

Small Satellite Passive Magnetic Attitude Control

by

David T. Gerhardt

B.S., Virginia Polytechnic Institute and State University, 2008

M.S., University of Colorado, 2011

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Aerospace Engineering Sciences

2014

This thesis entitled:
Small Satellite Passive Magnetic Attitude Control
written by David T. Gerhardt
has been approved for the Department of Aerospace Engineering Sciences

Professor Scott Palo

Professor Hanspeter Schaub

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Gerhardt, David T. (Ph.D., Aerospace Engineering Sciences)

Small Satellite Passive Magnetic Attitude Control

Thesis directed by Professor Scott Palo

Passive Magnetic Attitude Control (PMAC) is capable of aligning a satellite within 5 degrees of the local magnetic field at low resource cost, making it ideal for a small satellite. However, simulation attempts to date have not been able to predict the attitude dynamics at a level sufficient for mission design. Also, some satellites have suffered from degraded performance due to an incomplete understanding of PMAC system design. This dissertation alleviates these issues by discussing the design, inputs, and validation of PMAC systems for small satellites.

Design rules for a PMAC system are defined using the Colorado Student Space Weather Experiment (CSSWE) CubeSat as an example. A Multiplicative Extended Kalman Filter (MEKF) is defined for the attitude determination of a PMAC satellite without a rate gyro. After on-orbit calibration of the off-the-shelf magnetometer and photodiodes and an on-orbit fit to the satellite magnetic moment, the MEKF regularly achieves a three sigma attitude uncertainty of 4 degrees or less. CSSWE is found to settle to the magnetic field in seven days, verifying its attitude design requirement.

A Helmholtz cage is constructed and used to characterize the CSSWE bar magnet and hysteresis rods both individually and in the flight configuration. Fitted parameters which govern the magnetic material behavior are used as input to a PMAC dynamics simulation. All components of this simulation are described and defined. Simulation-based dynamics analysis shows that certain initial conditions result in abnormally decreased settling times; these cases may be identified by their dynamic response. The simulation output is compared to the MEKF output; the true dynamics are well modeled and the predicted settling time is found to possess a 20 percent error, a significant improvement over prior simulation.

Dedicated to Alana, my shelter in the storm.

Acknowledgements

Thank you to my advisor, Scott Palo. You have been a continual source of encouragement and guidance when it was sorely needed (which was always). Thank you Xinlin Li for always believing in my abilities (even when I didn't have them yet). Thank you Hanspeter Schaub for continuing to teach me attitude dynamics from Blacksburg to Boulder. Thanks as well to Albin Gasiewski, Paul Graf, and Delores Knipp for your guidance.

Thank you CSSWE team. I will never forget the time we spent working together to produce an operational(!) satellite. I enjoyed every minute of serving as your system engineer. I know our friendship will continue long after the beacons cease. Thank you Therese Jorgensen and the National Science Foundation for funding the little box that could.

Thank you to my parents, who set the bar high with their own example. You instilled the curiosity that prompts me to open the black box, as messy as the process may be.

Finally, the greatest thanks goes to my wife Alana. We did it. How it happened I'll never know, but we did it together. I can't wait to see where we end up next.

Contents

Chapter

1	Introduction	1
2	Basic Theory	5
2.1	Euler's Rotational Equation of Motion	5
2.2	Magnetic Theory	6
2.2.1	Magnetizing Field H vs. Magnetic Flux Density B	6
2.2.2	Hysteresis Loops	8
2.2.3	Magnetic Property Dependencies	10
2.2.4	Demagnetizing Fields	11
2.2.5	Magnetic Torques	16
2.2.6	Hysteresis Rods	17
3	Background	19
3.1	Mission History	19
3.1.1	Early History of Passive Magnetic Attitude Control	20
3.1.2	Modern Use of Passive Magnetic Attitude Control	22
3.1.3	CubeSats using Passive Magnetic Attitude Control	26
3.2	Analytical Models	29
3.2.1	Fischell Analytical Model (1961)	29
3.2.2	Mesch et al. Analytical Model (1966)	32

3.2.3	Kammüller Analytical Model (1971)	32
3.3	Numerical Simulations	33
3.3.1	Chen (1965)	34
3.3.2	Ovchinnikov & Penkov (2002) - Munin	34
3.3.3	CUBESIM (2004) and SNAP (2009)	35
3.3.4	Park et al. (2010) & Lee et al. (2011) - RAX	36
3.4	Hysteresis Measurement to Date	36
4	The Colorado Student Space Weather Experiment	38
4.1	Science Mission Success	38
4.2	Coordinate System	40
4.3	Sensors and Telemetry	43
4.3.1	Housekeeping	43
4.3.2	Attitude	43
4.4	Latch-up Anomaly	52
5	Control System Design	55
5.1	Maximum Expected Environmental Torques	55
5.2	Bar Magnet Design	56
5.3	Hysteresis Rod Design	57
6	Attitude Determination	61
6.1	Filter Design	61
6.2	Filter Tuning	67
6.2.1	Simulation-based Filter Tuning	67
6.2.2	Empirical Filter Tuning	70
6.3	CSSWE Attitude Determination	75
6.3.1	Filter Output	75

6.3.2	Attitude Determination Validation	89
7	Magnetic Measurement	98
7.1	Helmholtz Cage	98
7.1.1	Theory	99
7.1.2	Design	101
7.1.3	Assembly	104
7.1.4	Characterization	108
7.2	Bar Magnet Measurement	110
7.3	Hysteresis Measurement	112
7.3.1	Theory	112
7.3.2	Setup	115
7.3.3	Method	119
7.3.4	Results	121
8	Simulation	134
8.1	Components	135
8.1.1	Frames	135
8.1.2	Attitude Parameters	136
8.1.3	Equations of Motion	138
8.1.4	Orbit Propagation	138
8.1.5	Inertial Vector Models	139
8.1.6	External Torque Estimation	143
8.1.7	Numeric Integrators	151
8.2	Considerations	155
8.2.1	Torque-Free Rigid Body Motion	155
8.2.2	3D Pendulum Comparison	156
8.3	Simplified Simulation Analysis	159

8.3.1	Description	159
8.3.2	Energy Conservation Analysis	160
8.3.3	Angular Error Analysis	161
8.3.4	Settling Time Analysis	163
8.3.5	Summary	171
8.4	Results	174
8.4.1	Nominal Input	174
8.4.2	Nominal Output	176
8.4.3	High-Order Integrator Comparison	183
8.4.4	On-Orbit Data Comparison	183
9	Conclusion	192
9.1	Summary	192
9.2	Recommendations	194
	Bibliography	196
	Appendix	
A	Notation	202
B	Explicit Runge-Kutta Integrator Definitions	203

Tables

Table

4.1	CSSWE Housekeeping Sensors	44
5.1	Expected 3U CubeSat Environmental Torques	56
7.1	Helmholtz Cage Hardware	104
7.2	HyMu-80 Hysteresis Parameters	133
8.1	Simplified Simulation Input Sets	160
8.2	PMAC Simulation Inputs	175
B.1	General Butcher Tableau	204
B.2	RK2: Midpoint Method	204
B.3	RK3: Kutta Method	204
B.4	RK4: Runge-Kutta Method	204
B.5	RK5: Dormand-Prince Method	204
B.6	RK6: Hammund Scheme	205
B.7	RK7: Fehlburg Method	205

Figures

Figure

2.1	Magnetic Domains	7
2.2	Example B vs. H Hysteresis Loop	9
2.3	Effect of Magnetization Cycle Magnitude	12
2.4	Effect of Magnetization Cycle Magnitude	13
2.5	Hysteresis Loop Frequency Variation	14
2.6	Bar magnet \mathbf{H} and \mathbf{B} with no applied field	15
3.1	Transit 1B Satellite	21
3.2	Injun 3 Satellite	22
3.3	Azur Satellite	23
3.4	Munin Satellite	25
3.5	UNISAT-4 Satellite	25
3.6	QuakeSat CubeSat	27
3.7	RAX-1 CubeSat	28
4.1	CSSWE CubeSat & P-POD	39
4.2	REPTile Instrument	41
4.3	CSSWE Coordinate System	42
4.4	Magnetometer Position	46
4.5	Magnetometer Error	48

4.6	Photodiode Illumination	49
4.7	Photodiode Degradation and Scale Factor Correction	51
4.8	Investigation of TLE Position Error	53
4.9	ADC1 Latch-up Bit Damage	54
5.1	B -Field Error Angle	59
5.2	CSSWE PMAC Design	60
6.1	MEKF Simulation EA123 Error	68
6.2	MEKF Simulation Angular Velocity Error	69
6.3	MEKF Simulation Magnetometer Residuals	71
6.4	MEKF Simulation Photodiode Residuals	72
6.5	MEKF Simulation Error Angles	73
6.6	Magnetometer Position and Orientation	76
6.7	MEKF Empirical Magnetometer Residuals	77
6.8	MEKF Empirical Photodiode Residuals	78
6.9	MEKF Empirical Error Angles	79
6.10	First Month: Angular Velocity	81
6.11	First Month: Beta Angle	82
6.12	First Month: Energy	83
6.13	First Month: Histogram	84
6.14	First Month: Error Angle Magnitude	87
6.15	First Month: C&DH Temperature	88
6.16	Pre-settling Orbit: Beta	90
6.17	Pre-settling Orbit: Angular Velocity	91
6.18	Post-settling Orbit: Beta	92
6.19	Post-settling Orbit: Angular Velocity	93
6.20	Post-settling Orbit: Sun and Error Angles	94

6.21	Antenna Deployment Event	96
6.22	Solar Panel Temperature Distribution	97
7.1	Helmholtz Geometry	100
7.2	Helmholtz Coil Spacing Analysis	103
7.3	As-Built Theoretical Helmholtz Performance	105
7.4	Finished Helmholtz Cage	106
7.5	Helmholtz Support Hardware Chain	109
7.6	XY Plane Test	111
7.7	Bar Magnet Measurement Fit	113
7.8	Integrator Circuit for Magnetic Measurement	115
7.9	Hysteresis Measurement Setup Block Diagram	116
7.10	Hysteresis Measurement Sense Coil	118
7.11	Hysteresis Measurement Analysis Block Diagram	122
7.12	Fitted $\pm 100\text{A/m}$ Hysteresis Loop	123
7.13	Isolated Rod Performance	125
7.14	On-Orbit Hysteresis Rod Magnetizing Field Component	127
7.15	Simulated Loops from $\pm 100\text{ A/m}$ Measure	128
7.16	Simulated Loops from $\pm 10\text{A/m}$ Measure	129
7.17	System Hysteresis Rod Measurement	130
7.18	System Rod Performance	131
7.19	Bar Magnet Offset at Each Hysteresis Rod	132
8.1	ECI & ECEF Coordinate Frames	137
8.2	Global IGRF Magnitude	141
8.3	Hysteresis Model Output at ± 8 and $\pm 3\text{ A/m}$	147
8.4	Energy Ellipse and Momentum Sphere	157
8.5	Integrator Energy Conservation	162

8.6	Integrator Beta Angle Error	164
8.7	Integrator Euler Angle Error Breach Time	165
8.8	Settling Time Convergence	167
8.9	“Truth” Settling Time for Bar-Magnet-and-Hysteresis Case	169
8.10	“Truth” Settling Time for All-Torques Case	170
8.11	“Truth” Angular Velocity Response for Bar-Magnet-and-Hysteresis Case	172
8.12	“Truth” Angular Velocity Response for All-Torques Case	173
8.13	Nominal Output: Angular Velocity	177
8.14	Nominal Output: Energy	178
8.15	Nominal Output: Beta Angle	179
8.16	Nominal Output: External Torques	181
8.17	Nominal Output: Hysteresis Loops	182
8.18	RK4/7 Comparison: Angular Velocity	184
8.19	RK4/7 Comparison: Energy	185
8.20	RK4/7 Comparison: Beta	186
8.21	Orbital Data Comparison: Angular Velocity	188
8.22	Orbital Data Comparison: Energy	189
8.23	Orbital Data Comparison: Beta	191

Chapter 1

Introduction

Imagine lacking the ability to control your direction; it would be difficult to accomplish much of anything. The first US satellite, Explorer 1, suffered from this problem: due to flexible antenna on the craft, it began rotating orthogonal to its designed spin. The orientation, or attitude, control of a satellite can be difficult. However, satellite attitude control is crucial for satellite use in general; measurement, communication, propulsion, and much more are directly related to the satellite attitude. For these reasons, as satellites progressed beyond Explorer 1, so did attitude control methods. Early methods were passive: spin stabilization (spin about one axis fast enough so minor torques are negligible, much like a top), gravity gradient (aligns with the earth nadir direction due to differential gravity acting on the body), or passive magnetic attitude control, the subject of this dissertation.

Passive magnetic attitude control (PMAC) has been in use since the early 1960's. However, this does not mean it has outlived its usefulness. Far from it; PMAC remains a useful tool available to the spacecraft designer. As spacecraft decrease in size so does available power from the solar cells; the attitude control design space shrinks. Thus, many nanosatellite developers choose PMAC systems for the following reasons: the simplicity of installation (no processor running control laws is needed), low mass (often less than 5% of the satellite mass), zero power use, and alignment with the local magnetic field. However, PMAC is often little understood by many developers, especially student teams. This lack of understanding has led to poor design and inadequate performance prediction. This research grew from the desire to understand this control method to avoid design

pitfalls and allow for improved performance estimation.

Passive magnetic attitude control is the use of a magnet to torque the spacecraft towards earth's local magnetic field in conjunction with a dampening method. The dampening method most often used are hysteresis rods: soft magnetic material which is easily magnetized by the earth's local field. Because of satellite rotation, the direction of the local magnetic field relative to the hysteresis rod changes over time. This change in the direction of the local magnetic field changes the magnetization the rod, which decelerates the satellite angular velocity as rotational energy is converted to heat between the magnetic domains of the rod.

Research of the mechanics and simulation of Passive Magnetic Attitude Control is presented. A simulation of the rotational response of a satellite using a PMAC system is developed. This simulation is intended to predict the response of a PMAC system and thus is useful in design of future missions. Good PMAC system design can avoid negative consequences such as pointing offset error (a constant angular offset from the local magnetic field) or increased settling time (the time duration from initial orbit insertion to closely tracking the local magnetic field). The settling time is an important factor for small satellites which typically have a total mission lifetime of a few months; the success of such missions will be greatly hampered if they need to wait months or even years for attitude alignment.

Because an understanding of rotational motion and magnetic theory is crucial to an understanding of PMAC, these concepts are presented in Chapter 2. The difference between magnetizing field \mathbf{H} and magnetic flux density \mathbf{B} is described and applied to hysteresis rods to determine the relation between the hysteresis loop and the rod-based magnetic moment \mathbf{m}_{hyst} , which is crucial for the determination of hysteresis rod magnetic torque vs. the applied field. Although most small satellite developers to date have used material reference hysteresis loops, a vastly more realistic loop is based on empirical measurement of the flight hysteresis rods; this argument is developed throughout this research.

Chapter 3 presents an overview of PMAC history. A timeline of select missions is presented in an effort to understand which mission types are best for a PMAC system. The analytical models

which have been developed for some of these missions are shown, along with any assumptions which the models have used. The same treatment is given to the numerical models which have been used in the past. Finally, an overview of previous hysteresis rod measurement is presented.

Chapter 4 describes the Colorado Student Space Weather Experiment (CSSWE), a 3kg nanosatellite funded by the National Science Foundation (NSF) for space weather investigation. This CubeSat used a PMAC system for attitude control and serves as a concrete satellite example throughout this dissertation.

Chapter 5 presents the design of a PMAC system for the CSSWE CubeSat. The sizing and distribution of the bar magnet and hysteresis rods are discussed. Although applied to a CubeSat, design rules from this chapter are useful for any small satellite using a PMAC system.

Chapter 6 discusses a Multiplicative Extended Kalman Filter (MEKF) developed for attitude determination of a PMAC satellite. After the filter is defined, both simulation- and empirical-based tuning is performed using the CSSWE CubeSat. The on-orbit attitude performance of the CSSWE CubeSat is shown. The MEKF output is then verified by independently-measured telemetry.

Chapter 7 outlines the design of a Helmholtz cage and describes a variety of magnetic measurements applicable to a PMAC system. The cage is built such that a 3U CubeSat will experience a 99% uniform field across its length. After manufacturing, the Helmholtz cage setup is used to measure the magnetic moment of the CSSWE flight bar magnet. The Helmholtz cage setup is then used to measure CSSWE hysteresis rods. The hysteresis loops fitted to these empirical data are used as inputs to the dynamics simulation.

Chapter 8 presents a simulation developed to predict the attitude response of a satellite using a PMAC system. The components of the simulation are presented in detail such that mission teams may recreate the full simulation. This chapter also covers attitude dynamics which apply to a PMAC satellite; understanding these effects are helpful in interpreting the simulation output. A simplified simulation is developed to investigate the limitations and expected performance of the full simulation. The simplified simulation is also useful in selecting a numeric integrator and time step which achieve acceptable accuracy at a realistic computational cost. The results of the full

simulation are shown and compared to simulation output generated using a higher-order integrator.

The full simulation results are also compared to the on-orbit attitude data filter output.

Chapter 2

Basic Theory

Understanding Passive Magnetic Attitude Control (PMAC) begins with an overview of the underlying equations. This chapter serves as a review of the physics governing the components of a PMAC system. First, the basic equation for all rotational motion problems is defined. Next, ferromagnetic theory is reviewed; this is the foundation necessary for the design and study of PMAC systems. This dissertation uses the notation defined in Appedix A.

2.1 Euler's Rotational Equation of Motion

Regardless of attitude parameters used to describe rotational motion, the response of a satellite (or any rigid body) is given by Euler's rotational equations of motion [65]:

$$[I]\dot{\boldsymbol{\omega}} = -[\boldsymbol{\omega} \times][I]\boldsymbol{\omega} + \mathbf{L} \quad (2.1)$$

where $[I]$ is the 3×3 inertia matrix of the rigid body about its center of mass, $\boldsymbol{\omega}$ is the 3×1 body angular velocity vector, $\dot{\boldsymbol{\omega}}$ is the 3×1 derivative of the body angular velocity vector, \mathbf{L} is the 3×1 external torque vector, and $[\cdot \times]$ is the skew-symmetric matrix operator, defined as follows [65]:

$$[\mathbf{x} \times] = \begin{bmatrix} 0 & -x_3 & x_2 \\ x_3 & 0 & -x_1 \\ -x_2 & x_1 & 0 \end{bmatrix}. \quad (2.2)$$

Choosing a body-fixed coordinate system which aligns with the principal body axes results

in a diagonal inertia matrix, which simplifies Equation 2.1 as follows [65]:

$$\begin{aligned}
 I_{xx}\dot{\omega}_x &= -(I_{zz} - I_{yy})\omega_y\omega_z + L_x \\
 I_{yy}\dot{\omega}_y &= -(I_{xx} - I_{zz})\omega_z\omega_x + L_y \\
 I_{zz}\dot{\omega}_z &= -(I_{yy} - I_{xx})\omega_x\omega_y + L_z
 \end{aligned} \tag{2.3}$$

where the subscript represents the component aligning with a specific principal body axis. Note that Equation 2.3 shows that there will be angular velocity coupling for any non-symmetric rigid body. Equations 2.1 and 2.3 are the basis for all of the analytical models and simulations to follow. The difficulty in modeling is in correctly representing the external torque \mathbf{L} applied to the system.

2.2 Magnetic Theory

2.2.1 Magnetizing Field \mathbf{H} vs. Magnetic Flux Density \mathbf{B}


There exist two separate but closely related fields which both, at times, go by the name “magnetic field”: the magnetizing field \mathbf{H} and the magnetic flux density \mathbf{B} . In the SI system of units, \mathbf{H} is in units of A/m while \mathbf{B} is in units of Tesla (T). The relative permeability is defined as $\mu_r = \frac{\mathbf{B}}{\mu_0\mathbf{H}}$ where the permeability of free space $\mu_0 = 4\pi \cdot 10^{-7}$ T·m/A. For most materials, μ_r is very close to unity, meaning the material does not increase \mathbf{B} appreciably in response to \mathbf{H} . However, within ferromagnetic material the situation is much different. The following definition relates \mathbf{B} and \mathbf{H} : [17]

$$\mathbf{B} \equiv \mu_0(\mathbf{H} + \mathbf{M}) \tag{2.4}$$

where \mathbf{M} is the magnetization of the material within which the fields are present. Magnetization is defined as $\mathbf{M} = \mathbf{m}/V$ where \mathbf{m} is the magnetic moment and V is the magnetized volume; thus the magnetization \mathbf{M} is the magnetic moment \mathbf{m} density. For materials with μ_r close to unity (such as air), \mathbf{M} is very close to zero, and Equation 2.4 reduces to $\mathbf{B} = \mu_0\mathbf{H}$.

However, a ferromagnetic material has a non-zero magnetization which changes in response to an applied field. This change in \mathbf{M} is due to microscopic changes within the material. Figure 2.1

Figure 2.1: Example magnetic domains are shown. With zero applied field (left), the domain orientation of the magnetic material is such that the sum magnetization of the material is small. However, when a magnetizing field is applied to the material (right), the domains oriented parallel to that field grow as the out-of-alignment domains shrink.




is an example of multiple microscopic magnetic domains within a material. Each magnetic domain is composed of a group of atoms with parallel magnetic moments. When zero magnetizing field is present, the sum magnetization of the material is small because the magnetic moments of multiple domains are mostly canceled. However, when a magnetizing field is applied, a domain with a magnetic moment parallel to the field will grow as the magnetic moments of atoms close to its boundaries align to the applied field. This change in total magnetization \mathbf{M} is non-linear with changes in \mathbf{H} , and is responsible for the familiar hysteresis loop [8].

2.2.2 Hysteresis Loops

When a ferromagnetic material is subjected to a changing magnetizing field H and the magnetic flux density B is measured, plotting B vs. H will result in a hysteresis loop such as the one shown in Figure 2.2. A major hysteresis loop may be defined by three parameters: the coercivity H_c , the remanence B_r , and the saturation B_s . The coercivity is the applied field necessary to bring the B field to zero, or the x-axis intercept. The remanence B_r is the remaining B within the material when H has been decreased to zero, or the y-axis intercept. There exists a maximum value of M for a given material. If the applied field H is subtracted from B , the curve $B/\mu_0 - H$ vs. H will asymptotically approach this maximum magnetization, known as saturation [10]. At saturation (and only at saturation), the magnetization M within a bar or cylinder sample is constant and uniform. This is because all of the individual magnetic domains within the material have aligned in the same direction. The material cannot supply any more magnetization because there are no more domains to align. Thus, after saturation, the increase in B -field is solely due to the increase in magnetizing field, and thus has a slope of μ_0 . The point on the hysteresis curve at which B_s starts to increase with slope μ_0 is the saturation flux density. The area enclosed within the hysteresis loop is an important feature; it represents the energy absorbed by the magnetic material per unit volume as it completes one magnetization cycle.

A hysteresis loop may be split into a lower and upper curve which are generated by the increasing and decreasing sections of magnetizing field cycling, respectively. The hysteresis fitting

Figure 2.2: An example magnetic flux density B vs. magnetizing field H hysteresis loop. The coercivity H_c , remanence B_r , and saturation B_s are shown. The area encircled by the hysteresis loop is the energy loss per cycle per unit volume. After saturation B_s , the slope of the hysteresis loop is simply μ_0 .



described in Section 7.3 relies upon this bifurcation. The upper and lower curves of the hysteresis loop will be odd-symmetric if the cycle amplitude remains constant and there is no DC offset in the magnetizing field. This property will also be used in the hysteresis fitting.

2.2.3 Magnetic Property Dependencies

The shape of a magnetic hysteresis loop depends on many things, but some of the more important factors include: material composition, degree of heat treatment, applied H -field extrema, applied field offsets, frequency of H -field cycling, and sample dimensions. The material itself governs the saturation magnetization amplitude (the applied field at which saturation occurs and the shape of the curve from 0 A/m to saturation are not determined solely by the material as they are structure-specific); the inherent crystal structure of the material composition defines “easy” directions of magnetization [17]. Note that most nickel-iron alloys (such as HyMu-80) have low magnetocrystalline anisotropy after heat treatment, meaning the ease of magnetization is about the same regardless of direction [5].

Heat treatment can restore the crystalline structure of a material that is damaged during cold work, such as extruding, rolling or bending. Heat treatment also serves to break down the walls between magnetic domains, increasing the mean domain size and allowing the magnetic material to be magnetized to higher levels at lower magnetizing fields. Hysteresis loops before and after heat treatment will likely be very different.

The magnitude of the applied field cycle will change the resultant hysteresis loop [17]; this effect is included within the Flatley hysteresis model (described in Section 8.1.6.5). Figure 2.3 shows the output of the Flatley hysteresis model for cycle amplitudes of ± 2 A/m, ± 3 A/m, and ± 8 A/m. Note that the loop area decreases substantially as the applied field decreases. A constant offset in the applied field can also distort the hysteresis loop. In the case of a satellite PMAC system, such a constant offset may be provided by proximity to current loops or a bar magnet. As shown in Figure 2.4, an H -field offset pushes the hysteresis loop away from the origin. If the cycle amplitude approaches the material saturation, this magnetizing field offset can result in a smaller

loop area, and thus, decreased dampening.

The frequency of applied field cycling can effect the hysteresis loop measured. Figure 2.5 shows how an increase in cycle frequency tends to increase the coercivity H_c . More energy is used to switch the magnetic domains at higher cycle frequencies. However, a DC hysteresis curve (usually defined as an applied field cycle frequency of 10 Hz or less) is minimally affected by frequency. All hysteresis loops measured in Section 7.3 are produced by applying a field with a cycle frequency of less than 1 Hz. The hysteresis loop is also affected by the demagnetizing field of the test sample, which is further examined in the following section.

2.2.4 Demagnetizing Fields

Anyone who has handled magnets is familiar with the idea of a magnetic pole. Consider a bar magnet that has been magnetized by a magnetizing field H in the left to right axial direction. After H has been removed, there exist two magnetic poles, the south pole on the left and north pole on the right. Figure 2.6 shows the H and B fields resultant of the poles. As shown, there exists an H field outside and inside the magnet. Outside the magnet, the simple relation $B = \mu_0 H$ holds. However, inside the magnet, an H field opposes the B field and is termed the demagnetization field H_d . Equation 2.4 becomes $B = \mu_0(-H_d + M)$. If an external applied field H_a is present, Equation 2.4 becomes:

$$B = \mu_0(H_a - H_d + M). \quad (2.5)$$

Demagnetizing fields are difficult to calculate, but are directly proportional to the magnetization of the bar magnet: $H_d = N_d M$ where N_d is the demagnetization factor. The demagnetization factor varies mainly as a function of the length to diameter ratio L/D of the sample, but also varies as a function of magnetization: different N_d values are used at values close or far from saturation. These limitations make it especially complicated to calculate the hysteresis loop of a material; empirical determination is much more accurate.

One way to measure the hysteresis loop of a material without having to account for demagnetizing fields is to use a toroid-shaped sample. Lines of magnetic flux density B , which follow closed

Figure 2.3: Example B vs. H hysteresis loops for various applied field cycle magnitudes. Figure was generated using the Flatley hysteresis model and the HyMu-80 closed magnetic circuit hysteresis parameters (see Table 7.2).




Figure 2.4: Example B vs. H hysteresis loops for various applied field DC offsets. Figure was generated using the Flatley hysteresis model and the HyMu-80 closed magnetic circuit hysteresis parameters (see Table 7.2). All three datasets are generated using an AC magnetizing field cycle amplitude of ± 2 A/m.




Figure 2.5: The effect of applied field cycle frequency on the hysteresis loop. The coercivity H_c and the loop area increase as the hysteresis loop is cycled at increasing frequency. Image used from [8].







Figure 2.6: (a) \mathbf{H} -Field and (b) \mathbf{B} -Field of a bar magnet when there is zero applied field. Note that $\mathbf{M} > 0$ only within the magnet, and that outside the magnet, $\mathbf{B} = \mu_0 \mathbf{H}$. Figure adapted from [17].



loops, lie entirely within such a sample. This type of sample is known as a closed magnetic circuit. This means that, even when magnetized, no poles are present in such a sample, and thus no demagnetizing field H_d is present in the sample. The hysteresis loop measured from a toroidal sample would be B vs. H_{true} . The true magnetic field within the material is a combination of the applied magnetizing field and the demagnetizing field generated by the material: $H_{\text{true}} = H_a - H_d$. Of course, if $H_d = 0$ (as in a toroidal sample), then $H_{\text{true}} = H_a$. However, this is not the case for open magnetic circuit samples. Because the closed magnetic circuit hysteresis loop B vs. H_{true} is invariant with respect to the dimensions of the material, this hysteresis loop is generally what is quoted on many material data sheets.

However, the closed magnetic circuit hysteresis loop is not a good representation of the true open magnetic circuit hysteresis loop for the rods or strips typically used in a PMAC system. Testing to date has shown that the loop areas differ by one to two orders of magnitude [63]. As the loop area is directly related to the dampening provided by each hysteresis rod, this has vast implications for a numerical simulation of the attitude dynamics. Section 7.3 presents measurements of the true open magnetic circuit loops for hysteresis rods which are typically used in PMAC systems.

2.2.5 Magnetic Torques

All magnetic torques obey the following formula:

$$\mathbf{L} = \mathbf{m} \times \mathbf{B} \quad (2.6)$$

where \mathbf{m} is the magnetic moment vector and \mathbf{B} is the local magnetic flux density vector. In most situations, \mathbf{B} is due to earth's local magnetic field alone. Thus, the torque is based on the value of \mathbf{m} for various magnetic materials. The high coercivity of a permanent magnet prevents the earth field from changing its magnetization, thus for a bar magnet, \mathbf{m} is constant and may be measured (see Section 7.2). However, determining the magnetic moment of a hysteresis rod is more difficult.

2.2.6 Hysteresis Rods

A PMAC system necessarily uses bar- or cylinder-shaped hysteresis material. This is required because such samples are magnetized mainly in the axial direction, which produces a torque as defined by Equation 2.6. For these open magnetic samples the demagnetizing field H_d is unavoidable, and the L/D ratio is limited by the dimensions of the spacecraft and the necessary volume of hysteresis material. This means that B measured for an open magnetic circuit (with demagnetization) will be significantly less than B measured for a closed magnetic circuit.

As shown in Equation 2.6, the component of \mathbf{B} parallel to the magnetic moment \mathbf{m} does not produce a torque. Assuming the majority of uniform magnetization is parallel to each rod's axis, the rod produces a negligible B-field perpendicular to its axial direction. Thus, sets of rods which are co-planar but have some separation should have minimal interaction. The general rule of thumb for hysteresis rod placement is that the perpendicular distance between two rods should be greater than 30-40% of their length [57]. Given this separation, the assumption is made that the interaction of multiple hysteresis rods may be ignored in analysis. Thus, the magnetic flux density \mathbf{B} on the right side of Equation 2.6 is generated solely by the local earth B -field. This means that the magnetic moment \mathbf{m} is the only characteristic of the rod which contributes to the torque produced by the rod.

Therefore, in order to model the torque due to the hysteresis rods, the magnetic moment \mathbf{m} of the rod must be defined. First, it is assumed that the magnetic moment is entirely parallel to the rod. Using $m = M/V$, $H_d = -N_d M$ and Equation 2.5, the magnetic moment parallel to the rod for an open magnetic circuit may be defined as [17]:

$$m = V \left(\frac{B/\mu_0 - H_a}{1 - N_d} \right) \quad (2.7)$$

where B is the average parallel magnetic flux density within the rod and H_a is the applied field parallel to the rod. Many groups ([47],[61],[58],[63],[26]) simplify Equation 2.7 to $m = VB/\mu_0$. It is feasible to ignore N_d in the denominator for rods with $L/D > 30$ [17], as the error is $< 2\%$ (the demagnetizing field has not been ignored, it is taken into account by empirically measuring B).

However, the hysteresis curve must be measured to ensure H_a is negligible with respect to B/μ_0 . None of the groups reviewed in Section 3.3 state their assumptions in using the simplified formula. Some make the grave mistake of assuming B within Equation 2.7 is given by a material data sheet (usually B vs. H_{true}). This assumption saves one from having to measure the hysteresis loop, but also introduces serious errors because it ignores the effect of the demagnetizing field. The B used within Equation 2.7 is the average interior magnetic flux density (across the length of the sample), and must be measured for the open magnetic circuit to produce an accurate magnetic moment m . The B vs. H_a hysteresis loop is used in the PMAC dynamics simulation (see Section 8) because it directly relates B to the earth H -field encountered by the spacecraft, which is H_a .

Chapter 3

Background

Although difficult to model due to the nonlinear behavior of hysteresis, Passive Magnetic Attitude Control (PMAC) is simple to realize on a spacecraft: it only requires the installation of a bar magnet and a few hysteresis rods. Thus, early satellites made considerable use of PMAC. Small satellites echo this trend today as electronic components shrink faster than attitude control actuators. The previous PMAC analysis is traced through the mission history of PMAC satellites as well as the analytic and numeric models which have been developed in parallel with these missions. A review of PMAC hysteresis rod measurement to date is presented at the end of the chapter.

We find no previous work which verifies the performance of a PMAC attitude dynamics simulation through comparison to on-orbit attitude data from a PMAC satellite. Neither do we find previous work which measures hysteresis rod performance as affected by the system-level PMAC component magnetization. There is much to be learned in investigating these avenues.

3.1 Mission History

It is not surprising that small satellites today echo the development of large satellites in the early space race. Although satellite electronics shrink very quickly, it is hard for precision actuation devices to keep pace. PMAC is in use today just as it was when it was designed in the 1960's. Studying the missions which have used PMAC sheds light on the type of missions for which it is best suited, and tells the story of its development over time.

3.1.1 Early History of Passive Magnetic Attitude Control

Passive attitude control systems were used in early spacecraft because software and actuation hardware were not yet developed. As spacecraft developed, simple attitude control systems such as spin-stabilization gave way to more complicated (yet still passive) methods, such as gravity-gradient and passive magnetic attitude control.


Passive Magnetic Attitude Control (PMAC) was first used in space in April of 1960 [20]. Researchers at Johns Hopkins Applied Physics Laboratory came up with the method for the Navy's Transit experimental satellite program. The first satellite to use a passive magnetic attitude control system was Transit 1B (Transit 1A did not achieve orbit due to a launch failure); it used PMAC to point the spacecraft toward ground stations in the northern hemisphere. Transit 1B was a spherical satellite with a $10 \text{ A}\cdot\text{m}^2$ bar magnet and two sets of permeable rods, both arranged in a crosshatch pattern. Both sets of permeable rods had 4 rods each, with the second set rotated 45° from the first set, yet still within a plane perpendicular to the bar magnet. The satellite had an initial spin of 17.5 rad/s , which was reduced to 16.3 rad/s after 7 days (this decrease is likely due to the PMAC system). A mechanical de-spin via release of weights from the spin axis reduced the spin rate to 0.5 rad/s . The satellite spin rate decreased to 0.03 rad/s (1.8 deg/s) after another 7 days. The Transit 1B PMAC system was successful enough for Transit 2A to be launched without the mechanical de-spin included; it relied solely on the PMAC system. Transit 2A decreased from 5.0 rad/s to 0.13 rad/s (7.2 deg/s) within 24 days [25]. The analytical model used to analyze the PMAC performance of both Transit satellites is presented in Section 3.2.1.

Due to the success of the Transit 1B and 2A, PMAC was used for Injun 1, the first satellite built entirely by a university [27]. Launched in 1961, Injun 1 failed to separate from GREB, another satellite that was on the same launch [9]. A later satellite from the University of Iowa under the supervision of Dr. Van Allen, Injun 3, was launched in 1963 into a $237 \times 2785 \text{ km}$, 70.4° inclination orbit. Its PMAC system aligned to an average deviation of less than 10° from the magnetic field line after a period of 2 weeks [29].

Figure 3.1: The Transit 1B Satellite, the first satellite to use Passive Magnetic Attitude Stabilization. [1]



Figure 3.2: The Injun 3 Satellite, an early university satellite which successfully aligned to the magnetic field using PMAC. [56]




The first German satellite Azur was launched in 1969 and carried a $97 \text{ A}\cdot\text{m}^2$ bar magnet. Azur damped from an initial spin rate of 1.25 rad/s down to 0.01 rad/s (0.6 deg/s) and a magnetic field offset of less than 15° within 2.25 days [53].

This section is not an exhaustive list. Other early satellites using PMAC include Transit 2A (1960), ESRO 1A (1968), ESRO 1B (1969), Exos (1978) and Magion (1978) [57]. By the mid 1970's, analysis of PMAC was mostly comprised of analytical models. The use of PMAC begins to wane at this point, as active control systems enable specialized pointing methods. The difficulty of accurately predicting PMAC performance was likely a driver of its decreased use.

3.1.2 Modern Use of Passive Magnetic Attitude Control

As computers and actuators decrease in size, the satellite community has grown increasingly interested in small satellites due to their low launch costs and simplicity. While technology is in development to control these satellites using small actuators, the small satellite community has witnessed a return to passive methods. These methods are especially popular among student missions or technology demonstrations where component price and complexity join size as major

Figure 3.3: Azur, the first German satellite, used passive magnetic attitude control [53].



design factors. PMAC is also popular among science missions which benefit from alignment with the local magnetic field.

The 6kg Swedish *Munin* satellite contained a PMAC system. During design, the satellite used the following requirement: align to within 15° of the earth's magnetic field lines within three weeks. Ovchinnikov developed a numerical attitude simulation and predicted that in order to meet the setting time requirement, the initial angular velocity had to remain less than 10.5 deg/s on each axis [57]. The satellite launched in November 2000 but contact was lost in February 2001. Johnsson indicates that results from the attitude determination analysis are questionable [38]. Attitude determination analysis from *Munin* has not been released to date. Thus, the Ovchinnikov attitude simulation has not been validated; details of the simulation are described in Section 3.3.2.

UNISAT-3 was designed and built by students at the University of Rome. It was launched into a 710km x 790km, 98° inclination orbit on June 29, 2005, and used a PMAC system. The PMAC system used a permanent magnetic with magnetic moment $1 \text{ A}\cdot\text{m}^2$ and one hysteresis rod per axis with dimensions 15cm \times 0.1cm. The only dedicated attitude measurement device was a 3-axis magnetometer. The magnetometer and solar panel currents were used to piece together its three-axis attitude. The Z-axis of the magnetometer did not function on-orbit, so the total magnetic field was estimated using the spacecraft position and the IGRF magnetic field model [62]. UNISAT-3 oscillated about the earth magnetic field at an amplitude of about 30° . The team believed this response was due to under performing hysteresis rods [63].

UNISAT-4 was next in the series of educational satellites. Researchers at the University of Rome determined that the magnetic properties of the hysteresis rods must be measured, as specific dimensions and orientations of the rods could change their performance. They developed a rig to measure the hysteresis parameters of the rods. After careful design, they used a bar magnet of $1 \text{ A}\cdot\text{m}^2$ and eight hysteresis rods on both orthogonal axes to the bar magnet. The hysteresis rods had a square cross-section with dimensions 150mm \times 1mm \times 1mm and were composed of permalloy. After the rods were heat treated, measurements showed that their hysteresis parameters were well below the quoted hysteresis parameters for the material [63]. Unfortunately, due to the 2006 Dnepr

Figure 3.4: Artists conception of Munin, a Swedish small satellite which used PMAC [55].





Figure 3.5: The UNISAT-4 satellite, a student satellite built by the University of Rome [63].



launch failure, UNISAT-4 did not achieve orbit [32].

The next satellite from the University of Rome, EduSat, continued their work on PMAC development. Working with the University of Keldysh Institute of Applied Mathematics (KIAM), a new hysteresis parameter experimental set-up was developed. This set-up allows the hysteresis parameters to be measured along the length of the hysteresis rod. Researchers found that the maximum magnetic flux density was highest at the center of the hysteresis rod and decreased towards the ends [4]. The hysteresis rod measurements described in Section 7.3 use a sense coil with a length equivalent to the hysteresis rod length; this ensures that the average interior magnetic flux density of the rod is measured.


In 2001, the 23 kg Sapphire microsatellite was launched. Designed by Stanford university, it used a PMAC system to de-spin the satellite and ensure that an imaging sensor was pointed toward earth in the northern hemisphere. The communication antennas were also painted to impart a small radiation pressure torque which ensured a roll to prevent one side from always facing the sun (this attitude control is known as the "controlled tumble" and has been used for many AMSAT spacecraft). Sapphire was ejected from the launch vehicle with a tumble of multiple degrees per minute. This spin was reduced to 1.2 rpm about the major inertia axis with a few days due to the PMAC system. Radiation pressure caused the satellite to settle to 0.1 rpm [74].

3.1.3 CubeSats using Passive Magnetic Attitude Control

In summer of 2000, Bob Twiggs and researchers at Stanford university envisioned a new nanosatellite standard they called CubeSat [76]. This standard was soon accepted by universities across the country; PMAC was used as the stabilization method for many of these satellites. One of these early CubeSats that used PMAC was QuakeSat, built by Stanford university [52]. Launched in June 2003, QuakeSat relied upon solar panel currents and a single IR sensor to determine attitude. Unfortunately, the loss of a multiplexer early in the mission meant that the solar array currents were not available. QuakeSat used a $2.933 \text{ A}\cdot\text{m}^2$ bar magnet in combination with two $0.6\text{cm} \times 1.2\text{cm} \times 31\text{cm}$ rods of permalloy 49NM along the length of the satellite [69]. Using their single IR

sensor, the QuakeSat team estimated that their satellite was generally nadir pointing, with a roll rate of once every 15 - 20 min [7].

Figure 3.6: Artist's conception of the QuakeSat CubeSat. Built by Stanford University, QuakeSat used a PMAC system [52].




The Radio Aurora Explorer (RAX) was the first CubeSat funded by the National Science Foundation to study space weather [18]. Built by the University of Michigan, the RAX mission is actually composed of two satellites, RAX-1 and RAX-2; RAX-2 was launched to continue the science mission after the RAX-1 solar panels were found to be faulty [19].

The RAX team attempted to calculate the performance of its PMAC system using a dynamic model. Students at the University of Michigan developed the Lie Group Variational Integrator (LGVI). The LGVI is designed to model the rotation of a rigid body while conserving the constraints of the rotation matrix as well as the system energy [58]. More detail on this integrator is found in Section 8.1.7.2.

The RAX team used the LGVI to develop a simulation to predict the response of their CubeSat. However, they ignored the possibility of the satellite bar magnet saturating the hysteresis rods. To simplify integration, the team designed the PMAC system such that the bar magnet and hysteresis rods are on the same 10cm×10cm board which fits the form factor of other electronics

Figure 3.7: The RAX-1 3U CubeSat. Built by the University of Michigan, RAX-1 and RAX-2 used identical PMAC systems [19].



boards; this is not an optimal design as will be explained in Chapter 5. The team also used the closed-magnetic circuit hysteresis parameters as input to their numeric model; this results in a gross overestimate of the hysteresis dampening performance (see Section 7.3). These issues seriously degrade the predictive capability of the simulation.

The settling time of RAX-1 is not known as attitude data was only collected for three single-orbit periods, each 15 days apart [71]. However, RAX-2 attitude data shows that the satellite converges to within 20° of the local magnetic field two months after launch [73]. The predicted settling time of the RAX mission has not been published but personal communication with a member of the RAX team indicates the satellite was expected to align within days. This large discrepancy indicates that accurate PMAC dynamics simulation is both difficult and highly useful. Also, the RAX-2 satellite uses a bar magnet magnetic moment of $3.2 \text{ A}\cdot\text{m}^2$ [73]; this powerful magnet (relative to a CubeSat inertia matrix) may have introduced the high initial rotation rates experienced by RAX-2 (see Section 5.2).

3.2 Analytical Models

An analytical model of PMAC is complicated by the use of hysteresis rods, whose torque depends on both the current orientation and the previous magnetism induced within the rods. Further complication is introduced when considering a real earth field, which is difficult to model. To combat this, the analytical models derived below simplify the hysteresis effect by assuming some average damping, usually in the form of an angular velocity coefficient. The earth's magnetic field is simplified with either the dipole assumption or an average magnetic field strength.

3.2.1 Fischell Analytical Model (1961)

Robert Fischell derived the first analytical model of PMAC for the Transit satellite program [25]. He starts with the assumption of a completely symmetric satellite ($I = I_{xx} = I_{yy} = I_{zz}$) and the magnitude of magnetic torque $\tau = M\mu_0 H_0 \sin \theta$ where M is the magnetic moment of the bar magnet, μ_0 is the permeability of free space, H_0 is the earth's local magnetizing field, and θ

is the angle between the bar magnet and the earth field direction. Fischell chooses to ignore the hysteresis torque at first in order to form Euler's rotational equations of motion for an undamped satellite axis:

$$I \frac{d^2\theta}{dt^2} + M\mu_0 H_0 \sin \theta = 0. \quad (3.1)$$

Then, after making the small angle assumption for θ , the undamped angle relative to the magnetic field is given as $\theta_N = \theta_0 \cos(2\pi ft)$ where the natural frequency f is:

$$f = \frac{1}{2\pi} \sqrt{\frac{M\mu_0 H_0}{I}}. \quad (3.2)$$

Fischell, ignoring other disturbance torques, then defines the energy loss per time due to hysteresis cycling as:

$$\frac{dE}{dt} = -\frac{NV}{8\pi^2} \sqrt{\frac{M\mu_0 H_0}{I}} \oint H dB \quad (3.3)$$

where V is the volume of the hysteresis material, N is the number of rods, and $\oint H dB$ is the area of the hysteresis loop. Fischell obtained the hysteresis loop of the chosen rods experimentally and determined that the hysteresis loop area may be approximated as $\oint H dB = \alpha H_m^3$ where H_m is the peak magnetizing field of the hysteresis loop and α is an empirically-derived constant. The peak magnetizing field for a rod perpendicular to the local earth field is given as $H_m = H_0 \sin \theta_m$ where θ_m is the max angular displacement between the bar magnet and the earth field.

$$\frac{dE}{dt} = -\frac{\alpha NV}{8\pi^2} \sqrt{\frac{M\mu_0}{I}} H_0^{7/2} \sin^3 \theta_m. \quad (3.4)$$

Here Fischell chooses to define a constant $k = \frac{\alpha NV}{8\pi^2} \sqrt{\frac{M\mu_0}{I}} H_0^{7/2}$. Now, the potential energy of the bar magnet can be defined by integration of the magnitude of magnetic torque from equilibrium to the max displacement:

$$E(\theta) = \int_0^{\theta_m} M\mu_0 H_0 \sin \theta d\theta = M\mu_0 H_0 (1 - \cos \theta_m). \quad (3.5)$$

Now, taking the derivative of Equation 3.5 with respect to time, combining with Equation 3.4, and separating variables yields:

$$-\frac{d\theta_m}{\sin^2 \theta_m} = \frac{k dt}{M\mu_0 H_0} \quad (3.6)$$

which becomes the following after integration:

$$\cot \theta_m = \frac{k}{M\mu_0 H_0} t + C. \quad (3.7)$$

The constant can be solved for by defining an angle θ_0 to which the satellite is displaced at time $t = 0$. Thus, $C = \cot \theta_0$, and the following is the expression for the maximum angular displacement over time:

$$\theta_m(t) = \arccos \left(\frac{k}{M\mu_0 H_0} t + \cot(\theta_0) \right). \quad (3.8)$$

Finally, Fischell combines the undamped angle with the maximum damped angle to yield the angle of displacement from the magnetic field with respect to time:

$$\theta(t) = \arccos \left(\frac{k}{M\mu_0 H_0} t + \cot(\theta_0) \right) \cos 2\pi f t. \quad (3.9)$$

The settling time is easily found by modifying Equation 3.8:

$$t_{\text{settle}} = \frac{M\mu_0 H_0}{k} (\cot \theta_f - \cot \theta_0) \quad (3.10)$$

where θ_f is the final angular displacement. To recap, Fischell has defined an analytical solution with the following assumptions:

- Entirely symmetric satellite
- Small angle between B -field & bar magnet axis
- Orbit-average magnetic field strength used instead of position-dependent vector
- No other (non-magnetic) disturbance torques
- Cubic approximation of hysteresis area

The above assumptions are quite limiting. Even with a symmetric satellite, this analytical model is ineffective for the beginning of dampening, where high angular velocities are typical.

3.2.2 Mesch et al. Analytical Model (1966)

Mesch et al. developed a more comprehensive analytical model [22]. They start with Equation 2.3, then define the angular velocities in terms of Euler angle rates. However, here they make the assumption that the Euler angles are always a small angle for a satellite with attitude control. A unique aspect of this model is the development of the equations of motion in terms of the orbit true anomaly instead of time, allowing a dipole magnetic field to be included within the model. However, there are many assumptions made:

- Angle between B-field & bar magnet is a small angle
- Polar orbit (inclination = 90°)
- Dipole magnetic field
- Dampening torque is in constant proportion to angular velocity

Due to these assumptions, this model was generally used to determine the periodic motion of a satellite after it had settled to oscillating about the magnetic field.

3.2.3 Kammüller Analytical Model (1971)

Kammüller takes a different approach [39], [40]. Rather than starting from Euler’s rotational equations of motion (Equation 2.3), the Lagrangian is calculated assuming a 3-1-3 Euler angle set. By using the Lagrangian equations of motion, Kammüller is able to account for the gyroscopic torques of the spacecraft without solving six coupled equations of motion. For the Lagrangian, the potential energy is defined as Equation 3.5, but with a transformation used to convert θ_m to a function of Euler angles and the magnetic declination. A series expansion accounts for the magnetic field strength and the magnetic declination as a function of the orbital frequency and time. Here Kammüller introduces a “slow” time variable $\tau = \omega_e t$ where ω_e is the earth’s rotation rate. This assumption, based on the significant difference between the earth’s rotation rate and the orbital

frequency, allows Kammüller to treat some timescales as constant with respect to the “fast” time variable t .

Kammüller shows that, for near-polar orbits ($i \approx 90^\circ$), there exists a roll resonance for specific values of $\Delta = (I_{xx} - I_{yy})/I_{zz}$, where I_{xx} is the maximum moment of inertia, and I_{zz} is the minimum. He points out that there are three possible solutions to the roll equation (after transients have been damped): a) nonresonant rest-position, b) resonant oscillation, and c) resonant rotational solutions. The roll resonance is due to coupling between pitch and roll motions of the spacecraft, and thus changes depending on the spacecraft Δ . By changing the spacecraft Δ , the desired stability (a, b, or c) may be set.

Kammüller makes the following assumptions in his analysis:

- Circular polar orbit
- Dipole Earth field
- Pure pitch (major inertia axis) motion while following magnetic field lines
- “Slow” time used to consider diurnal rotation negligible with respect to orbital motion
- Hysteresis dampening described by matrix of dampening coefficients multiplied by Euler angles and Euler angle rates

Other analytical PMAC models could not be found in the literature. To date, none of the analytical models have solved for the settling time of a non-symmetric satellite.

3.3 Numerical Simulations

Numerical models have the advantage of not making the simplifying assumptions of the analytic models. As a result, numerical models have the potential to accurately predict the full dynamics of the system. However, numerical models have their own disadvantages; a balance must be sought between simulation accuracy and computational cost. Also, the model itself can introduce

errors if not properly defined. Numerical models which have been used in the past are presented below.

3.3.1 Chen (1965)

Chen [15] uses a 3-1-3 Euler angle set to describe the rotation of the body frame with respect to the inertial frame. An inclined dipole is used to model the earth field. An interesting note is the inclusion of another dampening torque. This “shorted coil” dampening torque is due to closed windings about each hysteresis rod which have current induced within them due to the earth field. This current, in turn, torques the satellite according to Equation 2.6. The hysteresis loop is modeled as a parallelogram, which does not account for minor hysteresis loops which occur as the satellite starts to track the earth field. After the torques are defined, Equation 2.3 is used to define the equations of motion. In order to avoid the singularity associated with Euler angles, quaternions are used to model the attitude during integration.

Chen uses the following assumptions:

- Dipole Earth field
- “Shorted Coil” dampening in addition to hysteresis rod dampening
- Parallelogram hysteresis loop

3.3.2 Ovchinnikov & Penkov (2002) - Munin

Ovchinnikov & Penkov investigate the motion of a 6 kg axisymmetric satellite with a bar magnet and hysteresis dampening [57]. First, a magnetic frame is defined by the direction of the local earth field and the orbit plane of the satellite. The earth field is modeled as a dipole, and a parallelogram model is used for the hysteresis loop. Equation 2.1 serves as the equation of motion. A 2-1-3 Euler angle is used as attitude parameters. The equations of motion are then written in dimensionless form and a few key assumptions are made: a strong bar magnet dominates the external torques and the initial angular velocity of the satellite about the symmetry axis is roughly

equivalent to the mean motion of the satellite. These assumptions allow the average equations of motion to be developed. These average equations of motion are investigated with numeric analysis.

Ovchinnikov & Penkov use the following assumptions:

- Strong bar magnet
- “Improved” Parallelogram model (able to generate minor loops near origin only)
- Averaged equations of motion

3.3.3 CUBESIM (2004) and SNAP (2009)

Both CUBESIM and the Smart Nano-satellite Attitude Propagator (SNAP) are PMAC attitude simulations developed by students (Levesque [47] and Rawashdeh [61], respectively) in an attempt to simulate the response of a PMAC system for satellites they were working on at the time. Both models use the Matlab Simulink environment along with the Dormand-Prince 45 variational numeric integrator (generally known as the DOPRI method, known in MATLAB as `ode45`), and both models use the parallelogram model to determine the hysteresis torque given the magnetic field input. Both models include the effect of gravity gradient torque, but ignore the other environmental torques (drag, solar pressure, magnetic residual, eddy current). Finally, both models use the closed magnetic circuit hysteresis parameters to form their parallelogram models. In both models, these assumptions and incorrect inputs add up to a simulation that converges very quickly (when CSSWE initial conditions are input, CUBESIM converges to the local magnetic field within a few orbits). Also, the CUBESIM output was not found to converge as lower time steps were used. Frustrations with CUBESIM inconsistencies was one of the motivations for work on a new PMAC simulation.

CUBESIM and SNAP use the following assumptions:

- Parallelogram hysteresis loop (no minor loops)
- Runge-Kutta45 numeric integrator

- No environmental torques included except gravity gradient
- Closed magnetic circuit hysteresis parameters

3.3.4 Park et al. (2010) & Lee et al. (2011) - RAX

The RAX CubeSat team recognized the faults with previous PMAC simulation tools [47]. They developed their own numeric integration model based on Equation 2.3. The RAX model includes magnetic and hysteresis torques, and rotation matrices are the chosen attitude coordinates. The RAX team uses the Lie Group Variational Integrator (LGVI) developed by Lee [45] at the University of Michigan. This energy-conserving numeric integrator is used in an attempt to limit system energy change due to the numeric integrator itself; the LGVI is described in Section 8.1.7 and tested versus other integrators in Section 8.3.

The RAX simulation uses the Flatley empirically-derived hysteresis loop model (described in Section 8.1.6.5). Unfortunately, the RAX simulation incorrectly uses closed magnetic circuit hysteresis parameters. Although students at the University of Michigan developed this PMAC simulator, only preliminary applications have been performed [58]. To date, the RAX team has not validated the performance of their simulation.

3.4 Hysteresis Measurement to Date

Only one other small satellite research group has made a priority of hysteresis measurement. Most teams incorrectly assume that the closed magnetic circuit hysteresis parameters of the hysteresis rod material may be used for simulation purposes, ([47],[61],[58]) but as is shown in Section 7.3, this is not the case. The very first PMAC mission, Transit 1B, measured the area of its hysteresis loops and used it as an input to its analytical model [25]. Of the modern missions, only the University of Rome (UNISAT-4, EduSat) have measured their rod hysteresis loops.

The UNISAT-4 team measured the hysteresis rod by placing a magnetometer close to one end of the rod and measuring the B-field resultant of magnetization changes within the rod [63].

However, their sense coil was not surrounding the hysteresis rod, they required the use of a scaling factor to translate from measured data away from the rod to some internal average B-field within the rod. Also, their measurement technique lacks the ability to perform system measurements. However, they did achieve measured hysteresis loops within an order of magnitude of those presented in Section 7.3.

Recently, the University of Rome built a new measurement system which uses a forcing coil and sense coil to determine the hysteresis parameters as a function of rod length (because the sense coil is much shorter than the hysteresis rod). The results from this measurement system yield hysteresis parameters ($H_c = 1.135$ A/m, $B_r = 0.0073$ Tesla, $B_s = 0.1315$ Tesla, Area=0.4263 J·m⁻³ for a 200 mm × 1 mm × 1 mm HyMu-80 rod [5]) which are vastly different from the closed magnetic circuit material hysteresis parameters (see Table 7.2). Other work measuring PMAC hysteresis rods is an analysis of various magnetic materials which could be used to fabricate hysteresis rods [24].

Chapter 4

The Colorado Student Space Weather Experiment


The Colorado Student Space Weather Experiment (CSSWE) is a 3U CubeSat [75] built at the University of Colorado Boulder as a joint project between the department of Aerospace Engineering Sciences (AES) and the Laboratory for Atmospheric and Space Physics (LASP). Involvement with CSSWE served as the motivation for this dissertation topic. CSSWE uses Passive Magnetic Attitude Control (PMAC) and we have full access to the satellite data. Thus, this CubeSat is an important feature of this dissertation.

CSSWE was selected for funding from the National Science Foundation in spring of 2010. Following a two year period of design, build, and test, CSSWE was delivered for PPOD integration in January 2012. Figure 4.1 shows an image of CSSWE and the Poly-Picosat Orbital Deployer (P-POD) launch device as captured during delivery. On September 13, 2012, CSSWE was inserted into a $478\text{km} \times 786\text{km}$, 64.7° inclination orbit as part of the NASA Educational Launch of Nanosatellites (ELaNa) VI launch [68]. CSSWE launched as a secondary payload aboard an Atlas V rocket operated by the United Launch Alliance (ULA) with a primary payload from the National Reconnaissance Office (NRO).

4.1 Science Mission Success

The science objectives of CSSWE are to investigate the relationship of the location, magnitude, and frequency of solar flares to the timing, duration, and energy spectrum of solar energetic particles reaching Earth and to determine the precipitation loss and the evolution of the energy spec-

Figure 4.1: The Colorado Student Space Weather Experiment (CSSWE, bottom left) CubeSat and its launch device, the Poly-Picosat Orbital Deployer (P-POD, upper right)



trum of radiation belt electrons [50]. To accomplish these objectives, CSSWE carries a miniaturized version of the Relativistic Electron and Proton Telescope (REPT), developed by LASP engineers for Van Allen Probes mission. CSSWE's lone science instrument, the Relativistic Electron and Proton Telescope *integrated little experiment* (REPTile, shown in Figure 4.2), uses 350V-biased silicon detectors to measure the directional differential flux of electrons from 0.5 to >3 MeV and protons from 10 to 40 MeV within a 52° field of view [66]. The data product of the mission is count rates for each particle within four energy bins, generated every six seconds.

The REPTile instrument measures charged particles which revolve around magnetic field lines as they travel. Thus, passive magnetic attitude control is beneficial because it results in higher particle count rates because the instrument field of view is oriented perpendicular to the local magnetic field direction. Also, the non-isotropic CSSWE antenna pattern favors alignment with the local magnetic field, which ensures an RF link can be established the majority of the time CSSWE is visible to the ground station in Boulder, CO.

The CSSWE mission has proven highly successful [49]. CSSWE proposed a mission lifetime of four months: one month of spacecraft checkout with full mission success defined by three months of science operations. As of Christmas Eve 2013, the satellite remains operational 466 days after launch with the science mission extended to over three times the 90 day full mission success duration. CSSWE is the quintessential proof that high-impact, journal-quality science can be accomplished with a low-cost CubeSat [48].

4.2 Coordinate System

CSSWE uses the body-fixed coordinate system shown in Figure 4.3. With the origin at the center of mass, the CSSWE X-, Y-, and Z-axes are aligned with the satellite major, intermediate, and minor inertia axes, respectively. The CSSWE principle inertia matrix about the center of mass is shown below; note that the satellite is close to symmetric about the X- and Y-axes.

Figure 4.2: The REPTile instrument collects particles in its 52° field of view through the collimator (C). Off-axis electrons are reflected by the collimator teeth away from the aperture. A beryllium window (F) absorbs electrons < 500 keV and protons < 10 MeV, preventing detector saturation. Particles travel through the detector stack (E), depositing energy on each detector as they travel. Binning logic allows the calculation of particle energy based on detector stack penetration. Aluminum (A) and Tungsten (B) shielding enclosing the detector stack limits noise due to particles not in the REPTile field of view. Three Tantalum alignment pins (D) provide rotational stability for the detector stack while providing the necessary shielding.

- A. Aluminum Outer Shielding
- B. Tungsten Inner Shielding
- C. Tantalum Collimator and Baffles
- D. Tantalum Alignment Pin
- E. Silicon Detectors
- F. Beryllium Window





Figure 4.3: The CSSWE coordinate frame is shown with the definitions for α and β error angles. The body X-axis ${}^B X$ is aligned with the REPTile aperture, the body Z-axis ${}^B Z$ is aligned with the deployed antenna, and the body Y-axis ${}^B Y$ is defined by the right-hand rule. The angle β exists between the body frame $+{}^B Z$ axis and magnetic flux density vector \mathbf{B} . The β angle is referenced throughout this document.



$$[I] = \begin{bmatrix} 0.0222 & 0 & 0 \\ 0 & 0.0218 & 0 \\ 0 & 0 & 0.0050 \end{bmatrix} \text{ kg} \cdot \text{m}^2.$$

4.3 Sensors and Telemetry

The satellite sensors are split into two types: housekeeping and attitude. The housekeeping sensors measuring temperature, current, and voltage throughout the spacecraft. The attitude sensors are used to measure the 3-axis local magnetic field vector and the partial sun position vector.

4.3.1 Housekeeping

CSSWE contains 38 housekeeping sensors spread throughout the satellite; Table 4.1 shows each sensor and its associated Analog to Digital Converter (ADC). Each ADC digitizes the analog readings with an 8-bit resolution. CSSWE queried each of these sensors once per minute and used these readings to compile ten-minute mean, maximum, and minimum values for each sensor. The sensors are detailed because some are used for magnetometer calibration (Section 4.3.2) while others are helpful for validating the attitude filter output (Section 6.3.2).

4.3.2 Attitude

The raw attitude measurements are not used by the satellite on-orbit; instead, they are transmitted to the ground for post-processing. Chapter 6 describes the ground-based satellite attitude determination using sensor output. This section focuses solely on the CSSWE attitude sensor hardware, output, and calibration.

CSSWE uses a three-axis magnetometer (Honeywell HMC5883L) and four photodiodes (Vishay TEMD6010FX01) to determine the full local magnetic flux density vector and the partial sun vector, respectively. The magnetometer is digitized with a 12-bit ADC, while the photodiodes share

Table 4.1: The 38 CSSWE housekeeping sensors are detailed below. Each sensor output is quantized by an 8-bit analog to digital converter before recording to memory. Attitude sensors are not included in this table, but ADC1 also digitizes the four photodiodes.

Analog to Digital Converter	Sensor	Quantity
ADC1	Solar Panel Temperature	4
ADC2	Solar Panel Voltage	4
	Solar Panel Current	4
ADC3	Battery Voltage	1
	Battery Temperature	1
	Battery Charge Current	1
	Battery Discharge Current	1
	5V Buck Voltage	1
	5V Buck Current	1
	3.3V Buck Voltage	1
	3.3V Buck Current	1
ADC4	REPTile Detector Voltage	4
	REPTile Detector Current	4
ADC5	REPTile Detector 1 Temperature	1
	REPTile Board Temperature	1
	REPTile Reference Voltage	3
	3.3V Buck Voltage at REPTile	1
Radio	Microcontroller Temperature	1
	Power Amp Temperature	1
	RSSI	1
C&DH	Microcontroller Temperature	1


the 8-bit ADC1 used by the solar panel temperature sensors (see Table 4.1). The instantaneous attitude sensors output values are recorded once every six seconds.

4.3.2.1 Magnetometer

The magnetometer was calibrated before flight using a time-invariant, attitude-independent method [31] which accounts for the offset, scale factor, and non-orthogonality biases. To account for temperature, current, and science instrument biases, a time-varying, attitude-independent calibration method was applied [72]. This calibration uses the first three days on-orbit (while the satellite is still covering the majority of the attitude sphere) to fit static calibration parameters based on time-varying telemetry. The Command and Data Handling (C&DH) board temperature telemetry is used as a proxy for the magnetometer temperature. Although a ten-minute average is the highest temporal resolution available for housekeeping data, it is found to be sufficient to correct for scale and offset magnetometer temperature errors. Post-launch analysis shows that the magnetometer temperature is the most significant source of time-varying error. The 10 minute average battery voltage telemetry is used as a proxy for system currents near the magnetometer which generate offset errors. The final magnetometer offset correction is based on the activation status of the REPTile instrument. The magnetometer is sensitive to REPTile currents because it is located on the REPTile electronics board, as shown in Figure 4.4. Also shown is the magnetometer proximity to the steel-tape communications antenna which deploys two hours after PPOD ejection. The antenna deployment changes the magnetic moment of the spacecraft (see Section 6.3.2.1) and has a significant effect on the magnetometer calibration.

Figure 4.5 shows the magnetometer B-field magnitude error over the first thirty days of on-orbit operations with various calibrations applied. The error is calculated as the difference between the measured B-field magnitude and the International Geomagnetic Reference Field (IGRF) magnetic model output magnitude (see Section 8.1.5.1). Relative to the daily average magnitude of the IGRF, the raw data has a daily average error of well over 400% and is thus unusable. The ground-based calibration reduces the daily average error to about 15%, or about 9° of attitude er-

Figure 4.4: Solid model of the magnetometer (green) position relative to the bar magnet (red) and the deployed antenna (copper). The magnetometer is sensitive to REPTile currents due to its placement on the REPTile electronics board.



ror. The on-orbit, time-varying calibration is the only dataset with a mean error below its standard deviation. The maximum net effect of the measurement error average and standard deviation is about a 2% error relative to the daily average IGRF magnitude, or about a 1.1° attitude error.

4.3.2.2 Photodiodes

CSSWE has photodiodes on each of its 3U faces ($+^BX$, $-^BX$, $+^BY$, $-^BY$); these allow for a partial sun vector measurement. Each photodiodes measures the sun direction by assuming 1) an output when the sun is perpendicular and 2) that the output decreases as a cosine with the sun direction. However, the second assumption is usually invalid at high incidence angles due to physical limitations (refraction, manufacturing imperfections, etc.) and secondary light sources. As a result, CSSWE does not use photodiode measurements beyond a 70° field of view. Figure 4.6 shows the relationship between the number of illuminated photodiodes and the sun direction vector in the body frame; this figure has slight inaccuracies due to the size of each grid element (which are large for clarity purposes).

The photodiodes were calibrated using an attitude-dependent, batch-based approach using on-orbit data which takes into account the effects of albedo. This calibration is largely based on the work of Springmann [70]. However, partial sun vector measurement based on the four CSSWE photodiodes is not sufficient to estimate the photodiode calibration parameters when directly included as filter states. Instead, a novel batch-based filtering approach is used to calibrate each photodiode (for more information, see Section 6.2.2). This calibration corrects for the scale factor and misalignment angles of each photodiode. The misalignment angles account for manufacturing and mounting defects; these parameters do not change with time.

The scale factor is defined as the output from a photodiode when it is perpendicular to the sun alone (no albedo); this parameter changes over time. Figure 4.7 shows that the photodiodes experience significant degradation over the first month on orbit; this degradation is believed to be due to UV light darkening the plastic covering of each photodiode. The figure also shows that the scale factor of each photodiode is recalculated every six hours to account for degradation as

Figure 4.5: The difference in magnetic flux density magnitude as measured by the CSSWE magnetometer vs. predicted by the International Geomagnetic Reference Field (IGRF, see Section 8.1.5.1) model. Both the daily mean error (filled squares) and the daily standard deviation of the error (open circles) are shown. The black lines are the maximum and minimum modeled B-field magnitude at the CSSWE orbit each day. Red is the error of the raw data from the magnetometer, blue is the error after the ground-based (static calibration parameters) calibration is applied, and green is the error after the on-orbit (dynamic calibration parameters) calibration is applied.

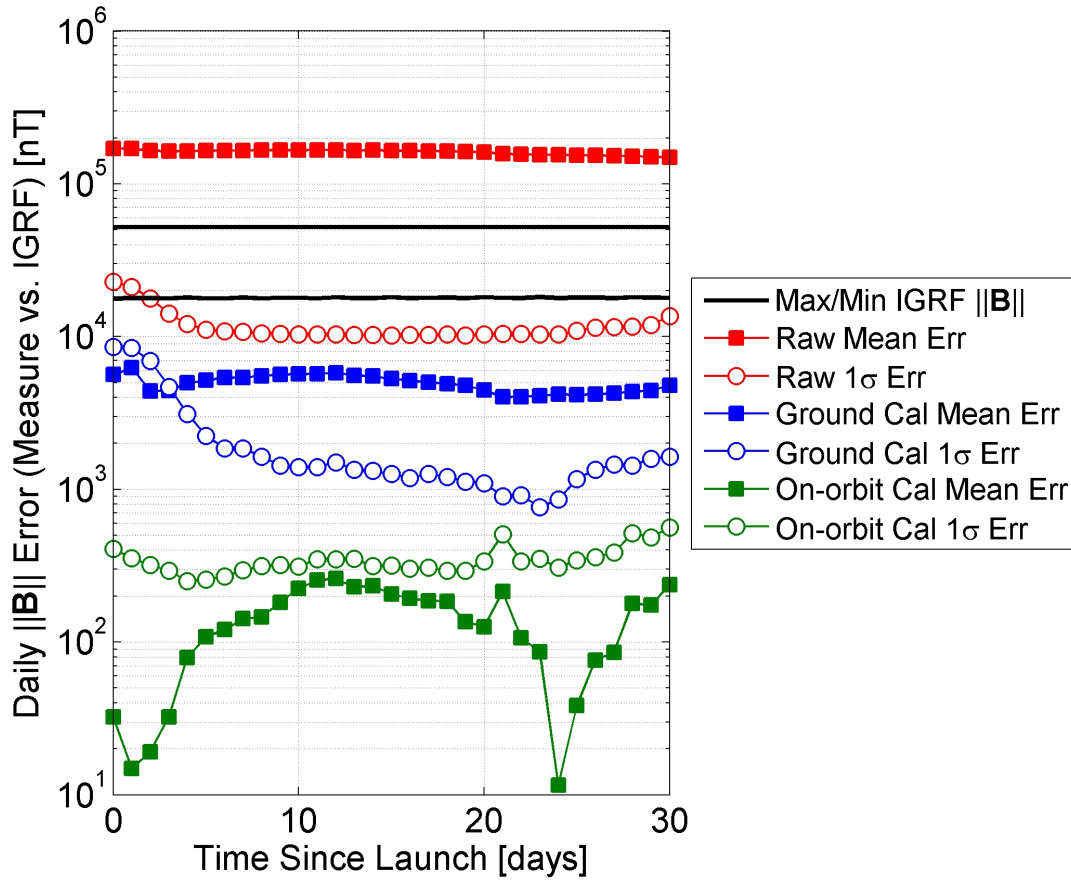
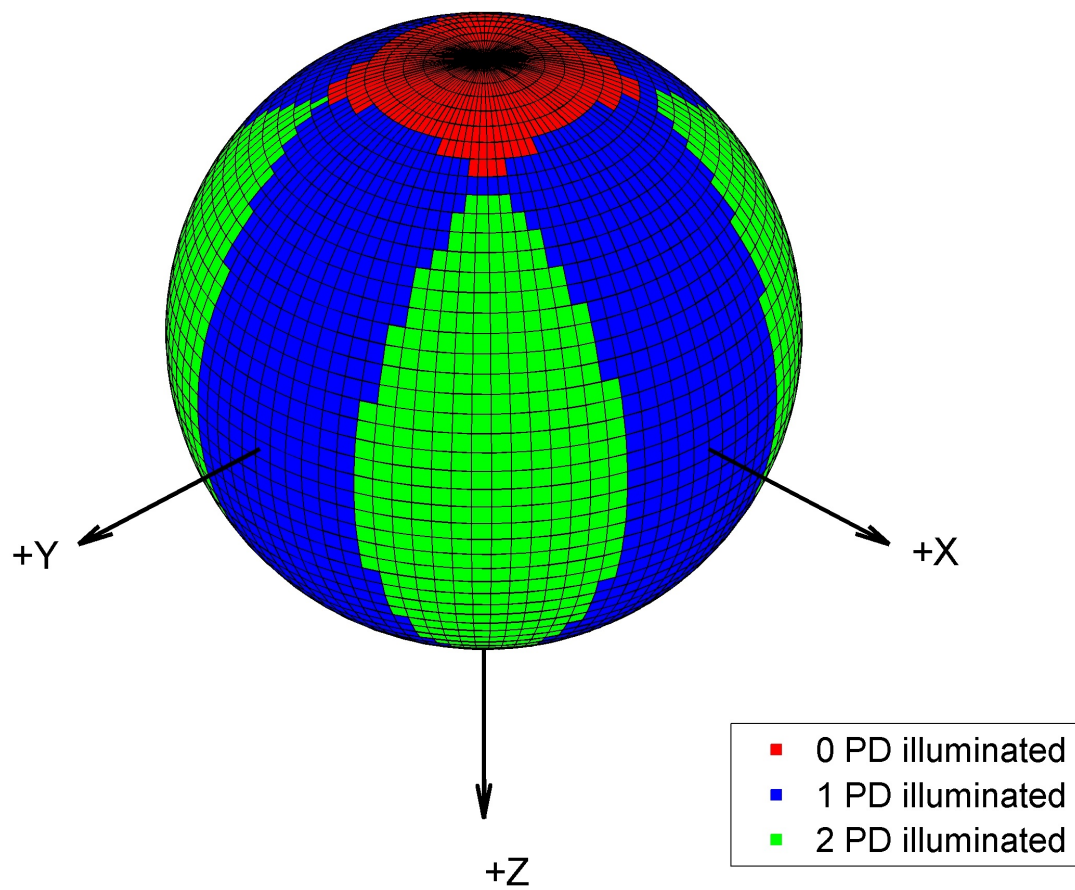


Figure 4.6: The number of illuminated photodiodes given a sun direction vector for the CSSWE CubeSat. This distribution assumes photodiodes with a 70° field of view aligned with the $+X$, $-X$, $+Y$, and $-Y$ axes.



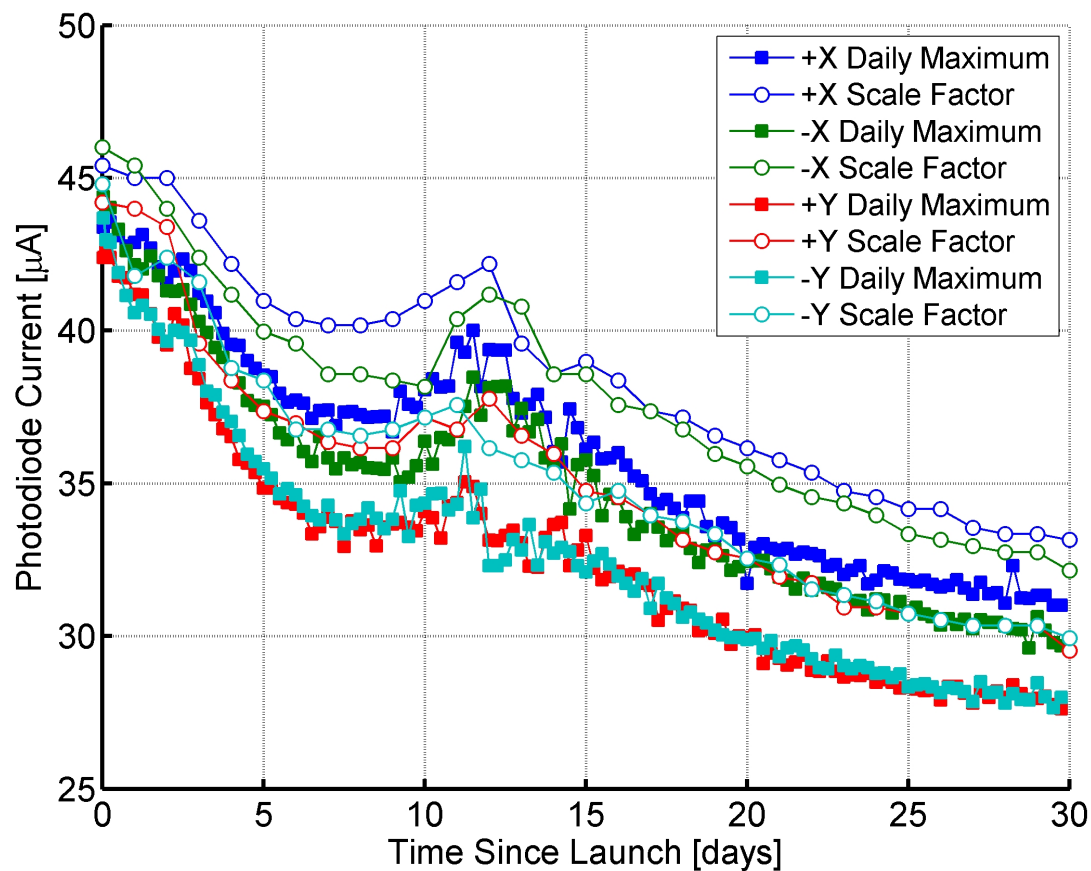
it occurs. After calibration, the photodiodes have a nominal 1-sigma standard deviation of $1.6\mu\text{A}$ due primarily to the uncertainty of the albedo model. In the early mission, the maximum sun-only current is about $45\mu\text{A}$ (equivalent to about 2.0° of attitude error) but by late mission this maximum sun-only current decreases to about $32\mu\text{A}$ (equivalent to about 2.8° of attitude error).

4.3.2.3 Inertial Models and Uncertainties

The measurements cannot be used to determine spacecraft attitude without some estimate of their inertial values. CSSWE uses a Two Line Element (TLE) set with the SPG4 propagator for position information (see Section 8.1.4). The satellite position is used to generate both the inertial sun vector (accurate to 0.01°) and the inertial magnetic field vector (accurate to about 1° as shown below); both models are explained in more detail in Section 8.1.5.

Figure 4.8 shows the results of an investigation of the Two Line Element (TLE) set uncertainty over the first 30 days on orbit. The TLE is propagated to determine the satellite position and velocity at a given time (see Section 8.1.4); thus an inaccurate TLE will lead to an inaccurate position estimate. This investigation was performed by simulating CSSWE dawn crossings and comparing each crossing with on-orbit photodiode data. Red lines have been added to the figure at ± 6 seconds to show the expected variation in dawn crossings as CSSWE records photodiode data once every six seconds. The figure shows that the CSSWE TLE-based position degrades over the first month on orbit. During this time period, CSSWE drifts apart from the other 10 CubeSats from the same launch; this may have caused issues for the tracking agency which supplies the TLEs. Regardless, the data shows that the simulated dawn crossings have errors as high as 18s compared to the measurements. Simulations of the first month's orbit show that an 18s in-track position error results in IGRF model errors with a 1σ standard deviation of 284nT. However, this is a worst case that is seldom experienced; a more reasonable assumption is a 10s in-track position error, resulting in IGRF model errors of 158nT. The IGRF model can also differ from truth in the presence of geomagnetic storms. Moderate storms can cause variations up to 100nT at low latitudes [51]. Recent analysis has shown high latitudes can experience variations of up to 1000nT

Figure 4.7: The maximum recorded output current for each photodiode is shown with the open circles. The filled squares show the scale factor for each photodiode, which was fit to the on-orbit data every six hours. The photodiode scale factor is the current that would be registered by the photodiode if it were perpendicular to the sun alone (no albedo).



during moderate storms [42]. Because the attitude determination is performed via post-processing, the storm-based magnetic field variation can be avoided by selectively processing datasets during which geomagnetic activity is low. For comparison, the lowest magnetic flux density magnitude experienced within the CSSWE orbit is about $18\mu\text{T}$, so 180nT is equivalent to 1% error or about 0.57° angular error.

4.4 Latch-up Anomaly

Although CSSWE met all of its goals for full mission success, it did have setbacks. The most severe anomaly (with implications for attitude determination) began on October 14, 2012, at 23:28:45 UTC. Based on received telemetry, it is believed that ADC4 (see Table 4.1) experienced a latch-up anomaly, likely due to a high-energy particle impact. The latch-up caused an undesired low impedance path to ground through ADC4. This short circuit eventually brought the battery voltage below a battery protection circuitry threshold, triggering a system reset which cleared the latch-up two hours after the anomaly began. As a result of the anomaly, ADC4 was destroyed and both ADC1 and ADC2 were damaged. Unfortunately, ADC1 digitizes the photodiode output, which was degraded by the anomaly. Figure 4.9 shows the raw bit-level output from ADC1 before and after the latch-up anomaly occurs in mid October. The +X, -X, and +Y photodiodes lose certain bits after the anomaly. As a result, there is increased quantization of the photodiode output, greatly reducing its usefulness in attitude estimation. For an 8-bit ADC (maximum output 255), losing bit 5 (+X, -X) or 7 (+Y) results in percent errors of 12.5% or 50.1%, respectively. Because the attitude determination relies upon accurate photodiode output, the CSSWE three-axis attitude is limited to the first month of on-orbit operations.

Figure 4.8: The difference between the dawn crossing predicted by the current Two Line Element (TLE) set and the photodiode-based dawn crossing observation. The green dots are the closest dawn crossing to an updated TLE; the blue dots are propagated forward or backward in time using data from the most recent TLE. Dotted red lines are visible at ± 6 sec to indicate the photodiode measurement period; differences beyond these boundaries are likely due to errors in TLE-based position. The increasing error is thought to be due to the ten CubeSats launched with CSSWE dispersing over time.

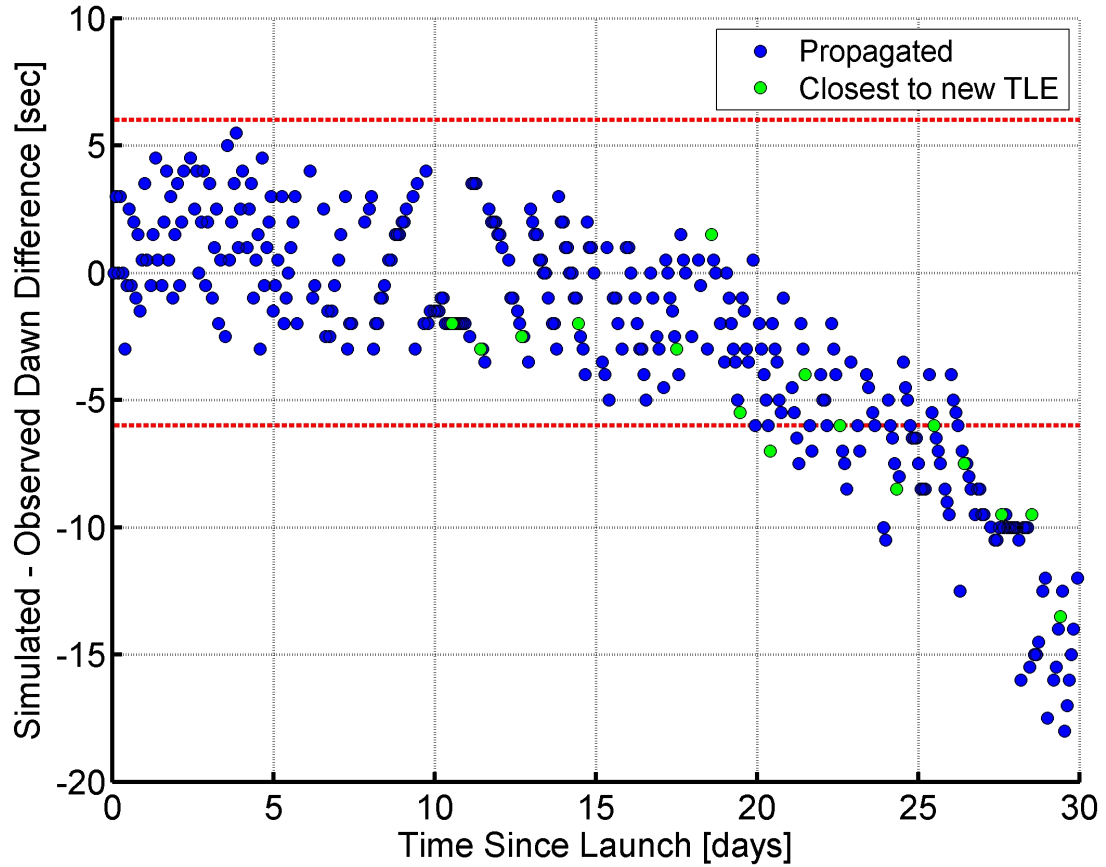
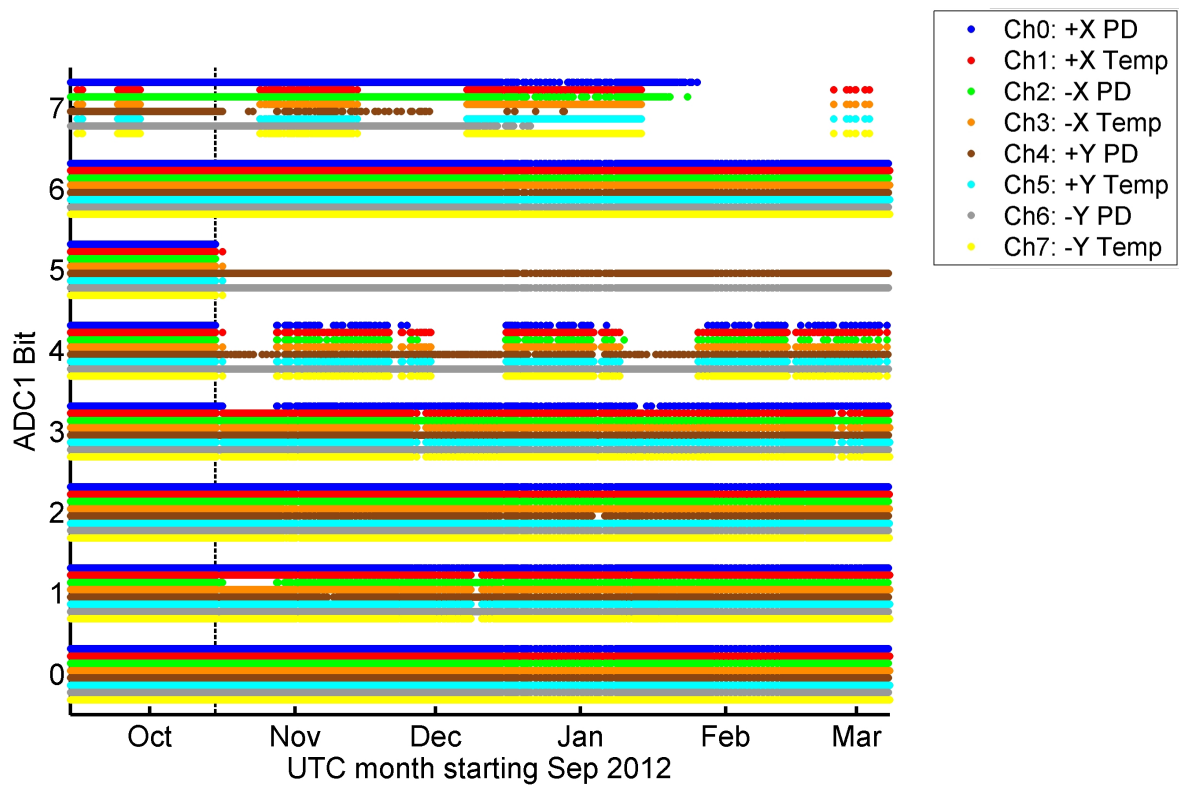


Figure 4.9: The output of ADC1 before and after the latch-up anomaly. Select bits are no longer active after the anomaly, causing increased quantization of the output data. The +X and -X photodiodes lose bit 5 while the +Y photodiode bit 7 becomes intermittent.



Chapter 5

Control System Design

We anticipate that satellite developers will refer to this dissertation when designing future PMAC systems. One example of good PMAC system design is the Colorado Student Space Weather Experiment (CSSWE) CubeSat (discussed in Chapter 4). This chapter details PMAC system design practices using CSSWE as a specific example.

The attitude control system of CSSWE has two performance requirements:

- (1) The attitude control system shall have a settling time of less than 7 days.
- (2) Once settled, the attitude shall stay aligned within 15° of the local magnetic field.

5.1 Maximum Expected Environmental Torques

A successful attitude system design begins with an analysis of maximum expected spacecraft environmental torques. Table 5.1 shows the expected environmental torques in the CSSWE environment. Methods to calculate these torques are explained in Section 8.1.6. The maximum expected non-magnetic torque total is used to determine the minimum acceptable bar magnet magnetic moment.

Table 5.1: Worst-case environmental torque magnitudes for the CSSWE 3U CubeSat. This analysis assumes a 3U CubeSat in a 480km×790km, 65° orbit and moderate solar input.

Torque	Maximum Value [N·m]
Magnetic Residual $ \mathbf{L}_R $	$4.3E - 7$
Aerodynamic $ \mathbf{L}_D $	$1.8E - 8$
Gravity Gradient $ \mathbf{L}_G $	$3.2E - 8$
Solar Pressure $ \mathbf{L}_{SP} $	$2.5E - 9$
Sum $ \mathbf{L} _{\text{sum}}$ (excluding $ \mathbf{L}_R $)	$5.3E - 8$

5.2 Bar Magnet Design

In the CSSWE orbit, $||\mathbf{B}||$ varies from 18 to 52 μTesla . We present a modified version of Santoni and Zelli’s [63] minimum recommended bar magnet strength:

$$m_{\min} = 15 \left(\frac{||\mathbf{L}_{\text{sum}}||}{||\mathbf{B}||_{\min} \cdot \sin(\beta_{\max})} \right) \quad (5.1)$$

where $||\mathbf{L}||_{\text{sum}}$ is the sum of the independent, non-magnetic environmental torque magnitudes, $||\mathbf{B}||_{\min}$ is the minimum magnetic flux density magnitude experienced by the satellite, and β_{\max} is the desired pointing accuracy. Although the required alignment with the magnetic field is 15°, the system is designed using $\beta_{\max} = 10^\circ$ to ensure there is adequate margin in the PMAC system design. Santoni and Zelli’s [63] version of Equation 5.1 defines $||\mathbf{L}||_{\text{sum}}$ as the sum of all independent environmental torques. However, $||\mathbf{L}||_{\text{sum}}$ is better defined as the sum of the non-magnetic torques. Instead of heavily weighting Equation 5.1 due to the maximum magnetic residual torque vs. the minimum bar magnet torque, the respective magnetic moments may be compared directly. Because any magnetic torque is given by Equation 2.6, as long as $m_{\text{bar}} \geq 15m_{\text{res}}$, the magnet will easily overpower the magnetic residual. This revised definition of $||\mathbf{L}||_{\text{sum}}$ produces a less extreme bar magnet moment, thereby reducing the necessary hysteresis damping material within the volume-limited CubeSat and lessening the initial magnetic potential energy of the satellite.

The bar magnet moment is also directly related to the initial energy which may be introduced to the system when a small satellite is released from its launcher. The rotational energy of a PMAC

satellite may be split into the kinetic and potential energy as shown below:

$$T_K = \frac{1}{2} \boldsymbol{\omega}^T [I] \boldsymbol{\omega} \quad (5.2)$$

$$T_P = -\mathbf{m} \cdot \mathbf{B} \quad (5.3)$$

where \mathbf{m} is the total satellite magnetic moment, $\boldsymbol{\omega}$ is the angular velocity vector, and $[I]$ is the satellite moment of inertia matrix. Usually, small satellites are secondary payloads and thus cannot control the initial angle versus the magnetic field β_0 at satellite deployment. Consider a satellite which initially possesses zero kinetic energy and starts with $\beta_0 = 180^\circ$. In this case, the initial rotational energy is directly related to the magnetic moment magnitude $\|\mathbf{m}\|$. Thus, the bar magnet moment must balance between dominating the disturbance torques and limiting the possible system energy which could be introduced at deployment.

For the considered conditions, Equation 5.1 yields an ideal bar magnet strength $m_{\min} = 0.25 \text{ A}\cdot\text{m}^2$, which is $25\times$ the expected residual magnetic moment. A bar magnet was ordered to meet this level, but when the bar magnet was measured (see Section 7.2), it was found to be $m_{\text{bar}} = 0.80 \pm 0.017 \text{ A}\cdot\text{m}^2$. Due to limited time before CubeSat delivery and because the measured $m_{\text{bar}} > m_{\min} > 15m_{\text{res}}$, the bar magnet magnetic moment was deemed acceptable and is used in the CSSWE PMAC system.

5.3 Hysteresis Rod Design

Once a bar magnet dipole moment has been chosen, the hysteresis rod dimensions and quantity should be determined. Usually in a PMAC system, hysteresis rods are mounted in pairs orthogonal to the bar magnet to reduce the offset of the applied field due to the bar magnet (see Section 2.2.3). Thus, the CSSWE bar magnet is aligned with the minor inertia axis (${}^B Z$) of the CubeSat and the hysteresis rods are mounted on orthogonal axes (${}^B X$ and ${}^B Y$, see Figure 4.3). The design below assumes that perpendicular rods never interact and that rod sets do not interact

when separated by more than 30% of their length [57].


As explained in Section 2.2.6, the length to diameter ratio of a hysteresis rod has a large effect on its performance. Higher length to diameter ratios result in higher permeability hysteresis rods. The maximum length of the hysteresis rod is limited by the dimensions of the spacecraft. For CSSWE, the interior spacecraft dimensions perpendicular to the bar magnet limit the rods to 9.5 cm in length. Typical values of the length to diameter ratios for hysteresis rods are on the order of 100 [57]. Thus, CSSWE used hysteresis rods of length 95 mm and diameter 1 mm.

Next comes the question of how many pairs of hysteresis rods to include in the system. First, there is a volume limitation; the planes of orthogonal rod sets must be separated by at least 30% of the length of one rod in order to ensure that the magnetization of one hysteresis rod set does not affect the other [57]. Also, it may be advantageous to separate the hysteresis rods from any spacecraft magnetometers. It is possible to calibrate a magnetometer to remove hysteresis effects, but most calibration methods assume a linear hysteresis curve [28] [72]; the magnetometer performance may degrade if the hysteresis rods are too close.

Regardless of the physical limitations of including hysteresis in the system, the optimal amount of hysteresis material is a question. If not enough hysteresis material is included, the system will take too long to converge, and PMAC design rule #1 will be violated. As dampening material is increased, the offset from the local magnetic field increases as well, because the total magnetic moment vector is what will align with the local magnetic field. Figure 5.1 defines the error angle γ between the total magnetic moment vector and the magnetic moment vector of the bar magnet. If the sum of the maximum hysteresis magnetization and the magnetic residual moment represent a significant fraction of the total magnetic moment vector, the error angle γ may cause the system to violate PMAC design rule #2.

Of course, the available locations within the satellite to affix hysteresis rods also determines the allowable number of rods. Figure 5.2 shows a solid model of the final design of the CSSWE PMAC system (the bar magnet and hysteresis rods are highlighted in red). CSSWE uses three hysteresis rods on each of the $^B X$ and $^B Y$ -axes. The rods are separated by a perpendicular distance


Figure 5.1: There exists an error angle γ due to magnetization on the satellite not parallel to the bar magnet.



of at least 3.25cm (34% of their length) and the bar magnet to ensure the hysteresis rods have minimal magnetic offset. The magnetometer (not shown) is located on an electronics board near the bottom of the satellite. The magnetometer was chosen to be separated from the hysteresis rods rather than the bar magnet because a constant offset from the bar magnet can be negated via calibration, but the non-linear variation of magnetic fields due to the hysteresis rods are more difficult to remove.

Accurate magnetic moment and hysteresis parameters input to a realistic PMAC simulation can be used to determine the volume of spacecraft hysteresis material needed to meet specific mission requirements. With the necessary hysteresis volume defined, multiple rods may be installed using the design rules outlined above. Thus, an accurate attitude dynamics simulation is a key part of PMAC system design. Chapter 7 ensures accurate input to this simulation through hysteresis rod measurement; it also tests the hysteresis design rules outlined above. Chapter 8 defines and tests a PMAC attitude dynamics simulation. This simulation is verified for future design use by comparing its output with the on-orbit attitude data from the CSSWE CubeSat. As such, Chapter 6 develops a filter for attitude determination of a PMAC satellite and applies it to the empirical CSSWE data.

Figure 5.2: A solid model highlighting the position of PMAC components. The hysteresis rods (top) have a large separation from the bar magnet (bottom left) in order to prevent magnetic offsets to the rod hysteresis loops. Each hysteresis rod position is labeled; the rod sets are separated by a minimum perpendicular distance of 3.25 cm.



Chapter 6

Attitude Determination

When developing a simulation, empirical data from the modeled system is useful for comparison. However, in the case of an attitude simulation, the raw observations from the satellite must be converted to an estimate of the satellite attitude before any simulation-to-empirical-data comparisons may be performed. This chapter develops an attitude determination filter which may be applied to any PMAC satellite. After the filter is defined, it is applied to on-orbit measurements to determine the attitude of a PMAC satellite. This provides the empirical data needed for comparison with the attitude simulation (developed in Chapter 8).

The filter is tuned using input from the CSSWE CubeSat (see Chapter 4) but is applicable to any satellite with a PMAC system when rate gyro data are not available. CSSWE is designed to perform attitude determination via post-processing using raw measurements transmitted to its ground station. However, given proper ground-based measurements before launch, the attitude determination developed in this chapter could occur in real-time on orbit.

6.1 Filter Design

The Extended Kalman Filter (EKF) is an established method of attitude determination [46] [16]. In general, an EKF is useful for estimating the state and covariance of a non-linear, discrete-time process. What follows is an overview of the general EKF. First, assume the state propagation for a given process is governed by the non-linear stochastic difference equation

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}, \mathbf{w}_{k-1}) \quad (6.1)$$

with measurements

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k, \mathbf{v}_k) \quad (6.2)$$

where \mathbf{x} is the true state vector, \mathbf{u} is the true control input, \mathbf{w} is the true process noise, \mathbf{v} is the true measurement noise, and k is the step number. Note that each measurement is an inseparable combination of the current state and the measurement noise. The true process noise and measurement noise cannot be observed directly for each time step; instead the Kalman filter assumes all noise is Gaussian, independent, and zero-mean. Because Equation 6.1 represents a non-linear process it is difficult to propagate directly. The EKF state propagation procedure is shown below. Note that for the EKF matrices, the normal convention within this dissertation of matrices being bracketed is ignored.

$$\hat{\mathbf{x}}_k^- = \mathbf{f}(\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) \quad (6.3)$$

$$P_k^- = F_k P_{k-1}^+ F_k^T + G_k Q_{k-1} G_k^T \quad (6.4)$$

$$K_k = P_k^- H_k^T (H_k P_k^- H_k + J_k R_k J_k^T)^{-1} \quad (6.5)$$

$$\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- + K_k (\mathbf{y}_k - \mathbf{h}(\hat{\mathbf{x}}_k^-, 0)) \quad (6.6)$$

$$P_k^+ = (I - K_k H_k) P_k^- \quad (6.7)$$

where a hatted variable ($\hat{\cdot}$) represents an estimate, a superscript minus ($-$) or plus ($+$) represents the *a priori* or *a posteriori* estimate before or after the measurement update, P is the state covariance

matrix, Q is the process noise covariance matrix, R is the measurement covariance matrix (not to be confused with the rotation matrix $[R]$), F is the state transition matrix, G is the process noise gain matrix, H is the measurement sensitivity matrix, and J is the measurement noise gain matrix.

The Jacobian matrices are defined as follows:

$$F_k \equiv \frac{\partial \mathbf{f}}{\partial \mathbf{x}} (\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) \quad (6.8)$$

$$G_k \equiv \frac{\partial \mathbf{f}}{\partial \mathbf{w}} (\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) \quad (6.9)$$

$$H_k \equiv \frac{\partial \mathbf{h}}{\partial \mathbf{x}} (\hat{\mathbf{x}}_k^-, 0) \quad (6.10)$$

$$J_k \equiv \frac{\partial \mathbf{h}}{\partial \mathbf{v}} (\hat{\mathbf{x}}_k^-, 0) \quad (6.11)$$

The EKF linearizes around the previous state estimate $\hat{\mathbf{x}}_{k-1}^+$ to generate an *a priori* estimate of the state vector $\hat{\mathbf{x}}_k^-$ and covariance matrix P_k^- . This *a priori* covariance matrix P_k^- is used to determine the gain K_k . Applying this gain results in the best fit state estimate $\hat{\mathbf{x}}_k^+$ for the current time step. The gain is also used to calculate the current best estimate of the *a posteriori* covariance matrix P_k^+ .

In most attitude-determination applications, the states modeled by the EKF are the attitude parameters and angular rates. However, the attitude parameterization choice is not trivial. All attitude parameterizations of three dimensions contain a singularity in their kinematic differential equation, while three-dimensional attitude coordinates expressed in four or more dimensions have dependent parameters, which can result in a singular covariance matrix after an EKF update because numerical errors can cause one or more of the parameter constraints to be violated. The Multiplicative Extended Kalman Filter (MEKF) resolves this issue by using a combination of quaternions and a three-dimensional attitude parameterization; the former is used as a non-singular reference and the latter is used to tabulate the attitude error at each time step. The MEKF was

originally developed assuming the use of a rate gyro [54]. Alternatively, the rate gyro can be omitted by modeling the dynamics of the spacecraft. For the purposes of the MEKF, the PMAC attitude dynamics can be modeled by including only the bar magnet torque (which dominates all other external torques) and treating all other external torques as Gaussian process noise [11]. A disadvantage of this method is the inclusion of the mass moment of inertia matrix and the bar magnet strength in the dynamics model; these must be accurately measured before launch or fitted to on-orbit data to achieve satisfactory results.

We use the scaled Gibbs vector as the attitude error parameterization and follow the convention of Markley [54] to define the scaled Gibbs vector:

$$\mathbf{a}_g \equiv 2 \frac{\mathbf{q}}{q_0} \quad (6.12)$$

where \mathbf{q} is the vector part of the quaternion and q_0 is the scalar part. The factor of two is included because it makes \mathbf{a}_g approximately equal to the yaw, pitch, and roll Euler angles for any rotation set, given that \mathbf{a}_g represents a small rotation. Thus, when the scaled Gibbs vector is used as the attitude error parameterization, the uncertainty estimate from the MEKF is directly applicable to satellite attitude uncertainty in yaw, pitch, and roll. Thus, the six-dimensional state vector for the PMAC MEKF is $\mathbf{x} = [\mathbf{a}_g \ \boldsymbol{\omega}]^T$. Although the reference quaternion is not technically a state, it is updated at the end of each filtering step as follows:

$$\hat{q}_k^+ = \delta q(\mathbf{a}_{g,k}) \otimes \hat{q}_k^- \quad (6.13)$$

where $\delta q(\mathbf{a}_g)$ is the error quaternion given by

$$\delta q(\mathbf{a}_g) \approx \begin{bmatrix} \mathbf{a}_g/2 \\ 1 - \mathbf{a}_g^2/8 \end{bmatrix} \quad (6.14)$$

where the scalar-last quaternion convention is used. In practice, the *a posteriori* quaternion should be renormalized after each use of Equation 6.13 to combat numerical error buildup. Combining Equation 6.14 with the definition of a rotation matrix in terms of quaternions and assuming \mathbf{a}_g is a

small angle and ignoring higher-order terms yields the error rotation matrix. This rotation matrix definition is useful in determining the MEKF Jacobian matrices.

$$[R(\delta q(\mathbf{a}_g))] \approx [I_{3 \times 3}] - [\mathbf{a} \times] - \frac{1}{2}(a_g^2 [I_{3 \times 3}] - \mathbf{a}_g \mathbf{a}_g^T) \quad (6.15)$$

In order to determine the Jacobian matrices that define the behavior of the MEKF, the state dynamics must be determined. By combining Equation 6.13 with the kinematic differential equation for quaternions and the time-derivative of Equation 6.12, the time-derivative of the scaled error Gibbs vector can be obtained as shown below. The time-derivative of the angular velocity vector is given by Equation 2.1 as torqued by the bar magnet alone (Equation 2.6). The dynamics model used by the PMAC MEKF is thus:

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{a}}_g \\ \dot{\boldsymbol{\omega}} \end{bmatrix} = \begin{bmatrix} ([I_{3 \times 3}] + \frac{1}{4} \mathbf{a}_g \mathbf{a}_g^T)(\boldsymbol{\omega} - \hat{\boldsymbol{\omega}}) - \frac{1}{2}(\boldsymbol{\omega} + \hat{\boldsymbol{\omega}}) \times \mathbf{a}_g \\ [I]^{-1} (\mathbf{m} \times ([R(\delta q(\mathbf{a}_g))][R(q_{\text{ref}})]^T \mathbf{B} + \boldsymbol{\eta}_2) - \boldsymbol{\omega} \times [I]\boldsymbol{\omega} + \boldsymbol{\eta}_1) \end{bmatrix} \quad (6.16)$$

where $\hat{\boldsymbol{\omega}}$ is the expected value of the angular velocity (as differentiated from the true angular velocity $\boldsymbol{\omega}$); these are analogous to the difference between \hat{q}_k^- and \hat{q}_k^+ , respectively. Note that we follow the method of Burton [11] and define the process noise vectors $\boldsymbol{\eta}_1$ and $\boldsymbol{\eta}_2$ as the inertial magnetic model error (including errors due to satellite position) and unmodeled external torques, respectively. With the state dynamics fully defined, the first two MEKF Jacobian matrices may be evaluated as follows:

$$F_k \equiv \frac{\partial \mathbf{f}}{\partial \mathbf{x}} (\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) = \begin{bmatrix} -[\boldsymbol{\omega}_{k-1} \times] & [I_{3 \times 3}] \\ [I]^{-1} [\mathbf{m} \times] [R(\hat{q}_{k-1}^+)]^T \mathbf{B}_{k-1} & [I]^{-1} (-[\boldsymbol{\omega}_{k-1} \times] [I] + [I \boldsymbol{\omega}_{k-1} \times]) \end{bmatrix} \quad (6.17)$$

$$G_k \equiv \frac{\partial \mathbf{f}}{\partial \mathbf{w}} (\hat{\mathbf{x}}_{k-1}^+, \mathbf{u}_{k-1}, 0) = \begin{bmatrix} [0_{3 \times 3}] & [0_{3 \times 3}] \\ [I]^{-1} & [I]^{-1} [\mathbf{m} \times] \end{bmatrix} \quad (6.18)$$

Before evaluating Equations 6.10 and 6.11, first note that a body frame observation vector can be expressed using inertial data as $\mathbf{h}(\mathbf{x}_k, \mathbf{v}_k) = [R(\delta q(\mathbf{a}_g))][R(\hat{q}_k^-)]^T \mathbf{b}_k + \mathbf{v}_k$ where ${}^I \mathbf{b}_k$ is the observation

vector in the inertial frame at this step. With this full measurement vector model in hand, the final two Jacobians are evaluated as:

$$H_k \equiv \frac{\partial \mathbf{h}}{\partial \mathbf{x}} (\hat{\mathbf{x}}_k^-, 0) = \begin{bmatrix} [R(\hat{q}_k^-)^T \mathbf{b}_k \times] & [0_{3 \times n}] \end{bmatrix} \quad (6.19)$$

$$J_k \equiv \frac{\partial \mathbf{h}}{\partial \mathbf{v}} (\hat{\mathbf{x}}_k^-, 0) = [I_{6 \times n}] \quad (6.20)$$

where Equation 6.15 has been substituted prior to integration. The number of measurements at a given step determines the size of the Jacobians H_k and J_k , as well as the size of the measurement covariance matrix R . When using the filter, the observation vector is approximated assuming zero noise and using the expected value of the state to rotate an inertial, model-based observation at each step:

$$\mathbf{h}(\hat{\mathbf{x}}_k, 0) = [R(\hat{q}_k^-)^T] \mathbf{b}_k \quad (6.21)$$

To be clear, the *a priori* values \hat{q}_k^- and $\hat{\omega}_k^-$ are obtained by numerically integrating the state dynamics given the *a posteriori* values from the previous step:

$$\begin{bmatrix} \hat{q}_k^- \\ \hat{\omega}_k^- \end{bmatrix} = \begin{bmatrix} \hat{q}_{k-1}^+ \\ \hat{\omega}_{k-1}^+ \end{bmatrix} + \int_{t_{k-1}}^{t_k} \begin{bmatrix} \dot{q} \\ \dot{\omega} \end{bmatrix} dt. \quad (6.22)$$

Note that the scaled error Gibbs vector \mathbf{a}_g is not iterated; \mathbf{a}_g is defined as the error between the *a priori* and *a posteriori* estimates. The expected value of the scaled error Gibbs vector expected value is always zero, but it does have a non-zero *a posteriori* estimate set by Equation 6.6. This is why the quaternion *a posteriori* update (Equation 6.13) is always the final action of each filtering time step.

The PMAC MEKF derivation assumes that the error Gibbs vector represents a small angle rotation (usually defined as $<5^\circ$). As a result, trust in the MEKF-output states and uncertainties should be tentative if components of the error Gibbs vector are consistently $>5^\circ$. The PMAC

MEKF filter is now fully defined. However, the filter must be tuned before its outputs may be trusted.

6.2 Filter Tuning

Although the measurement noise R can be usually be determined via sensor testing, the process noise Q is more difficult to set. Tuning is the process of varying Q and R as needed until the filter performance is acceptable. After tuning, the uncertainty given by the filter covariance should match the error of the state estimates. There are two tuning methods: simulating noisy input to the filter and empirical tuning using residuals; the CSSWE mission used both methods.

6.2.1 Simulation-based Filter Tuning

The simulation uses the SGP4 orbital propagator [79], the CSSWE TLE, an initial attitude, satellite properties, and environmental torques models to numerically determine the “true” satellite motion. The simulation shares many similarities with the one developed in Chapter 8. External torques which the MEKF does not model (gravity-gradient, drag, and solar pressure) are used in this “truth” simulation to determine if estimating these disturbance torques as Gaussian noise is acceptable. Gaussian noise of 15nT is added to the IGRF model before calculating the magnetic torque at each time step; this value models good orbital position knowledge with a quiet sun. Noise is also added to simulated body-frame measurements; the 1σ standard deviation of the photodiodes is set to $1.6\mu\text{A}$ which reflects the $50\text{ W}\cdot\text{m}^{-2}$ uncertainty in the albedo model used to correct the photodiode output. The 1σ magnetometer standard deviation is set to 400nT as it is based on the encountered \mathbf{B} -field magnitude error (see Figure 4.5). Noisy body-frame measurements generated by the simulation are then filtered by the MEKF. This filtering is repeated with a sensible range of Q (based on expected levels of the external torques which are not modeled and the input IGRF noise) until the uncertainty bounds correctly represent the errors (difference between “truth” and the MEKF-output) in filtered attitude and angular velocity.

Figures 6.1 and 6.2 show the properly-tuned MEKF output given simulated measurements

Figure 6.1: Angular error as output by the MEKF after tuning using a simulation of the attitude dynamics. The error is the difference between the simulated attitude and the MEKF fitted attitude generated from measurements corrupted with Gaussian noise. The red line shows the 3σ angular uncertainty, which should bound 99.73% of the angular error.





Figure 6.2: Error in angular velocity as output by the MEKF after tuning using a simulation of the attitude dynamics. The error is the difference between the simulated angular velocity and the MEKF fitted angular velocity generated from measurements corrupted with Gaussian noise. The red line shows the 3σ angular velocity uncertainty, which should bound 99.73% of the angular velocity error.



over a typical early-mission CSSWE orbit. Although each simulation is slightly different due to its stochastic nature, the 3σ bounds reliably encase roughly 99% of the angular error. The short-term spikes in uncertainty are when the satellite is aligned with the local magnetic field; this alignment reduces the signal to noise ratio on the other axes of the magnetometer, resulting in a less certain attitude estimate. These uncertainty spikes are especially pronounced from 30 to 60 minutes, when the photodiode measurements are ignored during eclipse.

As described in Section 4.3.2, CSSWE uses a magnetometer and four photodiodes to observe its attitude. Figures 6.3 and 6.4 show the measurement residuals for each magnetometer axis and photodiode. The simulation shows that 99% of the measurement residuals are also within the 3σ uncertainty bounds, which are calculated using both the assumed standard deviations and the MEKF state uncertainty output.

Figure 6.5 shows the components of the scaled error Gibbs vector for each step in the simulated. As described above, the MEKF output may be trusted if these errors remain below 5° . After a settling time of about seven minutes, the error angles remain well below the 5° mark. However, the error angles do approach 5° during the eclipse period (30 to 60 minutes). Thus, the filter may be operating close to its performance limits during eclipse. On-orbit data was used for the next stage in filter tuning.

6.2.2 Empirical Filter Tuning

When working with the on-orbit CSSWE attitude data, a timespan of interest was set from the first on-orbit measurement (September 14, 2012, at 00:51:50 UTC) to midnight the day before a latch-up resulted in degraded photodiode measurements (October 14, 2012 at 00:00:00 UTC). This month of time has much variation, from initial tumbling to settling, from safe mode to science mode, and from relatively quiet geomagnetic conditions to a significant storm.

The empirical tuning was complicated further by uncertainty in the physical properties of the satellite. As explained in Section 6.1, the bar magnet moment \mathbf{m} and the inertia matrix $[I]$ are key factors in the performance of the PMAC MEKF. The inertia matrix $[I]$ was numerically calculated

Figure 6.3: Magnetometer measurement residuals from the simulated-input PMAC MEKF output. The residual is the difference between the measured value and the inertial model rotated into the body frame. The red line shows the 3σ angular uncertainty, which should bound 99.73% of the residuals.




Figure 6.4: Photodiode measurement residuals from the simulated-input PMAC MEKF output. The residual is the difference between the measured value and the inertial model rotated into the body frame. The red line shows the 3σ angular uncertainty, which should bound 99.73% of the residuals.

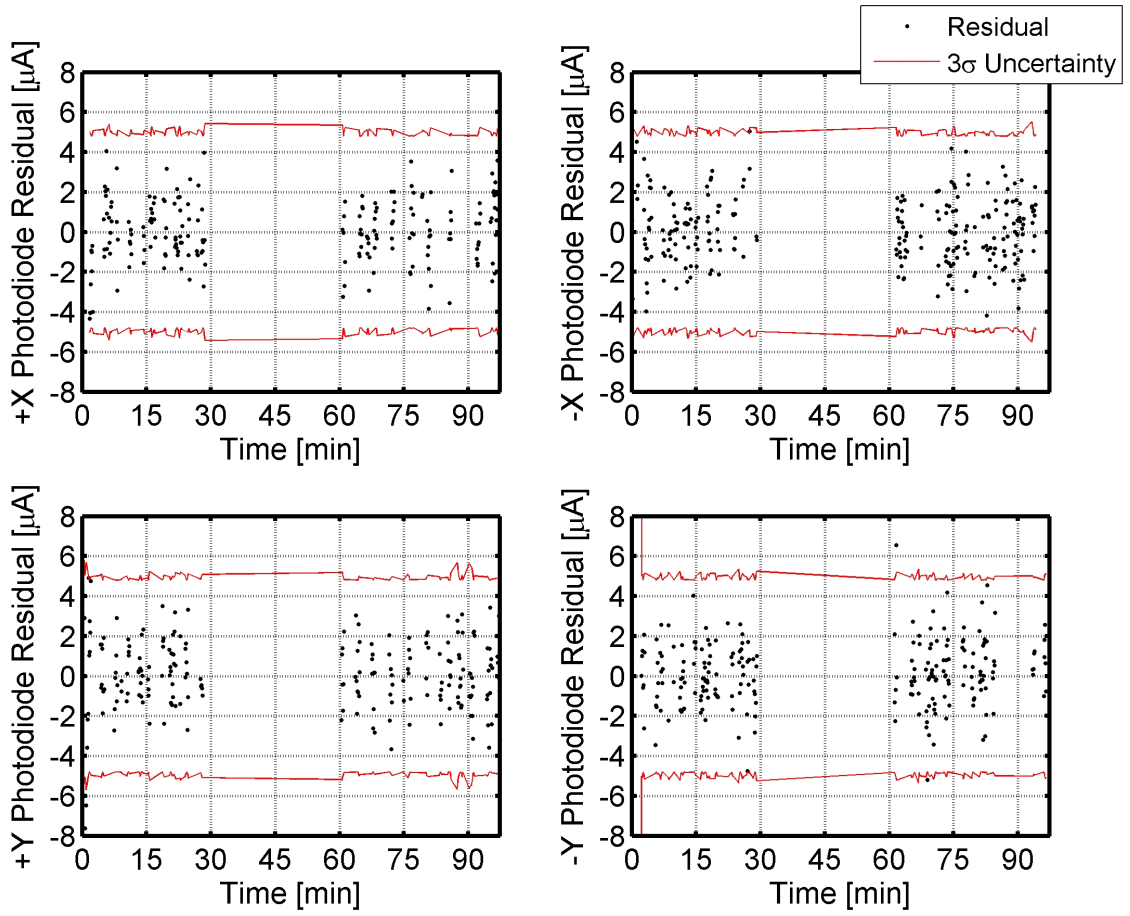



Figure 6.5: The scaled error Gibbs vector values from the simulated-input PMAC MEKF. When the error angle is $<5^\circ$, the MEKF output may be trusted.



from the spacecraft solid model. The uncertainty of this estimate is based on the completeness of the solid model; in the case of CSSWE, the $[I]$ uncertainty is estimated to be roughly 5%. The magnetic moment of the bar magnet was measured using the technique described in Section 7.2 which shows an uncertainty of 2%. The MEKF simulation (described in the previous section) was used to test filter performance given incorrect values of $[I]$ and \mathbf{m} at the maximum expected uncertainty; the results were not satisfactory.

However, both $[I]$ and \mathbf{m} should not change after all deployables are in their operational configuration; this makes on-orbit calibration ideal for these values. Constrained nonlinear optimization was used to determine the \mathbf{m} which minimized the filter measurement residuals. The fit was performed using early mission CSSWE data, while the spacecraft was still covering the majority of the attitude sphere. The inertia matrix was not fitted because it was found that \mathbf{m} and $[I]$ are not separable from the viewpoint of the CSSWE attitude response. Thus, the solid model $[I]$ and the fitted \mathbf{m} are used.

However, because the fit \mathbf{m} is based on minimizing the filter measurement residuals, the empirical tuning occurred simultaneously to the fit. A five hour timespan in the early mission is used to simultaneously fit \mathbf{m} , the photodiode and magnetometer alignment angles, and the photodiode scale factor. Longer fit times were attempted, but the results were degraded by the photodiode scale factor changes over their duration. The MEKF assumes photodiode 1σ standard deviations of $1.6 \mu\text{A}$, as outlined in Section 6.2.1. Each magnetometer axis 1σ standard deviation is set equal to the $\|\mathbf{B}\|$ error (measured vs. IGRF); this value varies over the timespan of interest as shown in Figure 4.5. The magnetometer X-axis contains a systematic error which is believed to be due to its proximity to a current-carrying wire within the satellite. Unfortunately, the time-varying calibration using the 10 minute average system current housekeeping telemetry lacks the temporal resolution necessary to adequately remove this error. This error is thus reflected in the tuned R matrix.

Starting with the simulation-based value of Q and expected level of R , the empirical filter-based fit is performed multiple times with various Q and R until the approximately 99% of the

measurement residuals are within the 3σ measurement uncertainty bounds. Empirical tuning shows that the best performance of the filter occurs when the magnetometer X-axis element of R is set to twice the nominal value (measured vs. IGRF $\|\mathbf{B}\|$) throughout the timespan of interest. Figure 6.6 shows the magnetometer position and orientation on the REPTile board within the satellite. The X-axis of the magnetometer is parallel to the currents that are generated by the REPTile ground plane.

Figures 6.7 and 6.8 show the tuned magnetometer and photodiode measurement residuals, respectively, for the approximately five hour early mission empirical fit timespan. As explained above, the magnetometer X-axis shows a systematic error which has been accounted for by doubling the standard deviation of that sensor. The photodiode residuals behave as expected. In both sensors, approximately 99% of the measurement residuals reside within the 3σ error bounds, proving that the empirical data has been properly fit. However, the decreased confidence in the magnetometer X-axis exacerbates the problem of high error angles during eclipse times, as shown in Figure 6.9. Although the figure shows about a dozen data points above the 5° threshold, the filter remains below this cutoff for the vast majority of the time. We have confidence in the filtered data during the timespan of interest because of the results of both tuning procedures.

6.3 CSSWE Attitude Determination

The tuned MEKF was applied to the first month of CSSWE on-orbit data, before the photodiodes were compromised by an on-orbit anomaly, as described in Section 4.4. The filter output is shown over this timespan, followed by a validation of the results.

6.3.1 Filter Output

The output from the MEKF using the on-orbit data is split into multiple sections. The first section shows data over the entire timespan of interest. The next two sections show 100 minute datasets representative of the PMAC performance before and after settling occurs.

Figure 6.6: A picture of the interior of the CSSWE as captured during final integration. The magnetometer is located on the corner of the REPTile electronics board (bottom left of the image); the system axes are indicated. The X-axis is parallel to the currents generated by the REPTile ground plane.




Figure 6.7: Magnetometer measurement residuals from the early mission on-orbit PMAC MEKF output. The residual is the difference between the measured value and the inertial model rotated into the body frame. The red line shows the 3σ angular uncertainty, which should bound 99.73% of the residuals. The systematic error in the X-axis magnetometer output is believed to be due to system currents which cannot be removed using the available telemetry. The expected standard deviation of the X-axis magnetometer data (within the R matrix) is increased to twice the nominal level to account for this error.

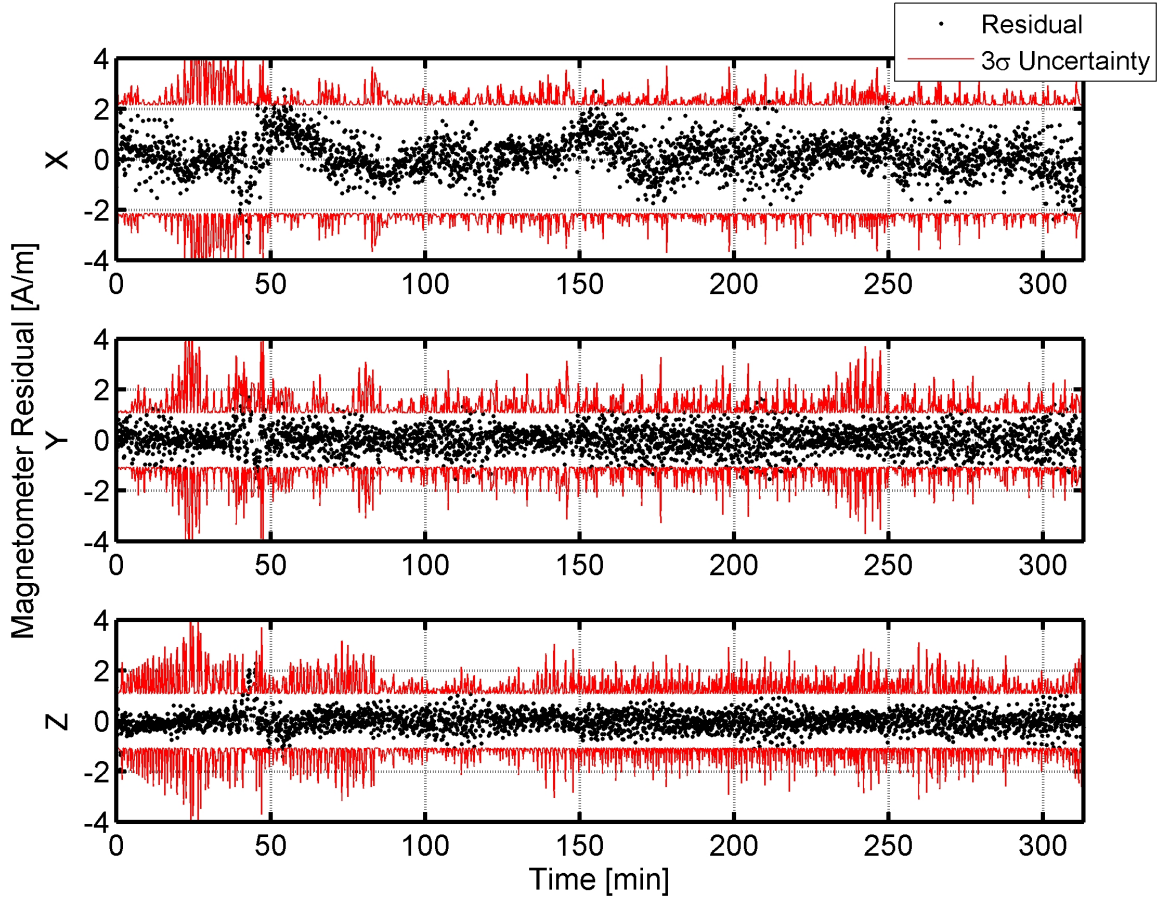


Figure 6.8: Photodiode measurement residuals from the early mission on-orbit PMAC MEKF output. The residual is the difference between the measured value and the inertial model rotated into the body frame. The red line shows the 3σ angular uncertainty, which should bound 99.73% of the residuals.

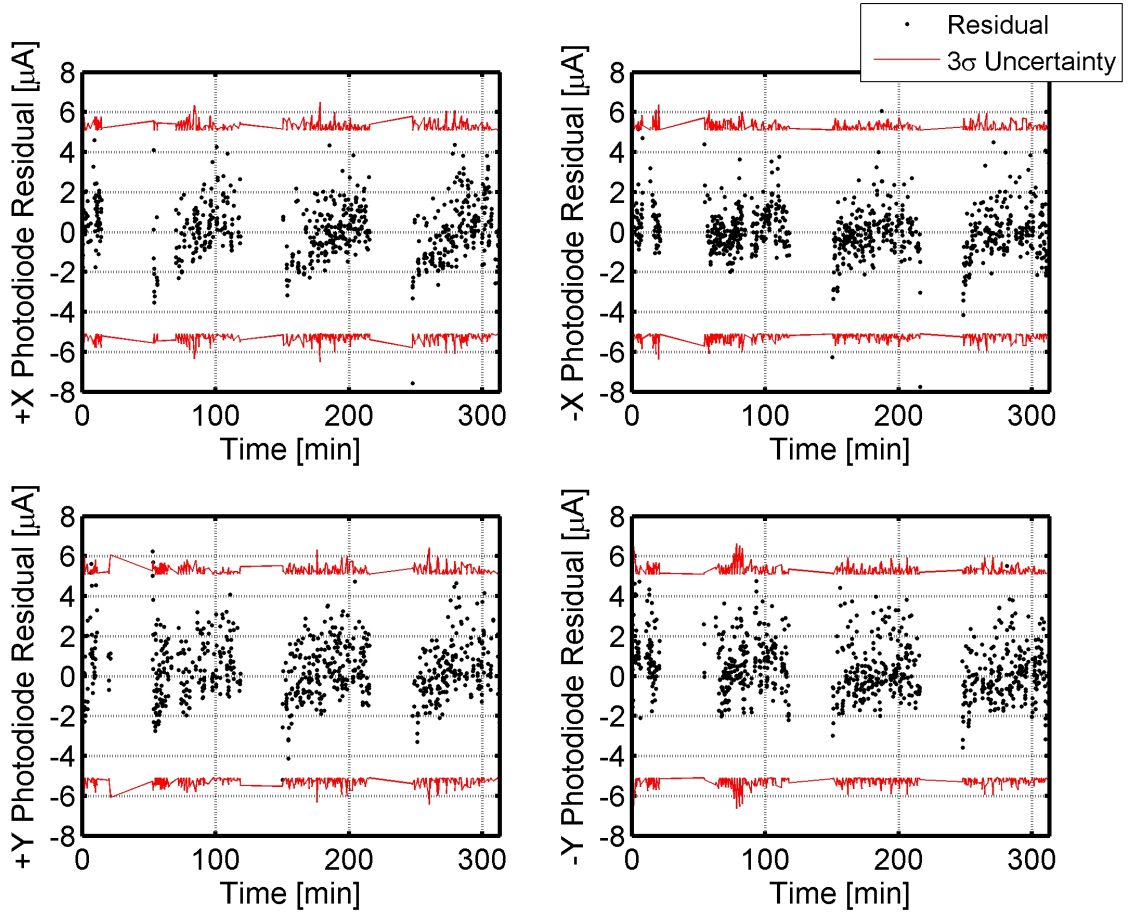



Figure 6.9: The scaled error Gibbs vector values from the early mission on-orbit PMAC MEKF. When the error angle is $<5^\circ$, the MEKF output may be trusted.



6.3.1.1 First 30 Days On-Orbit

Figure 6.10 shows the angular velocity components throughout the timespan of interest. The 3σ uncertainty bounds are included in the plot in red but there is no visible difference from the fitted values at the zoom level shown. The spikes in the fitted angular velocity and uncertainty bounds are due to gaps in the housekeeping data which is used to correct the magnetometer readings or gaps in the attitude data itself. The ${}^B X$ and ${}^B Y$ components of the angular velocity dampen within the first week, and remain at low values throughout the month. The roll about the ${}^B Z$ axis varies greatly within the first few days, then settles somewhat. The roll continues to evolve within a few degrees per second with a zero crossing about 13 days after launch. The zero roll rate causes the satellite exterior to experience a wider range of temperatures as detailed in Section 6.3.2.2.

Figure 6.11 shows the β angle (between ${}^B Z$ and the \mathbf{B} -field) with the 3σ uncertainty bounds. CSSWE settled to within 15° of the \mathbf{B} -field within seven days of launch. The increased magnitude of the 3σ uncertainty bound spikes in the middle of the dataset are likely related to a decreased roll rate combined with increased eclipse periods. As shown in Figure 4.6, the number of photodiodes visible to the sun (and thus the number of measurements available to the MEKF) is dependent on the roll angle. When the roll rate is slowed, the satellite has a longer timespan in which only one photodiode is viable for use. As the satellite eclipse period increases (see Figure 6.15), the photodiodes are not used for a longer period and the uncertainty bounds grow to a higher maximum as the MEKF is using less observations for a longer time.

The attitude estimate is further degraded due to the satellite orientation with respect to the local magnetic field. After the attitude has settled, the two magnetometer axes perpendicular to the \mathbf{B} -field have a much lower signal-to-noise ratio and the MEKF performance degrades as a result. This could be avoided by installing the magnetometer significantly off-axis from the satellite body frame.

The system energy evolution is a good “sanity check” for any dynamics problem. The rotational energy of a PMAC satellite may be split into the kinetic and potential energy. The rotational

Figure 6.10: The body to inertial angular velocity vector components are shown over the first month of CSSWE on-orbit operations. The 3σ uncertainty bounds are shown in red behind the fit data; their inclusion makes a negligible difference at this zoom level.




Figure 6.11: The angle β between the local magnetic field and the CSSWE $+^{\beta}Z$ axis is shown for the first month on-orbit with 3σ uncertainty bounds included in red. The 15° beta angle threshold indicating post-settling is denoted with a solid green line. The increased beta angles late in the dataset may be erroneous output due to satellite position error as explained in Section 4.3.2.3. The increased uncertainties in the middle of the dataset are likely due to decreased roll rates combined with increased eclipse times.




Figure 6.12: The rotational energy of the satellite over the first month on-orbit. The total energy is shown in black at bottom, while the top shows the kinetic (red) and magnetic potential (blue) energy components.





Figure 6.13: The histogram of the angle between the local magnetic field and the CSSWE + β X (REPTile aperture) axis. The histogram covers a ten day period starting ten days after launch (after the β angle has settled as shown in Figure 6.11). The mean angle of 89.5° is shown in green and a fitted normal distribution with a standard deviation of 1.87° is shown in red.



kinetic and potential energy equations are repeated below.

$$T_K = \frac{1}{2} \boldsymbol{\omega}^T [I] \boldsymbol{\omega} \quad (5.2)$$

$$T_P = -\mathbf{m} \cdot \mathbf{B} \quad (5.3)$$

The rotational potential energy is due to magnetic material within the satellite. Note that Equation 5.3 defines zero potential energy as when the bar magnet is perpendicular to the \mathbf{B} -field. The satellite possesses negative potential energy when the magnetic moment vector is less than 90° away from the \mathbf{B} -field. Note that the potential energy is directly related to the \mathbf{B} -field magnitude, which varies over the CSSWE orbit. Figure 6.12 shows the kinetic, potential, and total rotational energy over the first month on-orbit. The hysteresis rods decrease the rotational kinetic energy to nearly zero within the first week after launch. The potential energy decreases as the satellite aligns but the settled behavior echoes the orbital variations in the \mathbf{B} -field magnitude. The increased jitter in the kinetic energy after day twenty is believed to be an effect of the increased error in satellite position estimation (Figure 4.8) which similarly affects the beta angle and angular velocity estimates.

After settling, the angle between the REPTile aperture ($+^B X$) and the \mathbf{B} -field should be close to 90° . Figure 6.13 is a histogram of this REPTile alignment angle from days 10 to 20 after launch, after the attitude has settled and before the filter performance has degraded. As shown, the REPTile alignment angle relative to the magnetic field follows a normal distribution with a mean of 89.5° and a standard deviation of 1.87° . As mentioned in Chapter 4, this is advantageous as it provides a higher signal-to-noise ratio for the REPTile instrument.

Figure 6.14 shows MEKF scaled error Gibbs vector magnitude statistics over the first month on-orbit. The daily mean, standard deviation, and percentage of magnitudes greater than 5° are shown; each is a relative measure of the MEKF performance. By each measure, the MEKF performs best in the early mission once the initial angular velocity has settled somewhat. As noted

in Section 4.3.2, both the magnetometer and photodiodes require on-orbit calibration to achieve acceptable performance. The majority of the calibration parameters are set using early mission data; this is why the MEKF performance degrades after the first few days on-orbit. The REPTile instrument was activated for the first time on day twenty; this event sharply increases the error angle magnitude because the magnetometer is located on the REPTile electronics board (see Figure 4.4). From day 20 to 25, the REPTile detectors are being activated one by one. REPTile completes its commissioning on day 25; after this point, the magnetometer calibration better compensates for REPTile currents.

Figure 6.15 shows the daily maximum, mean, and minimum temperatures of the Command and Data Handling (C&DH) board in the interior of CSSWE. The orbit percent illuminated by the sun is also shown. It is easy to see that the interior temperature reflects the satellite insolation time. As explained in Section 4.3.2.1, the C&DH temperature is used as a proxy for the magnetometer temperature, which is responsible for the most significant time-varying magnetometer scaling and offset errors. The temperature correction is essentially extrapolating based on the first three days on-orbit; the magnetometer error can be expected to increase as the temperatures differ from the first three days. The magnetometer temperature has a nonlinear relationship to the magnetometer error which has consequences for the MEKF performance (Figure 6.14). Figure 6.15 shows a second source of MEKF uncertainty: the orbit eclipse period reaches its maximum near day 13, when the daily percent error angle magnitude is near its maximum; this is the effect of losing the photodiode measurements for a longer period.

6.3.1.2 Early Mission On-orbit Performance

Figure 6.16 and 6.17 show the β angle and angular velocity components with their associated 3σ bounds for a 100 minute period four hours after orbit insertion. Both plots show a satellite that has not yet settled to pointing parallel to the \mathbf{B} -field. The β angle is still varying up to almost 180° . The roll rate is about $4^\circ/\text{s}$ and the pitch and yaw rates are exchanging at $\pm 3^\circ/\text{s}$ due to gyroscopic torques induced by the satellite inertia matrix asymmetry. The fit uncertainty does not extensively

Figure 6.14: The daily mean and standard deviation of the scaled error Gibbs vector magnitude $\|\mathbf{a}_g\|$ are shown in blue squares and green circles, respectively. The daily percent of error angle magnitudes $\|\mathbf{a}_g\|$ which are greater than 5° is shown using the red triangles. Each dataset is a relative measure of the MEKF performance over the first thirty days on-orbit.





Figure 6.15: The daily maximum, mean, and minimum temperatures of the Command and Data Handling (C&DH) board (located within CSSWE) are represented by the red diamonds, green squares, and blue circles, respectively. The percentage of time the satellite is in the sun each day is represented by the brown dashed line.



increase during the eclipse period. In short, the MEKF is performing well.

6.3.1.3 Post-settling On-orbit Performance

Figure 6.18 and 6.19 show the β angle and angular velocity components with their associated 3σ bounds for a 100 minute period about ten days after orbit insertion. The angular velocity of the X and Y axes varies from $\pm 0.2^\circ/\text{s}$ and the roll rate is about $-0.3^\circ/\text{s}$. The roll rate estimate shows a minor discontinuity when the photodiodes enter use after the eclipse period. However, the discontinuity amplitude is within the uncertainty bounds of the fit, indicating that the filter is estimating uncertainty well.

CSSWE has now settled to within 5° of the **B**-field, though the 3σ uncertainty bounds of β sometimes reach large values, especially during eclipse. Note that the angular uncertainty is based on the assumption that the scaled error Gibbs vector uncertainty is a small angle and thus approximates the pitch, roll, and yaw uncertainties. Thus, uncertainties much greater than 5° should not be trusted. The uncertainty peaks during insolation (both β and angular velocity) are due to a temporary loss of the photodiodes due to the sun position in the body frame, as shown in Figure 6.20. The figure also shows the magnitude of the filter error angle; the small angle assumption is only broken once, during the transition from eclipse to insolation. Thus, even during the periods when the filter output uncertainty is suspect, the fitted attitude remains trustworthy.

6.3.2 Attitude Determination Validation

The CSSWE attitude determination results may be validated using independently-measured on-orbit data. Two events which validate the filtered attitude output are detailed below.

6.3.2.1 Antenna Deployment Event

Due to requirements set by the CubeSat Design Specification [13], the satellite was launched with the antenna stowed within the CubeSat. CSSWE was programmed to automatically deploy the antenna two hours after orbit insertion; the satellite event log vouches for the deployment. The

Figure 6.16: The angle between the local magnetic field and the $+^{\mathcal{B}}Z$ axis over a 100 minute period five hours after orbit insertion. The 3σ uncertainty bounds are shown in red and the insolation times are marked in blue.




Figure 6.17: The satellite angular velocity components over a 100 minute period five hours after orbit insertion. The 3σ uncertainty bounds are shown in red and the insolation times are marked in blue.




Figure 6.18: The angle between the local magnetic field and the $+^{\mathcal{B}}Z$ axis over a 100 minute period ten days after launch. The 3σ uncertainty bounds are shown in red and the insolation times are marked in blue.




Figure 6.19: The satellite angular velocity components over a 100 minute period ten days after launch. The 3σ uncertainty bounds are shown in red and the insolation times are marked in blue.





Figure 6.20: The top plot shows the angle between the sun vector and the CSSWE +^BZ axis (blue) and the high incidence angle cutoffs at 20° and 160° (dotted red lines). The uncertainty peaks during insolated times occur when the sun vector is within 20° of $\pm^B Z$, suggesting that the loss of the photodiodes causes the rapid increase in uncertainty. The bottom plot shows the magnitude of the scaled error Gibbs vector over the first 100 minutes (black circles) with the 5° limit to the small angle assumption denoted (red dotted line). Only the point directly following the transition back to insolation has an angular magnitude greater than five degrees.



spring-steel antenna deployment was expected to slightly change the magnetometer calibration parameters because of its ferromagnetic makeup. Instead, analysis of the on-orbit data shows that the deployment of the spring-steel antenna significantly changes both the magnetometer calibration and the magnetic moment of the satellite. The satellite motion difference was such that two different satellite magnetic moments were calculated during the fitting process described in Section 6.2.2: $0.84 \text{ A}\cdot\text{m}^2$ and $0.55 \text{ A}\cdot\text{m}^2$ before and after the antenna deployment, respectively. Figure 6.21 shows the system energy before and after the deployment event (denoted with the green line at 2 hours after orbit insertion); there is a significant decrease in the system rotational energy when the antenna deploys. Both the potential and kinetic rotational energy dynamics change significantly after the antenna deployment because the magnetic torque of the satellite dominates all other external torques. The antenna deployment does not create a visible difference in the filtered attitude or angular velocity estimates, indicating that the torque due to antenna deployment is insignificant.

6.3.2.2 Solar Panel Temperature Distribution

The satellite roll rate has many implications. Much like meat roasting on a spit, when the roll rate approaches zero one side of the satellite can burn while the opposite side freezes. Figure 6.22 shows that the solar panel temperatures can reach extreme values when the roll rate is near zero. The effect is so pronounced that the panels can reach nearly 100°C even during a period of maximum eclipse time. The independently-measured solar panel temperature data are in good agreement with the MEKF fit.

Figure 6.21: Satellite rotational energy before and after the antenna deployment event (represented by the green line). The top plot shows the kinetic (red) and potential (blue) rotational energy while the bottom plot shows the total rotational energy.





Figure 6.22: The top plot shows the maximum (open diamond), mean (filled square), and minimum (open circle) temperature of each of four solar panels; the orbit insolation percentage is overlaid (brown dotted line). The bottom plot shows the absolute value of the roll rate as estimated by the MEKF.



Chapter 7

Magnetic Measurement

As the old adage says, “garbage in equals garbage out”; inaccurate input translates to inaccurate output. The key inputs to PMAC simulation are magnetic torques; thus the parameters which govern these torques must be accurately measured to ensure that the simulation results are trustworthy. This chapter describes measurement techniques for both the static magnetic moment \mathbf{m} and the hysteresis rods. First, the design of a Helmholtz cage for accurate magnetic measurements is described (Section 7.1). Next, the static magnet magnetic moment \mathbf{m} is measured (Section 7.2). A Helmholtz-cage-based hysteresis rod measurement method is described (Section 7.3); this method is capable of measurement in the presence of other magnetic materials which could degrade the dampening capabilities of a hysteresis rod.

7.1 Helmholtz Cage

A Helmholtz cage is a set of six wire coils (two per orthogonal axis) which can be used to provide an arbitrary uniform field within the volume enclosed by the coils; the properties of the uniform field depend upon the cage dimensions and available current. The cage operates on the principle that a current-carrying wire will produce a magnetic field, and can be used to both cancel the inherent magnetic field and supply an arbitrary magnetic field. The use of two coils per axis allows for a uniform field to be generated between the coils. A Helmholtz cage was constructed to perform multiple magnetic measurements relevant to a PMAC system. The theory of a Helmholtz cage is presented, followed by the design of the cage and post-build test results.

7.1.1 Theory

The Helmholtz cage theory is based on manipulation of the Biot-Savart law: [81]

$$\mathbf{B} = \frac{\mu_0 n I}{4\pi} \int \frac{d\mathbf{l} \times \hat{\mathbf{r}}}{r^2} \quad (7.1)$$

where \mathbf{B} is the magnetic field vector, μ_0 is the permeability of free space, I is the current of the wire being analyzed, $d\mathbf{l}$ is an infinitesimal current-carrying wire element, r is the distance from element $d\mathbf{l}$ to the point of \mathbf{B} measurement, $\hat{\mathbf{r}}$ is a unit vector in the direction of r , and n is the number of turns in the current-carrying wire. The Biot-Savart law is used to determine the axial magnetic field due to two coils perpendicular to a single axis. First, the on-axis field due to a single wire segment on one of the coils is determined. A single coil is shown in Figure 7.1a. Here a is one-half of the side length A and the point of measurement P lies in the center of the coil. The distance z is measured from the center of the two coils, thus the -Z coil lies at $z = -h/2$, where h is the distance between the coils. Integration variables are shown as l and θ .



Figure 7.1b shows the relationships between the single wire segment and the field generated at on-axis location P . Note that the total field \mathbf{B} generated by the wire segment l is perpendicular to plane shown in Figure 7.1b. The trigonometric relationships shown in Figures 7.1a and 7.1b may be differentiated and combined with Equation 7.1 to arrive at Equation 7.2, which describes the infinitesimal magnetic flux generated by an infinitesimal length wire:

$$d\mathbf{B} = \frac{\mu_0 n I}{4\pi \rho} \cos \theta d\theta \hat{\mathbf{B}} \quad (7.2)$$

where $\hat{\mathbf{B}}$ is the magnetic flux unit vector and ρ is the axial distance from the center of the wire segment to the point P . Equation 7.2 may be integrated from $\theta_1 = -\sin^{-1}\left(\frac{a}{\sqrt{a^2 + \rho^2}}\right)$ to $\theta_2 = \sin^{-1}\left(\frac{a}{\sqrt{a^2 + \rho^2}}\right)$ to yield Equation 7.3, the magnetic flux generated by a single length of wire of length $2a$ as shown in Figure 7.1a. The direction of this field is perpendicular to the plane shown in Figure 7.1b:

$$B = \|\mathbf{B}\| = \frac{\mu_0 n I}{4\pi \rho} \left(\frac{2a}{\sqrt{a^2 + \rho^2}} \right). \quad (7.3)$$

Figure 7.1: Helmholtz geometry. (a) The $-Z$ current-carrying coil and the resultant field generated at point P . (b) The generated field is perpendicular to the surface from the wire segment to point P .



When four wire segments of equal length are combined with the same current traveling in a counterclockwise direction (as shown in Figure 7.1a), it is apparent that only the contribution of the magnetic field perpendicular to the coil will remain. Thus, Equation 7.3 is multiplied by $\sin(\phi) = \frac{a}{\rho}$ and a factor of 4 to yield the total axial field due to the -Z square coil:

$$B_{-z} = 4B \sin \phi = \frac{\mu_0 n I}{4\pi} \left(\frac{a}{\rho} \right) \left(\frac{2a}{\sqrt{a^2 + \rho^2}} \right) \quad (7.4)$$

where the direction of generated magnetic flux is given by the right hand rule based on the current flow. The field for the second coil can be calculated in much the same way, the only difference being $\rho_{-z} = \sqrt{z - h/2}$ and $\rho_{+z} = \sqrt{z + h/2}$. Equation 7.5 gives the total field produced by both square coils as a function of axial distance z from the center of the coils:

$$B_z = \frac{\mu_0 n I}{\pi} \left(\frac{2a^2}{(a^2 + (z - h/2)^2) \sqrt{2a^2 + (z - h/2)^2}} + \frac{2a^2}{(a^2 + (z + h/2)^2) \sqrt{2a^2 + (z + h/2)^2}} \right). \quad (7.5)$$

7.1.2 Design

The Helmholtz cage was designed for the testing of CubeSat spacecraft. The following requirements governed the design of the Helmholtz cage:

- (1) The Helmholtz cage test volume shall have dimensions of 30cm×30cm×30cm.
- (2) The Helmholtz cage test volume shall provide 1% theoretical field uniformity along each axis.
- (3) The Helmholtz cage shall provide a magnetic field strength range of ± 50 A/m on each axis.

The largest standard CubeSat is a 3U size, with dimensions 10cm×10cm×34cm [13]. Thus, requirement 1 ensures that the test volume is sufficient to measure any standard CubeSat (a 3U CubeSat easily fits when diagonally placed). Requirement 2 defines the expected uniformity of each axis of the Helmholtz cage. It is given in terms of theory because the empirical uniformity of each axis which may be measured by a magnetometer is dependent on the component of the earth

field along that axis. If the Helmholtz cage were placed in a magnetically noisy or non-uniform environment, the output field of the cage would be similarly noisy. However, a benefit of the small test volume is portability. If an environment is particularly noisy or non-uniform, the cage may be moved to a location that possesses a more stable field. Requirement 3 defines the configurable field within the Helmholtz cage after nullifying the earth-based local magnetic field.

In order to ensure these requirements were met, Equation 7.5 is used with various spacing between coils to determine the configuration that meets requirement 2. Figure 7.2 shows the results of this analysis. Spacing h is defined as the axial distance between the two coils of a single axis. The spacing h is given in terms of one half of the side length (a). Thus, the analysis shown in Figure 7.2 is independent of cage size. The analysis shows that the best combination of uniform field and uniform distance is given by $h = 1.2a$ which is the chosen spacing for the Helmholtz cage. Next, the side length of the Helmholtz cage is driven by requirement 1 above. In order to provide a 1% variation in field over a 30cm distance, the Helmholtz cage side length $2a$ is set to 62.25cm and the coil spacing is set to $h = 1.2a = 37.24$ cm. Figure 7.3 shows the normalized field resultant from these settings. The figure shows that along the axis from -15cm to 15cm, the theoretical flux density stays within 1% of the maximum value of the field, satisfying requirement 2.

There is a trade-off between the number of wire turns on each coil and the ability of the power supply to provide current to the coils. Each power supply has a maximum voltage and current it can deliver. The B&K Precision 9130 was chosen for use with the built Helmholtz cage, it has $V_{max} = 30$ V and $I_{max} = 3$ A. However, the relays which switch the polarity of the power supply output each have a current limit of $I_{max} = 2$ A. As shown in Equation 7.5, the current through the coils I is directly proportional to the magnetizing field H within the Helmholtz cage. As the number of turns per coil increases, so does the resistance of each axis. Ohm's law shows that as the resistance R increases, eventually the power supply current will be limited by maximum voltage output V_{max} . For the B&K Precision 9130 power supply, the total resistance per axis must be less than 15Ω to allow the power supply to provide $I_{max} = 2$ A. CSSWE's Helmholtz cage uses 30 turns of 22AWG magnet wire per coil, providing a total resistance of 8.5Ω per axis. Thus, the

Figure 7.2: The axial fields resultant of various coil spacings are shown. Here h is the axial distance between coils and $2a$ is the side length of each square coil.




Table 7.1: All purchases needed to complete the Helmholtz cage are shown.

Part	Quantity	Cost (USD)
Acrylic walls (24in×24in×0.5in)	4	207
500ft. 22AWG magnet wire	3	73
B&K Precision 9130 computer-controlled 3-axis power supply	1	895
Micromag3 magnetometer	1	60
PIC18F452 QwikFlash microcontroller	1	132
BNC connectors	3	22
DPDT relays	3	15
Total		1,404

power supply can easily provide 2A of current to each axis. Using Equation 7.5 with $I = 2\text{A}$ yields $B_z > 125\mu\text{Tesla}$ ($H_z > 100\text{ A/m}$) for $-15\text{cm} < z < 15\text{cm}$. The maximum earth-based magnetic field magnitude at the University of Colorado (40°N , 105°W , 1655m elevation) is $< 40\text{ A/m}$. Thus, even with the worst-case alignment the Helmholtz cage is able to provide a magnetic field strength of at least 60 A/m per axis, satisfying requirement 3.

7.1.3 Assembly

Table 7.1 shows the parts needed to complete the Helmholtz cage setup; the total hardware cost is about \$1400 with the vast majority of that cost being the computer-controlled power supply. The Helmholtz cage construction began by assembling the 0.5 in. thick acrylic walls. The top and bottom of the cage were left open to allow access to the interior volume. Next, plastic guide rails and wooden spacers were added to ensure all coils are tight and square. Magnet wire was then wrapped around each axis one coil at a time, with care to ensure that both coils were wrapped in the same direction. Wooden legs were added to keep the cage weight off the magnet wire coils. Finally, BNC connectors were added to each axis for ease of connection to the power supply. Figure 7.4 shows the finished Helmholtz cage.

Figure 7.3: The theoretical axial magnetic field of the as-built Helmholtz cage with a spacing of $h = 1.2a = 37.24\text{cm}$. The magnetic flux density B is normalized by the maximum flux density.





Figure 7.4: The finished Helmholtz cage with side length $2a=62.23\text{cm}$ and $h=1.2a=37.24\text{cm}$. The PNI Micromag 3-axis magnetometer and analysis grid are also shown.



Much of the difficulty involved in the use of a Helmholtz cage is controlling the current output to each coil. Figure 7.5 shows all support hardware and connections necessary for operation of the Helmholtz cage. A PIC18F452 microcontroller is programmed to listen for queries from the serial port and relay the PNI MicroMag 3-axis magnetometer readings when requested. The microcontroller is also responsible for controlling the relay state when commanded via the serial port. Because the programmable power supply is only capable of positive currents, three socket-mounted double pull double throw (DPDT) relays are used to invert the current on each axis as necessary.

A LabView software interface was coded to allow a user to avoid low-level tasks by simply setting a desired constant arbitrary field within the Helmholtz cage. When the software is initializing, it empirically determines the alignment and sensitivity of the magnetometer relative to the Helmholtz cage (described below). After initialization, the software waits for the user to input the desired magnetic field value on each axis. Once the input is collected, the software varies the magnitude and polarity of the current on each axis until the user settings have been achieved.

All magnetic testing described below assumes a constant proportion between current to each coil and the H -field produced perpendicular to that coil. The “Helmholtz constant” has units of m^{-1} , as it represents the magnetizing field (A/m) per current (A) through the Helmholtz coils. The Helmholtz constant theory is based on Equation 7.5; if μ_0 and I are moved to the left side of the equation, the right side is the theoretical Helmholtz constant $\text{HC}_{\text{theory}}$. In the center of the Helmholtz cage as built, the theoretical Helmholtz constant $\text{HC}_{\text{theory}}$ is $58.8 \pm 0.78 \text{ m}^{-1}$. This theory-based constant is equivalent on each axis. However, the Helmholtz constant may also be empirically derived through initial calibration. When the LabView software initializes with no hysteresis material present, each axis of the cage is calibrated separately by measuring the magnetic field at two different current values. The empirical calibration constant is then calculated for each axis as follows:


$$\text{HC}_{\text{empirical}} = \frac{B_2 - B_1}{\mu_0(I_2 - I_1)} \quad (7.6)$$

where B_1 is the magnetometer measurement when the coils have current I_1 and B_2 is the magnetometer measurement when the coils current I_2 . The constant μ_0 is needed because the magnetometer measures the B -field (not H -field) within the cage. Testing shows that the empirical Helmholtz constant can change from test to test depending on the alignment of the magnetometer relative to the coils. However, for a properly-aligned, calibrated magnetometer, the empirical Helmholtz constant $\text{HC}_{\text{empirical}}$ is $58.8 \pm 0.043 \text{ m}^{-1}$, which is within the uncertainty bounds of $\text{HC}_{\text{theory}}$.

7.1.4 Characterization

Much of the usefulness of the Helmholtz cage depends upon the uniformity of the post-nullification magnetic field within the cage. Thus, the planar field of the Helmholtz cage was mapped using a magnetometer. The magnetometer is used to map a 5cm resolution grid by measuring the total magnetic flux density at each point. This mapping was performed twice: once with no current through the Helmholtz coils and once with the Helmholtz cage set to nullify the local field. Note that the magnetometer is placed in the center of the Helmholtz cage during the nullification process and moved afterward to make the measurements. Figures 7.6a and 7.6b show the magnetic flux density magnitude as a function of planar position. At the time and location of the test, the earth field varied by $5\mu\text{Tesla}$ over the test area. This caused some variation in the zeroed magnetic field, as shown in Figure 7.6b. As shown, when the Helmholtz cage is zeroed, the test area magnetic field takes the shape of a bowl with a square base: the $\pm 10\text{cm}$ central area has a measured uniformity of $2 \mu\text{Tesla}$, but the field near the edges of the test area reaches $10 \mu\text{Tesla}$. The measured field strength variation after field nullification is 5%; this is roughly equivalent to the local field variation before nullification. Repeating the test in a more uniform local field may yield better results. However, a $20\text{cm} \times 20\text{cm} \times 20\text{cm}$ test volume maintained at $0\text{-}2\mu\text{Tesla}$ is sufficient for the magnetic testing described later in this chapter. This uniformity was measured for a nullified field but is expected that a Helmholtz-cage-produced arbitrary field of up to $100 \mu\text{Tesla}$ per axis

Figure 7.5: Helmholtz hardware chain showing signals from the lab computer converted to the appropriate current through the Helmholtz coils.



will possess equivalent uniformity.

7.2 Bar Magnet Measurement

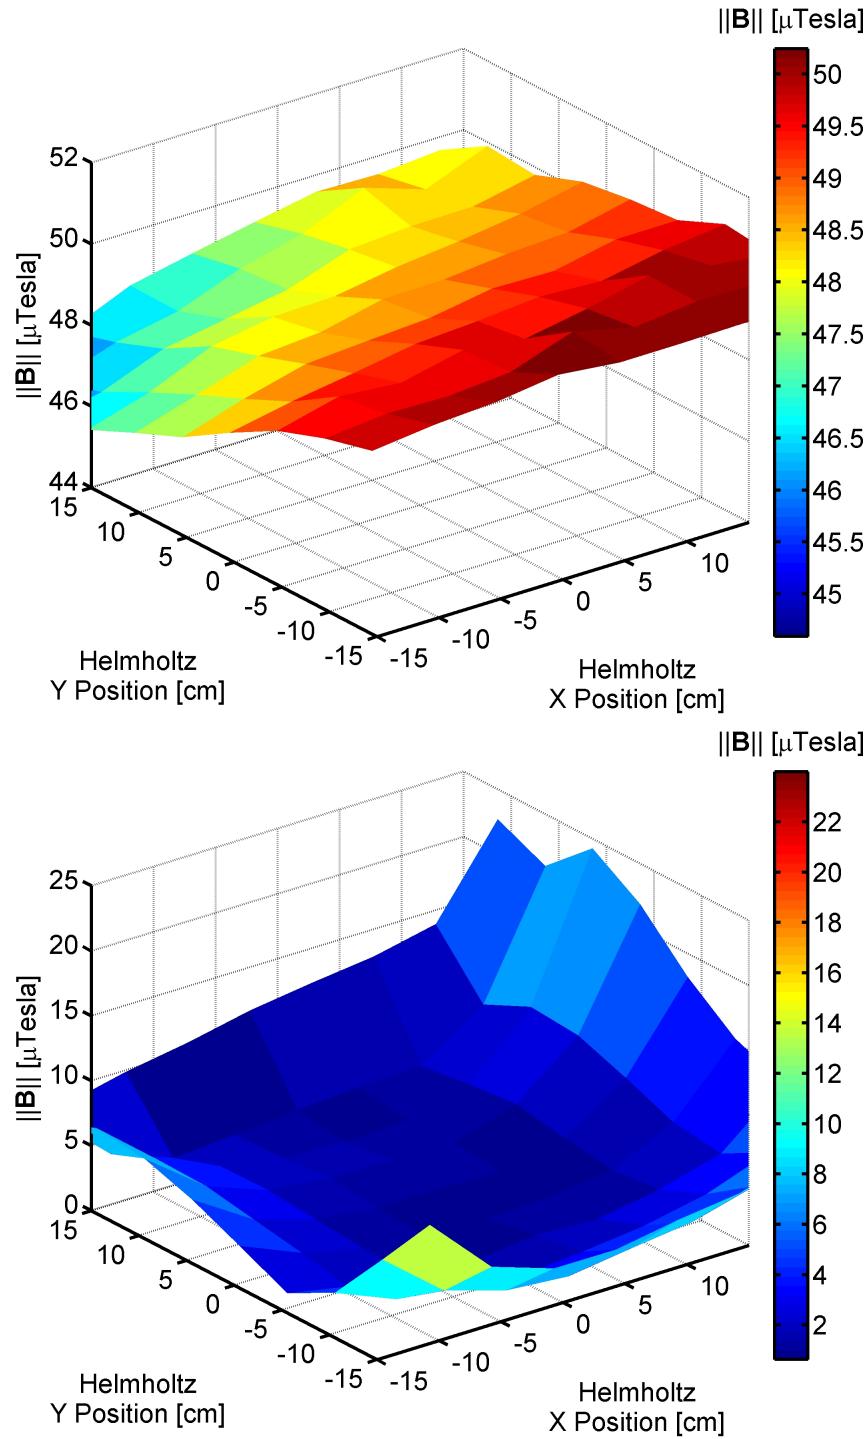
The Helmholtz cage was used to measure the magnetic moment of the bar magnet chosen for use in CSSWE. The steps for the bar magnet measurement are given below:

- (1) Ensure bar magnet is far from Helmholtz cage.
- (2) Take measurement of local magnetic field while magnetometer is at the center of the Helmholtz cage.
- (3) Supply current to the Helmholtz coils to nullify the magnetic field on all three axes. Record the power supply current provided to each axis.
- (4) Using the analysis grid (shown in Figure 7.4), move the PNI MicroMag 3-axis magnetometer, measuring the B -field magnitude in the 30 cm \times 30 cm area defined by the grid in steps of 5 cm.
- (5) Place the bar magnet in the center of the analysis grid.
- (6) Again using the analysis grid, record the B -field magnitude at the same positions as in step (4). Ignore grid positions less than 8cm from the bar magnet position as this may saturate the magnetometer and the dipole model is more accurate with greater distance from the dipole.

Once the two datasets are recorded and corrected for magnetometer calibration error, the B -field magnitude at each grid position due to the bar magnet alone may be calculated by subtracting the no-bar-magnet data from the bar-magnet data. The measured magnetic flux density magnitude $\|\mathbf{B}\|_{\text{measured}}$ data are fit to the magnetic dipole magnitude formula using a non-linear least squares fit. The fitted magnetic dipole magnitude equation is [60]:

$$\|\mathbf{B}\|_{\text{measured}} = \left(\frac{\mu_0 m_{\text{fit}}}{4\pi r^3} \right) \sqrt{1 + 3 \cos^2 \nu} \quad (7.7)$$

Figure 7.6: The magnetic flux density magnitude of the Helmholtz cage 30cm×30cm test area is shown with 5cm resolution. The upper plot (a) shows the distribution within the cage when there is no current in the coils; the lower plot (b) shows the distribution when the cage is set to nullify the magnetic field. Note the change in B -field scaling between a and b.



where m_{fit} is the fitted magnetic moment, r is the distance from the magnet to the measurement location, and ν is the magnetic co-latitude of the measurement location (90° represents a magnetometer position co-planar with the bar magnet position). Three independent measurements were performed, yielding a bar magnet magnetic moment $||\mathbf{m}|| = 0.80 \pm 0.017 \text{ A}\cdot\text{m}^2$ after a small sample size correction to the standard error of the mean has been applied [33]. Figure 7.7 shows the magnitude of the magnetic flux density and the magnetic dipole fit for the bar magnet measurement data for one of these measurements.

7.3 Hysteresis Measurement

The crux of a PMAC system is the hysteresis dampening. It is the most difficult torque to predict and can greatly effect the system settling time. The goal of hysteresis rod magnetic measurement is to determine the relationship between the earth-based H -field and the magnetic torque due to the hysteresis rods. As explained in Section 2.2.6, a relation from H_a to m is sufficient to determine the torque versus the earth's local magnetic field. Equation 2.7 translates the measurement of m to a measurement of the average magnetic flux density within the rod.


The use of a sense coil within a Helmholtz cage is a novel method of hysteresis rod measurement. The measurement method presented in this section is able to supply a magnetizing field in a large volume and in any direction. This allows multiple samples to be magnetized simultaneously to test their coupled magnetic performance.

7.3.1 Theory

A sense coil connected to an integrator is capable of directly measuring the average interior magnetic flux density of a sample. The Helmholtz cage enables the user to apply a varying magnetizing field over a large test volume. The addition of a sense coil to the Helmholtz cage results in a unique setup which is capable of system-level measurements at a nanosatellite scale.

The theory begins with Faraday's law, which states that a wire coil shall have a voltage generated within it proportional to the rate of change of magnetic flux through the coil [17]:

Figure 7.7: A single bar magnet measurement dataset with the fitted $0.80 \text{ A}\cdot\text{m}^2$ magnetic dipole overlaid. The B -field data as recorded by the PNI MicroMag 3-axis magnetometer are shown in black. The magnetic dipole field is shown using the colored surface plot. The dipole field within 8cm of the origin is omitted for clarity.



$$\mathcal{E} = -N \frac{d\phi}{dt} \quad (7.8)$$

where \mathcal{E} is the voltage induced in the coil, N is the number of turns of the search coil, ϕ is the magnetic flux in SI units of weber ($1 \text{ Wb} = 1 \text{ Tesla} \cdot \text{m}^2$), and t is time. Note that voltage is only induced by a changing magnetic flux through the coil. Equation 7.8 may be rearranged and integrated:

$$\int_0^t \mathcal{E} dt = -N \int_{\phi_1}^{\phi_2} d\phi = -N \Delta\phi \quad (7.9)$$

where $\Delta\phi$ represents the change in ϕ . Because the magnetic flux is directly related to the magnetic flux density by the area of the search coil A , Equation 7.9 can be written in terms of magnetic flux density:

$$\int_0^t \mathcal{E} dt = -NA\Delta B. \quad (7.10)$$

The integration in Equation 7.10 may be carried out one of two ways: through hardware or software. The hardware method uses an integrator circuit as shown in Figure 7.8. Commercial versions of this circuit (sometimes called a “fluxmeter”) are available for purchase but they were found to be prohibitively expensive. The hardware integrator circuit behaves as follows [17]:


$$\mathcal{E}_{out} = (RC)^{-1} \int_0^t \mathcal{E}_{in} dt \quad (7.11)$$

where R is the resistor value and C is the capacitor value. Thus, for a sense coil in series with the integrator:

$$\mathcal{E}_{out} = - \left(\frac{NA}{RC} \right) \Delta B. \quad (7.12)$$

The values of the resistor and capacitor may be calculated based on the materials to be measured. However, after building a hardware-based integrator, we found that temperature- and offset-voltage-based drift made measurements with the hardware integrator difficult. This drift caused the measured hysteresis loops to be significantly distorted and unsuitable for fitting with a hysteresis model. The hardware could be calibrated for one measurement at a time, but the system would drift too much over the course of multiple measurements to be useful for fitting purposes.

Figure 7.8: Hardware integrator circuit for magnetic measurement. Image from Cullity & Graham [17].



A commercial fluxmeter is expensive partially because it is built to correct for these drifts. The implemented low-cost solution is to perform the integration digitally.

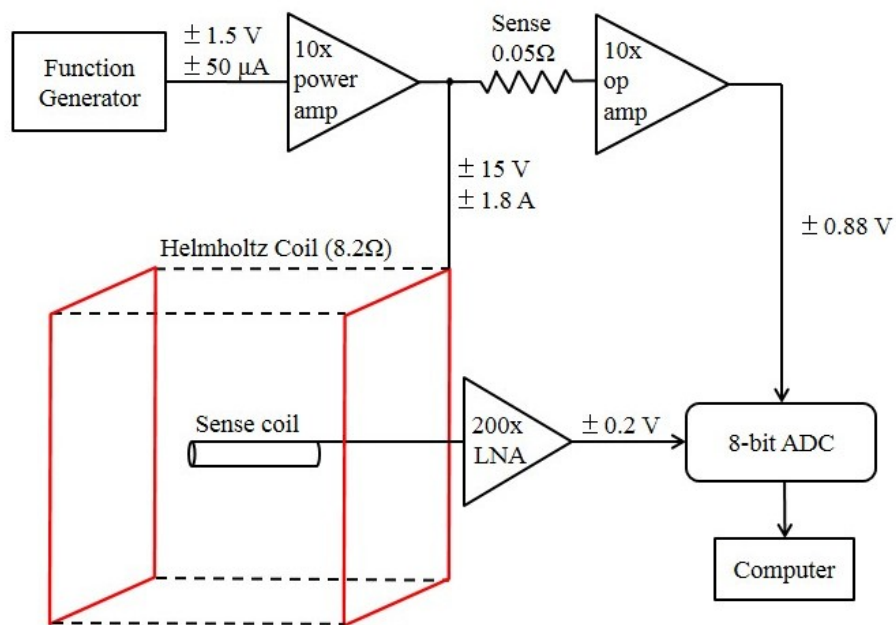
7.3.2 Setup

Figure 7.9 is a block diagram of the hardware setup for hysteresis measurement using the Helmholtz cage and a sense coil. Each element of the hardware setup is detailed below.

7.3.2.1 Sense Coil

The magnet wire cannot be wrapped around the hysteresis rod itself due to the minimum bend radius of the magnet wire and the risk of bending the hysteresis rod (thus damaging its magnetic properties). Instead, a sense coil is built such that magnetic samples can be slid within the coil when desired. In order to measure the average B-field of the magnetic sample, the sense coil length should extend the length of the sample. In this case, the sample is the hysteresis rod with a length of $97.17 \pm 0.03 \text{ mm}$ and cross-sectional area $A_m = 0.805 \pm 0.00064 \text{ mm}^3$. The sense coil is built by tightly winding 36 AWG magnet wire (manufacturer-listed diameter of $0.1400 \text{ mm} \pm 0.0013$) around a nonmagnetic aluminum tube with an inner/outer diameter of $5.00/5.30 \pm 0.03 \text{ mm}$ which yields a cross-sectional sensing area $A_s = 21 \pm 1.8 \text{ mm}^2$. The as-built sense coil (shown within the

Figure 7.9: The hysteresis rod measurement hardware setup block diagram. The labeled voltages and currents are for H-field cycling at ± 100 A/m amplitude; the amplitude (and subsequent voltages and currents) may be decreased as desired.



Helmholtz cage in Figure 7.10) has a wire-wrapped length of $96.48 \pm 0.03 \text{ mm}$; the number of turns is calculated to be $N = 691 \pm 6$.

Because the sense coil is wrapped around a hollow tube, the magnetic flux picked up by the sense coil is due to both the magnetic flux through the magnetic material and the magnetic flux through the air surrounding it (yet still within the coil). Thus, the magnetic flux density of each measurement is corrected as follows: [17]

$$B_{\text{true}} = B_{\text{apparent}} - \mu_0 H_a \left(\frac{A_s - A_m}{A_m} \right) \quad (7.13)$$

where B_{apparent} is the magnetic flux density as measured via Equation 7.10, H_a is the applied field, A_s is the cross-sectional area of the sense coil, and A_m is the cross-sectional area of the magnetic material.

7.3.2.2 Helmholtz Cage


The Helmholtz cage is used to provide the changing magnetizing field strength which causes the magnetic flux density within a magnetic sample to vary. The power supply / relay combination which drives each Helmholtz coil set is capable providing $\pm 2 \text{ A}$ current with a resolution of $10 \pm 1 \text{ mA}$. Using the Helmholtz constant developed in Section 7.1.3, this translates to an ability to control the static magnetizing field (including field nullification) in steps of $0.59 \pm 0.059 \text{ A/m}$.

The Helmholtz cage is small enough to allow for orientation changes which can result in improved performance. Before hysteresis rod testing begins, the Helmholtz cage is oriented such that the X-axis is perpendicular to the magnetic field. This allows for testing over the full $\pm 100 \text{ A/m}$ magnetizing field range which the Helmholtz cage can produce.

7.3.2.3 Other Hardware

The Agilent 33220A function generator is used to generate a 1 Hz sine wave with an amplitude of up to $\pm 1.5 \text{ V}$ for a desired H-field of $\pm 100 \text{ A/m}$. This low-current signal is fed to an Apex PA16 power operational amplifier, which multiplies the voltage by a factor of ten while acting as a current

Figure 7.10: The finished hysteresis measurement sense coil (A.) with approximately 691 turns of 36 AWG wire. A pen is shown next to the sense coil for length scale reference; the MicroMag3 3-axis magnetometer (B.) is also visible.



source for a single coil set of the Helmholtz cage. The maximum output current of $\pm 1.8\text{A}$ supplies an H-field of about $\pm 100\text{ A/m}$ to the center of the Helmholtz cage. The supplied H-field is recorded by measuring the voltage generated across a $50 \pm 0.05\text{ m}\Omega$ sense resistor. A general-purpose LM741 operation amplifier multiplies the sense resistor output by a factor of ten. Output from the op-amp is digitized by the Picoscope 2205A 8-bit ADC which is set to record a measurement once every $655.36\text{ }\mu\text{s}$. In many applications, the ADC cannot communicate with the computer fast enough to enable real-time measurements. Instead, the Picoscope has onboard memory with a maximum capacity of 8000 measurements. This means that the Picoscope can measure about 5.25 cycles at 1 Hz before stopping to send the dataset to the computer. The Picoscope has two inputs which it can measure nearly simultaneously and has a triggering feature which ensures that all measurements begin at the same phase of the H-field cycle.

The sense coil output is multiplied by a dual INA2126 low-noise operational amplifier. The first amplification is a factor of 100, the second is a factor of two. After this amplification, the typical signal from the tested hysteresis rod given $\pm 100\text{ A/m}$ is about $\pm 200\text{mV}$. This output is fed to the Picoscope ADC which digitizes the data for the computer.

7.3.3 Method

The hysteresis rod measurement method as performed by the user is outlined below.

- (1) Orient Helmholtz cage such that one axis is perpendicular to the local magnetic field.
- (2) Ensure no magnetic material is present in the Helmholtz cage test volume.
- (3) Place magnetometer in center of Helmholtz cage and align to Helmholtz cage axes (the magnetometer may be used to determine if the cage is properly aligned relative to the magnetic field).
- (4) Ensure that all axes of the Helmholtz cage are connected to the computer-controlled power supply outputs. The function generator / power amp output should not be connected yet.

- (5) Initialize the LabVIEW-based Helmholtz cage software. The software will now calculate the Helmholtz constant $HC_{\text{empirical}}$.
- (6) Set the software to nullify the magnetic field on all three axes. If the Helmholtz cage is properly aligned, zero current should be needed to nullify the field on one of the axes.
- (7) Detach the power supply output from the coil set which is perpendicular to the B-field. Attach this coil set to the function generator / power amp output.
- (8) Enable the function generator at 1 Hz and enable the power amp. The magnetometer output should show varying magnetic field along a single axis.
- (9) Remove the magnetometer from the Helmholtz cage. Place the sense coil perpendicular to the zero-current coil set.
- (10) Set the function generator to the desired output voltage, being careful not to exceed 1.5V.
- (11) Collect one 8000-point dataset.
- (12) Place the magnetic sample within the sense coil being careful not to disturb the orientation of the coil.
- (13) Collect one 8000-point dataset.
- (14) Repeat steps 10-13 until all samples have been measured at all desired H-field amplitudes.

The list above simply describes the physical process of collecting a hysteresis measurement; analysis occurs thereafter. Figure 7.11 shows a block diagram of the analysis process after collecting the digital dataset. The H-field is processed by removing the 10x amplification, converting to current, and multiplying by $HC_{\text{empirical}}$. The B-field measurement is more complicated. After removing the gain of the signal, the signal due to the changing flux of the local environment is removed by subtracting the measurement without the sample from the measurement with the sample. After background subtraction, the data is numerically integrated using the trapezoidal

method. The integrated voltage is converted to the B-field by application of Faraday's Law. The constant offset after numeric integration is determined by assuming that the hysteresis loop is symmetric; the constant offset of a linear fit to the hysteresis loop is removed from all B-field values.

The data is bifurcated into two distinct groups: data from the top/bottom hysteresis curves. The data grouping is based on whether the B-field is rising or falling (the sine wave behavior makes this easy to characterize). Further processing is possible due to the length of the each dataset (over five H-field cycles) and the digitization of the H-field measurements. The H-field resolution of the 8-bit ADC varies from 0.105 A/m (± 10 A/m) to 1.04 A/m (± 100 A/m). Thus, the measurements are further grouped into various B-field values at each unique H-field value and curve. The B-field average and standard deviation are calculated for the unique H-field values of each curve. This provides an estimate of the B-field uncertainty and is useful for further data processing.

7.3.4 Results

Two types of hysteresis measurements are performed: isolation and system testing. Isolation testing measures hysteresis rod performance with no other magnetic materials nearby. System testing measures the hysteresis rod performance with a bar magnet and other hysteresis materials distributed identically to the flight satellite. In each measurement type, a nonlinear weighted least-squares fit is used to determine the hysteresis parameters (H_c , B_r , and B_s) which best match the data. The data are weighted by the inverse of their uncertainty and the Flatley [26] hysteresis model (described in Section 8.1.6.5) is used for the fit. Figure 7.12 shows one such ± 100 A/m cycle amplitude dataset with its parameter-based fitted hysteresis loop.

Each experimental dataset was independently measured two times and independently fit to the Flatley model to develop an uncertainty estimate of the mean fit parameters. A small sample size correction to the standard error of the mean is applied to each of the uncertainty estimates [33]. Five CSSWE flight spare hysteresis rods were available for measurement. The measurement is performed after the launch of CSSWE and uses hysteresis rods from the same raw material order

Figure 7.11: Hysteresis measurement analysis block diagram. The top (A.) and bottom (B.) rectangles highlight the analysis procedure for the H-field and B-field, respectively. Multiplication is denoted by a triangle, subtraction is denoted by a rounded rectangle, and integration is denoted by a diamond.

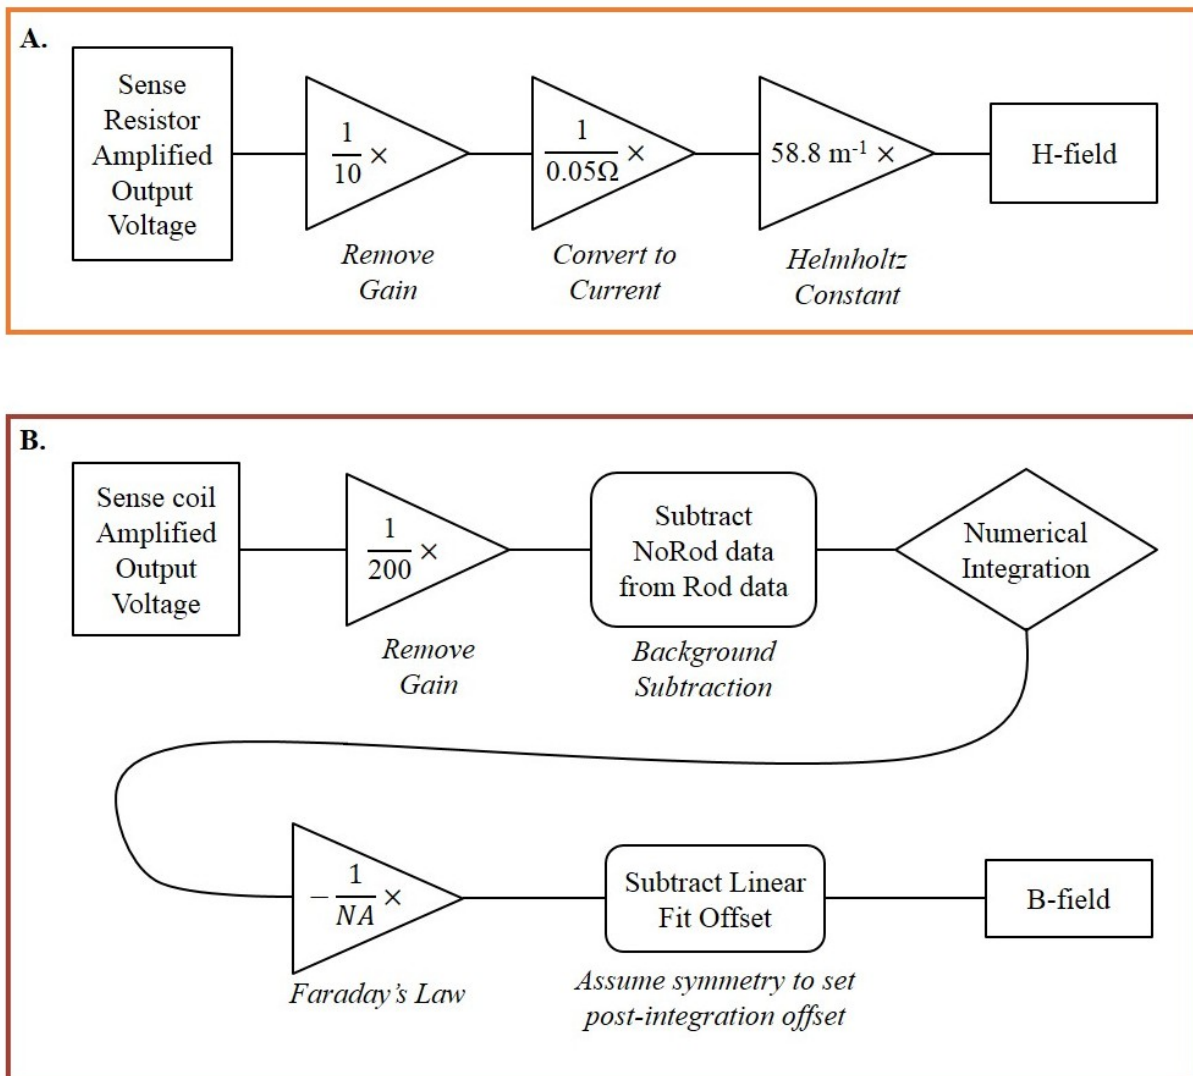



Figure 7.12: The $\pm 100\text{A/m}$ cycle amplitude measurement dataset of rod B with its fitted hysteresis loop. Error bars are included for each averaged flux density measurement.



and heat treatment set as the flight rods used on orbit.


7.3.4.1 Isolated Measurement

Each isolated rod measurement was fitted separately at each H-field cycle amplitude. The area of each simulated fit loop for each measurement is shown in Figure 7.13. The area of each simulated fit loop is calculated using an H-field cycle amplitude of ± 20 A/m. Note that the product of the hysteresis loop area and the hysteresis rod volume is the energy loss per H-field cycle. Thus, the dampening ability of each rod varies significantly. Surprisingly, the simulated ± 20 A/m loop area of each fit also varies significantly depending on the H-field cycle amplitude during the empirical measurement.

If the hysteresis model behaved perfectly, fit parameters based on a variety of empirical H-field cycle amplitudes would yield identical simulated cycle loop areas at the same simulated cycle amplitude. Figures 7.15 and 7.16 compare the hysteresis loops generated from a variety of simulated cycle amplitudes given hysteresis parameters fit data measured from ± 100 and ± 10 A/m empirical cycle amplitudes, respectively. The simulated loop areas for each fit at identical simulated cycle amplitudes are quite different. The fitted parameters best represent the true hysteresis loop when the measurement magnetizing field cycle amplitude is close to the simulated cycle amplitude. Thus, to achieve accurate simulation results, the measurement cycle amplitude should equal the range of magnetizing field which the true hysteresis rod is expected to experience the most frequently during dampening.

The extremes of the magnetization cycle amplitude are bounded by the magnitude of the on-orbit H-field but the extremes of the H-field component parallel to each rod will vary as the satellite attitude settles. Figure 7.14 shows the magnetizing field parallel to the hysteresis rods mounted along the $^B X$ and $^B Y$ axes. As shown, CSSWE experienced magnetizing field cycle amplitudes from ± 40 A/m to less than ± 5 A/m within the first seven days after launch. Note that smaller cycles due non-flat-spin motion occur but are difficult to see in this plot. For this reason, the fitted hysteresis parameters used by the simulation are based on measurement cycle amplitudes of ± 20

Figure 7.13: The simulated ± 20 A/m hysteresis loop area for hysteresis parameters fitted to a variety of H-field cycle amplitude datasets gathered for each of five measurement hysteresis rods. Each measurement was recorded with no other magnetic materials nearby.



and ± 10 A/m; these values are chosen in an attempt to use the most frequent cycle amplitudes experienced by each hysteresis rod during the settling period.

Figures 7.15 and 7.16 also show that the ± 100 A/m model degrades at especially small cycle amplitudes: each cycle is no longer a closed loop. This behavior is unrealistic and undesired. The parameters from the ± 10 A/m cycle perform much better at these low magnetizing field levels. Correctly representing these small cycle amplitudes is important because these loops are regularly encountered due to non-flat-spin motion throughout dampening. This is another reason to use a low cycle amplitude during measurement.

7.3.4.2 System-based Relative Position Measurement

The layout of PMAC components within a satellite may impact the dampening ability of each hysteresis rod. The net effect of the layout may be measured by placing magnetic components in their relative satellite positions and measuring individual hysteresis rod performance. Rods C, D, E, and F were placed at the X2, X3, Y2, and Y3 hysteresis rod positions within a CSSWE flight structure mock-up (see Figure 5.2). A flight spare bar magnet was also added to the mock-up in the flight position. The sense coil was then used to measure the performance of rod A at each of the hysteresis rod positions; the native rod at each position was removed during the measurement. Figure 7.17 shows the full setup for this measurement while Figure 7.18 shows the simulated area of the fitted hysteresis parameters.

There is some variation but the difference is only significant for the ± 100 A/m cycle amplitudes case. The hysteresis rod positions closest to the bar magnet (X3 and Y3) could be expected to possess decreased dampening abilities due to their increased H-field offset. The data shows that, for the CSSWE layout, interactions between rods are more important than the offset due to the bar magnet. Figure 7.19 shows the calculated H-field parallel to each rod position due to an $0.80 \text{ A}\cdot\text{m}^2$ dipole aligned with $+^B Z$ at the bar magnet location. Although the rods each experience a magnetizing field variation of up to 4.5 A/m along their length, the average H-field offset is 1-2 A/m.

Figure 7.14: The magnetizing field experienced by each of the body-frame axes with parallel hysteresis rods (${}^B X$ and ${}^B Y$) over the first week on-orbit. Each hysteresis rod experiences a wide variety of H-field cycle amplitudes.




Figure 7.15: A variety of simulated hysteresis loops generated using parameters fitted to measured output from an empirical cycle amplitude of ± 100 A/m. Each H-field amplitude is used to simulate 10 cycles of data. The top plot shows the simulated performance at bounds of ± 10 A/m; the bottom plot zooms to show the same data at a range of ± 2 A/m.




Figure 7.16: A variety of simulated hysteresis loops generated using parameters fitted to measured output from an empirical cycle amplitude of ± 10 A/m. Each H-field amplitude is used to simulate 10 cycles of data. The top plot shows the simulated performance at bounds of ± 10 A/m; the bottom plot zooms to show the same data at a range of ± 2 A/m.

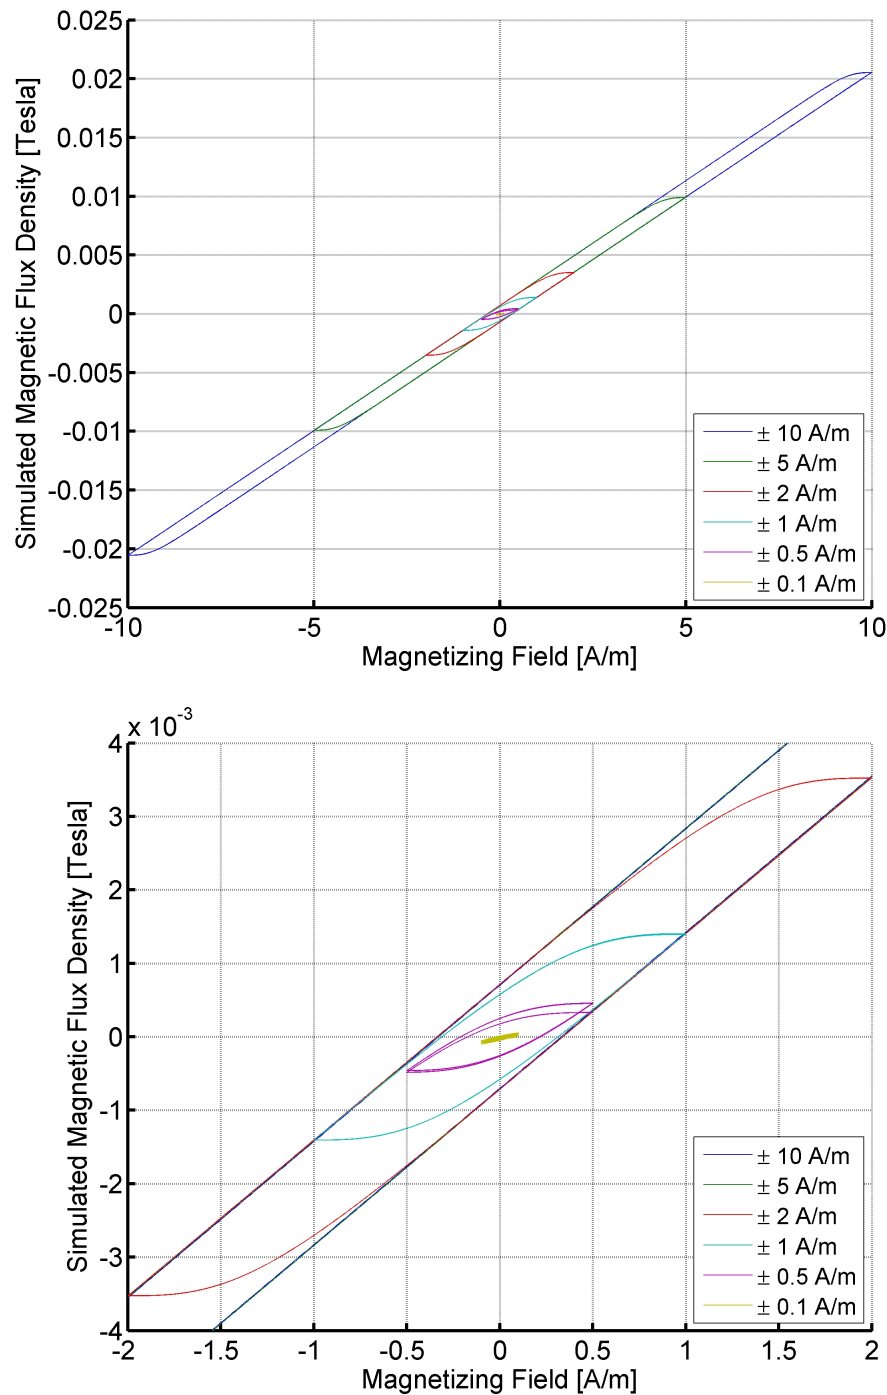


Figure 7.17: The CSSWE flight structure mock-up with hysteresis rods and bar magnet attached is used to measure the effect of other magnetic sources on a single hysteresis rod. The sense coil is in the X3 position (as defined by Figure 5.2).




Figure 7.18: The simulated hysteresis loop area using hysteresis parameters fitted to Rod A measurements performed at a variety of cycle amplitudes. The measurements are collected at each of the X2, X3, Y2, and Y3 hysteresis positions labeled in Figure 5.2. The Rod A Isolated Testing results are repeated here for comparison.




Figure 7.19: The magnetizing field offset parallel to each hysteresis rod position due to an $0.80 \text{ A}\cdot\text{m}^2$ dipole aligned with $+^B Z$ at the bar magnet location. The value of the mean H-field parallel to each rod is shown using open circles.




Table 7.2: HyMu-80 Hysteresis Parameters. The area was calculated using the Flatley hysteresis model [26] with a cycle amplitude of ± 20 A/m.

Hysteresis Parameter	Closed Magnetic Circuit Value (Material Datasheet)	Open Magnetic Circuit Value (Fitted to Measurement)
H_c (A/m)	1.59	0.3381
B_r (Tesla)	0.35	6.0618×10^{-4}
B_s (Tesla)	0.73	0.3000
Loop Area ($\text{J} \cdot \text{m}^{-3}$)	4.312	0.0448

7.3.4.3 Best-Fit Hysteresis Parameters

The individual testing shows that the hysteresis rod dampening ability varies significantly for each rod. However, the system testing shows that the position of a given rod within the CSSWE layout does not have a significant effect on the dampening ability of the rod at cycle amplitudes less than ± 50 A/m. Also, the hysteresis model has been found to work best when simulating data closest to its measurement cycle amplitude. To simplify the simulation, one set of parameters is used to model all six hysteresis rods within CSSWE. With all of the above considerations in mind, the best-fit parameters are found by fitting the individual measurement data from all five hysteresis rods over both tests at magnetizing field cycle amplitudes of both ± 10 and ± 20 A/m.

The fitted hysteresis loop parameters are collected in Table 7.2 and compared to the closed magnetic circuit values listed for the rod material. Note that the datasheet-based hysteresis parameters yield a loop area nearly 100 times greater than the measurement-based values. This difference in hysteresis dampening ability will have the profound effect of changing the simulated settling time by the same factor.

Chapter 8

Simulation

Inaccurate prediction of Passive Magnetic Attitude Control (PMAC) performance has hindered the use of such systems for some satellite missions. The purpose of the software work is to develop a numerical simulation which accurately describes the response of a satellite using a PMAC system. Such a simulation could be used to predict the settling time of a satellite using a PMAC system. The settling time is considered the most important system characteristic as it allows for mission planning using predictive simulation. To this end, a simulation is developed and defined below in terms of its major components; this will allow future mission teams to use it for predictive mission requirements verification.

The simulation is built within the MATLAB[®] environment, which is used to numerically integrate the equations of motion. Models are defined for each environmental torque at low earth orbit. An orbit propagation method is defined because many environmental torques are dependent on satellite position. With the simulation components fully defined, analysis is performed to answer basic questions about the expectations of the simulation. Finally, the simulation output is shown given nominal input; the output is compared to the on-orbit data analyzed in Chapter 6.

We seek to understand the PMAC system behavior where possible, but here we do not aim to fully describe the underlying dynamics. Instead, the PMAC simulation development and testing given select initial inputs are described. Whenever possible, simulation results which shed light on the underlying dynamics are discussed. However, the scope of this research is limited to the simulation development and validation. Further testing given a wider variety of initial input may

be helpful in a deeper understanding of the general dynamics of a PMAC satellite.

8.1 Components

The attitude simulation is built from many components. Each component is defined in the subsections below. The reader should be able to recreate the simulation using the information described in this section.

8.1.1 Frames

There are three frames, or coordinate systems, used in the PMAC simulation. The first of these is the Earth Centered Inertial (ECI) frame. The ECI frame $^{\mathcal{I}}\{X\ Y\ Z\}$ has its origin at the center of earth, with direction vectors defined as follows: $^{\mathcal{I}}X$ lies in earth's equatorial plane and is parallel to the vernal equinox direction, $^{\mathcal{I}}Z$ is parallel to the rotation axis of earth, and $^{\mathcal{I}}Y$ is defined by the right-hand rule.

The second frame used in the PMAC simulation is the body frame $^{\mathcal{B}}\{X\ Y\ Z\}$ of the spacecraft. The body frame is chosen to align with the principle inertia axes of the spacecraft. The body frame has its origin at the satellite center of mass with $^{\mathcal{B}}X$ parallel to the REPTile collimator, $^{\mathcal{B}}Z$ parallel to the satellite antenna, and $^{\mathcal{B}}Y$ defined by the right-hand rule; Figure 4.3 shows the spacecraft body frame. The 3×3 matrix $[R]$ is defined to rotate an arbitrary vector \mathbf{v} from the inertial frame to the body frame:

$$^{\mathcal{B}}\mathbf{v} = [R]^{\mathcal{I}}\mathbf{v}. \quad (8.1)$$

This rotation matrix will change at each step of the simulation as rotational dynamics change the attitude of the spacecraft relative to the inertial frame. More detail on the inertial to body frame rotation matrix is given in Section 8.1.2.

The final frame is the Earth Centered Earth Fixed (ECEF) frame which rotates with Earth. The ECEF frame $^{\mathcal{E}}\{X\ Y\ Z\}$ has its origin at the center of earth with direction vectors defined as follows: $^{\mathcal{E}}X$ lies in earth's equatorial plane and is parallel to the prime meridian (0°N, 0°E), $^{\mathcal{E}}Z$ is

parallel to the rotation axis of earth (90°N), and ${}^\mathcal{E}Y$ is defined by the right hand rule (0°N , 90°E). The conversion between the ECI and ECEF coordinate frames is not dependent on the satellite attitude; it is defined by the following rotation matrix:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{\mathcal{E}} = \begin{bmatrix} \cos(\lambda_0 + \omega_E t) & \sin(\lambda_0 + \omega_E t) & 0 \\ -\sin(\lambda_0 + \omega_E t) & \cos(\lambda_0 + \omega_E t) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{\mathcal{I}} \quad (8.2)$$

where λ_0 is the Greenwich Mean Sidereal Time (GMST) when the simulation starts, ω_E is the sidereal rotation rate of earth (rotation rate relative to fixed stars) and t is the time since simulation start. Figure 8.1 shows both the ECI and ECEF coordinate frames. The ECEF frame is useful because many inertial models are given in this frame.


8.1.2 Attitude Parameters

The rotation matrix is useful for converting vectors from one frame to another. Also, it is simple to work with; the reverse rotation matrix (body to inertial) is found by simply transposing the original rotation matrix: $[R] = [R]_{\text{BI}} = [R]_{\text{IB}}^T$. However, the rotation matrix is a nine-dimensional representation of a three-dimensional rotation; six elements are redundant [65]. These redundant elements can lead to inaccurate modeling as rounding errors soften the constraints.

Three dimensional attitude parameters (such as Euler Angles) always contain a singularity in their kinematic differential equation at some specific attitude, making them undesirable for numerical integration purposes. A good compromise is the quaternion, a four-dimensional attitude parameterization with a single redundant parameter. The relation between the quaternion and the rotation matrix is given by [65]:

$$[R]_{\text{BI}} = \begin{bmatrix} q_0^2 + q_1^2 + q_2^2 + q_3^2 & 2(q_1q_2 + q_0q_3) & 2(q_1q_3 - q_0q_2) \\ 2(q_1q_2 - q_0q_3) & q_0^2 - q_1^2 + q_2^2 - q_3^2 & 2(q_2q_3 + q_0q_1) \\ 2(q_1q_3 + q_0q_2) & 2(q_2q_3 - q_0q_1) & q_0^2 - q_1^2 - q_2^2 + q_3^2 \end{bmatrix} \quad (8.3)$$

Figure 8.1: The ECEF and ECI coordinate frames are shown.



where q_0 is the scalar part of the quaternion and $[q_1 \ q_2 \ q_3]^T$ is the vector part. The quaternion has a single constraint: $q_0^2 + q_1^2 + q_2^2 + q_3^2 = 1$. This constraint is easily satisfied by normalization at a certain interval. The simulation developed in this chapter re-normalizes the quaternion every 100 integration steps; this frequency is chosen as a balance between computational performance and error tolerance.

8.1.3 Equations of Motion

The core of the simulation is Euler's rotational equation of motion (Equation 2.1). However, this equation cannot be integrated alone; it must be combined with a kinematic differential equation which defines the relationship between the angular velocity and the rate of change of the attitude parameters. As such, the kinematic differential equation is dependent on the parameter set used to represent the attitude. The kinematic differential equation for quaternions is given as [65]:

$$\begin{bmatrix} \dot{q}_0 \\ \dot{q}_1 \\ \dot{q}_2 \\ \dot{q}_3 \end{bmatrix} = \begin{bmatrix} q_0 & -q_1 & -q_2 & -q_3 \\ q_1 & q_0 & -q_3 & q_2 \\ q_2 & q_3 & q_0 & -q_1 \\ q_3 & -q_2 & q_1 & q_0 \end{bmatrix} \begin{bmatrix} 0 \\ \omega_1 \\ \omega_2 \\ \omega_3 \end{bmatrix} \quad (8.4)$$

where the scalar-first quaternion convention is used.

8.1.4 Orbit Propagation

The PMAC simulation assumes that the satellite rotation and translation are not coupled; thus the attitude has no effect on the orbit of the spacecraft. This assumption results in inertial satellite position and velocity vectors which are directly related to the orbit elements at some epoch and time relative to this epoch. These inertial position and velocity vectors are used as inputs for the inertial models described in Section 8.1.5.

The CSSWE orbit mean elements at a given epoch are provided by a Two-Line Element (TLE) set. This TLE is provided by the Joint Space Operations Center (JSpOC) for use with

the CSSWE mission. The orbit elements contained within the TLE make specific assumptions about the orbit and are designed to be propagated using only select orbit propagation schemes [35]. The PMAC simulation uses an updated version of the SGP4 propagator [79] designed for TLE propagation; the same method is used to calculate the CSSWE position for real-time on-orbit operations.

The first usable TLE has an epoch over ten days after launch, as shown in Figure 4.8. The datasets shown later in this chapter simulate the first ten days on orbit and thus use one TLE over this time period. Although the TLE is most accurate close to the epoch, no other position estimate is available for the early mission and on-orbit telemetry has shown the TLE-based position to be accurate to at least ± 6 seconds (Figure 4.8). Using one TLE throughout the simulation has the advantage of avoiding discontinuities in satellite position and velocity which would be generated during the switch from one TLE to the next.

8.1.5 Inertial Vector Models

The simulation uses two models to generate the inertial vectors for both magnetic field and sun position. The magnetic field model is solely dependent on satellite position in the ECEF frame, while the sun position model is solely dependent on date. Both are defined below.

8.1.5.1 International Geomagnetic Reference Field Model

The PMAC simulation uses the eleventh generation International Geomagnetic Reference Field (IGRF-11) model [23] to generate the inertial magnetic field vector. The IGRF is based on empirical measurements and is widely used as a magnetic field model for many applications. The IGRF represents the magnetic flux density as the negative gradient of a scalar potential function ($\mathbf{B} = -\nabla V$) which is defined by a spherical harmonics series which is set by a maximum of 195 coefficients (order 13). These coefficients are updated every five years; the latest epoch is 2010. The model also predicts some coefficients for the next five years after the latest epoch at one year intervals; 80 coefficients are predicted (order 8). The model sets coefficients for times between the

prediction year epochs using linear interpolation.

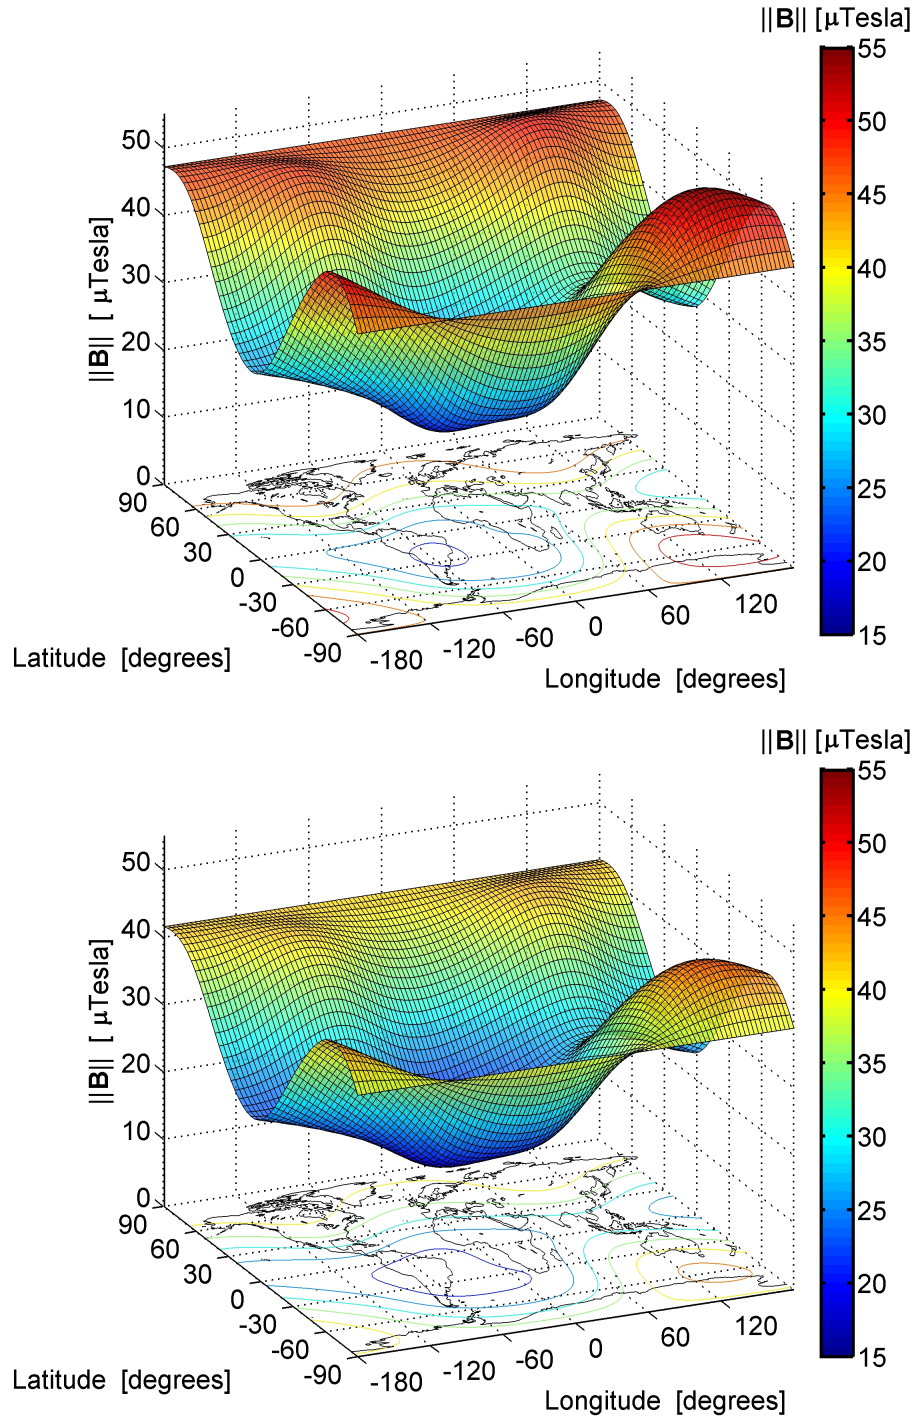
The IGRF is designed to provide an estimate of the magnetic field at the earth surface and above due to magnetic materials below the surface; it does not take into account variation due to space-based activity such as solar rotation, ionospheric currents, and geomagnetic storms. However, it does account for internal variations in the magnetic flux density which take place over a timespan of months to decades; this “secular variation” accounts for a global surface-level root mean squared magnitude change of about 80nT per year [6]. The IGRF-11 model predictions (for years after the 2010 epoch) of the core-based magnetic field are estimated to possess errors of up to 20nT per year [23]. Using the IGRF to model the inertial magnetic field experienced by a satellite can lead to higher errors, as discussed in Section 4.3.2.3.

Figure 8.2 shows the global magnetic flux density magnitude at altitudes of 450km and 770km (the extremes of the CSSWE orbit) as generated using the IGRF model with the 2012 predictive coefficients. This figure shows the range of magnetic flux densities which are used in the attitude simulation at various positions on-orbit. The CSSWE orbit experiences magnetic flux density magnitudes ranging from $17.6\mu\text{T}$ to $52.3\mu\text{T}$.

8.1.5.2 Inertial Sun Vector Model

The inertial sun position model is calculated using the method defined by Vallado [77]. A simplified version of the method which calculates the ECI frame unit vector from the center of earth to the sun $^{\mathcal{I}}\hat{\mathbf{s}}$ is shown below:

Figure 8.2: Earth global magnetic flux density magnitude at altitudes of 450km (top) and 770km (bottom) as calculated using the eleventh-generation International Geomagnetic Reference Field (IGRF-11) model with the 2012 predictive coefficients.



$$T_{\text{UT1}} = \frac{JD_{\text{UT1}} - 2,451,545.0}{36,525} \quad (8.5)$$

$$\lambda_{M_{\text{sun}}} = 280.4606184 + 36,000.77005361T_{\text{UT1}} \quad (8.6)$$

$$M_{\text{sun}} = 357.5277233 + 35,999.05034T_{\text{UT1}} \quad (8.7)$$

$$\lambda_{\text{ecliptic}} = \lambda_{M_{\text{sun}}} + 1.914666471 \sin(M_{\text{sun}}) + 0.918994643 \sin(2M_{\text{sun}}) \quad (8.8)$$

$$\varepsilon = 23.439291 - 0.0130042T_{\text{UT1}} \quad (8.9)$$

$$\mathcal{I}_{\hat{\mathbf{s}}} = \begin{bmatrix} \cos \lambda_{\text{ecliptic}} \\ \cos \varepsilon \sin \lambda_{\text{ecliptic}} \\ \sin \varepsilon \sin \lambda_{\text{ecliptic}} \end{bmatrix} \quad (8.10)$$

where JD_{UT1} is the Julian Date, T_{UT1} is the number of centuries since the epoch, $\lambda_{M_{\text{sun}}}$ is the mean longitude of the Sun, M_{sun} is the mean anomaly of the Sun, $\lambda_{\text{ecliptic}}$ is the ecliptic longitude of the Sun, and ε is the obliquity of the ecliptic. All angles ($\lambda_{M_{\text{sun}}}$, M_{sun} , $\lambda_{\text{ecliptic}}$, and ε) are in units of degrees. This method is valid from 1950 to 2050 and is accurate to 0.01° [77]. It is not necessary to convert the unit vector from earth to the sun $\mathcal{I}_{\hat{\mathbf{s}}}$ to the unit vector from the satellite to the sun; for a satellite at 1000km altitude, the angular difference between the center of the earth and the satellite position when perpendicular to the earth-sun vector (maximum error) is 0.0024° which is less than the 0.01° accuracy of the model.

It is useful to denote times at which the satellite is in eclipse. The simulation uses the method described by Kelso [41], outlined below. First, define the angular radii of the earth and sun as seen by the satellite:

$$\begin{aligned} \theta_E &= \sin^{-1} \left(\frac{R_E}{\rho_E} \right) \\ \theta_S &= \sin^{-1} \left(\frac{R_S}{\rho_S} \right) \end{aligned} \quad (8.11)$$

where R_E and R_S are the true radii of the earth and sun, respectively and ρ_E and ρ_S are the distances from the satellite to the earth and sun, respectively. Next, the angle from the center of

the earth and the center of the sun (as seen by the satellite) is calculated as:

$$\theta_{\text{ES}} = \cos^{-1}({}^{\mathcal{B}}\hat{\boldsymbol{\rho}}_{\text{E}} \cdot {}^{\mathcal{B}}\hat{\boldsymbol{\rho}}_{\text{S}}) \quad (8.12)$$

where ${}^{\mathcal{B}}\hat{\boldsymbol{\rho}}_{\text{E}}$ is the body-frame unit vector from the satellite to earth and ${}^{\mathcal{B}}\hat{\boldsymbol{\rho}}_{\text{S}}$ is the body-frame unit vector from the satellite to the sun. The earth-sun angle is used to determine when the sun is fully or partially eclipsed by the earth as shown below. Note that these equalities are only true when the earth appears larger than the sun (true for all satellite orbits within the moon's orbit):

$$\text{full eclipse (umbral):} \quad \theta_{\text{ES}} < \theta_{\text{E}} - \theta_{\text{S}} \quad (8.13)$$

$$\text{partial eclipse (penumbral):} \quad |\theta_{\text{E}} - \theta_{\text{S}}| < \theta_{\text{ES}} < \theta_{\text{E}} + \theta_{\text{S}}.$$

8.1.6 External Torque Estimation

A PMAC system relies upon two external torques for control: bar magnet and hysteresis torque. Other external torques are present due to the interaction of the spacecraft and the local environment: gravity gradient, aerodynamic (drag), solar pressure, magnetic residual, and eddy current. Each of these torques is described and modeled in the body frame for inclusion in Equation 2.1. The total external torque is simply:

$$\mathbf{L} = \mathbf{L}_{\text{B}} + \mathbf{L}_{\text{H}} + \mathbf{L}_{\text{G}} + \mathbf{L}_{\text{D}} + \mathbf{L}_{\text{SP}} + \mathbf{L}_{\text{R}} + \mathbf{L}_{\text{EC}} \quad (8.14)$$

8.1.6.1 Bar Magnet Torque

The bar magnet torque vector is given by Equation 2.6, repeated below.

$$\mathbf{L}_{\text{B}} = \mathbf{m}_{\text{bar}} \times \mathbf{B} \quad (2.6)$$

Note that there may be a large difference in the manufacturer-quoted value of bar magnet magnetic moment versus the true magnetic moment of the bar magnet \mathbf{m}_{bar} (see Section 7.2). The

local magnetic flux density \mathbf{B} is given by the IGRF model (see Section 8.1.5.1).

8.1.6.2 Hysteresis Torque

The hysteresis torque is the most difficult torque to model in the PMAC simulation due to the non-linear relationship between the local magnetizing field due to earth and the induced magnetization within the rod. Many models exist to attempt to predict this relationship; three such models are investigated in this research. As with the bar magnet torque, the IGRF model is used to calculate the local magnetic flux density vector \mathbf{B} . The component of the magnetizing field $\mathbf{H} = \mathbf{B}/\mu_0$ parallel to each hysteresis rod at the current time step is used as the applied field input to the hysteresis model. The component of the vector derivative of the magnetizing field as seen in the body frame $\frac{\mathcal{B}_d}{dt}H$ is used as an input for each of the investigated hysteresis models. The component is calculated as $\frac{\mathcal{B}_d}{dt}H = \left(\frac{\mathcal{B}_d}{dt}\mathbf{H}\right) \cdot \hat{\mathbf{n}}_{\text{rod}}$ where $\hat{\mathbf{n}}_{\text{rod}}$ is a unit vector describing the orientation of the hysteresis rod.

Each model seeks to represent the average induced magnetic flux density parallel to the rod. Equation 2.7 converts the average parallel magnetizing field within the rod into the magnetic moment parallel to the hysteresis rod m_{hyst} at the current time step. The vector magnetic moment is calculated using $\mathbf{m}_{\text{hyst}} = m_{\text{hyst}}\hat{\mathbf{n}}_{\text{rod}}$. The torque provided by the hysteresis rods is then given by the magnetic torque equation:

$$\mathbf{L}_H = \mathbf{m}_{\text{hyst}} \times \mathbf{B} \quad (8.15)$$

8.1.6.3 Parallelogram Model

The parallelogram hysteresis model is the simplest and easiest model to implement [47] [61]. It is defined by a parallelogram that: intersects the y-axis at $\pm B_r$, intersects the x-axis at $\pm H_c$, and has a maximum absolute B-field output of B_s . The top curve is used when the magnetizing field derivative $\frac{\mathcal{B}_d}{dt}H < 0$, while the bottom curve is used when $\frac{\mathcal{B}_d}{dt}H \geq 0$. Thus, the output B is dependent only on H , $\frac{\mathcal{B}_d}{dt}H$, H_c , B_r , and B_s and does not depend on cycle magnitude or frequency

or the previous magnetic flux density within the rod. The parallelogram model output is shown in blue in Figure 8.3.

8.1.6.4 Inverse Tangent Model

The inverse tangent hysteresis model [26] approximates the bounds of a typical hysteresis curve using the inverse tangent function. The model is defined as:

$$\begin{aligned}
 k &= \frac{1}{H_c} \tan \left(\frac{\pi B_r}{2B_s} \right) \\
 \text{if } \frac{B_d}{dt} H &\geq 0 : \\
 B &= \left(\frac{2B_s}{\pi} \right) \tan^{-1} (k(H - H_c)) \\
 \text{if } \frac{B_d}{dt} H &< 0 : \\
 B &= \left(\frac{2B_s}{\pi} \right) \tan^{-1} (k(H + H_c)).
 \end{aligned} \tag{8.16}$$

Again, the output B is dependent only on H , $\frac{B_d}{dt}H$, H_c , B_r , and B_s and does not depend on the magnetization cycle magnitude or frequency or the previous magnetic flux density within the rod. The inverse tangent model output is shown in green in Figure 8.3.

8.1.6.5 Flatley Model

The Flatley hysteresis model [26] is substantially different than the previous models, as it is defined in terms of a differential equation. While this adds complexity, the resultant loop is much more realistic as it can model minor hysteresis loops within the full loop as the hysteresis experiences lower cycle amplitudes. The Flatley model is defined as:

$$\begin{aligned}
k &= \frac{1}{H_c} \tan\left(\frac{\pi B_r}{2B_s}\right) \\
\text{if } \frac{B_d}{dt} H \geq 0 : \\
\dot{B} &= \left(q_0 + (1 - q_0) \left[\frac{1}{2H_c} \left(H - \frac{1}{k} \tan\left(\frac{\pi B}{2B_s}\right) + H_c \right) \right]^p \right) \left(\frac{2kB_s}{\pi} \right) \cos^2\left(\frac{\pi B}{2B_s}\right) \left(\frac{B_d}{dt} H \right)
\end{aligned} \tag{8.17}$$


$$\begin{aligned}
\text{if } \frac{B_d}{dt} H < 0 : \\
\dot{B} &= \left(q_0 + (1 - q_0) \left[\frac{1}{2H_c} \left(H - \frac{1}{k} \tan\left(\frac{\pi B}{2B_s}\right) - H_c \right) \right]^p \right) \left(\frac{2kB_s}{\pi} \right) \cos^2\left(\frac{\pi B}{2B_s}\right) \left(\frac{B_d}{dt} H \right)
\end{aligned}$$

where q_0 and p are selectable constants which are tailored to fit a given empirical dataset. The model is found to perform best (most realistic over the wide range of magnetizing field cycle amplitudes) when $q_0 = 0$ and $p = 2$; these values have been used by other groups as well [58] [12]. The magnetic flux density time derivative \dot{B} is thus dependent on B , H , $\frac{B_d}{dt}H$, H_c , B_r , and B_s . In practice, the Flatley model allows for the hysteresis rod magnetic flux densities to be integrated simultaneously with the attitude parameters and the angular velocity of the satellite. This means that the flux density is also dependent on the choice of numeric integrator, the simulation time step, and the hysteresis rod flux density of the previous integration step. Figure 8.3 shows a comparison of the output of the parallelogram, inverse tangent, and Flatley hysteresis models for two H-field cycles with amplitudes of ± 8 A/m and ± 3 A/m. The parallelogram and inverse tangent models are simply cutoff at lower cycle amplitudes, whereas the Flatley model actively adjusts to account for the changes in magnetizing field cycle amplitude.

The Flatley hysteresis model is chosen for this research because it was developed for numeric simulation of empirical datasets [26]. Also, the Flatley model has been used for a variety of past hysteresis dampening simulation [58], [43]. However, attempts to fit hysteresis measurement data using this model have shown its deficiencies (see Section 7.3).

The Flatley hysteresis model is numerically integrated and is not defined by two constant curves as the parallelogram and inverse tangent models are; as a result of this, some unique errors

Figure 8.3: Hysteresis loop output by three different hysteresis models with magnetizing field cycle amplitudes of ± 8 A/m (top) and ± 3 A/m (bottom). The parallelogram model (blue), inverse tangent (green), and Flatley (red) hysteresis models are shown. These loops were generated using closed magnetic circuit hysteresis parameters (see Table 7.2), two 1 Hz H -cycles, and a time step of 0.001s.



can occur. The Flatley model is designed to be bounded by the inverse tangent hysteresis model. Repeated simulation has shown that the Flatley-model-simulated hysteresis rod magnetic flux densities can sometimes exceed the bounds of the inverse tangent model for a time. In some extreme cases, the simulated rod magnetic flux density greatly exceeds the bounds of the inverse tangent model, corrupting the results of the PMAC simulation. This problem is resolved by checking the hysteresis rod magnetic flux density value after each simulation step. If the output exceeds the inverse tangent model bounds, the output is set to the top or bottom curve of the inverse tangent model, whichever is closer. This check prevents the erroneous behavior mentioned above.

8.1.6.6 Gravity Gradient Torque

The gravity gradient torque is given as follows [65]:

$${}^B\mathbf{L}_G = \left(\frac{3\mu_e}{R_c^5} \right) {}^B\mathbf{R}_c \times [I] {}^B\mathbf{R}_c \quad (8.18)$$

where μ_e is the earth gravitational parameter ($3.986 \cdot 10^5 \text{ km}^3\text{s}^{-2}$), ${}^B\mathbf{R}_c$ is the body-frame position vector from the center of earth to the center of mass of the spacecraft, and $[I]$ is, again, the spacecraft mass moment of inertia matrix. As shown, the gravity gradient torque is highest at low altitudes and is increased as the inertia matrix of the satellite becomes less symmetric. The torque acts to align the minimum inertia axis of a satellite with the nadir direction.

8.1.6.7 Aerodynamic Torque

The aerodynamic torque is calculated as follows [61]:

$${}^B\mathbf{L}_D = \frac{1}{2} \rho C_d ({}^B\mathbf{S} \cdot {}^B\mathbf{V}) ({}^B\mathbf{V} \times {}^B\mathbf{r}_d) \quad (8.19)$$

where ρ is the density of the medium in which the spacecraft is traveling, C_d is the spacecraft coefficient of drag (assumed to be 2.4 for CSSWE), ${}^B\mathbf{S}$ is a vector defining the surface area of each of the body axes of the spacecraft ($[0.03 \ 0.03 \ 0.01]^T \text{ m}^2$ for a 3U CubeSat), ${}^B\mathbf{V}$ is the spacecraft

velocity vector in the body frame, and ${}^B\mathbf{r}_d$ is the body-frame position vector from the satellite center of mass to the geometric center (can be calculated using a solid model). The density ρ is the most difficult to calculate, as it can vary greatly depending on solar input and spacecraft altitude. The NRLMSISE-00 empirical model [59] is used by the simulation to estimate the atmospheric density for the position of the satellite at each time step. Daily and 81-day average F10.7 indices, as well as the average 3-hour ap index for the last 48 hours are all inputs to the NRLMSISE-00 model. The following assumed inputs are used: A daily F10.7 flux of $128.7 \cdot 10^{-22} \text{J} \cdot \text{s}^{-1} \text{m}^{-2} \text{Hz}^{-1}$, an 81-day average F10.7 flux of $168.5 \cdot 10^{-22} \text{J} \cdot \text{s}^{-1} \text{m}^{-2} \text{Hz}^{-1}$, and an average 3-hour ap index of 48. These inputs are defined in this way to simulate moderate solar activity.

8.1.6.8 Solar Pressure Torque

When the satellite is insolated (see Section 8.1.5.2), the force vector due to radiative pressure may be defined in the body frame as [78]:

$${}^B\mathbf{F}_{\text{SP}} = -P_S c_R \text{diag}({}^B\mathbf{S}) {}^B\hat{\mathbf{s}} \quad (8.20)$$

where P_S is the solar radiation pressure at earth (set to $4.5 \cdot 10^{-6} \text{ Pa}$ [36]), c_R is the coefficient of reflectivity of the satellite (assumed as 0.8), and ${}^B\hat{\mathbf{s}}$ is the body-frame unit vector from the earth to the sun. The torque vector due to solar pressure is thus:

$${}^B\mathbf{L}_{\text{SP}} = {}^B\mathbf{r}_d \times {}^B\mathbf{F}_{\text{SP}}. \quad (8.21)$$

8.1.6.9 Magnetic Residual Torque

The magnetic residual torque is due to undesired magnetism that may be present in the satellite. These magnetic residuals may be due to current loops, hard magnetic material in the satellite, or a misalignment of the PMAC bar magnet. The PMAC simulation models the residual magnetism as a constant magnetic moment vector in the body frame \mathbf{m}_{res} (set to $[-0.0059 \ -0.0083 \ 0.0004]^T \text{ A} \cdot \text{m}^2$), and the torque vector due to magnetic residual is calculated as:

$${}^{\mathcal{B}}\mathbf{L}_R = \mathbf{m}_{\text{res}} \times {}^{\mathcal{B}}\mathbf{B} \quad (8.22)$$

The chosen magnetic residual of CSSWE is based on a fit which minimizes the MEKF residuals (see Section 6.2.2); the residual is simply the difference between the fit and the assumed magnetic moment of the bar magnet \mathbf{m}_{bar} . However, future missions could determine the magnetic residual before flight using a method similar to the bar magnet measurement (see Section 7.2) if the satellite is measured in its flight configuration.

8.1.6.10 Eddy Current Torque

Eddy currents are generated when a conductor experiences a changing magnetic field. These circular currents induce their own magnetic field which can torque the satellite. The magnetic fields generated by eddy currents within rotating satellites are usually negligible compared to the local magnetic field; the simulation makes this assumption. Eddy currents can be generated in magnetic or non-magnetic conducting material but torques generated by currents within magnetic material are usually dwarfed by magnetization-based torques. The eddy current torque is partially based on area available for loops to form. Thus, it is assumed that the eddy currents generated by CSSWE are formed within the solid aluminum shell of the 3U CubeSat. The eddy current torque is given by [64]:

$$\mathbf{L}_{\text{EC}} = \sum_{i=1}^n (\mathbf{k}_i \cdot \hat{\mathbf{B}})(\boldsymbol{\omega} \times \mathbf{B}) \times \mathbf{B} \quad (8.23)$$

where there are n surface elements which generate eddy currents and \mathbf{k}_i is a vector constant set by a given surface element [30]:

$$\mathbf{k}_i = \frac{\pi}{4} \sigma_i r_i^3 A_i \hat{\mathbf{n}}_i \quad (8.24)$$

where σ_i is the conductivity of the surface element, r_i is the maximum circular loop radius that can form within the surface element, A_i is the area of the surface element, and $\hat{\mathbf{n}}_i$ is the unit vector

normal to the surface element. Note that currents will be generated in the same direction and add together instead of nullify (consider opposite sides of the same CubeSat); the sign of $\hat{\mathbf{n}}_i$ should reflect this behavior. The aluminum 3U solid shell of CSSWE was calculated to have $\mathbf{k}_1 = [147.3 \ 0 \ 0]^T$, $\mathbf{k}_2 = [0 \ 147.3 \ 0]^T$, and $\mathbf{k}_3 = [0 \ 0 \ 49.3]^T$, all in units of $\text{A}^2 \cdot \text{s}^3 \cdot \text{m}^2 \cdot \text{kg}^{-1}$.

8.1.7 Numeric Integrators

The choice of numeric integrator can have a great effect on the results of the simulation. In fact, much analysis is performed to select a numeric integrator and time step for the simulation (see Section 8.3). This subsection defines the numeric integrators which are considered for simulation use.

The most basic explicit numeric integrator is Euler's method; it uses the value and derivative at the current step to determine the value at the next step using a time step duration h :

$$\begin{aligned} y_{n+1} &= y_n + h f(t_n, y_n) \\ t_{n+1} &= t_n + h \end{aligned} \tag{8.25}$$

where $f(t, y)$ is the rate of change of y at (t, y) . Decreasing the time step will yield better results, but Euler's method is often not the best choice for numerical integration. It is only accurate to first order. Also, it is possible that a higher-order integrator could yield more accurate results at a lower time step, thus saving computing resources. There are many higher-order explicit integrators; this research focuses on explicit Runge-Kutta integrators of order two through seven and the Lie Group Variational Integrator.

8.1.7.1 Runge-Kutta Integrators

Explicit Runge-Kutta methods of order greater than one calculate the next value using the current value, the current derivative, and the derivative at one or more points between each integration step. Explicit Runge-Kutta integration may be generalized as follows [34]:

$$y_{n+1} = y_n + h \sum_{i=1}^s b_i k_i \quad (8.26)$$

where

$$\begin{aligned} k_1 &= f(t_n, y_n) \\ k_2 &= f(t_n + c_2 h, y_n + a_{21} k_1) \\ k_3 &= f(t_n + c_3 h, y_n + a_{31} k_1 + a_{32} k_2) \\ &\vdots \\ k_s &= f(t_n + c_s h, y_n + a_{s1} k_1 + a_{s2} k_2 + \dots + a_{s,s-1} k_{s-1}) \end{aligned}$$

where the coefficients a_{ij} , b_i , and c_i are given by the Butcher tableau of the integrator in use. Appendix B explains the Butcher tableau and defines coefficients for integrators from order two to order seven.

This research investigates the performance of fixed time step integrators only; this limitation allows for improved data processing. Data from each simulation is saved at a rate of 1 Hz regardless of the integration time step; this reduces the size of save files and allows for direct comparison between datasets generated over a range of time steps. Adaptive step size Runge-Kutta methods do not possess these advantages.

8.1.7.2 Lie Group Variational Integrator

Researchers at the University of Michigan [44] have recently developed the Lie Group Variational Integrator (LGVI); it is designed to model long-term, low disturbance torques. LGVI is based on a discrete variational approach, and thus preserves the momentum of the system between time steps [45]. The Lie group integrator, created by Iserless et al. [37], is the basis of the LGVI. The Lie group numeric integrator uses the following integration rule:

$$[R]_{n+1} = \exp \left(-\frac{h}{2} ([\boldsymbol{\omega} \times]_n - [\boldsymbol{\omega} \times]_{n+1}) \right) [R]_n. \quad (8.27)$$

where n is the current integration step, $n + 1$ is the next integration step, and h is the constant integration time step duration. Equation 8.27 gives the updated rotation matrix but it requires the new angular velocity $\boldsymbol{\omega}_{n+1}$ as an input. The new angular velocity could be calculated using a Runge-Kutta method to integrate Equation 2.1, but that could corrupt the momentum conservation which Equation 8.27 was developed to avoid. Lee et al. obtained the discrete Lie group variational numeric integrator rules by discretizing Hamilton's principle [45]:

$$\mathbf{\Pi}_{n+1} = [F]_n^T \mathbf{\Pi}_n + \frac{h}{2} [F]_n^T \mathbf{L}_n + \frac{h}{2} \mathbf{L}_{n+1} \quad (8.28)$$

$$h[\mathbf{\Pi} \times] + \frac{h^2}{2} [\mathbf{L} \times]_n = [F]_n [I_d] - [I_d] [F]_n^T \quad (8.29)$$

$$[R]_{n+1} = [R]_n [F]_n \quad (8.30)$$

where $\mathbf{\Pi} = [I]\boldsymbol{\omega}$ is the angular momentum vector, $[I_d]$ is a non-standard inertia matrix defined by $[I] = \text{tr}([I_d])[I_{3 \times 3}] - [I_d]$, and $[F]_n$ is the 3×3 rotation matrix defining the relative attitude between t_n and t_{n+1} . Thus, given the angular velocity and external torque vectors for the present time $(\boldsymbol{\omega}_n, \mathbf{L}_n)$, $[F]_n$ may be used to determine the state at the next step. However, in practice solving Equation 8.29 is not a trivial task because $[F]_n$ is a 3×3 matrix.

In order to solve for $[F]_n$, Lee introduces the 3×1 vector \mathbf{f}_n , defined by [45]:

$$\begin{aligned} [F]_n &= \exp[\mathbf{f} \times]_n \\ &= [I_{3 \times 3}] + \frac{\sin \|\mathbf{f}_n\|}{\|\mathbf{f}_n\|} [\mathbf{f} \times]_n + \frac{1 - \cos \|\mathbf{f}_n\|}{\|\mathbf{f}_n\|^2} [\mathbf{f} \times]_n^2. \end{aligned} \quad (8.31)$$

When Equation 8.31 is substituted into Equation 8.29 and the definition of a skew-symmetric matrix is used, it becomes:

$$h\mathbf{\Pi}_n + \frac{h^2}{2} \mathbf{L}_n = \frac{\sin \|\mathbf{f}_n\|}{\|\mathbf{f}_n\|} [I] \mathbf{f}_n + \frac{1 - \cos \|\mathbf{f}_n\|}{\|\mathbf{f}_n\|^2} \mathbf{f}_n \times [I] \mathbf{f}_n. \quad (8.32)$$

where Equation 8.32 is solved using Newton's method. Thus, given initial inputs, the PMAC simulation is time-marched continually using the following process [45]:

- (1) Determine \mathbf{f}_n using Newton's method to solve Equation 8.32 given $\boldsymbol{\omega}_n$ and \mathbf{L}_n

- (2) Determine $[F]_n$ using Equation 8.31 given \mathbf{f}_n
- (3) Determine $[R]_{n+1}$ using Equation 8.30 given $[R]_n$ and $[F]_n$
- (4) Determine \mathbf{L}_{n+1} using the models developed in Section 8.1.6 given $[R]_{n+1}$
- (5) Determine $\boldsymbol{\omega}_{n+1}$ using Equation 8.28 given $\boldsymbol{\omega}_n$, \mathbf{L}_n , \mathbf{L}_{n+1} , and $[F]_n$.

However, there are several drawbacks to using LGVI to integrate a PMAC simulation. It is difficult to include the hysteresis magnetizations as additional integration states. The LGVI estimates the attitude and angular velocity states after each integration step; the hysteresis magnetization integration is thus limited to Euler's method. It is possible to use a Runge-Kutta integrator between each LGVI integration step to obtain a higher-order estimate of the hysteresis magnetizations based on intermediate attitude estimates, but this has been found to greatly impact the processing time of the simulation. Instead, the simulation uses Euler's method to integrate the hysteresis magnetization when using LGVI for attitude integration, as has been done in the past [58] [44].

LGVI was developed to integrate rigid body dynamics when the external torque is dependent on attitude alone; it is not designed for torques which are dependent on angular velocity. The hysteresis torque is dependent on the body-frame time derivative of the magnetizing field (see Section 8.1.6.2); this is calculated using the transport theorem below [65]:

$$\begin{aligned}
 \frac{{}^{\mathcal{B}}\mathrm{d}}{\mathrm{d}t}(\mathbf{H}) &= \frac{{}^{\mathcal{I}}\mathrm{d}}{\mathrm{d}t}(\mathbf{H}) - \boldsymbol{\omega} \times \mathbf{H} \\
 &= [R]^{\mathcal{I}}\dot{\mathbf{H}} - \boldsymbol{\omega} \times [R]^{\mathcal{I}}\mathbf{H}
 \end{aligned} \tag{8.33}$$

Thus, the hysteresis torque is dependent on the current rotation matrix and the current angular velocity. Previous work using the LGVI [44] assumes $\boldsymbol{\omega}_n \approx \boldsymbol{\omega}_{n+1}$ when calculating the hysteresis magnetization derivative; this work uses this approach for simulations using the LGVI with hysteresis torque. The explicit Runge-Kutta integration methods described in Section 8.1.7.1 do not possess the drawbacks listed above.

8.2 Considerations

It is helpful to briefly consider concepts of attitude dynamics which apply to a PMAC satellite. A basic grasp of these concepts will shed light on the simulation results presented later in this chapter.

8.2.1 Torque-Free Rigid Body Motion

A rigid body without any external torques acting upon it will behave according to gyroscopic motion alone (Equation 2.1 with $\mathbf{L} = 0$). With this special condition, the angular momentum vector $\mathbf{\Pi} = [I]\boldsymbol{\omega}$ is constant in the inertial frame and the magnitude of angular momentum is constant in any frame. Also, the total system energy (due to kinetic energy alone) is conserved. Taken together, this means the satellite motion is bounded by the intersection of the energy ellipsoid

$$T = \frac{1}{2}I_1\omega_1^2 + \frac{1}{2}I_2\omega_2^2 + \frac{1}{2}I_3\omega_3^2. \quad (8.34)$$

and the momentum ellipsoid

$$\Pi^2 = \mathbf{\Pi}^T \mathbf{\Pi} = I_1^2\omega_1^2 + I_2^2\omega_2^2 + I_3^2\omega_3^2. \quad (8.35)$$

The above equations may be manipulated such that the angular momentum of each axis Π_i are the independent coordinates; this results in a momentum sphere intersected by an energy ellipsoid. Figure 8.4 (modified from Fig. 4.6 of [65], used with permission) shows the momentum sphere and the energy ellipse assuming $I_1 > I_2 > I_3$ (the CSSWE inertia matrix follows this trend as well). When the system energy is at maximum, only rotation about the minimum inertia axis is possible ($\boldsymbol{\omega} = \pm\omega_3\hat{\mathbf{n}}_3$). However, if the system loses energy over time (due to structural flexing, hysteresis, etc.), the energy ellipsoid will shrink until the system energy is equivalent to motion about the intermediate axis. However, the motion is not necessarily about the intermediate axis because of the separatrix, the boundary between the high energy domain (wobble about $\pm\hat{\mathbf{n}}_3$) and the low energy domain (wobble about $\pm\hat{\mathbf{n}}_1$). The motion of a rigid body at the separatrix is inherently

chaotic as the energy ellipsoid intersects the momentum sphere in many places. As more energy is removed, the satellite will settle to a wobble about $\hat{\mathbf{n}}_1$ or $-\hat{\mathbf{n}}_1$; either is equally likely after motion along the sepratrix. Continued damping will cause the motion to be solely about the major inertia axis ($\boldsymbol{\omega} = \pm\omega_1\hat{\mathbf{n}}_1$) [65].

Torque-free motion has important implications for a PMAC satellite. This is motion which every satellite experiences to some extent as it is present in the gyroscopic term of Equation 2.1. The extent to which a given satellite will “feel” the gyroscopic motion (and thus behave in this manner) may be examined by calculating the “gyroscopic torque” as follows:

$$\mathbf{L}_{GY} = -[\boldsymbol{\omega} \times][I]\boldsymbol{\omega}. \quad (8.36)$$

This “torque” (which a rigid body experiences even in the absence of external torque) is taken directly from Euler’s rotational equation of motion. If the gyroscopic torque \mathbf{L}_{GY} is much greater than the external torque \mathbf{L} then the rigid body will experience mainly torque-free motion. If the gyroscopic torque is much less than the external torque, torque free motion is not dominant. If the gyroscopic torques is similar in magnitude to the external torque, some combination of torqued and torque-free motion will result.

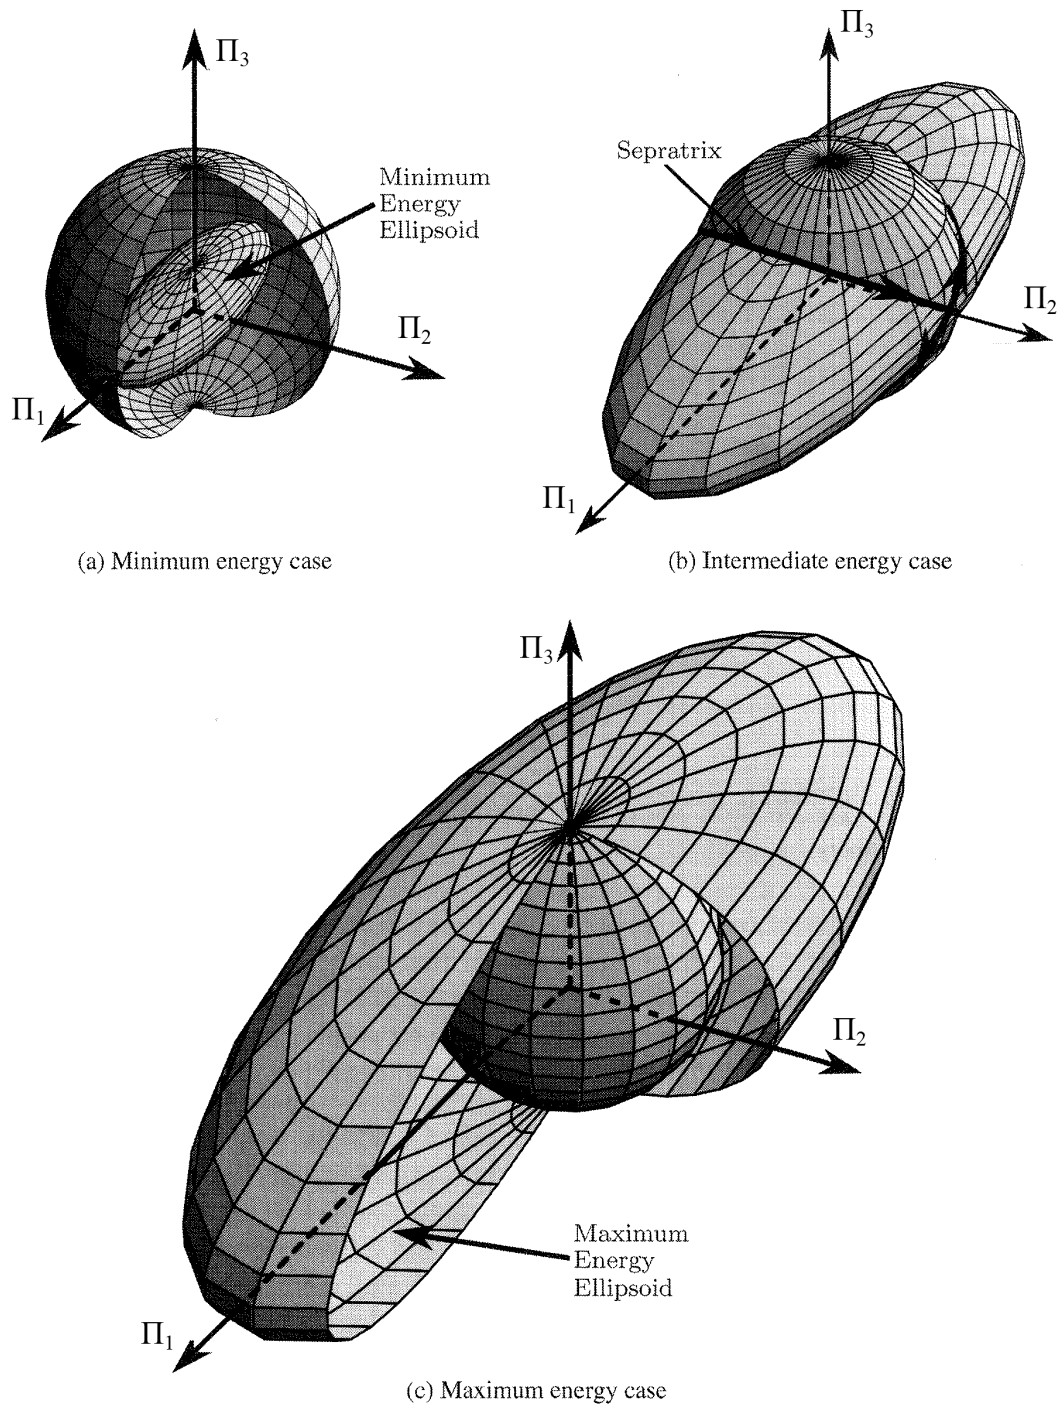
8.2.2 3D Pendulum Comparison

A satellite with a bar magnet in a constant magnetic field is dynamically equivalent to a rigid pendulum tethered at a distance to a fixed, frictionless point and acted upon by a constant gravitational acceleration. This special case is known as a 3D pendulum whose equations of motion are as follows [67]:

$$[I]\dot{\boldsymbol{\omega}} = -[\boldsymbol{\omega} \times][I]\boldsymbol{\omega} + \boldsymbol{\rho} \times m\mathbf{g} \quad (8.37)$$

where $\boldsymbol{\rho}$ is the distance vector from the pendulum center of mass to the pivot, m is the mass of the pendulum, and \mathbf{g} is the gravitational acceleration vector. Thus, the 3D pendulum $\boldsymbol{\rho}$ and $m\mathbf{g}$ are

Figure 8.4: Interactions of the momentum sphere and energy ellipsoid at (a) minimum energy, (b) intermediate energy, and (c) maximum energy conditions (modified from Fig. 4.6 of [65], used with permission).



analogous to the PMAC \mathbf{m}_{bar} and \mathbf{B} , respectively.

Because the equations of motion are numerically identical, conclusions drawn from previous studies of the 3D pendulum apply to the bar-magnet-only case of PMAC as well. What follows is a list of properties of the 3D pendulum which also apply to a PMAC satellite in certain conditions.

- (1) The 3D pendulum cannot be analytically solved [14]. If this is true of the PMAC bar-magnet-only case, it is also true of all other PMAC cases.
- (2) The 3D pendulum conserves the total energy of the system as well as the component of angular momentum about the axis parallel to $m\mathbf{g}$ (PMAC \mathbf{B}).
- (3) The system dynamics are unchanged by rotation about the axis parallel to $\boldsymbol{\rho}$ [67] (PMAC \mathbf{m}_{bar}).
- (4) The hanging equilibrium (with $\boldsymbol{\rho}$ parallel to $m\mathbf{g}$ and with the pendulum center of mass below the pivot) is Lyapunov stable [67]; initial states which are close to this point will remain close to this point. This is equivalent to a bar magnet which is aligned with the local magnetic field.
- (5) The inverted equilibrium (with $\boldsymbol{\rho}$ parallel to $m\mathbf{g}$ and with the pendulum center of mass above the pivot) is unstable [67]. This means that a small deviation from the equilibrium point could cause large changes in the state of the system. This is equivalent to a bar magnet anti-parallel to the local magnetic field.
- (6) Control in the form of $\mathbf{L} = \mathbf{u} \times m\mathbf{g}$ (where \mathbf{u} is the control input) preserves the conservation of angular momentum along the axis parallel to $m\mathbf{g}$ [67]. This is analogous to PMAC hysteresis torque (Equation 8.15).

The instability described in point (5) means that the PMAC dynamics are sensitive when $\beta = 180^\circ$; this should not be confused with chaotic motion due to the separatrix. Point (6) is important as it implies that different dynamics are in effect when non-magnetic environmental torques are included in a simulation.

8.3 Simplified Simulation Analysis

This section seeks to answer two major questions:

- (1) Can a simulation be used to represent PMAC dynamics? If so, to what extent?
- (2) If a simulation can be used to represent PMAC dynamics, what integrator and time step should be used to ensure accurate output?

A simplified simulation is developed to answer these questions at a reasonable computational cost. Because this simplified simulation maintains the key properties of the full simulation, it is assumed that the lessons learned can be applied to the full simulation. The assumptions and conditions of the simplified simulation are discussed before using the simulation for a variety of analysis. This study considers the LGVI and RK2-RK7 integrators as well as time step values of 1s, 0.1s, and 0.01s.

8.3.1 Description

A simplified simulation is developed in an attempt to understand the base properties of the full simulation. This simplified simulation can run much faster than the full simulation yet retains its key properties. The first simplification is an inertially-constant magnetizing field vector; this simulates a circular equatorial orbit if the earth has a perfect dipole magnetic field aligned with its poles with no offset. Second, when external torques are included in the model, only the gravitational gradient and solar pressure torques are included. The gravitational gradient assumes a constant inertial position of 450 km altitude above the surface of earth along $+^{\mathcal{I}}Y$. The solar pressure torque assumes a constant inertial unit vector from the earth to the sun $^{\mathcal{I}}\hat{\mathbf{s}} = [-1 \ 0 \ 0]^T$. Together, these are the gravity gradient and solar pressure torques for a prograde orbit dawn crossing at autumnal equinox; these constant inertial values are chosen for simple application and visualization. With these assumptions, only the ECI and body frames are required.

Two sets of initial inputs are run; each set is run through a variety of analysis. Both sets use a satellite magnetic moment $\mathbf{m} = [0 \ 0 \ 0.55]^T \text{ A}\cdot\text{m}^2$ and the principal moment of inertia matrix

Table 8.1: Nominal inputs for the simplified simulation sets are shown below. The derived values of β angle and initial system energy are also shown.

Parameter	Unit	Set 1	Set 2
Constant magnetizing field vector ${}^T\mathbf{H}$	A/m	$[0\ 0\ 20]^T$	$[25.18\ 2.76\ -8.59]^T$
Initial 1-2-3 Euler angles EA_{123}	deg	$[90\ 0\ 0]^T$	$[13.9\ -71.6\ 104.1]^T$
Initial angular velocity vector $\boldsymbol{\omega}_0$	deg/s	$[1\ 1\ 1]^T$	$[0.17\ -0.97\ 2.93]^T$
<i>(Derived Values)</i>			
Initial β angle	deg	90	178.1
Initial system energy	J	$7.46 \cdot 10^{-6}$	$2.82 \cdot 10^{-5}$

given in Section 4.2. Other parameters which define the hysteresis, gravity gradient, and solar pressure torques are given in Table 8.2.

The input values for both initial condition sets are defined in Table 8.1. The first set uses a 20 A/m constant magnetizing field and starts with the satellite perpendicular to the magnetic field and rotating at one degree per second in pitch, yaw, and roll. Initial input set two is defined by the CSSWE attitude and angular velocity as calculated by MEKF for September 14, 2012 at 00:59:48 UTC; this is eight minutes after PPOD deployment and shortly after the MEKF converges to an attitude solution (see Chapter 6).

8.3.2 Energy Conservation Analysis

The numeric integrator itself can cause the simulated system energy to drift over time. Although this behavior is undesired, every combination of numeric integrator and time step will have some energy drift; the key is to determine an acceptable level of this drift. Here the maximum allowable energy drift is set as the energy dissipated by a single ± 20 A/m cycle of one flight hysteresis rod over a 1000 hour simulation time; this duration (about 42 days) is considered the longest timespan over which the simulation will be used to calculate a settling time. Using the fitted hysteresis loop area shown in Table 7.2 and the volume of a single flight hysteresis rod, this amounts to a maximum allowable energy change of $3.3 \cdot 10^{-9}$ J over a 1000 hour simulation. It is assumed that energy differences below this threshold will have a negligible effect on the dynamics because

the rod-based energy loss will dominate the integrator-based energy drift.

The energy conservation of each integrator and time step is analyzed by running the simplified simulation for 1000 hours with a bar magnet only; no hysteresis, gravity gradient, or solar pressure torques are included. Ideally, a freely-floating bar magnet in a constant magnetic field will perfectly conserve the initial system energy as there are no dampening torques included in the simulation. Figure 8.5 shows the maximum energy drift of each integrator and time step combination over 1000 hours of simulation using input sets 1 and 2; an energy drift of zero is ideal. Input set 2 results in energy drifts as high as 100 times the level of input set 1. This may be due to motion near the instability point at $\beta = 180^\circ$, which is not seen by the dynamics when using input set 1. Also, set 2 has a higher initial energy than set 1. Due to computational cost, this analysis is not repeated for all possible initial inputs; instead set 2 is assumed to be the worst-case initial input. Thus, the integrator and time step combinations with output beyond the energy drift threshold are ruled out for use in the full simulation.


Although the LGVI is the quickest-running integrator for a given time step duration, its performance is often comparable to the RK4 at the same time step. Note that the higher-order integrators, when compared to the LGVI at a decreased time step, regularly show decreased energy drift at similar (or lower) computation times. This analysis is not in favor of using the LGVI for the attitude simulation as it is more complicated and more difficult to understand than Runge-Kutta methods while achieving similar performance.

8.3.3 Angular Error Analysis

The energy drift is not the only performance metric which can evaluate each integrator and time step combination; the beta angle error and the settling time error may also be used. However, the ideal beta angle and settling time are much more difficult to calculate than the bar-magnet-only system energy, which will ideally remain equal to the initial system energy. This problem is further complicated by the impossibility of an analytical solution for PMAC dynamics.

Here the approach is to compare each integrator and time step to the the highest-order,

Figure 8.5: The maximum energy change for each integrator and time step combination. The negligible energy change threshold of $3.3 \cdot 10^{-9} \text{ J}$ is denoted by the red dotted line. The normalized computation time of each simulation is also shown in green. The 0.01s time step is not shown due to the unreasonable computation times required to simulate 1000 hours.



lowest-time-step output, which is assumed to be “truth”. Figure 8.6 defines error relative to the RK7/0.01s output and shows the maximum beta angle error for simulations over a 30 hour duration using initial input sets 1 and 2. Three cases are run for each of the input sets: the bar-magnet-only case, the bar-magnet-and-hysteresis-only case, and the all-torque case. The output based on set 1 is expected; the error follows a downward trend to the right as the integrator order increases and the time step duration decreases. The inclusion of environmental torques results in slightly better performance in most cases. Note that the LGVI output usually exhibits the worst performance at each simulation time step. The LGVI performance is especially poor for the cases which include hysteresis torque; this is likely because the LGVI integrates the hysteresis magnetization using Euler’s method.

The output based on input set 2 is quite different; although the bar-magnet-only case performs better, the cases including dampening torques are much worse as they show little, if any, decrease in beta angle as the integrator order increases and the time step duration decreases. This behavior is believed to be due to the chaotic nature of the separatrix crossing which occurs for case 2 but not for case 1. Figure 8.7 shows the simulation time at which a 1-2-3 Euler Angle error magnitude of one degree is breached using input set 2. At a given time step, each integrator diverges from the “true” attitude at about the same time. This implies that the simulation has entered a chaotic region which causes this divergence. The chaotic behavior does not occur for the bar-magnet-only case (Figure 8.6). This implies that the chaotic region is at some intermediate state such as dampening through the separatrix.

8.3.4 Settling Time Analysis

The presence of a chaotic region does not rule out the possibility of accurate simulation. However, it does mean that the absolute attitude error is perhaps not the best metric for simulation performance. It is possible for the simulated beta angle to accumulate a large phase error without causing a significant change in the settling time. The settling time is chosen as the most important feature of the attitude simulation due to its importance in predicting on-orbit mission duration.

Figure 8.6: The maximum beta angle error for each combination of integrator and time step using input set 1 (top) and input set 2 (bottom) over a 30 hour simulation duration. Each plot shows the performance of the bar-magnet-only case (blue), the bar-magnet-and-hysteresis case (red), and the all-torques case (purple). The beta angle error is defined by comparing the output with the RK7/0.01s output of each case.





Figure 8.7: The time at which the magnitude of the 1-2-3 Euler angle error exceeds one degree is shown for each integrator and time step combination. These data are generated using initial input set 2.



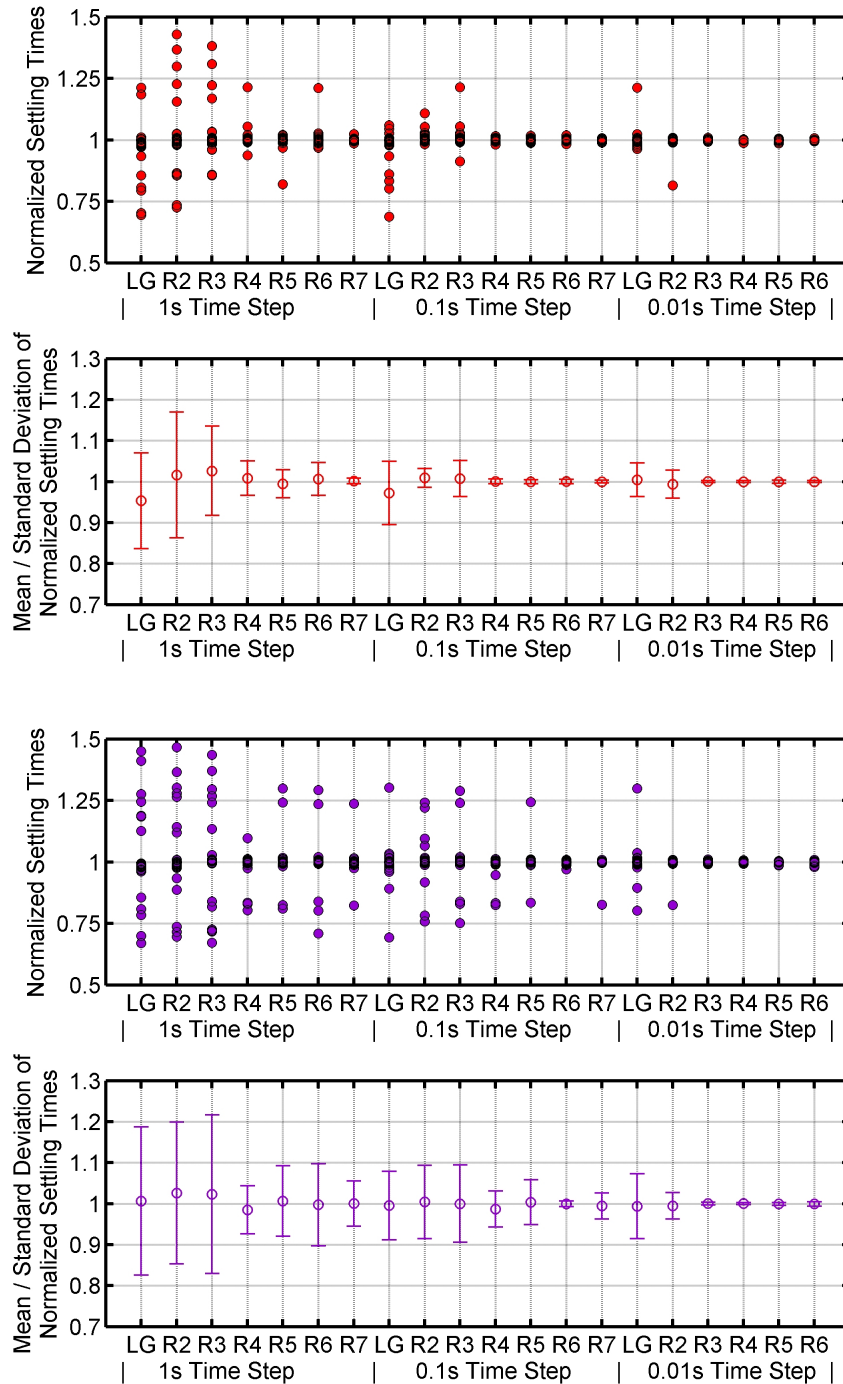
The presence of a chaotic region calls for a different analysis procedure. Initial input set 2 is used as the base input to the simplified simulation. The base input 1-2-3 Euler Angles and angular velocity vector are perturbed by random Gaussian noise. The standard deviation of this noise is set by the 1σ uncertainty of the MEKF fit which defines initial input set 2. The standard deviations are $\sigma_{EA123} = [0.3625 \ 0.6261 \ 0.9539]^T$ in units of degrees and $\sigma_{\omega 0} = [1.4950 \ 2.4140 \ 1.0829]^T \cdot 10^{-2}$ in units of degrees per second. An array of random Gaussian noise values is generated once and loaded before each simulation to ensure that a given run receives identical initial input over the full range of integrator and time step combinations. The simulation is run thirty times and the settling time for each run, integrator, and time step is calculated. Here the settling time is defined as the time after which the beta angle remains within 5° of its final value.

Figure 8.8 shows the settling time distribution, mean, and standard deviation when each run is normalized by its associated RK7/0.01s settling time. The normalized settling times converge as the integrator order increases and the time step duration decreases. This gives confidence that the simplified simulation is converging toward the true settling time for each initial input.

Figure 8.8 is also helpful when deciding upon an integrator and time step for the full simulation. Without non-magnetic torques, RK7 at 1s represents the true system behavior quite well. However, when non-magnetic environmental torques are included in the simulation, RK4 and above at 0.1s or below is necessary to achieve realistic results.

The true settling time is quite sensitive to the initial conditions. Figures 8.9 and 8.10 show initial system energy versus the settling times calculated using RK7/0.01s given normally-perturbed inputs from initial input set 2 for both the bar-magnet-and-hysteresis case and the all-torques case. As shown, the settling time can be quite sensitive to small changes in the initial inputs; the settling times appear to be mostly well-grouped with a few outliers. In order to investigate this behavior further, the sample median for each case is calculated, and bounds are defined for settling times within $\pm 3\%$ of the median; the $\pm 3\%$ threshold is chosen because it is the tightest bound which collects all of the grouped settling times for the bar-magnet-and-hysteresis case. The sample median is used because it is less sensitive to outliers than the sample mean. Pearson's product-moment

Figure 8.8: The simplified simulation normalized settling time, mean, and standard deviation for the bar-magnet-and-hysteresis case (top plots, red) and the all-torques case (bottom plots, purple) using initial input set 2. The settling time for each randomly-perturbed dataset is normalized by the RK7/0.01s settling time of that dataset. The individual settling times for each run are shown on plots 1 and 3 while the mean and standard deviation of the runs for each integrator and time step are shown on plots 2 and 4.



correlation coefficient is calculated for the total and selective samples of each case; the p-value of each correlation coefficient is also determined using a t -distribution with $n - 2$ degrees of freedom, where n is the number of samples in the dataset [80]. The p-value is the probability that the sample correlation could occur by random chance if there is truly no correlation; p-values lower than 0.05 typically represent a statistically significant correlation.

One would expect the initial energy of the system to be strongly correlated with the system settling time. However, the data shows that a significant correlation is only found using settling times within $\pm 3\%$ of the median. All of the calculated settling times appear to represent the true dynamics because the other integrators and time step values converge to the same result, even for the outlying runs (see Figure 8.8). There appears to be some true nonlinear behavior affecting the settling times. However, for all of the RK7/0.01s runs, the settling time does not exceed 110% of the sample median. Rather, the settling time is unexpectedly smaller than it should be. By discretizing the settling times into those above and below 110% of the 30-sample median, a binomial distribution fit finds that the probability of the population of all perturbed simplified simulation runs remaining below 110% has an upper bound of 1 and a lower bound of 0.8843 (with 95% confidence).

This implies that although the PMAC dynamics are very sensitive to initial conditions, the settling time is bounded on the high side. Thus, the nonlinear behavior of the PMAC dynamics is unlikely to cause the settling time to greatly increase. However, the PMAC dynamics may cause the settling time to be significantly smaller than expected. This is good news as it means a simulation can be used to determine the true worst-case settling time of a PMAC satellite.

However, if a simulation is to be used to determine the worst-case settling time, the abnormally low settling times should be avoided. If the simulation is run once and it happens to align with an abnormally-low settling time, a mission designer may incorrectly assume that the expected settling time is lower than it truly is; this incorrect interpretation could impact mission operations. Fortunately, the simplified simulation runs with abnormally low settling times possess a distinct dynamic response. Figures 8.11 and 8.12 show the angular velocity components versus time for

Figure 8.9: The RK7/0.01s settling times versus the initial energy for the simplified simulation perturbation runs of the bar-magnet-and-hysteresis case using the set 2 initial inputs. The left plot shows the settling times for all thirty cases run and uses the dotted black lines to bound values within 3% of the sample median shown by a solid red line; values outside these bounds are represented by empty circles while values within the bounds are represented by filled red circles. The right plot shows only those settling times within 3% of the entire dataset median. Pearson's product-moment correlation coefficient and the associated p-value are shown for the data within each plot.





Figure 8.10: The RK7/0.01s settling times versus the initial energy for the simplified simulation perturbation runs of the all-torques case using the set 2 initial inputs. The left plot shows the settling times for all thirty cases run and uses the dotted black lines to bound values within 3% of the sample median shown by a solid purple line; values outside these bounds are represented by empty circles while values within the bounds are represented by filled purple circles. The right plot shows only those settling times within 3% of the entire dataset median. Pearson's product-moment correlation coefficient and the associated p-value are shown for the data within each plot.



each of the 30 perturbed initial input set 2 simulations for the bar-magnet-and-hysteresis and the all-torques cases; the runs which have a settling time below 95% of the sample median are shaded.

The figures show that the normal response of the satellite is to achieve a mostly flat spin about the major inertia axis ${}^B X$ early in the simulation. The satellite remains in this configuration until ω_X has dampened to the level of ω_Y . At this point, the rotation is transferred solely to ${}^B X$ and ${}^B Z$ in the form of roll- and yaw-wobble about the constant magnetizing field. Note that the sign of the major inertia axis rotation ω_X flips at random; this is anticipated when traveling through the chaotic sepratrix. The angular velocity response of each simulation with low settling time is quite different and easy to distinguish from the normal case. Thus, the angular velocity response of the simplified simulation output can give clues to its reliability. The response of the full simulation may be equally helpful in determining its validity.

8.3.5 Summary

An analysis of the energy conservation of each integrator showed that, for PMAC simulation, higher order Runge-Kutta methods are comparable to LGVI in both energy conservation and computation time. Energy conservation analysis shows that acceptable numeric integrator performance is dependent on the time step. Generally, RK4 and above at a time step of 0.1s or smaller yields acceptable energy conservation at the worst-case initial input considered.

Beta angle analysis shows that Runge-Kutta methods clearly outperform LGVI in limiting absolute attitude error. However, this analysis also shows that the absolute attitude error is difficult to minimize for a PMAC simulation with hysteresis in certain initial conditions. This is likely due to the satellite traveling through the chaotic sepratrix with continued energy dissipation.

The settling time analysis shows that all integrators and time steps converge to one settling time for each run of identical perturbed initial conditions. This behavior is interpreted as the simulation converging to the true dynamics of the system. However, the true settling time is very sensitive to the initial conditions. The perturbation analysis shows that small changes in the initial state can lead to abnormally low settling times. For the simplified simulation, this abnormal

Figure 8.11: The RK7/0.01s angular velocity components versus time for each of the simplified simulation runs of the bar-magnet-and-hysteresis case using initial input perturbed from input set 2. The axis labels have been omitted from each plot for clarity. For each plot, the vertical axis ranges from -4 to 4 degrees per second while the horizontal axis ranges from 0 to 30 hours. The angular velocity components ω_X , ω_Y , and ω_Z are shown in blue, green, and red, respectively. The runs which result in a settling time less than 95% of the sample median are shaded.

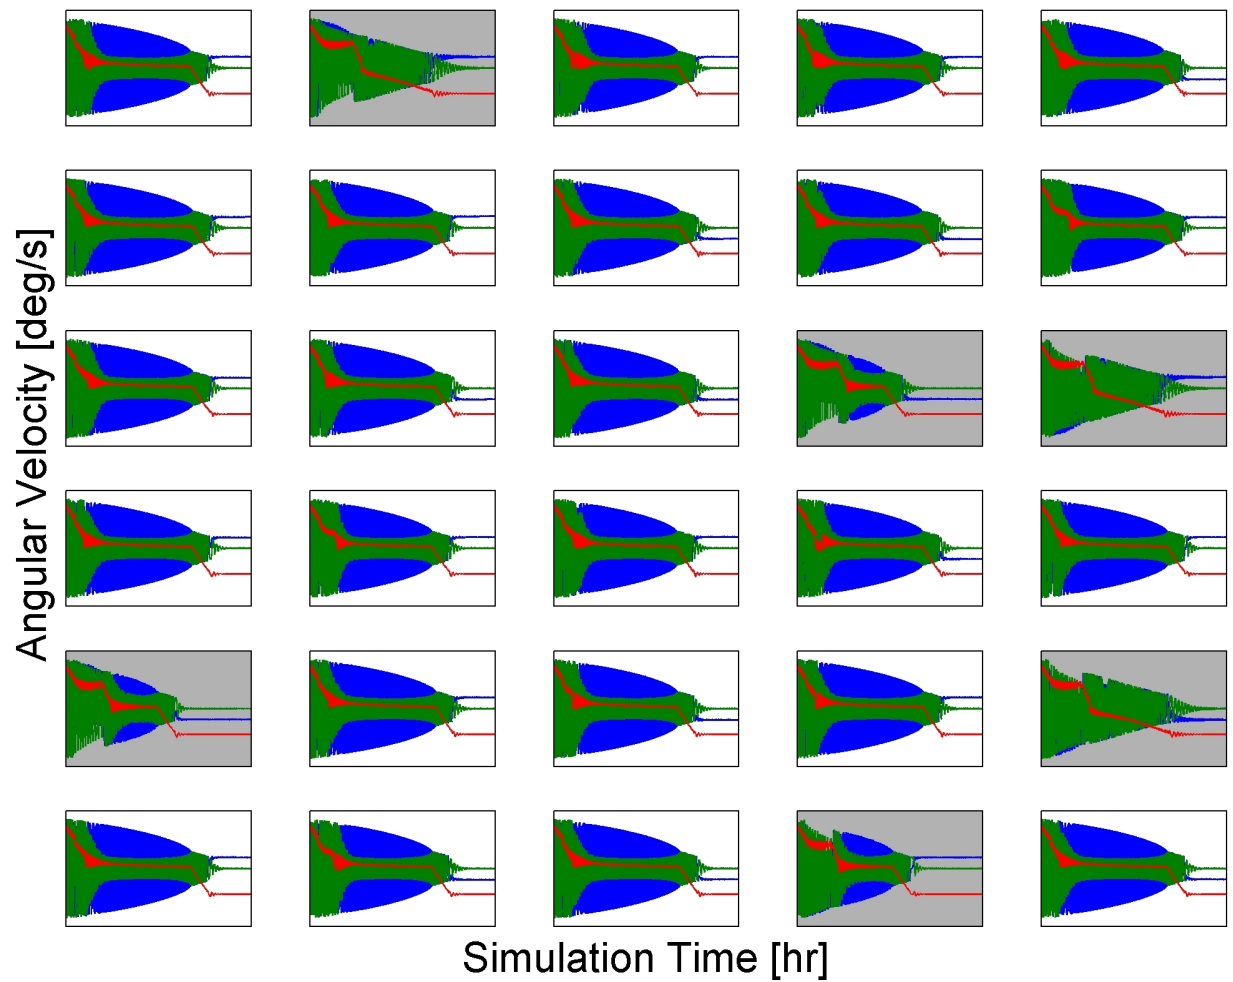
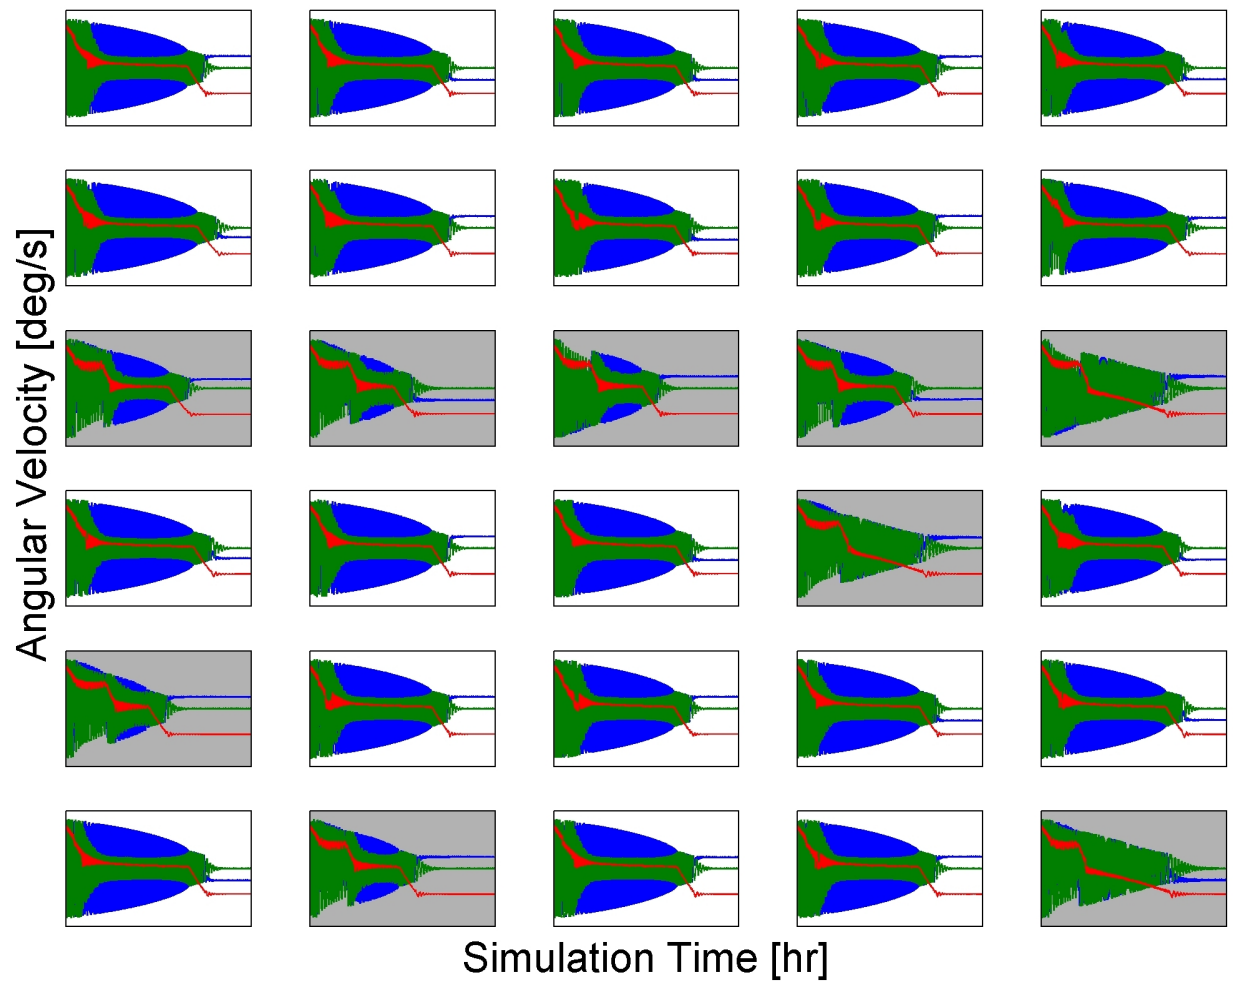


Figure 8.12: The RK7/0.01s angular velocity components versus time for each of the simplified simulation runs of the all-torques case using initial input perturbed from input set 2. The axis labels have been omitted from each plot for clarity. For each plot, the vertical axis ranges from -4 to 4 degrees per second while the horizontal axis ranges from 0 to 30 hours. The angular velocity components ω_X , ω_Y , and ω_Z are shown in blue, green, and red, respectively. The runs which result in a settling time less than 95% of the sample median are shaded.



behavior can be identified by the dynamic response of the simulation over time.

8.4 Results

Before the results are presented, the nominal simulation inputs are defined. Nominal case outputs such as beta angle, angular velocity, and system energy are presented. This nominal output is compared to on-orbit data and the magnetic-torque-only case.

8.4.1 Nominal Input

The full PMAC simulation nominal inputs are shown in Table 8.2. The RK4 integrator at 0.1s time step is chosen as the nominal case; simplified simulation showed this combination to have low rates of settling time error (see Figure 8.8) at moderate computational cost. Also, RK4 leaves higher orders of Runge-Kutta at the same time step for output comparison.

The hysteresis rod loop parameters are set by the best fit to the measured hysteresis rods (see Section 7.3.4.3). The moments of inertia are based on the SolidWorks model of the CSSWE spacecraft. The distance vector from the satellite CG to the geometric center \mathbf{r}_d is also given by the SolidWorks model. The ap index and F10.7 flux values are given by data from the previous solar cycle; the selected values overestimate the true solar activity which was actually experienced over the first ten days on orbit (the true 10-day average values of ap index and F10.7 flux were 6.1 and $113.7 \cdot 10^{-22} \text{ W} \cdot \text{m}^{-2} \text{ Hz}^{-1}$, respectively [2]). The early-mission CSSWE TLE is used as input to the SGP4 orbit propagator.

The base initial angular velocity vector and initial 1-2-3 Euler angles are equivalent to the simplified simulation initial input set 2, which is defined by early-mission CSSWE MEKF output. As with the analysis presented in Section 8.3.4, the initial attitude and angular velocity vector are perturbed by normal Gaussian noise with a standard deviation equivalent to the 1σ uncertainty of the MEKF filter output. This process allows the simulation performance to be examined over a number of perturbed initial inputs. However, the data from a single initial input set is used as the nominal output and presented below.

Table 8.2: Nominal inputs for the full simulation are shown below. Most inputs are based on CSSWE values; all inputs are given in the body frame when applicable.

Parameter	Value	Unit
Numeric Integrator	RK4	
Time step duration h	0.1	s
Simulation Start Date/Time	2012.09.14 00:59:48	UTC
Base Initial 1-2-3 Euler Angle EA_{123}	[13.9 -71.6 104.1]	deg
Base Initial angular velocity vector ω_0	[0.17 -0.97 2.93] ^T	deg/s
Initial magnetic field offset β	178.1	deg
X-axis moment of inertia I_{xx}	$2.22 \cdot 10^{-2}$	kg·m ²
Y-axis moment of inertia I_{yy}	$2.18 \cdot 10^{-2}$	kg·m ²
Z-axis moment of inertia I_{zz}	$5.00 \cdot 10^{-3}$	kg·m ²
Bar magnet magnetic moment vector \mathbf{m}_{bar}	[0 0 0.55] ^T	A·m ²
Number of hysteresis rods per body axis	[3 3 0] ^T	
Hysteresis rod length	95	mm
Hysteresis rod diameter	1	mm
Initial hysteresis rod magnetic flux density	0	Tesla
Hysteresis rod coercivity H_c	0.3381	A/m
Hysteresis rod remanence B_r	$6.0618 \cdot 10^{-4}$	Tesla
Hysteresis rod saturation B_s	0.3000	Tesla
Residual magnetic moment vector ${}^B\mathbf{m}_{\text{res}}$	[0.0059 0.0083 -0.0004] ^T	A·m ²
Satellite coefficient of drag C_d	2.4	
Distance vector from satellite CG to geometric center \mathbf{r}_d	[2.601 -0.218 -8.086] ^T	mm
Satellite surface area by body axis S	[0.01 0.03 0.03] ^T	m ²
Satellite coefficient of reflectivity c_R	0.8	
3-hour ap index average for the last 57 hours	48	
81-day average F10.7 flux	$168.5 \cdot 10^{-22}$	W·m ⁻² Hz ⁻¹
Daily F10.7 flux for previous day	$128.7 \cdot 10^{-22}$	W·m ⁻² Hz ⁻¹
Solar pressure at earth P_S	$4.5 \cdot 10^{-6}$	N·m ⁻²
TLE		
1 90039U 0	12268.58971383 +.00002482 +00000-0 +23852-3 0 00208	
2 90039 064.6731 007.9077 0219372 286.2692 203.1718 14.79135411001569		

8.4.2 Nominal Output

The nominal simulation output is shown using a variety of plots. Figure 8.13 shows the components of angular velocity simulating the first ten days after orbit insertion. Angular velocity is exchanged between satellite axes in accordance with the difference in mass moment of inertia of the axis (as expected from Equation 2.3). The roll rate is observed to rapidly oscillate in the early motion as the satellite is in a full tumble; no single axis dominates the angular velocity vector. The tumble becomes more controlled approximately 2.5 days after orbit insertion when the motion is mainly about the major inertia axis. However, the major inertia spin dampens to the level of intermediate axis while the roll rate steadily climbs. The final settling to a non-zero roll rate is intuitive; a PMAC satellite cannot rotate about any other axis while the bar magnet remains parallel to the magnetic field.

Figure 8.14 shows the kinetic, potential, and total rotational energy of the simulated satellite. A nearly linear decrease in energy is visible over the first four days. Immediately following day six, the satellite kinetic energy remains at a constant nonzero value. Post-settling variations in the potential energy are due to magnetic field amplitude changes throughout the satellite orbit.

Figure 8.15 shows the β angle between the BZ -axis and the local magnetic field vector. A green line has been added to represent settling at $\beta \leq 10^\circ$. Although the instantaneous beta changes rapidly, the maximum beta angle decreases nearly linearly over time, reflecting the system energy behavior. The simulation finds that the satellite settles to the magnetic field six days after PPOD deployment. Following settling, the beta angle remains within five degrees of the local magnetic field.

The next plots show the simulation output before and after settling occurs using 100-minute (about one orbit) datasets. The pre- and post-settling datasets are the simulation output at orbit 15 and 105, approximately 1 and 7 days after PPOD deployment. Figure 8.16 shows the relative magnitudes of the external torques acting on the satellite over both orbits.

The bar magnet dominates the other torques in the early mission; post-settling it is at the level

Figure 8.13: The nominal simulation output body-frame angular velocity vector components are shown. Yaw (blue) is about the ${}^B X$ (maximum inertia) axis, pitch (green) is about the ${}^B Y$ (intermediate inertia) axis, and roll (red) is about the ${}^B Z$ (minor inertia) axis.




Figure 8.14: The kinetic, potential, and total satellite rotational energy values are shown. The total energy settles to a constant offset from the potential energy due to a non-zero kinetic energy.





Figure 8.15: The β angle between $^B Z$ -axis and the local magnetic field vector is shown. A green line has been added at the value $\beta = 10^\circ$; this is used as the attitude settling threshold.



of the magnetic residual torque and the gyroscopic torque. Even before settling, the gyroscopic torque is within an order of magnitude of the bar magnet torque and at times is the dominant torque. This means that torque-free motion must be considered as it likely has a substantial effect on the system dynamics. The hysteresis rods lose an order of magnitude of torque from pre- to post-settling; they are demagnetized by decreasing H-field amplitude cycles as the satellite aligns. In the early mission, the hysteresis rods dominate the time-varying external torques but by late mission, the hysteresis torque is at the level of the gravity gradient and drag torques. After settling, the energy dissipation provided by hysteresis is in equilibrium with the energy input from the non-magnetic environmental torques. As expected, the eddy current torque is higher during the early mission which has higher rotation rates.

Figure 8.17 shows the hysteresis loops traced during the early- and late-mission orbits. The rods experience a much larger range of magnetizing fields before settling occurs. The hysteresis bounds remain the same throughout the simulation as they are directly set by the hysteresis parameters (H_c , B_r , and B_s) which do not change. The loops are thin and hard to distinguish in the early mission. After settling, the magnetization output is observed to produce small loops, always within the bounds of the inverse tangent loop. These “minor loops” are expected for ferromagnetic materials experiencing a magnetizing field insufficient to reach saturation [17]; the simulation is correctly modeling the low cycle amplitude hysteresis response. As mentioned previously, the traced hysteresis loop area is equivalent to the energy dissipated from the system. After settling, the hysteresis dampening is in equilibrium with the energy input by non-magnetic external torques; the constant magnetic external torques \mathbf{L}_B and \mathbf{L}_R cannot add or remove energy from the system.

Figure 8.17 shows a troubling characteristic of the Flatley model using the fitted hysteresis parameters. In the early mission orbit, an average of 34% of the hysteresis magnetizations need correction after each simulation time step. The late mission orbit does not require any corrections. It is difficult to determine the effect of the correction frequency; removing the correction in the early mission causes the simulated magnetization to exceed the inverse tangent bounds and become unrealistic. This is a drawback of the Flatley model, but the effect does decrease slightly for higher-

Figure 8.16: The magnitude of each external torque acting on the spacecraft is shown for the pre-settling orbit 15 (top) and the post-settling orbit 105 (bottom). Note the change of scale for the y-axis between the two plots. Here \mathbf{L}_{GY} is the gyroscopic motion torque ($[\boldsymbol{\omega} \times][I]\boldsymbol{\omega}$), \mathbf{L}_B is the bar magnet torque, \mathbf{L}_H is the hysteresis torque, \mathbf{L}_{GG} is the gravity gradient torque, \mathbf{L}_D is the drag torque, \mathbf{L}_{SP} is the torque due to solar pressure, \mathbf{L}_R is the magnetic residual torque, and \mathbf{L}_{EC} is the eddy current torque.

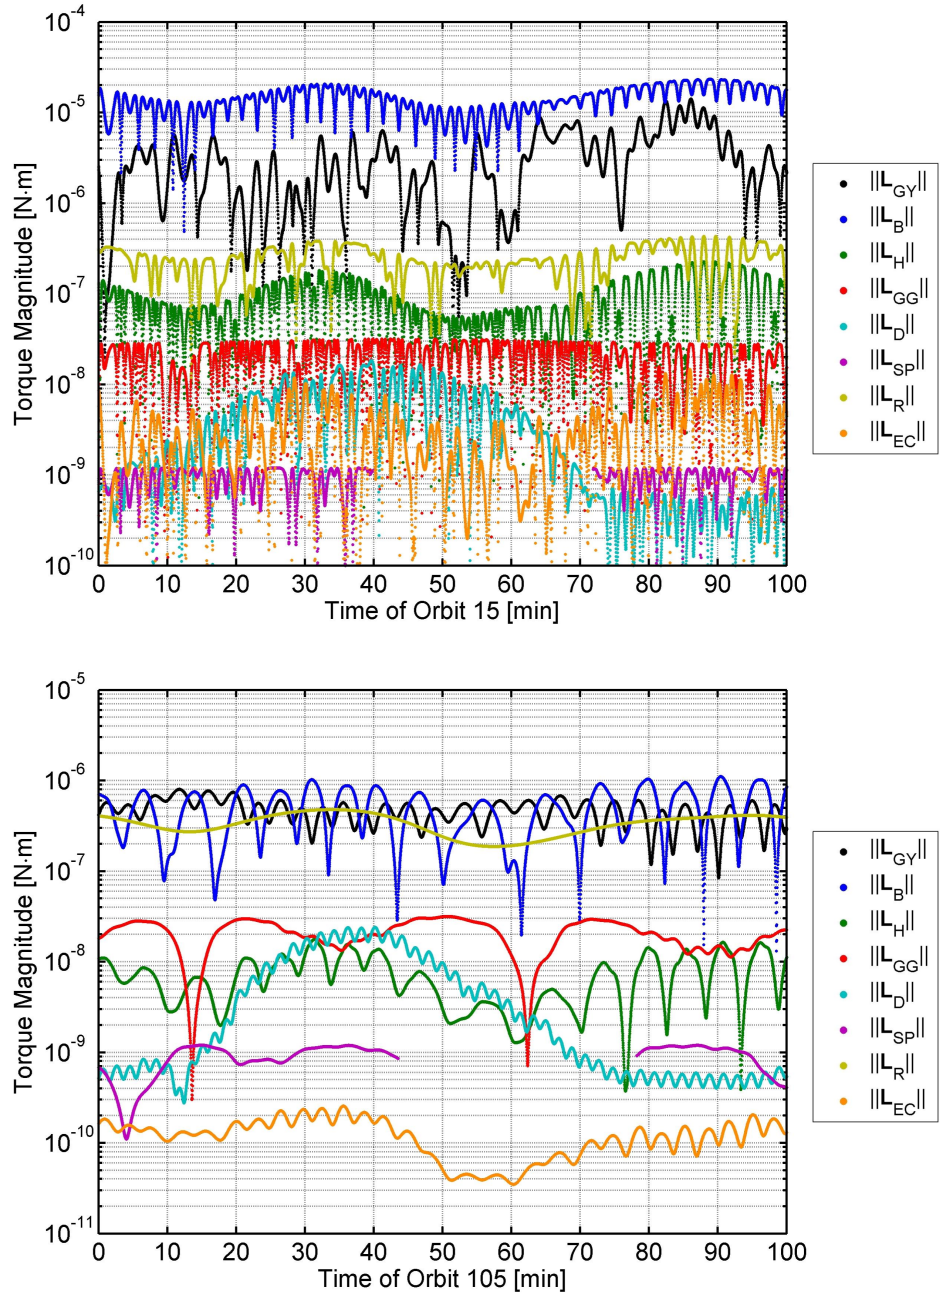
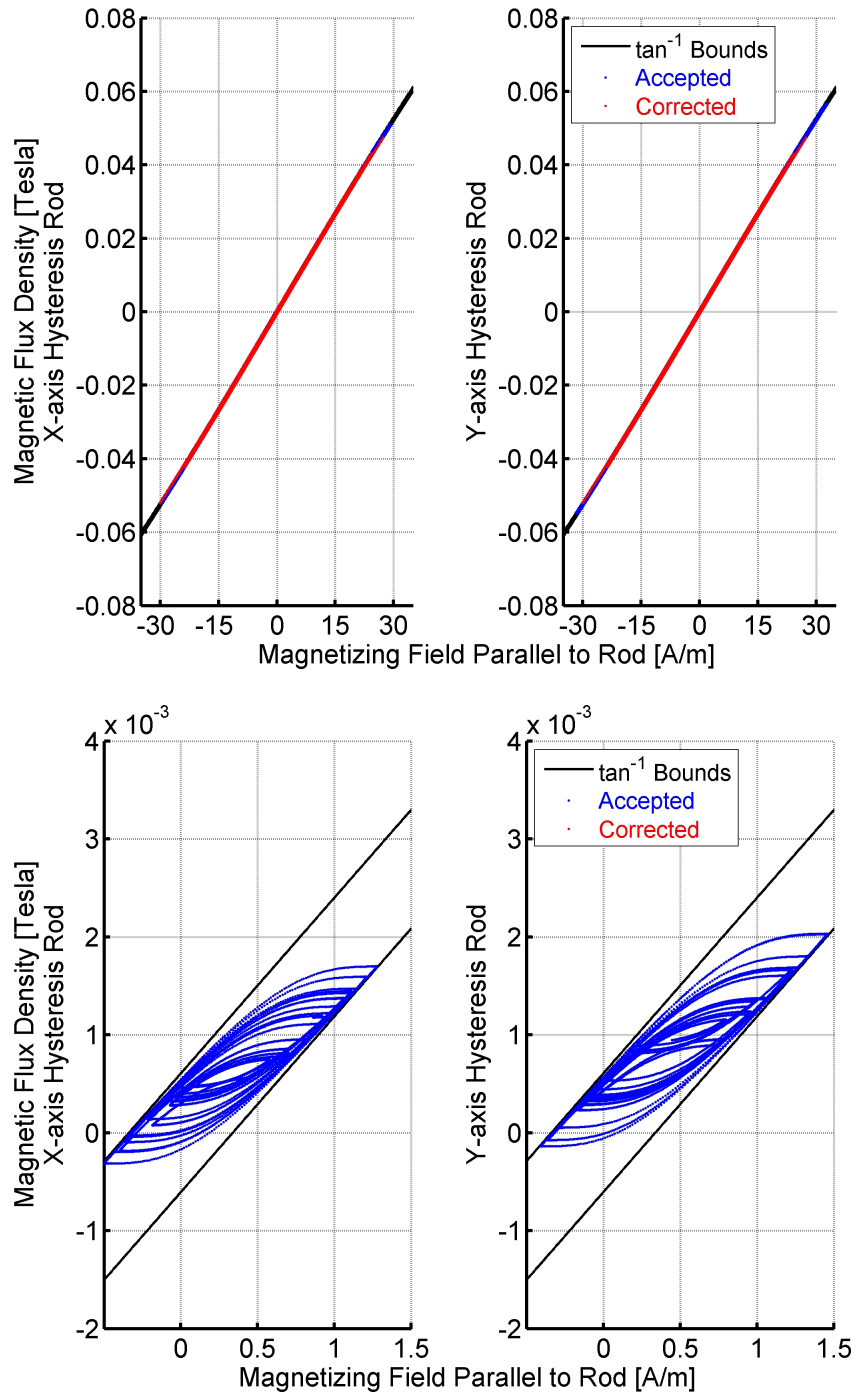


Figure 8.17: The single-orbit X- and Y-axis hysteresis loops simulated before (top) and after (bottom) the attitude settles. The inverse tangent bounds of the Flatley hysteresis model are shown. The magnetizations which have been corrected to remain within the inverse tangent bounds are shown in red; all other output is shown in blue.



order integrators at low simulation time step values.

8.4.3 High-Order Integrator Comparison

Although the simplified simulation showed the adequate performance of the RK4 at a 0.1s time step, confidence can be gained by comparison with a higher-order integrator such as the RK7 at the same time step. This section collects the results of this comparison analysis. Figure 8.18 compares the angular velocity components of both integrators. Both integrators display early-simulation rapid roll rate variation, mid-settling near-zero roll rate, and post-settling constant roll rate. The amplitude of the roll rate is inverted for the RK7 case; this may be due to chaotic dynamics during settling. Both datasets show most flat spin about the major inertia axis in the middle of attitude settling.

Figure 8.19 compares the energy response of the RK4 and RK7 integrators. The responses are quite similar, although the RK7 dampens the kinetic energy significantly quicker than the RK4 integrator. Note that the RK7 does not experience the jump in kinetic energy one day after PPOD deployment. This may represent a sensitive dynamics region which happened to increase the energy for RK4 but not for RK7, causing the RK4 simulation to have a longer settling time.

Figure 8.20 compares the beta angle as generated by the two integrators. The decreased energy of the RK7 simulation causes the attitude to settle in five days instead of six. The structure of the beta angle is similar for both cases; both show a linear decrease in the maximum beta angle over time. Both begin to track the magnetic field at an offset before slowly removing the offset over the course of about 36 hours. Further simulations are necessary to determine the cause of the difference in settling times.

8.4.4 On-Orbit Data Comparison

The output from the simulation may be verified by comparison with data filtered from the on-orbit attitude measurements of the CSSWE satellite. The initial conditions of the simulation are set based on the filtered data of the satellite close to its deployment from the PPOD. If the

Figure 8.18: Angular velocity components of the RK4 and RK7 integrators are shown. Yaw (blue) is about the ${}^B X$ (maximum inertia) axis, pitch (green) is about the ${}^B Y$ (intermediate inertia) axis, and roll (red) is about the ${}^B Z$ (minor inertia) axis.




Figure 8.19: The kinetic (blue), potential (green), and total (red) rotational energy as calculated by the RK4 (top) and RK7 (bottom) integrators.





Figure 8.20: The β angle between the local magnetic field and the $^B Z$ axis as calculated by the RK4 (top) and RK7 (bottom) integrators is shown. A green line at $\beta = 10^\circ$ has been added to show the settling threshold.



simulation is valid, it should roughly agree with the MEKF output.

The angular velocity is compared in Figure 8.21; similar behavior is observed. The roll rate shows equivalent variation in the first few days after deployment. The filtered data shows that the satellite maintains a near-zero roll rate for a short duration 2.5 days after PPOD deployment before approaching a roll rate of $-1^\circ/\text{s}$ until day six. The simulated data remains at a near-zero roll rate longer before settling to $1^\circ/\text{s}$ shortly before settling. The CSSWE attitude response does not experience a flat spin about any axis before settling. Instead, the angular velocity is equally shared between the intermediate- and major-inertia axes over the entire timespan.

Figure 8.22 compares the simulated and measured system rotational energy. The simulation does not model the antenna deployment event which occurs two hours after deployment and changes the bar magnet moment from $0.84 \text{ A}\cdot\text{m}^2$ to $0.55 \text{ A}\cdot\text{m}^2$ (see Section 6.3.2.1). This event causes the discontinuity seen in the filtered energy shortly after deployment. However, after the event, the total energy is roughly equal to the simulated condition. The simulated energy loss is more linear than the experimental data shows. This is likely because the Flatley hysteresis model accuracy degrades as the simulated magnetizing field cycle amplitude diverges from the experimental cycle amplitude used to generate the hysteresis fit parameters (see Section 7.3.4.1).

Because of the limitations of the model, the parameters were fitted to experimental data collected at $\pm 10 \text{ A/m}$ and $\pm 20 \text{ A/m}$. At cycle amplitudes larger than the fitted experimental datasets (such as shortly after PPOD deploy), the Flatley model will underestimate the hysteresis dampening. However, if the simulated cycle amplitude is much less than the fitted experimental datasets, the Flatley model may cease representing hysteresis as a closed loop (as shown Figures 7.15 and 7.16). The net effect of these errors is to linearize the simulated hysteresis dampening. The experimental data also shows that the energy dissipation continues after settling, when the majority of the rotation is about the roll axis; this simulation does not show this behavior.

Figure 8.23 compares the simulated and filtered β angle. CSSWE settles to the local magnetic field 6.5 to 7.5 days after PPOD deploy; this is comparable to the simulated six day settling time. The beta angle structure reflects the differences in energy dissipation shown previously; the

Figure 8.21: The angular velocity vector components as simulated by the RK4 at 0.1s (top) and as measured by the MEKF output of CSSWE data (bottom). Yaw (blue) is about the ${}^B X$ (maximum inertia) axis, pitch (green) is about the ${}^B Y$ (intermediate inertia) axis, and roll (red) is about the ${}^B Z$ (minor inertia) axis.

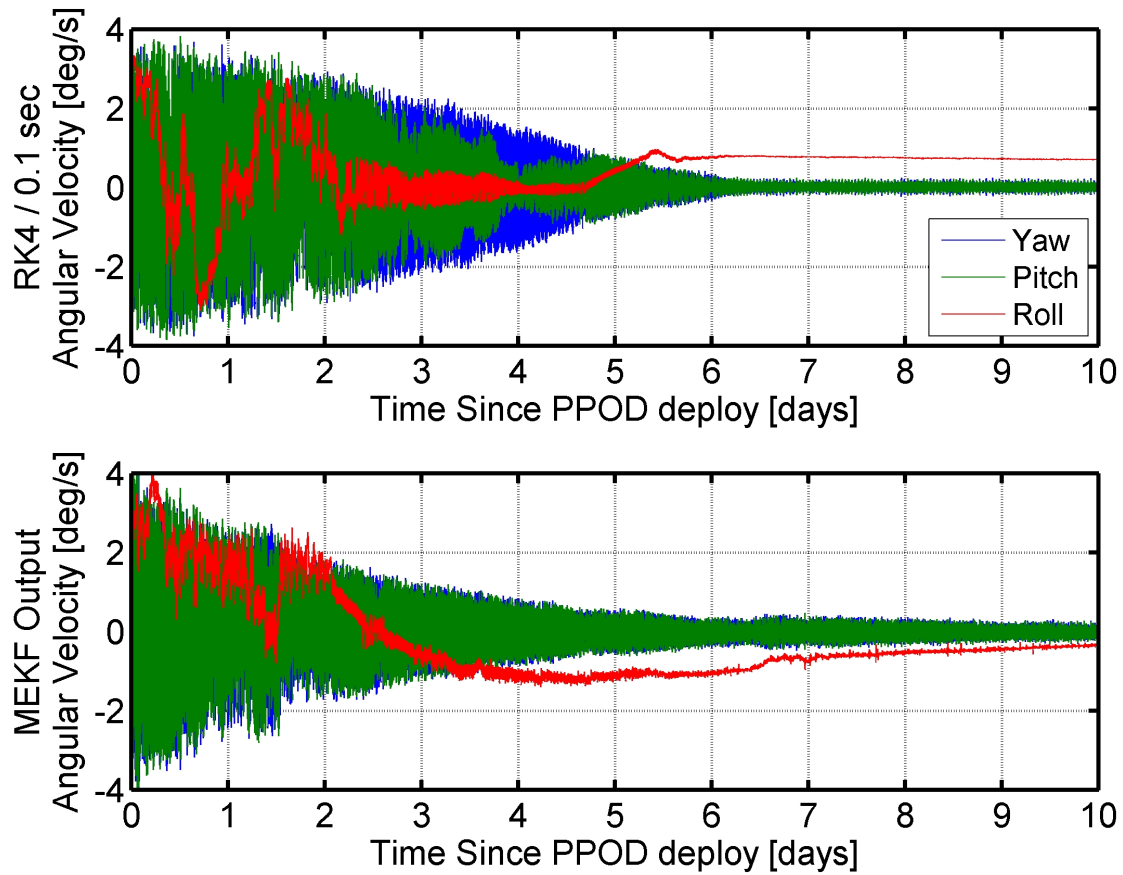
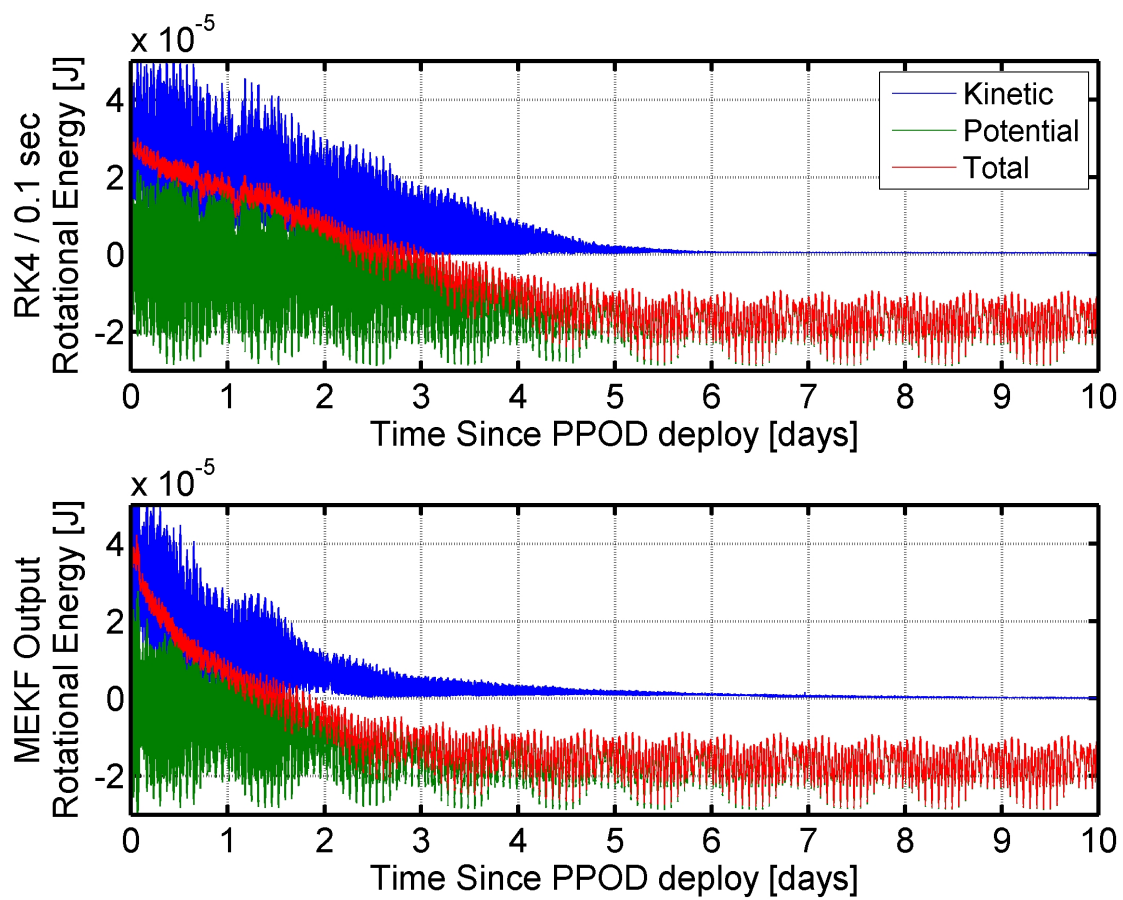


Figure 8.22: The satellite rotational kinetic (blue), potential (green), and total (black) rotational energy as simulated by the RK4 at 0.1s (top) and as measured by the MEKF output of CSSWE data (bottom).

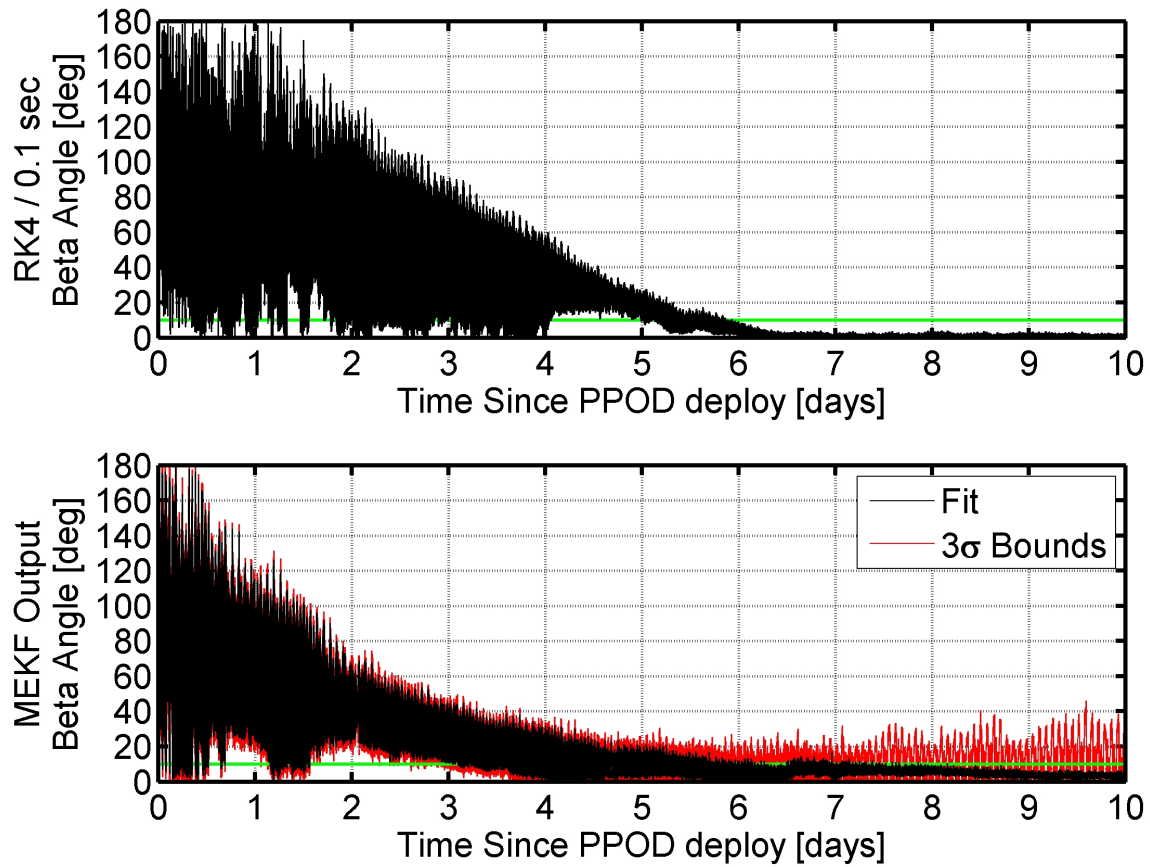


maximum beta angle decrease is mostly linear for the simulation but less so for the filtered data.

Overall, the filtered on-orbit attitude data is in good agreement with the simulation, which is able to realistically represent the true dynamics of a PMAC system. The RK4 and RK7 settling times are 6 and 5 days after PPOD deployment; these estimates are compared to the filtered data 7.5 day settling time and are found to possess errors of 20% and 33%, respectively. The simulation performance may be compared with the RAX-2 CubeSat mission, which settled after two months despite being predicted to settle within days (see Section 3.1.3).

Simulating PMAC dynamics is difficult but the research outlined in this dissertation represents a significant step forward. It is expected that simulation performance could be improved by using a hysteresis model which is better able to represent the measured hysteresis loops over a wider range of cycle amplitudes. Even at the current simulation performance levels, the settling time estimate is quite useful for satellite mission planning purposes.

Figure 8.23: The β angle between the local magnetic field and the $^B Z$ axis as simulated by the RK4 at 0.1s (top) and as measured by the MEKF output of CSSWE data (bottom). A green line is shown marking $\beta = 10^\circ$; when the angle remains below this threshold the system is considered settled. The 3σ uncertainty bounds of the β angle are shown in red for the MEKF output dataset.



Chapter 9

Conclusion

This dissertation discusses the design, inputs, and validation of a Passive Magnetic Attitude Control system for small satellites. Passive Magnetic Attitude Control (PMAC) is useful for a variety of satellite missions as it is simple to install, low-cost, and does not require on-orbit computation. Additionally, some science missions may be aided by alignment with the local magnetic field.

However, the drawbacks of this attitude control method are twofold. First, a lack of understanding of the basic concepts behind PMAC has sometimes led to poor control system design which negatively affects the attitude performance. Second, a lack of accurate simulation has artificially limited the satellite missions which may use PMAC. The ability to accurately predict settling times is key for a satellite mission which relies upon a stable attitude, especially for small satellites which typically have a short mission duration. This dissertation aims to counter the drawbacks listed above.

9.1 Summary

Chapter 2 outlined the basic theory of PMAC, which is a marriage of rigid body dynamics and ferromagnetism; an overview of both of these components is described. Chapter 3 outlined the development history of PMAC. The mission history covers PMAC satellites from 1960 to 2012. Previous attempts at analytical and numeric attitude models are reviewed.

Chapter 4 introduced the Colorado Student Space Weather Experiment (CSSWE), a 3U

CubeSat for space weather investigation which is an example of good PMAC design complementing a science mission. We have access to the on-orbit data; this is useful because CSSWE used a PMAC system. Chapter 5 discussed best practices for designing a small satellite PMAC system; the CSSWE design is used as an example when applicable. Chapter 6 developed a Multiplicative Extended Kalman Filter specially suited for attitude determination of PMAC systems. Simulation- and empirical-based filter tuning is performed before the filter is applied to the CSSWE attitude measurements. After on-orbit calibration and on-orbit magnetic moment fitting, the MEKF output regularly achieved a 3σ uncertainty of 4° or less using magnetometer and partial sun vector measurements at a six second period without a rate gyro. The filtered data validated the CSSWE PMAC design, showing attitude settling within 15° of the local magnetic field after 7 days.

Chapter 7 outlined magnetic measurements which are key inputs to an accurate dynamics simulation. A Helmholtz cage was designed and built for magnetic testing. The cage was used in parallel with other hardware and control software to characterize the CSSWE bar magnet and hysteresis rods. Testing showed that for the CSSWE magnetic material distribution, the system-level hysteresis rod positioning effects are negligible compared to the variation in individual rod dampening ability. Simulated hysteresis loop parameters were derived using a least-squares fit to the experimentally-measured hysteresis loop data collected at magnetizing field cycle amplitudes of ± 10 and ± 20 A/m. The best-fit hysteresis parameters defined a loop with a ± 20 A/m energy dissipation ability of $0.0448 \text{ J}\cdot\text{m}^{-3}$ per cycle, which is approximately 100 times less than the area based upon the material datasheet closed magnetic circuit parameters.

Chapter 8 described a simulation capable of accurately modeling PMAC dynamics. Simulation comparisons over a variety of time steps showed that a fourth-order fixed Runge-Kutta integrator at 0.1s time step has adequate energy conservation performance. Further analysis showed that for certain initial conditions, the absolute attitude cannot be modeled accurately at the considered time steps due to chaotic motion encountered during the settling period. However, the settling time was found to converge as simulation accuracy improved, though small variations in the initial conditions caused the true settling time to be decrease substantially for approximately 20% of the

considered sample size. However, the simulations with abnormally decreased settling times could be identified by their dynamic settling response. In this way, the true worst-case settling time may be predicted accurately by dynamic simulation. The full simulation was run given initial inputs from the CSSWE satellite properties and the early-mission filtered attitude; the predicted settling time was within 20% error of both the higher-order integrator simulation output and the filtered on-orbit attitude data.

9.2 Recommendations

The most obvious recommendation is to improve the hysteresis model. An ideal hysteresis model would be able to model low and high cycle amplitudes equally well for the same loop parameters. It may be that more hysteresis parameters are needed to characterize the empirical data at various magnetization cycle amplitudes. If the mission cannot afford to overestimate the predicted settling time by 20-30% to account for the simulation error, the hysteresis model is the simulation component which should be improved first. Any discontinuities which occur during the settling (such as mechanical deployment) should be investigated as another possibility to improve the simulation.

The satellite magnetic moment should be measured in the fully deployed configuration. This is especially important if the satellite does not have a rate gyro and a dynamic model will be used in the attitude determination. Measuring the hard magnet alone is not sufficient as there can be significant sources of magnetic offset within the satellite. In some cases, as with the CSSWE antenna, the magnetic moment of the spacecraft can change between the stowed and deployed configurations.

The uncertainty of the hysteresis measurement could be decreased by building a sense coil with more turns. We do not recommend using a wire gauge higher than 36 AWG; instead the wire turns should be layered two or more times. This decreased uncertainty will allow the testing to be performed at lower cycle amplitudes. Because of the limitations of the Flatley model, these low cycle amplitudes are key to accurate simulation.

If more computer processing power is available, further simulations are possible. A deeper understanding of the underlying dynamics could be discovered by using the described simulation to calculate the attitude performance when given initial conditions over the entire attitude sphere. Also, the full simulation could undergo the same perturbation analysis that was performed on the simplified simulation. Such studies would likely take computational years to run (the analysis of Section 8.3.4 alone took 68 computational days to run), but given enough processors (or more time for Moore's law to take effect), further analysis is possible.

Bibliography

- [1] 50th Anniversary of First Transit to Orbit. Online Resource. <http://forum.nasaspaceflight.com/index.php?topic=21227.0>, Accessed Jan 2012.
- [2] NOAA National Geophysical Data Center: Geomagnetic kp and ap Indices. Online Resource. ftp://ftp.ngdc.noaa.gov/STP/GEOMAGNETIC_DATA/INDICES/KP_AP/2012, Accessed Dec 2013.
- [3] E. A. Alshina, E. M. Zaks, and N. N. Kalitkin. Optimal first- to sixth-order accurate Runge-Kutta schemes. Computational Mathematics and Mathematical Physics, 48(3):395–405, February 2011.
- [4] M L Battagliere, F Santoni, F Piergentili, M Ovchinnikov, and F Graziani. Passive magnetic attitude stabilization system of the EduSAT microsatellite. Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering, 224(10):1097–1106, January 2010.
- [5] Maria L. Battagliere, Fabio Santoni, Michael Yu. Ovchinnikov, and Filippo Graziani. Hysteresis Rods In the Passive Magnetic Stabilization System for University Micro and Nanosatellites. In 59th International Astronautical Congress, pages 5203–5212, 2008.
- [6] Ciaran Beggan and Kathy Whaler. Forecasting secular variation using core flows. Earth Planets Space, 62(10):821–828, December 2010.
- [7] Tom Bleier, Paul Clarke, Jamie Cutler, Louis DeMartini, Clark Dunson, Scott Flagg, Allen Lorenz, and Eric Tapio. QuakeSat Lessons Learned : Notes from the Development of a Triple CubeSat. Technical report, 2004.
- [8] Jeff P. Anderson & Richard J. Blotzer. Permeability and Hysteresis Measurement. In The Measurement, Instrumentation and Sensors Handbook on CD-ROM, chapter Permeabili. CRC Press LLC, 1999.
- [9] Carl O Bostrom and Donald J Williams. The Space Environment. John Hopins APL Technical Digest, 19(1):43–52, 1998.
- [10] Richard M. Bozorth. Ferromagnetism. Wiley-IEEE Press, 1 edition, 1993.
- [11] Roland Burton, Stephen Rock, John Springmann, and James Culter. Online Attitude Determination of a Passively Magnetically Stabilized Spacecraft. In Proceedings of the 23rd AAS/AIAA Space Flight Mechanics Meeting, pages 1–18, 2013.

- [12] Roland Burton, Joseph Starek, and Stephen Rock. A New Method for Simulating the Attitude Dynamics of Passively Magnetically Stabilized Spacecraft. In Proceedings of the 22nd AAS/AIAA Space Flight Mechanics Meeting, pages 1–14, 2012.
- [13] California Polytechnic State University. CubeSat Design Specification. Rev. 12.
- [14] Nalin A. Chaturvedi, Taeyoung Lee, Melvin Leok, and N. Harris McClamroch. Nonlinear Dynamics of the 3D Pendulum. Journal of Nonlinear Science, 21(1):3–32, September 2010.
- [15] Y Chen. The damped angular motion of a magnetically oriented satellite. Journal of the Franklin Institute, 280(4):291–306, October 1965.
- [16] John L. Crassidis, F. Landis Markley, and Yang Cheng. Survey of Nonlinear Attitude Estimation Methods. Journal of Guidance, Control, and Dynamics, 30(1):12–28, January 2007.
- [17] B.D Cullity and C.D. Graham. Introduction to Magnetic Materials. Wiley-IEEE Press, 2 edition, 2008.
- [18] James Cutler, Matthew Bennett, Andrew Klesh, Hasan Bahcivan, and Rick Doe. The Radio Aurora Explorer: A Bistatic Radar Mission to Measure Space Weather Phenomenon. In AIAA/USU Conference on Small Satellites, pages 1–8, 2010.
- [19] James W. Cutler and Hasan Bahcivan. Radio Aurora Explorer: A Mission Overview. Journal of Spacecraft and Rockets, pages 1–9, June 2013.
- [20] Robert J Danchik. An Overview of Transit Development. John Hopkins APL Technical Digest, 19(1):18–26, 1998.
- [21] John Dormand. Numerical Methods for Differential Equations. CRC Press, Inc., 1 edition, 1996.
- [22] K. Stopfkuchen F. Mesch, G. Schweizer. Investigation of Earth Satellites with Magnetic Attitude Stabilization. In Proceedings of the First IFAC Symposium of Automatic Control in the Peaceful Uses of Space, pages 176–210, 1966.
- [23] C. C. Finlay, S. Maus, C. D. Beggan, T. N. Bondar, A. Chambodut, T. A. Chernova, A. Chuliat, V. P. Golovkov, B. Hamilton, M. Hamoudi, R. Holme, G. Hulot, W. Kuang, B. Langlais, V. Lesur, F. J. Lowes, H. Lühr, S. Macmillan, M. Manda, S. McLean, C. Manoj, M. Menvielle, I. Michaelis, N. Olsen, J. Rauberg, M. Rother, T. J. Sabaka, a. Tangborn, L. Tø ffer Clausen, E. Thébault, A. W. P. Thomson, I. Wardinski, Z. Wei, and T. I. Zvereva. International Geomagnetic Reference Field: the eleventh generation. Geophysical Journal International, 183(3):1216–1230, December 2010.
- [24] Fausto Fiorillo, Fabio Santoni, Enzo Ferrara, Maria Libera Battagliere, Oriano Bottauscio, and Filippo Graziani. Soft Magnets for Passive Attitude Stabilization of Small Satellites. IEEE Transactions on Magnetics, 46(2):670–673, February 2010.
- [25] Robert E. Fischell. Magnetic and Gravity Attitude Stabilization of Earth Satellites. Technical report, The Johns Hopkins University Applied Physics Laboratory, Silver Spring, MD, 1961.
- [26] Thomas W. Flatley and Debra A. Henretty. A magnetic hysteresis model. Technical Report N95-27801, NASA GSFC, 1995.

- [27] Abigail Foerstner. What Van Allen found in space. Bulletin of the Atomic Scientists, 63(4):58–65,78, July 2007.
- [28] C.C. Foster and G.H. Elkaim. Extension of a two-step calibration methodology to include nonorthogonal sensor axes. Aerospace and Electronic Systems, 44(3), 2008.
- [29] T.A. Fritz. The passive magnetic orientation of satellite injun 3. Technical Report 65-21, Iowa University, 1965.
- [30] Carolin Fruh, Thomas Kelec, and Moriba Jah. Attitude Dynamics Simulation of MLI Space Debris Objects in Geosynchronous Earth Orbits. In AIAA/AAS Astrodynamics Specialist Conference, number August, pages 1–18, 2012.
- [31] Demoz Gebre-Egziabher, Gabriel H. Elkaim, J. David Powell, and Bradford W. Parkinson. Calibration of Strapdown Magnetometers in the Magnetic Field Domain. Journal of Aerospace Engineering, 19(2):87, 2006.
- [32] Filippo Graziani, Fabrizio Piergentili, and Fabio Santoni. A space standards application to university-class microsatellites: The UNISAT experience. Acta Astronautica, 66:1534–1543, May 2010.
- [33] John Gurland and Ram Tripathi. A Simple Approximation for Unbiased Estimation of the Standard Deviation. The American Statistician, 25(4):30–32, 1971.
- [34] E. Hairer, S.P. Nørsett, and G. Wanner. Solving Ordinary Differential Equations I: Nonstiff Problems. Springer-Verlag, 2 edition, 1993.
- [35] Felix R. Hoots and Ronald L. Roehrich. Spacetrack report no. 3 models for propagation of norad element sets. Technical Report 3, 1980.
- [36] Peter Hughes. Spacecraft Attitude Dynamics. John Wiley & Sons, 1 edition, 1986.
- [37] Arieh Iserles, Hans Z Munthe-kaas, Syvert Norsett, and Antonella Zanna. Lie-group methods. Acta Numerica, pages 1–148, 2005.
- [38] Klas Johnsson. Munin - a swedish nanosatellite. In 4th IAA Symposium on Small Satellites for Earth Observation, number IAA-B4-0406P, 2003.
- [39] R. W. Kammuller. Nonlinear Resonant Roll Motion of Magnetically Oriented Satellites. AIAA Journal, 9(4):582–588, 1971.
- [40] R. W. Kammuller. Roll Resonance and Passive Roll Control of Magnetically Stabilized Satellites. AIAA Journal, 10(2):129–136, 1972.
- [41] T.S. Kelso. Visually Observing Earth Satellites. Satellite Times, 3(1):80–82, September 1996.
- [42] D. J. Knipp, T. Matsuo, L. Kilcommons, A. Richmond, B. Anderson, H. Korth, R. Redmon, B. Mero, and N. Parrish. Comparison of magnetic perturbation data from LEO satellite constellations: Statistics of DMSP and AMPERE. Space Weather, 12, January 2014.
- [43] Renjith Kumar, Daniel Mazanek, and Michael Heck. Simulation and Shuttle Hitchhiker Validation of Passive Satellite Aerostabilization. Journal of Spacecraft and Rockets, 32(5):806–811, 1995.

- [44] Daeyoung Lee, John C. Springmann, Sara C. Spangelo, and James W. Cutler. Satellite Dynamics Simulator Development Using Lie Group Variational Integrator. In Proceedings of the AIAA Modeling and Simulation Technologies Conference, pages 1–20, 2011.
- [45] Taeyoung Lee, N Harris Mcclamroch, and Melvin Leok. A Lie Group Variational Integrator for the Attitude Dynamics of a Rigid Body with Applications to the 3D Pendulum. IEEE Conference on Control Applications, pages 962–967, 2005.
- [46] E.J. Lefferts, F.L. Markley, and M.D. Shuster. Kalman Filtering for Spacecraft Attitude Estimation. Journal of Guidance, Control, and Dynamics, 5(5):417–429, 1982.
- [47] Jean-Francois Levesque. Passive magnetic attitude stabilization using hysteresis materials. Technical Report SIgMA-PU-006-UdeS, Université de Sherbrooke, 1995.
- [48] X. Li, Q. Schiller, L. Blum, S. Califf, H. Zhao, W. Tu, D. L. Turner, D. Gerhardt, S. Palo, S. Kanekal, D. N. Baker, J. Fennell, J. B. Blake, M. Looper, G. D. Reeves, and H. Spence. First results from CSSWE CubeSat: Characteristics of relativistic electrons in the near-Earth environment during the October 2012 magnetic storms. Journal of Geophysical Research: Space Physics, 118, October 2013.
- [49] Xinlin Li, Scott Palo, Rick Kohnert, Lauren Blum, David Gerhardt, Quintin Schiller, and Sam Califf. Small Mission Accomplished by Students-Big Impact on Space Weather Research. Space Weather, 11(2):55–56, February 2013.
- [50] Xinlin Li, Scott Palo, Rick Kohnert, David Gerhardt, Lauren Blum, Quintin Schiller, Drew Turner, Weichao Tu, Nathan Sheiko, and Chris S. Cooper. Colorado student space weather experiment: Differential flux measurements of energetic particles in a highly inclined low earth orbit. In Dynamics of the Earth’s Radiation Belts and Inner Magnetosphere, volume 199 of Geophysical Monograph Series, pages 385–404, 2012.
- [51] C. a. Loewe and G. W. Prölss. Classification and mean behavior of magnetic storms. Journal of Geophysical Research, 102(A7):14,209–14,213, 1997.
- [52] Matthew Long, Allen Lorenz, Greg Rodgers, Eric Tapio, Glenn Tran, Keoki Jackson, Robert Twiggs, and Thomas Bleier. A CubeSat Derived Design for a Unique Academic Research Mission In Earthquake Signature Detection. In USU Conference on Small Satellites, pages 1–17, 2002.
- [53] B Maag and K. Stopfkuchen. Flight Experience with the Passive Magnetic Attitude Control System of the First German Satellite Azur. In Proceedings of the 4th IFAC Symposium on automatic control in space, page 25, Friedrichshafen, Germany, 1971. Control Division, Dornier System GmbH.
- [54] F Landis Markley. Attitude Error Representations for Kalman Filtering. Journal of Guidance, Control, and Dynamics, 26(2), 2003.
- [55] Göran Marklund, Mats André, Rickard Lundin, and Sven Grahn. The Swedish Small Satellite Program for Space Plasma Investigations. Space Science Reviews, 111(3/4):377–413, April 2004.

- [56] B.J. O'Brien, C. D. Laughlin, and D. A. Guernett. High-Latitude Geophysical Studies with Satellite Injun 3 1. Description of the Satellite. Journal of Geophysical Research, 69(1):1–12, 1964.
- [57] M. Yu. Ovchinnikov and V.I. Penkov. Passive magnetic attitude control system for the munin nanosatellite. Cosmic Research, 40(2):142–156, 2002.
- [58] G. Park, S. Seagraves, and N. H. McClamroch. A Dynamic Model of a Passive Magnetic Attitude Control System for the RAX Nanosatellite. In AIAA Guidance, Navigation and Control Conference, number August, 2010.
- [59] J. M. Picone, A.E. Hedin, and D.P. Drob. NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues. Journal of Geophysical Research, 107(A12):1–16, 2002.
- [60] Gerd W. Prölss. Physics of the Earth's Space Environment. Springer Press, 1 edition, 2004.
- [61] Samir Rawashdeh. Passive Attitude Stabilization for Small Satellites. University of Kentucky Masters Thesis, 2009.
- [62] Fabio Santoni and Fabrizio Piergentili. Unisat-3 Attitude Determination Using Solar Panel and Magnetometer Data. In International Astronautical Federation - 56th International Astronautical Congress, pages 2812–2819, 2005.
- [63] Fabio Santoni and Mauro Zelli. Passive magnetic attitude stabilization of the unisat-4 microsatellite. Acta Astronautica, 65:792–803, 2009.
- [64] S. Schalkowsky and M. Harris. NASA Space Vehicle Design Criteria (Guidance and Control): Spacecraft Magnetic Torques. NASA SP-8018. Technical report, 1969.
- [65] Hanspeter Schaub and John L. Junkins. Analytical Mechanics of Space Systems. American Institute of Aeronautics and Astronautics, Inc., 1 edition, 2003.
- [66] Quintin G. Schiller, Abhishek Mahendrakumar, and Xinlin Li. REPTile : A Miniaturized Detector for a CubeSat Mission to Measure Relativistic Particles in Near-Earth Space. In USU Conference on Small Satellites, pages SSC10–VIII–1, 2010.
- [67] Jinglai Shen, Amit K. Sanyal, Nalin A. Chaturvedi, Dennis Bernstein, and Harris Mcclamroch. Dynamics and Control of a 3D Pendulum. In 43rd IEEE Conference On Decision and Control, pages 323–328, 2004.
- [68] Garrett Lee Skrobot and Roland Coelho. ELaNa - Educational Launch of Nanosatellite: Providing Routine RideShare Opportunities. In Small Satellite Conference, 2012.
- [69] Ahmet Sofyali and Rustem Aslan. Magnetic attitude control of small satellites: A survey of applications and a domestic example. Technical report, Istanbul Technical University, 2011.
- [70] John C Springmann. On-Orbit Calibration of Photodiodes for Attitude Determination. In Small Satellite Conference, number March, pages SSC13–VIII–1, 2013.
- [71] John C Springmann and James W Cutler. Initial Attitude Analysis of the RAX Satellite. In Proceedings of the AIAA/AAS Astrodynamics Specialist Conference, pages 1–20, 2011.

- [72] John C. Springmann and James W. Cutler. Attitude-Independent Magnetometer Calibration with Time-Varying Bias. Journal of Guidance, Control, and Dynamics, 35(4):1080–1088, July 2012.
- [73] John C Springmann, Benjamin P Kempke, and James W Cutler. Initial Flight Results of the RAX-2 Satellite. In Small Sat Conference, 2012.
- [74] M Swartwout, C Kitts, R Twiggs, T Kenny, B Raysmith, R Lu, K Stattenfield, and F Pranjaya. Mission results for Sapphire, a student-built satellite. Acta Astronautica, 62(8-9):521–538, April 2008.
- [75] Hank Heidt; Jordi Puig-Suari; Augustus Moore; Shinichi Nakasuka; Robert Twiggs. Cubesat: A new generation of picosatellite for education and industry low-cost space experimentation. In Small Satellite Conference, SSC00-V-5. AIAA/USU, 2000.
- [76] Robert Twiggs. Space System Developments at Stanford University - From Launch Experience of Microsatellites to the Proposed Future use of Picosatellites. In SPIE, volume 4136, pages 79–86, 2000.
- [77] David A. Vallado. Fundamentals of Astrodynamics and Applications. Microcosm Press, 2 edition, 2001.
- [78] David A. Vallado. Fundamentals of Astrodynamics and Applications. Microcosm Press, 3 edition, 2007.
- [79] David A. Vallado, Paul Crawford, Richard Hujak, and T. S. Kelso. Revisiting Spacetrack Report #3, volume 6753. American Institute of Aeronautics and Astronautics, 2006.
- [80] R.E. Walpole, R.H. Myers, S.L. Myers, and K. Ye. Probability and Statistics for Engineers and Scientists. Pearson Prentice Hall, 8 edition, 2007.
- [81] Hugh D. Young and Roger A. Freedman. University Physics. Addison Wesley, 11 edition, 2003.

Appendix A

Notation

- 3×1 vectors are represented in **bold**. Example: $\boldsymbol{\omega}$
- The scalar component of a vector is not in bold and has a subscript. Example: ω_x
- 3×3 matrices are shown in [brackets]. Example: $[R]$
- The absolute value of a scalar is shown using one vertical bar on either side of the |variable|. Example: $|B_x|$
- The magnitude of a vector is shown with two vertical bars on either side of the $||\mathbf{variable}||$. Example: $||\mathbf{B}||$
- The 3×3 identity matrix is represented by $[I_{3 \times 3}]$.
- The transpose of a matrix is represented by a superscript T. Example: $[R]^T$
- The trace matrix operation is represented by $\text{tr}()$. Example: $\text{tr}([R])$
- The reference frame of a vector is represented by superscript calligraphy letter before the vector. Example: ${}^{\mathcal{I}}\mathbf{r}$
- The skew-symmetric matrix operator is represented by brackets around the variable and a cross product within the brackets. Example: $[\boldsymbol{\omega} \times]$
- The inertial time derivative $\frac{{}^{\mathcal{I}}\text{d}}{\text{d}t}$ of a variable is represented by a dot above the variable. Example: $\frac{{}^{\mathcal{I}}\text{d}}{\text{d}t}\boldsymbol{\omega} = \dot{\boldsymbol{\omega}}$

Appendix B

Explicit Runge-Kutta Integrator Definitions

The family of explicit Runge-Kutta numeric integrators is generalized as follows (repeated from Section 8.1.7.1):

$$y_{n+1} = y_n + h \sum_{i=1}^s b_i k_i \quad (8.26)$$

where

$$\begin{aligned} k_1 &= f(t_n, y_n) \\ k_2 &= f(t_n + c_2 h, y_n + a_{21} k_1) \\ k_3 &= f(t_n + c_3 h, y_n + a_{31} k_1 + a_{32} k_2) \\ &\vdots \\ k_s &= f(t_n + c_s h, y_n + a_{s1} k_1 + a_{s2} k_2 + \dots + a_{s,s-1} k_{s-1}). \end{aligned}$$

A specific Runge-Kutta integrator are is given by its Butcher tableau, which is a standard form of presenting the coefficients used by Equation 8.26. The general form of a Butcher tableau for an explicit Runge-Kutta integrator is shown in Table B.1. The Butcher tableau of each Runge-Kutta integrator used within this dissertation is shown in Tables B.2 through B.7.

Table B.1: The general form of the Butcher tableau for explicit Runge-Kutta Methods [34].

0				
c_2	a_{21}			
c_3	a_{31}	a_{32}		
\vdots	\vdots	\vdots	\ddots	
c_s	a_{s1}	a_{s2}	\dots	$a_{s,s-1}$
	b_1	b_2	\dots	b_{s-1}

Table B.2: The Butcher tableau for explicit fixed RK2 (midpoint method) [34].

0		
1/2	1/2	
	0	1

Table B.3: The Butcher tableau for explicit fixed RK3 (Kutta method) [21].

0			
1/2	1/2		
1	-1	2	
	1/6	2/3	1/6

Table B.4: The Butcher tableau for explicit fixed RK4 (Runge-Kutta method) [34].

0				
1/2	1/2			
1/2	0	1/2		
1	0	0	1	
	1/6	2/6	2/6	1/6

Table B.5: The Butcher tableau for explicit fixed RK5 (fixed Dormand-Prince method) [34].

0						
1/5	1/5					
3/10	3/40	9/40				
4/5	44/45	-56/15	32/9			
8/9	19372/6561	-25360/2187	64448/6561	-212/729		
1	9017/3168	-355/33	46732/5247	49/176	-5103/18656	
	35/384	0	500/1113	125/192	-2187/6784	11/84

