


# GOCE Spacecraft Attitude Control

A magnetic attitude controller simulator

R. Maltha

Technische Universiteit Delft



**Front cover image taken from:**

[http://blogs.esa.int/rocketscience/files/2013/10/G0CE\\_in\\_orbit.jpg](http://blogs.esa.int/rocketscience/files/2013/10/G0CE_in_orbit.jpg)



# **GOCE Spacecraft Attitude Control**

## **A magnetic attitude controller simulator**

by

R. (Rutger) Maltha

to obtain the degree of Master of Science  
at the Delft University of Technology,  
to be defended publicly on Friday 24 August 2018, at 13:00.

Student number:	4523768
Project duration:	December 1, 2017 – August 24, 2018
Thesis committee:	Prof. dr. ir. P. N. A. M. Visser, TU Delft, chairman
	Dr. ir. E. N. Doornbos, TU Delft, supervisor
	Dr. ir. E. van Kampen, TU Delft, external member
	Ir. T. Visser, TU Delft, supervisor

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.





# PREFACE

With the master thesis finalised in this document, the end of my master studies is drawing near. It has not been easy, but looking forward to being able to call myself an aerospace engineer makes it worth it. Although I have always been amazed by space flight and satellite orbits in general – I don't see how someone can not be – my path to aerospace engineering has not been straightforward. I started my life as a student with a bachelor studies in civil engineering. Being initially attracted by the combination of engineering and management, I found that there was too much emphasis on the management part for me to fully enjoy it. Somewhere halfway through my bachelor studies I started thinking about going somewhere else for my master studies and had initially set my sight on offshore engineering. Around the same time I also found a computer game called 'Kerbal Space Program' in which I found a hobby that could satisfy my desire for space flight.

I still had not given any thought about doing a masters in aerospace engineering, but I did find myself spending more and more of my time thinking about satellite orbits, spacecraft rendezvous, and landing on the 'Mün'. At some point it started to get out of hand when I was making excel sheets in my spare time to help me optimise my spacecraft. It was around this time that my girlfriend, Vicky, suggested that I should consider doing my masters in aerospace engineering. Probably because she was getting sick of her being the only one that had to suffer through my "hobby". Initially I brushed it off, thinking I would have to do a pre-master program of more than a year. However, at some point she got me starting to consider it more seriously, which led me to looking up how much time the pre-master program would cost me. To my surprise there was no need for any pre-master program, which pushed me over the edge to give it a try.

Given my background in civil engineering, the switch to aerospace engineering was a particular difficult one. It initially probably cost me more effort than the average student to understand the lectures, complete assignments, and study for exams. However, I enjoyed studying for almost every single course, because of my interest in space flight. This interest has definitely helped me get through the hard work I had to put into studying, and I have never regretted making the switch to aerospace engineering. During the courses I started noticing that I was picking up an interest in control theory that was not satisfied by the curriculum. When Eelco, my main supervisor, brought up the subject of this master thesis I was thus immediately interested. Although control theory was again a new subject for me, by now I was confident that I would be able to learn to understand it, having gone through all the aerospace engineering courses successfully.

Similar to not regretting my switch to aerospace engineering, I also do not have a single regret for choosing a master thesis subject that was outside of my knowledge area. I enjoyed learning everything that I have learned about control theory and I have had great support from Eelco as well as Tim, my daily supervisor. They have both guided me through the thesis in a sympathetic, yet critical, manner, and have provided me with key insights when I needed them. I have been, and still am, genuinely appreciative of the way I was supervised, and for that I would like to extend my gratitude to Eelco and Tim. Next I would like to thank Vicky, not in the least for her perseverance in bringing up the idea to do aerospace engineering, but also for her support during the last few stressful months of the thesis. Finally I would like to thank my family and friends, for keeping me cheerful and sane during the thesis.

*R. Maltha  
Delft, August 2018*



# EXECUTIVE SUMMARY

The Gravity field and steady-state Ocean Circulation Explorer (GOCE) mission by ESA was a mission to measure the Earth's gravity field with unprecedented accuracy. To accomplish this mission, the spacecraft carried with it a scientific instrument for gravity determination that measured the accelerations acting on the spacecraft with extraordinary precision. For these measurements it flew at a very low Earth orbit, where aerodynamic forces are significant. In order to relate the acceleration measurements to gravitational accelerations, all non-gravitational accelerations acting on the spacecraft needed to be minimised. In order to achieve this the spacecraft used an ion engine to compensate linear accelerations caused by drag, and it controlled its attitude through the application of magnetorquers.

Thanks to the very accurate, common mode, acceleration measurements a new set of spacecraft torque models has been set-up. These torque models have provided an excellent opportunity to evaluate the possibility of applying a relatively new controller type for spacecraft attitude control. This controller type is called Incremental Nonlinear Dynamic Inversion (INDI) and promises to be easier to tune than conventional controllers, while simultaneously being more robust. During this thesis work it has been investigated whether INDI-control can be applied to a magnetically stabilized spacecraft, and keeping its attitude, angular rate, and angular acceleration within the same limits set by ESA for the GOCE mission. The shape model and properties of the GOCE spacecraft have been used in the simulation environment as the model for the simulated spacecraft.

This research has specifically focused on applying an INDI-controller with an NDI outer loop, hereafter called a NINDI-controller. This controller has been set-up in MATLAB™, as well as a more conventional P(I)D-controller that is used as a reference controller. The controllers have been tuned using a genetic algorithm, which has proven to be very effective. Not only has the genetic algorithm produced a set of gains for each controller that produced results within ESA limits for the GOCE mission. It also produced multiple set of gains for each controller that have been used to investigate how sensitive the controllers are to the values of the gains. After having tuned the controller, simulations were carried out for one validation day, and three other simulation periods during the GOCE mission. In these other three simulation periods, the spacecraft flew in different operating conditions. One was at the start of the mission where its average altitude was still relatively high, and there were little thermospheric winds to disturb the spacecraft. At the other end of the spectrum a period closer to the end of the mission was chosen, where the spacecraft flew at a very low average altitude of around 235 kilometres, and there were significant thermospheric winds.

The results in this research have demonstrated that it is in fact possible to use a NINDI-controller for magnetic control of a GOCE-type spacecraft, and meeting the ESA requirements in the time domain. However, it was also found that the requirements in the frequency domain were not met, although this has been identified to be caused by too aggressive inner loop gains. When lowering the gains, the frequency requirements could also be met. Hence it is expected that when the performance in the frequency domain is incorporated in the fitness function of the genetic algorithm, requirements in both domains can be met simultaneously. Furthermore, the NINDI-controller has proven to be more stable than a reference PD-controller, and even to keep the spacecraft state more stable in general than the actual GOCE spacecraft was during equal orbits. Finally, the NINDI-controller was still able to keep the spacecraft within the time domain requirements, when errors were introduced that are in line with what could be expected during the GOCE mission. In this thesis work, it has been shown that a GOCE-type spacecraft, with its accurate acceleration measurements and fine control actuation through the application of magnetorquers, is especially suited for use of a NINDI-controller.





# CONTENTS

<b>PREFACE</b>	<b>iii</b>
<b>EXECUTIVE SUMMARY</b>	<b>v</b>
<b>LIST OF ABBREVIATIONS AND SYMBOLS</b>	<b>ix</b>
<b>1 INTRODUCTION</b>	<b>1</b>
1.1 Problem statement . . . . .	2
1.2 Thesis goal and research questions . . . . .	3
1.3 Document outlook . . . . .	4
<b>2 THEORY</b>	<b>5</b>
2.1 Attitude dynamics . . . . .	5
2.1.1 Attitude representation . . . . .	5
2.1.2 Equations of motion . . . . .	7
2.1.3 Perturbations . . . . .	7
2.2 Magnetic attitude control . . . . .	9
2.3 Control theory . . . . .	12
2.3.1 Linear control theory . . . . .	13
2.3.2 Nonlinear dynamic inversion . . . . .	14
2.3.3 Incremental nonlinear dynamic inversion . . . . .	15
<b>3 METHODOLOGY</b>	<b>17</b>
3.1 Research set-up . . . . .	17
3.2 Model architecture . . . . .	18
3.2.1 Simulation . . . . .	19
3.2.2 Torque calculation . . . . .	20
3.2.3 Reference frames. . . . .	21
3.3 Controller design . . . . .	22
3.3.1 PD-controller . . . . .	22
3.3.2 LINDI-controller. . . . .	23
3.3.3 NINDI-controller . . . . .	24
3.4 Gain tuning . . . . .	25
<b>4 RESULTS AND DISCUSSION</b>	<b>29</b>
4.1 Model output verification . . . . .	29
4.1.1 Reference frames. . . . .	29
4.1.2 Torques acting on spacecraft. . . . .	34
4.2 Results . . . . .	38
4.2.1 Performance of the NINDI controller . . . . .	38
4.2.2 Comparison of the NINDI-controller with GOCE's observed behaviour . . . . .	42
4.2.3 Controller robustness to uncertainties and measurement errors. . . . .	48
4.2.4 Controller tuning difficulty and sensitivity to changes in gains. . . . .	53
4.3 Discussion . . . . .	57
<b>5 CONCLUSION AND RECOMMENDATIONS</b>	<b>59</b>
<b>REFERENCES</b>	<b>63</b>
<b>A STATE PLOTS</b>	<b>67</b>
<b>B DIFFERENCES BETWEEN MEASURED AND SIMULATED</b>	<b>71</b>
B.1 Attitude, angular rate, and angular acceleration differences. . . . .	71
B.2 Control torque differences . . . . .	73

---

<b>C</b>	<b>ALL DISTURBANCE TORQUES FOR 16 - 21 MARCH 2013</b>	<b>75</b>
<b>D</b>	<b>TIME SERIES STATE FOR 31 MAY TO 2 JUNE 2013</b>	<b>77</b>



# LIST OF ABBREVIATIONS AND SYMBOLS

## Abbreviations

BOL	Beginning Of Life
DFACS	Drag-Free Attitude and Orbit Control System
DFM	Drag Free Mode
EGG	Electrostatic Gravity Gradiometer
EMF	Earth Magnetic Field
EOL	End Of Life
EOM	Equation Of Motion
ESA	European Space Agency
GA	Genetic Algorithm
GOCE	Gravity field and steady-state Ocean Circulation Explorer
IGRF	International Geomagnetic Reference Field
INDI	Incremental Nonlinear Dynamic Inversion
IPA	Ion Propulsion Assembly
LEO	Low Earth Orbit
LINDI	Linear outer-loop INDI controller
LVLH	Local Vertical, Local Horizontal
MBW	Measurement Bandwidth
NDI	Nonlinear Dynamic Inversion
NINDI	NDI outer-loop INDI controller
PD	Proportional, Derivative (-controller)
PID	Proportional, Integral, Derivative (-controller)
PSD	Power Spectral Density
SAA	South Atlantic Anomaly
VLEO	Very Low Earth Orbit

## Greek symbols

$\beta_{AE}$	Aerodynamic side-slip angle
$\beta_{SR}$	Solar radiation pressure side-slip angle
$\dot{\omega}$	Angular acceleration
$v$	Virtual control input
$\omega$	Angular rate
$\theta$	Euler attitude angle

## Roman symbols

$J$	Inertia matrix
$T_{AE}$	Aerodynamic torque
$T_{FD}$	Fitted dipole torque
$T_{GG}$	Gravity gradient torque
$T_{RD}$	Residual dipole torque
$T_{SR}$	Solar radiation pressure torque
$T_{TM}$	Torque due to the IPA
$T_B$	Net body torque
$T_C$	Control actuation torque
$T_D$	Disturbance torque
$J_2$	The spherical harmonic term of degree 2
$K_D$	Derivative gain
$K_P$	Proportional gain
LB	Lower Bound of GA search space
u	argument of latitude
UB	Upper Bound of GA search space

# INTRODUCTION

Even before the first artificial satellite was launched into space, attitude control concepts that are still relevant for missions today had been investigated. In the 1950s space flight was starting to take off and thus the field of attitude control grew as well. At that time computing power was limited and so were the possibilities to do extensive simulations of spacecraft dynamics. However, the idea of rigid body motion and the kinematic and dynamic equations that follow from it were well understood. Using the mathematical expressions for spacecraft dynamics, engineers were able to control the attitude of even early satellites. The first spacecraft relied on either gyroscopic stabilisation or field stabilisation. That is, they were kept in a desired attitude by spinning them, or interact with Earth's gravity field or magnetic field, or a combination of both [27].

The early principles of attitude control, including gyroscopic, field, and servo stabilisation, are still fundamental for modern spacecraft [21, 27]. A common example is where one axis is stabilised through the use of gravity gradient, aligning the minimum moment of inertia axis with the gravitational field, and control the other two axes in another way [18, 47]. However, advancements in sensors, computers, and actuator hardware allow for three-axis stabilisation of any spacecraft [21]. Though the choice for three-axis stabilised spacecraft depends on the requirements, available budget, and desired simplicity and reliability. Hence it is still common for spacecraft today to have one axis stabilised through gyroscopic or field stabilisation, with the other two axes controlled actively.

Controlling the attitude of a spacecraft solely through interaction with the Earth Magnetic Field (EMF) is relatively new, especially for larger satellites. It is true that magnetorquers have been used in space for a long time, but they have often only been used for desaturating reaction wheels, instead of actually controlling the attitude [25, 47]. The reason why the use of magnetorquers is not as widespread is because using them poses a few challenges. First, to make use of the magnetorquers for control actuation, an external magnetic field needs to be present. While the Earth does possess a significant magnetic field, the intensity decreases by the square of the radius, rendering magnetorquers insufficient above a certain altitude [34, 47]. Secondly, since a torque that is parallel to the external magnetic field lines cannot be produced, any satellite using magnetorquers as a control actuator is instantaneously under-actuated [6, 19]. Still, for a satellite in an orbit that is both of low enough altitude and of sufficiently high inclination, magnetorquers as the sole control actuators are an interesting option. Magnetorquers are small, reliable because of the absence of moving parts, and do not consume any propellant [39]. Furthermore the components of a magnetorquer are inexpensive, and the magnetorquers themselves have low activation noise. This makes them ideal for missions with stringent pointing and acceleration requirements [31].

One such mission was the European Space Agency's (ESA) 'Gravity field and steady-state Ocean Circulation Explorer', more commonly known as GOCE, mission. The GOCE mission goal was to "measure the Earth's gravity field and to define a model of the geoid with an unprecedented accuracy" [11]. In practice this meant that the geoid height was determined with centimetre level accuracy and with a spatial resolution of about 100 kilometres on the Earth's surface. Interestingly, while GOCE focussed on maximum spatial resolution, the GRACE mission focused on the redistribution of masses in the Earth system, i.e. maximum temporal resolution. In other words, these two gravity field data sets are complementary to each other [10]. A few practical use cases of the produced data are specified in Ghisi et al. [11]: precise determination of ocean circulation, Earth lithosphere studies, and global unification of height systems, which is particular useful to establish a






Figure 1.1: Layout of the accelerometers in the electrostatic gravity gradiometer. [3]

global sea-level monitoring system. Early in the mission design phase it was decided that in order to meet the stringent linear and angular acceleration requirements, a set of magnetorquers were needed for controlling the attitude of the satellite [32].

The GOCE mission was launched in March 2009, and after a commissioning phase was planned for an operational time of 20 + 10 (nominal mission + estimated propellant reserve) months. However, experienced drag by the satellite turned out lower than expected, which meant the mission could be operated longer. This eventually led to the decision to extend the mission into 2013 and lower the orbit to improve the scientific data. The lowest average altitude GOCE flew at during this 'low orbit operations campaign' was about 230 kilometres. The re-entry campaign started on 23 October 2013, when propellant for the ion propulsion assembly was depleted [9, 36]. GOCE was deemed an extremely successful mission which gathered very accurate gravitational acceleration measurements with the Electrostatic Gravity Gradiometer (EGG) [11]. This device, shown in Figure 1.1, contains three pairs of three-axis accelerometers connected together on a thermally stable carbon honeycomb structure. Each single accelerometer contains a proof mass that is electrostatically suspended and kept at the centre of each cage. Keeping the proof mass at the centre is achieved by applying voltages. The applied voltages are proportional to the mass accelerations as the proof masses are displaced with respect to the cages because of accelerations acting on the spacecraft [3, 36].

The EGG was not only used as a scientific instrument for gravity determination, but also to measure accelerations that were used in the attitude control system [3, 32]. The differential signal between two proof masses is the scientific product, while the common mode is the average acceleration [32]. The common mode was used in the attitude control system. This resulted in a data set with very accurate acceleration measurements that lent itself well for modelling the torques acting on the spacecraft. Modelling of the torques acting on the GOCE spacecraft is discussed in a paper by Visser et al. [41]. The torque models that are discussed in the paper promise to be of higher fidelity than anything before. This means that the torque models can be used to create a simulation environment where the simulated torques closely match what would happen in real life.

### 1.1. Problem statement

It is generally accepted that the number of satellites that are going to be launched will increase rapidly. Some even discuss the potential for the total amount of spacecraft in orbit around Earth to double within a few years, and potentially doubling again thereafter. While this expected increase can for some part be attributed to the rise of cubesats [40], making satellites cheaper and therefore interesting for mass production and mass launch. More likely however is that this expected increase can be attributed to planned massive constellations aiming to provide global satellite internet coverage [37].

Unless these spacecraft are going to be passive spherical reflectors, their attitude will need to be controlled. Traditionally when tuning a controller, the controller has to be tuned for specific operating conditions. When these conditions change, the controller has to be tuned again [26]. In reality however the operating regimes and their corresponding gains are defined beforehand. This means that when a system switches into another operating regime, the controller switches to a new set of gains automatically. This process is called gain scheduling and is common practice. However, apart from gain scheduling being a tedious process, it can also cause instability when the controller switches a set of gains when the system moves from one pre-defined operating regime to the next one [26]. The difficulty in tuning a controller correctly also showed

up when the Drag-Free Attitude and Orbit Control System (DFACS) for the GOCE satellite was being tuned [11]. While the difficulties with tuning the GOCE DFACS were eventually overcome it highlights the problems associated with satellite attitude control. Combine this with the expected increase in the total number of satellites in space, and it can be concluded that there is a desire for a more flexible method of controlling the attitude of spacecraft.

One such method is called nonlinear dynamic inversion (NDI). This method linearises the system which allows to find a direct relation between the output of a system and the control input [35]. NDI was a first step into more advanced control topics, but there are some disadvantages associated to it. Requiring accurate model knowledge is one such disadvantage. To reduce the model dependency of the controller Incremental NDI, or INDI was developed [1, 33]. Combining the strengths of NDI, with the reduced dependency on accurate model knowledge is a powerful combination for spacecraft attitude control, but to this author's knowledge no such controller has yet flown in space.

The information presented in the introduction combined with the stated problems above are all pieces of a puzzle that leads to an interesting challenge. Magnetically actuated spacecraft attitude control is interesting for Low Earth Orbiting (LEO) satellites because of the low cost, reliable components that do not depend on any propellants. Controlling a spacecraft's attitude magnetically does bring forth some challenges, it is interesting nonetheless because the challenges are solvable. However, a new type of attitude control that has likely never flown in space promises to be more stable, and stable across a wider range of environments, than conventional controllers. It is interesting to test this new type of attitude control together with magnetic attitude control, to see if the combination of the two produce an alternative attitude control system that is both stable and reliable. To test it however, a simulation environment is needed in which torque models can be used that are as close to reality as possible. Coincidentally a set of torque models has recently been made available and promises to be very true to reality. So the GOCE mission, and more specifically the data it produced that led to the creation of the torque models, provides an excellent opportunity to accurately model the performance of a different attitude control system that also uses magnetic attitude control. Furthermore, given the LEO environment GOCE flew in, it is interesting to see whether the attitude control can be done in a control frame that tracks the direction of the linear accelerations acting on the spacecraft.

## 1.2. Thesis goal and research questions

The goal of this thesis is then to make use of the opportunity described in section 1.1, and use it to test a new attitude control architecture. The torque models that have been set-up together with the available data from the GOCE mission present an excellent opportunity to validate the feasibility of a new type of attitude control architecture, and it can be compared against a mission that actually flew and was deemed highly successful. Additionally given that the DFACS aimed to minimise linear accelerations acting on the spacecraft, it is interesting to investigate whether defining a control frame that tracks the direction of the linear acceleration will improve controller performance. Concretely, the goal of the thesis is then phrased as follows:

*The research objective is to contribute to the development of new attitude control systems for satellites, by designing an INDI-based control architecture for GOCE, evaluating its performance, showing the feasibility of such a controller for magnetic attitude control, and investigate whether controller performance improves by using a linear acceleration tracking control frame.*

The goal of the thesis was formulated based on the information gathered in the introduction and problem statement. In order to achieve the goal, one main research question is set-up that is necessary to answer in order to evaluate whether the goal has been met. In order to answer the main research question a series of sub-questions has also been formulated. These are all shown below.

- What is the overall performance of an INDI-type control architecture for the magnetically controlled attitude of the GOCE satellite?
  1. What is the performance of the controller with respect to the ESA requirements; i.e. the maximum values for attitude, angular rate, and angular acceleration?
  2. How do the simulated disturbance torques, control torque, and net body torque compare with the measured and modelled torques from the GOCE satellite during equal orbits?
  3. What is the change in controller performance when measurement errors, and increases in thermospheric properties like winds and air density are introduced?

4. *How difficult is it to find a stabilising set of gains, and how sensitive is the controller to changes in gains?*
- How does using a control frame that tracks the linear accelerations acting on the satellite affect the performance of the controller?

### **1.3. Document outlook**

This thesis report encompasses the process that was undertaken in order to answer the proposed research questions. The report will start with an overview of the relevant literature and theory in chapter 2. Then the methodology that was used to obtain the results is described in chapter 3. The results themselves are presented and discussed in chapter 4. Finally, the research questions will be reflected upon and answered in chapter 5, in which recommendations for further research are also given.



# 2

## THEORY

This chapter will lay the foundation for the required knowledge of understanding the problem at hand, and interpreting the obtained results. The chapter is split up in a section focusing on the equations of motion, perturbations, and attitude representation, section 2.1. In section 2.2 the theory of controlling a spacecraft using magnetic actuation only is discussed. Lastly, section 2.3 provides an introduction into control theory in general, and especially focuses on (I)NDI control theory.

### 2.1. Attitude dynamics

To describe the rotational dynamics of a spacecraft it is often assumed a valid approximation that the spacecraft is a rigid body [44]. This assumption results in the spacecraft having translational motion of its centre of mass in space, and rotational motion of the spacecraft about its own centre of mass. With the assumption of a rigid body spacecraft, its rotational dynamics can be expressed using Euler's rotational equation of motion. Before describing the equations of motion of the attitude of a spacecraft, first a general definition of the reference frames used for attitude dynamics is given. Reference frames particularly related to the GOCE spacecraft and this thesis research are discussed in subsection 3.2.3. Lastly, as part of the dynamics of a spacecraft the angular perturbations acting on a satellite will be discussed, focussing on the perturbing angular accelerations relevant to the GOCE satellite.

#### 2.1.1. Attitude representation

The expression of a spacecraft its attitude is done in terms of angles. However, to use or interpret these angles it has to be clear what these angles represent. In attitude control for spacecraft, the definition of attitude is commonly given as the angles between a body-fixed reference frame and an orbital, or control, reference frame. The exact definitions of these control frames can vary, but for rigid body spacecraft a body-fixed reference frame can be specified with the origin at the centre of mass of the spacecraft and a set of three orthogonal vectors as the definition of the three body axes. Similarly controlling a spacecraft can be done with respect to the orbital, or Local Vertical Local Horizontal (LVLH), frame. This is a common frame for spacecrafts in a circular orbit, with one axis pointing in the direction of the inertial velocity of the spacecraft, a second axis in the direction of the angular momentum vector, and the third axis completing a right-handed reference frame. These axes are shown in Figure 2.1, with the  $\bar{a}_i$ -axes being the control frame and the  $\bar{b}_i$ -axes being the body-fixed frame [44].

With this general definition of the control and the body-fixed reference frame, the attitude of a spacecraft can be defined as a set of angles between two axes of the two reference frames, or as a rotation about the Euler axis, or eigenaxis. Euler angles are successive rotations about the axes, and the rotation about the Euler axis is shown in Figure 2.2. Defining the attitude of a spacecraft in terms of a rotation about the Euler axis is known as representing attitude in terms quaternions. A set of three Euler angles can be used to describe the attitude of a rigid body craft in three dimensional space, but the definition of the three angles can vary. One commonly used definition of a set of three Euler angles is called roll, pitch, and yaw, but other definitions are also valid. The reason why Euler angles are often used is because they are an easy to grasp concept, and given an attitude representation in Euler angles one can visualise what that would look like, but they also have disadvantages. Euler angles can result in singularities, a phenomenon called gimbal lock, and they are




Figure 2.1: Rigid body in a circular orbit. [44]




Figure 2.2: Euler eigenaxis rotation. [4]

less accurate than quaternions when attitude changes are propagated over time. However, quaternions have their own disadvantages. The quaternion parameters do not have intuitive physical meaning, and therefore it can be hard to quickly visualise the meaning of their value. Quaternion attitude representation can also be difficult to include in an optimization [7].

The GOCE mission imposed strict requirements on the attitude of the spacecraft, therefore if everything goes well one should not expect large attitude angles. From here on out the thesis will thus focus on representing the attitude in terms of roll, pitch, and yaw Euler angles. However the code has been set-up to work with quaternions as well. Since an attitude is a rotation of one reference frame in another, the orientation of  $\bar{b}$  in  $\bar{a}$  can also be expressed as a rotation of  $\bar{a}$ . This can be done with a direction cosine matrix which is also used to make transformation between different reference frames. The direction cosine matrix of  $\bar{b}$  in  $\bar{a}$  is given as follows. Assume two reference frames, A and B, with each reference frame having a set of three right-hand orthogonal unit vectors,  $a_i$  and  $b_i$  respectively. Then the basis vectors of B can be expressed in terms of basis vectors of A using the relation shown in Equation 2.1. Which can also be written as Equation 2.2, in which  $\mathbf{C}^{B/A}$  is the direction cosine matrix expressing B in A [44].

$$\begin{bmatrix} \bar{b}_1 \\ \bar{b}_2 \\ \bar{b}_3 \end{bmatrix} = \begin{bmatrix} C_{11}\bar{a}_1 + C_{12}\bar{a}_2 + C_{13}\bar{a}_3 \\ C_{21}\bar{a}_1 + C_{22}\bar{a}_2 + C_{23}\bar{a}_3 \\ C_{31}\bar{a}_1 + C_{32}\bar{a}_2 + C_{33}\bar{a}_3 \end{bmatrix} \quad (2.1)$$

$$\begin{bmatrix} \bar{b}_1 \\ \bar{b}_2 \\ \bar{b}_3 \end{bmatrix} = \mathbf{C}^{B/A} \begin{bmatrix} \bar{a}_1 \\ \bar{a}_2 \\ \bar{a}_3 \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} \\ C_{21} & C_{22} & C_{23} \\ C_{31} & C_{32} & C_{33} \end{bmatrix} \begin{bmatrix} \bar{a}_1 \\ \bar{a}_2 \\ \bar{a}_3 \end{bmatrix} \quad (2.2)$$

The terms direction cosine matrix, rotation matrix, and coordinate transformation matrix are used interchangeably in literature [7, 44]. Since the direction cosine matrix was defined as the required transformation to express B in A, or the other way around, and both these reference systems consist of three sets of orthogonal unit vectors, the direction cosine matrix  $\mathbf{C}$  is an orthonormal matrix. As a consequence of this property the inverse and transpose of  $\mathbf{C}$  have the same result and both the inverse and transpose operation can be used to switch the expression from B in A to A in B or vice versa. This is shown in Equation 2.3 [44].

$$[\mathbf{C}^{B/A}]^{-1} = [\mathbf{C}^{B/A}]^T = \mathbf{C}^{A/B} \quad (2.3)$$

### 2.1.2. Equations of motion

The attitude equations of motion (EOM) are described in the kinematic equations and the dynamic equations. The dynamic equations relate the change in angular rate to the total torque acting on the spacecraft while the kinematic equations relate the angular rate to a change in attitude angles,  $\theta$ . The kinematic equations can both be expressed in terms of Euler angles and quaternions, while the dynamic equations are a function of the angular rates,  $\omega$ , and spacecraft inertia,  $\mathbf{J}$  [45]. The kinematic equations of motion in terms of Euler angles, shown in Equation 2.7, follow from the total angular velocity shown in Equation 2.4, and assuming a spacecraft is in a circular orbit, giving a constant orbital rate,  $n$  [44].

$$\omega = \omega^{B/N} = \omega^{B/A} + \omega^{A/N} \quad (2.4)$$

The angular rate of reference frame A relative to N depends on the definitions of the reference frames. However, if the frames are defined as shown in Figure 2.1, then the angular rate of A relative to N can be written as shown in Equation 2.5. In which  $n$  is the constant, assuming a circular orbit, orbital rate of the spacecraft, and  $\bar{a}_2$  follows from Equation 2.2, where  $\bar{a}_2$  is defined perpendicular to the orbital plane. The angular rate of reference frame B relative to reference frame A can be described by introducing the time dependence of Euler angles. This allows to introduce the time derivative of the Euler angles in the three successive rotation to go from the angular rate in A to the angular rate in B. If the angular rates are then known functions of time, or can be measured, then the orientation of B relative to A as a function of time can also be found. After some manipulation this produces Equation 2.6, which after isolating the time derivative of the Euler angles on one side of the equal side can be inverted to arrive at Equation 2.7.

$$\omega^{A/N} = -n\bar{a}_2 = -n \begin{bmatrix} \cos(\theta_2)\sin(\theta_3) \\ \sin(\theta_1)\sin(\theta_2)\sin(\theta_3) + \cos(\theta_1)\cos(\theta_3) \\ \cos(\theta_1)\sin(\theta_2)\sin(\theta_3) - \sin(\theta_1)\cos(\theta_3) \end{bmatrix} \quad (2.5)$$

$$\begin{aligned} \omega^{B/A} = \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{bmatrix} &= \begin{bmatrix} \dot{\theta}_1 \\ 0 \\ 0 \end{bmatrix} + C_1(\theta_1) \begin{bmatrix} 0 \\ \dot{\theta}_2 \\ 0 \end{bmatrix} + C_1(\theta_1)C_2(\theta_2) \begin{bmatrix} 0 \\ 0 \\ \dot{\theta}_3 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & -\sin(\theta_2) \\ 0 & \cos(\theta_1) & \sin(\theta_1)\cos(\theta_2) \\ 0 & -\sin(\theta_1) & \cos(\theta_1)\cos(\theta_2) \end{bmatrix} \begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \\ \dot{\theta}_3 \end{bmatrix} \end{aligned} \quad (2.6)$$

$$\begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \\ \dot{\theta}_3 \end{bmatrix} = \frac{1}{\cos(\theta_2)} \begin{bmatrix} \cos(\theta_2) & \sin(\theta_1)\sin(\theta_2) & \cos(\theta_1)\sin(\theta_2) \\ 0 & \cos(\theta_1)\cos(\theta_2) & -\sin(\theta_1)\cos(\theta_2) \\ 0 & \sin(\theta_1) & \cos(\theta_1) \end{bmatrix} \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{bmatrix} + \frac{n}{\cos(\theta_2)} \begin{bmatrix} \sin(\theta_3) \\ \cos(\theta_2)\cos(\theta_3) \\ \sin(\theta_2)\sin(\theta_3) \end{bmatrix} \quad (2.7)$$

Then lastly, the dynamic of the rotational motion are given by Euler's rotational EOM, shown in Equation 2.8. The dynamic EOM is a function of the inertia matrix of the rigid body spacecraft, shown in Table 2.1, the current (body-fixed) angular rate and angular acceleration  $\dot{\omega}$ , and the external moment,  $\mathbf{M}$ , acting on the spacecraft [44]. In literature the external moment,  $\mathbf{M}$  is often split up in a control torque,  $\mathbf{T}_C$ , and a collection of disturbance torques,  $\mathbf{T}_D$ , see for example [18, 34]. Although it is also possible to distinguish the disturbance torques some more and include additional terms, as was done in [25]. This is especially interesting when one wants to model all the different angular accelerations on a spacecraft, to improve the understanding and accuracy of the models. Separating the total moment acting on the spacecraft body in a control torque and several different disturbance torques was done in [41].

$$\mathbf{J}\dot{\omega} + \omega \times \mathbf{J}\omega = \mathbf{M} = \mathbf{T}_C + \mathbf{T}_D \quad (2.8)$$

### 2.1.3. Perturbations

For the GOCE mission it was extremely important that all accelerations were kept minimal, while it was orbiting the Earth at a Very Low Earth Orbit (VLEO), at an altitude lower than 300 kilometres. At VLEO altitudes, the aerodynamic forces and torques acting on a spacecraft are significant. Given GOCE's mission, it was necessary to reduce all accelerations as much as possible. To reduce the linear accelerations the GOCE spacecraft

Table 2.1: GOCE Inertia matrix from GOCE data file, data computed on: 2013-11-11 T06:05:02. data given in  $\text{kgm}^2$ , and (BOL; EOL). Matrix elements have been confirmed by Sechi to the author of the data file through email contact.

<b>J</b>	x	y	z
x	173.8; 173.0	20.1; 20.2	-1.8; -1.7
y	20.1; 20.2	2,823.9; 2,750.5	-1.1; -1.1
z	-1.8; -1.7	-1.1; -1.1	2,796.6; 2,723.2

carried with it an ion engine, or Ion Propulsion Assembly (IPA). The main goal of the IPA was to counteract drag forces, but there are also other linear perturbations acting on the spacecraft. The goal of this thesis is to design an attitude controller that performs equally or better than the GOCE attitude controller. Hence any possible differences in linear accelerations between the simulated spacecraft and the actual GOCE mission will be ignored, assuming the simulated attitude would produce similar linear accelerations as the actual GOCE mission. This section then focuses on the perturbing angular accelerations acting on the GOCE satellite. The nature of these accelerations has been intimately discussed already by Visser et al. [41].

The solar radiation pressure torque, and aerodynamic torque are similar in the sense that an important parameter in these torques is the attitude of the spacecraft. The torque model for the aerodynamic torque is shown in Equation 2.9, and the model for the solar radiation pressure torque is shown in Equation 2.10. In these equations  $\mathbf{r}_{\text{ref}}$  is the reference point around which the moments are calculated,  $\mathbf{r}_{\text{com}}$  is the centre of mass of the satellite. The force and moment coefficients for both the aerodynamic and solar radiation pressure are given in  $C_{F,x}$  and  $C_{M,x}$  respectively. The reference area for the spacecraft is given by  $A_{\text{ref}}$ , and the air density, satellite velocity, speed of light, and solar flux at the spacecraft position are given by  $\rho$ ,  $v$ ,  $c$ , and  $\phi$  respectively. Finally, a simple eclipse model is used to detect when the Earth blocks the solar radiation. This eclipse model is based on [8].

$$\mathbf{T}_{\text{AE}} = (C_{M,A} l_{\text{ref}} + (\mathbf{r}_{\text{ref}} - \mathbf{r}_{\text{com}}) \times C_{F,A}) \frac{1}{2} \rho |v|^2 A_{\text{ref}} \quad (2.9)$$

$$\mathbf{T}_{\text{SR}} = (C_{M,S} l_{\text{ref}} + (\mathbf{r}_{\text{ref}} - \mathbf{r}_{\text{com}}) \times C_{F,S}) A_{\text{ref}} \frac{\phi}{2c} \quad (2.10)$$

A significant external (geo)magnetic field was necessary to magnetically control the GOCE spacecraft. However, the presence of this magnetic field also means that any undesired dipole created in the bus due to electronics, magnets, or induced magnetism will also interact with the EMF, and thus cause a torque. In Visser et al. [41] it is mentioned that information for dipoles caused by the spacecraft bus is available, but not for the payload. The dipoles are split into two categories, one for (electro)magnets inside the spacecraft, and one for induced magnetic dipoles. These are the hard and soft magnetic effects respectively. The model for the spacecraft bus is given by Equation 2.11, which includes both hard and soft magnetic effects. The payload dipole is an estimate, based on the residual torque after all other modelled torques have been removed [41]. The payload dipole torque is shown in Equation 2.12. The dipoles of the payload and bus are represented by  $\boldsymbol{\mu}$ , with subscript P and B respectively. The subscripts H and S, refer to the hard and soft magnetic effects. The soft magnetic effects are captured in the matrices  $M$ ,  $\mathbf{I}_T$  is the current going to the magnetorquers, and finally  $\mathbf{B}_E$  is the EMF.

$$\mathbf{T}_{D,B} = (\boldsymbol{\mu}_{B,H} + (M_{B,S} + M_{B,T} |\mathbf{I}_T|) \mathbf{B}_E) \times \mathbf{B}_E \quad (2.11)$$

$$\hat{\mathbf{T}}_{D,P} = (\hat{\boldsymbol{\mu}}_{P,H} + \hat{M}_{P,S} \mathbf{B}_E + \hat{S}_T \boldsymbol{\mu}_T) \times \mathbf{B}_E \quad (2.12)$$

Activity of the IPA introduces angular perturbations as well. One part of the torque caused by the thruster is when the thrust does not point directly through the centre of mass of the spacecraft. This torque is shown in Equation 2.13. This is however not the only torque caused by the thruster. Since the IPA is an ion thruster, it includes a large electromagnet. Thus there is also a torque caused by the interaction with the EMF. The thruster magnet dipole is modelled as proportional to the current running through it, and a combination of the hard magnetic part induced by the electromagnet, and a soft magnetic part that depends on activation of the magnetorquers. The model for the IPA dipole is shown in Equation 2.14. The current running through

the electromagnet is denoted by  $I_I$ , and the currents as a consequence of torquer actuation by  $I_T$ . The hard magnetic part is captured in  $\boldsymbol{\mu}_{I,H}$ , and  $M_{I,S}$  is a matrix in which the soft magnetic part is captured.

$$\mathbf{T}_{TM} = (\mathbf{r}_T - \mathbf{r}_{com}) \times \mathbf{F}_I \quad (2.13)$$

$$\boldsymbol{\mu}_I = (\boldsymbol{\mu}_{I,H} + M_{I,S} \mathbf{I}_T) I_I \quad (2.14)$$

The last perturbing angular acceleration that Visser et al. [41] makes mention of is the gravity gradient torque. The expression for the gravity gradient torque is shown in Equation 2.15, and includes the  $J_2$  term on top of a spherical Earth model, to account for the Earth's oblateness. This method of including the  $J_2$  term was originally proposed by Roithmayr [28], and its significance for including it in the torque models of the GOCE spacecraft is shown in [41]. In Equation 2.15 the gravitational parameter of the Earth is given by  $\mu$ , the satellite inertia is given by  $J$ , and the radius of the Earth by  $R_E$ . The radial distance between GOCE and Earth's centre is given by  $r$ , and  $\mathbf{u}_r$  and  $\mathbf{u}_n$  are unit vectors pointing from the satellite's centre of mass to the Earth's centre of mass, and along the rotation axis of the Earth respectively.

$$\begin{aligned} \mathbf{T}_{GG} = & 3 \frac{\mu}{r^3} \mathbf{u}_r \times J \mathbf{u}_r + \\ & \frac{\mu J_2 R_E^2}{2r^5} \left( 30 [\mathbf{u}_r \cdot \mathbf{u}_n] [\mathbf{u}_n \times J \mathbf{u}_r + \mathbf{u}_r \times J \mathbf{u}_n] + \right. \\ & \left. [15 - 105(\mathbf{u}_r \cdot \mathbf{u}_n)^2] \mathbf{u}_r \times J \mathbf{u}_r + 6 \mathbf{u}_n \times J \mathbf{u}_n \right) \end{aligned} \quad (2.15)$$

The work done by Visser et al. [41] shows that in general the torque models that have been set-up using the GOCE data provide a good fit to the observed behaviour. The two largest error sources that are mentioned are the lack of a complete model of the magnetic properties of the spacecraft. The dipoles that are estimated contain a larger than necessary uncertainty because of the absence of a full magnetic characterisation of the spacecraft. The other uncertainty source is an apparent mismatch between the aerodynamic model from ANGARA and the thermospheric density and wind data derived using a panel model. This uncertainty has the largest effect on the yaw-axis.

## 2.2. Magnetic attitude control

For any LEO satellite, attitude control through magnetorquers is an interesting option, because magnetic attitude controllers are small, reliable because of the absence of moving parts, and do not consume any propellant [39]. Furthermore the components of a magnetorquer are inexpensive, which combined with the absence of propellant reduces spacecraft and launch costs, and they have low activation noise, making them ideal for missions with stringent pointing and/or acceleration requirements [31]. Although magnetorquers are found in a lot of spacecraft, they are often only used for desaturating the reaction wheels, instead of being used for the actual attitude control [25, 47]. The reason why magnetorquers are not used widely in spacecraft is because first, a significant magnetic field is required. This means that after a certain altitude it no longer makes sense to use magnetorquers [34, 47]. Secondly, the magnetorquers cannot be used to generate a torque that is parallel to the magnetic field lines, making the satellite instantaneously under-actuated [6, 18]. However, assuming a spacecraft with a sufficient inclination, the direction of the EMF changes enough over the orbit of a satellite that each axis is controllable at some point [18].

This change in magnetic field is illustrated in Figure 2.3. This figure shows the total intensity of the EMF, as well as its separate components in a North East Down frame. Figure 2.3 was made using the magnetic model from the International Geomagnetic Reference Field (IGRF). For this figure 1 November 2011 was chosen as the day to model the magnetic field on, at an altitude of 250 kilometres. In Figure 2.3, it can be seen that the total magnetic intensity shows differing behaviour over the North and the South Pole. Furthermore, there is a large depression in total magnetic intensity over South America. Most important, however, is that Figure 2.3 shows the significant changes in EMF direction over different latitudes.

The torque on a spacecraft that can be produced by interacting with the EMF is shown in Equation 2.16.  $\mathbf{B}_B$  is the EMF in the body frame, which is shortly written as,  $\mathbf{B}$ ,  $\boldsymbol{\mu}_T$  is the dipole generated by the magnetorquers, and  $\mathbf{T}_C$  is the torque as a consequence of the dipole interaction with the EMF. The consequence of having a fully magnetic attitude control system is highlighted in Figure 2.4. It can then be concluded that

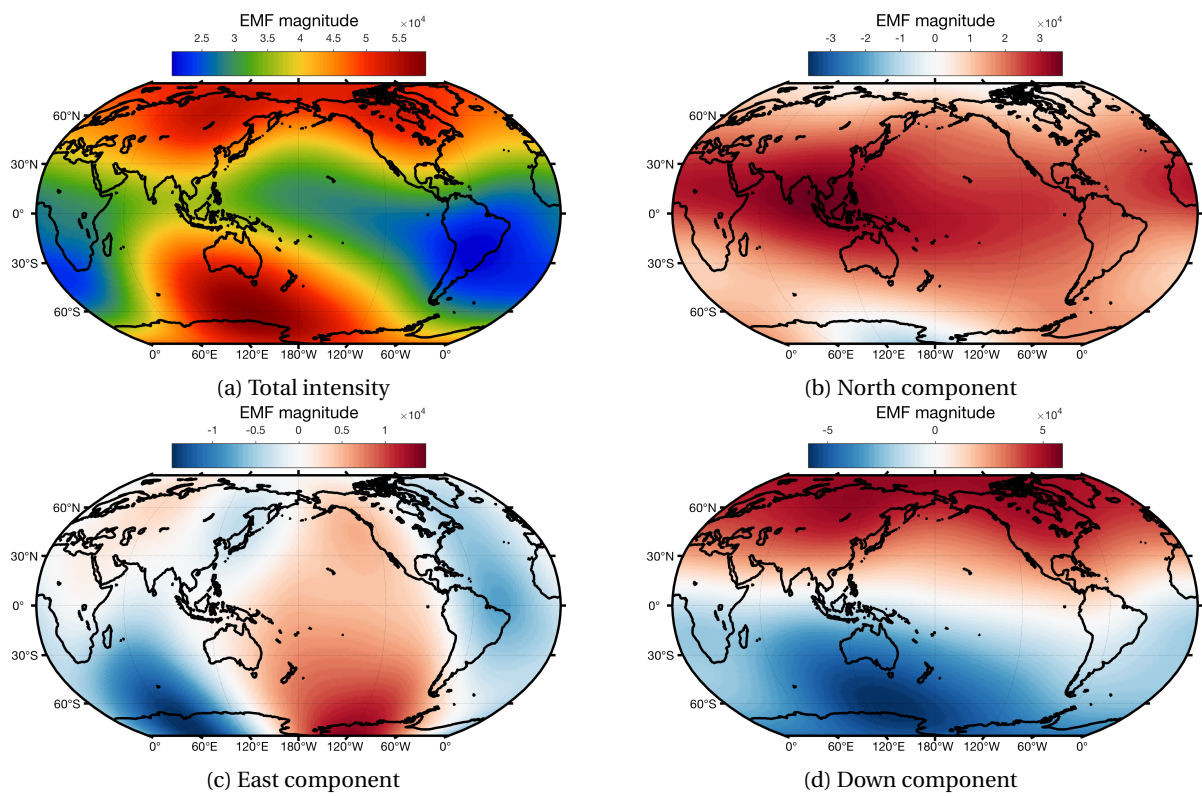


Figure 2.3: EMF shown in Robinson map projection as the main field total intensity (a), and the North, East, Down components, (b), (c), and (d) respectively. Modelled using IGRF-12 on 1 November 2011 at an altitude of 250 kilometres.




Figure 2.4: Mechanical torque generated by magnetorquers is always perpendicular to the geomagnetic field vector. This implies that yaw is not controllable over poles, and roll is not controllable over equator. [47]




Figure 2.5: Magnetic torquer applicable torque. [32]

to fully control the spacecraft using magnetic control only, the system relies on a variability of the direction of Earth's magnetic field. This entails that fully magnetic, 3-axes, attitude control is only possible when the orbit of a spacecraft is such, that it experiences a sufficient variation in the magnetic field to guarantee the stability of the spacecraft's attitude [34]. Since GOCE was in a near-polar orbit,  $\pm 96.7^\circ$ , the magnetic field it experienced did change sufficiently to guarantee attitude stability.

$$\mathbf{T}_C = \boldsymbol{\mu}_T \times \mathbf{B}_B = \boldsymbol{\mu}_T \times \mathbf{B} \quad (2.16)$$

However, the problem remains that the satellite is instantaneously under-actuated. Previously the magnetic torque on a satellite was written as Equation 2.16, but given that the torque will act on the plane perpendicular to the magnetic field vector the equation has to be rewritten. A torque aligned with the magnetic field vector is not available, and if the attitude control system requests a torque based on the perturbing torques acting on the satellite, then this required torque will have to be projected on the plane normal to the instantaneous magnetic field vector. In Figure 2.5, the line  $\mathbf{T}_{\perp B}$  shows the projection of a required torque,  $\mathbf{T}_R$ , onto the plane perpendicular to the magnetic field direction.

In Sechi et al. [32] the minimum Euclidean norm allocation rule is used to project the required torque along the plane orthogonal to the Earth magnetic field vector. After some manipulation, this allows to rewrite Equation 2.16 into Equation 2.17. In Equation 2.17 the required torque is commanded by the attitude control system is represented by  $\mathbf{T}_R$ ,  $\mathbf{B}$  represents the EMF, and  $\mathbf{B}(\mathbf{b})$ , is the cross-multiplication matrix for  $[\times \mathbf{B}]$ . Looking at Equation 2.17, it can be seen that above the equator where  $B_x^2 / \|\mathbf{B}\|^2 \rightarrow 1$  the torque  $T_{R,x}$  cannot be applied, and moreover, the pitch axis torque will affect the roll axis. Note, when above the equator, the Z-axis component of the Earth magnetic field on the body-axis reference frame is close to zero. Combine this with, in the case of the GOCE spacecraft, the large difference in inertia between the X- and Y-axis, shown in

Table 2.1, and it can be seen that this is a problematic situation.

$$\mathbf{T}_C = \left( \mathbf{I}_{3 \times 3} - \frac{\mathbf{B}\mathbf{B}^T}{\|\mathbf{B}\|^2} \right) \mathbf{T}_R = \frac{1}{\|\mathbf{B}\|^2} \mathbf{B}(\mathbf{b})\mathbf{B}(\mathbf{b})^T \mathbf{T}_R = \frac{1}{\|\mathbf{B}\|^2} \begin{bmatrix} 1-B_x^2 & B_x B_y & B_x B_z \\ B_y B_x & 1-B_y^2 & B_y B_z \\ B_z B_x & B_z B_y & 1-B_z^2 \end{bmatrix} \mathbf{T}_R \quad (2.17)$$

One proposed solution to this problem is to suspend control action on the Y-axis when the roll axis is not controllable [31, 32]. The proposed implementation for decoupling the roll and pitch axes is by implementing a time varying weighting matrix,  $\mathbf{S}(\mathbf{B})$ , that is used to pre-multiply the required torque vector with. Doing this changes Equation 2.17 to Equation 2.18. The weighting matrix is defined similarly in both Sechi et al. [31] and Sechi et al. [32], with the only difference being the absence of  $\gamma$  in Sechi et al. [31]. The constant weighting factor is within the range  $0 \leq \gamma \leq 1$ , Although Sechi et al. [32] makes no mention as to why the constant weighting factor was introduced, it is presumably to be able to fine tune the trade-off between X- and Y-axis control performance. The time varying weighting matrix itself is shown in Equation 2.19. Finally, it is clear that to be able to apply time dependent control, and project the required control torque on a plane perpendicular to the local magnetic field direction, information about the magnetic field is necessary. This information comes from the magnetometer readings, which have been calibrated and processed in Visser et al. [41].

$$\mathbf{T}_C = \frac{1}{\|\mathbf{B}\|^2} \begin{bmatrix} 1-B_x^2 & B_x B_y & B_x B_z \\ B_y B_x & 1-B_y^2 & B_y B_z \\ B_z B_x & B_z B_y & 1-B_z^2 \end{bmatrix} (\mathbf{S}(\mathbf{B})\mathbf{T}_R) \quad (2.18)$$

$$\mathbf{S}(\mathbf{B}) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 - \gamma \frac{B_x^2}{\|\mathbf{B}\|^2} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.19)$$

To control the spacecraft's attitude the magnetorquers torque against the local EMF. This is achieved by three rods parallel to the spacecraft X-, Y-, and Z-axis. Each torquer consists of a core material, and the windings around it [17]. Applying a current through these windings creates a magnetic field, which gets amplified by the core material. Although the generated dipole also depends on the core material, it is proportional to the number of windings, area of the windings, and current going through it [16, 22]. Since the core material, number of windings, and winding area cannot be changed for obvious reasons, the variable through which control has to be actuated is the amount of current going through the windings. This current can be found by using the commanded control torque, transforming it to a dipole, and calculating the amount of current necessary to generate that dipole based on the magnetorquer specifications. Finding the required dipole based on the commanding torque follows from Equation 2.16 and is shown in Equation 2.20. In which the Moore-Penrose pseudo-inverse (+) is used to minimise the norm of the calculated dipole vector. Using the Moore-Penrose pseudo-inverse was deemed necessary, in combination with applying weights to the EMF cross-multiplication matrices, to compensate for numerical instabilities when applying small roll torques. It is however computational intensive, and it was already found that the weights are no longer necessary. Thus when the work done in this thesis is continued it is advised to evaluate the stability of using Equation 2.21 instead of Equation 2.20.

$$\boldsymbol{\mu}_T = [\mathbf{B}(\mathbf{b})^T \mathbf{B}(\mathbf{b})]^+ \mathbf{B} \times \mathbf{T}_C \quad (2.20)$$

$$\boldsymbol{\mu}_T = \frac{\mathbf{B} \times \mathbf{T}_C}{\|\mathbf{B}\|^2} \quad (2.21)$$

### 2.3. Control theory

Throughout the history of control theory a lot of different controller types have been thought of and implemented. A small grasp from them all include controllers like  $H_\infty$ , Lyapunov, and linear quadratic regulator. Each of them has its own disadvantages and advantages, but since there are so many different types of controllers, we will not try to discuss them all in this section. Instead, this section will first focus on conventional, linear, PID-control theory. This type of control is still widely used today, and is useful as a stepping stone for a basic understanding of control theory in general. Then Nonlinear Dynamic Inversion (NDI) will be discussed, because NDI control theory is a critical introduction for the last control type to be discussed: Incremental Nonlinear Dynamic Inversion (INDI). The reason for choosing INDI as the main controller to discuss in more detail is that it is a relatively new control type and promises to be more flexible and more stable in controlling a system [1, 15, 33, 35].






Figure 2.6: General block diagram of a PID controller. [46]




Figure 2.7: (a) System showing input and output. (b) pole-zero plot of the system. [23]

### 2.3.1. Linear control theory

One of the classical examples of a controller is a PID-controller. This is a controller in which a reference signal is compared to a measured system signal and the difference between the two, the error, is multiplied with a gain, after which a control command is produced, see Figure 2.6. PID itself stands for Proportional, Integral, Derivative, and they refer to how the error signal is amplified. The proportional gain simply amplifies the error, the integral gain amplifies the time history of the error which can be used to remove an offset, and the derivative gain amplifies the change in error which can reduce overshoot. The difference between the reference signal and the output signal, the error, influences the controller output, which in turn affects the plant, which changes the error. In other words, Figure 2.6 is a feedback controller, whose goal it is to steer the error,  $e$ , to zero as the time,  $t$ , goes to infinity.

A linear controller does not have to be a PID-controller, but PID-controllers are useful to help explain basic control theory. While PID-controllers are a straightforward concept to understand, it can be difficult to tune the gains,  $K$ , of the controller depending on the system. Furthermore, once the gains have been tuned up to a satisfactory point, it is not necessarily guaranteed that the same gains will work in all conditions. Especially for nonlinear systems, when the state changes the gains that could once be used for stable control might suddenly result in an unstable system. This is in contrast to the most important system specification of any controller, that is to make it stable. A system is considered stable when the natural response of the system approaches zero as time approaches infinity [23]. To ensure the natural response of the system moves towards zero, it is required for the system poles to be in the left half-plane, as shown in Figure 2.7 with the cross.

The reason why poles need to be in the left half-plane is because that ensures that the natural response of a system is either an exponential decay or a damped sinusoidal. Which ensures a natural response moving towards zero for time going towards infinity. The interpretation of the poles being in the left half-plane is that the real part of the  $s$ -plane is negative. Figure 2.7 shows the poles (cross) and zeros (circle) of a system with transfer function denominator  $s + 5$ , and transfer function numerator  $s + 2$ . The poles of this system are the value of the Laplace transform variable,  $s$ , that causes the denominator of the transfer function to be zero. The zero of a transfer function does not affect the nature of the response, but instead only the residue or amplitude [23].

The transfer functions of a controller can be changed by changing system gains. The gains can move the system poles and as such change the damping of the system. This however means that when the gains are not chosen correctly the system can become unstable. Furthermore, even if the gains are chosen correctly, when something about the system changes (operating environment, a physical change, component malfunction), or in the case of nonlinear systems the state changes, the transfer functions change and the selected gains can become unsuitable. This is why for complex systems that operate in a large regime a technique called gain scheduling is often used where gains are tuned for specific operating ranges. Apart from gain scheduling being a tedious process, it can also cause instability when a set of gains is switched when the system moves from one pre-defined operating regime to the next one [26].




Figure 2.8: Illustration of outer- and innerloop control of an NDI controller. [1]

### 2.3.2. Nonlinear dynamic inversion

Nonlinear dynamic inversion (NDI) started to attract scientific attention around the 1980s. Initially it was researched as a means to make controllers more robust, while maintaining aggressive control in high performance manoeuvres and later on as well as a solution for maintained stability during actuator failure or damages [2, 14]. The way NDI control achieves this is by using state feedback to linearise the system. Below the way an NDI controller works will be described in a set-up that is commonly used in NDI control [1, 2, 14, 30, 33, 42].

The core of NDI control is to invert a nonlinear system using state feedback, so that a linear controller can be applied to the system. This results in the entire control system having two control loops. An outer linear control loop, and an inner linearisation loop. The system dynamics are expressed as shown in Equation 2.22 with the nonlinear state output shown in Equation 2.23 [14, 42]. The relative degree of the system is given by the number of differentiations needed for the input,  $\mathbf{u}$  to appear in the output,  $\mathbf{y}$ . When the output of the system is considered to be the state,  $\mathbf{x}$ , Equation 2.23 can be written as Equation 2.24. Since the output is a nonlinear function of the state and indirectly related to the physical control input,  $\mathbf{u}$ , the goal is to find a direct relation between  $\mathbf{y}$ , and  $\mathbf{u}$ . Finding this relation allows to invert the equation to find an expression for the physical control input that can be tuned to linearise the system. Which in turn is achieved by using the virtual control input  $\mathbf{v}$ , which can be achieved by using a linear controller [33].

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{G}(\mathbf{x})\mathbf{u} \quad (2.22)$$

$$\mathbf{y} = \mathbf{h}(\mathbf{x}) \quad (2.23)$$

$$\mathbf{y} = \mathbf{x} \quad (2.24)$$

In the general example shown in Equation 2.22 and Equation 2.23 taking the derivative with respect to time on Equation 2.25 allows one to find an expression for the physical control input. This is commonly written as shown in Equation 2.25, in which the virtual control input is inserted. As was mentioned, this virtual control input linearises the system, and can be chosen or set by a linear controller [1, 2, 33]. The smooth vector field  $\mathbf{f}(\mathbf{x})$ , in this case represents the angular rate dependent part of the dynamic EOM, while the matrix  $\mathbf{G}$  in this case represents the inverse of the inertia matrix of the spacecraft.

$$\mathbf{u} = \mathbf{G}^{-1}(\mathbf{x})[\dot{\mathbf{v}} - \mathbf{f}(\mathbf{x})] \quad (2.25)$$

$$\dot{\mathbf{y}} = \dot{\mathbf{v}} \quad (2.26)$$

Thanks to the inversion producing a control command, in Equation 2.25, the nonlinearities in the system are cancelled. A simple linear, input-output, relation is thus obtained between the output,  $\mathbf{y}$ , and the virtual control input,  $\mathbf{v}$ . This is shown in Equation 2.26. Another result from this relationship is the fact that the relation between input and output is now also decoupled, since the virtual control input only affects the output. With the formulas for NDI-control known, it is easier to understand what an NDI-controller looks like by having an example. Thus Figure 2.8 shows how the theory mentioned above fits into a controller layout.

The most important assumption underlying NDI-control is that the model of the system is assumed to be exactly known. This means that the control command generated by NDI assumes that the system will behave according to the model. Thus model uncertainties will make NDI-control unstable. Furthermore complete and accurate knowledge of the state of the system is also assumed to be available, and available at a sufficiently high rate. These two points emphasise that NDI-control is at a disadvantage when it comes to uncertainties. Therefore a new control type that depends less on accurate model knowledge is described below.

### 2.3.3. Incremental nonlinear dynamic inversion

The goal of Incremental Nonlinear Dynamic Inversion (INDI) is to reduce the dependency of the closed loop system on exact model and plant dynamics knowledge. INDI-control is thus a continuation on NDI-control, and has to this author's knowledge not been widely discussed as a means of spacecraft attitude control yet. In the INDI framework control is achieved by using the theory of NDI-control, but applying it to an incremental system, instead of directly. Meaning that the control commands are also specified in control increments, as opposed to the direct control input that follows from NDI. The expression in incremental form improves the robustness of the closed-loop system, because no attempt is made to capture the closed-loop dynamics in a model that is to be followed. The desired closed loop dynamics follow from closing the feedback loops [1, 15, 33, 35, 42]. As was done for NDI-control, the theory and formulas of a generalised system controlled using INDI will be provided below. The work done in Acquatella et al. [1] has proven to be especially useful in understanding INDI-control.

The derivation of INDI control starts similar as NDI control, i.e. with Equation 2.22 and Equation 2.24. However, to find an expression for  $\dot{\mathbf{x}}$ , a standard Taylor series expansion of this equation is taken. This provides a first-order approximation of  $\dot{\mathbf{x}}$ , for  $\mathbf{x}$  and  $\mathbf{u}$  in the neighbourhood of  $\mathbf{x}_0$  and  $\mathbf{u}_0$ . This is shown in Equation 2.27, in which higher order terms are ignored. Similar as was done for NDI control, the goal for INDI control is to invert Equation 2.27 so that the increment in control is a function of the output of the system.

$$\dot{\mathbf{y}} = \dot{\mathbf{x}}$$

$$= \mathbf{f}(\mathbf{x}_0) + \mathbf{G}(\mathbf{x}_0)\mathbf{u}_0 + \left. \frac{\partial}{\partial \mathbf{x}} [\mathbf{f}(\mathbf{x}) + \mathbf{G}(\mathbf{x})\mathbf{u}] \right|_{[\mathbf{x}=\mathbf{x}_0, \mathbf{u}=\mathbf{u}_0]} (\mathbf{x} - \mathbf{x}_0) + \left. \frac{\partial}{\partial \mathbf{u}} [\mathbf{G}(\mathbf{x})\mathbf{u}] \right|_{[\mathbf{x}=\mathbf{x}_0, \mathbf{u}=\mathbf{u}_0]} (\mathbf{u} - \mathbf{u}_0) \quad (2.27)$$

In Equation 2.27, the incremental control command is shown by,  $(\mathbf{u} - \mathbf{u}_0)$ . As this follows from a standard Taylor series expansion it suggests that in the neighbourhood of the reference state,  $\mathbf{x} = \mathbf{x}_0$  and  $\mathbf{u} = \mathbf{u}_0$ , the nonlinear system can be approximated by linearisation about that state. Inverting Equation 2.27 to isolate the control command increment gives Equation 2.28. Assuming time-scale separation is valid, i.e. the time increments are sufficiently small, and the control update rate is sufficiently high, implies that Equation 2.29 is true. Then Equation 2.28 can be rewritten to Equation 2.30, by using the definitions seen in Equation 2.22 and Equation 2.24.

$$(\mathbf{u} - \mathbf{u}_0) = \delta \mathbf{u} = \left. \frac{\partial}{\partial \mathbf{u}} [\mathbf{G}(\mathbf{x})\mathbf{u}] \right|_{[\mathbf{x}=\mathbf{x}_0, \mathbf{u}=\mathbf{u}_0]}^{-1} \left[ \mathbf{v} - \left( \dot{\mathbf{x}}_0 + \left. \frac{\partial}{\partial \mathbf{x}} [\mathbf{f}(\mathbf{x}) + \mathbf{G}(\mathbf{x})\mathbf{u}] \right|_{[\mathbf{x}=\mathbf{x}_0, \mathbf{u}=\mathbf{u}_0]} (\mathbf{x} - \mathbf{x}_0) \right) \right] \quad (2.28)$$

$$\left. \frac{\partial}{\partial \mathbf{x}} [\mathbf{f}(\mathbf{x}) + \mathbf{G}(\mathbf{x})\mathbf{u}] \right|_{[\mathbf{x}=\mathbf{x}_0, \mathbf{u}=\mathbf{u}_0]} (\mathbf{x} - \mathbf{x}_0) \ll \left. \frac{\partial}{\partial \mathbf{u}} [\mathbf{G}(\mathbf{x})\mathbf{u}] \right|_{[\mathbf{x}=\mathbf{x}_0, \mathbf{u}=\mathbf{u}_0]} (\mathbf{u} - \mathbf{u}_0) \quad (2.29)$$

$$\delta \mathbf{u} = \left. \frac{\partial}{\partial \mathbf{u}} [\mathbf{G}(\mathbf{x})\mathbf{u}] \right|_{[\mathbf{x}=\mathbf{x}_0, \mathbf{u}=\mathbf{u}_0]}^{-1} [\mathbf{v} - \dot{\mathbf{x}}_0] = \mathbf{G}(\mathbf{x}_0)^{-1} [\mathbf{v} - \dot{\mathbf{x}}_0] \quad (2.30)$$

The control law shown in Equation 2.30 is thus in incremental form, meaning that the provided increments in control output have to be added to the control output of the previous iteration. The theory of INDI-control allows for a control law that does not depend on exact knowledge of the plant dynamics, and decreases dependency on model knowledge. However, it does rely more on accurate state measurements, making it more sensitive to sensor aspects than NDI-control. Again, an overview of an INDI-control lay-out is shown to visualise how the above-mentioned theory can be put into an actual controller. This is done in Figure 2.9.

Additionally to the theory above, there is a variation on INDI-control that is relevant to discuss as well. The INDI-controller shown in Figure 2.9 sets the reference angular rate to track based on multiplying the attitude error with a linear gain. However, having discussed the advantages of NDI-controllers, and realising

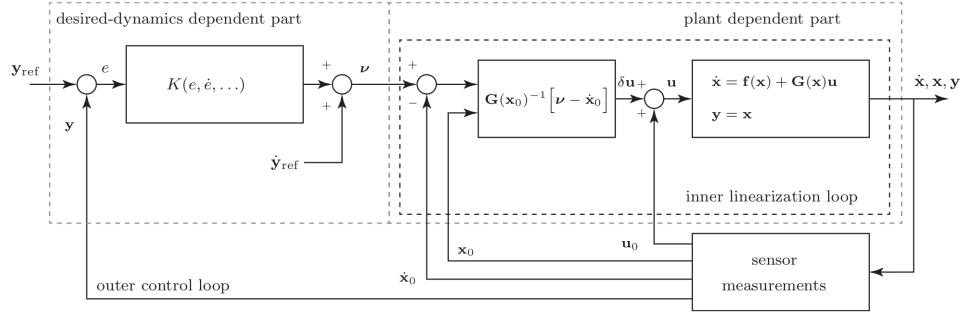


Figure 2.9: Illustration of outer- and innerloop control of an INDI controller. [1]




Figure 2.10: Illustration of outer- and innerloop control of an INDI controller, that uses NDI in its outer loop. [1]

that the kinematic EOM are free of uncertainties, given that they are a given geometrical representation, the reference angular rate can also be set using NDI. In this case the kinematic EOM then have to be inverted such that a reference angular rate is produced that achieves a desired rate of change of the Euler angles, based on the attitude error. An example of this INDI-controller, in which the outer-loop employs NDI to set a reference angular rate is shown in Figure 2.10.

For INDI-control, the most important assumptions are the following. 1) Complete and accurate knowledge about the state of the system. This can be achieved by measurements, but these will always to some extent contain errors (noise, bias, delay). 2) It is assumed to be valid that the angular velocities of the spacecraft only change by integrating the angular accelerations. The angular accelerations themselves change under influence of torques acting on the spacecraft. 3) Complementary to the assumption before, it is assumed that the dynamics of the control actuators evolve much faster than the angular rates. Based on the assumptions that underlie INDI-control, it could have been very well suited for attitude control of the GOCE spacecraft. GOCE carried accurate star trackers with it for attitude measurements, but more importantly is the accuracy of the angular acceleration measurements coming from the common mode product of the EGG. Additionally, the magnetorquers can be controlled very quickly, and have fine actuation control, thus they are appropriate for assuming fast control action.

# 3

## METHODOLOGY

In this chapter the methodology that is used for the thesis research is discussed. The chapter will start with explaining the research set-up in section 3.1, then the way individual pieces of the code are interconnected and the way the model works is discussed in section 3.2. The design of the main NINDI-controller, and the two reference controllers will be explained in section 3.3. Finally, the method for tuning the gains used in the controllers is explained in section 3.4.

### 3.1. Research set-up

The goal of the thesis research was mentioned in chapter 1, but is repeated below for clarity. To achieve the goal of the thesis a set of research questions was formulated, those are also shown in chapter 1. To answer the research questions, it is necessary to design a controller that steers the attitude to a reference point and stabilises the satellite, such that its angular rate and acceleration are minimal as well. Furthermore, the reason for using an INDI controller is that it promises to be more robust and stable than a more conventional controller, and also better than the original GOCE controller.

*The research objective is to contribute to the development of new attitude control systems for satellites, by designing an INDI-based control architecture for GOCE, evaluating its performance, showing the feasibility of such a controller for magnetic attitude control, and investigate whether controller performance improves by using a linear acceleration tracking control frame.*

During the literature study a trade-off was made on which software tools best to use for the research [20]. An overview and comparison of the tools that were considered for doing the simulations in, is given in Table 3.1. It should be noted that Table 3.1 reflects a personal evaluation of the available options, and does not hold for everyone since it is heavily influenced by the amount of experience with a certain tool. Based on this evaluation, MATLAB™ is the tool of choice to do the research in. Additionally, this choice also makes sense because the torque models that are going to be used have been set-up in MATLAB™.

Having decided on a simulation tool to work with, the overall plan to design and verify the attitude simulator will be laid out. The core piece of the research will consist of setting up, tuning, and validating the controller. The design of the model, and the way the controller fits into the model, will be described first, before discussing the design of the controller. Tuning the controller will be discussed in section 3.4. After

Table 3.1: Comparison of software/programming languages

	Experience	Availability pre-build functions	Documen- tation	Ease of use	Simulation speed	Overview	Support (incl. online)
MATLAB	+	+	+	+	-	0	+
MATLAB Simulink	+	+	+	0	-	+	0
C++	0	0	0	-	+	-	+
C#	-	-	0	-	0	-	+

having tuned the controller, the model will be used to simulate the behaviour of the satellite state, and the torque acting on the satellite. Since the disturbance torques acting on the satellite are highly dependent on the external environment, different date ranges will be chosen to verify the performance of the controller under different conditions. The date ranges are chosen such that the performance of the controller is evaluated under a wide variety of conditions. One chosen period is 1 – 9 January 2010, in which the spacecraft had a relatively high average altitude and thermospheric activity was low. Therefore the environment should be relatively calm. In contrast with this period is the period 15 – 21 March 2013, where the spacecraft's average altitude was considerably lower and there was a high thermospheric activity. Finally a period in between was chosen, 16 – 23 May 2011, in which there also was increased thermospheric activity.

These simulations will initially be done under nominal conditions, in which there are no error sources present in the model. Since all kinds of errors are present in real hardware, the performance of the controller also needs to be evaluated when there are errors present. The results obtained from the simulations will be processed and used to quantify the performance of the NINDI-controller. The results will be analysed and discussed, and then, based on the processed results the research sub-questions can be answered individually. The answers to the sub-questions are then used to answer the main research question, after which a conclusion can be drawn on the feasibility of using this new type of attitude control for spacecraft.

### 3.2. Model architecture

A high-level overview of the code is shown in Figure 3.1. At the highest level, the attitude simulator works by specifying settings in a standard MATLAB<sup>™</sup> script, when the script is run the desired settings are used in a function that does the actual attitude simulation. Before the function starts with the simulation, data is loaded based on the specified date(s) to do a simulation for. Most of the loaded data is stored as a time series, with time intervals of 10 seconds. This means that those data points have to be interpolated based on the specified simulation time settings. During the interpolation of the data the choice of reference frame influences the actions to be undertaken. This is because the options for the reference frames have a different definition of the control frame. See subsection 3.2.3 for the exact definition of the reference frames. The same also goes for the choice of attitude representation. However, in this case it affects the final preparation before the simulation. The preparation for the simulation ensures amongst others that all attitude angles are in the desired attitude representation.

Concerning the loading of the data itself, the external data files are processed data products from the actual GOCE mission, and also estimations of environmental parameters. For example, the position of the satellite above the Earth is derived from on-board GPS measurements, while the atmospheric density at that point is an estimation. The attitude of the spacecraft is also loaded, because some data files are still in the spacecraft frame. Therefore it is necessary to know in which orientation they were measured, so that they can be converted to Earth frame measurements in the simulation. The magnetic field measurements for example are measured by the on-board magnetometers, but the measurements depend on the attitude of the spacecraft. To be able to use the measured magnetic field values in the simulator they first need to be converted back from the body-frame to the Earth frame. This then allows to use them with a simulated attitude that likely is different from the measured attitude at that same point. As was mentioned before, linear accelerations on the spacecraft will be ignored, thus the simulated position of the spacecraft will always be the same as the measured GPS position of the actual spacecraft.

After the preparation for a simulation is done, all the data files that have been loaded are used in the simulation. The gains are specified in the main MATLAB<sup>™</sup> script and are then used in the controller. The controller is also specified in the main MATLAB<sup>™</sup> script. Finally the choice of reference frame also has an influence on what happens during the simulation. This is because, again, the choice of reference frame changes the definition of the control frame. The length of the simulation depends on the size of the time vector, and after the simulation is done all the data is saved in a data file. In Figure 3.1 a visual overview is given of what was described above. The subsections below will zoom in on the most relevant blocks.

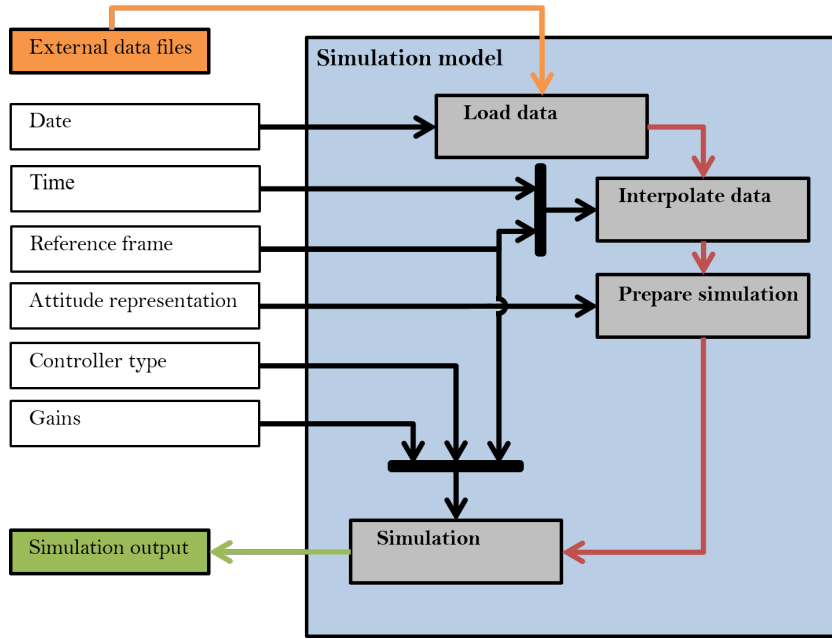


Figure 3.1: High-level code overview

### 3.2.1. Simulation

This subsection will focus on the simulation block. The internals of the simulation block is shown in Figure 3.2. The controller block is shown in a colour that stands out, because the design of the controller is a main focus of the thesis and will be discussed in depth in section 3.3. At the start of the simulation the state is compared against the reference state, and the differences are fed to the controller. It depends on the controller type which components of the state – attitude angles, angular rates, and angular accelerations – are used in the controller. The controller then outputs a required torque. This torque is not directly transformed to a current command that is to be send to the magnetorquers though, instead it is first fed through the time dependent control block. In this block the commanded torque is projected on the plane perpendicular to the local magnetic field direction which makes sure a required torque is produced that is feasible in the local magnetic field direction. The exact method of doing this is described in section 2.2.

The commanded torque is then pushed in a function that transforms the required torque to a vector of currents that are to be send to the magnetorquers. The currents are found by first calculating the dipoles required to produce the required torque and this can then be converted to a current assuming a linear slope. The assumption that the magnetorquers have a linear relationship between current and dipole comes from Kolkmeier et al. [17] and is valid up to at least 544 mA. This current together with the current satellite state, and a lot of variables from the loaded data, then go into the torque block. What goes on inside the torque block is discussed in subsection 3.2.2, so this will not be discussed here.

The torque block outputs a net torque, i.e. a sum of all the disturbance torques and the control torque. The goal of the control torque is to counteract the disturbance torque, so ideally the output from the torque block should be small. The torque is combined with the spacecraft state and inertia – which comes from the ‘prepare simulation’ block, although not explicitly shown – to update the state. The state update is done by using the kinematic and dynamic equations shown in subsection 2.1.2. After having calculated the time derivative of the attitude angles and the angular rate, the state is propagated using a simple Euler integration. After the state update the iteration starts again, until the final time has been reached.

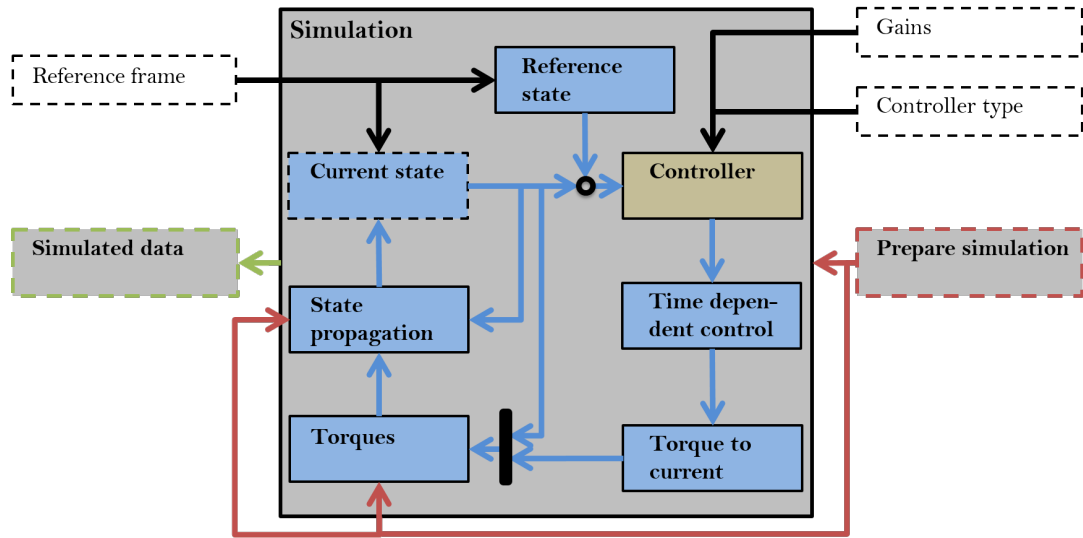


Figure 3.2: Simulation overview

### 3.2.2. Torque calculation

The ‘Torques’ block in Figure 3.2 is expanded in Figure 3.3. This figure does not show how exactly the torques are calculated, this would require to further zoom in on each individual torque block. However, the formulas for the disturbance torques are mentioned in subsection 2.1.3, the purpose of Figure 3.3 is to show some of the dependencies of the torque models, and how the torque models fit into the total model.

What is most notable in Figure 3.3 is how much dependencies there are from outside the simulation environment. These dependencies are loaded before the simulation starts, get interpolated, and in some cases they are first transformed to an Earth reference frame so that they can be used for attitude dependent torques. Two of these dependencies are the coefficients used to calculate the solar radiation pressure torque and the aerodynamic torque. For both cases there are force and moment coefficients. These coefficients can be based on a panel model that is able to provide them, but in this case are calculated using ANGARA Monte-Carlo simulation software. For future users of the code it is important to realise that the reference point for the moment coefficients for both the solar radiation pressure torque, and the aerodynamic torque are different. This difference stems from an early confusion in the solar radiation pressure model.

Note also that in Figure 3.3 the inputs for the control, residual dipole, fitted dipole, and thruster dipole torques appear equal. For one part this is because Figure 3.3 does not show all inputs in favour of preserving overview. For the other part it is because the inputs do have some corresponding dependencies. The control torque for example takes the current to the magnetorquers as an input, but the currents are then processed to generated dipoles, according to GOCE hardware specifications. The residual dipole on the other hand consists of a long term deterministic part, and soft magnetic effects as a result of magnetizable materials reacting to the external magnetic field of either the Earth or the magnetorquers. The fitted dipole torque, and thruster dipole torque are similar, but both also include current to the IPA as an input. The fitted dipole torque only contains an estimate for the hard magnetic dipole, while the no variable dipoles are estimated. This is to prevent fitting to the control algorithm of aerodynamic signals. The thruster dipole torque calculates the dipole as caused by the IPA magnetic field, based on the GOCE hardware specifications.



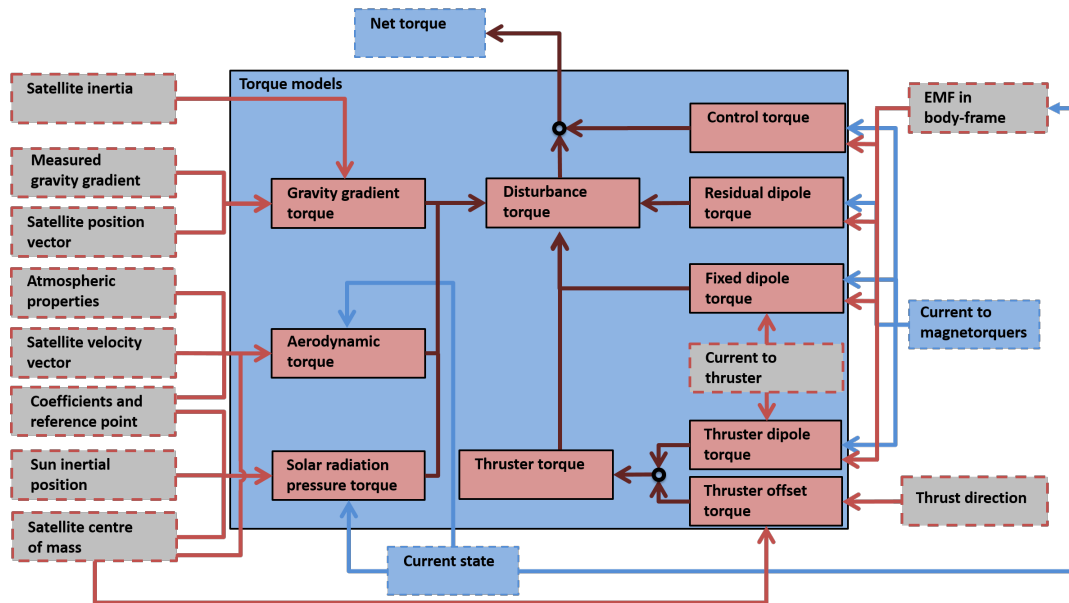


Figure 3.3: Overview of the torque calculation.

### 3.2.3. Reference frames

Some reference frames subscripts that are used throughout the code are shown in Table 3.2, the bold letters are the subscripts used. The table is shown here for clarity if one wants to work with the code. However the main reference frames that will be focused on here are the control reference frames. As discussed in section 2.1 the control of a spacecraft state is done as the movement of the body-axes frame – definition for GOCE in Table 3.2 – with respect to another frame. One of these frames, the orbital – or ‘O-frame’ – is already shown in Table 3.2. What is important to note is that the x-axis direction is defined as the velocity with respect to Earth’s surface. This means that the co-rotation of the Earth’s atmosphere is included in this frame. Two other frames that are also used as a control frame are discussed hereafter, these frames are the ‘L-frame’ and the ‘D-frame’.

The ‘L-frame’ is similar to the ‘O-frame’, except the x-axis direction is defined as the velocity in the inertial frame. So the x-axis will not take into account the co-rotation of the atmosphere. The origin of the reference frame is still on the satellite centre of mass, and the y-axis definition is also the same. The definition of the D-frame is different from the other two, because it serves a different purpose. The goal of the D-frame is to provide a reference frame in which small attitude angles result in pointing in the direction of the linear perturbations. This reference frame’s origin is the same as the other two. However its x-axis points in the direction of the linear acceleration. When the control frame tracks the linear acceleration and the spacecraft’s attitude in this frame is small, the linear accelerations on the (GOCE) spacecraft are minimised by providing thrust in the direction of the linear acceleration. In the implementation of this reference frame in the code, it is assumed that the aerodynamic drag and solar radiation pressure drag capture all the linear accelerations.

Having a definition for the x-axis, the other two axes still need to be defined. The choice for this reference frame is to have the z-axis also pointing more or less down. This is achieved by using the Gram-Schmidt method to project the down-facing body z-axis on the new x-axis and then make it perpendicular again to

Table 3.2: Reference frame definitions, all right-handed. [41]

Frame	Origin	x	y	z
<b>Inertial</b>	Earth’s centre of mass	Mean equinox	(-)	Rotation axis
<b>Earth fixed</b>	Earth’s centre of mass	Greenwich meridian	(-)	Rotation axis
<b>Spacecraft</b>	Reference point on back plate	Forward (sym. axis)	(-)	Main fin (nadir)
<b>Body</b>	Satellite centre of mass	Forward (sym. axis)	(-)	Main fin (nadir)
<b>Orbital</b>	Satellite centre of mass	Velocity (E surface)	Ang. moment.	(-)
<b>Geodetic</b>	Satellite centre of mass	North	East	Down

the x-axis. This ensures that the new reference frame is orthogonal, while the z-axis still points more or less down. Which in turn ensures the z-axis stays in the same plane relative to, in this case, the body frame x- and z-axes. After making the z-axis, the y-axis is found by using the cross product to complete a right handed reference frame.

### 3.3. Controller design

In section 2.3 the theory of a conventional linear controller as well as the theory behind NDI and INDI was discussed. Using those theories three controllers have been implemented in MATLAB<sup>™</sup> to control the spacecraft's attitude. How the controller fits into the overall model was shown in Figure 3.2. However, how the controller is exactly implemented from theory was not discussed yet, that will be done here. First the implementation of the P(I)D-controller will be discussed, then the LINDI-controller, which is the conventional INDI-controller, and finally the NINDI-controller, which is also an INDI-controller, but with an NDI outer loop. The figures in this section are meant to show how the controllers fit into the MATLAB<sup>™</sup> code, as such the controller lay-out may be different than the more conventional ways to display a controller lay-out, e.g. Figure 2.8 and Figure 2.9. It would have been nice to also be able to use the actual GOCE controller, but unfortunately the internals of the actual GOCE controller have to this author's knowledge not been publicised.

#### 3.3.1. PD-controller

The P(I)D-controller is implemented as shown in Figure 3.4. It can be seen that it does not include an integral term, hence it is actually a PD-controller. The PD-controller will be occasionally used throughout the report as a means of comparing the NINDI-controller against a conventional controller. The PD-controller is implemented as a cascaded controller, meaning that the outer loop has the attitude error as an input, and produces an angular rate to track. This commanded angular rate is then compared with the actual angular rate, resulting in a rate error which is the input for the inner loop. The output of the inner loop is then a required torque command. As was discussed in subsection 3.2.1, and shown in Figure 3.2, the required torque from the controller first goes through the time dependent control block, producing the actual commanded torque, which is then converted to a current. To more clearly show how Figure 3.4 fits into Figure 3.2, this time dependent control block is shown outside of the controller block. To close the feedback control loop, the dotted line is drawn from the 'Time Dependent Control' block, to the 'State propagation' block. The line is shown dotted, because other processes happen in between these two block, as shown in Figure 3.2.

Note that although the PD-controller was set-up as shown in Figure 3.4 it is not necessarily the case that a PD-controller is the most optimal controller for this problem. Though to be fair, PD-controllers are widely used in industry and it is therefore a solid choice for a reference controller. Furthermore, in section 3.4 it is discussed in depth how the gains of the controllers are tuned, thus for now it is only necessary to realise that the choice of which gains were used in the PD-controller came from the design of the NINDI-controller. The NINDI-controller itself was based on recent studies into INDI-based control architectures in which proportional and derivative gains were suggested. The code that was initially used to tune the NINDI-controller was also used to tune the PD-controller. To use the same code meant to have the same number of gains in the PD-controller as in the NINDI-controller. This thesis does not aim to further investigate whether including an integral term would improve the performance.

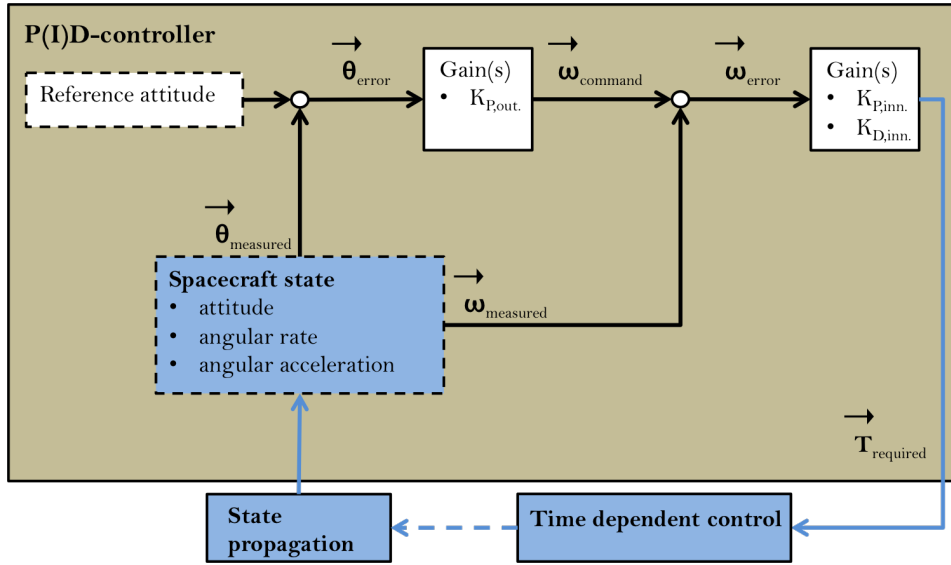


Figure 3.4: Lay-out of the reference P(I)D-controller, as implemented in MATLAB.

### 3.3.2. LINDI-controller

The LINDI-controller is a conventional INDI-controller. Meaning that the outer-loop sets the reference angular rate to track directly, by using a normal linear gain. Hence the name linear-INDI, i.e. LINDI. The implementation of the LINDI controller is based on the work by Acquatella et al. [1], also shown in section 2.3. How the controller is implemented in MATLAB<sup>TM</sup> is shown in Figure 3.5. The outer loop commands a rate, which is added to the commanded rate of the previous iteration in the sense of a feedforward term. These two together then produce the virtual control input for the middle loop. This gets subtracted by the measured angular rate, which then gives the rate error that is used in the inner loop. The inner loop produces the virtual control input for the INDI part of the controller. In this part, the measured angular acceleration is subtracted from the virtual control input, or the desired angular acceleration. This leads to a difference,

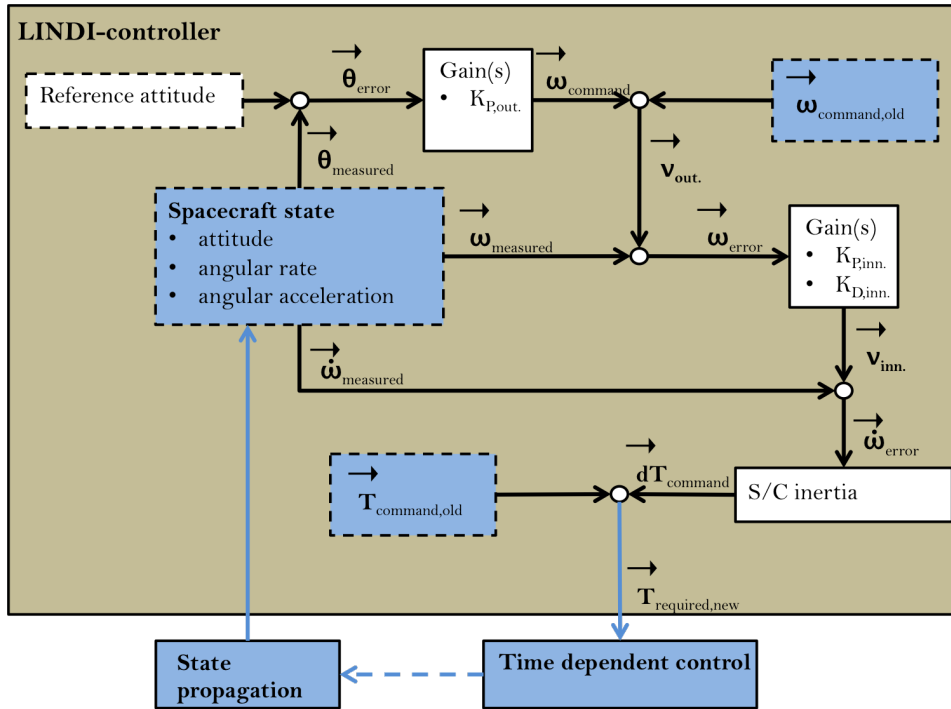


Figure 3.5: Lay-out of how the LINDI controller is implemented in MATLAB

or error, in angular acceleration, which is then inverted using the spacecraft (S/C) inertia matrix to produce a torque increment that steers the angular acceleration error to zero. The torque increment is added to the commanded torque of the previous iteration to produce the total commanded torque of the current iteration. Again, in Figure 3.5, the commanded torque is shown to go to the time dependent control block outside of the controller to highlight how the controller fits into the simulation environment shown in Figure 3.2.

According to theory an INDI controller with NDI in the outer loop instead of the more conventional linear outer loop is more promising. However, before this more promising controller could be made, the LINDI-controller was made first to gain experience with INDI-control architectures. The LINDI-controller is also included in this section, as well as in some parts of chapter 4, where it is compared against the NINDI controller. However, due to time constraints a complete evaluation whether the NINDI-controller is indeed better than the LINDI-controller could not be made.

### 3.3.3. NINDI-controller

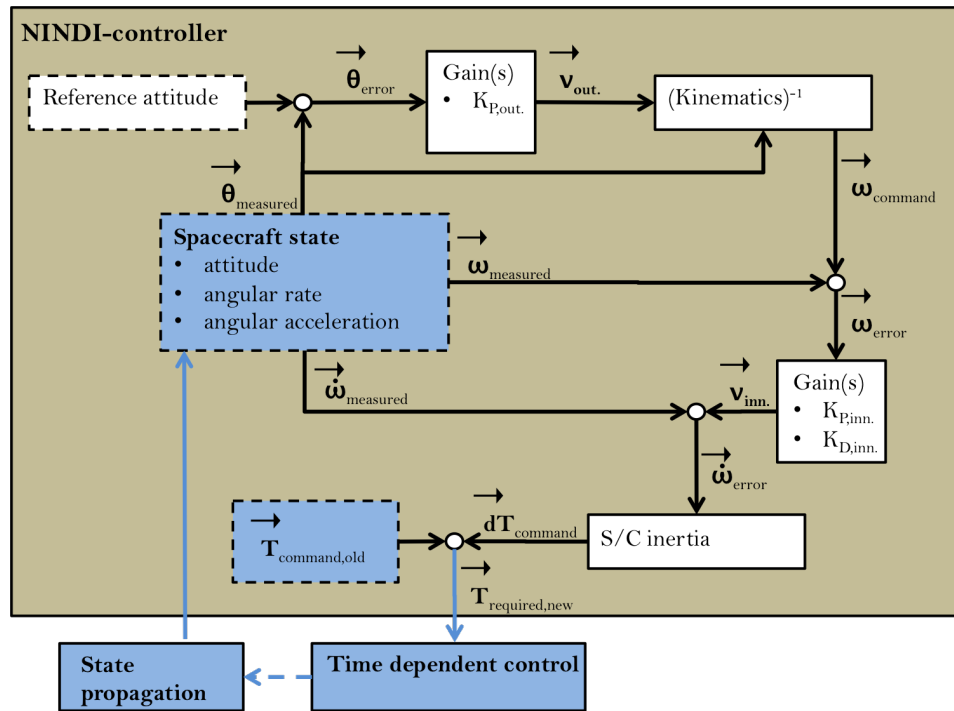


Figure 3.6: Lay-out of how the INDI controller with an NDI outer loop (NINDI-controller) is implemented in MATLAB.

Finally the INDI-controller with the NDI outer loop, or NINDI-controller, is shown in Figure 3.6. The NINDI-controller is the main controller under investigation in this thesis research. The biggest difference between this controller and Figure 3.5 is of course the inclusion of the NDI part in the outer loop. This controller is also based on the work done by Acquatella et al. [1], and the inversion of the outer loop is done by inverting the kinematic equations of motions, shown in Equation 2.7. The difference in the outer loop is that in the case of the NINDI-controller the error in attitude produces a virtual input to be used in the inverse of the kinematic EOM. This virtual input thus resembles a desired rate of change of the Euler angles,  $\dot{\theta}$  in Equation 2.7. This desired change in Euler angles is then used to produce a reference angular rate that would produce this change in Euler angles. The rest of the NINDI-controller is the same as the LINDI-controller, shown in Figure 3.5.

As mentioned in section 2.3 using NDI to invert the kinematics means that the outer loop equations are also linearised. Since the kinematic EOM are a given geometrical representation, they are free of uncertainties. This in turn means that the accurate model knowledge that is required for NDI is present, thus the reference angular rate that is produced can be expected to be less noisy than for the LINDI-controller, and thus easier to track. Additionally based on NDI theory one would expect the outer loop to behave similar in different operating conditions, i.e. it should behave similar at the start of the mission and at the end.

### 3.4. Gain tuning

The conventional way of tuning a controller is through setting up the transfer functions and finding the gains from there. However, given the extreme complexity and non-linearity of this control system a different method was necessary. The use of a relatively new controller type additionally meant that literature was not able to provide a straightforward way for tuning the controller. Tuning the gains was therefore done by selecting a specific day from the available data set, and choosing an interesting period during that day. The day of choice in this case was 1 November 2011, because there is some atmospheric activity on that day, without there being a storm. The period of choice was from the tenth orbit of the spacecraft on that day to the sixteenth. The reason this period was chosen was because controller activity was seen to increase, peak, and decrease again in that period. The method of tuning used was a genetic algorithm (GA). The GA is loosely based on how evolution works through natural selection. This thesis is not the right place to dive deep in the theory behind genetic algorithms, so instead the implementation of a GA that was used in this research is discussed.

The evolutionary part through natural selection in a GA works by running a simulation of every individual in a population of every generation. An individual in this case is a set of ones and zeros, bits, that represent the traits of an individual. A population is all individuals in one generation, and a generation is simply a specific population set. After each simulation of every individual, its performance is given a number, which is its fitness. After all individuals have been simulated, the next generation is made. There are many different ways to make a new generation [12, 13, 24, 43], but in this thesis three methods are used. The first is simply that a certain percentage of the individuals with the best fitness score go into the next generation without any modification. Then a certain percentage of the next generation will be filled with ‘children’, i.e. a new bit string will be made from the bit strings of two individuals, or parents, of the current generation. Again, there are different methods for choosing parents, but in this implementation of a GA the parents are selected semi-random. The reason why it is semi-random is because individuals with a better fitness score are slightly more likely to be selected as parents. The third method for creating the next generation that was used in this research is by selecting the individual with the single best fitness score as one parent, and the other parent is a newly generated individual that is thrown away afterwards.

Based on the above-mentioned it can be concluded that the fitness score has a lot of influence on how a next generation is made. This does make sense of course, because the end goal is to optimise something and therefore a certain performance metric will need to be either maximised or minimised. In this case that performance metric is the fitness score. It does however also pose a challenge, because there is no guarantee that an individual that has a good fitness score is close to a global optimum, it could instead also be near a local one. This is however one of the strongest points of GAs; not getting stuck in a local optimum [12, 24, 43]. As was mentioned, the next generation is filled with a certain percentage of the best performers, and the rest of the next generation is made by combining bit strings from individuals in the current generation, or newly generated outsiders. The reason for this is that it allows to mix bits, which can be considered analogues to genes in evolution, giving rise to new individuals that might perform better than individuals from the previous generation. However, the question still remains how ‘better’ exactly is defined. Better in this case, is captured in the fitness score, which is calculated using a fitness function. Since a GA is driven by its fitness function, thought has to be given to how the performance of an individual is captured in a number.

$$\text{error} = \text{reference} - \text{actual}$$

$$e_{\text{total}} = e_{\text{total}} + \sqrt{\left(\frac{\text{error}}{\text{requirement}}\right)^2} \quad (3.1)$$

$$F = W_{\theta} \cdot \|\theta_{e,\text{total}}\| + W_{\omega} \cdot \|\omega_{e,\text{total}}\| + W_{\dot{\omega}} \cdot \|\dot{\omega}_{e,\text{total}}\| \quad (3.2)$$

In this thesis the fitness function is defined as the sum of the norm of the normalised total error of each component in the state, as shown in Equation 3.1 and Equation 3.2. At every iteration during the generation the error in the attitude angle, angular rate, and angular acceleration is calculated with respect to their reference values. If the errors oscillate around zero, simply summing the errors could return a value close to zero, while the oscillation could have a large undesirable amplitude. Since this would falsely imply that the error is small, it was chosen to use the absolute value of the error in each iteration and add it to the previous value of the error. This is then called the total error. However, there are large order of magnitude differences between the errors in the attitude angle, angular rate, and angular acceleration. In order to sum them in a way that their




Figure 3.7: Anatomy of a bit string.

values are comparable they need to be normalised first. The normalisation of the errors is done with respect to the ESA requirements, shown in Table 4.2. The norm of the vector of each error component of the state is taken before they are summed together in a fitness score, so that the fitness score is represented by one single value.

Equation 3.2 also contains weights. These weights allow to increase or decrease how much influence the total error of each of the components of the state has on the fitness function. In other words, if the outcome of the genetic algorithm is not desirable because for example the errors for the angular rate are still too large, then the fitness function can be tuned to be more sensitive to the angular rate errors by increasing the weight,  $W_\omega$ . In this research however, the weights are all equal to 1. Similarly, although not explicitly shown in Equation 3.2, it is also possible to put more weight on the total error of an axis. This was however also deemed unnecessary to achieve desirable results.

Now that it is clear how the performance of an individual is calculated, an overview will be given of how the bit strings lead to the traits of an individual, and which traits are chosen to be represented by the bit string. First, that what needs to be captured in the bit strings is the values of the gains. This then allows to find a set of gains that optimises for the fitness function. To do this it must be known which gains are to be taken into account and which are not. In section 3.3 the design of the controller was discussed. The controllers all have the same gains active. Based on the controller design it can be concluded that nine parameters have to be captured in the bit string. These parameters are the outer loop proportional gains, 1 for each axis, and the inner loop proportional and derivative gains, 2 for each axis. However, as is shown in Table 4.1, the gains also include a factor that they are pre-multiplied with. So three additional parameters need to be included in the bit string, bringing the total to 12 parameters. There is a good reason for including these factors and that is that they vastly increase the search space for the gains. The factors are integers, so they only increase or decrease the order of magnitude of the gains, while the gains themselves take on a value between a set lower bound (LB) and upper bound (UB). Thus by incorporating the factors in the GA it is possible to increase the search space for suitable gains by orders of magnitude, without the need for very long bit strings. This simultaneously also enables the search space for the gains to have the same resolution for all order of magnitudes. In order to transform the bit string into these factors and gains, it is assumed to be known which positions in the bit string represent a factor and a gain, and how long these representations are. An example is shown in Figure 3.7, note that in this figure the length of the factors and gains are an example, and not representative of the actual implementation.

Converting a piece of the bit string into a number is not as simple as putting it into a 'binary2decimal' converter, since this only outputs integers. To overcome this problem a custom binary to decimal converter was made that does allow for decimal values, and also allows to specify minimum and maximum values. This custom binary2decimal converter is shown in Equation 3.3, and its based on [12]. In Equation 3.3, 'size' refers to the size of the bit string of the parameter that is to be converted. This then gets multiplied with either a 0 or a 1. This part of Equation 3.3 is similar to a conventional binary to decimal converter. However, left of the summation sign is the variable  $\Delta$ , calculated as shown in Equation 3.4. The  $\Delta$  variable allows the custom binary to decimal function to output decimal values. Lastly the minimum value is added to the summed part to ensure the decimal output matches the desired minimum value. Equation 3.4 also shows that the variable  $\Delta$  gets smaller as the number of bits for a parameter gets larger. In Equation 3.4, UB and LB are the desired maximum and minimum value respectively. So by increasing the number of bits for a parameter, the resolution of the available outputs increases. Or, to put it in other words, by increasing the number of bits, the significant number after the decimal point increases.

$$\text{decimal} = \text{LB} + \Delta \sum_{i=x}^{\text{size}} b_i 2^{\text{size}-i} \quad (3.3)$$

**P1:** 10111 1010010010 ..... 101 11111110 0010  
**P2:** 00000 1001010110 ..... 011 010011110100  
**C :** 10111 1001010110 ..... 011 11111110 0010

Figure 3.8: How children are made.

$$\Delta = \frac{UB - LB}{2^{\text{size}} - 1} \quad (3.4)$$

Lastly, it will be explained how the bit strings of two individuals in the current generation are used to create a new individual for the next generation. The implementation for doing this that was used in this thesis research is shown in Figure 3.8. The process is started at the first bit, it is then semi-randomly decided how much bits are to be taken. In the example shown in Figure 3.8 the first 5 bits are taken, then it is decided randomly from which of the two parents the bits will be taken. Again, in the example shown in Figure 3.8 the first five bits are taken from parent 1, *P1*. The child, *C*, is then given these five bits and the process starts over, beginning from the sixth bit. The length of the bits given to the child each time can vary, and each time it is randomly chosen which of the two parents the bits are coming from. It is possible that the child gets bits from one parent three times in a row, but since the choice for which parent is the donating party is random, it is highly unlikely that the child will be an exact copy of the parent.





# 4

## RESULTS AND DISCUSSION

In this chapter the results as produced by the simulations are presented. However, before presenting the results, a verification of the model output will be done. This is discussed in section 4.1. Then the results themselves are presented in section 4.2, and thereafter a discussion on the obtained results is done in section 4.3. All the results shown in this chapter have been produced by simulating one of three controllers. The set-up of these controllers has been discussed in section 3.3, the gains that are used in them are shown in Table 4.1.

Table 4.1: Value of the gains in the different controllers used.

			P(I)D	LINDI	NINDI
Outer					
	Factor <sub>p</sub>		1·10 <sup>-2</sup>	1·10 <sup>-3</sup>	1·10 <sup>-4</sup>
	K <sub>p</sub>	Roll	4.2754	0.9771	4.9962
		Pitch	3.0729	9.4433	9.6177
		Yaw	0.0167	0.4521	6.7243
Inner					
	Factor <sub>p</sub>		1·10 <sup>0</sup>	1·10 <sup>-1</sup>	1·10 <sup>-1</sup>
	K <sub>p</sub>	Roll	4.8004	7.4536	8.7036
		Pitch	9.7116	5.1541	6.1736
		Yaw	2.3137	0.0868	0.0643
	Factor <sub>d</sub>		1·10 <sup>-3</sup>	1·10 <sup>-3</sup>	1·10 <sup>-6</sup>
	K <sub>d</sub>	Roll	3.9595	8.6420	8.9128
		Pitch	5.9413	4.3870	7.9518
		Yaw	1.2552	0.2637	7.7017

### 4.1. Model output verification

In this section the produced output by the model will be verified. The verification of the model output is split up in a section verifying whether the reference frame work as expected, and then a section on whether the torque models yield realistic results. All the verification runs are done for the complete day of 1 November 2011. This day is used for the verification because the gains were optimised by the GA for a period during this day. It can then immediately be seen how the found gains behave outside the specific period they were optimised for. As was mentioned in section 3.4, the optimisation of the gains was done in the time period from the tenth orbit, to the sixteenth orbit. The verification of the output is done over the complete day of 1 November 2011.

#### 4.1.1. Reference frames

As discussed in subsection 3.2.3, there are multiple control frames defined, so the verification will be done for the three control frames; the 'O-frame', the 'L-frame', and the 'D-frame'. The model outputs shown here




Figure 4.1: Geomagnetic indices for the complete day of 1 November 2011.

will only encompass the NINDI-controller, since this is the main controller type under review in this thesis research. The gains that have been found by the GA for each controller are shown in Table 4.1. Each of the gains has a factor that they are pre-multiplied with, shown in bold in Table 4.1. The factors thus determine the order of magnitude of the gains themselves. Comparing the factors, i.e. the magnitude of the gains, already shows some of the differences that are crucial to the controllers. For example, the outer loop gains found for the PD-controller are significantly higher than the other two controllers, thus the PD-controller can be expected to more aggressively steer the spacecraft's attitude. The inner control loop gains show less of a difference, but are inevitably influenced by the output from the outer loop. It is thus difficult to say beforehand how this will affect controller performance. As an indication for expected atmospheric conditions, Figure 4.1 shows three geomagnetic indices that are indicative of thermospheric activity [8]. These indices are based on measurements that are done in magnetic observatories all over the Earth. However, since these observatories are unequally spread out, a mismatch can occur when a region of high geomagnetic activity is not in the neighbourhood of an observatory, which then incorrectly reports a too low value of geomagnetic activity [29]. However, since these indices do present a convenient indication for thermospheric activity, they will be used as such. Based on Figure 4.1, an increase in thermospheric activity can be expected, starting halfway through the day. It is likely that this will result in an increase of controller activity as well. The verification of the controller will be done for the complete day of 1 November 2011.

The verification of the reference frames will be split up over the three control reference frames that were discussed in subsection 3.2.3. A quick recap of what the reference frames are; the 'L-frame' is the local orbital frame with the x-axis pointing in the direction of the satellite inertial velocity, the z-axis is a radial down vector with respect to the Earth centre, and the y-axis completes a right-handed reference frame. The 'O-frame' is similar, but instead of the x-axis pointing in the inertial velocity direction, it points in the satellite velocity vector with respect to the surface of the Earth. The 'D-frame' is different, in that its x-axis direction depends on the direction that the vector sum of the disturbance forces is coming from. In other words, it points the satellite in the direction of the disturbance forces.

The verification output shown in Figure 4.2 to Figure 4.4 shows the simulated state in colour, the GOCE measured state in black, and the dotted red lines indicate the ESA requirements for GOCE in (a) to (c), (d) shows the disturbance torque in black and the achieved control torque in colour. The disturbance torque has been multiplied with  $(-1)$  to put it in phase with the achieved control torque. Looking at the performance of the NINDI-controller it is obvious that the controller has good performance when doing the simulation in the 'O-frame' and 'L-frame', visible in Figure 4.2 and Figure 4.3 respectively. The wave-like attitude on the yaw-axis, seen in Figure 4.3, is because of the 'L-frame' x-axis being defined as the inertial velocity of the satellite. Hence the yaw-axis attitude is offset from zero, because of the spacecraft's natural tendency to point in the direction of the incoming wind due to its aerodynamic properties. However, when looking at Figure 4.4 it can be seen that the controller performs very poorly, resulting in the simulation stopping due to too large errors.

The premise of the 'D-frame' is promising, steering the attitude in the direction of the linear (disturbance) acceleration. Therefore an attempt was made to make the 'D-frame' controller work. The attempts to resolve

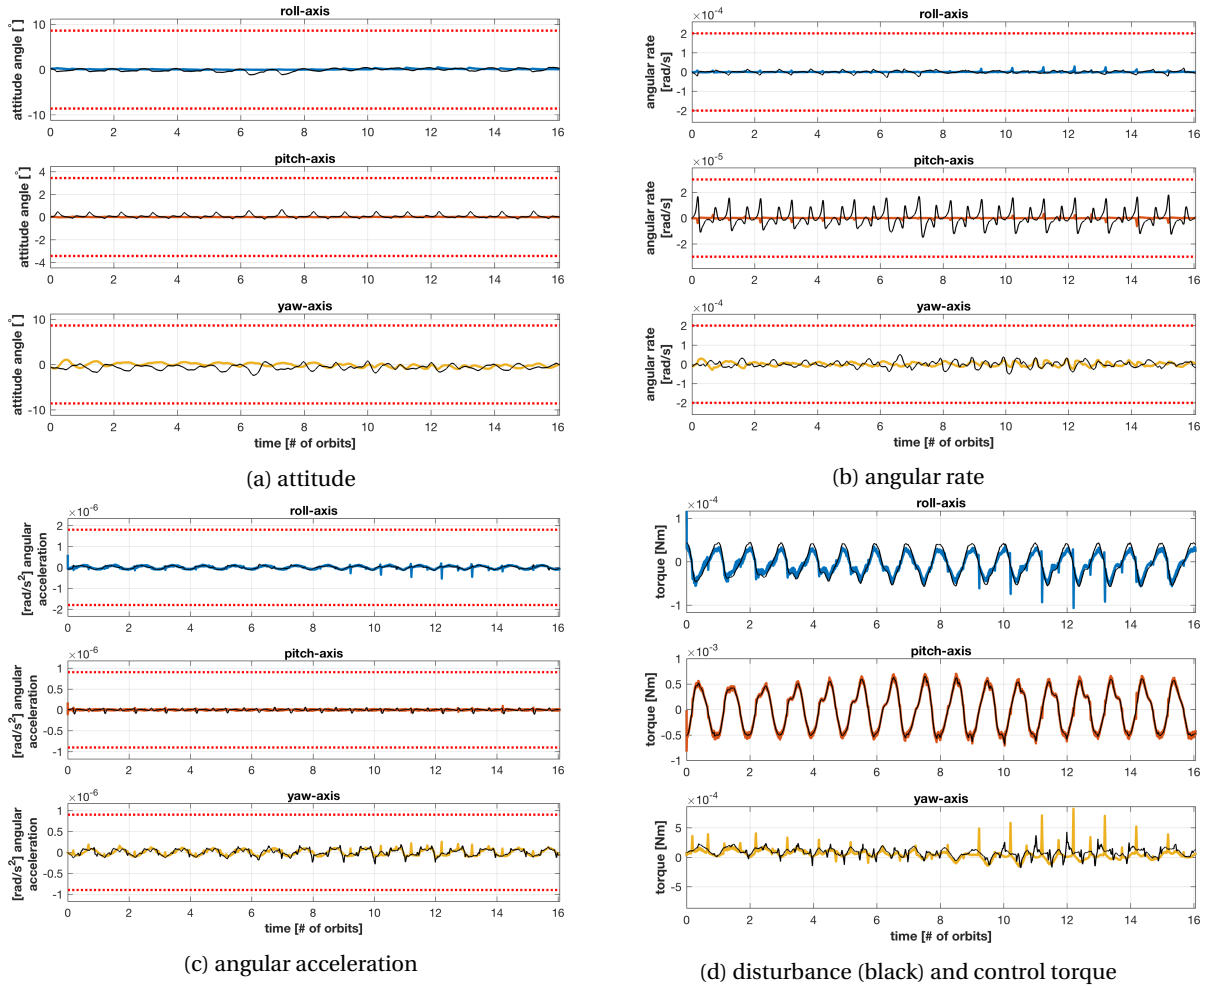


Figure 4.2: Simulated spacecraft attitude using the NINDI-controller, (a), angular rate, (b), angular acceleration, (c), and control torque versus the negative of the disturbance torque (d). Simulation is done in 'O-frame'. Except for (d), the black line shows the measured value of the GOCE satellite for comparison.

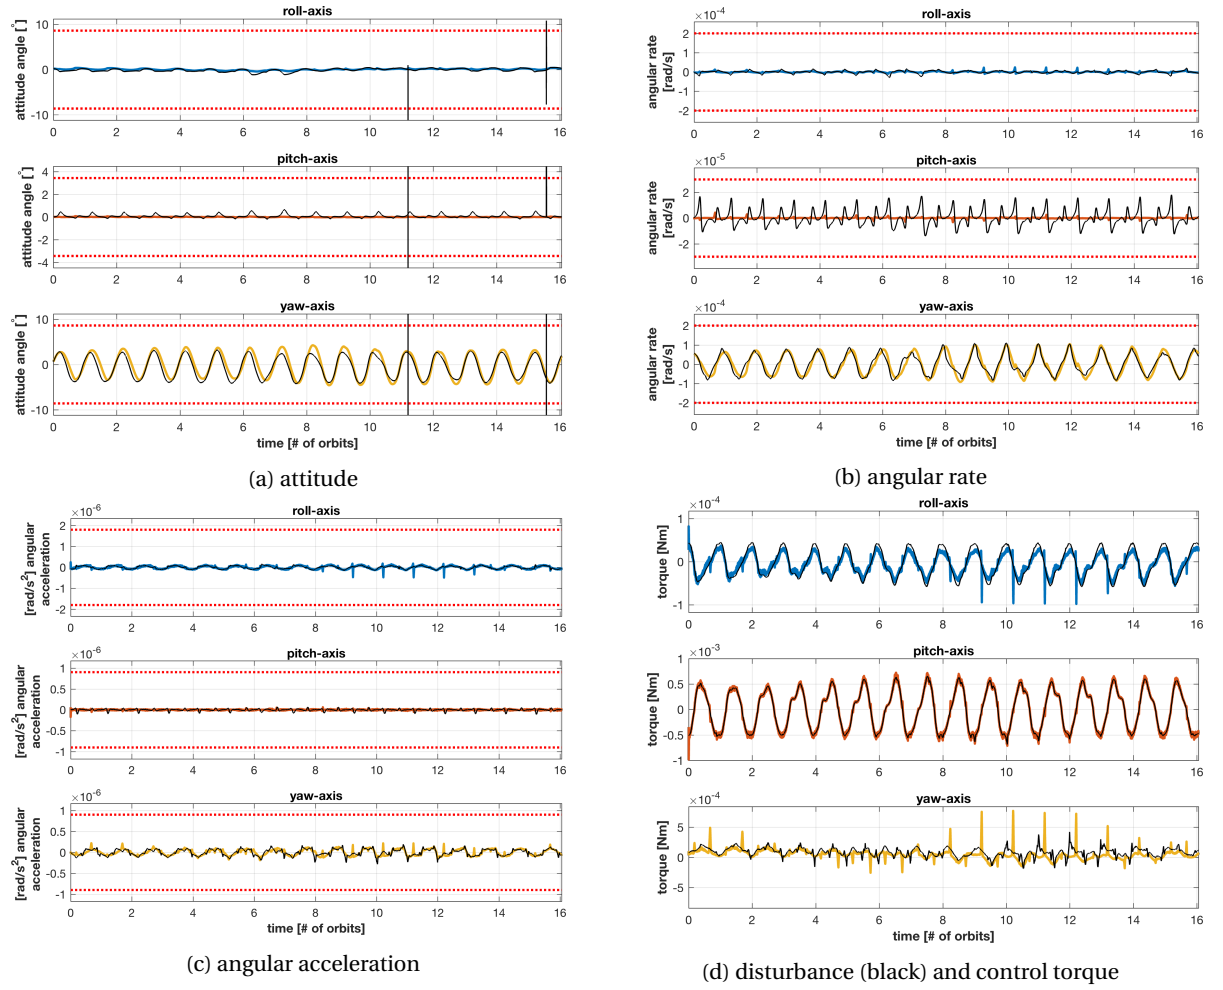


Figure 4.3: Simulated spacecraft attitude using the NINDI-controller, (a), angular rate, (b), angular acceleration, (c), and control torque versus the negative of the disturbance torque (d). Simulation is done in 'L-frame'. Except for (d), the black line shows the measured value of the GOCE satellite for comparison.

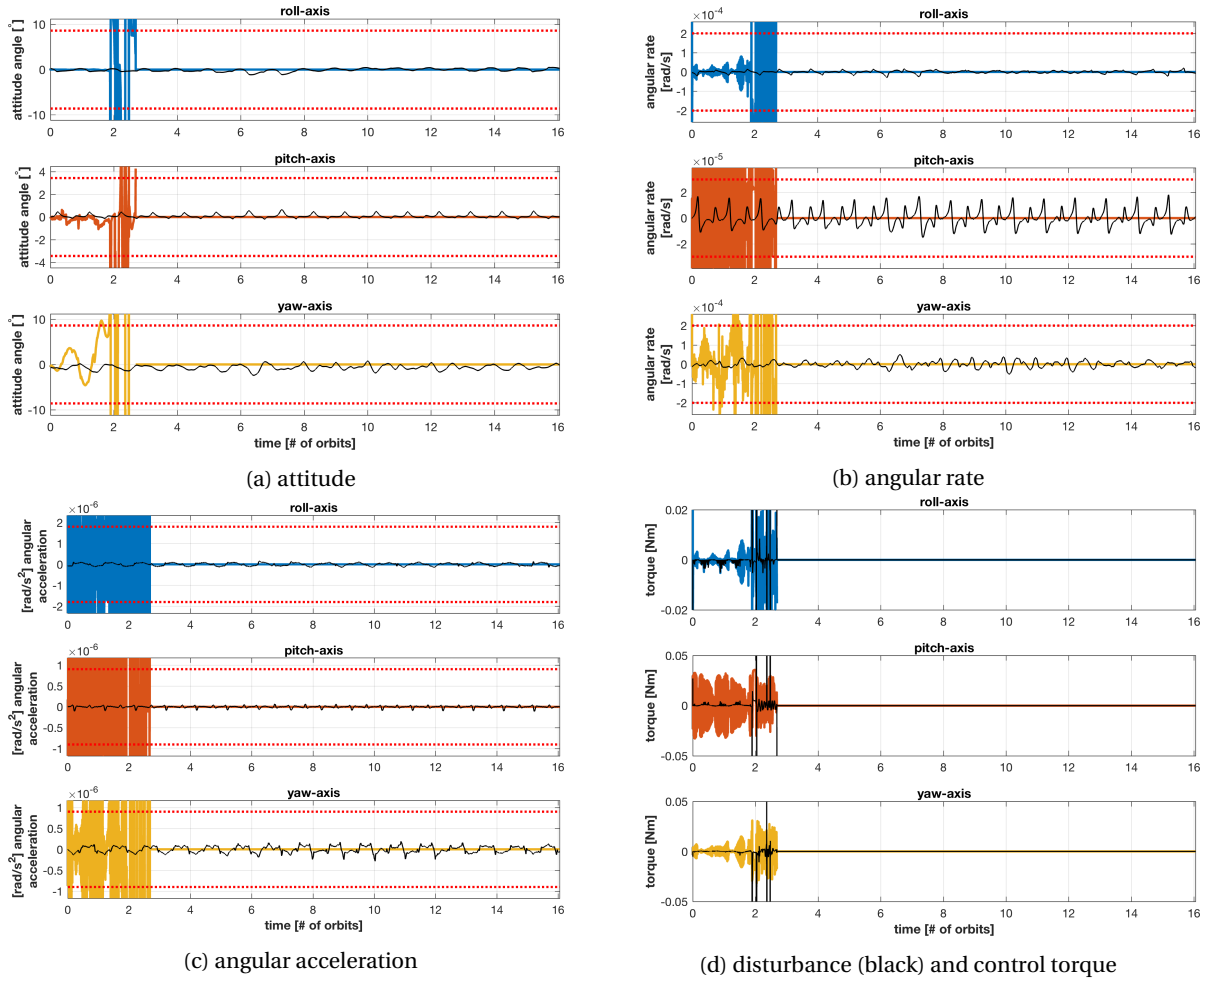


Figure 4.4: Simulated spacecraft attitude using the NINDI-controller, (a), angular rate, (b), angular acceleration, (c), and control torque versus the negative of the disturbance torque (d). Simulation is done in 'D-frame'. Except for (d), the black line shows the measured value of the GOCE satellite for comparison.




Figure 4.5: (a) Shows the forces in the body frame, as they occurred in the simulation of the NINDI-controller while simulating in the 'O-frame'. (b) Shows the forces in the body frame, as they occurred in the simulation of the NINDI-controller while simulating in the 'D-frame'.

the problems with the D-frame that have been tried include the following. Applying a moving average filter on the linear perturbing forces, and applying a moving average to the y-axis linear perturbing forces only. The idea behind applying moving averages was to smooth the reference attitude. The moving average was applied with different window sizes, 10 second windows up to windows equal to half the orbital period, but no window improved the performance significantly. Applying a moving average to the y-axis forces only was an attempt to filter out the wind noise, while maintaining the rest of the forces. However, this also did not improve stability. Due to a lack of time, further investigation in this anomaly was not possible. When investigation into this problem is picked up, one could start with removing the wind noise altogether, to see if that stabilises the control. It is also possible that it is simply a case of the controller not being appropriately tuned for control in the D-frame. Therefore running the GA to retune the gains specifically for control in this frame might help. Though this is deemed unlikely since the performance of the controller is very similar in the L-, and O-frame. Lastly, it cannot be ruled out that there simply is an error in the code that is used to define the D-frame.

To further investigate this claim, take a look at Figure 4.5. This figure shows the forces in the 'O-frame' as a reference, (a), and the forces in the 'D-frame', (b). It shows that the 'D-frame' starts promising, with similar forces on the x-axis, smaller forces on the z-axis, and slightly larger forces on the y-axis. Something then happens just before the second orbit is finished and control is lost. The timing of the start of the strange behaviour in forces corresponds to what is seen in Figure 4.4, but the reason why the forces go to zero afterwards is not known. However, the fact that they do go to zero, even though satellite parameters like velocity have not changed indicates there is an error in force modelling when using the D-frame. Possibly this only applies at large attitude angles, since the forces appear to be similar to the forces seen in the 'O-frame' simulation before attitude control is lost. This figure is thus a strong indication that a good starting point for getting the D-frame to work is to evaluate the code that leads to the definition of the D-frame for errors.

In conclusion, the model output for the two control reference frames, L and O, is now assumed to be verified. For the D-frame the conclusion is that even though attempts were made to solve the erratic behaviour, these were unsuccessful. Given that simulating attitude control in the D-frame clearly does not work, it will not be included in the results section. Furthermore, all the attitude simulations hereafter will be in the 'O-frame', because a  $0^\circ$  yaw axis attitude in the O-frame means the spacecraft is pointing in the direction of the, locally, co-rotating atmosphere meaning the thruster is aligned properly to counteract drag.

#### 4.1.2. Torques acting on spacecraft

That what ultimately drives the attitude of a spacecraft are the torques acting on it. In reality the external environment continuously exerts forces and torques on a spacecraft causing it to accelerate and rotate. Simultaneously, in the case of the GOCE spacecraft, the attitude control system discretely interacted with the EMF in order to keep the spacecraft's attitude within limits, while also minimising the angular rate and acceleration. In the case of the simulation that was set-up both the disturbance torques perturbing the spacecraft's

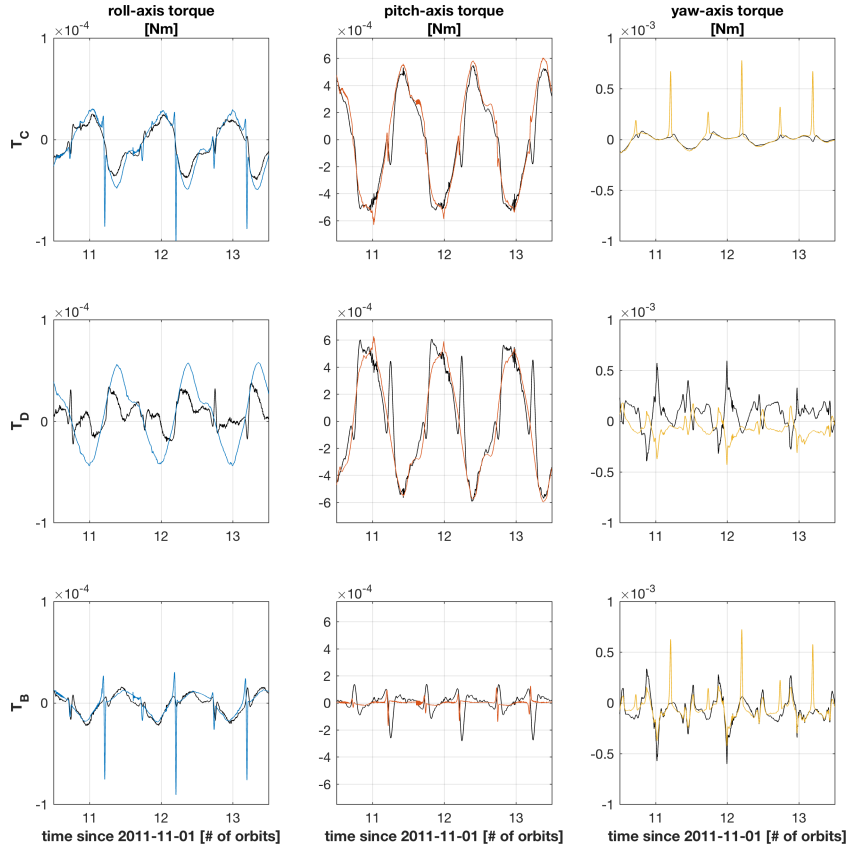


Figure 4.6: Measured (black) and simulated (colour) torque on the GOCE satellite, with the control torque,  $T_C$ , the total disturbance torque,  $T_D$ , and the net torque,  $T_B$ . Measured torques as modelled by [41].

rotational motion, and the control torque are discrete signals. And the perturbing torques are specified by models which attempt to mimic what happens in reality. Therefore it is interesting to see how the measured torques compare to the simulated ones. This comparison is first shown for a short time period in which the simulated controller was more active than on the rest of the day.

The torques in the comparison shown in Figure 4.6 are the control torque  $T_C$ , the total disturbance torque  $T_D$ , and the net torque on the body  $T_B$ . The net torque is the remaining torque after the controller has attempted to counteract the disturbance torque. As can be seen, the general behaviour of the control torque is similar for both the measured and the simulated case. However, in the simulated case all axes do show some spikes in the control torque, with the most prominent ones on the roll- and yaw-axis. The source of these spikes, and how they affect the controller performance is discussed in the next section. The pitch-axis disturbance torque is very similar for the measured and simulated case, in contrast to the roll- and yaw-axis torque in which there are differences. These differences are however within the same order of magnitude, and are therefore likely caused by attitude and control differences between the measured and simulated case. The behaviour in the net body torque is very similar again, apart from the spikes. In some components of the total disturbance torque there is a dependency on the attitude of the spacecraft, while in other components there is a dependency on the current to the magnetorquers, i.e. the control actuation. Furthermore, the internals of the actual GOCE attitude controller are unknown, and are likely different from the simulated controller. It is therefore considered explicable for the torques shown in Figure 4.6 to not be exactly equal.

The behaviour of the simulated torques for the complete day of 1 November 2011 is shown in Figure 4.7 and Figure 4.8. These figures show an envelope plot of the simulated torques. Thus the spikes that were visible in Figure 4.6, are wrapped inside the envelope, which is happening from around orbit 8. This does correspond nicely with the increase in thermospheric activity starting slightly earlier, shown in Figure 4.1. This increase in thermospheric activity is least visible on the pitch-axis, likely due to the requirements of the pitch-axis being most stringent, causing the pitch-axis to be weighted heavier in the GA. To complete the overview of the torques on 1 November 2011, Figure 4.8 shows the largest contributors of the total disturbance torque,

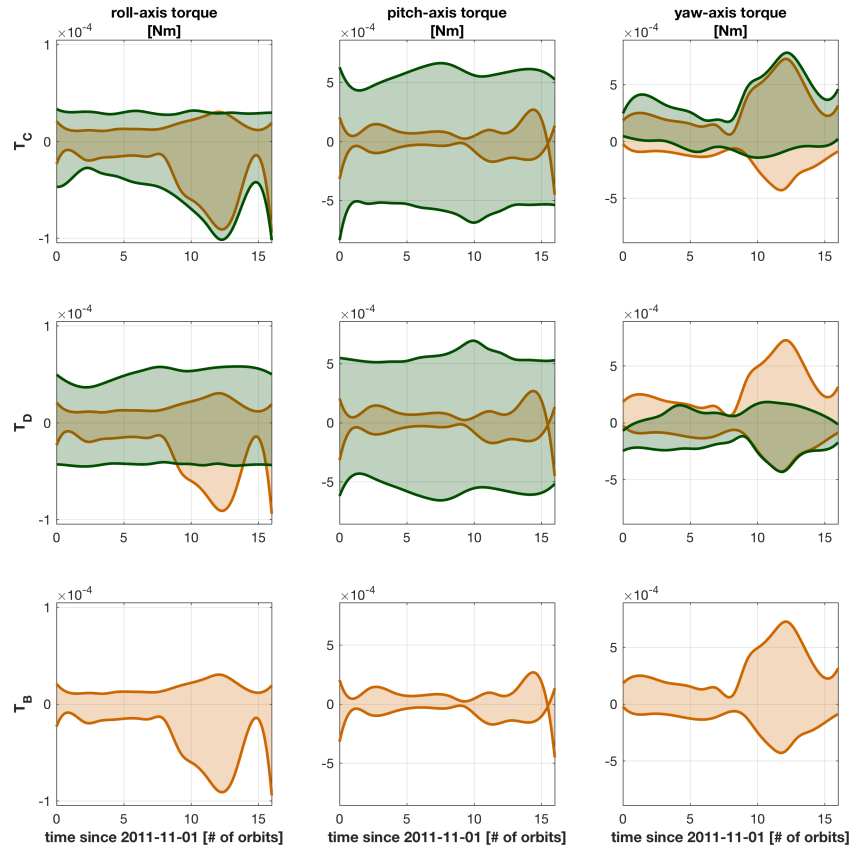


Figure 4.7: Envelope plot of the simulated torques acting on the GOCE spacecraft, with the control torque,  $T_C$ , the total disturbance torque,  $T_D$ , and the net torque,  $T_B$ . The net body torque on each axis is shown on all plots of that axis.

again in an envelope plot. The torques shown are the fitted dipole torque  $T_{FD}$ , the residual dipole torque  $T_{RD}$ , the aerodynamic torque  $T_{AE}$ , the torque due to the IPA  $T_{TM}$ , and the total disturbance torque itself. Figure 4.8 nicely shows that the yaw-axis disturbance torque is dominated by the aerodynamic torque on the spacecraft, while the roll-axis is dominated by magnetic effects, be it the fitted or the residual dipole. The pitch-axis disturbance torque is a combination of the residual dipole and the torque due to the IPA. Based on Figure 4.8 it appears that none of the largest four components of the total disturbance torque are coupled, which makes sense given their nature. An overview of all the modelled components of the disturbance torque is given in Figure C.1. However, note that this overview is for the period 15 – 21 March 2013.

Due to the shape of the GOCE spacecraft less control actuation is necessary on the yaw-axis. This in turn does also mean that the yaw-axis disturbance torque is more sensitive to the attitude of the spacecraft. The two torques on the yaw-axis that are highly dependent on the yaw-axis attitude angle are the solar radiation pressure torque and the aerodynamic torque. To show the relation between attitude and torque for these two disturbance the side-slip angles,  $\beta_{AE}$  and  $\beta_{SR}$ , together with the simulated aerodynamic and solar radiation pressure torque on the yaw-axis is shown in Figure 4.9. Additionally the measured solar radiation pressure torque and aerodynamic torque are also given. The dependency of the aerodynamic torque on the aerodynamic side slip angle is clearly visible in Figure 4.9. It is a little bit more difficult to see the same dependency for the solar radiation pressure torque, but a similar upward trend in this torque, corresponding to an increase in the solar radiation pressure side slip angle, can still be observed. The reason why the solar radiation pressure torque often jumps to zero is because the Sun is then eclipsed by the Earth. Looking at the y-axis scale of the torques, it can be seen that the solar radiation pressure torque is an order of magnitude smaller than the aerodynamic torque, this can also be seen in Figure C.1.

Now that the torques that are produced have also been compared with measured data, it can be concluded that the complete model is verified for use with both the L-frame and O-frame as a control reference frame. These two reference frames produce results that are within expected values, and the torques that were produced using the O-frame have been found to be similar to the observed torques. In the next section results



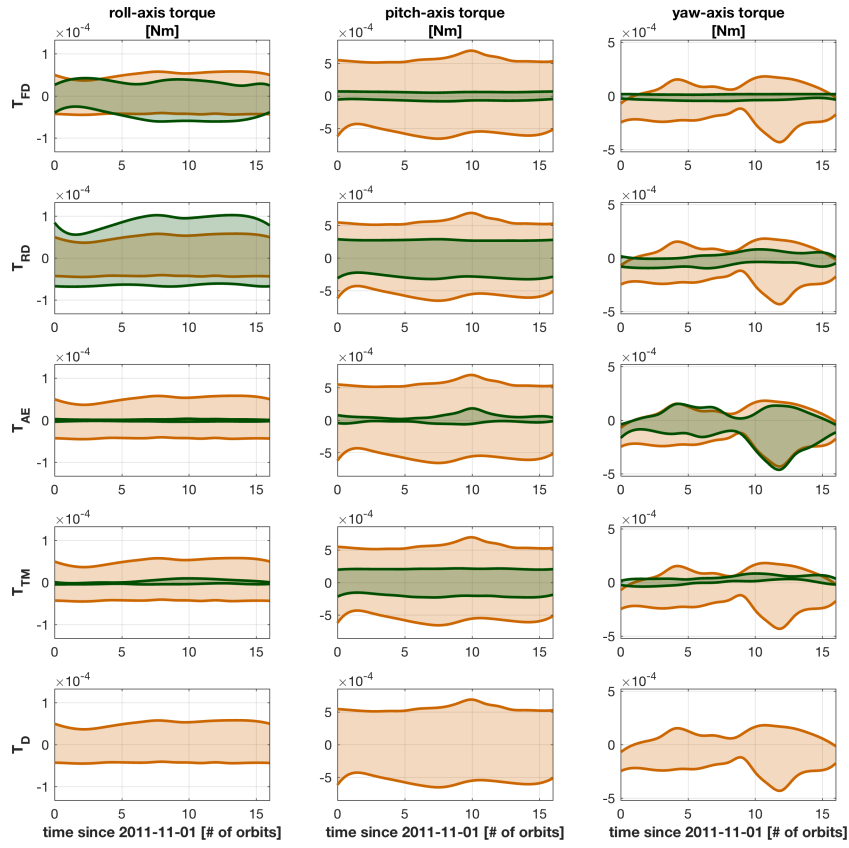


Figure 4.8: Envelope plot of the components of total disturbance torque, and total disturbance torque itself. Components shown here are the fitted dipole torque  $T_{FD}$ , the residual dipole torque  $T_{RD}$ , the aerodynamic torque  $T_{AE}$ , the torque due to the IPA  $T_{TM}$ , and the total disturbance torque,  $T_D$ .

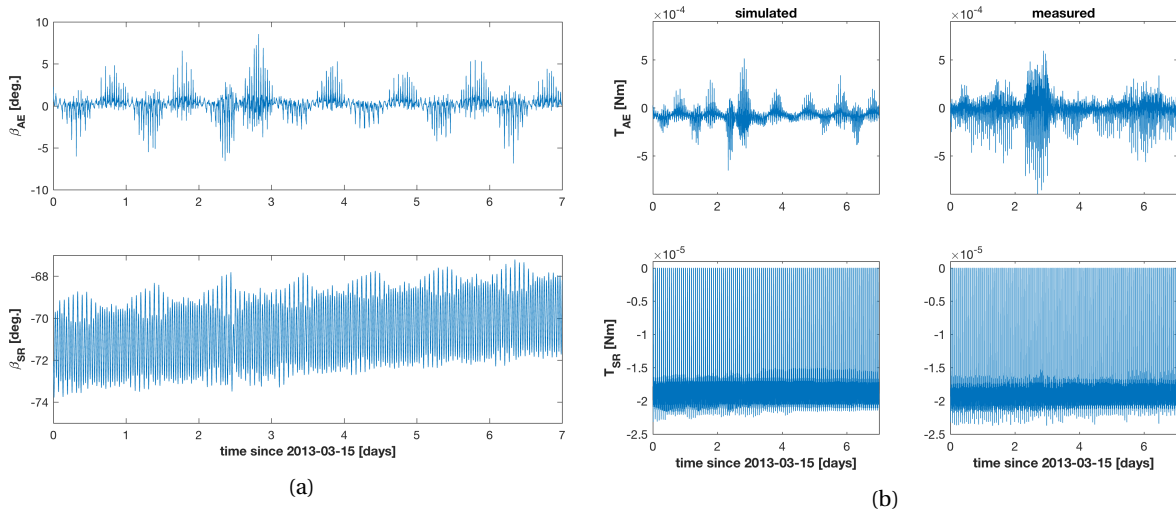


Figure 4.9: (a) Shows the simulated side slip angles for the aerodynamic torque and the solar radiation pressure torque in degrees. (b) shows the aerodynamic torque and the solar radiation pressure torque for the simulated case (first column) and the measured case (second column).

for different time periods will be produced in the O-frame, because for the GOCE mission it makes more sense to have a  $0^\circ$  yaw axis attitude mean that the spacecraft is pointing in the direction of the largest contributor to the linear forces.

## 4.2. Results

The results that will be presented in this section will be split up in the results evaluating the performance of the NINDI-controller in terms of meeting the requirements in subsection 4.2.1. Then how the controller affects the spacecraft will be compared with the measured behaviour of the spacecraft in subsection 4.2.2. The performance of the controller when errors and uncertainties are introduced is evaluated in subsection 4.2.3. Finally the sensitivity of the controller to changes in the gains is analysed in subsection 4.2.4. Different time spans may be used for the different results shown, but the controller settings (gains) are the same for all simulations. The gains that were used for producing the results in this section are shown in Table 4.1.

### 4.2.1. Performance of the NINDI controller

The period of 15 – 21 March 2013 is chosen to evaluate the performance of the NINDI-controller. Performance, in this case, is measured with respect to the requirements as set by ESA before the mission started. This period is chosen because at this time the satellite flew at a low altitude of  $\pm 245$  kilometres above the surface. This combined with increased geomagnetic activity after the second day, should result in an increase of thermospheric activity and thus requiring more controller effort to keep the satellite state stable. See Figure 4.10 for an indication of thermospheric activity in the specified period. Note how the behaviour of the values for  $K_p$  and  $A_p$  is different than what was seen in Figure 4.1. They are supposed to follow the block line shape, but they appear to have been interpolated in the provided data for Figure 4.10. What is important to see in this figure is that in between the second and third day all three indices rise, indicating that a storm in the upper atmosphere can be expected.

Before moving on to actually evaluating the results for the specified period it should be noted that it does not matter that much which period is chosen. Appendix A can be used to verify that, at least the LINDI- and NINDI-controller perform very similar in different periods, while the PD-controller shows itself to be not so flexible. This nicely lines up with the theory about PID-controllers and INDI-controllers. Lastly, Appendix A also shows how the measured GOCE values relate to the ESA requirements in the same periods as the controller. As can be seen the LINDI- and NINDI-controller perform similarly or better in this regard.

Moving on to the results for the specified period. It was already mentioned that the performance will be expressed in terms of the requirements that were set by ESA before the mission was launched. These requirements are shown in Table 4.2. The shown requirements are for the so-called ‘Drag Free Mode’ (DFM), which was the operational mode in which the scientific data was gathered. For the controller to be considered successful it needs to at the very least meet these requirements. Note that Table 4.2 shows requirements in both the time domain, and the frequency domain. Figure 4.11 shows a time series plot over the mentioned period. In this figure the attitude angle and the angular rate of the body axes with respect to the control frame,




Figure 4.10: Geomagnetic indices for the period of 13 – 21 November 2013.

Table 4.2: GOCE DFAC requirements for scientific observation phase. Attitude angle has been converted to degrees from original source. [32]

	Rotations Axis	Max Value		Max one-sided spectral density inside [5,100] mHz	
<b>Attitude w.r.t. Local Orbital Reference Frame</b>	Roll	8.60	°		
	Pitch	3.45	°		
	Yaw	8.60	°		
<b>Angular rate w.r.t. Local Orbital Reference Frame</b>	Roll	$2 \cdot 10^{-4}$	rad/s	$1 \cdot 10^{-6}$	rad/s/Hz <sup>1/2</sup>
	Pitch	$0.3 \cdot 10^{-4}$	rad/s	$0.5 \cdot 10^{-6}$	rad/s/Hz <sup>1/2</sup>
	Yaw	$2 \cdot 10^{-4}$	rad/s	$1 \cdot 10^{-6}$	rad/s/Hz <sup>1/2</sup>
<b>Angular acceleration</b>	Roll	$1.8 \cdot 10^{-6}$	rad/s <sup>2</sup>	$9 \cdot 10^{-8}$	rad/s/Hz <sup>1/2</sup>
	Pitch	$0.9 \cdot 10^{-6}$	rad/s <sup>2</sup>	$6.3 \cdot 10^{-8}$	rad/s/Hz <sup>1/2</sup>
	Yaw	$0.9 \cdot 10^{-6}$	rad/s <sup>2</sup>	$6.3 \cdot 10^{-8}$	rad/s/Hz <sup>1/2</sup>
<b>Linear acceleration along track</b>		$0.9 \cdot 10^{-6}$	m/s <sup>2</sup>	$2.3 \cdot 10^{-8}$	rad/s/Hz <sup>1/2</sup>

and the angular acceleration of the body axes with respect to the inertial frame have been normalised to the ESA requirements for the time domain. The figures show how close the simulated spacecraft state is to the requirements, with 100% being equal to the requirements.

Based on Figure 4.10, it can be expected that after two days an increase in thermospheric activity should be visible in Figure 4.11. This is indeed the case between day two and day three it can be seen that there is an increase of activity on the yaw-axis, and also the roll-axis has a bump in the acceleration. When looking at Table 4.1 it can be seen that the gains on the yaw axis are about an order of magnitude smaller than on the other two axes. This is why the increase in atmospheric activity leads to a bump in the attitude, angular rate, and angular acceleration on the yaw-axis. The spacecraft's yaw-axis acts, in a sense, like a wind vane. The bump on the roll-axis is then caused by the magnetic coupling when trying to enforce control on the yaw-axis. Apparently small yaw-axis attitude angles in Figure 4.11 do not guarantee stable control, even though the control reference frame in Figure 4.11 is the O-frame. However, a 0° yaw-axis attitude in the O-frame means the spacecraft is pointing in the direction of the co-rotating atmosphere. The thermospheric winds are modelled on top of the co-rotating atmosphere and can be significant, especially during periods of increased thermospheric activity. These winds are the source of the behaviour seen between day two and day three. For example, during a day in which thermospheric activity was low, and the spacecraft was at a higher average altitude, the state is controlled more stably, this is shown in Figure A.1. Another observation that can be made from Figure 4.11 is that the pitch axis angular rate appears to be very active. However, looking at the scale of the y-axis it can be seen that the angular rate on the pitch-axis only increases to just over 20% of the requirement. Therefore, it is clear that the spacecraft state is well within the requirements and thus this NINDI-controller meets the ESA requirements in the time domain. For proof that the NINDI-controller is not only able to meet the requirements in this period, but also during other periods with the same gains, see Appendix A. Having met the requirements in the time domain is only half the story however. The controllers' performance in the frequency domain will also need to be evaluated.

Table 4.2 shows the specified requirements in the frequency domain for the angular rate and angular acceleration. The performance of the NINDI-controller in the frequency domain with respect to the ESA requirements is shown in Figure 4.12. This figure shows the one-sided power spectral density (PSD) of the simulated angular acceleration on each axis. The ESA requirements within the Measurement Bandwidth (MBW) is also shown as a red line. The goal is for the controller to keep the PSD of the angular accelerations below the red line within the MBW. The MBW itself was defined as being in between 5 and 100 mHz,  $[5 \leq f \leq 100 \text{ mHz}]$  [31]. It is obvious that the ESA requirements are not met in the frequency domain. Moreover, similar results are obtained for the angular rate requirements in the frequency domain. This is a disappointing result, especially since the other requirements are so comfortably met. The question then becomes, how is it possible that the time domain requirements are met so comfortably, while the requirements in the frequency domain are missed entirely.

To further investigate what might influence the performance of the NINDI-controller in the frequency domain it was chosen to do a couple of single day simulations on the day that the controller was optimised for, thus 1 November 2011. Although the actual GOCE controller hardware ran at a frequency of 2 Hz [32] and all the simulations are done with a step size of one second, i.e. 1 Hz updates, lowering the simulation step

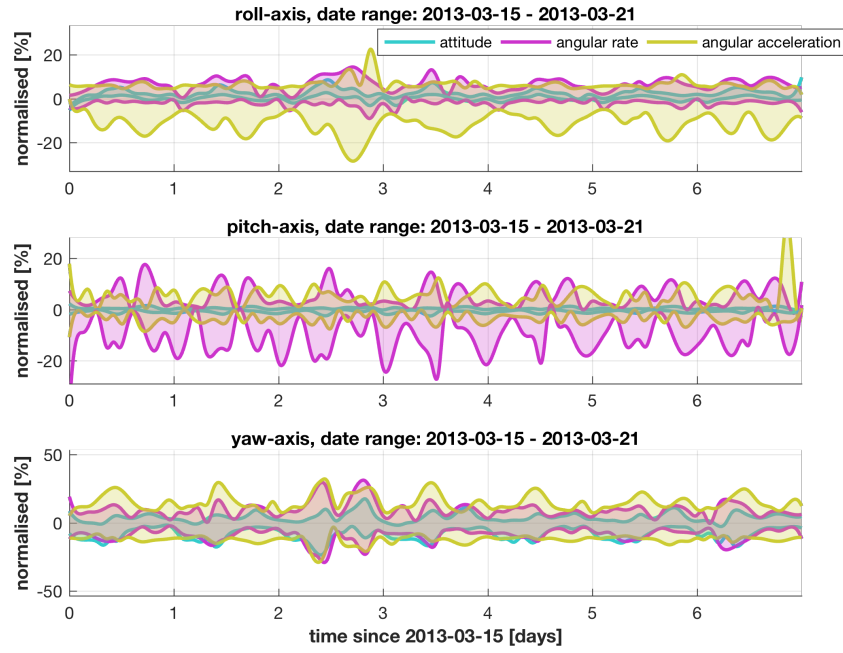


Figure 4.11: Envelope plot of the simulated attitude angle, angular rate, and angular acceleration of the GOCE satellite. All values have been normalised to the ESA requirements.

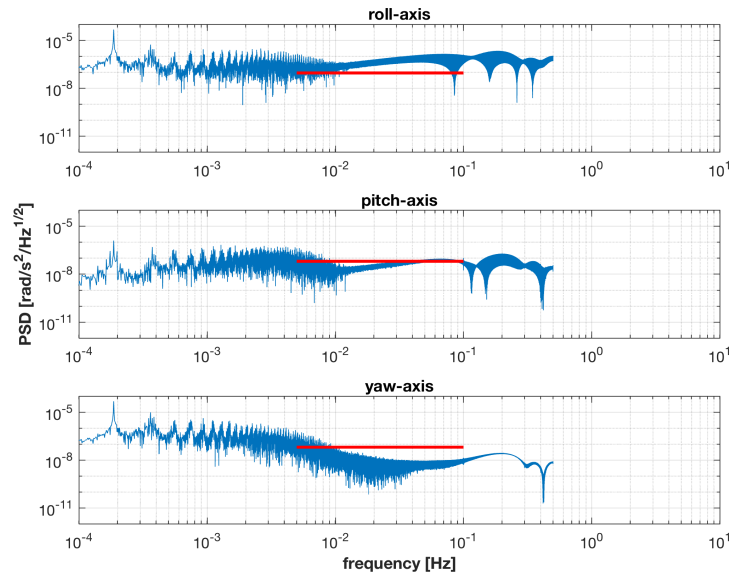


Figure 4.12: One-sided spectral density of the simulated angular accelerations over the period 15 – 21 March 2013. The red line indicates the ESA requirement for the maximum angular acceleration within the MBW.

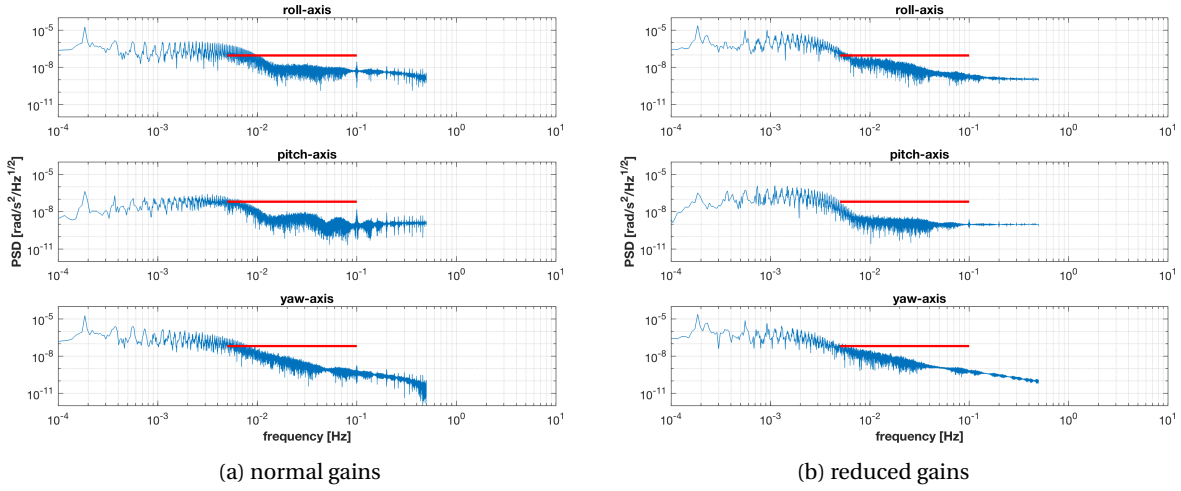


Figure 4.13: One-sided spectral density of the simulated angular accelerations during the complete day of 1 November 2011. The red line indicates the ESA requirement for the maximum angular acceleration within the MBW. (a) Shows the PSD using normal gains, while (b) shows the PSD of a simulation in which the inner loop gains of the NINDI-controller were reduced by a factor of 10.

size had no effect on the PSD performance. Remembering section 3.4, it is clear that the gains were optimised in the time domain. This is because the fitness function for the GA never took into account the performance of the controller in the frequency domain. If a simulation then is done with reduced gains for the inner loop, since those in the end lead to the commanded torque increment, it can be seen that this does improve the PSD performance. This is shown in Figure 4.13, where the gains for the inner loop of the NINDI-controller were reduced by a factor of 10.

First of all, even with the normal gains differences in the one-sided PSD can be observed between different periods (Figure 4.12 and Figure 4.13a). This can be explained by looking at the difference in geomagnetic activity for these two periods, as plotted in Figure 4.1 and Figure 4.10. During the simulation period in 2013 there is simply a lot more atmospheric activity, that needs to be compensated for. Since the NINDI-controller is somewhat aggressive in compensating the disturbance torques, this leads to larger magnitudes at the higher frequencies. Then Table 4.2 also shows frequency domain requirements for the angular rate. PSD plots for the angular rate have also been made, but similar results as for the angular acceleration were obtained, so these are not shown. Taking into account that Figure 4.13 are logarithmic plots, it can easily be seen that the difference between the two is actually quite large. This is a strong indication that not meeting the PSD requirements is caused by the controller gains being tuned too aggressive.

Considering that it seems likely that the tuning of the controller plays a major role in meeting the frequency domain requirements it is worth the effort to reflect on why this has such a large influence. The lower end of the MBW is 5 mHz, which corresponds to a period of 200 seconds. Exceeding the requirements at the lower end of the MBW means that that specific signal occurs roughly 27 times per orbit. This then indicates that it is unlikely that it has something to do with an Earth resonance, or other orbit related resonances. Meaning that it is most likely caused by the controller itself. Seeing that Figure 4.13 confirms this theory, it is necessary to consider what else changes when the gains of the inner loop are reduced. Figure 4.14 shows a small portion of the simulated angular acceleration on 1 November 2011. This figure shows that with the normal gains, the controller uses high frequencies to react quickly to changes. This is in contrast with how the measured angular acceleration behaves, shown in black. The spikes on the roll- and pitch-axis are most easily identified, and can be seen to have periods of 100 to 250 seconds. This corresponds to frequencies of 10 to 4 mHz which is where the MBW requirements are exceeded. In the figure with the reduced gains the measured angular acceleration is again shown in black. This emphasises that by reducing the gains of the inner control loop by a factor of 10, the spikes in acceleration have been made wider, i.e. the commanded control frequencies are lower.

Figure 4.14 is further confirmation that it is indeed most likely that exceeding the frequency domain requirements is caused by the controller commanding using high frequencies for control actuation. The reason why it causes the requirements to be exceeded frequency domain and not in the time domain, is because in tuning the controller no attempt was made to optimise for the PSD values. In section 3.4 it was explained

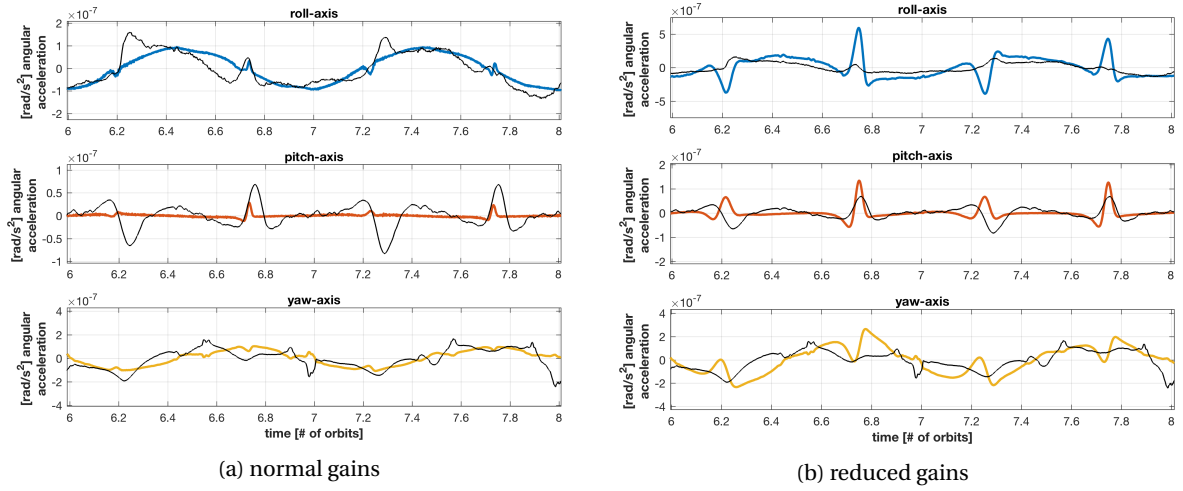


Figure 4.14: Time series of the simulated acceleration of two orbits during the complete day of 1 November 2011. (a) Shows the simulated acceleration using normal gains, while (b) shows the acceleration of a simulation in which the inner loop gains of the NINDI-controller were reduced by a factor of 10.

that the gains were optimised to minimise the attitude, angular rate, and angular acceleration error of the time series. So it is possible that by minimising these errors, the controller responds aggressively to disturbances which then produces the small and short (with respect to time), i.e. high frequency, spikes as seen in Figure 4.14. This then ultimately results in exceeding the PSD requirements. Lastly, the reduction by one order of magnitude that was applied to the inner gains was chosen at random, without any attempt to optimise the gains. This means that no problems are to be expected to further tune the gains to match the ESA requirements both in the time domain and in the frequency domain.

#### 4.2.2. Comparison of the NINDI-controller with GOCE's observed behaviour

In subsection 4.2.1 it was shown that the NINDI-controller is more than able to control the spacecraft's state and keep everything within the ESA requirements in the time domain. However, the actual GOCE mission performed well too, and thus it is interesting to also compare the performance of the controller with the actual performance of the satellite. The first step in doing this is to take the GOCE measured state as a reference, and then subtract the simulated state to see what the differences between the two are. The differences in the attitude angle, angular rate, and angular acceleration are shown in Figure 4.15, Figure 4.16, and Figure 4.17 respectively. During the comparison between the simulated and the measured behaviour four interesting points were observed. This section will go through each of these four points, and use the figures presented in this section to support the validity of the points.

Before starting with the main discoveries, the interpretation of the figures Figure 4.15 to Figure 4.17 is as follows. The y-axis shows the argument of latitude,  $u$ , of the spacecraft and the x-axis shows the time that has passed. The data of the simulation is then binned where the data points in each bin are combined by taking the mean. Because it is important to recognise what influence taking the mean has on the data, the standard deviation of the data in each bin has also been plotted. Almost all bins contain around 30 data points, with the exception being the bins around  $0^\circ$  and  $360^\circ$  argument of latitude. Around these argument of latitudes it can occur that a bin is empty, however since almost all other bins contain roughly 30 data points (with a standard deviation of 3), the standard deviation plots can be used to tell how justifiable it is to take the mean of a bin. The data in the bins themselves contain the measured values of the GOCE state, subtracted by the simulated values. Thus, the figures Figure 4.15 to Figure 4.17 show the difference between the measured and simulated state. The argument of latitude plots for the measured and simulated state separately have been included in Appendix B.

The first thing that stands out, and could also already be seen in Figure 4.6, is the coupling of the roll- and yaw-axis. In Figure 4.6, it can be seen that a control actuation on the roll- or yaw-axis also affects the other axis. This follows directly from Equation 2.16, the roll- and yaw-axis have the control dipole of the pitch-axis in common to exert a torque against the EMF. What changes however is the magnitude of the EMF vector components in the spacecraft body frame. This can result in strong, undesired, spikes in the control torque

depending on the control action and the position of the spacecraft above the Earth. This behaviour is not unexpected since time dependent control, discussed in section 2.2, was introduced to remove the coupling between the roll- and pitch-axis above the equator. The coupling between the two axes is most visible in Figure 4.17, since this figure shows the acceleration on the spacecraft. This is however caused by the simulated roll-axis attitude angle and angular rate being smaller than the measured, therefore Figure 4.15 and Figure 4.16 are dominated by the measured values. Another observation that emphasises the roll-yaw coupling can be made in Figure B.1. The roll-axis in these figures show a signal with a period of once per day, that is associated with flying over the South Atlantic Anomaly (SAA). Although the roll-axis itself might not be directly affected by flying over the SAA, it is affected indirectly through its coupling with the yaw-axis.

To get more insight in the source of the acceleration, geographic plots around the equator of the control actuation have been included in Figure 4.18 to Figure 4.20. While the geographic plots do not allow to investigate time related behaviour, they do allow for visualising longitude – latitude dependent behaviour. These figures also show the measured control torque as a reference, subtracted by the simulated control torque. Separate plots have again been included in Appendix B. In Figure 4.15 to Figure 4.17 dependency on ascending and descending tracks can be seen, therefore the geographic plots have been split up in ascending and descending tracks. Most noticeable in Figure 4.18 are the dark colours at lower and higher latitudes. These dark colours are caused by the fact that the direction of the control torque is reversed for the measured and simulated case. See section B.2 to see the control torque plots of the measured and simulated case separately. The next thing that is striking in Figure 4.18 is the dark band close to the equator on the descending track. This indicates a strong difference on the roll-axis between the measured and simulated torque. Looking at section B.2, it can be seen that this is in fact a result from a commanded control on the roll-axis. This appears to follow the magnetic equator, so every time the spacecraft passes the magnetic equator on the descending track, the NINDI-controller torques the roll-axis quickly. Above the magnetic equator control on the roll-axis is difficult. The magnetic field vector lies in the plane of the spacecraft x-axis above the magnetic equator, as seen in Figure 2.3, which results in little to no roll-axis control, shown in Figure B.4. Yaw-axis control for the simulated controller on the other hand is especially active around the magnetic equator, and the measured controller to a lesser extend as well. This has to do with yaw-axis control being particularly effective above the magnetic equator, but since the magnetic equator and geographical equator are close to each other, it is also where the (co-rotating) atmospheric velocity is large. Since the axes are coupled, and the roll-axis has a lower inertia (see Table 2.1), it makes sense for the roll-axis to respond aggressively after having passed the equator. Especially since it was already concluded that the NINDI-controller as it is tuned now responds quickly, with relatively large torques.

While explaining the behaviour of the coupling of the roll- and yaw-axis, the implementation of time dependent control to reduce a similar coupling between roll- and pitch-axis was mentioned. No attempt was made to reduce coupling between roll- and yaw-axis, instead time dependent control was specifically implemented to decouple the roll- and pitch-axis. The decoupling itself was successful, since the pitch-axis is controlled very stably without affecting the roll-axis much, see Figure 4.19 and Figure 4.18. In this case however, Figure 4.19 does not tell the whole story. Time dependent control was implemented according to theory, but there is a large difference between the observed control actuation on the pitch-axis of the GOCE spacecraft, and the simulated. This is shown in Figure B.5, where in the simulated case the pitch-axis control is indeed suspended close to the magnetic equator and ramps up again afterwards. However, the measured control actuation shows a longer pitch-axis control suspension. This is an indication that the time dependent control matrix as shown in Equation 2.19 is not exactly how it was implemented during the GOCE mission. It is possible that the actual mission included a control suspension band for the pitch-axis control, activated by a certain threshold instead of the continuous method shown in Equation 2.19. This implies that as the spacecraft gets close to the magnetic equator ( $B_x^2/||\mathbf{B}||^2 \rightarrow 1$ ), the pitch-axis control is suspended sooner, and also for a longer period of time, than what is expected based on the theory of time dependent control.

The third point to discuss is the pattern, or apparent lack thereof, on the yaw-axis in Figure 4.15 to Figure 4.17. This patterns might come off as anomalous, but bear in mind that the figures actually show the difference between the measured and simulated state. Based on how it looks, one might get the impression that something is going wrong in the binning process, but if that were the case than similar results are to be expected on the other two axes. What is interesting is that the standard deviation of the data points in a bin is quite small on the roll- and pitch-axis, but not always on the yaw-axis. Interestingly, Figure 4.10 shows multiple series of increases in the values of the geomagnetic indices. These values are indicative of increased thermospheric activity. Increased thermospheric activity means that more, noisy, torque can be expected on the yaw-axis. The comparison of the body torques shown in Figure 4.6 already showed that the simulated con-

troller is more active on the yaw-axis than the actual controller was, and more importantly when it responds it uses high frequencies. The use of more high frequencies was already shown in Figure 4.12 and Figure 4.13. Therefore, the noisy pattern, as well as the periods of high values for the standard deviation on the yaw-axis can be associated to the increases in geomagnetic activity. The aggressiveness of the simulated controller does not only present itself on the yaw-axis, but can also be seen on the other two axes in Figure 4.17. The strong signal on the roll-axis is most likely due to the yaw-axis coupling, and the pitch-axis suspends control for a short while when above the magnetic equator but then aggressively engages it again, as shown in Figure 4.19.

The fourth and final point has to do with the behaviour of the spacecraft during ascending and descending tracks. For example, the difference in the angular rate on the pitch-axis shown in Figure 4.16, shows a positive/negative difference during ascending and descending tracks. Different behaviour in ascending and descending tracks can of course also be seen in Figure 4.18 to Figure 4.20. Of these three figures, Figure 4.20 shows the strongest difference between measured and simulated state, and the strongest difference occurs around the magnetic equator. As mentioned before, it makes sense for the control on the yaw-axis to be active here, but it is interesting that the measured and simulated case are so different. Also, both in the ascending and descending tracks abnormal activity around South-America can be observed. This behaviour is most likely caused by the SAA. Another relevant structure in the EMF apart from the SAA is the fact that the magnetic North Pole is at high latitudes, while the magnetic South Pole is at lower (absolute) latitudes, these structures in the EMF were discussed in section 2.2. This means that the spacecraft is close to the magnetic North Pole on most orbits, and further away from the magnetic South Pole. Being closer to the magnetic North Pole complicates yaw-axis control, which in turn affects yaw-axis control during a large part of the descending tracks.

In general it can be said that the simulated controller performs equal or better than the measured performance of the GOCE spacecraft. There are however a few interesting differences that have been discussed above. An important difference that was already found before is the aggressiveness of the NINDI-controller, giving rise to short burst of control actuation. Furthermore it was found that the simulated yaw-axis is more active in control than the actual GOCE spacecraft appears to have been. This also has an adverse affect on the simulated roll-axis for which a strong coupling effect with the yaw-axis was observed. Lastly, comparing the obtained simulation results with the measurements from the GOCE spacecraft, it seems likely that the implementation of time dependent control in the actual GOCE controller is different than how it is described in theory. Taking these differences into account, it would be really interesting to know what the actual GOCE controller looked like to be able to further explain the differences. This is unfortunately not possible however, because as far as this author is aware the internal functioning of the GOCE controller is not publicly available.



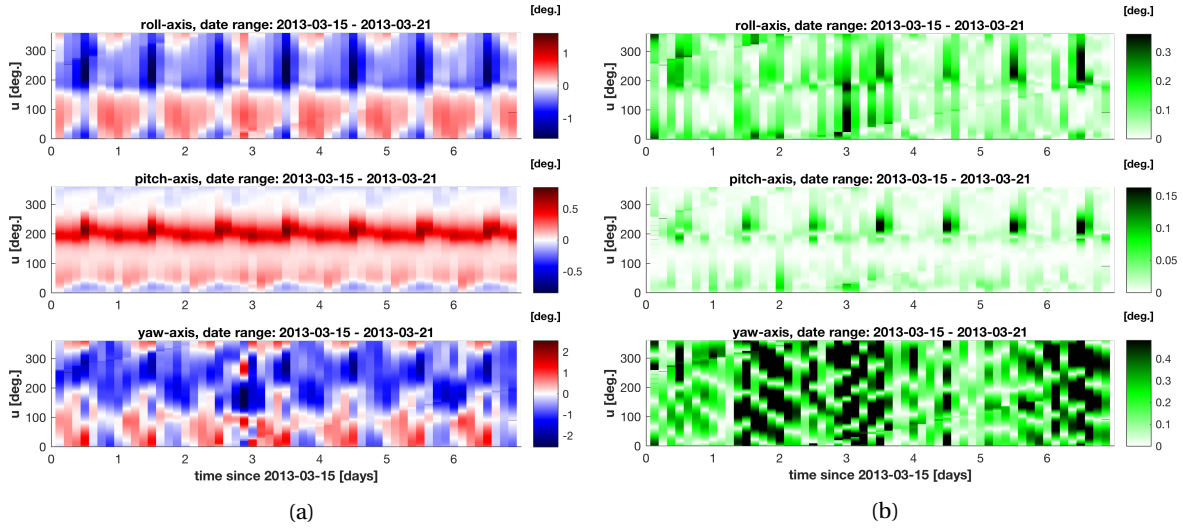


Figure 4.15: Argument of latitude plot of the measured attitude of the GOCE satellite subtracted by the simulated attitude. The actual differences are shown in (a), in which the colour bar scale is set to the maximum absolute values of the GOCE measured attitude on each axis. Since the data is binned together, (b) shows the standard deviation of each bin, where the colour bar scale is set to the standard deviation of the GOCE measured attitude on each axis.

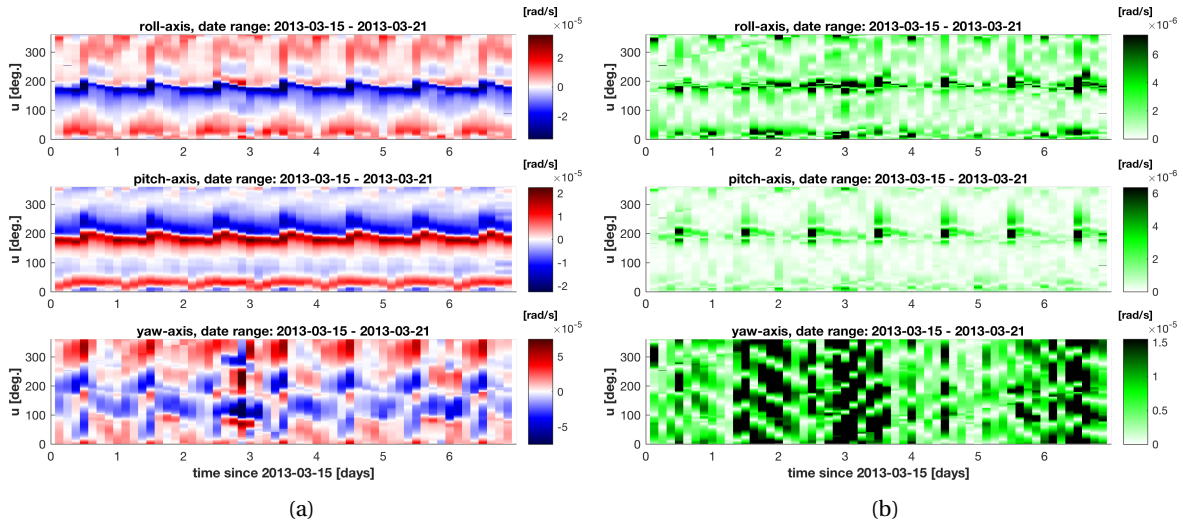


Figure 4.16: Argument of latitude plot of the measured angular rate of the GOCE satellite subtracted by the simulated angular rate. The actual differences are shown in (a), in which the colour bar scale is set to the maximum absolute values of the GOCE measured angular rate on each axis. Since the data is binned together, (b) shows the standard deviation of each bin, where the colour bar scale is set to the standard deviation of the GOCE measured angular rate on each axis.

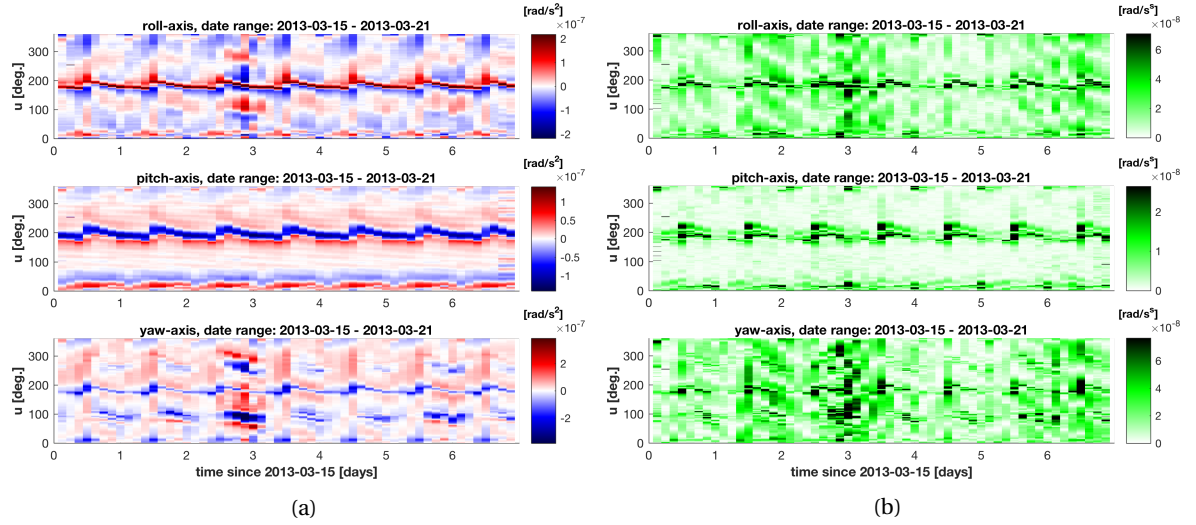


Figure 4.17: Argument of latitude plot of the measured angular acceleration of the GOCE satellite subtracted by the simulated angular acceleration. The actual differences are shown in (a), in which the colour bar scale is set to the maximum absolute values of the GOCE measured angular acceleration on each axis. Since the data is binned together, (b) shows the standard deviation of each bin, where the colour bar scale is set to the standard deviation of the GOCE measured angular acceleration on each axis.

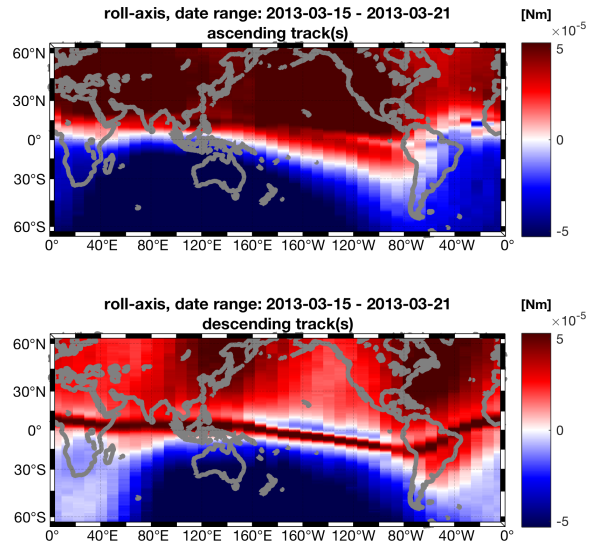


Figure 4.18: Miller map projection around the equator of the measured achieved control torque on the roll-axis of the GOCE satellite subtracted by the simulated achieved control torque. The colour bar scale is set to the maximum absolute value of the GOCE measured achieved control torque.




Figure 4.19: Miller map projection around the equator of the measured achieved control torque on the pitch-axis of the GOCE satellite subtracted by the simulated achieved control torque. The colour bar scale is set to the maximum absolute value of the GOCE measured achieved control torque.

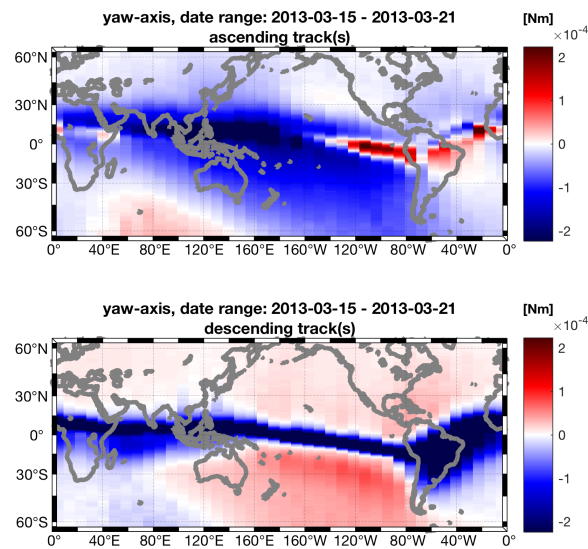


Figure 4.20: Miller map projection around the equator of the measured achieved control torque on the yaw-axis of the GOCE satellite subtracted by the simulated achieved control torque. The colour bar scale is set to the maximum absolute value of the GOCE measured achieved control torque.

### 4.2.3. Controller robustness to uncertainties and measurement errors

The results shown in this chapter so far have all been done by simulating the controller and torques on the spacecraft assuming no errors are present at all. In reality however, all measurements made by the satellite include some errors. These error sources can be anything, from uncertainties in the exact location of a measurement device, to random noise. Since dealing with errors is something that a controller will have to be able to do, this section will evaluate the performance of the NINDI-controller when errors are introduced in the simulation environment. The robustness of the controller will be evaluated on the day that was used for tuning the controllers, which is 1 November 2011. The errors that will be introduced focus on measurement errors for the attitude angles, angular rates, angular accelerations, and errors in the actuation of the magnetorquers. The actuation of the magnetorquers is done by sending a current to them, based on the required torque at a point in time. To capture the errors in actuation of the magnetorquers noise will be introduced on the current commands.

The magnitude of the errors both in measurements and actuation will first be of a size that is realistic for the GOCE mission. These values are based on a thesis written by Stummer [38], and values from the GOCE documentation [5]. The NINDI-controller is expected to only perform slightly worse when introducing errors that can realistically be expected. Therefore another round of simulation will be done with larger error sources to see when the controller is unable to reject the noise. The sizes of the realistic errors, or nominal errors, are shown in Table 4.3, this table also includes 'large errors'. Which show the sizes of the errors that cause the controller to not be able to keep the state within the requirements. The errors shown in Table 4.3 are used as the standard deviation for introducing white noise on top of the measurements. Two simulations are done. One with the nominal errors and one with the large errors. The results of introducing errors in the measurements and actuation are shown in Table 4.4. As a reference the mean and standard deviation of the simulation without any errors introduced are also shown. Table 4.4 also shows the mean and standard deviation for the three simulations using the PD-controller. This is done because it will be interesting to verify the difference in performance of the PD-, LINDI-, and NINDI-controller.

From Table 4.4 it can be concluded that the attitude and angular rate remain stable when introducing both the nominal and the large errors. The most interesting effect is seen in the angular acceleration, in which all controllers start to perform worse when introducing noise. The INDI-based control architectures rely on accurate acceleration measurements, thus when introducing noise both the LINDI- and NINDI-controller are less able to reject the noise than the PD-controller. Although the overall performance of both INDI-based controllers is still better, the PD-controller appears less sensitive to noise. This is further confirmed when introducing the large errors. Now the INDI-based controllers really start to perform poorly, while the PD-controller shows again a smaller decrease in performance. Especially noticeable in Table 4.4 is that the roll-axis suffers most from the introduction in errors. The reason for this is that the yaw-axis is somewhat passively stabilised and therefore less actively controlled, and the pitch-axis was weighted heavier in the GA due to the ESA requirements. On top of the pitch- and yaw-axis having these advantages over the roll-axis, the roll-axis' inertia is also a lot lower, see Table 2.1. This explains why the change in yaw-axis acceleration behaviour is the least when introducing errors, and the change in roll-axis behaviour is largest. Lastly, the PD-controller does not rely on acceleration measurements at all, and instead controls the attitude of the spacecraft by tracking a commanded rate, discussed in section 3.3. The PD-controller does however directly set a torque, which changes the acceleration, but the change in acceleration is filtered by the spacecraft inertia, before it changes the angular rate. In contrast, both the INDI-based controllers need the angular acceleration measurements

Table 4.3: Magnitude of errors introduced.

		Measurements			Actuation
		Attitude angle["]	Angular rate [rad/s]	Angular acceleration [rad/s <sup>2</sup> ]	Current to Magnetorquers [%]
Nominal errors	Roll	3	$5.66 \cdot 10^{-8}$	$1.41 \cdot 10^{-9}$	3
	Pitch	3	$4.95 \cdot 10^{-10}$	$6.36 \cdot 10^{-10}$	3
	Yaw	3	$6.36 \cdot 10^{-8}$	$2.83 \cdot 10^{-9}$	3
Large errors	Roll	50	$6.0 \cdot 10^{-7}$	$2.0 \cdot 10^{-8}$	5
	Pitch	50	$5.0 \cdot 10^{-9}$	$7.0 \cdot 10^{-9}$	5
	Yaw	50	$7.0 \cdot 10^{-7}$	$3.0 \cdot 10^{-8}$	5

Table 4.4: Results of introducing errors in the model for the PD-, NINDI, and LINDI-controller.

		Attitude [10 <sup>-1</sup> deg.]			Angular rate [10 <sup>-7</sup> rad/s]			Angular acceleration [10 <sup>-10</sup> rad/s <sup>2</sup> ]		
		Roll	Pitch	Yaw	Roll	Pitch	Yaw	Roll	Pitch	Yaw
<b>P(I)D-controller</b>										
<b>Mean</b>	<b>No errors</b>	-0.035	-0.013	-3.497	-0.545	0.201	-0.365	-1.472	-0.105	1.250
	<b>Nominal errors</b>	-0.035	-0.011	-3.498	-0.544	0.201	-0.367	-1.468	-0.092	1.245
	<b>Large errors</b>	-0.035	-0.013	-3.503	-0.550	0.202	-0.410	-1.558	-890.0	1.231
<b>Std.</b>	<b>No errors</b>	0.277	1.364	4.864	41.375	111.2	129.1	<b>1,108</b>	906.9	<b>672.9</b>
	<b>Nominal errors</b>	0.279	1.365	4.864	41.585	111.2	129.1	1,120	907.9	672.3
	<b>Large errors</b>	0.283	1.367	4.916	43.482	111.5	130.7	<b>2,523</b>	932.6	<b>676.7</b>
<b>NINDI-controller</b>										
<b>Mean</b>	<b>No errors</b>	0.863	-0.025	-1.473	-0.380	0.047	1.117	-3.186	0.064	0.503
	<b>Nominal errors</b>	0.865	-0.025	-1.472	-0.379	0.048	1.121	-3.192	0.064	0.514
	<b>Large errors</b>	0.874	-0.024	-1.464	-0.385	0.046	1.124	-3.152	0.055	0.496
<b>Std.</b>	<b>No errors</b>	0.946	0.096	4.267	25.877	4.678	103.3	<b>696.1</b>	63.703	<b>690.0</b>
	<b>Nominal errors</b>	0.946	0.096	4.270	25.893	4.683	103.3	879.4	108.9	690.6
	<b>Large errors</b>	0.943	0.095	4.273	26.454	4.841	103.4	<b>5,883</b>	1,274	<b>742.0</b>
<b>LINDI-controller</b>										
<b>Mean</b>	<b>No errors</b>	0.795	-0.003	-1.173	-0.433	0.057	1.098	-3.176	0.058	0.199
	<b>Nominal errors</b>	0.796	-0.003	-1.173	-0.432	0.057	1.097	-3.178	0.060	0.199
	<b>Large errors</b>	0.801	-0.003	-1.162	-0.433	0.059	1.125	-3.192	0.071	0.197
<b>Std.</b>	<b>No errors</b>	0.957	0.042	4.215	35.354	6.288	107.1	<b>724.8</b>	116.4	<b>699.7</b>
	<b>Nominal errors</b>	0.956	0.042	4.214	35.356	6.321	107.1	1,144	619.5	705.9
	<b>Large errors</b>	0.964	0.042	4.230	35.952	6.423	107.7	<b>4,908</b>	1,026	<b>741.0</b>

to stabilise the state. In conclusion when looking at the results of introducing the nominal errors specifically, it can be seen that the LINDI- and NINDI-controller are better able to stabilise the spacecraft than the PD-controller. It is only when introducing larger than expected errors that the PD-controller is more stable. This then emphasises that a GOCE-type platform, with its accurate measurements, is suitable for an INDI-type controller.

During this thesis research a modelling error in the aerodynamic torque was found. After having removed the modelling error the spacecraft was successfully stabilised by the NINDI-controller, even for storm periods. However, the initial modelling error has shown the sensitivity of the controller to aerodynamic torque. Furthermore, the torque models are based on the work by Visser et al. [41] in which it is concluded that “there is a mismatch between the aerodynamic model from ANGARA and the thermospheric density and wind data derived using a panel model”. Together with the low altitude the GOCE spacecraft flew at and the unpredictable, and volatile, nature of the atmosphere, there is a strong incentive to investigate how an increase in aerodynamic torque will affect the performance of the controller. The day for which the gains have been optimised is again chosen as a reference to investigate this. Then, using the data of this day as a basis, the modelled wind velocities and densities will be manually increased. Increasing these two variables is done separately, and they are increased with factors until control is completely lost.

The results of the investigation in how the controller performs during increases in both atmospheric density and modelled winds, is shown in Figure 4.21. The top row is the simulated output for when no increases are implemented. The other rows show increases in winds and atmospheric density separately. From Figure 4.21 it can be concluded that with a factor of five increase with respect to the modelled wind values for that day the controller is still able to stay within the time domain requirements. Increasing the wind velocities with a factor of 10 causes the state to occasionally go beyond the set requirements, and with a factor of 15 it can be seen that control is completely lost. This is especially visible on the pitch-axis angular rate. Increasing the density manually has similar effects, with an increase of 7.5 the controller is able to control the state within the set requirements, except on the yaw-axis from orbit 10 to 14, and recovers afterwards. However, increasing the density with a factor of 10 causes both more and stronger oscillations. It is important to keep in

mind that these factors were multiplied with the values on 1 November 2011. Multiplying the wind velocities with a factor of 10 for this day gives a maximum wind velocity of around 10,000 m/s and a mean of -1,300 m/s on the body y-axis.

These numbers can be put into perspective by comparing them to a day with a storm in it. Data is available for 31 May to 2 June 2013. The geomagnetic indices as well as the performance of both the actual GOCE spacecraft and the simulated spacecraft are shown in Figure D.1. The figure clearly shows the effects of the storm starting shortly after the start of the second day in the period. During this period the maximum wind velocity was around 1,000 m/s and a mean of 180 m/s. So multiplying the wind velocity with a factor of 10 yields physically not realistic results for the wind velocities, but the fact that the controller is able to almost keep the spacecraft state within the time domain requirements is all the more impressive. Doing the same comparison for the density it is found that at a factor of 7.5, the density has a mean of  $3.35 \cdot 10^{-10} \text{ kg/m}^3$ , and a maximum value of  $5.35 \cdot 10^{-10} \text{ kg/m}^3$ , while the storm day has a mean of  $7.75 \cdot 10^{-11} \text{ kg/m}^3$  and a maximum value of  $1.87 \cdot 10^{-10} \text{ kg/m}^3$ . Keep in mind that the satellite flew at an altitude that was around 30 kilometres lower during the storm day, than it was during the day in which the factors are applied. This provides some context on the difference in densities, but it still shows that the factor of 7.5 on the density, at which the controller is close to keeping the spacecraft state stable, is closer to being realistic than what was the case for the wind values. In fact, using the 'MSIS' class model, of the 'NRLMSISE-00' model implementation in MATLAB<sup>TM</sup>, it can be found that the density values used when manually increasing the density can be expected in an altitude range of 180 to 210 kilometres. This altitude value does depend on the day and solar activity, amongst other, but it gives an indication of what can be expected. Provided then that the controller is still able to control the spacecraft state reasonably at such low fictitious altitudes, it is interesting to see how the NINDI-controller performs during the re-entry phase of the GOCE mission.

Having evaluated the performance of the controller under various wind and density conditions, one important aspect that was not yet evaluated is what the actual currents are that are sent to the magnetorquers. This is relevant, because there is a maximum current that the spacecraft could apply during the science phase, and a total maximum current that could be sent to the torquers. The maximum current during the science phase was 50 mA, and the total maximum current was 544 mA [17]. The y-axis scale in Figure 4.22 has been chosen such that all spikes are captured, with the exception of outliers. Furthermore, the scales of each axis are kept identical, such that a comparison can be easily made. The largest current commands can be seen on the pitch-axis when the atmospheric density is increased by a factor of 10. However, even in this extreme case the current commands stay under 300 mA, with the exception of one outlier. However, it also shows that for the storm period, there are a few narrow spikes that cross the 50 mA. Since this was still a regular day in which scientific data was produced this should obviously not happen. It is again however likely that this has to do with the aggressiveness of the model, given how narrow the spikes are. Furthermore, Figure D.1 shows that the time domain requirements are still met. Regarding the factors, it was already observed that manually increasing the winds and the density results in values that are not typical, and in the case of the winds even physically impossible. Therefore it is not so strange to see the currents spike so high, and it is in fact remarkable that the point at which the torquers are saturated, according to the technical specifications of the GOCE spacecraft, has not been reached yet. This underlines that it would indeed be interesting to see how the NINDI-controller performs during the re-entry phase of the GOCE mission, given that there is room left in the torquers to increase the currents, and thus to apply more control torque.

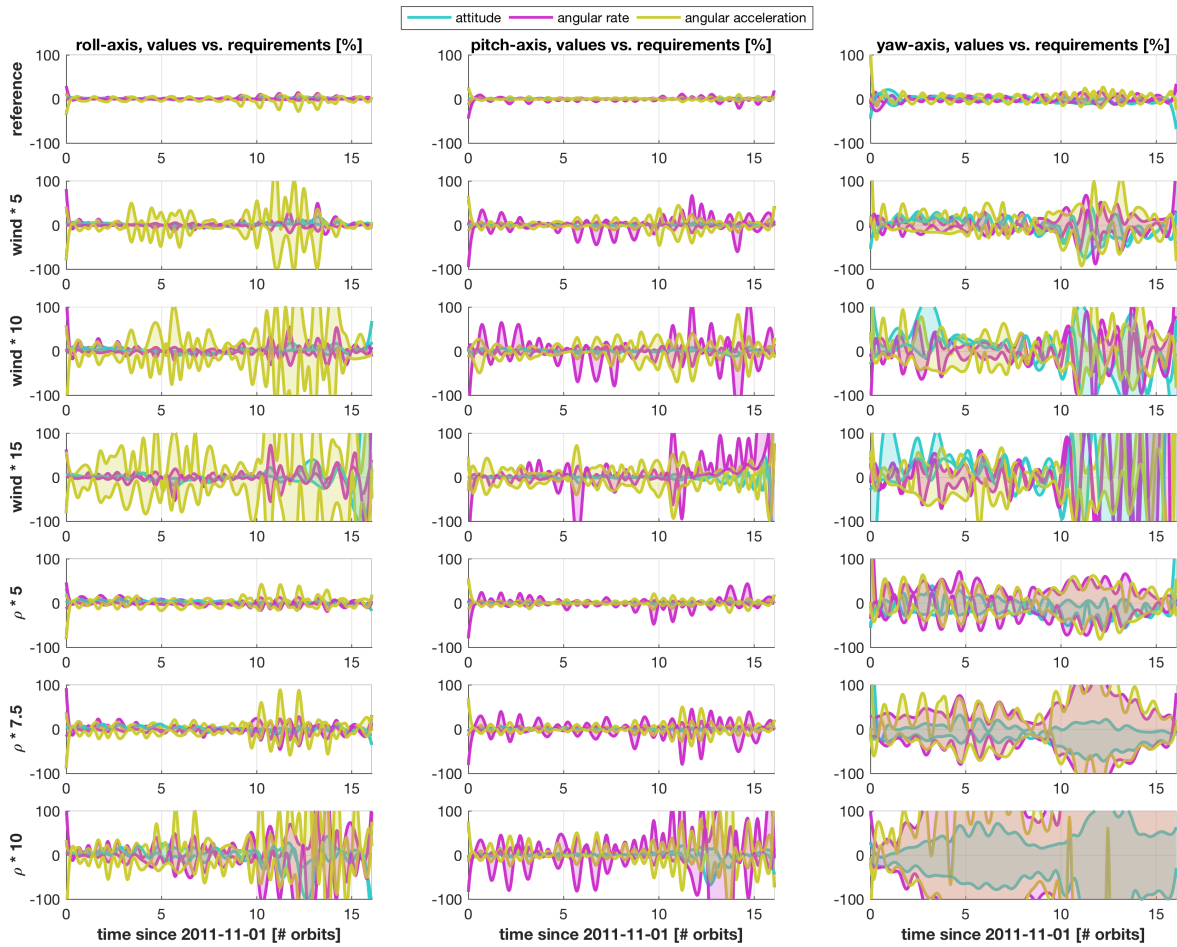


Figure 4.21: Influence of manually increasing the wind velocities and the atmospheric density in the simulator. All figures are simulations of a complete day simulation on 1 November 2011.

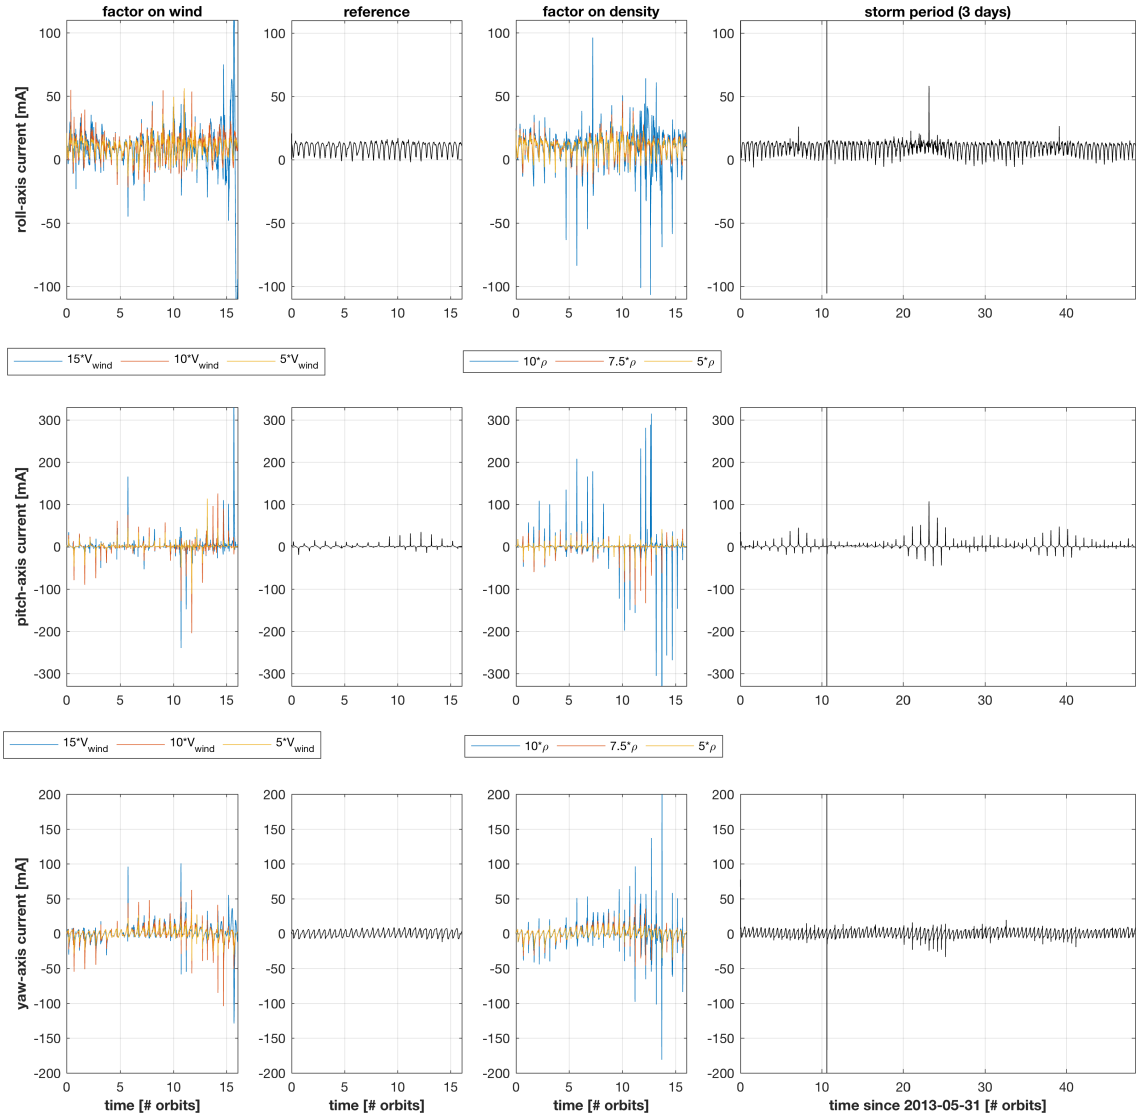


Figure 4.22: Commanded currents to magnetorquers to create torque on each axis. The first three columns are for the entire day of 1 November 2011. The last column shows the same plot for a comparable storm period. The period of choice is 31 May 2013 to 2 June 2013. The left column shows the currents for different factors applied to manually increase the wind speeds, the second column is the reference (no factors applied) for 1 November 2011, and the third column shows the currents for different factors applied to manually increase the atmospheric density.



#### 4.2.4. Controller tuning difficulty and sensitivity to changes in gains

Given the non-linear nature of this control problem, and the difficulty associated with tuning such control systems, it was chosen to use a GA to tune the controllers. How exactly the GA has been set up and how it is used to find a suitable set of gains was discussed in section 3.4. Having used the GA to tune the controllers provides us with an opportunity to quantify the effort required to find stabilising gains for them. To express the difficulty of tuning the controllers, Figure 4.23 was made. While this particular figure is for the NINDI-controller, similar figures for the other controllers show the same behaviour, though the fitness score of each generation as well as the final fitness score can be different. Figure 4.23 shows that most of the performance increase is achieved in the first five generations. Thereafter small improvements are made, but especially after 10 generations the improvements are negligible. An important additional remark to make here is that although a genetic algorithm is an optimiser, this thesis has made no attempt to guarantee that the found set of gains are in fact the single optimal solution. Instead, the focus of this thesis was to make the controller work.

In the end the genetic algorithm produces multiple individuals that have a high fitness score. These individuals are expressed as a series of 1's and 0's, i.e. a bit string. One would expect that after some generations a pattern should emerge in the distribution of these bits. Since the population starts out completely random, but is steered towards the situation where individuals have a favourable fitness score. To see whether such a pattern is visible, Figure 4.24 was included in this report. This figure displays how many 1's there are on each position along the bit string for the top 100 best individuals, in terms of their fitness score. The colours are used to highlight which bit indices belong to the same factor or gains, i.e. they visualise the genes of an individual. Also see section 3.4 for further explanation how this is done. As an example, the bits around index 80 to 95 can be traced back to the inner loop yaw proportional gain. Since the first bits from this set do not contain a lot of 1's, it is apparently beneficial to have a comparably low yaw inner loop gain. Then taking a look back at Table 4.1 and it can be seen that the inner loop proportional yaw gain for the best performing NINDI-controller is indeed two orders of magnitude smaller than the other two inner loop proportional gains. A similar example is that the first three indices often have 1's in them. These three indices are traced back to the outer loop factor, which mean they often return a higher value. This means the factor becomes smaller, i.e.  $10^{-4} < 10^{-2}$ . This can also be seen in Table 4.1. However, Table 4.1 only shows one set of gains, out of all the generations that were simulated using the GA. Both the examples mentioned above, the inner loop proportional yaw-axis gain and the outer loop factor, can also be observed in Figure 4.25. The spread of the inner loop yaw-axis proportional gains does indeed tend to smaller values than the spread of the other two axes. Similarly, the roll- and pitch-axis outer loop proportional gains are mostly displaced to the left of the inner loop proportional gains.

The results that have thus far been shown in this report were produced by taking the best performer for each controller in terms of the fitness score and running simulations with those gains. However, once a set of gains has been found, it is interesting to evaluate how sensitive the controller is to changes in these gains. It had already been seen in section 4.1, that the inner loop gains of the NINDI-controller can be reduced by a factor of 10, without any instabilities occurring afterwards. The question then is whether this is because the INDI control architecture allows for a wider range of gains to produce satisfactory results, or whether it was luck. In Figure 4.25 the spread of the gains for both the NINDI-, and PD- controller is shown. The gains shown here are again taken from the top 100 performers, in terms of the fitness function. The gains have been binned together in logarithmic bins due to possible magnitude differences between them. It is immediately clear that the PD-controller gains are clumped together, with the most extreme case being the pitch-axis inner loop proportional gains. The NINDI-controller also has clumps around which the gains are stable, but the clumps are in general wider than those of the PD-controller, and the NINDI-controller has outliers. Lastly, based on Figure 4.25 it can be concluded that both controllers are least sensitive to the inner loop derivative gains, and all the yaw-axis gains tend to be more spread out, highlighting the aerodynamic properties of the GOCE spacecraft and the reduced control effort on the yaw-axis. One last interesting observation that can be made is the extreme sensitivity of the inner loop proportional gain on the pitch-axis.

Although Figure 4.25 already gives an indication of the sensitivity of both controllers to gain changes, the sensitivity will be investigated further by decreasing and increasing all the gains simultaneously with a certain percentage. Doing this for both the NINDI- and the PD-controller allows to verify whether there is any difference in the sensitivity. For a reduction in gains three different percentages have been chosen, 12%, 24% and 48%. Thus in the case of the largest reduction the gains are almost halved. The reason why 48% was chosen and not 50% is because the PD-controller was so unstable at a reduction of 50% that the simulation threw an error. For the increase in gains also three percentages were chosen, 33%, 66%, and 100%. The




Figure 4.23: Fitness score of the best performer in each generation during a genetic algorithm run. The fitness score is the sum of the norm of the total error on each axis, for the attitude, angular rate, and angular acceleration. (a) Shows the fitness score from the beginning to the end of a genetic algorithm run. After 20 generations the genetic algorithm stopped, because no improvements had been found for five generations. (b) Shows a zoomed in version of how much (or little) the genetic algorithm improves after 5 generations.




Figure 4.24: Count of the amount of 1's on every bit index for the 100 individuals with the best fitness score of the entire NINDI-controller GA population. The colours show which bits are used to form one trait of an individual.

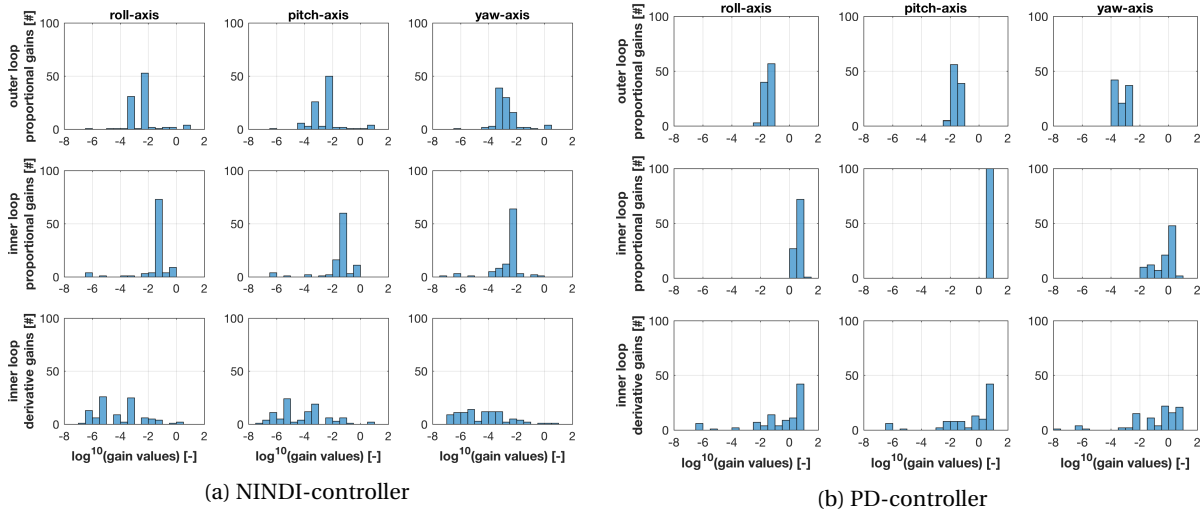


Figure 4.25: These two figures show the spread of the 100 best performing set of gains in terms of the fitness function score, of both the NINDI-controller (a), and the PD-controller (b).

largest increase, 100%, was chosen to evaluate the results of both doubling and splitting the gains in half. The simulations are once more done on the complete day of 1 November 2011, and are shown in Figure 4.26. The performance in this figure is expressed as the norm of the root mean square of each axis. As for example the proportional gain is increased, it can be expected that more oscillations occur. Therefore the root mean square produces a larger value, and a lower value for the performance is thus better. For the reduction in gains the results are as expected. The NINDI-controller scores very similar in all cases, with a reduction in performance visible as the gains are increased with respect to the reference set of gains. The PD-controller is especially sensitive for reducing the gains, while increasing them appears to improve the performance, at the cost of angular acceleration. This behaviour is caused by the controller becoming more aggressive controlling the attitude and angular rate better. However, because it is now more aggressive the angular acceleration becomes more noisy. It is interesting that something similar was also observed in Table 4.4, where the PD-controller was able to keep the attitude and angular rate stable when introducing errors, while stability for the angular acceleration was lost. It can also be seen that the NINDI-controller is apparently more sensitive to increases in gains than decreases. This can be explained by the fact that the NINDI-controller gains are already relatively large, as shown in Table 4.1. When these gains are increased they get into the region where there apparently are not that many stabilising gains for the NINDI-controller, based on Figure 4.25. This is why the performance of the NINDI-controller can be seen deteriorating at an increasing pace.

These findings do match with Figure 4.25, in which the distribution of the gains themselves was shown. The PD-controller gains have a tendency for higher values than the NINDI-controller gains, and as was observed in Figure 4.26 decreasing the PD-controller gains has a more negative effect on the performance of the controller than increasing them. Similarly, the NINDI-controller performance is largely flat in Figure 4.26, and it does have a wider spread in the gains than the PD-controller, as shown in Figure 4.25. It can also be seen that there are apparently not many stabilising inner and outer loop proportional gains around the  $10^0$  order of magnitude. Lastly, using a GA to tune the gains of the controllers has proven to be a valuable tool, as it has been successfully used to provide further insight in the observed behaviour of the controller.




Figure 4.26: The effect that increasing and decreasing all the gains with the same percentage has on the performance of the controller. All these simulations were done for the complete day of 1 November 2011. Performance in this figure is expressed as the norm of the root mean square of each axis, for the attitude angle, angular rate, and angular acceleration.

### 4.3. Discussion

In this chapter it was shown that a controller based on INDI-control architecture can successfully be used to control a LEO spacecraft's attitude using magnetic attitude control only. Multiple reference frames were defined to control the spacecraft in, and it was found that the controller was unable to stabilise the spacecraft in the D-frame. Even when the spacecraft is controlled in the O-frame, which tracks the co-rotating atmosphere, significant aerodynamic torques can still occur when the spacecraft's attitude is close to  $0^\circ$ . The premise of the D-frame is to track the linear acceleration acting on the spacecraft, which includes thermospheric winds. Attitude control in this reference frame therefore promises to be even more effective in counteracting drag with the IPA than in the O-frame. Although an error in the code that defines the D-frame was not found, it is likely that erroneous code is the reason for the D-frame not working.

The three controllers that were used in this thesis research all needed to be tuned. Finding stabilising sets of gains for the controllers was done by using a GA. This has worked really well in general, although there were some disadvantages as well. First, it can be difficult to include all relevant parameters in the fitness function, while the performance of the GA is completely governed by it. This was found out the hard way, since none of the controllers were able to meet the ESA frequency domain requirements. The reason that these were not met, is because they were not included in the fitness function. Second, given the apparent extreme sensitivity of the pitch-axis inner loop proportional gain for the PD-controller, and the fact that the performance of the PD-controller actually increased when its gains were increased, it is possible that the search space of the GA for gains was not optimal for the PD-controller. Although the improvement of the PD-controller when the gains were increased did come at the cost of a reduction in angular acceleration performance.

The simulated NINDI-controller performance in general was significantly better than the measured performance of the GOCE spacecraft. Although the NINDI-controller did turn out to be more aggressive, resulting in high frequency controller action and thus spikes of high acceleration. These spikes have been seen to result in higher magnitude values of the one-sided PSD on days with more thermospheric activity, and they were the cause of the controller not meeting the frequency domain requirements. By reducing the inner loop gains however, it was found that the PSD magnitudes decrease, allowing for the frequency domain requirements to be met. Although all axes were tuned too aggressive, the yaw-axis in particular was tuned more aggressive than the GOCE spacecraft appears to have been. The high frequency control action on the yaw-axis also caused spikes in acceleration on the roll-axis, through their magnetic coupling. This effect was especially prominent around the equator, which was found to be caused in part by the yaw-axis still recovering after having passed the poles, especially the North Pole. Additionally, when the spacecraft passed the SAA the effect became even more visible. Since reducing the inner loop gains has a positive effect on meeting the frequency domain requirements, and since it was found that the controllers are least sensitive to the yaw-axis controller gains because of the spacecraft's aerodynamic properties, it is likely that decreasing the weight of the yaw-axis performance in the GA improves overall controller stability.

Besides the simulated controller being more active on the yaw-axis than the GOCE spacecraft appeared to have been, another apparent difference between the actual GOCE spacecraft and the simulated controller was in the time dependent control. The time dependent control matrix that is presented in literature contains a variable,  $\gamma$ . This variable can be chosen between 0 and 1, but no mention is made about what value was used in the GOCE spacecraft. The behaviour on the pitch-axis of the simulated and measured spacecraft above the equator are so different that it is likely that in the real GOCE mission the time dependent control matrix was either implemented differently than it is described in literature, or possibly a value for  $\gamma$  of larger than 1 was used. It is a shame that the internals of the actual GOCE controller are not publicised, because it would have been interesting to investigate the source of the observed differences.

Finally, It was shown that the NINDI-controller is able to still perform adequate when errors are introduced that are in line with what can be expected based on the hardware present in the GOCE spacecraft. However, when errors are introduced further, the NINDI-controller starts to perform worse than the PD-controller. The same holds for the LINDI-controller, given that the inner loop of both the LINDI- and NINDI-controller are similar. This confirms literature on INDI based control architectures needing accurate acceleration measurements. An investigation into the stability of the NINDI-controller when increasing atmospheric winds and density was also done. This showed that the controller is able to control the spacecraft under demanding conditions as well. However, during this investigation the increase in the IPA corresponding to an increase in drag force was not taken into account. Combine this with the low sensitivity of the NINDI-controller to gain changes, and it can be concluded that NINDI-control is a powerful tool for attitude control, and can be successfully applied to control a spacecraft magnetically in demanding operating conditions.



## CONCLUSION AND RECOMMENDATIONS

The goal of the thesis was twofold, first first goal was to validate the feasibility of a NINDI-controller for a GOCE-type spacecraft, and the second goal was to investigate whether controller performance improves when a linear acceleration tracking control frame is used. First, to be able to answer the first main research question of: *What is the overall performance of an INDI-type control architecture for the magnetically controlled attitude of the GOCE satellite?*, four questions were proposed. In order to answer these questions three controller have been set-up. A PD-controller that is used as a reference controller, an INDI controller with a linear outer loop, or a LINDI-controller, and an INDI controller with an NDI outer loop, or a NINDI-controller. The NINDI-controller was the main controller under investigation during this thesis, because it is mentioned as a promising controller in literature.

1) *What is the performance of the controller with respect to the ESA requirements; i.e. the maximum values for attitude, angular rate, and angular acceleration?* Both ESA's time and frequency domain requirements were met by manually reducing the inner loop gains. It is expected that by taking the frequency domain requirements into account in the GA fitness function, feasible gains can be found automatically. More importantly, the time domain requirements were met with one single set of gains over different simulation periods in the GOCE mission, proving that a NINDI-controller is more flexible than a conventional PD-controller.

2) *How do the simulated disturbance torques, control torque, and net body torque compare with the measured and modelled torques from the GOCE satellite during equal orbits?* When compared to measurements of the GOCE spacecraft, the simulated controller was able to perform equal or better in most cases. Although the NINDI-controller periodically commands high frequency control actuation to make rapid changes in the acceleration, these events occurred mostly around the equator. The fact that these events occurred near the equator is due to have three main reasons. Firstly, time dependent control suspends or reduces pitch-axis control around the magnetic equator, which quickly ramps up again after crossing the equator. Secondly, the controller is first recovering from having passed one of the poles, while the velocity of the co-rotating atmosphere increases as the spacecraft moves towards the equator, and yaw-axis control is most effective above the magnetic equator. Finally due to the magnetic coupling of the roll- and yaw-axis, and the comparably low inertia of the roll-axis, the roll-axis showed spikes in acceleration around the equator. It is expected that by reducing the inner loop control aggressiveness, and the weight on the yaw-axis in the GA, these rapid changes around the equator will diminish.

3) *What is the change in controller performance when measurement errors, and increases in thermospheric properties like winds and air density are introduced?* Introducing errors in the simulation has shown that the NINDI-controller is able to still produce satisfactory results as long as the errors are in line with what can be expected based on the hardware present in the GOCE spacecraft. In fact, although the performance of the NINDI-controller decreases with the introduction of errors, it is still more stable than the reference PD-controller. Furthermore, the NINDI-controller performed well when the thermospheric winds and local atmospheric density were increased up to physically unrealistic values. The simulated state only occasionally went outside the limits of the ESA requirements in the time domain.

4) *How difficult is it to find a stabilising set of gains, and how sensitive is the controller to changes in gains?* It was shown that, as was expected based on theory about INDI-controllers, the NINDI-controller has a wider range of stabilising gains than the PD-controller. Moreover, the NINDI-controller produced more consistent results than the PD-controller when all the gains of the PD- and NINDI-controller were decreased and increased. This confirms that the NINDI-controller is a lot less sensitive to changes in gains than a conventional controller. The consistency of the NINDI-controller with respect to the PD-controller was also visible in the results of the simulation over different time periods for the controllers. Both the LINDI- and the NINDI-controller performed within the ESA requirements during all simulated periods, while the PD-controller performed adequate for the day it was optimised, but showed completely different results in other periods. An interesting example is a period in 2010, that was supposed to be an easy period for attitude control in terms of average spacecraft altitude and thermospheric properties. Due to the nature of PD-control and the gains of the PD-controller not being tuned for this day, it was unable to meet the requirements.

Combining the answers to these sub-questions leads to the conclusion that a NINDI-based controller is perfectly suited for fully magnetically actuated spacecraft attitude control. On top of that, a GOCE-type spacecraft is especially well suited for NINDI-based control, because of its accurate measurements and fine control actuation through the magnetorquers. The most important improvement that can be made to the work done in this thesis is to prove the frequency domain requirements can in fact be met while also still meeting the time domain requirements during different periods in the GOCE mission and using one set of gains. This has been shown to be achievable for one period, but it needs to be confirmed that it works in general. The suggested way to do this is to include the performance of the controller in the frequency domain in the fitness function of the GA, and use the GA to find a set of gains that accomplishes this. One way to incorporate the frequency domain requirements in the fitness function is by calculating the difference between the frequency domain requirements and the one-sided PSD at the end of a simulation for a few points along the MBW. These errors then need to be divided by the requirements to normalise them, similar to what is done for the attitude, angular rate, and angular acceleration errors. By doing it for multiple points along the MBW instead of along the entire MBW it is possible to put separate weights on certain sections along the MBW.

The second goal of the thesis was to investigate whether controller performance improves when a linear acceleration tracking control frame is used. The second main research question of this thesis was: *How does using a control frame that tracks the linear accelerations acting on the satellite affect the performance of the controller?* This linear acceleration tracking control frame was defined as the D-frame, but because it was found that there is likely an error in the lines of code that defines this reference frame, it was not possible to test this. However, it is an interesting reference frame to do control in, because it potentially allows to better compensate linear accelerations. Additionally, if the  $0^\circ$  yaw-axis attitude angle is in the direction of the linear acceleration then the yaw-axis disturbance torque will be minimised because the reference frame then not only tracks the co-rotating atmosphere, but also the thermospheric winds. This allows to reduce the yaw-axis gains further which likely improves roll-axis control, because of the coupling between the two axes. If one wants to continue with working on controlling in the D-frame it is advised to first exclude that the D-frame did not work because it is simply unable to track rapid changes in the thermospheric winds. Then eliminate the possibility that the controller simply needs to be retuned for use with the D-frame. If that also does not solve the problem then it is likely the case that there is indeed an error in the code.

The most exciting part about the results is that it shows that the NINDI-controller is very much capable of stabilising the spacecraft in a wide variety of operating ranges. Especially the performance of the controller when thermospheric winds and local densities were manually increased was impressive. This begs the question how the controller will have performed during the re-entry phase of the spacecraft, especially since it was found that the currents sent to the magnetorquers were not at the point of saturation yet. Unfortunately there was not enough time to investigate this during this thesis, but it is a very exciting research topic to be picked up by someone else. A second recommendation is that this thesis used high fidelity torque models to create a simulation environment that is as true to reality as possible. However, the linear accelerations acting on the spacecraft were ignored. Because these were ignored, the IPA thrust was taken to be equal to that of the actual GOCE spacecraft at the same point in time. While in reality the thrust should vary based on the drag encountered by the spacecraft, and the direction the spacecraft is pointing in. Thus a first step to further improve the simulation is to include a drag force and attitude dependent thrust. This will have consequences for the disturbance torque due to the IPA, which mainly affects the pitch-axis. Another question that is left unanswered in this thesis is how the LINDI- and NINDI-controller compare to each other. Although the LINDI-controller is used in this thesis, the controller under investigation was the NINDI-controller and



therefore an appropriate comparison has not been made. One would expect, based on the NDI outer loop, that the NINDI-controller has a wider range of stabilising gains for the GOCE attitude control problem than the LINDI-controller. This is for someone else to investigate further. Finally, given the effect that the GOCE spacecraft's shape and aerodynamic properties had on the stability of the spacecraft. It is worth the effort to see how a NINDI-controller in a spacecraft without these stabilising properties would behave and whether a NINDI-controller is still able to successfully control the attitude of a spacecraft magnetically in the demanding environment of VLEO.




# REFERENCES

- [1] P. Acquatella, W. Falkena, E.-J. van Kampen, and Q. P. Chu. Robust Nonlinear Spacecraft Attitude Control Using Incremental Nonlinear Dynamic Inversion. In *AIAA Guidance, Navigation, and Control Conference*. American Institute of Aeronautics and Astronautics, August 2012. doi: 10.2514/6.2012-4623.
- [2] B. Bacon and A. Ostroff. Reconfigurable Flight Control Using Nonlinear Dynamic Inversion with a Special Accelerometer Implementation. In *AIAA Guidance, Navigation, and Control Conference and Exhibit*, page 4565, 2000.
- [3] E. Canuto and L. Massotti. All-Propulsion Design of the Drag-Free and Attitude Control of the European Satellite GOCE. *Acta Astronautica*, 64(2-3):325–344, January 2009. doi: 10.1016/j.actaastro.2008.07.017.
- [4] Q. P. Chu. Spacecraft Attitude Dynamics and Control - Course Notes No. 1 Part 2. Delft University of Technology: Lecture slides of the course Spacecraft Attitude Dynamics and Control, February 2017.
- [5] F. Cometto. GOCE Models Data Bank. Technical Report GO-TN-AI-0081\_6, Alenia Spazio, April 2007.
- [6] B. DeKock, D. Sanders, T. Vanzwieten, and P. Capó-Lugo. Design and Integration of an All-Magnetic Attitude Control System for FASTSAT-HSV01’s Multiple Pointing Objectives. In *34th Annual Guidance and Control Conference*, February 2011. URL <https://ntrs.nasa.gov/search.jsp?R=20110009916>.
- [7] J. Diebel. Representing Attitude: Euler Angles, Unit Quaternions, and Rotation Vectors. *Matrix*, 58(15-16):1–35, 2006.
- [8] E. N. Doornbos. *Thermospheric Density and Wind Determination from Satellite Dynamics*. phdthesis, Technical University Delft, March 2011.
- [9] M. R. Drinkwater, R. Haagmans, D. Muzi, A. Popescu, R. Floberghagen, M. Kern, and M. Fehringer. The GOCE Gravity Mission: ESA’s First Core Earth Explorer. In *Proceedings of the 3rd international GOCE user workshop*, pages 6–8. European Space Agency Noordwijk, The Netherlands, 2006. URL [https://www.researchgate.net/profile/Rune\\_Floberghagen/publication/237628788\\_The\\_GOCE\\_gravity\\_mission\\_ESA’s\\_first\\_core\\_explorer/links/00b7d5275473595230000000.pdf](https://www.researchgate.net/profile/Rune_Floberghagen/publication/237628788_The_GOCE_gravity_mission_ESA’s_first_core_explorer/links/00b7d5275473595230000000.pdf).
- [10] R. Floberghagen, M. Fehringer, D. Lamarre, D. Muzi, B. Frommknecht, C. Steiger, J. Piñeiro, and A. da Costa. Mission Design, Operation and Exploitation of the Gravity Field and Steady-State Ocean Circulation Explorer Mission. *Journal of Geodesy*, 85(11):749–758, October 2011. doi: 10.1007/s00190-011-0498-3.
- [11] C. E. Ghisi, C. Steiger, M. Romanazzo, and P. P. Emanuelli. Drag-Free Attitude and Orbit Control System Performance of ESA’s GOCE Mission During Low Orbit Operations and De-Orbiting. In *SpaceOps 2014 Conference*. American Institute of Aeronautics and Astronautics, May 2014. doi: 10.2514/6.2014-1906.
- [12] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley Pub Co Inc, 1989. ISBN 0201157675.
- [13] J. Horn, N. Nafpliotis, and D. E. Goldberg. A Niche Pareto Genetic Algorithm for Multiobjective Optimization. In *Proceedings of the first IEEE conference on evolutionary computation, IEEE world congress on computational intelligence*, volume 1, pages 82–87, 1994.
- [14] C. Y. Huang and G. J. Knowles. Application of Nonlinear Control Strategies to Aircraft at High Angle of Attack. In *Decision and Control, 1990, Proceedings of the 29th IEEE Conference on Decision and Control*, pages 188–193. IEEE, 1990.
- [15] Y. Huang, D. M. Pool, O. Stroosma, and Q. P. Chu. Incremental Nonlinear Dynamic Inversion Control for Hydraulic Hexapod Flight Simulator Motion Systems. *IFAC-PapersOnLine*, 50(1):4294–4299, July 2017. doi: 10.1016/j.ifacol.2017.08.837.

- [16] L. J. Kamm. Magnetorquer - a Satellite Orientation Device. *ARS Journal*, 31(6):813–815, June 1961. doi: 10.2514/8.5645.
- [17] A. Kolkmeier, G. Präger, P. Möller, T. Strandberg, K. Kempkens, J. Stark, L. Gessler, and K. O. Hienerwadel. GOCE - DFAC Interface Control Document. Technical Report GO-IC-ASG-0005\_12, EADS Astrium, April 2008.
- [18] M. Lovera and A. Astolfi. Global Magnetic Attitude Control of Inertially Pointing Spacecraft. *Journal of Guidance, Control, and Dynamics*, 28(5):1065–1072, September 2005. doi: 10.2514/1.11844.
- [19] M. Lovera and A. Astolfi. Global Magnetic Attitude Control of Spacecraft in the Presence of Gravity Gradient. *IEEE Transactions on Aerospace and Electronic Systems*, 42(3):796–805, July 2006. doi: 10.1109/taes.2006.248214.
- [20] R. Maltha. GOCE Attitude Controller Simulator: Literature Study. resreport, TU Delft, 2017. TU Delft Course AE4020.
- [21] F. Landis Markley and J. L. Crassidis. *Fundamentals of Spacecraft Attitude Determination and Control (Space Technology Library)*. Springer, 2014. ISBN 978-1-4939-0801-1.
- [22] R. Merzouki, B. O. Bouamama, P. M. Pathak, and A. K. Samantaray. *Intelligent Mechatronic Systems*. Springer London, 2012. ISBN 1447146271.
- [23] N. S. Nise. *Control Systems Engineering*. John Wiley and Sons Ltd, sixth edition, 2011. ISBN 0470646128.
- [24] J. B. Park, Y. M. Park, J. R. Won, and K. Y. Lee. An Improved Genetic Algorithm for Generation Expansion Planning. *IEEE Transactions on Power Systems*, 15(3):916–922, 2000. doi: 10.1109/59.871713.
- [25] M. L. Psiaki. Magnetic Torquer Attitude Control via Asymptotic Periodic Linear Quadratic Regulation. *Journal of Guidance, Control, and Dynamics*, 24(2):386–394, March 2001. doi: 10.2514/2.4723.
- [26] K. J. Åström and R. M. Murray. *Feedback Systems*. University Press Group Ltd, electronic edition version v2.10b edition, 2008. ISBN 0691135762. URL <http://www.cds.caltech.edu/~murray/amwiki>.
- [27] R. Roberson. Two Decades of Spacecraft Attitude Control. In *16th Aerospace Sciences Meeting*. American Institute of Aeronautics and Astronautics, January 1978. doi: 10.2514/6.1978-7.
- [28] C. M. Roithmayr. Contribution of Zonal Harmonics to Gravitational Moment. *Journal of Guidance, Control, and Dynamics*, 14(1):210–214, January 1991. doi: 10.2514/3.20626.
- [29] G. Rostoker. Geomagnetic Indices. *Reviews of Geophysics*, 10(4):935–950, 1972. doi: 10.1029/RG010i004p00935.
- [30] C. Schumacher and P. P. Khargonekar. Missile Autopilot Designs Using  $H_\infty$  Control With Gain Scheduling and Dynamic Inversion. *Journal of guidance, control, and dynamics*, 21(2):234–243, 1998.
- [31] G. Sechi, G. André, D. Andreis, and M. Saponara. Magnetic Attitude Control of the GOCE Satellite. In *Proceedings of the 6<sup>th</sup> International ESA Conference on Guidance, Navigation and Control Systems*, October 2005.
- [32] G. Sechi, M. Buonocore, F. Cometto, M. Saponara, A. Tramutola, B. Vinai, G. André, and M. Fehringer. In-Flight Results from the Drag-Free and Attitude Control of GOCE Satellite. In *Preprints of the 18th IFAC World Congress*, August 2011.
- [33] S. Sieberling, Q. P. Chu, and J. A. Mulder. Robust Flight Control Using Incremental Nonlinear Dynamic Inversion and Angular Acceleration Prediction. *Journal of Guidance, Control, and Dynamics*, 33(6):1732–1742, November 2010. doi: 10.2514/1.49978.
- [34] E. Silani and M. Lovera. Magnetic Spacecraft Attitude Control: a Survey and Some New Results. *Control Engineering Practice*, 13(3):357–371, March 2005. doi: 10.1016/j.conengprac.2003.12.017.
- [35] P. Simplício, M. D. Pavel, E. van Kampen, and Q. P. Chu. An Acceleration Measurements-Based Approach for Helicopter Nonlinear Flight Control Using Incremental Nonlinear Dynamic Inversion. *Control Engineering Practice*, 21(8):1065–1077, August 2013. doi: 10.1016/j.conengprac.2013.03.009.

- [36] C. Steiger, N. Mardle, and P. P. Emanuelli GOCE Flight Control Team (HSO-OEG). GOCE End-of-Mission Operations Report. Technical Report GO-RP-ESC-FS-6268, ESA, February 2014. URL <https://earth.esa.int/documents/10174/85857/2014-GOCE-Flight-Control-Team.pdf>. Prepared by: GOCE Flight Control Team (HSO-OEG).
- [37] H. Stoewer. Future of Space. Technical report, Space Associates GmbH, June 2018. URL <https://vsv.tudelft.nl/society/activities/future-spaceflight-prof-heinz-stoewer>. Presentation at the TU Delft.
- [38] C. S. Stummer. *Gradiometer Data Processing and Analysis for the GOCE Mission*. phdthesis, Technische Universität München, 2012.
- [39] R. Sutherland, I. Kolmanovsky, and A. R. Girard. Attitude Control of a 2U Cubesat by Magnetic and Air Drag Torques. *IEEE Transactions on Control Systems Technology*, pages 1–13, 2018. doi: 10.1109/tcst.2018.2791979.
- [40] M. Swartwout. The First One Hundred Cubesats: a Statistical Look. *Journal of Small Satellites*, 2(2): 213–233, 2013.
- [41] T. Visser, E. N. Doornbos, C. C. de Visser, P. N. A. M. Visser, and B. Fritsche. Torque Model Verification for the GOCE Satellite. *Advances in Space Research*, 62(5):1114–1136, September 2018. doi: 10.1016/j.asr.2018.06.025.
- [42] Y.-C. Wang, D. Sheu, and C.-E. Lin. A Unified Approach to Nonlinear Dynamic Inversion Control with Parameter Determination by Eigenvalue Assignment. *Mathematical Problems in Engineering*, 2015:1–13, 2015. doi: 10.1155/2015/548050.
- [43] D. Whitley. A Genetic Algorithm Tutorial. *Statistics and Computing*, 4(2), June 1994. doi: 10.1007/bf00175354.
- [44] B. Wie. *Space Vehicle Dynamics and Control (AIAA Education Series)*. AIAA, 2008. ISBN 978-1-56347-953-3.
- [45] B. Wie, H. Weiss, and A. Arapostathis. A Quaternion Feedback Regulator for Spacecraft Eigenaxis Rotations. In *Guidance, Navigation and Control Conference*. American Institute of Aeronautics and Astronautics, August 1988. doi: 10.2514/6.1988-4128.
- [46] Wikimedia Commons. Figure: Closed Loop PID Controller. Online, October 2008. URL [https://en.wikipedia.org/w/index.php?title=PID\\_controller&action=history](https://en.wikipedia.org/w/index.php?title=PID_controller&action=history). Accessed: 2017-10-23.
- [47] R. Wiśniewski and M. Blanke. Fully Magnetic Attitude Control for Spacecraft Subject to Gravity Gradient. *Automatica*, 35(7):1201–1214, July 1999. doi: 10.1016/s0005-1098(99)00021-7.





## STATE PLOTS

In this appendix similar plots to Figure 4.11 are shown. However, the plots shown here are for different periods, and include the measured state and simulated state of the three controllers, NINDI-, LINDI-, and PD-controller. The two periods that are shown here are 1 – 9 January 2010, and 16 – 23 May 2011. For both periods, all y-axes are scaled to match that of the measured state. The January 2010 period was chosen because it was early in the GOCE mission, in which it flew relatively high still and thus the atmospheric density was lower than later in the mission. Additionally this is an extremely quiet period with little thermospheric wind activity. Even though the January 2010 period is very calm, the PD-controller is struggling, emphasising how a conventional linear controller tuned for a specific environment (1 November 2011) can perform poorly when used in a different environment. The spikes seen in Figure A.1 in the LINDI-controller are caused because for that simulation older data was used that still contained processing errors in the measured quaternions. In Figure A.2 similar behaviour can be seen. The PD-controller performs poor, and the LINDI- and NINDI-controller perform similar or better than the measured state.

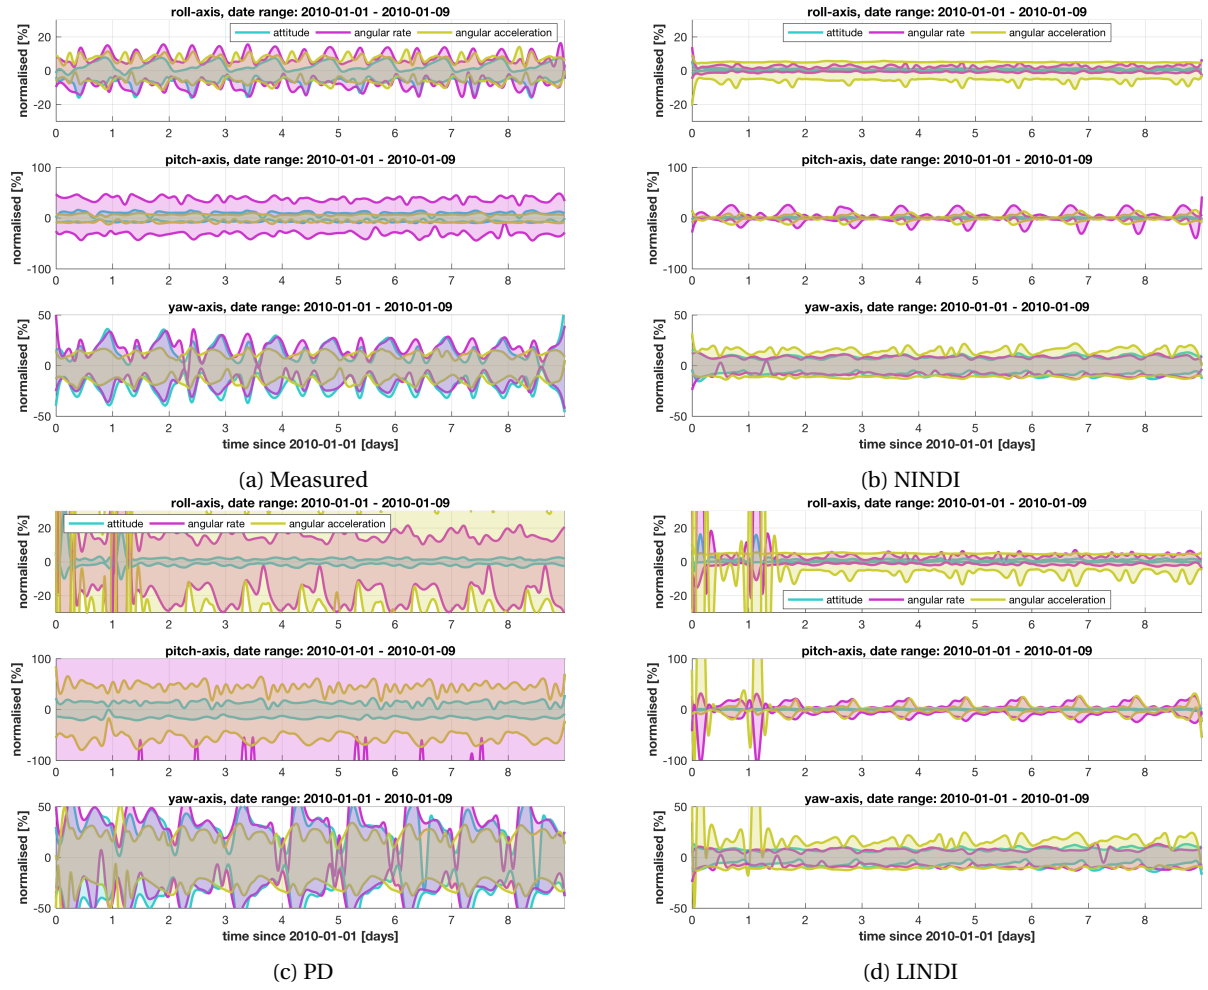


Figure A.1: Envelope plot of the simulated attitude angle, angular rate, and angular acceleration of the GOCE satellite for the period 1 – 9 January 2010. All values have been normalised to the ESA requirements. All y-axes have been scaled to be equal to the y-axis of the envelope plot of the measured state (a). Simulation was done in the 'O-frame'. (b) is the NINDI-controller, (c) is the reference PD-controller, and (d) is the LINDI-controller.



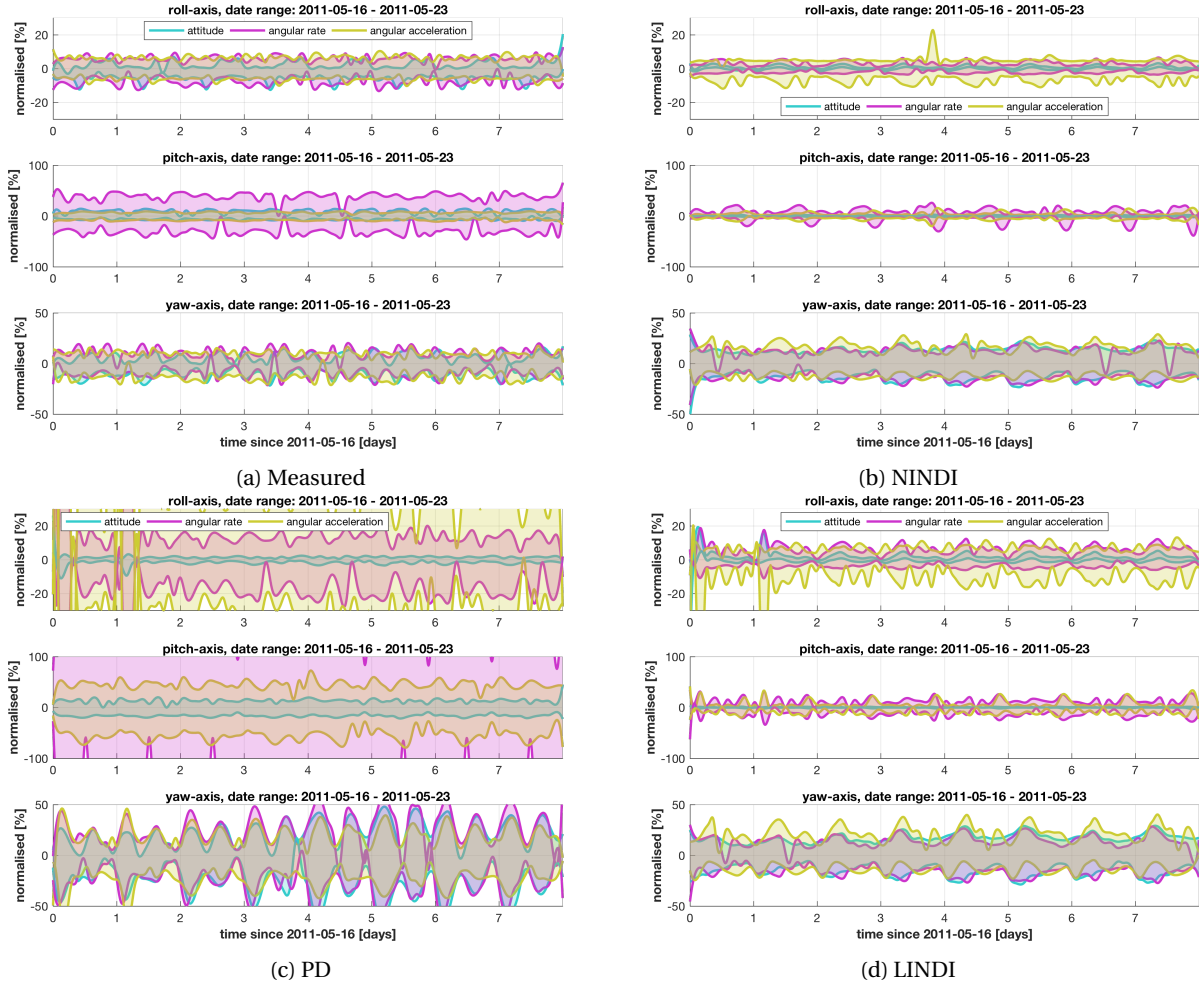


Figure A.2: Envelope plot of the simulated attitude angle, angular rate, and angular acceleration of the GOCE satellite for the period 16 – 21 May 2011. All values have been normalised to the ESA requirements. All y-axes have been scaled to be equal to the y-axis of the envelope plot of the measured state (a). Simulation was done in the ‘O-frame’. (b) is the NINDI-controller, (c) is the reference PD-controller, and (d) is the LINDI-controller.



# B

## DIFFERENCES BETWEEN MEASURED AND SIMULATED

### B.1. Attitude, angular rate, and angular acceleration differences

This section shows the data that is used in Figure 4.15 to Figure 4.17, but in this case the measured and simulated state are shown separately.

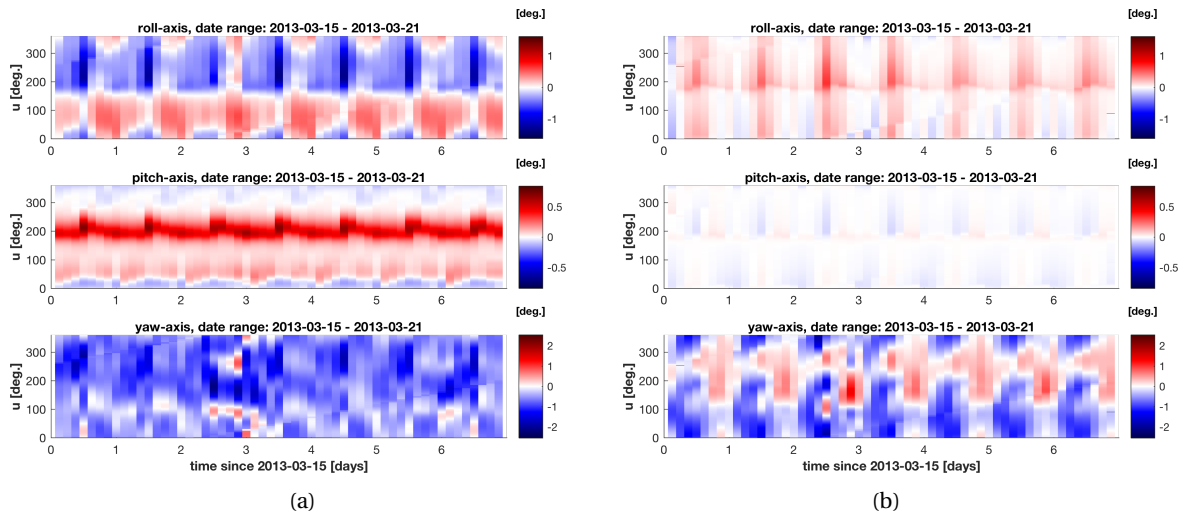


Figure B.1: Argument of latitude plot of the attitude of the spacecraft. Colour bar scale is set to the maximum absolute values of the GOCE measured attitude on each axis. (a) Shows the measured attitude. (b) Shows the simulated attitude.

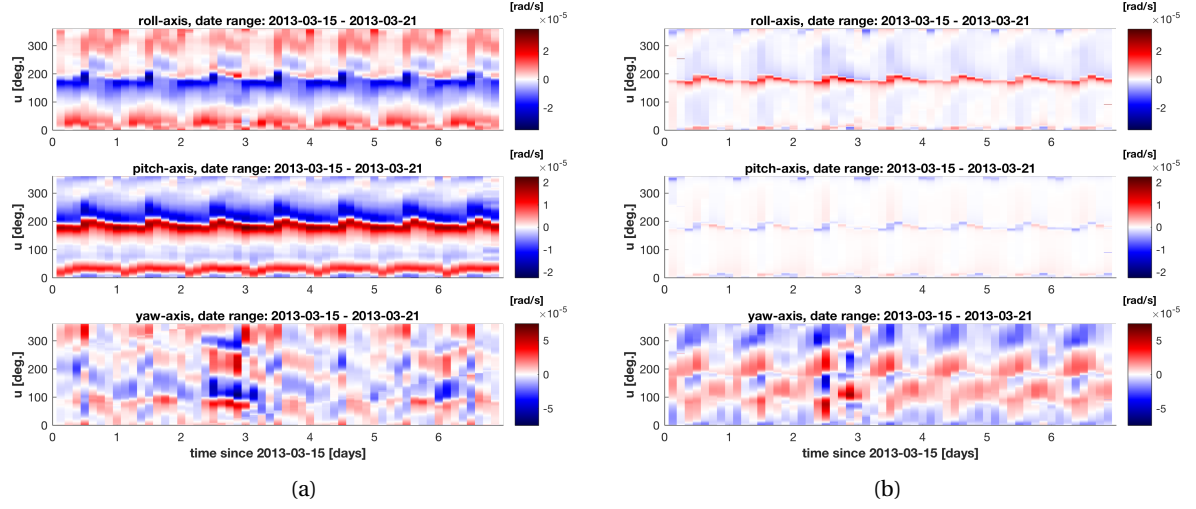


Figure B.2: Argument of latitude plot of the angular rate of the spacecraft. Colour bar scale is set to the maximum absolute values of the GOCE measured angular rate on each axis. (a) Shows the measured angular rate. (b) Shows the simulated angular rate.

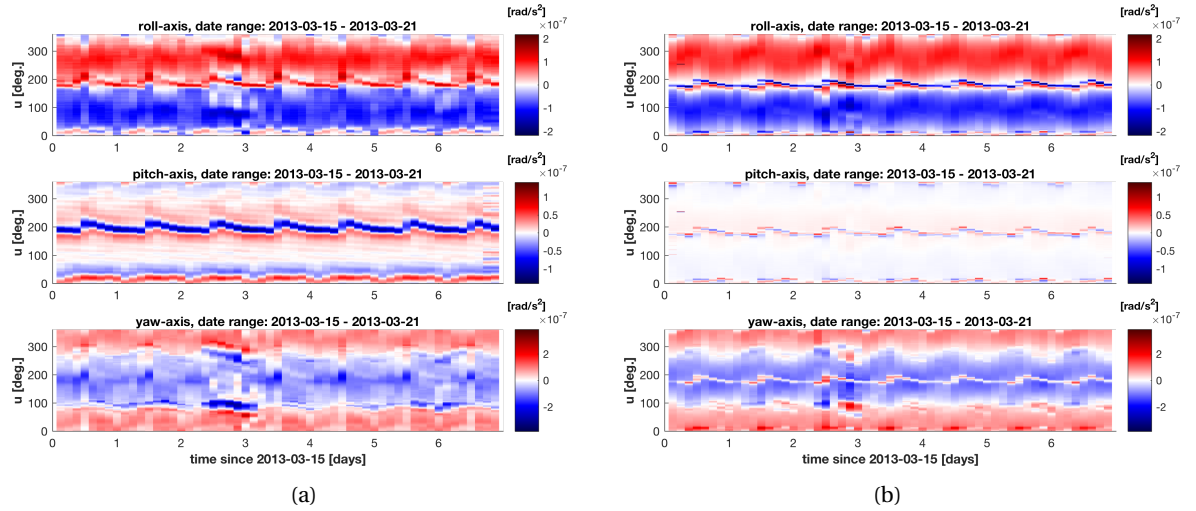


Figure B.3: Argument of latitude plot of the angular acceleration of the spacecraft. Colour bar scale is set to the maximum absolute values of the GOCE measured angular acceleration on each axis. (a) Shows the measured angular acceleration. (b) Shows the simulated angular acceleration.

## B.2. Control torque differences

This section shows the data that is used in Figure 4.18 to Figure 4.20, but in this case the measured and simulated state are shown separately.

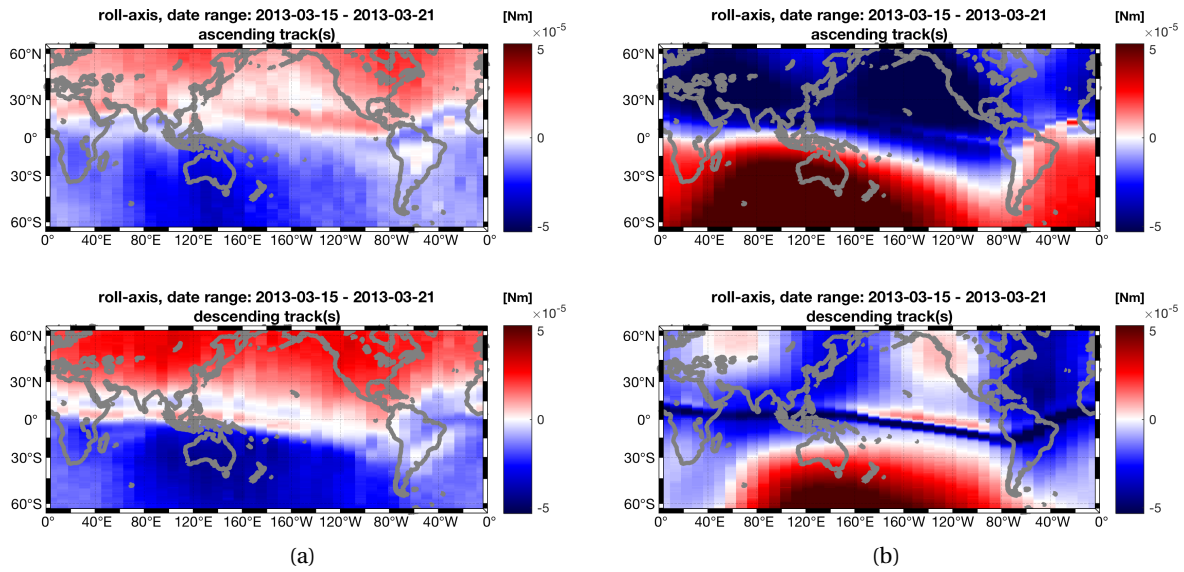


Figure B.4: Miller projection map plot around the equator for the measured (a), and simulated (b), achieved control torque on the roll-axis. The colour bar scale is set to the maximum absolute value of the GOCE measured achieved control torque.

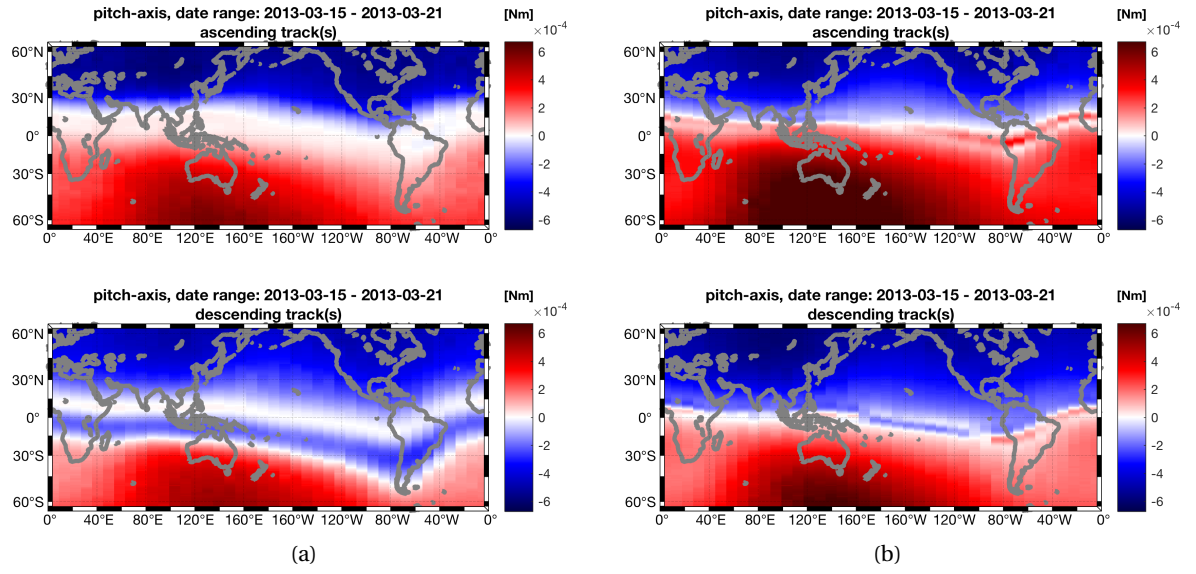


Figure B.5: Miller projection map plot around the equator for the measured (a), and simulated (b), achieved control torque on the pitch-axis. The colour bar scale is set to the maximum absolute value of the GOCE measured achieved control torque.

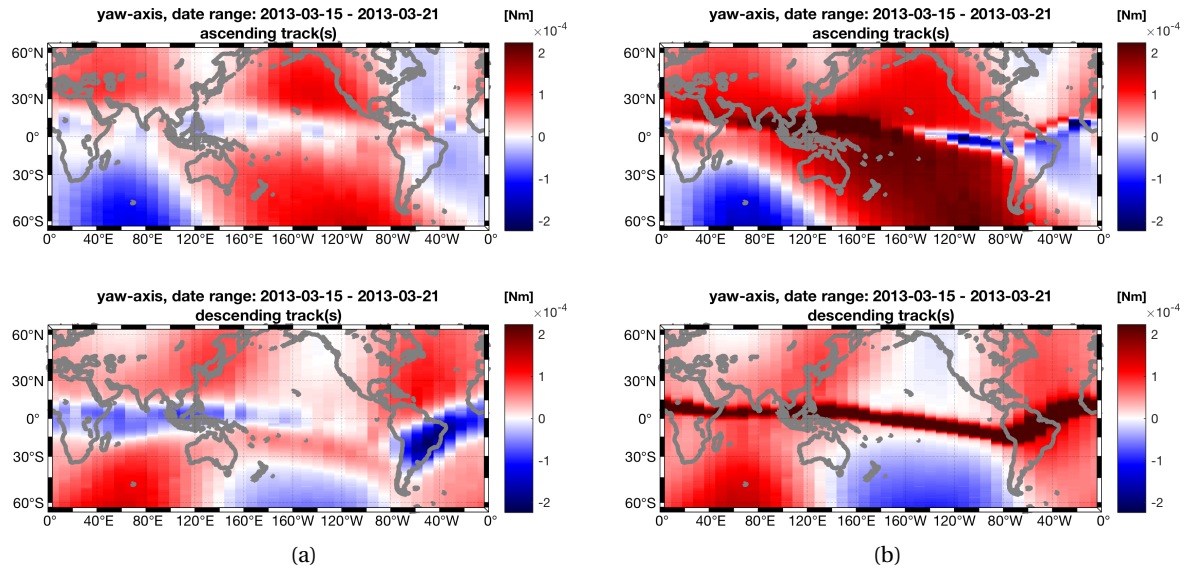


Figure B.6: Miller projection map plot around the equator for the measured (a), and simulated (b), achieved control torque on the yaw-axis. The colour bar scale is set to the maximum absolute value of the GOCE measured achieved control torque.

# C

## ALL DISTURBANCE TORQUES FOR 16 – 21 MARCH 2013

This appendix shows all the components of the total disturbance torque as they were simulated during the period of 16 – 21 March 2013. The spikes that are visible on all axes at the start of the time series, and on the roll- and pitch-axis at the end of the time series, are caused by a data processing error. Normally this would be a narrow vertical line, but due to the nature of an envelope plot the narrow line is stretched out somewhat.

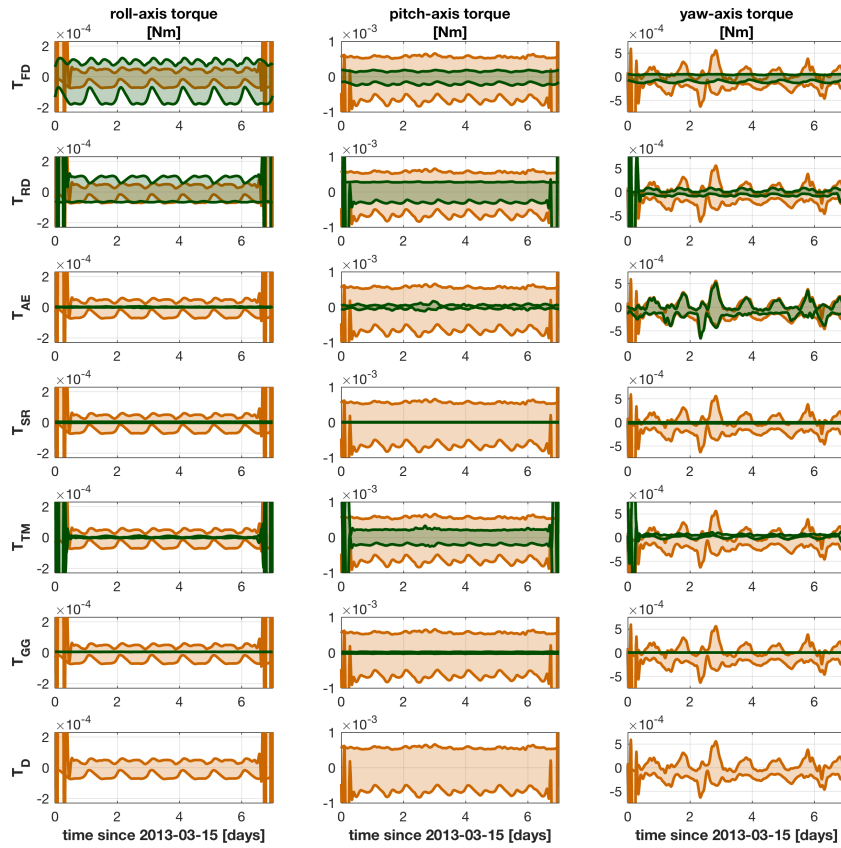


Figure C.1: Components of total disturbance torque, and total disturbance torque itself. Components shown here are the fitted dipole torque,  $T_{FD}$ , the residual dipole torque,  $T_{RD}$ , the aerodynamic torque,  $T_{AE}$ , the solar radiation pressure torque,  $T_{SR}$ , The torque due to the thruster  $T_{TM}$ , the gravity gradient torque  $T_{GG}$ , and the total disturbance torque,  $T_D$ . The total disturbance torque (orange) is repeated in the component plots as well, to see how they relate to each other.





# D

## TIME SERIES STATE FOR 31 MAY TO 2 JUNE 2013

This appendix shows the geomagnetic indices, as well as the simulated state for the storm period that is used to compare the manually increased atmospheric properties with.

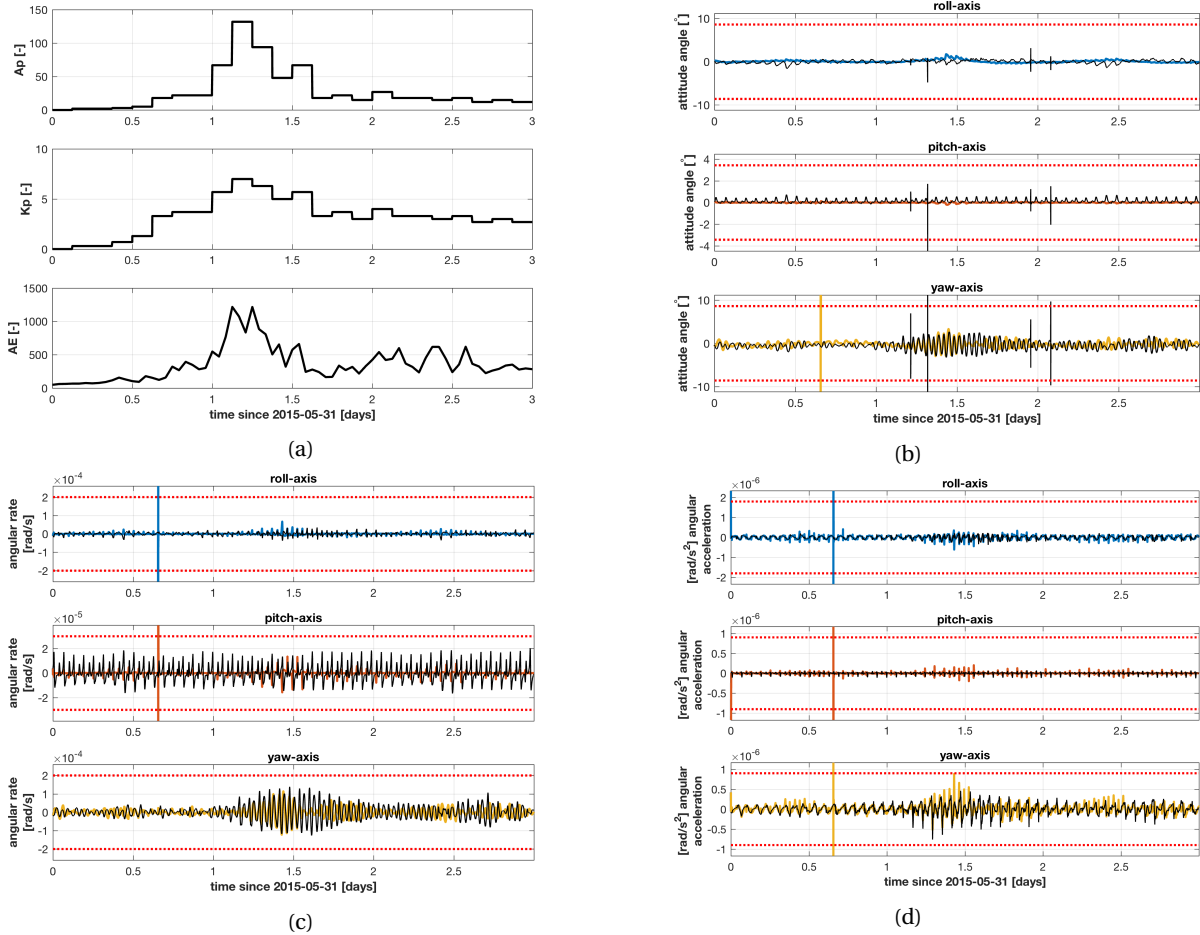


Figure D.1: Summary of the actual and simulated spacecraft behaviour during the period 31 May to 2 June 2013, with a mean altitude of 236 kilometres. Geomagnetic indices shown in (a), the storm peaks shortly after the start of the second day. (b) shows the simulated and measured (black) spacecraft attitude, (c) shows the angular rate, and (d) the angular acceleration. The red dashed lines in (b), (c), and (d) are the requirement limits.