

DESIGN OF A LOW-POWER MAGNETO-INDUCTIVE MAGNETOMETER

by

Rowan Antonuccio

A thesis submitted in partial fulfillment
of the requirements for the degree

of

MASTER OF SCIENCE

in

Computer Engineering

Approved:

Charles M. Swenson, Ph.D.
Major Professor

Jonathan Phillips, Ph.D.
Committee Member

Chris Winstead, Ph.D.
Committee Member

David F. Feldon, Ph.D.
Vice Provost of Graduate Studies

UTAH STATE UNIVERSITY
Logan, Utah

2025

Copyright © Rowan Antonuccio 2025

All Rights Reserved

ABSTRACT

Design of a Low-power Magneto-inductive Magnetometer

by

Rowan Antonuccio, Master of Science

Utah State University, 2025

Major Professor: Charles M. Swenson, Ph.D.
Department: Electrical and Computer Engineering

This thesis presents the design and implementation of a low-power magneto-inductive magnetometer for magnetic field measurements in the space environment. Building upon the principles of the PNI RM3100 sensor, the design focuses on power efficiency while incorporating two key innovations: physically isolated sensing coils mounted on a spacecraft boom to reduce electromagnetic interference, and a custom-designed FPGA firmware implementation optimized for both power consumption and measurement precision. The magnetometer operates by measuring the time differential between oscillation cycles in alternating current directions, achieving high resolution through high-speed clock cycle counting. Operating at configurable sample rates with 32-bit timing resolution per axis, this power-efficient design is specifically developed to support long-duration plasma bubble studies and other scientific objectives on upcoming NASA Low Earth Orbit missions where power and physical space constraints are critical. This design provides an innovative solution for space-based magnetic field measurements.

(101 pages)

PUBLIC ABSTRACT

Design of a Low-power Magneto-inductive Magnetometer

Rowan Antonuccio

Measuring magnetic fields in space helps scientists understand phenomena that can affect satellite communications and navigation systems on Earth. This research develops a new low-power magnetic field sensor for spacecraft that improves upon existing designs by moving the sensitive parts away from electrical interference and using energy-efficient digital electronics for precise measurements. The sensor's power-efficient design is particularly important for small satellites, where power is limited and must be carefully managed. It will fly on future NASA missions to study disturbances in Earth's upper atmosphere that can disrupt radio signals and GPS. This work contributes to our ability to better predict and understand space weather events that impact our everyday technology while advancing the development of power-efficient space instrumentation.

ACKNOWLEDGMENTS

I would like to thank Dr. Swenson for his insight, guidance, and patience throughout the development of this thesis. I would also like to thank the team at the Center for Space Engineering for their support and many conversations about “mag stuff” along the way. Justin Wellington, Kade Howes, and Josh Taylor worked hard to assist the analog, firmware, and mechanical design of this project.

I would also like to thank my family for supporting my continued education and helping my dreams of working on spacecraft come true.

Lastly, I would like to thank Ceaira for continuously putting up with my endless ramblings about the magnetometer.

Rowan M. Antonuccio

CONTENTS

	Page
ABSTRACT	iii
PUBLIC ABSTRACT	iv
LIST OF TABLES	viii
LIST OF FIGURES	ix
ACRONYMS	xiii
1 INTRODUCTION	1
2 LITERATURE REVIEW	3
2.1 Magneto-Inductive Effect	3
2.2 Development of Magneto-Inductive Magnetometers	3
2.2.1 Commercial Magnetometers for Space Applications	4
2.2.2 Limitations	4
2.2.3 Research Objectives	5
3 HARDWARE DESIGN	7
3.1 Circuit Design	7
3.1.1 Threshold Voltage Derivation	9
3.1.2 Hardware Development	12
4 FIRMWARE DESIGN	17
4.1 Firmware Development	17
4.1.1 Timing Optimization and Clock-Domain Management	19
4.1.2 State Machines and System Coordination	20

4.1.3	Polarity Switching, Measurement Sequencing, and Counter Coordination	21
4.1.4	Communication and Control	23
5	MECHANICAL INTEGRATION	26
5.1	Boom-Mounted Magnetometer Assembly	26
5.2	Sensor Head Enclosure Design Evolution	27
5.2.1	First Generation: Machined Delrin Enclosure	27
5.2.2	Second Generation: FR4 PCB-Based Enclosure	29
6	CALIBRATION	32
6.1	Calibration	32
6.1.1	Helmholtz Coil Characterization	32
6.1.2	Magnetometer Calibration Using the Calibrated HHC	34
6.1.3	Reporting Rate versus Integration Depth	36
6.1.4	Residuals Across Integration Depths	39
6.2	Noise Floor Analysis	41
6.2.1	Magnetic Shielding Chamber	41
6.2.2	Noise Floor Characterization	43
7	DISCUSSION	50
8	FUTURE WORK	53
	REFERENCES	55
	APPENDIX A Spice Models	58
	APPENDIX B HDL	60
	APPENDIX C Simulink Models	63
	APPENDIX D Calibration Plots for All Cycle-Count Configurations	72

LIST OF TABLES

Table	Page
6.1 Measured reporting rates as a function of integration depth.	37
6.2 Calibration results for all integration-time configurations.	38
6.3 Residual RMSE versus cycle count for each axis. Values in parentheses show percent of full scale.	40
6.4 Standard deviation (σ) and maximum absolute residual ($\max r $) for each integration depth and axis. All values in nT.	41
6.5 Median mid-band amplitude spectral density (ASD) for each integration-depth (cycle counts) configuration. Values represent the median ASD within the flat region of each spectrum (see Figs. 6.4–6.8).	44

LIST OF FIGURES

Figure	Page
3.1 Basic LR relaxation oscillator circuit using an inverting Schmitt trigger. . .	7
3.2 Square wave output of inverting Schmitt trigger inductor voltage.	8
3.3 Output of inverting Schmitt trigger inductor voltage.	9
3.4 Resistive network determining Schmitt trigger thresholds	10
3.5 Comparison of the inductor voltage waveform using (A) a switch with internal clamp diodes and (B) a switch supporting below ground signaling. The clamped device limited the negative excursion to approximately -0.4 V, distorting the LR discharge slope, while the proper switch preserved the full bipolar waveform required for accurate Schmitt-trigger timing. Image captured from Tektronix MDO3034.	13
3.6 Conceptual alignment of the sensing inductor, switching network, and inverting Schmitt trigger within the LR relaxation oscillator.	14
3.7 Prototype magnetometer electronics board incorporating the IGLOO2 FPGA, analog oscillator circuitry, polarity-control switches, and external sensing-head interface.	15
3.8 Power consumption with magnetometer in active sample.	16

- 4.1 Cycle count definition. (A) LR relaxation oscillation; (B) Schmitt-trigger square-wave output with one full cycle highlighted; (C) overlay of the 1 MHz (example) FPGA clock used for timing. A single *cycle count* corresponded to one complete charge–discharge period between two rising-edge transitions of the Schmitt output. 18
- 4.2 Top-level firmware flowchart. 19
- 4.3 Per-axis driver state machine controlling excitation, timing, and synchronization with the high-speed counter. 21
- 4.4 Handshake and polarity coordination between the axis driver and oscillation counter. Signals counting and delaying acted as busy indicators during measurement and delay intervals. 23
- 4.5 Data-Ready control logic. 23
- 4.6 Data-to-Send latching and concatenation logic. 24
- 5.1 Delrin boom enclosure with precision-machined internal coil channels and Langmuir Probe attached to boom. 28
- 5.2 Sensing head PCBs with two cabling configurations: (A) coaxial harness and (B) Gecko connector. 29
- 5.3 FR4 enclosure alongside sensor head PCB with surface-mounted inductors before assembly. 30
- 5.4 FR4 boom-mounted magnetometer enclosure without Langmuir Probe. 31
- 6.1 Helmholtz Coil (HHC) Calibration Setup 34
- 6.2 Sensitivity versus cycle counts plotted on a log–log scale. 39

6.3	Three-layer μ -metal magnetic shielding chamber used for noise testing with magnetometer attached. The magnetometer sensing head is located inside the center of the coil.	42
6.4	ASD at 100-cycle integration. One-sided ASD for X, Y, Z; $F_s \approx 836$ Hz with Nyquist frequency $f_N \approx 418$ Hz. The shaded region denotes the mid-band interval used to compute the median ASD values reported in Table 6.5.	45
6.5	ASD at 1,000-cycle integration. One-sided ASD for X, Y, Z; $F_s \approx 85.7$ Hz ($f_N \approx 42.8$ Hz). The shaded region indicated the frequency band over which the median ASD was computed (see Table 6.5).	46
6.6	ASD at 10,000-cycle integration. One-sided ASD for X, Y, Z; $F_s \approx 8.59$ Hz ($f_N \approx 4.30$ Hz). The shaded region marked the mid-band analysis interval used for the median ASD values in Table 6.5.	47
6.7	ASD at 20,000-cycle integration. One-sided ASD for X, Y, Z; $F_s \approx 4.30$ Hz ($f_N \approx 2.15$ Hz). The shaded region corresponded to the mid-band interval used to compute the median noise levels in Table 6.5.	48
6.8	ASD at 40,000-cycle integration. One-sided ASD for X, Y, Z; $F_s \approx 2.15$ Hz ($f_N \approx 1.07$ Hz). The shaded region represented the mid-band frequency range used for computing the median ASD values tabulated in Table 6.5.	49
7.1	Mid-band ASD versus cycle-count integration depth. Each point represents the median ASD within the shaded mid-band interval of Figs. 6.4–6.8, with numerical values listed in Table 6.5.	51
A.1	LTSpice Circuit Model.	59
C.1	Top Level Simulink.	64
C.2	High Level Simulink FSM.	65

C.3	Axis Driver Simulink Model.	66
C.4	Axis Driver Simulink FSM.	67
C.5	Data Ready FSM.	68
C.6	Data Ready Latching.	69
C.7	SPI Slave Mode Controller.	70
C.8	SPI FSM.	71
D.1	Counts vs. Field Calibration — 100cc.	73
D.2	Residuals vs. Fitted Field — 100cc.	74
D.3	Residual Histograms — 100cc.	75
D.4	Counts vs. Field Calibration — 1000cc.	76
D.5	Residuals vs. Fitted Field — 1000cc.	77
D.6	Residual Histograms — 1000cc.	78
D.7	Counts vs. Field Calibration — 10000cc.	79
D.8	Residuals vs. Fitted Field — 10000cc.	80
D.9	Residual Histograms — 10000cc.	81
D.10	Counts vs. Field Calibration — 20000cc.	82
D.11	Residuals vs. Fitted Field — 20000cc.	83
D.12	Residual Histograms — 20000cc.	84
D.13	Counts vs. Field Calibration — 40000cc.	85
D.14	Residuals vs. Fitted Field — 40000cc.	86
D.15	Residual Histograms — 40000cc.	87

ACRONYMS

- AC: Alternating Current
- ACMES: Active Cooling for Multispectral Earth Sensors
- ADC: Analog-to-Digital Converter
- ASD: Amplitude Spectral Density
- ASIC: Application-Specific Integrated Circuit
- CSE: Center for Space Engineering
- DAC: Digital-to-Analog Converter
- DARPA: Defense Advanced Research Projects Agency
- DC: Direct Current
- EMI: Electromagnetic Interference
- FFT: Fast Fourier Transform
- FPGA: Field-Programmable Gate Array
- FVM: Fluxgate Vector Magnetometer
- GEO: Geostationary Orbit
- GPS: Global Positioning System
- HDL: Hardware Description Language
- HHC: Helmholtz Coil
- HPF: High-Pass Filter
- Hz: Hertz
- IGLOO2: Microchip IGLOO2 FPGA family
- LEO: Low Earth Orbit
- LR: Inductor–Resistor
- MHz: Megahertz
- MI: Magneto-Inductive Effect

- NASA: National Aeronautics and Space Administration
- OP-AMP: Operational Amplifier
- PCB: Printed Circuit Board
- PSD: Power Spectral Density
- RF: Radio Frequency
- RM3100: PNI RM3100 MI magnetometer
- RMSE: Root Mean Square Error
- SPI: Serial Peripheral Interface
- SPORT: Scintillation Prediction Observations Research Task
- ULF: Ultra-Low Frequency
- VHDL: VHSIC Hardware Description Language
- VHSIC: Very High Speed Integrated Circuit

CHAPTER 1

INTRODUCTION

Magnetic fields arise from electric currents and from the intrinsic magnetic moments of elementary particles. They influence charged particle motion, electromagnetic wave propagation, and plasma behavior in both terrestrial and space environments. The Earth generates its magnetic field through dynamo action in the fluid outer core, where convective motion of electrically conducting molten iron produces a time-varying geomagnetic field. This field extends into space and forms the magnetosphere, a region dominated by Earth's magnetic influence. The magnetosphere protects the planet from the continuous flow of energetic particles emitted by the Sun, known as solar wind [1].

When solar wind interacts with the magnetosphere, it compresses the magnetic field on the dayside and stretches it into a long tail on the nightside. These interactions can produce geomagnetic storms, auroral substorms, and other variations in the near-Earth magnetic field. Such disturbances can disrupt navigation and communication systems, affect satellite performance, and induce currents in electrical infrastructure [2]. Precise in-situ magnetic field measurements are therefore important to the study of space weather and understanding the structure and dynamics of the magnetosphere.

Several magnetic sensing technologies have been developed for scientific, industrial, and aerospace applications: Hall-effect sensors are inexpensive and robust but have limited sensitivity. Fluxgate magnetometers offer more stability and accuracy than Hall-effect and are widely used on spacecraft but require more power and volume. Search-coil sensors work well for time-varying fields but cannot measure static or slowly changing fields. At the high end, optically pumped magnetometers and Superconducting Quantum Interference Device systems provide extremely high sensitivity but are complex, expensive, and impractical for small spacecraft.

Magneto-inductive (MI) sensors occupy a useful middle ground between the simplicity

of the Hall-effect and the sensitivity of fluxgate magnetic sensors. They require far less power than fluxgates, are simpler than optically pumped systems, and can achieve sensitivities well beyond Hall-effect devices. MI sensors operate by exploiting changes in the inductance of a magnetic core when exposed to an external field. This inductance change alters the oscillation period of an LR (inductor–resistor) relaxation oscillator. By driving the coil in alternating directions and measuring the time required to complete a set number of oscillation cycles, the sensor converts inductance changes into precise timing information. A well-known example of a commercial MI sensor is the PNI RM3100 magnetometer, which has demonstrated good radiation tolerance and reliable operation in space [3,4]. However, the RM3100 is limited to compassing applications due to its 13 nano-Tesla sensitivity [14]. Small size, digital time-domain readout, and low-power operation make MI sensors particularly attractive for CubeSats and other platforms where physical size, mass, and power are tightly constrained. The combination of efficiency and simplicity have driven growing interest in developing science-grade MI sensors with higher sensitivity for modern space missions [5].

This research aims to advance the application of MI sensors for small satellite missions by exploring their small-scale, space-grade, low-power, and robust design. Specifically, it focuses on optimizing the oscillator circuitry, the Field Programmable Gate Array (FPGA) firmware, and the mechanical design. Furthermore, the study investigates sensor calibration and characterization of the noise-floor for comparison with current commercial MI sensors. Together, these objectives aim to fill critical knowledge gaps and establish an innovative design for magnetic sensing.

CHAPTER 2

LITERATURE REVIEW

2.1 Magneto-Inductive Effect

The MI effect describes how the inductance of a magnetic material changes in response to an external magnetic field [6]. When the magnetic field modifies the permeability of the magnetic material, the inductance of the coil changes as well. Although the relationship between field and inductance is nonlinear, it is stable and repeatable within a given operating range which makes the MI effect suitable for precision sensing [7, 8].

MI magnetometers have gained attention in space applications due to their combination of low power consumption, simple construction, and good measurement performance [9]. Because they rely on timing measurements rather than analog signal levels, MI sensors avoid many of the noise issues common to analog front ends. The digital nature of the sensing method also improves radiation tolerance since timing-based measurements are generally more resistant to single-event effects compared to analog voltage or current readings [3].

2.2 Development of Magneto-Inductive Magnetometers

The MI principle has been used in a wide range of applications from scientific sensing to industrial automation. Sensitivities vary from nanoTesla (nT) to picoTesla(pT) depending on the design, but the shared advantages of low power and compact size make MI sensors strong alternatives to fluxgate and Hall-effect magnetometers. Several representative MI sensors include:

- **Liu et al. (2020) Modular MI Sensor:** A modular MI magnetometer for low-field measurements with a detection limit of 8.28 nT. Its modular design highlights adaptability and straightforward integration [10].

- **Aichi Steel AMI306:** A compact 3-axis MI sensor commonly used in automotive and industrial systems. It has a noise floor of about $20 \text{ pT}/\sqrt{\text{Hz}}$ at 1 Hz and a dynamic range of $\pm 100 \text{ } \mu\text{T}$ [11, 12].
- **Prance et al. (2000) Ultra-Low-Noise MI Magnetometer:** A laboratory-grade MI sensor with a $50 \text{ fT}/\sqrt{\text{Hz}}$ noise floor from 1–30 kHz using a Metglas 2714AF core and optimized analog processing. It operates stably across cryogenic and room-temperature conditions [13].
- **PNI RM3100:** A widely used MI magnetometer that employs an LR oscillator architecture where field-induced inductance changes alter the oscillation period. It achieves a noise floor of $4 \text{ pT}/\sqrt{\text{Hz}}$ and a dynamic range of $\pm 1100 \text{ } \mu\text{T}$, with very low power consumption [14, 15].

2.2.1 Commercial Magnetometers for Space Applications

Among the commercial MI sensors, the RM3100 and the Aichi Steel AMI306 are the most thoroughly characterized for space use as both have undergone radiation testing and space-environmental qualification [3].

The RM3100 has been tested up to 150 kRad(Si) total ionizing dose with no destructive single-event effects observed at linear energy transfer values up to $76.7 \text{ MeV}\cdot\text{cm}^2/\text{mg}$ [3]. It has been selected for several space missions, including the National Aeronautics and Space Administration’s (NASA’s) Artemis Gateway HERMES platform, demonstrating suitability for low Earth orbit (LEO), geostationary Earth orbit (GEO), and deep-space environments [3].

The Aichi Steel AMI306 is also designed for harsh space environments as it can withstand radiation and temperature extremes while maintaining stable performance. Its small size and low power draw make it well suited for CubeSats [12].

2.2.2 Limitations

Although the RM3100 and AMI306 are considered space-capable, each have limitations.

The RM3100 is cost-effective and easy to integrate, but its configuration options are constrained by its internal application specific integrated circuit (ASIC). While the device allows the user to select any cycle count between 1 and 65,535, the ASIC determines how the oscillation cycles are generated, accumulated, and timed [14]. This limits control over several important parameters, including fine control over the oscillation cycle count, the timing resolution and stability of the internal counters, and the digital output word length. Its PCB-mounted sensing coils can become misaligned during assembly, reducing orthogonality and introducing calibration uncertainty. In addition, the proximity of the coils to digital electronics can introduce electromagnetic interference.

The Aichi Steel sensor, while mechanically robust, has a higher noise floor of roughly $200 \text{ pT}/\sqrt{\text{Hz}}$ at 1 Hz, limiting its sensitivity for very weak signals [12]. Despite the appealing aspects of this sensor, it lacks documented data as a scientific grade MI sensor.

Therefore, these limitations motivate ongoing research into MI magnetometers with improved sensitivity, better mechanical alignment, and more flexible timing architectures. While the RM3100 and Aichi Steel MI sensors show potential for space mission integration, several critical knowledge and optimization gaps remain regarding their implementation. Specifically, there is a limited understanding of the Aichi Steel sensor’s true capabilities, promoting the RM3100 as the primary MI industry leader.

2.2.3 Research Objectives

This work aims to develop and evaluate a fully discrete component magnetometer based on the MI principles used by the RM3100, to target small satellite applications. The research objectives are outlined below:

The first objective was to develop a MI sensing architecture that allowed complete control over oscillator behavior, timing resolution, and sample integration depth. This included designing an LR relaxation oscillator that could be reconfigured for a range of operating conditions.

The second objective was to implement and verify a FPGA-based timing system capable of measuring LR oscillator periods with high precision. The FPGA architecture supported

configurable cycle counts, stable clocking, and reliable digital readout suitable for spacecraft instrumentation.

The third objective was to design a mechanically isolated, boom-mounted sensing head that improved axis orthogonality and reduced the influence of electromagnetic interference from nearby electronics. This mechanical separation aimed to enhance measurement accuracy and repeatability.

The fourth and final objective was to calibrate and characterize the magnetometer using a three-axis Helmholtz coil and a μ -metal magnetic shielding chamber. This calibration established linearity, sensitivity, and overall measurement accuracy, while the shielded chamber tests provided a detailed assessment of the sensor noise floor.

CHAPTER 3
HARDWARE DESIGN

3.1 Circuit Design

The magnetometer was designed following the MI sensing principle with an LR relaxation oscillator built around an inverting Schmitt trigger. As shown in Figure 3.1, the circuit consists of a sensing coil inductor (L_1), an op-amp (Operational Amplifier) based inverting Schmitt trigger (U_1), and a resistor network that defines the switching thresholds. The inductor, L_1 , is wound around a high-permeability mu-metal core, whose inductance varies in response to the external magnetic field strength.

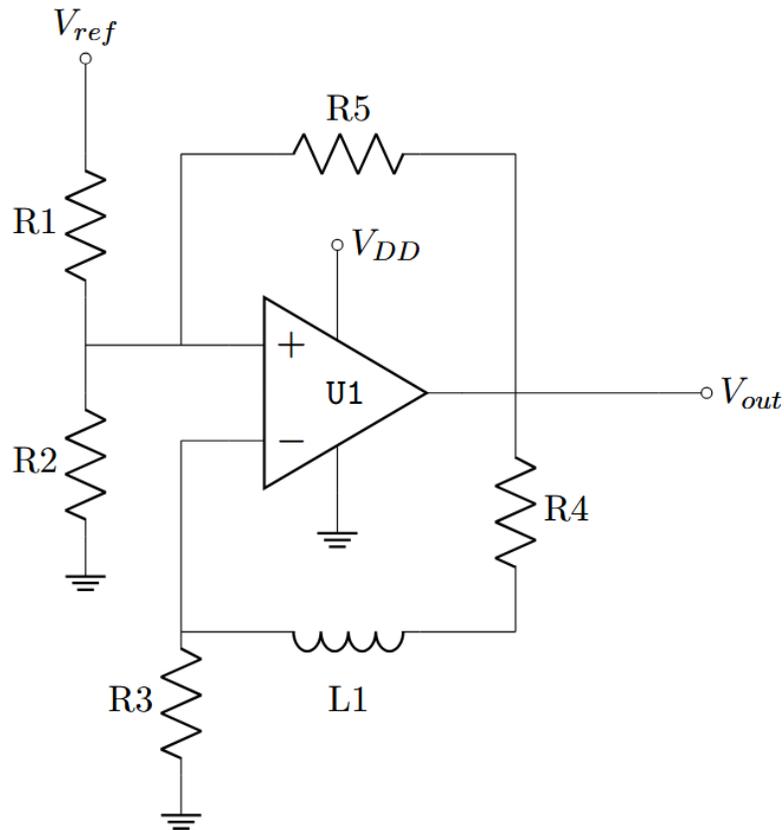


Fig. 3.1: Basic LR relaxation oscillator circuit using an inverting Schmitt trigger.

The inductor L_1 and resistor R_4 form a feedback path from the op-amp's output to its inverting input. When the circuit operates, current through L_1 charges the inverting input node. Due to the inductor's inherent delay, the voltage at this node ramps up gradually until it crosses a threshold, prompting the Schmitt trigger to switch states. The output then flips, and the inductor discharges in the opposite direction, repeating the cycle. This process creates a self-sustaining oscillation, shown in Figure 3.2, with a frequency determined by the LR time constant $\tau = L/R$ and the width of the Schmitt trigger's threshold band.

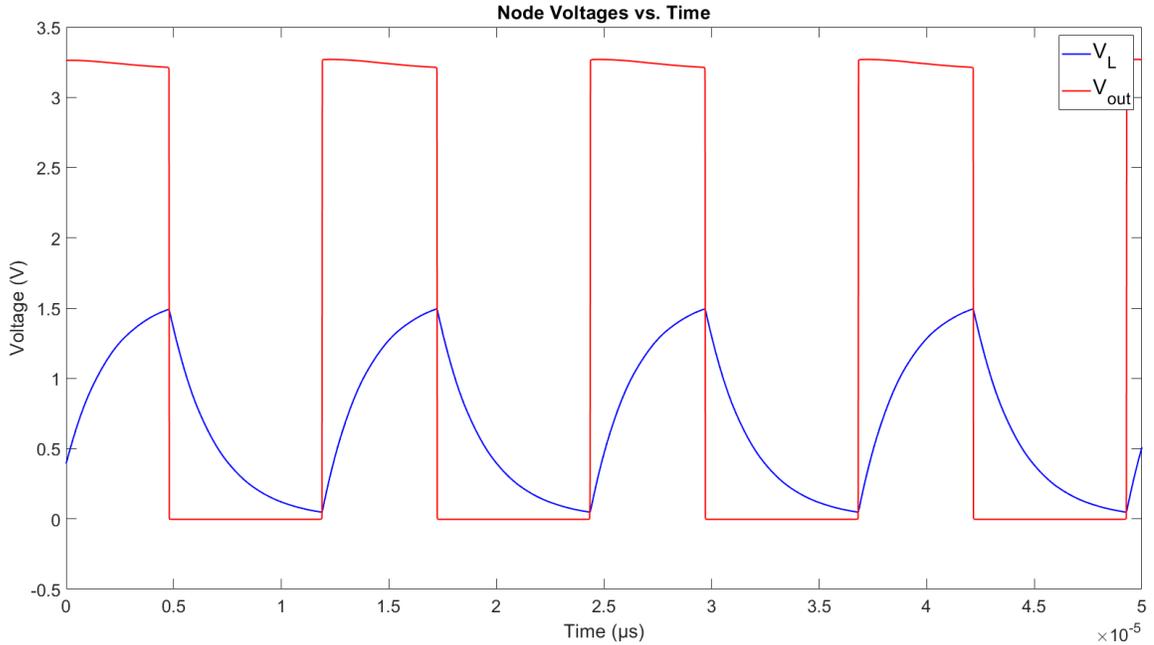


Fig. 3.2: Square wave output of inverting Schmitt trigger inductor voltage.

Resistors R_3 and R_4 serve essential roles in biasing and shaping the behavior of the oscillator. Resistor R_3 provides a DC return path from the inverting input through the inductor L_1 to ground, ensuring the input node remains properly referenced during both charging and discharging phases. Without this connection, the inverting input could float, leading to instability or undefined operation. Meanwhile, resistor R_4 , placed in series with L_1 , sets the effective resistance R that defines the LR time constant $\tau = L/R$, directly

controlling the voltage ramp rate toward the switching thresholds. Together, these resistors help ensure stable, predictable oscillation and proper interaction with the Schmitt trigger thresholds.

The oscillation frequency depends directly on the inductance of L_1 , which changes in response to the magnetic field. As the field increases or decreases, the permeability of the mu-metal core shifts, changing the inductance and, in turn, the rate at which the LR circuit charges. This affects how long it takes for the voltage to reach the switching threshold, altering the period of the Schmitt trigger's output waveform.

Since each oscillation produces a clean, sharp edge, the resulting square wave shown in Figure 3.2 is easily counted by the FPGA. By tracking the number of FPGA clock cycles during each oscillation, the system converts magnetic field changes into a time-based digital signal that can be analyzed. Figure 3.3 shows how the LR oscillation frequency changes in response to different magnetic fields with larger field strengths causing longer charge and discharge cycles.

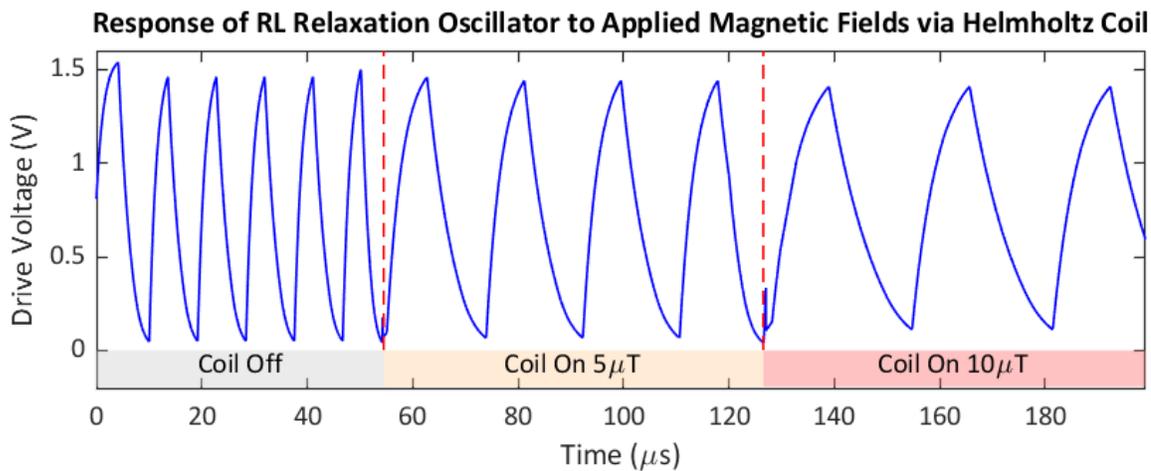


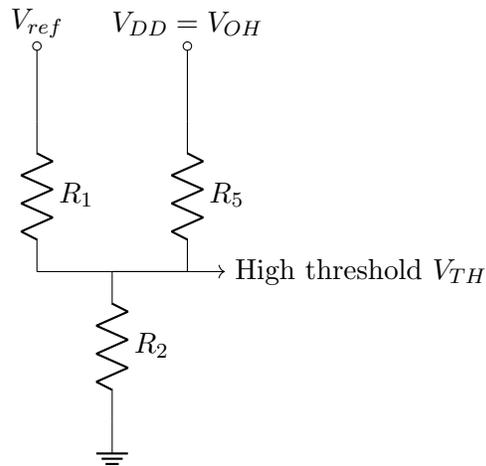
Fig. 3.3: Output of inverting Schmitt trigger inductor voltage.

3.1.1 Threshold Voltage Derivation

The clean switching behavior of the Schmitt trigger is defined by two voltage thresholds: an upper threshold V_{TH} and a lower threshold V_{TL} . These determine when the op-amp

output flips between its high and low states. The thresholds are controlled by a three-resistor network consisting of R_1 , R_2 , and the feedback resistor R_5 , driven by the stable voltage reference V_{ref} shown in Figure 3.4. V_{DD} is the supply voltage to the op-amp, which determines the voltage level of the output square wave V_{OH} for a logical high.

High Threshold State



Low Threshold State

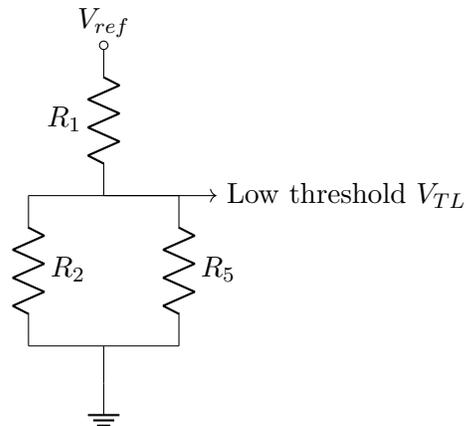


Fig. 3.4: Resistive network determining Schmitt trigger threshold: (top) high output pulls the node up via R_5 to set V_{TH} , (bottom) low output pulls the node down to set V_{TL} .

Resistor R_5 plays a critical role in establishing hysteresis by providing feedback from the output to the non-inverting input of the op-amp. When the output V_{out} in Figure 3.1

is high, current through R_5 biased the non-inverting input upward, increasing the effective switching threshold V_{TH} . This reinforces the current output state by requiring a larger opposing input signal to trigger a transition. Conversely, when V_{out} is low, the feedback current through R_5 lowered the input voltage, reducing the threshold to V_{TL} . This dynamic adjustment of the switching thresholds ensures that state transitions occurred only when the input voltage crossed well-defined boundaries, thereby improving noise immunity and suppressing false triggering.

To calculate these thresholds, the voltage V_+ at the non-inverting input is analyzed using superposition:

$$V_+ = \frac{R_2 R_5}{R_1 R_2 + R_1 R_5 + R_2 R_5} \cdot V_{\text{ref}} + \frac{R_1 R_2}{R_1 R_2 + R_1 R_5 + R_2 R_5} \cdot V_{\text{out}} \quad (3.1)$$

When $V_{\text{out}} = V_{OH}$, the output is high and $V_+ = V_{TH}$:

$$V_{TH} = \frac{R_2 R_5}{D} \cdot V_{\text{ref}} + \frac{R_1 R_2}{D} \cdot V_{OH} \quad (3.2)$$

When $V_{\text{out}} = V_{OL}$, the output is low and $V_+ = V_{TL}$:

$$V_{TL} = \frac{R_2 R_5}{D} \cdot V_{\text{ref}} + \frac{R_1 R_2}{D} \cdot V_{OL} \quad (3.3)$$

Here, $D = R_1 R_2 + R_1 R_5 + R_2 R_5$ is the common scaling denominator for both expressions.

The resulting hysteresis width was:

$$\Delta V_H = V_{TH} - V_{TL} = \frac{R_1 R_2}{D} \cdot (V_{OH} - V_{OL}) \quad (3.4)$$

This shows that the size of the hysteresis band is controlled by the output voltage swing and the resistor ratios but not by V_{ref} . The reference voltage instead determines the center point of the thresholds ensuring that they remained stable over time and temperature. Precision resistors and a low-drift voltage reference are used to guarantee long-term accuracy of the thresholds which directly influenced the stability of the oscillator frequency.

3.1.2 Hardware Development

With the theoretical operation of the oscillator established, the circuit was implemented in hardware and refined to ensure stable switching behavior, predictable interaction with the Schmitt trigger, and reliable manufacturability. Although the LR relaxation oscillator was conceptually simple, several practical considerations required specific design choices to preserve the correct waveform shape and threshold behavior.

A key constraint identified early in development was the behavior of the analog switches used to reverse the current direction through the sensing inductor. During each oscillation cycle, the inductor voltage naturally swung below ground as it discharged, and any device in series had to allow these bipolar waveforms to pass without distortion. Initial switch selections incorporated internal clamp diodes that limited the negative voltage to approximately -0.4 V, producing clipped waveforms and a slow discharge slope inconsistent with the theoretical LR response. Figure 3.5 illustrates this effect by comparing the distorted, clamped waveform produced by the original switch to the correct bipolar response. The clamped waveform exhibited a flattened negative waveform and a slow recovery slope while the proper bipolar waveform followed the expected exponential charge–discharge shape of the LR oscillator.

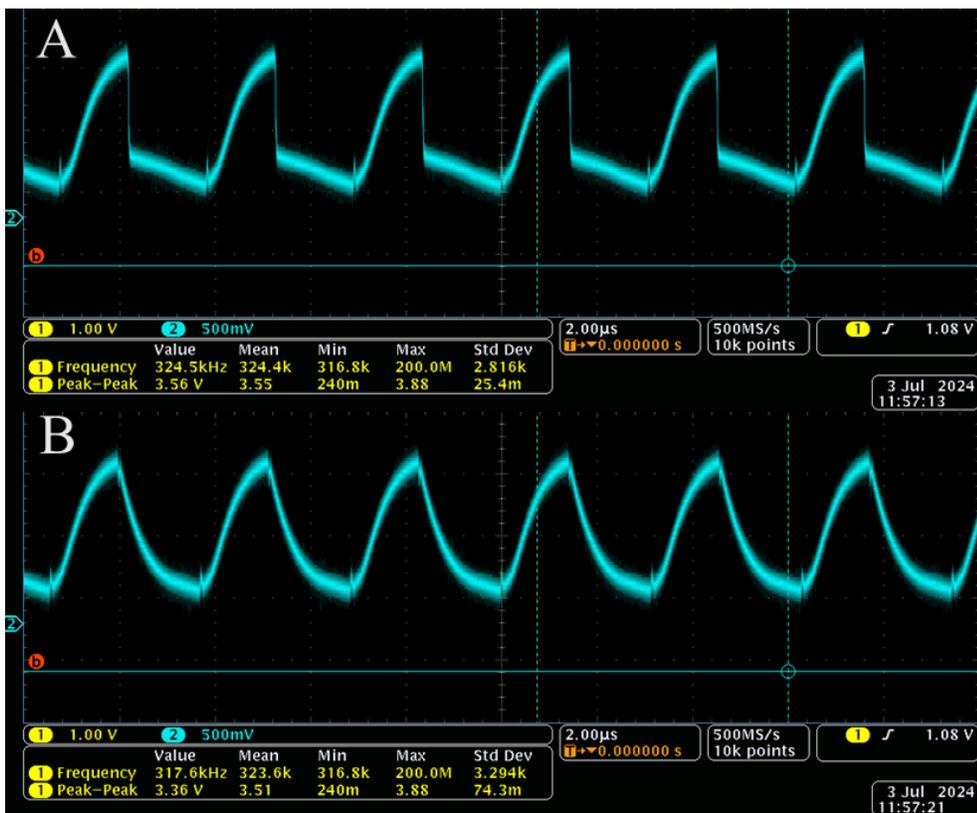


Fig. 3.5: Comparison of the inductor voltage waveform using (A) a switch with internal clamp diodes and (B) a switch supporting below ground signaling. The clamped device limited the negative excursion to approximately -0.4 V, distorting the LR discharge slope, while the proper switch preserved the full bipolar waveform required for accurate Schmitt-trigger timing. Image captured from Tektronix MDO3034.

A temporary solution using external flyback diodes protected the switch but altered the effective LR time constant, changing the oscillator dynamics. The problem was permanently resolved by selecting the Texas Instruments TS5A22362 analog switch which supported signals within $V_{CC} \pm 5$ V and therefore preserved the full bipolar ramp. This restored the expected exponential charge–discharge behavior and ensured that the Schmitt trigger transitioned cleanly at its defined thresholds.

Additional refinements improved the robustness and noise stability of the circuit. A small filtering capacitor was added at the inverting-input node to suppress minor ringing near the threshold crossings that could otherwise have caused unintended triggers of the Schmitt Trigger. The capacitor’s value was tuned to avoid excessive smoothing of the LR

waveform, which would have reduced sensitivity by masking small changes in the ramp slope. Because the sensing coils were mounted remotely on a boom and connected by long leads, each axis was provided with a high-value return resistor to ground during its inactive periods. This prevented the inductors from floating to unknown electrical potentials, reducing the risk of electrostatic charge buildup and ensuring consistent axis startup behavior.

In the final hardware, the additional circuit components shown in Figure 3.6—including the precision series resistors (R_6 , R_7), the high-value return resistor (R_8), the filtering capacitor (C_1), the polarity-control switches (SW_1 and SW_2), and the remotely mounted sensing inductor (L_1)—extended the basic oscillator of Figure 3.1 without altering its operating principle. These components enabled polarity reversal through , prevented floating nodes, suppressed ringing near the thresholds, and ensured stable operation when the coil was mounted on a boom.

These adjustments produced a stable and predictable analog oscillator suitable for precision timing measurements. Figure 3.6 summarized how the switching network, bias elements, and Schmitt-trigger core integrated to form the complete LR relaxation oscillator used in the magnetometer. The full simulation model in LTSpice is shown in Appendix A.1.

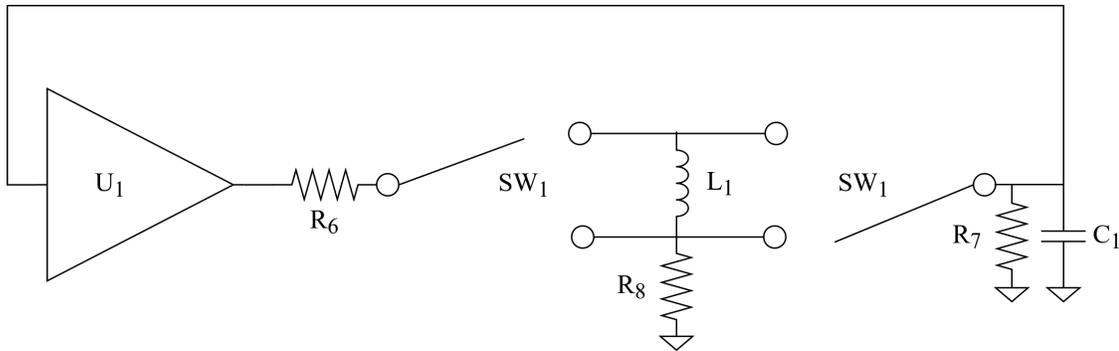


Fig. 3.6: Conceptual alignment of the sensing inductor, switching network, and inverting Schmitt trigger within the LR relaxation oscillator.

These refinements resulted in the complete hardware implementation of the magnetometer, integrating the analog front end, switching network, and FPGA timing circuitry

into a compact and manufacturable design. The final printed circuit board was shown in Figure. 3.7.

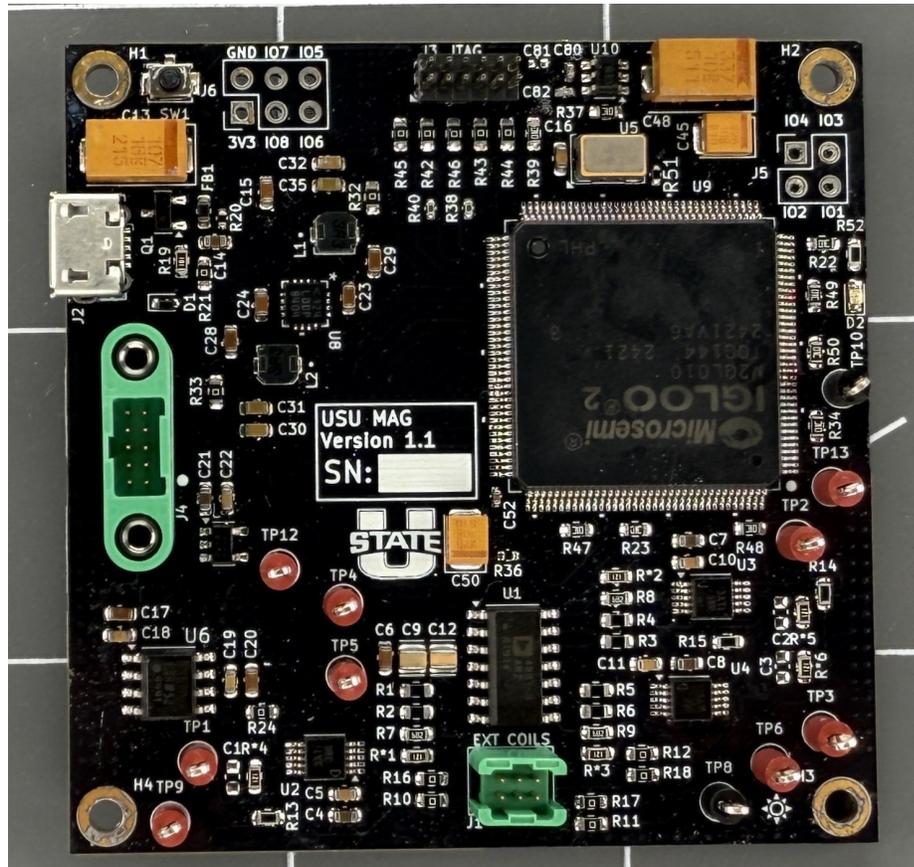


Fig. 3.7: Prototype magnetometer electronics board incorporating the IGLOO2 FPGA, analog oscillator circuitry, polarity-control switches, and external sensing-head interface.

The complete magnetometer electronics were implemented on a compact 60 mm × 60 mm six-layer PCB using only commercial off-the-shelf components. The final design operated at approximately 115 mW during continuous three-axis sampling, making it suitable for power-constrained CubeSat platforms. Powered by a DP832A (RIGOL) power supply, the power consumption is shown in Figure 3.8.



Fig. 3.8: Power consumption with magnetometer in active sample.

The next chapter describes the FPGA firmware that controls the oscillator, manages timing, and processes each magnetic measurement.

CHAPTER 4

FIRMWARE DESIGN

4.1 Firmware Development

The control and signal-processing functions were implemented on an IGLOO2 FPGA (MicroChip), which executed the sequencing and measurement logic required for the magnetometer. Firmware was developed in MATLAB Simulink using Hardware Description Language (HDL) Coder for automated VHSIC Hardware Description Language (VHDL) generation. The block-based workflow allowed rapid iteration and early functional testing while remaining compatible with the FPGA toolchain [16].

Automated HDL generation simplified the creation of timing control, state machines, and data-handling blocks while still allowing manual refinement after synthesis. Each functional unit, such as timing counters, supervisory logic, and serial interfaces, was tested individually in simulation before being integrated at the top level of the firmware. Although this approach accelerated development, it required careful verification of clock-domain boundaries and signal timing to ensure that the generated logic met performance constraints of the FPGA fabric.

One block in the firmware, the oscillation counter, was implemented entirely in hand-written VHDL instead of being generated from Simulink. This counter measured the period of the Schmitt-trigger output waveform by incrementing on each rising edge of the sensor oscillator, which acted as an external asynchronous clock. Because Simulink's fixed-step simulation environment could not accurately model asynchronous event-driven behavior, this component was developed manually to achieve precise and deterministic timing. It was then integrated with the Simulink-generated control logic as a stand-alone module within the top-level design. The Oscillation Counter VHDL is shown in Appendix [B.1](#).

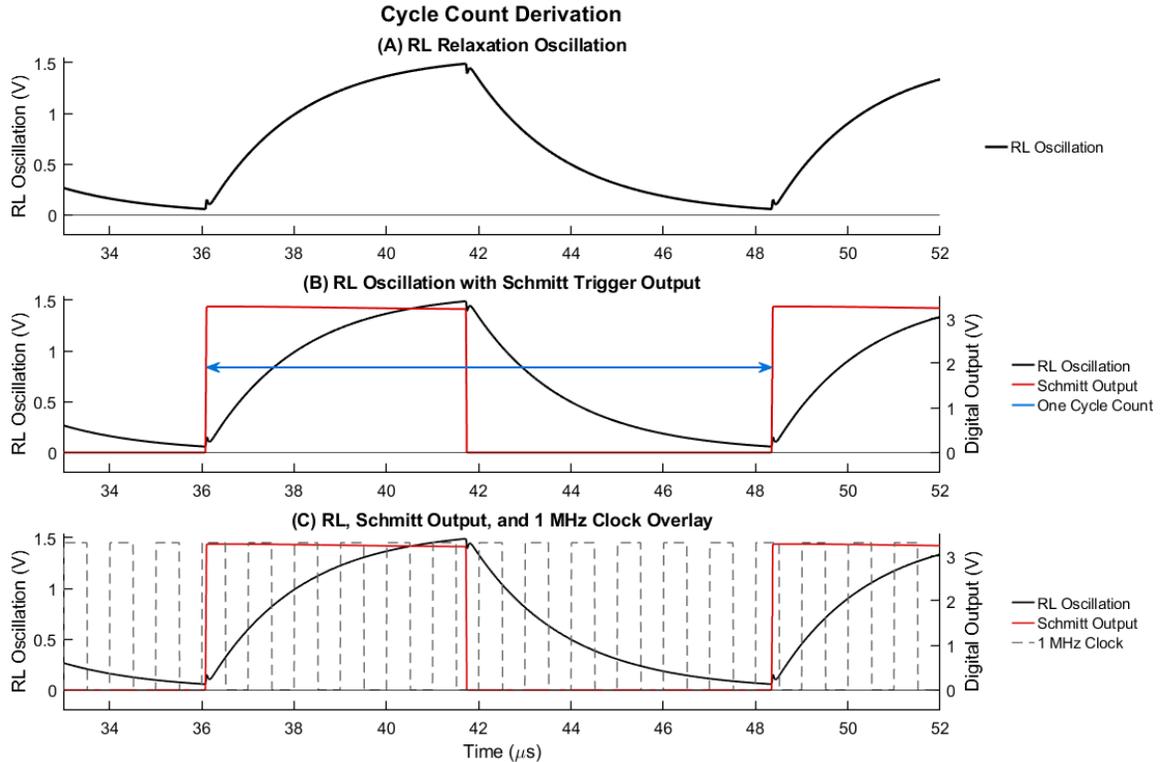


Fig. 4.1: Cycle count definition. (A) LR relaxation oscillation; (B) Schmitt-trigger square-wave output with one full cycle highlighted; (C) overlay of the 1 MHz (example) FPGA clock used for timing. A single *cycle count* corresponded to one complete charge–discharge period between two rising-edge transitions of the Schmitt output.

Figure 4.1 illustrates how a single *cycle count* was defined. The LR oscillator produced a repeating charge–discharge waveform, which the Schmitt trigger converted into a clean digital signal. In the visualization, a 1 MHz clock was overlaid to show how clock edges accumulated within one oscillation period. This lower-frequency clock was used only for clarity in the figure.

In actual operation, the FPGA measured the period using a 380 MHz high-speed counter (2.63 ns per tick). A cycle was defined as the interval between two consecutive rising edges of the Schmitt-trigger output. The counter incremented once per fast-clock tick during this interval, and the total number of increments between rising edges formed the cycle count, essentially a nanosecond-resolution measurement of the oscillator period.

This measurement period was the fundamental unit used by the magnetometer. Multiple cycle counts (e.g., 100–40,000) were accumulated during a measurement to improve noise performance. Increasing the number of cycles increased effective integration time, but also made the system more sensitive to oscillator drift and low-frequency noise. This was later examined in the calibration and noise-floor analysis.

The firmware was organized around a top-level state machine shown in Figure 4.2. The simulink model of the firmware that governed axis sequencing, polarity switching, and data flow is shown in Appendix C.1, The state machine driving the axis sequencer is shown in Appendix C.2.

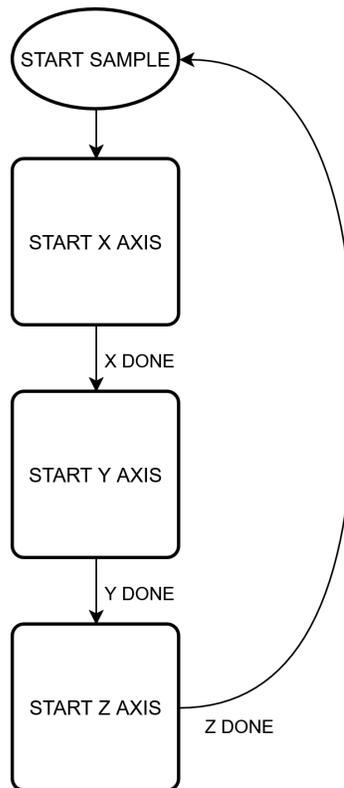


Fig. 4.2: Top-level firmware flowchart.

4.1.1 Timing Optimization and Clock-Domain Management

Logic element placement and signal propagation timing became a challenge during firmware implementation. The high-speed timing counter initially targeted a 400 MHz clock

to maximize resolution, but this frequency pushed the FPGA near its physical limit. This led to routing congestion and negative slack during place-and-route. Minor design changes often altered routing enough to reintroduce violations, making compilation unpredictable.

To improve reliability, the clock was reduced to 380 MHz, and key logic paths were optimized. Analysis of the critical path revealed that a subtraction between two 32-bit timestamp registers, representing positive and negative oscillation periods, was performed in the high-speed domain. Because both timestamp registers were buffered by a latching system immediately after being recorded, this arithmetic could safely be performed at a lower clock rate. Moving the subtraction to the 5 MHz control domain removed the long combinational path and resolved all remaining timing violations. After this change, the design met timing reliably at 380 MHz, and further increases were unnecessary for this sensor revision.

The resulting dual-clock architecture divided tasks cleanly between domains. The high-speed counter only performed timestamp capture while arithmetic, sequencing, and data packaging occurred in the slower control logic. This separation preserved timing precision and simplified synchronization between domains.

4.1.2 State Machines and System Coordination

The firmware architecture included three per-axis state machines that controlled excitation, timing, and measurement for the X, Y, and Z coils. Each state machine managed op-amp enable signals, polarity switching, and pre- and post-measurement delays to maintain stable operation. As shown in Figure 4.3, these state machines coordinated directly with the raw HDL oscillation counters which were the most timing-critical components of the system. The simulink model of the axis driver is shown in Appendix C.3, and the state machine controlling the axis driver is shown in Appendix C.4.

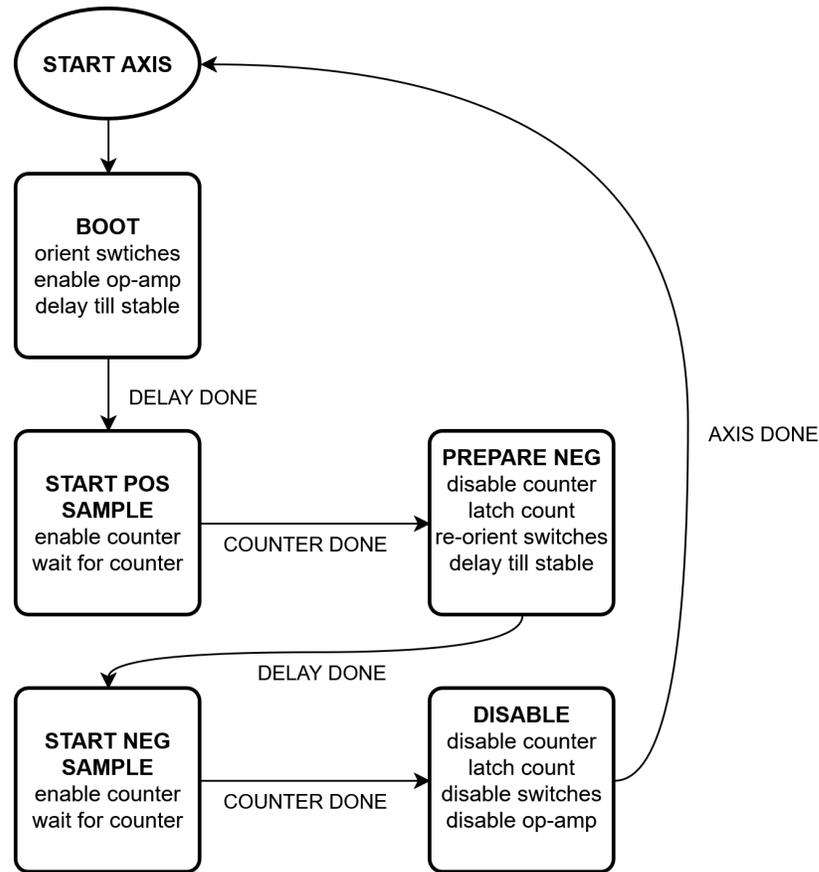


Fig. 4.3: Per-axis driver state machine controlling excitation, timing, and synchronization with the high-speed counter.

4.1.3 Polarity Switching, Measurement Sequencing, and Counter Coordination

Each axis driver coordinated timing handshakes with the oscillation counter and controlled the polarity of current through the sensor coil. Polarity reversal enabled differential magnetic-field measurement, canceling offset and thermal drift in the analog front end.

The driver followed a Moore style state machine alternating between positive and negative excitation. During each phase, the driver set control lines SW1_CONTROL and SW2_CONTROL to define current direction and asserted AMP_EN to enable the Schmitt trigger amplifier. After a short stabilization delay (delaying=1 until delay_done=1), the driver issued a single-cycle start pulse and asserted count_enable to begin measurement.

While the oscillation counter counted rising edges of the Schmitt-trigger waveform, the counting flag remained high. When the programmed `cycle_counts` value was reached, the counter asserted `counter_done`, prompting the driver to record the timing value into either the positive or negative timer register. The driver then disabled the amplifier, deasserted polarity control, and initiated another delay before advancing to the opposite phase. Once both polarities were measured, a differencing block computed the period delta proportional to the sensed magnetic field.

The axis driver connected the analog oscillator, the event-clocked oscillation counter, and the high-speed timing counter. Operating in the 5 MHz control domain, it managed amplifier enables, polarity switching, and measurement sequencing. The associated timing counter ran at 380 MHz, while the raw oscillation counter was clocked directly by the Schmitt-trigger output. This partitioned the measurement across three timing references: the physical oscillator for edge counting, the high-speed counter for duration measurement, and the control clock for sequencing.

During each measurement interval, the oscillation counter incremented once per rising edge until `cycle_counts` was reached, at which point it asserted `counter_done`. In parallel, the high-speed counter accumulated FPGA clock cycles to measure duration with nanosecond resolution. Handshake signals (`start`, `count_en`, `counting`, `counter_done`, `delaying`, `delay_done`) synchronized communication between domains.

When `counter_done` was detected, the driver latched the positive or negative timing value (`POS_TIMER` or `NEG_TIMER`), disabled the current polarity, applied a short post-measurement delay, and then moved to the next polarity or next axis. A supervisory controller ensured that only one axis was active at a time to prevent magnetic coupling between coils.

This multi-clock structure maintained clear domain isolation and deterministic sequencing: the oscillation counter provided cycle-accurate edge counting linked to the sensor output, the high-speed counter supplied nanosecond-level timing resolution and the axis driver synchronized and sequenced all events across domains.

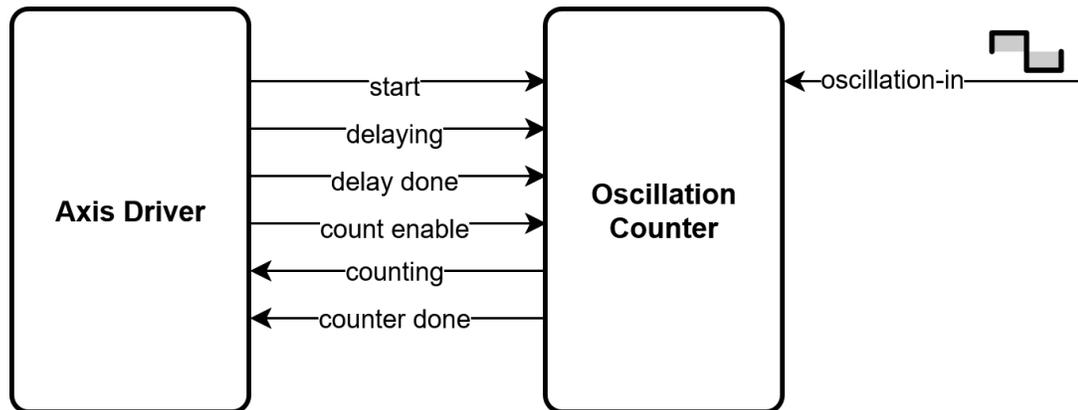


Fig. 4.4: Handshake and polarity coordination between the axis driver and oscillation counter. Signals counting and delaying acted as busy indicators during measurement and delay intervals.

4.1.4 Communication and Control

The magnetometer communicated through a Serial Peripheral Interface (SPI) interface with an additional Data-Ready line that pulsed high when a sample completed, signaling to the master that data was available. The Data-Ready logic, shown in Figure 4.5, was a simple state machine that monitored both data availability and chip-select activity to initiate each transfer. Appendix C.5 shows the full Data-Ready simulink model.

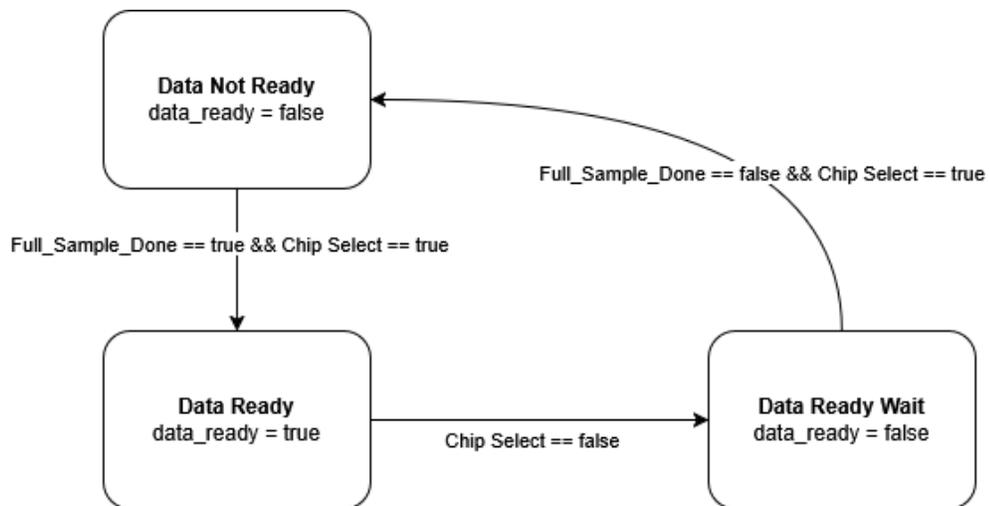


Fig. 4.5: Data-Ready control logic.

Including a Data-Ready-Wait state ensured that each dataset was sent only once. Before this addition, duplicate samples occasionally occurred when chip-select toggled multiple times in quick succession, most noticeable at lower cycle-count settings. Outside the state machine, the Data-Ready register (Figure 4.6) latched and concatenated each axis sample into a 96-bit **Data_To_Send** word for SPI transfer. Appendix C.6 shows the latching state machine.

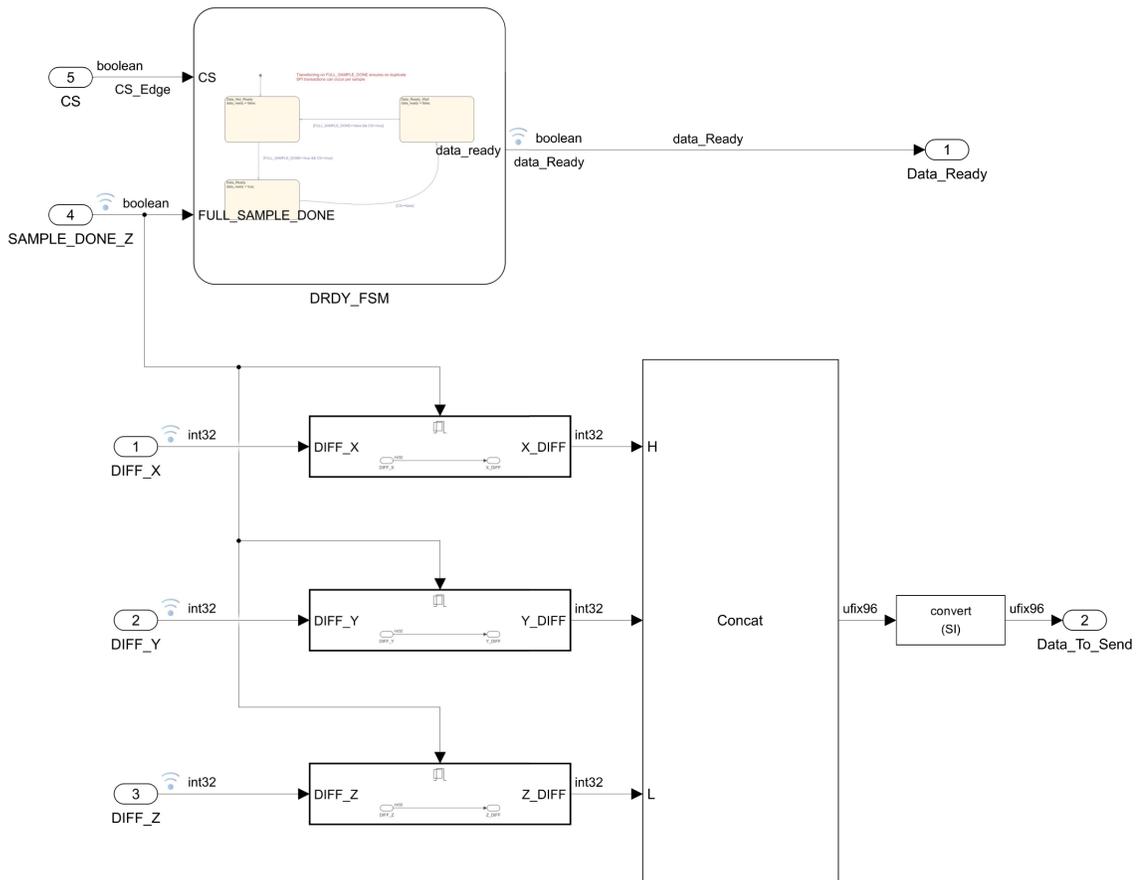


Fig. 4.6: Data-to-Send latching and concatenation logic.

Each SPI transaction transmitted 96 bits, 32 bits per axis sample. During readout, the master could also send configuration commands to adjust the cycle-count or delay settings, allowing on-the-fly reprogramming without downtime. The Data-Ready line was particularly useful because the reporting rate varied slightly with magnetic-field strength

rather than remaining fixed. Appendix C.7 shows the SPI simulink model and Appendix C.8 shows the SPI state machine. Overall, the firmware integrated high-speed timing capture, synchronized control logic, and configurable communication into a single, reliable system. The hybrid approach combining Simulink-generated HDL with custom VHDL provided both flexibility and precise timing control. By separating tasks across clock domains and refining critical logic paths, the design achieved stable operation at 380 MHz with consistent timing margins. This implementation formed a robust foundation for future versions of the magnetometer, supporting continued improvements in precision, stability, and flight integration.

The firmware described in this chapter established the digital foundation for the magnetometer, enabling precise timing capture, coordinated sequencing, and reliable data transfer. With the control and measurement architecture fully implemented in hardware, the mechanical design of the sensor head and supporting structures was next. The following chapter discusses the mechanical integration of the system, including the mounting strategy, enclosure materials, and methods used to minimize magnetic interference and maintain physical alignment between sensor axes.

CHAPTER 5

MECHANICAL INTEGRATION

5.1 Boom-Mounted Magnetometer Assembly

Space-based magnetometers operate in an electromagnetically complex environment where digital electronics, solar panels, and current-carrying structures generate fields that can mask the ambient geomagnetic signal. To minimize these effects, the magnetometer was mounted on a deployable boom that physically separated the sensing coils from the spacecraft bus and primary power electronics. This configuration allowed the instrument to measure the magnetic field with reduced interference while maintaining mechanical simplicity and low mass.

The boom assembly used heritage hinge hardware from the *Scintillation Prediction Observations Research Task* (SPORT) mission, originally developed at the Center For Space Engineering (CSE). A spring-loaded aluminum hinge interfaced with a lightweight carbon fiber boom that positioned the magnetometer approximately 15 cm from the spacecraft body. The carbon fiber structure provided high stiffness, low mass, and negligible magnetic susceptibility, making it ideal for magnetically sensitive instruments.

In addition to magnetic isolation, all structural and enclosure materials were selected for compatibility within vacuum and thermal environments. Both Delrin (polyoxymethylene) and woven Fiberglass (FR4) exhibited low outgassing and high dimensional stability, ensuring reliable operation under space-grade environmental conditions. The design achieved sufficient stiffness to withstand launch vibration and thermal cycling, while maintaining nominal coil alignment, as defined by the enclosure geometry. Although quantitative verification of orthogonality had not yet been performed at the time of this work, visual inspection and consistent calibration behavior suggested that any residual misalignment was small.

The magnetometer was located near the distal end of the boom, providing spatial separation from spacecraft electronics and enabling co-located measurements of plasma current and magnetic field fluctuations. This configuration allowed correlation of magnetic and plasma phenomena without additional deployable mechanisms.

5.2 Sensor Head Enclosure Design Evolution

The magnetometer's sensing head contained three orthogonally mounted inductors and their associated electrical connections. Its mechanical design evolved through two major iterations: a precision-machined Delrin enclosure and a more manufacturable FR4-based version. Each design balanced magnetic cleanliness, alignment accuracy, mechanical durability, and ease of assembly.

5.2.1 First Generation: Machined Delrin Enclosure

The first prototype used a custom-machined Delrin enclosure clamped directly to the carbon fiber boom. (Figure 5.1) Internal precision-machined channels were designed to hold each inductor in orthogonal alignment. The inductors were inserted into these channels and retained by epoxy. These channels were protected from damage by thin Delrin faceplates that were glued to the main Delrin structure. Delrin offered dimensional stability, low magnetic susceptibility, and very low outgassing, making it a suitable candidate for space qualification [17].

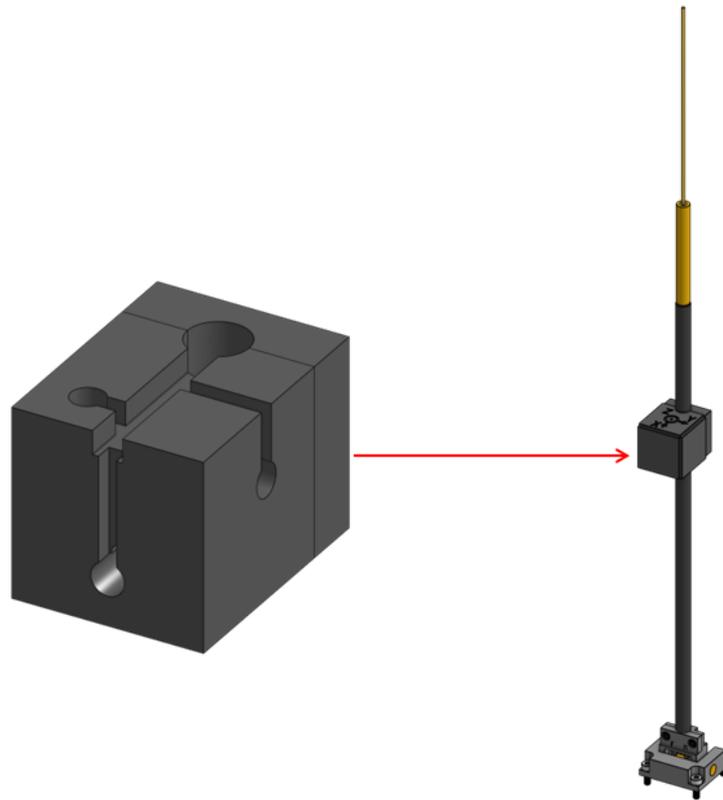


Fig. 5.1: Delrin boom enclosure with precision-machined internal coil channels and Langmuir Probe attached to boom.

This configuration was first implemented on the Defense Advanced Research Projects Agency (DARPA) *Kraken* spacecraft and performed successfully during laboratory and environmental testing. However, Delrin's low surface energy made the faceplate adhesive bonding difficult, requiring sanding and specialized epoxies. Therefore a new variant was developed using FR4. The multi-axis machining process also increased fabrication time and cost. Manual coil insertion carried a risk of damaging fine-gauge leads, and sealing required multiple steps since each coil channel needed to be covered by a separate faceplate.

The complete Delrin-based boom assembly from the initial prototype had a measured mass of 23.5 g with boom included, while the corresponding Delrin sensing head weighed

6 g and occupied $16 \text{ mm} \times 16 \text{ mm} \times 30.5 \text{ mm}$. The subsequent FR4 sensor head design occupied a smaller $16 \text{ mm} \times 16 \text{ mm} \times 13.5 \text{ mm}$ space and lower total mass of 2.8 g per head, 14.8 g per boom. Despite the limitations, the Delrin prototype validated the boom-mounted concept and provided a reference geometry for subsequent designs. However future versions will only use the FR4 variant because of its ease of manufacturing and low mass.

5.2.2 Second Generation: FR4 PCB-Based Enclosure

Building on lessons from the *Kraken* prototype, the sensing head was redesigned around an FR4 enclosure containing a small PCB that mechanically and electrically supported the three inductors. This approach simplified manufacturing and assembly while retaining low magnetic susceptibility and sufficient structural rigidity.

Each inductor was surface-mounted to the PCB which routed the LR circuit drive and sense lines to one of two board variants shown in Figure 5.2.

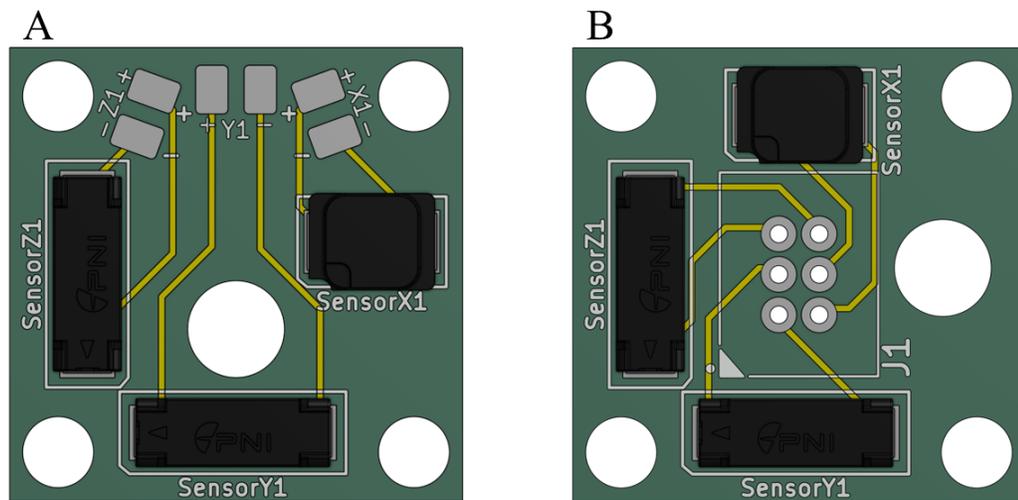


Fig. 5.2: Sensing head PCBs with two cabling configurations: (A) coaxial harness and (B) Gecko connector.

Variant A included a cutout with bare solder pads for direct wire attachment, enabling

the use of a coaxial cable harness and minimizing connector size and mass. Variant *B* employed a Harwin Gecko connector, allowing a direct 1:1 connection to the main electronics board for rapid prototyping and modular integration. This design eliminated the need for a custom cabling harness, reducing manufacturing time and cost. The PCB geometry defined the X-, Y-, and Z-axis orientations: the X and Y coils were mounted in-plane, and the Z coil was oriented perpendicular to the PCB surface.

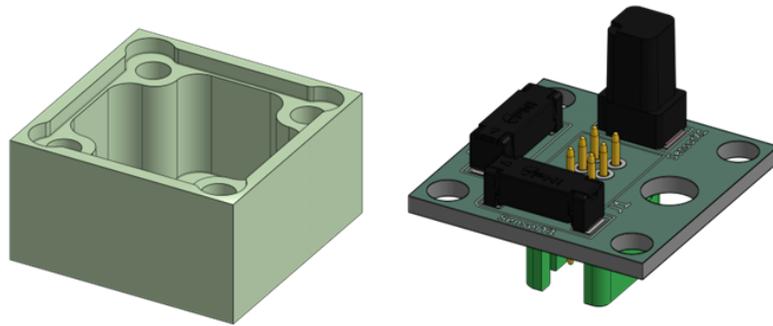


Fig. 5.3: FR4 enclosure alongside sensor head PCB with surface-mounted inductors before assembly.

The FR4 housing shown in Figure 5.4 enclosed the PCB and provided environmental protection. All metallic fasteners were aluminum to avoid ferromagnetic contamination.

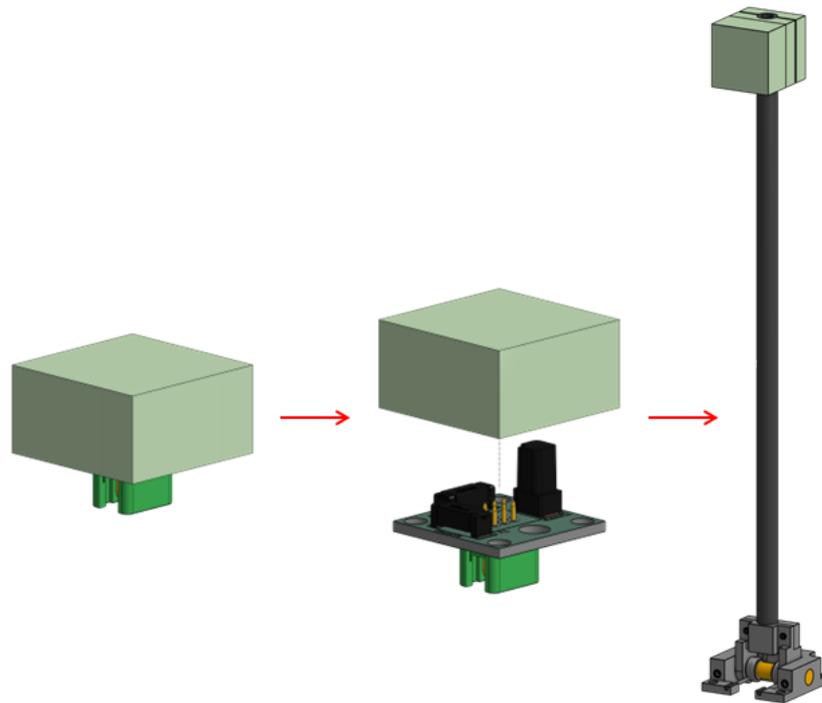


Fig. 5.4: FR4 boom-mounted magnetometer enclosure without Langmuir Probe.

This version reduced assembly time by more than half and improved unit-to-unit consistency. Minor angular deviations could occur between coils due to solder surface tension during reflow which could slightly reduce orthogonality compared to the machined Delrin version. However, these deviations remained small and within calibration compensation limits.

Future iterations are planned to incorporate a precision reflow stencil or mechanical alignment jig to preserve orthogonality while maintaining the simplicity and repeatability of the FR4 design.

CHAPTER 6

CALIBRATION

6.1 Calibration

To convert Schmitt trigger cycle counts into magnetic field values (in tesla), the magnetometer underwent a detailed calibration process using a three-axis Helmholtz coil (HHC) [18]. When driven with a known current, this coil pair generated a well-characterized, spatially uniform magnetic field based on the coil geometry and applied current.

A series of magnetic field strengths were generated by varying the current through the coil in known steps. For each step, the magnetometer recorded the differential cycle counts between the two current directions. A linear regression was then performed on the resulting dataset to establish a calibration curve mapping the measured cycle difference to the corresponding magnetic field value. The slope of this curve provided the calibration constant (units of tesla per count), while the R^2 value quantified the linearity of the response.

Repeatability and hysteresis were evaluated by sweeping the magnetic field both forward and backward. Any deviation in the measured cycle counts for the same field magnitude in opposite directions was noted as a potential source of systematic offset.

6.1.1 Helmholtz Coil Characterization

Prior to calibrating the magnetometer, the Helmholtz coil itself was characterized to establish the relationship between the digital-to-analog converter (DAC) drive-word and the generated magnetic field. Each coil axis was energized independently while a high-precision reference magnetometer was positioned at the coil center. The DAC-word was incremented from $-150\ \mu\text{T}$ to $150\ \mu\text{T}$ in steps of $25\ \mu\text{T}$, and the resulting field was recorded by the reference sensor.

The reference magnetometer used for these studies was a MEDA FVM-400 Vector

Magnetometer, which was loaned to the Center for Space Engineering through the Space Dynamics Lab [19]. The reference magnetometer had a resolution of 1 nT. A custom positioning plate was fabricated to mount the reference magnetometer within the CSE HHC.

To correct for ambient magnetic fields, a zero-field measurement was taken before each HHC field step. The calibration data were processed in MATLAB using a linear least-squares fit of the form

$$B_{\text{meas}} = G_{\text{coil}} \cdot D_{\text{word}} + B_{\text{offset}} \quad (6.1)$$

where B_{meas} was the field measured by the reference magnetometer (in nT), D_{word} was the applied DAC-word, G_{coil} was the coil gain (in nT per DAC-word), and B_{offset} was the residual offset. The fitted parameters allowed any commanded DAC-word to be translated directly into a magnetic field with known accuracy. Typical fits yielded residuals below 0.1 % of full scale, confirming linearity across the ± 150 μT range.

This calibration was completed in the Engineering-Laboratory sub-basement at Utah State University, which provided added magnetic isolation, eliminating the noise from the electronics in the CSE laboratory environment. Figure 6.1 shows the setup described above.



Fig. 6.1: Helmholtz Coil (HHC) Calibration Setup

Although a dedicated mounting plate was fabricated to position the magnetometer and reference sensor within the coil, the placement was performed manually. Mechanical tolerances in both the plate and the HHC frame allowed slight angular and translational offsets that could change between runs. Even small misalignments or wobble in the coil could introduce minor differences in the local field experienced by each axis, particularly at higher field steps where coil output increased and would cause a more noticeable shift.

6.1.2 Magnetometer Calibration Using the Calibrated HHC

Once the HHC was calibrated, the prototype magnetometer was placed at the coil center and aligned with the reference sensor's position. Using the DAC-to-field mapping, the coil was commanded to produce fields spanning $\pm 150 \mu\text{T}$ in $25 \mu\text{T}$ increments along

each axis. The following sequence was performed for each step:

1. **Zero-field measurement:** The coil was set to zero output, and the magnetometer's raw counts were recorded for all three axes.
2. **Field step:** The target magnetic field was applied, and the new steady-state counts were measured.
3. **Reference magnetometer measurement:** The target magnetic field was applied, and the reference magnetometer reading was recorded to correlate magnetic fields to magnetometer counts.
4. **Return to zero:** The coil was again set to zero to confirm stability and detect any drift.

This alternating zero-field-zero sequence ensured that each measurement pair was referenced to its own baseline, mitigating slow environmental magnetic field changes. The corresponding reference field values from the calibrated coil were then paired with the measured counts to produce the calibration datasets between the two magnetometers.

A linear regression was applied to each axis according to:

$$B_{\text{ref}} = G_{\text{mag}} \cdot (C_{\text{raw}} - C_0) + B_{\text{off}} \quad (6.2)$$

B_{ref} was the reference field (in nT), C_{raw} was the measured count, C_0 was the zero-field count, G_{mag} was the gain (in nT per count), and B_{off} was the fitted offset. The MATLAB processing script reported both the forward and inverse forms, enabling real-time field estimation from the raw counts in firmware:

$$C_{\text{raw}} = C_0 + \frac{B_{\text{ref}} - B_{\text{off}}}{G_{\text{mag}}} \quad (6.3)$$

The fitted R^2 values typically exceeded 0.9999, confirming linear behavior of the sensor across the full ± 150 μT range. All calibration constants (gain, offset, and zero counts) were saved in a structured MATLAB file and exported as integer arrays for use in future work.

Because the sensor’s resolution and noise characteristics were dependent on the Schmitt trigger’s integration time and the number of counted oscillation cycles, the calibration process was repeated across multiple operating points: 100, 1,000, 10,000, 20,000, and 40,000 cycles per measurement. Lower cycle counts yielded higher bandwidth but lower precision. Higher cycle count values improved sensitivity but at the cost of resolution. PNI Sensor recommends operating the RM3100 at one hundred counts, which was used as the low end cycle count for this test. The upper end of 40,000 was selected because the sensor quickly approached a 1 Hz reporting rate which was the maximum sample time allowed for the scope of this project. Each configuration underwent the same coil sweep procedure described above, producing a distinct calibration constant for each mode. This allowed quantitative mapping of sensitivity and noise performance as a function of integration time. Calibration plots are shown in Appendix [D.1–D.15](#).

6.1.3 Reporting Rate versus Integration Depth

Each integration setting defined the total duration of one full oscillation-count sequence, representing the maximum achievable reporting rate. The measured times for 100–40,000 cycle configurations are summarized in Table [6.1](#). For these tests, timing was measured as the duration required for one complete positive–negative measurement pair on a single axis. Because the magnetometer sampled the X, Y, and Z axes sequentially, the overall three-axis vector update rate was approximately one-third of the per-axis value.

Table 6.1: Measured reporting rates as a function of integration depth.

Cycle Count	Duration [μ s]	Total Rate [Hz]	3-Axis Vector Rate [Hz]
100	1199	834	278
1 000	11 692	85.5	28.5
10 000	116 625	8.57	2.86
20 000	233 230	4.29	1.43
40 000	466 400	2.14	0.71

The data showed a clear inverse relationship between integration depth and reporting rate. At low cycle counts (100–1,000), the system delivered rapid updates exceeding 200 Hz per axis, suitable for dynamic attitude determination or compassing applications. Increasing the cycle count to 40,000 extended the measurement duration to approximately 0.47 s per axis, producing a vector update rate near 1 Hz. This quantitative scaling highlighted the trade-off between bandwidth and sensitivity: longer integrations improved effective resolution and noise performance while shorter integrations favored responsiveness.

As expected, the measured gain decreased inversely with the number of counted cycles. This is consistent with the increased integration time per measurement and the corresponding enhancement in effective resolution. This scaling behavior confirmed that the LR relaxation oscillator maintained a consistent proportionality between inductance variation and timing period across multiple integration depths. Note that the Z axis showed a much larger offset due to the noisy electronics within the HHC being positioned directly in its sensing path but are filtered out of the measurement with the zero field baseline.

Table 6.2: Calibration results for all integration-time configurations.

Cycles	Axis	Gain [nT/count]	Offset [nT]	Zero Field Baseline [nT]	R^2
100	X	4.531	-194.510	-1300	0.999
	Y	4.624	28.000	2038	0.999
	Z	4.639	-187.810	9477	0.999
1 000	X	0.452	-14.740	-13 458	0.999
	Y	0.462	115.390	19 898	0.999
	Z	0.462	-344.040	93 119	0.999
10 000	X	0.0452	71.990	-133 443	0.999
	Y	0.0461	-7.280	202 805	0.999
	Z	0.0463	-163.980	931 506	0.999
20 000	X	0.0225	-87.790	-259 220	0.999
	Y	0.0230	65.760	409 341	0.999
	Z	0.0231	-144.100	1 878 703	0.999
40 000	X	0.0113	-26.000	-539 573	0.999
	Y	0.0115	53.290	805 989	0.999
	Z	0.0115	-230.700	3 788 098	0.999

The data in Table 6.2 shows a clear inverse relationship between magnetic field resolution and integration cycle count. Increasing the number of relaxation cycles improved resolution, achieving sub-nanotesla sensitivity at higher counts, and the near-linear scaling between 100, 1,000, and 10,000 counts confirmed a stable and predictable sensor response. This relationship was further illustrated in Figure 6.2, which shows the decrease in sensitivity with increasing cycle count. The counts measured with the HHC powered off were

also recorded as the zero field baseline.

In practice, the choice of cycle count represented a balance between measurement speed and sensitivity. Larger cycle counts provided higher resolution but at the cost of sampling rate. Smaller counts allowed faster updates at the expense of sensitivity. For applications such as compassing or attitude determination, lower counts were preferred for quicker response, whereas scientific measurements benefited from higher counts that delivered finer resolution.

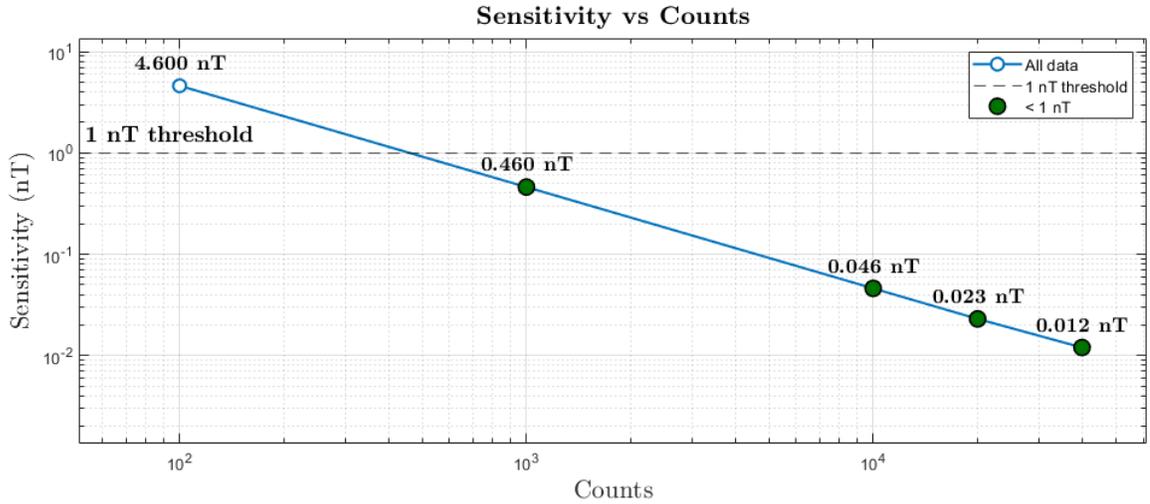


Fig. 6.2: Sensitivity versus cycle counts plotted on a log–log scale.

6.1.4 Residuals Across Integration Depths

To assess calibration quality as a function of integration depth, residuals were computed for each axis and cycle-count configuration using:

$$r = B_{\text{ref}} - \hat{B}_{\text{fit}} \quad (6.4)$$

\hat{B}_{fit} was the linear model prediction. For each dataset, the mean (bias), standard deviation, root-mean-square error (RMSE), and maximum absolute residual were summarized. Because the residual means were effectively zero in all cases (within numerical precision), Table 6.3 focuses on RMSE which captured residual fluctuations and calibration errors. Per-

cent full scale expresses how large the RMSE is compared to the magnetometer’s maximum measurement range $\pm 150 \mu\text{T}$ or 150,000 nT.

Table 6.3: Residual RMSE versus cycle count for each axis. Values in parentheses show percent of full scale.

Cycles	X RMSE [nT] (%FS)	Y RMSE [nT] (%FS)	Z RMSE [nT] (%FS)
100	136.247 (0.091%)	261.252 (0.174%)	112.426 (0.075%)
1 000	95.130 (0.063%)	161.318 (0.108%)	175.021 (0.117%)
10 000	70.654 (0.047%)	93.699 (0.062%)	143.314 (0.096%)
20 000	87.051 (0.058%)	127.042 (0.085%)	100.436 (0.067%)
40 000	80.330 (0.054%)	103.827 (0.069%)	89.311 (0.060%)

The residual means were effectively zero for all datasets, indicating that the fits were unbiased. RMSE decreased substantially between 100 and 10,000 cycles and then improved more gradually out to 40,000 cycles which is consistent with noise averaging at longer integration times. Minor non-monotonic steps in the trend (such as the increase between 10,000 and 20,000 cycles) are normal and reflect routine experimental variation rather than a change in sensor behavior. At higher cycle counts, each measurement spans a longer time window, making the fits more sensitive to small effects like coil-current drift, slight mounting shifts, or the influence of a single outlier in a smaller dataset. Even so, all RMSE values remained below about 0.20 % of full scale, with the best near 0.06 %, confirming that a first-order linear model was adequate across all operating modes. The Y axis showed the largest residuals at low cycle counts and remained the noisiest overall, consistent with the minor alignment issues noted during testing.

The corresponding standard deviations (σ) and worst-case deviations ($\max|r|$) for each axis and cycle-count configuration are summarized in Table 6.4.

Table 6.4: Standard deviation (σ) and maximum absolute residual ($\max|r|$) for each integration depth and axis. All values in nT.

Cycles	X σ	Y σ	Z σ	X $\max r $	Y $\max r $	Z $\max r $
100	144.5	279.3	119.2	210	468	222
1 000	100.9	172.5	185.6	175	217	287
10 000	74.9	100.2	152.0	142	152	228
20 000	92.3	135.8	106.5	179	189	149
40 000	85.2	111.0	94.7	151	169	189

These trends supported the expected trade-off: longer cycle counts yielded lower residual scatter (better effective resolution) at the cost of reporting rate. Axis-to-axis differences (e.g., larger Y residuals) were consistent with small alignment tolerances, local electromagnetic interference, or axis-dependent coupling to the coil electronics.

6.2 Noise Floor Analysis

6.2.1 Magnetic Shielding Chamber

Noise-floor measurements were performed using a three-layer μ -metal magnetic shielding chamber to establish a stable, low-field environment for the sensing head. Figure 6.3 shows the μ -metal chamber with the magnetometer inside. Each shield layer is a high-permeability μ -metal cylinder separated by non-magnetic dielectric spacers. The high static permeability of μ -metal creates a low-reluctance path for external magnetic flux, causing incident fields to be redirected around the interior volume rather than penetrating through it. When multiple layers are used, the residual field transmitted through each successive shell decreases exponentially, leading to substantial improvement in low-frequency shielding performance [20].

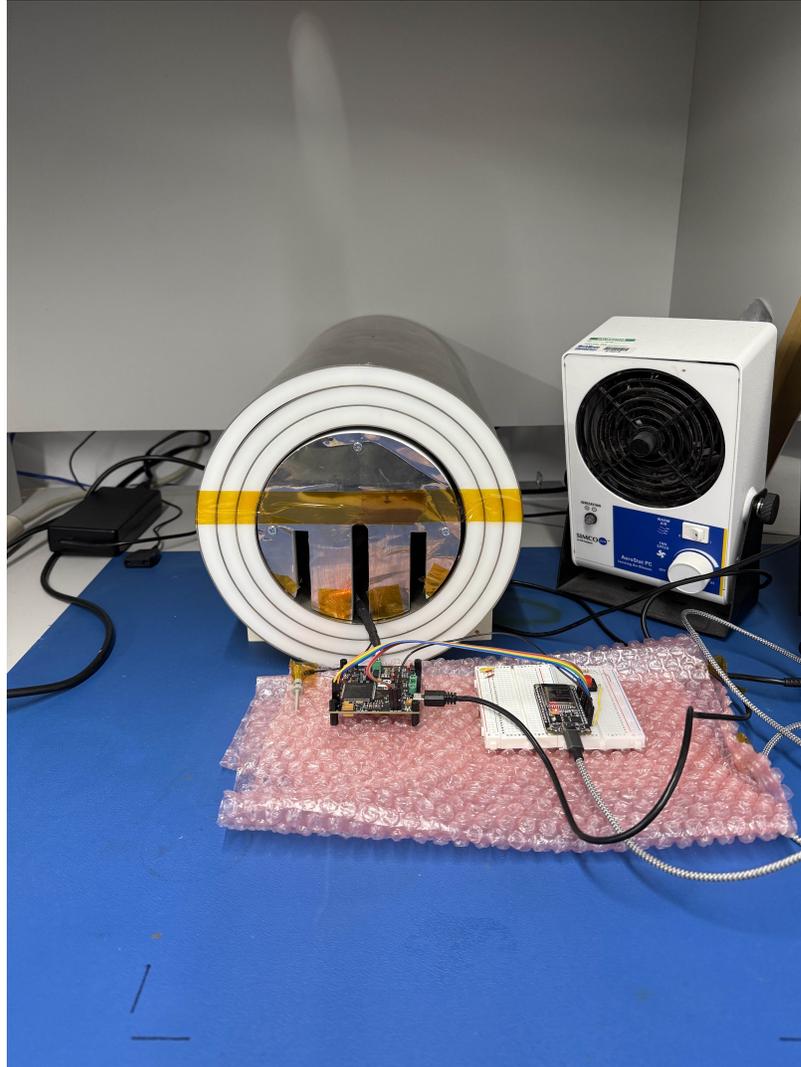


Fig. 6.3: Three-layer μ -metal magnetic shielding chamber used for noise testing with magnetometer attached. The magnetometer sensing head is located inside the center of the coil.

Magnetic shielding of this type is most effective for static and low frequency varying magnetic fields where the dominant attenuation is from the large Direct Current (DC) and quasi-static permeability. As frequency increases, eddy-current losses within the conductive μ -metal walls provide additional attenuation through induced opposing fields. This combination of high-permeability and eddy-current effects enables multilayer μ -metal enclosures to suppress geomagnetic fluctuations, building wiring fields, and environmental 50–60 Hz disturbances to levels that are effectively negligible inside the inner region of the chamber.

Only the magnetometer sensing head was placed inside the chamber during testing. All FPGA hardware, switching circuitry, and other electronics remained outside the shielded volume to prevent introducing magnetic artifacts or thermal fluctuations. The sensing head was positioned on a non-magnetic support near the geometric center of the innermost layer, and the cable bundle was routed through a shielded feedthrough to minimize leakage and preserve the shielding effectiveness.

With the chamber closed, the internal magnetic field is effectively static over the timescale of a noise-floor measurement. Any fluctuations observed in the recorded data therefore originate from the magnetometer system itself rather than from external environmental variations. This controlled low-field environment provides the appropriate reference for evaluating the intrinsic noise characteristics of the sensing head.

6.2.2 Noise Floor Characterization

The intrinsic noise characteristics of the magnetometer were evaluated for each integration-depth configuration following calibration. Measurements were performed inside the three-layer μ -metal chamber described previously, providing a magnetically quiet environment for assessing baseline performance across the 100–40,000 cycle range. Within this shielded environment, external magnetic variations were suppressed to the point that any remaining fluctuations in the data reflect the behavior of the sensor system itself rather than changes in the surrounding field.

Each dataset was converted from raw Schmitt-trigger cycle counts to magnetic field units (nT) using the calibration constants listed in Table 6.2. The resulting time series were analyzed in MATLAB to determine both time-domain variation and spectral noise density.

A one-sided power spectral density (PSD) was computed for each axis using Welch’s method with a Hanning window and 50 % overlap, from which the amplitude spectral density (ASD, $\text{nT}/\sqrt{\text{Hz}}$) was derived. Because the reporting rate depended on the number of counted oscillations, the valid spectral range varied with integration depth. To ensure consistent comparison, the mid-band noise floor was evaluated over a frequency band defined

individually for each dataset as:

$$[f_{lo}, f_{hi}] = [\max(0.05, 5\Delta f), 0.4F_s], \quad (6.5)$$

Δf was the spectral bin width, and F_s was the sampling frequency for that configuration. A second-order Butterworth high-pass filter with a 0.05 Hz cutoff was applied to remove slow baseline magnetic field drift prior to PSD estimation followed by a linear detrend to suppress any remaining bias.

Because the integration depth directly determined the sampling frequency, each configuration exhibited a different Nyquist limit and therefore a distinct valid spectral range. Shorter integrations yielded higher reporting rates and correspondingly broader spectra, whereas longer integrations provided slower sampling and restricted the analysis to lower frequencies. The frequency bounds listed in Table 6.5 were selected to isolate the flat, white-noise region of each spectrum, excluding low-frequency drift and high-frequency roll-off near the Nyquist limit edge.

Table 6.5: Median mid-band amplitude spectral density (ASD) for each integration-depth (cycle counts) configuration. Values represent the median ASD within the flat region of each spectrum (see Figs. 6.4–6.8).

Cycle Counts	F_s (Hz)	Band (Hz)	X (nT/ $\sqrt{\text{Hz}}$)	Y (nT/ $\sqrt{\text{Hz}}$)	Z (nT/ $\sqrt{\text{Hz}}$)
100	836.12	2–100	0.741	0.831	0.543
1 000	85.68	2–38.6	6.81	6.71	6.88
10 000	8.59	2–3.87	15.0	14.6	14.8
20 000	4.30	0.2–1.93	21.1	20.2	20.7
40 000	2.15	0.2–0.97	27.5	26.6	27.8

The 100-cycle configuration achieved the lowest overall mid-band noise, reaching sub-nanotesla levels across all three axes ($\sim 0.7\text{--}0.8$ nT/ $\sqrt{\text{Hz}}$) within the 2–100 Hz band. The 1,000-cycle mode showed a slightly higher ASD ($\sim 6\text{--}7$ nT/ $\sqrt{\text{Hz}}$) but maintained spec-

tral flatness, corresponding to the intrinsic electronic noise of the comparator and counter chain. At deeper integrations (10,000–40,000 cycles), the measured ASD increased rather than continuing to decrease. This behavior arises because long integrations increase the measurement interval to the point where slow variations in the oscillator period become visible in the sample-to-sample differences. Small timing fluctuations caused by noise in the LR waveform and finite clock jitter average down effectively at short integration depths. At long depths these effects become correlated over the duration of a single measurement. Combined with the reduced reporting rate and corresponding collapse of the Nyquist frequency, this causes slow drift inside the shielded chamber to fall within the analyzable spectrum, elevating the apparent mid-band noise. The higher ASD values at 20,000–40,000 cycles therefore reflect the increasing influence of the magnetometer’s intrinsic white-noise performance.

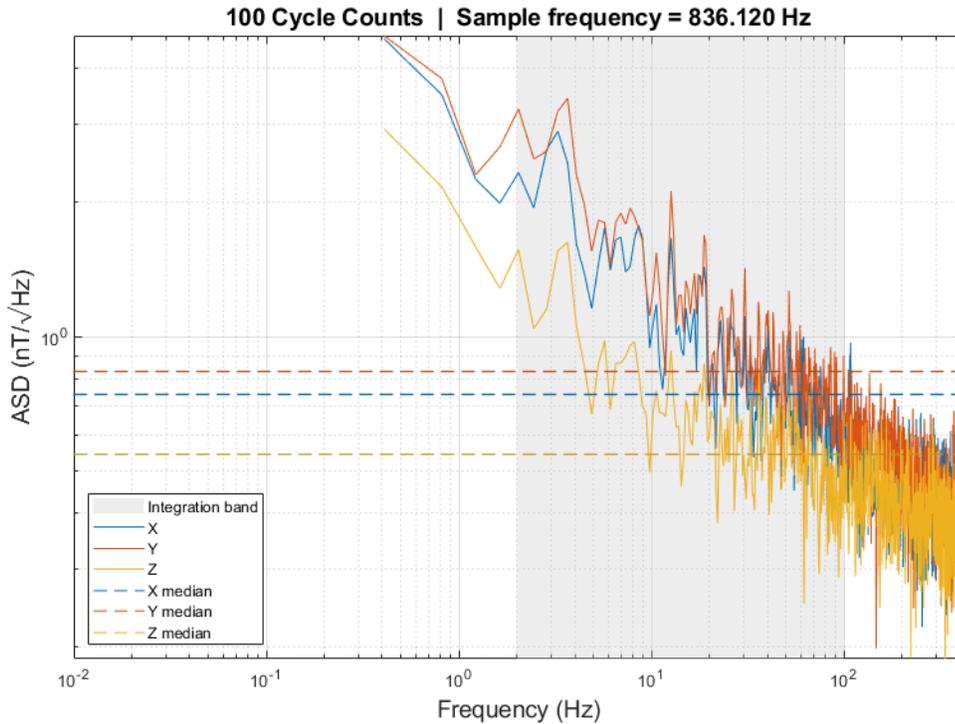


Fig. 6.4: ASD at **100-cycle** integration. One-sided ASD for X, Y, Z; $F_s \approx 836$ Hz with Nyquist frequency $f_N \approx 418$ Hz. The shaded region denotes the mid-band interval used to compute the median ASD values reported in Table 6.5.

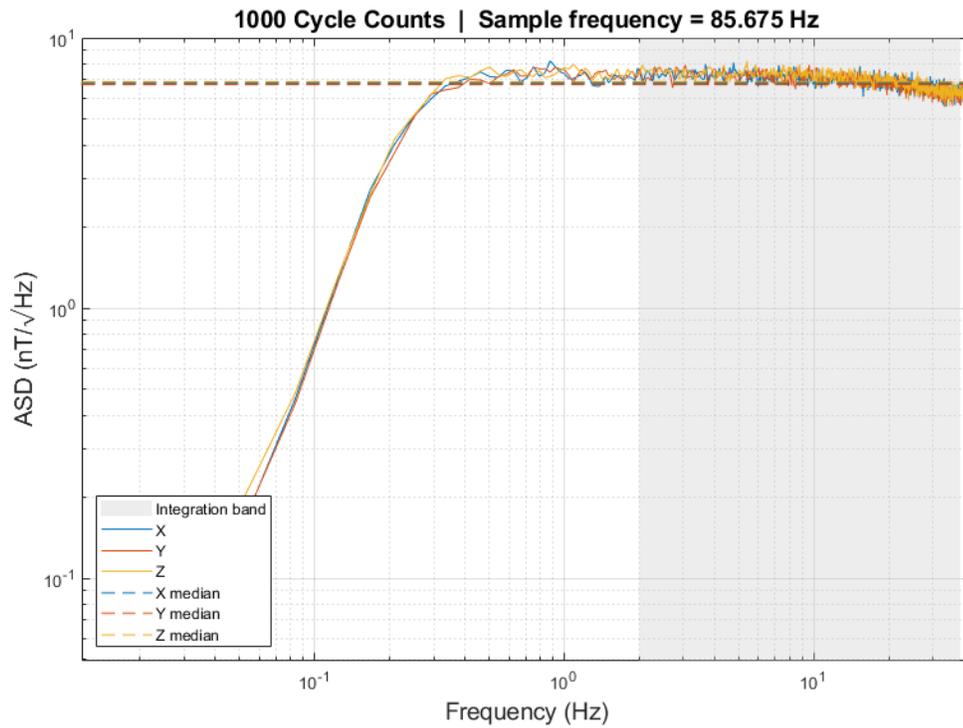


Fig. 6.5: ASD at **1,000-cycle** integration. One-sided ASD for X, Y, Z; $F_s \approx 85.7$ Hz ($f_N \approx 42.8$ Hz). The shaded region indicated the frequency band over which the median ASD was computed (see Table 6.5).

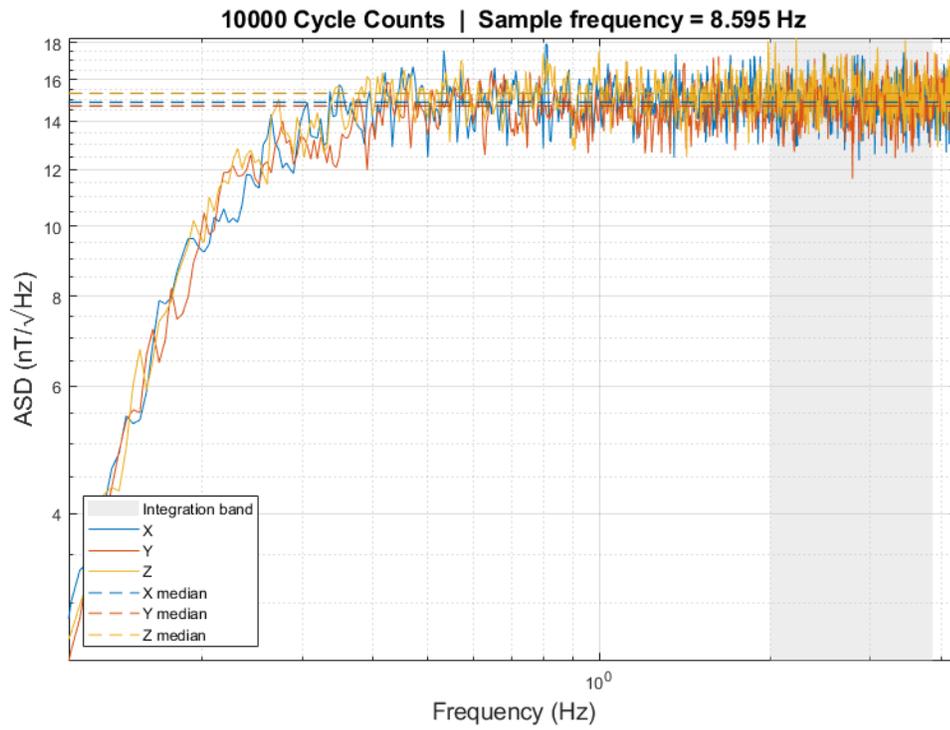


Fig. 6.6: ASD at **10,000-cycle** integration. One-sided ASD for X, Y, Z; $F_s \approx 8.59$ Hz ($f_N \approx 4.30$ Hz). The shaded region marked the mid-band analysis interval used for the median ASD values in Table 6.5.

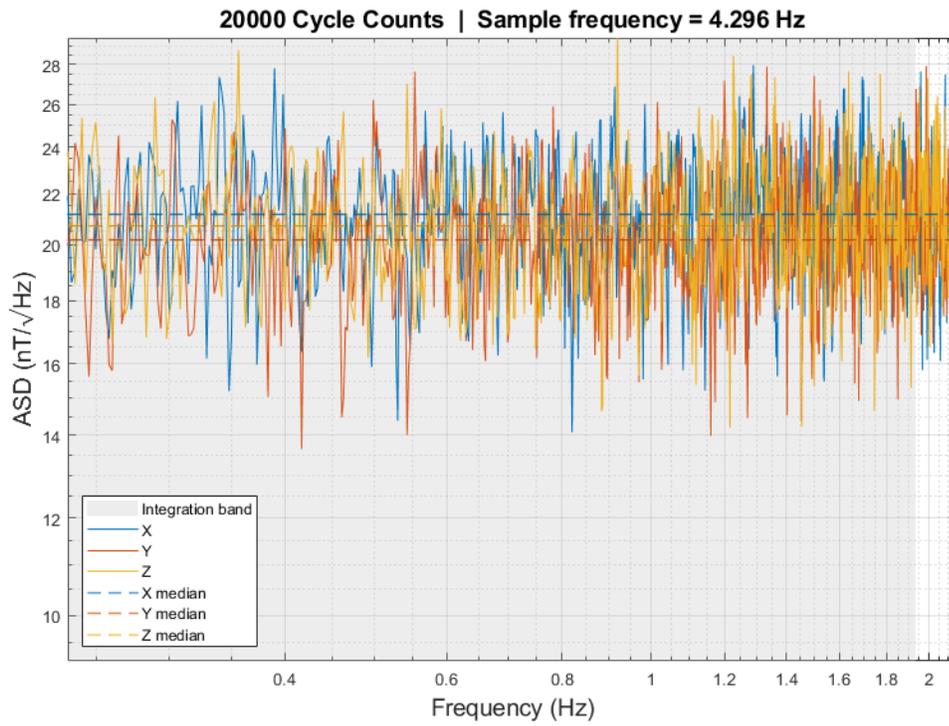


Fig. 6.7: ASD at **20,000-cycle** integration. One-sided ASD for X, Y, Z; $F_s \approx 4.30$ Hz ($f_N \approx 2.15$ Hz). The shaded region corresponded to the mid-band interval used to compute the median noise levels in Table 6.5.

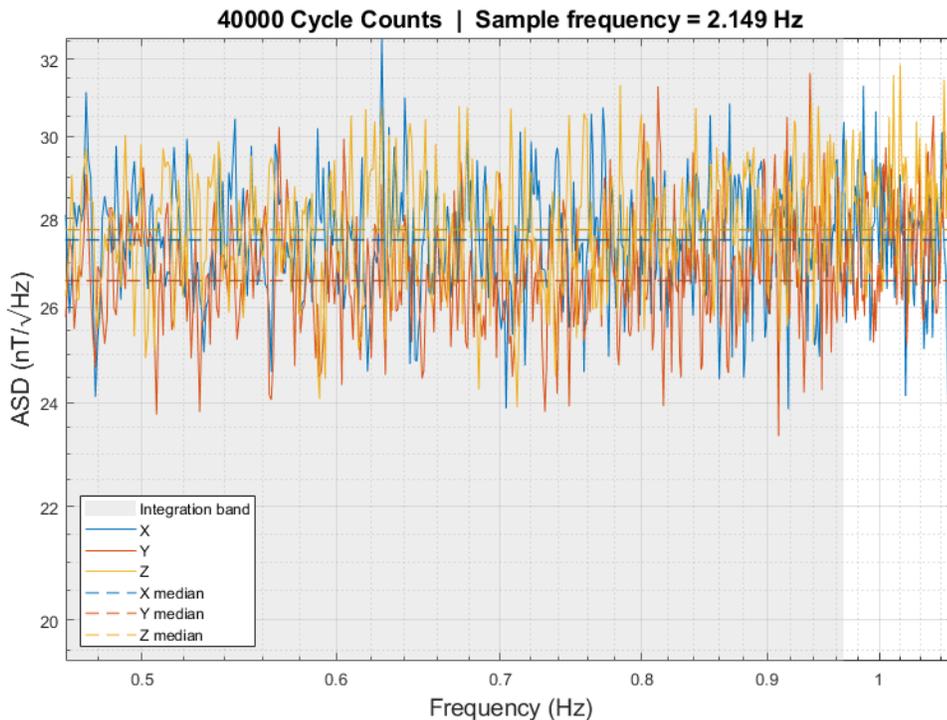


Fig. 6.8: ASD at **40,000-cycle** integration. One-sided ASD for X, Y, Z; $F_s \approx 2.15$ Hz ($f_N \approx 1.07$ Hz). The shaded region represented the mid-band frequency range used for computing the median ASD values tabulated in Table 6.5.

Taken together, these measurements established the baseline noise performance of the prototype magnetometer across all available integration-depth configurations. The ASD plots in Figures 6.4–6.8, along with the median mid-band values summarized in Table 6.5, provided a consistent characterization of the instrument’s intrinsic noise as a function of sampling rate and effective integration time.

CHAPTER 7

DISCUSSION

The results presented in this chapter demonstrate that the magneto-inductive sensor, its supporting electronics, and the boom-mounted configuration were all successfully implemented and operated as intended. The relaxation oscillator, polarity-reversal network, and FPGA-based timing architecture produced stable and repeatable measurements across a wide range of integration depths, while only consuming 115 mW, validating both the circuit design and the firmware responsible for cycle-count timing. The ability to deploy the sensing head at the end of a lightweight, non-magnetic boom, combined with the configurable control system implemented on the IGLOO2 FPGA, establishes a functional platform suitable for spacecraft-grade magnetic field measurements.

At the same time, the noise-floor analysis reveals a non-intuitive behavior in the present design. Increasing the cycle-count integration depth should improve field resolution because the quantization error of the period measurement decreases approximately as $1/N$, where N is the number of oscillation cycles. This expected trend was confirmed by the measured sensitivity curve which fell monotonically as the number of counted oscillations increased. However, the measured mid-band ASD exhibits the opposite trend: the noise floor increases with integration depth, rising from approximately $0.74 \text{ nT}/\sqrt{\text{Hz}}$ at 100 cycles to nearly $28 \text{ nT}/\sqrt{\text{Hz}}$ at 40,000 cycles. This behavior is summarized in Figure 7.1.

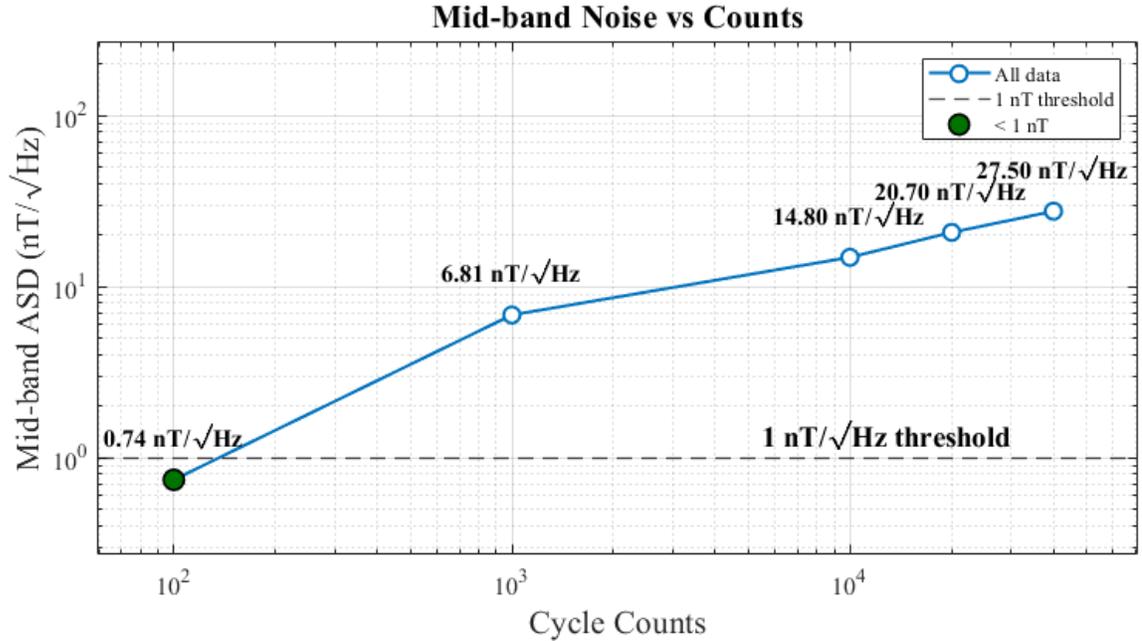


Fig. 7.1: Mid-band ASD versus cycle-count integration depth. Each point represents the median ASD within the shaded mid-band interval of Figs. 6.4–6.8, with numerical values listed in Table 6.5.

This divergence indicates that the instrument is not limited by quantization uncertainty at long integration intervals. Instead, the increasing ASD with cycle count shows that the measurements become masked by slow variations in the oscillator waveform or from off-by-one errors accumulating over time. Because the FPGA measures the total duration of the Schmitt-trigger output, any low-frequency timing jitter or threshold variation in the comparator directly maps into the accumulated cycle count. As the integration interval grows, this slow behavior is not averaged out but rather it is exaggerated, and it appears as an elevated mid-band noise floor.

Several mechanisms contribute to the low-frequency drift. Thermal (Johnson–Nyquist) noise in the circuit changes the LR ramp slope while small fluctuations in the mu-metal core’s inductance introduce slow permeability variations. The Schmitt trigger adds its own effects through threshold jitter, occasional metastability near its switching point, and gradual shifts in its resistor network. At short integration depths these sources are mostly uncorrelated from one oscillation to the next and tend to average out. At longer integrations, the same

variations persist across the entire measurement window, making them appear correlated. When this occurs, the noise falls within the mid-band region of the ASD and raises the apparent noise floor.

Overall, the combined sensitivity, calibration behavior, and noise-floor measurements show that the discrete magneto-inductive architecture developed in this work provides a stable and low-power platform for precise magnetic-field sensing. The sensor achieves a sub-nanotesla noise floor at short integration depths, and the full configurability of the timing, polarity-control, and cycle count register enable operation across a wide range of bandwidths from compassing to scientific grade measurements. Although long-integration modes are limited by the rising noise floor, these effects are well understood and point directly to the next refinements in circuit design, switching threshold stability, and analog layout. In its present form, the instrument establishes a solid technical foundation for continued development and future flight demonstrations on small-spacecraft missions.

CHAPTER 8

FUTURE WORK

The results of this study establish a strong foundation for continued development of the magneto-inductive magnetometer as a flight-qualified scientific instrument. While the current prototype has validated the discrete LR oscillator architecture and FPGA-based timing system, areas for refinement have already been identified that will extend performance, improve manufacturability, and enable integration into upcoming spacecraft missions.

Flight Integration and System Maturation

The initial prototype presented here has already been integrated on the DARPA *Kraken* mission as a technology-demonstration payload. Its successful operation in laboratory and environmental testing has provided valuable heritage for subsequent designs. A new, four-head version of the instrument is now being prepared for the NASA *Active Cooling for Multispectral Earth Sensors* (ACMES) mission. This version employs the same FR4-based sensor head geometry described in Chapter 5, mounted in a boomless configuration that interfaces directly with the spacecraft structure. The sensing heads retain the orthogonal coil layout of the boom-mounted model but use short, shielded harnesses to minimize parasitic coupling and facilitate distributed mounting across multiple spacecraft faces. This configuration allows simultaneous magnetic measurements at multiple locations, supporting noise-removal and field-gradient analyses in flight.

Hardware and Timing Improvements

The next revision of the electronics will target higher clock speeds and finer timing resolution. The current design runs reliably at 380 MHz, but with improved routing, tighter timing margins, and faster FPGA fabric, clock rates near 600 MHz should be achievable. A higher-speed counter shortens the effective period of each measured cycle, reducing timing

quantization error and lowering susceptibility to slow environmental drift. As a result, faster cycle completion directly translates to lower measurement noise and improved field resolution without increasing integration time. These enhancements will also support more precise characterization of the instrument’s intrinsic noise floor.

Mechanical Refinement and Magnetic Cleanliness

The FR4 sensor head design will continue to evolve toward improved mechanical alignment and magnetic cleanliness. While the current surface-mount assembly achieved acceptable orthogonality for lab testing, further improvement will be obtained through the use of a precision reflow stencil or mechanical alignment jig to fix each inductor’s position during soldering.

Expanded Measurement Capability

By extending the design to a multi-head configuration, the system will support advanced noise-mitigation and spatial-filtering methods such as blind-source-separation, leveraging simultaneous measurements from multiple sensors. This capability will allow correlated magnetic signatures to be identified across the spacecraft and filtered in-situ, improving immunity to spacecraft-generated fields.

Summary

Together, these efforts define a clear trajectory toward a compact, low-power, and scientifically capable MI magnetometer suitable for deployment on small-spacecraft missions. The work presented in this thesis establishes the technical foundation that creates a roadmap for maturing the instrument from a laboratory prototype to a flight-ready, multi-head magnetic-field sensing platform.

REFERENCES

- [1] J. A. Jacobs, *Geomagnetism: Volume 4*. Elsevier, Jan. 2016.
- [2] V. U. J. Nwankwo, N. N. Jibiri, M. T. Kio, V. U. J. Nwankwo, N. N. Jibiri, and M. T. Kio, “The Impact of Space Radiation Environment on Satellites Operation in Near-Earth Space,” in *Satellites Missions and Technologies for Geosciences*. IntechOpen, Jun. 2020. [Online]. Available: <https://www.intechopen.com/chapters/70180>
- [3] M. Moldwin, E. Wilcox, E. Zesta, and T. Bonalsky, “Single-event effect testing of the PNI RM3100 magnetometer for space applications,” *Geoscientific Instrumentation, Methods and Data Systems*, vol. 11, pp. 219–222, Jun. 2022.
- [4] L. H. Regoli, M. B. Moldwin, M. Pellioni, B. Bronner, K. Hite, A. Sheinker, and B. M. Ponder, “Investigation of a low-cost magneto-inductive magnetometer for space science applications,” *Geoscientific Instrumentation, Methods and Data Systems*, vol. 7, no. 1, pp. 129–142, Mar. 2018, publisher: Copernicus GmbH. [Online]. Available: <https://gi.copernicus.org/articles/7/129/2018/>
- [5] M. A. Khan, J. Sun, B. Li, A. Przybysz, and J. Kosel, “Magnetic sensors – a review and recent technologies,” *Engineering Research Express*, vol. 3, 06 2021.
- [6] K. Mohri, K. Bushida, M. Noda, H. Yoshida, L. Panina, and T. Uchiyama, “Magneto-Impedance Element,” *IEEE Transactions on Magnetics*, vol. 31, no. 4, pp. 2455–2460, 1995.
- [7] L. Panina and K. Mohri, “Magneto-impedance effect in amorphous wires,” *Applied Physics Letters*, vol. 65, no. 9, pp. 1189–1191, 1994.
- [8] S. Tumanski, “Induction coil sensors—a review,” *Measurement Science and Technology*, vol. 18, no. 3, pp. R31–R46, Mar. 2007. [Online]. Available: <https://iopscience.iop.org/article/10.1088/0957-0233/18/3/R01>
- [9] J. S. Bennett, B. E. Vyhnalek, H. Greenall, E. M. Bridge, F. Gotardo, S. Forstner, G. I. Harris, F. A. Miranda, and W. P. Bowen, “Precision Magnetometers for Aerospace Applications: A Review,” *Sensors*, vol. 21, no. 16, p. 5568, Jan. 2021, number: 16 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/1424-8220/21/16/5568>
- [10] Y. Liu, C. Zhang, and X. Huang, “A modular magneto-inductive sensor system for low vector magnetic field measurements,” *arXiv preprint arXiv:2007.09893*, 2020. [Online]. Available: <https://arxiv.org/abs/2007.09893>
- [11] Aichi Steel Corporation, “Ami306 3-axis magnetic sensor,” <https://www.aichi-steel.co.jp/ENGLISH/smart/mi/products/ami306.html>, 2023.

- [12] M. Nosé, T. Kawano, and H. Aoyama, “Application of Magneto-Impedance (MI) Sensor to Geomagnetic Field Measurements,” *Journal of Geophysical Research: Space Physics*, vol. 127, no. 10, p. e2022JA030809, 2022. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1029/2022JA030809>
- [13] R. Prance, T. Clark, and H. Prance, “Ultra low noise induction magnetometer for variable temperature operation,” *Sensors and Actuators A: Physical*, vol. 85, no. 1–3, pp. 361–364, 2000. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924424700003757>
- [14] “RM3100: Most Accurate Low-SWaP Tri-Axis Magnetometer.” [Online]. Available: <https://www.pnisensor.com/rm3100/>
- [15] A. Leuzinger and A. Taylor, “Magneto-Inductive Technology Overview,” Feb. 2010.
- [16] N. Wallace, “Developing Firmware for Space Weather Probes 2 Using HDL Coder,” *All Graduate Theses and Dissertations, Spring 1920 to Summer 2023*, Aug. 2023. [Online]. Available: <https://digitalcommons.usu.edu/etd/8800>
- [17] R. Marriott, “Materials Engineering Branch,” 2004.
- [18] U. S. University, “Three-Axis Magnetic Field Generator.” [Online]. Available: <https://engineering.usu.edu/ece/students/senior-projects/fall2023-spring2024/projects/three-axis-magnetic-field-generator>
- [19] “Meda | Products.” [Online]. Available: <http://www.meda.com/index.php/products?name=FVM400>
- [20] Y. Gao, X. Fang, D. Ma, B. Sun, K. Wang, S. Li, Y. Dou, and M. Zeng, “Analysis and measurement of magnetic noise in multilayer magnetic shielding system combined with mu-metal and ferrite for atomic sensors,” *Measurement*, vol. 233, p. 114745, Jun. 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0263224124006304>

APPENDICES

APPENDIX A: Spice Models

This appendix documents the LTSpice circuit models used during the initial design and verification of the magneto-inductive sensor front end. The simulations were used to validate the LR relaxation oscillator topology, explore the effects of different bias networks and threshold settings, and confirm that the selected component values produced stable oscillations over the intended operating range. The schematic shown in Figure [A.1](#) reflects the final component choices used in the hardware implementation and served as a reference when interpreting measured waveforms and debugging unexpected behavior during bench testing.

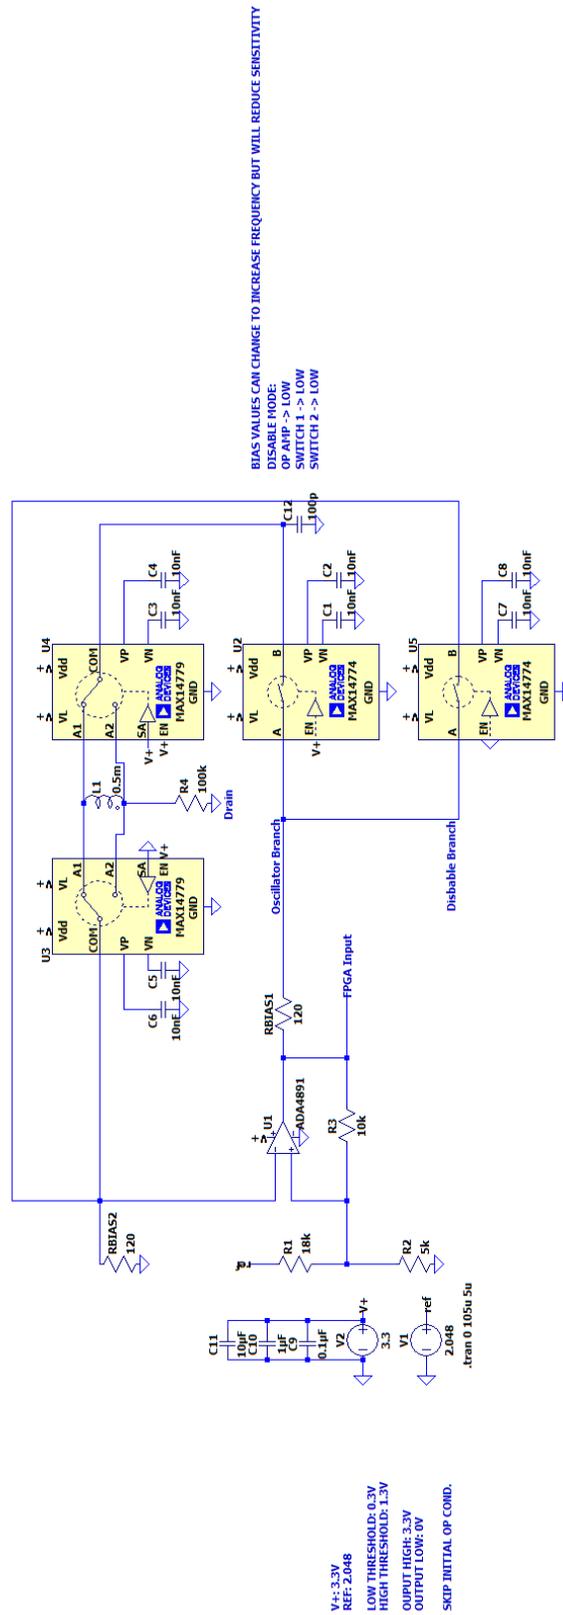


Fig. A.1: LTSpice Circuit Model.

APPENDIX B: HDL

This appendix provides the VHDL source code for the oscillation counter module used in the FPGA timing architecture. The `osc_counter` entity counts Schmitt-trigger oscillations on each axis, implements the programmable integration depth via `cycle_counts`, and enforces a configurable wait-time using `delay_counts`. Separate processes handle the oscillation-driven count and the system-clock-based delay, ensuring clean handoff between measurement windows and preventing overlap between successive acquisitions. The listing in Figure B.1 corresponds to the final version used for all laboratory measurements reported in this thesis.

```
1      library ieee;
2      use ieee.std_logic_1164.all;
3      use ieee.numeric_std.all;
4
5      --new counting method, & delay input, updated 11/14/2024
6
7      entity osc_counter is
8      port (
9          clk          : in std_logic;
10         rst          : in std_logic;
11         osc_in       : in std_logic;
12         count_en     : in std_logic;
13         delaying     : in std_logic;
14         cycle_counts : in std_logic_vector(31 downto 0);
15         --Set from SPI master
16         delay_counts : in std_logic_vector(31 downto 0);
17         --Set in matlab config file
18
```

```
19     counting      : out std_logic;
20     counter_done  : out std_logic;
21     delay_done    : out std_logic
22 );
23 end osc_counter;
24
25 architecture rtl of osc_counter is
26     signal counter : integer := 0;
27     signal d_counter : integer := 0;
28     signal counter_var : integer;
29     signal d_counter_var : integer;
30
31     begin
32
33     COUNT : process(rst, osc_in, count_en) --watch osc
34     begin
35         if rst = '1' then
36             counter <= 0;
37             counter_done <= '0';
38             counting <= '0';
39         elsif count_en = '0' then
40             counter <= 0;
41             counter_done <= '0';
42             counting <= '0';
43         elsif rising_edge(osc_in) then
44             counter <= counter_var;
45             counter_done <= '0';
46             counting <= '1';
47             if counter >= to_integer(unsigned(cycle_counts)) then
48                 counter <= 0;
49                 counter_done <= '1';
```

```
50     end if;
51     end if;
52     end process COUNT;
53
54     DELAY : process(rst, clk, delaying) --watch clk
55     begin
56         if rst = '1' then
57             d_counter <= 0;
58             delay_done <= '0';
59         elsif delaying = '0' then
60             d_counter <= 0;
61             delay_done <= '0';
62         elsif rising_edge(clk) then
63             d_counter <= d_counter_var;
64             delay_done <= '0';
65             if d_counter >= to_integer(unsigned(delay_counts)) then
66                 d_counter <= 0;
67                 delay_done <= '1';
68             end if;
69         end if;
70     end process DELAY;
71
72     counter_var <= counter + 1;
73     d_counter_var <= d_counter + 1;
74
75
76     end architecture;
```

Figure B.1. Oscillation Counter VHDL Source Code.

APPENDIX C: Simulink Models

This appendix summarizes the Simulink models used to develop, verify, and generate HDL for the IGLOO2-based control firmware. The top-level model (Figure C.1) implements the overall measurement sequence, including axis selection, polarity reversal, and timing coordination. Submodels implement the axis driver logic, data-ready handling, and SPI slave interface used to communicate with the external microcontroller.

The figures in this appendix illustrate the final block-level architecture:

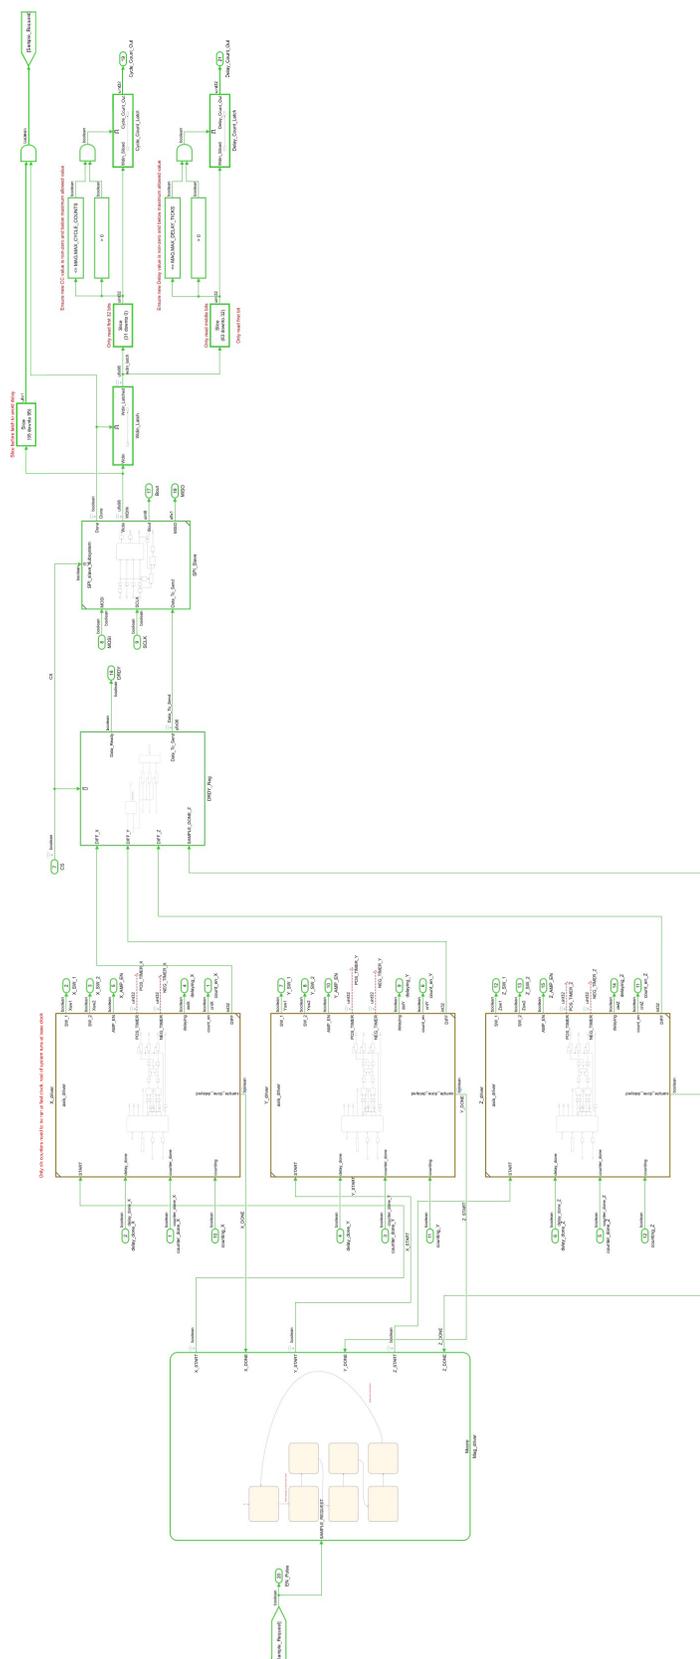


Fig. C.1: Top Level Simulink.

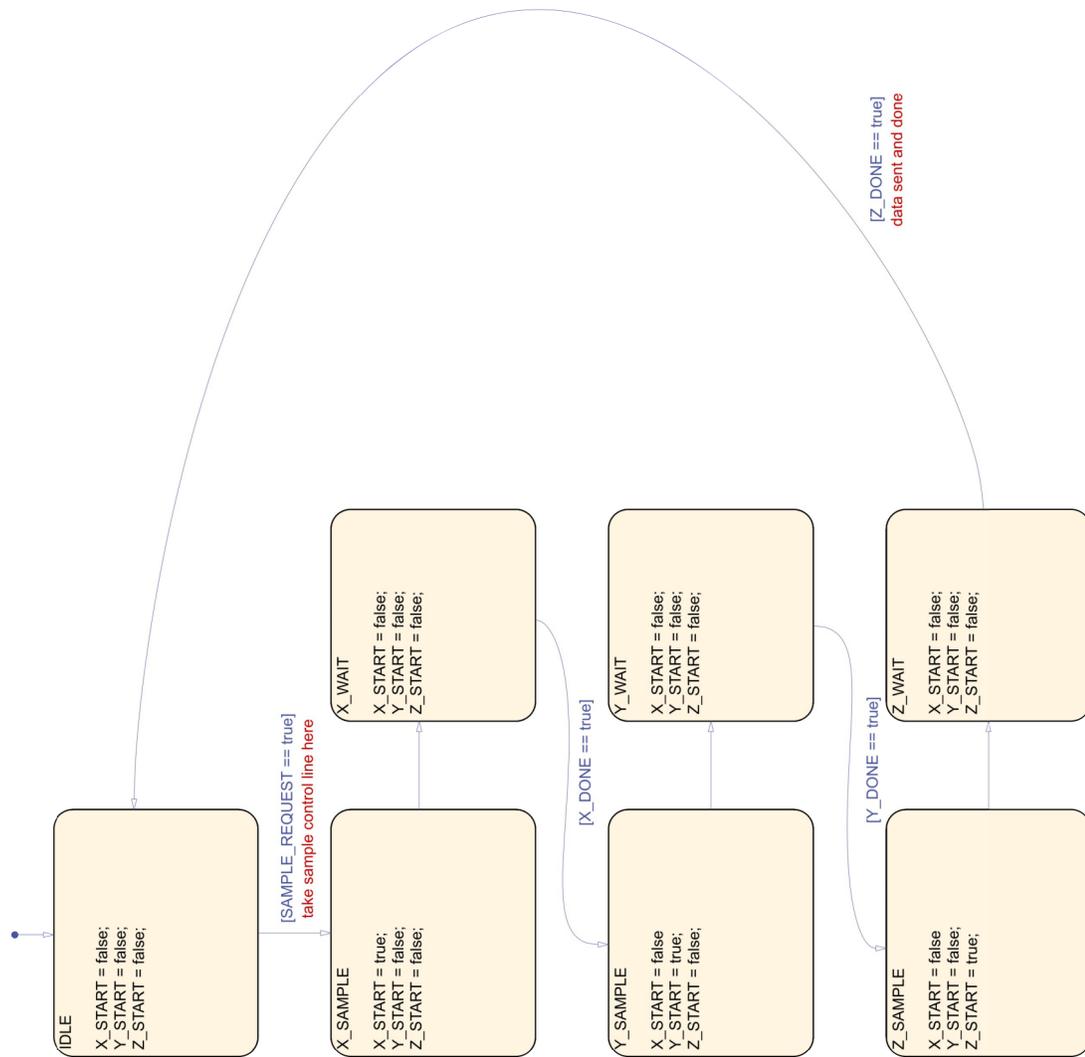


Fig. C.2: High Level Simulink FSM.

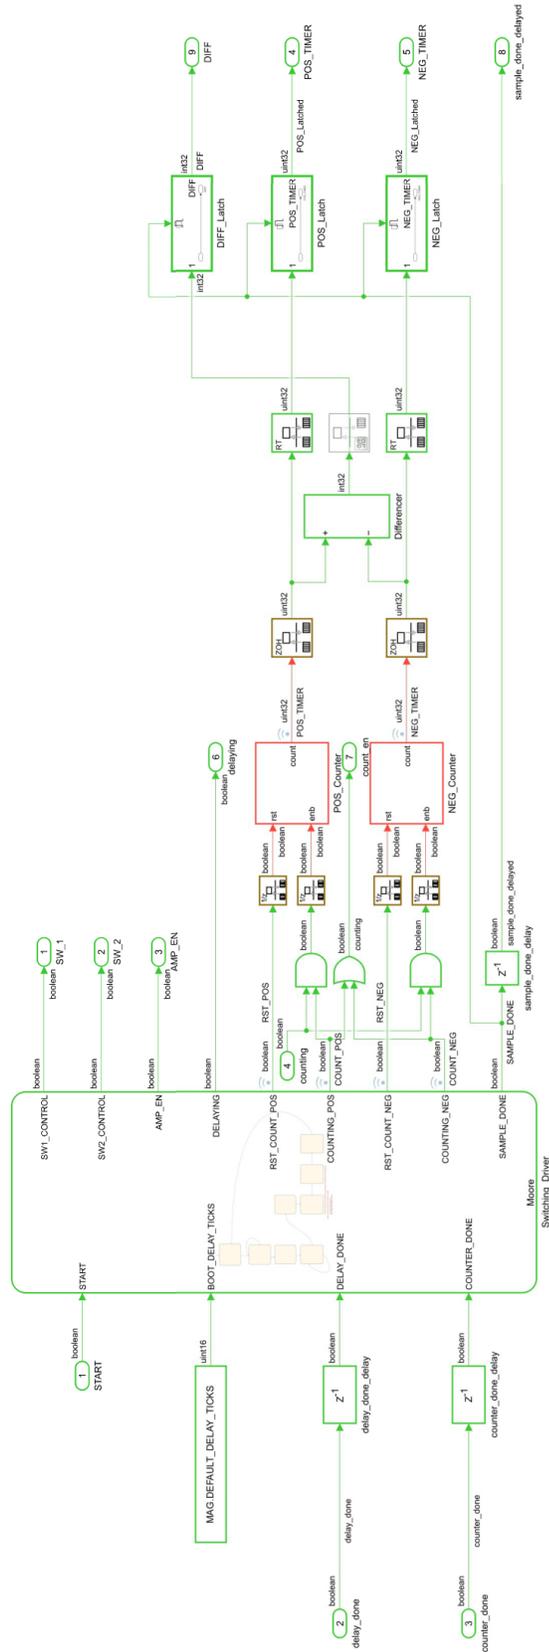
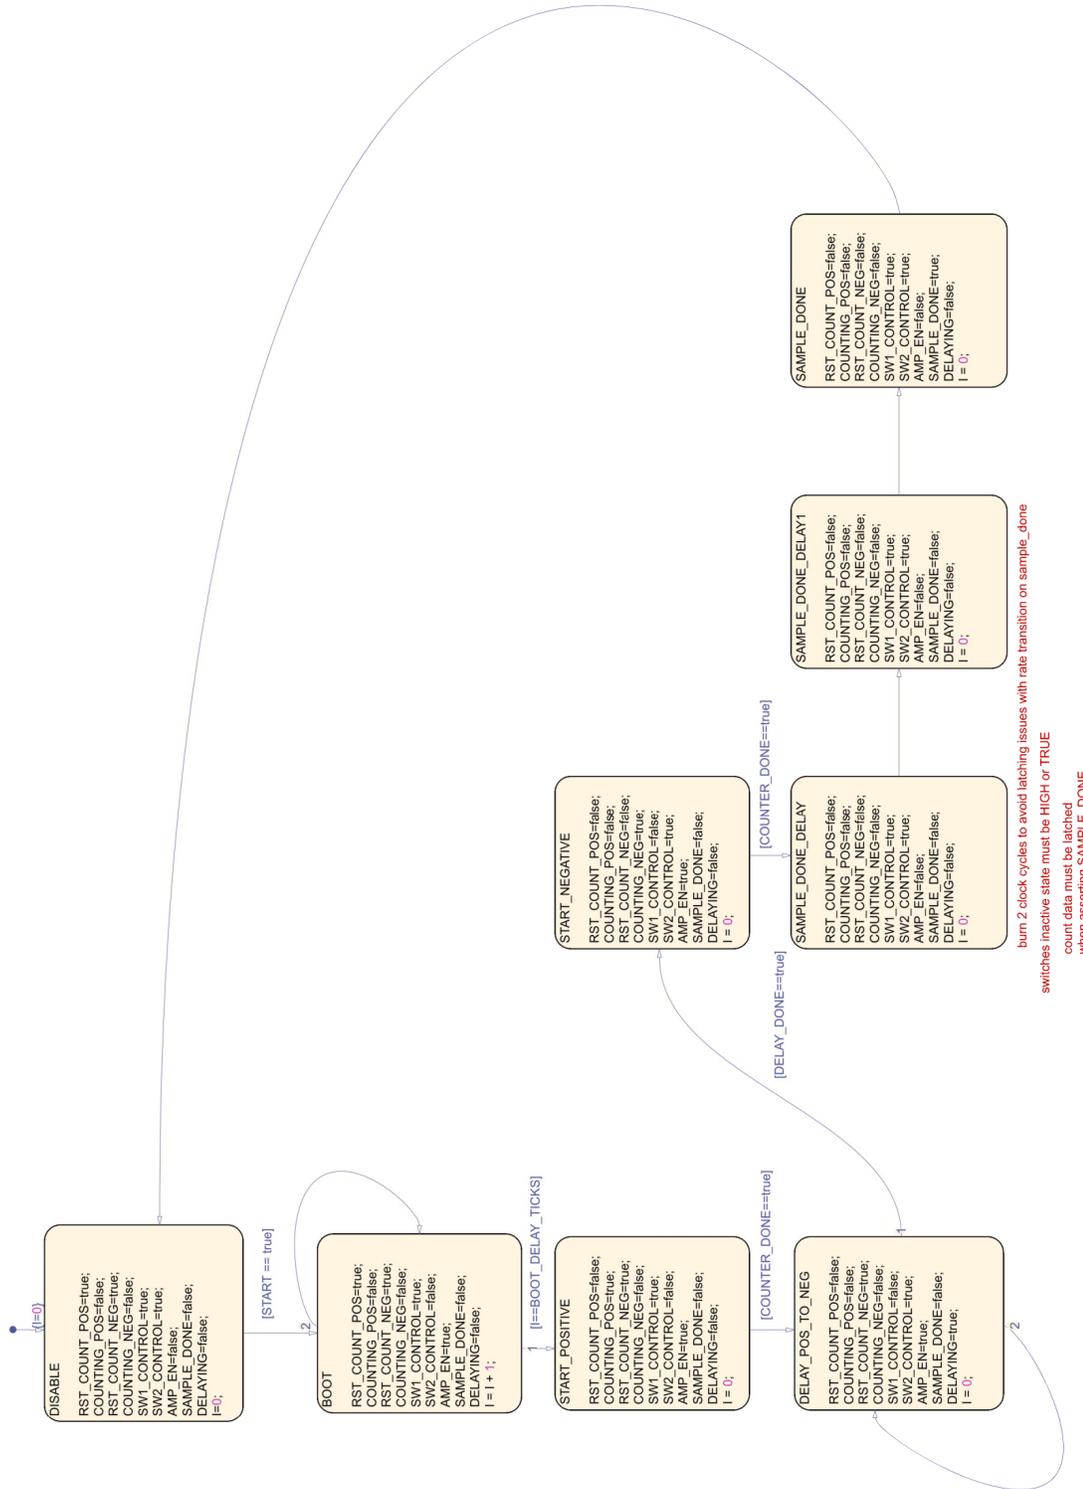


Fig. C.3: Axis Driver Simulink Model.



burn 2 clock cycles to avoid latching issues with rate transition on sample_done
 switches inactive state must be HIGH or TRUE
 count data must be latched
 when asserting SAMPLE_DONE

Fig. C.4: Axis Driver Simulink FSM.

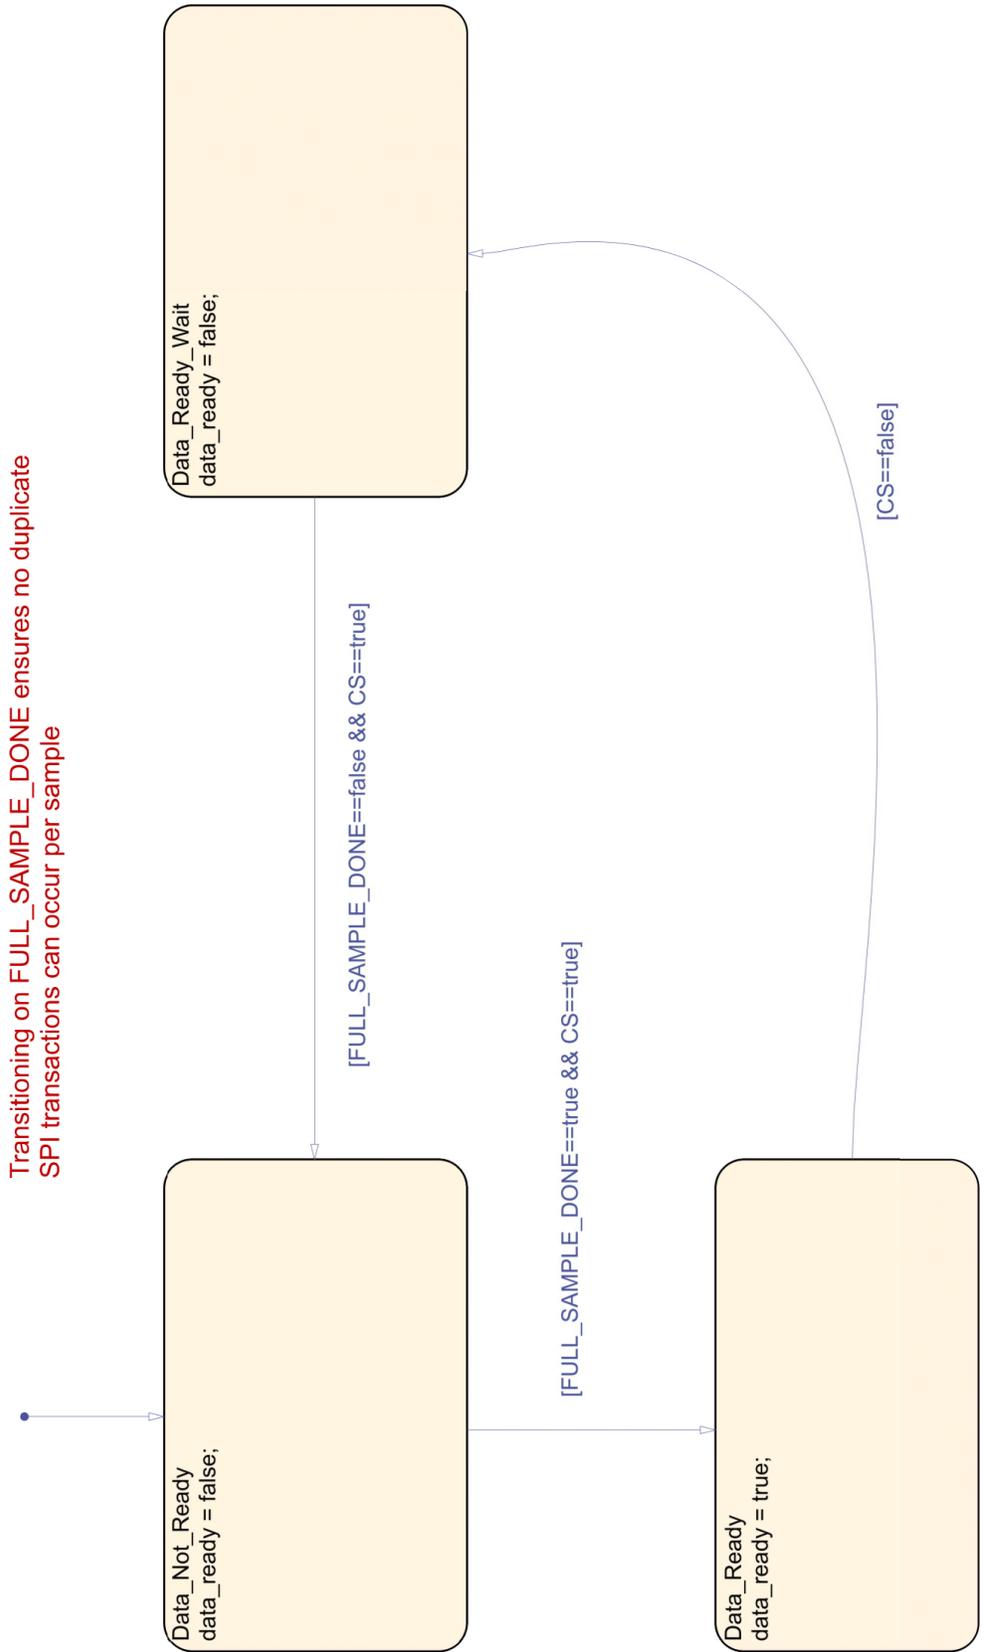


Fig. C.5: Data Ready FSM.

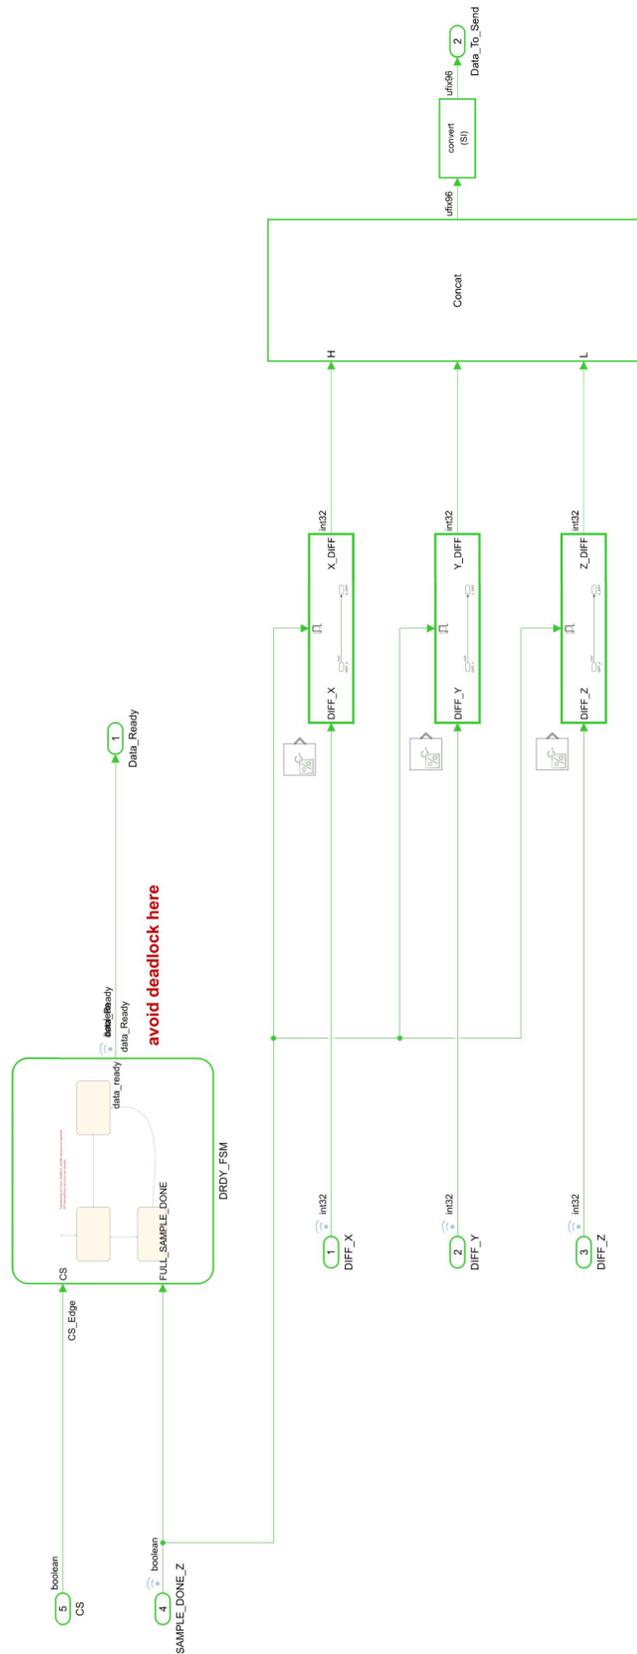


Fig. C.6: Data Ready Latching.

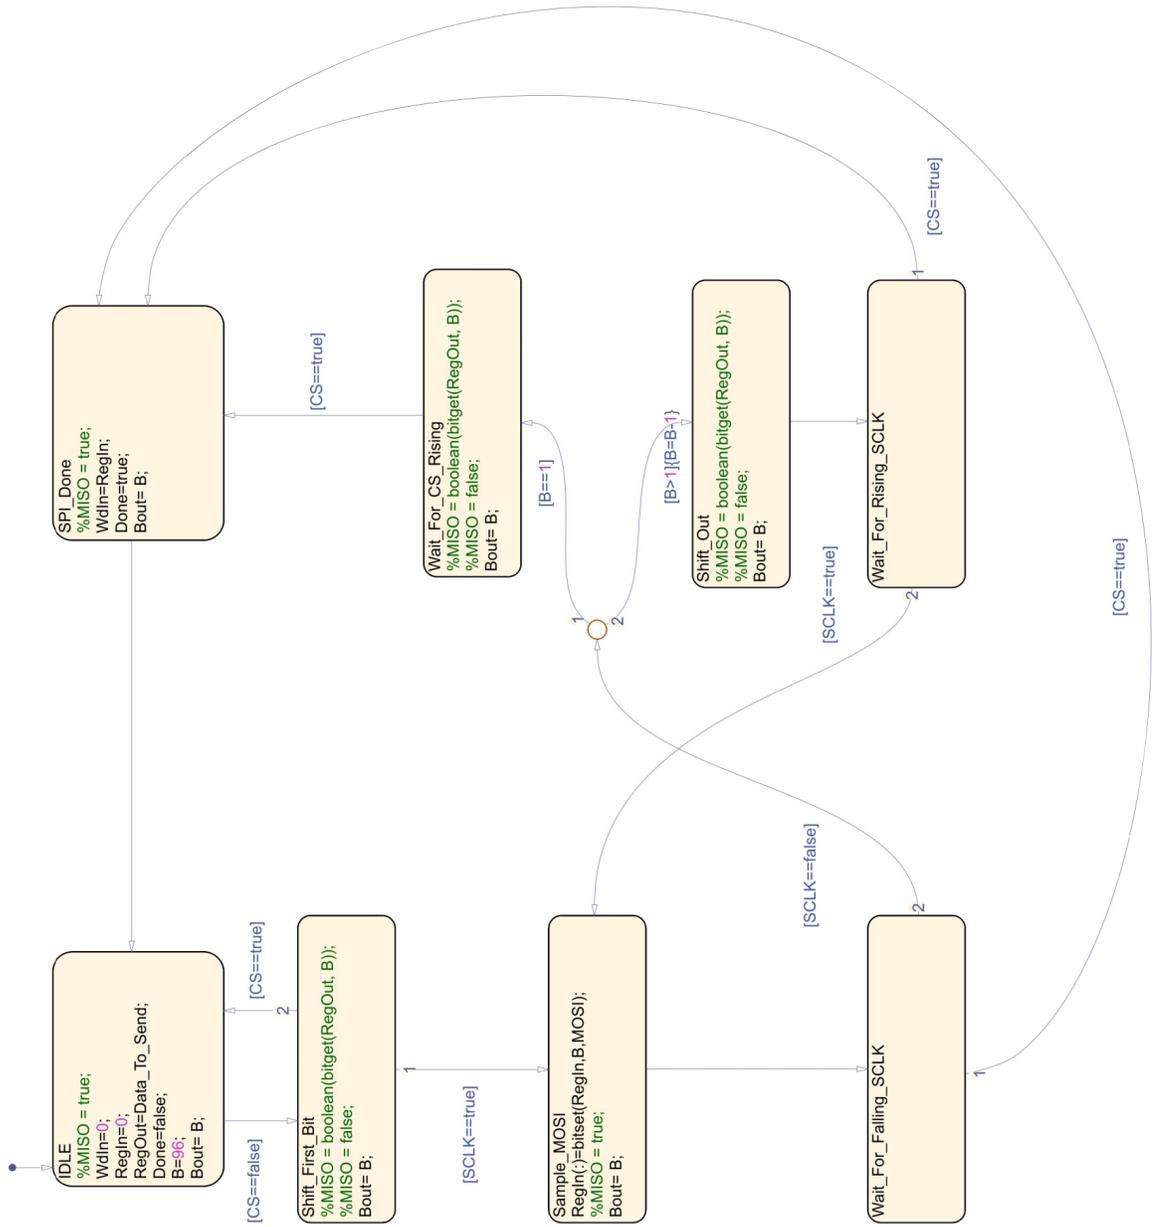


Fig. C.8: SPI FSM.

APPENDIX D: Calibration Plots for All Cycle–Count Configurations

This appendix compiles the full set of calibration and residual-analysis plots generated for each cycle–count configuration used during magnetometer characterization. For every configuration (e.g., 100cc, 1000cc, 10000cc, 20000cc, 40000cc), the MATLAB scripts automatically produced three diagnostic figures:

1. **Counts vs. Field Calibration** Displays the corrected Schmitt–trigger counts versus the reference magnetic field, along with a linear regression fit. The slope corresponds to the calibration gain (nT/count) and the intercept corresponds to the offset. Each plot includes the resulting R^2 value as a measure of linearity.
2. **Residuals vs. Fitted Field** Shows the difference between the measured reference field and the fitted model ($r = y - \hat{y}$) across the full measurement range. These plots illustrate the consistency of the fit and highlight nonlinearity or quantization irregularities. RMSE and standard deviation are reported directly in the figure titles.
3. **Residual Histograms** Presents the distribution of residuals for each axis, normalized as a probability density function. These histograms provide insight into noise characteristics, bias, and overall error behavior after calibration.

Magnetometer Counts to Field Calibration ($\pm 150 \mu\text{T}$ range)

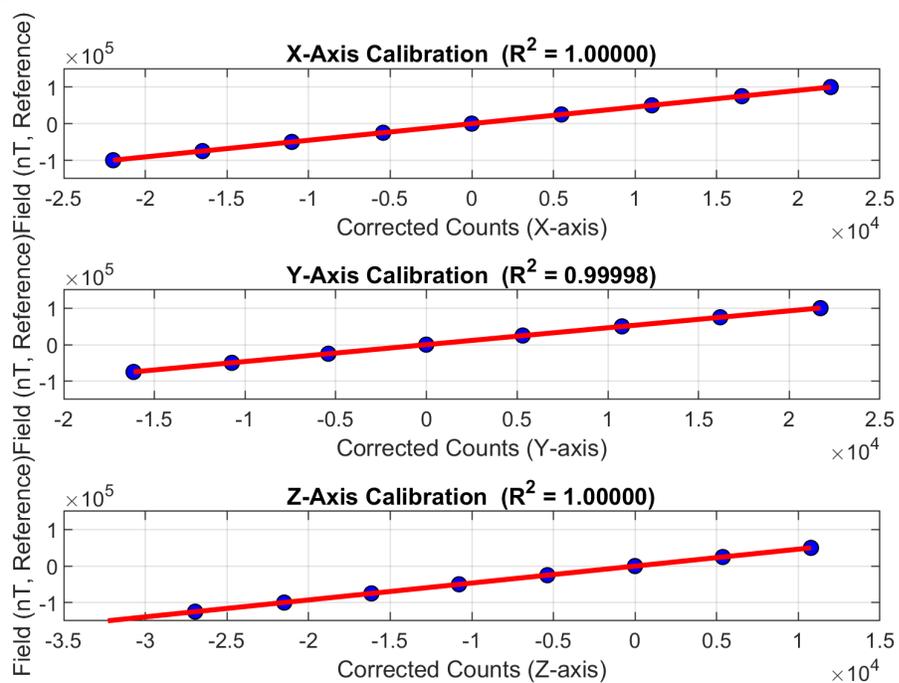


Fig. D.1: Counts vs. Field Calibration — 100cc.

Residuals vs Fitted ($r = y - \hat{y}$)

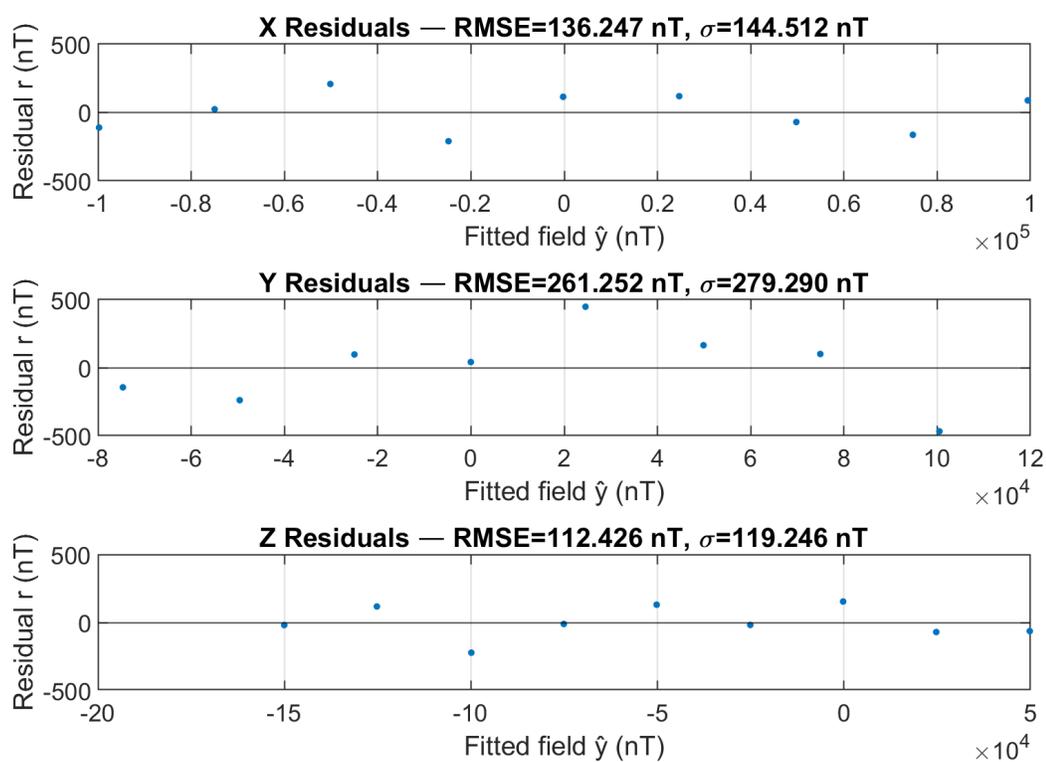


Fig. D.2: Residuals vs. Fitted Field — 100cc.

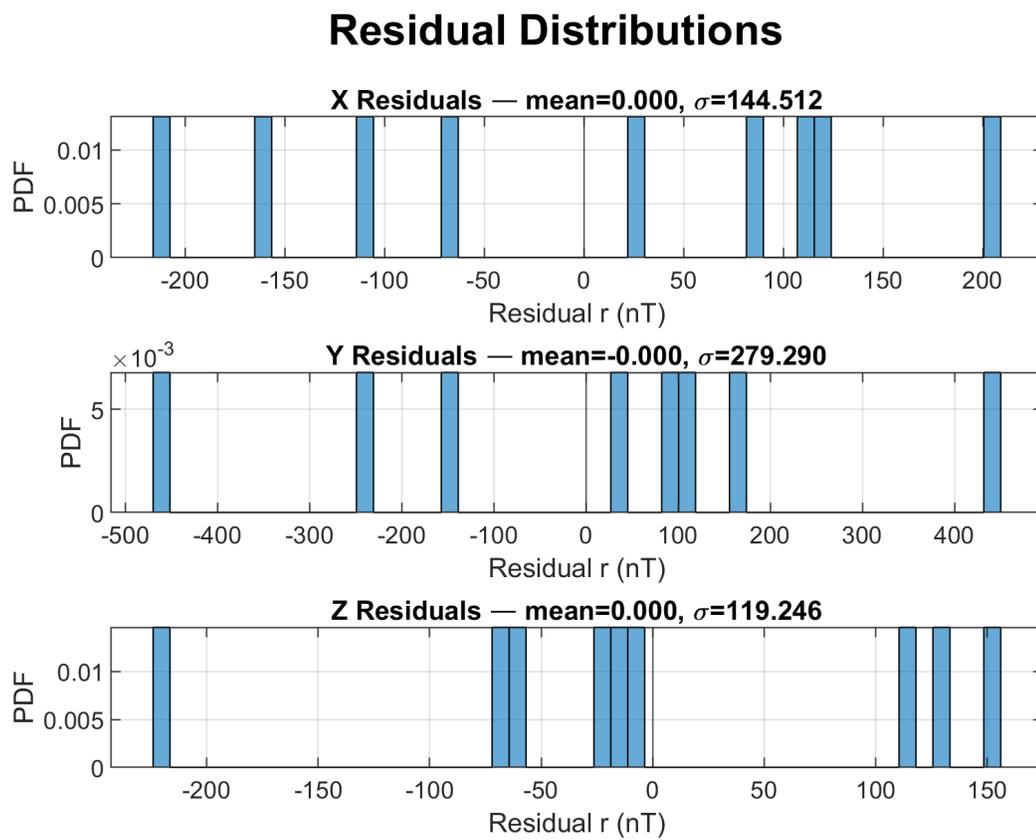


Fig. D.3: Residual Histograms — 100cc.

Magnetometer Counts to Field Calibration ($\pm 150 \mu\text{T}$ range)

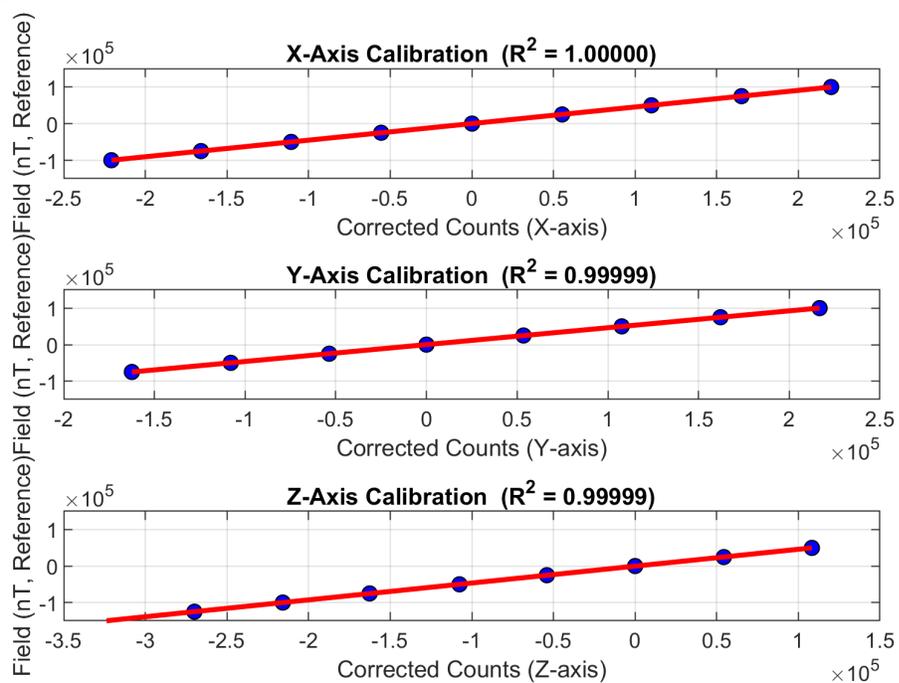


Fig. D.4: Counts vs. Field Calibration — 1000cc.

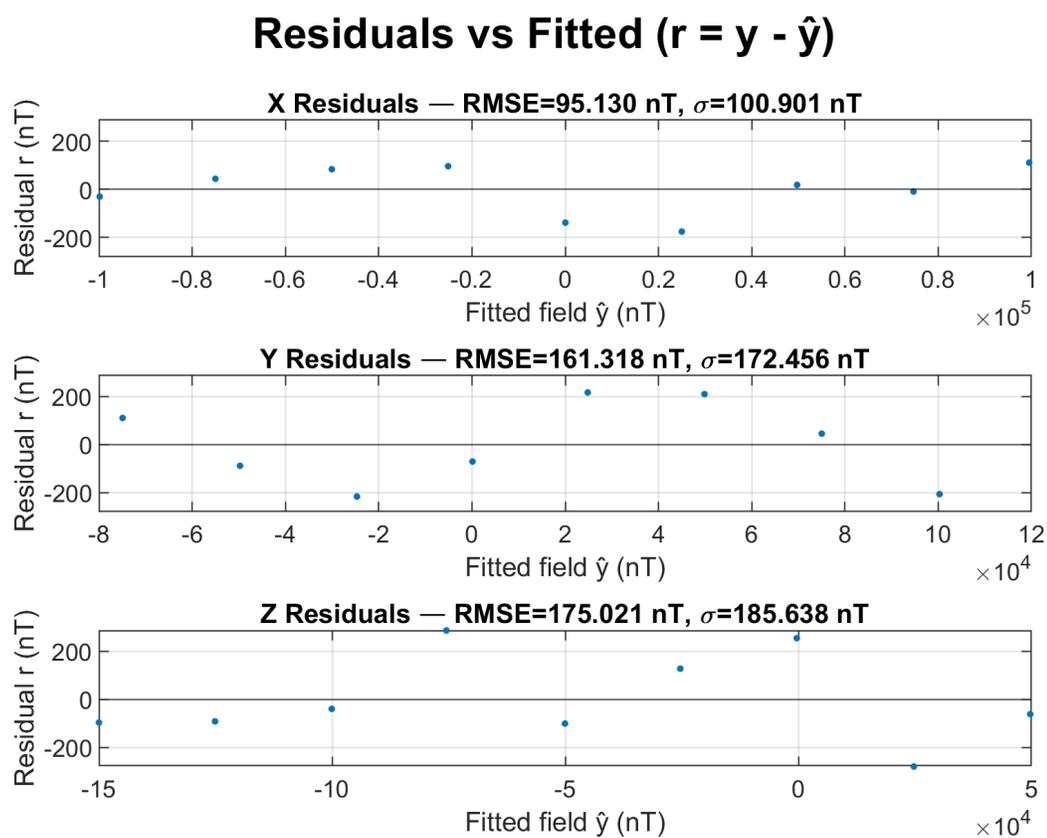


Fig. D.5: Residuals vs. Fitted Field — 1000cc.

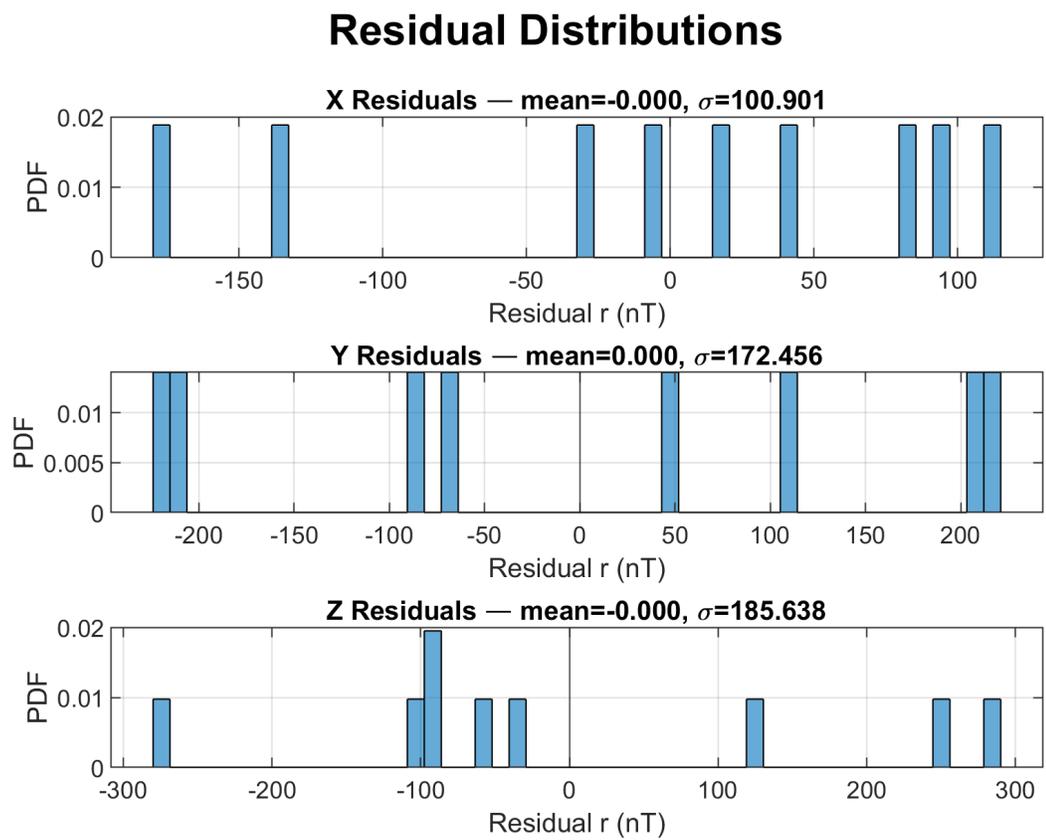


Fig. D.6: Residual Histograms — 1000cc.

Magnetometer Counts to Field Calibration ($\pm 150 \mu\text{T}$ range)

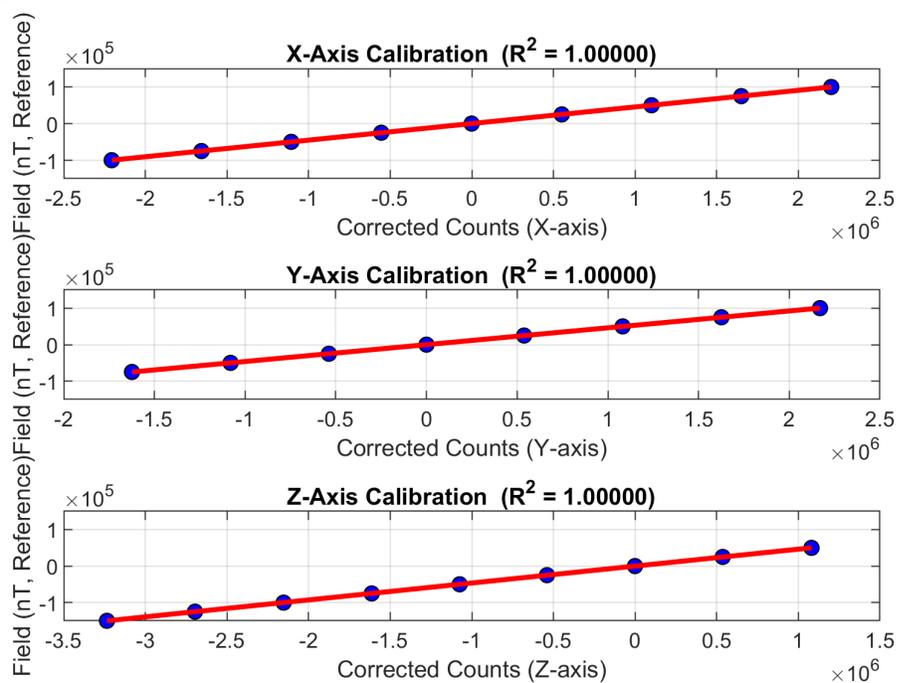


Fig. D.7: Counts vs. Field Calibration — 10000cc.

Residuals vs Fitted ($r = y - \hat{y}$)

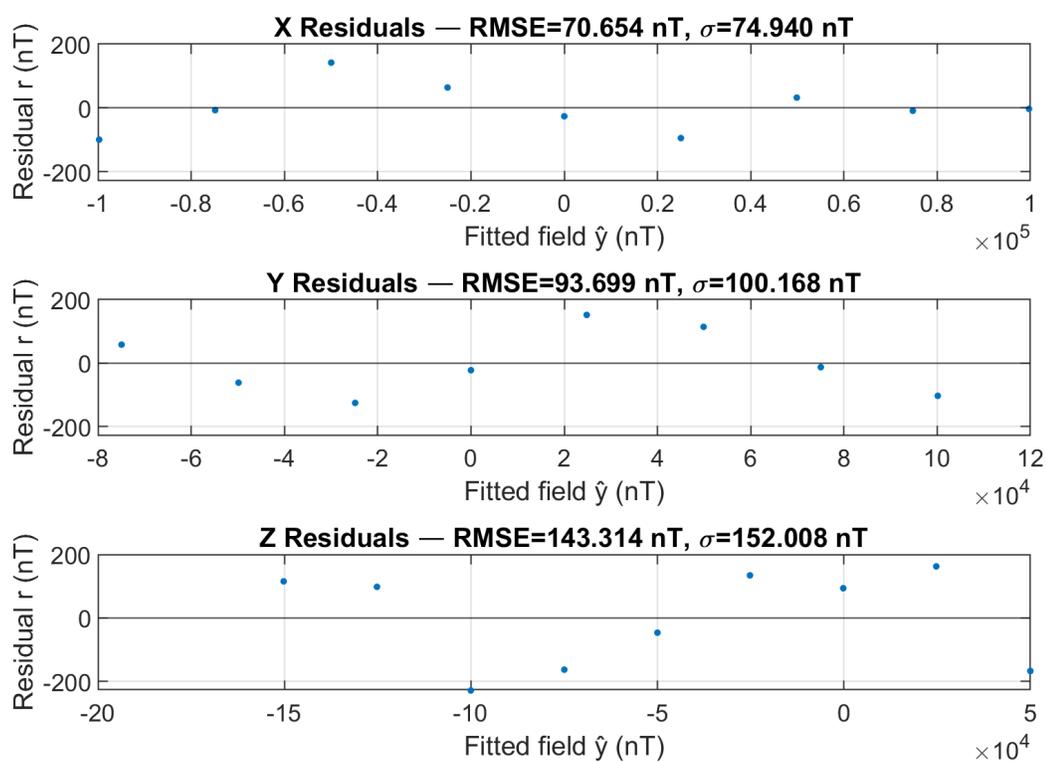


Fig. D.8: Residuals vs. Fitted Field — 10000cc.

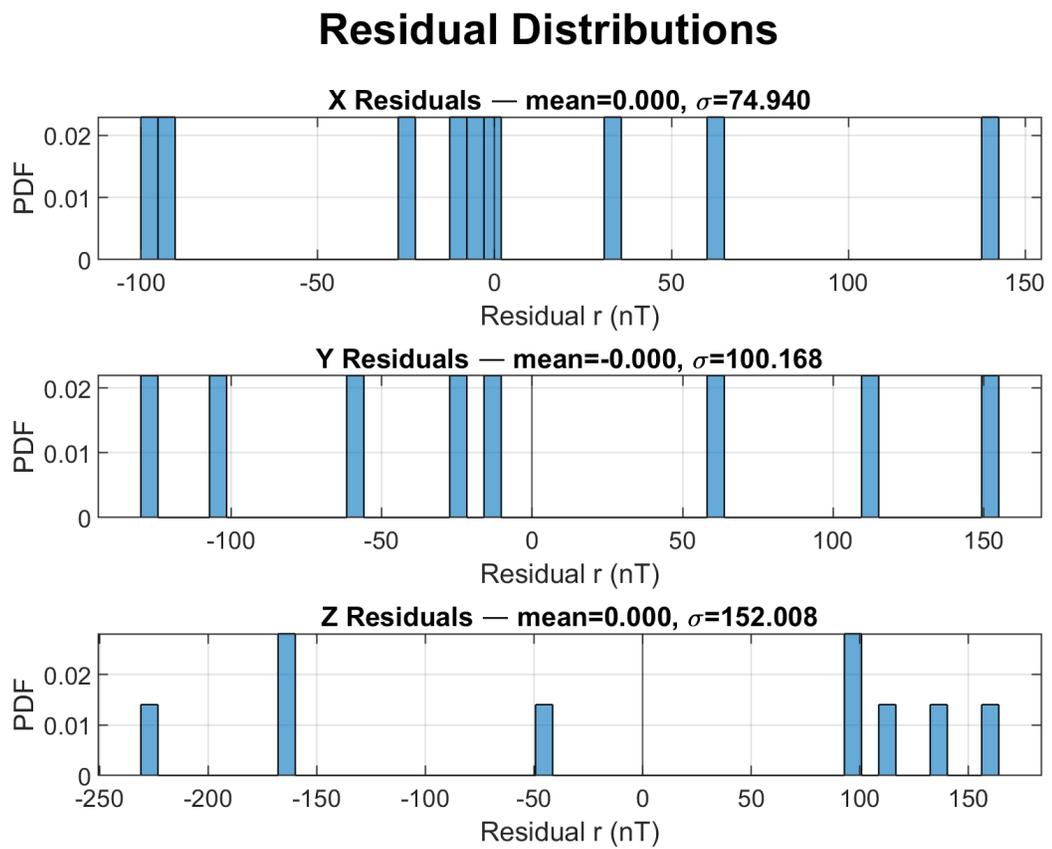


Fig. D.9: Residual Histograms — 10000cc.

Magnetometer Counts to Field Calibration ($\pm 150 \mu\text{T}$ range)

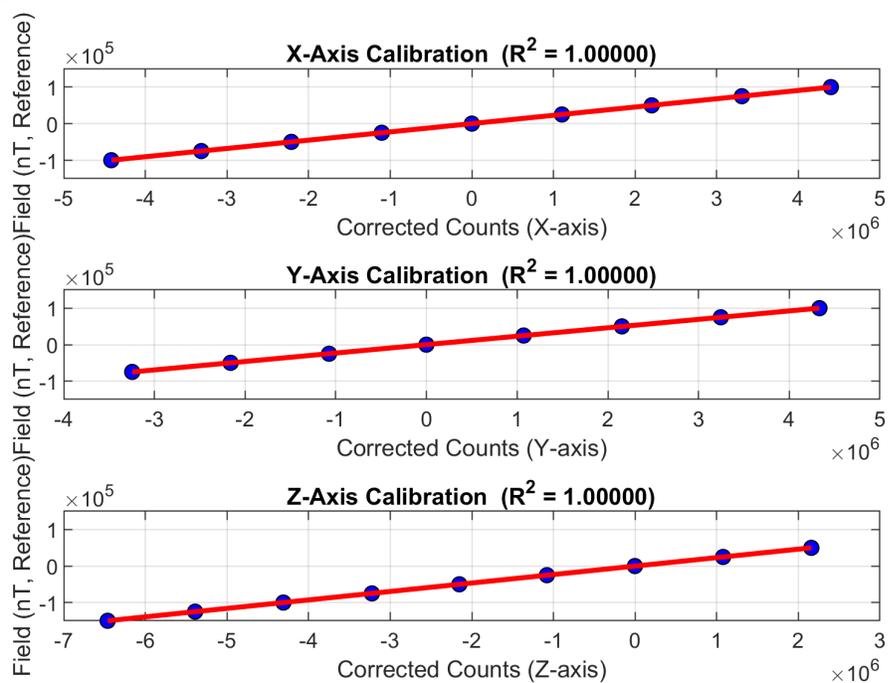


Fig. D.10: Counts vs. Field Calibration — 20000cc.

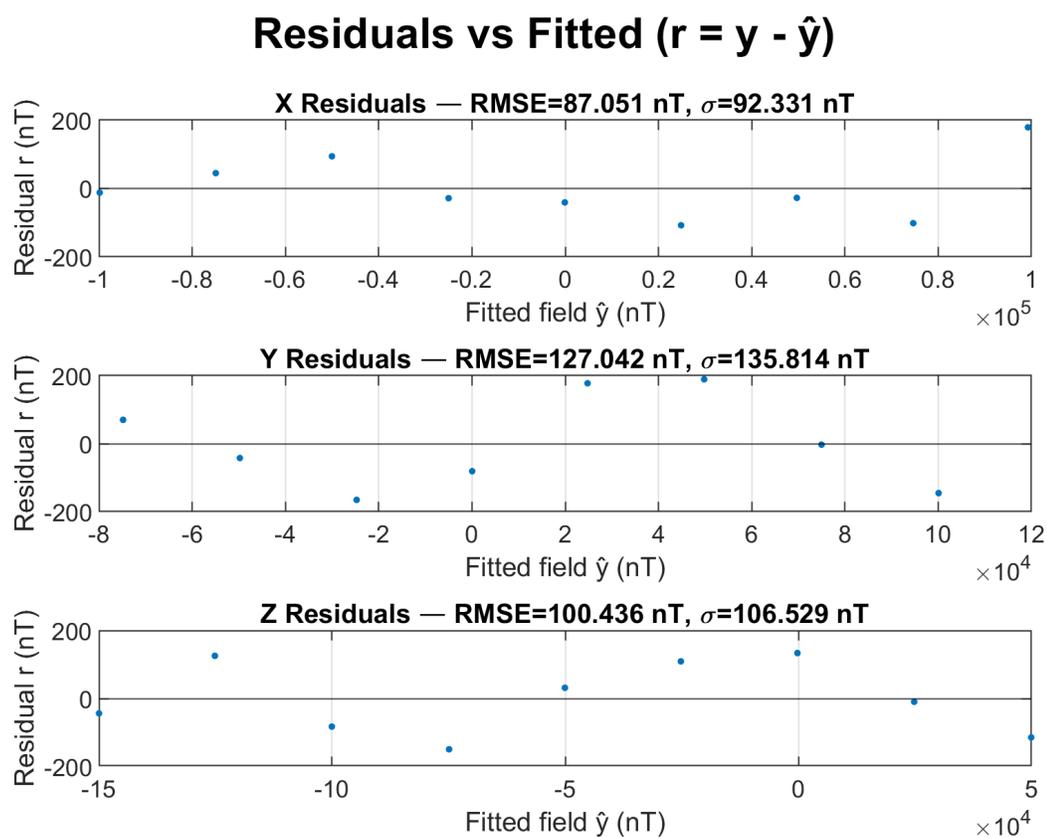


Fig. D.11: Residuals vs. Fitted Field — 20000cc.

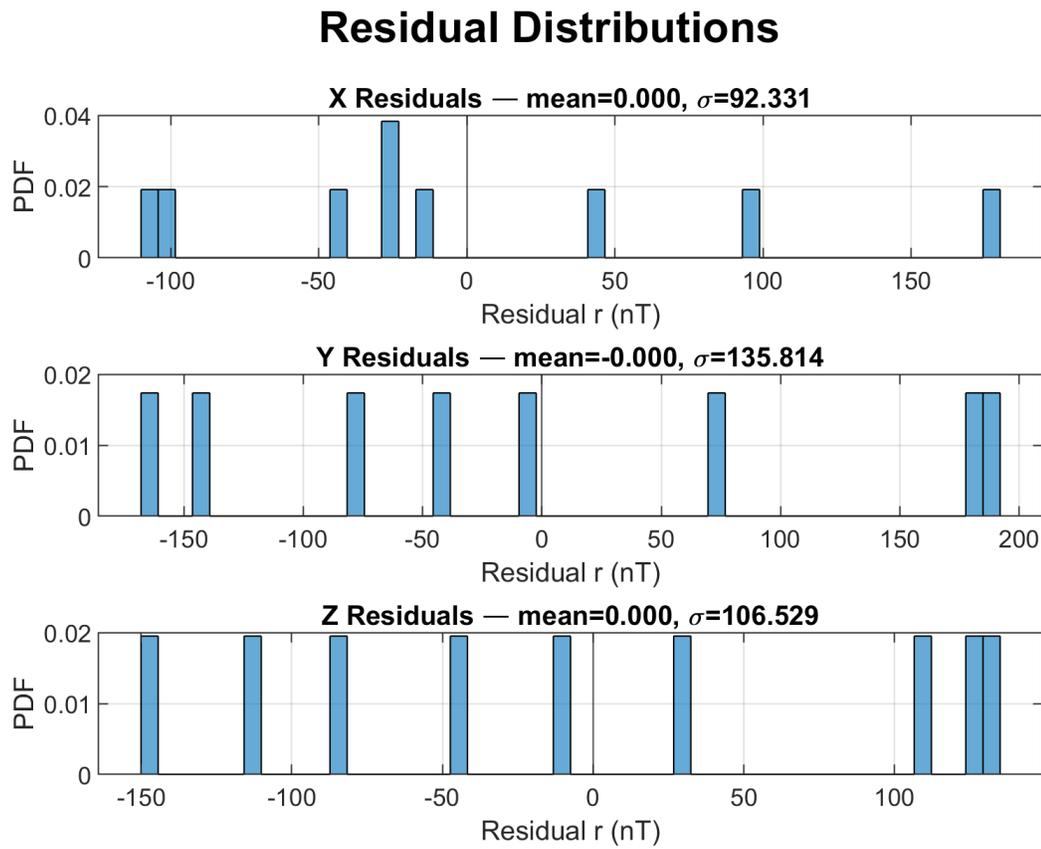


Fig. D.12: Residual Histograms — 20000cc.

Magnetometer Counts to Field Calibration ($\pm 150 \mu\text{T}$ range)

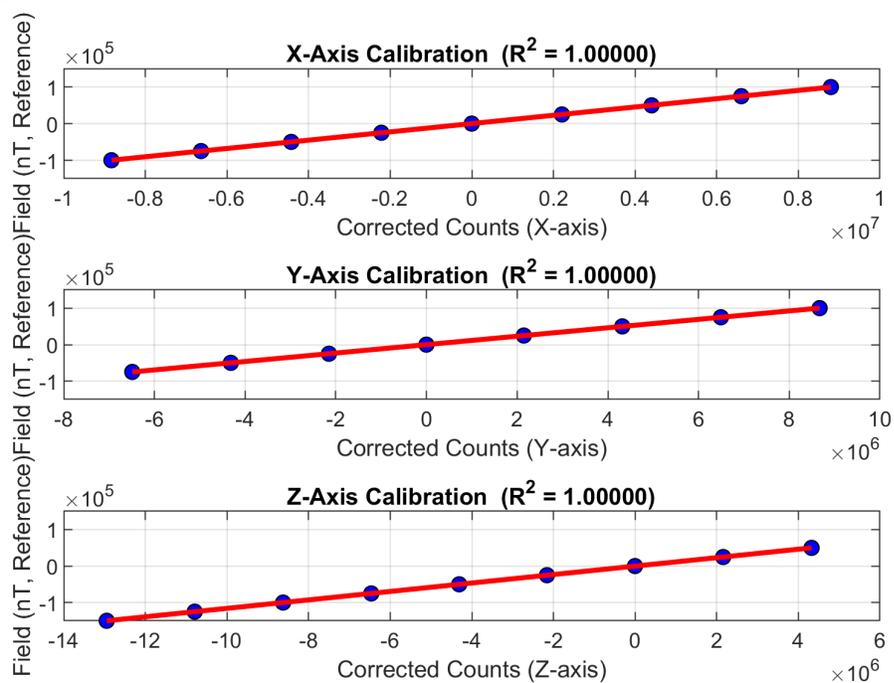


Fig. D.13: Counts vs. Field Calibration — 40000cc.

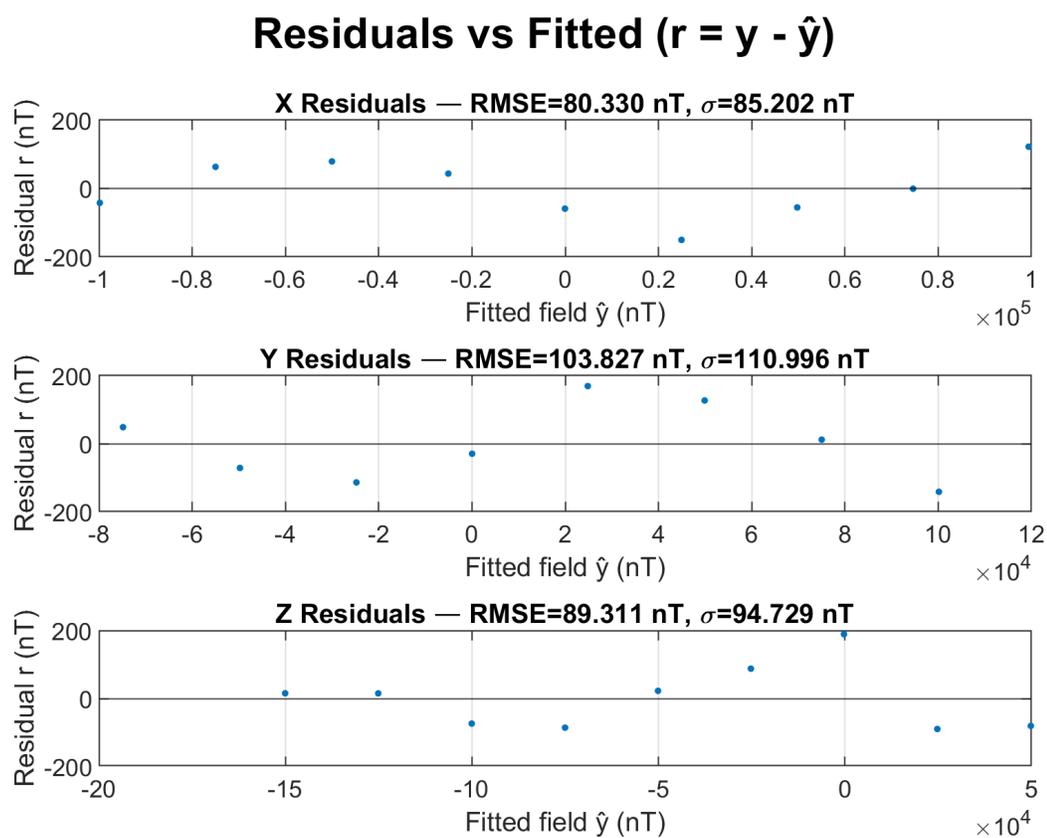


Fig. D.14: Residuals vs. Fitted Field — 40000cc.

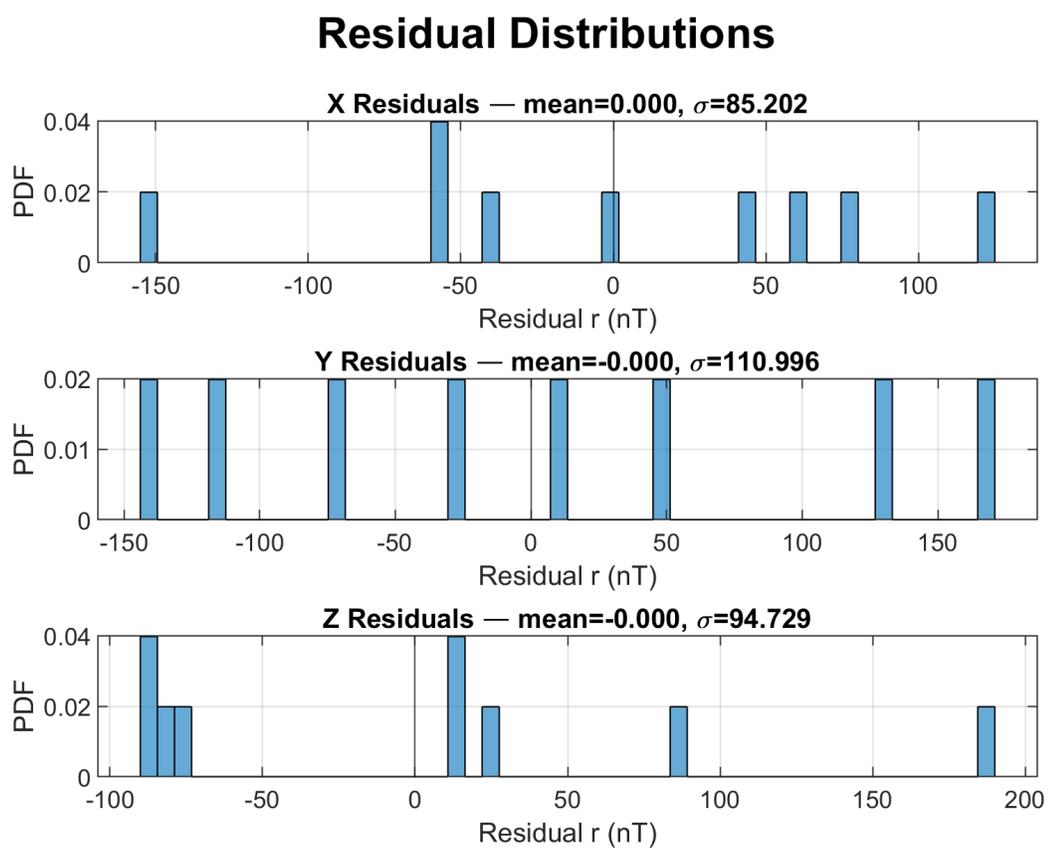


Fig. D.15: Residual Histograms — 40000cc.