MAGNETIC ATTITUDE CONTROL OF LAICE SATELLITE
WITH AERODYNAMIC STABILIZATION

BY

ERIK KROEKER

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor in Philosophy in Aerospace Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2017

Urbana, Illinois

Doctoral Committee:

    Professor Victoria Coverstone, Chair
    Assistant Professor Koki Ho
    Associate Professor Harrison Kim
    Professor Gary Swenson

ABSTRACT

The Lower Atmosphere Ionosphere Coupling Experiment (LAICE) is a NanoSatellite which will be performing in-situ measurements of neutral and ion densities in the mesosphere, lower thermosphere, and ionosphere and correlating them to measurements of gravity waves in the lower atmosphere. The satellite is based on a new 6U CubeSat form factor based on the IlliniSat-2 bus developed at the University of Illinois at Urbana-Champaign. The satellite's payloads need to be oriented such that three instruments are oriented along the velocity direction of the satellite, while a fourth instrument is pointed towards nadir. The attitude determination and control system must achieve the attitude pointing requirements (5° from nominal attitude) with minimum of cost and low power. The satellite will therefore make use of magnetic torque coils augmented with aerodynamic stabilization to accomplish the mission attitude control requirements. The proposed control method relies on passive aerodynamic stabilization of the spacecraft to maintain pointing in the satellite normal frame. The aerodynamic stabilization reduces the dimensionality of the magnetic attitude control to a one-dimensional problem. The major contributions of this work include the development of an object-oriented attitude control library with which to program flight code as well as simulate satellite dynamics to validate the flight code; the development and tuning of an Extended Kalman Filter for attitude determination using low-cost magnetometers and MEMS gyros; the development and tuning of a hybrid attitude control algorithm for magnetic torque coils which can reliably detumble and reorient the satellite; the development of an efficient graphical drag model for computing aerodynamic torques; and the novel use of a tri-axial Helmholtz cage for performing a Hardware-in-Loop optimization of the coupled attitude determination and control system.

ACKNOWLEDGMENTS

Finally, I would like to recognize the unending support of my family. My parents, from whom I have learned the value of hard work and perseverance; my brother and sister, who both inspire me on a daily basis; my wife Samantha, who has been my greatest friend and has never once failed to express her belief in me when I needed it most; this work, I dedicated to you all.

TABLE OF CONTENTS

CHAPTER 1
INTRODUCTION

1.1 Lower Atmosphere/Ionosphere Coupling Experiment

The Lower Atmosphere/Ionosphere Coupling Experiment (LAICE) CubeSat mission will study the interaction of atmospheric gravity waves generated by weather systems in the lower atmosphere with the mesosphere, lower thermosphere, and ionosphere (MLTI)[1]. Specifically, LAICE will observe the energy and momentum delivered by these waves and attempt to connect the wave sources and the wave effects in three widely different altitude ranges, substantially adding to our knowledge of coupling processes between these atmospheric regions. The mission is targeting a Low-Earth Orbit (LEO) deployment from the International Space Station (ISS). The deployment will occur from an altitude of approximately 400km altitude. As the spacecraft descends, it will be oriented with three of its four payloads oriented into the satellite velocity direction, with its remaining payload pointing along the nadir direction. To achieve this on a relatively small satellite bus, a magnetic-only control system will be used. At low altitudes, the small magnetic torque generated will be dominated by aerodynamic forces. As such, the satellite has been designed in such a way that it will be stabilized in the satellite velocity direction allowing the magnetic controlling to resolve the remaining component of the attitude. To validate the attitude determination and control system (ADCS), a hardware-in-loop simulation is conducted using a three-axis Helmholtz cage. In this work, the attitude control system will be examined and results will be shown for the hardware-in-loop validation of the ADCS.

1.2 LAICE Mission Overview

The LAICE satellite is built around a 6U CubeSat bus developed at the University of Illinois at Urbana-Champaign[2]. The satellite contains four-instrument payloads: the retarding potential

analyzer (RPA) will provide in-situ ion density and temperature measurements; a four-channel photometer will measure density and temperature variations in the mesosphere through observations of O2 (0, 0) atmospheric band and O2 Herzberg I band airglows; and two pressure sensors, the Space Neutral Pressure Instrument (SNeuPI) and the LAICE Ionization gauge Neutral Atmosphere Sensor (LINAS), will provide neutral density measurements[1]. The objective of LAICE is to provide a cost-effective approach to measuring these low altitude in-situ parameters.
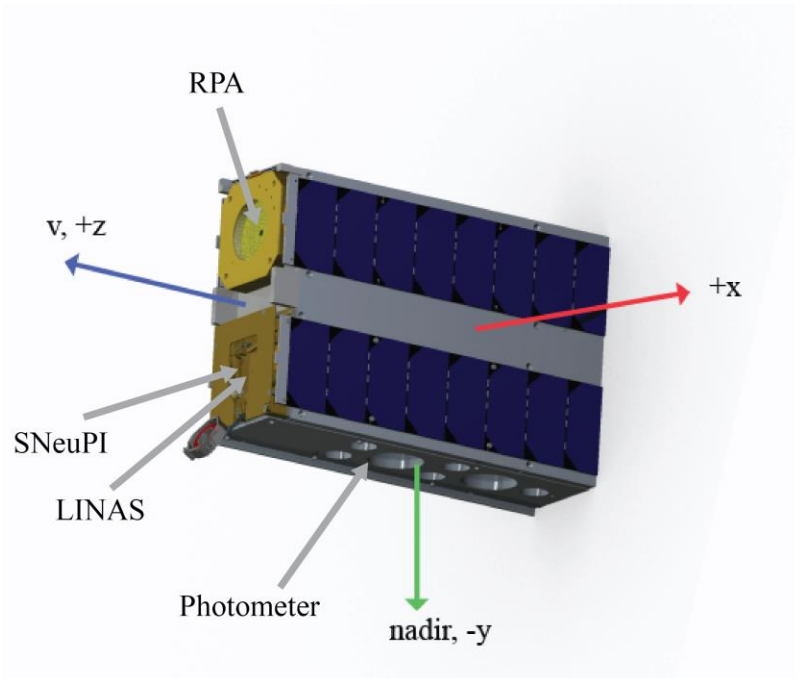
Figure 1. Rendering of the LAICE Satellite in its deployed configuration, aligned with the science orientation with the +z face in track with the velocity, and the –y face along the nadir direction.

As stated, the target deployment altitude is approximately 400km altitude. The satellite will be deployed from a CubeSat deployer mounted to the Kibo module of the ISS. Upon deployment, the satellite may be subjected to deployment induced angular tip-off rates. These rates are relatively small with a maximum expected tip-off rate of 5°/sec for a 6U satellite such as LAICE[3]. Even though the atmospheric density is low, the spacecraft will begin to descend due to atmospheric drag. LAICE will lose altitude more quickly than larger satellites owing to its low

mass (14kg) to surface-area ratio and thus poor ballistic coefficient. Over a period of 9-12 months, the satellite will descend to its target altitude for science operations (220-300km). The satellite will continue to lose altitude at an increasing rate as the density of the atmosphere increases exponentially; however, the stabilizing effect of the aerodynamic forces will improve the control over the spacecraft attitude during the peak science period of the mission. Once the satellite descends below the 220km altitude, the orbit decay becomes extremely rapid. At 190km altitude, simulations indicate that the satellite will not be able to complete additional orbits and reenters.

When recording science data, the instruments must be oriented such that 1) the RPA, SNeuPI, and LINAS instruments be oriented in track with the velocity, or "ram", direction, and 2) the photometer payload is oriented along the nadir direction. Figure 1 demonstrates the science orientation of the LAICE satellite. For a non-circular orbit, it is not possible to maintain both the ram and nadir facing orientation simultaneously. The ram-facing payloads (RPA, SNeuPI, and LINAS) are more sensitive to misalignment from the velocity direction; as such, the orientation of the ram-facing payloads is considered the principal axis for determining pointing accuracy. In order to perform adequate data collection, it is necessary to maintain a pointing accuracy of 5° from velocity direction; the nadir facing payloads must be oriented within 5° of the projected nadir vector onto the plane perpendicular to velocity vector. A pointing knowledge better than 2° is required in order to correlate data. These two requirements must only be met during science operation.
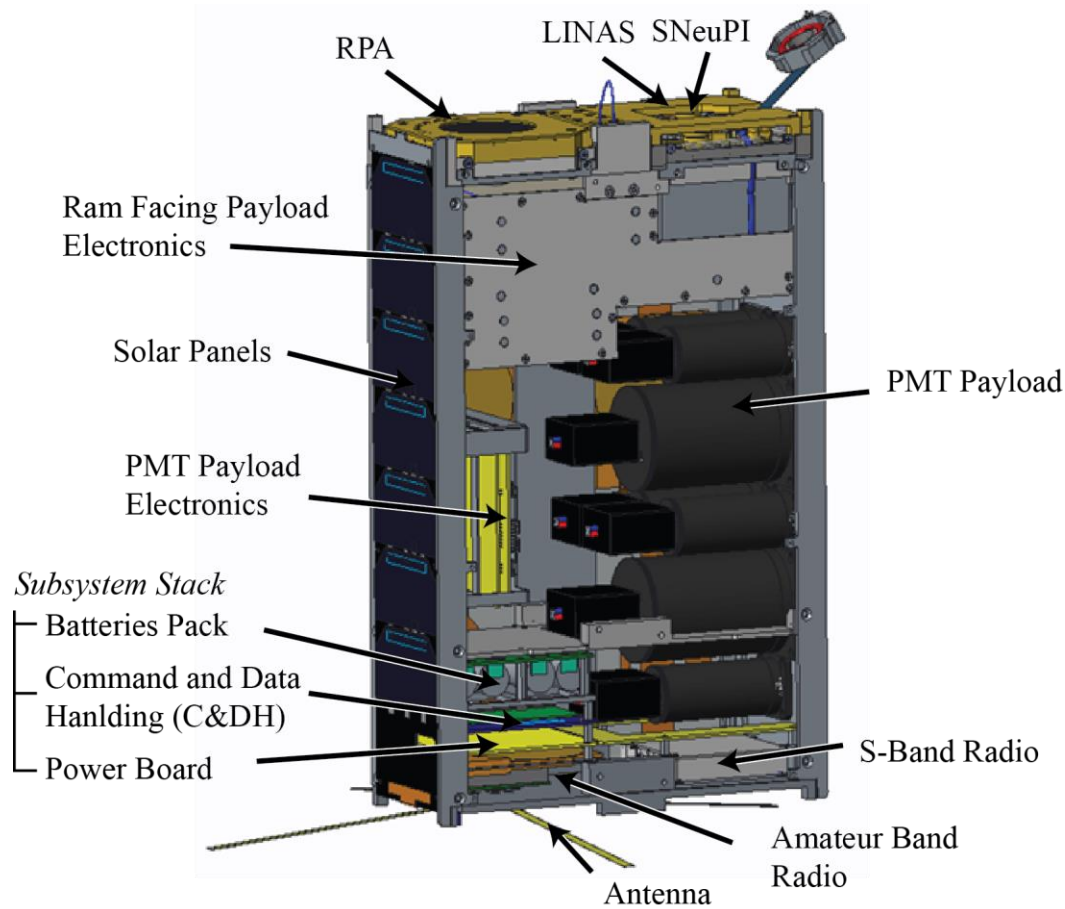
Figure 2: LAICE bus (with -X solar arrays and structure removed for clarity)

To facilitate attitude determination and control, a magnetic-only determination has been implemented. In addition, the center of mass of the spacecraft has been moved forward of the geometric center (+z direction) to give the satellite a positive static margin improving aerodynamic stability in the ram direction. These two design features, when used in concert, will provide the necessary mechanisms for three-axis attitude control about the science orientation.

The attitude determination system consists of a three systems, namely: magnetometers, coarse sun sensors, and rate gyros. The attitude determination and control algorithms are performed onboard on the satellite's command and data handling (C&DH) system. The algorithms use a custom object-oriented library for managing attitude transformations in a frame/format independent manner.

1.2.1   Design of Onboard Attitude Determination and Control Systems

The design of the onboard hardware is the result of many undergraduate and graduate student contributions over the course of the past 15 years. The majority of these contributions were from the course Introduction to NanoSatellite Design at the University of Illinois at Urbana-Champaign. Other contributions were made through undergraduate and graduate student independent study, aerospace and electrical engineering senior design projects, electrical engineering senior thesis projects, as well as undergraduate volunteer research. The attitude determination and control subsystems (magnetometers, rate gyros, and magnetic torque coils) were largely the work of students from semesters from 2012 to 2016. The design of the solar panels, which serve as both a coarse sun sensor and a harness for the ADCS, was the work of students from semesters from 2013 to 2016.

1.2.1.1   Magnetometers
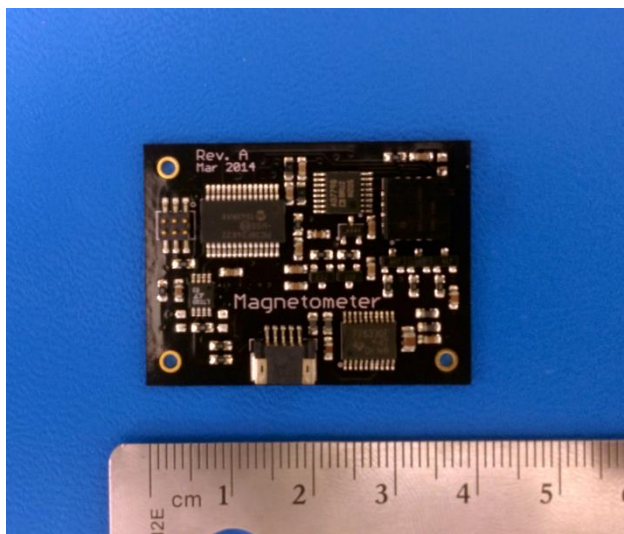
Figure 3. Flight HMC1053 Magnetometer Board.

The satellite is outfitted with up-to six three-axis Honeywell HMC1053 magnetometers[4]. Each magnetometers is installed on an individual board with its own dedicated 16-bit analogue-to-digital converter (ADC) and microprocessor to interface with the command and data handling

(C&DH) system using a serial connection (see Figure 3). When queried by the C&DH system, the magnetometer board measures the field strength of each axis sequentially (X-Y-Z). Each axis takes 5μs to measure; as such, the values are considered to be instantaneous in relation to the typical rotation rates of the satellite.

These inexpensive sensors are relatively imprecise when compared to scientific sensors; however, a calibration and registration algorithm is performed to improve the accuracy as well as determine the rotation from the sensor axes to the satellite body frame. Using the calibration algorithm, the accuracy can be improved to within 0.2% of full resolution, or approximately 0.5° of pointing of the true magnetic field[5]–[7].

The magnetometer boards are mounted on the interior face of the spacecraft solar panels (up to two per panel). Whilst only a single three-axis sensor is required for determination, multiple magnetometers will be used and a sensor fusion algorithm will be implemented to improve the magnetic field measurement accuracy in the presence of system electromagnetic interference.

1.2.1.2 Coarse Sun Sensors

Each solar panel has a small photoresistor mounted to its outer face. Assuming unobstructed light from the sun and a constant solar luminance, the measured intensity from the photoresistor can be correlated to the incident angle to the sun by Equation (1). By using several photoresistors, it is possible to determine the pose of the spacecraft with respect to the sun. The sensor is also used by other systems (such as payloads) to determine whether or not payloads should be shutdown to avoid oversaturation of sensors.

$$I_{observed} = I_{incident}\sin(\alpha) \tag{1}$$

As the photosensors are only mounted to the solar panel faces (and not the +z or –z faces), the satellite will only be able to determine the orientation of the satellite to two angles at best. While the magnitude of the third angle can be assumed based on the missing luminal intensity, it will be impossible for the spacecraft to resolve the sign of the angle. Due to variations in incident intensity (either due to variations in solar intensity or the addition of albedo off of other bodies such as the Earth or Moon), this magnitude of the third component is also less reliable. The lack of ability to resolve the third component makes deterministic determination methods less effective, but does lend itself well to augmenting an Extended Kalman Filter (EKF) determination algorithm if necessary[8]. In practice, the coarse sun sensing ability is used to validate attitude determination system performance in ground analysis and is not used as part of the onboard satellite determination algorithms.

1.2.1.3   Rate Gyros

For rate measurements, a MEMS gyro is installed on a dedicated board communicating over the same serial communication interface as the magnetometers. MEMS gyros are subject to drift, which causes the baseline value to deviate from zero[9]. As such, the reference needs to be calibrated against an external reference. For this purpose, a MEMS gyro is typically paired with an accelerometer[9]. The combined system of gyro and accelerometer is typically referred to as an Inertial Measurement Unit (IMU). Many commercial off-the-shelf (COTS) gyro and accelerometers exists. The challenge in selecting a COTS component for the LAICE satellite is the difference in required sensitivity from typical commercial applications. The most common use for these components is in handheld devices. Handheld devices are subject to much greater linear and angular acceleration and therefore measure over much larger ranges (typically 250°/s and greater full range for angular rate and 10g or greater for linear acceleration). The ±250°/s full

range is far greater than the anticipated 5°/s maximum tipoff rate. MEMS rate gyros specifically for space applications are commercially available, yet only improve the resolution by a factor of two for an order of magnitude greater cost[10]. The alternative, however, are non-COTS solutions such as mechanical gyros or ring-laser gyros which, albeit far more precise, are dramatically more expensive.

The LAICE satellite uses an MPU-6050 Motion Processing Unit. This sensor has 3-axis MEMS gyroscope, a 3-axis accelerometer, and a digital motion processor[11]. The sensor has a 250°/sec full range and, at 16bit resolution of these sensors, it provides a resolutions of up to 131 bits/°/sec (or approximately 0.0076°/sec); however, the reliability of the least significant bits (LSB) degrades due to power supply noise and system vibration. In practice, the resolution is more reliably on the order of 0.1°/s. If the noise is assumed to be zero-mean Gaussian, it is possible to filter the rate information based on the system dynamics[12]. For the purposes of simulation, the gyro is simulated using the attitude rate state with zero-mean, Gaussian noise. The value is also clipped to simulate the digital nature of the device (whereby a value less than the 0.0076°/sec rate would be interpreted as zero).

1.2.1.4   Magnetic Torqueing Coils

To generate the required magnetic moment, LAICE uses three orthogonally-mounted magnetic torqueing coils embedded into circuit boards with a dedicated ADC and microprocessor for command handling (see Figure 4). A magnetic torqueing coil generates a magnetic moment by passing a current through multiple aligned loops of wire. In the presence of an external magnetic field, the magnetic moment generates a torque according to the cross-product of the magnetic moment and the magnetic field. The coils are capable of producing a magnetic moment of 0.2744 Nm/T at peak[13]. Given an on-orbit maximum magnetic field of

4.90E-5 Tesla, this results in a maximum torque of 1.35E-5 Nm if the magnetic moment and magnetic field are orthogonal. The output of each coil is controllable in 128 even spaced steps and the direction of the output is controllable. The current driving the coils can be regulated to within 0.1%. The generated magnetic moment is estimated based on the coil geometry. To verify the generated magnetic moment, a spin table test will be conducted prior to launch.
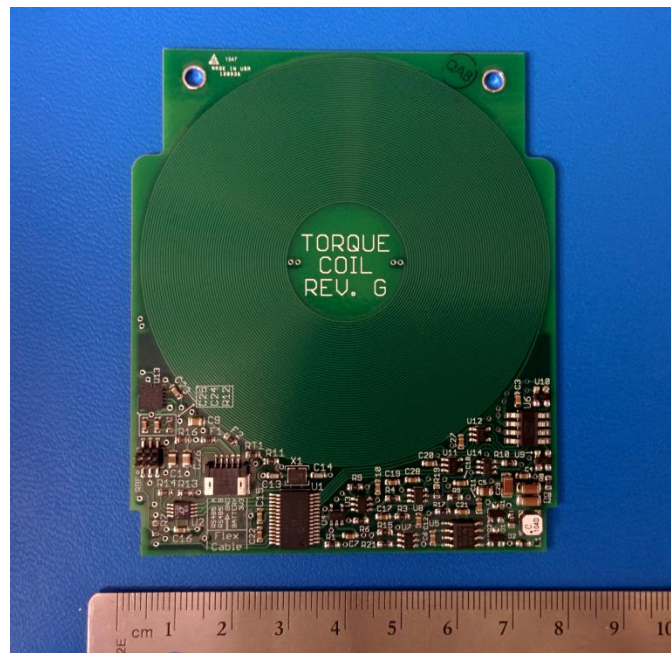
Figure 4. Magnetic Torqueing Coil Board[13].

1.2.1.5 Physical Arrangement of ADCS Components

The ADCS sensors and actuators are principally located on the inside face of the solar panels (see Figure 5 and Figure 6). The sole exception is the z-axis magnetic torque coil which is located in the service stack. The photodiodes are placed at the base of each panel (at the right extremum of the solar panel in Figure 5). The magnetic torque coils are placed centrally on the panels – aligned with the principal axes of the spacecraft's body (see Figure 6). The magnetometer is placed as far from the torque coil and science payloads as possible to avoid electromagnetic interference (EMI) effects. Even at this distance, the magnetometers will

become saturated when the torque coils are active. As a result, the magnetometers will only be used for sensing when the coils are inactive and will be degaussed between measurements. Five such panels will be mounted on the spacecraft (2 on the +x and –x, and one on the +y face).
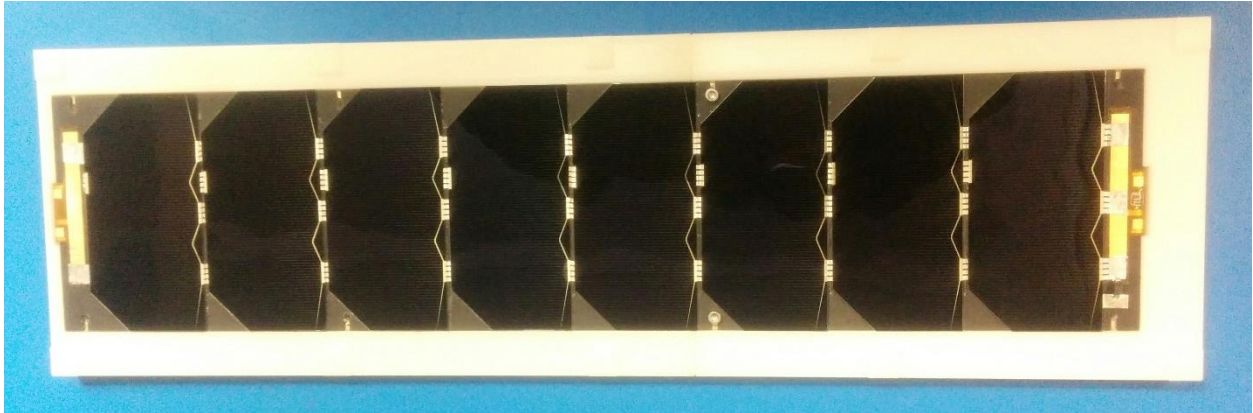
Figure 5: Outward facing side of single solar panel. The white border is a protective tray.
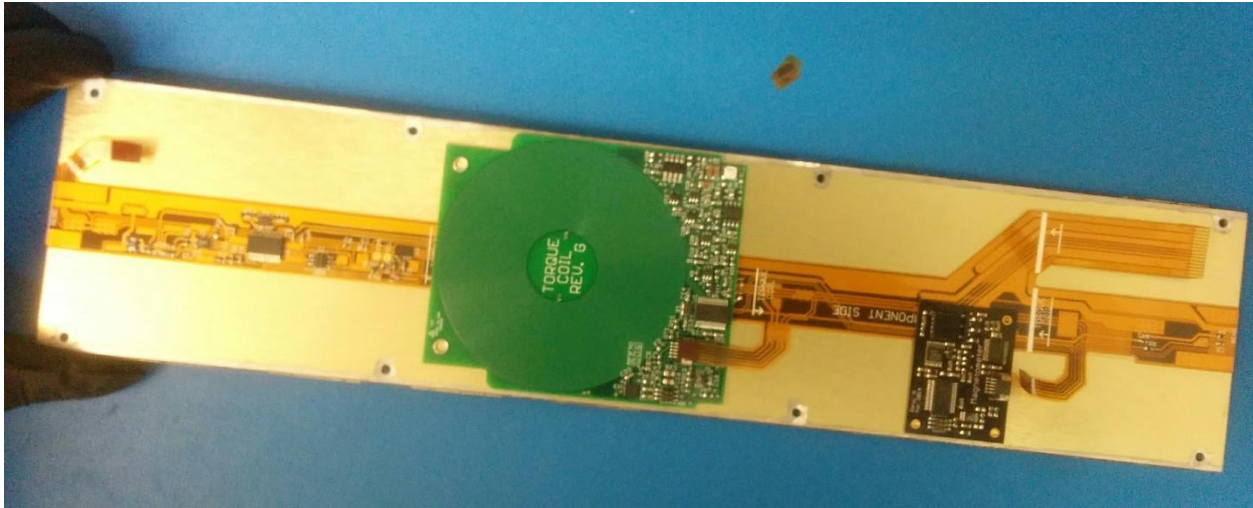
Figure 6: Inward facing side of solar panel with ADCS sensors and actuators. Protective tray has been removed for clarity.

1.3 Research Contributions

The presented work analyzes several aspects of the LAICE CubeSat's Attitude Determination and Control System. The most difficult component of the design, which also make this analysis unique and novel, is creating a high-reliability attitude determination system using low-cost,

COTS components to achieve the science attitude pointing requirements. The present work includes:

1. Design of a low-cost attitude determination and control system. The system aims to use primarily commercial off-the-shelf components.

2. Development of an attitude determination algorithm which is robust to sensor noise. An analysis of multiple magnetic-only approaches are presented and a unique system was designed to allow for globally convergent solutions. This effort required an Extended Kalman Filter (EKF) and subsequent tuning of the filter parameters.

3. Development of an efficient attitude control algorithm for magnetic torqueing control. A unique hybrid switching controller was used to accomplish efficient detumbling of the satellite post-deployment, followed by a proportional-integral control approach for fine pointing. This effort required gain tuning to obtain the desired response characteristics. An analysis of the control response is also presented.

4. Optimization of the coupled attitude determination and control system using a Particle Swarm Optimization (PSO).

5. Analysis of the system performance under external aerodynamic forces. An analytical model of the satellite was developed, as well as a novel GPU graphical method which is better able to deal with occluding geometry.

6. Measure the mass properties (GC and MOI) of the satellite and determine the uncertainty in the mass properties. This is essential to validate the robustness of the system to reasonable errors in the mass properties.

7. Validate the system using Hardware-in-Loop demonstrations of the attitude determination and control system using a novel three axis Helmholtz cage. This same

setup is also used to perform Hardware-in-Loop optimizations of the coupled attitude

determination and control system.

CHAPTER 2
MAGNETIC ATTITUDE DETERMINATION

2.1 Spacecraft Attitude Determination for Low-Earth Orbits

Attitude determination is the process of determining the orientation of a spacecraft with respect to some external frame of reference. It is a necessary step in determining the pointing of any satellite instruments and a prerequisite for reorienting the satellite to a desired orientation. The attitude state, in its simplest representation, can be thought of as the three body-fixed rotations required to transform the spacecraft body-frame of reference to some external reference frame. The attitude can be parameterized in several ways. For the purposes of this work, three parameterizations will be considered: Euler angles, direct cosine matrices (DCM), and quaternions.

2.1.1 Reference Frames for LAICE Mission

The attitude represents the transformation of the spacecraft body frame with respect to some external reference frame. Therefore, when discussing spacecraft attitude, it is necessary to specify the frame in which the attitude is being represented. For the majority of the discussions in this work, five reference frames are of importance: Earth-Center Inertial (ECI), Earth-Centered/Earth-Fixed (ECEF), Perifocal Frame (PWQ), Satellite Normal Frame (NTW), and Hill Frame (RSW). These reference frames are notional concepts, all having several specific implementations defined by either an epoch and/or an orbit. A brief overview of each is presented herein with a discussion of the significance with respect to the LAICE mission.

To define any attitude reference system, it is necessary to form a set of basis vectors. This basis vector set is with respect to reference features. In practice, these frames are defined as orthonormal basis vector sets. Therefore, two orthogonal vectors are sufficient to form a basis as

the third vector will be define by the cross-product of the two and the handedness of the system. Alternatively, it can be defined by a plane and principal direction; the plane normal and principal direction (which typically lies within the plane) are sufficient to define the basis vector set of the frame. For completeness, a reference frame is defined about an origin; however, as the location of the origin is decoupled from the orientation between frames, the origin is no significance to the attitude.
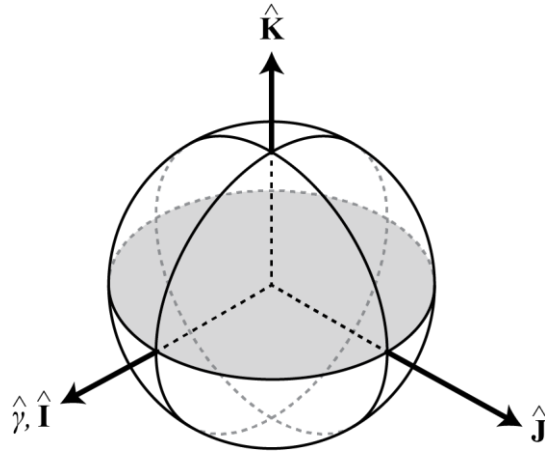
Figure 7: ECI Frame defined by the equatorial normal vector (K) and the First Point of Aries (γ).

The ECI frame is an inertial reference frame centered on the Earth. The Earth's Equator is the fundamental plane and the principal direction is the Vernal equinox. Since the Vernal equinox changes over time, the frame must be defined from an epoch[14]. For the purposes of this paper, the IAU-76/FK5 epoch[15] is used to define the ECI within the context of this paper. The ECI frame is of significance as it an inertia reference frame. Dynamic modelled in this reference, unlike the other frames discussed herein, are independent and separable from the motion of the frame. Unless otherwise specified, representations of the attitude in equations of spacecraft dynamics are with respect to the ECI frame.
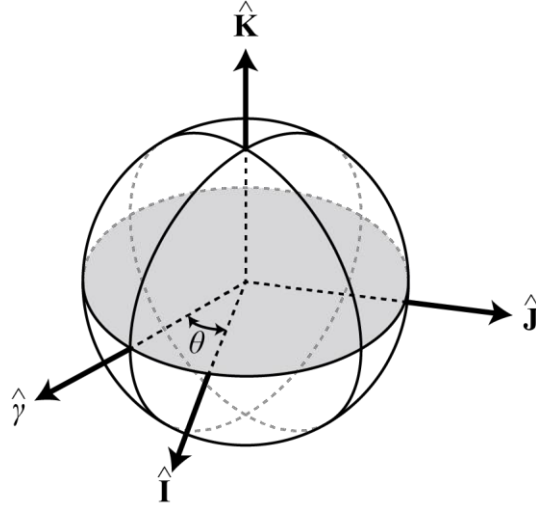
Figure 8: ECEF frame defined by the equatorial normal vector (K) and the Prime Meridian vector (I).

The ECEF frame is similar to the ECI frame in that the frame is centered on the Earth and the Earth's equator serves as the fundamental plane; however, the principal direction is defined through the Greenwich meridian. This couples the motion of the frame to the motion of the planet surface. The ECI and ECEF frames are related through a rotation about the Earth equator normal. The angle of this rotation is equal to the sidereal angle of the planet, θ, again as defined from an epoch. The ECEF frame is useful in that it readily represents values which are defined with respect to the surface of the Earth. The magnetic field, for example, is best modelled with respect to the Earth's surface. Similarly, ground stations and other terrestrial targets (either for purposes of radio downlink or scientific observation) are best defined in this frame.
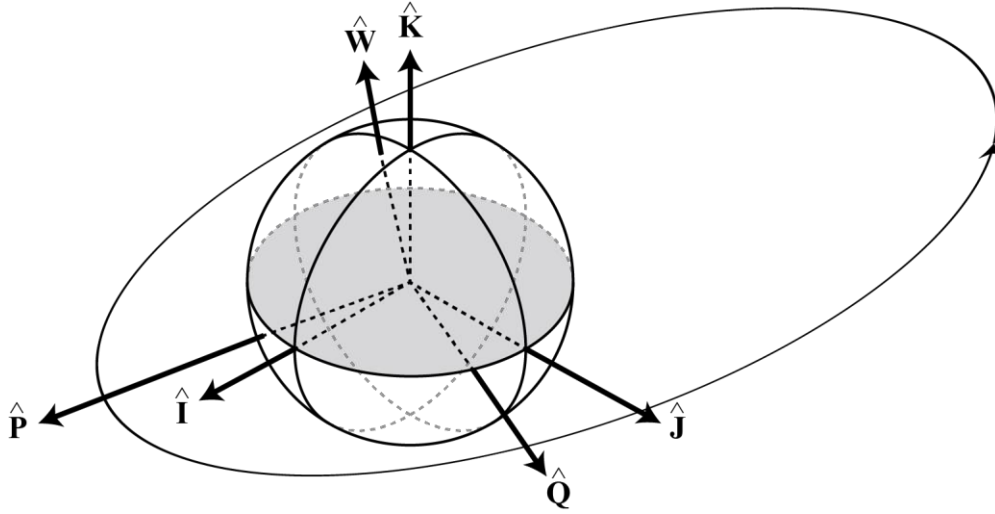
Figure 9: Perifocal Frame (PQW) defined by the orbit plane normal (W) and vector to periapsis (P). Direction of travel along the orbit is indicated by the arrow along the orbit path. As can be seen, the handedness of the orbit plane normal is consistent with the direction of travel.

The Perifocal frame (PQW) is an orbit centered frame. The frame is centered at the non-vacant foci of the orbit and the fundamental plane is defined by the plane swept out by the orbit. For practical purposes, this frame is defined for Keplerian two-body orbits. The principal direction of the orbit is defined by the vector to periapsis (the point of closest approach to the central body of the orbit). This frame is a useful intermediate frame for calculation. The transformation between the inertial ECI frame and the Perifocal frame is only a function of the classical orbital elements. Further, if the orbit does not itself precess, the orbit Perifocal frame is itself an inertial reference frame.
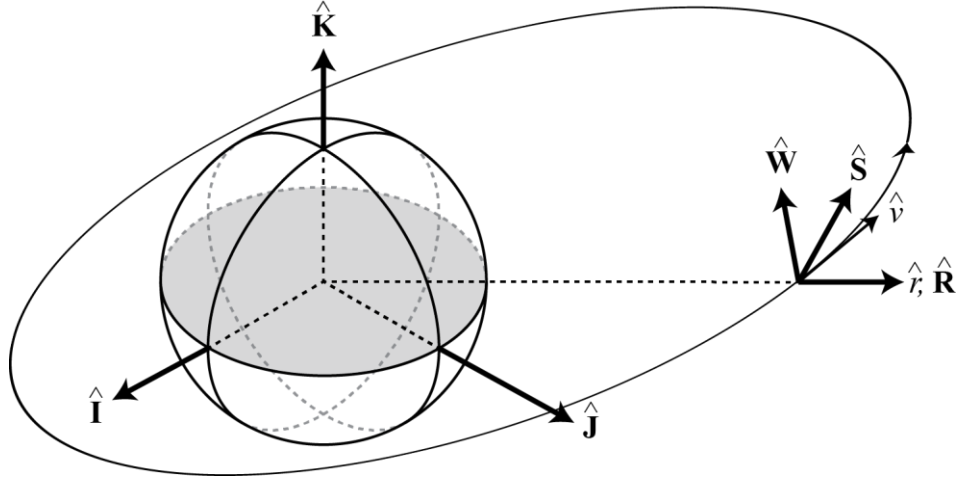
Figure 10: Satellite Radial (RSW) frame defined by the orbit plane normal (W) and the radial direction (R).
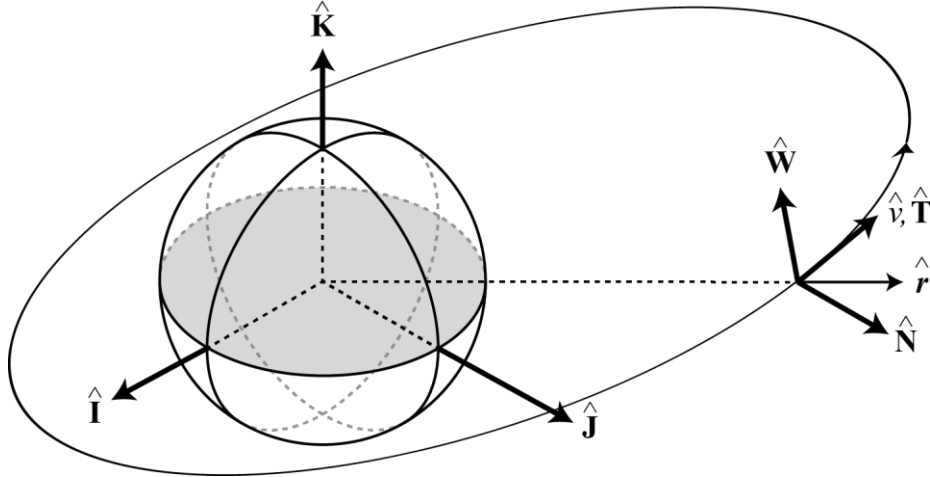
Figure 11: Satellite Normal (NTW) frame defined by the orbital plane normal (W) and the velocity track vector (T).

The Satellite Radial (RSW) and Satellite Normal (NTW) are both satellite centered frames. Like the Perifocal frame, the fundamental plane of the orbit is defined by plane swept out by the orbit; however, they differ from the Perifocal frame in the manner by which they define the principal vector. In the case of the RSW frame (Figure 10), the satellite radial vector (the vector in direct opposition to Nadir) is the principal direction. This is useful when defining attitude targets with respect to the nadir direction, such as for ground observing payloads. The NTW frame (Figure 11) defines its principal vector as the vector normal to the orbit plane and satellite velocity direction. In this manner, the velocity vector is aligned with the second basis of the frame and the handedness of the frame is consistent with the PQW and RSW frames. The NTW

17

frame is uniquely able to define constant attitude targets for payloads which need to be aligned with the velocity vector. The RSW and NTW frames are therefore the most important when defining attitude targets for the two category of payloads (the velocity-aligned RPA and SNeuPI, and the nadir aligned PMT). For circular orbits, the NTW and RSW frames are congruent. For elliptical orbits, the two frames will oscillate about the third axis (W) relative to each other. As such, these two frames are easiest to relate to other frames via the PQW frame.

It is simpler to represent certain quantities in specific frames (e.g. dynamics in an inertia frame, attitude targets in a satellite reference frame, geomagnetic field vectors in ECEF, etc.). As such, it is a best practice to create a map of frame adjacencies. Adjacent frames any two frames whose transformations from one to the other cannot be more simply represented by a sequence of other transformations. For example, ECI and PQW are adjacent but ECI and RSW are not (because ECI to RSW can be represented as the transformation from ECI to PQW followed by the transformation from PQW to RSW). For clarity, a diagram showing frame adjacencies is provided (see Figure 12). The mathematical representations of the transformations will be provided in the text where relevant. Frame transformations require different input information. In all cases presented herein, the orbit and current time provide sufficient information to determine frame transformations.
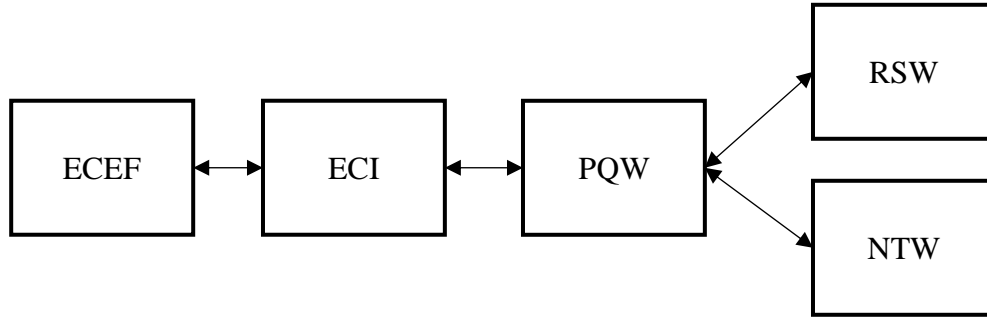
In this work, an attitude will commonly be represented as a quaternion or direct cosine matrix. In general, quaternions will be presented by a lower-case q with a right-superscript denoting the base frame and a left-superscript denoting the resulting frame. In this fashion, a quaternion representing the transformation from ECI to ECEF would be written as follows:

$$^{ECEF}q^{ECI} \tag{2}$$

Most commonly, the attitude will represent the orientation of the spacecraft body frame with respect to a reference frame. In such cases, either the right-superscript (SC) is used or the right-superscript is omitted for clarity. Further, when both superscripts are omitted, this always represents the attitude of the spacecraft with respect to the inertial (ECI) frame. This will most commonly be used when describing the attitude dynamics of the spacecraft.

In the case of a direct cosine matrix attitude representation, an upper-case letter (most often A) will be used to represent the transformation matrix in a similar fashion to the quaternion. This representation will often employ the same left- and right-superscript notation. The utility of this notation is most apparent when performing multiple successive rotations using a direct cosine matrix. When using vector notation where the vector reference is identified using a left-handed

superscripts, adjacent superscripts can be cancelled to determine the final frame representation (an example is provided in Eq. (3)).

$$^{NTW}\vec{v} = {}^{NTW}\boldsymbol{A}^{PQW\,PQW}\boldsymbol{A}^{ECI\,ECI}\boldsymbol{A}^{ECEF\,ECEF}\vec{v}$$

(3)

$$^{NTW}\vec{v} = {}^{NTW}\boldsymbol{A}^{\cancel{PQW}\,\cancel{PQW}}\boldsymbol{A}^{\cancel{ECI}\,\cancel{ECI}}\boldsymbol{A}^{\cancel{ECEF}\,\cancel{ECEF}}\vec{v}$$

In addition to attitude representations, reference frames are relevant to the representation of attitude rates. Whereas an attitude represents the orientation of one frame with respect to another, an attitude rate represents the rate of change of one frame with respect to a second frame, as seen from a third frame. In most cases, the third frame will be the same as one of the first two frames. The most useful representation from the perspective of this work is the rotation of the spacecraft body frame with respect to the inertial (ECI) frame, as seen in the inertial frame. Unless stated otherwise, the attitude rate representations will use this ordering of reference frames. The only notable deviation from this will be in defining attitude rate targets in the NTW frame which will be specifically addressed at the relevant point in the text (Section 3.3.2).

2.2 Spacecraft Attitude Dynamics

Spacecraft attitude dynamics are largely separable from spacecraft orbital dynamics. While orientation of the spacecraft can have an impact on satellite drag, which will be examined in greater detail in Chapter 5, the motion about the center of mass is separable from the motion of the center of mass[16]. The satellite attitude dynamics are derived from the law of the conservation of angular momentum and are best cast from that framework. The angular moment of the spacecraft, defined in an inertial reference frame, must be conserved unless acted upon by an external torque. If acted upon by an external torque, the angular moment must change in proportion to the amount of torque applied to the system and inversely proportional to the inertia

of the system. While angular momentum must be conserved at all times, this does not mean that the motion is constant. Angular momentum can be exchanged between axes, always in an attempt to increase the rotational stability of the craft.

The body angular rate with respect to the inertial frame can be represented by Equation (4) [17].

$$J\dot{\vec{\omega}} + \vec{\omega} \times J\vec{\omega} = \vec{M} \tag{4}$$

In this equation, J is the moment of inertia matrix, $\vec{\omega}$ is the angular rate, and $\vec{M}$ is the applied torque to the system. Without loss of generality, the spacecraft can be modelled as a rigid body with its body coordinates aligned with its principal axes[17]. If the principal axes are aligned with the body axes, the moment of inertia matrix becomes diagonal. As such, Equation (4) can be simplified to the system of Equations (5)-(7).

$$J_1\dot{\omega}_1 - (J_2 - J_3)\omega_2\omega_3 = M_1 \tag{5}$$

$$J_2\dot{\omega}_2 - (J_3 - J_1)\omega_3\omega_1 = M_2 \tag{6}$$

$$J_3\dot{\omega}_3 - (J_1 - J_2)\omega_1\omega_2 = M_3 \tag{7}$$

The moments, M_i, on the satellite are a combination of the external disturbance torques and the applied moments from the magnetic torqueing coils about the three principal body axes. This coupled set of non-linear ordinary differential equations has no closed-form solution though certain satellite configurations can exploit spacecraft symmetry to make the equations solvable in closed-form; however, this is not the case for LAICE. Instead, the system must be solved as a coupled system of equations using numerical methods.

The attitude quaternion is propagated by the infinitesimal state transition function defined by

$$\begin{bmatrix} \dot{q}_1 \\ \dot{q}_2 \\ \dot{q}_3 \\ \dot{q}_4 \end{bmatrix} = \begin{bmatrix} 0 & \omega_3 & -\omega_2 & \omega_1 \\ -\omega_3 & 0 & \omega_1 & \omega_2 \\ \omega_2 & -\omega_1 & 0 & \omega_3 \\ -\omega_1 & -\omega_2 & -\omega_3 & 0 \end{bmatrix} \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{bmatrix} = \mathbf{\Omega} q \tag{8}$$

This state transition function can be integrated forward. If the attitude rate is known to be constant, then the attitude at a future time can be explicitly determined without forward integration.

2.3 Deterministic Approaches to Magnetic/Sun Attitude Determination

In order to establish the orientation of a spacecraft, it is necessary to make observations of the surrounding environment to gain observability of the attitude state (i.e. orientation of the satellite with respect to a reference frame). This can be most easily accomplished by taking multiple independent measurements in the frame of the spacecraft and determining the transformation required to align the measurements with the same vectors modelled in a known reference frame[18]. As an example, if the spacecraft takes a measurement of the magnetic field using a magnetometer and the sun pointing vector using sun sensors, the direction of these measurements are related to the models of the magnetic field and sun vector in the Earth Centered Inertial (ECI) frame via the rotation. The rotation from the modelled vectors in the ECI frame to the measurements in the spacecraft-frame in this example is the attitude of the spacecraft with respect to the ECI frame. A single measurement cannot give the orientation of the spacecraft uniquely. This can be most easily described mathematically by a lack of determinism in the parameters. The attitude represents three rotations while a single three-axis measurement only provides two-unique angles; the magnitude of the magnetic field provides no useful information about the attitude, therefore it should only be considered to provide two pieces of information not three. Intuitively, this can be rationalized by considering a satellite which is spinning about the

magnetic field direction; as the spacecraft rotates about the magnetic field vector, the magnetic field measurement read by the magnetometer will remain unchanged thus indicating a single, unobserved degree of freedom. Hence, a second measurement is necessary to resolve the final degree of freedom.

The use of a second measurement presents its own challenges whereby a second, independent measurement would provide two additional pieces of information which causes the problem to become over-specified. This can be handled in a number of ways, one of which will be presented here in brief: the TRIAD method[19].

2.3.1 TRIAD Method

The TRIAD method for attitude determination forms two sets of orthonormal bases: 1) using the satellite body-fixed sensor values, and 2) using the reference-frame-fixed model values. Then, by taking advantage of the invertible properties of the attitude DCM representation, the attitude quaternion can be computed explicitly. To begin, one of the two vectors (sun or magnetic field) is selected as the principal direction. For the purposes of this example, the sun vector is used; in practice, it is preferable to use the vector which can be measured with higher precision. This vector is normalized and forms the principal axis of the basis for both the measurement and model (Eq. (9)). Next, the cross-product of the magnetic field and sun pointing vector is normalized to form the second component of the basis (Eq. (10)). The resultant vector will be orthogonal to the first vector of the basis by definition of the cross-product. This calculation should be made with consideration that if the two measurements happen to be aligned (i.e. the sun vector and magnetic field direction are aligned) it will be impossible to form an orthogonal basis; for the orbits considered for LAICE, this is a rare and transient condition. The final component of the bases is determined by taking the cross-product of the first and second

components of the bases (Eq. (11)). Similarly, this vector will be co-orthogonal to the first and second by construction. It is not necessary to normalize these vectors as the cross-product of two orthogonal normal vectors will itself be of unit length. Each of these bases (the reference basis and measurement basis) represent two identical sets of vectors represented in two different frames of reference (Eq. (12)) namely the reference frame and spacecraft body frame. The transformation between any pair of basis vectors from the reference and measurement base sets will be the rotation matrix representing the spacecraft attitude with respect to the reference frame. This can be written in the form of Equation (13), whereby the reference basis is equal to the matrix representation of the spacecraft attitude pre-multiplying the measurement basis. By post-multiplying both sides of the equation by the inverse of the measurement basis, ^{SC}S, and recalling that the inverse of an orthonormal matrix is equal to its transpose, it is possible to derive an equation for the spacecraft attitude explicitly in the form of the measurement and reference vectors (Eq. (14)).

$$\vec{R}_1 = \frac{^{ECI}\vec{r}_{sun}}{\left\|^{ECI}\vec{r}_{sun}\right\|}, \vec{S}_1 = \frac{^{SC}\vec{r}_{sun}}{\left\|^{SC}\vec{r}_{sun}\right\|} \tag{9}$$

$$\vec{R}_2 = \frac{^{ECI}\vec{r}_{sun} \times {^{ECI}\vec{r}_{mag}}}{\left\|^{ECI}\vec{r}_{sun} \times {^{ECI}\vec{r}_{mag}}\right\|}, \vec{S}_2 = \frac{^{SC}\vec{r}_{sun} \times {^{SC}\vec{r}_{mag}}}{\left\|^{SC}\vec{r}_{sun} \times {^{SC}\vec{r}_{mag}}\right\|} \tag{10}$$

$$\vec{R}_3 = \vec{R}_1 \times \vec{R}_2, \vec{S}_3 = \vec{S}_1 \times \vec{S}_2 \tag{11}$$

$$^{ECI}\boldsymbol{R} = \{\vec{R}_1|\vec{R}_2|\vec{R}_3\}, {^{SC}\boldsymbol{S}} = \{\vec{S}_1|\vec{S}_2|\vec{S}_3\} \tag{12}$$

$$^{ECI}\boldsymbol{R} = {^{ECI}\boldsymbol{A}^{SC}}({^{SC}\boldsymbol{S}}) \tag{13}$$

$$^{ECI}\boldsymbol{A}^{SC} = {^{ECI}\boldsymbol{R}}({^{SC}\boldsymbol{S}})^{-1} = {^{ECI}\boldsymbol{R}}({^{SC}\boldsymbol{S}})^{T} \tag{14}$$

This method is very useful for explicitly determining the attitude of the spacecraft using only two sets of vectors. It is, however, extremely susceptible to sensor noise and does not explicitly provide any information about the attitude rate. Finite-differencing of the attitude data can provide this information, though the effect of sensor noise is exacerbated when finite-differencing[20]. In general, this system is better suited to performing initial estimates of the attitude as a pre-step for a filtering of the measurements against a model of the satellite dynamics.

2.4 Filter Based Approaches to Magnetic-Only Attitude Determination

Magnetic-only attitude determination poses a challenge since there is insufficient information for full determination. As previously stated, a single magnetic field measurement vector only supplies two pieces of information in determining the attitude. In a conventional determination algorithm (TRIAD), two independent sets of reference and measurement vectors would be used. The two sets of vectors would provide complete observability of the orientation in SO(3) with redundancy. Filtering the two sets of measurements can be employed to refine the estimate of the measurement. In the case of magnetic-only attitude determination, full observability of the orientation in SO(3) cannot be obtained; however, by filtering the measurements and modelling the system dynamics, it is possible to make an estimate of the attitude[8], [21]–[25].

A single magnetic field measurement, made with arbitrarily high precision, is only capable of reducing the set of possible orientations to space of rotations about the magnetic field orientation. If the satellite body remains fixed relative to the local magnetic field, no additional information about the state would become available and it would therefore be impossible to ascertain the full attitude state of the spacecraft. Due to the spherical harmonic nature of the Earth's magnetic field, the direction of the magnetic field is not constant[26]. Therefore, as the

spacecraft traverses the orbit, the magnetic field vector will change direction with respect to an inertially-fixed reference frame. Considering an inertially-fixed satellite (i.e. non-spinning with respect to an inertial reference frame), any change in the direction of the measured magnetic field would resolve the missing information about the spacecraft information. If the spacecraft itself is spinning with unknown angular rate, it again becomes impossible to resolve the full orientation. Therefore, it becomes necessary to filter for both the attitude and attitude rate using the system dynamics and magnetic field models to obtain full observability of the orientation given knowledge of the system dynamics.

The system dynamics are highly nonlinear making linear filters, such as Kalman Filters unsuitable. Monte Carlo estimator and Particle Filters offer nonlinear filtering methods but are computationally expensive solutions[27], [28]. An Extended Kalman Filter (EKF) offers a compromise. A review of two types of attitude EKF's is presented herein: the 7-State EKF, and 6-State EKF.

2.4.1   7-State Extended Kalman Filter

The 7-state Extended Kalman Filter (EKF) employs an extended Kalman filter which predicts the new state based on the prior state and the satellite dynamics, and updates the estimation with the magnetic field measurement[29]. The state encompasses the attitude quaternion (4 components) and the attitude rate (3 components). While the four components of the attitude are not independent, they are not linearly dependent; therefore, the linear representation of the system is not singular.

The EKF, like the Kalman Filter, models the belief in the state estimate as a Gaussian distribution. The mean of the distribution is the best representation of the state estimate, and the covariance a measure of the certainty of this estimate. Unlike the Kalman Filter, the EKF allows

the state transition to be represented by nonlinear functions rather than linear ones. As a consequence of using a nonlinear function, the belief does not remain Gaussian in time. To remedy this, the EKF reconstructs a Gaussian approximation of the belief at each time step based on a linearization of the nonlinear function at each time step.

At each time step, the EKF updates the predictions of the state and covariance matrix. The predicted covariance matrix, $\bar{P}_{k+1}$, is formed by taking the covariance estimate from the previous iteration and the Jacobian of the nonlinear system dynamics (Equations (15)-(17)). If the covariance matrix goes to zero, the contribution from new measurements (as opposed to purely propagating the state via the dynamics) will similarly go to zero. To prevent this, a process noise covariance matrix, Q_{k+1}, is added to ensure the covariance matrix remains non-zero. The state estimate from the previous time step is integrated forward using the full nonlinear system dynamics according to Equations (18)-(20) to obtain the state prediction.

Prediction Step

$$\bar{P}_{k+1} = \Phi_{k+1}\hat{P}_k\Phi_{k+1}^{\mathrm{T}} + Q_{k+1} \tag{15}$$

$$\Phi_{k+1} \approx I_{7\times7} + \begin{bmatrix} \dfrac{\partial \dot{q}}{\partial q} & \dfrac{\partial \dot{q}}{\partial \omega} \\ \dfrac{\partial \dot{\omega}}{\partial q} & \dfrac{\partial \dot{\omega}}{\partial \omega} \end{bmatrix}(t_{k+1} - t_k) \tag{16}$$

$$= I_{7\times7} + F_k \cdot \Delta t \tag{17}$$

$$\bar{x}_{k+1} = \hat{x}_k + \int_{t_k}^{t_{k+1}} \dot{x}\,dt \tag{18}$$

27

$$\bar{q}_{k+1} = \hat{q}_k + \frac{1}{2} \int_{t_k}^{t_{k+1}} \mathbf{\Omega} q \, dt \qquad (19)$$

$$\bar{\omega}_{k+1} = \hat{\omega}_k + \int_{t_k}^{t_{k+1}} \mathbf{J}^{-1} \left[\vec{M} - \vec{\omega} \times \mathbf{J}\vec{\omega} \right] dt \qquad (20)$$

The state-prediction is then updated with the measurements z_k. In the original 7-State EKF based on the work in Reference [29], the measurement only consisted of the magnetic field measurement. Due to the nonlinear nature of quaternion attitude dynamics, such a filter must have a good guess of the attitude state (quaternion and body rate) in order to converge[22]. For even well behaved dynamical systems, a poor initial guess will cause the filter to diverge. By attempting to filter the rate without direct measurement of the angular rate, the filter will update the estimate of the rate in proportion to the partial of the attitude with respect to the rate and the measurement error of the magnetic field. In this manner, large errors in attitude will result in large updates to the rate even if the estimate of the rate was perfect. At the subsequent update step, the attitude will be propagated dramatically due to the large angular rate. This in turn causes the filter to again increase the angular rate causing the filter to become unstable.

In this work, the measurement is augmented with the body rate measurement from the MEMS gyros. This addresses the core challenge face by the previous filter. With knowledge of the angular rate, the update more heavily influences the update to the attitude. Even a noisy estimate of the angular rate is therefore able to allow a filter to converge – this will be demonstrated for initial attitude pointing errors of up to 180°.

For the proposed filter, the measurement $\vec{z}_k$ becomes the magnetometer measurement and the MEMS gyro angular rate measurement. The Kalman gain, $\mathbf{K}_{k+1}$, is a function of the predicted

covariance matrix, the derivative of the measurement with respect to the state, and the measurement noise covariance matrix, R_k (Equation (21)). The measurement noise covariance matrix ensures that the Kalman gain does not become very large when the predict covariance matrix is very small. The Kalman gain is then used with the measurement residual to update the state prediction, $\bar{x}_{k+1}$, to the state estimate, $\hat{x}_{k+1}$ (Equation (22)). Finally, the Kalman gain is then used to update the predicted covariance matrix to the estimated covariance matrix (Equation (23)).

Update Step

$$H_k = \frac{\partial \vec{z}_k}{\partial x}$$

$$= \begin{bmatrix} \dfrac{\partial A(\vec{q})}{\partial q_1}\vec{\beta} & \dfrac{\partial A(\vec{q})}{\partial q_2}\vec{\beta} & \dfrac{\partial A(\vec{q})}{\partial q_3}\vec{\beta} & \dfrac{\partial A(\vec{q})}{\partial q_4}\vec{\beta} & \dfrac{\partial A(\vec{q})}{\partial \omega_1}\vec{\beta} & \dfrac{\partial A(\vec{q})}{\partial \omega_2}\vec{\beta} & \dfrac{\partial A(\vec{q})}{\partial \omega_3}\vec{\beta} \\ \dfrac{\partial \vec{\omega}}{\partial q_1} & \dfrac{\partial \vec{\omega}}{\partial q_2} & \dfrac{\partial \vec{\omega}}{\partial q_3} & \dfrac{\partial \vec{\omega}}{\partial q_4} & \dfrac{\partial \vec{\omega}}{\partial \omega_1} & \dfrac{\partial \vec{\omega}}{\partial \omega_2} & \dfrac{\partial \vec{\omega}}{\partial \omega_3} \end{bmatrix}$$

$$K_{k+1} = P_k H_k^T [H_k P_k H_k^T + R_k]^{-1} \tag{21}$$

$$\hat{x}_{k+1} = \bar{x}_{k+1} + K_{k+1}\left(\vec{z}_k - \frac{A(\bar{q}_{k+1})\vec{\beta}}{\bar{\omega}_{k+1}} \right) \tag{22}$$

$$\widehat{P}_{k+1} = (I_{7\times7} - K_{k+1}H_{k+1})\overline{P}_{k+1} \tag{23}$$

The original 7-State does take advantage of the onboard processing capability to propagate the attitude state more precisely than a linear representation; however, convergence of the EKF is dependent on the accuracy of the initial guess. Since the attitude update is highly coupled to the attitude rate, noise on the attitude rate information can cause the estimate to become unstable. As stated, without direct observability of the attitude rate, it is easy for the filter to increase the angular rate unreasonably high. This results in the filtering being unable to correct the estimate

the attitude thus leading to filter instability. Additionally, high angular rates require more discretization of the forward integration step in Equations (18)-(20). The original 7-State Filter can therefore only be applied to a system with well understood and predictable dynamics (e.g. gravity gradient or spin-axis stabilized spacecraft)[22], [29], [30].

For this reason, the inclusion of the angular body rate measurements from the MEMS rate gyros has been added to the filter. While angular rates can be filtered from the changes in magnetic field data[30]–[32], it is highly susceptible to sensor noise and is best accomplished with additional independent vector measurements. Even though the MEMS gyros are themselves noisy, direct filtering of the MEMS gyros provides certainty in the attitude rate in a way that second order filtering cannot readily provide[33]. It also significantly simplifies the filter design and analysis of the filter's stability. The convergence of the original and augmented 7-State filter will be explored in more detail in Section 2.5.

2.4.2   6-State Extended Kalman Filter

The 7-State attitude filter described has been demonstrated to converge for large errors in the initial attitude estimate. This is an important feature as the satellite will have no information about its attitude state when deployed. It does, however, require the inversion of a 6x6 matrix which is computationally expensive (see Equation (21)).

To alleviate the computation burden of the attitude determination algorithm, a reduced order filter can be used once the covariance has converged. For this purpose, a 6-State Extended Kalman Filter is employed[29]. Rather than representing the attitude as the quaternion and attitude rate, this EKF employs an auxiliary state which represents small deviations about a nominal attitude and attitude rate. In the case of the angular body rates, the auxiliary state

represents a linear perturbation. For the attitude, the auxiliary state represents the attitude error $\Delta\vec{q}$ from the nominal attitude $\vec{q}_{nom}$ (see Equation (24)).

$$\vec{q} = \vec{q}_{nom} \otimes \begin{bmatrix} \Delta q_1 \\ \Delta q_2 \\ \Delta q_3 \\ 1 \end{bmatrix} = \vec{q}_{nom} \otimes \begin{bmatrix} \Delta\vec{q} \\ 1 \end{bmatrix} \tag{24}$$

The principal advantage is that for sufficiently accurate estimates of the initial attitude and attitude rate, the 6-state EKF is able to refine the estimates only using observations of the magnetic field. This reduces the size of the complexity of the matrix inversion step from 6x6 to 3x3 (See Equation (32)).

Prediction Step

$$\overline{\boldsymbol{P}}_{k+1} = \boldsymbol{\Phi}_{k+1}\widehat{\boldsymbol{P}}_k\boldsymbol{\Phi}_{k+1}^{\mathrm{T}} + \boldsymbol{Q}_{k+1} \tag{25}$$

$$\boldsymbol{\Phi}_{k+1} \approx \mathbf{I}_{6\times6} + \begin{bmatrix} \dfrac{\partial\delta\dot{q}}{\partial\delta q} & \dfrac{\partial\delta q}{\partial\delta\omega} \\ \dfrac{\partial\delta\dot{\omega}}{\partial\delta q} & \dfrac{\partial\delta\dot{\omega}}{\partial\delta\omega} \end{bmatrix} (t_{k+1} - t_k) \tag{26}$$

$$= \mathbf{I}_{6\times6} + \boldsymbol{F}_k \cdot \Delta t \tag{27}$$

$$\bar{x}_{k+1} = \hat{x}_k + \int_{t_k}^{t_{k+1}} \dot{x}\,dt \tag{28}$$

$$\bar{q}_{k+1} = \hat{q}_k + \frac{1}{2}\int_{t_k}^{t_{k+1}} \boldsymbol{\Omega}q\,dt \tag{29}$$

$$\bar{\omega}_{k+1} = \hat{\omega}_k + \int_{t_k}^{t_{k+1}} J^{-1}[N - \omega \times J\omega]\, dt \tag{30}$$

Update Step

$$H_k = \frac{\partial z_{by}}{\partial x} = \left[\frac{\partial A(q)}{\partial \delta q_1}\beta \quad \frac{\partial A(q)}{\partial \delta q_2}\beta \quad \frac{\partial A(q)}{\partial \delta q_3}\beta \quad \frac{\partial A(q)}{\partial \omega_1}\beta \quad \frac{\partial A(q)}{\partial \omega_2}\beta \quad \frac{\partial A(q)}{\partial \omega_3}\beta \right] \tag{31}$$

$$K_k = P_k H_k^T [H_k P_k H_k^T + R_k]^{-1} \tag{32}$$

$$\hat{y}_{k+1} = \bar{y}_{k+1} + K_{k+1}\left(\vec{z} - A(q)\vec{\beta} \right) \tag{33}$$

$$\hat{P}_{k+1} = (I_{6\times6} - K_{k+1}H_{k+1})\bar{P}_{k+1} \tag{34}$$

$$\delta q_4 = \sqrt{1 - \delta q_1^2 - \delta q_2^2 - \delta q_3^2} \tag{35}$$

$$\hat{q} = \delta q \otimes \bar{q} \tag{36}$$

$$\hat{\omega} = \bar{\omega} + \delta \omega \tag{37}$$

$$\hat{x} = [\hat{q}_1, \hat{q}_2, \hat{q}_3, \hat{q}_4, \hat{\omega}_1, \hat{\omega}_2, \hat{\omega}_3] \tag{38}$$

2.4.3   Particle Filter

This method employs a series of 7-state Extended Kalman Filters (described in the previous Section). Since the 7-state EKF has a convergence time that is highly dependent on the accuracy of its initial state estimate, it is necessary to make a set of initial guesses that guarantee at least one of the guesses is within the rapid-convergence region. As the EKF's evolve with updates, looking at the measurement residual as the state covariance matrix converges will indicate which of the EKF's has converged on the correct solution[34]. EKF's with high measurement residuals

can therefore be discarded as they have failed to converge to the correct solution[35]. Based on the measurement residual, it is therefore possible to reduce the number of EKF's.

2.4.4 Distributing Initial Attitude Estimates

In order to evenly distribute the initial guesses, a number of methods could be employed. First, a set of quaternions can be drawn at random from a uniform distribution spanning the space of all quaternions SO(3)[36]–[40]. The quaternion represents a point on a 4-dimensional unit sphere, thus sampling approaches can be thought of as extensions to uniformly sampling points on the surface of a 3-dimensional unit sphere. For a unit sphere in 3-dimensions, a uniformly distributed point (x, y, z) can be generate by choosing $z \in [-1,1]$ and x and y on a circle of radius $\sqrt{1 - z^2}$.[3] In the case of SO(3),

$$u_1, u_2, u_3 \in [0,1] \tag{39}$$

$$h = \left(\begin{matrix} \sqrt{1 - u_1} \sin(2\pi u_2), \sqrt{1 - u_1} \cos(2\pi u_2), \dots \\ \sqrt{1 - u_1} \sin(2\pi u_3), \sqrt{1 - u_1} \cos(2\pi u_3) \end{matrix} \right) \tag{40}$$

$$h = -h \ \forall \ h_4 < 0 \tag{41}$$

The number of samples must be large in order to guarantee the randomly sampled points span the space sufficiently well in order for at least one particle to converge. An ideal solution would be to evenly space the initial state estimate guesses. The first method would be to create a random sampling, then anneal the set to ensure maximum dispersion of the initial points. This can be achieved using a form of k-means dispersion in four dimensions[41]. This method does not guarantee uniform spacing between adjacent points, but does enforce some regularity to the mesh of points of initial estimates to within some desired value. The key benefit of this approach is that it reduces the number of sampled points. Another method would be to use regular 4-

polytopes to evenly distribute initial estimates[42]. A regular 4-polytope is a four dimensional extension of a regular polyhedra in three dimensions (or a regular polygon in two dimensions) in the sense that it consists of an evenly distributed set of vertices inscribed on a unit sphere of matching dimension (unit hypersphere in the 4-dimensional case, a unit sphere in the 3-dimensional case, or unit circle in two-dimensions). The principal advantage of using a 4-polytope is that it has the property that its vertices are guaranteed to be evenly spaced on the unit hypersphere, guaranteeing uniform coverage with a minimum of points.
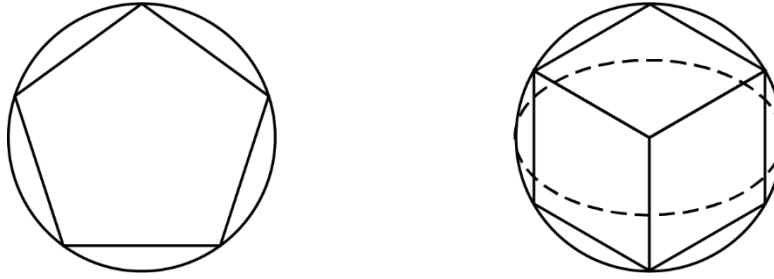
Figure 13: Regular geometries inscribed in n-Spheres of dimension 2 and 3.

A number of candidate 4-polytopes exist and are summarized in the table below[43]. In evaluating candidate geometries, the number of vertices of each polytope is most important feature. The 120-cell polytope is of particular interest due to its high number of vertices (total of 600 vertices). Due to the nature of quaternions whereby $\{q = -q\}$, the total number of candidate points that can be generated from a 120-cell is 300[42]. By distributing the points along the vertices of the upper-half of the 120-cell polytope, an initial distribution of 300 candidate quaternion estimates is obtained, each equally distributed from each other. A table of the quaternions generated by the 120-cell polytope is provided in Appendix B.

Table 1: List of Regular 4-Polytopes

| Name	Family	Vertices
5-Cell	n-simplex	5
8-Cell	n-cube	16
16-Cell	n-orthoplex	8
24-Cell	Fn family	24
120-Cell	n-pentagonal polytope	600
600-Cell	n-pentagonal polytope	120
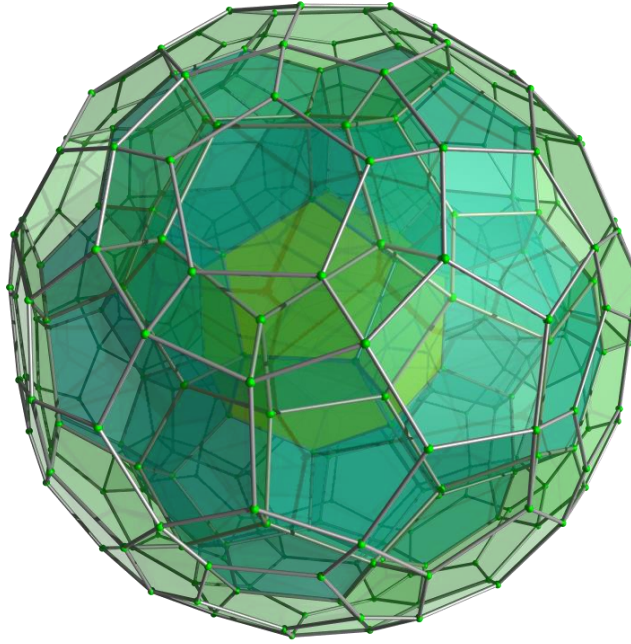
Figure 14: Three-dimensional projection of 120-cell 4-polytope.[44]

2.5 Software Simulation of Magnetic Attitude Determination

In order to simulate the magnetic attitude determination algorithms, a simulation test-bed needed to be created. The simulation incorporates three major types of components: orbit model, magnetic field model, and satellite dynamics model. These components are in addition to the various determination algorithms discussed. An object-oriented approach was used to allow the individual components to be easily replaced in the code without having to make significant changes to the program structure of the code. The purpose and implementation of each

component will be discussed, as well as the specific implementations used in the simulation. Validation of each model against accepted industry standards is also presented.

2.5.1   Implementation of Attitude Determination in Software

In order to code and test multiple attitude determination and control schemes, an object-oriented library was created in C++. As previously articulated, attitudes may be presented in a number of representations (quaternions, direct cosine matrices, Gibbs vectors, Euler angles, etc.) and can represent the transformation between any combinations of reference frames. Comprehensive libraries exist (for both C and Matlab) to transform between representations and frames[8]; however, the interfaces require programmers to keep track of the representation and frame of each variable throughout the code. For example, a Simulink function exists to convert between quaternions and direct cosine matrices; if, however, an attitude in the form of a set of Euler angles is input into the Simulink block, the code will not produce the correct direct cosine matrix output (or will throw an error during execution). Despite the Euler angles and quaternion representing the same attitude, the process to convert each to a direct cosine matrix (or any other representation) requires a unique, representation dependent function. While this can be mitigated with careful tracking of variables and extensive testing of code, it puts an unnecessary burden on the programmer and slows programming.

Figure 15: Simulink blocks for Quaternion to DCM (left) and Euler Angle Series to DCM (right)

The object-oriented attitude library presented herein addresses this problem by treating attitudes (and other types of germane variables such as angular velocity, velocity, and position vectors) as objects which store their representations and reference frames – in addition to the orientation information itself. In this way, the object can execute its own methods to change its

representation or reference frame on demand, eliminating the need for the programmer to track such information. Similarly, it allows programmers to code functions without placing restrictions on the incoming attitude representation or reference frame. Using the same example of converting an attitude from a quaternion to a direct cosine matrix, the object-oriented approach is agnostic to the original representation as shown in lines 3 and 4 of Code Block 1. Internally, the changeFormat command executes the appropriate frame and or format conversions.

Code Block 1: Converting Quaternion and Euler Angle Representations to DCM

```
1      //Create attitude as a Quaternion
2      Attitude a(Attitude::Quaternion);
3      //Create attitude as set of Euler Angles
4      Attitude b(Attitude::Euler);
5      //Convert attitudes a & b to Direct Cosine Matrices
6      a.changeFormat(Attitude::DCM);
7      b.changeFormat(Attitude::DCM);
8      //Attitudes a & b now have the same form (DCM)
```

In addition to simplifying coding, the generalization also allows for the construction of other more complex objects with standardized interfaces. The principle examples are sensor objects, actuator objects, attitude determination objects, and attitude control objects. By coding each sensor, actuator, determination method, or control law as an object with standard interfaces, it is easy to change control paradigms quickly with a couple lines of code. Similarly, it allows for the same paradigm to implement several versions of the same determination or control methods in parallel to test different weighting schemes. This approach significantly advantages the programmer by allowing for seamless parallel optimization of control gains.

Simplifying coding (specifically the function interfaces) is critical as there are in two parallel code bases: one flight code-base to be used by the satellite on-orbit, and a second simulation code-base for validating the flight code. Certain coding elements, such as the orbit ephemeris

model and magnetic field model, may be used for both the flight code-base and the simulation code-base. Rather than make two separate versions of each, a single code implementation with standardized interfaces is preferred. An object-oriented approach is essential to accomplishing this task.

2.5.2 CubeSim

In the following sections, the flight- and simulation-specific code will be called out. In all other instances, the code can be used for both the purpose of simulating the satellite or onboard the satellite. For brevity, the object-oriented implementations of the code which can be used for both flight and simulation code are referred to simply as the *CubeSim* implementation of the code. CubeSim, as a portmanteau, therefore refers to any code (objects, methods, interfaces, etc.) that was developed in the process of this work with the intent of being used for flight software and the validation of flight software through simulation.

2.5.3 Orbit Ephemeris Model

The purpose of the orbit model is to return the position and velocity state of the spacecraft at any arbitrary time within the orbit of the spacecraft. The orbit model is necessary for both the purpose of simulation as well as onboard the satellite. In practice, the orbit position is an input to magnetic field models which are necessary for simulation (to simulate magnetometer readings) and as a reference model onboard the satellite.

An orbit model can be either an analytical model, a semi-analytical model, or a propagation model. Internally, each operates differently; however, each is required to provide orbital parameters (as either position/velocity or classical orbital elements) for an input time. The orbit object prototype parses a Two-Line Element (TLE) into its constituent input variables (a TLE is a representation of the satellite orbital parameters at some reference time). The orbit models then

either propagate the satellite orbital parameters forward to the input time or return the orbital parameters for the input time based on an analytical model.

The Simplified General Perturbations 4 (SGP4) model is used as the principal model used onboard the satellite. It is an analytical orbital model which lends itself to quick computation and is accurate relative to high precision orbital propagators over small time scales[14]. The satellite will be receiving updated TLEs on a twice-daily basis, the SGP4 model remains very accurate over these times scales (~10m).

The simulated satellite ephemeris for the purpose of simulation is propagated using two main models: the SGP4 model and a high-precision orbital propagator (HPOP) model. The SGP4 analytic model is identical to the satellite's ephemeris model; it is principally used for creating idealized profiles for validating determination algorithms. The HPOP model is propagated forward from the TLE epoch based on a multi-body gravitational model[45]. The model is designed with a tunable precision by including more or fewer gravitational bodies. It also possesses the ability to optionally include a graphical drag and solar radiation pressure model. The HPOP model has higher accuracy for even small numbers of gravitational bodies though it is computationally more burdensome than the SGP4 model. For this reason, it is well suited for modelling the truth state of the orbital position. The discrepancy between the SGP4 satellite model and HPOP model used for simulating the truth state introduces uncertainties which are representative of the discrepancies between the satellite's orbit model and the truth state in orbit.

In order to validate the accuracy of the pseudo-analytical model, the CubeSim SGP4 ephemeris model was compared to the SGP4 model from a commercial software package (STK 10). The CubeSim implementation is based on the published code used for STK's SGP4

model[14]. As such, the two models only differ in how fractional seconds are handled. This

discrepancy leads to very small deviations in the orbital position (<3m at maximum).
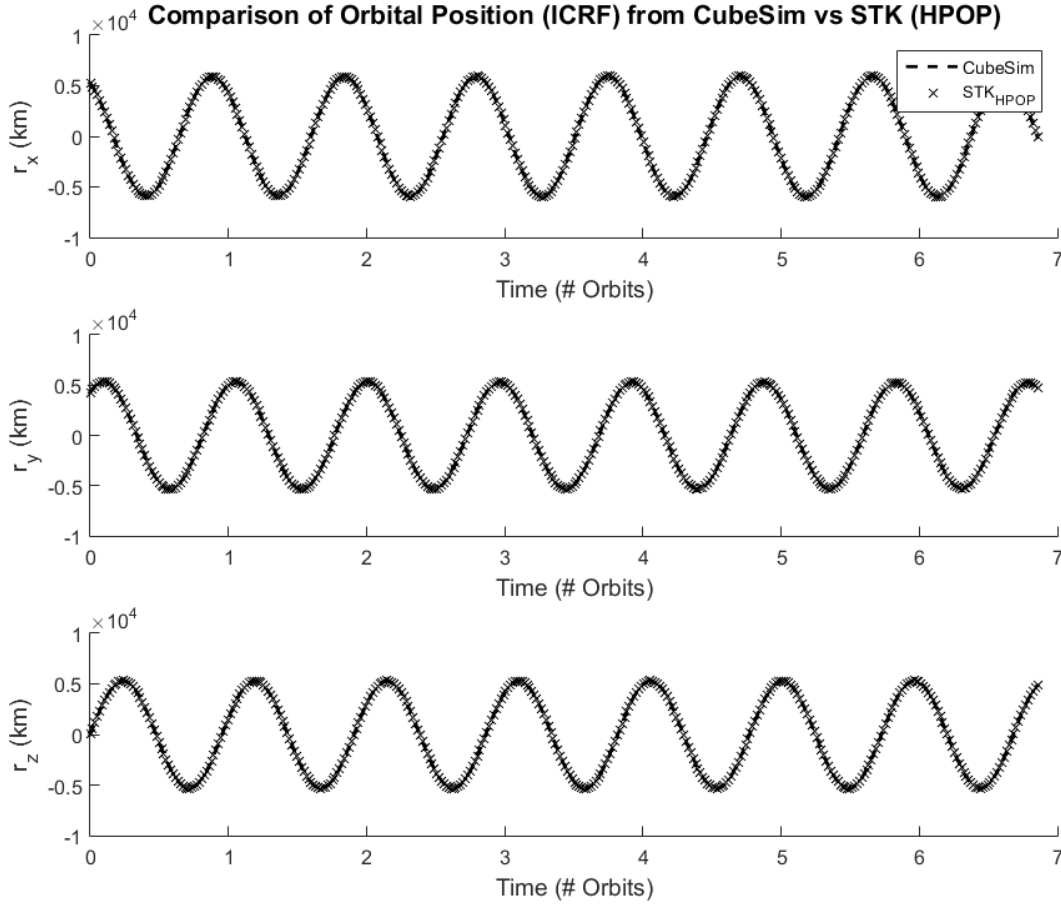
Figure 16: Demonstration of agreement in orbital position estimation between CubeSim orbital model and industry standard high-precision orbital propagator (STK).

In order to validate the CubeSim SGP4 model as an accurate model for the satellite, it is

compared to the HPOP model from STK 10. This high precision propagator is the best

representation of the orbit. To perform the comparison, both CubeSim SGP4 and STK HPOP

models were given identical TLE files (see Appendix A). Position and velocity samples were

taken for the same time points for a 12 hour period. Figure 16 and Figure 17 show the agreement

between the CubeSim and STK.
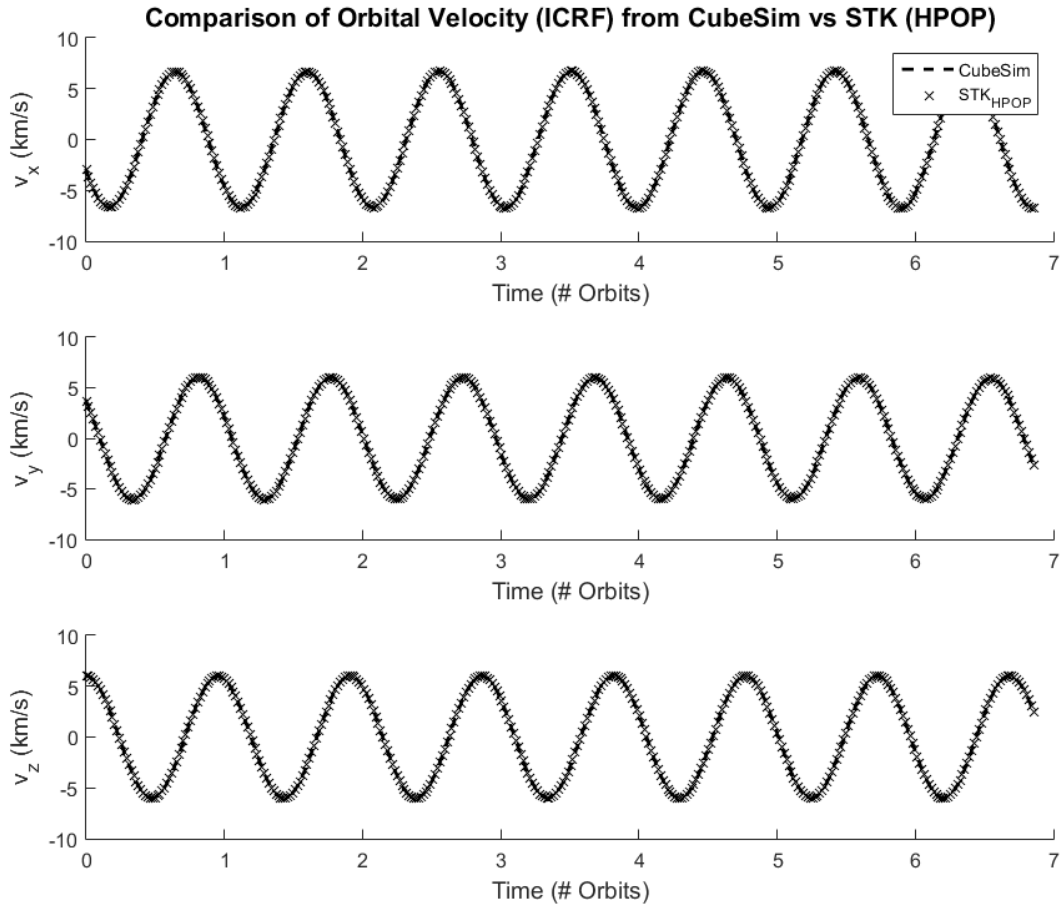
Figure 17: Demonstration of agreement between CubeSim and STK for orbital velocity.

When compared over a multi-day period, the STK HPOP and CubeSim SGP4 models do begin to deviate significantly in position. In practice, satellites will receive updated ephemeris data derived from ground ranging at least twice daily. If ground ranging is unavailable, high-precision propagators run on desktop system on the ground can be used to provide more accurate ephemeris estimates. Consequently, the deviations between the SGP4 model and HPOP will be relatively small by the time of update.

The CubeSim HPOP model was similarly compared to a commercial software package (STK 10) for the purposes of validating the model implementation. The number of gravitational bodies to include was kept consistent. In both cases, contributions from aerodynamic drag and solar

radiation pressure were ignored. The disagreement between the models is small and variations, again, are largely due to differences in how the models handle fractional seconds.

2.5.4   Magnetic Field Model

The satellite uses the magnetic field model as a reference vector for the purposes of determination. By comparing the expected magnetic field (based on estimated attitude) with the observed magnetic field (from onboard magnetometers), it is possible to observe errors in attitude knowledge. From a simulation standpoint, the magnetic field model is necessary to simulate satellite magnetometer data for input to the attitude determination algorithms.

The International Geomagnetic Reference Field (IGRF) model is used to model the magnetic field the satellite will experience on orbit[26]. The model is accurate to within 2% of the measured field intensity 98.1% of the time[46]. The model is updated every 5 years and has a set of time-varying parameters to improve the accuracy of modelled field as models become less current.

By way of validation, the CubeSim implementation of this model demonstrates good agreement with the model coded into AGI's STK. To validate the model, the CubeSim SPG4 model was used to determine the orbital position in ECEF, which was in turn used as an input to the CubeSim IGRF model for a set of time points over 12 hours. The output magnetic field was recorded in ECI coordinates. The same ECI magnetic field measurements were generated using AGI's STK using the same satellite TLE for the same time points. The STK satellite position was simulated using an HPOP propagator.
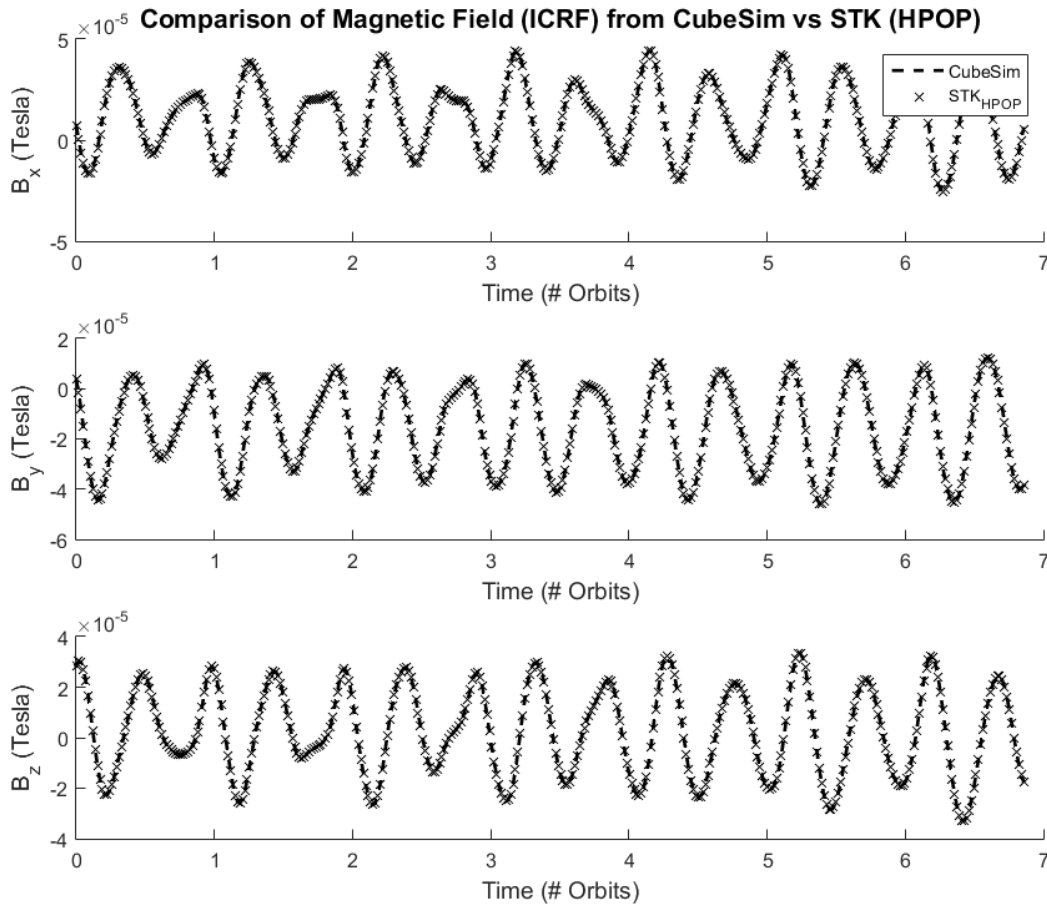
Figure 18: Demonstration of agreement between CubeSim model and STK in computing magnetic field strength.

Since the CubeSim IGRF model is used for both the onboard magnetic field model and simulating the magnetometer data, the perceived magnetic field will exactly match the model in the absence of sensor noise. This generates an artificially precise magnetic field measurement to model correspondence. To more realistically simulate the magnetic field sensors, Gaussian white noise is added to the measurement. While this improves the realism, the IGRF is only an approximation of the magnetic field; un-modelled variations in the magnetic field due to solar activity and other phenomena can create a variance of up to 5° in magnetic field direction based on measurement. This will significantly impact the performance of a real attitude determination

43

and control system and therefore it was desirable to have CubeSim model this variance in a physically realistic manner.

2.5.5 Satellite Dynamics Model

The satellite dynamic model propagates the satellite attitude state (attitude and attitude rate) based on provided external torques. The external torques may be solely from control torques (such as magnetic torqueing) or from disturbance torques (gravity-gradient, solar radiation pressure, aerodynamic drag, etc.) which would need their own separate models to compute. The model propagates the state based on the full non-linear dynamics for the rate given in Equations (4)-(7) and the infinitesimal quaternion update step given in Equation (8). An RK family propagator with an adaptive step-size is used to ensure that the error over the propagation duration does not become large. If the error is large, the step-size is reduced and the propagation is done in a piecewise manner over the original step-size. As such, the propagation step-sizes must be kept necessarily small to ensure that the input torque remains relatively constant over the entire step-size.

2.5.6 Simulation Results

The original 7-State EKF (without MEMS gyro readings) was simulated for a range of initial attitude estimates. For small angular errors ($<15°$) in quaternion estimates, the determination algorithm is able to converge within 0.5 orbits (Figure 19); however, for larger angular errors in quaternion estimate, the filter diverges (Figure 20). The performance of the filter is dependent on tuning. It is possible to obtain a larger region of convergence; however, this comes at a cost of time to convergence. As such, the original 7-State EKF demonstrates a need for improvement.
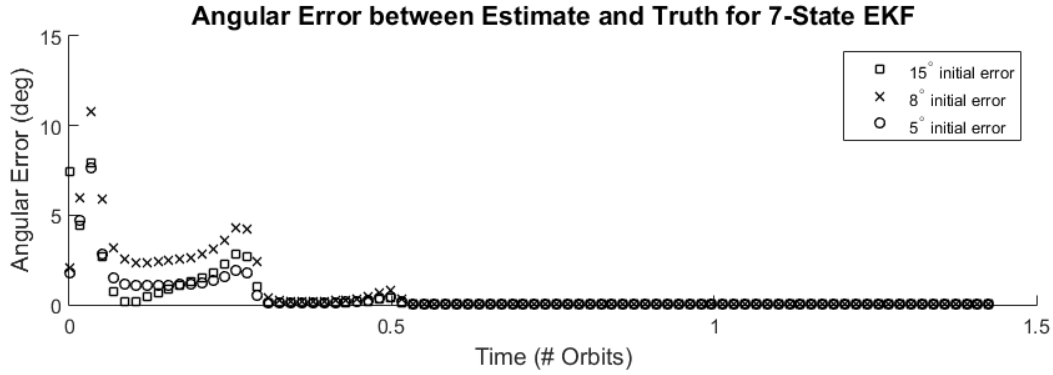
Figure 19: Convergence of the Original 7-State EKF for small initial attitude errors. The simulations undergo one update step prior to output values, so the first time points shown are after the first update step.
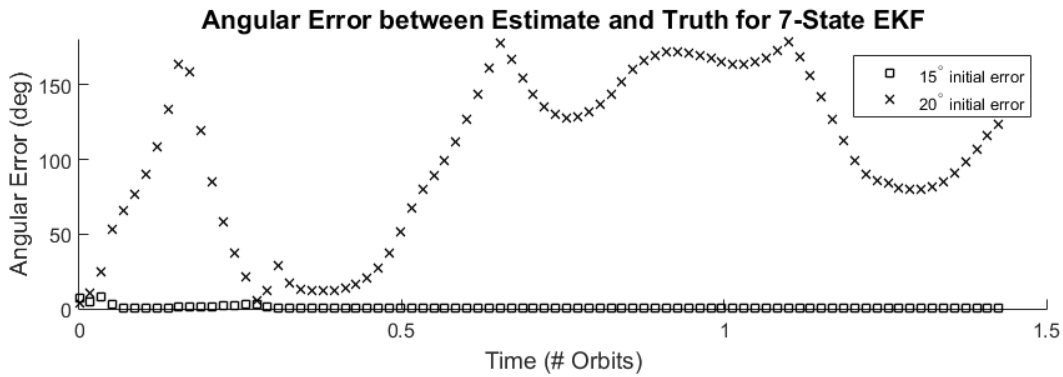
Figure 20: Demonstration of Original 7-State EKF divergence for larger initial attitude errors.

The augmented 7-State EKF (with MEMS gyro measurements) was simulated with initial angular errors up to 180° (the maximum possible disagreement in orientation). For the purposes of this simulation, the MEMS gyro measurements had Gaussian noise applied to it on the order of the noise described in Section 1.2.1.3. For even very large initial errors, the system was able to converge within 1.5 orbits (Figure 21). The angular rate estimate does deviate from zero, but the measurement updates prevent the filter from driving the rate estimate unreasonably high. Similar to the original 7-State EKF, the performance of the EKF is dependent on filter tuning. Tuning of this filter, along with tuning of the Attitude Control algorithms, will be revisited in Chapter 6.
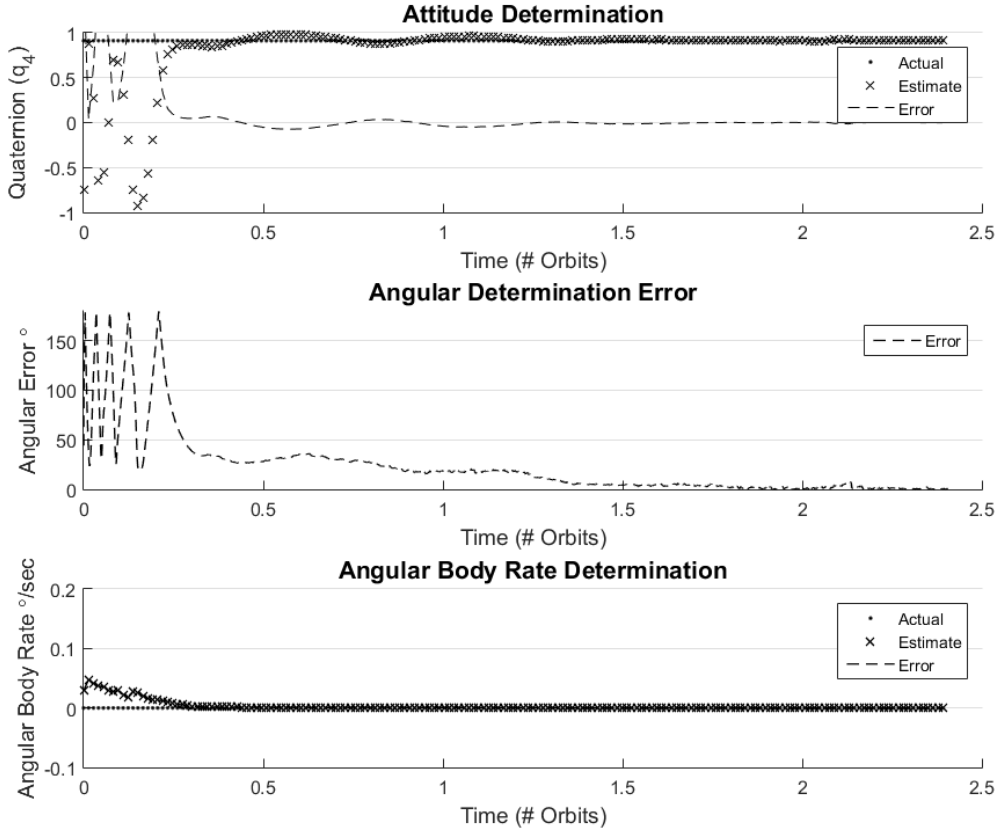
Figure 21: Performance of Augmented 7-State EKF for large initial pointing error (180°)

2.6 Summary of Contributions

In this section, an augmentation to an existing 7-State EKF was proposed as a means of improving the region of convergence for attitude determination algorithms relying on magnetic field measurements. The addition of MEMS rate gyro information significantly improved the performance of the filter. In order to improve the rate of convergence, a novel particle filter was proposed with equally distributed initial attitude estimates seeded by regular 4-polytopes. This approach deterministically seeds the space of initial attitude estimates for 7-State EKFs. Based on measurement error residuals, the particle filter can either resample state estimates about the best candidates to improve convergence, or eliminate poorly performing filters to save on computation.

46

CHAPTER 3
MAGNETIC ATTITUDE CONTROL

3.1 Modelling Spacecraft Dynamics

In order to control its attitude, the LAICE satellite uses magnetic torqueing coils. Typically, magnetic torqueing coils are used in concert with other reaction control devices such as reactional wheels, control moment gyros, or thrusters; however, in this application, they are used as the primary attitude control device. Due to the small size of the craft, relatively loose pointing requirements, and low-cost of the mission, magnetic control actuators are an ideal candidate.

In addition to magnetic control, the craft is given a positive static margin in the +z direction. As a result, any aerodynamic disturbance torques will cause the satellite to reorient in "ram" velocity direction naturally. This simplifies the controller considerably, especially as the mission nears its peak science period at lower altitudes near end of mission where the atmosphere is denser.

To define attitude targets (attitude and attitude rate), it is simplest to define the targets in a satellite fixed frame. As previously stated, the spacecraft is optimally oriented with its +z axis aligned into the orbit velocity direction and its –y face pointing towards nadir when collecting scientific data. As such, it is prudent to use the satellite normal (NTW) reference frame[14]. The reference frame originates at the satellite's position, with its secondary axis, $\hat{T}$, pointing in-track with the spacecraft velocity vector and the tertiary axis, $\hat{W}$, pointing in the direction of the normal vector to the orbital plane. The principal axis, $\hat{N}$, is the resultant co-orthogonal vector to the secondary and tertiary axes; nominally, the principal axis is aligned with the zenith direction for circular orbits, but diverges over the orbit for more eccentric orbits. Given this frame construction, LAICE's science pointing orientation is constant when defined in the NTW frame.

To represent the attitude target in the inertial frame, it is expedient to use the perifocal (PQW) frame as an intermediate transformation. The PQW frame shares the same fundamental plane as the NTW frame (i.e. orbit plane). The principle axis of the PQW frame is the vector originating from the Earth center extending to the point of periapsis. As such, the PQW and NTW are rotated by an angle defined by the velocity vector about the third axis, $\widehat{W}$, of both frames. This angle varies as the satellite moves along the orbit. To define this transformation for arbitrary points along the orbit, it is helpful to define the velocity in terms of classical orbital elements.

The velocity in the perifocal (PQW) frame can be defined by

$$\vec{v} = \frac{\mu}{h} \begin{bmatrix} -\sin f \\ e + \cos f \\ 0 \end{bmatrix} \tag{42}$$

where f is the true anomaly, e is the eccentricity, μ is the Gravitational parameter of Earth, and h is the angular momentum of the orbit[16]. As such, a rotation angle θ between the PQW frame and NTW frame can be defined by

$$\theta = atan2(-\sin f, e + \cos f) \tag{43}$$

Therefore, in order to rotate from the PQW to the NTW frames, a rotation matrix can be defined as

$$^{NTW}\boldsymbol{R}^{PQW} = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{44}$$

The transformation from the PQW frame to the Earth Centered Inertial (ECI) frame is similarly accomplished using rotations defined by classical orbital elements. The derivation of this rotation is beyond the scope of this paper, and well defined in Reference [14, p. 175].

Using the above rotations, it is possible to convert the constant target attitude in the satellite normal frame, to a time-dependent target in the ECI frame. This target is applied to a Quaternion-Rate Feedback (QRF) controller to produce a desired control torque.

3.2   Magnetic Torques

Applying a quaternion-rate controller, or similarly using an optimal rate controller such as a linear-quadratic regulator, to a rotating spacecraft will produce a desired control torque vector. Typically, a reaction control actuator would produce the desired control torque reorienting the spacecraft towards the target state. In the case of magnetic torqueing, the space of producible torques is restricted to a plane perpendicular to the local magnetic field vector[47]–[49].

Using orthogonal oriented magnetic torqueing coils, it is possible to generate an arbitrarily oriented magnetic moment. The torque generated by the coil will be defined by

$$\vec{T}_{mag} = \vec{\mu}_{sat} \times \vec{\beta}_{local} \tag{45}$$

where $\vec{\mu}_{sat}$ is the net magnetic moment produced by the satellite and $\vec{\beta}_{local}$ is the local magnetic field vector. As a result of this property, the space of possible torques consists of the space perpendicular to both $\vec{\mu}_{sat}$ and $\vec{\beta}_{local}$. As the magnetic moment can be oriented arbitrarily using a set of three orthogonal torqueing coils, the space of possible torque vectors is not restricted by magnetic moment; however, the direction of the Earth's magnetic field is fixed. Thus, the space of possible torques is restricted to the space of vectors perpendicular to the local magnetic field.

As such, the satellite will seldom be able to generate the optimal torque to perform a reorientation maneuver. If the desired torque lies near the plane perpendicular to the local magnetic field vector, is it possible to generate a torque very similar in direction to the desired torque; contrastingly, should the desired torque lie far from the plane (such as if it were aligned

near the magnetic field vector itself), any producible torque using magnetic torqueing coils would actuate a near right angles to the desired control effort.

Since a QRF and, similarly, linear quadratic regulators (LQR) controllers do not take into account the space of producible torques, the disparity between the producible and desired torques must be addressed[47], [50]. To do so, the satellite scales the desired torque based on a projection onto the plane perpendicular to the magnetic field, as in Figure 22. As such, the magnetic moment is scaled such that the output torque is minimized when the angular difference is large, and maximized when the angular difference is small. In such a way, a torque is only applied to full effect when the possible torque is primarily in the direction of the desired torque. This also guarantees that energy is not expended when it will do little to reduce the angular rates according to the control law.
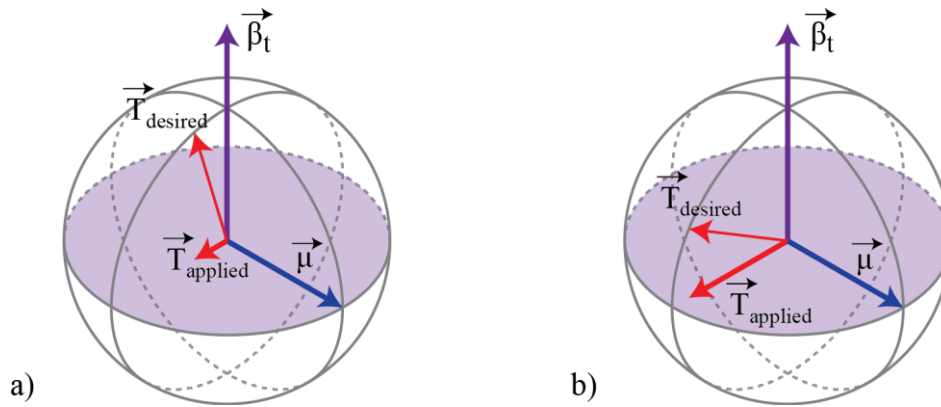
Figure 22: Two examples of difference between desired torque and possible torque based on orientation of local magnetic flux vector. The green disk represents the plane perpendicular to the local magnetic field flux vector. In the case where the desired torque.

3.3 Control Methods

Multiple different control methods are implemented to address the varying requirements during the mission. Most notably, the amount of information available about the spacecraft's

state differs over time. In general, the state information is more limited and less precise early in the mission and generally increases overtime; however, this trend is not without exception as is the event of a system reboot. Each control method will be described, along with an investigation of the limitations of each control system.

3.3.1   B-Dot and B-Cross Control

Upon deployment, the satellite does not possess knowledge of its orbital location, attitude, or angular velocity. Both the Triad and EKF attitude determination methods, discussed previously, require a reference sun vector and/or a reference magnetic field vector to estimate the attitude. Without knowledge of orbital location, reference magnetic field and sun vectors cannot be obtained making it difficult to estimate attitude. While attitude rate can be obtained from onboard gyroscopes without knowledge of orbital position, the attitude rate information is of limited use without attitude information.

During launch, a CubeSat is typically powered down and disconnected from any communication[51]. As a result, it is impossible for the satellite to possess accurate orbital parameters once deployed and powered on. The satellite must first make contact with a ground station and receive telemetry with its orbital parameters before it can determine its orbital position, then estimate attitude and attitude rate.

If the attitude rate is too high, however, it is difficult for the satellite to make contact with a ground station[52]. This is especially true for satellites with high-gain antenna systems which require accurate pointing to complete a downlink/uplink. In order to ensure contact can be made, the angular velocity must at a minimum be reduced prior to establishing communication. As such, a controller is needed that – by requirement – does not require knowledge of orbital

position or attitude. The controller must also be able to reliably and efficiently reduce the angular velocity using only the information that it can collect locally about its environment.
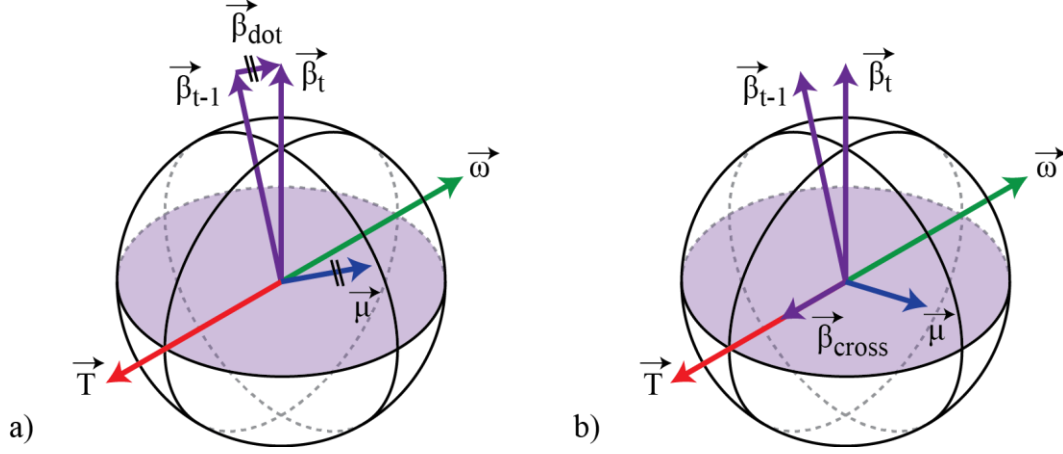
Figure 23: Comparison of B-Dot and B-Cross detumbling algorithms.

To address these requirements, a B-Dot controller is typically implemented. In its simplest definition, a B-Dot controller seeks to reduce the apparent rate of change in magnetic field direction[53]–[56]. If the apparent change in magnetic field direction is assumed to be purely a result of spacecraft angular rate (i.e. the direction of the magnetic field is assumed to not be changing rapidly with orbital position), then the angular velocity vector would be perpendicular to the place swept out by the magnetic field at two sequential measurements. To reduce the angular velocity, a torque vector must be generated opposite to the angular velocity vector. In order to determine the applied control, the magnetic field can be measured locally at any time using the magnetometers on board the spacecraft. The vector difference between the present magnetic field vector measurement, $\vec{\beta}_t$, and the previous magnetic field vector measurement, $\vec{\beta}_{t-1}$, is scaled by a proportionality gain, α, to produce the desired magnetic moment as expressed in Equation (46). For small changes in magnetic field direction, the resultant magnetic moment is nearly perpendicular to the current magnetic field direction. Since torques can only be

applied perpendicular to the current magnetic field direction when using magnetic torqueing, the desired torque can always be applied when using magnetic torqueing.

$$\vec{\mu}_{BDot} = \alpha\vec{\beta}_{dot} = \alpha(\vec{\beta}_t - \vec{\beta}_{t-1}) \tag{46}$$

$$\vec{T}_{BDot} = \vec{\mu}_{desired} \times \vec{\beta}_t \tag{47}$$

For larger changes in the magnetic field direction, the method of taking the linear difference of sequential magnetic field measurements becomes less efficient. In the extreme, if the measured magnetic fields from two subsequent measurements directly oppose one another, the resultant torque would be zero as the magnetic moment (determined by Equation (46) and the local magnetic field are aligned. In fact, for any angle of offset between magnetic field measurements, there exists another magnetic moment vector which will result in a larger torque than that produce by the magnetic moment vector specified in Equation (46).

To remedy this shortcoming, the B-Cross controller is proposed. Rather than taking the linear difference, the resultant desired torque is calculated as the cross-product of the two subsequent magnetic field measurements. To recover the magnetic moment needed to produce this torque, the present magnetic field vector is crossed with the torque. Since the magnetic moment and local magnetic field are by necessity perpendicular, the resultant torque will always be maximized. Like the B-Dot controller, the resultant torque will increase as the angular differences in magnetic field measurements increases. However, the commanded magnetic moment will be optimal for the resultant torque. This does come at the cost of a slightly more complex computation of the magnetic moment; however, the increase in complexity is not significant for modern processors.

$$\vec{T}_{BCross} = \gamma \frac{(\vec{\beta}_t \times \vec{\beta}_{t-1})}{\|\vec{\beta}_t\|} \tag{48}$$

$$\vec{\mu}_{BCross} = \frac{\vec{\beta}_t \times \vec{T}_{BCross}}{\|\vec{\beta}_t\|} \tag{49}$$

The controller, by requirement, only knows the change in magnetic field vector relative to the body of the craft. The controller assumes that, the change in magnetic field over the sample period is due predominantly to the rotation of the body relative to the inertial frame. Inherent to this assumption is that the local magnetic field is static relative to the inertial frame over that same period. In truth, the local magnetic field varies due to the satellite traversing its orbit and thus crossing through a non-rectilinear magnetic field. Over the course of the orbit, the magnetic field vector changes dramatically; however, over a sample period of 10 seconds, the field remains relatively stationary. Additionally, the magnetic field itself can fluctuate; however, below 300km, changes due to fluctuations in the magnetic field are less significant and can generally be ignored[46]. Therefore, the core assumption is that, over the angular velocities to which the controller is applied, the magnetic field is an inertial fixed vector about which the angular velocity of the spacecraft can be measured.

A further assumption made by the controller is that two successive magnetic field measurements sweep a plane whose surface normal is the angular velocity vector. As discussed, changes in the magnetic field orientation over the orbit will also contribute to the change in magnetic field measurements. Measurement error will also contribute to changes in the field measurement; at low rates of spin, this can substantially affect the perceived angular velocity direction. The assumption also requires that the same rate is sufficiently fast that two samples are guaranteed to be in the same half-plane to avoid aliasing. Should the craft be rotating fast enough

relative to the magnetic field that the craft spins greater than 180° between two samples, it is impossible to determine the direction of spin using magnetometers alone (this is true for both B-Dot and B-Cross controllers). As a result, there is an upper limit on angular rate which can be controlled based on the sampling frequency. For the assumed case of a maximum 5°/sec tip-off rate, a sampling frequency of no less than 0.028hz is required (36 second period between samples).

The limitations with B-Dot/B-Cross control lay with its core assumption that the magnetic field can be assumed as inertially fixed over two successive samples. For high inertial angular velocities, the measured changes in magnetic field direction are predominantly due to the rotation of the spacecraft in a relatively static magnetic field. Opposing this motion with a torque will decrease the inertial angular velocity of the satellite, resulting in negative feedback as desired. For very small spin rates, the measured changes in magnetic field direction result more from the magnetic field changing orientation as the satellite traverses its orbit than from the spin of the satellite relative to the inertial frame. From the perspective of the satellite, the change in magnetic field direction is perceived as the satellite rotating in the opposite direction relative to an assumed static magnetic field. In fact, there will be an instance whereby the observed magnetic field change becomes dominated by changes in magnetic field orientation along the orbit rather than by the spin of the satellite relative to the inertial frame. Due to the under-actuated nature of the magnetic torqueing system, this will always occur while there is still residual angular velocity. As a result, the controller applies a torque that drives the satellite in the direction of residual angular velocity, resulting in positive feedback from the perspective of the inertial frame. This phenomena is demonstrated in simulation in Figure 24 and Figure 25.
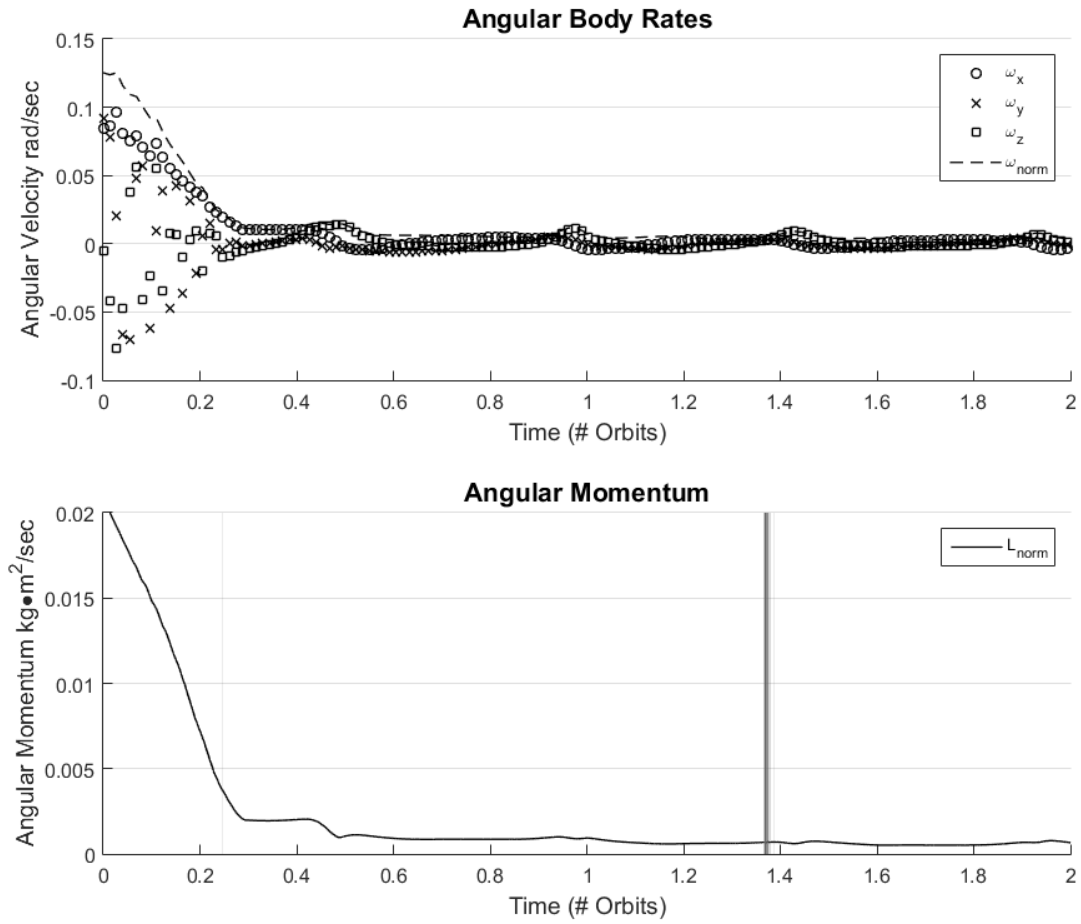
Figure 24: Control response of B-Cross rate control algorithm with zero magnetometer noise. Incidence of positive feedback shown as shaded bars (concentrated just prior to 1.4 orbits). Initial angular velocity is approximately 7°/sec, well in excess of the maximum expected tip-off rate of 5°/sec.
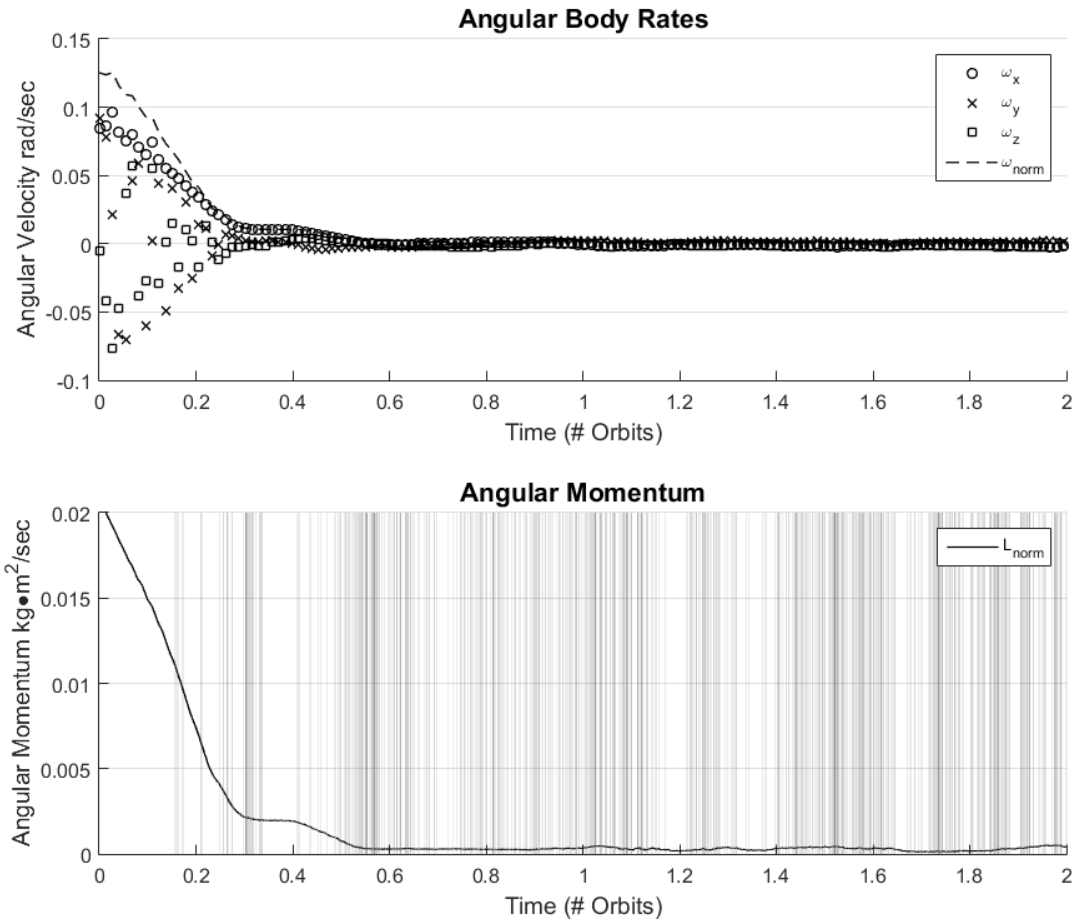
Figure 25: Control response of B-Cross rate control algorithm with simulated magnetometer noise. Incidence of positive feedback shown as shaded bars (widely distributed along the orbit). Same initial conditions used as for the no-noise case.

The positive feedback causes the satellite to accelerate in the direction of the residual spin, adding to the craft's angular momentum. This continues until the contribution to magnetic field vector change from the angular velocity of the craft begins to dominate again – either by accumulation of angular velocity or by a reversal in the direction of change of the magnetic field vector. The accumulation of angular velocity (and thereby angular momentum) by positive feedback requires additional torque to reduce the inertial spin rate to zero.

Positive feedback at low inertial angular velocities represents the inherent limitation to B-Dot and B-Cross control. This limitation results from the assumptions of the controller that the

Earth's magnetic field is static. Since, by requirement, the B-Dot and B-Cross controller possesses no knowledge of the orbital location, attitude, or inertial rates, it cannot be improved upon with additional knowledge of the state. Using measurements from gyroscopes, it is possible to discriminate whether small changes in magnetic field are in fact due to changes in the magnetic field or due to changes in angular momentum due to disturbance torques; however, the inherent limitation is easily mitigated by switching to a different type of controller before the onset of positive feedback is possible (so long as the satellite has been able to obtain additional telemetry to give it an enhanced knowledge of its state). Assuming a circular orbit of 300km, the rate of change of magnetic field contribution from progression along an orbit is 2.11E-3 radians/second at maximum (0.012 °/sec). As a result, the B-Cross controller should not be used when the perceived change in magnetic field direction drops below that rate.

Filtering the magnetic field measurements with respect to the predicted field based on the applied torque could improve the controller. By building a history of the orbit, and using simple assumptions about the size of the orbit, it could be possible to devise a method whereby the controller could 1) improve the rate at which it reduces its angular velocity, and 2) avoid instances of positive feedback. This controller would be more robust than the standard B-Dot controller and would be an improvement on the existing state-of-the-art. This concept will be discussed in greater detail in the research aims.

3.3.2 Quaternion-Rate Feedback

Quaternion-rate feedback (QRF) controller is a form of PID controller. The feedback law for QRF is stated as

$$\vec{M}_{desired} = -\kappa\vec{q}_e - \gamma\vec{\omega}_e \qquad (50)$$

where $\vec{M}_{desired}$ is the desired applied torque vector, $\vec{\omega}_e$ is the angular velocity error in the inertial frame, and $\vec{q}_e$ is the quaternion error[17]. κ and γ are positive scalar values. Quaternion error derived by Equation (51), wherein q is the current attitude and q_c is the desired attitude.

$$\begin{bmatrix} q_{1e} \\ q_{2e} \\ q_{3e} \\ q_{4e} \end{bmatrix} = \begin{bmatrix} q_{4c} & q_{3c} & -q_{2c} & -q_{1c} \\ -q_{3c} & q_{4c} & q_{1c} & -q_{2c} \\ q_{2c} & -q_{1c} & q_{4c} & -q_{3c} \\ q_{1c} & q_{2c} & q_{3c} & q_{4c} \end{bmatrix} \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{bmatrix} \tag{51}$$

The desired quaternion, q_c, and the angular rate target, ω_c from which the rate error ω_e is calculated, are constant in the NTW frame; however, the feedback law Eq. (50) is specified in an inertial frame (nominally the ECI frame). It is necessary to transform between frames these frames in order to derive the time-varying ECI targets based on the constant NTW targets. The transformations are derived in Appendix C.

As stated previously, the $\vec{M}_{desired}$ may not be achievable based on the orientation of the magnetic field. This is a consequence of fact that the satellite cannot directly generate a torque (such as through a reaction control system) but instead can only generate a magnetic moment to induce a torque on the spacecraft. Additionally, the coils have a maximum magnetic moment which can be generated, as such the controller may request a moment which must be scaled proportionally. The output magnetic moment is simulated by

$$\vec{\mu} = \frac{\vec{\beta}_{measured} \times \vec{M}_{desired}}{\left\| \vec{\beta}_{measured} \right\|^2} \tag{52}$$

In Equation (52), the distinction must be made that the magnetic field used to produce the magnetic moment is the measured magnetic field (that is the body fixed magnetic field measurement know to the satellite). It is also important to clarify that values in Equation (52) are all in the satellite body frame.

Since the torque coils have discrete outputs, the simulated magnetic moment is scaled uniformly if any axis exceeds the 0.2744 Nm/T at peak output. Additionally, the resultant moment is rounded to the nearest precision step to simulate the discrete output levels of the ADCs on the magnetic torqueing coil boards. The simulated satellite magnetic moment is then rotated into the inertial frame and the torque is calculated against the propagated magnetic field.

3.4 Hybrid Controller

Even when optimized, the QRF controller is not as efficient as reducing the angular rate as the B-Cross controller. In order for the controller to finely tune the pointing, it is difficult to simultaneously arrive at a set of weights which can efficiently reduce angular rate. One solution would be to use gain scheduling whereby the QRF controller could have one set of gains to reduce the angular rate then switch to another set of gains to achieve fine pointing once the angular rate has fallen below a prescribed angular rate threshold; however, this does not address the sub-optimal torque application of the QRF controller and adds complexity to the controller design.

Another approach is proposed to instead switch control algorithms. Since B-Cross is known to be optimal in angular rate reduction, it is an excellent controller to detumble the satellite. Once it falls below a prescribed threshold, the control algorithm can be changed to a QRF controller optimized for fine pointing.

The optimization of the hybrid controller requires tuning the quaternion error gains and rate error gains of the QRF controller, the gain of the B-Cross controller, as well as the angular rate threshold at which the controller switches over. Figure 26 shows a schematic representation of the hybrid controller. Depending on the satellite attitude state – specifically the attitude rate – as

determined by modelling of the satellite dynamics, the satellite will either apply the QRF controller or the B-Cross controller. Both systems will output a magnetic moment which is scaled and oriented according to Equation (52).
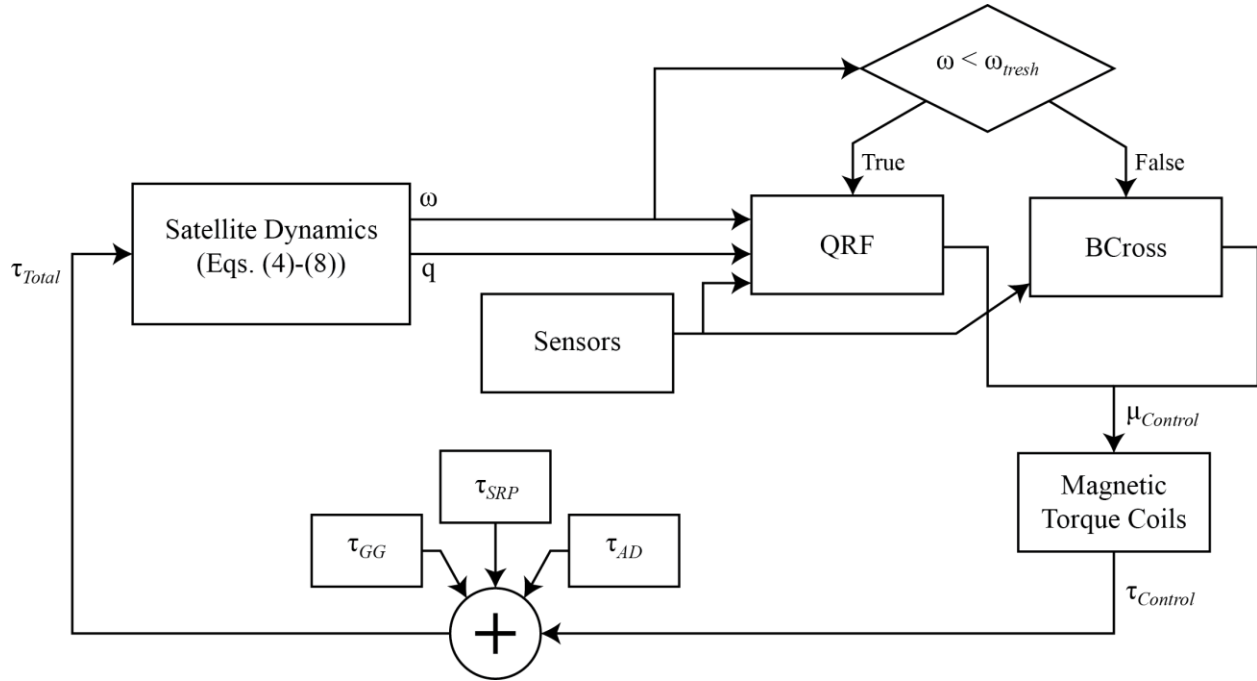
Figure 26: Attitude control diagram for Hybrid Controller.

A hysteresis threshold is also added to ensure that when the control switches from B-Cross to QRF, it does not immediately switch back if driven to a slightly higher angular rate by the QRF controller. These gains for the controller, as well as the low-rate control threshold value, were optimized simultaneously.

3.5   Controller Optimization

3.5.1   Particle Swarm Optimization

A particle swarm optimizer (PSO) was used to optimize the various gains of the QRF controller. The PSO functions by seeding a set of particles with initial gain parameters sampled from a uniform distribution bounded by upper and lower limits for each gain[57]. Each particles

is seeded with a randomly generated velocity which dictates by how much each parameter will vary at each successive iteration of the optimizer. At each iteration, the gain parameters are used to calculate a cost for each particle. Each particle tracks its own best set of gain parameters and cost observed to date, while the optimizer tracks the best set of gain parameters and its cost across all particles. The velocity of each particle is then modified based on its distance from its own best set of gains as well as its distance from the global best[58]. The velocity modifies the new gain parameter sample in the following iteration. In this manner, the PSO continually perturbs particles in the direction of the best solution to attempt to find better solutions. After all particles converge to a single value (to within some distance ϵ) or a specified number of iterations has elapsed, the PSO returns the gain parameters and cost of the best solution.

3.5.2 Cost Function

Each set of control gains was applied to a set of attitude control simulations. The attitude control simulations each have different initial conditions $(t_{init}, q_{init}, \omega_{init})$. The performance of each set of control gains was determined by an equal weighting of the costs of each simulation, with the weight of each simulation being computed as:

$$J(t_{init}, q_{init}, \omega_{init}) = t_{converged} + \sum_{t=0}^{T}(\|\vec{q}_{error,t}\| + \|\vec{\omega}_{error,t}\|)\Delta t \qquad (53)$$

The $t_{converged}$ term is the time required for the attitude to converge to within some small value ϵ of the target attitude state. This is the dominant term in the cost function. It is computed by keeping a running average of the attitude pointing error using a sliding window of 10 samples. At each update, the average is compared to threshold ϵ; when the average is less than the threshold, the value of $t_{converged}$ is updated. If the threshold is later exceeded, the value is

updated to the current time. In this fashion, if the determination algorithm remains below the threshold until the end of the simulation, then the time of convergence is returned in the cost; if the control algorithm is not able to converge (or only converges briefly before becoming unstable), the simulation end time will be reported.

Since the simulation runs for a finite time and makes discrete, constant time steps, the value of $t_{converged}$ will be identical for multiple solutions. Using $t_{converged}$ alone is disadvantageous as an optimization cost as it does not provide a way to discriminate between adjacent solutions. This is especially true when beginning the optimization process where the sampled control weights may not result in converged solutions. In such cases, all solutions would have a cost equal to the simulation length.

To remedy, a measure of the controller error over time is added. The terms $\vec{q}_{error,t}$ and $\vec{\omega}_{error,t}$ are the quaternion error and the angular rate error respectively. Both variables are the errors at simulation time *t*. The quaternion error is the first three components of the error quaternion, the norm of which is zero for perfect alignment with the target attitude. By summing the errors at each time step, the optimizer is able to discriminate between adjacent solutions for time convergence remedying the problem of using discrete time steps.

Since the quaternion error is itself a quaternion and therefore unity norm, the norm of the first three components must be less than or equal to unity. Unless pointed in direct opposition to the target attitude, the quaternion error norm will be less than unity. While the norm of the rate error is unbounded, practically it should not exceed the 5°/sec maximum tip-off rate which bounds $\left\| \vec{\omega}_{error,t} \right\|$ at 0.087 when expressed in radians/sec. As such, the size of the state error term in the

cost function will be in practice less than the size of $t_{converged}$ term. This is especially true for well controlled system as the state error term will be effectively zero for converged solutions.

No cost contribution from the control effort was considered for the control optimization. While the system is power restricted, the science requirements for accurate pointing and the under actuate nature of the magnetic torque control system require time optimal solutions. Post-optimization analysis of the attitude control system show that the depth of discharge for the attitude system do not exceed power budget allotment.

Multiple sets of initial condition for the attitude control simulations were used to avoid convergence around points which coincidentally perform well for a single set of initial conditions. For example, using B-Cross control alone, it is possible to reduce the angular rate to near zero while coincidentally reducing the attitude pointing error to near zero given a fortunate set of initial conditions. As discussed previously, no aspect of the B-Cross control guarantees a convergences to zero pointing error; thus, solutions of this nature are coincidentally optimal due to fortunate initial conditions. By choosing multiple sets of initial conditions and optimizing across all sets as a whole, convergence about coincidentally optimal cases can be mitigated. As described at the beginning of this section, the optimization cost is represented by the sum of n cost functions defined by Equation (53) using n sets of initial conditions (see Equation (54)).

$$J_{optim} = \sum_{x=1}^{n} J(t_{init,x}, q_{init,x}, w_{init,x}) \tag{54}$$

While the initial conditions vary, each set represents a reasonable set of initial conditions. Broadly, the initial conditions can be divided into two categories: detumbling scenarios and reorienting scenarios. Detumbling scenarios are typified by higher initial angular rates; initial

pointing error is less significant since the high angular rates will cause the initial pointing error to vary rapidly once propagated. Detumbling scenarios are most representative of the higher angular rates observed post-deployment. Reorienting scenarios, by contrast, are most representative of step changes in the desired pointing direction. The initial conditions for reorienting scenarios have low or zero angular rate with substantial attitude pointing error. Based on the concept of operations for the LAICE mission, both types of attitude control initial conditions will be observed. Optimizing for both types of initial conditions sets (detumbling and reorienting) guarantees a set of control weights which is robust.

Table 2: Initial conditions used for various sub-runs of optimization of the QRF controller.

Simulation Initial Condition	Case 1	Case 2	Case 3	Case 4	Case 5
t_0 (seconds Unix Time)	1453991110	1454575000	1456750042	1454572060	1454575000
ω_1 (rad/sec)	0.0872665	0.0726555	-0.079350	0.0	0.0
q_1	0.4	0.0	0.0	0.8	1.0
Scenario Description	Detumble 1	Detumble 2	Detumble 3	Reorient 1	Reorient 2

In addition to the initial attitude rate and pointing error, the initial time of the simulation was also varied. The orbit was kept consistent across all simulations, therefore changing the initial time of the simulation is analogous to changing where in the orbit the spacecraft is. As the magnetic field direction varies along the orbit, the space of controllable torques will also vary along an orbit. By changing the simulation state time and leaving the initial angular rate and pointing error unchanged, the control response will likely differ as the space of available torques will be different. As such, varying the initial time is necessary to ensure that coincidentally optimal solutions are avoided.

The International Space Station (ISS) orbit parameters were used for all simulation since this is the designated deployment orbit. As the deployment time, angle, or velocity will not be known until the LAICE spacecraft is already aboard the ISS, any attempt to modify the orbit in simulation would be speculative and no more representative than assuming the original ISS orbit.

3.5.3 Optimization Results

The optimization was conducted with 100 particles over 10 iterations. The bounds for each parameter were determined based on optimizing over multiple ranges with small samples of particles. The search regions were reduce to focus about the minima from previous optimizations until the structure of an optimization minima for each parameter was resolved in sufficient detail. This process was necessary to ensure the final search region encompassed the optimal solutions without containing extraneous regions which would slow the optimization. The resulting search regions are detailed in Table 3. The best particle at the end of 10 iterations is considered to be the optimized solution weights, though a more optimal solution may still exist.

Table 3: Optimization Parameter Search Regions

Simulation Parameter	Lower Bound	Upper Bound
Quaternion Error Gain, κ	0.0	1.0e-6
Angular Rate Gain, γ	0.0	4.0e-3
Low Rate Threshold,	0.0	0.02

3.5.3.1   Optimization of Control System with Perfect Determination

The first demonstration of controller optimization uses perfect determination. Rather than being fed the output of the determination algorithm, the controller uses the a priori attitude state provided by the simulation. Using the simulated attitude as an input is a best-case scenario as a flight control system would not have access to the true state information and therefore must rely on the value form the determination algorithm. This determination will be subject to both sensor

and process noise. The simulated satellite was, however, still subjected simulated sensor noise when determining the local magnetic field to calculate the commanded magnetic moments. Noise on the sensors was again simulated as zero-mean, Gaussian noise. The state was propagated against the truth state; therefore, the control torques were applied using the truth magnetic field and not the simulated sensor magnetic field values.
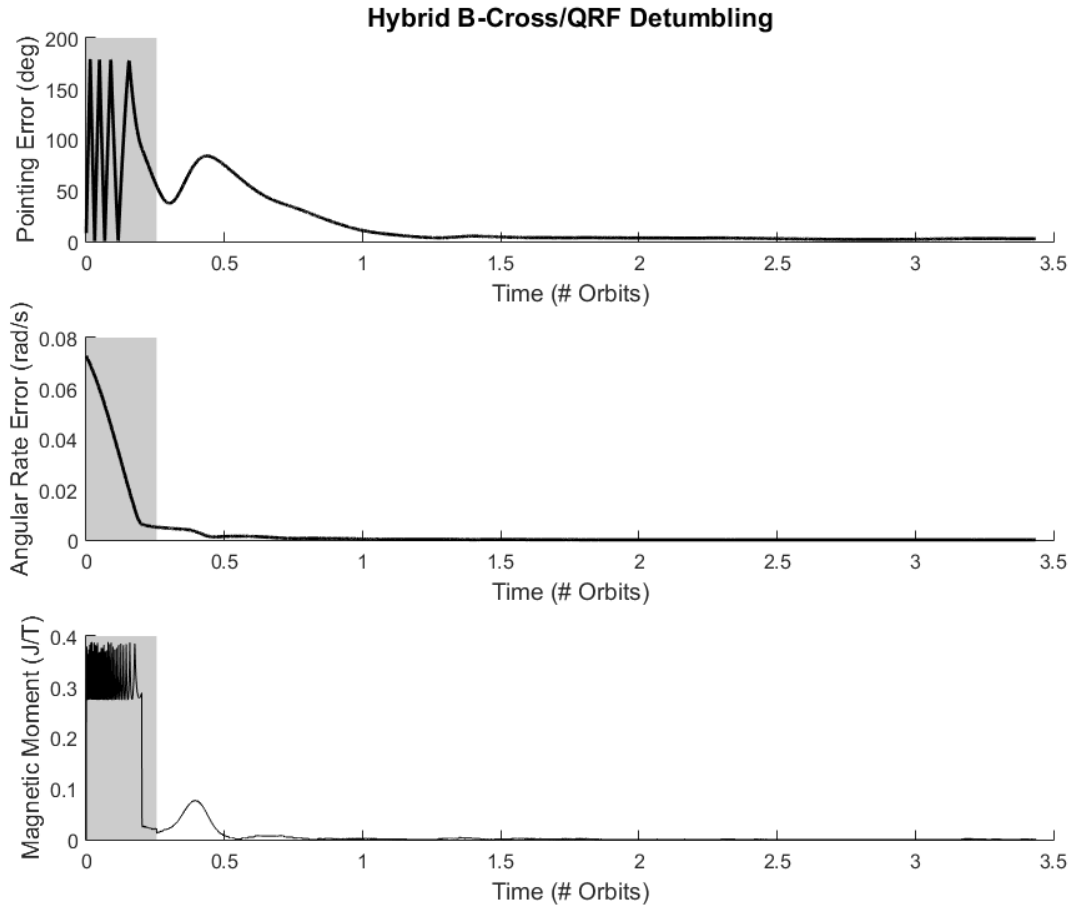
Figure 27: Control response for a detumbling maneuver using control gain set from Particle Swarm Optimization. Shaded background indicates when the system is using the B-Cross controller; White background indicates when the satellite is using the QRF controller.

Following the Particle Swarm Optimization, the best controller weights were applied to a new set of initial conditions. The first set of initial conditions (Figure 27) represent a detumbling maneuver with high initial angular rate. The second set of initial conditions (Figure 28) represent

67

a reorienting maneuver where the initial attitude rate is near zero, but the initial attitude error is large. In both cases, the optimized weights were able to reorient the spacecraft and maintain the target attitude.
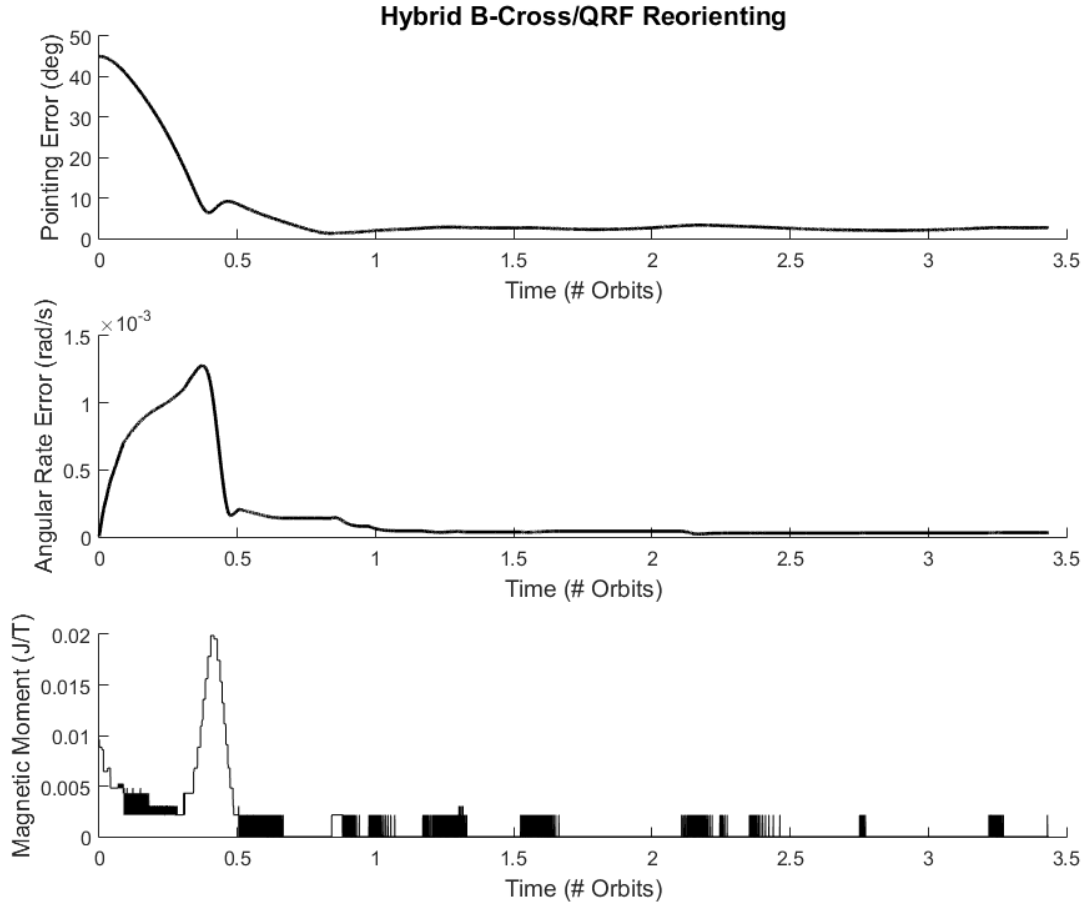
Figure 28: Control response for a reorienting maneuver using control gain set from Particle Swarm Optimization. In this instance, the system never switches to using the B-Cross controller due to the angular rate threshold having never been crossed.

In the case of the detumbling maneuver (Figure 27), the B-Cross controller is used for the first ¼ orbit (represented by the region of gray shading on the plot). The use of B-Cross is due to the high initial angular rates. During this period, the magnetic control moment is saturated as the satellite uses the maximum available torque to despin the satellite. Once the high angular rates drop below the threshold – a threshold determined by the optimization – the QRF controller is

used to reduce the attitude pointing error alongside the attitude rate. Within 1 orbit, the controller is able to pointing error to 10° and the rate error to near zero. By 1.5 orbits, the controller has achieved science pointing requirements of 5°. The control effort used by the QRF controller is much lower than that of the B-Cross controller. It also varies with not only pointing error but the orientation of the magnetic field. It is possible to see that between approximately 0.65 and 1.3 orbits, almost no magnetic moment is generated due to unfavorable alignment of the magnetic field with the desired torque vector. When the magnetic field is a near right angles to the desired torque vector, the controller quickly reduces the pointing error.

In the case of the reorienting maneuver (Figure 28), the attitude of the spacecraft is such that it is pointed approximately 45° from the target orientation; however, the initial angular rate is nearly zero. In order to reorient, the QRF controller increases the angular rate to adjust the pointing of the spacecraft. Since, the angular rates remain low, the system never crosses the threshold to trigger the B-Cross controller. Once the angular rate has sufficiently increased to adjust the attitude of the satellite towards the target attitude, the QRF controller quickly reduces the angular rate error to zero. The optimized system is able to complete the reorientation in slightly less than 1 orbit. During this entire operation, the magnetic moment applied is very low – so much so that the discrete levels of the simulated magnetic torqueing coils are visible in the plots. This demonstrates the necessity for simulating the discrete output levels of magnetic torqueing coils as a continuous output would allow for artificially rapid convergence.

3.5.3.2 *Optimization of Control System with Onboard Determination*

The same optimization was run, this time using the same determination scheme discussed in Chapter 2 as the input to the controllers. Without perfect knowledge of the satellite attitude state (quaternion error and attitude rate error), it is more difficult for the controller to properly reorient

the spacecraft. It should be noted that in the case of high initial angular rates, the B-Cross controller is unaffected by errors in the determination algorithm. Further, as previously discussed, higher initial angular rates assist in helping determination. As such, control cases with low initial angular rates are a greater indicator of system performance as they represent a worst-case scenario. The optimization was run against the same configuration of multiple initial conditions for a fixed attitude target in the ECI frame with a target of no angular body rate with respect to the ECI frame. The same cost function was applied. The optimization parameters were unchanged from the perfect determination control optimization. The optimization was run with 10 particles over 10 iterations.

While using the estimated covariance of the determination EKF could provide useful information about the certainty of the attitude state estimate, the control system does not make any use of this information. As such, the controllers consider the outputs of the EKF to be the truth values of the attitude state. While this has no impact on the B-Cross controller (which does not use the state estimate to condition its control), it does greatly affect the QRF controller which conditions its control output on the attitude error and angular rate error (both of which depend on the state information). As a result, the controller is expected to perform more poorly than in the case of perfect determination.

The resulting controller weights demonstrated that even with imprecise knowledge of the attitude state, the coupled attitude determination and control system was able to reorient within three orbits (Figure 29). The increased duration in the control response is in direct proportion to the additional time that it takes for the determination algorithm to converge (approximately 2.5 orbits). The determination algorithm, being itself not optimized, is therefore to blame for the increase in convergence time. To demonstrate this dependence, the same control gains

determined for the optimization of the coupled determination and control system were applied to a simulation with the same initial conditions but with perfect determination (similar to the previous optimization). The resulting plot (see Figure 30) demonstrates that the controller is able to reorient the simulated spacecraft within approximately ½ orbit.
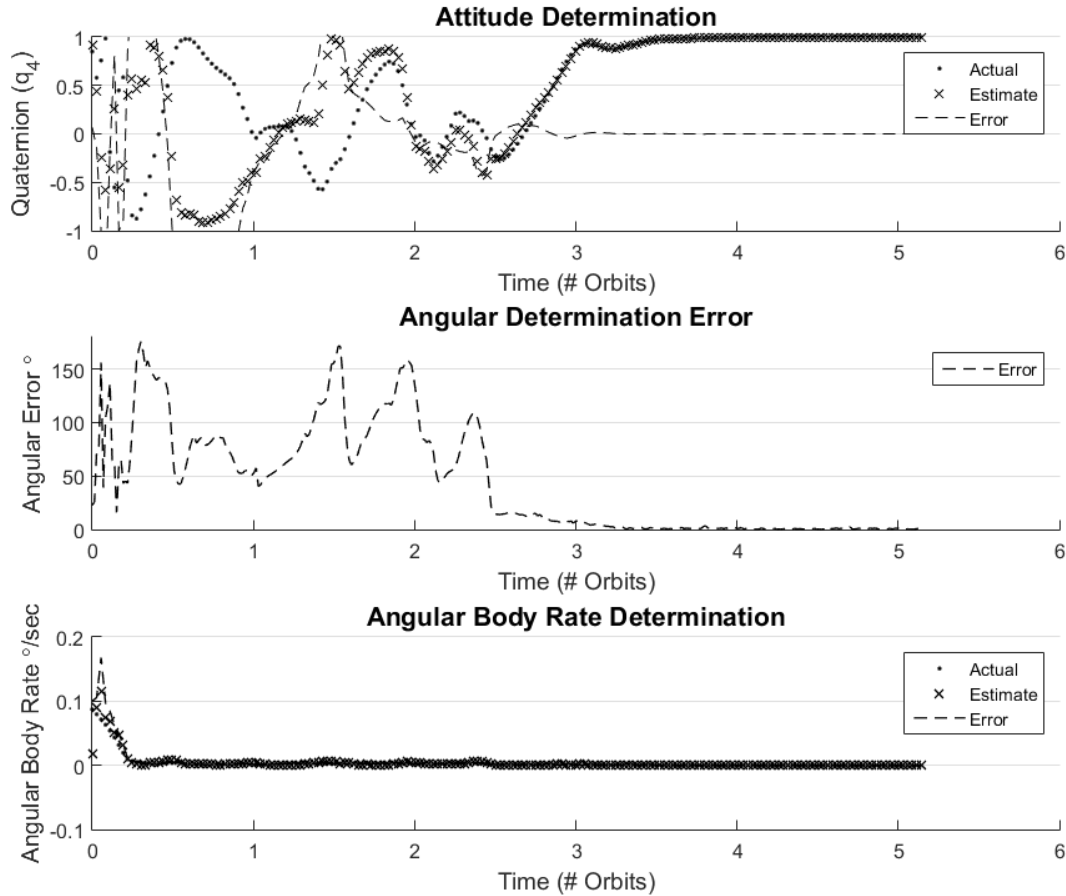
Figure 29: Results of simulation with coupled determination and control. This plot is the result of an optimization of the control system using the outputs of the determination system to condition the control output.

71
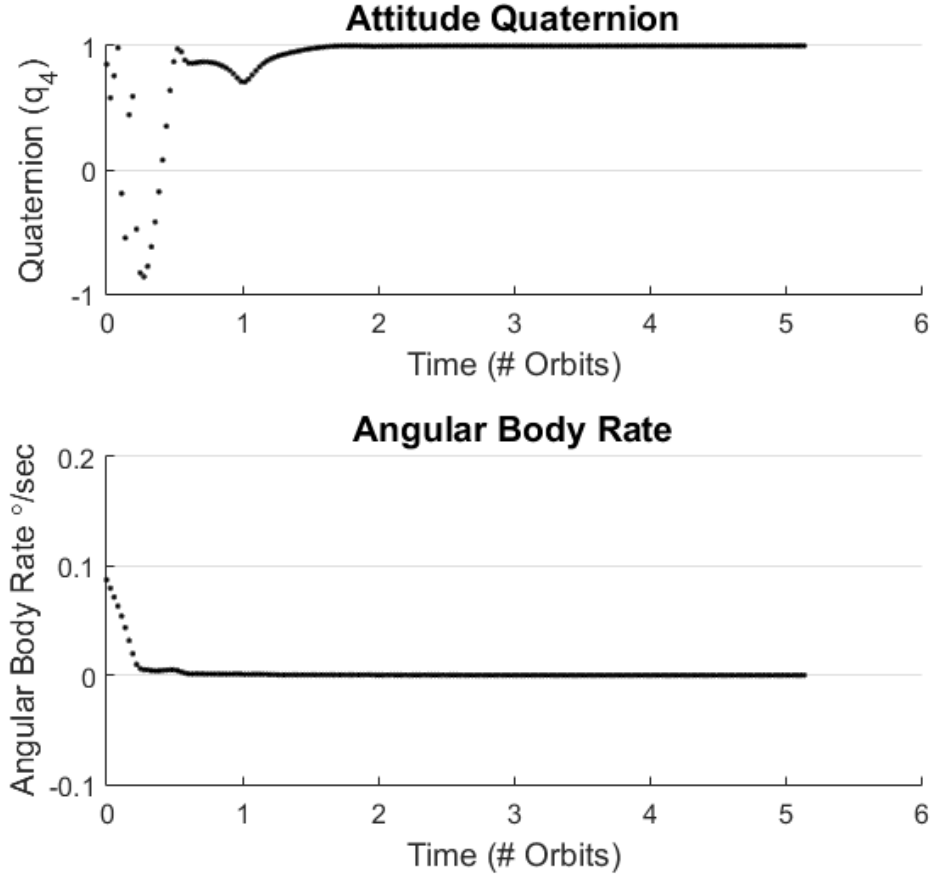
Figure 30: Demonstration of the effect of determination uncertainty on the control simulation. The same initial conditions and control gains, but applied to a system with perfect determination.

3.6 Summary of Contributions

The coupled control system, being robust to sensor noise and determination error, is capable of efficiently reducing the attitude error within 3 orbits even from high angular rates. This novel system accomplishes this by using an efficient B-Cross angular rate reduction algorithm until an optimized threshold is hit, whereby the system switches to a QRF control algorithm to reduce angular rate error and attitude pointing error. In terms of stability, the controller is piecewise stable[17], [53]. With the exception of oscillation between controllers at the threshold, which is mitigated by a hysteresis value, the hybrid controller is therefore stable as well. Further, the

novel application of a Particle Swarm Optimization gives credence that the controller gains are well tuned for the desired orbit.

CHAPTER 4
SATELLITE MASS PROPERTIES MEASUREMENT

Direct measurement of the satellite's mass properties is a difficult but necessary task for accurate modelling of the satellite dynamics. While CAD models provide estimates with sufficient accuracy to perform initial design studies, they lack the accuracy required for tuning control gains. Features which are difficult to accurately represent in CAD, such as cable harnesses, staking compounds, and conformal coatings, introduce error in the mass property estimates.

It is possible to measure the Gravitational Center (GC) and Moments of Inertia (MOI) on orbit during initialization of the satellite; however, this relies heavily on the accuracy of rate sensors on board. If the GC and MOI measured on orbit differ drastically from the ground based measurement, it would be a strong indicator that the configuration of the satellite has changed (likely due to damage during deployment). Therefore, measurement of the GC and MOI is critical to controller design. While commercial solutions exist for measuring both GC and MOI, they are expensive. In order to reduce mission cost while obtaining high precision mass property measurements, two instruments were designed to perform the measurements. The designs of these instruments help inform the precision of the GC and MOI values that can be obtained, and therefore the uncertainty which should be applied when simulating the satellite's attitude.

4.1   Gravitational Center Measurement

In order to measure the center of gravity, a four-point scale is used. The instrument (Figure 31 and Figure 32) consists of a large, flat aluminum tabletop that sits atop four strain-gauge load-cells (see Figure 33). The tabletop has a registration cross mark etched into the surface indicating the center of the table as well as the principal x-y axes. The component or system being massed

74

is placed on the surface, with its body axes aligned with the registration cross; the test subject can either be placed with its geometric center aligned with the origin (center of the registration cross) or aligned with its extremum along two edged of the registration cross. The table returns the value of the GC with respect to the center of the registration cross based on the differential voltage readings from the strain gauges.
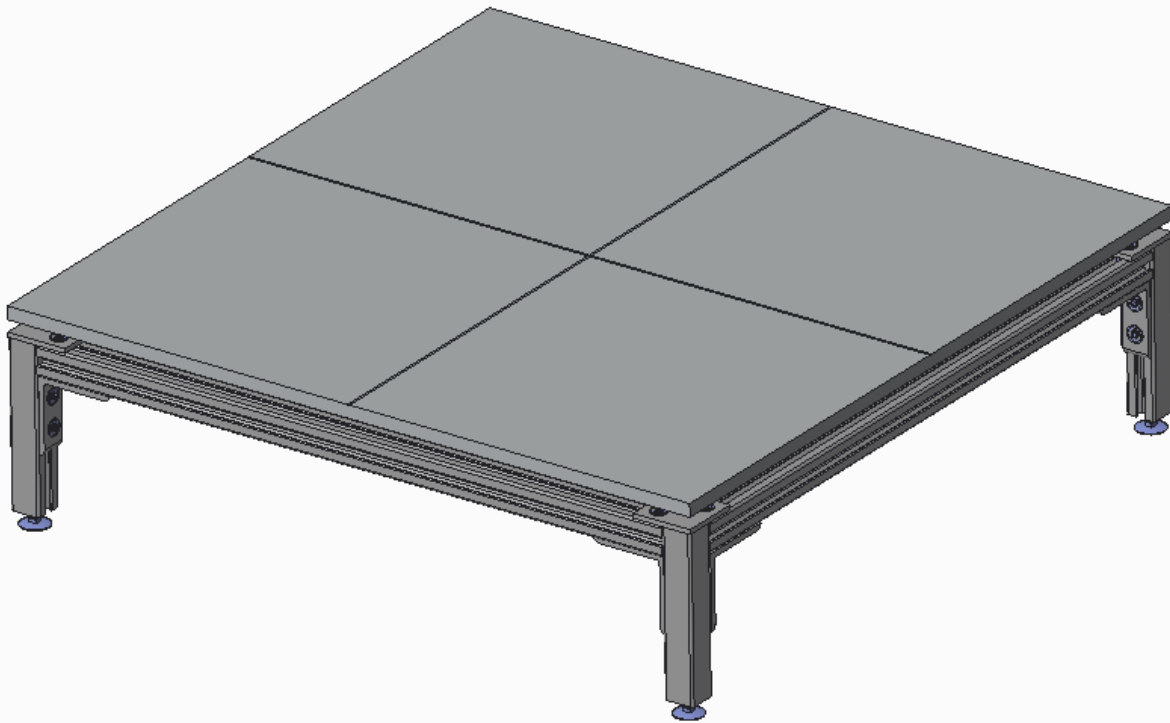
Figure 31: Gravitational Center (GC) Measurement Instrument CAD Assembly. For the purpose of computations, the corners are numbered sequentially from 1-4 counting clockwise from the left-most point in the above rendering.

Figure 32: GC Measurement Instrument in Laboratory
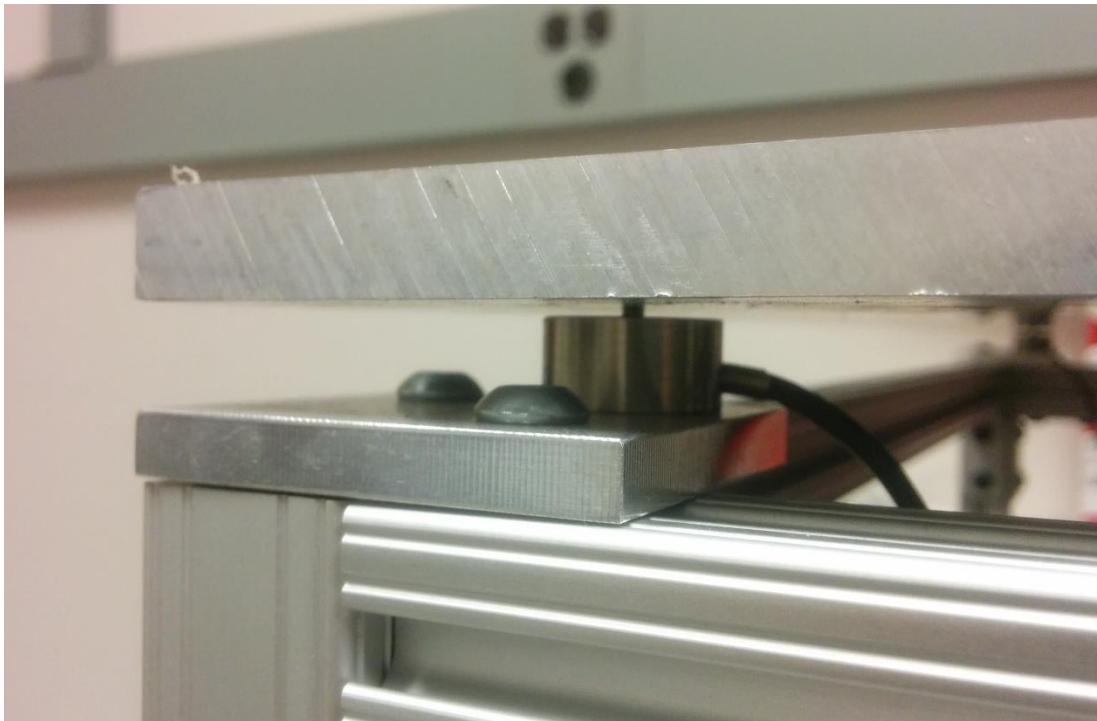
Figure 33: Close-up of one of the four load cells supporting the measurement platform.

4.1.1 GC Measurement Concept of Operation

Each of the four load cell includes a strain-gauge with variable resistance which is linearly proportional to its load. The strain-gauge is part of a Wheatstone bridge, which means the

voltage read from the sensor – the Wheatstone bridge differential voltage – is linearly proportional to the load[59]. The voltage response due to load must be determined through calibration for each load cell.

Prior to each measurement, the GC measurement instrument must be tared. During this step, the instrument will first determine if all four load cells are under load from the tabletop. If it is not, the user interface identifies which of the four corners is not under load (since at least three corners must be under load at all times) and instruct the operator to extend the foot on that corner thus raising the load-cell to meet the platform. Once the load cells are all under load, the instrument records the baseline voltage value. The baseline voltage represents the mass contribution of the aluminum plate. This must be removed from measurements off the instrument.

The center of mass is determined by evaluating the difference in load in the x-direction, then evaluating the difference in load in the y-direction. Using the total excess load, it is possible to evaluate the center of mass using Equation (55).

$$
\begin{bmatrix} x_{gc} \\ y_{gc} \\ M_{tot} \end{bmatrix} = \alpha_c \begin{bmatrix} l_c & l_c & -l_c & -l_c \\ l_c & -l_c & l_c & -l_c \\ 1/g & 1/g & 1/g & 1/g \end{bmatrix} \begin{bmatrix} V_1 - V_{1_0} \\ V_2 - V_{2_0} \\ V_3 - V_{3_0} \\ V_4 - V_{4_0} \end{bmatrix}
\tag{55}
$$

The coordinates of the center of mass of the test subject with respect to the center of the GC Measurement plane is defined by the above equation where α_c is the voltage proportionality constant in volts/N, l_c is the calibrated distance the sensors are from the plat sensor (nominally 30cm), g is the local gravitational acceleration (9.81m/s^2), V_n is the voltage read from the n[th] load cell, and V_{n_0} is the tare voltage for the n[th] load cell. The instrument returns the location of the

gravitational center in two axes $\{x_{gc}, y_{gc}\}$ with respect to the platform center, as well as the total mass (M_{tot}) of the system.

Measured in this manner, only two components of the gravitational center are resolved. If, for example, the LAICE spacecraft was placed with its x and y axes aligned with the plate x and y axes, the instrument would return the $\{x_{gc}, y_{gc}\}$ of the satellite (along with the satellite mass) though the z_{gc} would remain unknown. To remedy, a second measurement should be made with the satellite rotated 90° such that spacecraft's z-axis is aligned with the platforms x or y axis. If the satellite was placed with its x-axis along the platform axis, and the z-axis placed along the platform's y-axis, the instrument would return $\{x_{gc}, z_{gc}\}$ along with the satellite mass. The two measurements would completely resolve the satellite's GC, with redundancy in the x_{gc} and satellite mass measurements as validation.

4.1.2   GC Measurement Instrument Calibration

A series of proof masses are employed to calibrate the GC Measurement Instrument. The feet supporting the platform are adjusted to ensure that all load cells are in-plane with one another and level. Once the load cells are levelled, the output voltage of each load cell is measured as a baseline tare value. Once the tare values for each cell have been recorded, a cylindrical mass of known mass is placed in the center of the platform and the output voltage of each cell is measured again. This process is repeated for two additional proof masses. The output and baseline output voltages of the table are related to the force on the load cell. This relationship is assumed to be linear over the specified range of the table.

4.1.3 GC Measurement Instrument Accuracy Analysis

The full range of each load cell is 50kg. A 16 bit ADC will be used to measure the load cell resistance, which means that there will be a resolution of 131 LSB/kg. Conservatively assuming the 4 lowest bits of the ADC will be unreliable due to power supply noise, each load cell is capable of reliably resolving masses as small as approximately 6g. This translates to an error of ±0.006% of full scale. Based on the 600mm (±0.1mm) length between load cells, it is possible to resolve the center of mass of a 14kg satellite to within less than 0.26mm in each axis. This is similar to the accuracy of commercial instruments currently available for purchase. For a satellite the size of LAICE, it is reasonable that it could be located with 0.1mm accuracy on the platform, meaning the GC should be reliably resolvable to within 0.4mm. Half of this value will therefore be used as the precision with which the GC is known when determining how robust the system is to mass property uncertainty.

4.2 Moments of Inertia Measurement Instrument

The Moment of Inertia Measurement Instrument measures the moment of inertia (MOI) of the satellite or satellite components by using a torsion pendulum (Figure 34 and Figure 35). The unit consists of a circular platform of known inertia to support and retain the component being measured. The circular platform itself is supported by a rotational thrust bearing that allows the platform to rotate freely with minimal friction. A long, slender rod of uniform dimension is pinned at both ends: at the top, to the free-rotating circular platform, and at the bottom, to the fixed based of the instrument (see Figure 36). The components we designed and machined at the University of Illinois. The platform leverages extruded aluminum members for the majority of the structure. A circular optical breadboard was modified to receive the torsion rod fixture mount. The circular breadboard offers multiple ¼-20 screw attachment points at regular intervals

and its moment of inertia about its most significant axis of rotation is centered to within 0.001mm meaning.

4.2.1 MOI Measurement Concept of Operation

By rotating the circular platform, the slender rod undergoes torsion and stores up torsional energy[60].

$$\tau = -k\theta \tag{56}$$

$$I\ddot{\theta} + b\dot{\theta} + k\theta = 0 \tag{57}$$

$$\theta = e^{\beta t}\theta_0 \cos\left(\left(\sqrt{\omega_0^2 - \beta^2}\right)t - \phi\right) \tag{58}$$

$$\omega_0 = \sqrt{\frac{k}{I_{nn}}} \tag{59}$$

$$\beta = \frac{b}{2I_{nn}} \tag{60}$$
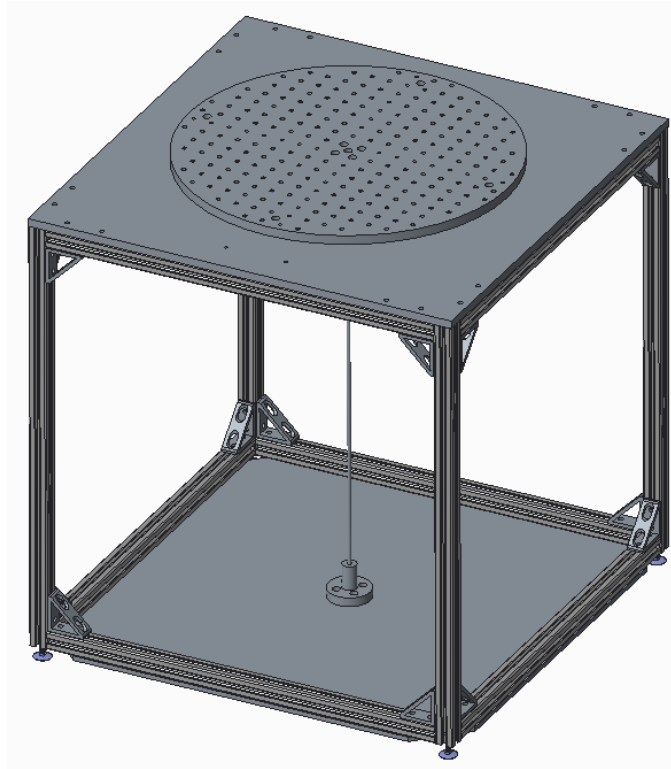
Figure 34: Moment of Inertia (MOI) Measurement Instrument CAD Assembly

Figure 35: MOI Measurement Instrument in laboratory.

Figure 36: Close-up of machined mount for retaining the torsion rod to the base. An identical mount is affixed to the measurement stage and rotates on top of a rotational thrust bearing.

The torque required to rotate the platform through θ degrees is linear with a proportionality constant known as the torque constant k (see Equation (56))[60]. When the platform is released, the torsion causes the platform to oscillate rotationally according to the ordinary differential equation described in Equation (57). The damping constant, β, as well as the torque constant, k, must be determined experimentally. Through calibration using proof-masses of uniform density and known mass and dimension, it is possible to determine these constants as well as the rotational inertia of the platform. The rotational inertia of the system, I_{nn}, is a composite of the rotational inertia of the platform and the rotational inertia of the test subject about the axis of rotation aligned with the platform axis of rotation. If the test subject's gravitational center is aligned with the axis of rotation ($R_{gc} = 0$), then the rotational inertia of the platform and the test subject are a linear sum (see Equation (61)). By measuring the angle of the platform over time, the moment of inertia of any test subject placed on the circular measurement platform can be determined by fitting Equation (58) to the data.

$$I_{nn} = I_{nn_{platform}} + I_{nn_{subject}} + m_{subject}R_{gc}^2 \qquad (61)$$

In order to measure the moment of inertia along any axis, it must be aligned with the center of gravity in that axis. A component mounted in such as fashion that the center of mass is not aligned with the axis of rotation will result in an eccentric force on the platform and rod. This is due to a Product of Inertia (POI) being created by offsetting the axis of rotation. It is possible to measure this force using strain gauges on the bearing mounts. Measurements of the rotational motion of the platform and the measurements from the strain gauges would allow for the simultaneous evaluation of the GC, MOI, and POI. This approach was deemed too expensive and unnecessarily complicated given that an instrument for evaluating the GC was already available.

To measure the angle, several methods may be used. The method presented is an optical angle measurement. A collimated light source is directed at a planar mirror affixed to the torsion rod just below the measurement stage. Being fixed to torsion rod, the planar mirror will reflect the light source based on its angle of rotation. For any rotation of the mirror θ, the light will be redirected 2θ from its original path[61]. By measuring the reflected light source using a CCD array, the angle of incidence can be determined.

4.2.2 MOI Measurement Instrument Accuracy Analysis

The first source of error in the determination of the MOI of a test subject comes from the alignment of the GC with the axis of rotation. The precision to which this value is known, as shown in Section 4.1.3, is ±0.2mm. If we assume that the test subject can be aligned to within 0.1mm, then there is a maximum of a 0.3mm deviation in each axis. As this can be off in two axes, the expected worst case deviation from the GC is 0.42mm given careful alignment. Based on Equation (61), a 0.42mm deviation in the GC would result in 2.4E-6 kg-m^2 error in the measurement of the moment of inertia. If the LAICE satellite is modelled as a rectangular prism

of constant density, its moment of inertia can be described by Equations (62)-(64) where m is the mass, and w, d, and h are the dimensions of the satellite. The smallest moment of inertia is about the z-axis at approximately 0.079 kg-m^2. Therefore, the error in the moment of inertia measurement due to misalignment is 0.003%.

$$I_{xx} = \frac{m}{12}(w^2 + h^2) \tag{62}$$

$$I_{yy} = \frac{m}{12}(h^2 + d^2) \tag{63}$$

$$I_{zz} = \frac{m}{12}(w^2 + d^2) \tag{64}$$

The second source of error is from incorrect measurement of the angle. This parameter is design dependent. A trade study was conducted to determine the best material to use for the torsion rod. A rod of high stiffness was desired for a fast response to mitigate the nonlinear effects of bearing damping. Simultaneously, the rod had to be sufficiently elastic to avoid plastic deformation under rotation. In general, the more slender the rod, the less its moment of inertia (which is variable with length under torsion) would contribute to the rotational inertia of the system. Therefore, a high Young's modulus per unit density was desired. While a comprehensive material selection analysis could have been conducted, limitations in the supply of cheap, high-dimensional tolerance rods resulted in it being more time efficient to evaluate the material options available.

Ultimately, a 416 stainless steel hardened tight tolerance rod of 3/16" diameter was selected. The length of the bar that undergoes torsion is 0.50m as measured between clamping surfaces. For a LAICE sized satellite, the undamped frequency response when spinning about its most significant moment of inertia would be 7.9 rad/s, meaning a single oscillation period would take

approximately 0.80 seconds. Nominally, this measure is independent of the initial angle of the oscillation. If the initial angle is 10° from the equilibrium point, the platform will rotate through just less than 20° on its first half period. When measured optically using the system described in Section 4.2.1, this will result in a 40° change measured by the sensor. Using a 1280 x 720 pixel CCD to record the sensor, it would be possible to record 0.031°/pixel; however, the electronic shutter means the light would be smeared (on average the smear would 0.83° or 26 pixels for a 138fps frame rate). As a result, the reasonable angular resolution would be just under 1°.

$$\theta = \theta_0 \cos(\omega_0 t - \phi) \tag{65}$$

In order to determine the error in MOI based on an inaccurate angular measurement, consider an undamped oscillator (Equation (65)). For a system described above, the torsion constant k is approximately 20.2Nm/rad. Assuming a measurement is made at the first quarter period with ±1° accuracy, then the error in the measurement of the inertia would be approximately 10.2% for a LAICE sized satellite; however, the fit is made across multiple data points not just a single point. The improvement in error goes as 1/n with increased measurements. Therefore, over ½ period of oscillation, the measurement error improves to 0.18% for a LAICE sized satellite (or an absolute error of 3.8E-4 kg-m^2). This renders negligible the misalignment error previously discussed. The total measurement error is significantly worse than commercially available solutions, though it comes at a greatly reduced cost. This estimation is also better than CAD estimates of the satellite MOI. As such, this error will be used when testing the robustness of the control system.

CHAPTER 5
CALCULATING DRAG INDUCED TORQUE

5.1 Modelling External Disturbance Forces

In order to accurately model the satellite dynamics for flight, it is necessary to understand the influence of external disturbance torques. Understanding the relative scale of the disturbances, as well as the scale with respect to the maximum control effort of the magnetic torqueing coils, will inform how to best design the control system for flight as well as inform how best to model the simulations. External disturbance torques will come from three possible sources: gravity-gradient, solar radiation pressure, and aerodynamic forces.

Aerodynamic torques will be the most considerable forces, especially during the peak science period at lower altitudes at which point they can even dominated the applied moment from magnetic torqueing[62]. This will be investigated in more detail in the following section. For the purpose of comparison with the other forces, the average aerodynamic torque (at 220km altitude) is on the order to 10^{-5} Nm. This is comparable in scale to the maximum magnetic torque generated by the coil.

The gravity gradient torques are defined by Equation (66)[16]. In that equation, $\hat{r}_{cm}$ is the distance from the center of the Earth to the GC of the spacecraft. $\boldsymbol{J}$ is the moment of inertia of the spacecraft and μ is the Earth's gravitational constant. For a circular orbit, both $\hat{r}_{cm}$ and ω_0 will be constant; for the purpose of this analysis, this assumption will be made. Given the small size of the spacecraft, gravity-gradient torques will be quite small. Based on the design of LAICE and expected orbit, the gravity gradient torques will be on the order of 10^{-8} Nm. This is negligible compared to the aerodynamic forces.

$$M_{GG} = 3\omega_0^2 \hat{r}_{cm} \times (J \cdot \hat{r}_{cm}) \tag{66}$$

$$\omega_0 = \frac{\mu}{r_{cm}^3} \tag{67}$$

Torques from solar radiation will be similarly small due to the small size of the spacecraft. The SRP was determined to be on the order of 10^{-9} Nm using a commercial software package (STK). This is also negligible when compared to the aerodynamic forces.

5.2 Aerodynamic Forces
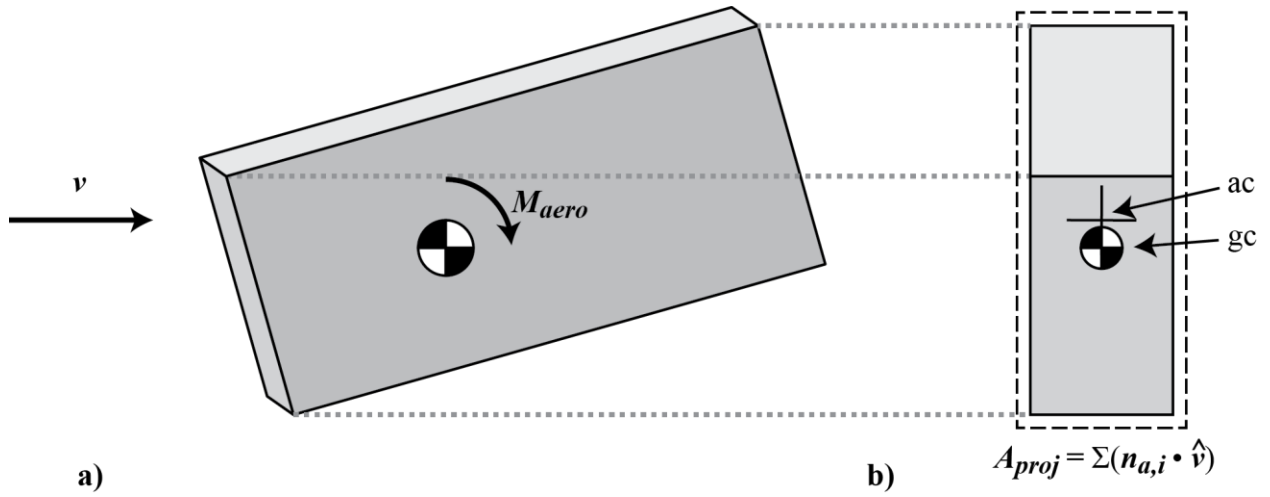
Figure 37: Schematic representation of aerodynamic torques resulting from offsets in the Center of Gravity (gc) and the Aerodynamic Center (ac). Subfigure a) shows the spacecraft encountering a left-to-right flow while tipped while b) shows the projection of the satellite normal to the velocity to demonstrate the offset in Center of Gravity and Aerodynamic Center.

The aerodynamic torque on a face of the satellite can be approximated by

$$\vec{M}_{aero} = -\int \frac{1}{2} c_d \rho (\hat{n} \cdot \vec{v})(\vec{r} \times \vec{v})dA \mid (\hat{n} \cdot \vec{v}) > 0 \tag{68}$$

where c_d is the coefficient of drag, ρ is the atmospheric density, $\hat{n}$ is the surface normal, $\vec{v}$ is the relative velocity of the satellite to the atmosphere, and $\vec{r}$ is the moment are of the face from the center of mass of the spacecraft[14]. The integral is only taken over the part of the face that is facing into the velocity direction as per the $(\hat{n} \cdot \vec{v}) > 0$ requirement.
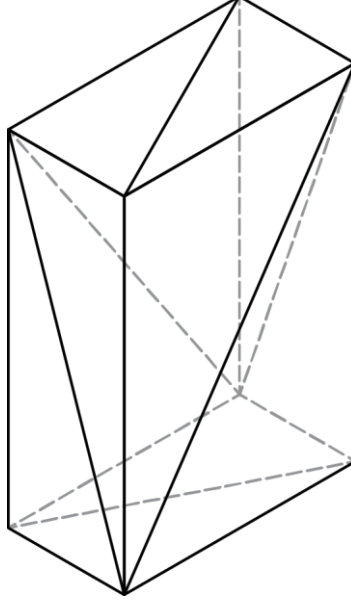
Figure 38: Modelling a simplified box version of a satellite as a set of triangular mesh faces.

To facilitate simulating this torque in software, a surface model of the satellite is used. In a computer aided design software package, a simplified stereolithographic model of the satellite is produced. The model consists of a series of triangular facets defined by the three vertices (see Figure 38). The ordering of the vertices in the file dictates the handedness of the surface normal representing of each facet. From the vertices of each facet, the facet centroid vector, $r_{c,i}$, is calculated as measured from the center of mass. Additionally, the normal vector of each facet is scaled such that its magnitude of the normal vector equals the facet area, thus creating the area normal, $\vec{n}_{a,i}$, for each facet. Using these vectors, Equation (68) can be rewritten as

$$\vec{M}_{aero} = -\frac{1}{2} c_d \rho \sum_{i=1} (\vec{n}_{a,i} \cdot \vec{v})(\vec{r}_{c,i} \times \vec{v}) \tag{69}$$

for all facets such that $(\vec{n}_{a,i} \cdot \vec{v}) > 0$ [63].

While the LAICE geometry is relatively simple, a protective cap for one of the ram oriented sensors does shadow other faces once deployed. As a result, in many orientation, there will be facets which occlude other facets. Assuming hypervelocity conditions, these shadowed regions

would not create drag nor drag-induced torques. Equation (69) does not take these occlusions into account, thus result in an overestimate of torques. While these occlusions are often small, they are significant enough to alter torque estimates and the simulated dynamics of the satellite.

It is possible to analytically compute the effect of occlusions by subdividing shadowed faces into sub-facets. As each facet could be shaded by one or more facets, it is necessary to check against every other face for shadowing. In practice, only facets with surface normals into the velocity direction and further forward along the velocity direction need to be checked for shadowing. Further, any fully shaded facet can be ignored when computing shading of subsequent facets.

Even with these refinements, the computational burden of this process is significant. Shadowing does not uniformly cover an occluded face; it is possible for a face to be unshaded, partially shaded, or completely shaded. The unshaded and completely shaded faces are readily dealt with by including their drag component in full or not-at-all respectively; however, the ways in which a single face can be partially occluded is difficult to implement. For a single triangular facet occluding another triangular facet, it is possible to have 16 shadowing configurations (see Figure 39). If a facet occludes another facet, it is possible to subdivide the occluded face into between 2 and 7 sub-faces. Between 1 and 4 of these faces could then be ignored as they would be fully shadowed. This process would need to be repeated for each subsequent shadowing face. As a result, a single partially shadowed triangular facet from the original model may undergo several subdivisions, increasing the faceting by a factor of 7 at each step. In turn, the model geometry which results from multiple subdivisions may become very complex for even relatively simple shadowing cases. This process entire process must be repeated at every time step as the orientation of the satellite changes, thus making this extremely computationally expensive.
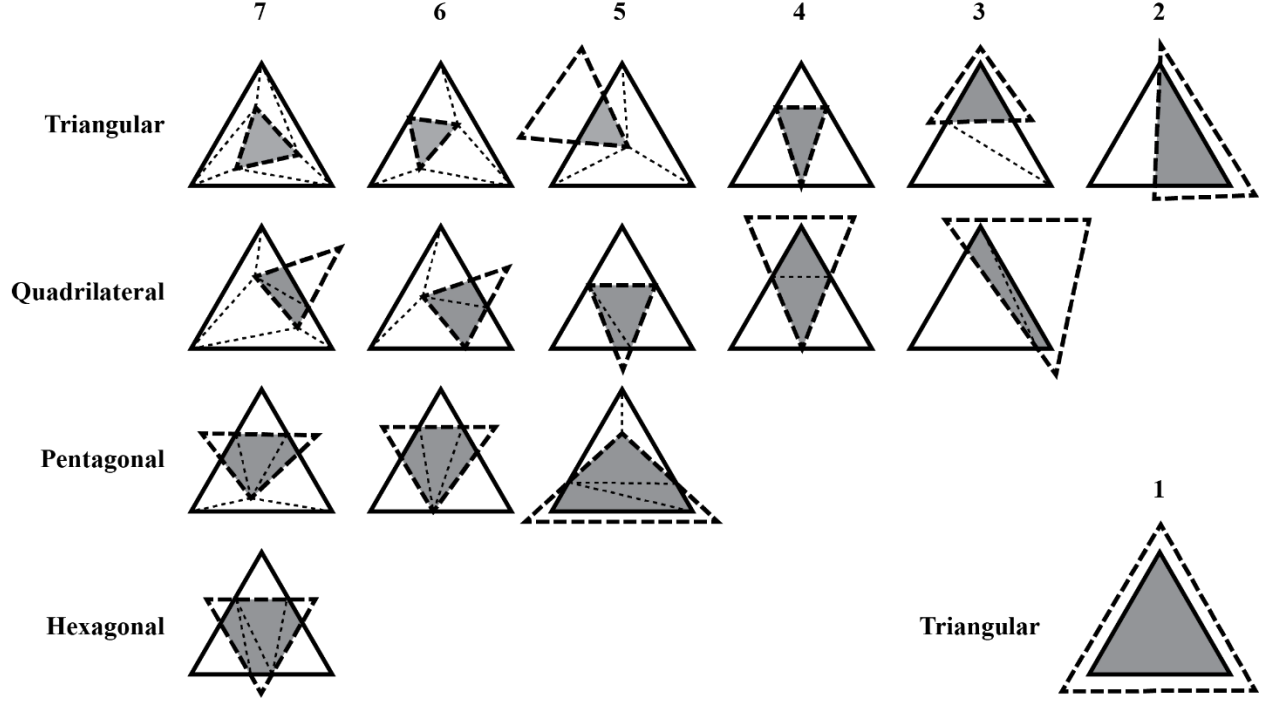
Figure 39: Set of 16 possible shadowing configurations of two triangular patches. The configurations are grouped by required triangular patches to refactor the occlusion (number along the top) and the geometry of the occlusion (identified by row at left). The degenerate case of a complete occlusion shown at bottom right.

5.3 Graphical Method

To reduce the computational burden, a graphical method may be employed whereby the satellite geometry is projected onto a plane perpendicular to the local velocity vector[64]. The area centroid is calculated from the resulting map on a pixel by pixel basis. By rendering the satellite facets, it is possible to determine the projected area centroid by inspection. As overlapping facets are rendered over the same pixels, the problem of occluding facets is handled natively. As the aerodynamic moment simulation is purely ground-based, it is possible to leverage graphics processing units (GPUs) to perform the computation quickly in simulation.

Rather than calculating the contribution from each face (or sub-face component), it is possible to estimate the contribution for the entire satellite. The graphical methods renders the satellite geometry to a frame buffer. The frame buffer is essentially an image of the satellite rendered in pixels. Each pixel represents a discretized portion of the satellite's projected area.

The simplest rendering method draws all faces to the frame buffer in a single color. Each facet of the satellite's geometry is orthographically rotated into the view plane. Every pixel inside the bounds of the rotated facet is then assigned the model color. This process is repeated for all model facets. In this framework, portions of facets which are shadowed do not add to the measurement of the satellite's projected area as the pixel value remains the model color regardless of how many times the pixel has been rendered. While this method also eliminates the need to determine whether a facet's surface normal is pointed away from the velocity vector, it is computationally cheaper to perform the check and not render those facets to the frame buffer. The computational expense for this method is time independent of the amount of shadowing and the memory footprint does not expand with shadowing.

Once the frame buffer has been rendered on the GPU, it is transferred back to the CPU. The CPU reads the value of each pixel. If the pixel value is 0.0f (black), the pixel is ignored; if the pixel value is 1.0f (white) the area is incremented. Additionally, a running totals of the x- and y-component values are calculated to determine the area centroid.

Since this method calculates area and area centroid from pixel values, the accuracy of the process is dependent on the frame buffer resolution and scaling of the geometry inside the frame buffer. In order to increase the accuracy for a given frame buffer size, the rendered satellite geometry should utilize the majority of the frame buffer. To avoid underestimating the geometry due to clipping at the edges of the buffer, the scaling of the rendered satellite much be such that the extremums of the satellite model remain inside the frame buffer for all orientations. This is accomplished simply by determining the radius of the most extreme vertex of the model. This can be accomplished when loading the model facet geometry. The scaling is then tuned such that the radius of the extremum is bound within the frame buffer. Applying this method, the satellite

will occupy the maximum proportion of the frame buffer without risking clipping. For simplicity, the scaling is kept consistent regardless of the satellite orientation. This also guarantees the accuracy remains fixed to the frame buffer resolution.

The above method is the computational simplest approach; however, its accuracy may be improved with certain refinements. At the edges of the projection, the accuracy of the torque addition is dependent on the resolution at which the model is rendered. Pixels which are ≥50% filled by the model are rendered as white, while pixels <50% are rendered as black. Therefore, the torque contribution of a pixel may be over- or underestimated by 50%. To mitigate this effect, the accuracy can be improved by increasing the size of the frame buffer. In this manner, features of the geometry are resolved across a larger pixel area. A pixel which was formerly 50% filled and overestimated by 50% could be divided into four pixels whereby the error is reduced to 25%. This improvement, however, comes at an n^2 increase in computational complexity and memory footprint, where n is the factor of resolution increase. For most of the satellite geometry, this increase in resolution provides no additional information. A fully filled pixel would simply result in four fully filled pixels after a 2x increase in the frame buffer size and matching scaling.

A multi-resolution approach could be used to mitigate the complexity penalty for higher resolution calculations. In this method, the satellite is first rendered at low resolution with only completely filled pixels being rendered white. The resolution is then doubled, copying all white pixels from the previous buffer to the matching quads of pixels on the new buffer. The rendering process is then repeated but only for faces whose projections lie in-full or in-part outside of the white areas from the previous map. This process is repeated until the desired resolution is achieved. This method still has an n^2 increase in memory footprint, but has an $n\log n$

computational complexity which represents a significant computational saving for high resolution calculation[65].

While this method is efficient, a third method is proposed herein whereby the resolution is kept smaller by encoding the pixel fill amount in the intensity value of the pixel. Rather than simply fill a pixel with white or black, each pixel is filled with a grayscale value (0-255) which represents the percentage of the pixel which is occluded (0 for 0%, 255 for 100%).

This approach, relative to increasing the resolution, does falter in that percentage of pixel fill does not solely dictate the torque contribution of the facets filling the pixel. The spatial distribution of the occluding facets within the pixel matter as well. A higher resolution rendering would resolve this distribution while the grayscale value does not. While this will not perfectly replicate the accuracy of a higher resolution model, it does reduce the error for lower resolution models – potentially offsetting the need for resolution improvements.
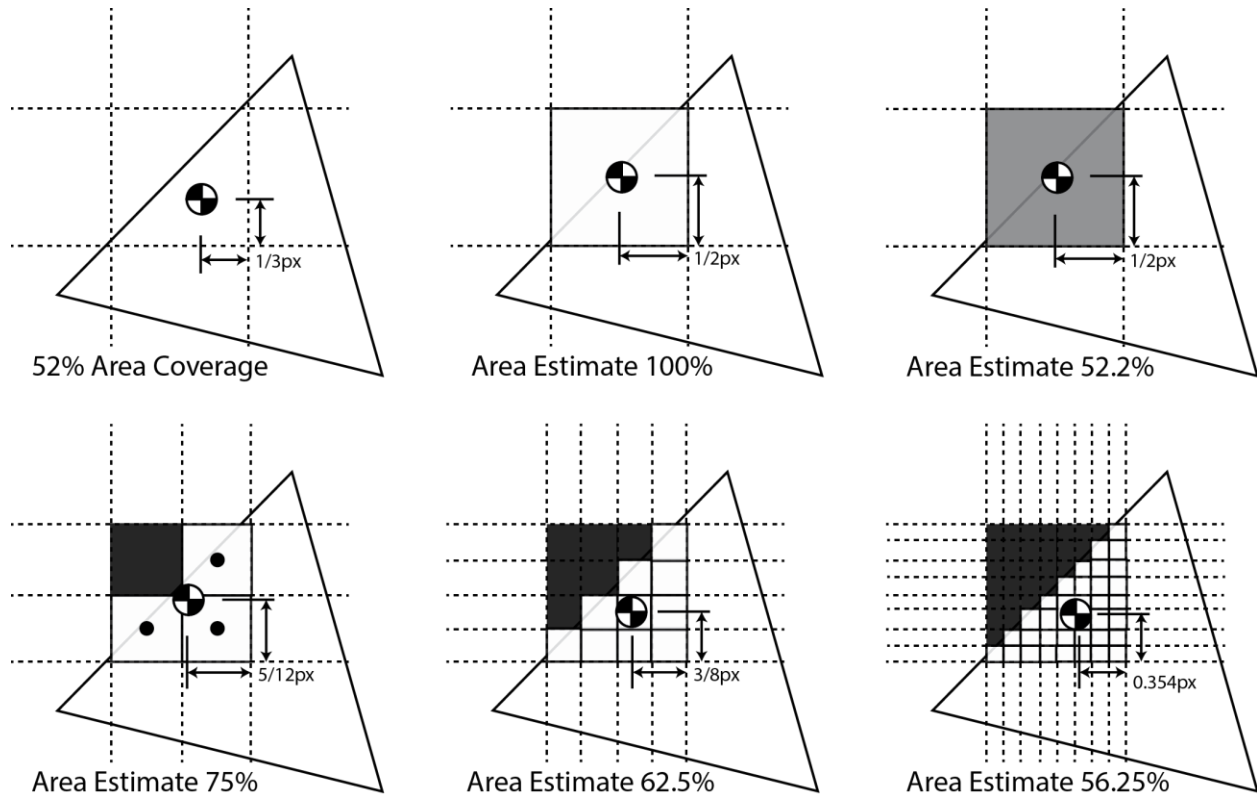
Figure 40: Accuracy of using grayscale to represent the percentage pixel fill versus increased resolution when determining torque contribution of a single pixel. The original coverage and centroid (top-left), the threshold filling method (top-center), the grayscale shading method(top-right), threshold filling at 2x resolution (bottom-left), threshold filling at 4x resolution (bottom-center), and threshold filling at 8x resolution (bottom-right).

A grayscale value can encode the amount of pixel fill to within ±0.2%. Assume conservatively that the error from misrepresenting the location of centroid is 20% of a pixel (the case in Figure 40 is 16.7%). If the pixel in question in Figure 40 is 50px from the projected center of mass (a small distance for a modestly sized frame buffer), the error in the torque contribution of the pixel is less than 1%. By contrast, the error for the threshold filling method is 93% at the original resolution, 45% at the 2x resolution, 20% at the 4x resolution, and 8.2% at the 8x resolution. In fact, a 32x increase in the resolution would be needed in order to match the accuracy of the grayscale value method in this instance.

A further refinement to the method could be in computing the surface normal. At the low pressures of the orbit, the mean-free path of the atmospheric particles increases significantly

94

relative to Earth. This is significant as the continuum mechanic assumptions for (68) depend on the characteristic length scales to be much larger than the mean-free path. A useful measure is the Knudsen number, which is the ratio of mean-free path to the characteristic length scale. For continuum systems, this value should be much less than 1[66]. For the LAICE satellite immediately following deployment, if the major dimension is used, the Knudsen number is only slightly less than one. When examining the protruding cap length, the Knudsen number is over 14. In the case of the cap, the continuum mechanic assumptions break down making the assumptions suspect[66].

In order to remedy this problem, a stochastic model of the drag can be used. In this instance, the interaction of individual particles, or groups of particles, are considered. In order to compute this, the particle-to-surface interaction must be examined, not merely the projected area. From a computational standpoint, an extension of the grayscale value is used to encode the surface normal into the frame buffer.

When the GPU encodes pixel values, it does so as an RGBA value (red, green, blue, and transparency). Formerly, the grayscale method would encode the same grayscale value across all four channels. In this method, when a facet is rendered to the buffer, the surface normal is encoded in the RGB color channel. In this manner, the angle of the surface normal can be stored to $\pm 0.35°$ for an 8-bit color value. The pixel fill factor is then encoded in the transparency channel of the pixel. If multiple facets of different surface normal write to the same pixel, the colors are averaged and the transparencies added (to a maximum of the fullscale value). This introduces some error as it has an effect of blunting edges. A rendering of the surface normal encoded frame buffer is shown in Figure 41
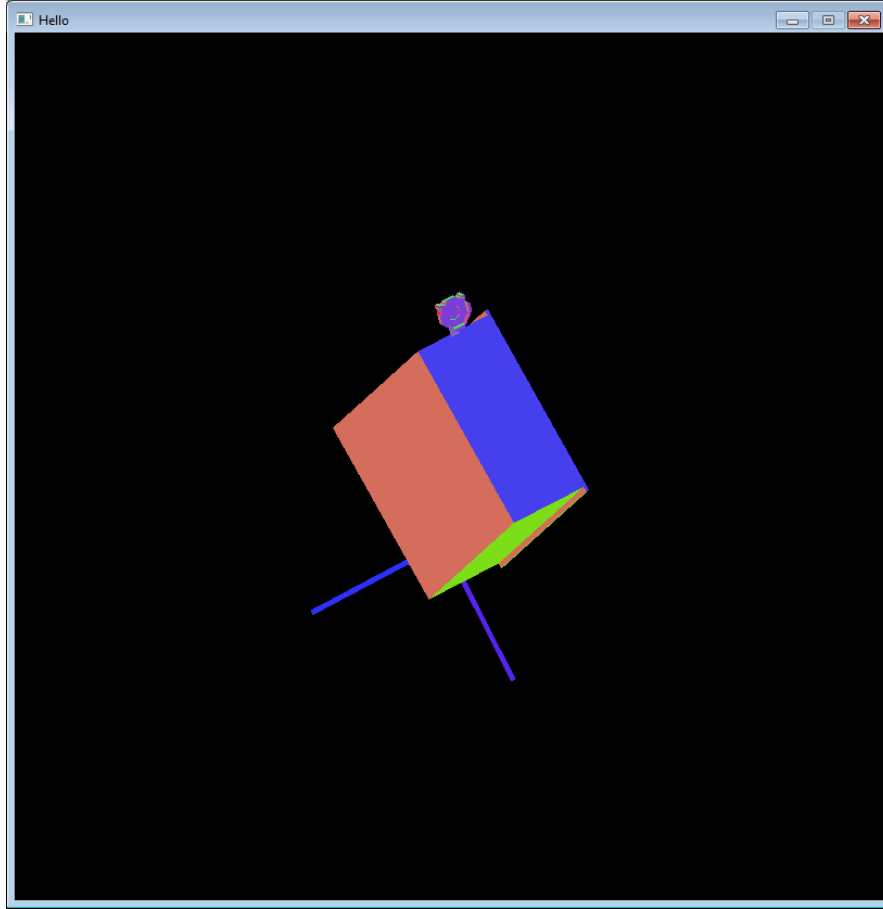
Figure 41: Screen-capture rendering of GPU frame-buffer for computing aerodynamic drag. Pixel color encodes surface normal direction. Faces pointed forwards are blue, pointed to the left and right are red, and up and down are green (all from the perspective of the viewer who looks along the anti-velocity vector with the orbit normal oriented up).

For the purposes of this simulation, the continuum aerodynamic model is used. The three-color, surface normal encoded graphical model will be used in future works after additional validation. The frame buffer resolution was selected as 800px by 800px, with 8-bit per channel (RGBA).

5.4 Calculating Velocity

In order to model the drag as well as determine the orientation at which to render the satellite for determining projected area, the relative velocity of the satellite with respect to the atmosphere must be calculated. The relative velocity is the vector difference of the satellite velocity and the atmospheric velocity.

$$\vec{v}_{rel} = \vec{v}_{orb} - \vec{v}_{atm} \tag{70}$$

The atmospheric velocity is modelled simply by a co-rotating atmosphere, whereby the velocity as measured in the Earth-Centered Earth-Fixed reference frame can be represented as

$$\vec{v}_{atm,ECEF} = \frac{2\pi}{T_{Earth}} \|\vec{r}_{orb}\| \cos(\phi) \begin{bmatrix} \cos(\frac{\pi}{2} + \lambda) \\ \sin(\frac{\pi}{2} + \lambda) \\ 0 \end{bmatrix} \tag{71}$$

where T_{Earth} is one sidereal day, $\vec{r}_{orb}$ is the vector to the satellite from the ECEF origin, ϕ is the spacecraft latitude, and λ is the spacecraft longitude[67]. For the simulation, velocity is rotated into the Earth Centered Inertial (ECI) frame.

The orbital velocity and position is modelled based on an orbit propagator. For the simulations, the SGP4 propagator is used to calculate the spacecraft velocity. The propagated velocity does not take into account the effect of the atmospheric drag on orbit decay; over the simulation time, the impact of drag on orbital velocity is considered negligible with respect to its impact on assessing the performance of the attitude control system. The orbital position is used to estimate atmospheric density based on the Jacchia model[68].

5.5   Simulation of Satellite Dynamics under Aerodynamic Forces

The principal reason for simulating the drag is to determine the performance of the determination and control system under external torques; specifically, the ability to control the satellite in the presence of these moments. The satellite does not model these forces, it merely considers them disturbances which must be rejected.

The simulation presented does not apply any magnetic control torques. By not applying control torques, it is possible to demonstrate the uncontrolled motion of the spacecraft under aerodynamic forces.
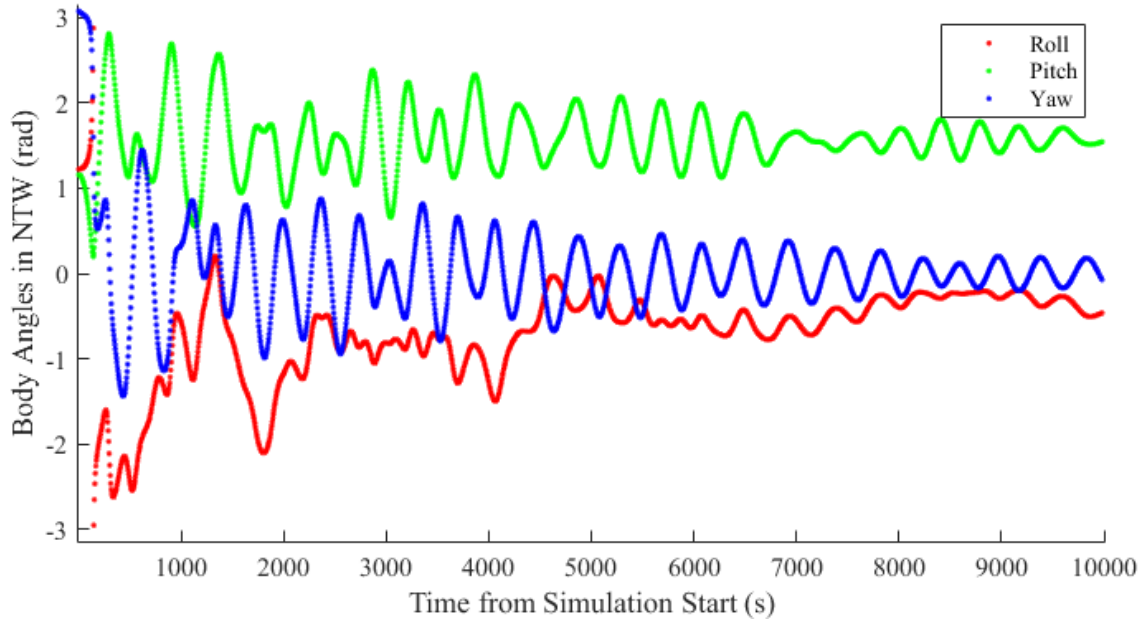
Figure 42. Body Euler angles with respect to the satellite normal (NTW) frame under aerodynamic forces without control applied.

As can be seen from Figure 42, the body angles start in the anti-velocity direction (the yaw angle shown in blue at π); however, due to the unstable initial configuration the axes quickly diverge from their starting position. After a period of oscillation, the satellite does begin to stabilize with the z-axis in the velocity direction (represented by the oscillation of the yaw about 0). The oscillation does not damp out over the time period of the simulation, nor does the z-axis stay within the 5° requirement for pointing. Most significantly, the orientations of the x and y axes oscillate about an arbitrary set of angles. For proper science alignment, the x and y axes should be centered about $\frac{\pi}{2}$. To more clearly demonstrate the final configuration of the satellite in the uncontrolled case, it is beneficial to view the trajectory of the satellite body vectors in the NTW frame.

98
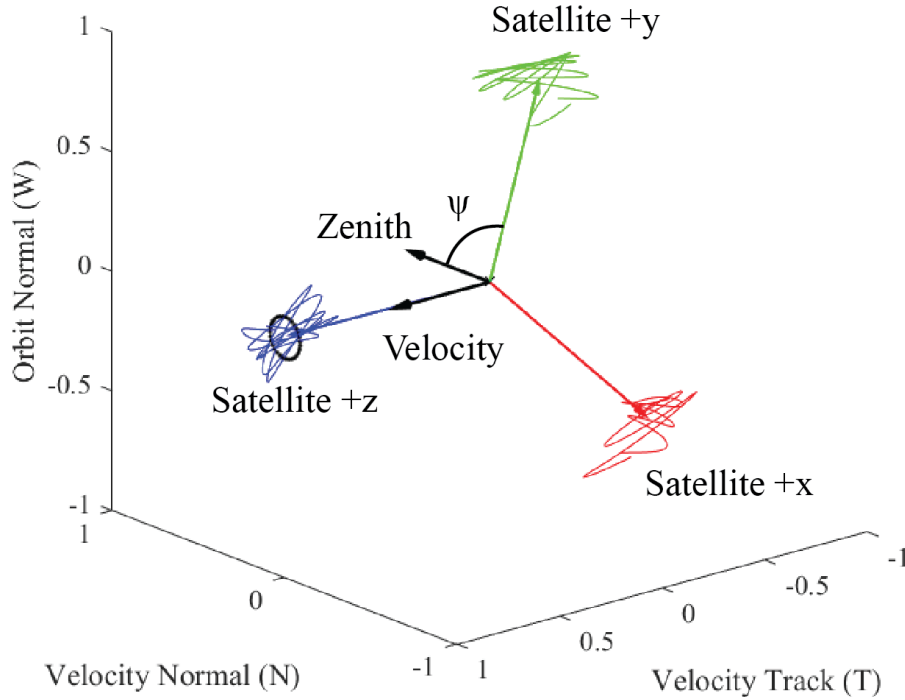
Figure 43. The trajectory of the spacecraft body unit axes shown in the velocity normal (NTW) frame with aerodynamic stabilization and no magnetic control plotted over the final 500 simulation steps (approximately 2500 sec). The velocity direction is shown as well as the zenith direction. The angle ψ represents the offset between the satellite's +y axis and the zenith direction. The time-averaged axes of the satellite are represented by the colored vector triad.

Figure 43 shows the satellite +z axis in track with the velocity vector, as expected due to the satellite's positive static margin in that direction. The +z axis oscillations exceed the 5° target, and clearly the +y axis is misaligned. As such, the pointing requirement of the three ram facing payloads are nearly satisfied by the aerodynamic stability of the satellite's design; however, the photometer payload (on the −y face) is not pointing along the nadir vector (the error is represented by the angle ψ).

Figure 43 also demonstrates oscillations back-and-forth about the satellite's z-axis represented in the traces of the +x and +y axes (as well as oscillations out of the velocity normal plane due to the deviations of the +z axis from the velocity direction as described above). Along the orbit, the satellite x and y axes rotate about the velocity direction due to variations in the local atmospheric velocity along the inclined orbit.

This result gives confidence in the satellite's robustness to aerodynamic forces. The system achieves a passive stability in one of the pointing axes by virtue of the aerodynamic design of the satellite. This stability can be exploited when optimizing for the coupled system. Control of the satellite's attitude under aerodynamic forces will be investigated further in the following Section.

5.6   Simulation of Attitude Control under Aerodynamic Forces

The simulation was rerun with the same initial conditions as for the control case described in Section 5.5, but for this simulation QRF control was applied. The control parameters for the simulation are defined in Table 4. The simulation was done using perfect attitude determination as a proof of concept. The commanded magnetic moment, which outputs in discrete levels, was used to calculate the generated torque (in the same manner as in Section 3.4). This torque was added to the aerodynamic moment generated using the graphical method with grayscale pixel values. The net torque was used to propagate the satellite dynamics according to Equations (4)-(7) (see Section 2.2).

Table 4. Control Parameters.

Simulation Parameter	Value
Quaternion Error Gain, κ	0.0004
Angular Rate Gain, γ	0.005

As demonstrated by Figure 44, the control was able to significantly reduce the final pointing error. The satellite more quickly orients the +z-axis into the velocity direction. Additionally, the oscillations in the yaw axis have been minimized. Finally, the +y axis is now aligned with the zenith vector ensuring the photometer payload (on the –y face of the spacecraft) is pointing towards nadir.

100
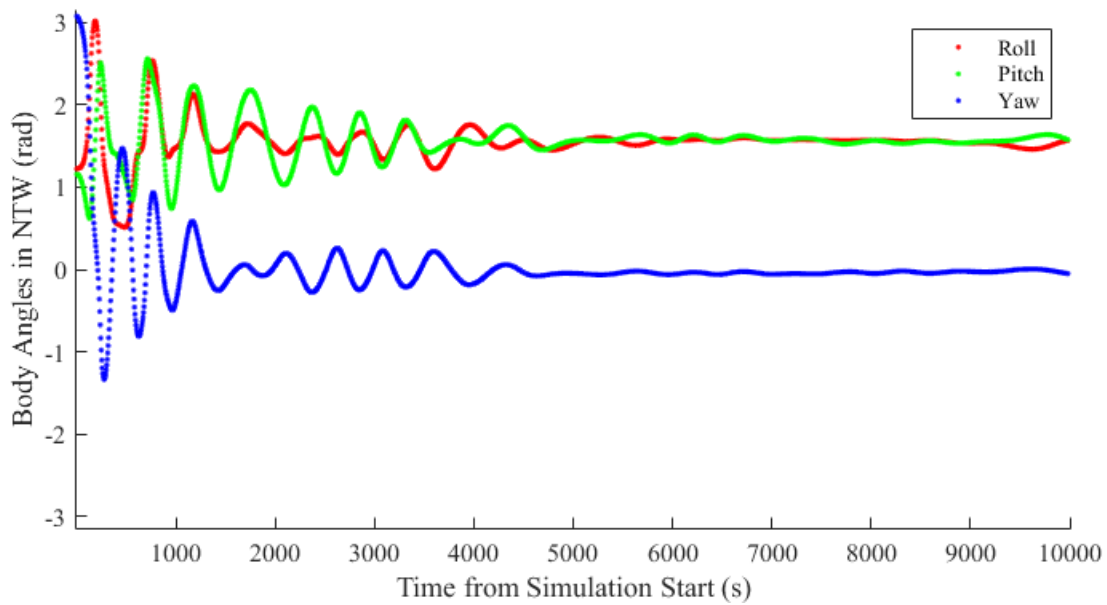
Figure 44. Body Euler angles with respect to the satellite normal (NTW) frame with QRF control applied.

By examining the orientation of the body axes in the NTW, as was done for the uncontrolled case, it is possible the see the improvement in the attitude control (see Figure 45). The orientation over the final 500 simulation steps shows the satellite body axes well aligned with the velocity and zenith vectors. The ram pointing vector also stays within the 5° pointing requirement (see Figure 46 and Figure 47).
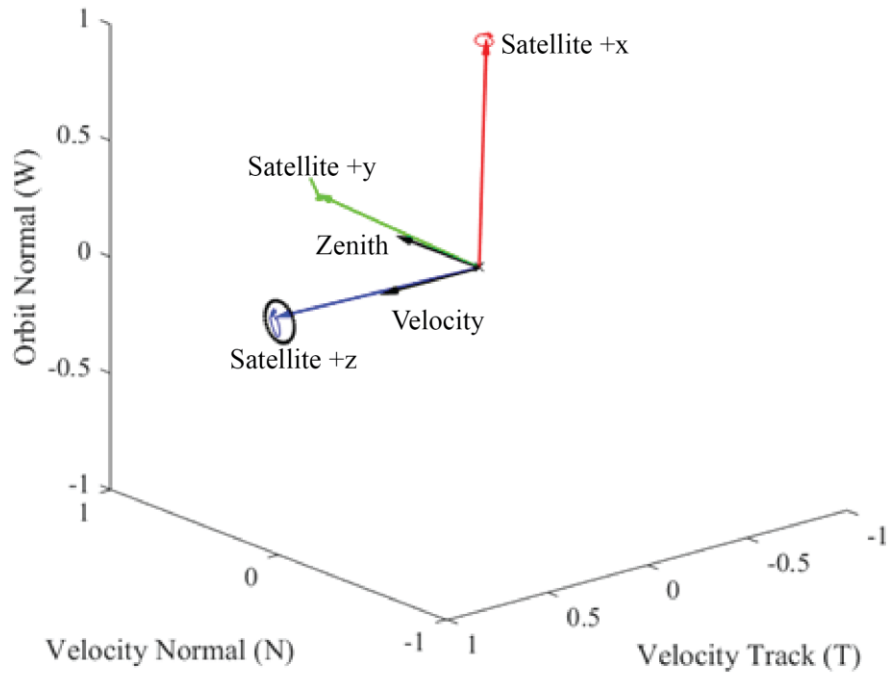
Figure 45: The trajectory of the spacecraft body unit axes shown in the velocity normal (NTW) frame with aerodynamic stabilization and with QRF control plotted over the final 500 simulation steps (approximately 2500 sec)
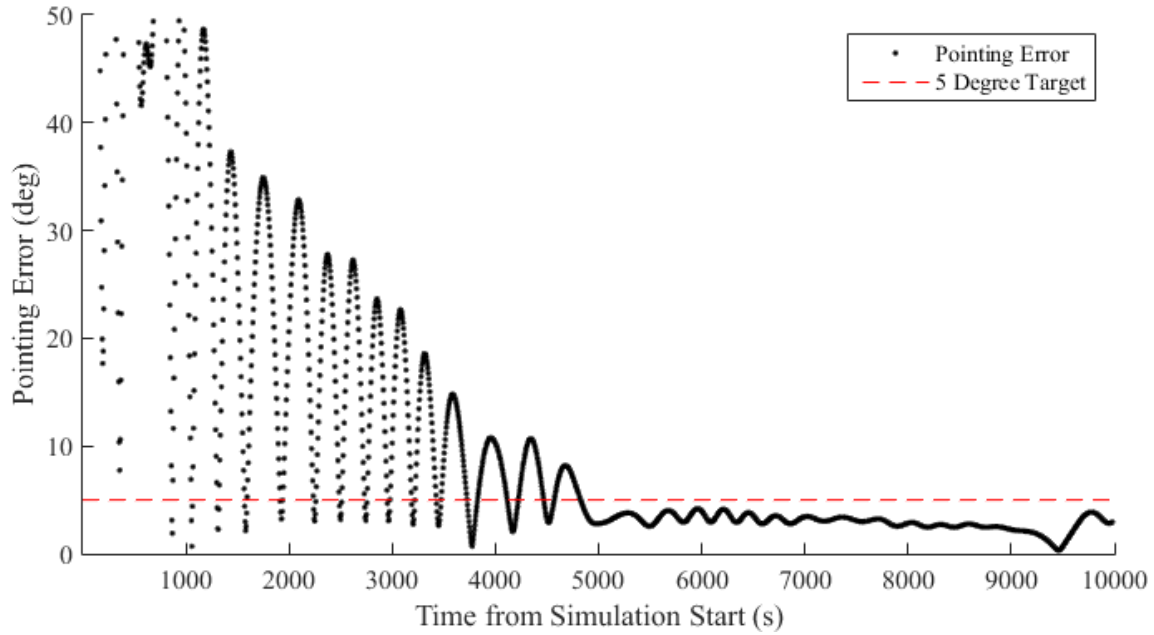
Figure 46: Plot of the pointing error for the case with active attitude control under aerodynamic forces.
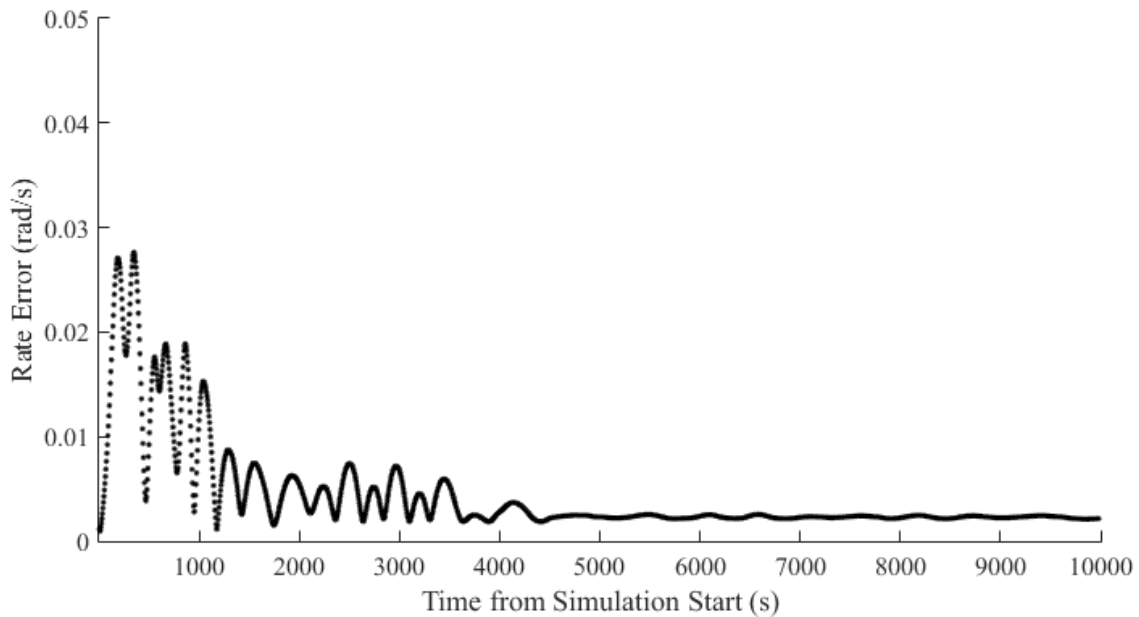
Figure 47: Plot of the norm of the body rate in the ECI frame for the case of active attitude control under aerodynamic forces. The steady state rate is approximately 0.0015 rad/s, which is slightly greater than 2π over the orbital period.

Examining the magnitude of the magnetic control moment and the aerodynamic moments gives insight into the behavior of the controller under aerodynamic forces. At the start of the simulation, the aerodynamic moments exceed the limits of the magnetic torque coils. As the spacecraft aligns its +z axis into the velocity direction, the magnetic torqueing coils are able to damp the oscillations previously observed, significantly decreasing the time to stabilized. Of note, as the system begins to reorient the x and y axes using magnetic torques, the spacecraft begins to turn its non-ram pointing faces into the freestream flow. As a result, the aerodynamic forces restore against the moments. This can be observed in Figure 48. The net effect is that the torques generated off the z-axis, which are a consequence of the ability to only generate torques perpendicular to the local magnetic flow, are cancelled by the aerodynamic forces. As a result, the satellite is able to reorient about the z-axis without pulling the satellite off of its ram orientation. This is a significant advantage.
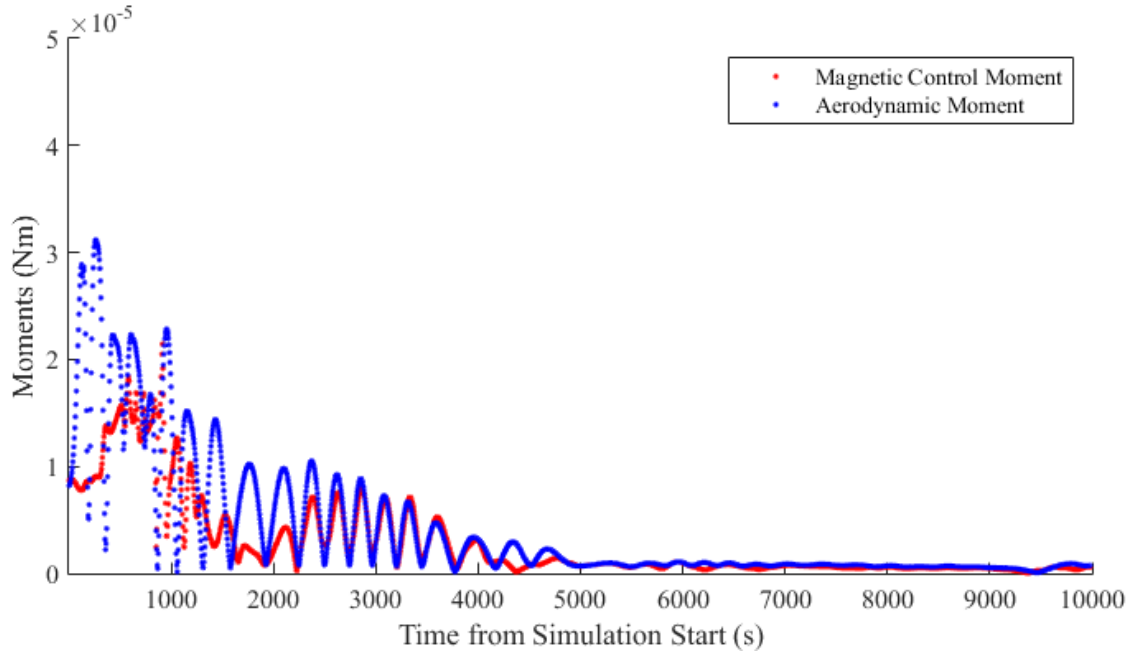
Figure 48: Plot of relative moments from Magnetic Control Moments and Aerodynamic Moments.

The consequence of the large aerodynamic restoring torques is that at low altitudes, intentionally attempting to reorient the satellite out of the most aerodynamically stable orientation is not possible. The control authority of the magnetic torqueing coils, even at peak, are not capable of overcoming the aerodynamic restoring torques.

5.7 Summary of Contributions

A graphical model was presented for calculating aerodynamic moments on NanoSatellites. The innovation of encoding pixel fill factor into the frame buffer alpha channel was demonstrated analytically to have upwards of a 32x resolution improvement per edge pixel over simple filling algorithms. The drag models were applied to magnetic attitude control cases of LAICE sized satellites and it was demonstrated that the spacecraft could be controlled for NTW frame targets by relying on aerodynamic torques to restore the satellite to mission pointing at low altitudes.

CHAPTER 6
HARDWARE-IN-LOOP SIMULATION

6.1 Software/Hardware Testbed for ADCS Testing

It is imperative to test the flight hardware prior to launch. The attitude determination and control for LAICE is a full three-axis system, and replicating this terrestrially presents significant challenges. By testing the system in pieces, it is possible to obtain confidence in the system. To facilitate this testing, a hardware-in-loop (HIL) test bed was developed[69]–[71].

The three-axis determination problem can be represented by a satellite moving through an approximately stationary magnetic field with respect to the ECEF frame, whilst also rotating relative to the ECI frame. From the perspective of the satellite's body-fixed frame, this results in a changing magnetic field resulting from the change in position along the orbit, the rotation of the satellite with respect to the ECI frame, and the rotation of the ECEF frame with respect to the ECEF frame. This could be simulated in two of ways: 1) generate a sequence of static magnetic fields which replicate the Earth-fixed reference frame representation of the magnetic field, then having the satellite rotate within that volume based on the spacecraft attitude, or 2) generate a sequence of dynamic magnetic fields which replicate the Body-Fixed reference frame representation of the magnetic field whilst keeping the satellite fixed.

In the first case, the satellite dynamics do not need to be modelled as long as the satellite can rotate freely in three dimensions without disturbance torques from gravity or dissipation due to friction. While not having to simulate the dynamics increases the fidelity of the simulation and would be a better demonstration of magnetic torqueing controls, it is extremely difficult to implement in practice. Gimbal systems would allow for three rotational degrees of freedom, but would need to be precisely aligned with the spacecraft center of mass to avoid torques due to

gravity. Simultaneously, a gimbal system is prone to gimbal lock as well as loss of energy due via friction in the bearings. For these reasons, a three rotational degree-of-freedom system is onerously difficult to implement.

In the second case, the satellite remains stationary, thus eliminating the significant mechanical challenges; however, the satellite dynamics do need to be modelled. This approach therefore requires more extensive testing of the control actuators to validate the accuracy of the model the dynamics prior to simulating. Since a thorough understanding of the systems response to control actuator activation is required for the tuning of control systems, this validation step does not represent any additional work. As such, the second approach of keeping the satellite stationary represents a much simpler technical approach with a quantifiably low reduction in fidelity.

6.2   Extending CubeSim from Software to Hardware

In Chapter 2, the concept of a joint flight and simulation codebase was introduced under the portmanteau of CubeSim. CubeSim represents a suite of tools which are able to accurately model CubeSat dynamics, magnetic fields, and aerodynamic forces in simulation, whilst also being interchangeably used as flight code for attitude determination and control. While the shared code base is useful for testing concepts and optimizing controllers, it has been further extended to serve as the architecture for HIL simulations. By leveraging the simulations and models already completed for software testing, development of HIL simulations was much more quickly accomplished. As a result, the moniker of CubeSim has been extended to include all hardware and software simulation capabilities developed as part of ADCS testing for the IlliniSat programs. Within this chapter, CubeSim will be used to refer to hardware specifically designed for this system, as well as the HIL testing suite as a system.

6.3 Hardware Setup

The hardware setup for the HIL simulations is based around a computer-controlled, tri-axial Helmholtz cage[69]–[71]. The Helmholtz cage is simultaneously able to cancel the local magnetic field and superimpose any arbitrary magnetic field that a satellite may experience on orbit. This is accomplished using a fast-switching power supply driving the current level of each of the three coil axes. The commanded current is calibrated to recreate directional magnetic fields at the cage center. The magnetic field specified is based on magnetic field models of the satellite.

The CubeSim hardware setup consists of four major components: a tri-axial Helmholtz cage; a fast-switching, computer-controlled power supply; desktop computer running the simulation software; and a lab-grade, high-precision magnetometer. Each component is described in brief.
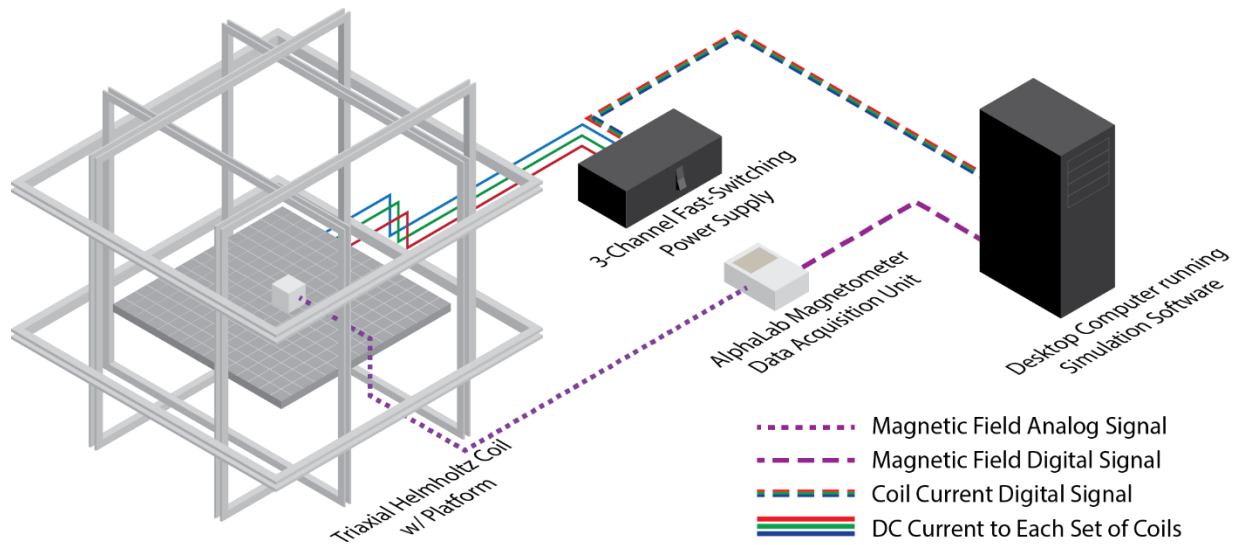
Figure 49: Experimental Setup for Tri-axial Cage Calibration and Simulation Model Validation

6.3.1   Tri-axial Helmholtz Cage

The Helmholtz cage consists of three pairs of wire coils mounted co-orthogonally (see Figure 50 and Figure 51). The copper wire is wound tightly around square frames made of aluminum

c-channel to form coils (see Figure 52). Each pair of coils is connected in series ensuring that the current is consistent between coils. When mounted, the handedness of the coils of wire is kept consistent; this ensures that the directions of the magnetic field generated by each coil are consistent. Each set of coils are mounted orthogonally. The configuration of the coils was designed to create a volume of roughly uniform magnetic field at the center of the cage when the coils were active. This dictated the size, number of windings, and spacing of the coils. A full analysis of the design was completed in Reference [69].
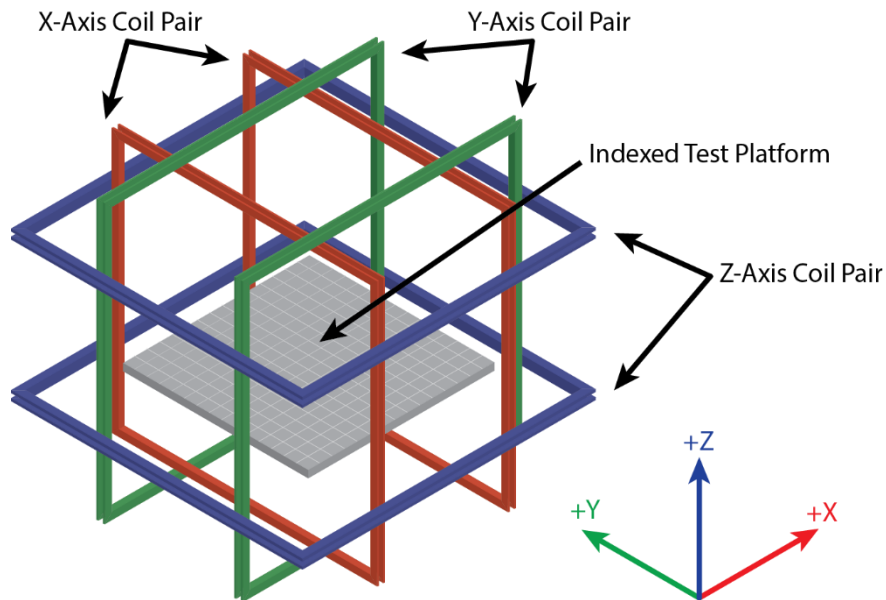
Figure 50: Tri-axial Helmholtz Cage (coil pairs colorized by axis for clarity).

Figure 51: Tri-axial Helmholtz Cage in laboratory.
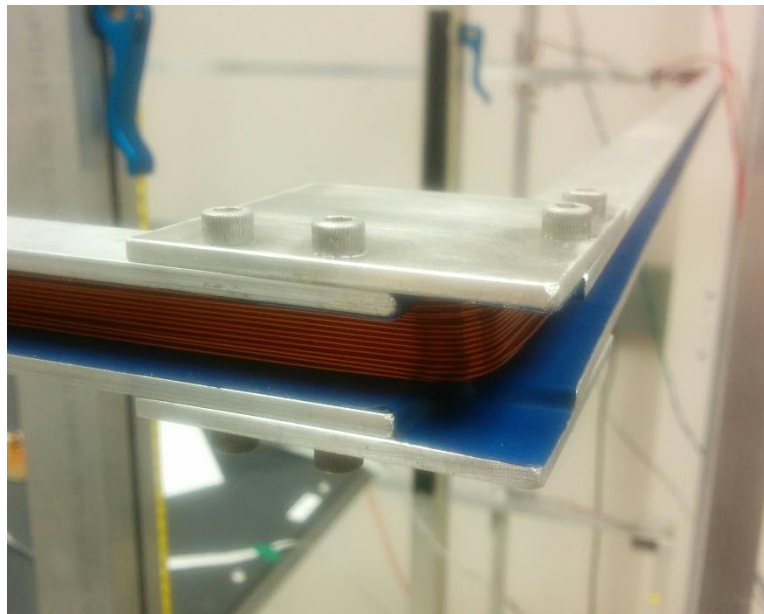
Figure 52: Copper coil winding around aluminum channel frames.

To support the satellite, an indexed platform was designed. The platform is made from ABS, a material with a negligible magnetic dielectric. The ABS platform is patterned with indexing grooves every 5mm in the X and Y axis centered along the cage center. The grooves allow specially designed sleds to be precisely and repeatedly placed on the platform. The platform is supported by an aluminum frame constructed from extruded dimensional rails fastened with aluminum, plastic, and brass fittings. The frame materials were selected for their low dielectric properties. A set of slotted aluminum L-channel (affixed to the X and Y axis coil frames) accommodate sliding supports which allow the platform assembly to be positioned at various heights. Translating the platform assembly within the cage allows sensors of various heights to be supported while ensuring the sensor is still located at the center of the test volume.

6.3.2 Fast-Switching, Computer-Controlled Power Supply

To control the current going through each set of coils, a power supply was designed (see Figure 53). The power supply regulates the current in each of the three sets of coils independently. Each current is commanded by providing a digital set-point between 0 and 16384 (2^{14}) which linearly correspond to a current between 0A and 1A. Separately, the direction in which current is being driven in each coil can be commanded. In effect, each set of coils can be driven with -1A to 1A of current. Communication with the power supply is accomplished via a serial communication protocol over USB.

Figure 53: CubeSim computer controlled, fast-switching power supply.

6.3.3 Simulation Computer

A simulation computer is needed to constantly vary the commanded current set point in order to produce the desired magnetic field in the Helmholtz Cage test volume. The simulation computer handles all serial communications with the power supply, as well as managing all serial communications with lab-grade magnetometer which is used for calibrating the system. The simulation computer runs the suite of CubeSim software which can manage the simulation environment or even emulate subsystems of the satellite if required.

6.3.4 High-Precision Magnetometer

CubeSim employs an AlphaLab Magnetometer for calibration of the chamber. The magnetometer is used primarily to determine the ambient magnetic field inside the chamber prior to beginning a simulation. The pre-simulation magnetic field value is first cancelled out before a new magnetic field is superimposed to simulate the on-orbit magnetic field.

6.4 Hardware Calibration

Prior to using Helmholtz cage for ADCS testing, a thorough calibration was necessary. The calibration correlated the current set-points of the power supply with the magnetic field values

measured inside the chamber. For the purposes of the present applications for CubeSim using the Helmholtz cage testing environment, only the magnetic field at the center of the test volume (position <0, 0, 0>) was of interest. To calibrate the system, the power supply was commanded to step through 16 current set-points over the full range (0A – 1A) in both directions, for each axis. For each set-point, the magnetic field at the center of the cage was measured using the AlphaLab magnetometer; the three-dimensional magnetic field vector was recorded in all three axes. First, an ambient magnetic field measurement was made with the coils de-energized to establish a baseline for the active measurement. The coil was then energized, and an active coil measurement was made. The system then proceeded to the next set-point.

Measurements were only made with a single coil active at a time. A linear regression was used to determine the correlation between current set-point and magnetic field strength in the direction of the coil axis. This linear relation was inverted to determine the current set-point for a desired magnetic field strength value. This process was repeated for each axis. The magnetic field contribution in the two off-axis directions were measured to determine the severity of cross-coupling induced magnetic fields.

6.5 Experimental Setup

The principal function of CubeSim is to perform hardware-in-loop magnetic attitude determination and control tests. In order to perform such a test, the satellite is placed inside the Helmholtz cage such that the flight magnetometer is located at the center of the test volume. If the satellite is not available, a single solar panel may instead placed inside the chamber as it contains all vital components for the test. In such tests, the satellite flight computer is emulated on the simulation computer (again, this is enabled by the fact that the simulation computer and flight computer run identical codebases). An ambient magnetic field reading is made using a lab-

grade magnetometer as close to the test volume center as possible; this ambient magnetic field must be cancelled through superposition of a static field generated by the Helmholtz cage. The CubeSim software then communicates directly with the satellite via a serial connection to provide initial TLE and the clock time. As such, both the satellite and CubeSim simulation begin with identical ephemeris data at initial time; based on potential differences in ephemeris models, they may diverge from that point. The CubeSim simulation then begins to vary the current set-points of the power supply according to the simulated attitude and magnetic model.
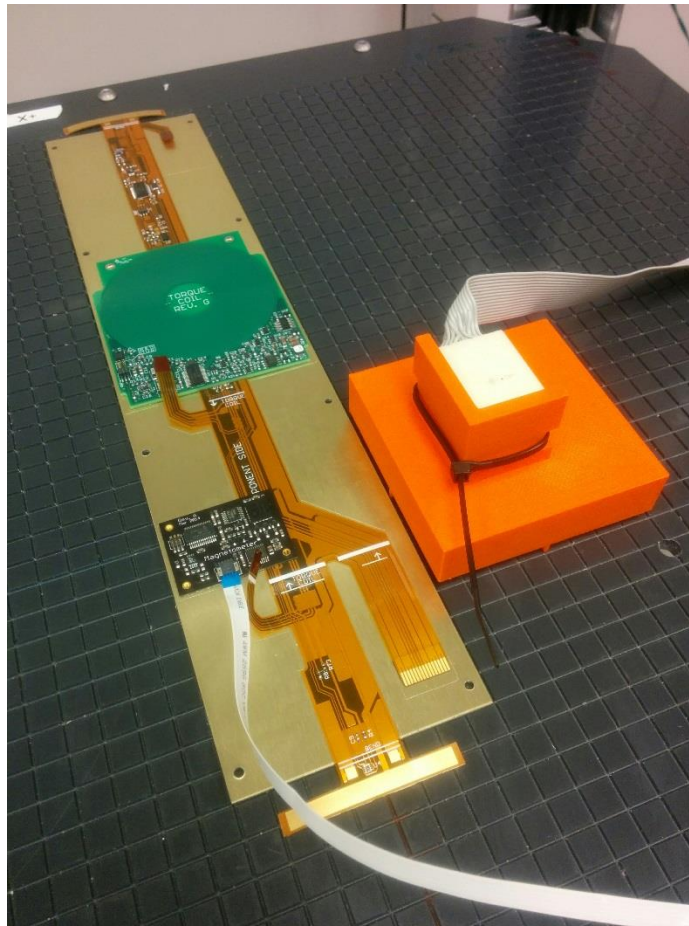
Figure 54: Flight solar panel placed with the magnetometer at the Helmholtz cage center with the AlphaLab high-precision magnetometer placed beside it (on orange registration sled).

The satellite internally runs its attitude determination and control algorithm based on its flight magnetometer inputs alone with no additional inputs from the CubeSim simulation package. As a

result, differences in ephemeris and magnetic field models will be representative of on-orbit variations. The attitude control signals from the satellite (namely, the currents supplied to each coil) are measured by CubeSim using a current sensor. The generated magnetic moment is modelled based on the supplied current and coil geometry. This magnetic moment is used to derive the induced torque on the satellite using the simulation truth magnetic field (which may differ from the satellite expected magnetic field) to propagated the satellite attitude dynamics. The propagated attitude state (along with the propagated ephemeris state) is used to determine the new power supply set-points, thus completing the loop.
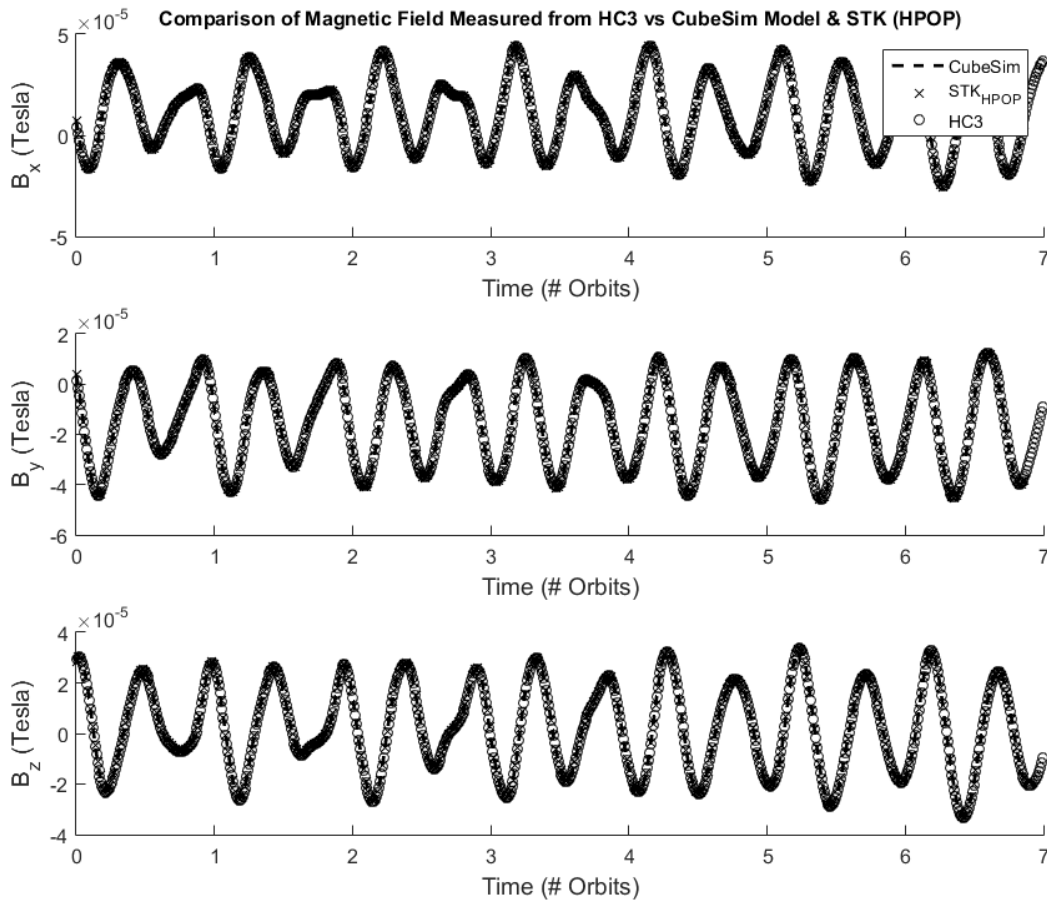
Figure 55: Plot showing the comparison of the measured magnetic field inside the HC3 with the CubeSim SGP4/IGRF Models and STK HPOP/IGRF Models.
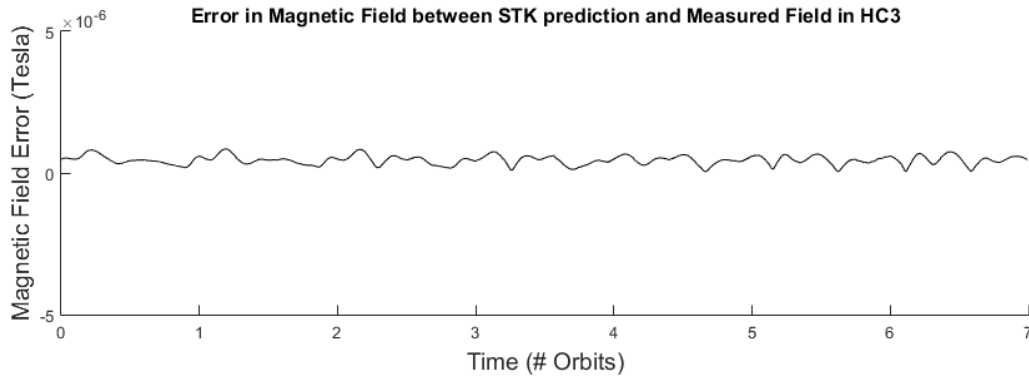
Figure 56: Demonstration of the error between the industry standard STK HPOP/IGRF Model and the field produced by the HC3 given identical TLE data.

The closed-loop duty cycle of the Helmholtz cage coil current control is less than 0.7 seconds. This is significantly faster than the 10 second closed-loop duty cycle of the satellite. The simulations may be run in real-time with no modification to the satellite flight software. Alternatively, the satellite can be commanded to update its attitude estimate by providing it with updates to its clock time. In this manner, it is possible to run the simulation faster than real-time.

6.6 HIL ADCS Simulation

To validate the attitude determination simulation, the HIL testbed was applied to a set of simulations. In all simulations, the 7-State Extended Kalman filter determination algorithm was applied. To facilitate early testing, a single EKF was employed. The initial state estimate was made sufficiently close to the simulated attitude quaternion value to ensure quick convergence. Due to the inability to use the flight hardware as it was not fully produced, only the flight sensors were used for this experiment. The rest of the satellite, including the magnetic torque coils, was emulated using the desktop computer running the simulation software. This significantly simplified the test procedure without affecting the parameters of the test. The hardware was arranged as shown in Figure 57.

115
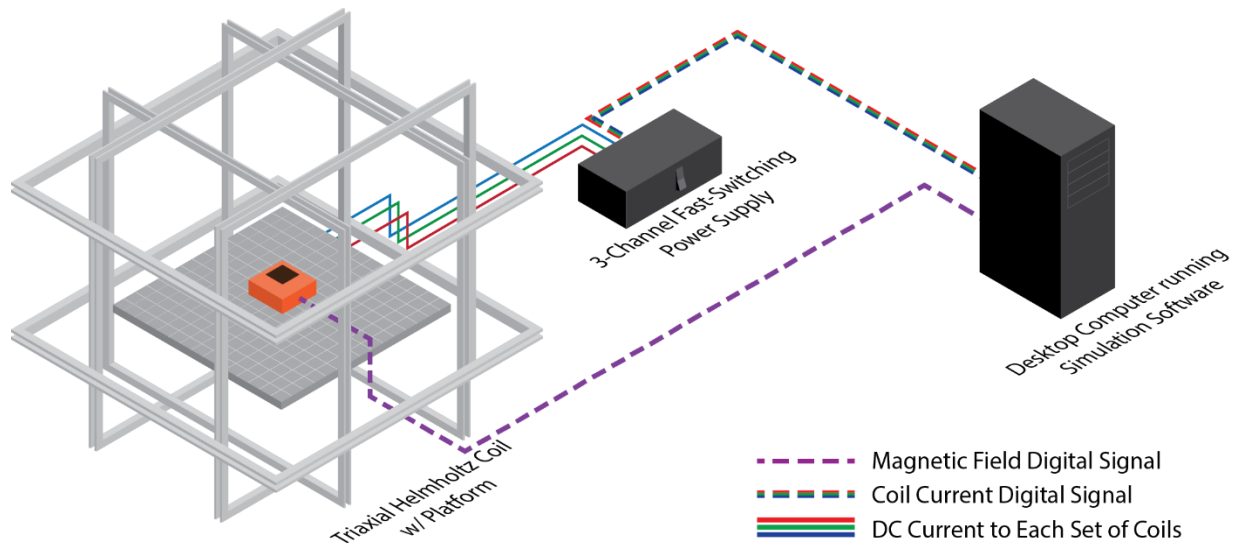
Figure 57: Hardware-in-Loop Attitude Determination Test with Flight Magnetometer

6.6.1 Emulated Satellite

The flight sensor communicated via serial communication to the simulation computer. The simulation computer ran an emulated version of the flight software. The emulated satellite did not have access to the truth values being propagated by the simulation computer, as such attitude determination based on the magnetic field values read from the flight magnetometer in the Helmholtz cage was the only way by which the emulated satellite could determine its orientation. As stated, the 7-State EKF was used with the filter weights from Section 2.4.

The Hybrid B-Cross/QRF controller was used with the weights resulting from the Particle Swarm Optimization discussed in Section 3.5.3. Of note, these weights were optimized for fixed inertial frame targets with no aerodynamic forces. The attitude target for this simulation is constant in the NTW frame, with the satellite's +z axis aligned with the orbit velocity and the –y axis oriented towards Nadir. In this configuration, the satellite should remain aerodynamically stable.

The emulated satellite's MOI estimate was increased by 0.2% relative to the MOI of the "truth" satellite model propagated in the simulation. This is greater than the expected uncertainty in the MOI measurement as discussed in Section 4.2.2.

6.6.2   Simulation Setup

At the beginning of the simulation, the satellite begins with the +z axis oriented at 90° with respect to local velocity vector. The –y axis is oriented towards Zenith. This is an aerodynamically unstable orientation and a 120° rotation about the body {1,1,1} axis from the target attitude. The satellite starts with zero initial angular rate with respect to the ECI (i.e. inertially fixed with respect to the ECI).

The satellite orbit was set to an inclination matching the International Space Station (ISS Zarya) at 200km altitude. This represents the peak science period of the mission and notably at which point aerodynamic moments will begin to dominate the magnetic control torques generated by the satellite.

6.6.3   Simulation Result
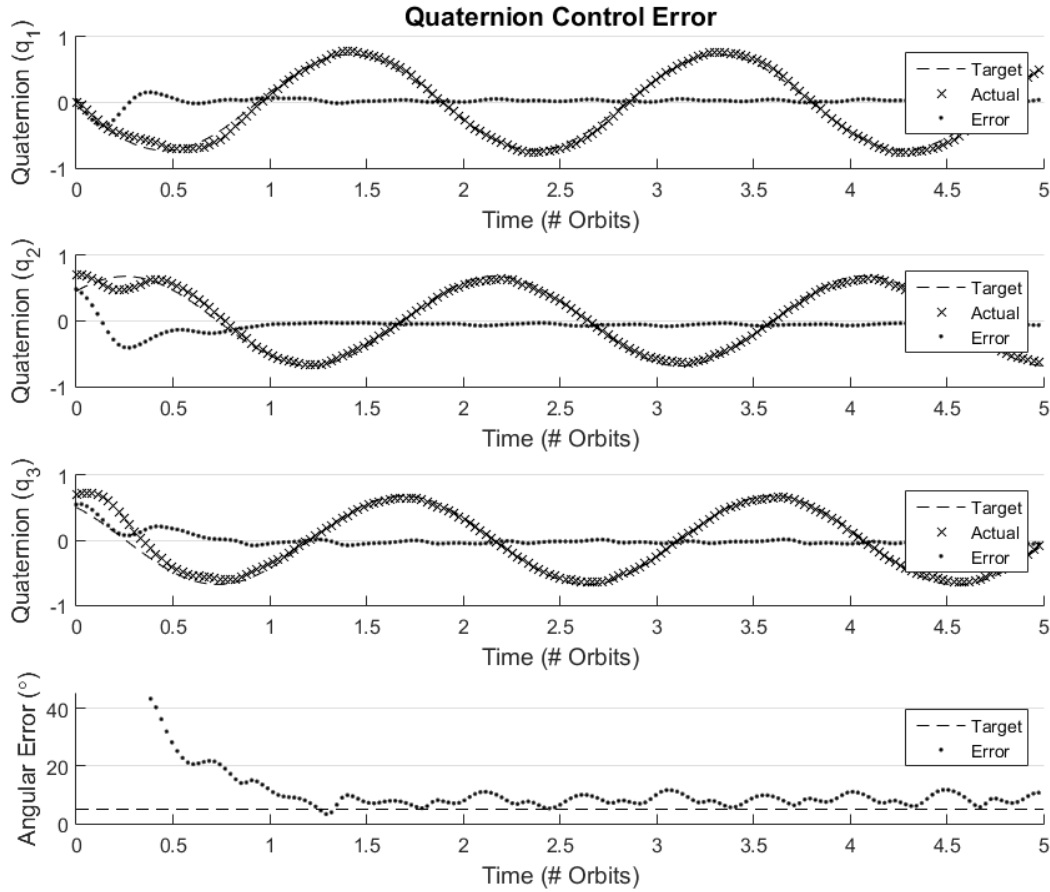
Figure 58: Attitude of the satellite with respect to the ECI frame for simulated satellite tracking a constant attitude target in the NTW frame. The attitude target aligned the +z satellite body axis with the satellite velocity vector and the –y satellite body axis towards Nadir. Simulation was for satellite with a 200km orbit averaged altitude to represent peak science operation.

The system slowly stabilizes into the attitude orientation; however, the attitude does not achieve the required 5° pointing requirement (see Figure 58). The angular rates converge towards the target; however, they do oscillate about the target rates (see Figure 59). The control moments remain very low (near the limit of the resolution of the torque coils) for the majority of the control case, despite the error in pointing. The plots demonstrate that these control gains are not well optimized for aerodynamic stabilization cases. Though the control effort is very low (which is desirable from a power management perspective), they fail to achieve the pointing

requirements of the science payloads. A re-optimization of these control gains with a HIL simulation is warranted. This simulation does, however, validate that the ADCS is able to track constant attitude and attitude rate targets in rotating reference frames (albeit poorly optimized).
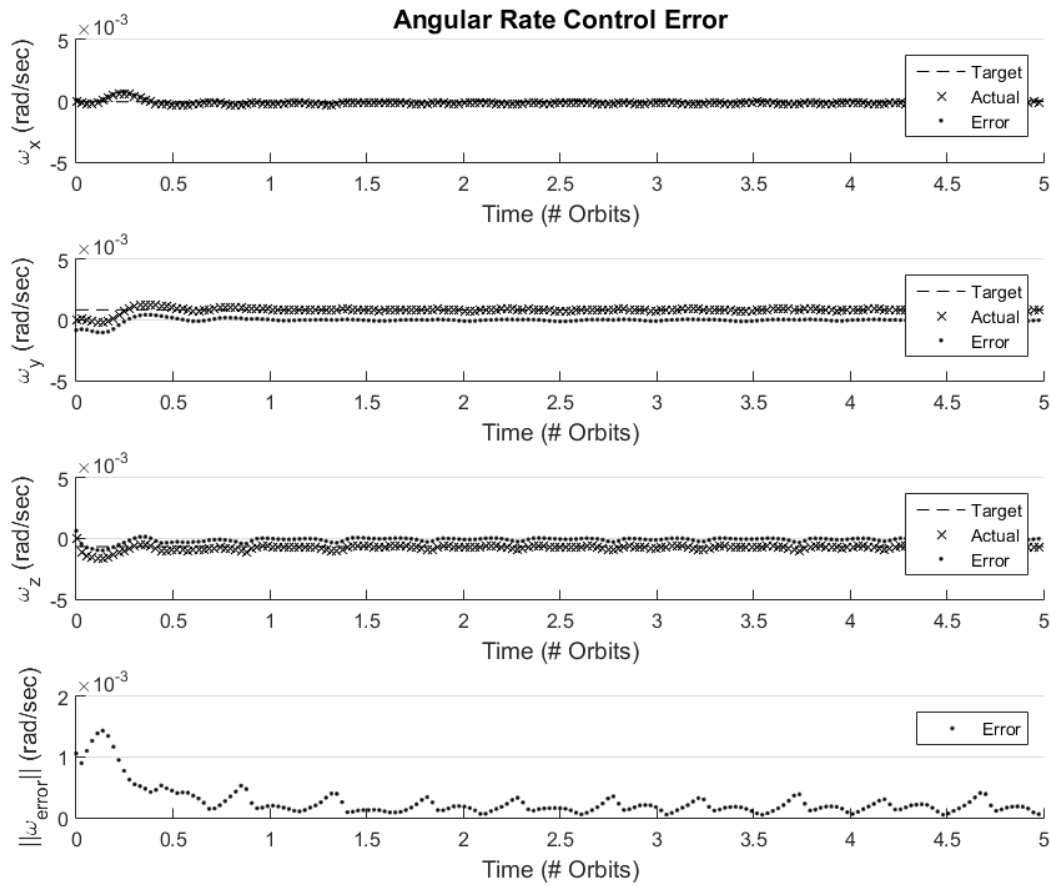
Figure 59: Angular rates of the satellite with respect to the ECI frame as seen from the ECI frame.
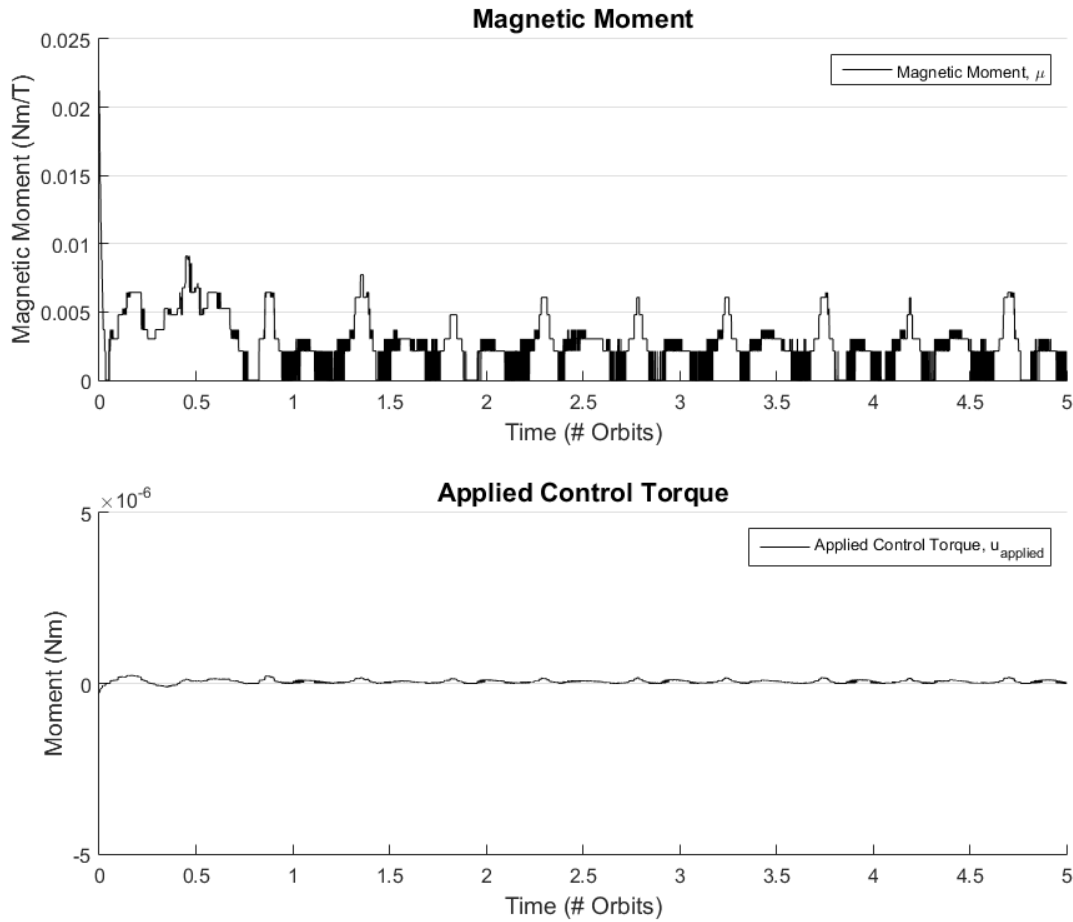
Figure 60: Plots of the norm of the applied magnetic moment and the resultant applied control torque.

6.7 HIL Optimization of Coupled ADCS under Aerodynamic Moments

With the ADCS validated using the HIL test platform, it was possible to re-optimize the ADCS filter and control gains. The advantage of performing the coupled optimization is that uncertainties in the attitude estimate no longer needed to be modelled using noise; the uncertainties are introduced via the determination system (from hardware/process noise) genuinely. This approach captures noise which is difficult to model, such as the uncertainties from magnetometers or sensor fusion algorithms.

The control gain parameters were seeded using the results of the Particle Swarm Optimization discussed in Chapter 3. The results shown in the previous subsection demonstrate the system

remains stable given perfect determination and is able to exceed the mission pointing requirements within an orbit. For the HIL simulation, 50 particles were seeded around this set of control gains. The span was determined based on the results of the Chapter 3 optimization. The number of iterations of the optimizer was increased to 15 iterations. The weights of the optimizer were altered to increase the velocity of each particle toward the globally optimally gain. This approach increases the rate of convergence to a solution at the cost of discounting sub-optimal multimodal peaks. Since the software simulations demonstrated a smooth, single peak minima, the assumptions implicit in the changes to the optimizer are justified. Notably, this change significantly reduces the amount of time required to refine the control gains.

6.7.1   Optimization Results

The optimization was able to significantly reduce the cost of the controller. By examining the control parameters with respect to the normalized simulation cost function values, it is possible to see the resolved minima which do not lie at the search boundary (see Figure 61 and Figure 62). If the minima exist at the boundaries, it is likely that the true minimum exists beyond the boundary. While this is not a guarantee of optimality, it improves confidence in the optimized gains. Similarly, while the minima are resolved, the minima exist in regions of similar cost. This result indicates that there is a region of control gains which will provide similar weighting which could be search for secondary optimality conditions such as energy cost.
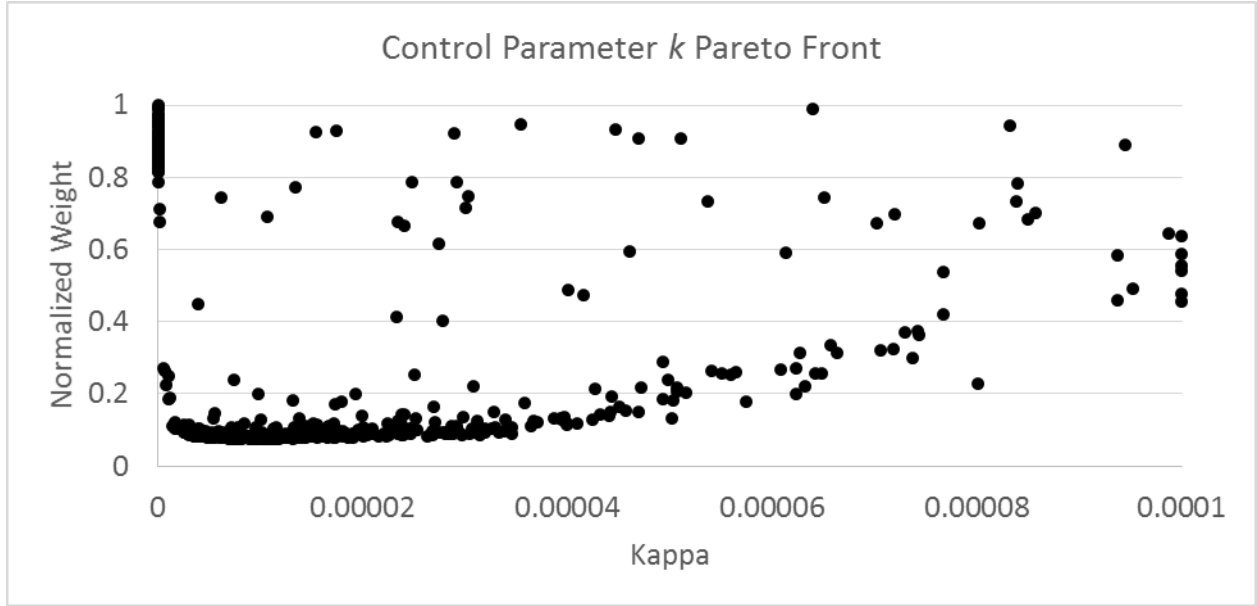
Figure 61: Optimization front of quaternion error control gain.
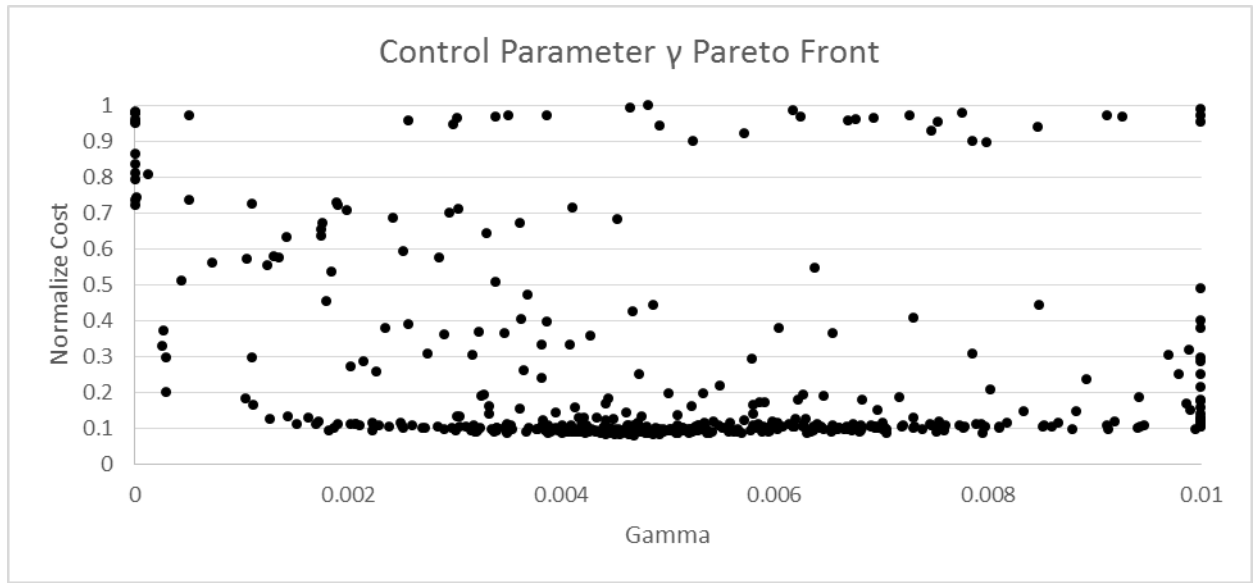
Figure 62: Optimization front of angular rate error control gain.

The control responses of the HIL simulations are present below. The simulations use the same initial conditions and attitude/angular rate targets as in Section 6.6. Within 1.5 orbits, the satellite is able to achieve better than the 5° pointing requirement and maintain the attitude target to better than 2° for the remainder of the simulation (Figure 63). The angular body rates also closely track the angular rate target to within an average norm error of 7.85E-5 rad/sec (Figure 64).
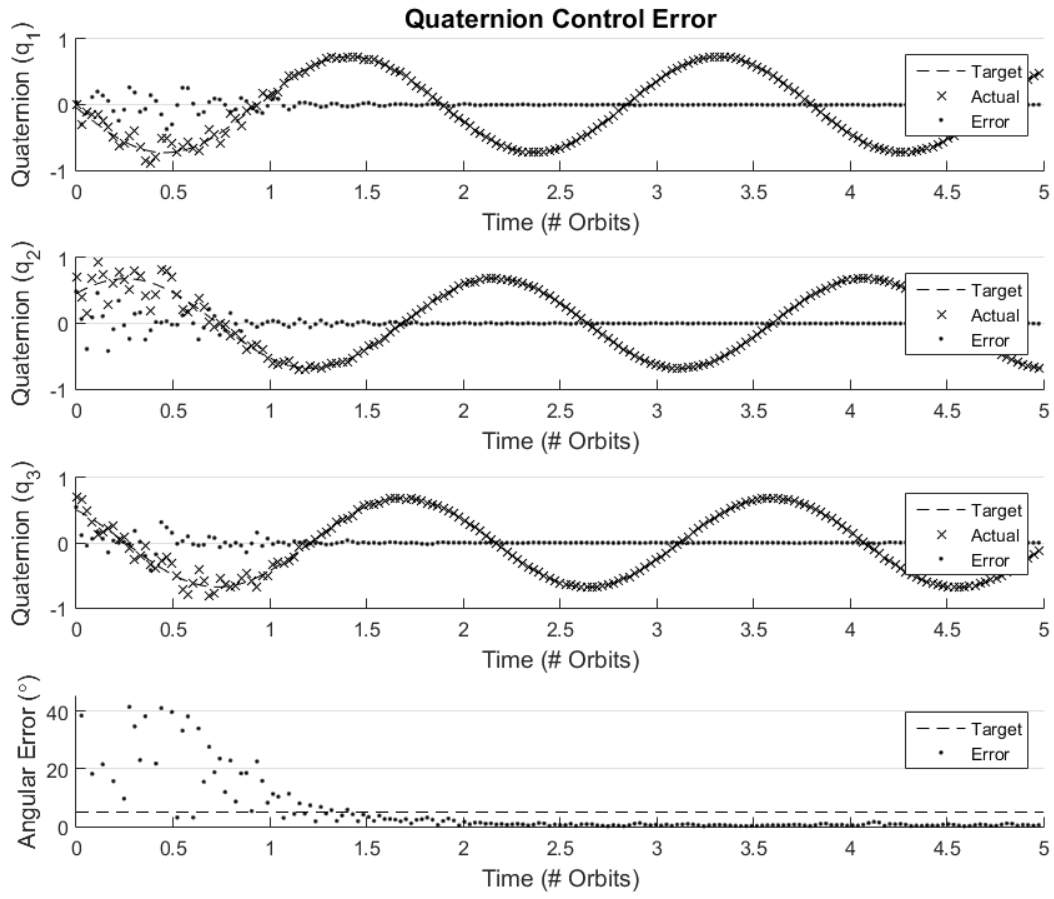
Figure 63: Attitude response of the satellite with respect to the ECI frame for simulated satellite tracking a constant attitude target in the NTW frame after optimization.
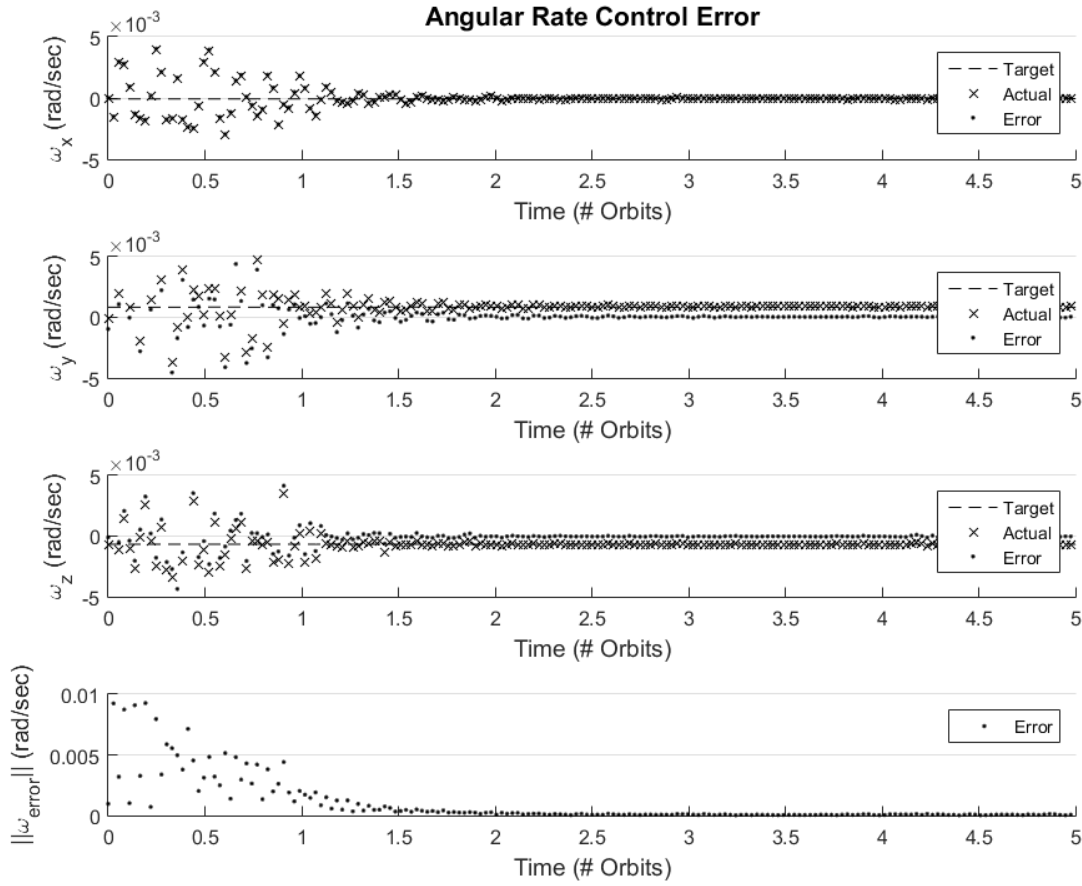
Figure 64: Post-optimization angular rates of the satellite with respect to the ECI frame as seen from the ECI frame.

The magnetic moments generated by the satellite to control the attitude do reach saturation at several instances within the first ¼ orbit (Figure 65). The applied magnetic moments then decrease on average to the 2 orbit mark. After that time, the applied magnetic moments become so small that they begin to be clipped by the discrete output levels of the magnetic torque coils. The applied torques, similarly, decrease in proportion to the applied magnetic moments. These figures demonstrate that once oriented into the aerodynamically stable configuration, with the appropriate angular body rate (which is approximately constant in the ECI frame), the satellite requires minimal actuation to remain stable.
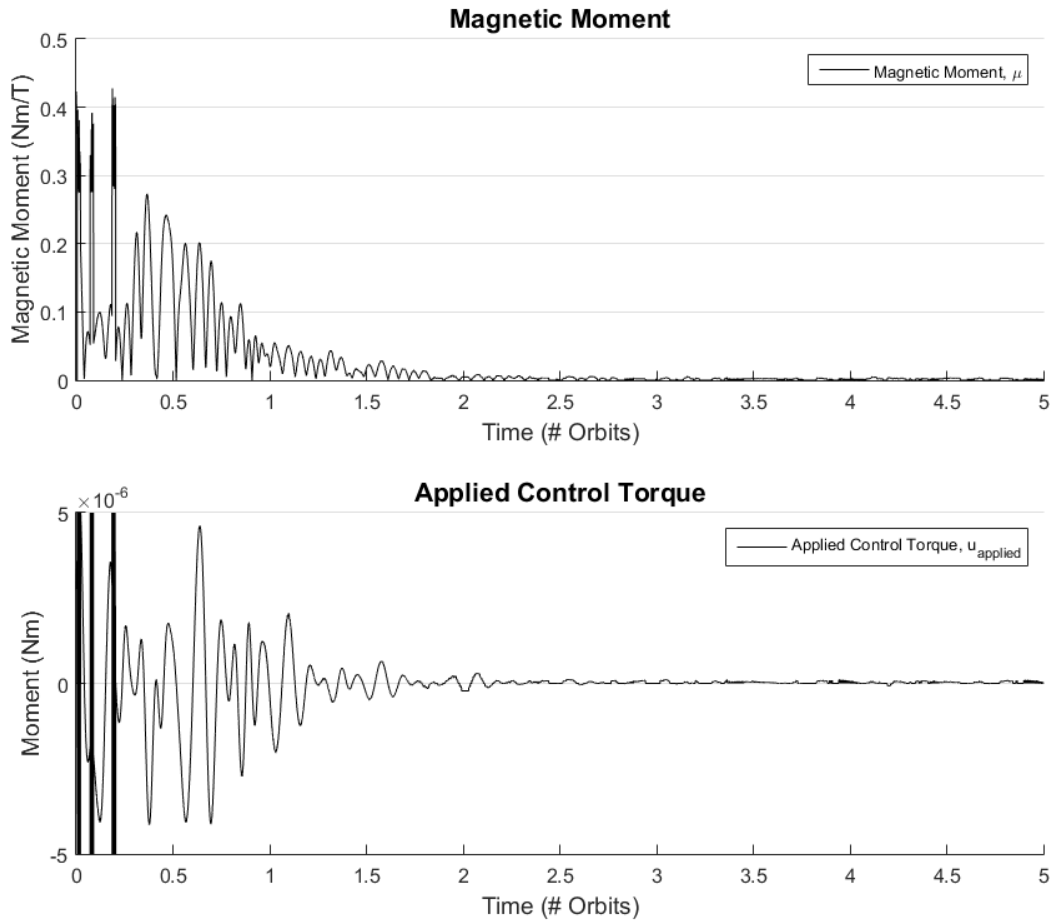
Figure 65: Plots of the norm of the applied magnetic moment and the resultant applied control torque after optimizing for attitude state error.

6.8   Summary of Contributions

In this chapter, a HIL system was developed to test the attitude determination and control systems. A test of the ADCS was conducted using flight magnetometers performing readings of the magnetic field generated by a computer controller Helmholtz coil. The satellite C&DH system was emulated. In this experiment, it was demonstrated that the attitude determination and control system could detumble and reorient the satellite to track a constant target in the NTW frame with no knowledge of its initial attitude state and using only low-cost magnetometer and MEMS gyro measurements as feedback for determination. Further, this system was employed to

perform HIL refinements to the optimization of attitude control gains. These experiments increase the reliability of low-cost, magnetic attitude control systems for NanoSatellites in a manner previously not possible.

CHAPTER 7
FUTURE WORK

The LAICE satellite is scheduled for delivery for launch in Spring 2017. The LAICE satellite and ADCS control system must continue to undergo preflight testing to improve the confidence in the design. To date, no test of the ADCS has been completed with the full satellite bus; only tests of the ADCS subsystem attached to inactive solar panels have been conducted. A test of the full system should be completed as soon as possible to verify that electromagnetic interference (EMI) from other subsystems do not generate noise that overwhelms the sensors. The discussion of future work will focus on three parts: pre-flight testing, on-orbit analysis, and refinements to CubeSim for future missions.

7.1 Pre-Flight Testing

Prior to flight, the LAICE ADCS must be repeatedly tested. To date, the attitude determination system has only been tested as a flat-sat. In the flat-sat configuration, the components are arranged flat and the interfaces are not finalized. In many cases, only the flight sensors have been located inside the HC3 when the simulations are running. This approach does not therefore include the additional of noise from nearby systems (e.g. payloads, batteries, solar panels, etc.). To date, this noise was modelled in the simulation though it may not accurately capture the noise characteristics of the fully assembled satellite. As of the time of this writing, the finalized LAICE payloads have not been available for such testing. Further, certain payloads cannot be operated in atmospheric conditions. As such, it may not be possible to capture and characterize the full noise environment of the LAICE satellite prior to launch. That said, a full test of the ADCS with the fully integrated satellite in the HC3 environment will provide strong validation of the system's performance.
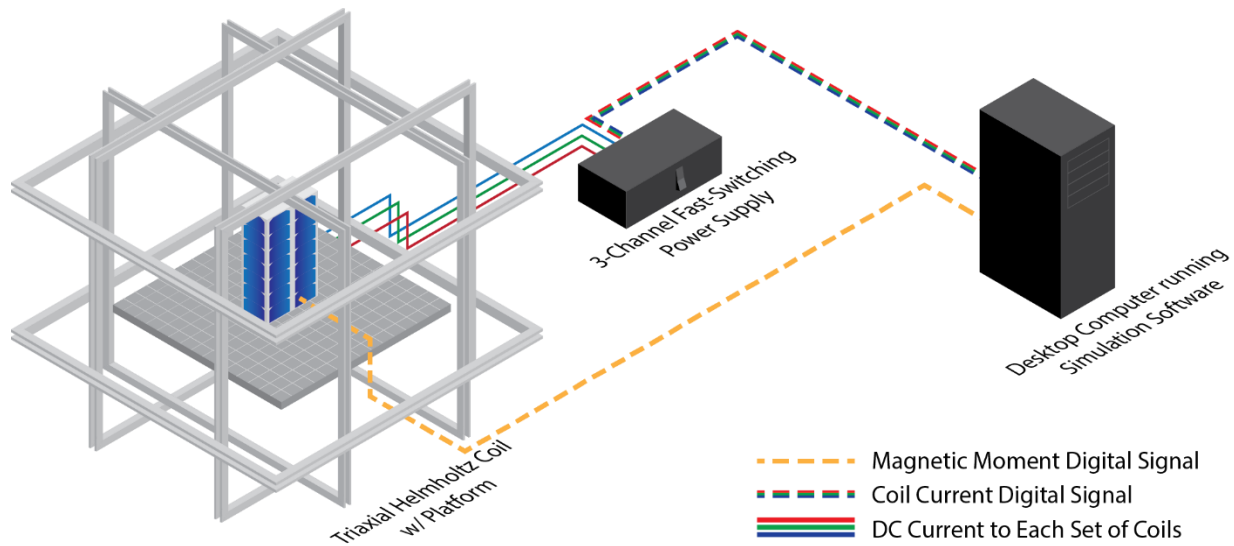
Figure 66: Hardware-in-Loop Attitude Determination and Control with LAICE Satellite

7.2 On-Orbit Analysis

Once launched, analysis of the LAICE attitude determination and control system is critically important. While ground simulation provides insight into how such a system may perform on orbit, the actual performance on orbit serves as a far more definitive analysis of the proposed system's performance. The LAICE satellite sensor readings should be collected along with the onboard attitude estimates for the initiation phase of the mission. The onboard sensors readings should be run using the emulated satellite software to verify that the emulated and real attitude estimates are consistent. Further, the filter residuals should be examined to ensure that the onboard determination algorithms have converged as expected.

One very important investigation will be studying the flight sensor performance, specifically with respect to sensor noise. Simulated sensors currently simulate the noise as Gaussian white noise of zero mean and uniform standard deviation. In practice, it is expected that onboard noise sources (e.g. radios, solar panels, payloads, etc.) will affect the sensor noise intermittently as they will not always be in operation. Some of these effects will be observed and modelled during the

pre-flight testing phase; however, the data garnered from LAICE's sensors will inform future simulations and improve the robustness of future IlliniSat ADCS.

7.3   Refinements to CubeSim

CubeSim as a concept can be extended to other existing CubeSat missions currently under development at the University of Illinois. Specifically, the CubeSail mission which is also scheduled for launch in 2017 will need to undergo simulated attitude determination and control testing. The CubeSim system can undergo a number of refinements to improve its performance and improve model fidelity.

The IGRF magnetic field model can be improved by implementing the more accurate HD Geomagnetic Model. This model is able to far more accurately replicate the Earth's magnetic field due to improvements in the model approach. Implementation of the model on CubeSim would be a simple extension and would not represent a significant computational burden for the simulation computer. This effort would be well suited for undergraduate/early-graduate independent study.

Refinements to the aerodynamic force model discussed in Chapter 5 should be adopted. Namely, the use of stochastic models using surface normal encoded renders. This improvement would increase the fidelity of the simulations by providing more realistic torque generation for the large mean-free path flows experienced on orbit. Investigation into performing the stochastic analysis directly on the GPU would significantly improve the performance of such a system. Further, taking advantage of the parallelism of GPUs, it is possible to have high bandwidth aerodynamic torque simulation. This research would be well suited to Master's level research.

The same GPU approach could similarly be used for calculating solar radiation pressure (SRP). While solar radiation pressure is relatively low order, the order of the refinements suggested warrant examining SRP effects. In principle, this extension is straightforward as it only requires rendering the satellite from a different view angle (i.e. along the sun vector as opposed to the anti-velocity vector).

Finally, refinements to the CubeSim software should be made to improve its utility in testing satellite hardware. Automatic magnetometer registration and calibration algorithms should be developed. These processes exist piecemeal across a variety of software packages, but an all-in-one solution would simplify the solution. The ultimate goal of such a system would be to place a satellite inside the Helmholtz cage, then run a program which calibrates each sensor, registers the orientation of each sensor with respect to the satellite body axes, then updates the onboard sensor correction values. This effort is well suited for undergraduate level research, specifically as an independent study project.

CHAPTER 8
CONCLUSIONS

The presented design for the LAICE attitude determination and control system offers a low-cost, highly reliable alternative to conventional attitude control system. The use of commercial off-the-shelf components significantly reduces the cost of the system. In order to reduce risk, extensive ground testing and system optimization using a novel hardware-in-loop (HIL) system was used. The HIL system includes novel approaches to the modelling of aerodynamic moments using a GPU accelerated approach. Further, both the flight and simulation software uses a custom object-oriented attitude library purpose built for this application. The library provides a straight forward system for programmers which alleviates the need for maintaining multiple code bases for flight and ground simulation.

The attitude determination system has been demonstrated to accurately determine the attitude and attitude rate even under external disturbances, large sensor noise, and control actuation. The attitude control system, taking advantage of passive stabilization built into the design of the LAICE bus structure, is able to reorient the LAICE spacecraft. Even from the largest expected deployment induced tumble rates, the LAICE spacecraft is able to detumble and reorient the spacecraft within a matter of 2-3 orbits. The spacecraft is also able to maintain two-point science attitude control using a combination of passive stabilization and active magnetic control.

APPENDIX A: TWO-LINE ELEMENT FOR SIMULATIONS

Two-Line Element (TLE) describing motion of LAICE Satellite in simulations. This is adapted from a TLE for the International Space Station (ISS-Zarya).

```
1    IL-LAICE
2    1 99999U 98067A   16028.60081312  .00014289  00000-0  21385-3 0
     9994
3    2 99999  51.6413  39.1283 0006529  51.2720 308.9374
     15.54305299983116
```

APPENDIX B: QUATERNIONS OF 120-CELL 4-POLYTOPE

Table of 300 Unique Quaternions Derived from the Positive Hemisphere of the 120-Cell 4-Polytope.

q_1	q_2	q_3	q_4	q_1	q_2	q_3	q_4
-0.20717	0.55115	0.03693	-0.80743	0.46964	-0.38127	-0.76071	-0.23538
-0.76071	-0.38127	0.46964	-0.23538	0.13042	0.59069	0.13042	0.78554
0.61712	-0.06026	0.70527	0.34369	0.03074	0.65751	-0.26215	-0.70570
-0.22267	0.57869	-0.00879	-0.78451	-0.60157	-0.26717	-0.51341	0.55059
-0.55348	-0.23925	-0.79757	-0.01703	0.23278	-0.60582	0.57127	0.50244
0.46189	-0.38761	0.34215	0.72066	0.43730	0.47956	0.19321	-0.73585
0.33567	0.44480	-0.78856	-0.26010	-0.38299	-0.90429	-0.04405	-0.18343
-0.54806	0.10079	0.12080	0.82151	0.98004	0.06279	-0.14425	0.12157
-0.72843	0.18813	0.65759	-0.03966	0.29114	0.50345	-0.46298	0.66890
0.15839	-0.73547	-0.33025	-0.57003	-0.58054	0.03455	0.80547	0.11399
-0.39951	0.33369	-0.39951	0.75461	0.85362	-0.50667	-0.11936	0.01913
-0.11936	-0.50667	0.85362	0.01913	0.07057	0.99016	0.07057	0.09815
-0.33603	-0.76950	-0.31157	-0.44484	-0.24717	0.06617	0.82754	0.49970
0.83849	0.04438	-0.25808	0.47786	0.05160	-0.25061	0.07593	-0.96372
0.39372	0.35077	-0.31571	-0.78885	-0.88164	-0.43940	-0.17217	0.00139
-0.49000	0.19481	-0.56377	0.63570	0.78786	-0.59130	0.07839	0.15330
0.42903	0.73971	0.40768	0.32024	0.09411	0.03311	-0.84397	-0.52704
-0.85354	0.05211	0.10362	-0.50796	-0.07177	0.06930	-0.09301	0.99066
-0.58036	-0.04169	-0.58036	0.56976	0.76938	-0.50019	0.29726	0.26362
0.29726	-0.50019	0.76938	0.26362	0.10252	0.91194	0.10252	-0.38386
-0.81699	-0.14923	0.26594	-0.48943	-0.07158	0.30461	0.01428	0.94968
0.49955	-0.66347	-0.55671	-0.01839	-0.24023	-0.20050	0.84256	-0.43838
0.52716	-0.10307	0.52716	-0.65847	-0.82486	0.51748	-0.14401	-0.17632
-0.14401	0.51748	-0.82486	-0.17632	-0.09312	-0.96928	-0.09312	0.20774
0.15663	0.76027	0.12889	0.61713	0.48157	-0.22536	-0.73745	-0.41650
-0.74986	-0.20063	0.49396	-0.39175	-0.01431	0.53150	-0.04192	0.84590
0.72198	-0.19962	0.64484	0.15197	-0.18746	0.66595	-0.11031	-0.71358
-0.14504	0.73489	-0.15308	-0.64457	-0.51136	-0.46599	-0.58851	0.41836
0.73508	-0.64919	-0.00004	-0.19545	-0.50396	0.29071	0.58047	-0.56971
0.25082	0.94810	-0.17470	0.08764	-0.03397	-0.58080	-0.76947	-0.26351
0.45948	-0.33105	-0.72147	-0.39844	-0.56203	-0.49682	0.61911	-0.23238
0.09008	0.55911	-0.03293	0.82353	0.69598	-0.27984	-0.48516	-0.44936
-0.28534	0.05163	-0.75371	-0.58978	-0.39954	-0.69702	0.29139	0.51926
0.08175	-0.27821	-0.18936	0.93812	0.80216	-0.04471	0.33355	-0.49323
-0.45839	0.63820	-0.45839	0.41529	0.57668	-0.69579	-0.36509	0.22366
-0.36509	-0.69579	0.57668	0.22366	0.36356	0.82735	0.36356	0.22613
-0.42020	-0.49270	-0.75588	0.09661	0.27468	-0.45491	0.77920	0.33234
0.62875	-0.15489	0.42506	0.63246	0.28361	0.44940	-0.05203	-0.84552
-0.20356	-0.63246	0.73444	0.13839	0.21494	0.94000	0.23624	0.11995
0.65663	0.10137	-0.20533	-0.71861	-0.93668	-0.22898	0.00140	-0.26495
-0.24827	0.69228	0.27475	-0.61938	-0.44243	-0.06464	-0.74200	0.49952
-0.50762	-0.53217	-0.67683	0.03193	0.25605	-0.36656	0.77906	0.43948
0.70890	-0.29975	-0.41524	0.48495	0.11569	-0.23680	0.45409	-0.85108
-0.04967	0.76807	0.61961	0.15393	0.15986	0.20953	-0.96433	-0.02461
0.65759	0.18813	-0.72843	-0.03966	-0.33025	-0.73547	0.15839	-0.57003
-0.46298	0.50345	0.29114	0.66890	0.80547	0.03455	-0.58054	0.11399
0.30670	0.62892	0.22954	-0.67654	-0.74224	-0.07673	-0.66509	0.02910
-0.69968	-0.00761	-0.70767	0.09793	0.40117	-0.62918	0.32402	0.58155
-0.00004	-0.64919	0.73508	-0.19545	-0.17470	0.94810	0.25082	0.08764
0.58047	0.29071	-0.50396	-0.56971	-0.76947	-0.58080	-0.03397	-0.26351

q_1	q_2	q_3	q_4	q_1	q_2	q_3	q_4
0.53925	-0.48960	-0.64204	-0.23937	-0.45569	-0.35446	0.72545	-0.37474
0.19608	0.70146	0.07369	0.68123	0.52107	-0.24856	-0.66007	-0.48064
-0.75371	0.05163	-0.28534	-0.58978	-0.18936	-0.27821	0.08175	0.93812
0.29139	-0.69702	-0.39954	0.51926	0.33355	-0.04471	0.80216	-0.49323
0.83460	-0.52873	0.12509	0.09079	-0.29287	0.34878	0.41659	-0.78678
0.09872	0.98304	0.02512	-0.15251	-0.18058	-0.41811	-0.89005	-0.01990
0.26594	-0.14923	-0.81699	-0.48943	-0.55671	-0.66347	0.49955	-0.01839
0.01428	0.30461	-0.07158	0.94968	0.84256	-0.20050	-0.24023	-0.43838
0.70527	-0.06026	0.61712	0.34369	-0.00879	0.57869	-0.22267	-0.78451
-0.26215	0.65751	0.03074	-0.70570	-0.51341	-0.26717	-0.60157	0.55059
-0.37127	-0.60281	-0.61538	0.34654	0.47617	-0.28000	0.81466	0.17662
0.70527	-0.06178	0.58555	0.39484	0.03693	0.55115	-0.20717	-0.80743
0.23022	-0.43427	0.23022	-0.83989	-0.85398	0.40046	0.28500	0.17068
0.28500	0.40046	-0.85398	0.17068	0.02482	-0.94288	0.02482	-0.33127
0.61063	0.66847	0.39037	0.16701	-0.12442	0.04932	-0.77253	-0.62071
-0.87121	0.24638	-0.02561	-0.42383	0.01166	-0.13333	-0.20876	0.96877
-0.01103	0.95972	-0.01103	-0.28053	-0.24586	-0.43980	-0.85309	0.13554
-0.85309	-0.43980	-0.24586	0.13554	0.50335	-0.02279	0.50335	0.70198
-0.16457	-0.07091	0.88411	-0.43155	-0.31210	0.82220	-0.33267	0.34046
0.13715	-0.11534	-0.78201	-0.59695	-0.59514	-0.64748	0.45345	0.14481
0.73444	-0.63246	-0.20356	0.13839	-0.20533	0.10137	0.65663	-0.71861
0.23624	0.94000	0.21494	0.11995	0.00140	-0.22898	-0.93668	-0.26495
0.27180	-0.34546	0.79481	0.41841	0.25229	0.86535	-0.04725	-0.43045
0.18685	0.39787	0.01747	-0.89804	-0.88658	-0.16264	-0.36353	0.23530
0.85090	0.00000	0.00000	0.52532	0.08837	0.00000	0.00000	-0.99609
-0.14553	0.61517	-0.14553	0.76106	0.74559	-0.40575	-0.48475	-0.21090
-0.48475	-0.40575	0.74559	-0.21090	0.02866	0.84836	0.02866	0.52786
-0.38127	0.46964	-0.76071	-0.23538	-0.03960	-0.89050	-0.07900	0.44632
-0.61891	0.19635	0.31914	-0.69032	-0.33978	0.19732	-0.31854	0.86264
0.10211	-0.64123	-0.76013	0.02438	-0.02396	-0.39219	0.91387	-0.10221
0.27475	0.69228	-0.24827	-0.61938	-0.67683	-0.53217	-0.50762	0.03193
-0.74200	-0.06464	-0.44243	0.49952	0.77906	-0.36656	0.25605	0.43948
0.49021	-0.04036	0.49021	0.71955	0.42355	0.45632	-0.18361	-0.76070
-0.18361	0.45632	0.42355	-0.76070	-0.59774	0.17413	-0.59774	0.50506
0.88411	-0.07091	-0.16457	-0.43155	-0.78201	-0.11534	0.13715	-0.59695
-0.33267	0.82220	-0.31210	0.34046	0.45345	-0.64748	-0.59514	0.14481
0.30029	-0.57412	0.30029	-0.70003	-0.76037	0.52571	0.37872	0.04521
0.37872	0.52571	-0.76037	0.04521	-0.12928	-0.91533	-0.12928	-0.35883
0.39037	0.66847	0.61063	0.16701	-0.02561	0.24638	-0.87121	-0.42383
-0.77253	0.04932	-0.12442	-0.62071	-0.20876	-0.13333	0.01166	0.96877
0.18984	0.79229	0.27801	-0.50887	-0.53999	-0.10647	-0.83287	0.05829
-0.79336	-0.18533	-0.57949	-0.02050	0.33716	-0.43503	0.42534	0.71845
-0.79757	-0.23925	-0.55348	-0.01703	0.34215	-0.38761	0.46189	0.72066
0.57127	-0.60582	0.23278	0.50244	0.19321	0.47956	0.43730	-0.73585
-0.31571	0.35077	0.39372	-0.78885	-0.56377	0.19481	-0.49000	0.63570
-0.17217	-0.43940	-0.88164	0.00139	0.07839	-0.59130	0.78786	0.15330
-0.57001	-0.64197	0.51278	0.00309	0.25817	0.74608	0.34416	0.50820
0.82933	-0.22194	-0.22700	-0.45981	-0.78267	-0.10349	0.30012	-0.53539
0.64484	-0.19962	0.72198	0.15197	-0.15308	0.73489	-0.14504	-0.64457
-0.11031	0.66595	-0.18746	-0.71358	-0.58851	-0.46599	-0.51136	0.41836
0.40592	0.00806	-0.32958	-0.85237	-0.94390	-0.29826	0.14047	0.01904
-0.18909	0.35928	-0.61434	0.67657	0.68958	-0.71020	-0.04592	-0.13411
-0.72147	-0.33105	0.45948	-0.39844	-0.03293	0.55911	0.09008	0.82353
0.61911	-0.49682	-0.56203	-0.23238	-0.48516	-0.27984	0.69598	-0.44936
0.18035	-0.87712	-0.28826	0.33918	0.22241	0.13531	0.91330	-0.31321
0.70333	0.55425	0.43239	0.10573	-0.22063	0.13821	-0.68924	-0.67615
-0.78856	0.44480	0.33567	-0.26010	0.12080	0.10079	-0.54806	0.82151
-0.04405	-0.90429	-0.38299	-0.18343	-0.14425	0.06279	0.98004	0.12157

q_1	q_2	q_3	q_4	q_1	q_2	q_3	q_4
-0.62300	-0.02262	0.76309	0.17046	0.43154	0.69233	-0.32170	0.48058
0.29937	-0.54683	-0.18953	-0.75858	-0.81231	0.07558	0.57378	0.07226
0.31914	0.19635	-0.61891	-0.69032	-0.76013	-0.64123	0.10211	0.02438
-0.31854	0.19732	-0.33978	0.86264	0.91387	-0.39219	-0.02396	-0.10221
0.79481	-0.34546	0.27180	0.41841	0.01747	0.39787	0.18685	-0.89804
-0.04725	0.86535	0.25229	-0.43045	-0.36353	-0.16264	-0.88658	0.23530
-0.18512	0.09284	-0.18512	0.96065	0.94176	-0.20710	-0.00005	-0.26494
-0.00005	-0.20710	0.94176	-0.26494	-0.23683	0.93473	-0.23683	0.11876
-0.75588	-0.49270	-0.42020	0.09661	0.42506	-0.15489	0.62875	0.63246
0.77920	-0.45491	0.27468	0.33234	-0.05203	0.44940	0.28361	-0.84552
0.22954	0.62892	0.30670	-0.67654	-0.70767	-0.00761	-0.69968	0.09793
-0.66509	-0.07673	-0.74224	0.02910	0.32402	-0.62918	0.40117	0.58155
-0.32958	0.00806	0.40592	-0.85237	-0.61434	0.35928	-0.18909	0.67657
0.14047	-0.29826	-0.94390	0.01904	-0.04592	-0.71020	0.68958	-0.13411
-0.64204	-0.48960	0.53925	-0.23937	0.07369	0.70146	0.19608	0.68123
0.72545	-0.35446	-0.45569	-0.37474	-0.66007	-0.24856	0.52107	-0.48064
-0.28826	-0.87712	0.18035	0.33918	0.43239	0.55425	0.70333	0.10573
0.91330	0.13531	0.22241	-0.31321	-0.68924	0.13821	-0.22063	-0.67615
-0.40766	0.34995	-0.40766	0.73835	0.84274	-0.52125	-0.13019	0.03361
-0.13019	-0.52125	0.84274	0.03361	0.08838	0.98697	0.08838	0.10133
-0.31157	-0.76950	-0.33603	-0.44484	-0.25808	0.04438	0.83849	0.47786
0.82754	0.06617	-0.24717	0.49970	0.07593	-0.25061	0.05160	-0.96372
0.12509	-0.52873	0.83460	0.09079	0.02512	0.98304	0.09872	-0.15251
0.41659	0.34878	-0.29287	-0.78678	-0.89005	-0.41811	-0.18058	-0.01990
0.40768	0.73971	0.42903	0.32024	0.10362	0.05211	-0.85354	-0.50796
-0.84397	0.03311	0.09411	-0.52704	-0.09301	0.06930	-0.07177	0.99066
-0.59220	-0.01807	-0.59220	0.54615	0.75360	-0.52133	0.28151	0.28471
0.28151	-0.52133	0.75360	0.28471	0.12839	0.90731	0.12839	-0.37924
0.51278	-0.64197	-0.57001	0.00309	-0.22700	-0.22194	0.82933	-0.45981
0.34416	0.74608	0.25817	0.50820	0.30012	-0.10349	-0.78267	-0.53539
0.53791	-0.12452	0.53791	-0.63702	-0.81052	0.53669	-0.12971	-0.19546
-0.12971	0.53669	-0.81052	-0.19546	-0.11663	-0.96508	-0.11663	0.20354
0.12889	0.76027	0.15663	0.61713	0.49396	-0.20063	-0.74986	-0.39175
-0.73745	-0.22536	0.48157	-0.41650	-0.04192	0.53150	-0.01431	0.84590
-0.41524	-0.29975	0.70890	0.48495	0.61961	0.76807	-0.04967	0.15393
0.45409	-0.23680	0.11569	-0.85108	-0.96433	0.20953	0.15986	-0.02461
0.76309	-0.02262	-0.62300	0.17046	-0.18953	-0.54683	0.29937	-0.75858
-0.32170	0.69233	0.43154	0.48058	0.57378	0.07558	-0.81231	0.07226
0.27801	0.79229	0.18984	-0.50887	-0.57949	-0.18533	-0.79336	-0.02050
-0.83287	-0.10647	-0.53999	0.05829	0.42534	-0.43503	0.33716	0.71845
-0.61538	-0.60281	-0.37127	0.34654	0.58555	-0.06178	0.70527	0.39484
0.81466	-0.28000	0.47617	0.17662	-0.20717	0.55115	0.03693	-0.80743

APPENDIX C: ECI TO NTW FRAME TRANSFORMATION

The frame transformation from ECI→NTW can be decomposed into two separate rotations: first from ECI→PQW, then from PQW→NTW. For simplicity of computing successive rotations, Direct Cosine Matrices (DCM) are used.

$$^{NTW}\boldsymbol{R}^{ECI} = {}^{NTW}\boldsymbol{R}^{PQW\,PQW}\boldsymbol{R}^{ECI}$$

The transformation from ECI to the perifocal frame is derived in Reference [14].

$$^{PQW}\boldsymbol{R}^{ECI} = \begin{bmatrix} \cos\Omega\cos\omega - \cos i \sin\Omega\sin\omega & \sin\Omega\cos\omega + \cos i\cos\Omega\sin\omega & \sin i\sin\omega \\ -\cos\Omega\sin\omega - \cos i\sin\Omega\cos\omega & -\sin\Omega\sin\omega + \cos i\cos\Omega\cos\omega & \sin i\cos\omega \\ \sin\Omega\sin i & -\cos\Omega\sin i & \cos i \end{bmatrix}$$

Where i is orbit inclination, Ω is the longitude of the ascending node, and ω is the argument of perigee. Assuming the orbit does not precess, then both frames are inertial and therefore the transformation is time independent.

The velocity of the spacecraft in the perifocal frame can be defined by:

$$\vec{v} = \frac{\mu}{h}\begin{bmatrix} -\sin f \\ e + \cos f \\ 0 \end{bmatrix}$$

Where f is the true anomaly, μ is the gravitational parameter for Earth, and h is the orbital energy. Using these terms, we may define the angle θ to be the angle between perigee and the satellite velocity vector, then

$$\theta \equiv atan2(-\sin f, e + \cos f)$$

The transformation from the perifocal frame to the satellite normal frame can then be defined as the rotation by θ about the $\widehat{W}$ axis.

$$
{}^{NTW}\boldsymbol{R}^{PQW} = \begin{bmatrix} \cos\theta & -\sin\theta & 0 \\ \sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}
$$

The two rotations can then be combined to obtain the transformation from ECI→NTW. The transpose of this rotation matrix will be the inverse rotation from NTW→ECI. Using this matrix, it is possible to obtain the time varying ECI frame attitude target for a constant attitude target in NTW.

To specify the angular body rate target of the spacecraft with respect to the ECI frame as seen in the ECI frame for a constant attitude target in the NTW frame, we will apply the vector addition formula for attitude rates.

$$
{}^{ECI}\omega_{sc/ECI} = {}^{ECI}\omega_{sc/NTW} + {}^{ECI}\omega_{NTW/ECI}
$$

Since the target attitude is fixed to the NTW frame, the target angular rate would be zero, therefore we can rewrite this equation as:

$$
{}^{ECI}\omega_{sc/ECI} = {}^{ECI}\omega_{NTW/ECI}
$$

To derive angular rate of the NTW frame with respect to the ECI frame, we will apply the rotation matrix derivative formula[17].

$$
\frac{d}{dt}\boldsymbol{C} = -[\omega_b \times]\boldsymbol{C}
$$

$$
[\omega_b \times] = \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix}
$$

Where $-[\omega_b \times]$ is the skew symmetric representation of the angular rate, and C is the rotation matrix representing the attitude of the rotating body. By post-multiplying both sides of the rotation matrix derivative by C^T, it is possible to isolate the angular body rate.

$$[\omega_b \times] = \frac{dC}{dt} C^T$$

In this application, the matrix C is the rotation from NTW→ECI and ω_b is the angular rate of the NTW frame with respect to the ECI frame. Recalling that the rotation from NTW→ECI can be decomposed into two rotations, we can rewrite this equation using the product of matrices. For clarity, we will refer to the rotation PQW→ECI as matrix A and the rotation NTW→PQW as B.

$$[\omega_{NTW/ECI} \times] = \frac{d(AB)}{dt}(AB)^T$$

Since the matrix A is constant with time, we can extract it from the derivative. Further, we can distribute the transpose recalling to invert the ordering when doing so.

$$[\omega_{NTW/ECI} \times] = A\frac{dB}{dt}B^T A^T$$

To take the derivative of matrix $B={}^{PQW}R^{NTW}$, we will simply take the derivative of the matrix elements.

$$B = {}^{PQW}R^{NTW} = ({}^{NTW}R^{PQW})^T = \begin{bmatrix} \cos\theta & \sin\theta & 0 \\ -\sin\theta & \cos\theta & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$\frac{dB}{dt} = \begin{bmatrix} -\sin\theta & \cos\theta & 0 \\ -\cos\theta & -\sin\theta & 0 \\ 0 & 0 & 0 \end{bmatrix}\frac{d\theta}{dt} = \begin{bmatrix} -\sin\theta & \cos\theta & 0 \\ -\cos\theta & -\sin\theta & 0 \\ 0 & 0 & 0 \end{bmatrix}\frac{(-e\cos f - 1)}{\sin^2 f + (e+\cos f)^2}\frac{df}{dt}$$

Unfortunately, there is no closed-form solution for the true anomaly f as a function of time for generic orbits. In a near circular orbit, the rate of change of true anomaly will remain roughly constant at a rate of 2π radians per orbit period. Using classical orbital elements to define orbital period, T_{orbit}, based on the orbit semi-major axis, a, and the gravitational parameter, μ.

$$T_{orbit} = 2\pi\sqrt{\frac{a^3}{\mu}}$$

Therefore, making the near circular orbit assumption, the derivative of matrix B becomes

$$\frac{dB}{dt} = \sqrt{\frac{\mu}{a^3}}\begin{bmatrix} -\sin\theta & \cos\theta & 0 \\ -\cos\theta & -\sin\theta & 0 \\ 0 & 0 & 0 \end{bmatrix}\frac{(-e\cos f - 1)}{\sin^2 f + (e + \cos f)^2}$$

Now, all values necessary to compute $\left[\omega_{NTW/ECI}\times\right]$ have been specified. The individual elements of $\omega_{NTW/ECI}$ can be extracted from the elements of the skew symmetric matrix. Therefore the angular body rate for a fixed attitude target in NTW can now be computed at any time step provided a model of the orbit (as provided by semi-analytic model such as SPG4 or a numerical orbit propagator).

REFERENCES

[1] J. Westerhoff, G. Earle, R. Bishop, and G. Swenson, "LAICE CubeSat Mission for Gravity Wave Studies," *Adv. Space Res.*, vol. 56, pp. 1413–1427, 2015.

[2] A. R. M. Ghosh, "Increasing cubesat form factor to 6U: The lower atmosphere/ionosphere coupling experiment," in *Proceedings of the International Astronautical Congress, IAC*, Toronto, ON, Canada, 2014, vol. 5.

[3] F. Azure, R. Hevner, and W. Holemans, "Lessons Learned Measuring 3U and 6U Payload Rotation and Velocity when Dispensed in Reduced Gravity Environment," presented at the 12th Annual CubeSat Workshop, 2015.

[4] "1, 2 and 3 Axis Magnetic Sensors HMC1051/HMC1052L/HMC1053," Honeywell, Datasheet, Jan. 2010.

[5] J. C. Springmann, "Attitude-Independent Magnetometer Calibration with Time-Varying Bias," presented at the AIAA/USU Small Satellite Conference, Logan, UT, 2011.

[6] R. Alonso and M. D. Shuster, "Complete linear attitude-independent magnetometer calibration.," *J. Astronaut. Sci.*, vol. 50, no. 4, pp. 477–490, 2003.

[7] R. Alonso and M. D. Shuster, "Attitude independent magnetometer-bias determination: A survey," *J. Astronaut. Sci.*, vol. 50, no. 4, pp. 453–475, 2003.

[8] E. P. Babcock and T. W. Bretl, "CubeSat Attitude Determination via Kalman Filtering of Magnetometer and Solar Cell Data," presented at the 25th Annual AIAA/USU Conference on Small Satellites, Logan, UT, 2011.

[9] Z. Diao, H. Quan, L. Lin, and Y. Han, "Analysis and compensation of MEMS gyroscope drift," in *2013 Seventh International Conference on Sensing Technology*, Massey University, New Zealand, 2013, pp. 596–600.

[10] "QRS11 MEMS Quartz Angular Rate Sensor," Systron Donner Inertial, Datasheet, 2016.

[11] "MPU-6000 and MPU-6050 Product Specification Revision 3.4," InvenSense Inc., Datasheet, Aug. 2013.

[12] M. Bezold, "An Attitude Determination System with MEMS Gyroscope Drift Compensation for Small Satellties," University of Kentucky, Lexington, KY, 2013.

[13] C. Haken, "A Wide Dynamic Range Switching Current Regulator for Cost Efficient Magnetic Attitude Control of NanoSatellites.," Senior Thesis, University of Illinois at Urbana-Champaign, Urbana, IL, USA, 2016.

[14] D. A. Vallado, *Fundamentals of Astrodynamics and Applications*, 3rd ed. Hawthorn, CA, USA: Microcosm Press, 2007.

[15] "The IAU Resolutions on Astronomical Constants, Time Scales, and the Fundamental Reference Frame.," US Naval Observatory, Washington D.C., USNO Circular No. 163, 1981.

[16] V. L. Pisacane, *Fundamentals of Space Systems*, 2nd ed. New York: Oxford Press, 2005.

[17] B. Wie, *Space Vehicle Dynamics and Control*. Reston, VA, USA: AIAA, 1998.

[18] G. Wahba, "A Least Squares Estimate of Spacecraft Attitude," *SIAM Rev.*, vol. 7, no. 3, p. 409, Jul. 1965.

[19] F. L. Markley, "Attitude Determination using Two Vector Measurements," in *Proceedings, Flight Mechanics Symposium*, NASA Goddard Space Flight Center, Greenbelt, MD, 1999, pp. 39–52.

[20] I. Y. Bar-Itzhack and Y. Oshman, "Attitude Determination from Vector Observations: Quaternion Estimation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-21, no. 1, pp. 128–135, Jan. 1985.

[21] M. L. Psiaki, "Nanosatellite Attitude Stabilization Using Passive Aerodynamics and Active Magnetic Torquing," *J. Guid. Control Dyn.*, vol. 27, no. 3, pp. 347–355, Jun. 2004.

[22] M. L. Psiaki, F. Martel, and P. K. Pal, "Three-Axis Attitude Determination via Kalman Filtering of Magnetometer Data," *J. Guid.*, vol. 13, no. 3, pp. 506–514, Jun. 1990.

[23] E. Silani and M. Lovera, "Magnetic spacecraft attitude control: a survey and some new results," *Control Eng. Pract.*, vol. 13, pp. 357–371, 2005.

[24] R. Yadlin, "Attitude Determination and Bias Estimation Using Kalman Filtering."

[25] M. J. Hale, P. Vergez, M. Meerman, and Y. Hashida, "Kalman Filtering and the Attitude Determination and Control Task," presented at the Space 2004 Conference, San Diego, CA, USA, 2004.

[26] E. Thébault *et al.*, "International Geomagnetic Reference Field: the 12th generation," *Earth Planets Space*, vol. 67, no. 1, p. 79, 2015.

[27] J. L. Crassidis, F. L. Markley, and Y. Cheng, "A Survey of Nonlinear Attitude Estimation Methods," *J. Guid. Control Dyn.*, vol. 30, no. 1, pp. 12–28, Jan. 2007.

[28] Y. Cheng and J. L. Crassidis, "Particle Filtering for Sequential Spacecraft Attitude Estimation," University of Buffalo.

[29] E. J. Lefferts, F. L. Markley, and M. D. Shuster, "Kalman Filtering for Spacecraft Attitude Estimation," *J. Guid.*, vol. 5, no. 5, pp. 417–429, Oct. 1982.

[30] M. L. Psiaki, "Spacecraft Attitude Rate Estimation from Geomagnetic Field Measurements," *J. Guid. Control Dyn.*, vol. 26, no. 2, pp. 244–252, Apr. 2003.

[31] R. Azor, I. Y. Bar-Itzhack, and R. R. Harman, "Satellite Angular Rate Estimation from Vector Measurements," *J. Guid. Control Dyn.*, vol. 21, no. 3, pp. 450–457, 1998.

[32] J. D. Searcy and H. J. Pernicka, "Magnetometer-Only Attitude Determination Using Novel Two-Step Kalman Filter Approach," *J. Guid. Control Dyn.*, vol. 35, no. 6, pp. 1693–1701, Dec. 2012.

[33] A. M. Sabatini, "Kalman-Filter-Based Orientation Determination Using Inertial/Magnetic Sensors: Observability Analysis and Performance Evaluation," *Sensors*, vol. 11, pp. 9182–9206, 2011.

[34] R. F. Stengel, *Stochastic Optimal Control*. New York: J. Wiley and Sons, 1986.

[35] S. Huang and G. Dissanayake, "Convergence and Consistency Analysis for Extended Kalman Filter Based SLAM," *IEEE Trans. Robot.*, vol. 23, no. 5, pp. 1036–1049, Oct. 2007.

[36] J. Arvo, "Fast Random Rotation Matrices," in *Graphic Gems III*, D. Kirk, Ed. Boston, MA, USA: Academic Press, 1992, pp. 117–120.

[37] K. Shoemake, "Uniform Random Rotations," in *Graphic Gems III*, D. Kirk, Ed. Boston, MA, USA: Academic Press, 1992, pp. 124–132.

[38] J. J. Kuffner, "Effective Sampling and Distance Metrics for 3D Rigid Body Path Planning," in *2004 IEEE International Conference on Robotics and Automation*, New Orleans, LA, USA, 2004.

[39] M. D. Shuster, "Uniform Attitude Probability Distributions," *J. Astronaut. Sci.*, vol. 51, no. 4, pp. 451–475, Dec. 2003.

[40] M. D. Shuster, "Maximum Likelihood Estimate of Spacecraft Attitude," *J. Astronaut. Sci.*, vol. 37, no. 1, pp. 79–88, Mar. 1989.

[41] J. A. Hartigan and M. A. Wong, "A K-Means Clustering Algorithm," *J. R. Stat. Soc. Ser. C Appl. Stat.*, vol. 28, no. 1, pp. 100–108, 1979.

[42] S. Mamone, G. Pileio, and M. H. Levitt, "Orientation Sampling Schemes Based on Four Dimensional Polytopes," *Symmetry*, no. 2, pp. 1423–1449, Jul. 2010.

[43] V. Capecchi, M. Buscema, P. Contucci, and B. D'Amore, Eds., *Applications of Mathematics in Models, Artificial Neural Networks and Arts*. New York: Springer, 2010.

[44] Tetracube, *120-Cell Perspective*. 2008.

[45] D. E. Gaylor, R. C. Page III, and K. V. Bradley, "Testing of the Java Astrodynamics Toolkit Propagator," presented at the AIAA/AAS Astrodynamics Specialist Conference and Exhibit, Keystone, CO, USA, 2006.

[46] N. A. Matteo and Y. T. Morton, "Ionosphere geomagnetic field: Comparison of IGRF model prediction and satellite measurements 1991–2010," *Radio Sci.*, vol. 46, 2011.

[47] M. L. Psiaki, "Magnetic Torquer Attitude Control via Asymptotic Periodic Linear Quadratic Regulation," *J. Guid. Control Dyn.*, vol. 24, no. 2, pp. 386–394, Apr. 2001.

[48] M. Y. Ovchinnikov, D. S. Roldugin, and V. I. Penkov, "Three-axis active magnetic attitude control asymptotical study," *Acta Astronaut.*, vol. 110, pp. 279–286, 2015.

[49] G. V. Smirnov, M. Ovchinnikov, and F. Miranda, "On the magnetic attitude control for spacecraft via the E-strategies method," *Acta Astronaut.*, vol. 63, pp. 690–694, May 2008.

[50] A. Pukniel, "Attitude Determination and Three-Axis Control System for Nanosatellites with Magnetic Torque Actuation," Master's, University of Illinois at Urbana-Champaign, Urbana, IL, USA, 206AD.

[51] "CubeSat Design Specification Rev. 13." The CubeSat Program, Cal Poly SLO, 20-Feb-2014.

[52] A. Pukniel, "The Dynamics and Control of the CubeSail Mission - A Solar Sailing Demonstration," Ph.D., University of Illinois at Urbana-Champaign, Urbana, IL, USA, 2010.

[53] M. Lovera, "Magnetic satellite detumbling: the b-dot algorithm revisited," presented at the 2015 American Control Conference, Chicago, IL, 2015, pp. 1867–1872.

[54] B. S. Leonard, "NPSAT1 Magnetic Attitude Control System," in *Proceedings of the 16th Annual AIAA/USU Conference on Small Satellites*, Logan, UT, 2002.

[55] A. Farhat, J. Ivase, Y. Lu, and A. Snapp, "Attitude Determination and Control System for CubeSat," Worcester Polytechnical Institute, Worcester, MA, 2013.

[56] P. Sekhavat, Q. Gong, and I. M. Ross, "Unscented Kalman Filtering: NPSAT1 Ground Test Results," presented at the AIAA Guidance, Navigation, and Control Conference, Keystone, CO, USA, 2006.

[57] J. Kennedy and R. Eberhart, "Particle Swarm Optimization," in *Proceedings., IEEE International Conference on Neural Networks*, 1995, pp. 1942–1948.

[58] D. P. Rini, S. M. Shamsuddin, and S. S. Yuhaniz, "Particle Swarm Optimization: Technique, System and Challenges," *Int. J. Comput. Appl.*, vol. 14, no. 1, pp. 19–27, Jan. 2011.

[59] J. W. Dally, W. F. Riley, and K. G. McConnell, Eds., *Instruments for Engineering Measurements*, 2nd ed. New York: J. Wiley and Sons, 1993.

[60] G. T. Gillies and R. C. Ritter, "Torsion balances, torsion pendulums, and related devices," *Rev. Sci. Instrum.*, vol. 64, no. 2, pp. 283–309, Sep. 1992.

[61] W. Chen, S. Zhang, and X. Long, "Angle measurement with laser feedback instrument," *Opt. Express*, vol. 21, no. 7, Apr. 2013.

[62] D. M. Schrello, "Passive Aerodynamic Attitude Stabilization of Near Earth Satellites," North American Aviation Inc., Columbus OH, Technical Report, 1961.

[63] E. I. Kroeker, A. R. M. Ghosh, and V. Coverstone, "Magnetic Attitude Control with Aerodynamic Stabilization for the LAICE Satellite," presented at the 26th AAS/AIAA Space Flight Mechanics Meeting, Napa, CA, 2016.

[64] S. Tanygin and G. M. Beatty, "GPU-Acceleratied Computation of SRP and Drag Forces and Torques with Graphical Encoding of Surface Normals," presented at the 26th AAS/AIAA Space Flight Mechanics Meeting, Napa, CA, 2016.

[65] E. H. Adelson, C. H. Anderson, J. R. Bergen, P. J. Burt, and J. M. Ogden, "Pyramid methods in image processing," *RCA Eng.*, vol. 29, no. 6, Dec. 1984.

[66] N. M. Laurendeau, *Statistical thermodynamics: fundamentals and applications*. Cambridge England: Cambridge University Press, 2005.

[67] E. M. Gaposchkin and A. J. Coster, "Analysis of Satellite Drag," *Linc. Lab. J.*, vol. 1, no. 2, 1988.

[68] L. G. Jacchia, "Atmospheric Models in the Region from 110 to 2000 km," in *COSPAR International Reference Atmosphere 1972*, Akademie-Verlag, Berlin, 1972.

[69] P. Haddox, "A Novel Magnetic Field Approach to Simulate Spacecraft Attitude Determination and Control," Master's Thesis, University of Illinois at Urbana-Champaign, Urbana, IL, USA, 2014.

[70] P. Haddox, A. R. M. Ghosh, E. I. Kroeker, and V. Coverstone, "Development of the IlliniSat-2 Attitude Determination and Control System Testing Suite," in *Advances in the Astronautical Sciences*, 2013.

[71] Vedant, E. I. Kroeker, P. Haddox, and A. R. M. Ghosh, "Hardware In Loop Simulation for Attitude Determination and Control of IlliniSat-2 Bus," presented at the 26th AAS/AIAA Space Flight Mechanics Meeting, Napa, CA, 2016.