

FINAL REPORT

Real-Time Hand-Held
Magnetometer Array

SERDP Project MR-2104

APRIL 2016

Dr. Mark Prouty
Geometrics, Inc.

Distribution Statement A

This document has been cleared for public release



Page Intentionally Left Blank

This report was prepared under contract to the Department of Defense Strategic Environmental Research and Development Program (SERDP). The publication of this report does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official policy or position of the Department of Defense. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the Department of Defense.

Page Intentionally Left Blank

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 29-04-16		2. REPORT TYPE Final		3. DATES COVERED (From - To) 15-08-11 - 30-04-29	
4. TITLE AND SUBTITLE Real-Time Magnetometer Array				5a. CONTRACT NUMBER W912HQ - 11-C-0004	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Mark Prouty				5d. PROJECT NUMBER MR-2104	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Geometrics, Inc. 2190 Fortune Drive San Jose, CA 95131				8. PERFORMING ORGANIZATION REPORT NUMBER GM16007F	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) SERDP 4800 Mark Center Drive Suite 17D08 Alexandria, VA 22350				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Our objective was to build, test, and characterize a diver-held ferrous metal detector array with real-time processing algorithms that accurately determines the position, depth, and size of magnetic anomalies from unexploded ordnance (UXO). This was accomplished by using an array of low power, miniature, total field atomic magnetometers, with a new miniaturized electronics design using a novel digital approach. We achieved order of magnitude reductions in the size and power consumption of total field magnetometers, and demonstrated the significant benefits these enable when used in arrays with real-time processing.					
15. SUBJECT TERMS magnetometer, unexploded ordnance, marine environment,					
16. SECURITY CLASSIFICATION OF: Unclassified			17. LIMITATION OF ABSTRACT Unlimited	18. NUMBER OF PAGES 108	19a. NAME OF RESPONSIBLE PERSON Mark Prouty
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code) 408-428-4212

Page Intentionally Left Blank

1 FRONT MATTER

1.1 TABLE OF CONTENTS

1	Front Matter	ii
1.1	Table of Contents.....	ii
1.2	Table of Figures.....	v
1.3	List of Acronyms	ix
1.4	Acknowledgements	x
2	Abstract.....	1
2.1	Objectives	1
2.2	Technical Approach.....	1
2.3	Results	1
2.4	Benefits	1
3	Objective.....	2
3.1	SERDP Relevance	2
3.2	Technical Objective	2
4	Background.....	3
5	Methods	5
5.1	Sensors.....	5
5.2	Processing Array Information	7
6	Results: Array demonstration	10
6.1	Overview	10
6.2	Design Sensor Array.....	10
6.2.1	Preparing Modeling Parameters	11
6.2.1.1	Modeling Fixture and Measurement Procedure	11
6.2.1.2	Measured Objects	13
6.2.1.3	Magnetic Signatures of the Area and the Objects	14
6.2.2	Simulation Results of Several Arrays.....	17
6.2.2.1	Background Field	21
6.2.2.2	Ambiguity Between Estimated Object and Background.....	23
6.2.2.3	Results for Mortar Shell	26
6.2.2.4	Conclusions and Considerations.....	27

6.3	Build Platform with Array of Sensors	29
6.4	Develop Real Time Processing Algorithm	30
6.4.1.1	Speed of computations.....	30
6.5	Testing and Demonstration.....	32
6.5.1	Testing System Hydrodynamics.....	32
6.5.2	Performance Simulations.....	34
6.5.2.1	Array Estimations in the Presence of Large Magnetic Bodies.....	34
6.5.2.2	Accuracy of the estimates.....	35
6.5.2.3	Performance in the presence of Multiple Objects	37
7	Results: MFAM Electronics Miniaturization	43
7.1	Overview	43
7.1.1	Approach	43
7.1.1.1	Magnetometer Instrument Architecture.....	47
7.1.1.2	Analog vs Digital Implementation	48
7.2	Sensor Model.....	50
7.2.1	Modulation: Pump laser.....	51
7.2.2	Atomic Absorption Resonance	52
7.2.3	Modulation Transfer: Magnetic Resonance	54
7.2.4	Detection: Probe Laser	55
7.2.5	Closed Loop Operation.....	55
7.2.6	Experimental Validation of Sensor Model	58
7.3	Implementation.....	60
7.3.1	System Architecture	60
7.3.2	Transimpedance Amplifier Front-End.....	61
7.3.3	Laser Driver Front-End	63
7.3.3.1	Laser Driver DC Signal Path.....	63
7.3.3.2	Laser Modulation Signal Path	65
7.3.4	Heater Driver Front-End.....	67
7.3.4.1	Vapor Cell Heater.....	67
7.3.4.2	Laser Heater.....	68
7.4	Firmware Programming.....	69

7.4.1	Model-Based Design	70
7.4.2	Larmor Loop.....	71
7.4.3	Magnetic Field Conversion	74
7.4.4	Laser Locking	74
7.4.5	MFAM Startup	77
7.4.5.1	Power On	77
7.4.5.2	Vapor Cell Heater Startup	78
7.4.5.3	Lock Lasers.....	78
7.4.5.4	Larmor Search / Lock.....	79
7.4.5.5	Data Out.....	79
7.5	Results	79
7.5.1	Sensitivity	79
7.5.2	Step Response.....	80
7.5.3	Operating Temperature.....	82
7.5.4	Gradiometer Common Mode Rejection	84
7.5.5	Eddy Current Measurement.....	88
8	Conclusion	89
8.1	Outlook: Multi-channel MFAM system	89
9	List of Publications	91
10	Appendix	92
10.1	MFAM Interface Details	92
10.1.1	Output Data Frame Format.....	93
10.1.1.1	Fid Data Format (16 bits):	95
10.1.1.2	System Status (16 Bits):	95
10.1.1.3	Magnetometer 1 Field Reading (32 bits):.....	95
10.1.1.4	Magnetometer 1 Status (32 bits):	96
10.1.1.5	Magnetometer 2 Field Reading:	96
10.1.1.6	Magnetometer 2 Status:.....	96
10.1.1.7	Auxiliary Channel X (16 bits):.....	96
10.1.1.8	Auxiliary Channel Y (16 bits):.....	96
10.1.1.9	Auxiliary Channel Z (16 bits):	96

10.1.1.10	Auxiliary Temperature (16 bits):	97
10.1.2	Ancillary Sensor Orientation and Axis Definition	97

1.2 TABLE OF FIGURES

Figure 1.	Concept of diver-held array with 12 sensors and real-time processing. 12 sensor locations are shown in blue.	3
Figure 2.	Technology roadmap indicating advances in four areas with the proposed system.	5
Figure 3.	Photo of an early prototype MFAM sensor.	6
Figure 4.	Magnetic anomaly map.....	8
Figure 5.	Analyzing anomalies by calculating dipole source that best fits the observed data. We propose an instrument performing these calculations automatically and in real time, using an array of sensors.	9
Figure 6.	Steps of this project.....	10
Figure 7.	Modeling fixture for accurately positioning sensor for reading above target.....	11
Figure 8.	Base station sensor with modeling fixture in background.....	12
Figure 9.	Google Earth image of test site.....	13
Figure 10.	Objects used. A: 10 cm metal ball; B: 80mm mortar shell; C: Neodymium magnet.	14
Figure 11.	Typical background for the area. Note that size of the computed field is 5 x 5 m while measuring frame is only 1x1 m (shown in the center). Magnetic field changes from 13 nT in the left bottom of the frame to corner to 14.4 nT in the right top corner.	15
Figure 12.	A (left): The field variation for the metal ball inside the figure is 14.5 to 30 nT . B (right): Mortar shell field variation inside the figure from 16 to 57 nT. Note the 1m measurement frame (rectangle) is a small portion of the figure.....	15
Figure 13.	The field variation of the Nd magnet inside the figure varies from -146 to 330 nT.....	16
Figure 14.	Different array configurations and their names used in the measurements.	17
Figure 15.	Overall position error for all tested configurations. The title denotes sensor count and geometry.....	18
Figure 16.	Horizontal position differences for all tested configurations.....	19
Figure 17.	Depth error distributions for tested configurations.....	20
Figure 18.	Results of the inversion in the absence of a target.....	22
Figure 19.	Results of inversion for the mortar shell in the nose down configuration.....	23
Figure 20.	Inversion results for the strong dipole with different orientations.....	25
Figure 21.	Inversion results for mortar shell object at different orientations. Estimated background values without object are D=14.6nT, C=3.4 nT.....	26
Figure 22.	Array used for experiments of real-time object tracking and localization.	29
Figure 23.	Electronics modules built for the array measurements.....	30
Figure 24.	Real time display of the fitting program.	31
Figure 25.	Mechanical mockup to test the system's maneuverability underwater.	32
Figure 26.	Enlisting the assistance of local divers in Monterey Bay.....	33
Figure 27.	Preparing to test the system underwater.	33
Figure 28.	Maneuvering the system underwater.	34
Figure 29.	Dipole matching with presence of systematic disturbance due to significant source of magnetic field.....	35
Figure 30.	Results of the real-time algorithm. Mortar shell is at 1.13m depth with moment 0.4 A m ² . 36	36

Figure 31. Results of the real-time algorithm. Coil simulator as target with moment 0.04 Am^2	36
Figure 32. Results of the real-time algorithm in estimating magnetic moment. Actual moment was 0.04 Am^2	37
Figure 33. Estimated X position at different locations of the array.	38
Figure 34. Estimated Y position, which should be at 0. Considerable error is observed when the array is between the targets (array position equal 0). The error is bigger for 2 m target separation. Again it shows that the operator would see Y position changing dramatically with array motion while in between the targets.....	38
Figure 35. Estimated Z position. True position for both targets is 1.2 m, and the system measures this value on either side of the targets. Instability of estimated Z is observed while operating inside targets.	39
Figure 36. Standard deviation of estimated X, in logarithmic scale. It can be seen that this values grows exponentially when system loses “lock” on either side of the targets. This is a clear indication that the system is not operating properly. All configurations deliver however reasonable standard deviation while in +/- 1 m from the targets.	39
Figure 37. Estimated magnetic moment, cgs. Note that real magnetic moment was 600 cgs for each object, and system converges to that value only with dipole separation of 2 m. In other cases higher values of 800 and 1000 cgs are delivered, which makes sense because of the presence of two objects.	40
Figure 38. Estimated position with two objects at 1m apart.	41
Figure 39. Estimated position with two objects at 0.5m apart.	41
Figure 40. Estimated moment with two objects 0.5m apart.....	42
Figure 41. Iterative design process in which both the physics package and electronics complexity are considered.....	43
Figure 42. Model-based Design Workflow.....	44
Figure 43. The first step in the physical electronics design was to prototype the system using programmable laboratory instrumentation.	45
Figure 44. The entire system was built in a bench top form factor for testing.	45
Figure 45. Knowing the exact design we needed, the final step was to package the circuits into the desired size and shape.	46
Figure 46. Completed electronics system.	46
Figure 47. Architecture of a magnetometer based on the MFAM Module.....	47
Figure 48. Physics package exploded view.....	47
Figure 49. Analog method of creating Larmor signal.....	48
Figure 50. Digital method of creating Larmor signal.....	49
Figure 51. The physics model of the sensor operation can be described in terms of cross pumping and probing on the cesium D1 transition. The modulated pump beam sets up Larmor oscillations in the F=4 hyperfine ground state and the probe beam is used to detect the oscillations.	50
Figure 52. Sensor Larmor signal path diagram. The system behaves as a resonant circuit to the electronics.....	51
Figure 53. Pump laser current waveform.....	52
Figure 54. Atomic absorption resonance.....	52
Figure 55. Instantaneous optical pumping rate.	53
Figure 56. Demodulated optical pumping rate.....	54

Figure 57. Phase and amplitude characteristics of the magnetic resonance.....	54
Figure 58. Low pass filtered probe laser signal.	55
Figure 59. Growth of the oscillatory signal with positive feedback through the resonance.	56
Figure 60. Simulink™ test bench used to model the closed loop magnetometer.	56
Figure 61. The results of our Simulink model predicted the proper operation and performance of the system.....	57
Figure 62. Even advanced performance characteristics, such as estimating the dead zone, could be explored in the simulation environment.....	57
Figure 63. Schematics of the experimental setup.....	58
Figure 64. Probe lock-in signal (rms amplitude) as a function of pump AM frequency f . The black dot (red solid) curve represents the in-phase (quadrature) component.	58
Figure 65. Open-loop spectral density of the magnetometer noise (Green curve). Inset shows the bandwidth of the magnetometer. The red curve is the sensitivity of the magnetometer, which represents the minimum signal level detectable by the magnetometer with a signal to noise ratio of 1:1.	59
Figure 66. MFAM System Architecture.	60
Figure 67. Distributed transimpedance amplifier topology.	61
Figure 68. Transimpedance amplifier transfer function characterization. Note the flat phase response in the Larmor band or 70kHz to 350kHz.	62
Figure 69. Block diagram of Larmor signal quantization.	62
Figure 70. Laser driver front-end topology for the pump laser. The probe laser only requires the lower signal path.	63
Figure 71. Frequency response of the low-pass filter for the laser driver low frequency analog signal path.....	64
Figure 72. The gain response of the output filter for the DAC that generates the Larmor signal.....	65
Figure 73. The yellow waveform shows an RF carrier at 10MHz modulated by the Larmor signal (blue waveform). Note the delay between the driving signal and the modulated signal.....	66
Figure 74. Topology of various heater control loops in the MFAM.....	66
Figure 75. Vapor cell heater driver block diagram.	67
Figure 76. Cell heater waveforms.	68
Figure 77. Laser heater driver block diagram.	69
Figure 78. Laser heater waveform.....	69
Figure 79. Model-based design workflow used in development of FPGA code.	70
Figure 80. Example of Model-based design used to implement pump laser current amplitude modulation.	71
Figure 81. Design of the digital Larmor loop.....	72
Figure 82. Digital PLL Design Concept.....	73
Figure 83. The hardware-in-the-loop system allowed us to verify performance.	73
Figure 84. Output filter chain used to report magnetic field reading.	74
Figure 85. Laser locking servo closed-loop model.	75
Figure 86. Laser locking concept. As the laser frequency is swept across the optical absorption curve, the light output follows the green curve shown above. If the laser frequency is modulated by a sinusoidal waveform, represented by the red curve, then the AC component of the green curve has characteristics shown by the blue curve and described in the text.	76

Figure 87. Demodulator output (red curve) used to lock the laser, derived from the blue curve shown in Figure 86.	77
Figure 88. MFAM power-on start-up sequence.	77
Figure 89. Start up operation of the final system.	78
Figure 90. Noise performance of the MFAM module as measured in a zero-gauss shield can. Note that the two sensors, shown by blue and red curves, have sensitivities below $1\text{pT}/\sqrt{\text{Hz}}$. The green curve is the amplitude noise spectral density of the gradient measurement between the two sensors. Here, the two sensors were placed in contact, with approximately 1-inch separation between sensing cells.	80
Figure 91. Response of the MFAM to a square wave magnetic field.	81
Figure 92. The bottom graph shows the operating temperature of the MFAM module. The temperature varies from -30 C to +50 C over 50 minutes. The top graph shows the readings of the two sensors from the MFAM module.	83
Figure 93. The common mode cancellation of the sensors is demonstrated by the suppression of common mode cultural noise measured in a typical commercial environment by more than an order of magnitude. The sensor performance can be verified without shielding the cultural noise in a ZGC.	85
Figure 94. Measurement of local gradients in presence of large disturbance far away. Here, the sensor separation was 10 inches.	87
Figure 95. Effect on the phase of 60Hz magnetic field component due to eddy currents induced in an aluminum plate by ubiquitous AC power line fields.	88
Figure 96. Multi-channel magnetic field acquisition system based on the MFAM technology.	89
Figure 97. Multi-chip module.	90
Figure 98. Straightforward Application Specific Integrated Circuit development from the RTL (Register Transfer Level) Design of the digital logic.	90
Figure 99. MFAM Module connector details.	92
Figure 100. SPI Frame Data Format	95
Figure 101. Position and orientation of ancillary sensors for orientation estimation	97

1.3 LIST OF ACRONYMS

ADC – Analog to Digital Converter
AFE – Analog Front End
AGC – Automatic Gain Control
AM – Amplitude Modulation
AOM – Acousto-Optic Modulator
ASIC – Application Specific Integrated Circuit
CAD – Computer Aided Design
CIC – Cascaded Integrator Comb
CMRR – Common Mode Rejection Ratio
CPU – Central Processing Unit
CSAC – Chip scale atomic clock
DAC – Digital to Analog Converter
DARPA – Defense Advanced Research Projects Agency
DC – Direct Current
DDS – Direct Digital Synthesis
DPLL – Digital Phase Locked Loop
EM – Electromagnetic
ESD – Electrostatic Discharge
FM – Frequency Modulation
FFT – Fast Fourier Transform
FIR – Finite Impulse Response
FPGA – Field Programmable Gate Array
FWHM – Full Width Half Maximum
GHz – Gigahertz
GPIO – General Purpose Input / Output
Hz – Hertz
HWHM – Half Width at Half Maximum
IC – Integrated Circuit
IIR – Infinite Impulse Response
LNA – Low Noise Amplifier
LPF – Low Pass Filter
MBD – Model Based Design
MEMS – Micro-Electro Mechanical System
MFAM – Microfabricated Atomic Magnetometer
MF – z-component Magnetic Quantum Number, MF
MHz –Megahertz
MIT – Massachusetts Institute of Technology
NIST – National Institute of Standards and Technology
nT – nano Tesla
PBS – Polarization Beam Splitter
PCB – Printed Circuit Board
PD -- Photodiode
Physics Package – The physical structure of the cell, laser diode, and light source.
PLL – Phase Locked Loop
PPS – Pulse Per Second
pT – pico Tesla
RF – Radio Frequency
RMS –Root Mean Square
RTL – Register Transfer Level

SERDP – Strategic Environmental Research and Development Program

SNR – Signal to Noise Ratio

SPI – Serial Peripheral Interface

SQUID – Superconducting Quantum Interference Device

TIA – Transimpedance Amplifier

TSV – Through-Silicon Vias

UXO – Unexploded Ordnance

VCSEL – Vertical-Cavity Surface-Emitting Laser

VCO – Voltage Controlled Oscillator

VCCS – Voltage Controlled Current Source

VHDL – VHSIC Hardware Design Language

VHSIC – Very High Speed Integrated Circuit

ZGC – Zero-Gauss Chamber

1.4 ACKNOWLEDGEMENTS

This material is based upon work supported by the Strategic Environmental Research and Development Program (SERDP) through the U.S. Army Corps of Engineers, Humphreys Engineering Center Support Activity under Contract No. W912HQ-11-C-0004. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of SERDP or the U.S. Army Corps of Engineers.

2 ABSTRACT

2.1 OBJECTIVES

Our objective was to build, test, and characterize a diver-held ferrous metal detector array with real-time processing algorithms that accurately determines the position, depth, and size of magnetic anomalies from unexploded ordnance (UXO). We accomplished this by using an array of prototype low-power, miniature, total field atomic magnetometers, designed using new miniaturized electronics and a novel digital approach. Specific technical objectives were as follows:

1. Design and build an optimal array of total field sensors
2. Design and implement real-time anomaly detection and analysis, yielding location, depth, and size of detected anomalies
3. Test, characterize, and optimize the performance of the system
4. Design digital electronics to drive the sensors

2.2 TECHNICAL APPROACH

Total field atomic magnetometer sensors are often used in making detailed surveys of underwater sites containing UXO. Such sensors are deployed on large wings or tow bodies. After the data are gathered, scientists typically plot and analyze them, then populate lists of targets requiring remediation. Finally, divers return to relocate these targets and the objects are removed. We previously developed atomic magnetometer sensors whose power consumption is an order of magnitude lower than previous magnetometers. This enabled the design of small systems containing two-dimensional sensor arrays that have the desirable sensitivity and noise immunity properties of existing magnetometer devices.

2.3 RESULTS

We achieved many significant results during this project. First of all, we determined an optimal number and positioning of sensors to be used to effectively analyze dipole targets in real time, using a reasonably sized array. Secondly, we developed a real-time inversion algorithm and user display interface allowing us to test the system. Thirdly, we built a prototype array of 12 sensors to test the performance of the system. Running the inversion algorithm in real-time worked as well as hoped, and actually yielded unexpected benefits in analyzing complex situations. Finally, in order to realize a portable system, we designed and built a miniature electronics system, using an all-digital approach to the signal extraction, which yields small size and low power consumption.

2.4 BENEFITS

With such an array of sensors, we now have the ability to run processing algorithms in real-time. This is made possible by gathering data suitable for dipole source analysis in a single instant. By analyzing the data on the spot, an instrument could guide a diver toward the target, while indicating the depth, size, and magnetic properties of the anomaly as well. Having such an instrument will allow the operator to better judge whether he or she has actually acquired the desired object, by comparing the object's size, depth, and signature with that obtained on the initial detailed survey. In addition, the diver may scan nearby to determine if other objects are present. Furthermore, after exposing the object under investigation, another scan may be performed to determine if additional UXO might have been misidentified as a single anomaly in the initial survey. Such an instrument could even eliminate the need for an initial detailed survey.

3 OBJECTIVE

3.1 SERDP RELEVANCE

This technology addresses MMSO-11-02: Improvements in the Detection and Remediation of Military Munitions Underwater. In particular, this project addresses the need for a tool to assist in discriminating or reacquiring small objects underwater.

Magnetometers are particularly useful in both the land and marine environments. Because they use the Earth's magnetic field as the transmitter, they have a large detection range. This is in contrast to Electromagnetic (EM) metal detectors, which have their own transmitter whose strength falls off rapidly with distance. However, individual magnetometer sensors have no directionality to locate the source of an anomalous field. Typically, therefore, surveys are undertaken to first gather the data, which is then plotted and analyzed later to determine the location of anomalies of interest.

Especially in the marine environment, this poses two challenges. First, there is considerable uncertainty in the positions of the readings. This not only makes it difficult to know precisely where an anomaly is located, but actually limits the accuracy of the discrimination algorithms as well. Secondly, once anomalies are chosen for remediation, reacquiring the location of the object of interest becomes a challenge. GPS systems do not work underwater, so the diver has a difficult task just knowing where he or she is. Therefore, having an instrument that can re-perform the target analysis done in processing the detailed survey data will be a huge advantage in the re-acquisition of anomaly locations.

The diver-held instrument would be able to provide the orientation, size, and depth of the magnetic dipole creating an anomaly. This could be directly compared with the parameters determined by processing the original survey data. Thus, the operator could be much more certain the object is the same one being sought, not a nearby piece of clutter. With a simple magnetic gradiometer, there is no way to determine this fact, and the diver might often miss the UXO by digging up a harmless piece of scrap nearby.

Such an instrument could even eliminate the need for a separate survey prior to remediation. With real-time analysis, a diver could potentially dig up targets of interest directly, possibly eliminating the need for separate detailed survey and remediation steps. As stated in the SERDP-ESTCP workshop in 2007, "The majority of participants agreed that real-time processing and classification of anomalies should be the ultimate goal because of the costs associated with detecting and reacquiring items in a dynamic underwater environment."

3.2 TECHNICAL OBJECTIVE

Our objective was to build, test, and characterize a diver-held ferrous metal detector array with real-time processing algorithms that accurately determine the position, depth, and size of magnetic anomalies from unexploded ordnance (UXO). We accomplished this by using an array of previously developed low power, miniature, total field atomic magnetometers. Specific technical objectives were as follows:

1. Design and build an optimal array of total field sensors
2. Design and implement real-time anomaly detection and analysis, yielding location, depth, and size of detected anomalies
3. Test, characterize, and optimize the performance of the system
4. Design integrated electronics to drive the sensors, yielding a light weight portable package with approximately 9-12 sensors

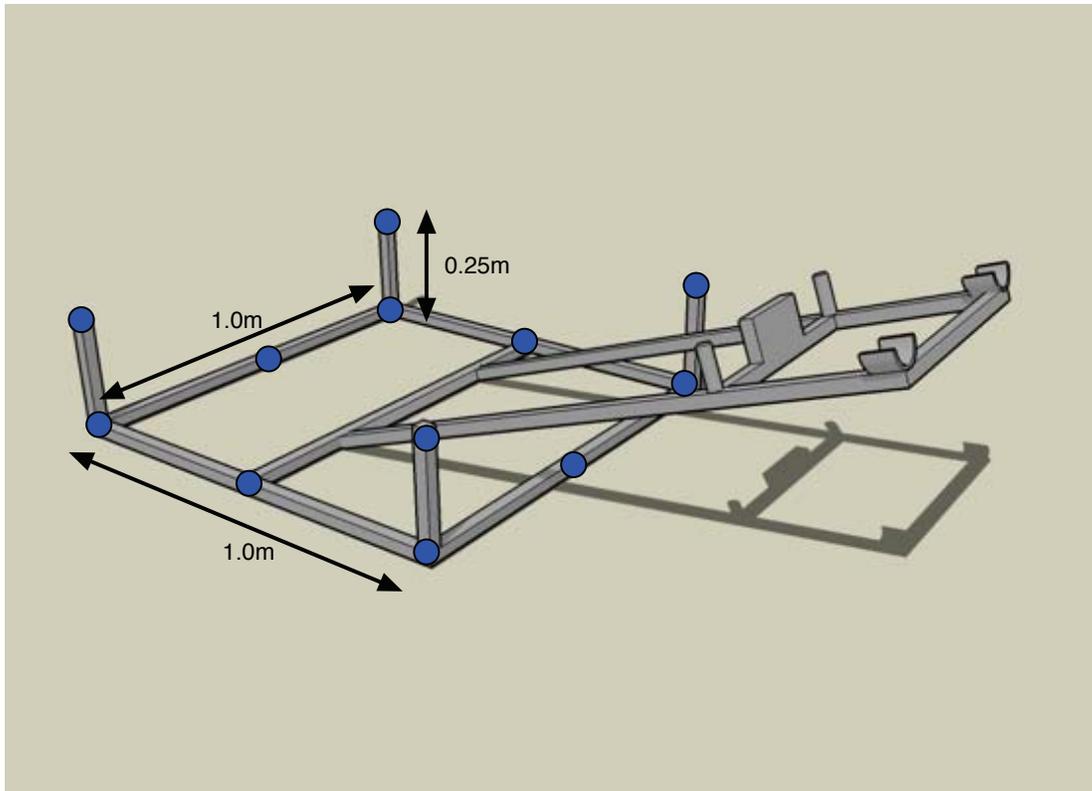


Figure 1. Concept of diver-held array with 12 sensors and real-time processing. 12 sensor locations are shown in blue.

Specific technical questions addressed in this report are as follows:

1. What sensor configuration will provide the optimal data set?
2. What optimization methods may be used in processing to provide the most accurate, fastest data analysis?
3. How small and lightweight may such a system be made?
4. What user interface would make the easiest to use system in an underwater environment?
5. Will performing a real-time dipole analysis yield more information to the operator than is possible in post processing?
6. How will integrated circuits improve the performance of the system beyond the obvious size reduction?

4 BACKGROUND

There are two main types of magnetometers in common usage today. Fluxgate magnetometers are low power and lightweight, and are used in wand-like gradiometers where an operator listens to sounds emitted by the device to locate anomalies in real time. In typical usage, the sensor is waved over a region while the operator attempts to determine the location of the object, which can be very confusing due to the vector nature of magnetic fields.

Squid magnetometers are sometimes used in arrays to measure the full tensor gradient of the magnetic fields. The complete vector components give useful information, but expensive cryogenics are required, and an array of total field measurement yields equivalent data.

However, total field magnetometers are very bulky, especially when considering the large battery packs required to power the sensors. Their advantage lies in their high sensitivity and the fact they measure the magnitude of the vector magnetic field, rather than a single component. These instruments are generally used in a radically different manner than fluxgates. With them, data is gathered over a region of interest, stored, and later mapped and analyzed. Only then can objects of interest be located, requiring the operators to return to the survey site to re-locate them.

Re-locating the anomalies presents a considerable challenge, especially in the underwater environment. Locations are very uncertain, and using only a magnetic gradiometer, the diver may easily acquire a nearby piece of scrap rather than the intended UXO identified in the original survey. If, however, an instrument were available that could perform the same anomaly analysis in real-time, this analysis could be compared to the original anomaly, and the diver would know for sure which target was the intended one. Conceivably, the process could allow for the initial mapping survey to be eliminated entirely. The diver could identify, interrogate, locate and remediate anomalies in one pass.

In order to perform this process in real-time operator's require a two- or three-dimensional array of magnetic sensors. However, finding suitable sensors has been a problem until now. Atomic magnetometers have the best magnetic field sensitivity of the types of magnetometers used. In addition, they are not sensitive to the sensor orientation relative to the vector magnetic field, giving them a large advantage over fluxgate and Giant Magneto-Restrictive (GMR) sensors. However, previous generations were extremely bulky and consumed a lot of power, which made an array of such devices impractical. In previous work, however, Geometrics with the support from the Strategic Environmental Research and Development Program (SERDP) developed low power atomic magnetometers. The project, MM-1512, won SERDP's Munitions Management Project of the Year Award for 2008. In this latest work, we utilized these sensors in an array capable of locating and analyzing magnetic anomalies automatically and in real time.

Linear arrays of magnetometer sensors have been deployed in many scenarios in unexploded ordnance (UXO) detection. These have an advantage over single sensors in their efficiency of gathering data, but still require data to be collected over time as the array is scanned past or over a target. This leads to various positioning and time changing errors, which limits the accuracy of the device. However, linear arrays have been the only feasible approach with the bulky sensors previously available.

There are four key areas where we made significant improvements in this project over previous generations of technology. On the sensing side, we replaced existing large, bulky sensors with low power, miniature sensors of similar (even better) sensitivity. On the electronics side, we converted large printed circuit boards to much smaller digital circuitry. On the array design, we moved from the linear arrays used in the past to scan over a region, to using 2-dimensional arrays to provide a series of complete pictures of the target of interest. Finally, we adapted algorithms used only in post processing of data in order to utilize them in real-time to provide an immediate indication of the properties and location of the target.

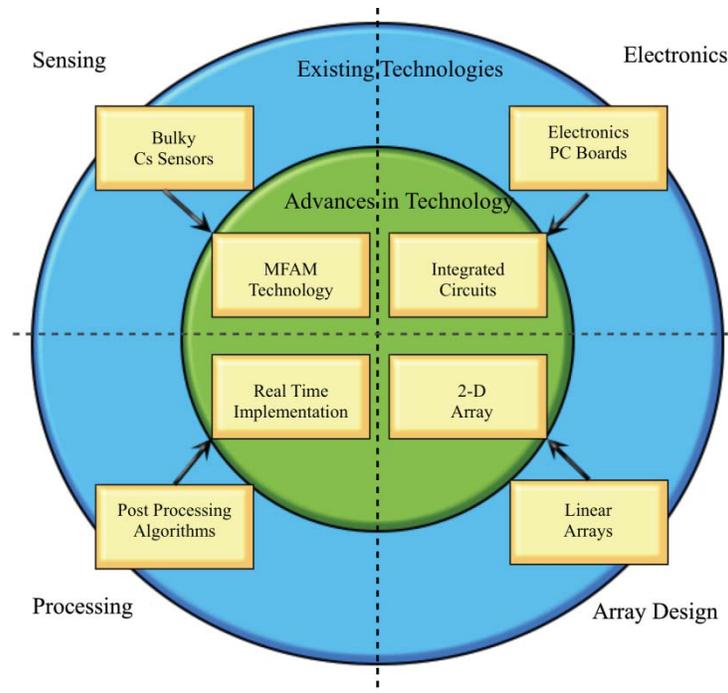


Figure 2. Technology roadmap indicating advances in four areas with the proposed system.

5 METHODS

5.1 SENSORS

Atomic magnetometers are the sensors of choice for ferrous metal detection. Such sensors are extremely sensitive (a few pico Tesla noise levels) and are insensitive to the orientation of the sensor direction relative to the vector magnetic field. Operators often try other technologies, such as magneto-restrictive devices and fluxgates¹ because of their much lower power consumption relative to atomic magnetometers. However, they often suffer from low performance due to their lower sensitivity, offset and drift problems, and orientation-dependent measurement. Even tiny vibrations create large enough noise to limit the sensors' usefulness.

Existing commercially available atomic vapor magnetometers (e.g., Geometrics G-858 sensors) are commonly used in Unexploded Ordnance (UXO) detection because of their sensitivity in spite of their large size and power consumption. Large batteries are required to power the system, which greatly increases the complexity of fielding these devices. By replacing the cesium (Cs) lamp with a low-power laser diode, and reducing the size of the Cs cell using micro-fabrication (MEMS) techniques, we reduced the power consumption by an order of magnitude. Previous research in creating such Micro-Fabricated Atomic Magnetometer (MFAM) sensors evolved from considering the feasibility of the devices to allowing detailed designs considering various tradeoffs and performance considerations (Prouty, 2007)². That work had been achieved through a collaboration between the National Institute of Standards and

¹ Munsch, et al, "Magnetic mapping for the detection and characterization of UXO: Use of multi-sensor fluxgate 3-axis magnetometers and methods of interpretation," *Journal of Applied Geophysics*, 61, 168-183, 2007.

² Prouty, M., 2007, *Progress in Chip-Scale Total Field Magnetometers, Battlefield Acoustic and Magnetic Sensors*, Laurel, MD, Aug 2007, BC01

Technology, Sandia National Laboratories, and Geometrics, Inc., supported by the Strategic Environmental Research and Development Program (SERDP).

These sensors achieve their performance advantages in three distinct ways. First of all, the Cs discharge lamp used in traditional sensors has been replaced with a laser diode light source. This greatly reduces the power consumption of the devices. However, this required research in the design of, and investment in the production facilities for, laser diodes which emit light with low amplitude and frequency noise and at the exact required wavelength. Secondly, we developed the techniques for manufacturing sensing cells of a much smaller size were, which further reduces the power requirements in heating the cell to the required operating temperature. Finally, we developed new techniques for polarizing the atoms in the Cs cell, and measuring the atomic precession frequency which increased the bandwidth and recovered the sensitivity, which otherwise would have been lost to the smaller cell volume³. These advances are now being commercialized by Geometrics.

With this new small size and low power consumption per sensor, a highly useful magnetometer system may be deployed, taking advantage of the ability of magnetometers to sense ferrous material at large distances and providing a series of snapshots with enough information to instantly determine the location and properties of the target.

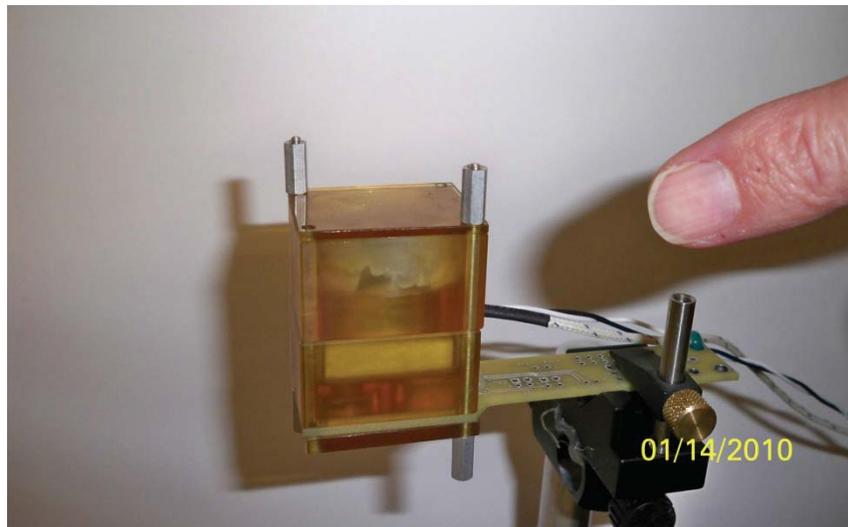


Figure 3. Photo of an early prototype MFAM sensor.

There are a number of potential advantages of using an array of MFAM sensors, as opposed to other technologies:

- 1) Small size, weight and low-power consumption
- 2) Highly stable readings, with no offset or drift
- 3) Readings immune to sensor orientation and vibration (as they are total-field sensors)
- 4) Sensitive enough to gather sufficient information to determine the location and properties of the magnetic anomaly being sensed in real-time

³ Schwindt, et. al., "A chip-scale atomic magnetometer with improved sensitivity using the Mx technique," Applied Physics Letters, 90, pp 081102, 2007.

- 5) No need for bulky and expensive cryogenics for cooling the sensors
- 6) A sufficiently large array of measurements may be made to characterize the location and properties of the target

Making an array of this type of sensor is much easier than making an array of vector sensors such as fluxgates⁴ or squids. Using those sensors, the platform must be highly rigid because of the orientation-dependent nature of the readings from vector sensors. Even acoustic vibrations produce enough angular change to induce significant noise. Total field sensors, on the other hand, may be freely rotated without changing their readings. Thus, the platform need not be rigid, and may even be something such as a flexible mat that could be laid on the ground.

In addition, the sensitivity of atomic magnetometers greatly exceeds that of fluxgates, and their readings do not drift with temperature. Operators using fluxgate sensors to analyze an anomaly must carefully measure and subtract offsets. Atomic magnetometer sensors do not have this difficulty.

While the basic sensor work had been completed previously, we needed to perform further work to miniaturize the electronics circuitry, which is required to extract the magnetometer signal from the Cs atoms in the atomic cell. Previously, the electronics had been reduced in volume by a factor of 3 over traditional commercial sensors. However, further reductions in size were required to do justice to the small size and power consumption of the sensing elements themselves.

In addition, we identified some very interesting interrogation techniques in this and earlier projects. The digital approach followed in this work enables the complicated functionality of the electronics to be performed more efficiently, with better performance, and with fewer numbers of devices achieving lower power and smaller size. This approach also allows considerable flexibility and upgrade of the features and performance of the devices, simply by downloading new firmware.

5.2 PROCESSING ARRAY INFORMATION

In this project Geometrics utilized our data processing and analysis experience⁵ to adapt those techniques to a real-time methodology. Techniques that were commonly used in post processing were tested using an array in real-time.

Typically, the method for processing magnetometer information is to calculate the dipole source that best reproduces the measured field readings over a certain region above the anomaly^{6,7}. This is always done after the fact, however. The data is first gathered and only later is it processed. The processing further comprises several steps. First, maps are plotted as shown in Figure 4 to identify anomalies that are likely to be objects of interest. This is typically done visually. Then, a region of data comprising the anomaly is selected. Next, an automatic optimization algorithm is used that computes the position, size, and orientation of a dipole magnetic source that best represents the data that is observed by the magnetometer.

⁴ Wiegert, U.S. Patent 6,841,994, "Magnetic anomaly sensing system" and SERDP project 1511.

⁵ Tchernychev and Snyder, "Open source magnetic inversion programming framework and its practical applications," *Journal of Applied Geophysics*, 61, p184-193, 2007.

⁶ Billings, "Discrimination and classification of buried unexploded ordnance using magnetometry," *IEEE Transactions on Geoscience and Remote Sensing*, v42, p1241-1251, 2004.

⁷ Bell, T., DeProspero, D., Prouty, M. MagAID: PC Based target characterization software for use with total field magnetometer survey data, SAGEEP, 1996.

Figure 5 shows the process as displayed by the software typically used to analyze anomalies. Plots of observed data, calculated data from the dipole chosen by the optimization algorithm, and the difference between the two are typically shown as 2D plots and in profiles. These graphs are typically produced as part of the processing for magnetic anomalies.

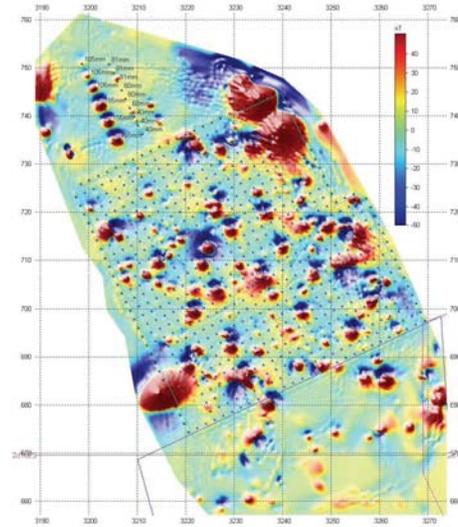


Figure 4. Magnetic anomaly map.

One of the greatest sources of error in this method is positioning errors in the readings and errors arising from the fact that the readings were taken at different times. The background fields vary with time, interference from other moving objects, varying power line influence, varying sensor orientation and position with respect to the operator. All this adds significantly to the error budget, and is what ultimately limits the performance of the location algorithms. Performing this data analysis with a single snapshot of data taken at the same moment greatly increases the accuracy of the results.

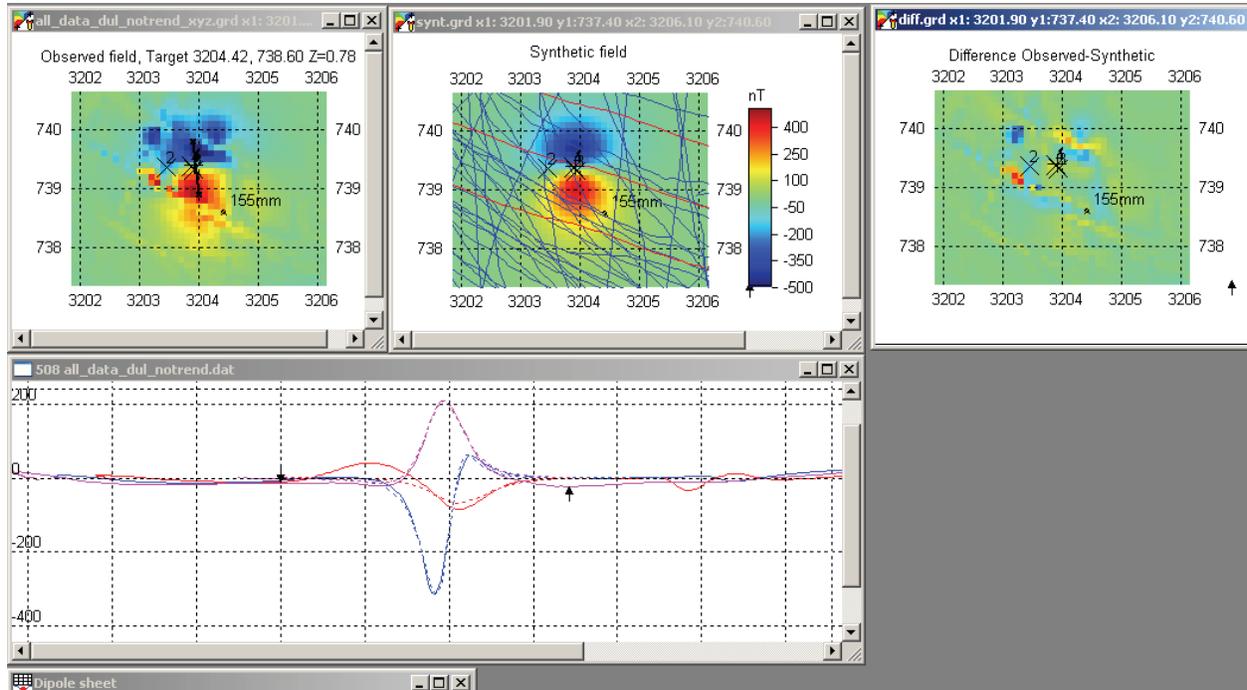


Figure 5. Analyzing anomalies by calculating dipole source that best fits the observed data. We propose an instrument performing these calculations automatically and in real time, using an array of sensors.

Therefore, we implemented this process in real time, using the simultaneous readings acquired from the array of sensors. There were several issues to be investigated in implementing this process in real time, including the speed of the algorithm, lag or delay in the resultant feedback to the operator, and the speed with which the sensor platform may be moved or, most significantly, rotated. The calculation speed of an optimization algorithm depends heavily on the speed of calculating the observed response of each sensor due to a dipole in a particular location. This must be done for the position of each sensor of the array. This lends itself to rapid parallel processing – the job of calculating the response at each sensor may be apportioned amongst several processors.

Such an array can be used for standalone detection of anomalies, or for reacquiring anomalies that have been identified in a high-resolution survey done previously using existing methods (such as MTADs and Marine MTADS systems). Operators using the array instrument would gather enough data to make an accurate comparison between the features of the previously identified target and the one being re-acquired in the remediation phase. Such features include the apparent magnetic size, degree of remnant and induced magnetization, and depth. Measuring these attributes with the portable instrument gives operators much greater confidence that the anomaly they are removing is the one that the data analyst intended when looking at the original survey data.

6 RESULTS: ARRAY DEMONSTRATION

6.1 OVERVIEW

Steps in this work, as indicated in Figure 6.

1. Design Sensor Array
 - a. Simulate sensor array to optimize array performance
2. Build Prototype Sensors
 - a. Build a number of prototype sensors for testing in array
3. Build Sensor Array
4. Develop Real-Time Processing
 - a. Modify post-processing algorithms to run in real-time and write real-time GUI
5. Test and Characterize Prototype Array
6. Electronics Design
 - a. Model sensor physics and electronics requirements
 - b. Design laboratory bench top brass boards
 - c. Design and build small form-factor final design

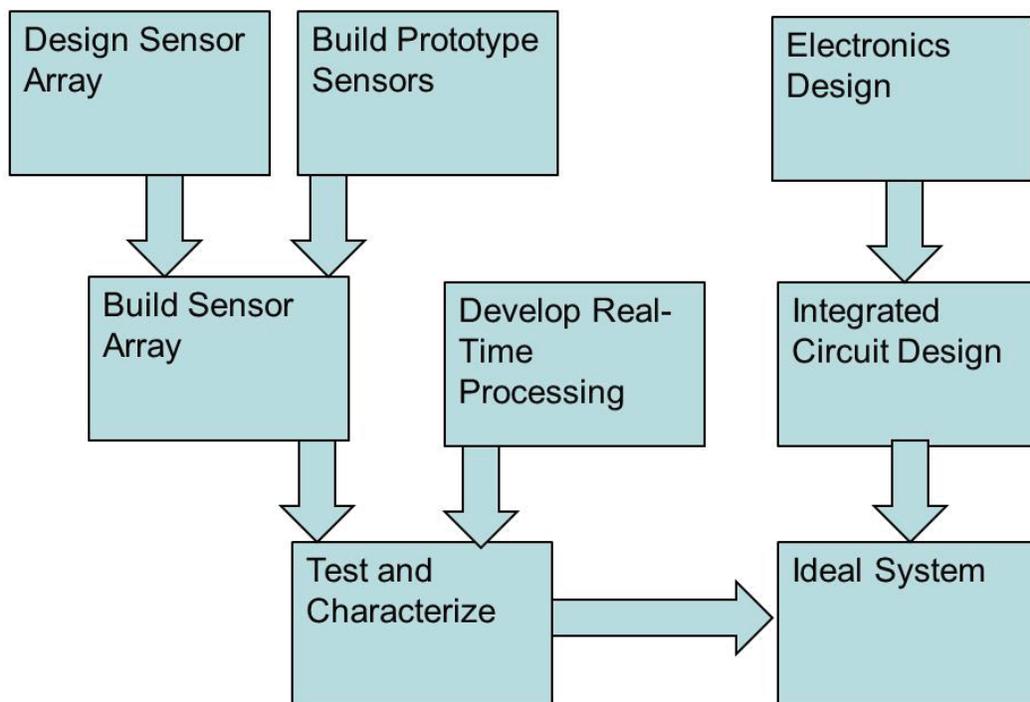


Figure 6. Steps of this project.

6.2 DESIGN SENSOR ARRAY

Our first step in building a real-time array was to prepare simulations of sensors in various positions. In order to be sure the instrument we built would perform usefully, we simulated data from chosen sensor

positions over various targets. We added noise of various sorts, both random electronics noise as well as geologic noise and signature from scrap metal to the simulation. This allowed us to analyze the performance of the array before building the actual array.

6.2.1 PREPARING MODELING PARAMETERS

In order to prepare our simulation algorithms for accurate simulations of real-world scenarios, we first gathered data using a single sensor over various targets to understand the signals and gradients that would be encountered due to geologic noise or distant objects.

6.2.1.1 MODELING FIXTURE AND MEASUREMENT PROCEDURE



Figure 7. Modeling fixture for accurately positioning sensor for reading above target.

Our modeling fixture consists of a non-magnetic plastic cart with two foam plates 1.24x1.24 meters installed on top of it, separated by 0.25 m vertically. Each plate had 25 holes separated by 0.2 m. The holes were used to position a commercial magnetometer sensor for accurate and repeatable positioning.

A 3-axis fluxgate magnetometer was mounted 0.5 m from the plate to measure the orientation of the array.

To avoid local time disturbances a second magnetometer was used as a base station. It was located about 7 m from the measurement plates and synchronized with the measurement sensor, allowing effective

removal of time variations in the magnetic field. Because all measurements using a real-time array will be taken at the same time, this second sensor will not be needed in the actual array. However, it was required for model measurements.

The operator sequentially moved the sensor into each of the 50 holes and recorded the data at a sample rate of 10 Hz for approximately 5-10 seconds at each position. The difference between two sensors was averaged for each location and represents the magnetic field value for this position.



Figure 8. Base station sensor with modeling fixture in background.

During data processing, the data was decimated in order to assemble different possible configurations, such as 18 sensors, 25 2-D sensors, etc.

The test site was located behind Geometrics' building, in a small, currently unused walnut orchard. Parking lots surround it, with the closest one about 25 m from the fixture. Our instrument detected cars moving in the parking lots, and we took care to ensure that measurements were not collected during these events.

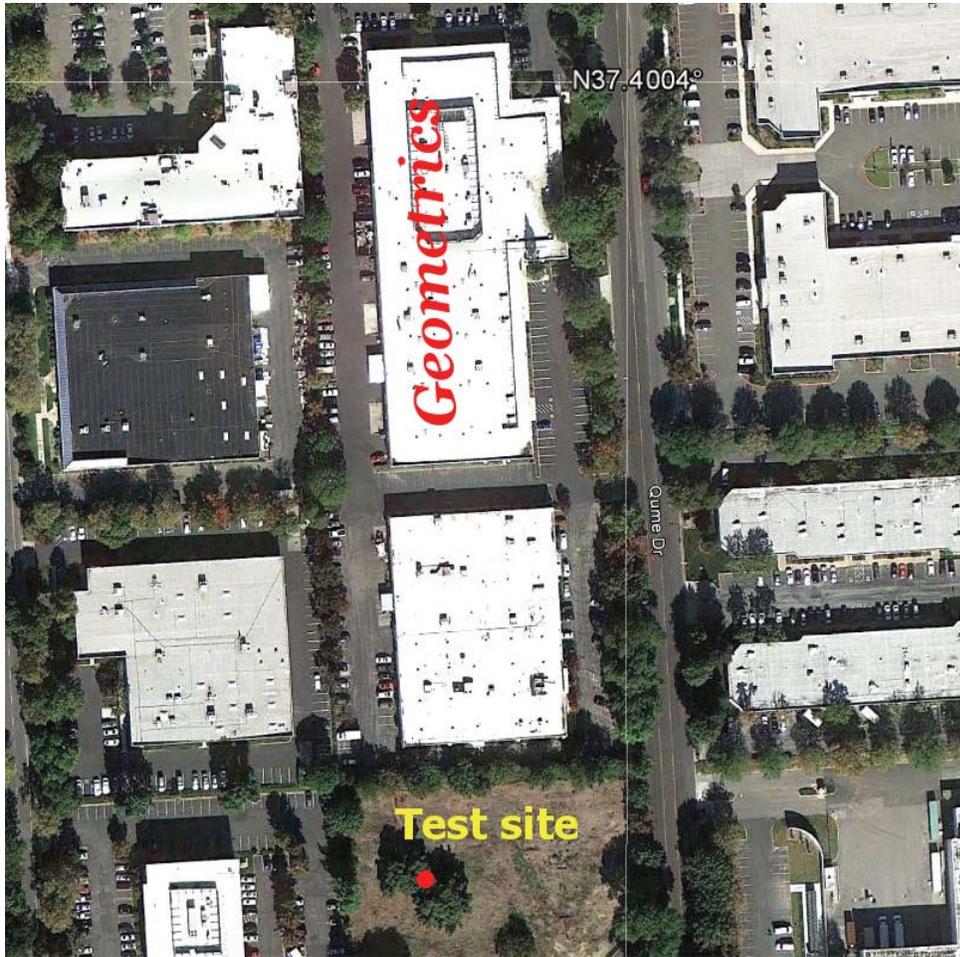


Figure 9. Google Earth image of test site.

6.2.1.2 MEASURED OBJECTS

We measured the following objects during this experiment:

1. Metal ball, diameter of 10 cm, weight of 3.86 kg (8.5 lb.). Based on these measurements, it appears that the object is weakly magnetic. It is possible that the filling is a non-magnetic material. This object is shown in Figure 10 A.
2. 80 mm mortar shell, weight of 3.86 kg (8.5 lb.), dimensions are 45x8 cm at the widest place. The shell is shown in Figure 10 B.
3. Strong Neodymium magnet, with magnetic moment, which corresponds approximately to a 100-200 lb. object. The specifications of the magnet are shown in Figure 10 C.

Each object was placed at different orientations and locations relative to the measurement plane, with depth (distance from the sensor plane) ranging from 1m or less.



A

B



C

Description

- Dimensions: 1" dia. x 7/8" thick
- Tolerances: ± 0.004 " x ± 0.004 "
- Material: NdFeB, Grade N42
- Plating/Coating: Ni-Cu-Ni (Nickel)
- Magnetization Direction: Axial (Poles on Flat Ends)
- Weight: 2.98 oz. (84.5 g)
- Pull Force, Case 1: 58.15 lbs
- Pull Force, Case 2: 60.95 lbs
- Surface Field: 5730 Gauss
- Max Operating Temp: 176°F (80°C)
- Brmax: 13,200 Gauss
- BHmax: 42 MGOe

Figure 10. Objects used. A: 10 cm metal ball; B: 80mm mortar shell; C: Neodymium magnet.

6.2.1.3 MAGNETIC SIGNATURES OF THE AREA AND THE OBJECTS

We surveyed the area with a Geometrics G-858 cesium magnetometer to ensure the absence of strong native anomalies. Background measurements, without a test object, were routinely performed as part of every experiment. Typical magnetic signatures are shown in this section to demonstrate the average magnetic field measurements obtained by a conventional survey. These plots are semi-analytic and we obtained them using the following procedure:

1. Using the dipole model, we performed magnetic field inversion for all 50 measurements. In cases where there is no target measurement, we instead used a linear model.
2. Calculate the magnetic field from the estimated target over the larger synthetic grid of 5 x 5 m around the array.

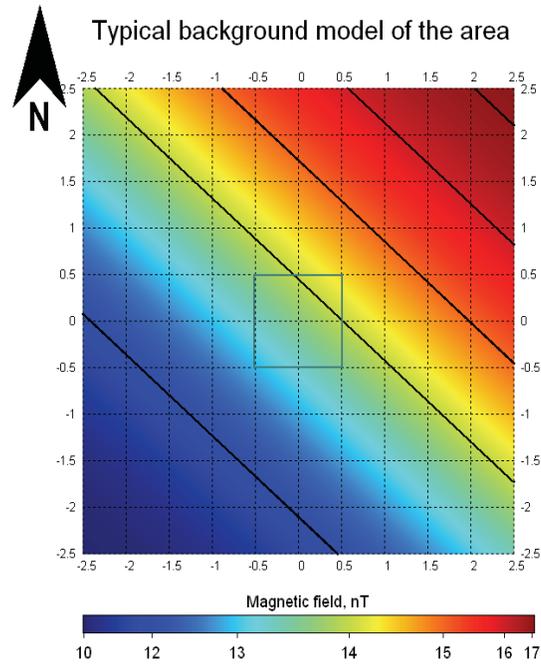


Figure 11. Typical background for the area. Note that size of the computed field is 5 x 5 m while measuring frame is only 1x1 m (shown in the center). Magnetic field changes from 13 nT in the left bottom of the frame to 14.4 nT in the right top corner.

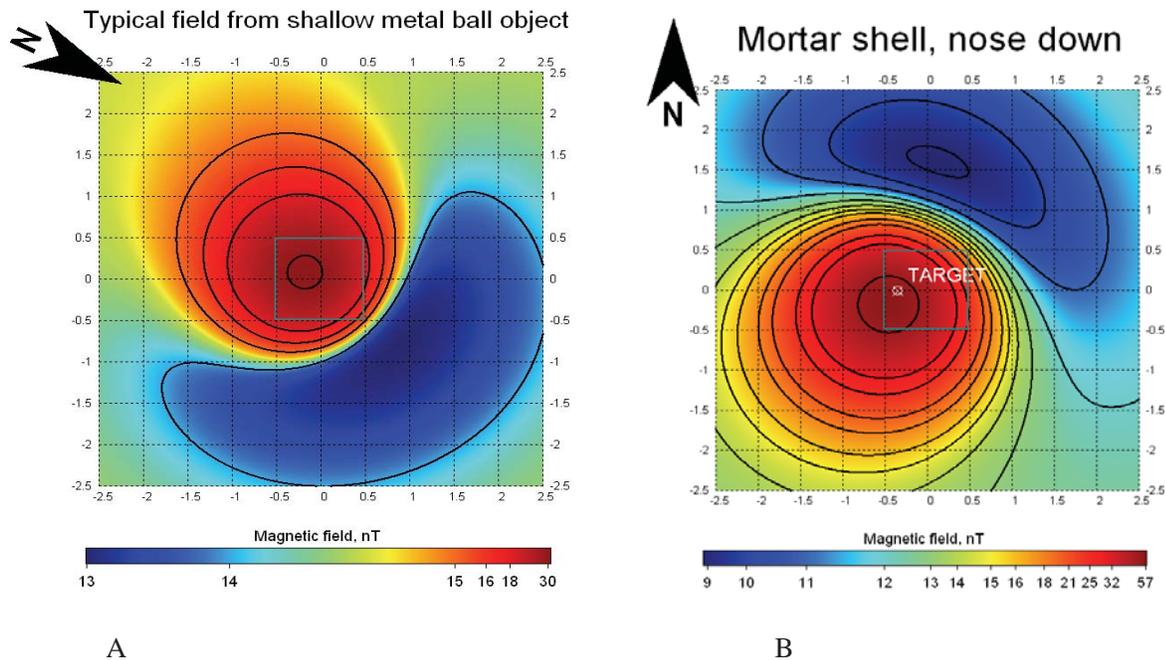


Figure 12. A (left): The field variation for the metal ball inside the figure is 14.5 to 30 nT . B (right): Mortar shell field variation inside the figure from 16 to 57 nT. Note the 1m measurement frame (rectangle) is a small portion of the figure.

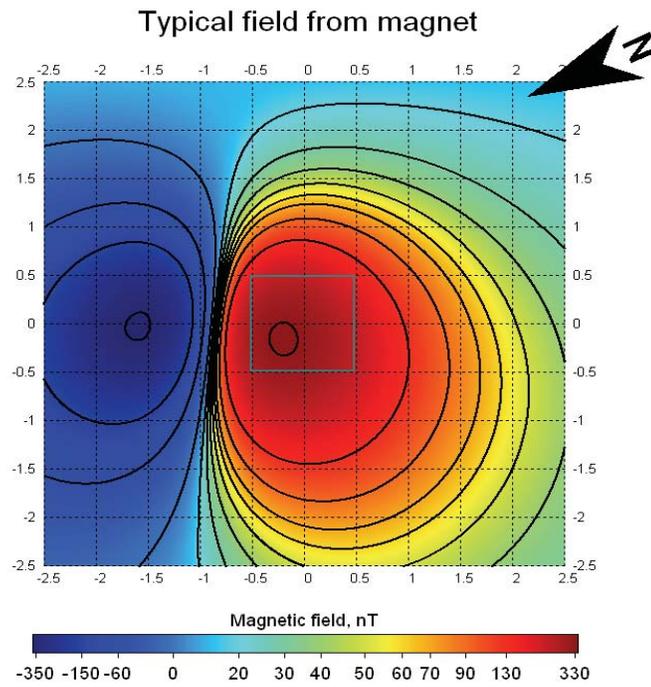


Figure 13. The field variation of the Nd magnet inside the figure varies from -146 to 330 nT.

6.2.2 SIMULATION RESULTS OF SEVERAL ARRAYS

Figure 14 shows the various potential array designs that we tested in our modeling experiments. We duplicated some of the geometries with different grid sizes, giving a total of ten different configurations.

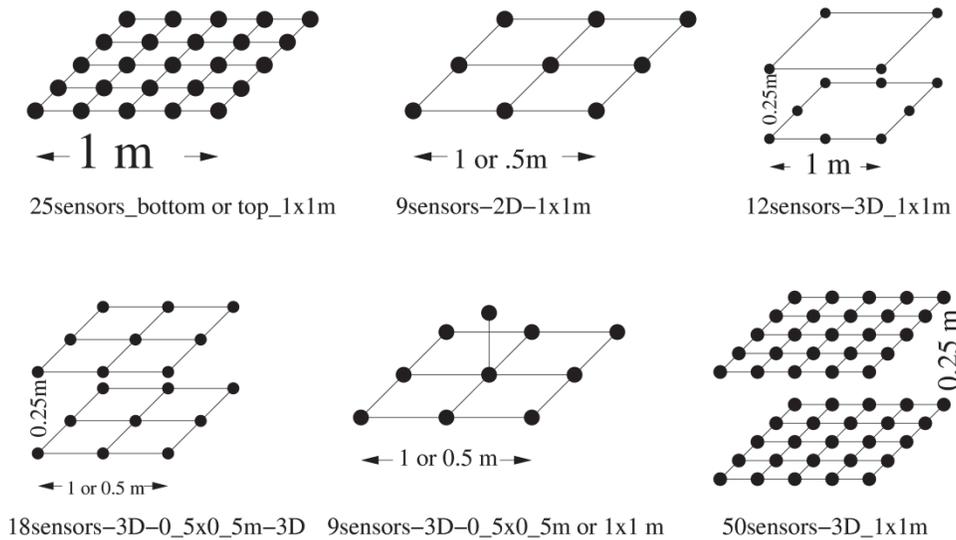


Figure 14. Different array configurations and their names used in the measurements.

We created a total of 11 measurements sets using various objects, which differed by changing the object type and its location relative to the frame. Typically, the object was not located right in the middle of the frame, but off center, and with the depth (distance from the sensors) ranging from 1 to 1.5 m. The summary of all measured cases is the following:

- 3 measurements using the metal ball, at a depth of ~1m.
- 3 measurements with the mortar shell, with nose down, up, and sideways with depth around ~1.2 m
- 5 measurements with the strong magnet, with depth ranging from 1 to 1.5 meters

These 11 measurements are not enough to well represent the statistical nature of the inversion. However, it allows us to draw some conclusions as to what type of accuracy a user could expect from the proposed device, assuming a “lock” on the target is obtained.

To analyze the results, the differences between the actual known positions of the object and the estimated positions given by the fitting algorithm are computed and plotted as histograms. Here we present histograms of the total, horizontal and vertical errors separately.

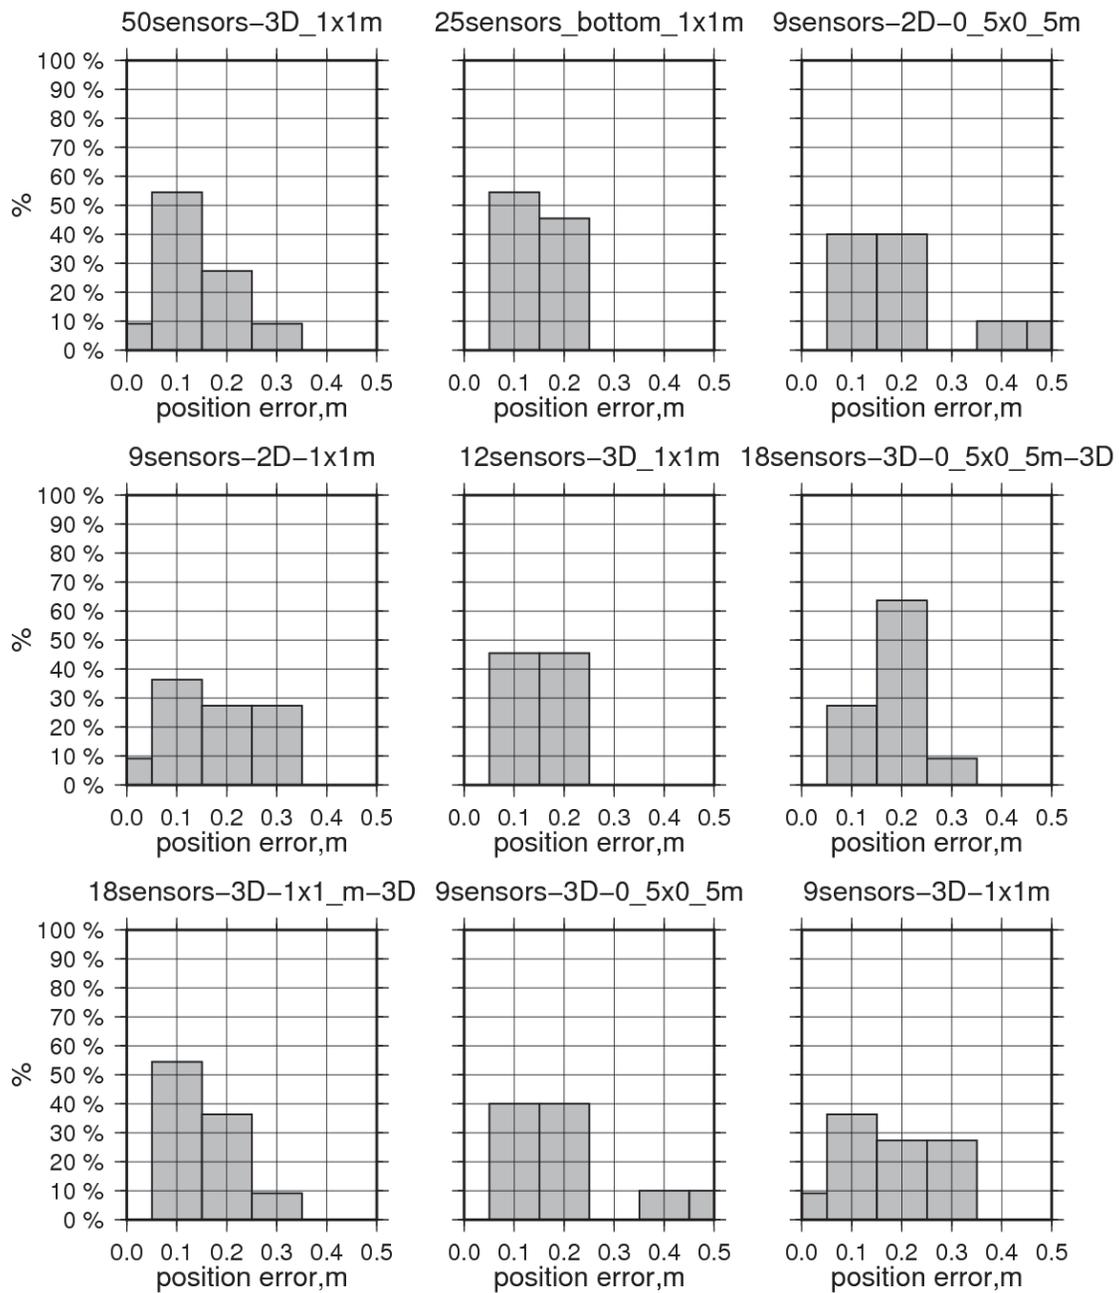


Figure 15. Overall position error for all tested configurations. The title denotes sensor count and geometry.

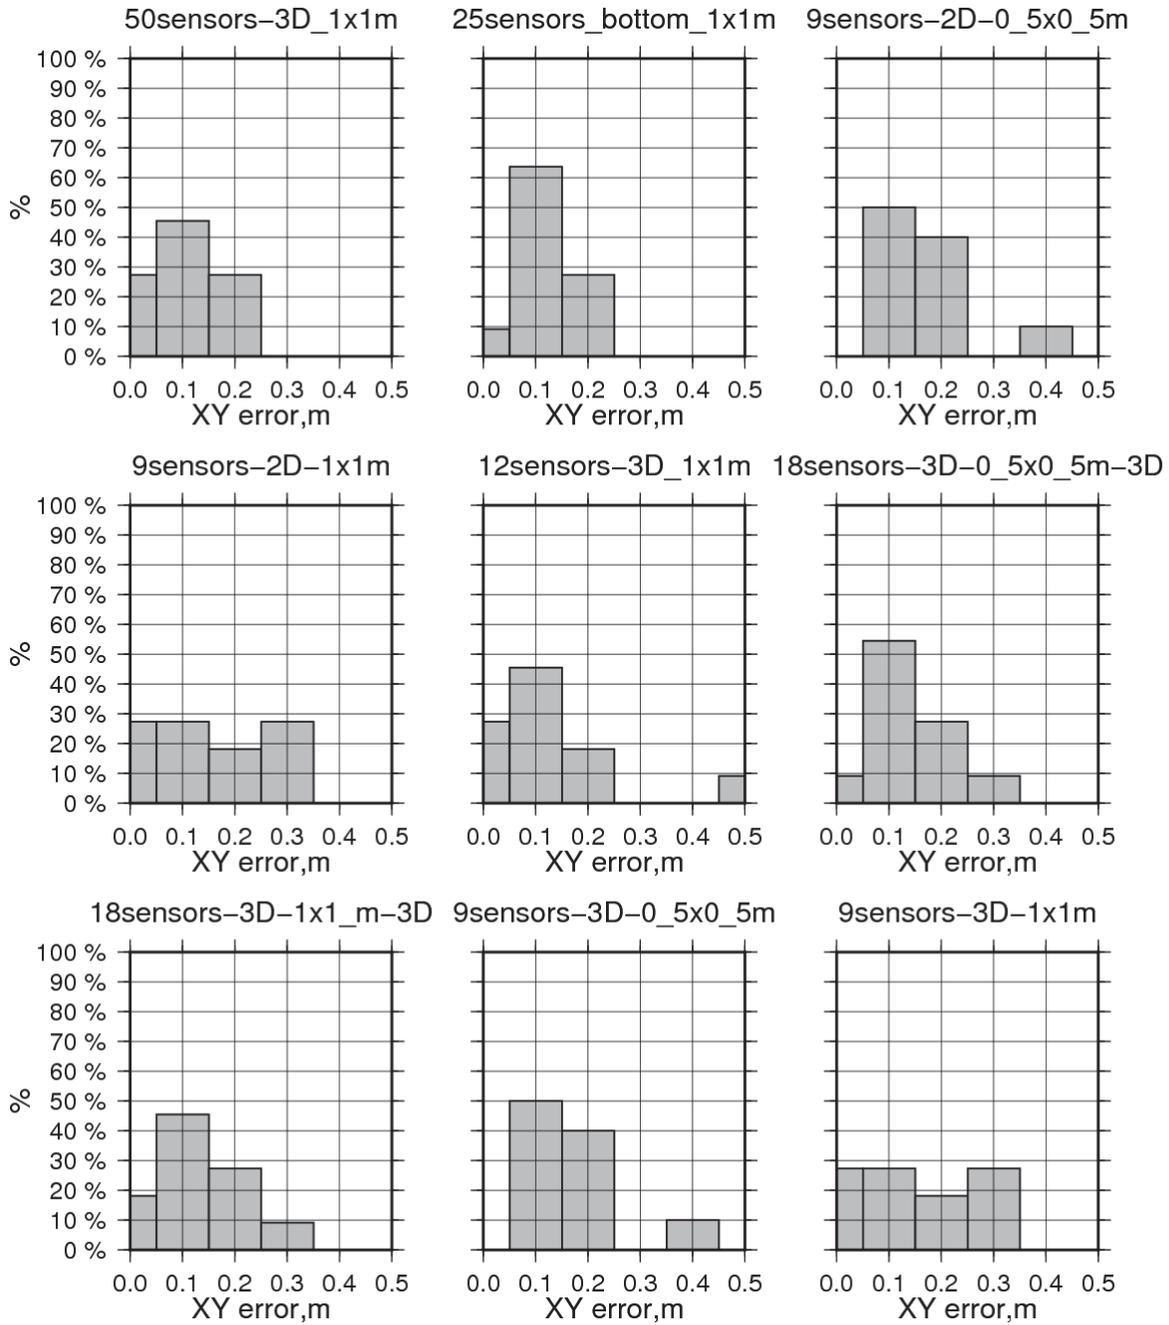


Figure 16. Horizontal position differences for all tested configurations.

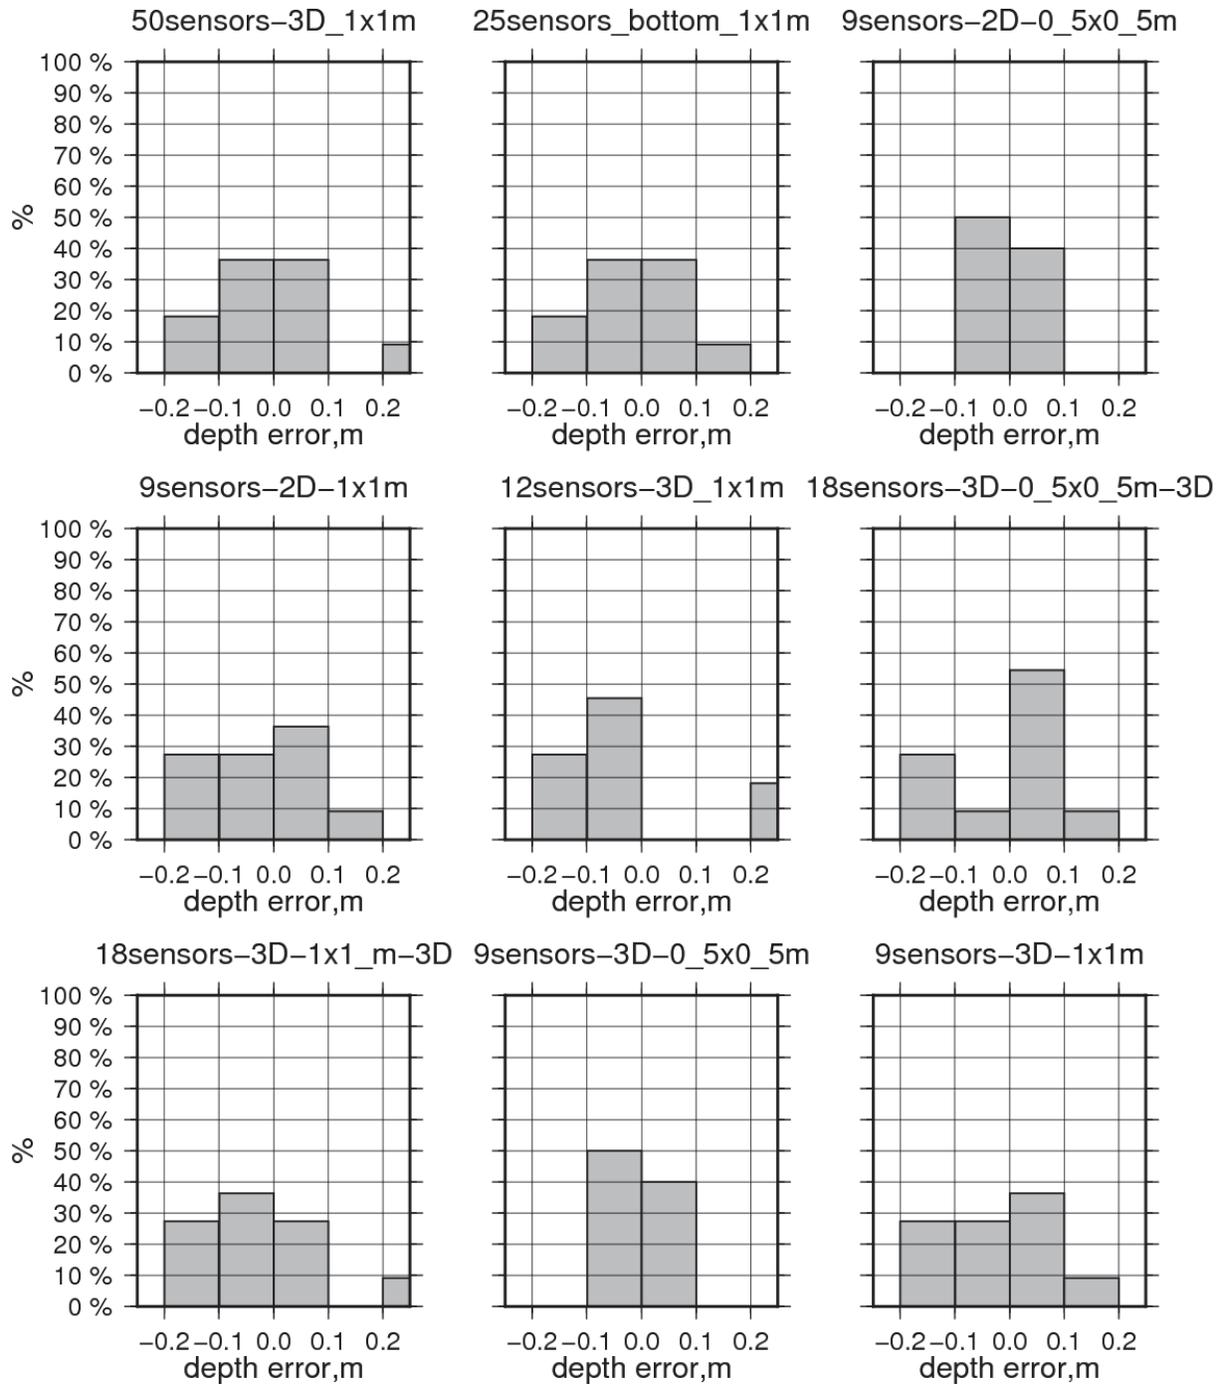


Figure 17. Depth error distributions for tested configurations.

Figure 15 and Figure 16 show typical distributions for 9 tested configurations (10th configuration, 25 sensors of top layer, is omitted because it is very similar to 25 bottom configuration). For actual geometry of the array, please see Figure 14. “9sensors-3D” configurations denote 9 sensor in plane + 1 sensor vertical in the middle. For “2D” configurations a DC term was estimated as part of the inversion, for “3D”

configurations the vertical gradient was estimated as well. Figure 17 shows the histograms for the depth error distribution.

The results can be summarized as follows:

1. As a rule of thumb, users could expect accuracy of 0.2 m or better while locating different objects at depths ranging from 1 to 1.5 m. It assumes that the system has locked on the target, i.e. inversions converge to a similar location regardless of the array motion. It is possible that the user could improve the survey's accuracy by properly locating the array relative to the target; it was not possible to test this idea in these simulations because of the considerable time involved into making a single measurement.
2. Apparently, there is no particular configuration that gives a considerable advantage in accuracy. For instance, comparing the full 50-sensor array with 18 sensors or even 12 sensors only showed slight, if any, improvement. Thus, we can utilize smaller arrays without restricting the performance of the system.
3. The most important factor appeared to be the spatial extent of the array. The smaller, 0.5m configurations appear to be marginal and should be avoided.

6.2.2.1 BACKGROUND FIELD

We investigated the question of how the system would behave in the absence of any magnetic object, and whether the operator would be able to know that no "lock" was obtained. To investigate this question, we performed a total of 4 background measurements at different times. The typical background model would be linear, as it is presented in Figure 18. However, to simplify the modeling and reduce the number of required sensors, we assumed that background field can be modeled by DC for 2D and DC and a vertical gradient term in case of 3D configurations.

It is possible to see that some results appear to be realistic, such as for 50 sensors ($X=0.2$, $Y=0.26$, $Z=1.95$). However most of the results show a huge magnetic moment and / or considerable standard variation for the estimates. These could be used to reject erroneous values.

The meaning of the columns in Figure 18 through Figure 21 are the following:

- $X(m)$, $Y(m)$, $Z(m)$ – estimated location;
- $D(m)$, $DH(m)$, $DZ(m)$ – difference from the true known location;
- $BkgC(nT/m)$ and $BkgD(nT)$ – estimated vertical gradient and DC term;
- J and $std.J$ – estimated magnetic moment and its standard deviation, cgs units;
- $max_fit(nT)$ maximum difference between observed and modeled field;
- $Fit(nT)$ RMS fit of model and measured data;
- inc , dec – estimated inclination and declination of the dipole in the coordinates of the frame, degrees;
- $angle$ – angle with Earth's magnetic field, degrees

Configuration	X(m)	Y(m)	Z(m)	std.X(m)	std.Y(m)	std.Z(m)	kgC(nT/m)) BkgD(nT)	J(cgs)	std.J(cgs)	av_fit(nT)) Fit(nT)	amp(nT)
25sensors_bot tom_1x1m	7.24	-1.12	3.65	27.2	6.99	15.68	0	3.62	1.36E+05	8.81E+05	0.21	0.11	2.15
25sensors_top _1x1m	51.78	3.39	33.14	1439.94	157.59	907.86	0	107.93	3.79E+09	3.05E+11	0.4	0.2	1.48
9sensors-2D- 0_5x0_5m	1.31	-0.44	6.26	87.51	64.71	424.71	0	-11.68	5.45E+04	4.95E+06	0.08	0.04	0.85
9sensors-2D- 1x1m	7.18	-1.44	10.27	1016.37	417.09	645.85	0	-6	6.16E+05	1.74E+08	0.23	0.13	2.15
12sensors- 3D_1x1m	0.18	0.25	1.82	0.26	0.31	0.78	3.76	13.57	1.07E+02	9.80E+01	0.27	0.13	2.55
18sensors-3D- 0_5x0_5m-3D	-8.35	25.61	3.46	114.22	337.49	48.6	-27.02	-26.52	7.75E+07	2.67E+09	0.28	0.13	1.9
18sensors-3D- 1x1_m-3D	0.18	0.07	1.77	0.26	0.25	0.77	3.25	13.19	9.50E+01	8.70E+01	0.37	0.17	2.79
50sensors- 3D_1x1m	0.2	0.26	1.95	0.26	0.29	0.64	2.9	12.87	1.13E+02	8.50E+01	0.41	0.21	2.79
9sensors-3D- 0_5x0_5m	1.31	-0.44	6.26	87.53	64.73	424.81	-6.2	-10.92	5.45E+04	4.95E+06	0.08	0.04	1.6
9sensors-3D- 1x1m	7.18	-1.44	10.27	1019.8	418.44	647.55	-1.01	-5.89	6.17E+05	1.75E+08	0.23	0.12	2.23

Figure 18. Results of the inversion in the absence of a target.

Figure 19 shows a similar table for the situation where an object is present. One can easily compare the values to determine whether an object actually exists.

Configuration	X(m)	Y(m)	Z(m)	std.X(m)	std.Y(m)	std.Z(m)	kgC(nT/m)	BkgD(nT)	J(cgs)	std.J(cgs)	max_fit(nT)	Fit(nT)	amp(nT)
25sensors_bottom_1x1m	-0.32	0.04	1.38	0.02	0.01	0.02	0	13.63	658	24	0.5	0.15	41.21
25sensors_top_1x1m	-0.33	0.03	1.37	0.05	0.03	0.05	0	12.86	649	41	0.22	0.11	23.13
25sensors-2D-1x1m	-0.32	0.04	1.38	0.02	0.01	0.02	0	13.63	658	24	0.5	0.15	41.21
9sensors-2D-0.5x0.5m	-0.16	0.07	1.23	0.05	0.02	0.1	0	22.74	419	51	0.1	0.06	27.39
9sensors-2D-1x1m	-0.33	0.03	1.38	0.03	0.02	0.03	0	13.28	665	19	0.07	0.04	39.76
12sensors-3D_1x1m	-0.33	0.03	1.38	0.03	0.02	0.03	1.88	13.11	657	19	0.12	0.05	39.76
18sensors-3D-0.5x0.5m-3D	-0.23	0.07	1.36	0.04	0.02	0.05	5.28	16.29	625	43	0.19	0.1	30.91
18sensors-3D-1x1_m-3D	-0.34	0.03	1.39	0.03	0.02	0.03	1.54	12.79	673	20	0.18	0.07	39.76
50sensors-3D_1x1m	-0.33	0.04	1.39	0.02	0.01	0.02	1.54	12.93	677	18	0.48	0.14	41.21
9sensors-3D-0.5x0.5m	-0.16	0.07	1.23	0.05	0.02	0.1	10.81	21.39	419	51	0.1	0.05	27.39
9sensors-3D-1x1m	-0.33	0.03	1.38	0.03	0.02	0.03	1.81	13.06	665	19	0.07	0.03	39.76

Figure 19. Results of inversion for the mortar shell in the nose down configuration.

6.2.2.2 AMBIGUITY BETWEEN ESTIMATED OBJECT AND BACKGROUND

One of the most difficult problems revealed during measurements and modeling is an ambiguity between the estimated parameters of the object and the estimated background. It appears that the solution is not unique in terms of explaining the observed field: the bulk of the field can be explained by the dipole object (which is desirable) or by large background DC term, which considerably distorts the estimated dipole parameters.

The physical nature of this phenomenon can be seen by looking at pictures similar to Figure 12. As indicated there, the sensor array frame is measuring only a small portion of the extent of the anomaly, compared to the entire dipole signature measured by a conventional survey. If this portion happens to sit

near the top of the anomaly and the object is deep enough, the field could be very well explained in terms of simple DC or linear function, thus eliminating much need for the actual dipole.

The problem could become severe for deep and strong magnetic objects. One of the methods to overcome this would be to increase the size of the frame, which might be impractical. Another would be to find the best position of the array to resolve the target. While this was not an option in the present measurements, it could be readily achieved with real-time data acquisition and inversion from multiple sensors.

The inversion results for a large, deep object are shown in Figure 20. When no source is present, the background field parameters are estimated to be $D=14.6\text{nT}$, and a vertical gradient, $C=3.5\text{ nT/m}$. The background values, estimated as result of joint inversion, are quite different and sometimes far from expected (highlighted in the Figure). In these cases it is not possible to estimate dipole location correctly.

Magnet North Pole Down, Y array axis facing South																
Initial seed: X=0.00 Y=0.00 Z=1.00																
True location: X=0.00 Y=-0.19 Z=1.55																
name	X(m)	Y(m)	Z(m)	D(m)	DH(m)	DZ(m)	BkgC(nTBkgD(nT))	J(cgs)	std.J(cgs)	max_fit(nFit(nT))	incl	decl	angle			
25sensors_bottom_1x1m	-0.05	-0.21	1.62	0.09	0.06	-0.07	0.0	3.9	12885	1345	2.1	283.02	78.9	-74.0	30.8	
25sensors_top_1x1m	0	-0.17	1.58	0.04	0.02	-0.03	0.0	8.4	12041	922	1.41	158.43	81.1	-33.2	35.8	
9sensors-2D-0_5x0_5m	-0.12	-0.16	1.02	0.54	0.13	0.53	0.0	306.0	1658	936	2.37	102.66	55.8	-88.4	35.9	
9sensors-2D-1x1m	-0.01	-0.2	1.66	0.11	0.02	-0.11	0.0	-19.3	14347	570	0.6	32	283.02	82.6	-47.9	33.3
12sensors-3D_1x1m	-0.01	-0.19	1.7	0.15	0.01	-0.15	-42.0	-32.1	15890	1636	0.82	0.53	272.89	81.9	-35.7	35.0
18sensors-3D-0_5x0_5m-3D	-0.15	-0.22	1.48	0.16	0.15	0.07	60.8	76.1	9032	2030	3.19	1.43	245.87	65.2	-91.2	30.0
18sensors-3D-1x1_m-3D	-0.02	-0.19	1.65	0.1	0.02	-0.1	-19.6	-11.9	14010	445	1.2	0.61	348.21	81.1	-42.8	34.7
50sensors-3D_1x1m	-0.04	-0.2	1.64	0.1	0.04	-0.09	-18.8	-8.1	13830	676	5.02	1.67	348.21	79.7	-61.7	32.7
9sensors-3D-0_5x0_5m	-0.12	-0.16	1.02	0.54	0.13	0.53	283.9	270.5	1658	936	2.37	1.2	189.78	55.8	-88.4	35.9
9sensors-3D-1x1m	-0.01	-0.2	1.66	0.11	0.02	-0.11	-24.5	-16.2	14347	570	0.6	0.3	283.02	82.6	-47.9	33.3
Magnet North Pole to West, Y array axis facing East																
Initial seed: X=0.00 Y=0.00 Z=1.00																
True location: X=-0.94 Y=0.05 Z=1.38																
name	X(m)	Y(m)	Z(m)	D(m)	DH(m)	DZ(m)	BkgC(nTBkgD(nT))	J(cgs)	std.J(cgs)	max_fit(nFit(nT))	incl	decl	angle			
25sensors_bottom_1x1m	-0.92	0.08	1.33	0.06	0.04	0.05	0.0	28.0	10200	488	2.52	1.18	183.29	-3.2	-84.9	90.4
25sensors_top_1x1m	-0.93	0.1	1.45	0.08	0.05	-0.07	0.0	3.3	13460	1122	1.71	0.66	109.15	-0.6	-85.2	88.3
9sensors-2D-0_5x0_5m	0.01	-0.01	1.03	1.01	0.95	0.35	0.0	199.4	990	28	0.1	0.07	104.19	46.4	82.0	46.8
9sensors-2D-1x1m	-0.93	0.09	1.38	0.04	0.04	0	0.0	15.9	11545	110	0.08	0.06	159.92	-2.4	-84.7	89.6
12sensors-3D_1x1m	-0.93	0.1	1.41	0.05	0.05	-0.03	-7.9	11.1	12259	302	0.46	0.26	190.01	-1.8	-84.6	89.0
18sensors-3D-0_5x0_5m-3D	-0.92	0.05	1.29	0.1	0.02	0.09	4.5	36.9	9881	1225	1.75	0.93	164.25	-5.4	-85.6	92.8
18sensors-3D-1x1_m-3D	-0.94	0.1	1.42	0.07	0.05	-0.04	-13.7	7.0	12853	433	1.54	0.54	209.21	-1.5	-84.6	88.8
50sensors-3D_1x1m	-0.93	0.09	1.38	0.04	0.04	0	-8.7	16.5	11577	368	3.21	1.13	232.58	-2.9	-84.8	90.1
9sensors-3D-0_5x0_5m	0.01	-0.01	1.03	1.01	0.95	0.35	204.7	173.8	990	28	0.1	0.07	128.91	46.4	82.0	46.8

Figure 20. Inversion results for the strong dipole with different orientations.

6.2.2.3 RESULTS FOR MORTAR SHELL

Nose down, system Y axis to the North																
Initial seed: X=0.00 Y=0.00 Z=1.00																
True location: X=-0.35 Y=0.01 Z=1.21																
Configuration	X(m)	Y(m)	Z(m)	D(m)	DH(m)	DZ(m)	BkgC	BkgD	J(cgs)	std.J	max_fit	Fit	ampl	incl	decl	angle
25sensors_bottom_1x1m	-0.32	0.04	1.38	0.18	0.05	-0.17	0	13.63	658	24	0.5	0.15	41.21	73.8	78.5	29.0
25sensors_top_1x1m	-0.33	0.03	1.37	0.16	0.03	-0.16	0	12.86	649	41	0.22	0.11	23.13	75.9	82.1	29.6
9sensors-2D-0_5x0_5m	-0.16	0.07	1.23	0.21	0.21	-0.02	0	22.74	419	51	0.1	0.06	27.39	48.4	73.2	38.3
9sensors-2D-1x1m	-0.33	0.03	1.38	0.18	0.04	-0.17	0	13.28	665	19	0.07	0.04	39.76	75.3	82.5	29.8
12sensors-3D_1x1m	-0.33	0.03	1.38	0.17	0.04	-0.17	1.88	13.11	657	19	0.12	0.05	39.76	75.0	83.4	30.1
18sensors-3D-0_5x0_5m-3D	-0.23	0.07	1.36	0.21	0.14	-0.15	5.28	16.29	625	43	0.19	0.1	30.91	59.8	72.3	31.2
18sensors-3D-1x1_m-3D	-0.34	0.03	1.39	0.18	0.03	-0.18	1.54	12.79	673	20	0.18	0.07	39.76	76.8	83.0	29.7
50sensors-3D_1x1m	-0.33	0.04	1.39	0.19	0.04	-0.18	1.54	12.93	677	18	0.48	0.14	41.21	75.3	78.4	28.7
9sensors-3D-0_5x0_5m	-0.16	0.07	1.23	0.21	0.21	-0.02	10.81	21.39	419	51	0.1	0.05	27.39	48.4	73.2	38.3
9sensors-3D-1x1m	-0.33	0.03	1.38	0.18	0.04	-0.17	1.81	13.06	665	19	0.07	0.03	39.76	75.3	82.5	29.8
Nose up, system Y axis to the North																
Initial seed: X=0.00 Y=0.00 Z=1.00																
True location: X=-0.21 Y=0.00 Z=1.21																
25sensors_bottom_1x1m	-0.37	-0.06	1.19	0.18	0.18	0.02	0	18.53	315	12	0.39	0.19	29.64	-43.0	-53.1	113.9
25sensors_top_1x1m	-0.38	-0.03	1.19	0.17	0.17	0.02	0	17.07	307	18	0.29	0.12	14.98	-45.4	-61.4	119.1
9sensors-2D-0_5x0_5m	-0.36	-0.04	1.15	0.17	0.16	0.06	0	16.93	279	45	0.12	0.06	18.76	-41.5	-57.9	114.3
9sensors-2D-1x1m	-0.43	-0.09	1.24	0.25	0.25	-0.03	0	19.56	370	30	0.16	0.09	29.64	-44.6	-45.5	112.8
12sensors-3D_1x1m	-0.42	-0.07	1.22	0.23	0.23	-0.01	5.39	18.45	348	52	0.4	0.25	29.64	-45.7	-49.5	115.1
18sensors-3D-0_5x0_5m-3D	-0.37	-0.04	1.17	0.17	0.17	0.04	1.2	16.76	299	21	0.23	0.11	18.76	-41.6	-59.3	115.0
18sensors-3D-1x1_m-3D	-0.40	-0.06	1.21	0.2	0.2	0	4.39	18.1	327	29	0.46	0.24	29.64	-45.3	-54.0	116.2
50sensors-3D_1x1m	-0.37	-0.05	1.19	0.17	0.17	0.02	3.91	17.86	311	11	0.67	0.24	29.64	-43.8	-55.6	115.5
9sensors-3D-0_5x0_5m	-0.36	-0.04	1.15	0.17	0.16	0.06	1.09	16.8	279	45	0.12	0.06	18.76	-41.5	-57.9	114.3
9sensors-3D-1x1m	-0.43	-0.09	1.24	0.25	0.25	-0.03	4.08	19.05	370	30	0.16	0.09	29.64	-44.6	-45.5	112.8
Nose west, system Y axis to the North																
Initial seed: X=0.00 Y=0.00 Z=1.00																
True location: X=-0.62 Y=0.06 Z=0.80																
25sensors_bottom_1x1m	-0.47	-0.01	0.84	0.16	0.16	-0.04	0	18.81	360	19	1.33	0.62	83.22	38.6	-77.0	54.9
25sensors_top_1x1m	-0.46	0.00	0.81	0.16	0.16	-0.01	0	17.67	337	17	0.44	0.24	34.42	36.5	-77.1	56.6
9sensors-2D-0_5x0_5m	-0.55	-0.06	0.81	0.14	0.14	-0.01	0	16	419	14	0.06	0.04	65.61	19.7	-87.6	75.6
9sensors-2D-1x1m	-0.50	0.01	0.85	0.14	0.13	-0.05	0	17.66	390	35	0.42	0.29	59.02	34.9	-76.1	57.4
12sensors-3D_1x1m	-0.52	0.00	0.81	0.11	0.11	-0.01	4.15	17.11	383	33	0.63	0.41	59.02	28.8	-79.2	63.9
18sensors-3D-0_5x0_5m-3D	-0.51	-0.06	0.83	0.16	0.16	-0.03	3.31	16.95	391	11	0.24	0.11	70.29	25.8	-87.3	67.0
18sensors-3D-1x1_m-3D	-0.52	0.00	0.82	0.12	0.11	-0.02	3.47	17.15	385	21	0.89	0.4	59.02	29.5	-78.9	63.1
50sensors-3D_1x1m	-0.50	-0.01	0.83	0.14	0.14	-0.03	4.78	17.48	377	12	1.64	0.59	83.22	33.4	-79.1	60.0
9sensors-3D-0_5x0_5m	-0.55	-0.06	0.81	0.14	0.14	-0.01	-0.54	16.07	419	14	0.06	0.03	65.61	19.7	-87.6	75.6
9sensors-3D-1x1m	-0.50	0.01	0.85	0.14	0.13	-0.05	2.98	17.29	390	35	0.42	0.27	59.02	34.9	-76.1	57.4

Figure 21. Inversion results for mortar shell object at different orientations. Estimated background values without object are $D=14.6\text{nT}$, $C=3.4\text{nT}$.

The most practical object used in these measurements was a demagnetized mortar shell. In this section results are presented and compared with the true location.

Figure 21 shows that the estimated magnetic moment is changing with orientation. This could be explained assuming permanent and induced magnetization of the object are similar. At a certain orientation these two could coincide, increasing the total magnetic moment. As it can be seen maximum magnetic moment is ~650 cgs for the “nose down” configuration and a minimum ~350 cgs for the “nose up” configuration. This implies that the permanent magnetization is from tail to nose, which adds to the induced one for “nose down”. Intermediate configuration “nose west” has its moment similar to “nose up”.

Figure 21 also demonstrates that no particular sensor configuration has a strong advantage over the others. In fact, 50 sensors, the most rigorous measurement possible, did not deliver an estimate closest to the ground truth.

6.2.2.4 CONCLUSIONS AND CONSIDERATIONS

We note the following conclusions and considerations from our results

1. A system must include a means for measuring its orientation relatively to the Earth's magnetic field. The simplest method would be to use a flux gate magnetometer, located approximately 0.5m from closest magnetic sensors. Designing a system without such orientation device would severely limit its applications.
2. The frame size should be at least 1x1 m to deliver estimates for targets located from 1 to 1.5 m from the array. A larger array is better for deeper targets. This leads to the idea of a modular design to allow user easy extension of the frame.
3. Because optically pumped magnetometers are not absolute instruments, measurement levels may change from sensor to sensor by 1nT or more. In the presented sets of measurements this was not an issue because only one actual sensor was used for measurements and another was used as a base station. Sensors in an array have to be pre-calibrated to remove DC offsets between them.
4. All sensors should be well synchronized to remove the effects of moving objects nearby.
5. The number of sensors does not appear to be as important as expected. Results from tests where we used a large number of sensors (25 or 50) did not show a particular advantage over the results from tests using 12-18 sensors. We conclude a 12 - sensor 3-D configuration should be sufficient.

We recommend the 12-sensor array, shown in Figure 1, and modeled as “12sensors-3D_1x1” in Figure 14. With this array we expect the system should deliver an accuracy of 0.2 m or better for estimating position. This assumes that the targets are similar to the mortar shell at a depth from 1 to 1.5 m. We further note the following:

1. Using a platform like a PC laptop, the efficiency of the computational code allows real time solutions.
2. Based on the modeling, it appears that the system should perform reasonably well in the presence of large objects such as cars and other machinery, assuming they are 10-15 m from the area.
3. Modeling shows that the operator should be able to resolve two or more objects, assuming they are about 1 to 1.5 m deep and separated by about two meters. Closer objects may not be resolved by the system and may appear as one object to the operator.

4. The result of the inversion does not depend much on the initial estimated position of the object (thanks to SLATEC numerical optimization routine), therefore there is no need to perform multiple seeds and inversions for the same set of data. It is reasonable to accept the initial location as middle of the array. Practice showed that in the presence of a strong enough object (such as the mortar shell), the inversion converges to exactly the same solution.

The presented measurements are valuable material in terms of developing rejection criteria for the system, i.e. identifying situations when system cannot “lock” to the object. This needs further analysis but it appears that the following could be used:

1. Overall low amplitude of the system, in a range of 1-3 nT. We define the amplitude as the difference between the largest and smallest sensor reading of the array.
2. Overall fit between observed and modeled data. This should not exceed 0.5 nT
3. Estimated magnetic moment and its standard deviation. Low values of magnetic moment such as below 50 cgs should be rejected. Another valuable criterion is the ratio of the magnetic moment standard deviation to the estimated moment. We have used rejection criteria of this value being above 30% and this sole condition allowed us to sort out all background measurements, and had only a few false rejections of a real object.
4. Estimated position. It is clear that position cannot be estimated reliably if it is further than 1 m from either side of the array, or if the estimated depth is greater than 2 m
5. Excessive standard deviations of estimated position, exceeding 0.5 m

Conditions opposite to the above indicate the system reliably estimates the position of the object.

One significant advantage of utilizing a real-time array is that one may observe how the estimated target parameters vary with the position of the array. The rule of thumb would be that all estimated parameters, including background field values, should not change much as the operator moves the frame. Based on sliding modeling we can see that it is possible that the system would “lock” on one object and then re-acquire lock on other objects. We will discuss this feature of the operation of our real time array in more detail in section 6.5.2.3

6.3 BUILD PLATFORM WITH ARRAY OF SENSORS

In order to continue our evaluation of the feasibility of the sensor array, we built a 12-sensor array to test our simulations. The final array, with sensors on a 1m x 1m square is shown in Figure 22:



Figure 22. Array used for experiments of real-time object tracking and localization.

The electronics modules, built using the electronics described in the final report of MM1512 are shown in Figure 23:

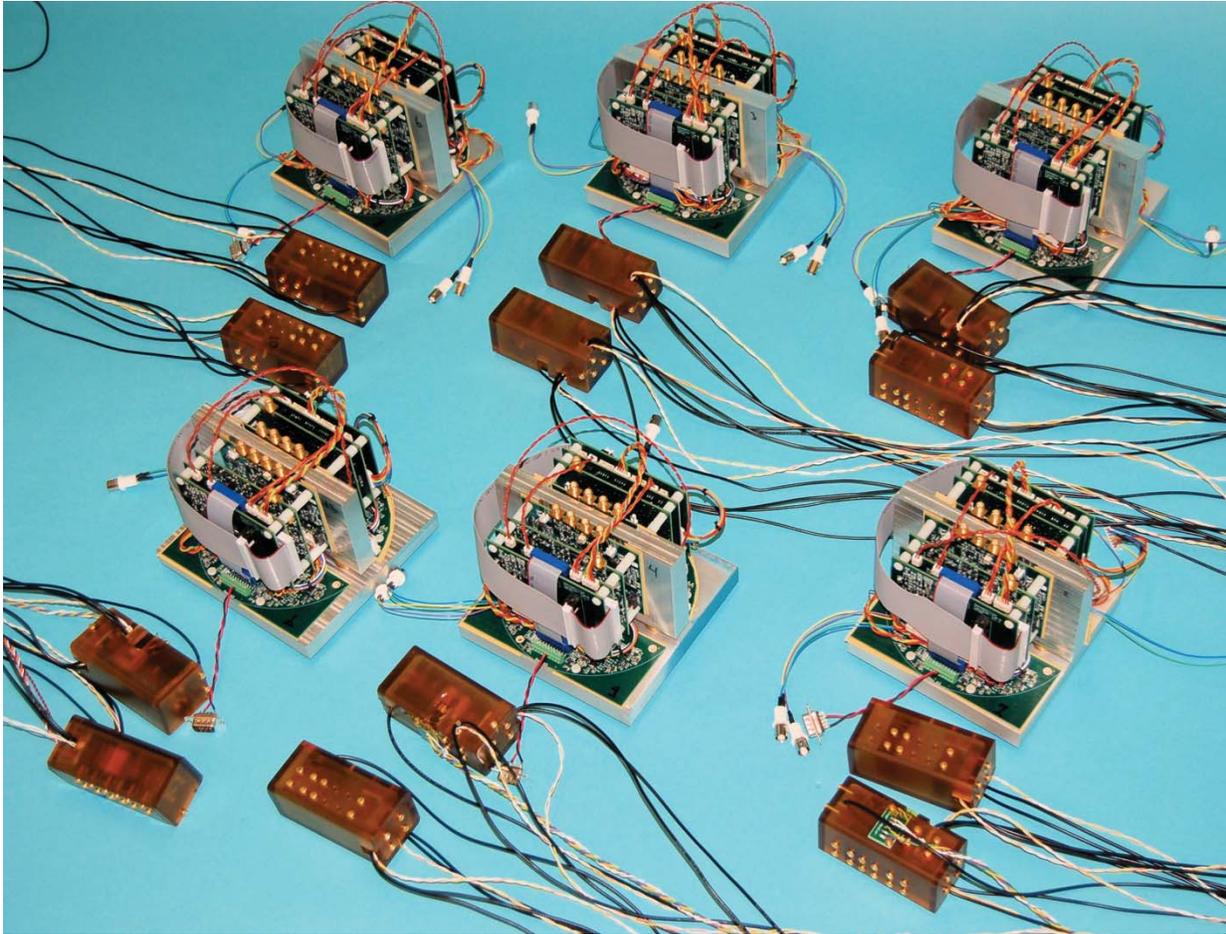


Figure 23. Electronics modules built for the array measurements.

These systems nicely served their purpose in evaluating the array, while also highlighting the need for smaller sensors and electronics systems. These were developed later in the project.

6.4 DEVELOP REAL TIME PROCESSING ALGORITHM

6.4.1.1 SPEED OF COMPUTATIONS

Because our array inversion system is designed to be operated in real time, it's important to know how much time could it take to perform an estimation, and what is its possible variation. We propose running the system at a 10 Hz sample interval, so ideally the software should be able to invert any data in 100 ms. Before undertaking this work, we anticipated difficulty due to the random nature of the data. Therefore, we utilized a multi-threaded model with separate data acquisition and inversion threads.

We estimated the speed of the modified program by running number consequent inversions (say 1000). From this, we measured the total time and then computed the CPU time for each inversion.

We did not use a state-of-the-art computer, even at the time this work was performed: Intel(R) Core(TM) Duo CPU E4500 @ 2.20 GHz, 4 GB of RAM, 32 bit Windows Vista.

We learned the following:

- Inversion can be carried out in two modes: with numerical approximation of derivatives and with analytical approximation. Measurements showed that (1) is approximately two times slower than (2).
- Assuming favorable initial inversion location near the center of the array, with real position shifted about 1 m from the center, one inversion takes about 5-6 ms.
- With the same initial location, but unrealistic seed depth of 10 m, inversion converges to the correct solution for 50 ms.
- With bad initial location (X and Y 4 m from the center) inversion does not converge to the correct solution, but still takes about 6 ms per inversion.
- All measurements made for 50-sensor array; using fewer sensors would be much faster.
- For the data without targets these times are even smaller, but it is difficult to estimate what type of noise could cause the inversion to take longer.

The conclusion is that there should be easily enough time to invert data assuming a 10 Hz acquisition sample rate.

The graphical user interface we developed is shown in Figure 24. The red dot tracks the location of objects passing over the array. While this display shows a plan view, the tabular results show additional estimates, including the height and angle of the target. We found the accuracy of these so be remarkable. For example, to make repetitive measurements, we swung a target, pendulum-style, from the ceiling above the array. We could easily observe that the height of the target was varying slightly, as well as its angle.

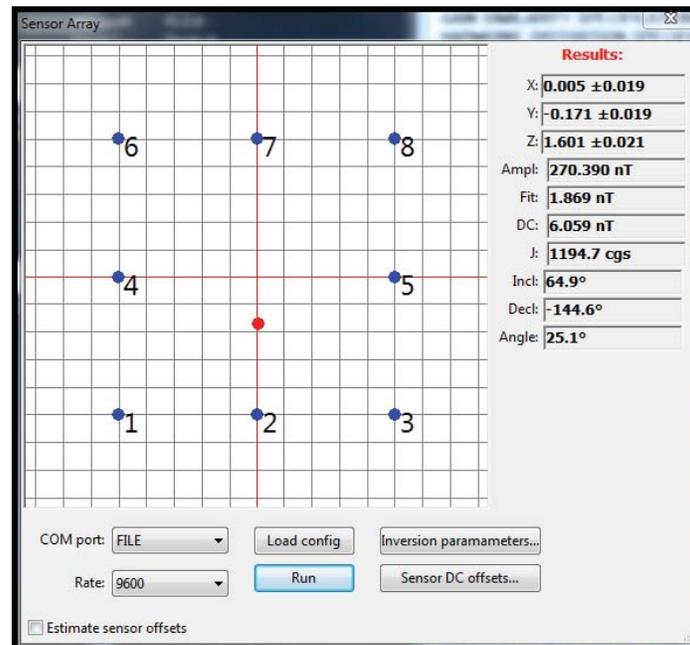


Figure 24. Real time display of the fitting program.

6.5 TESTING AND DEMONSTRATION.

In this section we present the results of our tests with this system. Measurements included the following:

Hydrodynamic characterization of the system. The system must be reasonable for a diver to maneuver.

Responsiveness of the system. The system must respond rapidly enough not to confuse the operator as they move about the area of interest seeking the target. There may be two modes of operation, a detection mode and a discrimination mode.

Capabilities in the presence of clutter. One of the most difficult tasks for any UXO detection system is to behave reliably in environments highly cluttered in potentially magnetic scrap material. Scrap metal can cause a response. This situation may still require some operator skill to recognize such situations. We expect that identifying multiple target situations will be much easier in the field than when post processing a given data set. How to best work in a cluttered environment will be studied in some detail.

Accuracy of target identification. Obviously a key feature is the ability to accurately measure the position of anomalies. This should be well characterized in our simulations and designs, and we wished to confirm and test our implementation in this task.

Performance of the system in the presence of multiple objects. One of the most difficult to analyze situations for an inversion technique is when the model being used in the inversion (in our case a single dipole target) does not fit the actual situation. We explored how the real-time system behaved in this situation.

6.5.1 TESTING SYSTEM HYDRODYNAMICS.

In order to most simply test the maneuverability of a 1x1m array underwater, we built an inexpensive mockup out of PVC pipe material. This device is shown in Figure 25.



Figure 25. Mechanical mockup to test the system's maneuverability underwater.



Figure 26. Enlisting the assistance of local divers in Monterey Bay.



Figure 27. Preparing to test the system underwater.

We enlisted the assistance of some divers at a local diving club. Cheerful volunteers tested the array mockup and reported that it was quite easy to maneuver. They even suggested the array could be made larger and still perform reasonably well. While this was hardly a rigorous test, the divers reported that the system was unexpectedly maneuverable underwater and even the presence of currents did not disrupt the system.



Figure 28. Maneuvering the system underwater.

6.5.2 PERFORMANCE SIMULATIONS

6.5.2.1 ARRAY ESTIMATIONS IN THE PRESENCE OF LARGE MAGNETIC BODIES

We purposely chose the test location in a relatively magnetically quiet area (Figure 8) to avoid possible complications at the testing and debugging stages. However the question arises how the array would perform in situations of close proximity to a large magnetic disturbance, such as an automobile, surface vessel, or other metal objects, assuming they are not moving. To address this question, semi-numerical modeling was done using as a reference the “nose down” mortar test measurements. Additional signals from about 1 T of steel, assuming a vertical magnetic moment was added to the recorded data, and the data re-inverted. This was repeated at different “big dipole” separations from the actual location of the mortar, and the position error was plotted as function of that distance.

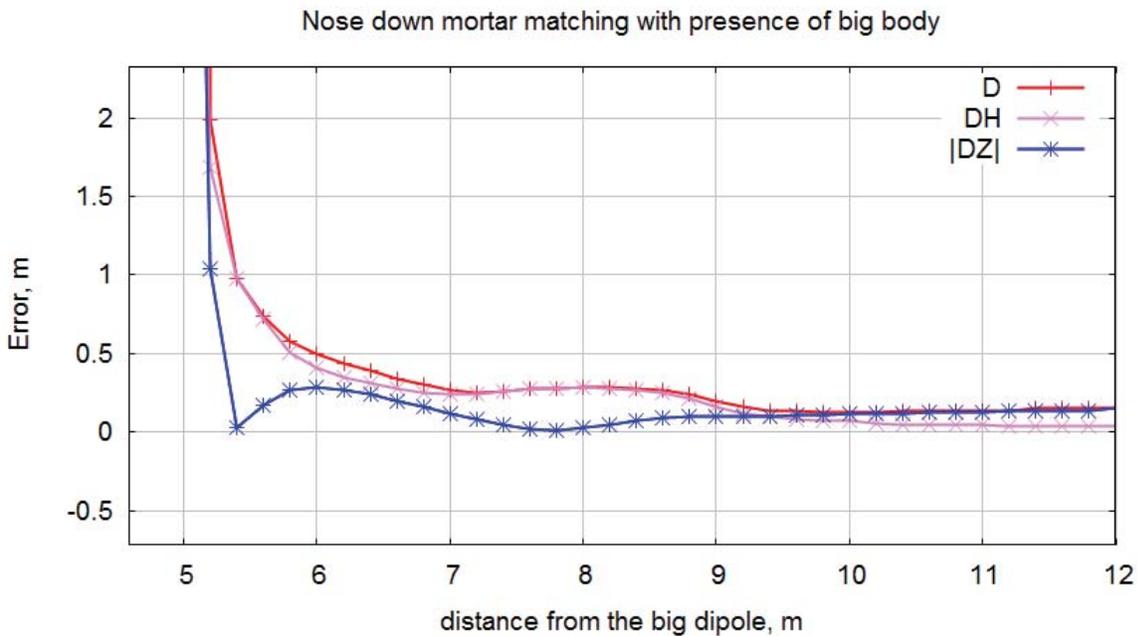


Figure 29. Dipole matching with presence of systematic disturbance due to significant source of magnetic field.

Results are presented in Figure 29. The errors are substantial at approximately 5-6 m from the big dipole (in fact, the software could not estimate parameters at a distance less than 5 m) but decay quickly as distance increases, becoming unimportant at range above 9-10 m.

6.5.2.2 ACCURACY OF THE ESTIMATES

Figure 30, Figure 31, and Figure 32 show highlights of our results from the real-time algorithm. The vertical axis shows either the x, y, or z estimated position, while the horizontal axis shows the actual x position. In each case, the object was moved across the array in the x-direction, at constant y and z coordinates.

In summary, the array performed extremely well, readily confirming the predictions we made in the array modeling and design phase of this project. The offset shown in the y-position of the mortar in Figure 30 is due to the fact the estimated magnetic dipole is not in the physical center of the mortar.

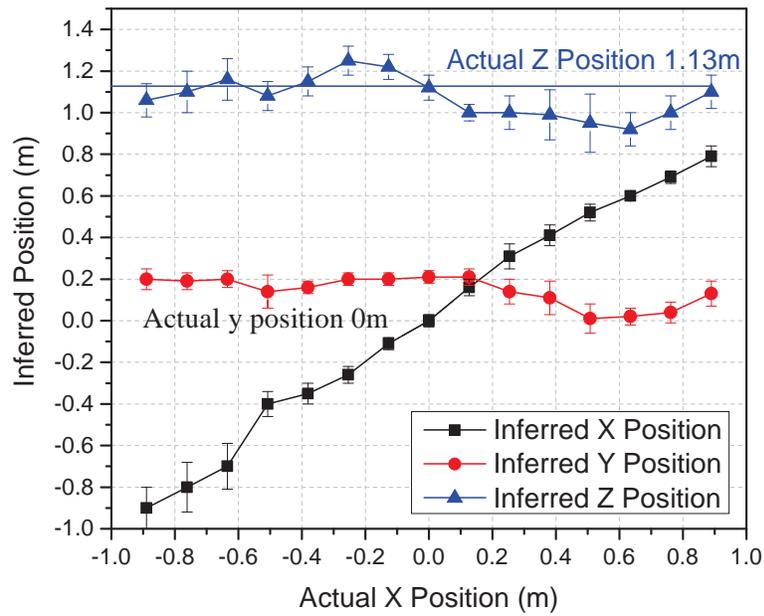


Figure 30. Results of the real-time algorithm. Mortar shell is at 1.13m depth with moment 0.4 A m^2 .

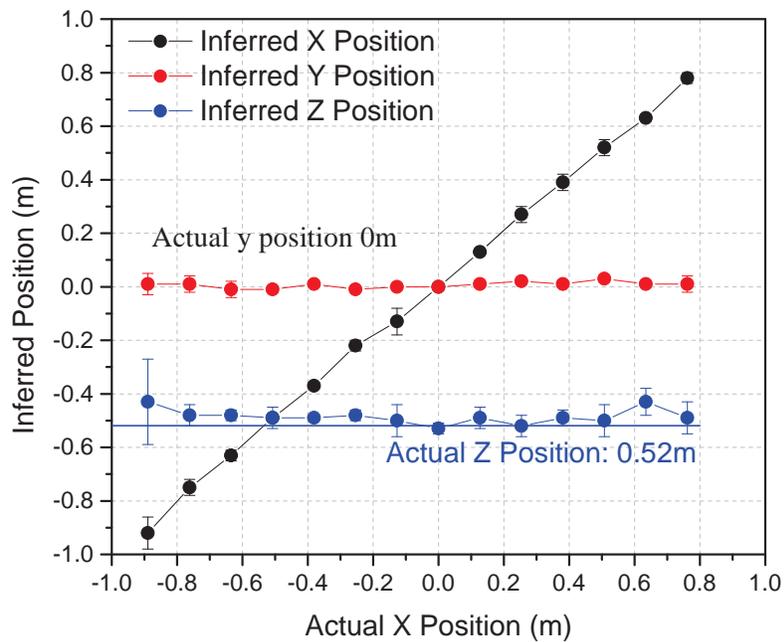


Figure 31. Results of the real-time algorithm. Coil simulator as target with moment 0.04 Am^2 .

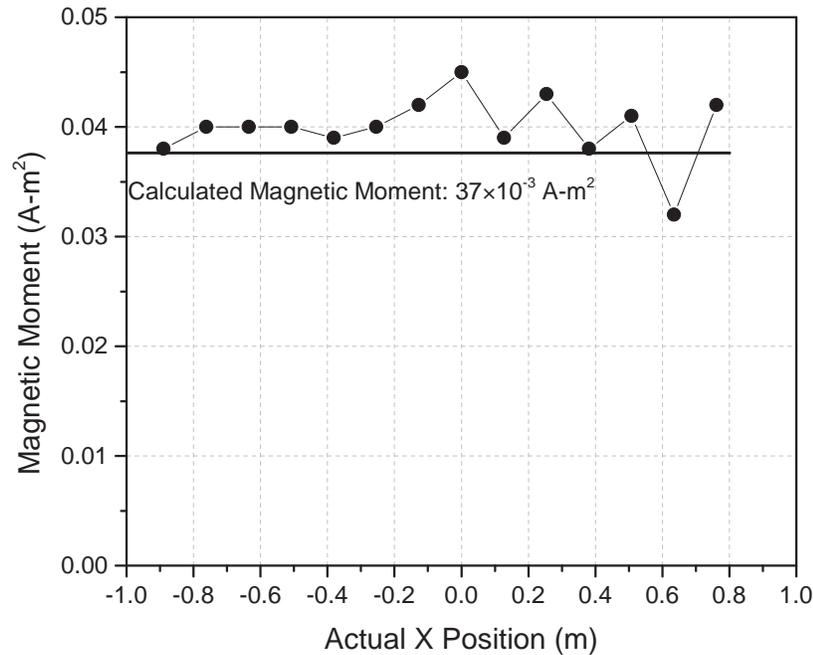


Figure 32. Results of the real-time algorithm in estimating magnetic moment. Actual moment was 0.04 Am².

6.5.2.3 PERFORMANCE IN THE PRESENCE OF MULTIPLE OBJECTS

In this experiment we asked how the system would behave in the case of the close proximity of another, similar object. To answer this, we used computer modeling as well as the results from our array.

As a reference, we took mortar “nose down” experiment magnetization data, and the object was believed to have only vertical magnetic moment of 600 cgs (0.6 A/m²). We placed the models sequentially at 2, 1 and 0.5 m separation and used an array of 1x1 m size, while numerically passing the 50-sensor array across that combined target between -5 to +5 m. Both targets had same Y position (Y=0) and the same depth (1.2 m).

Our objective was to see if the operator could detect a similar situation in the field and what, if any steps could be taken to resolve it.

Estimations for X, Y, Z, standard deviation of X and magnetic moment are plotted in Figure 33.

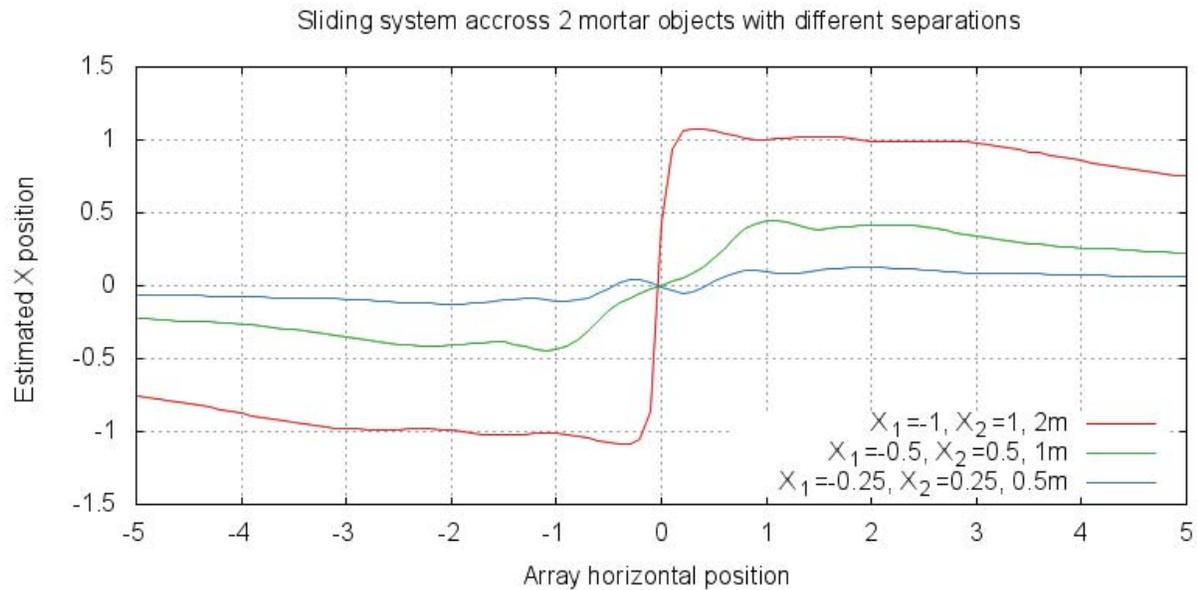


Figure 33. Estimated X position at different locations of the array.

One can see that with a separation of 2 m (red), it is possible to resolve two different objects when the array is located at -1 and +1 m, accordingly. There is a sharp transition zone around 0 m for red which indicates that operator would see unstable position estimations while in that zone. For 1 m separation (green) the curve is much smoother but it is still possible to get realistic estimate for X moving into -1, +1 m positions. For 0.5 m separation, objects tend to behave as if they were a single object, and no reliable position estimate is possible.

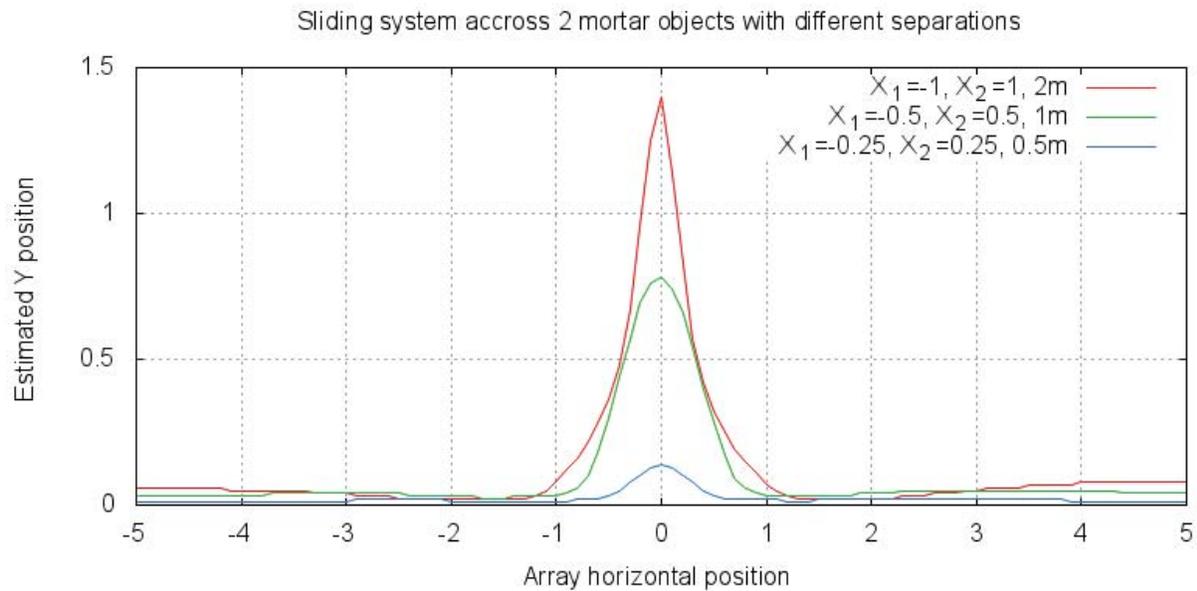


Figure 34. Estimated Y position, which should be at 0. Considerable error is observed when the array is between the targets (array position equal 0). The error is bigger for 2 m target separation. Again it shows that the operator would see Y position changing dramatically with array motion while in between the targets.

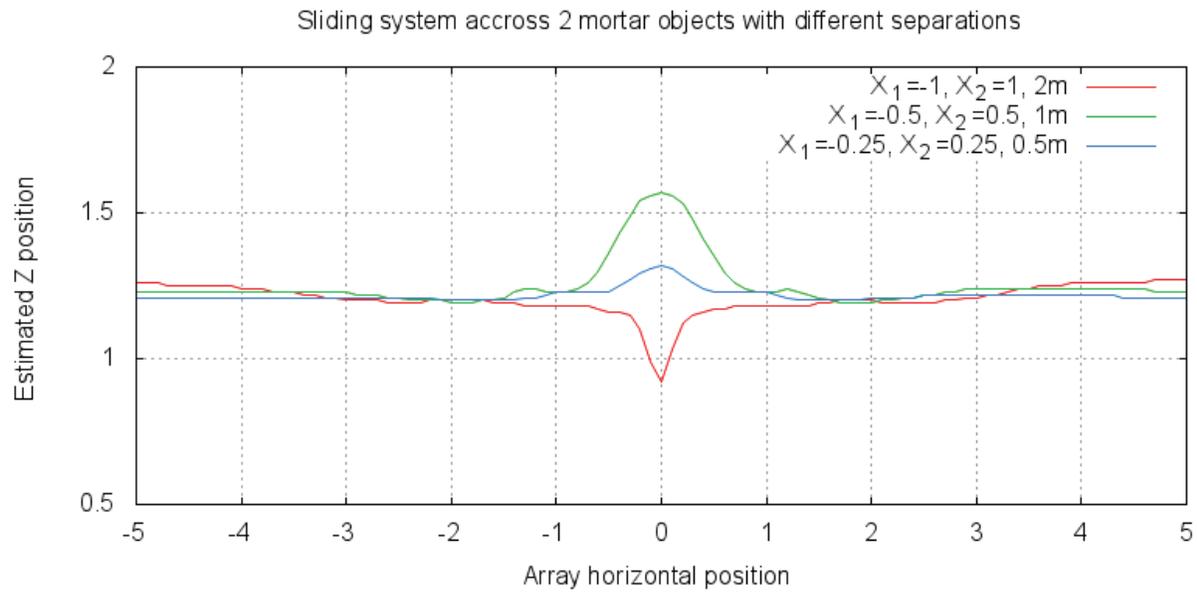


Figure 35. Estimated Z position. True position for both targets is 1.2 m, and the system measures this value on either side of the targets. Instability of estimated Z is observed while operating inside targets.

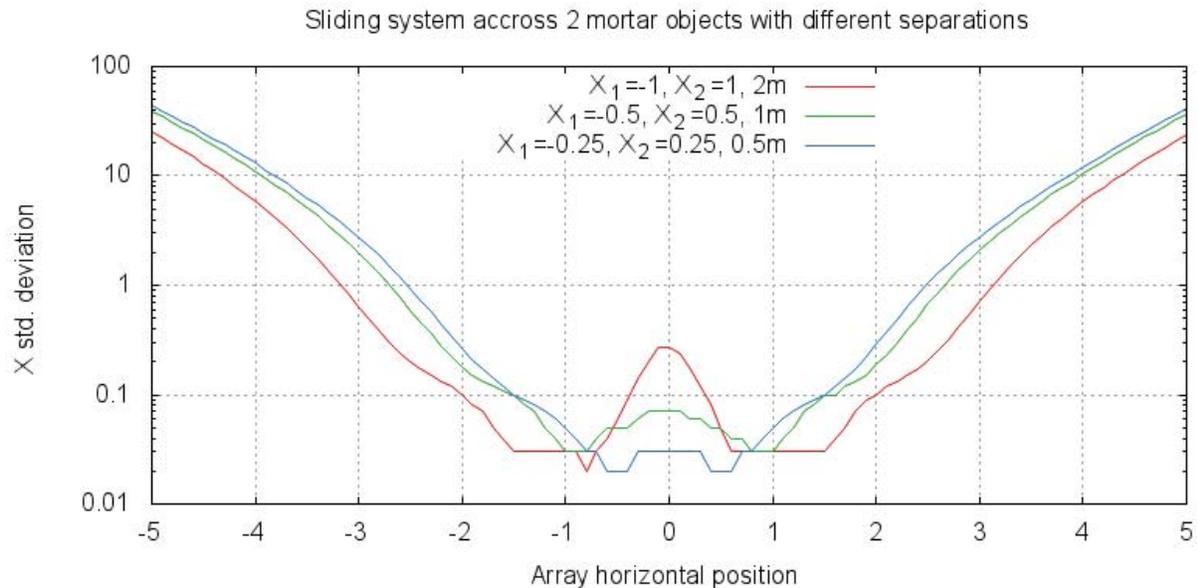


Figure 36. Standard deviation of estimated X, in logarithmic scale. It can be seen that this values grows exponentially when system loses “lock” on either side of the targets. This is a clear indication that the system is not operating properly. All configurations deliver however reasonable standard deviation while in +/- 1 m from the targets.

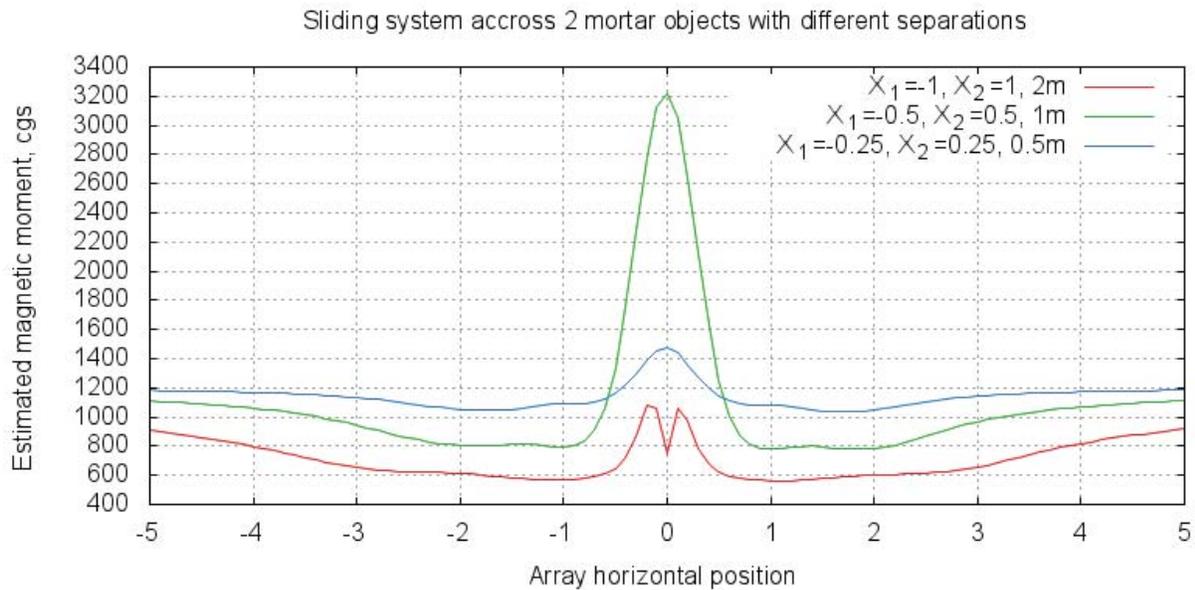


Figure 37. Estimated magnetic moment, cgs. Note that real magnetic moment was 600 cgs for each object, and system converges to that value only with dipole separation of 2 m. In other cases higher values of 800 and 1000 cgs are delivered, which makes sense because of the presence of two objects.

The most significant fact about the performance of the array in this adverse situation is that the results of the estimation depend radically on the position of the array. The simulation results shown in Figure 34, Figure 35, Figure 36, and Figure 37 show how dramatically the estimated position and moment of the target vary as the position of the array is changed. These simulations are confirmed with our experiments with the actual array, shown in Figure 38, Figure 39, and Figure 40.

These results show one of the most important features of the real time array. Because the results are shown in real time, the operator has the opportunity to observe the behavior of the system as it is moved. By being able to move the sensor array, see it lock onto one object, then have wildly varying estimates, then consistently lock onto another target, this complex situation may be readily interpreted in the field.

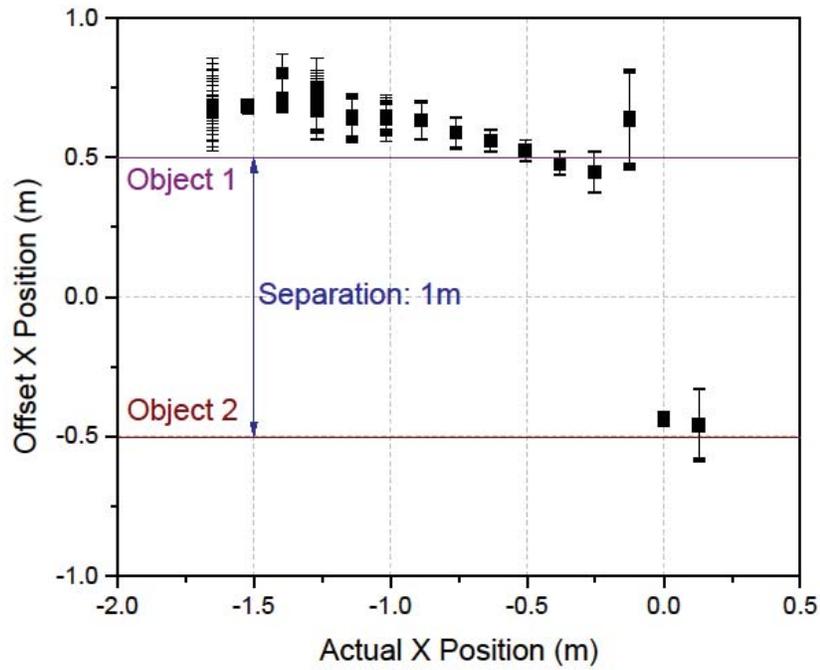


Figure 38. Estimated position with two objects at 1m apart.

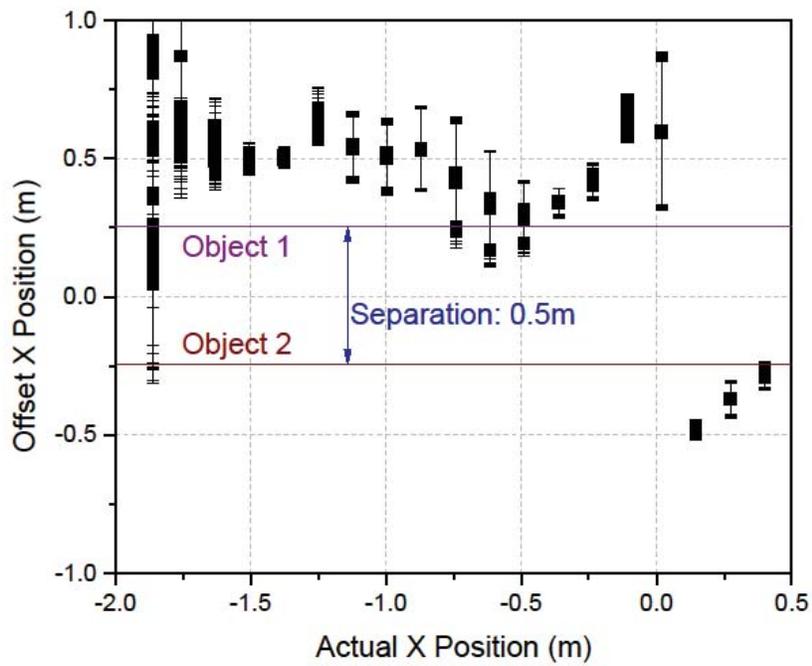


Figure 39. Estimated position with two objects at 0.5m apart.

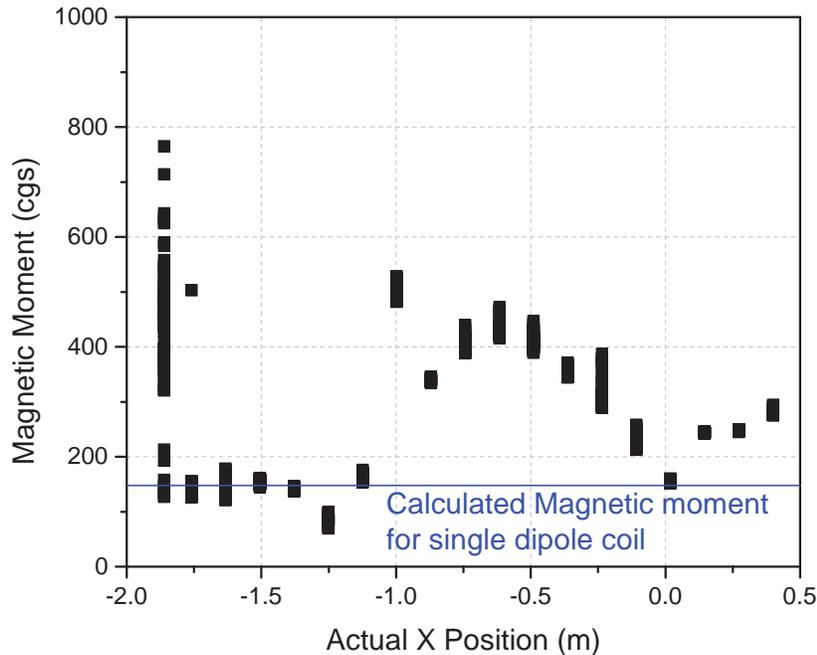


Figure 40. Estimated moment with two objects 0.5m apart.

The results we obtained in this section are summarized as follows:

1. As a rule of thumb, the estimated parameters (position, depth, magnetization and their standard deviations) should not vary more than about 20% when operator moves the array around the target. This would indicate that “lock” on the target is obtained.
2. In certain situations, it would be possible to resolve multiple targets by moving the array sideways, so it measures primarily field of one target. For the example, in this section this distance is about 2 m for that type of targets.
3. It is also possible that two targets close to each other would be treated as single object, which can be seen here at 0.5 m separation. The following rule could be proposed: targets could not be resolved separately if the distance between them is half of their depth or less.
4. One of the indications the array does not record reliable data would be a considerable increase in standard deviation of the estimates, as well as a decrease of amplitude. We define amplitude here as difference between minimum and maximum values recorded by the array sensors. Based on overall experience with field measurements, it is clear that minimal acceptable amplitude should be about 0.5 nT.

Our analysis of the performance and benefits of an array of sensors led to our desire to miniaturize the system electronics. Results of this are described in the next section.

7 RESULTS: MFAM ELECTRONICS MINIATURIZATION

7.1 OVERVIEW

7.1.1 APPROACH

Our approach in this project was somewhat different than in our previous projects. In previous work, we focused on building the simplest possible physics package, one that makes use of only a single laser diode. Then, we designed the electronics system to make use of the simple physics package. While this is easiest from that perspective, it makes the design of the electronics system much more complicated.

In this project, we approached the system design iteratively, as shown in Figure 41. This allowed us to design the simplest overall system, and allowed us to meet our performance objectives while maintaining the best balance between physics package and electronics complexity. In our latest design, we actually use two lasers, one for the pump function, and another for the probing function. While the pump and probe functions can be done with a single laser, utilizing two separate beams reduces the complexity of the electronics system considerably. We have thus chosen to utilize a two-laser scheme.

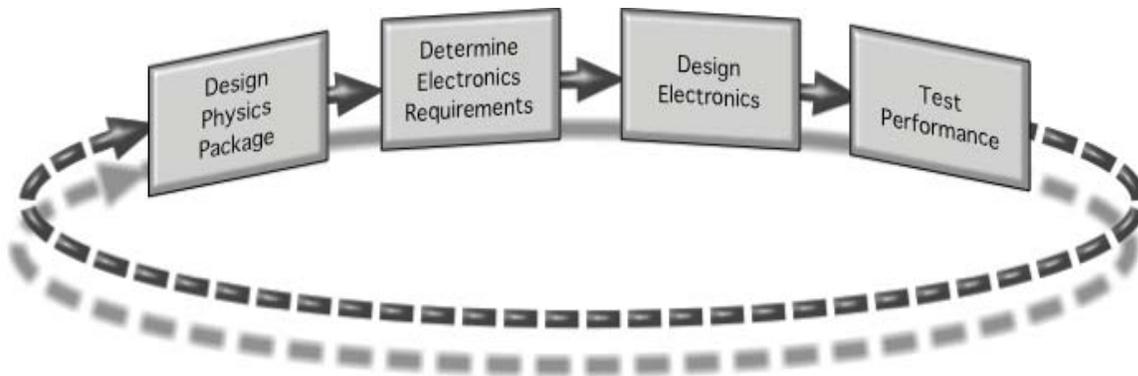


Figure 41. Iterative design process in which both the physics package and electronics complexity are considered.

In this section, we outline the design of the electronics system. We will discuss as follows:

- Approach – We chose a digital design approach, and this section describes the features of that approach and why we chose it. Most importantly, we took a careful and systematic approach to the design, using a Model, Implement, and Verify approach at each step (Figure 42).
- Instrument Architecture – This section describes the major building blocks of a magnetometer instrument and how they work together.
- Sensor Modeling – We began the design process by creating a software model of the sensor itself, and how it interacts with the electronics system. This model was detailed enough to allow us not only to test our design approaches, but even to simulate the detailed performance of the system and determine performance requirements for the inner blocks of circuitry.

- Implementation – Beginning with the high-level system block diagram, we designed the various blocks of circuitry. We then implemented these algorithms in a “hardware-in-the-loop” methodology, using programmable laboratory instrumentation (Figure 43). We then designed bench-top-testable brass boards implementing the desired circuits (Figure 44). Finally, we designed the final layouts, with the desired size and form factor (Figure 45). This approach allowed us to proceed step-by-step, addressing the major issues first, ensuring that each step would be successful. The final electronics board set is indicated in Figure 46.
- Firmware Programming – The crucial final step of the process was to implement the desired firmware, programming the digital control loops that perform the core magnetometer signal extraction. One major advantage of our architecture is that we can continually improve the system operation in firmware. Common to most modern digital systems, updates may be installed in the field by users.

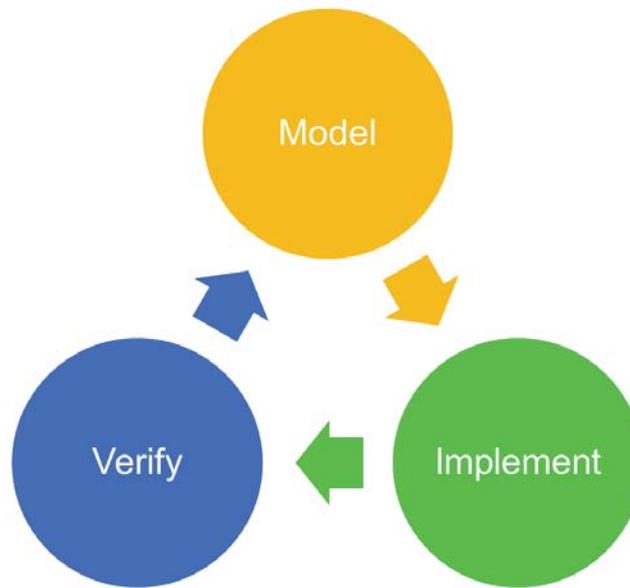
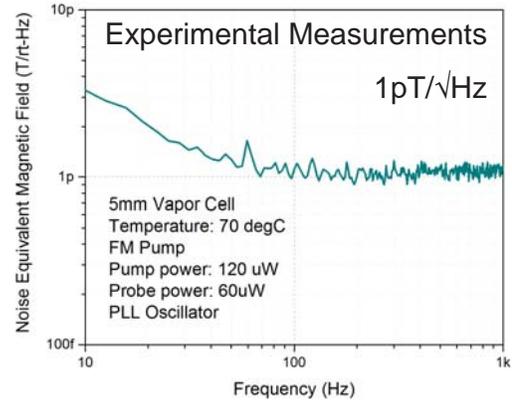
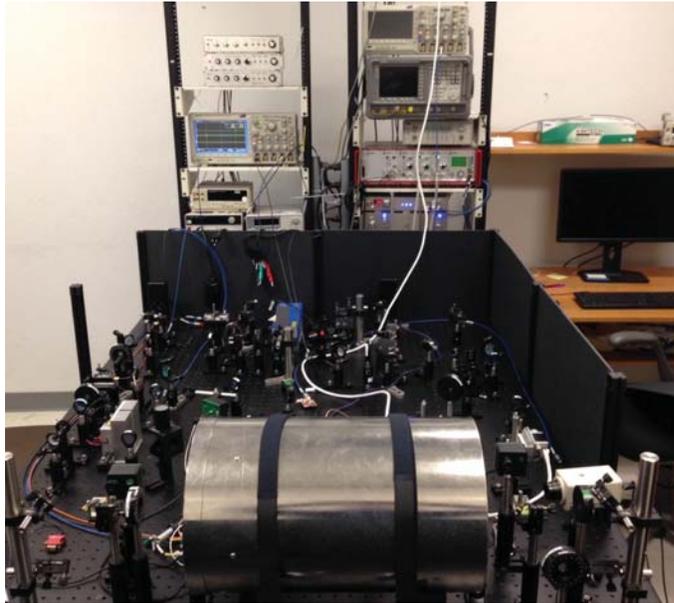


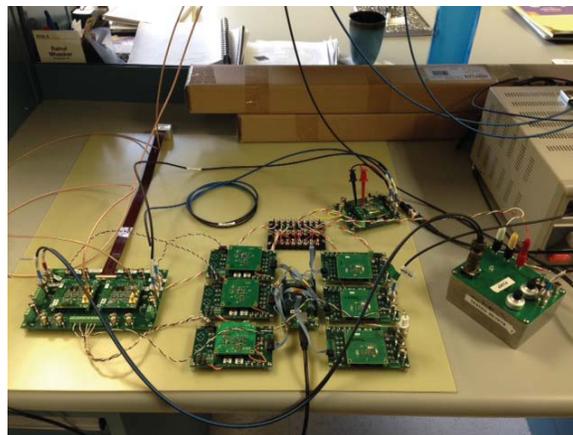
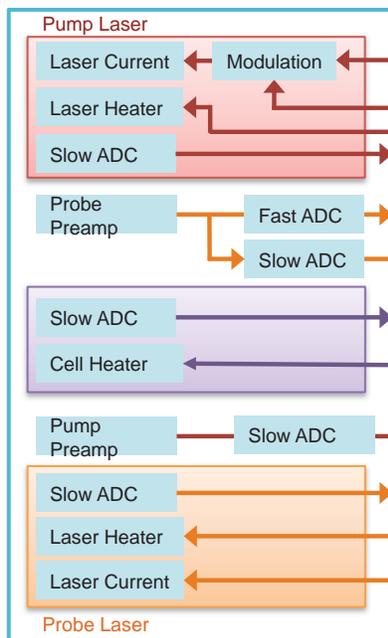
Figure 42. Model-based Design Workflow

Sensor Emulation in the Physics laboratory



Using commercial lock-in amplifier and PLL instrument and laboratory-grade lasers.

Figure 43. The first step in the physical electronics design was to prototype the system using programmable laboratory instrumentation.



- Each functional design tested individually
- Three designs updated for better performance / lower power consumption
- Sensor operation verified with these functional blocks for high performance
- $< 1\text{pT}/\sqrt{\text{Hz}}$ sensitivity observed with commercial PLL

Figure 44. The entire system was built in a bench top form factor for testing.

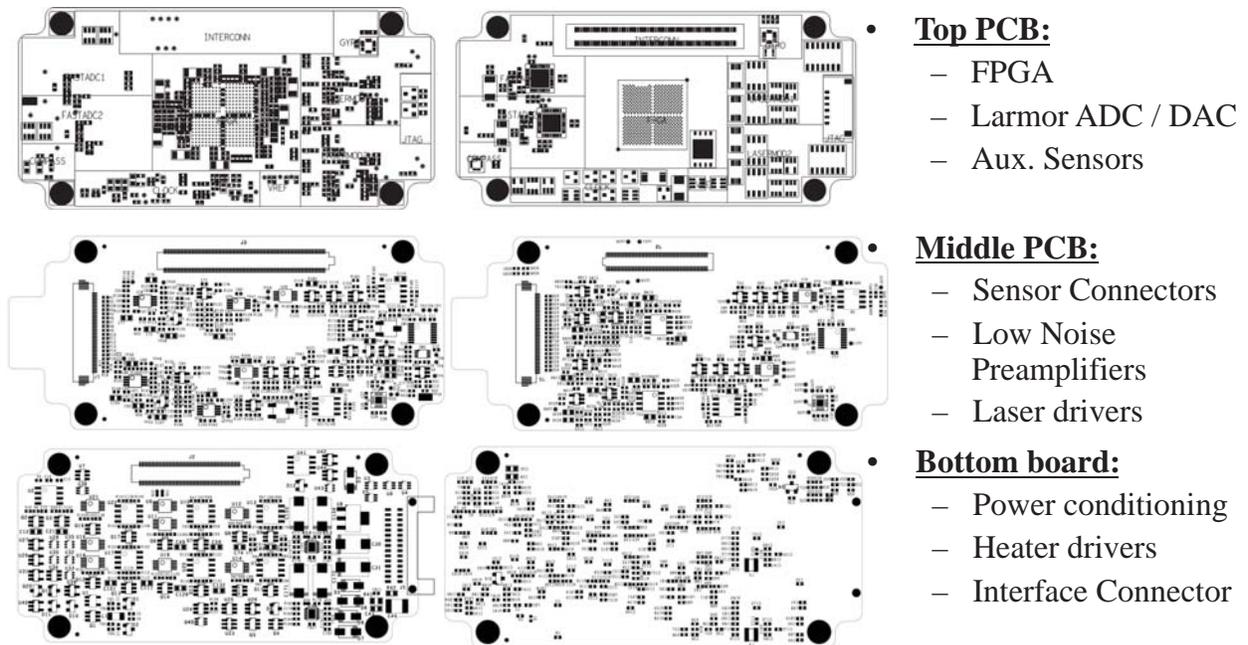


Figure 45. Knowing the exact design we needed, the final step was to package the circuits into the desired size and shape.

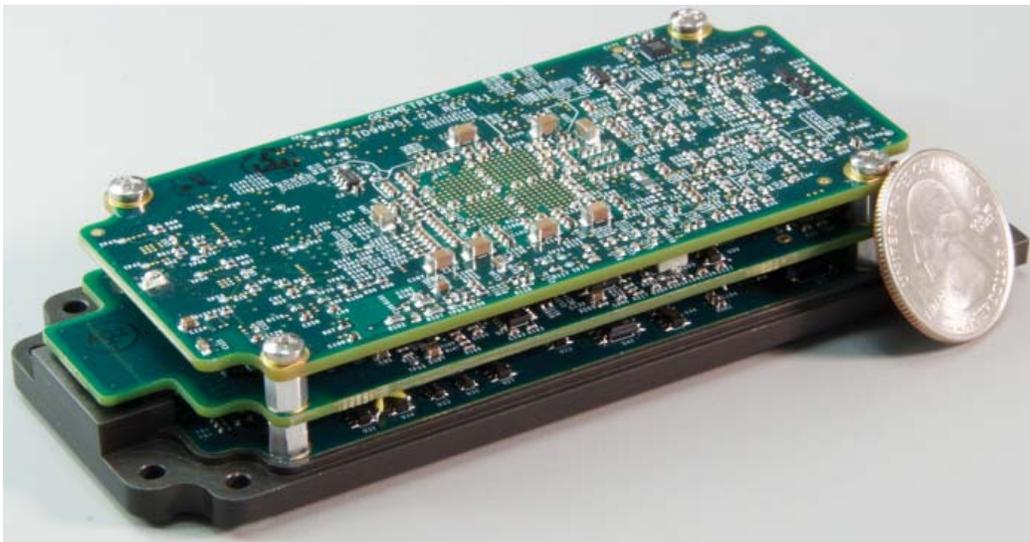


Figure 46. Completed electronics system.

7.1.1.1 MAGNETOMETER INSTRUMENT ARCHITECTURE

The overall system architecture is shown in Figure 47. The main system consists of the sensor, or physics package, the sensor driver, and the applications interface.

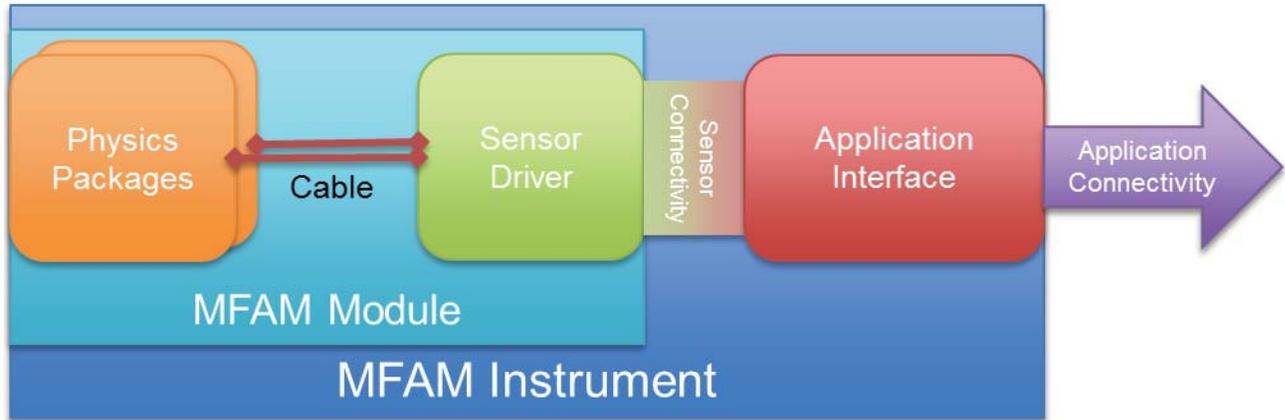


Figure 47. Architecture of a magnetometer based on the MFAM Module.

7.1.1.1.1 PHYSICS PACKAGE

The physics package contains the heaters, cesium cell, lasers, and optical components necessary to implement an atomic magnetometer. Designs of this were the focus of earlier SERDP-funded work (MR-1512). For reference, the diagram of a sensor physics package is shown in Figure 48

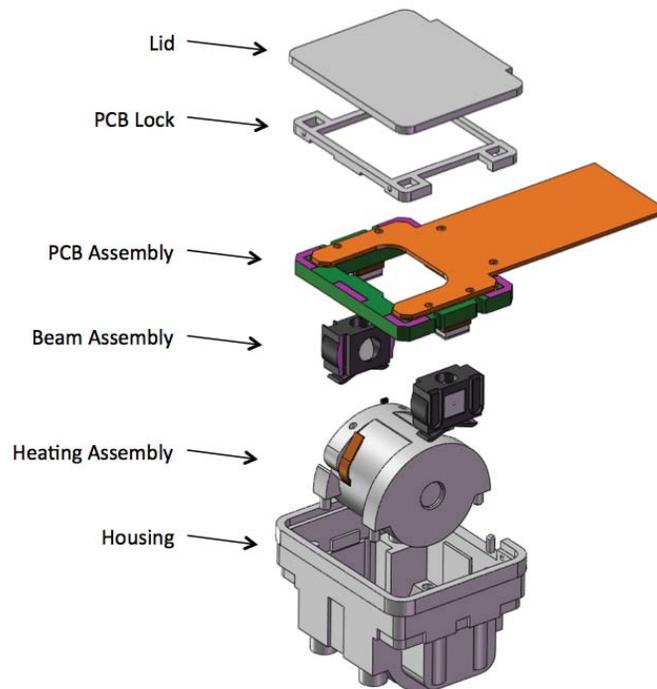


Figure 48. Physics package exploded view.

7.1.1.1.2 SENSOR DRIVER

The Sensor Driver is the electronics system described in this report. This system is responsible for running the physics package and deriving a magnetometer signal from it. The sensor driver system will operate two physics packages, which can be used either separately, as a gradiometer or dual sensor system, or in combination, to eliminate dead zones.

7.1.1.1.3 APPLICATIONS INTERFACE

The Applications Interface is responsible for collecting the data from the Sensor Driver, and packaging it for transmission to the user over convenient interfaces. This board will be specific to various applications, such as marine systems, land arrays, or, of course, the diver-held arrays towards which this project is directed.

7.1.1.2 ANALOG VS DIGITAL IMPLEMENTATION

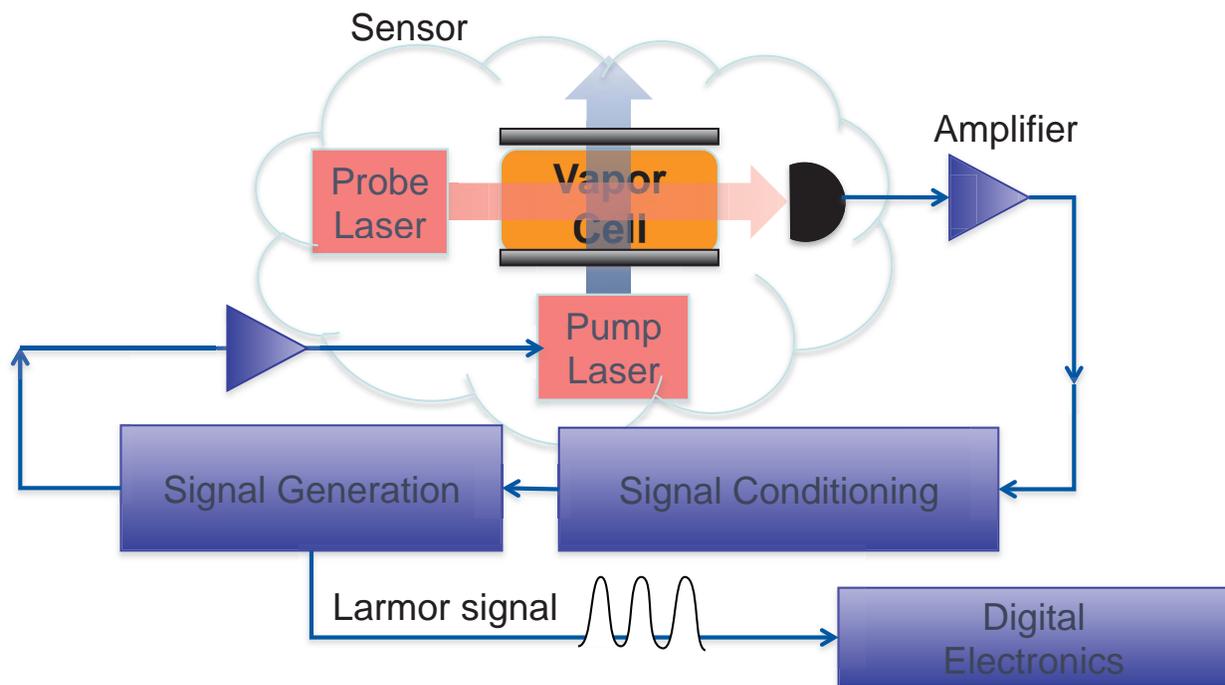


Figure 49. Analog method of creating Larmor signal.

The major elements of the electronics system are shown in Figure 49. The requirements for the basic elements are well understood, as they have all been prototyped and operated both in the lab and in outdoor testing. However, there are promising interrogation methods that are impractical to implement with off-the-shelf parts. In this task, we were able to implement both our existing interrogation methods (now in digital form), as well as enable additional methods, which further improve performance and lower power consumption, using custom integrated circuits.

Previously, using the analog methods, the electronics had been reduced in volume by a factor of 3 over previous commercial sensors. However, further reductions in size were desired to do justice to the small

size and power consumption of the sensing elements themselves. In this project, we designed and built a digital system (Figure 50) achieving an additional order-of-magnitude reduction in size and power consumption. Since the operation of the system is controlled digitally, the system is general purpose enough to operate the sensors in the full range of modes of operation we have explored.

The circuitry performs the following functions:

1. Convert the battery input voltage to the required voltages to operate the electronics components.
2. Heat the Cs cell and regulate its temperature.
3. Control the frequency of the laser diode.
4. Drive a signal into the laser diode with low noise circuitry.
5. Amplify the photodiode signal with low noise amplifiers.
6. Process the photodiode signal to extract the modulation components for controlling the laser frequency and the Larmor signal.
7. Measure the Larmor frequency and output the magnetometer reading determined by that frequency.

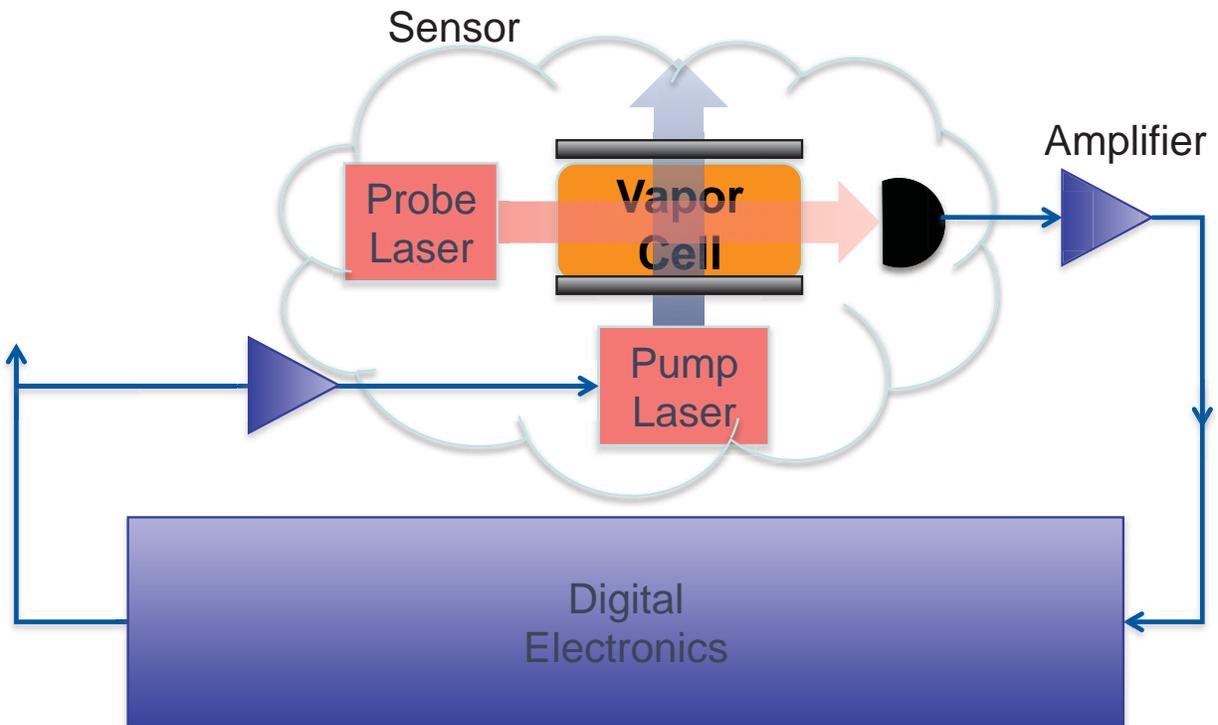


Figure 50. Digital method of creating Larmor signal.

One key advantage of the digital approach is that the magnetometer signal is actually very narrow band at any given instant. This signal is a sine wave at a frequency proportional to the magnetic field. However, the frequency may vary by a factor of five, tracking fields over a range from the lowest to the highest fields of interest. While analog circuitry would therefore need to cover this entire band, the digital filter loops may be adapted to the narrow-band requirement of the fields encountered at a given moment. This is a significant advantage of the digital approach.

The desired performance parameters were as follows:

Specification	Previous prototype	Design goal	Units
Size	100	10	In ²
Power consumption	3	0.5	W
Heading error	20	2	nT
Dead Zone Angle	30	0	Degrees
Sensitivity	15	1	pT / $\sqrt{\text{Hz}}$
Output sample rate	10	1000	Hz

7.2 SENSOR MODEL

The laser-pumped cesium magnetometer consists of two lasers (one pump and one probe) interrogating cesium atoms contained in a vapor cell. The transmitted light beams are detected using photodiodes.

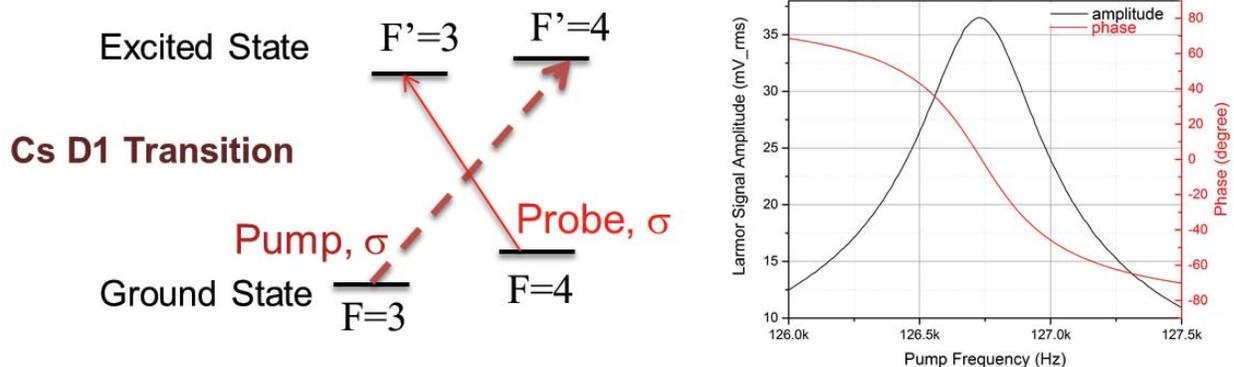


Figure 51. The physics model of the sensor operation can be described in terms of cross pumping and probing on the cesium D1 transition. The modulated pump beam sets up Larmor oscillations in the $F=4$ hyperfine ground state and the probe beam is used to detect the oscillations.

The externally accessible sensor parameters of interest are:

Inputs:

1. Pump laser current: I_p
2. Pump laser temperature: T_p
3. Probe laser current: I_r
4. Probe laser temperature: T_r
5. Vapor cell temperature: T_c

Outputs:

1. Probe photodiode current: I_o
2. Vapor cell thermistor: R_t
3. Pump photodiode current: I_d

The light from the pump laser, when it is tuned to the correct optical frequency, interacts with the atoms. The response of the atoms to the pump light depends on the magnetic field experienced by the atoms. The probe beam at a correct optical frequency detects the magnetic-field dependence of this response manifested as an oscillatory signal modulating its intensity as it propagates through the atoms.



Figure 52. Sensor Larmor signal path diagram. The system behaves as a resonant circuit to the electronics

Let us consider how the pump input current I_p modulates the probe photodiode output I_o depending on the magnetic field and the modulation frequency. The self-oscillating magnetometer architecture then includes a feedback system with I_o as the input and I_p as the output.

7.2.1 MODULATION: PUMP LASER

The injection current in a single mode diode laser changes both its optical power and optical frequency. Thus, a diode laser can be tuned to the atomic optical resonance by adjusting its DC current. An AC modulation can be added to the laser current to modulate the interaction of the laser light with the atoms. The total instantaneous pump laser current can be represented by the following equation:

$$I_p(t) = I_p^{dc} + I_p^{ac}(1 + m \cos(\omega t))\cos(\omega_c t)$$

Where:

I_p^{dc} : DC component of the pump laser current. At a particular laser temperature, this determines the optical frequency of the laser. The optical frequency of the laser needs to be precisely equal to the atomic absorption resonance frequency to within 10 parts per billion.

I_p^{ac} : AC component of the pump laser current. This AC amplitude needs to be set according to the optical resonance line width.

m: Modulation index

w: $2*\pi*$ AM frequency. This needs to equal the Larmor frequency for the sensor to resonate. Typical values range between 70kHz and 350kHz.

w_c : $2*\pi*$ Carrier frequency; an arbitrary value greater than 5MHz.

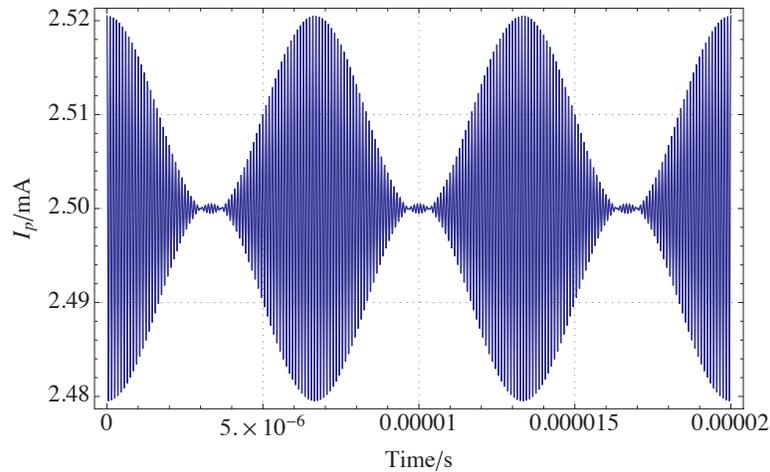


Figure 53. Pump laser current waveform.

7.2.2 ATOMIC ABSORPTION RESONANCE

The atomic optical absorption resonance has a damped oscillator characteristic. We can think of a virtual instantaneous optical pumping signal with the following relation to the instantaneous pump current:

$$R(\tau) = \frac{\Delta I^2}{(I_p(\tau) - I_c)^2 + \Delta I^2}$$

Where:

ΔI : The optical resonance half width at half maximum (HWHM) converted from laser frequency units to pump current units. This is possible because usually, the optical frequency of a diode laser is directly proportional to the injected current.

I_c : The pump current at which the optical frequency of the pump is equal to the atomic absorption resonance. I_p^{dc} is centered on this value using control loops. Because the optical frequency of the pump changes with the temperature of the laser, this number is strongly dependent on the laser temperature.

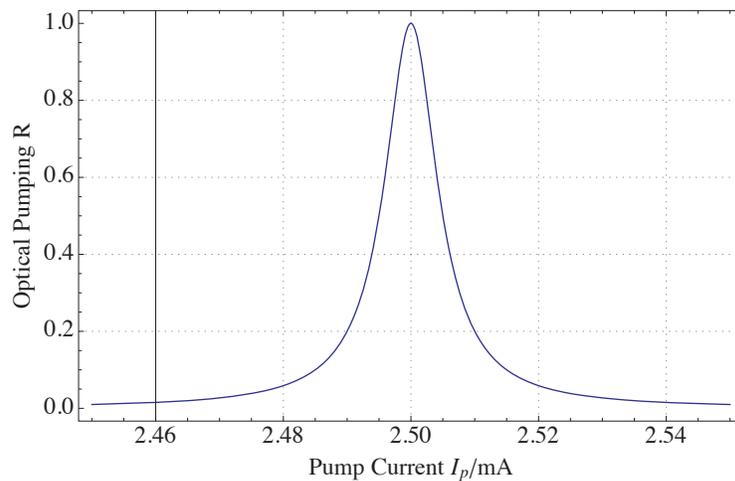


Figure 54. Atomic absorption resonance

A circularly polarized pump laser light tuned to the atomic optical resonance $F=3 \rightarrow F'=4$, as shown in Figure 51, maximally interacts with the atoms, thus transferring the atomic population to the non-interacting $F=4$ state. In this process called optical pumping, the light polarization is transferred to the atoms and the atoms become spin polarized. When the pump current is modulated with the waveform shown in Figure 53, the optical pumping process is instantaneously modulated as shown in Figure 55.

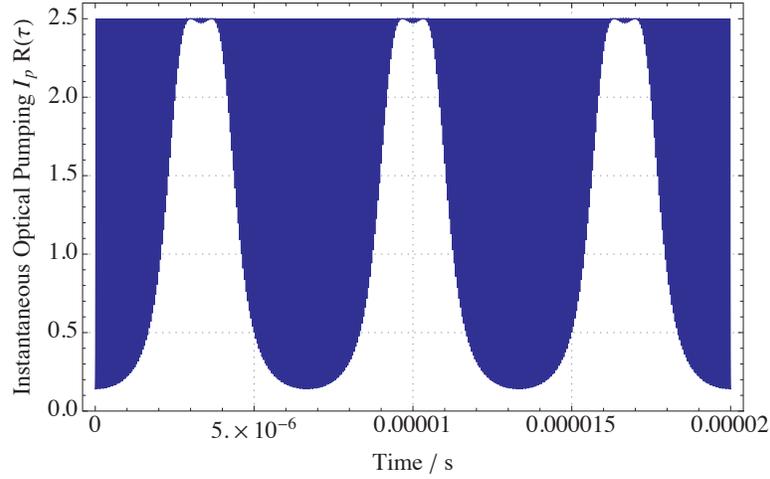


Figure 55. Instantaneous optical pumping rate.

However, the optical pumping process is not an instantaneous process and has an associated time constant that effectively acts as a low-pass filter for this waveform. Thus, the atomic absorption resonance demodulates the AM on the pump and gives rise to a virtual optical pumping signal $R_p(t)$ as follows:

$$R_p(t) = \frac{1}{T} \int_{t-T/2}^{t+T/2} I_p(\tau) R(\tau) d\tau = \frac{1}{T} \int_{t-T/2}^{t+T/2} I_p(\tau) \frac{\Delta I^2}{(I_p(\tau) - I_c)^2 - \Delta I^2} d\tau$$

The averaging time T is the optical pumping time constant, which depends primarily on the buffer gas composition and density.

It is instructive to look at the demodulated optical pumping waveform as the input to the magnetic resonance system as shown in Figure 56.

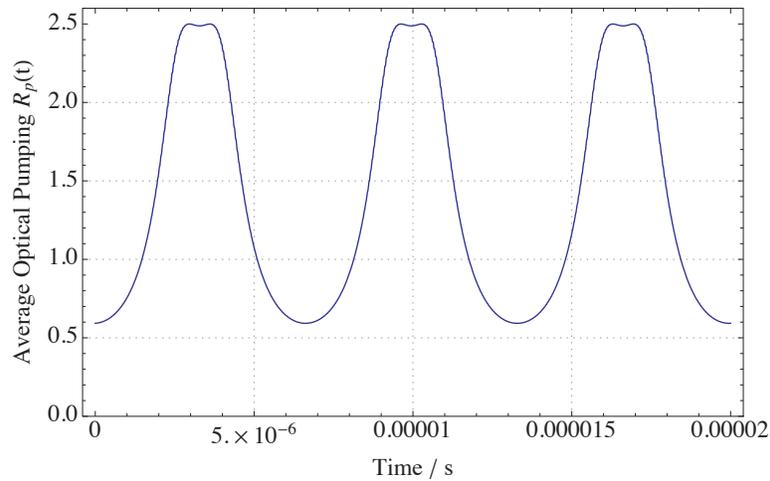


Figure 56. Demodulated optical pumping rate.

7.2.3 MODULATION TRANSFER: MAGNETIC RESONANCE

The optical pumping signal shown in Figure 56 drives the magnetic resonance. The magnetic resonance has the following s-domain and frequency-domain transfer functions:

$$G_m(s) = G_0 \frac{\frac{1}{Q} \frac{s}{\omega_B}}{\left(\frac{s}{\omega_B}\right)^2 + \frac{1}{Q} \frac{s}{\omega_B} + 1} \quad G_m(\omega) = G_0 \frac{\Gamma/2}{(\omega - \omega_B) + j\Gamma/2}$$

where

G_0 : Complex resonance gain.

ω_B : $2*\pi*$ Larmor frequency ~ 3.5 Hz/nT, with a range of $20 \mu\text{T}$ (70kHz) to $100 \mu\text{T}$ (350kHz).

Γ : $2*\pi*$ Full width at half maximum (FWHM) of the magnetic resonance, typically between 500Hz and 5000Hz. This is somewhat constant regardless of the magnetic field, which means the Q-factor of the resonance is proportional to the magnetic field.

Q: Q-factor of the resonance = ω_B/Γ . For the reason above, it is helpful to explicitly substitute Q with the more explicit expression in the above transfer function.

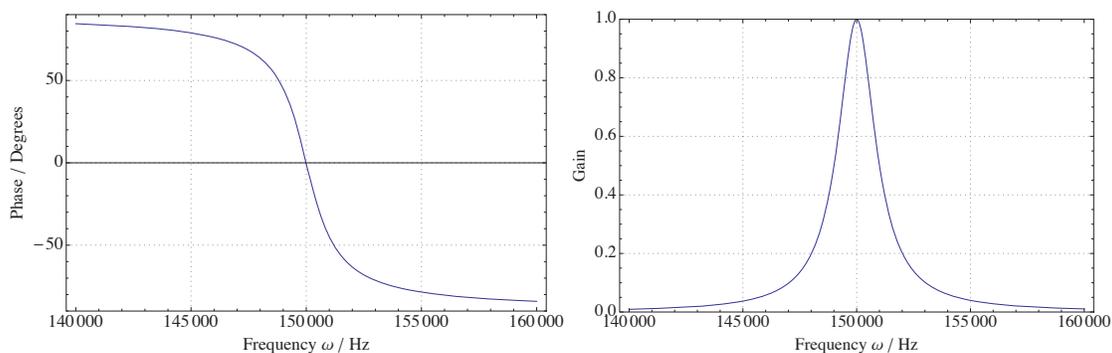


Figure 57. Phase and amplitude characteristics of the magnetic resonance.

7.2.4 DETECTION: PROBE LASER

Similarly to the pump laser, the probe laser is absorbed by the atoms when the probe optical frequency is equal to the optical frequency for the probe atomic transition as shown in Figure 51, which is $F=4 \rightarrow F'=3$ for Cs. The virtual optical pumping signal $R_p(t)$ is filtered by the magnetic resonance transfer function and modulates the absorption of the probe light. The probe light transmission is modulated such that when the optical pumping signal is high, less probe light is transmitted and vice versa. The transmitted probe light is detected on a silicon photodiode, thus generating a photocurrent described by:

$$I_o = AI_r[1 - \eta(\xi + G_m)R(\tau)] + I_n$$

A : probe beam attenuation, typically 0.5 to 0.8; multiplied by the probe laser current to power conversion ratio (0.06 mW/mA); multiplied by the photodiode quantum efficiency (0.65 mA/mW)

I_r : Probe laser current in mA (~ 2.5 mA)

η : Probe modulation efficiency

ξ : Cross-talk constant, typically 0.5 to 1

I_n : Output-referred current noise

A low-pass filtered output of the above waveform at the beginning of system startup is shown in Figure 58.

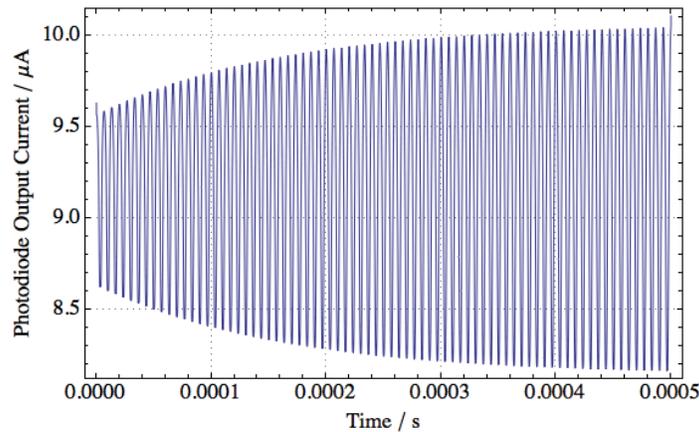


Figure 58. Low pass filtered probe laser signal.

7.2.5 CLOSED LOOP OPERATION

The probe photocurrent is conditioned by an analog front end and processed by a control loop to provide positive feedback to modulate the pump current. The pump laser is modulated with this signal, which is the $\cos(\omega t)$ term described in Section 7.2.1, thus closing the loop. In a given magnetic field, this closed-loop system rings up at the Larmor frequency as shown in Figure 59.

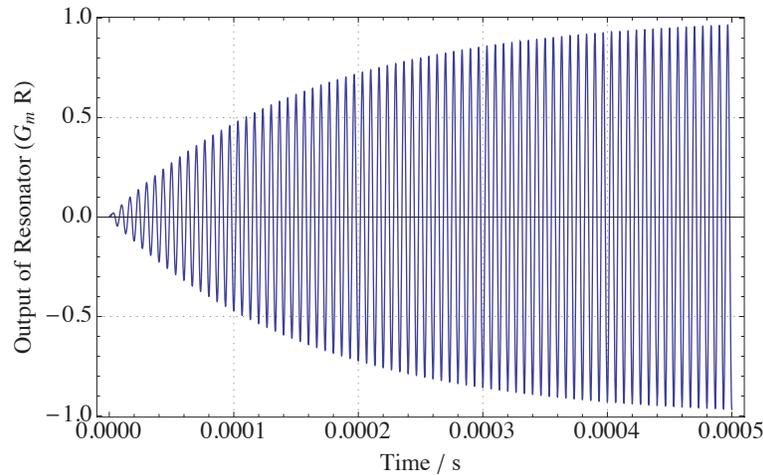


Figure 59. Growth of the oscillatory signal with positive feedback through the resonance.

This closed-loop model is then constructed as a test bench in the Simulink™ environment, as shown in Figure 60. This allows us to simulate the magnetometer operation and validate the model experimentally. The operational parameters, particularly the control loops, are then tuned using the model.

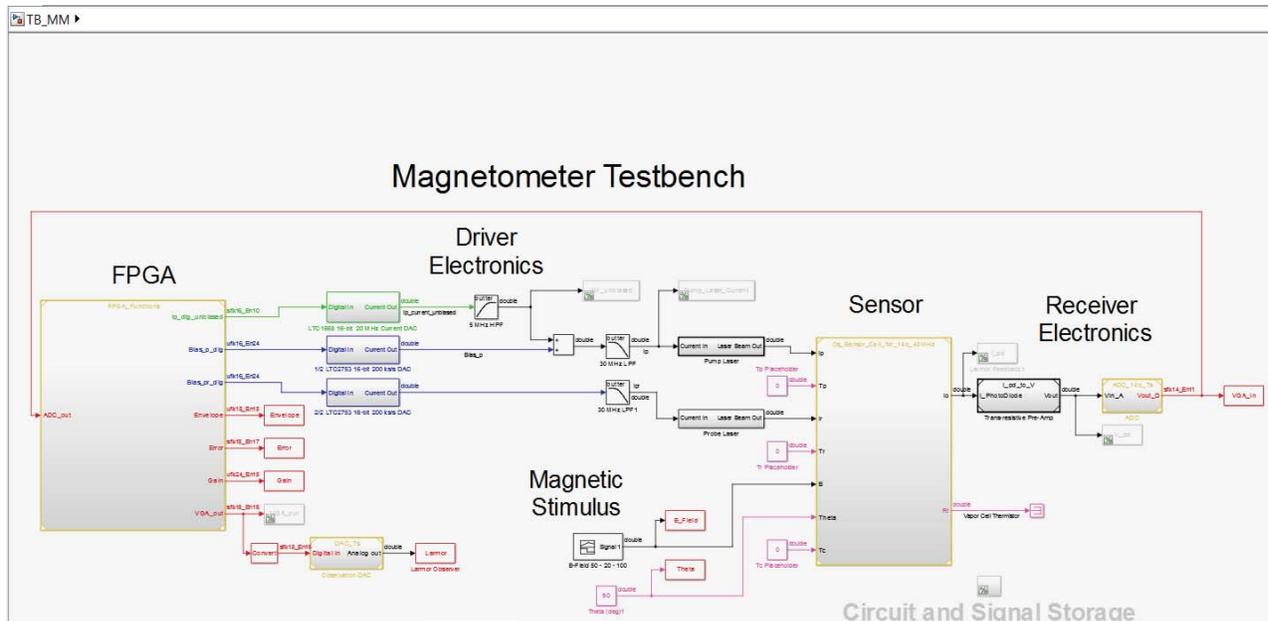


Figure 60. Simulink™ test bench used to model the closed loop magnetometer.

The model allows us to predict the operational envelope of the magnetometer as well, such as estimating the noise performance, as shown in Figure 61, and the size of the dead zones, as shown in Figure 62.

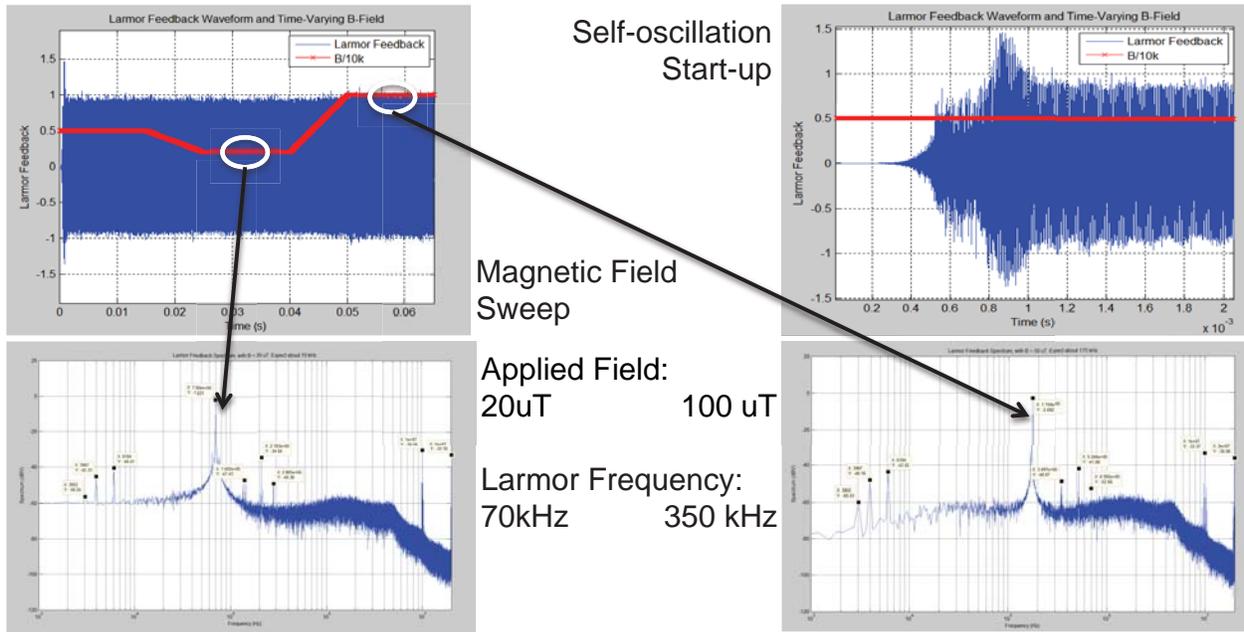


Figure 61. The results of our Simulink model predicted the proper operation and performance of the system.

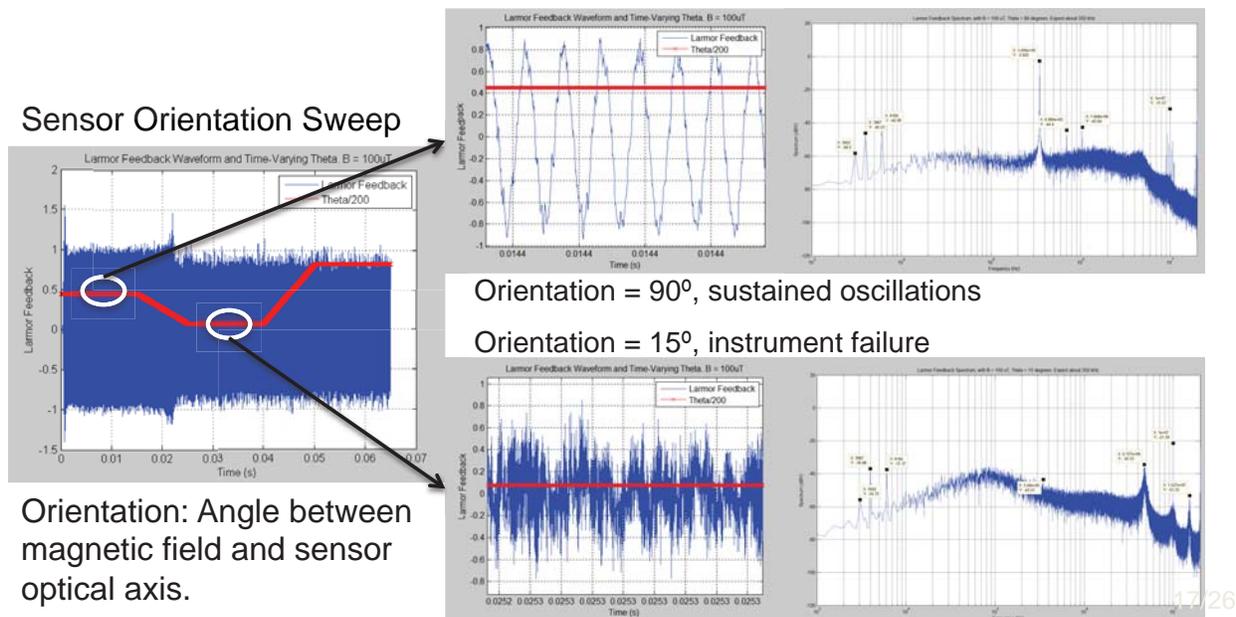


Figure 62. Even advanced performance characteristics, such as estimating the dead zone, could be explored in the simulation environment.

7.2.6 EXPERIMENTAL VALIDATION OF SENSOR MODEL

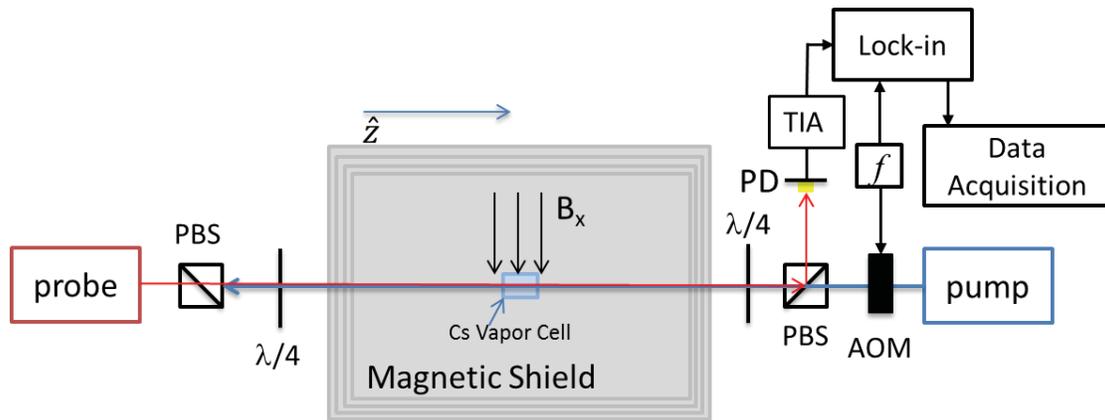


Figure 63. Schematics of the experimental setup.

In order to experimentally validate, implement, and extend the model, we simulated the sensor physics in a setup shown by the schematic in Figure 63. A cylindrical Cs vapor cell filled with N_2 buffer gas is placed inside a standard 4-layer magnetic shield can. The cell is heated to its operating temperature. A pair of Helmholtz coils generates a magnetic field perpendicular to the laser propagating direction. The pump and probe light from diode lasers is circularly polarized after going through a combination of polarization beam splitter (PBS) and quarter wave plate ($\lambda/4$). The pump laser passes through an acousto-optical modulator (AOM) to achieve amplitude modulation. After traversing the vapor cell, the probe light is detected by a photodiode (PD). The photocurrent is amplified by a transimpedance amplifier (TIA) and demodulated by a lock-in amplifier. The demodulated photocurrent amplitude and phase map the magnetic resonance as a function of the pump AM modulation frequency, f , as shown in Figure 64. The resonance has a full width at half maximum (FWHM) of about 500 Hz.

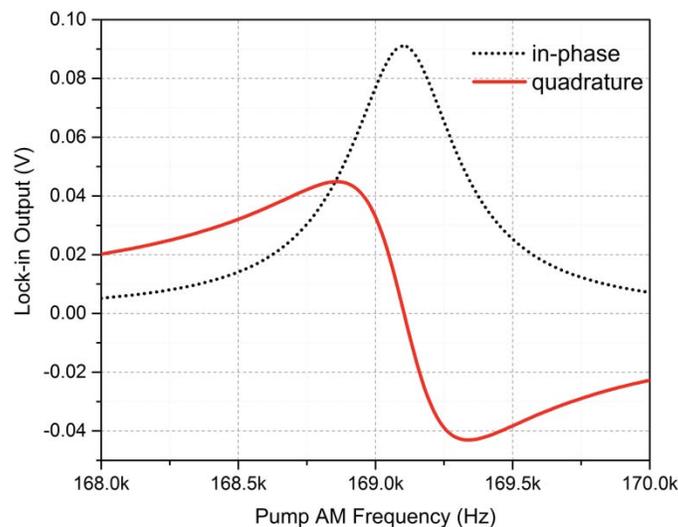


Figure 64. Probe lock-in signal (rms amplitude) as a function of pump AM frequency f . The black dot (red solid) curve represents the in-phase (quadrature) component.

The quadrature component displays a sharp slope, L , around Larmor frequency, f_L . When we set $f = f_L$, the output of the quadrature component is the measurement of the deviation of the magnetic field from the center field corresponding to f_L . Assuming a quadrature output of V (within the linear range of the slope), the measured magnetic field will be $B = (f_L - V/L)/k$, where $k=3.5$ Hz/nT is the Zeeman splitting factor between two adjacent magnetic sublevels. Accordingly, the noise equivalent magnetic field of the magnetometer in the open mode can be calculated through the noise spectrum of the quadrature output, δV , at $f = f_L$. After measuring δV , we calculated the magnetometer intrinsic noise according to the equation, $\delta B = (\delta V/L)/k$, represented by the green curve in Figure 65. As seen, above 30 Hz, the magnetometer noise is well below 200 fT/ $\sqrt{\text{Hz}}$. The increase of the noise level below 30 Hz is mainly due to air currents along the beam path causing the air's refractive index to fluctuate. In addition, room temperature fluctuation and laser instability (both lasers are NOT frequency-stabilized) also contribute to the higher noise at low frequency. The sharp spikes in Figure 65 are 60 Hz noise due to the power line and its harmonics.

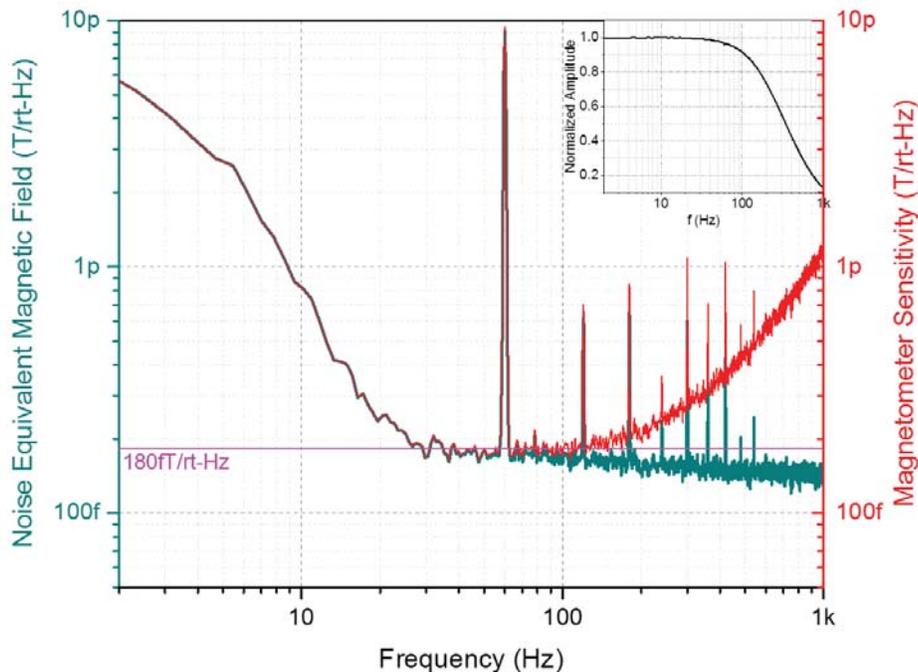


Figure 65. Open-loop spectral density of the magnetometer noise (Green curve). Inset shows the bandwidth of the magnetometer. The red curve is the sensitivity of the magnetometer, which represents the minimum signal level detectable by the magnetometer with a signal to noise ratio of 1:1.

Another important performance characteristic of the magnetometer is its bandwidth. To measure the bandwidth, we generate an oscillating magnetic field, $B_1 \sin(2\pi f_1 t)$, with fixed amplitude B_1 of about 3nT and varying frequency f_1 . The oscillating magnetic field can be detected by the magnetometer and the measured amplitude B_1 as a function of f_1 is recorded. The normalized B_1 vs. f_1 is shown in Figure 65 inset. As shown, the 3dB point of the bandwidth curve is about 210 Hz. The bandwidth is limited by the width of the magnetic resonance shown in Figure 64. By combining the magnetometer noise spectrum

and its bandwidth, we can calculate the sensitivity of the magnetometer. We define the sensitivity as the amplitude spectral density of the magnetic field detectable by the magnetometer with a signal to noise ratio (SNR) of 1:1. According to this definition, the sensitivity curve can be calculated by dividing the noise spectrum by the bandwidth curve. The sensitivity as a function of frequency is presented by the red curve in Figure 65. At low frequencies, the sensitivity and the noise curves overlap. At higher frequencies where the signal starts to drop due to the bandwidth, the sensitivity is getting worse since the noise remains almost the same. This, if a sensitivity of $1\text{pT}/\sqrt{\text{Hz}}$ is desired, the magnetometer bandwidth can be extended to 1kHz.

7.3 IMPLEMENTATION

This section describes the hardware implementation of the magnetometer. We designed the electronics with specifications derived from the system, which we modeled and simulated using the work described in the previous section.

7.3.1 SYSTEM ARCHITECTURE

The systems required to run the sensor are rather complicated. The sensor model and the experimental simulation enabled us to identify the critical design areas to deliver the required performance and partition the design to modularize the hardware development and testing. This effort led us to identify the MFAM system architecture as shown in Figure 66. The key building blocks in the electronics are the analog front-ends comprised of the laser current drivers, the transimpedance amplifiers and digitizers, the resistive heater drivers and the digital processing firmware programmed in the FPGA.

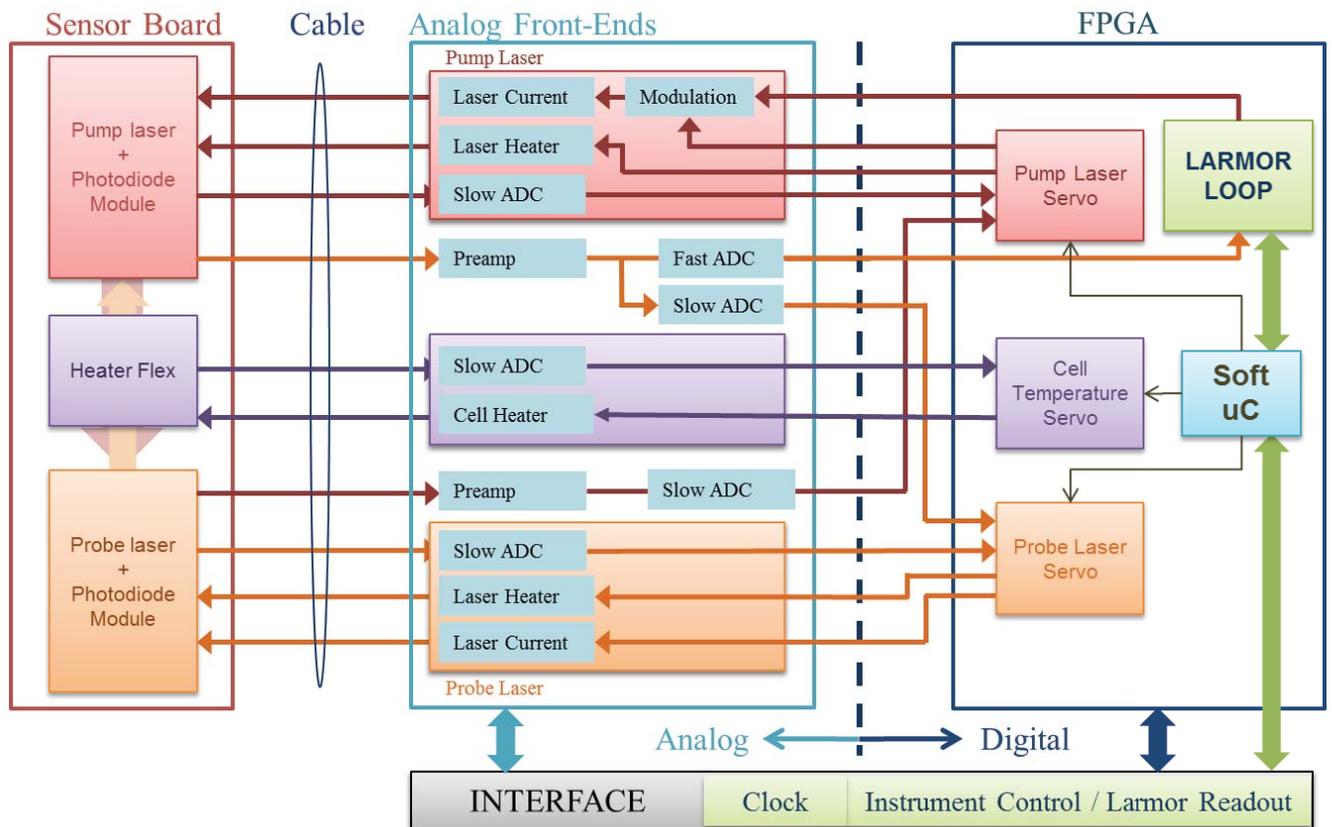


Figure 66. MFAM System Architecture.

The primary interdependencies of the various modules are highlighted in the system architecture shown in Figure 66. The sensor consists of three equivalent functional subsystems: lasers, atomic vapor cell and photodiodes, which are quite distinct from the physical subsystems used in assembling the sensors. For instance, the laser and photodiode module is implemented as a single physical subsystem in the sensor, whereas functionally they are on the opposing sides of the atomic vapor cell, which is the primary magnetic field sensing transducer. The analog front-ends (AFEs) drive the lasers and receive the photocurrents using low-noise amplifiers (LNAs). The AFEs also consist of power amplifiers (PAs) to drive the resistive heaters and preamplifiers to condition the signals from the temperature sensing elements. The AFEs interface with the field-programmable gate array (FPGA) through analog-to-digital converters (ADCs) or digitizers and digital-to-analog converters (DACs). The control loops and signal generation and conditioning algorithms are implemented in the FPGA and controlled by the MicroBlaze™ soft embedded microcontroller that is also programmed in the FPGA. The soft microcontroller is responsible for system management as well as interface communication.

7.3.2 TRANSIMPEDANCE AMPLIFIER FRONT-END

The key building block for the signal extraction is the transimpedance amplifier, which receives the signal from the probe photodiode and conditions it for further analysis. This topology of the transimpedance amplifier is shown in Figure 67. The amplifier design balances opposing design goals of ultra-low noise in a wide bandwidth and low power consumption. The amplifier is capable of delivering shot-noise limited photocurrent measurement in the Larmor band of 70kHz to 350kHz while avoiding signal attenuation or a significant group delay. The topology achieves this design goal by isolating cable capacitance from photodiode capacitance as shown in Figure 67.

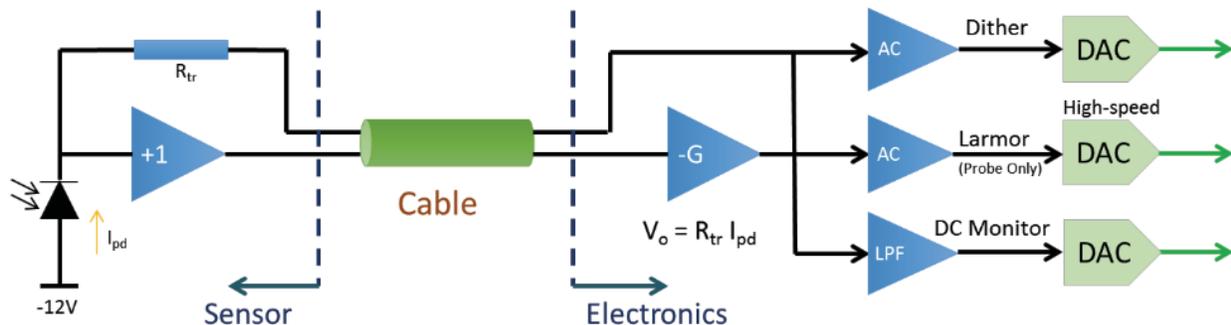


Figure 67. Distributed transimpedance amplifier topology.

The buffering amplifier, placed physically close to the photodiode situated in the sensor, provides a low impedance drive to the cable. This is a low power LNA with a bandwidth that extends well past 5MHz so that it doesn't introduce significant phase shift at Larmor frequencies. Using this topology allows the amplifier stage that is within the electronics to have design requirements that are practical for integration.

There is also a power efficiency advantage to this method. By isolating the photodiode from the capacitance of the cable, the circuit avoids signal attenuation by this capacitance. This removes the need for a high gain LNA that would otherwise be needed to compensate for this attenuation. Also, because the buffering amplifier is enclosed within the loop of the transimpedance circuit, it only needs to produce low amplitude at its output. This allows this stage to be powered at a reduced voltage.

With gain "G" sufficiently large, the gain of this circuit doesn't affect the overall performance. This high gain value only needs to be maintained up to the highest Larmor frequency expected. This stage is

receiving a low impedance input signal. These facts greatly ease the requirements on this amplifier making its eventual integration more practical. The gain and phase response of the transimpedance amplifier are shown in Figure 68. The gain axis is referenced to an arbitrary value. However, note that the -3dB point, i.e. the bandwidth, is 11 MHz. Additionally, the flat phase response allows straightforward Larmor loop implementation in the FPGA.

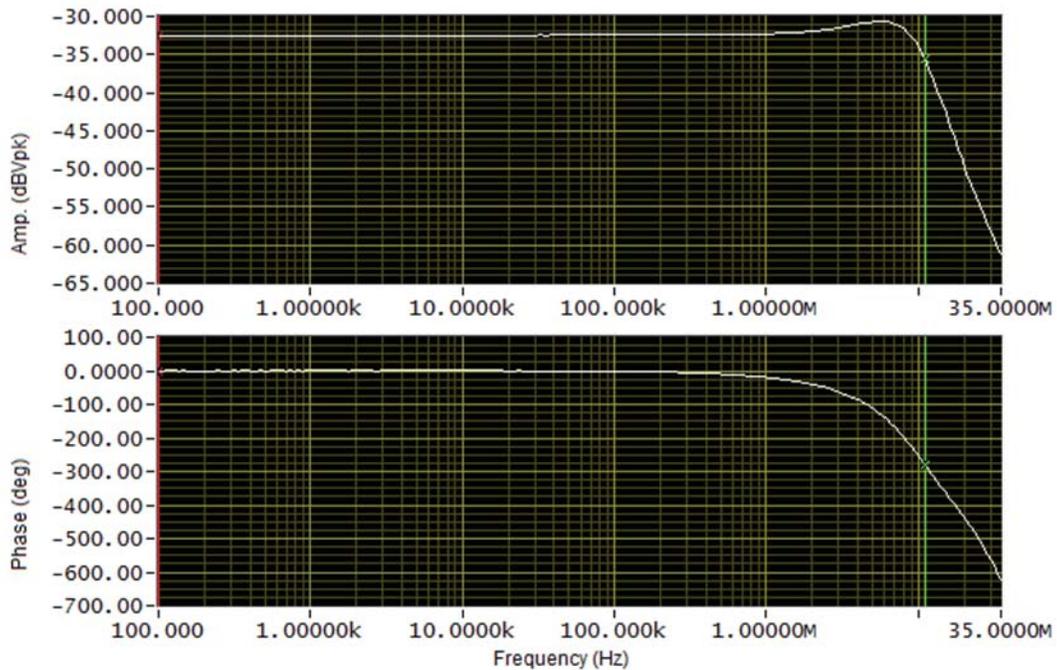


Figure 68. Transimpedance amplifier transfer function characterization. Note the flat phase response in the Larmor band or 70kHz to 350kHz.

The signal from the transimpedance stage has three components, with the Larmor signal being the most important component for the field measurement. There are signals in the few kHz range and others at low frequencies that are required for stabilizing the laser optical frequency, also referred to as laser locking. The components used for laser locking are filtered out of the Larmor signal before it is sent to the high performance ADC. This allows the full span of this ADC to be used for the Larmor signal.

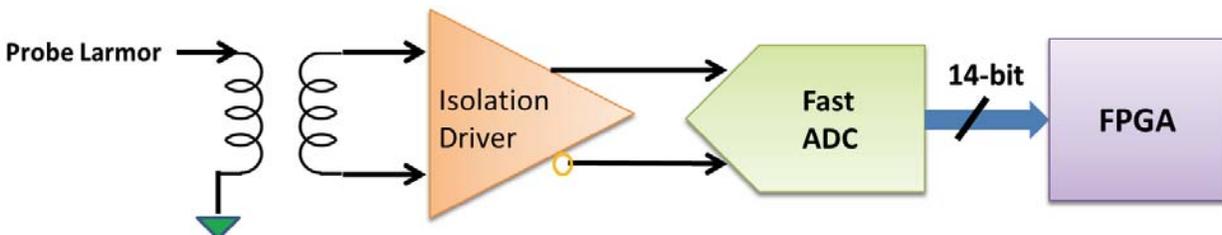


Figure 69. Block diagram of Larmor signal quantization.

There is one high-speed ADC circuit for each physics package with the topology shown in Figure 69. The high-speed ADC circuit takes the probe Larmor signal from the transimpedance amplifier, converts it from analog to a digital signal and then sends it to the FPGA to close the Larmor loop. The fast ADC circuit mainly consists of a low noise, rail-to-rail fully differential driver amplifier and a low power, 14-bit, 40Msps ADC.

As part of the laser locking method, the laser's wavelength is varied slightly at a rate of several kHz. This causes a very small variation in the optical absorption through the atomic vapor cell. This signal appears on the output of the transimpedance amplifier where it is amplified so that a low power ADC can be used to digitize this waveform. The Larmor signal is filtered out of this signal only to the degree required to allow the amplifier to not overload the ADC. Variations in the light reaching the photocell at frequencies below 100Hz are also monitored using another ADC.

7.3.3 LASER DRIVER FRONT-END

There are two lasers in every sensor, one acting as the pump and one as the probe. The pump laser is modulated as described in Section 7.2.1 and the probe laser light is modulated by the Larmor signal transferred through the atoms under the magnetic resonance conditions. Both the pump and probe laser optical wavelengths have to be stabilized to the appropriate atomic transitions for this process to function correctly. The lasers used in the sensors are semiconductor diode Vertical Cavity Surface Emission Lasers (VCSELs).

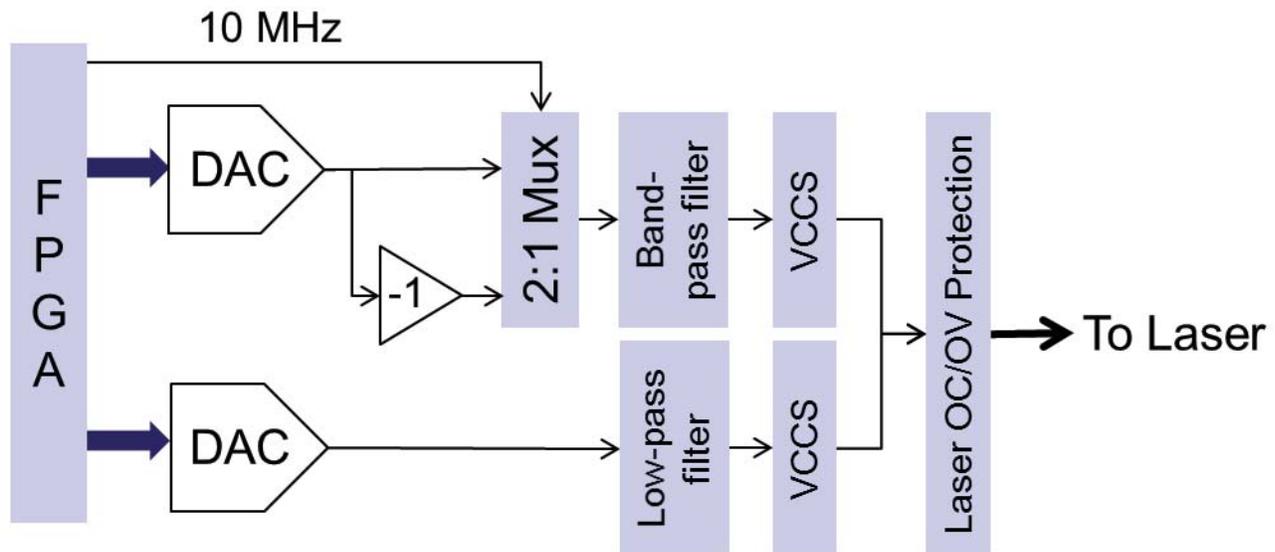


Figure 70. Laser driver front-end topology for the pump laser. The probe laser only requires the lower signal path.

The optical power output and the wavelength of the lasers depend strongly on the forward current injected in the diode. The wavelength of the lasers also strongly depends on the laser diode temperature. Thus, the laser temperature and current can be used in tandem to actively control both the laser power and wavelength. Thus, the laser current drive has to have enough bandwidth to be able to control the laser wavelength. At the same time, the circuit has to have extremely low noise in order to mitigate noise on the optical signal. Also, the drive has to exhibit high output impedance so that it is not affected by the laser forward voltage. The pump laser driver circuitry also has to provide the radio frequency modulation to the laser. The topology of the laser current drive section is shown in Figure 70.

7.3.3.1 LASER DRIVER DC SIGNAL PATH

The low frequency analog signal that drives the lasers is produced by a DAC driving the low pass filter. This signal is comprised of three components. The DC component sets the current to be applied to the Laser to obtain the desired optical power in the laser beam. A small variation around this optimum point

is used to keep the Laser at the correct wavelength if there is a sudden change in the ambient temperature. A small “dither” sine wave at several kHz is used to slightly modulate the laser wavelength as part of the laser locking method. This signal path circuitry is implemented for both the pump and probe lasers.

The low-pass filter is set to a bandwidth less than the several kHz frequency of the “dither” sine wave. This causes a significant attenuation at this frequency allowing the FPGA to construct the sine wave at larger amplitude to improve the signal fidelity of this waveform. The circuit design needs to account for the low power and low noise requirements of the system. The gain and phase response of the low frequency analog section of the laser driver are shown in Figure 71.

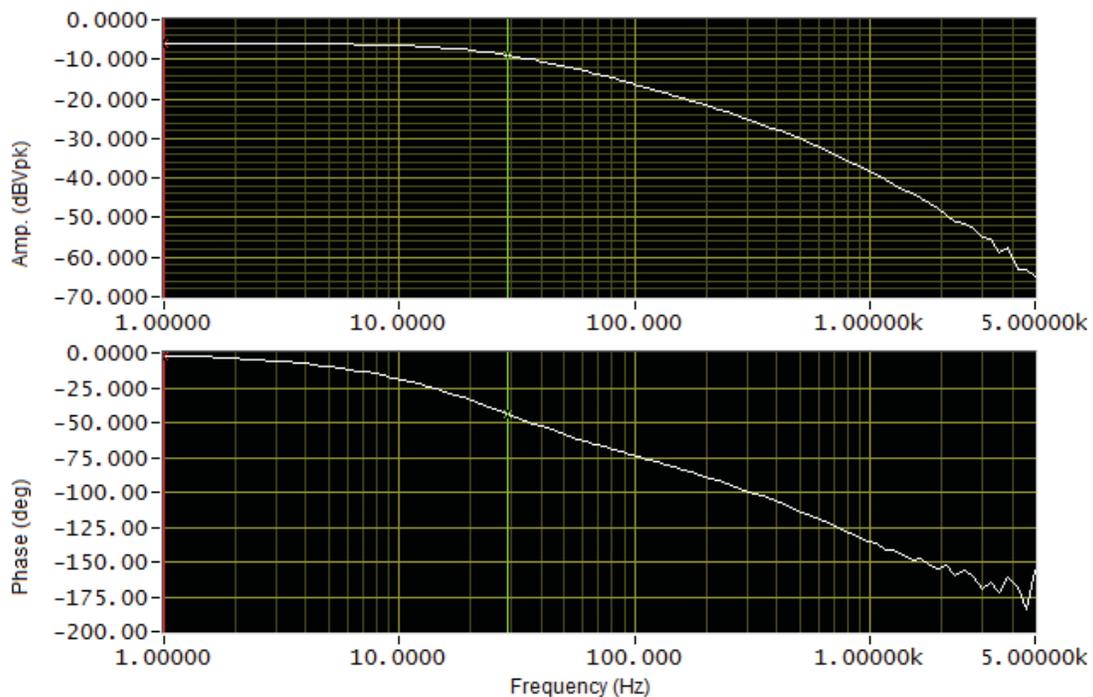


Figure 71. Frequency response of the low-pass filter for the laser driver low frequency analog signal path.

The Voltage Controlled Current Source (VCCS) is designed so that it maintains a very high output impedance between 0V and 4V so that the calibration of the applied laser current is independent of the laser used. The physics of the semiconductor laser assures that its forward voltage is less than 3V. The compliance range is maintained for currents from 100uA to the full scale current of 3mA. The laser protection circuitry is partly implicit in the VCCS circuits. The current source circuit normally maintains a very high output impedance. If, during operation, an electrostatic discharge attempts to apply a reverse voltage to the laser, the current source output impedance is reduced to clamp the voltage at a safe level. When the main power is off, an additional circuit applies a low resistance path shorting out the terminals of the laser diode.

7.3.3.2 LASER MODULATION SIGNAL PATH

The laser modulation circuitry provides amplitude-modulated radio-frequency current to the pump laser only. While this signal can be fully synthesized digitally in the FPGA, it would require a high speed DAC. Another design consideration was to ensure that transferring the implementation to Application-Specific Integrated Circuits (ASICs) should be fairly straightforward, to the extent that could be envisioned in the design phase. High speed DACs to generate RF in the 10 MHz frequency range typically have very high power consumption. These components are also not easy to integrate into ASICs. Additionally, the waveform synthesis puts computational burden on the FPGA. We developed the modulation topology shown in Figure 70 to use standard DAC integrated circuits (ICs), minimize RF electromagnetic interference (EMI), maintain signal integrity and provide a low power solution.

This modulator circuit needs both a positive and negative analog signal at the Larmor frequency. The DAC constructs the positive portions based on the digital signals from the FPGA. This signal is filtered with a low pass filter having the gain characteristics shown in Figure 72. The negative portion is created by having an amplifier with a gain of -1 and a bandwidth much greater than the Larmor frequency. This ensures that the two portions of the signal are matched to within 2%.

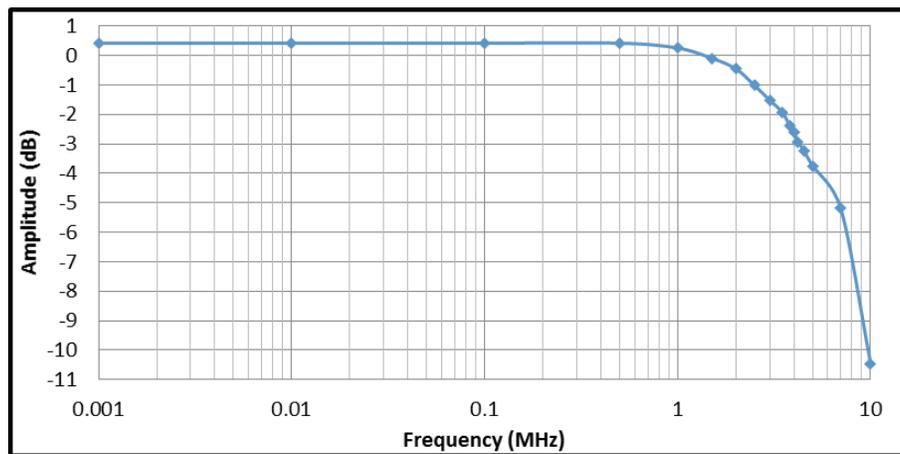


Figure 72. The gain response of the output filter for the DAC that generates the Larmor signal.

The FPGA circuits provide a 10MHz square wave with exactly 50:50 duty cycle and extremely low jitter. If the logical “true” of this signal is considered to mean “+1” and the logical false is considered to mean “-1”, this circuit can be seen to function as a multiplier. The 10MHz applied to the band-pass filter is amplitude modulated by the signal derived from the DAC. There will be less than 2% base band leakage in this modulator. The multiplexer-based modulator design has the additional advantages that it can easily be implemented in common semiconductor processes and is extremely power efficient. The waveform from the DAC and the AM RF carrier form the band pass filter are shown in Figure 73.

The band pass filter is designed to meet three main requirements:

1. It must block any base band signals.
2. It must pass the side bands of the 10MHz carrier.
3. It must attenuate the 3rd harmonic of the 10MHz sufficiently to allow a low power amplifier to be used in the Voltage Controlled Current Source (VCCS) circuit.

Keeping in mind the goal of a future further miniaturization, this filter is implemented without the use of inductive elements.

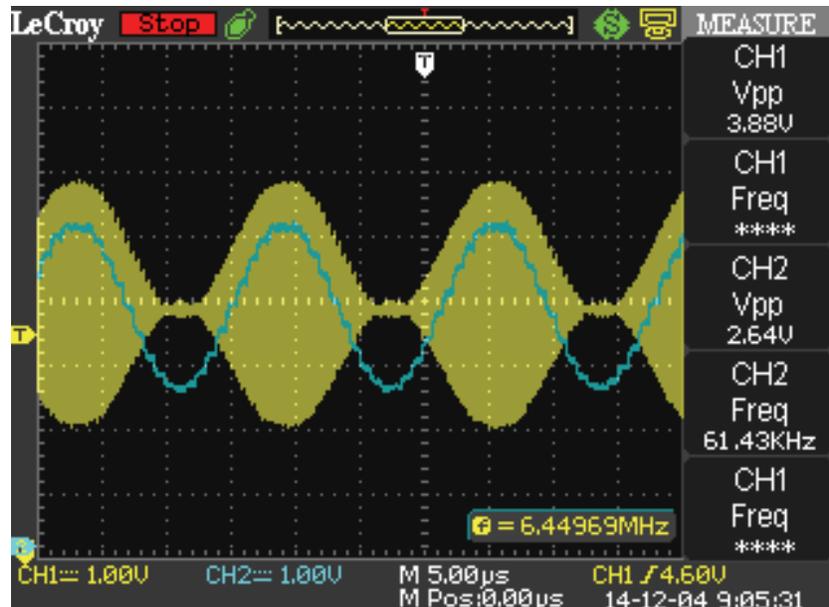


Figure 73. The yellow waveform shows an RF carrier at 10MHz modulated by the Larmor signal (blue waveform). Note the delay between the driving signal and the modulated signal.

The current source in the modulation path must provide both positive and negative currents with a bandwidth exceeding that of the modulated 10MHz while maintaining high output impedance. Because the two signals applied to the laser are in the form of currents from high impedance sources, they can be summed by simply connecting the two together.

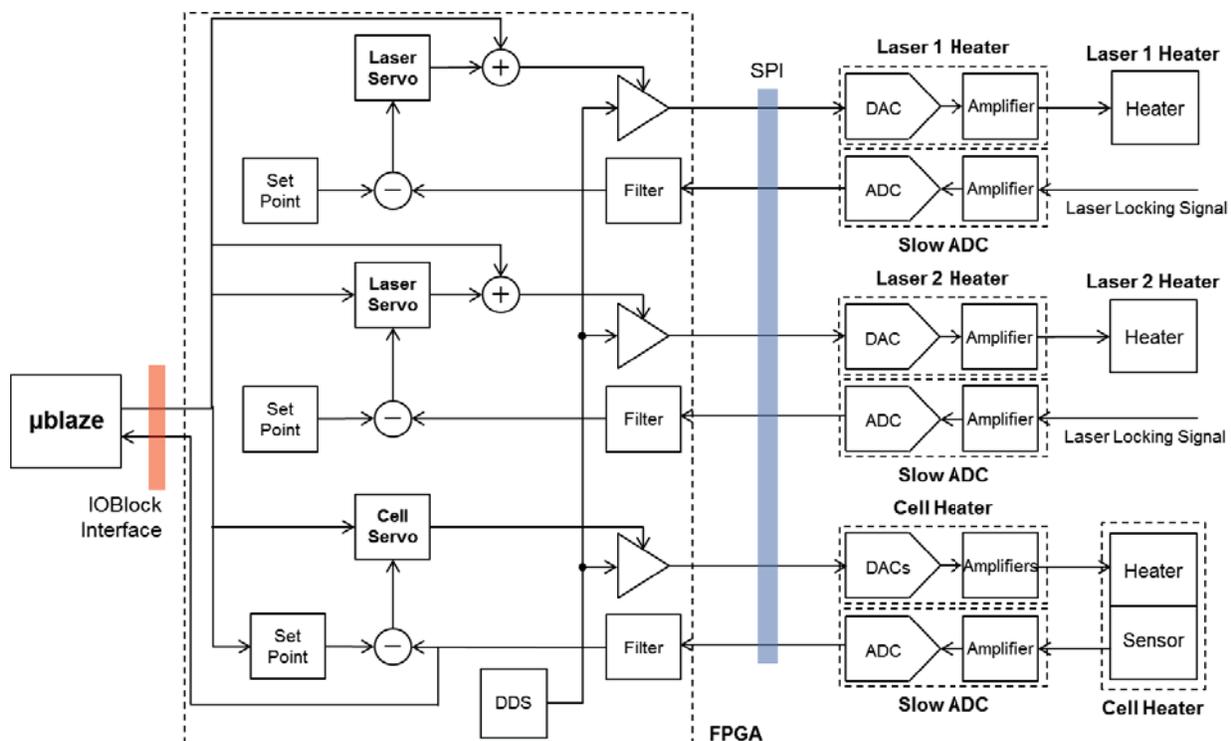


Figure 74. Topology of various heater control loops in the MFAM.

7.3.4 HEATER DRIVER FRONT-END

The lasers and the vapor cell are the critical components in the sensor. These components need to be maintained at a fixed temperature with very small tolerance to enable sensor operation. The temperature regulation needs to be effective in ambient operating temperatures ranging from -35°C to $+50^{\circ}\text{C}$. This requirement places very stringent constraints on the temperature regulation loops and the thermal design of the sensor itself. The thermal properties of the sensors play a crucial role in tuning the various control loops to ensure stable operation in the presence of the large temperature variation that can be encountered in the field, for instance, when the sensors are taken out of shade and exposed to the sun. The topology of the various heater control loops is illustrated in Figure 74.

7.3.4.1 VAPOR CELL HEATER

The atomic vapor cell is filled with nitrogen buffer gas and has a few micrograms of cesium in it. The cesium naturally evaporates and at any given moment, there is partial vapor pressure due to cesium gas in the cell. The density of cesium vapor is highly temperature dependent and doubles for every 10°C rise in temperature. Because the atoms are the transducing element that essentially converts the magnetic field to frequency, a higher density of atoms results in a higher signal level. On the other hand, the atomic vapor becomes opaque to resonant light at very high densities and attenuates the signal. The vapor cell has to be held at a temperature where these two effects are balanced, constant to within $\pm 0.1^{\circ}\text{C}$, in order to generate a stable signal.

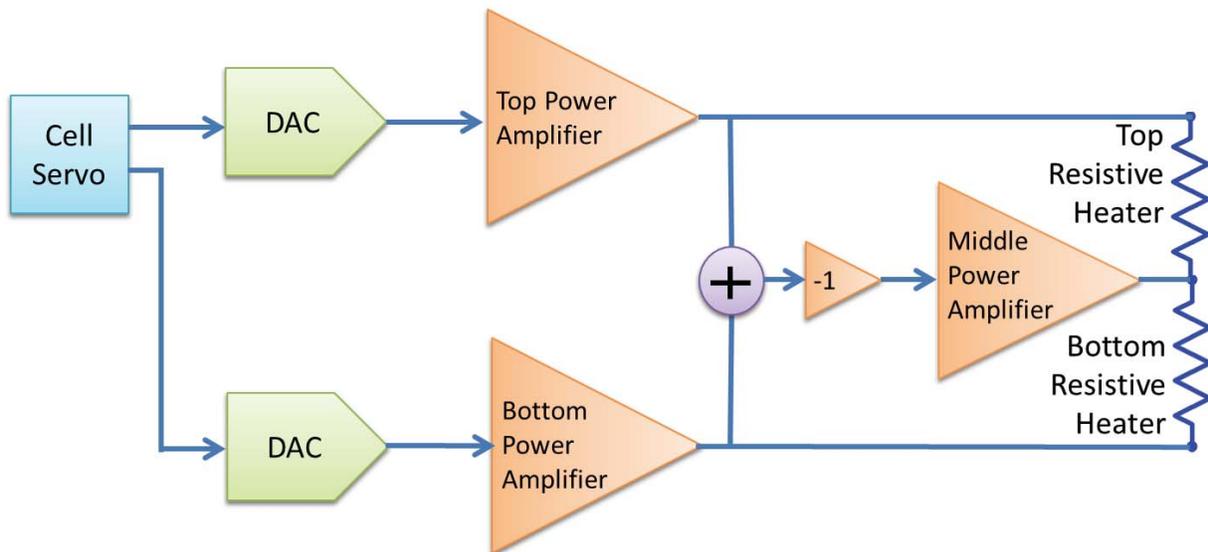


Figure 75. Vapor cell heater driver block diagram.

As shown in Figure 74, the control loop to regulate the vapor cell temperature, labeled as the “Cell Servo” is implemented in the FPGA. The output of a temperature sensor attached to the cell inside the magnetometer physics package is digitized, digitally filtered and compared to the set point to generate an error signal. The cell servo uses this error signal to communicate the desired power output to a digital amplifier within the FPGA.

The cell is heated with an AC waveform that is generated by Direct Digital Synthesis (DDS) embedded in the FPGA. The amplifier adjusts the gain and phase of this waveform according to the Cell Servo output.

This waveform is then fed into the cell heater driver circuit illustrated in Figure 75. The topology consists of a balanced driver that delivers power to two composite resistive heating elements in the sensor. The power can be coarsely adjusted using the waveform amplitude. The relative phase of the heater driver provides a finer control over the power delivered. This topology also ensures very high efficiency of power conversion from DC to heating power, thus lowering the overall power footprint of the system.

There is one cell heater driver circuit for each physics package, totaling two cell heater driver circuits. The cell heater driver circuit is located on the power board (the bottom board of the 3-board stack), and the 2 cell heater resistors are located in the sensor package. The cell heater is controlled by the cell temperature servo to maintain the cesium vapor cell at operation temperature, usually from 60 to 70 degree Celsius.

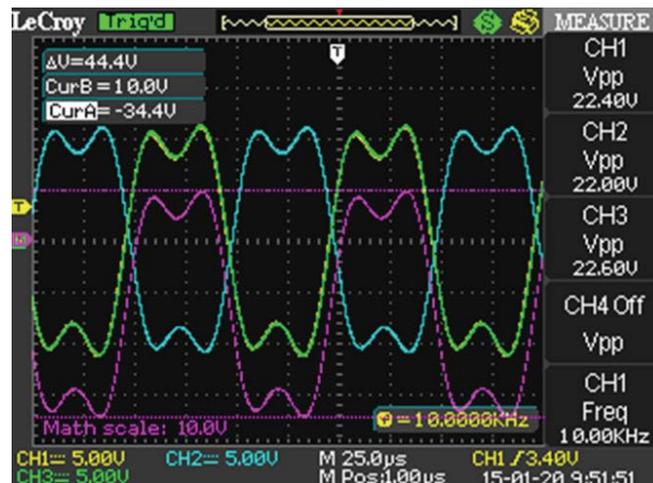


Figure 76. Cell heater waveforms.

The cell heater driver circuit has mainly two DACs and three power amplifiers with a ± 12 V swing. The three cell heater driver outputs are 10kHz AC waveforms, shown in Figure 76, in which the top and bottom outputs (overlapping green and yellow curves) are in phase and out of phase respectively relative to the middle output (blue curve) by 180 degrees. The effective voltage swing per resistive element is shown by the magenta curve. The three cell laser heater driver outputs go side by side from the power board all the way to the heater resistors at the sensor, such that there is no net magnetic field generated.

7.3.4.2 LASER HEATER

As mentioned previously, the optical power and wavelength of the laser are a function of both the injection current and the die temperature. Typically, for operating the laser at the atomic transition wavelength of cesium, the die temperature needs to be between 70°C and 90°C. Hence, the lasers are mounted on resistive heating elements. The wavelength of the laser responds to current tuning much faster than to temperature changes. Hence, current tuning is the preferred method to stabilize the laser wavelength to desired atomic absorption spectral feature. However, the optical power output of the laser changes with the current as well. To counteract that, the laser current actuation signal acts as the input to the laser heater servo. The laser heater servo, implemented in the FPGA, adjusts the die temperature such that the laser current actuation signal produces the desired optical power as well as wavelength.

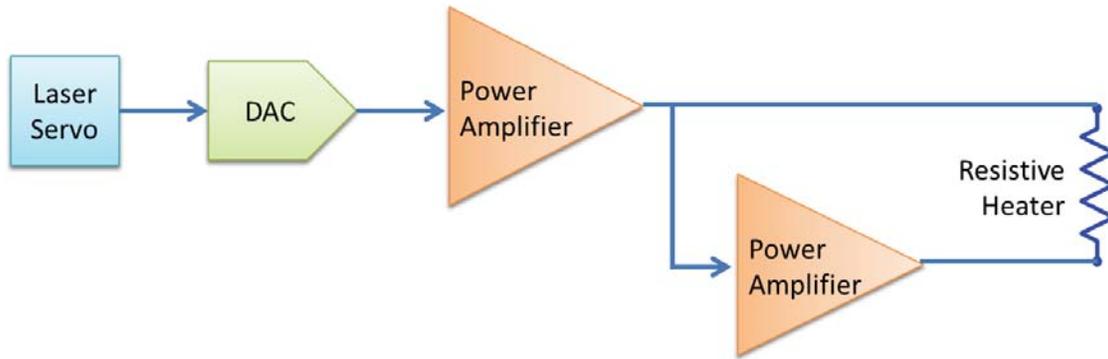


Figure 77. Laser heater driver block diagram.

There are two laser heater driver circuits for each physics package, one for the probe laser, one for the pump laser. The laser heater driver circuit has mainly a DAC and two power amplifiers with a $\pm 12\text{V}$ swing. The two laser heater driver outputs are 10kHz AC waveforms that are out of phase by 180 degrees, represented by the yellow and blue curves in Figure 78. The effective voltage swing across the heater resistor is shown by the magenta curve. The two laser heater driver outputs go side by side from the power board all the way to the heater resistor at the sensor, minimizing the magnetic field generated.

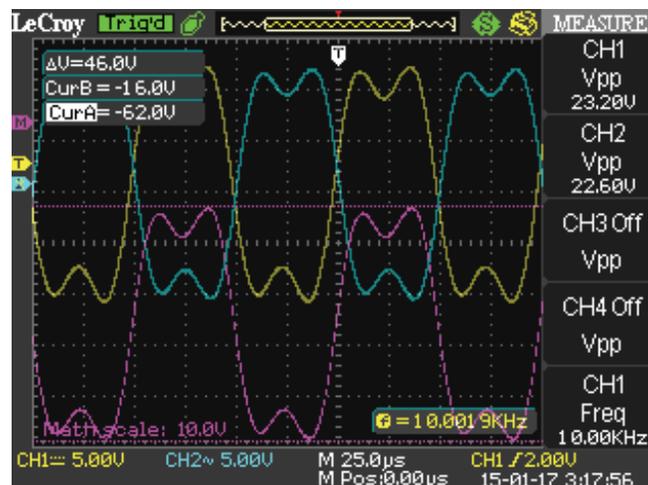


Figure 78. Laser heater waveform.

7.4 FIRMWARE PROGRAMMING

We chose a Xilinx™ Artix™ XC7A100T FPGA to implement the digital logic involved in the various control loops, signal generation and conditioning algorithms. These subsystems are controlled by the Microblaze™ soft embedded microcontroller that is also programmed in the FPGA. The digital logic in the FPGA is programmed with a low-level programming design abstraction called “Register Transfer Level” or RTL. The RTL can be coded in Verilog or VHDL language and is somewhat equivalent to programming in assembly language if the FPGA fabric primitives are used. As is seen from the magnetometer architecture of Figure 66, this is a very complex system with a number of interdependent and interacting subsystems.

Determining the interface protocols between these subsystems becomes a vital task in implementing this complex system. Further complicating matters is the fact that many subsystems have different optimal clock speeds. Hence, transferring signals across clock domain boundaries needs to be handled with

extreme care to prevent any glitches. Xilinx™ provides a number of modules for signal generation and conditioning to simplify this process. The Microblaze soft processor can be programmed in high-level C language. The embedded C code is independent of the FPGA fabric implementation and can be updated much quicker than the RTL code.

The RTL and the embedded code developments then proceed in parallel, with a well-defined interface between the two usually implemented in terms of a register map. Thus, the C code can control the RTL with standard data read/write transactions to memory-mapped addresses. Xilinx provides the IOBlock™ module to implement this functionality that is used in the MFAM system.

7.4.1 MODEL-BASED DESIGN

Model-based design (MBD) is used in FPGA design and verification. MBD is a mathematical and visual method of addressing problems associated with designing complex control, signal processing and communication systems. It is used in many motion control, industrial equipment, aerospace, and automotive applications. Tool chains in FPGA design are Simulink™ from Mathworks™ and System Generator™ and ISE™ from Xilinx™.

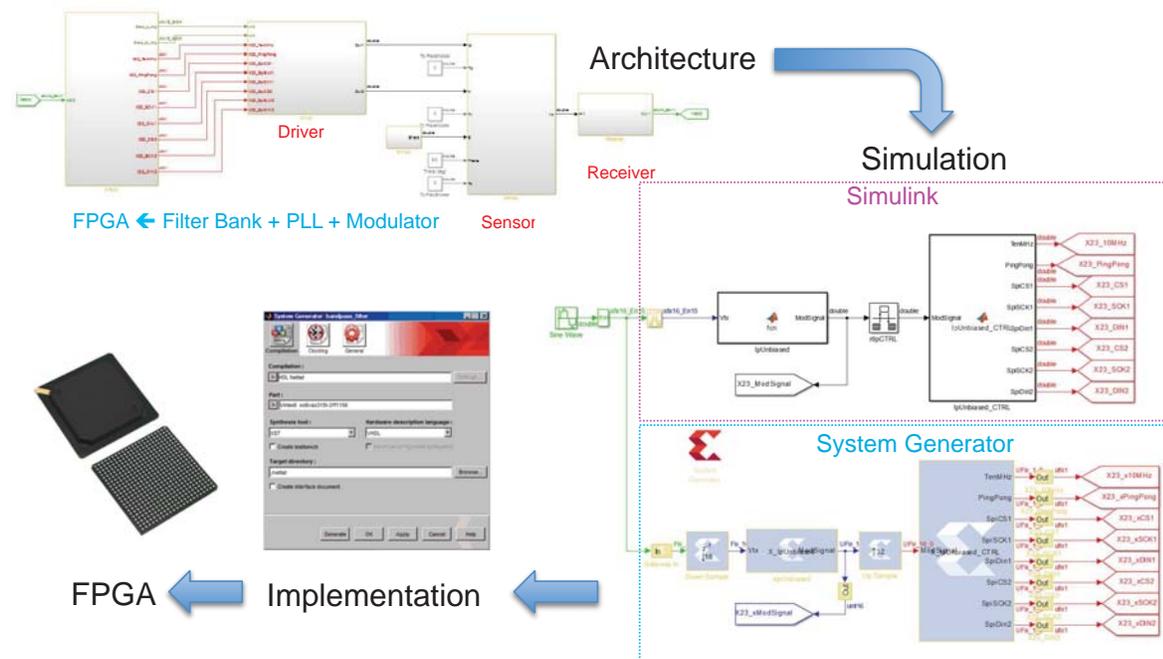


Figure 79. Model-based design workflow used in development of FPGA code.

Figure 79 shows the MBD workflow followed in the MFAM development. As discussed earlier, the MFAM modulation is divided into a digital part and an analog part. In the digital part, the synthesized Larmor output signal is sent to two fast DACs. The analog modulator uses this synthesized signal to modulate the RF carrier amplitude with analog filters and multiplexers, which are controlled by FPGA.

The analog part is functionally converted in the Simulink model delineated by the green box in the Figure 80. The FPGA design is modeled in System Generator and the Xilinx net list is automatically generated after functional and timing verifications through the computer simulation.

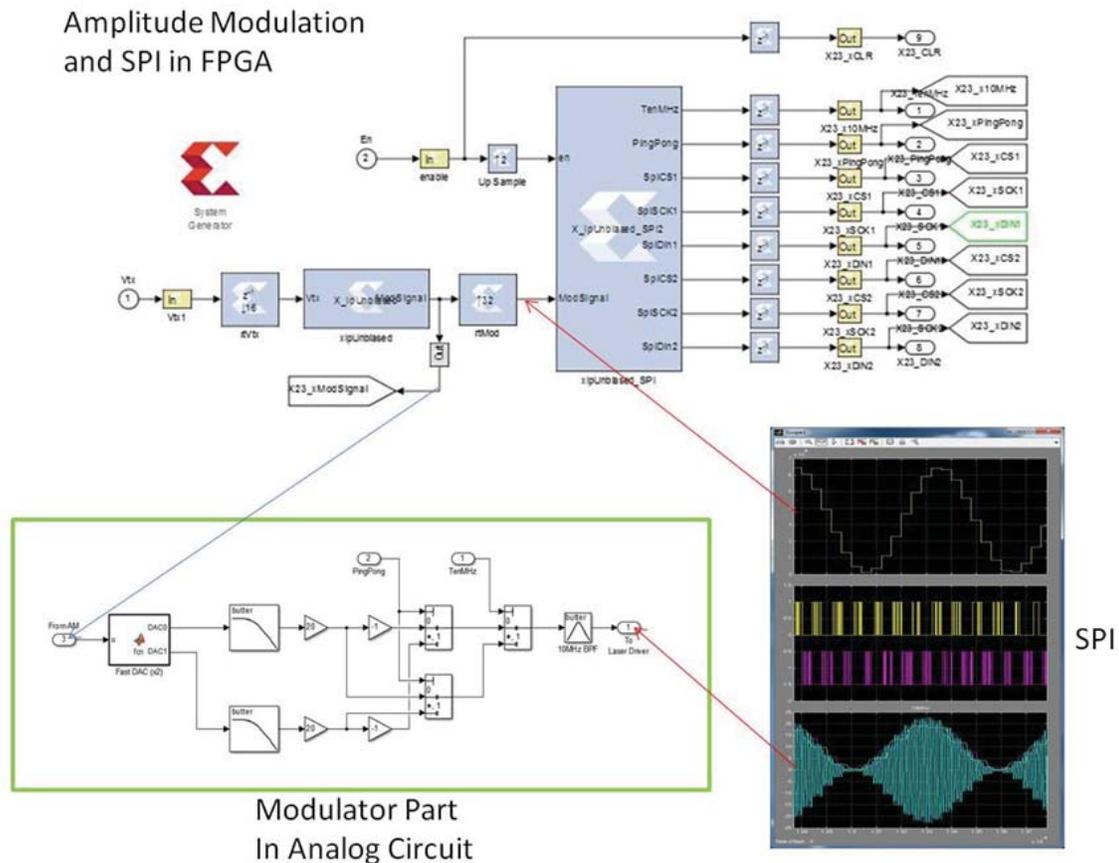


Figure 80. Example of Model-based design used to implement pump laser current amplitude modulation.

7.4.2 LARMOR LOOP

The sensor model, developed in Section 7.2, is introduced in the model-based design workflow to develop the digital processing required for closing the Larmor loop. The sensor, modeled as a resonator with a certain gain and phase response such that when the excitation frequency is equal to the resonant frequency, the gain is maximized and the response has zero phase shift relative to the excitation. Thus, the approximate center frequency of the magnetic resonance is determined by monitoring the amplitude of the response as the excitation frequency is swept through the entire Larmor band of 70kHz – 350kHz.

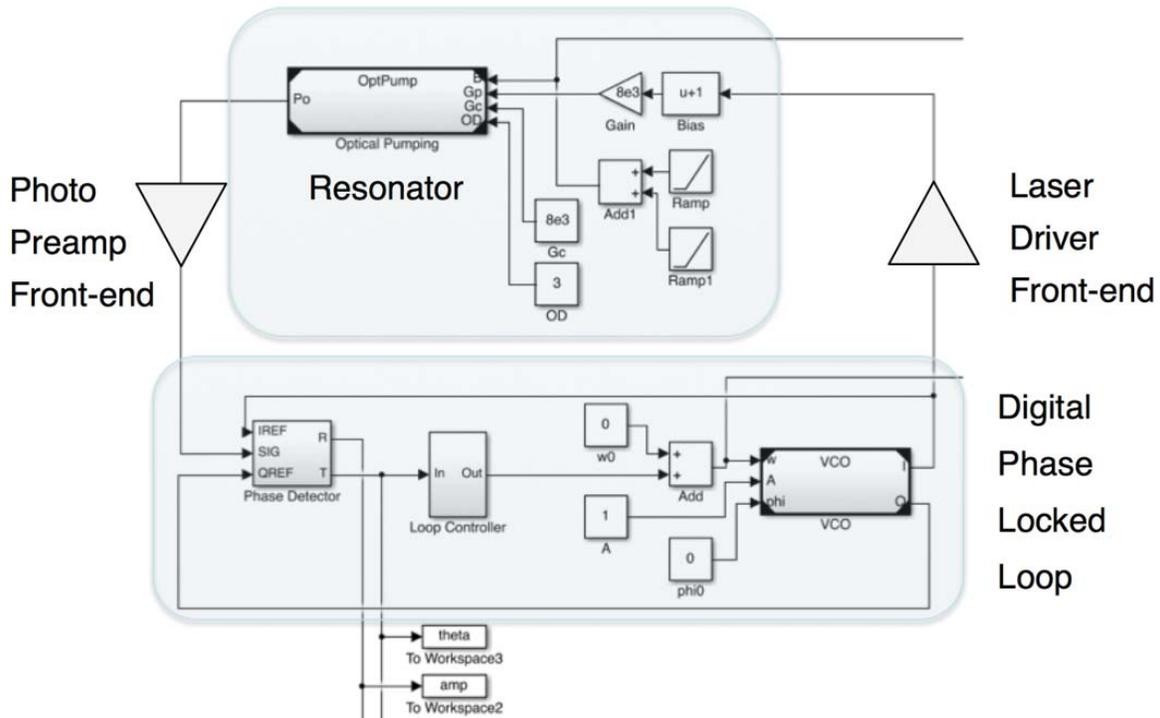


Figure 81. Design of the digital Larmor loop.

The resonator is then excited at a frequency close to the center frequency and the phase response is measured. This measurement is used to adjust the excitation frequency such that the resonator response is in phase with the excitation. The Larmor loop is said to be locked once this condition is met. This Phase-Locked Loop (PLL) methodology for operating a scalar atomic magnetometer is illustrated in Figure 81. The sensor model allows us to coarsely tune the PLL parameters, such as the control loop gains, for stability, bandwidth, noise and over/underflow avoidance within the simulation environment before even implementing the hardware.

The FPGA receives the noisy sinusoidal signal at the Larmor frequency, digitized by the high-speed ADC described in Section 7.3.2. Then, it may become the reference signal of the phase lock loop (PLL).

One possible Digital PLL (DPLL) concept is illustrated in Figure 82. A finite impulse response (FIR) Hilbert transform filter is used to generate an analytic signal from the input. A complex multiplier correlates the phase output from loop controller with input analytic signal and generates the relative phase error. The amplitude of the signal is adjusted using an automatic gain control (AGC) block that accepts an input corresponding to the signal level from the complex multiplier and adjusts the input gain to maintain a constant signal level. Two AGC blocks may be used to regulate the signal levels of the in-phase and quadrature components of the input analytic signal.

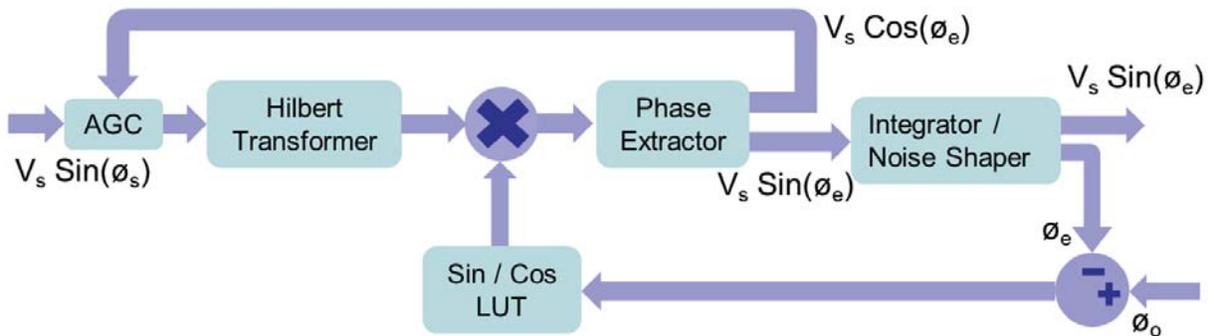


Figure 82. Digital PLL Design Concept.

A loop controller processes the phase error and adjusts the resonator excitation frequency to minimize it, this maintaining the loop in resonance. In the simplest scenario, the loop controller is simply an integrator. Typically though, a number of features such as higher-order terms, non-linear response, saturation handling and signal bracketing are used to improve the loop filter performance. As the magnetic field changes, the center frequency of the magnetic resonance changes and the loop filter follows by adjusting the pump amplitude modulation frequency to maintain the magnetic resonance condition. The DPLL in the MFAM operates at a clock frequency of 40MHz, which is more than a factor of 100 higher than the highest Larmor frequency. Figure 83 shows the digitized Larmor signal as it is detected at the input, reconstructed as it gets processed through the DPLL and then transmitted to the modulation section of the pump laser current driver.

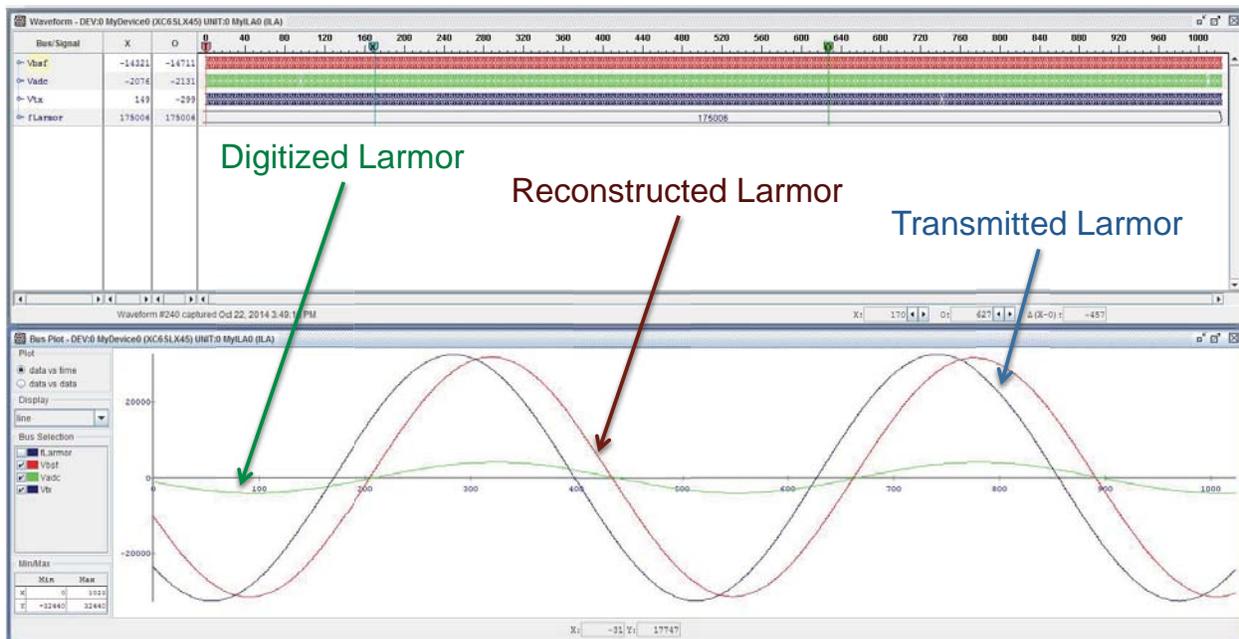


Figure 83. The hardware-in-the-loop system allowed us to verify performance.

There are a number of advantages in using a DPLL to close the Larmor loop as opposed to the traditional method of a self-oscillating magnetometer where an analog filter chain is used to provide the feedback.

- The design of the analog front end is significantly simplified as any phase error can be corrected for in the phase error calculation.

- A number of noise-shaping methods can be used to push the quantization noise out of the PLL bandwidth resulting in a high signal to noise ratio.
- Changing the bandwidth requires reengineering the loop controller, which can be implemented in firmware without any hardware change.
- The loop controller can be made adaptive to detect fast-slewing signals.
- A frequency counter is not required, as the excitation frequency is always exactly known.
- All the internal sensor signals are accessible and can be used to enhance the magnetic field measurements by detecting and compensating for the magnetometer orientation relative to the measured field.
- It is possible to extract vector component information using a scalar magnetometer.

7.4.3 MAGNETIC FIELD CONVERSION

The DPLL Larmor loop architecture does not require a frequency counter as the pump modulation frequency is synthesized, and hence always known. This frequency is sampled at the DPLL output at 40MHz. The resolution of the reported Larmor frequency can be significantly improved through multiple decimations. A 3-stage decimation scheme on the output-reporting path of the MFAM DSP is shown in Figure 84. At the first stage, the cascaded-integrator comb (CIC) decimator resamples the signal at 80Ksps. The first infinite impulse response (IIR) low pass filter (LPF) as the second stage can remove the noise beyond 400Hz (system bandwidth). The last LPF gets 1Ksps output and synchronizes output with a filtered 1 PPS.

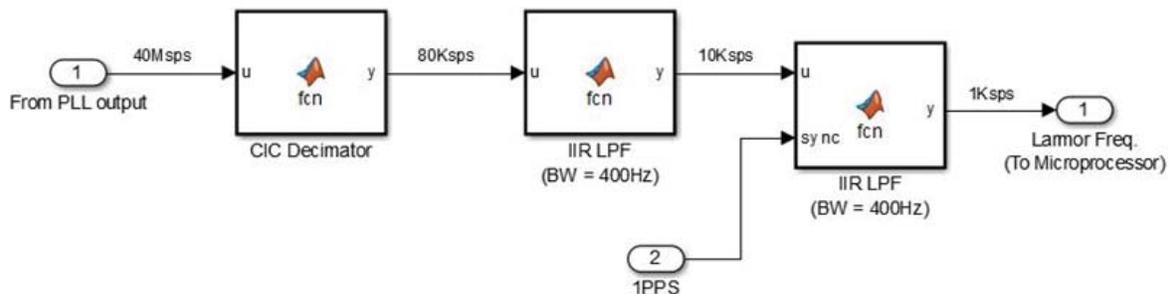


Figure 84. Output filter chain used to report magnetic field reading.

7.4.4 LASER LOCKING

The laser-locking servo stabilizes the laser frequency to the atomic absorption line. To tune and implement the servo we model the rest of the system and mathematically tune the servo with the obtained transfer function. As the servo is designed for a bandwidth of 500Hz, the two most critical parts are the laser current driver and the demodulator filter. The servo has to compensate for both. The sensor can be considered as a gain/mathematical function, although, being non-linear, will need a linear approximation around the absorption peak. The rest of the system can be modeled as simple gains if required, as they do not affect the operating frequency range.

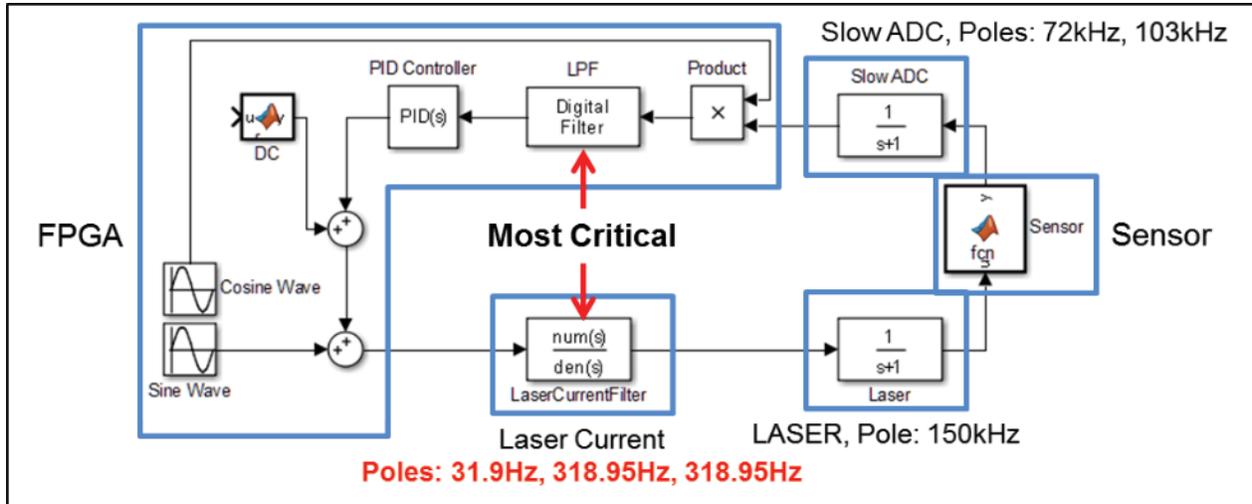


Figure 85. Laser locking servo closed-loop model.

The laser wavelength is dithered by modulating the laser current with a sinusoidal waveform at 3kHz or 4kHz frequency, represented by the red curve in Figure 86. To understand how the locking signal is generated, it is best to open the servo loop that holds the laser on the absorption and sweep the laser past the ideal value. As can be seen in Figure 86, there are two important components in the output of the ADC. The signal at the second harmonic of the applied sine wave indicates that the absorption line is detected. The signal at the fundamental is what is used for the locking method.

The output from the photodiode is an amplitude-modulated signal with the carrier being the 3kHz or 4 kHz dither. The amplitude of this signal can be said to go from a large positive value through zero and then to a large negative value, i.e. the signal appears inverted. For small offsets between the laser wavelength and the ideal, the variation in the amplitude is nearly linear with respect to the wavelength. This is the normal case when the servo is engaged.

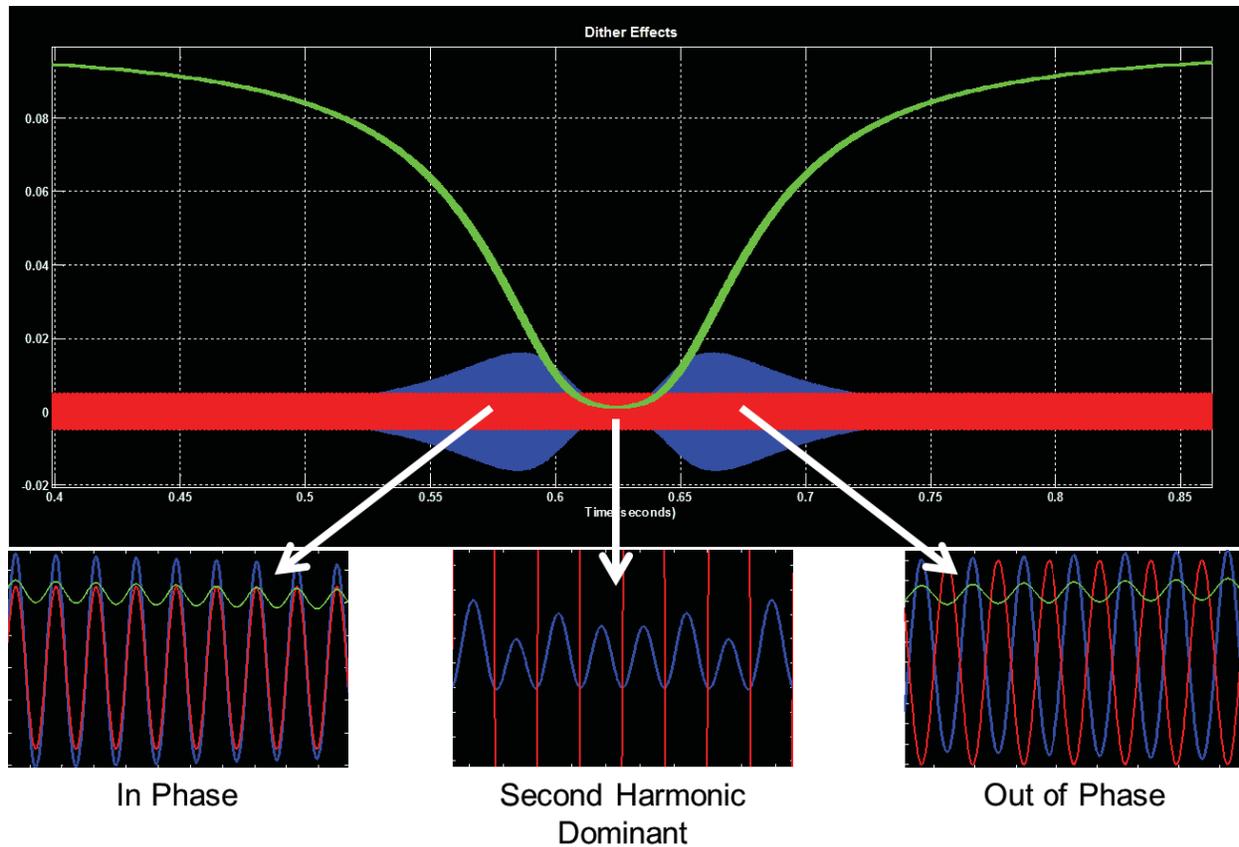


Figure 86. Laser locking concept. As the laser frequency is swept across the optical absorption curve, the light output follows the green curve shown above. If the laser frequency is modulated by a sinusoidal waveform, represented by the red curve, then the AC component of the green curve has characteristics shown by the blue curve and described in the text.

Demodulation can be implemented as the product of the photodiode signal with a phase coherent carrier followed by a low pass filter. Because this signal is demodulated within the FPGA, the noise level of amplifiers and the ADC in the several kHz range determine the performance of this section. The demodulator output, as illustrated by the red curve in Figure 87, follows the slope of the absorption curve. The laser frequency is stabilized to the zero crossing point of this curve, which corresponds to the zero-slope point or the peak of the optical absorption.

The decrease in the slope of the amplitude vs wavelength function at large offsets only matters during the system startup. With an unconditionally stable servo design, its only impact on the design is that it causes the servo to have a slightly longer settling time. We have explored using a conditionally stable servo and a compensating non-linear function as a means of improving the performance of the laser locking system. This appears to be unneeded in the current system but may be needed when the design is further miniaturized.

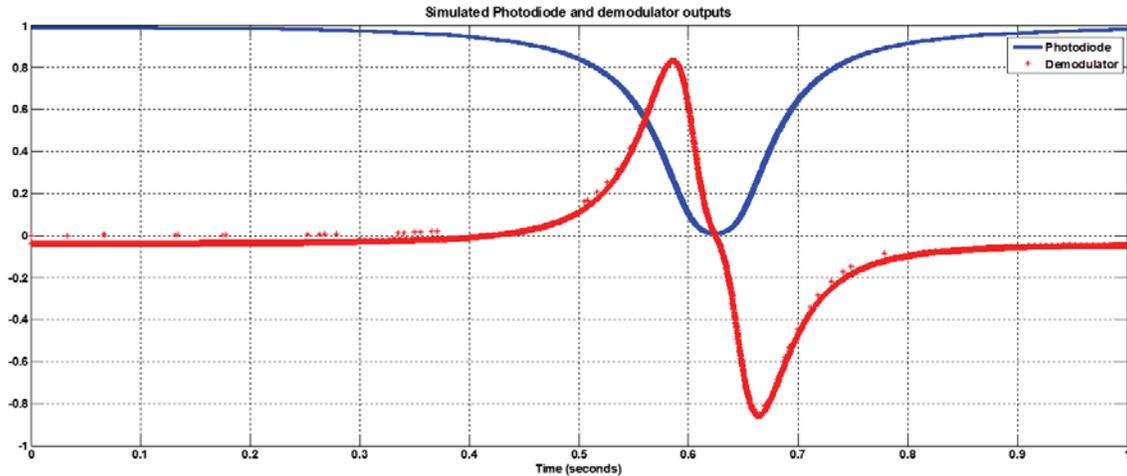


Figure 87. Demodulator output (red curve) used to lock the laser, derived from the blue curve shown in Figure 86.

The laser current controller uses the demodulator output as the error signal and adjusts the laser current around the nominal operating DC value, as described in Section 7.3.3. The laser heater controller minimizes the laser current actuation to maintain constant optical output power, as described in Section 7.3.4. These nested servos are tuned so as not to overlap in bandwidth and can be easily optimized further to improve the noise performance of the system. The laser stabilization has proven to be very robust in rejecting disturbances due to temperature change, which are usually the main cause of instrument failure.

7.4.5 MFAM STARTUP

The startup sequence of the MFAM takes place in five phases outlined in Figure 88. When started from room temperature, the entire procedure requires 3 minutes. Within about 5 minutes, the magnetometer reading is stable.

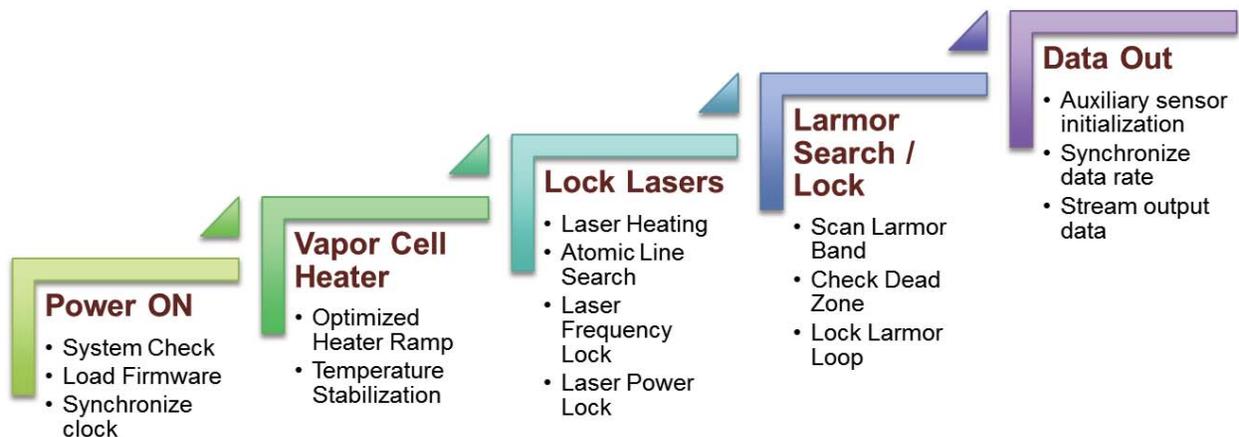


Figure 88. MFAM power-on start-up sequence.

7.4.5.1 POWER ON

On power up, the FPGA loads the configuration bit stream from the onboard SPI Flash. The internal PLL implemented in the RTL logic synchronizes the internal clock generator to the supplied crystal oscillator clock signal. The Microblaze processor boots up with the program already present in its RAM at startup.

The program initializes and self-tests all the processor peripherals. The general-purpose inputs and outputs (GPIO), timers, built-in ADC, communication interfaces and interrupt subsystem are initialized and self-tested. The program then initializes all the parameters required by the RTL logic. Next, the onboard inertial sensors are initialized and self-tested. Finally, the program waits for a predefined amount of time for all the hardware filters onboard to settle.

7.4.5.2 VAPOR CELL HEATER STARTUP

The vapor-cell heaters are engaged at maximum power to speed up the cell heating process. The cell temperature is continuously monitored until the set temperature has been reached. When the temperature read from the vapor cell reaches the set temperature, the cell temperature servo is initialized such that the temperature settles to the set point quickly, thus enabling the servo. Once enabled, we wait for a predefined amount of time to allow it to settle.

7.4.5.3 LOCK LASERS

Each of the pump and probe lasers follows the same procedure described here. The laser heater is slowly ramped up while scanning through a large range of the laser current and effectively laser wavelength. As the heaters are ramping, the program monitors for the optical absorption lines, as shown in the left panel in Figure 89, by checking if the laser locking signal has crossed a set threshold. The laser-locking signal exceeds the threshold only if the absorption lines are in range. The process from here on can be seen in the right panel of Figure 89 for both the pump and probe lasers and takes a few hundred milliseconds.

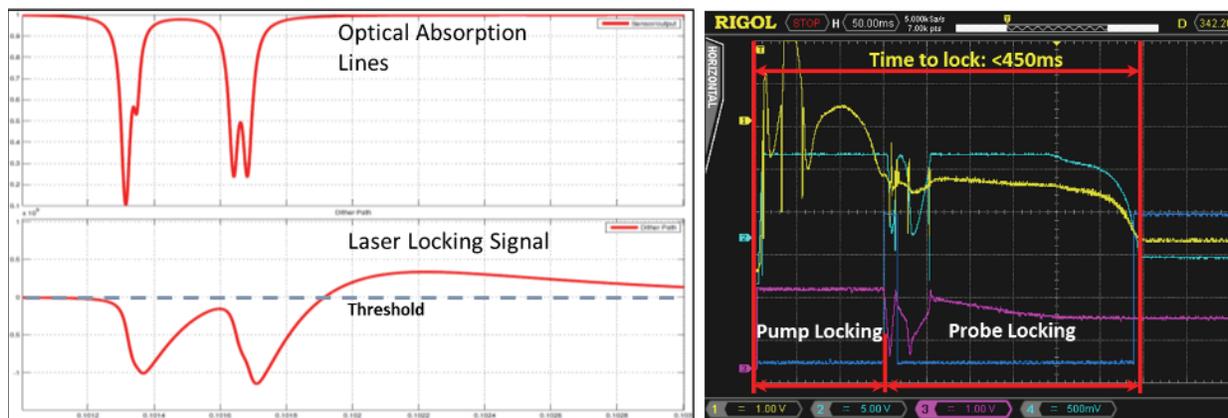


Figure 89. Start up operation of the final system.

Once the absorption lines within range, the laser heater ramp is stopped and the laser current range, in which the absorption lines are seen, is narrowed down. One can see this as the rapid, large transitions at the start of the yellow trace and the blue trace in Figure 89 when the pump locking is complete. After the narrower range of laser current is obtained, the program traces the absorption line until the peak is reached, at which point the laser-locking servo is enabled. This is seen on the yellow trace towards the end of the pump laser locking and the blue trace as the largely flat, last part of the probe laser locking in Figure 89. Once the wavelength control loop has settled, laser power control loop is enabled, which manipulates the laser heater to control the laser power. Once both control loops are running, the laser current and consequently the laser power are moved to the optimum operating point.

7.4.5.4 LARMOR SEARCH / LOCK

After both lasers are locked and moved to their corresponding optimum operating points, the Larmor loop RTL module is enabled. The Larmor loop performs a wide range search of the Larmor frequency and selects a frequency band based on the observed signal power as described in Section 7.4.2. If the signal power is observed to be smaller than a threshold value, it is assumed that the magnetometer is in a dead zone and the wide range search continues until the signal power exceeds the threshold. The frequency band estimation is used to initialize the Larmor loop DPLL, which is then enabled.

7.4.5.5 DATA OUT

An internal timer set to 1kHz frequency is used to time the data decimation and transmission described in Section 7.4.3. This timer is synchronized to an externally supplied 1 pulse per second (PPS) signal. The number of times the 1 kHz timer fires in between two PPS signals is used as the feedback to a servo adjusting the frequency of the timer. The synchronized 1 kHz is also used to regulate the output data rate of the Larmor loop. Using this timer, the magnetic field reading is collected from the Larmor loop every millisecond. Additional system status information is collected at each millisecond as well. At certain multiples of the 1-millisecond pulse, the auxiliary inertial measurement sensor data and other diagnostic data are collected. All the collected data are collated and transmitted over the external SPI interface.

7.5 RESULTS

The MFAM module, comprising the two physics packages and the sensor driver electronics, was tested to characterize the performance in terms of noise and temperature range. At this point, we used the MFAM to conduct a number of field tests that required the instrument to perform to its limits in terms of sensitivity and common mode noise cancellation. The results of the characterization and the tests are described below.

7.5.1 SENSITIVITY

The intrinsic noise of the sensor within its bandwidth sets the sensitivity of the magnetometer to externally applied magnetic fields. Great care has to be taken to measure the intrinsic noise of the sensor itself so as not to contaminate the measurement with ambient noise. The noise performance of the MFAM using the driver electronics described previously is shown in Figure 90.

The noise amplitude spectral density is measured in a zero gauss chamber (ZGC) with a setup similar to the one described in Section 7.2.6. The two MFAM sensors are placed vertically adjacent such that the sensing element separation is about 1 inch. A pair of coils in the Helmholtz configuration is used to generate a uniform magnetic field of magnitude 50,000 nT inside the ZGC at the location of the sensors. The coils are driven by an ultra-low-noise current source.

The red and green curves in Figure 90 show the amplitude spectral density of the individual sensors and the green curve represents the same for the difference of the two sensors, or the gradiometer. The noise of the two sensors is uncorrelated. Hence the gradiometer noise, which is the sum (in quadrature) of the noise of the two sensors, is higher than that of the individual sensors. This noise is attributed to the intrinsic sensitivity of the sensors. We can see that both the sensors have sensitivities well below $1\text{pT}/\sqrt{\text{Hz}}$!

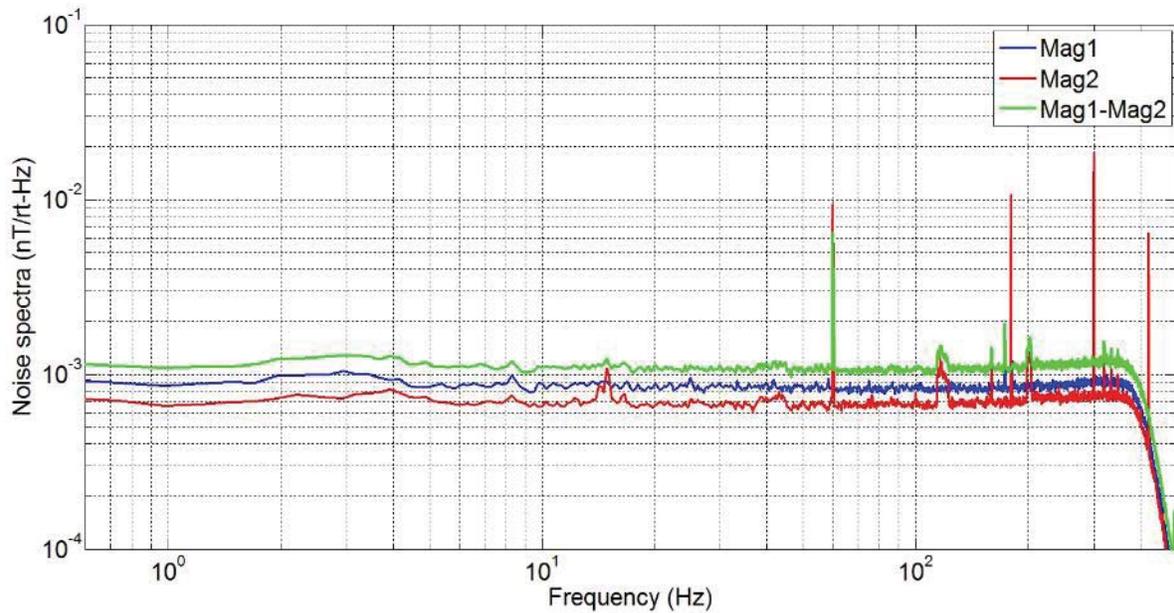


Figure 90. Noise performance of the MFAM module as measured in a zero-gauss shield can. Note that the two sensors, shown by blue and red curves, have sensitivities below $1\text{pT}/\sqrt{\text{Hz}}$. The green curve is the amplitude noise spectral density of the gradient measurement between the two sensors. Here, the two sensors were placed in contact, with approximately 1-inch separation between sensing cells.

The drop in the noise at frequencies above 350Hz is due to the bandwidth limit imposed on the sensor at the output decimation and anti-alias filtering stage. Thus, the 400Hz bandwidth of the sensors can be confirmed from the noise measurements as well.

7.5.2 STEP RESPONSE

The step response of the MFAM is measured in the same setup as described above with the addition of an auxiliary 3-turn coil to impose a square-wave magnetic field in addition to the background magnetic field. A function generator drives the auxiliary coil with a 10Hz, 10Vp-p square wave through a series 1kohm resistor. The recorded MFAM readout is shown in Figure 91, along with a zoomed-in view of the rising and the falling edges of the waveform.

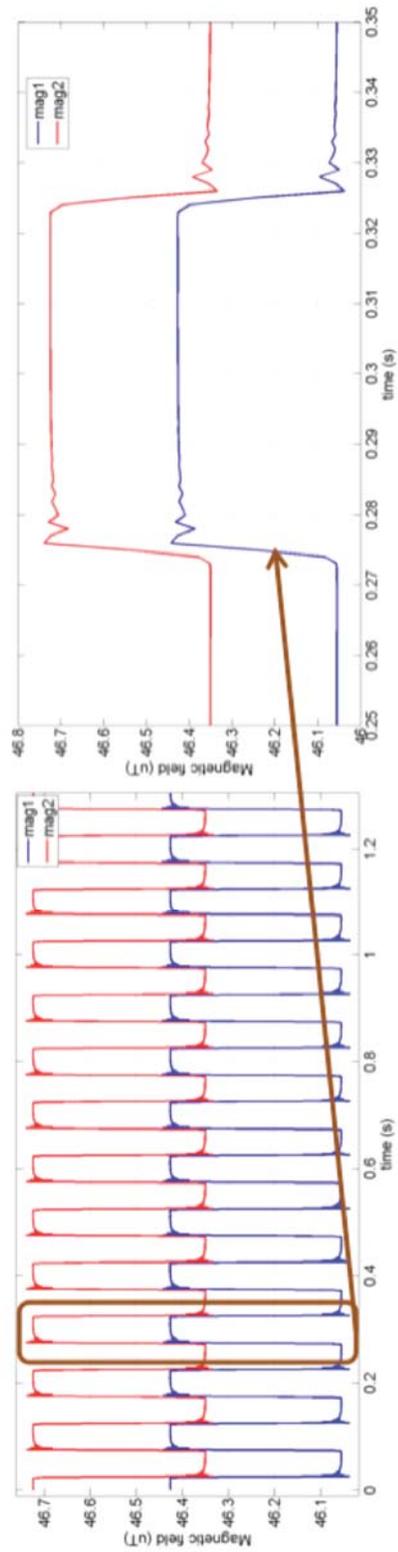


Figure 91. Response of the MFAM to a square wave magnetic field.

An important observation is that the edges are fairly well preserved. The MFAM can follow a magnetic field slew of 400nT in a millisecond. We can observe ringing and decay for a few milliseconds after the edges. We attribute this to the coil inductance and the bandwidth of the MFAM being limited by a linear phase filter, though further investigation is needed to confirm this hypothesis.

7.5.3 OPERATING TEMPERATURE

To verify the operation of the MFAM over the entire operating temperature range of -30°C to +50°C, we exposed an MFAM module to a temperature swing across this temperature range in an environmental chamber. Figure 92 shows that the MFAM module was operational throughout this temperature range. The magnetic environment inside the chamber is extremely noisy due to the steel construction and the close proximity to compressor and circulation fans. The MFAM measured this noise, as can be seen from the peak-to-peak variation of the magnetometer reading as a function of time.

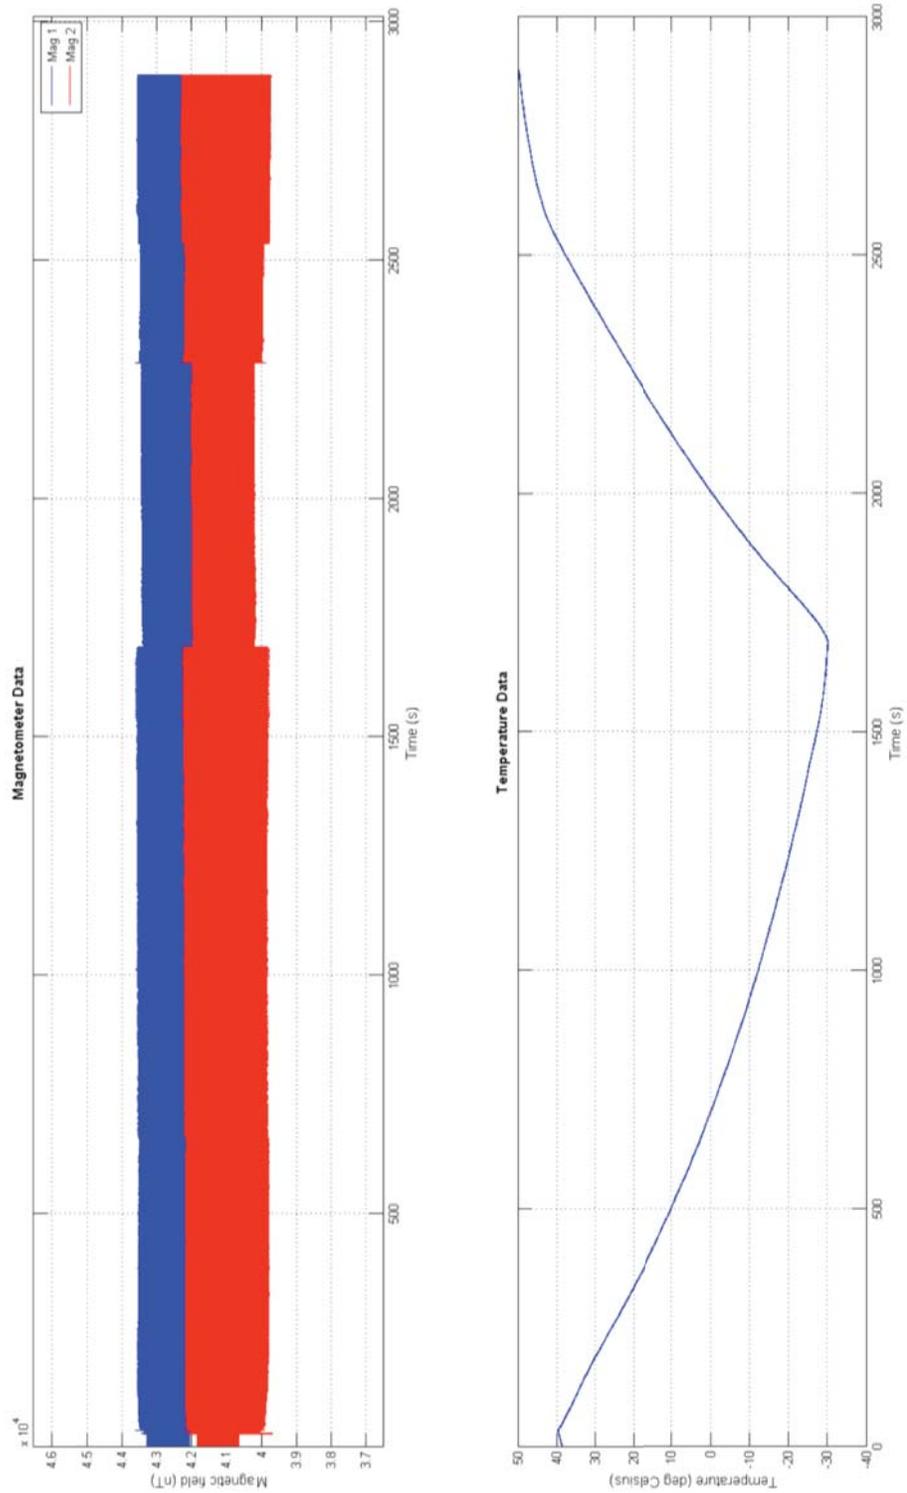


Figure 92. The bottom graph shows the operating temperature of the MFAM module. The temperature varies from -30 C to +50 C over 50 minutes. The top graph shows the readings of the two sensors from the MFAM module.

7.5.4 GRADIOMETER COMMON MODE REJECTION

Traditionally, measurement of the noise floor of a scalar atomic magnetometer has required elaborate setups including ZGCs as described in Section 0. This is because those sensors, when placed adjacent to each other, suffer from cross talk. The closest they may be placed is about 1 foot apart. The magnetic field gradient noise over that distance overwhelms the sensor noise, particularly in commercial settings.

The MFAM sensors, on the other hand, do not suffer from any direct cross talk. They may be placed touching each other, with the sensing volumes less than 1 inch apart. The noise in the gradient over such a small baseline is very small even in industrial settings. Thus, we can verify the noise performance of the MFAM even outside a ZGC. The right panel of Figure 93 shows the two MFAM sensors placed next to each other in a commercial laboratory. Both the sensors measure the same cultural noise, as seen from the overlapping red and blue noise spectral density curves in left panel of Figure 93. However, the spectral density of the gradiometer signal, shown by the green curve in Figure 93, is limited only by the sensor noise across the frequency spectrum. For instance, the common mode noise is suppressed by a factor of 10 at 10Hz and the suppression factor increases with decreasing frequency, reaching 50 at 0.5Hz. Only at a frequency higher than 150Hz is the cultural noise lower than the magnetometer noise.

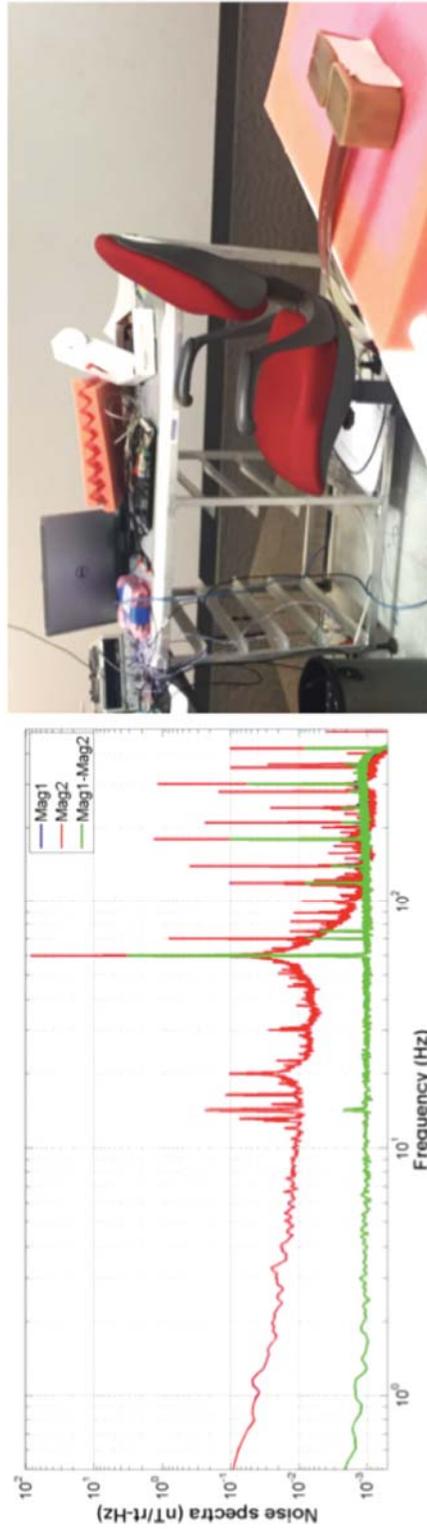


Figure 93. The common mode cancellation of the sensors is demonstrated by the suppression of common mode cultural noise measured in a typical commercial environment by more than an order of magnitude. The sensor performance can be verified without shielding the cultural noise in a ZGC.

This feature can be used to discriminate between magnetic objects by their distance from the sensors. For instance, in an experiment illustrated in Figure 94, the two MFAM sensors are placed about one foot apart. We waved screwdriver approximately six feet away from the gradiometer, and slightly closer to one sensor than the other. During the test, a car drove by about 30 feet away from the setup. The magnetic signature of the car can be prominently seen in the magnetometer data shown in the right panel of Figure 94. This signature, however, is highly suppressed in the gradiometer data, shown in the left panel of Figure 94. There, the motion of the screwdriver becomes very obvious. Such a setup can be very useful in a physical security system where the setup needs to be immune to cultural noise while being sensitive to local magnetic disturbances.



Figure 94. Measurement of local gradients in presence of large disturbance far away. Here, the sensor separation was 10 inches.

7.5.5 EDDY CURRENT MEASUREMENT

The MFAM sensors also provide exceptional spatial and temporal resolution. These sensors output field readings at 1000 Hz, leading to some fascinating observations. For instance, waveforms from power lines can be fully observed and filtered, if desired, not just aliased into a lower frequency. Such aliasing typically produces strange artifacts and noise. We can even measure eddy currents in non-magnetic materials using the ever-present 60 Hz noise as a transmitter.

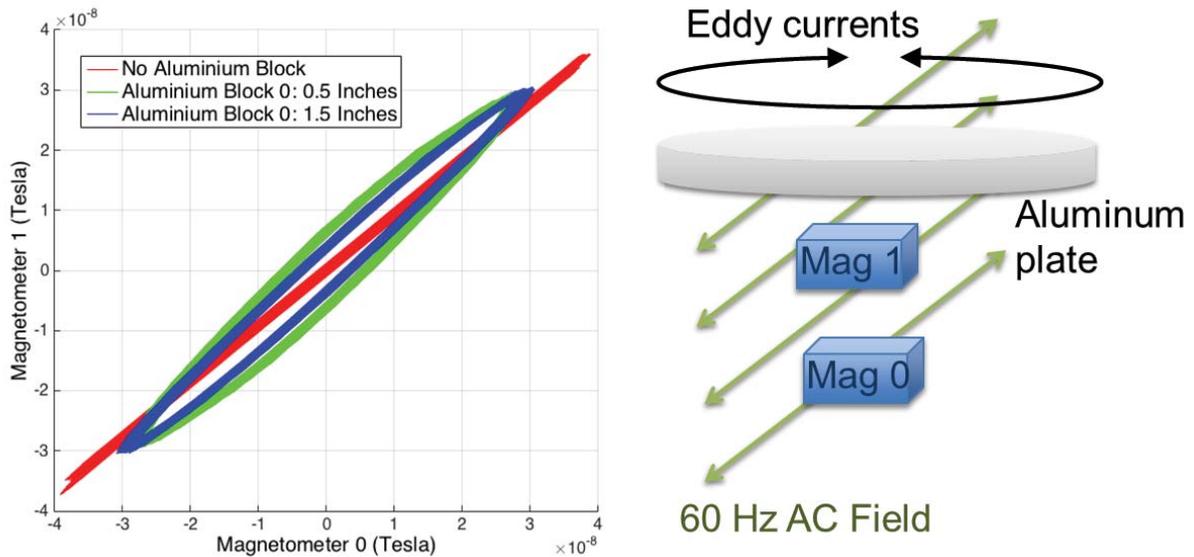


Figure 95. Effect on the phase of 60Hz magnetic field component due to eddy currents induced in an aluminum plate by ubiquitous AC power line fields.

Changing magnetic fields induce eddy currents in conductive objects. These induced currents produce magnetic fields that are out-of-phase with the field causing them. By making measurements of the out-of-phase component of the magnetic field, we can measure the eddy current. This allows for the sensing of even non-ferrous conductors, something not typically associated with magnetometers.

In Figure 95, the graph on the left shows the output of one sensor plotted on the horizontal axis, and the other sensor along the vertical axis. This produces Lissajous patterns according to the relative frequency and phase of the two readings. In this case the frequencies are the same, so we get a circular pattern if the two signals are perfectly out of phase, and a line if they are in phase. Elliptical patterns are the result of something in between. We can clearly see the effect as the phase difference between the two signals (caused by the eddy currents) grows larger.

8 CONCLUSION

We are very excited about the results of this project. During the project timeline, we analyzed, designed and built a prototype real-time array, and discovered exciting properties of its operation. We have also miniaturized the electronics' design, doing justice to the small size and low power consumption of the physics packages developed in previous projects.

We have shown that utilizing an array to acquire enough data to run inversion algorithms in real-time yields significant and exciting benefits. Such an instrument will be extremely valuable in operation, not only to confirm re-acquisition of the location of previously analyzed anomalies, but also to eliminate the two-step process of data acquisition and analysis. Situations that are difficult to interpret with a fixed set of data become much more manageable when analyzed in real time. This will result in a considerable savings in remediating UXO, especially in the underwater environment.

This work has significant benefits even beyond those envisioned in our initial objectives. The entire process of planning and executing surveys may be changed with the availability of small sensors. Arrays may be deployed on fixed platforms, as envisioned in this work, or from individual autonomous systems forming a swarm. This, in effect, creates an adaptive array, with a broad geometry for searching for anomalies, then converging to create a denser array for analysis. Even broader applications are possible, as described in the following section.

8.1 OUTLOOK: MULTI-CHANNEL MFAM SYSTEM

Historically, high-precision scalar magnetometry has been the realm of highly specialized scientific research and optimized, single-purpose geophysical surveys. The development of sensor driver electronics for the next-generation laser-based atomic magnetometers has enabled unprecedented performance for the size, weight and power of the system. The miniaturization of the magnetometer allows the MFAM to be used as a component in a multi-sensing application where magnetometry, unconstrained by payload limitations, can be easily integrated on to a highly configuration and versatile platform. Finally, high-precision scalar magnetometry can move from the realm of niche scientific instrumentation to ubiquitous industrial and commercial sensing with far-reaching applications ranging from non-destructive evaluation and structural health monitoring to physical security and even medical sensing. The MFAM architecture, with its high sample rate and unique specifications, is particularly amenable to multi-channel magnetometry, as illustrated in Figure 96.

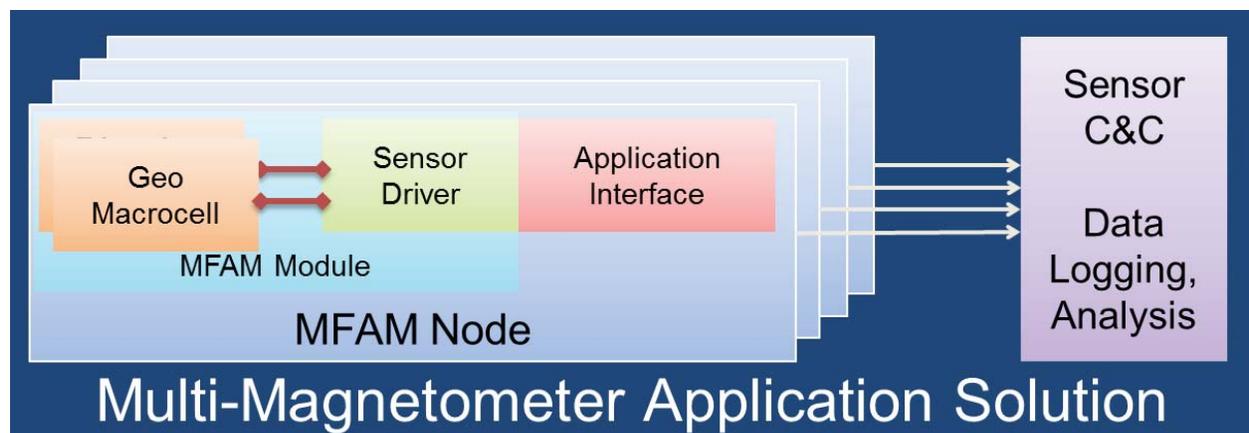


Figure 96. Multi-channel magnetic field acquisition system based on the MFAM technology.

In such an application, a number of MFAM modules are distributed over a large area and synchronized to each other within a microsecond. The data is collected from all locations simultaneously and processed to automatically detect a threat or to pinpoint a magnetic anomaly. The sensors can be deployed on a swarm of autonomous vehicles that gather to characterize the anomaly and spread out dynamically to cover large area. The system can integrate acoustic sensors to enable multi-modal passive threat detection or non-invasive imaging.

As the applications open up new markets, the technology can be easily upgraded by packaging the sensors using micro-electro-mechanical systems (MEMS) fabrication processes. At the same time, electronics drivers can be packaged into multi-chip modules, an example of which is shown in Figure 97.

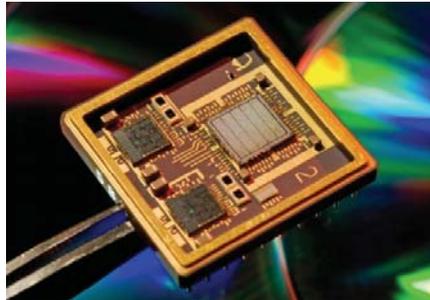


Figure 97. Multi-chip module.

Multi-chip modules (MCM) combine several chips with different functions in one module that is housed in a single package. The chips in a multi-chip module are adjacently placed or stacked together to form a 3D-IC. The chips are interconnected via a multi-chip substrate, which can be interconnected using associated through-silicon vias (TSV) to fabricate a stacked structure that is characterized by a small footprint. Before the multi-chip modules can be formed into a package, they may be encapsulated with a potting compound.

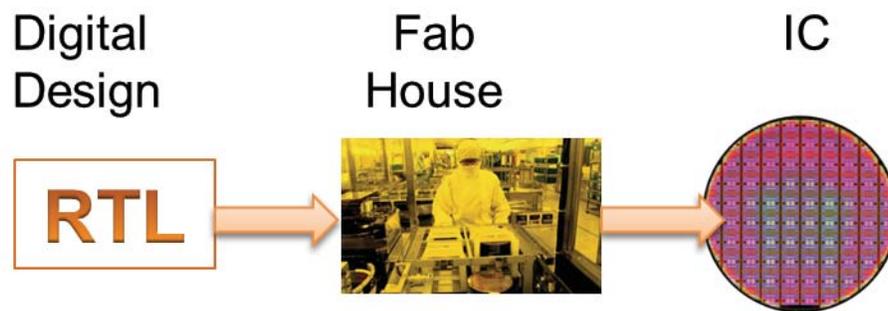


Figure 98. Straightforward Application Specific Integrated Circuit development from the RTL (Register Transfer Level) Design of the digital logic.

As the technology adoption matures, the electronics can be easily ported to application specific integrated circuits (ASICs) made possible by the straightforward partitioning of the design that was developed as a part of this project. A number of processor cores, such as the ARM[®] processor, are available that can be readily integrated into the logic fabric in an ASIC to provide field upgradability of the functioning of the MFAM. This not only has cost advantages once the production volume exceeds a few tens of thousands of units, but also lowers the power consumption and provides better signal integrity leading to better noise performance.

9 LIST OF PUBLICATIONS

Kitching, J., Mhaskar, R., “MEMS Revolution in Atomic Sensing – From Tiny Clocks to Ultra-sensitive Magnetometers,” IEEE SF Bay Area MEMS & Sensors Chapter meeting, Santa Clara, CA, Dec. 2015.

Prouty, M., “Applications of Microfabricated Atomic Magnetometers,” SAGEEP, Austin, TX, 2015.

Prouty, M., “Low Power Miniature Atomic Magnetometer,” MARELEC, Philadelphia, PA, 2015.

Prouty, M., Tchernychev, M., “Real-Time Threat Detection Using Magnetometer Arrays,” SPIE Defense and Security Conference, Baltimore, MD, April, 2016.

Zhang, R., “All Optical Pump Probe Magnetometer,” DAMOP, Columbus, OH, June 2015.

10 APPENDIX

10.1 MFAM INTERFACE DETAILS

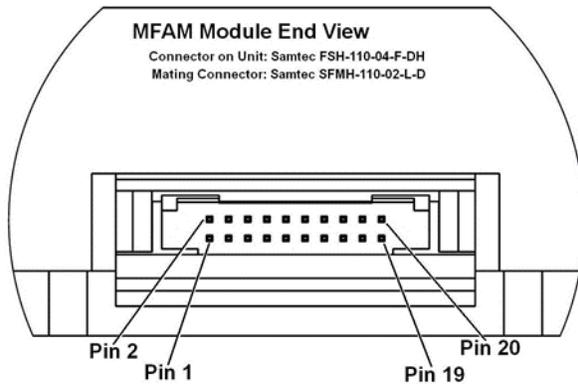


Figure 99. MFAM Module connector details

Pin	Signal	Description
1	CHASSIS	Chassis ground
2	GND	Power supply ground
3	Vin	Power supply (9.5V to 16V)
4	Vin	Power supply (9.5V to 16V)
5	GND	Power supply ground
6	GND	Power supply ground
7	MSPI_DOUT	Data output (TX), SPI protocol, unit is master
8	DNC	Do not connect, leave pin open.
9	MSPI_DIN	Data input (RX), SPI protocol unit is master
10	DNC	Do not connect, leave pin open.
11	MSPI_SCLK	Clock output, SPI protocol, unit is master
12	DNC	Do not connect, leave pin open.
13	MSPI_CSB	Chip select, active low signal, SPI protocol, unit is master
14	DNC	Do not connect, leave pin open.
15	CTS	Clear To Send handshake signal, INPUT to MFAM unit (not implemented)
16	REF10M	10 MHz reference clock input, OPTIONAL

17	RTS	Ready To Send handshake signal, OUTPUT from MFAM unit (not implemented)
18	1PPS	1 pulse per second input, positive edge triggered
19	GND	Power supply ground
20	GND	Power supply ground

10.1.1 OUTPUT DATA FRAME FORMAT

The MFAM module outputs the magnetometer data via an SPI bus at an output rate of 1000 samples per second. Each sample consists of 28 bytes of data as diagramed in Figure 100. Each sample consists of fourteen data blocks, which are further, described below.

For all 16 and 32 bit words not that the format lowest significant byte first, followed by the next significant byte, and on up to the most significant byte (for 32 bit words).

MFAM SPI Bus Data Transfer Format

Revision X3
03/07/16

One Data Sample (28 Bytes)



Fid Data Format

Aux Channel ID (bits 15:11): M2V, M1V, A2, A1, A0	FID Count (bits 10:0)
---	-----------------------

```

A2 A1 A0
---
0 0 0 Reserved
0 0 1 Compass
0 1 0 Gyro
0 1 1 Reserved
1 0 0 Accelerometer
1 0 1 Reserved
1 1 0 Reserved
1 1 1 Reserved
M2V: Mag 2 Data is Valid
M1V: Mag 1 Data is Valid
    
```

System Status Data Format

PPS Status (bits 15:14)	Diagnostic Stages 1 through 11 (bits 13-10)	Diagnostic Substate (bits 9-7)	Reserved for Future Use (bits 6-1)	Start Up Failed (bit 0)
-------------------------	---	--------------------------------	------------------------------------	-------------------------

Bit 15: IPPS Pulse is Being Received
 Bit 14: 1KHz Sample Rate is Locked to IPPS
 Bits 13-10: Start Up Diagnostic Stages 1-11 (internal use)
 Bits 9-7: Start Up Diagnostic Sub Stages (internal use)
 Bit 6-1: Reserved for Future Use
 Bit 0: Startup has Failed

All 32 Bit Words

Low Byte Address Data LSB	Low Address + 1 Data MSB - 2	Low Address + 2 Data MSB - 1	High Byte Address Data MSB
---------------------------	------------------------------	------------------------------	----------------------------

Example:

Byte Sequence: 0x16, 0x85, 0x86, 0x2F
 Equals: 0x2F868516
 Equals: 797345046 decimal

To Convert 32 bit value to magnetometer reading:

One LSB = 50 femtoTeslas

Therefore to convert 32 bit value to nanoTeslas:
 reading (in nT) = [value] * 50E-6

Using Example above:

$$797345046 * 50E-6 = 39867.2523 \text{ nT}$$

All 16 bit Words

Low Byte Address Data LSB	High Byte Address Data MSB
---------------------------	----------------------------

Example:

Byte Sequence: 0x05, 0xC9
 Equals: 0xC905

Figure 100. SPI Frame Data Format**10.1.1.1 FID DATA FORMAT (16 BITS):**

The Fid data block is subdivided into three sections.

1. The lowest 11 bits [10..0] are a binary number that increments by one each sample. At 2047 (0x7FF) the count rolls over to zero (0x00).
2. The next three bits [13..11] are Auxiliary Channel Identification bits. The auxiliary sensors update at a slower rate than 1000 samples per second. These auxiliary sensors get multiplexed into shared Auxiliary channels (described in more detail below). These three bits identify which auxiliary sensor data occupies the Auxiliary channels. Currently there are three auxiliary sensor channels being output: Compass, Gyro, and Accelerometer. Each of these sensors has 3 axis', plus a temperature channel used for sensor calibration. Thus there are four Auxiliary channels: X, Y, Z, and Temp.
3. The upper two bits [15..14] are flags used to indicate whether the magnetometers are outputting valid data. Bit 15 corresponds to Magnetometer 2, and bit 14 corresponds to Magnetometer 1. A one indicates valid data is being sent.

10.1.1.2 SYSTEM STATUS (16 BITS):

The System Status word is subdivided into four sections:

1. The lowest bit [0] indicates that the startup process has failed. A zero in this bit position indicates startup success, a one in this position indicates startup failure.
2. The next six bits [6..1] are not currently defined. They are reserved for future use.
3. The next seven bits [13..7] contain startup diagnostic information. When the MFAM module is first powered on, the unit sends out a different data set. In place of the magnetometer and auxiliary channels is internal diagnostic information documenting the start-up progress. There are 11 main stages during the startup process, and several substages. The Diagnostic Stage is documented System Status bits [13..10] (stages 1-11). The sub stage is documented in bits [9-7]. Currently these variables are not finalized, and reserved for internal use. After a successful startup the System Status Diagnostic Stage bits (13-10) are cleared indicating that the MFAM is outputting valid magnetometer data.
4. The upper two bits [15..14] document the 1 PPS phase lock status. If a 1PPS pulse is applied to the MFAM module, the 1 KHz cycle rate will phase lock onto the 1PPS pulse such that there will be an exact 1000 samples per second phase locked to the 1PPS pulse with a deviation of +/- 500 nS. When bit 15 is high, it indicates that the 1PPS pulse is being received. When bit 14 is high, it indicates that the 1 KHz sample is phase locked to the 1PPS input pulse.

10.1.1.3 MAGNETOMETER 1 FIELD READING (32 BITS):

The magnetometer 1 field value is located here. Each LSB is 50 femtoTeslas. Thus, for example, if the data bytes for this field were:

0x15, 0x85, 0x86, 0x2F

This would make the 32 bit number:

0x2F868515

Converting to decimal this would be:

797345045

Multiplying times 50E-6 to get nanoTeslas:

39867.2523 nT

During the power on startup process (when no magnetometer data is available) a different data set is sent in this slot. This contains diagnostic information when used with the startup diagnostic startup bits in the System Status data block (described above).

10.1.1.4 MAGNETOMETER 1 STATUS (32 BITS):

These bits are not currently defined.

10.1.1.5 MAGNETOMETER 2 FIELD READING:

The Magnetometer 2 field value is located here. Each LSB is 50 femtoTeslas. The field calculation proceeds as detailed in the Magnetometer 1 Field Reading calculation.

During the power on startup process (when no magnetometer data is available) a different data set is sent in this slot. This contains diagnostic information when used with the startup diagnostic startup bits in the System Status data block (described above).

10.1.1.6 MAGNETOMETER 2 STATUS:

These bits are not currently defined.

10.1.1.7 AUXILIARY CHANNEL X (16 BITS):

The three auxiliary sensor's X Axis data (Compass, Gyro, and Accelerometer) are multiplexed onto this channel. Which sensor data is being placed in this slot is defined by bits [13..11] in the Fid data block (described above).

During the power on startup process a different data set is sent in this slot. This contains diagnostic information when used with the startup diagnostic startup bits in the System Status data block (described above).

10.1.1.8 AUXILIARY CHANNEL Y (16 BITS):

The three auxiliary sensor's Y Axis data (Compass, Gyro, and Accelerometer) are multiplexed onto this channel. Which sensor data is being placed in this slot is defined by bits [13..11] in the Fid data block (described above).

During the power on startup process a different data set is sent in this slot. This contains diagnostic information when used with the startup diagnostic startup bits in the System Status data block (described above).

10.1.1.9 AUXILIARY CHANNEL Z (16 BITS):

The three auxiliary sensor's Z Axis data (Compass, Gyro, and Accelerometer) are multiplexed onto this channel. Which sensor data is being placed in this slot is defined by bits [13..11] in the Fid data block (described above).

During the power on startup process a different data set is sent in this slot. This contains diagnostic information when used with the startup diagnostic startup bits in the System Status data block (described above).

10.1.1.10 AUXILIARY TEMPERATURE (16 BITS):

The three auxiliary sensor's Temperature data (Compass, Gyro, and Accelerometer) are multiplexed onto this channel. Which sensor data is being placed in this slot is defined by bits [13..11] in the Fid data block (described above).

During the power on startup process a different data set is sent in this slot. This contains diagnostic information when used with the startup diagnostic startup bits in the System Status data block (described above).

10.1.2 ANCILLARY SENSOR ORIENTATION AND AXIS DEFINITION

The MFAM module contains an embedded Compass and Gyro/Accelerometer in the electronics housing. The Figures below define the position and axis orientation of these sensors. Note that the Compass is mounted on the back side of an internal PCB, so when looking down from the top the part is bottom side up. This will need to be taken into consideration when trying to match X/Y axis when comparing with the data sheet (which shows the part right side up).

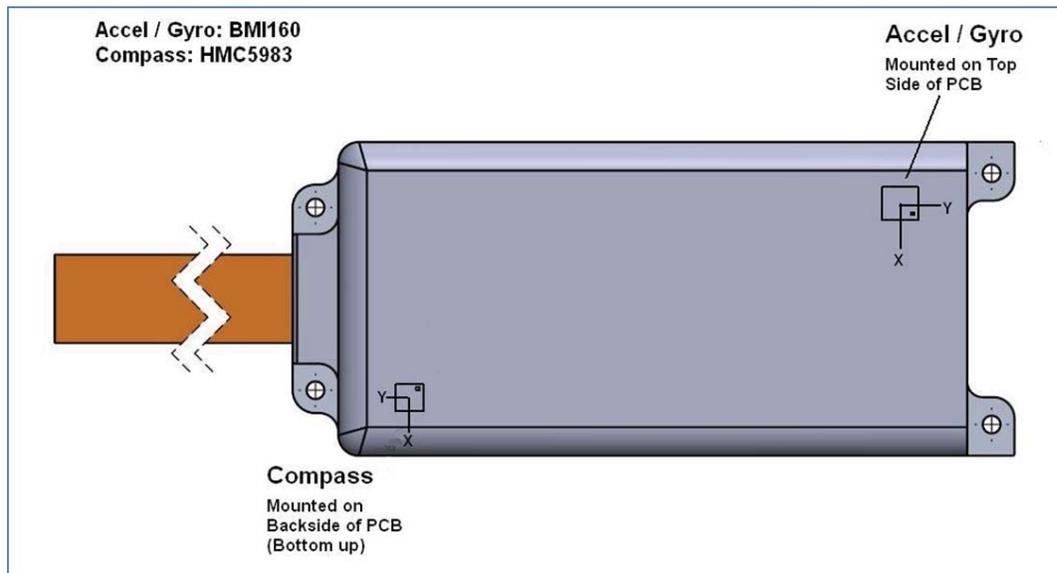


Figure 101. Position and orientation of ancillary sensors for orientation estimation