 5.7), we showed that simulations performed with $\rho^\star = 1$ trap pixel^{-1} reproduce the amplitude of location bias measured using experimental test data taken 1 s after a CI. We recall that the test data used for this model-to-data comparison was acquired using a Gaia irradiated CCD with a radiation dose of 4×10^9 protons cm^{-2} (10 MeV equivalent). This dose corresponds to the upper limit of the predicted accumulated radiation dose for the Gaia nominal life-time. For these reasons the values for δ and σ found in Chapter 3 correspond to $\delta_{\rm max}$ and $\sigma_{\rm max}$ for the Gaia life-time and nominal CI period, $T_{\rm CI} = 1$ s.

In Chapter 3 we also confirmed that the CTI effects on the image location are proportional to $\rho^\star$: for a lower accumulated dose of radiation and a smaller number of active traps, only a corresponding fraction of this maximum bias and increase in the standard errors should be applied. Moreover, it was shown that in the absence of radiation damage, the Gaia image location estimation procedure is unbiased. Hence, we formulate the following simple model for the image location bias of a particular CCD observation:

$$\delta = \delta_{\rm max}(G)\, f_\rho(t_{\rm m})\, f_{\rm IH}(t_{\rm CI}) \tag{4.4}$$

It depends on $t_{\rm m}$, the time into the science mission, and $t_{\rm CI}$, the time since the previous CI, through the functions f_ρ and $f_{\rm IH}$ discussed in Sections 4.2.2.3 and 4.2.2.4. The factor $f_{\rm IH}$, referred to as the 'illumination history' factor, in principle depends on the entire illumination history of the pixel column prior to the observation (cf. Section 4.2.2.5) and corresponds to the fraction of empty traps. For most observations it is almost completely set by the previous CI; thus

$$f_{\rm IH}(t_{\rm CI}) = 1 - \theta(t_{\rm CI}) = \rho^\star(t_{\rm CI})/\rho \tag{4.5}$$

The factor f_ρ, referred to as 'radiation dose fraction', is the fractional total density of traps at a particular time in the science mission $t_{\rm m}$, i.e., the number of accumulated

traps since launch divided by the total number of accumulated traps during the nominal life-time of the mission:

$$f_\rho(t_\mathrm{m}) = \rho(t_\mathrm{m})/\rho_\mathrm{max} \tag{4.6}$$

The associated image location standard errors can be modelled in a similar fashion:

$$\sigma = \sigma_0(G) + [\sigma_\mathrm{max}(G) - \sigma_0(G)]\, f_\rho(t_\mathrm{m})\, f_\mathrm{IH}(t_\mathrm{CI}) \tag{4.7}$$

where σ_0, the image location uncertainty in the absence of radiation damage, is set to the values found in Chapter 3.

4.2.2.2 Maximum location bias and standard errors

δ_max and σ_max were characterized in Chapter 3 as functions of magnitude for different mitigation schemes and for multiple values of w, r, β, and $\rho^\star_\mathrm{max}$. For the rest of this study we make use of only two cases, referred to as the full-damage case and the mitigated case, in addition to the CTI-free case which is used as a reference.
In the full-damage case no CTI mitigation procedure is assumed. The active trap density $\rho^\star_\mathrm{max} = 1$ trap pixel^{-1} corresponds to the maximum density of active traps for a the nominal CI period ($T_\mathrm{CI} = 1$ s). The selected image width corresponds to the width of the 'typical' reference image (cf. Chapter 3). The readout noise value is the measured value for the Gaia CCD ($r = 4.35$ e$^-$), and the background level corresponds to the average sky surface brightness ($\beta = 0.447$ e$^-$ pixel^{-1} s^{-1}).
In the mitigated case, we assume that the CTI calibration procedure presented in Chapter 3 is applied, viz., a forward modelling approach using a Charge Distortion Model (CDM) to modify the modelled CCD signal as part of the image parameter estimation procedure. We select the results for δ_max and σ_max obtained for the current best CDM candidate (Short et al. 2010; Prod'homme et al. 2010) and an associated optimal calibration of its parameters. The values for w, r, β, and $\rho^\star_\mathrm{max}$ are the same as in the full-damage case.
Figure 4.3 shows δ_max (left) and σ_max (right) as a functions of G in the two cases, as well as for the CTI-free case. As can be noticed, in particular for the mitigated case, the bias is a not a smooth function of magnitude. As a consequence we refrain from interpolating these results for intermediate values of G and use, in the generation of the astrometric solution input, only the values sampled in Chapter 3 (as in our star distribution model described in Section 4.2.1.3).

4.2.2.3 Radiation dose fraction

The radiation dose accumulated by the CCDs will evolve as a function of t_m, the time into the science mission. We assume a pre-science phase of half a year after the launch, followed by $T_\mathrm{m} = 5$ yr of scientific operations. Thus, $f_\rho(-0.5\ \mathrm{yr}) = 0$ and $f_\rho(T_\mathrm{m} = 5\ \mathrm{yr}) = 1$.
At L2 solar protons expelled during solar flares dominate the radiation environment. The intensity of the particle radiation is thus directly related to the solar activity, for which we may use sunspot numbers as a proxy for modelling over months and years.

To avoid any arbitrary choice for the modelling of f_ρ we do not use predictions for the next solar cycle, but make use of the smoothed monthly sunspot counts by SIDC-team (2011) during the last solar cycle, shifted forward in time by the average duration of a solar cycle duration, 11.0435 (Julian) years. f_ρ is then the normalized integral of the sunspot counts between the observation time and the launch date, which is taken to be 2013.5. The monthly sunspot counts and the normalized cumulative counts are shown in Fig. 4.4.

4.2.2.4 Illumination history factor

The location bias for a particular observation depends on the state of the traps prior to the stellar observation. As the trap states depend on the CCD illumination history, the accurate calibration of the CTI effects could in the general case become exceedingly complicated. The use of periodic charge injection has two big advantages: it fills a large fraction of the traps, and it resets the illumination history to a (relatively) well-defined state immediately after the CI. If the CI period is not too long, one can to first order neglect the few stars that accidentally come between the stellar observation of interest and the preceding CI. In this case the illumination history and the trap occupancy level, θ, depends only on $t_{\rm CI}$, the elapsed time since the last CI. Based on the Shockley–Read–Hall formalism (Shockley & Read 1952; Hall 1952) for a particular trap species, one can derive the expected functional relationship:

$$\theta(t_{\rm CI}) = 1 - p_{\rm r}(t_{\rm CI}) = \exp(-t_{\rm CI}/\tau_{\rm r}) \tag{4.8}$$

with $p_{\rm r}(t)$ the probability for a filled trap to release its electron within time t (Chapter 2) and $\tau_{\rm r}$ the release time constant associated with the trap species considered. Substituting the computed θ in Eq. (4.5) the illumination history factor, $f_{\rm IH}$, becomes:

$$f_{\rm IH}(t_{\rm CI}) = 1 - \exp(-t_{\rm CI}/\tau_{\rm r}) \tag{4.9}$$

In Chapter 3 we considered a single trap species with a release time constant $\tau_{\rm r} \simeq 90$ ms. For consistency we use here the same value of $\tau_{\rm r}$.
In this computation of $f_{\rm IH}$ we assume (i) that the re-capture of released electrons can be neglected and (ii) that a CI fills all the traps that are likely to interact with the electrons of a stellar image. The latter assumptions is not correct if the CI level (i.e., the number of electrons per pixel in the CI block) is lower than the peak value of the stellar image. The currently retained value for the CI level during operation is 17 000 e^-. This means that our assumption hold for stars with $G > 14.5$, and that ultimately δ and σ may be slightly underestimated for stars brighter than this.

4.2.2.5 Illumination history factor with disturbing stars

In the following we also want to assess the impact of 'disturbing' stars that happen to fall between the last CI and the 'target' star of interest. As disturbing stars change the illumination history and trap occupancy level they can potentially introduce an extra source of noise unless one has detailed information about the positions and magnitudes of all potentially disturbing objects even beyond $G = 20$. To simulate this effect we make use of the same star density function $A(G, p)$ that was described in Section 4.2.1.3, but this time using all the star counts down to $G = 21$ to calculate the star

density in the combined field of view at any time. From this we can easily compute the mean time between the disturbing events per pixel column for different G magnitudes. Since the scan angle is nearly constant during one field of view transit, the consecutive observations made during the transit will have almost identical disturbing star histories. We thus simulate the disturbing effect of more than 2 billion stars on the 1 million star dataset.

Let us now briefly describe how the trap occupancy level θ is estimated in the presence of disturbing stars. Let us assume that a star of a given magnitude would manage to fill a certain fraction ϕ of the traps, if they were all empty before encountering the star. If the star is encountered at time t, let $\theta(t^-)$ and $\theta(t^+)$ denote the trap occupancy level immediately before and after the encounter. Simplistically, one can derive the following relation:

$$\theta(t^+) = \theta(t^-) + [1 - \theta(t^-)]\phi = \phi + (1 - \phi)\theta(t^-) \tag{4.10}$$

assuming that the disturbing star fills a fraction ϕ of the traps that were empty immediately before the encounter. For example, if $\phi = 0$ (a very faint star), we have $\theta(t^+) = \theta(t^-)$, i.e., the star makes no difference to the θ. Conversely, if $\phi = 1$ (a very bright star, similar to a CI), we have $\theta(t^+) = 1$ independent of the θ prior to the transit. Now consider for example the scenario where we have first a CI at $t = 0$, then a disturbing star at time t_1 with 'trap filling potential' ϕ_1, then another disturbing star at time t_2 with ϕ_2, and finally our target star observed at time $t_{\rm CI}$ (i.e., at this time after charge injection). We would then calculate $\theta(t_{\rm CI})$, the trap occupancy level immediately before the observation of the target star, as follows:

$$\begin{aligned}
\theta(t_1^-) &= \exp(-t_1/\tau_{\rm r}) \\
\theta(t_1^+) &= \phi_1 + (1 - \phi_1)\theta(t_1^-) \\
\theta(t_2^-) &= \theta(t_1^+)\exp[-(t_2 - t_1)/\tau_{\rm r}] \\
\theta(t_2^+) &= \phi_2 + (1 - \phi_2)\theta(t_2^-) \\
\theta(t_{\rm CI}) &= \theta(t_2^+)\exp[-(t_{\rm CI} - t_2)/\tau_{\rm r}]
\end{aligned} \tag{4.11}$$

Figure 4.5 shows an example of the evolution of $\theta(t)$ for different $t_{\rm CI}$, including and excluding disturbing stars. A consequence of the relative short release time constant ($\tau_r \simeq 90$ ms) with respect to the 1 s CI interval is that a faint disturbing star long before the target star will hardly influence the trap occupancy level at the target star. In Section 4.3.4 we will therefore also consider the case $\tau_r \simeq 900$ ms.

It only remains to be determined how ϕ depends on the magnitude of the star. We use

$$\phi_i = \min(1, F_i/C) \tag{4.12}$$

where F_i is the maximum flux density (in e^- pixel^{-1}) in the image of disturbing star i, and C is the CI level, which is set to 17,000 e^-. The whole disturbing star model operates as follows:

1. For each field of view transit (i.e., 9 AF observations) of each target star we compute the combined field-of-view sky density in all 33 magnitude bins between $G = 4.5$ and 21.

2. For these densities a random 'disturbing scene' is created for a time interval of length T_{CI} before the observation of the target star. This disturbing scene consists of a list of disturbing stars for which two quantities are stored: the trap filling potential ϕ and the time difference to the target star.
3. For each observation in this field of view transit determine the time since charge injection t_{CI}, decide which disturbing stars were encountered after the CI, and compute θ by means of Eq. (4.11). This then replaces the expression in Eq. (4.8).

Although this model is probably too simplistic to be used for calibrating and correcting such effects in the processing of the real Gaia data, it should give a reasonable estimate of the level of these perturbations for the purpose of the present assessment.
As explained in Section 4.2.2.1 the periodic CI will be phased in such a way that each of the nine AF observations for the field-of-view transit of a target star will get a different t_{CI}, and possibly a different number of disturbing stars. For a cumulative sky density of 10^6 deg^{-2} there are on average only 2 disturbing stars in a full charge injection interval [1]. Moreover, since only $\sim 0.3\%$ of the sky contains regions with a cumulative density exceeding 10^6 deg^{-2} (for disturbing stars up to $G = 21$) the effect of the disturbing stars on the bulk of observations is expected to be very small, also because most of the disturbing stars are very faint (i.e. they have small ϕ and therefore induce only small θ deviation with respect to having no disturbing stars). For the highest-density regions of the sky it could however have an effect, which will be examined in Section 4.3.4.

4.2.2.6 Applied errors

We now have all ingredients to simulate the bias and increased uncertainty of the astrometric AL observation times due to the CTI. This is done by adding a Gaussian error to the noise-free observations generated by the scanner (Section 4.2.1.4):

$$t^{\mathrm{obs}} = t^{\mathrm{true}} + \mathcal{N}(\delta, \sigma^2) \tag{4.13}$$

with $\mathcal{N}(\delta, \sigma^2)$ denoting a normal random variable with mean δ (from Eq. 4.4) and standard deviation σ (from Eq. 4.7).
Figure 4.6 shows the mean applied AL errors (left) and their standard deviations (right) as functions of G. In Fig. 4.7 the statistics of the applied errors have been binned according to the time since charge injection (t_{CI}) and the time into the mission (t_{m}). Although the figures were computed from the actual random realizations of the applied errors, the statistical uncertainty in the displayed data is very small since each point or bin is based on several hundred thousands of observations. Note that the standard deviations in the two figures cannot be directly be compared because the standard deviation at each magnitude over all mission times and all times since CI (right plot of Fig. 4.6) necessarily includes the large bias fluctuation (left plot in Fig. 4.7), while this fluctuation is absent when computing the standard deviation in a bin for a particular time in mission and time since charge injection (right plot of Fig. 4.7).

1. This number takes into account the angular area covered in the interval T_{CI} =1 sec and the fact that the disturbing starts should lie on the same along-scan column as the target star, including a margin of 2 pixels in the AC direction matching the typical full width half maximum of the stars.

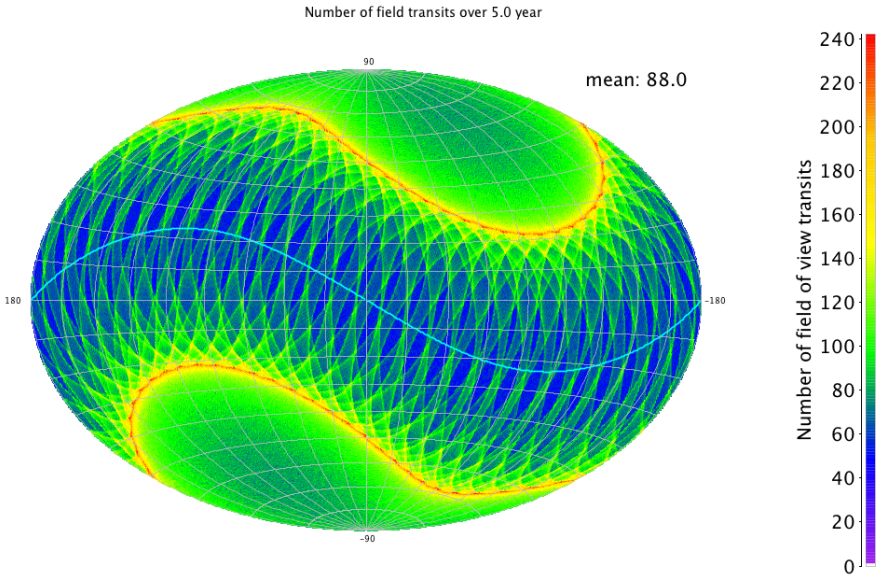


Figure 4.1 — Colour-coded map of the expected number of field of view transits experienced by sources at different celestial positions after a 5 year mission. The projection uses equatorial coordinates, with right ascension running from $-180°$ to $+180°$ right-to-left. The blue line is the ecliptic plane. The average number of field transits is 88 during 5 years, although the value that is normally used for performance evaluation is 72 (accounting for dead time). An over-abundance of transits occurs at $45°$ away from the ecliptic plane due to the difference between the $45°$ spin axis angle with respect to the sun and the $90°$ angle between spin axis and the fields of view.

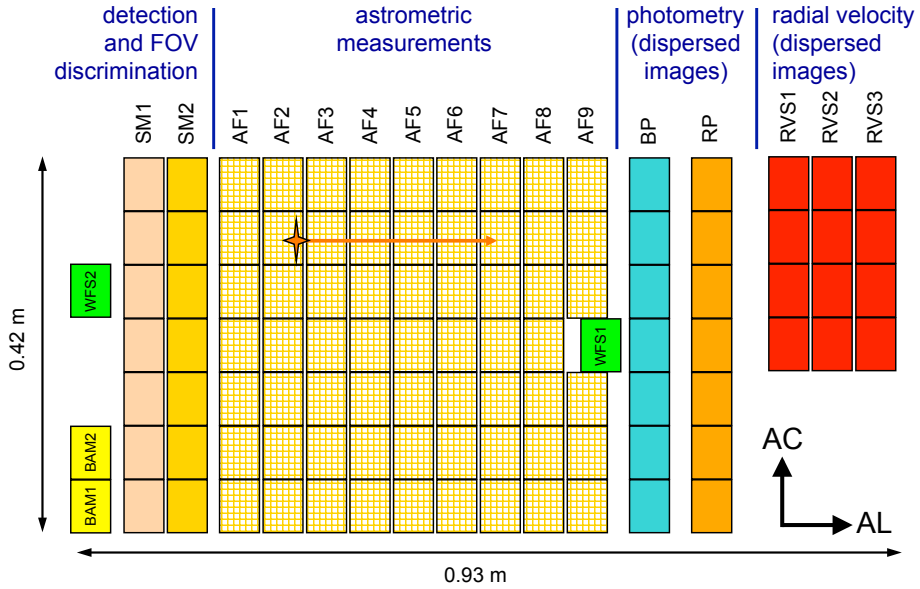


Figure 4.2 — Schematic layout of the CCDs in the focal plane of Gaia. Due to the satellite spin, a star enters the focal plane from the left in the along-scan (AL) direction. All stars brighter than $G = 20$ mag are detected by one of the sky mappers (SM1 or SM2, depending on the field of view) and then tracked over the subsequent CCDs dedicated to astrometry (AF1–9), photometry (BP and RP), and radial-velocity determination (RVS1–3). In addition there are special CCDs for interferometric Basic-Angle Monitoring (BAM), and for the initial mirror alignment using Wavefront Sensors (WFS).

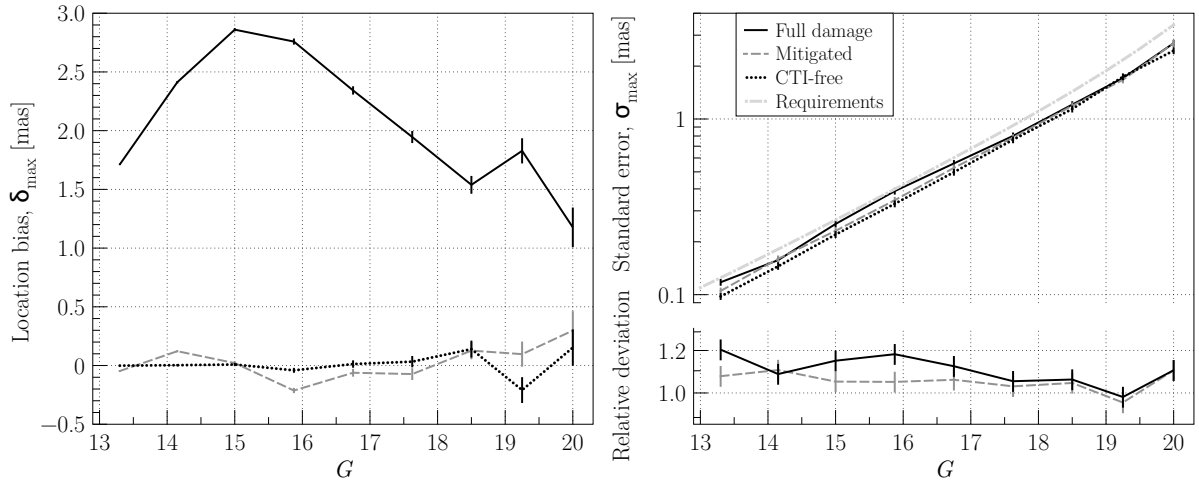


Figure 4.3 — $\delta_{\max}$ and $\sigma_{\max}$ in the 'full damage' (black continuous line) and 'mitigated' (grey dashed line) cases (see Section 4.2.2.2) as obtained in Chapter 3 for the Gaia operating conditions and for an active trap density of 1 trap pixel^{-1}. These values are used to construct our CTI model, shown in Fig. 4.6. **Left:** $\delta_{\max}$, maximum image location bias in milli-arcsecond (mas) as a function of magnitude (G-band) induced by CTI effects with (grey dashed line) and without (black continuous line) applying a CTI mitigation procedure. The CTI-free case is also shown (dotted line), but which has been set to strictly zero in our simulations. The error bars corresponds to the statistical uncertainties. **Right:** $\sigma_{\max}$, maximum image location standard errors in milli-arcsecond (mas) as a function of G induced by CTI effects with (grey dashed line) and without (black continuous line) applying a CTI mitigation procedure. The Gaia requirements (dash-dotted line) for a single CCD transit and the CTI-free case (dotted line) are also shown. The relative deviation from the later is computed (bottom panel) for the 'full damage' case (black continuous line) and the 'mitigated' case (grey dashed line).

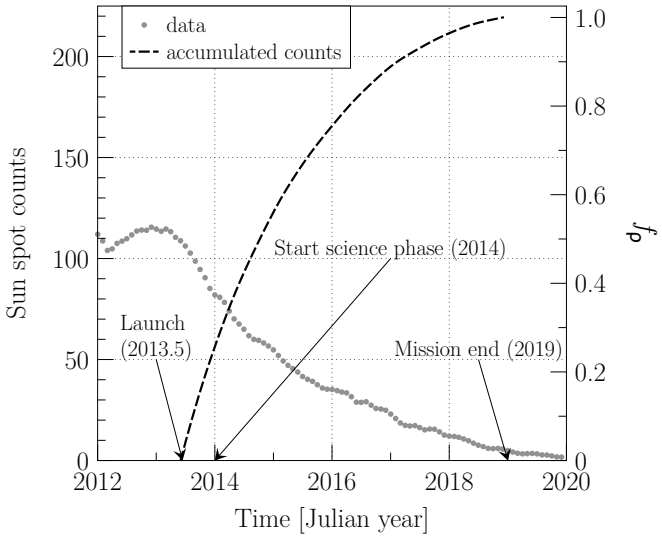


Figure 4.4 — The grey dots show the smoothed sunspot monthly counts (SIDC-team 2011) for the last solar cycle, shifted forward by one solar cycle. The black dashed line represents the cumulative sunspot counts from the assumed launch date. It is computed by numerical integration of a functional fit to the sunspot counts. f_ρ (right ordinate label) represents the cumulative counts, normalized to unity at the end of the nominal mission.

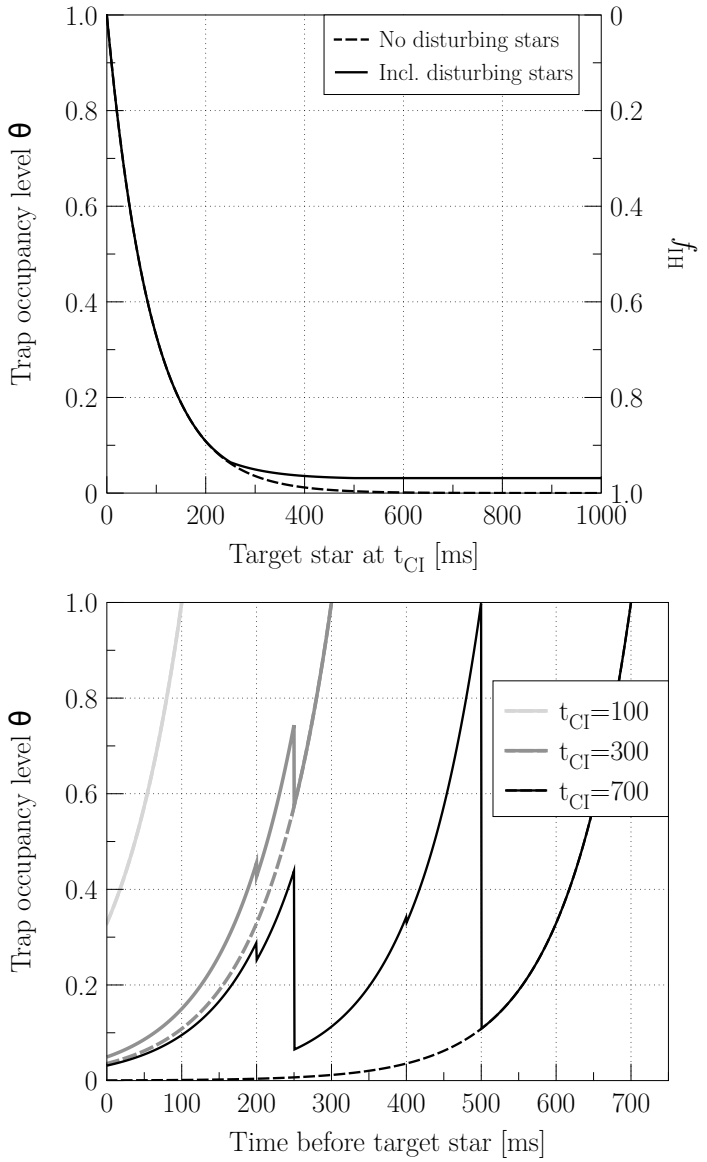


Figure 4.5 — Top: example of the trap occupancy level (θ) as a function of the time $t_{\rm CI}$ since the last change injection (CI), with (solid line) and without (dotted line) disturbing stars. The traps are all filled ($\theta = 1$ or $f_{\rm IH} = 0$) right after the CI (i.e., for short $t_{\rm CI}$), and they are all asymptotically empty ($\theta = 0$ or $f_{\rm IH} = 1$) for long CI delays. As expected from the Shockley–Read–Hall formalism, $\sim$67% of the traps have released their electron after $t_{\rm CI} = \tau_r = 90$ ms. The dotted curve deviates from the solid one due to the four disturbing stars (in this particular example) detailed below. **Bottom**: a detailed evolution of the trap occupancy level for three particular charge injection times before the target star: $t_{\rm CI} = 100, 300, 700$ ms. The trap occupancy levels at zero correspond to the levels displayed in the top diagram at these times. The disturbing stars are fixed with respect to the target star for a particular field of view transit, in this example they are located at: 200, 250, 400, and 500 ms before the target star, having $\phi = 0.05$, 0.4, 0.02, and 1.0 corresponding to $G = 17.6$, 15.4, 18.6, and $<$14.4 (for a CI of 17,000 e$^-$, see Eq. 4.12) , respectively. Because the 9 AF CCDs observations of a field of view transit have de-phased charge injections, the disturbing scene is sampled from 9 different times before the target star. For example, for $t_{\rm CI}$ =700 ms the trap occupancy is completely reset by the disturbing star at 500 ms (having ϕ=1.0), therefore the solid line in the top diagram is flat for $t_{\rm CI}$ $>$500 ms. From the bottom diagram it is also clear that the trap occupancy level as function of $t_{\rm CI}$ in the left diagram can only decrease.

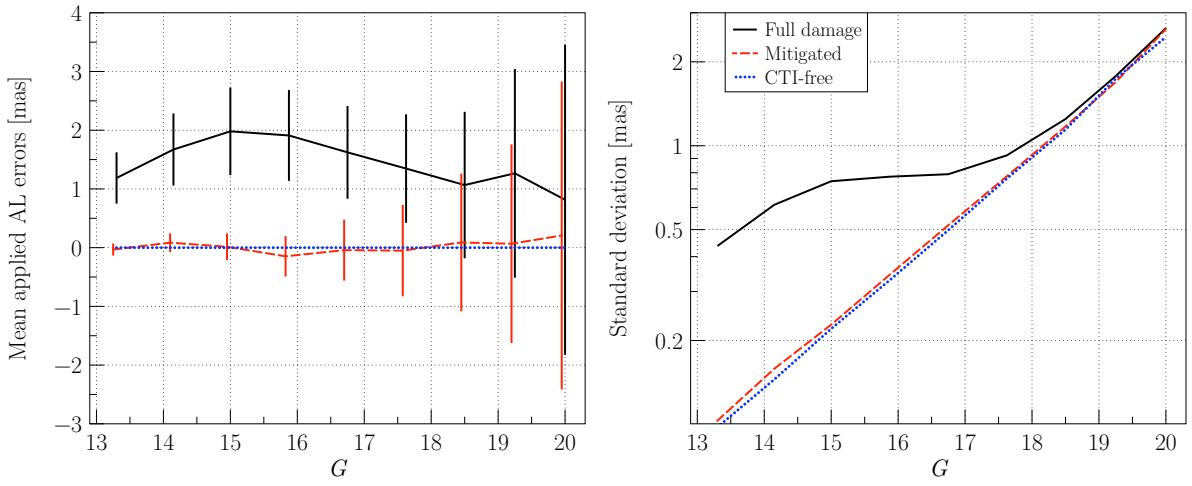


Figure 4.6 — Mean values (left) and standard deviations (right) of the applied along-scan errors for all the observations as functions of G. The curves are for the full damage (solid), mitigated (dashed), and CTI-free case (dotted). The vertical bars in the top diagram show the standard deviations from the right diagram in relation to the biases (for improved visibility the bars are omitted in the CTI-free case). The standard deviation necessarily includes the large bias fluctuation with mission time and time since CI (top plot in Fig. 4.7), and can therefore not be directly compared with the standard deviations shown in the bottom plot of Fig. 4.7.

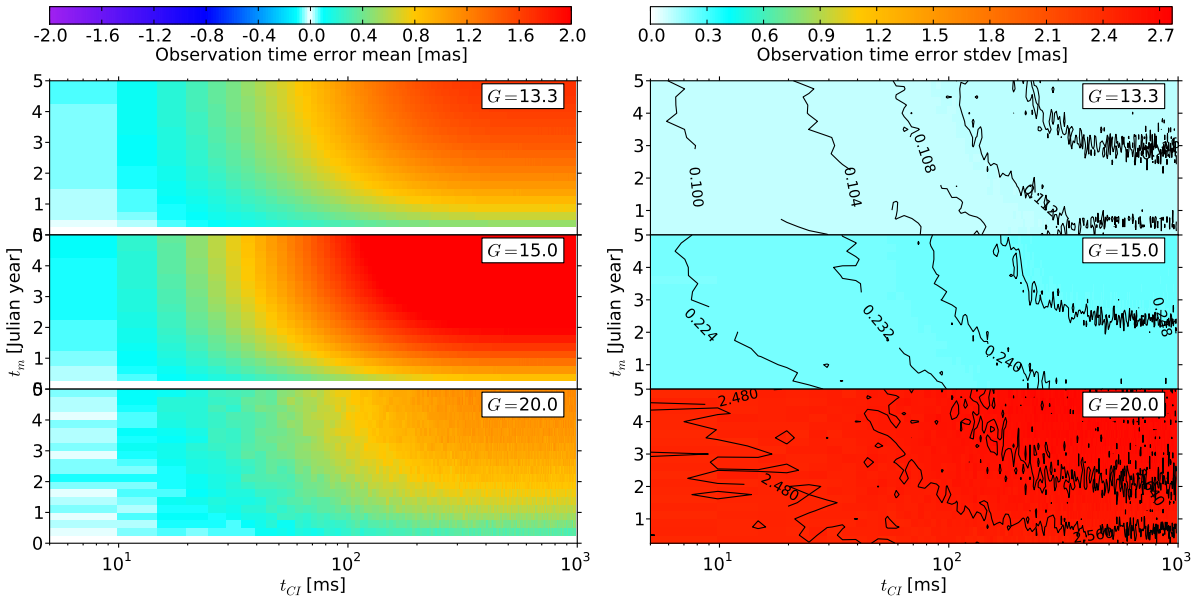


Figure 4.7 — Mean values (left) and standard deviations (right) of the applied along-scan error for the full damage case. These are similar to the statistics in Fig. 4.6, but subdivided according to time into the mission (t_{m}) and time since charge injection (t_{CI}), and only for the full-damage case and selected magnitudes. The applied bias is always positive in this case.

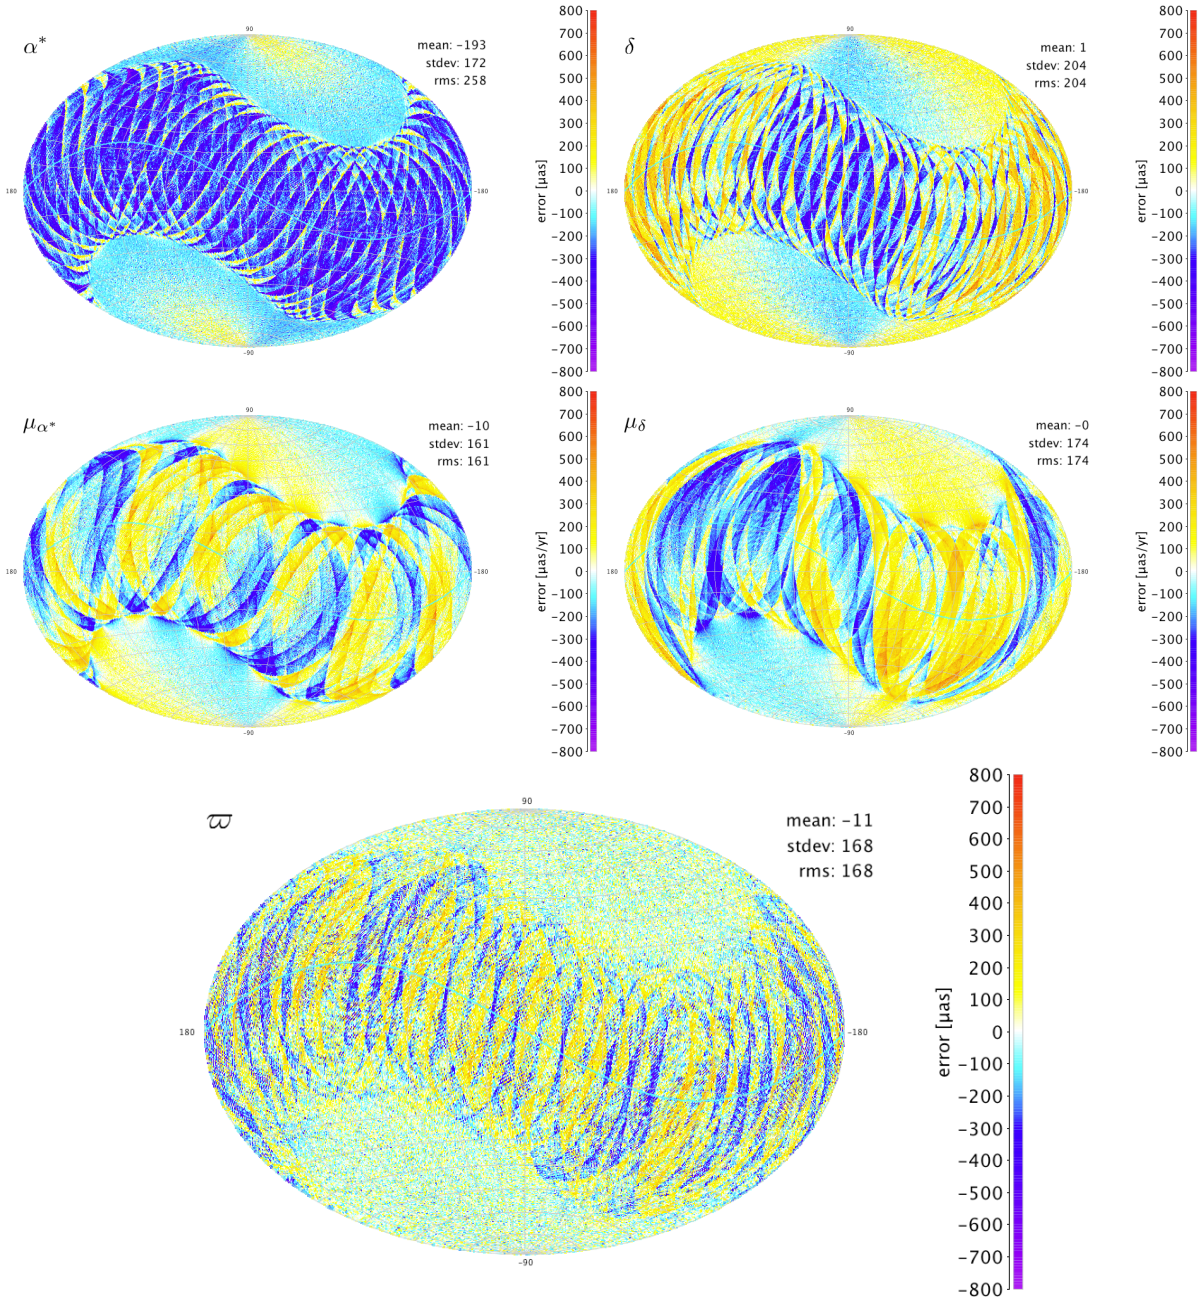


Figure 4.8 — Errors in position (α^*, δ), proper motion (μ_{α^*}, μ_δ), and parallax (ϖ) produced by a constant along-scan bias of 1000 μas (1 mas). See Section 4.3.1 for a detailed discussion of these plots.

The mean values of the applied errors in the left diagram of Fig. 4.6 can be understood, in relation to the maximum biases displayed in Fig. 4.3, by noting that the mean value of the product $f_\rho(t_\mathrm{m})f_\mathrm{IH}(t_\mathrm{CI})$ is $\simeq 0.69$, when averaged over the full mission length and all times since charge injection. The curves in Fig. 4.6 (left) are essentially the corresponding maximum biases multiplied by that factor. Similarly, since the RMS variation of $f_\rho(t_\mathrm{m})f_\mathrm{IH}(t_\mathrm{CI})$ around that mean value is $\simeq 0.25$, the standard deviations in the right diagram of Fig. 4.6 (and the error bars in the left diagram) contain a component that is about $0.25\delta_\mathrm{max}$. This explains the large standard deviation of the applied errors for the brighter stars in the full-damage case. In the mitigated case the bias is much smaller and the standard deviations are then dominated by the photon noise (as in the CTI-free case). In the right part of Fig. 4.7 the standard deviations are shown for fixed t_CI and t_m and therefore do not include the variation due to these factors.

4.2.3 Simulating the astrometric solution

4.2.3.1 The five astrometric parameters of a star

Similarly to what was done for the vast majority of stars in the Hipparcos Catalogue (ESA 1997; van Leeuwen 2007b), the geometric direction towards a single star, as seen from Gaia at time t, will be modelled by the unit vector

$$\mathbf{u}(t) = \langle \mathbf{r}_0 + \mathbf{p}_0\mu_{\alpha*}(t-t_0) + \mathbf{q}_0\mu_\delta(t-t_0) - \mathbf{b}(t)\varpi \rangle \tag{4.14}$$

in terms of five astrometric parameters α, δ, ϖ, $\mu_{\alpha*}$, and μ_δ. Here $\mathbf{r}_0$ is the barycentric direction towards the star at the agreed reference epoch t_0 (normally chosen to be halfway into the mission) and $\mathbf{p}_0$, $\mathbf{q}_0$ are unit vectors orthogonal to $\mathbf{r}_0$ in the directions of increasing α and δ, respectively. The so-called normal triad $[\mathbf{p}_0\ \mathbf{q}_0\ \mathbf{r}_0]$ is completely defined by the barycentric right ascension α and declination δ at epoch t_0. $\mu_{\alpha*}$ and μ_δ are the components of proper motion and ϖ is the parallax. $\mathbf{b}(t)$ is the barycentric position of Gaia, in astronomical units, and the angular brackets $\langle\,\rangle$ signify vector normalization. The geometric direction modelled by Eq. (4.14) is further modified by gravitational light deflection in the solar system and by stellar aberration (due to the velocity of Gaia in the barycentric frame), but as these effects are very well known and can be removed from the observations, they can be modelled in a simplified way (or even not at all), as long as the same model is used both for generating and analysing the observations.

The asterisk in $\mu_{\alpha*} \equiv \mu_\alpha \cos\delta$ signifies that the proper motion in right ascension is expressed as a true arclength on the sky (as opposed to $\mu_\alpha = \mathrm{d}\alpha/\mathrm{d}t$). Below, when discussing errors and uncertainties in right ascension, we similarly use an asterisk to denote the true angle; e.g., if $\Delta\alpha$ is the error in α, we may somewhat informally refer to $\Delta\alpha* \equiv \Delta\alpha\cos\delta$ as the error in $\alpha*$.

Considering that the parallax and the annual proper motion components, as well as the errors in $\alpha*$ and δ, are always very small angles, their effects on $\mathbf{u}$ can be considered independently by linear superposition. The proper motion components $\mu_{\alpha*}$ and μ_δ produce a uniform motion on the sky, typically by some mas yr^{-1}. Given that Gaia's orbit $\mathbf{b}(t)$ is nearly circular with a radius of about 1 AU, a non-zero parallax will cause

the star to move approximately in an apparent ellipse with a one-year period, semi-major axis $\simeq \varpi$, and semi-minor axis $\simeq \varpi \sin \beta_{\rm e}$, where $\beta_{\rm e}$ is the ecliptic latitude of the star. The combined effect of proper motion and parallax is a wiggly or spiral pattern on the sky. These simple geometrical considerations are helpful for interpreting the effects of the CTI bias on the astrometric errors (Section 4.3.1).

4.2.3.2 The Astrometric Global Iterative Solution

The baseline method that will be used to determine the astrometric parameters of stars observed by Gaia is the so-called Astrometric Global Iterative Solution (AGIS; Lindegren et al. 2011). This is an iterative least-squares estimation of the five astrometric parameters for a subset of $\sim 10^8$ well-behaved (apparently single) *primary stars*, with additional *nuisance parameters* for the instrument attitude, calibration and global parameters, resulting in a total number of $\sim 5 \times 10^8$ unknowns.
We can identify three main purposes of AGIS. The first is to estimate the nuisance parameters as well as possible using the observations of the well-behaved primary stars (which can be accurately modelled by Eq. 4.14). In this chapter we neglect the influence of the instrument calibration parameters and the global parameters on AGIS. The former are parameters that model, e.g., small changes in the CCD positions and orientations, and basic-angle variations, on time scales of days to years. Neglecting calibration and global parameters is motivated by their relatively small number ($\sim 10^6$) combined with the knowledge that each such parameter depends on a very large number of observations spread over many different primary stars over the whole celestial sphere. They are therefore not greatly affected by localized errors on the sky and can therefore be estimated relatively straightforward. (Depending on the choice of global parameters, they can sometimes have a profound effect on the astrometric solution, but a discussion of such effects is beyond the scope of this chapter.) In contrast, both the attitude and star parameters may have a very local influence across the sky, which could make their disentanglement from the star parameters much more difficult (cf. Sect. 1.4.6 in van Leeuwen 2007b). It is therefore essential that our simulation of the astrometric solution (Section 4.2.3.3) includes the simultaneous estimation of both star and attitude parameters (cf. Bombrun et al. 2010).
The second purpose of AGIS is to use the calibrated nuisance parameters to estimate the astrometric parameters of all stars. Since in this step the primary stars are treated no different than the rest of the stars, the results of the astrometric parameter estimates for the primary stars will be representative for all stars. In this chapter we will therefore only consider the primary stars.
The third purpose of AGIS is to tie the internally consistent astrometric solution to a global reference system. The solution provided by AGIS results in a reference frame (to which the positions and proper motions refer) which has in practice six degrees of freedom, corresponding to a solid-body rotation with fixed inertial spin (Lindegren et al. 2011). This is fixed using sources with a priori known astrometric parameters, including quasars, that define a kinematically non-rotating celestial frame.

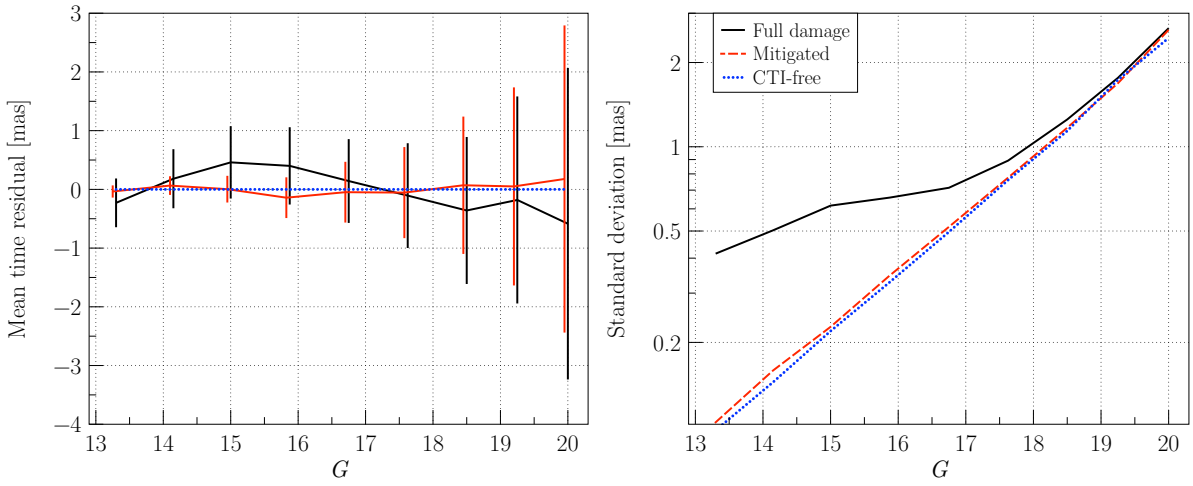


Figure 4.9 — Mean values (left) and standard deviations (right) of the along-scan residuals in the astrometric solution for all the observations as functions of G. The curves are for the full damage (solid), mitigated (dashed), and CTI-free case (dotted). These diagrams can be directly compared with the applied errors in Fig. 4.6. Note how the residuals are shifted down, with respect to the applied errors, so that the weighted mean bias is zero. The standard deviation necessarily includes the large residual fluctuation with mission time and time since CI (left plot in Fig. 4.10), and can therefore not be directly compared with the standard deviations shown in the right plot of Fig. 4.10.

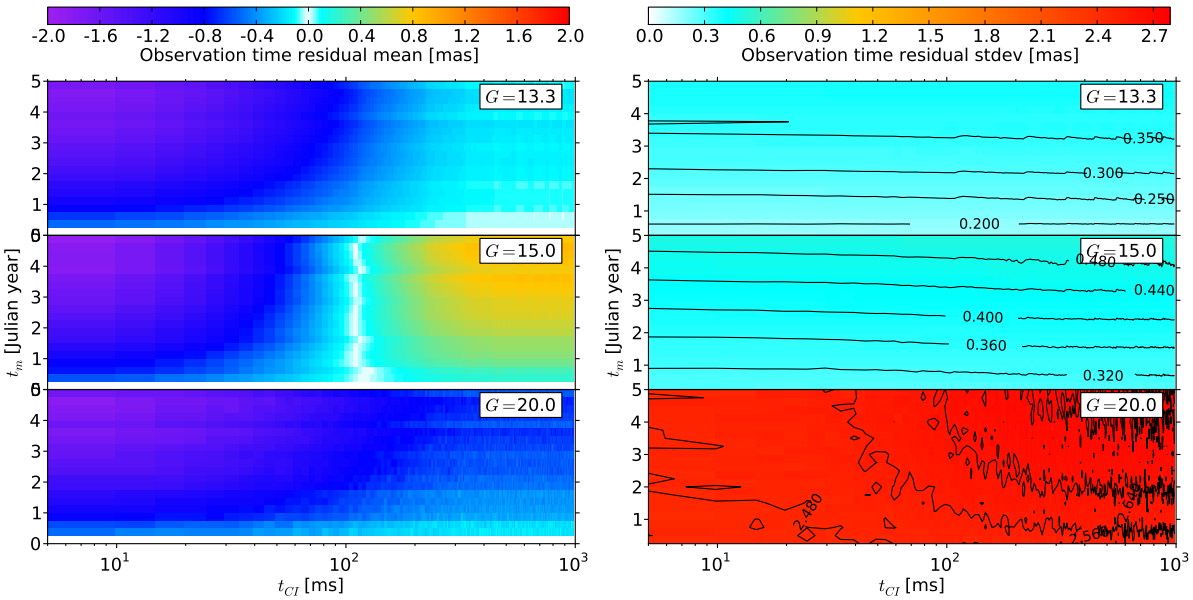


Figure 4.10 — Mean values (left) and standard deviations (right) of the along-scan residuals in the astrometric solution for the full damage case. These are similar to the statistics in Fig. 4.9, but subdivided according to time into the mission (t_{m}) and time since charge injection (t_{CI}), and only for the full-damage case and selected magnitudes. These diagrams can be directly compared with the applied errors in Fig. 4.7.

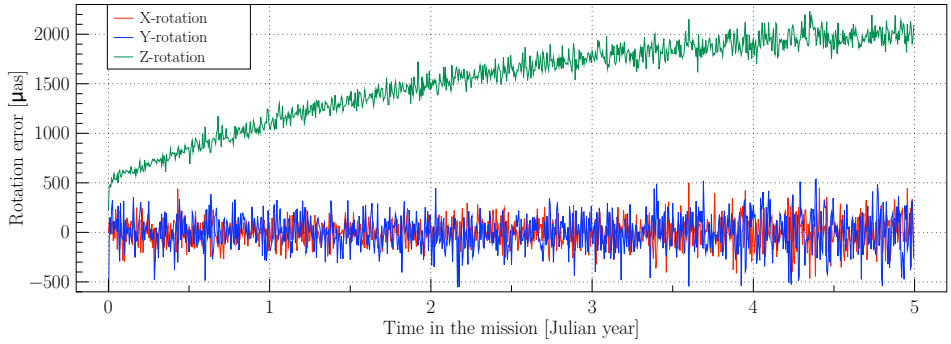


Figure 4.11 — Attitude errors around the satellite X, Y, Z axes as functions of time for the full damage case. The Z component is the AL rotation (spin) of the satellite. The applied errors are positive (a delay in time), which can be interpreted as a positive error in the direction of the rotation. The shape of the Z-component error mimics the shape of f_ρ in Fig. 4.4.

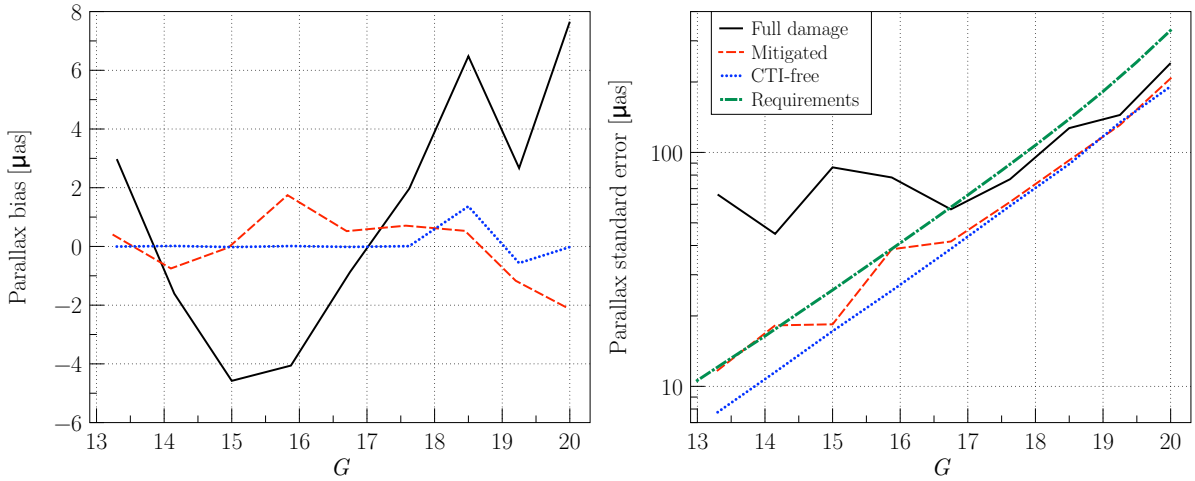


Figure 4.12 — Mean parallax error (left) and standard error (right) per magnitude. The Gaia requirements are plotted for comparison.

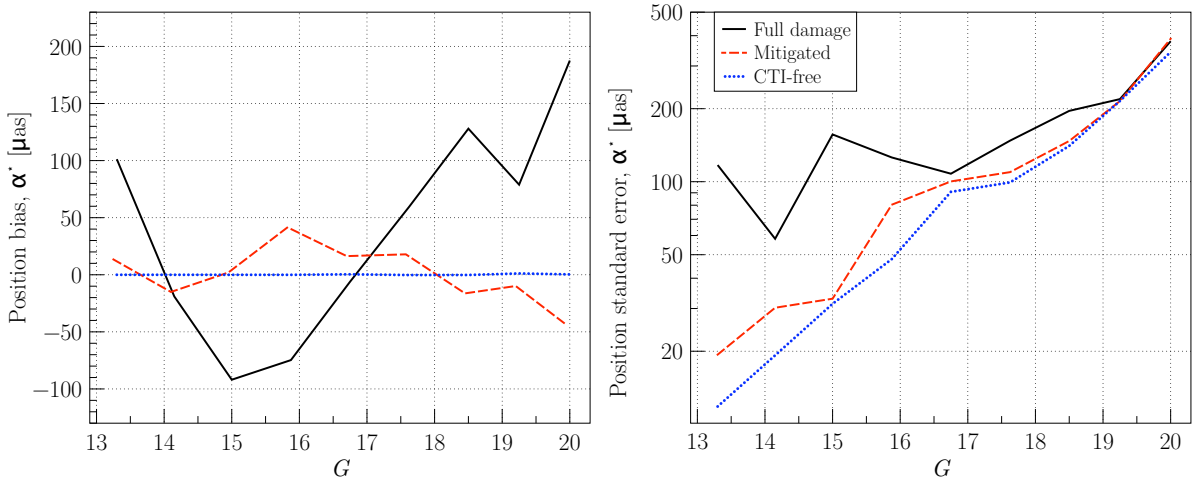


Figure 4.13 — Mean right ascension error (left) and standard error (right) per magnitude.

In principle AGIS solves the least-squares problem

$$\min_{\mathbf{s},\mathbf{a}} \sum_{l\in \mathrm{AL}} \left[\frac{t_l^{\mathrm{obs}} - t_l^{\mathrm{calc}}(\mathbf{s}_i,\mathbf{a})}{\sigma_l^{\mathrm{AL}}}\right]^2 + \sum_{l\in \mathrm{AC}} \left[\frac{\zeta_l^{\mathrm{obs}} - \zeta_l^{\mathrm{calc}}(\mathbf{s}_i,\mathbf{a})}{\sigma_l^{\mathrm{AC}}}\right]^2, \tag{4.15}$$

where t_l^{obs} is the observation time for an AL observation with index l, ζ_l^{obs} the corresponding AC measurement if it exists (cf. Section 4.2.1.1), and σ_l^{AL}, σ_l^{AC} their formal uncertainties from the image location estimator. The vectors $\mathbf{s}$ and $\mathbf{a}$ contain all the stars and attitude parameters, respectively, with $\mathbf{s}_i$ the subvector of $\mathbf{s}$ containing the five astrometric parameters of the primary star (i) to which the observation l refers. The predicted observations $(t_l^{\mathrm{calc}}, \zeta_l^{\mathrm{calc}})$ are calculated based on a composite model containing the stellar motion according to Eq. (4.14), the satellite attitude, and the geometry and orbit of the satellite, of which the first two are parametrized by the adjusted parameters. The sums in Eq. (4.15) extend over all the AL and AC observations of primary stars in the mission.

4.2.3.3 The simulation software AGISLab

To accurately characterize and interpret the results of the experiments in this chapter, it is essential to have complete control over the input, processing and output of a simulated AGIS solution. For this purpose we use the AGISLab software package (Holl et al. 2010). AGISLab allows the use of many different model implementations for the computation of t_l^{calc}, and ζ_l^{calc} in Eq. (4.15). As mentioned in Section 4.2.1.4, AGISLab is also used to generate noise-free observations using a 'scanner' function which can then be perturbed to mimic any type of noise and systematic errors. Having exact knowledge about the true model parameters allows us to follow the error propagation through AGIS and make a detailed characterization of the errors in the final parameter estimates.
We describe hereafter briefly the models and assumptions used in the simulations for this chapter. Unless otherwise noted, the same model was used for generating the observations and for solving the astrometric parameters, so that no additional modelling errors are introduced at the level of the astrometric solution.

Star and attitude model:
Our simulation of AGIS makes a simultaneous fit of all the star and attitude parameters. Given the size of the Gaia fields of view, it was found that a minimum of around 10^6 primary stars are needed for a robust solution, and this is also the number used, distributed as described in Section 4.2.1.3, and resulting in 5×10^6 astrometric parameters. The attitude parameters are cubic spline coefficients that describe the three-axis orientation of the satellite as a smooth function of time. The separation between spline knots is set to 120 s of time and there are 4 parameters per knot (corresponding to the four components of the attitude quaternion), resulting in 1.3×10^6 attitude parameters for the 5 yr mission.

Orbit and relativity model:
For the orbit of Gaia, a Keplerian model is assumed with a semi-major axis of 1.01 AU

(the real Lissajous orbit around L2 deviates from this by at most a few mAU). A simplified light bending model is used which only considers the gravitational deflection by the sun. Stellar aberration is rigorously computed using the velocity due to the Keplerian orbit. Since the observations are analysed using the same models, these simplifications have no impact on the conclusions.

Instrument geometry model:
The geometric model of the focal plane includes all SM and AF CCDs at their nominal positions (as in Fig. 4.2), plus an extra AF CCD at the position of the WFS CCD.

Mission time line:
We assume a launch date of 2013.5, a pre-science phase of 0.5 yr (relevant for the solar activity model in Section 4.2.2.3), and a science phase of 5 yr (thus 2014.0–2019.0).

Observation filtering:
Because in Chapter 3 we did not simulate the effect of radiation damage on the image location accuracy for the SM CCDs, and because the SM AL observations have a negligible contribution anyway (see Section 4.2.1.1), all SM AL observations were filtered out for the experiments of this chapter. All AC observations made by the SM and by the AF CCDs in the brightest magnitude bin ($12.45 \leq G \leq 13.3$) were kept, allowing a good three-axis attitude determination.

Observation perturbations:
The AF AL observations were perturbed using the CTI model in Section 4.2.2. Although it is expected that CTI effects will also be present in the AC measurements, they are not simulated. The AC observations are perturbed by Gaussian noise using a magnitude-dependent standard deviation based on a model by de Bruijne (2009a).

Number of iterations:
Because iterations are computationally expensive (on the present hardware system it typically takes one hour per iteration for 10^6 primary stars) we want to minimize the number of iterations needed to obtain a solution that is accurate enough for our purpose. Extensive experiments have shown that AGIS converges to a unique solution independent of the initial values for the star and attitude parameters (Bombrun et al. 2011), but the number of iterations required to reach the solution is of course larger if the initial values are far from the solution. Since the applied perturbations are small (at most a few mas), the solution will also be very close to the true parameters. We therefore minimize the number of iterations needed to reach the required level of accuracy by using the true parameters as the initial estimates. Even so, and using an efficient conjugate gradient algorithm, some 50 iterations are needed for a solution that is truly converged at the level of the numerical noise. However, after 30 iterations the updates are typically of the order of 0.001 μas or 0.001 μas yr^{-1} for the full-damage, mitigated and CTI-free case, which we consider 'good enough' (the astrometric biases and photon noise errors being typically 3 to 5 orders of magnitude larger); we therefore use 30 iterations for all our solutions, starting from the true star and attitude parameters.

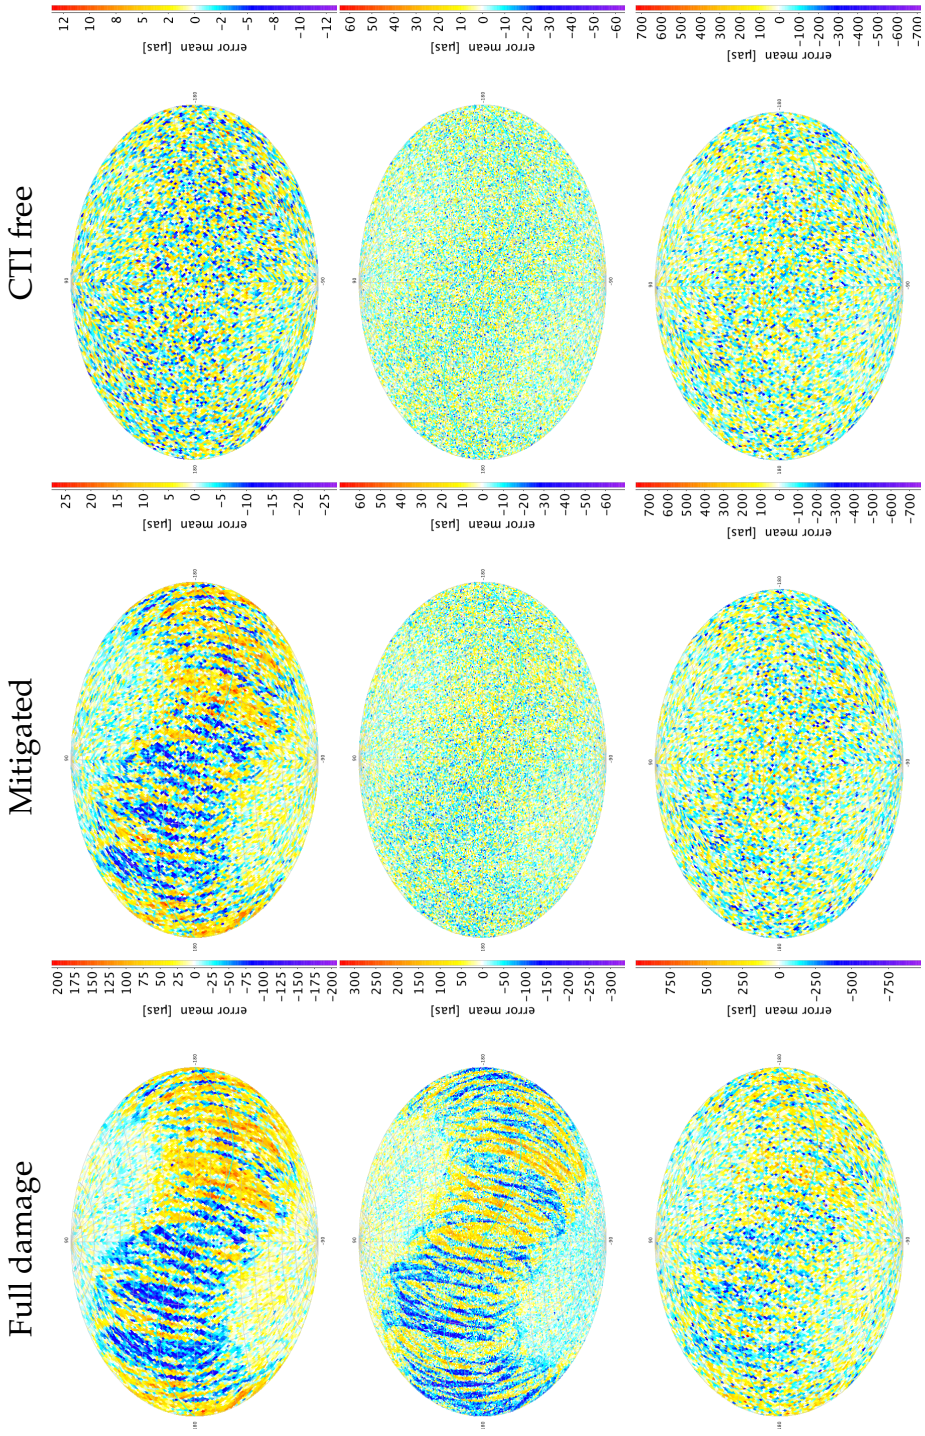


Figure 4.14 — Parallax error maps for the full damage (left), mitigated (middle), and CTI-free (right) case. Rows distinguish $G = 13.3$, 15, and 20. The CTI-free case demonstrates the error maps as it would look like based on pure photon noise. Depending on the strength of the observation bias with respect to the standard deviation, the bias pattern is seen more clearly towards brighter magnitudes. For the full damage $G = 20$ map the bias pattern is still clearly visible above the photon noise, while the mitigated map is nearly the same as the CTI-free case. Note that the mean bias as function of magnitude is different between the full damage and mitigated case (see Fig. 4.12). For example, the mitigated mean $G = 15$ residual is almost zero, resulting in no bias pattern here.

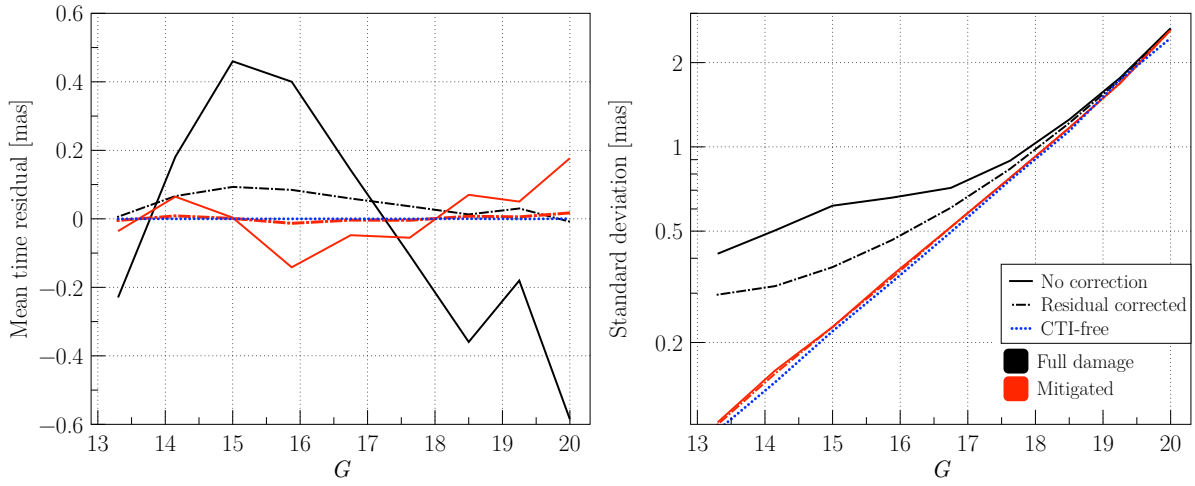


Figure 4.15 — Mean time residual (left) and standard deviation (right) per magnitude, with and without residual correction. The standard deviation necessarily includes the remaining residual fluction (left plot of Fig. 4.16), and can therefore not be directly compared with the standard deviations shown in the right plot of Fig. 4.16.

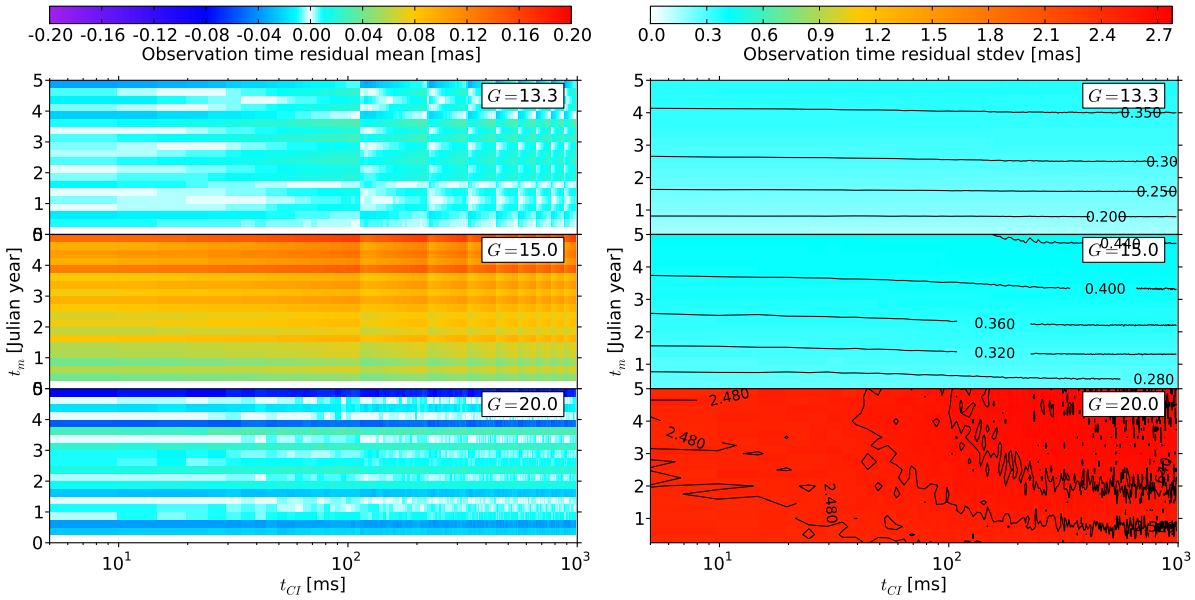


Figure 4.16 — AF AL residuals for the full damage case after applying a residual correction to the observations. Shown are the mean (left) and standard deviation (right) of the resulting residuals in 20 bins in mission time, t_m, and 200 bins in time since CI, t_{CI}. Compare this the left figure with the pre-correction mean residuals given in Fig. 4.10 (note that the color scale is 10× smaller!) to see how well this simple correction works.

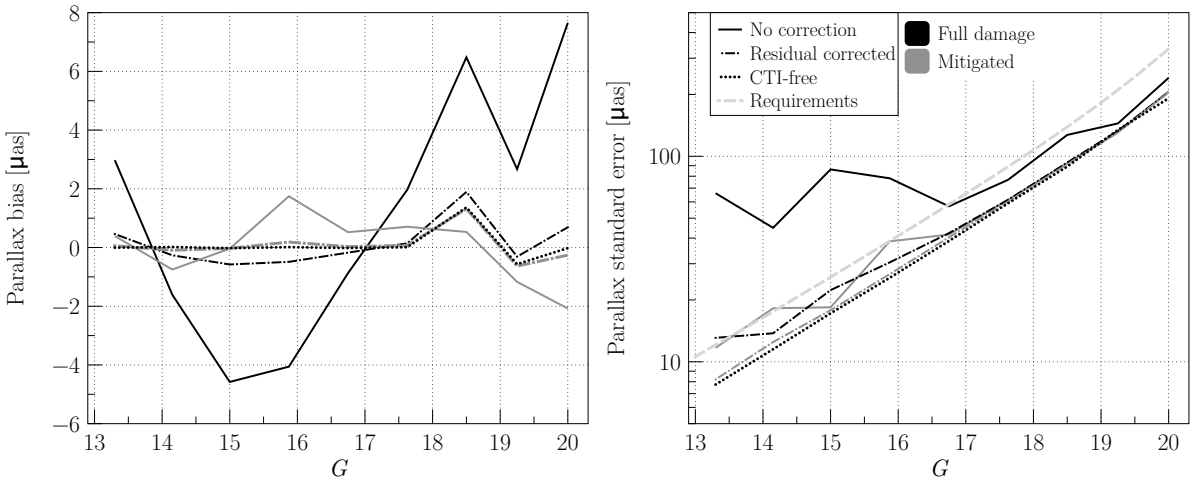


Figure 4.17 — Mean parallax error (left) and standard error (right) per magnitude, with and without the residual correction described in Section 4.3.3.

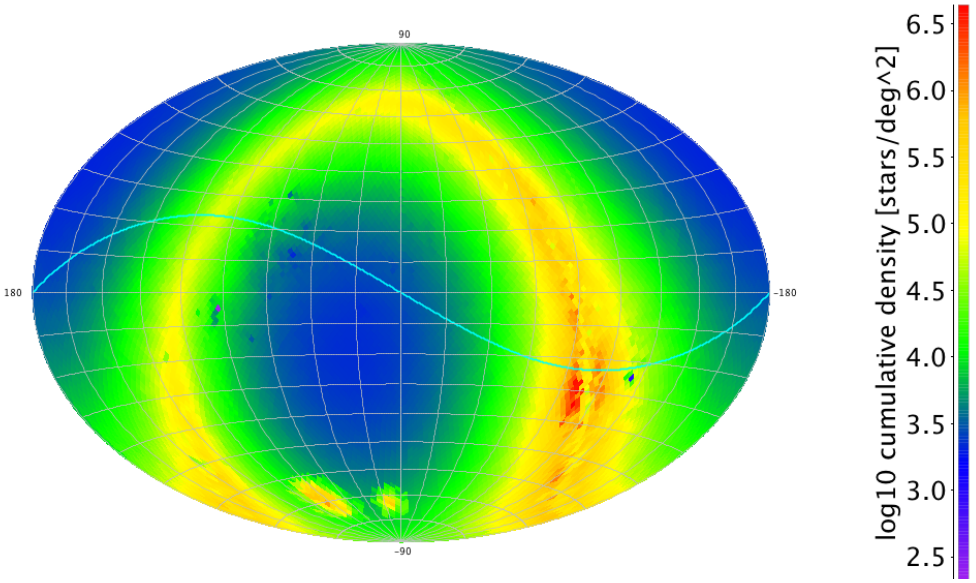


Figure 4.18 — Cumulative star density distribution $A(G, p)$ over $G = 4$–21.

Frame rotation:
As explained in Section 4.2.3.2, the astrometric solution produced by AGIS must be tied to a global reference system through the application of a frame rotation. We use the true positions and proper motions of the primary stars to fix the orientation and spin of the reference frame, so that the analysed astrometric errors are completely free of any effects due to frame misalignment.

4.2.3.4 Simulation output

As part of the standard output of the AGISLab solution, many automated log files and plots are generated concerning the convergence behaviour (the error and update distribution of each parameter for every iteration), histograms, sky maps, time plots, etc. What is most relevant for the post-processing needed in this chapter were the following: the true astrometric parameters, the estimated astrometric parameters together with an estimate of their standard errors, the true and estimated attitude, and the binned statistics of the AL residuals and of the errors in the astrometric parameters. The results shown below have been derived from these data.

4.3 Results

Using the CTI model and AGISLab described in the previous section, several simulations have been made to characterize the effects of radiation damage on the astrometric solution. The main difference between these simulations is the level of perturbations applied to the along-scan observations of the astrometric field CCDs according to Eq. (4.13). The following three cases are considered (cf. Section 4.2.2.2):

1. *The CTI-free case:* In this case no bias is applied ($\delta = 0$) and the magnitude-dependent standard deviation (σ) is entirely due to the photon noise; see the dotted curves in Fig. 4.6.
2. *The full-damage case*: The bias (δ) and standard deviation (σ) applied to each AL observation are functions of the magnitude (G), the time in the mission (t_{m}), and the time since the preceding charge injection (t_{CI}) as described in Section 4.2.2.1 and Fig. 4.7. Note that the bias varies from zero to δ_{max}, and the standard deviation from the photon-noise value to σ_{max} (the maximum levels are shown by the solid curves in Fig. 4.3).
3. *The mitigated case*: The model is the same as for the full-damage case, but the maximum levels δ_{max} and σ_{max} are reduced as shown by the dashed curves in Fig. 4.3, based on the CTI mitigation model in Chapter 3.

Before presenting the results of these detailed simulations it is useful to consider how a constant along-scan bias would affect the astrometric parameters. The outcome of this highly idealised experiment (in Section 4.3.1) helps to interpret the results of the more realistic CTI models in Section 4.3.2. The possibility to identify and partially correct residuals in the astrometric solution is discussed in Section 4.3.3. Finally, the astrometric error due to disturbing stars between the charge injection and the target star is discussed in Section 4.3.4.

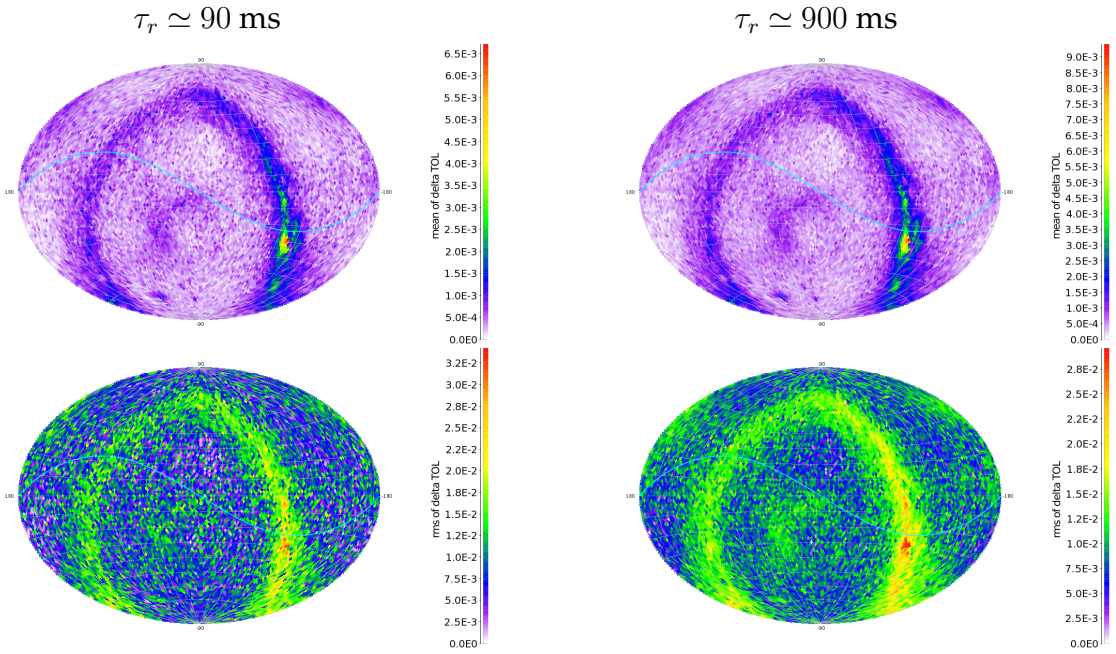


Figure 4.19 — Trap occupancy difference due to the inclusion of disturbing stars, showing the mean (top) and standard deviation (bottom) as function of position on the sky. The left maps are for the trap species with $\tau_r \simeq 90$ ms, as used in most of this investigation. The right maps are for a trap species with $\tau_r \simeq 900$ ms. The influence of the high density regions is visible over a large fraction of the sky due to coupling of the two fields of view in combination with the scanning law. The regions around the Galactic poles are enhanced because the other field of view must be at a low galactic latitude (as the basic angle is $106.^{\circ}5$).

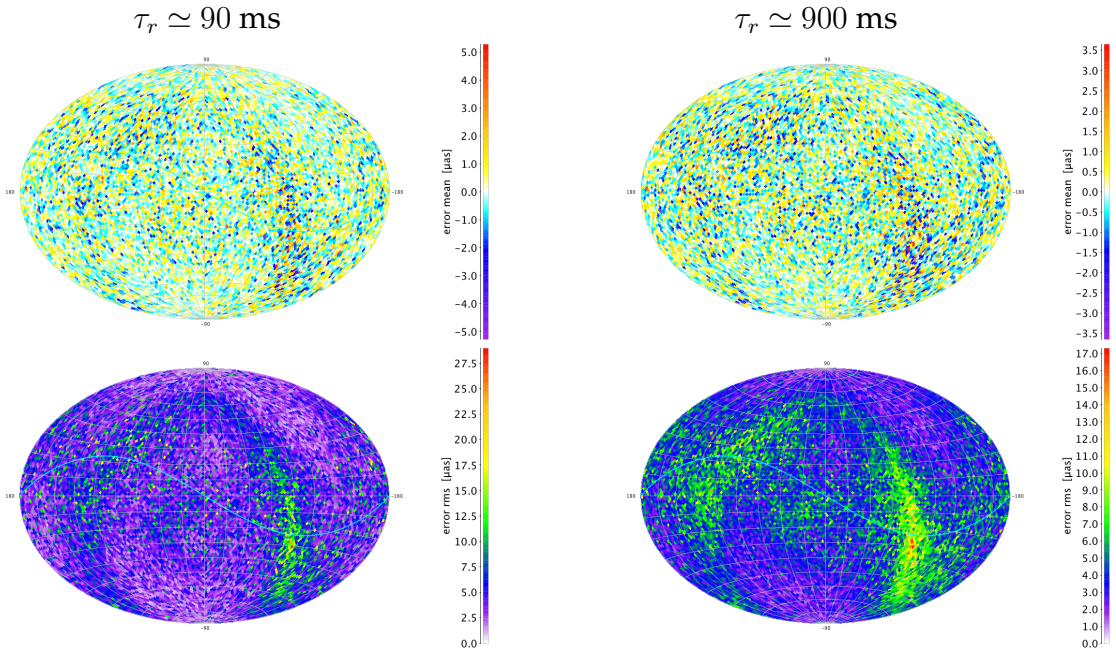


Figure 4.20 — Additional parallax error due to disturbing stars, showing the mean (top) and standard deviation (bottom) as function of position of the sky. The left maps are for the trap species with $\tau_r \simeq 90$ ms, as used in most of this investigation. The right maps are for a trap species with $\tau_r \simeq 900$ ms. No systematic bias patterns are present, only an increase in the random error level where the trap occupancy level rms was elevated due to disturbing stars.

4.3.1 How a constant bias would affect the astrometric solution

In this section we ask the question what would happen to the astrometric parameters, if all the observations had a constant, positive bias ($\delta = \text{const} > 0$).

The short answer is that there is no effect whatsoever on the astrometric parameters. The reason is that the observation bias is completely absorbed by the attitude estimate. A constant $\delta > 0$ means that all the measured observation times of all the stars are delayed by the same amount, which is equivalent to a constant orientation error of the instrument around the spin axis, i.e., to a certain offset in the attitude. Since the attitude is solved simultaneously with the astrometric parameters, and the reference frame is adjusted to the true astrometric parameters, the net effect on the astrometric parameters is zero while the full bias is absorbed by the attitude. The result will be the same if δ is a (smooth) function of time, but otherwise the same for all observations at a given time, for example if the bias increases gradually over the mission.

So why consider the effects of a constant bias? This is because in reality stars of different magnitude will in general have different biases (as we have simulated). As stars of different magnitudes are observed simultaneously all the time, a certain AL attitude rotation offset cannot compensate for the biases of all stars at the same time. It will however compensate such that the weighted sum-square of the observation time residuals in Eq. (4.15) is minimized, i.e., by making the weighted sum of residuals equal to zero. Effectively, this means that the (weighted) *mean* bias is absorbed by the AL attitude, subtracted from the residuals, and causing no harm to the astrometric parameters. Indeed, this behaviour is readily seen when the statistics of the residuals in Figs. 4.9–4.10 are compared with the actually applied errors in Figs. 4.6–4.7.

Although the overall mean bias thus magically disappears, there remains for almost every magnitude a non-zero residual bias (viz., the difference between the applied bias and the overall mean bias), which will affect the astrometric parameters in different ways depending on the position of the sky and the details of the scanning law. It is these patterns that we want to consider presently.

To examine the propagation of a fixed bias to the astrometric parameters, without having it trivially absorbed by the attitude, we make an astrometric solution in which the attitude parameters are not estimated but remain at their true values. Fortunately this is extremely simple in AGIS (or AGISLab): it simply requires that the iterations are stopped after the first update of the star parameters. Otherwise we can use exactly the same scanning law, mission parameters and model formalism as described in the previous sections, and the simulation output can be analysed in the same way. For this experiment we adopt $\delta = 1$ mas (i.e., 16.7 μs in the observation time) and $\sigma = 0$. We do not apply any frame rotation to this data, so any global rotation caused by the bias is preserved. The resulting error maps are shown in Fig. 4.8. A fixed color scale is used for all the maps to facilitate comparing the error levels for the different astrometric parameters (see also the discussion in Section 4.2.3.2). Some features of the error maps are briefly commented hereafter.

One of the most striking features of Fig. 4.8 is the relative uniformity of the parallax error map (the top right diagram, marked ϖ), with generally much smaller errors than in position ($\alpha*$, δ). This is probably related to the fact that the parallax signal is pe-

riodic, with a period (1 yr) much shorter than the mission (5 yr), meaning that the different phases are well sampled, in different scan angles, by the quasi-randomized scanning law. It is difficult to get the amplitude (i.e., the parallax) systematically wrong unless there happen to be many observations in a short time with the bias in the same direction. As was discussed in Section 4.2.1.2 this only happens for ecliptic latitude $|\beta_{\rm e}| < 45^\circ$ where the north–south arc-like structures are repeatedly scanned in the same direction within a few weeks, and this is precisely where the largest parallax errors are located. The other astrometric parameters correspond to the measurement of a linear motion and its origin on the sky, which is a lot more sensitive to systematics in the scanning law angles over time.
Another interesting feature is the predominantly negative errors in $\alpha*$ and the positive–negative dichotomy of the errors in δ (with the neutral meridian at $\alpha = \pm 90^\circ$), both of which can be explained by a negative rotation in the ecliptic plane by about 200 μas. This is probably explained by a subtle asymmetry of the nominal scanning law: because the precession of the spin axis is such that its speed with respect to the stars is approximately constant, the (positive) spin axis spends more time south of the ecliptic than in the northern hemisphere, and therefore the mean angular velocity of the satellite, when averaged over a year, is negative. A positive observation time bias then translates into a negative bias in ecliptic longitude. This global rotation is however easily offset locally; for example in the error map for δ the yellow arcs at $|\beta_{\rm e}| < 45^\circ$ indicate an over-abundance of 'upwards' scans, and the blue/magenta arcs an over-abundance of 'downwards' scans. For the errors in proper motion, the temporal distribution of the observations introduces a further level of complication, suggested by the large yellow and blue patches in the bottom diagrams of Fig. 4.8.
The overall RMS astrometric errors for the 1000 μas bias are 258 and 204 μas (in $\alpha*$ and δ), 168 μas (in ϖ), and 161 and 174 μas yr^{-1} (in $\mu_{\alpha*}$ and μ_δ). In other words, thanks to the clever way the scanning law has been defined, the propagation of an observational bias into the astrometric parameters is already 'mitigated' by a factor $\sim$5 due to the very efficient averaging of scans in different directions and at different times.

4.3.2 Detailed model results

In this section we analyse the results of the detailed CTI models described in earlier sections, and look in turn at the solution residuals, the attitude errors, and the astrometric errors.

4.3.2.1 Time residuals

In Figs. 4.9 and 4.10 we plot the statistics of the observation time residuals $t_l^{\rm obs} - t_l^{\rm calc}$ in the same manner as was done for the applied errors in Figs. 4.6 and 4.7.
Looking at the left diagram of Fig. 4.9, it is seen that the mean residual, when averaged over all magnitudes, is zero in all three cases (CTI-free, full-damage, and mitigated case). Comparing with the mean applied errors in Fig. 4.6 (left), the curves are virtually the same only that the full-damage curve is shifted to zero mean. This can be understood as the attitude solution absorbing the mean bias as a function of $t_{\rm m}$ as discussed in Section 4.3.1. The standard deviations in the right diagram of Fig. 4.9 are

also similar as for the applied errors, only slightly reduced in the full-damage case, which is explained by part of the bias variation, depending on the factor $f_\rho(t_\mathrm{m})$, being absorbed by the attitude. This is shown more clearly in Fig. 4.10, where the residuals in the full-damage case have been binned according to t_CI and t_m: each contour plot shows at least one particular t_CI where the mean residual is constant over the mission. A plot of the attitude errors (Fig. 4.11) shows the expected error pattern versus t_m in the along-scan attitude, exactly mirroring the assumed evolution of the accumulated radiation dose in Fig. 4.4.

4.3.2.2 Astrometric errors

Ultimately it is the effect of radiation damage on the astrometric parameters that is our main concern. In Figs. 4.12 and 4.13 we plot the sky-averaged astrometric errors in parallax and right ascension versus magnitude, while Fig. 4.14 shows the distribution of the parallax errors across the sky for selected magnitudes. The CTI-free case is included in all the figures as a reference; as expected, it shows negligible bias at all magnitudes and the fine-grained pattern in the CTI-free maps is entirely due to photon noise.

For the full-damage and mitigated cases there are magnitude-dependent biases in both parallax and right ascension, similar to the applied errors in Fig. 4.6, shifted to zero mean bias (and hence to the mean residuals in Fig. 4.9), but with the opposite sign. This can be understood in relation to the error maps in Fig. 4.8 for a constant bias: observations with a positive bias result in parallax and right ascension errors that are both negative, but the overall effect is much smaller in parallax than in right ascension. The behaviour of the other astrometric parameters can readily be predicted by similarly combining Figs. 4.6 and 4.8

When comparing the error maps of Fig. 4.14 with those in Fig. 4.8, it is seen that the error patterns due to the constant bias (at each magnitude) is clearly imprinted on top of the photon-noise error, except for the mitigated case at $G = 15$, where the bias is practically zero. Thus the level and the sign of the astrometric biases are related to the mean residuals found in Fig. 4.9.

Concerning the standard deviations of the astrometric parameters shown in the right diagrams of Figs. 4.12 and 4.13, it can be noted that the spatial variations of the biases at a particular magnitude increases the standard deviations. This is especially noticeable in right ascension, where the biases are generally much stronger than in parallax. In the full-damage case it even dominates the photon noise except for $G = 20$. Comparing with the residual plot in Fig. 4.9 we may conclude that the total standard error is composed of the CTI-free (photon-noise) component together with the (scaled) absolute value of the residuals.

A main conclusion here is that the complex processing through AGIS preserves the sign, amplitude, and spatial pattern of the astrometric biases expected from the simplified analysis in Section 4.3.1.

4.3.3 Residual based correction

In Section 4.3.2 we saw that the applied CTI biases translate into a similar but shifted residual pattern. The question then arises if we can use this residual information to 'correct' the input observations. In the Gaia forward modelling approach described in Chapter 3 we mentioned a feedback from AGIS to the image parameter extraction process to improve the modelling. The residual pattern we are discussing here is part of this feedback, as it contains information about the (relative) bias present in the observations. This feedback of the results from AGIS can be done in two different ways.

In the first way, the improved star parameters and attitude are used to compute the 'true' locations of the image centres in the CCD pixel data, which are then used to improve both the instrument model (PSF) and Charge Distortion Model (CDM). Ultimately the CTI calibration will thus be improved, and the next astrometric solution should give smaller residuals and improved parameter values, which are fed back to the CTI calibration. This iterative process can go on until the systematic pattern seen in the residuals (e.g., as a function of G, t_{CI} and t_{m}) is entirely removed. This procedure is the baseline adopted for the processing of the real Gaia data, and corresponds closely to the forward modelling approach discussed in Chapter 3. It has the important advantage that the CDM, once properly calibrated by means of the the primary stars, can then also be applied to more complex objects such as double and multiple stars. In a sense, the full-damage and mitigated cases considered above correspond to the first and last iterations in this process (short-cutting the intermediate iterations by using the true locations to calibrate the CDM in Chapter 3).

The second way is to use the mean residuals of AGIS directly as a calibration of the CTI effects. This is much simpler than the forward modelling involving the CDM, but has the disadvantage that it only works for simple objects like the (apparently single) primary stars. However, it worth investigating both as a possible fall-back solution (in case the CDM is not accurate enough) and as an exercise in how to interpret and make use of the astrometric residuals. The adopted correction procedure is very simple. After a first astrometric solution of the full-damage case, the residuals were binned exactly as shown in Fig. 4.10, using $9 \times 20 \times 200$ bins in G, t_{m}, and t_{CI}, respectively, and the mean residual was computed in each bin. A second astrometric solution was then obtained, using the same observations corrected by subtracting the mean residual of the corresponding bin. The same procedure was applied to the mitigated case as well. Figures 4.15 and 4.16 show the residual statistics after the second (corrected) astrometric solution. The mean residuals (left) are largely reduced and consequently the residual standard deviations are also considerably reduced. Figure 4.17 shows the resulting parallax bias (left) and standard errors (right) for the full-damage and mitigated cases, with and without the residual-based correction. The left diagram demonstrates that the combination of a CDM-based mitigation at the image parameters estimation level and a residual-based correction allows virtually unbiased estimates of the parallax (red dashed curve), barely deviating from the CTI-free case (dotted line). In this case, the parallax standard error lies just above the CTI-free case and this from bright to faint magnitudes. In Section 4.4 we discuss the agreement of these results with the requirements.

4.3.4 Disturbing stars

Recall that disturbing stars are images that happen to fall between the target stars and the preceding CI, and in practically the same pixel column as the target star (see Section 4.2.2.5). By changing the illumination history and consequently the trap occupancy level immediately prior to the observation of the target star, they introduce a source of noise in the CTI calibration and correction procedure. Due to the short CI period ($\sim$1 s) envisaged for Gaia, only a very small fraction of the sky has high enough density to introduce a significant number of disturbing stars (e.g., only $\sim$0.3 per cent of the sky has >2 disturbing stars brighter than $G = 21$ within a CI interval of 1 s).
To examine the additional effect of disturbing stars, a simulation was made in which the bias was calculated as

$$\delta = \delta_{\max}(G)\ f_\rho(t_{\rm m})\ \Delta f_{\rm IH} \tag{4.16}$$

where $\Delta f_{\rm IH}$ is the difference between the illumination history factor with disturbing stars (Section 4.2.2.5) and without (Eq. 4.9). Together with the suppression of photon noise, i.e. $\sigma = 0$, the result of this AGIS solution gives the astrometric error component due to disturbing stars for an instrument that is perfectly calibrated for the undisturbed case. Note that $\Delta f_{\rm IH} \equiv -\Delta\theta \leq 0$ because the trap occupancy level can only be increased by disturbing stars. The rest of the experimental setup was equivalent to the previous ones.
As was mentioned in Section 4.2.2.5, the computation of the effect of disturbing stars on the trap occupancy level depends strongly on the release time constant. Throughout this investigation we have used $\tau_r \simeq 90$ ms, which is relatively short with respect to the CI interval of 1 s. To get an indication of the astrometric parameter dependence on this release time we made an additional experiment with $\tau_r \simeq 900$ ms. This longer release time causes disturbing stars longer before the target star to be of significant influence, while still giving a significant variation in trap occupancy level due to electron release within a charge injection interval.
It is important to recall that the disturbing star scene of each field of view transit is based on the combined field of view star density. The amount of disturbing stars associated with a field of view transit of a particular star therefore depends on the stellar position on the sky and the location of the other field of view. Therefore it minimally contains the density at the star's position on the sky. A star located at a dense region in the Galactic plane will therefore always have a high number of disturbing stars, while a star away from the Galactic plane can have large variations in disturbing stars per transit depending on the location of the other field of view. Figure 4.18 shows the cumulative star density distribution $A(G, p)$ between $G =$ 4–21, which can be used as a basis for interpreting the following results.
In Fig. 4.19 the mean and rms value of the (always positive) trap occupancy level difference is shown for the two examined charge release time constants. As can be clearly seen, the high density regions on the sky have the largest mean and rms value. Also interesting to see is that the galactic pole regions are enhanced as well, because the other field of view must be at a low galactic latitude. The mean value is not particularly imporant as this will largely be absorbed by the attitude. It is observed that a

longer release time constant results in a higher trap occupancy difference and a lower, but more equally spread out rms level. Considering the ten times longer charge release time used, the difference is however very small.
In Fig. 4.20 we show the effect of the disturbing stars on the mean and rms of the parallax error. The top maps show that the average error is close to zero everywhere and that there are no large-scale patterns of systematic errors. The rms maps show that the disturbing stars do however increase the (random) error at the regions on the sky where the trap occupancy was elevated, as expected. The global parallax rms levels are 5.6 and 4.6 μas, for $\tau_r \simeq 90$ and 900 ms, respectively. The small difference in trap occupancy levels between the two release time constants is propagated to give a similarly small difference in final astrometric errors. When we compare the disturbing star rms levels to the parallax standard error for different magnitudes (Fig. 4.12) we can conclude that the disturbing stars, if unmodelled, would add an error which is small, but still significant for the brightest stars. Another important conclusion is that although the highest stellar density regions are concentrated along the Galactic plane, their influence is visible over a much larger part of the sky due to the coupling of the two fields of view in combination with the scanning law.

4.4 Discussion

4.4.1 Implications for the scientific performance of Gaia

The design and technical development of the Gaia instrument have, to a large extent, been dictated by specific mission requirements formulated already at an early stage of the project. In principle, they have been determined in such a way that their compliance should guarantee the feasibility of the main science goals of the mission from a technical viewpoint. In reality, the requirements must take into account numerous other constraints, including what is deemed technologically possible within a given cost envelope. A central part of the mission requirements specifies the astrometric accuracy that should be achieved. This is given in the form of maximum values for the sky-averaged standard errors of the parallaxes for unreddened stars of specific V magnitudes and spectral types. Translated to the G magnitude scale and interpolated, this corresponds to the dashed green curve in Fig. 4.17 (right). We emphasize that the requirement is for the sky-averaged standard errors, which admits some significant variation across the sky, and that it in practice only applies to apparently single and otherwise 'well-behaved' stars – indeed, it would be futile to try to take into account all levels of complexity in the real sky.
In spite of its coarseness, the formal accuracy requirement remains the standard against which the CTI errors must be assessed. In this section we discuss only the effects in parallax, as they are particularly important for the science goals of Gaia, and because the performance in the other astrometric parameters closely follow the performance in parallax (i.e., what is good for the parallaxes is most likely good for the other parameters as well).
At this stage it is important to summarize all the relevant simplifications that were necessary to perform this study. Regarding the first part of the study (Chapter 3) we assumed a unique trap species, a good knowledge of the instrument PSF, a perfect re-

moval of the background (including the CI background), and ignored the serial CTI. Considering the mitigated case only, the results were obtained for a very good (but yet not optimal) calibration of the mitigation procedure and in particular of the CDM parameters. In this chapter, we performed a simplified version of the astrometric solution (AGIS), only solving for the star astrometric parameters and the satellite attitude. The solution was performed on 1 million stars instead of the expected 100 million. This lead us to adopt a semi-realistic model for the star distribution in space and magnitude. We also ignored the observation dead-time induced by the periodic injection of charges in the CCDs and did not accounted for any other potential dead-time during the mission. Finally in our CTI-model we modelled disturbing stars only as a change in trap occupancy level based on a unique trap species. We constantly motivated those simplifications and tried to achieve the highest level of realism possible, nevertheless it is important to realize that these simplifications make the calibration of CTI simpler than it will be during the Gaia mission.

In this context, although we will discuss the Gaia requirements, it is always preferable to compare the results including radiation damage to the CTI-free case, and consider that the Gaia requirements are met if we are able to calibrate the CTI effects at the level of 10 per cent of the CTI-free case. We recall that for the level of damage considered in this chapter the maximum intrinsic and irreversible loss of accuracy in the image location estimation was found to be 6 per cent (see Chapter 3).

Figure 4.17 shows the sky averaged parallax bias (left) and standard error (right) as a function of G for the full damage and mitigated cases with and without an AGIS residual-based correction. From Fig. 4.17 (left) it is clear that without the residual-based correction the parallax estimation is biased both in the full damage and mitigated cases. In the mitigated case the bias however does not exceed 2 μas. As already discussed, when the residual-based correction is applied, the bias in the mitigated case reaches a satisfactory level (similar to the CTI-free case). Now looking at the right part of Fig. 4.17, it is clear that the full-damage case does not fulfil the requirements (green dashed line). The mitigated case without correction lies right below the requirements at bright signal levels. It is here however important to note that the requirements include dead-time in the observations. Because we did not include dead-time in our simulation we have a larger number of observations per star than the requirements assume, thus we find a better parallax standard error.

Figure 4.21 shows the relative deviation in parallax standard errors with respect to the CTI free case for the full-damage and mitigated cases, with and without the residual correction applied. As already mentioned we consider that the requirements are met if the relative deviation from the CTI-free case does not exceed 10 per cent. The 10 per cent accuracy loss interval is depicted by the shaded area. At bright signal levels ($G < 16$) the relative deviation between the mitigated case (without residual correction) and the CTI-free case reaches 50 per cent, and we can conclude that a mitigation at the level of the image parameter estimation only would not be enough to meet the requirements. Only a mitigation at the image parameter estimation level combined with a feedback mechanism from AGIS (grey dash-dotted line) allows the parallax accuracy to be recovered within 10 per cent of the CTI-free case.

The effect of disturbing stars in the time between the CI and the star has been stud-

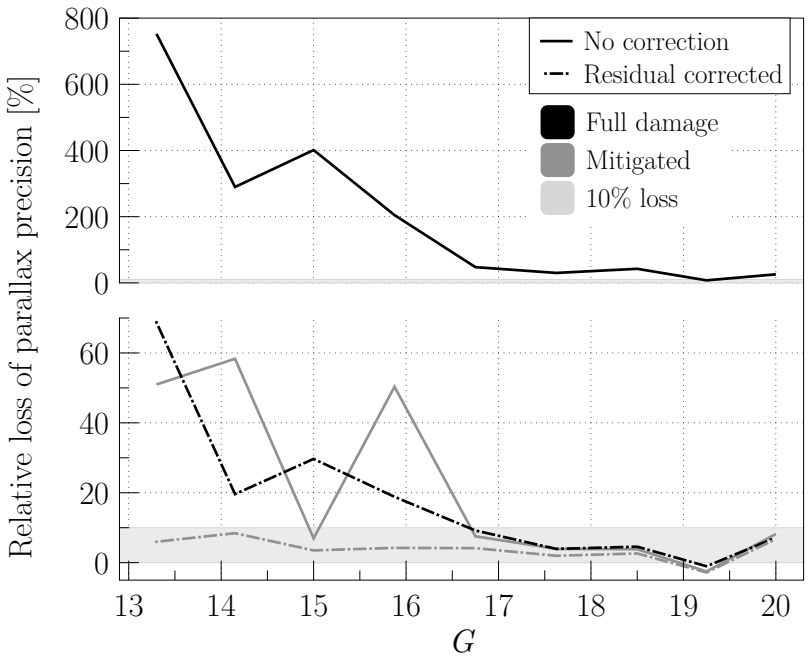


Figure 4.21 — Relative deviation in parallax standard errors with respect to the CTI-free case. This figure effectively represents the expected loss of accuracy for the full-damage (black) and mitigated (grey) case including (dashed-dotted line) or not (continuous line) a correction based on the solution residuals. Zero loss corresponds to the CTI-free case. The shaded area depicts the 10 per cent loss interval. Only a mitigation at the image parameter estimation level combined with a residual based correction (grey dashed-dotted line) allows the parallax accuracy to be recovered within 10 per cent of the CTI-free case for the whole magnitude range.

ies using an analytical trap occupancy model with a CI interval of 1 s, a CI level of 17,000 e^-, a realistic sky density model, and a trap species with $\tau_r = 90$ ms. The resulting parallax error does not show any systematic biases over the sky. The overall parallax error rms level found was $\sim$4 μas (i.e. the *additional* error due to disturbing stars), most strongly localized around the Galactic plane, somewhat weaker around the Galactic poles, and weakest in between, as shown in Fig. 4.20. In addition to the standard release time constant of 90 ms, also 900 ms was tested, showing only a reduction of $\sim$18% in the propagated error level. This suggests that the possible presence of other traps with longer charge release times in the Gaia CCDs does not significantly change the effect of disturbing stars on astrometry. The found rms level is a small but significant fraction of the photonstatistical error for the brightest stars, see Fig. 4.12, therefore it may be necessary to include the complex treatment of the illumination history at the level of the image parameter extraction for the brightest stars, especially at the highest density regions on the sky. We want to stress however that our model for computing the trap occupancy level and the corresponding location error due to disturbing stars has not been validated or tested against Monte-Carlo simulations or real experiments, therefore one should be careful not to over-interpret the presented astrometric errors associated with disturbing stars.

4.4.2 CTI mitigation in Gaia

In Chapter 3 and the present chapter we have studied the process of detecting individual photons in the presence of radiation damage of the CCD and the propagation of resulting CTI effects up to the final astrometric parameters. In this process we have been able to identify which mechanisms contribute to mitigate the CTI effects and investigate how effective they are. We have identified seven such mechanisms:

1. *Supplementary Buried Channel (SBC):*
This passive hardware mitigation comes from a doping profile that runs through the CCD and confines the volume of the charge package at low signal levels (i.e., $G > 15$), drastically reducing the number of traps encountered. Simulations without the SBC show location estimation biases that are 2–5 times higher for $15 \leq G \leq 20$ than including the SBC, and an even larger deterioration in location estimation precision for $G > 18.5$ (see Chapter 6).

2. *Charge injection (CI):*
This active hardware mitigation is a periodic injection of artificial charges which fills a large fraction of the empty traps and resets the illumination history. It is difficult to assign this an effective mitigation factor, but it is clear that the data processing is hugely simplified as the illumination history is dominated by the CI, making the effect of disturbing stars almost negligible (at least for most of the observations).

3. *Charge Distortion Model (CDM):*
This analytical non-linear distortion model is used in the forward modelling of the predicted observation counts. As seen in Chapter 3, an optimally calibrated CDM could potentially reduce the biases by a factor 10 and recover the location estimation accuracy to within the Gaia requirements. As we neglected some important but difficult to assess aspects like background subtraction and detailed line spread function calibration, it is not clear to what degree this will be possible with the real mission data.

4. *Sky background:*
Although the background level in the Gaia CCDs will remain very low due to their operation in TDI mode, the few background electrons that are constantly present will fill a fraction of the traps. For instance, in Chapter 3 we showed that the location bias is significantly reduced for faint stars by the background. Experiments also showed that a slight variation in the level of background illumination has a significant impact on the charge loss, e.g., from 0.3 e^- pixel^{-1} to 5 e^- pixel^{-1} reduces the measured charge-loss from $\sim$30 per cent to $\sim$10 per cent at 18th magnitude (Short et al. 2010; Brown 2009a).

5. *Scanning law:*
As described in Section 4.2.1.2, the scanning law has the unique property that it scans stars in different orientations, which already gives a bias reduction of about 5 times when analysed in terms of astrometric signal. This is most effective for parallax, which shows virtually no bias for ecliptic $|\beta_e| > 45°$.

6. *Attitude:*
As described in Section 4.3.2.1, the attitude determination absorbs any (slowly varying) rotation offset, including the CTI bias for a given magnitude. Because stars of different magnitude have different biases and these stars are observed simultaneously, the attitude offset cannot compensate for all stars at the same time, but it will remove the mean bias from the observations.

7. *Residual feedback:*
As described in Section 4.3.3, processing data of all magnitudes simultaneously allows us to accumulate the residuals in multiple dimensions (e.g., magnitude, time in mission, time since CI) and to identify systematic variations. During the data processing, this information will be fed back to the image parameter estimation to improve the calibration of the PSF and CDM models. Alternatively, or additionally, the residuals can be used directly as corrections on the data.

The last three mechanisms are only possible by processing all the observations and solving for all parameters together, as done in AGIS.

4.5 Conclusions

This chapter is the second and last part in a study that aimed at characterizing and quantifying the impact of CCD radiation damage on the final astrometric accuracy of Gaia. Here we focused on the effect of the image location errors induced by radiation damage on the Gaia astrometric solution, AGIS. To do so we applied a simplified version of AGIS, only solving for the astrometric parameters and the attitude of the satellite, to a set of synthetic Gaia-like observations (8×10^8), including CTI errors, generated for 1 million stars with a reasonable distribution in magnitude and on the celestial sphere. For most of the stars, we only conserved the along-scan CCD observations in the astrometric field (AF) due to their predominant weight in the solution. We modelled the radiation damage-induced bias and increased location uncertainty by adding Gaussian noise with non-zero mean and a widened standard deviation to the true observations. For this purpose, we developed a realistic and fast-to-apply model of the CTI errors based on the results from Chapter 3, considering a unique trap species and a maximum active trap density of 1 trap pixel^{-1}. This model determined the bias and standard error to be applied as a function of G for a particular observation accounting for the increasing accumulated radiation dose along the mission and the temporary mitigation of the CTI effects by the periodic injection of charges in the CCDs.
In this way and for the first time we rigorously propagated the image location bias as well as the increased random errors through a realistic astrometric solution, and this for two different levels of mitigation at the image processing stage. We also investigated whether the solution residuals can be used to improve the calibration of the CTI effects. This allows us to assess the impact of CCD radiation damage on Gaia astrometry, and we can draw the following conclusions:

While the mean of the CTI-induced bias is absorbed in the attitude modelling, the variation with magnitude and other factors (e.g., the illumination history) is propagated to the astrometric parameters. Thus, except for the trivial case of a slowly evolving but otherwise constant CTI bias, the astrometric results are biased unless some measure is taken to calibrate the effect.

The satellite scanning law and the distribution of the scan angles reduces the CTI-induced bias, which always occurs in the same direction with respect to the focal plane. This process also implies that a part of the residual bias appears as an increased standard deviation of the solution residuals and ultimately of the astrometric parameters.

The distribution on the sky of the CTI-induced errors for an astrometric parameter is not uniform but is imposed by the scanning law and the nature of the measurement for this particular parameter. The bias accumulates in regions of the sky for which the distribution of the scan angles is strongly anisotropic. For all the astrometric parameters, this could lead to significant systematic errors in particular in the ecliptic zone ($|\beta_\mathrm{e}| < 45°$) if the CTI effects are not fully calibrated out.

Among the astrometric measurements of Gaia, the parallax determination is least affected by these zonal errors and thus most robust against CTI. This is probably related to the periodic nature of the parallax signal, which makes it more like a repeated differential measurement than the determination of position and proper motion.

A CTI mitigation procedure at the level of the image location is necessary to reach the best agreement possible, set by the photon noise, between observations and AGIS predictions. Without applying the forward modelling approach as described in Chapter 3, the loss in parallax precision in unacceptable especially at bright magnitudes. This once again shows that hardware counter-measures, although needed, do no suffice, and that the CTI effects must be taken into account in the Gaia data processing.

The systematic variation of the CTI-induced bias with G and time since charge injection is conserved in the residuals of the astrometric solution. Thus one can use this imprint of the CTI effects left in the solution residuals to feed information back to the image parameter estimation and ultimately recover the astrometric accuracy. This works also for the errors remaining after the CTI mitigation procedure at the image location level.

For a charge injection period of 1 s the effect of disturbing stars on the CTI calibration was found to be small but non-negligible for the brightest stars, being spread out over a large part of the sky. Although the typical number of disturbing stars within a charge injection period is much smaller than one, stars with transits having (at least) one of the fields of view pointing close to the Galactic plane will experience fluctuations in the trap occupancy level that are enough to introduce additional astrometric

parallax errors on the order of a few μas in our model, which is a small but significant amount for the brighter stars. This suggests that a complex treatment of the illumination history at the level of the image parameter estimation might be needed for the brightest stars.

We have demonstrated that it is possible to calibrate the CCD radiation damage in the Gaia data processing, such that, despite the complexity and importance of the CTI effects on the stellar images, it is possible to recover a virtually bias-free estimation of the astrometric parameters of single stars and to achieve the intended astrometric accuracy for these objects. This is rendered possible by the joint actions of the hardware CTI counter-measures and the CTI mitigation approach at the image parameter estimation level (developed in Chapter 3). To preserve the Gaia astrometric accuracy and reach the scientific requirements for bright stars, the residuals from AGIS must be utilized to feed back information to the image parameter estimation for each CCD observation in order to improve the CTI mitigation at this level. In this chapter we have demonstrated that when taking into account all these CTI mitigation counter-measures the parallax standard errors for single stars can be preserved within 10 per cent from the CTI-free case, for all simulated magnitudes between $G = 13$ and 20.

Acknowledgments

The work of BH and TP was supported by the European Marie-Curie research training network ELSA (MRTN-CT-2006-033481). LL acknowledges support by the Swedish National Space Board. AB acknowledges support by the Netherlands Research School for Astronomy (NOVA). The authors would like to kindly thank Jos de Bruijne for the discussion about the set-up of our experiments.

Chapter 5

Stress-testing a fast analytical Charge Transfer Inefficiency model

ESA's Gaia mission aims to create a complete and highly accurate stereoscopic map of the Milky Way. The stellar parallaxes will be determined at the micro-arcsecond level, as a consequence the measurement of the stellar image location on the CCD must be highly accurate. During the mission solar wind protons will create charge traps in the CCDs of Gaia, thus drastically increasing the CCD Charge Transfer Inefficiency (CTI). CTI will distort the stellar images and induce a significant charge loss for all the Gaia measurements. If not properly mitigated, the CTI effects introduce a strong systematic bias in the image location estimation and cause a significant degradation of its precision. The Gaia Data Processing and Analysis Consortium chose to mitigate the CTI effects by an iterative forward modelling approach that requires the accurate modelling of the stellar image distortion and charge loss for each observation. In this scheme, the final astrometric accuracy of Gaia directly depends on the capability of a fast analytical model of CTI effects, a so-called Charge Distortion Model (CDM), to reproduce the image distortion. In this chapter we assess the capability of a fast analytical Charge Distortion Model, proposed by Short et al. (2010), to reproduce experimental data. We developed a rigorous procedure that compares at the sub-pixel level the model outcomes to images affected by CTI extracted from experimental tests. We show that the tested CDM candidate has the capability to accurately reproduce the test data acquired on a highly irradiated device operated in Time-Delayed Integration (TDI) mode for a wide range of signal levels and different CCD illumination histories. In particular we are able to demonstrate that the level of agreement obtained is enough to recover the CTI-induced image location bias by at least a factor ten, and thus enables the potential recovery of the required final astrometric accuracy. The calibration of such a model is however a complicated enterprise and potentially problematic. By comparing the performance of the model in different conditions of use, we discuss different calibration schemes and give recommendations for necessary improvements of this model that will ease the calibration process. The mitigation of CTI effects using a forward modelling approach, as done for Gaia, is also interesting for other space missions, in particular future ESA missions currently under consideration, Euclid and Plato.

T. Prod'homme, M. Weiler, S.W. Brown, A.D.T. Short, A.G.A. Brown, B. Holl
In preparation

5.1 Introduction

Gaia is a European Space Agency mission that aims to create the most complete and accurate stereoscopic map to date of the Milky Way by collecting parallaxes, proper motions, radial velocities, and astrophysical parameters for about one billion stars, one percent of the estimated stellar population in our galaxy (Perryman et al. 2001; Lindegren et al. 2008). Gaia, to be launched in 2013, will orbit around the second Lagrange point (L2) and constantly spin around its own axis such that its two telescopes scan a great circle on the sky several times a day. Due to the satellite spinning motion, the star images will not remain stationary during an observation but will transit across the focal plane. The stellar transits are recorded in a single focal plane covered by 106 CCDs. To integrate the stellar flux along the transits, the CCDs are operated in Time-Delayed Integration (TDI) mode. In this mode a CCD is constantly read out and the charge transfer rate is synchronized with the motion of the stellar images over the focal plane.

In order to reconstruct the astrometric motion on the sky for a particular star and estimate its astrometric parameters, the different times of observation must be determined very accurately. If the attitude of the satellite is known, determining a time of observation effectively corresponds to determining the stellar image location on each CCD before readout (Lindegren & Bastian 2011; Bastian & Biermann 2005). The anticipated Gaia astrometric accuracy is extreme and sets a stringent requirement on the estimation of the stellar image location per CCD star transit (see Chapters 3 and 4, and de Bruijne 2005). For instance the end-of-mission statistical parallax standard error is required to be 26 micro-arcseconds for a G2V star of magnitude 15, which corresponds to an uncertainty on the location of stellar images on the CCD of 0.0045 pixels (~ 0.27 milli-arcseconds) for a single transit.

The displacement of atoms in the silicon lattice of the Gaia CCDs by solar flare protons will result in the creation of energy levels in the semiconductor band gap. These energy levels trap the signal carriers (electrons in a n-type device) and consequently increase the CCD Charge Transfer Inefficiency (CTI). The stochastic capture and release by the traps of the charges result in a significant charge loss and distortion of the stellar images. Experimental tests conducted on Gaia irradiated CCDs (e.g., Hopkinson et al. 2005; Pasquier 2011) as well as in depth simulation-based studies (Chapters 3 and 4) have demonstrated that CTI effects will significantly affect the determination of the stellar image location and thus the final astrometric accuracy of Gaia, if not properly taken into account in both the Gaia data processing and the CCD design and operation. Indeed the charge loss induces an irreversible loss of accuracy (of the order of a few percent) that is independent of any image location estimator and can only be avoided by the use of hardware CTI countermeasures that physically prevent the trapping, such as periodical Charge Injections (CI). The stellar image distortion due to CTI introduces a bias in the image location determination, which can be as large as 0.2 pixels for a magnitude 15 star. This number should be compared to 0.0045 pixels, the required image location accuracy mentioned above. This important sensitivity of Gaia to CTI can be imputed to a number of factors (see Chapters 1 and 2), among which: a harsh radiation environment, a low level of shielding resulting from stringent weight constraints,

special CCD operating conditions (see Section 5.3), and a demanding measurement accuracy.

Hardware CTI countermeasures (Section 5.3) can only prevent a part of the trapping, and as a consequence CTI must be taken into account in the Gaia data processing. Correcting for CTI is a complicated task; CTI effects are stochastic, they depend non-linearly on the stellar brightness, and they are non-repeatable for a particular star or magnitude as they depend on the CCD illumination history prior to the stellar transit of interest. Besides it has to be noted that the calibration of any software-based procedure will be further complicated, first, by the lack of CTI-free measurements: as the Gaia launch date is close to the solar maximum activity, even the very first measurements will be affected by CTI to some degree. Secondly, by the lack of information regarding the CCD illumination history and the stellar observations themselves: due to a limited telemetry rate the Gaia observations will consist of truncated one-dimensional charge profiles of 6 to 12 pixels in the stellar transit direction (see Section 5.3).

For these reasons conventional CTI correction procedures cannot be applied and the Gaia Data Processing and Analysis Consortium (DPAC) has developed a novel CTI mitigation procedure that relies on the forward modelling of each observation including the radiation damage effects. In this scheme, the true image location and flux are estimated by comparing the actual observation to a predicted charge profile, for which the CTI-induced distortion is modelled by an analytical model of the CTI effects, a so-called Charge Distortion Model (CDM). This enables the mitigation for any kind of source including complex ones (such as binaries, triple stars, extended or resolved objects) and does not affect the noise properties of the observations by avoiding a direct correction of the raw data. In Chapter 3 we demonstrated that this approach allows for a precise and bias-free estimation of the true image location from an observation affected by CTI effects (referred to as damaged observations). The efficiency of such a scheme, and consequently the final Gaia astrometric accuracy, relies on the capability of a CDM to accurately reproduce the damaged observations. The distortion will have to be computed for each CCD observation (~ 700 over the full mission duration) of each star (~ 1 billion). Moreover the image, the Line Spread Function (LSF, the one-dimentional Point Spread Function), and the CDM parameters are obtained from the same data, thus necessitating an iterative procedure (i.e. repeated several times for each observation). Hence the CDM must at the same time realistically reproduce the radiation damage yet remain simple enough to be computationally inexpensive for use in the Gaia data processing pipeline.

The model presented by Short et al. (2010) constitutes the current best CDM candidate. Chapter 3 showed that this model is capable of reproducing simulated damaged observations accurately enough so that most of the CTI-induced bias in both the image location and flux estimation can be removed. For this purpose the damaged observations were simulated by a detailed and physical Monte Carlo model of CTI effects that simulates the transfer of charges at the CCD pixel electrode level (Chapter 2). The calibration of the model parameters was performed for each stellar magnitude and for a unique illumination history making use of the true instrument LSF and a high number

of observations. This chapter[1] goes one step further and assesses the model capability to reproduce *experimental data* under realistic calibration conditions and for a wide range of signal levels and different illumination histories. Making use of the results from Chapter 3 we are able to evaluate if the overall agreement between model predictions and experimental data reaches the requirements set by the Gaia astrometric accuracy. After a presentation of the CDM candidate (Section 5.2) and the description of the Gaia CCD design and operation (Section 5.3), we give an overview of the experimental data used for the purpose of this comparison exercise (Section 5.4). Section 5.5 provides a description of our comparison procedure and in particular how the CDM parameters can be calibrated to fit the damaged experimental data at the sub-pixel level. In Section 5.6, we present the highest level of agreement with the experimental data achievable by CDM and compare it to agreements obtained for a set of more realistic CDM parameter calibration schemes. Finally, in Section 5.7, we discuss the limits and performance of the tested CDM candidate in the context of the Gaia CTI mitigation procedure, and give recommendations regarding the use of CDM and its further improvement.

5.2 Charge Distortion Model

The current best CDM candidate for implementation in the Gaia data processing is the physically motivated fast analytical CTI effects model presented by Short et al. (2010), hereafter we directly refer to it as CDM. To cope with the computational speed requirement CDM suppresses the treatment of the numerous charge transfer steps required to transfer and build the signal from one end of the CCD to the other, and computes the signal transit in a single calculation. This implies several important simplifications, the most relevant for this chapter are detailed in the following (for a complete explanation see Short et al. 2010).

5.2.1 Principles

Statistical treatment of the trapping processes

The model is based on the common Shockley-Read-Hall formalism (Shockley & Read 1952; Hall 1952), which describes the capture and release of a charge by an individual trap as a decay process with an associated characteristic time constant. The Charge Distortion model does not compute the capture and release probabilities for individual traps. Instead the capture and release of electrons is treated statistically by integrating over the total amount of traps present in the CCD. It thus implicitly considers a homogeneous distribution of the traps within a particular geometrical confinement volume V. The latter corresponds to the volume occupied by an electron packet at Full Well Capacity (FWC).

Trap species

As a result of the displacement of silicon atoms in the CCD silicon lattice (caused by radiation damage), interstitial atom-vacancy pairs are created. These vacancies diffuse

1. The first results of this study were presented at the 2010 SPIE conference in San Diego (Prod'homme et al. 2010), this work now supersedes in methodology and scope the former results.

and bind with other vacancies or impurities (e.g., oxygen, carbon, phosphorus atoms) present in the lattice. A vacancy-impurity complex is usually referred to as a trap species as the energy level and capture cross-section σ vary from one type of complex to the other. In CDM an arbitrary number of trap species J can be used, and a set of three parameters is associated with each trap species: σ_j, τ_j the release time constant, and ρ_j the trap density. In principle the release time constant, τ, varies as a function of the trap energy level, the temperature T, and σ_j (see e.g., Chapter 2). However in CDM this parameter is directly set and thus independent from the temperature and the capture cross-section. As a consequence the trap species resulting from a fit of the CDM predictions to a particular dataset may not correspond to a real vacancy-impurity complex.

Constant electron density distribution within a growing signal confinement volume
The total number electrons trapped depends on the capture time constant τ_c. τ_c depends on the temperature T, the trap capture cross-section σ, and the electron density distribution n_e. In CDM, for a particular number of electrons N_e in a pixel, the density distribution is assumed to be constant within the charge cloud volume V_c:

$$n_e = N_\mathrm{e}/V_\mathrm{c} \tag{5.1}$$

The parameter β determines the growth of V_c as a function of signal level:

$$V_\mathrm{c} = V\left(\frac{N_\mathrm{e}}{\mathrm{FWC}}\right)^{\beta} \tag{5.2}$$

Despite the complexity of the Gaia pixel architecture (cf. Section 5.3), CDM considers only one parameter β for the entire range of signal levels: from 0 $\mathrm{e}^-$to FWC ($\sim$ 190 000 e^-). As we shall see in Section 5.6.3, this constitutes one of the main limitations of CDM when one aims to reproduce with a single set of parameters a set of stellar images including different magnitudes.

Illumination history
The CTI induced distortion and charge loss for a particular stellar image depends on the state of each encountered trap prior to the stellar transit of interest. The state of these traps is set by the CCD illumination history. As traps are not modelled individually in CDM but treated in a statistical way, for a particular illumination history CDM computes the total fraction of filled traps, the trap occupancy level. The effect on the trap occupancy level of both a discrete event (e.g., CI) and a continuous illumination (background light, S_DOB) can be taken into account. Depending on the use of CDM in image mode or TDI mode, the treatment of the illumination history involves different assumptions. In the following section along with the description of our specific use of CDM, we detail the associated assumptions.

5.2.2 Usage

In this chapter we assess the performance of CDM using its TDI mode only. This is mainly due to the fact that the experimental test data (Section 5.4) we aim to reproduce

Parameter	Description	Fixed value or [interval]
	CCD parameters	
T	temperature	163 K
V	geometrical confinement volume	24 μm $\times$ 11 μm $\times$ 0.75 μm
FWC	CCD full well capacity	190 000 e^-
P_{TDI}	TDI period	0.9828 ms
N_T	number of pixel transfers	4500
β	electron cloud growth parameter	$[10^{-5}$ - 1]
	Trap parameters	
J	number of trap species	[1 - 7]
ρ_j	number of traps per pixel	$[10^{-5}$ - $10^{2}]$
σ_j	capture cross-section	$[10^{-26}$ - $10^{18}]$ m^2
τ_j	release time constant	$[10^{-4}$ - $10^{2}]$ s
	Experiment parameters	
t_{CI}	time since last charge injection	$[3.0\,10^{-3}$ - $1.2\,10^{2}]$ s
S_{DOB}	diffuse optical background	[0 - 240] e^-sample^{-1} s$^{-1\star}$

Table 5.1 — This table presents the parameters necessary to perform a CDM simulation. We indicate the used value for the fixed parameters. And we present, for the free parameters, the interval over which they are allowed to vary during the calibration procedure.
*A sample corresponds to one pixel for a two-dimensional simulation and 12 pixels for a one-dimensional simulation.

with CDM were acquired with a CCD operated in TDI mode (to reproduce the in-flight Gaia operating conditions). However it is important to note that, although the Gaia CCDs will be operated in TDI mode, the imaging mode of CDM may also be used in the Gaia data processing: first, to model the effects induced by CTI in the serial register, and, second, to initialize in the most realistic way possible the trap occupancy level given a time since last CI. For this particular study the serial CTI can be neglected (see section 5.4.1). And, in order to compute the initial trap occupancy given a time since last CI, we assumed that all the traps are filled by a CI (Eq. (34) in Short et al. 2010). The validity of such an assumption is discussed in Section 5.7.2.
All the parameters of a CDM simulation are detailed in Table 5.1. During the whole comparison the parameters always kept fixed are: T, V, FWC, P_{TDI} the TDI period (i.e. pixel transfer period in the CCD image area), and N_T the total number of transfers in the image area. The fixed values are indicated along with the parameters in Table 5.1. From the test data we also set t_{CI} the time since last CI corresponding to each specific test (cf. Section 5.4). In principle the background light S_{DOB} could also be estimated from the RC data itself and set to a fixed value however we chose to let this parameter free as its estimation is not entirely reliable. The free, i.e. calibrated, parameters are thus the trap parameters for each trap species j: ρ_j, σ_j and τ_j, along with the background light S_{DOB}, and the electron cloud growth parameter β. The number of trap species J to be included in a CDM simulation is an important issue for the Gaia data processing. As we shall see in Section 5.6 the optimal number of trap species to be considered depends on the calibration scheme.
During the mission the Gaia CCDs will acquire two-dimensional stellar images, but

only one-dimensional binned images will be sent to ground (see Section 5.3). The question of using CDM to perform a 1D or a 2D forward modelling of the Gaia observations in the image parameter estimation and CTI mitigation procedures is thus a key issue. It is a problematic question in terms of calibration and instrument modelling, because the two-dimensional information about the instrument PSF will be very limited, only the brightest stars will not be binned. The use of relatively poorly characterized PSFs might introduce undesired effects in the sensitive and critical image location estimation. And it is also problematic in terms of computational power and time requirements, since the use of CDM to distort a 12 by 12 pixels CTI-free input signal requires 12 times the computational load of the same computation for 1 by 12 pixels signal. In the following we will thus asses the capability of CDM to reproduce experimental test data using it in 1D or 2D.

5.3 Gaia CCD design and operation

The Gaia CCDs (see Short et al. 2005) are back-illuminated and full frame devices. They are custom made by e2v technologies and referenced as CCD91-72. The CCD image area contains 4500×1966 pixels (parallel $\times$ serial). Among the 4500 pixel rows only 4494 are light-sensitive: six rows are blocked by an aluminum shield. Each pixel contains a supplementary buried channel (SBC, or notch) and an anti-blooming drain. The SBC corresponds to a particular doping implant that confines to a smaller volume (than the buried channel) small charge packets so that these small packets encounter less traps. As shown by Hopkinson et al. (2005) and as confirmed in Chapter 2, in Gaia CCDs, the SBC-induced CTI mitigation stops to be effective for charge packets greater than ~ 3000 e^-. Due to the spinning motion of the satellite and the operation of the CCD in TDI mode, it is important to realize that most of the Gaia measurements will be faint when the signal starts to be integrated (i.e. at the beginning of the star transit across the CCD). In fact, for more than half of their CCD transit, the entire image of stars fainter than magnitude [2] $G = 16$ are transferred in the SBC only. As a consequence, bright and faint stars will experience different levels of trapping. And the CTI effects cannot in principle be modelled across a wide range of signal levels without an explicit modelling of the SBC.

The Gaia CCDs allow the periodical injection of electrons. Along with the SBC, CIs constitute one of the main hardware CTI countermeasures to be used by Gaia. A CI consists of the injection of a certain number of electrically generated charges in the first CCD pixel row. The CI is then transfered through the CCD image area so that vacant traps throughout the CCD can be filled. A direct consequence is that stellar images transiting right after a CI experience less damage. The CIs also have the advantage of dominating and thus reseting the illumination history. This eases the process of understanding the CTI effects independently from the CCD illumination history. The traps with characteristic release time constants greater than the CI period are permanently filled by the artificially generated charges. However traps with shorter release time

2. The Gaia G-band magnitude is a broad-band, white-light magnitude in the wavelength range 300 – 1000 nm defined by the telescope transmission and CCD quantum efficiency. $G = V$ for an un-reddened A0V star (Jordi et al. 2010; Perryman et al. 2001).

constants remain active because they have the time to release their charges before the next CI. The expected CI period to be used on-board Gaia is 1 s.
As already mentioned, the Gaia CCDs will be operated in TDI mode to integrate the light along the star transits. Star transits will always be parallel to the CCD parallel transfer direction [3]. In the following we thus refer to it as the 'along-scan' direction or AL. Accordingly the CCD serial direction is referred to as the 'across-scan' direction, AC. Gaia will continuously produce a large amount of data on-board thanks to its 106 CCDs, but due to its limited telemetry capacity only a part of it will be downlinked. The transmitted data for each star corresponds to a window of pixels centered on the pixel containing the maximum flux. The AL width of a telemetry window will vary as a function of the star magnitude: 12 AL pixels for stars brighter than magnitude 16 and 6 AL pixels for the fainter stars. Also for most of the stars the telemetry window will be binned across-scan. As a result, only one-dimensional data will be available for these stars.

5.4 Experimental tests

To characterize and evaluate the impact of the CTI effects on the Gaia measurements and to test and optimize the different potential hardware CTI countermeasures, the industrial partners in the Gaia project performed several campaigns of laboratory experiments on irradiated CCDs with an architecture and parameters as anticipated for the Gaia flight CCDs. The prime contractor for Gaia, EADS Astrium, conducted up to now four of these test campaigns, from hereon referred to as Radiation Campaigns (RC). For these four RCs, Astrium used a test setup (Pasquier 2011) which closely reproduces the in-flight conditions and simulates the transit of stars over an irradiated Gaia CCD. In the following we give a brief description of the test setup. Then we present the subset of data we used in our comparison, which is extracted exclusively from Astrium's second RC (from hereon referred to as RC2). Finally we explain how the data acquired in the non-irradiated part of the tested CCD is used to build a CTI-free signal which serves as input to the CDM.

5.4.1 Experimental test setup

All the tests performed by Astrium were carried out on irradiated Gaia CCDs. The tested devices include an irradiated region with a radiation dose of 4×10^9 protons cm^{-2} (10 MeV equivalent), a non-irradiated region, and a transition region between them. The non-irradiated region is the closest to the readout node. In the irradiated region the serial register was also irradiated, but one can to first order ignore the effects of the radiation induced serial CTI. The stellar image located the furthest from the readout node will effectively undergo only a few hundreds of serial transfers in the irradiated region (to be compared to 4500 parallel transfers). The AC size of the window is large enough to contain the majority of the redistributed electrons, so that charge loss due to serial CTI can be ignored. Serial CTI distorts the shape of the AC profile, however this distortion cannot be observed due to the AC binning. For these reasons we chose to ignore the effect of serial CTI in this study.

3. The maximum deviation in the serial direction is expected to be of the order of 4 pixels at readout

The CCD was operated at different temperatures including the Gaia operating temperature of $163\,\mathrm{K}$. To produce point-like sources playing the role of artificial stars, a mask with punched holes was used. The artificial stellar brightness was determined by the illuminating source (LED) brightness. The mask was translatable to reproduce the star motion and to operate the CCD in TDI mode. Due to the continuous readout and the faint luminosity of the sky background, the background light in Gaia CCDs is expected to remain very low over the mission independent of the line of sight. As a consequence, particular care was taken to introduce as little background light as possible. To recreate in-flight conditions the stellar images were binned in the AC direction.

5.4.2 The test dataset

RC2 (see Pasquier 2011; Georges 2008; Brown 2009b) focused on the evaluation and characterization of repeated artificial charge injections as a potential CTI hardware countermeasure. The CI parameters are the level, duration and period. The CI level is the number of electrons injected in one pixel. The duration corresponds to the number of subsequent CCD pixel rows in which the charges are injected. Finally the period is the time separating two consecutive CIs. RC2 data proved to be particularly suitable for our comparison as each of the parameters was carefully investigated which provided us with a set of well defined and understood tests carried out at different temperatures, CI levels, durations, and delays between the charge injection and the first 'star' crossing the CCD.
To compare the CDM predictions and the experimental data we focus on a subset of the RC2 data at the operational temperature of Gaia for which the CI level and duration were kept fixed. The delay between the CI and the first star transit was varied as well as the level of illumination, or artificial star magnitude. This allows us to evaluate how well the CDM can reproduce with a single set of parameters a damaged star image or a set of images for different illumination levels ($13.3 < G < 20$) and for different time domains. The CI delays range from 30 ms to 120 s. Multiple scans were performed in each test and, as a result, for a single set of experimental parameters (temperature, star brightness, CI delay) we can extract an over-sampled damaged profile to compare with the CDM outcomes.

5.4.3 Modelling the CDM input signal

The accumulated data in the non-irradiated part of the CCD are used to create a reference curve. This curve is representative of the CCD illumination conditions and is used as input for the simulations with CDM. With only one dimensional stellar images available, the reference curves are actually LSFs (Line Spread Function, the PSF integrated in the across-scan direction), which are subsequently modelled using spline functions [4]. As already mentioned we want to investigate the performance of CDM in 1D and 2D. To generate a two-dimensional input signal for our simulations we use the original reference curve and assume that the AL and AC profiles are the same: $P(x, y) = L(x) \times L(y)$, where L is the undamaged reference curve and P the result-

4. The spline functions are characterised by the positions of a set of knots along the LSF profile and the parameters that describe the set of polynomials linked by the knots.

ing CDM input image. This process allows us to recover to some extent the original properties of the illumination set-up. For both one-dimensional and two-dimensional input, the integrated flux of the reference curve is scaled to produce a CDM input image for each artificial star magnitude.

5.5 Model-data comparison methodology

5.5.1 Comparison procedure

In this section we present the procedure we used for all the comparisons between the RC2 test data and the CDM outcomes. First a target charge profile and a corresponding undamaged reference curve are defined. The target charge profile is the experimental data including CTI effects that the CDM eventually aims to reproduce. Thanks to the damaged profile over-sampling, the comparison between the model and the data is better constrained and the derived parameters are independent of the sub-pixel position of the image profile. To perform a direct comparison between the experimental data points and the charge profile simulation with CDM, the sampling scheme specific to each of the target charge profiles is extracted. This sampling scheme is then applied to the PSF generated from the reference curve, in order to create the CDM image inputs. In this way the required number of CDM instances to generate an n times over-sampled simulated profile is n. Each of the n CDM instances is calculated with a single predefined set of parameters. Once the instances are completed, the individual CDM predictions are binned in the AC direction and each data point is then placed at the correct sub-pixel position according to the original sampling scheme so as to form an over-sampled predicted damaged profile. The actual comparison, i.e. the computation of the agreement (cf. Section 5.5.2), is performed using the over-sampled CDM prediction and target damaged profile.

5.5.2 Determination of the best-fitting CDM parameters

Our calibration procedure consists of finding a set of CDM parameters that optimizes the agreement between the CDM outcomes and a set of pre-selected data. The agreement between a predicted and an individual target over-sampled charge profile is quantified via the computation of a comparison criterion or goodness-of-fit parameter g. The latter is an altered version of the usual χ^2 and is defined as follows:

$$g = \frac{1}{S}\frac{1}{F}\sum_{i=0}^{S-1}\frac{\left(\lambda\left(k_i\right) - N\left(k_i\right)\right)^2}{\sigma_i^2} = \frac{\chi^2}{S\,F} \tag{5.3}$$

where λ is the simulated damaged charge profile, N the RC2 target charge profile, k_i a particular sub-pixel position, σ_i the noise, S the total number of data points, and F the total integrated flux. The noise is considered to be the quadratic sum of the photon-noise and the readout noise. The photon-noise is assumed to follow Poisson statistics with a standard deviation of $\sqrt{N}$ and r is assumed to have the constant value of 4.35 e^-(the typical measured value for the Gaia CCDs):

$$\sigma_i^2 = N\left(x_i\right) + r^2 \tag{5.4}$$

Depending on the chosen calibration scheme, one can attempt to fit with a single set of CDM parameters a variety of target profiles. But due to the fact that CDM does not provide a perfect representation of the CTI effects, we use two scaling factors, S and F, in our comparison criterion: S is required when one attempts to simultaneously fit different target charge profiles with different numbers of data points per pixel, and F to simultaneously fit charge profiles of different magnitudes. If these scaling factors are not present, the fit is biased towards the brightest magnitude and the most over-sampled profiles.
The calibration procedure involves two different optimization algorithms. In a first instance, we probe the entire parameter space (within the specified intervals in Table 5.1) by using an evolutionary algorithm [5] that includes two mechanisms: mutation and cross-over. It is applied on an initial population of 1 million parameter sets and evolves towards smaller g, generation after generation. Generally for such a large initial population, the best agreement does not change significantly after 10 generations. In a second step, the set of parameters found is further improved by using the downhill simplex minimization method (Nelder & Mead 1965). Due to the large number of parameter sets tested, this procedure ensures that the resulting parameter set does not correspond to a local minimum.

5.5.3 Calibration schemes

During the Gaia mission a number of parameters are expected to remain constant or vary only slightly for a particular CCD. These parameters are the number of trap species as well as the characteristic capture cross-section and release time constant for each trap species. These two last parameters are temperature dependent, however the temperature variation accross a single CCD is negligible. The trap density will vary as a function of the solar activity but if one considers short time scales (from several hours to a couple of days), this parameter can be assumed to remain constant. In principle, a CDM sufficiently accurate for describing the CTI effects at any signal level and for any kind of illumination history could be used with a single set of parameters for all the observations coming from one particular CCD, and thus the CDM parameters could be calibrated on a per CCD basis or at least on per CCD stitch block [6] basis. However as mentioned in Section 5.2, CDM uses a number of simplifications regarding in particular the modelling of the electron density distribution. For this particular reason, a calibration of the CDM parameters over a smaller set of observations selected on a per magnitude or a per illumination history basis is expected to give better results and should in principle remain practical for the Gaia data processing.
In Section 5.6 we present the resulting agreement between the CDM outcomes and the experimental data for four different calibration schemes of the CDM parameters; the CDM parameters are optimized to fit:

5. `http://watchmaker.uncommons.org/`

6. CCDs are manufactured using photo-lithographic masks. Due to practical manufacturing constraints, the masks are smaller than the image area of large format devices like the Gaia CCDs. As a result, a large format CCD consists in the assemblage of smaller 'CCD units' with slightly different parameters called stitch blocks.

1. a single target charge profile. This would correspond to a calibration of the CDM parameters for each observation of a particular star obtained under very similar conditions. Although impractical in the context of the Gaia data processing, this allows the assessment of the best agreement obtainable for this particular CDM candidate.
2. a set of target charge profiles with the same signal level but different CI delays. This corresponds to a calibration of the CDM parameters on a per stellar magnitude basis.
3. a set of target charge profiles with the same CI delay but different signal levels. This corresponds to a calibration of the CDM parameters on per CI delay basis.
4. the entire dataset including different signal levels and CI delays. This corresponds to a calibration of the CDM parameters on a per CCD basis.

Note that due to the nature of the experimental data, the disturbing effect of a star located in between the star of interest and the last CI cannot be accounted for in any of these calibration schemes. Chapter 4 showed that the effect of disturbing stars on the CTI calibration at the image processing stage is only significant for bright stars. This effect may be further investigated by, for instance, repeating the procedures described in this chapter on the 'sky-like' data acquired during RC3 and RC4 for which a mask with a pseudo-realistic sky-pattern was used.
In order to compare the agreements resulting from different calibration schemes at the level of individual predicted and target profiles, we make use of the reduced χ^2, χ^2_r:

$$\chi^2_r = \frac{\chi^2}{S - (3\,J + 2)} \tag{5.5}$$

where $3\,J + 2$ is the number of free parameters. χ^2_r is not used in the optimization of the CDM parameters, but only for the purpose of a final comparison between different calibration schemes.

5.6 Results

In the following we present the results of our comparison between the RC2 test data and the CDM predictions obtained using different calibration schemes. By establishing the performance of the tested CDM candidate, we are also able to identify some of its current limitations. Moreover we address for each calibration scheme two key issues for the Gaia data processing by characterizing the impact on the final agreement reached by the tested CDM candidate of: (i) the number of trap species included in a CDM simulation, and (ii) performing one-dimensional or two-dimensional CDM simulations.

5.6.1 The best agreement achievable

In this first step of our comparison, CDM is employed to reproduce each individual one-dimensional damaged profile from RC2; we thus proceed to an independent calibration of the CDM parameters for each available signal level and distance to the

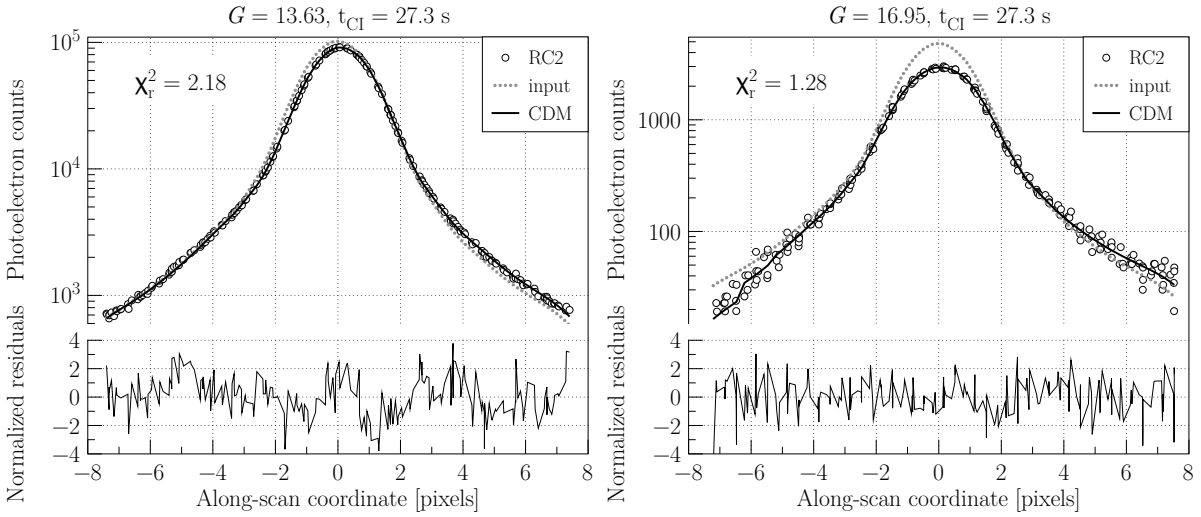


Figure 5.1 — Comparison between the RC2 test data (black circles) and the CDM predictions (black continuous lines) for a unique CI delay ($\sim$ 27 s) and two different signal levels: $G = 13.63$ (left) and $G =$16.95 (right). The grey dotted line represents the model input signal. While the top part of the figure allows for a direct and qualitative comparison between test data and CDM predictions, the bottom part shows the normalized residuals i.e. the difference between actual and predicted profiles divided by the photon noise. To quantify the agreement between model and data we also indicate the values of χ^2_r. The CDM prediction corresponds to the best fit to the shown RC2 charge profile. The best-fitting model parameters are presented in Table 5.2. The charge transfer in the CCD image area is from right to left (i.e. the leftmost samples are the first to be read out). The same convention is applied for all the figures representing charge profiles.

preceding CI (or time since last CI). In these conditions the best level of agreement between the CDM predictions and the test data can be obtained. Hence, in this way, we can assess the ultimate performance of the tested CDM candidate.

CDM can always achieve a very good fit to any specific damaged charge profile. Fig. 5.1 is an illustration of this capability to reproduce a particular damaged profile for two different signal levels (note that the ordinate scale is logarithmic). As illustrated in the bottom part of Fig. 5.1, no significant systematic variation of the normalized residuals can be observed across the profile, this means that the same level of agreement is achieved over the entire profile: peak, leading and trailing edges.

For these calibration conditions we were able to evaluate the number of trap species required to achieve the best agreement possible and wether the simulations must be performed in one dimension (1D) or two (2D). It is essential to address these two questions for the Gaia data processing and also in order to understand what level of realism is achieved by the tested CDM candidate. For each signal level and CI delay, the calibration procedure is performed as described in Section 5.5. And for each calibration instance, we use a different number of trap species ranging from 1 to 5. The obtained results are depicted in Fig. 5.2, which shows the resulting χ^2_r as a function of the number of trap species for a 1D and 2D simulations, for three different damaged profiles (two different signal levels and CI delays). The CDM parameters that realize the best achieved agreement are summarized in Table 5.2.

So far only one detailed study has been dedicated to the characterization of the trap

species present in an irradiated Gaia-like CCD. The results of this study are summarized in Hopkinson et al. (2005). They identified 7 different trap species with release time constants τ ranging from a few μs to several hundred seconds at the Gaia operational temperature. As already discussed, the CTI-induced distortion observed in the damaged charge profile from RC2 is mostly caused by the CTI occurring in the CCD image area. According to the SRH formalism, 67% of the traps release their charge after a time $t = \tau$. Most of the traps with a release time constant significantly shorter than the TDI period (or parallel transfer period), are thus expected to capture and release charges within the time of a single transfer event. As a consequence those traps have no net effect on the stellar image. Considering a Gaia CCD TDI period of 0.9892 ms and looking at Table 1 in Hopkinson et al. (2005), one can infer that 4 out of the 7 identified trap species will contribute to the CTI in the CCD image area. These 4 trap species have the following measured release time constants: 130 s, 2–4 s, 15 ms, and 0.8 ms.
Figure 5.2 clearly shows that 2 trap species are enough to obtain a good agreement between the model and the test data when the CDM parameters are optimized to a fit a single profile. One trap species may also be enough when focusing on shorter CI delays. This is less than the four trap species mentioned above, but this is in agreement with our expectations since an individual profile with a particular CI delay does not set a high constraint on the trap parameters and does not allow an efficient differentiation between trap species with similar effects on the stellar image:

1. The impact on the image profile of traps with a release time constant greater than the CI delay (i.e. $\tau >> t_{\rm CI}$) remains very little.
2. Traps with a release time constant smaller than the CI delay and greater than the temporal size of a window (here $P_{\rm TDI}\times$ 15 pixels) have very similar effects; they are empty prior to the star transit and capture electrons from the profile, but release only very few of them within the window. Two of the trap species mentioned above belong to that interval (for $t_{\rm CI} \sim 27$ s).
3. Traps with short release time constant compared to the window temporal size can capture electrons in the leading edge and release them within the profile, again two of the traps species from Hopkinson et al. (2005) study belong to that interval.

Figure 5.2 also shows that one can obtain the same level of agreement for CDM simulations performed in 1D or 2D. Apart from the necessary additional amount of time to perform one simulation in 2D, it also appears that in the 2D case the CDM calibration generally necessitates a greater amount of simulations to find a similar level of agreement as in the 1D case.
Finally, we investigated the impact of the size of the telemetry window on the best achievable agreement. Fig. 5.1 shows the best agreement obtained for the simulation of a 15 AL pixels window. For stars of magnitude 13.67 and 16.95, the actual telemetry windows for Gaia are composed of respectively 12 and 6 AL pixels. Performing the exact same calibration procedure using the proper size of the telemetry window we obtained a final agreement of $\chi^2_r = 3.78$ for the 13.67 magnitude (12 AL pixels and $t_{\rm CI} \sim 27$ s) and $\chi^2_r = 1.10$ for the 16.95 magnitude (6 AL pixels and $t_{\rm CI} \sim 27$ s). As a

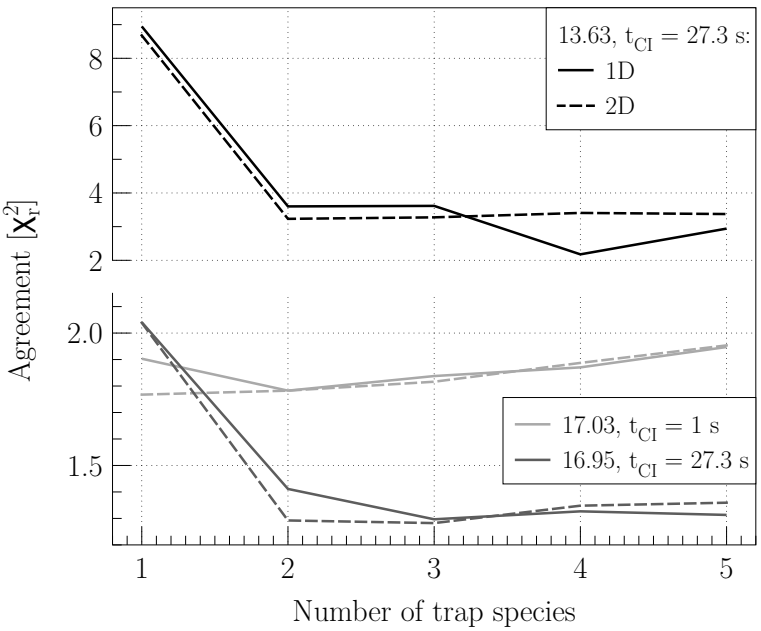


Figure 5.2 — Resulting agreement obtained after a calibration of the CDM parameters for a single charge profile using different number of trap species and performing the CDM simulation in 1 or 2D. The top part of the figure shows the level of agreement as a function of the number of trap species for a target charge profile for G =13.63 and a CI delay of ∼27 s, the bottom part shows the same results for two other target profiles with similar signal levels ($G = 17$) and different CI delays: 1 s (in light grey) and 27 s (in dark grey). The continuous lines depict the results obtained for 1D CDM simulations, and the dashed lines for 2D CDM simulations. While the best agreement possible can be obtained for CDM simulations performed with 2, 3 or 4 trap species depending on the target profile, generally only a minor improvement is achieved by introducing a third or a fourth trap species. Similarly one can notice that the same level of agreement between CDM predictions and test data is obtained for a CDM simulation performed in 1D or in 2D.

consequence the impact of the window size seems to be negligible, however this needs to be further investigated and confirmed.

5.6.2 Calibration on a per magnitude basis

Although a calibration of the CDM parameters performed for a particular damaged profile allow us to investigate the ultimate performance of a CDM candidate, it is unrealistic in the context of the Gaia data processing. We are thus now interested in studying the performance of CDM in more realistic conditions of calibration. In a calibration of the CDM parameter on a per magnitude basis, the resulting set of CDM parameters allows to describe the CTI effects for a particular magnitude (or similar signal levels) and different illumination histories. In this second step we thus fit simultaneously several damaged profiles having similar signal levels but different CI delays. As already mentioned, if a CI fills most of the traps in the CCD only traps with a release time constant similar to or shorter than the CI delay will be empty when the artificial star crosses the CCD. Therefore only these traps cause CTI effects, and with different delays different trap species are important. As a consequence this step also allows us to probe the occurence of trap species with very different release time constants.

Figure 5.3 illustrates the resulting agreement when fitting ten target profiles with G =15 and varying CI delays from 30 ms to 120 s. Fig. 5.3 (left) shows for a particular RC2 profile (G =15 and $t_{\rm CI}$ = 1 s) the difference between the best CDM prediction (see Section 5.6.1) and a CDM prediction based on parameters obtained from the per magnitude calibration. The agreement per profile is very good almost at the level of the best achievable although slightly worse especially in the region of the profile peak as can be seen from the normalized residuals (Fig. 5.3, left bottom). The average χ^2_r achieved for the ten profiles is 8.79.
A careful look at the normalized residuals for each profile that results from a per magnitude calibration (Fig. 5.3, right) shows a subtle systematic variation. It is not clear what causes this systematic variation of the residuals across the image location, yet we can propose some hypotheses. This could be the symptom of an intrinsic limitation of the tested CDM candidate. However such variation seems to be less severe when CDM is fitted to a single profile (Fig. 5.3 left) and absent in most cases (Fig. 5.1) meaning that this limitation is overcome by giving more flexibility to the model. This systematic variation across the image location may originate in the modelling of the CDM input signal (see Section 5.4.3). The construction of the input signal is based on the undamaged reference data accumulated in the non-irradiated part of the CCD. The illumination conditions could have been modified slightly between the acquisition of the CTI-free data and the damaged data. A careful study of the residuals between two CTI-free stellar images located at the top and the bottom of the non-irradiated region of the tested CCD showed that a variation of focus across the CCD produces a similar signature as the one we observed here (C. Crowley (ESA, ESTEC) and Astrium, priv. comm.).
Figure 5.4 shows the variation in mean agreement $\langle\chi^2_r\rangle$ as a function of the number of trap species included in the CDM simulations. As anticipated, to reproduce with a single set of parameters this large dataset that probes different temporal domains, CDM needs several trap species, with different release time constants that cover the CI delay range, and significant densities and cross sections (cf. Table 5.2 column 5). A reasonably good fit can be obtained with 2 or 3 trap species (respectively $\langle\chi^2_r\rangle = 11.39$ and 10.06), however a significant improvement can be obtained by including more trap species. As suggested by the experimental tests (cf. Section 5.6.1 and Hopkinson et al. (2005)), the best agreement is obtained for a number of 4 trap species ($\langle\chi^2_r\rangle = 8.79$). When repeating the exact same calibration procedure for a different magnitude and for the same range of CI delays, we obtain a different set of trap species. As will be further discussed in Section 5.7.1, this can partly be attributed to instabilities in the experimental conditions. Fig. 5.4 also shows that a better agreement can be obtained by performing 1D simulations instead of 2D.

5.6.3 Calibration on a per CCD basis

In this section we are interested in investigating the agreement between experimental data and CDM predictions that results from a calibration of the CDM parameters performed on a per CCD basis. In this case a single set of CDM parameters is used to describe the CTI effects occurring for all the observations acquired with a particular CCD i.e. for any CI delay and at any signal level. This would constitute the least

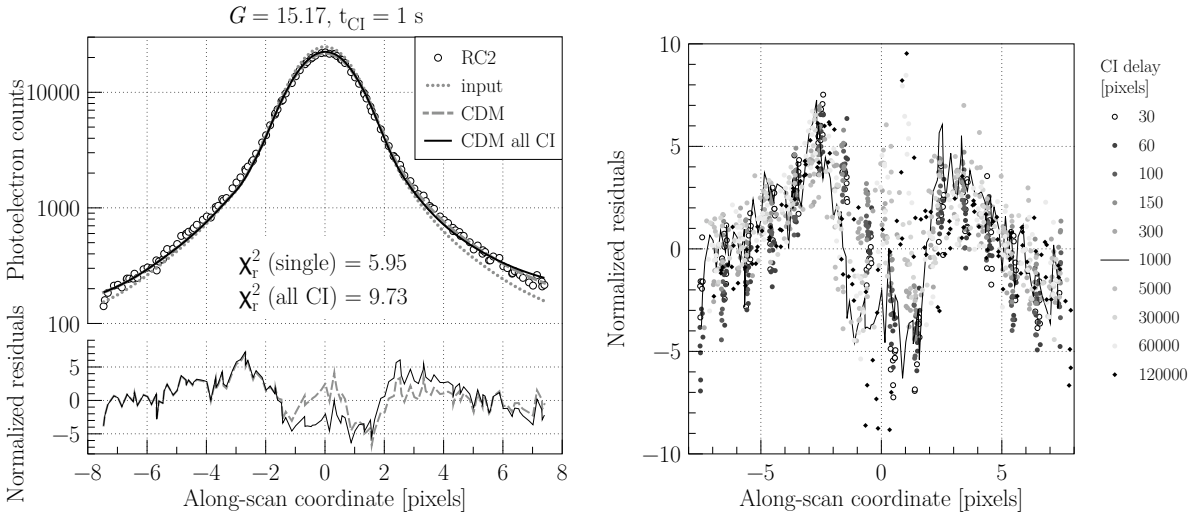


Figure 5.3 — Left: Comparison between the RC2 test data (black circles) and two CDM predictions for the same signal level ($G \sim 15$) and CI delay (~ 1 s). The dashed grey line is a CDM prediction for a parameter set (Table 5.2 column 4) obtained after an individual fit to the depicted RC2 profile, whereas the continuous black line is the CDM prediction for a parameter set (Table 5.2 column 5) obtained after a simultaneous fit to 10 damaged profiles with different CI delays including the one depicted. The dotted grey curve corresponds to the CDM input signal in both cases.
Right: The residuals normalized by the photon noise for each predicted damaged profile resulting from a per magnitude calibration of the CDM parameters.

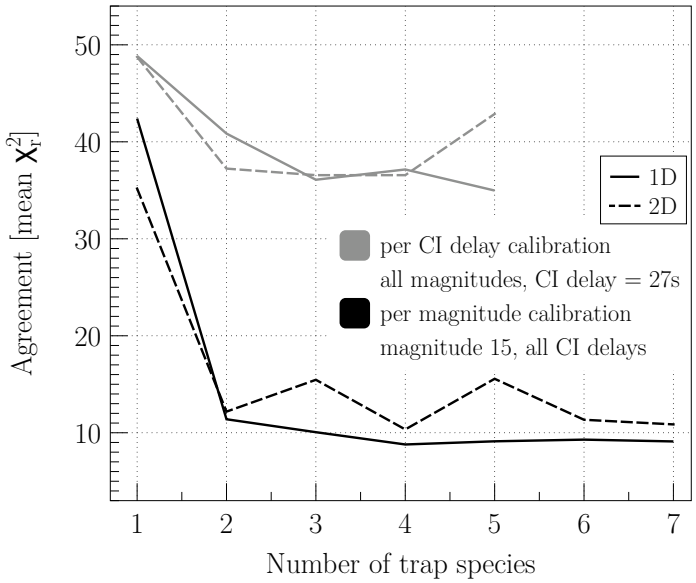


Figure 5.4 — Comparison between the mean agreement obtained after performing a per magnitude (black) and a per CI delay (grey) calibrations of the CDM parameters using different number of trap species and performing the CDM simulation in 1D (continuous line) or 2D (dashed line). In the per magnitude calibration case (black), the test data used to perform the calibration contains 10 damaged over-sampled profiles with $G = 15$ and different CI delays ranging from 30 ms to 120 s. In the per CI delay calibration case (grey), the test data used to perform the calibration contains 5 damaged over-sampled profiles for a CI delay of 27 s and for $G = 13.63, 15.29, 16.95, 18.65, 20.25$.

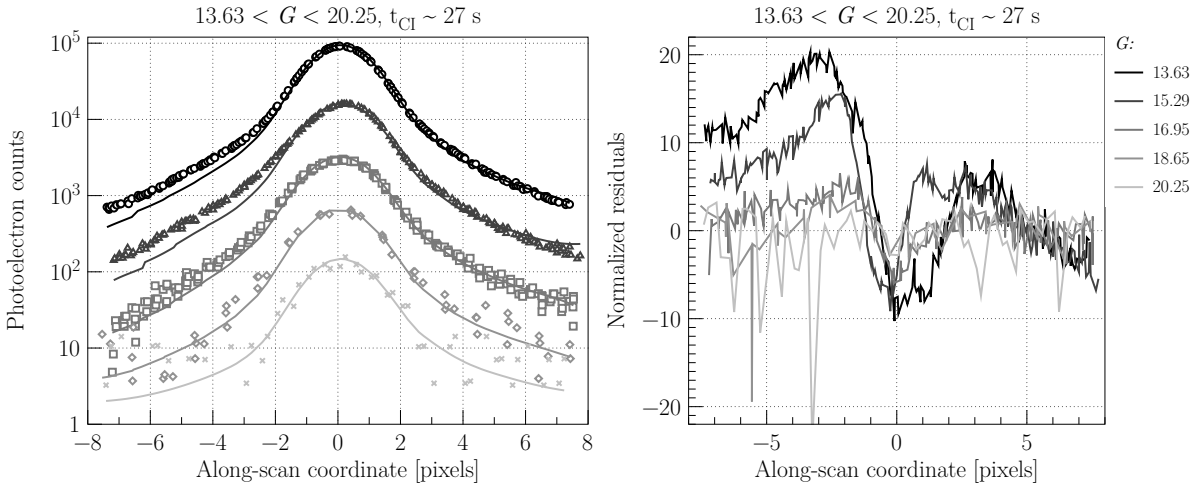


Figure 5.5 — Left: Comparison between the RC2 test data (symbols) and the CDM predictions (continuous lines) for a unique CI delay ($t_{\rm CI} \sim 27$ s) and several signal levels $G = 13.63$, 15.29, 16.95, 18.65, and 20.25. The CDM predictions result from a unique set of parameters (cf. Table 5.2 column 6) obtained by simultaneously fitting the five depicted RC2 damaged profiles.
Right: The residuals normalized by the photon noise for each predicted damaged profile resulting from a per CI delay calibration of the CDM parameters.

cumbersome calibration for the Gaia data processing.

In the previous section we have shown that CDM can fairly well reproduce a set of data including different CI delays at a particular signal level. In a per CCD calibration, CDM will have to deal with different signal levels. Hence we first need to explore how well CDM copes with different brightness levels independently from the illumination history. This is done by fitting with a single set of CDM parameters a set of damaged profiles comprising different magnitudes but similar CI delay. Fig. 5.5 shows a representative example of the level of agreement achieved for such calibration scheme. The overall agreement is clearly poorer than when the fitted dataset contains a unique signal level (Fig. 5.1 and 5.3). A systematic variation of the residuals (Fig. 5.5, right) across the image location similar to the one observed in Fig. 5.3 (right) may be distinguished. But the most prominent residual feature is the strong deviation occurring in the leading edge of the brightest stellar profiles (magnitudes 13.63 and 15.29).

As we shall further explain in Section 5.7.1, it is likely that a single set of trap parameters cannot describe the CTI effects for different signal levels in the context of the RC2 data due to changes in the experimental conditions between tests performed at different magnitudes. However this clear decrease in agreement (see also Fig. 5.4) can partly be imputed to the CDM limitations and in particular to the fact that CDM does not model explicitly the effect of the SBC on the growth of the electron cloud (see Sections 5.2 and 5.3). We indeed observe that reducing the range of signal levels included in our dataset of damaged profiles used in the calibration does improve greatly the resulting mean agreement. For instance, by including only the faintest magnitudes (16.95, 18.65, 20.25) that should experience a similar level of trapping, we obtained a mean agreement of $\langle \chi^2_r \rangle = 3.33$ that should be compared to $\langle \chi^2_r \rangle = 36.08$ when including the whole signal range available in RC2. We thus suggest a modification of CDM so

that the SBC is explicitly taken into account to ease the calibration process. As already discussed by Short et al. (2010), in order to model the effect of a SBC, one may consider to introduce a second parameter β for small charge packets. These two parameters can be constrained by performing experimental measurements on an irradiated CCD such as the fractional charge loss measurement (Chapter 2 and Hopkinson et al. 2005).
Finally we proceed to the calibration of the CDM parameters on a per CCD basis i.e. we perform a simultaneous fit of the CDM predictions to the entire dataset of over-sampled damaged profiles including different CI delays and different signal levels; the resulting χ^2_r for a particular profile is similar to the one resulting from a simultaneous fit to a dataset containing different signal levels only. The highest deviations are systematically obtained for the brightest magnitudes. And again one can obtain a significantly better agreement by removing the brightest damaged stellar images from the fitted dataset. By including only the faintest magnitudes, 16.95, 18.65, 20.25 and CI delays ranging from 30 ms to 30 s, we obtain a very good overall agreement that demonstrates the potential of the tested CDM candidate ($\langle\chi^2_r\rangle = 3.93$) in the context of a per CCD calibration of its parameters.

5.7 Discussion

In this section we interpret the results presented in Section 5.6 and discuss one of the main limits to the agreement between the CDM predictions and the test data: the experimental (Radiation Campaign test setup) instabilities (Section 5.7.1). In the Chapter 3 of this thesis we made use of synthetic data to study the impact of CTI on the image location accuracy. From this study, one can evaluate the level of agreement that any CDM must achieve to recover most of the CTI-induced bias in the image location estimation. Taking into account the experimental instabilities, we discuss the presented CDM performance in the context of these requirements (Section 5.7.2). Finally we discuss and give recommendations regarding the calibration of the CDM parameters during the Gaia mission (Section 5.7.3).

5.7.1 Experimental instabilities

We did not take into account the extra noise contribution induced by experimental instabilities neither in the calibration procedure nor in the comparison criteria, because its accurate measurement is a complicated task. However it is clear from the test data and our results that the experimental instabilities play an important role in limiting the agreement between the CDM predictions and the test data. This is particularly relevant as we aim at reproducing the test data at the noise level (i.e. χ^2_r = 1) and because such instabilities will not be present during the Gaia mission.
One can distinguish between two different types of instability or source of noise. The first is generated by the uncertainty related to the exact knowledge of the experimental conditions: operating temperature, illumination, background level. The CTI effects depend strongly on the temperature, the signal level, and the constant level of illumination or background level. The characteristic release time constant τ varies with the temperature. As a consequence, two tests carried out at different temperatures cannot be explained by CDM with the same value of τ even if the same traps are present in

the CCD. A slight variation in the level of background illumination affects the trap occupancy level prior to the transit of a star: an increase in the background level from 0.3 e^-pixel^{-1} to 5 e^-pixel^{-1} reduces the measured charge-loss from $\sim$ 30% to $\sim$ 10% at magnitude G =18 (Short et al. 2010; Brown 2009a). This variation thus directly affects the calibration of the background light level $S_{\rm DOB}$, and indirectly the trap density parameter for each trap species. To build over-sampled charge profiles we assumed that the experimental conditions remain exactly the same between different scans, i.e. repeated tests. Although this was not generally the case, we did observe in the test data, over-sampled profiles presenting fluctuations with an amplitude greater than the photon noise. Similarly when calibrating the CDM parameters using more than one over-sampled charge profile, we assumed (very) stable experimental conditions between tests carried out for different CI delays or for different signal levels.
The second source relates to the handling of the CCD in between tests carried out during a unique RC. The test data was acquired over a period of several months. A clear variation of the radiation damage effects during the course of a RC has been observed (Georges 2009). In particular, important discrepancies between similar tests carried out at different times have been explained by the fact that during the first three RC the tested CCDs were stored at room temperature in between tests (sometimes up to several days). This seems to affect the trap species characteristics.
Table 5.2 clearly shows that no particular trap species seems to be necessary to explain the data. Although the CDM approximations are not to be forgotten, experimental instabilities certainly play an important role in that matter. As illustrated in Fig. 5.3 (left), the agreement for a particular profile is almost comparable whether the CDM parameters are obtained by fitting a single over-sampled profile or a large set of profiles with the same brightness level (per magnitude calibration). If this noise contribution is one of the main limiting factors to this agreement, we can conclude that the experimental instabilities occurring between two tests with the same brightness and different CI delays have similar amplitudes as instabilities occurring between two scans within the same particular test. As we saw in Section 5.6.3, this agreement for a particular profile can however be considerably deteriorated when including profiles of different brightnesses in the calibration procedure. On top of the limited capabilities of CDM to reproduce CTI effects at different signal levels, this suggests that experimental instabilities between two tests carried out for different signal levels are the largest. This is corroborated by the fact that some of these tests were indeed performed on different days.

5.7.2 The CDM performance in the context of the Gaia approach to CTI mitigation

The Gaia data processing relies on a forward modelling approach to mitigate the CTI effects, in particular to recover the image location bias induced by the stellar image distortion. As a consequence the image location accuracy, and ultimately the Gaia final astrometric accuracy, depends on the capabilities of a CDM to reproduce observations affected by CTI (Chapters 3 and 4). In Section 5.6, we have evaluated the best agreement with the test data achievable by CDM and characterized the variation of this agreement as a function of the applied calibration scheme. Now it is important to evaluate if this agreement is good enough to recover most of the image location bias. To

do so we compare our results to those obtained in Chapter 3 using as test data a large set of synthetic damaged observations generated by a detailed Monte Carlo model of CTI effects (Chapter 2). In this study we applied the Gaia image location estimation procedure on a large set of CTI-free and damaged observations with and without applying the forward modelling CTI mitigation approach. Knowing the exact location of each observation, we are able to measure the residual image location bias after a CTI mitigation using CDM (Fig. 5.6). The original location bias (i.e. without any CTI mitigation) was also measured: for G =15 the bias is of the order of 0.2 pixels and 0.05 pixels for CI delays of respectively 27 and 1 s.

Figure 5.6 shows the agreement between the CDM predictions and synthetic test data as a function of magnitude and for two CI delays (right) and the residual bias when using the same CDM predictions to mitigate the CTI-induced bias (left). There is no clear correlation between agreement and image location bias. This is mainly because our comparison criterion allows a reasonably good agreement to be obtained between

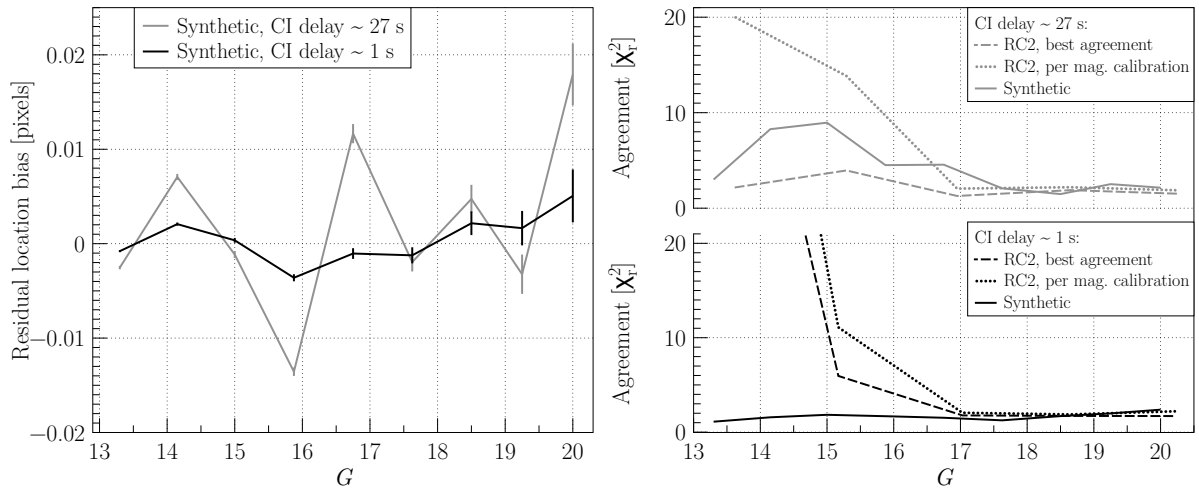


Figure 5.6 — Left: Figure reproduced from Chapter 3. Residual bias in the image location estimation when applying the Gaia image location estimation procedure on synthetic damaged observations including the forward modelling CTI mitigation approach using CDM. The synthetic damaged observations were generated using a detailed and physical Monte Carlo model of CTI effects that simulates the trapping at the trap level, and the charge transfer at the pixel-electrode level (Chapter 2). The residual bias is shown as a function of magnitude and for different levels of active trap densities in the simulated CCD. Active means that these traps are empty prior to the transit of a star. In grey the active trap density is 4 traps pixel^{-1}, in black 1 trap pixel^{-1}. The CTI effects amplitude (charge loss and image location bias) resulting from these particular active trap densities is representative of the CTI effects for a CCD irradiated at the level of 4×10^9 protons cm^{-2} (10 MeV equivalent) but for different CI delay: 27 s (4 traps pixel^{-1}) and 1 s (1 trap pixel^{-1}). To obtain these results the CDM parameters were calibrated to reproduce the damaged observations on a per magnitude basis using one-dimensional CDM simulations, the true instrument LSF was used as an input signal.
Right: Agreement as a function of magnitude between the CDM predictions and the damaged observations: synthetic (continuous line), and RC2 test data (dotted and dashed lines). This agreement is shown for two different CI delays: 1 s (black) and 27 s (grey). In the case of CDM agreement with the RC2 test data, we show the best agreement achieved (dashed line) corresponding to a calibration as described in Section 5.6.1 and the resulting agreement from a calibration on a per magnitude basis (dotted line) as described in Section 5.6.3.

a CDM prediction and a damaged observation, without a perfectly accurate reproduction of the CTI distortion or image shape. Yet the agreement shown guarantees a bias recovery of at least one order of magnitude. However the lack of correlation between the agreement and image location bias is an important indication of the current limits of our comparison criterion which is essentially photometric. We thus recommend the use of a criterion that would properly take into account the profile shape by introducing, for instance, the first derivative of the image profiles.
Figure 5.6 (right) also shows the agreements between the CDM predictions and the RC2 test data obtained for two different calibration schemes: (i) dashed lines, the CDM parameters are calibrated to reproduce a particular profile (Section 5.6.1), it corresponds to the best agreement achievable, (ii) dotted lines, the CDM parameters are calibrated on a per magnitude basis (Section 5.6.2), it corresponds to a realistic calibration scheme in the context of the Gaia data processing. We do not show the agreements resulting from a per CI delay and per CCD calibrations (Section 5.6.3), as the results obtained were not satisfactory. It is clear from Fig. 5.6 (right) that the tested CDM candidate has the capability to reproduce damaged observations at the required level of agreement that would enable a substantial, if not complete, recovery of the image location bias induced by CTI. This is especially true for the fainter magnitudes ($G > 16$) for which the agreement with the RC2 data is comparable to the one obtained with the synthetic data. For these magnitudes the agreement achieved using a realistic calibration scheme is similar to the best agreement possible. Nevertheless at brighter magnitudes, a realistic calibration scheme may not allow to reach a satisfactory agreement especially for a short CI delay (Fig. 5.6 right, bottom). Considering that the experimental noise is not taken into account in our agreement measurement, the results obtained for G =15 may still be acceptable. However for the brightest magnitude of our sample, G =13.6, there is a clear mismatch. This mismatch is present for a short CI delay even for the best agreement achievable. This is unexpected as the agreement obtained in the same calibration conditions and a longer CI delay (27 s) is very good (Fig. 5.6 right, grey dashed line, $\chi^2_r = 2.18$). This mismatch is thus present only when shorter CI delays are introduced in the dataset used in the calibration of the CDM parameters. In order to use the TDI mode only of the tested CDM candidate, we assume that a CI is filling all the traps present in the CCD. The number of traps filled by a CI depends on the CI level. The validity of such approximation breaks down if the CI level is lower than the signal level of interest. In our study, it is the case only for G =13 as for the RC2 data subset used in this study the CI level was 20 000e$^-$. As a consequence it is most likely that our approximation limits the performance of CDM in these particular conditions: short CI delays and signal level higher than the CI level. In any case this issue needs to be investigated further by finding out wether this limitation disappears when initializing the trap occupancy level with the image mode of CDM, prior to a CDM simulation in TDI mode.

5.7.3 Calibrating CDM

On top of its elaboration, the calibration of the chosen CDM is one of the main challenges faced by the Gaia data processing in the context of the CTI mitigation. In this study we have shown that, with the current version of CDM, a set of CDM parameters

per magnitude is necessary to obtain an acceptable reproduction of the damaged observations. In this per magnitude calibration scheme, four trap species must be included in the CDM simulation to obtain satisfactory results. Based on the results obtained for the faintest magnitude only, we also showed that a single set of CDM parameters for all the observations carried out with one CCD is conceivable, provided a modification is made of the way the current CDM is handling the electron cloud growth as a function of signal level.

Nevertheless it is not only the calibration approach that remains to be determined, but also the calibration method itself. Following the overall self calibrating approach chosen for the Gaia data processing, the image parameters, the instrument LSF, and the CDM parameters will be estimated from the same data in an iterative process (Chapter 3 and Prod'homme 2011). This means that such cumbersome calibration procedure as presented in this chapter (evolutionary algorithm probing the whole parameter space and then a fine tuning performed by the simplex downhill method) cannot be used in practice, at least not on a regular basis. It is thus expected that the CDM parameters will be estimated based on a maximum likelihood method, as for the image location estimation procedure. However the very degenerate nature of the CDM parameter space implies that the maximum likelihood method will have to be initialized with parameters already close to optimal. It is presently planned to use as a first guess the information that can be derived from the CI about the trap parameters. The trap density can be inferred from the charge loss occurring in the first pixels of the CI profile. Indications about the cross-section of the traps can be derived from the same measurement. And the characterization of the different trap species, as well as the determination of their release time constant can be achieved by analyzing the release of electrons after a CI. If the CDM parameter search is constrained to realistic values, a CDM candidate capable of very accurately reproducing the damaged observations must do so with a set of physically plausible parameter values.

In that respect the tested CDM candidate performs remarkably well. The interval for each of the free CDM parameters (cf. Table 5.1) was chosen to be physically plausible and also realistic regarding values available in the literature for these parameters. As can be seen from Table 5.2, good agreement with the test data were obtained with values well within those intervals. The diffuse optical background parameter, S_{DOB}, is generally found to be close to the one measured directly from the RC2 test data. Regarding the values of the β parameter, they were systematically found within the range [0, 0.3] this is in agreement with what is expected for a Gaia CCD (cf. Short et al. 2010). Nevertheless it is surprising that in some cases a better level of agreement can be obtained by performing one- rather than two-dimensional CDM simulations, since the CTI effects varies with the signal level and that the signal level vairation over a stellar profile in the AC direction is important. This may be partly attributed to the necessary assumptions required to build a two-dimensional input signal (Section 5.4.3). However similar levels of agreement for 1D or 2D CDM simulations were also found when using synthetic test data where there is no ambiguity regarding the input signal (Chapter 3). A more realistic modelling of the electron cloud growth with the signal level using two different values for β may also improve the results obtained for 2D simulations.

5.8 Conclusions and future work

The final astrometric accuracy of Gaia is conditioned on the capability of a CDM to reproduce observations affected by CTI. Using a detailed and rigorous comparison procedure, we established the current level of performance of the tested CDM candidate to reproduce experimental test data representative of the future Gaia observations. We showed that at low signal levels ($G > 15$) CDM is capable of reproducing damaged profiles accurately enough to enable a factor of ten recovery of the image location bias induced by CTI during the Gaia mission, provided that the CDM parameter calibration follows a certain scheme.

When the CDM parameters are calibrated to reproduce a single profile for a particular CI delay and signal level, the agreement between the CDM prediction and this profile is remarkable and this for the whole range of tested magnitudes ($13 < G < 20$) and CI delays (from 30 ms to 120 s). The worst agreement was obtained for the brightest magnitude and shortest CI delay. This may be explained by our assumption that a CI fills all the traps present in the CCD independently from its level. Such assumption can be avoided by a proper treatment of the effect of CIs on the trap occupancy level prior to the a stellar transit.

With a single set of parameters, CDM is capable of reproducing up to a certain level a set of experimental data including different CI delays and magnitudes. In the context of Gaia and considering a realistic calibration scheme, this level of agreement is only acceptable when the CDM parameters are calibrated on a per magnitude basis, that is to say if a different set of CDM parameters is used for each magnitude. In these calibration conditions, a CDM simulation can be performed in one or two dimensions and must include at least four different trap species.

Due to its way of handling the growth of the electron cloud as a function of signal level with a single parameter β, CDM is currently not able to reproduce the effect of the supplementary buried channel. This constitutes one of the main limitations of the current CDM when it is used to reproduce with a single set of parameters a set of experimental data that involves different brightness levels. It has to be noted that, in these calibration conditions, the instabilities in the experimental conditions also played an important role in limiting the best agreement achievable.

The CDM performance appears to be acceptable if the set of data that CDM aims to reproduce contains only signal levels for which the effect of the SBC is similar. This suggests that a calibration of the CDM parameters for signals levels of similar brightness on a per CCD basis remain a viable option for the Gaia data processing consortium.

The tested CDM can achieve the required level of agreement by using physical and realistic trap parameters. Hence we are confident that the CDM parameter calibration will be eased and benefit from a determination of the trap parameters using the CI onboard Gaia. However this remains to be studied in more detail, in particular by using the imaging mode of CDM to reproduce CIs and check that the trap parameters obtained are similar to the ones directly inferred from the CI. Currently the effects of the serial CTI on the image location estimation procedure and the CTI mitigation scheme remains unclear. The native serial CTI is expected to be the predominant source of image distortion at the beginning of the mission, when the CTI induced by radiation

damage is still mild. It is thus also important to evaluate the performance of CDM to reproduce stellar images affected by both serial and parallel CTI. Finally, in order to improve the image location bias recovery, we plan to investigate the use of a new comparison criterion in the CDM calibration procedure that would quantify not only the reproduction of the image flux but also of the image shape.

Acknowledgments

The authors would like to acknowledge EADS Astrium for kindly supplying the experimental test data necessary to accomplish the presented study and the participants of the Gaia Radiation Task Force for providing useful feedback. The work of TP and MW was funded by the European Community's sixth framework programme (FP6) through the Marie Curie research training network ELSA (European Leadership in Space Astrometry, MRTN-CT-2006-033481). AB acknowledges support by the Netherlands Research School for Astronomy (NOVA).

Calibration	per over-sampled profile				per magnitude	per CI delay
Section	5.6.1				5.6.2	5.6.3
Figure	5.1	5.1	-	5.3 (left)	5.3	5.5
Target data:						
Magnitude	13.67	16.96	17.03	15.17	~ 15	13 to 20
CI delay	27 s	27 s	1 s	1 s	30 ms to 120 s	27 s
Agreement $[\chi^2_r]$	2.18	1.28	1.77	5.95	8.79*	36.08*
CDM parameters:						
Simulation	1D	2D	2D	1D	1D	1D
β	0.21	0.28	0.23	0.36	0.15	$6.49\ 10^{-2}$
dob [$e^- s^{-1}$ sample^{-1}]	3.31	$9.28\ 10^{-4}$	46.53	240.0	24.54	1.43
ρ_1 [per sample]	2.37	9.09	21.87	2.35	4.35	5.33
σ_1 [m^2]	$1.48\ 10^{-23}$	$4.08\ 10^{-23}$	$4.67\ 10^{-22}$	$1.00\ 10^{-23}$	$4.71\ 10^{-23}$	$1.44\ 10^{-22}$
τ_1 [s]	$2.06\ 10^{-3}$	$4.47\ 10^{-3}$	$1.01\ 10^{-1}$	$2.58\ 10^{-3}$	$6.43\ 10^{-2}$	$9.64\ 10^{-2}$
ρ_2	7.52	$1.86\ 10^{-2}$	-	20.0	17.86	$7.37\ 10^{-2}$
σ_2	$1.23\ 10^{-23}$	$5.72\ 10^{-18}$	-	$3.69\ 10^{-23}$	$1.00\ 10^{-23}$	$1.73\ 10^{-18}$
τ_2	2.37	$3.09\ 10^{-1}$	-	0.60	0.91	$2.53\ 10^{-1}$
ρ_3	4.51	5.65	-	-	5.47	4.72
σ_3	$4.21\ 10^{-22}$	$9.40\ 10^{-22}$	-	-	$3.60\ 10^{-22}$	$2.42\ 10^{-18}$
τ_3	39.86	2.84	-	-	7.59	20.55
ρ_4	3.55	-	-	-	17.67	-
σ_4	$9.93\ 10^{-20}$	-	-	-	$9.28\ 10^{-22}$	-
τ_4	42.06	-	-	-	99.92	-

Table 5.2 — The different sets of CDM parameters that achieved the best agreement for the different calibration schemes presented.
* this value corresponds to the mean agreement $\langle \chi^2_r \rangle$ considering all the profiles used in this particular calibration scheme.

Chapter 6

Digging supplementary buried channels: Investigating the notch architecture within the CCD pixels on ESA's Gaia satellite

ESA's Gaia satellite has 106 CCD image sensors of a design that includes Supplementary Buried Channels (SBCs, otherwise known as 'notches') within the charge transfer channels or 'columns' to mitigate against the effects of radiation damage at low signal levels. The formally agreed CCD acceptance criterion in the Gaia contract between ESA/EADS Astrium and e2v only stipulated that SBCs should be present in Gaia CCDs but Kohley et al., in testing one Gaia Flight Model (FM), found that SBCs were present in only some of the pixels along every column. Missing SBCs will cause increased signal loss due to radiation damage and could thus degrade the predicted scientific performance of the mission. We re-analyse and model data from eight pre-2004 Gaia CCDs and find that one of these CCDs is missing some of its SBCs. Out of all the columns in these eight CCDs, less than 1% of columns are missing some SBCs. We obtain new measurements, re-analyse existing measurements and include a literature result for four post-2004 CCDs. Three of these have missing SBCs. Out of all the columns in these four CCDs, 68% of columns are missing some SBCs (all the FM CCDs are post-2004). The difference between these two samples is highly statistically significant and points to a change in e2v manufacturing of Gaia CCDs between 2003 and 2005. e2v predicts that all CCDs in the same batch should have the same SBC characteristics. By comparing the batch numbers of the three affected CCDs (three different batches) with those currently assigned to the Gaia satellite (it is too late to change them), we tentatively predict that a minimum of 17% of flight CCDs are likely to be affected by the SBC issue. In the absence of further testing, we predict that in the other 29 completely untested batches 69% of the CCDs may be affected (between 11 and 100% with a 99% confidence interval). Therefore it appears likely that the majority of Gaia's 106 CCDs have the same missing SBC issue as the two CCDs in Hubble's Advanced Camera for Surveys (ACS)/Wide Field Channel (WFC), albeit in only some of Gaia's pixels. We show that even with missing SBCs, Gaia astrometry still meets mission requirements. Simulations of whether Gaia photometry and spectroscopy still meet mission requirements in the absence of missing SBCs should be conducted. We recommend that SBC efficiency tests (described and used in this chapter) should be conducted: both pre-launch on-ground testing and post-launch in-flight testing.

G.M. Seabroke, T. Prod'homme, N.J. Murray, C. Crowley, G. Hopkinson, A.G.A. Brown, R. Kohley, A. Holland *To be submitted*

6.1 Introduction

ESA's Gaia satellite is a high-precision astrometric, photometric and spectroscopic ESA cornerstone mission, scheduled for launch in 2013, that will produce the most accurate stereoscopic map to date of the Milky Way. This will be achieved by measuring parallaxes, proper motions, radial velocities, and astrophysical parameters for one billion stars, one percent of the estimated stellar population in our Galaxy (Perryman et al. 2001). Figure 6 in Seabroke et al. (2008b) shows that Gaia observations will consist of charge packets within each CCD pixel ranging from one to the pixel full well capacity (FWC) of $190\,000$ electrons. This is due to Gaia's completeness range ($6 \leq V \leq 20$ mag) and the fixed exposure time for each CCD (4.4 s), due to operating in Time-Delayed Integration (TDI) mode in step with the satellite's spin rate. The vast majority of Gaia observations will be at the faint end of its magnitude range and correspondingly the vast majority of the charge packets that make up Gaia observations are expected to range from one to thousands of electrons. Gaia will operate in the radiation environment at L2 for at least five years. During the mission, radiation will generate 'traps' in the silicon which cause a loss of charge transfer efficiency (CTE) i.e. the fraction of the charge signal correctly transferred from pixel to pixel. The design of the Gaia CCDs includes Supplementary Buried Channels (SBCs) to confine charge signals to a small volume of silicon, thereby reducing the number of traps with which the signal can interact and maximising the CTE.

The formally agreed CCD acceptance criterion in the Gaia contract between ESA/EADS Astrium and e2v only stipulated that SBCs should be present in Gaia CCDs. This was demonstrated to be the majority case in CCDs built prior to 2004, Demonstrator Models (DMs, from the Technology Demonstration Activities (TDA) phase) and Engineering Models (EMs), by Hopkinson (2006) using the First Pixel Response (FPR) method, which can only be used on irradiated CCDs. Before manufacturing Flight Models (FMs), e2v changed their photo-lithographic mask set in 2004 but the set was meant to be identical to the one used to manufacture the DMs and EMs. FMs were not systematically tested for SBC efficiency. In an independent one-off test, Kohley et al. (2009) tested one close-reject FM and found SBCs are only present in some pixels along every CCD column. Note that this was not the reason for this particular FM not being selected for the actual mission. Also the method used here, pocket pumping, has the advantage of being able to be used on un-irradiated and irradiated CCDs.

Seabroke et al. (2010)'s 3D semi-conductor physics model of the Gaia pixel provided a manufacturing alignment explanation for why SBCs are only present in some pixels along a CCD column. An accumulation of nominal alignment errors, thought to be rare by e2v, can add up to the alignment error found in Seabroke et al. (2010) required to remove SBCs from some pixels along specific columns furthest from the readout register. However, Kohley et al. (2009) found SBCs are only present in some pixels along *every* CCD column. Given that the mask set changed in 2004, this chapter investigates whether there are systematic differences between the SBCs in pre- and post-2004 CCDs. The change in the mask set can only be considered as circumstantial evidence to explain any systematic changes between pre- and post-2004 CCDs because it would be very difficult to definitely prove the mask set was responsible due to a plethora of

other factors that go into manufacturing CCDs with complex pixel architectures.
The issue of missing or non-functional SBCs is not unique to Gaia. In spite of manufactured specifications and laboratory tests, SBCs in the two CCDs in Hubble's Advanced Camera for Surveys (ACS)/Wide Field Channel (WFC) are not behaving 'in-flight' as predicted. It was expected that below 100 electrons, Charge Transfer Inefficiency (CTI) in these detectors should have been mitigated. Instead CTI becomes worse with increasingly small charge packets (Anderson & Bedin 2010) suggesting that the SBCs are missing.
Similarly, missing SBCs would increase the effects of radiation damage on Gaia observations. This introduces a systematic bias in the measurements (e.g. image location on the CCD, radial velocities etc). As described in detail in Chapter 3 the Gaia data processing takes into account CTI effects via the forward modelling of the image distortion. This approach allows for the unbiased estimation of the stellar image parameters. However, the nature of Gaia's observations (TDI and windowed - only a small region around the image or spectrum is read out from the CCD and telemetered to the ground) means that radiation damage also induces signal loss. This signal loss is expected to degrade the predicted scientific performance of the mission. In Chapter 3, for CCDs with a fully functioning SBC, we predict a performance degradation of about 10% (over the CTI-free case). With missing SBCs this performance loss will be higher.
Gaia's CCDs have already been integrated on to the satellite and so it is too late to change which CCDs are selected to fly on Gaia. Nevertheless, the enormous amounts of data that Gaia will produce in its 5-year mission means it is important to establish, as much as possible before the mission starts, the potential impact on the mission of missing SBCs. This chapter addresses the following issues:

1. Whether SBCs were working properly in Gaia CCDs used to predict the mission performances and against which mitigation models are being developed and tested.
2. In the absence of testing whether FM CCDs have functional SBCs as a criterion for selecting which CCDs should fly on the Gaia satellite, is it possible to predict how many will have functional SBCs?
3. What is the impact of missing SBCs on the Gaia image location accuracy?

The chapter is organised as follows: the Gaia SBC is introduced in detail in Section 6.2. Section 6.3 reviews measurements of Gaia's SBC in the literature, addressing the first issue listed above. The next three main sections each present new results on whether SBCs are working or not in Gaia's CCDs: re-analysis of existing data (Sections 6.4 and 6.6) and new data obtained and analysed for this chapter (Section 6.5).
Section 6.4 re-analyses the Hopkinson (2006) data set, which includes SBC measurements of seven Gaia CCDs and identifies a probable new case of non-functional SBCs (Section 6.4.2). This case is subsequently modelled (Section 6.4.3) to investigate whether the measurements can be interpreted as evidence for non-functional SBCs, which is discussed in Section 6.4.5. The next main section (Section 6.5) presents new pocket pumping measurements and analysis of a Gaia CCD test structure. The final main section (Section 6.6) presents a re-intepretation of data obtained from irradiated devices

tested by EADS-Astrium, the Gaia main industrial partner, in terms of SBC efficiency, which addresses the first issue raised above.
Section 6.7.1 addresses the second issue raised above by tentatively predicting how many CCDs on the Gaia satellite will have fully functional SBCs. The third issue raised above is addressed by repeating the detailed characterization of the impact of CTI effects on the Gaia image location accuracy of Chapter 3, using synthetic data for which a non-functional SBC has been simulated. We conclude in Section 6.8 before presenting specific recommendations for future testing for SBC efficiency in Section 6.9, both pre-launch on-ground testing and post-launch in-flight testing.

6.2 The Gaia CCD pixel architecture

A modern CCD has an electrode structure with transfer channels or columns formed by implanted doping in the underlying silicon (see Fig. 6.1). The structure is termed Buried Channel (BC) because the design is such that the charge signals are collected and subsequently transferred within the silicon and away from the silicon/silicon dioxide interface states (the layer between the silicon and the electrodes). A pixel with only a BC and the potential distribution in the pixel in the across-channel direction are shown in Fig. 6.2. The stored charge generally occupies the whole width of the channel. However, by including a narrow region of additional channel doping (in Gaia's case at one side), as shown in Fig. 6.3, a small deeper potential well is formed and the volume occupied by a small charge signal can be minimised. The additional doping is

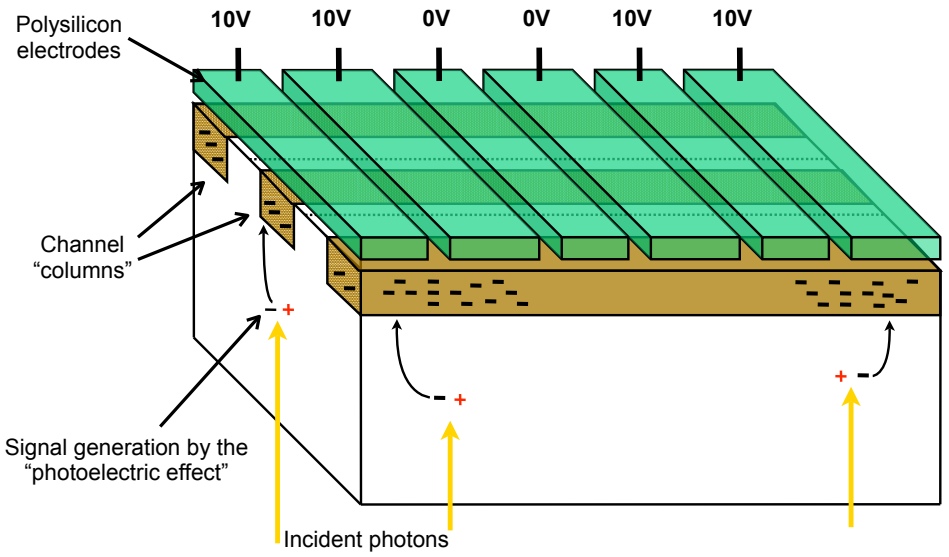


Figure 6.1 — Front-side schematic of a Gaia CCD showing the four-phase electrode structure. Gaia CCDs are back illuminated i.e. incident photons do not pass through the electrodes. Two consecutive electrodes are biased 'high' (10 V), to which photoelectrically generated signal electrons are attracted. Charge packets are separated in the charge transfer direction by biased 'low' (0 V) electrodes. The cube below the electrodes is a block of silicon and the darker regions within the silicon are implanted doping, which form channel columns. Charge packets are separated perpendicular to the charge transfer direction by column isolation regions of undoped silicon.

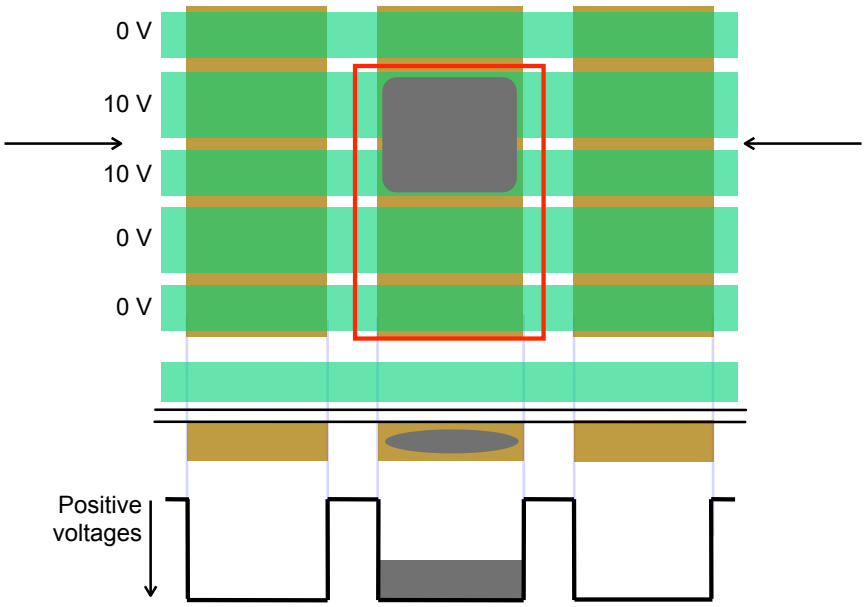


Figure 6.2 — Top: Front-side schematic of a CCD showing BCs with the vertical rectangles, electrodes with the horizontal rectangles, a charge packet in the rounded grey region with a pixel illustrated with a rectangular outline. **Bottom**: Cross-section through the arrows in the above schematic showing a single electrode above three BCs with the charge packet in the middle BC. A simplified potential distribution is below, showing that the charge packet sits within the potential maximum of a BC.

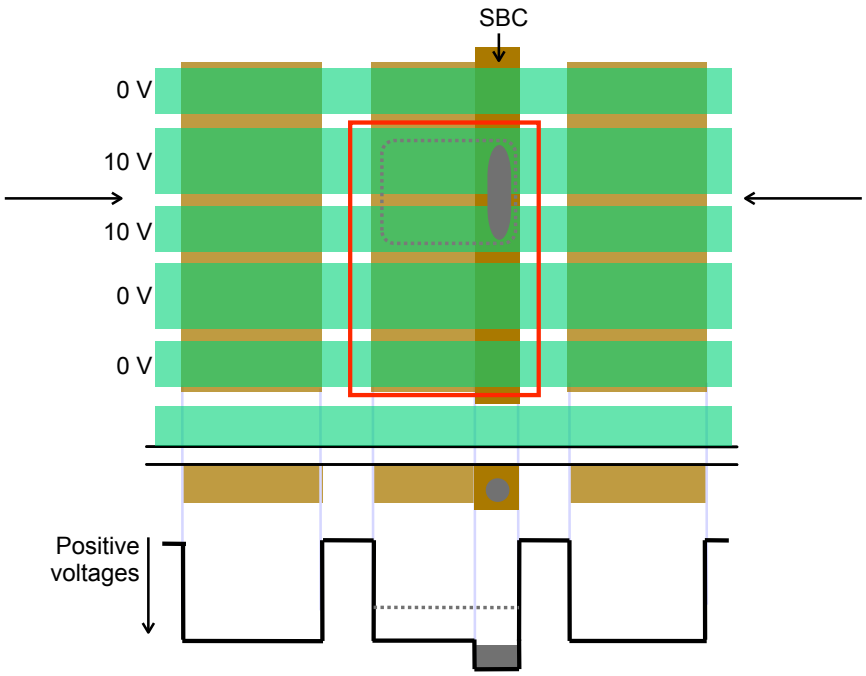


Figure 6.3 — Same as Fig. 6.3 but with SBC doping added to the central BC. The charge packet illustrated is smaller than the SBC capacity and so sits in the SBC volume, rather than the BC volume, (top) and sits in the higher SBC potential maximum, rather than the lower BC potential maximum.
Adapted from diagrams courtesy of D. Burt. (e2v technologies)

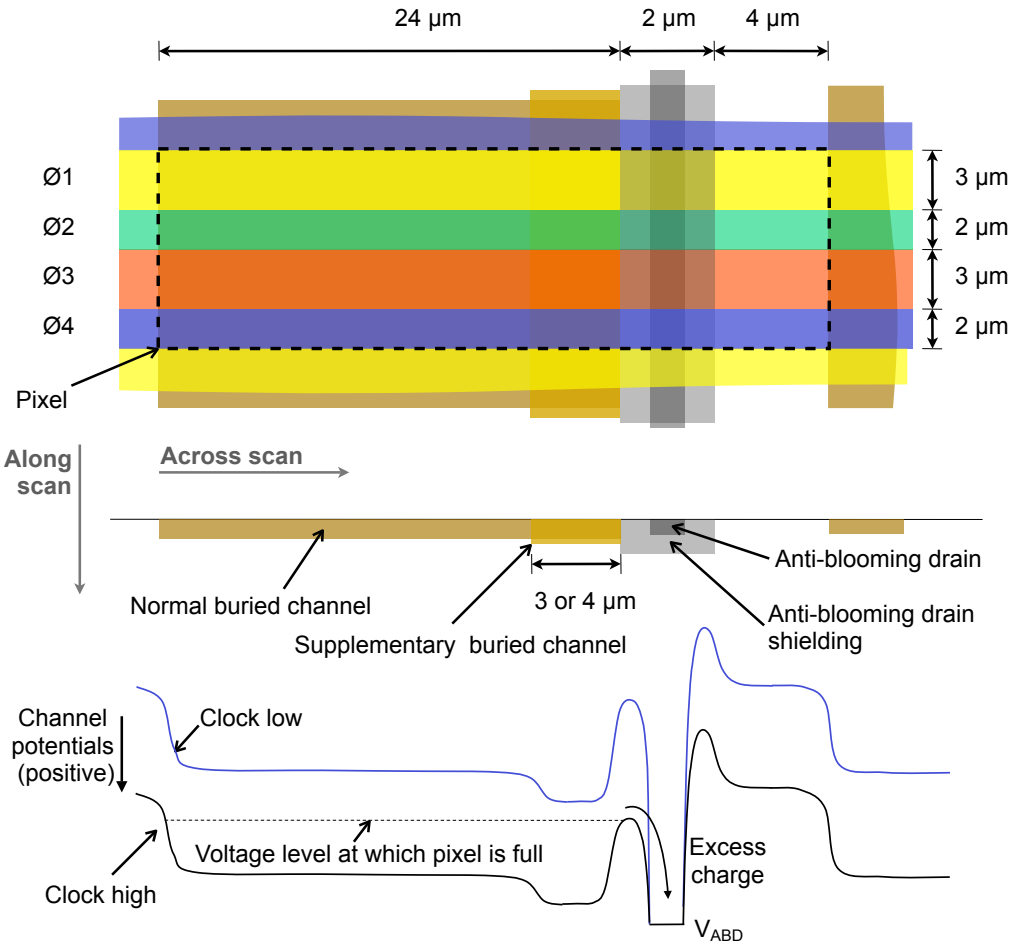


Figure 6.4 — Detailed schematic of the Gaia CCD pixel architecture. *Top*: Front-side schematic of a Gaia CCD pixel showing the four electrodes (ϕ) with pixel features labelled. *Middle*: Schematic of the vertical cross-section through the top schematic in the across-scan direction. *Bottom*: Channel potential profile in the same direction, resulting from different voltages being applied to the electrode (clock low $= 0$ V, clock high $= 10$ V). Diagram courtesy of D. Burt (e2v technologies)

termed a SBC or 'notch'. Once the capacity of the SBC is reached, any further charge spreads out to cover the whole BC volume, as shown dotted in Fig. 6.3.

The CCD91-72 was designed and manufactured by e2v especially for Gaia. Gaia has one of most complex pixel architectures built for astronomy (see Fig. 6.4). It includes the standard BC, the relatively rare SBC and, unusually, an Anti-Blooming Drain (ABD). Gaia is the only astronomical detector with an ABD. It prevents charge bleeding down columns from bright observations, allowing simultaneous faint observations. The ABD also removes excess charge from just upstream of TDI gates. These gates located at different positions within each CCD block charge packets of very bright observations that would otherwise saturate the pixels. This allows their integration to begin just downstream of the TDI gate at shorter distances from the readout register.

The doping that defines each pixel feature is implanted into the CCD silicon using a photo-lithographic mask. Each feature (e.g. BC or SBC) has its own mask. Due to practical manufacturing constraints, the masks are smaller than the image area of

large format devices like the Gaia CCD (4.5 cm $\times$ 5.9 cm). As only one mask can be used at a time, the mask area dictates the area of the CCD that can be implanted with dopant at any one time. This smaller sub-array is called a stitch block (de Bruijne 2008, see Fig. 6.5). Large format devices like the Gaia CCD are fabricated using a photo-lithographic 'step and repeat' process that implants a particular pixel feature simultaneously everywhere in a stitch block, one stitch block at a time. Both Gaia CCD's upper and lower halves are photo-lithographically stitched from 7 repeated stitch blocks and 2 end-termination stitch blocks. A Gaia CCD is thus composed of 18 different stitched blocks. Each termination stitch block contains 108 columns while each repeated stitch block section contains 250 columns: $2 \times 108 + 7 \times 250 = 1966$ columns in the ACross scan (AC) direction and 4500 TDI lines in the ALong scan (AL) direction. The AL direction refers to the travel direction of stars across the Gaia CCDs that is induced by the continuous spinning of the spacecraft around its own axis. And AC refers to the orthogonal direction to this travel direction.

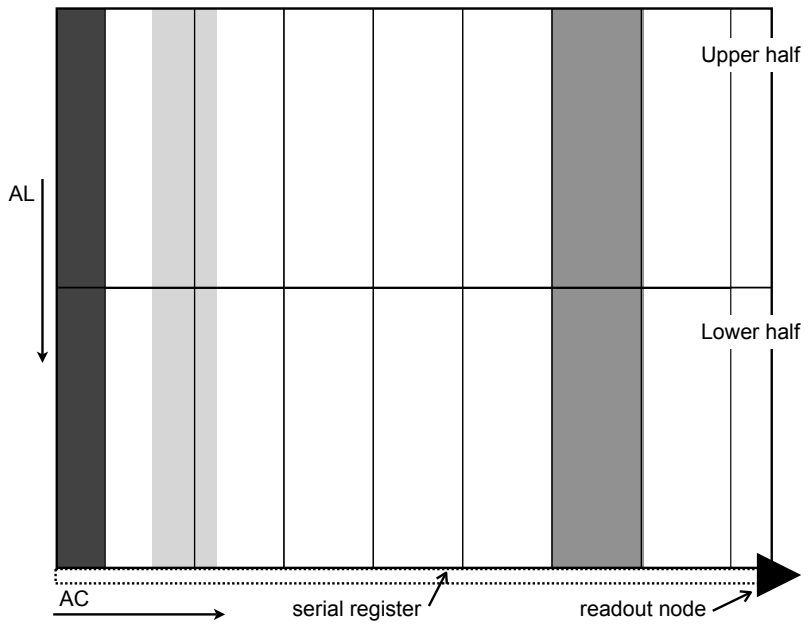


Figure 6.5 — Schematic of the Gaia CCD image area: CCD columns (TDI line ALong scan, AL) are vertical; ACross scan (AC) is horizontal; the readout node is at the bottom right corner. Each rectangle is a stitch block. All the internal lines in the schematic are stitch boundaries. The horizontal stitch boundary is at TDI line 2160 from the serial register (bottom box) and at TDI line 2340 from the start of the CCD (top line in the schematic). We colour-coded three different areas, reference to these areas will be made in the text, from left to right: (i) the dark grey (rightmost) area covers columns 359-608, corresponding to two stitch blocks in the AL direction (hereafter known as AL stitch block couples), (ii) the medium grey (leftmost) area covers columns 1859-1966, which corresponds to the AL stitch block couple identified as a candidate for missing SBCs in the Gaia CCD EM 03153-20-01. (iii) the light grey (central) area covers columns 1550-1755, which corresponds to the area over which the 05256-17-02 CCD was irradiated.

6.3 The SBC Full Well Capacity issue

All the Gaia CCDs (DMs, EMs and FMs) were built from the same nominal pixel architecture design, including SBCs. Therefore measuring the SBC FWC in as many of these devices as possible constrains the expected manufacturing spread of SBC FWCs. In the following we give a summary of the different measurements that have been carried out on irradiated CCDs (Section 6.3.1) and non-irradiated devices (Section 6.3.2). As a result this section is also a summary of the different methods that can be used to measure the SBC FWC and test the functionality of a SBC.

6.3.1 Measurements on irradiated CCDs

First Pixel Response:
The first published estimate of Gaia's SBC FWC was estimated from First Pixel Response (FPR, see Section 6.4.1 for details) measurements made on a Gaia EM CCD (Hopkinson et al. 2005). In their Fig. 11 Hopkinson et al. (2005) show FPR data from columns close to the readout node (although the column numbers are not specified). In Fig. 6.6 we reproduce the same results but only including measurements taken in columns 359-608 (the dark grey area in Fig. 6.5). This range of columns corresponds to two stitch blocks in the AL direction (hereafter known as AL stitch block couples). Hopkinson et al. (2005) estimated the SBC FWC to be $\sim$1400 e^-. However, this is a typographical error in their paper. The arrow defining their SBC FWC in their Fig. 11 is pointing to $\sim$1400 ADU, which corresponds to $\sim$1100 e^-. Fig. 6.6 shows that this signal level corresponds to the first inflexion point. It is at this signal level that the number of electrons in the SBC start to collapse the SBC potential so some of these electrons spill out of the SBC potential into the BC potential. No longer protected in the smaller volume of the SBC, these few electrons in the BC meet more traps, causing the fractional charge loss increase. At the second inflexion point, the signal level has completely collapsed the SBC potential into the BC potential so here the SBC no longer exists and all the electrons sit in the BC potential. Fig. 6.6 shows that without SBCs in each column, the BC fractional charge loss would continue to increase as signal level decreases. In other words, the presence of SBCs causes the fractional charge loss curve to shift left to smaller signal levels so that at a given signal level the fractional charge loss in the SBCs is less than in the BCs.
On its own, Fig. 6.6 only demonstrates that the stitch blocks from which the data was obtained have working SBCs. Because of the way e2v manufacture the CCDs (all the SBC doping implanted simultaneously, one stitch block at a time), intra-stitch block SBC FWCs should be the same. However, in the Gaia case, there is no guarantee that inter-stitch block SBC FWCs will be the same. This is due to the SBC doping abutting the Anti-Blooming Drain (ABD) doping (see Fig. 6.4 top schematic) and overlapping the ABD shielding doping (see Fig. 6.4 middle schematic). The ABD doping is the first pixel feature to be implanted into each pixel within a stitch block area of CCD silicon using its own photo-lithographic mask (the size of a stitch block). The ABD mask is aligned to the 'zero grid'. All the SBCs in the same stitch block have their doping implanted subsequently using a SBC photo-lithographic mask aligned to the same zero grid. Each mask alignment to the zero grid is subject to random alignment

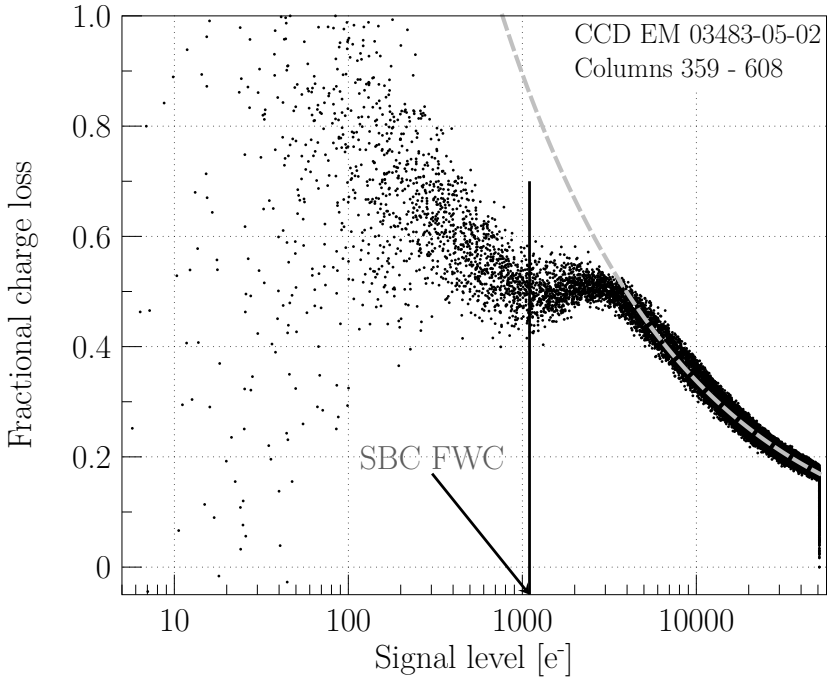


Figure 6.6 — First Pixel Response (FPR, see Section 6.4.1 for details) measurements made on a Gaia EM CCD: fractional charge lost due to radiation damage as a function of signal. In this case, the SBC FWC is evaluated to be ∼1100 e^- (black vertical line) following the definition of Hopkinson et al. (2005). In grey (dashed line), a power law is fitted to the fractional charge loss in the BC regime (signal level > 3500 e^-). The extrapolation of this fit in the SBC regime (signal level < 3500 e^-) shows the expected fractional charge loss for a CCD containing no SBCs.

(stitch) errors of $\leq$0.25 μm (Burt 2005a). Because the ABD shielding doping cancels out the SBC doping, the effective SBC doping width, consequent potential depth and resulting FWC is sensitive to the stitch error. Therefore, scatter in inter-stitch block SBC FWC (both intra-CCD and inter-CCD) is expected from the nominal e2v design of the Gaia CCDs.

The second published measurement of Gaia's SBC FWC can be read off from Fig. 5 in Chapter 2, which shows the fractional charge loss as a function of signal level for a different Gaia CCD, irradiated and tested by the prime contractor for Gaia EADS-Astrium in their second Radiation Campaign (RC2). This chapter reproduces the results of Chapter 2 in Fig. 6.10 (right). Following the SBC FWC definition of Hopkinson et al. (2005), the read off value is ~ 500 electrons. Note that the CCDs referred to in Hopkinson et al. (2005) and Chapter 2 were irradiated at different doses (respectively 1×10^{10} and 4×10^9 protons cm^{-2} 10 MeV equivalence), as seen by the different amplitudes of the fractional charge loss in their figures. Yet both CCDs exhibit the clear break in the increase of the fractional charge loss as the signal level decreases. This can only be explained by the presence of SBCs. While Hopkinson et al. (2005) fractional charge loss values are derived from different levels of uniform charge injections (CIs), Chapter 2 values are derived from one non-uniform CI. This means that the different data points in Chapter 2 come from different AL stitch block couples. The small scatter in their results indicates that the tested CCD has similar SBC FWCs in each of its AL

stitch block couples. This is important to establish as this CCD was used to establish the efficiency of CI as a CTI mitigation tool. RC2 data was since been used to further investigate CI and for verification of CTI modelling Chapter 2. Therefore all results based on this RC2 CCD are representative of the scenario where all Gaia CCD stitch blocks have working SBCs.
In the CCD upper half (cf. Fig. 6.5) the nominal SBC doping width is 3 μm, whereas in the CCD lower half the nominal SBC doping width is 4 μm (Burt 2003). This is because every CCD column and so every SBC crosses the horizontal stitch boundary. The SBC doping width is increased immediately after the stitch boundary to minimise the possibility of the SBC narrowing at the boundary as it is known that boundaries can cause the formation of a potential pocket, which acts like radiation traps, capturing and then releasing electrons and thus increasing CTI. This difference between the nominal SBC doping widths in the upper and lower halves of each Gaia CCD translates to different SBC FWCs in either half. All FPR data passes through both halves of the CCD so the fractional charge loss is the total over the entire column. Charge packets larger than the SBC FWC of the upper half and larger than the SBC FWC of the lower half will traverse the entire column in the BC. Charge packets smaller than the SBC FWC of the upper half will also be smaller than the SBC FWC of the lower half and so will traverse the entire column within the SBC. Charge packets larger than the SBC FWC of the upper half but smaller than the SBC FWC of the lower half will traverse the upper half of the column in the BC and the lower half within the SBC. In this case, the fractional charge loss will be greater than if the SBC doping width was 4 μm in both halves but less than the fractional charge loss if the SBC doping width was 3 μm in both halves. Therefore, the first inflexion point in Fig. 6.6 represents an average SBC FWC over the two halves of the CCD.

6.3.2 Measurements on non-irradiated CCDs

Pocket Pumping:
Unlike the FPR technique, pocket pumping can provide an independent measure of the SBC FWC in each half of a CCD. Charges from a flat field illumination are moved back and forth in the image area over a one pixel length. In pixels with traps, single electrons can be statistically trapped during one half-cycle and then released during the next half-cycle into the adjacent pixel. This produces bright-dark pairs around the mean flat field level at the trap position. Repeating the technique with increasing mean flat field signal levels identifies the number of traps per pixel as a function of signal size. This method was used by Kohley et al. (2009) to determine the number of fabrication-induced single electron traps in each pixel of a non-irradiated Gaia FM CCD. Their Fig. 9 shows that the SBC FWC (averaged over all the pixels in a stitch block) in each stitch block in the upper half of that CCD was measured to be less than 40 e^-. For a SBC FWC < 40 e^-, the depth of the SBC potential is very likely to be of the order of the thermal voltage, making thermal electron diffusion out of the SBC into the BC highly probable. In this case case the SBC is in practice absent. The SBC FWC (averaged over all the pixels in a stitch block) in each stitch block in the lower half of that CCD was measured to range from 1000 to 3000 e^-. Their Fig. 9 shows that they used the first inflexion point definition of SBC FWC. Kohley et al. (2009) is the first

published evidence that there is an issue with the Gaia SBCs. This chapter re-analyses the Hopkinson (2006) data set, which includes the Hopkinson et al. (2005) Gaia CCD as well as six others, to see how common working SBCs are in the upper halves of Gaia CCDs.

e2v's analytical Gaia CCD design predictions (published for the first time in Seabroke et al. 2010) are actually an order of magnitude higher than the SBC FWCs reviewed in this section: 7900 e^- in the upper half and $13\,000$ e^- in the lower half. Seabroke et al. (2010) provided an explanation for the discrepancy between these predicted and measured FWCs: none of the Gaia CCDs measured so far seem to have the nominal SBC doping widths. This explanation was investigated by simulating smaller SBC doping widths in a 3D semi-conductor physics model of the Gaia pixel. Seabroke et al. (2010) Fig. 5 shows that the simulated SBC FWCs agree with the Kohley et al. (2009) measured FWCs when the simulated SBC doping widths are $\leq$1.5 and 2 μm for the upper and lower halves of the CCD respectively. The resulting potential profiles in the $\leq$1.5 μm simulation showed no evidence for a SBC, only a BC. These simulated SBC doping widths are systematically 2 μm smaller than the nominal widths in both CCD halves. Worst-case random alignment stitch errors (-0.5 μm) applied to the nominal SBC doping widths cannot explain the offsets. The range of SBC FWCs in the lower half of the CCD measured by Kohley et al. (2009) agree with the simulated FWCs taking into account the worst-case random stitch error between two pixel features within a stitch block being $\pm$0.25 μm. Therefore these offsets appear to be uncalibrated systematic offsets in e2v photo-lithography, which could either be due to systematic stitch offsets or lateral ABD shield doping diffusion. The Seabroke et al. (2010) pixel model does not specifically simulate either of these scenarios, rather it simulates the same effective SBC doping width that could be produced by both or either of these scenarios.

Minimum Injection Method:

Figs. 6.7 and 6.8 illustrate different CI methods available on the Gaia CCDs: (top) the Voltage-Tunable Method (VTM) and (bottom) the Minimum Injection Method (MIM). The horizontal lines under the Injection Diode (ID), Injection Gate (IG) and electrodes (Iϕ1, Iϕ2 and Iϕ3) are representative of the voltage applied to each. Like in Figs. 6.2, 6.3 and 6.4, the convention of plotting low voltages levels as upper horizontal lines and high voltage levels as lower horizontal lines allows the voltage levels to also represent potential levels with electrons filling them up (grey regions) from the bottom (higher voltage) to the top (lower voltage), analogous to water filling up a bath. The top schematic in Fig. 6.7 shows the IG with a voltage level lower than ID and Iϕ1 so electrons (like bathwater) cannot flow from ID to Iϕ1. The ID fills up with electrons and the middle left schematic shows that when the level of electrons is higher than the IG, they can flow to Iϕ1. The bottom schematic in Fig. 6.7 shows the ID no longer filling up with electrons. This stops the flow of electrons across the IG, leaving electrons now occupying both the ID and Iϕ1. The voltages that are tunable in the VTM are IG and Iϕ1 (see the equation in Fig. 6.7).

The idea behind MIM is to transfer charge into Iϕ1 without the amount of charge stored in Iϕ1 depending on the voltages applied to IG or Iϕ1. This can be achieved by having

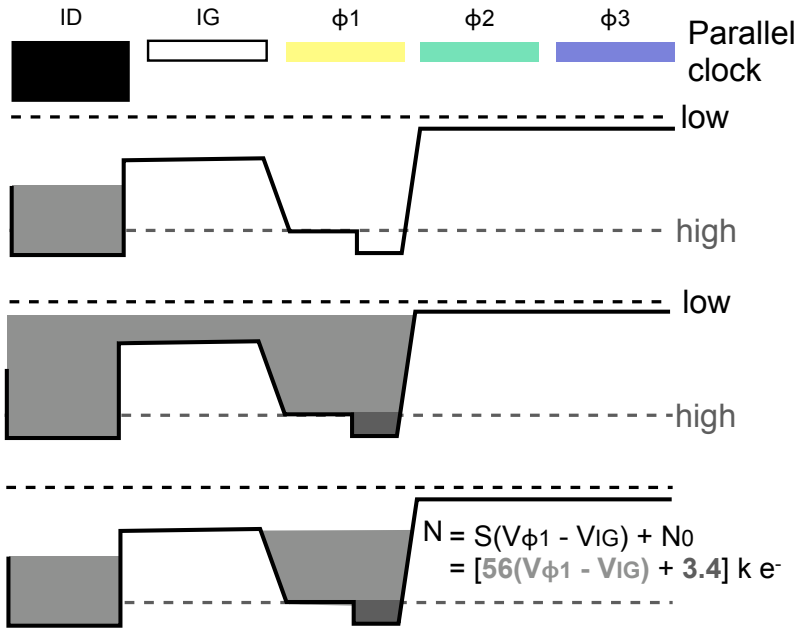


Figure 6.7 — Voltage-tunable method, where the quantity of charge depends on the difference of voltage between the 'high' level of the image section clock Iϕ1 and that on IG (Injection Gate). Note that the width of ID (Injection Diode) and IG are only schematic and so actually differ from one another. Also only 3 out of the 4 electrodes contained in a Gaia CCD pixel are depicted. Figure adapted from Burt (2003) Fig. 3 (i) and labelled with the latest e2v predicted levels of injected charge from Burt (2005b).

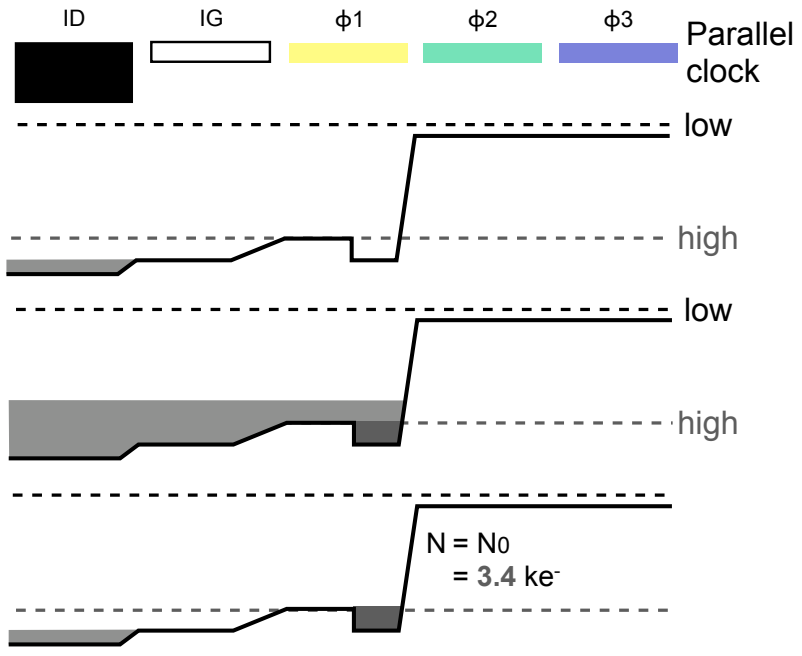


Figure 6.8 — Minimum injection method, where the quantity of charge is independent of the difference of voltage between the 'high' level of the image section clock Iϕ1 and that on IG and only depends on the SBC capacity under Iϕ1 (N_0). Note that we did not depict the case for which Iϕ2 is also biased high, such that electrons are also injected under Iϕ2.

ID > IG > Iϕ1 (see top schematic in Fig. 6.8). When the ID fills ups with electrons they can flow across the IG into Iϕ1 (see middle schematic in Fig. 6.8). As with the VTM, when the ID no longer fills up with electrons, the flow of electrons across the IG is stopped. As the electrons are attracted to the highest voltage, they drain back across IG to ID. However, because there is a SBC under Iϕ1 and its voltage is higher than Iϕ1, the electrons filling the SBC remain there and only the excess electrons drain out of Iϕ1. In theory MIM enables the precise injection of a very low level of charge independently of the voltage settings. The SBC FWC should be more reproducible column-to-column than using only the difference of voltage between the 'high' level of Iϕ1 and that on IG to inject the same small amount of charge. Without a SBC under Iϕ1, the equations in Fig 6.7 show that to inject 3400 e^- requires 0.06 V. Variations in oxide charge from pixel to pixel cause small fixed voltage offsets between columns at this level, which means that a single applied voltage to a single gate results in a non-uniform distribution of potential column-to-column and thus a non-uniform distribution of injected charge column-to-column. e2v expected MIM to have higher CI uniformity than VTM but this assumes only a small scatter in SBC FWCs. As explained in Section 6.2, SBC FWC should be uniform within a stitch block but can vary between stitch blocks. Therefore, MIM may only have higher CI uniformity than VTM within a stitch block but not necessarily over an entire CCD.
The original layout (pre-2004) of the CI structure included Iϕ1 with the nominal AL dimension of 3 μm and an underlying SBC foreshortened in AL by 1 μm (i.e. 3 $\times$ 2 μm). Burt (2005b) reports that e2v MIM testing of this structure did not find N_0: the electron capacity under Iϕ1 (see Fig. 6.8). This was because the SBC potential was not there as its electrical size was smaller than the geometrical size due to fringing fields and manufacturing tolerances. Therefore, the CI structure was modified such that Iϕ1's AL dimension was extended to 5 μm and the SBC increased in size to 3 $\times$ 3 μm to reduce fringing fields and increase the SBC FWC under Iϕ1 only. However, MIM testing of this new CI structure on a few devices still did not find N_0 (T. Eaton, e2v, private communication). This suggests that the post-2004 CCDs tested by e2v do not have working SBCs in their upper halves like the post-2004 CCD measured by Kohley et al. (2009). MIM testing was first suggested by Holland & Smith (2004) to investigate continuous CI. Now, MIM testing is an option for all Gaia FMs during in-flight commissioning to determine whether the FMs have working SBCs in their upper halves.

6.4 FPR measurements: Model to Data comparison

In this section we apply the First Pixel Response technique to laboratory measurements collected from irradiated Gaia CCDs. We then analyze the obtained results by using a detailed Monte Carlo simulation. This analysis is subsequently used to assess what fraction of the CCDs (or AL stitch block couples) are affected by the SBC issue.

6.4.1 Principles of the FPR measurement

Before explaining the details of the First Pixel Response measurements we first briefly review the physical process of electron trapping due to radiation damage. Non-ionising

displacement damage occurs when sufficiently energetic protons knock silicon atoms out of their lattice positions across the entire CCD. The resulting vacancies can move around the lattice until they combine with other atoms, e.g. phosphorous or oxygen, or combine with other vacancies in the BC and SBC to form electron traps, with different discrete energy levels between the valence and conduction bands.
Following the Shockley-Read-Hall formalism (Shockley & Read 1952; Hall 1952) that describes the capture and release of a charge by a trap as decay process, one can derive the capture and release probabilities (for a detailed derivation see Chapter 2). The probability that an empty trap captures an electron within a time interval t is:

$$p_{\mathrm{c}}(t) = 1 - e^{-r_{\mathrm{c}}t}\,, \tag{6.1}$$

$$r_{\mathrm{c}} = \frac{1}{\tau_{\mathrm{c}}} = \sigma_{\mathrm{t}} v_{\mathrm{th}} n_{\mathrm{e}}\,, \tag{6.2}$$

where r_{c} is the capture rate, τ_{c} the capture time constant, σ_{t} the capture cross section, v_{th} the electron thermal velocity, and n_{e} the electron density at the trap location. As a consequence the charge loss that is driven by the capture of charges is particularly sensitive to the electron density. The electron density distribution is shaped by the pixel architecture: as already mentioned, for high signal levels the electron packets sit in the BC and for low signal level in the SBC. The SBC confines the charge packet to a smaller volume and thus affects the electron distribution and the charge loss. Hence studying the charge loss and in particular the fractional charge loss allows us to probe the electron density and characterize the SBC.
The FPR measurement consists of the analysis of the charge loss (induced by trapping) that occurs in the first pixels of a well-characterized signal after its transfer across the full CCD image area. By well-characterized signal, we mean a signal for which the shape and the number of charges is known independently from the CTI effects.
As already explained, on top of the 4 electrodes required to transfer the charges from one pixel to its neighbour, in the first row only, the Gaia CCD pixels comprise a diode to generate artificial charges and a gate to control the number of artificial charges injected i.e. effectively transferred across the CCD (see Figs. 6.7 and 6.8). During the Gaia mission, CIs will be performed periodically (every ~ 1 s) by blocks of 4 to 20 lines to fill a large fraction of the traps prior to the stellar transits and thus mitigate the CTI effects.
CIs are particularly suitable to perform a FPR measurement. In a CI block of several tens of lines, the first line undergoes the most damage by encountering a certain fraction of the total amount of active traps (i.e. the empty traps) present in the signal confinement volume. As we shall see this fraction depends on the CI level (the number of electrons per pixel in the CI block) and the electron density distribution within a pixel. Depending on the capture cross-section of the trap species present and the clocking rate, the first CI line may not fill all the encountered traps. In this case the second CI line will also experience charge loss and subsequently for the other lines of the CI block. After a certain number of lines, no significant trapping can be measured. Using the last lines of a CI block one can measure a reference CI level, N_{ref}, and thus compute the charge loss Δ (i.e. the total number of trapped charges), and the fractional charge

loss δ:

$$\Delta = \sum_{i=0}^{S-1} N_{\text{ref}} - N_i \,, \tag{6.3}$$

$$\delta = \frac{\Delta}{S\, N_{\text{ref}}} \,, \tag{6.4}$$

where N_i is the number of e^- pixel^{-1} in the i^{th} line of the CI block. The traps filled by the CI with a release time constant shorter than the CI block duration will release electrons within the CI block and can thus bias the measurement of the CI reference level. To minimize this source of uncertainty only the first lines of the CI block should be considered in the charge loss measurement. The charge loss measured in the first lines of a CI block depends on the state of the traps at the time they are encountered by the CI signal. The trap state is set by the CCD illumination history. During a FPR experiment, it is thus important to have a good control (or at least a good knowledge) of the CCD illumination history. In the FPR experiments carried out by Hopkinson et al. (2005) to characterize the effect of the SBC, two CI blocks were performed with 100 lines of delay between them. The first CI block was performed to reset the illumination history, and only the second block was used to measure the charge loss. The level of the first block was fixed to $30\,000$ e^-, while the level of the second block was varied. In this way, one can study the charge loss variation as a function of signal level which enables the characterization of the notch architecture within a CCD pixel column. Although information from the CI will be available during the mission, it is not foreseen (at least during operational phase) to vary the CI level. This means that only very limited FPR information will be available (the minimum injection method discussed in Section 6.3.2 will likely only be used during the Gaia commissioning phase, if at all).

6.4.2 Selection of the experimental data

The Sira electro-optics tests were carried out on seven Gaia Astrometric Field (AF) DM and EM CCDs (see Table 6.1). Two of them were front-illuminated CCDs and the rest back-illuminated. This was the first set of CCD radiation test data obtained by an industrial partner in the Gaia project. The analyses described in Hopkinson et al. (2005) and Hopkinson (2006) focused on the determination of trap parameters and CCD characterisation. Figures 8-2-2 and 8-2-3 of Hopkinson (2006) show FPR data from different columns from different AL stitch block couples but not all groups of columns exhibit the characteristic SBC bump in the fractional charge loss. Here we test the hypothesis that these groups of columns consist of columns from the same stitch block AL couple and search for examples which do not exhibit the characteristic SBC bump in the fractional charge loss. These are candidates for CCDs with missing SBCs. Figure 6.9 shows that only two CCDs out of the seven tested by Sira electro-optics have AL stitch block couples that are candidates for missing SBCs and these only number three AL stitch block couples. It was decided to model the AL stitch block couples that had FPR data extending to the smallest signal levels (EM 03153-20-01 columns 1859–1966, see Fig. 6.5 medium grey area) as this provides more modelling constraints.

Table 6.1 — Summary of the irradiated AF CCDs tested by Sira. The first five digits of the CCD serial number form the batch number. The digits in the batch number refer to when the front-side processing occurred (explained in Section 6.7.1). The first two digits in the batch number are the year e.g. 03 refers to 2003. The second two digits in the batch number are the week of the year e.g. 15 out of 52. The fifth digit in the batch number is the number of batches in that week e.g. 3rd batch of week 15 in 2003. The middle set of two digits in the serial number gives the wafer number. The other set of two digits refers to the position of the CCD within the wafer.

Serial Number	Model (Extra information)
03153-05-02	DM (front illuminated)
03153-07-01	DM (front illuminated)
03153-16-02	DM
03153-20-01	DM
03442-11-01	EM
03483-05-02	EM
03483-06-02	EM

6.4.3 Generation of the synthetic data

In order to reproduce the FPR measurements, we simulate the transfer of a CI block over the image area of a virtual irradiated CCD using the detailed Monte Carlo model of charge transfer presented in Chapter 2[1]. This model was verified against experimental data acquired with Gaia irradiated CCDs. In particular, by using a flexible and analytical representation of the electron density distribution, the model is able to accurately reproduce the measurement of fractional charge loss as a function of signal level in the presence of an operational SBC (in both CCD halves). In the following we first describe the used electron density distribution model (Section 6.4.3.1). Then we present the simulation setup (Section 6.4.3.2) and the comparison procedure (Section 6.4.3.3) we used to reproduce the selected experimental data.

6.4.3.1 Modelling the electron density distribution

In the CTI effects model we used, electron packets are transferred at the pixel-electrode level and the trapping processes are simulated by computing the probabilities of the electron capture and release at the level of each individual trap. This necessitates the modelling of the electron density distribution n_{e} at the location $\mathbf{x} = (x, y, z)$ of each trap in the pixel signal confinement volume V. The latter corresponds to the volume occupied by an electron packet at the pixel FWC ($\sim 190\,000\ \mathrm{e}^-$ for a Gaia CCD). It is assumed to be a box for which the dimensions are set by the manufacturing characteristics of the CCD.

The electron density distribution is analytically described by a normalized Gaussian function in the three space directions. To take into account the effect of the SBC on the electron density distribution at low signal level, the distinction between two signal regimes is made and a smooth transition between these regimes is ensured:

$$n_{\mathrm{e}}(x, y, z) = S \times \rho(x, y, z) = S \times \rho(\mathbf{x})\,, \tag{6.5}$$

1. This model can be used as part of the Java package CEMGA available online `http://www.strw.leidenuniv.nl/prodhomme/cemga.php`

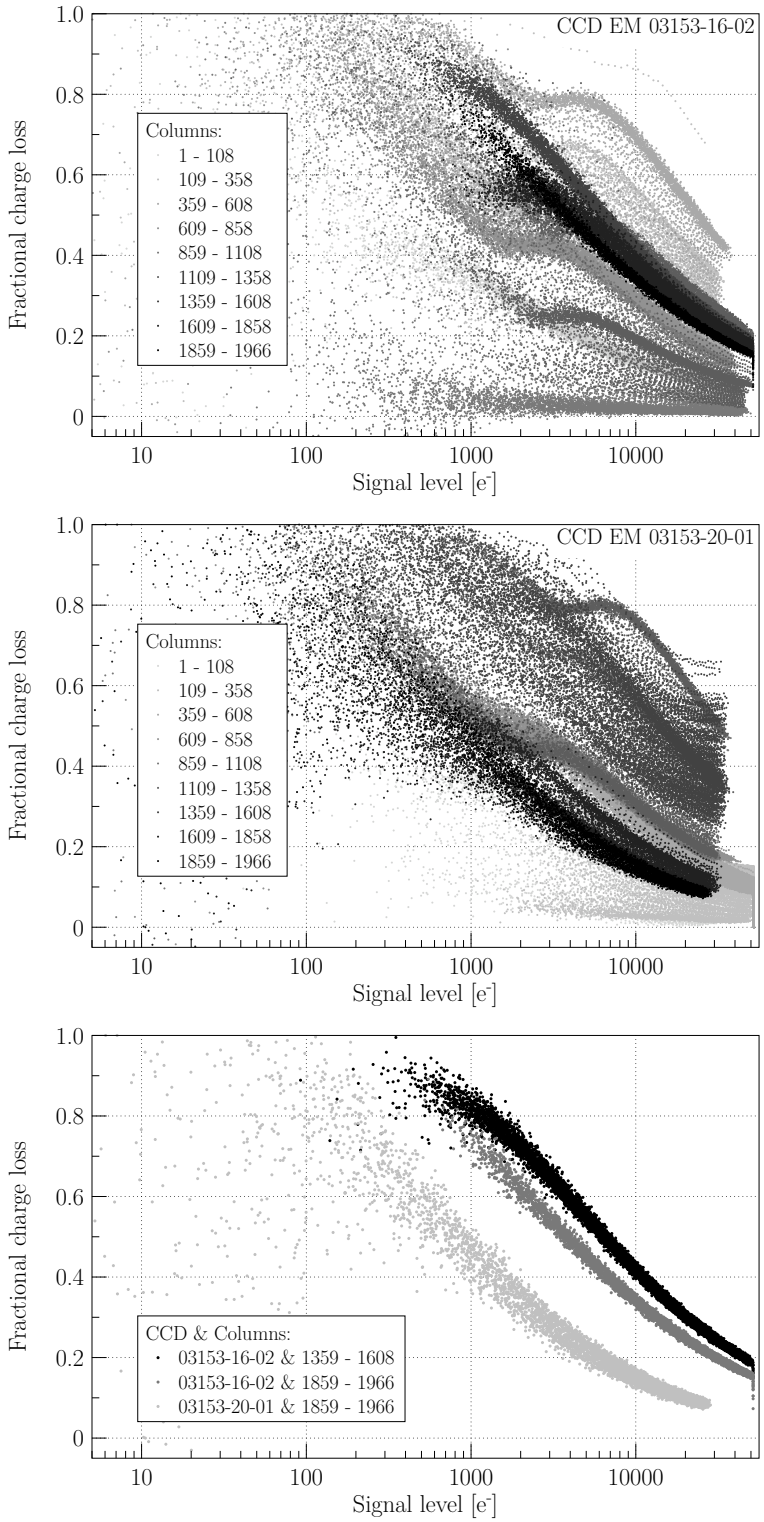


Figure 6.9 — Top and middle: FPR plots of the two CCDs out of the seven including AL stitch block couples that are candidates for missing SBCs. **Bottom:** FPR plots of the three AL stitch block couples that are candidates for missing SBCs.

$$\rho(\mathbf{x}) = \frac{\exp\left[-\frac{1}{2}\left(\left(\frac{x-x_0}{\sigma_x}\right)^2 + \left(\frac{y-y_0^\star}{\sigma_y^\star}\right)^2 + \left(\frac{z-z_0^\star}{\sigma_z^\star}\right)^2\right)\right]}{\left(\sqrt{2\pi}\right)^3 \sigma_x \sigma_y^\star \sigma_z^\star}\,, \tag{6.6}$$

where S is the signal level, and x, y, z respectively the CCD parallel and serial transfer direction, and depth. $\sigma_{\mathbf{x}}$ and $\mathbf{x}_0$ are the distribution standard width and centre coordinates along these directions. The parameters indicated with a $\star$ vary as a function of signal level:

$$P^\star = P_{\text{BC}}\left(1 - e^{-\frac{S}{S_{\text{BC}}}}\right) + P_{\text{SBC}} e^{-\frac{S}{S_{\text{SBC}}}}\,, \tag{6.7}$$

where P_{BC} refers to the parameter value in the BC and P_{SBC} to the corresponding value in the SBC. As already mentioned the SBC confines the electron packet in depth (z) and in the serial transfer direction (y), but the packet remains spread under two electrodes in the parallel transfer direction (x) independently of the signal regime. Thus, in the model, the SBC and BC signal confinement volumes have different dimensions along y and z ($y^\star_{\text{max}}$, $z^\star_{\text{max}}$), but the same size along x (x_{max}). Accordingly the centre and the standard widths of the density distribution along y and z are different for the BC and SBC regimes. For each regime, the distribution standard width in one direction is set to a fraction η of the signal confinement volume dimension in that direction, e.g.:

$$\sigma_y^\star = \eta^\star \times y^\star_{\text{max}}\,. \tag{6.8}$$

In this way the ratio between the dimensions of the signal confinement box is preserved in the electron density distribution. The greater is η the more spread is the electron distribution. η is thus later referred to as the electron density distribution spread factor. Saturation occurs at high signal close to FWC_{BC} and the modelling of this process is described in Chapter 2. The transition between the two signal regimes occurs at S_{SBC}. It is important to note that although in reality S_{SBC} would be compared to the FWC_{SBC}, in the presented model, due to the arbitrary nature of the chosen description for the transition between the BC and SBC regimes, S_{SBC} does not carry any real physical meaning.

6.4.3.2 Simulation setup

During the experiments carried out by Hopkinson, two CI blocks were transferred. The first CI block was used to reset the illumination history by filling a large fraction of the empty traps. The second block was used to perform the FPR measurement and its level varied. For most of the tested levels, the level of the first block ($\sim 30\,000$ e^- pixel^{-1}) was higher than the level of the second. As a consequence the population of traps probed by the second CI block was the same as for the first. However from this population the only traps capable of capturing electrons from the second block, are by definition the empty traps: the ones that were able to release their electron before the crossing of the second block. In our simulated experiment we can thus simulate the transfer of a unique CI block (of 50 lines) crossing the image area of a CCD containing only empty traps and perform the FPR measurement on this very same CI block. As for this particular experiment we are interested in the electron capture only, we used a unique

trap species with a common capture cross section ($\sigma_t = 5\times10^{-15}$ cm^2) and a long release time constant compared to the duration of the CI (50 $\times$ 0.9892 ms) such that the charge loss and the reference level measurements are not biased by a significant release of electrons. The simulated CCD image area consists in a single column of 4500 pixels. We did not simulate the transfer across the serial register, as one can to first order ignore the effects of the serial CTI. In order to investigate the potential difference of SBC implant between the first and second halves of the CCD, the standard widths of the electron distribution in the SBC regime and S_{SBC}, the SBC to BC regime transition signal, can be set to different values for the two CCD halves:

$$\eta_1^\star \leq \eta_2^\star\,, \tag{6.9}$$

$$S_{SBC_1} \leq S_{SBC_2}\,. \tag{6.10}$$

The BC and SBC signal confinement volumes as well as the standard widths of the electron density distribution in the BC regime remain the same for the two CCD halves. Note that the parameter label 1 refers to the CCD upper half i.e. the furthest away from the serial register. Once the CI block is transferred through the two CCD halves, the FPR measurement is performed in the same way as for the experimental test (Section 6.4.1). For each set of simulation parameters, the CI level is varied from 50 to 40 000 e$^-$ (with 200, 500, 2 000, 5 000, 20 000 e$^-$ as intermediate values).

6.4.3.3 Model to data comparison

In our simulations, there are 7 free parameters:

1. η_{BC}, η_1 and η_2 the electron density distribution spread factors, that set the standard widths of the electron distribution for each signal regime and CCD half according to the signal confinement volume dimensions,
2. S_{SBC_1} and S_{SBC_2} the signal level at which the transitions between the SBC and BC regimes is performed,
3. N_t the number of traps per pixel.

Table 6.2 details all the simulation parameters and indicates the values we used for the fixed parameters as well as the allowed intervals for the free parameters.
We are not interested in the exact reproduction of the data but rather in the study of the parameters of the electron density distribution model that lead to a reasonable agreement between the simulated and the experimental fractional charge loss measurements over the studied range of CI signal levels. We thus generate random sets of parameters in order to probe homogeneously the entire parameter space. Each parameter set results in a set of data points that sample the simulated fractional charge loss curve. These data points are then compared to the experimental measurements. To proceed to this comparison, we have first generated an analytical representation of the experimental data by the mean of a fit using spline functions. We use the analytical fit values F and the simulation results ϕ at each particular signal level n to compute the χ^2, that constitutes our comparison criterion:

Parameter	Description	Fixed value or [interval]
T	temperature	163 K
P_{TDI}	TDI period	0.9828 ms
N_t	number of traps per pixel	[0–20]
σ_t	capture cross section	5×10^{-20} m^2
τ_r	release time constant	9 s
	BC regime	
η_{BC}	distribution spread factor	[0.05–1]
FWC	BC full well capacity	190 000 e$^-$
	BC signal confinement volume size:	
x_{max}	in the parallel transfer direction	11 μm
y_{max}	in the serial transfer direction	24 μm
z_{max}	in depth	0.75 μm
	SBC regime	
	distribution spread factor:	
η_1	in the CCD upper half	[0.05–1]
η_2	in the CCD lower half	[0.05–1]
	SBC to BC regime transition signal:	
S_{SBC_1}	in the CCD upper half	[1–8000] e$^-$
S_{SBC_2}	in the CCD lower half	[1–8000] e$^-$
	SBC signal confinement volume size:	
$y_{SBC,max}$	in the serial transfer direction	2 μm
$z_{SBC,max}$	in depth	0.1 μm

Table 6.2 — Simulation parameters. Note that we use a non-nominal width for the SBC doping profile in the AC direction: $y_{SBC,max}$ = 2 μm. This value is given by Seabroke et al. (2010) to explain discrepancies between measured and predicted SBC FWC for functional SBCs (see Section 6.3.2).

$$\chi^2 = \sum_{n=0}^{N-1} \frac{\left(\phi\left(n\right) - F\left(n\right)\right)^2}{\sigma^2}\,, \tag{6.11}$$

where N is the total number of simulated CI levels and σ the noise. We use the formal errors associated to the analytical fit to the experimental data as σ values and thus assume that the formal errors encompass the experimental noise and the readout noise.

6.4.4 Comparison results

Figure 6.10 (left) shows the fractional charge loss measurement carried out by Hopkinson (black dots) on the Gaia AF CCD EM 03153-20-01 (columns 1859-1966) along with the result of the presented simulation setup that aimed to reproduce this measurement (red line). The simulation shown is representative of the best fit to the data that can be obtained following a random search in the parameter space. One can first notice a clear deviation at high signal levels ($> 10\,000$ e$^-$). Note that this deviation at high signal levels was already present in an attempt to fit the fractional charge loss

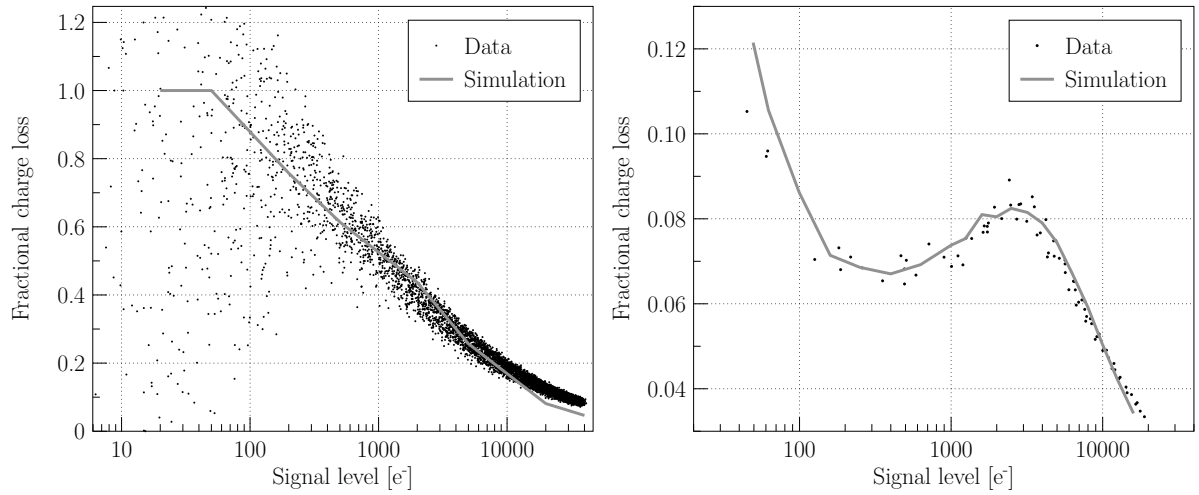


Figure 6.10 — Comparison between experimental (black dots) and simulated (grey line) fractional charge loss measurement as a function of CI level. The left panel shows a subset of Hopkinson's measurements for which the SBC of the CCD upper half is suspected to be not operational. The right panel shows the same measurement extracted from the second campaign of experimental tests carried out by Astrium on an irradiated Gaia CCD that demonstrates the expected behaviour for an operational SBC in both CCD halves. The detailed Monte Carlo model (Chapter 2) in combination with the analytical representation of the electron density distribution is able to qualitatively reproduce the fractional charge loss measurement in both cases.

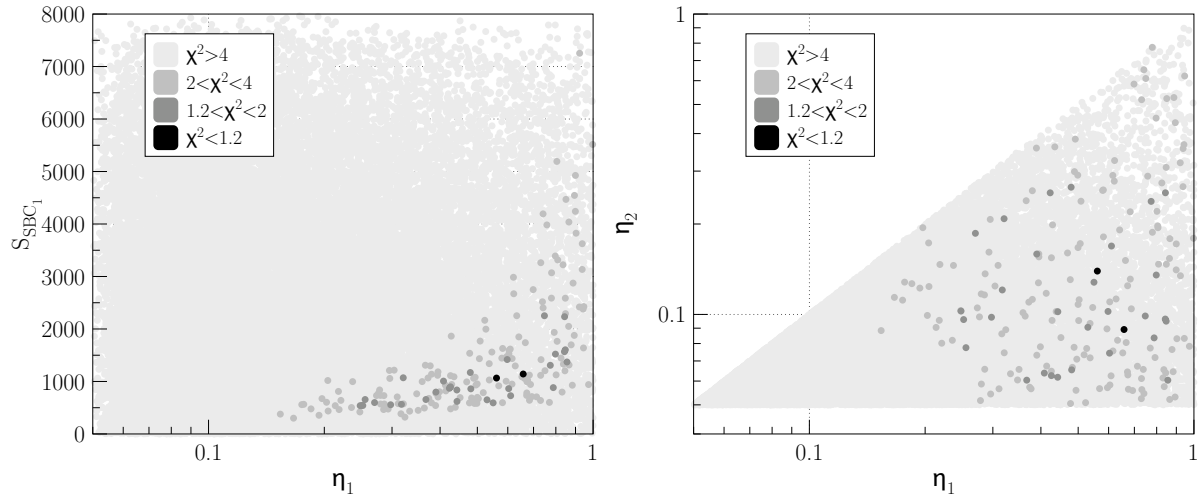


Figure 6.11 — Left: map of the resulting agreement between experimental (Fig. 6.10 left) and simulated fractional charge loss measurements as a function of the free parameters of the simulation for the modelled SBC in first half of the CCD: S_{SBC_1} and η_1. The plotted agreement corresponds to the described χ^2 (see Section 6.4.3.3) normalized by the best χ^2 achieved. This map shows that a reasonable agreement can only be obtained for: either (i) a narrow electron density distribution and a transition between SBC and BC regimes at a low signal level, or (ii) a wide electron density distribution and a transition between BC and SBC regimes at a higher signal level. **Right:** a similar map; this time the agreement is shown as a function of the electron density distribution spread factors for the SBC regimes in the first and in the second CCD halves. This map shows that it is not possible to obtain a good agreement for a narrow electron density distribution in both CCD halves in the SBC regime (bottom left corner of the map).

measurement carried out by Astrium on an irradiated Gaia CCD and using a more sophisticated optimization procedure of the simulation parameters (see Fig. 6.10 right and Chapter 2). This deviation is thus more likely due to the limitations of the model itself than inherent to the rough optimization procedure used for the purpose of this chapter. However at lower signal levels ($< 10\,000$ e$^-$), the experimental data is reproduced within the noise by the simulation. Note that at very low signal levels (i.e. for a CI level < 100 e$^-$), the fractional charge loss measurement can be greater than 1. This would mean that the charge loss is greater than the number of charges originally available for trapping. This can of course not be simulated, and it is at present impossible to know why the data include such values.
Figure 6.11 (left and right) shows how the agreement between model and data varies as a function of some of the free parameters of the simulation. This agreement is quantified by a normalized χ^2 that is the χ^2 as defined in Section 6.4.3.3 normalized by the best agreement obtained as a result of our optimization procedure. The free parameters of interest are the electron density distribution spread factors for the SBC regime in the first half of the CCD η_1 and the second half η_2, as well as the signal at which the transition between SBC and BC regimes occurs S_{SBC_1} in the first half of the CCD. The presented maps allows us to characterize the values that these parameters must take for a reasonable agreement to be achieved.
Figure 6.11 (left) shows that a good agreement cannot be obtained in general for low values of the spread factor ($\eta_1 < 0.2$) i.e. for a narrow electron distribution when the electron packet sits in the SBC of the CCD upper half. The best fit occurs at $\eta_1 \approx 0.65$, which gives $\sigma_{y,\text{SBC},1} \approx 1.3$ μm and $\sigma_{z,\text{SBC},1} \approx 0.065$ μm. Fig 6.11 (left) also shows that if the transition between SBC and BC regimes occurs at relatively high signal levels (> 2000 e$^-$) in the CCD upper half, the electron distribution must be wide in this same half ($\eta_1 > 0.8$). Fig. 6.11 (right) shows that as expected a good agreement cannot be obtained for a narrow electron density distribution in both CCD halves for the SBC regime. It also shows that as long as the distribution of the electron density is wide in the CCD upper half a good agreement can be obtained independently of the distribution spread factor in the lower half.

6.4.5 Discussion of FPR measurements

The best fit to data from all stitch blocks with suspected working SBCs in the CCD upper half (right side of Fig. 6.10) finds $\sigma_{y,\text{SBC},1} \approx 0.2$ μm. However, the best fit to data (left side of Fig. 6.10) from one stitch block AL couple (see Fig. 6.5) with suspected non-functional SBCs in the stitch block in the CCD upper half finds $\eta_1 = 0.65$ (Fig. 6.11 left), where $y_{\text{SBC,max}} = 2$ μm so $\sigma_{y,\text{SBC},1} = 1.3$ μm. The right side of Fig. 6.11 gives $\eta_2 = 0.089$ in the lower half stitch block so $\sigma_{y,\text{SBC},2} \approx 0.2$ μm. Given that the nominal doping widths are 3 and 4 μm in the upper and lower halves of the Gaia CCDs, it is unphysical that $\sigma_{y,\text{SBC},1}$ is nearly 7 times larger than $\sigma_{y,\text{SBC},2}$. We interpret this evidence as being due to non-functional or missing SBCs in the upper half stitch block but SBCs present in the lower half stitch block. If there are no SBCs or very small SBCs in the upper half stitch block, our model fit tries to find a no-SBC regime as soon as possible, which corresponds to the electron distribution being as large as possible (i.e. emulating the BC distribution). This is what our model is doing with this stitch block AL couple data,

supporting our hypothesis that the SBCs are absent in the upper half of this particular stitch block AL couple.
The alignment of the photo-lithographic masks to make each stitch block partly depends on the alignment of the previous stitch block in the AC direction. It is possible that the random alignment errors introduced in each stitch block affect the next stitch block in the AC direction so that stitch errors compound, reaching their maximum furthest from the readout node, i.e. the stitch error diverges as a function of AC like a zip coming undone. Gaia CCDs contain nine stitch blocks and eight inter-stitch block aligments in the AC direction so the worst-case stitch error could be $8 \times 0.25 = 2$ μm. Seabroke et al. (2010) Fig. 5 shows that an AC displacement of 2 μm in SBC doping is enough for SBCs to be absent from upper half stitch blocks and to explain SBC FWCs of a few thousand electrons in lower half stitch blocks. e2v consider it improbable that these tolerances would all add cumulatively and the square root of the sum of the squares assumption for uncorrelated error sources (0.7 μm) is believed to be more representative. This means that the worst-case stitch error can only occur in the stitch block furthest from where the stitching begins, which is at the readout node end.
The bottom plot of Fig. 6.9 shows the data from three candidate stitch block AL couples with suspected absent SBCs in the stitch block in the upper half of the CCD. Fig. 6.5 shows that the strongest and one other candidate are indeed the termination stitch block AL couples furthest from the readout node. The data of the other termination stitch block candidate does not extend to small enough signal levels to distinguish whether the trend will turn over, indicative of the presence of working SBCs, or continue to rise like the strongest missing SBCs candidate. The other candidate is the seventh stitch block AL couple from the readout node and so its worst-case stitch error could be $6 \times 0.25 = 1.5$ μm. Seabroke et al. (2010) Fig. 5 shows that this is enough for SBCs to be absent from the upper half stitch block. However, its trend may be turning over, indicative of the presence of working SBCs.

6.5 Pocket pumping measurements

Fig. 6.12 (left) shows that each e2v Gaia silicon wafer can only include two Gaia CCDs (CCD91-72). The room left on each wafer was used for test structures that were included to assess the radiation hardness of Gaia CCDs. In order for these test structures to be irradiated and tested, they have to be packaged. e2v packaged up two CCD221 test structures for the e2v centre for electronic imaging at The Open University to test (see Fig. 6.12 right). Because they have the same pixel architecture (and thus manufacturing processes) as Gaia CCDs, we have repeated the pocket pumping measurements conducted by Kohley et al. (2009) to test whether these test structures also have missing SBCs. Kohley et al. (2009) tested an AF FM CCD with serial number 05486-11-02. The test structures are also post-2004 but from two different batches: 06026-16-04 and 06095-09-04 [2]. The last two digits are the position of the CCD on the wafer: 03 and 04 are the two side positions for the CCD221 test structures.
All CCD221 test structures have 1440 TDI lines and 224 columns, which is smaller than

2. 06095-09-04 is still to be tested and its results will be included in the submitted version of this chapter.

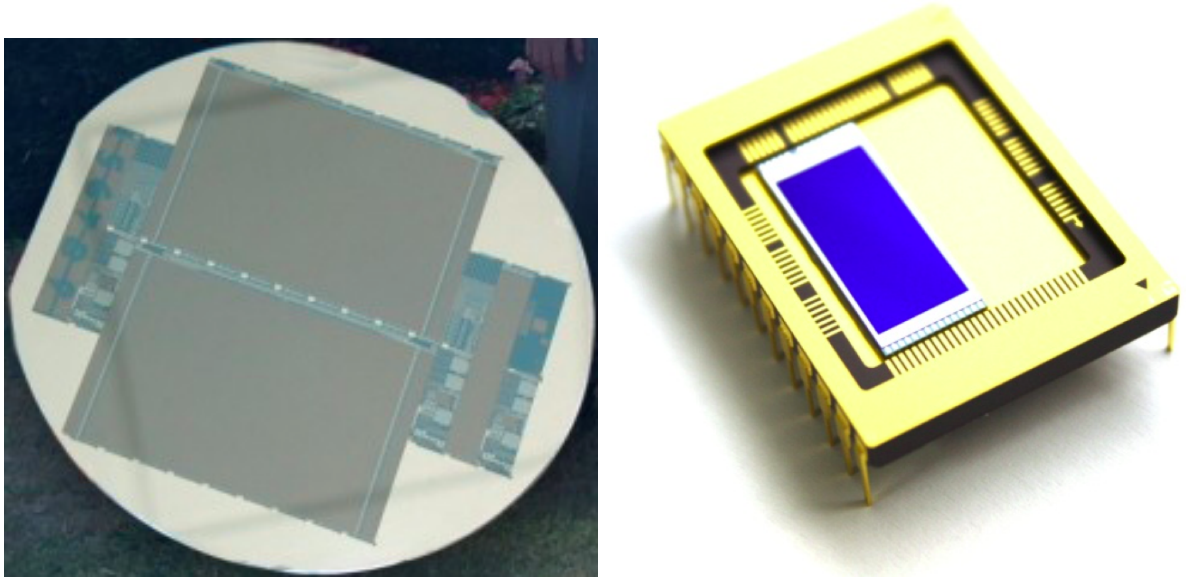


Figure 6.12 — Left: e2v Gaia silicon wafer with photoactive regions (dark grey) showing two CCD91-72s in the centre and small test structures on either side. The photoactive strip on the right side is two CCD221s abutted. **Right:** Gaia test structure CCD221 packaged for use. (Images courtesy of e2v technologies.)

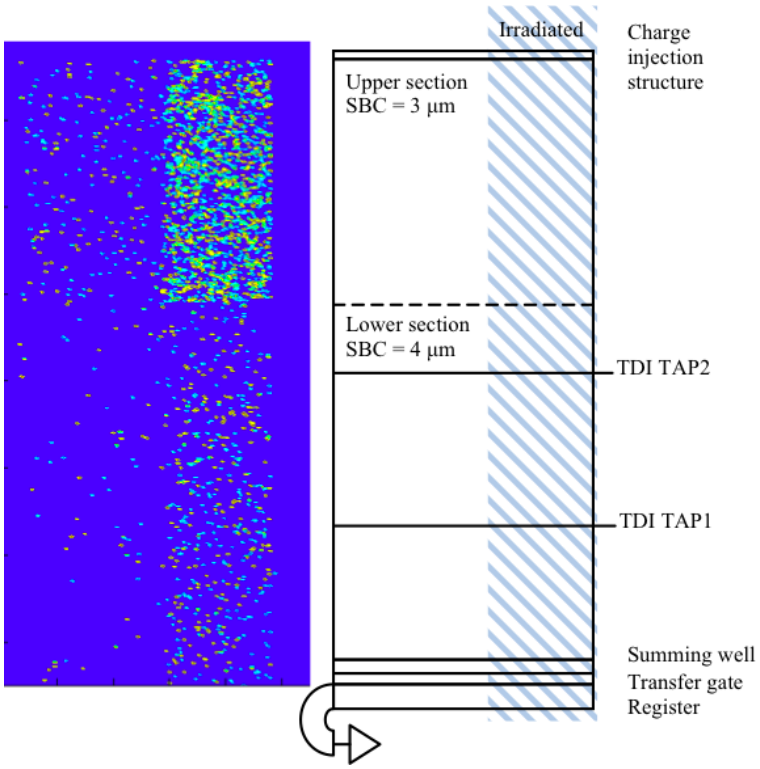


Figure 6.13 — Left: Image of a pocket pumping measurement of test structure CCD221 (serial number 06026-16-04) showing forward (yellow dots) and reverse (cyan dots) traps. **Right:** Layout of all CCD221 test structures and the area of 06026-16-04 that was subjected to proton irradiation to a 10 MeV equivalent dose of 1×10^9 protons cm^{-2}.

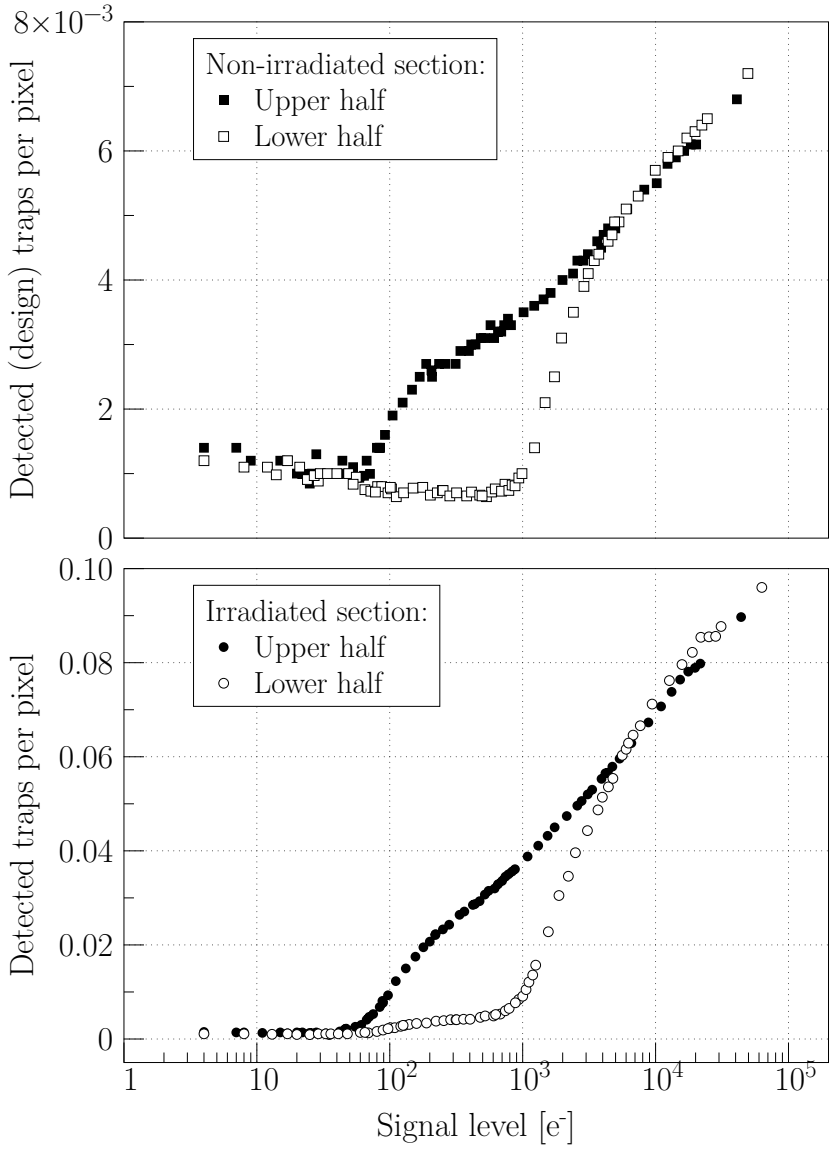


Figure 6.14 — Number of detected traps per pixel as a function of signal level. **Top:** Fabrication or design traps (see left side of pocket pumping image and left side of layout schematic in Fig. 6.13). **Bottom:** Radiation traps (see right side of pocket pumping image and right side of layout schematic in Fig. 6.13).

one Gaia stitch block (2160 or 2340 TDI lines and 250 or 108 columns). The dotted line in Fig. 6.13 is the only stitch boundary in the CCD221 test structure, showing it consists of two stitch blocks: two AL stitch blocks in one AC position (i.e. the CCD221 test structure cannot be used to compare stitch blocks in different AC positions as was done in Kohley et al. 2009).

Fig. 6.14 shows that the number of identified fabrication traps per pixel in CCD221 test structure 06026-16-04 is the same order of magnitude as seen in AF FM CCD 05486-11-02 (cf. Kohley et al. 2009 Fig. 9). It also shows that there are an order of magnitude more radiation traps per pixel as there are fabrication traps per pixel but that both types of trap follow the same trends as a function of signal size, depending on whether

they are in the upper or lower half of the test structure. These trends as a function of signal size are that in the lower half, signals < 1000 e$^-$ have a very different gradient to signals > 1000 e$^-$. We interpret the change in gradient as due to electrons spilling out of SBCs with FWCs of ≈ 1000 e$^-$. This is consistent with the spread of SBC FWCs found by Kohley et al. (2009). Like their Fig. 9, the upper half does not exhibit a change in gradient at ≈ 1000 e$^-$. We interpret this as strong evidence for no working SBCs in the upper half of this device. The number of traps in both halves of the device and both trap types converge at ≈ 65 e$^-$. Charge packets may need volumes greater than the volume of a charge packet containing ≈ 65 e$^-$ in order to be captured by more than the minimum trap density in this device. However, Kohley et al. (2009) Fig. 9 shows that their device has not reached its minimum trap density at 40 e$^-$.

6.6 Comparison of Radiation Campaigns 3 and 4 test data from the same CCDs

Both experimental and theoretical studies show that CTI drastically decreases the quality of stellar images as collected by the Gaia CCDs. On top of a significant charge loss, CTI induces an image distortion that introduces a systematic bias in the image location estimation. This affects the Gaia image location accuracy and consequently the final astrometric accuracy. It is for this reason that Astrium have carried out a series of 4 test campaigns on a small number of irradiated CCDs in order to evaluate the impact of radiation damage on the image parameter estimation and ultimately on the Gaia astrometric performances. Table 6.3 provides a summary of the CCDs tested during these campaigns. Each experimental test was carried out in order to study a range of differing effects (i.e., to examine the influence of a diffuse optical background on the CTI, to optimise CI parameters, to examine the CTI in the serial register, to address the reduction in CTI due to the recent transit of other stars etc.) and here we present a subset of the results from the analysis of these data. Note that a fifth campaign RC5 is currently in preparation.

The location biases of the images are generally calculated through an analysis of data obtained from a number of runs of an illuminated mask along the CCD which is operated in TDI mode. A number of different masks were used throughout the testing campaigns, each with a different pattern of pin-holes (of sometimes differing diameters). The goal is to simulate the accumulation of stellar images on the CCD as the mask is translated in synchronous timing with the clocking speed of the detector. The general procedure for deriving the location bias values due to radiation-induced CTI is as follows. Stellar images are accumulated from a portion of the CCD that has no radiation damage and the measured locations are then compared to the measured image locations derived from those data which originate from the irradiated region. A number of 'stars' on the mask are always detected in a damage-free region which permits the mask-displacements and illumination variations to be calibrated out. The charge loss measurements are calculated in a similar manner. An exception to this method of location bias determination occurs for a subset of data from RC3 and RC4 which was acquired using a 'sky-like' mask with a pseudo-realistic sky-pattern. This mask is passed along the detector at a number of differing AC locations and, crucially, these runs also were repeated with the mask rotated 180° with respect to the CCD. This al-

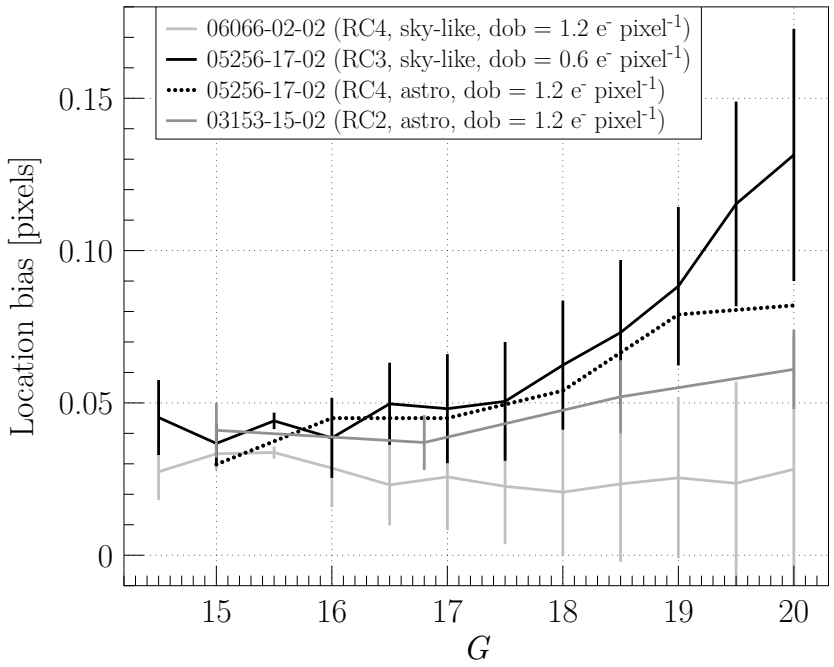


Figure 6.15 — Stellar image location bias values calculated through the analysis of data from a partially irradiated CCD using a mask with pseudo sky-patterns of pin-holes and CI interval of 1 s. A fully-functioning SBC should mitigate the location bias for faint stars, however a clear difference between the bias values obtained from the 05256-17-02 CCD (black line) and the two other CCDs used in campaigns 2 and 4 is apparent. This strongly suggests that SBCs in the 05256-17-02 CCD are not fully effective. See text for further details.

lows the location bias to be separated from calibration effects and for the bias values and calibration parameters to be solved for using a least-squares solution.

The resulting bias measurements shown in Fig. 6.15 differ between RC2, RC3 and RC4 at faint magnitudes: the RC3 data does not plateau to an approximately constant bias at faint magnitudes. This is in contrast to the data from the other campaigns. One possible reason for the RC3 bias values not plateauing is due to the presence of the constant diffuse optical background (DOB) being better controlled in the experimental set-up of RC3 compared to RC2 and RC4. When the DOB is very close to zero, then just a small amount of background light can vastly reduce image location bias due the resulting photoelectrons keeping many of the traps with long time constants (or 'slow' traps) filled. Indeed, the measured DOB levels are ~ 0.5 e$^-$ pixel^{-1} at readout lower in the RC3 data than in the RC4 data. However, data taken with the same CCD as was used in RC3 shows a similar bias function (i.e., no plateau at faint magnitudes) and in this case the DOB level is comparable to both the RC2 and RC4 data set. It thus appears that the DOB level is not the primary cause of the discrepancies between the location bias functions across the different test campaigns. This hypothesis is further strengthened when one considers that these data were obtained using a CI every second, which also has the effect of keeping the 'slow' traps filled, and making the effect of small levels of DOB much less significant. The most likely conclusion is that the measured differences are intrinsic to the CCDs and likely due to a lower functional ef-

ficiency of the SBCs in the upper half stitch blocks where the test data was taken from 05256-17-02.
Table 6.3 also summarises which of the RC CCDs are affected by the SBC issue. Fig. 6.15 shows that 03153-15-02's location bias as a function of brightness is in between the plateau of 06066-02-02 (likely functional SBCs) and 05256-17-02 (likely non-functional SBCs). However, the discussion in Section 6.3.1 of Fig. 6.10 (right) has already established that 03153-15-02 is not affected by the SBC issue. The RC data obtained from 06244-03-01 and 06273-08-01 was not designed to test SBC efficiency but it could potentially be modelled by CEMGA to determine whether these CCDs are affected by the SBC issue.

6.7 Discussion

6.7.1 Non-functional SBC statistics and predictions

e2v changed their photo-lithographic mask set in 2004. The mask set used to manufacture pre-2004 CCDs was meant to be identical to the mask set used to manufacture post-2004 CCDs. In order to test this assertion we consider evidence from the pre-2004 and post-2004 CCDs separately. Because of the low number of CCDs that have been tested, we augment our evidence found in this chapter with evidence from the literature.

Pre-2004 CCDs:
From the Sira test sample, out of the seven 2003 AF CCDs searched for evidence of missing SBCs in Section 6.4, we only find one CCD with one stitch block AL couple (termination stitch block with 108 columns) strongly suspected of not having SBCs in the upper half stitch block. We augment this pre-2004 CCD sample with the first device tested by Astrium, an AF DM (03153-15-02) that exhibits operational SBCs in all its AL stitch block couples (see Fig. 6.10). Thus, the SBC issue in the pre-2004 CCD sample of Gaia CCDs affects:
– 1 out 8 CCDs $= 12.5\%$ of CCDs;
– 1 out of $(9 \times 8 =)$ 72 stitch block AL couples in 8 CCDs $\sim 1.4\%$ of stitch block AL couples;
– 108 out of $[9 \times (2 \times 108 + 7 \times 250) =]$ $15\,728$ columns in 8 CCDs $\sim 0.7\%$ of columns.
This small number of columns with missing SBCs in their upper halves is consistent with accumulating stitch errors at each of the eight inter-stitch block AC alignments being a rare event, mainly in the termination stitch blocks furthest from the readout node.

Post-2004 CCDs:
In Section 6.5, we found one CCD221 test structure without SBCs in its upper stitch block (224 columns). In Section 6.6, two post-2004 devices have been tested by Astrium out of which one shows evidence for a non-functional SBC in its upper half (2 stitch block AL couples tested) and one appears not to be affected (5 stitch block AL couples tested). To this one case we add the device tested by Kohley et al. (2009) (9 stitch block AL couples affected). This means the SBC issue in the post-2004 sample of Gaia CCDs affects:
– 3 out 4 CCDs $= 75\%$ of CCDs;

– $(1 + 2 + 9 =)$ 12 out of $(1 + 2 + 5 + 9)$ 17 stitch block AL couples tested in 4 CCDs ∼71% of stitch block AL couples;
– $(224 + 500 + 1966 =)$ 2690 out of $(224 + 500 + 1250 + 1966 =)$ 3940 columns tested in 4 CCDs ∼68% of columns.

The difference between the number of columns with missing SBCs in their upper half in the pre- and post-2004 CCD samples was evaluated statistically using the χ^2 test. The difference between the two samples was shown to be highly statistically significant ($\chi^2 \sim 1{\times}10^4$, $p < 5{\times}10^{-8}$). Therefore the null hypothesis that the number of columns with missing SBCs in their upper half is the same in both the pre- and post-2004 CCD samples is rejected at the $>5\sigma$ significance level. This points to a change in e2v manufacturing of Gaia CCDs between 2003 and 2005. There are no post-2004 CCDs with some upper half stitch blocks affected by the SBC issue and others not. This suggests that a rare compounding of stitch errors is not responsible but a more systematic change in manufacturing has occurred since 2003. The different mask set is the only known change between 2003 and 2005. The two mask sets were meant to be identical but a large mask writing error may have occurred in the manufacture of the post-2004 mask set.

Gaia CCDs were front-face processed first and then back-face processed. Front-faced processing is when the electrodes are made and all doping implanted and is batch based. Because all the Gaia CCDs have the same pixel architecture, front-face processing is the same for all Gaia CCDs: AF, Blue Photometer (BP), Red Photometer (RP). e2v and Astrium refer only to these CCD variants as the Radial Velocity Spectrometer (RVS) CCDs are identical to RP CCDs. Table 6.4 shows that back-face processing is different for the different types of Gaia CCD and is wafer-based. Each CCD serial number encodes when the front-face processing occurred but not when the back-face processing occurred, which can be much later.

The fact that the SBC issue is only related to front-face processing and all the wafers in a batch have the same nominal front-face processing means that if a CCD of any type is found to be affected by the SBC issue, then all the CCD variants (AF, BP and RP) from that batch are likely to also be missing SBCs in their upper half stitch blocks. Table 6.5 summarizes the CCDs already tested for SBC efficiency and those that could be tested in the future. They are all from different batches so whether other CCDs from the same batch are also affected remains unproven (none of the DM or EM CCDs are from these three affected batches either). In the absence of evidence demonstrating that all CCDs from the same batch have the same SBC properties, we can only tentatively predict which FM CCDs are likely to be affected by the SBC issue, based on e2v's premise that all CCDs from the same batch are nominally the same (see Table 6.6), and which are currently assigned to Gaia's Focal Plane Array (FPA, see Fig. 6.16 and Table 6.7).

The four CCDs tested for the SBC issue in Table 6.5 are from four different batches, one of which does not include any FMs. Fig. 6.16 and Table 6.7 show that the remaining three batches that include FMs only include 19% of the CCDs that currently make up Gaia's FPA and for which tentative predictions of which are affected or not can be made. Fig. 6.17 shows the 106 CCDs that currently make up Gaia's FPA come from 32 different batches. Therefore, the 81% of CCDs that currently make up Gaia's FPA and for which tentative predictions cannot be made come from 29 completely untested

batches. Table 6.6 shows that some of the batches include flight spares. Like the FPA-assigned CCDs, the flight spare CCDs are also coupled to FM Proximity Electronics Module (PEM). The clock timing of PEMs does not allow modification for backwards shifts in the image area for pocket pumping. In case the currently FPA-assigned CCD and PEM couples need to be replaced, the flight spare CCD and PEM couples cannot be decoupled (and thus pocket pumped) until after Gaia is launched.
After launch, the CCDs on board can in principle be tested using CI MIM. If SBC FWC needs to be characterised before then, in principle all the test structures that were on the same wafers as the 106 Gaia CCDs that will actually fly could be tested using the pocket pumping technique. However, to obtain these may require paying e2v to package all these test structures to allow them to be tested. This would be the definitive method to resolve the SBC issue but could be very expensive and the testing could be very time consuming. A cheaper and quicker solution would be to establish how likely all the wafers in a single batch have the same SBC properties. This could be done by getting the test structures from the same batch but different wafer as the already tested CCDs in Table 6.5. If the test structure pocket pumping results agree with the results in Table 6.5 then e2v's premise that all CCDs from the same batch have the same SBCs will be supported by evidence and the tentative predictions of which FPA CCDs are likely to be affected by the SBC issue and which are not will be less tentative.
Table 6.7 shows that 17% of FPA CCDs are likely to be affected by the SBC issue. Given that only three out of the 32 FPA batches have been tested, 17% represents a likely minimum number of FPA CCDs affected. All the post-2004 evidence suggests that if a CCD is affected by the SBC issue, then all its upper half stitch blocks are affected. We can predict the fraction of CCDs that will be affected by continuing with the assumption that if a batch is affected by the SBC issue then all CCDs in that batch will be affected. We only have test data for four post-2004 batches, three of which were affected by the SBC issue. To predict the fraction of affected FM CCDs in the batches not examined so far and quantify the uncertainty on this number we use the following simple statistical model. We assume that there is a certain probability θ that a given batch suffers the SBC issue, and that the number of affected batches follows from a Bernouilli trial with probability θ. From the batches tested we know that θ is most likely 0.75, but the distribution of possible values of θ is rather wide because we only have a small number of tested batches. In this statistical model θ is given by a $B(4,2)$ distribution, with $B(x,y)$ the Beta function. From this distribution it follows that with 99% confidence $0.19 \leq \theta \leq 0.98$. We can now simulate a Bernouilli trial for the untested batches in order to predict the number of FM CCDs affected by the SBC issue. In the simulation we draw a random θ from the $B(4,2)$ distribution and then perform a Bernouilli trial for the 29 untested batches. In each trial we keep track of which batch from Fig. 17 was affected and add the corresponding number of CCDs to the total. In this way we can predict the distribution of the fraction of affected FM CCDs in the untested batches. Not surprisingly the median of the distribution is at 69%. However, the inter quartile range and the 99% range are 52–83% and 11–100% respectively, reflecting the small number of batches on which we base our extrapolation.

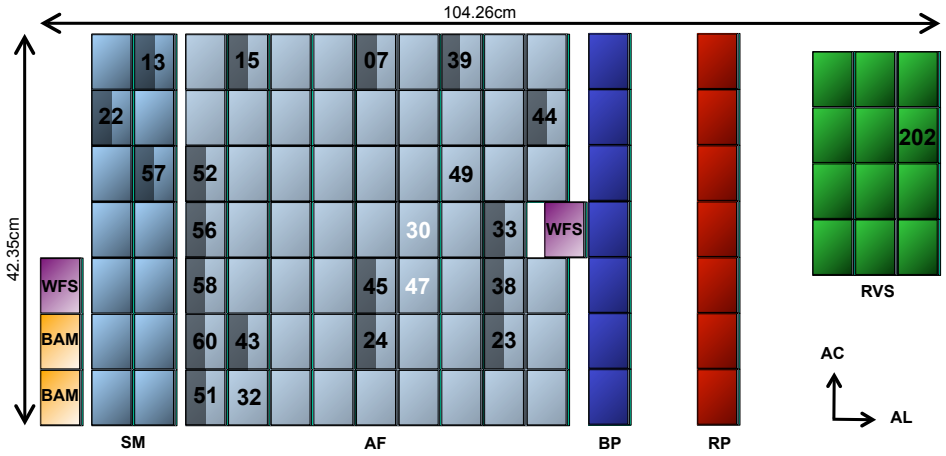


Figure 6.16 — Gaia's FPA shows the current assignment of FM CCDs (subject to on-ground testing or other Assembly, Integration, and Verification (AIV) procedures, CCDs can still be replaced by flight spares). The shaded regions indicate which CCDs from Table 6.6 are likely to be affected by the SBC issue with Astrium FM numbers in black. Numbers in black without shaded regions on AF CCDs indicate CCDs from a batch 06095-XX-0X which may or may not be affected by the SBC issue, depending on the pocket pumping results from test structure 06095-09-04, which will be included in the submitted version of this chapter. The number in black without a shaded region on one of the RVS CCDs indicates that this FM CCD is from the batch 06273-XX-0X, which includes 06273-08-01 (one of the RC CCDs, see Section 6.6). 06273-08-01 could either be pocket pumped in RC5 or previous RC data could be modelled to determine whether it, and so also the RVS FPA CCD from the same batch, are affected by the SBC issue. White numbers indicate CCDs currently integrated into the FPA are from the same batch (06066-XX-0X) as the FM CCD with likely working SBCs in its upper half (06066-02-02, see Section 6.6) are so are also likely to be fully functional. Background diagram courtesy of A. Short (ESA, ESTEC)

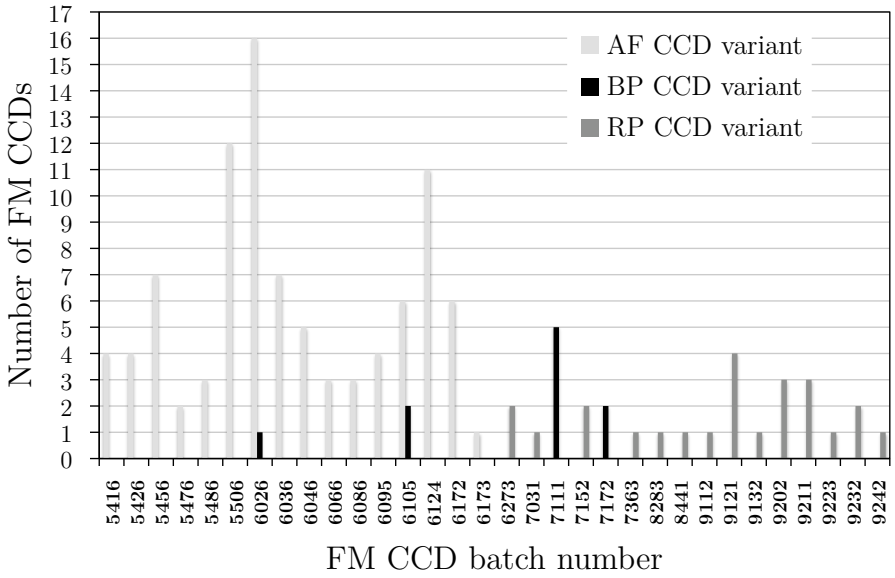


Figure 6.17 — Number of FM and spare CCDs (AF, BP and RP variants) chronologically plotted as a function of batch number.

Serial Number	Variant	Affected
03153-15-02	AF DM	N
05256-17-02	AF EM	Y
06066-02-02	AF EM	N
06244-03-01	RP DM	?
06273-08-01	RP EM	?

Table 6.3 — Summary of the CCDs tested during Astrium's RCs. Affected indicates whether the CCD is likely to be affected by the SBC issue (Y), likely to have fully functional SBCs (N) or yet to be tested for SBC efficiency (?).

CCD	d	Surface	λ	Instrument (number of CCDs)
AF	16	Basic	650	AF (62), SM (14), WFS (2)
BP	16	Enhanced UV	360	BP (7)
RP	40	Basic	750	BAM (2), RP (7), RVS (12)

Table 6.4 — Comparison of back-face processing between the Gaia CCDs. d is the depth the wafer was thinned to in μm. Surface passivation was annealing "dangling bonds" with hydrogen gas at high temperatures to reduce dark current. All the Gaia CCDs have a single layer Hafnium oxide anti-reflection coating but each variant has its coating centred on a different wavelength (λ in nm). The other Gaia instruments are Sky Mapper (SM), WaveFront Sensor (WFS) and Basic-Angle Monitoring (BAM).

Method	Variant	Serial Number	Batches	Affected
RC3&4 (this chapter: re-analysis)	AF EM CCD	05256-17-02	05256-XX-0X	Y
PP (Kohley et al. (2009): measurements & analysis)	AF FM CCD	05486-11-02	05486-XX-0X	Y
PP (this chapter: 1st measurement & analysis)	AF FM test structure	06026-16-04	06026-XX-0X	Y
RC4 (this chapter: re-analysis)	AF EM CCD	06066-02-02	06066-XX-0X	N
PP (undergoing: 1st measurement & analysis)	AF FM test structure	06095-09-04	06095-XX-0X	?
RC5 PP? (recommendation of this chapter)	RP DM CCD	06244-03-01	06244-XX-0X	?
RC5 PP? (recommendation of this chapter)	RP EM CCD	06273-08-01	06273-XX-0X	?

Table 6.5 — Measurement method used to test FM CCDs and test structures, where PP is pocket pumping and RC is analysis of Radiation Campaign data, e2v serial numbers of the tested devices and the FM batches likely to be affected by the SBC issue (Y), likely to have fully functional SBCs (N) or yet to be tested for SBC efficiency (?). Only FM CCDs will be flown on Gaia and all FM CCDs were built post-2004 so this table excludes 03153-15-02 (RC CCD shown not be affected by the SBC issue) and 03153-20-01 (the one pre-2004 CCD found with missing SBCs in one of its upper half stitch blocks).

6.7.2 Impact on the Gaia image location accuracy

Gaia has been designed to perform absolute astrometric measurements at very high accuracy. As the estimated image location for all CCD observations are ultimately used to derive the star astrometric parameters, the requirements on the image quality are very stringent. CTI distorts the image and decreases the signal-to-noise ratio. While the image distortion, if not properly taken into account, introduces a significant bias (e.g., Fig. 6.15) in the image location estimation, the decrease in signal-to-noise ratio implies an irreversible loss of accuracy independent of any estimator. Chapter 3 presents an accurate and detailed evaluation of the two effects. Their results were obtained by simulating numerous damaged and CTI-free stellar transits for different CCD operating conditions but assuming properly working SBCs; they used parameters for the electron density distribution derived from the FPR measurement based on the CCD 03153-15-02 with functional SBCs. We are now interested in repeating the same experiment with parameters for the electron density distribution in the SBC regime derived from the FPR measurement based on the CCD 03153-20-01 with non-functional SBCs in the upper CCD half (see Section 6.4.4).

Figure 6.18 (top) shows the relative loss of accuracy in image location as a function stellar magnitude. This is computed by comparing the theoretical limits (the Cramér-Rao bounds) to the image location accuracy for a damaged and a CTI-free Gaia-like image. Zero loss of accuracy corresponds to the CTI-free case (for more details see Chapter 3). Fig. 6.18 (bottom) shows the image location bias as a function stellar magnitude. It is obtained by applying the Gaia image location procedure to several hundreds of profiles without applying any CTI mitigation procedures, the bias itself is computed by subtracting the true from the estimated location. Fig 6.18 provide a comparison for three different cases: (i) the Gaia nominal case i.e. functional SBCs in both CCD halves (blue continuous line), (ii) non-functional SBCs in the CCD upper half (red dashed line), and (iii) no SBCs at all (black dotted line). The middle case corresponds to the most realistic case for Gaia CCDs containing non-functional SBCs as identified in this chapter. For the simulation we used an active trap density of 1 trap pixel^{-1}. This particular density was shown in Chapter 3 to reproduce the amplitude of location bias measured using experimental test data taken 1 s after a CI and for a Gaia irradiated CCD with a radiation dose of 4×10^9 protons cm^{-2} (10 MeV equivalent). This dose corresponds to the upper limit of the predicted accumulated radiation dose for the Gaia nominal lifetime.

From both figures it is clear that a non-functional SBC in the CCD upper half degrades the location precision and introduces a greater bias for the faint magnitudes ($G > 15$), however it is not as severe as for the no SBC (in both halves) case. At magnitude 20 both precision and bias are worse by a factor 2. A 10% decrease in image location accuracy is acceptable and does not threaten the Gaia requirements. The bias can in principle be calibrated out by using the same mitigation procedure as foreseen in the case of nominal CCD functioning. The modelling used in the mitigation procedure will however have to be modified so that a missing SBC in a fraction of the AL stitch block couples can be accounted for.

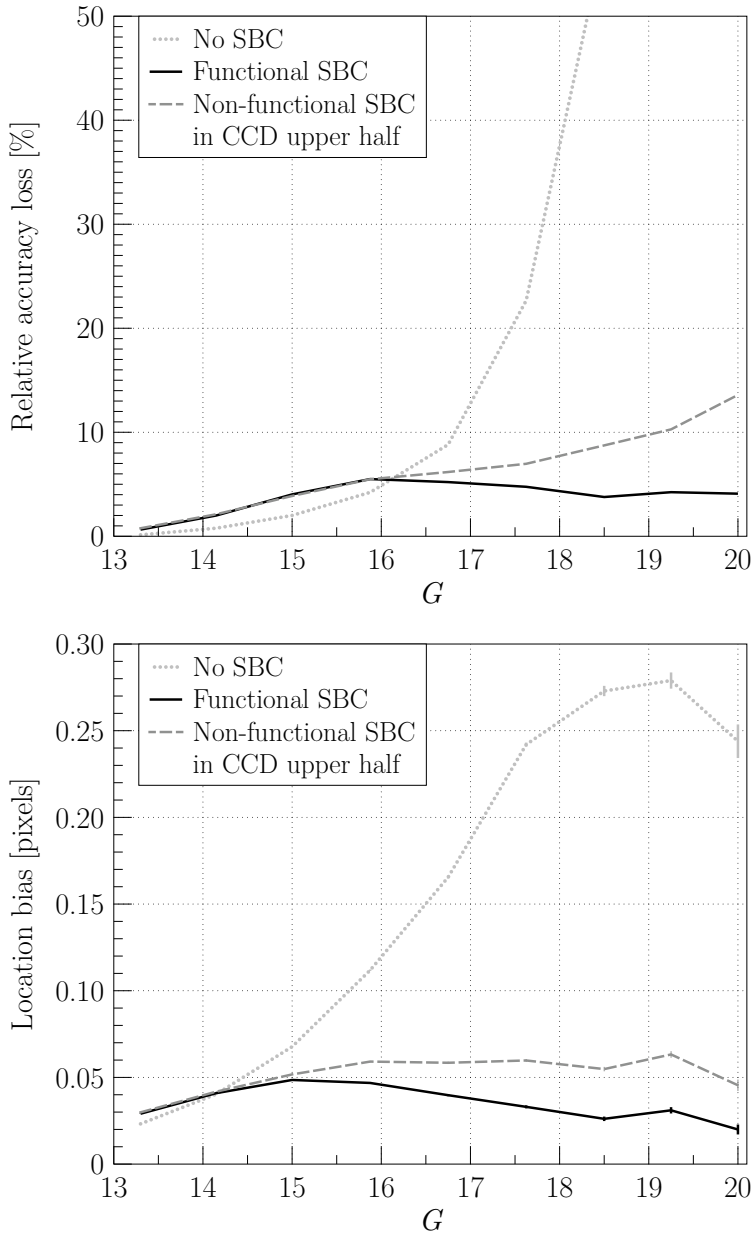


Figure 6.18 — Predicted end-of-life relative image location accuracy loss (top) and bias (bottom) as a function of magnitude (G-band) for a 1 s CI delay and three different pixel architecture cases: (i) functional SBCs in both CCD halves (blue continuous line), (ii) non-functional SBCs in the CCD upper half (red dashed line), and (iii) no SBCs at all in both CCD halves (black dotted line). **Top:** The non-functional SBC cases (ii) and (iii) only affects the faint signals ($G > 15$). For case (iii) the loss of accuracy is dramatic: 200% for $G = 20$ (not shown for readability), however, as explained, in the text this is not what is expected for Gaia. For cases (i) and (ii), representative of the Gaia CCD with respectively functional and non-functional SBCs the CTI-induced loss of accuracy does not exceed 10%, which is acceptable regarding the Gaia requirements (Chapter 3). **Bottom:** The CTI mitigation procedure enables a 90% recovery of the bias (Chapter 3), bringing back the residual bias to an acceptable level for cases (ii) and (iii). Error bars represent the statistical uncertainty, not the standard deviation. Note that here no mitigation procedure at the image processing level was applied.

6.8 Conclusions

We have re-analysed the Hopkinson (2006) data set of First Pixel Response (FPR) measurements of seven Gaia Engineering Model (EM) CCDs. This data is grouped according to how e2v manufactured the CCDs: stitch blocks, which correspond to the size of the photo-lithographic mask used to implant doping features that define the pixel architecture. Because Gaia operates in Time-Delayed Integration (TDI) mode, observations are transferred through two stitch blocks in the ALong scan (AL) direction before being read out. The FPR data therefore samples these stitch block AL couples. We only find one stitch block AL couple (the termination stitch block AL couple at the end of the CCD, consisting of 108 columns) in one of these CCDs where the FPR curve does not exhibit the characteristic bump of the supplementary buried channel (SBC, otherwise known as a notch or mini-channel).

Modelling this termination stitch block AL couple with a detailed electrode-level model finds a best fit where the electron distribution within the SBC is not confined to a narrow region as expected in the upper half of the CCD but is confined in the lower half. Its distribution in the upper half is more like a buried channel (BC), whereas it is like a SBC in the lower half. This suggests the stitch block in the upper half of the CCD does not have SBCs. The reason for missing SBCs and measured SBC Full Well Capacities has been explained by the SBC doping abutting the anti-blooming drain (ABD) and the ABD shielding doping cancelling out much of the SBC doping, reducing the SBC potential to much smaller than designed (Seabroke et al. 2010).

Evidence for missing or non-functional SBCs in the upper half has already been published (Kohley et al. 2009) but this was found in only one CCD. Based on our modelling results and the effects of CTI on artificial stellar images we can infer the presence of non-functional SBCs in the upper half of the CCDs tested by Astrium during their Radiation Campaigns (RCs). To the seven Astrometric Field (AF) CCDs tested by Hopkinson (2006), we add the results of one test structure tested by us and four more AF CCDs tested by Astrium and the results of Kohley et al. (2009) and thus present the largest sample of Gaia CCDs to be analysed for the presence of SBCs. The difference between the number of columns with missing SBCs in their upper half in the pre- and post-2004 CCD samples is highly statistically significant ($>5\sigma$), pointing to a change in e2v manufacturing of Gaia CCDs between 2003 and 2005. Different photo-lithographic mask sets is the only known change between 2003 and 2005.

The fact that the SBC issue is only related to front-face processing and all the wafers in a batch have the same nominal front-face processing means that if a CCD of any type is found to be affected by the SBC issue, then all the CCD variants, AF, Blue Photometer (BP) and Red Photometer (RP), from that batch are likely to also be missing SBCs in their upper half stitch blocks. By comparing the batch numbers of the three affected FM CCDs (three different batches) with those currently assigned to Gaia's Focal Plane Array (FPA), we tentatively predict that a minimum of 17% of FPA CCDs (three Sky Mappers and 15 AFs) are likely to be affected by the SBC issue. In the absence of further testing, we predict that in the other 29 completely untested batches 69% of the CCDs may be affected (between 11 and 100% with a 99% confidence interval). Therefore, despite different CCD manufacturers, it appears likely that the majority of Gaia's

106 CCDs have the same missing SBC issue as the two CCDs in Hubble's Advanced Camera for Surveys (ACS)/Wide Field Channel (WFC), albeit only in the upper half of Gaia's detectors.

By analyzing the data from Astrium RCs, this chapter has also addressed whether SBCs were working properly in Gaia CCDs used to predict the mission performances and against which CTI mitigation models are being developed and tested. We find that the AF CCD tested in RC3 appears to have non-functional SBCs in the upper half of the CCD region where the test data was obtained. The two other tested AF CCDs are not affected by this issue. This means that test data exists for predicting mission performances for CCDs with and without functioning SBCs in their upper halves and testing mitigation models. The limited test data available means we have only been able to demonstrate that non-operational SBCs are likely in AF FM CCDs. We show that for an AF CCD with non-operational SBCs in its upper half, the Gaia image location accuracy is affected by less than 10% if the CTI effects can be properly calibrated, which is within Gaia AF requirements. Future testing may reveal that other affected batches contain FM BP and RP variant CCDs. How this affects BP, RP and RVS science is beyond the scope of this chapter but should be investigated in case BP, RP or RVS Gaia science requirements are not met.

Gaia's CCDs have already been integrated into the Focal Plane Array by Astrium so it is too late to change which FM CCDs are selected to fly on Gaia.

If the SBCs need to be characterised before launch, then the likelihood of all the wafers in a single batch having the same SBC properties needs to be established. This could be done by getting the test structures from the same batch but different wafer as the already tested CCDs. If the test structure pocket pumping results agree with the extant results then e2v's premise that all CCDs from the same batch have the same SBCs will be supported by evidence and the tentative predictions of which FPA CCDs are likely to be affected by the SBC issue and which are not will be less tentative. If this is not possible, two of the three methods we have presented to measure the SBCs can be used in-flight: FPR and the minimum injection method (MIM). Pocket pumping cannot be used in flight because the clock timing in the Proximity Electronics Module does not allow modification for backwards shifts in the image area. In-flight tests only need to be conducted once for each CCD so it should be possible to schedule it before routine operations begin. FPR will be used anyway because it can measure radiation trap properties. However, it has the disadvantage of being model-dependent. MIM is a fast method that directly measures whether SBCs are present (see Recommendations in Section 6.9). The identification of CCD columns with non-operational SBCs will help the calibration of the CTI in all Gaia instruments and guarantee that at least the final astrometric accuracy of Gaia is preserved (the final photometric and spectroscopic accuracy still needs to be investigated).

It should be emphasized that SBC Full Well Capacity (FWC) has never been a formally agreed CCD acceptance criterion in the Gaia contract between ESA/EADS Astrium and e2v. For this reason, it was not tested very thoroughly in the Gaia development phase. A more detailed test plan that went beyond the CCD acceptance criteria would have found the SBC issue, maybe early enough to change the pixel design. The SBC issue is related to manufacturing alignment accuracy. We would recommend e2v to ei-

ther widen future SBCs if they are adjacent to ABDs to compensate for the SBC issue or place the SBC in the middle of the BC to avoid the issue completely but the technique used to manufacture Gaia's CCDs is a decade old and has been superseded by Nikon, which delivers five times better accuracy. The Nikon technique has not been used yet to manufacture SBCs, because the CCD industry now views SBCs as counter-productive. Nevertheless we demonstrate that the improvement in signal-to-noise ratio at low signal levels provided by functional or even partly-functional SBCs is critical to achieving the Gaia requirements; it is thus important that the industry pursues the development of such features to prepare for future space missions with stringent image quality requirement at low signal levels.

6.9 Recommendations

This paper can only tentatively predict how many CCDs on the Gaia satellite will be affected by the Supplementary Buried Channel (SBC) issue because the prediction uses e2v's assumption that all CCD batches will have nominally the same SBC characteristics. Therefore, the top priority is to directly measure Gaia CCDs for the SBC issue on the satellite prior to launch if possible. The quickest method to do this is the Minimum charge Injection Method (MIM). However, the efficiency of the method has not been demonstrated. In this context, we recommend the following:

1. Our highest priority recommendation is to test the efficiency of MIM. If our test structure that has not been pocket pumped yet is found to have working SBCs in its upper half, then we (NJM and AH) can test the efficiency of MIM with it. If the test structure is found not to have working SBCs in its upper half, then MIM should be tested on one of the CCDs shown in this paper to have working SBCs in their upper halves (pre- or post-2004 CCDs). It appears as though these tests will not be conducted by Astrium but we are investigating whether these CCDs can be loaned out from Astrium to allow these tests to take place.
2. If MIM efficiency is demonstrated, MIM should be included in on-ground testing or other Assembly, Integration, and Verification (AIV) procedures prior to launch.
3. If MIM efficiency is demonstrated but MIM tests cannot be performed prior to launch, they should be conducted post-launch during initial in-orbit calibration before nominal operations begin.
4. If MIM efficiency is not demonstrated and modelling mission data does not identify which CCDs are affected by the SBC issue but this information is required to mitigate the resulting increased radiation damage, then we recommend that First Pixel Response testing and modelling, as done in this paper, is included in-flight. When the CCDs are sufficiently irradiated, this could be scheduled during orbital maintenance manoeuvres during which nominal data acquisition is suspended for two hours once per month.

Acknowledgements

Dr. Gordon Hopkinson (1952-2010) of Surrey Satellite Technology Ltd., Sevenoaks, UK, kindly provided the data used in this chapter and was very supportive of the

project but sadly did not live to see it come to fruition. This chapter is dedicated to his enormous contribution to developing astronomical instrumentation in general and to measuring Gaia radiation effects in particular. We thank David Burt (e2v technologies plc, Chelmsford, UK) for Figures 6.1, 6.2, 6.3 and 6.4 and for his invaluable insight into the design and manufacturing processes of e2v's Gaia CCDs through his regular visits to the e2v centre for electronic imaging at The Open University, UK. GMS is funded by the UK VEGA Gaia Data Flow System grant. The work of TP was supported by the European Marie-Curie research training network ELSA (MRTN-CT-2006-033481).

Batch	FM	Serial number	Affected
05256-XX-0X	None	None	Y
05486-XX-0X	FM07	05486-04-02	Y
	FM06*	05486-12-02	Y
	FM13	05486-15-02	Y
06026-XX-0X	FM52	06026-02-02	Y
	FM15	06026-06-01	Y
	FM44	06026-07-02	Y
	FM51	06026-09-02	Y
	FM33	06026-11-01	Y
	BPFM101*	06026-12-02	Y
	FM38	06026-13-01	Y
	FM39	06026-14-01	Y
	FM22	06026-15-02	Y
	FM58	06026-17-02	Y
	FM43	06026-18-01	Y
	FM45	06026-18-02	Y
	FM57	06026-19-02	Y
	FM56	06026-20-02	Y
	FM60	06026-23-02	Y
	FM23	06026-24-01	Y
	FM24	06026-24-02	Y
06066-XX-0X	FM30	06066-03-01	N
	FM47	06066-12-02	N
06095-XX-0X	FM32	06095-10-01	?
	FM63*	06095-10-02	?
	FM49	06095-21-01	?
	FM31*	06095-22-01	?
06244-XX-0X	None	None	?
06273-XX-0X	RPFM202	06273-01-02	?
	RPFM203*	06273-09-01	?

Table 6.6 — Astrium FM numbers and e2v serial numbers of FM CCDs likely to be affected by the SBC issue due to coming from affected batches (Y), likely to have fully functional SBCs due to coming from a likely unaffected batch (N) or not known because no CCDs from the batch have been tested for SBC efficiency (?), where * indicates flight spares.

Instrument	Number	Total	%	Affected
WFS	0	2	0	Y
	0	2	0	N
	2	2	100	?
BAM	0	2	0	Y
	0	2	0	N
	2	2	100	?
SM	3	14	21	Y
	0	14	0	N
	11	14	79	?
AF	15	62	24	Y
	2	62	3	N
	45	62	73	?
BP	0	7	0	Y
	0	7	0	N
	7	7	100	?
RP	0	7	0	Y
	0	7	0	N
	7	7	100	?
RVS	0	12	0	Y
	0	12	0	N
	12	12	100	?
All	18	106	17	Y
	2	106	2	N
	86	106	81	?

Table 6.7 — Summary of numbers of FM CCDs currently assigned to Gaia's FPA (see Fig. 6.16) likely to be affected by the SBC issue due to coming from affected batches (Y), likely to have fully functional SBCs due to coming from an unaffected batch (N) or not known because no CCDs from the batch have been tested for SBC efficiency (?).

Bibliography

Anderson J., Bedin L. R., 2010, PASP, 122, 1035

Anderson J., King I. R., 2000, PASP, 112, 1360

Bastian U., Biermann M., 2005, Astronomy & Astrophysics, 438, 745

Bombrun A., Lindegren L., Holl B., Hobbs D., Lammers U., Bastian U., 2011, in preparation for Astronomy and Astrophysics

Bombrun A., Lindegren L., Holl B., Jordan S., 2010, A&A, 516, A77

Bristow P., 2003, Technical report, Application of Model Derived Charge Transfer Inefficiency Corrections to STIS Photometric CCD Data

Bristow P., Kerber F., Rosa M. R., 2005, in A. M. Koekemoer, P. Goudfrooij, & L. L. Dressel ed., The 2005 HST Calibration Workshop STIS Calibration Enhancement: Wavelength Calibration and CTI

Brown S. W., 2009a, Technical report, Further analysis of Radiation Campaign 1. DPAC

Brown S. W., 2009b, Technical report, Independent Analysis of Astriums Radiation Campaign 2. DPAC

Burt D., 2003, Technical report, AF CCD Design Report. e2v technologies plc

Burt D., 2005a, Technical report, AF CCD Design Justification. e2v technologies plc

Burt D., 2005b, Technical report, Re-Design of the AF CCD (CCD91). e2v technologies plc

de Bruijne J., 2008, Technical report, CCD pixel pitch, stitch boundaries, and supplementary buried channel. ESA/ESTEC

de Bruijne J., 2009a, Technical report, Along- and across-scan location-estimation performance. ESA/ESTEC

de Bruijne J., 2009b, Technical report, Gaia astrometric performance: summer -2009 status. ESA/ESTEC

de Bruijne J. H. J., 2005, in C. Turon, K. S. O'Flaherty, & M. A. C. Perryman ed., The Three-Dimensional Universe with Gaia Vol. 576 of ESA Special Publication, Accuracy Budget and Performances. pp 35–+

Dolphin A. E., 2009, PASP, 121, 655

Drimmel R., Bucciarelli B., Lattanzi M. G., Spagna A., Jordi C., Robin A. C., Reylé C., Luri X., 2005, in C. Turon, K. S. O'Flaherty, & M. A. C. Perryman ed., The Three-Dimensional Universe with Gaia Vol. 576 of ESA Special Publication, What Gaia Will See: All-Sky Source Counts from the GSC2. p. 163

Drimmel R., Spagna A., Bucciarelli B., et al., 2005, Technical report, What Gaia will see - source counts in G based on GSC2, http://www.rssd.esa.int/llink/livelink/open/452346. INAF

ESA 1997, The Hipparcos and Tycho Catalogues

Feynman J., Spitale G., Wang J., Gabriel S., 1993, Journal of Geophysical Research, 981, 13281

Fusero F., 2007, Technical report, Gaia Radiation Analysis Report. EADS Astrium

Georges L., 2008, Technical report, Tests Report of the Radiation Campaign 2. EADS Astrium

Georges L., 2009, Technical report, Radiation Campaign 3 Astrium astrometric regular tests report. EADS Astrium

Goudfrooij P., Kimble R. A., 2002, in S. Arribas, A. Koekemoer, & B. Whitmore ed., The 2002 HST Calibration Workshop : Hubble after the Installation of the ACS and the NICMOS Cooling System Correcting STIS CCD Photometry for CTE Loss. p. 105

Grant C. E., Bautz M. W., Kissel S. E., LaMarr B., 2004, in A. D. Holland ed., Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 5501, A charge transfer inefficiency correction model for the Chandra advanced CCD imaging spectrometer. pp 177–188

Hall D., 2010, Technical report, Analysis of the RVS tests in the 3rd Astrium Radiation Campaign. Open University

Hall R. N., 1952, Physical Review, 87, 387

Hardy T., Murowinski R., Deen M. J., 1998, IEEE Transactions on Nuclear Science, 45, 154

Høg E., 2008, http://www.astro.ku.dk/∼erik/History.pdf

Holl B., Hobbs D., Lindegren L., 2010, in S. A. Klioner, P. K. Seidelmann, & M. H. Soffel ed., IAU Symposium Vol. 261 of IAU Symposium, Spatial correlations in the Gaia astrometric solution. pp 320–324

Holl B., Lindegren L., Hobbs D., 2009, Technical report, AGISLab A facility for experimental astrometric solutions. DPAC

Holland A., Smith D., 2004, Technical report, Charge injection strategies for Gaia. Open University

Hopkinson G., 2006, Technical report, Final Report Radiation Testing of Gaia AF and MBP CCDs. Surrey Satellite Technology Ltd.

Hopkinson G., 2008, Technical report, Gaia Cold Irradiation Final Report. Surrey Satellite Technology Ltd.

Hopkinson G., Short A., Vetel C., Zayer I., Holland A., 2005, IEEE Transactions on Nuclear Science, 52, 2664

Janesick J., 2001, Scientific Charge-Coupled Devices. SPIE

Jordi C., Gebran M., Carrasco J. M., de Bruijne J., Voss H., Fabricius C., Knude J., Vallenari A., Kohley R., Mora A., 2010, Astronomy & Astrophysics, 523, A48

Katz D., et al., 2004, MNRAS, 354, 1223

Kohley R., 2011, in EAS Publications Series Vol. 45 of EAS Publications Series, Next Generation of Light Detectors in Astronomy. pp 73–80

Kohley R., Raison F., Martin-Fleitas J. M., 2009, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 7439, Gaia: operational aspects and tests of Gaia Flight Model CCDs

Lindegren L., 1978, in F. V. Prochazka & R. H. Tucker ed., IAU Colloq. 48: Modern Astrometry Photoelectric astrometry - A comparison of methods for precise image location. pp 197–217

Lindegren L., 1998, Technical report, Charge trapping effects in CCDs for Gaia astrometry, `http://www.rssd.esa.int/doc_fetch.php?id=2844371`. Lund Observatory

Lindegren L., 2003, Technical report, Representation of LSF and PSF for GDAAS-2. Lund Observatory

Lindegren L., 2005, in C. Turon, K. S. O'Flaherty, & M. A. C. Perryman ed., The Three-Dimensional Universe with Gaia Vol. 576 of ESA Special Publication, The Astrometric Instrument of Gaia: Principles. pp 29–+

Lindegren L., 2006, Technical report, Centroid definition for the Astro Line Spread Function. Lund Observatory

Lindegren L., 2008, Technical report, A general Maximum-Likelihood algorithm for model fitting to a CCD sample data. Lund Observatory

Lindegren L., 2009, Technical report, Minimum-dimension LSF modelling. Lund Observatory

Lindegren L., 2010, in S. A. Klioner, P. K. Seidelmann, & M. H. Soffel ed., IAU Symposium Vol. 261 of IAU Symposium, Gaia: Astrometric performance and current status of the project. pp 296–305

Lindegren L., Babusiaux C., Bailer-Jones C., Bastian U., Brown A. G. A., Cropper M., Høg E., Jordi C., Katz D., van Leeuwen F., Luri X., Mignard F., de Bruijne J. H. J., Prusti T., 2008, in W. J. Jin, I. Platais, & M. A. C. Perryman ed., IAU Symposium Vol. 248 of IAU Symposium, The Gaia mission: science, organization and present status. pp 217–223

Lindegren L., Bastian U., 2011, in EAS Publications Series Vol. 45 of EAS Publications Series, Basic principles of scanning space astrometry. pp 109–114

Lindegren L., O'Mullane W., Hobbs D., Lammers U., 2011, A& A, in preparation

Massey R., Stoughton C., Leauthaud A., Rhodes J., Koekemoer A., Ellis R., Shaghoulian E., 2010, MNRAS, 401, 371

Mignard F., 2005, in C. Turon, K. S. O'Flaherty, & M. A. C. Perryman ed., The Three-Dimensional Universe with Gaia Vol. 576 of ESA Special Publication, Overall Science Goals of the Gaia Mission. pp 5–+

Mignard F., Bailer-Jones C., Bastian U., Drimmel R., Eyer L., Katz D., van Leeuwen F., Luri X., O'Mullane W., Passot X., Pourbaix D., Prusti T., 2008, in W. J. Jin, I. Platais, & M. A. C. Perryman ed., IAU Symposium Vol. 248 of IAU Symposium, Gaia: organisation and challenges for the data processing. pp 224–230

Mignard F., Drimmel R., 2007, DPAC: Proposal for the Gaia Data Processing, GAIA-CD-SP-DPAC-FM-030-02

Nelder J., Mead R., 1965, Computer Journal, 7, 308

O'Mullane W., Banday A. J., Górski K. M., Kunszt P., Szalay A. S., 2001, in A. J. Banday, S. Zaroubi, & M. Bartelmann ed., Mining the Sky Splitting the Sky - HTM and HEALPix. p. 638

Pasquier J.-F., 2011, in EAS Publications Series Vol. 45 of EAS Publications Series, Native and irradiated Charge Transfer Inefficiency characterization. pp 61–66

Paulet P., 2009, Technical report, Video processing unit (VPU) specification. EADS Astrium

Perryman M. A. C., de Boer K. S., Gilmore G., Høg E., Lattanzi M. G., Lindegren L., Luri X., Mignard F., Pace O., de Zeeuw P. T., 2001, Astronomy & Astrophysics, 369, 339

Press W., B.P. F., Teukolsky S., Vetterling W., 1992, Numerical recipes in Fortran. Cambridge University Press

Prod'homme T., 2011, in EAS Publications Series Vol. 45 of EAS Publications Series, Radiation effects on Gaia CCDs, Modelling to mitigate the threat. pp 55–60

Prod'homme T., Weiler M., Brown S., Short A., Brown A., 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 7742, Comparison of a fast analytical model of radiation damage effects in ccds with experimental tests

Rhodes J., Leauthaud A., Stoughton C., Massey R., Dawson K., Kolbe W., Roe N., 2010, Publications of the Astronomical Society of the Pacific, 122, 439

Rhodes J. D., et al., 2007, ApJ, 172, 203

Robin A. C., Reylé C., Derrière S., Picaud S., 2003, A&A, 409, 523

Robin A. C., Reylé C., Grux E., The Gaia Dpac Consortium 2009, in M. Heydari-Malayeri, C. Reylé, & R. Samadi ed., SF2A-2009: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics Simulating Gaia observations using a "Universe Model". p. 79

Schrabback T., et al., 2010, A&A, 516, A63

Seabroke G., Holland A., Burt D., Robbins M., 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 7742, Silvaco ATLAS model of ESA's Gaia satellite e2v CCD91-72 pixels

Seabroke G., Holland A., Cropper M., 2008a, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 7021, Modelling radiation damage to ESA's Gaia satellite CCDs

Seabroke G., Holland A., Cropper M., 2008b, in Dorn D. A., Holland A. D., eds, High Energy, Optical, and Infrared Detectors for Astronomy III Vol. 7021 of Proc. SPIE (Bellingham, WA, USA), Modelling radiation damage to ESA's Gaia satellite CCDs

Seabroke G. M., Holland A. D., Burt D., Robbins M. S., 2009, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 7439, Modelling electron distributions within ESA's Gaia satellite CCD pixels to mitigate radiation damage

Shockley W., Read W. T., 1952, Physical Review, 87, 835

Short A., 2007, Technical report, A physical model of electron trapping effects in Gaia (TDI mode) CCDs. ESA

Short A., Hopkinson G., Laborie A., Pouny P., Vetel C., Eaton T., Steward R., Holland A., Hutchinson I., Smith D., de Bruijne J., Gare P., Perryman M., Sarri G., Zayer I., 2005, in T. J. Grycewicz & C. J. Marshall ed., Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 5902, Gaia astrometric CCDs and focal plane. pp 31–44

Short A., Prod'homme T., Weiler M., Brown S., Brown A., 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 7742, A fast model of radiation-induced electron trapping in ccds for implementation in the gaia data processing

SIDC-team 2011, Monthly Report on the International Sunspot Number, online catalogue

Townsley L. K., Broos P. S., Garmire G. P., Nousek J. A., 2000, ApJ, 534, L139

Townsley L. K., Broos P. S., Nousek J. A., Garmire G. P., 2002, Nuclear Instruments and Methods in Physics Research A, 486, 751

van Leeuwen F., 2007a, Technical report, Dealing with CCD radiation damage in the Gaia data processing. DPAC

van Leeuwen F., 2007b, Hipparcos, the New Reduction of the Raw Data. Astrophysics and Space Science Library Vol. 350

van Leeuwen F., Lindegren L., 2007, Technical report, Outline of the Radiation Calibration Strategy for Astrometry, Photometry and Spectroscopy. DPAC

List of acronyms

Acronym	Definition
ABD	Anti-Blooming Drain
ACS	Advanced Camera for Surveys
AC	ACross scan
AF	Astrometric Field
AGIS	Astrometric Global Iterative Solution
AIV	Assembly, Integration, and Verification
AL	ALong scan
AU	Astronomical Unit
BAM	Basic-Angle Monitor
BC	Buried Channel
BP	Blue Photometer
CCD	Charge-Coupled Device
CDM	Charge Distortion Model
CEMGA	CTI Effects Models for GAia
CI	Charge Injection
CMOS	Complementary Metal-Oxide-Semiconductor
CTE	Charge Transfer Efficiency
CTI	Charge Transfer Inefficiency
DM	Demonstration Model
DOB	Diffuse Optical Background
DPAC	Gaia Data Processing and Analysis Consortium
EADS	European Aeronautic Defence and Space company
ELSA	European Leadership in Space Astrometry
EM	Engineering Model
ESA	European Space Agency
ESAC	European Space Astronomy Centre
ESTEC	European Space Research and Technology Centre
eV	electronvolt
FM	Flight Model
FOV	Field Of View
FPA	Focal Plane Array
FPR	First Pixel Response
FWC	Full Well Capacity
FWHM	Full Width at Half Maximum

GASS	GAia System Simulator
GIBIS	GAia Instrument and Basic Image Simulator
GSC	Guide Star Catalogue
HST	Hubble Space Telescope
HTM	Hierarchical Triangular Mesh
ID	Injection Diode
IG	Injection Gate
JASMINE	Japan Astrometry Satellite Mission for INfrared Exploration
JPL	Jet Propulsion Laboratory
L2	Lagrangian Point 2
LED	Light-Emitting Diode
LOS	Line Of Sight
LSF	Line Spread Function
μas	micro-arcsecond
mas	milli-arcsecond
MeV	Mega-electronvolt
MIM	Minimum Injection Method
MOS	Metal-Oxide-Semiconductor
ML	Maximum-Likelihood
MNRAS	Monthly Notices of the Royal Astronomical Society
NIEL	Non-Ionizing Energy Loss
NOVA	Netherlands Research School for Astronomy
NSL	Nominal Scanning Law
PEM	Proximity Electronics Module
PKA	Primary Knock on Atom
PSF	Point Spread Function
QE	Quantum Efficiency
RC	Radiation Campaign
RP	Red Photometer
RVS	Radial Velocity Spectrometer
SEU	Single Event Upset
SIDC	Solar Influences Data Center
SM	Sky Mapper
SBC	Supplementary Buried Channel
SRH	Shockley-Read-Hall
SSTL	Surrey Satellite Technology Limited
TDI	Time-Delayed Integration
TID	Total Ionizing Dose
TDA	Technology Demonstration Activities
VTM	Voltage-Tunable Method
WFC	Wide Field Camera
WFS	WaveFront Sensor

Table 6.8 — List of acronyms used in this paper by alphabetical order.

Nederlandse samenvatting

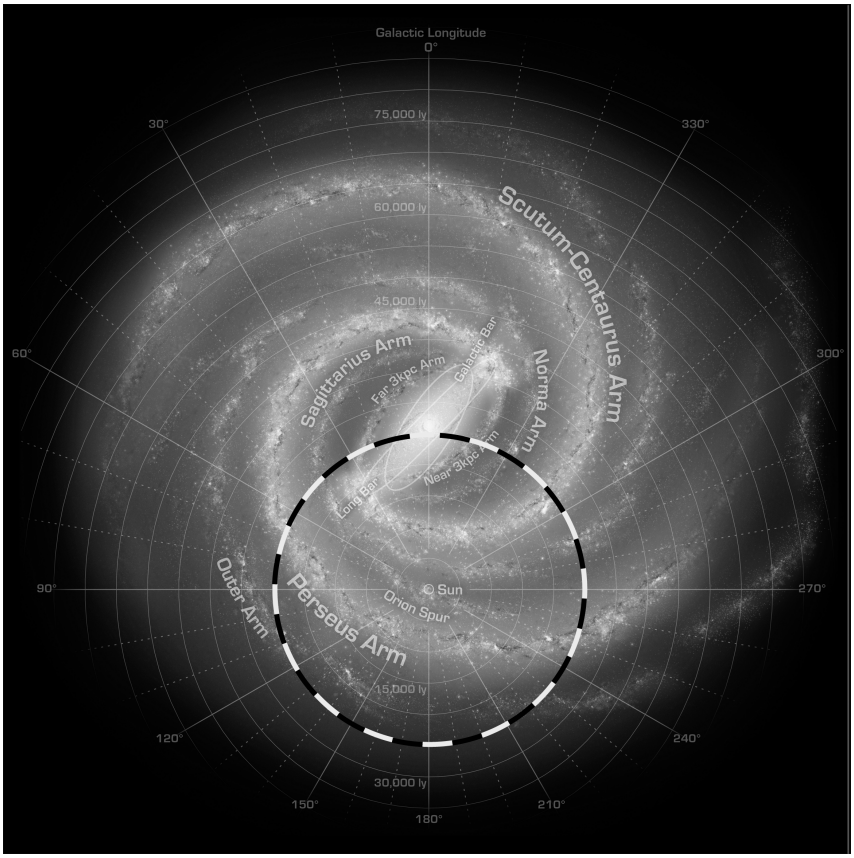
Met het vervangen van schepen door ruimtesondes en ogen door de moderne digitale tegenhanger — de CCD detector — volgen de wetenschappers van nu, die ruimtemissies bedenken, bouwen en gebruiken, in de voetsporen van de ontdekkingsreizigers uit de Renaissance. Net als hen onthullen zij de contouren van nieuwe continenten en stellen ze vragen over de ons bekende wereld. De ruimte biedt de mogelijkheid om de hele hemel waar te nemen bij alle golflengten van het licht zonder daarbij last te hebben van de verstorende effecten van de aardatmosfeer. De gewichtloosheid biedt wetenschappers ook betere controle en grotere stabiliteit van hun instrumenten. Het is dan ook niet verrassend dat veel van de recente vorderingen in onze kennis van het heelal te danken zijn aan het gebruik van ruimtetelescopen zoals de beroemde Hubble Space Telescope. Echter, de ruimte brengt niet alleen maar voordelen met zich mee. Het vacuüm, de hoge temperatuurgradiënten, en vooral de elementaire deeltjesstraling, maken de ruimte tot de meest onherbergzame omgeving die de mensheid heeft veroverd. Instrumenten aan boord van satellieten zullen onder invloed van straling van energetische deeltjes, zoals protonen afkomstig van de zon, langzamerhand slechter gaan functioneren of kunnen tijdelijk zelfs geheel uitgeschakeld worden. Het toepassen van maatregelen zoals extra bescherming tegen straling of het gebruik maken van materialen die goed tegen straling bestand zijn is vaak niet genoeg om heel nauwkeurige metingen te doen vanuit de ruimte. Het is dan ook noodzakelijk om bij het verwerken op aarde van de verzamelde gegevens rekening te houden met de effecten van de stralingsschade op de metingen. Detectoren van het type CCD zijn nu dusdanig efficiënt in de detectie van licht en de corresponderende beeldvorming dat ze het hart vormen van vele instrumenten aan boord van satellieten. Deze detectoren zijn vergelijkbaar met de detectoren die gebruikt worden in geavanceerde digitale camera's. Ze maken het mogelijk om de hemel waar te nemen, alsook de aarde en de vele lichamen in ons zonnestelsel, bij Röntgen tot infrarode golflengtes en vooral ook bij optische golflengtes. Het onderzoek beschreven in dit proefschrift is onderdeel van een wetenschappelijk programma om de effecten van stralingsschade op CCDs voor sterrenkundige toepassingen te begrijpen en de negatieve gevolgen ervan te verminderen. Mijn onderzoek vond plaats in de context van de Gaia astrometrische missie van het Europees Ruimtevaart Agentschap (ESA). Stralingsschade werd van het begin van de Gaia missie gezien als een van de belangrijkste bedreigingen voor het realiseren van de gewenste wetenschappelijke prestaties. In deze samenvatting zal ik eerst de Gaia ruimtemissie beschrijven en hoe daarmee het belangrijkste doel bereikt zal worden; het in drie dimensies nauwkeurig in kaart brengen van de plaats van de sterren in ons sterrenstelsel, de Melkweg. Ik zal dan kort uitleggen wat een CCD is en wat de effecten zijn van stralingsschade op het functioneren daarvan. Als laatste vat ik de inhoud van dit proefschrift en resultaten van mijn onderzoek samen.

De Melkweg in kaart brengen met Gaia

Waarom Gaia?
Gaia is een ruimtemissie met als doel het maken van de tot nu toe meest uitgebreide en nauwkeurige kaart van de Melkweg (zie figuur 7.1) door voor één miljard sterren de afstand tot de zon te bepalen, hun plaats aan de hemel, en hun beweging door de ruimte. We kunnen er alleen maar van dromen om ooit de gebieden in de Melkweg te bereiken die door Gaia in kaart gebracht zullen worden. Het doel van de Gaia missie is dan ook niet toekomstig ruimtetoerisme maar het doen van (wetenschappelijke) ontdekkingen. De precieze verdeling van sterren in de Melkweg is grotendeels nog onbekend en de exacte structuur en de bouwstenen van ons sterrenstelsel moeten nog ontdekt worden. We weten bijvoorbeeld dat ons sterrenstelsel, zoals vele andere, spiraalvormig is maar we weten niet precies hoeveel spiraalarmen er zijn. Een stereoscopische en dynamische kaart van de Melkweg stelt astronomen niet alleen in staat om exact de huidige structuur ervan te beschrijven maar ook om de vorming en evolutie van ons sterrenstelsel te verklaren. Bovendien zal Gaia ons in staat stellen om onze kennis van de vele soorten sterren aan te scherpen, om duizenden nieuwe buitenaardse planeten te vinden, de verdeling van asteroïden in het zonnestelsel in kaart te brengen, alsook om asteroïden te identificeren die op de aarde in kunnen slaan. Ten slotte kan Gaia ook de algemene relativiteitstheorie van Einstein verder op de proef te stellen.

De afstand van een ster tot de zon bepalen
Bij de zoektocht naar een beter begrip van ons heelal is het essentieel om de afstand te kennen die ons scheidt van de bronnen die we aan de hemel waarnemen. Het is bijvoorbeeld alleen door het meten van de afstand dat we kunnen bepalen hoe helder een ster intrinsiek is. Dit is het gevolg van het feit dat heldere sterren op grote afstand lichtzwakker lijken dan ze in werkelijkheid zijn. Echter, het meten van afstanden tussen ons (de zon) en onbereikbare objecten is een complexe aangelegenheid. De meest betrouwbare techniek die we hiervoor hebben is die van de trigonometrische parallax. Het is de meest betrouwbare techniek omdat er alleen van geometrie gebruik wordt gemaakt en er dus (bijna) geen aannames nodig zijn over ons begrip van het heelal. De parallax is de ogenschijnlijke beweging, veroorzaakt door een verandering van gezichtspunt, van een object op de voorgrond ten opzichte van objecten die verder weg in de achtergrond staan. U kunt makkelijk zelf het parallax effect waarnemen: (a) plaats een vinger vlak voor uw neus, (b) sluit één oog, (c) registreer de positie van de vinger ten opzichte van objecten in de achtergrond, (d) sluit het geopende oog en open het andere oog (verander dus van gezichtspunt), (e) zie hoe uw vinger zich lijkt te verplaatsen ten opzichte van objecten in de achtergrond. Door stappen (b) tot en met (e) kort na elkaar te herhalen kunt u de vinger zien bewegen ook al staat hij in werkelijkheid stil; het parallax effect. Door uw vinger nu verder weg te houden kunt u zien hoe de beweging van uw vinger ten opzichte van de achtergrond kleiner lijkt te worden. Hoe verder het waargenomen object hoe kleiner de parallax. Dit laatste illustreert waarom het lastig is om de afstand tot ver weg staande sterren te bepalen, hun verplaatsingen aan de hemel worden dan erg klein en dus moeilijk te meten. Een manier om de ogenschijnlijke verplaatsingen te vergroten is door de afstand tussen de

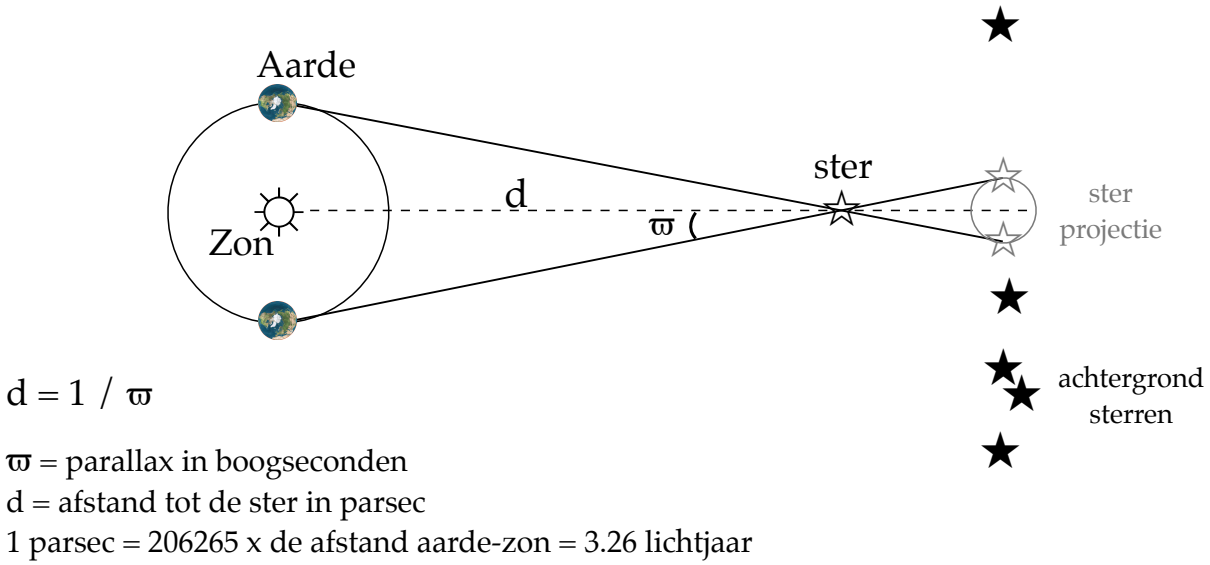
twee gezichtspunten, oftewel de basislijn, te vergroten. Sterren staan zó ver weg dat astronomen gebruik moeten maken van de omloop van de aarde rond de zon om zo hun (jaarlijkse) parallax te kunnen meten. De eerste meting in de 19e eeuw van de (allang voorspelde) parallax gebeurde ook daadwerkelijk door gedurende een jaar de posities van sterren aan de hemel te bepalen om zo gebruik te maken van de diameter van de aardbaan als basislijn. Figuur 7.2 laat het principe van deze meting zien. De eerste succesvolle parallax meting is gedaan door Friedrich Bessel in 1838 voor een van de helderste en dichtstbijzijnde sterren, 61 Cygni. De parallax wordt gewoonlijk als de hoek ϖ gegeven (zie figuur 7.2). Voor 61 Cygni is deze hoek $\varpi = 287$ milli-boogseconde (mas), ofwel minder dan $1/10\,000$ graad. Deze hoek is vergelijkbaar met de hoekmaat van de voetstap van Niel Armstrong op de maan gezien vanaf de aarde (ongeveer 150 mas). De parallax van 61 Cygni komt overeen met een afstand tot deze ster van 11.4 lichtjaar (1.07×10^{17} km!).



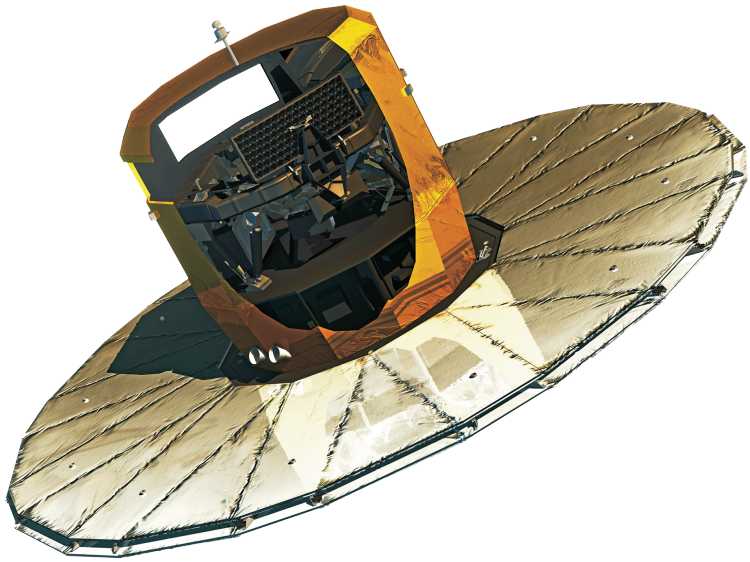
Figuur 7.1 — Schematische weergave van de verdeling van sterren in ons sterrenstelsel, Melkweg, zoals 'van boven af' gezien. De zon en de aarde bevinden zich in de buitendelen van ons sterrenstelsel vlak bij een van de spiraalarmen (midden onder in de illustratie). De gestreepte cirkel geeft het bereik aan tot waar de afstandsmetingen van Gaia nauwkeuriger dan 10% zullen zijn. De grens van dit bereik ligt op $30\,000$ lichtjaar van de zon (ofwel op 2 miljard keer de afstand aarde-zon). Buiten deze cirkel kunnen de afstanden van sterren met minder nauwkeurigheid bepaald worden zelfs tot aan de Andromeda Nevel, het dichtstbijzijnde grote spiraalstelsel. In totaal zullen één miljard sterren door Gaia waargenomen worden, één procent van het totaal aantal sterren in de Melkweg. De illustratie, die onze huidige kennis weergeeft, is gemaakt door R. Hurt (NASA/JPL-Caltech).

De noodzaak om de ruimte in te gaan
Het meten van de parallax van een ster (ofwel de afstand van de ster tot de zon) vereist: (i) een grote stabiliteit van het instrument en waarneemplatform in de tijd, omdat de metingen van een sterpositie aan de hemel herhaald moeten worden op verschillende tijdstippen om dan met elkaar verbonden te worden; (ii) precieze bepalingen van de plaats van een sterbeeld ten opzichte van het meetinstrument en de kijkrichting ervan, om zo op ieder moment nauwkeurig de plaats van de ster aan de hemel te meten; (iii) het zeer efficiënt op kunnen vangen van licht om op die manier een nauwkeurige plaatsbepaling van sterbeelden te kunnen doen en om zwakke sterren op grote afstand te kunnen waarnemen. Het waarnemen van de hemel met een ruimtesonde is de sleutel tot het voldoen aan deze drie voorwaarden. Bovendien maakt Gaia gebruik van twee telescopen, waarvan de kijkrichtingen een grote onderlinge hoek maken, het mogelijk om de parallaxen van de vele sterren aan de hemel van elkaar te onderscheiden om zo absolute parallaxen te bepalen die geen systematische fouten bevatten. Deze zogenaamde 'groothoek' techniek is alleen vanuit de ruimte mogelijk (zei hoofdstuk 1 voor meer details). Het aantal en de nauwkeurigheid van parallax metingen is gestaag toegenomen sinds de 19^{e} eeuw door het verbeteren van de meetmethoden, maar vooral door de technologische vooruitgang die gemaakt is met het detecteren en opvangen van licht. In 1989 werden de parallax metingen tot een voorlopig hoogtepunt gebracht met de lancering van de eerste satelliet bestemd voor dit soort metingen en uitgerust met fotoversterkerbuizen; de Hipparcos missie van ESA.

De Gaia missie
Gaia zal de tweede Europese ruimtemissie zijn die volledig gewijd is aan astrometrie, de kunde van het meten van posities van sterren, hun ruimtelijke bewegingen en hun parallaxen. Gaia is ontworpen met als doel het voorbijstreven van de prestaties van Hipparcos met een aantal ordes van grootte in parallax nauwkeurigheid (van millinaar micro-boogseconden) en aantal waargenomen objecten (van $100\,000$ tot 1 miljard). Het Gaia project is nu in de eindstadia van de voorbereidingen, zowel wat betreft het fabriceren van de satelliet als het voorbereiden van de gegevensverwerking, de twee grote uitdagingen van de missie. Gaia wordt in 2013 gelanceerd met een Soyuz raket vanaf de lanceerbasis vlakbij Kourou in Frans Guyana. De bestemming van de satelliet is niet een baan om de aarde maar een punt in ons zonnestelsel genaamd L2: het tweede Lagrange punt van het aarde-zon systeem. L2 ligt op een afstand van ongeveer 1.5 miljoen kilometer van de aarde op de zon-aarde as in de richting tegengesteld aan die van de zon (richting Mars zogezegd). Deze afstand is vier keer die van de aarde naar de maan. Gaia zal net als de aarde jaarlijks een baan om de zon maken gedurende de vijf jaar die de missie duurt. Gedurende deze tijd zal Gaia continu de hemel aftasten en de sterren waarnemen. Figuur 7.3 laat de Gaia satelliet zien. Alle apparatuur die nodig is voor het functioneren van de satelliet (de service module) en de wetenschappelijk metingen (de nuttige lading) bevindt zich onder de cilindervormige structuur, de thermische tent. De thermische tent bevindt zich op een zonnescherm van 10 meter doorsnede die de nuttige lading beschermt tegen direct zonlicht om zo de temperatuur stabiel te houden. De nuttige lading van Gaia bestaat uit twee telescopen en drie wetenschappelijke instrumenten die het licht naar één brandvlak leiden. Het brandvlak



Figuur 7.2 — Vereenvoudigde weergave van het principe van de afstandsbepaling voor een ster met behulp van de parallax. Deze techniek vereist het herhaaldelijk waarnemen van de positie van een ster aan de hemel gedurende een jaar. Zodoende kunnen we gebruik maken van de grootste mogelijke basislijn, de diameter van de baan van de aarde om de zon. De parallactische beweging van de ster wordt aangegeven met de kleine grijze cirkel (rechts). Het is de ogenschijnlijke beweging van de ster op de voorgrond ten opzichte van de achtergrondsterren. De parallactische beweging is periodiek en is een afspiegeling aan de hemel van de draaiing van de aarde om de zon.



Figuur 7.3 — Het Gaia ruimtevaartuig met het ontvouwen zonnescherm. De uitsnede in de thermische tent geeft zicht op (van boven naar beneden) (i) de nuttige lading met de twee hoofdspiegels die gemonteerd zijn op de mechanisch en thermisch uiterst stabiele torus van silicium-carbide, en (ii) de zogenaamde service module die het functioneren van de satelliet mogelijk maakt. De satelliet is drie meter hoog en het zonnescherm heeft een diameter van 10 meter. Illustratie met dank aan EADS Astrium.

bestaat uit 106 hoogwaardige CCDs die specifiek voor Gaia vervaardigd zijn. Deze detectoren zullen continu digitale beelden van sterren produceren, die vervolgens naar de aarde geseind worden voor een gedetailleerde analyse. Door voor iedere ster (ieder beeld) de exacte positie te bepalen kan de beweging van een ster aan de hemel gedurende 5 jaar bepaald worden. Hieruit volgen dan de plaats van de ster aan de hemel, de parallax, en de beweging van de ster door de ruimte.

De uitdaging van de gegevensverwerking voor Gaia
Het brandvlak van Gaia bevat meer dan 1 miljard beeldelementen; ook al wordt slechts een klein venster rond ieder sterbeeld doorgeseind zal de hoeveelheid data ontvangen op aarde aanzienlijk zijn. Bovendien worden de beelden op verschillende tijdstippen gemaakt voor verschillende soorten objecten (niet alleen sterren, maar ook sterrenstelsels, asteroïden, etc.) en met drie verschillende instrumenten. Vanwege de complexiteit en de omvang van de gegevensstroom zal de verwerking ervan een van de meest uitdagende onderdelen van de Gaia missie zijn. Daarom zijn meer dan 400 wetenschappers (astronomen, ingenieurs, ontwikkelaars van programmatuur) verdeeld over 24 landen al een aantal jaar bezig met het ontwikkelen van de algoritmes en programmatuur voor de verwerking van de toekomstige gegevensstroom van Gaia.
De uitzonderlijke prestaties en de verwachte oogst aan ontdekkingen met Gaia zijn grotendeels te danken aan het ruime gebruik van extreem efficiënte detectoren van licht, de CCDs. In het volgende deel wordt uitgelegd wat CCDs zijn en wat het effect van straling van energetische deeltjes op hun functioneren is.

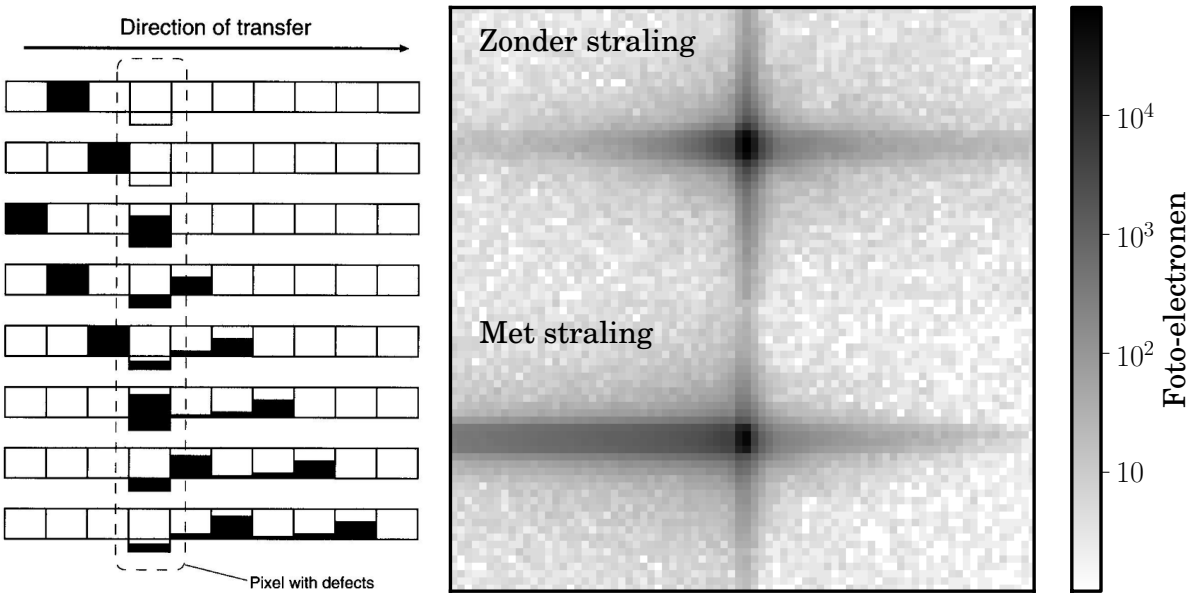
CCDs en de bedreiging van stralingsschade

Wat is een CCD?
Een CCD is een lichtsensor: een elektronisch apparaat dat licht kan detecteren en de corresponderende beelden kan digitaliseren. De CCD, ofwel 'Charge Coupled Device', werd in 1969 uitgevonden bij Bell Telephone Laboratories door Willard S. Boyle en George E. Smith. De mogelijkheid om kwantificeerbare en nauwkeurige metingen te doen van lichtdeeltjes (fotonen) bij een verscheidenheid aan golflengten bracht een revolutie teweeg in de sterrenkunde, maar ook in het dagelijks leven door het mogelijk maken van digitale fotografie. Als erkenning voor deze prestatie werd in 2009 de Nobelprijs in de natuurkunde toegekend aan Boyle en Smith.
Een CCD bestaat uit een matrix van discrete beeldelementen, pixels genaamd. De pixels bestaan uit drie lagen van verschillende materialen: (i) een geleidende laag (bijvoorbeeld een metaal), (ii) een isolerende laag, en (iii) een halfgeleider laag (meestal silicium). Het is in de halfgeleider laag dat de fotonen waar het licht uit bestaat omgezet worden in electronen. Door de voltages van het geleidende materiaal te manipuleren kunnen de electronen van een pixel in de detector naar de volgende verplaatst worden. Uiteindelijk worden de electronen naar een elektronisch onderdeel verplaats (het uitleesregister) dat in staat is om ladingspakketen in voltages om te zetten die vervolgens gedigitaliseerd worden. Op die manier wordt een digitaal beeld opgebouwd. De werking van een CCD berust dus op het vermogen om: (i) electronen te genereren uit het opvallende licht, (ii) die electronen te verplaatsen van het punt waar ze gege-

nereerd zijn tot aan het uitleesregister, en (iii) de lading bij het uitleesregister te meten met zo min mogelijk ruis. Stralingsschade leidt ertoe dat een CCD minder efficiënt de gegenereerde electronen kan verplaatsen.

Wat is stralingsschade?
Met de term stralingsschade wordt het negatieve effect bedoeld van de straling van energetische deeltjes in de ruimte op de werking van elektronische apparatuur aan boord van satellieten. Alhoewel het nabootsen van het vacuüm in de ruimte met kunstmatige middelen vrijwel onmogelijk is op aarde, is de ruimte verre van leeg, vooral niet in de buurt van sterren zoals de zon. De ruimte tussen de planeten bestaat uit neutrale deeltjes, plasma's, kosmische stralingsdeeltjes, micrometeoroïden, ruimteschroot, en vooral straling. Met straling wordt meestal het hele elektromagnetische spectrum bedoeld alsook energetische subatomaire deeltjes zoals electronen, protonen en neutronen. Fotonen kunnen ook schadelijk zijn maar het zijn vooral de subatomaire deeltjes die de stralingsschade veroorzaken. Deze deeltjes zijn afkomstig van de zon en vooral van de spectaculaire uitbarstingen aan het oppervlak. De constante stroom deeltje die het planetenstelsel doorkruist wordt de zonnewind genoemd. De omslag van dit proefschrift laat een kunstzinnige impressie zien van de zonnewind en zijn ontstaan in uitbarstingen die aan de buitenkant van onze ster plaatsvinden (de rode schijf onderaan). Het magnetische veld van de aarde (blauwe schijf in het midden) beschermt ons tegen de deeltjes in de zonnewind, maar Gaia (gele symbool bovenaan) is ver van deze beschermende omgeving verwijderd en bevindt zich bij L2 middenin deze vernietigende wind. Als de energetische deeltjes, vooral protonen, botsen op de CCDs van Gaia veroorzaken zij defecten in het halfgeleider materiaal door het verplaatsen van atomen uit hun normale positie in het kristalrooster. Deze defecten kunnen electronen invangen tijdens hun verplaatsing van de ene pixel naar de andere. Het invangen van electronen is aan het toeval onderworpen en is slechts tijdelijk. Het vrijlaten van het electron, ook een toevalsproces, kan gebeuren tijdens de verplaatsing van het beeld over het defect of als het beeld al voorbij gekomen is. In het laatste geval zal een deel van het signaal verloren gaan en zal de signaal-ruis verhouding afnemen. In beide gevallen zal de herverdeling van de electronen door de defecten leiden tot een vervorming van het digitale beeld. Deze vervorming is geïllustreerd in figuur 7.4. De defecten zullen zich in de loop van de missie opeenstapelen in iedere CCD en de gevolgen van stralingsschade zullen dan ook het grootst zijn tegen het einde van missie.

De effecten van stralingsschade op de Gaia metingen
De beelden die Gaia naar de aarde stuurt zullen door astronomen gebruikt worden om de positie van sterren aan de hemel te bepalen. Hiertoe moet met zeer grote precisie de exacte plaats van het digitale beeld op de CCD bepaald worden; tot op minder dan een duizendste pixel voor heldere sterren. De sterposities aan de hemel op verschillende tijdstippen worden dan gebruikt om de echte en schijnbare (parallax) verplaatsing van een ster aan de hemel te bepalen, om daar uiteindelijk de afstand uit af te leiden. De effecten van stralingsschade aan de metingen zijn tweeledig: (i) stralingsschade leidt tot een vermindering van de signaal-ruis verhouding. Dit is in feite een verlies aan informatie. Voor heldere sterren maakt het verlies van enkele electronen niet veel uit.



Figuur 7.4 — Links: Stralingsschade vermindert het vermogen van een CCD detector om electronen te verplaatsen van een pixel naar de andere als gevolg van het ontstaan van defecten in de het kristalrooster van het silicium waar de CCD uit bestaat. Iedere rij in de illustratie komt overeen met een pixelkolom in een verschillend stadium van het proces van ladingsoverdracht. Van boven naar beneden gezien wordt een ladingspakket verplaatst van de tweede naar de negende pixel. Het vierde pixel bevat een defect waar electronen door ingevangen kunnen worden en die vervolgens bij iedere volgende ladingsoverdracht geleidelijk weer vrijgegeven worden. Dit leidt tot het karakteristieke ladingsspoor dat men ziet in beelden die gemaakt zijn met door straling beschadigde CCDs. Illustratie met dank aan J. Walder (Universiteit van Lancaster). **Rechts:** Gesimuleerde beelden van een ster waargenomen met Gaia, vóór en na het optreden van stralingsschade. De positie van de ster aan de hemel wordt foutief bepaald als met dit effect geen rekening wordt gehouden. De vervorming van het sterbeeld is gesimuleerd met behulp van het fysische Monte Carlo model beschreven in hoofdstuk 2.

Daarentegen vertegenwoordigen de verloren electronen een flink deel van het signaal van zwakke sterren. Het gevolg is dat als de stralingsschade niet tegengegaan wordt, Gaia bijna blind zal zijn voor de verst verwijderde sterren in ons sterrenstelsel.
(ii) We hebben ook gezien dat het sterbeeld vervormd wordt door de effecten van stralingsschade. Als gevolg is de informatie over de plaats van de ster op de CCD verschillend en ingewikkelder. Wordt dit niet meegenomen in de analyse van het sterbeeld, onstaat er een systematische fout bij het bepalen van de plaats van het sterbeeld op de CCD. Deze fout kan oplopen tot een tiende pixel, dus honderd keer groter dan acceptabel is. Het is dus in de eerste plaats van belang om het invangen van electronen te voorkomen door fysieke maatregelen te treffen in de apparatuur en in de tweede plaats om rekening te houden met de vervorming van het sterbeeld tijdens de analyse ervan (tegenmaatregelen in de programmatuur voor de gegevensverwerking).

Dit proefschrift

Mijn onderzoek

Stralingsschade werd al vroeg herkend als mogelijk een belangrijke bedreiging voor de wetenschappelijke resultateni van Gaia. De industriële partners in het project kregen dus de opdracht om een aantal experimentele campagnes, bestaande uit proeven met bestraalde Gaia CCDs, uit te voeren met als doel: het effect van stralingsschade op de Gaia metingen te kenmerken, het identificeren en optimaliseren van fysieke tegenmaatregelen in de apparatuur, en het ondersteunen van de inspanningen van de astronomen om stralingsschade te modelleren en om een strategie te ontwikkelen om de effecten ervan in de gegevensverwerking zoveel mogelijk te verminderen.

In deze context was mijn onderzoek voornamelijk gericht op het modelleren van de effecten van stralingsschade aan CCDs, ondersteund door de resultaten van de analyse van de experimentele gegevens. Ik heb het tot nu toe meest gedetailleerde model ontwikkeld van stralingsschade in CCDs waarmee de werking van bestraalde CCDs nagebootst kan worden (**hoofdstuk 2**). Gebruik makende van dit model kon ik ons huidige begrip van de effecten van stralingsschade bevestigen en ook verbeteren, en kon ik ook het kenmerken van de beproefde CCDs ondersteunen en daarmee een beter begrip van de experimentele resultaten realiseren (**hoofdstukken 3 en 6**). Als onderdeel van dit onderzoek heb ik ook een gedetailleerde herbeoordeling gemaakt van de te verwachten nauwkeurigheid van de metingen van Gaia door de effecten van stralingsschade mee te rekenen (**hoofdstukken 3 en 4**). Ten slotte heb ik deelgenomen aan de inspanningen om de effecten van stralingsschade tegen te gaan door het uitwerken, het toetsen en het verbeteren van een nieuwe aanpak van de analyse van sterbeelden (**hoofdstukken 3 en 5**), alsmede aan het toetsen en verkennen van de mogelijkheden van een specifieke fysieke tegenmaatregel in de apparatuur (**hoofdstuk 6**).

Hoofdconclusies

Het belangrijkste resultaat in dit proefschrift is dat de stralingsschade aan CCDs niet langer een bedreiging vormt voor de Gaia missie omdat ik heb laten zien dat het mogelijk is om de negatieve effecten tegen te gaan door het treffen van een aantal fysieke maatregelen en door het ijken van de overgebleven effecten met behulp van specifieke algoritmes. Het is echter belangrijk om te bedenken dat dit proefschrift geen punt zet achter het stralingsschade probleem voor Gaia. De beschreven algoritmes moeten nog in de gegevensverwerking van Gaia opgenomen worden. We zullen slechts in 2013 bij het ontvangen van eerste ruwe gegevens van Gaia weten of de voorspellingen in dit proefschrift correct zijn. Ik beschrijf nu hoe de resultaten van ieder hoofdstuk in dit proefschrift bijdragen aan het tot stand komen van de hoofdconclusies.

Hoofdstuk 2 beschrijft het tot nu toe meest gedetailleerde model van stralingsschade aan CCDs. Dit model geeft uitvoering aan een nieuwe benadering van het beschrijven van de ladingsverdeling binnen een CCD pixel en het berekenen van de waarschijnlijkheden van het invangen en weer vrijgeven van electronen door defecten. Deze nieuwe aanpak maakt het nabootsen mogelijk van de effecten van stralingsschade voor verschillende soorten metingen (astrometrisch en spectroscopisch) over een groot signaalbereik en in het bijzonder voor signalen van de orde van slecht een paar electronen

(fotonen). De belangrijkste les uit dit onderzoek is dat geen enkel detail verwaarloosd mag worden om te komen tot een succesvolle nabootsing van de effecten van stralingsschade bij zeer zwakke signalen. De simulaties moeten zo realistisch mogelijk zijn, zelfs tot op het niveau van het verplaatsen van electronen van een elektrode naar de andere en op het niveau van de simulatie van ieder individueel defect.
In **hoofdstukken 3 en 4** beschrijf ik twee studies die gericht zijn op het kenmerken en kwantificeren van het effect van de stralingsschade aan de CCDs op de uiteindelijke astrometrische meetnauwkeurigheid van Gaia. De eerste studie is gericht op de analyse van de sterbeelden en de tweede studie op het voortplanten van de fouten, gemaakt bij de beeldanalyse, in de keten van algoritmen die uiteindelijk leiden tot de kaart van de Melkweg. In **hoofdstuk 3** laat ik zien dat het verlies aan signaal-ruis verhouding leidt tot een onherroepelijk verlies van de theoretisch haalbare nauwkeurigheid bij het meten van de positie van het sterbeeld, en dat dit verlies kan oplopen tot 6% aan het einde van de missie. Dit onomkeerbare verlies kan alleen tegengegaan worden door fysieke maatregelen in de apparatuur. Ik laat echter ook zien dat dit soort maatregelen niet voldoende is en dat een nieuwe aanpak van de sterbeeldanalyse nodig is om systematische fouten in het meten van de plaats van het beeld te voorkomen. Ik beschrijf de nieuwe aanpak en toets de prestaties ervan. Ik laat daarbij zien dat de systematische fouten als gevolg van de vervorming van het sterbeeld met een factor tien teruggebracht kunnen worden. In **hoofdstuk 4** worden de effecten van de overgebleven fouten op de uiteindelijke kaart van de Melkweg bestudeerd. Daarbij worden de onderliggende mechanismen gekenmerkt die leiden tot een opeenstapeling van fouten volgens bepaalde patronen aan de hemel. Ik laat zien dat dit het gevolg is van de manier waarop Gaia de hemel aftast en dat deze foutenpatronen ook in de bijproducten van de astrometrische algoritmes te herkennen zijn, zodanig dat een correctie toegepast kan worden. Uiteindelijk laat ik zien dat de gewenste astrometrische nauwkeurigheid van Gaia gewaarborgd is door de combinatie van maatregelen in de apparatuur, de nieuwe aanpak van de beeldanalyse, en het gebruik maken van de bijproducten van de astrometrische algoritmes.
Hoofdstuk 5 is gericht op een demonstratie van de prestaties van de nieuwe aanpak om vervormde beelden te analyseren. Dit wordt gedaan met behulp van beelden verkregen in experimenten met bestraalde CCDs. Ik laat zien dat het inderdaad mogelijk is om systematische fouten in de plaatsbepaling van de sterbeelden met een factor tien terug te brengen zoals voorspeld in hoofdstuk 3. Echter de ijking van de nieuwe methode zou wel eens moeilijk kunnen zijn als er geen oplossing wordt gevonden voor het probleem dat we nu niet met dezelfde model parameters de effecten bij verschillende signaalniveaus kunnen beschrijven. Ten slotte geef ik een aantal suggesties om de ijking van deze nieuwe methode te vergemakkelijken.
In **hoofdstuk 6** onderzoek ik een specifiek kenmerk van de Gaia CCD pixels dat erop gericht is om de effecten van stralingsschade te verminderen, het zogenaamde aanvullende verzonken kanaal. Ik laat zien dat dit kanaal inderdaad de stralingseffecten vermindert. Helaas is door fabricagefouten het aanvullende kanaal bij sommige Gaia CCDs niet aanwezig. Ik laat echter zien dat dit alleen het geval is in één helft van de betreffende CCDs, en dat de astrometrische nauwkeurigheid van Gaia binnen de gestelde eisen blijft zolang met dit probleem rekening wordt gehouden in de beeldanalyse.

Résumé en français

En remplaçant les navires par des satellites et nos yeux par leurs extensions modernes et numériques — les capteurs CCDs — les scientifiques qui conçoivent, construisent et utilisent les missions spatiales marchent dans la trace directe des explorateurs de la Renaissance. Comme eux, ils révèlent les contours de nouveaux continents et remettent en question le monde tel que nous le connaissons. Depuis l'espace, les astronomes peuvent observer le ciel dans son entièreté, dans toutes les longueurs d'onde de la lumière et sans les effets perturbateurs de l'atmosphère terrestre. L'absence de gravité leur offre aussi un meilleur contrôle et une plus grande stabilité de leurs instruments. Il n'est donc pas surprenant qu'un nombre important d'avancées récentes dans notre compréhension de l'univers aient été permises par l'utilisation de télescopes spatiaux tels que le télescope spatial international Hubble. Cependant l'espace ne comporte pas uniquement des avantages ; le vide, des gradients de température élevés, et surtout les radiations en font l'environnement le plus hostile que l'Homme n'ait jamais conquis. Des particules énergétiques, telles que les protons émis par le soleil, peuvent dégrader lentement les performances d'instruments embarqués à bord de satellites ou subitement interrompre leur fonctionnement. Protéger les satellites avec des matériaux résistants aux radiations ou utiliser des composants électroniques "durcis" ne suffit pas pour obtenir des mesures précises depuis l'espace ; bien souvent les effets des radiations sur ces mesures doivent être pris en compte dans leur analyse sur Terre. Les capteurs de type CCD sont aujourd'hui si performants en termes d'imagerie et de détection de la lumière qu'ils sont devenus le cœur de nombreux instruments à bord des satellites. Ces capteurs sont similaires aux capteurs présents dans les appareils photographiques numériques haut de gamme. Ils permettent l'observation de la Terre, du ciel et des corps du système solaire dans des longueurs d'ondes qui vont du proche infrarouge jusqu'aux rayons X et surtout dans le visible. Cette thèse s'inscrit dans un effort en cours de la communauté scientifique qui vise à comprendre et atténuer les effets des dommages dus aux radiations dans les CCDs à applications astronomiques. Ma recherche s'est déroulée dans le contexte de Gaia, mission astrométrique de l'Agence Spatiale Européenne (ESA). Les dommages dus aux radiations ont été identifiés comme une menace pour les performances scientifiques de Gaia. Dans ce résumé, je présente tout d'abord la mission et comment elle vise à accomplir son but premier : la cartographie de la position exacte des étoiles dans notre galaxie, la Voie Lactée. Ensuite j'explique brièvement ce qu'est un capteur CCD et quels sont les effets des radiations sur son fonctionnement. Enfin je résume le contenu de cette thèse et les résultats de mes recherches.

Cartographier la Voie Lactée avec Gaia

Pourquoi Gaia ?
Gaia est un satellite qui vise à créer la plus grande et la plus précise des cartes de la Voie Lactée (voir Figure 8.1) en mesurant pour un milliard d'étoiles leur distance par rapport au soleil, leur position dans le ciel et leur mouvement propre. Cette carte s'étendra jusqu'à des régions que nous pouvons seulement rêver d'atteindre physiquement un jour. Si elle n'est pas destinée à de futurs voyageurs de l'espace, son but est bien la découverte (scientifique). La distribution des étoiles dans la Voie Lactée est à présent peu connue et sa structure exacte ainsi que les éléments qui la composent reste à découvrir. Par exemple nous savons que notre galaxie se présente sous la forme d'une spirale comme de nombreuses autres, mais nous ignorons le nombre exact de bras qui la composent. Une carte stéréoscopique et dynamique de la Voie Lactée permettra aux astronomes non seulement de décrire son état actuel mais aussi d'expliquer sa formation et son évolution. De plus elle nous permettra de préciser notre compréhension des nombreuses classes d'étoiles, de découvrir des milliers de nouvelles exoplanètes à (c.-à-d. à l'extérieure du système solaire), d'établir une carte des astéroïdes dans notre système solaire et d'identifier celles qui pourraient croiser la trajectoire de la Terre. Enfin Gaia nous permettra de tester la théorie de la relativité générale d'Einstein.

Calculer la distance d'une étoile par rapport au soleil
Dans cette quête d'une meilleure compréhension de notre Univers, il est essentiel de connaître la distance qui nous sépare des objets célestes que nous observons. Par exemple, il est impossible de connaître la luminosité intrinsèque d'un objet sans en connaître la distance car plus un objet est éloigné plus sa luminosité apparente est faible. Cependant mesurer la distance entre nous et un objet inatteignable est une tâche complexe. La technique la plus fiable à notre disposition est celle dite de la parallaxe trigonométrique. Cette technique est la plus fiable car elle repose sur des principes géométriques et non sur des hypothèses relatives à notre compréhension de l'univers. La parallaxe correspond au mouvement apparent d'un objet situé au premier plan par rapport aux objets d'arrière-plan induit par le changement de position de l'observateur. Il vous est possible d'observer l'effet de parallaxe facilement grâce à une petite expérience : (a) dans un premier temps, placez un doigt en face de votre nez à quelques centimètres, (b) fermez un œil, (c) observez la position de ce doigt par rapport à des objets d'arrière-plan, (d) fermez l'œil ouvert et ouvrez l'œil fermé (c.-à-d. changez de point de vue), (e) vous pouvez observez que la position de votre doigt a changé par rapport aux objets d'arrière-plan. En répétant les opérations de (b) à (e), vous pouvez voir votre doigt se déplacer alors qu'il reste au même endroit, c'est l'effet de parallaxe. En répétant une seconde fois cette expérience tout en augmentant lentement la distance entre votre doigt et votre nez, vous observerez que l'amplitude du mouvement apparent de votre doigt se réduit. Plus l'objet est éloigné, plus l'effet de parallaxe s'atténue. Ceci illustre la difficulté à mesurer la distance d'objets très éloignés en utilisant leur parallaxe. Il est donc nécessaire d'augmenter l'amplitude de l'effet de parallaxe. Une façon d'augmenter cette amplitude est d'accroître la ligne de base c.-à-d. la distance entre deux positions de l'observateur. En observant tout au long de l'année la position

d'une étoile dans le ciel, les astronomes utilisent le voyage annuel de la Terre autour du soleil pour obtenir la plus grande ligne de base à notre disposition : deux fois la distance Terre-soleil (voir Fig. 8.2). Le mouvement apparent alors décrit par l'étoile dans le ciel est appelé la parallaxe stellaire (ou parallaxe annuelle). La Fig. 8.2 explique le principe de la mesure de la parallaxe stellaire pour obtenir la distance d'une étoile par rapport au soleil. Pourtant prédite de longue date, ce n'est qu'au début du XIX siècle que les astronomes ont, pour la première fois, détecté la parallaxe stellaire. Friedrich Bessel est le premier à avoir mesuré une parallaxe stellaire, celle d'une étoile très lumineuse et aussi l'une des plus proches du Soleil, l'étoile 61 du Cygne. La parallaxe est généralement exprimée par un angle : ϖ (voir Fig. 8.2). Pour l'étoile 61 du Cygne située à 11 années-lumière du soleil (1.07×10^{17} km !), ϖ est égale à 287 millisecondes d'arc (mas) c.-à-d. moins qu'un dix millième de degré. Cette valeur est comparable à la taille angulaire de l'empreinte de Niel Amstrong sur la lune vue de la Terre ($\varpi = 150$ mas).

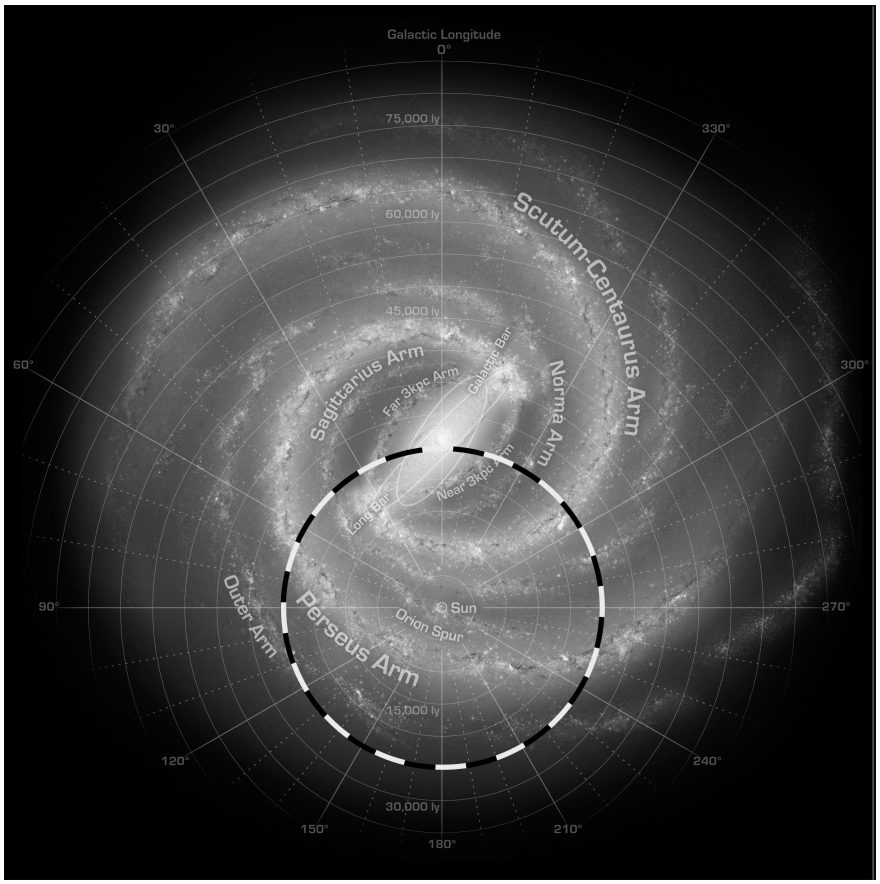


FIGURE 8.1 — Vue d'artiste de la distribution des étoiles dans la Voie Lactée telle qu'elle pourrait être observée vue d'en haut. Le soleil et la Terre se situent en périphérie dans un des bras de la spirale (au centre en bas). Le cercle en pointillé représente la limite au delà de laquelle la messure de la distance n'est plus assez précise (erreur supérieure à 10%). Cette limite est située à 30 000 années-lumière du soleil (presque 2 milliards de fois la distance Terre-soleil). Cependant la position des étoiles peut encore être mesurée jusqu'au sein même de la galaxie la plus proche : Andromède. En tout un milliard d'étoiles seront observées par Gaia, c'est à dire un pour cent des étoiles de notre galaxie. Cette illustration, représentative de nos connaissances actuelles, est fournie par R. Hurt (NASA/JPL-Caltech).

Pourquoi aller dans l'espace ?
La mesure de la parallaxe stellaire (c'est à dire de la distance d'une étoile par rapport au soleil) peut être effectuée à condition d'avoir :

1. une très grande stabilité d'instrument et de plateforme d'observation dans le temps, car les mesures de la position de l'étoile doivent être répétées à différentes époques et connectées ensemble par la suite,
2. une détermination précise de l'emplacement de l'image de l'étoile par rapport à l'instrument et de la direction de pointage de l'instrument pour connaître la position exacte de l'étoile dans le ciel,
3. une grande efficacité de détection et de collecte de la lumière afin d'obtenir une détermination précise de l'emplacement de l'image et de pouvoir observer les objets très lointains et donc peu lumineux.

Observer le ciel depuis l'espace grâce à un satellite est la clef pour satisfaire ces trois conditions. De plus il est possible, en utilisant deux télescopes dont le pointage est séparé d'un angle important, de dissocier la mesure de la parallaxe de l'observation des objets d'arrière plan et ainsi d'obtenir une parallaxe absolue qui s'affranchit de nombreuses erreurs. Cette dernière technique dite "à grand angle" n'est possible que depuis l'espace (voir l'introduction de cette thèse Chapitre 1 pour plus de détails). Le nombre de parallaxes stellaires mesurées a constamment augmenté depuis le XIX siècle grâce au perfectionnement des techniques utilisées, et surtout grâce aux bonds technologiques successifs dans la détection et la captation de la lumière. En 1989, la mission Hipparcos de l'ESA a porté à leur apogée les mesures de parallaxes stellaires de très hautes précisions grâce à la première utilisation d'un satellite consacré à ces mesures, embarquant des capteurs de type photomultiplicateur.

Le satellite Gaia
Gaia sera le deuxième satellite Européen voué à l'astrométrie. L'astrométrie est la discipline de l'astronomie dédiée à la mesure de la position des étoiles, de leur mouvement dans l'espace et de leur parallaxe. Le satellite Gaia a été conçu dans le but d'améliorer les résultats d'Hipparcos en termes de précision de mesure de parallaxe (de l'ordre de la microseconde d'arc au lieu d'une milliseconde d'arc pour Hipparcos) et en termes de nombre d'étoiles observées (un milliard d'étoiles au lieu d'un million pour Hipparcos). Actuellement le projet Gaia est en phase finale de préparation autant du point de vue de la construction du satellite que celui du développement des logiciels de traitement des données. Gaia sera lancé dans l'espace en 2013 par une fusée Soyouz depuis la base de lancement située près de Kourou en Guyane. La destination finale du satellite n'est pas en orbite autour de la Terre mais à un emplacement du système solaire appelé L2 : le deuxième point de Lagrange du système Terre-Soleil. L2 est situé à environ 1.5 millions de kilomètres de la Terre, sur l'axe Terre-Soleil dans la direction opposée au soleil (vers Mars en quelque sorte). Cette distance correspond à quatre fois la distance Terre-lune. Une fois L2 atteint, Gaia orbitera autour du Soleil au même rythme que la Terre et ce pendant au moins 5 ans. Durant ce temps, Gaia va balayer (scanner) le ciel et observer les étoiles sans interruption. La Figure 8.3 est une illustration représentant Gaia ; tous les équipements nécessaires, soit à son fonctionnement

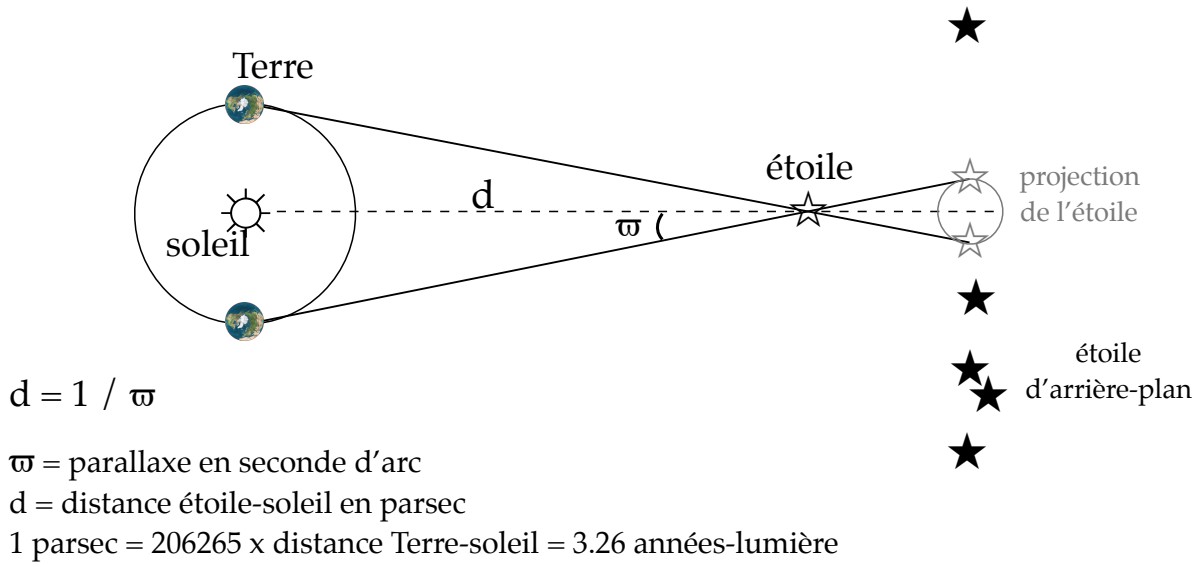


FIGURE 8.2 — Schéma représentant les principes de la mesure de la distance d'une étoile utilisant la parallaxe stellaire. Cette technique requiert des observations répétées de la position d'une étoile au cours de l'année. Ainsi la plus grande ligne de base à notre disposition, le diamètre de l'orbite terrestre autour du soleil, peut être utilisée. Le mouvement parallactique est illustré par le petit cercle gris (à droite) : c'est le mouvement apparent de l'étoile dont la distance est mesurée par rapport aux étoiles situées en arrière-plan. Ce mouvement reflète l'orbite terrestre telle qu'elle serait vue depuis l'étoile.

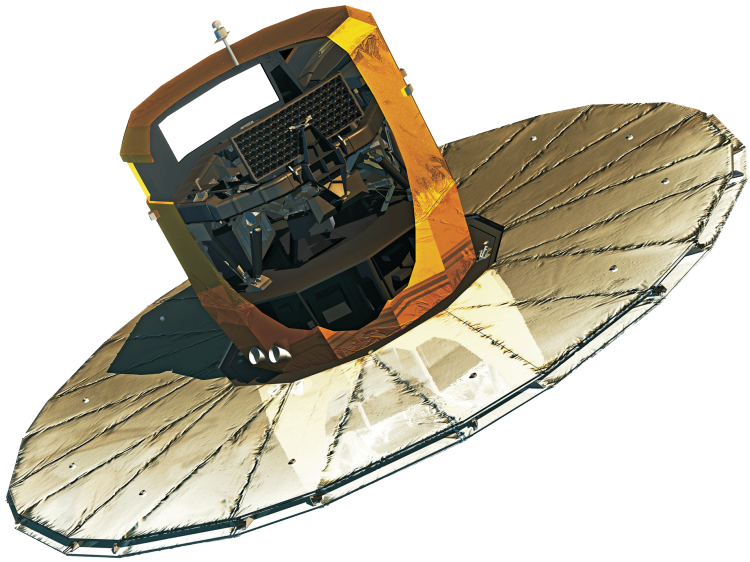


FIGURE 8.3 — Illustration représentant le satellite Gaia avec son bouclier solaire déployé. La coupe dans la tente thermique permet d'entrevoir (de haut en bas) : (a) la charge utile et ses deux miroirs primaires montés sur un tore de carbure de silicium mécaniquement et thermiquement ultra stable, et (b) le module de service qui sert au fonctionnement basique du satellite. Le satellite mesure environ trois mètres de hauteur et, une fois déployé, le bouclier solaire atteint les dix mètres de diamètre. Illustration fournie par EADS Astrium.

basique (le module de service), soit aux mesures scientifiques (la charge utile), sont positionnés à l'intérieur de la tente thermique : le cylindre du haut. Cette tente est située sur un bouclier solaire qui protège la charge utile de la lumière directe du soleil et la maintient thermiquement stable. La charge utile de Gaia est composée de deux télescopes et trois instruments qui dirigent la lumière des étoiles vers un plan focal unique. Ce plan focal est lui-même composé de 106 capteurs CCDs haute performance et fabriqués sur mesure pour Gaia. Ces capteurs produiront des images numériques de chaque étoile observée, par la suite envoyées sur Terre pour une analyse poussée. En calculant la position exacte des étoiles à chaque observation (pour chaque image), le mouvement des étoiles dans le ciel au cours des 5 années d'observation peut être ainsi reconstruit.

Le défi du traitement des données de Gaia
Le plan focal de Gaia contient plus d'un milliard de pixels, même si seulement une petite fenêtre autour de chaque étoile est transmise, la quantité de données reçues sur Terre va être considérable. De plus les images seront prises à différentes époques, pour différents types d'objets (pas seulement des étoiles, mais aussi des galaxies et des astéroïdes) et par trois types d'instruments. La complexité et la quantité des données reçues font du traitement des données l'un des défis de la mission. C'est pourquoi plus de 400 scientifiques (astronomes, ingénieurs, développeurs de logiciel) provenant de 24 pays différents participent à la préparation des logiciels de traitement des futures données envoyées par Gaia.
Les performances extraordinaires et la récolte de découvertes qui en découleront sont rendues possible en grande partie par l'utilisation d'un grand nombre de capteurs de lumière ultra-performants appelés CCD. La prochaine partie de ce résumé est consacrée aux CCDs et à l'explication des effets des radiations sur ceux-ci.

Les capteurs CCD et la menace des dommages dus aux radiations

Qu'est-ce qu'un CCD ?
Un CCD est un détecteur de lumière : un appareil électronique qui peut détecter et numériser la lumière sous la forme d'images. Le CCD, aussi appelé "dispositif à transfert de charge", a été inventé en 1969 au laboratoire Bell Telephone par Willard S. Boyle et George E. Smith. Cette invention a révolutionné l'astronomie et la science en général en rendant possible la mesure précise du nombre de particules de lumière (photons) émis par une source lumineuse dans une variété de longueurs d'onde. Elle a révolutionné aussi le monde en y introduisant la photographie numérique. C'est en reconnaissance de ces accomplissements que Boyle et Smith ont reçu le prix Nobel de physique en 2009.
Un CCD est un réseau bidimensionnel d'éléments distincts appelés pixels. Les pixels sont composés de trois couches de matériaux de différente nature : (a) un matériau conducteur (tel qu'un métal), (b) un isolant, et (c) un matériau semi-conducteur (en général du silicium). Dans le semi-conducteur, les photons qui composent la lumière sont convertis en électrons. En appliquant des tensions sur le matériau conducteur de chaque pixel, il est possible de transférer les électrons d'un pixel à l'autre jusqu'à un

composant électronique (broche de sortie) qui permet de lire et de convertir chaque paquet d'électrons en tension afin de les encoder en bits de données (numérisation). Les performances d'un capteur CCD dépendent donc de sa capacité à : (a) transformer la lumière incidente en électrons, (b) transférer efficacement les électrons depuis leur lieu de création jusqu'à la broche de sortie, (c) lire les paquets d'électrons dans la broche de sortie avec le moins de bruit possible c.-à-d. d'informations parasites produites par les circuits électroniques. Une fois utilisé dans l'espace, le capteur CCD va subir des dommages dus aux radiations, ces dommages vont essentiellement dégrader la capacité d'un CCD à transférer efficacement les électrons.

Quels sont ces dommages causés par les radiations ?
Dans le contexte de cette thèse, les dommages dus aux radiations sont la dégradation des performances des appareils électroniques embarqués à bord des satellites, dégradation induite par les particules énergétiques provenant de l'espace. Bien qu'il soit presque impossible d'atteindre le vide spatial par des moyens artificiels sur Terre, l'espace est en fait loin d'être vide et ce, particulièrement, à proximité d'étoiles telles que le soleil. L'espace interplanétaire est composé de particules neutres, de plasmas, de rayons cosmiques, de micrométéorites, de débris spatiaux, et surtout de radiations. Le terme radiation en science englobe le plus souvent l'ensemble du spectre lumineux ainsi que les particules subatomiques énergétiques telles que les électrons, les protons et les neutrons. La lumière peut aussi être nocive pour les composants électroniques mais c'est essentiellement les particules subatomiques qui sont la cause des dommages dus aux radiations. Ces particules ont pour origine le soleil, en particulier les éruptions spectaculaires qui ont lieu à sa surface. Le flot continuel de particules qui balaie le système solaire est appelé vent solaire. L'illustration en couverture de cette thèse, en particulier son arrière-plan, évoque ce vent en provenance des éruptions du soleil (le disque rouge du bas). Le champ magnétique de la Terre (illustrée par le disque bleu central) nous protège de ces particules et Gaia (en jaune et en haut), au deuxième point de Lagrange, se situera bien loin de cette protection subissant de plein fouet ce vent destructeur. Lorsque les particules énergétiques du vent solaire (majoritairement des protons) entreront en collision avec les CCDs de Gaia, ils vont créer des défauts dans le réseau cristallin du matériau semi-conducteur en y déplaçant les atomes de leur position de repos. Ces défauts peuvent capturer des électrons durant leur transfert d'un pixel à l'autre. La capture des électrons est aléatoire et temporaire. La libération (ou émission) d'un électron capturé, elle aussi aléatoire, peut se dérouler pendant ou après le transfert de l'image. Dans ce dernier cas, une partie du signal est donc perdue à jamais, ce qui conduit à une diminution du rapport signal sur bruit. Dans les deux cas la redistribution des électrons va déformer l'image. La Fig. 8.4 nous montre un exemple de déformation. Les défauts vont s'accumuler tout au long de la mission dans les CCDs et les effets des radiations sur les images vont donc être de plus en plus importants atteignant leur maximum en fin de mission.

Effets des dommages dus aux radiations sur les mesures de Gaia
Les astronomes vont utiliser les images d'étoiles de Gaia afin de connaître leur position dans le ciel. Pour chaque image, l'emplacement exact de celle-ci sur le CCD doit être

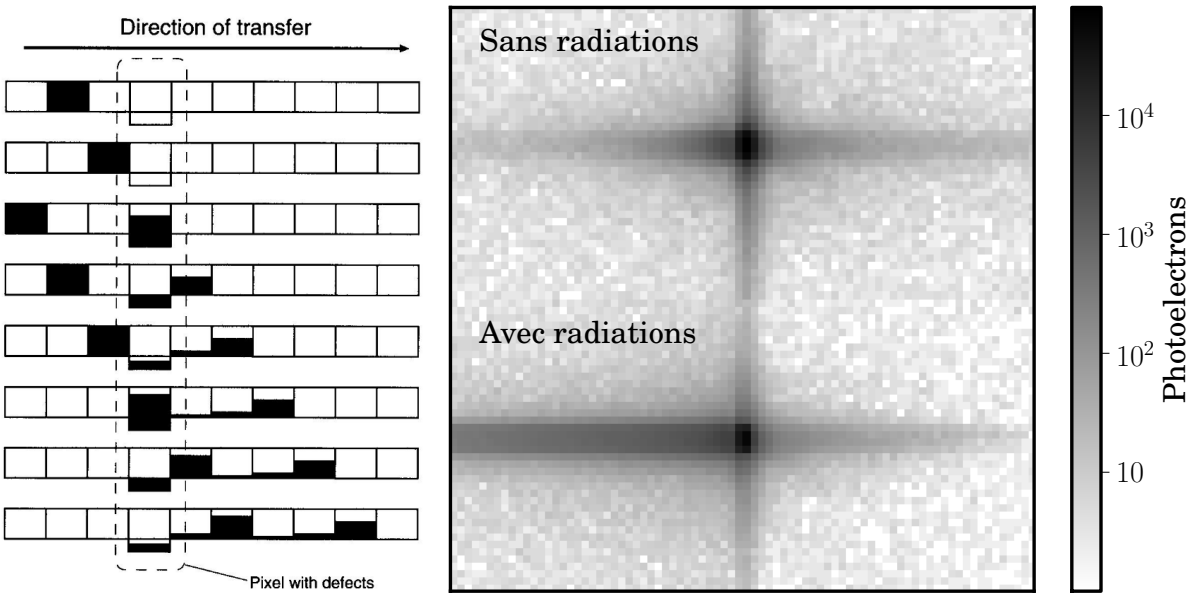


FIGURE 8.4 — A gauche : En créant des défauts à l'intérieur du CCD, les radiations dégradent sa capacité à transférer les paquets d'électrons d'un pixel à l'autre . Sur le schéma, chaque ligne correspond à une colonne de pixels à une étape différente du transfert de charges. Du haut vers le bas : un paquet est transféré du deuxième jusqu'au neuvième pixel. Un défaut est présent dans le quatrième pixel, il capture des électrons et les libère ensuite. Cela entraine la création d'une trainée d'électrons caractéristique des images acquises par des détecteurs irradiés. Illustration fournie par J. Walder (Lancaster University). **A droite :** Simulation d'une image d'étoile tel que l'observera Gaia, avant et après dommages dus aux radiations. Si cet effet n'ést pas pris en compte, la position et donc la distance de l'étoile est calculées de manière incorrecte. La déformation de l'image a été simulée en utilisant le modèle décrit au Chapitre 2.

connu avec une précision extrême ; l'erreur de localisation pour chaque observation doit être plus petite qu'un millième de pixel (pour une étoile lumineuse). Les positions d'une étoile dans le ciel (et sur le CCD) à différentes époques sont utilisées pour reconstruire le mouvement réel et apparent (parallaxe) de l'étoile en question afin d'en déduire sa distance. Les effets des dommages dus aux radiations sont de deux types : (a) comme je l'ai déjà établi, les radiations induisent une diminution du rapport signal sur bruit. Cette diminution est en fait une perte d'informations. Si une étoile est lumineuse, la perte de quelques électrons n'a pas d'importance. Par contre, pour une étoile à peine visible, ces quelques électrons correspondent à une fraction importante du signal total. Par conséquent les dommages dus aux radiations, s'ils ne sont pas évités, peuvent rendre Gaia aveugle aux objets de notre galaxie les plus lointains.
(b) Nous avons aussi vu que les radiations provoquent une déformation de l'image. Lorsqu'une image est déformée, elle transporte une information différente, plus complexe, sur la position de l'étoile dans le ciel. Si cet effet n'est pas pris en compte lors de l'analyse de l'image, une erreur systématique est introduite dans la mesure de l'emplacement de l'étoile sur le CCD. Cette erreur peut atteindre jusqu'à un dixième de pixel, c'est à dire cent fois le niveau acceptable. Il est donc primordial dans un premier temps d'éviter la capture des électrons grâce à des contremesures matérielles et dans un second temps de prendre en compte la déformation résiduelle de l'image lors de son analyse (contremesures logicielles).

Cette thèse

Ma recherche

Dès sa conception, les dommages dus aux radiations ont été identifiés comme une menace potentielle majeure pour les performances scientifiques de Gaia. Par conséquent les partenaires industriels de la mission ont reçu la consigne d'effectuer des campagnes de tests sur des détecteurs CCD de Gaia irradiés afin de caractériser les effets dus aux radiations sur les mesures de Gaia. Ces tests ont aussi été utilisés pour identifier et optimiser l'utilisation de contremesures matérielles et guider l'élaboration de modèles visant à reproduire ces effets. Ces efforts conjugués d'expérimentation et de modélisation ont permis aux scientifiques de Gaia d'élaborer une stratégie d'atténuation des dommages dus aux radiations. Dans ce contexte, ma recherche s'est concentrée sur la modélisation des effets des radiations sur les images acquises par les CCDs, soutenue par l'analyse des données expérimentales issues des campagnes de tests. J'ai conçu et développé le modèle le plus détaillé des effets des dommages dus aux radiations sur les CCDs. Ce modèle permet la simulation du fonctionnement d'un CCD irradié (**Chapitre 2**). Grâce à son utilisation, j'ai vérifié et approfondi notre compréhension actuelle des effets des dommages dus aux radiations, j'ai aussi apporté un soutien à la caractérisation des CCDs testés et permis une compréhension plus poussée des données expérimentales (**Chapitres 3 et 6**). Dans le cadre de cette recherche, j'ai mené la réévaluation détaillée des performances de Gaia en prenant en compte les effets des radiations (**Chapitres 3 et 4**). Enfin j'ai participé à l'effort commun pour contrer les effets des radiations en élaborant, testant, et améliorant une nouvelle approche d'analyse d'images (**Chapitres 3 et 5**), et en testant et explorant le potentiel d'une contremesure matérielle (**Chapitre 6**).

Conclusions principales

La conclusion principale de cette thèse est que les dommages dus aux radiations dans les CCDs de Gaia ne sont plus une menace pour la mission car j'ai démontré qu'il est possible de contrer ces effets néfastes en utilisant plusieurs types de contremesures matérielles et de calibrer les effets résiduels par l'emploi d'algorithmes spécifiques. Cependant, il est important de réaliser que cette thèse ne met pas un point final aux problèmes des radiations pour Gaia. En effet l'implémentation des algorithmes présentés n'a pas encore été réalisée dans l'ensemble de la chaîne de traitement de données. Et c'est seulement avec les premières données transmises par Gaia en 2013 que nous saurons si nos prévisions s'avèrent exactes. Ensuite je décris en quoi chaque chapitre de cette thèse a contribué à résoudre une partie du problème.

Chapitre 2 décrit le modèle le plus détaillé des dommages dus aux radiations dans les CCDs. Ce modèle utilise une approche originale de la distribution des électrons à l'intérieur d'un pixel et un nouveau calcul des probabilités de capture et libération d'électrons. Ces nouvelles approches permettent la reproduction des effets dus aux radiations sur une variété de mesures (astrométriques et spectroscopiques) pour une large gamme d'amplitude de signaux, en particulier pour des signaux constitués de quelques électrons seulement. L'enseignement le plus important apporté par ce chapitre est qu'aucun détail ne doit être négligé pour réussir à modéliser les effets des

dommages dus aux radiations sur les très faibles signaux : les simulations doivent être aussi réalistes que possible jusqu'au transfert des électrons au niveau de l'électrode dans chaque pixel et la simulation de chaque défaut dans le CCD.
Dans les **Chapitres 3 et 4**, je présente deux études qui ont pour but de caractériser et de quantifier l'impact des dommages dus aux radiations dans les CCDs sur la précision astrométrique finale de Gaia. La première étude se focalise sur l'analyse de l'image et la deuxième sur la propagation des erreurs au niveau de l'analyse de l'image dans l'ensemble de la chaîne d'algorithmes qui sera utilisée pour créer la carte de la Voie Lactée.
Dans le **Chapitre 3**, je montre que la diminution du rapport signal sur bruit induit une perte irréversible de précision théorique de localisation d'images qui peut atteindre au plus 6% (en fin de mission). La seule façon d'empêcher cette perte irréversible d'informations passe par le biais de contremesures matérielles. Néanmoins, dans cette étude, je confirme que ces solutions matérielles ne suffiront pas et qu'il est indispensable d'adopter une nouvelle approche de l'analyse des images afin d'éviter des erreurs de localisation. Je présente donc une nouvelle approche et démontre qu'en l'adoptant il est possible de diviser par dix les erreurs dues à la déformation de l'image. Dans le **Chapitre 4**, les effets des erreurs résiduelles de localisation sur la carte finale de la Voie Lactée sont étudiés ainsi que les principes qui mènent à l'accumulation d'erreurs et la formation de faux motifs sur cette carte. Je démontre que ces motifs proviennent de la façon dont Gaia balaye le ciel, et que ces motifs peuvent être identifiés dans les résidus des algorithmes. Ces résidus étant accessibles durant la mission, il est possible d'effectuer une correction en les utilisant. La conclusion finale de cette étude est qu'il est possible de préserver la précision astrométrique finale de Gaia en utilisant deux solutions matérielles distinctes, une nouvelle approche d'analyse des images déformées et les résidus des algorithmes de cartographie du ciel.
Le **Chapitre 5** démontre le potentiel de cette nouvelle approche de l'analyse des images déformées en utilisant des données expérimentales provenant des campagnes de tests. Je montre qu'il est possible de diminuer les erreurs jusqu'à un seuil acceptable mais que la calibration de cette approche restera problématique si une solution n'est pas trouvée pour simuler les effets des radiations à plusieurs niveaux de signaux avec les mêmes paramètres. Je propose des solutions pour faciliter cette calibration.
Dans le **Chapitre 6**, j'examine l'une des caractéristiques de l'architecture du pixel des CCDs de Gaia qui vise à réduire les effets des dommages dus aux radiations : le canal enterré supplémentaire. Je montre que ce canal remplit sa fonction lorsqu'il est présent, mais qu'en raison des erreurs dans le processus de fabrication du CCD, il semblerait qu'il manque dans certains des CCDs de Gaia. Je prouve aussi que lorsque ce canal est absent, il l'est uniquement dans la partie supérieure du CCD. Enfin je démontre que ce problème, s'il est pris en compte dès l'analyse de l'image, n'affectera pas outre mesure la précision astrométrique finale de Gaia.

Curriculum vitae

Born the 6th of December 1984 in Flers a small city of the north west of France, I spent most of my childhood in Normandy wandering in the countryside and occupying the numerous grey afternoons with books and computer related activities. From as far as I can remember, Isaac Asimov's popular astronomy books and a brand new Atari ST computer (a marvel at the time) my brother had received for Christmas are the most likely sparkles that triggered in me a profound interest in Science as well as a fascination for outer space and technology. Thanks to the subtle but constant support of my parents, Jean-Pierre and Isabelle, these passions are still alive today.

At the end of high-school, when the time of choosing a professional path arrived, I was still hesitating between studying physics or philosophy at the University or going to a two-year preparatory school to enter a prestigious engineering school. Over the summer 2002, I was selected in a nationwide contest to join a five-year program in one of the elite public-sector engineering schools of the INSA (National Institute for Applied Sciences) network. This offered a good compromise and my choice was made: I would go to Rennes in Brittany for five years. After two years of general courses in mathematics, physics, chemistry, electronics, and computer science, I chose the three-year specialization 'Materials and Nano-Technologies' where I studied the fundamentals of solid-state physics, material science, micro and optoelectronics. During my last year of study I also joined the MSc program in physics of the Rennes 1 University, and in June 2007 I graduated as both an Engineer in solid-state physics and a Master in physics.

During my studies I developed a particular interest in research by taking part as an intern to several projects such as the elaboration of novel satellite simulation tools at EADS Astrium Toulouse (France), the development of a new ion beam analysis method for a nuclear microprobe at the CEA Saclay (French atomic agency), and the study of radiation hard material in view of the nuclear fusion reactor ITER at the Cluj-Napoca Technical University (Romania). I also joined for one month the LATT (Toulouse-Tarbes Astrophysics Laboratory) as a graphic designer and got the chance to spend one night at the Pic du Midi observatory along with the astronomers mainly monitoring the humidity level.

In September 2007, I started a PhD in astronomy at the Leiden Observatory (The Netherlands) under the supervision of Dr. Anthony G.A. Brown and within the European Union Marie Curie research training network ELSA (European Leadership in Space Astrometry). During the past four years I focused on modelling radiation damage effects on astronomical CCDs and finding solutions to mitigate this threat to the science performance of the European Space Agency mission Gaia. My work consisted mainly in modelling, but I also participated to the analysis of tests carried out on irradiated Gaia CCDs by the Gaia main industrial partner EADS Astrium and took part to tests carried out on a non-irradiated device at the Instituto de Astrofísica de Canarias

(Spain) led by the Gaia calibration scientist Dr. Ralf Kohley (ESA/ESAC). Collaborations, visits and extended stays at several prestigious institutes across Europe have been a real catalyzer to this research in particular with Dr. Floor van Leeuwen and Dr. Scott Brown at the Institute of Astronomy in Cambridge (UK), with Prof. Dr. Lennart Lindegren and Berry Holl at the Lund Observatory (Sweden), and with Dr. Catherine Turon and Dr. Michael Weiler at the Paris Observatory in Meudon (France). I presented the results of my research in the form of talks and posters at meetings and conferences across the world in particular at the *Astronomical telescopes and instrumentation* SPIE conference in San Diego (USA), the *Gaia: at the frontiers of astrometry* ELSA conference in Paris (France), the *Milky Way and the local group - now and in the Gaia era* conference in Heidelberg (Germany), the IAU general assembly in Rio de Janeiro (Brazil) and at eight biannual Gaia Radiation Task Force meetings.

Over the past four years, I was given the opportunity to teach the exercise class of the gravitational lensing course by Prof. Dr. Koen Kuijken and to supervise the final research projects of two Master students of the Leiden Observatory dealing with CCD radiation damage in the context of Gaia and the ESA dark matter mission under study Euclid. These experiences greatly contributed to foster my education as a researcher as well as the participation to schools and workshops such as the project management course by Harrison and Associates at ESA/ESTEC (The Netherlands), the *First EIROforum school on instrumentation* at CERN (Switzerland), the *CCD modelling workshop* at Dutch Space in Leiden (The Netherlands), the *Scientific Writing for Young Astronomers* EAS school in Blankenberge (Belgium) and the ELSA schools in Leiden, Heidelberg and Barcelona (Spain). Finally I was a member of the scientific organizing committee of the *Gaia: at the frontiers of astrometry* ELSA conference.

I will defend my PhD thesis the 22nd of November 2011; in the future I intend to continue leading research and development activities related to detectors in the context of space-born experiments and space exploration.

Acknowledgments

Emilie you are the first I would like to acknowledge; without your constant love and support I would have not been able to succeed in completing this thesis. I am still amazed that four years ago when I suggested to embark on a Dutch adventure you did not think twice before saying yes with enthusiasm even if it meant giving up your current studies and professional life. Your personality, tenderness, and work as an artist are endless sources of inspiration. This thesis is also yours.

Secondly I would like to thank all my co-authors and close collaborators, it was always a pleasure to work with you and you substantially contributed to the work presented here. *Berry*, you made me discover that Science could be a great deal of fun; we make a fantastic team and I hope we will continue to collaborate on many projects. *Scott* and *Cian* thank you for answering my endless questions about Astrium's radiation campaigns and kindly providing your reduced data. *Lennart* your guidance and expertise have been essential; I learned a lot working along your side. *Michael* your subtle humour was a delight, I keep as an example your perfectionism in choosing the right words. *George* thank you for your enthusiasm, hard work, and for respecting my strict deadlines! To all the members of the *Gaia radiation task force* and in particular *Alex* and *Ralf*, thanks a lot for the feedback and nasty questions. To the *CTI dream team*, our collaboration was a success and our name is now more than ever legitimate!

The ELSA network has been one essential ingredient in my research. I first would like to acknowledge the EU and the network scientists in charge for giving me this fantastic opportunity and creating such a fostering environment. To the *ELSA fellows*, our meetings always felt like a family reunion; I truly enjoyed encountering such wonderful individualities. Our trips, parties, and passionate discussions are well placed among the highlights of these past four years. *Ester* and *Mihály* thanks for your friendship, joie de vivre, and persistent faith in my success.

A big thanks to all the people at the Leiden Observatory for making it a unique and multicultural professional environment. It was an honour to work with researchers always at the forefront of their discipline. Thanks my office mates and in particular *Michiel* for the numerous digressive philosophical chats! *Cyril* and *Erik*, thanks for smoothing my transition to the Dutch culture. *Giorgia*, *Tri*, *Daniel*, *Daniel & Carmen*, *Berni*, you are great persons I will always cherish the moments we spent together and thanks for tolerating my tempestuous temperament on a daily basis. *Atakan*, *Chael*, *Eva & Marcus*, *Jeroen*, *Michiko*, *Rafa* it was always a pleasure to share a break with you. *Hayden & Laura*, *Nikta*, I love you guys, I am looking forward to our next adventures.

Thank you the *LIsFE people*, the short film festival was the perfect break-out from this intense last year of work. *Christophe*, *Aude*, *Cheng Hua*, *Kristal*, *Wendy*, *Tilman*, *Ben & Anita*, and the *Voltron team*, thanks to you Holland felt like home.

Merci *Cyril*, *Fabien*, *Philippe & Camille*, *Hélène*, *Agathe*, *Vivi*, *Simon*, *Solenne*, *Loik* et *Matthieu* pour vos visites et votre amitié.

Merci *Maman*, *Papa*, *Jéjé* et *Lolo* pour votre amour et soutien.

Notes

STELLINGEN — PROPOSITIONS

associated with the thesis

From Electrons to Stars

Modelling and mitigation of radiation
damage effects on astronomical CCDs

1. CCD radiation damage induces an irreversible loss of accuracy in the image location measurement of at most 6 per cent in the Gaia operating conditions.

   Chapter 3

2. The use of hardware solutions to counter radiation induced charge transfer inefficiency is necessary but not sufficient to recover the required astrometric accuracy of Gaia.

   Chapters 3 and 4

3. It is possible to determine the true image location from an observation affected by charge transfer inefficiency without applying any direct correction of the raw data.

   Chapters 3 and 5

4. The Gaia charge transfer inefficiency mitigation strategy must take into account the possible lack of a functional supplementary buried channel in a Gaia CCD.

   Chapter 6

5. CCD manufacturers should pursue the development of supplementary buried channels despite the added complexity and risk in the CCD fabrication process.

   Chapter 6

6. Radiation damage is no longer a threat to the science performance of Gaia, its effects can be calibrated out.

   Chapters 3, 4, 5, 6

7. Art and science are the two driving forces of mankind's progress; budget cuts in these fields are a direct threat to our civilization.

8. In this time of environmental, economical, and financial crises, the more the immigration topic is present in a political speech the less attention the speaker deserves.

9. Societies driven by the economic growth dogma are living a short-term utopia.

10. Dictators who armed their people should always be judged together with the nations that provided the weapons used.

Thibaut Prod'homme
Leiden, 25 October 2011