Optical Attitude Estimation System with Low-Resolution Infrared Imaging

Low cost, miniaturized attitude sensing for microsatellites

Niklas Knöll

Delft University of Technology

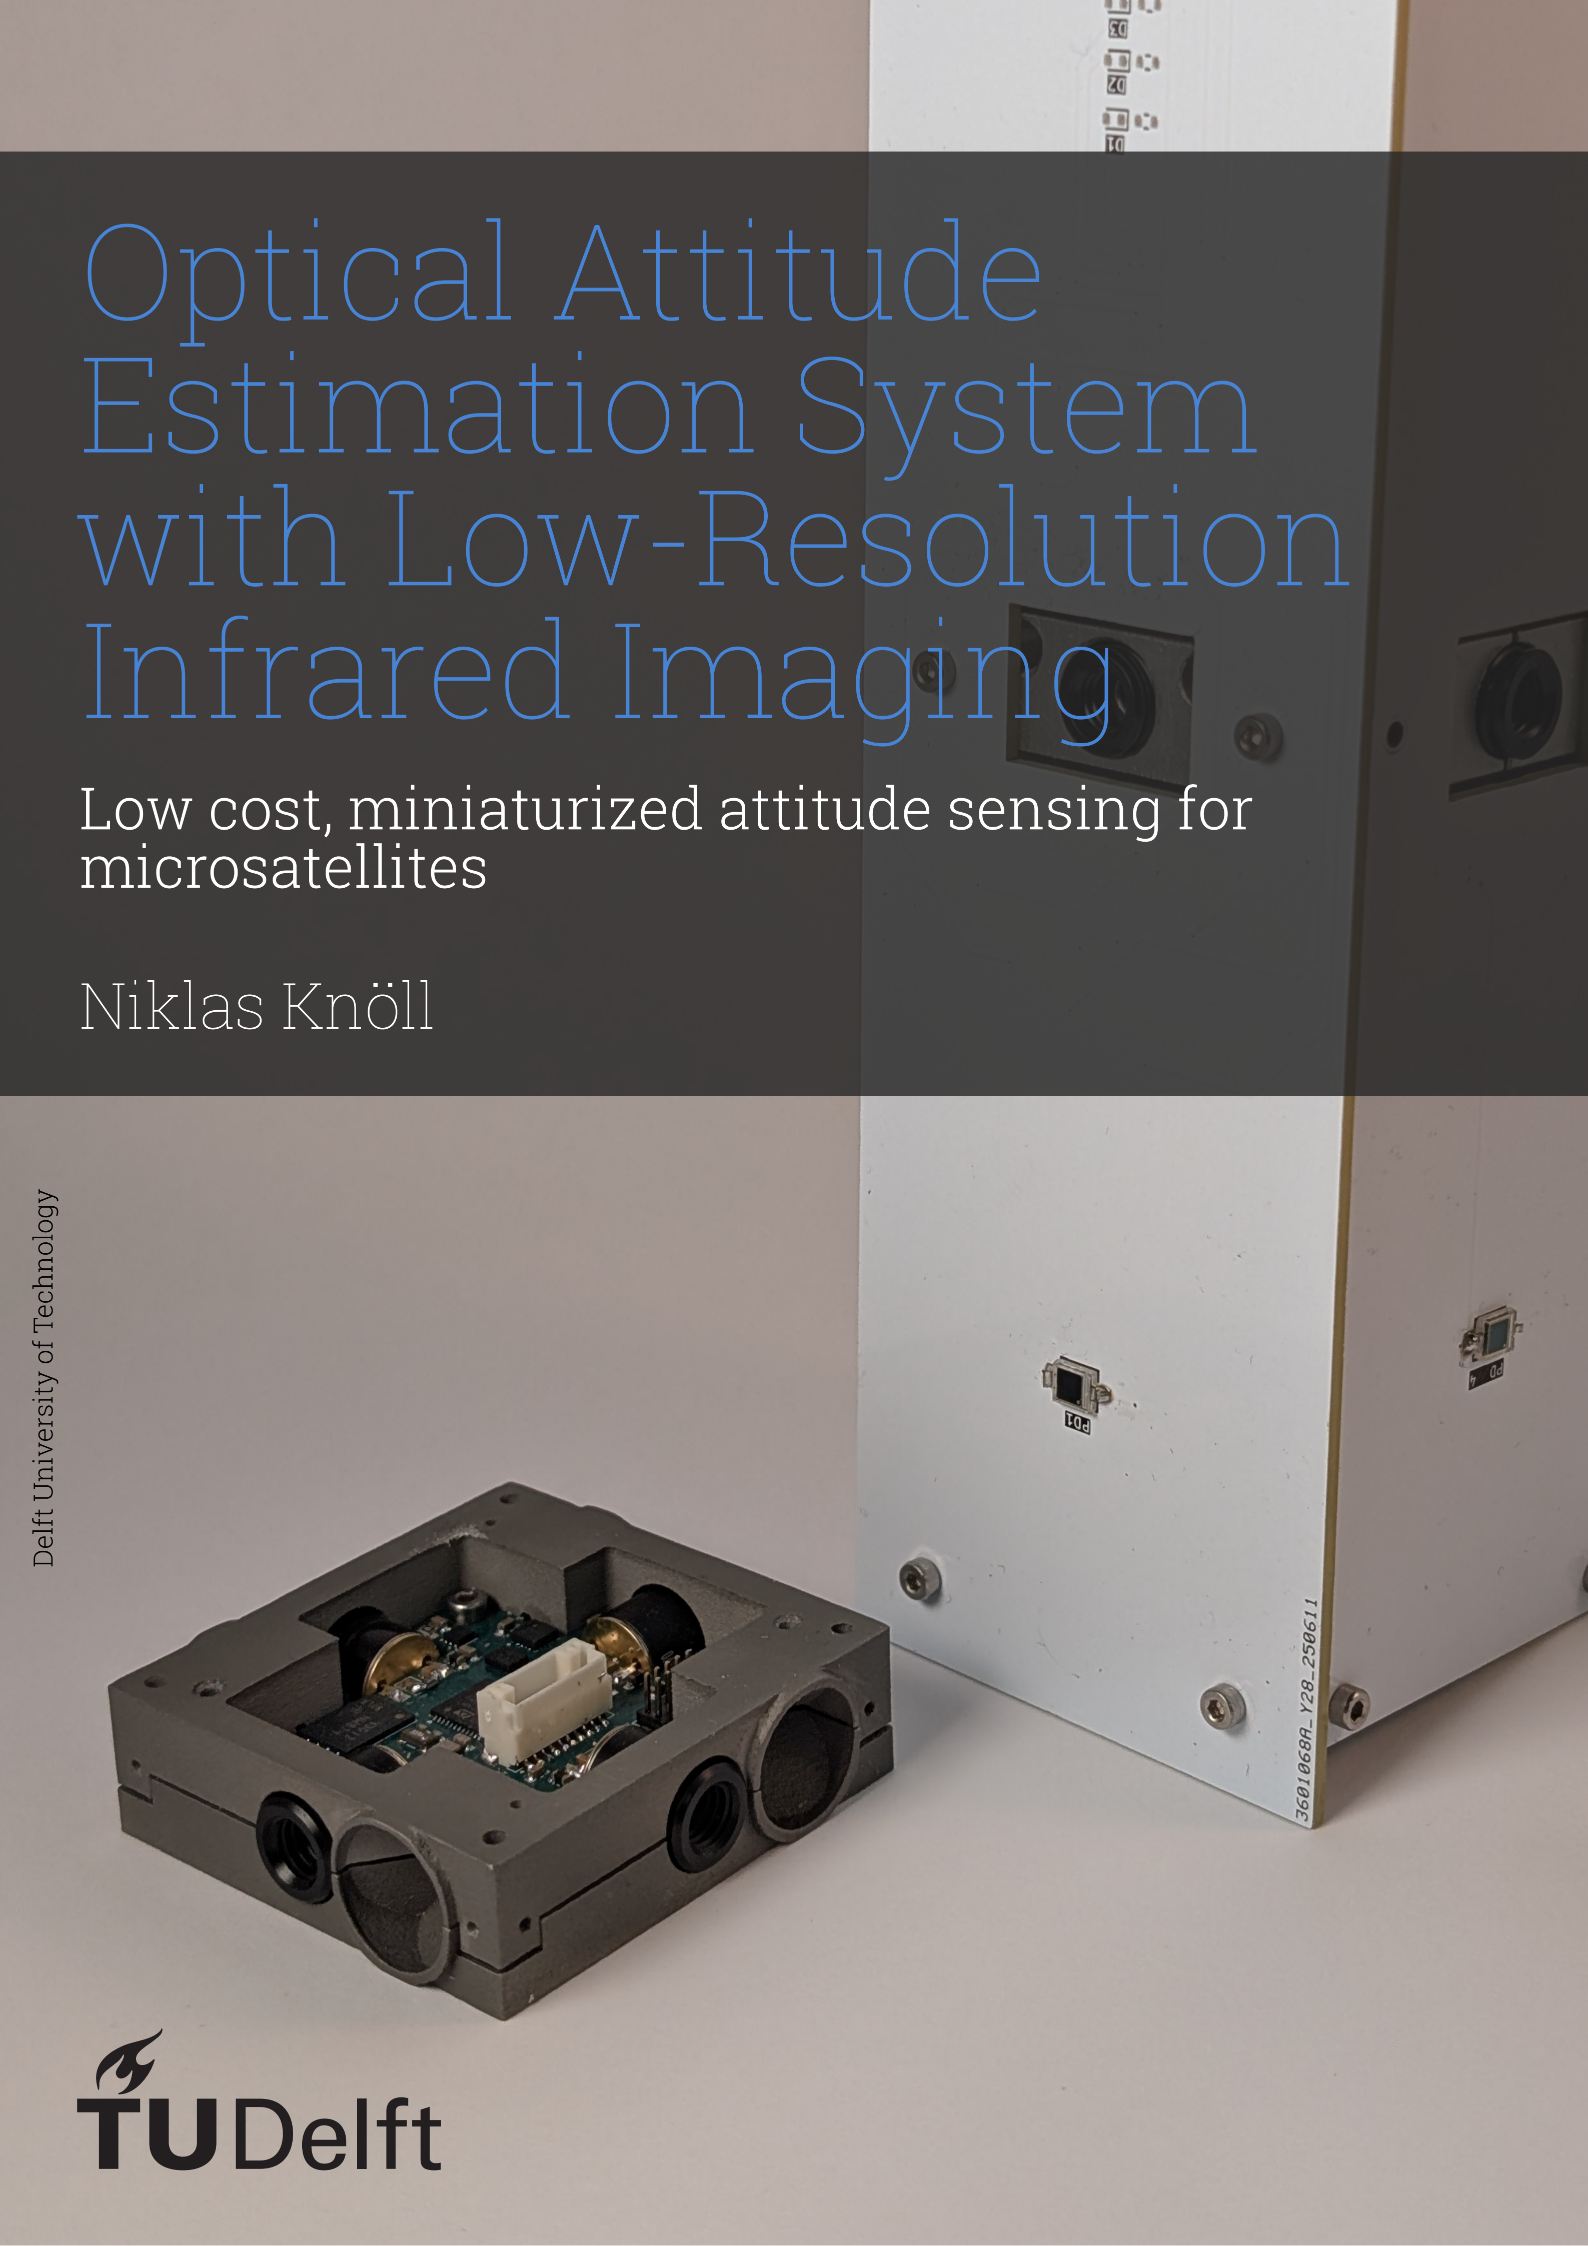

TU Delft

Optical Attitude Estimation System with Low-Resolution Infrared Imaging

Low cost, miniaturized attitude sensing for microsatellites

by

Niklas Knöll

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Monday March 23rd, 2026 at 10:00 AM.

Student number:	5006961	
Project duration:	April 2025 – March 2026	
Thesis committee:	Dr. S. Speretta,	TU Delft, supervisor
	Dr. ir. J. Bouwmeester,	TU Delft
	Dr. M.S. Uludağ	TU Delft

Cover:	Own work
Style:	TU Delft Report Style, with modifications by Daan Zwaneveld

Preface

DARE has strongly shaped my time at TU Delft. It influenced both my academic direction and my professional path, and introduced me to many people who have since become close friends. It is also where the foundation of this thesis was laid.

One of the projects I worked on in DARE was SHEAR, a heat shield experiment launched aboard REXUS, part of the ESA/DLR student programme that enables experiments on sounding rockets and stratospheric balloons. The launch campaign was an intense and formative experience, one of many during my time at DARE that shaped how I approach engineering and teamwork. I am grateful to DARE for providing an environment in which students are entrusted with ambitious projects and real responsibility.

After returning from Kiruna, I visited Stefano to thank him for his support and to discuss possibilities for my thesis. The first ideas for this work began to take shape during that conversation. Earlier, during the REXUS launch campaign, Esmée (who had started SHEAR and is now the manager of the REXUS/BEXUS programme) had encouraged our team to consider the balloon branch of the programme. That suggestion stayed with me. Back in Delft, I explored with Stefano how such an opportunity could align with my academic interests. His enthusiasm and the topic he proposed matched perfectly with the direction I hoped to pursue. Thank you for your guidance, trust, and steady support throughout this thesis.

Although the balloon flight ultimately crashed, working on SHAREon was an invaluable experience. This thesis would not have been possible without the team behind it. Jeff, Jonas, and Kristina — thank you for building this project together and for the experiences we shared along the way.

I would also like to thank the organisers of the REXUS/BEXUS programme for creating opportunities that allow students to take on real responsibility, especially Esmée for motivating us to apply again and Simon for supporting us through the programme.

I am also grateful to my colleagues at Dawn Aerospace. After my internship, I continued working part-time while completing my thesis. Contributing to hardware and software that is now flying in space was both motivating and grounding, and occasionally provided a welcome distraction from the inevitable challenges of this thesis.

Finally, I thank my family, who have supported me throughout my journey at TU Delft and beyond. Whenever I had doubts or needed motivation, I could turn to you for support and encouragement. Without you, I would never have had the opportunity or the courage to follow this path.

Niklas Knöll
Delft, March 2026

Summary

Accurate attitude determination is a fundamental requirement for spacecraft operation, enabling power generation, communication, payload pointing, and mission safety. Conventional attitude determination systems rely on star trackers, Sun sensors, magnetometers, and inertial measurement units. While effective, star trackers can be costly, computationally demanding, and sensitive to lighting or contamination. Compact infrared Infrared (IR) imaging sensors open the possibility for alternative optical sensors using naturally available references such as the Earth horizon.

This thesis investigates attitude determination for microsatellites using low-resolution infrared cameras. The goal is to assess whether a compact IR horizon sensor can provide a meaningful estimate of spacecraft orientation by observing the Earth's limb. The work evaluates the achievable accuracy, the main sources of error, and the practical limitations of such a system. The approach is based on detecting the Earth horizon and the Sun in infrared images and relating their observed geometry to the spacecraft attitude. In low Earth orbit, the projected Earth limb forms a conic section in the image plane. By extracting the edge points of the horizon and fitting an appropriate geometric model, the spacecraft's orientation with respect to the Earth can be reconstructed. When the Sun is visible, it provides an additional reference that further constrains the attitude solution.

The horizon and Sun measurements are fused in an Extended Kalman Filter (EKF) with gyro data to obtain a continuous attitude estimate. Reliable convergence requires initialization with two non-collinear reference vectors. In practice, the filter is initialized using any available optical vector (horizon normal or Sun direction) together with a magnetometer measurement, which is always available, albeit less accurate. This prevents convergence to an attitude solution rotated by 180 degrees, an ambiguity arising from the linearized measurement model.

To test the algorithms under controlled, repeatable conditions, a synthetic image generator was developed in conjunction with a simulation of the dynamics of a representative spacecraft. The simulator produces infrared horizon images for different spacecraft attitudes and orbital positions, while incorporating relevant sensor characteristics, including Field-of-View (FoV), resolution, and noise. This environment enables direct comparison between the estimated and true attitude, allowing evaluation of systematic effects and sensitivity to disturbances. Verification of the sensor concept was planned through a high-altitude balloon flight within the DLR/SNSA/ESA programme BEXUS to obtain real horizon imagery under near-space conditions. For this, an entire experimental setup to acquire this data was designed, built and verified. Due to a launch system failure, flight data could not yet be acquired, and the campaign will be repeated later. Nevertheless, the complete sensor system, including hardware and processing chain, was validated in laboratory conditions.

Simulation results show that, under realistic assumptions for a low-resolution infrared camera, attitude accuracies on the order of 2 degrees are achievable when a sufficient portion of the horizon is visible together with observations of the Sun. The accuracy degrades predictably in eclipse when only the horizon can be observed, and this degradation increases with reduced horizon visibility, which occurs at lower altitudes. Monte Carlo analysis shows that the expected accuracy does not degrade beyond approximately 20 degrees under worst-case conditions in eclipse. The results demonstrate that a compact infrared horizon sensor can provide meaningful attitude information for microsatellites and may serve as a complementary sensor in missions where cost, size, and power constraints exclude more complex systems.

In conclusion, a novel, fully integrated and miniaturized attitude determination system, dependent solely on small infrared imaging sensors, was developed that, unlike traditional small-satellite attitude determination systems, once initialized, can operate completely independently of magnetometers. In the future, this could enable simplified control using magnetorquers, as the system will not be disturbed by magnetic-field changes.

Contents

List of Figures

List of Tables

Acronyms

ADCS Attitude Determination and Control System 1, 5
AMR Anisotropic Magnetoresistance 6

BEXUS Balloon Experiments for University Students 48

CNN Convolutional Neural Networks 9, 12–15
COTS Commercial off-the-shelf 1, 7

DLR German Aerospace Center 48

EAH-Jena Ernst Abbe Hochschule Jena 48
EKF Extended Kalman Filter ii, 17, 45, 46
ERFA Essential Routines for Fundamental Astronomy 35

FoV Field-of-View ii, 6, 8, 10, 12, 22–25, 27, 29, 32, 67, 73

GNSS Global Navigation Satellite System 15, 42, 44, 56

IMU Inertial Measurement Unit 47, 50, 55
IoT Internet of Things 23
IR Infrared ii, viii, 19–22, 25, 26, 31, 34, 36
I^{2}C Inter-Integrated Circuit 22, 25, 70, 71

MEKF Multiplicative Extended Kalman Filter 18, 19, 42
MEMS Micro-Electromechanical Systems 8, 20, 25, 43, 47

NETD Noise-Equivalent Temperature Difference 24, 25
NIS Normalised Innovation Squared 46

OEM Original Equipment Manufacturer 23, 24

PCB Printed Circuit Board vi, 24, 25, 54–56

RAM Random Access Memory 55
RANSAC Random Sample Consensus 12, 14, 40, 41
REXUS/BEXUS Rocket/Balloon Experiments for University Students 21, 48
RMS Root mean square 36
ROM Read-Only Memory 55

SDK Software Development Kit 24
SHAREon Sensors for High Altitude Research Experiment – Optical Navigation 48
SPI Serial Peripheral Interface 24, 25

SSC Swedish Space Corporation 48, 50, 51

UART Universal Asynchronous Receiver Transmitter 23, 25
UAV Unmanned Aerial Vehicle 23
UKF Unscented Kalman Filter 17
USB Universal Serial Bus 23–25

WMM World Magnetic Model 6

1

Introduction

The rapid growth of the space sector, driven primarily by the deployment of large satellite constellations and the increasing accessibility of launch services, has intensified the demand for reliable, low-cost spacecraft subsystems. This trend is particularly evident in the small satellite domain, where CubeSat and PocketQube missions increasingly pursue more ambitious scientific and technological objectives. In this context, Attitude Determination and Control System (ADCS) play a critical role in mission success, enabling functions such as Earth observation or antenna pointing.

Within many educational and technology-demonstration programs, including the Delfi satellite program [1], a large fraction of CubeSat- and PocketQube-class missions operate without a dedicated attitude determination system [1]. Instead, these spacecraft often rely on passive stabilization methods, with no attitude knowledge at all, or only attitude-rate knowledge from inertial sensors. While such approaches are sufficient for basic missions, a growing subset of small satellites requires at least coarse attitude determination to enable payload operation, to ensure safe pointing of antennas or sensors, or to validate new technologies. Providing this capability without significantly increasing system complexity remains a key challenge. This is the case with the upcoming Delfi Twin satellite mission, which requires coarse attitude knowledge to allow for safe separation of the two twin satellites to remain in the correct orbits where the intended formation flying mission can be achieved.

Traditionally, attitude determination on small spacecraft relies on combinations of inertial and non-imaging reference sensors, such as gyroscopes, magnetometers, coarse Sun sensors, and accelerometers. These sensors are low-power, widely available as Commercial off-the-shelf (COTS) components, and well understood, making them attractive for resource-constrained platforms. However, even COTS implementations typically require multiple discrete sensor units, each with associated electronics, mounting requirements, and calibration considerations. In the tightly constrained volumes of CubeSat- and especially PocketQube-class spacecraft, this distributed sensor architecture complicates mechanical integration, wiring, and system-level design.

In addition, single non-imaging sensors provide limited direct observability of the complete attitude and/or are sensitive to the local environment in the satellite like magnetometers. Magnetic field disturbances, residual magnetic moments, and eclipse-induced loss of Sun observations can significantly degrade attitude knowledge. As a result, achievable attitude accuracy is often modest, and full observability cannot be guaranteed throughout all mission phases.

Optical attitude sensors offer an alternative by directly observing external reference objects. Star trackers provide the highest accuracy, but their mass, power consumption, cost, and sensitivity to bright objects typically preclude their use on CubeSat- and PocketQube-class spacecraft. Horizon and Sun sensors are simpler optical solutions; however, classical implementations often rely on dedicated optics, scanning mechanisms, or relatively high-resolution detectors, which limits their scalability to very small platforms.

---

[1] https://www.tudelft.nl/lr/delfi-space/delfi-program

Recent advances in uncooled infrared microbolometer technology have enabled compact, low-power thermal imagers with resolutions on the order of tens of pixels and prices below 50€. Despite their limited spatial resolution, these sensors can exploit the strong thermal contrast between the Earth and deep space to detect the Earth horizon, providing direct geometric information about spacecraft orientation in low Earth orbit. By combining multiple reference cues within a single imaging system, such sensors offer the potential to replace or augment several traditional attitude sensors with a more tightly integrated solution. Unlike visible-light sensors, infrared horizon sensors operate independently of solar illumination conditions and therefore provide continuous measurements during both sunlight and eclipse, a significant advantage for small satellites with limited sensor redundancy.

Furthermore, imaging-based sensors inherently offer access to multiple reference cues within a single measurement. Depending on the spacecraft orientation and orbital geometry, a single infrared image may contain information about the Earth's horizon as well as bright celestial bodies such as the Sun or Moon. This contrasts with traditional non-imaging sensors, which typically measure only a single physical quantity and must be combined to achieve observability. Imaging sensors, therefore, present an opportunity to increase attitude knowledge without a proportional increase in sensor count or system complexity.

Despite these advantages, the use of extremely low-resolution infrared imagery for attitude determination presents significant challenges. Classical image-processing and geometric reconstruction techniques were largely developed for higher-resolution sensors, and their applicability at very low pixel counts remains unestablished. Moreover, the performance of such systems is strongly influenced by sensor noise, orbital altitude, variability in thermal contrast, and partial visibility of reference objects.

This thesis investigates the feasibility of optical attitude determination using multiple low-resolution infrared cameras for small satellite applications. The focus lies on Earth-horizon-based attitude estimation, complemented by observations of the Sun where available, using physically motivated image-based processing techniques. Particular attention is given to configurations and algorithms suitable for CubeSat- and PocketQube-class platforms, where computational resources, power consumption, and system complexity are tightly constrained.

To support this investigation, a complete simulation framework is developed to generate synthetic infrared images for arbitrary spacecraft orientations and orbital states. This framework enables systematic evaluation of attitude determination performance under varying environmental and sensor conditions. Based on these simulations, an attitude estimation approach is implemented and assessed in terms of accuracy, robustness, and practical feasibility for embedded spacecraft systems.

The structure of this thesis is as follows. Chapter 2 reviews the fundamentals of spacecraft attitude determination and surveys existing optical and non-optical sensor approaches relevant to small satellites. Subsequently, the resulting research questions and the intended planning are presented in Chapter 3. Chapter 4 describes the implementation of the infrared sensor models, a synthetic image generation pipeline, and the algorithms used to perform the attitude estimation, as well as the work performed to verify these. Chapter 5 presents and discusses the results of the simulations, evaluates the expected performance, and examines the variables that affect it. Finally, Chapter 6 summarizes the main conclusions of this study and outlines recommendations for future work.

2

Context and Background of Microsatellite Attitude Determination

The rapid advancement of space technology has opened new frontiers in scientific research, Earth observation, and communication. Among the pioneering initiatives in this field is the Delfi Program at the Delft University of Technology (TU Delft), which has been at the forefront of developing nano- and picosatellites since 2004 [2]. This program exemplifies the potential of small satellites to revolutionize space missions, offering cost-effective and innovative solutions for technology demonstration and operational applications.

The Delfi Program was initiated with the goal of advancing space engineering education and technological innovation by developing very small satellites. Over the years, it has evolved into a comprehensive platform that not only educates future space engineers but also pushes the boundaries of satellite miniaturization and functionality. The program's success is evident in its series of satellite launches, each building upon the previous one to demonstrate new technologies and capabilities.

This literature study aims to provide a detailed context for the Delfi Program, highlighting its historical background, key achievements, and future directions. By examining the program's evolution, we can understand its impact on space engineering education, technological advancements, and the broader field of satellite development. This context will serve as a foundation for further analysis and discussion in the thesis.

2.1. Context of the Delfi Program

The Delfi Program, initiated in 2004 at the Space Engineering Department of the Delft University of Technology (TU Delft), represents a pioneering effort in the development and deployment of nano- and picosatellites. The program's primary objectives are to advance space engineering education, test and qualify novel space technologies, and enhance the nano-satellite platform to enable new applications in space exploration and Earth observation [3, 4].

2.1.1. Historical Background and Evolution

The Delfi Program began at a time when CubeSats, small 10x10x10 cm-sized satellites [5], were primarily used by a select group of universities for educational purposes. TU Delft sought to demonstrate that these small satellites could also serve as platforms for innovative technology and operational missions [3]. The program has since evolved to include not only CubeSats but also PocketQubes and picosatellites, pushing the boundaries of miniaturization and cost-effectiveness in space technology [3, 6].

2.1.2. Key Satellites and Achievements

- **Delfi-C3:** Launched in 2008, Delfi-C3 was the first Dutch university satellite and remained operational over 12.5 years after its launch, far exceeding its original one-year mission. It demonstrated the feasibility of CubeSats for technology demonstration and educational purposes, and its success inspired both academic and commercial initiatives in the Netherlands and beyond [2, 7]. Shown in Figure 2.1.

- **Delfi-n3Xt:** Launched in 2013, this successor to Delfi-C3 introduced active position control, allowing for more precise orientation and operations. Despite a temporary loss of contact, the satellite was successfully reactivated after seven years, showcasing the robustness of TU Delft's satellite designs [8, 9]. Shown in Figure 2.2.
- **Delfi-PQ:** Representing the next frontier in miniaturization, Delfi-PQ is a PocketQube, a new picosatellite standard with an even smaller base size of 5x5x5cm compared to the previous CubeSats, with a volume eight times smaller than its CubeSat predecessors. Weighing only 600 grams, it challenges conventional satellite design and demonstrates the potential for even smaller satellites to perform meaningful missions [6, 10]. It was launched in January 2022 and reentered in January 2024. Shown in Figure 2.3.

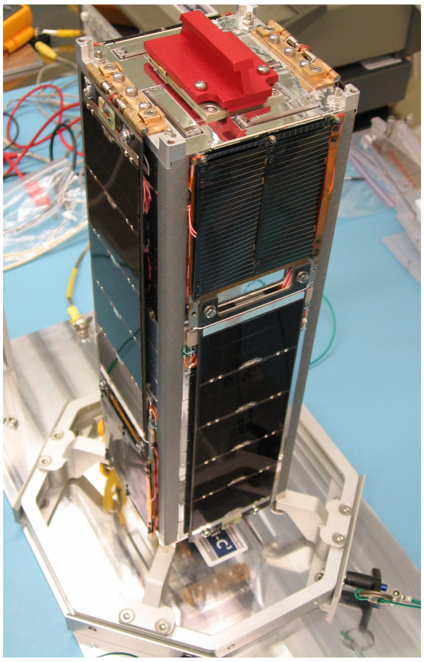
Figure 2.1: Delfi C3 [11]

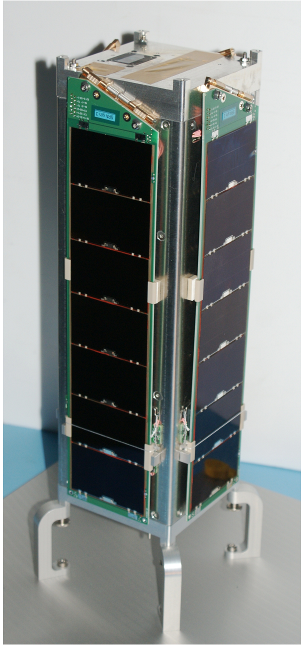
Figure 2.2: Delfi N3XT [12]


Figure 2.3: Delfi PQ [13]

2.1.3. Future Directions

The program continues to innovate, with plans to launch new satellites. Each mission builds on the previous one, incorporating lessons learned and new technologies. The focus on miniaturization, concurrent engineering, and rapid development cycles continues with Delfi Twin [14, 4].

The upcoming Delfi Twin mission aims to perform the first formation flying mission for both the Delfi Program and PocketQubes. The mission aims to launch two nearly identical satellites that are separated from the launch vehicle while still attached to one another. This coupled configuration will then be commissioned on orbit, and once it is confirmed that all systems are operational, the separation will be triggered. Once separated, the two satellites control their relative position by performing drag modulation with deployable, controllable solar flaps [15]. This separation is a critical part of the satellite's mission, as it must be performed at the correct attitude to maintain orbits close enough to enable drag modulation to control the separation between the satellites. This concept is shown in Figure 2.4.

2.2. Requirements for attitude sensing

Determining the attitude of a satellite can be critical to the functioning of a satellite. There are many applications where a known and stabilized attitude is required, for example, for imaging payloads to gather data on the desired targets, antennas to be pointed to the correct ground station, Sun pointing

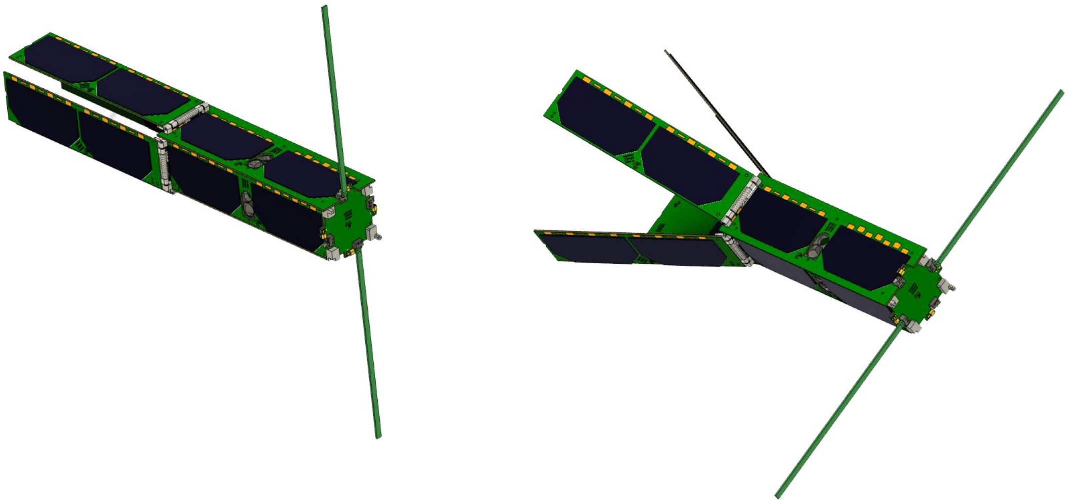

Figure 2.4: Delfi Twin concept with variable drag area. On the left, the solar panel flaps are entirely retracted, leading to the minimum drag, while on the right, they are deployed, generating increased drag [16].

for power generation or thermal reasons, and many more.

When CubeSats were first developed, many did not have any attitude control or even determination. However, in recent years, with the commercialisation of CubeSats, this has changed, as they are used for more and more applications that previously used only larger satellites [17].

For the smallest satellite segments, such as PocketQubes, a similar development is underway. A large part of the satellites in this size class are launched with purely passive means of attitude determination and control, or very rudimentary forms of it. One such example is the POQUITO satellite, which has active ADCS that only determines and controls tumbling rates using gyroscope, magnetometer, and magnetorquers [18]. Only a few more advanced ADCS using reaction wheels exist on board previously launched PocketQubes, such as AlbaOrbital's Unicorn 2 [1].

This is in large part due to the miniaturisation required, which makes the implementation of such systems challenging in the very limited volume. Combined with the mission requirements for these typically academic missions, this leads to very few highly developed attitude determination systems on board previous PocketQubes.

For the Delfi Twin mission, separation is required to occur within a bounded attitude range around the prograde or retrograde orbital direction. This constraint ensures that the post-separation trajectories of the two satellites do not interfere with one another, thereby avoiding collision risk or significant changes in orbital elements that would compromise the planned drag-modulation re-approach strategy [15].

Due to residual angular rates imparted by the launch vehicle and deployer, the spacecraft will naturally rotate after release. As a result, its attitude will periodically align with the required separation direction without requiring active three-axis control. However, accurate three-axis attitude determination remains essential to detect when the spacecraft orientation falls within the allowable separation region. This region is defined as a conical bound with a half-angle of 52.7°, illustrated in Figure 2.5 [19]. This geometric constraint directly drives the performance requirements of the attitude determination system. In particular, the estimation accuracy must be sufficiently high to ensure, with adequate confidence, that the true spacecraft attitude lies inside the prescribed cone at the time of separation. Any uncertainty in attitude knowledge effectively reduces the usable separation region, as a safety margin must be introduced to account for worst-case estimation errors. Consequently, improved attitude accuracy directly enlarges the effective operational cone, within which separation can be commanded while still guaranteeing compliance with the 52.7° requirement, thereby increasing separation opportunities and operational robustness.

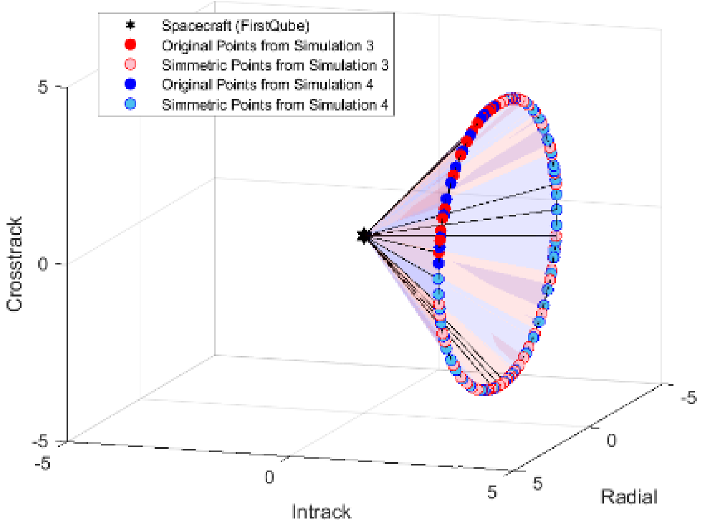

Figure 2.5: Half-cone within which the Delfi Twin satellites must separate for mission success [19].

2.3. Typical Attitude Sensors

Traditional techniques for attitude estimation typically use a combination of different sensors that measure both the absolute attitude and the rotation rates. At least two independent external vectors are required to be able to estimate the full attitude, which makes a combination of different sensors necessary. The most common of these sensors are the following [20]:

- **Sun sensors**: These measure the direction of incoming light from the Sun to determine the attitude relative to the Sun. Two types of Sun sensors are typically distinguished: Coarse Sun sensors are either a number of photodiodes arranged at various angles to measure the intensity of the Sun, while fine Sun sensors use either imaging sensors or linear photodiodes in combination with pinholes or slits.
- **Horizon sensors**: Horizon sensors detect the edge of the Earth, taking advantage of the large size of the Earth in a typical satellite's FoV. They also traditionally come in two types: static sensors with a number of simple sensing elements, but limited to smaller fields of view. Scanning sensors sweep a detector across the Earth, detecting the point at which the transition from the Earth to space occurs and provide a broader view. Horizon sensors typically use the infrared band due to better horizon definition in those wavelengths [20] and the ability to operate without sunlight.
- **Star sensors**: Star sensors (or star trackers) provide precise attitude determination by imaging the star field and comparing observed star patterns with an onboard star catalog. These sensors typically consist of a camera with a narrow FoV, a baffle to reduce stray light, and onboard processors to match star patterns to the catalog and compute the spacecraft's orientation. Due to their high accuracy—often in the arcsecond range—they are commonly used in high-precision attitude determination systems, although they are sensitive to sunlight and require careful thermal and optical shielding. However, in order to achieve a sufficient resolution and identify enough stars, the optical systems of the star sensors are usually too large for microsatellites [21].
- **Magnetometer**: Magnetometers measure the local magnetic field, which is dominated by the Earth's geomagnetic field in low Earth orbit. By comparing the measured vector to a geomagnetic field model such as the World Magnetic Model (WMM), the spacecraft's orientation relative to the field can be estimated. Common implementations include fluxgate magnetometers, which are reliable and sensitive, and Anisotropic Magnetoresistance (AMR) sensors for smaller platforms. While magnetometers on their own are not sufficient for full 3-axis attitude determination, as they provide only a single reference vector that must be augmented, they are often used in

combination with other sensors or in rate-only estimation algorithms, such as B-dot controllers, for detumbling [20, 22]. A major limitation of this type of sensor is its high sensitivity to magnetic disturbances. These can arise from interactions between the Earth's external magnetic field and the structure of the satellite, or from the electronics onboard the satellite. Some of these can be reduced by calibration; however, disturbances arising from nearby electric currents are particularly challenging to mitigate.

- **Gyroscope**: Gyroscopes measure the rotation rate of the spacecraft. Originally, gyroscopes used the conservation of angular momentum and spinning masses to measure the orientation. Alternatively, rate gyroscopes would take advantage of the precession properties of a gyroscope to directly measure the rotation rates. Traditionally, spacecraft have used either spinning mass gyroscopes or optical gyroscopes, relying on the Sagnac effect, where a rotation rate around a coiled optical fiber results in a phase shift of light travelling in opposite directions through that fiber [23]. With the miniaturisation of satellites with CubeSats and PocketQubes, the newest developments have produced semiconductor microelectromechanical systems based on the Coriolis effect, which have rapidly developed to approach the performances of optical gyros, while remaining more compact [24]. Gyroscopes suffer from various sources of error. One is noise, resulting in an angular random walk, where the integrated attitude results in a random walk about the true attitude. The most significant error source, however, is bias, bias instability, and bias temperature dependence. Bias is a constant offset of the sensor relative to the true values, whereas bias instability quantifies its stability. Bias can be compensated if the bias instability is sufficiently low. However, especially for MEMS sensors, temperature-dependent bias is harder to compensate for, especially in orbit, where temperature can swing significantly [25].

The sensors most commonly used on board microsatellites so far have been a mixture of Sun sensors, magnetometers, and gyroscopes, with occasional use of miniaturized star trackers or thermopile-based Earth horizon sensors. This is mainly due to the availability of Sun sensors, magnetometers, and gyroscopes as COTS components that are easily integrated on board a printed circuit board, allowing for good integration and miniaturisation.

The disadvantage of Sun sensors is the absence of the Sun during approximately half the time of each typical orbit, allowing full observation of the attitude only during that limited time. Earth horizon sensors, if based on infrared wavelengths, on the other hand allow continuous observation of the local nadir vector without interruption due to the constant emission of thermal infrared radiation from the Earth.

Magnetometers measure the Earth's magnetic field, which can be precisely determined as a function of location and time [22]. However, magnetometers are easily disturbed by other nearby magnetic components such as any electronics and electric actuators like motors for reaction wheels or magnetorquers. Hence, the magnetometer must be located far away from the rest of the satellite, often on a deployable boom. [26]. Additionally, any hard iron or soft iron effects from surrounding materials require calibration to ensure the measured magnetic field is correct [27].

2.4. Imaging sensors for attitude determination

Imaging sensors have not been widely utilized for spacecraft attitude determination; [28] however, they have been more frequently developed and adopted in unmanned aerial vehicles, which have seen a significant increase in usage recently [29]. In spacecraft, image based navigation is typically used to determine the relative position of spacecraft towards celestial bodies [28], where many methods have been developed to determine the accurate location of planetary limbs [30].

However, these methods typically have used visible light and are used for crewed missions, requiring celestial navigation as a self-contained contingency [31]. The reason behind using visible light is the higher resolution and quality of images available for this wavelength. This is critical for the high accuracy required for this type of celestial navigation, which aims to determine not the attitude but the position of a spacecraft.

The use of infrared imagery for attitude determination has been more extensively developed in the context of unmanned aerial vehicles (UAVs). At typical UAV operating altitudes, the Earth's curvature is negligible, and the horizon can be approximated as a straight line. Consequently, attitude estimation reduces primarily to robust horizon detection. However, unlike space-based applications, the visible horizon in UAV imagery is frequently distorted or partially obscured by terrain features such as vegetation, buildings, and local topography. The primary challenge in this context is therefore not geometric modeling of the horizon, but reliably distinguishing the true horizon from foreground obstructions [32, 33].

One concept combining these has been proposed for use in a PocketQube, using small infrared imagers in combination with a Micro-Electromechanical Systems (MEMS) gyroscope for attitude determination [34]. However, as far as publicly known, this has only been proposed with simulations that do not simulate the image processing, instead assuming a given accuracy of the reference vector determination for both Earth and Sun. Hence, continuing this effort and implementing the complete end-to-end image-to-attitude system is an area still to be explored.

2.5. Approach of determining attitude from imagery

Determining the attitude from imagery consists of multiple steps, each of which can be performed in a number of ways. In the following section, a review is done of the different general steps taken in the processing pipeline that is used for conventional image processing techniques.

2.5.1. Image distortion correction

Camera lenses inevitably deviate from the ideal pinhole projection model, introducing geometric distortions that must be modeled and corrected during calibration. Two prominent classes of distortion are *radial* distortion, caused by the symmetric bending of light through spherical lenses, and *tangential* (or decentering) distortion, resulting from lens misalignment or imperfect centering of the sensor relative to the lens.

Figure 2.6: Image distortion types [35]

Zhang [36] proposed a widely adopted distortion model used in planar pattern-based camera calibration. Radial distortion is represented as a polynomial in the squared radial distance $r^2 = x^2 + y^2$, while tangential distortion accounts for decentering effects:

$$x_d = x(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) + 2p_1 xy + p_2(r^2 + 2x^2) \tag{2.1}$$

$$y_d = y(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) + p_1(r^2 + 2y^2) + 2p_2 xy \tag{2.2}$$

Here, (x, y) are normalized image coordinates, (x_d, y_d) are the distorted coordinates, k_i are radial distortion coefficients, and p_i are tangential distortion coefficients. Heikkilä and Silvén [37] proposed a similar model with up to fourth-order radial terms, and used a structured four-step calibration approach, combining linear parameter estimation and nonlinear refinement via bundle adjustment. These polynomial distortion models are well-suited to narrow or moderate FoV lenses. One way to estimate these is to use genetic algorithms [38].

2.5.2. Image preclassification

Classifying images according to the presence of key features, such as the Earth's horizon or the Sun, can improve the efficiency and reliability of subsequent image processing algorithms. Simple methods for such classification include intensity-based thresholding and region statistics. For example, if the Sun is the brightest object in the image, a threshold on pixel intensity or a percentile of the brightest pixels can indicate its presence. Similarly, horizon detection can sometimes be approximated by analyzing intensity gradients or thresholding of images to separate dark space from the brighter Earth. These approaches are computationally cheap, require no training data, and can run in real time on resource-constrained hardware.

One method of thresholding that is useful for this type of application is the Otsu-threshold method, which dynamically adjusts the level at which the threshold for a specific image is chosen to optimize the segmentation of the intensity regions [39]. Based on the size of the regions resulting from the threshold and their intensity, various decisions can then be made, such as whether it is likely that the Earth and Sun are present. This step can be inaccurate, as further outlier removal and consistency checking can be applied later [40].

Convolutional Neural Networks (CNN)s offer a more robust approach by learning hierarchical spatial features directly from images. As demonstrated in foundational works like LeNet-5 [41] and explained in tutorials such as CS231n [42], CNNs use convolutional layers to detect edges and shapes, pooling layers to provide translation invariance, and dense layers to combine these features for classification. For tasks like horizon and Sun detection, small multi-label CNNs can output probabilities for each feature, allowing images to be classified as ”horizon visible” and/or ”Sun visible.” With proper data augmentation like varying brightness, contrast, rotation, and position, CNNs can generalize across a wide range of conditions and are less sensitive to noise, compared to simple thresholding methods. Lightweight architectures inspired by LeNet or MobileNets [43] can even be deployed on microcontrollers, enabling onboard preprocessing of low-resolution images in real time.

2.5.3. Edge detection methods

A number of different methods exist to detect the horizon of the Earth in images captured from low Earth orbit. The horizon appears as a sharp transition in pixel intensity, typically from the dark expanse of space to the brighter curved surface of the Earth. As such, it can be detected using classical edge detection algorithms or with more advanced techniques that offer subpixel accuracy. The general sequence from a raw image to edges is shown in Figure 2.7.

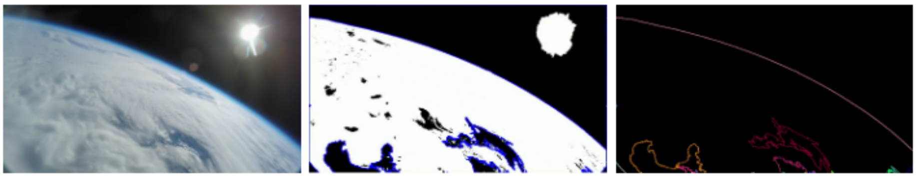

(a) Original (b) Binary (c) Edges

Figure 2.7: Example of edge detection on a RGB horizon image, converted to a binary image as a intermediate step [40].

Classical edge detection. One of the most common approaches is the Canny edge detector [44], which identifies intensity gradients in the image and highlights pixels where there is a strong change. It consists of several steps: smoothing the image with a Gaussian filter, computing the gradient magnitude and direction, applying non-maximum suppression to thin edges, and using double thresholding followed by edge tracking to produce a binary edge map. This method is simple and computationally efficient, making it attractive for real-time or onboard processing.

However, classical edge detectors are limited to pixel-level accuracy. When working with low-resolution imagery, this pixel-level uncertainty can translate into significant angular errors in the attitude determined from the detected edge.

Subpixel edge localization. To improve upon this, more sophisticated methods have been developed to locate the limb of the Earth with subpixel precision. These methods model the local shape and intensity variation of the horizon in a small window and extract geometric features, such as the precise position and curvature of the edge, from the phase and amplitude of the computed moments [45].

Zernike moments are a set of orthogonal basis functions defined on the unit disk that can compactly represent the shape and orientation of image structures. When applied to small patches along the detected edge, they allow refinement of the edge location to a fraction of a pixel. This can greatly enhance the accuracy of limb fitting, especially when working with small sensors or wide fields of view.

Trade-offs. While subpixel methods are more computationally demanding and may require careful calibration and tuning, they offer significant gains in attitude estimation accuracy. In particular, for applications involving low-resolution sensors or missions with stringent pointing requirements, the additional complexity can be well justified if the computational power required is available.

2.5.4. Earth direction finding from edges

Once the edge of the Earth is found, a circle, ellipsoid or hyperbolae can be fit to the edge [46] [47]. The ellipse or hyperbolae is the more general case, and in the case of non-spherical objects fitting an ellipse can provide attitude information beyond just the current vector to the center of the body. This is due to the possibility of determining the orientation of the ellipsoid relative to the imager, allowing the determination of the roll of the sensor relative to the imaged body. However, for spheroids (i.e., ellipsoids of revolution) with little flattening, such as the Earth, the possible accuracy of the attitude determination is very low and is seldom worth using [46]. This is especially true for the case investigated in this work, where the sensor resolution is very low. The decision between fitting a hyperbola or ellipse to the detected edges depends on the geometry of the problem, with the altitude and FoV determining what kind of curve is visible [47]. This is illustrated in Figure 2.8.

Figure 2.8: Earth edge curve type depending on sensor FoV and altitude [47]

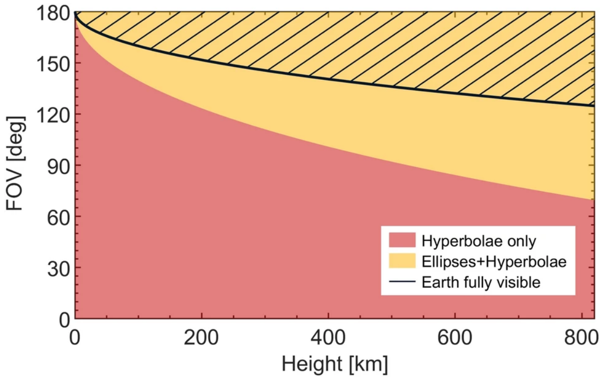

For all of these cases, the visible horizon can be described by a conic section of the form

$$Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0 \tag{2.3}$$

Hence, the most obvious course here is thus to fit these coefficients to the pixels identified as being on the edge of the Earth, which will require some constraints on the coefficients. However, various studies have shown this to be somewhat unstable and less accurate than other approaches [30]. Instead, the so-called Christian-Robertson algorithm can be used, which is a more direct least-squares method to determine the nadir vector rather than estimating the conic section shown in the image.

The Earth horizon as seen by a calibrated pinhole camera can be exploited to directly determine the nadir direction, i.e. the unit vector pointing from the spacecraft toward the centre of the Earth in the camera frame. In the ideal pinhole model, all light rays that form the image originate from straight lines through the camera aperture. If the Earth is approximated as a sphere, the set of rays corresponding to the observed limb forms a cone with apex at the camera centre and axis along the nadir vector $\mathbf{e}_c$ in the camera frame. The projection of this cone onto the image plane produces a conic section with coefficients that are directly related to the nadir direction.

Following the geometric fundamentals in Braun & Barf [47], let $\mathbf{p}_c = [p_{x,c}, p_{y,c}, f]^\top$ denote a point in the camera coordinate frame located on the projected horizon, where f is the focal length. The Earth nadir vector $\mathbf{e}_c = [e_{x,c}, e_{y,c}, e_{z,c}]^\top$ satisfies the cone equation

$$\mathbf{p}_c^\top \mathbf{e}_c = \cos\alpha \, \|\mathbf{p}_c\| \tag{2.4}$$

α is defined as

$$\alpha = \arcsin\left(\frac{r+d}{r+H}\right) \tag{2.5}$$

The geometry of this is shown in Figure 2.9

Figure 2.9: Geometry of the planetary limb determination [47]

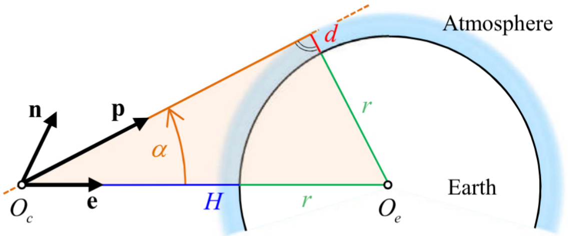

The projected image of the horizon therefore satisfies a conic of the form

$$A\,p_{x,c}^2 + B\,p_{x,c}p_{y,c} + C\,p_{y,c}^2 + D\,p_{x,c} + E\,p_{y,c} + F = 0, \tag{2.6}$$

with the conic coefficients A through F given explicitly in terms of the nadir components and $\cos\alpha$. Rather than performing an intermediate conic fit, the horizon extraction algorithm can treat each detected horizon pixel (u_i, v_i) as a line-of-sight direction.

A third and more direct approach is to estimate the Earth nadir vector $\mathbf{e}_c$ directly from the detected horizon points. Each of the m pixels on the extracted horizon curve corresponds, via camera calibration, to a point $\mathbf{p}_{c,i} = [p_{x,i},\, p_{y,i},\, f]^\top$ in the camera coordinate frame. For each such point, the linear horizon constraint given in (2.7) holds,

$$\mathbf{p}_{c,i}^\top \mathbf{e}_c = \cos\alpha \, \|\mathbf{p}_{c,i}\|, \tag{2.7}$$

which can be interpreted as a linear measurement equation in the unknown nadir vector $\mathbf{e}_c$. Stacking all m horizon points yields the linear system

$$\mathbf{H} = \begin{bmatrix} \mathbf{p}_{c,1}^\top \\ \mathbf{p}_{c,2}^\top \\ \vdots \\ \mathbf{p}_{c,m}^\top \end{bmatrix}, \tag{2.8}$$

and the corresponding right-hand side vector

$$\mathbf{y} = \begin{bmatrix} \cos\alpha \, \|\mathbf{p}_{c,1}\| \\ \cos\alpha \, \|\mathbf{p}_{c,2}\| \\ \vdots \\ \cos\alpha \, \|\mathbf{p}_{c,m}\| \end{bmatrix}. \tag{2.9}$$

The Earth nadir vector in the camera frame is then obtained by solving the overdetermined system in a least-squares sense,

$$\hat{\mathbf{e}}_c = (\mathbf{H}^\top \mathbf{H})^{-1}\mathbf{H}^\top \mathbf{y}, \tag{2.10}$$

followed by normalization to enforce $\|\hat{\mathbf{e}}_c\| = 1$.

2.5.5. Sun location determination

The FoV of the sensor and size of the Sun results in the Sun appearing as a point source in any images. For this type of source detection and location determination, a number of algorithms have been developed, one of which is the DAOPHOT method [48]. However, all of these methods have one thing in common, which is the subtraction of the image background to better identify any and all point sources in the image and to evaluate their strength. When identifying the Sun for attitude determination, this can become challenging because portions of the image may be occupied by the Earth, resulting in a non-uniform background behind potential point sources. Hence, it becomes necessary to mask off the Earth in the image. This is the approach taken for the calibration of a very similar sensor for the use on MRC-100, another PocketQube attempting a similar attitude determination system [38].

Other approaches that have been successfully used is to simply assume that the Sun will always be the brightest object in the image [49]. This requires some thresholding to be applied to prevent spurious detections in case the Sun is not visible in an image.

2.5.6. Outlier and noise removal

Any sensor will not only measure the desired signal, but also various noise. This can possibly be noise introduced from the sensors' working principle, or in the imaging case, simply some signal from an undesired source. Thus, to improve the function of the algorithm to be developed, the noise needs to be removed as far as possible.

For the case of the detection of the horizon line, multiple approaches are described in literature. One possible approach is the Random Sample Consensus (RANSAC) algorithm, which repeatedly selects random data points among those identified as the horizon by an edge detection method. To each of these, it fits a conic section and selects the conic section with the most candidate points in its vicinity as the correct point [40]. Due to the number of necessary iterations, this requires considerable computational power. Another approach is image thresholding before the edge detection, to remove possible candidate edges entirely within or outside the object of interest. This works best when there is a large contrast between these image regions, resulting in a clear edge between the Earth and space region [40].

A method to fit a coherent line onto the horizon line using the predetermined knowledge of what the horizon should look like are active contour models, or commonly called snakes in computer vision. These methods try to numerically minimize energy functionals of the form

$$E(C, c_1, c_2) = \mu \, \text{Length}(C) + \lambda_1 \int_{\text{inside}(C)} \big(I(x,y) - c_1\big)^2 \, dx \, dy + \lambda_2 \int_{\text{outside}(C)} \big(I(x,y) - c_2\big)^2 \, dx \, dy \quad (2.11)$$

where C is the curve between the regions, while c_1 and c_2 are the average intensities of the image inside and outside of C, respectively. Additionally, regularization terms are added, such as the length of the line or the area inside/outside the curve C to promote convergence. If C becomes the actual boundary between two regions in an image, it will minimize this energy functional, thus determining the desired horizon line[50].

This specific example is taken from a model that attempts to segment the image into two separate regions based on their contours [50]. This separation in this case is based not on the gradient between these regions, but on the best separation of intensity of the image into two regions. Other formulations of this approach exist where different functionals are used to detect different features. In fact, one of the original proposed functionals by Kass was the negative gradient of the image to detect edges[51]. These methods rely on gradient descent along the energy functional PDE, and thus require sizeable computation depending on the size of the image.

2.5.7. Neural Network approach

Deep learning approaches, and CNN in particular, have been successfully applied to a broad range of computer vision tasks such as classification, localization, and regression directly from image data. A CNN consists of layers of convolutional filters that learn spatial features from images, followed by nonlinear activations and pooling operations to reduce spatial redundancy [52, 53]. These learned features can be used not only for discrete classification but also for continuous regression, where an input image is mapped to a quantity of interest such as a direction vector or pose parameter [54].

In regression-based vision tasks, the network is trained end-to-end on pairs of images and ground-truth labels (e.g. direction vectors), with the training objective minimizing a loss function between predicted and true continuous values [54]. This framework removes the need for manually engineered feature extraction (e.g., edge detectors or thresholding), and instead leverages learned image representations that can implicitly encode complex geometric and photometric relationships. For spaceborne applications where computational resources are limited, lightweight architectures such as MobileNet variants have been developed to reduce parameter count and computation while retaining regression performance [55, 56].

Direct application of CNNs to spacecraft attitude and position estimation from Earth imagery has been demonstrated in recent literature. A method for simultaneous satellite position and attitude determination using a panoramic infrared Earth sensor and a deep feature–based registration framework has been developed using a convolutional neural network, with a modified architecture to retain more spatial feature accuracy [57]. In this approach, a CNN backbone extracts multiscale visual features from the Earth image, which are then matched to a reference Earth map to recover both orientation and position relative to the Earth model. This work shows that learned features can reliably capture the global structure of the Earth limb and surface patterns, enabling accurate pose estimation from a single view.

While methods specifically targeting Sun direction estimation with CNN are less common in the literature than Earth limb regression, related studies in robotics and autonomous vehicles have demonstrated that CNN can accurately regress Sun position from images when trained on appropriate datasets [58]. In the spacecraft context, the problem of Sun detection is simpler in principle: the Sun appears as a bright, compact source against a dark background and its image position can be mapped to a direction vector using camera calibration. CNN-based regression or heat-map localization methods thus present a promising alternative to traditional thresholding or centroiding when robust performance under varying illumination and partial occlusions is required.

Overall, these deep learning–based regression approaches replace individual hand-designed processing steps (e.g. Earth thresholding, limb fitting, Sun thresholding, centroiding) with a unified learned mapping, simplifying the pipeline and potentially improving robustness to noise, contrast variation, or partial feature visibility. However, the tradeoff for this increased accuracy and robustness is a large increase in computational resources and required memory.

2.5.8. Comparison of Processing Strategies

Although a wide range of algorithms has been proposed in literature, most imagery-based attitude determination solutions can be decomposed into three fundamental functional steps:

1. Image classification or region separation (e.g. Earth vs. space, Sun vs. background),
2. Feature localization (extraction of geometric primitives such as edges or points),
3. Geometric inference (estimation of physical direction vectors from image geometry).

Not all approaches separate these steps explicitly; some merge multiple stages into a single optimization or learning-based framework. For clarity, the reviewed methods are grouped below according to the primary function they perform within the pipeline.

1. Image classification or region separation

The purpose of this stage is to isolate meaningful regions of the image before geometric reasoning is applied. For Earth sensing, this typically corresponds to separating the illuminated Earth from space. For Sun detection, it corresponds to distinguishing bright point sources from the background.

Threshold-based techniques rely on sufficient intensity contrast between regions and produce binary masks that can be further processed. Adaptive thresholding improves robustness under varying illumination but remains intensity-driven. CNN-based segmentation replaces explicit intensity modeling with learned pixel-wise classification and can output masks, class probabilities, or directly refined regions depending on architecture design. Region-based active contour methods differ conceptually: instead of first producing a mask and then detecting edges, they directly optimize for a curve separating two

Table 2.1: Region separation approaches

Method	Computational Load	Illumination Robustness	Output Type
Fixed thresholding	Very Low	Low	Binary mask
Otsu adaptive threshold	Low	Medium-High	Binary mask
CNN segmentation	High	High	Implementation dependent
Active contour (region-based)	High	High	Closed boundary curve

intensity-consistent regions. As such, they partially combine region separation and boundary estimation into a single numerical problem.

2. Feature localization

After regions have been identified—explicitly through segmentation or implicitly through model-based methods—the next step is to extract geometric primitives that can be used for attitude estimation. For Earth sensing, this corresponds to the horizon edge while for Sun sensing it is typically the centroid of a bright region.

Table 2.2: Feature localization approaches

Method	Accuracy	Computational Load
Intensity based Centroiding (Sun)	Subpixel	Low
Pixel-level edge (Canny)	Pixel	Low–Med
Subpixel moment-based	Subpixel	Medium–High
CNN object detection	Subpixel	High

Centroiding of thresholded bright regions provides subpixel accuracy at very low computational cost and is therefore attractive for point-like sources such as the Sun. Pixel-level edge detectors such as Canny deliver computationally efficient horizon extraction but are limited by image resolution. Subpixel moment-based methods increase localization accuracy by modeling the local intensity distribution near the edge, at the expense of additional computation.

CNN-based object detection performs feature localization through learned bounding boxes or heatmaps, typically achieving robust and accurate localization but requiring substantially higher computational resources. Additionally, they can combine many, if not all, steps from image to reference vector in one net.

3. Geometric inference

In the final stage, localized image features are converted into physically meaningful direction vectors in the camera frame. This step links image geometry to spacecraft attitude.

Table 2.3: Geometric inference approaches

Method	Numerical Stability	Accuracy Potential	Noise and outlier tolerance
Conic coefficient fitting	Moderate	Medium	Medium
Direct nadir least-squares	High	High	Low
RANSAC based nadir least-squares	High	High	High
Learned end-to-end regression (CNN)	Data dependent	High	Training dependent

Classical conic coefficient fitting estimates the parameters of the observed limb curve and derives the nadir vector indirectly, which can lead to numerical sensitivity to noise and calibration errors. Direct least-squares estimation of the nadir vector removes the intermediate conic parameter step and improves numerical stability while retaining a physically interpretable model. RANSAC-based variants increase robustness against misdetected edge points but require increased computational effort.

End-to-end regression networks bypass explicit geometric modeling entirely and learn the mapping from image to attitude directly. While potentially accurate, their performance is inherently tied to training data representativeness and domain consistency.

Discussion

The reviewed approaches differ primarily in how explicitly they decompose the three fundamental tasks and how strongly they rely on geometric modeling versus learned representations. Classical pipelines perform region separation, feature localization, and geometric inference sequentially. This modular structure enables predictable computational cost and clear error attribution between stages. Active contour methods partially merge region separation and localization into a single optimization framework, increasing robustness but also computational demand.

CNN-based approaches may merge all three steps into a unified learned model, improving tolerance to complex illumination and noise conditions, but introducing training dependence and higher implementation complexity. For low-resolution imagery and embedded processors with limited resources, maintaining a clear separation between segmentation, feature extraction, and geometric inference provides predictable runtime behavior and interpretable performance trade-offs. More integrated approaches may offer improved robustness, but at increased computational and development cost.

2.6. Attitude Data Fusion

Attitude determination using vector observations is fundamentally limited by observability constraints. A single direction measurement constrains only two rotational degrees of freedom, leaving rotations about the measured direction unobservable. This ambiguity is not merely instantaneous, but structural: even when measurements are accumulated over time, full attitude observability is not guaranteed unless additional independent information is available. Interestingly, over time, observability of the entire attitude state can be achieved if the reference vector and attitude rates are known in the same frame, be it the body or inertial reference frame [59]. This is the case if an external system can be used to estimate the angular rates of a body in the reference frame (as opposed to gyroscopes expressing the angular rates in the body frame) or a known vector in the body frame is sensed in the inertial reference frame, as can be done via differential Global Navigation Satellite System (GNSS).

When angular velocity is measured in the body frame (as provided by gyroscopes) and the reference direction is defined in the inertial frame (such as the Sun or Earth nadir), the rotation about the reference direction remains unobservable. In this configuration, the ambiguity axis is fixed in inertial space, and the combination of gyroscope propagation and repeated observations of the same reference vector does not resolve the missing degree of freedom. However, if the reference vector is varying in the inertial reference frame, the attitude can still be obtained [60]. This requires sufficient dynamics of the reference to allow convergence and stability.

The attitude information acquired by the imaging sensor is only available at discrete time points and only considers the data at each specific moment. To improve the quality of attitude estimation, we can incorporate gyroscope data that provides continuous, albeit noisy and possibly biased, measurements of angular velocity. There are several approaches to combine these different types of measurements, each with its own advantages and disadvantages. Two fundamental filtering approaches are particularly relevant: complementary filters and Kalman filters.

2.6.1. Complementary Filter

The complementary filter combines measurements from different sensors using a weighted average, typically blending low-frequency information from one source (e.g., reference vector from a Sun sensor) with high-frequency information from another (e.g., gyroscopes). The basic formula for a first-order complementary filter is:

$$\theta = \alpha(\theta + \omega\Delta t) + (1 - \alpha)\theta_a \tag{2.12}$$

where:

- θ is the estimated angle

- ω is the angular velocity from the gyroscope
- θ_a is the angle determined from the example reference sensor
- α is the weighting factor (0 < α < 1)
- Δt is the time step

The weighting factor α determines the contribution of each sensor. A higher α gives more weight to the gyroscope data, which provides good short-term accuracy but drifts over time. A lower α gives more weight to the correcting reference sensor data, which is stable in the long term but noisy in the short term. This way, the sensor anchors the long-term measurements, while allowing short-term dynamics to be accurately resolved using the gyroscope.

This approach is computationally efficient and doesn't require complex modeling of sensor noise characteristics. It works particularly well when the noise properties of the sensors are known and relatively constant, while also containing a number of drawbacks:[61]

- It assumes constant noise characteristics for each sensor
- The optimal weighting factor depends on the specific application and sensor properties
- It doesn't provide a statistical measure of estimation error
- Performance degrades with varying noise conditions or dynamic maneuvers

2.6.2. Kalman Filter as Weighted Estimator

The Kalman filter provides an optimal solution for estimating the state of a system by weighting observations against the current state estimate based on their relative confidence (covariances). For linear systems with Gaussian noise, it's the theoretically optimal estimator [62].

The system is modeled as:

$$\mathbf{x}_k = \mathbf{F}_k\mathbf{x}_{k-1} + \mathbf{B}_k\mathbf{u}_k + \mathbf{w}_k \quad \mathbf{w}_k \sim \mathcal{N}(0, \mathbf{Q}_k) \tag{2.13}$$

$$\mathbf{z}_k = \mathbf{H}_k\mathbf{x}_k + \mathbf{v}_k \quad \mathbf{v}_k \sim \mathcal{N}(0, \mathbf{R}_k) \tag{2.14}$$

where $\mathbf{x}_k$ is the state vector, $\mathbf{z}_k$ is the measurement, $\mathbf{F}_k$ is the state transition matrix, $\mathbf{H}_k$ is the measurement matrix, and $\mathbf{Q}_k$ and $\mathbf{R}_k$ are the process and measurement noise covariances [63].

The filter operates in two steps:

1. **Predict step**:

$$\hat{\mathbf{x}}_k^- = \mathbf{F}_k\hat{\mathbf{x}}_{k-1} + \mathbf{B}_k\mathbf{u}_k \tag{2.15}$$

$$\mathbf{P}_k^- = \mathbf{F}_k\mathbf{P}_{k-1}\mathbf{F}_k^T + \mathbf{Q}_k \tag{2.16}$$

2. **Update step**:

$$\mathbf{K}_k = \mathbf{P}_k^-\mathbf{H}_k^T(\mathbf{H}_k\mathbf{P}_k^-\mathbf{H}_k^T + \mathbf{R}_k)^{-1} \tag{2.17}$$

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}_k^- + \mathbf{K}_k(\mathbf{z}_k - \mathbf{H}_k\hat{\mathbf{x}}_k^-) \tag{2.18}$$

$$\mathbf{P}_k = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_k^- \tag{2.19}$$

The key insight is that $\mathbf{K}_k$ acts as a weighting factor between the predicted state and the measurement, similar to the α in the complementary filter. However, unlike the fixed α, $\mathbf{K}_k$ is dynamically calculated based on the relative confidence in the prediction ($\mathbf{P}_k^-$) versus the measurement ($\mathbf{R}_k$) [64].

2.6.3. Nonlinear Variants

For systems with nonlinearities such as

$$\mathbf{x}_k^- = f(\mathbf{x}_{k-1}, \mathbf{u}_k) + \mathbf{w}_k, \tag{2.20}$$

$$\mathbf{y}_k = h(\mathbf{x}_k) + \mathbf{v}_k, \tag{2.21}$$

where $f(\cdot)$ and $h(\cdot)$ are nonlinear functions, and $\mathbf{w}_{k-1}$ and $\mathbf{v}_k$ represent zero-mean process and measurement noise, respectively, the linear Kalman filter cannot be used. In these cases, the following methods are used most frequently:

- **EKF**: Linearizes the system around the current estimate using Jacobian matrices [65]:

  $$\mathbf{F}_k = \left.\frac{\partial f}{\partial \mathbf{x}}\right|_{\hat{\mathbf{x}}_{k-1}} \tag{2.22}$$

  $$\mathbf{H}_k = \left.\frac{\partial h}{\partial \mathbf{x}}\right|_{\hat{\mathbf{x}}_k^-} \tag{2.23}$$

  The EKF uses these linearized matrices in the standard Kalman filter equations. This requires analytic computation of Jacobians and may perform poorly with significant nonlinearities [63]. The EKF also lacks guaranteed convergence, as the local linearization at the estimated state may be sufficiently different from the true system dynamics at the true state.

- **Unscented Kalman Filter (UKF)**: Uses a deterministic sampling approach with sigma points to propagate the mean and covariance through nonlinear transformations without explicit linearization [66]. A set of sigma points $\mathcal{X}_i$ in a pattern around the current estimate is chosen to represent the distribution, propagated through the nonlinear functions, and then used to compute the updated mean and covariance [67]. The UKF effectively evaluates the nonlinearity by transitioning the entire vicinity of the current state estimate through the propagation, capturing the nonlinear effects more accurately than the EKF. This approach generally provides better performance for highly nonlinear systems without requiring Jacobian calculations, instead requiring multiple evaluations of the propagation function [67].

The main difference from the linear Kalman filter is that both EKF and UKF attempt to handle nonlinearities, with UKF generally providing better performance for highly nonlinear systems while requiring more computational resources[67].

2.6.4. Consider State Approach

The consider state approach is used when certain parameters influence the system dynamics or measurements but are not estimated explicitly. These parameters are assumed to be uncertain, slowly varying, or poorly observable, and are therefore treated as nuisance quantities rather than solve-for states. Typical examples in attitude estimation include sensor calibration terms or environmental sensitivities.

The full system state is conceptually partitioned as

$$\mathbf{x}_a = \begin{bmatrix} \mathbf{x} \\ \mathbf{x}_c \end{bmatrix}, \tag{2.24}$$

where $\mathbf{x}$ denotes the primary states to be estimated (e.g., attitude and gyro bias), and $\mathbf{x}_c$ denotes the consider states. The consider states are assigned nominal values and associated uncertainties, but are not updated using measurements. The filter maintains the mean of the primary state $\hat{\mathbf{x}}$, along with the full covariance matrix

$$\mathbf{P}_a = \begin{bmatrix} \mathbf{P} & \mathbf{P}_{xc} \\ \mathbf{P}_{cx} & \mathbf{P}_c \end{bmatrix}, \tag{2.25}$$

where $\mathbf{P}$ is the covariance of the estimated states, $\mathbf{P}_c$ represents uncertainty in the consider states, and $\mathbf{P}_{xc}$ captures their coupling [68].

During time propagation, both the primary-state covariance and the cross-covariance terms are propagated through the system model. In the measurement update, the Kalman gain is computed such that the effect of consider-state uncertainty on the primary-state estimate is accounted for, while the consider-state estimates themselves remain unchanged. This way the impact of these nuisance quantities are taken into account during the covariance determination of the filter. This prevents overconfidence and thereby improves the stability of the filter while at the same time reducing the required computational resources.

This formulation, commonly referred to as the Schmidt–Kalman filter, improves the consistency of the estimated states by modeling the influence of uncertain parameters without increasing the number of estimated states or requiring their observability [64].

2.6.5. Attitude Representation and Nonlinearities

Attitude estimation is inherently nonlinear due to the geometry of three-dimensional rotations. Unit quaternions provide a compact and singularity-free representation of orientation and are therefore widely used in spacecraft attitude determination and inertial navigation systems [69]. A unit quaternion $\mathbf{q} \in \mathbb{H}$ represents an element of the rotation group SO(3) and is subject to the unit-norm constraint $\|\mathbf{q}\| = 1$.

The set of rotations forms a Lie group with the following relevant properties:

- Rotations can be smoothly composed and inverted, with quaternion multiplication acting as the group operation.
- The group SO(3) is a smooth manifold rather than a linear vector space.
- Small rotational perturbations are naturally represented in the associated Lie algebra $\mathfrak{so}(3)$, which is locally isomorphic to $\mathbb{R}^3$.

Because of this Lie group structure, attitudes cannot be manipulated using standard vector space operations. In particular, additive updates of quaternions are not geometrically meaningful and may violate the unit-norm constraint. This poses a problem for standard Kalman filter formulations, which assume that both the state and its estimation errors evolve in a linear space.

A naïve application of an Extended Kalman Filter using additive updates,

$$\mathbf{q}^+ = \mathbf{q}^- + \delta\mathbf{q}, \tag{2.26}$$

is therefore inconsistent: the updated quaternion may no longer represent a valid rotation, and the linearized error dynamics assumed by the filter do not respect the underlying manifold geometry [70]. One practical but ad hoc workaround is to renormalize the quaternion after the additive update:

$$\mathbf{q}^+ = \frac{\mathbf{q}^- + \delta\mathbf{q}}{\|\mathbf{q}^- + \delta\mathbf{q}\|}. \tag{2.27}$$

While this preserves the unit-norm constraint, it does not resolve the fundamental geometric inconsistency of the additive update. The applied correction does not correspond to a well-defined rotation in SO(3), and the resulting filter can suffer from poor covariance consistency and degraded long-term performance [70, 69].

A more principled solution is provided by the error-state formulation, commonly implemented in attitude estimation as the *Multiplicative Extended Kalman Filter (MEKF)*. In this approach, the attitude is decomposed into:

- a *nominal attitude*, represented by a unit quaternion $\hat{\mathbf{q}}$ and propagated using the full nonlinear dynamics, and
- a *small attitude error*, represented by a three-dimensional rotation vector in the tangent space of SO(3). As this only represents the error from the nominal attitude state with which the nominal state is updated, after which the error is reset to zero, this remains small and the linearisation errors are negligible.

The true attitude is expressed multiplicatively as

$$\mathbf{q} = \delta\mathbf{q} \otimes \hat{\mathbf{q}}, \tag{2.28}$$

where $\delta\mathbf{q}$ denotes a small rotation constructed from the error-state vector. The Kalman filter operates only on this error state, which evolves approximately linearly and is therefore compatible with standard linear estimation techniques. Measurement updates correct the error state, and the nominal quaternion is subsequently updated via quaternion multiplication, preserving both the unit-norm constraint and the geometric structure of SO(3).

This error-state formulation respects the nonlinear geometry of rotational motion, maintains numerical stability, and yields consistent covariance behavior. As a result, the MEKF and related error-state Kalman filters have become the standard approach for spacecraft and inertial attitude estimation [70, 71].

2.7. Conclusion and Discussion of literature found

From the literature found it has become clear, that an optical attitude sensor combining Earth horizon sensing and Sun sensing would allow the further miniaturisation of a satellite attitude determination system. Such a system, while already proposed and analysed [34], has not been used operationally on a satellite yet. Especially, the use of IR sensors promises to allow an uninterrupted attitude estimate throughout an entire orbit and will thus be used.

For the determination of the reference vectors from imagery, a number of different methods have been found and analysed. A classical image processing approach promises to be fairly simple to implement, consume only marginal computational resources and allow for a reasonable performance in terms of accuracy and robustness. If that is not found to be sufficient a number of noise rejection algorithms such as RANSAC are available for an increase in computational complexity. Another method to increase performance, if it is not found to be sufficient, could be the introduction of more advanced subpixel edge detectors for the detection of the Earth's horizon.

Additionally, if computational resources are available and sufficient training data can be obtained, neural network based methods can be introduced to allow for better robustness against noise and potentially better accuracy. Following the detection of the reference vectors, either complementary filters or an extended or unscented Kalman filter can be used, of which the Kalman filters are most optimal and used most frequently. Due to the most advantageous attitude representation being of the quaternion form, special formulations of these must be used.

3

Project goals and planning

This chapter outlines the objectives and planning of the research presented in this thesis. It first formulates the overarching goal and associated research questions, establishing the performance targets and constraints that guide the work. Subsequently, an overview of the adopted approach is provided, describing how simulation, algorithm development, and experimental validation are combined to address the defined objectives. Together, these elements frame the scope of the research and clarify how the individual contributions relate to the central research goal.

3.1. Research Objectives and Questions

The overarching objective of this work is to **demonstrate that a compact and low-cost combination of an IR camera and a MEMS gyroscope can provide sufficient attitude determination accuracy for small satellites.**

To achieve this, two main sub-objectives are defined:

1. **Develop a validated algorithm for an infrared horizon sensor** that estimates the spacecraft attitude based on the detection of the Earth limb and the Sun.
    - 1.1. **Develop an image generation framework** capable of producing synthetic IR images of the Earth limb and Sun under various orbital and thermal conditions to support algorithm development and analysis.
    - 1.2. **Validate the framework experimentally,** using real IR imagery to ensure representative simulation outputs, and quantify discrepancies.
2. **Integrate the horizon sensor with a MEMS gyroscope** to achieve robust attitude estimation even when no celestial or terrestrial object is visible.
    - 2.1. **Develop an attitude estimation filter** that fuses horizon/Sun-based measurements with gyroscope and possibly magnetometer data for continuous attitude tracking.
    - 2.2. **Characterize performance under partial and full obscuration** to determine achievable accuracy, stability, and robustness during eclipse or off-nadir pointing conditions.

From these objectives, the following **research questions** are formulated:

- What are the key parameters that govern the achievable attitude accuracy of an IR horizon sensor?
- What is the expected accuracy under nominal operating conditions? Can the proposed attitude determination system provide sufficient accuracy to ensure that separation of the Delfi Twin mission occurs within the prescribed 52.7° half-angle cone?
- What are the best-case and worst-case performance limits?

- How does the performance compare to alternative sensor combinations, such as a Sun, magnetometer, and gyroscope systems?
- Is an additional sensor, such as a magnetometer, necessary to have acceptable performance when no object is in view?

3.2. Overview of the Approach

The research follows a sequential, iterative approach that combines simulation, hardware development, and experimental validation.

Stage 1: Hardware Development and Data Collection
A compact hardware setup is developed to facilitate image acquisition and system characterization. The setup includes a low-cost IR camera and MEMS gyroscope mounted on a controllable platform. The hardware allows the capture of representative images of the Earth's horizon under various conditions, serving both as test data for algorithm validation and as a baseline for comparison with simulation outputs.

Stage 2: Simulation Framework Development
A simulation framework is developed to generate synthetic IR horizon images for arbitrary satellite attitudes, altitudes, and thermal conditions. The simulator models Earth's temperature distribution and radiance effects, producing images that approximate the sensor's output under both daylight and eclipse conditions. This framework enables algorithm testing over a wide range of scenarios that are difficult or costly to reproduce experimentally.

Stage 3: Algorithm Development
Using the simulated images, several image-processing algorithms are developed and tested to extract the Earth limb and estimate the spacecraft's attitude. The algorithms are designed to balance computational simplicity, robustness, and accuracy, considering the limitations of onboard microcontrollers. The image-derived attitude estimates are then combined with gyroscope data in an integrated estimation scheme (e.g., an Extended Kalman Filter) to assess achievable overall performance.

Stage 4: Experimental Validation
Initially, the experimental validation was planned through a high-altitude balloon flight in the DLR/SNSA/ESA programme Rocket/Balloon Experiments for University Students (REXUS/BEXUS) to expose the system to representative near-space conditions. Following the failure of the balloon, the validation is restructured for laboratory testing. A heated Earth simulator is constructed to reproduce realistic infrared contrast between Earth and space, allowing the acquisition of experimental data for direct comparison with simulation results. This stage verifies whether the simulator produces image statistics and horizon geometries comparable to those expected in flight.

Stage 5: Integrated Performance Evaluation
The final stage combines simulation and experimental data to estimate the achievable accuracy of attitude determination under different conditions. The performance is evaluated in terms of attitude error, stability over time, and robustness to environmental variations (e.g., partial horizon visibility or eclipse).

4

Implementation of Horizon Sensor

This chapter will describe the implementation of an attitude determination system that utilises a low-resolution infrared camera to detect the horizon and Sun, and determine the attitude from that. Chapter 2 has shown that this is a concept that has partially been explored in various modifications, but has not yet been implemented for this type of mission.

4.1. Hardware selection and design

The concept of this sensor is based on low-resolution IR cameras as the sole source of attitude data, with the augmentation of a gyro to allow both uninterrupted and filtered attitude data to be determined. Thus, a suitable IR sensor is to be selected, which needs to be combined into an entire system that can sensibly process the images as well as use a gyro to enhance these momentary data points. First, a IR sensor is selected, which is then integrated into a self contained system that includes all other required components. The hardware design of this system is detailed here.

4.1.1. IR Sensor selection

For this specific purpose, only the smallest of IR cameras will work. It was found that only very few of these exist, and even fewer exist as COTS parts, with a size and cost compatible with the type of satellite to be developed. For the complete size of such a camera, including the entire lens and mounting, it was assumed that approximately 30x30x25mm is an upper bound of what is feasible to implement in a PocketQube. This is based on the 50mm measurement of the PocketQube standard. Three different sensors with a size smaller than this were found and are presented here.

MLX90640

The MLX90640 is a low-resolution infrared thermal imaging array primarily designed for embedded systems and low-cost thermal sensing. With a 32 × 24 pixel grid and either a 55°×35° or 110°×75° FoV, it offers sufficient spatial resolution for coarse thermal mapping and target detection while maintaining minimal processing and power requirements.

Typical applications advertised by the manufacturer include human or animal presence detection, occupancy monitoring, home automation, and fire or hotspot detection. The sensor communicates via an Inter-Integrated Circuit (I^2C) interface and is supported by a range of open-source software libraries, making it popular in both prototyping and commercial embedded products.

Multiple versions exist, with different resolutions and FoV. Only the highest resolution sensor is considered here. The FoV of the sensor can be either 55°×35° or 110°×75°, resulting in different package sizes due to the differences in the lens used. Apart from this difference, these two versions are identical. Typical prices range around 30-40€. Additionally, this specific type of sensor has been previously used in a similar satellite mission aboard the MRC-100 3p PocketQube, and apparently operated without issue for this duration [73].

Figure 4.1: MLX90640 sensor [72]

Figure 4.2: TeraRanger Evo Thermal sensor [74]

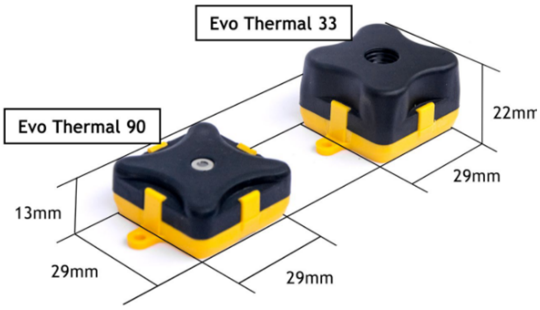

TeraRanger Evo Thermal

The TeraRanger Evo Thermal series (available in 90° and 33° FoV versions) is a compact, plug-and-play thermal module designed for robotics, automation, and environmental sensing. It uses a 32 × 32 pixel thermal sensor packaged with on-board processing and communication via Universal Asynchronous Receiver Transmitter (UART) or Universal Serial Bus (USB), allowing simple integration without dedicated image processing hardware.

The two variants target different use cases: the Evo Thermal 90 provides a wide coverage area suitable for general situational awareness, navigation, or occupancy mapping, while the Evo Thermal 33 offers a narrower FoV for temperature monitoring of specific targets or longer-range detection.

These sensors emphasize ease of integration, low weight, and low power consumption, making them attractive for Unmanned Aerial Vehicle (UAV)s, mobile robots, and Internet of Things (IoT) sensing platforms where visual privacy, environmental robustness, and fast prototyping are important. Typical prices range around 80 - 100€.

Seek Thermal Mosaic Core

Seek Thermal offers several Original Equipment Manufacturer (OEM) and consumer modules with comparable characteristics — typically in the 200×150 to 320×240 pixel range. These sensors are based on microbolometer arrays and are designed for portable imaging, industrial inspection, and security applications requiring higher resolution than compact embedded sensors like the MLX90640 or Evo Thermal.

Seek modules are commonly found in handheld thermal cameras, smartphone accessories, and in-

Figure 4.3: Seek Thermal Mosaic Core sensor [75]

dustrial monitoring systems, providing detailed thermal imagery at 15–25 Hz frame rates and Noise-Equivalent Temperature Difference (NETD) values down to tens of millikelvins. Variants differ mainly in resolution, lens type, and interface options (USB, Serial Peripheral Interface (SPI), or proprietary). According to the official Mosaic Core Engineering Datasheet [75], the module supports two main interface configurations:

- USB mode, when used together with a Seek coprocessor board (suffix P or PX), where the on-board processor handles non-uniformity correction, flat-field calibration, and thermal conversion.
- SPI mode, available on the bare sensor variant (suffix S or SX), which exposes a low-level digital video stream requiring the user to implement their own processing pipeline.

While this architecture offers flexibility, it also makes integration challenging. The sensor's 24-pin 0.4 mm-pitch Hirose DF40 connector is mechanically delicate and requires a custom flex cable or mating PCB. In SPI mode, the core outputs raw microbolometer data that must be corrected, linearized, and calibrated using proprietary algorithms not publicly documented in full detail. The datasheet provides only partial guidance, referring developers to Seek's Software Development Kit (SDK) (available only for the USB variant). Consequently, without the coprocessor board, users must design their own interface electronics, handle complex image calibration steps, and manage sensor pairing and EEPROM calibration data. Even in USB mode, integration is tied to Seek's SDK and drivers, limiting direct low-level access or use on embedded systems. In practice, this makes the S302 and related Mosaic Core modules less plug-and-play than the other thermal sensors, and more suitable for OEMs with experience in camera integration or access to Seek's developer support.

Overall, Seek's modules target use cases where fine spatial detail, accurate temperature measurement, and human-interpretable imaging are desired, rather than coarse thermal pattern recognition. Typical prices, depending on the specific models and capabilities, range between 500 and 650€.

Sensor selection

From Table 4.1 it is apparent that in the miniaturized form required for a sensor onboard a PocketQube, the sensor that balances its package best with the performance shall be selected. A limited trade-off using the criteria of size, cost, demonstrated space heritage, framerate and resolution is performed. For each category, a number between 1 and 3 is assigned to every sensor with equal weight for all criteria. The trade-off table is shown in Table 4.2.

In this case, the MLX90640 seems to be ideal, as it can be directly implemented onboard the same PCB that performs the processing and hosts the other Sensors onboard, such as the gyro. Additionally, the low cost is crucial for enabling the cheap development of such a sensor. The only question remaining is which version of the sensor shall be used. Two Field-of-View options exist, one with a large FoV with which almost the entire circumference of the satellite is covered and one with a smaller FoV where the view covered is smaller, though a higher accuracy of the horizon and Sun detection is expected.

To illustrate this further, the large FoV option would provide a view of 4 horizontal patches of 75 degrees, with vertical FoV of 110 degrees. This results in the coverage of 300 of 360 degrees within the horizontal

Property	**MLX90640**	**TeraRanger Evo Thermal 90/33**	**Seek Thermal Mosaic Core**
Resolution	32 × 24 px	32 × 32 px	200×150 – 320×240 px (typ.)
FoV (FOV)	55°×35° or 110°×75°	90° or 33° (model-dependent)	35°×27° to 57°×42° (model-dependent)
Frame Rate	0.5 – 64 Hz (programmable)	≈ 7 Hz	15 – 25 Hz (typ.)
NETD	≈ 0.1 K @ 1 Hz	≈ 330 mK @ 1 Hz	≈ 35 mK (model dependent)
Temperature Range	-40°C to 300°C	Wide range (model dependent; typically several hundred °C)	up to ∼330°C (typical for consumer models)
Accuracy	±1°C (typ., platform dependent)	Not specified (low-res, qualitative)	±5°C or ±5% (typ.)
Interface / Supply	I²C, 3.3 V	UART / USB, 5 V	USB / SPI / proprietary (model dependent)
Package/Size	TO39, 9.3mm Dia x 11.25/5.7mm	29mm x 29mm x 13/22mm	10mm x 20mm x 21mm, requires additional interface board
Price	30-40€	80-100€	500-650€
Space heritage	2 years aboard MRC-100 [73]	None found	None found
Notes	Easy to implement onboard PCB	Requires more complicated mechanical implementation	Requires more complicated mechanical implementation, additional interface board

Table 4.1: Comparison of key parameters for three IR sensors.

x-y-plane, with 4 gaps of 15 degrees not being covered. On the other hand, the smaller FoV would similarly provide 4 patches of 35 degree visibility horizontally, resulting in 55 degree gaps in between. To make this decision, the trade off between the FoV coverage over time, considering dead reckoning periods, and the accuracy difference needs to be evaluated.

Attitude accuracy degradation during dead reckoning periods

During the time where no images containing either Earth horizon or Sun are available, the attitude determination effectively becomes dead reckoning, purely integrating the gyro readings. The gyro selected is a recent smartphone-class MEMS sensor, of which a number exist from different manufacturers. These are all fairly equivalent in performance [76] in terms of gyro bias, bias instability, and noise, with sufficient performance for the expected accuracy of the reference vector determination. The largest difference in terms of performance is their stability over changing temperatures. Here, a comparison between three of the leading manufacturers was made as shown in Table 4.3, and the one with the best stability, the ICM45686 by InvenSense, was selected.

To effectively gain an understanding of the possible accuracy of this, the gyro to be used has been analyzed using the Allan variance method. For this, the gyro to be used, an ICM-45686, was configured to acquire data while sitting on a still desk for 8 hours overnight. This data is then analysed to determine the angular random walk and bias instability, which are the key parameters to determine the dead-reckoning errors. The method followed is described in [80].

This results in the Allan deviation plot shown in Figure 4.4, where a fairly common phenomenon for consumer MEMS gyroscopes can be observed, which is the difference in performance between the z

Sensor	Size	Cost	Space heritage	Frame Rate	Resolution	Total
MLX90640	3	3	3	3	1	13
TeraRanger Evo Thermal	1	3	2	1	1	8
Seek Thermal Mosaic Core	2	1	1	1	3	8

Table 4.2: Qualitative trade-off (1 = poor, 3 = excellent) based on size, cost, frame rate, and resolution for selected IR sensors.

Table 4.3: Gyroscope temperature stability characteristics of three manufacturers

Sensor	Invensense ICM45686	Bosch BMI270	STM LSM6DSV
Offset temperature stability [dps/C]	0.005 [77]	0.011 [78]	0.006 [79]

axis and the x-y axis.

Figure 4.4: Allan deviation of ICM-45686

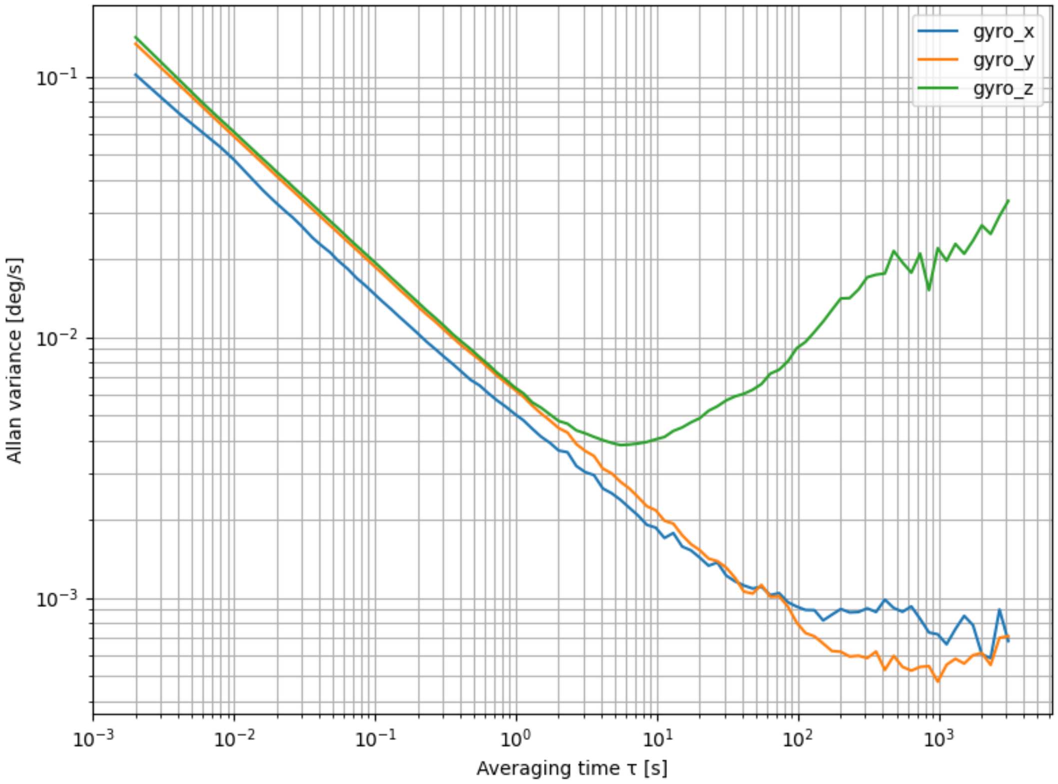

As explained by Woodman [81], the angular random walk and bias instability can be read from this diagram. The angular random walk represents the result from the white noise on the angular velocity readings, which is the slope with a gradient of -0.5, while the bias instability is represented as the flat line at the minimum. Other noise sources exist and are additive, resulting in other factors influencing the actual readings observed on the plot. However, for the purpose of estimating the dead reckoning accuracy, these two characteristics of the Gyro will suffice.

For the ICM-45686, the characteristics are found and presented in Table 4.4. The temperature dependence of these values could not be determined experimentally, as a temperature-controlled chamber for such an experiment is not available. Instead, the typical value of $18\frac{deg}{hr\cdot°C}$ provided by the manufacturer is used [77]. For simulations ,a margin of 50% is added to these, as these values are typical but may be worse.

Using these values, an analysis was performed on the expected attitude error for a given dead-reckoning

Table 4.4: Gyroscope noise characteristics

Parameter	**X-Axis**	**Y-Axis**	**Z-Axis**
Angle Random Walk (ARW)	0.285 [deg/hr$^{1/2}$]	0.343 [deg/hr$^{1/2}$]	0.347 [deg/hr$^{1/2}$]
Bias Instability (BI)	2.11 [deg/hr]	1.72 [deg/hr]	13.8 [deg/hr]

time. For this purpose, not only the gyro errors, but also the algorithmic errors for estimating the absolute attitude based on the imaging sensors, as well as the initial estimation quality of the gyro bias, are included. This initial bias error is especially important, as any bias that is not estimated and compensated will introduce an additional error. Additionally, due to the bias instability, an additional error from the change in bias is introduced [82], such that the error from these sources can be expressed as

$$\theta_E(t) = (b_{E_0} + b_{Instability}) \cdot t \tag{4.1}$$

The errors from the angular random walk are defined as the integration of white noise with zero mean, resulting in the variance given as

$$Var[\theta_E(t)] = \sigma^2_{ARW} * t \tag{4.2}$$

or the standard deviation of the error of the integrated estimated attitude being

$$\sigma_{\theta_E}(t) = \sigma_{ARW} * \sqrt{t} \tag{4.3}$$

An initial attitude error is assumed to correspond to the pixel angular resolution of the sensors, which is approximately 1.5 and 3 degrees for the small and large FoV, respectively. The initial gyro bias estimation error is assumed to be within 0.01deg/s. Based on these assumptions, the expected dead-reckoning error for an initial attitude error of 1.5 degrees is shown in Figure 4.5.

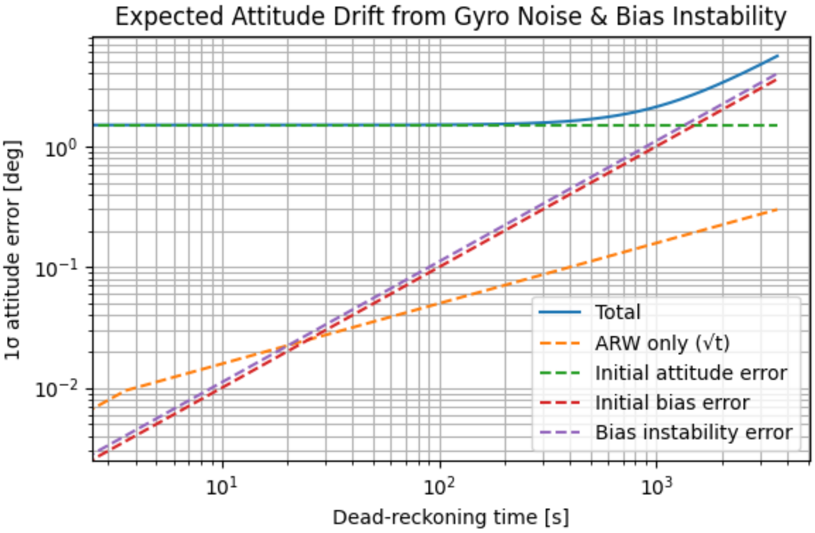

Figure 4.5: Expected attitude drift with initial attitude error of 1.5 degrees. Note that the point where the total expected drift exceeds 3 degrees is after more than 1000 seconds.

It can be seen that the smaller FoV sensor with a better expected initial attitude accuracy exceeds the 3 degree error assumed for the large FoV sensor only after more than 1000 seconds.

A caveat of this analysis is that it is performed only on a single axis, not considering the coupling of axes that occurs in the real dynamics. However, this is acceptable for this case, as all the angles are small, allowing the assumption that the attitude dynamics are linear and independent from each other. To sanity check this analysis, a section of the recorded data, from which the Allan variance analysis is performed, is used to integrate the attitude over time. However, the initial condition for the analysis is that the bias has been estimated using the sensor fusion. Instead, the bias of the three axis are estimated simply by taking the mean over a section of the acquired data. Following that, the attitude is integrated numerically over a different section of data. The resulting angular error is shown in Figure 4.6.

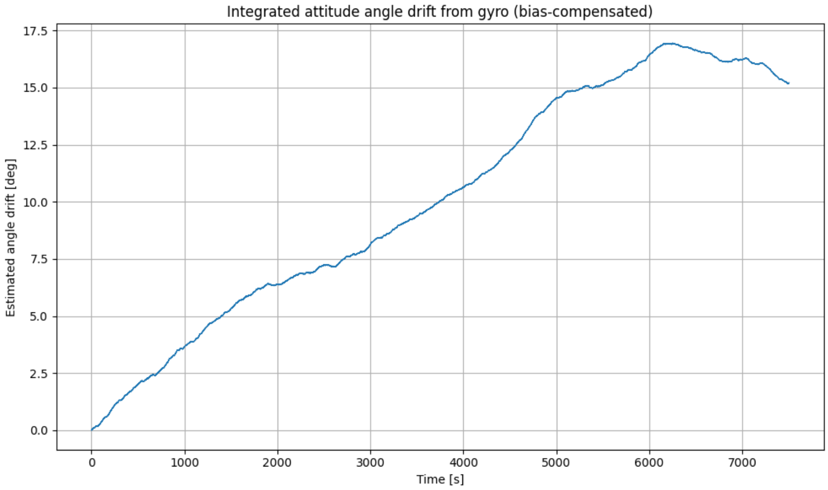

Figure 4.6: Real attitude drift from Allan variance dataset, bias estimated from a section of data not included in this timespan.

Additionally, this does not take into account the temperature variations expected during orbital conditions, and the corresponding change in bias due to that, as the underlying data was collected in an approximately isothermal indoor environment. With the worst case temperature gradients assumed to be around 20deg/hr based on similar missions [83] the expected temperature drift is added to the expected attitude drift in Figure 4.7, which shows that the temperature dependent drift dominates over the other error sources.

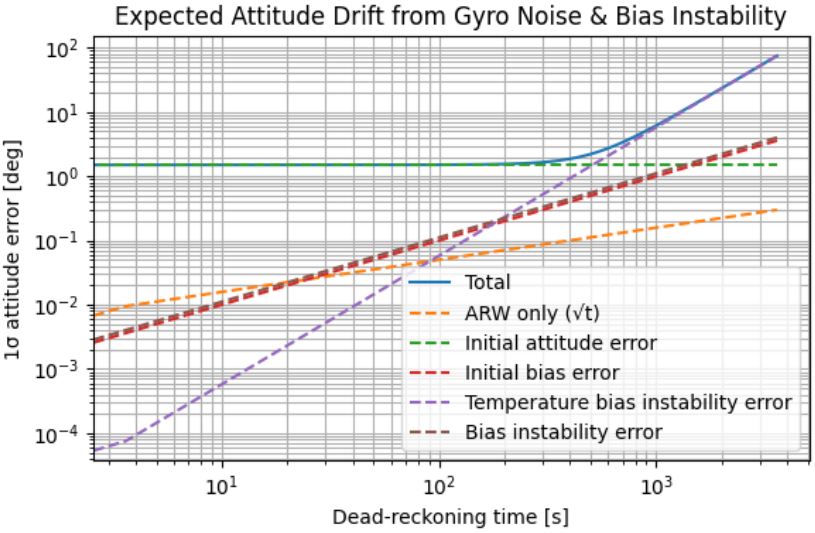

Figure 4.7: Expected attitude drift including temperature dependent bias instability at the maximum expected temperature gradient.

It is expected that the rotation rate of the satellite is approximately 10 deg/s, thus the horizon will be in view fairly often compared to these timescales. The only exception is when this axis of rotation is aligned such that it points towards the Earth within a certain tolerance, at which point there will be a larger section of the orbit where no horizon will be in view. This direction of the rotation axis where no horizon is visible has a half angle of 62.5 degrees.

With an orbit time of approximately 90 minutes or 5400s, the attitude relative to the Earth will change by 0.066deg/s. Thus, the time to traverse one of these regions within the orbit where no horizon can be seen is 1875s, which is the very worst case scenario, where the rotation axis is aligned in precisely the wrong direction. Consequently, the sensor with the lower FoV is selected.

4.2. Sensor calibration and distortion correction

As all the calculations to follow on from the horizon detection, Sun detection, and the translation of these detected objects from image space into the real world are based on a pinhole camera model, a method to determine possible deviations of the camera and correct for these is required.

For this purpose, first, the simple pinhole camera model is introduced, as well as a modification of this model to allow tailoring it to any possible deviations in the real camera. To evaluate and determine the required parameters, a process to acquire images using the cameras was built using facilities provided by the GNC Lab at TU Delft's aerospace faculty. Using this data, it was evaluated whether calibration was required, and if so, the parameters were determined using a genetic algorithm to optimize the parameters to the acquired dataset.

4.2.1. Pinhole Camera concept and calibration projections

For cameras with relatively low fields of view (below 180 degrees), the pinhole camera model can be used to describe the relationship between the location of an object in the world and where it appears on the image. To illustrate the principles of the camera model, consider Figure 4.8.

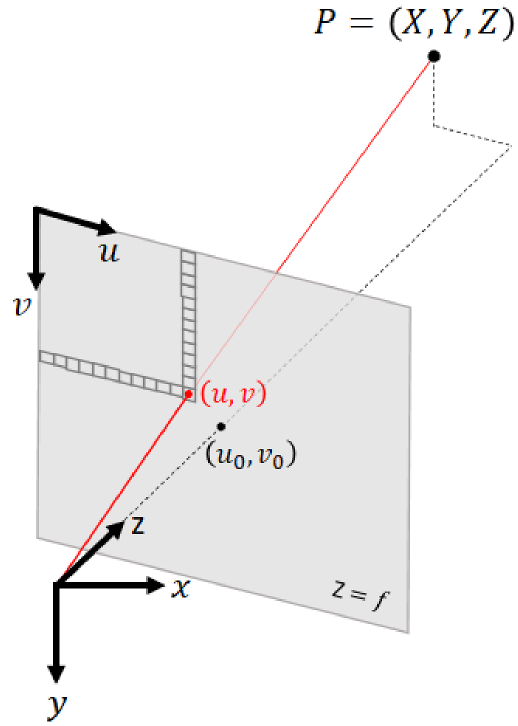

Figure 4.8: Illustration of the pinhole camera model, relating the position of point P in the world to the location (u, v) on the image plane. Illustration taken from [84]

The location of an object at the point $P = \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}$ can be related to the pixel where it appears on the image plane $p = \begin{pmatrix} u \\ v \end{pmatrix}$ on the image plane. This is done by moving the image plane coordinates to the homogeneous coordinates, resulting in the following equation.

$$pZ = KP \tag{4.4}$$

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} Z = \begin{bmatrix} f_x & 0 & u_0 \\ 0 & f_y & v_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{4.5}$$

Here, K is called the intrinsic matrix, describing the properties of the camera, containing the focal lengths in x and y, as well as the coordinates u_0 and v_0 of the principal point of the camera. Following from this, it is possible to calculate the pixel coordinates as follows:

$$\begin{aligned} u &= f_x \frac{X}{Z} + u_o \\ v &= f_y \frac{Y}{Z} + v_o \end{aligned} \tag{4.6}$$

All of these parameters can be determined from the datasheet, as the focal lengths are given indirectly in the datasheet of the chosen sensor through the field of view. The field of view, focal length, and the size of the image plane are related by the following equation.

$$FoV = 2\arctan(\frac{H}{2f}) \tag{4.7}$$

As both the focal length and the size of the image sensor are unknown, it is possible to define the image plane size to be of unit size, arriving at a focal length. However, due to the actual sensor not being an ideal pinhole camera but instead containing a finite sized sensor as well as lenses, distortions will occur that will change the location of objects in the image plane compared to their expected location to some degree.

For this purpose, Brown introduced a model to compensate for radial and tangential distortions [85], also called the Brown-Conrady model, based on earlier work by Conrady [86]. In this model, the corrected and ideal coordinates on the image plane are related by the following equations, where r is the radius on the image calculated as $r = \sqrt{x^2 + y^2}$

$$\begin{aligned} x_{\text{corrected}} &= x_{\text{ideal}} \left(1 + k_1 r^2 + k_2 r^4\right) + 2p_1 xy + p_2 \left(r^2 + 2x^2\right) \\ y_{\text{corrected}} &= y_{\text{ideal}} \left(1 + k_1 r^2 + k_2 r^4\right) + p_1 \left(r^2 + 2y^2\right) + 2p_2 xy \end{aligned} \tag{4.8}$$

Here the terms p_1 and p_2 specify the size of the tangential distortion, skewing the image. On the other hand, the coefficients k_1 and k_2 specify the radial distortion, which results in the image to bulge in- or outwards, also called barrel or pincushion distortion [87]. Identifying these parameters for the sensor allows the determination of the location of the objects in the world coordinates imaged on the image plane.

4.2.2. Setup to acquire data

The process of determining these parameters requires images on which an object at known locations relative to the sensor is imaged. Thus, it then becomes possible to compare the locations of these objects in the image plane to where they are expected to be based on the Brown-Conrady model. With this relationship, the coefficients can be varied until the deviations between the expected locations and acquired locations are minimized.

Usually, this is done for conventional cameras by taking images of checkerboard patterns and finding the parameters that undistort the image until the checkerboard pattern is as straight and undistorted as possible.

However, for the IR sensor, it is impractical to create a checkerboard pattern that is visible in the wavelengths relevant for the MLX90640. Instead, inspired by [38], the compositing of multiple images, which are acquired while varying the angle of the camera precisely, can achieve the same checkerboard pattern over multiple images.

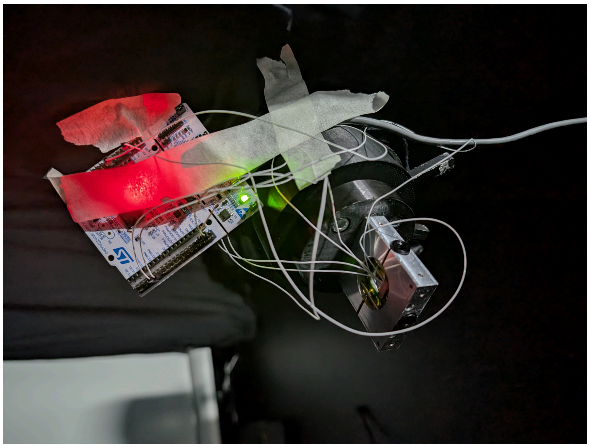

Figure 4.9: Setup of the sensor at the end of the robotic arm together with an external data acquisition board.

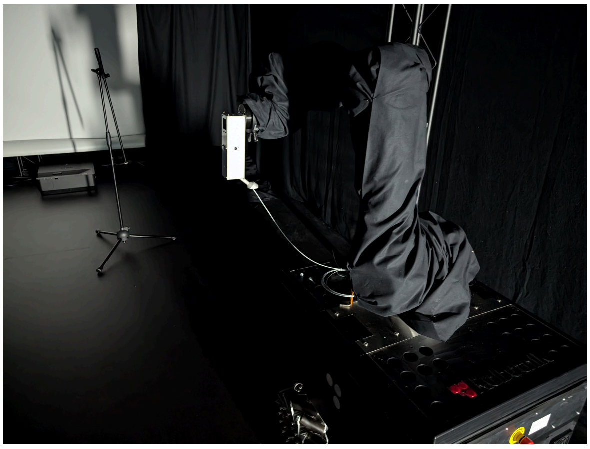

Figure 4.10: Overview of robotic arm showing representative payload with tripod in background to which the soldering iron is mounted.

For this purpose, a prototype of the sensor is attached to a robotic arm, which can precisely position the sensor. In about 1.5m distance from the location of the sensor, a soldering iron is positioned to act as a point source.

The sensor is rotated such that the point source is positioned across the entire image. This setup is shown in Figure 4.9 and Figure 4.10. As the inverse kinematics did not function properly when the data was acquired, instead the arm was positioned so that two of the six joints available are positioned such that they respectively exclusively rotate around the pitch and yaw axes of the sensor. A problem introduced through this method is that the center of rotation of these joints is offset from the center by a few tens of centimeters, resulting in a small amount of translation as well. However, this is only very small amounts on the order of centimeters compared to the approximately 1.5m distance to the soldering iron

On the resulting images, the point source is identified using the DAOPHOT algorithm [48] implemented in the photutils python package [88]. This results in a dataset where a large number of points are given across the entire image, each associated with a known angular position of the sensor. From this sensor location, together with the determined location of the soldering iron, it becomes possible to optimize the parameters k_1 and k_2 as well as p_1 and p_2 such that the distance between the expected location of all points and their real location in the image is minimized.

4.2.3. Calibration determination

A sample of the acquired images is shown in Figure 4.12 and Figure 4.11. Some images, like the one in Figure 4.12, were not of high enough quality for a hotspot to be clearly identified in the data. These are discarded, however, there is still plenty of imagery across the entire sensor area, allowing the calibration to be performed correctly.

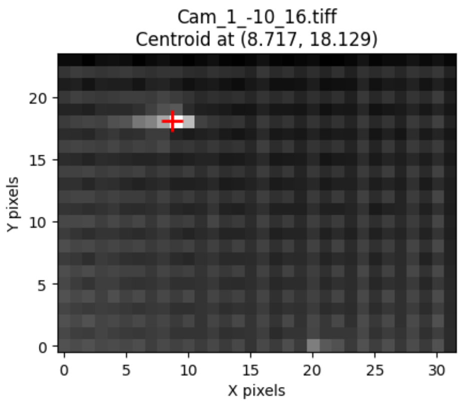

Figure 4.11: Centroid for soldering iron found

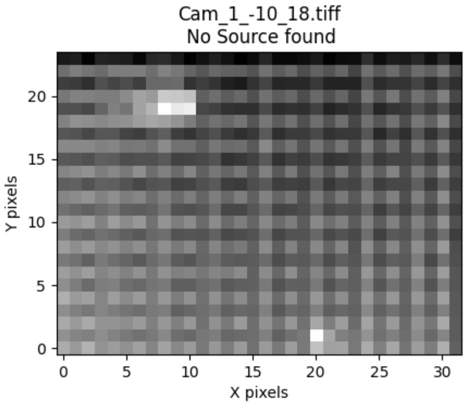

Figure 4.12: No Centroid found, as low contrast and background noise is too large

The collected points are shown in Figure 4.13.

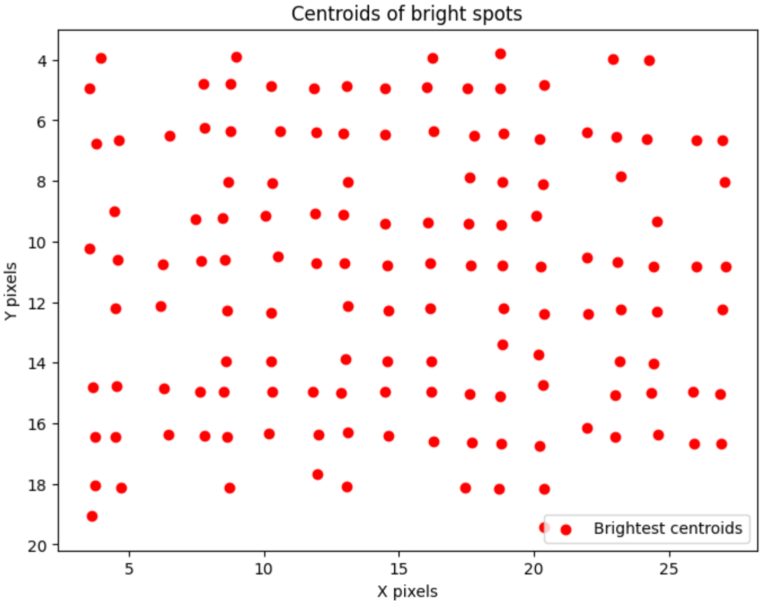

Figure 4.13: Points recorded on a sensor from the soldering iron, each with a 2 degree angular spacing from each other. Notice that gaps exist, where the recorded imagery did not give a sufficient image quality to determine an accurate brightspot estimate.

Plotting the expected locations of the bright spots compared to the actual detected locations is shown in Figure 4.14. These are determined by using an ideal pinhole camera model using the FoV specified by the datasheet of the device, determining where a point source directly in front of the sensor would land, given the same angular displacements as the ones performed by the robotic arm. It is immediately noticeable that the expected and detected points in the image are far offset from each other. This is due to the fact that the soldering iron acting as a target is misaligned in relation to the sensor being moved at its initial position. Additionally, it becomes obvious that on the x-axis, the spread of the expected points is smaller than the recorded ones, indicating that the real FoV is mismatched compared to the datasheet.

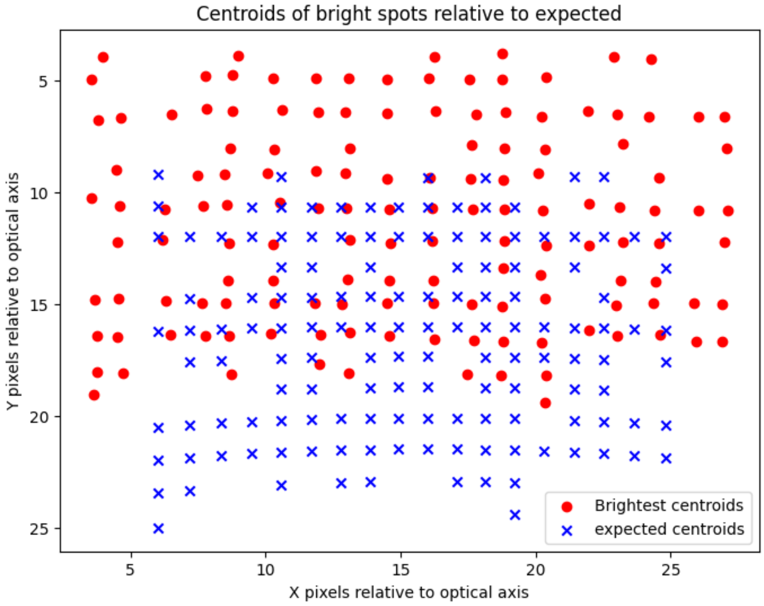

Figure 4.14: Acquired centroids of points in image compared to projected centroids using the expected pinhole projection function.

To correct for the lateral and vertical shifts between these two groups of points, the point source is offset from the optical axis, and the fields of view are adjusted accordingly. This initial adjustment is performed heuristically, still assuming an ideal pinhole camera model. Using this method, the calibration result shown in Figure 4.15 is obtained. The calibrated field of view is approximately 45×33 degrees and shows noticeable discrepancies compared to the specified values.

In pixel units, this corresponds to focal lengths of 38.6 and 40.5 pixels along the x and y axes, respectively. This heuristic calibration yields a mean squared error of about 0.23 pixels. At the image center, this error combined with the smaller focal length of 38.6 pixels corresponds to an angular error given by

$$\theta = \arctan\left(\frac{x}{f_x}\right) \tag{4.9}$$

$$\sqrt{0.23}, ^\circ \frac{d\theta}{dx} = 0.48, ^\circ \cdot \frac{f}{f^2 + x^2} \approx 0.71, ^\circ \tag{4.10}$$

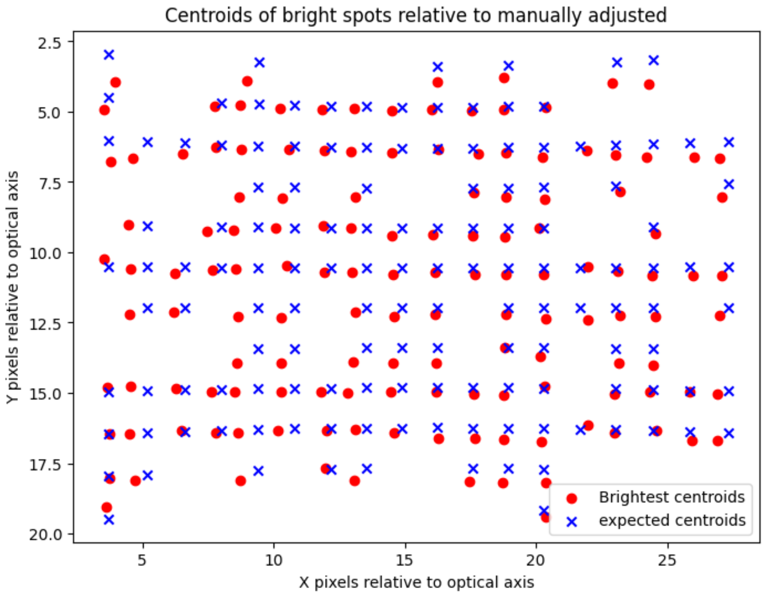

Figure 4.15: Acquired centroids of points in image compared to projected centroids using the a heuristically determined pinhole projection function with object offset.

However, this still assumes an ideal pinhole camera model, which leaves more optimisation possibilities. Using a genetic algorithm as proposed in [38], both the location offset of the point source, the focal lengths as well as the parameters k_1, k_2, p_1, p_2 are optimized with the goal of reducing the mean squared error as far as possible. To speed up the convergence, the initially manually tuned parameters for source location and focal lengths are used as an initial value, with the parameter ranges being fairly tightly constricted to a range around them. As described in [38], it was found that this type of IR sensor by this manufacturer has very similar projection parameters between different cameras of the same model. Hence, only one projection function is determined, which is used for all of the sensors used for flight. This would simplify the calibration of the system dramatically, as no calibration is required at all after one camera has been characterised. However, whether this actually holds for this specific model of the IR camera should be confirmed with a similar approach being performed for a sample of these cameras.

4.3. Simulated Image generation

As it is difficult to replicate the exact images that the sensor would acquire in orbit, a framework to generate simulated data is created. The MLX90640 provides a 32×24 pixel array of temperature readings, corresponding to a low-resolution thermal image. Direct collection of large training datasets in realistic scenarios is challenging, therefore, simulation was chosen as a complementary approach. For this purpose, a raytracing framework was used, coupled with the simulation of the positions and orientations of a satellite together with the positions of the Sun and Earth.

4.3.1. Scene creation

To be able to create any images, the scene to be imaged must be known. For this purpose, a number of variables need to be determined:

1. The location and orientation of the sensor, or more accurately, the direction of the principal optical axis of the camera
2. The locations of the various celestial objects, such as the Earth and Sun.
3. The amount of infrared light emitted by the various celestial objects.

Sensor location and orientation

To determine the location of the sensor, the approximate orbit the satellite is expected to be positioned in is used to determine a feasible location. For the orientation, any random orientation is taken.

Sun and Earth location

The Sun and Earth move about in the solar system in predictable ways. Thus, at any given time it is possible to determine their location if the time is known. Using the python package astropy's ephemerides functions [89] based on the Essential Routines for Fundamental Astronomy (ERFA) [90], which can later also be used to perform this computation on board the processor of the sensor, the location of the celestial objects are determined.

Infrared light radiation per object in scene

The intensity of the light coming off each celestial object is approximated here by its temperature to the power 4, as the infrared emissions follow the Stefan-Boltzmann law of thermal radiation. As the sensor is specified only for a target temperature range from -40 to 300 deg C, the outputs are rescaled such that the thermal emissions that would correspond to temperatures higher than 300 C are imaged as the maximum brightness in the image. One caveat for this adjustment is the fact that the Sun is far hotter than this, while being smaller than a single pixel. To counteract this, the rescaling is only done after anti-aliasing resolves this ambiguity, which will be explained in more detail in the following section.

The Sun's temperature is taken as the photospheric effective temperature of approximately 5778 K ($\approx 5505°$C) [91]. Because this is many times larger than the MLX90640's measurable range (sensor range: -40 to $300°$C ≈ 233.15–573.15 K) the raw solar radiance would otherwise dominate and saturate simulated pixels.

In the renderer, the Sun is modeled as a small emitting disk with the correct angular size rather than a uniform scene emitter. Radiance from each surface is computed using Planck's law and integrated over the MLX90640 spectral band (band-limited blackbody radiance), taking the sensor spectral response into account. To correctly capture a solar disk that is smaller than a pixel, per-pixel supersampling (anti-aliasing) is performed such that the solar contribution is resolved across sub-pixel samples instead of being treated as a single impulse.

For the Earth, a temperature of 0 degrees C on the non-illuminated side, and a temperature of 20 deg C is used, as the surface temperature can be taken as a reasonable estimate for the captured temperature from space, due to the largely IR transparent atmosphere of Earth. Only a few absorption bands of the atmosphere are present in the IR wavelengths, such as those of Ozone and CO2 [92]. However, these are fairly narrow compared to the typical microbolometer sensitivity bands, such that the surface temperature can be assumed to be seen from space.

4.3.2. Ray Tracing Framework

Image generation was implemented using the ray tracing framework from the open-source "Ray Tracing in One Weekend" guide by Shirley [93].

First, the scene was created as three spheres, corresponding to the locations of the Sun and Earth, obtained as previously explained. Afterwards, the location of the sensor, as well as the orientation of the optical principal axis, is initialised. Following this, the image plane is constructed, which is at a unit distance in front of the sensor. As explained in subsection 4.2.1, this results in an image size corresponding to the specified field-of-view of the camera. Following that, the centers of each pixel are identified. For each pixel, a ray is sent from the camera origin towards this pixel and the first object hit determines the value this pixel will display in the final image. For this, we construct the equation of that ray as

$$\mathbf{P} = \mathbf{Q} + t\mathbf{d} \tag{4.11}$$

while the equation of each celestial sphere is

$$(\mathbf{C} - \mathbf{P}) \cdot (\mathbf{C} - \mathbf{P}) = r^2 \tag{4.12}$$

Filling Equation 4.11 into Equation 4.12, it is possible to solve for t. If a solution exists, that pixel's ray intersects that sphere at the smallest positive t that solves the equation. At this point, that pixel could simply be assigned the temperature of the sphere associated with the smallest t which is important when one body occludes another. However, the Earth has differences in temperature depending on its illumination. For this reason, the value of the IR intensity to be assigned in the case of Earth is calculated based on the angle between the location of the hit on the sphere and the Sun.

For this, the surface normal vector at the intersection point is taken to be

$$I_{ray} = \begin{cases} I_{eclipse} + (I_{hot} - I_{eclipse}) \cdot cos(\frac{\mathbf{n} \cdot \mathbf{s}}{|\mathbf{n}||\mathbf{s}|}) & \text{if} cos(\frac{\mathbf{n} \cdot \mathbf{s}}{|\mathbf{n}||\mathbf{s}|}) > 0 \\ I_{eclipse} & \text{otherwise} \end{cases} \tag{4.13}$$

This yields a functional raytracer that can create simulated images.

However, as previously mentioned, for subpixel features such as the Sun or edges of the horizon, this will not suffice, as the one ray per pixel may narrowly miss an object. In reality, each pixel is not representative of just the center of its area, but acts more like an integration of all the light that is present in the section of the image it covers. Thus, more rays within the same pixel can be used. All of these go through the same process to acquire an intensity of the light, which is averaged for the final value of the pixel. This process is called antialiasing and results in softer edges, as well as the resolution of subpixel objects such as the Sun. This method results in images such as shown in Figure 4.16.

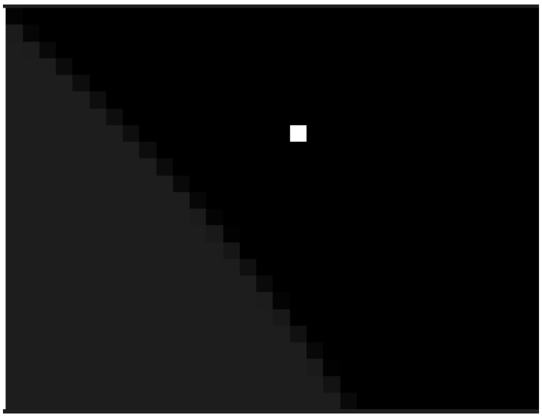

Figure 4.16: Example image generated using the raytracing program, showing the Earth horizon as well as the Sun

This is an ideal picture the sensor would acquire if the sensors were perfect. Some noise is to be expected on the images acquired, which is specified by the sensor in temperature Root mean square (RMS). This noise depends not only on the acquisition rate, but also on the specific pixel under consideration. This is shown in Figure 4.17, which can be used to generate realistic noise to be added on top of the generated images.

For each pixel, the equivalent temperature is calculated by taking the following formula, where r specifies the distance from the center pixel:

$$I_{noisy} = (\sqrt[4]{I_{pixel}} + T_{noise})^4 \tag{4.14}$$

where

$$\begin{aligned} T_{noise} &\sim \mathcal{N}(\mu, \sigma(r)^2) \\ \sigma(r) = \sigma_{min} &+ (\sigma_{max} - \sigma_{min}) * \frac{r}{r_{max}} \end{aligned} \tag{4.15}$$

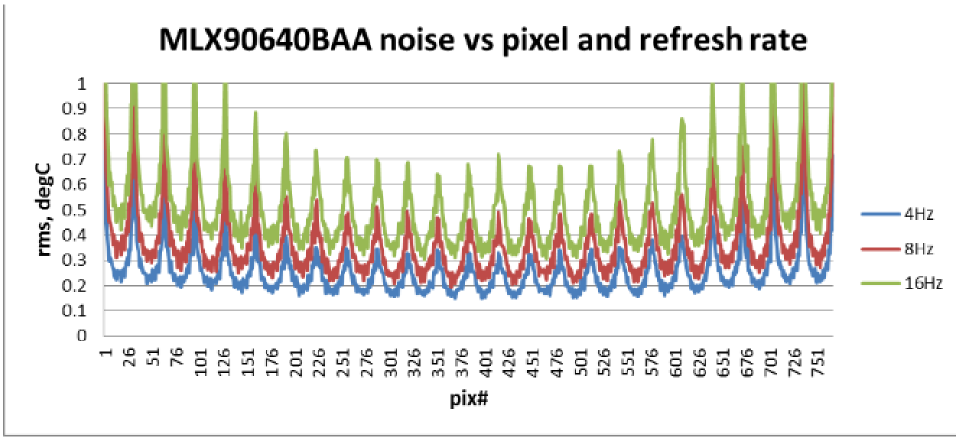

Figure 4.17: Noise performance of MLX90640 at different acquisition rates and pixel positions. Each line of pixels is the distance between the peaks, as the noise is larger near the edges of the sensor [72]

4.3.3. Use in Algorithm Development

The synthetic dataset allows controlled testing of orientation estimation algorithms under varying conditions (e.g., object temperature, geometry, and background). By providing known ground truth for camera pose and scene configuration, the simulated images facilitate quantitative evaluation that complements experimental data collected with the real sensor.

4.4. Image processing of IR Images

This section will treat the process of going from a raw image to determining the two vectors pointing towards the Sun and nadir to be used for attitude estimation.As an example a raw simulated image is used, shown in Figure 4.18.

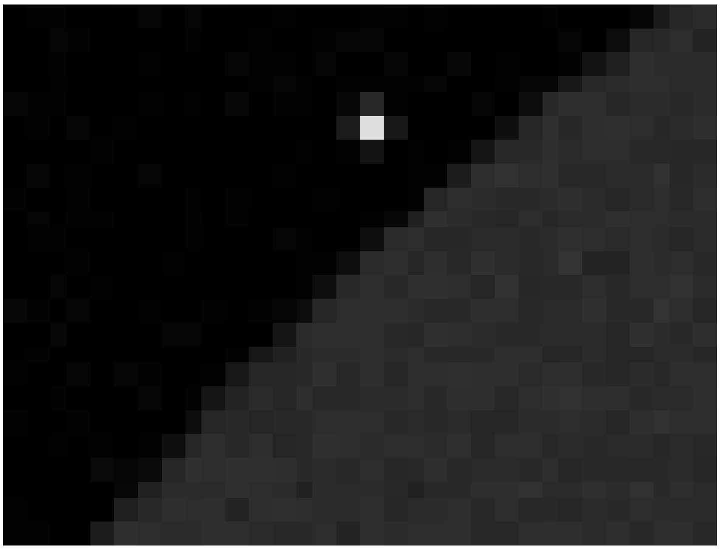

Figure 4.18: Initial example image generated with raytracing simulator.

4.4.1. Classification of present objects

Before the image processing is performed, a process is performed to detect whether the Sun or horizon or both are in the image. This process is somewhat coupled and computations are reused between the detection of the two objects, as well as the later direction finding for the Sun or NADIR vector. As the only objects that realistically can be recorded in orbit are those that we wish to classify, a fairly simple thresholding based process can be used.

Earth horizon presence detection

The first step is to perform an adaptive Otsu threshold operation on the image, as shown in Figure 4.19. The resulting threshold must be larger than the image background with some margin that is picked experimentally. Additionally, an effectiveness metric of the threshold algorithm is used, which specifies how closely the intensity histogram of the input image matches a purely bimodal distribution. This bimodality is additionally checked to be above a certain threshold for Earth presence, to ensure that approximately two types of areas are present in the image: background and the Earth.

If those conditions are matched, it was found that within the simulated images, the algorithm reliably determines the presence of the Earth within an image while being very simple as well as computationally efficient. However, this still provides a positive match if only small slivers of the Earth along one edge or corner are present. Especially in the case of the corners, it is challenging to correctly identify the NADIR vector, which is why a further check is performed: a minimum number of pixels must always be present that show the Earth and background, which guarantees that a large enough section of the horizon for attitude estimation is available. This is done by checking that the binary image output of the Otsu threshold algorithm has a minimum number of pixels for each of the two binary values.

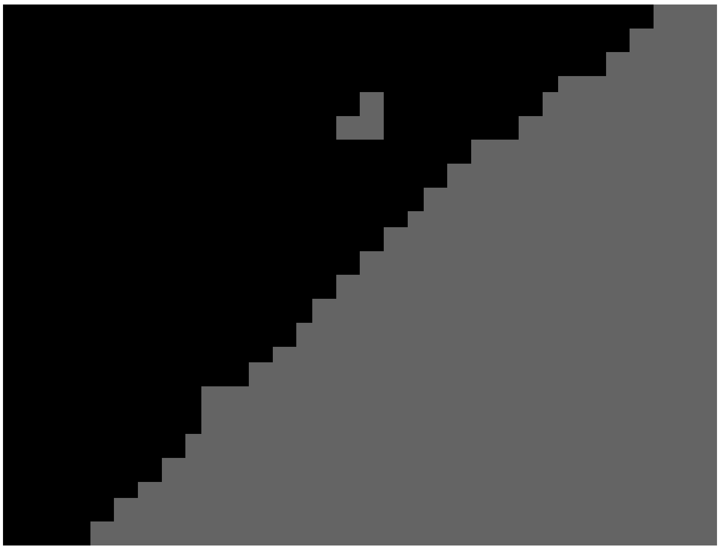

Figure 4.19: Otsu threshold performed on image

Sun presence detection

Two independent detection methods are implemented, and the final Sun detection result is obtained as the logical OR of their outputs. That is, if either method detects the Sun, the detection is considered positive.

The first method is the simplest method and is based on the maximum intensity recorded in the image. If this exceeds a threshold, which must be brighter than any of the other expected objects to be imaged, it is determined that the Sun is present. This relies on the fact that the Sun is the brightest object the sensor will experience by a large margin.

The other method is more involved and uses the same Otsu threshold as the Earth horizon detection as a basis. In fact, the same Otsu threshold and effectiveness thresholds must be met as for the Earth horizon limb, additionally, it is checked whether the bright pixels are in the vicinity of the borders of the image. This effectively prevents a positive Sun detection when the Earth is visible on only a small area of the image, as that small sliver of the Earth must be close to the edges, while the small area of the Sun can occupy any area of the image. This will only function if the Sun is visible in the image by itself, without the Earth being visible anywhere in the same capture. However, ideally, the Sun intensity threshold will perform the detection in most cases, with the second method being a backup in case the thresholds are set incorrectly.

Finally, if any of these conditions are true, some more computationally intensive processing is performed to determine whether the detected object matches the expected signature of the Sun. Using MATLAB's bwlabel, all connected regions in the thresholded image are identified. Any region larger than a certain number of pixels (it was found around 8 worked well for this case) is discarded. If any areas remain, this indicates the Sun is present.

4.4.2. Sun location detection

Once the Sun's presence is confirmed, the previously computed distinct areas within the thresholded images are used to isolate the region that represents the Sun. This is done by comparing the mean intensity in the original captured image within each region remaining. The Sun is expected to be by far the brightest object to be acquired, thus the region with the highest mean intensity is chosen. Finally, the specific location on the image plane in pixel coordinates is computed as the centroid of the selected region, weighted by the pixel intensity of the IR image within the region.

Because the Sun appears as a compact, high-intensity blob in the IR spectrum, this centroid-based localization is both robust and computationally efficient. It avoids the need for explicit shape fitting and performs reliably even when partial saturation occurs.

Using the camera model developed previously, it is possible to easily construct the vector from the camera's origin to the Sun. This will then be fed into the attitude estimation algorithms. This vector is always given in the local camera reference frame, and thus needs to be moved from the camera reference frame into a singular unified reference frame in which the attitude estimation can be performed.

4.4.3. Horizon location detection

Detection and modeling of the Earth's horizon require a more elaborate procedure: before proceeding with any edge detection, it was found that the Otsu threshold creates artifacts around the edge. Additionally, the Sun is not reliably detected as its own category, due to its small size. Thus, to remove those small outliers, a morphological transform is applied, which first erodes regions and thus removes all pixels that are not connected to larger geometry, followed by a dilation, which reverses this operation for pixels that survived the erosion. This way, only larger geometries, such as the Earth, remain in the image.

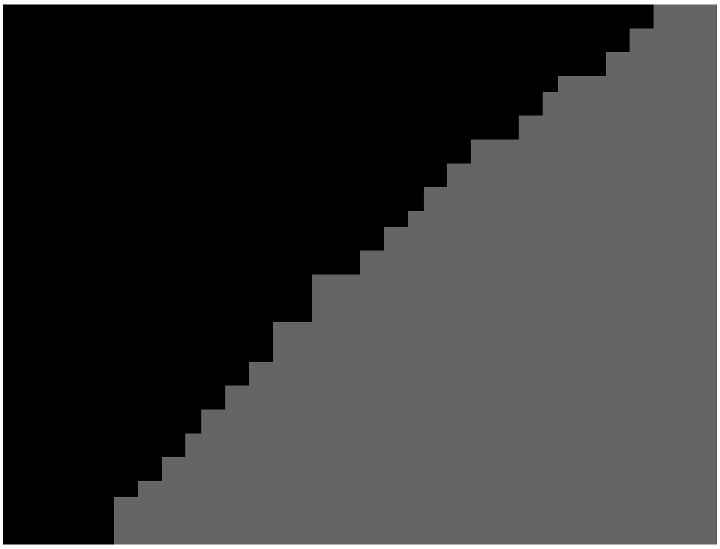

Figure 4.20: Removal of artifacts and Sun via morphological operation

As seen in Figure 4.20, this creates an image with only the Earth and space visible, with as much noise removed as possible. A canny edge detection is applied to this, revealing the horizon line as the only remaining geometry, as shown in Figure 4.21.

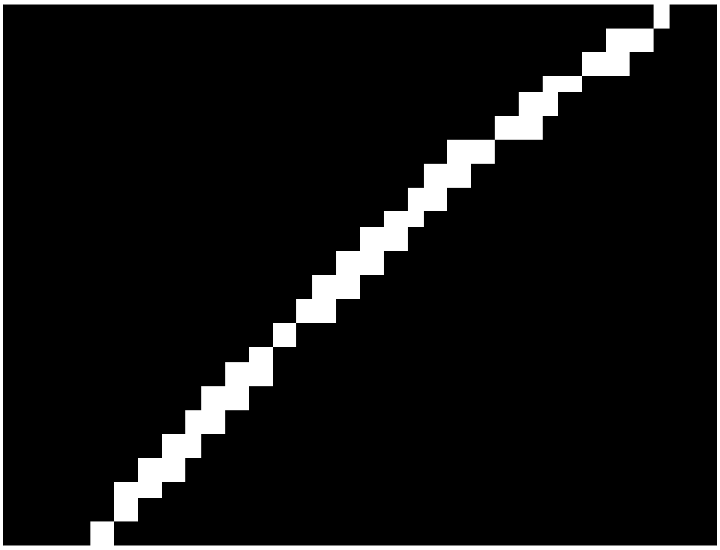

Figure 4.21: Remaining horizon line after Canny edge detection algorithm

Once the pixels that lie on the horizon line are determined, two approaches exist to determine the NADIR vector from it:

1. Fitting of an ellipse
2. Direct least squares computation of the nadir vector (also known as the Christian-Robertson algorithm)

Ellipse fitting

On this horizon line, an ellipse can be fit. An ellipse is a suitable geometric model because, under typical near-space viewing conditions, the Earth's curvature projects approximately as an ellipse in the image plane. The fitting process uses a RANSAC-based approach to ensure robustness against noise and spurious edges [40]. In each iteration, five random edge points are sampled, and the coefficients of the general conic equation

$$Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0 \tag{4.16}$$

are computed algebraically. The model's inliers are determined based on the algebraic residual between the predicted and observed points. The model with the highest inlier count is selected as the best-fitting ellipse. Alternatively, a constrained algebraic fit can be performed over all detected edge points to refine the coefficients. The fitted ellipse provides parameters such as the center, orientation, and curvature radius of the horizon, which can be mapped to the camera's attitude relative to the Earth's limb. Visual verification is performed by overlaying the fitted ellipse on the binary horizon mask, confirming accurate alignment between the detected edge and the fitted model. However, due to the low resolution and low field of view, combined with the low altitude, this method yielded ellipses that rarely if ever gave reasonable nadir vectors. This is due to only a small section of the ellipse being visible in each image. Thus, the following algorithm was favored.

Christian-Robertson algorithm

It has been shown that the Christian-Robertson algorithm performs better than direct ellipse fitting and additionally, is computationally more effective [94]. This approach creates a system of equations, with one equation for each point on the horizon. Each detected horizon pixel (x_i, y_i) is mapped to a camera-frame direction vector by inverting the camera intrinsic matrix,

$$\mathbf{p}_{c,i} = \mathbf{K}^{-1} \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix}, \tag{4.17}$$

where $\mathbf{K}$ denotes the camera intrinsic matrix. These vectors represent viewing rays originating from the camera center. The Earth horizon defines a right circular cone in the camera frame, whose axis aligns with the camera-frame nadir vector $\hat{\mathbf{e}}_c$ and whose half-angle α is determined by the spacecraft altitude,

$$\alpha = \arcsin\left(\frac{R_\oplus + d_{\text{atm}}}{R_\oplus + h}\right), \tag{4.18}$$

where $R_\oplus$ is the Earth radius, d_{atm} is an effective atmospheric offset, and h is the spacecraft altitude.

Hence, the spacecraft's altitude must be known a priori for this algorithm to perform well. Conversely, this algorithm is frequently used to estimate the inverse as well, determining the location relative to the body of known diameter if the attitude is known from separate sources.

For an ideal horizon pixel, the corresponding ray lies tangent to the Earth's surface, yielding the constraint

$$\hat{\mathbf{e}}_c^{\mathsf{T}} \mathbf{p}_{c,i} = \cos(\alpha)\,\|\mathbf{p}_{c,i}\|. \tag{4.19}$$

Collecting this constraint for all detected horizon pixels leads to an overdetermined linear system, which can be solved in a least-squares sense to estimate the nadir vector $\hat{\mathbf{e}}_c$. The resulting estimate is normalized to unit length. In practice, horizon edge detection may introduce outliers due to clouds, terrain features, image noise, or partial occlusions. These outliers violate the ideal cone constraint in Eq. (4.19) and can significantly degrade the least-squares estimate. To improve robustness, a RANSAC-based scheme is applied. In each iteration, a minimal subset of three horizon pixels is randomly selected and used to compute a candidate nadir vector by solving Eq. (4.19). Given this candidate estimate $\hat{\mathbf{e}}_c^{(k)}$, all horizon pixels are evaluated using the residual

$$r_i^{(k)} = \left|\hat{\mathbf{e}}_c^{(k)\mathsf{T}} \mathbf{p}_{c,i} - \cos(\alpha)\,\|\mathbf{p}_{c,i}\|\right|. \tag{4.20}$$

Pixels with residuals below a predefined threshold are classified as inliers. The candidate solution producing the largest inlier set is retained. After completion of the RANSAC iterations, a final least-squares nadir estimate is computed using only the inlier pixels associated with the best candidate. This was found to work reliably and accurately.

4.4.4. Reference frame alignment

The Sun and nadir vector are specifically only determined within the reference frame of the camera, due to a peculiarity of the MLX90640 camera output format, some manipulation is required to transform the vectors into the body frame. The primary change that needs to be performed is the transformation of the reference frame into a right-handed coordinate system, as the convention of the MLX90640 by default is that the first pixel is in the top-right corner of the image. Thus, with the convention that the optical axis is pointing out of the image and column-major ordering, the initial output is a left-handed coordinate system. To rectify that, the image is mirrored horizontally, such that the first pixel is at the top-left corner of the image. Following that, the determined vectors are rotated into the body frame, using simply the rotation matrices determined for each camera, depending on their specific mounting location.

4.5. Fusing of image and gyro observations

The previous chapters merely described the process of determining the direction of vectors within the camera reference frame, pointing towards either the Sun or Earth. However, from this point to establishing uninterrupted, accurate attitude knowledge within a relevant reference frame for a satellite, a number of steps need to be taken. This chapter will treat this process, starting from how the reference vectors are related to the attitude of the spacecraft, followed by the fusion of these measurements with the output of a MEMS gyroscope to allow both the uninterrupted knowledge of the attitude, as well as eliminating noise in the measurements when specific camera conditions occur.

4.5.1. Relation of reference vectors to attitude with satellite location knowledge

Solely being able to acquire a reference vector towards the Earth and/or Sun does not suffice to be able to fully determine the attitude of the spacecraft in an Earth-centered reference frame. For this purpose, the location of the satellite relative to the reference bodies is required. For the nadir vector, the only required datapoint is the current location of the satellite, which can easily be obtained by an on-board GNSS receiver, which at the same time also can provide precise timestamping. The true nadir vector in the body frame is then simply the unit vector pointing in the direction of the vector from the spacecraft towards the origin for any Earth-centered reference frame. The timestamp is important when considering the Sun ground-truth vector: for that vector, it is necessary to have knowledge of the relative position of the Sun to the spacecraft, which can be obtained if, as an intermediate, the relative positions of the Earth and Sun are known together with the GNSS-provided spacecraft location.

Thus, on board the satellite, the GNSS position and time solution is required, as well as ephemeris data that provides the Sun location in the Earth-centered reference frame. The accuracy of the position and time fix is directly related to the resulting attitude determination fix. For the Sun vector, any time error is important, while for the nadir vector, it is any position error. For the Sun vector, due to the use of the Earth-centered inertial reference frame, the angular errors are related to the fraction of the time error of a year. As an example, a time-error of one second is a negligible angular error of

$$\delta\theta = \frac{1s}{365 * 24 * 60 * 60s} * 360deg = 0.04arcsec \tag{4.21}$$

As the distance between the Sun and the spacecraft is so large, the positional error for the Sun is negligible as well, with an offset of 1km resulting in an angular error of

$$\delta\theta = arctan(\frac{1km}{1AU}) = 0.0014arcsec \tag{4.22}$$

In contrast, the nadir vector has non-negligible errors from time and position errors. Instead of a year being the time for one revolution, it is the orbital period of the satellite, and the distance between the center of the Earth and the spacecraft is far smaller. Here, the same time and positional errors (for an approximately 90-minute orbit) work out to 4 arcmin and 30 arcsec, respectively.

4.5.2. Kalman filter for sensor fusion

For the sensor fusion and attitude determination between the availability of reference vectors from the IR-sensors, an extended error state Kalman filter using a quaternion attitude representation is used. A conventional Kalman filter cannot be directly applied to attitude estimation due to the nonlinear structure of the rotation group. While linearization about the current state is possible in principle, difficulties arise when using quaternions as the attitude representation. The Kalman filter formulation relies on vector-space operations, in particular addition, whereas unit quaternions representing rotations do not form a vector space.

Rotation quaternions form a group under multiplication, not addition. Applying additive corrections to quaternion states, as done in a conventional Kalman filter, generally produces vectors that no longer satisfy the unit-norm constraint and therefore no longer represent valid rotations. A commonly used workaround is to renormalize the quaternion after each update; however, this approach is mathematically inconsistent and can introduce undesirable estimation artifacts.

Alternatively, an error-state formulation can be used, in which the attitude itself is propagated on the rotation group, while only a local, minimal representation of the attitude error is estimated by the Kalman filter. This eliminates the need for defining addition on the group of unit quaternions and leads to numerically robust behaviour. In this work, a MEKF is employed, using a quaternion to represent the nominal attitude and a three-dimensional small-angle error state to represent deviations from this nominal solution.

Quaternion kinematics

The attitude of the spacecraft is represented by a unit quaternion

$$\mathbf{q} = \begin{bmatrix} q_0 & \mathbf{q}_v^\top \end{bmatrix}^\top, \tag{4.23}$$

where q_0 is the scalar part and $\mathbf{q}_v \in \mathbb{R}^3$ is the vector part. The quaternion represents the rotation from an inertial reference frame to the spacecraft body frame.

The time evolution of the quaternion is governed by the measured angular rate of the spacecraft. Let the true body angular rate be denoted by $\boldsymbol{\omega}$. The quaternion kinematic equation is given by

$$\dot{\mathbf{q}} = \frac{1}{2}\,\boldsymbol{\Omega}(\boldsymbol{\omega})\,\mathbf{q}, \tag{4.24}$$

with

$$\boldsymbol{\Omega}(\boldsymbol{\omega}) = \begin{bmatrix} 0 & -\boldsymbol{\omega}^{\top} \\ \boldsymbol{\omega} & -[\boldsymbol{\omega}]_{\times} \end{bmatrix}, \tag{4.25}$$

where $[\boldsymbol{\omega}]_{\times}$ denotes the skew-symmetric matrix corresponding to the cross product.

In practice, the angular rate is obtained from a MEMS gyroscope, which provides a measurement

$$\boldsymbol{\omega}_m = \boldsymbol{\omega} + \mathbf{b}_{\omega} + \mathbf{n}_{\omega}, \tag{4.26}$$

where $\mathbf{b}_{\omega}$ is a slowly varying bias and $\mathbf{n}_{\omega}$ represents measurement noise. The measured angular rate is used to propagate the nominal quaternion forward in time by numerical integration, resulting in uninterrupted attitude propagation even in the absence of reference vector measurements.

Construction of the multiplicative extended Kalman filter

In the error-state formulation, the filter does not estimate the quaternion directly. Instead, the state is decomposed into a nominal part and a small error. The true attitude quaternion $\mathbf{q}$ is written as

$$\mathbf{q} = \delta\mathbf{q} \otimes \hat{\mathbf{q}}, \tag{4.27}$$

where $\hat{\mathbf{q}}$ is the nominal quaternion obtained from gyro integration, $\otimes$ denotes quaternion multiplication, and $\delta\mathbf{q}$ is a small error quaternion. For small attitude errors, $\delta\mathbf{q}$ can be approximated using a three-dimensional rotation vector $\delta\boldsymbol{\theta}$,

$$\delta\mathbf{q} \approx \begin{bmatrix} 1 \\ \frac{1}{2}\delta\boldsymbol{\theta} \end{bmatrix}. \tag{4.28}$$

This allows the attitude error to be represented in a minimal, Euclidean space suitable for Kalman filtering. The error-state vector is defined as

$$\delta\mathbf{x} = \begin{bmatrix} \delta\boldsymbol{\theta} \\ \delta\mathbf{b}_{\omega} \end{bmatrix}, \tag{4.29}$$

where $\delta\boldsymbol{\theta}$ is the attitude error and $\delta\mathbf{b}_{\omega}$ is the gyroscope bias error. Linearized error dynamics are derived around the nominal state. The continuous-time error dynamics are given by

$$\dot{\delta\boldsymbol{\theta}} = -[\boldsymbol{\omega}_m - \hat{\mathbf{b}}_{\omega}]_{\times}\,\delta\boldsymbol{\theta} - \delta\mathbf{b}_{\omega} - \mathbf{n}_{\omega}, \tag{4.30}$$

$$\dot{\delta\mathbf{b}}_{\omega} = \mathbf{n}_b, \tag{4.31}$$

where $\mathbf{n}_b$ is a bias driving noise, typically modelled as white noise. These equations are discretized to propagate the error covariance forward in time using standard Kalman filter prediction equations, while the nominal quaternion itself is propagated independently.

Measurement updates and bias estimation

The extended error-state Kalman filter separates attitude propagation from state correction by maintaining a nominal state and a small linear error state. The nominal state consists of the attitude quaternion $\hat{\mathbf{q}}$ and the gyroscope bias estimate $\hat{\mathbf{b}}_{\omega}$, while the Kalman filter operates on the error state

$$\delta\mathbf{x} = \begin{bmatrix} \delta\boldsymbol{\theta} \\ \delta\mathbf{b}_{\omega} \end{bmatrix}. \tag{4.32}$$

Propagation

The nominal quaternion is propagated using the bias-corrected angular rate

$$\boldsymbol{\omega}_c = \boldsymbol{\omega}_m - \hat{\mathbf{b}}_{\omega}, \tag{4.33}$$

according to the quaternion kinematic equation. This propagation is performed independently of the Kalman filter. The linearized continuous-time error dynamics can be written in standard state-space form as

$$\dot{\delta\mathbf{x}} = \mathbf{F}\,\delta\mathbf{x} + \mathbf{G}\,\mathbf{n}, \tag{4.34}$$

where

$$\mathbf{F} = \begin{bmatrix} -[\boldsymbol{\omega}_c]_\times & -\mathbf{I} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \qquad \mathbf{G} = \begin{bmatrix} -\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \tag{4.35}$$

and $\mathbf{n} = [\mathbf{n}_\omega^\top \; \mathbf{n}_b^\top]^\top$ represents the continuous-time gyroscope and bias driving noise. After discretization over the sampling interval Δt, the discrete-time state transition matrix is computed as

$$\boldsymbol{\Phi}_k \approx \exp(\mathbf{F}\Delta t) \approx \mathbf{I} + \mathbf{F}\Delta t. \tag{4.36}$$

In discrete time, the process noise is incorporated directly into the covariance propagation via the discrete process noise matrix $\mathbf{Q}_k$, giving

$$\mathbf{P}_{k+1}^- = \boldsymbol{\Phi}_k \mathbf{P}_k \boldsymbol{\Phi}_k^\top + \mathbf{Q}_k. \tag{4.37}$$

Effectively, $\mathbf{Q}_k$ plays the same role as $\mathbf{G}\mathbf{Q}\mathbf{G}^\top$ would in continuous time. By specifying the discrete noise covariance directly, the need to explicitly include $\mathbf{G}$ in the discrete-time covariance update is eliminated, while still correctly propagating the uncertainty in both the attitude error and the gyroscope bias.

Correction

When a reference vector measurement becomes available, a correction step is performed. Each measurement corresponds to a known inertial unit vector $\mathbf{r}^i$, obtained from ephemeris data combined with the GNSS-derived spacecraft position and time. Based on the nominal attitude, the predicted body-frame vector is

$$\hat{\mathbf{r}}^b = \mathbf{R}(\hat{\mathbf{q}})\,\mathbf{r}^i. \tag{4.38}$$

The measurement residual is defined as

$$\mathbf{y} = \mathbf{r}_m^b - \hat{\mathbf{r}}^b, \tag{4.39}$$

where $\mathbf{r}_m^b$ is the measured vector obtained from image processing.

Linearization of the measurement equation with respect to the attitude error yields

$$\mathbf{y} \approx -[\hat{\mathbf{r}}^b]_\times\,\delta\boldsymbol{\theta} + \mathbf{v}. \tag{4.40}$$

The measurement Jacobian with respect to the error state is therefore

$$\mathbf{H} = \begin{bmatrix} -[\hat{\mathbf{r}}^b]_\times & \mathbf{0} \end{bmatrix}. \tag{4.41}$$

Using this linearized model, the Kalman gain is computed as

$$\mathbf{K} = \mathbf{P}^-\mathbf{H}^\top \left(\mathbf{H}\mathbf{P}^-\mathbf{H}^\top + \mathbf{R}\right)^{-1}, \tag{4.42}$$

where $\mathbf{R}$ is the measurement noise covariance associated with the corresponding reference vector.

The error-state estimate and covariance are updated according to

$$\delta\hat{\mathbf{x}} = \mathbf{K}\,\mathbf{y}, \tag{4.43}$$

$$\mathbf{P}^+ = (\mathbf{I} - \mathbf{K}\mathbf{H})\,\mathbf{P}^-. \tag{4.44}$$

However, for numerical stability the Joseph's form of the update of the covariance is frequently recommended in literature [95]:

$$\mathbf{P}^{+} = (\mathbf{I} - \mathbf{K}\mathbf{H})\,\mathbf{P}^{-}\,(\mathbf{I} - \mathbf{K}\mathbf{H})^{\top} + \mathbf{K}\,\mathbf{R}\,\mathbf{K}^{\top} \tag{4.45}$$

The estimated attitude error $\delta\hat{\boldsymbol{\theta}}$ is then injected into the nominal quaternion via

$$\hat{\mathbf{q}} \leftarrow \delta\mathbf{q}(\delta\hat{\boldsymbol{\theta}}) \otimes \hat{\mathbf{q}}, \tag{4.46}$$

after which the attitude error state is reset to zero. Simultaneously, the gyroscope bias estimate is corrected using the corresponding components of $\delta\hat{\mathbf{x}}$. Through this mechanism, gyroscope drift and bias accumulation are corrected whenever reliable reference vectors are available, while uninterrupted attitude propagation is maintained during periods without valid image-based measurements.

Initialization of the Kalman filter

During initialisation after power-on, the spacecraft attitude is entirely unknown and therefore lies far outside the region where the EKF linearisation is valid. A dedicated initialisation procedure is therefore required to place the filter within the correct basin of attraction before normal operation can begin.

When only a single reference vector (e.g. Sun or Earth nadir) is available, the attitude is fundamentally ambiguous. Such a measurement constrains only two rotational degrees of freedom; rotations about the measured vector remain unobservable. On this axis, an error of up to 180 degrees is thus possible. More importantly, a single vector measurement allows the Kalman filter to initially converge on two attitude solutions that differ by an approximate 180° rotation. An EKF initialised using only one vector therefore cannot, in general, be guaranteed to converge to the correct attitude, regardless of the size of the initial covariance.

In theory, full attitude observability can be recovered over time if the inertial reference vector changes direction and the spacecraft undergoes sufficient rotational motion. Under these conditions, the combination of vector measurements and gyro propagation should, in principle, resolve the inherent directional ambiguity. In practice, however, it was observed that the EKF's local linearisation does not reliably address this global ambiguity. When initialised in the hemisphere opposite to the true attitude, the filter may interpret subsequent vector measurements as locally consistent with an attitude rotated by approximately 180 degrees. As a result, the EKF can converge to an inverted solution rather than correcting toward the true orientation.

For this reason, a magnetometer measurement is used during initialisation to provide a second, independent reference vector. The magnetometer is not required to be accurately calibrated for this purpose. Its role is solely to disambiguate the attitude and ensure that the initial estimate lies within the correct 180° basin. Even a magnetometer measurement with significant scale or bias errors is sufficient, provided its direction is approximately correct. This was confirmed by offsetting the magnetometer measurement by an angle of 45 degrees in random directions. An alternative is to allow two Kalman filter instances to run, which initialise in directions 180 degrees offset from each other. The filter with the lower covariances is likely the correct one. However, to confirm this a second reference vector is required like the sun outside of the eclipse. At this point, the correct filter could be identified. As at this moment, the single filter would be shown to be inconsistent and reset, a reliable method for faster convergence is not expected with two Kalman filters.

The initialisation procedure is therefore implemented as follows:

1. The gyroscope bias is initialised to zero. The attitude and bias covariance matrices are set to large values to reflect the high initial uncertainty. It was found a standard deviation of 35 degrees and 1 deg/s for the attitude and bias worked well.
2. As soon as one primary reference vector (Sun or Earth) is available, an attitude estimate is computed using the two-vector attitude determination method TRIAD with an auxiliary magnetometer measurement as a low-accuracy initial estimate.
3. The resulting attitude estimate is used to initialise the quaternion state of the Kalman filter.
4. Following the initialisation, the magnetometer is no longer required for nominal attitude estimation, and the filter operates using only the optical reference vectors and gyroscope measurements. This is specifically preferred, as magnetometer based attitude estimation is very sensitive to the

magnetic properties of the surrounding structures and electronics, making the determination of the accuracy challenging and requiring significant effort in calibration. However, for initialisation, the required accuracy is very relaxed, allowing this approach. For the simulation, a simulated offset of 45 degrees was used.

This approach guarantees that the filter is initialised in the correct global attitude hemisphere, preventing convergence to the 180°-inverted solution. Once this ambiguity is resolved, the EKF can reliably exploit the temporal observability provided by the spacecraft motion and gyroscope measurements to refine the attitude estimate.

Normalised innovation squared gating

In practice, not all image-based vector measurements are equally reliable. Partial occlusion, low contrast, or adverse illumination conditions can lead to erroneous vector estimates. Ideally these are prevented during the image processing, however some errors might still occur. To prevent such measurements from degrading the filter performance, a statistical consistency check based on the Normalised Innovation Squared (NIS) is employed.

The innovation is defined as

$$\mathbf{y} = \mathbf{z} - \hat{\mathbf{z}}, \tag{4.47}$$

where $\mathbf{z}$ is the measurement and $\hat{\mathbf{z}}$ its current value estimated by the Kalman filter prior to the measurement update. The NIS is then computed as

$$\epsilon = \mathbf{y}^\top \mathbf{S}^{-1} \mathbf{y}, \tag{4.48}$$

with $\mathbf{S}$, the innovation covariance, which was previously used in the Kalman gain computation and is defined as

$$\mathbf{S} = \left(\mathbf{H}\mathbf{P}^{-}\mathbf{H}^\top + \mathbf{R}\right) \tag{4.49}$$

Under the assumption of Gaussian noise and a correct model, ϵ follows a chi-squared distribution with degrees of freedom equal to the measurement dimension. A measurement is accepted only if ϵ lies below a predefined threshold corresponding to a chosen confidence level. Measurements exceeding this threshold are deemed to be statistically improbable and rejected and do not trigger a Kalman update.

This gating mechanism is particularly valuable for image-based sensors, as it allows the filter to automatically disregard unreliable vector observations while continuing to propagate the attitude using gyro data alone. As a result, robustness against transient sensor degradation is significantly improved without requiring explicit mode switching logic.

If a large number of measurements are flagged as outliers, this indicates that the filter is no longer consistent with its underlying assumptions. This situation may arise if the extracted reference vectors are corrupted by external disturbances, or if the filter has converged to an incorrect state estimate. In either case, the state estimate can no longer be trusted. The filter is therefore reset, allowing it to reconverge to a valid solution.

4.5.3. Propagation and measurement covariance determination

The Kalman filter's performance in estimating spacecraft attitude is dependent on the accurate specification of process and measurement noise covariances. For the integrated horizon and Sun sensor system, initial covariance values are derived from sensor datasheets and theoretical error models. The gyroscope's process noise is set according to its angle random walk and bias instability previously determined, while the measurement noise for the horizon and Sun sensors is initialized based on the approximate angular uncertainties of the measured vectors.

However, theoretical values alone do not yield optimal filter performance. The covariances are refined through an iterative tuning process: the filter is first tested in a simulation environment, where the true attitude is known, allowing for direct comparison between estimated and actual states. The residuals of the filter's estimates are analyzed to determine if the filter is overly optimistic or conservative.

If the residuals are larger than expected, the process or measurement noise is increased to better reflect the true uncertainty. If the filter is slow to correct errors, the measurement noise is reduced to give more weight to the sensor updates. Each covariance is adjusted manually, and the filter's performance is re-evaluated until the true errors align with the estimated covariances.

4.6. Simulation setup

To evaluate the combination of the various parts of the sensor and develop the end-to-end sensor algorithms, a MATLAB SIMULINK model is set up, which takes into account the entire chain of measurements to a final attitude estimate. This allows an evaluation of the achievable accuracy, as well as a determination of which conditions and variables impact this accuracy.

4.6.1. Simulation framework and model architecture

To develop, validate, and quantitatively evaluate the proposed attitude determination algorithms, a comprehensive simulation environment has been implemented in SIMULINK. This framework enables controlled testing of the complete processing chain, from orbital and attitude dynamics through sensor generation and onboard estimation, to the final attitude and bias error evaluation. Particular emphasis is placed on the realism of the sensor models and on modularity, allowing individual components to be modified or replaced without affecting the overall architecture.

The simulation framework is structured into three main parts: sensor and environment generation, attitude determination, and result derivation and evaluation. An overview of the data flow between these subsystems is shown in Fig. 4.22. More detailed views of this simulation are shown in Appendix A.

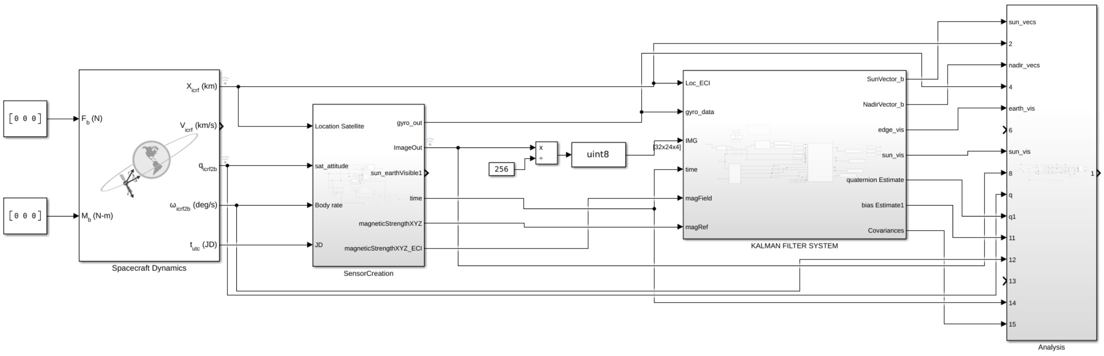

Figure 4.22: Overview of SIMULINK simulation

4.6.2. Sensor and environment generation

The first part of the simulation framework is responsible for generating the true spacecraft state and corresponding sensor observations. Orbital motion and attitude dynamics are modelled using standard SIMULINK blocks for rigid body dynamics, producing time histories of the spacecraft position, velocity, attitude quaternion, and angular rates in an Earth-centred reference frame.

Based on the simulated orbit and attitude, synthetic image data are generated by an external image rendering pipeline implemented in Python. At each simulation step, the SIMULINK model passes the current spacecraft pose, camera parameters, and time information to the image generator, which produces infrared images for each camera consistent with the geometric configuration of the Earth, Sun, and spacecraft. This coupling enables realistic modelling of visibility conditions, partial occlusion, and varying illumination while keeping the image generation independent of the estimation logic.

In parallel, inertial measurements are generated using a detailed MEMS gyroscope model. The angular rate measurements are constructed from the simulated true body rates and corrupted using stochastic error models derived from experimentally obtained Allan variance parameters. Bias instability, angle random walk, and rate noise are included using MATLAB Inertial Measurement Unit (IMU) modelling tools. Additionally, temperature-dependent effects are incorporated by modulating bias and noise parameters as a function of a simulated sensor temperature profile, enabling assessment of thermal sensitivity and bias drift over time. The output of this subsystem consists of synthetic camera images, gyroscope measurements, and ground-truth attitude and bias states used for later evaluation.

4.6.3. Attitude determination and sensor fusion

The second part of the simulation implements the onboard attitude determination pipeline. For each camera, the generated images are processed using the image-based algorithms described in the previous chapters to extract reference vectors, such as the Sun direction or Earth nadir vector, expressed in the camera or body frame. Intermediate quantities, including detected limb points and estimated reference directions, are retained for diagnostic purposes.

The extracted reference vectors are fused with gyroscope measurements using an extended error-state Kalman filter operating on a quaternion attitude representation. The filter propagates the nominal attitude using bias-corrected angular rate measurements and maintains estimates of both attitude error and gyroscope bias. Whenever valid image-based reference measurements are available, measurement updates are performed to correct accumulated drift and refine the bias estimate. Periods without valid image data are bridged through pure inertial propagation. The modular structure of this subsystem allows different image processing algorithms, sensor noise settings, or filter configurations to be evaluated within the same simulation framework.

4.6.4. Result derivation and performance evaluation

The final part of the simulation framework is dedicated to result derivation, visualisation, and quantitative performance assessment. Ground-truth states generated in the sensor simulation are compared against the estimated attitude and bias outputs of the Kalman filter to compute estimation errors over time.

Both intermediate and final results are recorded and visualised. Intermediate outputs include estimated reference vectors, innovation statistics, and filter covariance evolution, enabling detailed analysis of filter behaviour and observability under varying conditions. Final performance metrics include attitude error, bias estimation error, convergence time, and sensitivity to sensor noise and environmental effects.

This structured simulation framework enables systematic evaluation of the proposed attitude determination system under reproducible and realistic conditions, and provides a controlled environment for investigating the influence of sensor quality, image availability, and thermal effects on overall estimation performance.

4.7. Algorithm Verification on BEXUS balloon

The algorithm was developed targeting exclusively the images generated by the previously explained raytracing engine. Therefore, it is required to verify whether the algorithms perform reliably on real data acquired by the sensors.

To validate the accuracy of the image generation program as well as later verify the end-to-end sensor operation, a launch of the sensor aboard a stratospheric balloon was planned. This enables the acquisition of images using the used IR sensors from above the majority of the atmosphere, fairly closely replicating the conditions expected to be encountered in orbit. This verification flight was to be done within the REXUS/BEXUS programme, jointly organised by the German Aerospace Center (DLR) and Rymdstyrelsen, the Swedish Space Agency and executed by Swedish Space Corporation (SSC). This programme allows students to propose experiments to be performed either in microgravity onboard a sounding rocket or on a stratospheric balloon. For this thesis, an application was submitted and approved for a flight on Balloon Experiments for University Students (BEXUS). In addition to the attitude sensor treated in this thesis, a Sun sensor was to be tested with a high-performance 9-dof IMU acting as a reference attitude sensor.

The experiment was built within a team of students from both TU Delft and Ernst Abbe Hochschule Jena (EAH-Jena) in Germany, formed under the name Sensors for High Altitude Research Experiment – Optical Navigation (SHAREon) during a previous project, participating in this programme, and was done in parallel to this thesis. The intention was that this launch would provide data with which the design and assumptions used for the creation of this attitude sensor would be verified and validated. However, as such experiments are not guaranteed to work, especially if conducted by students, the design of the sensor and the resulting thesis were laid out such that the data would be a verification and validation case, but would not be required for the successful completion of the thesis. This section will first describe the justification for the similarity of the environments, followed by the actual implementation, followed by the implementation and its results.

4.7.1. Similarity of environments

Acquiring images of the horizon using an IR camera within the atmosphere will result in the background being at the approximate temperature of the atmosphere, as the atmosphere itself is only partially transparent to IR light. Instead, it also emits IR light with wavelengths correlated to the atmosphere's temperature. Thus, to accurately gain an understanding of the expected imagery to be acquired, the atmosphere needs to be removed when trying to take representative images. As a result, a high altitude flight can take such representative images.

Stratospheric balloons typically attain an altitude of around 25 to 30km, at which point more than 90% of the mass of the atmosphere is left below. This ensures the atmosphere has minimal influence on the images acquired by the sensors. Additionally, the horizon at this altitude approaches a distinct smooth line without any variations due to terrain or buildings while also revealing its curvature.

4.7.2. Experimental setup

The stratospheric balloon within the BEXUS programme provides a power and communication system to allow remote control as well as data transmission to the ground. Thus, the setup of the experiment requires a way to mount the sensor in a known orientation with an unobstructed view.

The stratospheric balloon platform consists of the balloon itself with the suspension lines, onto which the so-called gondola, which is constructed from aluminium profiles, is hung. On it, the various experiments, as well as power and communication systems, are mounted. The experiment was assigned a space in this construction, where all the desired parts could be mounted.

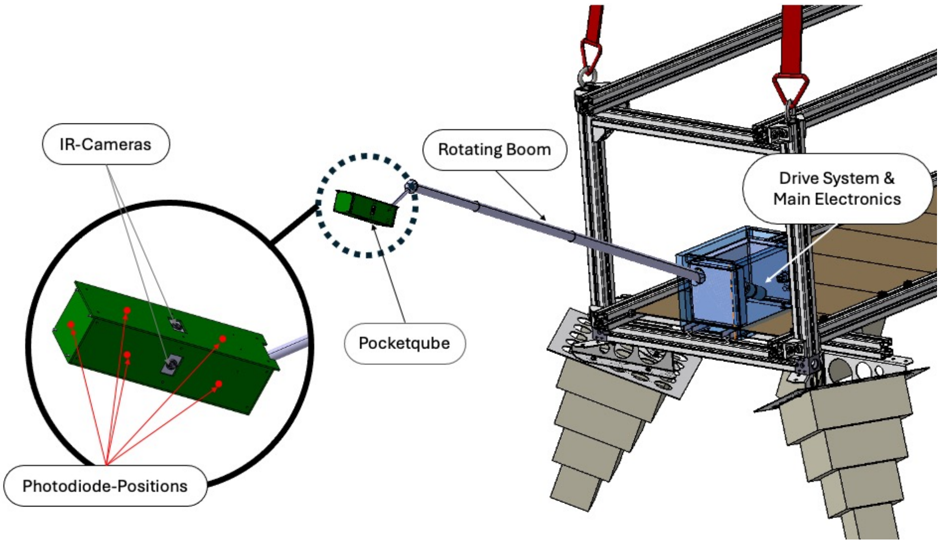

Figure 4.23: Experiment concept to acquire representative data for both the IR Sensors and gyro.

The sensor was integrated in a representative mockup of the final PocketQube satellite and mounted on a boom extending beyond the gondola. This ensures that all four sides of the PocketQube are unobstructed by both the rest of the gondola and the balloon floating above.

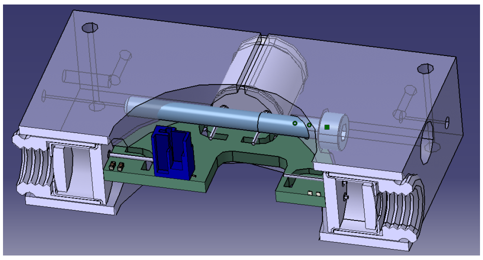

Figure 4.24: IR Sensor integration concept into structure for mounting in experiment PocketQube.

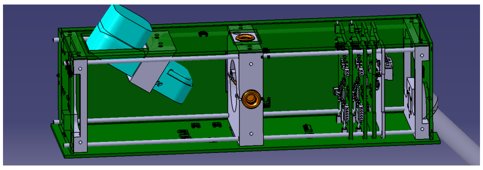

Figure 4.25: Layout of the internals of the PocketQube used for the high altitude flight. Note the angled attachment to a boom on the right end of the pseudo spacecraft, as well as the integrated camera to capture video during the flight.

The PocketQube itself, however, requires quite a few supporting systems to allow communication with the ground station during flight as well as the reference attitude determination.

Additionally, to increase the variation in the attitudes as well as angular rates the boom on which the PocketQube is mounted was supported between two bearings and motorized to allow for active control together with an encoder to accurately acquire the rotational state of the PocketQube. The pass-through of communications and power from the PocketQube to the rest of the systems thus had to be performed using a slip ring. The sensor system communicates to a Raspberry Pi mounted to the base of the experiment, which in turn communicates to a laptop on the ground via the system provided by the organiser SSC.

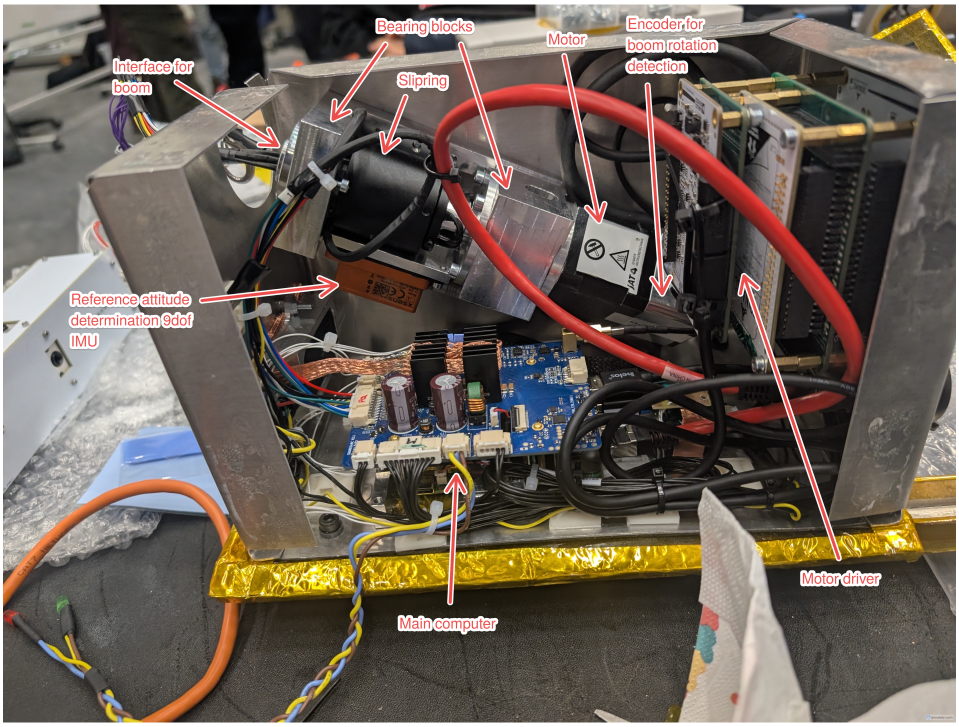

Figure 4.26: Integrated experiment main box, containing all systems to provide the required interfaces for the experiment to acquire the required data from the PocketQube mounted sensors.

A high-quality IMU (Xsens MTi-100) was mounted on the support of the bearings, which recorded angular velocities, accelerations, and the magnetic field, which, when fused, provides a high-accuracy reference attitude of the bearing mount. By combining this with the rotation of the boom from the encoder, a ground truth for the attitude of the sensor can be derived. In addition to the scientifically significant parts of the experiment, a number of sensors for health monitoring in the challenging environment of the stratosphere, as well as a camera within the PocketQube to capture pretty imagery of the experiment, were added. The entire system was integrated into the gondola, together with a number of other experiment teams, resulting in the final assembled gondola as shown in Figure 4.27.

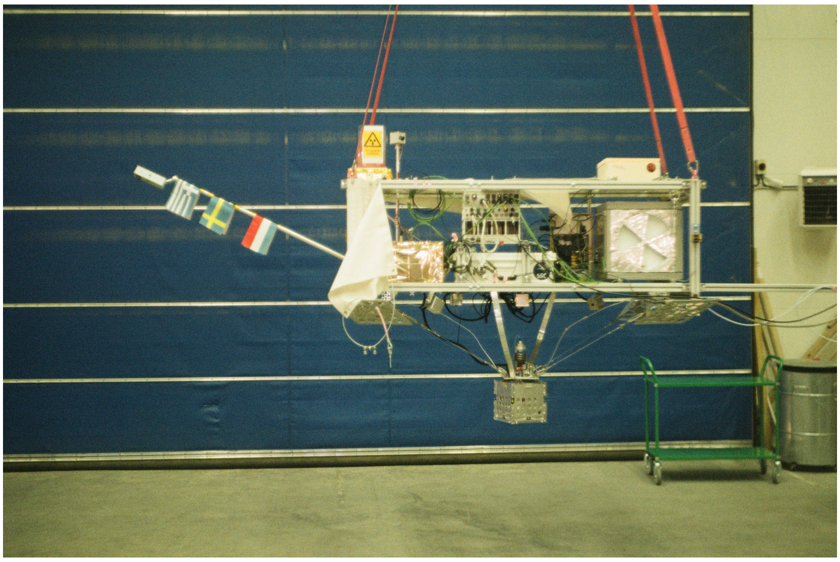

Figure 4.27: Fully integrated gondola with various other experiments and systems provided by SSC. Note the boom with the PocketQube mounted at the end that collects the required data for the experiment, as well as flags to increase visibility to prevent collisions with people working around the gondola.

4.7.3. BEXUS launch outcomes

The launch campaign took place in October 2025 at SSC's Esrange near Kiruna, Sweden, during which two gondolas called BEXUS 36 and BEXUS 37 were to be launched, with the horizon sensor being assigned to BEXUS 36. Due to issues with some experiments, the balloons were launched out of order, with BEXUS 37 performing a 12-hour flight.

Following that, the launch of the second balloon with the sensor onboard was attempted. However, an issue was encountered, and the balloon did not provide enough lift, resulting in the gondola dropping to the ground from the launch vehicle instead of ascending into the stratosphere. A sequence of these images is provided in Figure 4.28.

Two camera angles were available, one from an external viewpoint recording the gondola hanging off the launch vehicle, with the other being from within the PocketQube mounted on the rotating boom. The boom was commanded to view upwards to capture the launch of the balloon.

These images have been chronologically aligned, with each image showing the following:

1. Rise of the balloon after filling and release. The balloon is still rising above the launch vehicle as everything looks nominal.
2. The balloon has reached the point above the launch vehicle and, due to inertia, is moving towards the left of the image frame.
3. The balloon has passed the launch vehicle, at which point the driver starts backing up to allow the balloon to separate from its friction fit, holding it onto the launch vehicle.
4. After the balloon has not separated, the launch vehicle stops reversing and drives forward, making the gondola swing back and forth.
5. In tune with the swinging of the gondola, the vehicle starts backing up again, finally overcoming the friction holding on the gondola. However, instead of ascending into the stratosphere, the balloon drops on the ground.
6. The balloon drifts along without enough lift to ascend, however it does possess enough force to tug on the gondola, tipping it over.

Figure 4.28: Sequence of failed launch of BEXUS 36.

A second launch later during the campaign was not possible, as the cause of the failure could not immediately be determined and an investigation had to be completed, which at the time of writing is still ongoing. However, preliminary results indicate the balloon had a tear, which resulted in the lifting gas escaping. This reduced lift prevented the balloon from lifting the gondola off the launch vehicle, after which the release was forced through imparting the swing onto the gondola, after which it finally separated and crash-landed.

As a consequence, the acquisition of representative data in the desired approximate environment failed and could not be completed due to the external circumstances of the failure of the launch vehicle. However, the organizers of the BEXUS programme have indicated that a future replacement launch could be accommodated. This opportunity, however, falls in the summer of 2026, far beyond the intended duration of this thesis project. Consequently, the complete validation of the sensor will not be possible within the time frame of this thesis.

However, as it is intended to use the sensor onboard a real mission of Delfi Twin, this opportunity will still be useful for the qualification for orbital flight in the future.

4.7.4. Alternative Algorithm validation

Instead of being able to validate the entire algorithm end-to-end aboard the BEXUS balloon, an alternative became necessary. Instead of performing sensor measurements in situ in a similar environment, a number of images were taken that are in representative conditions on the ground instead. These cannot perform the same detailed validation, however can perform the validation of subsections of the algorithm. For this purpose, a number of images in two scenarios are taken:

The first scenario is images of the Sun from the ground, or alternatively, of a soldering iron to simulate the point source that the Sun represents. One such image is shown in Figure 4.29.

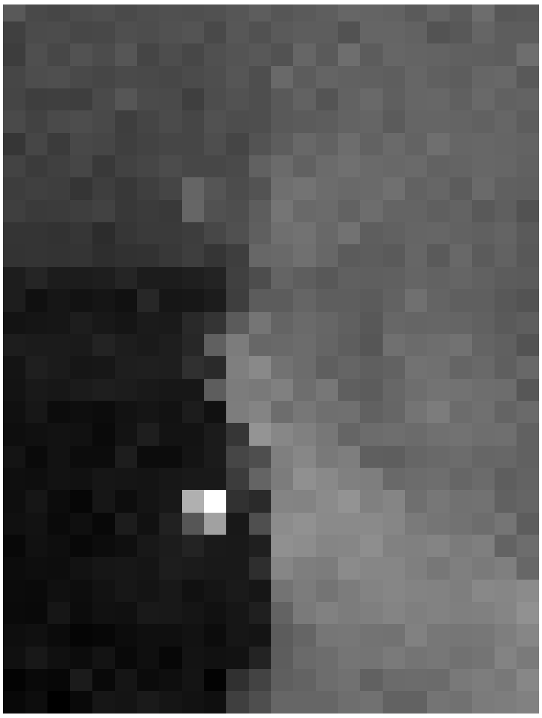

Figure 4.29: Example image of a soldering iron in IR to simulate Sun imaging

The other scenario is images of a specially constructed Earth simulator, which consists of an aluminium disk. The disk is painted black on one side for good IR emmissivity and has resistors mounted on the other side, as shown in Figure 4.30 and Figure 4.31. These heat up the disk so that it can be distinguished in the IR spectrum from the background at room temperature. This type of image is shown in Figure 4.32.

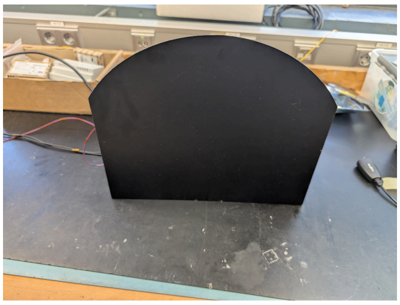

Figure 4.30: Front of Earth simulator showing black coating

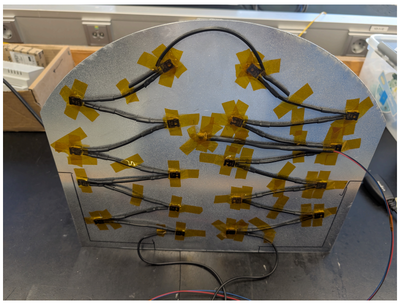

Figure 4.31: Backside of Earth simulator showing mounted resistors

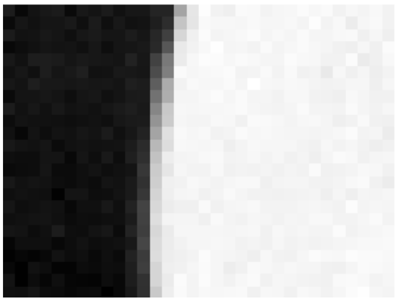

Figure 4.32: Example image of the Earth simulator in IR to simulate Earth horizon imaging

As accurate positioning of the sensors relative to the targets is challenging, and a setup to perform this task was not feasible to build within the time constraints after the balloon failure, the only way the algorithm could be judged is by confirming the presence classification of the two objects is correct and that the resulting vector if present is feasible and pointing in the approximately correct direction. Unfortunately, this cannot give certainty that the algorithms perform as desired and the accuracy is as simulated. That means this is an opportunity for further work to develop this type of sensor. The intended flight will be performed at a reflight event outside of this thesis in August of 2026.

4.8. Implementation of proto-flight sensor in hardware

During this thesis, a prototype sensor has been developed to allow for the BEXUS verification flight as well as allow software development to commence. This prototype is designed such that it theoretically should be suitable for flight, allowing all of the software development as well as fit checks to be performed.

4.8.1. Hardware design

The sensor must be integrated within the satellite structure such that the IR Cameras can see into four different unobstructed directions. Most of the surface area of the PocketQube satellite will be occupied by solar panels for power generation, thus only a small amount of area is available where IR cameras could be placed. In the center of the satellite, an aluminium structure is already required to mount laser retro reflectors, which are essential to allow accurate orbit determination from ground observations. These naturally also require an unobstructed field of view to the outside to be able to perform their function. Thus, it is desired to combine the sensor into this structure, which will allow the maximum number of solar panels to be used as well.

A structure similar to the one as flown on Delfi-PQ is used as a starting point for this design. Mechanically, this structure needs to mount a PCB large enough for all the power, processing, storage and communication components required in addition to the gyroscope.

The inner area of the laser reflector system available for the PCB is merely around 35x35 mm², while at the same time being at a point in the satellite where cabling between the top and bottom parts of the satellite has to pass through. As this cabling is a bus connector, carrying both power and a RS485 communications bus, that also needs to be connected to the attitude sensor, a similar approach to the Delfi-PQ stack connector is taken. Instead of adding both a connector and a cabling hole through the PCB, a 10-pin connector is added on both the top and bottom sides of the PCB. These connectors are directly connected to each other, acting as a pass-through. Additionally, the signals required for the horizon sensor to function are also forked out to the appropriate components on the PCB. The design is shown in Figure 4.33.

The PCB and its shape are entirely designed around the available space and the possible locations of the MLX90640 cameras. 4 of these find their place right next to the laser retroreflectors on each satellite face. This determines the location of the Cameras on the PCB and the rest of the components are placed around these.

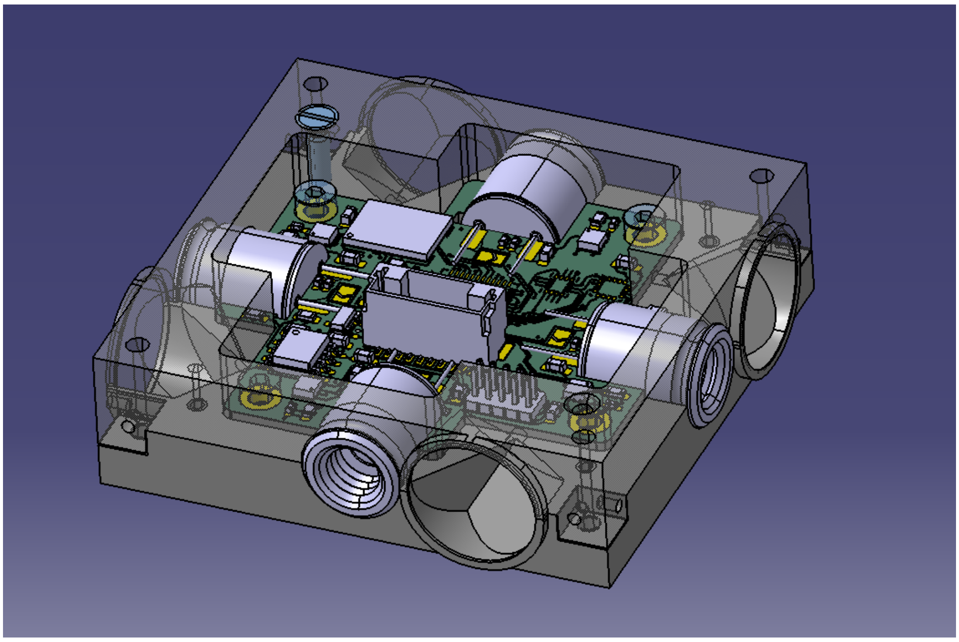

Figure 4.33: CAD of integrated PCB

At a high level these are:

- A STM32L452 microcontroller for the processing of the attitude sensor data and communication with the onboard computer. This is chosen as this series of microcontrollers is used across the satellite due to its high power efficiency and due to this specific model having a large amount of Random Access Memory (RAM) and Read-Only Memory (ROM) to allow the image processing algorithms to perform well.
- An InvenSense ICM-45686 IMU, which provides the interim data when no Earth horizon or Sun can be found in the imagery recorded. This is a high-performance, low-noise, and low-bias gyro, resulting in high performance even during times of no useful attitude data resulting from the imagery.
- A Micron 4Gb MT29F NAND flash chip. This allows the sensor to possibly store additional neural net parameters in case a neural net can be implemented with a size that exceeds the available ROM onboard the microcontroller. It also allows the recording of imagery that can be later retrieved for further analysis.
- A MMC5983MA magnetometer, which can be used as an auxiliary attitude sensor in case the algorithm does not perform as desired, using just observations from the imagery.
- A number of supporting components to regulate the power provided by the satellite bus and to allow communication via the RS485 interface.
- A load switch for each of the IR sensors to be able to switch on and off the power to these. The sensors are the components consuming the vast majority of the power of this attitude determination system, so as soon as an attitude is determined, it will be possible to shut off sensors that do not acquire any usable data.

4.8.2. Software implementation

The entire algorithm is to be implemented on the STM32 microcontroller mounted on the PCB designed previously. For this purpose, the algorithm developed in SIMULINK is used as a basis, which allows SIMULINK's Embedded Coder to generate the software running on the embedded system. This has the advantage that any possible errors in translating the algorithm can be avoided. The image processing algorithm and Kalman filter are generated as separate, standalone functions in C. However, this algorithm is not the only software that is required for the sensor to perform onboard the Delfi Twin mission. Thus, the generated functions can then be integrated into the manually programmed rest of the embedded software. A number of other tasks are required: the interfacing with the various sensors, such as the MLX90640 and gyro, the communication with the onboard computer of the satellite, as well as various housekeeping tasks that monitor the health and temperature of the sensor.

To allow for these various tasks, the general software architecture is built on FreeRTOS, which is an open-source real-time operating system that provides a structure that can allow many different tasks to run concurrently. FreeRTOS performs the scheduling of the various tasks according to pre-defined priorities, while also providing methods to coordinate the access to the hardware resources required by multiple tasks, as well as communication between tasks.

The software architecture of the prototype sensor can be split into the following tasks:

- **Communication and housekeeping task:** This task monitors the RS485 bus connected across the entire satellite, receiving required information for the sensor such as the GNSS location and time, transmitting the latest attitude estimates calculated by the other tasks.
- **Sensor readout tasks:** Reading the sensors at the appropriate times is essential. Especially for the gyro, it is crucial not to miss any measurements, as skipping measurements could lead to the attitude propagation becoming faulty. Therefore, a separate task is created for each and every IR Sensor as well as the gyro, where the gyro task is given the highest priority in the system. The read sensor data is then communicated to a separate task, which can asynchronously perform the more complex calculations that take up more time.
- **Image processing task:** This task is the most computationally heavy, as the image processing on a small embedded microcontroller requires a significant amount of both computational cycles and memory. However, as missing singular images acquired by the IR sensors is not immediately catastrophic for the attitude estimation, this task is given the lowest priority and thus runs whenever there is idle time from all of the other tasks which are at a higher priority. Additionally, the communication queues from the tasks providing the images are configured such that if the queue is full, the image is dropped and not further processed. This deals with the computational limits of the microcontroller while only slightly degrading the attitude determination performance, as the Kalman filter corrections are performed less frequently.
- **Kalman filter task:** This task receives the gyro readings as well as the processed reference vectors from the image processing task, and subsequently performs the actual attitude determination within the Kalman filter, performing both the propagation using the gyro measurements and the estimated bias, as well as the corrections which are received from the reference vectors.
- **Storage task:** A small flash memory chip is implemented onboard the PCB, which allows a large amount of the raw data to be stored during test measurement sessions or on the validation balloon flight. Thus, an additional task is created that receives the same data from the sensor readout tasks and writes it to this flash storage. This enables subsequent analysis of the data and evaluation of various settings and algorithms on the same data.

Not all of these tasks have been completed within the framework of this thesis, as the focus was on developing the software sufficiently to enable data collection and communication, while data processing and attitude estimation received lower priority. The communication and sensor readout, as well as the storage tasks, have been completed, which was required for all the raw data to be acquired and later analyzed during the BEXUS flight campaign. Further software development is required to complete the entire processing chain to determine the attitude.

5

Results

This chapter will present the results achieved using the previously described framework to develop the attitude determination system. First, the results achieved from the simulations performed in SIMULINK will be explained, showing the achievable accuracies of the reference vector determination during the simulation, followed by the achieved accuracy of the entire attitude determination system after filtering. Particular focus will be placed on the system initialisation and the performance of the integrated system in eclipse, where only a single reference vector is available. A montecarlo analysis is presented showing the robustness of the system, before exploring which component of the system is the limiting factor for better performance and should be the first target for improvement.

5.1. Accuracy of reference vector determination

The accuracy of the determination of the reference vectors is essential, as this is the only source for an absolute attitude knowledge in space. Two different reference vectors are used, the nadir vector estimated through the Earth's horizon line and the Sun vector, estimated by detecting the vector towards the bright point. For a representative simulation of the system in a 500km, 51.6 degree inclination orbit, the accuracies of these two vectors are shown in Figure 5.1 and Figure 5.2.

This simulation used arbitrary initial angular rates of $\omega_{b0} = \begin{bmatrix}1.2 & -2.5 & 3\end{bmatrix}$ [deg/s] as representative tumbling characteristic of unstabilised microsatellites.

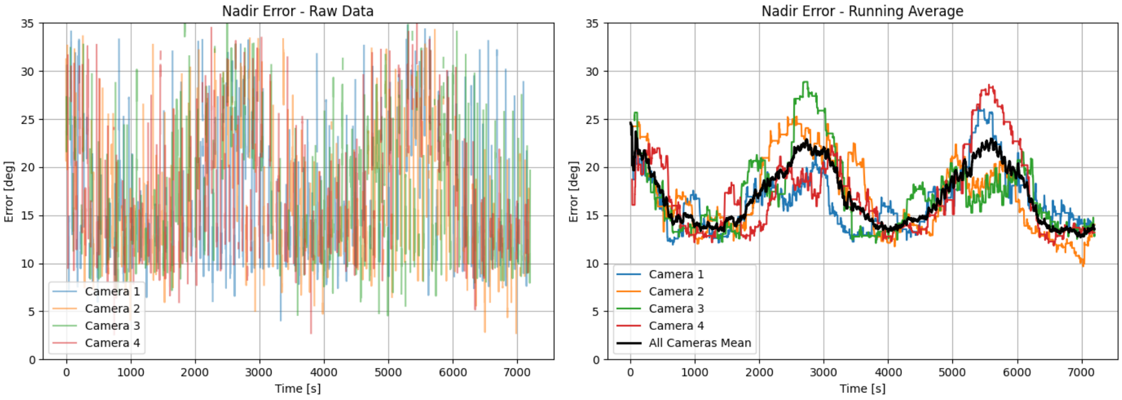

Figure 5.1: Accuracy of the raw reference nadir vector estimation over a 7200s simulation at an altitude of 500km

It can be observed that the accuracies of the Sun and nadir vector differ drastically. The Sun vector is significantly more accurate than the nadir vector, as the nadir vector is estimated from a small section of the horizon that is visible. The Sun vector error is on the order of single degrees, however, outliers can also be observed that have dramatically higher errors. These are observed during sunset or sunrise, as at this point, the Sun location algorithm interprets the combined Earth and Sun region as a singular region. The intensity weighted centroid of this region is thus far from the true Sun location. To eliminate these, more elaborate processing would be required, however these outliers are acceptable as they are caught by the innovation gating within the Kalman filter at a later stage in the process.

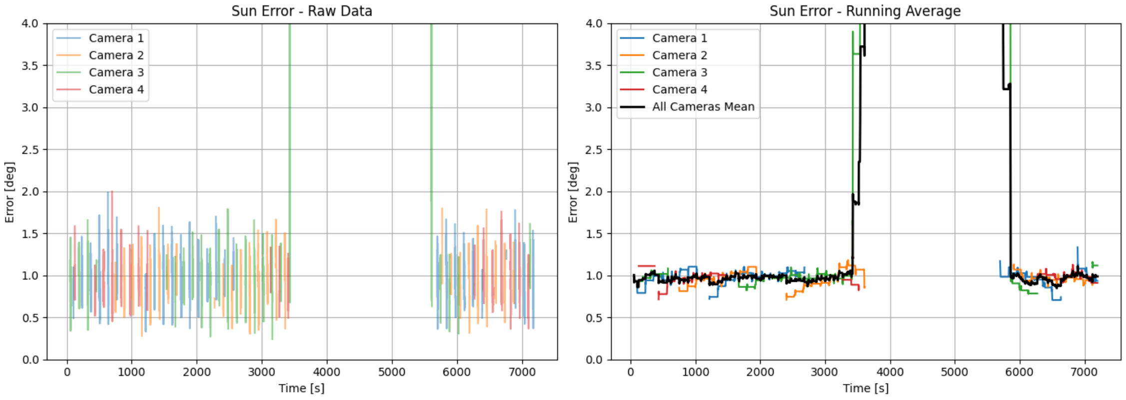

Figure 5.2: Accuracy of the raw Sun reference vector estimation over a 7200s simulation at an altitude of 500km. The region where the errors first increase and then disappear are the period of eclipse.

For the nadir vector determination, the key factor in the accuracy of the algorithm is the amount of the horizon visible in the acquired images. The algorithm already filters the imagery to always contain a minimum number of pixels of both the background and Earth, effectively enforcing a minimum length of horizon to be visible. However, as can be observed in Figure 5.1, the mean accuracy of the algorithm varies over time in a sinusoidal fashion. Initially, it could be believed this is correlated to the location in the orbit, however the period of this oscillation differs significantly from the duration of an orbit, and thus is more in line with the tumbling of the satellite.

5.2. Robustness to altitude variations

The altitude also dictates to a large extent how much of the horizon is visible. At a low altitude, only a small and straight section is visible, and the possible accuracy of the attitude determination is far worse compared to viewing a large section of the Earth's horizon from a large altitude.

For this reason, a study was performed where the same simulation was run with the only variation being the altitude of the orbit across a range of likely applications between 400 and 5000 km.

This is shown in Figure 5.3. The accuracy of the nadir vector depends on the altitude, however, the higher the altitude, the smaller the effect becomes as the fraction of the Earth in view does not change significantly.

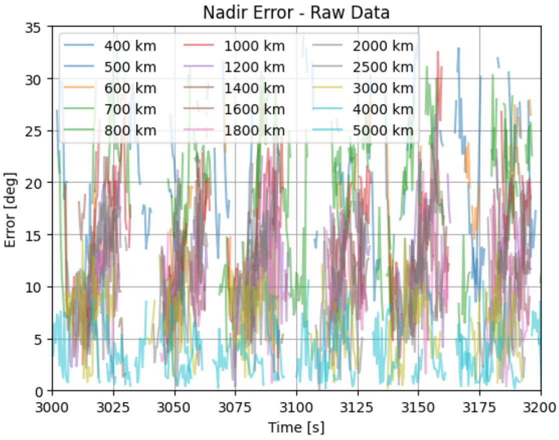

Figure 5.3: Accuracy of the raw nadir reference vector estimation compared to altitude. The NADIR vector improves at higher altitudes as more of the Earth's horizon becomes visible. The x axis is cropped to different timeframes to better visualise the periodic nature of the accuracy, with different amplitudes depending on the altitude.

As expected, the accuracy of the Sun reference vector determination is largely unaffected by the altitude, as shown in Figure 5.4. The only noticeable effect is the length of the eclipse period, which is reduced for the higher orbits. It can also be seen that around the entry and exit into the eclipse the accuracy of the Sun determination degrades, as a number of datapoints show far larger errors which are outliers produced by the algorithm.

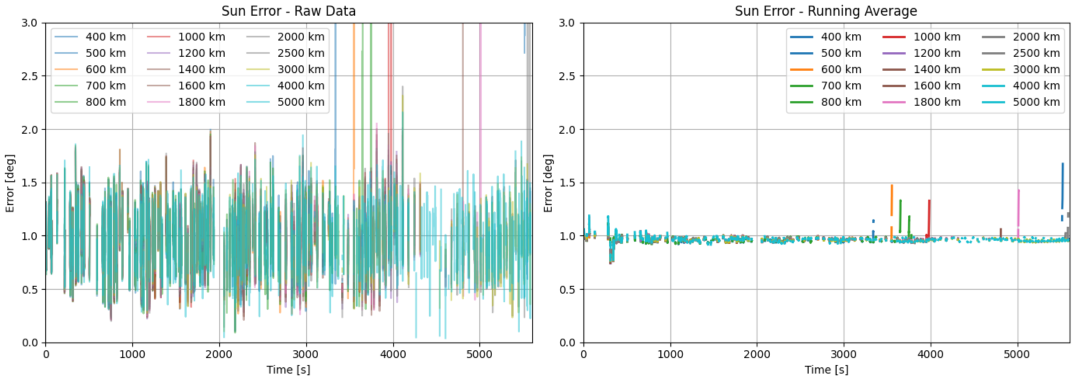

Figure 5.4: Accuracy of the raw Sun reference vector estimation compared to altitude. The Sun vector accuracy is independent of the altitude apart from the duration of eclipse.

5.3. Performance of Kalman filter

The previously treated performance of the reference vectors is crucial for the performance of the attitude determination system. However, with a well designed fusion system, the achieved attitude estimation can be improved which will be shown in this section. All of the following performances are shown for the same simulation case as used in section 5.1

5.3.1. Kalman filter intialisation

The Kalman filter is initialized and then run using the measurements acquired further on. As previously explained, the initialization merely aligns the reference vectors in the inertial and body frame, with an ambiguity around the axis of that vector. Hence, it can be observed during the initialisation that the intial errors of the Kalman filter are large, on the order of up to 90 degrees. It was found that this worked well most of the times, however in some circumstances the filter would converge to an attitude approximately 180 degrees away from the true attitude, essentially pointing in the exact opposite direction as the real attitude. With the Sun in view, the filter would reset once multiple observations of the Sun were made due to the repeated innovation gating triggering a reset. After a number of resets, where the covariance and attitude were reset, the filter would finally converge to the true attitude. As described in Equation 4.5.2, this was found to be an issue with the weak observability of the attitude combined with the linearisation of the error in the extended Kalman filter, leading to the filter converging to an inversed attitude estimate.

Once initialised, the errors drop quickly, as more measurements are acquired. Here two cases can be differentiated; initialisation in eclipse where only nadir measurements are available and the initialisation on the Sun side of the Earth where Sun vectors are immediately available as well.

Initialisation with observable Sun

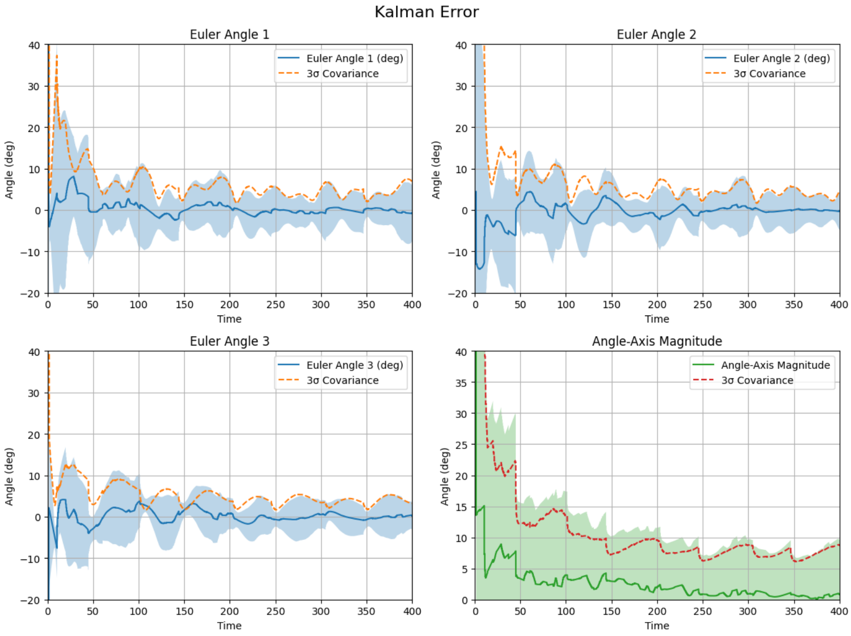

Figure 5.5: Initialisation of the Kalman filter with Sun visible. Quick convergence can be seen, with significant jumps at points where the Sun is visible.

The simplest case is when Sun reference measurements are available, since two independent reference vectors are provided, allowing full and immediate attitude determination. In this case, the Kalman filter converges quickly, as shown in Figure 5.6. Within tens of seconds, the estimation error decreases to the level of the accuracy of the nadir reference vector.

Within a minute or so, the attitude estimation error has dropped even further below approximately 5 degrees, as the state becomes more observable through the combination of the reference vector measurements together with the measured attitude dynamics by the gyro. Similarly, the gyro bias estimation, shown in Figure 5.6, converges as the attitude dynamics allow the disparity between the dynamics embedded in the filter together with the gyro measurements and the reference vectors to be observed. Thus the gyro bias estimation converges at a similar rate as the attitude estimate does to an accurate state.

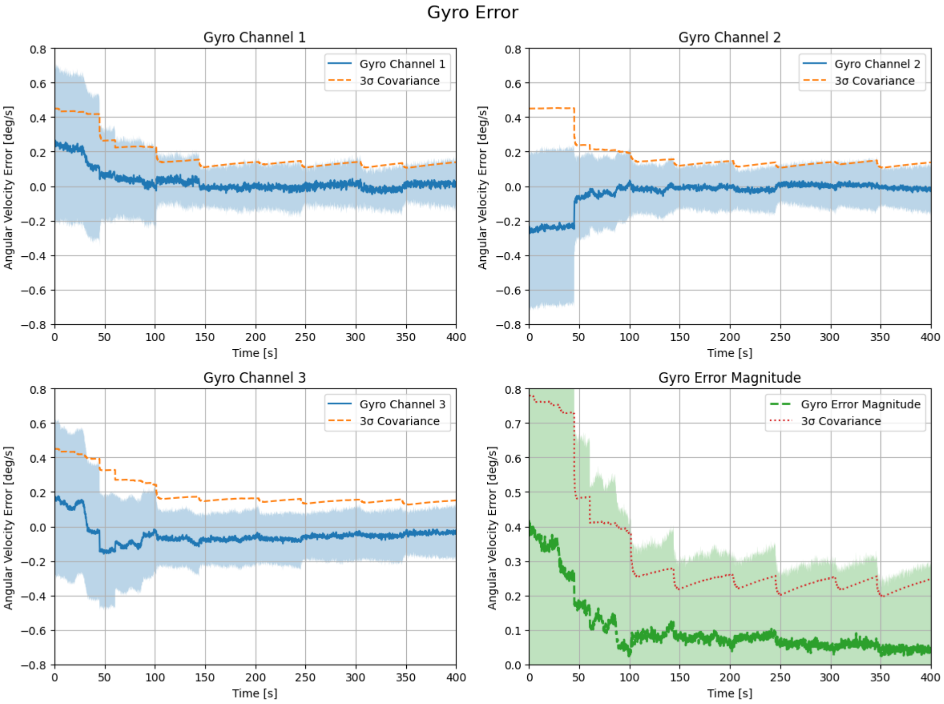

Figure 5.6: Initialisation of the Kalman filter's gyro bias with Sun visible. Quick convergence can be seen, with significant oscillatory jumps at points where the Sun is visible.

Initialisation in eclipse

In the case of the intialisation happening in eclipse as shown in Figure 5.7, the convergence does not progress past a certain point and the attitude errors converge only down to approximately between 10 to 20 degrees. The covariance output by the Kalman filter remains far larger, as the observability of the attitude is more limited, making the estimation worse. Especially during this phase, the initialisation with a magnetometer allows a full attitude estimate to be reached, even with the challenging observability of only one reference vector from the imaging data. It can be observed that the accuracy of the attitude first degrades after initialisation before converging again. This is due to the gyro error disturbing the attitude estimate while it converges. The attitude returns to the best estimate as soon as the gyro bias has reached a sufficient estimation accuracy.

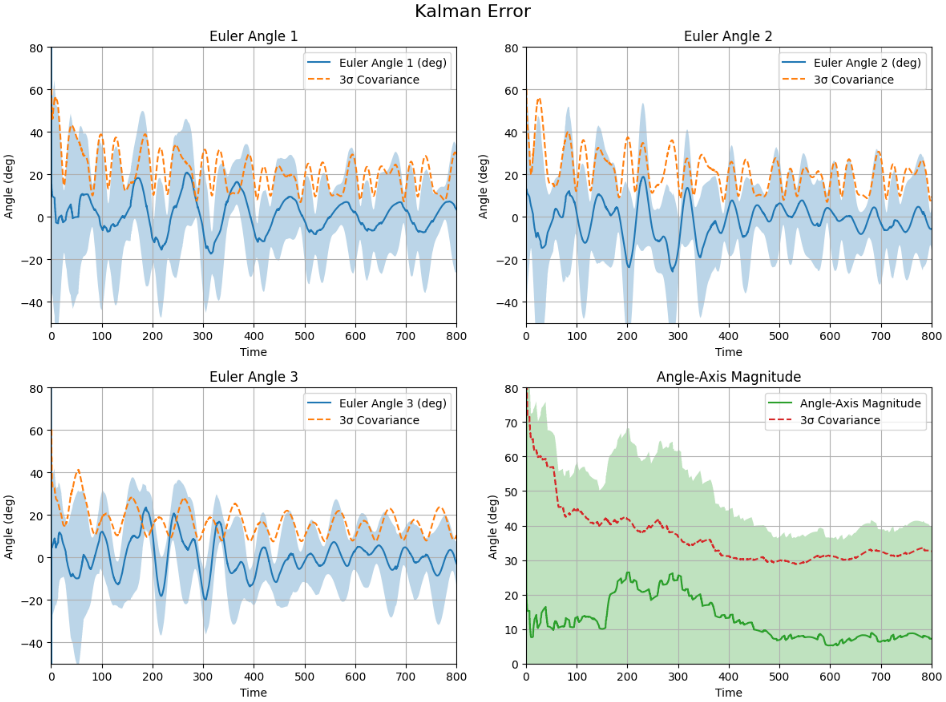

Figure 5.7: Initialisation of the Kalman filter's attitude estimate in eclipse. Errors do not decrease significantly from the initial estimate, while the covariances slowly converge showing increasing confidence in the estimate.

The convergence of the gyro bias, shown in Figure 5.8, is slower than in the case where the Sun is visible, and the error can only be reduced to to within around 0.1 deg/s. Again, the covariance estimated by the Kalman filter on these values is far larger due to the lower observability.

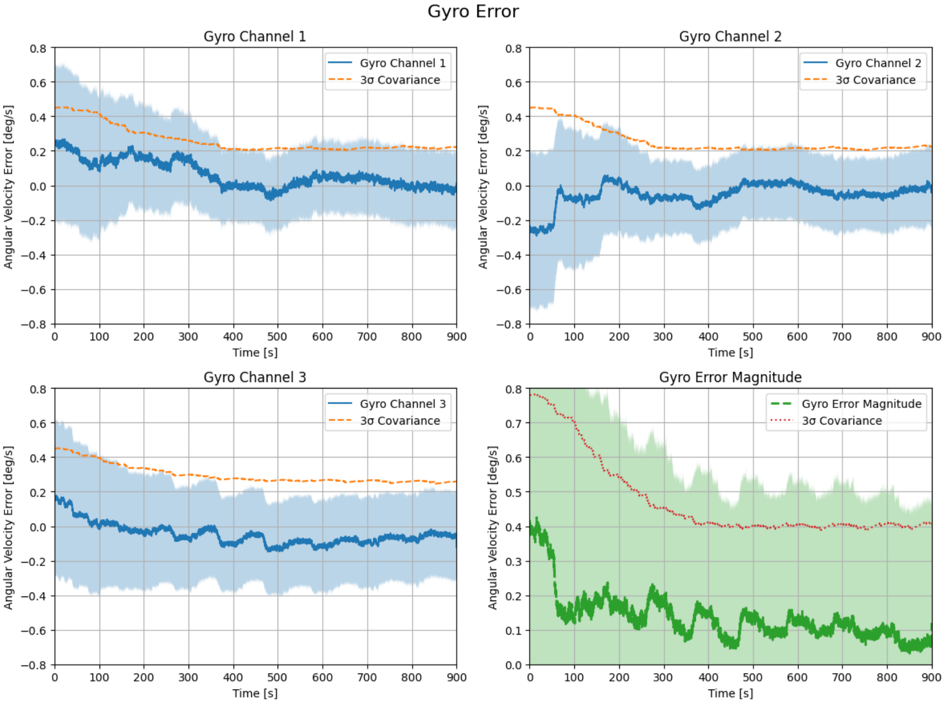

Figure 5.8: Slower convergence of the Kalman filter's estimated gyro bias, as the bias can not be observed as accurately through the noisy data provided by the nadir vector.

5.3.2. Steady state accuracy

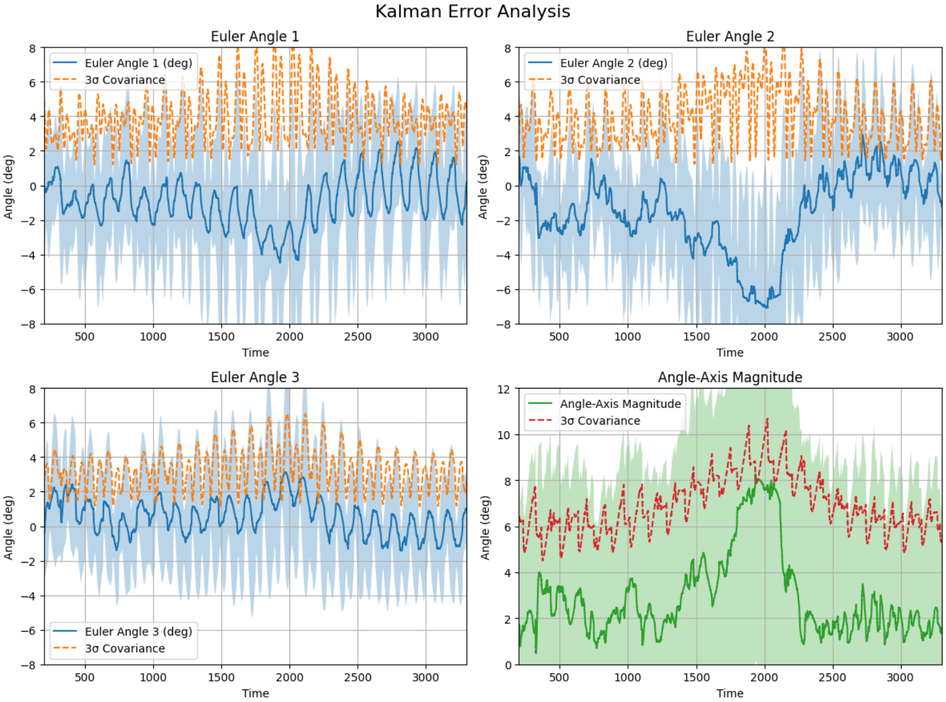

Figure 5.9: Steady state performance of the Kalman filter. The performance varies over time between approximately 4 degrees pointing knowledge up to approximately 8 degrees depending on the specific geometry of the estimation problem.

In steady state with the Sun visible as shown in Figure 5.9, it can be seen that the majority of the time the accuracy is between 2 and 4 degrees, with excursions up to 8 degrees. The variation in accuracy is presumed to be due to the varying geometry of the attitude estimation problem; the observability is degraded when the two reference vectors are pointing in directions that are approximately colinear. Additionally, as previously seen, the accuracy of the nadir reference vector varies with the attitude of the spacecraft.

Similar behaviour can be observed in the gyro error observation shown in Figure 5.10. The estimation accuracy varies between 0.05 and 0.075 deg/s. Additionally, in both the gyro bias and attitude estimation, the behaviour of the covariances is strongly oscillatory, which is explained by the intermittent observations of the Sun, which is not visible at all times but is only occasionally visible on the acquired imagery.

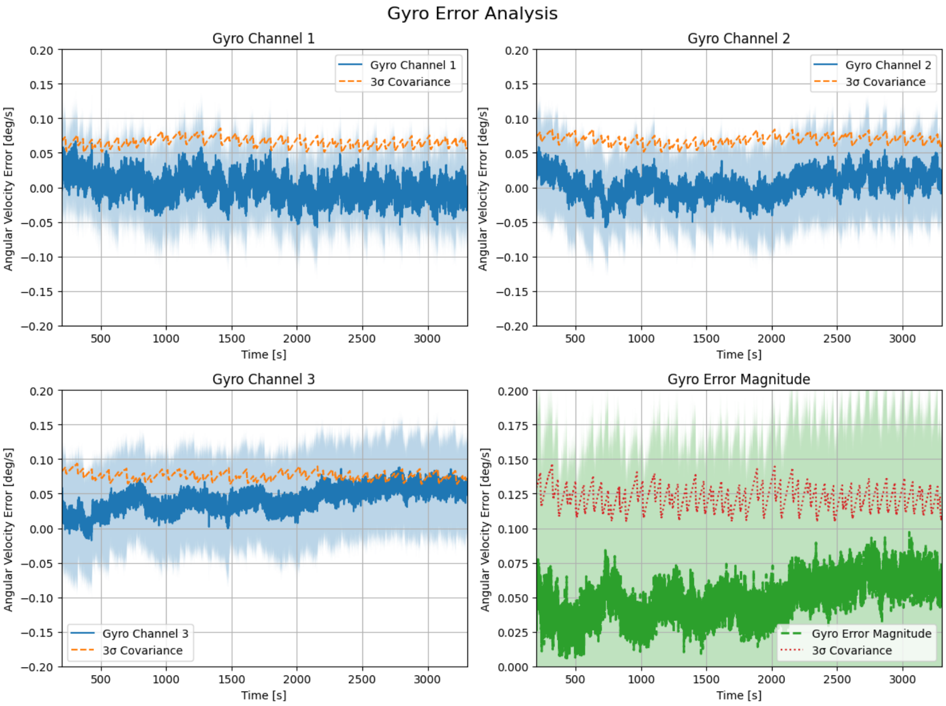

Figure 5.10: In steady state the bias gyro is estimated accurately to within 0.5 deg/s. It can be observed that the majority of this error stems from a single axis.

5.3.3. Performance in eclipse

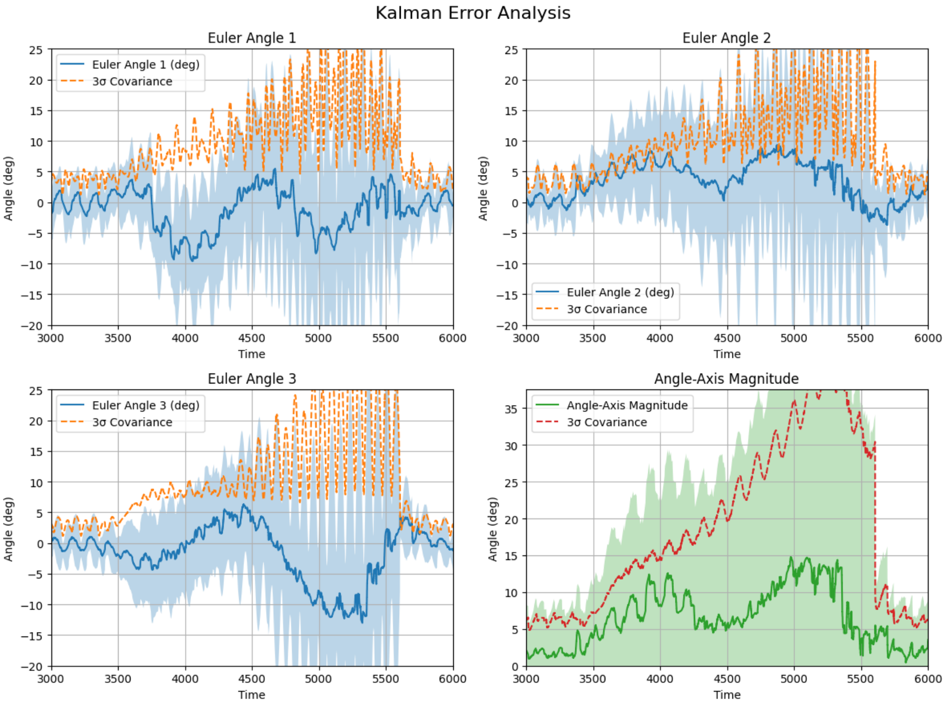

Figure 5.11: Attitude estimation performance as the spacecraft enters, transits, and eventually exits eclipse.

As expected, with only a single reference vector available, the performance of the estimation slowly degrades throughout the time in eclipse, as shown in Figure 5.11 where eclipse begins at approximately 3500s into this simulation. At this point, with the Sun available, the filter has converged to an accuracy of better than 5 degrees. Following the entry into eclipse, the error increases slowly, while the covariance increases significantly as well. The 3σ covariance increases up to approximately 40 degrees at the end of the eclipse, which remains below the allowed maximum error from the pointing knowledge requirement.

Interestingly, the estimates of the gyro biases shown in Figure 5.12 are somewhat more stable. The covariance does increase as well quite significantly, however the actual estimation error remains bounded to within 0.1 deg/s. Once the satellite exits eclipse it can be seen that the covariances and also the attitude errors rapidly converge back to the accuracy found in the original case with Sun reference vector measurements available.

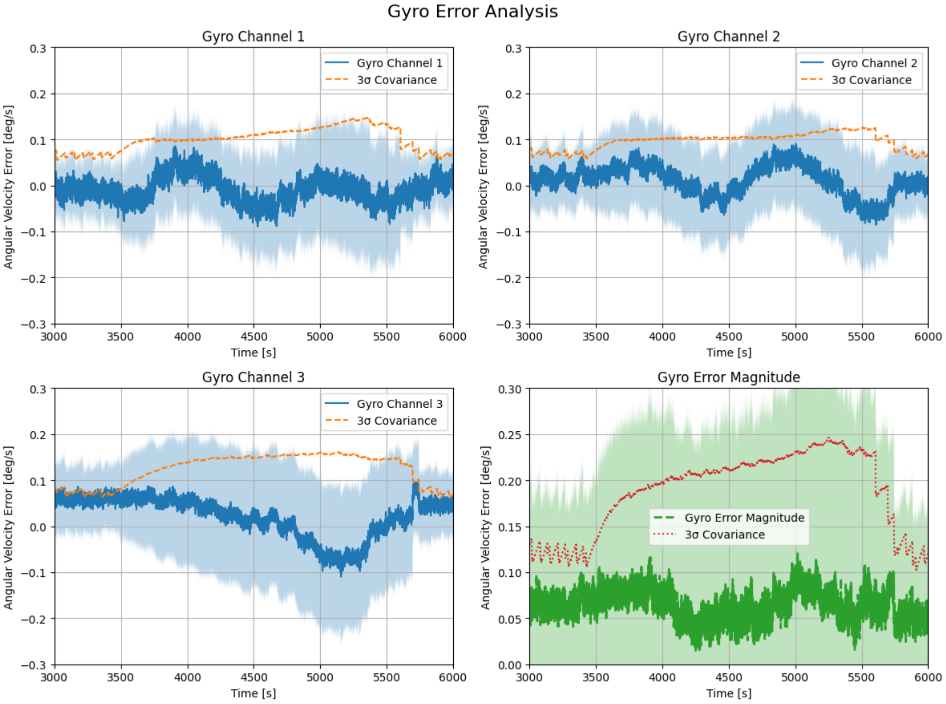

Figure 5.12: Gyro bias estimation performance as the spacecraft enters, transits, and eventually exits eclipse.

5.4. Monte Carlo Analysis

To evaluate the robustness and statistical performance of the proposed attitude determination approach, a Monte Carlo analysis was performed. Rather than relying on a single deterministic simulation, this analysis characterizes the sensitivity of the estimator to uncertainty in initial conditions and key system parameters by repeatedly executing the full simulation and estimation pipeline with randomized inputs.

For each Monte Carlo run, a set of parameters was sampled independently within predefined bounds around a nominal value. These bounds were chosen to represent realistic variations arising from initialization uncertainty, orbital state variability, sensor imperfections, and modelling errors in the image generation process. The complete list of varied parameters, together with their nominal values and sampling ranges, is provided in Table 5.1. All parameters were assumed to follow a uniform probability distribution.

The following parameters were varied in the Monte Carlo analysis:

- **Initial attitude**, capturing uncertainty in the spacecraft attitude at the start of the estimation process.
- **Body angular rates in 3 axes**, representing variations in the rotational motion during the observation interval.
- **True gyroscope bias**, modelling varying bias offsets to be estimated by the Kalman filter.

- **True anomaly**, defining the spacecraft position along its orbit and influencing illumination conditions and observation geometry.
- **Vertical/lateral focal length of the true image generator**, introducing uncertainty in the determined camera projection parameters.

Table 5.1: Monte Carlo parameters with nominal values and sampling ranges

Parameter	Nominal Value	Minimum	Maximum	Unit
Initial attitude (Euler angle)	[0,0,0]	[-90, -90, 90]	[90, 90, 90]	degrees
Body angular rates for each axis	0	-10	10	deg/s
True gyroscope bias for each axis	0	-1	1	deg/s
True anomaly	180	0	360	degrees
Vertical focal length (image generator)	40.5	41.65	42.5	pixels
Lateral focal length (image generator)	40.5	41.65	42.5	pixels

For each simulation run, the sampled parameters were used to generate synthetic sensor measurements and images. The attitude estimation algorithm was subsequently executed using a fixed estimator configuration and nominal model parameters. The resulting distributions of estimation errors and convergence behavior were then analyzed to assess estimator robustness and expected performance across a realistic operational envelope. The runs performed using these input parameters are shown in Figure 5.13.

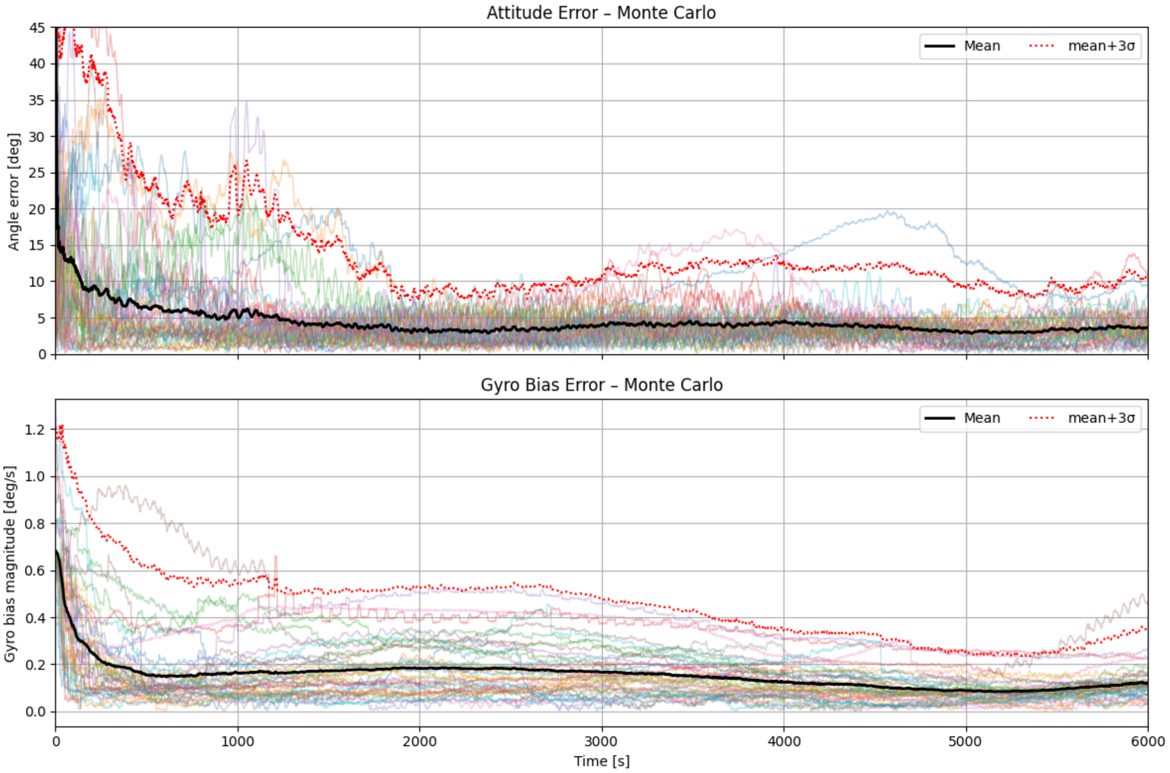

Figure 5.13: Performance of attitude determination sensor with varied inputs over 50 different runs.

It can be seen that even over this large variation of input parameters there is good performance of the attitude determination system, with the mean accuracy over all runs approaching and then remaining around 5 degrees. The 3σ accuracy also stays within 15 degrees once the attitude estimate has converged. The convergence of the attitude estimates has an interesting behaviour, as the convergence to a somewhat constant accuracy and variance takes quite a long time. This is explained by the fact that approximately half of the runs are initialised in eclipse, resulting in much slower convergence as previously seen in the observation of single runs previously, which eventually improves as the first observations of the Sun reference vector are made. At this point the standard deviation of the accuracy drops significantly with the 3σ accuracy approaching 10 degrees.

The bias estimation convergence behaves similarly to the attitude error, where convergence is partially slow for those runs initialised in eclipse. It can be however observed that the variance in the performance of the bias estimate is larger than the attitude estimate, and the mean performance of the bias estimate also converges to just below 0.2 deg/s.

5.5. Limiting Factors for Attitude Estimation Accuracy

In order to identify the dominant performance limitation of the proposed attitude determination chain, two additional simulation campaigns were conducted. The baseline configuration corresponds to the sensor and gyro models used throughout the previous sections. Starting from this baseline, two modified configurations were evaluated:

1. **Improved gyro model:** Angular random walk, bias instability, and temperature stability parameters were reduced by a factor of 10, while keeping all other components identical.
2. **Improved IR sensor model:** The baseline IR camera was replaced by the previously evaluated higher-performance alternative Seek Thermal sensor, with a resolution of 200x150 pixels and a FoV of 57x42 degrees. This was previously not selected in the nominal design due to size, mass, and cost constraints.

This approach isolates the contribution of inertial versus optical sensing to the final attitude accuracy.

5.5.1. Impact of Gyro Performance

Figure 5.14 shows the attitude error distribution obtained with the improved gyro model.

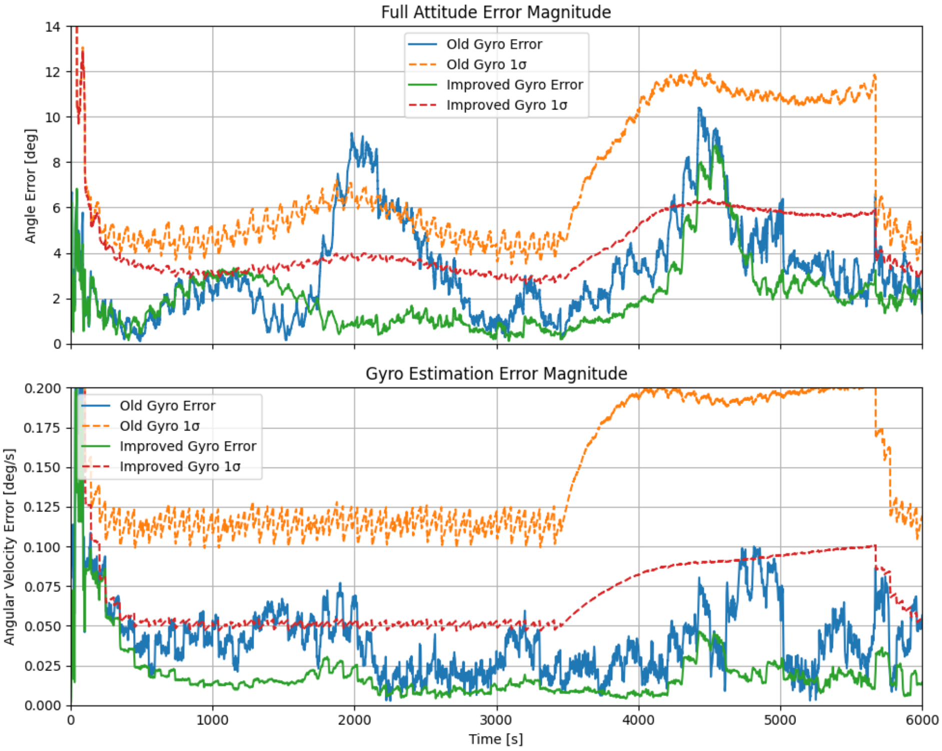

Figure 5.14: Attitude estimation and gyro bias error using a gyro model with noise and bias parameters reduced by a factor of 10.

Reducing gyro noise and bias terms leads to a measurable but limited improvement in overall performance, depending on the situation. In some phases, the attitude error improves significantly, however after sufficient time in eclipse the performance degrades in a similar manner as the baseline gyro.

The relative improvement is therefore not very significant, as the absolute attitude error still always depends on the absolute accuracy of the reference vectors.

5.5.2. Impact of IR Sensor Performance

The corresponding results for the improved IR sensor configuration are shown in Figure 5.15.

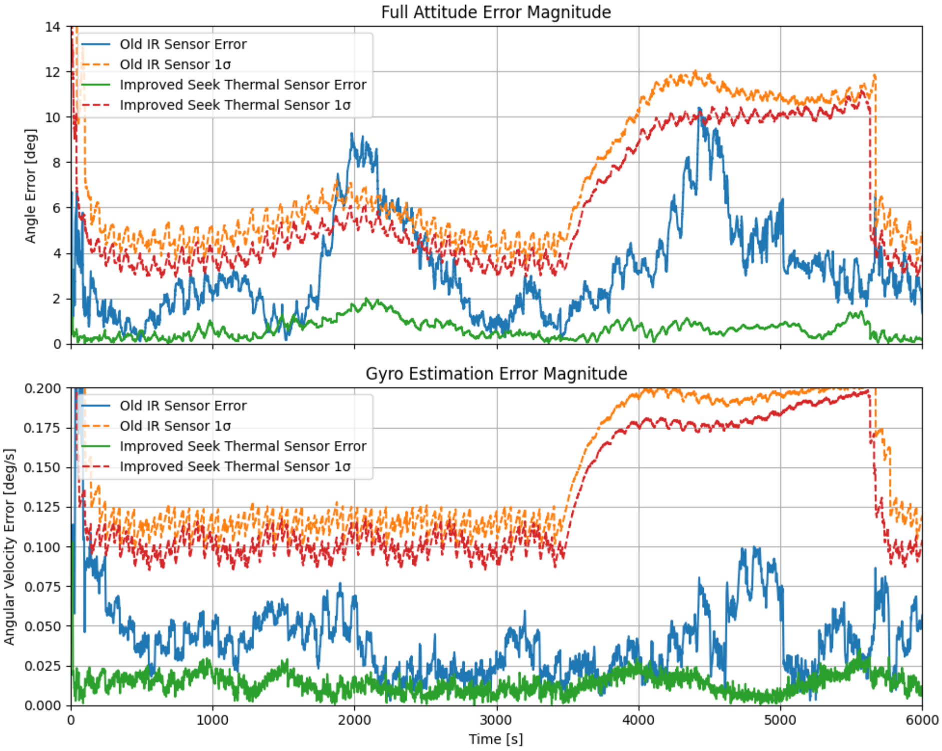

Figure 5.15: Attitude estimation error using the higher-performance IR camera model.

In contrast to the gyro improvement, replacing the IR sensor results in a substantial reduction in estimation error. The worst-case error decreases from approximately 10° to about 2°. The performance improves to errors below 30 arcmin for a large fraction of the time. Additionally, the performance in eclipse improves significantly.

5.5.3. Discussion

The comparison clearly indicates that the dominant limitation in the current system is the accuracy of the optical horizon-based measurement rather than the inertial sensing performance.

While reducing gyro noise improves short-term propagation and slightly tightens the estimation distribution, the overall achievable accuracy remains constrained by the quality of the image-based attitude information. In contrast, improving the IR sensor characteristics directly reduces measurement uncertainty, leading to a substantial improvement in both worst-case and nominal performance.

These results suggest that, within the investigated performance envelope:

- Further improvements in gyro performance yield diminishing returns.
- Optical sensor characteristics (resolution, noise-equivalent temperature difference, optical quality, and calibration stability) dominate the achievable accuracy.

From a system-level perspective, the current design is therefore measurement-limited rather than propagation-limited. Any significant improvement in final pointing accuracy will require improving the optical measurement quality or redesigning the image processing chain to extract more information from the existing sensor.

5.6. Robustness to Variations in Object Brightness

A critical aspect of the image processing algorithm is the selection of threshold values, which determine how objects are identified in the images. Accurately predicting the brightness levels recorded by the sensor is challenging, as the exact space environment cannot be fully replicated on Earth. An attempt to reproduce the closest available environmental analogue was unsuccessful due to the failure of the BEXUS flight (see section 4.7), making it difficult to define precise threshold values in advance.

Despite this uncertainty, testing showed that the algorithm is largely insensitive to the exact threshold values, provided that the Earth detection threshold is set above the background temperature. This remains complicated by the fact that neither the background temperature nor the apparent temperature of the Earth, as measured by the sensor, are precisely known without operating in the actual or a representative environment. Nevertheless, it is reasonable to assume that the Earth will appear approximately at its expected surface temperature, as discussed previously.

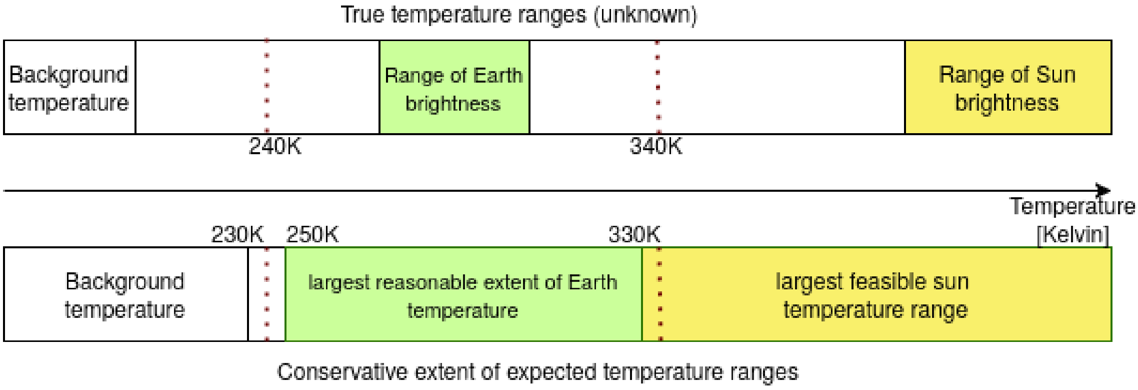

Figure 5.16: Illustration of expected and tested temperature ranges together with the applied thresholds. The green and yellow regions indicate the temperature ranges of the Earth and Sun in the acquired imagery, while the dotted red lines represent the thresholds used in the algorithm.

As illustrated in Figure 5.16, the expected temperature ranges of the Earth and Sun are distinct from each other and from the background. The threshold is therefore chosen between these expected ranges. However, as previously explained, the true temperature ranges are not known. To evaluate robustness, the simulation tool varied the temperatures recorded by the cameras for the background, Earth, and Sun within the ranges shown in Figure 5.16. Within these ranges, the algorithm operated reliably and produced no outliers, except in cases where the Sun appeared at the edge of the Earth's horizon, as previously observed. Outliers were only generated when the Earth detection threshold was lowered to the extent that background noise exceeded it.

5.7. Camera Calibration

As described in section 4.2, the camera distortions for the MLX90640 sensor were established. A set of parameters was found that reduces the mean square error of the acquired and projected points on the image down to 0.057 pixels, resulting in an angular error of $0.35°$ in angular terms. The reprojection of the points with this determined model is shown in Figure 5.17.

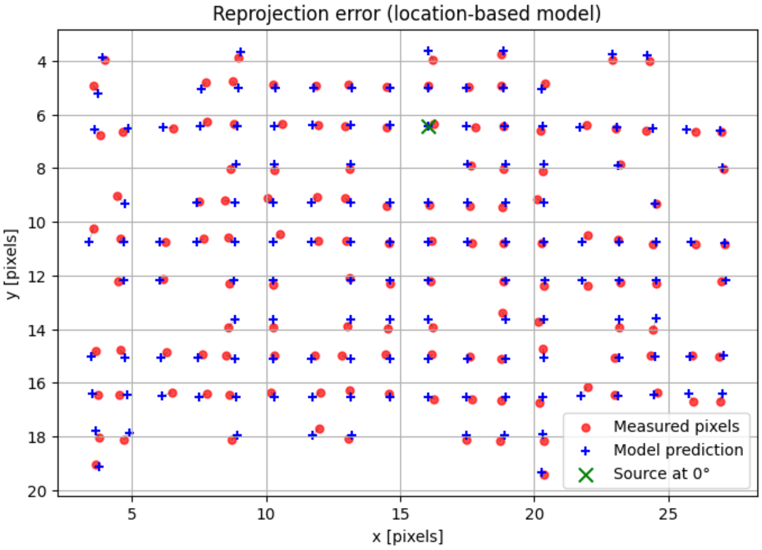

Figure 5.17: Acquired centroids of points in image compared to projected centroids using the optimised projection function including radial and tangential distortions.

It can also be noticed that even with this fit, the points do not perfectly overlap. The majority of this error is likely due to the inaccuracies in the data acquisition, which adds noise to the measured locations. This could be improved in the future by improving the test setup in terms of sensor and target mounting, as well as their alignment.

The resulting values are shown in Table 5.2.

Parameter	**Estimated value**	**Unit**
Focal length f_x	41.64	pixels
Focal length f_y	41.65	pixels
RMS error	0.057	pixels
RMS Angular error	0.35	deg
Radial distortion k_1	-0.6766	-
Radial distortion k_2	-0.6287	-
Tangential distortion p_1	-6.174×10^{-5}	-
Tangential distortion p_2	-1.374×10^{-2}	-
Source offset x_{src}	1.493×10^{-3}	focal lengths
Source offset y_{src}	0.1358	focal lengths

Table 5.2: Estimated intrinsic camera parameters and source offset obtained from geometric calibration.

It is easily noticed that the focal lengths in x and y have approached basically the same value, which is as expected when considering the drawings in the datasheet of the IR sensor, which show a spherical lens that is radially symmetric and assuming that the sensor has square pixels.

5.8. Implementation status of hard- and software

The hardware has been designed and produced for the test flight. Only small issues were discovered, such as a I²C bus for the IR sensors missing pull-up resistors, which have been added after manufacturing. The integration of this sensor into the satellite is to be done, but the mechanical mounting of the PCB in the satellite is designed and a prototype manufactured as shown in Figure 5.18.

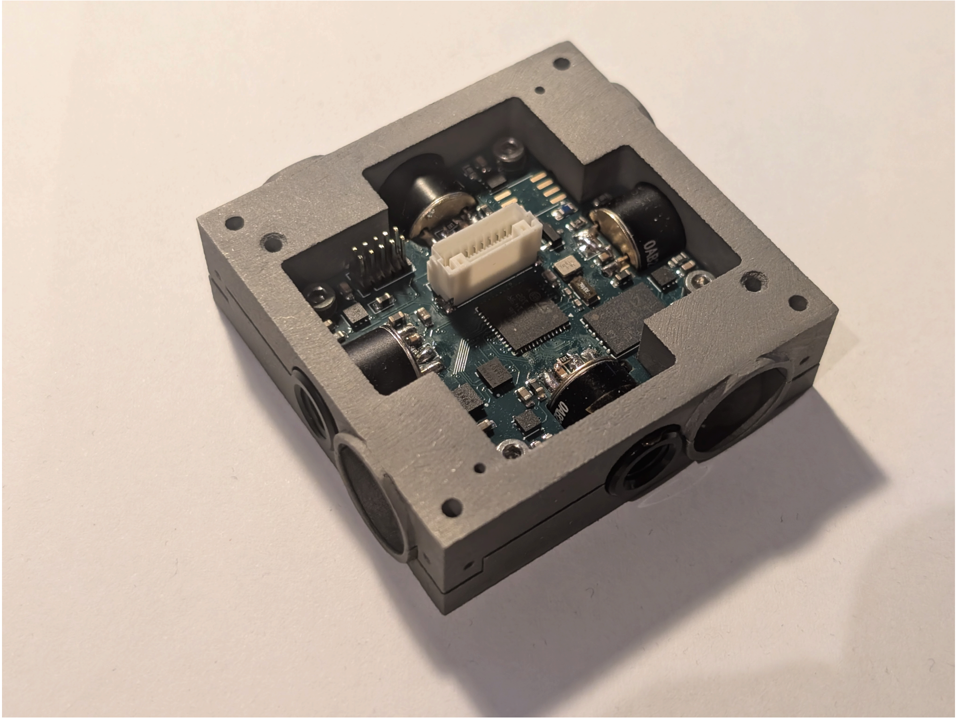

Figure 5.18: PCB integrated into satellite structure.

This hardware still requires another revision to fix these minor I²C bus issues, however, apart from that, it has shown to work well in the various testing done, such as the calibration of the sensors.

In terms of software, the embedded software to perform all the most crucial tasks, like reading out the various sensors, communicating through the RS485 bus as well as processing and storing the sensor data has been implemented. The attitude determination algorithm itself needs to be transferred into embedded C code from the SIMULINK model currently used. For this, the autocoding feature can be used to ensure a flawless implementation without bugs introduced during the development process. The interfaces between this generated code and the existing software then need to be setup, such that the end-to-end functionality is performed on board the embedded hardware. For this purpose, a structure using the FreeRTOS threading library has been setup to maintain the various tasks performing the attitude determination.

6

Conclusion and recommendations

This thesis set out to develop and validate an infrared horizon–based attitude determination approach, suitable for small spacecraft, with two primary objectives: (1) generating reliable attitude measurements from an infrared horizon sensor, supported by a validated synthetic image framework, and (2) integrating these measurements into a full attitude estimation filter running on embedded hardware capable of providing continuous attitude knowledge under realistic operational conditions.

Infrared-based attitude measurements.

The first objective focused on extracting attitude information from infrared imagery of the Earth and the Sun. To enable this, a physics-based image generation framework was developed to produce synthetic infrared images under varying orbital, geometric, and thermal conditions. Beyond image generation, dedicated image processing and estimation algorithms were developed to reliably detect the Earth horizon and Sun and to extract the corresponding reference vectors required for attitude determination. The synthetic framework enabled systematic testing and tuning of these algorithms across a wide range of scenarios, including cases that are difficult or costly to reproduce experimentally. Experimental validation in subsection 4.7.4 using real infrared imagery showed good agreement with the simulated data, confirming that the framework captures the dominant characteristics relevant for attitude estimation, while also revealing discrepancies arising from sensor noise, calibration uncertainty, and environmental variability. Based on these results, the developed algorithms were shown to consistently extract meaningful attitude-related measurements when sufficient thermal contrast and object visibility are present, as shown in section 5.6.

Full attitude estimation filter.

Building on the horizon- and Sun-based measurements, the second objective addressed continuous three-dimensional attitude estimation through sensor fusion. While infrared horizon and Sun measurements provide absolute reference information, they do not by themselves guarantee full instantaneous observability of the spacecraft attitude, particularly when only a single reference vector is available. An attitude estimation filter was therefore developed to combine infrared-derived measurements with inertial sensor data, and optionally magnetometer measurements, enabling full attitude observability through propagation and accumulation of information over time as demonstrated in section 5.3.

This fusion approach is essential for maintaining a consistent three-dimensional attitude estimate during nominal operation and allows the filter to bridge periods of partial or full obscuration. The results in subsection 5.3.3 show that the spacecraft attitude can be continuously estimated even when no objects are visible, and during eclipse phases or off-nadir pointing when only intermittent or single-vector measurements are available. In these cases, inertial propagation maintains attitude continuity, while occasional horizon or Sun observations provide corrective updates that bound drift and restore absolute attitude information. Together, this demonstrates that the filter together with the reference vectors, even if intermittent, provides robust and reliable attitude estimation across the full range of operational conditions considered in this thesis as shown by montecarlo simulations in section 5.4.

From these results, the research questions defined in section 3.1 can be addressed as follows.

What are the key parameters that govern the achievable attitude accuracy of an IR imaging attitude sensor?

The parameters determining achievable attitude accuracy can be grouped into two aspects: the accuracy of the extracted reference vectors and the performance of the attitude estimation filter that fuses them over time.

For Earth-based measurements, accuracy is dominated by the amount of horizon limb visible within the sensor field of view, which depends on field of view, spacecraft altitude, and pointing geometry. A larger visible limb improves geometric conditioning and reduces sensitivity to noise. Sun-based measurements, as effective point sources, are mainly limited by the sensor's angular resolution and signal-to-noise ratio. As shown in section 5.5, it is possible to drastically increase the accuracy of this system by increasing both the FoV as well as the resolution of the imaging sensor.

Filter performance depends not only on reference vector accuracy but also on the temporal availability of multiple reference vectors and the quality of inertial sensors. Periods when two reference vectors are simultaneously visible enable absolute attitude updates, whereas phases with a single or no visible object rely on gyroscope propagation to maintain observability. For spacecraft that are stabilized and not freely rotating, as expected for Delfi Twin, the pointing modes must be analysed such that observations of reference vectors can regularly be made. Consequently, achievable performance is determined both by instantaneous measurement quality and by the duration, frequency, and sequencing of available reference observations, as well as the quality of the intermediate propagation. However, section 5.5 showed that the gyro observations are not the limiting factor for the current configuration.

What is the expected accuracy under nominal operating conditions? Can the proposed attitude determination system provide sufficient accuracy to ensure that separation of the Delfi Twin mission occurs within the prescribed 52.7° half-angle cone?

Under nominal conditions, when the Earth horizon is clearly visible and the Sun is available, the system achieves attitude errors typically within 5°, and in favorable geometries can approach 2°.

During eclipse, the loss of Sun-based corrections reduces instantaneous observability, particularly for one attitude axis. In these periods, errors of approximately 15° are possible, but degradation occurs gradually due to gyroscope propagation with well-calibrated bias. The expected accuracy is reflected in the Kalman filter covariance, and remains well within the 52.7° requirement of Delfi-Twin. Nevertheless, critical operations, such as spacecraft separation, should be performed outside eclipse when possible, and filter covariances should be monitored to ensure confidence. If this requirement is met, the achieved accuracy permits a sufficiently large inner cone within which separation can be triggered, even when accounting for worst-case estimation errors.

What are the best-case and worst-case performance limits?

Best-case performance occurs under favorable observation geometry, with a large visible Earth limb, high thermal contrast, and the Sun simultaneously visible. Under these conditions, errors can drop below 2° as shown in subsection 5.3.2. Worst-case performance arises with limited visible horizon, low thermal contrast, or extended eclipse periods without reference vectors. In these cases, the attitude estimate relies heavily on inertial propagation, increasing sensitivity to gyroscope noise and bias drift, with errors reaching up to at most 20 degrees in the worst case of the montecarlo simulations, as shown in section 5.4. These limits show the operational envelope of the system and highlight the importance of the geometry of the problem andthe variance of the accuracy achievable.

How does the performance compare to alternative sensor combinations, such as Sun sensors, magnetometers, and gyroscope-based systems?

Compared to traditional systems using dedicated Sun sensors, magnetometers, or star trackers, the infrared imaging approach does not achieve the same absolute accuracy. Specialized hardware in these systems provides more precise and continuously observable references, particularly in eclipse. However, such systems come with higher complexity: magnetometers require careful calibration, and star sensors need larger optics and precise pointing.

The infrared-based system's main advantage is its compact size, low weight, and low power consumption of around 40x40x10mm, 12g, and at most 400mW in acquisition mode. The complete sensing and processing chain weighs only a few grams and integrates easily into miniaturized satellites, while providing sufficient accuracy for missions with relaxed pointing requirements, such as *Delfi-Twin*. For these applications, it offers an attractive trade-off between performance and system complexity.

Is an additional sensor such as a magnetometer necessary to have acceptable performance when no object is in view?

The system can operate without a magnetometer, but including one improves initialization reliability, particularly during eclipse when only one reference vector may be intermittently available. Once the filter has converged, performance during eclipse remains acceptable without a magnetometer. For this purpose, the magnetometer does not need high precision; it simply resolves initial ambiguities in poorly observable conditions.

6.1. Recommendations

Based on the results and limitations identified in this thesis, several recommendations are made to further mature the proposed infrared horizon-based attitude determination approach and to increase confidence in its performance under realistic mission conditions.

- **Validation using flight-like data.** While the algorithm has been developed and validated using synthetic imagery and limited real infrared data, a key next step is validation using fully representative flight data. This includes infrared imagery and gyroscope measurements collected simultaneously during a dedicated test flight. Such a campaign was originally planned and remains scheduled for a future flight opportunity in the summer period of 2026. This will be conducted still within the BEXUS programme as offered by the organizers, outside of this thesis, with the same equipment originally developed.
- **Refinement of thresholding and robustness analysis.** Several algorithm parameters, particularly threshold values used in image processing and feature extraction, were selected based on simulation results and limited experimental data. Further image acquisition under a wider range of real-world conditions is recommended to refine these thresholds and assess their sensitivity. The current simulation setup has proven slow, such that only 50 runs have been used. Additionally, more variance in the illumination conditions should be explored and adjusted to match the data collected in the testflight.
- **Embedded implementation and hardware-in-the-loop testing.** To assess practical feasibility for space applications, the algorithm should be implemented on an embedded processing platform representative of a small spacecraft avionics system. Hardware-in-the-loop (HIL) testing is recommended, in which externally simulated infrared imagery and inertial sensor data are streamed to the embedded system in real time. This would enable evaluation of numerical stability, timing behavior, and real-time performance constraints prior to any on-orbit or flight deployment.
- **Computational optimization and throughput analysis.** The current implementation has not been optimized for computational efficiency. Future work should quantify the maximum achievable image processing rate and identify bottlenecks in the pipeline. Targeted optimization of image processing, feature extraction, and filter update steps is recommended to maximize the number of images that can be processed per second within the available computational budget, thereby improving responsiveness and estimation accuracy during dynamic maneuvers.

References

[1] J. Bouwmeester et al. *Presentations of the 2nd PocketQube Workshop (Delft, 2018)*. Feb. 2024. DOI: `10.34894/J4KAQ3`. (accessed: 3.6.2025).

[2] TU Delft. *Delfi-C3*. Accessed: 2025-06-24. URL: `https://www.tudelft.nl/lr/delfi-space/delfi-c3`.

[3] TU Delft. *Delfi Program*. Accessed: 2025-06-24. URL: `https://www.tudelft.nl/lr/delfi-space/delfi-program`.

[4] TU Delft. *Advancing nano-satellite platforms: the Delfi program*. Accessed: 2025-06-24. URL: `https://repository.tudelft.nl/file/File_541244cf-a1de-4654-83f0-8863ec8bb7ea`.

[5] J. Puig-Suari, C. Turner, and W. Ahlgren. “Development of the Standard CubeSat Deployer and a CubeSat Class PicoSatellite”. In: *2001 IEEE Aerospace Conference Proceedings (Cat. No.01TH8542)*. Vol. 1. Mar. 2001, 1/347–1/353 vol.1. DOI: `10.1109/AERO.2001.931726`. (accessed: 20.2.2026).

[6] TU Delft. *The birth of the Delfi-PQ mini satellite*. Accessed: 2025-06-24. URL: `https://www.tudelft.nl/en/stories/articles/the-birth-of-the-delfi-pq-mini-satellite`.

[7] TU Delft. *University satellite in space for 12.5 years*. Accessed: 2025-06-24. URL: `https://www.tudelft.nl/en/2020/ewi/university-satellite-in-space-for-125-years`.

[8] TU Delft. *Satellite Delfi-n3Xt successfully launched*. Accessed: 2025-06-24. URL: `https://delta.tudelft.nl/en/article/satellite-delfi-n3xt-successfully-launched`.

[9] TU Delft. *Delfi-n3Xt back to life after 7 years of silence*. Accessed: 2025-06-24. URL: `https://www.tudelft.nl/en/2021/lr/delfi-n3xt-back-to-life-after-7-years-of-silence`.

[10] TU Delft Space Institute. *Delfi-PQ: The Delfi space program is preparing for launch of its third satellite*. Accessed: 2025-06-24. URL: `https://spaceinstitute.tudelft.nl/delfi-pq-the-delfi-space-program-is-preparing-for-launch-of-its-third-satellite/`.

[11] Erik Kulu. *Delfi-C3*. https://www.nanosats.eu/sat/delfi-c3.html. (accessed: 20.2.2026).

[12] Erik Kulu. *Delfi-n3Xt*. https://www.nanosats.eu/sat/delfi-n3xt.html. (accessed: 20.2.2026).

[13] Erik Kulu. *Delfi-PQ 1*. https://www.nanosats.eu/sat/delfi-pq-1.html. (accessed: 20.2.2026).

[14] TU Delft Space Institute. *TU Delft Space Institute – Top research for space*. Accessed: 2025-06-24. URL: `https://spaceinstitute.tudelft.nl/`.

[15] Simona Cundari. “Formation-Flying of Picosatellites using Differential Drag as a Control Methodology”. laurea. Politecnico di Torino, July 2024. (accessed: 1.3.2026).

[16] S. Speretta, R.K. van der Zwaard, and M.S. Uludag. “Autonomous Formation Flying in the Traffic: 15th IAA Symposium on Small Satellites for Earth Observation”. In: *Proceedings Small Satellites for Earth Observation*. DLR, 2025, pp. 1–8.

[17] Jian Guo and Congying Han. “Where Is the Limit: The Analysis of CubeSat ADCS Performance”. In: *30th Annual AIAA/USU Conference on Small Satellites*. 2016.

[18] Vittorio Franzese et al. “System Design of the POQUITO PocketQube Mission”. In: *CEAS Space Journal* (May 2025). ISSN: 1868-2510. DOI: `10.1007/s12567-025-00624-9`. (accessed: 3.6.2025).

[19] Marianna Centrella. “Mission and System Design of a Formation-Flying Picosatellites Cluster: A Technology Demonstration Mission for Space Situational Awareness Improvement”. laurea. Politecnico di Torino, Oct. 2023. (accessed: 20.2.2026).

[20] F. Landis Markley and John L. Crassidis. “Sensors and Actuators”. In: *Fundamentals of Spacecraft Attitude Determination and Control*. Ed. by F. Landis Markley and John L. Crassidis. New York, NY: Springer, 2014, pp. 123–181. ISBN: 978-1-4939-0802-8. DOI: `10.1007/978-1-4939-0802-8_4`. (accessed: 3.6.2025).

[21] James R. Wertz and Wiley J. Larson, eds. *Space Mission Analysis and Design*. 3rd ed. El Segundo, CA: Microcosm Press, 1999. ISBN: 978-0-7923-5901-2.

[22] National Centers for Environmental Information. *World Magnetic Model (WMM)*. 2025. URL: `https://www.ncei.noaa.gov/products/world-magnetic-model` (accessed: 3.6.2025).

[23] W. W. Chow et al. "The Ring Laser Gyro". In: *Reviews of Modern Physics* 57.1 (Jan. 1985), pp. 61–104. DOI: `10.1103/RevModPhys.57.61`. (accessed: 3.6.2025).

[24] N.F. de Rooij et al. "MEMS for Space". In: *TRANSDUCERS 2009 - 2009 International Solid-State Sensors, Actuators and Microsystems Conference*. June 2009, pp. 17–24. DOI: `10.1109/SENSOR.2009.5285575`. (accessed: 3.6.2025).

[25] Chandradip Patel et al. "Temperature Effects on the Performance and Reliability of MEMS Gyroscope Sensors". In: vol. 1. Jan. 2009. DOI: `10.1115/InterPACK2009-89370`.

[26] J. P. Eastwood et al. "In Flight Performance of the MAGIC Magnetoresistive Magnetometer on the RadCube CubeSat". In: *Space Science Reviews* 221.4 (May 2025), p. 45. ISSN: 1572-9672. DOI: `10.1007/s11214-025-01170-w`. (accessed: 3.6.2025).

[27] Bagus Riwanto. "CubeSat Attitude System Calibration and Testing". Master's thesis. Espoo, Finland: Aalto University, Aug. 2015.

[28] John A. Christian. "A Tutorial on Horizon-Based Optical Navigation and Attitude Determination With Space Imaging Systems". In: *IEEE Access* 9 (2021), pp. 19819–19853. ISSN: 2169-3536. DOI: `10.1109/ACCESS.2021.3051914`. (accessed: 12.6.2025).

[29] T.D. Cornall, G.K. Egan, and A. Price. "Aircraft Attitude Estimation from Horizon Video". In: *Electronics Letters* 42.13 (June 2006), pp. 744–745. ISSN: 0013-5194, 1350-911X. DOI: `10.1049/el:20060547`. (accessed: 15.3.2025).

[30] John A. Christian. "Accurate Planetary Limb Localization for Image-Based Spacecraft Navigation". In: *Journal of Spacecraft and Rockets* 54.3 (2017), pp. 708–730. ISSN: 0022-4650. DOI: `10.2514/1.A33692`. (accessed: 12.6.2025).

[31] John A. Christian et al. "Geometric Calibration of the Orion Optical Navigation Camera Using Star Field Images". In: *The Journal of the Astronautical Sciences* 63.4 (Dec. 2016), pp. 335–353. ISSN: 2195-0571. DOI: `10.1007/s40295-016-0091-3`. (accessed: 12.6.2025).

[32] S.M. Ettinger et al. "Vision-Guided Flight Stability and Control for Micro Air Vehicles". In: *IEEE/RSJ International Conference on Intelligent Robots and System*. Vol. 3. Lausanne, Switzerland: IEEE, 2002, pp. 2134–2140. ISBN: 978-0-7803-7398-3. DOI: `10.1109/IRDS.2002.1041582`. (accessed: 15.3.2025).

[33] Adrian Carrio, Hriday Bavle, and Pascual Campoy. "Attitude Estimation Using Horizon Detection in Thermal Images". In: *International Journal of Micro Air Vehicles* 10.4 (Dec. 2018), pp. 352–361. ISSN: 1756-8293. DOI: `10.1177/1756829318804761`. (accessed: 15.3.2025).

[34] János Takátsy et al. "Attitude Determination for Nano-Satellites – II. Dead Reckoning with a Multiplicative Extended Kalman Filter". In: *Experimental Astronomy* 53.1 (Feb. 2022), pp. 209–223. ISSN: 0922-6435, 1572-9508. DOI: `10.1007/s10686-021-09818-5`. arXiv: `2111.13193 [astro-ph]`. (accessed: 20.2.2026).

[35] GeeksforGeeks. *calibrateCamera() OpenCV in Python*. `https://www.geeksforgeeks.org/machine-learning/calibratecamera-opencv-in-python/`. Accessed 2026-02-20. 2024.

[36] Zhengyou Zhang. *A Flexible New Technique for Camera Calibration*. 2000. DOI: `10.1109/34.888718`.

[37] J. Heikkila and O. Silven. "A Four-Step Camera Calibration Procedure with Implicit Image Correction". In: *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. June 1997, pp. 1106–1112. DOI: `10.1109/CVPR.1997.609468`. (accessed: 15.6.2025).

[38] Kornél Kapás et al. *Attitude Determination for Nano-Satellites – I. Spherical Projections for Large Field of View Infrasensors*. Apr. 2021. DOI: `10.48550/arXiv.2104.10960`.

[39] Nobuyuki Otsu. “A Threshold Selection Method from Gray-Level Histograms”. In: *IEEE Transactions on Systems, Man, and Cybernetics* 9.1 (Jan. 1979), pp. 62–66. ISSN: 2168-2909. DOI: 10.1109/TSMC.1979.4310076. (accessed: 20.2.2026).

[40] Jochen Barf. “Development and Implementation of an Image-Processing-Based Horizon Sensor for Sounding Rockets”. In: (July 2017). Master’s thesis, Department of Computer Science, Electrical and Space Engineering; available at urlhttps://urn.kb.se/resolve?urn=urn:nbn:se:ltu:diva-65553.

[41] Yann LeCun et al. “Gradient-Based Learning Applied to Document Recognition”. In: *Proceedings of the IEEE* 86.11 (1998), pp. 2278–2324. DOI: 10.1109/5.726791.

[42] Andrej Karpathy. *CS231n: Convolutional Neural Networks for Visual Recognition*. Accessed: 2025-06-24. URL: http://cs231n.github.io/.

[43] Andrew G. Howard et al. “MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications”. In: *arXiv preprint arXiv:1704.04861* (2017).

[44] John Canny. “A Computational Approach to Edge Detection”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 8.6 (1986), pp. 679–698. DOI: 10.1109/TPAMI.1986.4767851. URL: https://doi.org/10.1109/TPAMI.1986.4767851.

[45] Devin T. Renshaw and John A. Christian. “Subpixel Localization of Isolated Edges and Streaks in Digital Images”. In: *Journal of Imaging* 6.5 (May 2020). ISSN: 2313-433X. DOI: 10.3390/jimaging6050033. (accessed: 20.2.2026).

[46] D. Modenini. “Attitude Determination from Ellipsoid Observations: A Modified Orthogonal Procrustes Problem”. In: *Journal of Guidance, Control, and Dynamics* 41.10 (Oct. 2018), pp. 2324–2326. ISSN: 0731-5090. DOI: 10.2514/1.G003610. (accessed: 12.6.2025).

[47] Benjamin Braun and Jochen Barf. “Image Processing Based Horizon Sensor for Estimating the Orientation of Sounding Rockets, Launch Vehicles and Spacecraft”. In: *CEAS Space Journal* 15.3 (2022), pp. 509–533. DOI: 10.1007/s12567-022-00461-0.

[48] Peter B. Stetson. “DAOPHOT: A Computer Program for Crowded-Field Stellar Photometry”. In: *Publications of the Astronomical Society of the Pacific* 99 (Mar. 1987), p. 191. ISSN: 0004-6280. DOI: 10.1086/131977. (accessed: 25.10.2025).

[49] C. Handgraaf. “Design and Testing of a Low-Budget 2□in□1 Star Tracker and Sun Sensor for CubeSats”. Master’s thesis, Aerospace Engineering. MA thesis. Delft, The Netherlands: Delft University of Technology, Jan. 2025. URL: https://resolver.tudelft.nl/uuid:8b1c675e-e5ec-4263-9ff2-2206820ec78a.

[50] T.F. Chan and L.A. Vese. “Active Contours without Edges”. In: *IEEE Transactions on Image Processing* 10.2 (Feb. 2001), pp. 266–277. ISSN: 1941-0042. DOI: 10.1109/83.902291. (accessed: 2.9.2025).

[51] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. “Snakes: Active Contour Models”. In: *International Journal of Computer Vision* 1.4 (Jan. 1988), pp. 321–331. ISSN: 1573-1405. DOI: 10.1007/BF00133570. (accessed: 2.9.2025).

[52] Y. Lecun et al. “Gradient-Based Learning Applied to Document Recognition”. In: *Proceedings of the IEEE* 86.11 (Nov. 1998), pp. 2278–2324. ISSN: 1558-2256. DOI: 10.1109/5.726791. (accessed: 20.2.2026).

[53] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. “ImageNet Classification with Deep Convolutional Neural Networks”. In: *Commun. ACM* 60.6 (May 2017), pp. 84–90. ISSN: 0001-0782. DOI: 10.1145/3065386. (accessed: 20.2.2026).

[54] Alex Kendall, Matthew Grimes, and Roberto Cipolla. *PoseNet: A Convolutional Network for Real-Time 6-DOF Camera Relocalization*. Feb. 2016. DOI: 10.48550/arXiv.1505.07427. arXiv: 1505.07427 [cs]. (accessed: 20.2.2026).

[55] Andrew G. Howard et al. *MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications*. Apr. 2017. DOI: 10.48550/arXiv.1704.04861. arXiv: 1704.04861 [cs]. (accessed: 20.2.2026).

[56] Benoit Jacob et al. *Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference*. Dec. 2017. DOI: `10.48550/arXiv.1712.05877`. arXiv: `1712.05877 [cs]`. (accessed: 20.2.2026).

[57] Wenyu Wu, Zhonghe Jin, and Chunye Wang. “Satellite Position and Attitude Estimation Using an Infrared Earth Sensor”. In: *Infrared Physics & Technology* 141 (Sept. 2024), p. 105381. ISSN: 1350-4495. DOI: `10.1016/j.infrared.2024.105381`. (accessed: 20.2.2026).

[58] Valentin Peretroukhin, Lee Clement, and Jonathan Kelly. “Reducing Drift in Visual Odometry by Inferring Sun Direction Using a Bayesian Convolutional Neural Network”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. May 2017, pp. 2035–2042. DOI: `10.1109/ICRA.2017.7989235`. (accessed: 20.2.2026).

[59] Weixin Wang, Kanishke Gamagedara, and Taeyoung Lee. *On the Observability of Attitude with Single Direction Measurements*. June 2021. DOI: `10.48550/arXiv.2008.13067`. arXiv: `2008.13067 [math]`. (accessed: 28.1.2026).

[60] Pedro Batista, Carlos Silvestre, and Paulo Oliveira. “GES Attitude Observers – Part II: Single Vector Observations”. In: *IFAC Proceedings Volumes* 44.1 (Jan. 2011), pp. 2991–2996. ISSN: 14746670. DOI: `10.3182/20110828-6-IT-1002.01114`. (accessed: 28.1.2026).

[61] D.H. Titterton and J.L. Weston. *Strapdown Inertial Navigation Technology*. 2nd ed. IET, 2004. Chap. 5.

[62] R.E. Kalman. “A New Approach to Linear Filtering and Prediction Problems”. In: *Journal of Basic Engineering* 82 (1960), pp. 35–45.

[63] N. Barbour and G.T. Schmidt. *Kalman Filtering: Theory and Practice Graduate*. Wiley, 2012.

[64] M.S. Grewal and A.P. Andrews. *Kalman Filtering: Theory and Practice with MATLAB*. Wiley, 2013.

[65] P.S. Maybeck. *Stochastic Models, Estimation, and Control*. Vol. 1. Academic Press, 1979.

[66] Simon J. Julier and Jeffrey K. Uhlmann. “A New Method for the Nonlinear Transformation of Means and Covariances in Filters and Estimators”. In: *Automatica* 34.11 (1997), pp. 1967–1973.

[67] E.A. Wan and R. Van der Merwe. “The Unscented Kalman Filter for Nonlinear Estimation”. In: *Proceedings of the IEEE* 94.3 (2000), pp. 497–522.

[68] J. Russell Carpenter and Christopher N. D’Souza. *Navigation Filter Best Practices*. NASA Technical Publication NASA/TP-2018-219822. Public domain technical report on navigation filter design and Kalman filtering best practices. Hampton, VA, USA: NASA Engineering, Safety Center, National Aeronautics, and Space Administration, Apr. 2018. URL: `https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20180003657.pdf`.

[69] Joan Solà. *Quaternion Kinematics for the Error-State Kalman Filter*. 2017. arXiv: `1711.02508`. URL: `https://arxiv.org/abs/1711.02508` (accessed: 19.10.2025).

[70] F. Landis Markley and Frank H. Bauer. *Attitude Error Representations for Kalman Filtering*. Tech. rep. NASA/TP-2002-60647. NASA, 2002.

[71] John L. Crassidis and F. Landis Markley. *Fundamentals of Spacecraft Attitude Determination and Control*. Springer, 2016.

[72] Melexis. *Far Infrared Thermal Sensor Array (32×24 RES)*. 2025. URL: `https://www.melexis.com/en/product/mlx90640/far-infrared-thermal-sensor-array` (accessed: 19.10.2025).

[73] Zachary Tolar. “Development of Small-Scale and Low-Power Attitude Determination System for Nanoscale Satellites by Infrared Earth-Imaging Sensors”. Undergraduate Honors thesis on IR-based attitude determination systems. Honors Thesis, Mechanical Engineering. Fayetteville, AR, USA: University of Arkansas, Fayetteville, May 2019. URL: `https://scholarworks.uark.edu/meeguht/82/`.

[74] Terabee SAS. *TR-EVO-T33-UART Thermal Sensor*. Product page on Digi-Key. 2025. URL: `https://www.digikey.nl/nl/products/detail/terabee-sas/TR-EVO-T33-UART/13682886` (accessed: 19.10.2025).

[75] Seek Thermal. *Mosaic Core OEM Thermal Camera*. 2025. URL: `https://www.thermal.com/mosaic-core.html` (accessed: 19.10.2025).

[76] Wenlin Yan, Maria Bastos, and A. Magalhães. “Performance Assessment of the Android Smartphone’s IMU in a GNSS/INS Coupled Navigation Model”. In: *IEEE Access* PP (Nov. 2019), pp. 1–1. DOI: 10.1109/ACCESS.2019.2956239.

[77] TDK InvenSense. *ICM-45686 6-Axis Motion Sensor Datasheet*. https://invensense.tdk.com/products/motion-tracking/6-axis/icm-45686/. Accessed 2026-02-20. 2024.

[78] Bosch Sensortec. *BMI423 Inertial Measurement Unit Datasheet*. https://www.bosch-sensortec.com/products/motion-sensors/imus/bmi423/. Accessed 2026-02-20. 2024.

[79] STMicroelectronics. *LSM6DSV Inertial Measurement Unit Datasheet*. https://www.st.com/en/mems-and-sensors/lsm6dsv.html. Accessed 2026-02-20. 2024.

[80] Jong Tai Jang et al. “Analysis of State Estimation Drift on a MAV Using PX4 Autopilot and MEMS IMU During Dead-reckoning”. In: *2020 IEEE Aerospace Conference*. Mar. 2020, pp. 1–11. DOI: 10.1109/AERO47225.2020.9172736. (accessed: 22.10.2025).

[81] Oliver J. Woodman. “An Introduction to Inertial Navigation”. In: *University of Cambridge Technical Report* (2007). Available as PDF from University of Cambridge.

[82] Tobias Hiller et al. “Origins and Mechanisms of Bias Instability Noise in a Three-Axis Mode-Matched MEMS Gyroscope”. In: *Journal of Microelectromechanical Systems* 28.4 (Aug. 2019), pp. 586–596. ISSN: 1941-0158. DOI: 10.1109/JMEMS.2019.2921607. (accessed: 22.10.2025).

[83] Róbert Kovács and Viktor Józsa. “Thermal Analysis of the SMOG-1 PocketQube Satellite”. In: *Applied Thermal Engineering* 139 (July 2018), pp. 506–513. ISSN: 1359-4311. DOI: 10.1016/j.applthermaleng.2018.05.020. (accessed: 20.2.2026).

[84] *Tutorial on Computer Vision with Fisheye Cameras*. https://plaut.github.io/fisheye_tutorial/. (accessed: 24.10.2025).

[85] Duane C. Brown. “Close-Range Camera Calibration”. In: *Photogrammetric Engineering* 37.8 (1971), pp. 855–866.

[86] A. E. Conrady. “Decentred Lens-Systems”. In: *Monthly Notices of the Royal Astronomical Society* 79.5 (Mar. 1919), pp. 384–390. ISSN: 0035-8711. DOI: 10.1093/mnras/79.5.384. (accessed: 25.10.2025).

[87] Faiz Muhammad Chaudhry et al. *Deep-BrownConrady: Prediction of Camera Calibration and Distortion Parameters Using Deep Learning and Synthetic Data*. Jan. 2025. DOI: 10.48550/arXiv.2501.14510. arXiv: 2501.14510 [cs]. (accessed: 25.10.2025).

[88] *Photutils — Photutils 2.3.0*. https://photutils.readthedocs.io/en/stable/. (accessed: 28.2.2026).

[89] Astropy Collaboration. *Solar System Ephemerides*. https://docs.astropy.org/en/latest/coordinates/solarsystem.html. Version v7.2.Dev771+gcba9f070b. 2025. (accessed: 25.10.2025).

[90] *Liberfa/Erfa*. liberfa. Oct. 2025. (accessed: 25.10.2025).

[91] A. T. Mecherikunnel and J. Richmond. *Spectral Distribution of Solar Radiation*. Tech. rep. NASA-TM-82021. NASA Goddard Space Flight Center, Sept. 1980. URL: https://ntrs.nasa.gov/citations/19810016493 (accessed: 25.10.2025).

[92] J. -P. Williams et al. “The Global Surface Temperatures of the Moon as Measured by the Diviner Lunar Radiometer Experiment”. In: *Icarus*. Lunar Reconnaissance Orbiter - Part II 283 (Feb. 2017), pp. 300–325. ISSN: 0019-1035. DOI: 10.1016/j.icarus.2016.08.012. (accessed: 25.10.2025).

[93] Peter Shirley. *Ray Tracing in One Weekend*. Online tutorial/book. 2016. URL: https://www.realtimerendering.com/books/RayTracingInOneWeekend.html (accessed: 4.3.2026).

[94] Christian Xianyang Wu et al. “Autonomous Optical Navigation for Small Spacecraft in Cislunar Space”. In: (Sept. 2022). (accessed: 7.1.2026).

[95] Robert Grover Brown. *Introduction to Random Signals and Applied Kalman Filtering*. New York : J. Wiley, 1992. ISBN: 978-0-471-52573-8. (accessed: 16.1.2026).

A

Algorithm SIMULINK screenshots

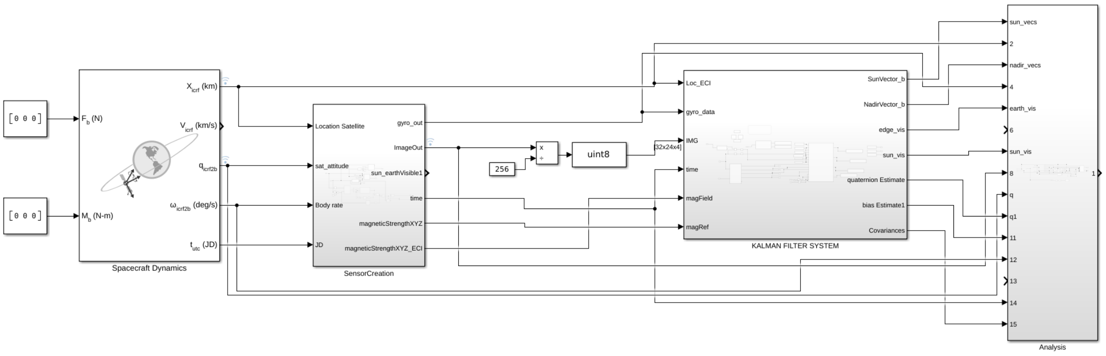

Figure A.1: Overview of SIMULINK simulation

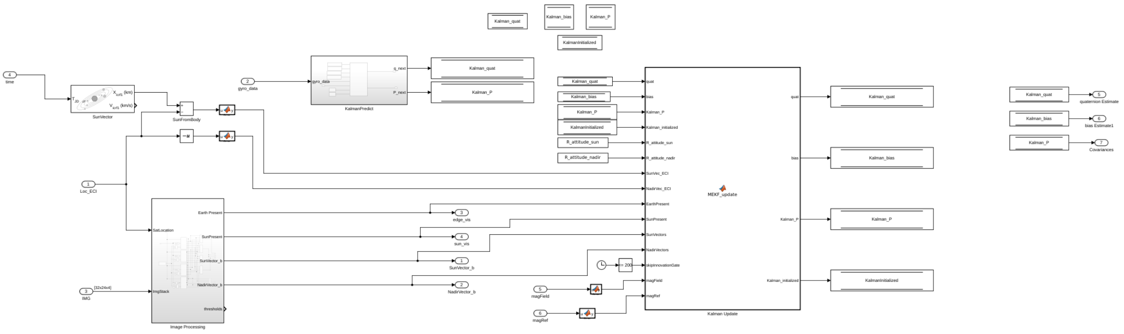

Figure A.2: Attitude determination subsystem

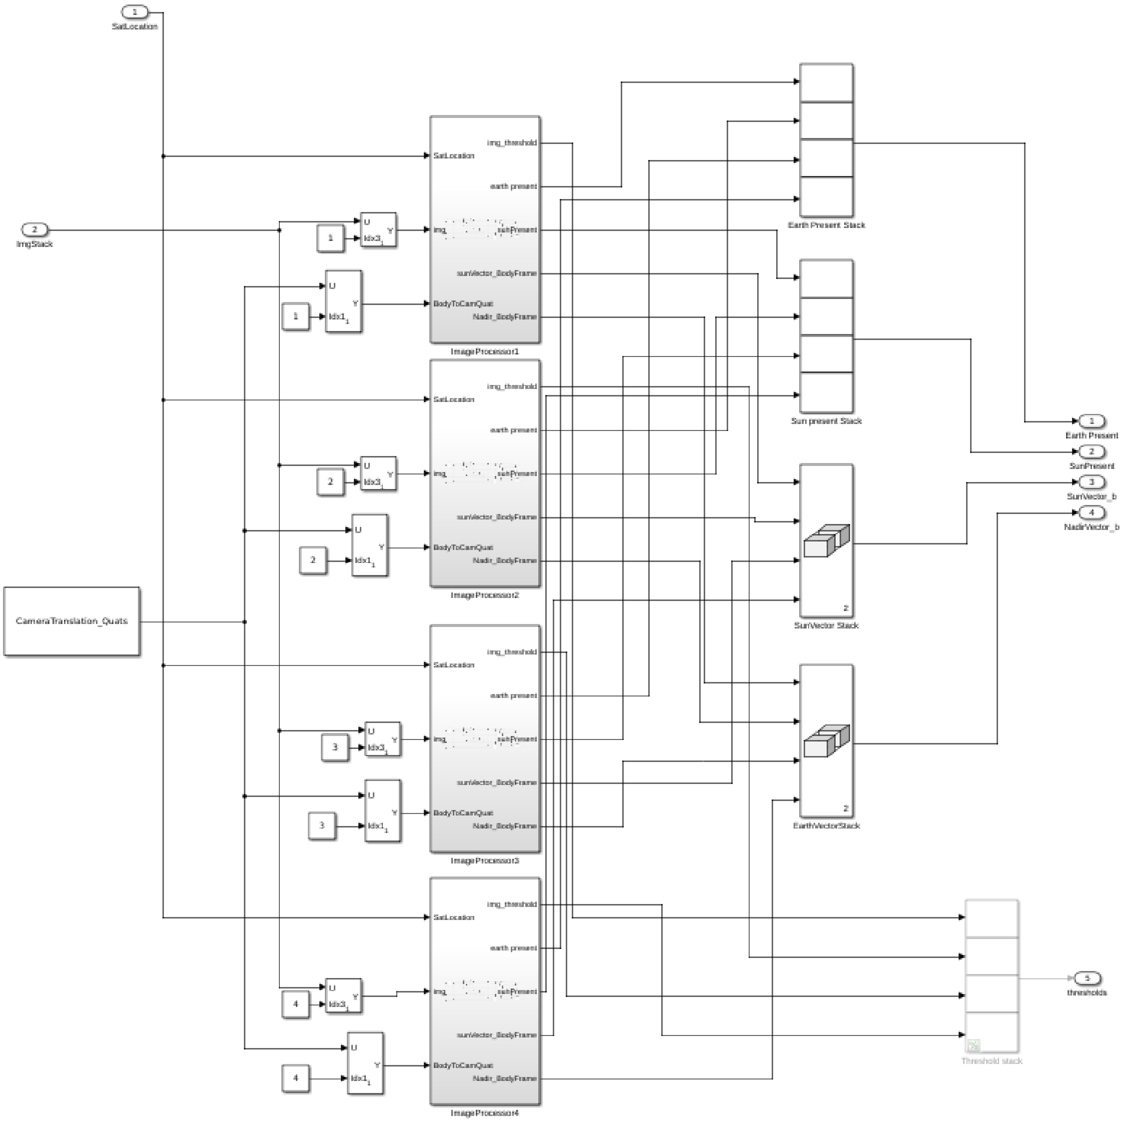

Figure A.3: Overview of image processing subsystem for all 4 image channels

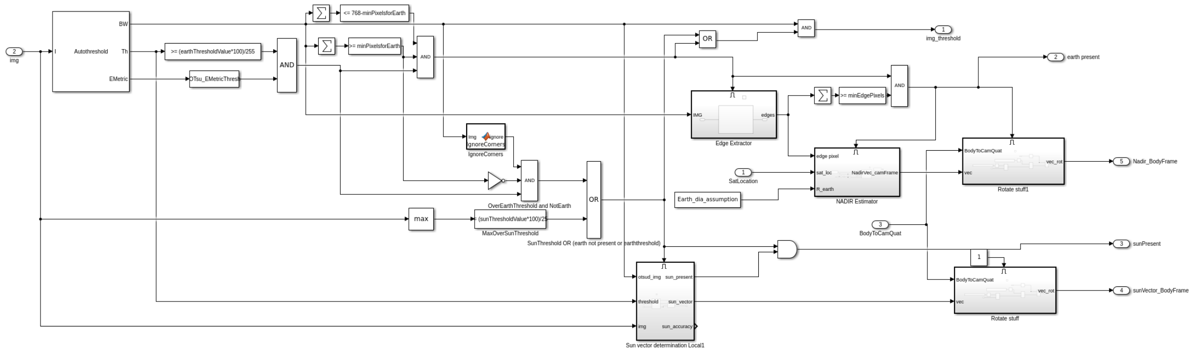

Figure A.4: Overview of image processing pipeline

B

Hardware Design Schematics and PCB Layout

B.1. PCB Schematics

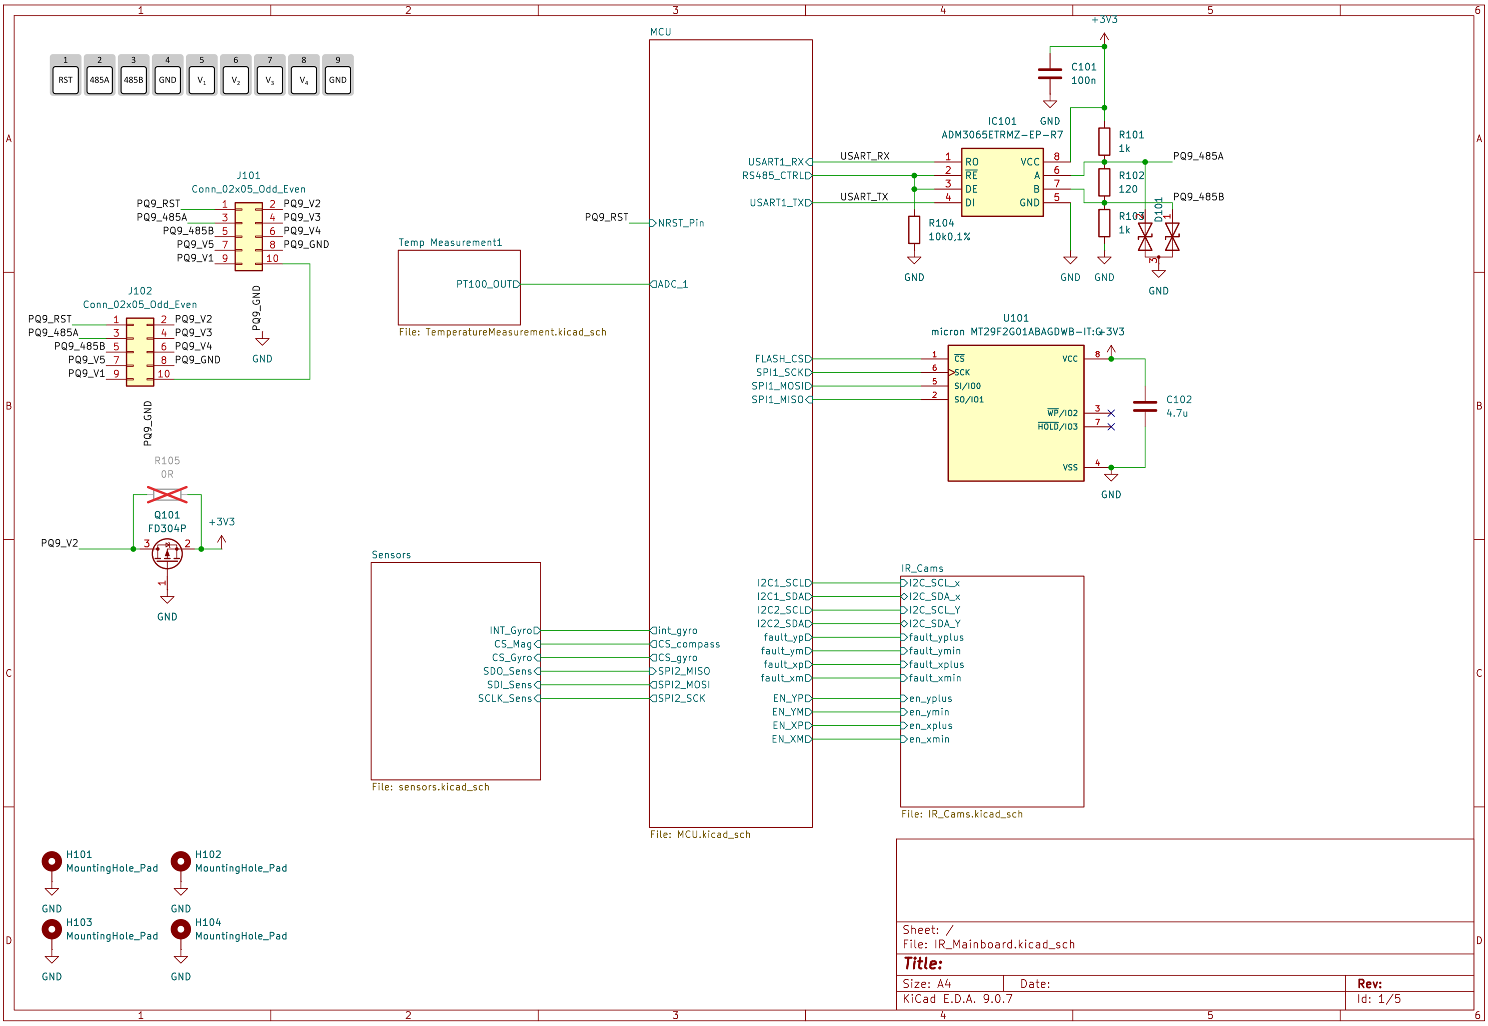

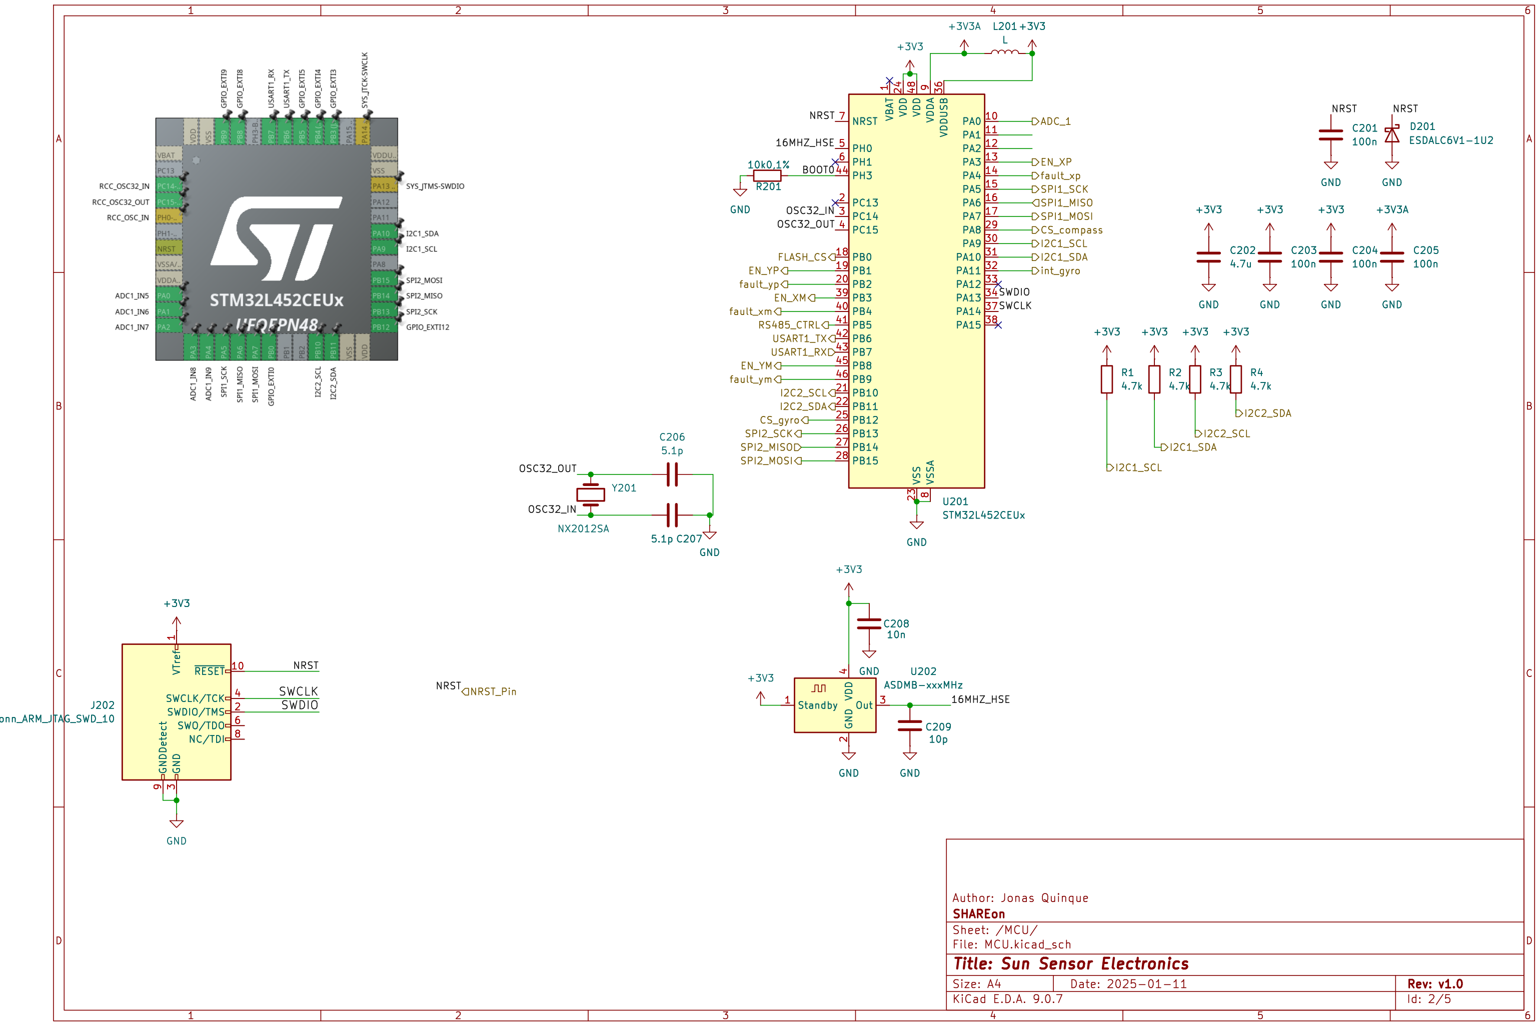
Author: Jonas Quinque
SHAREon
Sheet: /MCU/
File: MCU.kicad_sch
Title: Sun Sensor Electronics
Size: A4
Date: 2025-01-11
Rev: v1.0
KiCad E.D.A. 9.0.7
Id: 2/5
STM32L452CEUx
UFQFPN48
U201
STM32L452CEUx
U202
ASDMB-xxxMHz
16MHZ_HSE
Y201
NX2012SA
OSC32_OUT
OSC32_IN
C206
5.1p
C207
5.1p
C208
10n
C209
10p
R201
10k0.1%
BOOT0
L201
+3V3A
+3V3
GND
NRST
C201
100n
D201
ESDALC6V1-1U2
C202
4.7u
C203
100n
C204
100n
C205
100n
R1
4.7k
R2
4.7k
R3
4.7k
R4
4.7k
I2C1_SCL
I2C1_SDA
I2C2_SCL
I2C2_SDA
ADC_1
EN_XP
fault_xp
SPI1_SCK
SPI1_MISO
SPI1_MOSI
CS_compass
int_gyro
SWDIO
SWCLK
FLASH_CS
EN_YP
fault_yp
EN_XM
fault_xm
RS485_CTRL
USART1_TX
USART1_RXD
EN_YM
fault_ym
CS_gyro
SPI2_SCK
SPI2_MISO
SPI2_MOSI
J202
onn_ARM_JTAG_SWD_10
VTref
RESET
SWCLK/TCK
SWDIO/TMS
SWO/TDO
NC/TDI
GNDDetect
NRST_Pin
Standby
Out

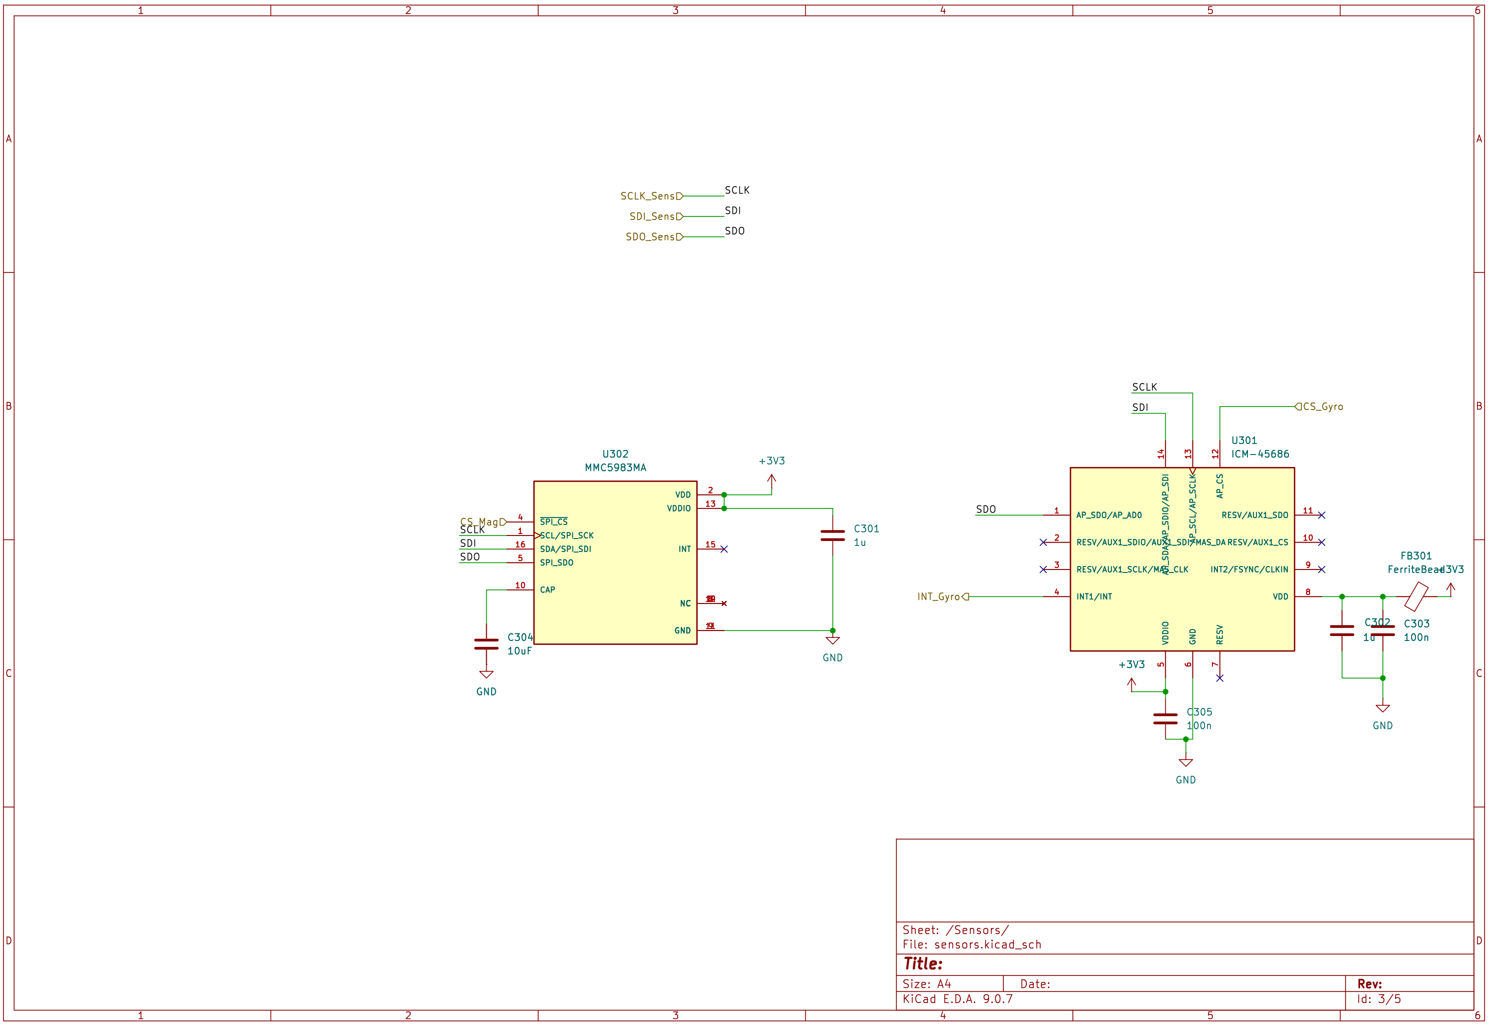
SCLK_Sens
SDI_Sens
SDO_Sens
SCLK
SDI
SDO
U302
MMC5983MA
+3V3
VDD
VDDIO
SPI_CS
SCL/SPI_SCK
SDA/SPI_SDI
SPI_SDO
CAP
INT
NC
GND
CS_Mag
C301
1u
C304
10uF
GND
U301
ICM-45686
CS_Gyro
AP_SDO/AP_AD0
RESV/AUX1_SDIO/AUX1_SDI/MAS_DA
RESV/AUX1_SCLK/MAS_CLK
INT1/INT
AP_SDA/AP_SDIO/AP_SDI
AP_SCL/AP_SCLK
AP_CS
RESV/AUX1_SDO
RESV/AUX1_CS
INT2/FSYNC/CLKIN
VDD
VDDIO
RESV
INT_Gyro
FB301
FerriteBead
C302
1u
C303
100n
C305
100n
Sheet: /Sensors/
File: sensors.kicad_sch
Title:
Size: A4
Date:
Rev:
KiCad E.D.A. 9.0.7
Id: 3/5

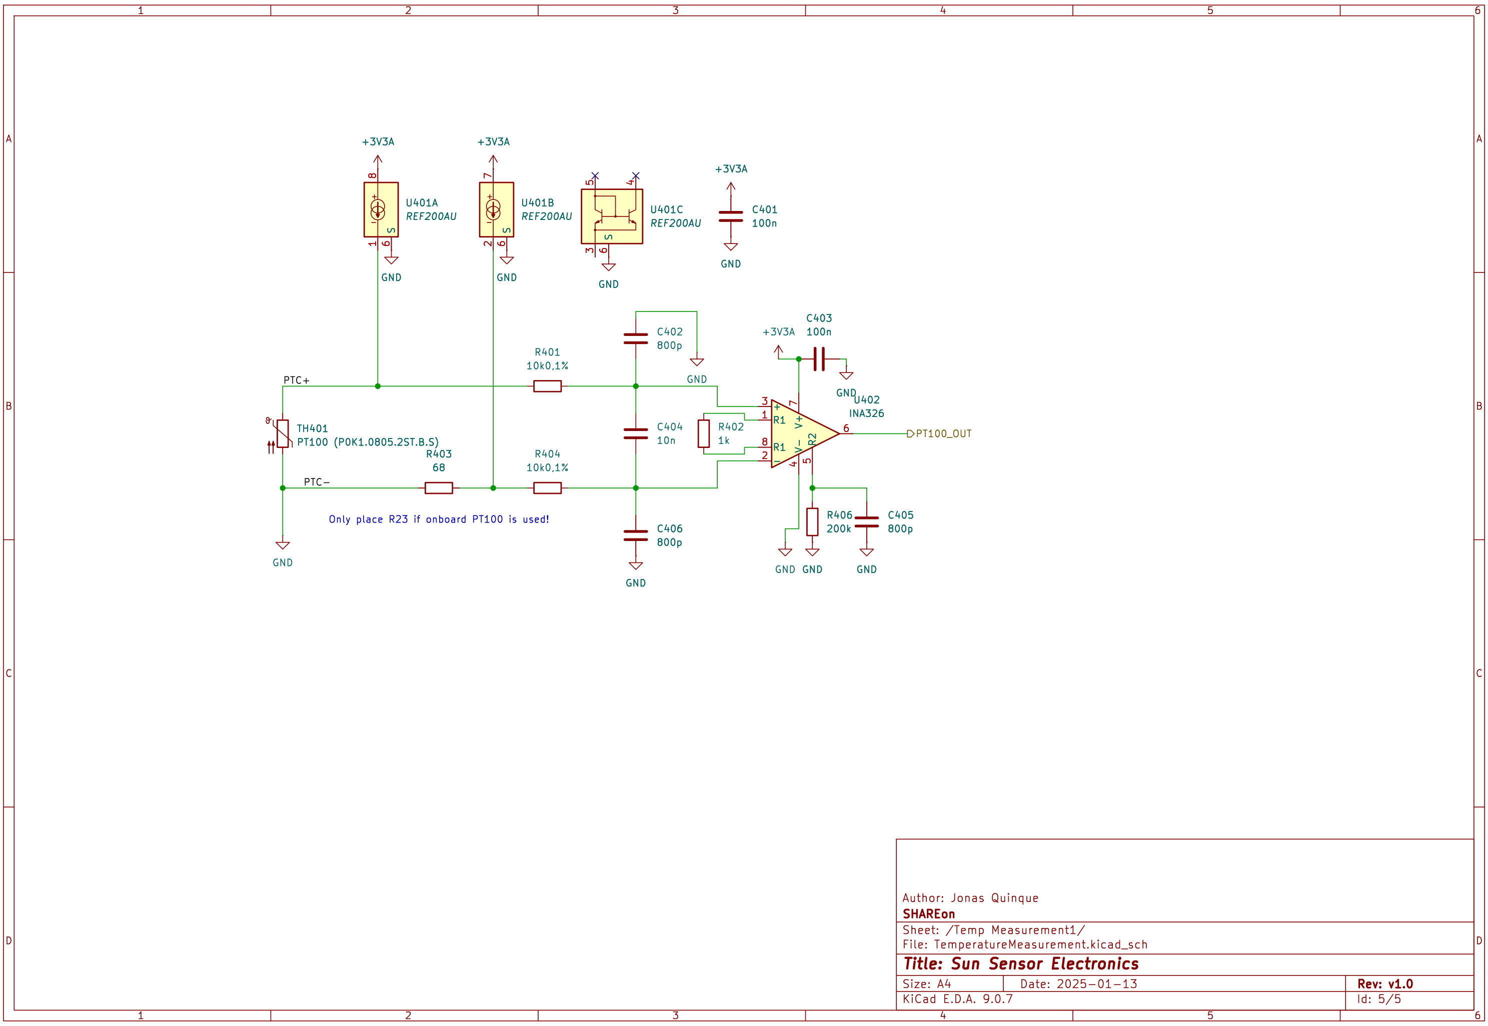
+3V3A
U401A
REF200AU
GND
+3V3A
U401B
REF200AU
GND
U401C
REF200AU
GND
+3V3A
C401
100n
GND
C402
800p
GND
+3V3A
C403
100n
GND
R401
10k0,1%
PTC+
TH401
PT100 (P0K1.0805.2ST.B.S)
R403
68
R404
10k0,1%
PTC−
C404
10n
R402
1k
U402
INA326
PT100_OUT
R406
200k
C405
800p
C406
800p
GND
Only place R23 if onboard PT100 is used!
Author: Jonas Quinque
SHAREon
Sheet: /Temp Measurement1/
File: TemperatureMeasurement.kicad_sch
Title: Sun Sensor Electronics
Size: A4
Date: 2025−01−13
Rev: v1.0
KiCad E.D.A. 9.0.7
Id: 5/5

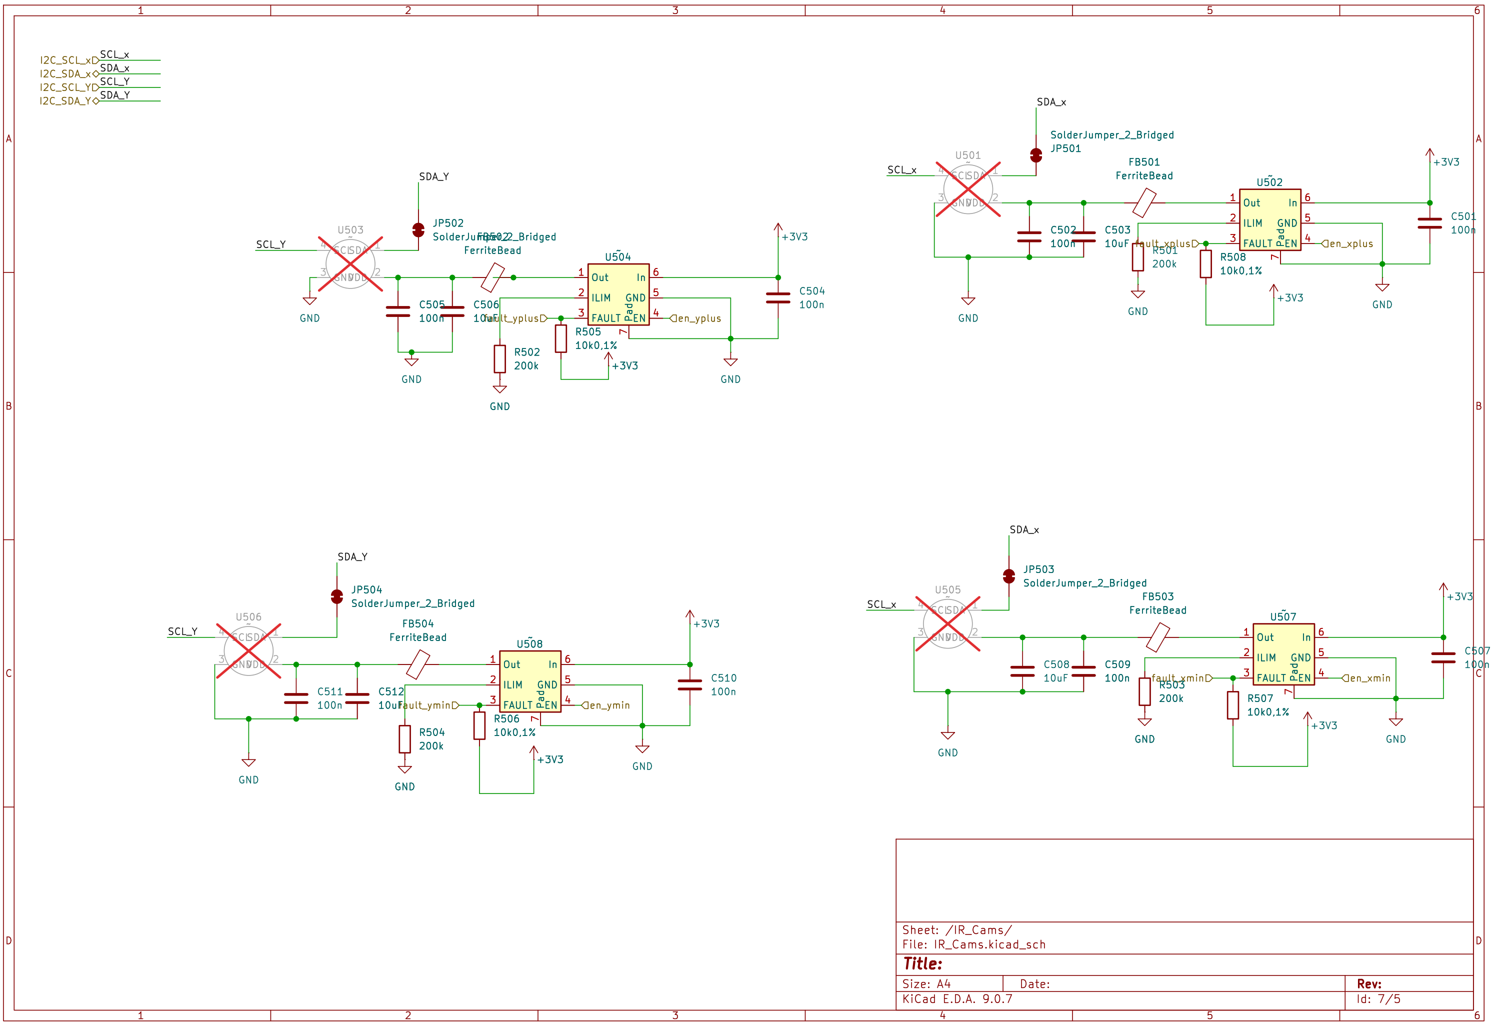
I2C_SCL_x
I2C_SDA_x
I2C_SCL_Y
I2C_SDA_Y
SCL_x
SDA_x
SCL_Y
SDA_Y
U501
U502
U503
U504
U505
U506
U507
U508
JP501
JP502
JP503
JP504
SolderJumper_2_Bridged
FB501
FB502
FB503
FB504
FerriteBead
C501
C502
C503
C504
C505
C506
C507
C508
C509
C510
C511
C512
100n
10uF
R501
R502
R503
R504
200k
R505
R506
R507
R508
10k0,1%
Out
In
ILIM
GND
FAULT
EN
Pad
+3V3
GND
fault_xplus
fault_yplus
fault_xmin
fault_ymin
en_xplus
en_yplus
en_xmin
en_ymin
Sheet: /IR_Cams/
File: IR_Cams.kicad_sch
Title:
Size: A4
Date:
Rev:
KiCad E.D.A. 9.0.7
Id: 7/5

B.2. PCB

PCB 2D screenshot

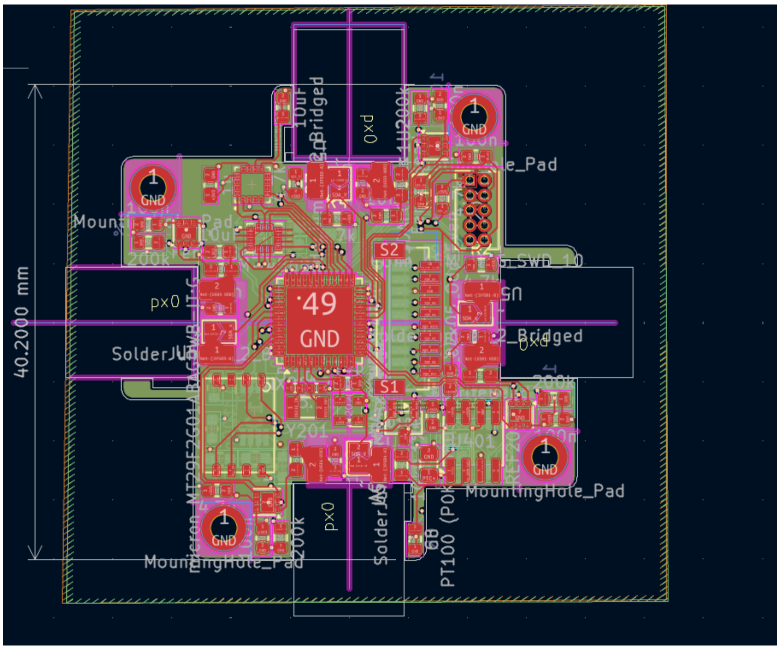

Figure B.1: PCB Layout (Top view).

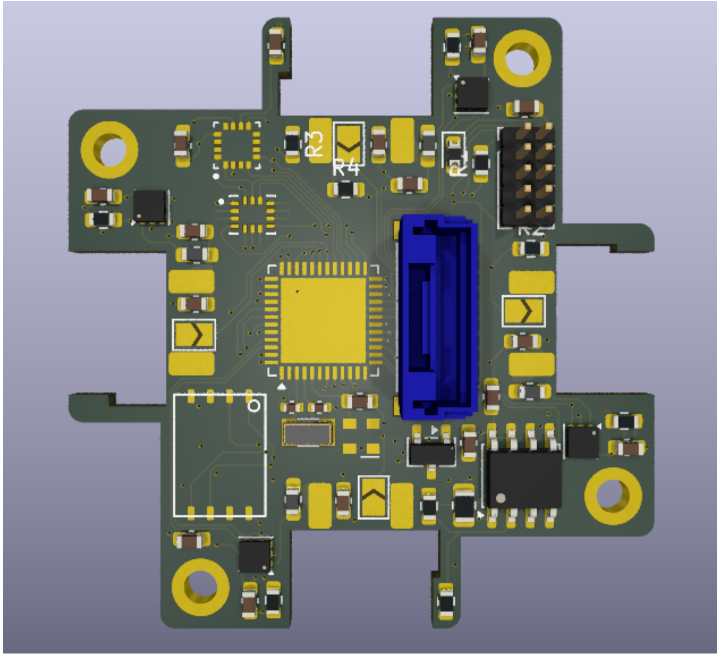

Figure B.2: PCB Layout in 3D view.