



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).

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, biasing 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 $106^{\circ}5$ (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 \times 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 $|\beta_e| > 45^\circ$ the scan angles have a rather uniform distribution over 360° , while for $|\beta_e| < 45^\circ$ 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 $|\beta_e| < 45^\circ$ 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 $\text{mag}^{-1} \text{deg}^{-2}$ as function of G and position p . The adopted model has a spatial resolution of about 1deg^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]		$ b > 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 $5.7 \leq G \leq 12.45$ ($\sim 0.5\%$ 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 ,

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

for different fixed values of the image width w , level of background β , readout noise r , and density of active traps ρ^* .

The active trap density ρ^* 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:

$$\rho^* = (1 - \theta) \rho \quad (4.2)$$

Both theory and experiments show that the CTI effects increase monotonically with ρ^* . 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, $T_{\text{CI}} = 1$ 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_{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 ρ^* a maximum, for a CI delay equal to the CI period:

$$\rho^*(t_{\text{CI}} = T_{\text{CI}}) = \rho_{\text{max}}^* \quad (4.3)$$

Consequently, δ and σ also reach a local maximum for $t_{\text{CI}} = T_{\text{CI}}$. The global maximum values of CTI-induced bias and standard errors, δ_{max} and σ_{max} , will thus be reached for $t_{\text{CI}} = T_{\text{CI}}$ 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 $\rho^* = 1$ 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 δ_{max} and σ_{max} for the Gaia life-time and nominal CI period, $T_{\text{CI}} = 1$ s.

In Chapter 3 we also confirmed that the CTI effects on the image location are proportional to ρ^* : 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:

$$\delta = \delta_{\text{max}}(G) f_{\rho}(t_{\text{m}}) f_{\text{IH}}(t_{\text{CI}}) \quad (4.4)$$

It depends on t_{m} , the time into the science mission, and t_{CI} , the time since the previous CI, through the functions f_{ρ} and f_{IH} discussed in Sections 4.2.2.3 and 4.2.2.4. The factor f_{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

$$f_{\text{IH}}(t_{\text{CI}}) = 1 - \theta(t_{\text{CI}}) = \rho^*(t_{\text{CI}})/\rho \quad (4.5)$$

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_{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:

$$f_\rho(t_m) = \rho(t_m)/\rho_{\max} \quad (4.6)$$

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

$$\sigma = \sigma_0(G) + [\sigma_{\max}(G) - \sigma_0(G)] f_\rho(t_m) f_{\text{IH}}(t_{\text{CI}}) \quad (4.7)$$

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 ρ_{\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 $\rho_{\max}^* = 1$ trap pixel⁻¹ corresponds to the maximum density of active traps for a the nominal CI period ($T_{\text{CI}} = 1$ 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 ρ_{\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 $T_m = 5$ yr of scientific operations. Thus, $f_\rho(-0.5$ yr) = 0 and $f_\rho(T_m = 5$ 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_{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:

$$\theta(t_{\text{CI}}) = 1 - p_r(t_{\text{CI}}) = \exp(-t_{\text{CI}}/\tau_r) \quad (4.8)$$

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

$$f_{\text{IH}}(t_{\text{CI}}) = 1 - \exp(-t_{\text{CI}}/\tau_r) \quad (4.9)$$

In Chapter 3 we considered a single trap species with a release time constant $\tau_r \simeq 90$ ms. For consistency we use here the same value of τ_r .

In this computation of f_{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 assumption 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 17000 e^- . This means that our assumption holds 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:

$$\theta(t^+) = \theta(t^-) + [1 - \theta(t^-)]\phi = \phi + (1 - \phi)\theta(t^-) \quad (4.10)$$

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_{CI} (i.e., at this time after charge injection). We would then calculate $\theta(t_{\text{CI}})$, the trap occupancy level immediately before the observation of the target star, as follows:

$$\begin{aligned} \theta(t_1^-) &= \exp(-t_1/\tau_r) \\ \theta(t_1^+) &= \phi_1 + (1 - \phi_1)\theta(t_1^-) \\ \theta(t_2^-) &= \theta(t_1^+) \exp[-(t_2 - t_1)/\tau_r] \\ \theta(t_2^+) &= \phi_2 + (1 - \phi_2)\theta(t_2^-) \\ \theta(t_{\text{CI}}) &= \theta(t_2^+) \exp[-(t_{\text{CI}} - t_2)/\tau_r] \end{aligned} \quad (4.11)$$

Figure 4.5 shows an example of the evolution of $\theta(t)$ for different t_{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

$$\phi_i = \min(1, F_i/C) \quad (4.12)$$

where F_i is the maximum flux density (in $e^- \text{ pixel}^{-1}$) 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 interval¹. 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):

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

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 $T_{\text{CI}} = 1$ sec and the fact that the disturbing stars 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.

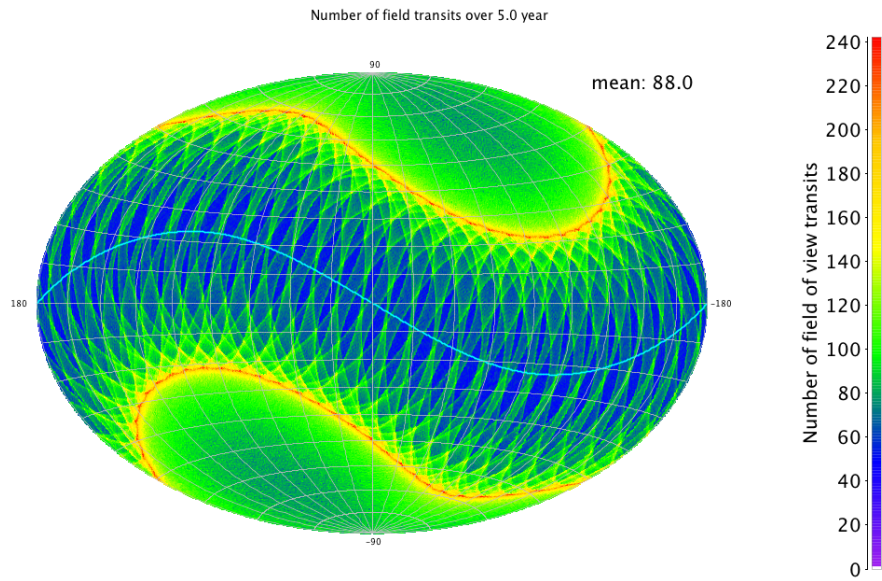


Figure 4.1 — Colour-coded map of the expected number of field of view transits experienced by sources at different celestial positions after a 5 year mission. The projection uses equatorial coordinates, with right ascension running from -180° to $+180^\circ$ right-to-left. The blue line is the ecliptic 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 — δ_{\max} and σ_{\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:** δ_{\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:** σ_{\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_{CI} since the last change injection (CI), with (solid line) and without (dotted line) disturbing stars. The traps are all filled ($\theta = 1$ or $f_{\text{IH}} = 0$) right after the CI (i.e., for short t_{CI}), and they are all asymptotically empty ($\theta = 0$ or $f_{\text{IH}} = 1$) for long CI delays. As expected from the Shockley–Read–Hall formalism, $\sim 67\%$ of the traps have released their electron after $t_{\text{CI}} = \tau_r = 90$ 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: $t_{\text{CI}} = 100, 300, 700$ 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 $t_{\text{CI}} = 700$ ms the trap occupancy is completely reset by the disturbing star at 500 ms (having $\phi = 1.0$), therefore the solid line in the top diagram is flat for $t_{\text{CI}} > 500$ ms. From the bottom diagram it is also clear that the trap occupancy level as function of t_{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 $1000 \mu\text{as}$ (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 $f_\rho(t_m)f_{\text{IH}}(t_{\text{CI}})$ 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 $f_\rho(t_m)f_{\text{IH}}(t_{\text{CI}})$ 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 $0.25\delta_{\text{max}}$. 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

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

in terms of five astrometric parameters α , δ , ϖ , μ_{α^*} , 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 . μ_{α^*} 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 $\mu_\alpha = d\alpha/dt$). 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 α^* .

Considering that the parallax and the annual proper motion components, as well as the errors in α^* and δ , are always very small angles, their effects on \mathbf{u} can be considered independently by linear superposition. The proper motion components μ_{α^*} 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 $\simeq \varpi \sin \beta_e$, where β_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.

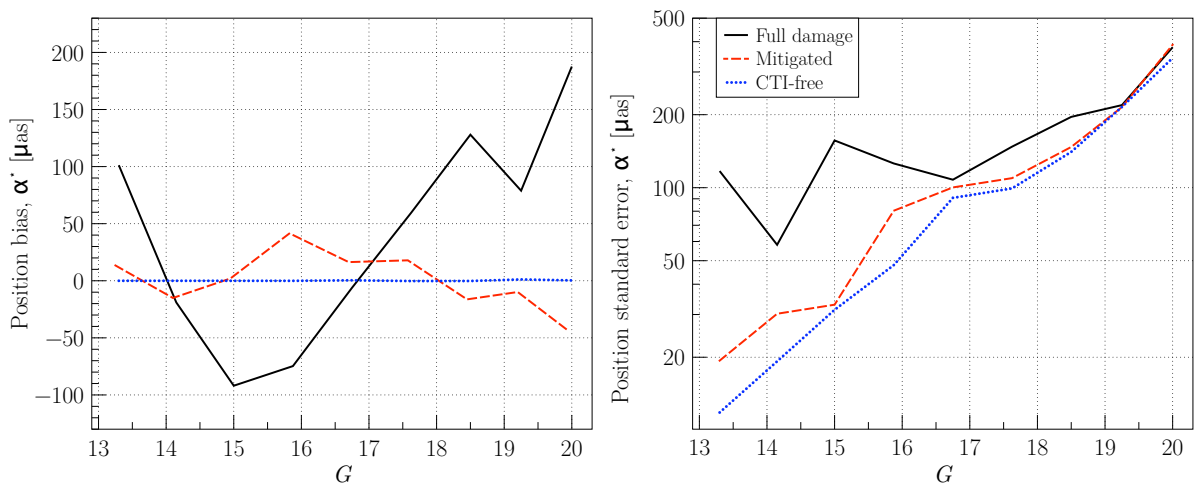


Figure 4.13 — Mean right ascension error (left) and standard error (right) per magnitude.

In principle AGIS solves the least-squares problem

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

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^{calc} , ζ_l^{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 $0.001 \mu\text{as}$ or $0.001 \mu\text{as yr}^{-1}$ 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 fluctuation (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.

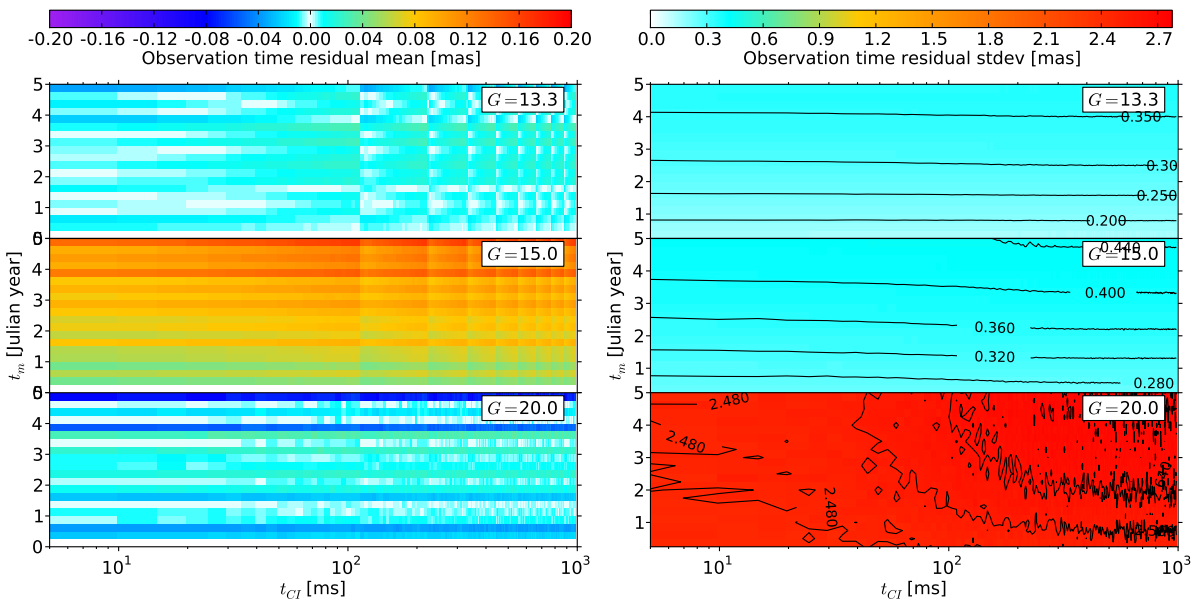


Figure 4.16 — AF AL residuals for the full damage case after applying a residual correction to the observations. Shown are the mean (left) and standard deviation (right) of the resulting residuals in 20 bins in mission time, t_m , and 200 bins in time since CI, t_{CI} . Compare this the left figure with 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.

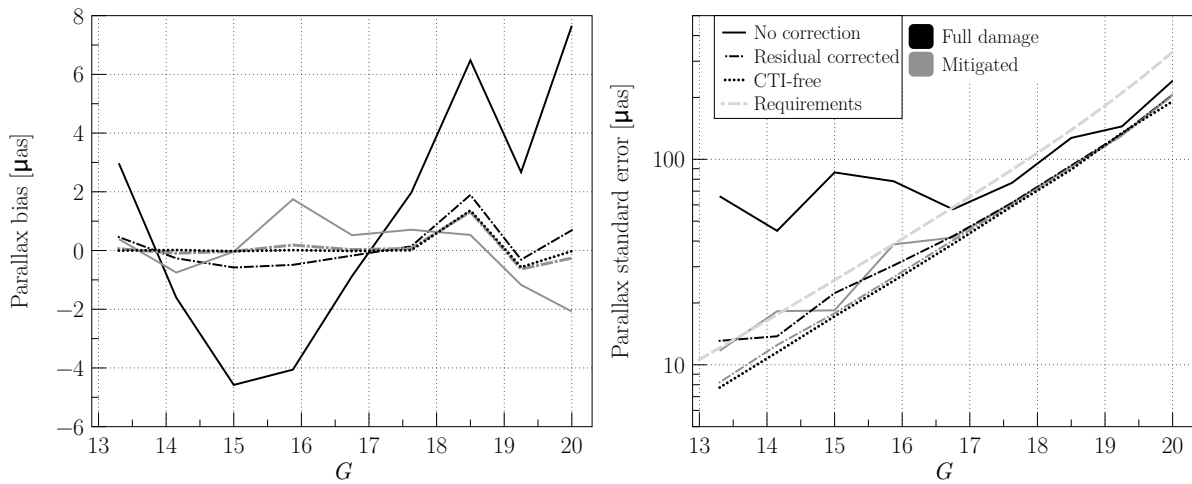


Figure 4.17 — Mean parallax error (left) and standard error (right) per magnitude, with and without the residual correction described in Section 4.3.3.




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.

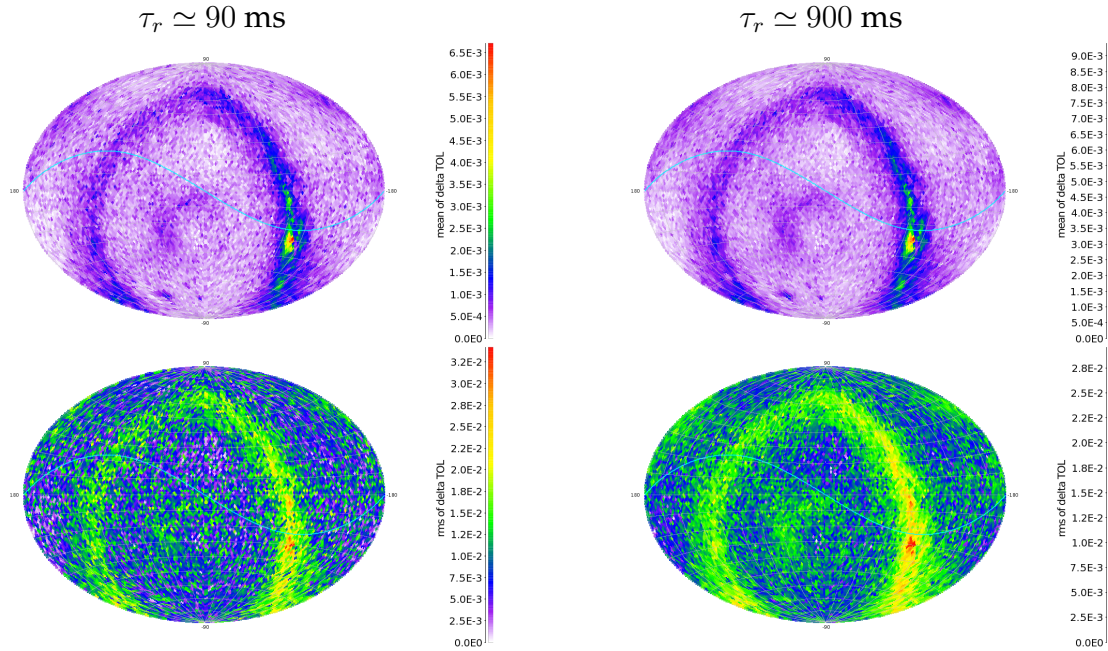


Figure 4.19 — Trap occupancy difference due to the inclusion of disturbing stars, showing the mean (top) and standard deviation (bottom) as function of position on the sky. The left maps are for the trap species with $\tau_r \simeq 90$ ms, as used in most of this investigation. The right maps are for 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., $16.7 \mu\text{s}$ 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 (α^* , 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 $|\beta_e| < 45^\circ$ 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 α^* 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 $200 \mu\text{as}$. 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 $|\beta_e| < 45^\circ$ 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 $1000 \mu\text{as}$ bias are 258 and $204 \mu\text{as}$ (in α^* and δ), $168 \mu\text{as}$ (in ϖ), and 161 and $174 \mu\text{as yr}^{-1}$ (in μ_{α^*} 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 ~ 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 $t_i^{\text{obs}} - t_i^{\text{calc}}$ 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_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 $f_\rho(t_m)$, 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 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 is 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 (~ 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 ~ 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

$$\delta = \delta_{\max}(G) f_{\rho}(t_m) \Delta f_{\text{IH}} \quad (4.16)$$

where Δf_{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 $\Delta f_{\text{IH}} \equiv -\Delta\theta \leq 0$ 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 important 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-

removal 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 led 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 account 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 $2 \mu\text{as}$. 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-

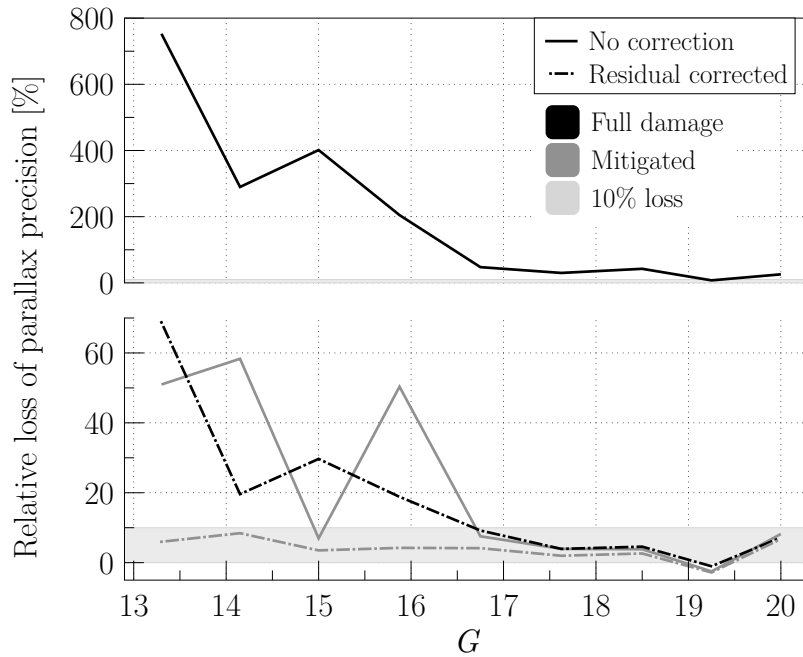


Figure 4.21 — Relative deviation in parallax standard errors with respect to the CTI-free case. This figure effectively represents the expected loss of accuracy for the full-damage (black) and mitigated (grey) case including (dashed-dotted line) or not (continuous line) a correction based on the solution 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 $\sim 4 \mu\text{as}$ (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 $0.3 \text{ e}^- \text{ pixel}^{-1}$ to $5 \text{ e}^- \text{ pixel}^{-1}$ reduces the measured charge-loss from ~ 30 per cent to ~ 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 $|\beta_e| > 45^\circ$.

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 $1 \text{ trap pixel}^{-1}$. 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 ($|\beta_e| < 45^\circ$) 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 is unacceptable especially at bright magnitudes. This once again shows that hardware counter-measures, although needed, do not 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.