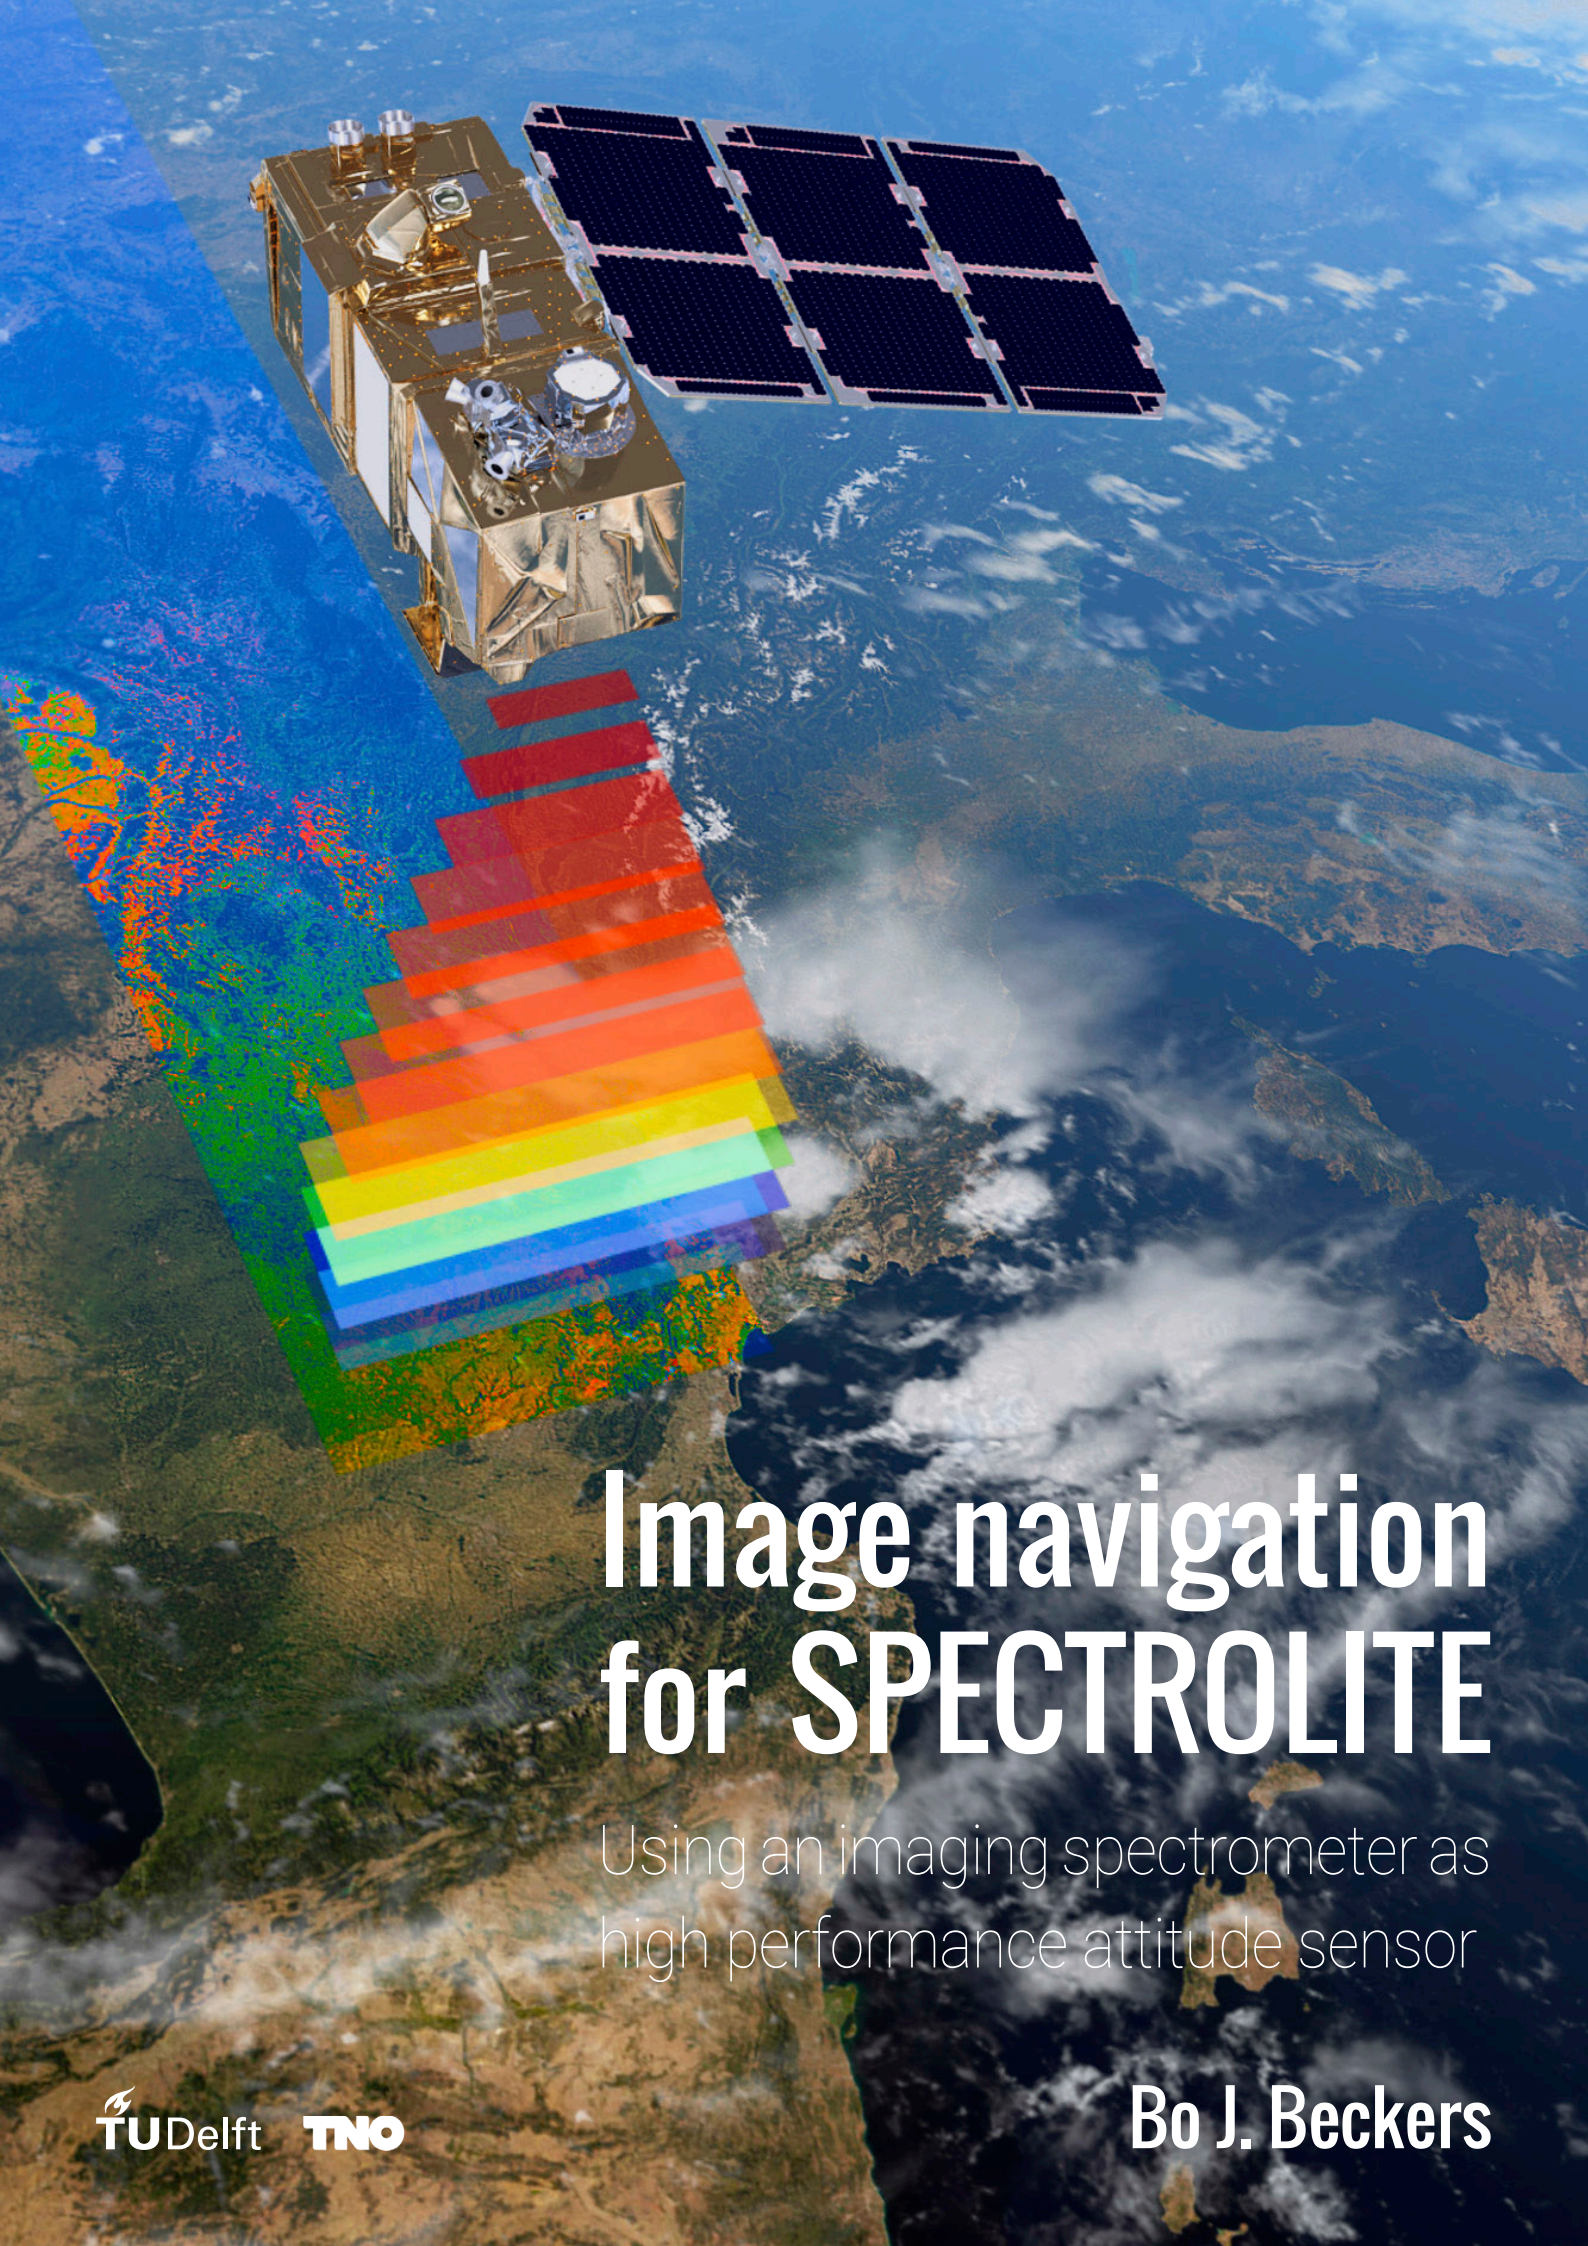

Image navigation for SPECTROLITE

Using an imaging spectrometer as high performance attitude sensor

TU Delft TNO

Bo J. Beckers

Image navigation for SPECTROLITE

Using an imaging spectrometer as high performance attitude sensor

by

Bo J. Beckers

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday August 16, 2017 at 13:00.

Student number:	4098315	
Project duration:	January 2, 2017 – August 3, 2017	
Thesis committee:	Dr. A. (Angelo) Cervone	TU Delft, Space systems engineering, Chairman
	Dr. ir. J. M. (Hans) Kuiper	TU Delft, Space systems engineering, Supervisor
	Ir. B. T. G. (Bryan) de Goeij	TNO, External supervisor
	Ir. B. C. (Bart) Root	TU Delft, Astrodynamics and space missions

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

Cover image courtesy of ESA.

TU Delft

Preface

Before you lies my thesis on the subject of image navigation; using captured Earth imagery for satellite navigation purposes. The research described in this document has been conducted within the context of the SPECTROLITE mission at the Netherlands Organisation for Applied Scientific Research (TNO).

Performing my thesis research has been a challenging, but fulfilling ride. Having no prior knowledge in the field of image navigation, or image processing in general, meant I had a great deal to learn. I remember having little faith in the possibility of designing an image navigation system to the requirements set for the SPECTROLITE mission. Learning more about the complexity of image navigation and the registration of satellite imagery, only meant additional challenges started to pop up. However, slowly I was able to start applying my newly acquired knowledge and start tackling these problems one by one. This process has immensely motivated me throughout the project and made performing the research into a very rewarding experience.

I would like to use this opportunity to thank TNO for providing this great research opportunity. Being able to perform research within the context of a real project made it all the more interesting. Furthermore, the TNO office was not simply a place to perform my research, but working there also allowed me to get a glimpse into the next phase of my life. It provided a great place to talk to fellow students, PhD candidates and other colleagues about their work, work-life and future plans. I want to express my sincere gratitude to Bryan de Goeij, my supervisor at TNO, for his support throughout my thesis project. Bryan, I want to thank you not only for the interesting discussions we had and great feedback you provided, but also for the way you helped me learn more about myself. Furthermore, I want to thank Hans Kuiper, my supervisor at the TU Delft, for his critical eye and constructive criticism. Looking back, I'm grateful for you giving me the freedom to approach my research in my own way, which has allowed me to make the most of this learning experience.

Of course I also want to thank my parents for their support, not only during my time in Delft, but during my entire education. Learning, in the broadest sense of the word, has always been a subject close to your hearts and you have played a huge part in allowing me to get to where I am now. Lastly, I want to thank my girlfriend Judit. You have really helped me stay on track during the last part of my thesis work, especially when summer came and the holidays had started. Thank you for being with me throughout this period and supporting me in so many ways.

Bo J. Beckers
Delft, August 2017

Summary

In recent years, efforts to lower the cost of space assets have resulted in a trend towards the development of ever smaller satellites. This trend has led to miniaturisation of all satellite subsystems and components, including payloads. Following this trend, TNO has developed SPECTROLITE, a new miniaturised spectrometer instrument for the measurement of trace gasses in the Earth's atmosphere at a spatial resolution of one kilometre. This instrument is to fly on board a small low-cost satellite with a total mass of under 25 kilograms. TNO envisions the creation of a constellation of these satellites, offering unprecedented temporal resolution and providing near real-time information on the concentration of greenhouse gasses in the atmosphere.

However, miniaturisation of the instrument alone is not sufficient to allow the design of this small Earth observation satellite. One of the main design challenges for a small Earth observation satellite lays in the design of its attitude determination and control system (ADCS). In this field too, miniaturisation has led to significant reductions in mass and cost of components. However, the ADCS remains one of the most costly satellite sub-systems, putting significant strain on satellite budgets. Additional miniaturisation of attitude sensors without the loss of performance is starting to be limited by the laws of physics. This is especially true for star trackers, currently the reference sensor for meeting high attitude determination requirements. TNO has opted for a different path to meet the high attitude determination requirement of the SPECTROLITE instrument under stringent mass constraints. Together with Delft University of Technology they have started the research on an image navigation system capable of simultaneously using the SPECTROLITE payload instrument as a high performance attitude sensor. This would limit the need of carrying additional attitude sensors on board the satellite. Use of such an image navigation system could increase attitude determination accuracy currently achievable by small Earth observation satellites, while at the same time reducing the cost and especially mass of their ADCS.

The SPECTROLITE imaging spectrometer, which is of the push-broom type, captures images at a rate of around 7 Hz in one spatial dimension. For each spatial element, the intensity of the reflected sunlight is measured as a function of wavelength, allowing the concentrations of selected trace-gasses to be determined. For use in the image navigation process, the intensity from the entire instrument wavelength spectrum is averaged to create a single grayscale image line. Propagation of previous satellite states using a state augmented Kalman filter is then used to coarsely determine the attitude corresponding to the captured image line. Based on this initial attitude estimate, an image correlation method is used to register the captured image line to an onboard database of geo-located reference imagery. The registration is used to correct the initial attitude estimate supplied by the Kalman filter and creates a new attitude measurement.

To limit non-rigidity effects due to surface elevation variations in the captured scene, data from a digital elevation model is used to compute an estimate of the coordinates of individual pixels in the captured image line, based on the initial attitude estimate. Although the error in the initial attitude estimation causes an offset in the computed image line coordinates, the error in the relative location between pixels is small. Creating reference image lines based on the computed image line coordinates allows use of rigid registration method, simplifying the registration task.

Still, the computational cost of performing the image registration task can be significant. To allow the image registration to be performed in real-time on board the SPECTROLITE satellite, a new computational efficient registration algorithm was designed. This algorithm is based on the incremental sign distance correlation method. This method creates a binary series based on a positive or negative change in the intensity between neighbouring pixels. Since it only uses relative pixel intensities between neighbouring pixels, this method allow registration of imagery captured with different instruments under varying lighting conditions. To reduce the number of correlations needed to measure the roll, pitch and yaw angle, three separate one-dimensional registrations are performed on three different

areas of the captured image line. Two along-track registrations, one at each edge of the image line, are used to compute pitch and yaw angles. A third cross-track registration at the centre of the image line is used to compute the roll angle.

Simulations have been performed using imagery from the MODIS imaging spectrometer to mimic the future SPECTROLITE imagery. The extensive database of Landsat images has been used to create a database of reference imagery needed for the registration process. These simulations show the one-dimensional registration method is capable of registering individual image lines with accuracy (1σ) of under 0.15 the image resolution in cross-track direction and under 0.20 times the image resolution in along-track directions. This allows a roll, pitch and yaw angle measurement accuracy (1σ) of respectively 0.014, 0.016 and 0.022 degrees. The high measurement frequency and excellent measurement trueness allow the Kalman filter to improve the real-time available attitude knowledge with respect to the attitude measurement accuracy. Combined with rotational rate measurements from an onboard gyroscope, filtering improves the real-time attitude determination accuracy to 0.0073, 0.0076 and 0.012 degrees for the roll, pitch and yaw angle, respectively.

By tuning of registration algorithm parameters and trading registration accuracy for a decrease in computational cost, an image navigation system for use on board the SPECTROLITE satellite was designed. This system uses currently available commercial off-the-shelf components to meet the SPECTROLITE attitude determination requirements. It provides a real-time attitude determination accuracy of 0.0082, 0.0082 and 0.013 degrees for roll, pitch and yaw angles, respectively. The image processing is performed by three ARM based processing units with a combined mass of only 84 grams. The combined attitude determination system components have a mass of only 236 grams and a cost not exceeding €50,000. This shows an image navigation for SPECTROLITE can indeed be feasible and provide mass and cost benefits over a more traditional star-tracker based system.

Contents

List of Figures

List of Tables

1

Introduction

In recent years, efforts to lower the cost of space assets have resulted in a trend towards the development of ever smaller satellites. This has led to miniaturisation of all satellite subsystems and components, including payloads. Following this trend, TNO is currently developing SPECTROLITE, a new miniaturised spectrometer instrument for the measurement of trace gasses in the Earth's atmosphere. This instrument is to fly on-board a small low-cost satellite, allowing a constellation of satellites to be launched, providing near real-time information on the concentration of greenhouse gasses in the atmosphere with relatively high spatial resolution.

Determining atmospheric particle concentrations from SPECTROLITE measurements with a high spatial resolution requires accurate knowledge of the satellite's attitude during the measurement. This leads to requirements on the attitude determination accuracy previously only achievable by larger heavier satellites. These requirements mean the SPECTROLITE satellite will have to carry a set of high accuracy attitude determination sensors, putting significant strain on the satellite's mass and cost budgets.

To limit the mass and cost of the SPECTROLITE satellite, TNO wants to investigate the possibility of simultaneously using their SPECTROLITE spectrometer instrument as a high performance attitude sensor, thereby eliminating the need of carrying traditional attitude sensor on board. Using Earth observation imagery to determine a satellites location and attitude is a process called satellite image navigation. Although this idea exists in theory, practical implementation has proven to be difficult, mainly due to the significant amount of processing power required to process the imagery.

Research objective
This research project is a first feasibility study into the use of an image-based attitude determination system on board the SPECTROLITE satellite. It aims to assess this feasibility by investigating whether an image navigation system can meet the attitude determination requirements of the SPECTROLITE satellite. Furthermore, it will investigate how the properties of such an image navigation system compare to those of a traditional star-tracker attitude determination system and determine possible advantages and disadvantages.

Structure
This report is divided into four main parts. Part one will focus on the design of an image navigation concept and the algorithms used to measure satellite attitude from captured imagery. This process starts with a description of the SPECTROLITE instrument, satellite and mission in chapter two. Chapter three will focus on the requirements placed on the SPECTROLITE attitude determination system. We will identify custom needs and requirements, and from it we will determine a set of attitude determination system requirements. Chapter three will focus on concept design. First, the overall attitude determination system concept will be described. Then, a concept for the image-based attitude determination

method is selected. Based on this selected concept, the required image navigation algorithms are designed in chapter five.

Part two of the report focuses on verifying the design of the image-based attitude determination method and determining its achievable performance. Chapter six describes the image registration simulation used to verify the attitude measurement method and finish its design. The results of this simulation are discussed in chapter seven. Chapter eight discusses the design of the image navigation simulation, where the image registration simulation is combined with a simulation of the satellite's attitude to determine the achievable attitude determination accuracy. The results of this simulation are discussed in chapter nine.

Part three focuses on finalising the design of the image-based attitude determination for SPECTROLITE, determining its final performance and assessing its feasibility. In chapter ten, the attitude determination system hardware is selected to determine the physical properties of this system. In chapter eleven, the final performance of this system is analysed and its system properties are compared to the current star-tracker system designed for SPECTROLITE, allowing assessment of the feasibility of an image navigation system.

Part four contains the conclusion and recommendations, available in chapters twelve and thirteen, respectively.

Part I

Concept and algorithm design

2

SPECTROLITE mission concept and figures

TNO, the Netherlands Organisation for Applied Scientific Research, has a long history of developing Earth observation instruments, like spectrometers for the measurement of trace gasses in the atmosphere. SPECTROLITE is to be the newest member in a line of Dutch designed space instruments, which includes instruments like SCIAMACHY, OMI and TROPOMI [49].

2.1. The SPECTROLITE mission

The SPECTROLITE instrument is designed to measure the concentrations of air pollutants in the atmosphere, like NO_2 and SO_x, with a relatively high spatial resolution of 1x1 km. A leap forward compared to the TROPOMI spatial resolution of 7x7 km [48].

2.1.1. Measuring greenhouse gasses from space

Trace gas concentrations in the atmosphere can be determined by pointing an imaging spectrometer at the Earth and measuring the intensity of the reflected sunlight as a function of wavelength. The optics inside a spectrometer disperse the light into different wavelengths before it reaches the imaging sensor, as visualised in Fig. 2.1. This allows the sensor to measure the intensity of small ranges of wavelength for each spatial ground element.

When plotting the intensity versus wavelength relationship for a spatial ground element, a figure like the one in Fig. 2.2 is created. In this figure, clear intensity peaks can be identified. Different particles in the atmosphere absorb or reflect light at their own specific wavelengths. Since the specific wavelengths at which certain gasses absorb or reflect light is known, each peak or valley identified can be matched to an accompanying trace gas [34].

2.1.2. Novel instrument approach

Typically, TNO developed space-born spectrometers like SCIAMACHY, OMI and TROPOMI have masses on the order of a hundred kilograms [49] [48]. This makes both the development of the instrument, but also the development, production and launch of the carrying satellite, very costly. The development and production of the TROPOMI instrument, without accompanying satellite, was budgeted at 78 million euros [13].

Recent advancements in optics, especially the development of freeform mirrors, have opened the way for smaller and cheaper instruments with sufficient performance for the trace-gas measurement task. With SPECTROLITE, TNO aims to surpass the spatial resolution of previous instruments at an

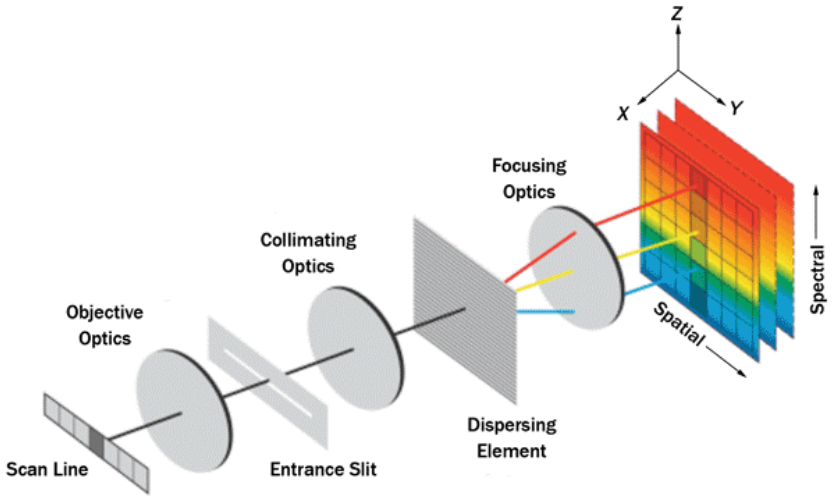

Figure 2.1: Schematic representation of the spectrometer components [11].

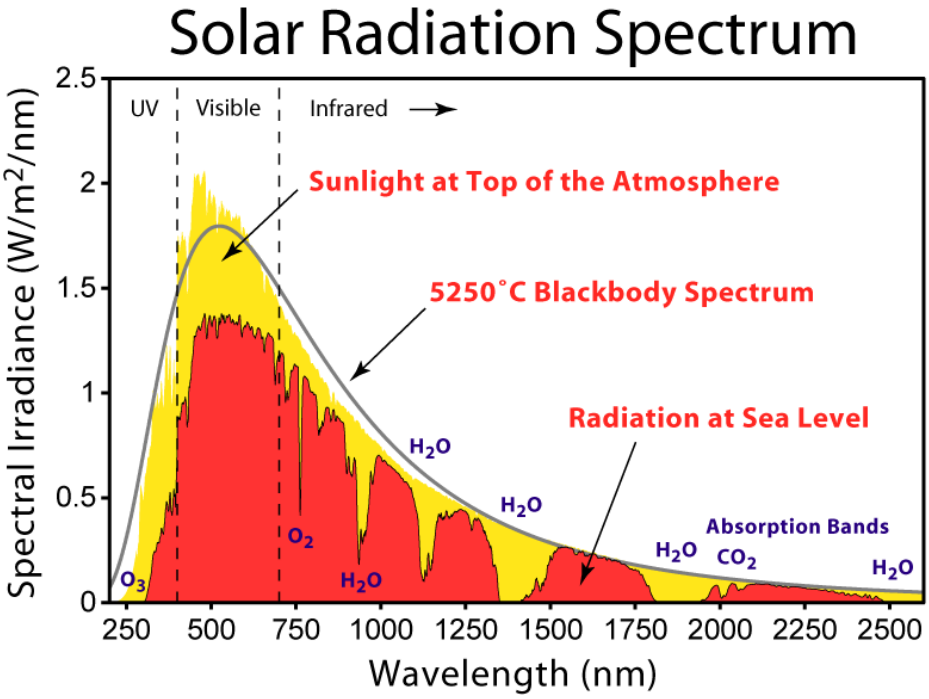

Figure 2.2: Solar radiation spectrum above and below atmosphere [5].

instrument mass of less than 8 kilograms. Preliminary specifications for the SPECTROLITE instrument can be found in Table 2.1.

2.1.3. Measurement constellation

The low recurring cost of the SPECTROLITE instrument could allow the creation of a constellation of Earth observation satellites measuring greenhouse emissions in near real-time. However, creating such a constellation will also require a low recurring cost for the carrying satellite. TNO and its partners therefore aim to develop such a satellite with a recurring cost below 5 million euros and a mass below 25 kilograms. To do so, the satellite will be sized according to the CubeSat standard. Total size will be a 12 unit CubeSat, measuring around 20 by 20 by 30 centimetres. A volume of 10 by 20 by 20 centimetres will be allocated to the SPECTROLITE instrument itself, limiting the available volume for other satellite subsystem. This means both size and mass of all satellite subsystems must be reduced compared to traditional satellites. Preliminary SPECTROLITE satellite specifications are summarised in Table 2.2. A total of nine satellite is envisioned, providing trace gas measurements with an unprecedented temporal resolution.

The focus of the measurement constellation is placed on Earth's most polluting regions: North-America, Europe and East-Asia. The goal is to cover all major cities with latitudes between 29° north and 56°. This means the observation region is limited to an area much smaller than the entire globe.

2.2. Satellite and instrument specifications

In the design of the attitude determination system, the current design specifications of both the SPECTROLITE instrument and satellite will be used. Note, however, that these specifications can still be subject to change as the design progresses.

2.2.1. Instrument specifications

The design specifications of the instrument are tabulated in Table 2.1. Recently, flight-tests with a prototype instrument have been conducted, the results of which assure the feasibility of these specifications according to TNO.

Table 2.1: SPECTROLITE instrument specifications

Specification	Value
Instrument mass	<8 kg
Instrument volume	10 x 20 x 20 cm (4 CubeSat units)
Instrument swath	900 km
Detector pixel size	20 μm
Ground sampling distance (GSD)	1 km in spatial direction
	1 km in flight direction
Spectral channels	UV (270-320 nm)
	UVIS (320-500 nm)
	NIR (675-775 nm)
Spectral resolution	<0.5 nm
Integration time	72.3 ms (TBC)
Signal to noise ratio (SNR)	>800
Modulation transfer function (MTF)	>25% at Nyquist frequency
Power consumption	3.5 W average/7 W peak

2.2.2. Satellite concept specifications

To keep the recurring cost of the SPECTROLITE satellites low, the design of the satellite is based on standardisation and Commerical Off-The-Shelf (COTS) components. Whenever possible, flight-proven

COTS subsystems are used, or available designs are adjusted to meet the SPECTROLITE needs. The preliminary specifications of the satellite are tabulated in Table 2.2.

Table 2.2: SPECTROLITE satellite concept specifications

Specification	**Value**
Mass	<25 kg
Volume	20 x 30 x 20 cm (12 CubeSat units)
Recurring production cost	<€ 5 million
Orbital height	600 km
Orbit type	Sun-synchronous
Lifetime	3 years
Average power	20 W

3

Requirements analysis and definition

To ensure the design of the SPECTROLITE attitude determination system will fulfil the needs of the customer and the overall mission, it is of great importance to first properly identify these needs. Based on these needs, a set of formal customer requirements is defined. These customer needs and requirements would normally form the basis for a system concept selection. However, since the top-level system concept of using an image based attitude determination system has already been selected by the customer, this step is omitted. Instead, the next step is the discovery of the full set of system requirements for the attitude determination system. By analysing the various requirements a set of formal system requirements is then defined, which will form the basis of the further design process.

3.1. Customer needs and requirements

3.1.1. Needs analysis

The Attitude Determination and Control System (ADCS) on board satellites allows pointing of the satellite and on board instrument at a target area. Furthermore, it provides the stability required by orbiting optical instruments to achieve their required imaging quality. The ADCS is therefore key in achieving the required performance of an optical space instrument.

No more than two decades ago, the possibilities of incorporating an ADCS on board a small satellite were limited. Satellites in the micro satellite class (10-100 kg), generally used only passive attitude control [7]. This prevented their use in Earth observation missions. However, miniaturisation of ADCS components has greatly increased the pointing accuracy and stability of of these micro satellites. Especially the development of MEMS sensors and actuators meant a huge leap forward in ADCS performance for small satellites. Together with the miniaturisation of other satellite subsystems and payloads, this has opened the way for small Earth observation satellites with performances, previously only achievable by their larger siblings. As we have seen in the previous chapter, SPECTROLITE is a good example of such a compact high performance satellite, achieving a huge decrease in satellite size compared to its predecessors.

A typical ADCS consists of sensors to measure the attitude, actuators to alter the attitude and a controller to transform the attitude measurements into actuator commands. Current high performance micro satellites for Earth observation missions use star-trackers for fine attitude measurements and reaction wheels as attitude actuators. These star-trackers image the celestial sphere and use a star catalogue to determine where they are pointing.

Need identified
Even though miniaturisation has made high accuracy attitude control for micro satellites possible, the ADCS requirements often put significant strain on mass, volume and cost budgets of these satellites. For larger traditional satellite, the ADCS mass is around 5-10% of the satellite mass [29]. If we look at

small satellites, the mass portion of the ADCS can be much larger. For the BRITE series of satellites, a prime example of a small satellite with high pointing accuracy capabilities, the mass of the ADCS is more than 20% of the total satellite mass [7]. This relatively high ADCS mass is partially due to the high mass of the star-tracker used, compared to the total mass of the satellite. Further miniaturisation of these star trackers is becoming increasingly difficult. Since they are optical instruments, their performance is physically bound by the size of their optics [35]. Performance increases are still being achieved by improving star-tracking algorithms and improving the electronics design. However, satellite integrators are starting to look at other methods of achieving high performance attitude determination allowing further reduction in satellite mass. The need for high performance attitude sensors in smaller packages will only increase as we move from micro satellites to nano (1-10 kg) and pico (0.1-1 kg) satellites.

To determine the current state-of-the-art in commercial off-the-shelf attitude sensors for small satellites, a market analysis was performed. The results of this market analysis is presented in Table 3.1. Only sensors with performances proven in orbit are included in this table.

Opportunity identified
In 2003, the MOST micro satellite was launched. It achieved an unprecedented micro satellite ADCS performance and has one of the most accurate attitude determination systems to date [7]. The MOST satellite carries a miniaturised space telescope to measure variations in star light. It achieves the high attitude determination accuracy by subsequently using its telescope as star-tracker. This allowed the creation of a high accuracy star-tracker without any additional sensor mass [7]. The pointing accuracy it achieved is of order 3-5 arcsec, barely achievable by the most accurate micro satellite star tracker found in Table 3.1.

It might be possible to transfer this principle to Earth observation satellites. Therefore, TNO wants to investigate whether the captured imagery from the SPECTROLITE instrument can also be used to determine the satellite's attitude. This could allow using SPECTROLITE as high performance attitude sensor and achieving a reduction in satellite mass and cost compared to using a more traditional star-tracker system.

Table 3.1: Market analysis of high accuracy attitude sensors for small satellites. (ST: star tracker, SS: sun sensor)

Brand	Name	Type	Accuracy (3σ) [arcsec]	Mass [g]	Average power [mW]	Cost per unit [k€]
SpaceMicro	uStar-400M	ST	3	3300	18000	
SpaceMicro	uStar-200M	ST	3	2100	8000	
SSTL	Rigel-L	ST	9	2600	6500	430
Hyperion	ST400	ST	10	280	700	
SpaceMicro	uStar-100M	ST	15	1800	5000	
SSTL	Procyon	ST	15	1200	6500	330
Sinclair	ST-16RT2	ST	15	185	500	120
Blue Canyon Tech	NSC-1	ST	21	500	500	
SpaceMicro	MIST	ST	90	500	3000	
Maryland Aerospace	MAI-SS	ST	140	282	1500	32.5
NewSpace	Fine SS	SS	360	35	210	12
CubeSpace	CubeSense	SS	720	40	50	4.8
SolarMEMS	SSOC-A60	SS	1080	25	36	
GOMspace	NanoSense FSS	SS	1800	3	11	
SolarMEMS	nanoSSOC-A60	SS	1800	4	10	2.2
NewSpace	Cubesat SS	SS	1800	5	50	
NanoAvionics	DSS1	SS	1800	15	50	
Lens R&D	BiSon64	SS	1800	22	50	4.1
Orbital Systems	SS	SS	3600	10	22	

3.1.2. Customer requirements

The capabilities of the attitude determination system as required by TNO are tabulated in Table 3.2. The customer has provided a pointing accuracy requirement of 0.01 degree per axis (REQ-CUST-01). To assess the validity of this requirement, a pointing budget analysis is performed in Section 3.3. This analysis will also allow the definition of an attitude determination accuracy requirement for the image-based attitude determination system under design.

The customer has also defined the required Modulation Transfer Function (MTF) value at the Nyquist sampling rate (REQ-CUST-02). The MTF is an image quality parameter. Although an MTF at Nyquist of 0.25 is relatively low compared to other hyper-spectral Earth observation instruments, it was determined by TNO to suffice for performing the trace-gas measurements with the required resolution. The MTF requirement will form the basis of the pointing budget analysis in Section 3.3 in which the required pointing and pointing knowledge accuracies are determined. As we will discuss in detail in Section 6.1.2, the MTF of the instrument also effects the performance of image-based attitude determination methods. Therefore, this MTF requirement will also serve as a boundary condition for the quality of the imagery used in simulating the image-based attitude determination performance.

Attitude measurements provided by the image based attitude determination system will only be available when the instrument is capturing imagery of a landmass. At the start of an imaging run, the satellite's attitude might have drifted and instrument might not be nadir pointing. In the time needed to correct the attitude error using measurements provided by the image based attitude determination system, the captured imagery might not be as required. Image quality can be deteriorated due to higher rotational rates and consecutive image lines might not be placed next to each other. This can lead gaps in the data provided by the instrument. To limit this effect, the customer requires the ADCS to provide stable nadir pointing within 30s of reaching an observable region (REQ-CUST-03).

Apart from use in the control of the satellite's attitude, instrument pointing knowledge is also needed in the processing of instrument data on ground. To accurately compute trace-gas concentrations from the spectrometer measurements, the surface reflectance value of each pixel in an image line is needed. To accurately determine this surface reflectance value, the location of each individual pixel must be known with an accuracy of 100m (REQ-CUST-04). Note that this requirement does not necessarily have to be fulfilled by the attitude determination system on board the satellite, but can also be achieved by post-processing of the instrument data on ground. However, in Section 11.1.2 we will also investigate the image registration accuracy of the designed algorithms when used for on ground post-processing of the instrument data.

Table 3.2: Attitude determination system capabilities required by the customer.

ID	Requirement
REQ-CUST-01	The attitude determination system shall allow an absolute instrument pointing accuracy of 0.01 degree (1σ)
REQ-CUST-02	The attitude determination system shall allow an instrument MTF of 0.25 at Nyquist rate in along- and cross-track direction
REQ-CUST-03	The attitude determination system shall enable the satellite to achieve stable nadir pointing within 30s after reaching a region of interest (TBC)
REQ-CUST-04	Pixel geo-location with an accuracy of 0.1 times the GSD (3σ) shall be available for on-ground data processing

Since this project involves a feasibility study for a technology with a low Technology Readiness Level (TRL) no requirements are placed on the system characteristics like mass, volume, power and cost at this point. Instead, the properties of the image based attitude determination system will be compared to the current design of the SPECTROLITE star-tracker based attitude determination system to assess its feasibility. This comparison is performed in Section 11.3.

3.2. Requirement discovery and definition

The next step in the design process is to define a set of system requirements. As we will discuss in more detail in Section 4.1, a generic COTS ADCS will be used as a basis for further design work. This system will be augmented by an image based attitude determination system during observation mode. In the scope of this research project, only the image based attitude determination system is investigated in detail. However, to assess the feasibility of the image based attitude determination system, we want to compare the properties of this system to those of a star-tracker based system. To allow a valid comparison between the two, any additional sensors incorporated in the image based attitude determination system must be taken into account. Therefore, a full set of attitude determination system requirements is defined.

3.2.1. Pointing mode identification

The attitude determination system requirement definition starts with an identification of the satellite pointing modes. The pointing modes are identified using an analysis of the satellite's operational flow. This operational flow is visualised in Fig. 3.1. Functions of importance to the attitude determination design are highlighted.

Based on the functional analysis, the following pointing modes are identified:

- Detumbling
- Injection error correction
- Calibration
- Antenna pointing
- Fine attitude determination acquisition
- Instrument pointing
- Solar panel pointing
- Power saving
- Safe mode

While most of these pointing modes are commonly found in satellite design, one of them is not. The "Fine attitude determination acquisition" mode is specific to an image based attitude determination system. As we will discuss in more detail in Chapter 5, for an image based system to acquire an attitude fix, some initial knowledge of the satellite's attitude is needed. This initial attitude knowledge must therefore be achieved by some other means.

Also notice the propulsion system will only be used to correct orbit injection errors and for constellation phase acquisition. Given the limit lifetime of the satellite, it is assumed by the customer that no further orbit corrections will be necessary.

3.2.2. Pointing mode requirements

For each of the pointing modes identified, the attitude determination accuracy requirements are determined. These accuracy requirements per pointing mode are tabulated in Table 3.3. Note that not all pointing modes require knowledge of the absolute attitude.

Table 3.3: Attitude determination accuracy requirements per pointing mode.

Pointing mode	**Required accuracy 1(σ) [$^\circ$]**	**Based on**
Injection error correction	>1	Reference satellite
Calibration	TBD	Calibration procedure not yet fully known
Antenna pointing	1 (TBC)	GOMX-3 reference satellite
Fine AD acquisition	0.1	Determined via simulation in Section 9.3
Instrument pointing	0.01	Analysis in Section 3.3
Momentum dumping	>1	Reference satellite
Solar panel pointing	>1	SPECTROLITE power budget analysis

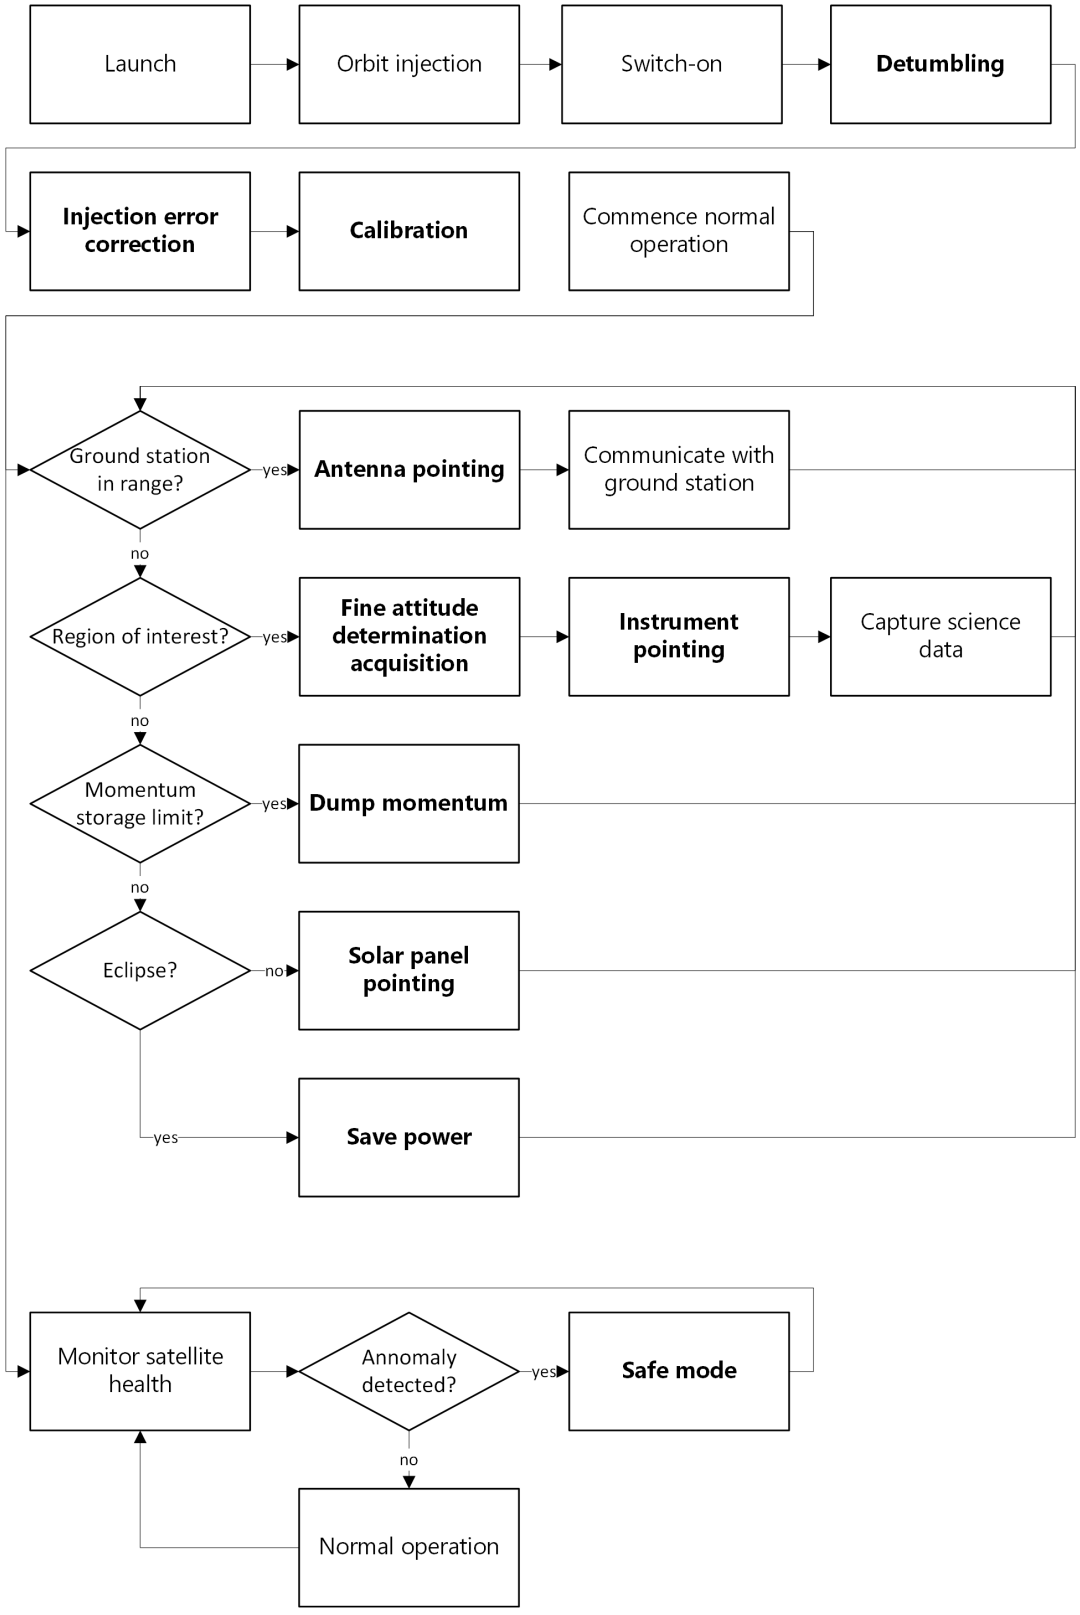

Figure 3.1: Satellite operational flow diagram.

3.3. Pointing error budget analysis

The customer has expressed a requirement on the satellite pointing accuracy of 0.01 degree (1σ) during observation mode. The pointing accuracy affects the motion of the instrument during the in-

tegration time and thereby has a large effect on the instruments MTF. A lower pointing stability will decrease the instrument's MTF and thereby limit the achievable resolution. To determine the pointing stability required, the effect of movement during the integration time of the instrument on the MTF is investigated. Based on the allowed movement to meet the MTF requirement, a pointing error budget is made.

3.3.1. Instrument MTF analysis

The MTF of an optical instrument is a convolution of its individual contributions. In this case we distinguish the contributions due to the instrument's optics, detector and motion. The MTF is a function of the spatial frequency u. The spatial frequencies measured are in this case limited by the detector pixel size. The Nyquist frequency is defined, according to the Nyquist criterion, as twice the limiting frequency.

To determine the MTF contribution of the instrument's optics, we make use of verification measurements determining the instrument's spot size. If we assume a Gaussian spot distribution, the optics MTF contribution can be computed using [12]

$$MTF_{optics} = exp(-3.56u^2\alpha_{spot}^2) \tag{3.1}$$

Where α_{spot} is the Full Width at Half Maximum (FWHM) angular spot size in degrees. This RMS FWHM spot size was measured to be $6\mu m$. We assume this spot size is circular symmetric. In reality, the SPECTROLITE instrument uses a diffraction grating to decompose incoming rays based on wavelength along one dimension. This can cause a non circular symmetric spot size.

Next, we determine the MTF contribution of the detector. The detector MTF is itself a combination of the detector footprint MTF and the sampling MTF. If the sampling distance used is equal to the detector footprint, the MTF contribution of the detector can be computed using [6]

$$MTF_{detector} = sinc(u\alpha_{sampling})^2 \tag{3.2}$$

Where $\alpha_{sampling}$ is the detector angular sampling distance.

Lastly, the MTF contribution of the instruments motion during the integration time is computed using [6]

$$MTF_{motion} = sinc(u\alpha_{motion}) \tag{3.3}$$

Since the parameters like spot size, sampling distance and motion are not necessarily equal in all directions, we make a distinction between along-track and cross-track MTF. The along-track MTF will determine the pointing accuracy required for the pitch and yaw angles, while the cross-track MTF will determine the pointing accuracy for the roll angle. Meeting the along-track MTF requirement will be most challenging, since the motion in along-track directions has contributions from the roll and yaw angles, as well as from the motion of the satellite in orbit.

Table 3.4: Along-track MTF contributions and total MTF at Nyquist.

Contribution	**Angular value**	**MTF at Nyquist**
Spot size	0.0286°	0.923
Detector footprint	0.0955°	0.405
Motion	0.0908°	0.668
Total		**0.250**

To meet the along-track MTF requirements, the total angular motion during the integration time of 0.0723 seconds cannot exceed 0.0908°. The angular motion due to the satellite's velocity during the integration time is equal to 0.0478°, thus leaving room for 0.0430° of additional motion. Since the image-based attitude determination method will provide absolute attitude measurements, we turn this maximum allowable motion during the integration time into an absolute attitude measurement accuracy, instead of setting a requirement for the allowable rotational rates of the satellite.

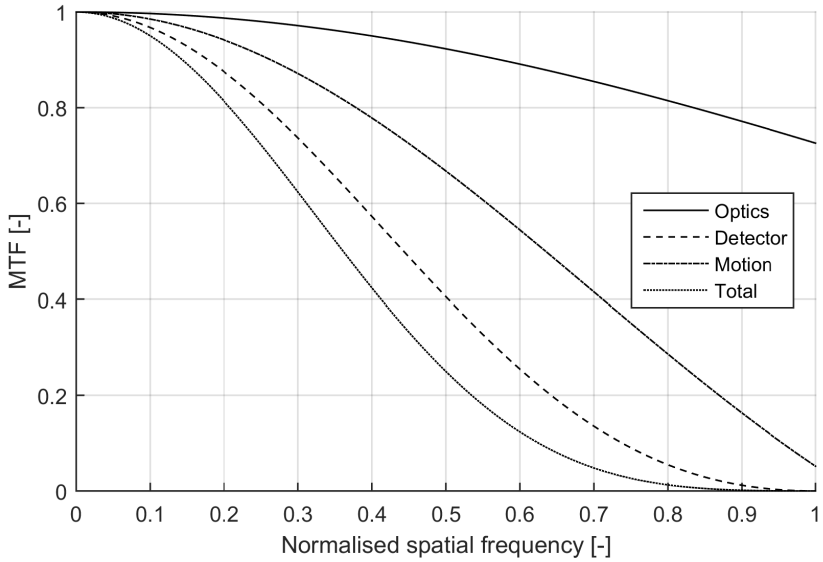

Figure 3.2: Along-track MTF contributions and total MTF.

We model the attitude controller as a Proportional-Differential (PD) controller. We assume the controller tries to move the entire distance of the accumulated pointing error in a single control cycle [23]. For this analysis we also assume the control cycle frequency is equal to the imaging frequency. Currently, the integration time is estimated to be equal to half of the imaging period. This means the imaging period and thus the control cycle period take twice as long as the integration time. If we assume the rotational rates are constant over a control cycle, the total allowable motion of the satellite in along-track direction during a control cycle is twice the allowable motion during the integration time of 0.43° and is therefore equal to 0.086°.

This maximum allowable pointing error value of 0.086° represents the maximum pointing error at any location on the image line. The worst case pointing error occurs at the outer edges of the image line, as this is the location where the arm of the yaw angle causes the largest error contribution. This maximum allowable pointing error will form the basis of our pointing error budget as discussed in the next section.

3.3.2. Pointing error contributors

The pointing error budget analysis is based on the method used in [23]. The reference analysis forms the basis of the pointing budget for the BRITE satellites, a multi-university CubeSat project. These satellites use star-trackers to achieve arc-minute pointing accuracies. Recently, the pointing accuracy of one of the BRITE satellites has been validated by analysis of the captured imagery and was shown to be within the pointing budget [36].

We start by identifying all contributors to the pointing error budget [29] [23]:

- Attitude determination error
- Reaction wheel imbalance jitter
- Reaction wheel control jitter
- Reaction wheel cross-coupling
- Environmental disturbance torques
- Damping error

Since the instrument itself is also used as the attitude sensor, there are no sensor to instrument misalignment error. The propulsion system will only used for orbit insertion error correction and to get the satellite in its correct position within the constellation. Due to the limit lifetime of the satellite, the current design assumes no orbit corrections are needed after the satellite becomes fully operational.

Therefore propellant slushing during the operational lifetime of the satellite is neglected.

Reaction wheel imbalance jitter
Mass imbalance causes an asymmetric momentum distribution around the rotational axis of the reaction wheels. As the reaction wheel is spinning, the mass imbalance induces jitter on the satellite body. The magnitude of this jitter can be computed, starting with the Fourier transform of the equation of motion for a single axis reaction wheel[23]

$$\mathcal{F}\left\{I\ddot{\theta}(t)\right\} = \mathcal{F}\left\{g(t)\right\} \tag{3.4}$$

This equation represents the angular acceleration $\ddot{\theta}$ caused by an applied torque g. Written in the frequency domain, this equation becomes

$$-I\omega_f\theta(\omega_f) = g(\omega_f) \tag{3.5}$$

This can be rewritten as

$$\theta(\omega_f) = \frac{g(\omega)}{-I\omega_f^2} \tag{3.6}$$

θ represents the angular wander due to the disturbing torque g as a function of ω. Note that ω_f is equal to $2\pi f$ and represents the frequency of the jitter, not a rotational rate. The mass imbalance in the reaction wheel will cause a centrifugal force when the reaction wheel is spinning. The magnitude of this centrifugal force exerted on the satellite body is given by [30]

$$\overline{F}_{Wb} = \omega_W^2 \cdot \overline{d}_W \tag{3.7}$$

Where $\overline{d}_W$ describes the direction and magnitude of the reaction wheel imbalance

$$\overline{d}_W = d_W \begin{bmatrix} 0 \\ sin(\omega_W t + \theta_d) \\ cos(\omega_W t + \theta_d) \end{bmatrix} \tag{3.8}$$

If the reaction wheel centre of rotation has an offset from the centre of mass of the satellite, the centrifugal force will induce a torque on the satellite body. The average torque induced by the centrifugal force is equal to

$$\overline{\tau}_{Wb} = \sum_w \overline{r}_W \overline{C}_{bW} \overline{F}_{Wb} \tag{3.9}$$

Eq. (3.8) show the frequency of the induced jitter occilates with the rotational rate of the reaction wheel. We can therefore combine Eq. (3.6) to Eq. (3.9) to estimate the pointing error caused by the reaction wheel imbalance using [23]

$$\theta_{Wb} \approx \frac{r_W \cdot d_W}{I} \tag{3.10}$$

As a reference for the SPECTROLITE reaction wheels, we use the Sinclair Nanosatellite Reaction Wheels (RW-0.03) designed for 6-12 unit CubeSats. Testing of a number of reaction wheels shows they suffer from a maximum imbalance of $0.51 \times 10^{-6} kg \cdot m$ [23].

Reaction wheel control jitter
To accurately control the attitude of the satellite, we want to be able to control the torque applied by the reaction wheels. Therefore the reaction wheel assembly is equipped with a torque controller. Because of imperfections in the controller, the torque applied is not perfectly constant. Noise in the torque levels causes additional jitter on the satellite.

Measurements on the reference Sinclair reaction wheels show RMS value of the torque error is under $0.5\mu Nm$ [23]. The angular distance travelled due to a torque in general is

$$\theta_\tau = \iint_{t_O} \frac{\Delta\tau}{I} d^2t \tag{3.11}$$

When assuming constant torque over the control time-step this reduces to

$$\theta_\tau = \frac{\Delta\tau}{2I} t_c^2 \tag{3.12}$$

where t_c is the control period. As a precaution we use triple the measured torque RMS error to compute the estimated pointing error due to reaction wheel control jitter.

Reaction wheel cross-coupling
The moments of inertia of the satellite will not be perfectly symmetric. Additionally the reaction wheels will never be perfectly aligned with the principle axes of the satellite. These effects cause cross-coupling of the reaction wheel momentum, leading to an internal disturbance torque. The exact value of this torque is dependent on the rotational rates of the satellite, but also on the exact satellite inertia tensor. For the BRITE satellite the worst case disturbance torque was found to be 1.75 μNm. Since the mass of the BRITE satellite is around a factor four lower, we expect an increase of this disturbance torques for the SPECTROLITE satellite by a factor of four as well. With some additional margin we estimate this disturbance torque by a value of 10 μNm.

Environmental disturbance torques
A number of external disturbance can cause a torque on the satellite body. These environmental disturbance torque sources are:

- Aerodynamic drag
- Solar pressure
- Residual magnetic dipole
- Gravity gradient

The worst case aerodynamic torque can be computed using [29].

$$T_a = \frac{1}{2}\rho V^2 C_d A_a (C_{pa} - cg) \tag{3.13}$$

Since the satellite is nadir pointing, the aerodynamic torque will only affect the pitch angle. For the pointing error budget a typical C_d of 2.2 is used and the density is estimated to be around $1.02 \times 10^{-11} kg/m$ [45]. The centre of pressure is assumed to be 0.1 m off from the centre of mass.

Solar pressure torque can be estimated using the equation [29]

$$T_s = \frac{F_s}{c} A_s (1+q) cos(i_s)(C_{ps} - cg) \tag{3.14}$$

Where F_s is the solar intensity and c the speed of light. The q value represents the reflectivity of the satellite estimate at 0.6 [29]. We use a worst case situation where the sun directly hits the front side of the satellite at an angle of incidence i_s of zero degrees. The centre of pressure is assumed to be 0.1 m off from the centre of mass for the pitch axis and 0.05 m off for the yaw axis.

The worst case residual magnetic dipole torque for a satellite in a circular orbit can be computed using

$$T_m = \frac{2MD}{R^3} \tag{3.15}$$

For an uncompensated CubeSat, the residual dipole D can vary greatly. For a number of investigated CubeSats, the maximum residual dipole was found to be around 0.01 Am^2 [40]. The magnetic moment of the Earth M is equal to $7.96 \times 10^{15} Tm^3$. R is the radius of the satellite's orbit.

The worst case gravity gradient torque experienced by the satellite can be computed using [29]

$$T_r = \frac{3\mu}{2R^3}|I_z - I_{x/y}|sin(2\theta) \tag{3.16}$$

μ and R are the Earth's gravitational parameter and the radius of the satellite's orbit respectively. I_a is the moment of inertia around axis a. θ is the attitude of the satellite relative to nadir. We assume a maximum deviation from nadir of 0.1 degrees.

Damping error
The PD controller determines the applied reaction wheel torque based on the measured pointing error and rotational rates. The dependence on rotational rate is used to damp out oscillations around the satellite's steady state. Although it improves the stability of the satellites, error in the rotational rate knowledge will cause an additional pointing error source. The erroneous torque applied due to an error in rate knowledge is equal to [23].

$$\Delta\tau_{control} = K_d \Delta\omega \tag{3.17}$$

where K_d is the differential controller coefficient. This coefficient is determined in Section 8.1.2.

3.3.3. Pointing error budget

The values of the pointing error contributions found using the described estimation equations are tabulated in Table 3.5. The total error per axis for the pitch and yaw angle is selected such that the along-track movement of the image line outer pixels is equal to the maximum allowable value of 0.086°.

Table 3.5: Pitch and yaw axis pointing budgets.

Source	Pitch error [°]	Yaw error [°]
Determination noise	0.0474	0.0474
Reaction wheel controller jitter	3.24×10^{-6}	5.26×10^{-6}
Reaction wheel imbalance jitter	1.62×10^{-5}	1.75×10^{-5}
Reaction wheel cross-coupling	2.16×10^{-5}	3.51×10^{-5}
Environmental disturbance torque	1.62×10^{-5}	1.75×10^{-5}
Damping error	0.00150	0.00182
Total error per axis	**0.0489**	**0.0493**

We see the pointing error is dominated by the error in attitude determination. The determination noise values for pitch and yaw angle are currently set to an equal value. However, since both values contribute to the maximum along-track motion at the outer image line pixel location, we can trade measurement accuracy values between the two angles.

The allowed attitude determination noise per axis is found to be 0.0474°. We interpret this noise level as the 3σ value to ensure the MTF requirement is met with a high level of certainty. This gives an allowed per axis attitude determination noise of 0.0156° 1σ, a value close to the 0.01 ° 1σ value prescribed by the customer. Given the current maturity of the SPECTROLITE design, incorporation of some additional margin in the attitude determination requirement might be needed. We therefore stick to the customer requirement stating an attitude determination accuracy of 0.01 ° 1σ per axis.

This requirement will cover both absolute pointing, as well as pointing stability. In Section 11.2, we will analyse whether use of an image-based attitude determination method with the required attitude determination accuracy, indeed leads to sufficient pointing stability and sufficiently low rotational rates.

3.4. List of system requirements

Based on the analyses described in this chapter, a formal set of attitude determination system requirements is formulated. These requirements are tabulated in Table 3.6.

Table 3.6: Attitude determination system requirements.

ID	Requirement
ADS-REQ-1	The attitude determination system shall provide 0.01 degree (1σ) per axis absolute pointing knowledge during observation mode
ADS-REQ-2	The attitude determination system shall allow stable nadir pointing within 30 seconds after reaching an observable zone (TBC)
ADS-REQ-3	The attitude determination system shall provide 0.1-1 degree of absolute pointing knowledge outside of eclipse
ADS-REQ-4	The attitude determination system shall provide attitude measurement with a frequency equal to imaging frequency (TBC)

The current set of attitude determination system requirements is small. This is partly due to the immaturity of the SPECTROLITE satellite design. As the design matures, the list of attitude determination system requirements will grow.

4

Attitude determination system conceptual design

Now the requirements for the attitude determination system are set, we can move on to the design of a system concept. As we have seen in the previous chapter, multiple pointing modes exist for the SPECTROLITE satellite. The image-based attitude determination method we will design might not suffice for all of these pointing modes. We will therefore first select an overall attitude determination concept and select which method of attitude determination is used during which pointing mode. We will then move on to selecting the image-based attitude determination concept, after a short summary of the current state-of-the-art in this field.

4.1. Attitude determination system concept

The conceptual design process starts with the selection of an overall attitude determination system concept. As determined in the previous section, it is the observation mode which induces the most stringent requirements on the attitude determination accuracy. The purpose of this research project is to investigate the possibility of using an image based attitude determination method to meet these requirements. However, in the design of the overall attitude determination concept, it is important to also look at the other pointing modes identified. Although the accuracy requirements during these modes might be looser, also using the image based attitude determination method during these secondary points modes might not lead to the best overall system design.

Limitations on image-based attitude determination
Existing high accuracy image based attitude determination concepts, as we will later discover, have significant limitations on their usability. They require an unobstructed view of land or coastal areas for captured imagery to have the level of detail required to determine attitude with sub-pixel accuracy. This limits their functionality in case of ocean crossing or when a significant portion of the instruments field of view is obstructed by clouds. This is especially true for satellites in low Earth orbit with a limited swath width, an important reason why image based attitude determination has mainly been used on-board geostationary weather satellites.

Furthermore, if the on-board science instrument is used to provide the required images, attitude determination accuracy is bound by this instrument's performance. For the SPECTROLITE mission, this is an instrument designed for capturing reflected sunlight, meaning image based attitude determination is unavailable during eclipse.

Alternative methods of instrument based attitude determination
Still, it might be possible to devise a new lower accuracy attitude determination method employing the

on-board instrument as a sensor, for use during other pointing modes. Such methods could include using the instrument as a horizon or sun sensor, or using imagery of captured thermal radiation emitted by the Earth's surface during eclipse. However, such methods would induce additional requirements on the instrument design in terms of detectable wavelength, radiometric resolution, etc. The push-broom design of the instrument also creates additional complexity in the design of a horizon or sun-sensor-like determination method, given the one-dimensionality of its sensor array.

Incorporating secondary attitude sensors
Based on the market analysis provided in the previous chapter, it can be concluded that there are many COTS attitude determination sensors available meeting the accuracy requirements outside of observation mode. The effect on the attitude determination system budgets of incorporating these additional sensors is limited. The additional mass of a sensors bench meeting the attitude determination requirement of 1° outside of observation mode is estimate at around 100 gram. The power consumption of these sensors is significantly lower compared to the average power used by the SPECTROLITE instrument of 3.5 W.

Final concept
Given the challenges of using the instrument as attitude sensor outside of observation mode and lack of a clear advantage in doing so, secondary attitude sensors will be incorporated in the full attitude determination concept. The final concept is build around an existing CubeSat Commercial off-the-shelf (COTS) sun-sensor-based attitude determination system. The sun-sensors will provide attitude determination capability during all pointing modes except for observation mode. The COTS system will be augmented by image based attitude determination capabilities for use during observation mode. We introduce the term "image navigation" to indicate the image based attitude determination.

The use of sun-sensors to provide attitude measurements is limited to periods when the satellite is outside of eclipse. We therefore assume the various pointing modes identified can all be performed when the sun is in view of the satellite. However, during eclipse we must ensure the satellite's rotational rates do not become excessively large. Therefore, a gyroscope is added to allow monitoring and adjustment of these rotational rates during the eclipse phase. Additionally, the gyroscope will be used to improve the image-based attitude determination performance, by combining absolute attitude measurements with rotational rate measurements in a filter.

The assumption that there is no pointing requirement during eclipse must be validated using a detailed analysis. If pointing is indeed needed during eclipse, a change of COTS attitude sensor might be needed. However, this design change will likely have a negligible effect on the design of the image navigation system. Therefore, this detailed analysis is not yet performed.

Sizing of the COTS attitude determination system will be performed after the design of the image navigation method, in Chapter 10.

4.2. Existing image based attitude determination concepts

A number of different methods for image-based attitude determination have already been conceived. Two distinctively different methods can be distinguished, base on either photogrammetry or direct image co-registration.

4.2.1. Photogrammetry based

The principles of photogrammetry allow the determination of the satellite's attitude by tracking the motion of subsequent images taken by an on-board camera. As visualised in Fig. 4.1, the movement of the satellite in orbit allows an on-board camera to capture a series of partially overlapping images. Image processing algorithms can then be used to determine the motion of the Earth's surface with respect to the camera's imaging plane from these images. This motion is dependent on the Earth geoid, and the satellite's location and attitude. The use of an Earth geoid model and accurate knowledge of the

satellite's location acquired using for example a Global Navigation Satellite System (GNSS) receiver, therefore allows the calculation of the satellite's attitude [21].

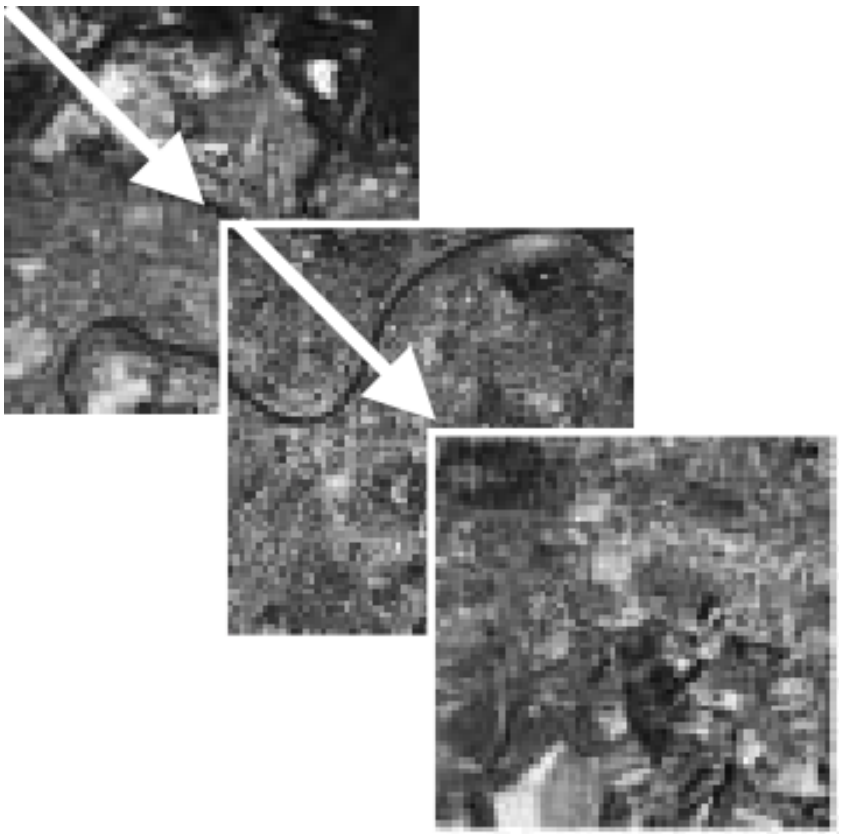

Figure 4.1: Overlapping subsequent images [21].

Unfortunately, the need of capturing subsequent overlapping images requires the satellite to carry a two-dimensional imager. SPECTROLITE on the other hand is a one-dimensional push broom imager. This means it captures only a single image line at a time. It is therefore not possible to use SPECTROLITE to capture partially overlapping images without actively controlling the satellite's attitude for a longer period of time. Doing so will require a constant knowledge of the satellite's attitude, which is not available prior to the processing of the captured images. This inhibits the use of photogrammetry as an image navigation method for the SPECTROLITE instrument.

There is, however, a notable exception to this rule. Recently, data from a push-broom imager on board a Disaster Monitoring Constellation (DMC) satellite was used to perform image based attitude determination on-ground. The imager has a dual bank configuration, allowing an increase in swath width. This essentially means two imagers are combined with their respective swaths placed next to each other to form a single straight imaging line. For this specific instrument, the two arrays have a partial overlap in cross-track direction to ensure continuity. Additionally, the two arrays have a separation in along-track direction. This causes a time delay for features at the fore array to be captured by the aft array. The region of overlap between the two arrays can therefore be used to compute image motion and, by extension, the satellite's attitude [3]. The SPECTROLITE satellite will house between one and three separate channels with individual optics and telescopes. For a satellite with two or more channels it is possible to employ the photogrammetry method by allowing an along-track separation between the swaths of at least two channels. However, even if multiple channels are used, separation between the swaths might not be acceptable as it can cause increased difficulty in co-registration of the image data.

4.2.2. Direct co-registration based

Co-registration is the process of determining the coordinates of an image by aligning it to another image with known coordinates. In this image registration process, the image captured by the on-board instrument is known as the sensed image, while the image with known coordinates is referred to as the reference image [16]. The reference image can be created from previously co-registered satellite imagery stored on-board the satellite and represents the region where the sensed image could be located. Direction co-registration methods can be split into feature based and intensity based methods.

Feature based
Feature based image registration is a two-step process. The first step is the feature detection. In this step, features like intensity maxima, edges or corners of homogeneous regions, or other distinctive features are detected in both images. In the second step, feature matching, the detected features in the sensed image are matched to features in the reference image. Matching at least three features between both images allows the computation of the satellite's position and attitude [16].

Similarly to the photogrammetry method, the feature matching process requires the sensed image to be a two-dimensional image. However, in this case the two-dimensional image can be created using a push broom instrument by addition of multiple subsequent one-dimensional image lines. It does mean it is impossible to determine the location of each image line separately. The satellite's attitude can only be computed every time a new group of subsequent image lines is available and a propagation technique will be required to fill in the gaps in attitude knowledge. This will lower the real the accuracy of the real-time attitude knowledge available to the attitude controller.

Intensity based
A third way to perform co-registration is to make use of an intensity matching method. This method registers the sensed image to a reference image, similarly to the feature detection method. However, the method of registration differs. Intensity based methods use algorithms to compute the similarity or dissimilarity between the sensed and reference image at various locations. The location with highest correlation between the two images should correspond to the actual location of the sensed image within the reference image [16].

Once again an on-board catalogue of co-registered satellite imagery is needed for constant attitude determination. Furthermore, the need to move the sensed image around within the reference image and perform the correlation at each point means the amount of correlations needed grows rapidly when the size of the reference image is large. The amount of processing power required to locate the sensed image can therefore be highly dependent on the approximate initial knowledge of the sensed image's location [52]. An advantage of this image correlation method over the previously described methods is that for many of the correlation algorithms, the shape of the image does not influence the performance of the correlation. This means that it is possible to register an image with only few or even a single pixel line, allowing for more frequent updates of the satellites attitude.

4.2.3. Image based attitude determination in practice

In the literature study conducted as part of this research project, no satellite ever launched was found, using the processing of Earth imagery as its primary method of attitude determination. Instead, practical implementations of image based attitude determination have only been used to improve the performance of the main attitude determination system, mainly by removing bias or by providing improved knowledge of orbital parameters.

The highest theoretical and practical accuracies found in literature are tabulated in Table 4.1. The highest practical accuracies are currently achieved by China's National Satellite Meteorological Center (NSMC) and the US National Oceanic and Atmospheric Administration (NOAA) in collaboration with the National Aeronautics and Space Administration (NASA). Both of their satellites, FY2 and GOES respectively, are series of geostationary weather satellites.

Table 4.1: Image based attitude determination accuracies in literature

Institution	Satellite	Method	Accuracy [km]/[pixels]	Note
TU Dresden	-	Photogrammetry	5/19.1	Theoretical accuracy [21]
TU Dresden	-	Correlation	0.5/50	Theoretical accuracy [22]
Surray Satellite Centre	DMC	Photogrammetry	0.4/18.1	Theoretical, but achieved with unprocessed instrument data [3]
NSMC	FY2	Feature detection	5/4	Practical accuracy with on-ground processing [54]
NOAA/NASA	GOES	Feature detection	1.25/1.25	Practical accuracy with on-ground processing [14]

4.3. Image navigation mode concept

As part of this research project, a literature study was performed into the current state of image navigation research, as well as into the state-of-the-art in image registration in general. An important conclusion of this literature study is the limited number of practical implementations of an image navigation system found. Furthermore, these practical implementations are limited to geostationary weather satellites. For the purpose of image navigation, the imagery supplied by these satellites is vastly different from the imagery to be captured by SPECTROLITE. The large stationary two-dimensional swath of these weather satellites allows an image navigation system to use surface characteristics like shorelines and river patterns. For the much smaller one-dimensional swath of the SPECTROLITE instrument, the surface characteristics imaged can vary significantly from time to time and might not always be so clear.

Another important conclusion is that no image navigation design was found, practical nor theoretical, that achieves a sub-pixel attitude determination accuracy. The SPECTROLITE instrument, however, requires attitude determination with an accuracy of a tenth of its pixel size. This requirement is not just limited to the SPECTROLITE satellite. In general, the attitude determination requirement for an Earth observation satellite scales with its Ground Sampling Distance (GSD) and sub-pixel accuracy is needed to achieve the stability required for an optical instrument. This means that none of the existing image navigation concept are capable of functioning as primary attitude determination method during observation mode.

4.3.1. Concept generation

The image navigation concepts considered will not be limited to concepts already in existence, since none of these concepts already meet the requirements for SPECTROLITE. Image registration, which will form the basis of our image navigation concept, is a widely researched topic and its applications can be found in a large number of scientific fields. This includes the processing of satellite imagery, but registration is also used in the analysis of medical images and to aid autonomous navigation of various vehicle types like cars or drones. To fully benefit from the available research, new image navigation concepts will be generated based on image registration methods and implementations found in all relevant research fields.

Within these relevant research fields, a large number of image registration algorithms exist, each tailored to a specific design problem. It is impossible to look at all these implementations in detail to determine which is best suited for this specific design case. Instead, all of these algorithms are grouped into a number of higher level concepts based on key aspects of their method of operation. To aid in this process, a design options tree is created, shown in Fig. 4.2. This tree shows image navigation concepts on a number of different levels. Some branches expand further than others, as they represent a larger group of concepts.

A number of options within the design option tree correspond to methods described in the previous chapter. However, in the selection of an image navigation concept, we considered all possible image

registration implementations within their category, not simply the ones previously used for image navigation. Therefore, if a design option is selected corresponding to one of the categories in the previous section, the specific implementation used might be new, leading to a completely new image navigation system design. We will now shortly discuss the various design options in their broadest sense.

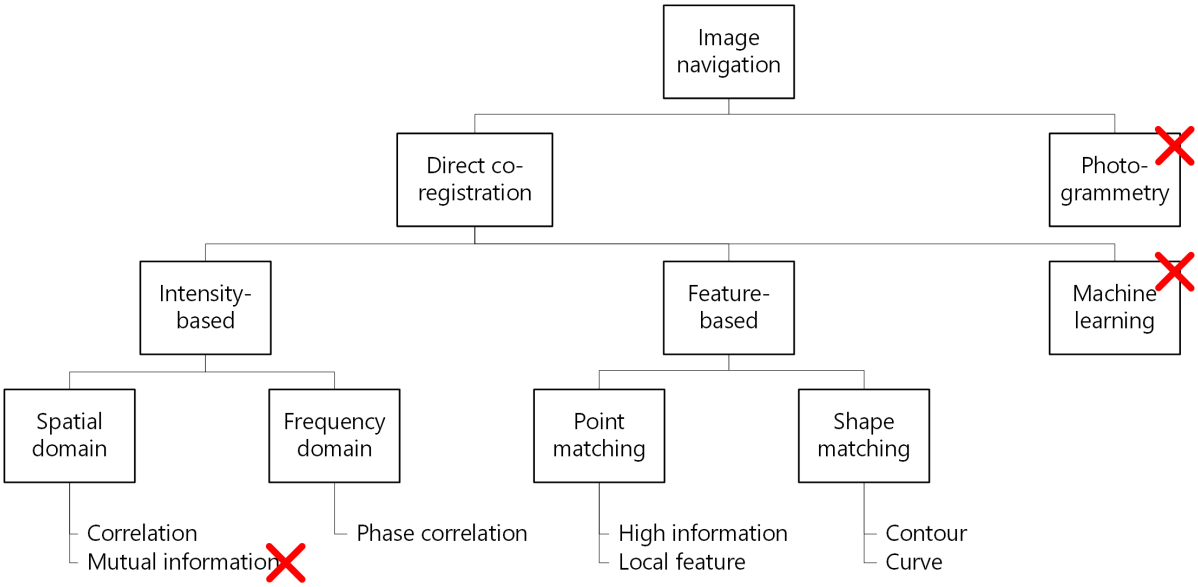

Figure 4.2: Design options tree.

Feature-based

Feature-based registration techniques match the sensed and reference images by individually matching detected features like points, lines or contours. By matching multiple points between the two images, the location of the sensed image within the reference image can be determined. Within the group of feature-based registration techniques we make a distinction between shape matching and point matching. Shape matching techniques match features such as lines or contours, not simply based on the location of the feature, but also by comparing a description of the shapes. An example of a shape matching technique commonly used for satellite imagery is the matching of coastal line patterns. Shape matching methods can achieve high registration accuracies, but require large enough two-dimensional images to accurately identify a shape.

Feature-based registration can also be performed by point matching. This registration technique can be split into high information points and local feature points. Local feature points can be found by detecting edges, local intensity maxima or any other type of local feature. Matching of points between the sensed and reference image is performed based on a description of the point or its surrounding, or simply by matching point patterns in both images. High information points, also known as ground control points, are easily distinguishable points on the Earth surface for which the location has been determined with high accuracy. This method generally uses less points, but the points used are very clearly defined. High information point matching methods allow a set of previously identified points to be used as reference imagery, without the need of analysing the reference imagery on board the satellite.

Intensity-based

Intensity-based methods exist in both the spatial and frequency domain. In the spatial domain we distinguish intensity correlation and mutual information methods. Intensity correlation treats the images as a number sequence. Many general correlation methods have found their way into image registration and can be used to determine the similarity between a sensed and reference image. By correlating the sensed image at multiple locations within the reference image, registration can be performed. The mutual information method is a statistical method. It provides a more general measure of the

association between two images. This allows it to have a high resistance against geometric changes within the sensed and reference images. In the frequency domain we can perform phase correlation. This method uses the Fourier transform of the sensed and reference image to determine the correlation between the two. Its response spikes at the location of highest correlation within the images. It has the advantage that a single computation of the phase correlation provides the correlation values at all locations of the sensed image within the reference image, limiting the need of performing individual correlations for each possible location as required for the intensity-correlation and mutual information methods.

Photogrammetry
The photogrammetry method was already discussed in the previous section.

Machine-learning
Machine-learning and the use of neural networks is a relatively new and promising image registration method, never before used for the purpose of image navigation. Instead of designing the exact registration algorithm, machine-learning systems create their own functionality by training. This training is performed by presenting the system with a large set of example registrations. However, properly training a neural network is a trivial task. These types of registration systems can easily start to show unexpected behaviour and validating their performance is challenging.

4.3.2. Concept elimination

As a next step in the concept selection process, clearly unfeasible concepts and concepts that cannot be developed yet are eliminated.

Photogrammetry
The achievable measurement accuracy of the photogrammetry method is dependent on the accuracy with which the movement of one image with respect to the next can be measured. This also requires image registration. It has been shown by [3] that when a registration accuracy of 0.1 pixel is achieved, attitude determination accuracy does not exceed 1.5 degrees for the roll axis, 2 degrees for the pitch axis and 1 degree for the yaw axis. Since achieving a registration accuracy of 0.1 pixel is already a very challenging task, it is believed to be impossible to meet the required attitude determination accuracy using this method.

Deep-learning
Deep-learning methods and neural networks is a relatively new and promising field within image registration. However, implementing such a method is highly complex. It is believed this method cannot be developed yet for the purpose of measuring attitude on board the SPECTROLITE satellite.

Mutual information
The mutual information method is incapable of accurately registering images with significant intensity variations [16]. Since the lighting conditions at which the SPECTROLITE imagery is captured can vary significantly, this is an important requirement on the registration method. Therefore, the mutual information method is eliminated.

4.3.3. Concept trade-off

To perform the image navigation concept trade-off, four trade-off criteria are defined. These trade-off criteria are based on what are considered the two main challenges of designing and implementing an image navigation system on board the SPECTROLITE satellite: achieving the required attitude determination accuracy and performing the image processing using on board hardware.

For both of these challenges, two driving concept parameters were selected as trade-off criteria. For the achievable attitude determination accuracy, the selected criteria are the individual registration accuracy of the image registration method and the number of image lines needed to achieve a successful registration. These two parameters drive the real-time attitude determination accuracy achievable using the image registration method. For the challenge of performing the image processing on board the satellite, the criteria are the computational complexity of the registration method and the amount of reference data it requires. These two properties drive the computing power and data storage capability to be carried on board the satellite. We will now discuss each trade-off criterion in more detail and provide a rationale for the scoring of each concept.

Registration accuracy
The accuracy of the image registrations performed is one of the drivers of the overall attitude determination accuracy. Specific implementations of the various techniques considered can have a large influence on the accuracy achieved. However, some general conclusions about the accuracy of the concepts considered can be drawn.

Both phase correlation and intensity correlation provide roughly the same levels of accuracy [16]. Both methods are not geometric invariant, meaning that small changes in image geometry can lower the registration accuracy. This is especially important for satellite imagery, as it is prone to geometric variations due to image motion and variations in surface elevation. Feature matching methods have a better inherent resistance against these geometry changes as they register individual image points. This gives them a higher inherent accuracy for satellite image registration purposes [31].

Even better accuracy is achieved by curve matching methods. Using the gradients of curves and edges allows registration with very high resistance against noise and smoothing of the image. These methods achieve the highest registration accuracy for satellite imagery [31]. Therefore, shore-lines and rivers are commonly used to register satellite imagery. However, these kinds of features might not always available.

Image lines needed for successful registration
In the selection of an image navigation concept it is important to consider that we are designing for a certain real-time attitude knowledge, not for a certain registration accuracy. The achievable real-time attitude determination error available to an attitude controller is not only dependent on the accuracy of individual registration. Significant time might be required to perform the image processing needed to register an image. Furthermore, due to the one-dimensional nature of a push-broom imager, multiple image lines need to be combined to create a two-dimensional image. The larger the image size needed for a registration method to achieve accurate registration, the longer it takes to acquire this image.

The time needed to acquire the imagery lowers the frequency with which attitude measurements are performed. A low attitude measurement frequency causes gaps in attitude knowledge, which have to be filled by propagation of a previous attitude state. The propagation causes a decrease in the real-time available attitude knowledge. To see how the measurement frequency affects the real-time attitude knowledge we model the propagation of attitude measurement using basic Kalman filter. An explanation about this filter type can be found in Section 5.4. We model a single attitude angle of a satellite body representative of the SPECTROLITE satellite, oscillating around nadir with an amplitude equal to the allow 3 sigma pointing error of 0.03 degrees. We combine the attitude measurements with simulate rotational rate measurements provided by a gyroscope, representative of the reference gyroscope described in Section 8.2.1. Per measurement period value, we determine the attitude measurement accuracy required to achieve a real-time attitude determination accuracy fulfilling the SPECTROLITE requirement. This relationship between attitude measurement accuracy and measurement period is shown in Fig. 4.3.

The relationship in Fig. 4.3 shows an exponential increase of the required registration accuracy with the measurement period. For high measurement frequency, the allow measurement error standard deviation can be significantly larger than the attitude determination requirement value of 0.01°. This shows that the number of image lines needed to

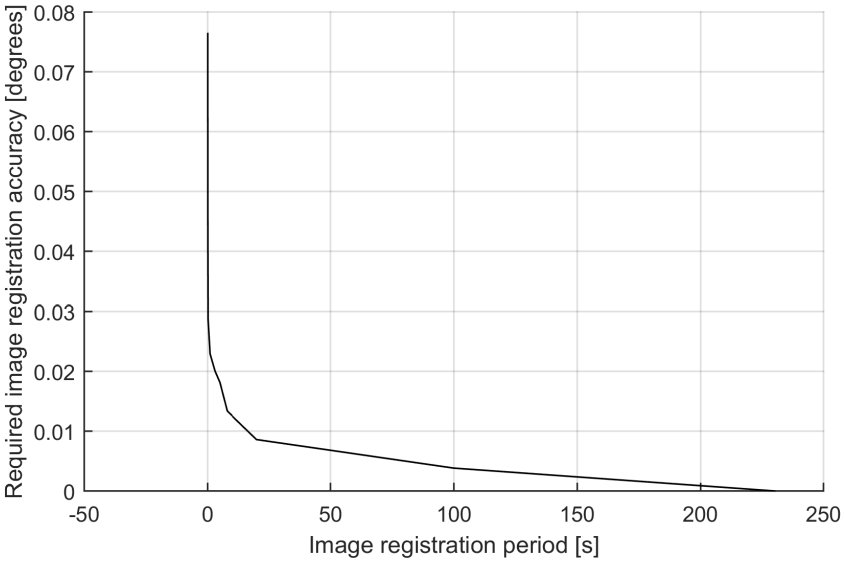

Figure 4.3: Required measurement standard deviation versus measurement update period.

For now we have neglected the processing time required for registration of the imagery. However, as processing time required to register an image scales with image size, image registration methods requiring large images will need a longer processing time per attitude measurement. This adds to the delay caused by the time needed to acquiring sufficient image lines and will strengthen the exponential relationship seen in Fig. 4.3.

Computational complexity

The computational complexity is always a key aspect of image registration techniques. However, given the limited processing power that can be carried on board a satellite, the computational complexity of the image navigation concept is of even greater importance.

Phase correlation methods are generally considered to most computationally efficient registration methods. This is especially true when large images are registered. However, when the width of the sensed image has fewer pixels than the number of correlations needed to perform the registration, zero padding of the sensed image is needed to perform the phase correlations. This makes the intensity correlation method more computational efficient when a sensed image with a limited number of image lines is registered. Feature matching methods generally have a higher computational complexity compared to correlation methods. [16].

Reference database size

The size of the reference database can grow significantly if imagery of the entire globe is needed with sufficient resolution. To achieve sub-pixel registration, correlation methods often use reference imagery with smaller resolution and resampling of the sensed imagery. As resampling of reference imagery on board will require additional processing power, it makes sense to assume reference imagery with a resolution below that of the sensed imagery is stored on board.

For feature detection methods, sub-pixel imagery is not needed to achieve sub-pixel registration. This allows a smaller reference imagery database size. Further reduction of the reference imagery database is possible when using the high information point method. This method uses a description of specific landmarks with associated locations, decreasing database size with respect to a full global set of reference imagery.

Scoring
Based on the rationale provided above, scoring of the individual concepts is performed. Quantifying the performance of each concept for the selected trade-off criteria is difficult, as the performance of each concept is highly dependent on their specific implementation. Therefore, concepts are scored with respect to phase correlation method.

Table 4.2: Image navigation concept trade-off scores.

	Computational complexity	**Registration accuracy**	**Image lines needed**	**Reference database size**	**Final score**
Weight	0.4	0.2	0.3	0.1	1
Phase correlation	0	0	0	0	0
Intensity correlation	+1	0	0	0	+0.4
Contour matching	-1	+1	-2	+1	-0.7
Curve matching	-1	+2	-1	+1	-0.2
Local point matching	-1	+1	-1	+1	-0.4
High information point	-1	+1	-1	+2	-0.3

4.3.4. Final concept

The final image navigation concept selected is based on intensity correlation of the captured image lines. The sensed imagery will be correlated against an on board database of reference imagery. The selected concept still has considerable freedom in the design of, for example, the exact correlation algorithm or in the way the registrations are converted to attitude measurements. The exact implementation used in the image navigation mode is designed in detail in the next chapter.

The design philosophy behind the selected concept is to limit the number of image lines needed to perform a registration as much as possible. This allows a high measurement frequency, which allows a significant improvement in real-time attitude determination accuracy by filtering of the attitude measurements. The goal is to design a registration method capable of registering individual image lines, thus achieving the highest attitude measurement frequency possible.

5

Image navigation mode algorithm design

The algorithms needed to perform the image registration and attitude measurements will be designed in detail in the course of this chapter. Furthermore, the filtering method needed to improve the achievable attitude determination accuracy and provide real-time attitude knowledge is selected.

5.1. Correlation algorithm

The basis of the image navigation design is the correlation algorithm used to perform the image registration.

5.1.1. Correlation algorithm selection

The number of correlation algorithms in existence is vast. A large number of general correlation methods have found their way into the image registration domain. At the same time dedicated correlation algorithms have been designed often tailored to a specific registration task. Implementing all possible candidates to perform tests on relevant imagery was deemed impossible given the scope of this project and its schedule constraints. Instead, an extensive literature study was performed to investigate the operating principle and performance of a large number of correlation algorithms.

Feasible candidate selection
As a first step, a list of feasible candidates was made, by eliminating algorithms incapable of registering multi-modal imagery with significant intensity variations. The ability to register multi-modal images and possessing a high resistance against intensity variations in an image are both key functionalities required for this design. Yet, they are difficult to achieve, meaning only a small minority of algorithms possesses this functionality.

Multi-modallity, as previously discussed, is needed since SPECTROLITE imagery will be registered against an on-board database of geo-located imagery captured by another space-born instrument. Furthermore, SPECTROLITE will capture imagery under a large range of lighting conditions. It is undesirable to represent all of these lighting conditions in the database of reference imagery, as this would increase the complexity of the system and the database size would grow to an unacceptable size. Instead, the aim is to use a single set of reference imagery.

Trade-off
To select a final correlation method for use in the image navigation mode, a trade-off was performed. The feasible candidates were judged based on four criteria:

- computational complexity
- resistance to intensity variations

- resistance to noise
- required image size for successful correlation

The performance of candidates in each of these categories was judged using a number of comparative studies. These studies determine performance by means of image registration simulations using various types of satellite imagery with different levels of noise and intensity variation between the reference and sensed image. However, no comparative study was found including all feasible candidates. Instead, candidate performance was judged compared to the performance of the Pearson correlation coefficient algorithm, a popular method found in a large number of comparative studies.

In this trade-off, the four top-ranking candidates turned out to have very similar scores. A significant uncertainty exists in the scoring of candidates and analysis showed a large sensitivity of the overall winner to the exact criteria weights. Therefore, these four candidates are all included in the image registration simulation presented in Chapter 7, allowing further analysis of their performance under relevant conditions. A more extensive summary of the correlation algorithm trade-off, including a list of algorithms considered and the final trade-off scoring, is available in Appendix A.

5.1.2. Selected algorithms

We will now briefly discuss each of the selected correlation methods. These selected methods are:

- Pearson correlation
- Incremental sign distance
- Kendall's tau
- Spearman's rho

Pearson correlation coefficient

This correlation coefficient has been proven by Pearson in [33] to provide the best possible correlation between two sequences of numbers. Although first used by Pearson in the mathematical theory of evolution, this correlation method is also widely used in image registration.

Given the pixel intensity sequences $X = \{x_i : i = 1, \dots n\}$ and $Y = \{y_i : i = 1, \dots n\}$, the Pearson correlation coefficient is defined by:

$$r = \frac{\sum_{i=1}^{n}(x_i - \overline{x})(y_i - \overline{y})}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_i - \overline{x})^2}\sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \overline{y})^2}} \tag{5.1}$$

Its value will range from -1 to 1, with 1 being a perfect positive correlation and -1 a perfect negative correlation [16].

Due to the simplicity of the computation, performing a Pearson correlation between two images requires only a small number of additions and multiplications per pixel pair. Therefore, given two images with n pixels each, the computational complexity of a single correlation is also of order n [16].

Because the pixel sequences of both the sensed and reference image are normalised in the computation of the Pearson coefficient, this correlation method performs very well when correlating images with varying intensities. However, this is only true when the intensities in both images have some linear relationship. Highly non-linear variations in image intensities between two images can lead to significantly lower values for the correlation coefficient [16]. Noisy and smoothed images also produce lower correlation coefficients, yet this algorithm still performs relatively well compared to many of the other methods described in this section [16].

Incremental sign distance

Given intensity sequences $X = \{x_i : i = 1, \dots n\}$ and $Y = \{y_i : i = 1, \dots n\}$, let binary sequence $B_x = \{x_i : i = 1, \dots n-1\}$ be the sign of the difference between the intensity of a pixel and the intensity of its neighbouring pixel, i.e. the sign of $x_i - x_{i+1}$. Similarly, create binary sequence B_y. We can then use the Hamming distance between B_x and B_y to determine the similarity between X and Y

[32]. The Hamming distance between two sequences is the minimum number of substitutions required to change one sequence into the other. For the two binary series the Hamming distance is the number of ones in B_x XOR B_y.

The incremental sign distance method once again requires on the order of n comparisons, additions and subtractions [16]. A great advantage of this method is that by using the sign of intensity changes from one pixel to the next, this method is virtually insensitive to intensity variations. This is true for both linearly and non-linearly related intensity variations, but also for local variations caused by shadows or occlusions. Furthermore, this method works well with low contrast scenes, meaning highly smoothed scenes can still be registered [32]. However, this method is sensitive to noise in areas where the intensity variations between neighbouring pixels are small [16].

Kendall's tau

Given intensity sequences $X = \{x_i : i = 1, \dots n\}$ and $Y = \{y_i : i = 1, \dots n\}$, then for $i \neq j$ either $sign(x_j - x_i) = sign(y_j = y_i)$ or $sign(x_j - x_i) = -sign(y_j = y_i)$. The first case is called concordance and the second case discordance. Having more concordance than discordance is an indication of a high similarity between images, as intensity pairs change together [24].

Kendall's tau is a very computationally intensive similarity measure. If the correlated images each have n pixels, it requires the calculation of the concordants and discordants for $\frac{n(n-1)}{2}$ combinations of pixel intensity pairs. This gives a required number of operations of order n^2, even higher than the order $nlog_2 n$ operations for Spearman's rho, which we will later discuss [16].

However, this method also has some advantages. It has been shown that if bivariate (X, Y) is normally distributed, Kendall's tau τ is related to the Pearson correlation coefficient r by $r = \sin(\pi \frac{\tau}{2})$. This sinusoidal relationship means Kendall's tau can better distinguish images with high similarity as τ moves more rapidly away from values close to one (perfect correlation) compare to r [27]. Although Kendall's tau's performance is similarly affected by intensity variations, noise and smoothing, its performance compared to the Pearson method is superior for highly correlated images. It has been shown by [16] that this causes Kendall's tau method to also outperform the Pearson coefficient method as the size of the sensed and reference image become smaller. This may it to perform registration of images using only a subset of the two images, which could reduce the number of computations required. Similarly, it may allow for more frequent updates of the satellite's attitude as images with fewer image lines can be registered.

A method to further improve the performance of Kendall's method with noise images is to apply an intensity based weighting factor. When low intensity values are more affected by noise than high intensity values or vice-versa, more weight can be placed on the concordance of either high or low noise intensities [38]. This can be especially useful for the TROPOLITE instrument, as the signal power received by the instrument is relatively low, meaning noise factors independent of input power are dominant. This can cause low intensity measurements to be especially noisy.

Spearman's rho

The Spearman rank correlation or Spearman's rho is another similarity measure that related to the Pearson coefficient. Spearman's rho is obtained when the intensities in intensity sequences $X = \{x_i : i = 1, \dots n\}$ and $Y = \{y_i : i = 1, \dots n\}$ are replaced with their respective ranks $R(X)$ and $R(Y)$ and the Pearson coefficient of these rank vectors is computed. It was first proven to be a valid correlation method by Spearman in [39]. Mathematically, this can be represented as

$$\rho = 1 - \frac{6 \sum_{i=1}^{n} [R(x_i) - R(y_i)]^2}{n(n^2 - 1)} \tag{5.2}$$

n

With $R(x_i)$ and $R(y_i)$ representing the ranks of pixel intensities x_i and y_i respectively and n being the number of pixels being correlated [9].

Although its computational complexity is much higher, it was shown in [2] that Spearman's rho performs consistently better than Pearson's method when correlating images with large intensity differences, shadowing, smoothing, noise and other random differences. This is especially true when intensity variations are non-linearly related and when impulse noise is applied.

It has also been shown that when X and Y are independent, the ratio between Pearson's rho and Kendall's tau $\frac{\rho}{\tau}$ approaches $\frac{3}{2}$ as the number of pixels in an image n approaches infinity [25]. This result implies that the Pearson coefficient and Spearman's rho have the same discrimination power when comparing large images. However, as image size decreases, Spearman's rho will outperform Kendall's tau, which in turn was shown to outperform Pearson's method when image size is small. However, each time at the cost of an increase in computational complexity.

5.2. Image rigidity

Image rigidity is an important concept in image registration. For this purpose, an image is considered rigid if it is only subject to affine transformations. Affine transformations, like translation, rotation and scaling, are linear in nature and uniform over an entire image. SPECTROLITE imagery is subject to affine transformations due to the movement of the satellite in orbit (translation), yawing (rotation) and changes in local altitude (scaling).

Non-rigid imagery is subject to non-rigid or elastic transformations. Elastic transformations are non-linear and local in nature. They will therefore severely increase the complexity of the registration task. To understand why this is the case, we look at correlation based image registration in its most basic form. This is essentially a brute force method, computing the correlation between a sensed image and a number of reference images. Each reference image represents a specific value for each degree of freedom of the registered imagery. An image only subjected to a transformation will have two degrees of freedom for the entire image. Both rotation and scaling will add another degree of freedom to the registration problem. If the number of values to be tested is equal for each degree of freedom, we see the number of correlations to be performed increases quadratically with an increase in degrees of freedom. However, since elastic transformations are not uniform, they do not have a single degree of freedom for the entire image. Instead, they allow relative movement of pixels within the image, providing each individual pixel with its own degrees of freedom. However, generally the relative movement of pixels is limited. Therefore, transformation models like surface splines or deomorphisms are used to limit the total degrees of freedom of the registration problem. Still, non-rigid imagery will have significantly more degrees of freedom compared to rigid imagery. Owing to the quadratic increase in correlations with degrees of freedom, this has a large impact on the number of correlations required for successful registration and thus on the computational complexity of the registration task.

5.2.1. Rigidity of satellite imagery

Although the Earth's surface can be considered rigid for the scale and time-span of this problem, satellite Earth imagery is not. The cause of this non-rigidity lies in the range of off-nadir viewing angles of the imager in combination with the curvature and irregularity of the Earth's surface. Although the curvature effect is small can be estimate based on available knowledge of the Earth's curvature, the effect due to surface irregularity is more signification. Due to off-nadir viewing angles, a vertical variation in surface elevation can cause a horizontal displacement of point on Earth from the perspective of the imaging sensor. This effect is visualised in Fig. 5.1.

Given the effect of non-rigid imagery on the computational complexity of the registration task, we create artificial rigidity by determining the absolute location of each captured image pixel based on the instrument geometry and knowledge of the surface elevation.

5.2.2. Linear pushbroom model

The linear pushbroom model is used to compute the location of each individual pixel in an Earth Centred Earth Fixed (ECEF) reference frame. This is a process we will refer to as pixel coordinate mapping. The equations used are a simplification of the Toutin model, originally developed for the

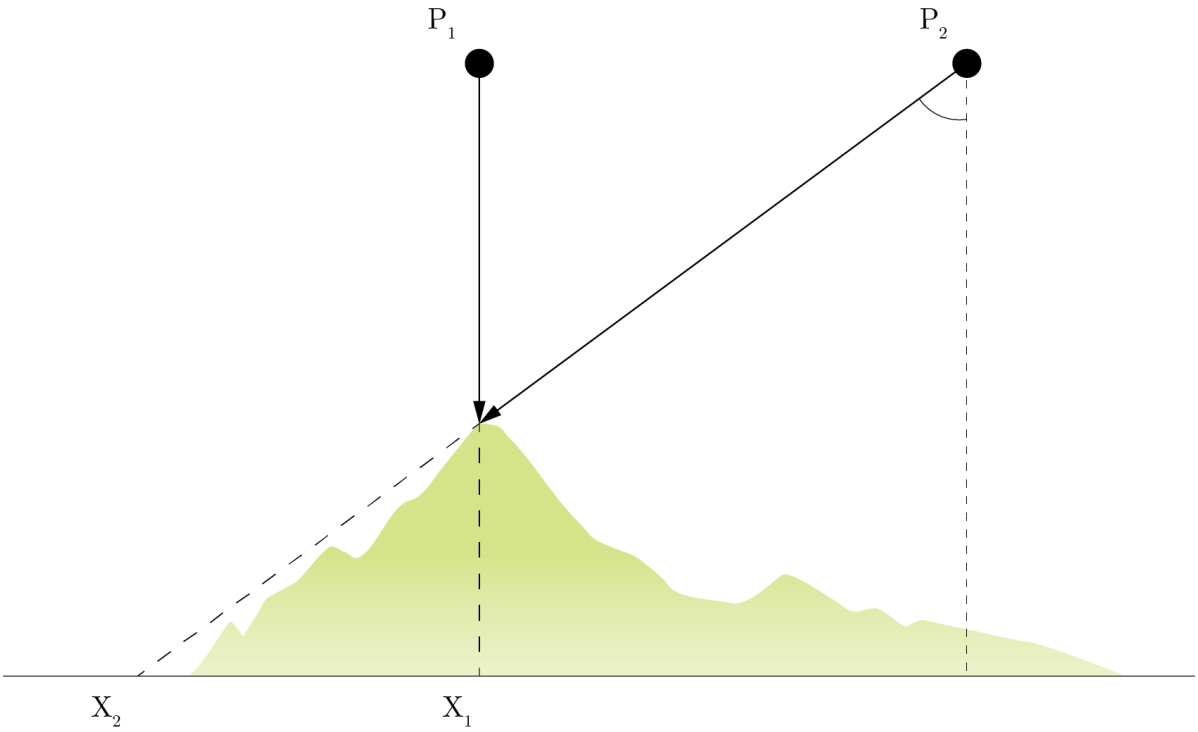

Figure 5.1: Visualisation of the Earth curvature effect.

SPOT-1 satellite [44]. Its goal was to use ground control points and the post-processing of images to improve a-posteriori knowledge of the satellite's attitude and orbit. The exact equations used are for a simplification of the original model described by [17]. This simplification assumes the satellite is in an unperturbed circular orbit. The equations have been simplified further for use in an instantaneous moment in time, assuming knowledge of the satellite's position.

Camera frame

The first reference frame we define is the camera frame, as can be seen in Fig. 5.2. This reference frame has the camera optical centre as its origin. The focal plane of the camera is located at the focal distance f from the optical centre. The z-axis runs perpendicular to the focal plane and crosses the focal plane at the principal point. Assuming symmetry this is the middle of the detector array. The y-axis runs parallel to focal plane along the direction of imaging line. The x-axis also runs parallel to focal plane, but in the direction of flight. Together these axes form an orthonormal right-handed coordinate frame.

The 3D point X corresponds to the imaged pixel, located in the camera frame at $X = (0, Y, Z)^T$. The pixel is mapped to point $(0, fY/Z, f)^T$ in the focal plane, as can be seen in Fig. 5.3.

We now define a sensor frame, as can be seen in Fig. 5.4. We define y as the distance from the start of the sensor, measured in pixels, and w as the pixel width. The distance of the principle point from the start of the sensor array in pixels is defined as y_0. This gives fY/Z is equal to $w(y - y_0)$ in the sensor frame.

This means the 3D point X describing the location of the pixel in the sensor frame can be found along the vector $u = (0, w(y - y_0), f)^T$. This is only true if there is no distortion of the camera geometry due to instrument optics. Although this distortion is minimised in the instrument design, some distortion will likely be present. However, this distortion can accurately be measured once the instrument is build. A calibration term can then be added to account for this distortion.

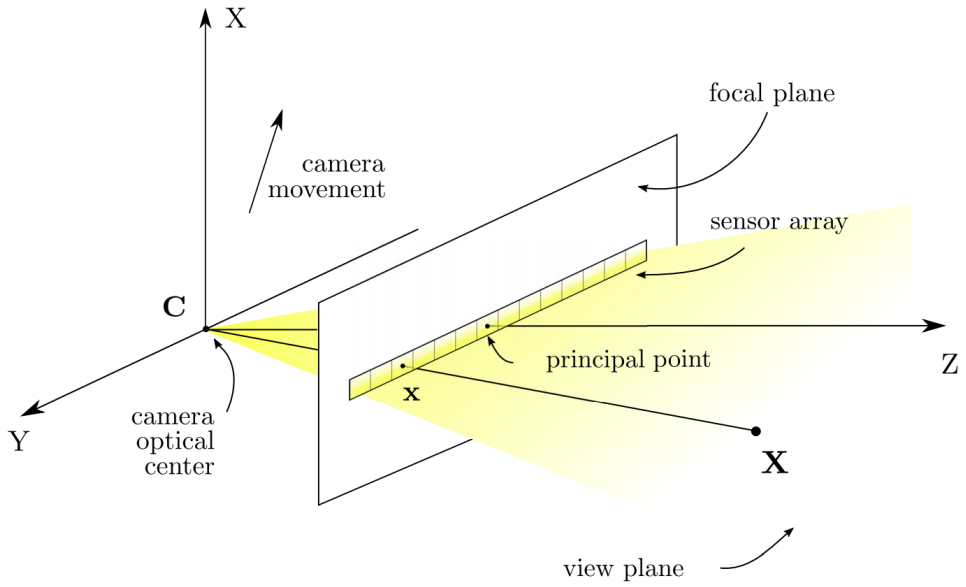

Figure 5.2: Camera reference frame [10].

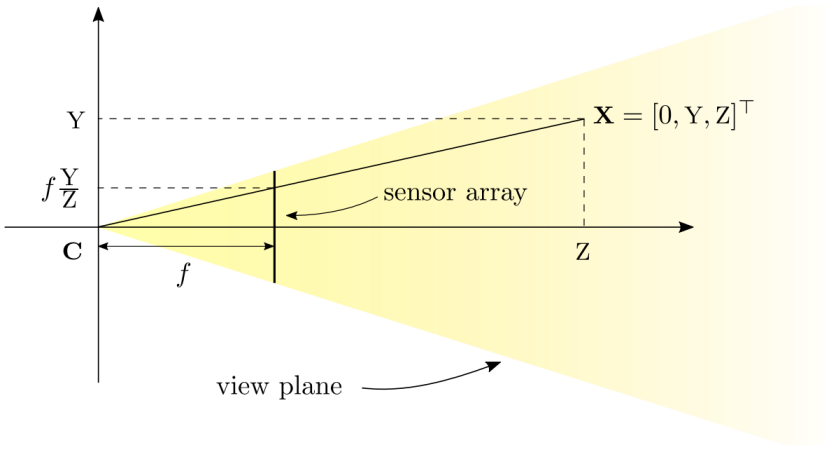

Figure 5.3: Camera reference frame - side view[10]

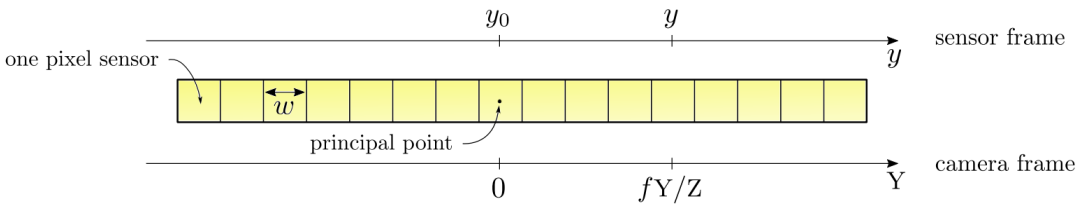

Figure 5.4: Camera and sensor frame [10].

LVLH frame

We now apply a change of reference frame to determine the vector u in the Local Vertical Local Horizontal (LVLH) frame. The origin of the LVLH frame is the satellite's centre of mass. Its z-axis points towards the centre of the Earth, while the x-axis once again runs along the direction of flight of the satellite. The y-axis is defined such that an orthonormal right-hand coordinate frame is formed.

By assuming the optical centre coincides with the satellite's centre of mass, the origins of the two reference frames coincide. This means a change between the camera frame and orbital frame is achieved using the standard rotation matrices for roll angle ϕ, pitch angle θ and yaw angle ψ, as can be seen in Fig. 5.5.

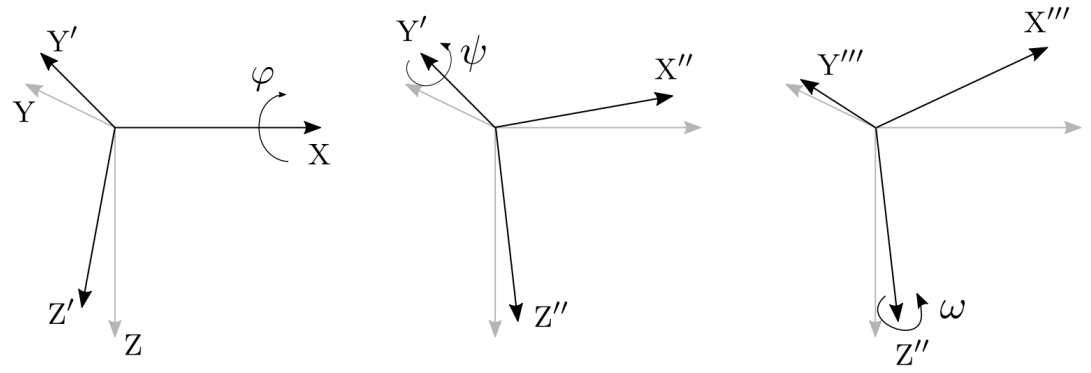

Figure 5.5: Rotation from camera frame to LVLH frame[10].

$$R(\phi,\theta,\psi) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_\phi & -s_\phi \\ 0 & s_\phi & c_\phi \end{bmatrix} \begin{bmatrix} c_\theta & 0 & s_\theta \\ 0 & 1 & 0 \\ -s_\theta & 0 & c_\theta \end{bmatrix} \begin{bmatrix} c_\psi & -s_\psi & 0 \\ s_\psi & c_\psi & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{5.3}$$

We can now define u in the LVLH frame.

$$u = R(\phi,\theta,\psi) \begin{pmatrix} 0 \\ w(y - y_0) \\ f \end{pmatrix} \tag{5.4}$$

ECEF frame

Our last change of reference frame is from the local orbital to the ECEF frame, as can be seen in Fig. 5.6. We assume the satellite's orbit is circular around the centre of mass of Earth which coincides with origin or ECEF frame. The z-axis of the ECEF frame is the Earth's rotational axis. The x-axis runs through the location of zero degrees latitude and zero degrees longitude on the Earth's surface. Th y-axis is once again defined such that an orthonormal right-hand coordinate frame is formed.

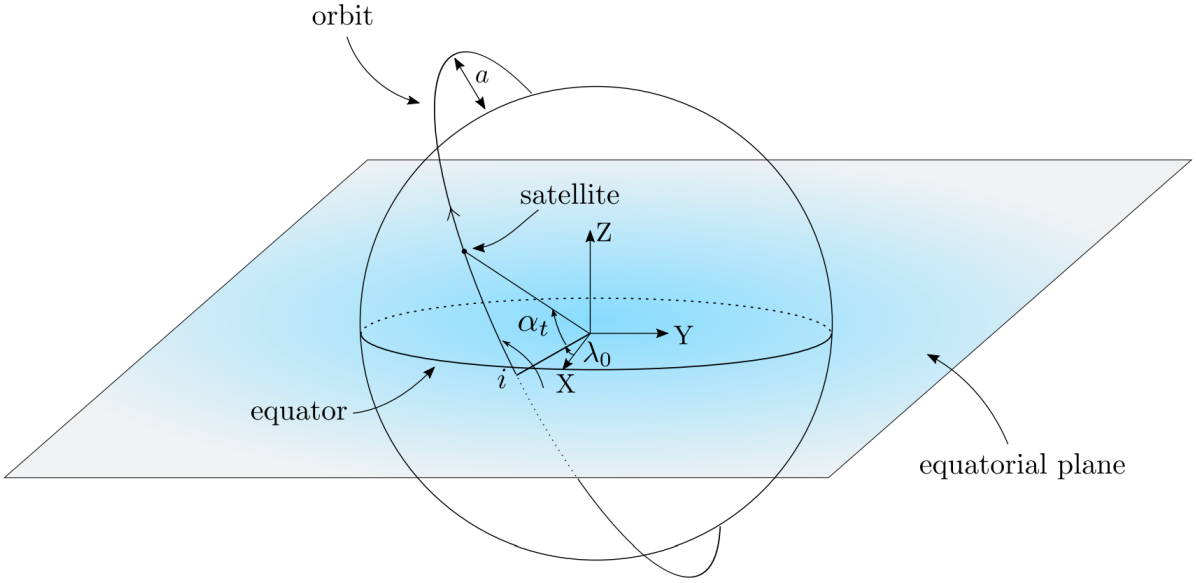

Figure 5.6: Definition of the ECEF frame [10].

A circular orbit around the Earth is described by three parameters. The first parameter is the orbit radius $R_E + a$, with R_E the Earth's radius and a the altitude of the satellite. The second parameter

is the longitude λ, defined with respect to the x-axis of the ECEF frame. The third parameter is the inclination i, which is the angle between the equatorial plane and the orbital plane. A fourth parameter can now be used to locate the satellite in its orbit. This parameter α describes the angular position of the satellite with a zero value at the equator.

All four parameters can be computed from a known satellite position determined by means of, for example, a Global Navigation Satellite System (GNSS) or by propagating a Two-Line Element set (TLE). However, this does mean that an error in the knowledge of the satellite position will lead to an error of the absolute satellite attitude determined by the image navigation system. For small satellites, real-time orbit determination can be performed with an accuracy of under 2 m [15]. This will cause a negligibly small error in the attitude measurement using the image navigation system. In the course of this document we will therefore assume the satellite's position in orbit is known exactly. Still, the need for accurate satellite position measurements will be important to consider in a future detailed design of an image navigation system.

The rotation matrix P from local orbital to ECEF frame is now defined as

$$P_{ECEF,orbital}(\lambda, i, \alpha) = \begin{bmatrix} c_\lambda & -s_\lambda & 0 \\ s_\lambda & c_\lambda & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{i-\frac{\pi}{2}} & -s_{i-\frac{\pi}{2}} \\ 0 & s_{i-\frac{\pi}{2}} & c_{i-\frac{\pi}{2}} \end{bmatrix} \begin{bmatrix} c_{-\alpha-\frac{\pi}{2}} & 0 & s_{-\alpha-\frac{\pi}{2}} \\ 0 & 1 & 0 \\ -s_{-\alpha-\frac{\pi}{2}} & 0 & c_{-\alpha-\frac{\pi}{2}} \end{bmatrix} \tag{5.5}$$

This allows us to define u in the ECEF frame.

$$u = P_{ECEF,orbital}(\lambda, i, \alpha) R(\phi, \theta, \psi) \begin{pmatrix} 0 \\ w(y - y_0) \\ f \end{pmatrix} \tag{5.6}$$

Furthermore, the location of the satellite in the ECEF frame S can be defined

$$S = P_{ECEF,orbital}(\lambda, i, \alpha) \begin{pmatrix} 0 \\ 0) \\ R_E + a \end{pmatrix} \tag{5.7}$$

Localisation
We can now compute the 3D point X corresponding to the location of the imaged pixel. X is located at a height h above the Earth's radius. This means the location of X can be found at the intersection of the line passing through S and directed by u, with the sphere of radius $R_E + h$ coinciding with the Earth's centre. This intersection point is found at

$$||S + ru||^2 = (R_E + h)^2 \tag{5.8}$$

This equation has two positive roots, with the smallest found at

$$r = \frac{-S^T u - \sqrt{(S^T u)^2 - u^T u (S^T S - (R_E + h)^2)}}{u^T u} \tag{5.9}$$

The location of X can now be computed using

$$X = S + ru \tag{5.10}$$

The correct location X cannot be evaluated directly as the local value of h is needed. Therefore, an iterative method must be used. By starting at some value for h a first estimate of X can be computed. The local value of h at this estimated located can then be used to perform another iteration. This process is then repeated until convergence is reached. To limit the number of iterations required, the

location of the centre pixel can be computed first. Since the local viewing angle of this pixel is near zero, the dependence of its computed location on h is negligibly small. The value of h found for this pixel can then be used in the first evaluation of the location of its neighbouring pixel. This process is then repeated working outwards along the detector length.

Digital elevation model
It is important to note that knowledge about surface elevation must be available on-board to perform the pixel coordinate mapping. Luckily, surface elevation is well documented. Digital Elevation Models (DEM) provide gridded elevation data with respect to a reference geoid. Highly accurate commercial DEMs exists with a grid size of 12 by 12 meters and an absolute elevation accuracy of 4 meters, providing coverage of nearly the entire globe [50]. Free alternative exist like the ASTER DEM with a 30 by 30 meter grid and an absolute vertical accuracy of 17 meters [42]. Depending on the resolution of the DEM used, its storage capacity needed can be significant and will likely exceed that of the reference imagery. Therefore, the size of the DEM will be taken into account in the sizing of the image navigation system in Section 10.2.

We will evaluate the effect of both the DEM accuracy and the number of iterations in Section 7.1.

5.3. Registration to attitude measurement

The pixel coordinate mapping equations can now be used to determine the location of each captured pixel based on some initial knowledge of the three attitude angles, provided by propagation of a previous satellite state. We will call this location the initially estimated location of the image line. This initially estimate location can now be used to create a reference image line corresponding to this estimated attitude. Additional reference image lines can be created by varying the attitude over a range of angles around the initial estimate. We will call the range of angles the correlation window. If each of these image lines is then correlated to the sensed image line, the attitude corresponding to the reference image with highest correlation can be used as a new attitude measurement.

5.3.1. Minimising correlations

By means of the coordinate mapping equations, rigid-like imagery was created, limiting the degrees of freedom of the registration task to the three attitude angles. However, three degrees of freedom can still lead to a large number of required correlations, due to the quadratic relationship previously described. In case we have a total of ten values we want to test for each separate attitude angle, the total number of correlations required to test all possible attitude angle combinations is $10^3 = 1000$.

To further limit the number of correlations needed to measure the three attitude angles, we try to uncouple the three angles. We do so by using tow trigonometric properties of small angles: $cos(\psi)d \approx d$ and $sin(\psi)d \ll d$.

Imagine a triangle ABC with AB being half of an image line and AC being the same half image line after some yaw angle ψ, see Fig. 5.7. We then add another triangle BCD adjacent to the original triangle. The length BC corresponds to the movement of a pixel from the original image line AB to the image line BC. BD now corresponds to the pixels vertical movement and CD to its horizontal movement.

Since ABC is an isosceles triangle, $\angle ABC = \angle ACB$. This means

$$\angle ABC = \frac{180^\circ - \psi}{2} = 90^\circ - \frac{\psi}{2} \tag{5.11}$$

Since alternate angles in parallel lines are equal (the z-rule), $\angle BCD = \angle ABC$. As BCD is a right-angled triangle, we find

$$\angle CBD = \frac{\psi}{2} \tag{5.12}$$

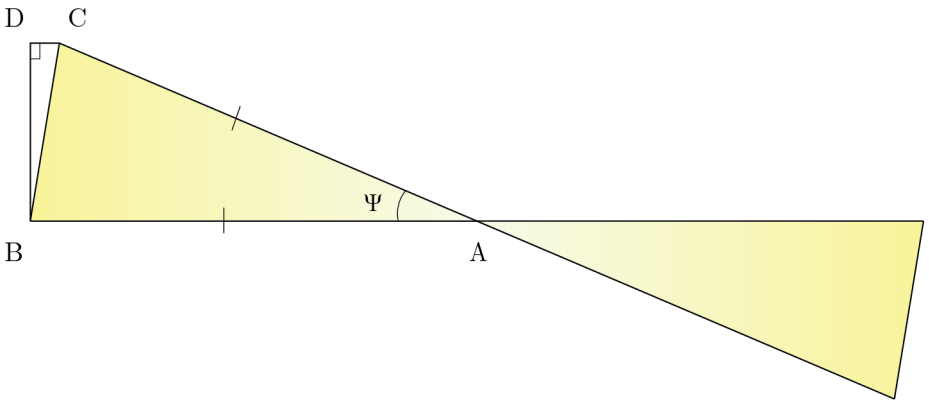

Figure 5.7: Yaw angle induced geometry change.

Imagine a maximum tested yaw angle value ψ of $0.5°$. This roughly corresponds to a deviation of 4 km at the outer pixel of the image line, which is significant given the required accuracy of 100 m. If the pixel at B moves to location C by some dimensionless distance of 100 we find

$$BD = cos(\psi/2)100 = 100.00 \tag{5.13}$$

$$CD = sin(\psi/2)100 = 4.36 \tag{5.14}$$

We can therefore try to estimate the yaw angle based on the vertical movement of a pixel section at the outer edge of the image line. In fact, a one-dimensional correlation perpendicular to the image line at each image line end could serve as an estimate of both yaw and pitch angle, see Fig. 5.8.

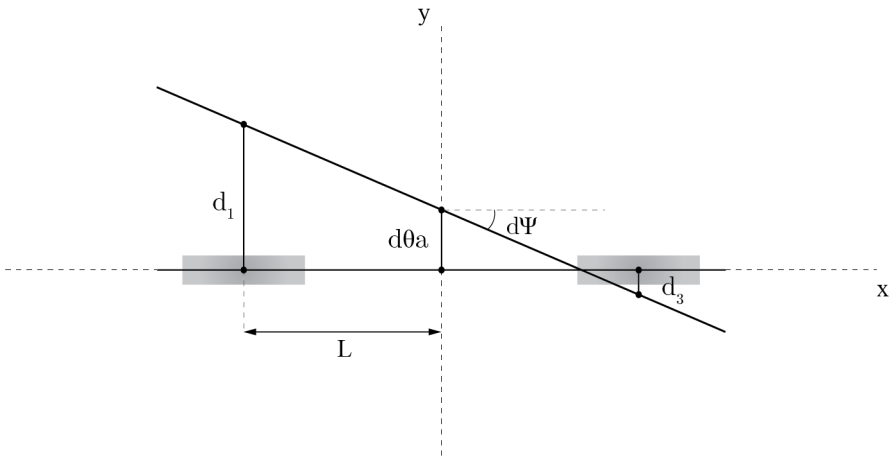

Figure 5.8: Pitch and yaw angle measurement using one-dimensional correlation.

We define d_1 to be the distance between the centre of the left-hand side correlation area, at the location of highest correlation, and its initially estimated location. Similarly, we define d_3 for the right-hand side correlation area. This gives

$$\Delta\theta a = \frac{d_1 + d_3}{2} \tag{5.15}$$

$$\Delta\psi = tan^{-1}(\frac{d_1 - d\alpha a}{L}) = tan^{-1}(\frac{d_1 - d_3}{2L}) \tag{5.16}$$

The measured pitch and yaw angles are therefore given by

$$\theta_{measured} = \theta_{initial} + \frac{d_1 + d_3}{2a} \tag{5.17}$$

$$\psi_{measured} = \psi_{initial} + tan^{-1}(\frac{d_1 - d_3}{2L}) \tag{5.18}$$

The roll angle could then be estimated based on another one-dimensional correlation in horizontal direction at the centre of the image line. If we define a distance d_2 similar to d_1 and d_3 we find

$$\phi_{measured} = \phi_{initial} + \frac{d2}{a} \tag{5.19}$$

5.3.2. Simplified coordinate mapping

Since the pixel coordinate mapping equations have to be evaluated for each individual pixel within the image line, the total computational effort to compute all pixel coordinates for a single image line is already high. Repeating this process at each correlation step will require a computational effort far exceeding that of the correlation algorithm itself, as we will later discover. We therefore use the coordinate mapping equations to determine the coordinates at the initially estimated image line location only.

Instead, for the left and right-hand side correlation areas, we compute the pixel coordinates for each correlation step reference image based on a certain displacement perpendicular to the image line from their initially estimated location. This displacement corresponds to the angular displacement $\Delta\theta$. This essentially means that, for each correlation step, we move each pixel coordinate in the correlation areas by a varying distance d_1 or d_3 as defined in Fig. 5.8. Similarly, we vary the pixel coordinates in the central correlation area by a distance d_2 corresponding to the angular displacement $\Delta\phi$.

The main source of error caused by this simplification is due to a change in surface elevation along the direction of correlation. However, since our correlation window is small, this error is expected to be sufficiently small to still allow accurate registration. The error caused by this simplified coordinate mapping will be quantified in Section 7.1, where we will also verify its negligible effect on registration performance.

5.3.3. Sub-pixel registration

Registering a sensed image to a reference image without addition of a sub-pixel registration scheme limits the achievable accuracy to one pixel. However, for many applications, sub-pixel registration is needed. Therefore, a large number of methods have been created to improve the registration accuracy beyond the one pixel mark. For feature detection methods with multiple detected features, we can use geometric methods and cost function optimisation. These methods combine the registration of individual features in an image to achieve a sub-pixel registration for the complete image. However, for correlation methods, the options are somewhat limited. Roughly three methods of reaching sub-pixel registration exist: intensity interpolation, correlation interpolation and differential methods [37] [43]. Differential methods required multiple images of the same scene and are therefore generally applied to registration in video feeds [4]. In case of a two-dimensional imager, it might be possible to use the overlap of subsequent images in a differential method. However, differential methods cannot be applied to a one-dimensional push-broom imager. This leaves us with the two types of interpolation methods.

Intensity interpolation

In an intensity interpolation method, the sensed image is resampled with a higher spatial frequency to increase the image resolution. Registering the resampled image with a one pixel accuracy results

in a sub-pixel accuracy for the original image. However, resampling of the image will only work if the original image is sampled at a frequency higher than the Nyquist frequency. Otherwise, the high spatial frequency information is not stored. The sampling resolution for SPECTROLITE is already on or close to the limiting resolution of the optics and detector. This means the registration accuracy improvement using an intensity interpolation method will be limited. To make the most of the available spatial frequency of the instrument, an advanced interpolation method like bi-cubic interpolation is needed [43]. This will require a significant amount of additional processing power to perform the registration task.

Correlation interpolation
The other interpolation method to achieve sub-pixel registration is to interpolate the correlation values found at each pixel location. As we will discover in Section 7.3.1, the correlation values will form a Gaussian distribution peaking at the location matching the sensed to the reference image. With this knowledge, we can interpolate the correlation values found with a Gaussian distribution and determine its peak with sub-pixel accuracy. However, as we will also see, the correlation values found for the correlation of a single image line contain a lot of noise. This severely limits the accuracy of the interpolation, making sub-pixel registration using correlation interpolation impossible.

Alternative method
Instead of the commonly used interpolation methods, an alternative sub-pixel registration scheme was designed, especially useful in cases where minimising the computational cost is essential. This method can be described as a combination of the intensity interpolation and a geometrical methods. The main advantage of this method is that no interpolation of the captured imagery is needed.

Using the pixel coordinate mapping equations we determine an estimated location for each individual pixel in the captured image line. We match these sensed image pixels to the closest positioned reference image pixels by comparing pixel coordinates. This allows us to build the first reference image corresponding to the initially estimate image line location. The sensed image pixel grid is not equal to the reference image pixel grid, caused by a varying sensed image pixel sizes and a rotation of the sensed image line due to the inclination of the satellite. Therefore, the distance between the centre of a sensed pixel and and the centre of its matching reference pixel differs from pixel to pixel. If we now add a sub-pixel offset to the sensed pixel location, the distances between the sensed pixel centres and their matching reference pixel centres changes. Depending on the initial distance between them, some sensed pixels will have a new closest matching reference pixel, while other sensed pixel will retain their original matching reference pixel. This way, the reference image will slowly start to change when the offset to the initial estimated image coordinates is increased. This process is visualised in Fig. 5.9, where we see the partial change in reference image with a sub-pixel offset in the along-track sensed image line location. The two thick black lines represent typical sensed image line shapes in an equirectangular projection, separated by a sub-pixel offset. The underlying grid represents the reference imagery, with the shaded boxes indicating the reference pixels matching to the sensed image line pixels. The boxes with a red outline indicate the partially changed reference image caused by the sub-pixel image line offset.

By using reference imagery with a pixel pitch and/or resolution smaller than the sensed imagery will increase the change in reference image with an increasing offset to the original estimated image pixel coordinates. Using this method, sub-pixel registration accuracy can be achieved without the need for any interpolation procedure on board the satellite, resulting in a very computationally efficient sub-pixel registration method.

5.4. Filtering and propagation

Generally, measurements provided by an attitude sensor are not directly inserted as measurements into a control loop. Instead, filtering is applied to improve the quality of measured state knowledge. For image based attitude determination methods, another reason exists why attitude measurements

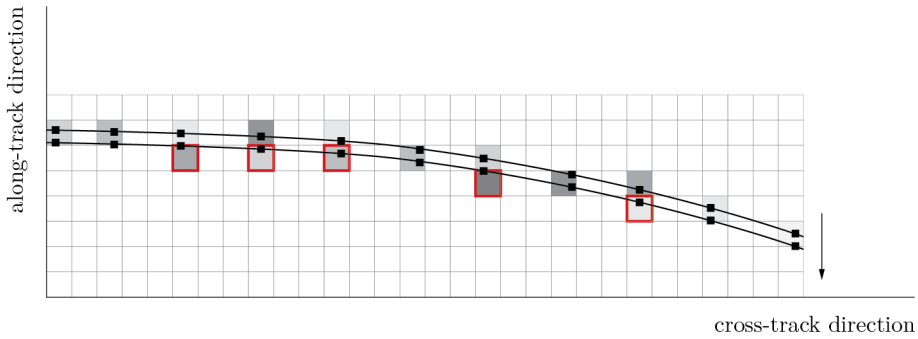

Figure 5.9: Schematic representation of the sub-pixel registration method used.

cannot be directly inserted into the control loop. The time required for the correlation process is significant, leading to a delay between the availability of the measurement and the time of its corresponding state. This generates a control error, since the controller is responding to an outdated measurement, which can deviate significantly from the true state. State propagation, estimation of a state based on information about previous states, is therefore needed to improve the real-time attitude knowledge available to the attitude controller.

5.4.1. Kalman filter principle

The current de facto standard for filtering of attitude measurements is the Kalman filter, otherwise known as the linear quadratic estimation technique. It combines sensor measurements with theoretical state estimations based on a physical model. This physical model can include the dependencies between individual state variables.

This property allows knowledge of one state variable to be improved using measurements of another state variables, as long as the two are physically dependent. The attitude knowledge can therefore be improved using rotational rate measurements, allowing improved attitude determination accuracy when the limits of the image based method are reached. Incorporating cheap, small, COTS sensors as secondary measurement devices, can relax the requirements on the image navigation system.

Many versions of the Kalman filter exist, each one tailored to specific filtering and propagation requirements. As a basis for discussing more elaborate Kalman filters capable of incorporating delayed measurements, a basic linear Kalman filter is discussed. This example is created for a single axis, in this case the pitch axis θ .

First, the satellite's state at time t is estimated by propagating the previous state at $t-1$ using a description of the satellite's dynamics. Vector x is the state vector containing pitch angle θ and pitch rate $\dot{\theta}$.

$$x = \begin{bmatrix} \theta \\ \dot{\theta} \end{bmatrix} \tag{5.20}$$

Matrix A is the state transition model containing the satellite's dynamics. Propagation is achieved by multiplying the two.

$$\hat{x}(t_k) = Ax(t_{k-1}) \tag{5.21}$$

$\hat{x}$ is an a priori update of the state vector, meaning that it is based on measurements up to time $t-1$. Similarly an a priori update of the covariance matrix $\hat{P}$ is computed using covariance matrix P at time $t-1$.

$$\hat{P}(t_k) = AP(t_{k-1})A^T + Q \tag{5.22}$$

Matrix Q is the system noise covariance matrix, providing information on how well the filter is able to propagate a previous state using A. $\hat{P}$ is then used to compute the Kalman gain factor K. The gain factor determines to what extend the estimated state at time t is based on both propagation of the

previous state at $t-1$ and the state measurements at time t.

$$K = \hat{P}(t_k)H^T[H\hat{P}(t_k)H^T + R]^{-1} \tag{5.23}$$

Matrix H is the measurement equation matrix, relating the measurements in z to the state in x. If z contains the measurements of the states in x, H is an identity matrix. R is the measurement noise covariance matrix, containing the variances of measurements in z. The a priori state update is now combined with the state measurements at time t to compute the a posteriori state estimate.

$$x(t_k) = \hat{x}(t_k) + K[z(t_k) - H\hat{x}(t_k)] \tag{5.24}$$

Similarly, the a posteriori covariance matrix is computed, again based on the Kalman gain K.

$$P(t_k) = [I - KH]\hat{P}(t_k) \tag{5.25}$$

5.4.2. Fusing time-delayed measurements

As stated before, due to the processing time required for the registration of an image line, the delay between the availability of an attitude measurements and its corresponding state is significant. There are various methods to deal with this problem. The first option is to halt the filter until the measurement is available. For a significant delay, this means the attitude controller will constantly be acting on outdated attitude information. Therefore, a method must be found to keep the Kalman filter running, but still fuse the delayed measurements. This problem is referred to as an out of sequence measurement problem.

Fusing a time-delayed measurement is significantly more complex when the measurement delay is unknown [28]. In this case, the moment an image line is captured is time stamped by the instrument. For now we assume the exact moment of capturing the image is known. However, due to integration time required, image is not captured at an instantaneous point in time, but over a time-span equal to the integration time. We assume the state corresponding the the captured image is the state at a point in time in the centre of the integration time time-span. This is a valid assumption if the slew rates of the satellite during the integration time are constant and the jitter is symmetric.

Additional complexity occurs when measurements have to be fused corresponding to states in between the filter updates that have already been performed [28]. Given a constant imaging frequency, the time between capturing of the image lines is constant. This allows us to synchronise the Kalman filter updates to the moment an image line is captured.

Three methods were found to adjust the Kalman filtering algorithm to allow fusion of time-delayed measurements with a known time-delay at a discrete time corresponding to a filter update:

- **Augmented state Kalman filter:** This filter augments the original state vector with a number of previous states. This also augments the measurement vector, allowing fusion of measurements corresponding to all states present in the augmented state vector. This method provides an optimal solution to the time-delay problem, but its computational complexity increases rapidly with the number of augmented states [8].
- **Measurement extrapolation Kalman filter:** This filter provides a close to optimal solution for the time-delay problem by computing a correction factor for the time-delayed measurement. The correction factor is computed based on the gain matrices of previous filter updates. Using the correction factor, the delayed measurement can be incorporated into the filter as a non-delayed measurement. Computation of the correction factor is faster than augmenting the filter when the delay consists of a large number of filter updates [28].
- **Zhang Kalman filter:** This filter reduces computational complexity of fusing time-delayed measurements even further, however, at the cost of a loss in accuracy [55].

Timing sequence

To achieve the highest attitude determination performance, we want to be able to process all incoming image lines. This means the processing time required for a single image line cannot exceed the

imaging period. As we have seen before, this imaging period is equal to twice the instrument integration time.

Four steps are identified in the entire image navigation process. The first step is the capturing of the image line, governed by the integration time of the instrument. Next, the instrument data will be processed by the payload data handling system. This processing will include creating the intensity averaged image line and might also include additional processing tasks like computing the percentage of cloud coverage in the image line. From the perspective of the image navigation system, we will refer to this as pre-processing of the image line. The third step is the actual image registration and computation of the attitude angles. The fourth and last step is running the filtering algorithm, computing the real-time satellite attitude and controlling the reaction wheels accordingly.

As we discussed before, the state corresponding to an image line is the satellite state at the centre of the integration time time-span. Therefore, we synchronise the filter updates to the centre of the integration time, so the filter states correspond to the image line states. We now get an image navigation timing sequence as visualised in Fig. 5.10. The imaging period, being equal to the time available to perform the image registration, is indicated by dt. We make the assumption that the pre-processing time required is equal to the instrument integration time and that the filtering and control time is equal to half of this integration time. However, it is possible to trade computation time between these two processes at will in the detailed system design.

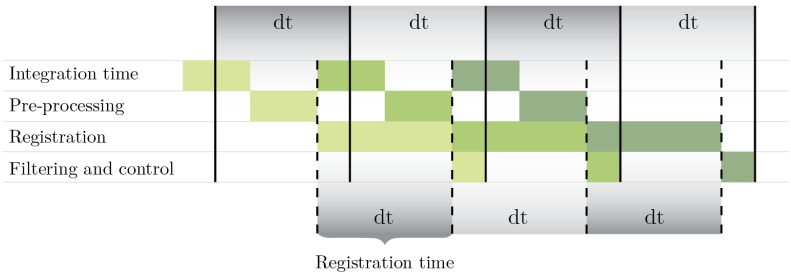

Figure 5.10: Visualisation of the timing sequence for imaging, pre-processing, registration and filtering.

From Fig. 5.10 we can also determine that the delay between the occurrence of the image line state and the availability of this state measurement is equal to twice the imaging period. Since we have synchronised the filter frequency to the imaging frequency, a total of three satellite states is cover this delay. Given the limited number of additional states needed, the augmented state Kalman filter is selected as filtering algorithm, given its optimal state estimation performance.

Augmented state Kalman filter equations

Augmenting the basic linear Kalman filter with a number of historic states is relatively simple. If the delayed measurement corresponding to the state at $t = \tau$ is received at the current time $t = t_k$, the augmented state is $[x(t_k), x(\tau)]^T$. The augmented state equations become [8]

$$\begin{bmatrix} x(t_k) \\ x(\tau) \end{bmatrix} = \begin{bmatrix} A_{t_k|\tau} & 0 \\ I & 0 \end{bmatrix} \begin{bmatrix} x(t_{k-1}) \\ x(\tau) \end{bmatrix} + \begin{bmatrix} w(t_{k|\tau}) \\ 0 \end{bmatrix} \tag{5.26}$$

$$\begin{bmatrix} y(t_k) \\ y(\tau) \end{bmatrix} = \begin{bmatrix} H_k & 0 \\ 0 & H_k \end{bmatrix} \begin{bmatrix} x(t_k) \\ x(\tau) \end{bmatrix} + \begin{bmatrix} v(t_k) \\ v(t_\tau) \end{bmatrix} \tag{5.27}$$

The noise covariance matrix matrix becomes

$$R_k = \begin{bmatrix} R_k & 0 \\ 0 & R_\tau \end{bmatrix} \tag{5.28}$$

The standard Kalman filter recursion can then be employed to obtain the augmented state estimate which is updated using both the delayed measurent $y(\tau)$ and the current measurement $y(t_k)$.

Part II

Simulation and verification

6

Image registration simulation data sources and set-up

The image registration simulation is the first of two simulations performed. It will be used to verify the performance of the one-dimensional registration method designed and select a final correlation method out of the four methods described in Section 5.1. In this chapter, the set-up of the image registration simulation will be discussed. The results of this simulation will be presented in Chapter 7. The image registration simulation will also form the basis for the image navigation simulation described in Chapter 8. This simulation will be used to assess the achievable attitude determination accuracy using the designed registration method in a realistic scenario. The results of this second simulation will be presented in Chapter 9.

6.1. Instrument data generation

To accurately assess the performance of the designed image registration method, instrument data closely resembling the data that will be produced by SPECTROLITE is needed. Therefore, this simulation will use data from an instrument closely resembling the SPECTROLITE instrument. The data from this source instrument is processed to further improve the resemblance to the data produced by SPECTROLITE.

6.1.1. Source instrument

To select a suitable source instrument for simulation data, a number of requirements for the source data are set. Based on these requirements a source instrument is selected.

Source instrument requirements

The following requirements are set for the source instrument.

- **Instrument type:** A key property of the SPECTROLITE instrument is its push-broom design. This means each individual captured image line pixels has a nadir viewing angle in along-track direction. A two dimensional sensor array causes increasingly off-nadir viewing angles for image lines further from the centre, similar to the off-nadir viewing angles experienced in cross-track direction.
- **Spectral band:** The spectral band of the imagery can potentially have an affect on the registration performance, since surface reflectance values are wavelength dependent. Therefore, a source instrument with at least one spectral band overlapping with a SPECTROLITE spectral band is needed.
- **Ground sampling distance and swath size:** Many of the imaging spectrometers currently orbiting the Earth have a Ground Sampling Distance (GSD) larger than SPECTROLITE. The

resolution of an image will have a large effect on the registration accuracy. Although the GSD of an image can easily be lowered, increasing the GSD is not trivial. Images with a native GSD equal or smaller than SPECTROLITE's 1 km value are therefore needed.

- **Noise and MTF:** The noise level and smoothing of an image are key image parameters influencing image registration performance. In optical terms, the noise and smoothing can be represented by respectively the Signal to Noise ratio (SNR) and Modulation Transfer Function (MTF). To achieve results representative for the SPECTROLITE instrument, the SNR and MTF levels of the instrument data used in the simulation must match those of SPECTROLITE. However, finding perfectly matching values is likely not possible. Instead, a source instrument with an SNR or MTF higher than SPECTROLITE can be used. The quality of the simulation imagery can then be worsened by applying artificial noise and smoothing functions.
- **Processing level:** Images with low levels of processing are needed, since processing performed on board the satellite will be limited. Especially orthorectified imagery must be avoided. Orthorectified images are resampled to create a constant pixel pitch, removing the non-rigidity effect discussed in Section 5.2.1.
- **Meta-data:** Additional image data is needed for both the simulation itself, as well as for verification and validation purposes. Satellite position and attitude are needed in the coordinate mapping process. Pixel level geo-location data is needed to validate the coordinate mapping process. Preferably additional meta-data like water and cloud masks, and terrain type information is available to further investigate their effect on the registration algorithm performance.
- **Availability of data:** Data produced by the instrument must be freely available. Preferable a large number of images taken at different locations during different days of year are available, so the effects of surface types, weather, seasonal effects, etc. can be tested.

Selected source instrument

The selected source instrument is the Moderate Resolution Imaging Spectroradiometer (MODIS) instrument carried on board NASA's Terra satellite. The exact data product used is MOD021KM level 1B. This is the freely distributed data product closest representing raw instrument data and has only undergone radiometric calibration. The specifications of the MODIS instrument are compared to the current SPECTROLITE specifications in Table 6.1. Both instruments have a significant overlap in spectral bands, allowing the use of a MODIS band corresponding to one of the SPECTROLITE bands. Depending on the MODIS band, three different spatial resolutions are available. For the 1 km resolution corresponding to SPECTROLITE, the SNR value is close to that of SPECTROLITE. The MTF an Nyquist is, however, significantly higher for MODIS.

Table 6.1: MODIS instrument parameters versus SPECTROLITE design requirements

Parameter	MODIS value	SPECTROLITE value
Spectral bands	36 spectral bands ranging from 400 nm to 14400 nm	270-775 nm
Spatial resolution	250 m (band 1–2) 500 m (band 3–7) 1 km (band 8–36)	1 km
Swath size	2330 km	900 km
SNR	750-850 for relevant 1km bands	800
MTF at Nyquist	0.57-0.58 for relevant 1km bands	0.25

The specifications of the Terra satellite are compared to preliminary SPECTROLITE satellite in Table 6.2. Notice that the satellites have very similar orbits. The descending node of the SPECTROLITE satellite is still unknown. If this descending node is close to that of the Terra satellite, lighting conditions experienced by both instruments over an orbit will be very similar. The descending node around 10:00 a.m. is commonly used, as it provides maximum illumination with minimum water vapour haze and cloud build-up [26].

The MODIS instrument is not of the push-broom type. Instead it uses a scanning mirror array, capturing a total of ten along-track lines at a time. This means that the outer image lines have a small off-nadir angle in along-track direction. However, since the number of along-track lines on the array are limited, the along-track viewing angles are still close to nadir.

Table 6.2: Terra orbit vs. SPECTROLITE preliminary orbit

Parameter	MODIS value	SPECTROLITE value
Orbit type	Circular sun-synchronous	Circular sun-synchronous
Altitude	705 km	600 km
Descending node	10:30 a.m	-

Another difference compared to a push-broom instrument is that the geometry of the incoming rays is different for a scanning mirror system in cross-track direction. Therefore, the coordinate mapping equations cannot be used directly. Instead a calibration is performed to determine the location of each cross-track pixel on a theoretical push-broom array. Using this theoretical array geometry instead of a push-broom geometry with constant pixel pitch, allows use of the described push-broom pixel coordinate mapping equations with the MODIS instrument data. Another effect of the scanning mirror geometry is a larger increase of pixel size with distance from image centre. We still use a total of 900 pixels for the SPECTROLITE simulation, however the 900 MODIS pixels correspond to a swath width of around 1000 km. We assume this increased pixel size at the outer edges of the imaging line has a small negative effect on the performance of the registration method, meaning it will not cause an overestimation of the measurement accuracy.

MODIS geo-location files are available, providing individual pixel coordinates in latitude/longitude format, as well as satellite position and attitude corresponding to each image line. The pixel coordinates provided have an RMS error of 50 meters, sufficient for validation purposes.

6.1.2. MTF adjustment

As we previously discussed, the imagery MTF level has an effect on the registration performance. The MTF of the 1 km GSD MODIS imagery is significantly higher than the MTF of SPECTROLITE. To better represent the SPECTROLITE performance, the MTF level of the MODIS imagery is artificially lowered.

As already discussed in Section 3.3, the total MTF is a convolution of its individual contributions. To lower the MTF we can apply a filter with an MTF contribution. We do not know the full MTF as a function of spatial frequency for MODIS nor SPECTROLITE, but only their values at Nyquist frequency. Furthermore, the difference in MTF between MODIS and SPECTROLITE can have a number of contributions. Instrument optics, detector sampling and platform stability, each have a different MTF contribution at a different spatial frequency. Since we cannot model the simulation imagery MTF for the full range of spatial frequencies, we choose to convolute the MODIS imagery with a Gaussian smoothing filter to lower the MTF at Nyquist to the SPECTROLITE level. The two-dimensional Gaussian has the form

$$G(x,y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \tag{6.1}$$

To lower the MTF from the MODIS value of 0.57 to the SPECTROLITE value of 0.25, we need a filter with an MTF of around 0.44 at the Nyquist frequency of the image. Since an image has discrete pixel sizes, we need a discrete approximation of the Gaussian to convolute with the image. However, the response of a Gaussian is effectively zero more than three standard deviations away from its centre, meaning the filter will only have a non-zero value at its centre pixel. Therefore, the image is first resampled with an increased number of samples before applying the filter and resampling it to the original sample size. To ensure the MTF response of the applied filter has the correct value, unit testing of the filter code was performed and is described in Section 7.1.

Validity of MTF adjustment method

The MTF adjustment method used will not produce imagery perfectly representative for the imagery

produced by SPECTROLITE. The Gaussian shape of the filter contribution is not perfectly representative for all contributions to the instrument MTF. Futhermore, the discrete pixel size of the imagery means that especially high spatial frequency image patterns cannot be perfectly reconstructed.

However, it is believed that the smoothing effect, responsible for the loss in registration accuracy, is still accurately represented by the MTF adjustment method. This smoothing effect causes a drop in the contrast between neighbouring pixels and has a large effect on the performance of the correlation methods. It is therefore believe that the image registration results achieved using the MTF adjusted imagery will be representative for the results achieved using SPECTROLITE imagery. However, to ensure the achievable registration accuracy for SPECTROLITE imagery, additional simulations using other source instruments can be performed. By checking whether the results of these simulations are in line with simulation results for MODIS imagery with an MTF level adjusted to this secondary source instrument imagery, validity of the MTF adjustment method can be assured.

6.2. Reference data

Apart from the instrument data, the registration simulation requires reference data to perform the registrations. The reference data consists of reference imagery and elevation data provided by a digital elevation model.

6.2.1. Reference imagery source

To select a suitable reference imagery source, a number of requirements for the reference data are defined. Based on these requirements a the reference image source is selected.

Reference imagery requirements
The following requirements are set for the reference imagery source.

- **Ground sampling distance:** Ground sampling distance must be sufficiently below the 1 km value, to investigate methods of sub-pixel registration with small pixel pitch values, as discussed in Section 5.3.3
- **Spectral band:** Since surface reflectance values can vary for different wavelengths, the reference imagery spectral band used should overlap or be close to the SPECTROLITE spectral band.
- **Geo-location meta-data:** The performance of the registration method is directly influenced by the quality of the reference image geo-location information. We therefore want geo-location imagery with an accuracy far below 100 m.
- **Cloud coverage:** To provide image registration capability for all captured image pixels, cloud free reference imagery is needed.
- **Availability of data:** Reference imagery with global coverage must be freely available, to allow a functioning image navigation system at all times.

Selected reference imagery source
The selected reference image source is the Landsat 8 satellite, the latest in the Landsat series of satellites. Landsat imagery is orthorectified and resampled to a 30 meter GSD. Landsat imagery uses the Worldwide Reference System (WRS) to catalogue imagery in fixed tiles with specified boundaries. The WRS catalog allows us to select a reference imagery based on the coordinates of the sensed image by means of a look-up table. Geo-location information for Landsat images uses the Universal Transverse Mercator (UTM) projection. UTM uses a two-dimensional cartesian coordinate system, similar to a standard Mercator projection. However, UTM does not use a single projection for the entire globe. Instead, the globe is divided into sixty UTM zones, with a specific secant transverse Mercator projection used in each zone. Use of the UTM projection requires a change of reference frame from the latitude/longitude projection used in the coordinate mapping equations and the MODIS geo-location data.

Based on the specifications of SPECTROLITE, MODIS and Landsat 8, three overlapping bands have been selected for use in the image registration simulation. This allowed investigation of the effect of wavelength on the registration accuracy, especially the difference between the UVIS and NIR bands. However, no clear differences in registration accuracy were found between the three wavelength bands.

Table 6.3: Overlapping SPECTROLITE, MODIS and Landsat 8 bands.

SPECTROLITE	MODIS	Landsat 8
UVIS (320-500 nm)	Band 9 (438-448 nm)	Band 1 (430-450 nm)
UVIS (320-500 nm)	Band 10 (483-493 nm)	Band 2 (450-510 nm)
NIR (675-775 nm)	Band 14 (673-683 nm)	Band 4 (640-670 nm)

Pre-processing
The database of accumulated Landsat imagery is extensive, with imagery of several years still available. Using the available imagery, a global database of cloud free images was created. These images are processed and resampled to decrease the reference image database size. The sampling distance used is 500 m, as decreasing this value further showed no improvement in registration accuracy.

6.2.2. Elevation data source

The elevation data needed for the image line coordinate mapping is provided by NASA's Shuttle Radar Topography Mission (SRTM). This data was captured by a dedicated radar instrument carried on board NASA's shuttle. The exact Digital Elevation Model (DEM) version used is SRTM3, which represents the finished data product after additional processing by NASA. This DEM has a resolution of 90 m, leading to a database with acceptable size for use on board the SPECTROLITE satellite. The processing performed to create SRTM3 fills voids in the original measurements. This makes it more reliable compared to some other freely available DEMs like the ASTER GDEM [18].

6.3. Set-up description

A step by step explanation of the simulation steps by means of a schematic view of the simulation is provided in Fig. 6.2. An explanation of the meaning of the blocks used in this diagram is available in Fig. 6.1. Note that the simulation and all its code is not explained in detail in this document. Given the size and complexity of the simulation created, a separate simulation manual was written. The goal of this manual is to aid in the continuation of the image navigation research by other authors.

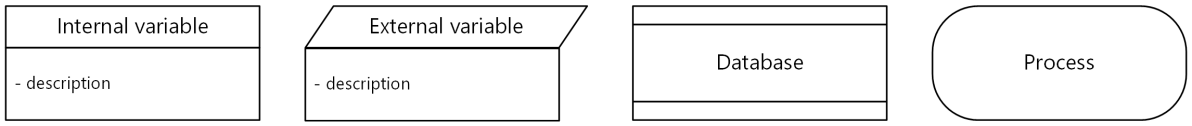

Figure 6.1: Registration simulation architecture legend.

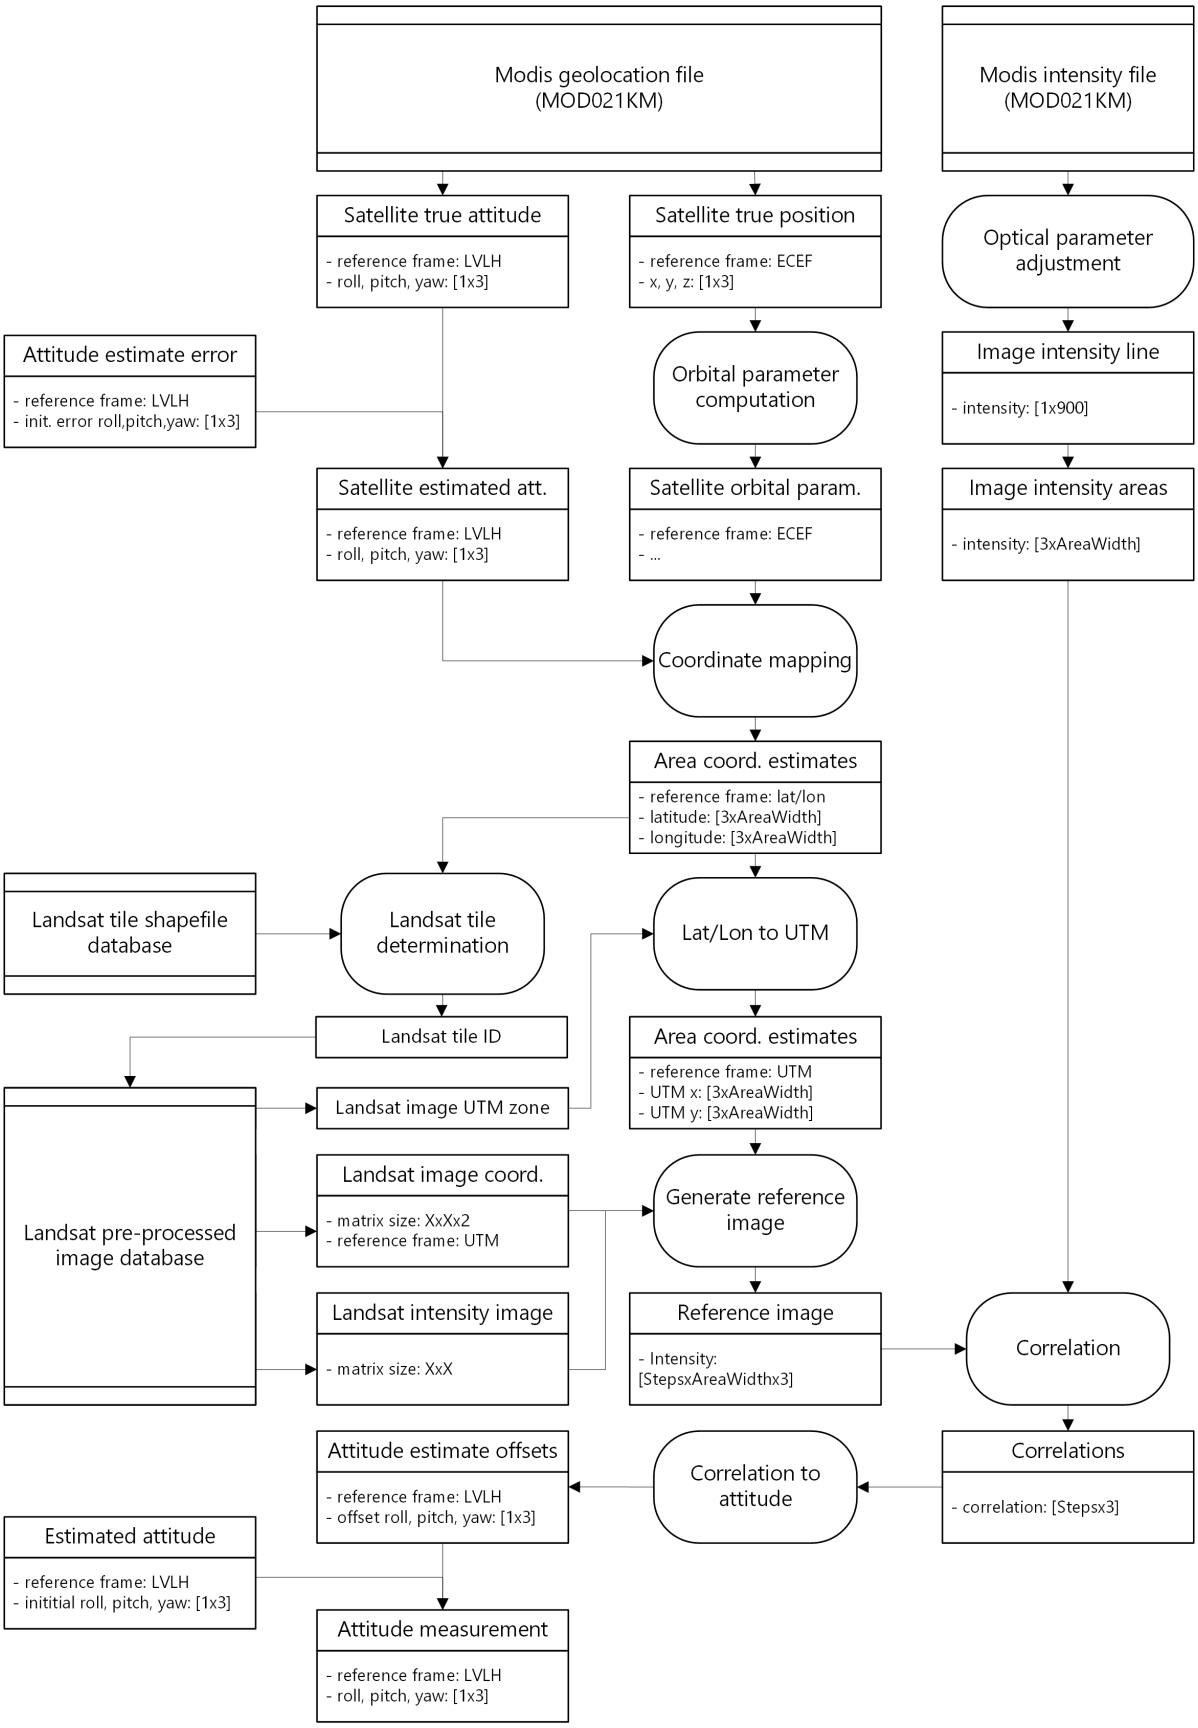

Figure 6.2: Registration simulation architecture

7

Image registration algorithm verification and performance comparison

The image registration simulation described in the previous chapter is used to check whether the algorithms described in Chapter 5 perform as expected. As a first step, unit testing of key simulation components is performed. Next, the performances of the four algorithms described are compared and a final correlation method is selected. Additional simulations results are analysed to get a better understanding of the selected algorithm's response to verify its performance and allow tuning to increase its accuracy.

7.1. Unit testing

Unit testing has been performed for the entire simulation code. The results for the coordinate mapping and MTF adjustment code will be presented in this section.

7.1.1. Coordinate mapping

Coordinate mapping of the captured image line is performed using the equations described in Section 5.2.2. This process corresponds to the "coordinate mapping" block in Fig. 6.2.

Accurate coordinate mapping is key in achieving high accuracy image registration. We therefore validate the coordinate mapping equations using the geo-location information provided for the MODIS imagery. Furthermore, we check how quickly the solution of these mapping equations converge and what the error induced by the simplified coordinate mapping method is.

Coordinate mapping equation accuracy
If we allow the coordinate mapping process to fully converge, the RMS error per pixel reduces to within the accuracy of the geo-location information provided by the MODIS team, as can be seen in Fig. 7.1.

We do, however, notice a sinusoidal distribution of the RMS error over the image line. It is unclear what the source of this error is, but it could be caused by an error source in the production of the geo-location information, the digital elevation model used, the reference geoid used, or perhaps the coordinate system transformations between ECEF frame and latitude/longitude. However, since the error lies within the accuracy of the geo-location information, its source is not further investigated. If, in future research, an effort is made to improve the image registration accuracy beyond the levels achieved in the course of this project, higher accuracy coordinate mapping might be needed. In such a case, it might be of interest to investigate the cause of the discrepancy between the image coordinates computed and those supplied by the geo-location files.

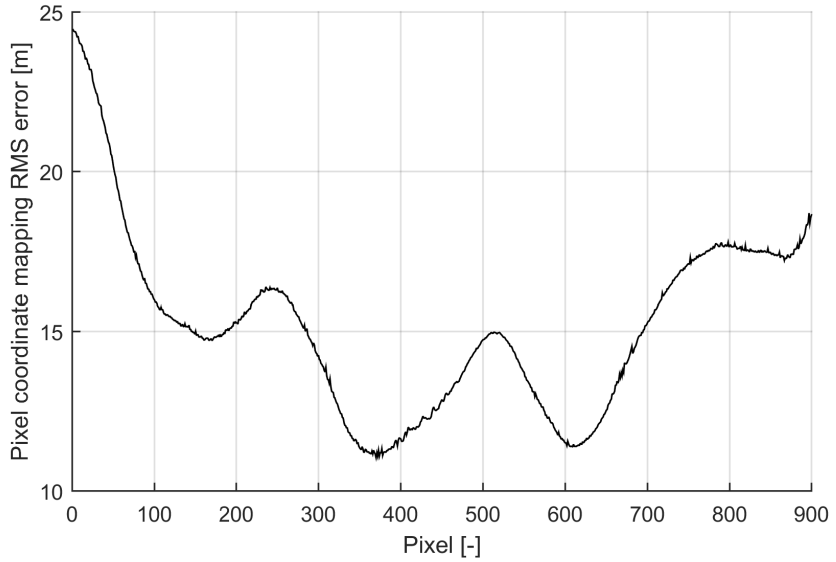

Figure 7.1: Per pixel coordinate mapping RMS error over 200 image lines after maximum convergence

Speed of convergence
In Fig. 7.2, the RMS error of the coordinate mapping process is plotted for a number of iterations. By using the elevation of the neighbouring pixel as first estimation of a pixels elevation, the location estimate converges quickly. After three iterations, the RMS errors have converged to within 2% of the error values after full convergence. Since further convergence was found to have no effect on the image registration accuracy, the number of iterations will be limited to three.

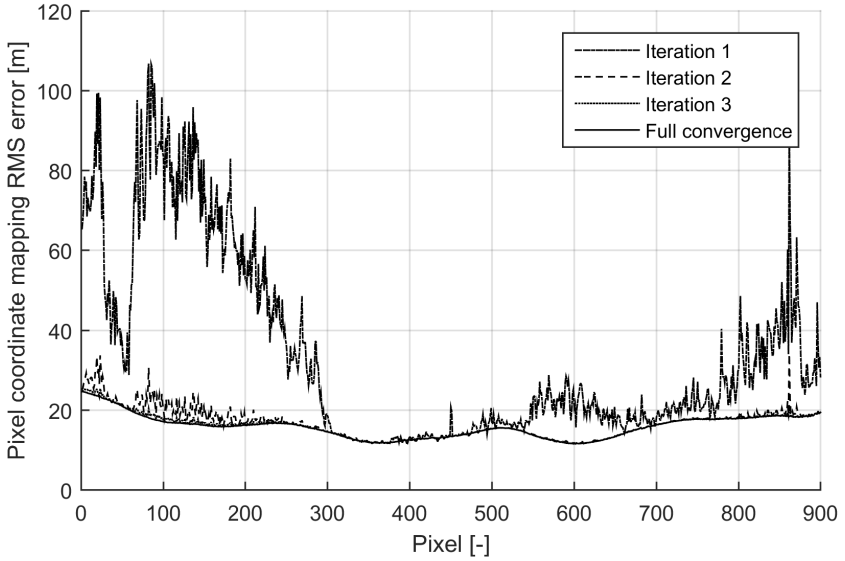

Figure 7.2: Per pixel coordinate mapping RMS error after iterating

7.1.2. MTF adjustment

MTF measurements can be performed relatively simple if the right imagery is available. To measure instrument MTF in a static set-up, target imagery with a clear edge is used. A commonly used method is the slanted edge method [20]. An image is captured of a slanted edge target, like the one in Fig. 7.3. At a significant distance from the edge, intensity values should either be very high (the white side) or

very low (the black side). However, in the edge region, instrument imperfections will cause a gradient from low to high intensity values. Based on this gradient the instrument MTF can be computed.

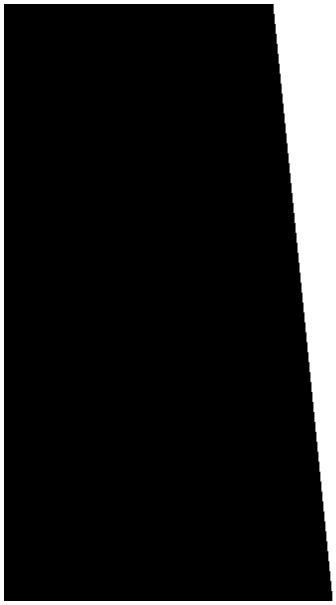

Figure 7.3: Example slanted edge image.

This method can also be used to measure the MTF of instruments in orbit . However, the slanted edge test target used in static testing might not be available in orbit. Instead, other landmarks with a clearly defined edge must be found. To determine the MODIS instrument in-orbit MTF, Lunar observations have been used [51]. These images provide the clear edge of the Moon on a deep space background. For in-orbit MTF determination of Landsat imagery, long straight landmarks like bridges have been used [41].

Still, the imperfections in the test targets make accurate MTF measurements for satellite imagery significantly more challenging. Therefore, for the purpose of this research, validation of the MODIS imagery MTF after application of the Gaussian filter will not be performed. Instead, we will validate the MTF contribution of the filter using the slanted edge method and assume its contribution will be the same when applied to the MODIS imagery.

To perform the validation, a slanted edge image is created with an MTF at Nyquist equal to that of the MODIS instrument. Measuring of the slanted edge image MTF is performed using the ImageJ software, an open platform for scientific image analysis [1]. Next, the filter is applied to the slanted edge image, after which its MTF is once again determined. This process is repeated while tuning the Gaussian filter to achieve the required MTF adjustment. The MTF measurements of the original and adjusted slanted edge images, together with the MTF contribution of the Gaussian filter, are available in Fig. 7.4. In this figure, we see the MTF at Nyquist (a normalised spatial frequency of 0.5) of the filter and adjusted image are equal to 0.44 and 0.25 respectively.

Although this method proves the validity of the MTF adjustment to some extend, it is recommended that additional validation on the MTF adjusted MODIS imagery is performed.

7.2. Algorithm performance comparison

To analyse and compare the performance of the four selected correlation algorithms, we look at two distinct performance characteristics determining the measurement accuracy: trueness and precision. Trueness refers to the systematic error or bias present in a series of measurements, while precision refers to random errors or the variability of a series of measurements [19]. Trueness is assessed by the ability of the algorithms to retrieve the error in the initial estimation of the attitude. This ability is

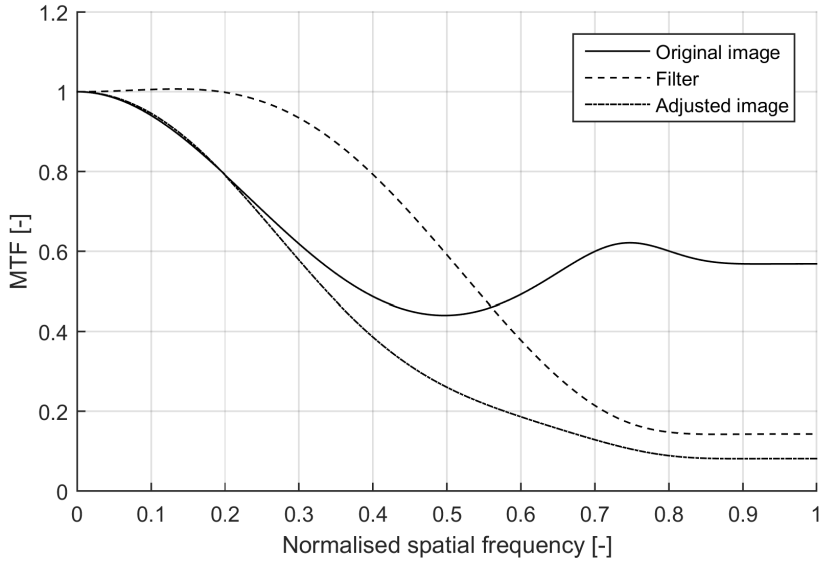

Figure 7.4: MTF measurements of original and MTF adjusted slanted edge image.

key in achieving bias free measurements and thus achieving a good trueness. Precision is assessed by looking at the distribution of the correlation values for a registration.

7.2.1. Registration trueness

The registration method designed uses an initial estimate of the satellite's attitude. A correlation window is set-up around this initial estimate. Next, the location with highest correlation within this correlation window is used to compute and offset corresponding to the error in the initial estimate. This offset is applied to the initial attitude estimate to create an absolute attitude measurement.

To assess the trueness of the measurements provided by the four algorithms, a constant initial error value is fed into the simulation. This initial error is the difference between true attitudes corresponding to an image line state and the attitude angles used as an initial guess in the registration process. To distinguish between systematic and random errors, the mean of the offsets computed by the registration algorithm is taken. This value should correspond to the initial error value fed into the simulation. If the mean offset computed is not equal to the constant initial error applied, the registration suffers from a bias.

Fig. 7.5 show the mean estimated initial error as a function of the true initial error for the roll axis. Note that initial errors for the other two attitude angles are kept at zero. We see good trueness for the incremental sign distance, Kendall's tau and Spearman's rho methods, with differences between the true and estimated initial errors not exceeding 0.002°. The trueness for the Pearson coefficient algorithm is less good, with differences between true and estimated initial error exceeding 0.01°.

If we look at the initial error estimation results for the pitch angle in Fig. 7.6, we notice the difference between algorithms increases. We see the Pearson coefficient method shows particularly poor trueness. A possible explanation for this poor trueness is the coupling between the pitch and yaw measurements. Both measurements are based on the same two along-track correlations.

The trueness of the Kendall's tau and Spearman's rho methods is still good, but becomes poorer at higher applied initial error. The incremental sign change method shows good trueness for the entire range of applied initial errors.

Fig. 7.7 shows the trueness for the yaw angle is significantly poorer than for the roll and pitch angles. Once again, the Pearson coefficient method shows the poorest trueness. Both Kendall's tau and Spearman's rho show reasonable trueness, with the difference between true and estimated initial error

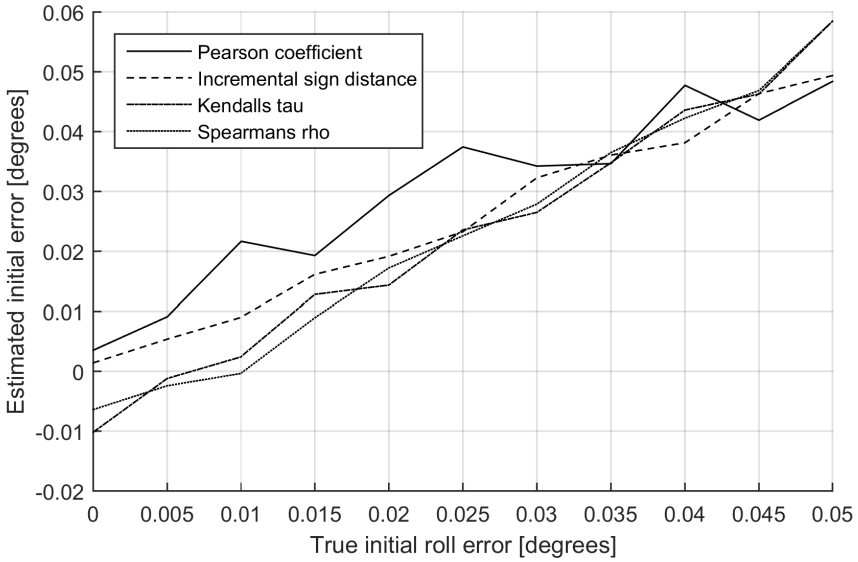

Figure 7.5: Roll angle initial error estimation results.

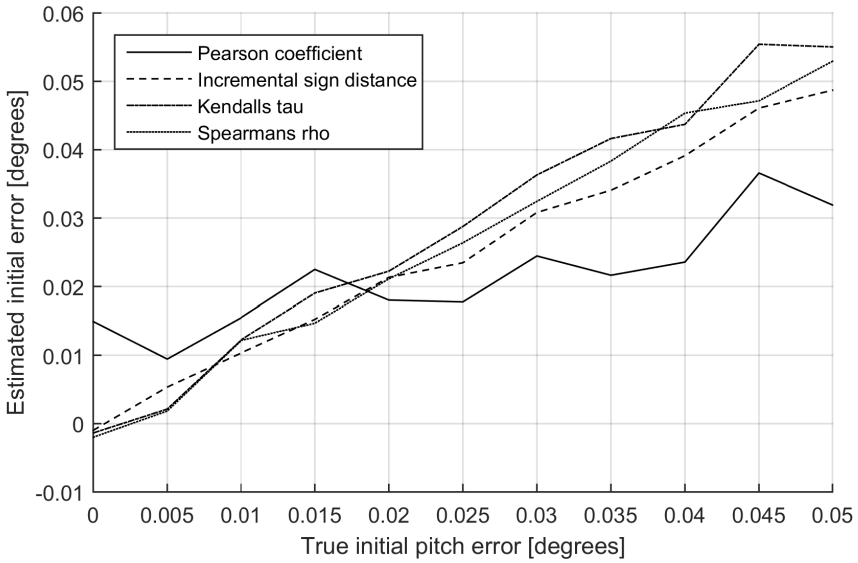

Figure 7.6: Pitch angle initial error estimation results.

at times exceeding 0.01°. The incremental sign change methods shows the best trueness, with the difference between true and estimated initial error not exceeding 0.006°.

For all three measured attitude angles, the Pearson coefficient shows the poorest trueness. The Kendall's tau and Spearman's rho methods show similarly good trueness for all three angles. However, they are outperformed by the incremental sign change methods, which shows the best trueness for all three angles.

7.2.2. Registration precision

To analyse and assess the registration precision for the four algorithms, we look at both the individual registrations of the three correlation areas, as well as the measurements for the three attitude angles computed based on these registrations. We feed the simulation with a random distribution of initial

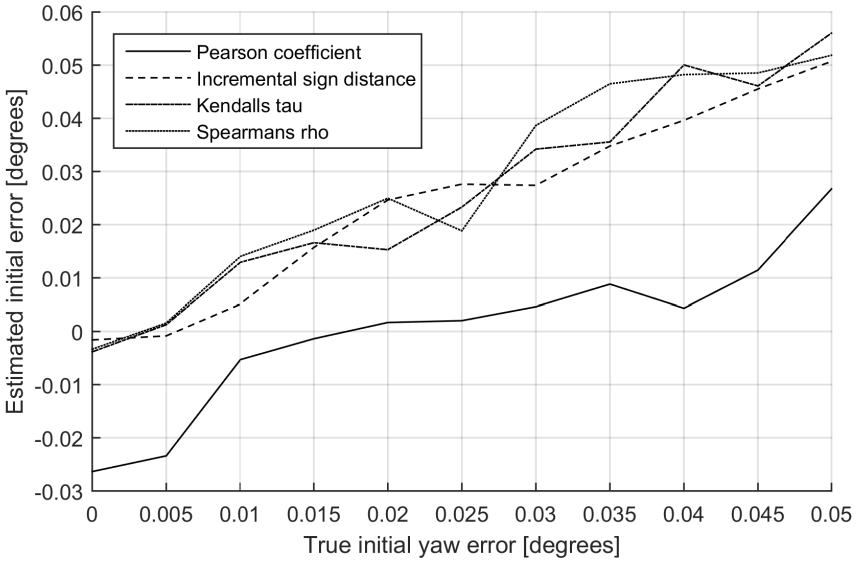

Figure 7.7: Yaw angle initial error estimation results.

errors between -0.01° and 0.01°.

Fig. 7.8 and Fig. 7.9 show the mean normalised correlation distributions for the along-track and cross-track registration, respectively. We see that, for both registration directions, the mean normalised correlation values form a Gaussian distribution. The width of this bell curve shape is representative of the precision of the registration. A small width corresponds to a high accuracy. We see the incremental sign change methods displays the highest precision for both along- and cross-track registration. Kendall's tau and Spearman's rho show similar precision, with former showing a slightly outperforming the latter. The Pearson coefficient method shows the worst precision.

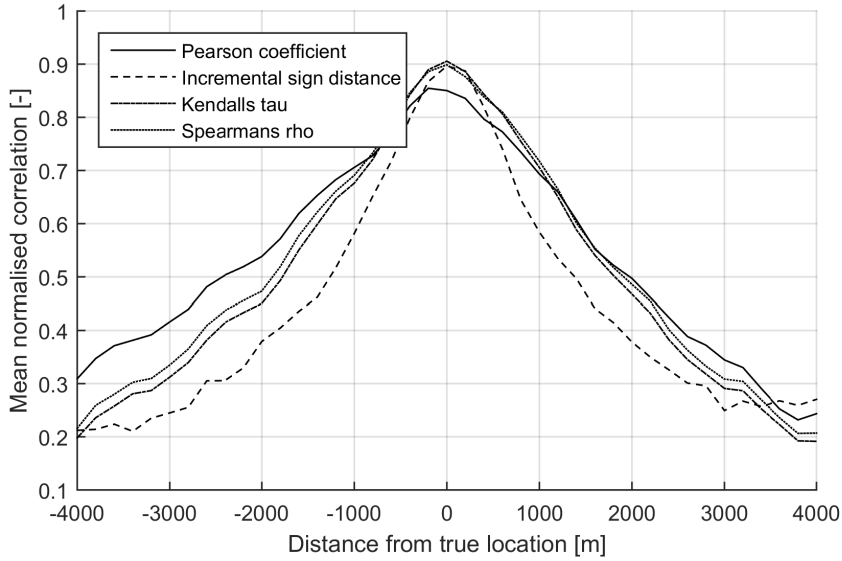

Figure 7.8: Along-track mean normalised correlation distribution for considered correlation methods.

Looking at the standard deviations of the attitude angle measurements, tabulated in Table 7.1, we see these results are in line with the correlation distributions found in Fig. 7.8 and Fig. 7.9. The Pearson coefficient method shows the worst precision, followed by the Spearman's rho method. The latter is outperformed slightly by the Kendall's tau algorith. The incremental sign change method shows by far

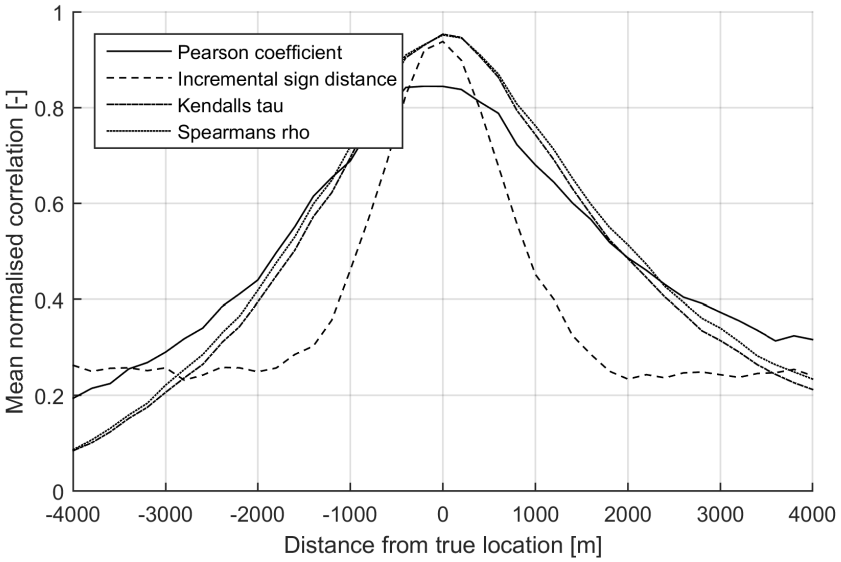

Figure 7.9: Cross-track mean normalised correlation distribution for considered correlation methods.

the highest precision.

Table 7.1: Attitude measurement accuracy for considered correlation algorithms.

	Roll measurement accuracy (1σ) [°]	**Pitch measurement accuracy (1σ) [°]**	**Yaw measurement accuracy (1σ) [°]**
Pearson coefficient	0.054	0.056	0.11
Incremental sign distance	0.018	0.018	0.043
Kendall's tau	0.020	0.021	0.056
Spearman's rho	0.025	0.028	0.068

Roll measurements perform consistently better than pitch and yaw measurements. However, in Fig. 7.8 and Fig. 7.9, we also notice that for a given correlation algorithm, the correlation distribution for the along-track registrations is wider than for the cross-track registrations. This effect is investigated in more detail in Section 7.3.1.

7.2.3. Selected algorithm

The incremental sign change algorithm outperforms the other three tested algorithms in all aspects. Analysis shows the algorithm achieves both the best trueness and highest precision. Furthermore, the simulation shows the incremental sign change method requires the least computation time. Therefore, the this algorithm is selected for use in the SPECTROLITE image navigation system.

7.3. Correlation algorithm verification

Following the selection of the incremental sign change method as the final correlation algorithm, further verification steps are performed.

7.3.1. Along and cross-track precision variation

As we have already touched upon in Section 7.2.2, roll angle measurements show a higher accuracy compared to pitch and yaw angle measurements. For the yaw angle measurements this is expected,

since the sum of two registrations at both edges of the image line is used to create its measurement. Therefore, errors in both registrations contribute to the total measurement error of the yaw axis.

However, this effect cannot explain the reduced accuracy of the pitch angle measurements. In fact, since a pitch angle measurement is created by taking the mean of two independent registrations, we expect its accuracy to exceed that of the roll angle measurements. Therefore, the precision of the two individual registrations used to create the pitch angle measurement must be lower compared to the registration used to create the roll angle measurement. This hypothesis is supported by the mean normalised correlation value distribution for the along and cross-track registrations, as shown in Fig. 7.10. The cross-track registration, used to create the roll angle measurement, shows a significantly narrower bell curve compared to the along-track registration used to create the pitch angle measurement. This result is unexpected, since the MODIS MTF values in along-track direction are slightly higher than in cross-track direction.

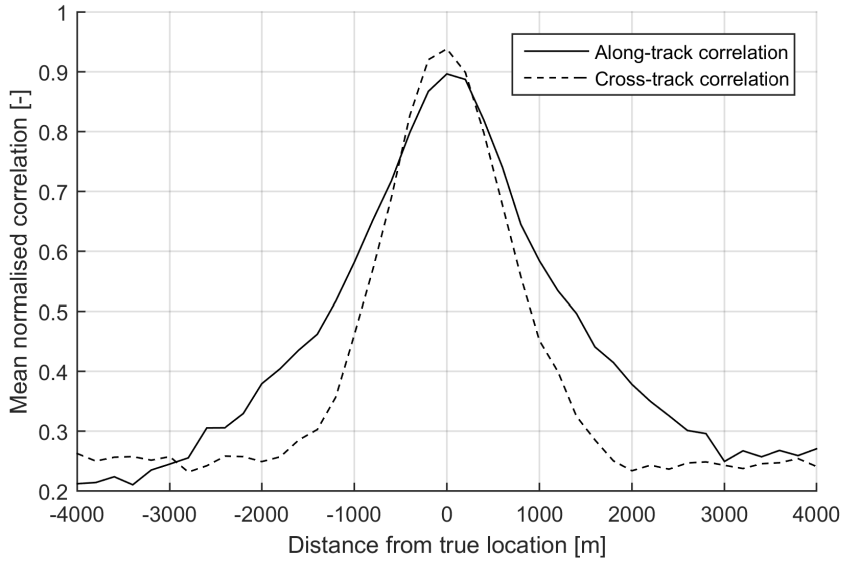

Figure 7.10: Mean normalised correlation values for along and cross-track registration.

An possible explanation for this phenomenon lies in the image line regions used to perform the individual registrations. To achieve the highest possible yaw angle measurement accuracy, the along-track registrations used to measure the pitch and yaw angles are performed on image regions at the edges of the image line. The cross-track registration used to measure the roll angle is performed on an image region in the centre of the image line. As we discussed in Section 6.1.1, the scanning mirror nature of the MODIS instrument causes a decrease in pixel resolution when moving from the centre pixel to the edges of the image line. This means the registrations used to measure the pitch and yaw angles are performed uses lower resolution imagery.

To further investigate this phenomenon, we perform along and cross-track registrations at both the centre and edge image line regions. Fig. 7.11 and Fig. 7.12 show the mean normalised correlation values at an edge and centre region for along and cross-track registration, respectively. We see that the registration precision is consistently higher for a registration at the centre region. We therefore conclude that, indeed, the lower pixel resolution at the edges of the image line causes a decrease in the registration precision. This will in turn cause a decrease in the accuracy of pitch and yaw measurements. The decreasing resolution effect is less severe for a pushbroom imager, like SPECTROLITE, than it is for a scanning mirror instrument. Therefore, the accuracy of pitch and yaw measurements based on the registration of SPECTROLITE imagery could be higher than shown by this simulation.

However, the resolution is not the only source of the precision difference between along and cross-track registration. When comparing Fig. 7.11 and Fig. 7.12, we see that even when the same image region is registered, along-track registrations have a lower precision compared to cross-track registrations. An obvious explanation for the reduced precision in along-track direction might be that this is

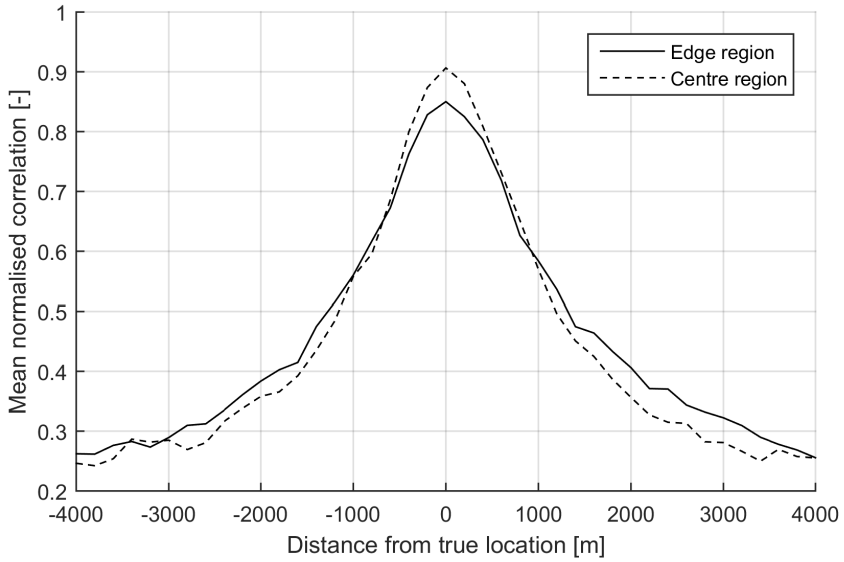

Figure 7.11: Mean normalised correlation values for along-track registration at edge and centre locations.

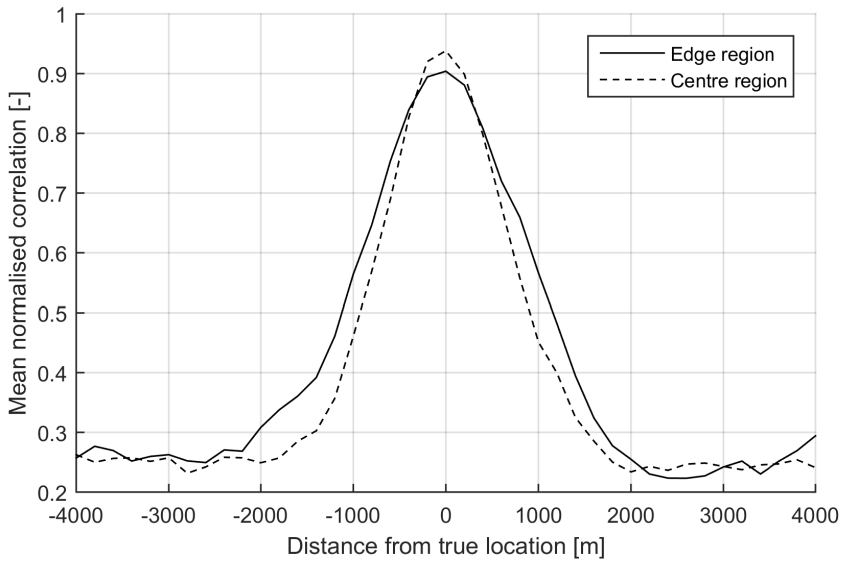

Figure 7.12: Mean normalised correlation values for cross-track registration at edge and centre locations.

also the direction of flight of the satellite. The motion in along-track direction might cause additional image blurring worsening the registration performance in this direction. However, this explanation is not necessarily supported by in-orbit MTF measurements performed for SPECTROLITE, which show similar MTF values in both along and cross-track direction [53].

Another possible explanation for the reduced precision of along-track registrations compared to cross-track registrations can be found in the one-dimensionality of the registered image. Although the along-track registration is performed by correlating the reference and sensed images at various locations in along-track direction, the correlations are based on the intensity variations between pixels in cross-track direction. If this hypothesis is true, using an image line in along-track direction could increase the along-track registration accuracy. This is not possible for a pushbroom imager like SPECTROLITE. However, it could allow an increase in pitch and yaw measurement accuracy to be achieved when a two-dimensional imager is used. Still, further investigation is required to prove this hypothesis and determine possible benefits of a two-dimensional imager.

7.3.2. Cross-coupling

In Section 7.2.1 we assess the trueness of the registration algorithms by their ability to estimate the error in the initial attitude estimate. However, we have only investigated the estimation of an initial error for a single attitude angle, with the other two angles having zero initial error. In practical situations, an initial attitude estimation error for all three axis will be present. We will therefore also analyse how an initial error in the attitude of one angle, influences the estimation of the other two angles.

For the roll angle, the results are visualised in Fig. 7.13. We see the mean of the estimated initial error is indeed close to the true initial error applied. Furthermore, we see that the estimation of the other two angles remains accurate until a 0.03 degree initial roll error, when the coupling between pitch and yaw angle results in erroneous measurements.

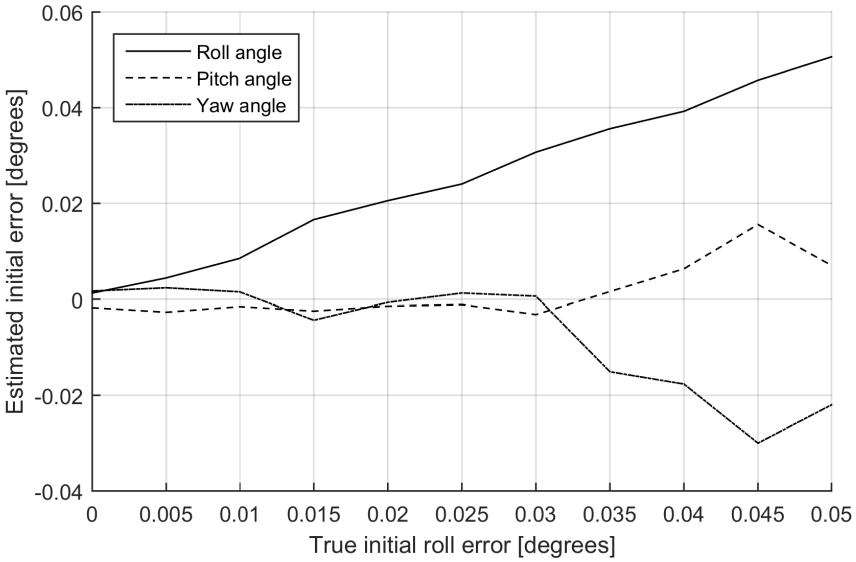

Figure 7.13: Estimation of initial roll angle error.

We do the same for the pitch angle as shown in Fig. 7.14, and for the yaw angle as shown in Fig. 7.15. For these angles, we see no

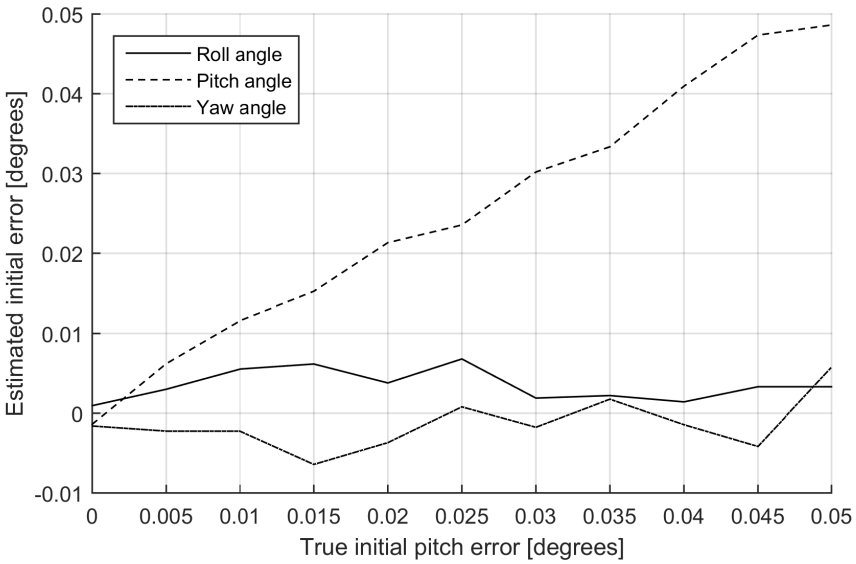

Figure 7.14: Estimation of initial pitch angle error.

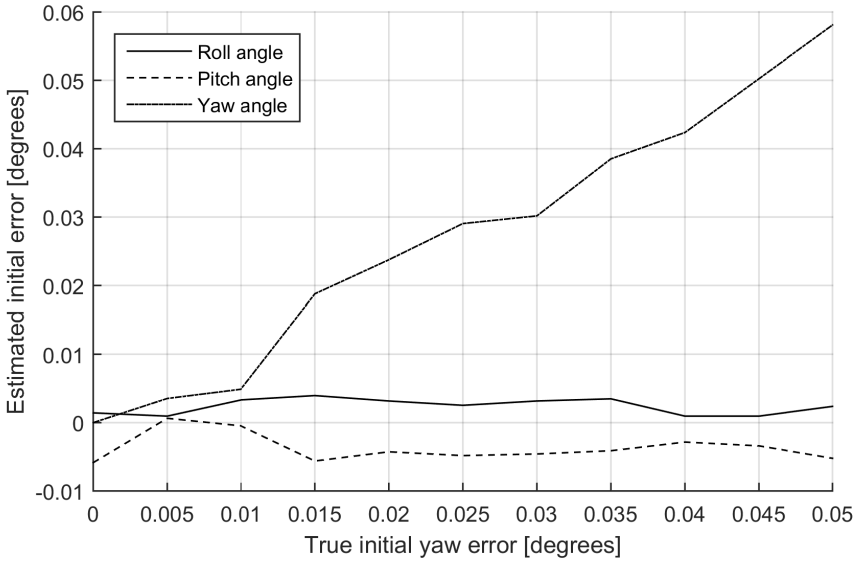

Figure 7.15: Estimation of initial yaw angle error.

As expected, we see the one-dimensionallity of the designed registration method causes some cross-coupling between measurement accuracies of the three attitude angles. An initial error for the roll angle causes a decrease in the measurement accuracy of the other angles. However, at small initial errors of under 0.03 degrees, the effect is minimal. Since the attitude determination requirement prescribes a one sigma accuracy of 0.01 degree, the cross-coupling effect will be minimum once this accurate attitude knowledge has been achieved. However, from these results it is unclear whether the one-dimensional registration method is capable of converging the attitude knowledge from a larger initial error at the start of an observation run. This fine attitude acquisition phase will be investigate in more detail using the full image navigation simulation in Section 9.3

7.3.3. Correlation window size

As we shortly touched upon in Section 5.3, the correlation window is the range of distances from the initial estimate, for which the sensed and reference image are correlated. Clearly, the true position of the sensed image with respect to the reference image must lie within this correlation window for registration to occur. Therefore, the correlation window must be larger than the maximum expected error in the initial estimate.

However, even if the correlation window is larger than the maximum error in the initial estimate, its size can have an effect on the registration trueness. If we look at the initial error estimation results for a range of correlation window sizes, we notice decreasing the window size below a certain value starts to worsen the attitude measurement trueness. For the roll angle measurements, trueness worsens at a correlation window size below 0.20°, as can be seen in Fig. 7.16.

Similar results are observed for the pitch angle measurements in Fig. 7.17, although here the drop in trueness already occurs for a larger correlation window size.

We believe the drop in trueness can be attributed to clipping of the Gaussian correlation distributions found in Section 7.3.1. An error in the initial estimate for a registration will cause the correlation window to not be placed around the centre of the correlation distribution. This means part of the bell shape at one end of the distribution will fall outside the correlation window bounds. This moves the distribution average towards the centre of the correlation window, causing the drop in the estimated initial error visible in Fig. 7.16 and Fig. 7.17.

This loss in measurement trueness, can in turn cause a loss in the measurement accuracy achieved using the registration method. It means the correlation window size should not simply be determined

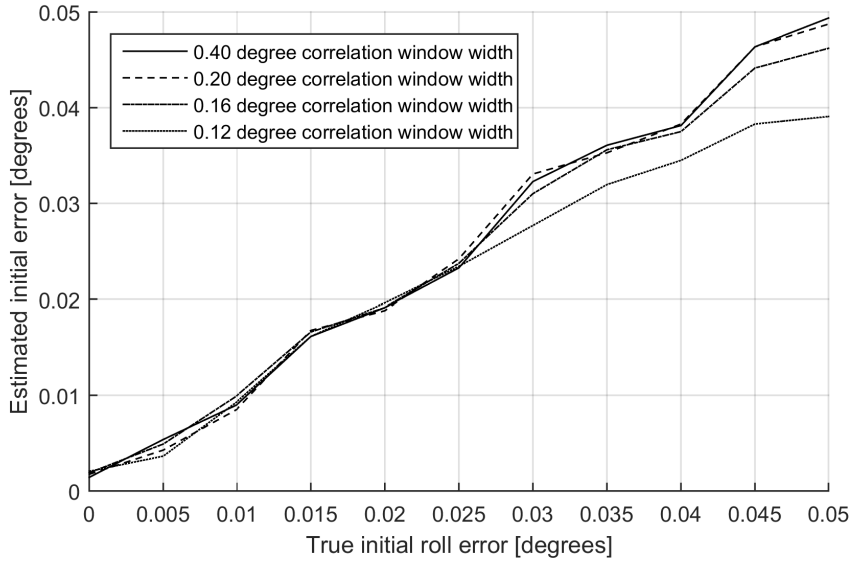

Figure 7.16: Effect of correlation window size on roll axis initial error estimation.

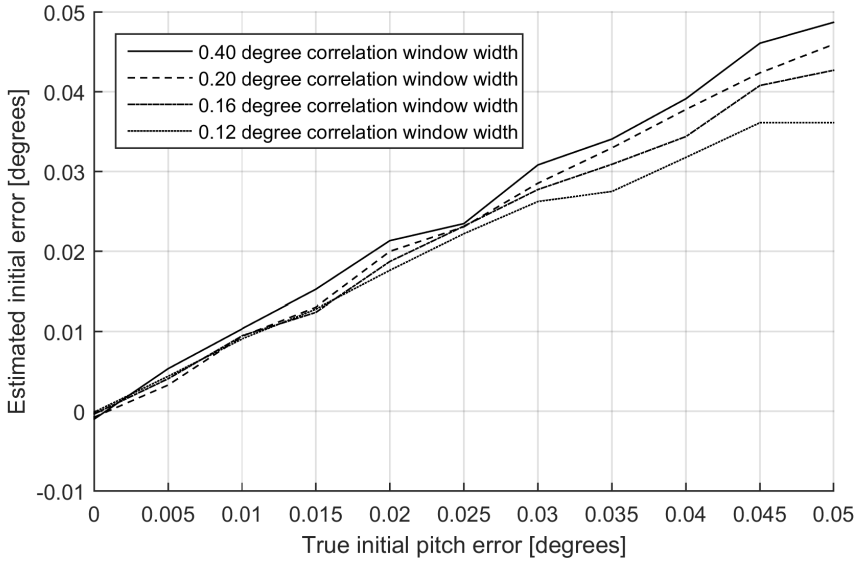

Figure 7.17: Effect of correlation window size on pitch axis initial error estimation.

based on the largest error expected in the initial error. To achieve accurate results, the width of the correlation distribution should be added to the correlation window size.

The effect of the correlation window size on the measurement accuracy will be investigated further using the image navigation simulation in Section 9.2.1.

8

Image navigation mode simulation set-up

The image registration simulation is extended with a simulation of the satellite's attitude. This allows an investigation of the image registration process in-orbit, as well as of the achievable real-time attitude determination accuracy after filtering. Furthermore, this simulation allows us to determine the time needed to achieve stable nadir pointing at the start of an observation run. Combining the image registration simulation with a simulation of the attitude dynamics, also allows a more detailed analysis of the rotational rates experienced by the satellite and the pointing stability it can achieve.

8.1. Attitude dynamics and control

In the modelling of the satellite's attitude dynamics we assume the satellite is equipped with three reaction wheels, with each wheel's axis of rotation aligned with one of the satellite's principle axes. Any additional torques induced by misalignment of the wheels and cross-coupling due to non-symmetric moments of inertia are modelled as internal disturbance torques, as described in Section 8.2.2.

8.1.1. Attitude dynamics

We will now derive the equations of motion used in the image navigation simulation, starting with the rotational equations of motion of a rigid body about its centre of mass

$$\overline{M} = \frac{\mathrm{d}\overline{H}}{\mathrm{d}t}|_B + \overline{\omega} \times \overline{H} \tag{8.1}$$

Note that $|_B$ indicates the satellite body frame. The moment vector $\overline{M}$ has multiple terms. We distinguish the control moments $\overline{M}_c$ and disturbance torques $\overline{M}_d$.

$$\overline{M} = \overline{M}_c + \overline{M}_d \tag{8.2}$$

Similarly the angular moment vector $\overline{H}$ has multiple terms. We distinguish the momentum of the satellite body $\overline{H}_B$ and the momentum of the reaction wheels

$$\overline{H} = \overline{H}_B + \overline{H}_W \tag{8.3}$$

Substitution of these vectors in equation Eq. (8.1) yields

$$\overline{M}_c + \overline{M}_d = \frac{\mathrm{d}(\overline{H}_B + \overline{H}_W)}{\mathrm{d}t}|_B + \overline{\omega} \times (\overline{H}_B + \overline{H}_W) \tag{8.4}$$

The angular momentum vectors of the satellite body is defined as

$$\overline{H}_B = I_B \overline{\omega} \tag{8.5}$$

The angular momentum vector of the reaction wheels is defined as

$$\overline{H}_W = I_W(\overline{\omega}_W + \overline{\omega}) \tag{8.6}$$

Generally, reaction wheels are spun up to create a significant momentum biased. This circumvents the additional jitter experienced by the wheels for angular rate values close to zero. We can therefore assume the rotational rate of the reaction wheels is always much larger than the rotational rate of the satellite. This allow us to approximate the angular momentum vector of the reaction wheels by

$$\overline{H}_W \approx I_W \overline{\omega}_W \tag{8.7}$$

The total angular momentum vector is now equal to

$$\overline{H}_B + \overline{H}_W = I_B \overline{\omega} + I_W(\overline{\omega}_W + \overline{\omega}) \approx I_B \overline{\omega} + I_W \overline{\omega}_W \tag{8.8}$$

Substitution in Eq. (8.4) yields

$$\begin{aligned} \overline{M}_c + \overline{M}_d = I_B \frac{\mathrm{d}\overline{\omega}}{\mathrm{d}t}|_B + I_W \frac{\mathrm{d}\overline{\omega}_W}{\mathrm{d}t}|_B + \overline{\omega} \times (I_B \overline{\omega} + I_W \overline{\omega}_W) \\ = I_B \frac{\mathrm{d}\overline{\omega}}{\mathrm{d}t}|_B + \overline{\omega} \times I_B \overline{\omega} + \frac{\mathrm{d}\overline{H}_W}{\mathrm{d}t}|_B + \overline{\omega} \times \overline{H}_W \end{aligned} \tag{8.9}$$

LVLH reference frame
We now assume a satellite in a circular orbit in the Local Vertical Local Horizontal (LVLH) reference frame define in Section 5.2.2. The attitude kinematics of the satellite in this reference frame in terms of rotational rates in the satellite's body frame can be defined using the reference frame transformation matrices

$$\begin{bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{bmatrix} = \frac{1}{cos(\psi)} \begin{bmatrix} c_\psi & -c_\phi s_\psi & s_\phi s_\psi \\ 0 & c_\phi & -s_\phi \\ 0 & s_\phi c_\psi & c_\phi c_\phi \end{bmatrix} \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{bmatrix} + \begin{bmatrix} 0 \\ n \\ 0 \end{bmatrix} \tag{8.10}$$

Note that ϕ, θ and ψ are the roll, pitch and yaw angles, and $\dot{\phi}$, $\dot{\theta}$ and $\dot{\psi}$ are the roll, pitch and yaw rates, both in the LVLH frame. ω_x indicates the rate in the satellite's body frame, where axes 1, 2 and 3 correspond to the roll, pitch and yaw principle axes. The n term is the satellite's mean motion, induced by the fact that the LVLH is not an inertial reference frame, but follows the rotation of the satellite around the Earth's centre of mass.

The torque provided by the reaction wheel motors does not result in a control moment $\overline{M}_c$, but rather in an increase in reaction wheel angular momentum over time. Therefore, the term $\overline{M}_c$ is zero. The dynamics of the satellite in the LVLH frame can now be defined as

$$\begin{bmatrix} I_{11} & I_{12} & I_{13} \\ I_{21} & I_{22} & I_{23} \\ I_{31} & I_{32} & I_{33} \end{bmatrix} \begin{bmatrix} \dot{\omega_1} \\ \dot{\omega_2} \\ \dot{\omega_3} \end{bmatrix} = - \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix} \begin{bmatrix} I_{11} & I_{12} & I_{13} \\ I_{21} & I_{22} & I_{23} \\ I_{31} & I_{32} & I_{33} \end{bmatrix} \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{bmatrix} + 3n^2 \begin{bmatrix} 0 & -C_{13} & C_{12} \\ C_{13} & 0 & -C_{11} \\ -C_{12} & C_{11} & 0 \end{bmatrix} \begin{bmatrix} I_{11} & I_{12} & I_{13} \\ I_{21} & I_{22} & I_{23} \\ I_{31} & I_{32} & I_{33} \end{bmatrix} \begin{bmatrix} C_{11} \\ C_{12} \\ C_{13} \end{bmatrix} + \begin{bmatrix} -\dot{H}_{W1} + M_{d1} \\ -\dot{H}_{W1} + M_{d2} \\ -\dot{H}_{W1} + M_{d3} \end{bmatrix} \tag{8.11}$$

Where we defined

$$C_{11} = -sin(\theta)cos(\psi) \tag{8.12}$$

$$C_{12} = cos(\phi)sin(\theta)sin(\psi) + sin(\phi)cos(\theta) \tag{8.13}$$

$$C_{13} = -sin(\phi)sin(\theta)sin(\psi) + cos(\phi)cos(\theta) \tag{8.14}$$

Linearisation

Since the satellite is nadir pointing, we assume deviations from nadir are small. As we will later discover, even during fine attitude acquisition, the maximum deviations from nadir will remain below 0.5°. This allows linearisation of the equations of motion around the nadir position. The attitude kinematics are now reduced to

$$\dot{\phi} - n\psi = \omega_1 \tag{8.15}$$

$$\dot{\theta} - n = \omega_2 \tag{8.16}$$

$$\dot{\psi} + n\phi = \omega_3 \tag{8.17}$$

The dynamics of the satellite are reduced to

$$\begin{bmatrix} I_{11} & I_{12} & I_{13} \\ I_{21} & I_{22} & I_{23} \\ I_{31} & I_{32} & I_{33} \end{bmatrix} \begin{bmatrix} \dot{\omega_1} \\ \dot{\omega_2} \\ \dot{\omega_3} \end{bmatrix} = n \begin{bmatrix} I_{31} & 2I_{32} & I_{33} - I_{22} \\ -I_{32} & 0 & I_{12} \\ I_{22} - I_{11} & -I_{12} & -I_{13} \end{bmatrix} \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{bmatrix} + 3n^2 \begin{bmatrix} I_{33} - I_{22} & I_{21} & 0 \\ I_{12} & I_{33} - I_{11} & 0 \\ -I_{13} & -I_{23} & 0 \end{bmatrix} \begin{bmatrix} \phi \\ \theta \\ \psi \end{bmatrix} + n^2 \begin{bmatrix} -2I_{23} \\ 3I_{13} \\ -I_{12} \end{bmatrix} + \begin{bmatrix} -\dot{H}_{W1} + M_{d1} \\ -\dot{H}_{W1} + M_{d2} \\ -\dot{H}_{W1} + M_{d3} \end{bmatrix} \tag{8.18}$$

Furthermore, since the satellite is a rectangular cuboid, we can assume the moments of inertia are symmetric. The dynamics of the satellite are now further reduced to

$$I_1\dot{\omega}_1 + n(I_2 - I_3)\omega_3 + 3n^2(I_2 - I_3)\phi = -\dot{H}_{W1} + M_{d1} \tag{8.19}$$

$$I_2\dot{\omega}_2 + 3n^2(I_1 - I_3)\theta = -\dot{H}_{W2} + M_{d2} \tag{8.20}$$

$$I_3\dot{\omega}_3 - n(I_2 - I_1)\omega_1 = -\dot{H}_{W3} + M_{d3} \tag{8.21}$$

If we combine this with the attitude kinematic equations we get a set of linearised equations of motion describing the satellite's attitude in the LVLH frame.

$$I_1\ddot{\phi} + 4n^2(I_2 - I_3)\phi - n(I_1 - I2 + I3)\dot{\psi} = -\dot{H}_{W1} + M_{d3} \tag{8.22}$$

$$I_2\ddot{\theta} + 3n^2(I_1 - I_3)\theta = -\dot{H}_{W2} + M_{d2} \tag{8.23}$$

$$I_3\ddot{\psi} + n^2(I_2 - I_1)\psi + n(I_1 - I2 + I3)\dot{\phi} = -\dot{H}_{W3} + M_{d3} \tag{8.24}$$

These differential equations are evaluated using the ODE45 solver in the MATLAB environment.

8.1.2. PD controller

The Proportional Different (PD) controller controls the reaction wheel torque based on measurements of both the attitude angle and attitude rate. For the pitch axis, the control law is formalised as

$$\dot{H_{W2}} = K_{p2}\Delta\theta + K_{d2}\Delta\dot{\theta} \tag{8.25}$$

This closed-loop control system contains the two gain factors K_p and K_d. The disturbance rejection method is used to determine the value of each of these controller gains. This method ensures a stable system, although the achieved pointing accuracy might not be optimal. However, for the purpose of this research, this method will suffice.

For the pitch angle, derivation of the control gains starts by writing Eq. (8.23) in the Laplace domain s

$$(I_{22}s^2 + K_{d2}s + K_{p2})\theta(s) = M_{d2} \tag{8.26}$$

We assume the disturbance torque is equal to its worst case value, as determine using the equations in Section 3.3.2

$$M_{d2}(s) = \frac{M_{d2,max}}{s} \tag{8.27}$$

The final value theory is used to compute the proportional gain K_p

$$\lim_{t\to\infty}\theta(t) = \lim_{s\to 0}\theta(s)s = \lim_{s\to 0}\frac{M_{d2,max}}{I_{22}s^2 + K_{d2}s + K_{p2}} = \frac{M_{d2,max}}{K_{p2}} \tag{8.28}$$

The we find the gain is equal to the worst case disturbance torque over the steady state pointing error θ_{ss}.

$$K_{p2} = \frac{M_{d2,max}}{\theta_{ss2}} \tag{8.29}$$

For this steady state error value we use the pointing requirement of 0.01°. The derived transfer function is now compared to a standard second order system

$$H(s) = \frac{\theta_2(s)}{M_{d2}(s)} = \frac{1}{I_{22}s^2 + K_{d2}s + K_{p2}} = \frac{\frac{1}{I_{22}}}{s^2 + \frac{K_{d2}}{I_{22}}s + \frac{K_{p2}}{I_{22}}} = \frac{\frac{1}{I_{22}}}{s^2 + 2\zeta\omega_n s + \omega_n} \tag{8.30}$$

If we compare the coefficients we find

$$K_{d2} = 2\zeta\omega_W I_{22} = 2\zeta\sqrt{K_{p2}I_{22}} \tag{8.31}$$

For the damping ratio ζ we use the optimal value of $\sqrt{2}/2$.

8.2. Modelling error sources and disturbances

A number of measurement and control error sources and disturbances must be modelled for the image navigation simulation to better represent reality. This includes modelling of the gyroscope, attitude controller, reaction wheels and external disturbance sources.

8.2.1. Modelling gyroscope measurements

As with all practical measurement devices, gyroscopes are prone to measurement noise. Additionally, gyroscopes suffer from measurement bias. This bias is equal to the output at zero rate and fluctuates over time. The angular rates measured by a gyroscope can be modelled by

$$\omega_{measured} = \omega_{true} + b + n \tag{8.32}$$

Where b represents the bias value and n is the measurement noise modelled as zero mean white noise. The standard deviation of the noise level can be considered constant and can be determined by testing the measurement device. The bias value, however, is not constant. Over longer periods of time, the bias value starts to drift. The bias drift is also not constant and the severity of the drift can be increased by external effects like vibrations and temperature variations. The bias can be estimated by integrating gyroscope measurements over a longer period of time. In case of a nadir pointing satellite, integrated gyroscope measurements should equal zero for the roll and yaw axes, and should be equal to the accumulated rotation in inertial space due to the mean motion for the pitch axis.

However, a sufficiently long integration time is needed to accurately determine the gyroscope bias. If external effects influence the bias during the integration period, error in bias estimation occurs. A satellite in sun-synchronous orbit is prone to significant temperature variations, as it moves in and out of eclipse during each orbit. We therefore assume the bias over temperature is the main source of error in the bias estimation. The error in bias estimation used in the image navigation simulation is therefore equal to the total bias change over a single temperature cycle, specified by manufacturers as the "bias over temperature".

An additional error source for the attitude rate measurements is the resolution of the measurement device. Especially when the magnitude of the measured rotational rates is low, this error can be significant. To account for the measurement resolution, the value found for Eq. (8.32) is rounded to the nearest resolution value.

As an initial reference gyroscope we use the SiREUS MEMS rate sensors jointly developed by SEA, Atlantic Inertial Systems, and SELEX GALILEO. This gyroscope was developed within the framework of an ESA research program and first flown on board the SENTINEL-3 satellite. The gyroscope is now available as COTS component. It has the highest performance found for any gyroscope with a mass under 1kg. Its performance will therefore serve as an upper bound in the initial analysis of the image navigation system. In Section 10.2 the gyroscope used in the final image navigation system design will be selected. The specifications of the SiREUS gyroscope are tabulated in Table 8.1.

Table 8.1: Performance specifications of the SiREUS gyroscope.

Specification	**Value**
Angular random walk (1σ)	0.04 $^\circ/\sqrt{hr}$
Bias over temperature (RMS)	10 $^\circ/hr$
Resolution	<1 arcsec/s

8.2.2. Modelling reaction wheel output

As discussed in Section 3.3.2, the reaction wheel torque output does not perfectly follow Eq. (8.25). Instead, imperfections in the reaction wheel controller cause a noise in the torque levels provided. For the Sinclair Nanosatellite Reaction Wheels (RW-0.03) used as reference, the controller noise is equal to 0.5μNm. A zero mean white noise contribution with a standard deviation of 0.5μNm is therefore added to Eq. (8.25).

Imperfections in the reaction wheel assembly and its alignment cause additional jitter and a cross-coupling component to the torque output. Instead of modelling these disturbance sources, the worst case values according to the equations in Section 3.3.2 are added to the disturbing moment M_d in Eq. (8.22), Eq. (8.23) and Eq. (8.24).

8.2.3. Modelling environmental disturbances

Similarly to the internal reaction wheel disturbances, the worst case environmental disturbances are used in the image navigation simulation. These worst case disturbances are computed according to the equations in Section 3.3.2. In reality, it might not be possible for the worst case disturbance values of the various disturbance sources to occur at the same time. However, since we are not simulation entire satellite orbits we would still like to know how the satellite responds to these worst case disturbances. Once again, the worst case values are added to the disturbing moment M_d in Eq. (8.22), Eq. (8.23) and Eq. (8.24).

8.3. Set-up description

A step by step explanation of the simulation steps by means of a schematic view of the simulation is provided in Fig. 8.1. The image registration block, indicated in Fig. 8.1, uses the set-up for the image registration simulation described in Section 6.3.

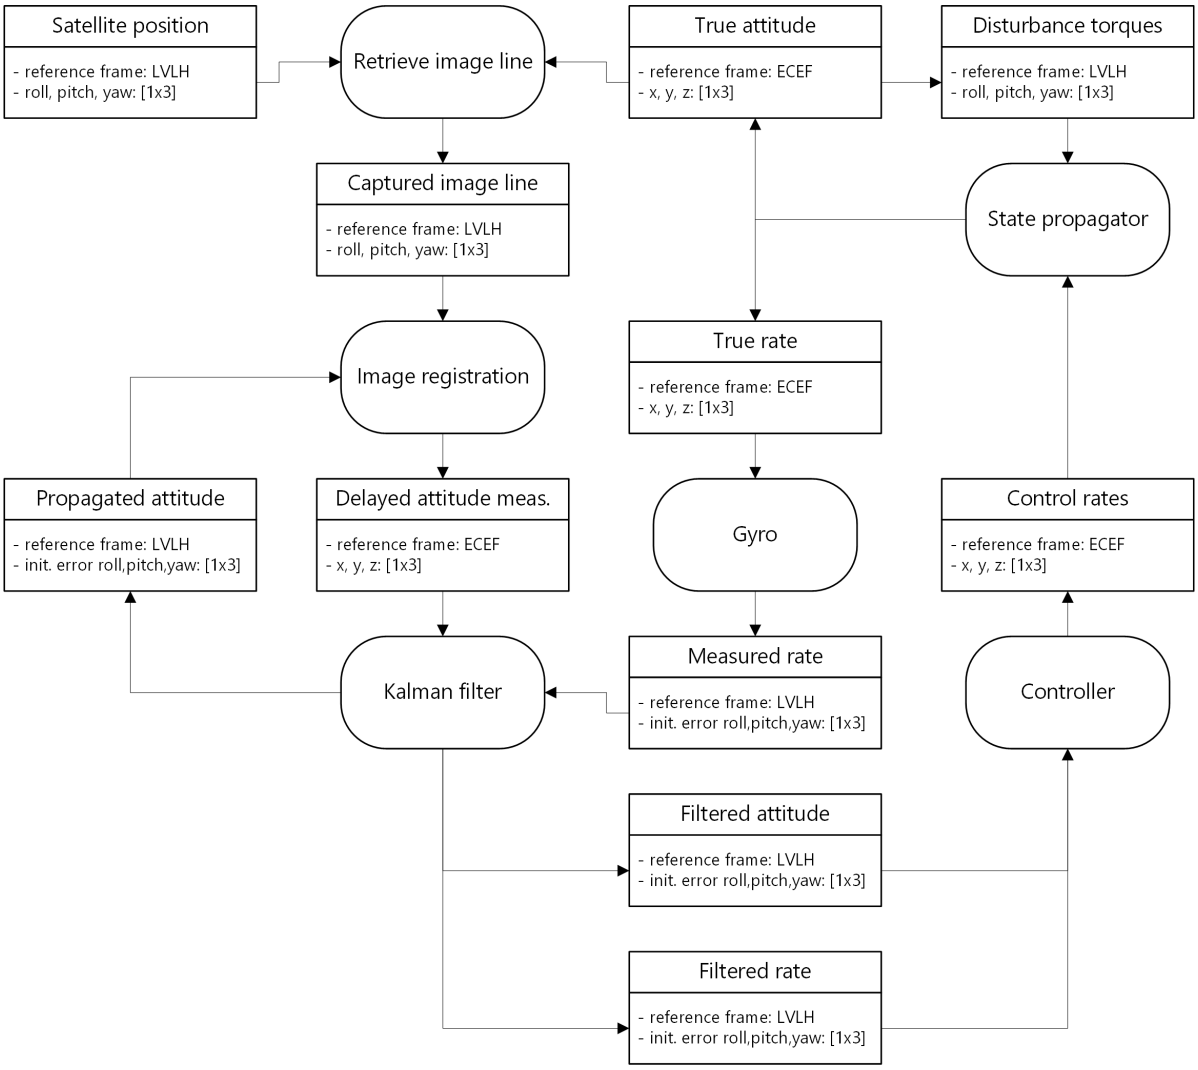

Figure 8.1: Image navigation simulation architecture

9

Image navigation performance and sensitivity study

The image navigation simulation allows a more in depth investigation of the attitude determination accuracy achievable using the image registration algorithms. We will first look at the achievable image navigation performance using the reference gyroscope and discuss the performance of the attitude filter and controller. We will then investigate the sensitivity of the image navigation performance to the algorithm parameters and gyroscope properties. Next we will look at the acquisition phase and the ability of the image navigation system to achieve fine attitude determination starting at a large initial attitude knowledge error. Lastly, we will discuss the measurement accuracy sensitivity to occlusion effects, like cloud formations, snow and ice.

In this and following sections, we make a distinction between the attitude measurement accuracy and the real time attitude determination accuracy. The former is the accuracy of the attitude measurements provided by the image registration method. The later is the accuracy of the attitude determination after filtering of the attitude measurements combined with gyroscope rate measurements, taking into account the delay in the availability of the attitude measurements due to the required processing time.

9.1. Achievable performance

Fig. 9.1 shows the achievable attitude determination and control performance over an occlusion free region.

The accuracies achieved are tabulated in Table 9.1. We see a significant improvement of the real-time determination accuracy compared to the attitude measurement accuracy. The improvement in real-time determination accuracy is more than a factor two higher compared attitude measurement accuracy.

Table 9.1: Maximum achievable attitude measurement and real-time attitude determination accuracy.

Attitude angle	**Measurement accuracy (1σ) [°]**	**Attitude determination accuracy (1σ) [°]**
Roll	0.016	0.0063
Pitch	0.019	0.0079
Yaw	0.041	0.015

As expected, based on the results of the image registration simulation. Highest accuracy is achieved for the roll angle, followed by the pitch angle. Attitude measurement accuracy for the yaw angle is significantly worse. However, the improvement in real-time accuracy achieved after filter is greater for

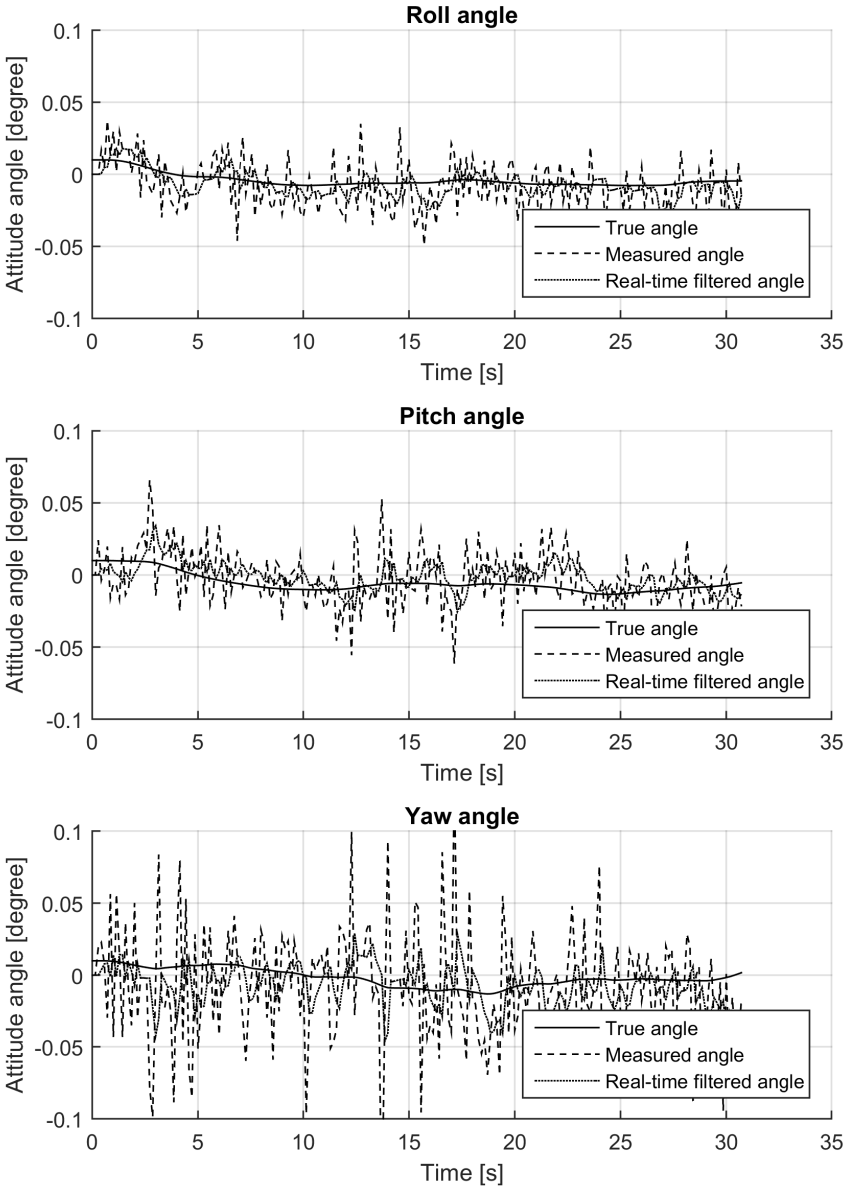

Figure 9.1: Simulation results for maximum achievable attitude determination accuracy.

the yaw angle. This leads to a less smaller absolute and relative difference in accuracy between the three angles.

We also notice a delay in the response of the satellite to the attitude measurements. The inertia of the satellite seems to act as another filter on the real-time attitude measurements. This further improves the stability of the satellite and reduces rotational rates. The magnitude of the rotational rates does not exceed 0.01$^\circ/s$ for all three axes.

As we have already seen in Section 7.2.1, the attitude measurements have a good trueness. The mean error for all three axis over a period of 30s is below 0.002 degree. Because of the good trueness, the smoothing effect caused by the satellite's inertia improves the pointing accuracy to beyond the accuracy of the real-time attitude determination. This effect is especially visible for the pitch and yaw

axes, which suffer from a lower determination accuracy compared to the roll angle. The Root Mean Square (RMS) error of the pointing accuracy for all three axis is tabulated in Table 9.2.

Table 9.2: Maximum achievable pointing accuracy.

Attitude angle	Pointing error (RMS) []
Roll	0.0062
Pitch	0.0065
Yaw	0.0066

9.2. Performance sensitivity

To come to more optimal overall image navigation design we investigate the effect of algorithm parameters and gyroscope performance on the achieved attitude determination accuracy. Algorithm parameters allow tuning of the algorithm to achieve attitude determination accuracy with a minimal computational cost. Gyroscope improves filtering and relaxes the required attitude measurement accuracy.

9.2.1. Algorithm parameter sensitivity

The registration algorithm allows three main parameters to be adjust. These are the width of the correlation area, the number of correlation steps per area and the size of each of these steps. The algorithm parameters are key in achieving the required performance in a computationally efficient way. The width of the correlation areas determines the amount of image pixels processed. Therefore, the computational cost of the registration task scales one on one with the correlation area size. The number of correlation steps determines the total number of correlations performed and therefore also influences the computational cost of the registration. However, computational cost of the coordinate mapping process is not dependent on number of correlations performed, hence the number of correlations influences the total computational cost of the registration task to a lesser extend.

Correlation area width

Fig. 9.2 shows the influence of the correlation area width on the registration accuracy. As expected, the attitude measurement accuracy worsens with a decrease in the correlation area. However, when moving from an area width of 300 pixels to 150 pixels, the measurement accuracy decrease is around 20% for the pitch and yaw angles and no more than 5% for the roll angle. The computational cost, however, decreases by a factor of two. This means the computational cost decreases relatively to the achievable accuracy.

Correlation step size and number of steps

The product of the number of correlation steps and the step size of each correlation determines the size of the correlation window. The effect of the correlation window size on the trueness of the registrations was discussed in Section 7.3.3. Here, we saw a decrease in trueness when the size of the correlation window was lowered to below a certain value. This was attributed to clipping of the Gaussian correlation distribution found for each of the individual registrations.

We investigate whether the same effect is visible for the attitude measurement accuracy in the image navigation simulation. Fig. 9.3 shows the roll angle attitude measurement accuracy as a function of correlation window size for a range of correlation step sizes. For all four step size settings, the measurement accuracy displays a minimum near a correlation window size of 3500 to 4000 meters. This corresponds to the width of the Gaussian distribution found for the cross-track registration in Section 7.3.3. For smaller correlation window sizes, clipping of the Gaussian once again lowers the trueness and therefore the measurement accuracy. For a larger correlation window, accuracy also lowers. This is likely due to a decrease in precision, since the registration is prone to an increased magnitude of random errors.

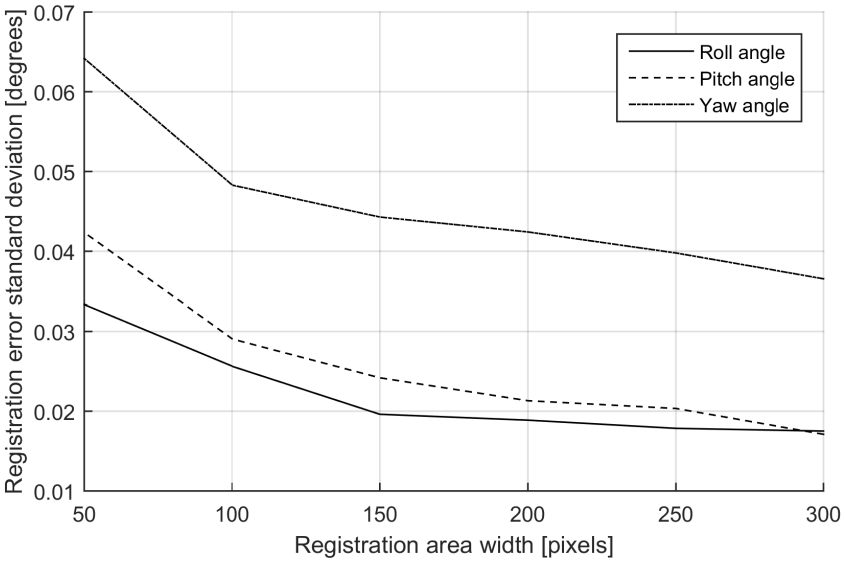

Figure 9.2: Sensitivity of attitude measurement accuracy to correlation area width.

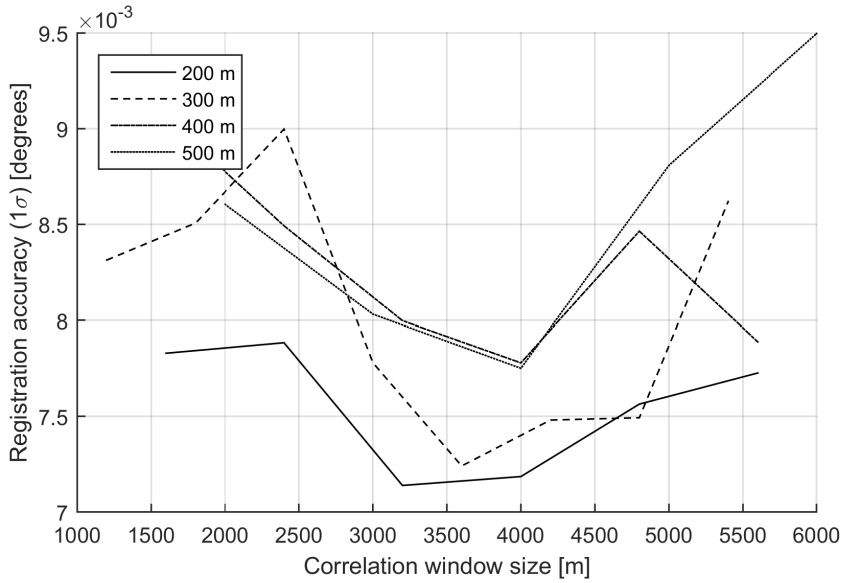

Figure 9.3: Sensitivity of roll angle attitude determination accuracy to correlation window size.

In Fig. 9.4 we see the attitude measurement accuracy as a function of correlation window size for the pitch angle. We once again see an optimum correlation window size. For the pitch angle the optimum occurs at a correlation window size between 4500 and 5000 meters. This corresponds to the larger width of the Gaussian distribution found for the along-track registrations in Section 7.3.3.

If we look at the various step sizes presented in Fig. 9.3 and Fig. 9.4, we see that a lower correlation step size also increases the measurement accuracy. However, a lower step size for a given correlation window size results in an increased number of correlations and thus computational cost. This means the optimal correlation window size gives the highest possible attitude measurement accuracy, but is not necessarily the most computationally efficient window size to meet the required attitude determination accuracy. A smaller correlation window with a decreased correlation step size might allow the required accuracy to be met using fewer correlations. This optimisation is performed in Chapter 10.

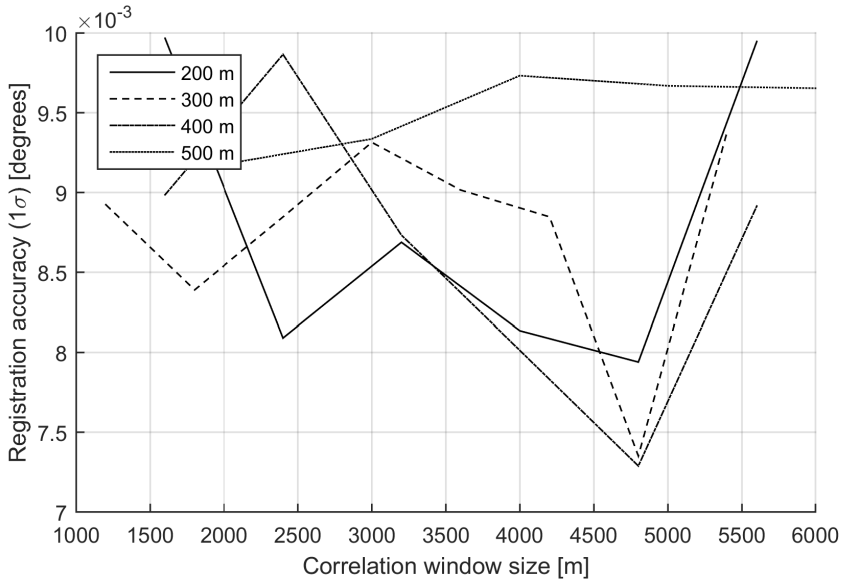

Figure 9.4: Sensitivity of pitch angle attitude determination accuracy to correlation window size.

9.2.2. Gyroscope performance sensitivity

The addition of rotational rate measurements provided by a gyroscope allows an improvement in the attitude determination accuracy after filtering of the attitude measurements. This increases the maximum attitude determination accuracy achievable by the image navigation system. For a given required determination accuracy, a higher gyroscope performance can relax the requirement on the accuracy of the attitude measurements. This means higher gyroscope performance can reduce the computational power needed for the registration process. To come to a more optimal overall image navigation system design, we investigate how the gyroscope performance affects the attitude determination accuracy after filtering. We investigate the effect of gyroscope noise and bias separately.

Fig. 9.5 shows the effect of the gyroscope angular random walk on the filtered attitude determination accuracy. Zero gyroscope bias is assumed. We see a gradual increase in attitude determination accuracy with a lower gyroscope ARW.

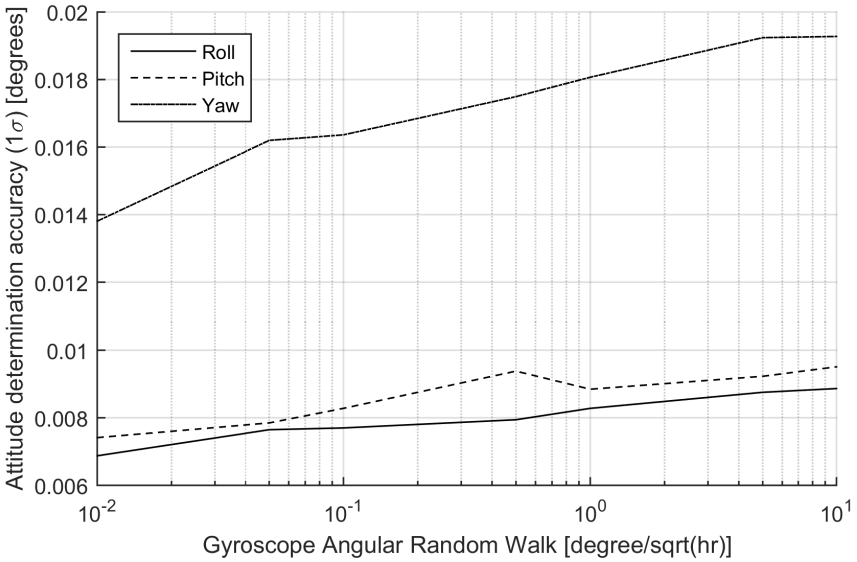

Figure 9.5: Sensitivity of attitude determination accuracy to gyroscope ARW.

Effect of the gyroscope bias on the filtered attitude determination accuracy is presented in Fig. 9.6. We see a slow but gradual decrease of accuracy up to a bias value of around $10^{\circ}/hr$. At higher bias values, the determination accuracy decreases rapidly. This is caused by the gyroscope bias exceeding the order of magnitude of the measured satellite rates. It is easily understood that, to achieve the added value of rate measurements in the filtering of the satellite's attitude, gyroscope bias cannot exceed the magnitude of the measured rotational rates. This puts a maximum on the bias level of the selected gyroscope of around $10^{\circ}/hr$.

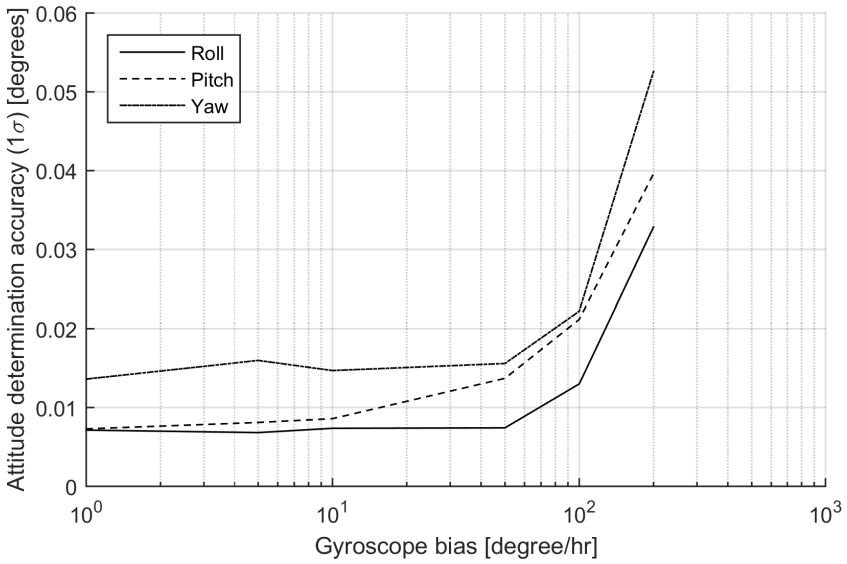

Figure 9.6: Sensitivity of attitude determination accuracy to gyroscope bias.

9.3. Image navigation acquisition

At the start of an observation run, initial knowledge of the satellite's attitude will be provided by the sun sensors. We must ensure this initial knowledge is sufficiently accurate for the image navigation system to acquire successful attitude measurements and allow the attitude knowledge to converge to a steady-state.

Clearly, the maximum initial error for which registration is possible is dependent on the size of the correlation window. If the true attitude value lies outside the correlation window, successful registration cannot occur. However, there is a maximum initial error for which one-dimensional correlation is possible, regardless of the correlation window size.

Worst case initial error combination

To determine the maximum initial error for which convergence is possible, we must first determine what the worst case combination of initial errors for the three attitude angles is. Interestingly, it is found that the worst case combination of initial errors is not a maximum error for each of the three attitude angles. When we set the initial error to 0.4 degree for each of the three attitude angles, we still achieve consistent convergence, as can be seen in Fig. 9.7. However, Fig. 9.8 shows that if we set both roll and yaw initial error to 0.4 degree, but yaw initial error to zero, convergence does not occurs.

This effect is due to the roll angle measurement performed using registration of the central correlation area. A pitch angle initial error causes a along-track offset between the captured image line location and the location of the reference images created, as visualised in top image in Fig. 9.9.

The yaw angle initial error causes a rotation of the constructed reference image line locations. This rotation has a displacement component in along-track direction, partially compensating for the error in pitch angle in one half of the central correlation area. In the bottom image in Fig. 9.9, we see the partial

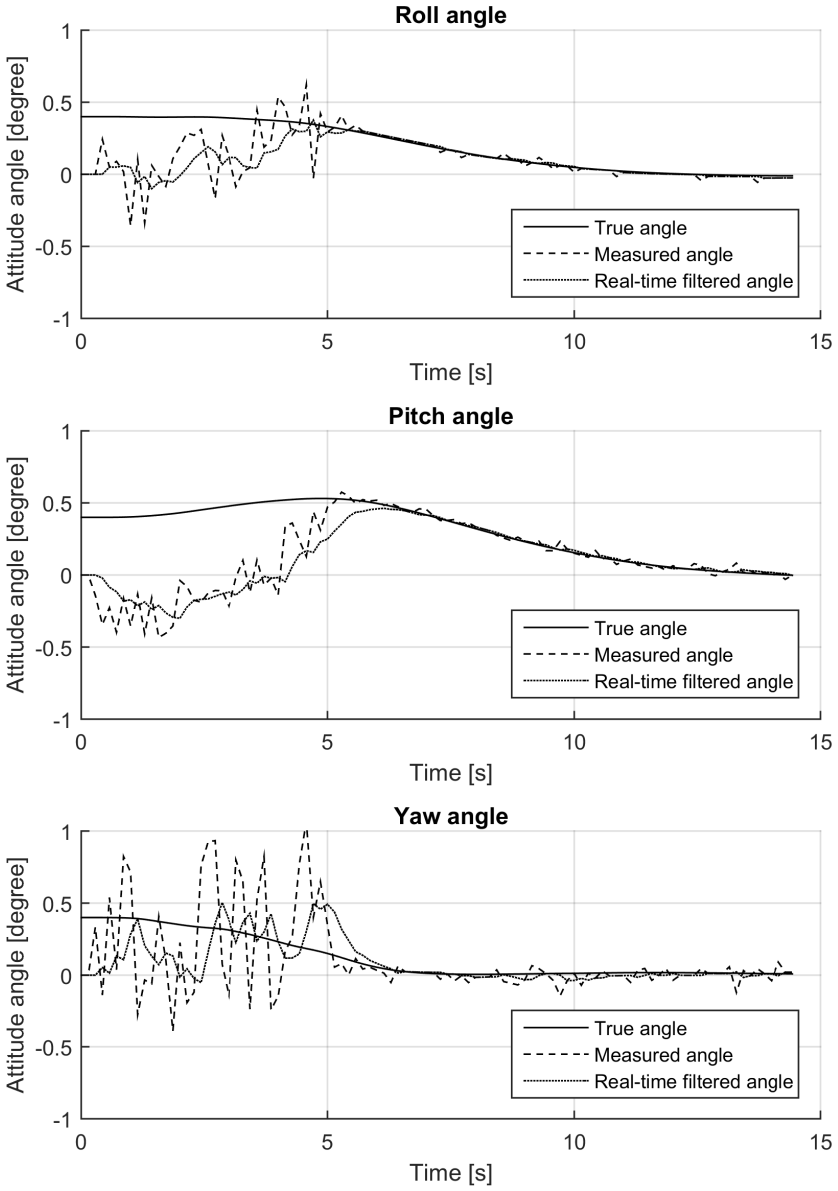

Figure 9.7: Convergence results for roll, pitch and yaw initial error of 0.4 degree.

overlap between sensed and reference image created by the yaw angle initial error contribution. This partial overlap now allows measurement of the roll angle. Following the roll angle acquisition, pitch angle and yaw angles can be measured as well. This sequence can be seen in Fig. 9.7.

Maximum initial error

We define the maximum allowed initial error as the combined roll and pitch angle for which consistent instantaneous convergence of the attitude knowledge is achieved.

We find this maximum initial error is equal to 0.3 degrees. It is a maximum initial error inherent to the one dimensional registration method and cannot be improved by the tuning of the algorithm parameters or by improved gyroscope performance. Both algorithm parameters and especially gyroscope perfor-

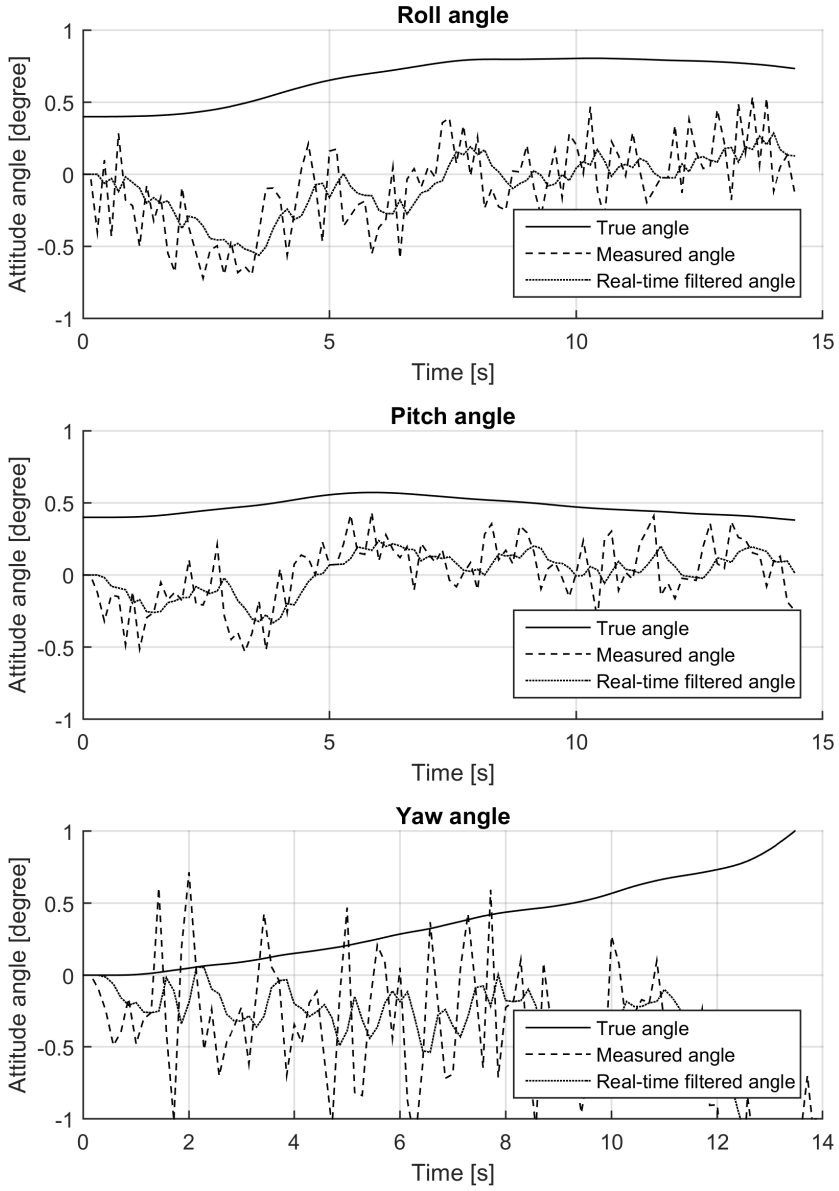

Figure 9.8: Convergence results for roll and pitch initial error of 0.4 degree, and zero yaw initial error.

mance do, however, influence the time needed for the satellite to achieve stable nadir pointing. Given the maximum initial error and reference gyroscope performance, Fig. 9.10 shows time needed to reach stable nadir pointing is within 10 seconds. This is well within the requirement of 30 seconds.

9.4. Performance under occlusion

A major drawback of an image navigation system is its reliance on the successful capture of Earth surface imagery. Even during observation mode, usable imagery is not always available. This can be due to the crossing of a large body of water or occlusion due to cloud coverage, or snow and ice sheets. Since trace gas measurements techniques also require an unobstructed view of the Earth's

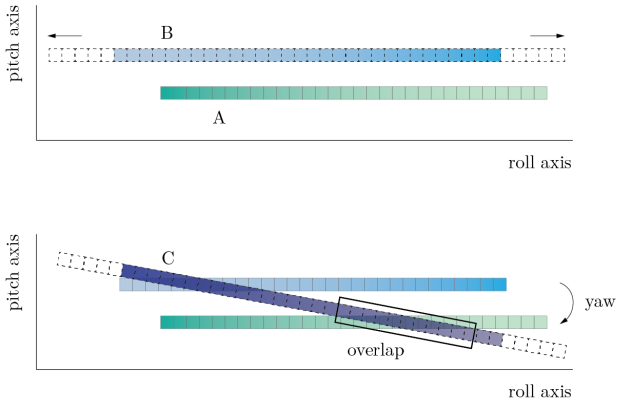

Figure 9.9: Visual explanation of worst case initial attitude knowledge scenario.

surface, fully occluded imagery will be discarded on board the satellite and will not be down-linked for further processing. Therefore, fine attitude determination is not needed during occlusion of the entire instrument swath, allowing fallback to gyroscope integration or sun sensor measurements. However, this changes when only part of the instrument's swath is occluded.

Therefore, we investigate the achievable performance of the image registration algorithm using imagery with up to 50% occlusion. Occluded pixels are generally not randomly distributed over an image line. Instead, cloud formations, water bodies and other sources of occlusion tend to obstruct entire regions from view. This increases the difficulty of measuring the satellite's attitude, since it is easily possible for an entire correlation area to be occluded. To take this grouping effect into account, we use imagery of various regions captured at different instances. Each of these instances are prone to a variation of realistic occlusion. Using water, cloud, snow and ice masks provided by the MODIS team, we can match an occlusion percentage to each registration. Table 9.3 shows the attitude measurement accuracy for different ranges of occlusion. We see the attitude measurement accuracy drops rapidly with the occlusion factor. this is especially true for the yaw angle measurements.

Table 9.3: Sensitivity of attitude measurement accuracy to image line occlusion.

Occlusion [%]	**Roll measurement accuracy (1σ) [°]**	**Pitch measurement accuracy (1σ) [°]**	**Yaw measurement accuracy (1σ) [°]**
0-10	0.019	0.020	0.052
10-20	0.025	0.029	0.065
20-30	0.037	0.042	0.087
30-40	0.054	0.049	0.11
40-50	0.070	0.064	0.13

To assess performance of the image navigation system, we are not so much interested in the registration performance, but rather in the real-time attitude determination accuracy we can achieve using these registrations. However, at this point, the simulation is not capable of providing this information. False registrations due to large occlusion percentages are not detected. Incorporating these false registrations in the attitude filter causes large errors in the estimated attitude. Multiple false registrations can cause an estimation error exceeding the initial error for which successful registration is possible, meaning the system is unable to recover when unobstructed image lines are once again available. Furthermore, the attitude knowledge available during the capture of an image line is not dependent on the registration accuracy for this image line, but rather on the registration accuracy of previous image lines fed into the attitude filter.

For these reasons, estimation of the quality of a registration is key in both the design of a successful image navigation system, but also in determining how this system performs under occlusion and how

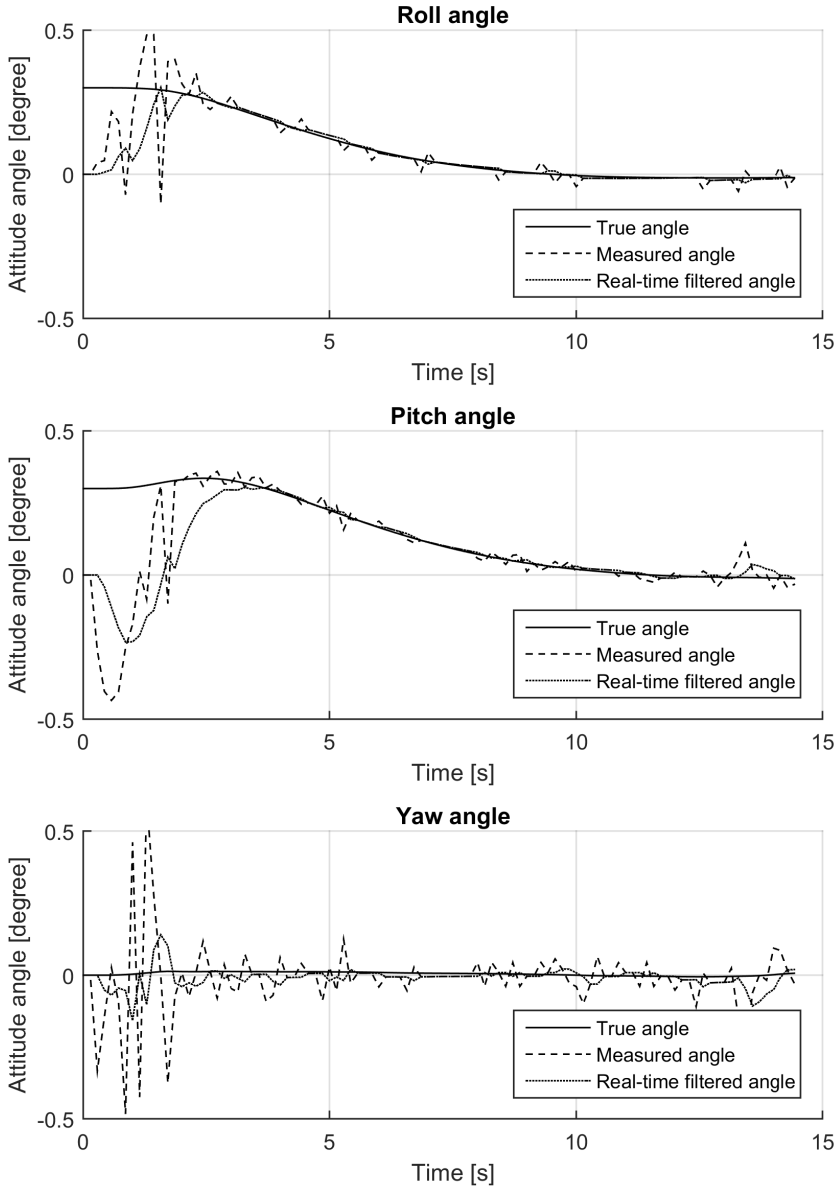

Figure 9.10: Convergence to stable nadir viewing with an initial error of 0.3 degrees in roll and pitch angles.

much potentially usable imagery is lost due to a lack of attitude knowledge accuracy. Three methods of determining the registration quality were investigated. The first method is based on computing the occlusion factor of an image line before it is registered. Major advantage of this method is that the occlusion factor is already computed by the payload data handling system, so image lines with an excessive amount of occlusion can be rejected on board and do not have to be stored and down-linked. If the occlusion factor is considered too high for successful registration, further processing of the image line is not needed. A second method investigated is to use the distribution of the correlation values for a registration to determine the registration quality. As we have seen in Section 7.2.2, the correlation distribution of a successful registration should form a Gaussian distribution around the true attitude. Other distribution shapes can be a sign of a false registration. The last method investigated was the use of automated Kalman filter tuning based on the variance of the attitude measurements.

If the measurement variance is large, less trust will be placed on these measurements by altering the co-variance matrix coefficients. Disadvantage of this method is that it will need multiple consecutive measurements to estimate the variance, leading to a delayed response once an occluded region is encountered. Registration quality estimation and especially the switching between attitude measurement sources is a complex task. Successful implementation of such a method was considered beyond the scope of this research project. Therefore, the results of the registration quality estimation simulations are not presented in this document, but remain available for any future researcher continuing the work on an image navigation system.

Part III

Detailed design and design analysis

10

Detailed design

To assess the feasibility of implementing an image navigation system on board SPECTROLITE, the design of such a system is worked out in more detail. Following the design philosophy of the SPECTROLITE satellite, only space proven COTS hardware components are used.

As we have seen in Chapter 9, both the available processing power and the performance of the selected gyroscope have an effect on the achievable accuracy of the image navigation system. There is significant design freedom in the processing capacity on board the satellite. By selecting one or more available processing units and varying their processor clock frequencies, the available processing power can be tailored to the need of the image navigation algorithm. However, the choice of high performance gyroscope for SPECTROLITE class satellites is far more limited, especially when a bias over temperature value below the $10^{\circ}/hr$ limit is required. Therefore, the gyroscope used was selected first, after which the algorithm parameters were optimised and the processing units were selected. The algorithm parameter optimisation will be described in the next section, followed by the selection of hardware components.

10.1. Algorithm parameter optimisation

The goal of optimising the registration algorithm parameters is to meet the attitude knowledge requirements with minimum processing power. As we have seen in the previous chapter, the three parameters which can be varied are the correlation area width, number of correlation steps and the correlation step size. Of these three parameters, only the first two influence the processing power required to run the algorithm within the time limit set by the imaging frequency of the instrument.

The results of the sensitivity analysis in Section 9.2.1 has given us some insight into the relationship between the determination accuracy and the algorithm parameters. Still, the accuracy under varying algorithm parameters can sometimes show unexpected behaviour. Although we have some idea of the mechanisms at play in this optimisation problem, we cannot explain all results with certainty. It was therefore decided to use a brute-force optimisation method to select the final algorithm parameters.

To limit the computational cost of the optimisation, parameter value limits were set based on the results of the sensitivity study. Furthermore, the number of candidate values per parameter was limited. This also limits the level to which the final solution is optimised. However, based on the optimisation results, it is believed the effect of further optimisation is limited. The optimisation limits are tabulated in Table 10.1.

The selected algorithm parameters are a correlation area width of 200 pixels and a step size of 0.02° with a total of 11 steps.

Table 10.1: Image navigation algorithm parameter optimisation limits.

Parameter	Lower limit	Upper limit	Step size	Unit
Area width	50	300	25	[pixels]
Number of steps	7	41	2	[-]
Step size	0.005	0.08	0.005	[°]

10.1.1. Acquisition

The size of the correlation window, given the selected algorithm parameters, is 0.1°. This means initial attitude knowledge with an accuracy smaller than 0.05° is needed for successful registration to occur. Although propagation of the Kalman filter can easily provide this level of accuracy once the attitude knowledge has converged, this requirement puts significant strain on the attitude knowledge accuracy needed at the beginning of an observation run. At the beginning of an observation run, the initial attitude knowledge is provided by the secondary attitude determination sensors.

The accuracy requirement placed on the secondary attitude determination sensors can be loosened by increasing the correlation step size during image navigation acquisition. As we have seen in Section 9.3, the maximum initial attitude knowledge error for which immediate and consistent convergence is reached is 0.3°. This value can be achieved by increasing the step size during acquisition to 0.07°. Once the attitude knowledge and satellite pointing has converged to the steady state, the step size is lowered to the steady state value of 0.02°.

Table 10.2: Final image navigation algorithm parameters.

Parameter	Steady state value	Acquisition value	Unit
Area width	200	200	[pixels]
Number of steps	11	11	[-]
Step size	0.02	0.07	[°]

10.2. Hardware selection

For each of the required hardware components needed in the image navigation system, an extensive market search was performed. To achieve a fair comparison, only COTS components with flight proven performance were considered. This severely limited the number of available choices in each component selection process.

10.2.1. Processing unit

The main consideration in selecting the processing unit was the processing power required to run the image navigation algorithms in real-time. Therefore, a method of determining the time required to run the algorithms on a given processing unit was needed. Roughly two methods can be used to determine this processing time. The more theoretical method is to compute the amount of instructions and floating point operations needed to run the entire set of algorithms. This theoretical value can then be compared to the Instructions Per Second (IPS) or Floating Point Operations Per Second (FLOPS) achievable by the processing unit. However, the IPS and FLOPS values reported for processing units are measured using artificial instruction sets, generally leading to the peak achievable performance. The practical instruction sets needed to run an algorithm, especially the more complex ones with multiple branches, lead to a significantly lower achievable number of IPS or FLOPS.

The alternative is a more practical approach using a synthetic benchmark. This synthetic benchmarks performs a range of different computations and memory transfers to mimic a realistic algorithm implementation. By running the benchmark algorithm on multiple processing units, the relative computing power of these processing units can be determined. Using the computation time on one computing unit

and the relative benchmark performance, the computation time on a second computing unit can be estimated. The most commonly used benchmark is the Dhrystone algorithm, reported as Dhrystone MIPS or DMIPS. The DMIPS values for both the simulation computer, as well as the considered on board processing units were found and used to determine the on board processing time required.

The selected processing unit is based on an ARM processor commonly found in modern smartphones. This allowed additional comparison of simulation computer and on-board processing unit performance using a generic multi-platform benchmarking algorithm. Performance values for both processing units were found for the Geekbench 3 benchmark. For the on-board processing unit, the reported value used is not for the specific chipset selected, but rather an average value found for multiple different chipsets with the same ARM processor and clock speed. The relative performance between the simulation computer and on-board processing unit found for both benchmarks differ by only 10%.

Although it is believed these benchmarks provide a good first order estimation of the processing power required on board the satellite, only running the algorithm on the selected processing unit will provide definitive proof on whether the selected processing power is sufficient to run the image navigation algorithms in real-time. It must be noted, however, that the simulation code is in no way fully optimised. Especially due to the large amount of data processed by the image navigation algorithms, it is believed that code optimisation can lead to a significant reduction in computation time.

Selected processing unit

No available processing units were found capable of running the image navigation algorithms in real time on a single processor. Instead, three GOMspace NanoMind Z7000 computers were selected to perform the image navigation tasks. One of these processing units has sufficient processing power and memory to perform the registration of a single correlation area in the available time span, with some additional computation time left to perform the registration to attitude computations on a single processor. By performing each correlation area registration on a dedicated processing unit, parallel computing issues will be limited. The specifications of the selected processing unit are tabulated in Table 10.3.

Table 10.3: Selected processing unit specifications.

Specification	Value
Processor	Dual ARM Cortex A9
Core frequency	800 MHz
Memory	1 GB DDR3L
Storage	4 GB eMMC
Mass	28 g
Volume	65x40x6.5 mm
Maximum power	2.3 W
Cost	$\sim$ €5000

We assume the processing units run at maximum power during observation mode, as nearly the full extend of their processing speed is used for the image navigation tasks.

Storage capacity

Significant storage capacity is needed for both the reference imagery and the digital elevation model data. Given a reference imagery resolution of 500 m and a DEM resolution of 90 m, each square kilometre contains 4 intensity values and 124 elevation values. If we use one byte per data point, the reference database will require 128 bytes per square kilometre.

To determine the total storage capacity needed, the landmass surface area in the region of interest between 29° and 56° is computed. The surface are is found to be equal to 8% of the entire Earth's surface or 408×10^5 km^2. This results in a total reference database size of 4.9 GB. This is larger than the reported storage capacity available on a single processing unit chipset of 4 GB. However, these

units allow allow additional flash storage to be added up several 100 GB. Alternatively, the data can be distributed among the three processing units, since communication between them is possible.

10.2.2. Gyroscope

The selected gyroscope is the Sensonor STIM 210, which has been extensively used on board micro and nano satellites. It is the only ITAR free COTS gyroscope found providing a bias over temperature value not exceeding the $10^{\circ}/hr$ limit at a mass of under 400 g. The relevant specifications of the gyroscope are tabulated in Table 10.4.

Table 10.4: Selected gyroscope specifications.

Specification	Value
Angular Random Walk	0.15 $^{\circ}/\sqrt{hr}$
Bias stability	0.3 $^{\circ}/hr$
Bias over temperature	10 $^{\circ}/hr$
Mass	52 g
Volume	38x45x22 mm
Maximum power	1.5 W
Cost	$\sim$ €5000

It is assumed the gyroscope will keep running during an entire orbit to limit rotational rates during eclipse and to allow accurate bias estimation. It is believed that bias estimation with an accuracy below the bias over temperature value is possible, however, further research is needed to determine what the achievable bias estimation accuracy is.

10.2.3. Sun-sensor

The sun-sensor accuracy needed is based on the maximum initial error requirement of the registration method. The SolarMEMS SSOC-A60 sensor is selected, which has a 3σ error equal to the maximum allow initial error of 0.3 degrees. The relevant specifications of the sun-sensor are tabulated in Table 10.5.

Table 10.5: Selected sun sensor specifications.

Specification	Value
Accuracy	0.3° (3σ)
Mass	25 g
Volume	30x30x12 mm
Maximum power	36 mW
Cost	€7200

10.2.4. Attitude determination system properties

Combining the properties of individual attitude determination system components, provides an indication of the properties of the attitude determination system. These system properties are tabulated in Table 10.6. Note that this only includes the combined properties of the system components and additional mass and volume will be needed to integrate the components into a system in the satellite.

10.3. Physical architecture and timing

An example physical architecture of the image navigation system is presented in Fig. 10.1. This architecture shows how the image navigation system components are embedded with in the traditional

Table 10.6: Attitude determination system specifications.

Specification	Value
Total mass	236 g
Maximum power	8.4 W
Average power over orbit	2.0 W
Cost	$\sim$ €48800

structure of payload and attitude determination system.

10.3.1. Physical architecture

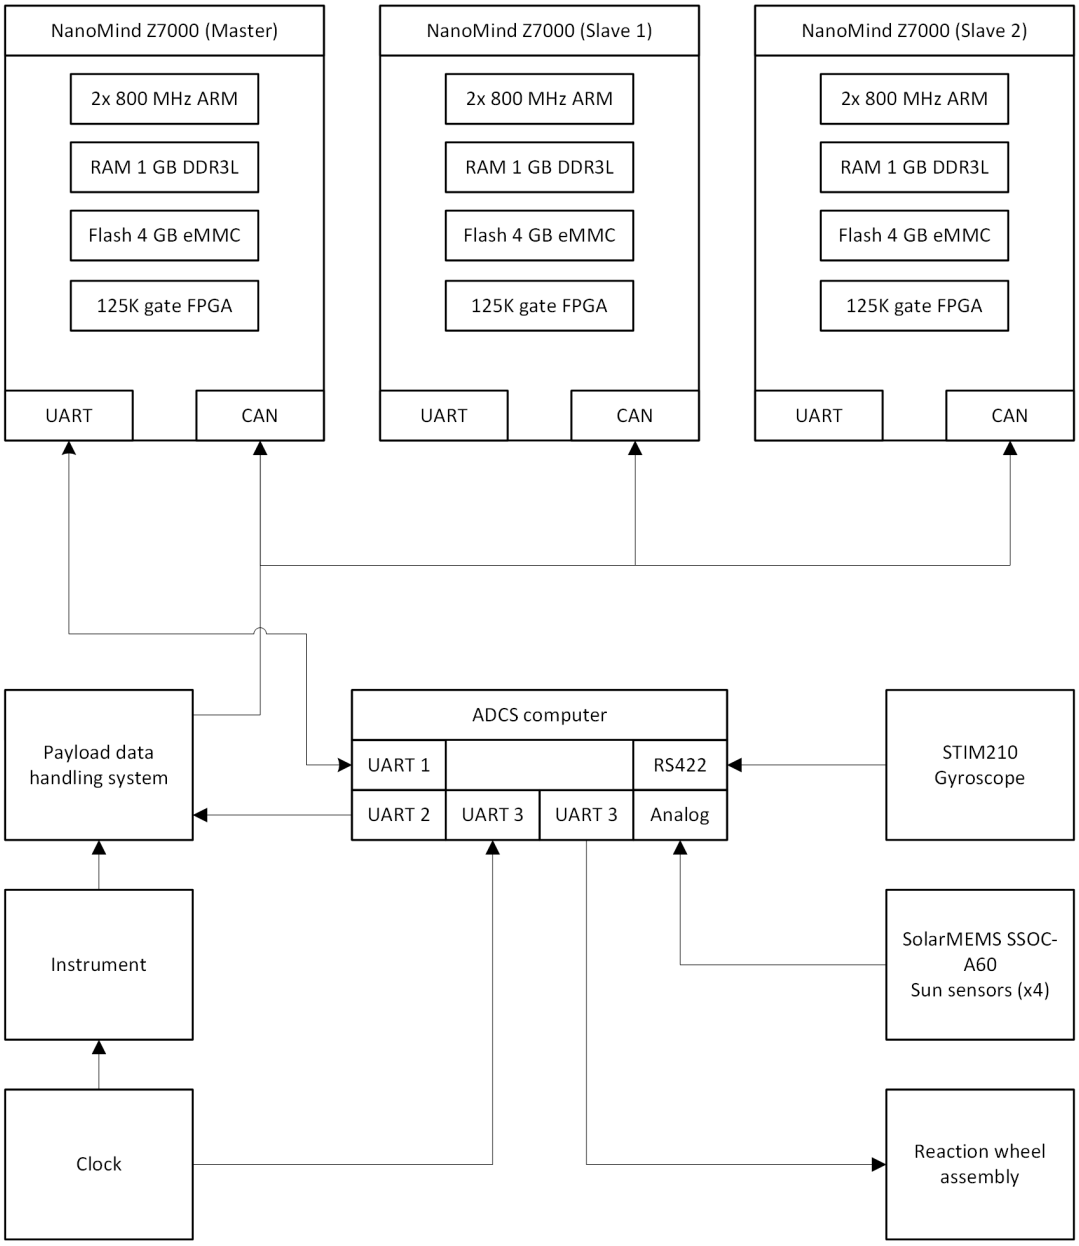

Figure 10.1: Visualisation of the attitude determination system physical architecture including connected systems.

10.3.2. Image navigation timing sequence

The timing sequence for the image navigation process is once again visualised in Fig. 10.2.

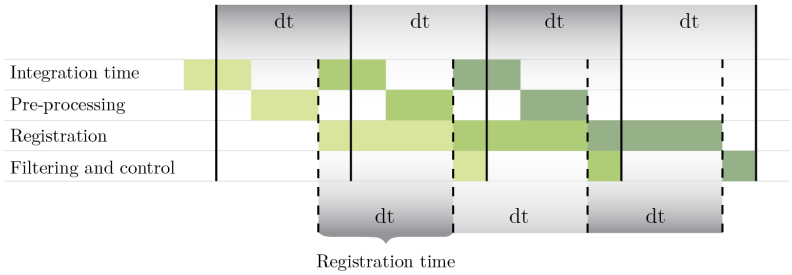

Figure 10.2: Visualisation of the timing sequence for imaging, pre-processing, registration and filtering.

The time available for each sub-process is dependent on the image period of the instrument. The timing budget in Table 10.7 shows the time available for each sub-process for the current SPECTROLITE imaging period.

Table 10.7: Image navigation timing budget.

Task	Allocated time [ms]	Required time [ms]	Margin [%]
Integration time	72	-	-
Pre-processing	72	-	-
Registration	144	132	8.3
Filtering and control	36	-	-

The time allocated for registration task is 144 ms. The estimated time required for the image processing based on the simulations performed is 132 ms, leaving a margin of around 8%. However, as discussed before, significant uncertainty is present in this estimation. The design of the payload data handling system will determine whether the allocated time of 72 ms is sufficient to store and pre-process a captured image line. Similarly, the selection of ADCS computer will determine whether the allocated time of 36 ms is sufficient to perform the state augmented Kalman filter update.

11

Design analysis

After selection of the hardware components, we can analyse the final performance and specifications of image navigation system designed. We determine the performance achievable on board the SPECTROLITE satellite, as well as the performance achievable by post-processing of the data on ground. Furthermore, we analyse the pointing stability and rotational rates achieved for the SPECTROLITE satellite in the image navigation simulation and determine whether the required MTF values can be met. Lastly, we compare the properties of the image navigation system designed to the specifications of commercially available star-trackers found in our market analysis.

11.1. Attitude determination performance

The final registration and attitude determination accuracies are determined by performing image navigation simulations using a range of MODIS imagery. The imagery used is captured at various different times throughout the year to minimise any seasonal effects. Furthermore, imagery from various locations within the region of interest is used to ensure the performance of the image navigation system over various terrain types. In the analysis of the final image navigation performance we look at both the available accuracies on board the satellite, as well as the accuracies achievable by post-processing on ground. The accuracies found were consistent, with only one notable exception found. This anomaly will be discussed at the end of this section.

11.1.1. Real-time on-board

The attitude measurement accuracy and real-time attitude determination accuracy achievable on board the SPECTROLITE satellite are tabulated in Table 11.1.

Table 11.1: Attitude measurement and real-time attitude determination accuracy of SPECTROLITE image navigation design.

Attitude angle	**Measurement accuracy (1σ) [°]**	**Attitude determination accuracy (1σ) [°]**
Roll	0.018	0.0082
Pitch	0.019	0.0082
Yaw	0.035	0.013

Image registration method provides measurements with a 3 sigma accuracy of 200 arcsec for the cross boresight axes. This places its accuracy right in between the most accurate sun sensor and least accurate star tracker found in the market analysis in Table 3.1. However, the image navigation method has a high measurement frequency of around 7Hz and a near perfect trueness. This allows filtering to achieve an attitude determination accuracy of more than a factor two higher. The real-time attitude determination accuracy for the combined image navigation system is around 90 arcsec for the cross boresight axes.

Nadir pointing settling time
The image navigation system allows convergence of the attitude knowledge from a maximum initial error of 0.03° for each axis. The time required for the satellite to achieve stable nadir pointing for a range of initial attitude knowledge errors is tabulated in Table 11.2.

Table 11.2: Settling-time to stable nadir pointing for SPECTROLITE image navigation desing.

Initial pointing error [°]	Roll axis settling time [s]	Pitch axis settling time [s]
0.05	5.2	5.9
0.10	6.8	7.3
0.15	6.3	8.0
0.20	7.0	9.4
0.25	8.1	9.9
0.30	11	13.4

We see that the requirement of achieving stable nadir pointing with 30 s after reaching an observable region is easily met.

11.1.2. Post-processed

Post-processing of instrument and attitude determination data can improve the attitude knowledge available on ground. Since the image navigation algorithms do not have to be run in real-time, the algorithm parameters optimised for processing speed are not used. Instead, the full number of available pixels is used in the correlations, and the number of correlation steps is increased.

Apart from the algorithm parameters, we also use different values for the gyroscope performance. We assume that, by post-processing, we are able to accurately estimate the gyroscope bias at every instance. To determine the maximum achievable accuracies by post-processing, we assume the gyroscope bias is zero. These results are tabulated in Table 11.3.

Table 11.3: Attitude measurement and real-time attitude determination accuracy after post-processing.

Attitude angle	Registration accuracy (1σ) [°]	Attitude determination accuracy (1σ) [°]
Roll	0.014	0.0073
Pitch	0.016	0.0076
Yaw	0.022	0.012

We notice only limited improvement over the achievable real-time on board accuracy. The combined error in roll, pitch and yaw angle determination causes a 3σ geo-location error for the outer image line pixels of 480 m. This value is significantly more than the 100 m requirement set by the customer.

However, additional improvements to the post-processing algorithms might be possible. The image registration technique using one-dimensional correlations is computationally efficient, but likely does not provide the highest possible accuracy. It would be interesting to investigate the effect of using the designed algorithms, but abandoning the one-dimensional correlation technique. Still, the registration method used is specifically designed for use in real-time on board a satellite. Other, more accurate and efficient registration techniques might be available for use in post-processing.

11.1.3. United States High Plains anomaly

Within the tested imagery, a single anomaly in the image navigation performance was found. This anomaly occurs within the High Plains region in the middle of the United States. The location of this region can be seen in Fig. 11.1. The designed algorithm is unable to perform accurate image registration at multiple locations within this region. Closer investigation shows that these locations correspond to agricultural areas where centre-pivot irrigation is extensively used. The High Plains region has a semi-arid climate with extreme temperature variations, meaning irrigation is vital for crop

growing. The availability of underground water from the Great Plains aquifer, has led to this centre-pivot irrigation technique, where crops are watered by rotating sprinklers around a central pivot. This results in a pattern of circular fields as can be seen in Fig. 11.2.

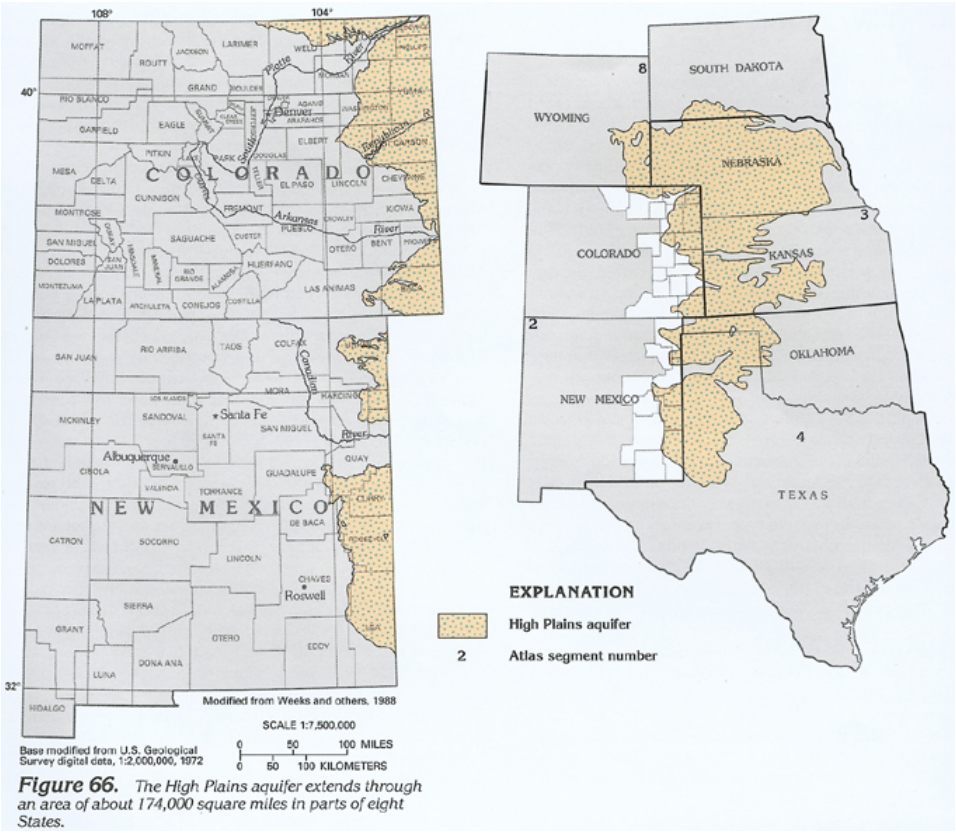

Figure 66. *The High Plains aquifer extends through an area of about 174,000 square miles in parts of eight States.*

Figure 11.1: High plains region [46].

It is not fully understood what causes the difficulty in registering imagery with these circular field patterns. However, it is believed that the size of the fields being roughly the same size as the ground sampling distance plays an important role. Although this is currently the only anomaly in registration performance found, other anomalies might exist. Therefore, it is recommended to further investigate surface type and seasonal effects on the attitude measurement performance.

11.2. Pointing accuracy and stability

Simulation of the satellite's attitude dynamics also allow us to investigate whether the required pointing accuracy and stability to allow the required MTF levels is met. We see that over the range of MODIS imagery simulated, the pointing Root Mean Square (RMS) error for each individual axis is well below the pointing accuracy requirement. The pointing RMS error values are tabulated in Table 11.4.

Table 11.4: Pointing accuracy achievable using SPECTROLITE image navigation design.

Attitude angle	**Pointing error (RMS) []**
Roll	0.0055
Pitch	0.0059
Yaw	0.0079

The pointing error is smaller than the attitude determination accuracy available to the controller. We believe this is due to the satellite inertia in combination with a high measurement trueness. Due to the satellite's inertia, the attitude controller does not immediately move the satellite's attitude to its target

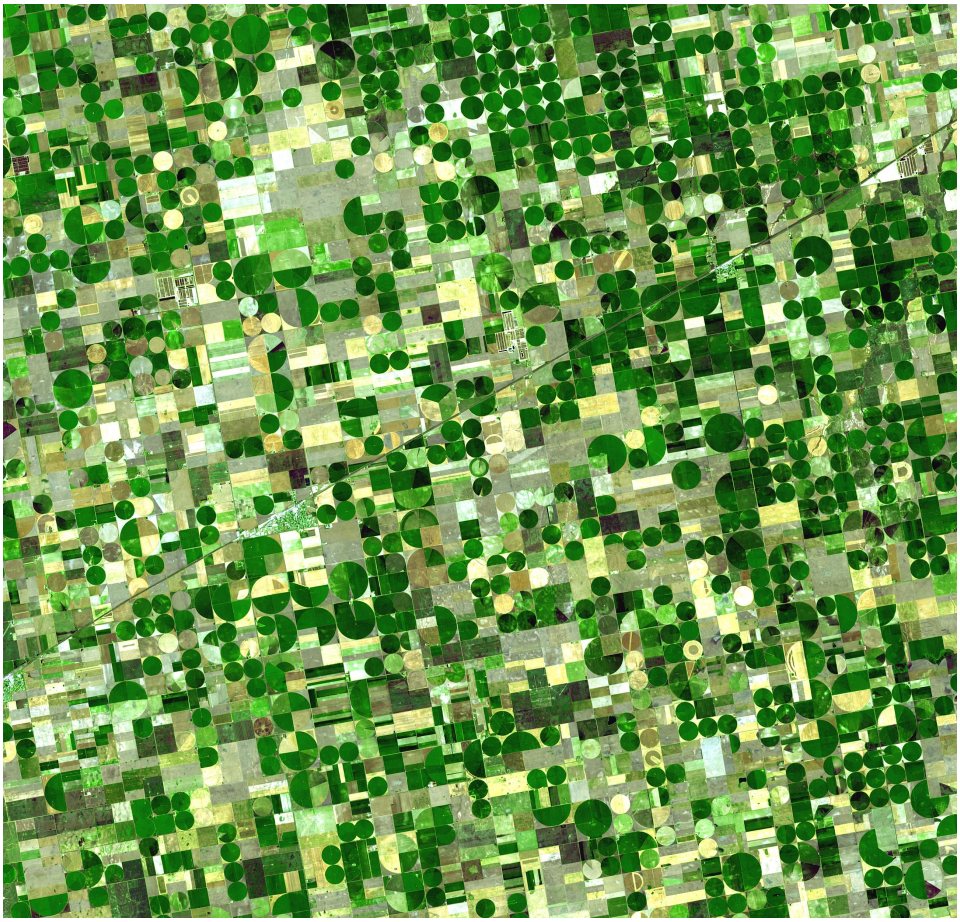

Figure 11.2: Circular field pattern caused by centre-pivot irrigation [47].

value based on the current attitude knowledge. Given the high frequency and high trueness of the measurements, it allows the satellite's inertia to further smooth the output of the Kalman filter.

The rotational rates experienced by the satellite are determine for both a steady-state, as well as during the convergence to the steady-state at the start of an imaging run. The rotational rates during the steady state and during convergence are tabulated in Table 11.5 and Table 11.6, respectively.

Table 11.5: Maximum steady-state attitude rates using SPECTROLITE image navigation design.

Attitude angle	**Maximum rotational rate [$^\circ/s$]**
Roll	0.0076
Pitch	0.011
Yaw	0.032

Table 11.6: Maximum convergence attitude rates using SPECTROLITE image navigation design.

Attitude angle	**Maximum rotational rate [$^\circ/s$]**
Roll	0.051
Pitch	0.061
Yaw	0.083

We see that even during convergence, the rotational rates are surprisingly low, lower than estimated during the pointing budget analysis in Section 3.3. The 0.0430° of motion allowed during the instrument integration time of 72.3 ms, means the maximum total rate of the outer image line pixels cannot

exceed $0.60^\circ/s$. During convergence the maximum rotational for these outer pixels as determine by the simulation is less than a quarter of this requirement. It is believed these low rates can also be attributed to the additional smoothing of attitude measurements caused by the satellite's inertia.

11.3. Star-tracker system comparison

The reported image navigation measurement accuracy 0.018° means it has a 3σ accuracy of around 200 arcsec. This places it at the low end of the small satellite COTS star-tracker accuracies tabulated in Table 3.1. However, a major difference between the image navigation system designed and many of the low end star-trackers is the high trueness and high measurement frequency of the image navigation method. This means that, after filtering, the attitude determination accuracy achieved is lowered to a 3σ value of around 90 arcsec. This performance is on par with the SpaceMicro MIST, weighing in at 500 grams for a single unit. Remember that SPECTROLITE will require two star-tracker units to provide attitude measurements during the entire observation period.

Looking at components sizes of the image navigation system, we see that a significant reduction in volume can be achieved compared to a star-tracker system. Including baffle and support structure, a single star-tracker unit will often house an entire CubeSat unit (10x10x10 cm). The image navigation system on the other hand requires no more than 7x5x5 cm or less than a quarter unit.

The combined mass of the image navigation system components of 236 g is already lower than the combined mass of two units of the lightest star-tracker found (Sinclair ST-16RT2) of 370 g. Furthermore, it is believed that the lower volume and the lack of a need for accurate alignment of the image navigation system will results in a further decrease in the overall system mass, compared to a star-tracker system. However, a detailed design of the image navigation system is needed to provide definitive numbers for the mass of the entire system.

A significant drawback of the image navigation system is its large power usage. By only using the image navigation system during observations, the orbit averaged power consumption of the attitude determination system can be lowered to an acceptable level. However, 2.0 W is still twice as large as the power consumption of two Sinclair star-trackers active throughout the orbit. Reduction of the image navigation system's power consumption will therefore be an important aspect of increasing its feasibility for use on board SPECTROLITE or any other small satellite. Reducing the required processing power for the image navigation algorithms and the use of more power efficient processing hardware can be methods of achieving this goal.

The component cost of the image navigation system is low, at less than half of the cost of a single star-tracker. However, development of an operational image navigation system will likely be costly. Still, an image navigation system has the potential to significantly reduce the recurring costs of the attitude determination system for SPECTROLITE or other small satellites for Earth observation. This could be interesting especially for large Earth observation constellations like the Dove constellation flown by Planet Labs.

A last important comparison we make is the availability of attitude measurements. As we discussed before, image navigation can only be used when the payload instrument is captures usable Earth imagery. Therefore the image navigation system is suited to perform attitude determination in observation mode, but not in other modes. If high pointing accuracy or stability is needed in other modes, image navigation might not be a suitable attitude determination method. Although this is not the case for SPECTROLITE, it could limit the use of image navigation for other satellites. Even in observation mode, the image navigation system will not always be able to provide accurate attitude measurements. Occlusion due to cloud coverage or water bodies can significantly reduce the attitude measurement availability. Star-trackers can have similar occlusion issues when their view of the celestial sphere is blocked by the Earth or by blinding of the Sun. Extension of the image navigation system design and simulations, as discussed in Section 9.4, is required to assess its overall availability and allow comparison to that of a star-tracker system.

Part IV

Conclusions and recommendations

12

Conclusions

This report shows the design of an image navigation system capable of meeting the attitude determination requirements for the SPECTROLITE satellite. We will summarise the most important conclusions found for this design, the feasibility of implementing image navigation in the SPECTROLITE satellite design and the feasibility of image navigation for small satellite missions in general.

Image navigation design
Simulations using MTF adjusted MODIS imagery to represent imagery captured by the SPECTROLITE instrument, show the ability of the incremental sign distance correlation method to register multi-modal satellite imagery captured under varying lighting conditions with high accuracy. Using a simplification of the Toutin push-broom model to compute the estimated coordinates of captured image line pixels, rigidity between a captured image line and the reference imagery is created. This allows sub-pixel registration with a maximum registration accuracy of achieved under 0.15 times the pixel size (1 sigma).

It is also shown that, given a sufficiently accurate initial estimate of the satellites attitude, three separate one-dimensional registrations can be used to measure roll, pitch and yaw angles. The one-dimensional registration method greatly reduces the computational cost of the registration task. This allows the image processing to take place in real-time on board the satellite.

The designed method allows accurate registration of a single image line and therefore allows attitude measurements with a frequency equal to the imaging frequency of the instrument. Together with the high trueness of the image-based attitude measurements and the addition of rotational rate measurements provided by an on board gyroscope, this allows a large improvement in the real-time available attitude knowledge accuracy after filtering.

Attitude measurements with an accuracy (1 sigma) of 0.016, 0.019 and 0.040 for respectively the roll, pitch and yaw angles were achieved. After filtering, the real-time attitude determination accuracy (1 sigma) achieved is 0.0063, 0.0079 and 0.015 for the roll, pitch and yaw angles.

Image navigation for SPECTROLITE
The current design of an image navigation system for SPECTROLITE shows that this type of attitude determination system can indeed be feasible. However, its superiority over a more traditional star-tracker system is not obvious. Although a significant reduction in terms of volume and mass of the attitude determination system might be possible by using image navigation, the power consumption of such a system is estimated to be very large. An important aspect of its feasibility is also the cost associated with an image navigation system. Although component cost is relatively low, the cost to develop an operational image navigation system can be significant. An inherent disadvantage of an image based attitude determination method is the worsened attitude determination performance around areas occluded by for example clouds or water bodies.

At this point, it is still difficult to confidently state whether the cost and risk of developing an image navigation system are outweighed by its possible advantages. However, we are confident that a few years ago, development of such a system would not have been possible due to a lack of powerful processing units for small satellites. At this point, image navigation is right on the edge of feasibility. If we assume the processing power for satellites follows Moore's law, this might be the perfect moment to start the development of an image navigation system. By the time its development is complete, advancements in computing units for space applications could have tipped the scales in favour of using image navigation over other attitude determination methods.

Image navigation for other small satellite missions

We can also draw some preliminary conclusions about the feasibility of image navigation for small satellites in general. An important conclusion is that for the SPECTROLITE satellite, the achievable attitude determination accuracy relative to its imaging resolution is sufficient to allow this resolution to be achieved. The absolute accuracy of the image registration process, and therefore the accuracy of the image-based attitude determination, is highly dependent on the resolution of the images used. If we assume the attitude determination accuracy indeed scales with the resolution of the instrument used, the image navigation system designed could be used for a large range of small Earth observation satellites. If applied to satellites with a higher resolution instrument, the advantages of using an image navigation system over a star-tracker system could improve, as the absolute attitude determination performance achievable increases.

It must be noted, however, that the reference image database grows quadratically with resolution, so reference database size can be one of the limiting factors in the overall achievable accuracy. Therefore, there might be some optimal instrument resolution for which image navigation has the highest feasibility. In general, the trade-off between image navigation and other attitude determination types will be largely dependent on the resolution of the imaging instrument.

13

Recommendations

A number of recommendations will be made based on the research performed. First, we will discuss some future work which can be performed to prove the feasibility of image navigation for SPECTROLITE. Next, we will recommend some future steps in continuing the design of the image navigation system. We will also make some recommendations on possible methods of improvement both the performance achieved using the image navigation system and improving the physical properties of the system, both of which can improve the feasibility of image navigation for SPECTROLITE as well as other satellites.

Proving feasibility
The main goal of this research project was to investigate the feasibility of using an image navigation system for the SPECTROLITE satellite. Proving the feasibility of the image navigation system consists of two main parts. The first part is to prove the required attitude determination accuracy can be indeed be achieved using image based attitude measurements. The image navigation simulation created indicates this accuracy can indeed be achieved. However, some further investigation can be performed to increase the certainty we have in the validity of the simulation results.

Image registration performance, cornerstone of the achievable accuracy of the image navigation system, is highly dependent on the quality of the imagery used. Therefore, we must ensure the MTF adjusted MODIS imagery used is representative for the SPECTROLITE instrument. Measure MTF of adjusted imagery directly. Use imagery from other source instrument with lower quality.

The second key element in determining the feasibility of image navigation for SPECTROLITE is the hardware needed. Especially the processing power needed on board the satellite greatly influences the feasibility. A benchmark comparison method was used to estimate what processing hardware would suffice to run the image navigation algorithms in real-time. However, the only way to actually determine whether the selected hardware is capable of running it in real-time is to perform the simulation on the hardware. Luckily, COTS micro and nano satellite processing power is largely based on the ARM processors found in many other devices. This means a relevant processing unit can easily be acquired.

The High Plains anomaly also shows that surface regions can exist for which registration performance is worse. A number of different terrain types at different times of capture were already investigated. However, more imagery at more locations and more times of capture must be simulated to investigate surface type, lighting and seasonal effects.

The accuracy of the rotational rate measurements provided by the on board gyroscope also have a significant influence on the real-time attitude determination achievable. Currently, gyroscope measurements are simulation using zero-mean white noise and a constant bias value. However, true measurement noise might have a negative impact on the filter performance. Therefore, simulations with

realistic gyroscope measurement noise will improve the realism of the image navigation simulation and its results.

Continuing image navigation system design
An important next step in continuing the design of the image navigation system is devise a quality parameter for the image-based attitude measurements. This quality parameter can be used to determine whether an image-based attitude measurement should be accepted or rejected, limiting the effect of bad registrations due to occlusion and other effects. Next, a method of switching between the image navigation and secondary attitude determination system, based on the availability of registerable imagery should be designed. This will allow the simulation of entire satellite orbits, providing more definitive values for the attitude determination accuracy achievable using an image navigation system.

Improving performance and system properties
It is believed that the real-time attitude determination performance achievable using the designed image navigation system can further be improved by improving the performance of the Kalman filter. Improved filter performance can be achieved by improved tuning of the Kalman filter coefficients. Furthermore, improving the gyroscope bias estimation accuracy will have a positive effect on attitude determination performance.

Optimisation of the image processing code can significantly reduce the computing power needed on board the satellite. This will have a major impact on the feasibility of the image navigation system, as it can significantly reduce its mass and especially power usage. Another method of lowering the image navigation system mass is by integrating the system components in the instrument design. This may allow use of the instrument structure as radiation shielding for the image navigation components. Furthermore, the thermal mass of the instrument may be used to decrease the temperature variations experienced by the gyroscope, lowering the bias drift.

Lastly, methods of using the instrument as attitude sensor during satellite modes other than the observation mode can be investigated. If a coarser instrument-based attitude determination method can be designed for these modes, it might be possible to drop the secondary attitude sensors from the attitude determination system design.

Image navigation for micro and nano satellites
To assess the feasibility of using image navigation for other micro and nano satellite designs, it is important to investigate the image navigation performance for other instrument resolutions. Especially the effect of higher instrument resolutions is of interest, as this might increase the achievable absolute attitude determination accuracy. This might increase the advantages of image navigation over a star-tracker system for higher resolution instruments. Furthermore, a possible image navigation performance increase when using a two-dimensional imager can be investigated.

A

Correlation algorithm trade-off

The full set of correlation algorithms investigated in the course of the literature study is tabulated in Table A.1. Out of the 15 candidates, 4 candidates were deemed unfeasible as they are unable to register images with significant intensity variations.

Table A.1: List of correlation algorithms investigated in literature.

Correlation method	**Feasibility assessment**
Pearson correlation	Feasible
Tanimoto measure	Feasible
Stochastic sign change	*Unable to register images with intensity variations*
Deterministic sign change	Feasible
Incremental sign distance	Feasible
Kendall's Tau	Feasible
L1 norm	*Unable to register images with intensity variations*
Square L2 norm	Feasible
Normalised square L2 norm	Feasible
Median of difference	*Unable to register images with intensity variations*
Minimum ratio	*Unable to register images with intensity variations*
Intensity ratio variance	Feasible
Spearman's Rho	Feasible
Rank distance	Feasible
Mutual information	Feasible

Trade-off
The literature studied provides a great deal of information about the performance of various correlation methods. Various authors provide quantitative values for the performance of a correlation algorithm in combination with a certain set of images. However, this data shows the large dependency of algorithm performance on the specific set of images used. Assessing algorithms based on the qualitative results found by different authors on different image sets therefore does not make for a fair comparison. Instead, only studies comparing multiple registration methods using the same relevant imagery are used. The Pearson correlation coefficient method is used as baseline correlation algorithm, as it is the one most commonly encountered methods in literature. Furthermore, it can be mathematically related to a lot of the other correlation methods described, meaning some relative performance parameters can be mathematically proven.

To aid in this process, a Pugh matrix is used. The baseline algorithm receives a zero, or neutral, score for each of the trade-off criteria. Outperforming of the baseline algorithm can lead to a maximum score of plus two, while under performing can lead to a minimum score of minus two. The scores awarded

in each of the trade-off criteria is multiplied by the criteria's weight factor before being summed to compute the total score of an algorithm.

The trade-off criteria used in both trade-offs are computational cost, resistance to intensity variations, resistance to noise and number of image lines required:

- **Computational cost:** Key parameter in selecting image registration methods is the computational cost of the registration. Although the exact number of computations required for a single registration may vary depending on the exact image used, the order of magnitude for a given image size can easily be determined and is often available in literature.
- **Resistance to intensity variations:** Consecutive images taken over a single location on Earth will be taken at different times of day. This will cause differences in the lighting of these scenes. Registration algorithms will therefore need to be able to cope with intensity variations in the images being registered. These intensity variations are largely linear; for surfaces with constant albedo the intensity is only dependent on the incoming light flux. However, other surfaces exist for which the albedo is dependent on the angle of incidence of the incoming sunlight. This can cause non-linear intensity variations. Registration algorithms capable of dealing with non-linear intensity variations are therefore assumed to outperform those algorithms only capable of dealing with linear variations.
- **Resistance to noise:** Noise is an inherent part of any image captured by real sensor. Although the exact noise level of the SPECTROLITE instrument is still unknown, the signal to noise ratio could be relatively low due to the small instrument size and large distance of the instrument to the scene. This means it is important for the registration algorithms to be able to cope with noisy images.
- **Number of image lines required:** As discussed before, a lower number of pixel lines required in an image to perform a registration, leads to a higher update frequency. This higher update frequency in turn positively influences the accuracy of the location knowledge. Some registration algorithms require a larger number of pixel lines simply because they need more pixels to perform an accurate correlation, while other methods require a specific minimum number of pixel lines to function.

The trade-off criterion values and total scores of the feature detection methods considered are tabulated in Table A.2.

Table A.2: Correlation algorithm trade-off scores.

Correlation method	**Computational cost**	**Resistance to intensity var.**	**Resistance to noise**	**Number of lines required**	**Final score**
Weight	2	2	1	1	
Pearson correlation	0	0	0	0	0
Tanimoto measure	0	0	0	-	-1
Deterministic sign change	0	- -	+	0	-3
Incremental sign distance	0	+ +	-	0	3
Kendall's tau	- -	+	+	+	0
Square L2 norm	+	- -	-	0	-3
Normalised square L2 norm	-	0	0	0	-2
Intensity ratio variance	0	-	+	0	-1
Spearman's rho	- -	+	+	+ +	1
Rank distance	- -	+	+	0	-1
Mutual information	- -	+	-	-	-4

Bibliography

[1] Michael D. Abràmoff, Paulo J. Magalhães, and Sunanda J. Ram. Image processing with imageJ. *Biophotonics International*, 11(7):36–41, 2004.

[2] Olugbenga Ayinde and Yee-hong Yang. Face recognition approach based on rank correlation of Gabor-altered images. *Pattern Recognition*, 35:1275–1289, 2002.

[3] David Bamber, Phil Palmer, and Stephen Mackin. High Performance Attitude Determination through Analysis of Geometric Distortions within Earth Observational Satellite Imagery. *AIAA/USU Conference on Small Satellites*, pages 1–12, 2006.

[4] C.A. Berenstein, L.N. Kanal, D. Lavine, and E. Olson. A geometric approach to subpixel registration accuracy. *Computer Vision, Graphics and Image Processing*, 40:334–360, 1987.

[5] R.E. Bird, R.L. Hulstrom, and L.J. Lewis. Terrestrial Solar Spectral Data Sets. *Solar Energy*, 30: 563, 1983.

[6] Glenn D. Boreman. *Modulation Transfer Function in Optical and Electro-Optical Systems*. SPIE Press, 2001. ISBN 0-8194-4143-0.

[7] Kieran a Carroll, Slavek Rucinski, and Robert E Zee. Arc-Minute Nanosatellite Attitude Control : Enabling Technology for the BRITE Stellar Photometry Mission. *Proceedings of the AIAA/USU Conference on Small Satellites*, 2004.

[8] Subhash Challa, Robin J. Evans, and Xuezhi Wang. A Bayesian solution and its approximations to out-of-sequence measurement problems. *Information Fusion*, 4(3):185–199, 2003.

[9] Richard W. Conners and Charles A. Harlow. A theoretical comparison of texture algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2(3):204–222, 1980.

[10] Carlo de Franchis, Enric Meinhardt-Llopis, Daniel Greslou, and Gabriele Facciolo. Attitude Refinement for Orbiting Pushbroom Cameras: a Simple Polynomial Fitting Method. *Image Processing On Line*, 5:328–361, 2015.

[11] Francesco Dell'Endice, Jens Nieke, Benjamin Koetz, Michael E. Schaepman, and Klaus Itten. Improving radiometry of imaging spectrometers by using programmable spectral regions of interest. *ISPRS Journal of Photogrammetry and Remote Sensing*, 64(6):632–639, 2009.

[12] Ronald G. Driggers. *Encyclopedia of Optical Engineering*. CRC Press, 2003. ISBN 0-8247-0940-3.

[13] ESA. Agreement Between The Netherlands And ESA Signed For Sentinel-5 Precursor Instrument, 2009. URL `http://www.esa.int/Our_Activities/Observing_the_Earth/Copernicus/Agreement_between_the_Netherlands_and_ESA_signed_for_Sentinel-5_Precursor_instrument`.

[14] Lu Feng, W Paul Menzel, Christopher S Velden, Xu Jianmin, Steve Wanzong, and Li Jun. Attempts to Improve GOES Image Navigation. In *8th International Winds Workshop*, 2006.

[15] E Gill, O Montenbruck, K Arichandran, SH Tan, and others. High-precision onboard orbit determination for small satellites-the GPS-based XNSon X-SAT. In *Small Satellites, Systems and Services*, volume 571, page 47, 2004.

[16] A. Ardeshir Goshtasby. *Image Registration: Principles, Tools and Methods*. Springer-Verlag London, 1 edition, 2012. ISBN 978-1-4471-2457-3.

[17] Rajiv Gupta. Linear pushbroom cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9):963–975, 1997.

[18] Yuichi S Hayakawa, Takashi Oguchi, and Zhou Lin. Comparison of new and existing global digital elevation models: ASTER G-DEM and SRTM-3. *Geophysical Research Letters*, 35(17):1–5, 2008.

[19] International Organization for Standardization. ISO 5725-1: Accuracy (trueness and precision) of measurement methods and results. Technical Report C, 2012.

[20] International Organization for Standardization. ISO 12233-2017: Slanted-edge method. Technical report, 2017.

[21] K Janschek and S Dyblenko. Satellite Autonomous Navigation Based on Image Motion Analysis. *15th IFAC Symposium on Automatic Control in Aerospace*, 2001.

[22] K Janschek, S Dyblenko, and T Boge. Image Based Attitude Determination Using an Optical Correlator. In *ESA International Conference on Spacecraft Guidance, Navigation and Control Systems*, number October, pages 18–21, 1999. ISBN 4935146370.

[23] B. Johnston-Lemke, K. Sarda, C. C. Grant, and R. E. Zee. Arc-Minute Attitude Stability on a Nanosatellite: Enabling Stellar Photometry on the Smallest Scale. *Proceedings of the 25th Annual Small Satellite Conference*, pages 1–13, 2011.

[24] Maurice G Kendall. A New Measure of Rank Correlation. *Biometrika*, 30(1):81–93, 1938.

[25] M.G. Kendall and J.D. Gibbons. *Rank correlation methods*. Oxford University Press, 5th edition, 1948.

[26] Edward J. Knight and Geir Kvaran. Landsat-8 operational land imager design, characterization and performance. *International Symposium on Remote Sensing*, 6(11):10286–10305, 2014.

[27] W. H. Kruskal. Ordinal measures of association. *Journal of the American Statistical Association*, 53(298):814–861, 1958.

[28] T.D. Larsen, N.A. Andersen, O. Ravn, and N.K. Poulsen. Incorporation of time delayed measurements in a discrete-time Kalman filter. In *IEEE Conference on Decision and Control*, volume 4, pages 3972–3977, 1998.

[29] Wiley J Larson and James R Wertz. *Space Mission Analysis and Design*. Microcosm, 4 edition, 2004. ISBN 1881883108.

[30] R.A. Masterson, D.W. Miller, and R.L. Grogan. Development and Validation of Reaction Wheel Disturbance Models: Empirical Model. *Journal of Sound and Vibration*, 249(3):575–598, 2002.

[31] Jacqueline Le Moigne and Nathan S Netanyahu. *Image Registration for Remote Sensing*. Cambridge University Press, 2011.

[32] Ichiro Murase and Satoru Igarashi. Robust Image Registration by Increment Sign Correlation. *Pattern Recognition*, 35:2223–2234, 2002.

[33] K. Pearson. Contributions to the Mathematical Theory of Evolution. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 185:71–110, 1894.

[34] U Platt and D Perner. Measurements of Atmospheric Trace Gases by Long Path Differential UV/Visible Absorption Spectroscopy. In Dennis K Killinger and Aram Mooradian, editors, *Optical and Laser Remote Sensing*, pages 97–105. Springer Berlin Heidelberg, 1983. ISBN 978-3-540-39552-2.

[35] Bernard Polle, Julien Morand, and EADS Astrium. Study of new architecture needs for AOCS/Avionics. *ESA General Studies Program*, 2006.

[36] Karan Sarda, C. Cordell Grant, Monica Chaumont, Seung Yun Choi, Bryan Johnston-Lemke, and Robert E. Zee. On-Orbit Performance of the Bright Target Explorer (BRITE) Nanosatellite Astronomy Constellation. *28th Annual AIAA/USU Conference on Small Satellites*, 2014.

[37] H Shekarforoush, M. Berthod, and J. Zerubia. Subpixel Image Registration by Estimating the Polyphase Decomposition of Cross Power Spectrum. In *Proceedings of the Computer Society Conference on Computer Vision and Pattern Recognition*, pages 532–537, 2004.

[38] Grace S. Shieh. A Weighted Kendall's Tau Statistic. *Statistics & Probability Letters*, 39(1):17–24, 1998.

[39] C. Spearman, C Spearman Con, and Ntents Page. The Proof and Measurement of Association Between Two Things. *The American Journal of Psychology*, 15(1):72–101, 1904.

[40] John C. Springmann, James W Cutler, and H Bahcivan. Magnetic Sensor Calibration and Residual Dipole Characterization for Application to Nanosatellites. *AIAA/AAS Astrodynamics Specialist Conference*, pages 1–14, 2010.

[41] James C. Storey. Landsat 7 on-orbit modulation transfer function estimation. *International Symposium on Remote Sensing*, 4540:50–61, 2001.

[42] Tetushi Tachikawa, Manabu Kaku, Akira Iwasaki, Dean B Gesch, Michael J Oimoen, Z Zhang, Jeffrey J Danielson, Tabatha Krieger, Bill Curtis, Jeff Haase, Michael Abrams, and C Carabajal. ASTER Global Digital Elevation Model Version 2 - summary of validation results. Technical report, NASA, 2011.

[43] Qi Tian and Michael N. Huhns. Algorithms for Subpixel Registration. *Computer Vision and Pattern Recognition*, 233:220–233, 1986.

[44] Thierry Toutin. Three-dimensional topographic mapping with ASTER stereo data in rugged topography. *IEEE Transactions on Geoscience and Remote Sensing*, 40(10):2241–2247, 2002.

[45] B Udrea, M. Nayak, and F. Ankersen. Analysis of the Pointing Accuracy of a 6U CubeSat Mission for Proximity Operations and Resident Space Object Imaging. In *5th International Conference on Spacecraft Formation Flying Missions and Technologies*, 2013.

[46] United States Geological Survey. High Plains Aquifer Figures. URL `https://pubs.usgs.gov/ha/ha730/ch_c/C-text5.html`.

[47] United States Geological Survey. Irrigation water use: Center-pivot irrigation, 2016. URL `https://water.usgs.gov/edu/gallery/wuir-centerpivot-aerial.html`.

[48] J.P. Veefkind, I. Aben, K. McMullan, H. Förster, J. de Vries, G. Otter, J. Claas, H.J. Eskes, J.F. de Haan, Q. Kleipool, M. van Weele, O. Hasekamp, R. Hoogeveen, J. Landgraf, R. Snel, P. Tol, P. Ingmann, R. Voors, B. Kruizinga, R. Vink, H. Visser, and P.F. Levelt. TROPOMI on the ESA Sentinel-5 Precursor: A GMES mission for global observations of the atmospheric composition for climate, air quality and ozone layer applications. *Remote Sensing of Environment*, 120:70–83, 5 2012.

[49] J. de Vries;, R. Hoogeveen, R. Voors, Q. Kleipool, P. Veefkind, I. Aben, R. Snel, N. van der Valk, H. Visser, and G. Otter. Technology evolution of the TROPOMI instrument, 2012.

[50] Jeffrey P Walker and Garry R Willgoose. A comparative study of Australian cartometric and photogrammetric digital elevation model accuracy. *Photogrammetric Engineering and Remote Sensing*, 72(7):771–779, 2006.

[51] Zhipeng Wang, Xiaoxiong Xiong, Taeyoung Choi, and Daniel Link. On-orbit characterization of MODIS modulation transfer function using the Moon. *IEEE Transactions on Geoscience and Remote Sensing*, 52(7):4112–4121, 2014.

[52] Medha V Wyawahare, Pradeep M Patil, and Hemant K Abhyankar. Image Registration Techniques : An overview. *International Journal of Signal Processing, Image Processing and Pattern Recognition*, 2(3):11–28, 2009.

[53] X. Xiong, K. Chiang, J. Sun, W. L. Barnes, B. Guenther, and V. V. Salomonson. NASA EOS Terra and Aqua MODIS on-orbit performance. *Advances in Space Research*, 43(3):413–422, 2009.

[54] L Yang. FY-2 Automatic Landmark Positioning for Image Navigation and its Applications in FY-2D VISSR imagery. In *International geoscience and remote sensing symposium*, pages 2–5, 2010.

[55] Huanshui Zhang, Xiao Lu, Weihai Zhang, and Wei Wang. Kalman Filtering for Linear Time-Delayed Continuous-Time Systems with Stochastic Multiplicative Noises. *International Journal of Control, Automation and Systems*, 5(4):355–363, 2007.