Satellite Attitude Determination with Low-Cost Sensors

by

John C. Springmann

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Aerospace Engineering)
in the University of Michigan
2013

Doctoral Committee:

Assistant Professor James W. Cutler, Chair
Professor Dennis S. Bernstein
Professor Ilya Kolmanovsky
Professor Euisik Yoon

This dissertation is dedicated to my parents.

ACKNOWLEDGMENTS

I'd like to thank my advisor, Professor Cutler, for his guidance, mentoring, and inspiration throughout graduate school. I am grateful for the opportunity to work with him and his lab. Thanks also to the rest of my committee, Professor Bernstein, Professor Kolmanovsky, and Professor Yoon, for their time, feedback, and support.

Much of the work in this dissertation was motivated by and demonstrated through application to the design and flight data analysis of the RAX satellites developed here at the University in collaboration with SRI International. This would not have been possible without the awesome accomplishments of the RAX team. I'd like to thank the RAX team members, past and present, and in particular Matt Bennett, Andy Klesh, Alex Sloboda, and Ben Kempke, for their support of this work. It was a pleasure working with the RAX team throughout graduate school. Thanks also to Alex, Matt, and Andy for their contributions to the content of Chapter 2, which is drawn largely from a paper we co-authored on the topic.

This work has been funded by the Department of Defense through a National Defense Science and Engineering Graduate (NDSEG) fellowship, by the University of Michigan Rackham Graduate School through a Rackham Predoctoral Fellowship, by the Department of Aerospace Engineering through a departmental fellowship, and by the Air Force Research Laboratory through an SBIR in collaboration with Creare, Inc. RAX was funded by the National Science Foundation. Thanks to these funders for enabling this work.

And of course, thanks to my family and friends for their encouragement, inspiration, and camaraderie throughout graduate school.

TABLE OF CONTENTS

Dedication . . . . **ii**

Acknowledgments . . . . **iii**

List of Figures . . . . **vi**

List of Tables . . . . **xi**

List of Appendices . . . . **xiii**

List of Abbreviations . . . . **xiv**

Abstract . . . . **xvi**

Chapter

1 Introduction **1**

1.1 Small Satellites 1
1.2 Attitude Determination 4
1.3 Approaches to Attitude Determination 6
1.4 Contributions and Outline 8

2 A Low-Cost Attitude Determination System **11**

2.1 RAX Background 11
2.2 Attitude Determination and Control Requirements 13
2.3 Sensor Selection and Specifications 14
2.4 Integrated Hardware 17
    2.4.1 ADB & Peripheral Sensors 17
    2.4.2 Timing 20
2.5 Photodiode Configurations 21
2.6 Summary 25

3 On-Orbit Magnetometer Calibration **26**

3.1 Introduction 26
3.2 Magnetometer Errors 28
3.3 Review of Existing Methods 29
3.4 Inclusion of Time-Varying Bias 32
3.5 Application to Flight Data 34
    3.5.1 RAX-1 and IGRF Data 35

- 3.5.2 Calibration with Time-Invariant Parameters . . . 36
- 3.5.3 Application of Time-Varying Calibration . . . 37
- 3.5.4 Convergence . . . 41
- 3.5.5 Calibration Accuracy . . . 41
- 3.6 Conclusions . . . 51

4 Directional Sensor Optimization . . . 52

- 4.1 Introduction . . . 52
- 4.2 Use of the Attitude Sphere . . . 53
- 4.3 Formulating the Objective Function for Sun Sensor Optimization . . . 54
  - 4.3.1 Sun Vector Estimation from Multiple Photodiodes . . . 57
  - 4.3.2 Formulation of the Objective Function to Minimize Angular Uncertainty . . . 60
- 4.4 Application . . . 61
  - 4.4.1 Weighting the Attitude Sphere . . . 61
  - 4.4.2 Design Examples . . . 63
- 4.5 Generalization the Optimization Method . . . 69
- 4.6 Conclusions . . . 71

5 On-Orbit Photodiode Calibration . . . 73

- 5.1 Introduction . . . 73
- 5.2 Existing Calibration Techniques . . . 75
- 5.3 Formulation of the Attitude Estimation and Calibration Filter . . . 76
  - 5.3.1 MEKF and USQUE Summary . . . 78
  - 5.3.2 Photodiode Measurement Model . . . 84
  - 5.3.3 Integration into MEKF and USQUE . . . 87
  - 5.3.4 Observability of the Photodiode Normal Directions . . . 90
- 5.4 Application and Results . . . 91
  - 5.4.1 RAX Data . . . 91
  - 5.4.2 Simulated Testing . . . 92
  - 5.4.3 Application to Flight Data . . . 97
  - 5.4.4 Impact on Sun Vector Angular Accuracy and Attitude Determination . . . 106
- 5.5 Discussion . . . 109
- 5.6 Conclusion . . . 111

6 Resulting Attitude Determination Accuracy and the New Approach to Attitude Determination . . . 112

- 6.1 Sensor Calibration and Alignment . . . 112
- 6.2 Attitude Estimation Methods and Results . . . 114
- 6.3 Discussion and the Approach to Low-Cost Attitude Determination . . . 120

7 Conclusions and Areas for Future Work . . . 124

Appendices . . . 126

Bibliography . . . 133

LIST OF FIGURES

1.1 Examples of a star tracker and rate gyroscope. 7
1.2 Examples of a magnetometer and MEMS gyroscope suitable for use within a nanosatellite ADS. 8

2.1 An illustration of the RAX concept of operations. A ground-based incoherent scatter radar station (PFISR) transmits pulses (red cone) into the ionosphere. The signals scatter from the irregularities, which are aligned with Earth's magnetic field, into space. The loci of scatter is represented by the cones. Irregularities are shown at two different altitudes, h1 and h2, and the two cubes represent two different times during a single satellite pass. 12
2.2 The flight-ready RAX satellites. The body-fixed coordinate system is shown and the locations of photodiodes are circled. Photodiodes are located on all six sides of the spacecraft, but only the locations visible in the photos are indicated. 13
2.3 Block diagram and actual RAX Attitude Determination Board (ADB). 18
2.4 Block diagram and actual hardware showing the peripheral clusters on the $\pm x$ and $\pm y$ faces (solar panels) of RAX-1. 19
2.5 Block diagram and a picture of determination hardware located on the +z face of RAX. 19
2.6 Block diagram and a picture of determination hardware located on the -z face of RAX. 20
2.7 An illustration of photodiode coverage for a two-dimensional example. Two orthogonal sensors, labeled 1 and 2, with a $\alpha = 70°$ field of view (FOV) from the normal direction of each sensor are shown. The sun must be in the field of view of both sensors for a two-component vector measurement, limiting the sun sensor coverage to the 50° shaded area. 22
2.8 Photodiode coverage over the body-fixed frame for RAX-1 and RAX-2. Each symbol on the plot is a direction in the body-fixed frame, and the symbols have a uniform angular distribution of approximately 4°. The symbols are shaped and color-coded by the number of sun vector components measured when the sun is in the direction of the symbol. On RAX-1 the three component availability corresponds to exactly three photodiodes being illuminated, but on RAX-2, up to seven photodiodes are illuminated for portions of the body frame. 24
2.9 Two angled photodiodes located near the -z edge of the -x surface (see coordinate system in Figure 2.2(b)). 24

3.1 The convention used to define the sensor non-orthogonality and the sensor frame. The axes of the right-handed sensor frame are denoted by x, y, and z, and the directions of the actual sensing elements of the sensor are $\tilde{x}$, $\tilde{y}$, and $\tilde{z}$. The x-axis of the sensor frame is aligned with $\tilde{x}$, the y-axis is in the $\tilde{x}$-$\tilde{y}$ plane, and the z-axis completes the right-handed frame. The non-orthogonality angles are ρ, measured from the y-axis; ϕ, measured from the x-z plane; and λ, measured from the y-z plane. 31
3.2 The inside of the lower portion of RAX-1 is shown. The arrow points to the PNI MicroMag3 magnetometer. The circuit board located directly above the magnetometer is the satellite power regulation system. The fully integrated RAX-1 satellite is shown in Figure 2.2(a) 36
3.3 Data from the RAX-1 PNI MicroMag3 magnetometer. The x-axis of each plot shows time elapsed since the start of the data set, 01-Dec-2010 08:30:46 UTC. 37
3.4 Difference between the magnitude of the corrected measurements using time-invariant calibration, and the expected field magnitude, as shown in Figure 3.3(b). 38
3.5 Results of the calibration to estimate both constant errors and time-varying magnetometer bias. 40
3.6 The normalized vector components of the corrected measurements plotted in three dimensions on top of a unit sphere in the body frame for each data set, which shows the coverage of the *attitude sphere*. 45
3.7 Mean and standard deviation of the error in scale factor estimates from the 1000 simulated trials. The true parameters are shown in Table 3.3, and the mean errors and standard deviations as a percentage of the true parameters are given in Table 3.6. 47
3.8 Mean and standard deviation of the error in constant bias estimates from the 1000 simulated trials. The true parameters are shown in Table 3.3, and the mean errors and standard deviations as a percentage of the true parameters are given in Table 3.6. 47
3.9 Mean and standard deviation of the error in non-orthogonality angle estimates from the 1000 simulated trials. The true parameters are shown in Table 3.3, and the mean errors and standard deviations as a percentage of the true parameters are given in Table 3.6. 48
3.10 Mean and standard deviation of the error in the current coefficients estimates for the magnetometer $\tilde{x}$-axis from the 1000 simulated trials. The mean error and standard deviations as a percentage of the true parameters are given in Table 3.6. 48
3.11 Mean and standard deviation of the error in the current coefficients estimates for the magnetometer $\tilde{y}$-axis from the 1000 simulated trials. The mean error and standard deviations as a percentage of the true parameters are given in Table 3.6. 49
3.12 Mean and standard deviation of the error in the current coefficients estimates for the magnetometer $\tilde{z}$-axis from the 1000 simulated trials. The mean error and standard deviations as a percentage of the true parameters are given in Table 3.6. 49

3.13 Measurements of current from the December 1 data set that were used in the magnetometer calibration. Currents 1-4 are solar panels, and the fifth is current draw from the battery. 50

4.1 Unit spheres with two different methods of distributing points shown with their 2D projections. 55
4.2 The first two steps of the icosahedron-based approach to sphere pixelization are shown. 56
4.3 Number of pixels N and corresponding angular resolution versus the resolution r. 56
4.4 An illustration of two photodiodes in a single plane. The solid rectangles represent the two photodiodes, and their normal directions are shown by the unit vectors $\hat{\boldsymbol{n}}_i, i \in \{1, 2\}$. The angles θ_i define the sun vector direction relative to the photodiode normal directions, as in Eq. (4.2), and the sun vector $\hat{\boldsymbol{s}}$ corresponds to the intersection of the two planes. The planes are represented by the dotted lines, are perpendicular to $\hat{\boldsymbol{n}}_i$, and their location along $\hat{\boldsymbol{n}}_i$ is given by $I_i/I_{0,i}$ 58
4.5 The directions on the attitude sphere color-coded by the number of occurrences of the sun vector for the nadir-pointing spacecraft example of Section 4.4.4.4.1. Directions with zero occurrences are not shown. 62
4.6 The photodiode configuration for Example 1 reduced to a single design variable. There are two photodiodes per face of the cube and the directions normal to the photosensitive planes are labeled with $\hat{n}$. The photodiodes on the $+x$, $+y$, and $+z$ surfaces are in the x-y, y-z, and x-z planes, respectively, and each photodiode is tilted by angle γ from the surface. γ is the single design variable. Photodiodes on the $-x$, $-y$, and $-z$ faces are not shown but have an identical configurations as the opposite faces of the cube. 64
4.7 The objective value and sun vector angular uncertainty as a function of photodiode mounting angle relative to the spacecraft surface in the geometry of Figure 4.6. The minimum objective values are achieved with mounting angles of 57° and 47° for the identical and nadir weighting schemes, respectively. Calculation of the sun vector angular uncertainty is discussed in the Appendix. 65
4.8 Photodiode orientations and the resulting sun vector angular uncertainties over the attitude sphere for Example 1. In Figures 4.8(a) and 4.8(c), the azimuth and elevation angles are those of the direction normal to the photosensitive plane in the body-fixed frame. 66
4.9 The photodiode configuration for Example 2 reduced to a single design variable, which is the mounting angle of the photodiodes on the x/y surfaces. Single photodiodes are placed on the +/- z surfaces of a cubical spacecraft and orientated normal to the surface (not shown). The three photodiodes on each of the remaining surfaces are shown with their normal directions denoted by $\hat{n}$. They are oriented in increments of 120° in the plane of the surfaces and are tilted from the surfaces (out of the page) by the single design angle. 67

4.10 The objective value and sun vector angular uncertainty as a function of photodiode mounting angle relative to the spacecraft surface in the geometry of Figure 4.9. The minimum objective values are achieved at the same mounting angle for both weighting schemes. Unlike Example 1, the angular uncertainty of the estimated sun vector varies significantly over the range of mounting angles that provide 4π steradian sun sensing coverage. 68
4.11 Photodiode orientations and the resulting sun vector angular uncertainties over the attitude sphere for Example 2. In Figure 4.11(a), the azimuth and elevation angles are those of the direction normal to the photosensitive plane in the body-fixed frame. 70

5.1 The mean and standard deviation of Earth reflectivity measurements from the year 2005. 86
5.2 Two different azimuth/elevation systems used in the photodiode calibration. . . 91
5.3 The average error in each state estimate over the 50 trials with the initial state estimates of Simulation 1. In each plot, the UF estimates are shown in solid blue lines and the EKF estimates are shown in dotted red lines. 95
5.4 The average error in each state estimate over the 50 trials with the initial state estimates of Simulation 2. In each plot, the UF estimates are shown in solid blue lines and the EKF estimates are shown in dotted red lines, but the estimates from the two filters are not distinguishable; the difference between the EKF and UF errors is nearly zero. 96
5.5 Raw photodiode measurements (voltage) from RAX-2 data set one. For this data set, the times of the data used for calibration are 56-118 minutes. 97
5.6 Scale factor estimates and ±3-σ bounds for photodiodes from RAX-1 data set one. 99
5.7 Azimuth estimates and ±3-σ bounds for photodiodes from RAX-1 data set one. 100
5.8 Elevation estimates and ±3-σ bounds for photodiodes from RAX-1 data set one. 101
5.9 Scale factor estimates and ±3-σ bounds for photodiodes from RAX-2 data set one. 102
5.10 Azimuth estimates and ±3-σ bounds for photodiodes from RAX-2 data set one. 103
5.11 Elevation estimates and ±3-σ bounds for photodiodes from RAX-2 data set one. 104
5.12 Histograms of the magnitude of the measured sun vector from RAX-1 data set one using the initial (left) and calibrated (right) sensor parameters. 106
5.13 Histograms of the magnitude of the measured sun vector from RAX-2 data set two using the initial (left) and calibrated (right) sensor parameters. 106
5.14 Magnitude of the measured sun vector over time for the two RAX-2 data sets using both the uncompensated photodiode measurements (red circles) and the albedo-compensated measurements (blue dots). These magnitudes were calculated with the estimated calibration parameters. The albedo-compensated series of data is the same data as the calibrated histograms of Figures 5.12-5.13. 107
5.15 Histograms of the angular difference between the measured sun vector using the initial and calibrated photodiode parameters from the RAX-2 data sets. The mean from data set one is 9.1° and the mean from data set two is 10.2°. . . 108

5.16 Attitude accuracy (1-σ) achieved in applying the attitude estimation/calibration filter to data sets one and two. 109
5.17 Attitude accuracy of a third RAX-2 data set, which was collected December 9, 2011 16:00:00 UTC. 109

6.1 Attitude estimation accuracy of the simultaneous photodiode calibration and attitude estimation filter using RAX-1 measurements collected Dec. 1, 2010 08:30:46 UT. 117
6.2 Attitude estimation accuracy of the simultaneous photodiode calibration and attitude estimation filter using RAX-2 measurements collected Nov. 4, 2011 18:29:45 UT. 117
6.3 Measured angular velocities during the data sets shown in Figures 6.1 and 6.2. 118
6.4 Attitude determination accuracy of RAX-2 data collected Dec. 9, 2011 16:00:00 UT using the photodiode calibration filter. 120
6.5 Attitude determination accuracy of RAX-2 data collected Dec. 9, 2011 16:00:00 UT using the 6-state multiplicative extended Kalman filter (MEKF) with an Earth albedo model to compensate the photodiode measurements. 121
6.6 Attitude determination accuracy of RAX-2 data collected Dec. 9, 2011 16:00:00 UT using the 6-state MEKF without an albedo model. Albedo is treated as noise in the filter. 121

A.1 Histogram of $\tilde{\mathbf{y}}$ calculated from 10,000 random sensor measurements with Gaussian distribution and the calibration parameters for the December 1 data set, shown in Table 3.3. The standard deviations used are $\sigma_x = \sigma_y = \sigma_z = 128$ nT, and $\sigma_1 = \cdots = \sigma_5 = 5$ mA. 128

B.1 Histograms of the simulated measurements for the five sun sensors for which the simulated sun vector (direction on the attitude sphere) is in the field of view. The peak current output of the photodiodes is 1 and the standard deviation is 0.02. 132
B.2 Histogram of the angular error in the sun vector estimate. The mean error is 1.3°. 132

LIST OF TABLES

1.1 Satellite categories by mass and approximate cost. 2
1.2 Typical performance, mass, and power of ADS sensors [1]. 6

2.1 Characteristics of the RAX magnetometers. 16
2.2 Characteristics of the rate gyro, part of the Analog Devices ADIS16405 IMU. . 16
2.3 Characteristics of the Osram SFH 2430 photodiodes. 16
2.4 Characteristics of the temperature sensors. The maximum error of the ADIS16405 is not given because it is not reported on the manufacture-provided specification sheet and the error was not experimentally characterized. 17
2.5 The intended azimuth and elevation angles of the photodiode normal directions on RAX-1. The side panel is the surface to which the sensors are mounted (coordinate system shown in Figure 2.2(a)). 22
2.6 The intended azimuth and elevation angles of the photodiode normal directions on RAX-2. The side panel is the surface to which the sensors are mounted (coordinate system shown in Figure 2.2(b)). 23

3.1 The time difference between the epochs of the TLEs used and the time of the magnetometer data sets. Three data sets are discussed: December 1, December 15, and December 30, 2010. 35
3.2 The range of uniformly distributed initial conditions used in the Monte Carlo simulations (column two) and the final estimate of the calibration parameters (column three). There are 15 total mapping coefficients, $s_{i,j}$. The range of the final coefficients is given rather than each individual parameter for simplicity. . 42
3.3 Estimates of the calibration parameters for each data set. Nine of the 24 parameters are shown; the remaining are omitted from the table for simplicity. . . 44
3.4 The root mean squared error (RMSE) (nT) of the corrected measurements after calibration using the parameter estimates from each data set. Measurements from the data sets listed in the left column are corrected using parameters from the data sets listed across the top. 45
3.5 The RMSE (nT) of the corrected measurements after time-invariant calibration using the parameter estimates from each data set. This is analagous to Table 3.4, but the calibration method is the existing time-invariant method described in Section 3.3. Measurements from the data sets listed in the left column are corrected using parameters from the data sets listed across the top. 46

3.6 Mean and standard deviation of the error in the estimates of scaling, bias, and non-orthogonality resulting from the 1000 simulated trials. The mean and standard deviations are given as a percentage of the true parameter. 46

5.1 Sensor error parameters and initial angular velocity for the simulated data. . . . 93
5.2 Initial state estimate accuracy for the simulated testing. 94
5.3 Scale factor estimates from the RAX-1 data sets. 107

6.1 Current prices of the commercial off-the-shelf (COTS) sensors utilized in the Radio Aurora Explorer (RAX) attitude determination system (ADS). 122

A.1 Approximate lower bound of the 3-σ values of the calibration parameter estimates for each data set. Nine of the 24 parameters are shown; the remaining are omitted from the table for simplicity. 129

LIST OF APPENDICES

A Uncertainty of the Magnetometer Calibration Parameters 126

B Uncertainty of the Sun Vector Estimate 130

LIST OF ABBREVIATIONS

ADB Attitude Determination Board

ADC analog-to-digital converter

ADCS attitude determination and control system

ADS attitude determination system

COTS commercial off-the-shelf

EKF extended Kalman filter

EPS electrical power system

FAI field-aligned irregularities

FCPU flight central processing unit

FOV field of view

IGRF International Geomagnetic Reference Field

IMU inertial measurement unit

LEO low-Earth orbit

MEKF multiplicative extended Kalman filter

MEMS microelectromechanical systems

MXL Michigan Exploration Laboratory

NSF U.S. National Science Foundation

PTB position and time board

RAX Radio Aurora Explorer

RMSE root mean squared error

STK Satellite Toolkit

TLE	two-line element
TOMS	Total Ozone Mapping Spectrometer
UF	unscented filter
UHF	ultra-high frequency
USQUE	unscented quaternion estimator

ABSTRACT

Satellite Attitude Determination with Low-Cost Sensors

by

John C. Springmann

Chair: James W. Cutler

This dissertation contributes design and data processing techniques to maximize the accuracy of low-cost attitude determination systems while removing pre-flight calibration requirements. This enables rapid development of small spacecraft to perform increasingly complex missions. The focus of this work is magnetometers and sun sensors, which are the two most common types of attitude sensors.

Magnetometer measurements are degraded by the magnetic fields of nearby electronics, which traditionally limit their utility on satellites unless a boom is used to provide physical separation between the magnetometer and the satellite. This dissertation presents an on-orbit, attitude-independent method for magnetometer calibration that mitigates the effect of nearby electronics. With this method, magnetometers can be placed anywhere within the spacecraft, and as demonstrated through application to flight data, the accuracy of the integrated magnetometer is reduced to nearly that of the stand-alone magnetometer.

Photodiodes are light sensors that can be used for sun sensing. An individual photodiode provides a measurement of a single sun vector component, and since orthogonal photodiodes do not provide sufficient coverage due to photodiode field-of-view limitations, there is a tradeoff between photodiode orientation and sun sensing angular accuracy. This dissertation presents a design method to optimize the photodiode configuration for sun sensing, which is also generally applicable to directional sensors. Additionally, an on-orbit calibration method is developed to estimate the photodiode scale factors and orientation, which are critical for accurate sun sensing.

Combined, these methods allow a magnetometer to be placed anywhere within a spacecraft and provide an optimal design technique for photodiode placement. On-orbit calibration methods are formulated for both types of sensors that correct the sensor errors on-orbit without requiring pre-flight calibration. The calibration methods are demonstrated by application to on-orbit data, and attitude determination accuracies of 0.5° 1-σ are achieved with commercial-off-the-shelf magnetometers, photodiodes, and a MEMS rate gyroscope, which to the author's knowledge, is the best accuracy reported in the literature for this class of sensors.

CHAPTER 1

Introduction

The *attitude* of a spacecraft is its orientation in space, and the term *attitude determination* refers to the entire process of estimating attitude, which includes the use of directional sensors in conjunction with state estimation techniques. Attitude determination is a critical aspect of most satellite missions, and a wide variety of sensors and estimation algorithms are readily available for use in an attitude determination system (ADS) [2]. The traditional approach to attitude determination utilizes relatively large sensors in conjunction with high-tolerance pre-flight alignment and calibration procedures. In this dissertation, we develop new design and sensor calibration techniques that can be used to maximize attitude determination accuracy with a given set of sensors while at the same time reducing the pre-flight integration and testing requirements. These techniques result in improved attitude determination accuracy and are particularly enabling for small spacecraft. Compared to traditional large satellites, small satellites can be advantageous because of their lower development time and costs, but they face significant power, mass, and volume constraints, limiting their use of high-performance attitude determination sensors. The sensors used in this work are magnetometers and sun sensors, both extremely common types of attitude sensors suitable for spacecraft of any size, but the approach to attitude determination resulting from the contributions of this dissertation is generally applicable to all types of attitude sensors on any vehicle. This chapter provides an overview and summarizes the contributions of this dissertation.

1.1 Small Satellites

Small Satellites are a class of spacecraft that generally have lower costs and faster development times than traditional large satellites. While the exact interpretation of *small* varies in the community, a generally-accepted definition is that small satellites are those with a mass of under 500 kg. Within the category of small satellites, sizes are further classified into sub-

Table 1.1: Satellite categories by mass and approximate cost.

Category	**Mass, kg**	**Cost, USD**
Large satellite	> 1000	0.1-2 B
Medium satellite	500-1000	50-100 M
Minisatellite	100-500	10-50 M
Microsatellite	10-100	2-10 M
Nanosatellite	1-10	0.2-2 M
Picosatellite	0.1-1	20-200 K
Femtosatellite	< 0.1	0.1-20 K

classes such as microsatellite and nanosatellite. The masses and approximate development costs of the various size satellites is shown in Table 1.1 [3].

Since the launch of the first satellite in 1957, increasing mission requirements have driven up satellite mass – to over 6000 kg for some systems – as well as cost, complexity, and management overhead [3]. At the same time, the growing small satellite industry has developed increasingly-capable and lower-cost missions by simplifying requirements and utilizing the latest available technology. The recent growth in the small satellite industry has been catalyzed by CubeSats, an effort started in 1999 [4]. A CubeSat, short for *cube satellite*, is a standardized cubic nanosatellite form factor that is ten centimeters on each side with a mass of approximately one kilogram. Larger CubeSats are multiples of a this single CubeSat unit (U). For example, a 3U CubeSat, also referred to as a triple CubeSat, is a common configuration that is approximately $10 \times 10 \times 30$ cm^3 and 3.3 kg. The standardization facilitates relatively frequent and low-cost access to space through launch as secondary payloads. While CubeSat development has traditionally been rooted in academia, the rapid growth in CubeSat development has spread beyond academia, and the satellites are currently being developed by universities, private industry, and both civil and defense-related government organizations. 80 individual CubeSats are scheduled to launch in 2013 [5].

While small satellites will likely never completely replace their larger counterparts due limitations on aspects such as power generation and aperture size, small satellites are replacing large satellites for certain types of missions and complementing large satellites in other areas [6]. For example, small satellites have been used for space weather monitoring [7], Earth observation [8], remote sensing [9], and even astrophysics [10]. The U.S. National Science Foundation (NSF) has funded a series of CubeSats at approximately \$1 M per mission to study space weather[1], and the first of these missions, the Radio Aurora

---

[1] The NSF missions launched to date include RAX (Radio Aurora Explorer), DICE (Dynamic Ionosphere

Explorer (RAX) (which will be discussed further in Chapter 2), has provided unprecedented measurements of electron density irregularities in Earth's Ionosphere [11, 12]. This has demonstrated the ability of nanosatellites to perform focused science missions at a cost orders of magnitude lower than their larger counterparts. The success of nanosatellites in low-Earth orbit (LEO) has gained the attention of the space science community, and small satellite missions in both LEO and beyond are currently being proposed in the areas of astrophysics, heliophysics, and planetary sciences [13]. The Jet Propulsion Laboratory, in collaboration with university partners, is currently developing INSPIRE, the first interplanetary CubeSat mission [14]. This technical demonstration mission is a precursor to future scientific nanosatellite missions beyond LEO.

In addition to their utility in single-point measurements and observations, the low cost of small satellites facilitates the deployment of constellations that can be used for global in-situ observations [15], and examples of small satellite constellations that are in development include QB50[2], a constellation of 50 CubeSats for lower thermosphere and re-entry research, and CYGNSS [16], a constellation of eight microsatellites being developed to study the formation of hurricanes. Clusters of small satellites can also be used to disperse capabilities required for a single mission among multiple platforms, referred to as disaggregation. Recently, the U.S. Air Force announced that they will pursue disaggregation rather than single, monolithic platforms in future space acquisitions [17].

Small satellites are also frequently used for technical demonstrations and education. Given their relatively low costs, CubeSats have been used to test payloads in the space environment, validating their performance for use on future large-scale missions (for example, [18]). Many CubeSats, especially those developed in academia, utilize new design methods, algorithms, and components that would be considered too risky for traditional, more expensive and necessarily more risk-averse missions. But once these methods are demonstrated on small satellites, they can be used on larger systems, making CubeSats ideal platforms for innovation. This dissertation is an example of such innovation: the methods presented are motivated largely by small satellite development, but the methods are applicable to attitude determination on any spacecraft as well as water-, ground-, and air-based vehicles. Additionally, small satellite development at Universities provides students the opportunity to design, build, and operate satellites, resulting in an unprecedented educational impact.

Small satellites can be advantageous due to their low costs and fast development times[3],

---

CubeSat Experiment), CSSWE (Colorado Student Space Weather Experiment), and CINEMA (Cubesat for Ion, Neutral, Electron, Magnetic fields). Additional missions are scheduled to launch in late 2013 and 2014.

[2] `https://www.qb50.eu/`, accessed September 2013.

[3] The initial time between the start of the RAX development and its launch was one year. The launch

but they present significant technical challenges. For example, the decreased surface area results in a relatively low available power[4], and the use of deployables to increase solar panel area comes with added development times and costs. Power constraints, combined with mass and volume constraints, results in the individual subsystems being tightly coupled. For example, power generation and storage is a function of battery sizing, solar panel sizing, and the attitude determination and control system (ADCS), tightly coupling subsystems that have traditionally been treated separately in the design phase [19, 20]. Additionally, the close physical proximity of the various subsystems leads to the increased potential of electromagnetic interference onboard the satellite, such as the interaction of magnetometers used for attitude determination with nearby electronics, which addressed in Chapter 3. The impact on attitude determination specifically is discussed further in Section 1.3

The methods presented in this dissertation were developed in parallel with small satellite design, development, and operations in the Michigan Exploration Laboratory (MXL) at the University of Michigan. This research began in Fall 2009, and MXL's RAX-1 launched November 2010, RAX-2 and MCubed-1 launched October 2011, and three more satellites are scheduled to launch in late 2013 and 2014. These are all CubeSats developed for either scientific or technology demonstration missions. Some of the methods developed in this dissertation will be demonstrated through application to flight data from RAX-1 and RAX-2, and some of the methods of this dissertation [21, 22] are also already in use by satellite developers at the University of Colorado, the University of Texas, and Boston University.

1.2 Attitude Determination

Spacecraft attitude is quantified by the 3×3 proper orthogonal rotation matrix A that defines the orientation of the satellite body-fixed frame relative to a reference frame [23]. This is shown by Eq. (1.1), where $\boldsymbol{r}|_B$ is a 3×1 vector resolved in the body-fixed frame B, and $\boldsymbol{r}|_R$ is the same vector resolved in the reference frame R.

$$\boldsymbol{r}|_B = A\, \boldsymbol{r}|_R \tag{1.1}$$

The traditional types of sensors that can be used for attitude determination are sun sensors, star trackers, horizon sensors, magnetometers, and rate gyroscopes [1, 2]. Each type of sensor is described briefly below, and performance, mass, and power requirements are summarized in Table 1.2.

---

schedule slipped and the development ultimately ended up taking approximately 1.5 years.

[4] A 3U CubeSat with body-mounted solar panels has an average available power of less than 10 W.

Sun Sensors provide a measurement of one or two angles between the sensor boresight direction and the sun, providing information about the line-of-sight vector to the sun in the satellite body-fixed frame. Various types of sun sensors exist, with accuracies ranging from better than $0.01°$ to a few degrees.

Star trackers, also referred to as star cameras, are the most accurate type of attitude sensor. Most modern star trackers work much like digital cameras in that they take a picture of the star field. A full attitude solution can be derived by comparing the star field image to a star catalog stored onboard. Star tracker accuracies range from one arc-second to one arc-minute ($0.0003°$ to $0.01°$), depending on the quality of the sensor.

Horizon sensors, also known as Earth sensors, are infrared sensors that detect the contrast between cold of space and the heat of Earth's atmosphere. Fixed-head horizon sensors can be used on spinning spacecraft, or scanning horizon sensors that utilize a rotating mirror or lens can be used on non-spinning spacecraft. Through detection of the edges of the earth, the sensors measure Earth phase and chord angles to derive a measurement of the nadir direction in the body-fixed frame. Accuracies between $0.1°$ and $0.25°$ are typical of horizon sensors.

Magnetometers measure the direction and magnitude of Earth's magnetic field. Given the decreasing strength of the geomagnetic field as distance from the Earth increases (the field decreases proportionally to $1/r^3$), they are generally used only at altitudes below 1000 km. Attitude-grade magnetometer accuracies typically range from $0.5°$ to $3°$.

Rate gyroscopes, also referred to simply as gyros, do not provide a measure of attitude, but are commonly used in spacecraft ADS as they provide a measurement of angular velocity. They are combined with other directional or attitude measurements in recursive attitude estimators. Fusing the gyro with attitude measurements provides more accurate attitude estimates than the attitude measurements alone, with the degree of improvement dependent on a number of factors such as the sensor sampling frequency, spacecraft dynamics, and gyro noise characteristics.

Of these sensors, sun sensors, horizon sensors, and magnetometers provide directional measurements – measurements of $\boldsymbol{r}|_B$ in Eq. (1.1) – and star trackers generally provide a direct measurement of attitude. For the directional measurements, a model is used to obtain the corresponding vector in the reference frame, $\boldsymbol{r}|_R$. With attitude and/or directional measurements in-hand, many methods exist for attitude estimation [24], including point-by-point methods, which estimate attitude given at least two non-parallel measured vectors at a given time instant, or recursive methods, which utilize a time series of measurements to estimate attitude. Of the available methods, the extended Kalman filter (EKF) is the most commonly used method and has been referred to as "the workhorse of real-time spacecraft

Table 1.2: Typical performance, mass, and power of ADS sensors [1].

Sensor	**Typical Performance Range**	**Mass (kg)**	**Power (W)**
Sun Sensor	Accuracy of 0.005° to 3°	0.1 to 2	0 to 3
Star Tracker	Accuracy of 0.0003° to 0.01°	2 to 5	5 to 20
Horizon Sensor	Accuracy:		
Fixed head	$<$ 0.1° to 0.25°	0.5 to 3.5	0.3 to 5
Scanning	0.05° to 1°	1 to 4	5 to 10
Magnetometer	Accuracy of 0.5° to 3°	0.3 to 1.2	$<$1
Rate gyroscope	Drift rate of 0.003 deg/hr to 1 deg/hr	$<$ 0.1 to 15	$<$ 1 to 200

attitude estimation [24].” EKFs fuse attitude sensor measurements with either rate gyroscope measurements or a dynamic model, with the former option much more common [25]. Kalman filtering for spacecraft attitude estimation will be discussed further in Section 5.3.

1.3 Approaches to Attitude Determination

Of the available attitude determination sensors, utilizing a combination of star trackers and rate gyroscopes is the most common approach to meeting the most stringent attitude accuracy requirements on large spacecraft [1]. Examples of such sensors are shown in Figures 1.1(a)[5] and 1.1(b)[6]. This conventional approach is not possible on nanosatellites given the power and physical size requirements of these high fidelity sensors. As seen in Table 1.2, the power and mass requirements of the various attitude determination sensors can be significant, precluding the use of some of these sensors on small satellites.

The types of attitude sensors that are commonly flown on nanosatellites are magnetometers and sun sensors. These sensors have power requirements on the order of milliwatts, mass requirements on the order of hundreds of grams, volume requirements on the order of cubic centimeters or less, and have hardware costs on the order of tens of dollars for digital magnetometers and on the order of dollars for the most basic type of sun sensors. As for rate gyroscopes, miniature gyros based on microelectromechanical systems (MEMS) technology are the only type of gyro currently suitable for nanosatellites. Examples of a magnetometer and rate gyroscope that could be used in a nanosatellite ADS

---

[5] `http://www.ballaerospace.com/file/media/D0399_CT-633.pdf`, accessed September 2013.

[6] `http://www2.l-3com.com/spacenav/space_and_nav/space_products/pdfs/SellSheet_RGA_20_Apr11.pdf`, accessed September 2013.

(a) Ball Aerospace CT-633 star tracker, which has a mass of 2.5 kg without the baffle, 13.4 cm diameter, requires 8 W of power, and provides an attitude accuracy of better than 0.01°.

(b) L-3 Space and Navigation RGA-20 High Performance Gyro Assembly, which has a mass of 5.6 kg, $15 \times 24 \times 20$ cm^3 physical size, and requires 25 W of power.

Figure 1.1: Examples of a star tracker and rate gyroscope.

are shown in Figures 1.2(a)[7] and 1.2(b)[8]. Compared to the high performance gyro shown in Figure 1.1(b), this MEMS gyro has noise characteristics and bias stability that are factors of 500 and 2520 worse, respectively. A magnetometer- and MEMS gyro- based ADS on a nanosatellite could provide attitude accuracies of approximately two degrees, whereas the hardware of Figure 1.1 would provide an attitude accuracy of better than 0.0011°.

Star trackers and Earth horizon sensors have also been miniaturized for use on small spacecraft, but to date, they have been utilized less frequently than magnetometers and sun sensors predominately due to their higher cost, higher power and mass requirements, and more recent availability compared to magnetometers and sun sensors. A currently-available nanosatellite star tracker with an attitude accuracy of better than 0.02° can be purchased for $100,000[9], and additional star trackers at lower costs as well as horizon sensors are becoming available as the demand for more accurate nanosatellite ADCS continues to grow.

The term *low cost attitude determination* relates to both the sensor hardware costs as well as the sensor calibration and integration costs. The traditional approach to satellite development consists of thorough calibration of the individual sensors before integration into the satellite, followed by high-tolerance alignment procedures during spacecraft integration [26]. An alternate approach is to perform sensor calibration once the spacecraft is

---

[7] `https://www.sparkfun.com/datasheets/Sensors/MicroMag3%20Data%20Sheet.pdf`, accessed September 2013.

[8] `http://www.analog.com/static/imported-files/data_sheets/ADIS16400_16405.pdf`, accessed September 2013.

[9] `http://www.sinclairinterplanetary.com/`, accessed September 2013.

(a) PNI MicroMag3 three-axis magnetometer. The physical dimensions are $2.54 \times 2.54 \times 1.9$ cm^3 and the sensor requires 1.5 mW.

(b) Analog Devices ADIS13405 inertial measurement unit, which includes a three-axis rate gyroscope, three-axis accelerometer, and a three-axis magnetometer. The physical dimensions are $2.3 \times 2.3 \times 2.3$ cm^3 and the sensor requires 225 mW.

Figure 1.2: Examples of a magnetometer and MEMS gyroscope suitable for use within a nanosatellite ADS.

on-orbit. This reduces satellite development time and cost through reduction in pre-flight testing and integration requirements. In this dissertation, we develop methods for on-orbit sensor calibration of low-cost sensors that enables unprecedented attitude determination accuracies with these sensors while at the same time minimizing the pre-flight testing and integration requirements.

1.4 Contributions and Outline

There are three main contributions of this dissertation. The contributions are described below, and further details on how these methods improve on the current state of the art are given in their corresponding chapters.

1. A new magnetometer calibration technique has been developed to compensate for the impact of surrounding electronic components on magnetometers embedded within a satellite. Current-carrying wires create magnetic fields, which subsequently bias magnetometer measurements. The traditional approach to mitigate this is through careful design and manufacturing practices to reduce the magnetic disturbance of nearby electronics, or by physically separating the magnetometer from the rest of the satellite by placing it on the end of a boom. The new calibration technique compensates for the measurement errors due to both the surrounding components as well as those inherent to the sensor (such as scale factors), enabling the magnetometer to be placed anywhere within a satellite. This calibration is an attitude-independent method that requires only on-orbit measurements, mitigating the need for pre-flight

calibration.

2. A method is developed to optimize the orientation of directional sensors, and this is applied to optimize the configuration of photodiodes for sun sensing. Photodiodes, also referred to as coarse sun sensors, generate current as a function of incoming light and are the most basic type of sun sensor. They can be used to provide a measurement of a single component of the line-of-sight vector to the sun. Three non-parallel photodiodes are required for a three-component sun vector measurement, but FOV limitations preclude the use of an orthogonal three-sensor configuration. The angular accuracy of the sun vector measurement is dependent on the orientation of each photodiode, and the new formulation for directional sensor optimization provides a method to optimize the photodiode configuration subject to the resulting angular accuracy. The optimization method is generally applicable to other sensors or instruments that have either a limited FOV or performance that varies over the FOV, such as body-mounted solar panels or antennas.

3. An on-orbit photodiode calibration method is developed. This method estimates the sensor parameters that are required for photodiode-based attitude determination – the photodiode orientation on the spacecraft and the sensor scale factor – and does not require any pre-flight calibration or alignment. The calibration is accomplished through simultaneous estimation of spacecraft attitude and the sensor calibration parameters, and formulations that utilize both an EKF and unscented filter (UF) are developed and compared.

Combined, these methods allow a magnetometer to be placed anywhere within a spacecraft, provide an optimal design technique for photodiode orientation, and enable in-flight sensor calibration that replaces pre-flight calibration and alignment. This improves the performance of the ADS while lowering the development time and cost. The sensor calibration methods are demonstrated to flight data from the RAX satellites, and 1-σ attitude accuracies of 0.5° are demonstrated. To the author's knowledge, this is the highest accuracy attitude determination that has been reported with this class of sensors.

The remainder of this dissertation is organized as follows. In Chapter 2, we describe the design of the ADS developed for the RAX satellites. The ADS design is not a unique contribution of this dissertation, but rather, it is provided as an example of a low-cost attitude determination system that is typical of nanosatellites. Additionally, the calibration methods of this dissertation are demonstrated through application to flight data from the RAX ADS, so Chapter 2 provides information on the subsystem that will be used throughout the dissertation. The main contributions of this dissertation are given in the subsequent individual chapters: the on-orbit magnetometer calibration is presented in Chapter 3, the directional sensor optimization is given in Chapter 4, and the on-orbit photodiode calibration is developed in Chapter 5. The results of RAX attitude determination after application of these methods as well as the resulting approach to attitude determination is then summarized in Chapter 6. Finally, conclusions and future work are discussed in Chapter 7.

CHAPTER 2

A Low-Cost Attitude Determination System

In this chapter, we describe the design of the ADS of the RAX satellites, RAX-1 and RAX-2. This is an example of a low-cost attitude determination system, and flight data from these satellites will be used in Chapters 3, 5, and 6. Although this system was developed for RAX, it could be used for attitude determination in any LEO satellite. This chapter is organized as follows. We first give an overview of the RAX mission in Section 2.1 and discuss its attitude determination and control requirements in Section 2.2. The sensor selection and sensor integration into the RAX ADS are the topics of Sections 2.3 and 2.4. The sun sensor configurations utilized on RAX-1 and RAX-2 are then discussed in Section 2.5, and a summary is given in Section 2.6.

2.1 RAX Background

The Radio Aurora Explorer (RAX) is a CubeSat mission developed to study space weather in Earth's ionosphere [27, 28]. It is the first in a series of CubeSats funded by the U.S. National Science Foundation (NSF) to study space weather. The mission was co-developed by the University of Michigan and SRI International.

The primary objective of the RAX mission is to study the formation of magnetic field-aligned irregularities (FAI) of electron density in the polar lower ionosphere (80-300 km), an aspect of space weather. These irregularities are known to scatter radio signals, potentially interfering with space-based resources such as GPS and global communication, but their formation is not well understood. The scientific payload is an ultra-high frequency (UHF) radar receiver that works in conjunction with ground-based transmitters in a bistatic configuration: a ground-based transmitter illuminates the FAI while the satellite-based receiver passes overhead, recording both the direct signals and those scattered from the irregularities. This measurement concept is illustrated in Figure 2.1. The radar measurements collected by the RAX receiver, combined with other measurements of ionospheric

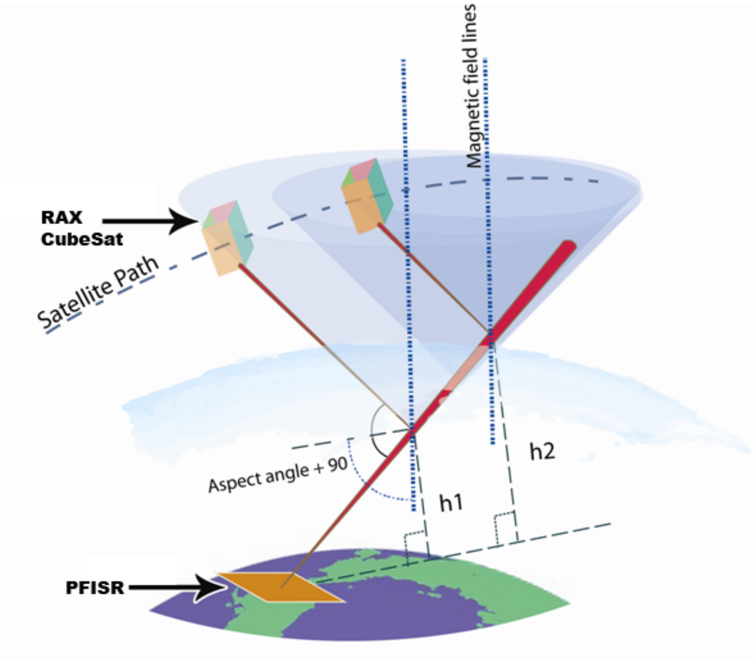

Figure 2.1: An illustration of the RAX concept of operations. A ground-based incoherent scatter radar station (PFISR) transmits pulses (red cone) into the ionosphere. The signals scatter from the irregularities, which are aligned with Earth's magnetic field, into space. The loci of scatter is represented by the cones. Irregularities are shown at two different altitudes, h1 and h2, and the two cubes represent two different times during a single satellite pass.

conditions provided by the ground-based radar, provide information on the formation and structure of the FAI. The ultimate goal of the mission is to enable the development of short-term forecast models for the ionospheric irregularities.

There are currently two satellites in the RAX mission; RAX-1 launched November 2010 and RAX-2 launched October 2011. Both satellites are 3U CubeSats and the flight satellites are shown in Figure 2.2. RAX-1 successfully completed science experiments (radar measurements), demonstrating the payload and bus capabilities, but the mission ended prematurely after three months of operations due to a flaw in the solar panel design [29, 30]. RAX-2 corrected the solar panel issue and operated for 1.5 years on-orbit, successfully completing its planned one year mission and providing unprecedented measurements of the ionospheric irregularities [11]. More information on the scientific mission and satellite development can be found in References 27 and 28.

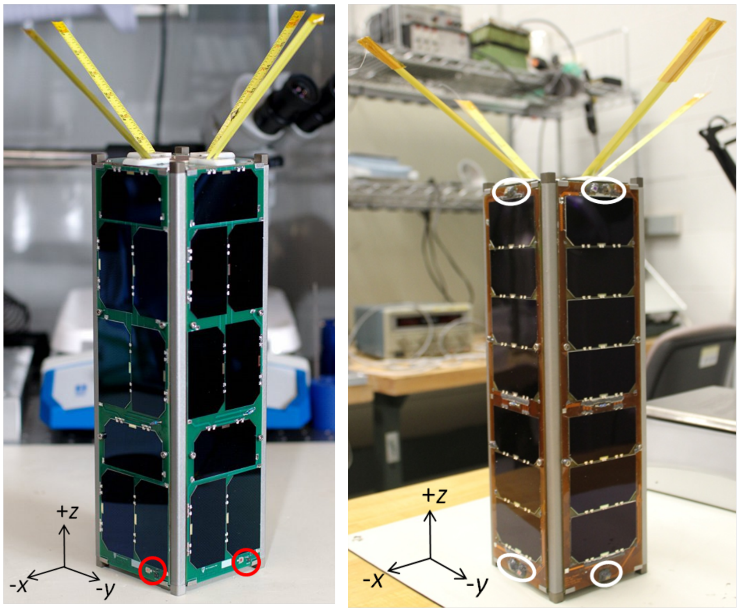

(a) RAX-1 (b) RAX-2

Figure 2.2: The flight-ready RAX satellites. The body-fixed coordinate system is shown and the locations of photodiodes are circled. Photodiodes are located on all six sides of the spacecraft, but only the locations visible in the photos are indicated.

2.2 Attitude Determination and Control Requirements

The RAX attitude determination and control requirements are derived from the requirements of the scientific mission, which state that the gain of the UHF antenna system of the radar receiver shall be known to 1 dB. From the antenna gain pattern, this translates to an attitude determination requirement of 5° 1-σ in each axis [31]. This relatively coarse requirement, coupled with the initial satellite development time constraints of one year between the start of the project and delivery of the flight vehicle, led to the utilization of a magnetometer- and photodiode-based attitude determination sensor suite.

Attitude control is also driven by the 1 dB antenna gain requirement. Given the high-latitude locations of the ground-based transmitters used for the mission as well as the near-omnidirectional antenna gain pattern in its $+z$ hemisphere (see coordinate system in Figure 2.2), a passive magnetic attitude control system meets the attitude control requirements. This type of attitude control system utilizes permanent magnets to align the spacecraft with the geomagnetic field, and utilizes soft magnetic material to dissipate rotational kinetic energy. This passive magnetic scheme is a common on CubeSats due to its simplicity and low cost [32, 33]. The permanent magnets of the RAX control system are aligned such that the RAX $+z$-axis is Earth-pointing at high latitudes.

Since the spacecraft attitude is controlled passively and scientific data processing is done on the ground, RAX has no real-time attitude determination requirement. Instead, the attitude sensor measurements during the times of interest are downloaded for analysis, and these measurements are processed on the ground to estimate attitude. The methods and results of attitude estimation from the sensor measurements will be discussed in Chapter 6.

2.3 Sensor Selection and Specifications

In addition to the angular accuracy requirements for spacecraft attitude determination, factors such as power, volume, mass, and temperature drive the sensor selection, which is described in this section. In general, and as described in Section 1.2, spacecraft attitude is estimated using directional measurements (vector measurements) and corresponding reference directions. For example, magnetometers are used to measure the local magnetic field vector in the satellite body-fixed frame. The corresponding reference direction is the expected magnetic field vector resolved in the reference frame, and is given by a model such as the International Geomagnetic Reference Field (IGRF) coupled with spacecraft location. Both the measured and reference directions can be used to estimate attitude using a variety of methods [24]. Other typical attitude determination sensors were discussed in Section 1.2 and include sun sensors, horizon sensors, and star trackers. Typical accuracy ranges of the different types of sensors are given in Table 1.2.

The attitude sensors on RAX-1 and RAX-2 are identical with one exception: redundant magnetometers on RAX-1 were replaced with additional sun sensors on RAX-2. This is described in the applicable portions of this section as well as Section 2.4, and the difference in photodiode configurations is the topic of Section 2.5. Given the similarities of the ADS on RAX-1 and RAX-2, the satellites are simply referred to as RAX in many cases.

The RAX ADS employs magnetometers and sun sensors because of their ability to meet the angular attitude determination requirements in a relatively simple and low-cost manner. The magnetometers are used to measure the local geomagnetic field vector in the satellite body-fixed frame, and the sun sensors are used to measure the line-of-sight vector to the sun, herein referred to as the sun vector, in the body-fixed frame. In addition to these directional sensors, RAX employs a MEMS three-axis rate gyroscope. The rate measurements are fused with the directional measurements using an EKF, which is discussed in Section 5.3.

Multiple magnetometers are included in the RAX ADS. On RAX-1, this includes two three-axis magnetometers embedded within the satellite as well as four two-axis magnetometers on the solar panels. On RAX-2, only the two three-axis magnetometers are used;

the two-axis magnetometers on the solar panels were not included. The magnetometers were distributed throughout the spacecraft on RAX-1 with the idea of using the distributed measurements to characterize the spacecraft-generated magnetic fields. But this idea was ultimately replaced by the magnetometer calibration of Chapter 3, which enables accurate measurement of the geomagnetic field using a single magnetometer embedded within the spacecraft. For the three-axis magnetometers, sensors included in the initial trade study were the Honeywell HMR2300, HMC2003, HMC1053, and the PNI MicroMag3. The PNI MicroMag3 was chosen because of its relatively high resolution per unit power and its digital interface, which reduces the supporting circuitry required. Specifications of the PNI magnetometer are provided in Table 2.1.

The second three-axis magnetometer is part of the inertial measurement unit (IMU), an Analog Devices ADIS16405. Other IMUs in the initial trade study included the Microstrain 3DM-GX2, InterSense InertiaCube3, O-Navi FalconGX, O-Navi GyroCube3F, Honeywell HG1930, and Micro Aerospace Solutions MASIMU02. An IMU was used rather than a stand-alone rate gyro because the particular IMU chosen included a three-axis magnetometer. Additionally, the ADIS16405 had a low volume and cost compared to others under consideration, and has built-in temperature correction in that sensor measurements are automatically compensated such that they are not adversely affected by temperature. Basic specifications of the rate gyro are given in Table 2.2.

In addition to the two three-axis magnetometers, two-axis magnetometers are located on four sides of the spacecraft of RAX-1 (these were not included on RAX-2). Because the sensors are located on external faces of RAX and must meet clearance requirements for the P-POD launcher[1], size is critical. Honeywell HMC1052L magnetometers were used due to their small size ($3.0 \times 3.0 \times 0.90$ mm^3) and analog interface. The analog interface minimized wiring by utilizing existing connections between peripheral analog-to-digital converters (ADCs) and the Attitude Determination Board[2] (the Attitude Determination Board is described in Section 2.4). Specifications of the magnetometers are provided in Table 2.1.

---

[1] P-POD, or Poly-PicoSatellite Orbital Deployer, is a standardized CubeSat deployment system [34].

[2] Here, we say an analog interface is advantageous, but for the PNI magnetometer, the digital interface was advantageous. This is due to the location of the HMC1052Ls. They are located near the sun sensors, which are also analog, so an ADC is already available and use of analog magnetometers reduces the harness size.

[3] Chosen to decrease latency, which is approximately proportional to the resolution.

[4] The poor resolution is due to quantization error from the analog to digital converters (ADCs) selected. The resolution should be 218 nT, but an incorrect ADC was used due to an error in the interface control document.

[5] The RAX default dynamic range was chosen because it is the smallest possible dynamic range, which corresponds to the best resolution, and, based on predictions of the dynamics of satellite deployment into orbit, RAX angular rates should remain well below 75 deg/s.

Table 2.1: Characteristics of the RAX magnetometers.

Characteristic	**PNI MicroMag3**	**ADIS 16405 IMU**	**HMC1052L**
Dynamic Range	Adjustable; ± 1100 μT max	250 μT	600 μT
Resolution	Adjustable; 15 nT max	50 nT	12 nT max
RAX Default Resolution	128 nT [3]	50 nT	872 nT [4]

Table 2.2: Characteristics of the rate gyro, part of the Analog Devices ADIS16405 IMU.

Characteristic	**Value**
Dynamic Range	Adjustable; 300 deg/s max
Resolution	Adjustable; 0.0125 deg/s min
RAX Default Dynamic Range[5]	75 deg/s
RAX Default Resolution	0.0125 deg/s

Photodiodes, also referred to as coarse sun sensors, were chosen instead of higher accuracy two-component sun sensors because of their simplicity and low cost, and coarse sun sensors coupled with magnetometers would meet the attitude determination requirement. Photodiodes provide a measurement of a single component of the sun vector. Measurements from multiple photodiodes can be combined for a three-component sun vector measurement, which will be discussed in Section 2.5 and Chapter 4, and an in-depth discussion of photodiode measurements as well as their measurement model is discussed in Chapter 5.

From an initial trade study, the Vishay TEMD6010FX01 and Osram SFH-2430 were selected as two possible photodiodes due to their relative wide field of view and operating temperature. The Osram photodiode was chosen over the Vishay because testing showed that the Vishay had a slower response time, which could result in the photodiode measurements lagging behind other sensors. Specifications of the Osram sun sensors are included in Table 2.3.

To account for possible sensor dependence on temperature, National Semiconductor LM20 temperature sensors were placed around the spacecraft to monitor temperatures

Table 2.3: Characteristics of the Osram SFH 2430 photodiodes.

Characteristic	**Value**
Wavelength of peak sensitivity	570 nm
Spectral range of sensitivity [6]	400-900 nm
Half-angle FOV	60° [7]

---

[6] Wavelength of the visible spectrum is approximately 400-700 nm. The remaining 700-900 is infrared.

[7] Sensor testing showed the sensor is accurate up to a half-angle of 70°; see Reference 31.

near the various attitude sensors. The desire for a common temperature sensor across RAX subsystems and CubeSat P-POD size requirements led to the selection of LM20 sensors. Other sensors under consideration included the Sensirion SHT1x, Maxin DS18B20 and DS1620, SMSCEMC1403, Microchip TC72-2.8MUA, Honeywell 777-A-U-0, and Melexis MLX90601EZABAA. In addition to the LM20s, the ADIS16405 includes a temperature sensor. Basic characteristics of the temperature sensors are given in Table 2.4.

Table 2.4: Characteristics of the temperature sensors. The maximum error of the ADIS16405 is not given because it is not reported on the manufacture-provided specification sheet and the error was not experimentally characterized.

Characteristic	**ADIS 16405 IMU**	**LM20**
Dynamic Range	-40 to $+85$ °C	-55 to $+130$ °C
Resolution	0.14 °C	0.4°
Maximum Error	n/a	±5°C

2.4 Integrated Hardware

The RAX attitude determination hardware consists of a central computer and sensor board inside the spacecraft, referred to as the Attitude Determination Board (ADB), with clusters of peripheral sensors located on each of the six external faces of the spacecraft. Each peripheral sensor cluster communicates directly with the ADB through an individual wire harness. This configuration allows for a distributed sampling of sensors across the spacecraft, which is required since sun sensors are mounted on the external faces of the spacecraft. We describe the ADB and peripheral sensors in Section 2.4.1, and discuss timing in 2.4.2.

2.4.1 ADB & Peripheral Sensors

The main components of the ADB include one Texas Instruments MSP430F1611 microprocessor, one external watchdog timer (Maxim MAX6814XK) and reset supervisor for the microprocessor, one Secure Digital (SD) memory card, one PNI MicroMag3 three-axis magnetometer, one Analog Devices ADIS16405 IMU, which contains a three-axis magnetometer, a three-axis gyroscope, and a three-axis accelerometer (not used for attitude determination), and a National Semiconductor LM20 temperature sensor. A block diagram of the ADB along with the actual ADB is shown in Figure 2.3. The ADB interfaces with three other RAX subsystems: the flight central processing unit (FCPU), electrical power system (EPS), and position and time board (PTB).

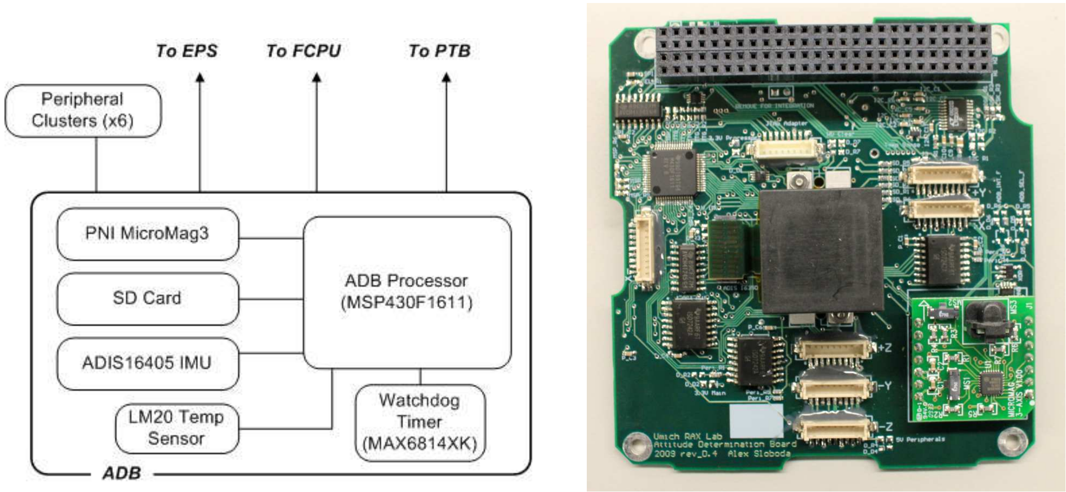

(a) Simplified block diagram of the RAX ADB. The major components are two three-axis magnetometers, a Secure Digital (SD) card, the processor, and watchdog timer. The board also interfaces the three other subsystems: the electrical power system (EPS), flight central processing unit (FCPU), and position and time board (PTB).

(b) The flight unit ADB. The black box in the center of the board is the ADIS16405 IMU, and the square chip on the bottom right is the PNI MicroMag3. The physical dimensions are approximately 10 cm × 10 cm × 2.4 cm.

Figure 2.3: Block diagram and actual RAX Attitude Determination Board (ADB).

The external clusters of sensors each have an ADC (Maxim MAX1249). Nine-wire harnesses connecting these clusters to the ADB carry both power and data. Four of the peripheral sensor clusters reside on the solar panels, and there is a single cluster each on the +z and -z surfaces (coordinate system shown in Figure 2.2). Each of the RAX-1 solar panel clusters is identical: they each contain one Honeywell HMC1052L two-axis magnetometer, one Osram SFH 2430 photodiode, and one LM20 temperature sensor. These clusters are shown in Figure 2.4. On RAX-2, the two-axis magnetometer was replaced with two additional photodiodes, and unlike the photodiodes of the RAX-1 solar panels which are mounted flat on the surfaces, the RAX-2 photodiodes are mounted at angles; this is discussed in Section 2.5.

The peripheral sensor clusters on the +z and -z surfaces are the same for RAX-1 and RAX-2. The cluster on the +z face of the spacecraft contains three SFH 2430 photodiodes and one LM20 temperature sensor. This cluster is shown in Figure 2.5. The cluster on the -z face of the spacecraft contains two photodiodes, one temperature sensor, and one radiation-hardened voltage reference (Linear Technology RH1009), and is shown in Figure 2.6. The redundant photodiodes on the +z and -z faces were introduced to enable sensing despite shadowing from the radar receiver and GPS antennas.

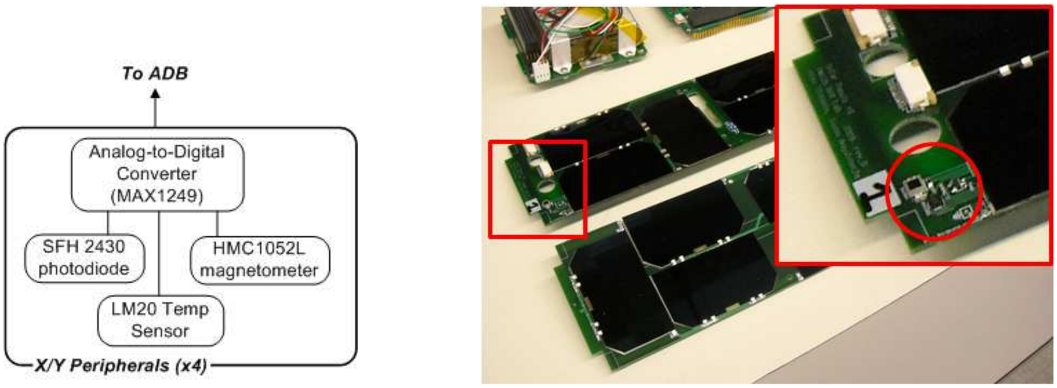

(a) Simplified block diagram of the attitude determination hardware on the RAX-1 solar panels. The major components include an ADC, photodiode, two-axis magnetometer, and temperature sensor. The cluster is connected to the ADB via a wire harness (nine wires).

(b) The determination circuitry is located in a corner of each solar panel. The components are circled in the image in the top right.

Figure 2.4: Block diagram and actual hardware showing the peripheral clusters on the $\pm x$ and $\pm y$ faces (solar panels) of RAX-1.

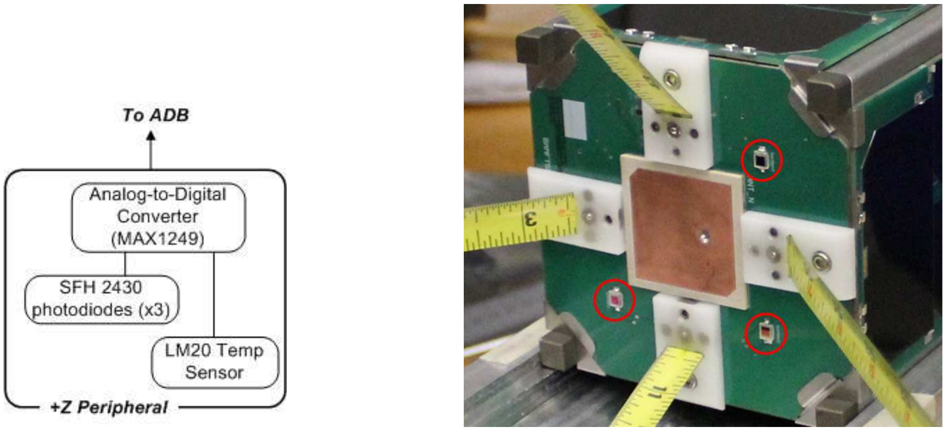

(a) Simplified block diagram of the components in the peripheral circuitry located on the $+z$ face of RAX-1 and RAX-2. The hardware includes three photodiodes, a temperature sensor, and an ADC. The cluster is connected to the ADB via a wire harness (nine wires).

(b) The $+z$ face of RAX-1. The three photodiodes are circled, and the other components are not visible. Redundant photodiodes are used because of possible shadowing from the radar receiver antenna elements seen extending from the $+z$ face.

Figure 2.5: Block diagram and a picture of determination hardware located on the $+z$ face of RAX.

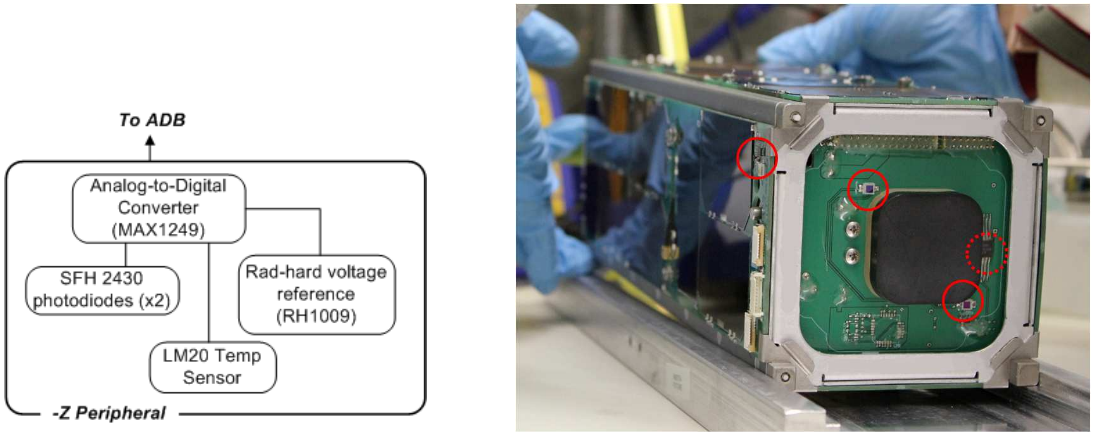

(a) Simplified block diagram of the components in the peripheral circuitry located on the -z face of RAX-1 and RAX-2. The hardware includes two photodiodes, a temperature sensor, a radiation-hardened voltage reference, and an ADC. The cluster is connected to the ADB via a wire harness (nine wires).

(b) The -z face of RAX-1. The two photodiodes and the peripheral cluster on the +x face are circled with a solid line, and the radiation-hardened reference voltage is circled with the dashed line. The temperature sensor cannot be seen clearly in the figure. Redundant photodiodes are used because of possible shadowing from the GPS antenna, which is the black protrusion in the middle of the -z face.

Figure 2.6: Block diagram and a picture of determination hardware located on the -z face of RAX.

The radiation-hardened voltage reference on the exterior of -z face is included to characterize total radiation dose effects on the non-radiation-hardened ADCs. The radiation-hardened voltage reference is designed to output a constant 2.5 V throughout the mission. The ADC measurement of this 2.5 V will drift over time, quantifying the ADC error induced due to radiation. This error can then be accounted for in processing all ADC data, which mitigates radiation errors in all data sets by using only one radiation-hardened part (instead of several expensive radiation-hardened ADC). A similar radiation-hardened voltage reference is included inside the spacecraft to compare total dose effects between the exterior and interior.

2.4.2 Timing

The time of the ADS sensor data is recorded by time-tagging all collected sensor measurements. The ADB monitors the passage of time in two ways: by (1) using the spacecraft's digital pulse-per-second signal on the PTB[8] to increment a seconds counter for coarse tim-

---

[8] The PTB maintains the time for all spacecraft subsystems, facilitating an accurate comparison of data when characterizing FAI. It includes a GPS receiver, and the time is accurate to less than 1 microsecond when the receiver is locked to the GPS constellation. Time synchronization on RAX is discussed further in

ing, and by (2) using the microprocessor to maintain its own milliseconds counter for fine timing. Each sensor datum can then be associated with a timestamp that has one millisecond resolution. At any point in the mission, the FCPU may command an arbitrary value to the seconds counter on-board the ADB, thus changing the ADB's local time. These counters are not operational when the ADB is powered off, thus it is necessary for the FCPU to command a given time to the ADB whenever the ADB is power cycled.

2.5 Photodiode Configurations

The only difference between the attitude determination hardware configuration of RAX-1 and RAX-2 is the number and configuration of the photodiodes. RAX-1 utilized one two-axis magnetometer and one photodiode on each of its four solar panels, as is shown in Figure 2.4. On RAX-2, the two-axis magnetometers were removed from the solar panels and replace with two additional photodiodes on each solar panel. This was done for two reasons: (1) the two-axis magnetometers on the RAX-1 solar panels were not utilized – with the methods of Chapter 3, only a single three-axis sensor is needed – and (2) to increase the photodiode coverage over the body-fixed frame.

On RAX-1, photodiodes are mounted flat on each of the six sides of the spacecraft. The orientation of each photodiode is given in Table 2.5. The azimuth and elevation angles given in the table define the direction normal to the photosensitive plane (analogous to the borseight direction), which we also refer to as the photodiode normal direction. This configuration is fairly typical of CubeSats, but it does not provide three-component sun vector coverage over the body-fixed frame. Since each photodiode provides one component of sun vector information (discussed further in Chapter 4), and the FOV of the individual photodiodes is less than 180°, this configuration provides three-component sun vector coverage over a very limited portion of the body-fixed frame. This is illustrated with a two-dimensional example in Figure 2.7. The map of sun sensor coverage of the satellite body frame that results from this configuration is shown in Figure 2.8(a).

On RAX-2, the additional photodiodes were mounted at angles on the $\pm x/y$ surfaces of the satellites, while the number and orientation of photodiodes on the $\pm z$ surfaces was left the same as RAX-1. A picture of two of the angled photodiodes on a RAX-2 prototype solar panel is shown in Figure 2.9. The orientation of each photodiode on RAX-2 is listed in Table 2.6, and some of their physical locations on the spacecraft are indicated in Figure 2.2(b). The resulting sun sensor coverage over the body-fixed frame is shown in Figure 2.8(b). This configuration was chosen because it provides three-component coverage over

---

Reference 35.

Table 2.5: The intended azimuth and elevation angles of the photodiode normal directions on RAX-1. The side panel is the surface to which the sensors are mounted (coordinate system shown in Figure 2.2(a)).

Photodiode	**Side Panel**	**Azimuth (deg)**	**Elevation (deg)**
1	+x	0	0
2	-x	180	0
3	+y	90	0
4	-y	270	0
5	+z	0	90
6	+z	0	90
7	+z	0	90
8	-z	0	-90
9	-z	0	-90

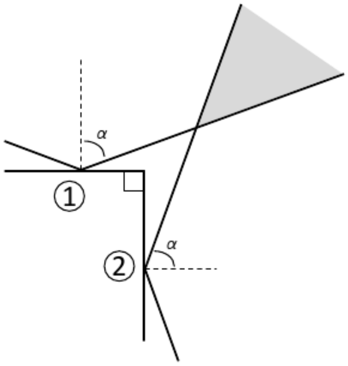

Figure 2.7: An illustration of photodiode coverage for a two-dimensional example. Two orthogonal sensors, labeled 1 and 2, with a $\alpha = 70°$ FOV from the normal direction of each sensor are shown. The sun must be in the field of view of both sensors for a two-component vector measurement, limiting the sun sensor coverage to the 50° shaded area.

Table 2.6: The intended azimuth and elevation angles of the photodiode normal directions on RAX-2. The side panel is the surface to which the sensors are mounted (coordinate system shown in Figure 2.2(b)).

Photodiode #	Side panel	Azimuth (deg)	Elevation (deg)
1	+x	17	-10
2	+x	0	20
3	+x	-17	-10
4	-x	-162	-10
5	-x	180	20
6	-x	162	-10
7	+y	72	10
8	+y	107	10
9	+y	90	-20
10	-y	-107	10
11	-y	-72	10
12	-y	90	-20
13	+z	0	90
14	+z	0	90
15	+z	0	90
16	-z	0	-90
17	-z	0	-90

96% of the body frame – a significant improvement over the RAX-1 coverage shown in Figure 2.8(a) – while adhering to the CubeSat specifications that limit the height of objects mounted to the CubeSat surfaces [36]. It was chosen from manual iterations of different possible configurations. Since this design, a design optimization method to provide an optimal photodiode configuration has been developed, providing an optimal design techniques that replaces manual, iterative-based approaches to sensor placement; this is the topic of Chapter 4.

In addition to the orientation of the photodiodes, a second difference in the photodiode configuration of RAX-2 compared to RAX-1 is the added use of solar cell coverglass on RAX-2. On RAX-1, the photodiodes were utilized with no additional shielding. On RAX-2, solar cell coverglass was added to the surface of each photodiode to shield from ultraviolet radiation, which caused some degradation of the sensors on RAX-1. This degradation is shown in Section 5.4.4.

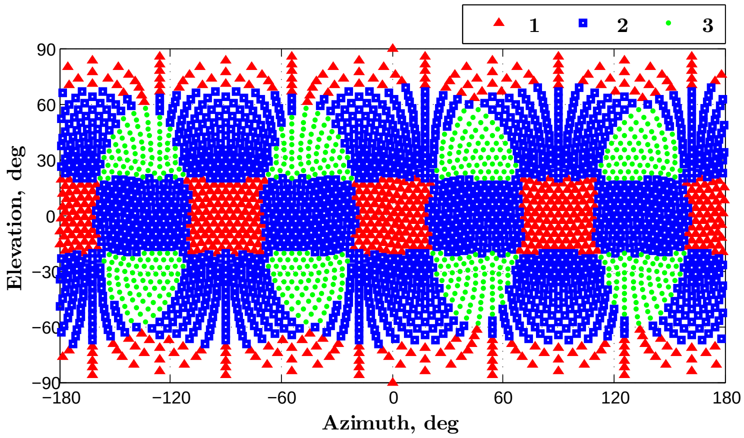

(a) RAX-1 photodiode coverage.

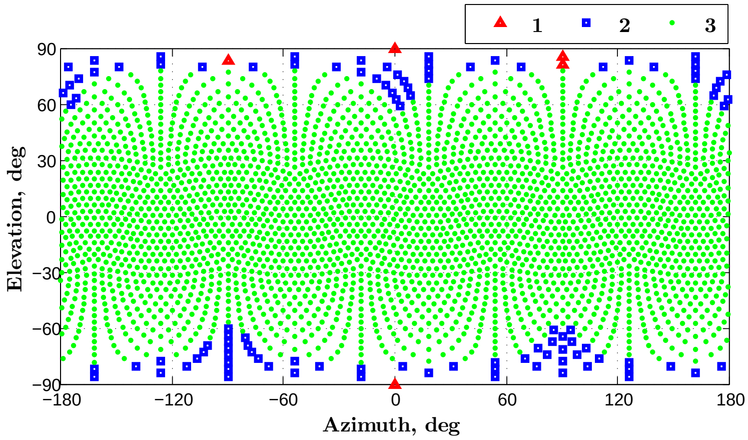

(b) RAX-2 photodiode coverage.

Figure 2.8: Photodiode coverage over the body-fixed frame for RAX-1 and RAX-2. Each symbol on the plot is a direction in the body-fixed frame, and the symbols have a uniform angular distribution of approximately 4°. The symbols are shaped and color-coded by the number of sun vector components measured when the sun is in the direction of the symbol. On RAX-1 the three component availability corresponds to exactly three photodiodes being illuminated, but on RAX-2, up to seven photodiodes are illuminated for portions of the body frame.

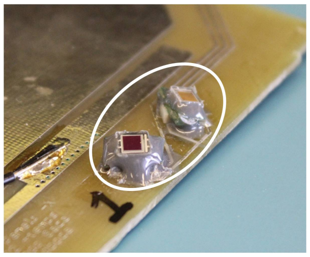

Figure 2.9: Two angled photodiodes located near the -z edge of the -x surface (see coordinate system in Figure 2.2(b)).

2.6 Summary

The RAX ADS utilizes magnetometers, photodiodes, and a MEMS gyroscope for attitude determination. This is an example of a low-cost attitude determination system, and although it was developed for the RAX mission, it could be used for attitude determination on any LEO satellite. The system is implemented using a central attitude determination processing and sensor board in combination with clusters of sensors located on the external faces of the spacecraft. Sensor calibration was not discussed in this section; some pre-flight calibration was done on RAX-1 [31], but this was ultimately replaced by the on-orbit calibration that is presented in Chapters 3 and 5. The sensor calibration methods used and the resulting attitude determination accuracy achieved with this system is discussed in Chapter 6.

CHAPTER 3

On-Orbit Magnetometer Calibration

Attitude sensor calibration is critical for accurate attitude determination. Magnetometers, which are used in an attitude determination system to measure Earth's magnetic field, are subject to errors both inherent to the sensor and due to the surrounding spacecraft environment. In this chapter, we present a method for attitude-independent, on-orbit magnetometer calibration that mitigates the effect of time-varying magnetic fields produced by electronics on-board a spacecraft. This method is attitude-independent in the sense that it neither requires attitude knowledge nor estimates attitude during the calibration process. The calibration significantly increases the accuracy of measurements from magnetometers embedded within a satellite, enabling them to make accurate measurements of the geomagnetic field without imposing location constraints on the sensor. Additionally, since the calibration requires only on-orbit data, it mitigates the need for pre-flight calibration.

3.1 Introduction

Three-axis magnetometers are a common sensor on spacecraft in low-Earth orbit (LEO) because they are reliable, lightweight, have low power requirements, and have no moving parts [37]. They are typically used for attitude determination and scientific measurements. In this work, calibration was motivated by improved magnetometer-based attitude determination on satellites, but the calibration is applicable to other magnetometer-based sensing systems on a variety of platforms, such as air-, ground-, or water-based vehicles [38].

The purpose of calibration is to quantify the statistical properties of the magnetometer errors [39, 40]. Magnetometers are subject to both constant and time-varying errors, including errors caused by hard iron, soft iron, null shift, scale factors, non-orthogonality, and nearby electronics. These sources of error will be discussed in Section 3.2. Various algorithms exist to estimate the time-invariant errors, which are captured by three general error types: bias, scale factors, and non-orthogonality. For example, TWOSTEP estimates magnetometer bias [41,42] and has been extended to estimate scale factors and non-orthogonality [43]. The first step is a centering approximation to provide an initial estimate of the calibration parameters, and the second step is a Gauss-Newton method to iteratively refine the parameters. In similar work, [44] utilizes a geometric approach to formulate the problem of compensating for magnetometer errors by estimating parameters lying on an ellipsoid. A simpler method that utilizes least-squares minimization to estimate bias, scale factors, and non-orthogonality is developed in [45,46]. These algorithms [41–46] are batch methods, but real-time magnetometer correction can be implemented by uploading the calibration parameters to the spacecraft after the batch calibration is completed. Real-time implementations of [43] using non-linear Kalman filtering techniques are developed in [47].

The above calibration techniques are attitude-independent, meaning no attitude knowledge is required for the calibration. This is critical when magnetometers are used to estimate attitude, and thus attitude measurements are not available before calibration. In lieu of knowledge of the vector components of the ambient magnetic field, attitude-independent calibration algorithms rely on knowledge of the expected magnitude of the geomagnetic field, which is obtained from a trusted sensor regarded as truth or a model such as the IGRF [48]. The sensor calibration is typically carried out in a coordinate system attached to the magnetometer. However, the calibration methods presented in [44] and [49], an extension of [46], include an additional step to estimate the alignment of the magnetometer relative to the vehicle body frame.

The existing algorithms have been shown to work well in compensating for constant sources of magnetometer error [43, 44, 46]. In practice, the time-varying bias caused by nearby electronics can result in additional magnetometer errors. Traditionally, this bias is mitigated by either using a boom to physically separate the magnetometer from the spacecraft, or by using design and manufacturing practices to minimize the effect of electronic components (for example, [50]). These design practices increase the satellite development time and cost, and with the trend toward smaller spacecraft with reduced development times and costs, it may be impossible to physically separate a magnetometer from other spacecraft electronics.

The original contribution of this work is an attitude-independent, on-orbit method to estimate magnetometer bias caused by nearby electronics. In particular, the work of Foster and Elkaim [46] is expanded to include time-varying bias in the calibration. This is accomplished by including measurements of spacecraft electric currents in the sensor model and estimating constant parameters that map time-varying current to magnetometer bias. In similar work, Kim et al [51] estimate the bias resulting from a magnetic torque coil by expanding the model of [43] to include the magnetic dipole moment produced by the torque coil. Our method is more general in that the bias caused by any electronic component can be estimated. The current measurements used in the calibration are typically already part of spacecraft health monitoring, so no additional sensors are needed. This method simplifies the spacecraft design process by replacing magnetic cleanliness, magnetometer location, and pre-flight calibration requirements with on-orbit calibration. The effectiveness of the calibration is demonstrated by application to flight data from RAX-1.

The remainder of this chapter is organized as follows. In Section 3.2, we discuss the various sources of magnetometer error before presenting a sensor error model and reviewing an existing calibration method in Section 3.3. The calibration method is then extended to include time-varying bias in Section 3.4. In Section 3.5, we apply both the time-invariant and time-varying calibrations to on-orbit data, and we discuss the improvement over time-invariant methods, convergence of the calibration parameters, and the accuracy of the calibration. Conclusions are given in Section 3.6.

3.2 Magnetometer Errors

Three-axis magnetometer measurements may be corrupted by both constant [45, 46] and time-varying sources of error. These error sources, which are described below, include sources inherent to the sensor, which can be caused by factors such as manufacturing tolerances in the construction of the sensor, as well as errors resulting from the surrounding environment, such as magnetic materials in close physical proximity to the magnetometer.

Hard iron error is a constant magnetic field bias that is caused by unwanted magnetic fields near the magnetometer, typically due to ferromagnetic (hard iron) materials. The permanent magnetization of these materials perturbs the magnetic field at the location of the sensor, resulting in an offset between the ambient and measured magnetic fields. Hard iron errors are independent of the ambient magnetic field vector, meaning that this error can be parametrized by a constant bias in each magnetometer axis.

Null shift error, also known as dc offset or zero bias, is a constant offset that shifts the output of the sensor. Like hard iron error, null shift also results in a constant magnetometer bias, but null shift error is inherent to the sensor.

Soft iron materials generate magnetic fields in response to externally applied fields. Like hard iron, this results in a magnetometer bias, but soft iron error depends on the externally applied field. In this work, we assume a linear (non-hysteretic) between an external magnetic field and the resulting magnetic field produced by a soft iron material. This assumption, which is commonly used in the literature, is justified because the magnitude of the fields produced by soft iron materials are insignificant compared to the ambient (geomagnetic) field. Under the linearity assumption, these errors are parameterized by a combination of a scale factor and misalignment terms in the model [46], which is discussed in Section 3.3.

Scale factor errors are inherent to the sensor and are caused by uncertainty in the constant of proportionality between the local magnetic field (sensor input) and the sensor output. Soft iron errors also cause a scale factor error, but soft iron errors are caused by soft iron materials, whereas the scale factor error is inherent to the sensor.

Non-linearity error includes any deviation from a linear relationship between the sensor input and output. In modeling the sensor, which is discussed in Sections 3.3 and 3.4, non-linear behavior is assumed to be negligible and thus is not included in the sensor model. This is a valid assumption because, similar to the soft iron error, deviations from linear sensor behavior are negligible ($< 1\%$), even for low-cost magnetometers. For example, for two sensors used in this work, a PNI MicroMag3[1]and Analog Devices ADIS16405[2], the manufacturers report non-linearity error of 0.6% and 0.5% of the field strength, respectively, from a best-fit straight line relating the sensor input and output.

Non-orthogonality error is deviation from orthogonal sensor axes. It can be caused by manufacturing error, thermal strain, or mechanical strain.

Sensor noise is the stochastic component of the sensor output. It is mean zero and typically has a Gaussian distribution, but in practice, the distribution of sensor noise should be characterized. For the sensors used as application examples in this work, we have verified through testing in a controlled environment that the assumption of Guassian zero mean noise is valid.

Time-varying bias refers to the magnetic field bias caused by nearby electronics. Current-carrying wires generate magnetic fields, resulting in a time-varying offset between the measured and geomagnetic field.

3.3 Review of Existing Methods

The new method for estimating time-varying magnetometer bias is an extension of an existing magnetometer calibration algorithm that estimates time-invariant errors [46]. For completeness, the existing method is reviewed before presenting the extension.

---

[1] Specification sheet available online: `https://www.sparkfun.com/datasheets/Sensors/MicroMag3%20Data%20Sheet.pdf`, accessed September 2013.

[2] Specification sheet available online: `http://www.analog.com/static/imported-files/data_sheets/ADIS16400_16405.pdf`, accessed September 2013.

In general, there are two steps in sensor calibration: (1) model the sensor, and (2) estimate parameters of the model. A right-handed orthogonal coordinate frame is used to model the sensor; this is referred to as the *rectified frame*. The *sensor frame* is defined by the actual axes of the sensor, which are generally not orthogonal due to manufacturing error, thermal strain, or mechanical strain. No knowledge of the orientation of these frames relative to the spacecraft frame is required during calibration. The axes of the rectified frame are denoted by x, y, and z, whereas the sensor axes are denoted by $\tilde{x}$, $\tilde{y}$, and $\tilde{z}$. The rectified frame is shown in Figure 3.1 and is defined by the sensor frame: the x-axis of the rectified frame is coincident with $\tilde{x}$, the y-axis is in the plane of $\tilde{x}$ and $\tilde{y}$, and the z-axis completes the right-handed frame. The calibration is carried out in the rectified frame, and estimation of the relative alignment between the rectified frame and the frames of other attitude determination sensors can be carried out after the magnetometer calibration using existing methods [52, 53]. Attitude estimation after the magnetometer calibration without additional alignment estimation would yield an estimate of the attitude of the rectified frame relative to the reference frame.

The three-axis sensor model that includes the time-invariant errors is [46]

$$\tilde{B}_{\tilde{x}} = aB_x + x_0 + \eta_{\tilde{x}} \ , \tag{3.1}$$

$$\tilde{B}_{\tilde{y}} = b\left(B_y \cos(\rho) + B_x \sin(\rho)\right) + y_0 + \eta_{\tilde{y}} \ , \tag{3.2}$$

$$\tilde{B}_{\tilde{z}} = c(B_x \sin(\lambda) + B_y \sin(\phi)\cos(\lambda) + B_z \cos(\phi)\cos(\lambda)) + z_0 + \eta_{\tilde{z}} \ , \tag{3.3}$$

where $\tilde{B}_{\tilde{x}}$, $\tilde{B}_{\tilde{y}}$, and $\tilde{B}_{\tilde{z}}$ are the measured magnetic fields along each sensing element (each axis of the sensor frame); B_x, B_y, and B_z are the components of the ambient geomagnetic field resolved in the rectified frame; the magnetometer errors are parametrized by scale factors in each axis (a, b, c), constant bias in each axis (x_0, y_0, z_0), and the sensor non-orthogonality angles (ρ, λ, ϕ) modeled with the geometry shown in Figure 3.1; and $\eta_{\tilde{x}}$, $\eta_{\tilde{y}}$, and $\eta_{\tilde{z}}$ are zero-mean measurement noise in sensor axis. These parameters quantify the hard iron, null shift, soft iron, scale factor, non-orthogonality errors, and sensor noise described in Section 3.2.

The purpose of the calibration is to estimate a, b, c, x_0, y_0, z_0, ρ, λ, and ϕ, herein referred to as the calibration parameters. There are three assumptions in the magnetometer model:

1. The magnetometer and the sources of error are rigidly attached to the same platform, meaning that the position and orientation of the sources of error included in the model are fixed relative to the sensor. This assumption is applicable to the hard iron and soft

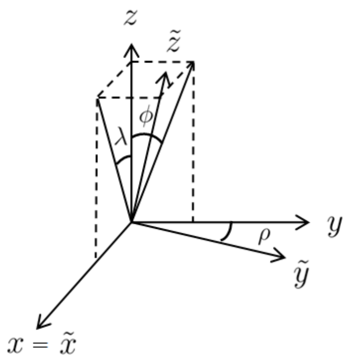

Figure 3.1: The convention used to define the sensor non-orthogonality and the sensor frame. The axes of the right-handed sensor frame are denoted by x, y, and z, and the directions of the actual sensing elements of the sensor are $\tilde{x}$, $\tilde{y}$, and $\tilde{z}$. The x-axis of the sensor frame is aligned with $\tilde{x}$, the y-axis is in the $\tilde{x}$-$\tilde{y}$ plane, and the z-axis completes the right-handed frame. The non-orthogonality angles are ρ, measured from the y-axis; ϕ, measured from the x-z plane; and λ, measured from the y-z plane.

iron errors.

2. Responses to soft iron materials are linear (without hysteresis), which corresponds to the use of scale factors to parametrize errors caused by soft iron material.

3. Other non-linearity errors are also negligible (they are not included in the model).

The justification for assumptions 2 and 3 were given in Section 3.2.

The magnetometer model of Eqs. (3.1)-(3.3) captures the total bias, scaling, and non-orthogonality errors, but it is not possible to mathematically separate each individual error source. For example, both hard iron error and null shift error contribute to the bias in each magnetometer axis, x_0, y_0 and z_0. Similarly, both soft iron and scaling inherent to the sensor contribute to the scale factor errors a, b, and c, and both soft iron errors and physical non-orthogonality errors are manifested in the angles ρ, λ, and ϕ. The mapping from soft iron error to non-orthogonality occurs because soft iron can cause cross-axis gain terms [45,46]. For example, an ambient field in the magnetometer z-axis could cause a soft-iron-induced magnetic field in the x-axis. Even though this ambiguity of physical error sources exists, estimation of the total bias, linear scaling, and non-orthogonality parameters is sufficient to correct the magnetometer measurements [46].

The calibration parameters are estimated using only the expected magnitude of the geomagnetic field as well as the three-axis magnetometer measurements. Previous work and numerical studies show that the calibration parameters are identifiable through the magnetic field magnitude when using measurements that sufficiently cover the sensor field of view, and this is discussed further in Section 3.5.5. The calibration is independent of spacecraft attitude since only the field magnitude, rather than vector components, is required. To formulate the minimization problem and estimate the calibration parameters, Eqs. (3.1)-(3.3) are re-written to give the geomagnetic field components as functions of the measurements and calibration parameters. The squared magnitude of the measured geomagnetic field is then given by

$$B^2 = B_x^2 + B_y^2 + B_z^2 = f(\tilde{B}_{\tilde{x}}, \tilde{B}_{\tilde{y}}, \tilde{B}_{\tilde{z}}, a, b, c, x_0, y_0, z_0, \rho, \phi, \lambda)\ . \tag{3.4}$$

The expected magnetic field magnitude, B_E, is obtained from a model of the geomagnetic field, such as the IGRF [48], coupled with knowledge of spacecraft location. The calibration parameters are estimated by minimizing the difference between the measured and expected geomagnetic field magnitudes, $\min(B_E^2 - B^2)$.

Various methods can be used to minimize and estimate the calibration parameters. In [46], least-squares minimization is used. If the geomagnetic field magnitude is constant, such as during ground-based testing in a fixed location, Eq. (3.4) can be re-written in a linear form using intermediate variables that are functions of the calibration parameters. This facilitates the use of batch linear least squares to estimate the parameters [46]. If the geomagnetic field is time-varying (the vehicle is moving sufficient distances through the geomagnetic field), the minimization can be iteratively processed by a non-linear least squares batch algorithm. In other work, maximum likelihood estimation [43,44] as well as non-linear Kalman filtering [47] have been used to estimate the calibration parameters. In this work, we utilize an iterative non-linear least squares minimization.

3.4 Inclusion of Time-Varying Bias

Time-varying bias due to on-board electronics is estimated by including measurements of electric current in the magnetometer model. Measurements of electric current within the satellite power system and other main components are typically part of satellite health monitoring and thus are included in the spacecraft telemetry. We include these telemetered current measurements in the magnetometer model and estimate constant parameters that map the time-varying current to magnetometer bias. In general, the magnetic field produced by a current-carrying wire is a function of the wire geometry and the current magnitude. This field is given by the Biot-Savart law, which can be integrated in closed-form for only specific wire configurations, such as a circular loop or straight infinite length [54]. However, assuming that the wire geometry and the location and orientation of the magnetometer relative to the wire are constant, the magnetic field produced at the location of the magnetometer can be written as a linear function of the current, where a single coefficient maps the current to the resulting magnetic field component. This yields the following magnetometer model:

$$\tilde{B}_{\tilde{x}} = aB_x + x_0 + \sum_{i=1}^{m_c} s_{i,\tilde{x}}\tilde{I}_i + \eta_{\tilde{x}} \ , \tag{3.5}$$

$$\tilde{B}_{\tilde{y}} = b\left(B_y \cos(\rho) + B_x \sin(\rho)\right) + y_0 + \sum_{i=1}^{m_c} s_{i,\tilde{y}}\tilde{I}_i + \eta_{\tilde{y}} \ , \tag{3.6}$$

$$\tilde{B}_{\tilde{z}} = c(B_x \sin(\lambda) + B_y \sin(\phi)\cos(\lambda) + B_z \cos(\phi)\cos(\lambda)) + z_0 + \sum_{i=1}^{m_c} s_{i,\tilde{z}}\tilde{I}_i + \eta_{\tilde{z}} \ . \tag{3.7}$$

The model is identical to Eqs. (3.1)-(3.3) with one additional term in each axis: $\sum_{i=1}^{m_c} s_{i,j}\tilde{I}_i$, $j \in \{\tilde{x}, \tilde{y}, \tilde{z}\}$, where $s_{i,j}$ is the coefficient that maps the i-th current measurement, $\tilde{I}_i$, to the magnetic field in the j-th magnetometer axis, and m_c is the total number of current measurements included in the model. Even though current measurements are required for the calibration, on-board current sensors are typically part of spacecraft health-monitoring, so inclusion of the current measurements in the calibration does not necessarily require current sensors to be added to a spacecraft solely for the purpose of calibration. The model does not require any knowledge of the layout of the electronic components; the calibration requires only current measurements, magnetometer measurements, and the expected magnitude of the geomagnetic field.

Care must be taken when selecting current telemetry to include in the magnetometer model. Use of the model is effective only when the included currents affect the magnetometers. If not, physically meaningful mapping coefficients, $s_{i,j}$, are not observable, and the accuracy of the calibration will be degraded. Evidence to decide which current measurements to include is discussed in Section 3.5.3. Additionally, the current measurements must be linearly independent to obtain unique mapping coefficients. If linearly dependent measurements are used, the unobservability will be manifested by a rank deficient Jacobian matrix that is required for the estimation ($\frac{\partial \mathbf{f}\left(\tilde{\mathbf{B}},\tilde{\mathbf{I}},\mathbf{x}\right)}{\partial \mathbf{x}}$, see Eq. (3.8)).

In the magnetometer model of Eqs. (3.5)-(3.7), there are $9{+}3m_c$ calibration parameters, and the parameters are estimated in the same manner described in the previous section: the difference between the measured and expected geomagnetic field magnitude is minimized. Non-linear least squares is used for the minimization. Also known as Gaussian least squares differential correction, this method is a generalization of Newton's root solving method that iteratively minimizes the loss function J [55], given by

$$J = \frac{1}{2}\left[\mathbf{B}_E^2 - \mathbf{f}(\tilde{\mathbf{B}}, \tilde{\mathbf{I}}, \mathbf{x})\right]^T \left[\mathbf{B}_E^2 - \mathbf{f}(\tilde{\mathbf{B}}, \tilde{\mathbf{I}}, \mathbf{x})\right] \ . \tag{3.8}$$

In Eq. (3.8), $\mathbf{B}_E^2$ is an $m \times 1$ vector of the expected geomagnetic ield magnitudes squared, where m is the number of measurements, and $\mathbf{f}(\tilde{\mathbf{B}}, \tilde{\mathbf{I}}, \mathbf{x})$ is an $m \times 1$ vector of the field magnitudes squared as a function of the $m \times 3$ magnetometer measurements, $\tilde{\mathbf{B}}$, the $m \times m_c$ current measurements, $\tilde{\mathbf{I}}$, and the calibration parameters, $\mathbf{x}$, at the current iteration. The loss function J is quartic, which can result in the convergence to incorrect local minima. In some minimization techniques, such as that used in [41], centering is used to reduce the loss function from quartic to quadratic. In non-linear least squares, $\mathbf{f}(\tilde{\mathbf{B}}, \tilde{\mathbf{I}}, \mathbf{x})$ is linearized at each iteration, which results in a quadratic loss function without the need for centering. A derivation of non-linear least squares is not included here, but rather, the reader is directed to existing references [55]. In general, convergence of non-linear least squares is not guaranteed, but after numerical testing and application to flight data, we have found no convergence problems. This is discussed further in Section 3.5.4.

In minimizing of the loss function of Eq. (3.8), the calibration parameters are estimated while assuming the measurements are perfect. In practice, the estimated calibration parameters inherit errors due to measurement error, and their accuracy is also dependent on the amount of information available from the measurements (the “richness” or “excitation” of the data). Magnetic field measurements spanning the entire field of view of the magnetometer produce the most accurate calibration parameters. The current measurements must also be non-constant and linearly independent to obtain unique and accuracy current mapping coefficients. These aspects are discussed further in Section 3.5.5.

3.5 Application to Flight Data

The calibration algorithm has been tested through application to both simulated and actual on-orbit data. In this section, the algorithm is applied to on-orbit data from the RAX-1 ADS, discussed in Chapter 2. In Subsection 3.5.1, the on-orbit data and use of the IGRF are discussed. In Subsection 3.5.2, an existing time-invariant calibration method is applied to the on-orbit data, and the results demonstrate the need for calibration with time-varying bias. Application of the time-varying calibration is presented in Subsection 3.5.3, and accuracy of the calibration is discussed in Subsection 3.5.5.

Table 3.1: The time difference between the epochs of the TLEs used and the time of the magnetometer data sets. Three data sets are discussed: December 1, December 15, and December 30, 2010.

Data Set	Time Difference (days)
Dec. 1	2.24
Dec. 15	0.16
Dec. 30	1.08

3.5.1 RAX-1 and IGRF Data

RAX-1 [27], discussed in Section 2.1 and shown in Figure 2.2(a), is a CubeSat (a specific satellite form-factor) with physical dimensions of $30 \times 10 \times 10$ cm^3 and mass of 2.8 kg. It launched November 2010 into a 650 km altitude, 72° inclination circular orbit. As described in Chapter 2, commercial off-the-shelf magnetometers are used in conjunction with photodiodes and a three-axis rate gyroscope for attitude determination. The magnetometers are embedded within the satellite and subject to spacecraft-generated magnetic fields. The data used in this section is from one of the on-board three-axis magnetometers, the PNI Sensor Corporation MicroMag3, which is shown in Figure 3.2. The sensor resolution is 128 nanotesla (nT) in each axis. From pre-flight testing, the noise floor of the stand-alone sensor was found to be below the sensor resolution, so the resolution itself provides a metric for the effectiveness of the on-orbit calibration.

The IGRF [48] is used for the expected field magnitude. The IGRF is a geomagnetic field model that provides geomagnetic field components in an Earth-fixed frame as a function of location. Spacecraft position is obtained from a two-line element (TLE) set. The time differences between the TLE epochs and each data set discussed in this section are given in Table 3.1. AGI's Satellite Toolkit (STK)[3] was used to obtain IGRF data from TLEs. The TLEs are propagated using STK's SGP4 propagator, and the complete (order 13, degree 13) eleventh generation IGRF model [48] is built in to STK, so the expected field magnitude corresponding to each sensor reading is obtained directly from STK based on the time of the sensor reading and the TLE.

The accuracy of the expected field magnitude is a critical aspect of the magnetometer calibration. Accuracy of the IGRF is approximately 10 nT during nominal space weather conditions [56], which is better than typical attitude-grade magnetometers, and is an order of magnitude better than the resolution of the magnetometer used in this section. Space weather conditions must be considered since space weather affects the geomagnetic field,

---

[3] `http://www.agi.com/products/by-product-type/applications/stk/`, accessed September 19, 2011.

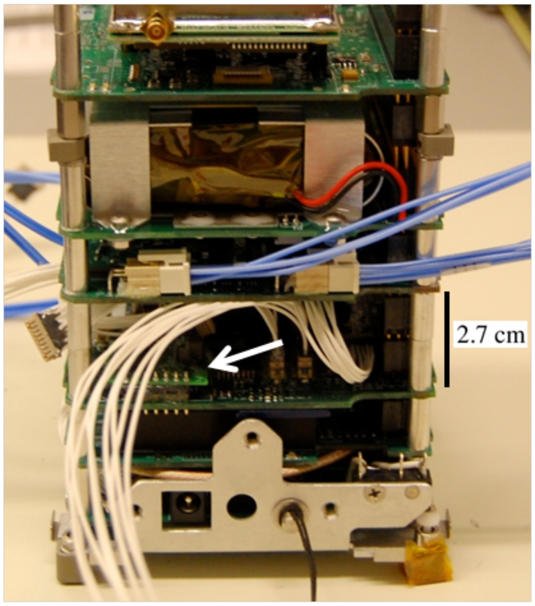

Figure 3.2: The inside of the lower portion of RAX-1 is shown. The arrow points to the PNI MicroMag3 magnetometer. The circuit board located directly above the magnetometer is the satellite power regulation system. The fully integrated RAX-1 satellite is shown in Figure 2.2(a)

and the calibration should only be carried out when the geomagnetic field fluctuations are sufficiently below the resolution of the magnetometer. The Kp index[4], a measure of geomagnetic field activity, was less than 2+ for each data set used for the calibration. This indicates that geomagnetic field fluctuations are expected to be below 20 nT[5] and validates the assumption that the IGRF will provide accurate data for the time periods of interest.

3.5.2 Calibration with Time-Invariant Parameters

To demonstrate the need for calibration that includes time-varying bias, time-invariant calibration is first applied to the on-orbit data before application of the new algorithm. Figure 3.3 shows 112 minutes of on-orbit magnetometer data taken at approximately 1 Hz on December 1, 2010. Figure 3.3(a) shows the raw, uncalibrated magnitude of the measured data, which is overlaid with the expected field magnitude. The calibration method discussed in Section 3.3 is applied to estimate the time-invariant errors, and the magnitude of the corrected measurements are shown in Figure 3.3(b). Compared to the raw data of

---

[4] Kp index retrieved from the NOAA National Geophysical Data Center, `http://www.ngdc.noaa.gov/stp/geomag/kp_ap.html`, accessed September 19, 2011.

[5] The approximate magnetic field fluctuations as a function of Kp index are provided by the NOAA Space Weather Prediction Center, `http://www.swpc.noaa.gov/info/Kindex.html` (accessed December 17, 2011).

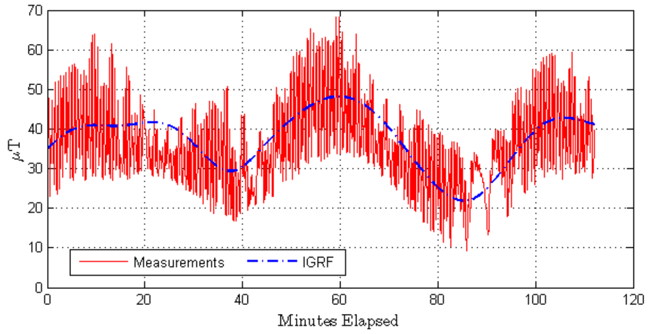

(a) Magnitude of the raw, uncalibrated measurements (μT) overlaid with the expected field magnitude using the IGRF model.

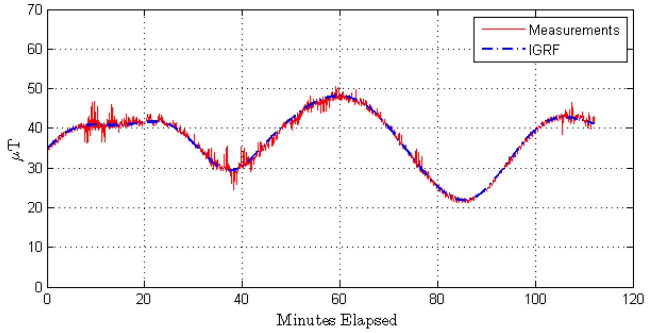

(b) Magnitude of the measurements after correcting for time-invariant errors (μT). The differences between the measured and IGRF magnitudes are shown in Figure 3.4.

Figure 3.3: Data from the RAX-1 PNI MicroMag3 magnetometer. The x-axis of each plot shows time elapsed since the start of the data set, 01-Dec-2010 08:30:46 UTC.

Figure 3.3(a), the measured magnitude is closer to the expected magnitude, but there are still discrepancies of up to 23% of the expected field magnitude. The difference between the corrected measured magnitude and the expected field magnitude is shown in Figure 3.4. In Figure 3.4(a), an indicaor is overlaid to show when the solar panels are illuminated and generating current. The indicator is based on sun sensor readings. There is a clear increase in the magnetometer errors when the satellite is in the sun, which suggests that currents flowing in the solar panels are affecting the magnetometer. A histogram of the differences is shown in Figure 3.4(b).

3.5.3 Application of Time-Varying Calibration

The model used to estimate time-varying bias (Eqs. (3.5)-(3.7)) can include any number of current measurements. Data presented in Figure 3.3 suggests that currents produced by the solar panels degrade the magnetometer measurements. There are four body-mounted

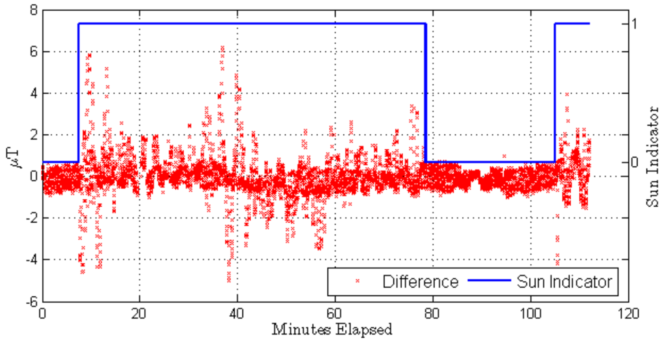

(a) The difference (μT) versus time. The sun indicator takes the value of one when RAX-1 is in the sun, and zero when in eclipse, which shows when the solar panels are illuminated and generating current.

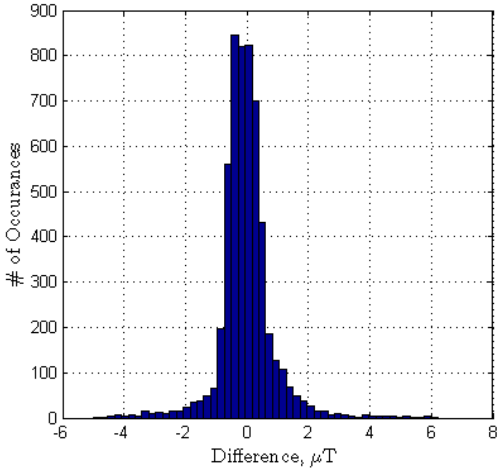

(b) A histogram of the differences (μT).

Figure 3.4: Difference between the magnitude of the corrected measurements using time-invariant calibration, and the expected field magnitude, as shown in Figure 3.3(b).

solar panels on RAX-1, and measurements of the current in each panel are included. From further experimentation with the on-orbit data, a fifth element, the current drawn from the battery, has been found to degrade the measurements [57]. Therefore, these five current measurements are included in the magnetometer model.

For the datasets presented in this chapter, the current and magnetic sensors throughout the spacecraft were sampled simultaneously. In general, the frequency of sensor sampling is mission dependent: magnetometer and current sensors can be sampled at different rates based on the scientific and engineering requirements. The model of Eqs. (3.5)-(3.7) assumes that the current measurements were taken at the same time as the magnetometer measurements. If not, interpolation could be used as an approximation, but this will degrade the results. For general use of this algorithm, the ability to sample sensors simultaneously for the purposes of calibration should be considered in the design phase of the spacecraft.

The results of the calibration are shown in Figure 3.5. Figure 3.5(a) shows the magnitude of the corrected measurements overlaid with the expected field magnitude, Figure 3.5(b) shows the difference between the measured and expected magnitudes, and Figure 3.5(c) is a histogram of the difference. Figures 3.5(a) and 3.5(b) show the significant improvement of the magnetometer data compared to the corresponding time-invariant calibration results shown in Figures 3.3(b) and 3.4(a). The root mean squared error (RMSE) of the measurements after calibration for time-invariant errors is 903 nT, where error is defined as the difference between the expected magnitude and the magnitude of the corrected measurements. After calibration with time-varying bias, the RMSE is reduced to 174 nT, an improvement factor of 5.2. This corresponds to an order of magnitude improvement in angular accuracy. Accuracy is discussed further in Section 3.5.5.

The resolution of the raw magnetometer measurements is 128 nT along each axis. The resolution of the corrected sensor readings in the rectified frame is transformed by the scale factors and non-orthogonality angles. After transformation of the 128 nT resolution with the calibration parameters, the resolution of the calibrated sensor is 144, 143, and 111 nT along the x, y, and z axes, respectively. This results in a 231 nT resolution on the magnitude of the corrected measurements. The RMSE of the calibrated data is only 174 nT, which is below the sensor resolution of 231 nT and indicates that the accuracy of the calibrated measurements has approached the fundamental accuracy limit of the sensor.

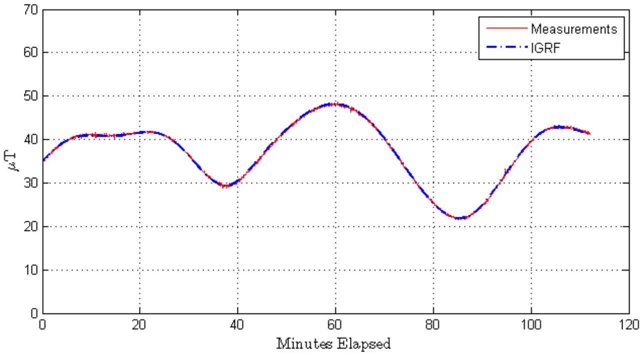

(a) The magnitude of the corrected measurements (μT) versus time (minutes). The magnitude of the expected magnetic field is overlaid.

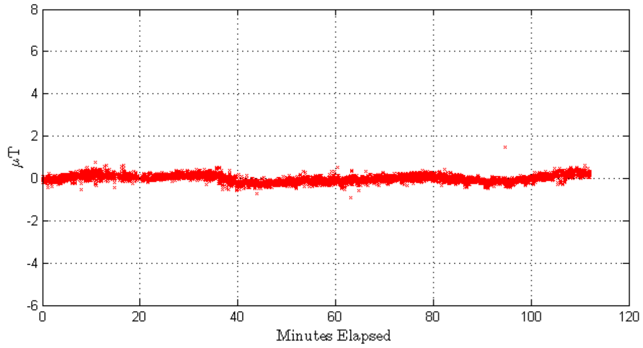

(b) Difference between the corrected measured field magnitude and the expected magnitude (μT)

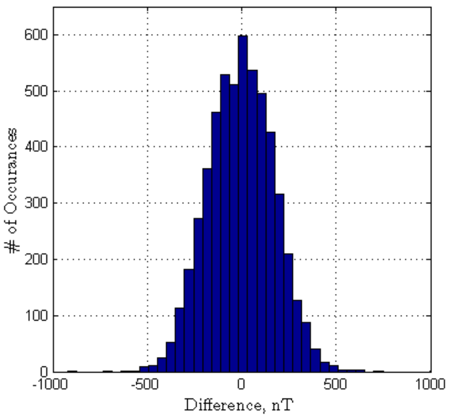

(c) Histogram of the difference between the measured and expected magnitudes (nT). There are 5,405 total data points.

Figure 3.5: Results of the calibration to estimate both constant errors and time-varying magnetometer bias.

3.5.4 Convergence

The calibration parameters are estimated using non-linear least-squares minimization as described in Section 3.4. An initial estimate of the calibration parameters is required to begin the iterative minimization, and in general, convergence to a global minimum is not guaranteed. A Monte Carlo simulation using 1000 different initial conditions was used to investigate the convergence. The initial conditions were uniformly distributed over the ranges shown in the second column of Table 3.2. The parameter estimates from each trial converged to the single local minimum shown in the third column of the table. The tolerance used in the minimization is

$$|\Delta y_l - \Delta y_{l-1}| < 1 \text{ nT}^2, \tag{3.9}$$

where l is the iteration number and

$$\Delta y_l = \frac{1}{m} \sum_{k=1}^{m} \left(B_{E,k}^2 - f(\tilde{\mathbf{B}}_k, \tilde{\mathbf{I}}_k, \mathbf{x}_l) \right) \tag{3.10}$$

is the mean difference between the expected and measured geomagnetic field magnitudes squared. In Eq. (3.10), the quantity $B_{E,k}^2 - f(\tilde{\mathbf{B}}_k, \tilde{\mathbf{I}}_k, \mathbf{x}_l)$ is defined as in Eq. (3.8), and k is the data index. Qualitatively, the tolerance of Eq. (3.9) means that the mean difference between the expected and measured geomagnetic field magnitudes squared changes by less than 1 nT^2 with each iteration. Although 1 nT is below the resolution of the measurements, the expected geomagntic field magnitude is obtained from a continuous model, so Δy can take any value. From the Monte Carlo simulations, a 1 nT^2 tolerance was found to be sufficient: the parameter estimates from each trial converged to a single local minimum within seven iterations. This demonstrates that there are no problems with convergence and suggests that a global minimum exists. A formal study of convergence is left for future work.

3.5.5 Calibration Accuracy

In this subsection, the accuracy of both the corrected magnetic field measurements and the calibration parameters is discussed.

The direction of all possible measurements taken by a three-axis vector sensor can be represented by a sphere. That is, if an ideal (perfectly calibrated, no noise) three-axis magnetometer is rotated in a constant magnetic field, the measurements, plotted in three dimensions, lie on a sphere. This sphere is referred to as the *attitude sphere*. Data from

Table 3.2: The range of uniformly distributed initial conditions used in the Monte Carlo simulations (column two) and the final estimate of the calibration parameters (column three). There are 15 total mapping coefficients, $s_{i,j}$. The range of the final coefficients is given rather than each individual parameter for simplicity.

Parameter	**Range**	**Estimate**
a	[-4, 4]	0.89
b	[-4, 4]	0.90
c	[-4, 4]	1.13
x_0 (μT)	[-20, 20]	-0.69
y_0 (μT)	[-20, 20]	9.91
z_0 (μT)	[-20, 20]	-7.70
ρ (°)	[-20, 20]	-1.04
ϕ (°)	[-20, 20]	-3.97
λ (°)	[-20, 20]	5.02
$s_{i,j}$ (μT/mA)	[-1, 1]	[-0.02, 0.01]

an uncalibrated magnetometer forms an ellipsoid [44, 46]. Estimates of the calibration parameters are most accurate if the measurements used in the calibration are distributed over the entire surface of the ellipsoid [44,46], and this is discussed further in the following paragraphs.

Figure 3.6(a) shows the components of the normalized measurements after calibration. This is the same data as Figures 3.3-3.5. For this data set, the measurements are distributed over the entire attitude sphere. The components from two other data sets, taken December 15 and December 30, 2010, are shown in Figures 3.6(b) and 3.6(c). RAX-1 utilizes a passive magnetic control system (see Section 2.2), which gradually aligns the satellite with Earth's magnetic field. This is evident in Figure 3.6. On December 1, the satellite is still tumbling after deployment from the launch vehicle. As time passes, the satellite aligns with the geomagnetic field, resulting in the reduced sphere coverage seen in Figures 3.6(b) and 3.6(c).

Nine of the 24 calibration parameters estimated from each data set are shown in Table 3.3. There is variation between the parameters, and to investigate the accuracy of the parameters, each set of measurements is corrected with each set of calibration parameters. The resulting RMSEs after calibration are given in Table 3.4. The table shows the RMSE after each set of measurements is corrected by each set of calibration parameters. For example, the RMSE of the December 1 corrected measurements using the parameters estimated from the December 1 measurements is 174 nT. The RMSE of the December 15 corrected measurements using the parameters estimated from the December 30 measurements is 490 nT. Table 3.5 is analogous to Table 3.4, but the measurements in Table 3.5 are corrected using the time-invariant method described in Section 3.3. Table 3.5 is included to demonstrate the improvement of the new calibration method relative to existing methods.

There are two significant aspects of the errors shown in Table 3.4. First, the RMSEs of the December 1 data set corrected with parameters from the December 1, 15, and 30 data sets (first row) indicate that parameter estimates from the December 1 data set are most accurate. Additionally, the fact that the RMSEs of each data set corrected with the December 1 parameters are all within one sensor resolution (first column) indicates that the parameters are constant over time.

A second approach is taken to study the differences in the parameters of Table 3.3. The accuracy of the individual calibration parameter estimates are subject to measurement noise and the level of excitation of the data. To approximate the accuracy of the individual parameter estimates, we have performed a simulation of 1000 trials in which the conditions of the three data sets were used to simulate measurements, and the calibration parameters were estimated from the simulated measurements. 1000 trials was used because it is a large enough number to produce repeatable average results. In these simulations, the actual estimates of the parameters from each data set (Table 3.3) were taken to be the true parameters, the calibrated measurements were taken to be the true magnetic field components in the rectified frame, and the current measurements were taken to be the true current values. Measurements were then simulated with these parameters and the model of Eqs. (3.5)-(3.7), and measurement noise was simulated by adding zero mean Gaussian noise to the magnetometer measurements (128 nT 1-σ) and current measurements (5 mA 1-σ)[6]. The calibration parameters were then estimated from these simulated measurements.

The mean and standard deviation of the errors in the resulting parameter estimates from each trial are shown in Figures 3.7-3.12, and the mean and standard deviation of these errors are also shown as a percentage of the true calibration parameters in Table 3.6. From the resulting mean error in the parameters estimates for each data set, we see that the accuracy of the parameters estimated from the December 1 data set is significantly better than the accuracy of the parameter estimates resulting from the December 15 and 30 data sets. For the December 1 data, the mean error in scaling, constant bias, and non-orthogonality angles is less than 1%. The relatively high accuracy of the December 1 data set compared to the others is due to the reduced coverage of the measurements for the data sets as shown in Figure 3.6. Additionally, as seen in Figures 3.10-3.12 as well as Table 3.6, the coeffi-

---

[6] Through experimental characterization, the noise in the PNI measurements was determined to be below the resolution in each axis (128 nT). Therefore, the resolution itself was used as the standard deviation of the measurement noise for the purposes of simulated testing. The noise of the current sensors was not explicitly characterized. Their resolution is 1 mA, and the 5 mA standard deviation used is assumed to be a reasonable approximation.

Table 3.3: Estimates of the calibration parameters for each data set. Nine of the 24 parameters are shown; the remaining are omitted from the table for simplicity.

Data Set	a	b	c	x_0 (μT)	y_0 (μT)	z_0 (μT)	ρ (°)	ϕ (°)	λ (°)
Dec. 1	0.890	0.910	1.130	-0.687	9.909	-7.700	-1.039	-3.974	5.019
Dec. 15	0.893	0.913	1.136	-0.581	10.014	-7.419	-1.085	-4.310	5.178
Dec. 30	0.900	0.911	1.170	-0.944	9.478	-6.220	-1.437	-4.160	4.498

cients mapping the fourth current measurement (one of the solar panels) have significantly higher error than the other current mapping coefficients. From the December 1 data, the mean error of the current mapping coefficients is less than 4%, except for the fourth current measurement, for which the error is 98% for the z-axis of the magnetometer. This is due to relative lack of variation in the fourth current measurement compared to the other current measurements, as seen in the time series of the measurements shown in Figure 3.13. This analysis shows the dependence of the parameter estimates on the coverage of the magnetometer measurements as well as the variation in the current measurements, and highlights the fact that although the calibrated magnitude of the measurements has mean zero error, the error in the individual calibration parameters is not necessarily mean zero. Nonetheless, the estimated parameters are sufficient to correct the magnetometer measurements, as is shown in Table 3.4.

A second and similar analysis of the relative accuracy of the calibration parameters is also carried out using the Fisher Information matrix. The Fisher information matrix [55] is used to approximate the lower bound on the covariance of the calibration parameters. This is shown in Appendix A. As discussed in the Appendix, the lower bound is optimistic, but it shows a decrease in the parameter estimation accuracy as the sphere coverage decreases, which matches the results of the simulated testing. A formal study on the sufficient conditions for accuracy calibration is left for future work. In practice, actively-controlled spacecraft can perform calibration maneuvers to ensure sufficient excitation of the data.

The resulting accuracy of the corrected measurements (within one sensor resolution) and the simulated estimation of the calibration parameters indicate that the calibration parameter estimates from the Dec. 1 data set are accurate. Through application to flight data, it has been shown that the calibration with time-varying bias significantly improves the accuracy of the calibration. For example, the RMSE of the December 15 data set after correction with constant parameters obtained from the December 1 data set is 792 nT. With time-varying calibration, this is improved to 219 nT, an angular improvement[7] from 1.3° to 0.36° in an ambient field of 35μT.

---

[7] To approximate the angular accuracy, we assume the error is orthogonal to the magnetic field vector.

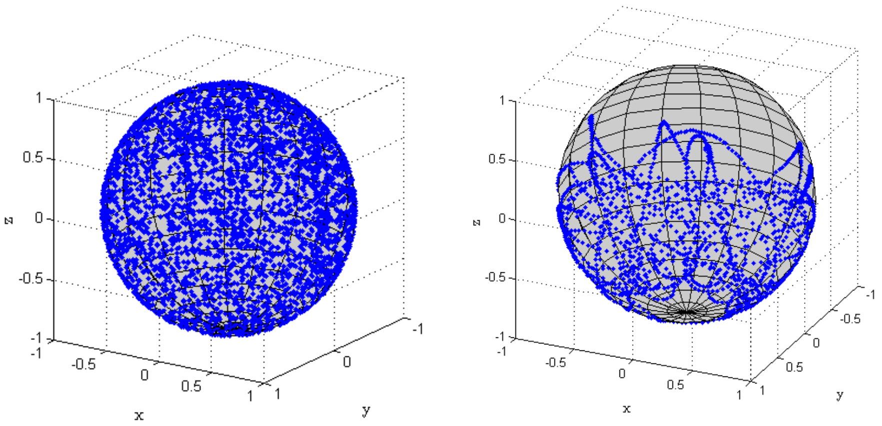

(a) December 1, 2010 data set (b) December 15, 2010 data set

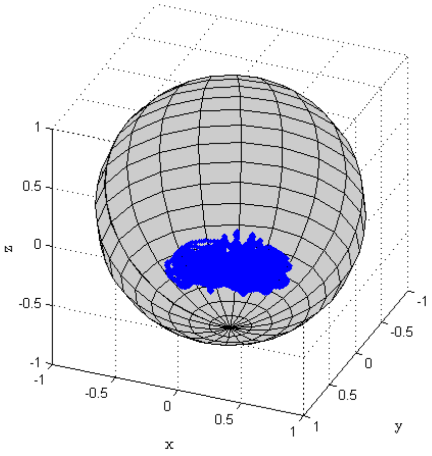

(c) December 30, 2010 data set

Figure 3.6: The normalized vector components of the corrected measurements plotted in three dimensions on top of a unit sphere in the body frame for each data set, which shows the coverage of the *attitude sphere*.

Table 3.4: The RMSE (nT) of the corrected measurements after calibration using the parameter estimates from each data set. Measurements from the data sets listed in the left column are corrected using parameters from the data sets listed across the top.

		Calibration Parameters Used		
		Dec. 1	Dec. 15	Dec. 30
Corrected Measurements	Dec. 1	174	320	1,074
	Dec. 15	219	160	490
	Dec. 30	210	200	153

Table 3.5: The RMSE (nT) of the corrected measurements after time-invariant calibration using the parameter estimates from each data set. This is analagous to Table 3.4, but the calibration method is the existing time-invariant method described in Section 3.3. Measurements from the data sets listed in the left column are corrected using parameters from the data sets listed across the top.

		Calibration Parameters Used		
		Dec. 1	Dec. 15	Dec. 30
Corrected Measurements	Dec. 1	903	2,064	3,351
	Dec. 15	792	725	1,266
	Dec. 30	677	825	391

Table 3.6: Mean and standard deviation of the error in the estimates of scaling, bias, and non-orthogonality resulting from the 1000 simulated trials. The mean and standard deviations are given as a percentage of the true parameter.

Parameter	**Mean (%)**			**Standard Deviation (%)**		
	Dec. 1	Dec. 15	Dec. 30	Dec. 1	Dec. 15	Dec. 30
a	0.01	0.01	0.12	0.02	0.98	0.21
b	0.01	0.05	0.01	0.02	0.04	0.02
c	0.00	0.12	0.01	0.01	0.12	0.03
x_0	0.04	0.44	1.08	0.76	7.66	6.18
y_0	0.02	0.08	0.07	0.05	0.11	0.34
z_0	0.04	0.31	0.62	0.09	0.39	0.51
ρ	0.73	15.75	173.66	1.39	535.22	1389.03
ϕ	0.09	0.70	8.26	0.34	6.63	246.13
λ	0.03	1.35	2.37	0.28	1.21	3.07
$s_{1,x}$	0.77	0.93	7.34	0.47	3.20	2.50
$s_{2,x}$	0.34	6.70	19.15	1.56	97.36	23.33
$s_{3,x}$	1.44	7.68	17.53	3.53	49.96	26.95
$s_{4,x}$	15.55	85.19	99.12	7.05	22.43	3.12
$s_{5,x}$	3.11	0.82	5.81	21.41	52.78	76.67
$s_{1,y}$	1.38	2.66	5.24	0.60	1.98	5.88
$s_{2,y}$	0.32	1.43	12.13	1.82	5.69	25.30
$s_{3,y}$	0.73	5.98	7.22	0.70	5.82	13.93
$s_{4,y}$	34.80	142.88	103.42	4.91	203.05	82.06
$s_{5,y}$	2.98	4.55	5.81	18.72	42.00	227.27
$s_{1,z}$	0.70	1.50	2.52	0.33	0.96	2.63
$s_{2,z}$	1.64	15.95	123.48	5.76	55.38	129.42
$s_{3,z}$	0.70	2.65	0.99	7.18	16.11	45.08
$s_{4,z}$	98.10	91.88	98.10	31.51	38.55	8.32
$s_{5,z}$	0.23	13.19	2.79	12.85	47.37	70.05

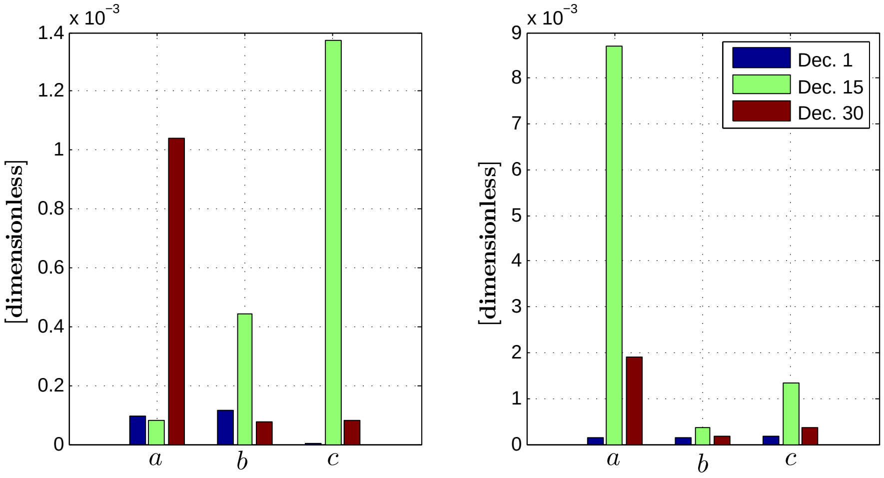

(a) Mean (b) Standard deviation

Figure 3.7: Mean and standard deviation of the error in scale factor estimates from the 1000 simulated trials. The true parameters are shown in Table 3.3, and the mean errors and standard deviations as a percentage of the true parameters are given in Table 3.6.

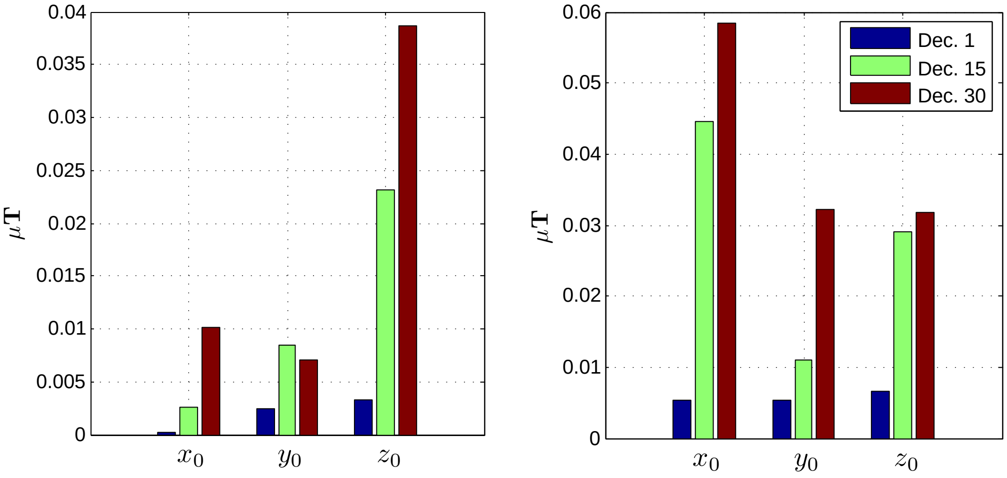

(a) Mean (b) Standard deviation

Figure 3.8: Mean and standard deviation of the error in constant bias estimates from the 1000 simulated trials. The true parameters are shown in Table 3.3, and the mean errors and standard deviations as a percentage of the true parameters are given in Table 3.6.

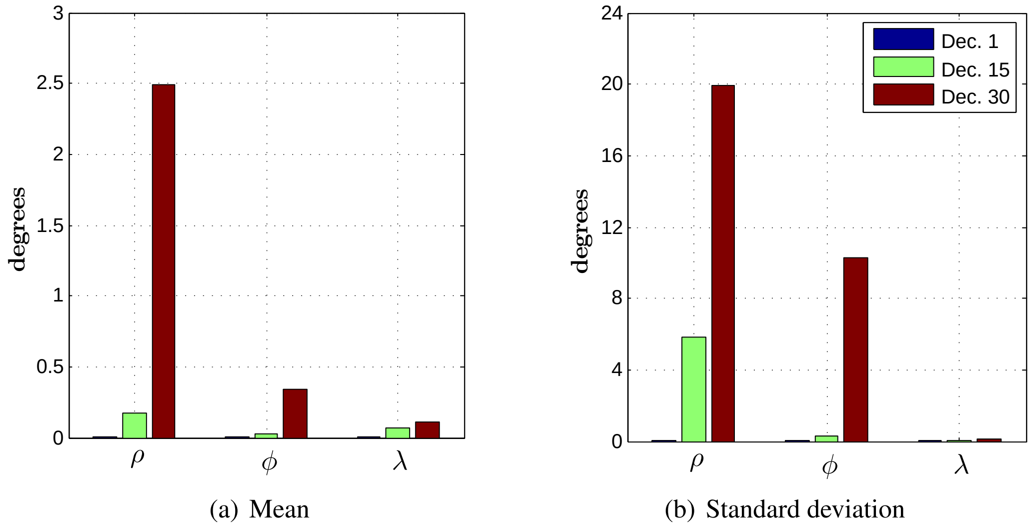

(a) Mean

(b) Standard deviation

Figure 3.9: Mean and standard deviation of the error in non-orthogonality angle estimates from the 1000 simulated trials. The true parameters are shown in Table 3.3, and the mean errors and standard deviations as a percentage of the true parameters are given in Table 3.6.

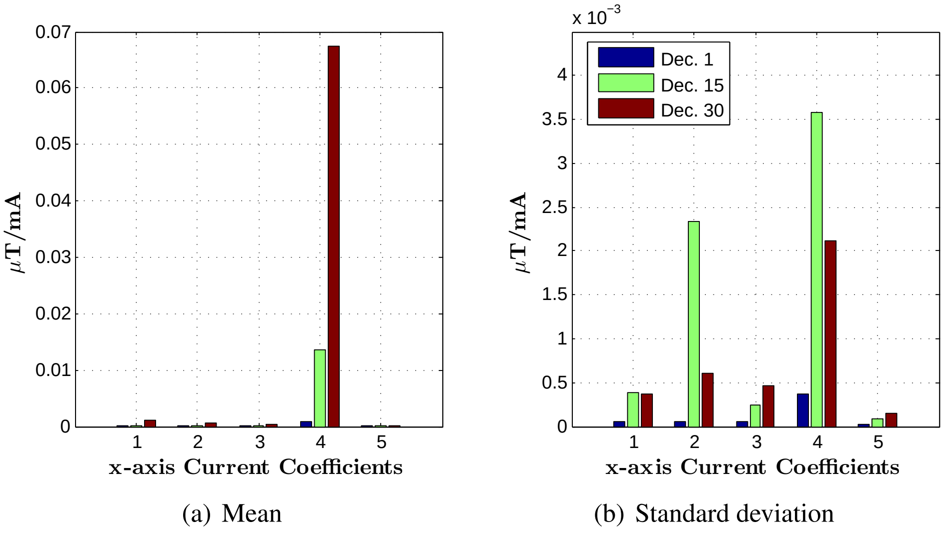

(a) Mean

(b) Standard deviation

Figure 3.10: Mean and standard deviation of the error in the current coefficients estimates for the magnetometer $\tilde{x}$-axis from the 1000 simulated trials. The mean error and standard deviations as a percentage of the true parameters are given in Table 3.6.

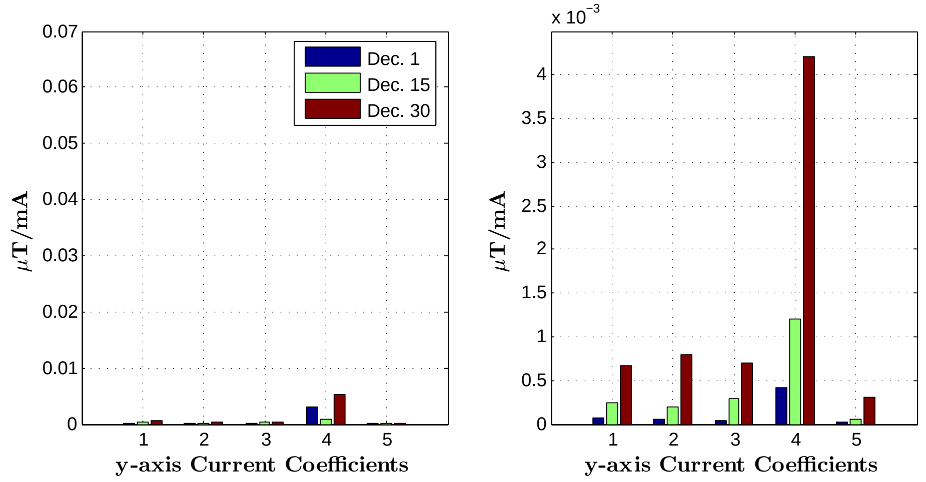

(a) Mean (b) Standard deviation

Figure 3.11: Mean and standard deviation of the error in the current coefficients estimates for the magnetometer $\tilde{y}$-axis from the 1000 simulated trials. The mean error and standard deviations as a percentage of the true parameters are given in Table 3.6.

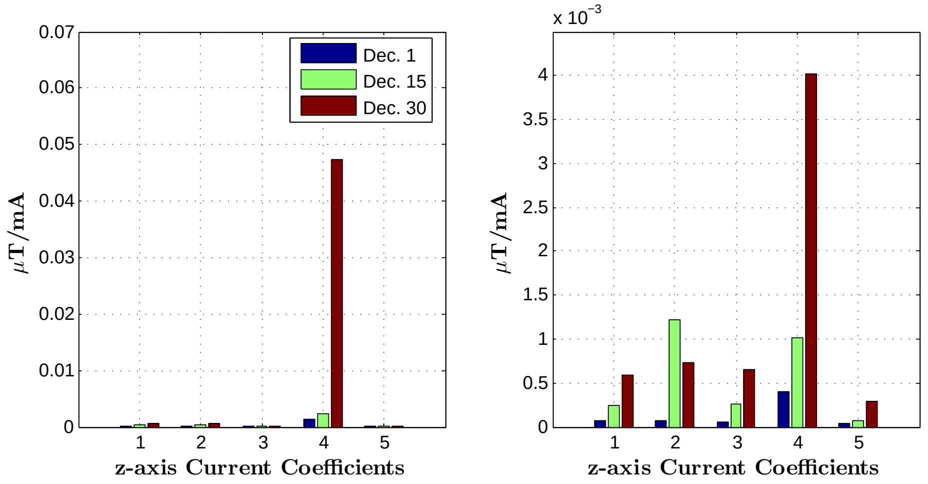

(a) Mean (b) Standard deviation

Figure 3.12: Mean and standard deviation of the error in the current coefficients estimates for the magnetometer $\tilde{z}$-axis from the 1000 simulated trials. The mean error and standard deviations as a percentage of the true parameters are given in Table 3.6.

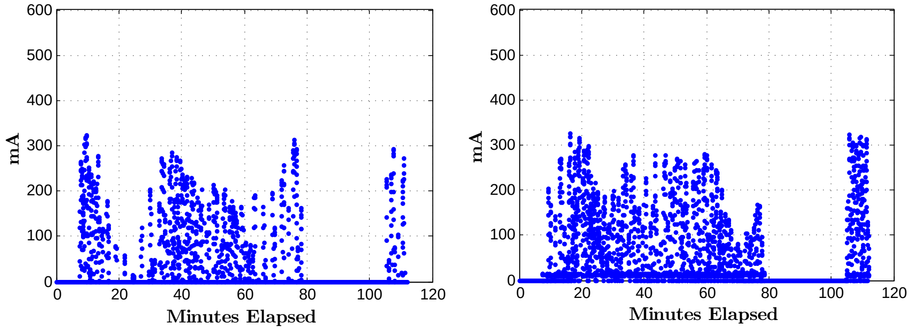

(a) Current 1 (b) Current 2

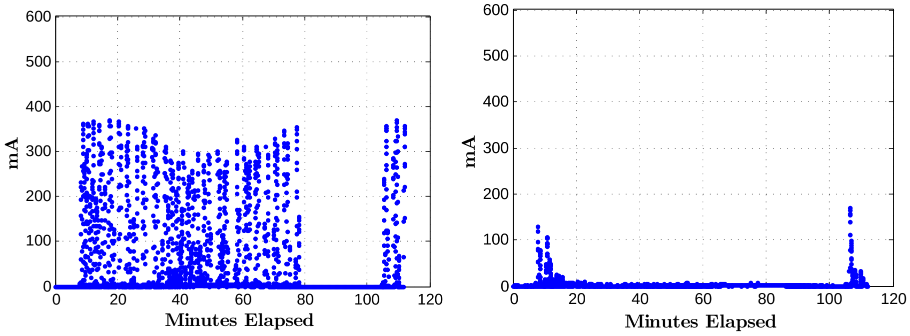

(c) Current 3 (d) Current 4

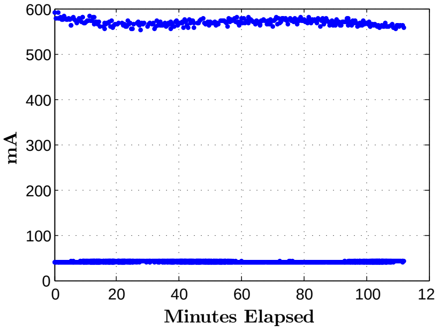

(e) Current 5

Figure 3.13: Measurements of current from the December 1 data set that were used in the magnetometer calibration. Currents 1-4 are solar panels, and the fifth is current draw from the battery.

3.6 Conclusions

A method for on-orbit, attitude-independent magnetometer calibration that includes time-varying bias due to nearby electronics has been developed. Existing time-invariant calibration methods have been extended by including current measurements in the sensor model, and the resulting calibration method estimates both time-invariant errors and time-varying bias due to on-board electronics. The time-varying bias is estimated by mapping current to magnetometer bias through constant parameters. The calibration parameters are estimated using an iterative least-squares minimization, and only the magnetometer measurements, current measurements, and the expected magnitude of the geomagnetic field are required for the calibration.

The utility of the calibration has been demonstrated by application to on-orbit data. In application to the RAX-1 satellite, the calibration successfully mitigates magnetometer bias due to the satellite power system. The calibration improved the RMSE of magnetometer measurements from 792 nT to 219 nT, corresponding to an order of magnitude increase in angular accuracy.

Traditionally, time-varying bias is mitigated by either using a boom to physically separate the magnetometer from the spacecraft, or by using design and manufacturing practices to minimize the influence of electronic components on magnetometers. Such design practices increase design time and cost, and physical separation of a magnetometer from on-board electronics may not be possible as satellites continue to decrease in size. The algorithm presented in this work effectively replaces such design practices with improved processing of the sensor measurements. The algorithm utilizes current sensors throughout the spacecraft that are sampled at the same time as the magnetometers, so inclusion of current sensors and the ability to sample the sensors simultaneously for the purpose of calibration should be considered in the design phase of the vehicle. Although the calibration is a batch method, parameters can be uploaded to the spacecraft for real-time magnetometer correction, and real-time implementation will be developed in future work. The calibration has been applied to satellite-based magnetometers in this work, but this algorithm is applicable to other magnetometer applications on a variety of platforms.

CHAPTER 4

Directional Sensor Optimization

In this chapter, we present a method to optimize the orientation of directional sensors. This is applicable to sensors or instruments with either a limited field of view (FOV) or an uncertainty that varies over the sensor FOV. Examples of these types of sensors and instruments include sun sensors, solar panels, and antennas. In application to a low-cost attitude determination system, the method is used to optimize the configuration of photodiodes used for sun sensing. Whereas the magnetometer calibration of the previous chapter enables magnetometers to be placed anywhere within a satellite and the orientation of tri-axial sensors does not affect their accuracy, this chapter provides an technique to optimize the orientation of photodiodes during the design phase of spacecraft development.

4.1 Introduction

Photodiodes, also referred to as cosine detectors [58], are a common method of sun sensing on small spacecraft because of their simplicity and low cost (for example, [31, 59–62]). Multiple photodiodes can be combined to estimate the line-of-sight vector to the sun, which is subsequently used for attitude determination or instrument pointing. The uncertainty of the estimated sun vector depends on both the performance of the individual photodiodes as well as the orientation of each sensor. In this chapter, the covariance of the sun vector estimate is derived as a function of the photodiode configuration, and the orientation of the sensors is optimized to minimize the angular uncertainty of the resulting sun vector estimate.

The traditional spacecraft design process relies on iterating to satisfy design requirements and constraints [63]. In the design process, sensor placement is typically performed manually based on prior experience and similarity to previous designs [64]. Although optimization is becoming more prevalent in spacecraft design [65, 66], there are very few methods for the optimization of directional sensors in the literature. In the only paper found on optimal sun sensor configurations, Jackson and Carpenter optimize the orientation of sun sensors by parameterizing the orientation as placement of the sensors on a spherical spacecraft and using genetic algorithms and simulated annealing to maximize coverage of the sensors over the spacecraft body frame [64]. In this work, the angular uncertainty of the sun vector estimates are minimized subject to the orientation of the sun sensors, rather than maximizing the sun sensor coverage, which results in the best attitude determination accuracy with the given hardware. The formulation is general and can be used for the optimization of various directional sensors or instruments in the body-fixed frame.

The orientation optimization problem is formulated using a unit sphere. With the spacecraft located at the center, the surface of the sphere represents all directions in the body-fixed coordinate system. This sphere is commonly referred to as the *attitude sphere* or *spacecraft-centered celestial sphere*. Directions in the body-fixed frame are created by discretizing the attitude sphere, and these directions are used in the objective function. Directions on the attitude sphere are weighted to account for vehicle- and mission-specific parameters such as orbital inclination and attitude. The nature of the objective function is problem-specific and can be minimized using a suitable algorithm. The unique contributions of this chapter are the formulation for optimization that uses the attitude sphere as well as the application to photodiodes to minimize the uncertainty of the measured sun vector. This new formulation enables the application of existing minimization techniques to a broad range of directional-dependent problems.

The remainder of this chapter is organized as follows. In Section 4.2, the attitude sphere is discretized to create the directions used to formulate the objective function. The objective function for the sun sensor optimization problem is derived in Section 4.3 and applied to example design problems in Section 4.4. The method is then summarized for general application in Section 4.5 and conclusions are given in Section 4.6.

4.2 Use of the Attitude Sphere

A unit sphere can be used to define all possible directions from the spacecraft to an object of interest in the body-fixed frame. This is referred to as the attitude sphere. This sphere is also commonly referred to as the spacecraft-centered celestial sphere [67]. The attitude sphere is discritized, or pixelized, to create directions over the entire body frame for use in formulating the objective function. Points on the sphere represent the end of unit vectors originating at the center of the sphere.

A pixelization with constant angular resolution over the entire attitude sphere should be used so that portions of the sphere are not artificially weighted due to a higher density of points in the region. Because of this, the typical method of varying the azimuth and elevation evenly to create points on a sphere is not used; it results in a distribution with higher density near the poles as seen in Figure 4.1(a). Uniform pixelization of the surface of a sphere is not a new problem. It has been used for mapping the celestial sphere for astronomical applications [68] as well as mapping Earth for remote sensing applications [69]. Attitude sphere pixelization has also been used to formulate path-planning for spacecraft attitude control [70].

To pixelate the attitude sphere, the icosahedron-based approach of Reference 68[1] is used because it produces a near uniform distribution of points on the sphere as seen in Figure 4.1(c). The pixelization method begins by inscribing an icosahedron inside a unit sphere. Points are then distributed evenly on each triangular face of the icosahedron. This is shown in Figure 4.2 [68]. Each point in Figure 4.2(b) is the center of the hexagonal pixel. The pixel centers are then projected from the faces of the icosahedron onto the sphere and shifted slightly to give all pixels approximately equal area.

It is not possible to place an arbitrary number of pixels on the sphere while maintaining the near uniform distribution. The number of pixels on the sphere is given by

$$N = 40r(r-1) + 12 \ , \tag{4.1}$$

where $r = 1, 2, 3, \ldots$ is the resolution. A plot of the number of pixels N as well as the average angular distance between pixel centers is shown in Figure 4.3.

4.3 Formulating the Objective Function for Sun Sensor Optimization

An ideal photodiode produces current I as a function of incoming light according to the model [71]

$$I = I_0 \cos(\theta) \ , \tag{4.2}$$

where θ is the angle between the direction normal to the photosensitive plane and the line-of-sight vector to the sun (herein referred to as *sun vector*), and I_0 is the maximum current output of the sensor, corresponding to $\theta = 0$.

Multiple photodiodes can be combined to provide an estimate of the sun vector. The uncertainty of the sun vector estimate is a function of the performance of the individual

---

[1] FORTRAN code for implementation is available from the author: `http://space.mit.edu/home/tegmark/icosahedron.html`, accessed September 12, 2012.

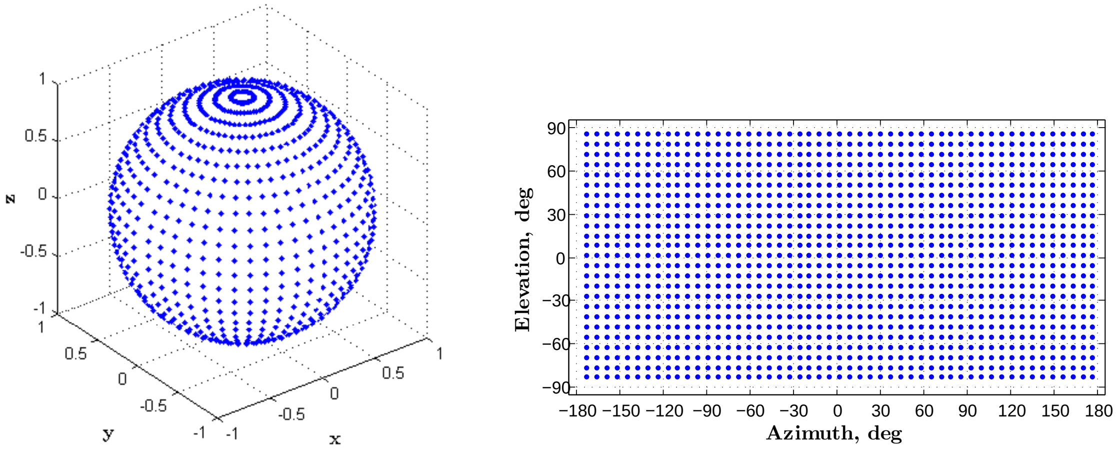

(a) Unit sphere with points genearated using constant azimuth and elevation increments.

(b) 2D projection of Figure 4.1(a).

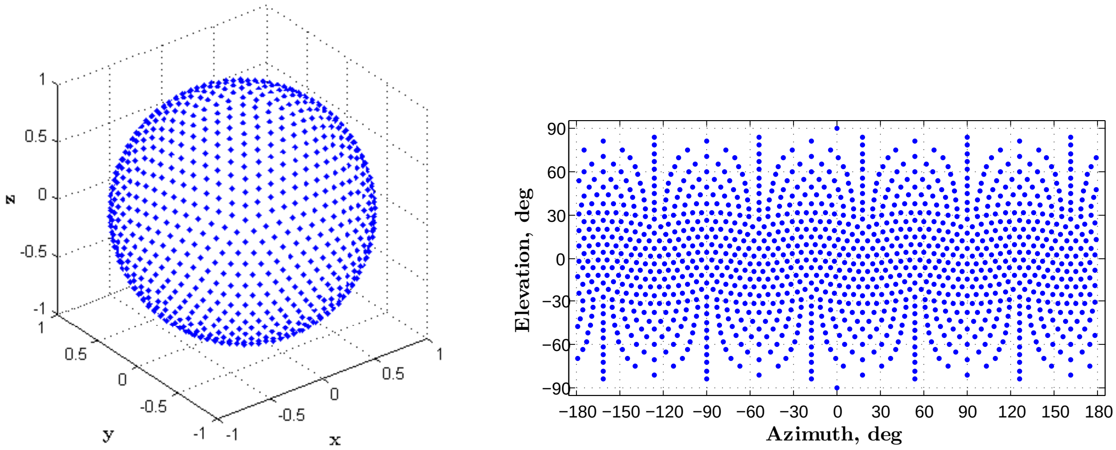

(c) Unit sphere with points evenly distributed over the surface using the icosahedron approach.

(d) 2D projection of Figure 4.1(c).

Figure 4.1: Unit spheres with two different methods of distributing points shown with their 2D projections.

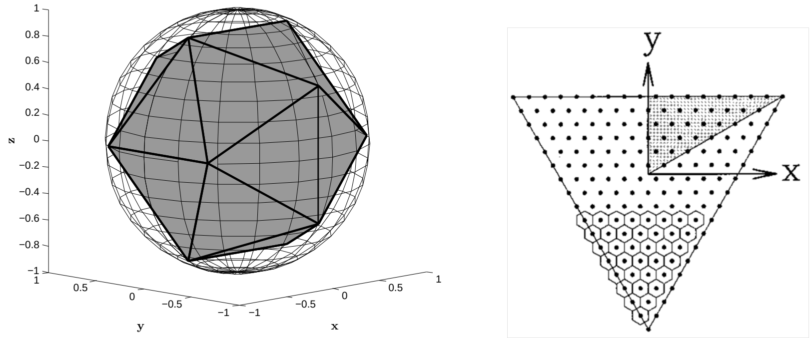

(a) An icosahedron with a unit sphere overlaid.

(b) The triangular icosahedron faces are pixelated with a regular triangular grid. The points shown are the centers of the hexagonal pixels.

Figure 4.2: The first two steps of the icosahedron-based approach to sphere pixelization are shown.

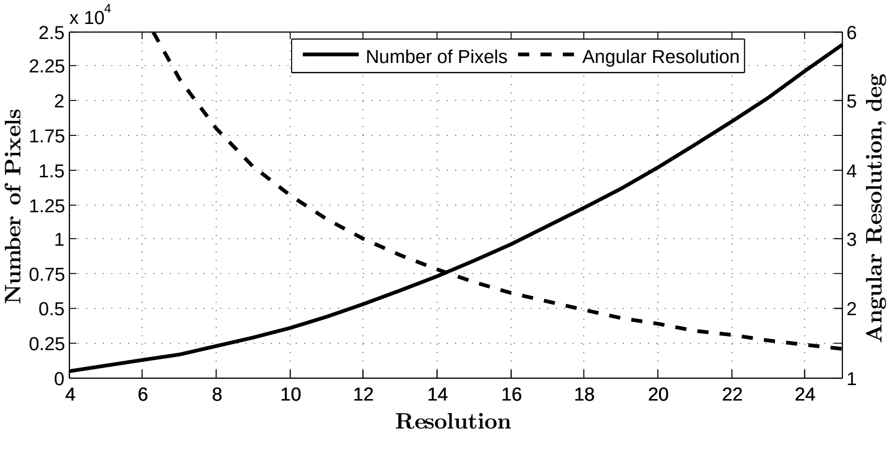

Figure 4.3: Number of pixels N and corresponding angular resolution versus the resolution r.

photodiodes as well as the photodiode configuration. In this section, the objective function to minimize the angular uncertainty of the sun vector estimate is formulated as a function of the photodiode orientation. In Section 4.3.1, the covariance of the sun vector estimate is derived, and this is used with the directions on the attitude sphere to formulate the objective function in Section 4.3.2.

4.3.1 Sun Vector Estimation from Multiple Photodiodes

A single photodiode provides one dimension of sun vector information; multiple photodiodes can be combined to provide a full sun vector. Two common photodiode configurations are to use six photodiodes that are orthogonal [31, 59], or to place pairs of photodiodes angled in a single plane [58]. The former method does not provide photodiode coverage sufficient for sun vector estimation over the entire attitude sphere because the field of view of individual photodiodes is generally less than 180° (this was illustrated in Section 2.5). In the latter configuration, the two photodiode readings can be used to estimate the sun vector component in the common plane of the photodiodes, thus achieving a one-axis sun measurement [58]. For unconstrained optimization, a method to estimate the sun vector from an arbitrary photodiode configuration is required, which is derived in this section.

The sun vector estimation is formulated as finding the intersection of multiple planes, where the planes are defined by the photodiode normal directions and the measured currents. Intersection of planes is used rather than intersection of cones to estimate the sun vector from the photodiode measurements because the former is a linear function of the photodiode measurements. The formulation is illustrated for a two-dimensional case in Figure 4.4. In the figure, the rectangles represent two photodiodes, the dashed arrows show the directions normal to the photosensitive plane, $\hat{\boldsymbol{n}}_i$ (where $i \in \{1, 2\}$ is the photodiode index), the solid arrow shows the sun vector, $\hat{\boldsymbol{s}}$, and θ_i is the angle between the photodiode normal directions and the sun vector, as in Eq. (4.2). The sun vector $\hat{\boldsymbol{s}}$ corresponds to the intersection point of the two planes, shown by the dotted lines, and $\hat{\boldsymbol{s}}$ is estimated by finding the intersection point of the two planes. The planes are perpendicular to the photodiode normal direction, and, referring to Eq. (4.2), their location along the normal direction $\hat{\boldsymbol{n}}_i$ is $I_i/I_{0,i}$.

In general, a plane can be defined by a known point on the plane, $\boldsymbol{r}_0$, and a normal vector to the plane, $\boldsymbol{p}$, according to Eq. (4.3), where $\boldsymbol{r}$ is the location of any point on the plane.

$$\boldsymbol{p}^T (\boldsymbol{r} - \boldsymbol{r}_0) = 0 \tag{4.3}$$

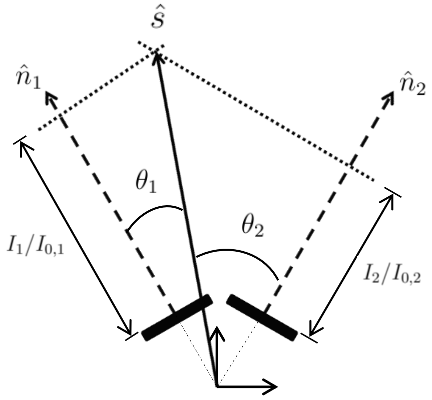

Figure 4.4: An illustration of two photodiodes in a single plane. The solid rectangles represent the two photodiodes, and their normal directions are shown by the unit vectors $\hat{\boldsymbol{n}}_i, i \in \{1, 2\}$. The angles θ_i define the sun vector direction relative to the photodiode normal directions, as in Eq. (4.2), and the sun vector $\hat{\boldsymbol{s}}$ corresponds to the intersection of the two planes. The planes are represented by the dotted lines, are perpendicular to $\hat{\boldsymbol{n}}_i$, and their location along $\hat{\boldsymbol{n}}_i$ is given by $I_i/I_{0,i}$

In application to the photodiode configuration of Figure 4.4, Eq. (4.3) becomes

$$\hat{\boldsymbol{n}}_i^T \left(\hat{\boldsymbol{s}} - \frac{I_i}{I_{0,i}} \hat{\boldsymbol{n}}_i \right) = 0 \; , \tag{4.4}$$

With $\hat{\boldsymbol{n}}_i = \begin{bmatrix} n_{i,1} & n_{i,2} & n_{i,3} \end{bmatrix}^T$ and $\hat{\boldsymbol{s}} = \begin{bmatrix} s_1 & s_2 & s_3 \end{bmatrix}$, and using the fact that $\|\hat{\boldsymbol{n}}\| = 1$, Eq. (4.4) can be re-written as

$$\frac{I_i}{I_{0,i}} = \begin{bmatrix} n_{i,1} & n_{i,2} & n_{i,3} \end{bmatrix} \begin{bmatrix} s_1 \\ s_2 \\ s_3 \end{bmatrix} \; . \tag{4.5}$$

Eq. (4.5) can also be obtained directly by rewriting Eq. (4.2) as $I = I_0 \cos(\theta) = I_0 \hat{\boldsymbol{n}}^T \hat{\boldsymbol{s}}$, but derivation with Figure 4.4 is included to provide further insight into the vector measurement formulation. For sun sensing in three dimensions, at least three non-parallel and non-coplanar photodiodes are required. With photodiodes satisfying these conditions, the sun vector corresponds to the intersection point of the planes from each photodiode.

In practice, a single intersection point generally does not exist due to measurement error. The ideal model of Eq. (4.2) is corrupted by sensor noise, and the measurement model is given by Eq. (4.6), where $\tilde{I}$ is the measured current, and η is zero-mean measurement noise [2].

$$\tilde{I} = I_0 \cos(\theta) + \eta \tag{4.6}$$

Given at least 3 non-parallel and non-coplanar photodiodes, the sun vector estimate is the solution to

$$\boldsymbol{y} = H\boldsymbol{s} + \boldsymbol{\eta}\ , \tag{4.7}$$

where, from Eq. (4.5),

$$\boldsymbol{y} = \begin{bmatrix} \frac{\tilde{I}_1}{I_{0,1}} \\ \frac{\tilde{I}_2}{I_{0,2}} \\ \vdots \\ \frac{\tilde{I}_k}{I_{0,k}} \end{bmatrix}\ ,\ H = \begin{bmatrix} n_{1,1} & n_{1,2} & n_{1,3} \\ n_{2,1} & n_{2,2} & n_{2,3} \\ \vdots & \vdots & \vdots \\ n_{k,1} & n_{k,2} & n_{k,3} \end{bmatrix}\ ,\ \boldsymbol{\eta} = \begin{bmatrix} \frac{\eta_1}{I_{0,1}} \\ \frac{\eta_2}{I_{0,2}} \\ \vdots \\ \frac{\eta_k}{I_{0,k}} \end{bmatrix} \tag{4.8}$$

and k is the number of photodiodes illuminated. This is the classic linear least-squares problem. Using maximum likelihood estimation, the best estimate of $\boldsymbol{s}$ is found by minimizing the objective function given by Eq. (4.9) with measurement covariance $R = \mathbb{E}\left(\boldsymbol{\eta}\,\boldsymbol{\eta}^T\right)$[3]. From the Gauss-Markov Theorem, the unconstrained optimal estimate is given by Eq. (4.10), and the covariance of the estimate is given by Eq. (4.11) [72].

$$J(\boldsymbol{s}_{est}) = \frac{1}{2}\left[\boldsymbol{y} - H\boldsymbol{s}_{est}\right]^T R^{-1}\left[\boldsymbol{y} - H\boldsymbol{s}_{est}\right] \tag{4.9}$$

$$\boldsymbol{s}_{est} = \left(H^T R^{-1} H\right)^{-1} H^T R^{-1}\boldsymbol{y} \tag{4.10}$$

$$P = \left(H^T R^{-1} H\right)^{-1} \tag{4.11}$$

The sun vector estimate resulting from the linear least-squares solution of Eq. (4.10) is not constrained to be a unit vector, and the covariance P is full-rank. For an optimal unit vector estimate, the objective function $J(\boldsymbol{s}_{est})$ should be minimized with the nonlinear constraint $\boldsymbol{s}_{est}^T\boldsymbol{s}_{est} = 1$. The constrained minimization can be carried out with various numerical techniques. Alternatively, a unit vector estimate can be obtained through brute-force normalization of Eq. (4.10), but while the angular difference between this and the optimal unit vector estimate may be negligible for a given photodiode configuration, the difference can be on order of tenths of degrees depending on the photodiode configuration.

---

[2] Photodiodes in LEO are also affected by Earth albedo, and the measurement model that captures this is presented in Chapter 5. Since the albedo contribution does not affect the optimization, it is left off here for simplicity.

[3] It is common to assume the uncertainty in each photodiode has a Gaussian distribution, and this assumption is used in the simulations presented in this paper. In practice, characteristics of the sensor should be measured by calibration.

Therefore, constrained minimization of the objective function should be used[4].

The covariance P is that of the unconstrained solution. The covariance of a unit vector is necessarily rank-deficient and can be difficult to calculate. A covariance approximation commonly used for attitude determination sensors is the QUEST measurement model, for which the covariance of a measured unit vector $\hat{\boldsymbol{u}}$ is given by

$$R_{QMM} = \sigma^2 \left(I_{3\times3} - \hat{\boldsymbol{u}}\hat{\boldsymbol{u}}^T \right) \ , \tag{4.12}$$

where σ^2 is the variance of the components normal to the estimated direction [73]. For tri-axial sensors, an assumption of uniform uncertainty in all directions is typically valid, so σ^2 can be taken to be the variance in one of the sensor axes. This approximation is not applicable to photodiodes in arbitrary configurations because the uncertainty in each direction is not equal and depends on the photodiode configuration. This is quantified by the vector covariance P of Eq. (4.11).

The covariance of the optimal unit vector estimate can be approximated by the linearized transformation of the unconstrained covariance that is equivalent to brute-force normalization of the unconstrained sun vector estimate. That is, the covariance of the unit vector estimate, P', is given by Eq. (4.13), where $f(\boldsymbol{x})$ is given by Eq. (4.14) and the subscript $\boldsymbol{x} = \hat{\boldsymbol{s}}_{est}$ indicates that the Jacobian is evaluated at the optimal unit vector estimate.

$$P' = \left. \frac{\partial f}{\partial \boldsymbol{x}} \right|^T_{\boldsymbol{x}=\hat{\boldsymbol{s}}_{est}} P \left. \frac{\partial f}{\partial \boldsymbol{x}} \right|_{\boldsymbol{x}=\hat{\boldsymbol{s}}_{est}} \tag{4.13}$$

$$\hat{\boldsymbol{x}} = f(\boldsymbol{x}) \equiv \frac{\boldsymbol{x}}{\sqrt{\boldsymbol{x}^T\boldsymbol{x}}} \tag{4.14}$$

4.3.2 Formulation of the Objective Function to Minimize Angular Uncertainty

From Eqs. (4.11) and (4.13), the sun vector uncertainty is a function of the sensor noise, manifested in R, and the orientation of each illuminated photodiode, manifested in H. The goal of the optimization is to determine the photodiode normal directions that minimize the unit sun vector covariance, P'. The photodiode normal directions (originating at the center of the attitude sphere, which coincides with the center of the body-fixed frame) are

---

[4] One drawback of the constrained minimization is that any errors in the photodiode parameters (their orientation and maximum current output) will not be evident. When working with flight data, the magnitude of the unconstrained sun vector estimate is a useful verification of the calibration parameters. Even the unconstrained estimate should yield a vector with unit magnitude on average, and a consistently biased vector magnitude indicates that sensor calibration is required. This sensor calibration is the topic of Chapter 5.

parametrized by their azimuth, α, and elevation, β, in the spacecraft body-fixed frame, and it is assumed the measurement covariance is known. To optimize the configuration over the entire attitude sphere, the sum of the weighted covariances of each direction on the attitude sphere is minimized. The objective function is

$$J = \sum_{j=1}^{N} W_j \mathrm{tr}\Big(P'_j(\boldsymbol{\alpha}, \boldsymbol{\beta}\,)\Big) \,, \tag{4.15}$$

where j is the index of the direction on the attitude sphere, N is the total number of directions on the sphere (Eq. (4.1)), W_j is the weight of the j-th direction, P'_j is the unit sun vector covariance matrix given by Eq. (4.13) when the sun is in the j-th direction, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are column vectors of the azimuths and elevations of the photodiode normal directions, and $\mathrm{tr}\,(\cdot)$ denotes the trace of a matrix. W_j is discussed in the examples of Section 4.4. P'_j depends on which photodiodes are illuminated for the j-th sun vector direction; only the rows of H and R that correspond to illuminated photodiodes are used in the calculations.

4.4 Application

In this section, the formulation of Section 4.3 is applied to optimize the photodiode configuration in design examples. This demonstrates the optimization method and shows its utility in providing a design technique to maximize performance of sensors. In Section 4.4.1, the utility of weighting in the objective function is discussed, and the optimization technique is applied to two design examples in Section 4.4.2.

4.4.1 Weighting the Attitude Sphere

The directions on the attitude sphere have been used to formulate the objective function. In some applications, certain directions are more important than others. For example, Earth-facing regions of the attitude sphere for an Earth-pointing spacecraft will never contain a sun vector direction during nominal operations, thus those directions would carry less weight in the objective function. To account for this in the optimization, unique weights can be assigned to each direction. These weights are given by W_j in Eq. (4.15). For spacecraft, the weighting is a function of the expected orbital parameters and spacecraft attitude.

As an example application, consider a a nadir-pointing spacecraft in a 400 km, 90° inclination circular orbit. The spacecraft attitude is such that the body z-axis is aligned with the nadir direction and the rotation rate is one degree per second in the z-axis. The relative

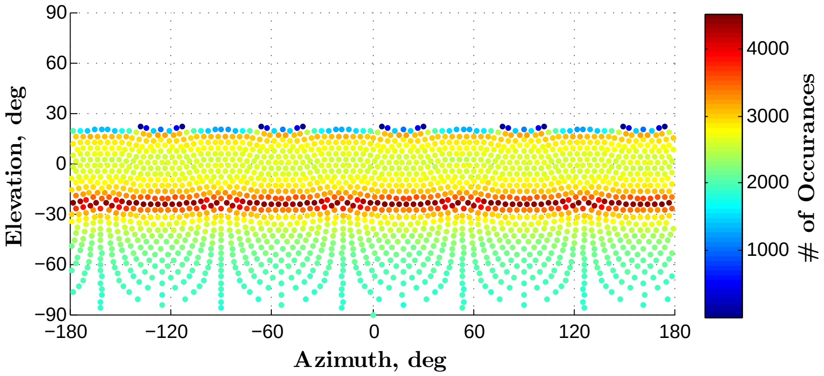

Figure 4.5: The directions on the attitude sphere color-coded by the number of occurrences of the sun vector for the nadir-pointing spacecraft example of Section 4.4.4.4.1. Directions with zero occurrences are not shown.

importance of the j-th direction on the attitude sphere for sun sensor optimization can be quantified by the amount of time that the sun vector is expected to be in the j-th direction. To determine these times, the spacecraft was simulated for a one year mission using AGI's Satellite Toolkit (STK)[5]. In the simulation, the sun vector resolved in the spacecraft body frame was polled once every four seconds, and after eliminating data points in eclipse, this results in 5,343,459 sun vectors which are subsequently binned to the nearest direction on the attitude sphere.

Figure 4.5 shows the number of occurrences of the sun vector at every direction on the attitude sphere (directions with zero occurrences are not shown). In this example, 2892 directions on the attitude sphere are used for an average angular resolution of four degrees. The simulated spacecraft spin rate and sun vector sampling frequency results in the sun vector being polled after approximately four degrees of spacecraft rotation, which is sufficient for the angular resolution of the attitude sphere. In photodiode orientation optimization for a mission with this orbit and attitude, the relative weighting shown in Figure 4.5 can be used to improve the accuracy of the sun vector estimates for the most common sun vector directions expected during the mission. The most common sun vector direction for this scheme is at an elevation of -25° in the body-fixed frame, and regions above an elevation of 30° have zero sun vector occurrences.

---

[5] `http://www.agi.com/products/by-product-type/applications/stk/`, accessed May 2012.

4.4.2 Design Examples

With identically-weighted directions on the attitude sphere, there are an infinite number of optimal configurations for a given number of photodiodes. In practice, the configuration will be constrained by factors such as a lack of available surface area on the spacecraft to place sensors (due to solar cells, instruments, etc) or orientations that are not useful due to shadowing by other spacecraft components. These factors should be used to provide an initial feasible configuration to be optimized. In this section, two examples demonstrating the procedures and potential utility of the optimization method are provide, and remarks on global versus local solutions are given at the end of the section. In both examples, the photodiodes have a conical field of view of 140°, a maximum current output of 1, and a measurement uncertainty with a standard deviation of 0.02. These parameters are representative of sensors used by the authors on past missions.

4.4.2.1 Example 1

Consider a cubical spacecraft on which, given the constraints of the telemetry system, it is possible to place two photodiodes on each face of the cube. This results in 24 design variables (azimuth and elevation for each sensor). As a first step in the optimization process, the number of design variables is first reduced to explore the solution space and determine an initial condition for optimization that meets design constraints. The problem is reduced to a single design variable by using the geometry of Figure 4.6, where the elevation of each sensor relative to the surface to which it is mounted (γ) is the single design variable. The objective value as a function of the elevation angle is shown in Figure 4.7(a). The elevation angle has been varied over the entire range which maintains photodiode coverage sufficient for sun vector estimation over the entire attitude sphere (4π steradian coverage). Two weighting schemes are used and shown in Figure 4.7: identical weighting, where each direction on the attitude sphere is weighted equally, and nadir-weighting, which is the weighting example given in Section 4.4.1. The former weighting would be used if the attitude control scheme is not finalized at this time in the design scheme or so that the design is robust to any attitude control scheme, and the nadir-weighting would be used if the design was for an Earth-pointing spacecraft. The objective values have been normalized by their maximum value from each weighting scheme to facilitate comparison on the same plot. The two different weighting schemes result in different optimal photodiode configurations. The optimal mounting angle for the identical weighting is 57°, and the optimal angle for the nadir weighting is 47°.

The total angular uncertainty over the attitude sphere is shown as a function of the eleva-

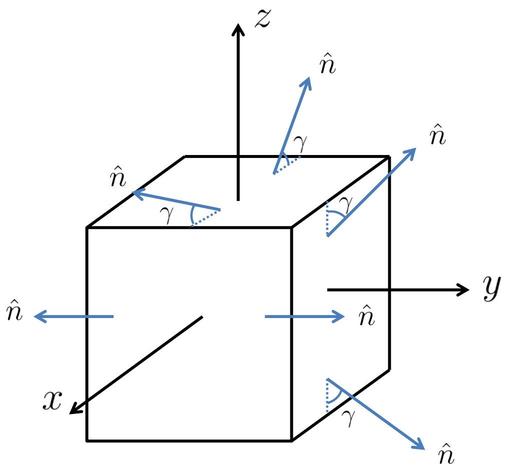

Figure 4.6: The photodiode configuration for Example 1 reduced to a single design variable. There are two photodiodes per face of the cube and the directions normal to the photosensitive planes are labeled with $\hat{n}$. The photodiodes on the $+x$, $+y$, and $+z$ surfaces are in the x-y, y-z, and x-z planes, respectively, and each photodiode is tilted by angle γ from the surface. γ is the single design variable. Photodiodes on the $-x$, $-y$, and $-z$ faces are not shown but have an identical configurations as the opposite faces of the cube.

tion angle in Figure 4.7(b). Calculation of this angular uncertainty is discussed in Appendix B. As seen in the figure, for this example, varying the elevation angle over the range of angles which maintain 4π steradian coverage does not result in a significant improvement in angular uncertainty.

The angles that achieve the minimum objective value shown in Figure 4.7(a) are optimal given the constraints of the geometry shown in Figure 4.6. To improve the configuration, these configurations are used as initial conditions to minimize the objective value with the full 24 design variables. MATLAB's optimization toolbox[6] is used to minimize the objective function subject to the design variables. The only constraint imposed during the optimization is to maintain sufficient photodiode coverage for sun sensing over 4π steradians.

The initial and optimized photodiode normal directions for both weighting schemes are shown in Figures 4.8(a) and 4.8(c), and the angular uncertainty of the optimized configurations are shown in Figures 4.8(b) and 4.8(d). The effect of the weighting is evident in Figure 4.8(d), where the uncertainty in the weighted region (where the elevation is less than 25°, see Figure 4.5) is less than in the region with zero weight. But for the case of

---

[6] *fmincon*, an algorithm to find the minimum of constrained nonlinear multivariable functions, is used to carry out the minimization. Details of the algorithm can be found in MATLAB's documentation: `http://www.mathworks.com/help/optim/ug/fmincon.html`, accessed November 2012

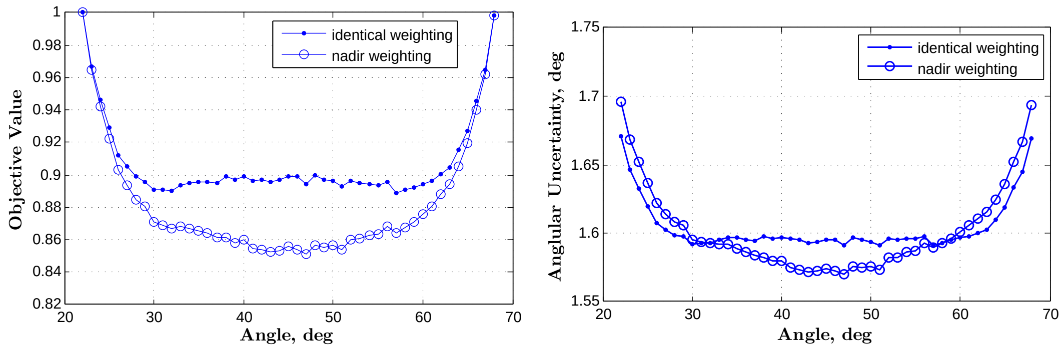

(a) Objective value (b) Sun vector angular uncertainty

Figure 4.7: The objective value and sun vector angular uncertainty as a function of photodiode mounting angle relative to the spacecraft surface in the geometry of Figure 4.6. The minimum objective values are achieved with mounting angles of 57° and 47° for the identical and nadir weighting schemes, respectively. Calculation of the sun vector angular uncertainty is discussed in the Appendix.

identical weighting over the attitude sphere, the uncertainty is mirrored about an elevation of 0°. For comparison to Figure 4.7(b), the total angular uncertainties of the identically- and nadir-weighted optimizations are 1.59° and 1.57°, respectively, which are insignificant improvements over the single variable optimization. The impact of the optimization will vary with the type of problem (number of sensors, design constraints, weighting scheme); Example 2 demonstrates the utility of the optimization in achieving significant angular accuracy improvements.

4.4.2.2 Example 2

As a second example, consider a cubical spacecraft with the geometry shown in Figure 4.9, which is derived from the actual design constraints of the RAX-2 spacecraft, shown in Figure 2.2(b) (see Sections 2.1 and 2.5). On RAX-2, antennas extend from the +/- z surfaces, potentially shadowing photodiodes mounted to these surfaces. To provide coverage in case of shadows, multiple sensors are placed in different locations with the same orientation. This orientation is normal to the surface; other orientations are not possible due to satellite structural requirements. Three photodiodes per x/y surface is the maximum allowable number of sensors constrained by the number of analog-to-digital converted channels in the telemetry system. These sensors can be tilted relative to the surface to which they are mounted, as in the photo of actual RAX-2 photodiodes in Figure 2.9.

The same procedure as Example 1 is used, where the problem is first reduced to a single design variable in order to choose feasible initial conditions for the constrained mul-

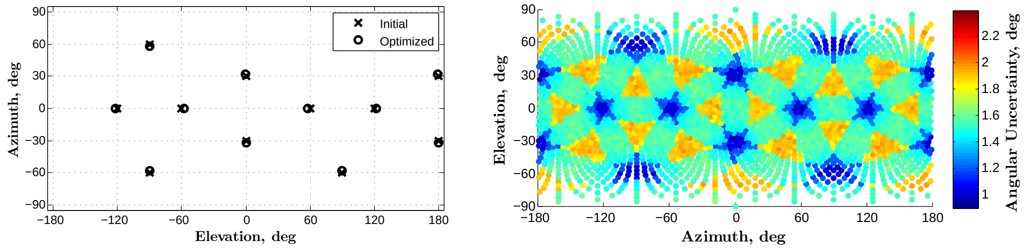

(a) Initial and optimized photodiode configurations for the case of identical weighting over the attitude sphere.

(b) Angular uncertainty of the sun vector estimates that result from the optimized configuration for the identical weighting scheme.

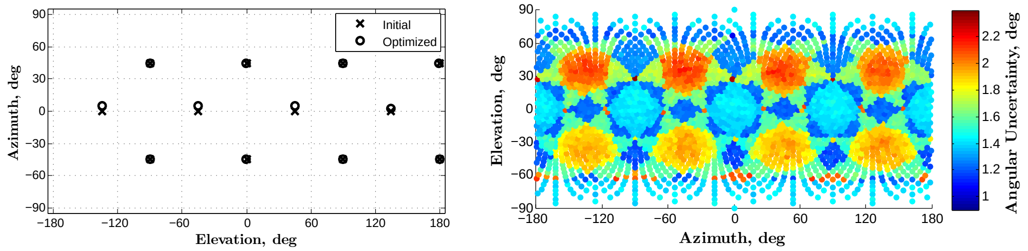

(c) Initial and optimized photodiode configurations for the nadir-pointing weighting scheme.

(d) Angular uncertainty of the sun vector estimates that result from the optimized configuration for the nadir-pointing weighting scheme.

Figure 4.8: Photodiode orientations and the resulting sun vector angular uncertainties over the attitude sphere for Example 1. In Figures 4.8(a) and 4.8(c), the azimuth and elevation angles are those of the direction normal to the photosensitive plane in the body-fixed frame.

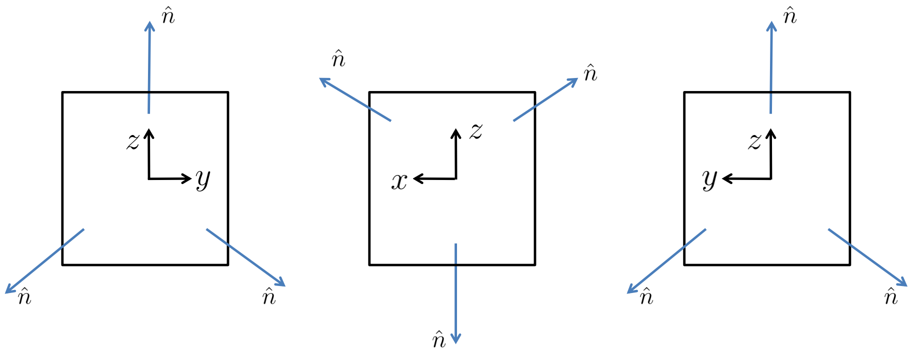

(a) +x surface (b) +y surface (c) -x surface

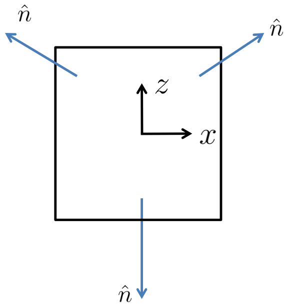

(d) -y surface

Figure 4.9: The photodiode configuration for Example 2 reduced to a single design variable, which is the mounting angle of the photodiodes on the x/y surfaces. Single photodiodes are placed on the +/- z surfaces of a cubical spacecraft and orientated normal to the surface (not shown). The three photodiodes on each of the remaining surfaces are shown with their normal directions denoted by $\hat{n}$. They are oriented in increments of 120° in the plane of the surfaces and are tilted from the surfaces (out of the page) by the single design angle.

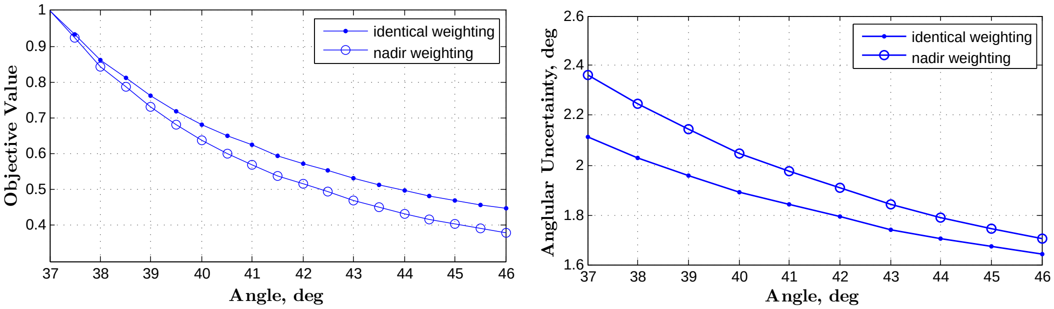

(a) Objective value (b) Sun vector angular uncertainty

Figure 4.10: The objective value and sun vector angular uncertainty as a function of photodiode mounting angle relative to the spacecraft surface in the geometry of Figure 4.9. The minimum objective values are achieved at the same mounting angle for both weighting schemes. Unlike Example 1, the angular uncertainty of the estimated sun vector varies significantly over the range of mounting angles that provide 4π steradian sun sensing coverage.

tivariable optimization. The single design variable is the angle of the photodiode normal directions from the surface to which they are mounted as shown in Figure 4.9. The objective values and resulting total sun vector angular uncertainties for both weighting schemes as a function of mounting angle is shown in Figure 4.10. The angle was varied over the maximum range that maintained sufficient coverage for sun sensing over 4π steradians. In this example, the optimal angle for both weighting schemes is the same. Unlike Example 1, the angular accuracy of the estimated sun vector as a function of the photodiode configuration varies significantly. In the nadir weighting scheme, angular accuracy improvement from 2.4° to 1.7° can be achieved by varying the single mounting angle.

The best angle obtained from varying the single design variable is used as the initial condition for the full multivariable optimization problem. MATLAB's *fmincon* is again used to optimize the configuration subject to 28 design variables (3 sensors per x/y surface and one each on the z surfaces). The only constraint imposed during the optimization is that the configuration must provide 4π steradian coverage. The resulting optimal orientation of all sensors is shown in Figure 4.11(a) and the angular uncertainty of the optimal configuration for each weighting scheme is shown in Figures 4.11(b) and 4.11(c). The total angular uncertainty of the optimized configurations achieved with both weighting schemes is 1.5°. Compared to the worst-case accuracy from variation of the single design variable shown in Figure 4.10 (2.1° and 2.4°), this is an improvement of 40% and 60% for the identical and nadir weighting schemes, respectively. This highlights the utility of optimizing the configuration for maximum angular accuracy rather then simply choosing a configuration with 4π steradian coverage.

4.4.2.3 Global and Local Solutions

In both examples, the problem was first reduced to single-variable optimization. The single-variable optimal solution was then used as the initial condition for the full multivariable optimization, which was carried out with MATLAB's *fmincon*. *fmincon* finds the nearest local minimum and is not a global optimal solution. In the photodide optimization, there are a large number of solutions with nearly the same objective value (for example, see Figure 4.10(a)), and in practice, there are additional factors that drive the design and provide an initial configuration for optimization. As demonstrated with Example 2, finding a local optimum near the initial condition can result in significant accuracy improvements, even if it is not proven to be a globally optimal solution. The approach presented in this chapter enables the use of standard minimization techniques for directional sensor optimization. Techniques such as genetic algorithms or simulated annealing could be used instead of *fmincon* to find globally-optimal solutions, but a detailed discussion on the available optimization techniques and comparison of their results is beyond the intended scope of this chapter.

4.5 Generalization the Optimization Method

The optimization method has been applied to photodiode configurations to minimize the angular uncertainty in sun sensing, but it is generally applicable to optimize the configuration of directional sensors and instruments. Potential additional applications include optimizing the orientations of body-fixed solar panels, antennas, and other types of attitude sensors that either have a limited field of view or do not have a uniform accuracy over the field of view. Here, the main steps of the optimization method are summarized.

The first step is to discretize the attitude sphere. The attitude sphere represents all directions in the vehicle body-fixed frame, and it is discretized to form a finite set of directions over which to optimize. The discretization is also referred to as pixelization. The icosahedron-based approach to discretization is used because it produces points that have a uniform angular distribution over the sphere [68]. Uniform distribution is important so that regions of the sphere are not artificially weighted by a higher density of points. The user must select the resolution of directions on the sphere based on the application. The relationship between the number of directions on the sphere and angular resolution of the directions is shown in Figure 4.3.

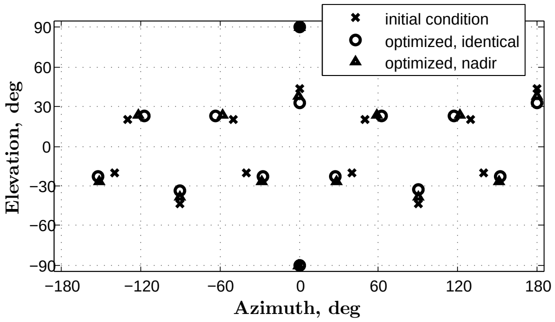

(a) Initial and optimized photodiode configurations.

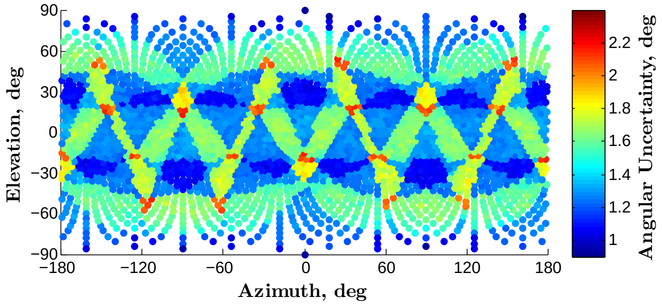

(b) Sun vector angular accuracy resulting from the configuration optimized using the identical weighting scheme.

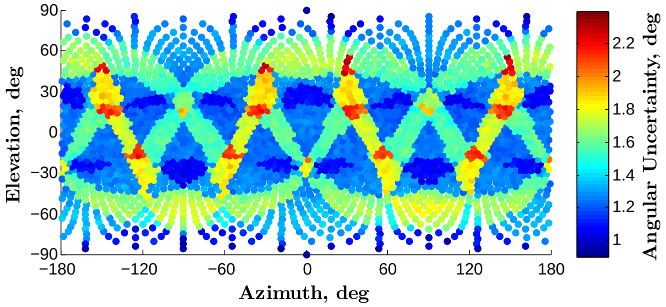

(c) Sun vector angular accuracy resulting from the configuration optimized using the nadir-pointing weighting scheme.

Figure 4.11: Photodiode orientations and the resulting sun vector angular uncertainties over the attitude sphere for Example 2. In Figure 4.11(a), the azimuth and elevation angles are those of the direction normal to the photosensitive plane in the body-fixed frame.

The second step is to formulate the objective function. To do so, the quantity to be optimized must be derived as a function of directions on the attitude sphere and the design variables. In application to photodiodes, the covariance of the sun vector was minimized as a function of photodiode orientation. If the relative importance of the directions on the sphere are not identical, each direction can be given a unique weight in the objective value. This can be used to account for vehicle- or mission-specific parameters such as orbit and attitude. In the example of Section 4.4.1, dwell time of the sun in each direction of the attitude sphere was used as the weight. Abstracting orbit and attitude as weights is advantageous because for a given sensor configuration, the component of the objective function in each direction on the attitude sphere needs to be calculated just once and then can be weighted for various potential mission parameters. This facilitates analysis of the robustness of a configuration over different mission parameters.

The final step is to minimize the objective value as a function of the design parameters. The behavior of the objective function will be problem specific. The photodiode optimization example was nonlinear and discontinuous, the problem was first reduced to a single design variable by limiting the configuration. Given that there are an infinite number of optimal photodiode configurations for an identically-weighted attitude sphere (free rotation around the sphere is possible), the initial reduction of the problem is useful to provide a feasible configuration that meets design constraints. MATLAB's optimization toolbox was then used to carry out the constrained minimization will the full set of variables to find locally optimal solutions. In general, the user must determine a suitable minimization algorithm based on the specific objective function.

4.6 Conclusions

A method to optimize the orientation of directional sensors in a vehicle body-fixed frame has been presented. The attitude sphere is used to formulate the optimization problem, and vehicle- and mission-specific parameters such as orbit and attitude are accounted for by weighting directions on the sphere. This formulation provides a method generally applicable to directional sensors and instruments.

The utility of the method has been demonstrated by application to photodiodes for sun sensing in attitude determination subsystems, where it was used to find photodiode configurations that minimize uncertainty in the resulting sun vector estimate. This maximizes the performance of the photodiodes used for attitude determination, and an improvement of 60% in angular accuracy was demonstrated in one of the examples. The method is useful not only because it maximizes the performance of a given set of hardware, but also because it provides a design technique to determine sensor orientation that replaces traditional methods in which placement is manually iterated to meet requirements. This can result in both improved subsystem performance and reduced design time and cost.

CHAPTER 5

On-Orbit Photodiode Calibration

In this chapter, we present a method for on-orbit calibration of photodiodes for attitude determination. While the previous chapter provides an optimal design method to orient the photodiodes, this chapter provides a method to estimate the actual orientation of the photodiodes in flight such that no high tolerance or alignment procedures are required when installing the photodiodes. This method also estimates the sensor scale factors, mitigating the need to calibrate the individual sensors before flight.

5.1 Introduction

Sun sensors are the most widely used sensor type in attitude determination systems [71]. They are used to measure the vector from the spacecraft to the sun (herein referred to as the sun vector) in the spacecraft body-fixed frame, and their angular accuracies range from several degrees to less than an arc-second. Photodiodes, which generate current as a function of incoming light [74], are the most basic type of sun sensor. As discussed in Chapter 4, a stand-alone photodiode provides a measurement of the angle between the sun vector and the direction normal to the photosensitive plane, effectively measuring one component of the sun vector. For multi-axis sun sensing with photodiodes, two common schemes exist: (1) individual photodiodes can be mounted in different orientations – either within a single sensor package [75] or distributed over the spacecraft body (see Chapter 4) – for up to three-component sun sensing, or (2) multiple photodiodes and a mask can be combined within a single sensor package for two-axis sun sensing. The angular accuracy of the former scheme is on the order of degrees, and the angular accuracy of the latter, which can be as good as arcseconds, depends on the complexity of the sensor [76, 77]. A photodiode-based sensor with accuracy on the order of degrees is typically referred to as a coarse sun sensor, whereas sensors with angular accuracies on the order of tenths of degrees or better are referred to as fine sun sensors. Fine sun sensors commonly utilize more advanced components than photodiodes, such as CMOS sensors [78].

This work focuses on coarse sun sensors composed of individual photodiodes mounted at different angles. This photodiode scheme is extremely common on CubeSats and other small spacecraft [31,59–62] because of its simplicity and low cost[1]. One configuration that is typical of CubeSats is to mount the photodiodes orthogonal to each other by placing one on each surface of the spacecraft [31,59,62]. However, this configuration does not provide three-component sun sensing in all directions because the conical field of view of photodiodes is typically less than 180° (see Section 2.5). A natural improvement to this is to use additional photodiodes mounted in various orientations over the spacecraft body to achieve three-component sun sensing in all directions. In addition to enabling three-component sun sensing, this configuration provides more information for the attitude determination system, resulting in a potential angular accuracy improvement. This photodiode scheme has been used on various spacecraft [60,75,79], and a design method to determine optimal photodiode orientations for sun sensing was the topic of Chapter 4.

There are two parameters critical for accurate sun sensing with photodiodes: the scale factor – which relates the measured current to the sun vector component – and the orientation of each photodiode on the spacecraft. The scale factor is dependent on the characteristics of both the photodiode and surrounding circuitry, and it will be discussed further in Section 5.3.2. Although the scale factor can be estimated from pre-flight calibration, photodiodes are known to degrade on-orbit due to radiation, and previous flight experience demonstrates that this has a significant effect on the scale factor [80]. Additionally, thorough pre-flight calibration requires a light source that is calibrated to match the characteristics of sunlight in orbit. On-orbit estimation of the scale factor provides the best estimate of the on-orbit sensor characteristics and lowers spacecraft development time and cost by mitigating rigorous pre-flight calibration requirements. Similarly, the pre-flight orientation of the photodiodes is known if sufficiently high tolerance procedures are used during spacecraft integration, but any high-tolerance procedures that add cost to the integration defeat the low-cost benefits of photodiodes, and the orientation may change during launch regardless of the initial tolerance. Therefore, the photodiode orientation is also estimated in flight.

The contributions of this chapter are the development and application of a new, on-orbit photodiode calibration method to estimate the photodiode scale factor and orientation. Implementations of the calibration within both an extended Kalman filter (EKF) and

---

[1] Osram SFH 2430 photodiodes, which flew on the RAX satellites [31], can be purchased for $1.84 per sensor (`http://www.digikey.com/product-detail/en/SFH%202430-Z/475-2579-1-ND/1228076`, accessed March 2013).

an unscented filter (UF) are presented and compared. This recursive, attitude-dependent approach enables the inclusion of an Earth albedo model, which is demonstrated to be an important aspect of the calibration. In application to flight data from the RAX satellites, an average of 10° improvement in accuracy of the photodiode-based sun vector measurement is achieved with the on-orbit calibration.

The remainder of this chapter is organized as follows. Existing calibration techniques and the motivation for formulating the photodiode calibration as a recursive attitude-dependent method is discussed in Section 5.2. The formulation of the calibration within both the EKF and UF is discussed in Section 5.3, and both filters are compared via simulation and applied to flight data in Section 5.4. The impact on sun vector angular accuracy and the importance of including the Earth albedo model is also discussed in Section 5.4. A discussion of some of the assumptions and techniques used in the calibration is presented in Section 5.5, and conclusions are given in Section 5.6.

5.2 Existing Calibration Techniques

In general, sensor measurements are corrupted by errors such as scale factors, bias, and angular mis-alignments. Calibration is used to estimate and subsequently compensate for the sensor errors, and it is critical for accurate attitude estimation. Many calibration techniques have been presented in the literature; this section is not meant to include an exhaustive overview of existing calibration techniques, but rather to summarize the most relevant types of methods in the context of the new photodiode calibration presented in this paper.

The calibration process consists of modeling the sensor of interest and estimating the parameters of the model using sensor measurements from either ground-based testing or on-orbit operations. Calibration using only on-orbit data is referred to as *on-orbit* calibration, and it is advantageous over ground-based calibration because it accounts for any changes in sensor characteristics once they are in orbit, and it reduces satellite development time and cost by mitigating the pre-flight calibration requirements.

On-orbit calibration techniques can be categorized as either attitude-independent or attitude-dependent. Attitude-independent calibration does not require attitude knowledge and is accomplished by minimizing a scalar objective function that is dependent on the calibration parameters. This was the approach taken for magnetometer calibration in Chapter 3, which utilized the magnitude of the measured vector as the scalar objective function. Another typical attitude-independent objective function is the angle between two vectors in the same frame [52, 53]. Attitude-independent methods can be applied to various types of three-axis sensors. Attitude-dependent techniques are recursive methods that use attitude estimates for the sensor calibration, and thus simultaneously estimate attitude and sensor calibration parameters. Examples of this approach include References 81 and 82, which utilize an EKF and a UF, respectively, to estimate sensor mis-alignments, rate gyroscope scaling and bias, and attitude.

Regarding sun sensors specifically, on-orbit calibration of coarse sensors is not prevalent in the literature. Furthermore, the models of fine sun sensors are dependent on the sensor design, so existing calibration methods are typically applicable to only specific sensors [76, 77]. While both the attitude-independent and -dependent calibration approaches mentioned previously can be applied to on-orbit sun sensor calibration, direct application of these methods requires a sun sensor that provides a three-component sun vector measurement [83,84]. Since photodiodes provide a measurement of a single sun vector component, a three-component measurement is not always available, making application of these methods non-trivial.

The new, on-orbit calibration method presented in this chapter has been developed for stand-alone photodiodes in any orientation and does not require simultaneous illumination of multiple sensors, which would be required for a three-component vector measurement. This method enables the calibration of an arbitrary number of illuminated sensors and facilitates the use of an attitude-dependent Earth albedo model. Earth albedo is reflection of sunlight from Earth's surface, and it can significantly degrade the accuracy of photodiode-based sun sensing. Both of these aspects are discussed in more detail with the calibration methodology in the next section.

5.3 Formulation of the Attitude Estimation and Calibration Filter

On-orbit photodiode calibration is achieved by estimating the calibration parameters and spacecraft attitude simultaneously. We formulate the estimation problem by extending existing recursive attitude estimation methods to include the calibration parameters as states, which is a similar approach to that of existing attitude-dependent on-orbit calibration techniques that deal with star tracker and rate gyroscope misalignment [81,82]. The calibration has been implemented within both an EKF and UF. The EKF is a widely-used approach for non-linear state estimation and has become a standard method for satellite attitude determination [24,25]. The UF is a more recent estimation method that uses a different propagation technique than the standard EKF [85], and it can be used as an alternative, potentially more accurate estimator for attitude estimation [24, 86]. In general, UFs can be advantageous over the standard EKF because the expected error is lower, they can be applied to non-differentiable functions, Jacobian matrices are not required, and the UF is valid to higher order expansions than the standard EKF [86]. But the degree of improvement over the EKF depends on the nature of estimation problem and UFs are more computationally intensive than the EKF, particularly if parallel processing is not available. A comparison of the photodiode calibration accuracy when using an EKF and UF is discussed in Section 5.4.2, where we show that under initial conditions that are sufficiently accurate, the difference between EKF and UF accuracy is negligible, but under higher initial uncertainty, the UF provides significantly more accurate state estimates than the EKF.

Both EKFs and UFs can be used to estimate the states of the system with the general form given by Eq. (5.1), where $\boldsymbol{x}_k$ is the state of the system at time k, $\tilde{\boldsymbol{y}}_k$ is the measurement vector at time k, $\boldsymbol{f}()$ and $\boldsymbol{h}()$ are nonlinear functions, G_k is a gain, and $\boldsymbol{w}_k \sim N(0, Q_k)$ and $\boldsymbol{\nu}_k \sim N(0, R_k)$ random vectors that quantify model and measurement uncertainty. Q_k and R_k are referred to as the process covariance and measurement covariance, respectively.

$$\boldsymbol{x}_{k+1} = \boldsymbol{f}\left(\boldsymbol{x}_k, k\right) + G_k \boldsymbol{w}_k \tag{5.1a}$$

$$\tilde{\boldsymbol{y}}_k = \boldsymbol{h}\left(\boldsymbol{x}_k, k\right) + \boldsymbol{\nu}_k \tag{5.1b}$$

Direct application of both the EKF and UF to attitude estimation is non-trivial due to the orthogonality constraint of the attitude matrix, or equivalently, the unity norm constraint when quaternions are used to parametrize attitude; the norm constraint can be violated by the linear measurement updates of the standard EKF. This shortfall has been widely studied, and a common implementation of an EKF for attitude estimation has become known as the multiplicative extended Kalman filter (MEKF). The term *multiplicative* stems from the fact that quaternion multiplication supplements the standard additive state update within the EKF [25, 87]. The multiplicative approach can also be used to satisfy the quaternion norm constraint in unscented filtering, and the specific implementation of this UF for attitude estimation is called the unscented quaternion estimator (USQUE) [86].

In the remainder of this section, the conventional forms of the MEKF and USQUE are first presented before extending the estimators to include the photodiode calibration states. A complete derivation and explanation of the MEKF and USQUE are beyond the scope of this chapter, but we provide an overview of each while assuming the reader is familiar with Kalman filtering, unscented filtering, and spacecraft attitude dynamics [24, 25, 85–88]. The equations for both attitude estimation filters are given in Section 5.3.1 with sufficient detail for the reader to implement them, and this provides the background information that is necessary to extend both estimators for recursive photodiode calibration. The photodiode measurement model and Earth albedo are discussed in Section 5.3.2, and the filters are extended to include the photodiode calibration in Sections 5.3.3-5.3.4.

5.3.1 MEKF and USQUE Summary

5.3.1.1 MEKF

The MEKF is a widely-used approach to spacecraft attitude estimation for vehicles with a three-axis rate gyroscope and at least one attitude sensor [25, 87, 88]. The filter estimates six states: a 3×1 attitude error vector $\delta \boldsymbol{p}$ and the 3×1 rate gyroscope bias $\boldsymbol{\beta}$. The attitude error vector, which parametrizes error quaternions $\delta \boldsymbol{q} = \begin{bmatrix} \delta \boldsymbol{\varrho} & \delta q_4 \end{bmatrix}$, is combined with quaternions being propagated within the filter to estimate spacecraft attitude. Various three-component attitude error vectors can be used, such as the Gibb's vector, and each results in the same approximate relationship to the error quaternion, given in Eq. (5.2), that is valid to second-order (sufficient for the EKF) [87].

$$\delta \boldsymbol{q}\left(\delta \boldsymbol{p}\right) \approx \begin{bmatrix} \delta \boldsymbol{p}/2 \\ 1 - \delta \boldsymbol{p}^2/8 \end{bmatrix} \tag{5.2}$$

The state vector is given by $\boldsymbol{x} = \begin{bmatrix} \delta \boldsymbol{p}^T & \boldsymbol{\beta}^T \end{bmatrix}^T$. Unlike the general form of an EKF which combines a dynamic model with sensor measurements, the MEKF does not utilize a dynamic model. Instead, the bias-corrected gyroscope measurements are used directly to propagate the attitude kinematics, which alleviates the difficulties of accurately modeling spacecraft dynamics [25].

The following procedure is used for attitude estimation with the MEKF. First, the state vector estimate is initialized to $\hat{\boldsymbol{x}}_0 = \begin{bmatrix} \boldsymbol{0}_{1\times 3} & \hat{\boldsymbol{\beta}}_0^T \end{bmatrix}^T$, the state error-covariance is initialized to P_0, and the quaternion estimate is initialized to $\hat{\boldsymbol{q}}_0$. The attitude error vector is initialized to zero because the attitude is quantified completely by the quaternion estimate. Then, the Kalman gain is computed using Eqs. (5.3)-(5.4). Throughout the equations, the superscripts $^-$ and $^+$ denote quantities before and after the state update, respectively. In Eq. (5.4), $A(\boldsymbol{q})$ is the attitude matrix corresponding to the quaternion $\boldsymbol{q}$, $\boldsymbol{r}_i$ is the i-th vector measurement, and $\times$ denotes the skew-symmetric cross-product matrix. The state estimate and error-covariance are updated via the conventional EKF update of Eqs. (5.5)-(5.7). After the state update, the attitude error portion of the updated state, $\delta \hat{\boldsymbol{p}}_k^+$, is used to update the quaternion estimate using Eq. (5.8). Eq. (5.8) is where the term *multiplicative* comes from: the attitude error that was updated via the traditional additive update was passed to the quaternion estimate using quaternion multiplication. Next, the bias estimate is used to correct the rate gyro measurement as in Eq. (5.9), and the attitude estimate is then propagated using Eqs. (5.10)-(5.12). Since the attitude error has been passed to the quaternion estimate, the attitude error estimate is reset to zero as in Eq. (5.13). The bias estimate remains constant until the next time step as in Eq. (5.14). Finally, the state error-covariance is propagated using Eqs. (5.15)-(5.18), where Δt is the time increment until the next available measurement. In Eq. (5.18), σ_u and σ_v are characteristics of the rate gyroscope, discussed in the next paragraph. After propagation to the next time step, the process is repeated with the next available measurement starting with the Kalman gain computation of Eq. (5.3).

$$K_k = P_k^- H_k^T\left(\hat{\boldsymbol{x}}_k^-\right)\left[H_k\left(\hat{\boldsymbol{x}}_k^-\right)P_k^- H_k^T\left(\hat{\boldsymbol{x}}_k^-\right) + R_k\right]^{-1} \tag{5.3}$$

$$H_k\left(\hat{\boldsymbol{x}}_k^-\right) \equiv \left.\frac{\partial \boldsymbol{h}}{\partial \boldsymbol{x}}\right|_{\hat{\boldsymbol{x}}_k^-} = \left.\begin{bmatrix} \left[A\left(\hat{\boldsymbol{q}}^-\right)\boldsymbol{r}_1\times\right] & \mathbf{0}_{3\times 3} \\ \vdots & \vdots \\ \left[A\left(\hat{\boldsymbol{q}}^-\right)\boldsymbol{r}_N\times\right] & \mathbf{0}_{3\times 3} \end{bmatrix}\right|_{t_k} \tag{5.4}$$

$$P_k^+ = \left[I - K_k H_k\left(\hat{\boldsymbol{x}}_k^-\right)\right]P_k^- \tag{5.5}$$

$$\hat{\boldsymbol{x}}_k^+ = \hat{\boldsymbol{x}}_k^- + K_k\left[\tilde{\boldsymbol{y}}_k - \boldsymbol{h}\left(\hat{\boldsymbol{x}}_k^-\right)\right] \tag{5.6}$$

$$\boldsymbol{h}\left(\hat{\boldsymbol{x}}_k^-\right) = \left.\begin{bmatrix} A\left(\hat{\boldsymbol{q}}^-\right)\boldsymbol{r}_1 \\ \vdots \\ A\left(\hat{\boldsymbol{q}}^-\right)\boldsymbol{r}_N \end{bmatrix}\right|_{t_k} \tag{5.7}$$

$$\hat{\boldsymbol{q}}_k^+ = \hat{\boldsymbol{q}}_k^- + \frac{1}{2}\Xi\left(\hat{\boldsymbol{q}}_k^-\right)\delta\hat{\boldsymbol{p}}_k^+ \tag{5.8}$$

$$\hat{\boldsymbol{\omega}}_k^+ = \tilde{\boldsymbol{\omega}}_k - \hat{\boldsymbol{\beta}}_k^+ \tag{5.9}$$

$$\hat{\boldsymbol{q}}_{k+1}^- = \Omega\left(\hat{\boldsymbol{\omega}}_k^+\right)\hat{\boldsymbol{q}}_k^+ \tag{5.10}$$

$$\Omega\left(\hat{\boldsymbol{\omega}}_k^+\right) \equiv \begin{bmatrix} \cos\left(0.5\left\|\hat{\boldsymbol{\omega}}_k^+\right\|\Delta t\right) I_{3\times 3} - \left[\hat{\psi}_k^+\times\right] & \hat{\psi}_k^+ \\ -\hat{\psi}_k^{+T} & \cos\left(0.5\left\|\hat{\boldsymbol{\omega}}_k^+\right\|\Delta t\right) \end{bmatrix} \tag{5.11}$$

$$\hat{\psi}_k^+ \equiv \sin\left(0.5\left\|\hat{\boldsymbol{\omega}}_k^+\right\|\Delta t\right)\hat{\boldsymbol{\omega}}_k^+ / \left\|\hat{\boldsymbol{\omega}}_k^+\right\| \tag{5.12}$$

$$\delta\hat{\mathbf{p}}_{k+1}^- = \mathbf{0}_{3\times 1} \tag{5.13}$$

$$\hat{\boldsymbol{\beta}}_{k+1}^- = \hat{\boldsymbol{\beta}}_k^+ \tag{5.14}$$

$$P_{k+1}^- = \Phi_k P_k^+ \Phi_k^T + G_k Q_k G_k^T \tag{5.15}$$

$$\Phi = \begin{bmatrix} \Phi_{11} & \Phi_{12} \\ \Phi_{21} & \Phi_{22} \end{bmatrix} \tag{5.16a}$$

$$\Phi_{11} = I_{3\times3} - [\hat{\boldsymbol{\omega}}\times] \frac{\sin\left(\|\hat{\boldsymbol{\omega}}\| \Delta t\right)}{\|\hat{\boldsymbol{\omega}}\|} + [\hat{\boldsymbol{\omega}}\times]^2 \frac{\{1 - \cos\left(\|\hat{\boldsymbol{\omega}}\| \Delta t\right)\}}{\|\hat{\boldsymbol{\omega}}\|^2} \tag{5.16b}$$

$$\begin{aligned} \Phi_{12} = {} & [\hat{\boldsymbol{\omega}}\times] \frac{\{1 - \cos((\|\hat{\boldsymbol{\omega}}\| \Delta t)\}}{\|\hat{\boldsymbol{\omega}}\|^2} - I_{3\times3}\Delta t \\ & - [\hat{\boldsymbol{\omega}}\times]^2 \frac{\{\|\hat{\boldsymbol{\omega}}\| \Delta t - \sin\left(\|\hat{\boldsymbol{\omega}}\| \Delta t\right)\}}{\|\hat{\boldsymbol{\omega}}\|^3} \end{aligned} \tag{5.16c}$$

$$\Phi_{21} = 0_{3\times3} \tag{5.16d}$$

$$\Phi_{22} = I_{3\times3} \tag{5.16e}$$

$$G_k = \begin{bmatrix} -I_{3\times3} & 0_{3\times3} \\ 0_{3\times3} & I_{3\times3} \end{bmatrix} \tag{5.17}$$

$$Q_k = \begin{bmatrix} \left(\sigma_v^2 \Delta t + \frac{1}{3}\sigma_u^2 \Delta t^3\right) I_{3\times3} & \left(\frac{1}{2}\sigma_u^2 \Delta t^2\right) I_{3\times3} \\ \left(\frac{1}{2}\sigma_u^2 \Delta t^2\right) I_{3\times3} & \left(\sigma_u^2 \Delta t\right) I_{3\times3} \end{bmatrix} \tag{5.18}$$

A widely-used model for a rate gyroscope is given by

$$\tilde{\boldsymbol{\omega}}(t) = \boldsymbol{\omega}(t) + \boldsymbol{\beta}(t) + \boldsymbol{\eta}_v(t) \tag{5.19a}$$

$$\dot{\boldsymbol{\beta}}(t) = \boldsymbol{\eta}_u(t) \tag{5.19b}$$

where $\tilde{\boldsymbol{\omega}}(t)$ is the continuous-time measured angular rate and $\boldsymbol{\eta}_v(t)$ and $\boldsymbol{\eta}_u(t)$ are independent mean-zero Gaussian white-noise processes with

$$E\left\{\boldsymbol{\eta}_v(t)\boldsymbol{\eta}_v^T(\tau)\right\} = I_{3\times3}\sigma_v^2\delta(t-\tau) \tag{5.20a}$$

$$E\left\{\boldsymbol{\eta}_u(t)\boldsymbol{\eta}_u^T(\tau)\right\} = I_{3\times3}\sigma_u^2\delta(t-\tau) \tag{5.20b}$$

where $\delta(t-\tau)$ is the Dirac delta function [89]. In practice, σ_v and σ_u are typically referred to as angular random walk and rate random walk, respectively.

5.3.1.2 USQUE

The unscented quaternion estimator (USQUE) is an implementation of the UF developed for attitude estimation [86]. Like the MEKF, it requires a three-axis rate gyroscope and at least one attitude sensor, and it maintains the quaternion norm constraint using the same multiplicative approach. In this section, the notation is largely the same as that used for the MEKF, such as P representing state error-covariance, but the equations in this subsection are unique to USQUE and should not be combined with the MEKF equations of Section 5.3.1.1 except where noted.

USQUE utilizes generalized Rodrigues parameters to define the attitude error vector as in Eq. (5.21), where a is a parameter from 0 to 1 and f is a scale factor. This is analogous to Eq. (5.2), but is an exact relationship and facilitates the use of any parametrization of the error quaternion [86]. $\delta\mathbf{p}$ is equivalent to the Gibb's vector when $a = 0$ and $f = 1$. The authors of USQUE use $f = 2(a + 1)$ such that $\|\delta\boldsymbol{p}\| = \vartheta$, where ϑ is the rotation angle of the attitude quaternions.

$$\delta\mathbf{p} \equiv f\left[\delta\boldsymbol{\varrho}/(a + \delta q_4)\right] , \tag{5.21}$$

The main difference between the UF and EKF lies in the manner in which Gaussian random variables are quantified and propagated through the dynamic model. In the EKF, the state error covariance is propagated linearly via first-order linearization of the nonlinear system. The UF quantifies the state distribution by using carefully chosen sample points, called *sigma points*. The sigma points completely capture the mean and covariance of the state and are propagated through the true nonlinear system [90]. Given an $n \times n$ state covariance matrix P, a set of $2n$ sigma points are generated from the columns of the matrices $\pm\sqrt{(n+\lambda)P}$, where $\sqrt{M}$ is shorthand notation for a matrix Z such that $ZZ^T = M$. λ is a scalar parameter that can be used to exploit knowledge about higher order moments for the given distribution, if available [85]. For more details, see References 86, 90, and the references therein.

The procedure for attitude estimation with USQUE is the following. The estimated quaternion, estimated state, and state error covariance are first initialized to $\hat{\boldsymbol{q}}_0^+$, $\hat{\boldsymbol{x}}_0^+ = \begin{bmatrix} \mathbf{0}_{1\times 3} & \hat{\boldsymbol{\beta}}_0^{+T} \end{bmatrix}^T$, and P_0^+, respectively. Next, the sigma points are calculated using Eqs. (5.22)-(5.23). The process covariance can be accounted for as in Eq. (5.22) because it is purely additive [90][2]. The sigma points are partitioned as in Eq. (5.24), where $\boldsymbol{\chi}_k^{\delta p}$ is the attitude error portion and $\boldsymbol{\chi}_k^{\beta}$ is the gyro bias portion. The corresponding error quaternions and sigma point quaternions are calculated with Eqs. (5.25) and (5.26), respectively. The quaternions are subsequently propagated using Eqs. (5.27)-(5.28), where $\Omega()$ is given by Eq. (5.11). The propagated error quaternions are then calculated with Eq. (5.29) and the sigma points are propagated with Eqs. (5.30)-(5.31). The predicted state and state error covariance can now be calculated using Eqs. (5.32)-(5.33). Using the predicted quaternion

---

[2] The authors of USQUE derived a different treatment of Q which approximates the integration of the process noise over the sampling interval, but this has been shown to have a negligible impact on the performance of the filter, so it is not utilized here [82, 86].

from Eq. (5.27), the mean observation is calculated with Eqs. (5.34)-(5.35). The output covariance, innovation covariance, and cross-correlation matrix are computed with Eqs. (5.36)-(5.38). The gain is then computed with Eq. (5.39) and the state and error covariance are updated with Eqs. (5.40)-(5.42). The updated state is used to calculate the updated error quaternion with Eq. (5.43), which is subsequenty used to update the attitude estimate with Eq. (5.44). Finally, the attitude error vector $\delta\hat{\boldsymbol{p}}_{k+1}^{+}$ is reset to zero for the next propagation. A detailed derivation and explanation of USQUE can be found in Reference [86] and the references therein.

$$\boldsymbol{\chi}_k(0) = \hat{\mathbf{x}}_k^+ \tag{5.22a}$$

$$\boldsymbol{\chi}_k(i) = \boldsymbol{\sigma}_k(i) + \hat{\mathbf{x}}_k^+ \ , \ i = 1 \ldots 12 \tag{5.22b}$$

$$\boldsymbol{\sigma}_k \leftarrow 2n \text{ columns from } \pm\sqrt{(n+\lambda)\left[P_k^+ + Q_k\right]} \tag{5.22c}$$

$$Q_k = \begin{bmatrix} \left(\sigma_v^2 \Delta t + \frac{1}{3}\sigma_u^2 \Delta t^3\right) I_{3\times 3} & -\left(\frac{1}{2}\sigma_u^2 \Delta t^2\right) I_{3\times 3} \\ -\left(\frac{1}{2}\sigma_u^2 \Delta t^2\right) I_{3\times 3} & \left(\sigma_u^2 \Delta t\right) I_{3\times 3} \end{bmatrix} \tag{5.23}$$

$$\boldsymbol{\chi}_k(i) \equiv \begin{bmatrix} \boldsymbol{\chi}_k^{\delta p}(i) \\ \boldsymbol{\chi}_k^{\beta}(i) \end{bmatrix} \ , \ i = 0, 1, \ldots, 12 \tag{5.24}$$

$$\delta\boldsymbol{\varrho}_k^+(i) = f^{-1}\left[a + \delta q_{4_k}^+(i)\right]\boldsymbol{\chi}_k^{\delta p}(i) \ , \ i = 1, 2, \ldots, 12 \tag{5.25a}$$

$$\begin{aligned} \delta q_{4_k}^+(i) = & \frac{-a\left\|\boldsymbol{\chi}_k^{\delta p}(i)\right\|^2 + f\sqrt{f^2 + (1-a^2)\left\|\boldsymbol{\chi}_k^{\delta p}(i)\right\|^2}}{f^2 + \left\|\boldsymbol{\chi}_k^{\delta p}(i)\right\|^2} \ , \\ & i = 1, 2, \ldots, 12 \end{aligned} \tag{5.25b}$$

$$\hat{\boldsymbol{q}}_k^+(0) = \hat{\boldsymbol{q}}_k^+ \tag{5.26a}$$

$$\hat{\boldsymbol{q}}_k^+(i) = \delta\boldsymbol{q}_k^+(i) \otimes \hat{\boldsymbol{q}}_k^+ \ , \ i = 1, 2, \ldots, 12 \tag{5.26b}$$

$$\hat{\boldsymbol{q}}_{k+1}^-(i) = \Omega\left[\hat{\boldsymbol{\omega}}_k^+(i)\right]\hat{\boldsymbol{q}}_k^+(i) \ , \ i = 0, 1, \ldots, 12 \tag{5.27}$$

$$\hat{\boldsymbol{\omega}}_k^+(i) = \tilde{\boldsymbol{\omega}}_k - \boldsymbol{\chi}_k^{\beta}(i) \ , \ i = 0, 1, \ldots, 12 \tag{5.28}$$

$$\delta\boldsymbol{q}_{k+1}^-(i) = \hat{\boldsymbol{q}}_{k+1}^-(i) \otimes \left[\hat{\boldsymbol{q}}_{k+1}^-(0)\right]^{-1} \ , \ i = 0, 1, \ldots, 12 \tag{5.29}$$

$$\boldsymbol{\chi}_{k+1}^{\delta p}(0) = \mathbf{0} \tag{5.30a}$$

$$\boldsymbol{\chi}^{\delta p}_{k+1}(i) = f\,\frac{\delta\boldsymbol{\varrho}^{-}_{k+1}(i)}{a + \delta q^{-}_{4_{k+1}}(i)}\ ,\ i = 1, 2, \ldots, 12 \tag{5.30b}$$

$$\boldsymbol{\chi}^{\beta}_{k+1}(i) = \boldsymbol{\chi}^{\beta}_{k}(i)\ ,\ i = 0, 1, \ldots, 12 \tag{5.31}$$

$$\hat{\boldsymbol{x}}^{-}_{k+1} = \frac{1}{n+\lambda}\left\{\lambda\boldsymbol{\chi}_{k+1}(0) + \frac{1}{2}\sum_{i=1}^{2n}\boldsymbol{\chi}_{k+1}(i)\right\} \tag{5.32}$$

$$\begin{aligned} P^{-}_{k+1} = \frac{1}{n+\lambda}\Bigg\{ & \lambda\left[\boldsymbol{\chi}_{k+1}(0) - \hat{\boldsymbol{x}}^{-}_{k+1}\right]\left[\boldsymbol{\chi}_{k+1}(0) - \hat{\boldsymbol{x}}^{-}_{k+1}\right]^{T} \\ & + \frac{1}{2}\sum_{i=1}^{2n}\left[\boldsymbol{\chi}_{k+1}(i) - \hat{\boldsymbol{x}}^{-}_{k+1}\right]\left[\boldsymbol{\chi}_{k+1}(i) - \hat{\boldsymbol{x}}^{-}_{k+1}\right]^{T}\Bigg\} + Q_k \end{aligned} \tag{5.33}$$

$$\hat{\boldsymbol{y}}^{-}_{k+1} = \frac{1}{n+\lambda}\left\{\lambda\boldsymbol{\gamma}_{k+1}(0) + \frac{1}{2}\sum_{i=1}^{2n}\boldsymbol{\gamma}_{k+1}(i)\right\} \tag{5.34}$$

$$\boldsymbol{\gamma}_{k+1}(i) = \begin{bmatrix} A\left[\hat{\boldsymbol{q}}^{-}(i)\right]\boldsymbol{r}_1 \\ \vdots \\ A\left[\hat{\boldsymbol{q}}^{-}(i)\right]\boldsymbol{r}_N \end{bmatrix}_{k+1} \tag{5.35}$$

$$\begin{aligned} P^{yy}_{k+1} = \frac{1}{n+\lambda}\Bigg\{ & \lambda\left[\boldsymbol{\gamma}_{k+1}(0) - \hat{\boldsymbol{y}}^{-}_{k+1}\right]\left[\boldsymbol{\gamma}_{k+1}(0) - \hat{\boldsymbol{y}}^{-}_{k+1}\right]^{T} \\ & + \frac{1}{2}\sum_{i=1}^{2n}\left[\boldsymbol{\gamma}_{k+1}(i) - \hat{\boldsymbol{y}}^{-}_{k+1}\right]\left[\boldsymbol{\gamma}_{k+1}(i) - \hat{\boldsymbol{y}}^{-}_{k+1}\right]^{T}\Bigg\} \end{aligned} \tag{5.36}$$

$$P^{vv}_{k+1} = P^{yy}_{k+1} + R_{k+1} \tag{5.37}$$

$$\begin{aligned} P^{xy}_{k+1} = \frac{1}{n+\lambda}\Bigg\{ & \lambda\left[\boldsymbol{\chi}_{k+1}(0) - \hat{\boldsymbol{x}}^{-}_{k+1}\right]\left[\boldsymbol{\gamma}_{k+1}(0) - \hat{\boldsymbol{y}}_{k+1}\right]^{T} \\ & + \frac{1}{2}\sum_{i=1}^{2n}\left[\boldsymbol{\chi}_{k+1}(i) - \hat{\boldsymbol{x}}^{-}_{k+1}\right]\left[\boldsymbol{\gamma}_{k+1}(i) - \hat{\boldsymbol{y}}^{-}_{k+1}\right]^{T}\Bigg\} \end{aligned} \tag{5.38}$$

$$K_k = P^{xy}_{k}\left(P^{vv}_{k}\right)^{-1} \tag{5.39}$$

$$\hat{\boldsymbol{x}}^{+}_{k} = \hat{\boldsymbol{x}}^{-}_{k} + K_k\boldsymbol{v}_k \tag{5.40}$$

$$\boldsymbol{v}_k \equiv \tilde{\boldsymbol{y}}_k - \hat{\boldsymbol{y}}^{-}_{k} = \tilde{\boldsymbol{y}}_k - \boldsymbol{h}\left(\hat{\boldsymbol{x}}^{-}_{k}, k\right) \tag{5.41}$$

$$P^{+}_{k} = P^{-}_{k} - K_k P^{vv}_{k} K^{T}_{k} \tag{5.42}$$

$$\delta \boldsymbol{\varrho}_{k+1}^{+} = f^{-1}\left[a + \delta q_{4_{k+1}}^{+}\right]\delta \hat{\boldsymbol{p}}_{k+1}^{+} \tag{5.43a}$$

$$\delta q_{4_{k+1}}^{+} = \frac{-a\left\|\delta \hat{\boldsymbol{p}}_{k+1}^{+}\right\|^{2} + f\sqrt{f^{2} + (1-a^{2})\left\|\delta \boldsymbol{p}_{k+1}^{+}\right\|^{2}}}{f^{2} + \left\|\delta \boldsymbol{p}_{k+1}^{+}\right\|^{2}} \tag{5.43b}$$

$$\hat{\boldsymbol{q}}_{k+1}^{+} = \delta \boldsymbol{q}_{k+1}^{+} \otimes \hat{\boldsymbol{q}}_{k+1}^{-}(0) \tag{5.44}$$

5.3.2 Photodiode Measurement Model

The foundation of inclusion of photodiode measurements in recursive estimation is their measurement model (the specific form of Eq. (5.1b)). The measurement model for a photodiode illuminated by only the sun is given by Eq. (5.45), where $\tilde{I}$ is the measured current output; E_{AM0} is the irradiance of the sunlight[3]; I_0 is the maximum current output of the photodiode; E_{cal} is a scaling parameter that relates the current output to the incident irradiance and is dependent on the specific photodiode and surrounding circuitry; $\mathbf{n}$ is the 3×1 unit vector that defines the direction normal to the photosensitive plane, herein referred to as the photodiode normal direction; $\mathbf{s}$ is the 3×1 unit sun vector; and η is zero mean Gaussian measurement noise [59]. Even though both $\mathbf{n}$ and $\mathbf{s}$ are unit vectors, they are not denoted with the common unit vector notation of $\hat{\mathbf{n}}$ and $\hat{\mathbf{s}}$ because the hat symbol $\hat{\cdot}$ is reserved to denote estimated quantities.

$$\tilde{I} = \frac{E_{AM0} I_0}{E_{cal}} \mathbf{n}^T \mathbf{s} + \eta \tag{5.45}$$

The quantity $\mathbf{n}^T\mathbf{s}$ is equivalent to $\cos(\theta)$, where θ is the angle between the two vectors. $\cos(\theta)$ is commonly used instead of $\mathbf{n}^T\mathbf{s}$ in photodiode models, and photodiodes are sometimes referred to as cosine detectors [71]. Phototiodes generate current anytime $0 < \theta \leq 90°$, but the output deviates from the cosine model of Eq. (5.45) at high angles. The photodiode FOV is the conical region over which the cosine model is valid. The exact FOV varies by sensor, but a half angle of 60°-70° is typical.

Photodiodes in low-Earth orbit are also subject to illumination from sunlight reflected by the Earth, which is called Earth albedo [91]. The irradiance of Earth albedo can be up to 30-40% of the solar irradiance, resulting in a potentially significant degradation of a sun vector measurement if it is not accounted for. Albedo is typically treated as noise in attitude estimation, and this approach can be successful particularly when other attitude

---

[3] The subscript AM0 denotes air mass zero, meaning zero atmospheres. E_{AM0} is the solar irradiance at 1 AU with no loss due to Earth's atmosphere.

sensors are available to complement the photodiodes. But we have found that this approach of simply increasing the measurement covariance is not sufficient for photodiode calibration because the resulting state error covariance of the calibration parameters is too high; the confidence in the calibration parameter estimates is not a significant improvement over the pre-calibration state estimates. Therefore, we include a model of albedo in the filter to remove its contribution from the measurements. The albedo model used was developed by Bhanderi using Earth reflectivity measurements from NASA's Total Ozone Mapping Spectrometer (TOMS), and the model has been verified through comparison to flight data from the Ørstead satellite [59,91,92]. The photodiode measurement model with the inclusion of Earth albedo is

$$\tilde{I} = \frac{I_0}{E_{cal}} \left(E_{AM0} \mathbf{n}^T \mathbf{s} + E_a \right) + \eta \,, \tag{5.46}$$

where the albedo irradiance E_a is given by Eq. (5.47).

Earth albedo is a function of the relative positions of the Sun, Earth, and satellite, and the Earth reflectivity varies with longitude, latitude, and atmospheric conditions such as cloud coverage. To account for the varying reflectivity in the albedo model, Earth's surface is partitioned into cells and the reflectivity of each cell is the average of measurements over a one year time period. The latitude and longitude of the center of each cell is denoted by $(\phi_g, \theta_g) \in \mathcal{D}$, where $\mathcal{D}$ is the set of all cell locations. The incident Earth albedo irradiance on a photodiode with normal direction $\mathbf{n}$ is

$$E_a = \sum_{\mathcal{V}_{sun} \cap \mathcal{V}_{sat}} E_c \left(\phi_g, \theta_g \right) \mathbf{n}^T \mathbf{r}_g \,, \tag{5.47}$$

where $\mathcal{V}_{sun} \subset \mathcal{D}$ and $\mathcal{V}_{sat} \subset \mathcal{D}$ are the cells in the field of view of the Sun and satellite, respectively; E_c is the irradiance reflected by the g-th cell in the direction of $\mathbf{r}_g$, which is a function of the reflectivity of the cell, direction of incoming solar irradiance, and direction to the satellite; and $\mathbf{r}_g$ is the vector from the satellite to the center of the cell. In this work, average Earth reflectivity data from the year 2005 is used for the Earth albedo model, which is the latest available TOMS data. A map of the average reflectivity over the one year duration is shown in Figure 5.1(a), and the standard deviation of those measurements are shown in Figure 5.1(b). Although the reflectivity is nearly 100% over some portions of the Earth, the albedo is dependent on the reflectivity as well as the angle of incoming sunlight, and the maximum albedo over the poles is between 30% and 40% of direct sun irradiance. Further details of the albedo model development and calculation of E_c can be found in the existing literature [59,92]. It is not written explicitly, but E_a is a function of spacecraft attitude; the attitude is required to compute $\boldsymbol{n}$ and $\boldsymbol{r}_g$ in the same frame for Eq.

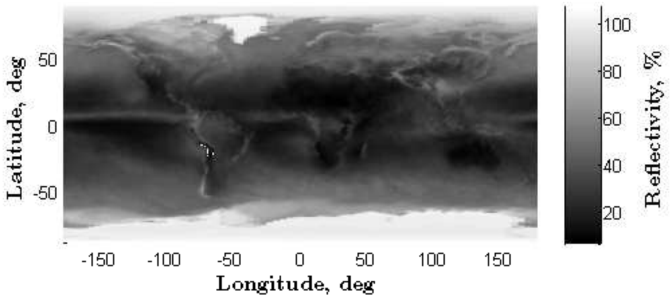

(a) Mean

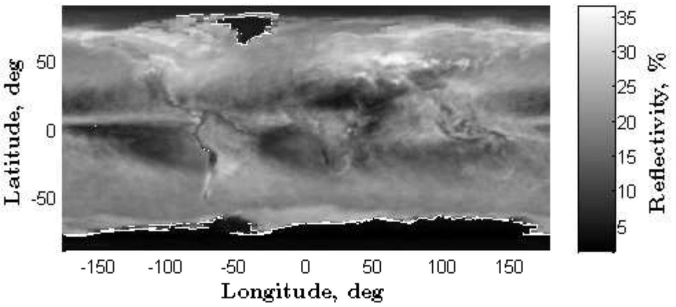

(b) Standard deviation

Figure 5.1: The mean and standard deviation of Earth reflectivity measurements from the year 2005.

(5.47).

The information available from a photodiode measurement for attitude determination is the sun vector component along the photodiode normal direction, $\mathbf{n}^T\mathbf{s}$. This quantity can also be written as a function of spacecraft attitude and the sun vector in the inertial reference frame, which is a known function of time and satellite position. This is given by Eq. (5.48), where $|_B$ indicates that the preceding quantity is resolved in the satellite body-fixed frame, $|_R$ indicates that the quantity is resolved in the inertial reference frame, and A is the 3×3 proper orthogonal attitude matrix defining the orientation of the body-fixed frame relative to the reference frame. The dependence of the photodiode normal direction and spacecraft attitude on the calculation of Earth albedo has also been noted explicitly in Eq. (5.49).

$$\left(\mathbf{n}^T\mathbf{s}\right)\big|_B = \mathbf{n}^T\big|_B \, A \, \mathbf{s}|_R \tag{5.48}$$

The measurement model that relates the photodiode measurement to spacecraft attitude (analagous in Eqs. (5.7) and (5.35) of the MEKF and USQUE) is obtained by substituting Eq. (5.48) into Eq. (5.46). The result is given by Eq. (5.49), where the j subscripts indicate that the terms are specific to the j-th photodiode and the substitution $C_j = \frac{E_{AM0} I_{0,j}}{E_{cal,j}}$ has been used. C_j is the sensor scale factor that will be estimated with on-orbit calibration. It is a dimensional scale factor that is equivalent to the maximum current output caused by irradiance of only direct sunlight.

$$\tilde{I}_j = C_j\, \mathbf{n}_j^T\big|_B\, A\, \mathbf{s}|_R + C_j \frac{E_{a,j}(\boldsymbol{n}_j, A)}{E_{AM0}} + \eta\ . \tag{5.49}$$

As discussed in Section 5.1, C_j and $\mathbf{n}_j$ in Eq. (5.49) are critical for accurate attitude determination and are the calibration parameters that will be estimated.

5.3.3 Integration into MEKF and USQUE

With the conventional attitude estimators and photodiode measurement model in-hand, the MEKF and USQUE are extended to include the photodiode calibration. MEKF and USQUE estimate six states: the 3×1 attitude error vector $\delta\boldsymbol{p}$, which in combination with quaternions being propagated within the filter, quantifies spacecraft attitude, and the 3×1 rate gyroscope bias $\boldsymbol{\beta}$. For photodiode calibration, we use these same states and include the calibration parameters as additional states. The photodiode normal directions $\mathbf{n}_j$ are parametrized by their corresponding azimuth and elevation, α_j and ϵ_j, respectively. This is given by Eq. (5.50) (there is an observability problem with this parametrization that is discussed in Section 5.3.4).

$$\mathbf{n}_j = \begin{bmatrix} \cos(\epsilon_j)\cos(\alpha_j) & \cos(\epsilon_j)\sin(\alpha_j) & \sin(\epsilon_j) \end{bmatrix}^T \tag{5.50}$$

Azimuth and elevation are used rather than the three-component normal vector because three components of a unit vector are not linearly independent. The full state vector is then

$$\mathbf{x}_{(6+3m_p)\times 1} = \begin{bmatrix} \delta\boldsymbol{p}^T & \boldsymbol{\beta}^T & \mathbf{C}^T & \boldsymbol{\alpha}^T & \boldsymbol{\epsilon}^T \end{bmatrix}^T\ , \tag{5.51}$$

where m_p is the total number of photodiodes included in the attitude determination system and

$$\mathbf{C} = \begin{bmatrix} C_1 & \dots & C_{m_p} \end{bmatrix}^T\ , \tag{5.52a}$$

$$\boldsymbol{\alpha} = \begin{bmatrix} \alpha_1 & \dots & \alpha_{m_p} \end{bmatrix}^T\ , \tag{5.52b}$$

$$\boldsymbol{\epsilon} = \begin{bmatrix} \epsilon_1 & \dots & \epsilon_{m_p} \end{bmatrix}^T\ . \tag{5.52c}$$

The azimuth and elevation of each sensor is expected to remain constant over time (assuming they are not mounted on actuated surfaces). The scale factor may decrease over time due to radiation, but since the degradation is much slower than the frequency of measurements, process noise is sufficient to capture the degradation. The dynamic models of the calibration states are therefore given by Eq. (5.53), where $\mathbf{w}_C$, $\mathbf{w}_\alpha$, and $\mathbf{w}_\epsilon$ are each mean zero Gaussian random vectors with covariance matrices Q_C, Q_α, and Q_ϵ, respectively.

$$\dot{\mathbf{C}}(t) = \mathbf{w}_C \tag{5.53a}$$

$$\dot{\boldsymbol{\alpha}}(t) = \mathbf{w}_\alpha \tag{5.53b}$$

$$\dot{\boldsymbol{\epsilon}}(t) = \mathbf{w}_\epsilon \tag{5.53c}$$

As seen in Eq. (5.4), partial derivatives of the measurement model are used in the MEKF. The partial derivatives that make up the photodiode portion of the sensitivity matrix H are

$$\frac{\partial \tilde{I}_j}{\partial \delta \boldsymbol{p}} = C_j \mathbf{n}_j^T \left[(A\, \mathbf{s}|_R) \times \right] \ , \tag{5.54}$$

$$\frac{\partial \tilde{I}_j}{\partial \boldsymbol{\beta}} = \mathbf{0}_{1\times 3} \ , \tag{5.55}$$

$$\frac{\partial \tilde{I}_j}{\partial C_j} = \mathbf{n}_j^T A\, \mathbf{s}|_R \ , \tag{5.56}$$

$$\frac{\partial \tilde{I}_j}{\partial \alpha_j} = C_j \begin{bmatrix} -\cos(\epsilon_j)\sin(\alpha_j) \\ \cos(\epsilon_j)\cos(\alpha_j) \\ 0 \end{bmatrix}^T A\, \mathbf{s}|_R \ , \tag{5.57}$$

$$\frac{\partial \tilde{I}_j}{\partial \epsilon_j} = C_j \begin{bmatrix} -\sin(\epsilon_j)\cos(\alpha_j) \\ -\sin(\epsilon_j)\sin(\alpha_j) \\ \cos(\epsilon_j) \end{bmatrix}^T A\, \mathbf{s}|_R \ . \tag{5.58}$$

Eq. (5.54) is derived using the same methods as other vector measurements in the MEKF [87]. These partial derivatives ignore the albedo contribution in Eq. (5.49), which is a reasonable approximation since the sensitivity due to direct sunlight dominates the sensitivity due to Earth albedo. These equations are used in the MEKF and not USQUE since the UF does not require linearization of the measurement model.

With the new states, measurement model, and partial derivatives in-hand, the calibration is integrated into the forms of the MEKF and USQUE presented in Section 5.3.1. The measurement model (Eq. (5.7)) and sensitivity matrix of the MEKF (Eq. (5.4)) become

$$\boldsymbol{h}_k\left(\hat{\boldsymbol{x}}_k^-\right) = \left.\begin{bmatrix} A\left(\hat{\boldsymbol{q}}^-\right)\boldsymbol{r}_1|_R \\ \vdots \\ A\left(\hat{\boldsymbol{q}}^-\right)\boldsymbol{r}_N|_R \\ \hat{C}_1^-\ \hat{\mathbf{n}}_1^{-T}\big|_B A\left(\hat{\boldsymbol{q}}^-\right)\mathbf{s}|_R + \hat{C}_1^- \frac{E_{a,1}(\hat{\boldsymbol{n}}_1, A(\hat{\boldsymbol{q}}^-))}{E_{AM0}} \\ \vdots \\ \hat{C}_{m_p}^-\ \hat{\mathbf{n}}_{m_p}^{-T}\Big|_B A\left(\hat{\boldsymbol{q}}^-\right)\mathbf{s}|_R + \hat{C}_{m_p}^- \frac{E_{a,m_p}(\hat{\boldsymbol{n}}_{m_p}, A(\hat{\boldsymbol{q}}^-))}{E_{AM0}} \end{bmatrix}\right|_{t_k} \tag{5.59}$$

and

$$H_k(\hat{\boldsymbol{x}}_k^-) = \left.\begin{bmatrix} [A(\hat{\boldsymbol{q}})\boldsymbol{r}_1\times] & 0_{3\times3} & 0_{3\times m_p} & 0_{3\times m_p} & 0_{3\times m_p} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ [A(\hat{\boldsymbol{q}})\boldsymbol{r}_N\times] & 0_{3\times3} & 0_{3\times m_p} & 0_{3\times m_p} & 0_{3\times m_p} \\ \frac{\partial\tilde{I}_1}{\partial\delta\hat{\boldsymbol{p}}} & 0_{1\times3} & \frac{\partial\tilde{I}_1}{\partial\hat{\boldsymbol{C}}} & \frac{\partial\tilde{I}_1}{\partial\hat{\boldsymbol{\alpha}}} & \frac{\partial\tilde{I}_1}{\partial\hat{\boldsymbol{\epsilon}}} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{\partial\tilde{I}_{m_p}}{\partial\delta\hat{\boldsymbol{p}}} & 0_{1\times3} & \frac{\partial\tilde{I}_{m_p}}{\partial\hat{\boldsymbol{C}}} & \frac{\partial\tilde{I}_N}{\partial\hat{\boldsymbol{\alpha}}} & \frac{\partial\tilde{I}_N}{\partial\hat{\boldsymbol{\epsilon}}} \\ & & & & \end{bmatrix}\right|_{t_k} \tag{5.60}$$

where the partial derivatives are given by Eqs. (5.54)-(5.58). Only illuminated photodiodes are used in the state update; the rows of $\boldsymbol{h}_k$ and $H_k(\hat{\boldsymbol{x}}_k^-)$ that correspond to photodiodes for which the sun is outside the FOV are omitted. Hence the dimension of $\boldsymbol{h}_k$ and $H_k(\hat{\boldsymbol{x}}_k^-)$ may be different at each update. The logic used to determine if a photodiode is sufficiently illuminated can be based on either the illumination level of the photodiode and/or the estimated attitude and photodiode normal direction combined with the reference sun vector.

From Eq. (5.53), the state propagation equations of the MEKF are extended to

$$\hat{\boldsymbol{C}}_{k+1}^- = \hat{\boldsymbol{C}}_k^+ \tag{5.61a}$$

$$\hat{\boldsymbol{\alpha}}_{k+1}^- = \hat{\boldsymbol{\alpha}}_k^+ \tag{5.61b}$$

$$\hat{\boldsymbol{\epsilon}}_{k+1}^- = \hat{\boldsymbol{\epsilon}}_k^+ \tag{5.61c}$$

and the equations for state error covariance propagation are extended to

$$\Phi = \begin{bmatrix} \Phi_{11} & \Phi_{12} & \\ \Phi_{21} & \Phi_{22} & \\ & & I_{m_p\times m_p} \end{bmatrix} \tag{5.62}$$

$$G_k = \begin{bmatrix} -I_{3\times3} & 0_{3\times3} & \\ 0_{3\times3} & I_{3\times3} & \\ & & I_{3m_p\times3m_p} \end{bmatrix} \tag{5.63}$$

$$Q_k = \begin{bmatrix} \left(\sigma_v^2\Delta t + \frac{1}{3}\sigma_u^2\Delta t^3\right) I_{3\times3} & \left(\frac{1}{2}\sigma_u^2\Delta t^2\right) I_{3\times3} & & & \\ \left(\frac{1}{2}\sigma_u^2\Delta t^2\right) I_{3\times3} & \left(\sigma_u^2\Delta t\right) I_{3\times3} & & & \\ & & Q_C & & \\ & & & Q_\alpha & \\ & & & & Q_\epsilon \end{bmatrix} \tag{5.64}$$

The UF implementation is accomplished by extending USQUE in the same manner as the MEKF was extended. There are now $12 + 6m_p$ sigma points due to the additional calibration states. The sigma points are partitioned for the additional states and propagated as in Eqs. (5.65)-(5.66).

$$\boldsymbol{\chi}_k(i) \equiv \begin{bmatrix} \boldsymbol{\chi}_k^{\delta p}(i) \\ \boldsymbol{\chi}_k^{\beta}(i) \\ \boldsymbol{\chi}_k^{C}(i) \\ \boldsymbol{\chi}_k^{\alpha}(i) \\ \boldsymbol{\chi}_k^{\epsilon}(i) \end{bmatrix}, \quad i = 0, 1, \ldots, 12 + 6m_p \tag{5.65}$$

$$\boldsymbol{\chi}_{k+1}^{C}(i) = \boldsymbol{\chi}_k^{C}(i), \quad i = 0, 1, \ldots, 12 + 6m_p \tag{5.66a}$$

$$\boldsymbol{\chi}_{k+1}^{\alpha}(i) = \boldsymbol{\chi}_k^{\alpha}(i), \quad i = 0, 1, \ldots, 12 + 6m_p \tag{5.66b}$$

$$\boldsymbol{\chi}_{k+1}^{\epsilon}(i) = \boldsymbol{\chi}_k^{\epsilon}(i), \quad i = 0, 1, \ldots, 12 + 6m_p \tag{5.66c}$$

The mean observation equations (Eqs. (5.34)-(5.35) are extended by appending the photodiode measurement equation as in Eq. (5.59), and only photodiodes for which the sun is in the FOV are used in the state update. The process covariance matrix is also extended in the same manner as for the MEKF.

With the MEKF and USQUE extended to include the calibration parameters, the process to implement each filter is the same as was described in Section 5.3.1.

5.3.4 Observability of the Photodiode Normal Directions

Azimuth and elevation angles have been used to parametrize the photodiode normal directions. The azimuth/elevation system utilized in Eq. (5.50) and subsequently in Eqs. (5.57)-(5.58) is shown in Figure 5.2(a). It is a common system where azimuth is about

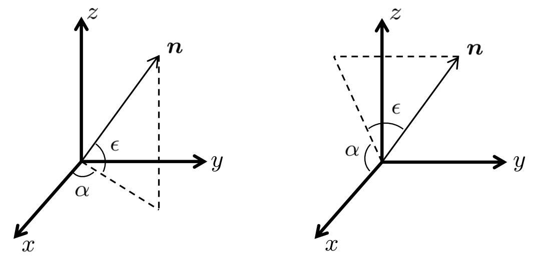

(a) Traditional az/el coordinates (b) Modified az/el coordinates

Figure 5.2: Two different azimuth/elevation systems used in the photodiode calibration.

the +z axis and elevation is measured from the x-y plane. With this system, the azimuth of $\boldsymbol{n}$ when the elevation is ±90° is undefined and thus unobservable. This prevents accurate estimation of the normal directions of photodiodes mounted parallel to the z axis. To overcome this unobservability, the azimuth/elevation system of Figure 5.2(b) is used with photodiodes that are intended to be mounted parallel or near-parallel to the z axis. In Figure 5.2(b), azimuth is about the +y axis and elevation is measured from the x-z plane. This overcomes the observability problem for photodiodes along the z axis and is utilized only for those photodiodes (it creates the same observability problem for photodiodes along the y axis). In implementation of the calibration, Eqs. (5.50), (5.57), and (5.58) are re-derived to use this system for the z photodiodes.

5.4 Application and Results

Both the UF- and EKF-based approaches to photodiode calibration have been tested through simulation, and their utility has been demonstrated by application to flight data from the RAX satellites. An overview of the RAX data used in this section is given in Section 5.4.1, and the simulation results are presented in Section 5.4.2. The calibration is applied to flight data in Section 5.4.3, and its impact on sun vector angular accuracy and the resulting attitude determination accuracy is presented in Section 5.4.4.

5.4.1 RAX Data

The calibration filter is applied to data from the RAX satellites, which were discussed in Section 2.1. The attitude determination sensors are commercial off-the-shelf components and include an Analog Devices ADIS16405 inertial measurement unit (which includes a three-axis gyroscope), a PNI MicroMag3 three-axis magnetometer, and Osram SFH2430 photodiodes. The attitude determination system and sensors are described in detail in Chapter 2.

The satellites are pictured in Figure 2.2 with their photodiodes circled. The photodiodes on RAX-1 are soldered to the side panels of the spacecraft. On RAX-2, the angled photodiodes are mounted on Delrin wedges, soldered to the solar panels, and staked to the spacecraft. No high tolerance practices were used in mounting the photodiodes; the wedges were laser cut and the photodiodes were mounted to the spacecraft by hand. Therefore, it is expected that the photodiode orientations are not well known, particularly for the angled sensors on RAX-2. The intended mounting angles of the RAX-1 and RAX-2 photodiodes are listed in Tables 2.5 and 2.6 (using the azimuth-elevation system of Figure 5.2(a)).

Both satellites utilize a passive magnetic attitude control system, a common low-cost approach to attitude control of nanosatellites (see Section 2.2). After deployment from the launch vehicle, the satellites settle to a rotation rate of 1-2 deg/s about the local magnetic field vector. The time between deployment from the launch vehicle and steady-state dynamics about the local magnetic field vector is typically referred to as *tumbling*. The RAX-1 and RAX-2 measured angular velocity following deployment into orbit was approximately 20 deg/s. In general, sensor calibration requires sufficient excitation of the sensors of interest. Since RAX spins about the magnetic field vector even when reaching steady-state dynamics, all photodiodes are typically illuminated during a single orbit, and the steady-state dynamics are sufficient for photodiode calibration. Flight data from both the tumbling and steady-state portions of flight is discussed this section. For actively controlled spacecraft, calibration maneuvers can be performed to achieve sufficient excitation for sensor and actuator calibration.

Flight data used in this section was downloaded from the spacecraft for health assessment and sensor calibration, and the calibration was performed during ground-based analysis of the flight data. A discussion of real-time calibration and calibration via post-processing the data is included in Section 5.5.

5.4.2 Simulated Testing

Before application to flight data, the calibration filters were tested and compared through application to simulated data. The simulated satellite dynamics mimic the flight data: the initial conditions are actual estimated attitude and angular rates from a portion of the flight data, and the only torque included is the control torque provided by permanent magnets, which is the dominant torque. Magnetometer and photodiode measurements are simulated by adding zero mean Gaussian noise to simulated body-frame magnetic and sun vector

Table 5.1: Sensor error parameters and initial angular velocity for the simulated data.

Parameter	**Value**
Magnetometer standard deviation	100 nT
Photodiode standard deviation	0.05 V
Photodiode FOV half-angle	70°
Angle random walk (σ_v)	4.89×10^{-4} rad/s$^{1/2}$
Rate random walk (σ_u)	3.14×10^{-5} rad/s$^{3/2}$
Initial angular rate	$\begin{bmatrix} -0.89 \\ 0.11 \\ -2.05 \end{bmatrix}$ deg/s

components, and the rate gyroscope measurements are simulated using the model of Eq. (5.19).

The sensor noise parameters and initial angular rates used in the simulations are shown in Table 5.1. The uncertainty of the photodiode measurement is given in volts because in practice, voltage rather than current is measured within the photodiode circuit. Voltage is directly proportional to current and is used for the remainder of this chapter. The corresponding photodiode angular uncertainty depends on the incidence angle and photodiode scale factor, and the output uncertainty of 0.05 V corresponds to an angular uncertainty of between 1.0° and 5.7° for a scale factor of 3.0 V (the actual scale factors of the RAX-1 and RAX-2 photodiodes are between 2.5 and 4 V, and a range between 2.5 and 3.5 V is used in the simulations). The initial angular rates are taken from a portion of the flight data after the post-deployment angular rates have dissipated and the spacecraft dynamics have settled about the geomagnetic field. Compared to the tumbling portion of the flight, these initial conditions are a worst-case for sensor excitation, and the simulations demonstrate that this is still sufficient for calibration. The sensor sampling frequency is 1 Hz for both the gyroscope and attitude sensors. The tuning parameter λ of the UF is set to $\lambda = 1$ since this value has been demonstrated to provide accurate estimates in previous applications to attitude determination [82, 86]. Additionally, like USQUE, we use $a = 1$ and $f = 4$. [86].

To compare EKF and UF performance, each filter was applied to 50 trials of simulated data with the initial state estimates set to Gaussian random vectors with mean of the true states and standard deviations given in Table 5.2. Two simulations of 50 trials each were run to compare the filters under different levels of initial condition accuracy. As seen in Table 5.2, Simulation 1 utilizes relatively poor initial estimates of the calibration states, while Simulation 2 utilizes better initial conditions. We will show in Section 5.4.3 that the better initial conditions of Simulation 2 are more representative of the initial conditions of the flight data. The simulated true values of the photodiode scale factor were set to

Table 5.2: Initial state estimate accuracy for the simulated testing.

State	**Standard Deviation**	
	Sim. 1	**Sim. 2**
δp (deg)	5	5
β (deg/s)	5	5
C (V)	0.5	0.2
α (deg)	10	2
ϵ (deg)	10	2

uniformly-distributed random values between 2.5 and 3.5 V, and the simulated true values of the photodiode orientation were set the expected states for RAX-2 shown in Table 2.6. Earth albedo was not included in the simulated testing of the two filters because simply adding the modeled albedo to the simulated measurements and compensating for it within the filter does not affect the relative performance of the two filters. The true accuracy of the albedo model compared to the albedo during the flight data sets is not known, and this is addressed in filter tuning that is discussed in Section 5.4.3.

The error in each state estimate averaged over the 50 trials is shown in Figures 5.3 and 5.4. As seen in Figure 5.3, with the initial conditions of Simulation 1, the UF provides significantly higher accuracy state estimates than the EKF. The differences in accuracy are quantified in the caption of each plot. An additional aspect that is not shown in the plots is that the EKF state estimation error falls outside of the state-error covariance for a significant portion of the states and trials, whereas the UF estimates are within bounds for all states and trials. This shows that under these initial conditions, the EKF is inconsistent and does not provide reliable state estimates.

As seen in Figure 5.4, with the better initial conditions of Simulation 2, the difference in state estimation accuracy between the two filters is negligible. Additionally, both the EKF and UF provide state estimates that are consistent with the state-error covariance. Therefore, we conclude that if the initial conditions are sufficiently accurate, the accuracy of the UF- and EKF-based calibration is the same, but under poor initial conditions, the UF outperforms the EKF. This is consistent with accuracy comparisons between the UF and EKF for various non-linear estimation problems. Since the initial states of the flight data are unknown, the accuracy of the initial conditions is also unknown, and we apply both the EKF and UF to the flight data. We will show that the initial conditions derived from the measurements are sufficient for use with the EKF.

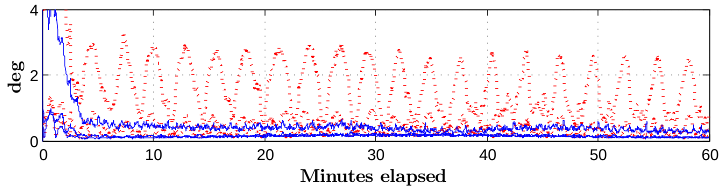

(a) Error in the three attitude components. The EKF attitude error exceeds 2°, while the UF error does not exceed 0.6° after convergence.

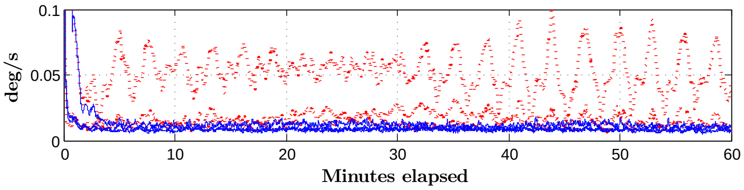

(b) Error in the three gyro bias components. The EKF error exceeds 0.05 deg/s, while the UF error does not exceed 0.02 deg/s after convergence.

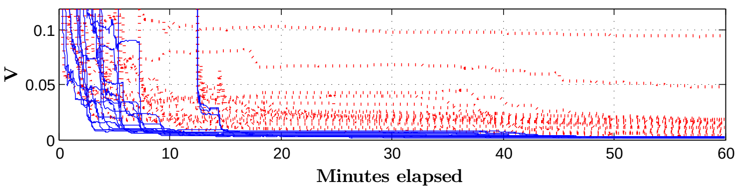

(c) Error in the 17 scale factor estimates. The worst-case EKF error is 0.09 V at the end of the simulation, whereas the worst UF error is below 3×10^{-3} V.

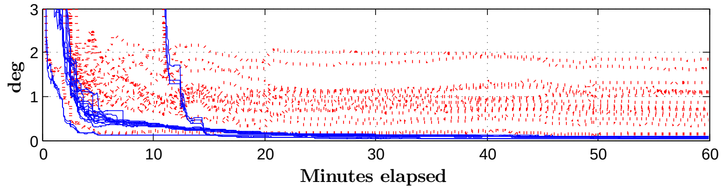

(d) Error in the 17 azimuth angle estimates. The worst-case EKF error is 1.8° at the end of the simulation, whereas the UF error is below 0.1°.

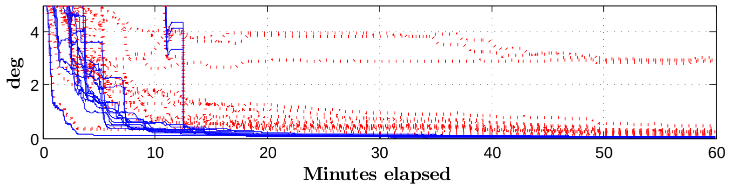

(e) Error in the 17 elevation angle estimates. The worst-case EKF error is 3.0° at the end of the simulation, wherease the UF error is below 0.1°.

Figure 5.3: The average error in each state estimate over the 50 trials with the initial state estimates of Simulation 1. In each plot, the UF estimates are shown in solid blue lines and the EKF estimates are shown in dotted red lines.

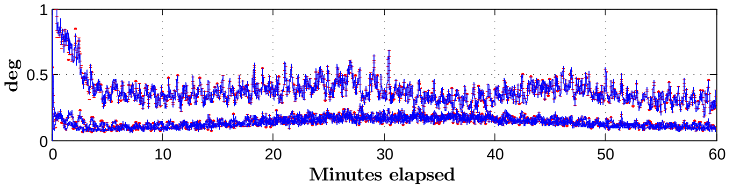

(a) Error in the three attitude components.

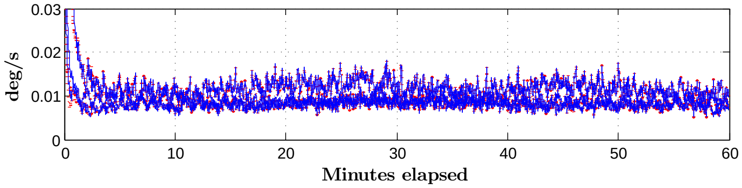

(b) Error in the three gyro bias components.

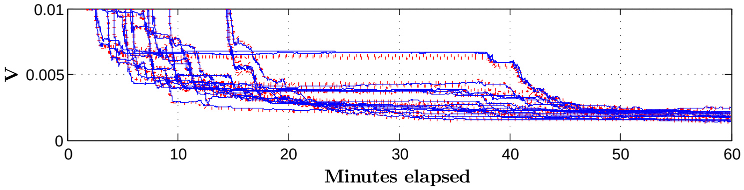

(c) Error in the 17 scale factor estimates.

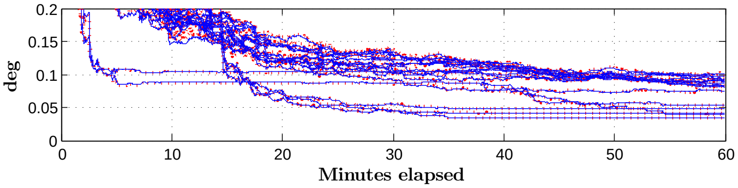

(d) Error in the 17 azimuth angle estimates.

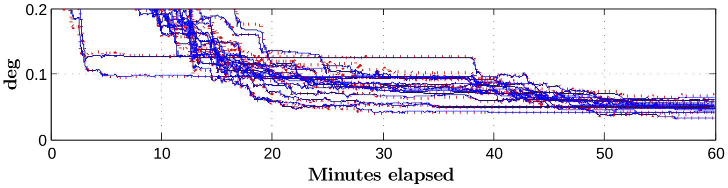

(e) Error in the 17 elevation angle estimates.

Figure 5.4: The average error in each state estimate over the 50 trials with the initial state estimates of Simulation 2. In each plot, the UF estimates are shown in solid blue lines and the EKF estimates are shown in dotted red lines, but the estimates from the two filters are not distinguishable; the difference between the EKF and UF errors is nearly zero.

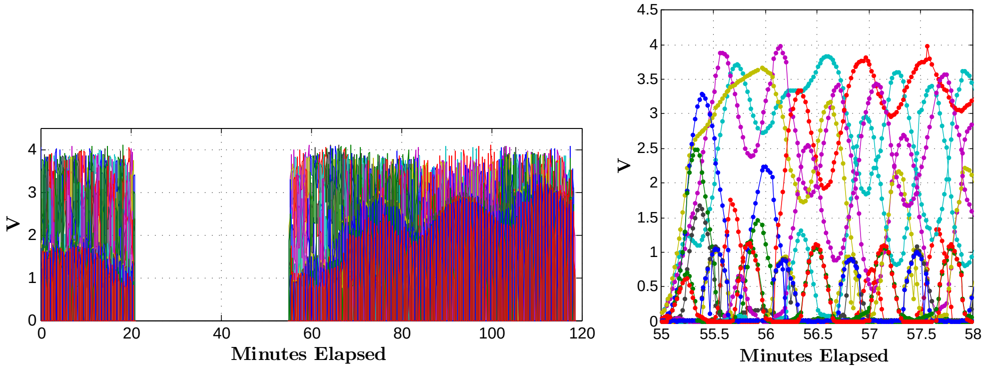

(a) Full portion of data.

(b) Zoomed-in portion.

Figure 5.5: Raw photodiode measurements (voltage) from RAX-2 data set one. For this data set, the times of the data used for calibration are 56-118 minutes.

5.4.3 Application to Flight Data

Both the UF- and EKF-based methods are applied to estimate the attitude, gyroscope bias, photodiode orientation, and photodiode scale factors using flight data from RAX-1 and RAX-2. Two data sets from RAX-1 and two data sets from RAX-2 are discussed in this section. Each data set has 1 Hz sample frequency and a duration of approximately an orbit, and only the portion of the data during which the spacecraft is in the sun is used with the filters (see Figure 5.5). The RAX-1 data sets begin Dec. 1, 2010 08:30:46 UT and Dec. 30, 2010 14:28:15 UT, and are herein referred to as RAX-1 data set one and two, respectively. The RAX-2 data sets begin Nov. 4, 2011 18:29:45 UT and Nov 12, 2011 18:33:00 UT, and are herein referred to as RAX-2 data set one and two, respectively. As a sample that is representative of each data set, photodiode measurements from RAX-2 data set one are shown in Figure 5.5.

Initial state estimates are required for both estimators. The initial estimate of the scale factor is taken to be the maximum measured output of each photodiode while the spacecraft is tumbling. Recall that this dimensional scale factors is equivalent to the maximum current output caused by direct sunlight (see Eq. (5.49)), so the maximum measured output is a reasonable initial estimate. The initial estimate of azimuth and elevation are the intended mounting angles, which are given in Tables 2.5 and 2.6. The initial attitude estimate is calculated from the measured magnetic and sun vectors using the q-method [24]. The magnetometers were first calibrated using an on-orbit, attitude-independent method [21], and the sun vector measurement for the initial attitude calculation utilizes the initial calibration parameters. The initial estimate of rate gyroscope bias is zero. The magnetic and sun reference vectors were exported from AGI's Satellite Toolkit (STK) using two-line elements and the SGP4 propagator to estimate spacecraft position.

Since the filters utilize the rate gyroscope measurements directly in the attitude propagation, an inherent assumption is that the angular rate is constant between measurements. If the sampling frequency is not sufficient to capture the dynamics, this assumption does not hold and the process covariance must be increased to account for the uncertainty in the dynamics. The data sets shown in this section are during the tumbling phase of the RAX-1 and RAX-2 missions, and the 1 Hz gyro sampling frequency is not sufficient to completely capture the dynamics. Thus the process covariance is increased to rely more heavily on the vector measurements than the gyroscope measurements. This process of adjusting the covariance is referred to as tuning. When working with flight data, filters typically need to be tuned to provide accurate state estimates. The criteria for an accurate and near-optimal filter is that the state error covariance must accurately quantify the state estimation error. During simulated testing, the true states are known (simulated) so the state error covariance P can be compared directly to the true estimation error. Since the true estimation error is unknown when using flight data and there is no other sensor available as an independent verification of the state estimates, tuning is based on the measurement residuals. The residuals are the differences between the measured vector components and the expected body-frame components calculated by rotating the reference vectors with the estimated attitude. Given the assumption of Gaussian measurement and process noise, the residuals are expected to be zero mean and within the predicted 3-σ bounds of the combined state error and measurement covariance. Starting with the initial process and measurement covariances, the covariances are then tuned to meet these criteria. The initial process covariance Q is obtained from the assumed rate gyroscope parameters and simulated testing, and the initial measurement covariance is composed of the individual sensor uncertainties. The standard deviation of the PNI magnetometer measurements after on-orbit calibration [21] is approximately 200 nT and 320 nT for RAX-1 and RAX-2, respectively, and from pre-flight testing, the approximate standard deviation of the individual photodiode measurements is 0.015 V. To account for uncertainty in the Earth albedo model, an additional uncertainty corresponding to 50 W/m^2 is added to the photodiode measurement uncertainty. This uncertainty was found to work well during filter tuning with the flight data[4]. Using the initial estimate of the scale factors, 50 W/m^2 irradiance uncertainty corresponds to a photodiode output of 0.13 V, an

---

[4] The average Earth reflectivity is 30.4%, with the variation over Earth's surface shown in Figure 5.1(a). As seen in Figure 5.1(b), the standard deviation in reflectivity measurements over a year is up to 30%. The albedo (as opposed to reflectivity) is less than this due to the high incident angles of sunlight at the poles. 30% of 30% of E_{AM0} is 123 W/m^2, and the impact of albedo also depends on spacecraft attitude. Therefore, the 50 W/m^2 that was found to work well through filter tuning is on the same order of magnitude as the expected uncertainty from variation in Earth reflectivity measurements.

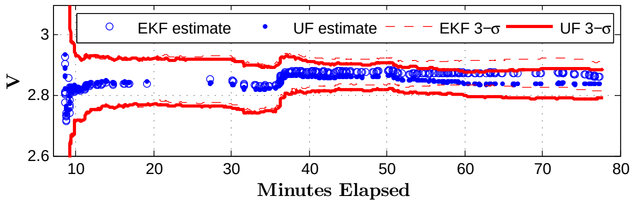

(a) $+x$ photodiode

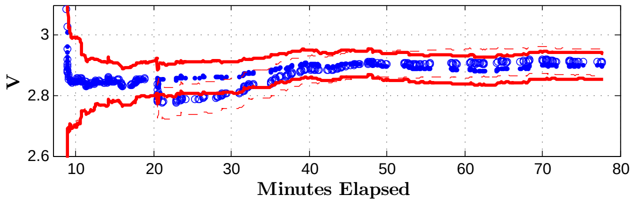

(b) $+y$ photodiode

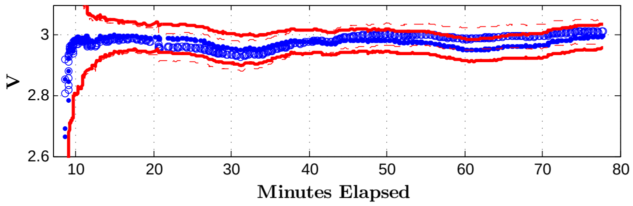

(c) $+z$ photodiode

Figure 5.6: Scale factor estimates and ± 3-σ bounds for photodiodes from RAX-1 data set one.

order of magnitude greater than the individual photodiode uncertainty.

Plots of the state estimates over time from the tuned filter applied to RAX-1 and RAX-2 data are shown in Figures 5.6-5.11. These show the calibration parameter estimates for three photodiodes in each data set, and the behavior is representative of all the photodiodes on each spacecraft. The state estimates resulting from both the UF and EKF, as well as their 3-σ bounds from the state error covariance, are shown. In each plot, the estimate is shown only when a measurement is available for the state update.

In the RAX-1 state estimates (Figures 5.6-5.8), we see that there is a discontinuity in the state estimates at approximately 20 minutes elapsed. This is due to a pause in the nominal 1 Hz sampling frequency and gap in the measurements[5]. The EKF takes longer to re-converge than the UF for some states, which is particularly evident in Figure 5.7(b), but both filters converge to consistent state estimates by 50 minutes elapsed (20 minutes after the discontinuity). The scale factor estimates are not exactly constant; oscillations in

---

[5] This pause is specific to the RAX flight data and is an artifact of the flight software. The pause in measurements is when the spacecraft is writing the data to flash memory.

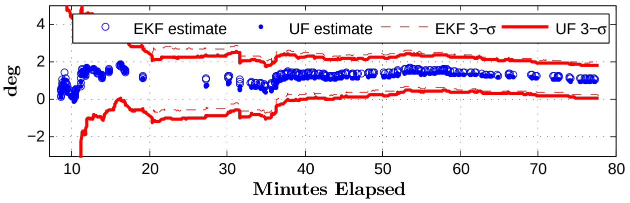

(a) $+x$ photodiode

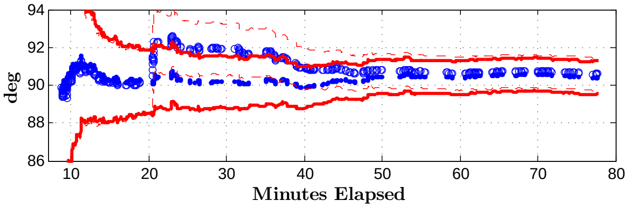

(b) $+y$ photodiode

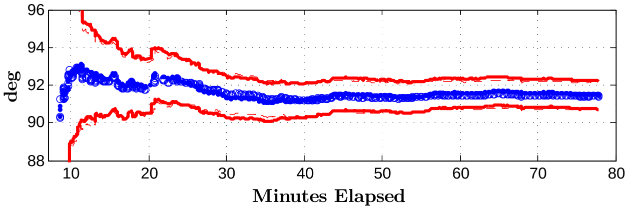

(c) $+z$ photodiode

Figure 5.7: Azimuth estimates and ± 3-σ bounds for photodiodes from RAX-1 data set one.

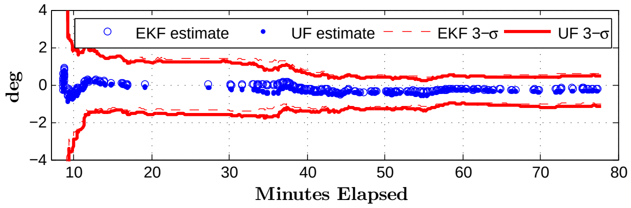

(a) $+x$ photodiode

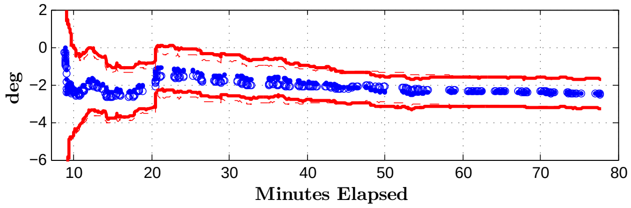

(b) $+y$ photodiode

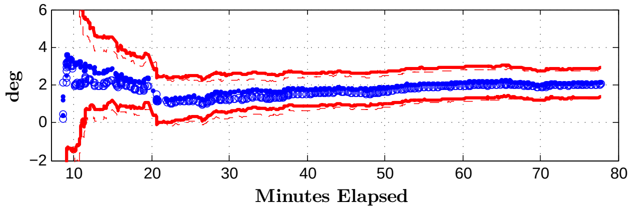

(c) $+z$ photodiode

Figure 5.8: Elevation estimates and ± 3-σ bounds for photodiodes from RAX-1 data set one.

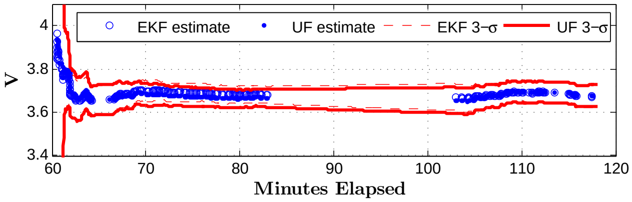

(a) Photodiode #1

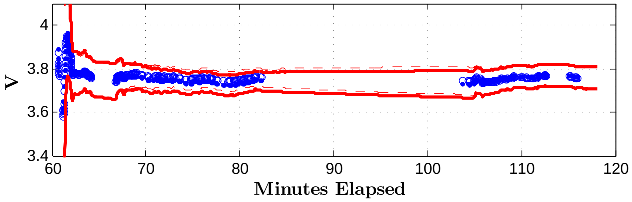

(b) Photodiode #2

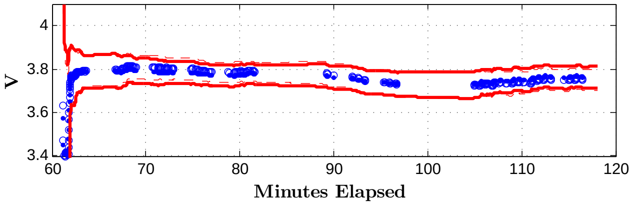

(c) Photodiode #3

Figure 5.9: Scale factor estimates and ±3-σ bounds for photodiodes from RAX-2 data set one.

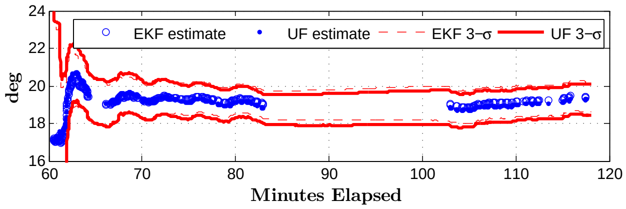

(a) Photodiode #1

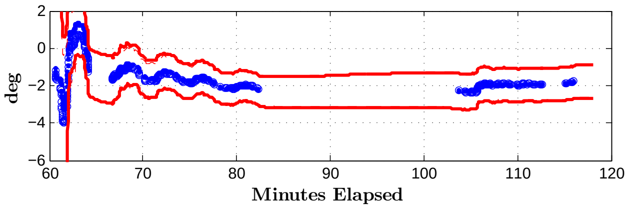

(b) Photodiode #2

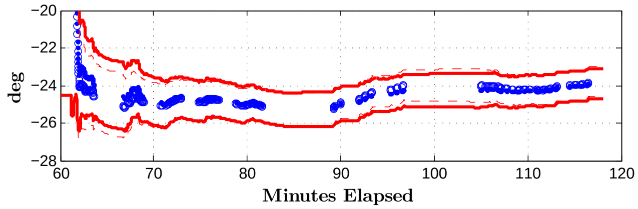

(c) Photodiode #3

Figure 5.10: Azimuth estimates and ±3-σ bounds for photodiodes from RAX-2 data set one.

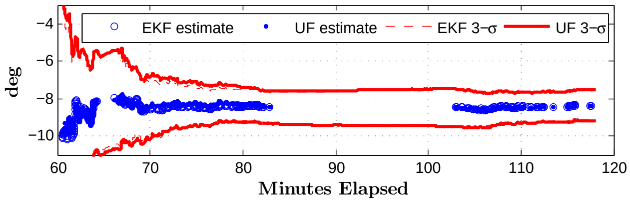

(a) Photodiode #1

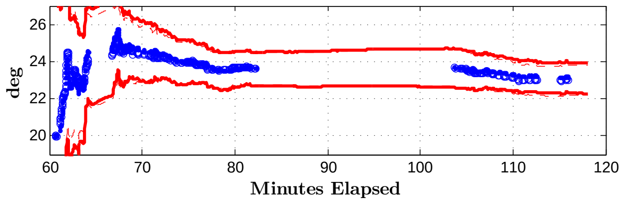

(b) Photodiode #2

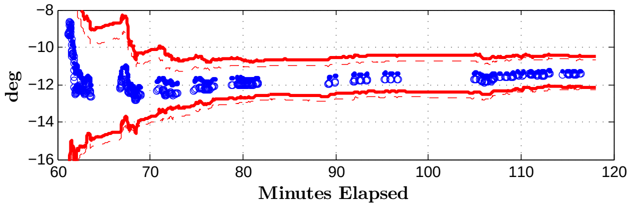

(c) Photodiode #3

Figure 5.11: Elevation estimates and ±3-σ bounds for photodiodes from RAX-2 data set one.

the estimates even after convergence are due to the uncertainty in the Earth albedo model[6]. From the plots, we see that when the filters are converged, the differences in the state estimates of the EKF and UF are negligible.

In the RAX-2 state estimates (Figure 5.9-5.11), we see that the estimates converge to a near-constant value at approximately 80 minutes elapsed time (20 minutes into the sun-illuminated portion of the data). Like the RAX-1 estimates, the differences between the EKF- and UF-based state estimates is negligible, and a slightly faster convergence time is evident in the UF estimates.

In each plot, we see that the estimates converge to a consistent values. This, as well as the residuals criteria discussed previously, indicate that the filter is providing accurate state estimates. We use two metrics as an additional verification of the accuracy of the state estimates. First, comparing the estimates of the azimuth and elevation from the different data sets from each satellite, which we expect to be constant over time, the estimated 3-σ bounds overlap, which demonstrates consistency in the estimates. Additionally, the effectiveness of the calibration can be seen in comparing the measured sun vector magnitude using the initial and final calibration parameters, which is an attitude-independent verification metric. The magnitude of the measured sun vector should be one. Histograms of this vector magnitude[7] from the RAX-2 data sets[8], which was calculated with the albedo-corrected photodiode measurements and both the initial and calibrated photodiode parameters, are shown in Figures 5.12-5.13. In the first data set, the mean vector magnitude improves from 0.925 to 0.993, and the standard deviation decreases from 0.029 to 0.022 when using the on-orbit-estimated parameters rather than initial parameters. In the second data set, the mean improves from 0.932 to 0.999 and the standard deviation decreases from 0.032 to 0.021.

The importance of including Earth albedo in the calibration is evident by comparing the magnitude of the measured sun vector calculated with the raw photodiode measurements to the magnitude calculated with albedo-compensated measurements. This is shown in Figure 5.14 for the RAX-2 data sets. There are clear deterministic trends in the magnitude

---

[6] The deviations from a constant estimate could also be due to temperature dependence, but from a visual comparison of the measurements of nearby temperature sensors to the state estimates, consistent correlation between scale factor estimate and temperature was not evident. A thorough characterization of temperature dependence of these photodiodes was not preformed pre-flight, and this is an area that could be addressed in future work.

[7] The vector magnitude does not have a Gaussian distribution. The sum of the squares of components with Gaussian distributions has a chi-square distribution.

[8] RAX-2 data is used to show the attitude-independent verification metrics because a three-component sun vector is required for the attitude-independent metrics, and with the angled photoidoes on RAX-2, sun vectors can be calculated for a much higher percentage of data than RAX-1. But the improvement in vector magnitude is evident in all data sets.

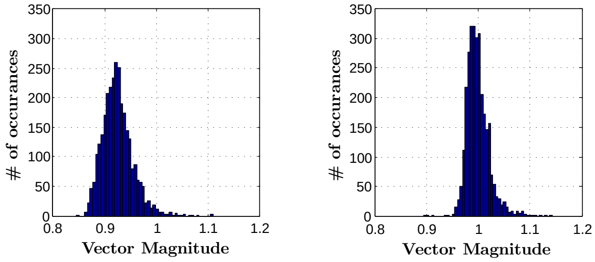

Figure 5.12: Histograms of the magnitude of the measured sun vector from RAX-1 data set one using the initial (left) and calibrated (right) sensor parameters.

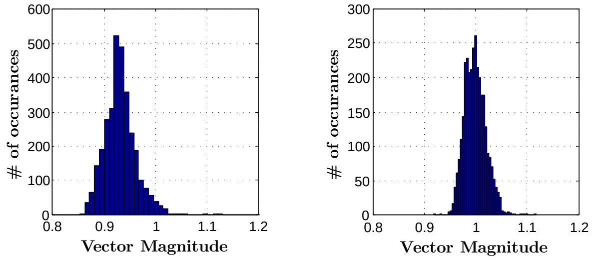

Figure 5.13: Histograms of the magnitude of the measured sun vector from RAX-2 data set two using the initial (left) and calibrated (right) sensor parameters.

calculated from the uncompensated measurements that have been removed by subtracting the modeled albedo contribution from the measurements. Since the albedo contribution is attitude-dependent, this demonstrates the need for simultaneous attitude estimation and calibration as opposed to an attitude-independent calibration. These trends were also seen in an initial attitude-independent attempt at calibrating the RAX-1 photodiodes [80], which required three simultaneously-illuminated sensors and did not account for Earth albedo. The new recursive method presented in this paper overcomes both the adverse affect of albedo and the need for multiple illuminated sensors in the calibration.

5.4.4 Impact on Sun Vector Angular Accuracy and Attitude Determination

The difference between the initial calibration parameters – which were the intended mounting angles and the maximum measured output of each photodiode while the spacecraft was tumbling – and the estimates from on-orbit calibration is significant. In the RAX-2 data, the improvement in elevation and azimuth estimates ranged from 0° - 9°, and the improvement in scale factor ranged from 0.17 - 0.46 V, which is equivalent to 4.4% - 12.7% of the calibrated scale factor. For RAX-1, the azimuth and elevation improvement ranged from

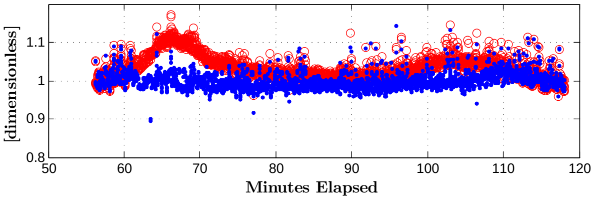

(a) Data set one

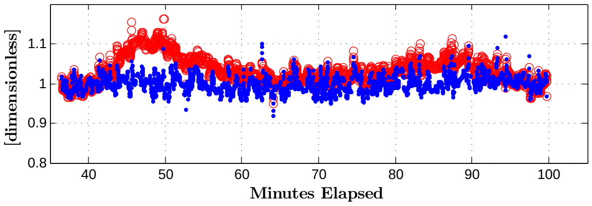

(b) Data set two

Figure 5.14: Magnitude of the measured sun vector over time for the two RAX-2 data sets using both the uncompensated photodiode measurements (red circles) and the albedo-compensated measurements (blue dots). These magnitudes were calculated with the estimated calibration parameters. The albedo-compensated series of data is the same data as the calibrated histograms of Figures 5.12-5.13.

Table 5.3: Scale factor estimates from the RAX-1 data sets.

Photodiode number	**1**	**2**	**3**	**4**	**5**
Maximum measurement (V)	3.02	3.12	3.00	2.98	3.15
Data set one estimate (V)	2.84	2.95	2.90	2.89	3.00
Data set two estimate (V)	2.76	2.92	2.81	2.79	2.86

0° - 4°, and the improvement in scale factor ranged from 0.09 - 0.29 V, which is equivalent to 3.4% - 10.2% of the calibrated scale factor. The scale factor estimates from the RAX-1 data sets are shown in Table 5.3[9]. A decrease in the photodiode scale factor as the satellite spends more time on-orbit is evident. This is hypothesized to be due to ultra-violet (UV) radiation. Photodiodes are known to degraded due to UV radiation, and no protective measures were taken to prevent this on RAX-1. This highlights the utility of the calibration in tracking sensor parameters on-orbit. On RAX-2, solar cell coverglass was applied to the photodiodes, and the degradation of RAX-1 is not evident.

Histograms of angular differences between the measured sun vector using the initial and calibrated parameters for the RAX-2 data sets are shown in Figure 5.15. The mean improvement in sun vector angular accuracy is 9.1° and 10.2° from the two data sets, which

---

[9] There are nine photodiodes on RAX-1, but only five are shown in the table. This is because measurements from the redundant photodiodes on the z surfaces were not downloaded for analysis. Additionally, the -z photodiode is shadowed for a significant portion of the measurements and is also affected by reflections from the antenna when it is not shadowed, and therefore, its measurements were not included in the calibration.

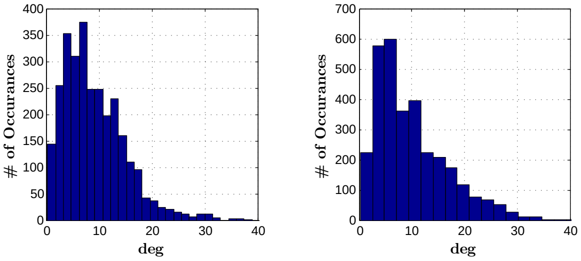

(a) RAX-2 data set one (b) RAX-2 data set two

Figure 5.15: Histograms of the angular difference between the measured sun vector using the initial and calibrated photodiode parameters from the RAX-2 data sets. The mean from data set one is 9.1° and the mean from data set two is 10.2°.

is quite significant given that the angular accuracy of photodiodes is on the order of degrees. The poor alignment of the photodiodes is not surprising given the manual process used to integrate them to the spacecraft. The mean improvement in sun vector angular accuracy for the RAX-1 data sets is 5.6° and 6.4°. In application to other spacecraft, the degree of improvement resulting from the on-orbit calibration will be dependent on the quality of the pre-flight calibration parameters.

The 1-σ attitude estimation accuracy during the simultaneous attitude estimation and photodiode calibration from RAX-1 data set one and RAX-2 data set one is shown in Figure 5.16. The three components are the uncertainty in rotation about the x, y, and z axes of the body-fixed frame. We see that the accuracy of the three components is better than 2° 1-σ and 1° 1-σ for RAX-1 and RAX-2, respectively, for most of the duration that the spacecraft is in the sun. The higher uncertainty between 90-100 minutes for RAX-2 is caused by the relative alignment of the sun and magnetic vectors in the body frame. During this time period, the vectors are nearly parallel, significantly reducing the amount of information available for attitude determination. This condition does not arise during the RAX-1 data set (the spacecraft are in different orbits). Both of these data sets are taken during the tumbling portion of the flight, but the angular velocities of RAX-1 and RAX-2 during these data sets is nearly the same and the same process covariance is used for each data set. The cause of the better attitude accuracy of RAX-2 compared to RAX-1 is the additional photodiodes in the attitude determination system.

In the data sets used so far, the 1 Hz gyro sampling frequency is not sufficient to capture the spacecraft dynamics, and the process covariance matrix is inflated to account for the dynamics uncertainty as discussed in Section 5.4.3. Higher accuracy attitude estimation is achieved when the gyro measurements sufficiently capture the dynamics and the process covariance does not have to be increased significantly beyond that of Eqs. (5.18) and (5.23).

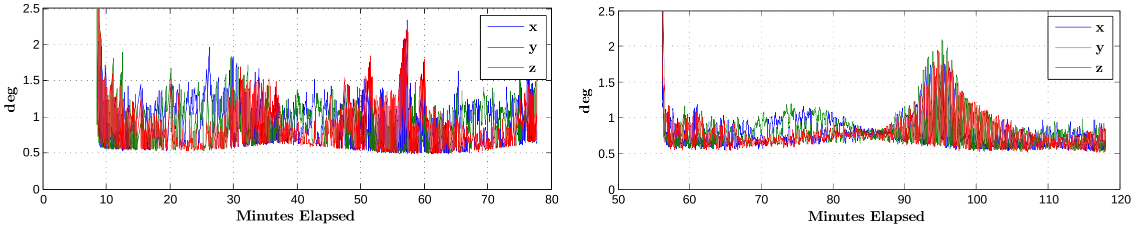

(a) RAX-1 data set one

(b) RAX-2 data set one

Figure 5.16: Attitude accuracy (1-σ) achieved in applying the attitude estimation/calibration filter to data sets one and two.

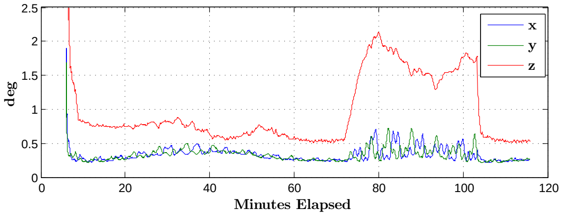

Figure 5.17: Attitude accuracy of a third RAX-2 data set, which was collected December 9, 2011 16:00:00 UTC.

This is demonstrated by Figure 5.17, which shows the attitude accuracy during a data set from later in the RAX-2 mission (Dec 9, 2011) after the passive magnetic control system had dampened the rotational kinetic energy and aligned the spacecraft z-axis with Earth's magnetic field. We see that the 1-σ attitude accuracy in the x- and y-axes is better than 0.5° when the spacecraft is in the sun. The z-axis (spin axis) has the least accuracy because the spacecraft is spinning about the magnetic vector, which minimizes the amount of non-redundant information that the gyroscope and magnetometer provide. The decrease in accuracy of all three components at 72-105 minutes is when the spacecraft enters eclipse and the photodiode measurements are no longer available. This plot demonstrates the full potential of the low-cost, commercial-off-the-shelf attitude sensors used in conjunction with the on-orbit calibration methods of this paper. Overall attitude determination results are discussed further in Chapter 6.

5.5 Discussion

We have utilized a three-axis rate gyroscope and three-axis magnetometer in addition to the photodiodes for the photodiode calibration. Since the calibration is a recursive method, either a rate gyroscope or a dynamic model is required for fusion with the vector measurements. Use of a rate gyroscope instead of a dynamic model has become a common approach for attitude estimation [25], but use of a dynamic model with no rate gyroscope has also been demonstrated [93, 94] and can be utilized for photodiode calibration with the formulation presented in Section 5.3. The calibration can also be accomplished without a magnetometer or other vector sensor, but the state estimation accuracy will be worse than with the additional attitude sensor since the sun sensors alone provide only two axis information (spin about the sun vector is the third axis). We included the magnetometer since it is available and results in significantly better accuracy compared to that using only the photodiodes. We have done initial testing of the calibration filter with only the gyroscope and photodiode measurements and have demonstrated convergence, but a thorough study on the accuracy when using photodiodes as the only vector measurement is left for future work.

Application of the EKF- and UF-based calibration to both simulated and flight data showed that the UF is required when the initial state estimates are relatively poor, but the accuracy of the EKF matches that of the UF if the initial state estimates are sufficiently accurate. The initial state estimates used with the flight data – which for the attitude and photodiode scale factors, were derived from the measurements, and for the photodiode orientation, were assumed to be the designed values – were sufficiently accurate for use with the EKF. Therefore, it is reasonable to assume that the EKF is sufficient for photodiode calibration with other spacecraft of comparable or better specifications. This would facilitate real-time implementation, whereas the UF is not well-suited for real-time implementation on small spacecraft without significant parallel-processing capability. In general, if the processing power is available to facilitate the UF, such as during ground-based analysis, the UF should be used over the EKF since its performance is better over a broader range of initial conditions.

In application of this calibration method to RAX, the flight data was downloaded from the spacecraft and processed on the ground. Real-time attitude estimation is not required on RAX, and downloading batches of data periodically throughout the mission is part of normal RAX operations. Nonetheless, extended Kalman filters are well-suited for real-time implementation. Murrell's version of the EKF can be used to reduce the size of matrices required for inversion to 3×3 and discrete attitude propagation can be used to reduce the computational requirements [88]. The only aspect of the EKF-based calibration method that is not well suited for on-line implementation in its current form is the Earth albedo model. For on-line implementation, the albedo model could be simplified [95] or precalculated and tabulated on-board. Investigation into these options is left for future work. Off-line calibration can be advantageous since it allows for a thorough inspection of the measurement residuals and tuning parameters, and real-time sensor correction can be accomplished even with off-line calibration by uploading the calibration parameters to the spacecraft.

Two other areas for further investigation include the degradation in state estimation accuracy as the number of photodiodes increase, and use of a colored noise model. The maximum number of photodiodes utilized for the calibration was 17, dictated by the number of photodiodes used on RAX-2. It is reasonable to assume that increasing the number of photodiodes will begin to degrade the accuracy of the state estimates due to information dilution. Additionally, given that the uncertainty in Earth albedo varies over Earth's surface (see Figure 5.1(b)), a colored noise model or a Gaussian noise model with time-varying covariance may provide improved results instead of the Gaussian noise model of constant covariance used in this work. This was not pursued since the current implementation already provides a significant improvement in sun vector estimation compared to the use of pre-flight calibration. Both of these areas of study are left for future work.

In addition to attitude determination, the photodiode calibration can be used to track the orientation of actuated surfaces on a spacecraft. For example, if a photodiodes are placed on actuated solar panels, then the filter presented in this paper can be used to estimate the orientation of the actuated panels relative to the body-mounted attitude sensors.

5.6 Conclusion

We have developed a method for on-orbit photodiode calibration to estimate the orientations and scale factors of photodiodes in an attitude determination system. The calibration has been formulated for use with either an extended Kalman filter or an unscented filter to simultaneously estimate spacecraft attitude and the calibration parameters, and it can be applied to any number of photodiodes in an arbitrary configuration on the spacecraft. The importance of the attitude-dependent approach and the use of an Earth albedo model was demonstrated. In application to RAX-2, which utilizes photodiodes, magnetometers, and a three-axis rate gyroscope for attitude determination, the calibration improved the accuracy of the measured sun vector by an average of 10°. This calibration enables the most accurate performance of the attitude determination system with the given hardware. With the combination of calibrated photodiodes as well as a low-cost magnetometer and gyroscope, attitude accuracies of better than 1° 1-σ have been demonstrated. Overall attitude determination results using are discussed further in Chapter 6.

CHAPTER 6

Resulting Attitude Determination Accuracy and the New Approach to Attitude Determination

In this chapter, we discuss the attitude determination accuracy of the RAX ADS that is achieved as a result of the new on-orbit sensor calibration methods, and we summarize the new approach to low-cost attitude determination that results from this dissertation. The chapter is organized as follows. First, in Section 6.1, we summarize the methods for sensor calibration and estimation of relative sensor alignment. The methods used and resulting accuracies of the RAX ADS (described in Chapter 2) are then presented in Section 6.2, and the approach to attitude determination is summarized in Section 6.3.

6.1 Sensor Calibration and Alignment

New methods for on-orbit calibration of magnetometers and photodiodes were the topics of Chapter 3 and 5, and these methods have been applied for on-orbit calibration of these sensors on both RAX-1 and RAX-2. The only RAX attitude sensor for which calibration has not yet been discussed is the three-axis rate gyroscope. That, as well as relative sensor alignments, are discussed in this section.

Three-axis rate gyroscope measurements are subject to drifting bias as well as scale factor and non-orthogonality errors. Drifting bias is inherent to all inertial sensors and the drift rate characteristics are a function of the quality of the sensor. For example, the gyroscope of Figure 1.1(b) has a bias stability that is a factor of 5040 better than that of the MEMS gyro of Figure 1.2(b). As discussed briefly in Section 5.3.1.1, a commonly-used model for gyroscope drift is given by Eq. (5.19) [89], and this model can be utilized for estimation of the gyro bias within an attitude estimator such as the MEKF. Gyro scale factor and non-orthogonality errors (analogous to the magnetometer scale factor and non-orthogonality error described in Section 3.2), can be estimated via sensor calibration.

Various on-orbit rate gyroscope calibration methods exist to estimate and compensate for these errors [81, 82, 96], but no on-orbit gyroscope calibration was applied to the RAX data. Pre-flight calibration[1] was used to estimate the gyroscope scale factors, and no scale factor errors were detected [31]. Additionally, since the magnetometer and gyroscope are packaged together in the ADIS16405, we assume that the non-orthogonality error of the gyroscope is the same as the magnetometer, which is estimated with the on-orbit magnetometer calibration. This is a reasonable assumption because within the ADIS16405, the channels of the two sensors are mounted to the same structure. After application of the MEKF for attitude estimation (discussed in Section 6.2), no additional gyroscope errors are evident, and any unmodeled non-orthogonality or scale factor errors are captured in the process noise.

The on-orbit magnetometer calibration method of Chapter 3 is carried out in an orthogonal sensor-fixed frame, which in Chapter 3, was referred to as the rectified frame. Relative alignment of attitude sensors is also critical for accurate attitude estimation. After magnetometer calibration, which resolves the non-orthogonality error in the rectified frame, the relative alignment of two magnetometers can be estimated by applying any single-point attitude estimation method, such as the q-method [24], to a batch of measurements from both magnetometers. In application of the attitude estimation method, one magnetometer-fixed rectified frame is treated as the reference frame, and the other is treated as the measurement frame. The attitude is then the orientation of the measurement frame relative to the reference frame. Using this method, the estimated attitude of the PNI magnetometer relative to the ADIS16405 magnetometer on RAX-2 is

$$\hat{A}_{PNI/IMU} = \begin{bmatrix} 0.9996 & 0.0090 & 0.0280 \\ -0.0083 & 0.9997 & -0.0249 \\ -0.0282 & 0.0247 & 0.9993 \end{bmatrix}, \tag{6.1}$$

which corresponds to rotation angles of 0.5°, 1.6°, and -1.4° about the x, y, and z axes. For comparison, the estimated non-orthogonality errors of the PNI magnetometer are -0.5°, -4.0°, and 2.2°, and the estimated non-orthogonality errors of the ADIS16405 magnetometer are -0.6°, -2.4°, and -1.7°. From this analysis, we see that angular errors of up to 4° were present in the RAX-2 magnetometers, due to both sensor non-orthogonality and relative mis-alignments. These errors have been estimated through on-orbit calibration and are subsequently used to correct the magnetometer measurements before use in attitude

---

[1] This relatively-simple calibration procedure consists of collecting rate gyro measurements with the sensor on a rate table and spinning at a series of known rates. Linear least-squares minimization is then used to estimate the scale factor relating the measured and true angular rates [31].

determination.

In the photodiode calibration method of Chapter 5, the orientation of each photodiode is estimated relative to the other attitude determination sensors. The relative alignment of these other sensors is assumed to be known during the photodiode calibration. For the RAX data, the relative magnetometer alignment was discussed in the previous paragraph, and since one of the magnetometers and gyroscopes are packaged together, they are also assumed to be aligned. If this is not the case, an existing on-orbit rate gyroscope calibration method can be used to estimate the relative alignment between the gyroscope and other sensors [81, 82, 96], and then the photodiode calibration is carried out within the frames of the other sensors.

In summary, the procedure to calibrate and align each sensor using only on-orbit measurements is the following.

1. Calibrate the magnetometers using the method of Chapter 3. When complete, the relative alignment of multiple magnetometers can be estimated from a batch of measurements using any single-point attitude method, such as the q-method [24].

2. If needed, calibrate and estimate the alignment of the gyro relative to the magnetometer using one of the existing gyroscope calibration and alignment methods [81,82,96].

3. Calibrate the photodiodes, which includes estimating the alignment of the photodiodes relative to the other sensors, using the method of Chapter 5.

Attitude can then be estimated from the calibrated sensor measurements. This attitude is that of the common attitude sensor frame relative to the reference frame. The relative alignment between the common sensor frame and any other frame, such as a separate user-defined satellite body-fixed frame or a payload-fixed frame, has not been addressed in this work. In many cases, the relative alignment of the attitude sensors is sufficient to complete the mission [52], and both relative and absolute alignment of body-fixed frame are discussed further in References 52, 53. Attitude estimation results from the calibrated RAX-1 and RAX-2 sensor measurements are discussed in the next section.

6.2 Attitude Estimation Methods and Results

With on-orbit sensor calibration complete, the rate gyroscope, magnetometer, and photodiode measurements are used to estimate satellite attitude. Two methods have been used for attitude estimation. The first method is the photodiode calibration filter of Chapter 5, which estimates the photodiode calibration parameters and spacecraft attitude simultaneously, and the second is the conventional 6-state MEKF [25, 87], which was discussed in Section 5.3.1.1. Of the many methods available for attitude estimation, the MEKF has become a widely-used approach [24]. The MEKF estimates six states: three components of attitude and three components of rate gyroscope bias. This filter alleviates the difficulties of modeling spacecraft dynamics by using bias-corrected rate gyroscope measurements directly in the attitude propagation, and it maintains attitude orthogonality constraints through a multiplicative quaternion update. Although attitude estimation for RAX is done off-line, the MEKF is well suited for real-time implementation [97], and we utilize the MEKF even for off-line processing so that the results are comparable to what can be achieved during real-time estimation.

In general, non-linear recursive estimators require initial conditions for the state estimates. The attitude is initialized to the estimate provided by a single-point estimator using the magnetometer and photodiode measurements. Since both directional measurements are required, the estimator is initialized at a time when the spacecraft is in the sun so the sun vector measurement is available. The gyro bias states are initialized to zero, and the initialization of the photodiode calibration states was discussed in Section 5.4.3.

Attitude determination accuracy is quantified by the state error covariance of the tuned filter. Filter tuning is the process of adjusting the process and measurement covariance to produce accurate and consistent results, with consistent results meaning that the state error covariance accurately quantifies the state estimation error. During simulated testing, the true states are known (simulated) so the state error covariance can be compared directly to the true estimation error. When working with flight data, the true error is unknown, and tuning is based on the measurement residuals. The residuals are the difference between the measured magnetic and sun vector components and the corresponding reference components after rotation into the body-fixed frame with the estimated attitude. Given the assumption of Gaussian measurement and process noise, the residuals are expected to be zero mean and accurately bounded by the measurement and process covariance. In application to flight data, the covariances are tuned to meet this criteria.

Throughout this section, we discuss accuracy in terms of the 1-σ (one standard deviation) bounds on the three attitude components, as well as through an approximation of the "total" attitude uncertainty. The 1-σ bounds are taken directly from the diagonal elements of the attitude portion of the state error covariance matrix of the filters and, under the small angle approximation, quantify uncertainty in rotation about each axis of the body-fixed frame. Although no single number adequately represents attitude error (the uncertainty is generally not identical in all directions), a single number representing total attitude error is often desired. The principle axis of the error ellipsoid provides a conservative approximation of this total attitude error [98]. The total attitude error has a chi-square probability distribution, and thus is discussed in terms of confidence bounds rather then 1-σ or 3-σ bounds that are typically used to quantify a Gaussian distribution. The 68% and 99% confidence bounds used to provide an approximate total attitude uncertainty are derived from the attitude state error covariance; they are given by 1.872 and 3.368 times the square root of the maximum element of the diagonalized covariance matrix (the major axis of the error ellipsoid), respectively [98].

The attitude estimation accuracy from the earliest available RAX-1 and RAX-2 data sets are shown in Figures 6.1-6.2. These results are from application of the simultaneous photodiode calibration and attitude estimation EKF. The RAX-1 data was recorded 12 days after launch, and the RAX-2 data was recorded seven days after launch. Both data sets are 1 Hz measurements from the tumbling phase of the mission. Since the EKF uses bias-corrected rate gyroscope measurements directly to propagate attitude, an inherent assumption is that the angular velocity of the spacecraft is constant between measurements. The 1 Hz gyro sampling frequency is not sufficient to validate this assumption during the tumbling phase of the mission, so the process covariance is increased to rely more heavily on the directional measurements. The angular velocities of the spacecraft during both data sets, shown in Figure 6.3, are comparable such that the same process covariance is used for both data sets. From comparison of Figures 6.1(a) and 6.2(a), we see that the uncertainty in estimating the RAX-1 attitude is greater than the uncertainty during the RAX-2 data set. The 1-σ accuracy of the RAX-1 data varies between 0.5° and 1.5°, where as the 1-σ accuracy of the RAX-2 data varies between 0.5° and 1.0° for the majority of the data set. This difference is due to the increased number of photodiodes in the RAX-2 ADS (see Section 2.5). In both data sets, the dramatic increase in uncertainty for approximately 35 minutes is when the spacecraft is in eclipse and the photodiodes are not illuminated. The decrease in accuracy in the RAX-2 data around 95 minutes corresponds to a satellite position where the magnetic and sun vectors are approximately parallel, minimizing the amount of information provided by the sensors.

Attitude measurements taken after the satellite has reached steady-state dynamics about the geomagnetic field demonstrate the full potential of the attitude determination system (it takes a few weeks for the spacecraft to dissipate rotational kinetic energy and settle to steady-state dynamics; see Section 5.4.1). During this time, as opposed to the tumbling phase of the mission, the 1 Hz sampling frequency more accurately captures the dynamics and direct gyro-based attitude propagation within the filter can be carried out with a smaller process covariance. This process covariance is dictated largely by the noise characteristics

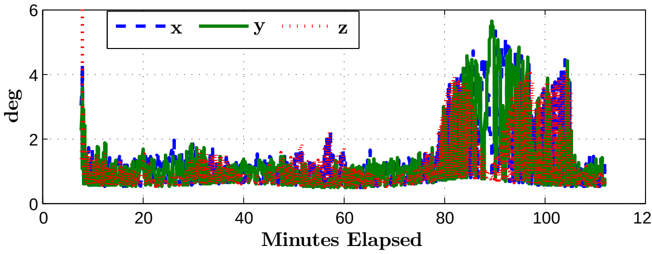

(a) 1-σ bounds.

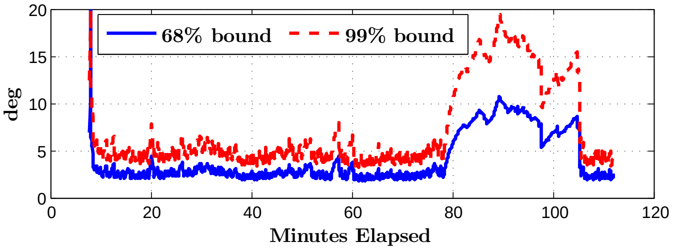

(b) Approximate total accuracy.

Figure 6.1: Attitude estimation accuracy of the simultaneous photodiode calibration and attitude estimation filter using RAX-1 measurements collected Dec. 1, 2010 08:30:46 UT.

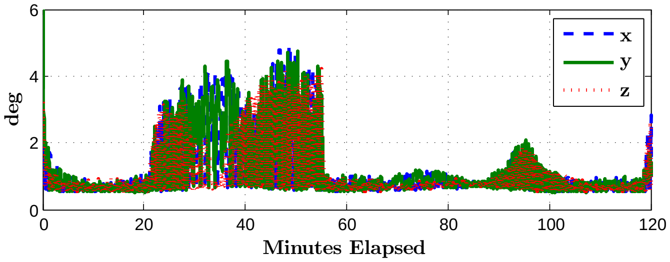

(a) 1-σ bounds.

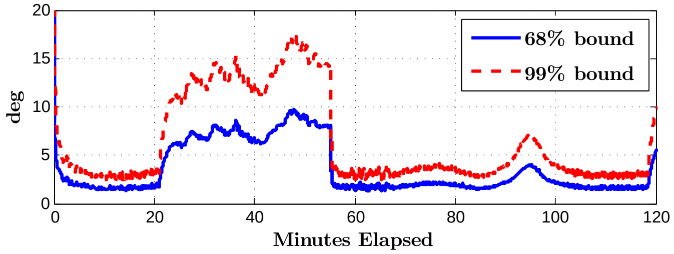

(b) Approximate overall accuracy.

Figure 6.2: Attitude estimation accuracy of the simultaneous photodiode calibration and attitude estimation filter using RAX-2 measurements collected Nov. 4, 2011 18:29:45 UT.

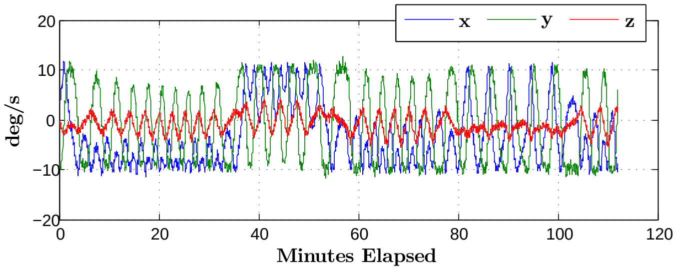

(a) RAX-1 measurements, Dec. 1, 2010 08:30:46 UT.

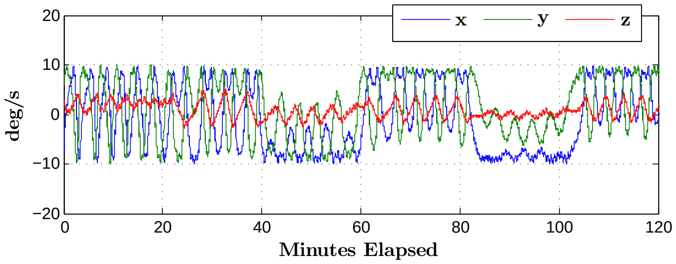

(b) RAX-2 measurements, Nov. 4, 2011 18:29:45 UT.

Figure 6.3: Measured angular velocities during the data sets shown in Figures 6.1 and 6.2.

of the gyros, rather than additional inflation due to the sampling frequency. Attitude estimation accuracy for a representative data set during this phase of the mission is shown in Figures 6.4-6.6 using three different methods: the full photodiode calibration filter, the conventional MEKF with an Earth albedo model to correct the photodiode measurements, and the conventional MEKF without an Earth albedo model, each of which are discussed in the following paragraphs. This data was collected Dec. 9, 2011 16:00:00 UT.

Figure 6.4 shows the attitude determination accuracy when using the full photodiode calibration and attitude estimation EKF. The x- and y-components of the 1-σ attitude error remain below 0.5° when in the sun, and below 0.75° when in eclipse. The 1-σ accuracy of the z-component remains below 0.75° in the sun and below 2.1° in eclipse. The uncertainty is higher about the z-axis because it is aligned with the geomagnetic field and is also the predominant spin axis of the spacecraft. Since the spacecraft is spinning about the magnetic field, the magnetometer and gyroscope provide little non-redundant information about that axis, increasing the uncertainty.

The attitude accuracy of the same data set using the 6-state MEKF with an albedo model (the same albedo model used in the photodiode calibration filter; see Section 5.3.2) to compensate the photodiode measurements is shown in Figure 6.5. The photodiode calibration parameters used in the measurement model are not taken from calibration with this data set; the calibration parameters used are those estimated from the first available data set, discussed previously in this section. The calibration parameters remain constant over time[2], and this highlights the utility of batch calibration used in conjunction with real-time attitude estimation: the on-orbit calibration can be applied during ground processing of an initial set of data collected for spacecraft health assessment and calibration, and the parameters can be uploaded for real-time sensor correction and attitude estimation using existing methods. This can simplify the algorithms running onboard the spacecraft when real-time estimation is needed. Compared to the performance of the photodiode calibration filter, the MEKF converges faster, as seen particularly in the z component accuracy of the two filters. In the photodiode calibration filter, the uncertainty continues to decrease throughout nearly the entire initial illuminated portion of the data, whereas with the MEKF, the filter converges almost immediately. This is due to the higher number of state variables in the photodiode calibration filter ($6 + 3N$, where $N = 17$ is the number of photodiodes) and is also dependent on the accuracy of state estimates used to initialize the filter. In addition to the difference in convergence time, the MEKF provides slightly higher accuracy attitude estimates throughout the data set. The 1-σ accuracy of all three components remains below 0.5° for the majority of the time in the sun, and the 99% bound on total attitude accuracy remains below 2° during the time in the sun. The increase in accuracy of the MEKF compared to the photodiode calibration filter also stems from the reduced number of states: the same amount of information (measurements) are used to estimate fewer state variables, resulting in overall slightly higher accuracy.

A third method is shown in Figure 6.6, which is again the 6-state MEKF but without inclusion of an Earth albedo model to compensate the photodiode measurements. The covariance of the photodiode measurements is increased significantly to account for the albedo contribution while utilizing the same magnetometer and process covariance as the MEKF with albedo compensation. The measurement covariance tuning is done by inspection of the residuals as discussed previously. The measurement covariance of the tuned filter corresponds to an irradiance uncertainty of 395 W/m^2, 29% of the direct sunlight. Although accounting for the albedo by simply increasing the measurement covariance is not a mathematically correct treatment of the albedo contribution in an EKF (since it does not have a Gaussian distribution), it is a widely-used approach that works well in practice. As seen by comparing Figures 6.5 and 6.6, Earth albedo significantly degrades the attitude determination accuracy. Without albedo compensation, the z-component uncertainty reaches 1.5° in the sun, and the x- and y- components reach 0.75°. The 68% confidence

---

[2] On RAX-2, the scale factor remains constant over time. On RAX-1 the scale factor did decrease over time, and this decrease can be compensated for with periodic calibration. The decrease was attributed to the effects of ultra-violet radiation, and was mitigated on RAX-2 by utilizing solar cell coverglass to shield the photodiodes.

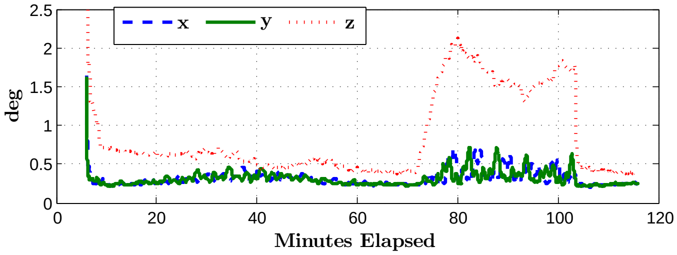

(a) 1-σ bounds.

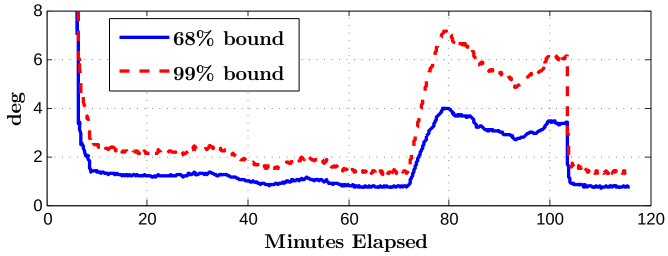

(b) Approximate total accuracy.

Figure 6.4: Attitude determination accuracy of RAX-2 data collected Dec. 9, 2011 16:00:00 UT using the photodiode calibration filter.

bound on total attitude determination accuracy is between 2° and 3° in the sun, whereas it remains below 1.1° with albedo compensation. While these accuracies may still satisfy the requirements of many missions, including the RAX mission, this highlights the impact of Earth albedo on the ADS.

6.3 Discussion and the Approach to Low-Cost Attitude Determination

We have demonstrated 0.5° 1-σ attitude accuracy using COTS magnetometers, photodiodes, and a MEMS gyroscope, and without the use of high tolerance integration/alignment procedures or thorough pre-flight calibration. The current prices of theses attitude sensors are listed in Table 6.1, and the total sensor cost is less than \$700. Although the low-cost COTS sensors typically have inherent scaling, bias, and nonorthogonality errors, and in the case of magnetometers, are also affected by the surrounding spacecraft electronics, on-orbit calibration enables accurate attitude determination with these sensors. This approach of using on-orbit calibration in place of ground-based calibration and alignment enables the rapid development and integration of low-cost attitude determination and control subsystems for small satellite missions with modest determination requirements. This

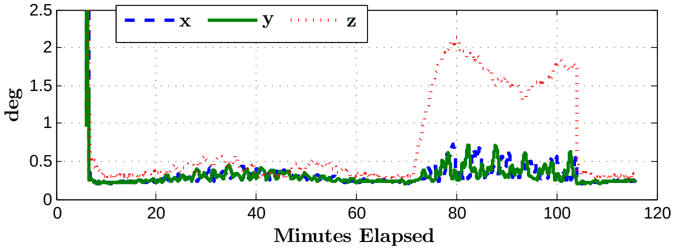

(a) 1-σ bounds.

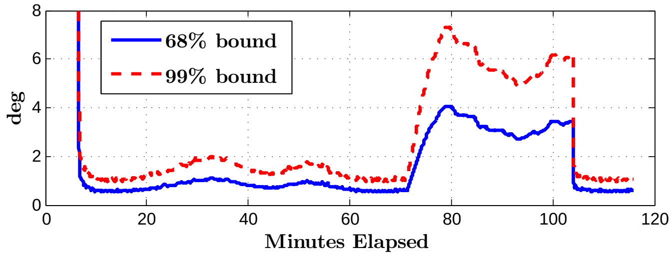

(b) Approximate total accuracy.

Figure 6.5: Attitude determination accuracy of RAX-2 data collected Dec. 9, 2011 16:00:00 UT using the 6-state MEKF with an Earth albedo model to compensate the photodiode measurements.

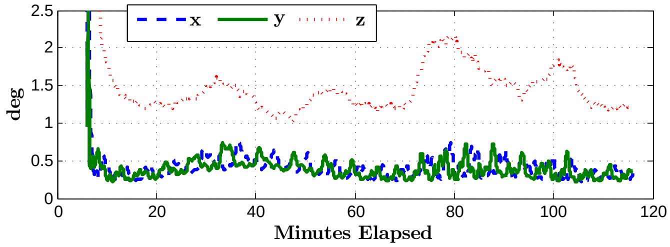

(a) 1-σ bounds.

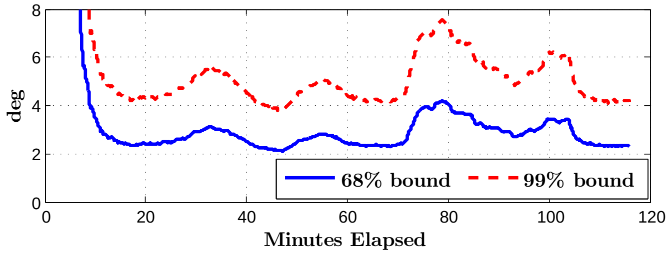

(b) Approximate total accuracy.

Figure 6.6: Attitude determination accuracy of RAX-2 data collected Dec. 9, 2011 16:00:00 UT using the 6-state MEKF without an albedo model. Albedo is treated as noise in the filter.

Table 6.1: Current prices of the COTS sensors utilized in the RAX ADS.

Sensor	**Price, USD**
Osram SFH430 photodiodes[3]	1.82 per photodiode
PNI MicroMag3 Magnetometer[4]	49.95
Analog Devices ADIS16405 IMU[5]	610.96

contrasts with the traditional spacecraft integration process that includes high-tolerance relative sensor alignment during bus-level integration [99]. In the remainder of this section, we summarize the approach taken for RAX attitude determination and its impact on the subsystem design.

By utilizing the on-orbit magnetometer calibration of Chapter 3, the magnetometer can be placed anywhere in the spacecraft without careful consideration of nearby electronics. The calibration method requires measurements of the electric current of nearby components that have a significant impact on the magnetometers, and it estimates constant parameters that map the time-varying current to magnetometer bias. In application to RAX, the current sensors that were included for general spacecraft health monitoring were sufficient for magnetometer calibration; no additional sensors were used. This included monitoring the currents in each of the four solar panels as well as the bus power lines. For real-time magnetometer correction, the ADS needs access to these current measurements, which should be considered in flight software development.

Photodiodes can be used to provide measurements of the sun vector component, and on-orbit photodiode calibration is accomplished with the method of Chapter 5. The calibration works with any number photodiodes in an arbitrary configuration. Since photodiodes provide a single component of sun vector information, multiple photodiodes are required for a full sun vector estimate. On RAX-2, photodiodes were mounted at various angles over the body frame to enable three-component sun sensing for nearly all attitudes. Since RAX-2 development, a design optimization technique was developed to determine photodiode configurations that minimize the sun vector angular accuracy, and this can be used to provide photodiode configurations for future missions. This design optimization method was the topic of Chapter 4.

Although pre-flight calibration was performed to characterize rate gyroscope scale factor error, various methods exist for on-orbit rate gyro calibration [81, 82, 96]. Small gy-

---

[3] `http://www.digikey.com/product-detail/en/SFH%202430-Z/475-2579-1-ND/1228076`, accessed September 2013.

[4] `https://www.sparkfun.com/products/244`, accessed September 2013.

[5] `http://www.digikey.com/product-search/en/sensors-transducers/multifunction/1967155?k=adis16405`, accessed September 2013.

roscope errors can also be accounted for in the process covariance of the EKF. Relative attitude and rate sensor alignment needs to also be considered. The magnetometer calibration resolved non-orthogonality errors in the magnetometer-fixed frame. Relative alignment between two identical three-axis sensors can be accomplished with any single-point attitude determination method, such as the q-method or QUEST [24]. Since the RAX gyroscope has a magnetometer in the same package, this was also used to estimate the relative gyro-magnetometer alignment under the assumption that the axes of the packaged gyro and magnetometer are coincident. The photodiode calibration estimates the alignment of the photodiodes relative to the other sensors, which in the formulation of Chapter 5, are assumed to be already aligned. Additional methods also exist for estimating the relative alignments of attitude and angular rate sensors [52, 53, 81, 82]. The sensor calibration methods, combined with relative alignment estimation methods, enable the assembly of the ADS without regard to high-tolerance integration or alignment procedures.

The RAX mission has no real-time attitude determination requirement, and all sensor processing and attitude determination was done offline after downloading sensor measurements from the spacecraft. For real-time attitude determination with these same methods, correction factors can be uploaded to the spacecraft after initial on-orbit checkout and calibration. This can also be advantageous over purely real-time on-orbit calibration as it allows for a thorough inspection of the data and filter tuning. For real-time on-orbit calibration, the photodiode calibration is well-suited for real time use with the exception of the built-in Earth albedo model. The model should be simplified [95] or pre-calculated and tabulated on-board for real-time use. The magnetometer calibration is a bath method in its current form, and extension to real-time implementation is left for future work.

CHAPTER 7

Conclusions and Areas for Future Work

We have developed methods for magnetometer and photodiode calibration that require only on-orbit measurements. These methods maximize the accuracy of the individual sensors while mitigating the need for pre-flight calibration and alignment. This reduces the cost of attitude determination by both increasing the accuracy of low-cost sensors and reducing spacecraft development time.

We have also developed a design optimization method that can be used to optimize the configuration of sensors or instruments with a limited field of view (FOV) or uncertainty that varies over the FOV, and this can be applied in the spacecraft design phase to provide an optimal photodiode configuration for sun sensing. Whereas the on-orbit magnetometer calibration method allows magnetometers to be placed anywhere within the spacecraft and their angular accuracy is not affected by sensor orientation, the angular accuracy of photodiode-based sun vector measurements is dependent on the photodiode configuration. This optimization method can be used to provide a photodiode configuration that minimizes the resulting sun vector angular uncertainty, and this replaces the traditional manual and iterative-based approaches to sensor placement in the satellite design phase, decreasing the time required for design as well as maximizing the sensor performance.

These methods have been demonstrated through application to flight data from the Radio Aurora Explorer (RAX) satellites, which utilize commercial off-the-shelf (COTS) magnetometers, photodiodes, and a MEMS gyroscope for attitude determination. After application of the on-orbit calibration, attitude determination accuracies of 0.5° 1-σ have been demonstrated. To the author's knowledge, this is the best accuracy that has been reported in the literature for this class of sensors, and it has been accomplished without the use of pre-flight calibration or high-tolerance sensor alignment during spacecraft integration. These calibration and design methods enable the rapid development of low-cost attitude determination systems. Additionally, although the calibration methods were developed for magnetometers and photodiodes, the same approach of replacing pre-flight calibration and high-tolerance integration with on-orbit calibration is applicable to all types of attitude determination sensors.

We suggest two areas for future research related to this work. First, the calibration methods can be extended for both batch and real-time implementation. The magnetometer calibration was formulated as a batch method (non-linear least squares was used), and this could be developed into a recursive implementation. Similarly, the photodiode calibration was formulated as a recursive method, but implementation with a batch implementation technique could provide a slight improvement in state estimation accuracy. In general, batch calibration is generally sufficient for spacecraft on-orbit calibration as the sensor parameters are not expected to change over time, but nonetheless, real-time implementation of in-flight sensor calibration may be desired for some applications.

The second area for future research is a comparison of using gyro measurements versus a dynamic model in recursive attitude estimation. The traditional and widely-used method for fusing gyroscope and attitude/directional measurements is to use the bias-corrected rate gyroscope measurements directly to propagate attitude [25]. This removes the need to model spacecraft dynamics, which on traditional large spacecraft, can be difficult due to factors such as large deployable structures or sloshing fuel that impact the spacecraft mass properties and disturbance torques. This approach of direct attitude propagation with the rate gyro measurements is utilized within the conventional form of the multiplicative extended Kalman filter (MEKF) and was used in the photodiode calibration developed in this dissertation. However, MEMS gyroscopes that are suitable for small spacecraft have noise characteristics that are orders of magnitude worse than larger and higher fidelity gyroscopes that are not suitable for small spacecraft (when using gyro measurements directly for attitude propagation, the gyro noise and stability characteristics are manifested in the process covariance of the filter). Additionally, since small satellites generally do not utilize complex deployable structures and in many cases do not carry fuel onboard, utilizing a dynamic model in attitude estimation may provide more accurate results than use of rate gyroscopes. This suggested area of research should answer the question of under what conditions should a gyroscope be used rather than a dynamic model, and in what cases should the two methods be combined.

APPENDIX A

Uncertainty of the Magnetometer Calibration Parameters

The Fisher information matrix [55] is used to demonstrate observability of the calibration parameters and to investigate the degradation of the calibration accuracy as the coverage of the attitude sphere decreases (this was discussed in Section 3.5.5). For a linear estimation problem of the form $\tilde{\mathbf{y}} = H\mathbf{x} + \mathbf{v}$, where $\tilde{\mathbf{y}}$ is the $m \times 1$ measurement vector, $\mathbf{x}$ is the $n \times 1$ state vector, and $\mathbf{v}$ is the $m \times 1$ zero mean measurement noise vector of covariance R, the Fisher information matrix is

$$F = H^T R^{-1} H \ . \tag{A.1}$$

The Cramér-Rao Inequality provides a lower bound for the state covariance matrix, P [55]:

$$P \geq F^{-1} \ . \tag{A.2}$$

The magnetometer parameter estimation is linearized in order to use this inequality to determine the approximate lower bound on the uncertainty of the calibration parameters. The measurement in the calibration problem, the squared magnitude of the corrected magnetic field measurements, is given by

$$\tilde{\mathbf{y}} = \mathbf{f}(\tilde{\mathbf{B}}, \tilde{\mathbf{I}}, \mathbf{x}) \ , \tag{A.3}$$

as in Eq. (3.8). H is approximated by linearizing at the final state estimate, $\mathbf{x}_f$:

$$H = \left. \frac{\partial \mathbf{f}(\tilde{\mathbf{B}}, \tilde{\mathbf{I}}, \mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_f} \ . \tag{A.4}$$

The raw measurements for a single point in time are

$$\tilde{\mathbf{y}}' = \left[\begin{array}{cccccccc} \tilde{B}_x & \tilde{B}_y & \tilde{B}_z & \tilde{I}_1 & \tilde{I}_2 & \tilde{I}_3 & \tilde{I}_4 & \tilde{I}_5 \end{array}\right]^T , \tag{A.5}$$

where $\tilde{B}_x$, $\tilde{B}_y$, and $\tilde{B}_z$ are the magnetometer measurements (modeled in Eqs. (3.5)-(3.7)), and $\tilde{I}_1$-$\tilde{I}_5$ are the five current measurements used in the calibration.

Let the covariance of the error in $\tilde{\mathbf{y}}$ be given by R and the covariance of $\tilde{\mathbf{y}}'$ be given by $R_{\tilde{B},\tilde{I}}$. It is assumed that the sensor measurement covariance is a diagonal matrix of the variances of each measurement,

$$R_{\tilde{B},\tilde{I}} = \text{diag}\left[\begin{array}{cccccccc} \sigma_x^2 & \sigma_y^2 & \sigma_z^2 & \sigma_1^2 & \sigma_2^2 & \sigma_3^2 & \sigma_4^2 & \sigma_5^2 \end{array}\right] . \tag{A.6}$$

Since $\tilde{\mathbf{y}}$ is a non-linear function of the sensor measurements and calibration parameters, it is not necessarily a Gaussian random vector. A Monte Carlo simulation was performed to investigate the probability density function of $\tilde{\mathbf{y}}$. $\tilde{\mathbf{y}}$ was calculated using 10,000 random Gaussian distributed sensor measurements and the calibration parameters from the December 1 data set, given in Table 3.3. The histogram of the results is shown in Figure A.1. This histogram indicates that it is reasonable to assume that the error in $\tilde{\mathbf{y}}$ is a Gaussian random vector. Therefore, the covariance matrix R is calculated after linearization as it would be for linear systems:

$$R = \text{diag}\left[\begin{array}{ccc} R_1 & \dots & R_m \end{array}\right] , \tag{A.7}$$

where

$$R_i = \left[\mathbf{f}(\tilde{\mathbf{B}}_i, \tilde{\mathbf{I}}_i, \mathbf{x})\right] R_{\tilde{B},\tilde{I}} \left[\mathbf{f}(\tilde{\mathbf{B}}_i, \tilde{\mathbf{I}}_i, \mathbf{x})\right]^T . \tag{A.8}$$

The necessary tools to approximate the Fisher information matrix are now available. From pre-flight testing, the uncertainty in the magnetometer measurements was found to be less than the resolution of the sensor, so the standard deviation is approximated by the resolution, 128 nT in each axis. The expected uncertainty of the IGRF magnitude is an order of magnitude less than the sensor resolution, so it is not explicitly included in the magnetic field uncertainty. The uncertainty of the current sensors throughout the spacecraft was not thoroughly characterized before flight. The standard deviation is assumed to be 5 mA (the resolution is 1 mA). Using these standard deviations, the estimated lower bound of the 3-σ values of the calibration parameters, obtained from the diagonal elements of the state covariance matrix of Eq. (A.2), are given in Table A.1.

The 3-σ values are optimistic. The Cramér-Rao Inequality provides a *lower bound* of the state covariance, not the actual covariance. In the magnetometer calibration, non-linear least-squares minimization was used, but it is not an optimal estimator and does

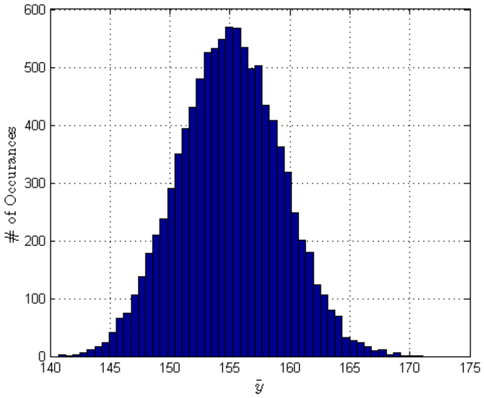

Figure A.1: Histogram of $\tilde{\mathbf{y}}$ calculated from 10,000 random sensor measurements with Gaussian distribution and the calibration parameters for the December 1 data set, shown in Table 3.3. The standard deviations used are $\sigma_x = \sigma_y = \sigma_z = 128$ nT, and $\sigma_1 = \cdots = \sigma_5 = 5$ mA.

not consider the probability density function of the errors. This, coupled with the fact that linearized approximations of nonlinear functions were used, indicates that the actual 3-σ values of our state estimates are worse than the lower bound. This is confirmed by observing that the parameters estimated from each data set do not match within the 3-σ bound. However, the lower bound provides insight into the accuracy of the parameters and the relationship to coverage of the attitude sphere. The lower bound of the 3-σ values increases as the coverage of the attitude sphere decreases (sphere coverage for each data set is shown in Figure 3.6). This is expected, and is an example of one metric that quantifies the decrease in observability of the calibration parameters as the sphere coverage decreases. Ultimately, if the data points were available from less than n independent data points on the sphere, then H would become rank deficient and the parameters would be unobservable.

Table A.1: Approximate lower bound of the 3-σ values of the calibration parameter estimates for each data set. Nine of the 24 parameters are shown; the remaining are omitted from the table for simplicity.

Data Set	Dec. 1	Dec. 15	Dec. 30
a	4.0×10^{-4}	8.8×10^{-4}	7.3×10^{-3}
b	4.2×10^{-4}	1.1×10^{-3}	1.8×10^{-2}
c	4.6×10^{-4}	3.3×10^{-3}	8.9×10^{-3}
x_0 (μT)	0.015	0.030	0.22
y_0 (μT)	0.015	0.031	0.31
z_0 (μT)	0.019	0.081	0.32
ρ (deg)	0.041	0.092	1.18
ϕ (deg)	0.040	0.14	1.31
λ (deg)	0.040	0.13	0.78

APPENDIX B

Uncertainty of the Sun Vector Estimate

The objective function used in the photodiode orientation optimization of Chapter 4 was derived using the covariance matrix of the sun vector estimate. In application to attitude determination, knowledge of the angular accuracy of the sun vector estimate is desired. But since the error bounds on the components of the sun vector estimate form an ellipsoid[1] and not a sphere, there is no single number quantifying the angular error that corresponds to the sun vector covariance matrix.

Two possible approximations of a single angular error corresponding to the sun vector covariance matrix are to use the radius of a sphere with equivalent volume as the error ellipsoid, or to use the semi-major axis of the error ellipsoid [98]. The former approximation is most accurate when the principal axes of the error ellipsoid are approximately equal, and the latter approximation is best when the semi-major axis of the error ellipsoid is much greater than the other two axes. Neither of these assumptions are generally valid for the optimization examples presented in this paper. To quantify the angular accuracy of the sun vector estimate without making any assumptions on the shape of the error ellipsoid, Monte Carlo simulations are performed to estimate the angular uncertainty resulting from a given photodiode configuration. The simulations consist of the following steps:

1. For each direction on the attitude sphere, calculate the ideal measurement of each photodiode.
2. For each trial, add pseudo-random numbers from a mean zero normal distribution to simulate the measurements. 1000 trials for each sun vector direction were used.
3. Estimate the sun vector using the simulated measurements and constrained minimization of Eq. (4.9).

---

[1] This is under the assumption that the errors in the photodiode measurements have a Gaussian distribution.

4. For each sun vector estimate, calculate the angular error using the dot product. That is,
$$\delta\theta = \cos^{-1}\left(\hat{\boldsymbol{s}}_{est}^{T}\hat{\boldsymbol{s}}_{true}\right) \ , \tag{B.1}$$
where $\hat{\boldsymbol{s}}_{est}$ is the unit vector estimate and $\hat{\boldsymbol{s}}_{true}$ is the true unit sun vector.

5. The uncertainty in each sun vector direction is then taken to be the average angular uncertainty from the 1000 trials. For the j-th sun vector direction, this is denoted as $\delta\bar{\theta}_j$. This is the quantity shown in Figures 4.8(b), 4.8(d), 4.11(b), and 4.11(c).

Example histograms resulting from the simulation for a single sun vector direction are shown in Figures B.1 and B.2. Figure B.1 shows the histograms of the photodiode measurements for all five photodiodes for which the sun was in the field of view for the given sun vector direction. The histogram of the resulting angular uncertainty is shown in Figure B.2. These results are from the photodiode configuration of Example 2 optimized using identical weighting over the attitude sphere.

The angular uncertainty over the attitude sphere is taken to be the weighted average of the angular uncertainty in each sun vector direction on the sphere. That is,

$$\delta\theta_{total} = \sum_{j=1}^{N}\left(\frac{W_j}{\sum_{k=1}^{N} W_k}\delta\bar{\theta}_j\right) \ . \tag{B.2}$$

This is the quantity given in Figures 4.7(b) and 4.10(b).

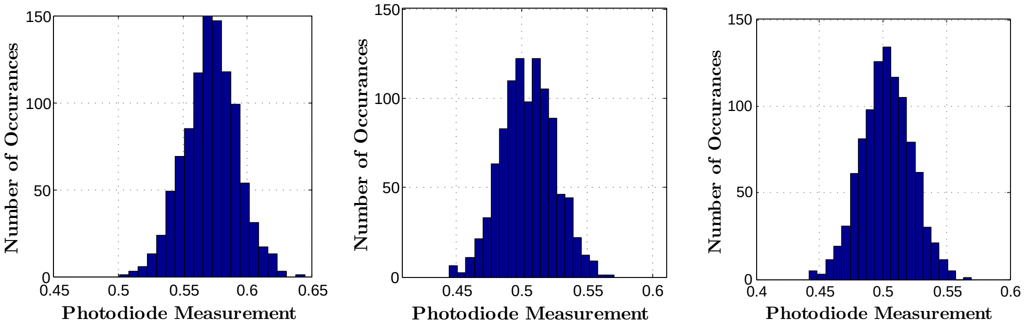

(a) Sensor 1 (b) Sensor 2 (c) Sensor 3

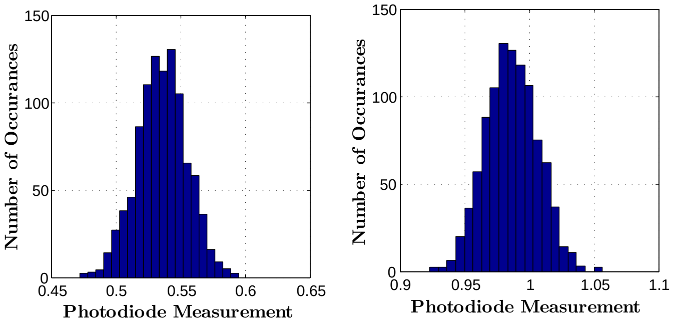

(d) Sensor 4 (e) Sensor 5

Figure B.1: Histograms of the simulated measurements for the five sun sensors for which the simulated sun vector (direction on the attitude sphere) is in the field of view. The peak current output of the photodiodes is 1 and the standard deviation is 0.02.

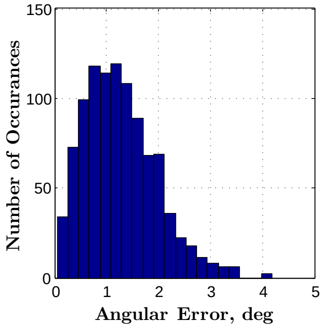

Figure B.2: Histogram of the angular error in the sun vector estimate. The mean error is 1.3°.

BIBLIOGRAPHY

[1] Starin, S. R. and Eterno, J., *Space Mission Engineering: The New SMAD*, chap. 19: Spacecraft Subsystems II – Control Systems, Microsm Press, 2011, pp. 565–600.

[2] Wertz, J. R., editor, *Spacecraft Attitude Determination and Control*, Kluwer, 1978.

[3] Barnhart, D. J., Vladimirova, T., and Sweeting, M. N., "Very-Small-Satellite Design for Distributed Space Missions," *AIAA Journal of Spacecraft and Rockets*, Vol. 44, No. 6, November-December 2007, pp. 1294–1306, doi: 10.2514/1.28678.

[4] "CubeSat Design Specification," Availailable online: `http://cubesat.org/images/developers/cds_rev12.pdf`, accessed August 2013.

[5] Swartwout, M., "The Long-Threatened Flood of University-Class Spacecraft (and CubeSats) Has Come: Analyzing the Numbers," *Proceedings of the AIAA/USU Small Satellite Conference*, Utah State University, Logan, UT, August 2013, Paper number SSC13-IX-01.

[6] Baker, D. N. and Worden, S. P., "The Large Benefits of Small-Satellite Missions," *EOS, Transactions, American Geophysical Union*, Vol. 89, No. 33, August 2008, pp. 301–303, doi:10.1029/2008EO330001.

[7] Moretto, T. and Robinson, R. M., "Small Satellites for Space Weather Research," *Space Weather*, Vol. 6, No. 5, May 2008, doi: 10.1029/2008SW000392.

[8] Sandau, R., "Status and trends of small satellite missions for Earth observation," *Acta Astronautica*, Vol. 66, No. 1-2, 2010, pp. 1 – 12, doi: 10.1016/j.actaastro.2009.06.008.

[9] Kramer, H. J. and Cracknell, A. P., "An overview of small satellites in remote sensing," *International Journal of Remote Sensing*, Vol. 29, No. 15, 2008, pp. 4285–8337, doi: 10.1080/01431160801914952.

[10] Smith, M. W., Seager, S., Pong, C. M., Villasenor, J. S., Ricker, G. R., Miller, D. W., Knapp, M. E., Farmer, G. T., and Jensen-Clem, R., "ExoplanetSat: Detecting transiting exoplanets using a low-cost CubeSat platform," *Proceedings of SPIE–the International Society for Optical Engineering*, Vol. 7731, 2010, doi: 10.1117/12.856559.

[11] Bahcivan, H., Cutler, J. W., Bennett, M., Kempke, B., Springmann, J. C., Buonocore, J., Nicolls, M., and Doe, R., "First Measurements of Radar Coherent Scatter by

the Radio Aurora Explorer CubeSat," *Geophysical Research Letters*, Vol. 39, 2012, pp. L14101, doi:10.1029/2012GL052249.

[12] Bahcivan, H., Cutler, J. W., Springmann, J. C., and Doe, R., "Magnetic Aspect Sensitivity of High-Latitude E Region Irregularities Measured by the RAX CubeSat," *Journal of Geophysical Research Space Physics*, Submitted for publication, January 2012.

[13] Norton, C., Pellegrino, S., and Johnson, M., "Findings of the KECK Institute for Space Studies Program on Small Satellites: A Revolution in Space Science," *Proceedings of the AIAA/USU Conference on Small Satellites*, Logan, UT, August 2013.

[14] Klesh, A., Baker, J., Bellardo, J., CastilloRogez, J., Cutler, J., Halatek, L., Lightsey, G., and Murphy, N., "NSPIRE: Interplanetary NanoSpacecraft Pathfinder in Relevant Environment," *Proceedings of the AIAA/USU Small Satellite Conference*, Logan, UT, August 2013.

[15] Sandau, R., Brieß, K., and D'Errico, M., "Small Satellites for Global Coverage: Potential and Limits," *ISPRS Journal of Photogrammetry and Remote Sensing*, Vol. 66, No. 6, November 2010, pp. 492–504, doi:10.1016/j.isprsjprs.2010.09.003.

[16] Ruf, C., Unwin, M., Dickinson, J., Rose, R., Rose, D., Vincent, M., and Lyons, A., "CYGNSS: Enabling the Future of Hurricane Prediction [Remote Sensing Satellites]," *Geoscience and Remote Sensing Magazine, IEEE*, Vol. 1, No. 2, 2013, pp. 52–67, doi: 10.1109/MGRS.2013.2260911.

[17] Air Force Space Command, "Resiliency and Dissaggregated Space Architectures," 2013, Available online: `http://www.afspc.af.mil/shared/media/document/AFD-130821-034.pdf`.

[18] Bekker, D. L., Pingree, P. J., Werne, T. A., Wilson, T. O., and Franklin, B. R., "The COVE Payload A Reconfigurable FPGA-Based Processor for CubeSats," *Proceedings of the AIAA/USU Small Satellite Conference*, Logan, UT, August 2011.

[19] Spangelo, S. C. and Cutler, J. W., "Analytical Modeling Framework and Applications for Space Communication," *AIAA Journal of Aerospace Computing, Information, and Communication*, Accepted for publication, 2012.

[20] Hwang, J. T., Lee, D. Y., Cutler, J. W., and Martins, J. R. R. A., "Large-Scale Multidisciplinary Optimization of a Small Satellites Design and Operation," *AIAA Journal of Spacecraft and Rockets*, 2013, Under review.

[21] Springmann, J. C. and Cutler, J. W., "Attitude-Independent Magnetometer Calibration with Time-Varying Bias," *Journal of Guidance, Control, and Dynamics*, Vol. 35, No. 4, June-July 2012, pp. 1080–1088, doi: 10.2514/1.56726.

[22] Springmann, J. C., "On-Orbit Calibration of Photodiodes for Attitude Determination," *Proceedings of the USU/AIAA Small Satellite Conference*, Utah State University, Logan, UT, August 2013, Paper number SSC13-VIII-1.

[23] Crassidis, J. L. and Junkins, J. L., *Optimal Estimation of Dynamic Systems*, Chapman and Hall/CRC, 2004, Section 3.7.

[24] Crassidis, J. L., Markley, F. L., and Cheng, Y., "Survey of Nonlinear Attitude Estimation Methods," *Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 1, January-February 2007, pp. 12–28, doi: 10.2514/1.22452.

[25] Lefferts, E., Markley, F., and Shuster, M., "Kalman Filtering for Spacecraft Attitude Estimation," *Journal of Guidance, Control, and Dynamics*, Vol. 5, No. 5, September-October 1982, pp. 417–429, doi: 10.2514/6.1982-70.

[26] Lee, D. E., *Space Mission Engineering: The New SMAD*, chap. 23.3: Spacecraft Manufacturing, Integration, and Test, Microsm Press, 2011, pp. 711–718, J.R. Wertz, D.F. Everett, and J.J.Puschell, editors.

[27] Cutler, J. W. and Bahcivan, H., "The Radio Aurora Explorer – A Mission Overview," *AIAA Journal of Spacecraft and Rockets*, 2013, In-print. doi: 10.2514/1.A32436.

[28] Bahcivan, H. and Cutler, J. W., "Radio Aurora Explorer: Mission Science and Radar System," *Radio Science*, Vol. 47, No. RS2012, 2012, pp. 1–12, doi:10.1029/2011RS004817.

[29] Cutler, J. W., Springmann, J. C., Spangelo, S., and Bahcivan, H., "Initial Flight Assessment of the Radio Aurora Explorer," *Proceedings of the 25th Annual Small Satellite Conference*, Utah State University, Logan, Utah, August 2011, Paper number SSC11-VI-6.

[30] Springmann, J. C., Cutler, J. W., and Bahcivan, H., "Initial Flight Results of the Radio Aurora Explorer," *Proceedings of the 62nd International Astronautical Congress*, Cape Town, South Africa, October 2011.

[31] Springmann, J. C., Sloboda, A. J., Klesh, A. T., Bennett, M. W., and Cutler, J. W., "The Attitude Determination System of the RAX Satellite," *Acta Astronautica*, Vol. 75, 2012, pp. 120–135, doi: 10.1016/j.actaastro.2012.02.001.

[32] Park, G., Seagraves, S., and McClamroch, N. H., "A Dynamic Model of a Passive Magnetic Attitude Control System for the RAX Nanosatellite," *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Toronto, August 2010.

[33] Gerhardt, D. T., "Passive Magnetic Attitude Control for CubeSat Spacecraft," *Proceedings of the AIAA/USU Conference on Small Satellites*, Logan, UT, August 2010.

[34] Nason, I., Puig-Suari, J., and Twiggs, R., "Development of a family of picosatellite deployers based on the CubeSat standard," *IEEE Aerospace Conference Proceedings*, Vol. 1, Big Sky, MT, USA, March 2002, pp. 1 – 457.

[35] Springmann, J. C., Kempke, B. P., Cutler, J. W., and Bahcivan, H., "Clock Synchronization on the RAX Spacecraft," *Acta Astronautica*, 2013, Under review.

[36] "CubeSat Design Specifications," August 2009, Available online: `http://cubesat.org/images/developers/cds_rev12.pdf`, accessed September 2013.

[37] Blaylock, B. T., *Spacecraft Attitude Determination and Control*, chap. 6.3: Magnetometers, Kluwer Academic Publishers, 1978, pp. 180–184.

[38] Ripka, P. and Acuna, M. H., *Magnetic Sensors and Magnetometers*, chap. 10: Applications of Magnetic Sensors, Artech House, 2001, pp. 369–402.

[39] Kabamba, P. T., "Coursenotes for AE584: Avionics, Navigation and Guidance of Aerospace Vehicles. Chapter 4: Nagivation," University of Michigan, Ann Arbor, October 12, 2010.

[40] van der Horn, G. and Huijsing, J. L., *Integrated Smart Sensors Design and Calibration*, chap. 1.4: Measurement Errors and Correction, Kluwer Academic Publishers, 1998, pp. 6–16.

[41] Alonso, R. and Shuster, M. D., "TWOSTEP: A fast robust algorithm for attitude-independent magnetometer-bias determination," *Journal of the Astronautical Sciences*, Vol. 50, No. 4, 2003, pp. 433 – 451.

[42] Alonso, R. and Shuster, M. D., "Attitude-independent magnetometer-bias determination: A survey," *Journal of the Astronautical Sciences*, Vol. 50, No. 4, 2003, pp. 453 – 475.

[43] Alonso, R. and Shuster, M. D., "Complete linear attitude-independent magnetometer calibration," *Journal of the Astronautical Sciences*, Vol. 50, No. 4, 2003, pp. 477 – 490.

[44] Vasconcelos, J. F., Elkaim, G., Silvestre, C., Oliveira, P., and Cardeira, B., "A Geometric Approach to Strapdown Magnetometer Calibration in Sensor Frame," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 47, No. 2, April 2011, pp. 1293–1306, doi: 10.1109/TAES.2011.5751259.

[45] Gebre-Egziabher, D., Elkaim, G. H., Powell, J. D., and Parkinson, B. W., "Calibration of Strapdown Magnetometers in the Magnetic Field Domain," *ASCE Journal of Aerospace Engineering*, Vol. 19, No. 2, April 2006, pp. 87–102, doi: 10.1061/(ASCE)0893-1321(2006)19:2(87).

[46] Foster, C. C. and Elkaim, G. H., "Extension of a two-step calibration methodology to include nonorthogonal sensor axes," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 44, No. 3, July 2008, pp. 1070 – 1078, doi: 10.1109/TAES.2008.4655364.

[47] Crassidis, J., Lai, K.-L., and Harman, R., "Real-time attitude-independent three-axis magnetometer calibration," *Journal of Guidance, Control, and Dynamics*, Vol. 28, No. 1, 2005, pp. 115 – 20.

[48] Finlay, C. C., Maus, S., Beggan, C. D., Bondar, T. N., Chambodut, A., Chernova, T. A., Chulliat, A., Golovkov, V. P., Hamilton, B., Hamoudi, M., Holme, R., Hulot, G., Kuang, W., Langlais, B., Lesur, V., Lowes, F. J., Luehr, H., Macmillan, S., Mandea, M., McLean, S., Manoj, C., Menvielle, M., Michaelis, I., Olsen, N., Rauberg, J., Rother, M., Sabaka, T. J., Tangborn, A., Toffner-Clausen, L., Thebault, E., Thomson, A. W. P., Wardinski, I., Wei, Z., Zvereva, T. I., and Int Assoc Geomagnetism Aeronomy, Working Group V-MOD, "International Geomagnetic Reference Field: the eleventh generation," *Geophysical Journal International*, Vol. 183, No. 3, December 2010, pp. 1216–1230, doi: 10.1111/j.1365-246X.2010.04804.x.

[49] Ghanbarpour Asl, H., Pourtakdoust, S., and Samani, M., "A new non-linear algorithm for complete pre-flight calibration of magnetometers in the geomagnetic field domain," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Vol. 223, No. 6, 2009, pp. 729 – 739, doi: 10.1243/09544100JAERO485.

[50] Acuna, M. H., Scheifele, J. L., Stella, P., Kloss, C., Smith, B., Heinshohn, G., and Sharmit, K., "Magnetic field cancellation techniques for the Mars Global Surveyor solar array," *Conference Record of the Twenty Fifth IEEE Photovoltaic Specicialtes Conference*, IEEE, Washington, D.C., May 1996, pp. 325 – 8.

[51] Kim, E., Bang, H., and Lee, S.-H., "Engineering notes: Attitude-independent magnetometer calibration considering magnetic torquer coupling effect," *Journal of Spacecraft and Rockets*, Vol. 48, No. 4, 2011, pp. 691 – 694, doi: 10.2514/1.52634.

[52] Shuster, M. D., Pitone, D. S., and Bierman, G. J., "Batch Estimation of Spacecraft Sensor Alignments I. Relative Alignment Estimation," *Journal of the Astronautical Sciences*, Vol. 39, No. 4, October-December 1991, pp. 519–546.

[53] Shuster, M. D. and Pitone, D. S., "Batch Estimation of Spacecraft Sensor Alignments II. Absolute Alignment Estimation," *Journal of the Astronautical Sciences*, Vol. 39, No. 4, October-December 1991, pp. 547–471.

[54] Wolfson, R. and Pasachoff, J. M., *Physics for Scientists and Engineers*, chap. 30, Addison Wesley Longman, Inc., 3rd ed., 1999, pp. 770–802.

[55] Crassidis, J. L. and Junkins, J. L., *Optimal Estimation of Dynamic Systems*, Chapman and Hall/CRC, 2004, Sections 1.4 and 2.4.

[56] Maus, S., Macmillan, S., Lowes, F., and Bondar, T., "Evaluation of candidate geomagnetic field models for the 10th generation of IGRF," *Earth, Planets, and Space*, Vol. 57, No. 12, 2005, pp. 1173–1181.

[57] Springmann, J. C., "Attitude-Independent Magnetometer Calibration with Time-Varying Bias," *Proceedings of the 25th Annual Small Satellite Conference*, Logan, UT, August 2011.

[58] Sidi, M. J., *Spacecraft Dynamics and Control, a Practical Engineering Approach*, Cambridge University Press, 1997, Appendix B.3: Sun Sensors.

[59] Bhanderi, D. D. V., *Spacecraft Attitude Determination with Earth Albedo Corrected Sun Sensor Measurements*, Ph.D. thesis, Aalborg University, Aalborg, Denmark, August 2005.

[60] Taraba, M., Rayburn, C., Tsuda, A., and MacGillivray, C., "Boeing's CubeSat TestBed 1 Attitude Determination Design and On-Orbit Experience," *Proceedings of the AIAA/USU Conference on Small Satellites*, Logan, UT, August 2009, Paper number SSC09-X-6.

[61] Ovchinnikov, M., Ilyin, A., Kupriynova, N., Penkov, V., and Selivanov, A., "Attitude dynamics of the first Russian nanosatellite TNS-0," *Acta Astronautica*, Vol. 61, No. 1-6, 2007, pp. 277 – 285, doi: 10.1016/j.actaastro.2007.01.006.

[62] Allgeier, S. E., Mahin, M., and Fitz-Coy, N. G., "Design and analysis of a coarse sun sensor for pico-satellites," *Proceedings of the AIAA Infotech at Aerospace Conference and Exhibit and AIAA Unmanned...Unlimited Conference*, AIAA, Seattle, WA, United states, April 2009, doi: 10.2514/6.2009-1837.

[63] Reeves, E. I., *Spacecraft mission analysis and design*, chap. 10: Spacecraft Design and Sizing, Microsm, 1999, pp. 301–352.

[64] Jackson, B. and Carpenter, B., "Optimal placement of spacecraft sun sensors using stochastic optimization," *2004 IEEE Aerospace Conference Proceedings*, Big Sky, Montana, March 2004, pp. 3916–3923.

[65] Mosher, T., "Spacecraft design using a genetic algorithm optimization approach," *1998 IEEE Aerospace Conference Proceedings*, Aspen, CO, March 1998, pp. 123–134.

[66] Taylor, E., "Evaluation of multidisciplinary design optimization techniques as applied to spacecraft design," *2000 IEEE Aerospace Conference Proceedings*, Vol. 1, March 2000, pp. 371–384.

[67] Wertz, J. R., *Spacecraft Attitude Determination and Control*, chap. 2: Attitude Geometry, Kluwer Academic Publishers, 1978, pp. 22–35.

[68] Tegmark, M., "An Icosahedron-based Method for Pixelizing the Celestial Sphere," *Astrophysical Journal, Letters*, Vol. 470, No. 2, October 1996, pp. L81–L84, doi: 10.1086/310310.

[69] Teanby, N. A., "An icosahedron-based method for even binning of globally distributed remote sensing data," *Computers & Geosciences*, Vol. 32, No. 9, November 2006, pp. 1442 – 1450, doi: 10.1016/j.cageo.2006.01.007.

[70] Kjellberg, H. C. and Lightsey, E. G., "A Constrained Attitude Control Module for Small Satellites," *Proceedings of the 26th Annual AIAA/USU Conference on Small Satellites*, Logan, Utah, August 2012.

[71] Lerner, G. M., *Spacecraft Attitude Determination and Control*, chap. 6.1: Sun Sensors, Kluwer Academic Publishers, 1978, pp. 155–166, ISBN 9027712042.

[72] Crassidis, J. L. and Junkins, J. L., *Optimal Estimation of Dynamic Systems*, Chapman and Hall/CRC, 2004, Chapter 2.

[73] Shuster, M. D., "Constraint in Attitude Estimation Part I: Constrained Estimation," *The Journal of the Astronautical Sciences*, Vol. 51, No. 1, January-March 2003, pp. 51–74.

[74] Dado, S. and Fischer, J., *Modern Sensors Handbook*, chap. 2.2.2: Light Detectors, ISTE Ltd, 2007, pp. 54–62, ISBN 1441964657.

[75] Tambo, T., Shibata, M., Mizuno, Y., and Yamauchi, T., "Search Method of Sun Using Fixed Five Photodiode Sensor," *IEEJ Transactions on Sensors and Micromachines*, Vol. 129, No. 2, 2009, pp. 53–59, doi: 10.1541/ieejsmas.129.53.

[76] Ortega, P., Lopez-Rodriguez, G., Ricart, J., Dominguez, M., Castaner, L., Quero, J., Tarrida, C., Garcia, J., Reina, M., Gras, A., and Angulo, M., "A Miniaturized Two Axis Sun Sensor For Attitude Control of Nano-satellites," *IEEE Sensors Journal*, Vol. 10, No. 10, October 2010, pp. 1623 – 32, doi: 10.1109/JSEN.2010.2047104.

[77] Wu, S.-F. and Steyn, W. H., "Modelling and in-orbit calibration practice of a miniature 2-axis analogue Sun sensor," *Aerospace Science and Technology*, Vol. 6, No. 6, October 2002, pp. 423 – 33, doi: 10.1016/S1270-9638(02)01187-2.

[78] Rufino, G., Grassi, M., and Rolfi, M., "Preliminary calibration results for a high-precision CMOS sun sensor," *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Honolulu, HI, August 2008, doi: 10.2514/6.2008-7319.

[79] Springmann, J. C., Kempke, B. P., Cutler, J. W., and Bahcivan, H., "Initial Flight Results of the RAX-2 Satellite," *Proceedings of the 26th Annual AIAA/USU Conference on Small Satellites*, Logan, Utah, August 2012, Paper number SSC12-XI-5.

[80] Springmann, J. C. and Cutler, J. W., "Initial Attitude Analysis of the RAX Satellite," *Advances in the Astronautical Sciences: Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, Girdwood, Alaska, 2011, Paper number AAS 11-600. ISSN 00653438.

[81] Pittelkau, M. E., "Kalman Filtering for Spacecraft System Alignment Calibration," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 6, November-December 2001, pp. 1187–1195, doi: 10.2514/2.4834.

[82] Lai, K.-L., Crassidis, J., and Harman, R., "In-Space Spacecraft Alignment Calibration Using the Unscented Filter," *AIAA Guidance, Navigation, and Control Conference*, Austin, TX, August 2003, doi: 10.2514/6.2003-5563.

[83] Duan, F., Liu, J., Li, R., and Li, D., "Real-time attitude-independent sun-sensor/magnetometer calibration algorithm for micro-satellite," *Proceedings of the SPIE - The International Society for Optical Engineering*, Vol. 5985, Wuhan, China, October 2005, pp. 598548–1–5, doi: 10.1117/12.658334.

[84] Reid, D., "Calibration of the Milstar attitude determination system," *Proceedings of the 1997 American Control Conference*, Albuquerque, NM, June 1997, doi: 10.1109/ACC.1997.611861.

[85] Julier, S. J., Uhlmann, J. K., and Durrant-Whyte, H. F., "A New Method fo the Nonliear Transformation of Means and Covariances in Filters and Estimators," *IEEE Transactions on Automatic Control*, Vol. AC-45, No. 3, 2000, pp. 477–482, doi: 10.1109/9.847726.

[86] Crassidis, J. L. and Markley, F. L., "Unscented filtering for spacecraft attitude estimation," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 4, July-August 2003, pp. 536–542, doi: 10.2514/2.5102.

[87] Markley, F. L., "Attitude Error Representations for Kalman Filtering," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 2, March-April 2003, pp. 311–317, doi: 10.2514/2.5048.

[88] Crassidis, J. L. and Junkins, J. L., *Optimal Estimation of Dynamic Systems*, Chapman and Hall/CRC, 2004, ISBN 1-58488-371-X.

[89] Farrenkopf, R. L., "Analytic Steady-State Accuracy Solutions for Two Common Spacecraft Attitude Estimators," *Journal of Guidance, Control, and Dynamics*, Vol. 1, No. 4, 1978, pp. 282–284, doi: 10.2514/3.55779.

[90] Wan, E. A. and van der Merwe, R., *Kalman Filtering and Neural Networks*, chap. 7. The Unscented Kalman Filter, Wiley, 2001, pp. 221–280, edited by S. Haykin. ISBN: 0471369985.

[91] Bhanderi, D. D. V., "Modeling earth albedo currents on sun sensors for improved vector observation," *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Keystone, CO, August 2006, doi: 10.2514/6.2006-6592.

[92] Bhanderi, D. D. V. and Bak, T., "Modeling earth albedo for satellites in earth orbit," *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, San Francisco, CA, August 2005, doi: 10.2514/6.2005-6465.

[93] Psiaki, M. L., Martel, F., and Pal, P., "Three-axis attitude determination via Kalman filtering of magnetometer data," *Journal of Guidance, Control, and Dynamics*, Vol. 13, No. 3, May-June 1990, pp. 506–514, doi: 10.2514/3.25364.

[94] Burton, R., Rock, S., Springmann, J., and Cutler, J., "Online Attitude Determination of a Passively Magnetically Stabilized Spacecraft," *Proceedings of the 23rd AAS/AIAA Spaceflight Mechanics Meeting*, Kauai, HI, Februrary 2013, Paper number AAS 13-364.

[95] Appel, P., "Attitude estimation from magnetometer and earth-albedo-corrected coarse sun sensor measurements," *Acta Astronautica*, Vol. 56, No. 1-2, January 2005, pp. 115–126, doi: 10.1016/j.actaastro.2004.09.001.

[96] Lai, K.-L., Crassidis, J., and Harman, R., "Real-Time Attitude-Independent Gyro Calibration from Three-Axis Magnetometer Measurements," *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Providence, RI, August 2004, Paper number AIAA-2004-4855.

[97] Crassidis, J. L. and Junkins, J. L., *Optimal Estimation of Dynamic Systems*, Chapman and Hall/CRC, 2004, Chapter 7.2.

[98] Lerner, G. M., *Spacecraft Attitude Determination and Control*, chap. 12.3: Covariance Analysis, Kluwer Academic Publishers, 1978, pp. 429–434.

[99] Lee, D. E., *Space Mission Engineering: The New SMAD*, chap. 23.3: Spacecraft Manufacturing, Integration and Test, Microsm Press, 2011, pp. 711–718, J.R. Wertz, D.F. Everett, and J.J. Puschell, ed.