



LUND UNIVERSITY

Characterization and analysis of the astrometric errors in the global astrometric solution for Gaia

Holl, Berry

2012

[Link to publication](#)

Citation for published version (APA):

Holl, B. (2012). *Characterization and analysis of the astrometric errors in the global astrometric solution for Gaia*. [Doctoral Thesis (compilation)]. Department of Astronomy and Theoretical Physics, Lund University.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Characterization and analysis
of the astrometric errors
in the global astrometric
solution for Gaia

Berry Holl

Lund Observatory
Department of Astronomy and Theoretical Physics
Lund University
Sweden

Thesis for degree of Doctor of Philosophy

Thesis advisors:
Prof. Lennart Lindegren
Dr. David Hobbs

To be presented, with the permission of the Faculty of Science of Lund University,
for public criticism in the Lundmark lecture hall (Lundmarksalen) at the Department
of Astronomy and Theoretical Physics on Friday, the 17th of February 2012, at 10.15.

Faculty Opponent

Prof. David W. Hogg
Center for Cosmology and Particle Physics
Department of Physics
New York University

Evaluation Committee

Dr. Timo Prusti
ESTEC, Noordwijk
European Space Agency

Dr. Andreas Korn
Department of Physics and Astronomy
Uppsala University

Dr. Mette Owner-Petersen
Department of Astronomy and Theoretical Physics
Lund University

Cover: A graphical summary of the thesis. It illustrates the apparent motion of a source on the sky, the detection of its light, the estimation of image parameters, and the resulting observation times being fed into the Astrometric Global Iterative Solution (AGIS). AGIS estimates the source (S), attitude (A), calibration (C) and global (G) parameters, leading to the astrometric parameters coming out. On the left side the error e in each astrometric parameter, which can be decomposed according to the used models in AGIS.

LUNFD6/(NFAS-1042)/1-190/(2012)

© Berry Holl 2012

ISBN 978-91-7473-233-7

Printed by Media-Tryck, Lund 2012

*To the wonderful people in my life,
who give meaning to my universe*

This thesis is based on the following publications:

- 1 **Spatial correlations in the Gaia astrometric solution**
B. Holl, D. Hobbs, L. Lindegren (2010)
IAU Symp., Vol. 261, *Relativity in Fundamental Astronomy: Dynamics, Reference Frames, and Data Analysis*, ed. S. A. Klioner, P. K. Seidelmann, and M. H. Soffel, pp. 320–324
- 2 **Characterizing the Astrometric Errors in the Gaia Catalogue**
B. Holl, L. Lindegren, D. Hobbs (2011)
EAS Publications Series, Vol. 45, *Gaia: at the frontiers of astrometry*, ed. C. Turon, F. Meynadier, and F. Arenou, pp. 117–122
- 3 **Efficient calculation of covariances for astrometric data in the Gaia Catalogue**
B. Holl, L. Lindegren, D. Hobbs (2012)
Workshop *Astrostatistics and Data Mining in Large Astronomical Databases*, La Palma, 30 May-3 June 2011, ed. L. Sarro, J. De Ridder, L. Eyer, and W. O’Mullane, in press
- 4 **Error characterization of the Gaia astrometric solution**
I. Mathematical basis of the covariance expansion model
B. Holl, L. Lindegren (2012)
Astronomy & Astrophysics, submitted.
- 5 **Error characterization of the Gaia astrometric solution**
II. Validating the covariance expansion model
B. Holl, L. Lindegren, D. Hobbs (2012)
Astronomy & Astrophysics, submitted.
- 6 **The impact of CCD radiation damage on Gaia astrometry**
I. Image location estimation in the presence of radiation damage
T. Prod’homme, **B. Holl**, L. Lindegren, A. G. A. Brown (2011)
Monthly Notices of the Royal Astronomical Society, in press,
DOI: 10.1111/j.1365-2966.2011.19934.x

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

B. Holl, T. Prod'homme, L. Lindegren, A. G. A. Brown (2012)
 Monthly Notices of the Royal Astronomical Society, accepted.

Publications not included in this thesis:

- i **Gaia and the Astrometric Global Iterative Solution**
 D. Hobbs, L. Lindegren, **B. Holl**, U. Lammers, W. O'Mullane (2008)
 IAU Symp., Vol. 248, *A Giant Step: from Milli- to Micro-arcsecond Astrometry*, ed.
 W. J. Jin, I. Platais, and M. A. C. Perryman, pp. 268–269
- ii **Determining PPN γ with Gaia's astrometric core solution**
 D. Hobbs, **B. Holl**, L. Lindegren, F. Raison, S. Klioner, A. Butkevich (2010)
 IAU Symp., Vol. 261, *Relativity in Fundamental Astronomy: Dynamics, Reference
 Frames, and Data Analysis*, ed. S. A. Klioner, P. K. Seidelmann, and M. H. Soffel,
 pp. 315–319
- iii **Complexity of the Gaia astrometric least-squares problem and the
 (non-)feasibility of a direct solution method**
 A. Bombrun, L. Lindegren, **B. Holl**, S. Jordan (2010)
 Astronomy & Astrophysics, 516, A77
- iv **A conjugate gradient algorithm for the astrometric core solution of Gaia**
 A. Bombrun, L. Lindegren, D. Hobbs, **B. Holl**, U. Lammers, U. Bastian (2011)
 Astronomy & Astrophysics, in press, DOI: 10.1051/0004-6361/201117904

Contents

1	The Gaia mission	1
1.1	Gaia and its scientific output	1
1.2	How Gaia works	4
1.3	The observations	6
2	Modelling the observations	11
2.1	Astrometric parameters	11
2.2	Reference systems	14
2.3	Modelling stellar motion on the sky	16
2.4	Forward modelling of the observations	16
3	Characterizing the astrometric errors	23
3.1	Errors and uncertainties	23
3.2	Why we need to understand errors	25
3.3	Correlations	26
3.4	Random errors in the Gaia data	28
3.5	Systematic errors in the Gaia data	30
4	The evolution of the thesis work	31
4.1	The history of the papers	31
4.2	AGISLab development	36
4.3	Addressing the main questions	37
5	The Papers	39
	Popular Summaries (English, Swedish and Dutch)	49
	Acknowledgements	69
	References	72

Chapter 1

The Gaia mission

This thesis is dedicated to the field of research called astrometry, which aims to measure the motions of celestial bodies with the highest possible accuracy. The Gaia mission is likely to be the most important resource for astrometric research in the near future, and the main part of the thesis is a discussion of the expected errors in the future Gaia catalogue. But let us start by describing Gaia itself and what it is supposed to produce in terms of scientific output.

1.1 Gaia and its scientific output

Gaia is an astrometric space mission that is planned for launch in 2013 by the European Space Agency (ESA). It will provide the most comprehensive and accurate catalogue of astrometric data for galactic and astrophysical research in the coming decades. For roughly 1 billion stars, quasars and other point-like objects (which we refer to as ‘sources’) having a magnitude¹ between $G = 5.7$ and 20, it will determine the five astrometric parameters, i.e., the two components of the position, the trigonometric parallax, and the two components of the proper motion. The magnitudes and colours of all the sources are simultaneously measured by the photometric instruments on board the satellite. Accuracies of 8–25 micro-arcsec (μas) are expected for the trigonometric parallaxes, positions at mean epoch and annual proper motions of sources with $G \leq 15$ and with lower accuracy down to $G = 20$. For sources with $G \lesssim 17$ radial velocities will also be measured, giving the full 6-dimensional position and velocity components. Compared with the Hipparcos

¹The Gaia G-band magnitude is a broad-band, white-light magnitude in the wavelength range 300 – 1000 nm defined by the telescope transmission and CCD quantum efficiency. $G = V$ for an un-reddened A0V star (Jordi et al. 2010; Perryman et al. 2001).




Figure 1.1: Comparison of the Hipparcos and Gaia catalogues. **Left:** The first volume (out of five) of astrometric data in the Hipparcos Catalogue, in total containing the five astrometric parameters of nearly 120 000 sources. **Right:** (Hypothetical) Gaia astrometric data, volume 1 out of 50 000, in total containing the astrometric parameters of $\sim 1\,000\,000\,000$ sources.

Catalogue (ESA 1997) the Gaia catalogue will contain roughly 10 000 times more astrometric data with 10–100 times smaller standard errors, for a mainly complementary set of fainter sources. While the main astrometric data of the Hipparcos Catalogue was contained in just five volumes which comfortably fit on a bookshelf (15 cm to be precise), the Gaia catalogue would need about 50 000 volumes (or 1.5 km of shelf) in the unlikely event that it is ever published on paper in its complete form, see Fig. 1.1. The astrometric data are complemented by photometric and spectroscopic information that will fill even larger data volumes. The resulting catalogue will become available to the scientific community around 2020. Gaia was originally envisioned as an interferometer and its name stood for *Global Astrometric Interferometer for Astrophysics* (see Lindegren & Perryman 1996). The concept soon evolved thanks to the industrial studies and the incorporation of new ideas (Høg 2008) and the current design is no longer an interferometer. Gaia has retained the name ‘Gaia’ although it is not an acronym².

Gaia is a survey mission, meaning that its observing program is not pre-defined as it was for Hipparcos, but will encompass all objects that are sufficiently point-like and bright at

²In Greek mythology, Gaia was the Earth goddess and the great mother of all, including Uranus, the sky. Thus the name is not inappropriate.




Figure 1.2: A simulation of the expected 3D distribution of the one billion galactic stars in the Gaia catalogue. In the background an artistic top view of our Galaxy (top image, copyright: NASA/JPL-Caltech/R. Hurt) and an actual picture of the Milky Way on the sky (bottom image, copyright: ESO/S. Brunier). The colours show the density of the stars, ranging from very high close to the Sun to low at greater distances. The ‘spikes’ pointing away from the Sun are due to windows in the interstellar extinction. The overlaid simulation has been produced by the DPAC-CU2 at the MareNostrum supercomputer (copyright: X. Luri & the DPAC-CU2) and are based on an adaptation of the Besançon Galaxy model (Robin et al. 2003) for Gaia. Adapted from the ESA Gaia web site at <http://www.rssd.esa.int/GAIA/>

the time of observation. This includes a very large number of solar-system objects (mainly asteroids) and extragalactic objects (e.g., quasars and supernovae), but the vast majority of objects will be stars in our own Galaxy, the Milky Way (Fig. 1.2). The scientific goals of Gaia have been reviewed in a number of publications (e.g., Perryman et al. 2001; Jordi 2011) and are comprehensively described in the proceedings of the conference *The Three-Dimensional Universe with Gaia* (Turon et al. 2005).

A unique feature of Gaia is its ability to measure trigonometric distances for a huge number of stars out to distances of many kpc. The measurement of distances is extremely important for deriving many stellar quantities like the absolute luminosity, mass, and age, but also for calibrating methods that can be used to measure distances to other galaxies (e.g., the period-luminosity relation). The astrometric data themselves contain kinematic information about the dynamical structure and formation history in our Galaxy. With the Gaia catalogue we will be able to study not only the dynamics in our own Galaxy but also in the satellite galaxies surrounding it, allowing us for example to accurately estimate their masses, disentangle their internal structure, and probe the mass distribution in the outer regions of our Galaxy. Astrometry is also very useful within our solar system. For example, by measuring precise orbits of asteroids it is possible to model close passages from which their masses can be estimated. Additionally the observations are so precise that General Relativity is needed to correctly model the direction of sources (e.g., the light deflection by the Sun is still about $4000 \mu\text{as}$ at 90° from the Sun), therefore we can also model these effects as part of the data processing and test possible deviations from General Relativity. Finally, Gaia provides an extremely accurate and dense reference frame of importance for many kinds of ground-based observations, including deep surveys and detection of near-Earth objects.

1.2 How Gaia works

The main instrument of Gaia is an optical telescope with two fields of view which are imaged on the same focal plane, see Figs. 1.3 and 1.4. Each astrometric field of view covers an approximately square area of $\sim 0.5 \text{ deg}^2$ on the sky. The two fields of view are separated by an angle of 106.5° , known as the basic angle. Having two fields of view separated by such a large angle allows Gaia to measure *absolute* parallaxes to sources without external calibration. As explained in Lindegren & Bastian (2011) this is fundamentally different from the measurement of relative parallaxes that can be made using just a single (narrow) field of view. The spacecraft will orbit around the second Lagrangian point (L2) of the Sun–Earth system, located 1.5 million kilometres behind the Earth as seen from the Sun (about four times the Earth–Moon distance) at which an object will orbit the Sun with the same angular velocity as the Earth. This location is ideal for a satellite that is supposed to map the whole sky because light from the Sun and infrared radiation from the Earth come from the same direction and can be simply shielded off, while the orbit around




Figure 1.3: An impression of the Gaia spacecraft with the sun shield deployed. The cut-out in the thermal tent allows one to see the payload and its two primary mirrors mounted on the mechanically and thermally stable torus (top, see also Fig. 1.4), and the service module (bottom). The height of the satellite is about 3 m and the diameter of the deployed sun shield about 10 m. Image courtesy of EADS Astrium.

the Sun allows the satellite to observe every position on the sky in a period of half a year. The currently available communication data-rate over this distance is large enough ($\sim 3 \text{ Mbit s}^{-1}$ for 8 h day^{-1}) to transmit about 70 TeraByte of raw data over its nominal science mission duration of five years.

Gaia will continuously spin around its axis with a period of six hours, causing the two fields of view to scan the sky approximately along a great circle at a rate of 60 arcsec s^{-1} . The spin axis of the satellite is constantly pointed 45° away from the Sun and precesses 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), which is illustrated in Fig. 1.5. The spacecraft is commanded to follow the NSL to within 1 arcmin in all three axes. The precession of the spin axis changes the orientation of the consecutive great-circles that are observed, allowing the whole sky to be covered in about six months. For a mission lifetime of five years the evolution of the number of field-of-view transits as function of position on the sky is shown in Fig. 1.7. A given point on the sky will transit the combined fields of view on average 88 times with a minimum and maximum of about 40 and 240. When accounting for mission dead-time the average is 72 times. The observation time sampling is highly irregular and strongly dependent on the position on the sky.




Figure 1.4: Overview of the Gaia payload. The primary mirrors M1 and M'1 of the two telescopes can be seen on the top left, each $1.46 \times 0.51 \text{ m}^2$. The torus supports the entire set of optical elements and the focal plane (bottom right). The focal plane will be the largest ever flown in space, having a size of $0.42 \times 0.93 \text{ m}$ and containing 106 CCDs (see also Fig. 1.6). The optical path for each telescope is indicated by the two lines of sight LOS1 and LOS2. Image courtesy of EADS Astrium.

1.3 The observations

Each time a source transits the field of view it is observed by several CCDs as shown in Fig. 1.6. This typically results in nine accurate one dimensional (6×1 pixels) observations in the astrometric field, and one less accurate (6×12 pixels) observation in the skymapper. An observation consists of the number of photo-electrons measured in each pixel, together with the time at which the observation was made. This is the raw observational data that Gaia will transmit to the ground. Because the satellite is rotating with respect to the sources, all source images will move in the along-scan (AL) direction over the CCDs. To be able to image the moving images the CCDs are operated in Time-Delayed Integration mode, meaning that the charges in the CCDs are shifted through the pixels at exactly the same rate as the satellite is spinning (about 1 pix ms^{-1}). The integration time per CCD is about 4.4 s, in which the faintest sources ($G = 20$) will have generated less than a thousand electrons in the CCD. The CCDs are therefore designed to have a low readout




Figure 1.5: Figures illustrating the nominal scanning law for Gaia. The left image shows the scanning of great circles on the sky by the two fields of view due to the six hour spin period. The slow precession of the spin axis will change the orientation of the scanned great circles allowing coverage of different areas on the sky. The right image shows the precession of the spin axis at 45° around the Sun with a period of 63 days. This period gives the depicted overlap which ensures that each position on the sky is observed in at least three distinct epochs each half year. The combined motion allows the complete coverage of the sky illustrated in Fig. 1.7. Figures based on Lindegren et al. (2008) and Lindegren (2010).

noise of about $4 e^-$ rms while having a full well capacity of about $200\,000 e^-$ meaning that sources with $G \gtrsim 13$ will not be saturated. For $G < 13$ gates are activated that reduce the effective integration time such that these sources can still be observed unsaturated down to the lower magnitude limit of $G = 5.7$. The uncertainty in the number of electron counts per pixel will be dominated by the random detection statistics of the photons (described by Poisson statistics), therefore being largely unbiased, uncorrelated, and with well-defined uncertainties. In Sect. 2.4 I will discuss the effect of radiation damage to the CCDs which introduces biases to the photo-electron counts as illustrated in Fig. 2.7.

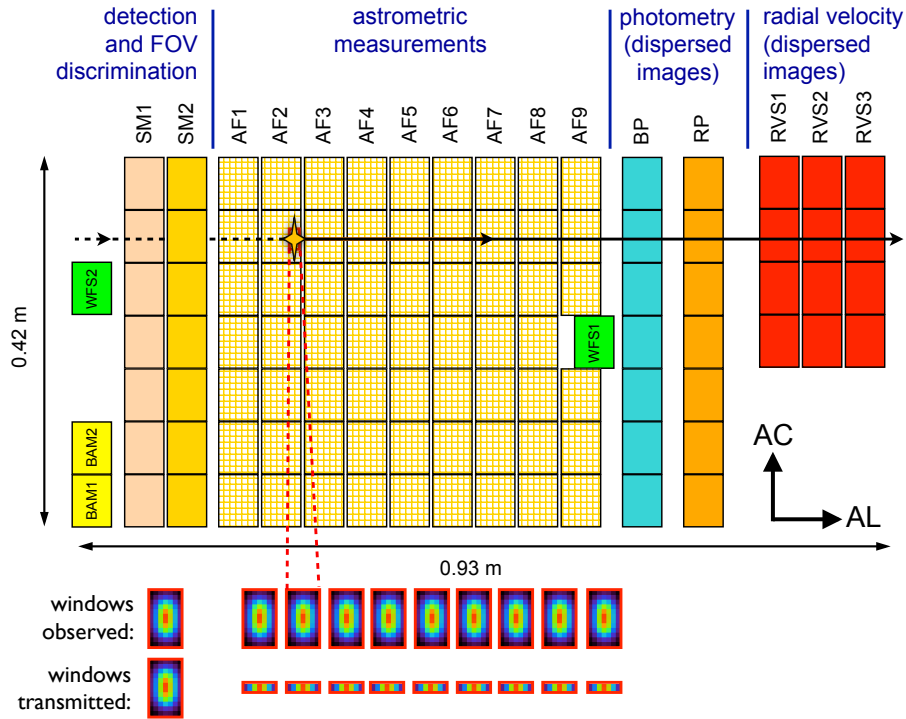


Figure 1.6: Schematic layout of the CCDs in the focal plane of Gaia and the observation process for a source transiting one of the fields of view. Due to the satellite spin, a source enters the focal plane from the left in the along-scan (AL) direction. All sources brighter than $G = 20$ mag are detected 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). A window of typically 6×12 pixels is read out around each source resulting in the ten observations per field-of-view transit that are used for astrometry (one SM and nine AF CCDs) shown in the bottom part of the diagram. For most sources the AF observations are binned in the across-scan direction (AC) resulting in one-dimensional photo-electron counts. Shown are also the additional CCDs used for the interferometric Basic-Angle Monitor (BAM), and the Wavefront Sensors (WFS) used for the initial mirror alignment. Illustration adapted from the focal plane image by A. Short (ESA/ESTEC).

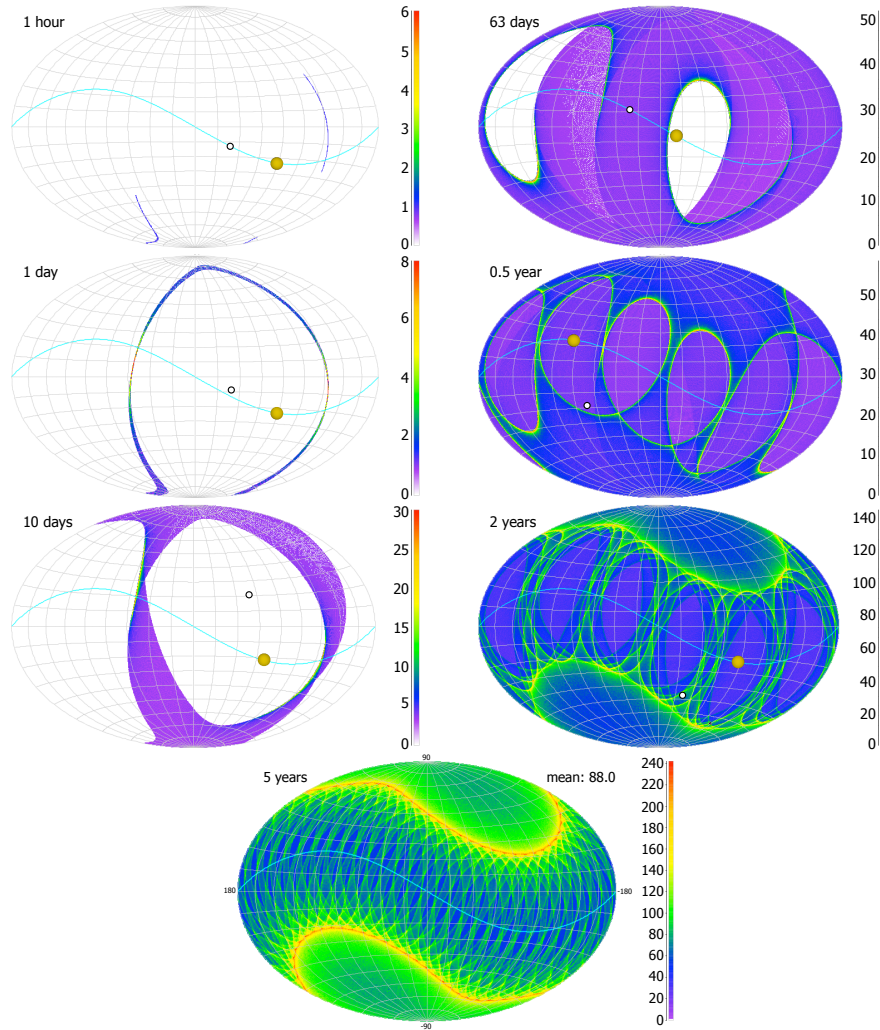


Figure 1.7: The expected number of field-of-view transits experienced by sources at different celestial positions due to the Gaia nominal scanning law. In the top six snapshots the location of the Sun is indicated by the yellow circle and the spin axis by the open circle. The projection uses equatorial coordinates, with right ascension running from -180° to $+180^\circ$ right-to-left. The blue line is the ecliptic (the plane in which Gaia orbits around the Sun together with the Earth). The average number is 88 field-of-view transits, although normally an average value of 72 transits is quoted (accounting for dead time). An over-abundance of transits occurs at 45° from the ecliptic 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.

Chapter 2

Modelling the observations

A fairly accurate definition of astrometry would be ‘the scientific measurement of the positions and motions of celestial bodies’. Since all heavenly bodies move with respect to one another (and are observed from a moving viewpoint as well), astrometry essentially comes down to the measurement of *relative* positions and motions at specific times. It is important to realize that these relative measurements, together with their uncertainty estimates, are the most fundamental and ‘true’ measurements that astrometry can provide. Nevertheless, before they can be used in astrophysical investigations, these measurements must be condensed into a smaller set of variables, more easily interpreted because they refer to commonly accepted systems of coordinates, units, and other models of the real world. This condensation process is what is meant by ‘data reduction’. Successful data reduction relies on the formulation of good models of the data, using the smallest possible set of model parameters. This chapter is about the formulation of such models for describing and reducing the Gaia observations.

2.1 Astrometric parameters

The International Celestial Reference System (ICRS) is the commonly used coordinate system for expressing positions of objects on the celestial sphere. The position of an object at a certain reference epoch is expressed by the two angles α , δ and the parallax ϖ , as shown in Fig. 2.1. The trigonometric parallax is used to express the distance to an object far outside of the solar system, and it is the maximum angular displacement of an object observed from a circular orbit of radius 1 AU. When measuring the position of a source from a viewpoint orbiting the Sun the parallactic offset of the source will trace an ellipse on the celestial sphere during the time it takes to complete one orbit. It




Figure 2.1: Overview of the astrometric parameters. The position and velocity of a source are parametrized by α , δ , ϖ , and $\mu_{\alpha*}$, μ_δ , μ_r , respectively. The radial proper motion component μ_r can normally not be measured from the astrometric observations but is derived for part of the Gaia sources ($G \lesssim 17$) using the Radial Velocity Spectrometer (see Fig. 1.6). The astrometric parameters are expressed with respect to the Barycentric Celestial Reference System (BCRS) represented by the vector triad $[\mathbf{X}, \mathbf{Y}, \mathbf{Z}]$ (explained in Fig. 2.4), which is aligned with the non-rotating International Celestial Reference System (ICRS).

is interesting to note that if Gaia would (hypothetically) be located at the solar-system barycentre from which we express the astrometric parameters, it would not observe any parallax displacements and distances to sources could not be derived, see Fig. 2.2. The parallax¹ ϖ is related to the distance d by

$$\sin \varpi = \frac{1 \text{ AU}}{d} \quad \Rightarrow \quad \varpi [\text{arcsec}] \simeq 1/d [\text{pc}] \quad (2.1)$$

Even for the closest star $\varpi < 10^{-5}$ rad, causing the approximation $\sin \varpi \simeq \varpi$ to be in error by less than 10^{-15} rad ($0.0002 \mu\text{as}$), which is completely negligible even for Gaia. Expressing the parallax in units of arcseconds (arcsec) gives rise to the definition of the astronomical distance unit of *parsec* (pc): an object at a distance of one parsec (about 3.26 light years) will have a parallax of exactly one arcsecond.

For all sources beyond our solar system the motion on the sky is very small, even over

¹The symbol ϖ is a cursive form of π with its legs bent inwards till they meet. It is frequently used for parallax to avoid confusion with the mathematical constant.




Figure 2.2: Example of the observed motion during 3 years for a source with a proper motion of $\mu = \sqrt{\mu_{\alpha*}^2 + \mu_{\delta}^2} = 250 \text{ mas yr}^{-1}$ and parallax $\varpi = 100 \text{ mas}$ (10 pc distance). In the left diagram, the motion as seen from the solar-system barycentre (no parallax effect), in the right diagram as seen from Gaia. $\Delta\alpha \cos \delta$ and $\Delta\delta$ are the offsets from an arbitrary reference point. Image adapted from Perryman (2004).

centuries, and can therefore be accurately expressed as a perturbation to the position. The proper motion of a source is expressed by three components² $\mu_{\alpha*}$, μ_{δ} , and μ_r , as shown in Fig. 2.1. Because astrometric observations can only measure the (tangential) motion on the celestial sphere the radial component is usually determined by spectroscopy using the Doppler shift of stellar spectral lines. Gaia has an instrument that measures radial velocities for part of the sources, the Radial Velocity Spectrometer (RVS), but because this is independent of the astrometric measurements we will usually only refer to the two tangential components $\mu_{\alpha*}$ and μ_{δ} when speaking about proper motion. For most sources the radial proper motion is computed from the parallax and the spectroscopic radial velocity v_r as $\mu_r = v_r \varpi / (1 \text{ AU})$.

The proper motion components $\mu_{\alpha*}$, μ_{δ} , and μ_r (expressed in $\mu\text{as yr}^{-1}$) can be transformed into linear velocity components (in km s^{-1}) if they are multiplied by the distance to the source. However, since the distance is often poorly known, this would sometimes result in huge uncertainties in the velocities, even when the proper motion is very well determined. It is therefore much better to model the observations by means of the proper motion components. In fact, this model works very well even when the parallax is essentially zero, as it is for the quasars.

The five components α , δ , ϖ , $\mu_{\alpha*}$ and μ_{δ} allow us to *parametrize* (describe) the motion of sources in space and we will therefore refer to them as *astrometric parameters*. Together with the coordinate transformations that will be described in the next section we can construct a *model* that is able to predict the position and the times at which a source with given astrometric parameters will be observed (by Gaia). This modelling of stellar motion on the sky will be discussed in Sect. 2.3.

²The notation $\mu_{\alpha*} = \mu_{\alpha} \cos \delta$ signifies that the proper motion in right ascension is expressed as a true arc length on the sky (as opposed to $\mu_{\alpha} = d\alpha/dt$).




Figure 2.3: In General Relativity the description of the motion of the source, observer, and light connecting the two is highly complex and needs additional information describing the space-time metric encountered by the light which is not *a priori* known. The solution is therefore to model just the observations, i.e., the direction of the light ray in a sphere located around the observer. Illustration by L. Lindegren adapted from Klioner (2003).

2.2 Reference systems

Astrometry is about measuring the (relative) direction of light rays reaching us from distant astronomical sources. This observation is the result of a sequence of events (Fig. 2.3): the emission of the light ray from a moving source, the propagation of the ray through space towards the observer, and the reception of the light ray by the observer, which is also moving. A precise description of these events, to the accuracy of Gaia, requires the use of General Relativity. However, it is only in a small volume around the observer (suggested by the grey circle in Fig. 2.3) that the space-time metric can be modelled well enough to trace the light rays to the full accuracy of Gaia. The relativistic modelling of astrometric observations needed for Gaia has been worked out by Klioner (2003) and essentially consists of the rigorous specification of several different reference systems and the transformations between them.

For practical uses it is necessary to translate the relative measurements into absolute positions with respect to a particular reference system, with a well-defined space-time origin, axis directions, and coordinate units. A useful origin in space is the centre of mass of our solar system, called the *solar-system barycentre*, around which all planets and the Sun rotate. This barycentre is typically located within two solar radii from the Sun's centre, mainly depending on the relative positions of Jupiter and Saturn. Because the masses, positions and velocities of the planets in our solar system are known to very high precision we can compute the position of the barycentre very accurately with respect to all the planets. The Barycentric Celestial Reference System (BCRS) has its origin at the solar-system barycen-




Figure 2.4: Overview of the reference systems used to model Gaia observations. The astrometric parameters (see also Fig. 2.1) are expressed with respect to the non-rotating Barycentric Celestial Reference System (BCRS) represented by the vector triad $[\mathbf{X}, \mathbf{Y}, \mathbf{Z}]$. The attitude (orientation) of the satellite is modelled in the non-rotating Centre-of-Mass Reference System (CoMRS) represented by the vector triad $[\mathbf{X}', \mathbf{Y}', \mathbf{Z}']$. The CoMRS has the same orientation as the BCRS but its origin is fixed to the centre of mass of Gaia instead of our solar system. Finally there is the rotating Scanning Reference System (SRS), represented by the vector triad $[\mathbf{x}, \mathbf{y}, \mathbf{z}]$ whose spatial axes are defined with respect to the two viewing directions of Gaia, \mathbf{f}_P (pointing to the centre of the preceding field of view) and \mathbf{f}_F (pointing to the centre of the following field of view). \mathbf{x} points halfway between \mathbf{f}_P and \mathbf{f}_F , and together with \mathbf{y} and the nominal spin axis \mathbf{z} (normal to \mathbf{f}_P and \mathbf{f}_F) defines the right-handed SRS.

tre and is aligned with the non-rotating International Celestial Reference System (ICRS), which is defined by the positions of a few hundred extragalactic sources. The ICRS is approximately aligned with the equatorial system but is not linked to the motion of the Earth and therefore unaffected by the Earth's precession and nutation. For Gaia we need to introduce at least two more systems; see Fig. 2.4 for a graphical overview. The first one is the non-rotating Centre-of-Mass Reference System (CoMRS) which is aligned with the BCRS, but has its origin fixed to the centre of mass of Gaia. This system is used to model the orientation of the spinning satellite. Secondly there is the rotating Scanning Reference System (SRS) which also has its origin at the centre of mass of Gaia, but with its axes fixed to the instrument, causing it to rotate with the satellite. Translating an observation made by the spinning Gaia satellite to the solar-system barycentre (or the other way around) therefore takes at least three coordinate transformations.

2.3 Modelling stellar motion on the sky

A model is often, but not necessarily, a simplified representation of reality that is able (and often tuned) to accurately describe the observed data. In principle there is always an infinitude of models that can describe the observed data well within the uncertainties, but most of them will have no predictive value. A *useful* model is therefore also able to predict observations outside the available data (or accurately predicts all of the data when it is constructed from a subset of it). Coming up with a good model is not always straightforward and the real scientific insight is therefore often gained when one is constructing and testing out different models. A model that has a high physical relevance (i.e., it behaves according to the known laws of physics) is often preferred. In this respect the model with the astrometric parameters introduced earlier is not the most straightforward because it describes the space motion of objects from a solar-system centred coordinate system: why do we put ourselves in the centre while we know that the sources in the universe literally do not revolve around us? Imagine for a moment that we would describe the position and velocity of sources in just plain Cartesian coordinates, for example expressed with respect to the BCRS. Now think about what it means to measure the angle between two sources: it is the angle between the light rays from source i and j measured at position \mathbf{x} at time t . Source i might for example be 10 light-years away, while source j is 10 000 light-years away, meaning that you see the light emitted from the location where source i was 10 years ago, and where source j was 10 000 years ago! The question is: how are you going to assign the Cartesian coordinates \mathbf{x}_i and \mathbf{x}_j for these sources? One way is to refer each position to the time at which the light was emitted, another is to use the measured speed to estimate where the source was or will be at some global point in time. Both of these Cartesian model parametrizations perhaps make sense physically, but they are highly impractical mainly because the result depends critically on the distance (or light-time), which is the least accurately determined coordinate. The adopted solar-system centred model is preferred because it is *practically much more usable*. Additionally the adopted model matches much better the actual observations (locally observed relative angles on the sky) which reduces the modelling errors and correlations between the estimated astrometric parameters. Given that the typical mission time span is of the order of a decade, and the history of accurate astrometric measurements is perhaps a century old, the current model will likely stay the only practical way of expressing positions and velocities of objects in our Galaxy until we figure out how to build warp-engines.

2.4 Forward modelling of the observations

In Sect. 2.1 we introduced the five astrometric parameters that will be determined for every source in the Gaia catalogue: the spherical angles α and δ , the parallax ϖ , and the tangential proper motion components μ_{α^*} and μ_δ , shown in Fig. 2.1. Additional instru-

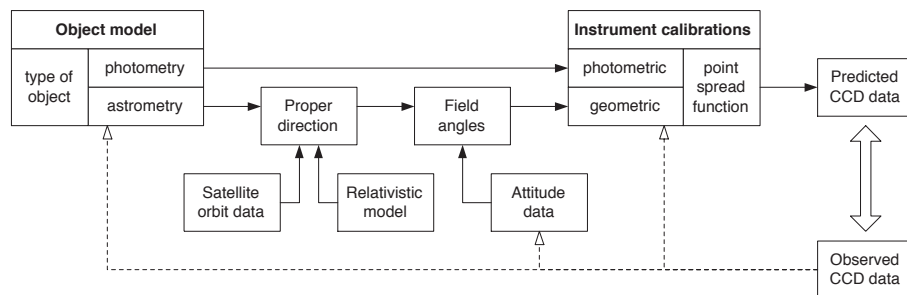


Figure 2.5: The forward modelling approach for Gaia. The dotted lines indicate feedback loops in which the model parameters are iteratively improved. Figure taken from Perryman et al. (2001).

ments in Gaia will help to determine some intrinsic properties of each source as well, like luminosity, colour, age, metallicity, and more. The way all of these source parameters (astrometric and intrinsic) are being estimated is by *forward modelling* of the observed photo-electron counts (Fig. 2.5). **This means that, given a set of source parameters, we need models that allow us to predict every single pixel value ever observed by Gaia!** These models in principle need to describe every process from light being emitted by a certain source to the detection of the photons on the CCDs in Gaia. They contain a large number of ‘nuisance parameters’ that need to be estimated from the observations as well. The term ‘nuisance parameters’ is used because, although they are astronomically completely uninteresting, they are still needed, and they complicate the modelling and introduce additional errors and correlations between the source parameters. The basic angle, focal plane geometry and the photometric calibration of the CCDs are a few examples of nuisance parameters, but there are millions of them in the complete set of models. A lot of effort in the preparation of the data processing for Gaia has been directed towards developing models that can represent the data adequately with the smallest possible number of nuisance parameters.

So, how do we estimate all of these source and nuisance parameters? As mentioned before, given a set of parameter values we can use the models to predict for each source when it is observed and the number of photo-electrons in each pixel. Comparing this prediction with the actual observations we can determine the *residual* difference between the two. The goal is to find a set of parameter values that minimizes these residuals for all observations of all sources. Using the residuals and parameter dependencies of the models we can compute a set of updated parameters for which the residuals will be smaller, for example using the Newton-Raphson method. If it is computationally feasible we can compute the updates for a large set of parameters at the same time, for example using Least Squares when the observation errors can be assumed to have a normal distribution, or Maximum Likelihood in the more general case. The change of the predicted obser-

vations is not necessarily linearly dependent on the update of the parameters, therefore the update process typically needs to be iterated. In the end one finds a set of optimal parameters that give the global minimum of residuals. For a well-designed experiment (which Gaia is), and provided that we use models that are good enough representations of reality, the resulting source parameters should be close to their ‘true’ values. Exactly how ‘close’ they are is a very important question to which I will return in the next chapter.

Since the total number of source and nuisance parameters is extremely large ($\sim 5 \times 10^9$), it is not realistic to update all of them in one go. One way of iteratively updating a large set of (correlated) parameters is to compute the updates for each model separately by keeping the parameters of the other models fixed (effectively assuming they are correct). The total set of parameters can then be iteratively improved by sequentially updating the parameters of each model, every time using the latest parameter estimates of the other models. In the data processing of Gaia there are many loops within loops that use this technique to iteratively improve the parameter estimates. Since I am here only discussing the astrometric parameter estimation I will describe the two iteration loops used for that purpose: image parameter estimation and AGIS.

Image parameter estimation

Image parameter estimation is the process by which the location and flux of an image is determined from the observed one- or two-dimensional photo-electron counts. For each individual observation the centroid of an image must be located with a precision of a few hundredth to a thousandth of a pixel, depending on how faint or bright the source is. This needs extremely accurate modelling of the centroid location and instrument response. Fig. 2.6 shows the scheme that is used to predict the photo-electron counts. It involves three different models: the scene, the point/line spread function, and the charge distortion model explained below. The parameters of each model are updated separately by keeping the others fixed. To find the optimal set of parameters (that minimizes the residuals) the procedure is iterated several times. More details of this process can be found in Paper 6.

An important effect that needs to be modelled is radiation damage. During solar flares the Sun expels fast protons that follow non-straight trajectories through our solar system. Gaia will thus experience a proton flux coming from virtually all directions. These protons will not only enter the satellite through the telescope openings but can also fly through (thin) parts of the satellite like the sunshield. When they collide with the atomic lattice of a CCD the deposited energy can cause dislocations in the lattice structure. These dislocations effectively act as ‘traps’ that can capture and release electrons. Because the photo-electrons are moved through the CCD in the along-scan direction, the photo-electrons that got trapped and released during the integration time will be redistributed in the negative along-scan direction. The typical time for a captured electron to be released depends

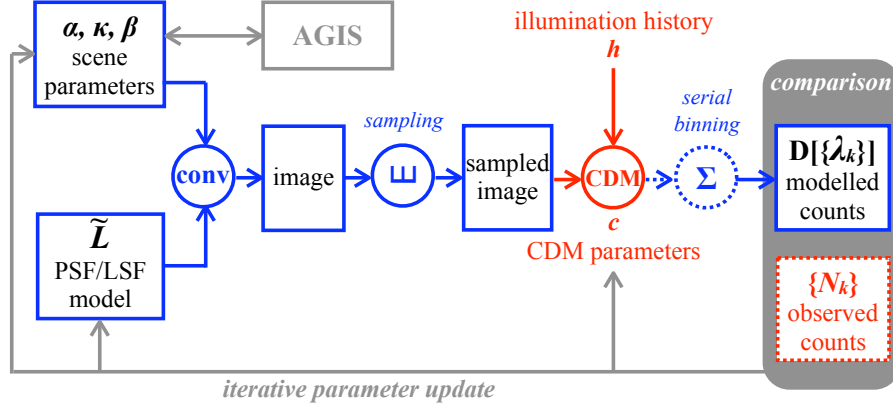


Figure 2.6: Forward modelling approach in the image parameter estimation (Lindgren 2008): an ‘image’ unaffected by radiation damage is generated from the ‘scene’ by combining the flux (α), (sub)pixel location (κ) and sky background (β), together with the point/line spread function (PSF/LSF) that characterizes the instrument response. This ‘image’ is then sampled into individual pixels. Subsequently it is distorted by a charge distortion model (CDM) mimicking the effect of radiation damage in the CCD. If the observation was binned on board to one dimension, the same is done in this model. The predicted distorted counts are compared to the observed counts. In an iterative procedure the scene, LSF, and CDM parameters are successively improved. The astrometric solution (AGIS) uses the scene parameters to estimate the astrometric parameters. AGIS also provides an update to the parameters that correct for residual modelling errors that can only be detected by combining all the observations. Diagram extracted from Paper 6.

on the type of trap. If the release time is longer than the width of the window that will be read out around the source, the charge is ‘lost’ from the window and the total number of detected photo-electron is reduced. If the trap is fast (e.g. $\lesssim 1$ millisecond, the time it takes to move photo-electrons from one pixel to the next) the charge is moved within the window, causing a shift (delay) of the image centroid. See Fig. 2.7 for examples of both cases. More traps cause more distortion in the observed photo-electron counts. Because the damage cannot be undone, the total damage in the Gaia CCDs will only increase during the mission. A major role in the modelling of the observations is therefore taken by the so-called Charge Distortion Model (CDM). The CDM is basically a non-linear transformation of the undamaged photo-electron counts that is (loosely) based on the theoretical expected trapping behaviour over the whole transfer width of the CCD. With a suitable parametrized (and calibrated) CDM one can in principle predict the non-linear distortion of any (extended) source observed by Gaia. Finding a suitable parametrization and calibration of the CDM has been a continuous endeavour over the last several years.




Figure 2.7: The effect of different trap release times on the observed photo-electron counts with respect to the true LSF. Because the image (true LSF) moves over the CCD due to the rotation of the satellite, temporary trapping of electrons during the integration time over a CCD causes the traps to effectively move some of the electrons in the direction opposite to the readout, resulting in a distortion of the observed electron counts. The images move at a speed of about 1 pixel ms^{-1} , which allows us to make a distinction between *fast* traps that release electrons after $\lesssim 1 \text{ ms}$, and *slow* traps that release after (much) longer times. Fast traps effectively move some of the electrons towards the ‘backside’ of the window that is read out, while slow traps will effectively reduce the amount of electrons in the window. Both effects are unfortunately highly non-linear. Gaia is expected to get a mixture of fast and slow traps. The traps themselves are induced by fast protons from the Sun causing dislocations in the atomic lattice of the CCD. This is what we call radiation damage.

To reduce the effect of radiation damage on the observations there are two hardware mitigation strategies. One of them is a positively charged ‘doping profile’ called supplementary buried channel (SBC) that is present along each column of pixels in the CCD. It can keep a few thousand electrons closeby, effectively shielding them from being captured by traps located in the rest of the pixel. Sources of the faintest magnitudes observed by Gaia ($G \sim 20$) will only produce a few hundred electrons during the CCD integration time and are therefore less affected by radiation damage. Another strategy is to artificially inject a large amount of electrons at the first line of pixels of a CCD every second or so, this process is called Charge Injection. The electrons will be moving over the CCD along with the source images, and likely fill many of the empty traps. The slow traps will stay filled for quite some time and sources observed just after a charge injection will have less distortion in their observations. The fact that the captured charge injection electrons will also be released after some time needs some additional modelling of the background. In practice it is however expected that this modelling will be much easier than trying to model the continuous filling and emptying of all the traps in the CCD, which would be

the alternative. Also, the charge injection ‘resets’ the illumination history of the pixels, so that the CDM does not have to take into account the images crossing the pixels before the injection.

The last two papers in this thesis contain a detailed study of how the radiation damage will affect the astrometric observations of Gaia. In particular, in Paper 6 we study the image parameter estimation, including the modelling in Fig. 2.6 and the different hardware mitigation strategies.

The Astrometric Global Iterative Solution (AGIS)

The image parameter estimation described above will provide us with the subpixel image location of the source in all sets of photo-electron counts (i.e., the observations). Because the satellite spins with a very precise rate of $0.982 \text{ pix ms}^{-1}$ the estimated subpixel location is equivalent to a time correction. Applying this correction to the time at which the observation was read out we can express all observations as ‘observation times’³. Again we can now use models to predict the times and minimize the residuals with respect to the observation times. Figures 2.2 and 2.4 explain how part of the modelling works. A source parametrized by astrometric parameters in the BCRS will move on the celestial sphere with a certain proper motion vector with respect to its position at a given reference epoch, see the left panel in Fig. 2.2. Given the parallax of the source and our knowledge of the position of Gaia with respect to the solar-system barycentre we can predict the motion as seen from the non-rotating CoRMS at the position of Gaia, see the right panel in Fig. 2.2. This adds the parallax effect explained in Sect. 2.1.

The next step is to translate the observed direction of the source to the rotating SRS system, for which we need a model of the orientation of the satellite at any moment in the mission. Additionally (not shown in Fig. 2.4) there is a coordinate system associated with each field of view in which minuscule but significant variations in the position and orientation of the CCDs is modelled. This finally allows us to predict at which pixel position and at what time the source will be observed by Gaia. Because Gaia makes such precise measurements relativistic effects like light bending by the Sun and planets, aberration due to the motion of the satellite, and the time difference between BCRS and CoMRS (the Römer delay) need to be taken into account when modelling the observed direction to a source (see Klioner 2003). For example, in order to make the Lorentz transformation that corrects for the aberration, the barycentric velocity of Gaia (which is about 30 km s^{-1}) must be known to a few mm s^{-1} , which is far from trivial.

Given the observation times for the observations of all the sources, we can distinguish

³Note that although these times are not strictly what Gaia observed we still refer to them as ‘observation’ times. When the parameter values of the models in the image parameter extraction are updated the observation times will change as well, therefore it is necessary to iterate between the image parameter extraction and AGIS.




Figure 2.8: The four basic models that are iteratively updated in AGIS: the astrometric model having the astrometric parameters \mathbf{s}_i of each source (i), the attitude model describing the orientation of the satellite with parameters \mathbf{a} , the geometric calibration model describing the location and orientation of the CCDs with parameters \mathbf{c} , and the global model containing parameters \mathbf{g} that affect all observations, such as the parametrization of relativistic light bending.

four different models in which the parameters can be iteratively improved: the astrometric model (S) having the astrometric parameters of each source, the attitude model (A) describing the combination of the optics and the orientation of the satellite as a continuous function of time, the geometric calibration model (C) describing the location and orientation of the CCDs, and the global model (G) containing parameters that affect all observations, see Fig. 2.8. The iterative procedure to improve the parameters of each model is called the Astrometric Global Iterative Solution (AGIS). A comprehensive overview of AGIS is given in Lindegren et al. (2012).

All the papers in this thesis, except Paper 6, use a simplified version of AGIS, called AGISLab (Sect. 4.2), to study the propagation of errors through these iterations. In the first five papers we consider only the propagation of the random errors due (mainly) to the Poisson statistics of the photo-electrons. In Paper 7 we propagate the results from Paper 6 through AGIS to study the effect of the systematic errors caused by radiation damage on the astrometric parameter.

Chapter 3

Characterizing the astrometric errors

Based on the title of this thesis you could have expected this introduction to be filled with equations related to error calculation¹. As you have seen, my choice has been to focus on the *context* of astrometric parameter estimation for Gaia in order to understand from where the errors originate. In this Chapter I will briefly describe why it is important to understand these errors and how they can be described.

3.1 Errors and uncertainties

We first need to sort out some terminology. By ‘error’ I mean the deviation of a measured or derived quantity (such as parallax) from its true value. Normally you do not know the error – if you did, you should just subtract it from the measured value and use the corrected (true) value instead! What you *can* know, and very often need to know, is the *uncertainty* of the measured quantity. This is usually given in the form of an estimated standard deviation, but often (confusingly) called ‘error’, ‘mean error’, or (slightly better) ‘standard error’, ‘RMS error’, etc. A better term could be ‘standard uncertainty’, where the ‘standard’ signals its relation to the standard deviation of the (implicitly assumed) Gaussian error distribution. So, when I talk about the need to understand errors, I actually mean (among other things) that we need to know the standard uncertainties. But why is that important?

¹If you find the lack of equations in this introduction disappointing have a look at Paper 4.

Every measurement has some random variability that is inherent to the thing being measured and the measurement process. Because it is random, the exact value cannot be predicted and therefore not modelled, and the variability will define the uncertainty of the measurement. This uncertainty can however be described in terms of *probabilities*, as illustrated in Fig. 3.1. The *probability distribution* of the error can be characterized with quantities like the mean, standard deviation, and skewness. The measurement together with an estimate of its uncertainty will be referred to as an *observation*, and any set of observations or values estimated from them as *data*.

The presentation of data do not have scientific meaning without a specification of their uncertainties. Although it can sometimes be difficult to specify the exact probability distribution of the measurement errors, often a reasonable guess of the statistics describing it will suffice, e.g., the standard deviation. The reason is that most measurement processes are affected by (many) different random processes, and according to the central limit theorem the combination of them will result in a probability distribution approaching a Gaussian (which is completely specified by the mean value and standard deviation). A non-zero mean of the error distribution is referred to as *bias*. In case it stays undetected (i.e., if it is assumed to be zero) it leads to systematic errors. The uncertainty resulting from the random errors of measurements can sometimes be estimated quite well, often because one can repeat a measurement with the same set-up several times to sample the

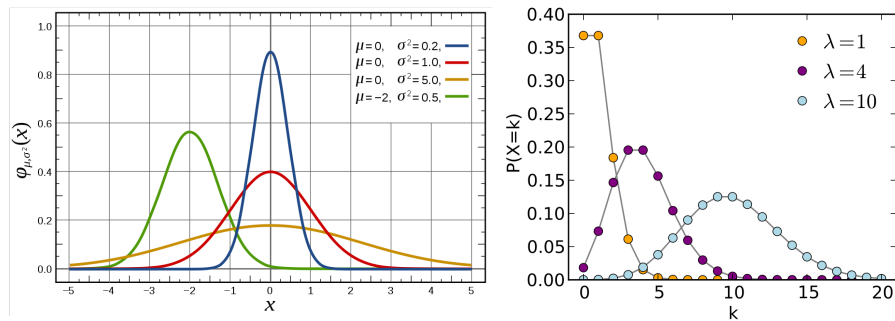


Figure 3.1: Figures illustrating the probability distributions of Gaussian and Poisson random variables. **Left:** A Gaussian is completely described by its mean value μ and standard deviation σ . Its skewness is zero because it is symmetric. The Gaussian is often used to model measurement errors in a continuous variable. **Right:** A Poisson distribution gives the rate of occurrences of a discrete random variable k for a mean rate of λ (the function is only defined for non-negative integer values of k , the connecting lines are only guides for the eye). It is completely described by the parameter λ . Its standard deviation is $\lambda^{1/2}$ and skewness $\lambda^{-1/2}$. The Poisson distribution is often used to model photo-electron counts in a detector. Copyright images: Inductiveload and Skbkekas (Wikimedia Commons).




Figure 3.2: A set of data points with their standard uncertainties, illustrating the importance of knowing the true uncertainties (see discussion in text).

random variability of the measurement (this implicitly assumes that the experiments are identical). Determining the systematic errors is the most difficult part because it will not reveal itself by repeating the experiment (e.g., observing the wrong star).

3.2 Why we need to understand errors

Knowledge of errors, as described for example by the standard uncertainty, is essential for interpreting data. To illustrate this I will use the sample of data in Fig. 3.2, and assume that it represents the proper motions of stars in a stellar cluster. Say that we want to determine cluster membership from this sample, how important are then the individual uncertainties? We can for example use a criterion based on the difference between the proper motion of a single star and the mean of the other stars divided by the standard uncertainty of the difference. Using this criterion star number 4 would clearly stand out and is likely to be excluded. The example illustrates the need to know the uncertainties, but it is also an example where the outcome does not critically depend on this information, i.e., a 10% smaller or larger uncertainty would give a very similar membership selection.

But what if the data would be used to estimate the velocity dispersion of the cluster? Star 4 will probably be rejected, and a dispersion computed from the remaining data. This is then the quadratic sum of the intrinsic dispersion (which we want to know) and the uncertainty of the data. If the two are of similar size (as is often the case in science when trying to quantify effects at the border of what is measurable) the uncertainties need to be known very reliably in order to estimate the intrinsic dispersion. Depending on the application the result depends more or less critically on the assumed standard uncertainties. Understanding the errors (or uncertainties) is at least sometimes of crucial importance.




Figure 3.3: The left diagram shows a simulation of 500 parallax pairs (ϖ_1, ϖ_2) where ϖ_1 and ϖ_2 are drawn from independent Gaussian distributions with mean value $100 \mu\text{as}$ and standard deviation is $\sigma = 10 \mu\text{as}$. The parallax difference $d = \varpi_2 - \varpi_1$ is constant along any 45° line running from bottom left to upper right, while $m = (\varpi_1 + \varpi_2)/2$ is constant along the lines in the perpendicular direction. In the right diagram the same points are plotted versus these transformed variables. The standard deviation of m is $\sigma/\sqrt{2} \simeq 7.1 \mu\text{as}$ and that of d is $\sigma \times \sqrt{2} \simeq 14.1 \mu\text{as}$. Figures courtesy of L. Lindegren.

3.3 Correlations

Another important aspect of errors, when interpreting data, is the *correlation* between the errors of data points. Let us start by explaining what correlations are using the left panels of Figs. 3.3 and 3.4. Figure 3.3 shows 500 possible outcomes of measuring the parallaxes of two stars 1 and 2, where for each outcome the two parallax values have been plotted against each other. The pair of values obtained in a real experiment could, with equal probability, be represented by any of the points. Both measurements have a Gaussian distribution with mean $100 \mu\text{as}$ and standard deviation $10 \mu\text{as}$ (which can be seen by projecting the data points onto one of the axis and making a histogram of the distribution along the line). As can be seen from the left panel, the particular value of the error in one star is not related to the error in the other (e.g., the error distribution for star 1 is the same along all horizontal lines representing different values for star 2), thus the correlation of the errors is zero. The situation is very different in the left panel of Fig. 3.4. Here there is clearly a very strong correlation between the measurement errors in star 1 and 2. The strength by which the errors are coupled is expressed by the *correlation coefficient*. A correlation coefficient of $+1$ or -1 means that the errors are linear related, which would




Figure 3.4: Similar plots as in the previous figure, except that the parallax pairs (ϖ_1, ϖ_2) are drawn from a 2D Gaussian distribution with correlation coefficient 0.8. In this case the standard deviation of m is $9.5 \mu\text{as}$ (i.e., *larger* than in the uncorrelated case) and the standard deviation of d is $6.3 \mu\text{as}$ (i.e., *smaller* than in the uncorrelated case). Figures courtesy of L. Lindegren.

cause the error pairs in the left panel of Fig. 3.4 to form a straight line. The sign of the slope of the line determines the sign of the correlation coefficient (e.g., the correlation coefficient in Fig. 3.4 is positive).

Correlations become important when combining different data points. For example, the two parallax values discussed above (for stars 1 and 2) could be combined into a mean parallax if the two stars are components in a binary and we want to get the distance to the binary with the best possible accuracy. On the other hand, if we suspect that the two stars are actually an optical pair, we might be more interested in the difference between the parallaxes. The distributions of the mean and difference of the two parallaxes are shown in the right panels of Figs. 3.3 and 3.4. Although the data are plotted on the same scale, the axes have changed due to the transformation. In case of the zero correlation coefficient in Fig. 3.3, the standard deviation in the transformed quantities is $7.1 \mu\text{as}$ for the the mean, and $14.1 \mu\text{as}$ for the difference. This is just an effect of the different statistics that are computed. If we now however look at the standard deviation of the positively correlated data in Fig. 3.4 the standard deviation in the transformed quantities is $9.5 \mu\text{as}$ for the mean, and $6.3 \mu\text{as}$ for the difference. This means that the standard deviation of the mean is *larger* than in the uncorrelated case, while the standard deviation of the difference is *smaller*! For a negative correlation coefficient this result is reversed. This leads to the (perhaps surprising) result that correlations can work in your favour or against you depending on what kind of statistic you are interested in. Positive correlation

among the parallax values is bad for determining the mean distance to a cluster, but good for studying its internal structure. As found in Paper 1 and Paper 5, the astrometric parameters are (weakly) positively correlated for sources at small separations (of the order of, or less than, the field-of-view size).

The correlations and standard uncertainties among data points are conveniently expressed by (co)variances. The covariance between two data points p_i and p_j is expressed as $\text{Cov}[p_i, p_j] = E[e_i e_j]$, with e_i and e_j the (unknown) error in each data point, and E the expectation operator. We assume here that the errors are unbiased, $E[e_i] = E[e_j] = 0$. The variance of p_i is $\sigma_i^2 = E[e_i^2]$ and the variance of p_j is $\sigma_j^2 = E[e_j^2]$. The covariance can also be written $E[e_i e_j] = \rho_{ij} \sigma_i \sigma_j$ with ρ_{ij} being the correlation coefficient between the data points, having a value in the range $[-1, 1]$. For a set of data the covariances between all pairs of data points can be expressed in a single matrix where each row and column signify a different data point. The diagonal will therefore contain the variances of the data points, and the off-diagonal elements the covariances between the data points. In this way the standard uncertainties and correlations of all 5×10^9 astrometric parameters for Gaia can be expressed in a single covariance matrix having 2.5×10^{19} elements.

3.4 Random errors in the Gaia data

As was mentioned in Sect. 1.3 the raw observations for Gaia are the number of photo-electrons measured in each pixel, together with the time at which the observation was read out from the CCD. In analogy with the measurement description of Sect. 3.1 the ‘thing’ being measured here is the light from the source plus background and scattered light. This detection follows the random statistics of the photons described by the Poisson distribution (the right panel of Fig. 3.1). The ‘measurement process’ includes effects like diffraction of light through the telescope, the effective integration time, motion of the satellite during the integration time, charge transfer inefficiencies in the CCD (e.g. due to radiation damage), and readout noise. Although many of the latter effects have a significant effect on the measurement, the inherent random variability (the part of the measurement that cannot be predicted/modelled) will be completely dominated by the random variability of the photon statistics (the CCD readout noise also has a random variability but it is much smaller than the statistical photon fluctuations and can be included in the Poisson model as a slightly increased background intensity). The raw number of photo-electrons measured in each pixel therefore will very accurately follow a random variable with Poisson statistics. Note however that for $\lambda > 10$ the Poisson distribution already looks similar to a Gaussian distribution sampled at discrete points (right panel of Fig. 3.1). Even for the faintest sources that Gaia observes ($G \sim 20$) the number of photo-electron counts per pixel will be $\gg 10$. An important property of this dominant photo-statistic nature of the observations is that they will be (largely) unbiased and uncorrelated. Because the Gaia instrument will be very well calibrated and the readout

noise is small and well-known, the (standard) uncertainties of the observations will also be well-defined.

The centroiding process will combine the photo-electron counts from several pixels, effectively resulting in a Gaussian error probability distribution of the image location estimates. In the absence of radiation damage these are unbiased. The transformation to ‘observation times’ as described in Sect. 2.4 propagates the errors of the image parameters into the estimated observation times, which are then also Gaussian and unbiased. Under these assumptions one only needs to study the propagation of random errors through AGIS in order to characterize the final astrometric errors. Therefore, the first question that needs to be answered is: **how do random errors propagate in the astrometric solution?** This is addressed to various degrees in Papers 1–5. In the papers we only study the propagation of second-order moments of the errors (which can be conveniently expressed as covariances), not the probability density functions. The fact that we always use Gaussian errors is a realistic but not necessary assumption.

As explained in the previous section it is very important to have information about the correlations between the parameters. An important practical question is therefore: **how can the covariance between any pair of astrometric parameters be estimated?** This question is discussed in Paper 3, 4 and 5. It is a very legitimate question as the full covariance matrix for the expected 5×10^9 astrometric parameters would be $\sim 10^8$ TeraByte, which is totally impractical to compute, store and query efficiently with current techniques. We found a way around this problem by using a series expansion of the covariances for which we would need only some ~ 20 TeraByte of uncompressed data for the full mission.

Through the modelling of the observations, the error in each single astrometric parameter is in some degree linked to the errors in all other astrometric and nuisance parameters. However the largest contribution to the error in an astrometric parameter comes from the attitude parameters, in particular those describing the orientation of the satellite at the times a source was observed. This can be understood by the nature of the attitude modelling, which is very flexible on timescales longer than 5–30 s. The standard uncertainty of the attitude in each such interval will be dominated by the number and uncertainty of the observations made of sources in the two fields of view. The resulting (unknown) random error in the attitude will affect the modelled observations of all these sources in each field of view in the same way (e.g., an error in the opposite direction of rotation will cause a delay in the modelled observation times). Because the attitude is used to estimate the source parameters (remember the AGIS scheme in Sect. 2.4) this common error will produce correlations between the parameters of the observed sources. Sources that have a separation on the sky that is similar or smaller than the field-of-view size will be observed together in a large majority of the transits, therefore being the most strongly correlated. The instrument calibration parameters do not have this ‘local’ dependence in time and position on the sky, and their contribution to the error of an astrometric pa-

parameter is therefore significantly smaller than that of the attitude. In all the papers that involve the estimation of astrometric parameters using AGIS we therefore neglect other nuisance parameters than the attitude.

Because information about the covariances of astrometric parameters is needed to do science with the Gaia data, the work described in this thesis is in principle relevant for any scientist that will use the astrometric Gaia catalogue.

3.5 Systematic errors in the Gaia data

In the previous section we deliberately specified the absence of radiation damage effects on the observations. This is mainly because it is known from tests that the resulting distortions of the observations are difficult to model and calibrate to the levels needed for Gaia. It has been observed that radiation damage introduces biases (i.e., non-zero mean in the error distributions) as well as non-linear distortions of the photo-electron counts (Fig. 2.7). In case the distortions cannot be modelled well enough in the image parameter estimation they will lead to modelling errors propagating into the observation times used in AGIS, which subsequently propagate into the astrometric parameters that will make up the astrometric catalogue. A very important question that needs to be addressed is therefore: **what is the impact of CCD radiation damage on Gaia astrometry?** This question has two parts. First we need to estimate the effect of radiation damage on the image parameter estimation, and how it affects the observation time estimates that will be provided to AGIS. This is extensively discussed in Paper 6 where we made detailed Monte-Carlo simulations of the observation process in the astrometric CCDs. Secondly we need to study how these errors propagate through AGIS, which is done in Paper 7.

Chapter 4

The evolution of the thesis work

*'Once you defined the problem precisely
the rest is just a matter of mechanics.'*

— Lennart Lindgren

This chapter describes the history of the seven papers in this thesis and explains how they are connected to each other. The first three papers are conference proceedings given in their original publication order. They are included in this thesis because they give a good insight in the development of our understanding and contain some interesting experimental results which are not repeated in the later papers. Papers 4 and 5 can be considered ‘matured’ versions of these proceedings and extend or supersede the simplified models that we introduced in them. While Papers 1 to 5 analyse the error propagation of random errors only, Papers 6 and 7 discuss the error propagation of systematic (modelling) errors introduced by radiation damage to the astrometric instrument of Gaia. The fact that the four journal papers were (mostly) written and published in the last year of my PhD studies illustrates that it took quite some time to develop the software and the understanding needed to write them, and above all: to properly define what the problems precisely are. Let us discuss the evolution of all the papers in a bit more detail now.

4.1 The history of the papers

Paper 1: *Spatial correlations in the Gaia astrometric solution*

This paper was written in the spring of 2009 (about two years into my PhD studies) and introduces our scalable AGISLab software together with a first characterization study of correlations between astrometric parameters due to the estimation of attitude parameters.

At the time we could only run simulations up to some $\sim 50\,000$ sources, but since these small scale solutions exaggerate the correlations between sources and the spatial scale on which they appear on the sky, this was an ideal starting point for exploring how these correlations behave. To get good correlation statistics we found that the number of source pairs obtained in a single simulation was not enough. Therefore we combined results from many (Monte-Carlo) simulations with different noise realisations.

Paper 2: *Characterizing the Astrometric Errors in the Gaia Catalogue*

The identification of correlation structures using Monte-Carlo simulations was interesting, but we ultimately would like to *predict* the correlation (and in general the covariance) between any pair of astrometric parameters. At this point we started developing analytical covariance models, which lead to the initial version of the covariance expansion model that was presented in Paper 2 in the spring of 2010. It discusses a truncated and simplified two-term expansion for sources uniformly distributed over the sky and all having the same magnitude. This simple model was able to explain some results from Paper 1 and also qualitatively describe the outcome of Monte-Carlo experiments in which the source and attitude contributions to the errors were separately studied. We knew we were on the right track.

Paper 3: *Efficient calculation of covariances for astrometric data in the Gaia Catalogue*

This paper was written in the spring of 2011. As the research progressed we started thinking about how the covariance model would be used in practice. This led to the more general description of how to compute the variances and covariances of astrophysical quantities that are functions of any number of astrometric parameters. The main question was: how can it be done in a way that is both accurate and computationally feasible? This question has two parts: accuracy and feasibility. The accuracy depends on a number of simplifying assumptions and approximations that were later discussed in Paper 4. The implementation of the covariance model we had at the time was not efficient enough to allow the computation of covariances for large (realistic) numbers of sources, but we could use it to study how sources are connected to each other through the attitude parameters at the times the sources were observed together. Because the series expansion of the covariance model iterates between the attitude intervals at which a source was observed, and the sources that were observed in the observed attitude intervals (and so on), sources can be connected through the attitude parameters even if they are not directly observed together. Because this ‘connectivity’ between source and attitude parameters grows quickly for each new term in the expansion we decided to test the practical feasibility of the model by examining how the number of connections grows as function of expansion term for a realistic simulated set of source observations. This was the most quickly produced paper: the time from the first idea to the final paper was only two weeks.

Paper 4: *Error characterization of the Gaia astrometric solution**I. Mathematical basis of the covariance expansion model*

From August 2011 and onwards we finally had time to work full-time on Paper 4 which contains the mathematical formulation of the covariance model. This has really felt as the most rewarding period of my PhD studies where all the pieces of the last years finally started to come together. At coffee breaks the discoveries of the day before were discussed and ideas exchanged. There were many small ‘breakthroughs’ in this process: the (quite long) expressions we had been using for over two years were reformulated which considerably improved the readability and understanding, the ‘kinematographic’ (step-and-stare) approximation for the attitude model was introduced, some compressed quantities were introduced which save considerable computations in the model, and it was realized that the symmetry of the subsequent expansion terms allowed the sequential computation of higher order expansion terms using the data from the previous term in a computational efficient way (Sect. 5.1 in Paper 4). The whole covariance model was recoded in a multi-threaded efficient way, and clever array indexing was invented to do all the structural array referencing without expensive lookup functions. All of the papers described up till now can be explained using the covariance model that is presented in Paper 4. A small but significant discrepancy between our kinematographic model and the (in simulations and reality used) spline model was considered to be a subject for later study. However, on the side we kept working on an attitude correction term that could be estimated from the data. Within three weeks (only three weeks before my thesis should be handed in to the printer) we had worked out this attitude ‘fudge factor’, which seems to give a significant (if not complete) correction!

Paper 5: *Error characterization of the Gaia astrometric solution**II. Validating the covariance expansion model*

This paper contains a numerical validation of the covariance model presented in Paper 4. It demonstrates that the model works remarkably well. When applying the fudge factor, which we estimated from the data itself, the accuracy of the covariance prediction improves significantly and we find that in the worst simulated case the variance estimation errors are $< 2.5\%$ if a second-order expansion is used ($\alpha = 2$) and $< 1\%$ for a third-order expansion. In the final Gaia solution, using many more primary sources, the attitude errors should be much smaller resulting in estimation errors that are at least 10 times smaller. Within the statistical uncertainties of the experiments, the model also correctly estimates the covariances (or correlations) for pairs of sources separated by $\simeq 0.2^\circ$. As found in Paper 1, these close separations have the highest correlations and are therefore the most important to model correctly. This paper has been written in the last two months of my PhD studies and was only finished and submitted on the day this thesis went to the printer.

Paper 6: *The impact of CCD radiation damage on Gaia astrometry**I. Image location estimation in the presence of radiation damage*

The discussion of Papers 6 and 7 brings us to my involvement in the radiation damage calibration problem. When I started my PhD studies the problem of radiation damage to the Gaia CCDs was just recognized as a serious threat to the (astrometric) performance of Gaia and led to the formation of the Radiation Task Force (RTF) within the Gaia Data Processing and Analysis Consortium (DPAC). In the spring of 2008 I worked for seven weeks at the Institute of Astronomy in Cambridge with Thibaut Prod'homme (Leiden Observatory), Scott Brown (Institute of Astronomy), Michael Weiler (Paris Observatory), Juan Carlos Terrazas and Deborah Busonero (both from Osservatorio Astronomico di Torino) on the analysis of the test results of the first Radiation Test Campaign, under guidance of Floor van Leeuwen (Institute of Astronomy). Halfway through that period the first RTF meeting was held where the results and problems were discussed. We have been to the follow-up meetings every half year ever since, often spending some time working at Cambridge, although at the last meeting only Thibaut and me were left from our initial team. It is very interesting to see the development of the understanding of radiation damage over the last four years: starting from the first radiation damage bias estimates at the first meeting, to a full blown CTI mitigation strategy being implemented in the Gaia pipeline in the last meeting. A major role in this process is taken by the development of the charge distortion model (CDM) where crucial contributions were made by Alex Short (ESA). The CDM is basically a non-linear distortion model that is (loosely) based on the theoretically expected trapping behaviour of electrons over the whole transfer width of the CCD. For a suitable parametrized (and calibrated!) CDM one can in principle compute the non-linear distortion of any (extended) source observed by Gaia. Finding a suitable parametrization and calibration of the CDM has been a continuous endeavour over the last few years. Currently we have a fourth version 'CDM04', but it is likely that improvements can still be made. This was however enough to start on two papers about the impact of CCD radiation damage on Gaia astrometry jointly with the group at Leiden Observatory. Paper 6 has evolved into a very extensive work describing the effect on the image location estimation (i.e., the determination of the observation time) using detailed electron-level Monte-Carlo simulations of the charge transport in the presence of radiation damage in the astrometric CCDs of Gaia. The simulations were carried out using Charge Transfer Inefficiency Effects Models for Gaia (CEMGA, see Prod'homme 2011) software developed by Thibaut Prod'homme (Leiden Observatory). that can simulate the statistical propagation of individual electrons through a Gaia CCDs in very high detail. In the paper we have replicated a large part (but not all) of the Gaia forward modelling process used in the image parameter estimation. We derived the theoretical variance limits and checked that our data indeed comply with them. Then we introduced radiation damage without any attempt to calibrate it and determined what biases this introduced in our estimates (i.e., due to modelling errors). Next we included the CDM model in the forward modelling loop, tried to optimize its parameters for minimal residuals, and examined the residual bias. Based on this study we expect that the current CDM model can perform to requirements

if it is well enough calibrated. The work for this study was effectively done during several intense periods over about two years in which the principal authors have frequently visited each other. Working with Thibaut on this and the next paper has certainly been one of the most fun and exciting parts of my PhD!

Paper 7: *The impact of CCD radiation damage on Gaia astrometry*

II. Effect of image location errors on the astrometric solution

This paper discusses the propagation of the image location errors derived in Paper 6 into the astrometric parameters. This involves creating a model for how the image location errors evolve over time in the mission and time since charge injection (CI). Then an astrometric AGISLab solution is made in which the observations are perturbed by the model and rigorously propagated through the astrometric solution after which the effect on the estimated astrometric parameters can be estimated. To create a sufficiently realistic impression of these effects it was necessary to not scale-down our simulations. Although at the start of my PhD we could only run simulations with at most 50 000 sources, a new computer with 72 Gigabyte (GB) of memory and two super-fast (at the time) dual-core processors, together with my improvements in the memory storage of observations, now allowed us to run simulations for a million sources. This is about the minimum size of a solution that allows for a proper attitude determination using the real Gaia field-of-view size, scan rate, and a reasonable attitude model. We had never tried this number before, and it turned out to be close to the limit of what the machine could handle memory wise, and for us time wise, taking about one hour per iteration and 1.3 days for a complete solution. When combined in the right way, the residuals revealed the systematic variations from the radiation damage model. This allowed us to make a ‘correction’ to the observations based only on the residual data (i.e., on information available in the real mission as well). Although in reality this information would be fed back to improve the calibration of the image location estimation this direct correction method indicates that the astrometric accuracy can likely be recovered within the requirements for the expected radiation dose levels. Additionally we studied the effect of unmodelled ‘disturbing stars’ on the parameter estimates. The idea is that these (mainly faint) stars, that fall in between the CI and the star that is processed, alter the trap occupancy level of the pixels and therefore add additional distortions to the location estimation, and subsequently the astrometric parameters. Our simple model predicts an additional error of the order of $4 \mu\text{as}(\text{yr}^{-1})$ to the astrometric parameters, which would be insignificant for the faint stars, but would for the bright stars require modelling of the disturbing stars to avoid small but significant additional errors. This paper was written in the relatively short time of about four months in the spring and summer of the last year of my PhD studies. With much work still to be done on Papers 4 and 5 there was unfortunately no time to further investigate some of the small but unexplained effects in the data of the residual-based correction, such as the consistently positive mean time residual of the full damage case in Fig. 15 and its systematic pattern in Fig. 16 which seems aligned with the CI phasing. The final words about this paper I want to dedicate to Fig. 20 because to produce it, almost all of the code I have developed over the last four years was needed!



Figure 4.1: The AGISLab logo.

4.2 AGISLab development

AGISLab was initiated by Lennart Lindegren in 2007 (although it was called ‘GaiaLab’ at the time) just before I arrived to Lund in September that year. David Hobbs started in the spring of the same year, and has been developing and programming essential parts of code in both AGISLab and AGIS every since. Almost all numerical (Monte-Carlo) simulations produced for the papers made use of the conjugate gradient algorithm (Bombrun et al. 2011), which significantly reduces the number of iterations needed to converge an AGIS solution. This algorithm was initially developed in AGISLab by Alex Bombrun (ARI, Heidelberg) in collaboration with Lennart Lindegren and David Hobbs. Since 2010 PhD student Daniel Michalik contributes by introducing ways to load and combine previous catalogues with the Gaia results, improving the data compatibility between AGIS and AGISLab, and using AGISLab to simulate the expected astrometric results of the Japanese Nano-Jasmine satellite. My contributions to AGISLab make up a significant part of the code, ranging from object models, interfaces, data structures, statistical analysis tools, monitoring and visualization tools, the accuracy model, an efficient sky density query method, to multi-threading time critical methods, propagating seed values in a multi-threaded system, introducing a flexible properties system, the observation modifier framework, and providing a logo (Fig. 4.1). It is probably fair to say that I spent about half of my time in Lund developing AGISLab, always trying to add efficient, well documented, and user-friendly pieces of code. Note that Lennart also contributed code over the years (in particular the essential ‘scanner’ class), although his biggest contribution is through his ideas given during many great discussions about how to design the different parts of AGISLab. Whenever AGISLab is mentioned in the papers one should be aware that this builds on the work of the above mentioned persons.

AGISLab has been used as a tool for several investigations. Sergei Klioner and Alexey Butkevich (both at the Lohrmann Observatory, Dresden) have used it in the study of estimating the satellite velocity from the observations. Studies of the effect of polarisation on the astrometric parameters have been made by Frederic Raison (ESTEC) and master student Chris Skoog (Lund Observatory). The Mono- and Quadrupole Gravitational

Light Deflection by Jupiter has been investigated by master student Adriaan Ludl (Lund Observatory). The AGISLab scanner is extensively used for time-series analyses by the Gaia photometric calibration unit (CU5) at the Observatory of the University of Geneva (e.g. Varadi et al. 2009). The study of estimating the relativistic PPN γ (light-bending) parameter from the observations has been done by David Hobbs and others (e.g., Hobbs et al. 2010; Lindegren et al. 2012).

4.3 Addressing the main questions

Before going to the summary of each paper I would like to come back to the three main questions that were outlined in Sects. 3.4 and 3.5 and indicate in which papers they are discussed:

- **How do random errors propagate in the astrometric solution?**
This is discussed in Papers 1, 2, and 3 using Monte-Carlo simulations backed up by (simplified) theoretical considerations. Paper 4 provides a much more complete theoretical description and derivation.
- **How can the covariance between any pair of astrometric parameters be estimated?**
Mathematical and practical models are derived in Paper 4. In Paper 5 the practical model is validated.
- **What is the impact of CCD radiation damage on Gaia astrometry?**
This is discussed in two parts. In Paper 6 we estimate the effect on the image location estimation based on detailed electron-level Monte-Carlo simulations of the observation process. In Paper 7 these errors are propagated through the astrometric solution in numerical experiments designed to determine the impact on the astrometry.

Chapter 5

The Papers

The next pages summarize the individual papers and my contribution to each of them.

Paper 1: **Spatial correlations in the Gaia astrometric solution**

B. Holl, D. Hobbs, L. Lindegren (2010)

IAU Symp., Vol. 261, *Relativity in Fundamental Astronomy: Dynamics, Reference Frames, and Data Analysis*, ed. S. A. Klioner, P. K. Seidelmann, and M. H. Soffel, pp. 320–324

In Paper 1 we introduce our scalable AGISLab Java software. We describe its basic scaling properties and how that allows us to characterize Gaia-like simulations with only 3000–30 000 sources in an astrometric solution.

Also we introduce the importance of correlations in the astrometric catalogue, and demonstrate how the average parallax of a cluster of stars is affected in the presence of spatial correlations between the astrometric errors. We continue by explaining that an error in the attitude at a particular time will affect all observations in both fields of view, producing correlations among sources both for small angular separations and for separations of about 106.5° . This is confirmed by Monte-Carlo simulations, from which we estimate the correlation as function of separation on the sky. It is found that (i) the (small angle) correlation length scales with the size of the field-of-view, and (ii) the maximum correlation depends mainly on the number of sources in the field.

My contribution:

The AGISLab Monte-Carlo simulations and data processing have been done by me. The set-up and interpretation of the simulation results have been done in collaboration with Lennart Lindegren and David Hobbs. The text has been written by me with the help from the co-authors.

Paper 2: Characterizing the Astrometric Errors in the Gaia Catalogue

B. Holl, L. Lindegren, D. Hobbs (2011)

EAS Publications Series, Vol. 45, *Gaia: at the frontiers of astrometry*, ed. C. Turon, F. Meynadier, and F. Arenou, pp. 117–122

In Paper 2 we describe our initial version of a covariance model. We divide up the normal matrix in two diagonal source and attitude blocks, and two very sparse off-diagonal blocks that link the source and attitude parameters as defined by the scanning law. Based on this block division we introduce a series expansion of the source covariance block. The first two terms of this expansion are discussed for a simplified case where sources are uniformly distributed over the sky and all have the same magnitude. We find that this is consistent with the result of Paper 1 saying that the maximum correlation (at smallest separations) scales inversely to the number of sources observed per attitude parameter.

Additionally we show that astrometric errors can be separated in a source and an attitude part, due to the estimation of the source and attitude parameters respectively. Hence the covariances can be separated in a source, an attitude and a cross term. Using AGISLab Monte-Carlo simulations we demonstrate that this is indeed the case, and additionally, that by combining many simulations the source and attitude standard errors seem to converge to the same but differently scaled pattern, as predicted by the truncated and simplified two-term expansion.

My contribution

The AGISLab Monte-Carlo simulations and data processing have been done by me. The set-up and interpretation of the simulation results have been done in collaboration with Lennart Lindegren and David Hobbs. The covariance model derivations were done by me and Lennart Lindegren. The text has been written by me with some help from the co-authors.

Paper 3:

Efficient calculation of covariances for astrometric data in the Gaia Catalogue

B. Holl, L. Lindegren, D. Hobbs (2012)

Workshop *Astrostatistics and Data Mining in Large Astronomical Databases*, La Palma, 30 May-3 June 2011, ed. L. Sarro, J. De Ridder, L. Eyer, and W. O'Mullane, in press

In Paper 3 we (partly) address the question of how to compute the covariances between astrophysical quantities, that are a function of a (large) number of astrometric parameters, in a way that is both accurate and computationally feasible. This question has two parts: accuracy and feasibility. The accuracy depends on a number of simplifying assumptions and approximations that are discussed in Paper 4, the practical feasibility of the computation is examined based on the connectivity between source and attitude parameters for an unscaled dataset of sampled source observations at nearly 200 000 positions on the sky. This demonstrates that all sources will be connected to each other (by common field-of-view transits, i.e., common attitude parameters) within three steps. We also include two examples, explaining how to practically compute the covariance for the average parallax of a star cluster, and the acceleration of the solar-system barycentre in a cosmological frame.

We also stress why one would need a model for computing the covariances in the first place. First of all there is data volume: it would take 10^8 TeraByte (TB) to store the full covariance matrix for 10^9 sources. Additionally it is not feasible to compute the elements by a direct inversion of the full normal matrix. Finally, it is anyway clearly desirable that the covariance between any pair of source parameters can be computed from a reduced amount of data (e.g., the final catalogue values themselves complemented with some additional observation statistics). We estimate the minimum amount of data needed to be ~ 2 TB of storage, containing per source and field-of-view transit: (i) the partial derivatives of the along-scan observations with respect to the source parameters (typically 5), (ii) the observation time, and (iii) the combined weight of the observations.

My contribution:

The (initial version of the) covariance model was coded by me, and I did the simulations and data analyses. The set-up and interpretation of the simulation results have been done in collaboration with Lennart Lindegren and David Hobbs. The text was written by me and Lennart Lindegren with help of David Hobbs.

Paper 4:**Error characterization of the Gaia astrometric solution****I. Mathematical basis of the covariance expansion model**

B. Holl, L. Lindegren (2012)

Astronomy & Astrophysics, submitted.

In Paper 4 we provide a mathematical basis for estimating the variance-covariance of any pair of astrometric parameters, and more generally the covariance matrix for multi-dimensional functions of the astrometric parameters. Based on simplifying assumptions (in particular that calibration errors can be neglected), we derive and analyse a series expansion of the covariance matrix of the least-squares solution. A recursive relation for successive terms is derived and interpreted in terms of the propagation of errors from the sources to the attitude and back. We argue that the expansion should converge rapidly to useful precision. The recursion is vastly simplified by using a kinematographic (step-wise) approximation of the attitude model. We find that low-order approximations of arbitrary elements from the covariance matrix can be computed efficiently in terms of a limited amount of pre-computed data representing compressed observations and the structural relationships among them. It is proposed that the user interface to the Gaia Catalogue should provide the tools necessary for such computations.

My contribution:

The content of this paper has been developing over a period of perhaps three years. The structure and text was initially created by me and Lennart Lindegren although the (semi-)final paper that is presented in the thesis was largely re-formulated by Lennart. For Sect. 3.2 I made some initial derivations which did not end up in the final paper. At the time of writing Sect. 5.1 and 5.2 I was heavily coding the covariance model, and it describes the result of our discussions and mutual insights. Concerning the appendices, Lennart introduced the graph theorem in Appendix B, and did the numerical (Matlab) experiments that lead to the data in Appendix C. The development and testing of the fudge factor discussed in Appendix C (in relation to the validation experiments of Paper 5) was however a joint endeavour. I produced the scanning law and matrix figures in the paper.

Paper 5:

Error characterization of the Gaia astrometric solution

II. Validating the covariance expansion model

B. Holl, L. Lindegren, D. Hobbs (2012)
Astronomy & Astrophysics, submitted.

In Paper 5 we aim to determine to what extent the covariance model introduced in Paper 4 provides an accurate representation of the expected random errors in the astrometric solution. We simulate the astrometric core solution for Gaia by making least-squares solutions of the astrometric parameters for one million stars and the attitude parameters for a five-year mission, using nearly one billion simulated elementary observations for a total of 26 million unknowns. Two cases are considered: one in which all stars have the same magnitude, and another with 30% brighter and 70% fainter stars. The resulting astrometric errors are statistically compared with the model predictions. In all cases considered, and within the statistical uncertainties of the numerical experiments (typically below 0.4%), the theoretically calculated variances and covariances are consistent with the simulations. To achieve this it is however necessary to expand the covariances to at least third or fourth order, and to apply a (theoretically motivated and derived) “fudge factor” in the kinematic model. We find that the model provides a feasible method to estimate the covariance of arbitrary astrometric data, accurate enough for most applications, and as such it should be available as part of the user’s interface to the Gaia Catalogue. A main assumption in the current model is that the observational errors are uncorrelated (e.g., photon noise), and further studies are needed of how correlated modelling errors, in particular in the attitude, can be taken into account.

My contribution:

Most of the paper was written by Lennart Lindegren, although parts of Sect. 2 and the complete Appendix C were written by me, and Appendix B by David Hobbs. Bits and pieces of an original draft by me survived in many places. For this paper I implemented and tested the covariance model and analyses methods, which required a very substantial amount of work. The development of the theoretical covariance model and the practical implementation for Papers 4 and 5 progressed in parallel, as well as the final introduction of the fudge factor. On this level there was intensive collaboration over a period of several months between Lennart Lindegren and me. Most of the data generation and reduction was done by me.

Paper 6:
**The impact of CCD radiation damage
on Gaia astrometry I. Image location estimation
in the presence of radiation damage**

T. Prod'homme, **B. Holl**, L. Lindegren, A. G. A. Brown (2011)
Monthly Notices of the Royal Astronomical Society, in press,
DOI: 10.1111/j.1365-2966.2011.19934.x

In Paper 6 we discuss the effect of radiation damage on the image location estimation process. The requirements set a stringent constraint on the accuracy of the estimation of the location of the stellar image on the CCD for each observation: e.g., 0.3 milli-arcseconds (mas) or 0.005 pixels for the same $V = 15$ G2V star. However the Gaia CCDs will suffer from charge transfer inefficiency (CTI) caused by radiation damage that will degrade the stellar image quality and may degrade the astrometric performance of Gaia if not properly addressed. For the first time at this level of detail, the potential impact of radiation damage on the performance of Gaia is investigated. In this first paper we focus on the evaluation of the CTI impact on the image location accuracy using a large set of CTI-free and damaged synthetic Gaia observations supported by experimental test results. We show that CTI decreases the stellar image signal-to-noise ratio and irreversibly degrades the image location estimation precision. As a consequence the location estimation standard errors increase by up to 6% in the Gaia operating conditions for a radiation damage level equivalent to the end-of-mission accumulated dose. We confirm that in addition the CTI-induced image distortion introduces a systematic bias in the image location estimation (up to 0.05 pixels or 3 mas in the Gaia operating conditions). Hence a CTI mitigation procedure is critical to achieve the Gaia requirements. We present a novel approach to CTI mitigation that enables, without correction of the raw data, the unbiased estimation of the image location and flux from damaged observations. We show that its current implementation reduces the maximum measured location bias for the faintest magnitude to 0.005 pixels ($\sim 4 \times 10^{-4}$ pixels at magnitude 15) and that the Gaia image location estimation accuracy is preserved.

My contribution:

The content of this paper has been developing during several intense periods spread over about two years. Most of the text was written by Thibaut Prod'homme. I have written large parts of the Sect. 1 (especially the Gaia performance part and Table 1), 3.1–3.3, and 5 (also some initial versions of the calibration diagrams shown in it), have helped to improve other parts of the paper, and had a large influence on the structure of the paper. Additional improvements and suggestions were made by Lennart Lindegren and Anthony

Brown. The amount of programming needed to do the analyses was quite substantial and I estimate that I did about half of it (this only refers to the additional coding in CEMGA that was needed for this paper). I coded a large part of the methods related to LSF modelling, sampling, fitting and reconstruction, and the methods needed for the computation of the Cramér-Rao bound for different profiles. I also worked on a maximum likelihood fitting method for the CDM parameters, but it was not stable enough to be used in the paper. All coding development was discussed between Thibaut and me, and the final data processing methods were developed together. The set-up and interpretation of the experiments were intensely discussed with Lennart and Anthony.

Paper 7:
**The impact of CCD radiation damage
on Gaia astrometry II. Effect of image location errors
on the astrometric solution**

B. Holl, T. Prod'homme, L. Lindegren, A. G. A. Brown (2012)
Monthly Notices of the Royal Astronomical Society, accepted.

Paper 7 is the second and last in a study aiming at characterizing and quantifying the impact of CCD radiation damage on Gaia astrometry. 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 Paper 6. 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.

My contribution:

This paper and its data were created in a time-span of about four months. A large part of the text was initially written by me, except for the Abstract and Sects. 2.2.1–2.2.4, part of 4.1, and 5 which were written by Thibaut Prod'homme. The overall text has occasionally been re-formulated and extended by Lennart Lindegren, with additional improvements and suggestions from Anthony Brown. All the line-plots were created by Thibaut except for Fig. 5 which was created by me. The contour-plots and the Hammer-Aitoff all-sky plots were created by me. The amount of programming for the analyses was quite substantial and I estimate that I did about 90% of it (this only refers to the additional coding in AGISLab for this paper). All steps of the experiment were intensely discussed among the co-authors.

Popular Summary

From Hipparchus to Gaia

Astronomy must be one of the oldest sciences known to mankind given the artefacts left behind by prehistoric cultures that mark and predict celestial events. The measurement of the positions and motions of celestial bodies on the sky is now called *astrometry*, but most observational activities in *astronomy* up until the mid-19th century would be classified as *astrometry* in the modern usage of the word. Some of the oldest records of astrometric observations date back to Hipparchus (ca 130 BC) and Ptolemy (ca 150 AD), who measured the positions of about 1000 stars with an accuracy of about 0.5° (roughly the angle of the full Moon on the sky). This accuracy was only improved after more than half a millennium by Islamic astronomers and later by William IV, Landgrave of Hesse-Kassel and Tycho Brahe, who catalogued some 1000 stars around 1600 AD with an accuracy of about 1 arcmin (1/30th of the full Moon). This accuracy was gradually improved in the 400 years that followed, until the introduction of space techniques with the Hipparcos satellite mission, see figure 5.3. This caused an increase in accuracy by a factor of 100 and allowed for the first time the accurate measurement of distances to many thousand nearby stars. Nearby here means up to about 500 light years distance from the Sun, which compared to the 25 000 light years to the centre of our Milky Way is indeed ‘nearby’. The measurement of distances is extremely important for deriving many stellar quantities like the total amount of light a star emits, its mass, and age, but also for calibrating methods that can be used to measure the much larger distances to other galaxies. Besides positions and distances, also velocities of stars were measured by the Hipparcos satellite, which allows us to study the forces that move stars around in the Galaxy. The apparent motion of stars on the sky is so small that the displacement of stars during a century is still indistinguishable with the naked eye, except for a handful of nearby high-velocity stars.

The next big jump in astrometric accuracy and number of stars will come with the launch of the European Gaia satellite in 2013. It will again improve the accuracy by a factor of 100 and measure the positions, motions, and distances for a staggering 1 000 000 000 stars




Figure 5.1: An historical overview of the errors of star positions in the most accurate catalogues. Tycho Brahe achieved a jump in accuracy followed by four centuries with more gradual improvement. Another much larger jump in accuracy was obtained by the ESA satellite giving the Hipparcos and Tycho-2 Catalogues containing a total of 2.5 million stars. Parallaxes (a measure for stellar distances) are also measured by Hipparcos and Gaia with the same accuracy as positions. Courtesy: Erik Høg, 1995/2005, Copenhagen University Observatory. The historical account in the first paragraph of this summary is based on Høg (2009).

(more than one star per European citizen) in our Milky Way and the small nearby galaxies that surround it. Effectively 1% of all the stars in our Galaxy will be mapped which will allow us to create a three-dimensional map of the Galaxy than can be run both forward and backward in time to study its dynamical evolution. The accuracy that Gaia can reach is of the order of 10 *micro*-arcsecond which is about the angular size of a Swedish Krona on the Moon as seen from the Earth! To compare all this to what Tycho Brahe achieved with his naked-eye observations: Gaia will measure a million times more stars, a million times fainter, and a million times more accurate.

Gaia as a merry-go-round

All of these numbers are impressive, but you might wonder how hard (or easy) it is to map a billion stars with a satellite. Let us look at a simple analogue. Imagine one evening you are going to the local fair and taking a seat on a merry-go-round with a pair of binoculars in one hand and a stopwatch in the other. Pointing the binoculars outwards and looking through it, you will see everything move by in one direction. The motion of the merry-go-round allows you to scan the area around it along a large circle. The binoculars will limit the area you can see, but what you see is greatly magnified and you see a lot of detail. Now every time you see a person through the binoculars you write down the exact time he/she crosses the centre of your field of view. We assume you are good at remembering faces and group the timings per person¹. After you finished the ride you have a list with many timings per person. Now the question is: can we reconstruct the paths of all the listed persons based on these observations? The answer is yes, although it depends on the assumption that most people are moving in straight lines.

Modelling the observations

To reconstruct the movements of people we make use of *models*: descriptions of reality that allow us to *predict* when a person would be observed. Let us for simplicity assume that everyone is walking around the merry-go-round at his own speed (or standing still). According to this model every person should be clocked at regular time intervals, for each turn of the merry-go-round, but this interval will be slightly different depending on how the person moves. By comparing the time intervals for different persons, you will be able to determine how they move relative to each other. However, in order to make precise predictions about the timings we need to model a few more things. First of all the rotation of the merry-go-round: for example, we need to know its orientation at an agreed instant of time, as well as the rotation speed and perhaps whether it is speeding up or slowing down. We also need to model the pointing of the binoculars with respect to the merry-go-round. Although you probably tried to hold the binoculars steady, it is likely that there was some variation in the pointing that needs to be taken into account.

We have now specified the problem in terms of three different models: how people are moving around the merry-go-round ('person model parameters'), the rotation of the merry-go-round ('orientation model parameters'), and how your binoculars were wobbling over time ('binocular model parameters'). We use the word *parameters* here to indicate the different elements of a model that can have a certain value, for example the rotation speed in the orientation model. So, how will these models help us solve our problem?

¹The fact that you are observing people with a pair of binoculars and a stopwatch from a merry-go-round will probably make people look at you which will make this task easier.

The iterative merry-go-round solution

Let's assume we start with a set of parameters that predicts observation times that are not too far off from your observed list. The challenge is now to improve these parameters such that the differences between the predicted and observed times get as small as possible. One way of solving this problem is to improve the parameters for each model separately and to repeat (iterate) this process many times. That is, we first assume that we know exactly how the merry-go-round rotates and how the binoculars are pointed. Using the observed timings this allows us to reconstruct the movement of each person separately. Then, assuming that we know the movements of all persons and the pointing of the binoculars, we can determine how the merry-go-round rotated. Finally, assuming that we know the movements of people and the rotation of the merry-go-round, we can reconstruct the wobbles of the binoculars. This whole process must be iterated, because the assumptions made in each step are obviously not quite correct, but they get better with each iteration.

By doing many more of these cycles (person – orientation – binoculars parameter updates) we will reach a point when adjusting the parameters does not further reduce the differences between the predicted and observed timings. If our models are a good description of what happened in the real world these time differences should not be larger than expected from the uncertainty in your observations. The amount by which you can change each parameter without significantly changing these differences is a measure of its *uncertainty* (how much you can expect the parameter to differ from its true value). This means that the uncertainty of each model parameter depends on the uncertainty in your observations. Unless you can measure the timings with infinite precision (and your models are perfect) you cannot determine the path of each person exactly.

What this says about Gaia

The part that you probably guessed is that the merry-go-round is Gaia, constantly spinning in space. In reality however it is also changing its spin axis such that it is sometimes even rotates upside down to talk in fair-terms. It is actually more like one of those rotating dare-devil attractions, although in super slow motion since Gaia goes round only once every six hours. Its motion is modelled in three dimensions instead of the one we used here.

The binoculars represent the telescopes inside Gaia that can observe an area slightly larger than the full moon at each moment, and the binocular model in reality is an instrument model that describes the minute movements of the two telescopes and the about 100 digital cameras inside Gaia.

The moving people are the stars, of which about 80% effectively move in straight lines

during Gaia's mission duration of five years, so that model is very similar to what we used for the people. Many of the stars will however have a close partner causing them to make an additional continuous wobble, which on the fairground could be young in-love couples circling each other's neck. Nearby stars having multiple (heavy) planets will also show a complicated motion that can be compared to that of break dancers on the fair (you need to observe them for quite some time before you can follow their movements). In the iterative merry-go-round solution it will improve the estimated orientation and binocular model parameters if we leave out the observations of the persons that do not seem to follow a straight line. In the same way only the stars with simple straight motions are used to determine the orientation and instrument model of Gaia.

Radiation damage and alcohol

Up till now I described some elements of how Gaia will function under ideal conditions, but I did not get to the most difficult problem yet: radiation damage. Particles from the sun (fast protons expelled during solar flares) affect Gaia as alcohol affects your vision: it doesn't make it any better. Starting to drink a beer while doing your observations on the merry-go-round will have you make more errors, but above all: your reaction time will get longer. It is also likely that your reaction time will depend on how visible (brightly lit) the person is. You can imagine that it will be a lot harder to use this set of observations to accurately model the rotation of the merry-go-round, binocular direction, and the path of each person.

The beer in this example stands for radiation damage to the cameras of Gaia which also causes the observations to be delayed differently for bright and faint stars. Radiation damage in Gaia will be handled in much the same way as the other effects discussed above: it needs a model that can predict how each image is affected, and the parameters of this model then have to be adjusted to minimize the differences between the observed and predicted timings. The main problem here is that the model is very complicated and that we do not completely understand how the radiation damage works. And, unlike the effects of beer, there is no way for Gaia to 'sober up': the radiation damage will keep accumulating during the whole mission.

What I have been doing

Contrary to what you might think, I did not spend my PhD studies riding a merry-go-round and drinking beer (well, not all of the time at least). What I did was trying to figure out how all these different models and parameters affect the accuracy by which we can estimate the positions, distances and velocities of the one billion stars that Gaia will observe. To do this we use the Astrometric Global Iterative Solution (AGIS) to improve the model parameters, which works in a similar way as the iterative merry-go-round solution.

A lot of my work has involved programming models and doing experiments with AGIS. Because Gaia will only be launched in 2013 we have to simulate observations from which we then estimate the parameters for each star using AGIS. For example, to know how accurately the distance to stars can be estimated we can make many different simulations (which would be similar to having many Gaia missions) and computing for each simulation how large the error is between the estimated and the true distance to each star. Looking at the spread of these errors for the many simulations tells you how accurately the distance can typically be measured for the single real Gaia mission. Also I have worked on a model that can compute the accuracies directly (without using simulations). When the Gaia catalogue will be published around 2020 it will be a main source of data for astronomers that want to study the structures in our Milky Way, for example by estimating distances and motions of clusters of stars. In collaboration with Thibaut Prod'homme at Leiden Observatory I have also made detailed studies of the effects of radiation damage, where we examined how well the Gaia results can be corrected for these effects.

I could certainly not have done all this work by myself. Besides the Gaia team in Lund I have collaborated with many people in Europe, in particular at Leiden Observatory, the Institute of Astronomy in Cambridge, and the ESA centres in Madrid and Noordwijk. From 2007 to 2010 I was part of a Marie Curie research training network called 'European Leadership in Space Astrometry' (ELSA). Through the network and the collaborations I made many friends with which I shared the occasional beer.

Next time you find yourself slightly tipsy on a merry-go-round, look up to the stars, and realize that what you are doing is not *that* different from mapping a billion stars with Gaia...

Populärvetenskaplig sammanfattning

Från Hipparchos till Gaia

Astronomi måste vara en av de äldsta vetenskaperna att döma av förhistoriska artefakter som visar och förutsäger skeenden på himlen. Mätningar av himlakropparnas positioner och rörelser kallas numera astrometri, men de flesta astronomiska observationer fram till mitten av 1800-talet skulle klassificeras som astrometri i ordets moderna betydelse. Några av de äldsta bevarade astronomiska observationerna går tillbaka till Hipparchos (ca 130 f.Kr.) och Ptolemaios (ca 150 e.Kr.), vilka mätte positionerna för omkring 1000 stjärnor med en noggrannhet av $0,5^\circ$ (ungefär lika med den vinkel som fullmånen upptar). Denna noggrannhet överträffades inte förrän efter ett halvt årtusende av arabiska astronomer och senare av lantgreve Wilhelm IV av Hesse-Kassel samt av Tycho Brahe, som omkring år 1600 katalogiserade drygt 1000 stjärnor med en noggrannhet på omkring 1 bågminut ($1/30$ av fullmånens diameter). Noggrannheten förbättrades gradvis under de kommande 400 åren, till dess att rymdtekniken togs i bruk genom Hipparcos-satelliten, se Fig. 5.2. Denna ökade noggrannheten 100-falt och gjorde det för första gången möjligt att mäta noggranna avstånd till många tusen närbelägna stjärnor. Närbelägna betyder här upp till ett avstånd på ca 500 ljusår, vilket verkligen är "nära" jämfört med de 25 000 ljusåren till Vintergatans centrum. Avståndsbestämningar är mycket viktiga för att härleda många av stjärnornas egenskaper, exempelvis hur mycket ljus en stjärna utstrålar, dess massa och ålder, men även för att kalibrera metoder som kan mäta de mycket större avstånden till andra galaxer. Hipparcos mätte, förutom stjärnornas positioner och avstånd, även deras rörelser. Detta gör det möjligt att undersöka de krafter som sätter stjärnorna i rörelse i vår galax. Stjärnornas skenbara rörelser på himlen sker så långsamt att det knappast är möjligt att uppfatta lägesförändringen med blotta ögat ens under ett århundrade, utom för en handfull närbelägna höghastighetsstjärnor.

Nästa stora språng i astrometrisk noggrannhet kommer efter uppskjutningen av den eu-




Figure 5.2: En historisk översikt av positionsfelen i de mest noggranna stjärnkatalogerna. Den horisontella skalan är årtalet, den vertikala visar felens storlek uttryckta i bågsekunder. Tycho Brahe åstadkom ett språng i noggrannheten vilket följdes av en mera gradvis förbättring under fyra århundraden. Ett nytt, ännu större språng åstadkoms av ESA-satelliten som gav oss Hipparcos-katalogen och Tycho-2 med totalt 2,5 miljoner stjärnor. Hipparcos och Gaia ger även parallaxer (ett mått på stjärnors avstånd) med samma noggrannhet som positionerna. Diagram av Erik Høg, 1995/2005, Copenhagen University Observatory. Den historiska bakgrunden i sammanfattningens första stycke baseras på Høg (2009).

ropeiska Gaia-satelliten 2013. Denna kommer att förbättra noggrannheten med ytterligare en faktor 100 och mäta positioner, rörelser och avstånd för så många som 1 000 000 000 stjärnor (mer än en för varje medborgare i Europa) i vår Vintergata och i de små galaxerna i dess omedelbara närhet. Nästan 1% av alla stjärnor i vår galax kommer att kartläggas, vilket gör det möjligt att skapa en tredimensionell bild av Vintergatan som kan köras både fram- och baklänges i tid för att undersöka dess dynamiska utveckling. Gaia kan uppnå en noggrannhet av storleksordningen 10 *mikro*-bågsekunder, vilket ungefär är vinkelstorleken av en svensk enkrona på månen, sedd från jorden! För att jämföra allt detta med vad Tycho Brahe åstadkom genom observationer med blotta ögat: Gaia kommer att mäta en miljon gånger så många stjärnor, en miljon gånger ljussvagare, och en miljon gånger noggrannare.

Gaia som en karusell

Alla dessa tal är imponerande, men läsaren kanske undrar hur svårt (eller lätt) det är att kartlägga en miljard stjärnor med en satellit. Låt oss göra en enkel liknelse. Förställ dig att du besöker den lokala nöjesparken en kväll och tar plats på en karusell med en kikare i ena handen och ett stoppur i den andra. Om du riktar kikaren utåt ser du allting passera förbi i samma riktning. Karusellens rörelse låter dig svepa med blicken längs en stor cirkel. Kikaren begränsar ditt synfält men det du ser är starkt förstorat och du kan se en massa detaljer. Tänk dig nu att, varje gång du ser en person genom kikaren, du antecknar den exakta tidpunkten när personen passerar genom mitten av synfältet. Vi antar att du har ett gott minne för ansikten och kan gruppera tidpunkterna för varje person². Efter avslutad tur har du en lista med många tidpunkter för varje person. Frågan är nu: kan vi rekonstruera hur varje person på listan har rört sig, baserat på dessa iakttagelser? Svaret är ja, fast det hänger på antagandet att folk i allmänhet rör sig i räta linjer.

Modellering av observationerna

För att rekonstruera folks rörelser använder vi oss av *modeller*: förenklade beskrivningar av verkligheten som låter oss *förutsäga* när en person borde observeras. Antag för enkelhetens skull att var och en promenerar omkring karusellen i sin egen hastighet (eller står stilla). Enligt denna modell borde varje person bli synlig i jämna tidsintervall, för varje varv karusellen gör, men tidsintervallen kommer att vara något olika beroende på personens rörelse. Genom att jämföra tidsintervallen för olika personer, är det möjligt att bestämma hur de rör sig i förhållande till varandra. För att förutsäga de exakta tidpunkterna måste vi dock modellera ytterligare några saker. Först och främst karusellens rotation: vi behöver t.ex. veta dess orientering vid en viss, överenskommen tidpunkt, liksom dess rotationshastighet och eventuellt om den ökar eller minskar i hastighet. Vi behöver också modellera hur kikaren är riktad i förhållande till karusellen. Även om du försökte hålla kikaren stadigt, är det troligt att dess inriktning varierade något, vilket vi får ta hänsyn till.

Vi har nu beskrivit problemet med hjälp av tre olika modeller: hur folk rör sig kring karusellen (“personmodellen”), karusellens rotation (“orienteringsmodellen”) samt hur kikaren vinglade i förhållande till karusellen (“kikarmodellen”). Varje modell har sina “parametrar”, storheter som kan tilldelas ett visst värde, t.ex. rotationshastigheten i orienteringsmodellen. Hur kan dessa modeller hjälpa oss att lösa problemet?

²Det faktum att du observerar folk med kikare och stoppur från en karusell får dem antagligen att titta extra noga på dig, vilket bara gör uppgiften lättare.

Den iterativa karusellösningen

Antag att vi börjar med en uppsättning parametervärden som förutsäger observations-tidpunkter som någorlunda överensstämmer med dina anteckningar. Utmaningen är att förbättra dessa parametervärden så att skillnaderna mellan de förutsagda och observerade tidpunkterna blir så små som möjligt. En metod för detta är att förbättra parametrarna för en modell i taget och sedan upprepa (iterera) denna process många gånger. Vi antar således först att vi känner exakt hur karusellen rör sig och hur kikaren riktas. Från de observerade tidpunkterna kan vi då rekonstruera folks rörelser, en person i taget. Genom att sedan anta att vi känner folks rörelser och kikarens riktning, kan vi bestämma hur karusellen roterade. Slutligen, genom att anta att vi känner folks rörelser och karusellens rotation kan kikarens vinglande rekonstrueras. Hela proceduren måste itereras, eftersom antagandena i varje steg uppenbarligen inte är helt korrekta, men de blir bättre för varje iteration.

Efter många fler sådana cykler (uppdatering av person – orientering – kikaparametrar) kommer vi till en punkt där skillnaderna mellan de förutsagda och observerade tiderna inte går att minska ytterligare. Om våra modeller är en bra beskrivning av verkligheten bör dessa skillnader inte överstiga vad vi kan förvänta oss på grund av osäkerheten i observationerna. Genom att undersöka hur mycket det går att ändra parametrarna utan att dessa skillnader blir nämnvärt större får man ett mått på *osäkerheten* i varje parametervärde (hur mycket det kan tänkas avvika från det sanna värdet). Detta innebär att osäkerheten i varje parameter beror på osäkerheten i observationerna. Om inte tidpunkterna har mätts med oändlig noggrannhet (och modellerna är perfekta) går det inte att exakt fastställa hur folk har rört sig.

Vad detta säger oss om Gaia

Du har antagligen redan gissat att Gaia är karusellen, ständigt roterande i rymden. I verkligheten är Gaias rotationsaxel inte fast, utan ändrar sig med tiden så att satelliten ibland snurrar upp-och-ner (i jämförelse med nöjesparken). Den är mer lik de vilt snurrande attraktioner som finns på mer avancerade tivolin, fast i "ultra-slow motion" eftersom Gaia behöver sex timmar för att göra ett varv.

Kikaren motsvarar teleskopen inuti Gaia som i varje ögonblick observerar ett område något större än fullmånen, och kikarmodellen är i verkligheten en instrumentmodell som beskriver mikroskopiska förskjutningar av de två teleskopen och de mer än 100 digitala kamerorna i Gaia.

Människorna som promenerar omkring på nöjesfältet är stjärnorna, av vilka omkring 80% i praktiken rör sig i räta linjer under Gaias femåriga livslängd. Många stjärnor har dock en nära följeslagare, som ger dem en extra svängande rörelse, likt nyförälskade par som

vrivlar runt i varandras armar. Närbelägna stjärnor med flera (tunga) planeter uppvisar ännu mer komplicerade rörelser, att jämföras med utövare av breakdance (man får observera dem en stund innan man kan följa rörelserna). I den iterativa karusellösningen förbättras bestämningen av orienterings- och kikarparametrarna om man utelämnar observationerna av personer som inte tycks följa en rät linje. På samma sätt används bara stjärnor med enkla, rätlinjiga rörelser för bestämningen av Gaias orienterings- och instrumentparametrar.

Strålningsskador och alkohol

Hittills har jag beskrivit vissa delar av hur Gaia fungerar under idealiska förhållanden, utan att komma in på det svåraste problemet: strålningsskador. Partiklar från solen (snabba protoner från häftiga utbrott på solytan) påverkar Gaia ungefär som alkohol påverkar synen: den blir inte bättre. Att börja dricka öl under observationerna på karusellen leder till att du gör fler fel, men framförallt blir din reaktionsförmåga långsammare. Det är också troligt att reaktionstiden beror på hur väl synlig (upplyst) personen är. Det är begripligt om det blir betydligt svårare att använda sådana observationer för att noggrant beskriva karusellens rotation, kikarens inriktning och rörelsen hos varje människa.

Ölen står här för solpartiklarnas påverkan på Gaias kameror, vilken också orsakar att observationerna fördröjs olika mycket för ljusa och svaga stjärnor. Effekterna av partikelstrålningen kommer att hanteras på liknande sätt som de andra effekter som diskuterats ovan: det behövs en modell som kan förutsäga hur varje bild påverkas, och modellens parametrar får sedan justeras så att skillnaderna mellan observerade och beräknade data minimeras. Huvudproblemet är att modellen är mycket komplicerad och att vi inte förstår fullt ut hur strålningseffekterna verkar. Dessutom, i motsats till ölpåverkan, finns det inget sätt för Gaia att "nyktra till": det tillkommer hela tiden nya strålningsskador under Gaias livslängd.

Vad jag har gjort

I motsats till vad man kan tro har jag inte tillbringat min doktorandtid med att åka karusell och dricka öl (i varje fall inte hela tiden). Vad jag har gjort är att försöka förstå hur alla dessa olika modeller och parametrar påverkar den noggrannhet med vilken vi kan uppskatta positionerna, avstånden och hastigheterna för den miljard stjärnor som Gaia kommer att observera. För att förbättra modellparametrarna har vi använt AGIS (Astrometric Global Iterative Solution), som fungerar på ett liknande sätt som den iterativa karusellösningen. Mycket av mitt arbete har gällt programmering av modeller och experiment med AGIS. Eftersom Gaia inte kommer att skjutas upp förrän 2013 måste vi simulera observationer från vilka stjärnparametrarna kan uppskattas med hjälp av AGIS. För att exempelvis ta reda på hur noggrant man kommer att kunna bestämma avstånden till stjärnor, genomför

vi många olika datorsimuleringar (vilket motsvarar att vi hade många Gaia-satelliter), och beräknar för varje simulering hur stora felen är mellan de uppskattade avstånden och de verkliga. Jag har också arbetat på en modell för direkt beräkning av osäkerheterna (utan att använda simuleringar). När Gaia-katalogen publiceras omkring 2020 kommer den att vara en huvudkälla för astronomer som vill studera Vintergatans strukturer, exempelvis genom att uppskatta stjärnhopars avstånd och rörelser. I samarbete med Thibaut Prod'homme vid Leiden Observatory har jag även i detalj studerat effekterna av strålningsskador, varvid vi undersökt hur väl Gaia-resultaten kan korrigeras för dessa effekter.

Jag skulle knappast kunna utföra allt detta ensam. Förutom med Gaia-gruppen i Lund har jag samarbetat med ett antal personer i Europa, särskilt vid Leiden Observatory, Institute of Astronomy i Cambridge och ESA-anläggningarna i Madrid och Noordwijk. Från 2007 till 2010 var jag en del av ett Marie Curie forskningsnätverk, European Leadership in Space Astrometry (ELSA). Tack vare dessa nätverk har jag fått många vänner med vilka jag emellanåt tagit en öl.

Nästa gång du åker karusell, kasta en blick upp mot stjärnorna och tänk på att det inte är alltför olikt att kartlägga en miljard stjärnor med Gaia...

Med tack till Lennart Lindegren för översättningen.

Populair wetenschappelijke samenvatting

Van Hipparchus tot Gaia

Sterrenkunde is waarschijnlijk één van de oudst bekende wetenschappen van de mensheid. Prehistorische culturen waren al gefascineerd door de manier waarop de sterren, planeten en kometen zich aan de hemel voortbewegen. Dit kan worden afgeleid uit bewaard gebleven bouwwerken die hemelse gebeurtenissen markeren en voorspellen zoals de piramides in Egypte en de Stonehenge steenformatie in Engeland. Het meten van de posities en bewegingen van hemellichamen aan de hemel wordt tegenwoordig aangeduid met *astrometrie*, al kunnen de meeste observationele activiteiten in de *astronomie* tot het midden van de 19e eeuw worden aangemerkt als *astrometrie* in de moderne zin van het woord. De oudste vermeldingen van astrometrische metingen zijn van Hipparchus (ca. 130 v.Chr.) en Ptolemaeus (ca. 150 na Chr.), die de positie van ongeveer 1000 sterren gemeten hebben met een nauwkeurigheid van ongeveer een halve booggraad (ruwweg de hoek van de volle maan aan de hemel). Deze nauwkeurigheid zou pas echt worden verbeterd na meer dan een half millennium door Islamitische astronomen en later door Willem IV, Graaf van Hessen-Kassel en Tycho Brahe, die (ca. 1600 na Chr.) de ongeveer 1000 sterren catalogiseerden met een nauwkeurigheid van ongeveer één boogminuut (1/30 van de volle maan). Deze nauwkeurigheid is in de 400 jaar die volgden geleidelijk aan verbeterd, tot het moment dat ruimte technologie zijn intrede deed met het verschijnen van de Hipparcos satelliet, zie afbeelding 5.3. Dit veroorzaakte een verhoging in nauwkeurigheid met een factor van 100 en maakt het mogelijk om nauwkeurige metingen te doen van de afstanden tot vele duizenden nabijgelegen sterren. Nabijgelegen betekent hier tot op ongeveer 500 lichtjaren afstand van de zon, dit is in vergelijking met de 25 000 lichtjaar naar het centrum van onze Melkweg inderdaad ‘dichtbij’. Het meten van afstanden is uiterst belangrijk voor het afleiden van eigenschappen van sterren, zoals de totale hoeveelheid licht die een ster uitzendt, zijn massa en leeftijd, maar ook voor het kalibreren




Figure 5.3: Een historisch overzicht van de nauwkeurigheid waarmee sterposities zijn gemeten in de meest accurate catalogi. Tycho Brahe maakte een sprong voorwaarts in nauwkeurigheid gevolgd door vier eeuwen met meer geleidelijke verbetering. Een andere veel grotere sprong in nauwkeurigheid werd verkregen door de ESA satelliet die de Hipparcos en de Tycho-2 catalogi voortbracht met daarin een totaal van 2,5 miljoen sterren. Parallaxen (een maat voor stellaire afstanden) werden ook gemeten door de Hipparcos satelliet met een zelfde nauwkeurigheid als de ster posities aan de hemel. Dit zal ook gelden voor de opvolger van Hipparcos: de Gaia satelliet. Afbeelding: Erik Høg, 1995/2005, Copenhagen University Observatory. De historische introductie in de eerste paragraaf is gebaseerd op Høg (2009).

van methoden die kunnen worden gebruikt voor het meten van de veel grotere afstanden tot andere sterrenstelsels.

Naast posities en afstanden, werden ook snelheden van de sterren gemeten door de Hipparcos satelliet. Hiermee kunnen we de krachten bestuderen die de beweging van de sterren in onze Melkweg bepalen. De schijnbare beweging van de sterren aan de hemel is zo klein dat de verplaatsing van de sterren tijdens een eeuw nog steeds niet te onderscheiden is met het blote oog, met uitzondering van een handvol ‘nabijgelegen’ sterren die met grote snelheid bewegen.

De volgende grote sprong voorwaarts in astrometrische nauwkeurigheid en aantal sterren

zal komen met de lancering van de Europese Gaia satelliet in 2013. Het zal opnieuw de nauwkeurigheid een factor 100 verhogen bij het meten van de sterposities, beweging en afstanden voor een ongelofelijk aantal van één miljard (1 000 000 000) sterren (meer dan één ster per Europese burger) in onze Melkweg en de kleine nabijgelegen sterrenstelsels die onze Melkweg omringen. Effectief worden 1% van alle sterren in onze Melkweg vastgelegd, wat het mogelijk maakt om een driedimensionale 'kaart' van de Melkweg te maken. Deze kaart kan, in de tijd, zowel voor- als achteruit gespoeld worden waardoor we kunnen bestuderen hoe de beweging van de sterren in onze Melkweg is geëvolueerd. De nauwkeurigheid die Gaia kan bereiken is in de orde van 10 *micro*boogseconde, dit is de hoekgrootte van een Euro op de maan, gezien vanaf de aarde! Om dit te vergelijken met alles wat Tycho Brahe met zijn blote-oog-waarnemingen bereikte: Gaia zal een miljoen keer meer sterren, tot een miljoen keer zwakker, en een miljoen keer nauwkeuriger meten.

Gaia als een draaimolen

Al deze getallen klinken misschien indrukwekkend, maar u vraagt zich mogelijk af hoe moeilijk (of eenvoudig) het is om met een satelliet de positie, afstand, en snelheid van een miljard sterren te bepalen. Laten we eens kijken naar een eenvoudige analogie. Stelt u zich het volgende voor: op een avond gaat u naar de kermis en neemt plaats op een draaimolen met een verrekijker in de ene hand en een stopwatch in de andere. Als de draaimolen gaat draaien richt u de verrekijker naar buiten en ziet nu alles verplaatsen in één richting. Door de beweging van de draaimolen kunt u het gebied eromheen waarnemen in een grote cirkel. De verrekijker zal het gebied beperken dat u kunt zien, maar wat u ziet is sterk uitvergroot en bevat veel details. Elke keer dat u een persoon door de verrekijker ziet noteert u de exacte tijd. Wij veronderstellen dat u goed bent in het herinneren van gezichten en groepeert de tijdsmetingen per persoon³. Nadat u klaar bent met de rit heeft u een lijst met vele tijdsmetingen per persoon. Nu is de vraag: kunnen we de beweging van alle gemeten personen op basis van deze observaties reconstrueren? Het antwoord is 'ja', hoewel het zal afhangen van de veronderstelling dat de meeste mensen in een rechte lijn bewegen.

Modelleren van de waarnemingen

Om het reconstrueren van de bewegingen van de mensen mogelijk te maken, maken we gebruik van *modellen*: beschrijvingen van de werkelijkheid om te *voorspellen* wanneer een persoon waargenomen wordt. Laten we om het eenvoudig te houden veronderstellen dat iedereen rond de draaimolen loopt op zijn eigen snelheid (of stil staat). Volgens dit

³Het feit dat u mensen aan het observeren bent met een verrekijker en een stopwatch vanaf een draaimolen zal waarschijnlijk zorgen dat mensen naar u kijken wat deze taak gemakkelijker zal maken.

model zal elke persoon worden geklokt op regelmatige tijdsintervallen, voor elke omwenteling van de draaimolen, maar dit interval zal lichtelijk verschillen afhankelijk van hoe de persoon zich verplaatst. Door het vergelijken van de tijdsintervallen voor verschillende personen, is het mogelijk om te bepalen hoe ze zich verplaatsen ten opzichte van elkaar. Voor een nauwkeurige voorspelling van de tijdsintervallen moeten we echter nog een aantal andere zaken modelleren. Ten eerste de rotatie van de draaimolen. Bijvoorbeeld: we moeten weten wat de oriëntatie op een overeengekomen tijdstip is, evenals de rotatiesnelheid en misschien ook of die versnelt of vertraagt. We moeten ook de richting van de verrekijker ten opzichte van draaimolen weten. Hoewel u vermoedelijk probeerde de verrekijker stil te houden is het waarschijnlijk dat er wat variatie was in de richting waarheen u keek ten opzichte van de draaimolen, waarmee we ook rekening dienen te houden.

We hebben nu het probleem in de vorm van drie verschillende modellen gespecificeerd: hoe mensen bewegen om de draaimolen ('persoon model parameters'), de rotatie van de draaimolen ('rotatie model parameters'), en hoe uw verrekijker wiebelt na verloop van tijd ('verrekijker model parameters'). We gebruiken het woord *parameters* hier om de verschillende onderdelen van een model aan te geven die een bepaalde waarde kunnen hebben, bijvoorbeeld de rotatie snelheid van de draaimolen in het rotatie model. De vraag is: hoe kunnen deze modellen u helpen om op basis van uw lijst met observaties de beweging van alle gemeten personen te reconstrueren?

De iteratieve draaimolen oplossing

Laten we aannemen dat we beginnen met model parameters die resulteren in voorspelde observatie tijden die niet te ver af liggen van uw waargenomen tijden. De uitdaging is nu om deze parameters zodanig te verbeteren dat de verschillen tussen de voorspelde en waargenomen tijden zo klein mogelijk worden. Een methode om dit probleem op te lossen is om de parameters voor elk model afzonderlijk te verbeteren en dit proces vele malen te herhalen (itereren). Dat wil zeggen, we nemen eerst aan dat we precies weten hoe snel de draaimolen roteert en hoe de verrekijker wiebelde in uw hand. Met behulp van de tijdswaarnemingen is het nu mogelijk om de beweging van elk persoon afzonderlijk te reconstrueren. Vervolgens, ervan uitgaande dat de gereconstrueerde beweging van alle personen correct is, en aannemende dat we weten hoe de verrekijker wiebelde in uw hand, kunnen we bepalen wat de rotatie van de draaimolen is. Tot slot, ervan uitgaande dat we de bewegingen van mensen en de rotatie van de draaimolen weten, kunnen we het wiebelen van de verrekijker reconstrueren. Dit hele proces moet worden herhaald, omdat de veronderstellingen die elke stap gemaakt zijn uiteraard niet helemaal correct zijn, maar ze worden steeds beter met elke iteratie.

Door deze cyclus vele malen te herhalen (persoon–rotatie–verrekijker parameter verbeteringen) zullen we een punt bereiken waarbij het aanpassen van de parameters de verschillen tussen de waargenomen en voorspelde tijden niet verder vermindert. Als onze

modellen een goede beschrijving zijn van wat er zich in de echte wereld heeft afgespeeld zullen deze tijdsverschillen niet groter zijn dan de verwachte onnauwkeurigheid van uw waarnemingen. De waarde waarmee u elke parameter kunt aanpassen zonder dat dit een significante verandering geeft in de tijdsverschillen is een maat voor de *onzekerheid* in deze parameter (hoeveel u kunt verwachten dat de parameters afwijkt van de werkelijke waarde). Dit betekent dat de onzekerheid van elke model parameter afhangt van de onnauwkeurigheid van uw waarnemingen. Tenzij u de tijdsmetingen met oneindige precisie kunt maken (en de modellen perfect zijn) kunt niet u *exact* bepalen hoe elke persoon zich voortbewogen heeft.

Wat dit zegt over Gaia

Het deel dat u waarschijnlijk wel geraden heeft is dat de draaimolen staat voor de Gaia satelliet die constant ronddraait in de ruimte. In werkelijkheid verandert ook de as waarom zij draait zodanig dat ze soms zelfs ondersteboven draait (om in kermis termen te spreken). Gaia kun je in die zin misschien beter vergelijken met een van die misselijkmakende attracties die alle kanten op roteert, al gebeurt dat dan wel in ‘super slow motion’ aangezien Gaia er zes uur over doet om een keertje rond te draaien. Haar beweging wordt dan ook gemodelleerd in drie dimensies in plaats van de éne die we hier gebruikten.

De verrekijker vertegenwoordigt de twee telescopen in Gaia die continu een gebied aan de hemel iets groter dan de volle maan kunnen waarnemen, en het verrekijker model is in werkelijkheid de beschrijving van kleine bewegingen en rotaties van de 100 digitale camera’s waarmee de exacte passage-tijd van sterren gemeten wordt.

De bewegende personen zijn de sterren, waarvan ongeveer 80% zich in praktijk in een rechte lijn voorbeweegt gedurende de Gaia missie duur van vijf jaar, dus dat model is zeer vergelijkbaar met wat we voor de mensen hebben gebruikt. Veel van de sterren hebben een dicht bijzijnde partner waardoor ze om elkaar heen draaien en een extra wiebelende beweging aan de hemel maken. Dit is te vergelijken met verliefde jonge stelletjes op de kermis die om elkaars nek heen cirkelen. Nabije sterren met meerdere (zwarte) planeten hebben ook ingewikkelde bewegingen die mogelijk het best vergeleken kunnen worden met die van breakdancers (u moet hen geruime tijd observeren voordat u hun bewegingen kunt volgen). In de iteratieve draaimolen oplossing zullen de geschatte rotatie en verrekijker model parameters verbeteren als we de observaties weglaten van personen die zich niet in een rechte lijn voort hebben bewogen. Op dezelfde manier worden alleen de sterren met eenvoudige rechte bewegingen gebruikt om het drie dimensionale oriëntatie model en het camera model van Gaia te bepalen.

Stralingsschade en alcohol

Tot nu toe heb ik beschreven hoe Gaia onder ideale omstandigheden zal functioneren, maar ik heb nog niet het lastigste aspect naar voren gebracht: stralingsschade. Deeltjes van de zon (snelle protonen, uitgestoten tijdens zonnevlammen) beïnvloeden de waarnemingen van Gaia zoals alcohol uw zicht beïnvloedt: het wordt er niet beter op. Wanneer u een biertje drinkt tijdens het uitvoeren van de waarnemingen op de draaimolen neemt de kans toe dat u fouten maakt, maar vooral: uw reactietijd zal langer worden. Het is ook waarschijnlijk dat uw reactietijd zal afhangen van hoe (fel) verlicht de persoon is. U kunt zich voorstellen dat het een stuk moeilijker is om met deze set van tijdsmetingen een nauwkeurig model van de rotatie van de draaimolen, richting van de verrekijker en het pad van elke persoon te maken.

Het bier in dit voorbeeld staat voor stralingsschade aan de camera's van Gaia waardoor ook haar tijdsmetingen verschillend worden vertraagd voor heldere dan wel zwakkere sterren. Schade door straling in Gaia zal op dezelfde manier worden behandeld als de andere effecten die we besproken hebben: door het maken van een model dat kan voorspellen hoe elke tijdswaarneming wordt beïnvloed, en de parameters van dit model moeten net zoals de andere besproken model parameters worden aangepast opdat het verschil tussen de voorspelde en waargenomen tijden minimaal is. Het grootste probleem hier is dat het model zeer ingewikkeld is en dat we niet volledig begrijpen wat de stralingsschade voor invloed op de camera's heeft. En, in tegenstelling tot de effecten van bier, is er geen manier voor Gaia om weer nuchter te worden: de stralingsschade zal tijdens de hele missie alleen maar blijven toenemen.

Waar ik me mee bezig heb gehouden

In tegenstelling tot wat u zou denken bestond mijn PhD studie niet uit het drinken van bier op een draaimolen (nou ja, in ieder geval niet de gehele tijd). Waar ik me mee bezig hield was er proberen achter te komen hoe al deze verschillende modellen en parameters van invloed zijn op de nauwkeurigheid waarmee we kunnen schatten wat de posities, afstanden en snelheden van de één miljard sterren zijn die Gaia zal waarnemen. Om dit te doen gebruiken we de Astrometrische Globale Iteratieve Oplossing (in het engels afgekort als AGIS) om de model parameters te schatten, welke werkt op een soortgelijke manier als de iteratieve draaimolen oplossing. Veel van mijn werk betreft het programmeren van modellen en het doen van experimenten met AGIS. Omdat Gaia pas zal worden gelanceerd in 2013 moeten we waarnemingen simuleren waaruit we vervolgens de parameters voor elke ster met behulp van AGIS schatten. Bijvoorbeeld, om te weten hoe nauwkeurig de afstand tot sterren kan worden bepaald kunnen we vele verschillende simulaties maken (die vergelijkbaar zijn met vele Gaia missies) en voor elke simulatie bepalen hoe groot de fout is tussen de geschatte en de ware afstand tot elke ster. De spreiding

van de fouten in de vele simulaties is dan een maat voor hoe nauwkeurig de afstand kan worden gemeten gedurende de enige echte Gaia missie. Ook heb ik gewerkt aan een model dat de nauwkeurigheden direct (zonder het gebruik van simulaties) kan berekenen. Wanneer de Gaia catalogus rond 2020 uit komt zal dit model een belangrijke bron van informatie zijn voor astronomen die de structuren in onze Melkweg willen bestuderen, bijvoorbeeld door afstanden en bewegingen van groepen van sterren te schatten.

In samenwerking met Thibaut Prod'homme, werkende aan de Sterrewacht in Leiden, heb ik ook gedetailleerde studies gemaakt van de effecten van stralingsschade op Gaia, en hebben we onderzocht hoe dit de uiteindelijke resultaten van Gaia beïnvloed. Ik had al dit werk zeker niet alleen kunnen doen. Naast het Gaia team in Lund heb ik samengewerkt met veel mensen in Europa, met name van Sterrewacht Leiden, het Instituut van de Astronomie in Cambridge, en de ESA-centra in Madrid en Noordwijk. Van 2007 tot 2010 maakte ik deel uit van een Marie Curie onderzoekers trainingsnetwerk 'Europees Leaderschap in Ruimte Astrometry' (in het engels afgekort als ELSA). Via dit netwerk heb ik veel vrienden gemaakt waarmee ik dikwijls een biertje heb gedeeld.

De volgende keer dat u zich realiseert dat u licht aangeschoten op een draaimolen heeft plaats genomen, kijk dan omhoog naar de sterren, en besef dan dat wat u doet niet héél veel verschilt van het in kaart brengen van een miljard sterren met Gaia...

Met dank aan Jan Holl en Brenda Vos voor de vertaling.

Acknowledgements

*'A happy person has mastered the art of satisfaction.
A good scientist has mastered the art of perfectionism.
A happy scientist has managed to balance the two opposites.'*
— Rutger van Haasteren

When I was given the opportunity to start working in Lund on the calibration of five billion astrometric parameters I had no idea what was about to happen, I only knew that it *felt really right*.

Looking back I can without doubt say that these four and a half years have been among the best years of my life! Not only have I been able to work on an amazingly interesting project that is literally out of this world, but I also discovered how much I value *sharing* with people in collaborations, in defining problems, in finding solutions, in travelling, in friendships, and in just being among others. I therefore dedicate this theses to the many wonderful people I have been fortunate to meet, and hereby would like to address a few words to some of you directly.

There is one person in particular that has been crucial, and very dear to me in these respects: Lennart Lindegren, you are the greatest tutor I ever had. I deem it unlikely that I will ever work again with someone who will have such an everlasting effect on how I work, think, and collaborate with people. I am deeply grateful for the many opportunities you have given me over these years, and how I have been able to experience and learn from such an honest, intelligent, modest and friendly person as yourself, it has been an honour, but above all, a great pleasure to have been working with you. Writing my thesis has been among the most rewarding experiences in my life and it is of great value to me how much time and energy you put into finishing it together with me. I sincerely hope that we will be able to collaborate also in the future!

David Hobbs, I am very happy about having worked with you all these years. I always enjoyed our discussions and I am very proud of how we programmed the majority of AGISLab together. I have seen you working on fixing software 'bugs' that sometimes took

days to track down, but you always managed to find the problem by carefully checking everything that could be wrong. I have great respect for the dedication you have to your work and the persistence by which you do it. I hope that I will be able to keep working on AGISLab in my future job so that I do not even have to find an excuse to visit you and Lennart in Lund!

Then there is ELSA (European Leadership in Space Astrometry), the Marie-Curie Research Training Network which I gratefully acknowledge for the many opportunities it has given me to Experience, Learn, and Share through Astronomy. It resulted in many friendships and I hereby want to thank you all for that! In particular I'd like to mention: Thibaut, you have been my partner in crime on many great scientific and non-scientific adventures these years and I wish that we will keep on going! Ester, thanks for sharing the mountains of the Dolomiti in search for a destination, never stop smiling from your heart. Mihály, you always amaze me in some way: with your stories, with your presence, keep on amazing! Umami, you are really one of the kindest persons I know and I always really enjoyed the times we met. Marysia, you are the most free spirit I ever met, you inspired me to embrace life! Aidan, although you started later than the rest, when you arrived it was as if you have always been there, I am happy to know you. Alex, it was great to have worked with you, you are among the brightest people I know. Thanks also for the Gaudí tour in Barcelona! Daniel, it has been great travelling and working with you, I think the way you explain the motions of Gaia with your hips should be part of the outreach material. Michael, I always enjoyed your dry humor, and your passion when speaking about your collection of rare meteorites. Scott, it has always been a pleasure to work and hang out with you; by the way, when is the next Cambridge Open Golf tournament? Carlos, I still have great memories of the night you gave me a private guitar concert on your balcony in Torino, I hope to see you someday soon again. Paola, we should definitely dance again when we meet next time! Dagmara, I think you are a great navigator, by car and in multi-dimensional parameter space, and it's great fun to discover new places with you. Tenay, I really love your energy and enthusiasm: Si Si Si Si! Maya, quiet and thoughtful you appear to me, very worth to get to know as was so greatly worded in your beautiful *Secrets*. Luca, we haven't spoken much but I always enjoyed your company.

I also want to thank the *many* great scientists I have been able to work or interact with over the last years through ELSA or in working groups in Gaia. A lot of thanks to Anthony, it was great to work with you on the papers and discuss research when I was visiting Leiden. Also a big thanks to Floor, I still remember the first time I walked into your office and you gave me a private demonstration of the Hipparcos re-reduction, that was great! Thanks also for refereeing our paper, and for supervising us during the visits in Cambridge that lead to the paper; the *CTI dream team* would not have existed without you. In particular I would also like to mention: Uwe, Ulrich, Jos, Sergei, Stefan, and Catherine, you have been a great inspiration to me!

To all friends and colleagues at Lund Observatory: many thanks, I have always felt really

at home and I will miss you! A special thanks to Anna and Chiara, it was so much fun to share an office with you! Thanks also the Stiftelsen Walter Gyllenbergs fond that gave me a 120 000 kr grant to buy some super-duper computers which have been essential for the computations in most of my papers. I also want to express my gratitude to Lunds Aikidoklubb and the wonderful people there. I discovered you two years ago and since then Aikido has become an essential part of my life. I plan to visit you whenever I am around in Lund!

During my time here in Sweden the contact with my friends in The Netherlands has only improved for which I am very grateful. Whenever I visited Leiden I never had a problem finding a place to sleep (or finding people that were in for a beer night): thank you all!

I would like to conclude by thanking my parents Jan en Wil, and my sister Anneke. You have always been there for me, supporting me in any choice I made, unconditionally. You have also seen how I have been struggling to find my own place in this world and accept myself. To me this thesis means much more than the words, figures and equations it contains, to me it is also a celebration of having found peace within myself, and having found my place among the billions of stars with the people I love on this amazing living planet called Gaia...

References

- Bombrun, A., Lindegren, L., Hobbs, D., et al. 2011, *A&A*, in press (DOI: 10.1051/0004-6361/201117904)
- ESA. 1997, *The Hipparcos and Tycho Catalogues*, ESA SP-1200
- Hobbs, D., Holl, B., Lindegren, L., et al. 2010, in *IAU Symposium*, ed. S. A. Klioner, P. K. Seidelmann, & M. H. Soffel, Vol. 261, 315–319
- Høg, E. 2008, in *IAU Symposium No. 248*, ed. W. J. Jin, I. Platais, & M. A. C. Perryman, 300
- Høg, E. 2009, *Experimental Astronomy*, 25, 225
- Jordi, C. 2011, *ArXiv e-prints*, arXiv:1105.6166v1
- Jordi, C., Gebran, M., Carrasco, J. M., et al. 2010, *A&A*, 523, A48
- Klioner, S. A. 2003, *AJ*, 125, 1580
- Lindegren, L. 2008, A general Maximum-Likelihood algorithm for model fitting to a CCD sample data, Tech. note, Tech. rep., Lund Observatory, GAIA-C3-TN-LU-LL-078-01
- Lindegren, L. 2010, in *IAU Symposium*, ed. S. A. Klioner, P. K. Seidelmann, & M. H. Soffel, Vol. 261, 296–305
- Lindegren, L., Babusiaux, C., Bailer-Jones, C., et al. 2008, in *IAU Symposium*, ed. W. J. Jin, I. Platais, & M. A. C. Perryman, Vol. 248, 217–223
- Lindegren, L. & Bastian, U. 2011, in *EAS Publications Series*, Vol. 45, 109–114
- Lindegren, L., Lammers, U., Hobbs, D., et al. 2012, *A&A*, in press (DOI: 10.1051/0004-6361/201117905)
- Lindegren, L. & Perryman, M. A. C. 1996, *A&A*, 116, 579
- Perryman, M. A. C. 2004, in symposium *Exploring the Cosmic Frontiers*, from presentation *Detection and Characterization of Extra-Solar Planets*
- Perryman, M. A. C., de Boer, K. S., Gilmore, G., et al. 2001, *A&A*, 369, 339
- Prod'homme, T. 2011, in *EAS Publications Series*, ed. C. Turon, F. Meynadier and F. Arenou, Vol. 45, 55–60
- Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, *A&A*, 409, 523

Turon, C., O’Flaherty, K. S., & Perryman, M. A. C., eds. 2005, ESA Special Publication, Vol. 576, The Three-Dimensional Universe with Gaia

Varadi, M., Eyer, L., Jordan, S., Mowlavi, N., & Koester, D. 2009, in AIP Conference Series, ed. J. A. Guzik & P. A. Bradley, Vol. 1170, 330–332