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

Prod'homme, T.

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

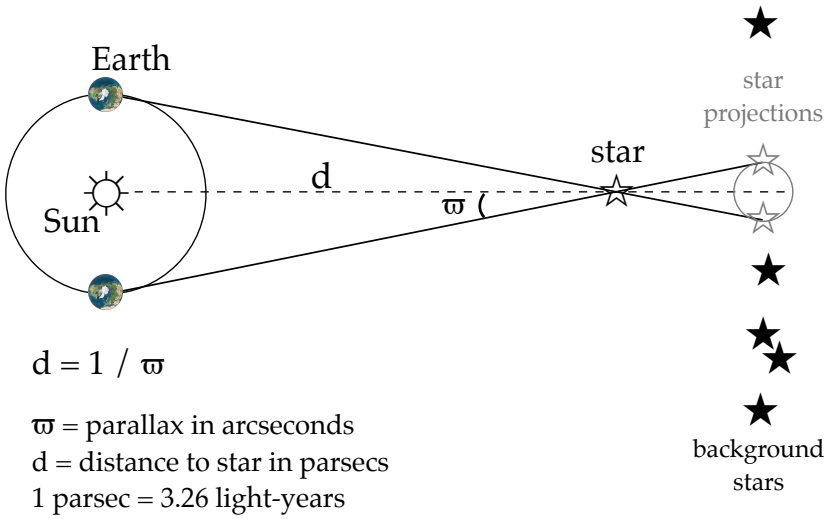

Figure 1.1 — Schematic representation of the stellar parallax and its measurement using the diameter of the Earth's as the baseline. The parallactic motion is depicted by the small grey circle. It is the apparent motion of the star of interest with respect to background (more distant) stars. The parallactic 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 $\mu_\alpha = d\alpha/dt$).

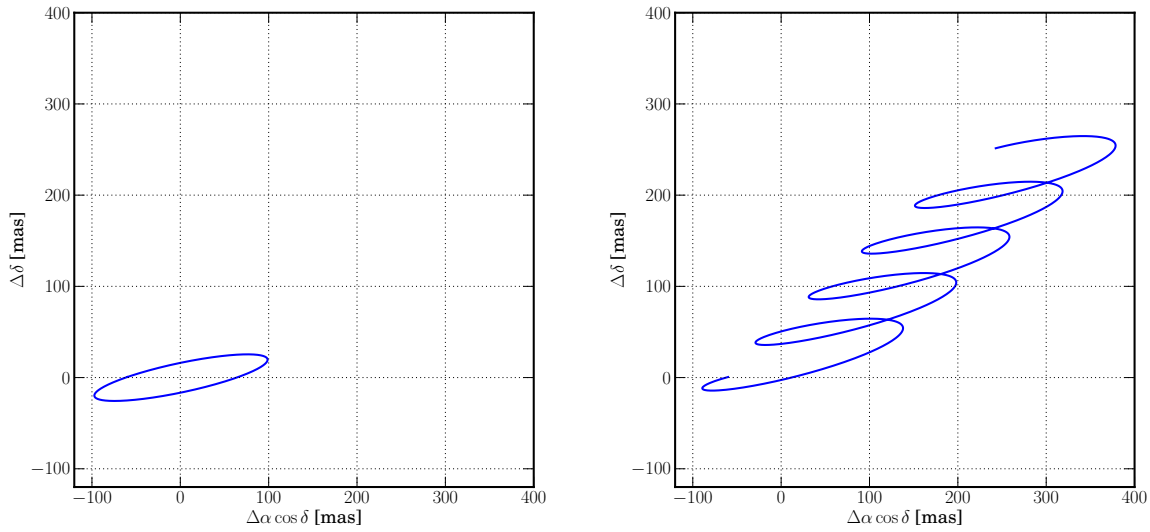

Figure 1.2 — Simulation of the motion of a star on the sky as observed from Earth in the Equatorial coordinate system. The star's astrometric parameters are: Position $(\alpha, \delta) = (60°, 30°)$, Parallax $\varpi = 100$ mas, and Proper motion $(\mu_{\alpha*}, \mu_\delta) = (60, 50)$ mas yr^{-1} **Left:** 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).

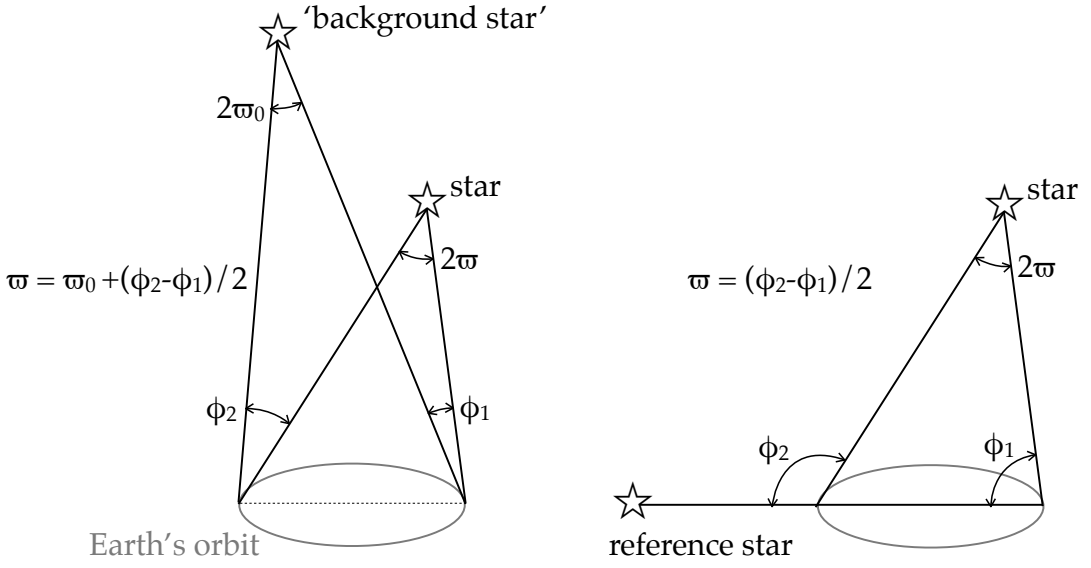

Figure 1.3 — Parallax measurement principle. **Left:** Differential small field measurement, the 'background' star is in the same field of view as the star of interest. **Right:** Wide-angle or absolute measurement, the reference star is separated by a large angle from the star of interest. This technique 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

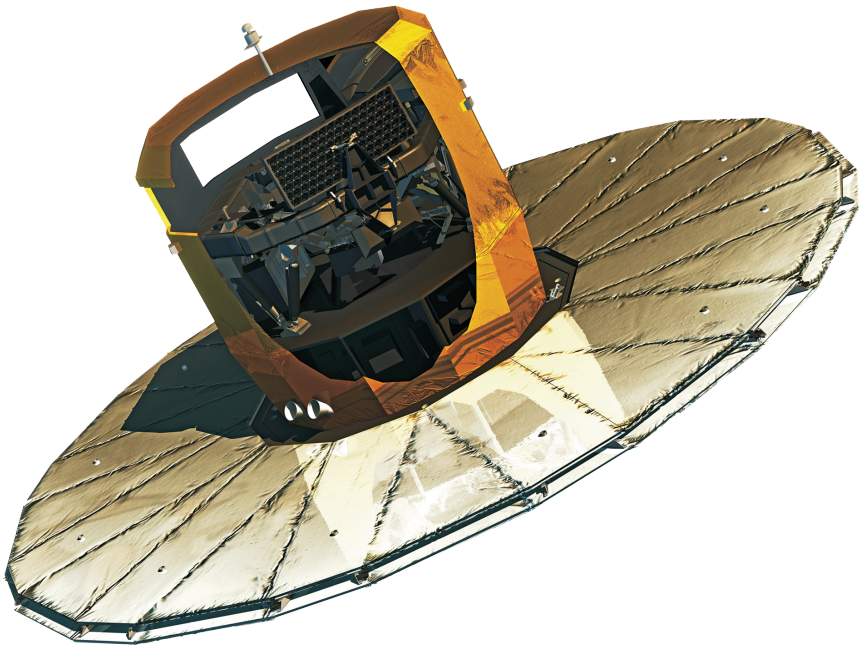

Figure 1.4 — General view of the Gaia spacecraft with the deployed sun shield. The cut-out in the thermal tent allows one to see: the payload and its two primary mirrors mounted on the mechanically and thermally stable torus (top), the service module (bottom). The satellite height is ~ 3 m and, once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

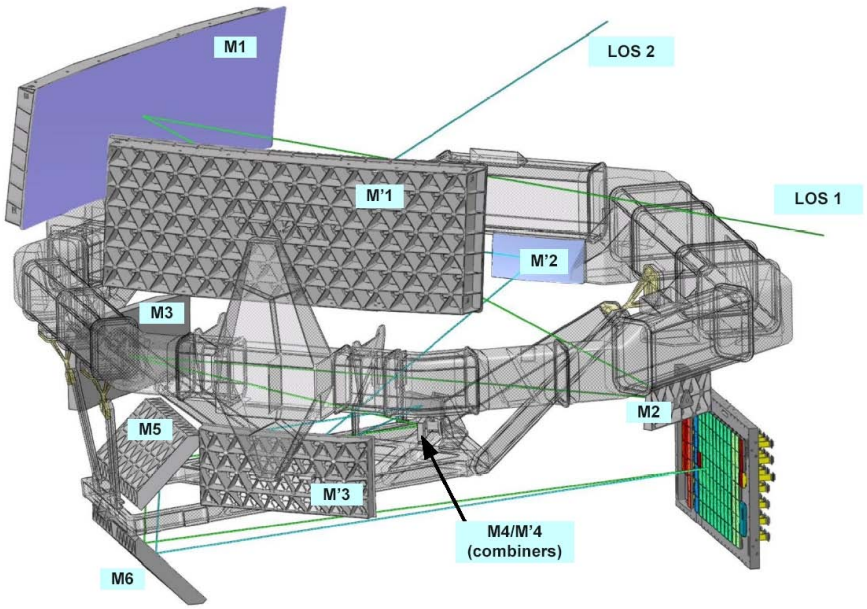

Figure 1.5 — Overview of the Gaia payload. One can recognize the multiple mirrors, including the primary mirrors M1 and M1′ (top left) of the two telescopes, the torus that supports the entire set of optical elements, and the focal plane (bottom right). The optical path for each telescope is also depicted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

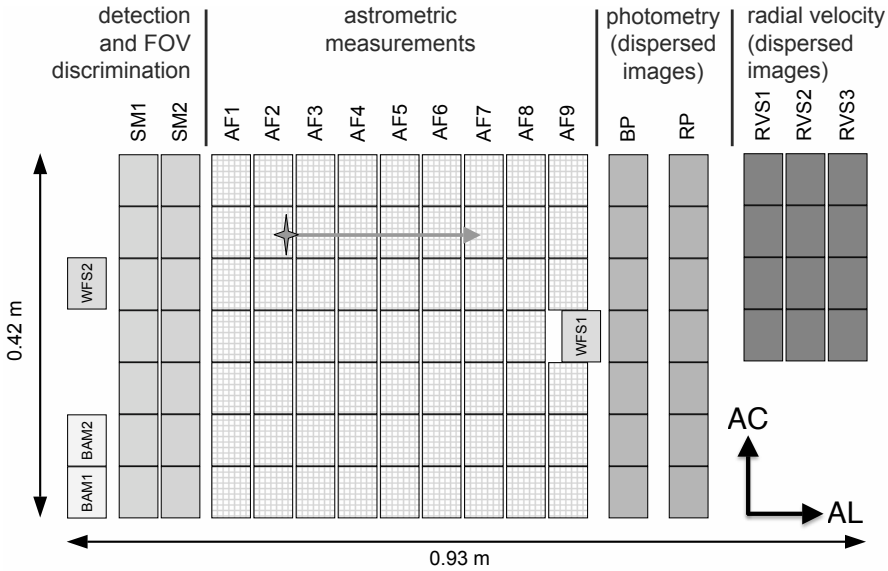

Figure 1.6 — Schematic layout of the CCDs in the focal plane of Gaia. The field of view for one telescope corresponds to 0.74 square degrees. Due to the satellite's spinning motion, a star enters the focal plane from the left and travels across it in the along-scan (AL) direction. The orthogonal direction 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-

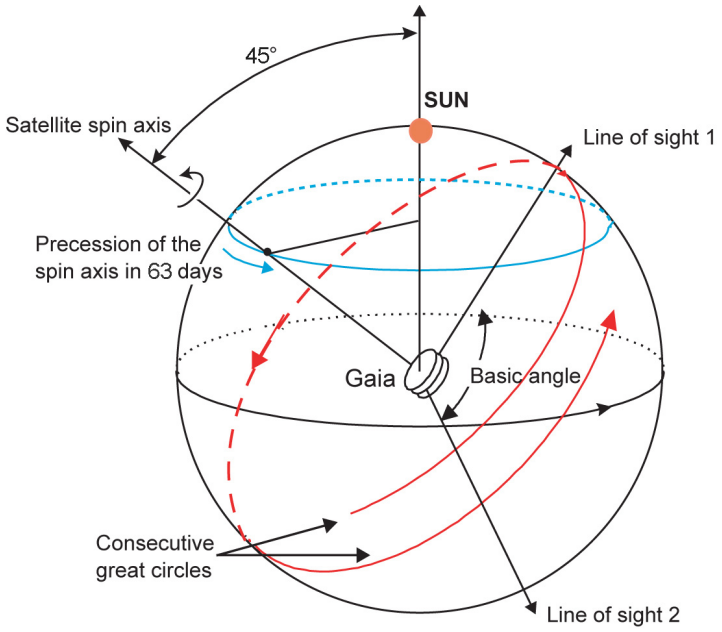

Figure 1.7 — Main principles of the Gaia scanning law. The spacecraft spin axis is perpendicular to the two fields of view and makes an angle of $45°$ with respect to the direction to the sun. The precession of the spin axis around the spacecraft-sun direction is completed in about 63 days. As a result 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

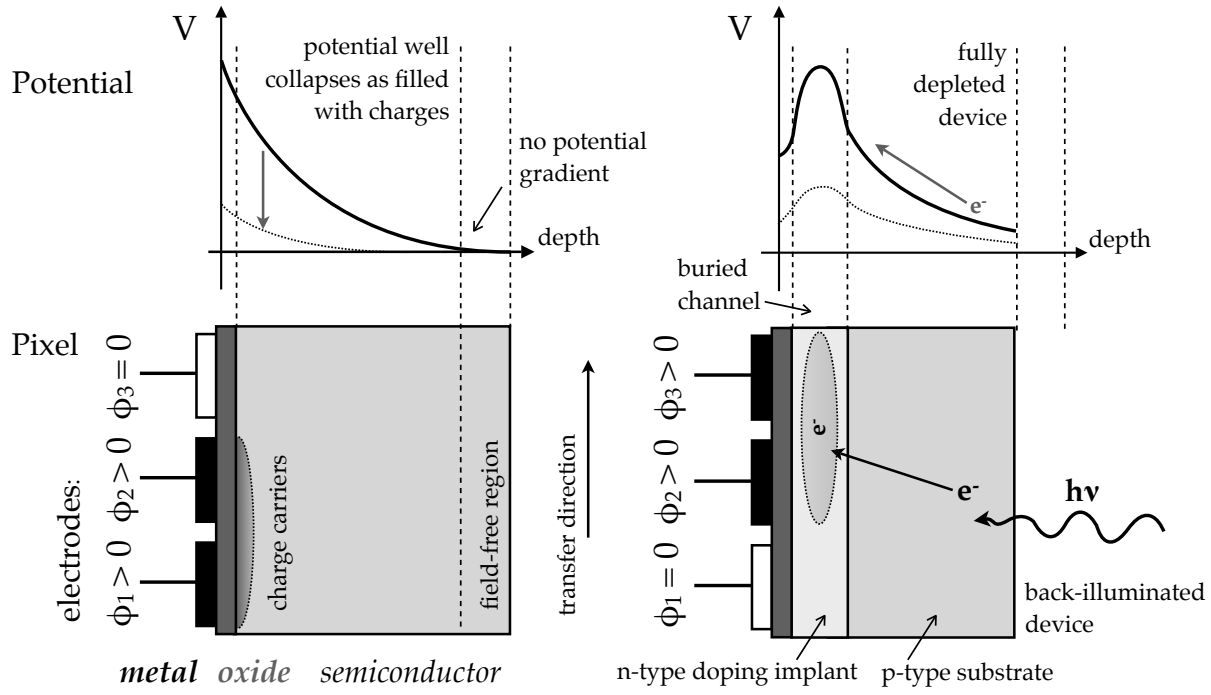

Figure 1.8 — Schematic of two different CCD pixel architectures (bottom panels) with integrating electrodes depicted in black and corresponding electrostatic potentials for these integrating electrodes (top panels). **Left:** Surface device: the CCD pixel is a simple MOS structure with a p-type substrate. 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.

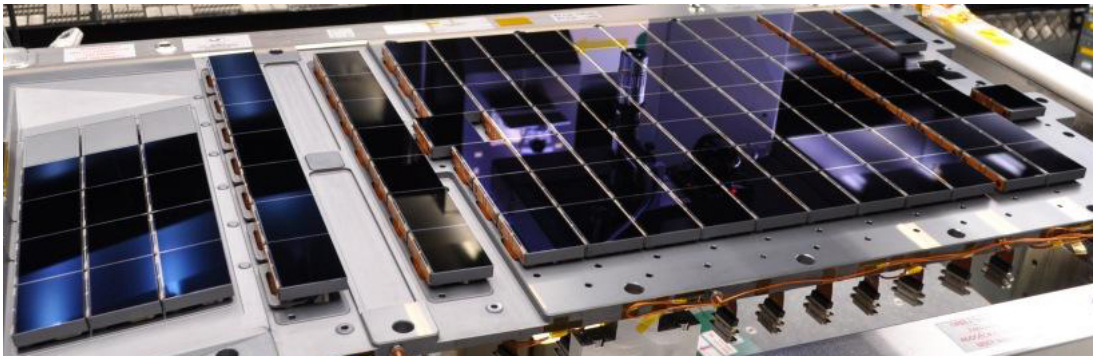

Figure 1.9 — The 106 Gaia CCDs integrated onto the CCD support structure. The columns of CCDs corresponding to the different instruments of the focal plane can be recognized: (from left to right) RVS (first three columns), RP, BP, AF9–1, SM2, SM1, and in the last column BAM1, BAM2 and WFS. A single CCD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

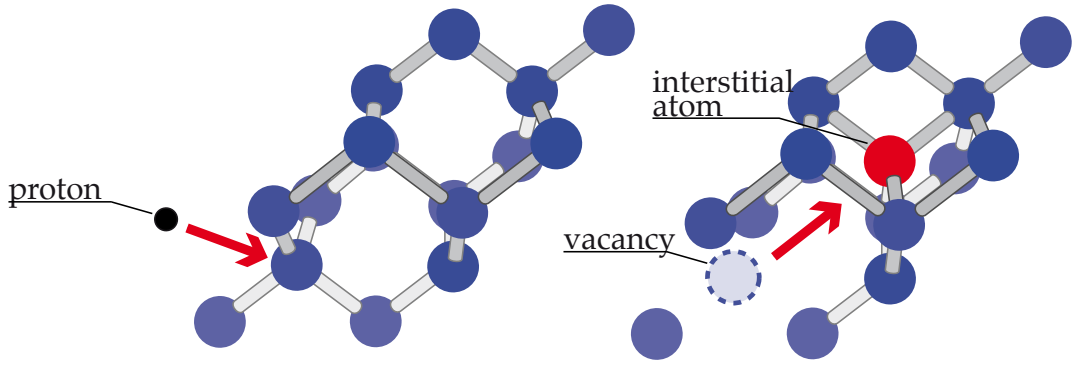

Figure 1.10 — Displacement damage in the silicon lattice. An energetic proton collides with a silicon atom (left). The energy transferred by the proton to the target atom is enough to displace it into an interstitial location (requiring more potential energy than the normal lattice location of the Si 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

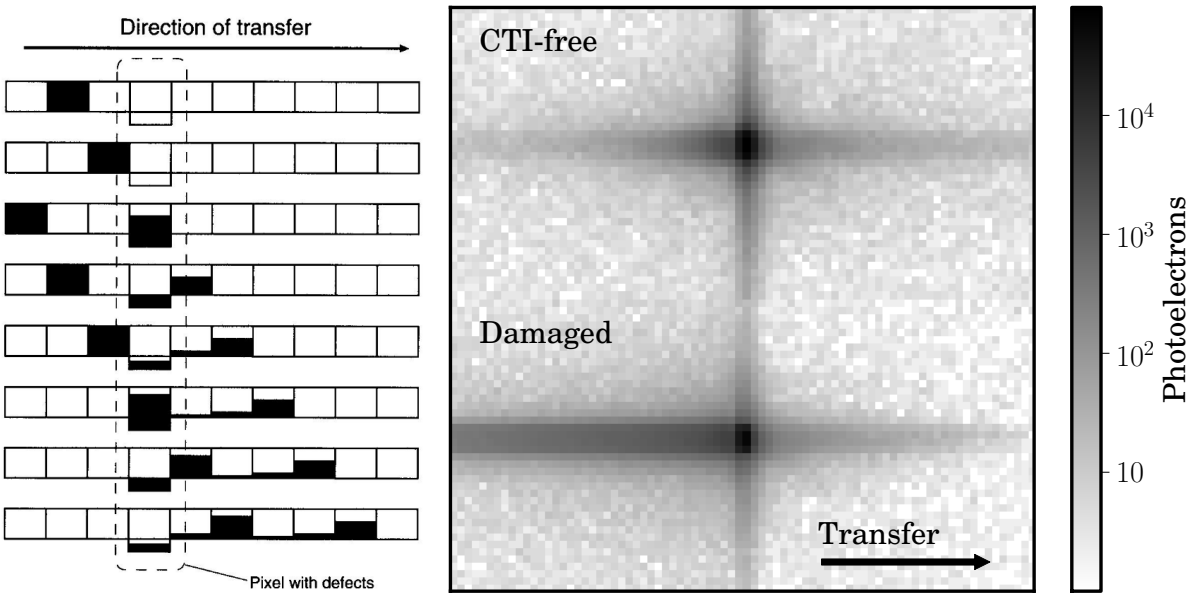

Figure 1.11 — Left: CCD Charge Transfer Inefficiency induced by bulk traps. In the schematic, rows correspond to a pixel row at different step of the charge transfer process. From top to to bottom, a charge packet is transferred from the second to ninth pixel. The fourth pixel contains bulk traps that 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.

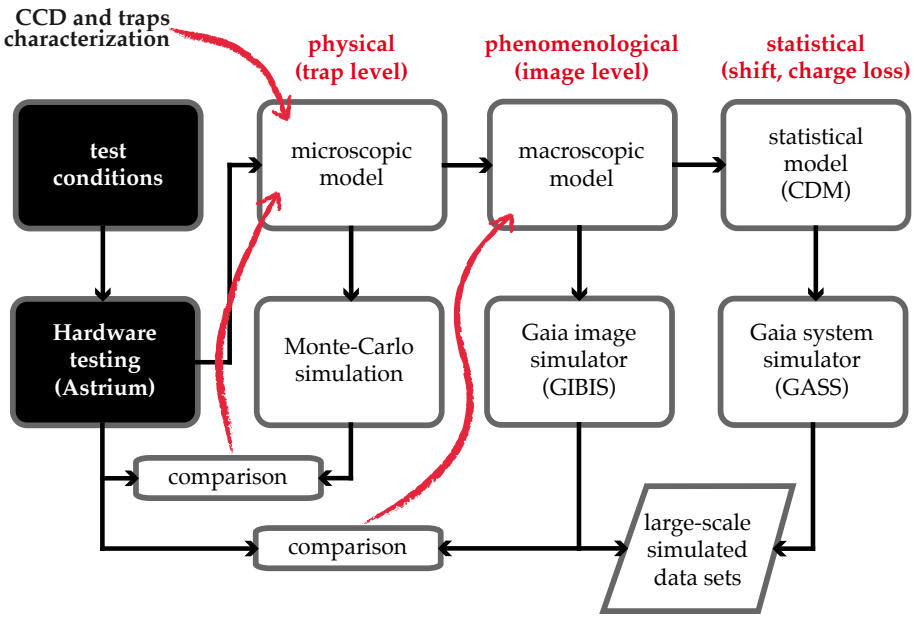

Figure 1.12 — The elaboration and implementation of the Gaia CTI mitigation strategy requires both experimental studies and modelling efforts at different levels of detail. This diagram outlines the purpose and represents the interactions between the experimental tests performed on Gaia irradiated CCDs 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.